Minimax mean squared error for low-order network interference under Bernoulli assignment

CausalSmith*

August 25, 2026

Abstract

For fixed interaction order $\beta \geq 1$ and fixed assignment probability $p \in (0, 1)$, this paper establishes a finite-population design-based minimax mean squared error frontier, up to constants depending on (β, p), for estimating the all-treated-versus-all-control effect under known bounded-degree low-order polynomial interference and common-probability Bernoulli assignment. The constants depend only on (β, p) and are uniform in n, d, and B; they may deteriorate as p approaches 0, 1, or a zero of the Bernoulli contrast. For radius B, population size n, and degree bound d, the paper's coefficient-mass and uniformly bounded-outcome minimax risks share the scale

$$B^2 \min\left\{1, \frac{dA_d}{n}\right\},$$

where A_d, the complete-block score energy, is the design second moment of the complete-block score used by the score-weighted neighborhood inverse-probability estimator (SNIPE, the shifted-neighbourhood inverse-probability estimator introduced by Cortez-Rodriguez et al. [2023]). Equivalently,

$$dA_d \asymp_{\beta,p} d\binom{d}{k_\star(d, \beta, p)},$$

where $k_\star(d, \beta, p)$, the largest interaction order with nonzero all-one-versus-all-zero Bernoulli contrast, is the exposed order in that contrast.

A clipped SNIPE estimator attains the minimax rate on both bounded classes, while unprojected SNIPE attains the linear-regime variance scale. Complete directed d-blocks, with $\lfloor n/d \rfloor$ active blocks in the block construction of Definition 15, supply the matching lower-bound construction. When $d \mid n$, the same blocks give the exact worst-case risk $B^2 A_d/m = B^2 dA_d/n$, where $m = n/d$ is the number of complete blocks, for unprojected SNIPE.

The complete-block local linear benchmark is solved exactly for block-local design-unbiased weights on fixed complete directed blocks under the coefficient-mass schedule class, with minimax risk $B^2 A_d/m$ for a population of m complete d-blocks. For the fair-coin design, even-order Bernoulli contrasts cancel, $A_d = 4\sum_{1 \leq r \leq \min\{\beta,d\},\, r \text{ odd}} \binom{d}{r}$, and first-order interference yields the frontier $B^2 \min\{1, d^2/n\}$.

*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.

1 Introduction

Randomized experiments on networks raise a basic design question: how much Bernoulli assignment variation is available for estimating a population-level contrast when each unit's outcome may depend on its neighbors' assignments? In the finite-population design-based tradition, potential outcomes are fixed and the assignment mechanism supplies the probability law for estimation [Horvitz and Thompson, 1952, Rubin, 1978, Imbens and Rubin, 2015]. Interference changes the problem by making the all-treated-versus-all-control total treatment effect depend on exposure patterns induced by the graph, so the risk of an estimator reflects both local treatment-contrast geometry and overlap among neighborhoods.

This paper studies that question for known directed interference graphs with bounded in- and out-degree, common-probability Bernoulli assignment, and low-order polynomial potential outcomes. The estimand is the finite-population all-treated-versus-all-control effect τ_n, the average difference between the all-one and all-zero assignments. The estimator observes the graph, the treatment assignment, and one realized outcome per unit. The two bounded classes considered here impose either a unitwise ℓ_1 coefficient-mass radius B or a uniform potential-outcome radius B; Definitions 6 and 12 define these classes.

The main result is a matched minimax frontier. For fixed interaction order $\beta \geq 1$ and treatment probability $p \in (0, 1)$, Theorems 1 and 2 establish

$$R^\star_{n,\ell_1}(d, \beta, p, B) \asymp_{\beta,p} R^\star_{n,\infty}(d, \beta, p, B) \asymp_{\beta,p} B^2 \min\left\{1, \frac{dA_d}{n}\right\}.$$

Here A_d, the complete-block score energy in Definition 10, is the design second moment of the Bernoulli score representing the local all-treated-versus-all-control contrast on a d-coordinate block. The same theorems compare this energy to the exposed binomial scale:

$$dA_d \asymp_{\beta,p} d\binom{d}{k_\star(d, \beta, p)},$$

where $k_\star(d, \beta, p)$ is the largest exposed interaction order with nonzero Bernoulli contrast as defined in Definition 1.

The rate separates two sources of design difficulty. The first is local: A_d is the squared norm of the Bernoulli representer for the block contrast. The second is graph-combinatorial: the extra factor d is the out-degree overlap charge, arising because one assignment coordinate can enter the scores of as many as d units under Assumption 2. The overlap count in Lemma 2 converts shared assignment monomials across neighborhoods into this single additional degree factor.

The estimator attaining the frontier is SNIPE with clipping. The unprojected statistic averages observed outcomes weighted by the unit score $g_{i,\beta,p}$ in Definitions 7 and 8. On the coefficient-mass class, projection to $[-B, B]$ supplies the saturated branch of the minimax bound; on the uniformly bounded-outcome class, projection to $[-2B, 2B]$ gives the corresponding bounded-outcome estimator in Definition 14. The unprojected estimator is design-unbiased on both bounded classes and has worst-case mean squared error at most $B^2 dA_d/n$, giving the linear branch directly.

The lower bound comes from complete directed d-blocks. The block family in Definition 15 perturbs compactly supported baseline outcomes in the direction of the normalized representer $h_d = g_d/A_d$. The perturbation changes the total treatment effect by a controlled amount while the Hellinger distance between the two sign-indexed prior predictive laws remains bounded through the energy identity for h_d. When $d \mid n$, Theorem 2 gives the exact complete-block worst-case risk of canonical unprojected SNIPE:

$$\frac{B^2 A_d}{m} = \frac{B^2 d A_d}{n},$$

where $m = n/d$ is the number of complete blocks.

The complete-block local linear benchmark fixes complete directed block graphs, restricts attention to block-local linear weights satisfying the unbiasedness equations in Definition 17, and evaluates them over the coefficient-mass schedules in Definition 16. Theorem 3 shows that the minimax risk over this class is exactly

$$\frac{B^2 A_{d_t}}{m_t}.$$

Its blockwise criterion $\Psi_{b,t}$ gives an exact worst-case risk decomposition, and the complete-block score g_{d_t} attains the benchmark. The same result characterizes asymptotic optimality through blockwise excess risk and gives a sufficient representer-closeness condition in $L^2(P_Z)$.

The fair-coin specialization makes the role of the assignment probability explicit. At $p = 1/2$, $\Delta_r(1/2)$ vanishes for even r and equals 2^{1-r} for odd r. Consequently, Theorem 4 gives

$$A_d = 4 \sum_{\substack{1 \le r \le \min\{\beta, d\} \\ r \text{ odd}}} \binom{d}{r}.$$

For first-order interference, this reduces to $A_d = 4d$, so the minimax scale becomes $B^2 \min\{1, d^2/n\}$, with exact complete-block unprojected SNIPE risk $4B^2 d^2/n$ when $d \mid n$.

In words, each neighborhood creates a local Bernoulli contrast problem, and A_d measures its design cost. The graph then charges one additional factor d because a treatment coordinate can enter many neighboring outcome equations. The frontier says that these two quantities, together with n and the envelope B, determine the finite-population design difficulty under the stated bounded low-order model.

The results connect to work on design-based causal inference under interference, including partial-interference, exposure-mapping, and network-experiment formulations [Sobel, 2006, Hudgens and Halloran, 2008, Bowers et al., 2013, Ugander et al., 2013, Aronow and Samii, 2017, Eckles et al., 2017, Ogburn and VanderWeele, 2017]. They are closest to low-order SNIPE analysis under Bernoulli assignment [Cortez-Rodriguez et al., 2023] and to Riesz and pseudoinverse perspectives on linear estimation [Swaminathan et al., 2017, Harshaw et al., 2022]. The frontier here gives the known-graph Bernoulli benchmark for bounded low-order polynomial interference, providing a calibration point for settings with graph uncertainty, clustered assignment, and covariate adjustment [Yu et al., 2022, Eichhorn et al., 2024, Wang and Li, 2026].

The remainder of the paper is organized as follows. The related-work section positions the result within design-based causal inference under interference and the SNIPE literature. The setup section fixes the finite population, Bernoulli design, low-order interference model, bounded model classes, SNIPE scores, minimax risks, and block quantities. The main-results section states the two-class minimax frontier, the complete-block lower construction and clipped-SNIPE upper bounds, the bounded-outcome counterpart, the exact local linear benchmark, and the fair-coin specialization. The discussion interprets the block score energy and out-degree overlap charge and records limitations and future work. The appendices collect the mathematical proofs; the cited literature supplies econometric framing, terminology, and positioning.

2 Related work

This paper belongs to the design-based tradition in which potential outcomes are fixed features of a finite population and treatment assignment supplies the probability law for estimation. The Horvitz–Thompson construction [Horvitz and Thompson, 1952] and the potential-outcome formulation of randomized causal effects [Rubin, 1978, Imbens and Rubin, 2015] provide the baseline language for the finite-population Bernoulli randomization framework used here. Relative to that tradition, interference changes the estimand and the feasible information in the design problem: outcomes may depend on neighbors' assignments, and the assignment mechanism must be read through the exposure structure induced by the graph.

A large literature studies causal inference when one unit's treatment can affect another unit's outcome. Early and foundational contributions include partial-interference and spillover formulations [Manski, 1993, Sobel, 2006, Hudgens and Halloran, 2008], network and peer-effect models [Graham, 2008, Bramoullé et al., 2009, Tchetgen and VanderWeele, 2010, Manski, 2013], and design-based estimands for network experiments [Bowers et al., 2013, Ugander et al., 2013, van der Laan, 2014, Liu et al., 2016, Aronow and Samii, 2017, Eckles et al., 2017, Ogburn and VanderWeele, 2017]. This work shares their finite-population orientation while specializing the outcome model to fixed low-order polynomial interference on a known bounded-degree graph. In that setting, the treatment assignments Z_j are independent Bernoulli draws with common probability p, the treatment probability, and the graph enters through the degree bound d and neighborhoods N_i.

The closest point of contact is low-order SNIPE estimation under Bernoulli assignment [Cortez-Rodriguez et al., 2023]. In the present notation, their SNIPE analysis supplies design-unbiasedness and a stated worst-case worst-case variance upper bound whose degree dependence enters through d^β, with constants depending on the interaction order and assignment probabilities, for bounded low-order polynomial interference. The present analysis characterizes the bounded-class design mean squared error in the low-order known-graph Bernoulli model through the complete-block score energy A_d in Definition 10: for interaction order β, the minimax scale is $B^2 \min\{1, dA_d/n\}$, equivalently $B^2 \min\{1, d\binom{d}{k_\star}/n\}$ up to constants depending on (β, p), where $k_\star$ is the exposed nonzero Bernoulli contrast order in Definition 1. Within this bounded-class setting, Theorems 1 and 2 provide the matched minimax lower construction, the attaining clipped SNIPE upper bound, the uniformly bounded-outcome extension, the exposed-binomial calibration, and exact complete-block constants. The comparison to the unrestricted class studied in Cortez-Rodriguez et al. [2023] is through that paper's stated theorem, while the present frontier supplies the bounded-class calibration in the notation and model class used here. The attaining procedure is a clipped SNIPE estimator from Definitions 8 and 14, built from the unit score in Definition 7, under the finite-population Bernoulli model.

It is worth being precise about the scope of the comparison. Cortez-Rodriguez et al. [2023] study low-order neighborhood interference under Bernoulli assignment and establish design-unbiasedness of SNIPE together with a worst-case variance upper bound, in a setting that leaves the graph and the coefficient magnitudes free. The present paper fixes a smaller class — known graph, in- and out-degree bounded by d, interaction order at most β, and unitwise coefficient mass at most B — and asks for the minimax risk over that class. Within it, the exposed-binomial scale $d\binom{d}{k_\star}$ takes the place of the d^β factor carried by their stated bound, and the matching lower bound shows this rate is the correct one. The sharpening thus applies to the stated worst-case bound read on this restricted class; their own theorem, stated for their broader setting, stands as published, and the estimator compared against here is the same SNIPE family.

Sparse and local interference results provide a second comparison point. Work on exposure mappings, neighborhood interference, and asymptotic regimes with growing graphs develops conditions under which design-based estimators remain informative as the interference structure grows [Aronow and Samii, 2017, Athey et al., 2017, Baird et al., 2018, Basse and Airoldi, 2018, Basse and Feller, 2018, Jagadeesan et al., 2020, Sävje et al., 2021, Leung, 2022, Hu et al., 2022, Vazquez-Bare, 2023, Gao and Ding, 2025]. The rate characterization here gives a finite-population minimax calibration for known bounded-degree low-order polynomial interference: the complexity contribution of the graph is summarized by the degree bound and the complete-block energy, while overlap across neighborhoods contributes one further degree factor. This connects local-interference intuition to a decision-theoretic benchmark for design mean squared error.

The complete-block calculations also relate to Riesz representer and pseudoinverse perspectives on linear estimation. The block-local problem can be viewed in the finite design space P_Z, the Bernoulli assignment law, through a polynomial subspace and a contrast functional, with the normalized representer h_d attaining the relevant block criterion. This geometry parallels the role of orthogonal scores and linear representers in semiparametric and design-based estimation, while remaining fully finite-dimensional under Bernoulli assignment. The exact complete-block local linear benchmark in the paper identifies the risk of block-local unbiased linear rules and supplies a sharp comparison class for SNIPE-type weights.

Recent work on unknown interference, covariate adjustment, and clustered or structured designs clarifies adjacent sources of statistical difficulty. Unknown or partially observed interference graphs introduce learning and robustness components [Yu et al., 2022, Eichhorn et al., 2024], covariate adjustment changes the information used by the estimator [Lin, 2013, Wang and Li, 2026], and clustered designs modify the assignment law and exposure variation [Basse and Airoldi, 2018, Basse and Feller, 2018, Harshaw et al., 2022]. The results here give the corresponding known-graph Bernoulli benchmark for bounded low-order polynomial interference, against which those additional design and information structures can be compared.

3 Setup and assumptions

All expectations in the main text are design expectations over the random assignment, with potential outcomes held fixed. Generic constants may depend on the fixed interaction order and assignment probability when explicitly subscripted by (β, p), and otherwise are numerical. Norms on finite coefficient arrays are the usual ℓ_1 or sup norms indicated by the displayed definitions. Asymptotic notation in the results keeps β and p fixed while allowing the population size and degree bound to vary; throughout this section, unit indices i, j range over the finite population V_n, and subset indices S range over finite subsets of the relevant interference neighborhood.

The design is the finite-population Bernoulli randomization framework in Assumption 1: the graph and potential-outcome schedule are fixed, and the assignment supplies the probability law.

Assumption 1. [`ass:bernoulli-design`] (Bernoulli assignment) *Let V_n be a finite population of $n := |V_n|$ units, indexed by $i, j \in V_n$, and let $p \in (0,1)$. For every unit $j \in V_n$,*

$$\Pr(Z_j = 1) = p,$$

and the variables $\{Z_j : j \in V_n\}$ are mutually independent. The assignment law P_Z on $\{0,1\}^{V_n}$ is

$$P_Z(z) = \prod_{j \in V_n} p^{z_j}(1-p)^{1-z_j},$$

and $\mathbb{E}_Z$ denotes expectation under P_Z.

Assumption 1 is the standard unit-level Bernoulli randomization condition used in low-order SNIPE analysis [Cortez-Rodriguez et al., 2023]. It fixes the common treatment probability and the product assignment law, so all risk calculations below are design-based calculations under P_Z.

Assumption 2. [`ass:bounded-degree`] (Bounded interference degree) *Let $G_n \subseteq V_n \times V_n$ be a directed interference graph, and write $N_i := \{j \in V_n : (j,i) \in G_n\}$ for the interference neighborhood of unit i. Self-loops are allowed and count in both the in-degree and out-degree. The interference neighborhoods N_i satisfy*

$$\max_{i \in V_n} |N_i| \le d \qquad \textit{and} \qquad \max_{j \in V_n} \big|\{i \in V_n : j \in N_i\}\big| \le d.$$

Assumption 2 is the standard bounded in- and out-degree neighborhood interference condition, with self-loops counted, as in Cortez-Rodriguez et al. [2023]. The in-degree bound limits how many assignment coordinates can enter a unit's potential outcome, while the out-degree bound controls how many unit-level scores can share any one assignment coordinate.

Assumption 3. [`ass:low-order`] (Low-order polynomial outcomes) *For every unit $i \in V_n$ and every finite subset $S \subseteq V_n$, the coefficient schedule satisfies*

$$c_{i,S} = 0 \qquad \textit{whenever} \qquad |S| > \beta.$$

Assumption 3 is the standard bounded polynomial interaction order condition [Cortez-Rodriguez et al., 2023]. It restricts the model class in Definition 6 to monomials of degree at most the interaction order β, so the relevant Bernoulli contrasts are finite-order neighborhood contrasts.

The model is an exact finite-population polynomial schedule: once the assignment vector z, graph, and coefficient schedule are fixed, $Y_i(z)$ is fixed. The reported risk is therefore the assignment-design mean squared error for that schedule. The procedure is permitted to use the interference graph G_n, the interaction order β, the assignment probability p, and the radius B; the radius determines the projection interval in the clipped estimators.

Known directed graphs with self-loops are natural when interference neighborhoods are measured before assignment and a unit's own treatment is allowed to affect its outcome. A common Bernoulli probability matches unit-randomized experiments run through a single assignment rule, while fixed β describes regimes where the analyst treats low-order interactions as the structural complexity constraint as n and d vary.

Assumption 4. [`ass:bounded-coefficient-mass`] (Bounded coefficient mass) *For every unit $i \in V_n$, the coefficient mass satisfies*

$$\sum_{S \subseteq N_i} |c_{i,S}| \le B.$$

Assumption 4 is the standard SNIPE $Y_{\max}$ coefficient-mass normalization condition [Cortez-Rodriguez et al., 2023]. The radius B fixes the unitwise scale of the coefficient schedule and therefore the scale of the minimax mean squared error.

The next definitions collect the model class, the estimand, and the estimators. The exposed-order notation records which Bernoulli contrast orders contribute to all-treated-versus-all-control estimation under the interaction and degree restrictions.

Definition 1. [`def:exposed-order`] (Exposed order $\bar{\beta}_d$ and $k_\star(d,\beta,p)$) *For $r \geq 1$, define the order-r all-one-versus-all-zero Bernoulli contrast coefficient by*

$$\Delta_r(p) := (1-p)^r - (-p)^r.$$

For every $d \geq 0$, define

$$\bar{\beta}_d := \min\{\beta, d\}.$$

In particular,

$$\bar{\beta}_0 = 0.$$

If $\bar{\beta}_d \geq 1$, define

$$k_\star(d,\beta,p) := \max\big\{r \in \{1,\ldots,\bar{\beta}_d\} : \Delta_r(p) \neq 0\big\}.$$

Since $\Delta_1(p) = 1$, this maximum is well defined. If $\bar{\beta}_d = 0$, equivalently $d = 0$ or $\beta = 0$, define

$$k_\star(d,\beta,p) := 0$$

as a totality convention, rather than as an attained interaction order. Thus the maximum case applies exactly when $d \geq 1$ and $\beta \geq 1$; subsequent statements take $\beta \geq 1$.

The order-r Bernoulli contrast coefficient $\Delta_r(p)$ is the all-one-versus-all-zero contrast induced by a centered r-fold Bernoulli monomial. The exposed order $\bar{\beta}_d$ combines the polynomial order and the neighborhood size, while $k_\star(d,\beta,p)$ identifies the largest exposed order with a nonzero Bernoulli contrast.

Definition 2. [`def:zero-degree-conventions`] (Zero-degree conventions $\binom{0}{0}$, g_0, A_0, and h_0) *For every $\beta \in \mathbb{N}$ and $p \in \mathbb{R}$, the degree-zero conventions are*

$$\bar{\beta}_0 = 0, \qquad k_\star(0,\beta,p) = 0, \qquad g_0 = 0, \qquad A_0 = 0, \qquad h_0 = 0, \qquad h_{0,T} = 0 \text{ for every } T \subseteq \{1,\ldots,0\}.$$

Here h_0 is imposed directly as the zero representer, rather than formed as the quotient g_0/A_0.

These conventions make the degree-zero boundary compatible with the score and energy notation used for positive degree. They keep later statements total in d while the substantive interaction results take $\beta \geq 1$.

Definition 3. [`def:graph-class`] (Graph class $\mathcal{G}_{n,d}$)

$$\mathcal{G}_{n,d} := \left\{G_n \in 2^{V_n \times V_n} : G_n \text{ satisfies Assumption 2}\right\}.$$

The graph class $\mathcal{G}_{n,d}$ collects the directed interference graphs satisfying the degree restriction in Assumption 2.

Definition 4. [`def:coefficient-class`] (Coefficient class $\mathcal{C}_{n,d,\beta}(B)$)

$$\mathcal{C}_{n,d,\beta}(B) := \{c = (c_{i,S})_{i,S} : c_{i,S} = 0 \text{ whenever } S \not\subseteq N_i,\ c \text{ satisfies Assumptions 3and 4}\}.$$

The coefficient class $\mathcal{C}_{n,d,\beta}(B)$ combines neighborhood support, low-order polynomial structure, and the unitwise ℓ_1 envelope. It is the coefficient-mass class used for the first minimax risk.

Definition 5. [`def:potential-outcome`] (Potential outcomes $Y_i(z)$ and observed outcomes Y_i^{obs}) *For a coefficient schedule c and an assignment $z \in \{0,1\}^{V_n}$, the potential outcome of unit $i \in V_n$ is*

$$Y_i(z) := \sum_{S \subseteq N_i} c_{i,S} \prod_{j \in S} z_j.$$

Evaluating the schedule at the realized assignment gives the observed outcome $Y_i^{\mathrm{obs}} := Y_i(Z)$.

The potential outcome $Y_i(z)$ is a raw-monomial polynomial in the assignments inside the unit's neighborhood, and the observed outcome Y_i^{obs} is its value at the realized assignment.

Definition 6. [`def:model-class`] (Model class $\mathcal{M}_{n,d,\beta}(B)$) *Under Assumptions 2 to 4, with coefficient schedules supported on the in-neighborhoods in the sense that $c_{i,S} = 0$ whenever $S \not\subseteq N_i$, define*

$$\mathcal{M}_{n,d,\beta}(B) := \{(G_n, c) : G_n \in \mathcal{G}_{n,d},\ c \in \mathcal{C}_{n,d,\beta}(B)\}.$$

The model class $\mathcal{M}_{n,d,\beta}(B)$ pairs a known bounded-degree graph with an admissible low-order coefficient schedule. The estimator may use the graph, so the decision problem is graph-aware.

Definition 7. [`def:snipe-score`] (SNIPE score $g_{i,\beta,p}$)

$$g_{i,\beta,p}(z) := \sum_{r=1}^{\bar{\beta}_d} \frac{\Delta_r(p)}{[p(1-p)]^r} \sum_{\substack{S \subseteq N_i \\ |S|=r}} \prod_{j \in S} (z_j - p).$$

The unit-level score $g_{i,\beta,p}$ is built from centered Bernoulli monomials on the interference neighborhood. Its coefficients use the contrast factors in Definition 1, matching the all-treated-versus-all-control target for low-order monomials. When $|N_i| < \bar{\beta}_d$, the inner subset sum is empty for every order $r > |N_i|$. Equivalently, the same score can be read as summing over $1 \le r \le \min\{\beta, |N_i|\}$ for unit i.

Definition 8. [`def:snipe-estimator`] *For each unit $i \in V_n$, the observed outcome under the realized assignment is*

$$Y_i^{\mathrm{obs}} := Y_i(Z).$$

The unprojected SNIPE estimator is

$$\widehat{\tau}_n^{\mathrm{SNIPE}} := \frac{1}{n} \sum_{i \in V_n} Y_i^{\mathrm{obs}} g_{i,\beta,p}(Z).$$

The projected SNIPE estimator is

$$\widehat{\tau}_n^{\mathrm{up}} := \Pi_{[-B,B]} \left(\widehat{\tau}_n^{\mathrm{SNIPE}} \right),$$

where $\Pi_{[-B,B]}$ is Euclidean projection onto $[-B, B]$.

The estimator $\widehat{\tau}_n^{\mathrm{SNIPE}}$ averages observed outcomes weighted by their neighborhood scores. The projected version $\widehat{\tau}_n^{\mathrm{up}}$ keeps the estimate in the natural coefficient-mass range for the total effect.

Definition 9. [`def:minimax-risk`] *The finite-population all-treated-versus-all-control effect is*

$$\tau_n := \frac{1}{n} \sum_{i \in V_n} \big(Y_i(\mathbf{1}) - Y_i(\mathbf{0})\big),$$

where $\mathbf{1}$ *and* $\mathbf{0}$ *denote the all-one and all-zero assignments. The minimax risk is*

$$R_n^\star(d, \beta, p, B) := \inf_{\widehat{\tau}_n} \sup_{(G_n, c) \in \mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z \left[(\widehat{\tau}_n - \tau_n)^2 \right].$$

The infimum ranges over estimators $\widehat{\tau}_n$ *that are measurable functions of* G_n*,* Z*, and* $\{Y_i^{\mathrm{obs}} : i \in V_n\}$*.*

The estimand τ_n, the finite-population all-treated-versus-all-control total treatment effect, is evaluated on the fixed potential-outcome schedule. The risk $R_n^\star(d, \beta, p, B)$ is the design-based minimax mean squared error over graph-aware estimators. Because the supremum ranges over the graph class $\mathcal{G}_{n,d}$, the frontier characterizes the hardest bounded-degree graph; for a given graph, achievable risk is governed by that graph's own local energies and overlap counts, which the upper bound tracks unit by unit.

The block quantities used in the main theorems are the complete-block analogues of the unit score. They isolate the local Bernoulli energy before network overlap contributes additional degree dependence.

Definition 10. [`def:block-score-energy`] (Block score energy g_d, A_d, h_d) *For* $d \geq 1$ *and a* d*-coordinate Bernoulli assignment vector* z*, define*

$$g_d(z) := \sum_{r=1}^{\bar{\beta}_d} \frac{\Delta_r(p)}{[p(1-p)]^r} \sum_{\substack{S \subseteq \{1,\ldots,d\} \\ |S| = r}} \prod_{j \in S} (z_j - p).$$

The complete-block score energy is

$$A_d := \sum_{r=1}^{\bar{\beta}_d} \binom{d}{r} \frac{\Delta_r(p)^2}{[p(1-p)]^r}.$$

The normalized block representer is

$$h_d := \frac{g_d}{A_d},$$

with raw-monomial expansion

$$h_d(z) = \sum_{T \subseteq \{1,\ldots,d\}} h_{d,T} \prod_{j \in T} z_j.$$

Here g_d is the complete-block score on d Bernoulli coordinates, A_d is its design second moment, and h_d is the normalized representer used in the block calculations. The raw-monomial coefficients $h_{d,T}$ also enter the least-favourable block construction below.

The paper also studies a uniformly bounded-outcome class. It uses the same graph and low-order structure, but the envelope is imposed directly on the potential-outcome functions.

Definition 11. [`def:bounded-outcome-coefficient-class`] (Bounded-outcome coefficient class $\mathcal{C}^{\infty}_{n,d,\beta}(B)$)

$$\mathcal{C}^{\infty}_{n,d,\beta}(B) := \left\{ c : c_{i,S} = 0 \text{ whenever } S \not\subseteq N_i, \quad c_{i,S} = 0 \text{ whenever } |S| > \beta, \quad \sup_{z \in \{0,1\}^{V_n}} |Y_i(z)| \le B \text{ for every } i \in V_n \right.$$

The low-order restriction is the one in Assumption 3.

The class $\mathcal{C}^{\infty}_{n,d,\beta}(B)$ places the radius B on the realized range of each unit's potential-outcome polynomial. This gives a parallel bounded class for comparing coefficient-mass and outcome-bounded formulations.

Definition 12. [`def:bounded-outcome-model-class`] (Bounded-outcome model class $\mathcal{M}^{\infty}_{n,d,\beta}(B)$)

$$\mathcal{M}^{\infty}_{n,d,\beta}(B) := \left\{ (G_n, c) : G_n \in \mathcal{G}_{n,d}, \quad c \in \mathcal{C}^{\infty}_{n,d,\beta}(B) \right\}.$$

The graph component satisfies Assumption 2, and the coefficient component satisfies the low-order condition in Assumption 3.

The model class $\mathcal{M}^{\infty}_{n,d,\beta}(B)$ is the bounded-outcome counterpart of Definition 6. Both classes retain the same known graph and Bernoulli assignment structure.

Definition 13. [`def:two-class-minimax-risks`] (Two minimax risks $R^{\star}_{n,\ell_1}$ and $R^{\star}_{n,\infty}$)

$$R^{\star}_{n,\ell_1}(d, \beta, p, B) := R^{\star}_n(d, \beta, p, B),$$

and

$$R^{\star}_{n,\infty}(d, \beta, p, B) := \inf_{\widehat{\tau}_n} \sup_{(G_n, c) \in \mathcal{M}^{\infty}_{n,d,\beta}(B)} \mathbb{E}_Z\left[(\widehat{\tau}_n - \tau_n)^2 \right],$$

where $\widehat{\tau}_n$ is measurable with respect to $(G_n, Z, \{Y_i^{\mathrm{obs}} : i \in V_n\})$.

The notation $R^{\star}_{n,\ell_1}(d, \beta, p, B)$ names the coefficient-mass minimax risk, while $R^{\star}_{n,\infty}(d, \beta, p, B)$ names the bounded-outcome minimax risk. The estimator class is the same graph-aware design-based class in both risks.

Definition 14. [`def:bounded-outcome-clipped-snipe`] (Bounded-outcome clipped SNIPE $\widehat{\tau}^{\mathrm{up}}_{n,\infty}$)

$$\widehat{\tau}^{\mathrm{up}}_{n,\infty} := \Pi_{[-2B,2B]}\left(\widehat{\tau}^{\mathrm{SNIPE}}_n \right),$$

where $\Pi_{[-2B,2B]}$ is Euclidean projection onto $[-2B, 2B]$.

The bounded-outcome clipped estimator $\widehat{\tau}^{\mathrm{up}}_{n,\infty}$ uses the same unprojected SNIPE statistic as Definition 8 and projects it onto the total-effect range induced by uniformly bounded outcomes.

Finally, the lower-bound and block-calibration arguments use a complete-block least-favourable family. The construction fixes the block count, active population fraction, cosine-squared baseline density, and score-aligned perturbation used later in the proofs.

Definition 15. [`def:block-family`] (Block priors Π_σ, $\delta(n, d, \beta, B, p)$, and H^2) *For $n, d, \beta \in \mathbb{N}$, $B > 0$, $p \in \mathbb{R}$, $1 \le d \le n$, and $\sigma \in \{-1, 1\}$, define*

$$m := \left\lfloor \frac{n}{d} \right\rfloor, \qquad q := md, \qquad \rho := \frac{q}{n}.$$

On the first q units, form consecutive blocks V_b, each of size d. Each block contains every within-block directed arrow, including loops, and the remaining $n-q$ units are isolated.

Define

$$s := \frac{B}{2}$$

and

$$f_s(u) := s^{-1}\cos^2\left(\frac{\pi u}{2s}\right)\mathbf{1}\{|u| \le s\}.$$

Let $U_1, \dots, U_m$ be independent real draws with density f_s. Define

$$H_{\beta,p} := \sup_{r\ge 1}\sum_T |h_{r,T}|$$

and

$$\delta(n,d,\beta,B,p) := B\min\left\{(2H_{\beta,p})^{-1},(4\pi)^{-1}\right\}\min\left\{1,\sqrt{\frac{A_d}{m}}\right\}.$$

Write $\delta := \delta(n,d,\beta,B,p)$. The prior Π_σ is the pushforward of the law of $(U_1,\dots,U_m)$ under the following coefficient-schedule map. For each active unit $i \in V_b$ and within-block subset $T \subseteq V_b$,

$$c_{i,T} := U_b\mathbf{1}\{T=\varnothing\} + \sigma\delta h_{d,T}, \qquad |T| \le k_\star.$$

All coefficients for subsets not contained in the active block, and all coefficients of inactive units, are zero.

Definition 15 is the least-favourable block construction used to state the complete-block lower-bound construction. The quantities m, q, and ρ describe how many units are active in complete directed d-blocks; f_s supplies a compactly supported baseline density; and the sign-indexed prior Π_σ perturbs each active block in the direction of the normalized representer h_d.

3.1 Discussion of the assumptions

Bounded degree with self-loops retained asks that each unit interfere with, and be interfered with by, at most d others; it is satisfied by bounded-degree contact networks, by classroom, village, or market designs with capped group sizes, and by any graph obtained by truncating a sparse network at a known degree cap. It controls topology, while the coefficient-mass bound controls interference strength. A common assignment probability p is normally a design choice: the experimenter selects it, and the frontier's dependence on p enters through the Bernoulli contrasts $\Delta_r(p)$, so the fair-coin case is a specialization of the common-probability design. The low-order condition fixes the functional form; the exposed order $k_\star$ then records which interaction orders the all-treated-versus-all-control contrast can detect.

4 Main results

The main results characterize the design-based mean squared error in terms of the complete-block score energy A_d and the neighborhood-overlap factor d. Before the formal statements, the theorem hierarchy is:

- **Coefficient-mass minimax frontier.** Theorem 1 gives matching lower and clipped-SNIPE upper bounds over $\mathcal{M}_{n,d,\beta}(B)$ at the exposed-binomial scale.

- **Bounded-outcome extension.** Theorem 2 compares the coefficient-mass and uniformly bounded-outcome risks and gives the attaining clipped estimator for the bounded-outcome envelope.
- **Exact complete-block SNIPE risk.** Theorem 2 also evaluates canonical unprojected SNIPE exactly on complete directed blocks and globally for that estimator when $d \mid n$.
- **Local-linear benchmark.** Theorem 3 solves the block-local unbiased linear problem on fixed complete directed blocks and characterizes its blockwise risk criterion.

The following map keeps those roles visible while also locating the fair-coin specialization.

Role	**What is calibrated**	**Formal location**
Minimax frontier	coefficient-mass lower bound and clipped-SNIPE upper bound	Theorem 1
Two-envelope comparison	coefficient-mass and uniformly bounded-outcome classes	Theorem 2
Estimator-risk equality	canonical unprojected SNIPE on complete directed blocks	Theorem 2
Linear benchmark	block-local design-unbiased linear weights	Definitions 16 to 18 and Theorem 3
Fair-coin calibration	odd-order block energy and first-order frontier at $p = 1/2$	Theorem 4

Theorem 1. [`thm:degree-frontier`] (Degree frontier bounds) *Fix an interaction order $\beta \in \mathbb{N}$ and an assignment probability $p \in \mathbb{R}$. Assume:*

- ***(Order.)*** $1 \le \beta$.
- ***(Design probability.)*** $0 < p < 1$.

Then there exist constants $c_{\beta,p}, C_{\beta,p} \in \mathbb{R}$ such that

$$0 < c_{\beta,p} \le C_{\beta,p}.$$

For every $n, d \in \mathbb{N}$ and $B \in \mathbb{R}$ satisfying $1 \le n$, $d \le n$, and $0 \le B$, with $k_\star(d,\beta,p)$ as in Definition 1 and $R_n^\star(d,\beta,p,B)$ as in Definition 9, the following hold:

$$c_{\beta,p} B^2 \min\left\{1, \frac{d\binom{d}{k_\star(d,\beta,p)}}{n}\right\} \le R_n^\star(d,\beta,p,B).$$

Moreover,

$$R_n^\star(d,\beta,p,B) \le \sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}_n^{\mathrm{up}} - \tau_n\right)^2\right] \le C_{\beta,p} B^2 \min\left\{1, \frac{d\binom{d}{k_\star(d,\beta,p)}}{n}\right\},$$

where $\mathcal{M}_{n,d,\beta}(B)$ is the model class in Definition 6 and $\widehat{\tau}_n^{\mathrm{up}}$ is SNIPE projected to $[-B, B]$. The unprojected SNIPE estimator also satisfies

$$\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}_n^{\mathrm{SNIPE}} - \tau_n\right)^2\right] \le C_{\beta,p} B^2 \frac{d\binom{d}{k_\star(d,\beta,p)}}{n}.$$

For every $(G,c) \in \mathcal{M}_{n,d,\beta}(B)$ and unit i, writing $N_i = \{j : (j,i) \in G\}$,

$$\sum_{r=1}^{\min\{\beta,|N_i|\}} \binom{|N_i|}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r} \le A_d,$$

where

$$A_d = \sum_{r=1}^{\min\{\beta,d\}} \binom{d}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r}.$$

For every $(G,c) \in \mathcal{M}_{n,d,\beta}(B)$, every $r \in \mathbb{N}$, and every unit i, if $1 \le r$, then

$$\sum_{\ell \in V_n} \binom{|N_i \cap N_\ell|}{r} = \sum_{\substack{S \subseteq N_i \\ |S|=r}} \#\{\ell \in V_n : S \subseteq N_\ell\} \le d\binom{|N_i|}{r} \le d\binom{d}{r}.$$

Finally, if $0 < B$ and $1 \le d$, let

$$m = \left\lfloor \frac{n}{d} \right\rfloor, \qquad \rho = \frac{md}{n}, \qquad \delta = \delta(n,d,\beta,B,p).$$

For every $U : \{1,\dots,m\} \to \mathbb{R}$ satisfying $|U_b| \le B/2$ for all b, there are models M_+ and M_- in $\mathcal{M}_{n,d,\beta}(B)$ whose graph is the complete d-block graph, whose coefficient schedules are the corresponding block schedules with signs $+1$ and -1, and whose total treatment effects obey

$$\tau_n(M_+) = \rho\delta, \qquad \tau_n(M_-) = -\rho\delta.$$

The two block-prior densities with signs $+1$ and -1, relative to the block dominating measure, satisfy

$$H^2(\Pi_+, \Pi_-) \le \frac{4\pi^2 m\delta^2}{B^2 A_d}.$$

The active fraction and block energy satisfy

$$\frac{1}{2} \le \rho \le 1, \qquad \frac{A_d}{m} = \frac{dA_d/n}{\rho}.$$

The rate in Theorem 1 separates two forces. The local Bernoulli part is the complete-block energy A_d, equivalently the exposed binomial scale in the theorem statement; the global network part is the single additional overlap charge d, coming from the number of unit scores that can share a given r-fold assignment set. The projection onto $[-B, B]$ supplies the saturated branch of the bound, while the unprojected estimator carries the linear-regime variance scale.

The next theorem states the simultaneous comparison between the coefficient-mass class and the uniformly bounded-outcome class. It uses the two minimax risks from Definition 13 and the bounded-outcome projection from Definition 14. The statement should be read in five parts: inclusion of the coefficient-mass class in the bounded-outcome class, the two-class minimax frontier, SNIPE unbiasedness and risk bounds, exact complete-block risks for unprojected SNIPE, and the block lower-bound construction with the exposed-binomial energy comparison.

Theorem 2. [`thm:bounded-outcome-degree-frontier`] (Bounded-outcome degree frontier) *Fix an interaction order $\beta \ge 1$ and a treatment probability $p \in (0,1)$. Then there are constants $c_{\beta,p}$ and $C_{\beta,p}$ such that $0 < c_{\beta,p} \le C_{\beta,p}$ and the following statements hold for every $n, d \in \mathbb{N}$, every $B \in \mathbb{R}$, and every finite design D on $\{0,1\}^{V_n}$ satisfying the itemized conditions:*

- ***(Population and degree.)*** *The population has size $n \geq 1$, $V_n = \{1, \dots, n\}$, and the degree index satisfies $d \leq n$.*
- ***(Radius.)*** *The radius satisfies $B \geq 0$.*
- ***(Design.)*** *The design D is the product Bernoulli design with common treatment probability p, as in Assumption 1.*

Let

$$A_d = \sum_{r=1}^{\bar{\beta}_d} \binom{d}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r},$$

with $\bar{\beta}_d$ and $k_\star(d, \beta, p)$ as in Definition 1. The coefficient-mass model $\mathcal{M}_{n,d,\beta}(B)$ of Definition 6 is contained in the uniformly bounded-outcome model $\mathcal{M}^\infty_{n,d,\beta}(B)$ of Definition 12. If $B > 0$ and $d \geq 1$, this inclusion is strict: some graph and bounded-outcome coefficient schedule in $\mathcal{M}^\infty_{n,d,\beta}(B)$ has no realization with the same graph and coefficient schedule in $\mathcal{M}_{n,d,\beta}(B)$.

The minimax risks of Definition 9 satisfy

$$c_{\beta,p} B^2 \min\left\{1, \frac{dA_d}{n}\right\} \leq R^\star_{n,\ell_1}(d, \beta, p, B) = R^\star_n(d, \beta, p, B) \leq R^\star_{n,\infty}(d, \beta, p, B).$$

For the bounded-outcome clipped SNIPE estimator $\widehat{\tau}^{\mathrm{up}}_{n,\infty}$ of Definition 14,

$$R^\star_{n,\infty}(d, \beta, p, B) \leq \sup_{(G,c) \in \mathcal{M}^\infty_{n,d,\beta}(B)} \mathbb{E}_D\left[\left(\widehat{\tau}^{\mathrm{up}}_{n,\infty} - \tau_n(G, c)\right)^2\right] \leq C_{\beta,p} B^2 \min\left\{1, \frac{dA_d}{n}\right\}.$$

The coefficient-class clipped SNIPE estimator $\widehat{\tau}^{\mathrm{up}}_n$ also satisfies

$$\sup_{(G,c) \in \mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_D\left[\left(\widehat{\tau}^{\mathrm{up}}_n - \tau_n(G, c)\right)^2\right] \leq C_{\beta,p} B^2 \min\left\{1, \frac{dA_d}{n}\right\}.$$

For every $(G, c) \in \mathcal{M}_{n,d,\beta}(B)$, the unprojected SNIPE estimator $\widehat{\tau}^{\mathrm{SNIPE}}_n$ is D-unbiased for $\tau_n(G, c)$. The same D-unbiasedness holds for every $(G, c) \in \mathcal{M}^\infty_{n,d,\beta}(B)$. Its worst-case mean squared error obeys

$$\sup_{(G,c) \in \mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_D\left[\left(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n(G, c)\right)^2\right] \leq \frac{B^2 dA_d}{n},$$

and

$$\sup_{(G,c) \in \mathcal{M}^\infty_{n,d,\beta}(B)} \mathbb{E}_D\left[\left(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n(G, c)\right)^2\right] \leq \frac{B^2 dA_d}{n}.$$

For every unit i, on both $\mathcal{M}_{n,d,\beta}(B)$ and $\mathcal{M}^\infty_{n,d,\beta}(B)$, the local score energy is at most A_d.

The complete-block energy is comparable to the exposed binomial term:

$$c_{\beta,p}\, d \binom{d}{k_\star(d, \beta, p)} \leq dA_d \leq C_{\beta,p}\, d \binom{d}{k_\star(d, \beta, p)}.$$

If $d \geq 1$ and $d \mid n$, let $m = n/d$ and let the graph be the disjoint union of m complete directed d-blocks, including loops. Then canonical unprojected SNIPE has exact worst-case risk on the fixed block graph and globally on both classes:

$$\sup_{\mathcal{M}_{n,d,\beta}(B) \text{ on the fixed block graph}} \mathbb{E}_D\left[\left(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n\right)^2\right] = \frac{B^2 A_d}{m},$$

$$\sup_{\mathcal{M}^{\infty}_{n,d,\beta}(B) \textit{ on the fixed block graph}} \mathbb{E}_D\Big[\big(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n\big)^2\Big] = \frac{B^2 A_d}{m},$$

and

$$\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_D\Big[\big(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n(G,c)\big)^2\Big] = \sup_{(G,c)\in\mathcal{M}^{\infty}_{n,d,\beta}(B)} \mathbb{E}_D\Big[\big(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n(G,c)\big)^2\Big] = \frac{B^2 A_d}{m} = \frac{B^2 d A_d}{n}.$$

If $B > 0$ and $d \geq 1$, put

$$m = \lfloor n/d \rfloor, \qquad \rho = \frac{md}{n}, \qquad \delta = \delta(n, d, \beta, B, p) \;\; \textit{as in Definition 15.}$$

For every vector $U = (U_b)_{b=1}^m$ with $|U_b| \leq B/2$ for all b, there are two coefficient-mass models on the block graph, with signs $+1$ and -1, whose coefficient schedules are the corresponding block schedules and whose total treatment effects are

$$\tau_n^+ = \rho\delta, \qquad \tau_n^- = -\rho\delta.$$

For the two associated block priors Π_+ and Π_-,

$$H^2(\Pi_+, \Pi_-) \leq \frac{4\pi^2 m \delta^2}{B^2 A_d},$$

and the active share satisfies

$$\frac{1}{2} \leq \rho \leq 1, \qquad \frac{A_d}{m} = \frac{dA_d/n}{\rho}.$$

Theorem 2 shows that the outcome-bounded envelope shares the same design difficulty as the coefficient-mass envelope, up to constants depending only on (β, p). The strict inclusion clause clarifies that this is a genuine comparison between different bounded classes. The complete-block equality identifies the variance contribution $B^2 A_d/m = B^2 d A_d/n$ exactly for canonical unprojected SNIPE, so the global rate can be read as complete-block energy multiplied by the block-to-population conversion.

The block-local benchmark restricts attention to fixed complete directed block graphs, coefficient-mass schedules on those graphs, and linear block-measurable weights satisfying exact design-unbiasedness equations. Fix a sequence indexed by t with positive integers n_t and d_t satisfying $d_t \mid n_t$, put $m_t := n_t/d_t$, and let G_t be the complete directed d_t-block graph on V_{n_t}. The following definitions name the fixed-graph coefficient class, the admissible local linear estimators, and their minimax risk before the sharp benchmark theorem.

Definition 16. [`def:local-linear-class`] (Local class $\mathcal{M}_t(G_t)$) *Under Assumptions 3 and 4, on the fixed graph G_t, define*

$$\mathcal{M}_t(G_t) := \left\{ c : c_{i,S} = 0 \textit{ unless } S \subseteq N_i^{G_t} \textit{ and } |S| \leq \bar{\beta}_{d_t}, \quad \sum_{S\subseteq N_i^{G_t}} |c_{i,S}| \leq B \textit{ for every } i \in V_{n_t} \right\}.$$

The class in Definition 16 fixes the complete-block graph and varies only the admissible coefficient schedule. This isolates the cost of block-local unbiased weighting from graph selection or graph heterogeneity.

Definition 17. [`def:local-linear-estimator-class`] (Block-local estimators $\mathcal{E}_t^{\mathrm{loc,lin}}$) *For each unit $i \in V_{n_t}$, let $b(i)$ be the index of the complete block containing i, and let $Z_{b(i)}$ be the assignment vector restricted to that block:*

$$b(i) := \text{the unique } b \text{ such that } i \in V_b, \qquad Z_{b(i)} := (Z_j)_{j \in V_{b(i)}}.$$

Let $L^2(P_Z)$ denote the finite-design square-integrable functions under P_Z. Define

$$\mathcal{E}_t^{\mathrm{loc,lin}} := \left\{ \widehat{\tau}_{w,t} = \frac{1}{n_t} \sum_{i \in V_{n_t}} w_{i,t}(Z_{b(i)}) Y_i(Z) : \begin{array}{l} w_{i,t} \in L^2(P_Z) \text{ is measurable with respect to } Z_{b(i)}, \\ \mathbb{E}_Z[w_{i,t}] = 0, \\ \mathbb{E}_Z\left[w_{i,t} \prod_{j \in S} Z_j\right] = 1 \text{ for every nonempty } S \subseteq N_i^{G_t} \text{ with } |S| \le \bar{\beta}_{d_t} \end{array} \right.$$

The class in Definition 17 encodes block-locality and unbiasedness through moment equations against every nonempty raw monomial up to $\bar{\beta}_{d_t}$ in the unit's block. The finite-design space $L^2(P_Z)$ is the natural Hilbert space for these block weights under the Bernoulli assignment law from Assumption 1.

Definition 18. [`def:local-linear-risk`] (Local linear risk $R_t^{\mathrm{loc,lin}}$)

$$R_t^{\mathrm{loc,lin}} := \inf_{\widehat{\tau}_{w,t} \in \mathcal{E}_t^{\mathrm{loc,lin}}} \sup_{c \in \mathcal{M}_t(G_t)} \mathbb{E}_Z\left[(\widehat{\tau}_{w,t} - \tau_{n_t})^2\right].$$

The supremum ranges over coefficient schedules in $\mathcal{M}_t(G_t)$ on the fixed graph G_t.

The risk in Definition 18 is the complete-block linear benchmark against which SNIPE-type weights can be compared.

Theorem 3. [`thm:sharp-local-linear-constant-and-representers`] (Sharp local linear representers) *Fix an integer β, a common Bernoulli treatment probability p, and a radius B. Suppose that:*

- ***(Smoothness order.)*** *$\beta \ge 1$.*
- ***(Assignment probability.)*** *$p \in (0, 1)$.*
- ***(Radius.)*** *$B > 0$.*
- ***(Block sequence.)*** *For each index t, n_t and d_t are positive integers with $d_t \mid n_t$, and*

$$m_t := \frac{n_t}{d_t}.$$

- ***(Complete-block graph.)*** *For each t, G_t is the directed graph on V_{n_t} in which $j \to i$ exactly when j and i are active units in the same quotient class after division by d_t; equivalently, the active units form complete directed blocks with self-loops and inactive units are isolated.*
- ***(Assignment design.)*** *The assignment design D_t is the product Bernoulli design with common probability p, as in Assumption 1.*

Define the complete-block score energy by

$$A_d := \sum_{r=1}^{\bar{\beta}_d} \binom{d}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r}.$$

Then, for every t,

$$R_t^{\mathrm{loc,lin}} = \frac{B^2 A_{d_t}}{m_t} = \frac{B^2 d_t A_{d_t}}{n_t}.$$

For local linear weights $w \in \mathcal{E}_t^{\mathrm{loc,lin}}$, put $Z_S := \prod_{j\in S} Z_j$. For each active block V_b, define

$$\Psi_{b,t}(w) := \max_{\substack{(\varepsilon_i, S_i)_{i\in V_b} \\ \varepsilon_i \in \{-1,1\},\ S_i \subseteq V_b,\ |S_i| \le \bar{\beta}_{d_t}}} \mathbb{E}_Z\left[\left(\sum_{i\in V_b} \varepsilon_i \left\{w_{i,t}(Z_{b(i)}) Z_{S_i} - \mathbf{1}\{S_i \neq \varnothing\}\right\}\right)^2\right].$$

For every t and every $w \in \mathcal{E}_t^{\mathrm{loc,lin}}$,

$$\sup_{c\in\mathcal{M}_t(G_t)} \mathbb{E}_Z\left[(\widehat{\tau}_{w,t} - \tau_{n_t})^2\right] = \frac{B^2}{n_t^2} \sum_{b=1}^{m_t} \Psi_{b,t}(w),$$

and, for every active block V_b,

$$d_t^2 A_{d_t} \le \Psi_{b,t}(w).$$

For any sequence $w_t \in \mathcal{E}_t^{\mathrm{loc,lin}}$,

$$\frac{\sup_{c\in\mathcal{M}_t(G_t)} \mathbb{E}_Z\left[(\widehat{\tau}_{w_t,t} - \tau_{n_t})^2\right]}{B^2 A_{d_t}/m_t} \longrightarrow 1$$

if and only if

$$\frac{1}{m_t d_t^2 A_{d_t}} \sum_{b=1}^{m_t} \left\{\Psi_{b,t}(w_t) - d_t^2 A_{d_t}\right\} \longrightarrow 0.$$

Moreover, the representer condition

$$\frac{1}{n_t A_{d_t}} \sum_{i\in V_{n_t}} \mathbb{E}_Z\left[\left\{w_{i,t}(Z_{b(i)}) - g_{d_t}(Z_{b(i)})\right\}^2\right] \longrightarrow 0$$

implies

$$\frac{\sup_{c\in\mathcal{M}_t(G_t)} \mathbb{E}_Z\left[(\widehat{\tau}_{w_t,t} - \tau_{n_t})^2\right]}{B^2 A_{d_t}/m_t} \longrightarrow 1.$$

If $d_t \to \infty$, then there exists a sequence $w_t \in \mathcal{E}_t^{\mathrm{loc,lin}}$ such that

$$\frac{\sup_{c\in\mathcal{M}_t(G_t)} \mathbb{E}_Z\left[(\widehat{\tau}_{w_t,t} - \tau_{n_t})^2\right]}{B^2 A_{d_t}/m_t} \longrightarrow 1,$$

and, for all sufficiently large t,

$$\frac{1}{n_t A_{d_t}} \sum_{i\in V_{n_t}} \mathbb{E}_Z\left[\left\{w_{i,t}(Z_{b(i)}) - g_{d_t}(Z_{b(i)})\right\}^2\right] = 2.$$

For every permutation $\pi : V_{n_t} \to V_{n_t}$ that preserves block membership, meaning $b(\pi(i)) = b(i)$ for every i, the corresponding relabeled normalized average squared distance from the complete-block SNIPE score is also equal to 2 for all sufficiently large t.

[Theorem 3](#) gives an exact finite benchmark for local linear rules on complete blocks. $\Psi_{b,t}$ is the exact blockwise worst-case quadratic form: averaging it across blocks reproduces the worst-case risk, and its lower bound $d_t^2 A_{d_t}$ yields the sharp risk $B^2 A_{d_t}/m_t$. The representer condition shows that average $L^2(P_Z)$ closeness to the complete-block SNIPE score is sufficient for asymptotic optimality. The distance-2 construction matters because optimality can occur through the blockwise quadratic form even away from the representer in normalized $L^2(P_Z)$ distance.

The fair-coin specialization makes the exposed-order calculation transparent. At $p = 1/2$, even-order Bernoulli contrasts cancel and odd orders determine the energy; for first-order interference this gives the particularly simple d^2/n rate.

Theorem 4. [`thm:fair-coin-energy-frontier`] (Fair-coin energy frontier) *For the fair-coin design $p = 1/2$, the following two conclusions hold.*

- ***(Exact fair-coin contrasts.)*** *For every $n, d, \beta \in \mathbb{N}$ and $B \in \mathbb{R}$ with $n \geq 1$, $d \geq 1$, $d \leq n$, $\beta \geq 1$, and $B \geq 0$, the Bernoulli contrast coefficients satisfy, for every $r \geq 1$,*

$$\Delta_r(1/2) = \begin{cases} 2^{1-r}, & r \text{ odd}, \\ 0, & r \text{ even}. \end{cases}$$

  Consequently the complete-block score energy at interaction order β is

$$A_d = 4 \sum_{\substack{1 \leq r \leq \bar{\beta}_d \\ r \text{ odd}}} \binom{d}{r}, \qquad \bar{\beta}_d = \min\{\beta, d\},$$

  and the exposed order $k_\star(d, \beta, 1/2)$ from [Definition 1](#) is the largest odd integer at most $\bar{\beta}_d$, with value 0 when no such odd integer exists.

- ***(First-order frontier.)*** *There exist constants $c_{\text{lower}}, c_{\text{upper}} \in \mathbb{R}$ with*

$$0 < c_{\text{lower}} \leq c_{\text{upper}}$$

  such that, for every $n, d \in \mathbb{N}$ and $B \in \mathbb{R}$ with $n \geq 1$, $d \geq 1$, $d \leq n$, and $B \geq 0$, the order-one energy is

$$A_d = 4d,$$

  and the minimax risks of [Definition 9](#) satisfy

$$c_{\text{lower}} B^2 \min\left\{1, \frac{d^2}{n}\right\} \leq R_n^\star(d, 1, 1/2, B),$$

$$R_n^\star(d, 1, 1/2, B) = R_{n,\ell_1}^\star(d, 1, 1/2, B) \leq R_{n,\infty}^\star(d, 1, 1/2, B) \leq c_{\text{upper}} B^2 \min\left\{1, \frac{d^2}{n}\right\}.$$

  If additionally $d \mid n$, then the unprojected SNIPE estimator $\widehat{\tau}_n^{\text{SNIPE}}$ has exact worst-case mean squared error

$$\sup_{M \in \mathcal{M}_{n,d,1}(B)} \mathbb{E}_{1/2,M}\left[\left(\widehat{\tau}_n^{\text{SNIPE}} - \tau_n\right)^2\right] = \frac{4B^2 d^2}{n},$$

  and, over the uniformly bounded-outcome class,

$$\sup_{M \in \mathcal{M}_{n,d,1}^\infty(B)} \mathbb{E}_{1/2,M}\left[\left(\widehat{\tau}_n^{\text{SNIPE}} - \tau_n\right)^2\right] = \frac{4B^2 d^2}{n}.$$

Theorem 4 gives the most legible additive-interference calibration of the general theory. For $p = 1/2$ and $\beta = 1$, the energy identity $A_d = 4d$ turns the general scale $B^2 \min\{1, dA_d/n\}$ into $B^2 \min\{1, d^2/n\}$, with exact unprojected SNIPE risk $4B^2d^2/n$ on complete blocks when $d \mid n$. This benchmark is the known-graph Bernoulli counterpart to problems where the interference structure, clustering, or adjustment information is part of the statistical design [Yu et al., 2022, Eichhorn et al., 2024, Wang and Li, 2026].

5 Discussion and extensions

The rate characterization in Theorems 1 and 2 separates two sources of design difficulty. The first is the local Bernoulli score energy A_d, which is determined by the assignment probability p, the exposed interaction order, and the number of coordinates in one neighborhood. The second is the additional out-degree overlap charge d, which enters because one assignment coordinate can appear in the scores of as many as d units under Assumption 2. Together they yield the scale

$$B^2 \min\left\{1, \frac{dA_d}{n}\right\},$$

with the equivalent exposed-order form

$$B^2 \min\left\{1, \frac{d\binom{d}{k_\star(d,\beta,p)}}{n}\right\}$$

up to constants depending on (β, p). The constants obtained from the block-energy comparison are interpreted for a fixed (β, p) with $p \in (0, 1)$. Under that fixed pair, the frontier is uniform in n, d, and B.

For design planning, A_d is the local price of extracting the all-treated-versus-all-control contrast from Bernoulli variation inside one neighborhood, while dA_d/n is the population-level burden after neighborhood overlap. The linear branch $dA_d/n \ll 1$ is the regime where unprojected SNIPE variance is informative. In the stated bounded classes, when dA_d/n is bounded away from zero at the saturation scale, the minimax mean squared error remains of order B^2 up to constants depending on (β, p).

The energy A_d gives the local price of recovering the all-treated-versus-all-control contrast from Bernoulli variation within a neighborhood. Its definition in Definition 10 sums the squared Bernoulli contrast coefficients across raw monomials up to the available order $\bar{\beta}_d$. The comparison in Theorem 2 shows that, for fixed (β, p), this energy has the same order as $\binom{d}{k_\star(d,\beta,p)}$. Thus the exposed nonzero contrast order determines how the local polynomial dimension enters the risk, while the assignment probability determines which orders are visible through the coefficients $\Delta_r(p)$.

The overlap factor has a different interpretation. It is a graph-combinatorial charge rather than a local representer charge. The covariance terms in SNIPE are organized by common assignment coordinates across neighborhoods, and Theorem 1 bounds the number of such overlaps by $d\binom{d}{r}$ at each order r. This is the point at which bounded out-degree matters for mean squared error: it controls how many unit-level score contributions can share a given monomial. Related local-interference analyses also emphasize neighborhood growth and overlap as central design features [Sävje et al., 2021, Leung, 2022, Hu et al., 2022, Gao and Ding, 2025]; the characterization here converts that feature into the finite-population minimax factor dA_d/n for known bounded-degree low-order polynomial interference.

Complete directed blocks provide the sharp calibration for the same expression. In the block construction of Definition 15, each active unit has the full d-coordinate neighborhood, and the normalized representer direction produces least-favourable perturbations whose separation and affinity are governed by A_d/m. When $d \mid n$, Theorem 2 gives the exact unprojected SNIPE worst-case risk $B^2 A_d/m = B^2 d A_d/n$ on complete blocks. The complete-block calculation matches the global upper bound's degree dependence and gives the exact unprojected-SNIPE worst-case risk on the block graph.

The local linear benchmark in Theorem 3 gives a complementary finite-dimensional interpretation. On fixed complete directed block graphs under the coefficient-mass schedule class, the minimax risk over block-local unbiased linear estimators is exactly $B^2 A_{d_t}/m_t$. This exact constant characterizes the block-local unbiased linear class on fixed complete directed blocks under the coefficient-mass schedule class: it optimizes over weights satisfying the block-local moment equations against every nonempty raw monomial up to $\bar{\beta}_{d_t}$. The clipped estimator, which uses projection, is the procedure that attains the saturated branch of the bounded-class frontier. The criterion $\Psi_{b,t}(w)$ identifies the blockwise extremal quantity that any proposed local linear weight must control, and the representer condition gives a sufficient route to the same benchmark through average $L^2(P_Z)$ closeness to the complete-block SNIPE score. This places SNIPE within a broader local Riesz geometry while preserving the exact blockwise minimax constant.

The comparison between the coefficient-mass and uniformly bounded-outcome classes in Theorem 2 shows that the same degree dependence governs both boundedness formulations. The bounded-outcome class is larger, but the clipped SNIPE estimator projected to $[-2B, 2B]$ attains the same $B^2 \min\{1, dA_d/n\}$ scale. This matters for applications in which the primitive restriction is naturally stated as a uniform potential-outcome envelope rather than an ℓ_1 coefficient-mass envelope: the design-based rate remains calibrated by the same local energy and overlap factor.

The fair-coin specialization in Theorem 4 makes the role of assignment probability especially transparent. At $p = 1/2$, the even-order Bernoulli contrasts vanish and the energy is the sum over odd orders. For first-order interference, $A_d = 4d$, so the minimax scale becomes $B^2 \min\{1, d^2/n\}$, with exact complete-block unprojected SNIPE risk $4B^2 d^2/n$ when $d \mid n$. This gives a simple known-graph Bernoulli benchmark for comparisons with clustered assignment, unknown interference, and covariate-adjusted designs [Eichhorn et al., 2024, Wang and Li, 2026].

For fixed (β, p) sequences with eventual positive degree and stable exposed order $k_\star$, the comparison $dA_d \asymp_{\beta,p} d\binom{d}{k_\star}$ identifies a sufficient and rate-equivalent vanishing condition for the displayed minimax scale. The heuristic condition is $d\binom{d}{k_\star} = o(n)$. Under that fixed-pair interpretation, this corresponds to degree growth on the order $d = o(n^{1/(k_\star+1)})$. For fair-coin assignment with first-order interactions ($\beta = 1$), the condition reads $d = o(\sqrt{n})$. At the saturation scale, for fixed (β, p), the frontier identifies B^2 as the constant-order minimax scale.

5.1 Limitations and future work

The results establish design-based minimax mean squared error rates for known bounded-degree low-order polynomial interference under common-probability Bernoulli assignment, together with exact complete-block constants for the stated local linear class.

The framework treats potential outcomes as exact low-order polynomials of the assignment vector, so the reported risk is assignment-design mean squared error for the fixed schedule. Open directions include additive idiosyncratic outcome noise and the associated variance floor, exact leading constants over all measurable estimators beyond the complete-block and local linear benchmarks, heterogeneous treatment probabilities, variance estimation and studentization for the SNIPE family, limit laws under growing-degree regimes, unknown or partially observed graphs, clustered assignment, and covariate adjustment. These extensions would connect the benchmark developed here to settings studied in recent work on local interference, graph uncertainty, and adjusted experimental estimators [Sävje et al., 2021, Leung, 2022, Hu et al., 2022, Gao and Ding, 2025, Eichhorn et al., 2024, Wang and Li, 2026].

Appendices

A Proofs and auxiliary lemmas

This appendix collects the auxiliary objects used in the minimax and complete-block arguments. The first group of statements isolates the one-block Hilbert-space calculations: a polynomial perturbation program for least-favourable directions, a dual weight program for block-local unbiased scores, and the representer identity tying both programs to the complete-block energy.

The block-local Riesz formulation works in the finite design space induced by the product Bernoulli law in Assumption 1. For a complete d-block, the low-order polynomial subspace $\mathcal{P}_d$ contains raw monomials up to the exposed order, and the contrast L_d evaluates the all-one-versus-all-zero difference on that block.

Definition 19. [`def:perturbation-program`] (Low-order block space $\mathcal{P}_d$, contrast L_d, and program $\mathsf{PerturbProg}_{d,\beta,p}$) *For $d \geq 1$, let $\mathcal{P}_d$ be the low-order polynomial subspace on d Bernoulli coordinates,*

$$\mathcal{P}_d := \left\{ h(z) = \sum_{\substack{S \subseteq \{1,\ldots,d\} \\ |S| \leq \bar{\beta}_d}} a_S \prod_{j \in S} z_j \right\}.$$

Define the all-one-versus-all-zero block contrast functional $L_d : \mathcal{P}_d \to \mathbb{R}$ by

$$L_d(h) := h(1, \ldots, 1) - h(0, \ldots, 0).$$

The perturbation program is

$$\mathsf{PerturbProg}_{d,\beta,p} := \min_{h \in \mathcal{P}_d} \mathbb{E}_Z[h(Z)^2]$$

subject to

$$L_d(h) = 1.$$

Definition 19 identifies the minimum-energy direction that moves the block contrast by one unit. This is the direction used by the least-favourable complete-block construction in Definition 15: controlling its $L^2(P_Z)$ energy controls the affinity between nearby prior-predictive laws while preserving a fixed separation in total treatment effects.

The dual calculation is a minimum-norm problem over square-integrable block weights. Its moment equations impose exact unbiasedness against every nonempty raw monomial that can enter a low-order block outcome.

Definition 20. [def:weight-program] (Weight program $\mathsf{WeightProg}_{d,\beta,p}$) *For $d \geq 1$, define*

$$\mathsf{WeightProg}_{d,\beta,p} := \min_{w \in L^2(P_Z)} \mathbb{E}_Z[w(Z)^2]$$

subject to

$$\mathbb{E}_Z[w(Z)] = 0$$

and, for every nonempty $S \subseteq \{1, \ldots, d\}$ with $|S| \leq \bar{\beta}_d$,

$$\mathbb{E}_Z\left[w(Z) \prod_{j \in S} Z_j\right] = 1.$$

Definition 20 is the block analogue of the SNIPE unbiasedness equations. The feasible weights reproduce the all-treated-versus-all-control contrast on every exposed raw monomial, so the value of the program is the smallest block variance compatible with those unbiasedness restrictions.

The next lemma supplies the exact representer calculation behind both programs. It also records the energy scale and a uniform raw-coefficient-mass bound for the normalized representer coefficient $h_{d,T}$, which is used in the block perturbations.

Lemma 1. [lem:block-energy-representer] (Block representer energy) *Fix an integer $\beta \geq 1$ and a Bernoulli treatment probability $p \in (0, 1)$. Let $k_\star(d, \beta, p)$ be the exposed order in Definition 1. For each $d \geq 1$, write*

$$A_d = \sum_{r=1}^{\bar{\beta}_d} \binom{d}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r}.$$

Let g_d denote the complete-block centered contrast score, $h_d := g_d/A_d$, and let $h_{d,T}$ be the coefficient of the raw monomial indexed by $T \subseteq \{1, \ldots, d\}$ in h_d.

Then there exist constants $c_1, c_2, H \in \mathbb{R}$, depending only on β and p, such that $0 < c_1 \leq c_2$, $0 < H$, and for every $d \geq 1$:

- ***(Energy scale.)***

$$c_1 \binom{d}{k_\star(d, \beta, p)} \leq A_d \leq c_2 \binom{d}{k_\star(d, \beta, p)}.$$

- ***(Product Bernoulli identities.)*** *For every finite design D on $\{0, 1\}^d$ that is product Bernoulli with common probability p as in Assumption 1, and every $f \in \mathcal{P}_d$,*

$$\mathbb{E}_D[g_d(Z)f(Z)] = L_d(f).$$

Moreover,

$$L_d(h_d) = 1, \qquad \mathbb{E}_D[h_d(Z)^2] = A_d^{-1}.$$

- ***(Perturbation program.)*** *The representer h_d is feasible for the perturbation program in Definition 19. For every feasible perturbation h,*

$$A_d^{-1} \leq \mathbb{E}_D[h(Z)^2],$$

with equality if and only if $h = h_d$. Consequently, for any proof witnesses of $0 \leq p$ and $p \leq 1$,

$$\mathsf{PerturbProg}_{d,\beta,p} = A_d^{-1}.$$

- ***(Weight program.)*** *For every weight w feasible for the weight program in Definition 20 under D,*

$$A_d \le \mathbb{E}_D[w(Z)^2],$$

with equality if and only if $w = g_d$. Consequently, for any proof witnesses of $0 \le p$ and $p \le 1$,

$$\mathsf{WeightProg}_{d,\beta,p} = A_d.$$

- ***(Coefficient mass.)***

$$\sum_{T \subseteq \{1,\dots,d\}} |h_{d,T}| \le H.$$

Proof of Lemma 1. Fix $\beta \ge 1$ and $p \in (0,1)$. For every integer $r \ge 1$ write

$$e_r(p) := \frac{\Delta_r(p)^2}{[p(1-p)]^r}, \qquad m_r(p) := \frac{|\Delta_r(p)|}{[p(1-p)]^r}(1+p)^r,$$

and define the three constants of the statement by

$$c_1 := \min\{e_r(p) : 1 \le r \le \beta,\ \Delta_r(p) \ne 0\}, \qquad c_2 := 2^\beta \sum_{r=1}^{\beta} e_r(p), \qquad H := c_1 + \frac{2^\beta \sum_{r=1}^{\beta} m_r(p)}{c_1}.$$

Since $\Delta_1(p) = (1-p) - (-p) = 1 \ne 0$, the index set defining c_1 contains $r = 1$ and is therefore a nonempty finite set; on it $\Delta_r(p)^2 > 0$ and $[p(1-p)]^r > 0$, so c_1 is a minimum of finitely many strictly positive numbers and $c_1 > 0$. Each $m_r(p) \ge 0$ for the same reason, so $H \ge c_1 > 0$.

Step 1 (exposed orders). Fix $d \ge 1$ and write $k_\star := k_\star(d,\beta,p)$. Since $\beta \ge 1$ and $d \ge 1$ we have $\bar\beta_d = \min\{\beta, d\} \ge 1$, and $\Delta_1(p) = 1 \ne 0$ shows that the set $\{r \in \{1,\dots,\bar\beta_d\} : \Delta_r(p) \ne 0\}$ of Definition 1 is nonempty. Hence $k_\star$ is the maximum of that set, so

$$1 \le k_\star \le \bar\beta_d, \qquad \Delta_{k_\star}(p) \ne 0, \qquad \text{and} \qquad r \le k_\star \ \text{ for every } r \in \{1,\dots,\bar\beta_d\} \text{ with } \Delta_r(p) \ne 0.$$

In particular $k_\star \le \beta$ and $k_\star \le d$.

Step 2 (binomial domination at fixed order). For all integers $r \le k \le d$ with $k \le \beta$,

$$\binom{d}{r} \le 2^\beta \binom{d}{k}.$$

Indeed, counting the pairs $R \subseteq K \subseteq \{1,\dots,d\}$ with $|R| = r$ and $|K| = k$ in two ways gives $\binom{d}{r}\binom{d-r}{k-r} = \binom{d}{k}\binom{k}{r}$; the left factor $\binom{d-r}{k-r} \ge 1$ and $\binom{k}{r} \le 2^k \le 2^\beta$.

Step 3 (energy scale). By Definition 10,

$$A_d = \sum_{r=1}^{\bar\beta_d} \binom{d}{r} e_r(p),$$

a sum of nonnegative terms. Discarding all summands except $r = k_\star$, which is legitimate by Step 1,

$$A_d \ \ge\ \binom{d}{k_\star} e_{k_\star}(p) \ \ge\ c_1 \binom{d}{k_\star},$$

the second inequality because $1 \le k_\star \le \beta$ and $\Delta_{k_\star}(p) \ne 0$, so $e_{k_\star}(p)$ is one of the numbers over which c_1 is the minimum. For the upper bound fix $r \in \{1,\dots,\bar\beta_d\}$. If $\Delta_r(p) = 0$ then $e_r(p) = 0$

and both sides of the next display vanish; if $\Delta_r(p) \neq 0$ then $r \le k_\star \le \min\{\beta, d\}$ by Step 1 and Step 2 applies. Either way

$$\binom{d}{r} e_r(p) \le 2^\beta \binom{d}{k_\star} e_r(p).$$

Summing over $r \in \{1, \dots, \bar{\beta}_d\}$ and then enlarging the index range from $\{1, \dots, \bar{\beta}_d\}$ to $\{1, \dots, \beta\}$, which only adds nonnegative terms since $\bar{\beta}_d \le \beta$,

$$A_d \le 2^\beta \binom{d}{k_\star} \sum_{r=1}^{\beta} e_r(p) = c_2 \binom{d}{k_\star}.$$

This is the energy-scale claim.

Step 4 ($c_1 \le c_2$). Apply Step 3 at $d = 1$ and set $K := \binom{1}{k_\star(1,\beta,p)}$. Step 1 gives $k_\star(1, \beta, p) \le \bar{\beta}_1 \le 1$, so $K \ge 1 > 0$, and

$$c_1 K \le A_1 \le c_2 K.$$

Dividing by $K > 0$ gives $c_1 \le c_2$.

Step 5 (design moments). Fix $d \ge 1$ and let D be a finite design on $\{0,1\}^d$ satisfying the common-p product Bernoulli condition of Assumption 1. Then D assigns to each $z \in \{0,1\}^d$ the mass $\prod_{j=1}^d p^{z_j}(1-p)^{1-z_j} > 0$, so for every $f : \{0,1\}^d \to \mathbb{R}$,

$$\mathbb{E}_D[f(Z)^2] = 0 \qquad \text{if and only if} \qquad f(z) = 0 \;\text{ for every } z \in \{0,1\}^d.$$

Moreover, factorizing the expectation over the independent coordinates, for every nonempty $S \subseteq \{1, \dots, d\}$ and every $T \subseteq \{1, \dots, d\}$,

$$\mathbb{E}_D\left[\prod_{j\in S}(Z_j - p)\prod_{j\in T} Z_j\right] = \begin{cases} [p(1-p)]^{|S|}\, p^{|T|-|S|}, & S \subseteq T, \\ 0, & \text{otherwise,} \end{cases}$$

$$\mathbb{E}_D\left[\prod_{j\in S}(Z_j - p)\prod_{j\in T}(Z_j - p)\right] = \begin{cases} [p(1-p)]^{|S|}, & S = T, \\ 0, & \text{otherwise.} \end{cases}$$

In the first identity a coordinate in $S \cap T$ contributes $\mathbb{E}[(Z_j - p)Z_j] = p(1-p)$, a coordinate of T outside S contributes $\mathbb{E}[Z_j] = p$, a coordinate of S outside T contributes $\mathbb{E}[Z_j - p] = 0$, and all remaining coordinates contribute 1; in the second identity a coordinate in $S \cap T$ contributes $\mathbb{E}[(Z_j - p)^2] = p(1-p)$ and a coordinate of the symmetric difference contributes 0.

Step 6 (the score represents L_d). For every integer $t \ge 1$,

$$\sum_{r=1}^{t} \binom{t}{r} \Delta_r(p)\, p^{t-r} = 1,$$

because $\sum_{r=0}^t \binom{t}{r}(1-p)^r p^{t-r} = \big[(1-p)+p\big]^t = 1$ and $\sum_{r=0}^t \binom{t}{r}(-p)^r p^{t-r} = (-p+p)^t = 0$, while the two $r = 0$ terms cancel since $\Delta_0(p) = 0$. Let $T \subseteq \{1, \dots, d\}$ be nonempty with $|T| \le \bar{\beta}_d$. Inserting the definition of g_d from Definition 10 into the first moment identity of Step 5, and using that exactly $\binom{|T|}{r}$ of the sets S with $|S| = r$ satisfy $S \subseteq T$,

$$\mathbb{E}_D\left[g_d(Z)\prod_{j\in T} Z_j\right] = \sum_{r=1}^{\bar{\beta}_d} \frac{\Delta_r(p)}{[p(1-p)]^r}\binom{|T|}{r}[p(1-p)]^r p^{|T|-r} = \sum_{r=1}^{\bar{\beta}_d}\binom{|T|}{r}\Delta_r(p)\, p^{|T|-r} = 1,$$

where the last equality is the previous display with $t = |T| \ge 1$: the summands with $|T| < r \le \bar\beta_d$ vanish because $\binom{|T|}{r} = 0$, so the sum runs effectively over $1 \le r \le |T|$. Taking $T = \varnothing$ in the same moment identity, where no nonempty S satisfies $S \subseteq T$,

$$\mathbb{E}_D[g_d(Z)] = 0.$$

Now $L_d\big(z \mapsto \prod_{j\in T} z_j\big) = 1$ for nonempty T and $L_d(1) = 0$, so the two displays state precisely that $\mathbb{E}_D[g_d(Z)f(Z)] = L_d(f)$ whenever f is a raw monomial of degree at most $\bar\beta_d$. Both sides are linear in f, and Definition 19 defines $\mathcal{P}_d$ as the span of these raw monomials; hence for every $f \in \mathcal{P}_d$,

$$\mathbb{E}_D[g_d(Z)f(Z)] = L_d(f).$$

Step 7 (energy identities and positivity). Evaluating g_d at the all-treated and at the all-control assignment, each of the $\binom{d}{r}$ sets S with $|S| = r$ contributes $(1-p)^r$ and $(-p)^r$ respectively, so

$$L_d(g_d) = \sum_{r=1}^{\bar\beta_d} \frac{\Delta_r(p)}{[p(1-p)]^r}\binom{d}{r}\big[(1-p)^r - (-p)^r\big] = \sum_{r=1}^{\bar\beta_d}\binom{d}{r} e_r(p) = A_d.$$

By the second moment identity of Step 5 the centered monomials appearing in g_d are pairwise orthogonal and the one indexed by S has second moment $[p(1-p)]^{|S|}$, so

$$\mathbb{E}_D[g_d(Z)^2] = \sum_{r=1}^{\bar\beta_d}\binom{d}{r}\frac{\Delta_r(p)^2}{[p(1-p)]^{2r}}[p(1-p)]^r = A_d.$$

Keeping only the $r = 1$ summand of A_d and using $\Delta_1(p) = 1$,

$$A_d \;\ge\; \binom{d}{1}\frac{\Delta_1(p)^2}{p(1-p)} = \frac{d}{p(1-p)} > 0.$$

Since $h_d = g_d/A_d$ with $A_d > 0$, dividing the two preceding identities by A_d and by A_d^2 gives

$$L_d(h_d) = 1, \qquad \mathbb{E}_D[h_d(Z)^2] = A_d^{-1}.$$

Step 8 (the score lies in $\mathcal{P}_d$). Expanding the product coordinate by coordinate, for every $S \subseteq \{1,\dots,d\}$ and every z,

$$\prod_{j\in S}(z_j - p) = \sum_{T\subseteq S}(-p)^{|S|-|T|}\prod_{j\in T} z_j.$$

Every raw monomial occurring on the right has degree $|T| \le |S| \le \bar\beta_d$ when $|S| \le \bar\beta_d$, hence lies in $\mathcal{P}_d$. As g_d is a finite linear combination of centered monomials with $|S| \le \bar\beta_d$, it follows that $g_d \in \mathcal{P}_d$ and therefore $h_d = g_d/A_d \in \mathcal{P}_d$. Together with $L_d(h_d) = 1$ from Step 7, h_d is feasible for the program of Definition 19.

Step 9 (perturbation optimality). Let $h \in \mathcal{P}_d$ satisfy $L_d(h) = 1$. Step 6 applied to $f = h$ gives $\mathbb{E}_D[g_d(Z)h(Z)] = 1$, and Step 7 gives $\mathbb{E}_D[g_d(Z)^2] = A_d$. Expanding the square,

$$\mathbb{E}_D\big[(A_d h(Z) - g_d(Z))^2\big] = A_d^2\,\mathbb{E}_D[h(Z)^2] - 2A_d\,\mathbb{E}_D[g_d(Z)h(Z)] + \mathbb{E}_D[g_d(Z)^2] = A_d^2\,\mathbb{E}_D[h(Z)^2] - A_d.$$

The left-hand side is nonnegative and $A_d > 0$, so $A_d^2\,\mathbb{E}_D[h(Z)^2] \geq A_d$, that is,

$$A_d^{-1} \leq \mathbb{E}_D[h(Z)^2].$$

Equality holds if and only if the displayed square has expectation zero, which by the faithfulness statement of Step 5 happens if and only if $A_d h(z) = g_d(z)$ for every z, that is, if and only if $h = h_d$. Conversely $h = h_d$ attains the value A_d^{-1} by Step 7.

Step 10 (weight optimality). Let w be feasible for Definition 20 under D, that is, $\mathbb{E}_D[w(Z)] = 0$ and $\mathbb{E}_D\big[w(Z)\prod_{j\in S} Z_j\big] = 1$ for every nonempty $S \subseteq \{1,\dots,d\}$ with $|S| \leq \bar\beta_d$. Since $L_d(1) = 0$ and $L_d\big(z \mapsto \prod_{j\in S} z_j\big) = 1$ for nonempty S, these constraints say that $\mathbb{E}_D[w(Z)f(Z)] = L_d(f)$ on every raw monomial f of degree at most $\bar\beta_d$; both sides are linear in f, and Definition 19 defines $\mathcal{P}_d$ as the span of those raw monomials, so for every $f \in \mathcal{P}_d$,

$$\mathbb{E}_D[w(Z)f(Z)] = L_d(f).$$

Taking $f = g_d$, which lies in $\mathcal{P}_d$ by Step 8, and using $L_d(g_d) = A_d$ and $\mathbb{E}_D[g_d(Z)^2] = A_d$ from Step 7,

$$\mathbb{E}_D\big[(w(Z) - g_d(Z))^2\big] = \mathbb{E}_D[w(Z)^2] - 2\,\mathbb{E}_D[w(Z)g_d(Z)] + \mathbb{E}_D[g_d(Z)^2] = \mathbb{E}_D[w(Z)^2] - A_d.$$

Nonnegativity of the left-hand side gives

$$A_d \leq \mathbb{E}_D[w(Z)^2],$$

and equality forces the displayed square to have expectation zero, hence $w = g_d$ by the faithfulness statement of Step 5. Conversely, the two identities of Step 6 are exactly the feasibility constraints for $w = g_d$, and $\mathbb{E}_D[g_d(Z)^2] = A_d$, so g_d is feasible and attains the value A_d.

Step 11 (program values). Both programs are infima of sets of expectations of squares, hence of sets of nonnegative reals, and both sets are nonempty: h_d is feasible for the perturbation program with value A_d^{-1} by Steps 7 and 8, and g_d is feasible for the weight program with value A_d by Step 10. Therefore

$$\mathsf{PerturbProg}_{d,\beta,p} \leq A_d^{-1}, \qquad \mathsf{WeightProg}_{d,\beta,p} \leq A_d,$$

while Steps 9 and 10 bound every value in the two feasible sets below by A_d^{-1} and by A_d respectively, giving the reverse inequalities. Hence

$$\mathsf{PerturbProg}_{d,\beta,p} = A_d^{-1}, \qquad \mathsf{WeightProg}_{d,\beta,p} = A_d.$$

Step 12 (coefficient envelope). Expanding each centered monomial of g_d by the identity of Step 8, collecting the coefficient of $\prod_{j\in T} z_j$, and dividing by A_d, the raw coefficients of h_d in Definition 10 are

$$h_{d,T} = A_d^{-1} \sum_{r=1}^{\bar\beta_d} \sum_{\substack{S\subseteq\{1,\dots,d\},\ |S|=r \\ T\subseteq S}} \frac{\Delta_r(p)}{[p(1-p)]^r} (-p)^{r-|T|}.$$

Applying the triangle inequality inside each $h_{d,T}$ and then exchanging the order of summation, so that each S with $|S| = r$ is paired with the subsets $T \subseteq S$,

$$\sum_{T\subseteq\{1,\dots,d\}} |h_{d,T}| \;\leq\; A_d^{-1} \sum_{r=1}^{\bar\beta_d} \sum_{\substack{S\subseteq\{1,\dots,d\} \\ |S|=r}} \sum_{T\subseteq S} \frac{|\Delta_r(p)|}{[p(1-p)]^r}\, p^{r-|T|} \;=\; A_d^{-1} \sum_{r=1}^{\bar\beta_d} \binom{d}{r} m_r(p),$$

using $\sum_{T\subseteq S} p^{|S|-|T|} = (1+p)^{|S|} = (1+p)^r$ and the fact that there are $\binom{d}{r}$ sets S with $|S| = r$. Exactly as in Step 3 — orders with $\Delta_r(p) = 0$ contribute zero, orders with $\Delta_r(p) \neq 0$ satisfy $r \le k_\star$ and hence $\binom{d}{r} \le 2^\beta \binom{d}{k_\star}$ by Step 2, and the index range may then be enlarged to $\{1, \dots, \beta\}$ —

$$\sum_{r=1}^{\bar\beta_d} \binom{d}{r} m_r(p) \le 2^\beta \binom{d}{k_\star} \sum_{r=1}^{\beta} m_r(p).$$

Combining this with the lower energy bound $c_1\binom{d}{k_\star} \le A_d$ of Step 3, with $A_d > 0$ and $\binom{d}{k_\star} > 0$,

$$\sum_{T\subseteq\{1,\dots,d\}} |h_{d,T}| \;\le\; \frac{2^\beta\binom{d}{k_\star}\sum_{r=1}^{\beta} m_r(p)}{A_d} \;\le\; \frac{2^\beta\sum_{r=1}^{\beta} m_r(p)}{c_1} \;\le\; H,$$

the last step because H exceeds that quotient by $c_1 > 0$. $\square$

Lemma 1 is the algebraic core of the appendix. It establishes that the normalized complete-block score is the unique minimum-energy perturbation with unit block contrast, while the unnormalized score is the unique minimum-energy unbiased weight. The same calculation gives $A_d \asymp_{\beta,p} \binom{d}{k_\star(d,\beta,p)}$, which is the local energy comparison used in Theorems 1 and 2.

The global SNIPE variance calculation needs one graph-combinatorial ingredient. The covariance terms are organized by shared assignment subsets across neighborhoods, and the next lemma converts those shared subsets into a single out-degree charge.

Lemma 2. [`lem:overlap-count`] (Overlap count bound) *Let V_n be the finite population and unit index set, let G be a directed interference graph on V_n, and write*

$$N_i = \{j \in V_n : G(j,i)\}.$$

Suppose:

- ***(Bounded degree.)*** *The graph G satisfies the bounded-degree condition with bound d in Assumption 2.*
- ***(Order.)*** *The integer r satisfies $r \ge 1$.*
- ***(Unit.)*** *The unit i belongs to V_n.*

Then

$$\sum_{l\in V_n} \binom{|N_i\cap N_l|}{r} = \sum_{\substack{S\subseteq N_i\\|S|=r}} \#\{l\in V_n : S\subseteq N_l\}$$

and

$$\sum_{\substack{S\subseteq N_i\\|S|=r}} \#\{l\in V_n : S\subseteq N_l\} \le d\binom{|N_i|}{r} \le d\binom{d}{r}.$$

Proof of Lemma 2. Let

$$\mathcal{F} := \{S\subseteq N_i : |S| = r\}$$

be the family of r-element subsets of the neighborhood N_i, so that the middle expression in the statement is $\sum_{S\in\mathcal{F}} \#\{l\in V_n : S\subseteq N_l\}$.

Step 1: a pointwise counting identity. Fix $l \in V_n$. A finite set S satisfies $S \subseteq N_i \cap N_l$ if and only if $S \subseteq N_i$ and $S \subseteq N_l$, so the r-element subsets of $N_i \cap N_l$ are exactly the members of $\mathcal{F}$ that are contained in N_l:

$$\{S \subseteq N_i \cap N_l : |S| = r\} = \{S \in \mathcal{F} : S \subseteq N_l\}.$$

The number of r-element subsets of a finite set A is $\binom{|A|}{r}$. Taking cardinalities of the two sides and writing the cardinality of the right-hand side as a sum of indicators gives

$$\binom{|N_i \cap N_l|}{r} = \#\{S \in \mathcal{F} : S \subseteq N_l\} = \sum_{S \in \mathcal{F}} \mathbf{1}\{S \subseteq N_l\}.$$

Step 2: the double-count identity. Summing the display of Step 1 over $l \in V_n$ and exchanging the two finite sums,

$$\sum_{l \in V_n} \binom{|N_i \cap N_l|}{r} = \sum_{l \in V_n} \sum_{S \in \mathcal{F}} \mathbf{1}\{S \subseteq N_l\} = \sum_{S \in \mathcal{F}} \sum_{l \in V_n} \mathbf{1}\{S \subseteq N_l\} = \sum_{S \in \mathcal{F}} \#\{l \in V_n : S \subseteq N_l\},$$

the last equality because an indicator sum over $l \in V_n$ counts the l satisfying the condition. This is the equality asserted in the statement.

Step 3: each subset is reused at most d times. Fix $S \in \mathcal{F}$. Since $|S| = r$ and $1 \leq r$, the set S is nonempty; fix any $j \in S$. If $S \subseteq N_l$ then in particular $j \in N_l$, so

$$\{l \in V_n : S \subseteq N_l\} \subseteq \{l \in V_n : j \in N_l\}.$$

The out-degree half of Assumption 2, applied at the coordinate j, states that $\#\{l \in V_n : j \in N_l\} \leq d$. Combining this with monotonicity of cardinality under inclusion,

$$\#\{l \in V_n : S \subseteq N_l\} \leq \#\{l \in V_n : j \in N_l\} \leq d.$$

Step 4: the first inequality. The family $\mathcal{F}$ is precisely the family of r-element subsets of N_i, hence

$$|\mathcal{F}| = \binom{|N_i|}{r}.$$

Summing the bound of Step 3 over $S \in \mathcal{F}$ therefore gives

$$\sum_{S \in \mathcal{F}} \#\{l \in V_n : S \subseteq N_l\} \leq \sum_{S \in \mathcal{F}} d = d\,|\mathcal{F}| = d\binom{|N_i|}{r}.$$

Step 5: the second inequality. The in-degree half of Assumption 2, applied at the unit i, gives $|N_i| \leq d$. Since $m \mapsto \binom{m}{r}$ is nondecreasing in m, we get $\binom{|N_i|}{r} \leq \binom{d}{r}$, and multiplying this inequality by the nonnegative integer d yields

$$d\binom{|N_i|}{r} \leq d\binom{d}{r}.$$

Chaining the identity of Step 2 with the inequalities of Steps 4 and 5 proves the claim. $\square$

Lemma 2 explains where the extra factor d enters the upper bounds. The binomial term $\binom{|N_i|}{r}$ is already represented in the local block energy, while the count of units whose neighborhoods contain the same r-set is controlled by the out-degree part of Assumption 2. This is the covariance-counting step used to pass from block energy to the global dA_d/n scale.

The remaining auxiliary statements support the lower bound. They place the two sign-indexed complete-block priors from Definition 15 on a common dominating measure and bound the unhalved squared Hellinger distance between their densities.

Lemma 3. [`lem:block-prior-dominating-measure`] (Block prior integrates to one) *Use the notation for the Bernoulli assignment law and finite population from Assumption 1, the block size from Assumption 2, the interaction order from Assumption 3, and the radius from Assumption 4. Let $n, d, \beta \in \mathbb{N}$ and $B, p, \sigma \in \mathbb{R}$, and set $m = \lfloor n/d \rfloor$. Suppose that:*

- ***(Population and blocks.)*** $1 \le n$ *and* $1 \le d \le n$.
- ***(Radius.)*** $B > 0$, *and let* $s = B/2$.
- ***(Design probability.)*** $p \in (0, 1)$.
- ***(Prior sign.)*** $\sigma \in \{-1, +1\}$.

Let f_s be the cosine-squared baseline density from Definition 15, let δ be the perturbation amplitude from Definition 15, and let h_d be the normalized complete-block representer from Definition 10. Define the block dominating measure by

$$\mu_{n,d} = \mathrm{count}_{\{0,1\}^{V_n}} \otimes \lambda^m.$$

For $(z, u) \in \{0,1\}^{V_n} \times \mathbb{R}^m$, define the sign-$\sigma$ block prior density by

$$\pi_\sigma(z, u) = \Big(\prod_{j \in V_n} p^{z_j} (1-p)^{1-z_j} \Big) \prod_{b=1}^{m} f_s\big(u_b - \sigma\delta\, h_d(z_{V_b})\big),$$

where $V_1, \dots, V_m$ are the complete d-blocks of Definition 15. Then

$$\int_{\{0,1\}^{V_n} \times \mathbb{R}^m} \pi_\sigma(z, u)\, d\mu_{n,d}(z, u) = 1.$$

Thus π_σ is a probability density with respect to the common block dominating measure $\mu_{n,d}$.

Proof of Lemma 3. Put

$$a(z) := \prod_{j \in V_n} p^{z_j}(1-p)^{1-z_j}, \qquad c_b(z) := \sigma\delta\, h_d(z_{V_b}), \qquad s := B/2.$$

Then $s > 0$. The cosine-squared baseline satisfies

$$\int_{\mathbb{R}} f_s(u)\, du = s^{-1} \int_{-s}^{s} \cos^2\Big(\frac{\pi u}{2s}\Big)\, du = 1,$$

by the change of variables $x = \pi u/(2s)$ and $\int_{-\pi/2}^{\pi/2} \cos^2 x\, dx = \pi/2$. Translation invariance of Lebesgue measure gives, for every real c,

$$\int_{\mathbb{R}} f_s(u - c)\, du = 1.$$

For a fixed assignment z, the conditional density in the block coordinates is

$$u \longmapsto \prod_{b=1}^{m} f_s(u_b - c_b(z)).$$

Each coordinate factor is integrable, and finite-product Fubini gives

$$\int_{\mathbb{R}^m} \prod_{b=1}^{m} f_s(u_b - c_b(z))\, du = \prod_{b=1}^{m} \int_{\mathbb{R}} f_s(u - c_b(z))\, du = 1.$$

Consequently,

$$\int_{\mathbb{R}^m} \pi_\sigma(z, u)\, du = a(z).$$

The measurability and integrability needed to apply product Fubini follow from measurability of f_s, finiteness of the assignment space, and integrability of the translated coordinate densities. Hence

$$\int_{\{0,1\}^{V_n} \times \mathbb{R}^m} \pi_\sigma(z, u)\, d\mu_{n,d}(z, u) = \sum_{z \in \{0,1\}^{V_n}} \int_{\mathbb{R}^m} \pi_\sigma(z, u)\, du = \sum_{z \in \{0,1\}^{V_n}} a(z).$$

Finally, the Bernoulli product weights sum to one:

$$\sum_{z \in \{0,1\}^{V_n}} \prod_{j \in V_n} p^{z_j} (1-p)^{1-z_j} = \prod_{j \in V_n} \{p + (1-p)\} = 1.$$

This proves that the displayed integral is equal to 1. □

Lemma 3 makes the prior comparison a calculation between two ordinary densities. The product Bernoulli assignment component is the same under both signs, while the observed block coefficients are shifted by the score-aligned perturbation from Definition 15.

Lemma 4. [`lem:block-prior-hellinger-bound`] (Block prior Hellinger bound) *Let n, d, β be integers and let $B, p \in \mathbb{R}$. Suppose that:*

- ***(Design probability.)*** *The Bernoulli design notation and treatment probability p are as in Assumption 1, with $0 < p < 1$.*
- ***(Block size.)*** *The degree parameter d is as in Assumption 2, with $1 \le d \le n$ and $1 \le n$.*
- ***(Interaction order.)*** *The interaction-order parameter β is as in Assumption 3, with $\beta \ge 1$.*
- ***(Radius.)*** *The radius B is as in Assumption 4, with $B > 0$.*
- ***(Block priors.)*** *Let $\mu_{n,d}$, π_{+1}, and π_{-1} be the block dominating measure and the two signed block prior densities from Lemma 3.*
- ***(Amplitude and energy.)*** *Let*

  $$m := \left\lfloor \frac{n}{d} \right\rfloor, \qquad \delta := B \min\{(2H_{\beta,p})^{-1}, (4\pi)^{-1}\} \min\left\{1, \sqrt{\frac{A_d}{m}}\right\},$$

  where m and δ are the complete-block count and perturbation amplitude from Definition 15, and A_d is the complete-block score energy from Definition 10.

Then the unhalved squared Hellinger distance

$$H^2(\pi_{+1}, \pi_{-1}) := \int \big(\sqrt{\pi_{+1}} - \sqrt{\pi_{-1}}\big)^2 \, d\mu_{n,d}$$

satisfies

$$H^2(\pi_{+1}, \pi_{-1}) \leq \frac{4\pi^2 m \delta^2}{B^2 A_d}.$$

Proof of Lemma 4. Let

$$s := B/2, \qquad \delta := B \min\{(2H_{\beta,p})^{-1}, (4\pi)^{-1}\} \min\left\{1, \sqrt{\frac{A_d}{m}}\right\},$$

and, for an assignment z, write

$$r_b(z) := h_d(z_{V_b}).$$

For the two conditional block-coordinate densities define

$$C_z^+(u) := \prod_{b=1}^{m} f_s\big(u_b - \delta r_b(z)\big), \qquad C_z^-(u) := \prod_{b=1}^{m} f_s\big(u_b + \delta r_b(z)\big),$$

so that

$$\pi_{+1}(z,u) = P_Z(z) C_z^+(u), \qquad \pi_{-1}(z,u) = P_Z(z) C_z^-(u).$$

Since $P_Z(z) \geq 0$,

$$\big(\sqrt{\pi_{+1}(z,u)} - \sqrt{\pi_{-1}(z,u)}\big)^2 = P_Z(z)\big(\sqrt{C_z^+(u)} - \sqrt{C_z^-(u)}\big)^2.$$

Applying product Fubini over the counting measure in z and Lebesgue measure in u gives

$$H^2(\pi_{+1}, \pi_{-1}) = \sum_{z \in \{0,1\}^{V_n}} P_Z(z) \, H^2(C_z^+, C_z^-).$$

We next bound the conditional Hellinger distance. For fixed z, set

$$f_b(u) := f_s(u - \delta r_b(z)), \qquad g_b(u) := f_s(u + \delta r_b(z)).$$

The coordinate densities are nonnegative, integrable, and normalized:

$$f_b(u) \geq 0, \qquad g_b(u) \geq 0, \qquad \int f_b(u) \, du = \int g_b(u) \, du = 1.$$

Let

$$\alpha_b := \int_{\mathbb{R}} \sqrt{f_b(u) g_b(u)} \, du.$$

The identity

$$H^2(C_z^+, C_z^-) = 2\left(1 - \prod_{b=1}^{m} \alpha_b\right)$$

follows by expanding $(\sqrt{\prod_b f_b} - \sqrt{\prod_b g_b})^2$, using $\sqrt{\prod_b f_b g_b} = \prod_b \sqrt{f_b g_b}$, and applying finite-product Fubini. Moreover $0 \le \alpha_b \le 1$, where the upper bound is Cauchy–Schwarz together with the two normalization identities. Therefore

$$H^2(C_z^+, C_z^-) \le 2\sum_{b=1}^{m}(1-\alpha_b).$$

For any $a, b \in \mathbb{R}$, the cosine translate affinity obeys

$$1 - \int_{\mathbb{R}} \sqrt{f_s(u-a) f_s(u-b)}\, du \le \frac{\pi^2(a-b)^2}{8s^2}.$$

Indeed, if $s \le |a-b|$, the left side is at most 1, while $\pi^2(a-b)^2/(8s^2) \ge 1$. If $|a-b| < s$, order the shifts as $x \le y$, put

$$t := y - x, \qquad q := \frac{\pi t}{2s},$$

and compute the overlap integral as

$$\int \sqrt{f_s(u-x) f_s(u-y)}\, du = \left(1 - \frac{t}{2s}\right)\cos q + \frac{\sin q}{\pi}.$$

Since $0 \le q < \pi/2$ and $q \le \tan q$, the right-hand side is at least $\cos q$. Thus

$$1 - \int \sqrt{f_s(u-x) f_s(u-y)}\, du \le 1 - \cos q \le \frac{q^2}{2} = \frac{\pi^2(x-y)^2}{8s^2},$$

and symmetry gives the displayed bound.

Applying this with $a = \delta r_b(z)$ and $b = -\delta r_b(z)$ yields

$$1 - \alpha_b \le \frac{\pi^2(2\delta r_b(z))^2}{8s^2}.$$

Therefore, since $s = B/2$,

$$H^2(C_z^+, C_z^-) \le 2\sum_{b=1}^{m} \frac{\pi^2(2\delta r_b(z))^2}{8s^2} = \frac{4\pi^2\delta^2}{B^2}\sum_{b=1}^{m} r_b(z)^2.$$

Substituting the conditional bound into the outer sum gives

$$H^2(\pi_{+1}, \pi_{-1}) \le \frac{4\pi^2\delta^2}{B^2} \sum_{z\in\{0,1\}^{V_n}} P_Z(z) \sum_{b=1}^{m} h_d(z_{V_b})^2.$$

It remains to evaluate the Bernoulli average. The restriction Z_{V_b} has the d-coordinate product Bernoulli law with probability p, because P_Z is the product law from Assumption 1. By the energy identity in Lemma 1,

$$\mathbb{E}\big[h_d(Z_{V_b})^2\big] = A_d^{-1}.$$

Summing over the m blocks gives

$$\sum_{z\in\{0,1\}^{V_n}} P_Z(z) \sum_{b=1}^{m} h_d(z_{V_b})^2 = \frac{m}{A_d}.$$

Combining the last two displays,

$$H^2(\pi_{+1}, \pi_{-1}) \leq \frac{4\pi^2\delta^2}{B^2} \frac{m}{A_d} = \frac{4\pi^2 m\delta^2}{B^2 A_d}.$$

This is the claimed bound. □

Lemma 4 supplies the affinity control for the two-point prior comparison. Combined with the total-effect separation stated in Theorems 1 and 2, it yields the saturated lower-bound branch through the complete-block construction and the linear branch through the choice of perturbation amplitude.

Together, these auxiliary statements organize the theorem proofs. The upper bounds in Theorems 1 and 2 use the product-Bernoulli orthogonality in Lemma 1 to identify SNIPE unbiasedness and then apply Lemma 2 to bound the covariance sum by dA_d/n. The lower bounds use the complete-block family in Definition 15, the coefficient-mass control in Lemma 1, and the density and Hellinger calculations in Lemmas 3 and 4. The complete-block local linear theorem is the finite-dimensional dual counterpart: Definition 20 and Lemma 1 identify the least possible block energy for unbiased weights, and Theorem 3 aggregates that block criterion across complete blocks.

B Verification note

This appendix records the Lean verification scope for the mathematical results stated in the paper.

The Lean 4 development verifies the finite-design algebra, displayed formal statements, and displayed proofs used in the paper. The checked portion includes the Bernoulli design algebra in Assumption 1, the bounded-degree and low-order polynomial constructions in Assumptions 2 and 3 and Definitions 5 and 6, the SNIPE score and estimator identities in Definitions 7 and 8, the minimax-risk definitions in Definitions 9 and 13, the complete-block score energy in Definition 10, and the least-favourable block family in Definition 15. It also verifies the displayed theorem and auxiliary statements in Lemmas 1 to 4 and Theorems 1 to 4.

The checked statements include the SNIPE unbiasedness claims, the displayed worst-case mean squared error and risk bounds, the exact complete-block risk identities, the block-prior density and Hellinger calculations in Lemmas 3 and 4, and the perturbation and weight programs in Definitions 19 and 20 through the representer result in Lemma 1. The verification contract records no external formal dependencies for the displayed objects. The cited literature supplies the econometric framing, terminology, and scholarly positioning around these formal statements, including the randomized-experiment, potential-outcomes, design-based, SNIPE, and low-order interference contexts [Horvitz and Thompson, 1952, Rubin, 1978, Imbens and Rubin, 2015, Aronow and Samii, 2017, Cortez-Rodriguez et al., 2023, Harshaw et al., 2022].

The Lean development accompanying this paper is the module tree `CausalSmith/Experimentation/EXP_Snip` of the CausalSmith repository, at commit `3e2bcb6`. It builds with `lake build` against the pinned toolchain recorded there, and every displayed statement carries the declaration name it was checked against, so each can be located and rechecked individually.

C Proofs of the main results

Proof of Theorem 1. 1. Apply Theorem 2 with the fixed β and p. It gives constants $\alpha_{\beta,p}, \Gamma_{\beta,p} \in \mathbb{R}$, depending only on (β, p), such that

$$0 < \alpha_{\beta,p} \leq \Gamma_{\beta,p}.$$

Set

$$c_{\beta,p} := \min\{1, \alpha_{\beta,p}^2\}, \qquad C_{\beta,p} := \max\{1, \Gamma_{\beta,p}^2\}.$$

Then $0 < c_{\beta,p}$ and $c_{\beta,p} \leq C_{\beta,p}$. Now fix $n, d \in \mathbb{N}$ and $B \in \mathbb{R}$ with $1 \leq n$, $d \leq n$, and $0 \leq B$. Let $D_{n,p}$ be the product Bernoulli law on $\{0,1\}^{V_n}$,

$$D_{n,p}(z) = \prod_{j \in V_n} p^{z_j}(1-p)^{1-z_j}.$$

This is the product Bernoulli design with common probability p, so Theorem 2 applies to $D_{n,p}$. Write

$$x_{n,d} := \frac{d\binom{d}{k_\star(d,\beta,p)}}{n}, \qquad y_{n,d} := \frac{dA_d}{n}.$$

The same conclusion supplies

$$\alpha_{\beta,p}\, d\binom{d}{k_\star(d,\beta,p)} \leq dA_d \leq \Gamma_{\beta,p}\, d\binom{d}{k_\star(d,\beta,p)}.$$

Since $n \geq 1$, division by n gives

$$\alpha_{\beta,p} x_{n,d} \leq y_{n,d} \leq \Gamma_{\beta,p} x_{n,d}.$$

The case $d = 0$ is included in these displayed inequalities; then $x_{n,d} = y_{n,d} = 0$.

2. The elementary rescalings used to pass from $y_{n,d}$ to $x_{n,d}$ are

$$c_{\beta,p} \min\{1, x_{n,d}\} \leq \alpha_{\beta,p} \min\{1, \alpha_{\beta,p} x_{n,d}\},$$

and

$$\Gamma_{\beta,p} \min\{1, \Gamma_{\beta,p} x_{n,d}\} \leq C_{\beta,p} \min\{1, x_{n,d}\}.$$

For the first inequality, split into $x_{n,d} \leq 1$ and $x_{n,d} > 1$, and inside each case split again according as $\alpha_{\beta,p} x_{n,d} \leq 1$ or $\alpha_{\beta,p} x_{n,d} > 1$; use $c_{\beta,p} \leq \alpha_{\beta,p}^2$, $c_{\beta,p} \leq \alpha_{\beta,p}$, and $x_{n,d} \geq 0$. The second inequality follows from the same two splits, using

$$\Gamma_{\beta,p} \leq C_{\beta,p}, \qquad \Gamma_{\beta,p}^2 \leq C_{\beta,p}, \qquad x_{n,d} \geq 0.$$

Together with $\alpha_{\beta,p} x_{n,d} \leq y_{n,d} \leq \Gamma_{\beta,p} x_{n,d}$, these imply

$$c_{\beta,p} \min\{1, x_{n,d}\} \leq \alpha_{\beta,p} \min\{1, y_{n,d}\},$$

and

$$\Gamma_{\beta,p} \min\{1, y_{n,d}\} \leq C_{\beta,p} \min\{1, x_{n,d}\}.$$

3. Let $R^\star_{n,\ell_1}(d,\beta,p,B)$ denote the coefficient-mass minimax risk in Definition 13. Theorem 2 gives

$$\alpha_{\beta,p}B^2\min\{1,y_{n,d}\}\le R^\star_{n,\ell_1}(d,\beta,p,B),\qquad R^\star_{n,\ell_1}(d,\beta,p,B)=R^\star_n(d,\beta,p,B).$$

   Combining this with the first rescaling yields

$$c_{\beta,p}B^2\min\{1,x_{n,d}\}\le R^\star_n(d,\beta,p,B),$$

   which is the lower bound.

4. The same application of Theorem 2 gives the projected-SNIPE upper bound at the energy scale,

$$\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)}\mathbb{E}_Z\left[(\widehat{\tau}^{\mathrm{up}}_n-\tau_n)^2\right]\le\Gamma_{\beta,p}B^2\min\{1,y_{n,d}\}.$$

   Using the second rescaling gives

$$\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)}\mathbb{E}_Z\left[(\widehat{\tau}^{\mathrm{up}}_n-\tau_n)^2\right]\le C_{\beta,p}B^2\min\{1,x_{n,d}\}.$$

5. It remains to place the minimax risk below the projected-SNIPE risk. The model class is nonempty: the graph with no edges and the identically zero coefficient schedule belongs to $\mathcal{M}_{n,d,\beta}(B)$, since its degree and coefficient mass are both zero and $0\le B$. Also, for every $(G,c)\in\mathcal{M}_{n,d,\beta}(B)$, every unit i, and every assignment z,

$$\begin{aligned}|Y_i(z)|&=\left|\sum_{S\subseteq N_i}c_{i,S}\prod_{j\in S}z_j\right|\\&\le\sum_{S\subseteq N_i}|c_{i,S}|\left|\prod_{j\in S}z_j\right|\le\sum_{S\subseteq N_i}|c_{i,S}|\le B,\end{aligned}$$

   because each $z_j\in\{0,1\}$. Hence, since $n\ge1$,

$$|\tau_n|=\left|\frac{1}{n}\sum_{i\in V_n}\{Y_i(\mathbf{1})-Y_i(\mathbf{0})\}\right|\le 2B.$$

   The projection defining $\widehat{\tau}^{\mathrm{up}}_n$ gives

$$|\widehat{\tau}^{\mathrm{up}}_n|\le B,$$

   and therefore

$$(\widehat{\tau}^{\mathrm{up}}_n-\tau_n)^2\le 9B^2$$

   pointwise. Finite-space measurability is immediate, so $\widehat{\tau}^{\mathrm{up}}_n$ is an admissible estimator for the coefficient-mass minimax problem.

6. Since $R^\star_n(d,\beta,p,B)$ is the infimum of worst risks over admissible estimators and the preceding step shows that $\widehat{\tau}^{\mathrm{up}}_n$ is admissible,

$$R^\star_n(d,\beta,p,B)\le\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)}\mathbb{E}_Z\left[(\widehat{\tau}^{\mathrm{up}}_n-\tau_n)^2\right].$$

7. For the unprojected SNIPE estimator, Theorem 2 gives

$$\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}_n^{\mathrm{SNIPE}}-\tau_n\right)^2\right] \le B^2 y_{n,d}.$$

Using $y_{n,d} \le \Gamma_{\beta,p} x_{n,d}$ and $\Gamma_{\beta,p} \le C_{\beta,p}$,

$$\sup_{(G,c)\in\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}_n^{\mathrm{SNIPE}}-\tau_n\right)^2\right] \le C_{\beta,p} B^2 x_{n,d},$$

which is the stated unprojected bound.

8. For the unitwise energy statement, define, for $(G,c)\in\mathcal{M}_{n,d,\beta}(B)$ and $i\in V_n$,

$$A_i := \sum_{r=1}^{\min\{\beta,|N_i|\}} \binom{|N_i|}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r}.$$

The local-energy conclusion of Theorem 2 gives

$$A_i \le A_d,$$

which is exactly

$$\sum_{r=1}^{\min\{\beta,|N_i|\}} \binom{|N_i|}{r} \frac{\Delta_r(p)^2}{\{p(1-p)\}^r} \le A_d.$$

9. Fix $(G,c)\in\mathcal{M}_{n,d,\beta}(B)$, $i\in V_n$, and $r\in\mathbb{N}$ with $1\le r$. Put

$$\mathcal{S}_{i,r} := \{S\subseteq N_i :\ |S| = r\}.$$

Since membership in $\mathcal{M}_{n,d,\beta}(B)$ includes the bounded-degree condition with bound d, Lemma 2 gives

$$\sum_{\ell\in V_n} \binom{|N_i\cap N_\ell|}{r} = \sum_{S\in\mathcal{S}_{i,r}} \#\{\ell\in V_n : S\subseteq N_\ell\},$$

and

$$\sum_{S\in\mathcal{S}_{i,r}} \#\{\ell\in V_n : S\subseteq N_\ell\} \le d\binom{|N_i|}{r} \le d\binom{d}{r}.$$

10. Finally assume $0<B$ and $1\le d$, and set

$$m=\left\lfloor \frac{n}{d}\right\rfloor, \qquad \rho=\frac{md}{n}, \qquad \delta=\delta(n,d,\beta,B,p).$$

For every $U:\{1,\ldots,m\}\to\mathbb{R}$ satisfying $|U_b|\le B/2$, the block-mixture conclusion of Theorem 2 supplies models M_+ and M_- in $\mathcal{M}_{n,d,\beta}(B)$ with complete d-block graph and the corresponding sign-$+1$ and sign--1 block schedules. Their total treatment effects satisfy

$$\tau_n(M_+)=\rho\delta, \qquad \tau_n(M_-)=-\rho\delta.$$

The same conclusion gives the Hellinger bound

$$H^2(\Pi_+,\Pi_-) \le \frac{4\pi^2 m\delta^2}{B^2 A_d},$$

and the active-fraction identities

$$\frac{1}{2} \le \rho \le 1, \qquad \frac{A_d}{m} = \frac{dA_d/n}{\rho}.$$

These are precisely the remaining assertions.

□

Proof of Theorem 3. Throughout, when t is fixed we write $n = n_t$, $d = d_t$, $m = m_t$, $A = A_{d_t}$ and $\bar{\beta} = \bar{\beta}_{d_t} = \min\{\beta, d\}$. The population V_n is partitioned into the blocks $V_1, \dots, V_m$, each of size d; $b(i)$ is the block containing i, $Z_b = (Z_j)_{j \in V_b}$, and $Z_S = \prod_{j \in S} Z_j$. Since G_t is the directed complete-block graph,

$$N_i^{G_t} = V_{b(i)} \quad \text{for every } i \in V_n, \qquad n = md.$$

The clauses are proved in the order in which their ingredients are established: the exact worst-case formula, then the per-block lower bound, then the minimax constant, then the two asymptotic criteria, then the witness sequence.

We first record the three block-score identities that every step uses. The block score g_d of Definition 10 lies in $\mathcal{P}_d$, and Lemma 1 gives, under Assumption 1,

$$\mathbb{E}_Z[g_d(Z_b)] = 0, \qquad \mathbb{E}_Z[g_d(Z_b) Z_S] = 1 \quad \text{for every } \varnothing \neq S \subseteq V_b \text{ with } |S| \le \bar{\beta},$$
$$\mathbb{E}_Z[g_d(Z_b)^2] = L_d(g_d) = A > 0.$$

The first two displays are the representer identity $\mathbb{E}_Z[g_d(Z) f(Z)] = L_d(f)$ at $f \equiv 1$ and at $f = Z_S$; the third combines that identity at $f = g_d$ with $\mathbb{E}_Z[h_d(Z)^2] = A^{-1}$ and $g_d = A h_d$, and the strict positivity is the energy-scale bound $A \ge c_1 \binom{d}{k_\star(d,\beta,p)} > 0$.

1. **The exact worst-case risk formula.** Fix t and $w \in \mathcal{E}_t^{\mathrm{loc,lin}}$. For $c \in \mathcal{M}_t(G_t)$ the potential outcomes are $Y_i(Z) = \sum_{S \subseteq V_{b(i)},\, |S| \le \bar{\beta}} c_{i,S} Z_S$, and the target is the all-treated versus all-control contrast

$$\tau_n = \frac{1}{n} \sum_{i \in V_n} \{Y_i(\mathbf{1}_n) - Y_i(\mathbf{0}_n)\} = \frac{1}{n} \sum_{i \in V_n} \sum_{\substack{S \subseteq V_{b(i)} \\ |S| \le \bar{\beta}}} c_{i,S} \mathbf{1}(S \neq \varnothing).$$

Subtracting this from $\widehat{\tau}_{w,t} = n^{-1} \sum_i w_{i,t}(Z_{b(i)}) Y_i(Z)$ and grouping units by block,

$$\widehat{\tau}_{w,t} - \tau_n = \frac{1}{n} \sum_{b=1}^{m} e_b(c; Z), \qquad e_b(c; Z) := \sum_{i \in V_b} \sum_{\substack{S \subseteq V_b \\ |S| \le \bar{\beta}}} c_{i,S} \big\{ w_{i,t}(Z_b) Z_S - \mathbf{1}(S \neq \varnothing) \big\}.$$

Each $e_b(c; Z)$ is a function of Z_b alone, and the moment restrictions defining $\mathcal{E}_t^{\mathrm{loc,lin}}$ give $\mathbb{E}_Z[e_b(c; Z)] = 0$ termwise. Distinct blocks use disjoint assignment coordinates, so under Assumption 1 the variables $e_1(c; Z), \dots, e_m(c; Z)$ are independent and the cross terms vanish:

$$\mathbb{E}_Z\big[(\widehat{\tau}_{w,t} - \tau_n)^2\big] = \frac{1}{n^2} \sum_{b=1}^{m} \mathbb{E}_Z\big[e_b(c; Z)^2\big].$$

For a fixed block, $c \mapsto e_b(c; Z)$ is linear in the rows $(c_{i,S})_S$ and the only constraint $\mathcal{M}_t(G_t)$ places on a row is $\sum_S |c_{i,S}| \le B$ with $S \subseteq V_b$ and $|S| \le \bar{\beta}$. The criterion $c \mapsto \mathbb{E}_Z[e_b(c; Z)^2]$ is convex in each row, so its maximum over the row simplex of radius B is attained at a vertex $c_{i,S} = \varepsilon_i B \mathbf{1}(S = S_i)$ with $\varepsilon_i \in \{-1, 1\}$, for which $e_b(c; Z) = B \sum_{i \in V_b} \varepsilon_i \{w_{i,t}(Z_b) Z_{S_i} - \mathbf{1}(S_i \neq \varnothing)\}$. Hence

$$\sup_{c \in \mathcal{M}_t(G_t)} \mathbb{E}_Z\big[e_b(c; Z)^2\big] = B^2 \Psi_{b,t}(w).$$

The row constraint is unitwise, so the maximizing vertices may be chosen block by block and assembled into one schedule in $\mathcal{M}_t(G_t)$ that attains every block maximum simultaneously. Therefore

$$\sup_{c \in \mathcal{M}_t(G_t)} \mathbb{E}_Z\big[(\widehat{\tau}_{w,t} - \tau_n)^2\big] = \frac{B^2}{n^2} \sum_{b=1}^{m} \Psi_{b,t}(w),$$

which is the asserted worst-case identity.

2. **The per-block lower bound.** Fix t, $w \in \mathcal{E}_t^{\mathrm{loc,lin}}$ and a block V_b. Every $i \in V_b$ has $N_i^{G_t} = V_b$, so the single function $g_d(Z_b)$ is the common block score for all of them. Expand it in the raw monomials it is spanned by,

$$g_d(Z_b) = \sum_{\substack{T \subseteq V_b \\ |T| \le \bar{\beta}}} a_T Z_T, \qquad \sum_{\substack{T \subseteq V_b \\ |T| \le \bar{\beta}}} a_T = g_d(\mathbf{1}_d), \qquad a_\varnothing = g_d(\mathbf{0}_d).$$

Applying $\mathbb{E}_Z[w_{i,t}] = 0$ to the term $T = \varnothing$ and $\mathbb{E}_Z[w_{i,t} Z_T] = 1$ to every other term gives, for every $i \in V_b$,

$$\mathbb{E}_Z[w_{i,t}(Z_b) g_d(Z_b)] = \sum_{\substack{\varnothing \neq T \subseteq V_b \\ |T| \le \bar{\beta}}} a_T = g_d(\mathbf{1}_d) - g_d(\mathbf{0}_d) = L_d(g_d) = A.$$

Summing over the d units of V_b and using $\mathbb{E}_Z[g_d(Z_b)^2] = A$, the expansion of a nonnegative square gives

$$0 \le \mathbb{E}_Z\left[\left\{\sum_{i \in V_b} w_{i,t}(Z_b) - d g_d(Z_b)\right\}^2\right] = \mathbb{E}_Z\left[\left\{\sum_{i \in V_b} w_{i,t}(Z_b)\right\}^2\right] - 2d(dA) + d^2 A,$$

that is,

$$d^2 A \le \mathbb{E}_Z\left[\left\{\sum_{i \in V_b} w_{i,t}(Z_b)\right\}^2\right].$$

The choice $\varepsilon_i = 1$ and $S_i = \varnothing$ for every $i \in V_b$ is admissible in $\Psi_{b,t}$, since $\varnothing \subseteq V_b$ and $|\varnothing| = 0 \le \bar{\beta}$, and it produces exactly that expectation; as $\Psi_{b,t}$ is a maximum over admissible choices,

$$\mathbb{E}_Z\left[\left\{\sum_{i \in V_b} w_{i,t}(Z_b)\right\}^2\right] \le \Psi_{b,t}(w), \qquad \text{hence} \qquad d_t^2 A_{d_t} \le \Psi_{b,t}(w).$$

3. **The exact minimax constant.** Define the canonical block-local weights
$$w^0_{i,t}(Z) := g_d(Z_{b(i)}),$$
which are the SNIPE scores of Definition 7 on G_t because $N_i^{G_t} = V_{b(i)}$ has exactly d elements. They depend on Z only through $Z_{b(i)}$, and the three block-score identities recorded above are precisely the centering and unit raw-moment restrictions, so $w^0 \in \mathcal{E}_t^{\mathrm{loc,lin}}$.

   Fix a block V_b and an admissible choice $(\varepsilon_i, S_i)_{i\in V_b}$ in $\Psi_{b,t}$. All units of V_b carry the same weight $g_d(Z_b)$, so
$$\sum_{i\in V_b} \varepsilon_i\big\{w^0_{i,t}(Z_b)Z_{S_i} - \mathbf{1}(S_i \neq \varnothing)\big\} = g_d(Z_b)F(Z_b) - \Lambda,$$
$$F(Z_b) := \sum_{i\in V_b} \varepsilon_i Z_{S_i}, \qquad \Lambda := \sum_{i\in V_b} \varepsilon_i \mathbf{1}(S_i \neq \varnothing).$$
   The block-score identities give $\mathbb{E}_Z[g_d(Z_b)F(Z_b)] = \Lambda$, whence
$$\mathbb{E}_Z\Big[\big(g_d(Z_b)F(Z_b) - \Lambda\big)^2\Big] = \mathbb{E}_Z\Big[\big(g_d(Z_b)F(Z_b)\big)^2\Big] - \Lambda^2 \le \mathbb{E}_Z\Big[\big(g_d(Z_b)F(Z_b)\big)^2\Big].$$
   Since $0 \le Z_{S_i} \le 1$, $|\varepsilon_i| = 1$ and $|V_b| = d$, we have $|F(Z_b)| \le d$ pointwise, so
$$\mathbb{E}_Z\Big[\big(g_d(Z_b)F(Z_b)\big)^2\Big] \le d^2\,\mathbb{E}_Z[g_d(Z_b)^2] = d^2 A.$$
   Maximizing over admissible choices gives the upper bound, and the empty-monomial construction gives the reverse inequality for the same canonical weights, so
$$\Psi_{b,t}(w^0) \le d^2 A, \qquad \text{and} \qquad \Psi_{b,t}(w^0) = d^2 A.$$
   By the worst-case identity and the per-block lower bound, every $w \in \mathcal{E}_t^{\mathrm{loc,lin}}$ satisfies
$$\sup_{c\in\mathcal{M}_t(G_t)} \mathbb{E}_Z\big[(\widehat{\tau}_{w,t} - \tau_n)^2\big] = \frac{B^2}{n^2}\sum_{b=1}^{m} \Psi_{b,t}(w) \ge \frac{B^2}{n^2} m d^2 A = \frac{B^2 A}{m},$$
   the last equality because $n = md$; and the same computation with $\Psi_{b,t}(w^0) = d^2 A$ turns the inequality into an equality for w^0. Taking the infimum over $\mathcal{E}_t^{\mathrm{loc,lin}}$,
$$R_t^{\mathrm{loc,lin}} = \frac{B^2 A_{d_t}}{m_t} = \frac{B^2 d_t A_{d_t}}{n_t},$$
   where the second equality is again $n_t = m_t d_t$.

4. **The risk ratio equals one plus the normalized block excess.** For a sequence $w_t \in \mathcal{E}_t^{\mathrm{loc,lin}}$ set
$$Y_t := \frac{1}{m_t d_t^2 A_{d_t}} \sum_{b=1}^{m_t} \big\{\Psi_{b,t}(w_t) - d_t^2 A_{d_t}\big\}.$$
   Dividing the worst-case identity by $B^2 A_{d_t}/m_t$ and using $n_t^2 = m_t^2 d_t^2$,
$$\frac{\sup_{c\in\mathcal{M}_t(G_t)} \mathbb{E}_Z[(\widehat{\tau}_{w_t,t} - \tau_{n_t})^2]}{B^2 A_{d_t}/m_t} = \frac{m_t}{n_t^2 A_{d_t}} \sum_{b=1}^{m_t} \Psi_{b,t}(w_t) = \frac{1}{m_t d_t^2 A_{d_t}} \sum_{b=1}^{m_t} \Psi_{b,t}(w_t) = 1 + Y_t.$$
   Adding, respectively subtracting, the constant sequence 1 therefore shows that the left-hand ratio converges to 1 if and only if $Y_t \to 0$, which is the asserted equivalence.

5. **Score closeness implies attainment.** Put $\delta_{i,t} := w_{i,t} - g_{d_t}$ on the block of i, and set

$$E_{b,t} := \sum_{i \in V_b} \mathbb{E}_Z\left[\delta_{i,t}(Z_b)^2\right], \qquad E_t := \sum_{b=1}^{m_t} E_{b,t} = \sum_{i \in V_{n_t}} \mathbb{E}_Z\left[\left\{w_{i,t}(Z_{b(i)}) - g_{d_t}(Z_{b(i)})\right\}^2\right],$$

$$X_t := \frac{E_t}{n_t A_{d_t}} = \frac{1}{n_t A_{d_t}} \sum_{i \in V_{n_t}} \mathbb{E}_Z\left[\left\{w_{i,t}(Z_{b(i)}) - g_{d_t}(Z_{b(i)})\right\}^2\right],$$

so that the hypothesis of this clause is $X_t \to 0$. The per-block lower bound gives $Y_t \geq 0$ for every t.

Fix t, a block V_b and an admissible choice $(\varepsilon_i, S_i)_{i \in V_b}$. Splitting each weight into its canonical part and its deviation,

$$\sum_{i \in V_b} \varepsilon_i \left\{w_{i,t}(Z_b) Z_{S_i} - \mathbf{1}(S_i \neq \varnothing)\right\} = f(Z_b) + v(Z_b),$$

$$f(Z_b) := \sum_{i \in V_b} \varepsilon_i \left\{g_{d_t}(Z_b) Z_{S_i} - \mathbf{1}(S_i \neq \varnothing)\right\}, \qquad v(Z_b) := \sum_{i \in V_b} \varepsilon_i \delta_{i,t}(Z_b) Z_{S_i}.$$

For the canonical part, the calculation with the weights w^0 gives $\mathbb{E}_Z[f(Z_b)^2] \leq d_t^2 A_{d_t}$. For v, the Cauchy–Schwarz inequality over the d_t summands together with $0 \leq Z_{S_i} \leq 1$ and $|\varepsilon_i| = 1$ gives $v(Z_b)^2 \leq d_t \sum_{i \in V_b} \delta_{i,t}(Z_b)^2$ pointwise. Thus

$$\mathbb{E}_Z[f(Z_b)^2] \leq d_t^2 A_{d_t}, \qquad \mathbb{E}_Z[v(Z_b)^2] \leq d_t E_{b,t}.$$

Bounding the cross term by Cauchy–Schwarz, $\mathbb{E}_Z[(f+v)^2] \leq \mathbb{E}_Z[f^2] + 2\sqrt{\mathbb{E}_Z[f^2]}\sqrt{\mathbb{E}_Z[v^2]} + \mathbb{E}_Z[v^2]$, and maximizing over admissible choices,

$$\Psi_{b,t}(w_t) - d_t^2 A_{d_t} \leq 2\sqrt{d_t^2 A_{d_t}}\sqrt{d_t E_{b,t}} + d_t E_{b,t}.$$

Summing over b and applying Cauchy–Schwarz once more, $\sum_{b=1}^{m_t} \sqrt{d_t E_{b,t}} \leq \sqrt{m_t}\sqrt{d_t E_t}$, so

$$\sum_{b=1}^{m_t} \left\{\Psi_{b,t}(w_t) - d_t^2 A_{d_t}\right\} \leq 2\sqrt{d_t^2 A_{d_t}}\sqrt{m_t}\sqrt{d_t E_t} + d_t E_t.$$

Dividing by $m_t d_t^2 A_{d_t}$ and using $n_t = m_t d_t$,

$$\frac{2\sqrt{d_t^2 A_{d_t}}\sqrt{m_t}\sqrt{d_t E_t} + d_t E_t}{m_t d_t^2 A_{d_t}} = 2\sqrt{\frac{E_t}{n_t A_{d_t}}} + \frac{E_t}{n_t A_{d_t}} = 2\sqrt{X_t} + X_t,$$

so that

$$0 \leq Y_t \leq 2\sqrt{X_t} + X_t.$$

If $X_t \to 0$ then the right-hand bound tends to 0 by continuity of the square root, so $Y_t \to 0$ by squeezing. Since the risk ratio is $1 + Y_t$, this is the asserted convergence of the risk ratio to 1.

6. **An optimal sequence at normalized distance two.** Assume $d_t \to \infty$. For every t and each block V_b, fix a distinguished unit $i_b^\star \in V_b$, and define the perturbation by

$$\psi_b(Z_b) := \begin{cases} \left\{ \dfrac{2d_t A_{d_t}}{[p(1-p)]^{d_t}} \right\}^{1/2} \prod_{j \in V_b} (Z_j - p), & \text{if } \beta < d_t, \\ 0, & \text{if } \beta \geq d_t. \end{cases}$$

Set, for every t,

$$w_{i,t}(Z) := g_{d_t}(Z_{b(i)}) + \mathbf{1}\big(i = i^\star_{b(i)}\big)\, \psi_{b(i)}(Z_{b(i)}).$$

Thus the sequence is defined on the whole index set; at indices with $\beta \geq d_t$ it coincides with the canonical block-score weights.

Admissibility. Each ψ_b is a function of Z_b alone. If $\beta \geq d_t$, then $\psi_b = 0$, and the block-score identities verify the centering and unit raw-moment conditions. If $\beta < d_t$, then $\bar\beta_{d_t} = \beta$, and for every $S \subseteq V_b$ with $|S| \leq \beta$, including $S = \varnothing$, at least one index $j \in V_b$ outside S occurs in the product. Under Assumption 1, that coordinate is independent of Z_S and of the remaining factors and satisfies $\mathbb{E}_Z[Z_j - p] = 0$; hence

$$\mathbb{E}_Z\left[\Big(\prod_{j \in V_b} (Z_j - p) \Big) Z_S \right] = 0 \qquad \text{for every } S \subseteq V_b \text{ with } |S| < d_t.$$

So the perturbation changes neither the centering nor the unit raw moments inherited from g_{d_t}, and $w_t \in \mathcal{E}_t^{\mathrm{loc,lin}}$ for every t.

Since $d_t \to \infty$, the inequality $\beta < d_t$ holds for all sufficiently large t. All remaining computations are on this tail, where the first branch in the definition of ψ_b applies.

Perturbation energy. By independence and $\mathbb{E}_Z[(Z_j - p)^2] = p(1-p)$,

$$\mathbb{E}_Z\left[\Big(\prod_{j \in V_b} (Z_j - p) \Big)^2 \right] = [p(1-p)]^{d_t}, \qquad \text{hence} \qquad \mathbb{E}_Z\big[\psi_b(Z_b)^2\big] = 2d_t A_{d_t}.$$

Block bound and the risk ratio. Fix a block V_b and an admissible choice, and split the inner sum into canonical and perturbation parts. Because the perturbation is carried by the single unit $i_b^\star$, the deviation part reduces to one summand, $v(Z_b) = \varepsilon_{i_b^\star} \psi_b(Z_b) Z_{S_{i_b^\star}}$, and $0 \leq Z_{S_{i_b^\star}} \leq 1$ gives $\mathbb{E}_Z[v(Z_b)^2] \leq \mathbb{E}_Z[\psi_b(Z_b)^2] = 2d_t A_{d_t}$; the canonical part obeys $\mathbb{E}_Z[f(Z_b)^2] \leq d_t^2 A_{d_t}$. The same Cauchy–Schwarz bound on the cross term therefore yields, for all sufficiently large t and every b,

$$\Psi_{b,t}(w_t) \leq d_t^2 A_{d_t} + 2\sqrt{d_t^2 A_{d_t}}\sqrt{2d_t A_{d_t}} + 2d_t A_{d_t}.$$

Dividing the excess by $d_t^2 A_{d_t}$ gives $2\sqrt{2/d_t} + 2/d_t$ for each block, hence the same bound for Y_t, while $Y_t \geq 0$ by the per-block lower bound. Since the risk ratio is $1 + Y_t$,

$$1 \leq \frac{\sup_{c \in \mathcal{M}_t(G_t)} \mathbb{E}_Z[(\widehat{\tau}_{w_t,t} - \tau_{n_t})^2]}{B^2 A_{d_t}/m_t} \leq 1 + 2\sqrt{\frac{2}{d_t}} + \frac{2}{d_t}$$

for all sufficiently large t. Since $d_t \to \infty$, squeezing proves the asserted convergence of the risk ratio to 1.

Normalized distance. For the same sufficiently large t, the deviation $w_{i,t} - g_{d_t}$ equals ψ_b at $i = i_b^\star$ and vanishes at every other unit, so

$$\sum_{i \in V_{n_t}} \mathbb{E}_Z\Big[\big\{w_{i,t}(Z_{b(i)}) - g_{d_t}(Z_{b(i)})\big\}^2\Big] = \sum_{b=1}^{m_t} \mathbb{E}_Z\big[\psi_b(Z_b)^2\big] = 2 m_t d_t A_{d_t},$$

and dividing by $n_t A_{d_t} = m_t d_t A_{d_t}$ gives the value 2.

Invariance under block-preserving relabeling. Let π be a bijection of V_{n_t} with $b(\pi(i)) = b(i)$ for every i, and take t in the same sufficiently large tail. Then π maps each block onto itself and therefore permutes the coordinates of Z_b among themselves. The relabeled distance moves the score term only: at the unit i it compares the weight $w_{i,t}$, read at the unit i and at the original assignment, with the canonical score of the relabeled unit $\pi(i)$, whose neighborhood is again $V_{b(i)}$, read at the relabeled assignment whose jth coordinate is $Z_{\pi^{-1}(j)}$. Since g_{d_t} is a symmetric function of the coordinates of its block, that relabeled score is the original one,

$$g_{d_t}\big((Z_{\pi^{-1}(j)})_{j \in V_{b(i)}}\big) = g_{d_t}(Z_{b(i)}) \qquad \text{for every } i \in V_{n_t}.$$

Each summand of the normalized distance is therefore unchanged, and substituting this equality into the normalized-distance expression gives

$$\frac{1}{n_t A_{d_t}} \sum_{i \in V_{n_t}} \mathbb{E}_Z\bigg[\Big\{w_{i,t}(Z_{b(i)}) - g_{d_t}\big((Z_{\pi^{-1}(j)})_{j \in V_{b(i)}}\big)\Big\}^2\bigg] = 2,$$

which is the asserted invariance.

□

Proof of Theorem 2. 1. **The two constants.** By Lemma 1 there are constants c_1, c_2, K, depending only on (β, p), with $0 < c_1 \le c_2$ and $0 < K$, such that for every $d \ge 1$

$$c_1 \binom{d}{k_\star(d, \beta, p)} \le A_d \le c_2 \binom{d}{k_\star(d, \beta, p)}, \qquad \sum_{T \subseteq \{1, \ldots, d\}} |h_{d,T}| \le K.$$

In particular the quantity $H_{\beta,p} = \sup_{d \ge 1} \sum_T |h_{d,T}|$ of Definition 15 is finite, with $H_{\beta,p} \le K$. It is also at least one: evaluating the raw expansion of Definition 10 at the all-treated and the all-control assignment of a single-coordinate block and subtracting gives

$$1 = L_1(h_1) = \sum_{\varnothing \ne T \subseteq \{1\}} h_{1,T} \le \sum_{T \subseteq \{1\}} |h_{1,T}| \le H_{\beta,p},$$

where $L_1(h_1) = 1$ is the representer identity of Lemma 1. Now set

$$\kappa := \min\big\{(2H_{\beta,p})^{-1}, (4\pi)^{-1}\big\}, \qquad c_{\beta,p} := \min\left\{c_1, \frac{\kappa^2}{16}\right\}, \qquad C_{\beta,p} := \max\{16, c_2\}.$$

Since $H_{\beta,p} \geq 1 > 0$ and $\pi > 0$, both entries in the definition of κ are positive, so $\kappa > 0$ and therefore $c_{\beta,p} > 0$. Also

$$c_{\beta,p} \leq c_1 \leq c_2 \leq C_{\beta,p}.$$

With this κ, the perturbation amplitude of Definition 15 reads

$$\delta = B\kappa\eta, \qquad \eta := \min\Big\{1, \sqrt{A_d/m}\Big\}.$$

2. **Standing data.** Fix $n \geq 1$, a design D satisfying Assumption 1 with probability p, an integer d with $d \leq n$, and $B \geq 0$. Since $p(1-p) > 0$, every summand of A_d in Definition 10 is nonnegative, so $A_d \geq 0$; and by Definition 2 $A_0 = 0$. For $d \geq 1$ the order $r = 1$ contributes $\binom{d}{1}\Delta_1(p)^2/[p(1-p)] = d/[p(1-p)] > 0$ to A_d, so

$$0 < A_d \qquad (d \geq 1).$$

3. **Model comparison.** Let $(G_n, c) \in \mathcal{M}_{n,d,\beta}(B)$. Every raw monomial takes values in $[0,1]$, so for every unit i and every assignment z

$$|Y_i(z)| = \left|\sum_{S \subseteq N_i} c_{i,S} \prod_{j \in S} z_j\right| \leq \sum_{S \subseteq N_i} |c_{i,S}| \leq B,$$

the last step being Assumption 4. The remaining requirements of Definition 11 — neighborhood support and Assumption 3 — are already part of Definition 4, and the graph requirement is unchanged. Hence the same pair (G_n, c) lies in $\mathcal{M}^{\infty}_{n,d,\beta}(B)$: the containment holds carrier by carrier.

Now let $B > 0$ and $d \geq 1$. Pick a unit $i_0 \in V_n$ (possible since $n \geq 1$), let G^0 consist of the single loop $i_0 \to i_0$, and set

$$c^0_{i_0,\varnothing} := B, \qquad c^0_{i_0,\{i_0\}} := -2B,$$

with all other coefficients zero. Then $N_{i_0} = \{i_0\}$ and $N_i = \varnothing$ for $i \neq i_0$, so both degrees in Assumption 2 are at most $1 \leq d$, and the only nonzero coefficients are indexed by sets of cardinality at most $1 \leq \beta$. The induced outcomes are

$$Y_{i_0}(z) = B - 2Bz_{i_0} \in \{B, -B\}, \qquad Y_i(z) = 0 \quad (i \neq i_0),$$

so $(G^0, c^0) \in \mathcal{M}^{\infty}_{n,d,\beta}(B)$. Its coefficient mass at i_0 is

$$\sum_{S \subseteq N_{i_0}} |c^0_{i_0,S}| = B + 2B = 3B > B,$$

so no member of $\mathcal{M}_{n,d,\beta}(B)$ has graph G^0 and schedule c^0, and the containment is strict.

4. **Local energy.** Let (G_n, c) belong to either model class and let $i \in V_n$. By Assumption 2, $|N_i| \leq d$, hence $\min\{\beta, |N_i|\} \leq \bar{\beta}_d$ and $\binom{|N_i|}{r} \leq \binom{d}{r}$ for every r. Since each term $\Delta_r(p)^2/[p(1-p)]^r$ is nonnegative,

$$A_i = \sum_{r=1}^{\min\{\beta,|N_i|\}} \binom{|N_i|}{r} \frac{\Delta_r(p)^2}{[p(1-p)]^r} \leq \sum_{r=1}^{\bar{\beta}_d} \binom{d}{r} \frac{\Delta_r(p)^2}{[p(1-p)]^r} = A_d.$$

The estimate uses only the degree bound of the carrier graph, so it holds verbatim on $\mathcal{M}_{n,d,\beta}(B)$ and on $\mathcal{M}^{\infty}_{n,d,\beta}(B)$.

5. **Unbiasedness of SNIPE.** Under Assumption 1 the design D is the common-p product Bernoulli law, so the centered monomials in Definition 7 obey, for every unit i and every nonempty $S \subseteq N_i$ with $|S| \le \min\{\beta, |N_i|\}$,

$$\mathbb{E}_Z\left[g_{i,\beta,p}(Z)\prod_{j\in S} Z_j\right] = 1, \qquad \mathbb{E}_Z[g_{i,\beta,p}(Z)] = 0.$$

The first identity collapses the double sum in $g_{i,\beta,p}$ to the binomial contrast $\sum_{r=1}^{|S|}\binom{|S|}{r}\Delta_r(p)p^{|S|-r} = 1$; the second is the vanishing design mean of every non-constant centered monomial. Expanding Y_i over its raw monomials, using Assumption 3 to discard the terms with $|S| > \beta$, and applying the two identities term by term gives

$$\mathbb{E}_Z[g_{i,\beta,p}(Z)Y_i(Z)] = \sum_{\varnothing\neq S\subseteq N_i} c_{i,S} = Y_i(\mathbf{1}) - Y_i(\mathbf{0}).$$

Averaging over $i \in V_n$ yields $\mathbb{E}_Z[\widehat{\tau}_n^{\text{SNIPE}}] = \tau_n$ at every model of $\mathcal{M}_{n,d,\beta}(B)$ and of $\mathcal{M}^{\infty}_{n,d,\beta}(B)$, since only the low-order and support properties of the schedule were used.

6. **Unclipped SNIPE risk.** Fix a model in either class and write

$$W_i(z) := Y_i(z)g_{i,\beta,p}(z), \qquad F_i(z) := W_i(z) - \mathbb{E}_Z[W_i(Z)].$$

Each F_i has design mean zero and depends on the assignment only through $\{Z_j : j \in N_i\}$, because both Y_i and $g_{i,\beta,p}$ do. Expanding each F_i in the orthogonal basis of centered monomials $\prod_{j\in S}(Z_j - p)$, only subsets $S \subseteq N_i$ receive a nonzero coefficient, and the empty subset receives none. For a fixed nonempty S, the units that can contribute are those with $S \subseteq N_i$, and there are at most d of them: choosing $j \in S$ places every such unit in the out-neighborhood of j, which is the out-degree count in Lemma 2. Applying Cauchy–Schwarz to each such group of at most d coefficients and summing over S gives

$$\mathbb{E}_Z\left[\Big(\sum_{i\in V_n} F_i(Z)\Big)^2\right] \le d\sum_{i\in V_n}\mathbb{E}_Z\left[F_i(Z)^2\right].$$

For each unit, $|Y_i(z)| \le B$ — by the coefficient-mass envelope derived from Definitions 5 and 6 on $\mathcal{M}_{n,d,\beta}(B)$, and by Definition 12 on $\mathcal{M}^{\infty}_{n,d,\beta}(B)$ — while the exact score energy is $\mathbb{E}_Z[g_{i,\beta,p}(Z)^2] = A_i$. Hence, using the local energy comparison above,

$$\mathbb{E}_Z\left[F_i(Z)^2\right] \le \mathbb{E}_Z\left[W_i(Z)^2\right] \le B^2\,\mathbb{E}_Z\left[g_{i,\beta,p}(Z)^2\right] = B^2 A_i \le B^2 A_d.$$

The unbiasedness identity above makes the mean squared error equal to the variance, and

$$\mathbb{E}_Z\left[\big(\widehat{\tau}_n^{\text{SNIPE}} - \tau_n\big)^2\right] = \frac{1}{n^2}\mathbb{E}_Z\left[\Big(\sum_{i\in V_n} F_i(Z)\Big)^2\right] \le \frac{d}{n^2}\cdot n\cdot B^2 A_d = B^2\frac{dA_d}{n}.$$

Taking the supremum over each class,

$$\sup_{\mathcal{M}_{n,d,\beta}(B)}\mathbb{E}_Z\left[\big(\widehat{\tau}_n^{\text{SNIPE}} - \tau_n\big)^2\right] \le B^2\frac{dA_d}{n}, \qquad \sup_{\mathcal{M}^{\infty}_{n,d,\beta}(B)}\mathbb{E}_Z\left[\big(\widehat{\tau}_n^{\text{SNIPE}} - \tau_n\big)^2\right] \le B^2\frac{dA_d}{n}.$$

7. **Target envelopes and clipping.** On $\mathcal{M}_{n,d,\beta}(B)$ the all-treated versus all-control contrast of a single unit is $Y_i(\mathbf{1}) - Y_i(\mathbf{0}) = \sum_{\varnothing \neq S \subseteq N_i} c_{i,S}$, so by Assumption 4

$$|\tau_n| \leq \frac{1}{n} \sum_{i \in V_n} \sum_{\varnothing \neq S \subseteq N_i} |c_{i,S}| \leq B,$$

whereas on $\mathcal{M}^{\infty}_{n,d,\beta}(B)$ the uniform outcome envelope gives

$$|\tau_n| \leq \frac{1}{n} \sum_{i \in V_n} \big(|Y_i(\mathbf{1})| + |Y_i(\mathbf{0})|\big) \leq 2B.$$

For $R \geq 0$ and any target t with $|t| \leq R$, Euclidean projection onto $[-R, R]$ is a contraction towards t:

$$\big(\Pi_{[-R,R]}(x) - t\big)^2 \leq (x-t)^2 \qquad \text{for every } x \in \mathbb{R}.$$

Applying this with $R = B$, $t = \tau_n$ on $\mathcal{M}_{n,d,\beta}(B)$, and with $R = 2B$, $t = \tau_n$ on $\mathcal{M}^{\infty}_{n,d,\beta}(B)$, gives modelwise

$$\mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{up}}_n - \tau_n\big)^2\Big] \leq \mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n\big)^2\Big], \qquad \mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{up}}_{n,\infty} - \tau_n\big)^2\Big] \leq \mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{SNIPE}}_n - \tau_n\big)^2\Big].$$

The projections also give saturated bounds that do not involve the design at all: $|\widehat{\tau}^{\mathrm{up}}_n| \leq B$ together with $|\tau_n| \leq B$ yields $(\widehat{\tau}^{\mathrm{up}}_n - \tau_n)^2 \leq 4B^2$ pointwise, and $|\widehat{\tau}^{\mathrm{up}}_{n,\infty}| \leq 2B$ together with $|\tau_n| \leq 2B$ yields $(\widehat{\tau}^{\mathrm{up}}_{n,\infty} - \tau_n)^2 \leq 16B^2$ pointwise. Using the unprojected SNIPE risk bound above when $dA_d/n \leq 1$ and the saturated bounds when $dA_d/n > 1$,

$$\sup_{\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{up}}_n - \tau_n\big)^2\Big] \leq 4B^2 \min\Big\{1, \frac{dA_d}{n}\Big\},$$

$$\sup_{\mathcal{M}^{\infty}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{up}}_{n,\infty} - \tau_n\big)^2\Big] \leq 16B^2 \min\Big\{1, \frac{dA_d}{n}\Big\}.$$

Since $4 \leq 16 \leq C_{\beta,p}$ and $B^2 \min\{1, dA_d/n\} \geq 0$, both right-hand sides are at most $C_{\beta,p} B^2 \min\{1, dA_d/n\}$, which proves the coefficient-class clipped bound of the theorem.

8. **Ordering of the two minimax risks.** The identity

$$R^{\star}_{n,\ell_1}(d, \beta, p, B) = R^{\star}_n(d, \beta, p, B)$$

is the definition of the coefficient-mass risk in Definition 13.

The carrier-preserving inclusion established from Definitions 6 and 12 sends each $(G_n, c) \in \mathcal{M}_{n,d,\beta}(B)$ to a member of $\mathcal{M}^{\infty}_{n,d,\beta}(B)$ with the same graph and the same schedule, hence with the same observed outcomes and the same target; the modelwise risk of any estimator is therefore unchanged by the embedding. Consequently, for every graph-aware estimator $\widehat{\tau}_n$,

$$\sup_{\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n - \tau_n\big)^2\Big] \leq \sup_{\mathcal{M}^{\infty}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n - \tau_n\big)^2\Big].$$

The infima in Definitions 9 and 13 range over the same graph-, assignment-, and observed-outcome-measurable estimators. Taking the infimum over that common estimator class gives

$$R^{\star}_{n,\ell_1}(d, \beta, p, B) \leq R^{\star}_{n,\infty}(d, \beta, p, B).$$

The estimator $\widehat{\tau}^{\mathrm{up}}_{n,\infty}$ of Definition 14 is measurable in the graph, assignment, and observed outcomes. Therefore it is one of the estimators over which the bounded-outcome infimum in Definition 13 is taken. The infimum is at most the worst-case risk of this estimator, and the clipped-risk bound above gives

$$R^{\star}_{n,\infty}(d,\beta,p,B) \le \sup_{\mathcal{M}^{\infty}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}^{\mathrm{up}}_{n,\infty} - \tau_n\big)^2\Big] \le C_{\beta,p} B^2 \min\Big\{1, \frac{dA_d}{n}\Big\}.$$

9. **The least-favourable block family.** Assume $B > 0$ and $d \ge 1$, and use the notation of Definition 15: $m = \lfloor n/d \rfloor$, $q = md$, $\rho = q/n$, $s = B/2$, and $\delta = B\kappa\eta$, as specified in Definition 15 with the constant κ fixed above. Since $1 \le d \le n$ we have $m \ge 1$. Fix $U = (U_1, \dots, U_m)$ with $|U_b| \le B/2$ for all b, and let $\sigma \in \{-1, +1\}$.

   Membership. The signed schedule of Definition 15 assigns to each active unit $i \in V_b$ the coefficients $c_{i,T} = U_b \mathbf{1}\{T = \varnothing\} + \sigma\delta h_{d,T}$ for $T \subseteq V_b$, and zero to every other coefficient. It is supported in $N_i = V_b$ on the complete-block graph, and it vanishes for $|T| > \beta$ because $h_{d,T} = 0$ there, so Assumption 3 holds; the complete-block graph has both degrees equal to d, so Assumption 2 holds. For the mass, $\eta \le 1$ and $\kappa \le (2H_{\beta,p})^{-1}$ give

$$\sum_{T \subseteq N_i} |c_{i,T}| \le |U_b| + |\delta| \sum_T |h_{d,T}| \le \frac{B}{2} + B\kappa H_{\beta,p} \le \frac{B}{2} + \frac{B}{2} = B.$$

   Hence both signed schedules define models M_+ and M_- in $\mathcal{M}_{n,d,\beta}(B)$ carried by the complete-block graph.

   Targets. On an active unit the baseline U_b cancels in the all-treated versus all-control contrast, and $L_d(h_d) = h_d(\mathbf{1}_d) - h_d(\mathbf{0}_d) = 1$ by Lemma 1, so

$$Y_i(\mathbf{1}) - Y_i(\mathbf{0}) = \sigma\delta \qquad (i \text{ active}), \qquad Y_i(\mathbf{1}) - Y_i(\mathbf{0}) = 0 \qquad (i \text{ inactive}).$$

   There are $q = md$ active units, so

$$\tau_n(M_\sigma) = \frac{q}{n}\sigma\delta = \sigma\rho\delta, \qquad \text{i.e.} \qquad \tau_n(M_+) = \rho\delta, \quad \tau_n(M_-) = -\rho\delta.$$

   Hellinger distance. Every unit of an active block shares the block's outcome value, so the observation reduces to the pair $\big(Z, (Y_b)_{b \le m}\big)$, whose law under Π_σ has density

$$\Pr(Z = z) \prod_{b=1}^{m} f_s(y_b - \sigma\delta\, h_d(z_{V_b}))$$

   against counting measure on assignments times m-dimensional Lebesgue measure. That reference measure is the block dominating measure $\mu_{n,d}$ and the displayed function is the sign-σ block prior density π_σ of Lemma 3, which integrates to one against $\mu_{n,d}$; hence Π_+ and Π_- are probability laws with the common dominating measure $\mu_{n,d}$, and their squared Hellinger distance may be computed from the two densities. Conditionally on $Z = z$, the two laws Π_+ and Π_- are products of m translates of f_s with separations $2\delta h_d(z_{V_b})$, and the cosine-squared density has quadratic affinity defect

$$1 - \int \sqrt{f_s(u-a) f_s(u-a')}\, \mathrm{d}u \le \frac{\pi^2 (a-a')^2}{8s^2}.$$

Summing the defect over the m coordinates and using $H^2 = 2(1 - \text{affinity})$ with $s = B/2$ gives the conditional bound

$$H^2\big(\Pi_+(\cdot \mid z), \Pi_-(\cdot \mid z)\big) \le \frac{4\pi^2\delta^2}{B^2} \sum_{b=1}^{m} h_d(z_{V_b})^2.$$

Because the assignment marginal is common to the two laws, the unconditional squared Hellinger distance is the design average of the conditional ones. Each block sees an independent copy of the Bernoulli law, so $\mathbb{E}_Z[h_d(Z_{V_b})^2] = A_d^{-1}$ by Lemma 1, whence $\mathbb{E}_Z\big[\sum_b h_d(Z_{V_b})^2\big] = m/A_d$ and

$$H^2(\Pi_+, \Pi_-) \le \frac{4\pi^2 m\delta^2}{B^2 A_d}.$$

This is exactly the bound of Lemma 4, whose amplitude δ and complete-block count m are the ones fixed above.

Active share. From $q = md \le n$ we get $\rho \le 1$. Writing $n = q + (n - q)$ with $n - q < d \le q$ (the second inequality because $m \ge 1$) gives $n < 2q$, hence

$$\frac{1}{2} \le \rho \le 1.$$

Finally, $q = md$ and $\rho = q/n$ give $\rho A_d/m = dA_d/n$, that is,

$$\frac{A_d}{m} = \frac{dA_d/n}{\rho}.$$

10. **Lower frontier.** If $d = 0$ then $A_0 = 0$ by Definition 2 and the left side of the claimed bound is 0; if $B = 0$ it is 0 as well. In both cases the bound holds because the minimax risk is nonnegative: every modelwise risk is a mean squared error, the zero estimator is admissible so the infimum is over a nonempty set, and the empty graph with the zero schedule is a member of $\mathcal{M}_{n,d,\beta}(B)$.

    Assume now $B > 0$ and $d \ge 1$, and keep the notation of Definition 15. Since $\eta \le \sqrt{A_d/m}$ we have $\eta^2 \le A_d/m$, i.e.

    $$m\eta^2 \le A_d,$$

    and since $\kappa \le (4\pi)^{-1}$ we have $16\pi^2\kappa^2 \le 1$, i.e. $4\pi^2\kappa^2 \le \frac{1}{4}$. Substituting $\delta = B\kappa\eta$ into the preceding Hellinger bound,

    $$H^2(\Pi_+, \Pi_-) \le \frac{4\pi^2 m\delta^2}{B^2 A_d} = 4\pi^2\kappa^2 \frac{m\eta^2}{A_d} \le \frac{1}{4} \cdot \frac{m\eta^2}{A_d} \le \frac{1}{4}.$$

    In the unhalved convention the total variation distance is bounded by the Hellinger distance, so

    $$\sup_A \big|\Pi_+(A) - \Pi_-(A)\big| \le \sqrt{H^2(\Pi_+, \Pi_-)} \le \frac{1}{2}.$$

    The two displayed block targets $\pm\rho\delta$ are $2\rho\delta$ apart. Let $\widehat{\tau}_n$ be any admissible estimator and let T be its value read on the block construction, a measurable function of $\big(Z, (Y_b)_{b \le m}\big)$. Le Cam's two-point bound gives

    $$\max\Big\{\Pi_+\big(|T - \rho\delta| \ge \rho\delta\big),\ \Pi_-\big(|T + \rho\delta| \ge \rho\delta\big)\Big\} \ge \frac{1 - \frac{1}{2}}{2} = \frac{1}{4}.$$

Each of the two probabilities is dominated by the worst-case risk. Indeed, for each sign σ and each baseline vector U in the support of the prior, the block schedule described in Definition 15 is a member of $\mathcal{M}_{n,d,\beta}(B)$ whose target is $\sigma\rho\delta$, and pointwise in the assignment

$$(\rho\delta)^2\,\mathbf{1}\big\{|\widehat{\tau}_n - \sigma\rho\delta| \geq \rho\delta\big\} \leq \big(\widehat{\tau}_n - \sigma\rho\delta\big)^2,$$

so taking design expectations and then averaging over U — a probability law — gives

$$(\rho\delta)^2\,\Pi_\sigma\big(|T - \sigma\rho\delta| \geq \rho\delta\big) \leq \sup_{\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n - \tau_n\big)^2\Big].$$

Combining the last two displays, every admissible estimator has worst-case risk at least $(\rho\delta)^2/4$, and taking the infimum over admissible estimators,

$$\frac{(\rho\delta)^2}{4} \leq R^\star_{n,\ell_1}(d,\beta,p,B).$$

It remains to convert the scale. By the displayed identity $A_d/m = (dA_d/n)/\rho$ and $\rho \leq 1$,

$$\frac{dA_d}{n} = \rho\,\frac{A_d}{m} \leq \frac{A_d}{m}, \qquad \text{hence} \qquad \min\Big\{1, \frac{dA_d}{n}\Big\} \leq \min\Big\{1, \frac{A_d}{m}\Big\} = \eta^2,$$

the last equality because $\eta = \min\{1, \sqrt{A_d/m}\}$. Also $\rho \geq \frac{1}{2}$ gives $\rho^2/4 \geq 1/16$. Therefore, using $c_{\beta,p} \leq \kappa^2/16$,

$$\begin{aligned}
c_{\beta,p} B^2 \min\Big\{1, \frac{dA_d}{n}\Big\} &\leq \frac{\kappa^2}{16} B^2\eta^2 \\
&\leq \frac{\rho^2}{4}\,B^2\kappa^2\eta^2 \\
&= \frac{(\rho B\kappa\eta)^2}{4} = \frac{(\rho\delta)^2}{4} \leq R^\star_{n,\ell_1}(d,\beta,p,B).
\end{aligned}$$

11. **Binomial scale.** Let $d \geq 1$ and write $k_\star = k_\star(d,\beta,p)$. Multiplying the energy scale of Lemma 1 by $d \geq 0$ and using the inequalities $c_{\beta,p} \leq c_1$ and $c_2 \leq C_{\beta,p}$ fixed with the constants above,

$$c_{\beta,p}\,d\binom{d}{k_\star} \leq c_1\,d\binom{d}{k_\star} \leq dA_d \leq c_2\,d\binom{d}{k_\star} \leq C_{\beta,p}\,d\binom{d}{k_\star}.$$

For $d = 0$ all three quantities vanish, since the factor d is zero.

12. **Exact complete-block risk.** Assume $d \geq 1$ and $d \mid n$, put $m = n/d$, and let G be the disjoint union of m complete directed d-blocks with loops. Since $n = md$,

$$\frac{B^2 A_d}{m} = B^2\frac{dA_d}{n}.$$

Suppose first $B > 0$. On G every unit has $N_i = V_{b(i)}$ with $|N_i| = d$, so the SNIPE score of Definition 7 is exactly the complete-block score of Definition 10 read on that unit's block,

$$g_{i,\beta,p}(z) = g_d\big(z_{V_{b(i)}}\big),$$

and consequently $\widehat{\tau}_n^{\mathrm{SNIPE}} = \widehat{\tau}_{w,t}$ for the block-local weights $w_{i,t} := g_d$, which satisfy the mean-zero and raw-moment restrictions of Definition 17 by Lemma 1. Apply Theorem 3 at the index with $n_t = n$, $d_t = d$, $m_t = m$ and $G_t = G$. It gives the exact blockwise representation of the fixed-graph worst-case risk and the lower bound $d^2 A_d \le \Psi_{b,t}(w)$ for every admissible w. For the canonical choice $w_{i,t} = g_d$ the matching upper bound

$$\Psi_{b,t}(w) \le d^2 A_d$$

holds on every block V_b. Indeed, fix a block V_b and a signed choice $(\varepsilon_i, S_i)_{i \in V_b}$ admissible in the maximum defining $\Psi_{b,t}$. The d summands are not copies of one another — they carry different subsets S_i and different signs ε_i — but every unit of the block carries the same weight $g_d(Z_b)$, which therefore factors out:

$$\sum_{i \in V_b} \varepsilon_i \big\{ g_d(Z_b) Z_{S_i} - \mathbf{1}(S_i \neq \varnothing) \big\} = g_d(Z_b)\Phi(Z_b) - \Lambda,$$

$$\Phi(Z_b) := \sum_{i \in V_b} \varepsilon_i Z_{S_i}, \qquad \Lambda := \sum_{i \in V_b} \varepsilon_i \mathbf{1}(S_i \neq \varnothing).$$

The representer identities of Lemma 1 give $\mathbb{E}_Z[g_d(Z_b) Z_S] = 1$ for every nonempty $S \subseteq V_b$ with $|S| \le \bar{\beta}_d$ and $\mathbb{E}_Z[g_d(Z_b)] = 0$, so $\mathbb{E}_Z[g_d(Z_b)\Phi(Z_b)] = \Lambda$ and expanding the square removes the centering term,

$$\mathbb{E}_Z\left[\big(g_d(Z_b)\Phi(Z_b) - \Lambda\big)^2\right] = \mathbb{E}_Z\left[\big(g_d(Z_b)\Phi(Z_b)\big)^2\right] - \Lambda^2 \le \mathbb{E}_Z\left[\big(g_d(Z_b)\Phi(Z_b)\big)^2\right].$$

Each of the d terms of Φ has modulus at most one, since $0 \le Z_{S_i} \le 1$ and $|\varepsilon_i| = 1$, so $|\Phi(Z_b)| \le d$ pointwise and therefore

$$\mathbb{E}_Z\left[\big(g_d(Z_b)\Phi(Z_b)\big)^2\right] \le d^2\, \mathbb{E}_Z\left[g_d(Z_b)^2\right] = d^2 A_d,$$

the last equality being the exact block-score energy $\mathbb{E}_Z[g_d(Z)^2] = A_d$ of Lemma 1. Maximizing over admissible choices gives the displayed upper bound. Hence $\Psi_{b,t}(w) = d^2 A_d$ on every block, and

$$\sup_{c \in \mathcal{M}_t(G)} \mathbb{E}_Z\left[\big(\widehat{\tau}_n^{\mathrm{SNIPE}} - \tau_n\big)^2\right] = \frac{B^2}{n^2} \sum_{b=1}^{m} \Psi_{b,t}(w) = \frac{B^2 m d^2 A_d}{n^2} = \frac{B^2 A_d}{m}.$$

If instead $B = 0$, the only admissible schedule on any graph has zero coefficient mass, so $\tau_n = 0$, while the unprojected SNIPE risk bound above bounds every modelwise SNIPE risk by $B^2 d A_d / n = 0$; as design mean squared errors are nonnegative, the fixed-graph coefficient-class risk again equals $0 = B^2 A_d / m$.

For all $B \ge 0$ the four risks now coincide. Restricting the supremum to a single graph can only decrease it, so the fixed-graph coefficient-class risk is at most the global coefficient-class risk, and the fixed-graph bounded-outcome risk is at most the global bounded-outcome risk. The carrier-preserving inclusion above embeds each coefficient-mass model on G in the bounded-outcome class without changing the graph, the schedule, or the risk, so the fixed-graph coefficient-class risk is at most the fixed-graph bounded-outcome risk. The unprojected SNIPE risk bound above bounds both global risks by $B^2 d A_d / n = B^2 A_d / m$, which is the value just computed for the fixed-graph coefficient-class risk. The chain of inequalities therefore closes, and

$$\sup_{c\in\mathcal{M}_t(G)} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n^{\mathrm{SNIPE}}-\tau_n\big)^2\Big] = \sup_{\substack{(G_n,c)\in\mathcal{M}^{\infty}_{n,d,\beta}(B)\\ G_n=G}} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n^{\mathrm{SNIPE}}-\tau_n\big)^2\Big] = \sup_{\mathcal{M}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n^{\mathrm{SNIPE}}-\tau_n\big)^2\Big]$$

$$= \sup_{\mathcal{M}^{\infty}_{n,d,\beta}(B)} \mathbb{E}_Z\Big[\big(\widehat{\tau}_n^{\mathrm{SNIPE}}-\tau_n\big)^2\Big] = \frac{B^2A_d}{m} = B^2\frac{dA_d}{n}.$$

The model comparison, the unbiasedness identity, the worst-case bound B^2dA_d/n for unclipped SNIPE, the clipped and minimax frontier bounds, the local energy comparison, the binomial scale, the exact complete-block risk, and the least-favourable block family together give the six assertions of the theorem. The constants $c_{\beta,p}$ and $C_{\beta,p}$ were fixed above independently of n, D, d and B. ☐

Proof of Theorem 4. 1. **Fair-coin contrasts.** Fix $n,d,\beta\in\mathbb{N}$ and $B\in\mathbb{R}$ with $n\geq 1$, $1\leq d\leq n$, $\beta\geq 1$, and $B\geq 0$. The contrast, energy, and exposed-order identities involve only d and β. The contrast coefficient is $\Delta_r(p)=(1-p)^r-(-p)^r$, so at $p=1/2$ every $r\geq 1$ satisfies

$$\Delta_r(1/2)=\left(\frac{1}{2}\right)^r-\left(-\frac{1}{2}\right)^r=2^{-r}-(-1)^r2^{-r}=2^{-r}\big\{1-(-1)^r\big\}.$$

Since $1-(-1)^r$ equals 2 for odd r and 0 for even r,

$$\Delta_r(1/2)=\begin{cases}2^{1-r}, & r\text{ odd},\\ 0, & r\text{ even}.\end{cases}$$

2. **Fair-coin energy.** By Definition 10 at $p=1/2$, the design variance factor is $[p(1-p)]^r=(1/4)^r$, so

$$A_d=\sum_{r=1}^{\bar\beta_d}\binom{d}{r}\frac{\Delta_r(1/2)^2}{\left[\frac{1}{2}\left(1-\frac{1}{2}\right)\right]^r}.$$

For every r with $1\leq r\leq\bar\beta_d$, the fair-coin contrast formula gives

$$\binom{d}{r}\frac{\Delta_r(1/2)^2}{\left[\frac{1}{2}\left(1-\frac{1}{2}\right)\right]^r}=\begin{cases}4\binom{d}{r}, & r\text{ odd},\\ 0, & r\text{ even},\end{cases}$$

because for odd r the numerator is $\binom{d}{r}(2^{1-r})^2=4\binom{d}{r}4^{-r}$ while the denominator is $(1/4)^r=4^{-r}$, and for even r the numerator vanishes. Summing the surviving terms,

$$A_d=4\sum_{\substack{1\leq r\leq\bar\beta_d\\ r\text{ odd}}}\binom{d}{r}.$$

3. **Fair-coin exposed order.** The index set in Definition 1 satisfies

$$\big\{r\in\{1,\dots,\bar\beta_d\}:\Delta_r(1/2)\neq 0\big\}=\big\{r\in\{1,\dots,\bar\beta_d\}:r\text{ odd}\big\},$$

since $\Delta_r(1/2)=2^{1-r}\neq 0$ for odd r and $\Delta_r(1/2)=0$ for even r. Hence $k_\star(d,\beta,1/2)$ is the maximum of the odd integers in $\{1,\dots,\bar\beta_d\}$ when that set is nonempty, and equals 0 otherwise, which is the value assigned by the zero-degree convention in Definition 1.

4. **Choice of constants.** Apply Theorem 2 with interaction order $\beta = 1$, which satisfies $\beta \geq 1$, and treatment probability $p = 1/2$, which satisfies $0 < p < 1$. It supplies constants $c_{1,1/2}$ and $C_{1,1/2}$ with
$$0 < c_{1,1/2} \leq C_{1,1/2}.$$
Set
$$c_{\text{lower}} := c_{1,1/2}, \qquad c_{\text{upper}} := 4C_{1,1/2}.$$
Then $c_{\text{lower}} > 0$, and since $0 < C_{1,1/2}$ gives $C_{1,1/2} \leq 4C_{1,1/2}$,
$$0 < c_{\text{lower}} \leq C_{1,1/2} \leq 4C_{1,1/2} = c_{\text{upper}}.$$

5. **Design.** Fix $n, d \in \mathbb{N}$ and $B \in \mathbb{R}$ with $n \geq 1$, $1 \leq d \leq n$, and $B \geq 0$, and take for P_Z the product Bernoulli law on V_n with common treatment probability $1/2$, that is,
$$\Pr(Z_j = 1) = \frac{1}{2} \qquad \text{for every } j \in V_n,$$
with $\{Z_j : j \in V_n\}$ mutually independent. Since $0 < 1/2 < 1$, this design satisfies Assumption 1 at $p = 1/2$, so the conclusions of Theorem 2 are available at (n, d, B) with $\beta = 1$ and $p = 1/2$.

6. **Degree-one energy.** With $\beta = 1$ and $d \geq 1$ the exposed order of Definition 1 is
$$\bar{\beta}_d = \min\{1, d\} = 1,$$
so the odd integers in $\{1, \ldots, \bar{\beta}_d\}$ are the single index $r = 1$. The fair-coin energy formula gives
$$A_d = 4\binom{d}{1} = 4d.$$

7. **Reduction of the rate argument.** Put
$$x := \frac{d^2}{n}.$$
Since $n \geq 1$, we have $x \geq 0$, and the identity $A_d = 4d$ gives
$$\frac{dA_d}{n} = \frac{d(4d)}{n} = 4x.$$
Two elementary comparisons of truncated arguments are used below. First, $x \leq 4x$ because $x \geq 0$, and truncation at 1 is monotone, so
$$\min\{1, x\} \leq \min\{1, 4x\}.$$
Second,
$$\min\{1, 4x\} \leq 4\min\{1, x\}.$$
Indeed, if $x \leq 1$ then $\min\{1, x\} = x$, and either $4x \leq 1$, in which case $\min\{1, 4x\} = 4x = 4\min\{1, x\}$, or $4x > 1$, in which case $\min\{1, 4x\} = 1 \leq 4x = 4\min\{1, x\}$; while if $x > 1$ then $\min\{1, x\} = 1$ and $\min\{1, 4x\} \leq 1 \leq 4 = 4\min\{1, x\}$.

8. **Lower bound.** Because $c_{\text{lower}} > 0$ and $B \geq 0$,

$$0 \leq c_{\text{lower}} B^2,$$

so multiplying $\min\{1, x\} \leq \min\{1, 4x\}$ by this factor preserves the inequality. The lower frontier bound of Theorem 2 at $\beta = 1$, $p = 1/2$ reads

$$c_{\text{lower}} B^2 \min\left\{1, \frac{dA_d}{n}\right\} \leq R^\star_{n,\ell_1}(d, 1, 1/2, B),$$

and the same result identifies

$$R^\star_{n,\ell_1}(d, 1, 1/2, B) = R^\star_n(d, 1, 1/2, B).$$

Using $x = d^2/n$ and $dA_d/n = 4x$,

$$c_{\text{lower}} B^2 \min\left\{1, \frac{d^2}{n}\right\} = c_{\text{lower}} B^2 \min\{1, x\} \leq c_{\text{lower}} B^2 \min\{1, 4x\} = c_{\text{lower}} B^2 \min\left\{1, \frac{dA_d}{n}\right\} \leq R^\star_n(d, 1, 1/2, B$$

9. **Class comparison.** The carrier inclusion recorded in Theorem 2 supplies

$$R^\star_{n,\ell_1}(d, 1, 1/2, B) \leq R^\star_{n,\infty}(d, 1, 1/2, B).$$

10. **Upper bound.** The clipped bounded-outcome estimator $\widehat{\tau}^{\text{up}}_{n,\infty}$ of Definition 14 is one admissible competitor in the bounded-outcome minimax problem, and Theorem 2 bounds its worst-case risk, giving

$$R^\star_{n,\infty}(d, 1, 1/2, B) \leq \sup_{(G,c) \in \mathcal{M}^\infty_{n,d,1}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}^{\text{up}}_{n,\infty} - \tau_n(G, c)\right)^2\right] \leq C_{1,1/2} B^2 \min\left\{1, \frac{dA_d}{n}\right\}.$$

Since $C_{1,1/2} > 0$ and $B \geq 0$, the factor $C_{1,1/2} B^2$ is nonnegative, so multiplying $\min\{1, 4x\} \leq 4\min\{1, x\}$ by it preserves the inequality. With $dA_d/n = 4x$, this yields

$$C_{1,1/2} B^2 \min\left\{1, \frac{dA_d}{n}\right\} = C_{1,1/2} B^2 \min\{1, 4x\} \leq 4C_{1,1/2} B^2 \min\{1, x\} = c_{\text{upper}} B^2 \min\left\{1, \frac{d^2}{n}\right\},$$

and therefore

$$R^\star_{n,\infty}(d, 1, 1/2, B) \leq c_{\text{upper}} B^2 \min\left\{1, \frac{d^2}{n}\right\}.$$

Combining the lower bound, the equality of the two coefficient-class risk notations, the class comparison, and the upper bound gives the asserted chain

$$c_{\text{lower}} B^2 \min\left\{1, \frac{d^2}{n}\right\} \leq R^\star_n(d, 1, 1/2, B) = R^\star_{n,\ell_1}(d, 1, 1/2, B) \leq R^\star_{n,\infty}(d, 1, 1/2, B) \leq c_{\text{upper}} B^2 \min\left\{1, \frac{d^2}{n}\right\}.$$

11. **Exact complete-block identities.** Assume in addition that $d \mid n$ and set $m := n/d$, the number of complete d-blocks. The exact-block-risk clause of Theorem 2 gives, for the unprojected SNIPE estimator of Definition 8,

$$\sup_{(G,c) \in \mathcal{M}_{n,d,1}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}^{\text{SNIPE}}_n - \tau_n(G, c)\right)^2\right] = \frac{B^2 A_d}{m} = \sup_{(G,c) \in \mathcal{M}^\infty_{n,d,1}(B)} \mathbb{E}_Z\left[\left(\widehat{\tau}^{\text{SNIPE}}_n - \tau_n(G, c)\right)^2\right],$$

together with the identity

$$\frac{B^2 A_d}{m} = \frac{B^2 d A_d}{n}.$$

Substituting $A_d = 4d$,

$$\frac{B^2 d A_d}{n} = \frac{B^2 d(4d)}{n} = \frac{4B^2 d^2}{n},$$

which is the asserted common value of both worst-case risks.

□

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.

Susan Athey, Dean Eckles, and Guido W. Imbens. Exact¡i¿p¡/i¿-values for network interference. *Journal of the American Statistical Association*, 113(521):230–240, 2017. doi: 10.1080/01621459.2016.1241178.

Sarah Baird, J. Aislinn Bohren, Craig McIntosh, and Berk Özler. Optimal design of experiments in the presence of interference. *The Review of Economics and Statistics*, 100(5):844–860, 2018. doi: 10.1162/rest_a_00716.

Guillaume Basse and Avi Feller. Analyzing two-stage experiments in the presence of interference. *Journal of the American Statistical Association*, 113(521):41–55, 2018. doi: 10.1080/01621459.2017.1323641.

Guillaume W. Basse and Edoardo M. Airoldi. Model-assisted design of experiments in the presence of network-correlated outcomes. *Biometrika*, 105(4):849–858, 2018. doi: 10.1093/biomet/asy036.

Jake Bowers, Mark M. Fredrickson, and Costas Panagopoulos. Reasoning about interference between units: A general framework. *Political Analysis*, 21(1):97–124, 2013. doi: 10.1093/pan/mps038.

Yann Bramoullé, Habiba Djebbari, and Bernard Fortin. Identification of peer effects through social networks. *Journal of Econometrics*, 150(1):41–55, 2009. doi: 10.1016/j.jeconom.2008.12.021.

Mayleen Cortez-Rodriguez, Matthew Eichhorn, and Christina Lee Yu. Exploiting neighborhood interference with low-order interactions under unit randomized design. *Journal of Causal Inference*, 11(1):20220051, 2023. doi: 10.1515/jci-2022-0051.

Dean Eckles, Brian Karrer, and Johan Ugander. Design and analysis of experiments in networks: Reducing bias from interference. *Journal of Causal Inference*, 5(1):1–23, 2017. doi: 10.1515/jci-2015-0021.

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.

Mengsi Gao and Peng Ding. Causal inference in network experiments: Regression-based analysis and design-based properties. *Journal of Econometrics*, 252:106119, 2025. doi: 10.1016/j.jeconom.2025.106119.

Bryan S. Graham. Identifying social interactions through conditional variance restrictions. *Econometrica*, 76(3):643–660, 2008. doi: 10.1111/j.1468-0262.2008.00850.x.

Christopher Harshaw, Fredrik Sävje, and Yitan Wang. A general design-based framework and estimator for randomized experiments, 2022.

Daniel G. Horvitz and Donovan J. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952. doi: 10.1080/01621459.1952.10483446.

Yuchen Hu, Shuangning Li, and Stefan Wager. Average direct and indirect causal effects under interference. *Biometrika*, 109(4):1165–1172, 2022. doi: 10.1093/biomet/asac008.

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.

Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, Cambridge, 2015. ISBN 9780521885881. doi: 10.1017/CBO9781139025751.

Ravi Jagadeesan, Natesh S. Pillai, and Alexander Volfovsky. Designs for estimating the treatment effect in networks with interference. *The Annals of Statistics*, 48(2):679–712, 2020. doi: 10.1214/18-AOS1807.

Michael P. Leung. Causal inference under approximate neighborhood interference. *Econometrica*, 90(1):267–293, 2022. doi: 10.3982/ECTA17841.

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

Lan Liu, Michael G. Hudgens, and Sylvia Becker-Dreps. On inverse probability-weighted estimators in the presence of interference. *Biometrika*, 103(4):829–842, 2016. doi: 10.1093/biomet/asw047.

Charles F. Manski. Identification of endogenous social effects: The reflection problem. *The Review of Economic Studies*, 60(3):531–542, 1993. doi: 10.2307/2298123.

Charles F. Manski. Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1):S1–S23, 2013. doi: 10.1111/j.1368-423X.2012.00368.x.

Elizabeth L. Ogburn and Tyler J. VanderWeele. Vaccines, contagion, and social networks. *The Annals of Applied Statistics*, 11(2):919–948, 2017. doi: 10.1214/17-AOAS1023.

Donald B. Rubin. Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6(1):34–58, 1978. doi: 10.1214/aos/1176344064.

Michael E Sobel. What do randomized studies of housing mobility demonstrate? *Journal of the American Statistical Association*, 101(476):1398–1407, 2006. doi: 10.1198/016214506000000636.

Adith Swaminathan, Akshay Krishnamurthy, Alekh Agarwal, Miroslav Dudík, John Langford, Damien Jose, and Imed Zitouni. Off-policy evaluation for slate recommendation. In *Advances in Neural Information Processing Systems 30*, pages 3632–3642, 2017.

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.

Eric J Tchetgen Tchetgen and Tyler J VanderWeele. On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1):55–75, 2010. doi: 10.1177/0962280210386779.

Johan Ugander, Brian Karrer, Lars Backstrom, and Jon Kleinberg. Graph cluster randomization. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 329–337. ACM, 2013. doi: 10.1145/2487575.2487695.

Mark J. van der Laan. Causal inference for a population of causally connected units. *Journal of Causal Inference*, 2(1):13–74, 2014. doi: 10.1515/jci-2013-0002.

Gonzalo Vazquez-Bare. Identification and estimation of spillover effects in randomized experiments. *Journal of Econometrics*, 237(1):1–26, 2023. doi: 10.1016/j.jeconom.2021.10.014.

Xinyi Wang and Shuangning Li. Covariate adjustment cannot hurt: Treatment effect estimation under interference with low-order outcome interactions, 2026.

Christina Lee Yu, Edoardo M. Airoldi, Christian Borgs, and Jennifer T. Chayes. Estimating the total treatment effect in randomized experiments with unknown network structure. *Proceedings of the National Academy of Sciences*, 119(44):e2208975119, 2022. doi: 10.1073/pnas.2208975119.