Graph-Adaptive Bernoulli Design for Bipartite Interference

CausalSmith*

August 25, 2026

Abstract

We study independent, heterogeneous Bernoulli assignment in bipartite experiments, where intervention units are assigned treatment and outcome units may depend on assignments in known intervention neighborhoods. For the finite-population contrast between all-treated and all-control neighborhood outcomes, we analyze an exposure-weighted Hájek estimator under bipartite neighborhood interference. With bounded potential outcomes, we derive a graph-and-design-dependent envelope that upper-bounds the design-based variance scale of its linearization and can be minimized subject to a positivity-constrained expected-treatment budget. Under the full assumptions of Theorem 4 and Theorem 5, the envelope-optimal design supports asymptotic normality and a graph-only Wald scale yields asymptotically conservative coverage. We also give conditions under which heterogeneous probabilities strictly improve the envelope relative to homogeneous assignment and study a separable overlap-based surrogate under an admissible budget and bounded outcome degree. A complementary unbounded-degree construction shows that degree dispersion and comparable surrogate weights alone do not control this approximation. Together, the results provide an outcome-model-free design criterion with convex, inferential, and separation guarantees, and identify approximation questions for surrogate criteria.

1 Introduction

Interference often separates the units receiving assignments from the units whose outcomes are measured. In bipartite experiments, treatments are assigned to one set of units, while outcomes are recorded for another set whose members may each depend on several assigned interventions. This structure arises in settings involving suppliers and customers, advertisers and users, providers and patients, or policies and exposed populations. It also creates a design problem: even when assignments are independent, overlapping intervention neighborhoods induce dependence among outcome-level exposure terms.

This paper considers a finite-population bipartite experiment with known graph structure and independent Bernoulli assignment probabilities that may differ across intervention units. The estimand is the average contrast between each outcome unit's all-treated and all-control neighborhood potential outcomes. Under the bipartite interference restriction in Assumption 1, the exposure-weighted Hájek estimator in Definition 3 uses realized all-treated and all-control

*Machine-generated by the CausalSmith research pipeline. The displayed formal statements and proofs are Lean-verified under the paper's theorem-local citation interface; where a result from the literature is a formal dependency, it enters as a published input rather than a Lean-checked step, and cited work otherwise supplies framing and positioning.

Approach	Assignment law	Outcome information in design	Objective and deliverable
Homogeneous bipartite Bernoulli analysis [Lu et al., 2025]	Common Bernoulli probability	None	Design-based estimation, asymptotic inference, and conservative variance analysis under homogeneous assignment.
Graph-cluster and bipartite-cluster designs [Ugander and Yin, 2020, Brennan et al., 2022]	Correlated assignment induced by clusters	Depends on the design objective and model restrictions	Cluster construction to alter exposure probabilities, spillovers, bias, power, or minimax performance.
Exposure-reweighted or model-based bipartite methods [Harshaw et al., 2021, Doudchenko et al., 2020]	Varies by method	Exposure or response-model structure in the cited approaches	Estimation or design criteria tailored to those specified structures.
This paper	Independent heterogeneous Bernoulli probabilities	None beyond the observed bipartite graph	Convex minimization of an analytically convenient bounded-outcome variance upper bound, a strict objective-separation diagnostic, and bounded-degree surrogate analysis.

Table 1: Scope of the graph-only heterogeneous-probability criterion relative to related approaches.

neighborhood exposures. The design question is how to allocate a fixed expected number of treated intervention units when potential outcomes are unavailable at the design stage.

The answer studied here is a conservative design criterion. Theorem 2 gives an outcome-model-independent graph envelope under heterogeneous Bernoulli assignment with $0 < p_k < 1$ and bounded potential outcomes. When the budget is admissible, we minimize this upper bound over the feasible class in Definition 1, which imposes both a positivity floor and a fixed expected-treatment budget. This yields a graph-adaptive allocation with an auditable robustness interpretation before outcomes are observed.

The criterion targets an analytically convenient uniform variance upper bound over the bounded-outcome class. Its value is that the objective is observable before randomization, permits an exact convex formulation, and comes with formal separation and inference guarantees. It provides a robust graph-only benchmark for subsequent numerical comparisons with homogeneous, clustered, and outcome-informed alternatives.

The paper's closest literatures differ in either the assignment law, the information used to choose a design, or the target of optimization.

The first contribution is a covariance representation that separates graph-and-design loads from unknown centered potential outcomes. Theorem 1 recovers the homogeneous Bernoulli case, in which same-exposure loads depend on the size of shared neighborhoods. Theorem 2 extends this representation to heterogeneous independent probabilities and establishes the corresponding envelope under bounded potential outcomes. This extension captures an intervention's full role in outcome-pair overlap beyond its intervention-side incidence.

The second contribution concerns the allocation problem induced by that envelope. Under an admissible positivity floor and budget, Theorem 3 establishes existence of an envelope minimizer and characterizes it with first-order conditions. The associated gradient score aggregates the overlap contribution of an intervention across outcome pairs. Theorem 6 shows that, when homogeneous-point scores differ and the homogeneous rate is strictly interior and unequal to one half, a heterogeneous direction strictly improves the envelope and admits a quantitative lower bound. In the special case of singleton outcome neighborhoods, unequal intervention-side incidences are sufficient for this score heterogeneity.

The third contribution studies a computational simplification and its formal scope. The envelope couples probabilities within shared neighborhoods, motivating a separable surrogate under an admissible positivity-constrained budget and bounded outcome degree. Theorem 8 complements this analysis with an unbounded-degree construction in which the approximation ratio diverges despite dispersed intervention-side degrees and comparable surrogate weights. Together, the results motivate reporting outcome-neighborhood and overlap diagnostics as checks when assessing surrogate reliability.

The inferential results remain design based throughout: potential outcomes are fixed, and assignment is the source of randomness. For sequences satisfying the full assumptions of Theorem 4, including outcome indexing and feasible envelope-optimal designs, Theorem 4 establishes a Hájek central limit theorem. Under the full assumptions of Theorem 5, the deterministic graph-only scale in Definition 9 yields asymptotically conservative coverage. The stated bounded-degree regime supplies the local-dependence control supporting both guarantees.

The paper is related to design-based causal inference under interference [Hudgens and Halloran, 2008, Aronow and Samii, 2017, Sävje et al., 2021], to approximate or structured neighborhood restrictions [Leung, 2019], and to bipartite causal designs [Zigler and Papadogeorgou, 2018, Chattopadhyay et al., 2023]. Relative to homogeneous Bernoulli bipartite analysis [Lu et al., 2025], the contribution is the graph-dependent heterogeneous-probability envelope and the associated optimization and approximation results. Relative to clustered designs, the assignment law remains independent across intervention units; whether that operational distinction is preferable is application dependent.

The verification appendix describes the precise scope of the formal verification support. The remainder of the paper sets out the experiment, estimator, and feasible design class; develops the envelope, inference, separation, and surrogate results; and discusses the resulting scope and limitations.

2 Setup and assumptions

At stage n, the experiment assigns interventions to a finite set I_n, with $|I_n| = m_n$, and records outcomes for the outcome-unit set O_n of size n. A known bipartite graph G_n links intervention units to outcome units. For outcome unit i, its intervention neighborhood $N_i(G_n)$ collects adjacent intervention units and has degree d_i. Conversely, let $M_k(G_n)$ denote the outcome units adjacent to intervention unit k, and write s_k for its cardinality. These two degree concepts can differ: d_i governs the number of assignments entering one outcome's exposure, whereas s_k summarizes the outcomes potentially affected by a given intervention.

For $i, j \in O_n$, the shared neighborhood $S_{ij}(G_n)$ equals $N_i(G_n) \cap N_j(G_n)$. Distinct outcomes are adjacent in the usual loop-free overlap graph when this set is nonempty. The formal quantity Δ_n is instead the maximum cardinality of the closed overlap neighborhood $\{j \in O_n : S_{ij}(G_n) \neq \varnothing\}$, so it includes i whenever $N_i(G_n) \neq \varnothing$. In the nonempty-neighborhood case it is the

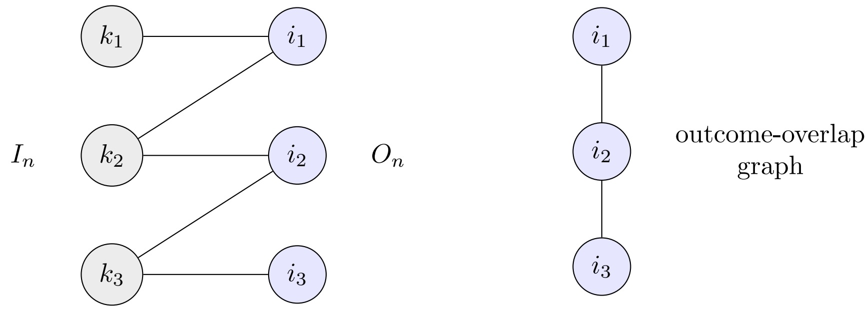

Figure 1: The bipartite graph G_n and its loop-free outcome-overlap graph. Two distinct outcome units are adjacent in the latter exactly when their intervention neighborhoods overlap. The formal dependence bound uses the corresponding closed overlap-neighborhood cardinality, which also counts the outcome itself.

conventional overlap-graph degree plus one. Thus, s_k measures intervention-side incidence, while Δ_n bounds the number of outcome-level exposure terms that can depend on a given term.

We work with a fixed finite-population potential-outcome schedule Y. Specifically, $Y_i(z_{N_i})$ is the potential outcome induced by a neighborhood assignment vector, and the all-treated and all-control potential outcomes are Y_i^1 and Y_i^0, respectively. The observed outcome is $Y_i^{\text{obs}} = Y_i(Z_{N_i(G_n)})$. Let $\mu_1 = n^{-1}\sum_{i\in O_n} Y_i^1$ and $\mu_0 = n^{-1}\sum_{i\in O_n} Y_i^0$, so that the finite-population estimand is $\tau_n = \mu_1 - \mu_0$.

The potential-outcome restriction is as follows.

Assumption 1. [`ass:bipartite-interference`] (Bipartite interference restriction) *For every $i \in O_n$ and every $z, z' \in \{0,1\}^{m_n}$,*

$$z_{N_i(G_n)} = z'_{N_i(G_n)} \quad \Longrightarrow \quad Y_i\big(z_{N_i(G_n)}\big) = Y_i\Big(z'_{N_i(G_n)}\Big).$$

Assumption 1 is the usual neighborhood-interference restriction for a bipartite experiment: outcome i depends exclusively on assignments in $N_i(G_n)$. It permits arbitrary interference within a neighborhood while making the exposure mapping graph-based and known to the researcher [Lu et al., 2025].

Assignment follows a finite-population Bernoulli randomization design. Write $Z = (Z_k)_{k\in I_n}$ for the assignment vector, with Z_k denoting its coordinate for intervention unit k, and let $p = (p_k)_{k\in I_n}$ denote the probability vector.

Assumption 2. [`ass:independent-heterogeneous-bernoulli`] (Independent heterogeneous Bernoulli assignment) *For every $k \in I_n$, Z_k is Bernoulli with success probability p_k, and the collection $\{Z_k : k \in I_n\}$ is mutually independent.*

Assumption 2 allows assignment probabilities to vary across intervention units while retaining independent randomization. This is the design feature that permits adaptation to the observed graph, in the spirit of heterogeneous exposure-based designs [Leung, 2019].

Fix a common positivity margin $0 < \epsilon < 1/2$ and an expected-treatment budget B_n. The design is subject to both restrictions.

Assumption 3. [`ass:positivity-floor`] (Assignment positivity) *For every $k \in I_n$,*

$$\epsilon \le p_k \le 1 - \epsilon.$$

Assumption 3 imposes uniform overlap through the positivity floor ϵ. It keeps every relevant exposure probability away from zero once outcome degrees are bounded, preventing unstable inverse-probability weights [Leung, 2019].

Assumption 4. [ass:budget-balance] (Assignment budget balance)

$$\sum_{k \in I_n} p_k = B_n.$$

Assumption 4 fixes the expected number of treated intervention units through the budget B_n. This restriction is specific to the present design problem: heterogeneous probabilities may be reallocated across intervention units, but their total must remain fixed.

Together, Assumption 3 and Assumption 4 define the admissible design set.

Definition 1. [def:feasible-designs] (Feasible design class $\mathcal{P}_{n,B_n,\epsilon}$) *Under Assumption 3 and Assumption 4, with $0 < \epsilon < 1/2$, define*

$$\mathcal{P}_{n,B_n,\epsilon} = \left\{ p \in [0,1]^{m_n} : \epsilon \le p_k \le 1 - \epsilon \text{ for every } k \in I_n, \quad \sum_{k \in I_n} p_k = B_n \right\}.$$

The class in Definition 1 is the intersection of a probability box and a budget hyperplane. It therefore isolates the design trade-off of interest: probabilities can respond to graph structure while preserving the experiment's expected treatment intensity.

For each outcome unit, define the treated exposure indicator $T_i(Z)$ as the indicator that every intervention in $N_i(G_n)$ is treated, and define the control exposure indicator $C_i(Z)$ analogously. Their design probabilities are $\pi_i^1(p)$ and $\pi_i^0(p)$. The inverse-probability denominators are

Definition 2. [def:hajek-denominators] (Hájek denominators $D_1(p, Z)$ and $D_0(p, Z)$) *Define*

$$D_1(p, Z) = \sum_{i \in O_n} \frac{T_i(Z)}{\pi_i^1(p)}, \qquad D_0(p, Z) = \sum_{i \in O_n} \frac{C_i(Z)}{\pi_i^0(p)}.$$

Here, $D_1(p, Z)$ and $D_0(p, Z)$ normalize the treated- and control-exposure weighted sums. The indicator convention in the next definition assigns value zero to a ratio whose corresponding exposure denominator is zero.

Definition 3. [def:hetero-hajek-estimator] (Hájek estimator $\widehat{\tau}_H(p)$) *Define*

$$\widehat{\tau}_H(p) = \mathbf{1}\{D_1(p, Z) > 0\} \frac{\sum_{i \in O_n} \frac{T_i(Z) Y_i^{obs}}{\pi_i^1(p)}}{D_1(p, Z)} - \mathbf{1}\{D_0(p, Z) > 0\} \frac{\sum_{i \in O_n} \frac{C_i(Z) Y_i^{obs}}{\pi_i^0(p)}}{D_0(p, Z)}.$$

Definition 3 is the exposure-weighted Hájek estimator. Normalization makes the estimator less sensitive to random fluctuations in the number of realized all-treated and all-control exposures than its unnormalized counterpart.

Its first-order behavior is described by the following centered summands.

Definition 4. [def:first-order-linearization] (Linearization summand $\eta_i(p, Z)$) *For each $i \in O_n$, define*

$$\eta_i(p, Z) = \left(\frac{T_i(Z)}{\pi_i^1(p)} - 1 \right) (Y_i^1 - \mu_1) - \left(\frac{C_i(Z)}{\pi_i^0(p)} - 1 \right) (Y_i^0 - \mu_0).$$

The summands in Definition 4 inherit a dependency graph from outcome-neighborhood overlap. Their scaled average has variance

Definition 5. [`def:variance-scale`] (Variance scale $\sigma^2_{G_n,p}(Y)$) *For a working finite design D, define*

$$\sigma^2_{G_n,D,p}(Y) = n \ \mathrm{Var}_D\left(\frac{1}{n}\sum_{i\in O_n}\eta_i(p,Z)\right).$$

When D is the heterogeneous Bernoulli design indexed by p, abbreviate this quantity as $\sigma^2_{G_n,p}(Y)$.

The variance scale $\sigma^2_{G_n,p}(Y)$ is design based: potential outcomes are held fixed and variation is taken only over the assignment mechanism.

We next construct a graph-only upper envelope for this scale. The treated and control covariance loads, $r^1_{ij}(G_n,p)$ and $r^0_{ij}(G_n,p)$, quantify the inverse-probability amplification arising from the shared neighborhood of i and j. The cross-exposure load $r^{10}_{ij}(G_n)$ records whether that shared neighborhood is nonempty. These quantities depend only on G_n and, for the same-exposure loads, on p.

Definition 6. [`def:graph-envelope`] (Graph envelope $V_{\mathrm{env}}(G_n,p)$) *Define*

$$V_{\mathrm{env}}(G_n,p) = \frac{4}{n}\sum_{i,j\in O_n}\left\{r^1_{ij}(G_n,p) + r^0_{ij}(G_n,p) + 2r^{10}_{ij}(G_n)\right\}.$$

The graph envelope $V_{\mathrm{env}}(G_n,p)$ replaces unknown outcome contrasts by a uniform bound, thereby making the design criterion observable before randomization. The envelope-optimal design is

Definition 7. [`def:optimal-design`] (Envelope-optimal design $p^*_n(G_n)$) *Whenever the minimizer set is nonempty, define*

$$p^*_n(G_n) \in \mathrm{argmin}_{p\in\mathcal{P}_{n,B_n,\epsilon}} V_{\mathrm{env}}(G_n,p).$$

Definition 7 selects $p^*_n(G_n)$ from the feasible envelope minimizers. To characterize its allocation of probability mass, we use the coordinatewise envelope score.

Definition 8. [`def:kkt-gradient`] (Envelope gradient score $g_k(G_n,p)$) *Define*

$$g_k(G_n,p) = \frac{1}{n}\sum_{i,j\in O_n:\ k\in S_{ij}(G_n)}\left\{-\left(\prod_{\ell\in S_{ij}(G_n)}p_\ell^{-1}\right)p_k^{-1} + \left(\prod_{\ell\in S_{ij}(G_n)}(1-p_\ell)^{-1}\right)(1-p_k)^{-1}\right\}.$$

Definition 8 gives the partial derivative of one quarter of the envelope. Its dependence on all pairs sharing intervention unit k shows that design value incorporates the full overlap pattern beyond direct degree.

For implementation, we also consider a separable surrogate. Define the overlap-derived first-order weight

$$h_k(G_n) = \frac{1}{n}\sum_{\substack{i,j\in O_n:\\ k\in S_{ij}(G_n)}}|S_{ij}(G_n)|^{-1}.$$

These weights average pairwise overlaps and reduce to degree-based summaries under additional graph structure.

Definition 9. [def:conservative-variance-estimator] (Conservative variance estimator $\widehat{V}_{\text{cons}}(G_n, p)$) *Define*

$$\widehat{V}_{\text{cons}}(G_n, p) = V_{\text{env}}(G_n, p).$$

The retained notation $\widehat{V}_{\text{cons}}(G_n, p)$ denotes a deterministic graph-and-design scale known before randomization.

Definition 10. [def:surrogate-design] (Surrogate design $p_n^{\text{deg}}(G_n)$) *Whenever the minimizer set is nonempty, define*

$$p_n^{\text{deg}}(G_n) \in \operatorname{argmin}_{p \in \mathcal{P}_{n,B_n,\epsilon}} \sum_{k \in I_n} h_k(G_n) \left\{ p_k^{-1} + (1 - p_k)^{-1} \right\}.$$

Definition 10 selects the computationally separable design $p_n^{\text{deg}}(G_n)$. Its performance relative to the envelope optimum is measured by

Definition 11. [def:approximation-ratio] (Approximation ratio $\alpha_{\text{cert}}(G_n)$) *Define*

$$\alpha_{\text{cert}}(G_n) = \begin{cases} \dfrac{V_{\text{env}}\left(G_n, p_n^{\text{deg}}(G_n)\right)}{V_{\text{env}}(G_n, p_n^*(G_n))}, & \textit{if } 0 < V_{\text{env}}(G_n, p_n^*(G_n)), \\ 1, & \textit{otherwise.} \end{cases}$$

Here, each design selector equals its respective feasible minimizer when one exists, and equals the homogeneous budget vector otherwise.

The approximation ratio $\alpha_{\text{cert}}(G_n)$ adopts a zero-loss convention when the optimal envelope is zero.

For later comparison with homogeneous assignment, let $\rho = B_n/m_n$ and set $p_k^{\text{hom}} = \rho$ for every $k \in I_n$. Write Δ_g for the difference between the largest and smallest values of $g_k(G_n, p^{\text{hom}})$, and let a and b index the corresponding extremal scores. The largest symmetric budget-feasible displacement from p^{hom} in direction $e_b - e_a$ is denoted by η_{box}, where e_k is the kth coordinate vector. The associated directional second-order modulus L_{ab} is

$$L_{ab} = \sup_{q \in \mathcal{P}_{n,B_n,\epsilon}} (e_b - e_a)^\top \operatorname{Hess}\{V_{\text{env}}(G_n, q)/4\} (e_b - e_a).$$

The supremum is explicitly over the feasible probability domain. Positivity therefore supplies a compact domain on which application-specific finite bounds on this directional curvature can be obtained.

The large-sample results impose regularity on potential outcomes and graph dependence.

Assumption 5. [ass:bounded-outcomes] (Bounded potential outcomes) *For every $i \in O_n$,*

$$|Y_i^1| \le 1 \qquad \textit{and} \qquad |Y_i^0| \le 1.$$

Assumption 5 is a standard bounded finite-population potential-outcome condition. It keeps the potential-outcome component of exposure-weighted summands uniformly bounded in the design-based model.

Assumption 6. [ass:bounded-outcome-degree] (Bounded outcome degree)

$$0 < \bar{d} \quad \textit{and} \quad \max_{i \in O_n} d_i \le \bar{d}.$$

Assumption 6 keeps every outcome neighborhood uniformly small through the positive constant $\bar{d}$. Together with uniform positivity, it limits the number of assignment factors entering all-treated and all-control exposure probabilities.

Assumption 7. [`ass:bounded-overlap-dependency`] (Bounded overlap dependency)

$$0 < \bar{D} \quad \textit{and} \quad \Delta_n \le \bar{D}.$$

Assumption 7 bounds outcome-level dependence by the positive constant $\bar{D}$. It is the standard dependency-graph sparsity condition that permits a central-limit argument despite interference among overlapping neighborhoods [Aronow and Samii, 2017].

Finally, inference requires an asymptotically positive variance scale under the optimal design.

Assumption 8. [`ass:variance-nondegenerate`] (Nondegenerate variance scale)

$$\liminf_{n\to\infty} \sigma^2_{G_n, p_n^*(G_n)}(Y) > 0.$$

Assumption 8 excludes sequences in which the design-based fluctuation of the linearized estimator degenerates. Such a lower-bound condition is standard for asymptotic normal approximation under interference [Lu et al., 2025].

3 Main results

This section characterizes the design-based variance of the Hájek linearization, then studies the graph-only envelope as a conservative allocation criterion. The outcome-agnostic criterion protects against the bounded potential-outcome schedules permitted by Assumption 5. It provides an exploratory design-stage rule for reallocating a fixed expected treatment budget using observed overlap structure. Related design-based analyses under interference likewise distinguish the randomization distribution from an outcome model [Aronow and Samii, 2017, Sävje et al., 2021].

We first recover the familiar homogeneous Bernoulli expressions as a special case. When all intervention units receive the same probability, only the size of the shared neighborhood determines the same-exposure covariance loads.

Theorem 1. [`prop:homogeneous-reduction`] (Homogeneous Bernoulli Reduction) *Under Assumption 2, suppose that the assignment probabilities are constant, $p_k = p$ for every $k \in I_n$, for some $p \in (0,1)$. Then, for every $i, j \in O_n$,*

$$r^1_{ij}(G_n, p) = \begin{cases} p^{-|S_{ij}(G_n)|} - 1, & |S_{ij}(G_n)| > 0, \\ 0, & |S_{ij}(G_n)| = 0, \end{cases} \qquad r^0_{ij}(G_n, p) = \begin{cases} (1-p)^{-|S_{ij}(G_n)|} - 1, & |S_{ij}(G_n)| > 0, \\ 0, & |S_{ij}(G_n)| = 0. \end{cases}$$

Moreover,

$$\sigma^2_{G_n,p}(Y) = \frac{1}{n} \sum_{i \in O_n} \sum_{j \in O_n} \left[r^1_{ij}(G_n, p)(Y^1_i - \mu_1)(Y^1_j - \mu_1) + r^0_{ij}(G_n, p)(Y^0_i - \mu_0)(Y^0_j - \mu_0) + 2r^{10}_{ij}(G_n)(Y^1_i - \mu_1)(Y^0_j - \right.$$

Theorem 1 shows that homogeneous assignment is governed by overlap counts beyond the marginal degrees of outcome units. It also supplies the benchmark against which the heterogeneous criterion below is compared.

3.1 Variance representation and envelope optimization

The next result gives the pairwise covariance representation under heterogeneous assignment and bounds the resulting variance scale by the graph envelope.

Theorem 2. [`thm:hetero-envelope`] (Heterogeneous envelope kernel) *Suppose that:*

- ***(Assignment probabilities.)*** $0 < p_k < 1$ *for every* $k \in I_n$*.*
- ***(Design.)*** *Assumption 2 holds with assignment-probability vector* p*.*
- ***(Outcomes.)*** *Assumption 5 holds.*

Then, for every $i, j \in O_n$*,*

$$\mathbb{E}_p[\eta_i(p,Z)\eta_j(p,Z)] = r^1_{ij}(G_n,p)(Y_i^1-\mu_1)(Y_j^1-\mu_1)+r^0_{ij}(G_n,p)(Y_i^0-\mu_0)(Y_j^0-\mu_0)+r^{10}_{ij}(G_n)\{(Y_i^1-\mu_1)(Y_j^0-\mu_0)+($$

Moreover,

$$\sigma^2_{G_n,p}(Y) \le V_{\mathrm{env}}(G_n,p).$$

Theorem 2 separates the schedule-dependent covariance terms from their graph-and-design loads. Bounded outcomes make the displayed envelope uniform over the admissible schedule class and give it a conservative design interpretation.

The envelope's inputs are known before randomization, and its minimization has a convex formulation. The following result establishes existence and provides the corresponding optimality characterization.

Theorem 3. [`thm:convex-design`] (Convex feasible design) *Let* G_n *be a bipartite experiment with intervention-unit set* I_n*, and let* $\epsilon, B_n \in \mathbb{R}$*. Suppose that*

- ***(Positivity margin.)*** $0 < \epsilon < 1/2$*;*
- ***(Feasible budget.)*** $m_n\epsilon \le B_n \le m_n(1-\epsilon)$*.*

Then $\mathcal{P}_{n,B_n,\epsilon}$ *is nonempty, compact, and convex, and* $p \mapsto V_{\mathrm{env}}(G_n,p)$ *is convex on* $\mathcal{P}_{n,B_n,\epsilon}$*. Moreover, there exists a minimizer* $p^*_n(G_n) \in \mathcal{P}_{n,B_n,\epsilon}$ *such that*

$$V_{\mathrm{env}}(G_n,p^*_n(G_n)) \le V_{\mathrm{env}}(G_n,q) \qquad \text{for every } q \in \mathcal{P}_{n,B_n,\epsilon}.$$

For every such minimizer $p^*_n(G_n)$*, there exist* $\lambda_n \in \mathbb{R}$ *and functions* $\nu^+_{\cdot,n}, \nu^-_{\cdot,n} : I_n \to \mathbb{R}$ *such that, for every* $k \in I_n$*,*

$$\nu^+_{k,n} \ge 0, \qquad \nu^-_{k,n} \ge 0, \qquad g_k(G_n,p^*_n(G_n)) = \lambda_n - \nu^+_{k,n} + \nu^-_{k,n},$$

and

$$\nu^+_{k,n}\big((p^*_n(G_n))_k-(1-\epsilon)\big)=0, \qquad \nu^-_{k,n}\big(\epsilon-(p^*_n(G_n))_k\big)=0.$$

Theorem 3 places the allocation problem in a standard finite-dimensional convex-optimization form [Boyd and Vandenberghe, 2004, Chs. 4–5]. Positivity makes every reciprocal-product term differentiable on the feasible box; convexity, the affine budget equality, and the two coordinate bounds supply the objective and constraint structure. At an interior coordinate, the envelope gradient score equals the common budget multiplier; at a probability bound, complementary slackness gives the corresponding one-sided condition. The formally verified statement covers feasible endpoint budgets directly. Because the objective couples probabilities across shared neighborhoods, the allocation depends on the full overlap score beyond s_k, intervention degree, or any single local statistic.

3.2 Large-sample inference and design-stage coverage

We next turn to inference for the envelope-optimal design. The central-limit result combines the Hájek linearization with the bounded-degree dependency structure induced by overlapping neighborhoods.

Theorem 4. [`thm:hetero-clt`] (Heterogeneous Hájek CLT) *Consider a sequence of bipartite experiments indexed by n. Suppose:*

- ***(Outcome indexing.)*** *For all sufficiently large n, $|O_n| = n$.*
- ***(Assignment and interference.)*** *Assumption 2 and Assumption 1 hold at every n for the assignment design D_n and probabilities p_n.*
- ***(Outcome and graph regularity.)*** *Assumption 5, Assumption 6, and Assumption 7 hold at every n, with constants $\bar{d}$ and $\bar{D}$, respectively.*
- ***(Feasibility and optimality.)*** *For every n, $p_n \in \mathcal{P}_{n,B_n,\epsilon_n}$ as in Definition 1, and*

$$p_n = p_n^*(G_n)$$

as in Definition 7.

- ***(Uniform positivity.)*** *There exists $\epsilon_0 > 0$ such that $\epsilon_0 \le \epsilon_n$ for all sufficiently large n.*
- ***(Variance.)*** *Assumption 8 holds for $\sigma^2_{G_n,p_n}(Y)$ as in Definition 5.*

Then, for every $\delta > 0$,

$$\Pr_{D_n}\left(\left|\sqrt{n}\big\{\widehat{\tau}_H(p_n^*(G_n)) - \tau_n\big\} - n^{-1/2}\sum_{i\in O_n}\eta_i(p_n^*(G_n), Z)\right| \ge \delta\right) \longrightarrow 0,$$

and, for every $s \in \mathbb{R}$,

$$\Pr_{D_n}\left(\frac{\sqrt{n}\big\{\widehat{\tau}_H(p_n^*(G_n)) - \tau_n\big\}}{\sqrt{\sigma^2_{G_n,p_n^*(G_n)}(Y)}} \le s\right) \longrightarrow \Pr\{N(0,1) \le s\}\,.$$

Theorem 4 is a design-based result: the potential-outcome schedule remains fixed throughout. Uniform positivity and bounded outcome degree control inverse-probability weights, bounded overlap dependency limits the number of dependent exposure terms associated with any outcome, and nondegeneracy supplies positive fluctuation. These familiar dependency-graph conditions define a sparse local-dependence regime [Aronow and Samii, 2017, Lu et al., 2025].

For applied use, the relevant graph diagnostics are the empirical distributions of d_i, intervention incidences s_k, and degrees in the outcome-overlap graph. Sequences with uniformly bounded values of the first and third quantities fall within the theorem's graph scope. These diagnostics therefore make the bounded-local-dependence requirement directly auditable.

The next result translates the same regularity conditions into a Wald interval based on the deterministic graph-only scale of Definition 9.

Theorem 5. [`thm:postdesign-wald`] (Post-design Wald coverage) *For a sequence of bipartite experiments G_n with outcome sets O_n, finite assignment designs, probability vectors p_n, positivity margins ϵ_n, and budgets B_n, suppose:*

- ***(Outcome indexing.)*** *$|O_n| = n$ for all sufficiently large n.*
- ***(Assignment law.)*** *The assignment design satisfies Assumption 2 with $0 \le p_{n,k} \le 1$ for every n and $k \in I_n$.*
- ***(Interference and boundedness.)*** *Assumption 1, Assumption 5, Assumption 6, and Assumption 7 hold with constants $\bar{d}$ and $\bar{D}$.*
- ***(Feasibility and optimality.)*** *For every n, p_n is a feasible design in the sense of Definition 1, $\epsilon_n \in (0, 1/2)$, and*

  $$p_n = p_n^*(G_n)$$

  in the sense of Definition 7.
- ***(Uniform positivity.)*** *There exists $\epsilon_0 > 0$ such that $\epsilon_0 \le \epsilon_n$ for all sufficiently large n.*
- ***(Nondegenerate variance.)*** *Assumption 8 holds for $\sigma^2_{G_n,p_n}(Y)$.*
- ***(Wald critical value.)*** *$\alpha_{\text{cov}} \in (0, 1)$, $z_{1-\alpha_{\text{cov}}/2} \ge 0$, and*

  $$\Phi\big(z_{1-\alpha_{\text{cov}}/2}\big) = 1 - \frac{\alpha_{\text{cov}}}{2}.$$

Then, for every n,

$$\sigma^2_{G_n,p_n}(Y) \le \widehat{V}_{\text{cons}}(G_n, p_n),$$

where $\sigma^2_{G_n,p_n}(Y)$ and $\widehat{V}_{\text{cons}}(G_n, p_n)$ are as in Definition 5 and Definition 9, respectively; moreover,

$$1 - \alpha_{\text{cov}} \le \liminf_{n\to\infty} \Pr_{p_n}\left(|\tau_n - \widehat{\tau}_H(p_n)| \le z_{1-\alpha_{\text{cov}}/2} \sqrt{\frac{\widehat{V}_{\text{cons}}(G_n, p_n)}{n}} \right).$$

Theorem 5 yields asymptotically conservative coverage using a scale determined entirely by the graph and selected probabilities, giving an outcome-model-free guarantee.

3.3 When heterogeneous probabilities improve on homogeneity

The envelope can strictly favor heterogeneity even under a fixed expected treatment budget. To state a quantitative comparison, recall the directional curvature quantity introduced in Section 2. Equivalently, along the budget-preserving direction from the highest-score coordinate a to the lowest-score coordinate b,

$$L_{ab} = \sup_{q \in \mathcal{P}_{n,B_n,\epsilon}} (e_b - e_a)^\top \operatorname{Hess}\{V_{\text{env}}(G_n, q)/4\}\,(e_b - e_a).$$

The supremum ranges only over the feasible probability domain. Since the positivity floor makes that domain compact and keeps every reciprocal probability factor finite, L_{ab} admits finite graph-specific bounds whenever the displayed derivatives are evaluated on that domain.

Theorem 6. [`thm:heterogeneity-separation`] (Heterogeneity separation) *Let G_n be a bipartite experiment with $m_n = |I_n| > 0$, and let $\epsilon \in (0, 1/2)$. Suppose that:*

- ***(Admissible budget.)*** *$m_n\epsilon \le B_n \le m_n(1 - \epsilon)$.*

- ***(Homogeneous rate.)*** $\rho = B_n/m_n$, *with* $\epsilon < \rho < 1-\epsilon$ *and* $\rho \neq 1/2$.
- ***(Homogeneous design.)*** $p_k^{\text{hom}} = \rho$ *for every* $k \in I_n$.

Then all of the following hold:

- *If there exist* $a, b \in I_n$ *such that*

$$g_a(G_n, p^{\text{hom}}) \neq g_b(G_n, p^{\text{hom}}),$$

then p^{hom} *is not a minimizer of* $V_{\text{env}}(G_n, \cdot)$ *over* $\mathcal{P}_{n,B_n,\epsilon}$. *Moreover, there exist* $a, b \in I_n$ *attaining, respectively, the maximum and minimum homogeneous-point gradient scores, such that*

$$\Delta_g := g_a(G_n, p^{\text{hom}}) - g_b(G_n, p^{\text{hom}}) > 0.$$

For every $p_n^*(G_n) \in \operatorname{argmin}_{p \in \mathcal{P}_{n,B_n,\epsilon}} V_{\text{env}}(G_n, p)$,

$$p_n^*(G_n) \neq p^{\text{hom}}, \qquad V_{\text{env}}(G_n, p_n^*(G_n)) < V_{\text{env}}(G_n, p^{\text{hom}}),$$

and

$$V_{\text{env}}(G_n, p^{\text{hom}}) - V_{\text{env}}(G_n, p_n^*(G_n)) \geq 2\Delta_g \begin{cases} \min\{\rho - \epsilon,\, 1 - \epsilon - \rho\}, & L_{ab} = 0, \\ \min\{\min\{\rho - \epsilon,\, 1 - \epsilon - \rho\},\, \Delta_g/L_{ab}\}\,, & L_{ab} \neq 0, \end{cases}$$

where L_{ab} *is the directional modulus along* $e_b - e_a$.

- *If* $|N_i(G_n)| = 1$ *for every* $i \in O_n$*, then, for every* $k \in I_n$,

$$g_k(G_n, p^{\text{hom}}) = n^{-1} s_k^2 \left((1-\rho)^{-2} - \rho^{-2} \right).$$

- *If* $|N_i(G_n)| = 1$ *for every* $i \in O_n$ *and there exist* $a, b \in I_n$ *such that* $s_a^2 \neq s_b^2$*, then the preceding non-homogeneity, strict-improvement, and gap conclusions hold.*

The singleton-neighborhood case gives a transparent diagnostic: when $\rho \neq 1/2$, heterogeneity in intervention-side incidence produces heterogeneous homogeneous-point scores and a strict envelope-improving direction away from homogeneous assignment. More generally, the score spread captures overlap patterns beyond s_k alone.

3.4 A bounded-degree surrogate and its limitation

The exact envelope couples probabilities across shared neighborhoods. The additive surrogate in Definition 10 removes that coupling and is therefore attractive for implementation. Under bounded outcome degree, its loss relative to the envelope optimum is uniformly controlled.

Theorem 7. [`thm:surrogate-certificate`] (Surrogate approximation certificate) *For a bipartite experiment on* G_n*, suppose:*

- ***(Admissible floor.)*** $0 < \epsilon < 1/2$.
- ***(Bounded outcome degree.)*** *Assumption 6 holds with bound* $\bar{d}$.
- ***(Admissible budget.)*** $m_n\epsilon \leq B_n \leq m_n(1-\epsilon)$.

Then, for every $p \in \mathcal{P}_{n,B_n,\epsilon}$ of Definition 1,

$$\sum_{k\in I_n} h_k(G_n)\{p_k^{-1} + (1-p_k)^{-1}\} \le \frac{V_{\mathrm{env}}(G_n,p)}{4} \le \max\Big\{1, \epsilon^{-(\bar{d}-1)}\Big\} \sum_{k\in I_n} h_k(G_n)\{p_k^{-1} + (1-p_k)^{-1}\}\,.$$

Moreover, the approximation ratio of Definition 11 satisfies

$$\alpha_{\mathrm{cert}}(G_n) \le \max\Big\{1, \epsilon^{-(\bar{d}-1)}\Big\}\,.$$

The certificate constant deteriorates with the allowed outcome degree and with a small positivity floor. Reporting $\max_i d_i$, ϵ, and the resulting bound therefore makes clear when the surrogate guarantee is informative and when it is effectively vacuous.

A natural conjecture is that broad dispersion in intervention-side degrees, combined with the one-sided conditional weight-ratio restriction used below, might itself control the surrogate's envelope approximation quality. The following result rules out such a conclusion; its counterexample can satisfy equal positive h-weights even though the formal condition is stated one-sidedly.

Theorem 8. [`thm:dispersion-certificate-unbounded`] (Unbounded dispersion certificate) *For every $\epsilon \in (0,1/2)$, $c_{\mathrm{disp}} > 0$, and $C_{\mathrm{disp}} \ge 1$, there exist sequences of finite intervention-unit sets I_n, finite outcome-unit sets O_n, finite bipartite experiments G_n, and budgets B_n such that, for every n,*

- ***(Admissible budget.)*** *With $m_n = |I_n|$,*

$$m_n\epsilon \le B_n \le m_n(1-\epsilon).$$

- ***(Positive energy.)***

$$0 < \sum_{k\in I_n} s_k^2.$$

Moreover, for all sufficiently large n,

- ***(Degree dispersion.)*** *For every $k \in I_n$,*

$$s_k^2 \le c_{\mathrm{disp}} \sum_{l\in I_n} s_l^2.$$

- ***(h-weight ratio.)*** *For every $k, l \in I_n$,*

$$h_l(G_n) > 0 \implies h_k(G_n) \le C_{\mathrm{disp}}\, h_l(G_n).$$

Finally, the approximation ratio of Definition 11 satisfies

$$\alpha_{\mathrm{cert}}(G_n) \longrightarrow \infty.$$

Theorem 8 shows that degree dispersion and the conditional restriction $h_l(G_n) > 0 \implies h_k(G_n) \le C_{\mathrm{disp}} h_l(G_n)$ cannot replace the bounded-degree condition behind Theorem 7. Thus, the separable rule is best treated as a bounded-degree approximation with an explicit certificate, not as a generally reliable proxy for the full overlap criterion.

4 Discussion and extensions

The proposed allocation rule uses the observed bipartite graph to choose independent assignment probabilities before outcomes are observed. Its objective is deliberately conservative: $V_{\text{env}}(G_n, p)$ bounds the design-based variance scale uniformly over the bounded potential-outcome schedules allowed by Assumption 5. The resulting design gives researchers a transparent, outcome-model-free way to reallocate a fixed expected treatment budget.

4.1 A prespecified design checklist

The following checklist turns the theoretical scope into a reproducible design-stage comparison using graph and design inputs alone.

1. **Fix capacity and candidate floors.** Set $\rho = B_n/m_n$ from the substantive treatment-capacity constraint. Any candidate floor must satisfy $0 < \epsilon \le \min\{\rho, 1-\rho\}$. Prespecify a small grid of substantively acceptable floors inside this interval to insulate the reported gain from floor selection.

2. **Audit graph scope.** Report the maximum and selected quantiles of d_i, s_k, and the closed outcome-overlap-neighborhood sizes. Flag persistent hubs or graph sequences whose maxima grow with n for analysis under growing-degree methods.

3. **Audit exposure support.** For every candidate design, compute

$$\pi_i^1(p) = \prod_{k \in N_i(G_n)} p_k, \qquad \pi_i^0(p) = \prod_{k \in N_i(G_n)} (1 - p_k),$$

   and report their minima and selected quantiles. An experiment-specific minimum acceptable exposure probability should be chosen before optimization. A candidate that violates it is rejected even if its envelope is smaller.

4. **Compare transparent benchmarks.** For each floor, report the homogeneous vector, the exact envelope minimizer when numerically available, and the separable surrogate as an exploratory benchmark. Report V_{env}, the relative envelope reduction from homogeneity, the number of coordinates at each probability bound, the KKT residual for the exact program, and the surrogate ratio.

5. **Apply fallback rules.** Retain homogeneous assignment when the envelope reduction is no larger than the declared numerical tolerance, changes sign or is highly unstable across acceptable floors, or requires unacceptable exposure probabilities. Treat the surrogate as exploratory when its empirical performance is too unstable or too weak for the application's prespecified tolerance. Evaluate clustering as a separate design candidate when cohesive graph regions, hubs, or exposure-support failures make independent assignment unattractive.

4.2 Limitations and future work

The numerical tolerances and minimum exposure probability in this checklist are application-specific design inputs, not universal constants supplied by the theory. Publishing the full floor-sensitivity table prevents them from becoming hidden researcher choices. Sparse implementations, computation benchmarks, and simulations relating the envelope to realized variance and coverage remain necessary before recommending routine use.

This perspective differs from clustering-based approaches to interference. Cluster randomization and graph-cluster constructions alter the dependence structure of assignment in order to increase the probability of informative exposures or reduce cross-cluster spillovers [Eckles et al., 2017, Ugander and Yin, 2020, Brennan et al., 2022]. Here the assignment law remains independent across intervention units; adaptation occurs through the probability vector. That distinction can be useful when independent implementation is operationally attractive, when clusters would be difficult to define, or when a fixed expected treatment total is an important design constraint. At the same time, clustering may be preferable when the graph contains cohesive regions for which correlated assignment produces substantially more relevant exposure variation.

The surrogate analysis motivates questions about computational simplification. Under an admissible positivity-constrained budget and bounded outcome degree, the surrogate can be compared with the full overlap criterion using constants determined by the positivity floor and degree bound. Thus, the additive rule should be treated as a substitute for the full overlap criterion only when its reported performance meets the application's prespecified tolerance.

Theorem 8 characterizes the unbounded-degree behavior of this surrogate perspective. The unbounded-degree construction shows that dispersed intervention-side incidences and the stated one-sided conditional h-weight-ratio restriction do not by themselves prevent the approximation ratio from diverging. In particular, the construction shows that intervention-side degree dispersion need not control higher-order overlap patterns that enter the exact criterion through shared neighborhoods. It therefore suggests reporting outcome-neighborhood and overlap diagnostics, rather than relying on intervention-side degree summaries alone. Related work on low-order and graph-aware experimental designs likewise emphasizes that graph structure beyond marginal degree can matter for design performance [Eichhorn et al., 2024, Harshaw et al., 2021].

Several extensions are natural but are not developed here. First, auxiliary pre-treatment outcomes or credible structural restrictions could support an outcome-informed criterion that targets an estimated variance rather than its uniform envelope. Such an approach would trade the present robustness for dependence on an outcome model and its validation. Second, clustered or dependent assignment mechanisms could be incorporated by replacing the independent-Bernoulli exposure probabilities and covariance loads with those induced by the chosen randomization scheme. Finally, graphs with growing degrees or overlap dependence may require asymptotic conditions tailored to their growth rates, beyond the bounded-degree regime used for Theorem 4. These extensions would connect the present design problem to broader work on network experimentation and optimized randomization [Ugander and Yin, 2020, Viviano, 2020].

Appendices

A Proofs and auxiliary lemmas

This appendix records the auxiliary probability result used by the large-sample analysis: the normal approximation for bounded summands indexed by a sparse dependency graph.

A.1 Dependency-graph normal approximation

The linearization summands in Definition 4 are dependent only when their outcome neighborhoods overlap. The next lemma states the bounded-degree central-limit result used to convert that structure into the asymptotic normality asserted in Theorem 4; it is the bounded-summand, bounded-dependency-degree specialization of the dependency-graph normal approximation of Baldi and Rinott [1989].

Lemma 1. [`lem:bounded-degree-dependency-clt`] (Bounded-Degree Dependency CLT) *Let $\{X_{n,i} : i \in \iota_n\}$ be real-valued random variables, where $|\iota_n| = n$, under the usual measurability and integrability conditions. Suppose:*

- ***(Centered summands.)*** *For every n and $i \in \iota_n$, $\mathbb{E}[X_{n,i}] = 0$.*
- ***(Dependency degree.)*** *The variables admit a dependency graph whose neighborhood of every $i \in \iota_n$ has cardinality at most D, for a fixed $D \in \mathbb{N}$.*
- ***(Uniform bound.)*** *There is an $M \geq 0$ such that $|X_{n,i}| \leq M$ for every n, $i \in \iota_n$, and outcome.*
- ***(Variance growth.)*** *Writing $S_n = \sum_{i \in \iota_n} X_{n,i}$, suppose $\mathbb{E}[S_n^2] = v_n$ for every n, and, for some $c > 0$, $v_n \geq cn$ for all sufficiently large n.*

Then, for every $s \in \mathbb{R}$,

$$\lim_{n\to\infty} \Pr\left(\frac{S_n}{\sqrt{v_n}} \leq s\right) = \Pr(Z \leq s)\,, \qquad Z \sim \mathcal{N}(0,1).$$

Proof of Lemma 1. Write $i \sim j$ for adjacency in the dependency graph carried by the array, and write

$$S_n = \sum_{i \in \iota_n} X_{n,i}.$$

The two structural properties of that graph that the argument uses are

$$|\{j \in \iota_n : \ j \sim i\}| \leq D \qquad \text{for every } n \text{ and every } i \in \iota_n,$$

and: whenever $A, B \subseteq \iota_n$ satisfy $a \not\sim b$ for all $a \in A$ and all $b \in B$, the tuples $(X_{n,a})_{a\in A}$ and $(X_{n,b})_{b\in B}$ are independent.

Step 1 (passage to a tail on which the variance floor holds). The convergence to be proved is a convergence of a sequence indexed by n, so it is unaffected by discarding a finite initial segment. By the variance-growth hypothesis and $|\iota_n| = n$, there is an N such that

$$v_n \geq cn \qquad \text{and} \qquad |\iota_n| = n \geq 1 \qquad \text{for every } n \geq N,$$

and it suffices to prove the asserted limit along $n \geq N$. Fix such an n. Then

$$v_n \geq cn \geq c > 0,$$

so $\sqrt{v_n} > 0$ and division by $\sqrt{v_n}$ is legitimate.

Step 2 (standardization). Define

$$Y_{n,i} = \frac{X_{n,i}}{\sqrt{v_n}} \qquad (i \in \iota_n), \qquad \text{so that} \qquad \sum_{i \in \iota_n} Y_{n,i} = \frac{S_n}{\sqrt{v_n}}.$$

Each $Y_{n,i}$ is the image of $X_{n,i}$ alone under the fixed measurable map $x \mapsto x/\sqrt{v_n}$, which does not depend on the outcome; applying this map coordinatewise to a tuple preserves independence, so the very same graph $\sim$ is a dependency graph for $\{Y_{n,i} : i \in \iota_n\}$. In particular the degree bound is unchanged:

$$|\{j \in \iota_n : \; j \sim i\}| \le D \qquad \text{for every } i \in \iota_n.$$

The standardized summands are still centered, and their sum has second moment one:

$$\mathbb{E}[Y_{n,i}] = \frac{\mathbb{E}[X_{n,i}]}{\sqrt{v_n}} = 0, \qquad \mathbb{E}\left[\left(\sum_{i\in\iota_n} Y_{n,i}\right)^2\right] = \frac{\mathbb{E}[S_n^2]}{v_n} = \frac{v_n}{v_n} = 1,$$

the last equality by the hypothesis $\mathbb{E}[S_n^2] = v_n$ and $v_n > 0$.

Step 3 (summand bound and Lyapunov ratio). Put

$$b_n = \frac{M/\sqrt{c}}{\sqrt{n}}.$$

Since $M \ge 0$ and $c > 0$ we have $b_n \ge 0$, and the uniform bound $|X_{n,i}| \le M$ together with $v_n \ge cn$ gives, for every $i \in \iota_n$ and every outcome,

$$|Y_{n,i}| = \frac{|X_{n,i}|}{\sqrt{v_n}} \le \frac{M}{\sqrt{v_n}} \le \frac{M}{\sqrt{cn}} = b_n.$$

Moreover

$$b_n \longrightarrow 0 \qquad \text{and} \qquad |\iota_n|\, b_n^3 = n \cdot \frac{(M/\sqrt{c})^3}{n^{3/2}} = \frac{(M/\sqrt{c})^3}{\sqrt{n}} \longrightarrow 0.$$

Step 4 (the dependency-graph normal approximation). The array $\{Y_{n,i} : i \in \iota_n\}$ now meets, term by term, the hypotheses of the following dependency-graph central limit theorem: if the variables of a triangular array are centered, admit a dependency graph all of whose closed neighborhoods have cardinality at most a fixed D, satisfy a uniform bound $|Y_{n,i}| \le b_n$ with $b_n \ge 0$, $b_n \to 0$ and $|\iota_n| b_n^3 \to 0$, and have $\mathbb{E}[(\sum_{i\in\iota_n} Y_{n,i})^2] = 1$, then, for every $s \in \mathbb{R}$,

$$\Pr\left(\sum_{i\in\iota_n} Y_{n,i} \le s\right) \longrightarrow \Pr(Z \le s)\,, \qquad Z \sim \mathcal{N}(0,1).$$

This is the Stein-method estimate for locally dependent summands: the dependency graph turns the Stein equation into a sum of neighborhood-localized error terms, and the two displayed limits $b_n \to 0$ and $|\iota_n| b_n^3 \to 0$ drive those error terms to zero.

Since $\sum_{i\in\iota_n} Y_{n,i} = S_n/\sqrt{v_n}$, this reads

$$\lim_{n\to\infty} \Pr\left(\frac{S_n}{\sqrt{v_n}} \le s\right) = \Pr(Z \le s)\,, \qquad Z \sim \mathcal{N}(0,1),$$

for every $s \in \mathbb{R}$, which is the assertion on the tail $n \ge N$ and hence, by Step 1, the assertion itself. □

The imported hypotheses map directly to the paper's primitives. Bounded outcomes, bounded outcome degree, and the positivity floor give a uniform bound on each centered linearization summand. The outcome-overlap graph is a dependency graph because disjoint intervention neighborhoods depend on disjoint independent assignments, and Assumption 7 bounds its closed-neighborhood size. Assumption 8, together with the paper's normalization, supplies variance growth of order n. These are precisely the boundedness, local-dependence, and nondegeneracy inputs used in the displayed design-based specialization of Baldi and Rinott [1989].

A.2 Convex-program provenance

The KKT characterization used in Theorem 3 is the standard convex-program optimality system for a differentiable objective with affine budget and box constraints [Boyd and Vandenberghe, 2004, Chs. 4–5]. The positivity box keeps the reciprocal objective differentiable; the theorem supplies convexity and existence; the budget equality and coordinate bounds are affine. Its displayed multipliers state dual feasibility, stationarity, and complementary slackness coordinate by coordinate. The verified theorem directly includes feasible endpoint budgets and the corresponding optimality conditions.

Verification note. The theorem and lemma statements, together with their proofs for the finite-design and independent Bernoulli-randomization components, are machine-checked in Lean 4. The formal development takes as hypotheses the experiment-specific bipartite graph and neighborhood structure, the fixed potential-outcome schedule and interference restriction, feasibility and positivity conditions, and the stated asymptotic regularity conditions on outcomes, degrees, overlap dependence, indexing, and variance nondegeneracy. Classical mathematical and probability-theory inputs used by the underlying libraries form the assumed formal substrate.

B Proofs of the main results

Proof of Theorem 1. Recall the three covariance loads. For $i, j \in O_n$ and an assignment-probability vector $p = (p_k)_{k \in I_n}$ with $0 < p_k < 1$,

$$r^1_{ij}(G_n, p) = \mathbf{1}\{|S_{ij}(G_n)| > 0\} \left(\prod_{k \in S_{ij}(G_n)} p_k^{-1} - 1 \right), \qquad r^0_{ij}(G_n, p) = \mathbf{1}\{|S_{ij}(G_n)| > 0\} \left(\prod_{k \in S_{ij}(G_n)} (1 - p_k)^{-1} - 1 \right)$$

$$r^{10}_{ij}(G_n) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}.$$

Throughout the proof, p denotes the probability vector and $\rho \in (0, 1)$ denotes the common homogeneous coordinate, so $p_k = \rho$ for every $k \in I_n$. We write $\rho^{-\ell}$ as shorthand for $(\rho^\ell)^{-1}$, and similarly for $(1 - \rho)^{-\ell}$. We also adopt throughout the reciprocal convention

$$1/0 = 0,$$

so that the symbol $1/n$, with $n = |O_n|$, is defined for every finite outcome-unit set: it is the ordinary reciprocal when $O_n \neq \varnothing$ and equals 0 when $O_n = \varnothing$. Every identity displayed below is asserted under this convention, so the degenerate case $O_n = \varnothing$ needs no separate hypothesis.

1. For every $k \in I_n$, homogeneity and $\rho \in (0, 1)$ give

$$0 < p_k = \rho < 1,$$

so all exposure probabilities $\pi_i^1(p) = \prod_{k \in N_i(G_n)} p_k$ and $\pi_i^0(p) = \prod_{k \in N_i(G_n)}(1 - p_k)$ are strictly positive and the reciprocals above are defined.

2. Fix $i, j \in O_n$. If $|S_{ij}(G_n)| > 0$, substituting $p_k = \rho$ into the displayed definition of the treated load yields

$$r_{ij}^1(G_n, p) = \prod_{k \in S_{ij}(G_n)} \rho^{-1} - 1 = (\rho^{|S_{ij}(G_n)|})^{-1} - 1 = \rho^{-|S_{ij}(G_n)|} - 1,$$

because the product has $|S_{ij}(G_n)|$ equal factors. If $|S_{ij}(G_n)| = 0$, the indicator in that definition vanishes and $r_{ij}^1(G_n, p) = 0$.

3. The same split applies to the control load, with ρ replaced by $1 - \rho$. When $|S_{ij}(G_n)| > 0$,

$$r_{ij}^0(G_n, p) = \prod_{k \in S_{ij}(G_n)} (1 - \rho)^{-1} - 1 = ((1 - \rho)^{|S_{ij}(G_n)|})^{-1} - 1 = (1 - \rho)^{-|S_{ij}(G_n)|} - 1;$$

when $|S_{ij}(G_n)| = 0$, it equals 0.

4. It remains to prove the variance-scale identity. Write, for $i \in O_n$,

$$\delta_i^1 = Y_i^1 - \mu_1, \qquad \delta_i^0 = Y_i^0 - \mu_0, \qquad A_i = \frac{T_i(Z)}{\pi_i^1(p)} - 1, \qquad B_i = \frac{C_i(Z)}{\pi_i^0(p)} - 1,$$

for the centered potential-outcome deviations and the centered inverse-probability weights, so that Definition 4 reads $\eta_i(p, Z) = A_i \delta_i^1 - B_i \delta_i^0$. Under Assumption 2 the coordinates of Z are independent, so

$$\mathbb{E}_p[T_i(Z)] = \prod_{k \in N_i(G_n)} p_k = \pi_i^1(p), \qquad \mathbb{E}_p[C_i(Z)] = \prod_{k \in N_i(G_n)} (1 - p_k) = \pi_i^0(p),$$

both strictly positive by the coordinate bounds above; hence $\mathbb{E}_p[A_i] = \mathbb{E}_p[B_i] = 0$ and

$$\mathbb{E}_p[\eta_i(p, Z)] = 0.$$

Unfolding Definition 5 and applying the finite-design variance formula to the linear combination $\sum_{i \in O_n} \frac{1}{n} \eta_i(p, Z)$ gives

$$\sigma_{G_n,p}^2(Y) = n \sum_{i \in O_n} \sum_{j \in O_n} \frac{1}{n} \cdot \frac{1}{n} \operatorname{Cov}_p(\eta_i(p, Z), \eta_j(p, Z)).$$

The centering just proved turns each covariance into the corresponding raw pair moment. If $O_n \neq \varnothing$, then $n > 0$ and the scalar identity $n \cdot (1/n)^2 = 1/n$ applies; if $O_n = \varnothing$, both the displayed right-hand side and the double sum below are empty, hence 0, and the two sides again agree under the reciprocal convention. In either case

$$\sigma_{G_n,p}^2(Y) = \frac{1}{n} \sum_{i \in O_n} \sum_{j \in O_n} \mathbb{E}_p[\eta_i(p, Z) \eta_j(p, Z)].$$

Next evaluate those pair moments. For any random variables U, V with $\mathbb{E}_p[U] = \theta_U \neq 0$ and $\mathbb{E}_p[V] = \theta_V \neq 0$, $\mathbb{E}_p[(U/\theta_U - 1)(V/\theta_V - 1)] = \mathbb{E}_p[UV]/(\theta_U\theta_V) - 1$. Independence gives $\mathbb{E}_p[T_i(Z)T_j(Z)] = \prod_{k\in N_i(G_n)\cup N_j(G_n)} p_k$, while $\pi_i^1(p)\pi_j^1(p) = \left(\prod_{k\in N_i(G_n)\cup N_j(G_n)} p_k\right)\left(\prod_{k\in S_{ij}(G_n)} p_k\right)$, so

$$\mathbb{E}_p[A_iA_j] = \prod_{k\in S_{ij}(G_n)} p_k^{-1} - 1 = r_{ij}^1(G_n, p).$$

The same argument with $1 - p_k$ in place of p_k gives

$$\mathbb{E}_p[B_iB_j] = \prod_{k\in S_{ij}(G_n)} (1 - p_k)^{-1} - 1 = r_{ij}^0(G_n, p),$$

with the empty-overlap case reducing to $1 - 1 = 0$, the value selected by the indicators in r_{ij}^1 and r_{ij}^0. For the mixed moments, if $S_{ij}(G_n) \neq \varnothing$, then $T_i(Z)C_j(Z) = 0$ identically, since a shared intervention cannot be both treated and untreated, so $\mathbb{E}_p[A_iB_j] = -1$. If $S_{ij}(G_n) = \varnothing$, then the two exposure events involve disjoint assignment coordinates, and independence gives $\mathbb{E}_p[T_i(Z)C_j(Z)] = \pi_i^1(p)\pi_j^0(p)$, hence $\mathbb{E}_p[A_iB_j] = 0$. In both cases, and likewise with the roles of the two outcome units exchanged,

$$\mathbb{E}_p[A_iB_j] = \mathbb{E}_p[B_iA_j] = -r_{ij}^{10}(G_n).$$

Expanding $\eta_i(p, Z)\eta_j(p, Z) = A_iA_j\delta_i^1\delta_j^1 - A_iB_j\delta_i^1\delta_j^0 - B_iA_j\delta_i^0\delta_j^1 + B_iB_j\delta_i^0\delta_j^0$ and inserting the four displayed moments gives

$$\mathbb{E}_p[\eta_i(p, Z)\eta_j(p, Z)] = r_{ij}^1(G_n, p)\delta_i^1\delta_j^1 + r_{ij}^0(G_n, p)\delta_i^0\delta_j^0 + r_{ij}^{10}(G_n)\left(\delta_i^1\delta_j^0 + \delta_i^0\delta_j^1\right).$$

Finally, $S_{ij}(G_n) = S_{ji}(G_n)$, so $r_{ij}^{10}(G_n) = r_{ji}^{10}(G_n)$; relabelling $(i, j) \mapsto (j, i)$ in the double sum gives

$$\sum_{i\in O_n}\sum_{j\in O_n} r_{ij}^{10}(G_n)\delta_i^0\delta_j^1 = \sum_{i\in O_n}\sum_{j\in O_n} r_{ij}^{10}(G_n)\delta_i^1\delta_j^0.$$

The two mixed contributions therefore combine into a single doubled term. Substituting the pair-moment identity into the pair expansion of $\sigma^2_{G_n,p}(Y)$ yields

$$\sigma^2_{G_n,p}(Y) = \frac{1}{n}\sum_{i\in O_n}\sum_{j\in O_n}\left[r_{ij}^1(G_n, p)(Y_i^1 - \mu_1)(Y_j^1 - \mu_1) + r_{ij}^0(G_n, p)(Y_i^0 - \mu_0)(Y_j^0 - \mu_0) + 2r_{ij}^{10}(G_n)(Y_i^1 - \mu_1\right.$$

as asserted.

$\square$

Proof of Theorem 2. Recall the three covariance loads: for $i, j \in O_n$,

$$r_{ij}^1(G_n, p) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}\left(\prod_{k\in S_{ij}(G_n)} p_k^{-1} - 1\right), \quad r_{ij}^0(G_n, p) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}\left(\prod_{k\in S_{ij}(G_n)} (1 - p_k)^{-1} - 1\right),$$

$$r_{ij}^{10}(G_n) = \mathbf{1}\{|S_{ij}(G_n)| > 0\},$$

all well defined under $0 < p_k < 1$.

1. Temporarily write

$$A_i = \frac{T_i(Z)}{\pi_i^1(p)} - 1, \qquad B_i = \frac{C_i(Z)}{\pi_i^0(p)} - 1.$$

Then Definition 4 gives

$$\eta_i(p, Z) = A_i(Y_i^1 - \mu_1) - B_i(Y_i^0 - \mu_0).$$

Under Assumption 2, the assignment coordinates are independent, and $0 < p_k < 1$ makes the exposure probabilities positive. Hence

$$\mathbb{E}_p[T_i(Z)] = \prod_{k \in N_i(G_n)} p_k = \pi_i^1(p), \qquad \mathbb{E}_p[T_i(Z)T_j(Z)] = \prod_{k \in N_i(G_n) \cup N_j(G_n)} p_k,$$

because $T_i(Z)T_j(Z)$ is the indicator that every intervention in $N_i(G_n) \cup N_j(G_n)$ is treated. For random variables U, V with $\alpha = \mathbb{E}_p[U] \neq 0$ and $\beta = \mathbb{E}_p[V] \neq 0$,

$$\mathbb{E}_p\left[\left(\frac{U}{\alpha} - 1\right)\left(\frac{V}{\beta} - 1\right)\right] = \frac{\mathbb{E}_p[UV]}{\alpha\beta} - 1.$$

Applying this identity to $T_i(Z)$ and $T_j(Z)$ gives

$$\mathbb{E}_p[A_i A_j] = \frac{\prod_{k \in N_i(G_n) \cup N_j(G_n)} p_k}{\pi_i^1(p)\pi_j^1(p)} - 1.$$

The last ratio is matched to $r_{ij}^1(G_n, p)$ by the two overlap branches. If $S_{ij}(G_n) \neq \varnothing$, then the nonzero factors satisfy

$$\frac{\prod_{k \in N_i(G_n) \cup N_j(G_n)} p_k}{\pi_i^1(p)\pi_j^1(p)} = \prod_{k \in S_{ij}(G_n)} p_k^{-1},$$

so $\mathbb{E}_p[A_i A_j] = \prod_{k \in S_{ij}(G_n)} p_k^{-1} - 1 = r_{ij}^1(G_n, p)$. If $S_{ij}(G_n) = \varnothing$, then $N_i(G_n)$ and $N_j(G_n)$ are disjoint, the union product factors as $\pi_i^1(p)\pi_j^1(p)$, and $\mathbb{E}_p[A_i A_j] = 0 = r_{ij}^1(G_n, p)$. Therefore

$$\mathbb{E}_p[A_i A_j] = r_{ij}^1(G_n, p).$$

The control-control calculation first gives

$$\mathbb{E}_p[B_i B_j] = \frac{\prod_{k \in N_i(G_n) \cup N_j(G_n)} (1 - p_k)}{\pi_i^0(p)\pi_j^0(p)} - 1.$$

It has the same two branches with $1 - p_k$ in place of p_k: on the nonempty-overlap branch the ratio equals $\prod_{k \in S_{ij}(G_n)} (1 - p_k)^{-1}$, while on the empty-overlap branch the disjoint union product factors and the centered ratio is zero. Thus

$$\mathbb{E}_p[B_i B_j] = r_{ij}^0(G_n, p).$$

For the mixed moment, if $S_{ij}(G_n) \neq \varnothing$, then a shared intervention is required to be both treated and untreated in $T_i(Z)C_j(Z)$, so $T_i(Z)C_j(Z) = 0$ for every assignment and

$\mathbb{E}_p[A_iB_j] = -1$. If $S_{ij}(G_n) = \varnothing$, the two exposure events use disjoint assignment coordinates, so $\mathbb{E}_p[T_i(Z)C_j(Z)] = \pi_i^1(p)\pi_j^0(p)$ and $\mathbb{E}_p[A_iB_j] = 0$. Thus

$$\mathbb{E}_p[A_iB_j] = -r_{ij}^{10}(G_n).$$

The same argument with i and j exchanged, using $S_{ij}(G_n) = S_{ji}(G_n)$, gives

$$\mathbb{E}_p[B_iA_j] = -r_{ij}^{10}(G_n).$$

Expanding the product of the two linearization summands,

$$\begin{aligned}\eta_i(p,Z)\eta_j(p,Z) = A_iA_j(Y_i^1-\mu_1)(Y_j^1-\mu_1) - A_iB_j(Y_i^1-\mu_1)(Y_j^0-\mu_0)\\ - B_iA_j(Y_i^0-\mu_0)(Y_j^1-\mu_1) + B_iB_j(Y_i^0-\mu_0)(Y_j^0-\mu_0),\end{aligned}$$

and inserting the four moment identities yields

$$\begin{aligned}\mathbb{E}_p[\eta_i(p,Z)\eta_j(p,Z)] = r_{ij}^1(G_n,p)(Y_i^1-\mu_1)(Y_j^1-\mu_1) + r_{ij}^0(G_n,p)(Y_i^0-\mu_0)(Y_j^0-\mu_0)\\ + r_{ij}^{10}(G_n)\left\{(Y_i^1-\mu_1)(Y_j^0-\mu_0) + (Y_i^0-\mu_0)(Y_j^1-\mu_1)\right\}.\end{aligned}$$

2. The one-score centered-ratio identities give $\mathbb{E}_p[A_i] = \mathbb{E}_p[B_i] = 0$, hence $\mathbb{E}_p[\eta_i(p,Z)] = 0$ for every i. Starting from Definition 5, the finite-design variance of the linear combination $n^{-1}\sum_i \eta_i(p,Z)$ expands as a double sum of covariances, and the zero means identify those covariances with joint moments. The empty-outcome branch is evaluated separately, where both finite sums and the scaled variance term are zero; on the nonempty branch, this is the ordinary pair average. Thus

$$\sigma_{G_n,p}^2(Y) = \frac{1}{n}\sum_{i\in O_n}\sum_{j\in O_n}\mathbb{E}_p[\eta_i(p,Z)\eta_j(p,Z)].$$

Substituting the preceding pair-moment identity and combining the two mixed double sums by the symmetry $S_{ij}(G_n) = S_{ji}(G_n)$ gives

$$\begin{aligned}\sigma_{G_n,p}^2(Y) = \frac{1}{n}\sum_{i\in O_n}\sum_{j\in O_n}\big\{r_{ij}^1(G_n,p)(Y_i^1-\mu_1)(Y_j^1-\mu_1) + r_{ij}^0(G_n,p)(Y_i^0-\mu_0)(Y_j^0-\mu_0)\\ + 2r_{ij}^{10}(G_n)(Y_i^1-\mu_1)(Y_j^0-\mu_0)\big\}.\end{aligned}$$

3. By Assumption 5, $|Y_i^1| \le 1$ and $|Y_i^0| \le 1$ for every i. Applying the triangle inequality to the finite-population means gives

$$|\mu_1| \le 1, \qquad |\mu_0| \le 1,$$

with the empty-population case evaluating both means as 0. Therefore, for every $i \in O_n$,

$$|Y_i^1 - \mu_1| \le 2, \qquad |Y_i^0 - \mu_0| \le 2.$$

The overlap loads are nonnegative under $0 < p_k < 1$. Indeed, when $S_{ij}(G_n) = \varnothing$ all three loads are 0; when $S_{ij}(G_n) \ne \varnothing$, the factors p_k^{-1} and $(1-p_k)^{-1}$ are at least 1, so the corresponding products minus 1 are nonnegative, and $r_{ij}^{10}(G_n) = 1$. Thus

$$r_{ij}^1(G_n,p) \ge 0, \qquad r_{ij}^0(G_n,p) \ge 0, \qquad r_{ij}^{10}(G_n) \ge 0.$$

4. For a nonnegative weight r and real numbers u, v with $|u| \le 2$ and $|v| \le 2$,

$$ruv \le r|u||v| \le 4r.$$

Applying this bound to the treated, control, and mixed terms using the bounds and non-negativity just established gives, for every ordered pair (i, j),

$$\begin{aligned} &r_{ij}^1(G_n, p)(Y_i^1 - \mu_1)(Y_j^1 - \mu_1) + r_{ij}^0(G_n, p)(Y_i^0 - \mu_0)(Y_j^0 - \mu_0) \\ &\quad + 2r_{ij}^{10}(G_n)(Y_i^1 - \mu_1)(Y_j^0 - \mu_0) \\ &\qquad \le 4\left\{r_{ij}^1(G_n, p) + r_{ij}^0(G_n, p) + 2r_{ij}^{10}(G_n)\right\}. \end{aligned}$$

5. Summing the pairwise inequality over $i, j \in O_n$ and multiplying by the nonnegative inverse n^{-1} gives

$$\begin{aligned} \sigma^2_{G_n,p}(Y) &\le \frac{1}{n} \sum_{i \in O_n} \sum_{j \in O_n} 4\left\{r_{ij}^1(G_n, p) + r_{ij}^0(G_n, p) + 2r_{ij}^{10}(G_n)\right\} \\ &= V_{\text{env}}(G_n, p), \end{aligned}$$

by Definition 6.

□

Proof of Theorem 3. Throughout write $m_n = |I_n|$ and $S_{ij} = S_{ij}(G_n)$, and interpret the factors n^{-1} in Definitions 6 and 8 as $(|O_n|)^{-1}$, with the total reciprocal convention $0^{-1} = 0$. Thus when $O_n = \varnothing$ the corresponding finite sums and scaled gradients are zero, and otherwise the notation is ordinary division by the outcome count. Recall that the three covariance loads entering Definition 6 are

$$r_{ij}^1(G_n, p) = \mathbf{1}\{|S_{ij}| > 0\}\left(\prod_{k \in S_{ij}} p_k^{-1} - 1\right), \qquad r_{ij}^0(G_n, p) = \mathbf{1}\{|S_{ij}| > 0\}\left(\prod_{k \in S_{ij}} (1 - p_k)^{-1} - 1\right),$$

$$r_{ij}^{10}(G_n) = \mathbf{1}\{|S_{ij}| > 0\}.$$

Note that every $p \in \mathcal{P}_{n,B_n,\epsilon}$ satisfies $0 < \epsilon \le p_k \le 1 - \epsilon < 1$, so all coordinate reciprocals in these loads are defined on the feasible class.

1. First suppose that I_n is nonempty and define the constant design

$$p_k = \frac{B_n}{m_n} \qquad (k \in I_n).$$

Since $m_n > 0$, the budget bounds $m_n\epsilon \le B_n \le m_n(1 - \epsilon)$ give

$$\epsilon \le \frac{B_n}{m_n} \le 1 - \epsilon,$$

and $0 < \epsilon < 1/2$ makes this interval a subset of $[0, 1]$. Thus p satisfies the probability and positivity requirements, and its budget is $\sum_{k \in I_n} p_k = m_n(B_n/m_n) = B_n$. If instead I_n is empty, the two budget inequalities read $0 \le B_n \le 0$ and force $B_n = 0$, so the unique empty design is feasible. Hence $\mathcal{P}_{n,B_n,\epsilon} \ne \varnothing$.

2. The feasible class is closed. Indeed, under $0 < \epsilon < 1/2$, Definition 1 describes it as the intersection of the coordinatewise closed probability constraints $0 \leq p_k \leq 1$, the coordinatewise closed bounds $\epsilon \leq p_k \leq 1-\epsilon$, and the closed affine hyperplane $\sum_{k\in I_n} p_k = B_n$; each of these is the preimage of a closed set under a continuous coordinate map or under the continuous finite sum. It is contained in the finite product $[0,1]^{I_n}$, which is compact; a closed subset of a compact set is compact, so $\mathcal{P}_{n,B_n,\epsilon}$ is compact.

3. Let $p, q \in \mathcal{P}_{n,B_n,\epsilon}$, and let $\theta, \theta' \geq 0$ with $\theta + \theta' = 1$. Coordinatewise,

$$\epsilon = \theta\epsilon + \theta'\epsilon \leq \theta p_k + \theta' q_k \leq \theta(1-\epsilon) + \theta'(1-\epsilon) = 1-\epsilon,$$

and by the same convex combination $0 \leq \theta p_k + \theta' q_k \leq 1$. Its budget is

$$\sum_{k\in I_n} (\theta p_k + \theta' q_k) = \theta \sum_{k\in I_n} p_k + \theta' \sum_{k\in I_n} q_k = \theta B_n + \theta' B_n = B_n.$$

Thus $\theta p + \theta' q \in \mathcal{P}_{n,B_n,\epsilon}$, proving convexity of the feasible class.

4. The envelope is continuous on the feasible class. Fix $i, j \in O_n$. If $S_{ij} \neq \varnothing$, then the maps $p \mapsto \prod_{k\in S_{ij}} p_k$ and $p \mapsto \prod_{k\in S_{ij}} (1-p_k)$ are continuous, and the bounds $\epsilon \leq p_k \leq 1-\epsilon$ with $0 < \epsilon < 1/2$ keep both products nonzero on $\mathcal{P}_{n,B_n,\epsilon}$; reciprocals of nonvanishing continuous functions are continuous, so $r^1_{ij}(G_n,\cdot)$ and $r^0_{ij}(G_n,\cdot)$ are continuous there. If $S_{ij} = \varnothing$, both are identically zero. Adding the constant $2r^{10}_{ij}(G_n)$, summing over the finitely many pairs (i,j), and multiplying by the fixed scalar $4(|O_n|)^{-1}$ shows that $V_{\mathrm{env}}(G_n,\cdot)$ is continuous on $\mathcal{P}_{n,B_n,\epsilon}$.

5. For each nonempty S_{ij}, the maps

$$p \longmapsto \left(\prod_{k\in S_{ij}} p_k\right)^{-1}, \qquad p \longmapsto \left(\prod_{k\in S_{ij}} (1-p_k)\right)^{-1}$$

are convex on $\mathcal{P}_{n,B_n,\epsilon}$. Indeed, on that class every p_k and every $1-p_k$ is strictly positive, so

$$\left(\prod_{k\in S_{ij}} p_k\right)^{-1} = \exp\left(\sum_{k\in S_{ij}} -\log p_k\right), \qquad \left(\prod_{k\in S_{ij}} (1-p_k)\right)^{-1} = \exp\left(\sum_{k\in S_{ij}} -\log(1-p_k)\right).$$

Each $p \mapsto -\log p_k$ is convex on the feasible class as a convex one-dimensional function composed with a linear coordinate map; likewise $p \mapsto -\log(1-p_k)$, with the affine coordinate map $p \mapsto 1-p_k$. Finite sums of convex functions are convex, and exp is convex and nondecreasing, so composing preserves convexity. Subtracting the constant 1 gives convexity of $r^1_{ij}(G_n,\cdot)$ and $r^0_{ij}(G_n,\cdot)$; when $S_{ij} = \varnothing$ both are constant, hence convex. Consequently each summand

$$r^1_{ij}(G_n,p) + r^0_{ij}(G_n,p) + 2r^{10}_{ij}(G_n)$$

is convex on the feasible class. Finite summation preserves convexity, and so does multiplication by the nonnegative scalar $4(|O_n|)^{-1}$. Thus $V_{\mathrm{env}}(G_n,\cdot)$ is convex on $\mathcal{P}_{n,B_n,\epsilon}$.

6. The preceding nonemptiness and compactness conclusions, together with the continuity of $V_{\mathrm{env}}(G_n,\cdot)$ on $\mathcal{P}_{n,B_n,\epsilon}$, put the envelope in the finite-dimensional extreme-value theorem. Hence it attains its minimum on $\mathcal{P}_{n,B_n,\epsilon}$. Denote one minimizer by $p_n^*(G_n)$. Then

$$V_{\mathrm{env}}(G_n, p_n^*(G_n)) \le V_{\mathrm{env}}(G_n, q)$$

for every $q \in \mathcal{P}_{n,B_n,\epsilon}$.

7. It remains to establish the multiplier statement, and for that we first identify the coordinate derivatives of the envelope. Let $\chi : \mathbb{R} \to \mathbb{R}$ be the standard smooth transition function with $\chi(t) = 0$ for $t \le 0$ and $\chi(t) = 1$ for $t \ge 1$. Define the floored reciprocal

$$R_\epsilon(x) = \begin{cases} \chi\left(\dfrac{x-\epsilon/4}{\epsilon/4}\right) x^{-1}, & x \neq 0, \\ 0, & x = 0. \end{cases}$$

Because $\chi((x-\epsilon/4)/(\epsilon/4))$ is identically zero on a neighborhood of the singular point $x = 0$, this function is continuously differentiable on all of $\mathbb{R}$. If $x \ge \epsilon/2$, then $(x-\epsilon/4)/(\epsilon/4) \ge 1$, so $R_\epsilon(x) = x^{-1}$. Replace every reciprocal factor in the two overlap kernels by this R_ϵ, obtaining

$$\widetilde{r}_{ij}^1(p) = \mathbf{1}\{|S_{ij}| > 0\} \left(\prod_{k \in S_{ij}} R_\epsilon(p_k) - 1 \right), \qquad \widetilde{r}_{ij}^0(p) = \mathbf{1}\{|S_{ij}| > 0\} \left(\prod_{k \in S_{ij}} R_\epsilon(1-p_k) - 1 \right),$$

$$\widetilde{V}(p) = 4(|O_n|)^{-1} \sum_{i,j \in O_n} \left\{ \widetilde{r}_{ij}^1(p) + \widetilde{r}_{ij}^0(p) + 2r_{ij}^{10}(G_n) \right\}.$$

Then $\widetilde{V}$ is continuously differentiable on all of $\mathbb{R}^{I_n}$, and

$$\widetilde{V}(p) = V_{\mathrm{env}}(G_n, p) \qquad \text{whenever } p_k \ge \epsilon/2 \text{ and } 1 - p_k \ge \epsilon/2 \text{ for every } k \in I_n.$$

Transporting $\widetilde{V}/4$ along a bijection between I_n and $\{1, \ldots, m_n\}$ to Euclidean coordinates gives a differentiable objective on $\mathbb{R}^{m_n}$.

For a feasible p, a coordinate k, and a shared neighborhood S_{ij}, differentiation along the k-th coordinate line gives

$$\frac{d}{dt}\bigg|_{t=0} \prod_{\ell \in S_{ij}} \left(p_\ell + \mathbf{1}\{\ell = k\}t\right)^{-1} = \begin{cases} -\left(\prod_{\ell \in S_{ij}} p_\ell^{-1}\right) p_k^{-1}, & k \in S_{ij}, \\ 0, & k \notin S_{ij}, \end{cases}$$

and

$$\frac{d}{dt}\bigg|_{t=0} \prod_{\ell \in S_{ij}} \left(1 - p_\ell - \mathbf{1}\{\ell = k\}t\right)^{-1} = \begin{cases} \left(\prod_{\ell \in S_{ij}} (1-p_\ell)^{-1}\right) (1-p_k)^{-1}, & k \in S_{ij}, \\ 0, & k \notin S_{ij}. \end{cases}$$

When $k \notin S_{ij}$ the product does not depend on t, and when $k \in S_{ij}$ only the factor indexed by k varies, contributing $\frac{d}{dt}(p_k + t)^{-1}|_{t=0} = -p_k^{-2}$ and $\frac{d}{dt}(1 - p_k - t)^{-1}|_{t=0} = (1-p_k)^{-2}$, respectively. Subtracting the constant 1 leaves these derivatives unchanged, so the same formulas hold for the treated and control overlap loads, including the case $S_{ij} = \varnothing$, where both derivatives vanish. Summing over (i,j) and multiplying by $(|O_n|)^{-1}$ gives exactly Definition 8:

$$\frac{d}{dt}\bigg|_{t=0} \frac{V_{\mathrm{env}}(G_n, p + te_k)}{4} = g_k(G_n, p).$$

Since $\epsilon \le p_k \le 1 - \epsilon$, every $|t| < \epsilon/2$ keeps $p + te_k$ in the region where $\widetilde{V}$ agrees with $V_{\mathrm{env}}(G_n, \cdot)$. Therefore the k-th Euclidean partial derivative of the transported objective at any feasible p is $g_k(G_n, p)$.

8. Fix any minimizer $p_n^*(G_n)$. In Euclidean coordinates, impose the equality constraint

$$h(x) = \sum_{u=1}^{m_n} x_u - B_n = 0$$

and the $2m_n$ inequality constraints

$$c_u^-(x) = x_u - \epsilon \ge 0, \qquad c_u^+(x) = 1 - \epsilon - x_u \ge 0 \qquad (u = 1, \dots, m_n).$$

The correspondence between the Euclidean points satisfying these constraints and $\mathcal{P}_{n,B_n,\epsilon}$ is exact. The displayed agreement of $\widetilde{V}$ with $V_{\mathrm{env}}(G_n, \cdot)$ on the feasible box shows that the transported smooth objective agrees on that set with $V_{\mathrm{env}}(G_n, \cdot)/4$; hence the Euclidean image x^* of $p_n^*(G_n)$ is a local minimum of the transported objective subject to those constraints. The objective is differentiable, the Euclidean domain is all of $\mathbb{R}^{m_n}$, and h, c_u^-, c_u^+ are affine, so the affine linear constraint qualification holds at x^*.

The finite-dimensional first-order necessary condition under this affine constraint qualification states: if x^* is a local minimum of a differentiable f subject to $h(x) = 0$ and $c_j(x) \ge 0$ with all constraints affine, then there exist $\lambda_n \in \mathbb{R}$ and $\nu_j \ge 0$ with

$$\nabla f(x^*) - \lambda_n \nabla h(x^*) - \sum_j \nu_j \nabla c_j(x^*) = 0 \qquad \text{and} \qquad \nu_j c_j(x^*) = 0 \quad \text{for every } j.$$

Write $\nu_{k,n}^-$ for the multiplier of c_u^- and $\nu_{k,n}^+$ for the multiplier of c_u^+ at the coordinate u corresponding to $k \in I_n$. Since ∇h is the all-ones vector, $\nabla c_u^- = e_u$ and $\nabla c_u^+ = -e_u$, the k-th coordinate of the stationarity condition reads, using the coordinate-gradient identity just established, whose right-hand side is the score in Definition 8,

$$g_k(G_n, p_n^*(G_n)) - \lambda_n - \nu_{k,n}^- + \nu_{k,n}^+ = 0,$$

or equivalently

$$g_k(G_n, p_n^*(G_n)) = \lambda_n - \nu_{k,n}^+ + \nu_{k,n}^-.$$

Dual feasibility gives $\nu_{k,n}^+ \ge 0$ and $\nu_{k,n}^- \ge 0$. Complementary slackness for the upper and lower constraints gives $\nu_{k,n}^+ c_u^+(x^*) = 0$ and $\nu_{k,n}^- c_u^-(x^*) = 0$, equivalently

$$\nu_{k,n}^+\big((p_n^*(G_n))_k - (1-\epsilon)\big) = 0, \qquad \nu_{k,n}^-\big(\epsilon - (p_n^*(G_n))_k\big) = 0.$$

This holds for every $k \in I_n$, as required.

∎

Proof of Theorem 4. Write $p_n = p_n^*(G_n)$. The proof uses the feasibility, positivity, assignment, interference, graph-regularity, and variance hypotheses attached to this sequence. Throughout, put

$$K(\epsilon, \bar{d}) = \epsilon^{-1} \max\{1, \epsilon^{-(\bar{d}-1)}\}.$$

1. Outcome indexing gives $|O_n| = n$ eventually. Hence $|O_n| \to \infty$, and, on this tail, $\sqrt{|O_n|} = \sqrt{n}$ and $|O_n|^{-1/2} = n^{-1/2}$. All limit statements below may therefore be proved after discarding the finitely many earlier stages.

2. Uniform positivity gives $\epsilon_0 > 0$ with $\epsilon_0 \le \epsilon_n$ eventually. Since each $\epsilon_n \in (0, 1/2)$, the denominator-kernel monotonicity bound gives

$$K(\epsilon_n, \bar{d}) \le K(\epsilon_0, \bar{d}) \qquad \text{eventually.}$$

   Together with $|O_n| \to \infty$, this fixed upper bound implies

$$\frac{K(\epsilon_n, \bar{d})}{|O_n|} \longrightarrow 0.$$

3. Define the two centered numerators and the two centered denominator ratios,

$$U_1(Z) = \sum_{i \in O_n} \frac{T_i(Z)}{\pi_i^1(p_n)} \left(Y_i^1 - \mu_1\right), \qquad U_0(Z) = \sum_{i \in O_n} \frac{C_i(Z)}{\pi_i^0(p_n)} \left(Y_i^0 - \mu_0\right),$$

$$R_{1n} = |O_n|^{-1/2} U_1(Z), \qquad R_{0n} = |O_n|^{-1/2} U_0(Z),$$
$$Q_{1n} = |O_n|^{-1} D_1(p_n, Z) - 1, \qquad Q_{0n} = |O_n|^{-1} D_0(p_n, Z) - 1.$$

   The scaled numerators are bounded in probability. Indeed, $\mathbb{E}_{D_n}[T_i(Z)/\pi_i^1(p_n)] = 1$ and $\mu_1 = |O_n|^{-1} \sum_i Y_i^1$ give $\mathbb{E}_{D_n}[U_1(Z)] = 0$, hence $\mathbb{E}_{D_n}[R_{1n}] = 0$; and expanding the variance into pair covariances, using Assumption 5 and the row bound of Assumption 7 together with the kernel bound supplied by Assumption 6 and the eventual floor $\epsilon_0 \le \epsilon_n$, one gets eventually

$$\mathrm{Var}_{D_n}(R_{1n}) \le 4\bar{D} K(\epsilon_0, \bar{d}), \qquad \mathrm{Var}_{D_n}(R_{0n}) \le 4\bar{D} K(\epsilon_0, \bar{d}).$$

   A sequence with mean zero and uniformly bounded variance is bounded in probability, by Chebyshev's inequality.

   The centered denominator ratios vanish in probability. The treated denominator has mean $|O_n|$, so $\mathbb{E}_{D_n}[Q_{1n}] = 0$. The same covariance expansion gives

$$\mathrm{Var}_{D_n}(D_1(p_n, Z)) \le |O_n| \bar{D} K(\epsilon_n, \bar{d}),$$

   and therefore

$$\mathrm{Var}_{D_n}(Q_{1n}) \le \frac{\bar{D} K(\epsilon_n, \bar{d})}{|O_n|} \longrightarrow 0$$

   by the denominator-kernel rate displayed above. Chebyshev's inequality then yields

$$Q_{1n} \xrightarrow{\mathrm{Pr}} 0, \qquad Q_{0n} \xrightarrow{\mathrm{Pr}} 0.$$

A product of a sequence bounded in probability with one vanishing in probability vanishes in probability, so

$$2|R_{1n}||Q_{1n}| + 2|R_{0n}||Q_{0n}| \xrightarrow{\mathrm{Pr}} 0.$$

For all sufficiently large n, split according as

$$D_1(p_n, Z) \geq \frac{|O_n|}{2} \quad \text{and} \quad D_0(p_n, Z) \geq \frac{|O_n|}{2}.$$

When both inequalities hold, both denominators are positive, the Hájek estimator in Definition 3 takes its ratio form in both arms, and expanding each ratio around its denominator mean gives the capped remainder bound

$$\left| \sqrt{|O_n|}\{\widehat{\tau}_H(p_n) - \tau_n\} - |O_n|^{-1/2} \sum_{i \in O_n} \eta_i(p_n, Z) \right| \leq 2|R_{1n}||Q_{1n}| + 2|R_{0n}||Q_{0n}|.$$

Here Assumption 1 identifies the observed outcome Y_i^{obs} with Y_i^1 on $\{T_i(Z) = 1\}$ and with Y_i^0 on $\{C_i(Z) = 1\}$, so that the two arm numerators are $U_1(Z) + \mu_1 D_1(p_n, Z)$ and $U_0(Z) + \mu_0 D_0(p_n, Z)$, and the cap $D_a(p_n, Z) \geq |O_n|/2$ bounds the factor $|O_n|/D_a(p_n, Z)$ by 2. If instead $D_1(p_n, Z) < |O_n|/2$, then $|O_n|^{-1} D_1(p_n, Z) < 1/2$, so $Q_{1n} < -1/2$ and $|Q_{1n}| > 1/2$; likewise for the control arm. Thus, for every $\delta > 0$, the probability that the absolute remainder exceeds δ is bounded by the probability of this denominator-ratio bad event plus

$$\Pr_{D_n}(2|R_{1n}||Q_{1n}| + 2|R_{0n}||Q_{0n}| \geq \delta)\,,$$

and both terms converge to zero.

4. Consequently, for every $\delta > 0$ the probability that the absolute remainder is at least δ tends to zero, that is

$$\sqrt{|O_n|}\{\widehat{\tau}_H(p_n) - \tau_n\} - |O_n|^{-1/2} \sum_{i \in O_n} \eta_i(p_n, Z) \xrightarrow{\mathrm{Pr}} 0.$$

Replacing $|O_n|$ by n on the outcome-indexing tail gives the first conclusion.

5. Set $X_{n,i} = \eta_i(p_n, Z)$ for $i \in O_n$, and declare i and j adjacent when $i = j$ or $S_{ij}(G_n) \neq \varnothing$. This is a dependency graph for $\{X_{n,i}\}$: the summand $\eta_i(p_n, Z)$ of Definition 4 is assembled from $T_i(Z)$ and $C_i(Z)$, together with $\pi_i^1(p_n), \pi_i^0(p_n), Y_i^1, Y_i^0, \mu_1, \mu_0$, so it depends on Z only through the coordinates in $N_i(G_n)$. If no outcome in $A \subseteq O_n$ is adjacent to any outcome in $B \subseteq O_n$, then $\bigcup_{i \in A} N_i(G_n)$ and $\bigcup_{j \in B} N_j(G_n)$ are disjoint, and Assumption 2 makes the two assignment-coordinate blocks independent. The closed neighborhood at i is contained in $\{i\} \cup \{j : S_{ij}(G_n) \neq \varnothing\}$, so Assumption 7 gives cardinality at most $\lceil \bar{D} \rceil + 1$.

   These summands have mean zero, because $\mathbb{E}_{D_n}[T_i(Z)/\pi_i^1(p_n)] = \mathbb{E}_{D_n}[C_i(Z)/\pi_i^0(p_n)] = 1$. Moreover Assumption 5 gives $|Y_i^1 - \mu_1| \leq 2$ and $|Y_i^0 - \mu_0| \leq 2$, while Definition 1, Assumption 6, and the eventual floor $\epsilon_0 \leq \epsilon_n$ give

$$\left(\pi_i^1(p_n)\right)^{-1} \leq K(\epsilon_0, \bar{d}) + 1, \qquad \left(\pi_i^0(p_n)\right)^{-1} \leq K(\epsilon_0, \bar{d}) + 1.$$

   Thus $|T_i(Z)/\pi_i^1(p_n) - 1| \leq K(\epsilon_0, \bar{d}) + 2$ and likewise for the control arm, and the eventual uniform bound

$$|X_{n,i}| \leq 4\{K(\epsilon_0, \bar{d}) + 2\}$$

holds for every $i \in O_n$ and every assignment.

Finally, with

$$v_n = \mathbb{E}_{D_n}\left[\left(\sum_{i \in O_n} X_{n,i}\right)^2\right],$$

expanding the square into pair moments and comparing with Definition 5 gives

$$v_n = \sum_{i \in O_n} \sum_{j \in O_n} \mathbb{E}_{D_n}[\eta_i(p_n, Z)\eta_j(p_n, Z)] = |O_n|\, \sigma^2_{G_n,p_n}(Y).$$

Assumption 8 supplies $c > 0$ with $\sigma^2_{G_n,p_n}(Y) \geq c$ eventually. Hence $v_n \geq c|O_n|$ eventually, and on the outcome-indexing tail this is the variance-growth condition $v_n \geq cn$. Applying Lemma 1 with $D = \lceil \bar{D} \rceil + 1$, $M = 4\{K(\epsilon_0, \bar{d}) + 2\}$, and this c yields, for every $s \in \mathbb{R}$,

$$\Pr_{D_n}\left(\frac{\sum_{i \in O_n} \eta_i(p_n, Z)}{\sqrt{v_n}} \leq s\right) \longrightarrow \Pr\{N(0,1) \leq s\}.$$

Since $\sqrt{v_n} = \sqrt{|O_n|}\sqrt{\sigma^2_{G_n,p_n}(Y)}$, the statistic inside is $\big(|O_n|^{-1/2} \sum_i \eta_i(p_n, Z)\big)/\sqrt{\sigma^2_{G_n,p_n}(Y)}$, so

$$\Pr_{D_n}\left(\frac{|O_n|^{-1/2} \sum_{i \in O_n} \eta_i(p_n, Z)}{\sqrt{\sigma^2_{G_n,p_n}(Y)}} \leq s\right) \longrightarrow \Pr\{N(0,1) \leq s\}.$$

6. The difference between the studentized Hájek statistic and the studentized linear-score statistic is

$$\frac{\sqrt{|O_n|}\{\widehat{\tau}_H(p_n) - \tau_n\} - |O_n|^{-1/2} \sum_{i \in O_n} \eta_i(p_n, Z)}{\sqrt{\sigma^2_{G_n,p_n}(Y)}}.$$

By Assumption 8, eventually $\sqrt{\sigma^2_{G_n,p_n}(Y)} \geq \sqrt{c}$. Hence, for every $\delta > 0$, the probability that the absolute value of the displayed difference is at least δ is eventually bounded above by

$$\Pr_{D_n}\left(\left|\sqrt{|O_n|}\{\widehat{\tau}_H(p_n) - \tau_n\} - |O_n|^{-1/2} \sum_{i \in O_n} \eta_i(p_n, Z)\right| \geq \delta\sqrt{c}\right),$$

which tends to zero by the linearization established above. The converging-together principle in CDF form then transfers the standard-normal CDF limit from the studentized linear-score statistic to the studentized Hájek statistic: for each x, sandwich the event for the Hájek statistic between the corresponding linear-score events at $x - \delta$ and $x + \delta$, up to the small-difference event, and let $\delta \downarrow 0$ using continuity of the standard-normal CDF. Therefore, for every $s \in \mathbb{R}$,

$$\Pr_{D_n}\left(\frac{\sqrt{|O_n|}\{\widehat{\tau}_H(p_n) - \tau_n\}}{\sqrt{\sigma^2_{G_n,p_n}(Y)}} \leq s\right) \longrightarrow \Pr\{N(0,1) \leq s\}.$$

Replacing $|O_n|$ by n on the outcome-indexing tail proves the second conclusion.

$\square$

Proof of Theorem 5. The eventual identity $|O_n| = n$ implies that $|O_n| \to \infty$; in particular $|O_n| > 0$ on an eventual tail.

1. Fix n. The admissibility condition $\epsilon_n \in (0, 1/2)$ and the feasibility floor give
$$0 < \epsilon_n \le p_{n,k} \le 1 - \epsilon_n < 1 \qquad (k \in I_n).$$
This display supplies the assignment-probability hypothesis $0 < p_{n,k} < 1$ in Theorem 2. Together with Assumption 2 and Assumption 5, the second conclusion of Theorem 2 gives
$$\sigma^2_{G_n,p_n}(Y) \le V_{\text{env}}(G_n, p_n).$$
By Definition 9, $\widehat{V}_{\text{cons}}(G_n, p_n) = V_{\text{env}}(G_n, p_n)$, and hence
$$\sigma^2_{G_n,p_n}(Y) \le \widehat{V}_{\text{cons}}(G_n, p_n).$$
Since n was arbitrary, this proves the variance-domination assertion for every n.

2. By Assumption 8 there are $c > 0$ and an eventual tail on which
$$c \le \sigma^2_{G_n,p_n}(Y).$$
Hence $\sigma^2_{G_n,p_n}(Y) > 0$ eventually. The eventual identity $|O_n| = n$ also gives the positive normalization $|O_n| > 0$ eventually. Under the assumptions listed in Theorem 4 – outcome indexing, assignment and interference, bounded outcomes, bounded degree, bounded overlap, feasibility, optimality $p_n = p_n^*(G_n)$, uniform positivity, and variance nondegeneracy – Theorem 4 gives, for every $s \in \mathbb{R}$,
$$\Pr_{p_n}\left(\frac{\sqrt{n}\,[\widehat{\tau}_H(p_n) - \tau_n]}{\sqrt{\sigma^2_{G_n,p_n}(Y)}} \le s \right) \longrightarrow \Phi(s).$$

3. Write
$$W_n = \frac{\sqrt{n}\,[\widehat{\tau}_H(p_n) - \tau_n]}{\sqrt{\sigma^2_{G_n,p_n}(Y)}}, \qquad z = z_{1-\alpha_{\text{cov}}/2}.$$
Since $z \ge 0$, the event $\{W_n \le -z\}$ is contained in $\{W_n \le z\}$, and the finite-design probability split gives
$$\Pr_{p_n}(-z < W_n \le z) = \Pr_{p_n}(W_n \le z) - \Pr_{p_n}(W_n \le -z)\,.$$
The convergence supplied by Theorem 4 at $s = z$ and $s = -z$, together with $\Phi(-z) = 1 - \Phi(z)$ and the calibration $\Phi(z) = 1 - \alpha_{\text{cov}}/2$, yields
$$\Pr_{p_n}(W_n \le z) - \Pr_{p_n}(W_n \le -z) \longrightarrow \Phi(z) - \Phi(-z) = 1 - \alpha_{\text{cov}}.$$
On the eventual tail where $|O_n|$, n, and $\sigma^2_{G_n,p_n}(Y)$ are positive, the strict interval event is contained in the Wald coverage event. Indeed, if $-z < W_n \le z$, then $|W_n| \le z$, so
$$\sqrt{n}\,|\widehat{\tau}_H(p_n) - \tau_n| \le z\sqrt{\sigma^2_{G_n,p_n}(Y)}.$$

Dividing by $\sqrt{n} > 0$ and using the variance-domination bound $\sigma^2_{G_n,p_n}(Y) \le \widehat{V}_{\text{cons}}(G_n, p_n)$, monotonicity of the square root, and $z \ge 0$, gives

$$|\tau_n - \widehat{\tau}_H(p_n)| \le z\sqrt{\frac{\sigma^2_{G_n,p_n}(Y)}{n}} \le z\sqrt{\frac{\widehat{V}_{\text{cons}}(G_n, p_n)}{n}}.$$

Consequently, on that tail,

$$\Pr_{p_n}\left(|\tau_n - \widehat{\tau}_H(p_n)| \le z\sqrt{\frac{\widehat{V}_{\text{cons}}(G_n, p_n)}{n}}\right) \ge \Pr_{p_n}(W_n \le z) - \Pr_{p_n}(W_n \le -z)\,.$$

Taking lower limits and using the displayed convergence gives

$$1 - \alpha_{\text{cov}} \le \liminf_{n\to\infty} \Pr_{p_n}\left(|\tau_n - \widehat{\tau}_H(p_n)| \le z_{1-\alpha_{\text{cov}}/2}\sqrt{\frac{\widehat{V}_{\text{cons}}(G_n, p_n)}{n}}\right),$$

which is the claimed coverage bound.

□

Proof of Theorem 8. Fix $\epsilon \in (0, 1/2)$, $c_{\text{disp}} > 0$, and $C_{\text{disp}} \ge 1$. Put

$$t_n = n + 1, \qquad d_n = t_n^2, \qquad \rho = \frac{\epsilon + 1/2}{2}, \qquad \varphi = \frac{3\rho - 1/2}{2}, \qquad R(x) = x^{-1} + (1 - x)^{-1},$$

so that $\epsilon < \rho < 1/2$ and $\epsilon < \varphi < 1/2$, both by direct substitution of $0 < \epsilon < 1/2$. Recall also the objects the argument evaluates: for $i, j \in O_n$,

$$r^1_{ij}(G_n, p) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}\left(\prod_{k\in S_{ij}(G_n)} p_k^{-1} - 1\right), \quad r^0_{ij}(G_n, p) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}\left(\prod_{k\in S_{ij}(G_n)} (1 - p_k)^{-1} - 1\right),$$

$$r^{10}_{ij}(G_n) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}.$$

1. Let I_n be the disjoint union of a core block of d_n interventions and a filler block of $2d_n$ interventions, so $m_n = 3d_n$. Let O_n consist of d_n core outcomes together with t_n filler outcomes for each filler intervention, so

   $$|O_n| = d_n + 2d_nt_n.$$

   Each core outcome is adjacent to every core intervention, while each filler outcome is adjacent only to its designated filler intervention; take all potential-outcome schedules to be zero. Explicitly, $N_i(G_n)$ is the whole core block when i is a core outcome, and $N_{(f,r)}(G_n) = \{f\}$ when (f, r) is the r-th filler outcome of the filler intervention f. Counting the outcomes adjacent to a given intervention,

   $$s_k = d_n \ \ (k \text{ core}), \qquad s_k = t_n \ \ (k \text{ filler}).$$

2. Set $B_n = m_n\rho$. Since $\epsilon < \rho < 1/2$, one has $\epsilon \le \rho \le 1 - \epsilon$, and multiplying by $m_n \ge 0$,

   $$m_n\epsilon \le B_n \le m_n(1 - \epsilon).$$

3. Summing the two degree values over the d_n core and $2d_n$ filler interventions, the degree energy is

$$\sum_{k\in I_n} s_k^2 = d_n\, d_n^2 + 2d_n\, t_n^2 = d_n^3 + 2d_n^2,$$

using $t_n^2 = d_n$; since $d_n \geq 1$ this is strictly positive.

4. Since $d_n \to \infty$ and $c_{\text{disp}} > 0$, for all sufficiently large n,

$$1 \leq c_{\text{disp}} d_n.$$

Fix such an n. If k is a core intervention, then $s_k^2 = d_n^2$ and, multiplying the displayed inequality by $d_n^2 > 0$,

$$s_k^2 = d_n^2 \leq c_{\text{disp}} d_n^3 \leq c_{\text{disp}}\left(d_n^3 + 2d_n^2\right) = c_{\text{disp}} \sum_{l\in I_n} s_l^2.$$

If k is a filler intervention, then $s_k^2 = t_n^2 = d_n$, and since $d_n \geq 1$,

$$s_k^2 = d_n \leq d_n^2 \leq c_{\text{disp}} d_n^3 \leq c_{\text{disp}} \sum_{l\in I_n} s_l^2.$$

Thus the degree-dispersion condition holds eventually.

5. Every intervention has the same strictly positive h-weight. Indeed, the ordered pairs (i, j) with $k \in S_{ij}(G_n)$ are: for a core k, the d_n^2 core-core pairs, each with $|S_{ij}(G_n)| = d_n$; and for a filler $k = f$, the t_n^2 pairs of filler outcomes both carrying the label f, each with $|S_{ij}(G_n)| = 1$. A core outcome and a filler outcome share no intervention, and two filler outcomes with distinct labels share none. Hence

$$h_k(G_n) = \frac{d_n^2 \cdot d_n^{-1}}{|O_n|} = \frac{d_n}{|O_n|} \quad (k \text{ core}), \qquad h_k(G_n) = \frac{t_n^2 \cdot 1}{|O_n|} = \frac{d_n}{|O_n|} \quad (k \text{ filler}),$$

again by $t_n^2 = d_n$. Consequently, for every $k, l \in I_n$,

$$h_k(G_n) = h_l(G_n) = \frac{d_n}{|O_n|} > 0 \qquad \text{and hence} \qquad h_k(G_n) \leq C_{\text{disp}} h_l(G_n),$$

the last step because $C_{\text{disp}} \geq 1$. So the h-weight-ratio condition holds for every n.

6. *The surrogate design is homogeneous.* The equal-weight calculation above makes the surrogate objective of Definition 10 a positive multiple of a separable barrier sum,

$$\sum_{k\in I_n} h_k(G_n)\left\{p_k^{-1} + (1-p_k)^{-1}\right\} = \frac{d_n}{|O_n|} \sum_{k\in I_n} R(p_k).$$

For $0 < r < 1$, the barrier R satisfies

$$R(x) - R(r) - R'(r)(x-r) = \frac{(x-r)^2\{r^2 + 2rx - 2r - x + 1\}}{r^2 x(1-r)^2(1-x)} \qquad (0 < x < 1).$$

The denominator is positive. The bracket is also positive: if $r \leq 1/2$, then

$$r^2 + 2rx - 2r - x + 1 = r^2 + (1-2r)(1-x) > 0,$$

and if $1/2 < r$, then

$$r^2 + 2rx - 2r - x + 1 = (1-r)^2 + (2r-1)x > 0.$$

Therefore

$$R(r) + R'(r)(x-r) \le R(x) \qquad (0 < x < 1),$$

and equality holds exactly when $x = r$, since the displayed gap then has positive denominator and positive bracket, leaving only the factor $(x-r)^2$ to vanish. Summing the tangent inequality over $k \in I_n$ at $r = \rho$ and using the budget $\sum_{k \in I_n} p_k = B_n = m_n \rho$, the linear terms cancel and

$$m_n R(\rho) \le \sum_{k \in I_n} R(p_k) \qquad \text{for every } p \in \mathcal{P}_{n,B_n,\epsilon},$$

with equality only if $p_k = \rho$ for every k. The homogeneous vector $p_k^{\text{hom}} = \rho$ is feasible, because $\epsilon < \rho < 1/2 < 1 - \epsilon$ and $\sum_{k \in I_n} \rho = m_n \rho = B_n$. Hence the surrogate minimizer is unique and

$$p_n^{\text{deg}}(G_n) = p^{\text{hom}}.$$

7. *Two envelope evaluations.* For any design p, Definition 6 splits over the outcome blocks. A core-core ordered pair has $S_{ij}(G_n)$ equal to the whole core block, of size d_n; a pair of filler outcomes with the same label f has $S_{ij}(G_n) = \{f\}$; all remaining ordered pairs have $S_{ij}(G_n) = \varnothing$ and contribute 0. There are d_n^2 pairs of the first kind and $2d_n t_n^2$ of the second. At the homogeneous design p^{hom} this gives

$$\frac{V_{\text{env}}(G_n, p^{\text{hom}})}{4} = \frac{1}{|O_n|} \left\{ d_n^2 \left(\rho^{-d_n} + (1-\rho)^{-d_n} \right) + 2d_n t_n^2 R(\rho) \right\},$$

since a core-core pair contributes $(\rho^{-d_n} - 1) + ((1-\rho)^{-d_n} - 1) + 2 = \rho^{-d_n} + (1-\rho)^{-d_n}$ and a same-label filler pair contributes $(\rho^{-1} - 1) + ((1-\rho)^{-1} - 1) + 2 = R(\rho)$.

For comparison, let q be the design assigning $1/2$ to every core intervention and φ to every filler intervention. It is feasible: $\epsilon < \varphi < 1/2$ and $\epsilon < 1/2 < 1 - \epsilon$ give the box constraints, and

$$\sum_{k \in I_n} q_k = d_n \cdot \frac{1}{2} + 2d_n \varphi = d_n \left(\frac{1}{2} + 3\rho - \frac{1}{2} \right) = 3d_n \rho = B_n.$$

The same block split, with a core-core pair now contributing $2^{d_n} + 2^{d_n} = 2^{d_n+1}$, gives

$$\frac{V_{\text{env}}(G_n, q)}{4} = \frac{1}{|O_n|} \left\{ d_n^2 \, 2^{d_n+1} + 2d_n t_n^2 R(\varphi) \right\}.$$

8. *The envelope minimum is positive.* The bounded-outcome-degree condition holds here with $\bar{d} = d_n$, since core outcomes have degree d_n and filler outcomes degree 1. Hence the lower half of the sandwich in Theorem 7, applied at the envelope-optimal design $p_n^*(G_n)$, gives

$$0 < \frac{d_n}{|O_n|} \sum_{k \in I_n} R\big((p_n^*(G_n))_k\big) = \sum_{k \in I_n} h_k(G_n) \left\{ (p_n^*(G_n))_k^{-1} + (1 - (p_n^*(G_n))_k)^{-1} \right\} \le \frac{V_{\text{env}}(G_n, p_n^*(G_n))}{4},$$

the strict positivity because $d_n/|O_n| > 0$, $I_n \neq \varnothing$, and $R > 0$ on $(0,1)$. Thus $V_{\text{env}}(G_n, p_n^*(G_n)) > 0$.

9. *Lower bound on the ratio.* The positivity just proved selects the positive branch of Definition 11, and the surrogate-design calculation identifies its numerator with the homogeneous envelope. Since q is feasible, $V_{\text{env}}(G_n, p_n^*(G_n)) \leq V_{\text{env}}(G_n, q)$, and $V_{\text{env}}(G_n, p^{\text{hom}}) \geq 0$, so

$$\alpha_{\text{cert}}(G_n) = \frac{V_{\text{env}}(G_n, p^{\text{hom}})}{V_{\text{env}}(G_n, p_n^*(G_n))} \geq \frac{V_{\text{env}}(G_n, p^{\text{hom}})}{V_{\text{env}}(G_n, q)}.$$

Substituting the two envelope evaluations above, cancelling the common factor $|O_n|^{-1}$, and using $2d_n t_n^2 = 2d_n^2$ so that d_n^2 cancels as well,

$$\frac{V_{\text{env}}(G_n, p^{\text{hom}})}{V_{\text{env}}(G_n, q)} = \frac{\rho^{-d_n} + (1-\rho)^{-d_n} + 2R(\rho)}{2^{d_n+1} + 2R(\varphi)} \geq \frac{\rho^{-d_n}}{2^{d_n+1} + 2R(\varphi)},$$

the inequality because $(1-\rho)^{-d_n} \geq 0$ and $R(\rho) > 0$. Hence

$$\alpha_{\text{cert}}(G_n) \geq \frac{\rho^{-d_n}}{2^{d_n+1} + 2R(\varphi)}.$$

10. Since $2^{d_n} \to \infty$, eventually $2R(\varphi) \leq 2^{d_n}$, so the denominator is at most $2^{d_n+1} + 2^{d_n} = 3 \cdot 2^{d_n}$ and

$$\alpha_{\text{cert}}(G_n) \geq \frac{\rho^{-d_n}}{3 \cdot 2^{d_n}} = \frac{1}{3}(2\rho)^{-d_n}.$$

Finally, $\rho < 1/2$, so $(2\rho)^{-1} > 1$; as $d_n \to \infty$, the right-hand side tends to $+\infty$. Therefore $\alpha_{\text{cert}}(G_n) \to \infty$.

□

Proof of Theorem 6. Throughout, e_k is the k-th coordinate vector on $\mathbb{R}^{I_n}$, and we write

$$\eta_{\text{box}} = \min\{\rho - \epsilon,\ 1 - \epsilon - \rho\} > 0,$$

which is positive because $\epsilon < \rho < 1-\epsilon$. Recall from Definition 8 the summation formula defining the envelope gradient score $g_k(G_n, p)$. Throughout we use the reciprocal convention

$$1/0 = 0,$$

so that the normalizing factor $1/n$, with $n = |O_n|$, is defined for every finite outcome-unit set: it is the ordinary reciprocal when $O_n \neq \varnothing$ and equals 0 when $O_n = \varnothing$. Under this convention the gradient formula of Definition 8 and the identities derived from it below hold verbatim, with no separate nonemptiness hypothesis.

1. Suppose that the homogeneous-point gradient scores are not all equal. First, $p^{\text{hom}} \in \mathcal{P}_{n,B_n,\epsilon}$: each coordinate equals $\rho \in (\epsilon, 1-\epsilon) \subseteq [0,1]$, and

$$\sum_{k \in I_n} p_k^{\text{hom}} = m_n \rho = B_n.$$

We next show that p^{hom} is not a minimizer. Otherwise, Theorem 3 supplies a budget multiplier λ_n and nonnegative upper- and lower-box multipliers $\nu_{k,n}^+$ and $\nu_{k,n}^-$ satisfying

$$g_k(G_n, p^{\text{hom}}) = \lambda_n - \nu_{k,n}^+ + \nu_{k,n}^-,$$

together with the complementary-slackness identities

$$\nu_{k,n}^{+}\big(p_k^{\mathrm{hom}} - (1-\epsilon)\big) = 0, \qquad \nu_{k,n}^{-}\big(\epsilon - p_k^{\mathrm{hom}}\big) = 0.$$

Since $p_k^{\mathrm{hom}} = \rho$ and $\epsilon < \rho < 1-\epsilon$, both factors multiplying the respective multipliers are nonzero. Thus $\nu_{k,n}^{+} = \nu_{k,n}^{-} = 0$ for every k, so all gradient scores equal λ_n, a contradiction.

Choose a attaining the maximum and b attaining the minimum of $k \mapsto g_k(G_n, p^{\mathrm{hom}})$, which exist because I_n is finite and nonempty. If these extrema were equal, all scores would be equal, contrary to hypothesis. Hence

$$\Delta_g = g_a(G_n, p^{\mathrm{hom}}) - g_b(G_n, p^{\mathrm{hom}}) > 0.$$

For these extremal indices, set

$$d = e_b - e_a, \qquad \gamma(s) = \frac{1}{4} V_{\mathrm{env}}(G_n, p^{\mathrm{hom}} + sd).$$

The segment $p^{\mathrm{hom}} + sd$ is feasible for $s \in [-\eta_{\mathrm{box}}, \eta_{\mathrm{box}}]$. Indeed $|d_k| \le 1$, so for $0 \le s \le \eta_{\mathrm{box}}$ every coordinate satisfies

$$\epsilon \le \rho - s \le \rho + sd_k \le \rho + s \le 1 - \epsilon$$

by the two bounds defining η_{box}; the budget is preserved because

$$\sum_{k \in I_n} d_k = 1 - 1 = 0,$$

so $\sum_k (p_k^{\mathrm{hom}} + sd_k) = B_n$. For $s < 0$ the same argument applies with the direction reversed. Moreover, on $0 \le s \le \eta_{\mathrm{box}}$,

$$p_k^{\mathrm{hom}} + sd_k \ge \rho - s \ge \epsilon \ge \epsilon/2, \qquad 1 - (p_k^{\mathrm{hom}} + sd_k) \ge 1 - \rho - s \ge \epsilon \ge \epsilon/2.$$

On this region the envelope agrees with the globally twice continuously differentiable cutoff extension obtained by replacing each reciprocal factor by a C^2 reciprocal cutoff that agrees with x^{-1} for $x \ge \epsilon/2$. Consequently γ is twice continuously differentiable on $[0, \eta_{\mathrm{box}}]$ and differentiable at 0. The line-calculus identity for $V_{\mathrm{env}}(G_n, \cdot)/4$ along d gives

$$\gamma'(0) = \sum_{k \in I_n} d_k \, g_k(G_n, p^{\mathrm{hom}}) = g_b(G_n, p^{\mathrm{hom}}) - g_a(G_n, p^{\mathrm{hom}}) = -\Delta_g.$$

Define the directional second-order modulus for this same pair by

$$L_{ab} = \sup_{q \in \mathcal{P}_{n, B_n, \epsilon}} \frac{d^2}{dt^2} \left\{ \frac{1}{4} V_{\mathrm{env}}(G_n, q + t(e_b - e_a)) \right\} \bigg|_{t=0}.$$

The supremum is finite: the feasible class is compact by Theorem 3, and the displayed directional curvature is continuous in the base point through the same C^2 cutoff extension. For $0 \le t \le \eta_{\mathrm{box}}$, the shifted base point $p^{\mathrm{hom}} + td$ is feasible, and the second derivative of γ at t is precisely the directional curvature at that base point. Hence

$$\gamma''(t) \le L_{ab} \qquad (0 \le t \le \eta_{\mathrm{box}}).$$

Convexity of $V_{\text{env}}(G_n, \cdot)$ on $\mathcal{P}_{n,B_n,\epsilon}$, supplied by Theorem 3, makes γ convex on the feasible line segment. Thus the curvature at the homogeneous base point is nonnegative, and since that curvature is one of the feasible values bounded by the displayed supremum,

$$0 \leq L_{ab}.$$

Define

$$s_* = \begin{cases} \eta_{\text{box}}, & L_{ab} = 0, \\ \min\{\eta_{\text{box}},\ \Delta_g / L_{ab}\}, & L_{ab} \neq 0. \end{cases}$$

Then $s_* \in [0, \eta_{\text{box}}]$. The one-dimensional second-order descent bound applied to γ, with initial slope at most $-\Delta_g$ and second derivative bounded above by L_{ab} on the segment, gives

$$\gamma(0) - \gamma(s_*) \geq \frac{\Delta_g}{2} s_*.$$

The design $q = p^{\text{hom}} + s_* d$ is feasible by the segment argument above. Therefore, for every $p_n^*(G_n) \in \operatorname{argmin}_{p \in \mathcal{P}_{n,B_n,\epsilon}} V_{\text{env}}(G_n, p)$,

$$V_{\text{env}}(G_n, p_n^*(G_n)) \leq V_{\text{env}}(G_n, q).$$

Multiplying the preceding descent inequality by 4, and using $4\gamma(0) = V_{\text{env}}(G_n, p^{\text{hom}})$ and $4\gamma(s_*) = V_{\text{env}}(G_n, q)$, gives

$$V_{\text{env}}(G_n, p^{\text{hom}}) - V_{\text{env}}(G_n, p_n^*(G_n)) \geq V_{\text{env}}(G_n, p^{\text{hom}}) - V_{\text{env}}(G_n, q) \geq 2\Delta_g s_*,$$

which is the asserted bound after substituting η_{box} and s_*.

Finally, $p_n^*(G_n) \neq p^{\text{hom}}$, since otherwise p^{hom} would be a minimizer. Optimality and feasibility of p^{hom} give $V_{\text{env}}(G_n, p_n^*(G_n)) \leq V_{\text{env}}(G_n, p^{\text{hom}})$; equality would again make p^{hom} a minimizer. Thus

$$V_{\text{env}}(G_n, p_n^*(G_n)) < V_{\text{env}}(G_n, p^{\text{hom}}).$$

2. Now suppose that $|N_i(G_n)| = 1$ for every $i \in O_n$, and fix $k \in I_n$. If $k \in S_{ij}(G_n)$, then $S_{ij}(G_n)$ is a nonempty subset of $N_i(G_n)$, and hence has cardinality one. Since p^{hom} is constant, every nonzero summand in the gradient formula of Definition 8, evaluated at p^{hom}, equals

$$-\rho^{-1}\rho^{-1} + (1-\rho)^{-1}(1-\rho)^{-1} = ((1-\rho)^2)^{-1} - (\rho^2)^{-1}.$$

For the counting term, the intervention-side degree is

$$s_k = \left|\{i \in O_n :\ k \in N_i(G_n)\}\right|.$$

The condition $k \in S_{ij}(G_n)$ is exactly the conjunction $k \in N_i(G_n)$ and $k \in N_j(G_n)$. Hence the ordered pairs (i, j) contributing to the sum are precisely

$$\{i \in O_n :\ k \in N_i(G_n)\} \times \{j \in O_n :\ k \in N_j(G_n)\},$$

and their number is s_k^2. Substitution into the gradient formula gives, for $\rho \in (0,1)$ and $n = |O_n|$,

$$g_k(G_n, p^{\text{hom}}) = \frac{s_k^2}{n}\left((1-\rho)^{-2} - \rho^{-2}\right),$$

which is the displayed formula. When $O_n \neq \varnothing$ both sides are read in the ordinary way; when $O_n = \varnothing$ the defining double sum is empty and $s_k = 0$ for every k, so both sides are 0 under the reciprocal convention fixed above.

3. Retain the singleton-neighborhood condition and suppose that $s_a^2 \neq s_b^2$ for some $a, b \in I_n$. Put
$$C = ((1-\rho)^2)^{-1} - (\rho^2)^{-1}.$$
This scalar is nonzero. Indeed, if $C = 0$, then $(1-\rho)^2 = \rho^2$; thus either $1-\rho = \rho$, contradicting $\rho \neq 1/2$, or $1-\rho = -\rho$, which would imply $1 = 0$.

   Also $n = |O_n| > 0$. Otherwise O_n would be empty, so $\{i \in O_n : k \in N_i(G_n)\}$ would be empty and $s_k = 0$ for every k, making $s_a^2 \neq s_b^2$ impossible. Hence $C/n \neq 0$, and the formula of the preceding step gives
$$g_a(G_n, p^{\text{hom}}) - g_b(G_n, p^{\text{hom}}) = \frac{C}{n}\left(s_a^2 - s_b^2\right) \neq 0,$$
so that
$$g_a(G_n, p^{\text{hom}}) \neq g_b(G_n, p^{\text{hom}}).$$
The conclusions of the first step therefore apply.

□

Proof of Theorem 7. Put
$$A(q) := \sum_{k\in I_n} h_k(G_n)\{q_k^{-1} + (1-q_k)^{-1}\}, \qquad C := \max\{1, \epsilon^{-(\bar{d}-1)}\}.$$
Then $C \geq 1$, hence $C \geq 0$. Recall also the three covariance loads, for $i, j \in O_n$:
$$r_{ij}^1(G_n, q) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}\left(\prod_{k\in S_{ij}(G_n)} q_k^{-1} - 1\right), \quad r_{ij}^0(G_n, q) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}\left(\prod_{k\in S_{ij}(G_n)} (1-q_k)^{-1} - 1\right),$$
$$r_{ij}^{10}(G_n) = \mathbf{1}\{|S_{ij}(G_n)| > 0\}.$$

1. By Theorem 3, whose positivity-margin and feasible-budget hypotheses are exactly those assumed here, the envelope-optimal design $p_n^*(G_n)$ is feasible and satisfies
$$V_{\text{env}}(G_n, p_n^*(G_n)) \leq V_{\text{env}}(G_n, q) \quad \text{for every } q \in \mathcal{P}_{n,B_n,\epsilon}.$$

2. On $\mathcal{P}_{n,B_n,\epsilon}$, every q_k and $1-q_k$ lies in $[\epsilon, 1-\epsilon]$, hence is strictly positive. Thus A is continuous there, being a finite sum of reciprocals of nonvanishing continuous coordinate functions. Nonemptiness and compactness from Theorem 3 therefore yield a feasible surrogate design $p_n^{\text{deg}}(G_n)$ such that
$$A(p_n^{\text{deg}}(G_n)) \leq A(q) \quad \text{for every } q \in \mathcal{P}_{n,B_n,\epsilon}.$$

3. Fix $q \in \mathcal{P}_{n,B_n,\epsilon}$, and write
$$K_{ij} := r_{ij}^1(G_n, q) + r_{ij}^0(G_n, q) + 2r_{ij}^{10}(G_n).$$
Substituting the definition of $h_k(G_n)$ into $A(q)$ and exchanging the order of the finite summations,
$$A(q) = \sum_{i\in O_n}\sum_{j\in O_n}\sum_{k\in S_{ij}(G_n)} n^{-1}|S_{ij}(G_n)|^{-1}\{q_k^{-1} + (1-q_k)^{-1}\},$$

and Definition 6 gives

$$\frac{V_{\text{env}}(G_n,q)}{4} = \sum_{i\in O_n}\sum_{j\in O_n} n^{-1}K_{ij}.$$

For a fixed pair (i,j), let $S = S_{ij}(G_n)$. If $S=\varnothing$, then $K_{ij}=0$ and the corresponding triple-sum contribution to $A(q)$ is an empty sum. Otherwise, since $0<q_k<1$ for each $k\in S$, the two reciprocal factors q_k^{-1} and $(1-q_k)^{-1}$ are at least one. Hence the product of the remaining treated reciprocal factors is at least one, so

$$q_k^{-1} \le \prod_{\ell\in S} q_\ell^{-1} \qquad (k\in S),$$

and averaging over $k\in S$ gives

$$|S|^{-1}\sum_{k\in S} q_k^{-1} \le \prod_{k\in S} q_k^{-1}.$$

The same argument applied to the control reciprocal factors gives

$$|S|^{-1}\sum_{k\in S}(1-q_k)^{-1} \le \prod_{k\in S}(1-q_k)^{-1}.$$

For nonempty S, the displayed definitions of the loads give

$$K_{ij} = \left(\prod_{k\in S} q_k^{-1} - 1\right) + \left(\prod_{k\in S}(1-q_k)^{-1} - 1\right) + 2 = \prod_{k\in S} q_k^{-1} + \prod_{k\in S}(1-q_k)^{-1}.$$

Adding the preceding two inequalities yields

$$|S|^{-1}\sum_{k\in S}\{q_k^{-1} + (1-q_k)^{-1}\} \le K_{ij}.$$

Multiplication by $n^{-1}\ge 0$ and summation over (i,j), with the empty-shared-neighborhood pairs contributing zero on both sides, prove

$$A(q) \le \frac{V_{\text{env}}(G_n,q)}{4}.$$

4. For the reverse bound, retain $S\neq\varnothing$. Since $S\subseteq N_i(G_n)$, Assumption 6 gives $|S|\le d_i\le \bar d$. Moreover, the feasibility bounds $\epsilon\le q_k\le 1-\epsilon$ give

$$0\le q_k^{-1}\le \epsilon^{-1}, \qquad 0\le (1-q_k)^{-1}\le \epsilon^{-1}.$$

After deleting any single index $k\in S$, each treated or control reciprocal product has $|S|-1$ factors and is therefore bounded by

$$(\epsilon^{-1})^{|S|-1} = \epsilon^{-(|S|-1)} \le \epsilon^{-(\bar d-1)} \le C,$$

the first inequality because $0<\epsilon\le 1$ makes $t\mapsto \epsilon^{-t}$ nondecreasing and $|S|\le \bar d$. For nonnegative reciprocal factors whose product after removal of any one factor is at most C, the full product is at most C times the factors' average: for each $k\in S$,

$$\prod_{\ell\in S} q_\ell^{-1} = q_k^{-1}\prod_{\ell\in S\setminus\{k\}} q_\ell^{-1} \le C q_k^{-1},$$

and averaging over $k \in S$ gives the treated reciprocal claim. Applying the same argument to the control reciprocal factors gives

$$\prod_{k\in S} q_k^{-1} \le C|S|^{-1}\sum_{k\in S} q_k^{-1}, \qquad \prod_{k\in S}(1-q_k)^{-1} \le C|S|^{-1}\sum_{k\in S}(1-q_k)^{-1}.$$

Adding these and using the preceding identity for K_{ij},

$$K_{ij} \le C|S|^{-1}\sum_{k\in S}\{q_k^{-1} + (1-q_k)^{-1}\}.$$

If $S = \varnothing$, then $K_{ij} = 0$ by the displayed load definitions, and the associated triple-sum contribution on the right-hand side of the final summed inequality is empty. Multiplying the nonempty-pair inequality by $n^{-1} \ge 0$, summing over all nonempty (i, j), and adding the empty pairs with zero contribution on both sides, the displayed expansions above yield

$$\frac{V_{\mathrm{env}}(G_n, q)}{4} \le C\,A(q).$$

This proves the asserted sandwich for every feasible q.

5. Let $V_* := V_{\mathrm{env}}(G_n, p_n^*(G_n))$. If $V_* > 0$, the two sandwich bounds and the minimizing property of $p_n^{\mathrm{deg}}(G_n)$ established above imply

$$\frac{V_{\mathrm{env}}\Big(G_n, p_n^{\mathrm{deg}}(G_n)\Big)}{4} \le C\,A(p_n^{\mathrm{deg}}(G_n)) \le C\,A(p_n^*(G_n)) \le C\,\frac{V_*}{4},$$

the middle step because $p_n^*(G_n)$ is feasible and $C \ge 0$. Thus

$$V_{\mathrm{env}}\Big(G_n, p_n^{\mathrm{deg}}(G_n)\Big) \le C V_*,$$

and division by $V_* > 0$, together with Definition 11, gives $\alpha_{\mathrm{cert}}(G_n) \le C$. If $V_* \not> 0$, that definition assigns $\alpha_{\mathrm{cert}}(G_n) = 1$, which is at most C. The claimed approximation-ratio bound follows.

□

References

Peter M. Aronow and Cyrus Samii. Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11(4):1912–1947, 2017. doi: 10.1214/16-AOAS1005.

Pierre Baldi and Yosef Rinott. On normal approximations of distributions in terms of dependency graphs. *The Annals of Probability*, 17(4):1646–1650, 1989. doi: 10.1214/aop/1176991178.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

Jennifer Brennan, Vahab Mirrokni, and Jean Pouget-Abadie. Cluster randomized designs for one-sided bipartite experiments. In *Advances in Neural Information Processing Systems*, volume 35, pages 37962–37974, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/f7f043c3438ba9e385c51bcf50ed007e-Abstract-Conference.html.

Ambarish Chattopadhyay, Kosuke Imai, and José R. Zubizarreta. Design-based inference for generalized network experiments with stochastic interventions. 2023. URL https://arxiv.org/abs/2312.03268.

Nick Doudchenko, Minzhengxiong Zhang, Evgeni Drynkin, Edoardo Airoldi, Vahab Mirrokni, and Jean Pouget-Abadie. Causal inference with bipartite designs. 2020. URL https://arxiv.org/abs/2010.02108.

Dean Eckles, Brian Karrer, and Johan Ugander. Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 5(1), 2017. doi: 10.1515/jci-2015-0021. URL https://arxiv.org/abs/1404.7530.

Matthew Eichhorn, Samir Khan, Johan Ugander, and Christina Lee Yu. Low-order outcomes and clustered designs: Combining design and analysis for causal inference under network interference. 2024. URL https://arxiv.org/abs/2405.07979.

Christopher Harshaw, Fredrik Sävje, David Eisenstat, Vahab Mirrokni, and Jean Pouget-Abadie. Design and analysis of bipartite experiments under a linear exposure-response model. 2021. URL https://arxiv.org/abs/2103.06392.

Michael G. Hudgens and M. Elizabeth Halloran. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482):832–842, 2008. doi: 10.1198/016214508000000292.

Michael P. Leung. Causal inference under approximate neighborhood interference. 2019. URL https://arxiv.org/abs/1911.07085.

Sizhu Lu, Lei Shi, Yue Fang, Wenxin Zhang, and Peng Ding. Design-based causal inference in bipartite experiments. 2025. URL https://arxiv.org/abs/2501.09844.

Fredrik Sävje, Peter M. Aronow, and Michael G. Hudgens. Average treatment effects in the presence of unknown interference. *The Annals of Statistics*, 49(2):673–701, 2021. doi: 10.1214/20-AOS1973. URL https://arxiv.org/abs/1711.06399.

Johan Ugander and Hao Yin. Randomized graph cluster randomization. 2020. URL https://arxiv.org/abs/2009.02297.

Davide Viviano. Experimental design under network interference. 2020. URL https://arxiv.org/abs/2003.08421.

Corwin M. Zigler and Georgia Papadogeorgou. Bipartite causal inference with interference. 2018. URL https://arxiv.org/abs/1807.08660.