Exact Randomized Designs for Two-Block Interference Experiments

CausalSmith*

August 25, 2026

Abstract

This paper studies when a covariance relaxation for interference-aware randomized design is exactly implementable by a finite $\{\pm 1\}$ assignment law. The setting is a design-based finite-population model with two equal homophilous communities, sign-symmetric assignment, and a block-weighted graph. The design objective combines a graph Laplacian term, a Laplacian-pseudoinverse term, a Schatten–2 robustness penalty, and an aggregate balance penalty. A symmetry reduction shows that the relaxed problem is exactly the optimization over a two-parameter block elliptope, while the implementable problem is the same objective restricted to covariances induced by block-exchangeable sign designs. The relaxation is exact in a strict cut region, where the unique relaxed optimum is generated by a two-point randomized design. Independent fair assignment is implementable throughout the model and is a finite-robustness relaxed optimum precisely on the affine locus $a + 3b = 2m$ and $r = 2b(a + b)$; elsewhere it emerges as the asymptotic target as the robustness weight diverges. For odd community size, parity creates an open parameter region with a uniquely optimal relaxed covariance and strictly positive implementability loss. Finally, an exact finite active-set formula computes the loss for all admissible m, a, b, r, κ, with zero loss for even community size.

1 Introduction

Randomized experiments under interference require choosing both marginal treatment probabilities and the dependence structure of the assignment. In a design-based finite-population analysis, the units, potential outcomes, and network are fixed, and randomization comes from the assignment law. When interference is mediated by a graph, the covariance matrix of the assignment can therefore become a design object: it determines how treatment contrasts align with graph smoothness, aggregate balance, and robustness criteria. This perspective is closely related to the randomization-based tradition in causal inference [Horvitz and Thompson, 1952, Rubin, 1974], to design and inference under interference [Hudgens and Halloran, 2008, Aronow and Samii, 2013, Sävje et al., 2017, Leung, 2019, Li and Wager, 2020], and to recent covariance-based design formulations for network experiments [Thiyageswaran et al., 2026].

A covariance relaxation is attractive because it replaces a finite randomized assignment problem by a convex matrix problem. Its feasible set can be strictly larger than the covariances

*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.

generated by actual $\{\pm 1\}$ assignment laws, producing the familiar gap between semidefinite or elliptope relaxations and discrete correlation structures [Goemans and Williamson, 1995, Deza and Laurent, 1997]. In experimental design, this gap has an operational meaning: implementation requires a covariance prescription induced by a randomization mechanism.

This paper isolates that issue in a finite two-community model. There are $n = 2m$ fixed units, split into two equal communities. The graph has within-community edge weight a/m and across-community edge weight b/m, with $a > b > 0$. Assignments are sign vectors $Z \in \{-1, 1\}^n$, and assignment laws are required to be sign-symmetric. The objective $F_{r,\kappa}$ in Definition 6 combines four terms: graph Laplacian exposure, a Laplacian-pseudoinverse component weighted by r, a Schatten–2 robustness penalty weighted by κ, and an aggregate balance penalty. The object of interest is the implementability loss $\Delta_m^{\pm}(r, \kappa)$, the difference between the best implementable block-exchangeable sign design and the relaxed block-elliptope benchmark. Proposition 1 establishes equality with the infimum over the broader sign-symmetric class $\mathcal{P}_m^{\mathrm{bal}}$.

The first step is a symmetry reduction. Under Assumption 1, Proposition 1 shows that averaging over graph symmetries weakly decreases the objective for both relaxed covariance matrices and sign-symmetric assignment laws. Hence the relaxed infimum is exactly the infimum over $\mathcal{E}_m^{\mathrm{blk}}$, and the implementable infimum is exactly the infimum over $\mathcal{P}_m^{\mathrm{sym}}$. This reduction makes the implementability question finite-dimensional.

The paper then gives two exactness results. Under Assumption 1, Theorem 1 shows that, for $r, \kappa \geq 0$ satisfying the displayed strict cut-region inequality, the cut covariance X_{cut} is the unique relaxed minimizer. It is also generated by the symmetric two-point cut design, so the implementability loss is zero throughout that region. This exact finite-sample result supplies a transparent sufficient region for literal implementation of the relaxed optimum.

The second exactness result concerns iid assignment. Under Assumption 1 and $r \geq 0$, Theorem 2 shows that I_n is a finite-robustness relaxed minimizer precisely when $a + 3b = 2m$ and $r = 2b(a + b)$. On that affine-balanced locus, every positive κ makes I_n the unique relaxed minimizer. Elsewhere, relaxed minimizers converge entrywise to I_n as $\kappa \to \infty$. Thus iid assignment is always implementable in this symmetric class, with exact finite-κ optimality characterized by the affine-balanced locus.

The main positive-loss result is arithmetic. When m is odd, each realized block sum has odd parity and therefore absolute value at least one. Under the homophily, low-scale, and odd-community conditions, Theorem 3 identifies a nonempty robustness range and an open interval of pseudoinverse weights in which the spread covariance X_{spread} is the unique relaxed minimizer, lies outside $\mathcal{C}_m^{\pm}$, and yields $\Delta_m^{\pm}(r, \kappa) > 0$. The same theorem also shows that every covariance in the block elliptope is implementable when m is even.

For computation outside that local positive-gap window, Theorem 4 gives an exact finite formula. Under Assumption 1, for every $r, \kappa \geq 0$, the loss equals a sharp reduced quantity $\rho_\star(m, a, b, r, \kappa)$. The relaxed problem becomes a weighted simplex problem in three spectral coordinates, and implementability adds only the parity truncation $d_m \leq y + z_{\mathrm{sp}}$, where $d_m = 0$ for even m and $d_m = 2/m$ for odd m. With positive robustness, the relaxed minimizer is unique and is obtained by an active-set calculation; if it violates the parity truncation, the implementable value is obtained on a one-dimensional boundary segment. With $\kappa = 0$, the theorem gives the corresponding exposed-face condition. The displayed assumptions, definitions, lemmas, and theorems delimit the paper's mathematical scope, while the motivating discussion, related-work positioning, and interpretive remarks are ordinary prose.

The contribution is an exact finite-sample study of the covariance criterion $F_{r,\kappa}$ in a homophilous two-block design. Within that model it characterizes literal implementability, exhibits an open parameter region of strictly positive implementability loss when the community size is odd, and gives an exact finite formula for the objective-value price of returning from the relaxation to a feasible randomized assignment.

2 Setup and assumptions

We work in a design-based finite-population assignment model: the units are fixed, and the only random object is the treatment assignment. There are $n = 2m$ units partitioned into two equal communities. An assignment is represented by the sign vector $Z \in \{-1, 1\}^n$, with a realized assignment denoted by z, and an assignment law P is a probability mass function on this finite assignment space. This formulation follows the randomization-based tradition in which the assignment mechanism is the source of sampling variation [Horvitz and Thompson, 1952, Rubin, 1974], while the network structure enters through the design criterion rather than through a superpopulation model.

The graph used by the criterion is a two-community weighted graph. The analysis rests first on the following homophily and scale conventions.

Assumption 1. [`ass:two-block-homophily`] (Two-Block Homophily) *The block size and edge intensities satisfy*

$$m \geq 2 \qquad \text{and} \qquad a > b > 0.$$

The constants a and b encode stronger interaction intensity within a community than across communities. This homophily restriction is specific to this analysis: it gives the two-block model a nondegenerate within-versus-between contrast while keeping both edge intensities positive.

Definition 1. [`def:two-block-graph`] (Two-Block Graph Weights W_{ij}) *The vertex set is the disjoint union $A_m \cup B_m$, where*

$$|A_m| = |B_m| = m.$$

The graph weights W_{ij} are given as follows: if $i \neq j$ and i and j lie in the same block, then

$$W_{ij} = \frac{a}{m};$$

if i and j lie in opposite blocks, then

$$W_{ij} = \frac{b}{m};$$

and for every vertex i,

$$W_{ii} = 0.$$

Thus the communities A_m and B_m have equal size, and all heterogeneity in exposure strength is summarized by whether a pair lies within or across communities. Let L_m denote the graph Laplacian associated with these weights, with diagonal entries equal to weighted degrees and off-diagonal entries $-W_{ij}$. We write $L_m^\dagger$ for its Moore–Penrose pseudoinverse. Network-interference applications often use such weighted graphs to encode exposure or spillover structure [Hudgens and Halloran, 2008, Aronow and Samii, 2013, Sävje et al., 2017, Leung, 2019, Li and Wager, 2020]; here the two-block graph isolates the exact finite-sample geometry of a homophilous special case.

The sign representation is invariant to a global relabeling of treatment and control. We impose that invariance directly on the assignment law.

Assumption 2. [ass:balanced-sign-design] (Balanced Sign Symmetry) *The assignment law is sign-symmetric: for every $z \in \{-1, 1\}^n$,*

$$P(Z = z) = P(Z = -z).$$

This symmetry condition is specific to this analysis. It expresses balance between the two treatment labels at the law level, yields zero one-point margins, and permits many forms of dependence and varying treated counts across realizations.

Definition 2. [def:balanced-design-class] (Balanced sign class $\mathcal{P}_m^{\mathrm{bal}}$) *Under Assumption 2, the balanced sign class is*

$$\mathcal{P}_m^{\mathrm{bal}} = \{P \text{ on } \{-1, 1\}^n : P(Z = z) = P(Z = -z) \text{ for every } z \in \{-1, 1\}^n\}.$$

The class $\mathcal{P}_m^{\mathrm{bal}}$ contains all sign-symmetric laws. For the two-block graph, the relevant symmetric subclass also treats units within a block exchangeably and treats the two equal blocks symmetrically.

Definition 3. [def:block-exchangeable-design-class] (Block-exchangeable class $\mathcal{P}_m^{\mathrm{sym}}$) *The block-exchangeable class is the set of all assignment laws*

$$P \in \mathcal{P}_m^{\mathrm{bal}}$$

that are invariant under permutations within A_m, invariant under permutations within B_m, and invariant under swapping the two equal communities.

For optimization of $F_{r,\kappa}$ over the sign-symmetric class $\mathcal{P}_m^{\mathrm{bal}}$, the block-exchangeable class $\mathcal{P}_m^{\mathrm{sym}}$ preserves the optimum: averaging over the two-block automorphism group weakly decreases this invariant objective. This reduction makes the maintained implementability question finite-dimensional.

For any assignment law, write $X(P)$ for the second-moment matrix,

$$X(P) = \mathbb{E}_P[ZZ^\top].$$

This is a unit-diagonal second-moment (correlation) matrix; under the sign symmetry of Assumption 2 the one-point margins vanish, so for implementable laws it coincides with the covariance of Z. We refer to X as a second-moment or correlation matrix throughout, reserving "covariance" for the sign-symmetric case; the relaxed matrices below are unit-diagonal correlation matrices that need not arise from any law. The covariance relaxation replaces such implementable second moments by block-symmetric correlation matrices. For scalars u and v, $X(u, v)$ denotes the matrix with diagonal entries equal to one, common off-diagonal covariance u within each block, and common covariance v across the two blocks.

Definition 4. [def:block-elliptope] (Block Elliptope $\mathcal{E}_m^{\mathrm{blk}}$) *Under Assumption 1, the block elliptope is*

$$\mathcal{E}_m^{\mathrm{blk}} = \{X(u, v) : 1 - u \geq 0,\ 1 + (m-1)u - mv \geq 0,\ 1 + (m-1)u + mv \geq 0\}.$$

The relaxed covariance set $\mathcal{E}_m^{\mathrm{blk}}$ is the block-symmetric slice of the elliptope. The three displayed inequalities are the positive-semidefiniteness restrictions in the invariant directions induced by the two equal blocks. This convex relaxation is the object optimized before imposing that the covariance arise from an actual $\{\pm 1\}$ assignment law [Boyd and Vandenberghe, 2004, Thiyageswaran et al., 2026].

Definition 5. [`def:implementable-covariance-class`] (Implementable covariance class $\mathcal{C}_m^{\pm}$) *The implementable covariance class is*

$$\mathcal{C}_m^{\pm} = \{X(P) : P \in \mathcal{P}_m^{\mathrm{sym}}\}.$$

The set $\mathcal{C}_m^{\pm}$ records the discrete feasibility constraint that remains after symmetry reduction: a relaxed covariance matrix is useful for exact randomized design only when it is generated by some block-exchangeable sign law.

The objective combines graph smoothness, a pseudoinverse term, robustness, and a balance penalty. The scalar r weights the pseudoinverse term, κ weights the Schatten–2 norm penalty, and J_n denotes the $n \times n$ all-ones matrix.

Definition 6. [`def:design-objective`] (Design objective $F_{r,\kappa}$) *For $r, \kappa \geq 0$, the design objective evaluated at a covariance matrix X is*

$$F_{r,\kappa}(X) = \mathrm{tr}(L_m X) + r\ \mathrm{tr}(L_m^{\dagger} X) + \kappa \|X\|_{S_2} + \mathrm{tr}(J_n X).$$

The first term is the expected weighted disagreement of the sign assignment across graph edges: because $\mathrm{tr}(L_m X) = \sum_{i<j} W_{ij}\, \mathbb{E}\big[(Z_i - Z_j)^2\big] = 2 \sum_{i<j} W_{ij}\big(1 - X_{ij}\big)$ for a unit-diagonal second moment, minimizing it penalizes sign dispersion across high-weight edges and rewards positive treatment-sign alignment there. The second term uses the Laplacian pseudoinverse to weight low-frequency graph directions, the third is a Frobenius-norm robustness penalty, and the final term $\mathrm{tr}(J_n X) = \mathbb{E}\big[(\sum_i Z_i)^2\big]$ penalizes imbalance in the aggregate sign vector. We treat $F_{r,\kappa}$ as a covariance design criterion in the spirit of the network-experiment SDP formulation of Thiyageswaran et al. [2026]. The paper compares this criterion over the relaxed block elliptope and implementable sign designs and characterizes exactly when the optima coincide. The implementability loss measures their objective-value gap.

Definition 7. [`def:implementability-gap`] (Implementability loss $\Delta_m^{\pm}(r, \kappa)$) *The implementability loss is*

$$\Delta_m^{\pm}(r, \kappa) = \inf_{P \in \mathcal{P}_m^{\mathrm{sym}}} F_{r,\kappa}(X(P)) - \inf_{X \in \mathcal{E}_m^{\mathrm{blk}}} F_{r,\kappa}(X).$$

The quantity $\Delta_m^{\pm}(r, \kappa)$ is therefore the exact price of requiring a covariance matrix to be induced by a finite randomized assignment law rather than merely satisfying the block elliptope constraints.

A useful handle on implementability is provided by the two block sums. For a given law P, these sums retain the information relevant to the block-symmetric second moments.

Definition 8. [`def:block-sum-handle`] (Block-sum law $\mathcal{L}_P(S_A, S_B)$) *For an assignment law P, the block-sum law is the law under P of*

$$(S_A, S_B) = \left(\sum_{i \in A_m} Z_i,\ \sum_{i \in B_m} Z_i \right).$$

The block-sum law $\mathcal{L}_P(S_A, S_B)$ will be used to express the discrete restrictions that separate $\mathcal{C}_m^{\pm}$ from $\mathcal{E}_m^{\mathrm{blk}}$.

Two additional restrictions are stated here because later exact-loss results invoke them explicitly.

Assumption 3. [ass:low-scale-two-block] (Low-Scale Two-Block Regime) *The two-block edge intensities satisfy*

$$a + 3b < 2m.$$

The low-scale regime is specific to this analysis. It is a finite-sample scale restriction comparing graph intensity with block size; for fixed a and b, it is automatically satisfied once m is large enough.

Assumption 4. [ass:odd-community-size] (Odd Community Size) *The block size m is odd.*

The odd-block condition is also specific to this analysis. It isolates the parity case in which every realized block sum has absolute value at least one, which is exactly where the later implementability loss result has content.

The following reduction justifies working with the block elliptope and the block-exchangeable assignment class while preserving the relevant optima.

Proposition 1. [prop:symmetry-reduction] (Symmetry Reduction Principle) *Assume Assumption 1. Let $r, \kappa \in [0, \infty)$. Then:*

- ***(Relaxed matrices.)*** *For every positive semidefinite matrix $X \in \mathbb{R}^{n\times n}$ with $X_{ii} = 1$ for every i, let u and v be the averages of the off-diagonal within-community and across-community entries of X, respectively. The block matrix $X(u, v)$ of Definition 4 has feasible block-elliptope parameters, and*

$$F_{r,\kappa}\big(X(u, v)\big) \le F_{r,\kappa}(X),$$

where $F_{r,\kappa}$ is Definition 6.

- ***(Assignment laws.)*** *For every assignment law $P \in \mathcal{P}_m^{\mathrm{bal}}$ in the sense of Definition 2, its orbit average over the full two-block automorphism group belongs to $\mathcal{P}_m^{\mathrm{sym}}$ in the sense of Definition 3, has zero one-point margins,*

$$\mathbb{E}[Z_i] = 0 \qquad \textit{for every } i,$$

and satisfies

$$F_{r,\kappa}\big(X(P^{\mathrm{orb}})\big) \le F_{r,\kappa}\big(X(P)\big).$$

- ***(Relaxed infimum.)***

$$\inf_{\substack{X\succeq 0\\ X_{ii}=1\ \forall i}} F_{r,\kappa}(X) = \inf_{X\in\mathcal{E}_m^{\mathrm{blk}}} F_{r,\kappa}(X).$$

- ***(Implementable infimum.)***

$$\inf_{P\in\mathcal{P}_m^{\mathrm{bal}}} F_{r,\kappa}\big(X(P)\big) = \inf_{P\in\mathcal{P}_m^{\mathrm{sym}}} F_{r,\kappa}\big(X(P)\big).$$

Proposition 1 is the bridge between the original finite assignment problem and the two-parameter geometry used in the rest of the paper. It says that, for the objective in Definition 6, symmetrizing either a relaxed covariance matrix or an assignment law is weakly improving. Consequently, the relaxed benchmark is exactly the infimum over $\mathcal{E}_m^{\mathrm{blk}}$, and the implementable benchmark is exactly the infimum over $\mathcal{P}_m^{\mathrm{sym}}$.

3 Exactness at implementable designs

The symmetry reduction in Proposition 1 leaves a precise question: when does the relaxed optimum over $\mathcal{E}_m^{\mathrm{blk}}$ coincide with a covariance matrix generated by an actual block-exchangeable sign design? This section records two exact cases. We use three block covariances: the cut covariance $X_{\mathrm{cut}} = X(1, -1)$, the identity $I_n = X(0, 0)$ corresponding to independent fair Bernoulli signs, and the spread covariance $X_{\mathrm{spread}} = X\left(-\frac{1}{m-1}, 0\right)$. The implementing laws are the two-point cut law P_{cut} and the independent fair-sign law P_{iid}. These designs represent different design instincts: one separates the homophilous communities, while the other spreads randomization independently across all units [Efron, 1971, Morgan and Rubin, 2012, Harshaw et al., 2019].

The cut design is optimal in a strict small-r region. There, the graph and balance terms favor separating the two blocks strongly enough to keep the relaxed optimum at the cut corner after including the pseudoinverse and robustness terms.

Theorem 1. [`thm:cut-corner-exactness`] (Cut-Corner Exactness) *Assume the two-block homophily condition in Assumption 1. Let $r, \kappa \in \mathbb{R}$ satisfy*

- ***(Nonnegative tradeoff.)*** $0 \le r$.
- ***(Nonnegative robustness.)*** $0 \le \kappa$.
- ***(Cut region.)***

$$r < r_{\mathrm{cut}}(m, a, b, \kappa) := \max\left\{0,\ \min\left(2b(a+b)\left(1 - \frac{\kappa}{a-b}\right), 2b(2m - 2b - \kappa)\right)\right\}.$$

Then $X_{\mathrm{cut}} \in \mathcal{E}_m^{\mathrm{blk}}$, and for every $X \in \mathcal{E}_m^{\mathrm{blk}}$ with $X \neq X_{\mathrm{cut}}$,

$$F_{r,\kappa}(X_{\mathrm{cut}}) < F_{r,\kappa}(X),$$

where $F_{r,\kappa}$ is the objective in Definition 6. Moreover,

$$\Delta_m^{\pm}(r, \kappa) = 0,$$

with $\Delta_m^{\pm}$ as in Definition 7. Finally, $X_{\mathrm{cut}} \in \mathcal{C}_m^{\pm}$, and for every $X \in \mathcal{C}_m^{\pm}$ with $X \neq X_{\mathrm{cut}}$,

$$F_{r,\kappa}(X_{\mathrm{cut}}) < F_{r,\kappa}(X),$$

where $\mathcal{C}_m^{\pm}$ is the class in Definition 5; the cut design satisfies

$$P_{\mathrm{cut}} \in \mathcal{P}_m^{\mathrm{sym}} \qquad \text{and} \qquad X(P_{\mathrm{cut}}) = X_{\mathrm{cut}},$$

with $\mathcal{P}_m^{\mathrm{sym}}$ as in Definition 3.

Theorem 1 is an exact attainability statement. The displayed strict inequality involving $r_{\mathrm{cut}}(m, a, b, \kappa)$ gives a sufficient region in which the relaxed block-elliptope optimum is unique, implementable, and has zero implementability loss. The implementing law P_{cut} is the symmetric two-point law on the cut assignment and its global sign reversal. This result characterizes optimality for the stated covariance criterion $F_{r,\kappa}$.

The second exactness result concerns the robust corner I_n. Independent fair assignment is always implementable, but the theorem below shows that its relaxed optimality is much more restrictive: a finite positive robustness weight can make I_n optimal only on a particular affine relation among the graph scale and the pseudoinverse weight.

Theorem 2. [`thm:robust-corner-exactness`] (Robust Corner Exactness) *Let $m \in \mathbb{N}$ and $a, b, r \in \mathbb{R}$. Suppose that*

- ***(Two-block homophily.)*** *Assumption 1 holds for (m, a, b).*
- ***(Nonnegative pseudoinverse weight.)*** *$r \geq 0$.*

For the objective $F_{r,\kappa}$ of Definition 6, the following statements hold.

First, there exists a finite nonnegative robustness weight κ such that $I_n \in \mathcal{E}_m^{\text{blk}}$ and

$$F_{r,\kappa}(I_n) \leq F_{r,\kappa}(X) \qquad \text{for every } X \in \mathcal{E}_m^{\text{blk}},$$

if and only if

$$a + 3b = 2m \qquad \text{and} \qquad r = 2b(a+b).$$

Second, on this affine-balanced locus, every positive robustness weight makes I_n the unique relaxed minimizer: if

$$a + 3b = 2m \qquad \text{and} \qquad r = 2b(a+b),$$

then for every $\kappa > 0$,

$$I_n \in \mathcal{E}_m^{\text{blk}}$$

and, for every $X \in \mathcal{E}_m^{\text{blk}}$ with $X \neq I_n$,

$$F_{r,\kappa}(I_n) < F_{r,\kappa}(X).$$

Third, the iid design satisfies

$$P_{\text{iid}} \in \mathcal{P}_m^{\text{sym}}$$

for the block-exchangeable class of Definition 3, and

$$X(P_{\text{iid}}) = I_n.$$

Finally, off the affine-balanced locus, I_n is never a minimizer at any positive finite robustness weight, and relaxed minimizers converge entrywise to I_n as $\kappa \to \infty$. Precisely, if

$$\neg\big(a + 3b = 2m \text{ and } r = 2b(a+b)\big),$$

then for every $\kappa > 0$ it is not the case that

$$F_{r,\kappa}(I_n) \leq F_{r,\kappa}(X) \qquad \text{for every } X \in \mathcal{E}_m^{\text{blk}}.$$

Moreover, for any family $(X_\kappa)_{\kappa \in \mathbb{R}}$ such that, for every $\kappa > 0$,

$$X_\kappa \in \mathcal{E}_m^{\text{blk}} \quad \text{and} \quad F_{r,\kappa}(X_\kappa) \leq F_{r,\kappa}(X) \ \text{ for every } X \in \mathcal{E}_m^{\text{blk}},$$

each matrix entry converges to the corresponding entry of I_n:

$$(X_\kappa)_{ij} \longrightarrow (I_n)_{ij} \qquad \text{as } \kappa \to \infty$$

for every $i, j \in \{1, \ldots, n\}$.

Theorem 2 separates implementability from optimality. The iid law P_{iid} belongs to the block-exchangeable design class and attains I_n, but this implementable covariance is a relaxed minimizer at finite robustness only when $a+3b=2m$ and $r=2b(a+b)$. Off that affine-balanced locus, relaxed minimizers converge entrywise to I_n as $\kappa \to \infty$, yet no positive finite value of κ makes I_n itself optimal. Thus robustness can make iid assignment an asymptotic limiting target without making it an exact finite-κ solution.

Together, Theorem 1 and Theorem 2 identify two implementable covariances whose roles are sharply different. The cut covariance is exactly optimal throughout the stated strict cut region, yielding zero implementability loss there. The iid covariance is implementable everywhere under the maintained homophily conditions, with exact relaxed optimality on the stated affine-balanced locus and limiting optimality elsewhere as robustness diverges. The next section turns to parameter regions with strictly positive implementability loss.

4 Positive loss and sharp computation

The preceding section identifies two implementable covariances that can solve the relaxed problem exactly. We now study relaxed points in $\mathcal{E}_m^{\text{blk}}$ lying beyond the covariances attainable by balanced $\{\pm 1\}$ assignment laws. This is the same basic tension that appears in correlation and semidefinite relaxations more broadly: the convex benchmark can contain exposed points beyond the original discrete objects [Goemans and Williamson, 1995, Deza and Laurent, 1997]. In the present two-block design problem, the obstruction is finite and arithmetical. When m is odd, each realized block sum has odd parity, imposing a positive-mass restriction away from the pure within-block spectral direction.

The next result gives a parameter region in which this obstruction is binding. It uses two-block homophily from Assumption 1, the low-scale condition from Assumption 3, and the odd-block condition from Assumption 4. For robustness weights in the stated range, the relevant relaxed covariance is the spread covariance X_{spread}, the vertex with spectral coordinates $(m/(m-1), 0, 0)$. The theorem states that, in a nonempty interval of pseudoinverse weights, this relaxed point is uniquely optimal but not implementable.

Theorem 3. [`thm:gap-window`] (Positive-Gap Window) *Suppose Assumption 1 and Assumption 3 hold for the two-block model with block size m and weights a, b. Let*

$$q = 2(m-1), \qquad \kappa_{\text{gap}}(m,a,b) = \frac{(2m-a-3b)(a-b)\sqrt{q}}{a+b},$$

and

$$r_{\text{cut,gap}}(m,a,b,\kappa) = \max\left\{0,\ 2b(a+b)\left(1-\frac{\kappa}{a-b}\right)\right\}.$$

Define

$$R_x^-(m,a,b,\kappa) = 2b(a+b)\left(1+\frac{\kappa}{(a-b)\sqrt{q}}\right), \qquad R_x^+(m,a,b,\kappa) = (a+b)\left(2m-a-b-\frac{\kappa}{\sqrt{q}}\right),$$

and

$$r_{\text{gap}}^-(m,a,b,\kappa) = \frac{R_x^-(m,a,b,\kappa)+R_x^+(m,a,b,\kappa)}{2}, \qquad r_{\text{gap}}^+(m,a,b,\kappa) = R_x^+(m,a,b,\kappa).$$

Then the following two statements hold.

- ***(Odd block size.)*** *If Assumption 4 holds, then for every κ satisfying*
$$0 \le \kappa < \kappa_{\mathrm{gap}}(m,a,b),$$
one has
$$r_{\mathrm{cut,gap}}(m,a,b,\kappa) < r^{-}_{\mathrm{gap}}(m,a,b,\kappa) < r^{+}_{\mathrm{gap}}(m,a,b,\kappa).$$
Moreover, for every
$$r^{-}_{\mathrm{gap}}(m,a,b,\kappa) < r < r^{+}_{\mathrm{gap}}(m,a,b,\kappa),$$
the spread covariance X_{spread} belongs to $\mathcal{E}^{\mathrm{blk}}_m$, and it is the unique strict minimizer of the design objective $F_{r,\kappa}$ from Definition 6 over $\mathcal{E}^{\mathrm{blk}}_m$: for every $X \in \mathcal{E}^{\mathrm{blk}}_m$ with $X \ne X_{\mathrm{spread}}$,
$$F_{r,\kappa}(X_{\mathrm{spread}}) < F_{r,\kappa}(X).$$
At the same time,
$$X_{\mathrm{spread}} \notin \mathcal{C}^{\pm}_m,$$
where $\mathcal{C}^{\pm}_m$ is the implementable covariance class of Definition 5, and the implementability loss of Definition 7 satisfies
$$\Delta^{\pm}_m(r,\kappa) > 0.$$

- ***(Even block size.)*** *If m is even, then every covariance in the block elliptope is implementable:*
$$\mathcal{E}^{\mathrm{blk}}_m \subseteq \mathcal{C}^{\pm}_m.$$

Theorem 3 is the paper's positive-loss statement. Under its two-block homophily, low-scale, and odd-community hypotheses, the upper bound $\kappa_{\mathrm{gap}}(m,a,b)$ defines a robustness range in which the interval $(r^{-}_{\mathrm{gap}}, r^{+}_{\mathrm{gap}})$ is nonempty and lies strictly beyond the cut-side comparison threshold $r_{\mathrm{cut,gap}}$. Inside that interval, the relaxed optimum is not merely one unattainable minimizer among many; it is the unique strict minimizer over $\mathcal{E}^{\mathrm{blk}}_m$. Since $X_{\mathrm{spread}} \notin \mathcal{C}^{\pm}_m$, imposing implementability necessarily raises the objective value, giving $\Delta^{\pm}_m(r,\kappa) > 0$.

The even-m clause marks the opposite parity case. When each block can have zero sign sum, every relaxed covariance in the block elliptope slice is generated by some block-exchangeable sign law. Thus Theorem 3 ties positive loss specifically to the odd-block parity restriction.

The interval result is deliberately local. It identifies one open region where the spread covariance is the relaxed optimum and the loss is strictly positive, but it does not assert a closed-form global boundary for exact attainability outside that region. For computation at arbitrary r and κ, the relevant object is instead the sharp finite active-set formula below.

The formula works in the reduced spectral simplex T_m, obtained from Lemma 1: the block symmetry diagonalizes every $X(u,v)$ into three eigendirections, with coordinates $x = 1-u$ on the within-block contrast space (multiplicity $q = 2(m-1)$), $y = 1+(m-1)u-mv$ on the between-block contrast direction, and $z_{\mathrm{sp}} = 1+(m-1)u+mv$ on the all-ones (aggregate) direction; the subscript "sp" marks this last, spread-across-all-units coordinate and distinguishes it from an ambient sign index. In these coordinates $F_{r,\kappa}$ becomes the scalar objective $\phi_{r,\kappa}(x,y,z_{\mathrm{sp}}) = c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sqrt{qx^2+y^2+z_{\mathrm{sp}}^2}$, and implementability adds only the parity truncation $d_m \le y + z_{\mathrm{sp}}$, with $d_m = 0$ for even block sizes and $d_m = 2/m$ for odd block sizes. The theorem's remaining ingredients are the following finite objects, established in Section C (Lemma 6 and Lemma 7). For linear coefficients $\alpha^\sharp = (c_x/q, c_y, c_z)$ and norm weights $\beta^\sharp = (1/q, 1, 1)$, the weighted-simplex objective is $\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}(t) = \sum_i \alpha_i^\sharp t_i + \kappa\sqrt{\sum_i \beta_i^\sharp t_i^2}$ on the simplex of mass $2m$; an *admissible* active support–multiplier pair (S, λ) is a nonempty $S \subseteq \{x, y, z_{\rm sp}\}$ and $\lambda \in \mathbb{R}$ with $\sum_{i\in S}(\lambda - \alpha_i^\sharp)^2/\beta_i^\sharp = \kappa^2$, $\alpha_i^\sharp < \lambda$ on S, and $\lambda \le \alpha_j^\sharp$ off S; the associated *active-set point* $t^{S,\lambda}$ puts mass $2m\,((\lambda - \alpha_i^\sharp)/\beta_i^\sharp)/\sum_{h\in S}(\lambda - \alpha_h^\sharp)/\beta_h^\sharp$ on each $i \in S$ and zero off S; truncSelector returns the endpoint or interior optimizer of Φ restricted to the truncation segment $\{y + z_{\rm sp} = d_m\}$; and $\rho_\star(m, a, b, r, \kappa)$ is the resulting constrained-minus-unconstrained reduced value, i.e. the minimum of $\phi_{r,\kappa}$ over $T_m \cap \{y + z_{\rm sp} \ge d_m\}$ minus its minimum over T_m. The theorem shows this $\rho_\star$ equals the implementability loss and gives its finite case split.

Theorem 4. [`thm:sharp-rho-star`] (Sharp Rounding Loss) *Under Assumption 1, fix $r, \kappa \in [0, \infty)$. Let*

$$q = 2(m-1), \qquad c_x = q\Big((a+b) + \frac{r}{a+b}\Big), \qquad c_y = 2b + \frac{r}{2b}, \qquad c_z = 2m,$$

let

$$T_m = \{(x, y, z_{\rm sp}) : x, y, z_{\rm sp} \ge 0,\ qx + y + z_{\rm sp} = 2m\},$$

let

$$\phi_{r,\kappa}(x, y, z_{\rm sp}) = c_x x + c_y y + c_z z_{\rm sp} + \kappa\sqrt{qx^2 + y^2 + z_{\rm sp}^2},$$

and let $d_m = 0$ if m is even and $d_m = 2/m$ if m is odd. Then:

- ***(Sharp equality and range.)***

$$\rho_\star(m, a, b, r, \kappa) = \Delta_m^\pm(r, \kappa), \qquad \rho_\star(m, a, b, r, \kappa) \ge 0, \qquad \Delta_m^\pm(r, \kappa) \ge 0,$$

  where $\Delta_m^\pm$ is the implementability gap of Definition 7.

- ***(Zero loss.)***

$$\rho_\star(m, a, b, r, \kappa) = 0$$

  if and only if there exists $(x, y, z_{\rm sp}) \in T_m$ such that

$$\phi_{r,\kappa}(x, y, z_{\rm sp}) \le \phi_{r,\kappa}(x', y', z'_{\rm sp}) \quad \textit{for every } (x', y', z'_{\rm sp}) \in T_m,$$

  and

$$d_m \le y + z_{\rm sp}.$$

- ***(Strict robustness.)*** *If $\kappa > 0$, then there is a unique relaxed minimizer of $\phi_{r,\kappa}$ over T_m. Moreover,*

$$\rho_\star(m, a, b, r, \kappa) = 0$$

  if and only if every relaxed minimizer $(x, y, z_{\rm sp}) \in T_m$ satisfies

$$d_m \le y + z_{\rm sp}.$$

- ***(Linear case.)*** *If $\kappa = 0$, then for every $(x, y, z_{\rm sp}) \in T_m$, the point $(x, y, z_{\rm sp})$ minimizes $\phi_{r,0}$ over T_m if and only if*

$$x \ne 0 \Rightarrow \frac{c_x}{q} \le c_y \textit{ and } \frac{c_x}{q} \le c_z,$$

$$y \neq 0 \Rightarrow c_y \leq \frac{c_x}{q} \text{ and } c_y \leq c_z,$$

and

$$z_{\mathrm{sp}} \neq 0 \Rightarrow c_z \leq \frac{c_x}{q} \text{ and } c_z \leq c_y.$$

- ***(Active-set formula.)*** *If $\kappa > 0$, then there exist an active support $S \subseteq \{x, y, z_{\mathrm{sp}}\}$ and a multiplier $\lambda \in \mathbb{R}$ such that S and λ are admissible for*

$$\alpha^\sharp = \Big(\frac{c_x}{q}, c_y, c_z\Big), \qquad \beta^\sharp = \Big(\frac{1}{q}, 1, 1\Big), \qquad \kappa.$$

Writing $t^{S,\lambda} = (t_x^{S,\lambda}, t_y^{S,\lambda}, t_z^{S,\lambda})$ for the associated active-set point on the simplex of mass $2m$, the reduced point

$$\Big(\frac{t_x^{S,\lambda}}{q}, t_y^{S,\lambda}, t_z^{S,\lambda}\Big)$$

belongs to T_m, minimizes $\phi_{r,\kappa}$ over T_m, and the relaxed reduced value is

$$2m\,\lambda.$$

- ***(Truncation correction.)*** *If no relaxed minimizer $(x, y, z_{\mathrm{sp}}) \in T_m$ satisfies $d_m \leq y + z_{\mathrm{sp}}$, then the implementable reduced value equals the weighted-simplex objective with coefficients $\alpha^\sharp$, weights $\beta^\sharp$, and robustness weight κ, evaluated at the point on the truncation segment*

$$H_{d_m} = \{(2m - d_m, s, d_m - s) : 0 \leq s \leq d_m\}$$

selected by $\mathrm{truncSelector}(2m, d_m, \alpha^\sharp, \beta^\sharp, \kappa)$*:*

$$\mathrm{val}_{\mathrm{impl}}(m, a, b, r, \kappa) = \Phi_{\alpha^\sharp,\beta^\sharp,\kappa}(2m - d_m,\ s_\star,\ d_m - s_\star)\,,$$

where

$$s_\star = \mathrm{truncSelector}(2m, d_m, \alpha^\sharp, \beta^\sharp, \kappa).$$

- ***(Even block size.)*** *If m is even, then*

$$\rho_\star(m, a, b, r, \kappa) = 0.$$

Theorem 4 turns the implementability question into a finite calculation. The sharp loss $\rho_\star(m, a, b, r, \kappa)$ is exactly the implementability loss $\Delta_m^\pm(r, \kappa)$, so the formula is not an approximation to the rounding penalty. Zero loss occurs precisely when at least one relaxed minimizer satisfies the parity-induced lower bound $d_m \leq y + z_{\mathrm{sp}}$. When $\kappa > 0$, the active-set result yields a unique relaxed minimizer, and the zero-loss condition can be checked at that point alone.

The active-set part supplies the promised finite case split. The coefficients $\alpha^\sharp = (c_x/q, c_y, c_z)$ and weights $\beta^\sharp = (1/q, 1, 1)$ convert the reduced objective into a weighted simplex problem over total mass $2m$. For positive robustness, an admissible support S and multiplier λ identify the relaxed minimizer and its value. If that minimizer violates the implementability truncation, the remaining optimization is one-dimensional: the implementable value is obtained on the segment H_{d_m} by the selector appearing in Theorem 4. The theorem thereby gives a complete finite procedure for computing $\Delta_m^\pm(r, \kappa)$ for any admissible m, a, b, r, κ.

The two main results in this section therefore play complementary roles. Theorem 3 provides an interpretable open region in which the relaxation has a strictly unattainable optimum for odd block sizes. Theorem 4 gives the exact loss everywhere under Assumption 1, including the linear case $\kappa = 0$, the unique-minimizer case $\kappa > 0$, and the even-block case where the loss vanishes. For interference-aware randomized design, the practical implication is a two-step calculation: solve the covariance relaxation, then check its optimizer against the finite sign-assignment restrictions and apply the exact active-set adjustment when needed.

A worked example and parameter map. Take $m = 3$, $a = 2$, $b = 1$, and $\kappa = 0$; these satisfy homophily ($a > b > 0$) and the low-scale regime ($a + 3b = 5 < 2m = 6$). Here $q = 4$, the parity truncation is $d_3 = 2/3$, and the reduced coefficients are $c_x = q\big((a + b) + r/(a + b)\big) = 4\big(3 + r/3\big)$, $c_y = 2b + r/(2b) = 2 + r/2$, and $c_z = 2m = 6$. The gap-window robustness bound is $\kappa_{\text{gap}} = (2m - a - 3b)(a - b)\sqrt{q}/(a + b) = 2/3$, so $\kappa = 0$ lies in the admissible range. Mapping over $r \geq 0$ then exhibits all three regimes of the theory. For $r < r_{\text{cut}} = \min\{2b(a + b), 2b(2m - 2b)\} = \min\{6, 8\} = 6$, the cut covariance $X_{\text{cut}} = (0, 6, 0)$ is the unique relaxed optimum and is implementable, so $\Delta_3^{\pm} = 0$ (Theorem 1). For r in the gap window $(r^-_{\text{gap}}, r^+_{\text{gap}}) = (7.5, 9)$, the spread covariance $X_{\text{spread}} = (3/2, 0, 0)$ is the unique relaxed optimum but is *not* implementable (Theorem 3). At $r = 8$, for instance, the unconstrained reduced optimum is at X_{spread} with value $c_x \cdot \frac{3}{2} = 34$, while the implementable optimum lies on the parity slice $y + z_{\text{sp}} = 2/3$ at $(x, y, z_{\text{sp}}) = (4/3,\ 2/3,\ 0)$ with value $308/9$; hence

$$\Delta_3^{\pm}(8, 0) = \tfrac{308}{9} - 34 = \tfrac{2}{9} \approx 0.222,$$

a strictly positive but, in this instance, modest price relative to the optimal value 34 (about 0.65%). For even community size the picture collapses: $d_m = 0$, the parity truncation is vacuous, and $\Delta_m^{\pm} \equiv 0$ throughout. This single example only illustrates the three regimes for one parameter choice; it does not establish that the loss is generally small or generally negligible, and Theorem 4 should be used to compute $\Delta_m^{\pm}$ at any specific (m, a, b, r, κ) of interest. How large the loss can become as a, b, r, κ scale is not settled here and is left open.

Table 1 summarizes the regimes and the result that governs each.

Regime	**Relaxed optimum**	**Implementable?**	**Loss $\Delta_m^{\pm}$**
Cut ($r < r_{\text{cut}}$)	$X_{\text{cut}} = (0, 2m, 0)$	yes (via P_{cut})	0 (Thm Theorem 1)
Robust locus ($a + 3b = 2m$, $r = 2b(a + b)$)	$I_n = (1, 1, 1)$	yes (via P_{iid})	0 (Thm Theorem 2)
Odd-parity window (m odd, low-scale)	$X_{\text{spread}} = (\frac{m}{m-1}, 0, 0)$	no	> 0 (Thm Theorem 3)
Even m	—	every $X \in \mathcal{E}_m^{\text{blk}}$	0 (Thm Theorem 4)
General (m, a, b, r, κ)	active-set point	iff parity slice met	$\rho_\star$ (Thm Theorem 4)

Table 1: Regimes of the two-block implementability problem. Coordinates are the reduced spectral coordinates (x, y, z_{sp}); "low-scale" abbreviates $a + 3b < 2m$ with the robustness range of Theorem 3.

5 Discussion and extensions

Relation to covariance-based network design. The design criterion and elliptope relaxation build on covariance-based experimental design for network interference, particularly the semidefinite formulation of Thiyageswaran et al. [2026]; Boyd and Vandenberghe [2004] supplies convex-optimization background, and Goemans and Williamson [1995], Deza and Laurent [1997] provide canonical elliptope and cut-polytope context. This paper contributes an exact implementability theory for the homophilous two-block family: a joint symmetry reduction of relaxed covariances and implementable laws, an odd-parity obstruction producing positive loss, and a finite active-set correction computing $\Delta_m^{\pm}$ exactly.

Viewed abstractly, $\mathcal{C}_m^{\pm}$ is a moment set for block-exchangeable Rademacher vectors and a two-coordinate projection of the correlation polytope. The paper pins down the exact boundary of this low-dimensional exchangeable slice through the odd-block parity truncation $y + z_{\mathrm{sp}} \geq 2/m$, then converts that boundary into the exact rounding loss. Future work can connect this family-specific characterization more fully to general correlation-polytope and exchangeable-moment results.

The results above give a finite-sample account of when a covariance relaxation can be read literally as a randomized design. In the two-block homophily model, the relaxed block elliptope $\mathcal{E}_m^{\mathrm{blk}}$ serves both as a computational device and as a benchmark for exact implementability by a sign-symmetric, block-exchangeable $\{\pm 1\}$ assignment law. Theorem 1 gives a region in which the relaxed optimum is generated by an assignment law, whereas Theorem 3 gives a region with a uniquely optimal relaxed covariance and strictly positive implementability loss.

For experimental design under interference, the covariance matrix is a property of the assignment mechanism, and implementation requires a covariance induced by a randomization law. The implementability loss $\Delta_m^{\pm}(r, \kappa)$ therefore has a direct design interpretation: it is the exact objective-value price of returning from the relaxed covariance problem to feasible randomized assignment. Theorem 4 computes that price through a finite active-set calculation.

The iid result points in the opposite direction. Independent fair assignment is always available as a randomized design and has a familiar role in design-based inference. Theorem 2 identifies exact finite-robustness optimality on the affine-balanced locus $a+3b = 2m$ and $r = 2b(a+b)$, while relaxed minimizers elsewhere converge to I_n as κ grows. Thus robustness rationalizes iid assignment either as an exact solution on the locus or as a limiting target elsewhere.

The positive-loss result is arithmetic. For odd m, each realized block sum has absolute value at least one. Under the hypotheses and robustness range stated in Theorem 3, this parity restriction binds and places the uniquely optimal spread covariance outside $\mathcal{C}_m^{\pm}$. For even m, every relaxed block covariance is implementable in the maintained symmetric model. The theorem therefore isolates odd-block parity as the source of positive loss.

5.1 Limitations and future work

The two-block structure is deliberately stylized, but it captures a common empirical feature of networked populations: ties and exposures are often stronger within social groups than between them. Homophily and community structure are central themes in social network analysis [McPherson et al., 2001, Newman, 2002], and stochastic block models provide a canonical probabilistic language for such clustered networks [Holland et al., 1983, Abbe, 2018]. Field and policy experiments with interference often face analogous clustering through villages, schools, households, online neighborhoods, or other exposure groups [Banerjee et al., 2013, Baird et al., 2018, Ugander and Yin, 2020, Li et al., 2021]. The present model abstracts from those empirical details in order to isolate an exact design question: after symmetry reduction, which relaxed covariance matrices are generated by actual sign assignments?

Several extensions follow naturally, but they require additional arguments beyond the statements in this paper. The present results do not cover unequal block sizes, where both the spectral reduction and the block-sum lattice would require new arguments. The present results also do not cover more than two communities; the three-coordinate reduction (x, y, z_{sp}) and the implementability characterization would require modification. In those settings, one should not expect the parity correction d_m or the active-set formula in Theorem 4 to carry over without modification.

A second extension concerns additional restrictions on the randomization scheme. The class $\mathcal{P}_m^{\mathrm{sym}}$ imposes sign symmetry and block exchangeability, but it does not require a fixed treated count in every realization, nor does it impose covariate balance, rerandomization acceptance rules, or cluster-level assignment constraints. Such restrictions are common in empirical practice and can be important for finite-sample precision and credibility [Sävje, 2021, Zigler and Papadogeorgou, 2018]. They would shrink the implementable covariance class and could introduce losses even in cases where $\mathcal{E}_m^{\mathrm{blk}} \subseteq \mathcal{C}_m^{\pm}$ under the present assumptions.

A third issue is that the graph itself is treated as fixed. In applications, measured networks may be incomplete, exposure weights may be estimated, and the relevant interference neighborhood may be uncertain. The objective $F_{r,\kappa}$ can be interpreted conditionally on the chosen graph weights, but the results here do not address robustness to network measurement error or to alternative exposure mappings. Those questions are important for empirical deployment because different plausible graphs can imply different design priorities, especially when homophily is strong.

Finally, the analysis is finite-sample throughout. The exact formulas are indexed by m, a, b, r, κ, and the odd-block obstruction is stated at the finite block size used by the design. Although the parity truncation becomes numerically smaller as m grows, the consequences for objective values depend on the full set of coefficients in Theorem 4. Any asymptotic simplification should therefore be derived from the finite active-set expression rather than assumed from the geometry alone.

Appendices

A Reduced geometry and proof ingredients

This appendix makes visible the auxiliary geometry behind the two main optimization statements: the reduced coordinates, vertex certificates, parity restriction, and finite active-set calculation used in Theorem 3 and Theorem 4. The reduction is a low-dimensional convex-analytic one: the block symmetry turns the covariance problem into a weighted simplex problem, in the same general spirit as convex optimization [Boyd and Vandenberghe, 2004] and elliptope relaxations for discrete optimization [Goemans and Williamson, 1995, Deza and Laurent, 1997].

The first ingredient identifies the spectral coordinates in which the block elliptope becomes a simplex and the objective becomes a scalar function.

Lemma 1. [`lem:block-spectral-coordinates`] (Block Spectral Coordinates) *For any real numbers r, κ, u, v, assume:*

- ***(Homophily.)*** *Assumption 1 holds for the block size m and weights a, b.*
- ***(Spectral coordinates.)*** *Set*

$$q = 2(m-1), \qquad x = 1 - u, \qquad y = 1 + (m-1)u - mv, \qquad z_{\mathrm{sp}} = 1 + (m-1)u + mv.$$

- ***(Reduced coefficients.)*** *Set*

$$c_x = q\left((a+b) + \frac{r}{a+b}\right), \qquad c_y = 2b + \frac{r}{2b}, \qquad c_z = 2m.$$

Then the block covariance matrix $X(u,v)$ belongs to the block elliptope $\mathcal{E}_m^{\mathrm{blk}}$ if and only if

$$x \geq 0, \qquad y \geq 0, \qquad z_{\mathrm{sp}} \geq 0, \qquad qx + y + z_{\mathrm{sp}} = 2m.$$

Moreover, for the design objective of Definition 6,

$$F_{r,\kappa}\big(X(u,v)\big) = c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}.$$

Finally,

$$X_{\mathrm{cut}} = X(1,-1) \quad \textit{and} \quad (1-1,\ 1+(m-1)\cdot 1 - m(-1),\ 1+(m-1)\cdot 1 + m(-1)) = (0,2m,0),$$

while

$$I_n = X(0,0) \quad \textit{and} \quad (1-0,\ 1+(m-1)\cdot 0 - m\cdot 0,\ 1+(m-1)\cdot 0 + m\cdot 0) = (1,1,1).$$

Proof of Lemma 1. Throughout write $n = 2m$, $q = 2(m-1)$, and

$$x = 1-u, \qquad y = 1+(m-1)u - mv, \qquad z_{\mathrm{sp}} = 1+(m-1)u+mv.$$

Assumption 1 gives $m \geq 2$, hence $q > 0$ and $m > 0$, and it gives $a > b > 0$, hence $a+b > 0$ and $2b > 0$; these are the only nonvanishing denominators used below.

Step 1 (the two parameters are read off from the matrix). Since $|A_m| = m \geq 2$, we may pick two distinct indices $i,j \in A_m$, and since $|B_m| = m \geq 1$ we may pick $k \in B_m$. By the entry description of $X(u,v)$,

$$X(u,v)_{ij} = u, \qquad X(u,v)_{ik} = v.$$

Consequently $X(u,v) = X(u',v')$ forces $u' = u$ and $v' = v$: the block-symmetric parameterization is injective. Therefore the only representation of $X(u,v)$ as a member of $\mathcal{E}_m^{\mathrm{blk}}$ in the sense of Definition 4 is the one with parameters (u,v), and

$$X(u,v) \in \mathcal{E}_m^{\mathrm{blk}} \iff 1-u \geq 0, \quad 1+(m-1)u - mv \geq 0, \quad 1+(m-1)u+mv \geq 0,$$

that is, $x \geq 0$, $y \geq 0$, $z_{\mathrm{sp}} \geq 0$. Moreover the affine identity

$$qx + y + z_{\mathrm{sp}} = 2(m-1)(1-u) + \big(1+(m-1)u - mv\big) + \big(1+(m-1)u+mv\big) = 2m$$

holds for all u,v, the u-terms and the v-terms cancelling. Hence block-elliptope membership is equivalent to $(x,y,z_{\mathrm{sp}}) \in T_m$, which is the first assertion.

Step 2 (two pair-counting identities). Call a matrix $w \in \mathbb{R}^{n\times n}$ *block-patterned with data* (d,σ,c) if $w_{ii} = d$ for every i, $w_{ij} = \sigma$ whenever $i \neq j$ lie in the same community, and $w_{ij} = c$ whenever i and j lie in opposite communities. Among the n^2 ordered index pairs there are $n = 2m$ diagonal pairs, $2m(m-1)$ same-community off-diagonal pairs, and $2m\cdot m$ opposite-community pairs, so

$$\sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij} = 2m\,d + 2m(m-1)\,\sigma + 2m^2\,c.$$

With the community sign vector $s \in \mathbb{R}^n$ defined by $s_i = 1$ for $i \in A_m$ and $s_i = -1$ for $i \in B_m$, the sign product $s_i s_j$ equals $+1$ on diagonal and same-community pairs and -1 on opposite-community pairs, so the same count with the sign weight gives

$$\sum_{i=1}^{n}\sum_{j=1}^{n} s_i s_j\, w_{ij} = 2m\, d + 2m(m-1)\,\sigma - 2m^2\, c.$$

Step 3 (the four terms of the objective). The matrix $X(u,v)$ is block-patterned with data $(1,u,v)$. Applying the unsigned count with $(d,\sigma,c) = (1,u,v)$ gives the balance term

$$\operatorname{tr}\bigl(J_n X(u,v)\bigr) = \sum_{i,j} X(u,v)_{ij} = 2m + 2m(m-1)u + 2m^2 v = 2m\bigl(1 + (m-1)u + mv\bigr) = 2m\, z_{\mathrm{sp}}.$$

The entrywise square of $X(u,v)$ is block-patterned with data $(1,u^2,v^2)$, so the same count gives $\|X(u,v)\|_{S_2}^2 = 2m + 2m(m-1)u^2 + 2m^2v^2$; on the other hand, expanding the three spectral coordinates,

$$qx^2 + y^2 + z_{\mathrm{sp}}^2 = 2(m-1)(1-u)^2 + 2\bigl(1 + (m-1)u\bigr)^2 + 2m^2v^2 = 2m + 2m(m-1)u^2 + 2m^2v^2,$$

the linear-in-u terms $-4(m-1)u$ and $+4(m-1)u$ cancelling. Hence

$$\|X(u,v)\|_{S_2} = \sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}.$$

Next, $\operatorname{tr} X(u,v) = n = 2m$ because the diagonal is constant equal to one. Let

$$\Pi_1 = \frac{1}{n} J_n, \qquad \Pi_s = \frac{1}{n}\, ss^\top$$

be the orthogonal projections onto the all-ones direction and onto the community-sign direction. Dividing the two displays of Step 2 (applied with $(d,\sigma,c) = (1,u,v)$) by $n = 2m$ gives

$$\operatorname{tr}\bigl(\Pi_1 X(u,v)\bigr) = z_{\mathrm{sp}}, \qquad \operatorname{tr}\bigl(\Pi_s X(u,v)\bigr) = \frac{2m + 2m(m-1)u - 2m^2 v}{2m} = y.$$

The pseudoinverse $L_m^\dagger$ enters the objective only through its spectral action, which on the two-block graph is

$$L_m^\dagger = \frac{1}{a+b}\bigl(I_n - \Pi_1 - \Pi_s\bigr) + \frac{1}{2b}\,\Pi_s,$$

i.e. eigenvalue $1/(a+b)$ on the within-community contrast space, $1/(2b)$ on the community-sign direction, and 0 on the all-ones direction. Therefore, using $\operatorname{tr} X(u,v) = 2m$ and the affine identity $2m - y - z_{\mathrm{sp}} = qx$ of Step 1,

$$\operatorname{tr}\bigl(L_m^\dagger X(u,v)\bigr) = \frac{2m - z_{\mathrm{sp}} - y}{a+b} + \frac{y}{2b} = \frac{qx}{a+b} + \frac{y}{2b}.$$

For the graph term, every vertex has the same weighted degree

$$\sum_j W_{ij} = (m-1)\frac{a}{m} + m\,\frac{b}{m} = \frac{(m-1)a}{m} + b,$$

because each vertex has $m-1$ same-community neighbours of weight a/m and m opposite-community neighbours of weight b/m. Thus L_m is block-patterned with data $\big((m-1)a/m+b,\ -a/m,\ -b/m\big)$, and the matrix with entries $(L_m)_{ij}X(u,v)_{ij}$ is block-patterned with data $\big((m-1)a/m+b,\ -(a/m)u,\ -(b/m)v\big)$. Since $X(u,v)$ is symmetric, the unsigned count gives

$$\operatorname{tr}\big(L_m X(u,v)\big) = 2m\Big(\frac{(m-1)a}{m}+b\Big) - 2m(m-1)\frac{a}{m}u - 2m^2\frac{b}{m}v = 2(m-1)a(1-u)+2mb(1-v),$$

and expanding both sides shows that this equals the reduced form

$$2(m-1)a(1-u)+2mb(1-v) = q(a+b)x+2b\,y.$$

Step 4 (assembling the objective). Substituting the four displays of Step 3 into Definition 6,

$$\begin{aligned}
F_{r,\kappa}\big(X(u,v)\big) &= \big(q(a+b)x+2b\,y\big) + r\Big(\frac{qx}{a+b}+\frac{y}{2b}\Big) + \kappa\sqrt{qx^2+y^2+z_{\mathrm{sp}}^2} + 2m\,z_{\mathrm{sp}} \\
&= q\Big((a+b)+\frac{r}{a+b}\Big)x + \Big(2b+\frac{r}{2b}\Big)y + 2m\,z_{\mathrm{sp}} + \kappa\sqrt{qx^2+y^2+z_{\mathrm{sp}}^2} \\
&= c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sqrt{qx^2+y^2+z_{\mathrm{sp}}^2},
\end{aligned}$$

which is the asserted reduced form.

Step 5 (the two corner readings). The cut covariance is $X_{\mathrm{cut}} = ss^\top$, whose entries are $s_i^2=1$ on the diagonal, $+1$ on same-community pairs, and -1 on opposite-community pairs; hence $X_{\mathrm{cut}} = X(1,-1)$, and substituting $u=1$, $v=-1$,

$$\big(1-1,\ 1+(m-1)-m(-1),\ 1+(m-1)+m(-1)\big) = (0,2m,0).$$

Likewise I_n has entries 1 on the diagonal and 0 off it, so $I_n = X(0,0)$, and substituting $u=v=0$,

$$\big(1-0,\ 1+(m-1)\cdot 0 - m\cdot 0,\ 1+(m-1)\cdot 0+m\cdot 0\big) = (1,1,1).$$

□

Lemma 1 is the coordinate change used throughout the main text. It replaces the two covariance parameters u, v by nonnegative spectral coordinates satisfying one affine constraint. In these coordinates, the cut covariance is the vertex $(0,2m,0)$, the iid covariance is the center-like point $(1,1,1)$, and the objective separates into a linear part plus a weighted Euclidean norm.

The next two statements are vertex certificates. They give simple coefficient inequalities under which either the cut vertex or the spread vertex is the unique minimizer of the reduced problem.

Lemma 2. [`lem:cut-vertex-certificate`] (Cut Vertex Certificate) *Let m be fixed and write $q=2(m-1)$. Suppose that*

- ***(Positive spectral multiplicity.)*** *$q>0$.*
- ***(Nonnegative robustness.)*** *$\kappa \ge 0$.*
- ***(Cut-side coefficient separation.)*** *The coefficients satisfy*

$$\frac{c_x}{q} > c_y+\kappa \qquad \text{and} \qquad c_z > c_y+\kappa.$$

Then $(0, 2m, 0) \in T_m$*, where*

$$T_m = \{(x, y, z_{\text{sp}}) : x \geq 0,\ \ y \geq 0,\ \ z_{\text{sp}} \geq 0,\ \ qx + y + z_{\text{sp}} = 2m\}.$$

Moreover, for every $(x, y, z_{\text{sp}}) \in T_m$ *with* $(x, y, z_{\text{sp}}) \neq (0, 2m, 0)$*,*

$$c_x \cdot 0 + c_y(2m) + c_z \cdot 0 + \kappa\sqrt{q \cdot 0^2 + (2m)^2 + 0^2} < c_x x + c_y y + c_z z_{\text{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\text{sp}}^2}.$$

Thus $(0, 2m, 0)$ *is the unique minimizer on* T_m *of the reduced objective.*

Proof of Lemma 2. Write $q = 2(m-1)$ and $M = 2m$, and abbreviate the reduced objective by

$$\phi(x, y, z_{\text{sp}}) = c_x x + c_y y + c_z z_{\text{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\text{sp}}^2},$$

so that the claim is that $(0, M, 0)$ is the unique minimizer of ϕ on T_m. The hypothesis $q > 0$ means $m > 1$, so $M = 2m > 0$.

Step 1 (feasibility). The point $(0, M, 0)$ has nonnegative coordinates and satisfies $q \cdot 0 + M + 0 = M = 2m$; hence $(0, M, 0) \in T_m$.

Step 2 (the value at the cut vertex). Since $M > 0$, $\sqrt{q \cdot 0^2 + M^2 + 0^2} = \sqrt{M^2} = M$, so

$$\phi(0, M, 0) = c_y M + \kappa M = (c_y + \kappa)M.$$

Step 3 (a lower bound for the norm term). Let $(x, y, z_{\text{sp}}) \in T_m$. Because $q > 0$, both $qx^2 \geq 0$ and $z_{\text{sp}}^2 \geq 0$, so

$$y^2 \leq qx^2 + y^2 + z_{\text{sp}}^2, \qquad \text{hence} \qquad y \leq \sqrt{qx^2 + y^2 + z_{\text{sp}}^2},$$

and multiplying by $\kappa \geq 0$ gives $\kappa y \leq \kappa\sqrt{qx^2 + y^2 + z_{\text{sp}}^2}$. Therefore

$$c_x x + c_y y + c_z z_{\text{sp}} + \kappa y \leq \phi(x, y, z_{\text{sp}}).$$

Step 4 (the coefficient separation in cleared form). Multiplying the hypothesis $c_x/q > c_y + \kappa$ by $q > 0$, and taking the hypothesis $c_z > c_y + \kappa$ as it stands, gives

$$c_x - qc_y - q\kappa > 0, \qquad c_z - c_y - \kappa > 0.$$

Step 5 (some off-y mass is present). Suppose now in addition $(x, y, z_{\text{sp}}) \neq (0, M, 0)$. If $x = 0$ and $z_{\text{sp}} = 0$, the simplex equation $qx + y + z_{\text{sp}} = M$ would force $y = M$, i.e. the point would be $(0, M, 0)$; hence $x > 0$ or $z_{\text{sp}} > 0$. Since $x \geq 0$ and $z_{\text{sp}} \geq 0$ and both coefficients in Step 4 are strictly positive, the sum of the two products has a strictly positive term and a nonnegative one:

$$0 < (c_x - qc_y - q\kappa)x + (c_z - c_y - \kappa)z_{\text{sp}}.$$

Step 6 (the difference identity). Substituting $M = qx + y + z_{\text{sp}}$ into $(c_y + \kappa)M$ and cancelling the y-terms gives the identity

$$\big(c_x x + c_y y + c_z z_{\text{sp}} + \kappa y\big) - (c_y + \kappa)M = (c_x - qc_y - q\kappa)x + (c_z - c_y - \kappa)z_{\text{sp}}.$$

Combining it with Step 5,

$$(c_y + \kappa)M < c_x x + c_y y + c_z z_{\text{sp}} + \kappa y.$$

Step 7 (conclusion). Chaining Steps 2, 6 and 3,

$$\phi(0, M, 0) = (c_y + \kappa)M < c_x x + c_y y + c_z z_{\text{sp}} + \kappa y \le \phi(x, y, z_{\text{sp}}),$$

which is the asserted strict inequality

$$c_x \cdot 0 + c_y(2m) + c_z \cdot 0 + \kappa\sqrt{q \cdot 0^2 + (2m)^2 + 0^2} < c_x x + c_y y + c_z z_{\text{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\text{sp}}^2}$$

for every $(x, y, z_{\text{sp}}) \in T_m$ other than $(0, 2m, 0)$. In particular no other point of T_m attains the value $\phi(0, 2m, 0)$, so the cut vertex is the unique minimizer. □

Lemma 2 is the reduced certificate underlying the cut exactness result. The coefficient separation says that moving mass away from the y-coordinate toward either the within-block spectral coordinate or the aggregate coordinate is too costly, even after accounting for the robustness term.

Lemma 3. [`lem:spread-vertex-certificate`] (Spread Vertex Certificate) *Let $m \in \mathbb{N}$, let $c_x, c_y, c_z, \kappa \in \mathbb{R}$, and set $q = 2(m-1)$. Suppose that*

- ***(Positive multiplicity.)*** $q > 0$.
- ***(Nonnegative robustness.)*** $\kappa \ge 0$.
- ***(The y-coefficient is large.)*** $c_y > c_x/q + \kappa/\sqrt{q}$.
- ***(The z_{sp}-coefficient is large.)*** $c_z > c_x/q + \kappa/\sqrt{q}$.

Define

$$T_m = \{(x, y, z_{\text{sp}}) : x \ge 0,\ y \ge 0,\ z_{\text{sp}} \ge 0,\ qx + y + z_{\text{sp}} = 2m\}$$

and

$$\phi(x, y, z_{\text{sp}}) = c_x x + c_y y + c_z z_{\text{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\text{sp}}^2}.$$

Then $(2m/q, 0, 0) \in T_m$, and for every $(x, y, z_{\text{sp}}) \in T_m$ with $(x, y, z_{\text{sp}}) \ne (2m/q, 0, 0)$,

$$\phi(2m/q, 0, 0) < \phi(x, y, z_{\text{sp}}).$$

Proof of Lemma 3. Put $q = 2(m-1)$, $M = 2m$ and $\sigma = \sqrt{q}$. Since $q > 0$ we have $\sigma > 0$, $\sigma^2 = q$, and $M = q + 2 > 0$.

Step 1 (feasibility of the spread vertex). The point $(M/q, 0, 0)$ has nonnegative coordinates and satisfies $q(M/q) + 0 + 0 = M = 2m$, so $(M/q, 0, 0) \in T_m$.

Step 2 (the value at the spread vertex). Since $M/q \ge 0$,

$$\sqrt{q(M/q)^2 + 0^2 + 0^2} = \sqrt{\big(\sigma\,(M/q)\big)^2} = \sigma\,\frac{M}{q},$$

hence

$$\phi\Big(\frac{M}{q}, 0, 0\Big) = c_x\frac{M}{q} + \kappa\sigma\frac{M}{q}.$$

Step 3 (a lower bound for the norm term). Let $(x, y, z_{\text{sp}}) \in T_m$. From $\sigma^2 = q$ and $y^2, z_{\text{sp}}^2 \ge 0$,

$$(\sigma x)^2 = qx^2 \le qx^2 + y^2 + z_{\text{sp}}^2, \qquad \text{hence} \qquad \sigma x \le \sqrt{qx^2 + y^2 + z_{\text{sp}}^2},$$

and multiplying by $\kappa \geq 0$ gives $\kappa\sigma x \leq \kappa\sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}$. Therefore

$$c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sigma x \leq \phi(x, y, z_{\mathrm{sp}}).$$

Step 4 (the coefficient separation in cleared form). Multiplying the hypotheses $c_y > c_x/q + \kappa/\sigma$ and $c_z > c_x/q + \kappa/\sigma$ by $\sigma > 0$ gives

$$\Big(c_y - \frac{c_x}{q}\Big)\sigma - \kappa > 0, \qquad \Big(c_z - \frac{c_x}{q}\Big)\sigma - \kappa > 0.$$

Step 5 (some off-x mass is present). Assume in addition $(x, y, z_{\mathrm{sp}}) \neq (M/q, 0, 0)$. If $y = z_{\mathrm{sp}} = 0$, the simplex equation $qx + y + z_{\mathrm{sp}} = M$ would force $x = M/q$, i.e. the point would be the spread vertex; hence $y > 0$ or $z_{\mathrm{sp}} > 0$. Since $y, z_{\mathrm{sp}} \geq 0$ and both coefficients of Step 4 are strictly positive,

$$0 < \Big(\big(c_y - \tfrac{c_x}{q}\big)\sigma - \kappa\Big)y + \Big(\big(c_z - \tfrac{c_x}{q}\big)\sigma - \kappa\Big)z_{\mathrm{sp}}.$$

Step 6 (the difference identity). Write $M/q = x + (y + z_{\mathrm{sp}})/q$, which is the simplex equation divided by q. Then

$$\big(c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sigma x\big) - \Big(c_x\frac{M}{q} + \kappa\sigma\frac{M}{q}\Big) = c_y y + c_z z_{\mathrm{sp}} - \frac{c_x}{q}(y + z_{\mathrm{sp}}) - \frac{\kappa\sigma}{q}(y + z_{\mathrm{sp}}),$$

and multiplying by σ and using $\sigma^2 = q$ turns the last term into $\kappa(y + z_{\mathrm{sp}})$, giving

$$\begin{aligned}\Big(c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sigma x - c_x\frac{M}{q} - \kappa\sigma\frac{M}{q}\Big)\sigma \\ = \Big(\big(c_y - \tfrac{c_x}{q}\big)\sigma - \kappa\Big)y + \Big(\big(c_z - \tfrac{c_x}{q}\big)\sigma - \kappa\Big)z_{\mathrm{sp}}.\end{aligned}$$

By Step 5 the right-hand side is strictly positive, and $\sigma > 0$, so the bracket on the left is strictly positive:

$$c_x\frac{M}{q} + \kappa\sigma\frac{M}{q} < c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sigma x.$$

Step 7 (conclusion). Chaining Steps 2, 6 and 3,

$$\phi\Big(\frac{M}{q}, 0, 0\Big) = c_x\frac{M}{q} + \kappa\sigma\frac{M}{q} < c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sigma x \leq \phi(x, y, z_{\mathrm{sp}})$$

for every $(x, y, z_{\mathrm{sp}}) \in T_m$ other than $(2m/q, 0, 0)$, which is the asserted strict inequality. □

Lemma 3 plays the corresponding role for the spread vertex. In the notation of Theorem 3, this vertex is the reduced representation of X_{spread}. The coefficient inequalities make the x-direction uniquely cheapest, so the relaxed optimizer concentrates all simplex mass there.

The iid result uses a different geometric fact: the point $(1, 1, 1)$ is the unique minimizer of the Frobenius part on the simplex, but it minimizes the full objective only when the linear coefficients are balanced across the three spectral directions.

Lemma 4. [`lem:frobenius-center-certificate`] (Frobenius Center Certificate) *Let $m \in \mathbb{N}$, set $q = 2(m-1)$, and suppose $q > 0$. For any real coefficients c_x, c_y, c_z and any real κ, define*

$$T_m = \{(x, y, z_{\mathrm{sp}}) : x, y, z_{\mathrm{sp}} \geq 0,\ qx + y + z_{\mathrm{sp}} = 2m\}$$

and

$$\phi_{r,\kappa}(x, y, z_{\mathrm{sp}}) = c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}.$$

Then:

- ***(Frobenius center.)*** *The point* $(1,1,1)$ *belongs to* T_m*, and for every* $(x,y,z_{\mathrm{sp}}) \in T_m$ *with* $(x,y,z_{\mathrm{sp}}) \neq (1,1,1)$*,*

$$\sqrt{q \cdot 1^2 + 1^2 + 1^2} < \sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}.$$

- ***(Certificate necessity.)*** *If* $(1,1,1)$ *minimizes* $\phi_{r,\kappa}$ *over* T_m*, that is, if*

$$\phi_{r,\kappa}(1,1,1) \leq \phi_{r,\kappa}(x,y,z_{\mathrm{sp}}) \quad \textit{for all } (x,y,z_{\mathrm{sp}}) \in T_m,$$

then

$$\frac{c_x}{q} = c_y \quad \textit{and} \quad c_y = c_z.$$

- ***(Strict optimality.)*** *If* $\kappa > 0$*,* $\frac{c_x}{q} = c_y$*, and* $c_y = c_z$*, then for every* $(x,y,z_{\mathrm{sp}}) \in T_m$ *with* $(x,y,z_{\mathrm{sp}}) \neq (1,1,1)$*,*

$$\phi_{r,\kappa}(1,1,1) < \phi_{r,\kappa}(x,y,z_{\mathrm{sp}}).$$

Proof of Lemma 4. Throughout $q = 2(m-1) > 0$, so that

$$q + 2 = 2m.$$

Step 1 (the center is feasible). The point $(1,1,1)$ has nonnegative coordinates and satisfies $q \cdot 1 + 1 + 1 = q + 2 = 2m$, hence $(1,1,1) \in T_m$.

Step 2 (a bias–variance decomposition of the weighted norm on T_m*).* Let $(x,y,z_{\mathrm{sp}}) \in T_m$. Expanding the squares,

$$q(x-1)^2 + (y-1)^2 + (z_{\mathrm{sp}}-1)^2 = \big(qx^2 + y^2 + z_{\mathrm{sp}}^2\big) - 2\big(qx + y + z_{\mathrm{sp}}\big) + (q+2),$$

and substituting the simplex equation $qx + y + z_{\mathrm{sp}} = 2m$ together with $q + 2 = 2m$ gives $-2(2m) + (q+2) = -(q+2)$. Hence

$$qx^2 + y^2 + z_{\mathrm{sp}}^2 = (q \cdot 1^2 + 1^2 + 1^2) + q(x-1)^2 + (y-1)^2 + (z_{\mathrm{sp}}-1)^2.$$

Step 3 (strict minimality of the center for the norm). Suppose $(x,y,z_{\mathrm{sp}}) \in T_m$ with $(x,y,z_{\mathrm{sp}}) \neq (1,1,1)$. Each of $q(x-1)^2$, $(y-1)^2$, $(z_{\mathrm{sp}}-1)^2$ is nonnegative, and if their sum vanished then, since $q > 0$, each would vanish and the point would be $(1,1,1)$. Therefore

$$0 < q(x-1)^2 + (y-1)^2 + (z_{\mathrm{sp}}-1)^2,$$

so by Step 2, $q \cdot 1^2 + 1^2 + 1^2 < qx^2 + y^2 + z_{\mathrm{sp}}^2$, and taking square roots (both arguments being nonnegative and the square root strictly increasing)

$$\sqrt{q \cdot 1^2 + 1^2 + 1^2} < \sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}.$$

This is the Frobenius-center assertion.

Step 4 (necessity of $c_x/q = c_y$*).* Assume $\phi_{r,\kappa}(1,1,1) \leq \phi_{r,\kappa}(x,y,z_{\mathrm{sp}})$ for all $(x,y,z_{\mathrm{sp}}) \in T_m$, and consider the curve

$$t \longmapsto (1+t,\ 1-qt,\ 1).$$

It is feasible for $|t| < \min\{1, 1/q\}$: the coordinates $1+t$ and $1-qt$ are then positive, the third is 1, and

$$q(1+t) + (1-qt) + 1 = q + 2 = 2m$$

for every t, so the curve stays in T_m. Consequently $t=0$ is a local minimum of

$$g(t) = \phi_{r,\kappa}(1+t, 1-qt, 1) = c_x(1+t) + c_y(1-qt) + c_z + \kappa\sqrt{q(1+t)^2 + (1-qt)^2 + 1}.$$

The radicand has derivative

$$\frac{d}{dt}\Big(q(1+t)^2 + (1-qt)^2 + 1\Big)\Big|_{t=0} = 2q(1+t) - 2q(1-qt)\Big|_{t=0} = 2q - 2q = 0,$$

and its value at $t=0$ is $q+2>0$, so the chain rule makes the derivative of the square-root term vanish at $t=0$. Hence $g'(0) = c_x - qc_y$, and local minimality at an interior point forces

$$c_x - qc_y = 0, \qquad \text{i.e.} \qquad \frac{c_x}{q} = c_y.$$

Step 5 (necessity of $c_y = c_z$). Consider likewise

$$t \longmapsto (1,\, 1+t,\, 1-t),$$

feasible for $|t|<1$ since then all three coordinates are nonnegative and $q\cdot 1 + (1+t) + (1-t) = q+2 = 2m$. Writing $h(t) = \phi_{r,\kappa}(1, 1+t, 1-t)$, the radicand $q + (1+t)^2 + (1-t)^2$ has derivative $2(1+t) - 2(1-t)$, which vanishes at $t=0$, and its value there is $q+2>0$; so again the square-root term contributes nothing to $h'(0)$ and

$$h'(0) = c_y - c_z = 0, \qquad \text{i.e.} \qquad c_y = c_z.$$

Step 6 (strict optimality under the balance conditions). Suppose finally $\kappa>0$, $c_x/q = c_y$ and $c_y = c_z$; thus $c_x = qc_y$ and $c_z = c_y$. For every $(x, y, z_{\mathrm{sp}}) \in T_m$ the linear part is then constant:

$$c_x x + c_y y + c_z z_{\mathrm{sp}} = c_y\big(qx + y + z_{\mathrm{sp}}\big) = c_y \cdot 2m = c_y(q+2) = c_x\cdot 1 + c_y \cdot 1 + c_z \cdot 1.$$

For $(x, y, z_{\mathrm{sp}}) \neq (1,1,1)$, Step 3 gives a strictly larger norm term, and multiplying that strict inequality by $\kappa>0$ preserves strictness. Adding the equal linear parts to the strict norm comparison yields

$$\phi_{r,\kappa}(1,1,1) < \phi_{r,\kappa}(x, y, z_{\mathrm{sp}}),$$

so $(1,1,1)$ is the strict minimizer. $\square$

Lemma 4 explains why the iid covariance has a knife-edge exactness condition in Theorem 2. The norm term alone favors the point $(1,1,1)$, but the linear coefficients must be equal after accounting for the q-multiplicity of the x-coordinate. When that equality fails, the linear part tilts the objective away from the iid point at every finite positive robustness weight.

The next statement is the implementability restriction in reduced form. It is where the parity of the block size enters the geometry.

Lemma 5. [`lem:pm-reduced-slice-characterization`] (Parity-Truncated Reduced Slice) *Let $m \geq 2$ be an integer and let $u, v \in \mathbb{R}$. Write*

$$x = 1-u, \qquad y = 1 + (m-1)u - mv, \qquad z_{\mathrm{sp}} = 1 + (m-1)u + mv,$$

and let $d_m = 0$ when m is even and $d_m = 2/m$ when m is odd. Then the block covariance matrix $X(u,v)$ belongs to the implementable covariance class $\mathcal{C}_m^{\pm}$ of Definition 5 if and only if

$$x \geq 0, \qquad y \geq 0, \qquad z_{\mathrm{sp}} \geq 0, \qquad 2(m-1)x + y + z_{\mathrm{sp}} = 2m,$$

and

$$y + z_{\mathrm{sp}} \geq d_m.$$

Proof of Lemma 5. Write $q = 2(m-1)$ and keep the abbreviations $x = 1-u$, $y = 1+(m-1)u - mv$, $z_{\mathrm{sp}} = 1+(m-1)u+mv$ of the statement, so that

$$y + z_{\mathrm{sp}} = 2 + 2(m-1)u, \qquad z_{\mathrm{sp}} - y = 2mv, \qquad qx + y + z_{\mathrm{sp}} = 2m,$$

the last identity holding for all u, v because the u- and v-terms cancel. Throughout, (S_A, S_B) are the two community sums of Definition 8.

Part I: necessity. Suppose $X(u,v) \in \mathcal{C}_m^{\pm}$, i.e. $X(u,v) = X(P) = \mathbb{E}_P[ZZ^\top]$ for some $P \in \mathcal{P}_m^{\mathrm{sym}}$. Reading the entries of $X(P)$ off the block pattern of $X(u,v)$ gives $\mathbb{E}[Z_iZ_j] = 1$ if $i = j$, $= u$ if $i \neq j$ lie in the same community, and $= v$ if they lie in opposite communities.

(I.1) The three second-moment identities. Expanding $S_A^2 = \sum_{i,j \in A_m} Z_iZ_j$, there are m diagonal pairs and $m(m-1)$ off-diagonal same-community pairs, and likewise for S_B^2; expanding $S_AS_B = \sum_{i \in A_m}\sum_{j \in B_m} Z_iZ_j$, all m^2 pairs are cross-community. Hence

$$\mathbb{E}[S_A^2] = \mathbb{E}[S_B^2] = m + m(m-1)u, \qquad \mathbb{E}[S_AS_B] = m^2v.$$

(I.2) Nonnegativity of y and z_{sp}. Combining the three identities,

$$\mathbb{E}\big[(S_A - S_B)^2\big] = 2m + 2m(m-1)u - 2m^2v = 2m\,y, \qquad \mathbb{E}\big[(S_A + S_B)^2\big] = 2m + 2m(m-1)u + 2m^2v = 2m\,z_{\mathrm{sp}}.$$

Both left-hand sides are expectations of squares, hence nonnegative, and $m > 0$; therefore $y \geq 0$ and $z_{\mathrm{sp}} \geq 0$.

(I.3) Nonnegativity of x. Since $|A_m| = m \geq 2$, pick distinct $i, j \in A_m$. Because $Z_i^2 = Z_j^2 = 1$ pointwise,

$$0 \leq \mathbb{E}\big[(Z_i - Z_j)^2\big] = 2 - 2\,\mathbb{E}[Z_iZ_j] = 2 - 2u,$$

so $u \leq 1$, i.e. $x = 1 - u \geq 0$.

Together with the identity $qx + y + z_{\mathrm{sp}} = 2m$ recorded above, this gives the reduced-triangle conditions

$$x \geq 0, \qquad y \geq 0, \qquad z_{\mathrm{sp}} \geq 0, \qquad 2(m-1)x + y + z_{\mathrm{sp}} = 2m.$$

(I.4) The parity bound. If m is even then $d_m = 0 \leq y + z_{\mathrm{sp}}$ by (I.2), and there is nothing more to prove. If m is odd, then S_A is a sum of m terms each equal to ± 1 with m odd, so S_A is an odd integer; in particular $S_A \neq 0$ and, being an integer, $S_A^2 \geq 1$. The same holds for S_B. Taking expectations,

$$\mathbb{E}[S_A^2] \geq 1, \qquad \mathbb{E}[S_B^2] \geq 1.$$

On the other hand, by (I.1),

$$\mathbb{E}[S_A^2] + \mathbb{E}[S_B^2] = 2m + 2m(m-1)u = m\,(y + z_{\mathrm{sp}}),$$

so $m(y + z_{\mathrm{sp}}) \geq 2$ and hence $y + z_{\mathrm{sp}} \geq 2/m = d_m$.

Part II: sufficiency. Conversely, assume $x, y, z_{\mathrm{sp}} \geq 0$, $qx + y + z_{\mathrm{sp}} = 2m$ and $y + z_{\mathrm{sp}} \geq d_m$. We exhibit a law in $\mathcal{P}_m^{\mathrm{sym}}$ whose second moment is $X(u,v)$.

Two general facts are used. First, the block parameterization is affine: for weights $w_k \geq 0$ with $\sum_k w_k = 1$,

$$\sum_k w_k\, X(u_k, v_k) = X\Big(\sum_k w_k u_k,\ \sum_k w_k v_k\Big),$$

because every diagonal entry of the left side is $\sum_k w_k = 1$ and the within- and across-community entries average separately. Second, if $P = \sum_k w_k P_k$ is a convex mixture of laws $P_k \in \mathcal{P}_m^{\mathrm{sym}}$, then $P \in \mathcal{P}_m^{\mathrm{sym}}$ (sign symmetry and invariance under the two-block automorphism group are preserved by convex combinations) and, by linearity of expectation, $X(P) = \sum_k w_k X(P_k)$.

The building blocks are uniform laws on supports that are invariant under global sign reversal and under the two-block automorphism group; each such law lies in $\mathcal{P}_m^{\mathrm{sym}}$, and its second moment is therefore block-symmetric, say $X(u_\bullet, v_\bullet)$, with $u_\bullet, v_\bullet$ determined through (I.1) by $\mathbb{E}[S_A^2]$ and $\mathbb{E}[S_A S_B]$:

$$u_\bullet = \frac{\mathbb{E}[S_A^2] - m}{m(m-1)}, \qquad v_\bullet = \frac{\mathbb{E}[S_A S_B]}{m^2}.$$

(II.1) Even m: a three-vertex mixture. Take

- P_{cut}, uniform on the cut assignment ($+1$ on A_m, -1 on B_m) and its global sign reversal, for which $S_A^2 = m^2$ and $S_A S_B = -m^2$, hence $X(P_{\mathrm{cut}}) = X(1, -1)$;
- P_{all}, uniform on the two constant assignments (all $+1$, all -1), for which $S_A^2 = m^2$ and $S_A S_B = m^2$, hence $X(P_{\mathrm{all}}) = X(1, 1)$;
- P_{spr}, uniform on $\{z : S_A(z) = S_B(z) = 0\}$, a nonempty support when m is even, for which $\mathbb{E}[S_A^2] = \mathbb{E}[S_A S_B] = 0$ and hence $X(P_{\mathrm{spr}}) = X\Big(-\frac{1}{m-1}, 0\Big)$.

Assign the weights

$$w_{\mathrm{cut}} = \frac{y}{2m}, \qquad w_{\mathrm{all}} = \frac{z_{\mathrm{sp}}}{2m}, \qquad w_{\mathrm{spr}} = \frac{(m-1)x}{m}.$$

They are nonnegative because $x, y, z_{\mathrm{sp}} \geq 0$ and $m \geq 2$, and they sum to one because

$$\frac{y}{2m} + \frac{z_{\mathrm{sp}}}{2m} + \frac{(m-1)x}{m} = \frac{y + z_{\mathrm{sp}} + 2(m-1)x}{2m} = \frac{2m}{2m} = 1.$$

Their barycentre in the (u, v) parameters is

$$w_{\mathrm{cut}} \cdot 1 + w_{\mathrm{all}} \cdot 1 + w_{\mathrm{spr}} \cdot \Big(-\frac{1}{m-1}\Big) = \frac{y + z_{\mathrm{sp}}}{2m} - \frac{x}{m} = \frac{2m - 2(m-1)x}{2m} - \frac{x}{m} = 1 - x = u$$

and

$$w_{\mathrm{cut}} \cdot (-1) + w_{\mathrm{all}} \cdot 1 + w_{\mathrm{spr}} \cdot 0 = \frac{z_{\mathrm{sp}} - y}{2m} = v.$$

Hence the mixture $P = w_{\mathrm{cut}} P_{\mathrm{cut}} + w_{\mathrm{all}} P_{\mathrm{all}} + w_{\mathrm{spr}} P_{\mathrm{spr}}$ lies in $\mathcal{P}_m^{\mathrm{sym}}$ and has $X(P) = X(u, v)$, so $X(u, v) \in \mathcal{C}_m^{\pm}$.

(II.2) Odd m: a four-vertex mixture. For odd m the support $\{S_A = S_B = 0\}$ is empty, so P_{spr} is unavailable and the two parity vertices replace it:

- $P_{\mathrm{cut}}^{\mathrm{par}}$, uniform on $\{z : (S_A, S_B) = (1, -1) \text{ or } (-1, 1)\}$ — a nonempty support for odd m — for which $S_A^2 = 1$ and $S_A S_B = -1$, hence
$$X\big(P_{\mathrm{cut}}^{\mathrm{par}}\big) = X\Big(-\frac{1}{m}, -\frac{1}{m^2}\Big);$$
- $P_{\mathrm{all}}^{\mathrm{par}}$, uniform on $\{z : (S_A, S_B) = (1, 1) \text{ or } (-1, -1)\}$, for which $S_A^2 = 1$ and $S_A S_B = 1$, hence
$$X\big(P_{\mathrm{all}}^{\mathrm{par}}\big) = X\Big(-\frac{1}{m}, \frac{1}{m^2}\Big).$$

In reduced coordinates these two points are $\left(\frac{m+1}{m}, \frac{2}{m}, 0\right)$ and $\left(\frac{m+1}{m}, 0, \frac{2}{m}\right)$, i.e. they sit on the parity boundary $y + z_{\mathrm{sp}} = 2/m$, whereas $X(1,-1)$ and $X(1,1)$ are the outer vertices $(0, 2m, 0)$ and $(0, 0, 2m)$.

Put $s = y + z_{\mathrm{sp}}$ and

$$\mu = \frac{y}{s}, \qquad \lambda = \frac{s - 2/m}{2m - 2/m},$$

which are well defined because $s \geq 2/m > 0$ and $2m - 2/m > 0$ for $m \geq 2$. Then $0 \leq \mu \leq 1$ since $0 \leq y \leq s$, and $0 \leq \lambda \leq 1$ since $2/m \leq s \leq 2m$, the upper bound coming from $s = 2m - qx \leq 2m$ with $qx \geq 0$. Take the four weights

$$\big(\lambda\mu,\ \lambda(1-\mu),\ (1-\lambda)\mu,\ (1-\lambda)(1-\mu)\big)$$

on $\left(P_{\mathrm{cut}}, P_{\mathrm{all}}, P_{\mathrm{cut}}^{\mathrm{par}}, P_{\mathrm{all}}^{\mathrm{par}}\right)$; they are nonnegative and sum to one. Their barycentre in the u-parameter is

$$\lambda \cdot 1 + (1-\lambda)\Big(-\frac{1}{m}\Big) = \frac{(m+1)\lambda - 1}{m} = u,$$

where the last equality is the definition of λ rewritten through $s = 2m - 2(m-1)(1-u)$, which gives $\lambda = (1 + mu)/(m+1)$. For the v-parameter, the two μ-pairs collapse to

$$\lambda\mu(-1) + \lambda(1-\mu)(1) + (1-\lambda)\mu\Big(-\frac{1}{m^2}\Big) + (1-\lambda)(1-\mu)\frac{1}{m^2} = (1-2\mu)\Big(\lambda + \frac{1-\lambda}{m^2}\Big),$$

and the two factors evaluate to

$$1 - 2\mu = \frac{z_{\mathrm{sp}} - y}{s}, \qquad \lambda + \frac{1-\lambda}{m^2} = \frac{1}{m^2} + \lambda\,\frac{m^2-1}{m^2} = \frac{s}{2m},$$

the second using $\lambda = (ms-2)/\big(2(m^2-1)\big)$. Their product is $(z_{\mathrm{sp}} - y)/(2m) = v$. Hence the four-point mixture lies in $\mathcal{P}_m^{\mathrm{sym}}$ and has second moment $X(u,v)$, so $X(u,v) \in \mathcal{C}_m^{\pm}$.

Parts I and II are the two implications of the asserted equivalence. $\square$

Lemma 5 gives implementability a transparent low-dimensional form. Relative to the relaxed simplex, the implementable slice adds the lower bound $y + z_{\mathrm{sp}} \geq d_m$. For even m, $d_m = 0$, so this lower bound is already implied by nonnegativity. For odd m, the positive lower bound records the positive squared magnitude of the realized block sums.

It remains to solve the reduced optimization problem. The next lemma gives the unconstrained weighted-simplex solution as a finite active-set calculation.

Lemma 6. [`lem:weighted-simplex-active-set`] (Weighted Simplex Active Set) *Let $m \geq 2$, let $M > 0$ satisfy $M = 2m$, and let $\mathcal{I} = \{x, y, z_{\mathrm{sp}}\}$. Let $\alpha = (\alpha_i)_{i\in\mathcal{I}}$, let $\kappa \geq 0$, and let $\beta = (\beta_i)_{i\in\mathcal{I}}$ have strictly positive entries with*

$$\beta_x = \frac{1}{2(m-1)}, \qquad \beta_y = \beta_{z_{\mathrm{sp}}} = 1.$$

For $t \in \mathbb{R}^{\mathcal{I}}$, write

$$\Delta_M = \Big\{t :\ t_i \geq 0 \text{ for all } i \in \mathcal{I},\ \sum_{i\in\mathcal{I}} t_i = M\Big\},$$

and

$$\Phi(t) = \sum_{i\in\mathcal{I}} \alpha_i t_i + \kappa\sqrt{\sum_{i\in\mathcal{I}} \beta_i t_i^2}.$$

Then the following hold.

- ***(Positive robustness.)*** *If $\kappa > 0$, there is a unique pair (S, λ), where $S \subseteq \mathcal{I}$ is nonempty and $\lambda \in \mathbb{R}$, such that*

$$\sum_{i \in S} \frac{(\lambda - \alpha_i)^2}{\beta_i} = \kappa^2, \qquad \alpha_i < \lambda \text{ for every } i \in S, \qquad \lambda \le \alpha_j \text{ for every } j \notin S.$$

For this unique pair, define

$$t_i^\star = \begin{cases} M \dfrac{(\lambda - \alpha_i)/\beta_i}{\sum_{h \in S} (\lambda - \alpha_h)/\beta_h}, & i \in S, \\ 0, & i \notin S. \end{cases}$$

Then $t^\star \in \Delta_M$, and $t^\star$ is the unique minimizer of Φ over Δ_M in the strict sense that, for every $s \in \Delta_M$ with $s \neq t^\star$,

$$\Phi(t^\star) < \Phi(s).$$

Its value is

$$\Phi(t^\star) = M\lambda.$$

- ***(Zero robustness.)*** *If $\kappa = 0$, then for every $t \in \mathbb{R}^{\mathcal{I}}$,*

$$\Big(t \in \Delta_M \text{ and } \Phi(t) \le \Phi(s) \text{ for every } s \in \Delta_M\Big) \quad \Longleftrightarrow \quad t \in \Delta_M \text{ and whenever } t_i \neq 0,\ \alpha_i \le \alpha_j \text{ for every } j \in$$

Equivalently, the minimizer set is exactly the exposed α-minimizing face of Δ_M.

Proof of Lemma 6. Throughout $M = 2m > 0$ and every $\beta_i > 0$.

Part A: $\kappa > 0$.

Step A1 (an admissible pair exists). Define the threshold function

$$G(\lambda) = \sum_{i \in \mathcal{I}} \frac{\big(\max\{\lambda - \alpha_i, 0\}\big)^2}{\beta_i}, \qquad \lambda \in \mathbb{R},$$

which is continuous because it is a finite sum of continuous functions of λ. Let $\alpha_{\min} = \min_i \alpha_i$ and $\alpha_{\max} = \max_i \alpha_i$, attained at indices $k_{\min}$ and $k_{\max}$. At $\lambda = \alpha_{\min}$ every positive part vanishes, so

$$G(\alpha_{\min}) = 0 \le \kappa^2.$$

At $\lambda_+ = \alpha_{\max} + \kappa\sqrt{\beta_{k_{\max}}}$ the single summand at $k_{\max}$ already equals

$$\frac{\big(\kappa\sqrt{\beta_{k_{\max}}}\big)^2}{\beta_{k_{\max}}} = \kappa^2,$$

and all summands are nonnegative, so $G(\lambda_+) \ge \kappa^2$. Since $\alpha_{\min} \le \lambda_+$, the intermediate value theorem yields $\lambda \in [\alpha_{\min}, \lambda_+]$ with $G(\lambda) = \kappa^2$. Put

$$S = \{i \in \mathcal{I} : \alpha_i < \lambda\}.$$

Off S the positive part $\max\{\lambda - \alpha_i, 0\}$ vanishes, so

$$\sum_{i \in S} \frac{(\lambda - \alpha_i)^2}{\beta_i} = G(\lambda) = \kappa^2,$$

and S is nonempty (otherwise the left side would be $0 \neq \kappa^2$). By construction $\alpha_i < \lambda$ on S and $\lambda \leq \alpha_j$ off S, so (S, λ) is admissible.

Step A2 (the active-set point lies in Δ_M). Write

$$D = \sum_{h \in S} \frac{\lambda - \alpha_h}{\beta_h}.$$

Each summand is strictly positive ($\lambda > \alpha_h$ on S, $\beta_h > 0$) and $S \neq \emptyset$, so $D > 0$. Hence each $t_i^\star = M\big((\lambda - \alpha_i)/\beta_i\big)/D$ with $i \in S$ is nonnegative — indeed strictly positive — while $t_j^\star = 0$ off S; and

$$\sum_{i \in \mathcal{I}} t_i^\star = \frac{M}{D} \sum_{i \in S} \frac{\lambda - \alpha_i}{\beta_i} = \frac{M}{D} \cdot D = M.$$

Thus $t^\star \in \Delta_M$, and moreover

$$t_i^\star > 0 \iff i \in S.$$

Step A3 (norm and stationarity at $t^\star$). Using $t_i^\star = 0$ off S and the admissibility equation,

$$\sum_{i \in \mathcal{I}} \beta_i (t_i^\star)^2 = \frac{M^2}{D^2} \sum_{i \in S} \beta_i \frac{(\lambda - \alpha_i)^2}{\beta_i^2} = \frac{M^2}{D^2} \sum_{i \in S} \frac{(\lambda - \alpha_i)^2}{\beta_i} = \Big(\frac{M\kappa}{D}\Big)^2,$$

so, writing $N = \sqrt{\sum_i \beta_i (t_i^\star)^2}$,

$$N = \frac{M\kappa}{D} > 0.$$

Consequently, for every $i \in S$,

$$\frac{\kappa\, \beta_i t_i^\star}{N} = \frac{\kappa\, \beta_i}{N} \cdot \frac{M(\lambda - \alpha_i)}{\beta_i D} = \frac{\kappa M}{ND}(\lambda - \alpha_i) = \lambda - \alpha_i,$$

which is the stationarity relation on the active support.

Step A4 (the optimal value). Multiplying the stationarity relation of Step A3 by $t_i^\star$, summing over $i \in S$ (equivalently over all of $\mathcal{I}$, since $t^\star$ vanishes off S), and using $\sum_i \beta_i (t_i^\star)^2 = N^2$ and $\sum_i t_i^\star = M$,

$$\sum_{i \in \mathcal{I}} \alpha_i t_i^\star + \kappa N = \sum_{i \in \mathcal{I}} \alpha_i t_i^\star + \frac{\kappa}{N} \sum_{i \in \mathcal{I}} \beta_i (t_i^\star)^2 = \lambda \sum_{i \in \mathcal{I}} t_i^\star = M\lambda,$$

i.e.

$$\Phi(t^\star) = M\lambda.$$

Step A5 (strict optimality). Let $s \in \Delta_M$ with $s \neq t^\star$, and write $N(s) = \sqrt{\sum_i \beta_i s_i^2}$. Since β is strictly positive, $\sum_i s_i = \sum_i t_i^\star = M \neq 0$ and $s \neq t^\star$, the strict weighted Cauchy–Schwarz inequality gives

$$\sum_{i \in \mathcal{I}} \beta_i t_i^\star s_i < N\, N(s).$$

Subtracting $\sum_i \beta_i (t_i^\star)^2 = N^2$ from both sides and multiplying by $\kappa/N > 0$,

$$\frac{\kappa}{N} \sum_{i \in \mathcal{I}} \beta_i t_i^\star\, (s_i - t_i^\star) < \kappa\big(N(s) - N\big).$$

By the stationarity relation of Step A3 (and $t^\star_j = 0$ off S) the left side equals $\sum_{i\in S}(\lambda-\alpha_i)(s_i-t^\star_i)$, so

$$\Phi(s) - \Phi(t^\star) = \sum_{i\in\mathcal{I}} \alpha_i(s_i - t^\star_i) + \kappa\big(N(s) - N\big) > \sum_{i\in\mathcal{I}} \alpha_i(s_i - t^\star_i) + \sum_{i\in S}(\lambda - \alpha_i)(s_i - t^\star_i).$$

Setting $\gamma_i = \lambda$ for $i \in S$ and $\gamma_i = \alpha_i$ for $i \notin S$, the right-hand side is $\sum_i \gamma_i(s_i - t^\star_i)$; since $\gamma_i t^\star_i = \lambda t^\star_i$ for every i (on S by definition of γ, off S because $t^\star_i = 0$) and $\sum_i t^\star_i = M$, it equals

$$\sum_{i\in\mathcal{I}} \gamma_i s_i - \lambda M.$$

Finally $\gamma_i \geq \lambda$ for every i — with equality on S and by admissibility $\alpha_j \geq \lambda$ off S — and $s_i \geq 0$, so $\sum_i \gamma_i s_i \geq \lambda \sum_i s_i = \lambda M$ and the last display is nonnegative. Hence

$$\Phi(t^\star) < \Phi(s) \qquad \text{for every } s \in \Delta_M \text{ with } s \neq t^\star.$$

Step A6 (uniqueness of (S, λ)). Let (S', λ') be another admissible pair, with active-set point $t^{\star\prime}$. Steps A2 and A5 apply to it as well, so both $t^\star$ and $t^{\star\prime}$ are strict minimizers of Φ over Δ_M; if they differed, each would have strictly smaller value than the other, which is impossible. Hence $t^\star = t^{\star\prime}$. By the support identity of Step A2,

$$i \in S \iff t^\star_i > 0 \iff t^{\star\prime}_i > 0 \iff i \in S',$$

so $S = S'$. By Step A4, $M\lambda = \Phi(t^\star) = \Phi(t^{\star\prime}) = M\lambda'$, and $M > 0$ gives $\lambda = \lambda'$.

Part B: $\kappa = 0$. Here $\Phi(t) = \sum_i \alpha_i t_i$ is linear. Fix k with $\alpha_k = \min_j \alpha_j$.

If $t \in \Delta_M$ minimizes Φ over Δ_M, compare it with the vertex of Δ_M that places all the mass on coordinate k, i.e. the point w with $w_k = M$ and $w_j = 0$ for $j \neq k$, whose value is $\Phi(w) = M\alpha_k$:

$$\sum_{i\in\mathcal{I}} \alpha_i t_i \leq M\alpha_k = \alpha_k \sum_{i\in\mathcal{I}} t_i, \qquad \text{i.e.} \qquad \sum_{i\in\mathcal{I}} (\alpha_i - \alpha_k) t_i \leq 0.$$

Every summand is nonnegative ($\alpha_i \geq \alpha_k$, $t_i \geq 0$), so every summand vanishes; hence $t_i \neq 0$ forces $\alpha_i = \alpha_k \leq \alpha_j$ for every j. That is exactly membership in the exposed α-minimizing face.

Conversely, if $t \in \Delta_M$ is supported only on α-minimizing coordinates, then $\alpha_i t_i = \alpha_k t_i$ for every i and

$$\Phi(t) = \alpha_k \sum_{i\in\mathcal{I}} t_i = M\alpha_k,$$

whereas every $s \in \Delta_M$ satisfies $\Phi(s) = \sum_i \alpha_i s_i \geq \alpha_k \sum_i s_i = M\alpha_k$. So t is a minimizer. The two implications give the stated equivalence, and the minimizer set is precisely the exposed α-minimizing face of Δ_M. □

Lemma 6 is the finite case split behind the active-set formula in Theorem 4. For $\kappa > 0$, the support and multiplier determine both the optimizer and the value. For $\kappa = 0$, the norm term drops out and the minimizers form the exposed face associated with the smallest linear coefficient.

The implementable problem may add the parity lower bound. The following truncation statement records how the solution changes when the unconstrained active-set point lies outside that truncated simplex.

Lemma 7. [`lem:weighted-simplex-truncation`] (Weighted Simplex Truncation) *Let $m \ge 2$, let $M, d \in \mathbb{R}$, and let $\alpha = (\alpha_x, \alpha_y, \alpha_z)$, $\beta = (\beta_x, \beta_y, \beta_z)$, and $\kappa \in \mathbb{R}$. Suppose:*

- ***(Mass and truncation.)*** *$M > 0$, $M = 2m$, and $0 \le d \le M$.*
- ***(Weights.)*** *Each β_i is strictly positive and*

$$\beta = \left(\frac{1}{2(m-1)}, 1, 1\right),$$

  so in particular $\beta_y = \beta_z = 1$.
- ***(Robustness.)*** *$\kappa \ge 0$.*

Write

$$K_d = \{t \in \Delta_M : t_y + t_z \ge d\}, \qquad \Phi(t) = \sum_i \alpha_i t_i + \kappa \sqrt{\sum_i \beta_i t_i^2}.$$

For $\kappa > 0$, for every nonempty active support $S \subseteq \{x, y, z\}$ and multiplier λ satisfying

$$\sum_{i \in S} \frac{(\lambda - \alpha_i)^2}{\beta_i} = \kappa^2, \qquad \alpha_i < \lambda \text{ for } i \in S, \qquad \lambda \le \alpha_j \text{ for } j \notin S,$$

let t^{rel} be the active-set point

$$t_i^{\mathrm{rel}} = \begin{cases} \dfrac{M((\lambda - \alpha_i)/\beta_i)}{\sum_{h \in S} (\lambda - \alpha_h)/\beta_h}, & i \in S, \\ 0, & i \notin S. \end{cases}$$

If $t^{\mathrm{rel}} \in K_d$, then

$$\Phi(t^{\mathrm{rel}}) \le \Phi(s) \qquad \text{for every } s \in K_d.$$

If $t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} < d$, then

$$t^\star = (M - d, s^\star, d - s^\star) \in K_d$$

and

$$\Phi(t^\star) \le \Phi(s) \qquad \text{for every } s \in K_d,$$

where $s^\star = \mathrm{truncSelector}(M, d, \alpha, \beta, \kappa)$, namely with

$$\delta = \alpha_y - \alpha_z, \qquad A = \beta_x (M - d)^2,$$

$$s^\star = \begin{cases} 0, & \delta \ge \dfrac{\kappa d}{\sqrt{A + d^2}}, \\ d, & \delta \le -\dfrac{\kappa d}{\sqrt{A + d^2}}, \\ \dfrac{d - \delta\sqrt{(A + d^2/2)/(\kappa^2 - \delta^2/2)}}{2}, & \text{otherwise}. \end{cases}$$

For $\kappa = 0$, let t^{rel} be a point of the exposed α-minimizing face

$$\{t \in \Delta_M : \ t_i \ne 0 \Rightarrow \alpha_i \le \alpha_j \text{ for every } j\}.$$

Assume the tie convention that, if this exposed face contains at least one point of K_d, then the selected point t^{rel} lies in K_d. If $t^{\mathrm{rel}} \in K_d$, then

$$\Phi(t^{\mathrm{rel}}) \le \Phi(s) \qquad \textit{for every } s \in K_d.$$

If $t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} < d$, then

$$t^\star = (M-d, s^\star, d-s^\star) \in K_d$$

and

$$\Phi(t^\star) \le \Phi(s) \qquad \textit{for every } s \in K_d,$$

where $s^\star = \mathrm{truncSelector}(M, d, \alpha, \beta, 0)$, equivalently, in the ordered selector convention, $s^\star = 0$ if $\alpha_y - \alpha_z \ge 0$ and $s^\star = d$ if $\alpha_y - \alpha_z < 0$.

Proof of Lemma 7. Write $\delta = \alpha_y - \alpha_z$, $A = \beta_x(M-d)^2 \ge 0$, and for $\sigma \in \mathbb{R}$ let

$$h(\sigma) = (M-d, \sigma, d-\sigma)$$

be the point of the boundary segment H_d with second coordinate σ. Since $\beta_y = \beta_z = 1$, evaluating the objective there gives the one-dimensional slice

$$g_d(\sigma) := \Phi\big(h(\sigma)\big) = \alpha_x(M-d) + \alpha_y\sigma + \alpha_z(d-\sigma) + \kappa\, r(\sigma), \qquad r(\sigma) = \sqrt{A + \sigma^2 + (d-\sigma)^2}.$$

Step 1 (two-point convexity of Φ). For $u, w \in \mathbb{R}^{\mathcal{I}}$ and $\theta \in [0,1]$,

$$\Phi\big((1-\theta)u + \theta w\big) \le (1-\theta)\Phi(u) + \theta\Phi(w).$$

Indeed the linear part is exactly additive, and $t \mapsto \sqrt{\sum_i \beta_i t_i^2}$ is the Euclidean norm of the vector with coordinates $\sqrt{\beta_i}\, t_i$, so it obeys the triangle inequality and positive homogeneity; multiplying by $\kappa \ge 0$ preserves the inequality.

Step 2 (the dichotomy from a relaxed minimizer). Let $t^{\mathrm{rel}} \in \Delta_M$ satisfy $\Phi(t^{\mathrm{rel}}) \le \Phi(s)$ for every $s \in \Delta_M$.

If $t^{\mathrm{rel}} \in K_d$, then since $K_d \subseteq \Delta_M$ the same inequality holds for every $s \in K_d$; this is the first alternative in both robustness regimes.

Assume from now on $t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} < d$. Because $t_y^{\mathrm{rel}}, t_z^{\mathrm{rel}} \ge 0$, this forces $d > 0$.

Step 3 (every feasible point is dominated on the face H_d). Let $s \in K_d$ and put $\Sigma = s_y + s_z \ge d$ and $\Sigma^{\mathrm{rel}} = t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} < d$. If $\Sigma = d$ then s already lies on H_d: its first coordinate is $M-d$ by the mass constraint, so $s = h(s_y)$ with $0 \le s_y \le d$, and we may take $\sigma = s_y$. If $\Sigma > d$, set

$$\theta = \frac{\Sigma - d}{\Sigma - \Sigma^{\mathrm{rel}}} \in [0,1], \qquad t^{\mathrm{b}} = (1-\theta)s + \theta\, t^{\mathrm{rel}},$$

which is well defined since $\Sigma - \Sigma^{\mathrm{rel}} > 0$. Then $t^{\mathrm{b}} \in \Delta_M$ (a convex combination of simplex points) and

$$t_y^{\mathrm{b}} + t_z^{\mathrm{b}} = (1-\theta)\Sigma + \theta\Sigma^{\mathrm{rel}} = \Sigma - \theta(\Sigma - \Sigma^{\mathrm{rel}}) = d,$$

so $t^{\mathrm{b}} = h(\sigma)$ with $\sigma = t_y^{\mathrm{b}} \in [0,d]$. By Step 1 and $\Phi(t^{\mathrm{rel}}) \le \Phi(s)$,

$$\Phi\big(h(\sigma)\big) = \Phi(t^{\mathrm{b}}) \le (1-\theta)\Phi(s) + \theta\Phi(t^{\mathrm{rel}}) \le \Phi(s).$$

In either case there is $\sigma \in [0,d]$ with $\Phi(h(\sigma)) \le \Phi(s)$.

Step 4 (the selector lies in $[0,d]$*).* In the two endpoint branches $s^\star \in \{0,d\} \subseteq [0,d]$ because $d > 0$. In the interior branch both guards fail,

$$|\delta| < \frac{\kappa d}{\sqrt{A+d^2}}, \qquad \text{equivalently} \qquad \delta^2(A+d^2) < \kappa^2 d^2,$$

which forces $\kappa > 0$ and, since $\delta^2 d^2 \le \delta^2(A+d^2)$, also $\delta^2 < \kappa^2$, so $\kappa^2 - \delta^2/2 > 0$ and

$$W := \sqrt{\frac{A+d^2/2}{\kappa^2-\delta^2/2}}$$

is well defined with $s^\star = (d-\delta W)/2$. Moreover

$$\delta^2 W^2 = \frac{2\delta^2 A + \delta^2 d^2}{2\kappa^2-\delta^2} < d^2 \iff \delta^2(A+d^2) < \kappa^2 d^2,$$

so $|\delta W| < d$ and hence $0 < s^\star < d$. In all branches $h(s^\star) = (M-d, s^\star, d-s^\star)$ has nonnegative coordinates (using $d \le M$), total mass M, and $s^\star + (d-s^\star) = d$; therefore

$$t^\star = h(s^\star) \in K_d.$$

Step 5 (the selector's sign condition). We claim that for every $\sigma \in [0,d]$,

$$0 \le (\sigma - s^\star)\Big(r(s^\star)\,\delta + \kappa\,(2s^\star - d)\Big).$$

If $s^\star = 0$, the branch guard reads $\delta\sqrt{A+d^2} \ge \kappa d$, and $r(0) = \sqrt{A+d^2}$, so the bracket $\delta\sqrt{A+d^2} - \kappa d$ is nonnegative while $\sigma - 0 \ge 0$. If $s^\star = d$, the guard reads $\delta\sqrt{A+d^2} \le -\kappa d$, and $r(d) = \sqrt{A+d^2}$, so the bracket $\delta\sqrt{A+d^2} + \kappa d$ is nonpositive while $\sigma - d \le 0$. In the interior branch the bracket vanishes: with W as in Step 4 we have $2s^\star - d = -\delta W$ and

$$r(s^\star)^2 = A + (s^\star)^2 + (d-s^\star)^2 = A + \frac{d^2}{2} + \frac{\delta^2 W^2}{2} = \Big(\kappa^2 - \frac{\delta^2}{2}\Big)W^2 + \frac{\delta^2 W^2}{2} = (\kappa W)^2,$$

using $W^2(\kappa^2 - \delta^2/2) = A + d^2/2$; both $r(s^\star)$ and κW are nonnegative, so $r(s^\star) = \kappa W$ and

$$r(s^\star)\,\delta + \kappa(2s^\star - d) = \kappa W\delta - \kappa\delta W = 0.$$

Step 6 (the selector minimizes the slice). Fix $\sigma \in [0,d]$. By Cauchy–Schwarz applied to the vectors $(\sqrt{A}, s^\star, d-s^\star)$ and $(\sqrt{A}, \sigma, d-\sigma)$,

$$A + s^\star\sigma + (d-s^\star)(d-\sigma) \le r(s^\star)\,r(\sigma).$$

Subtracting $r(s^\star)^2 = A + (s^\star)^2 + (d-s^\star)^2$ from both sides and using the algebraic identity

$$A + s^\star\sigma + (d-s^\star)(d-\sigma) - r(s^\star)^2 = (\sigma - s^\star)(2s^\star - d),$$

we obtain $(\sigma - s^\star)(2s^\star - d) \le r(s^\star)\big(r(\sigma) - r(s^\star)\big)$; multiplying by $\kappa \ge 0$ and adding $r(s^\star)\delta(\sigma - s^\star)$ to both sides gives, together with Step 5,

$$0 \le (\sigma - s^\star)\Big(r(s^\star)\delta + \kappa(2s^\star - d)\Big) \le r(s^\star)\Big(\delta(\sigma - s^\star) + \kappa\big(r(\sigma) - r(s^\star)\big)\Big).$$

Since $r(s^\star) > 0$ (because $d > 0$ makes $(s^\star)^2 + (d - s^\star)^2 > 0$),

$$\delta(\sigma - s^\star) + \kappa\big(r(\sigma) - r(s^\star)\big) \geq 0, \qquad \text{i.e.} \qquad g_d(s^\star) \leq g_d(\sigma),$$

because $g_d(\sigma) - g_d(s^\star) = \delta(\sigma - s^\star) + \kappa(r(\sigma) - r(s^\star))$.

Combining Steps 3 and 6: for every $s \in K_d$ there is $\sigma \in [0, d]$ with $\Phi(t^\star) = g_d(s^\star) \leq g_d(\sigma) = \Phi(h(\sigma)) \leq \Phi(s)$, so

$$t^\star = (M - d, s^\star, d - s^\star) \in K_d \qquad \text{and} \qquad \Phi(t^\star) \leq \Phi(s) \ \text{ for every } s \in K_d.$$

Step 7 (the two robustness regimes). For $\kappa > 0$, an admissible pair (S, λ) has active-set point t^{rel} lying in Δ_M and minimizing Φ over Δ_M by Lemma 6; Steps 2–6 then give exactly the two stated alternatives.

For $\kappa = 0$, a point t^{rel} of the exposed α-minimizing face minimizes the linear objective Φ over Δ_M, again by Lemma 6, so Steps 2–6 apply verbatim. Under the stated tie convention the first alternative is selected whenever the face meets K_d. Finally, at $\kappa = 0$ the two selector guards read $0 \leq \delta$ and $\delta \leq 0$, so the interior branch is never taken and

$$s^\star = \begin{cases} 0, & \alpha_y - \alpha_z \geq 0, \\ d, & \alpha_y - \alpha_z < 0, \end{cases}$$

which is the ordered endpoint rule. □

Lemma 7 reduces the constrained implementable calculation to a one-dimensional boundary segment whenever the relaxed optimizer violates the lower bound. The selector chooses either an endpoint of that segment or the unique interior point identified by the displayed expression. This is the source of the truncation correction in Theorem 4.

The final ingredient ties the reduced optimization problems back to the implementability loss defined in the main text.

Lemma 8. [`lem:rounding-gap-reduction`] (Reduced Gap Identity) *For any real r and any $\kappa \geq 0$, under Assumption 1, the implementability loss of Definition 7 satisfies*

$$\Delta_m^{\pm}(r, \kappa) = \inf_{\substack{(x,y,z_{\mathrm{sp}}) \in T_m \\ y + z_{\mathrm{sp}} \geq d_m}} \phi_{r,\kappa}(x, y, z_{\mathrm{sp}}) - \inf_{(x,y,z_{\mathrm{sp}}) \in T_m} \phi_{r,\kappa}(x, y, z_{\mathrm{sp}}).$$

Proof of Lemma 8. Set $q = 2(m-1)$; Assumption 1 gives $m \geq 2$, so $m \neq 0$ and $q > 0$. By Definition 7 it suffices to identify the two image sets

$$F_{r,\kappa}(\mathcal{E}_m^{\mathrm{blk}}) = \{F_{r,\kappa}(X) : X \in \mathcal{E}_m^{\mathrm{blk}}\}, \qquad F_{r,\kappa}(\mathcal{C}_m^{\pm}) = \{F_{r,\kappa}(X) : X \in \mathcal{C}_m^{\pm}\}$$

with the corresponding sets of reduced objective values, since an infimum depends only on the set of attained values.

Step 1 (the coordinate change is invertible on T_m). Given $(x, y, z_{\mathrm{sp}}) \in T_m$, put

$$u = 1 - x, \qquad v = \frac{z_{\mathrm{sp}} - y}{2m}.$$

Then $1 - u = x$ by construction. The simplex equation $qx + y + z_{\mathrm{sp}} = 2m$ gives $(m-1)x = m - \frac{1}{2}(y + z_{\mathrm{sp}})$, hence $1 + (m-1)u = m - (m-1)x = \frac{1}{2}(y + z_{\mathrm{sp}})$, and therefore

$$1 + (m-1)u - mv = \frac{y + z_{\mathrm{sp}}}{2} - \frac{z_{\mathrm{sp}} - y}{2} = y, \qquad 1 + (m-1)u + mv = \frac{y + z_{\mathrm{sp}}}{2} + \frac{z_{\mathrm{sp}} - y}{2} = z_{\mathrm{sp}}.$$

So every point of T_m is the spectral-coordinate image of the block matrix $X(u,v)$ with these parameters.

Step 2 (the relaxed image). We claim

$$\{F_{r,\kappa}(X) : X \in \mathcal{E}_m^{\mathrm{blk}}\} = \{\phi_{r,\kappa}(x,y,z_{\mathrm{sp}}) : (x,y,z_{\mathrm{sp}}) \in T_m\}.$$

For the inclusion $\subseteq$: every $X \in \mathcal{E}_m^{\mathrm{blk}}$ is of the form $X(u,v)$ by Definition 4, and Lemma 1 places its spectral coordinates in T_m and identifies $F_{r,\kappa}(X(u,v)) = \phi_{r,\kappa}(x,y,z_{\mathrm{sp}})$. For $\supseteq$: given $(x,y,z_{\mathrm{sp}}) \in T_m$, Step 1 produces (u,v) with these spectral coordinates, and Lemma 1 then gives both $X(u,v) \in \mathcal{E}_m^{\mathrm{blk}}$ and $F_{r,\kappa}(X(u,v)) = \phi_{r,\kappa}(x,y,z_{\mathrm{sp}})$.

Step 3 (the implementable image). We claim

$$\{F_{r,\kappa}(X) : X \in \mathcal{C}_m^{\pm}\} = \{\phi_{r,\kappa}(x,y,z_{\mathrm{sp}}) : (x,y,z_{\mathrm{sp}}) \in T_m,\ y + z_{\mathrm{sp}} \geq d_m\}.$$

For $\subseteq$: let $X = X(P)$ with $P \in \mathcal{P}_m^{\mathrm{sym}}$. Because P is invariant under the two-block automorphism group and $m \geq 2$, all within-community off-diagonal second moments coincide and all across-community ones coincide, so there are u, v with

$$X(P) = X(u,v);$$

Lemma 5 then gives $(x,y,z_{\mathrm{sp}}) \in T_m$ together with $y + z_{\mathrm{sp}} \geq d_m$, and Lemma 1 identifies $F_{r,\kappa}(X) = \phi_{r,\kappa}(x,y,z_{\mathrm{sp}})$. For $\supseteq$: given $(x,y,z_{\mathrm{sp}}) \in T_m$ with $y + z_{\mathrm{sp}} \geq d_m$, Step 1 produces (u,v) realizing these coordinates, Lemma 5 places $X(u,v) \in \mathcal{C}_m^{\pm}$, and Lemma 1 again matches the objective values.

Step 4 (conclusion). Substituting the two image identities of Steps 2 and 3 into Definition 7,

$$\Delta_m^{\pm}(r,\kappa) = \inf_{\substack{(x,y,z_{\mathrm{sp}}) \in T_m \\ y + z_{\mathrm{sp}} \geq d_m}} \phi_{r,\kappa}(x,y,z_{\mathrm{sp}}) - \inf_{(x,y,z_{\mathrm{sp}}) \in T_m} \phi_{r,\kappa}(x,y,z_{\mathrm{sp}}),$$

as claimed. □

Lemma 8 is the bridge from geometry back to design. The relaxed benchmark is the infimum over T_m, while the implementable benchmark is the same reduced objective with the parity truncation imposed. Consequently, the loss is exactly the value difference between these two scalar problems, which is the quantity evaluated by the active-set and truncation formulas above.

B Verification note

This appendix records the verification boundary for the mathematical claims used in the paper. The formal statements take the assignment law, its induced second moments, the two-block weighted graph, the block elliptope relaxation, and the finite comparison between relaxed and implementable covariance optima as their load-bearing objects. The design-based finite-population potential-outcome framework supplies the surrounding motivation.

The verified mathematical layer consists of the assumptions, definitions, lemmas, and theorems stated in the paper. In particular, it includes the two-block homophily, sign-symmetry, low-scale, and odd-block conditions in Assumption 1–Assumption 4; the graph, design-class, covariance, objective, loss, and block-sum definitions in Definition 1–Definition 8; the symmetry reduction in Proposition 1; the exactness results in Theorem 1 and Theorem 2; and the positive-loss and sharp-computation results in Theorem 3 and Theorem 4. The auxiliary reductions collected in Appendix A are part of the same mathematical layer.

The econometric motivation, related-work positioning, and interpretation are prose surrounding that layer. The displayed assumptions in the paper give exactly the conditions used to derive the formal conclusions, which are finite-dimensional and design based.

For reproducibility, the theorem statements and proofs are machine-checked in Lean 4. The assumed inputs to the checked arguments are the explicitly stated mathematical hypotheses, including the two-block homophily, sign-symmetry, low-scale, odd-block, and nonnegativity conditions where they appear; standard finite-dimensional real matrix primitives for traces, positive semidefiniteness, the Moore–Penrose pseudoinverse, and the Schatten–2 norm; and the usual classical real-analysis and finite-set principles available in the underlying library. The Lean verification discharges the algebraic, convex-geometric, spectral, parity, and active-set claims stated in the formal environments, while the historical, empirical, and interpretive material remains prose.

The formalization was checked against Lean toolchain `leanprover/lean4:v4.29.0-rc3` with `mathlib` at revision `bf8875c7`. All declarations live under the module namespace `CausalSmith.Experimentation.EXP_DesignPm1ExactnessBoundaryV1_Research`, and each carries a `@node` tag matching its manuscript label; every stated result is `sorry`-free. The five main results map to Lean declarations as follows.

Manuscript label	**Lean file**	**Declaration**
Proposition 1	`Helpers/SymmetryReduction.lean`	`symmetry_reduction`
Theorem 1	`Tcut.lean`	`cut_corner_exactness`
Theorem 2	`Trobust.lean`	`robust_corner_exactness`
Theorem 3	`Tgap.lean`	`gap_window`
Theorem 4	`Tsharp.lean`	`sharp_rho_star`

The auxiliary lemmas of Appendix A are declared in the sibling `Helpers/` files under the same namespace, with declaration names matching their manuscript labels (for example `block_spectral_coordinates`, `weighted_simplex_active_set`, `weighted_simplex_truncation`, and `rounding_gap_reduction`).

C Proofs of the main results

Proof of Proposition 1. Write $n = 2m$ and, for a matrix $X \in \mathbb{R}^{n\times n}$, introduce the two off-diagonal block sums and the two pair counts

$$S_{\mathrm{w}}(X) = \sum_{\substack{i\neq j \\ i,j \text{ same community}}} X_{ij}, \qquad S_{\mathrm{a}}(X) = \sum_{\substack{i,j \\ \text{opposite communities}}} X_{ij},$$

$$N_{\mathrm{w}} = 2m(m-1), \qquad N_{\mathrm{a}} = 2m^2,$$

the counts being the numbers of ordered index pairs of each kind. Call $G(d, w, c)$ the *block-constant* matrix with d on the diagonal, w on within-community off-diagonal entries, and c on across-community entries. Assumption 1 gives $m \geq 2$, so $N_{\mathrm{w}}, N_{\mathrm{a}} > 0$.

Step 1 (a master trace identity). If X is symmetric, then grouping the n^2 ordered pairs into diagonal, within-community, and across-community pairs gives

$$\operatorname{tr}\big(G(d, w, c)\, X\big) = d \sum_i X_{ii} + w\, S_{\mathrm{w}}(X) + c\, S_{\mathrm{a}}(X).$$

The three matrices L_m, $L_m^\dagger$ and J_n are block-constant: L_m has the constant weighted degree $(m-1)a/m+b$ on its diagonal and entries $-a/m$, $-b/m$ off it; $J_n = G(1,1,1)$; and $L_m^\dagger$ is a linear combination of the identity and the two rank-one projections onto the all-ones and community-sign directions, all three of which are block-constant. Consequently, if X and Y are symmetric with unit diagonal and

$$S_{\mathrm{w}}(Y) = S_{\mathrm{w}}(X), \qquad S_{\mathrm{a}}(Y) = S_{\mathrm{a}}(X),$$

then $\mathrm{tr}(L_m Y) = \mathrm{tr}(L_m X)$, $\mathrm{tr}(L_m^\dagger Y) = \mathrm{tr}(L_m^\dagger X)$ and $\mathrm{tr}(J_n Y) = \mathrm{tr}(J_n X)$; if in addition $\|Y\|_{S_2} \le \|X\|_{S_2}$, then, since $\kappa \ge 0$,

$$F_{r,\kappa}(Y) \le F_{r,\kappa}(X).$$

This comparison is the common engine of the two symmetrization claims.

Step 2 (relaxed matrices: the block average is feasible). Let X be positive semidefinite with $X_{ii} = 1$ for all i, and set

$$u = \frac{S_{\mathrm{w}}(X)}{N_{\mathrm{w}}}, \qquad v = \frac{S_{\mathrm{a}}(X)}{N_{\mathrm{a}}},$$

the averages of the within- and across-community off-diagonal entries. By construction $X(u,v)$ is symmetric with unit diagonal and

$$S_{\mathrm{w}}\big(X(u,v)\big) = N_{\mathrm{w}} u = S_{\mathrm{w}}(X), \qquad S_{\mathrm{a}}\big(X(u,v)\big) = N_{\mathrm{a}} v = S_{\mathrm{a}}(X).$$

Membership in $\mathcal{E}_m^{\mathrm{blk}}$ requires the three inequalities of Definition 4. First, testing positive semidefiniteness against $e_i - e_j$ and $e_i + e_j$ and using the unit diagonal gives

$$0 \le (e_i - e_j)^\top X (e_i - e_j) = 2 - 2X_{ij}, \qquad 0 \le (e_i + e_j)^\top X (e_i + e_j) = 2 + 2X_{ij},$$

so $|X_{ij}| \le 1$ for all i,j. Averaging the bound $X_{ij} \le 1$ over the N_{w} within-community off-diagonal pairs gives $u \le 1$, i.e. $1 - u \ge 0$. Second, with s the community sign vector ($+1$ on A_m, -1 on B_m) the sign product $s_i s_j$ is $+1$ on diagonal and within-community pairs and -1 on across-community pairs, so

$$0 \le s^\top X s = \sum_i X_{ii} + S_{\mathrm{w}}(X) - S_{\mathrm{a}}(X) = 2m + 2m(m-1)u - 2m^2 v = 2m\big(1 + (m-1)u - mv\big),$$

and dividing by $2m > 0$ gives $1 + (m-1)u - mv \ge 0$. Third, testing against the all-ones vector, whose quadratic form is the total entry sum,

$$0 \le \sum_{i,j} X_{ij} = \sum_i X_{ii} + S_{\mathrm{w}}(X) + S_{\mathrm{a}}(X) = 2m\big(1 + (m-1)u + mv\big),$$

giving $1 + (m-1)u + mv \ge 0$. With Assumption 1 this is exactly $X(u,v) \in \mathcal{E}_m^{\mathrm{blk}}$.

Step 3 (relaxed matrices: the objective does not increase). The block sums already match by Step 2, so by Step 1 only the Frobenius term has to be compared. Squaring entrywise and grouping pairs as before,

$$\|X\|_{S_2}^2 = 2m + \sum_{\text{within}} X_{ij}^2 + \sum_{\text{across}} X_{ij}^2, \qquad \|X(u,v)\|_{S_2}^2 = 2m + N_{\mathrm{w}} u^2 + N_{\mathrm{a}} v^2.$$

Cauchy–Schwarz over each of the two pair families gives $\left(\sum_{\text{within}} X_{ij}\right)^2 \le N_{\mathrm{w}} \sum_{\text{within}} X_{ij}^2$, i.e.

$$N_{\mathrm{w}} u^2 \le \sum_{\text{within}} X_{ij}^2, \qquad \text{and likewise} \qquad N_{\mathrm{a}} v^2 \le \sum_{\text{across}} X_{ij}^2.$$

Hence $\|X(u,v)\|_{S_2} \le \|X\|_{S_2}$, and Step 1 yields

$$F_{r,\kappa}\big(X(u,v)\big) \le F_{r,\kappa}(X).$$

Step 4 (assignment laws: the orbit average). Let H be the finite, nonempty group of two-block automorphisms: permutations σ of the units that either preserve both communities setwise or swap them. For $P \in \mathcal{P}_m^{\mathrm{bal}}$ define the orbit average

$$P^{\mathrm{orb}}(Z = z) = \frac{1}{|H|} \sum_{\sigma \in H} P\big(Z = z \circ \sigma\big).$$

It is a probability law, it is sign-symmetric because each summand is, and it is H-invariant because precomposing z with $\tau \in H$ permutes the summation over H; hence $P^{\mathrm{orb}} \in \mathcal{P}_m^{\mathrm{sym}}$. Sign symmetry gives zero one-point margins: replacing z by $-z$ leaves the law unchanged and flips the sign of Z_i, so $\mathbb{E}[Z_i] = -\mathbb{E}[Z_i]$ and therefore

$$\mathbb{E}[Z_i] = 0 \qquad \text{for every } i.$$

Step 5 (assignment laws: the objective does not increase). The same change of variables in the expectation gives the entrywise orbit-averaging identity

$$X\big(P^{\mathrm{orb}}\big)_{ij} = \frac{1}{|H|} \sum_{\sigma \in H} X(P)_{\sigma(i)\sigma(j)}.$$

Every $\sigma \in H$ maps within-community ordered pairs to within-community ordered pairs and across-community pairs to across-community pairs, and does so bijectively; hence summing the display over each family gives

$$S_{\mathrm{w}}\big(X(P^{\mathrm{orb}})\big) = S_{\mathrm{w}}\big(X(P)\big), \qquad S_{\mathrm{a}}\big(X(P^{\mathrm{orb}})\big) = S_{\mathrm{a}}\big(X(P)\big).$$

For the Frobenius term, Cauchy–Schwarz applied to the average over H gives entrywise

$$X\big(P^{\mathrm{orb}}\big)_{ij}^2 \le \frac{1}{|H|} \sum_{\sigma \in H} X(P)_{\sigma(i)\sigma(j)}^2,$$

and summing over all (i,j), each inner sum being a permutation of the full square sum, yields

$$\big\|X(P^{\mathrm{orb}})\big\|_{S_2} \le \big\|X(P)\big\|_{S_2}.$$

Both second moments are symmetric with unit diagonal ($X(P)_{ii} = \mathbb{E}[Z_i^2] = 1$), so Step 1 applies and gives

$$F_{r,\kappa}\big(X(P^{\mathrm{orb}})\big) \le F_{r,\kappa}\big(X(P)\big).$$

Step 6 (an infimum-reduction principle). Let f be a real function on a set $\mathcal{S}$, let $\mathcal{T} \subseteq \mathcal{S}$ be nonempty, suppose f is bounded below on $\mathcal{S}$, and suppose every $x \in \mathcal{S}$ admits $y \in \mathcal{T}$ with $f(y) \le f(x)$. Then

$$\inf_{\mathcal{S}} f = \inf_{\mathcal{T}} f.$$

Indeed $\inf_{\mathcal{S}} f \le \inf_{\mathcal{T}} f$ because $\mathcal{T} \subseteq \mathcal{S}$, and conversely each value $f(x)$, $x \in \mathcal{S}$, dominates some $f(y) \ge \inf_{\mathcal{T}} f$.

Step 7 (the two infimum identities). On the relaxed side take

$$\mathcal{S}_{\text{rel}} = \{X : X \succeq 0,\ X_{ii} = 1\ \forall i\}, \qquad \mathcal{T}_{\text{blk}} = \mathcal{E}_m^{\text{blk}}.$$

The inclusion $\mathcal{T}_{\text{blk}} \subseteq \mathcal{S}_{\text{rel}}$ follows from the projection decomposition used for block matrices. Let $P_{\mathbf{1}} = J_n/(2m)$, let $P_s = ss^\top/(2m)$, and let $P_\perp = I_n - P_{\mathbf{1}} - P_s$. For every u, v,

$$X(u,v) = (1-u)P_\perp + \big(1 + (m-1)u - mv\big)P_s + \big(1 + (m-1)u + mv\big)P_{\mathbf{1}}.$$

The matrices $P_{\mathbf{1}}$ and P_s are positive semidefinite rank-one projections. Also, for any vector ξ, if

$$A = \sum_{i \in A_m} \xi_i, \qquad B = \sum_{i \in B_m} \xi_i, \qquad Q_A = \sum_{i \in A_m} \xi_i^2, \qquad Q_B = \sum_{i \in B_m} \xi_i^2,$$

then

$$\xi^\top P_\perp \xi = Q_A + Q_B - \frac{A^2 + B^2}{m} \ge 0$$

by Cauchy–Schwarz on the two communities, so $P_\perp$ is positive semidefinite. Thus the three nonnegative coordinates in Definition 4 make $X(u,v)$ positive semidefinite, and the definition of $X(u,v)$ gives unit diagonal. The set $\mathcal{T}_{\text{blk}}$ is nonempty because $I_n = X(0,0) \in \mathcal{E}_m^{\text{blk}}$.

It remains to record the lower bound on the relaxed objective. Choose block-constant representations

$$L_m = G(d_L, w_L, c_L), \qquad L_m^\dagger = G(d_\dagger, w_\dagger, c_\dagger), \qquad J_n = G(1,1,1),$$

and set

$$B_L = |d_L|\, 2m + |w_L| N_{\text{w}} + |c_L| N_{\text{a}}, \qquad B_\dagger = |d_\dagger|\, 2m + |w_\dagger| N_{\text{w}} + |c_\dagger| N_{\text{a}},$$
$$B_J = 2m + N_{\text{w}} + N_{\text{a}}.$$

For $X \in \mathcal{S}_{\text{rel}}$, Step 2 gives $|X_{ij}| \le 1$; hence $|S_{\text{w}}(X)| \le N_{\text{w}}$ and $|S_{\text{a}}(X)| \le N_{\text{a}}$. The trace identity in Step 1 then gives, for every block-constant $G(d,w,c)$,

$$\big|\operatorname{tr}(G(d,w,c)X)\big| \le |d|\, 2m + |w| N_{\text{w}} + |c| N_{\text{a}}.$$

Therefore

$$F_{r,\kappa}(X) \ge -B_L - |r| B_\dagger + \kappa \|X\|_{S_2} - B_J \ge -\big(B_L + |r| B_\dagger + B_J\big),$$

because $\kappa \ge 0$ and $\|X\|_{S_2} \ge 0$. Steps 2–3 supply, for every $X \in \mathcal{S}_{\text{rel}}$, the point $X(u,v) \in \mathcal{T}_{\text{blk}}$ with no larger objective. Step 6 gives

$$\inf_{\substack{X \succeq 0 \\ X_{ii} = 1\ \forall i}} F_{r,\kappa}(X) = \inf_{X \in \mathcal{E}_m^{\text{blk}}} F_{r,\kappa}(X).$$

On the implementable side take

$$\mathcal{S}_{\text{bal}} = \mathcal{P}_m^{\text{bal}}, \qquad \mathcal{T}_{\text{sym}} = \mathcal{P}_m^{\text{sym}}, \qquad f(P) = F_{r,\kappa}\big(X(P)\big).$$

Then $\mathcal{T}_{\text{sym}} \subseteq \mathcal{S}_{\text{bal}}$ by Definition 3; $\mathcal{T}_{\text{sym}} \ne \emptyset$ because the iid fair-sign law is block-exchangeable; and f is bounded below because every second moment $X(P) = \mathbb{E}_P[ZZ^\top]$ is positive semidefinite with unit diagonal, hence belongs to $\mathcal{S}_{\text{rel}}$ and satisfies the lower bound just displayed. Steps 4–5 supply, for every $P \in \mathcal{S}_{\text{bal}}$, the law $P^{\text{orb}} \in \mathcal{T}_{\text{sym}}$ with no larger objective. Step 6 gives

$$\inf_{P \in \mathcal{P}_m^{\text{bal}}} F_{r,\kappa}\big(X(P)\big) = \inf_{P \in \mathcal{P}_m^{\text{sym}}} F_{r,\kappa}\big(X(P)\big).$$

□

Proof of Theorem 1. 1. By the two-block homophily assumption, $m \geq 2$, $0 < b < a$. Hence

$$q = 2(m-1) > 0, \qquad a - b > 0, \qquad a + b > 0, \qquad 2b > 0.$$

Set

$$A = 2b(a+b)\left(1 - \frac{\kappa}{a-b}\right), \qquad B = 2b(2m - 2b - \kappa).$$

From $r < \max\{0, \min(A, B)\}$ and $0 \leq r$, we first get

$$r < \min(A, B).$$

Indeed, if $\min(A, B) \leq 0$, then the outer maximum is 0, contradicting $0 \leq r < 0$. Otherwise the outer maximum is $\min(A, B)$. Thus

$$r < A, \qquad r < B.$$

2. We now convert these two inequalities into the coefficient separations needed at the cut vertex. From $r < A$, division by the positive number $2b(a+b)$ gives

$$\frac{r}{2b(a+b)} < 1 - \frac{\kappa}{a-b}.$$

Multiplying by $a - b > 0$ and rearranging gives

$$\kappa < (a-b)\left(1 - \frac{r}{2b(a+b)}\right).$$

Using

$$(a+b) + \frac{r}{a+b} - \left(2b + \frac{r}{2b} + \kappa\right) = (a-b)\left(1 - \frac{r}{2b(a+b)}\right) - \kappa$$

and $c_x/q = (a+b) + r/(a+b)$, $c_y = 2b + r/(2b)$, we obtain

$$\frac{c_x}{q} > c_y + \kappa.$$

Similarly, from $r < B$, division by $2b > 0$ gives

$$\frac{r}{2b} < 2m - 2b - \kappa,$$

and hence

$$2b + \frac{r}{2b} + \kappa < 2m.$$

Since $c_z = 2m$, this is

$$c_z > c_y + \kappa.$$

3. The hypotheses of Lemma 2 are therefore satisfied: $q > 0$, $\kappa \geq 0$, and the two coefficient separations above hold. Hence
$$(0, 2m, 0) \in T_m,$$
and every other point of T_m has strictly larger reduced objective value.

   By Lemma 1, the cut covariance is
$$X_{\mathrm{cut}} = X(1, -1),$$
and its spectral coordinates are
$$(1 - 1,\ 1 + (m-1)\cdot 1 - m(-1),\ 1 + (m-1)\cdot 1 + m(-1)) = (0, 2m, 0).$$
Thus $X_{\mathrm{cut}} \in \mathcal{E}_m^{\mathrm{blk}}$, and
$$F_{r,\kappa}(X_{\mathrm{cut}}) = \phi_{r,\kappa}(0, 2m, 0).$$

4. Let $X \in \mathcal{E}_m^{\mathrm{blk}}$ and suppose $X \neq X_{\mathrm{cut}}$. Write $X = X(u, v)$. By Lemma 1, the associated coordinates
$$x = 1 - u, \qquad y = 1 + (m-1)u - mv, \qquad z_{\mathrm{sp}} = 1 + (m-1)u + mv$$
belong to T_m, and
$$F_{r,\kappa}(X) = \phi_{r,\kappa}(x, y, z_{\mathrm{sp}}).$$
These coordinates cannot equal $(0, 2m, 0)$: if $x = 0$, then $u = 1$, and if also $z_{\mathrm{sp}} = 0$, then
$$0 = 1 + (m-1)\cdot 1 + mv = m(1 + v),$$
so $v = -1$ since $m > 0$, giving $X = X_{\mathrm{cut}}$, a contradiction. The strict part of Lemma 2 therefore gives
$$F_{r,\kappa}(X_{\mathrm{cut}}) = \phi_{r,\kappa}(0, 2m, 0) < \phi_{r,\kappa}(x, y, z_{\mathrm{sp}}) = F_{r,\kappa}(X).$$
Thus X_{cut} is the unique minimizer over $\mathcal{E}_m^{\mathrm{blk}}$.

5. The cut design P_{cut} puts mass $1/2$ on the cut assignment z^+ ($+1$ on A_m, -1 on B_m) and mass $1/2$ on its global sign reversal $z^- = -z^+$. It is therefore sign-symmetric, i.e. $P_{\mathrm{cut}}(Z = z) = P_{\mathrm{cut}}(Z = -z)$ for all z. Its two-point support is invariant under the two-block automorphism group: a permutation preserving both communities fixes each of z^+ and z^-, and the community swap exchanges them. Hence
$$P_{\mathrm{cut}} \in \mathcal{P}_m^{\mathrm{sym}}.$$
Moreover $z_i^+ z_j^+ = z_i^- z_j^-$ for all i, j, so both atoms contribute the same product and
$$X(P_{\mathrm{cut}})_{ij} = \mathbb{E}\big[Z_i Z_j\big] = z_i^+ z_j^+ = (X_{\mathrm{cut}})_{ij}, \qquad \text{i.e.} \qquad X(P_{\mathrm{cut}}) = X_{\mathrm{cut}}.$$
Consequently $X_{\mathrm{cut}} \in \mathcal{C}_m^{\pm}$.

   Next, every $X \in \mathcal{C}_m^{\pm}$ is the second moment of a design $P \in \mathcal{P}_m^{\mathrm{sym}}$. Because P is invariant under the two-block automorphism group and $m \geq 2$, all within-community off-diagonal second moments coincide and all across-community ones coincide, so
$$X = X(u, v) \qquad \text{for some } u, v,$$

and Lemma 5 places its spectral coordinates in the reduced triangle T_m. By Lemma 1 this membership is exactly block-elliptope membership, so

$$\mathcal{C}_m^{\pm} \subseteq \mathcal{E}_m^{\mathrm{blk}}.$$

Therefore, if $X \in \mathcal{C}_m^{\pm}$ and $X \neq X_{\mathrm{cut}}$, the strict relaxed optimality already proved yields

$$F_{r,\kappa}(X_{\mathrm{cut}}) < F_{r,\kappa}(X).$$

Thus X_{cut} is also the unique minimizer over $\mathcal{C}_m^{\pm}$.

6. It remains to identify the gap. From the strict inequalities above, with equality allowed only at X_{cut}, we have the non-strict minimizing inequalities

$$F_{r,\kappa}(X_{\mathrm{cut}}) \le F_{r,\kappa}(X) \quad \text{for all } X \in \mathcal{E}_m^{\mathrm{blk}},$$

and

$$F_{r,\kappa}(X_{\mathrm{cut}}) \le F_{r,\kappa}(X) \quad \text{for all } X \in \mathcal{C}_m^{\pm}.$$

Since X_{cut} belongs to both feasible classes, each of the two infima is attained at X_{cut}:

$$\inf_{X \in \mathcal{E}_m^{\mathrm{blk}}} F_{r,\kappa}(X) = F_{r,\kappa}(X_{\mathrm{cut}}) = \inf_{X \in \mathcal{C}_m^{\pm}} F_{r,\kappa}(X).$$

Hence, by the definition of the implementability gap in Definition 7,

$$\Delta_m^{\pm}(r,\kappa) = F_{r,\kappa}(X_{\mathrm{cut}}) - F_{r,\kappa}(X_{\mathrm{cut}}) = 0.$$

This proves all asserted conclusions.

$\square$

Proof of Theorem 2. 1. We first prove the finite-weight equivalence. Write

$$q = 2(m-1), \qquad c_x = q\left((a+b) + \frac{r}{a+b}\right), \qquad c_y = 2b + \frac{r}{2b}, \qquad c_z = 2m.$$

Suppose that some $\kappa \ge 0$ makes I_n a relaxed minimizer over $\mathcal{E}_m^{\mathrm{blk}}$. For an arbitrary $(x, y, z_{\mathrm{sp}}) \in T_m$, set

$$u = 1 - x, \qquad v = \frac{z_{\mathrm{sp}} - y}{2m}.$$

The simplex equation $qx + y + z_{\mathrm{sp}} = 2m$ gives $1 + (m-1)u = \frac{1}{2}(y + z_{\mathrm{sp}})$, so these formulas recover

$$1 - u = x, \qquad 1 + (m-1)u - mv = y, \qquad 1 + (m-1)u + mv = z_{\mathrm{sp}}.$$

Using Lemma 1, this inverse change of variables sends (x, y, z_{sp}) to a feasible block matrix $X(u,v)$, and the matrix objective equals the reduced objective at (x, y, z_{sp}). Since $I_n = X(0,0)$ corresponds to $(1,1,1)$, so that

$$F_{r,\kappa}(I_n) = \phi_{r,\kappa}(1,1,1),$$

the assumed matrix minimality says that $(1,1,1)$ minimizes

$$c_x x + c_y y + c_z z_{\mathrm{sp}} + \kappa\sqrt{qx^2 + y^2 + z_{\mathrm{sp}}^2}$$

on T_m. The certificate-necessity part of Lemma 4 therefore gives

$$\frac{c_x}{q} = c_y, \qquad c_y = c_z.$$

The first equality is

$$a + b + \frac{r}{a+b} = 2b + \frac{r}{2b},$$

and by the identity

$$\left(a + b + \frac{r}{a+b}\right) - \left(2b + \frac{r}{2b}\right) = (a-b)\left(1 - \frac{r}{2b(a+b)}\right)$$

this is equivalent to

$$(a-b)\left(1 - \frac{r}{2b(a+b)}\right) = 0.$$

By two-block homophily, $a > b > 0$, so $a - b \neq 0$ and the second factor vanishes:

$$r = 2b(a+b).$$

Substituting this into $c_y = c_z$ gives

$$2b + \frac{2b(a+b)}{2b} = 2m,$$

that is $a + 3b = 2m$. Thus the affine-balanced locus holds. Conversely, if

$$a + 3b = 2m, \qquad r = 2b(a+b),$$

then substituting $r = 2b(a+b)$ into the three coefficients gives

$$\frac{c_x}{q} = (a+b) + \frac{2b(a+b)}{a+b} = a + 3b, \qquad c_y = 2b + \frac{2b(a+b)}{2b} = a + 3b, \qquad c_z = 2m = a + 3b,$$

so $c_x/q = c_y = c_z$. Taking $\kappa = 1$, the strict optimality part of Lemma 4, transported through Lemma 1, gives $I_n \in \mathcal{E}_m^{\mathrm{blk}}$ and strict optimality against every $X \neq I_n$. For $X = I_n$ the desired inequality is equality, while for $X \neq I_n$ it follows from strict optimality. Hence a finite nonnegative robustness weight exists exactly on the stated locus.

2. On the affine-balanced locus, let $\kappa > 0$. As above,

$$\frac{c_x}{q} = c_y = c_z.$$

The point $I_n = X(0,0)$ has reduced coordinates $(1,1,1)$ by Lemma 1. The Frobenius-center certificate, Lemma 4, gives strict reduced optimality of $(1,1,1)$ on T_m. Translating back by Lemma 1, $I_n \in \mathcal{E}_m^{\mathrm{blk}}$, and for any $X = X(u,v) \in \mathcal{E}_m^{\mathrm{blk}}$ with $X \neq I_n$, the associated reduced coordinates are not $(1,1,1)$: if they were, then $x = 1$ gives $u = 0$, and then $z_{\mathrm{sp}} = 1$ reads $1 + mv = 1$, so $v = 0$ since $m > 0$, i.e. $X = I_n$. Therefore

$$F_{r,\kappa}(I_n) < F_{r,\kappa}(X).$$

3. The iid law assigns mass 2^{-n} to every sign vector, so it is invariant under sign flip, under within-block permutations, and under the block swap, whence

$$P_{\mathrm{iid}} \in \mathcal{P}_m^{\mathrm{sym}}.$$

Its second moments are diagonal: the diagonal entries are $\mathbb{E}[Z_i^2] = 1$, and for $i \neq j$ the involution $z \mapsto z'$ that flips only the i-th sign is a bijection of the assignment space preserving the uniform law and sending $Z_i Z_j$ to $-Z_i Z_j$, so

$$\mathbb{E}[Z_i Z_j] = -\mathbb{E}[Z_i Z_j] = 0 \qquad (i \neq j).$$

Thus all off-diagonal moments vanish and

$$X(P_{\mathrm{iid}}) = I_n.$$

4. Now assume the affine-balanced locus fails. If for some $\kappa > 0$ the matrix I_n were a relaxed minimizer, then the same center-coefficient necessity argument from the first step would give

$$\frac{c_x}{q} = c_y, \qquad c_y = c_z,$$

and the same algebra would force

$$a + 3b = 2m, \qquad r = 2b(a+b),$$

contradicting the assumption. Hence I_n is not a relaxed minimizer at any positive finite robustness weight off the locus.

5. Finally, let $(X_\kappa)_{\kappa \in \mathbb{R}}$ be any family of relaxed minimizers for positive κ. Since $r \geq 0$, $a > b > 0$ and $m \geq 2$, the reduced linear coefficients satisfy $c_x, c_y, c_z \geq 0$; and every point of T_m has nonnegative coordinates, so the linear part $c_x x + c_y y + c_z z_{\mathrm{sp}}$ is nonnegative on T_m. Put

$$L = c_x + c_y + c_z.$$

The reduced coordinates of I_n are $(1,1,1)$, whose weighted squared norm is $q \cdot 1 + 1 + 1 = q + 2 = 2m$; hence, by Lemma 1,

$$F_{r,\kappa}(I_n) = L + \kappa\sqrt{2m} \qquad \text{for every } \kappa.$$

Comparing each minimizer with the feasible point I_n gives

$$F_{r,\kappa}(X_\kappa) \leq F_{r,\kappa}(I_n) = L + \kappa\sqrt{2m}.$$

Writing $X_\kappa = X(u_\kappa, v_\kappa)$ and setting

$$N_{\mathrm{w}} = 2m(m-1) > 0, \qquad N_{\mathrm{a}} = 2m^2 > 0,$$

the reduced squared norm decomposes as

$$q(1-u_\kappa)^2 + \big(1 + (m-1)u_\kappa - mv_\kappa\big)^2 + \big(1 + (m-1)u_\kappa + mv_\kappa\big)^2 = 2m + N_{\mathrm{w}} u_\kappa^2 + N_{\mathrm{a}} v_\kappa^2.$$

Suppose some within-community off-diagonal entry stayed at distance at least $\varepsilon > 0$ from its target for arbitrarily large κ; that entry equals u_κ, so $|u_\kappa| \geq \varepsilon$ and the displayed squared norm is at least $2m + N_{\mathrm{w}}\varepsilon^2$. Setting

$$\delta = \sqrt{2m + N_{\mathrm{w}}\varepsilon^2} - \sqrt{2m} > 0,$$

the robustness term at X_κ is therefore at least $\kappa(\sqrt{2m} + \delta)$, and, the linear part being nonnegative,

$$\kappa\big(\sqrt{2m} + \delta\big) \leq F_{r,\kappa}(X_\kappa) \leq L + \kappa\sqrt{2m},$$

which forces $\kappa\delta \leq L$; this fails for every $\kappa > L/\delta$, a contradiction. Hence $u_\kappa \to 0$. The same argument with N_{a} in place of N_{w}, applied to an across-community entry (which equals v_κ), shows $v_\kappa \to 0$. Diagonal entries of X_κ are identically 1, matching I_n exactly, while off-diagonal entries are either u_κ or v_κ and the corresponding entries of I_n are 0. Hence every entry of X_κ converges to the corresponding entry of I_n as $\kappa \to \infty$.

□

Proof of Theorem 3. 1. Assumption 1 gives $m \geq 2$, $b < a$, and $b > 0$. Hence

$$q = 2(m-1) > 0, \qquad \sqrt{q} > 0, \qquad a + b > 0, \qquad a - b > 0.$$

Assumption 3 gives

$$2m - a - 3b > 0.$$

These inequalities are the standing sign facts used below for division and multiplication by positive quantities.

2. Assume now that Assumption 4 holds, and fix

$$0 \leq \kappa < \kappa_{\mathrm{gap}}(m,a,b) = \frac{(2m-a-3b)(a-b)\sqrt{q}}{a+b}.$$

Multiplying this strict inequality by the positive factor $(a+b)/((a-b)\sqrt{q})$,

$$\frac{\kappa(a+b)}{(a-b)\sqrt{q}} < 2m - a - 3b.$$

Next, from $1 + \dfrac{2b}{a-b} = \dfrac{a+b}{a-b}$ we get the identity

$$R_x^+ - R_x^- = (a+b)\Big(2m - a - b - \frac{\kappa}{\sqrt{q}}\Big) - 2b(a+b)\Big(1 + \frac{\kappa}{(a-b)\sqrt{q}}\Big) = (a+b)\left((2m-a-3b) - \frac{\kappa(a+b)}{(a-b)\sqrt{q}}\right),$$

whose right-hand side is positive by the previous display and $a + b > 0$. Thus $R_x^- < R_x^+ = r_{\mathrm{gap}}^+$, and since

$$r_{\mathrm{gap}}^- = \frac{R_x^- + R_x^+}{2}$$

is the midpoint of R_x^- and R_x^+,

$$R_x^- < r_{\mathrm{gap}}^- < r_{\mathrm{gap}}^+.$$

Also $R_x^- \geq 0$, because $2b(a+b) > 0$ and $\kappa/((a-b)\sqrt{q}) \geq 0$; and since $1 - \kappa/(a-b) \leq 1 \leq 1 + \kappa/((a-b)\sqrt{q})$ and $2b(a+b) > 0$,

$$2b(a+b)\left(1 - \frac{\kappa}{a-b}\right) \leq R_x^-,$$

so, the maximum of two quantities each at most R_x^- being at most R_x^-,

$$r_{\mathrm{cut,gap}}(m,a,b,\kappa) = \max\left\{0,\ 2b(a+b)\left(1-\frac{\kappa}{a-b}\right)\right\} \le R_x^-.$$

Combining these inequalities yields

$$r_{\mathrm{cut,gap}}(m,a,b,\kappa) < r_{\mathrm{gap}}^-(m,a,b,\kappa) < r_{\mathrm{gap}}^+(m,a,b,\kappa).$$

3. Let $r_{\mathrm{gap}}^- < r < r_{\mathrm{gap}}^+$. From the preceding step,

$$R_x^- < r.$$

After division by $2b(a+b) > 0$, this says

$$1 + \frac{\kappa}{(a-b)\sqrt{q}} < \frac{r}{2b(a+b)}.$$

Multiplying by $a-b > 0$ gives

$$\frac{\kappa}{\sqrt{q}} < (a-b)\left(\frac{r}{2b(a+b)} - 1\right).$$

Using the identity

$$2b + \frac{r}{2b} - \left(a+b+\frac{r}{a+b}\right) = (a-b)\left(\frac{r}{2b(a+b)} - 1\right),$$

and $c_x/q = a+b+r/(a+b)$, this proves

$$c_y > \frac{c_x}{q} + \frac{\kappa}{\sqrt{q}}.$$

Similarly, the upper bound $r < r_{\mathrm{gap}}^+ = R_x^+ = (a+b)\big(2m-a-b-\kappa/\sqrt{q}\big)$, divided by $a+b>0$, gives

$$\frac{r}{a+b} < 2m - a - b - \frac{\kappa}{\sqrt{q}},$$

hence, since $c_z = 2m$ and $c_x/q = a+b+r/(a+b)$,

$$c_z > \frac{c_x}{q} + \frac{\kappa}{\sqrt{q}}.$$

By Lemma 3, the point

$$\left(\frac{2m}{q}, 0, 0\right) = \left(\frac{m}{m-1}, 0, 0\right)$$

belongs to T_m and is the strict unique minimizer of the reduced objective on T_m.

4. The spread covariance is $X(-1/(m-1), 0)$, whose spectral coordinates are

$$x = 1 - \left(-\frac{1}{m-1}\right) = \frac{m}{m-1} = \frac{2m}{q}, \qquad y = 1 - \frac{m-1}{m-1} = 0, \qquad z_{\mathrm{sp}} = 0.$$

By Lemma 1, the membership of this triple in T_m gives

$$X_{\mathrm{spread}} \in \mathcal{E}_m^{\mathrm{blk}},$$

and the value of $F_{r,\kappa}$ at X_{spread} is the reduced objective at $(2m/q, 0, 0)$. If $X \in \mathcal{E}_m^{\mathrm{blk}}$, write $X = X(u, v)$. Lemma 1 sends it to a point of T_m and identifies $F_{r,\kappa}(X)$ with the reduced objective there. If this reduced point equals $(2m/q, 0, 0)$, then $y = z_{\mathrm{sp}} = 0$, so

$$z_{\mathrm{sp}} - y = 2mv = 0 \quad \text{and} \quad y + z_{\mathrm{sp}} = 2 + 2(m-1)u = 0,$$

forcing $v = 0$ and $u = -1/(m-1)$, hence $X = X_{\mathrm{spread}}$. Therefore every $X \neq X_{\mathrm{spread}}$ corresponds to a different point of T_m, and the strict reduced optimality just proved gives

$$F_{r,\kappa}(X_{\mathrm{spread}}) < F_{r,\kappa}(X).$$

5. Since m is odd, Lemma 5 gives $d_m = 2/m > 0$ for the implementable reduced slice. The spread point has

$$y + z_{\mathrm{sp}} = 0 < d_m,$$

so

$$X_{\mathrm{spread}} \notin \mathcal{C}_m^{\pm}.$$

For the loss, pass from $(x, y, z_{\mathrm{sp}}) \in T_m$ to the simplex variable

$$t = (qx,\ y,\ z_{\mathrm{sp}}), \qquad t_x + t_y + t_z = qx + y + z_{\mathrm{sp}} = 2m = M,$$

with inverse $(t_x, t_y, t_z) \mapsto (t_x/q, t_y, t_z)$. With the coefficients and weights

$$\alpha = \left(\frac{c_x}{q}, c_y, c_z\right), \qquad \beta = \left(\frac{1}{q}, 1, 1\right),$$

this change of variables matches the two objectives exactly:

$$\sum_i \alpha_i t_i = \frac{c_x}{q}\, qx + c_y y + c_z z_{\mathrm{sp}} = c_x x + c_y y + c_z z_{\mathrm{sp}}, \qquad \sum_i \beta_i t_i^2 = \frac{(qx)^2}{q} + y^2 + z_{\mathrm{sp}}^2 = qx^2 + y^2 + z_{\mathrm{sp}}^2,$$

so $\Phi_{\alpha,\beta,\kappa}(t) = \phi_{r,\kappa}(x, y, z_{\mathrm{sp}})$, and the parity constraint reads

$$t_y + t_z \geq d_m.$$

Under this dictionary the relaxed minimizer displayed above is

$$t^{\mathrm{rel}} = (2m, 0, 0), \qquad t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} = 0 < d_m.$$

It is in Δ_M and minimizes $\Phi_{\alpha,\beta,\kappa}$ over Δ_M, because the corresponding reduced point minimizes $\phi_{r,\kappa}$ over T_m. The data for the relaxed-minimizer truncation argument are therefore in place: $d_m \leq M = 2m$, $\beta_x = 1/q \geq 0$, $\beta_y = \beta_z = 1$, $\kappa \geq 0$, and t^{rel} is a global minimizer on Δ_M. That argument says the following. If a relaxed minimizer satisfies $t_y + t_z \geq d_m$, it remains minimizing after truncation; if it satisfies $t_y + t_z < d_m$, then the boundary-selector point

$$t^\star = (M - d_m,\ s^\star,\ d_m - s^\star), \qquad s^\star = \mathrm{truncSelector}(M, d_m, \alpha, \beta, \kappa),$$

belongs to $K_{d_m} = \{t \in \Delta_M : t_y + t_z \geq d_m\}$ and minimizes $\Phi_{\alpha,\beta,\kappa}$ over K_{d_m}. Indeed, the feasible case is immediate from relaxed optimality; in the infeasible case, every truncated feasible point is reduced to a point on the boundary segment $(M - d_m, s, d_m - s)$, and the selector $s^\star$ minimizes the resulting one-dimensional objective on $[0, d_m]$.

Since $t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} < d_m$, the second alternative supplies such a minimizer $t^\star$. Translating back,

$$(x^\star, y^\star, z_{\mathrm{sp}}^\star) = \Big(\frac{t_x^\star}{q}, t_y^\star, t_z^\star\Big) \in T_m$$

satisfies $y^\star + z_{\mathrm{sp}}^\star \geq d_m$ and minimizes $\phi_{r,\kappa}$ over the parity-truncated slice. In particular both infima in Lemma 8 are attained:

$$\inf_{(x,y,z_{\mathrm{sp}})\in T_m} \phi_{r,\kappa} = \phi_{r,\kappa}\Big(\frac{2m}{q}, 0, 0\Big), \qquad \inf_{\substack{(x,y,z_{\mathrm{sp}})\in T_m \\ y+z_{\mathrm{sp}}\geq d_m}} \phi_{r,\kappa} = \phi_{r,\kappa}(x^\star, y^\star, z_{\mathrm{sp}}^\star).$$

The truncated minimizer is feasible for the relaxed problem but differs from the relaxed minimizer, because $y^\star + z_{\mathrm{sp}}^\star \geq d_m > 0$ while the relaxed minimizer has $y + z_{\mathrm{sp}} = 0$; the strict reduced uniqueness above therefore gives

$$\phi_{r,\kappa}\Big(\frac{2m}{q}, 0, 0\Big) < \phi_{r,\kappa}(x^\star, y^\star, z_{\mathrm{sp}}^\star).$$

Finally Lemma 8 identifies $\Delta_m^{\pm}(r,\kappa)$ with the truncated reduced infimum minus the relaxed reduced infimum, so

$$\Delta_m^{\pm}(r,\kappa) > 0.$$

6. It remains to prove the even case. Assume m is even and let $X \in \mathcal{E}_m^{\mathrm{blk}}$. Write $X = X(u,v)$, and let (x, y, z_{sp}) be its spectral coordinates. By Lemma 1, this point lies in T_m, so $y, z_{\mathrm{sp}} \geq 0$. For even m, the parity threshold in Lemma 5 is $d_m = 0$, and hence $y + z_{\mathrm{sp}} \geq 0$ automatically. The same lemma therefore gives $X \in \mathcal{C}_m^{\pm}$. Thus

$$\mathcal{E}_m^{\mathrm{blk}} \subseteq \mathcal{C}_m^{\pm}.$$

□

Proof of Theorem 4. Throughout Assumption 1 gives $m \geq 2$, hence $q = 2(m-1) > 0$, $a+b > 0$, and $2b > 0$; together with $r \geq 0$, this makes the reduced linear coefficients nonnegative,

$$c_x = q\Big((a+b) + \frac{r}{a+b}\Big) \geq 0, \qquad c_y = 2b + \frac{r}{2b} \geq 0, \qquad c_z = 2m \geq 0.$$

Write $M = 2m$ and record the change of variables used below:

$$t = (qx,\, y,\, z_{\mathrm{sp}}), \qquad \alpha^\sharp = \Big(\frac{c_x}{q}, c_y, c_z\Big), \qquad \beta^\sharp = \Big(\frac{1}{q}, 1, 1\Big).$$

Since $q > 0$, this map carries T_m to the simplex Δ_M and the inverse formula $t \mapsto (t_x/q, t_y, t_z)$ carries Δ_M back to T_m. The mass equation $qx + y + z_{\mathrm{sp}} = 2m$ becomes $t_x + t_y + t_z = M$, the parity constraint becomes $t_y + t_z \geq d_m$, and the objectives match:

$$\sum_i \alpha_i^\sharp t_i = \frac{c_x}{q}\, qx + c_y y + c_z z_{\mathrm{sp}} = c_x x + c_y y + c_z z_{\mathrm{sp}}, \qquad \sum_i \beta_i^\sharp t_i^2 = \frac{(qx)^2}{q} + y^2 + z_{\mathrm{sp}}^2 = qx^2 + y^2 + z_{\mathrm{sp}}^2,$$

so that

$$\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}(t) = \phi_{r,\kappa}(x, y, z_{\mathrm{sp}}).$$

We also use $0 \le d_m \le M$, since $d_m \in \{0, 2/m\}$ and $m \ge 2$.

Equality and range. By definition, $\rho_\star(m, a, b, r, \kappa)$ is the reduced implementable-minus-relaxed value,

$$\rho_\star(m, a, b, r, \kappa) = \inf_{\substack{(x,y,z_{\rm sp}) \in T_m \\ y+z_{\rm sp} \ge d_m}} \phi_{r,\kappa}(x, y, z_{\rm sp}) - \inf_{(x,y,z_{\rm sp}) \in T_m} \phi_{r,\kappa}(x, y, z_{\rm sp}).$$

The reduced gap identity in Lemma 8 identifies $\Delta_m^{\pm}(r, \kappa)$ with this same difference. Hence

$$\rho_\star(m, a, b, r, \kappa) = \Delta_m^{\pm}(r, \kappa).$$

For the range, every point of T_m has nonnegative coordinates, and $\kappa \ge 0$; with the coefficient inequalities above,

$$\phi_{r,\kappa}(x, y, z_{\rm sp}) = c_x x + c_y y + c_z z_{\rm sp} + \kappa\sqrt{qx^2 + y^2 + z_{\rm sp}^2} \ge 0.$$

The truncated feasible set is nonempty because $(1, 1, 1) \in T_m$ and $1{+}1 \ge d_m$. Since the truncated feasible set is contained in T_m, the constrained reduced infimum is at least the unconstrained one, and therefore

$$\rho_\star(m, a, b, r, \kappa) \ge 0.$$

Together with the equality just proved, this gives $\Delta_m^{\pm}(r, \kappa) \ge 0$.

Zero loss. The reduced minima used in the zero-loss argument are attained by the simplex construction above. For the relaxed problem, if $\kappa > 0$, the positive-robustness clause of Lemma 6 supplies an active-set minimizer of $\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}$ on Δ_M; if $\kappa = 0$, the zero-robustness clause identifies the exposed $\alpha^\sharp$-minimizing face, and placing the mass M on a coordinate where $\alpha_i^\sharp$ is minimal gives a relaxed minimizer. Transporting this point back gives $\xi^{\rm rel} \in T_m$ such that

$$\phi_{r,\kappa}(\xi^{\rm rel}) \le \phi_{r,\kappa}(x, y, z_{\rm sp}) \quad \text{for all } (x, y, z_{\rm sp}) \in T_m.$$

For the truncated problem, apply the truncation dichotomy in Lemma 7 to the image of a relaxed minimizer. If that image lies in $K_{d_m} = \{t \in \Delta_M : t_y + t_z \ge d_m\}$, it remains minimizing on K_{d_m}; otherwise the selector point on the segment $t_y + t_z = d_m$ minimizes on K_{d_m}. Transporting the selected point back gives $\xi^{\rm impl} \in T_m$ with $d_m \le \xi_y^{\rm impl} + \xi_z^{\rm impl}$ and

$$\phi_{r,\kappa}(\xi^{\rm impl}) \le \phi_{r,\kappa}(x, y, z_{\rm sp}) \quad \text{whenever } (x, y, z_{\rm sp}) \in T_m \text{ and } d_m \le y + z_{\rm sp}.$$

Thus the two reduced values are the objective values at $\xi^{\rm rel}$ and $\xi^{\rm impl}$, and

$$\rho_\star = \phi_{r,\kappa}(\xi^{\rm impl}) - \phi_{r,\kappa}(\xi^{\rm rel}).$$

If $\rho_\star = 0$, the two displayed objective values are equal. Since $\xi^{\rm impl}$ is feasible for the relaxed problem and attains the relaxed value, it is a relaxed minimizer satisfying $d_m \le \xi_y^{\rm impl} + \xi_z^{\rm impl}$. Conversely, if some $(x, y, z_{\rm sp}) \in T_m$ minimizes $\phi_{r,\kappa}$ over T_m and satisfies $d_m \le y + z_{\rm sp}$, then it is feasible for the truncated problem, so

$$\phi_{r,\kappa}(\xi^{\rm rel}) \le \phi_{r,\kappa}(\xi^{\rm impl}) \le \phi_{r,\kappa}(x, y, z_{\rm sp}) = \phi_{r,\kappa}(\xi^{\rm rel}),$$

forcing equality throughout and $\rho_\star = 0$. This proves

$$\rho_\star = 0 \iff \text{some relaxed minimizer of } \phi_{r,\kappa} \text{ on } T_m \text{ satisfies } d_m \le y + z_{\rm sp}.$$

Strict robustness. Assume $\kappa > 0$. Under the same change of variables, the positive-robustness clause of Lemma 6 gives a strictly unique minimizer of $\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}$ on Δ_M. Transporting it back gives exactly one minimizer of $\phi_{r,\kappa}$ on T_m. With uniqueness in hand, the assertion that some relaxed minimizer satisfies $d_m \le y + z_{\mathrm{sp}}$ is equivalent to the assertion that every relaxed minimizer satisfies it. The zero-loss criterion therefore becomes

$$\rho_\star = 0 \iff \text{every minimizer } (x, y, z_{\mathrm{sp}}) \text{ of } \phi_{r,\kappa} \text{ on } T_m \text{ has } d_m \le y + z_{\mathrm{sp}}.$$

Linear case. Assume $\kappa = 0$, and fix $(x, y, z_{\mathrm{sp}}) \in T_m$. Passing to $t = (qx, y, z_{\mathrm{sp}})$, the reduced objective is the linear weighted-simplex objective $\sum_i \alpha_i^\sharp t_i$. The zero-robustness clause of Lemma 6 says that t minimizes it over Δ_M exactly when

$$t_i \neq 0 \implies \alpha_i^\sharp \le \alpha_j^\sharp \ \text{ for every } j.$$

Since $q > 0$, $t_x \neq 0$ is equivalent to $x \neq 0$, while $t_y = y$ and $t_z = z_{\mathrm{sp}}$; and $\alpha^\sharp = (c_x/q, c_y, c_z)$. Reading the three coordinates in turn, this condition is exactly

$$x \neq 0 \Rightarrow \frac{c_x}{q} \le c_y \text{ and } \frac{c_x}{q} \le c_z, \qquad y \neq 0 \Rightarrow c_y \le \frac{c_x}{q} \text{ and } c_y \le c_z, \qquad z_{\mathrm{sp}} \neq 0 \Rightarrow c_z \le \frac{c_x}{q} \text{ and } c_z \le c_y,$$

as asserted.

Active-set formula. Assume $\kappa > 0$. The active-set construction from Lemma 6 supplies an admissible support $S \subseteq \{x, y, z_{\mathrm{sp}}\}$ and multiplier λ for $\alpha^\sharp, \beta^\sharp, \kappa$. Write the associated simplex point as $t^{S,\lambda}$. It is nonnegative and satisfies $t_x^{S,\lambda} + t_y^{S,\lambda} + t_z^{S,\lambda} = M = 2m$; therefore its preimage has nonnegative coordinates and satisfies

$$q\Big(\frac{t_x^{S,\lambda}}{q}\Big) + t_y^{S,\lambda} + t_z^{S,\lambda} = 2m, \qquad \Big(\frac{t_x^{S,\lambda}}{q}, t_y^{S,\lambda}, t_z^{S,\lambda}\Big) \in T_m.$$

The weighted-simplex minimizer property transfers back through the objective identity, so this reduced point minimizes $\phi_{r,\kappa}$ over T_m. The value formula in Lemma 6 gives $\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}(t^{S,\lambda}) = M\lambda$, and hence the relaxed reduced value is

$$\inf_{(x,y,z_{\mathrm{sp}}) \in T_m} \phi_{r,\kappa} = 2m\,\lambda.$$

Truncation correction. Suppose no relaxed minimizer in T_m satisfies $d_m \le y + z_{\mathrm{sp}}$. Let ξ^{rel} be a relaxed minimizer supplied by the simplex construction, and let $t^{\mathrm{rel}} = (q\xi_x^{\mathrm{rel}}, \xi_y^{\mathrm{rel}}, \xi_z^{\mathrm{rel}})$. Then $t^{\mathrm{rel}} \in \Delta_M$, it minimizes $\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}$ over Δ_M, and the hypothesis gives $t_y^{\mathrm{rel}} + t_z^{\mathrm{rel}} < d_m$. In the case $\kappa = 0$, the tie convention in Lemma 7 is vacuous under this hypothesis, because the exposed relaxed face is disjoint from K_{d_m}. Since $0 \le d_m \le M$, Lemma 7 puts a minimizer of $\Phi_{\alpha^\sharp,\beta^\sharp,\kappa}$ over

$$K_{d_m} = \{t \in \Delta_M : t_y + t_z \ge d_m\}$$

on the boundary segment

$$H_{d_m} = \{(2m - d_m, s, d_m - s) : 0 \le s \le d_m\}$$

at

$$s_\star = \mathrm{truncSelector}(2m, d_m, \alpha^\sharp, \beta^\sharp, \kappa).$$

Define the reduced implementable value by

$$\mathrm{val}_{\mathrm{impl}}(m,a,b,r,\kappa) := \inf_{\substack{(x,y,z_{\mathrm{sp}})\in T_m \\ d_m \le y+z_{\mathrm{sp}}}} \phi_{r,\kappa}(x,y,z_{\mathrm{sp}}).$$

The change of variables carries K_{d_m} onto the parity-truncated slice of T_m and preserves objective values. Therefore

$$\mathrm{val}_{\mathrm{impl}}(m,a,b,r,\kappa) = \Phi_{\alpha^\sharp,\beta^\sharp,\kappa}(2m-d_m, s_\star, d_m - s_\star).$$

Even block size. If m is even, then $d_m = 0$. A relaxed minimizer exists by the simplex construction, and every point of T_m has $y, z_{\mathrm{sp}} \ge 0$, hence $0 = d_m \le y + z_{\mathrm{sp}}$. The zero-loss criterion gives

$$\rho_\star(m,a,b,r,\kappa) = 0.$$

□

References

Emmanuel Abbe. Community detection and stochastic block models. *Foundations and Trends in Communications and Information Theory*, 14(1-2):1–162, 2018. doi: 10.1561/0100000067.

Peter M. Aronow and Cyrus Samii. Estimating average causal effects under general interference, with application to a social network experiment, 2013. URL https://arxiv.org/abs/1305.6156.

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

Abhijit Banerjee, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013. doi: 10.1126/science.1236498.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, 2004. ISBN 978-0-521-83378-3.

Michel M. Deza and Monique Laurent. *Geometry of Cuts and Metrics*, volume 15 of *Algorithms and Combinatorics*. Springer, Berlin, 1997. ISBN 978-3-540-61611-6.

Bradley Efron. Forcing a sequential experiment to be balanced. *Biometrika*, 58(3):403–417, 1971. doi: 10.1093/biomet/58.3.403.

Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. doi: 10.1145/227683.227684.

Christopher Harshaw, Fredrik Sävje, Daniel Spielman, and Peng Zhang. Balancing covariates in randomized experiments with the gram-schmidt walk design, 2019. URL https://arxiv.org/abs/1911.03071.

Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983. doi: 10.1016/0378-8733(83)90021-7.

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.

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

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

Shuangning Li and Stefan Wager. Random graph asymptotics for treatment effect estimation under network interference, 2020. URL https://arxiv.org/abs/2007.13302.

Wenrui Li, Daniel L. Sussman, and Eric D. Kolaczyk. Causal inference under network interference with noise, 2021. URL https://arxiv.org/abs/2105.04518.

Miller McPherson, Lynn Smith-Lovin, and James M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001. doi: 10.1146/annurev.soc.27.1.415.

Kari Lock Morgan and Donald B. Rubin. Rerandomization to improve covariate balance in experiments. *The Annals of Statistics*, 40(2):1263–1282, 2012. doi: 10.1214/12-AOS1008.

M. E. J. Newman. Assortative mixing in networks. *Physical Review Letters*, 89(20):208701, 2002. doi: 10.1103/PhysRevLett.89.208701.

Donald B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688–701, 1974. doi: 10.1037/h0037350.

Fredrik Sävje. Causal inference with misspecified exposure mappings: separating definitions and assumptions, 2021. URL https://arxiv.org/abs/2103.06471.

Fredrik Sävje, Peter M. Aronow, and Michael G. Hudgens. Average treatment effects in the presence of unknown interference, 2017. URL https://arxiv.org/abs/1711.06399.

Vydhourie Thiyageswaran, Alex Kokot, Jennifer Brennan, Marina Meila, Christina Lee Yu, and Maryam Fazel. Optimal design under interference, homophily, and robustness trade-offs, 2026. URL https://arxiv.org/abs/2601.17145.

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

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