Second-order minimax risk in multi-arm binary randomization

CausalSmith*

September 12, 2026

Abstract

This paper studies design-based minimax risk for randomized experiments with a fixed number K of treatment arms, n labeled units, binary potential outcomes, and a prespecified nonzero zero-sum treatment contrast. The game, which optimizes jointly over the assignment law and an estimator clipped to the estimand's own range against all complete labeled schedules, has the same value as an explicit finite game whose states are the response-type count vectors over the 2^K binary response types. For every fixed nonzero zero-sum contrast c, the first-order minimax constant is $C_0(c) = L_c^2/4$, where L_c is the ℓ_1 norm of the contrast, and the contrast-weighted allocation $q_a^\star \propto |c_a|$ attains the finite-sample upper bound $C_0(c)/n$ with a projected Horvitz–Thompson rule. The same allocation, paired with an explicit clipped-shrinkage estimator, improves on that envelope by a positive multiple of $n^{-4/3}$ for all sufficiently large n, while an embedded two-arm Bayesian information argument gives a matching $O(n^{-4/3})$ converse. Thus $4/3$ is the universal second-order exponent for every fixed nonzero zero-sum binary contrast. For rational contrasts the paper builds a finite linear program whose primal and dual solutions supply a feasible procedure and a prior — a matching pair of computable upper and lower bounds on the minimax risk, with all data exactly rational — and transfers those brackets to every real contrast by a square-root risk continuity bound. A three-arm example shows that at $n = 3$, for the contrast $(1, -1/2, -1/2)$ under one fixed allocation, estimators using the full observed arm label and outcome attain a strictly smaller worst-case risk than estimators using only a scalar signed score.

1 Introduction

Randomized experiments with several active arms often target a prespecified comparison among treatment means: one regimen against another, one regimen against an average of alternatives, or a signed contrast across several arms. In a finite-population analysis, the potential outcomes attached to the units are fixed and the assignment mechanism supplies the randomness. This paper studies the corresponding design problem when each potential outcome is binary, the number of arms K is fixed, and the target is a nonzero zero-sum contrast c. The object of interest is the minimax design-based mean squared error obtained by optimizing jointly over the randomization law and a clipped estimator, then taking the largest risk over all complete labeled binary schedules.

The first contribution is an exact finite reduction. The unrestricted labeled-schedule minimax risk $\rho_n(c)$, the worst-case risk over all labeled binary potential-outcome schedules, is equal to the

*Machine-generated by the CausalSmith research pipeline. Each displayed formal statement is either mapped to a Lean declaration or marked as a presentation-level definition, and each displayed proof renders a Lean proof, as recorded in the verification appendix (scope, toolchain, commit, and any statement conditional on a published input). Cited work otherwise supplies framing and positioning.

finite response-type orbit value $G_n(c)$ in Definition 7. The reduction in Theorem 1 symmetrizes over unit labels, preserves the finite-population contrast target, and identifies a finite game whose states are response-type count vectors. For fixed K, the response-type space has size 2^K, so the adversary ranges over the integer simplex of counts of binary response types and the statistician ranges over allocation-count distributions and orbit estimators.

The second contribution is the first-order minimax benchmark. Let $L_c = \sum_a |c_a|$, the contrast ℓ_1 norm, and let $S_c = \{a : c_a \neq 0\}$, the active support of the contrast. Theorem 3 shows that

$$\lim_{n\to\infty} n\rho_n(c) = C_0(c), \qquad C_0(c) = \frac{L_c^2}{4}.$$

It also gives a finite-sample procedure attaining the upper envelope $C_0(c)/n$: assign each unit independently among active arms with probabilities $q_a^\star = |c_a|/L_c$, and estimate the centered contrast by the projected contrast-weighted Horvitz–Thompson rule in Definition 14. Thus the first-order allocation depends on the signed contrast through absolute weights, matching the range-weighted logic familiar from finite-population sampling and Horvitz–Thompson estimation [Horvitz and Thompson, 1952, Hansen et al., 1953, Särndal et al., 1992, Aronow and Lopatto, 2026].

The third contribution is the second-order rate. The improvement

$$d_n(c) = \frac{C_0(c)}{n} - \rho_n(c)$$

is positive at the $n^{-4/3}$ scale for every fixed nonzero zero-sum contrast. Theorem 5 constructs an explicit clipped-shrinkage estimator under the same contrast-weighted allocation and proves a uniform risk bound of the form

$$C_0(c)\{n^{-1} - \kappa_c n^{-4/3}\}$$

for a positive contrast-dependent constant κ_c. The converse side in Theorems 4 and 6 embeds the two-arm difficulty inside any fixed contrast and yields the upper bound $d_n(c) \leq 43C_0(c)n^{-4/3}$ for large n. Together these results show that if a regularly varying normalization makes the improvement converge to a positive finite limit, then its index must be $4/3$.

The analysis builds on the design-based causal-inference tradition initiated by Fisher, Neyman, and Rubin and developed in modern finite-population treatments [Fisher, 1935, Splawa-Neyman, 1990, Rubin, 1974, 1978, Holland, 1986, Imbens and Rubin, 2015, Imbens and Wooldridge, 2009, Athey and Imbens, 2017, Ding, 2024]. It is also connected to optimum design, minimax decision theory, Bayesian information inequalities, and second-order minimax analysis [Smith, 1918, Kiefer, 1959, 1974, Atkinson et al., 2007, Wald, 1950, Le Cam, 1986, Ibragimov and Has'minskii, 1981, Rao, 1945, Van Trees, 1968, Gill and Levit, 1995, Levit, 1981, Bickel, 1981, Pinsker, 1980, Johnstone and MacGibbon, 1992]. The nearest causal comparisons are sharp two-arm and bounded-outcome minimax analyses [Sudijono et al., 2026, Hull, 2026], together with recent minimax and optimal randomization work in experimental design [Kallus, 2018, 2020, Bai, 2023, Harshaw et al., 2024, Basse et al., 2023, Dereziński et al., 2019, Hu et al., 2024, Kandiros et al., 2024, 2026]. The present results give a fixed-K binary finite-population counterpart: an exact orbit game, a design and estimator class as large as the estimand's own range allows, and computable finite bounds on the minimax value.

Computable bounds are the fourth contribution. By a *certificate* we mean a finite object whose risk can be evaluated exactly and which therefore bounds $\rho_n(c)$ on one side: a feasible procedure bounds it above, a prior bounds it below. For rational contrasts, Definition 19 and Theorem 7 define a grid-action linear program over allocation-count orbits and estimator-grid actions. Its primal solution is an invariant randomized procedure, its dual response-count multipliers form a prior, and the gap between the two bounds is controlled by the grid mesh. Theorems 8 and 9 transfer these rational certificates to real contrasts through a Lipschitz bound for the square root of minimax risk. At the $n^{4/3}$ scale, the mesh can be chosen so that the certificate width is asymptotically negligible for the rational certificates, and for real contrasts after rational approximation and continuity transfer.

The three-arm diagnostic illustrates why the orbit experiment retains more structure than a scalar contrast score. For $c^{\dagger} = (1, -1/2, -1/2)$, Proposition 2 compares the full arm-label-and-outcome experiment with a signed-score reduction under the same contrast-weighted allocation. At $n = 3$, a full-data clipped estimator attains a strictly smaller worst-case risk than any rule based on the scalar score. The support-specific results in Theorems 11 and 12 then separate two-active-arm contrasts, where the value is exactly the two-arm value scaled by $C_0(c)$, from contrasts with at least three active arms, where the same order sandwich holds and the finite orbit programs supply the brackets.

The results apply directly to prespecified contrasts in multi-regimen binary-endpoint trials. In an ACTG 175-style four-regimen setting [Hammer et al., 1996], a contrast comparing one regimen with another falls under the support-two reduction, while a contrast comparing one regimen with an average of alternatives uses the support-at-least-three conclusions and the finite-program certificates. Section B records the scope of the mathematical results, the presentation-only notation, and the boundary of the formal claims.

Section 2 places the analysis in the literatures on design-based causal inference, finite-population sampling, optimum design, information inequalities, and minimax randomization. Section 3 defines the labeled schedule game, the response-type orbit likelihood, and the exact orbit reduction. Section 4 establishes the first-order constant, the clipped-shrinkage improvement, and the embedded two-arm converse. Section 5 develops rational grid certificates and the continuity transfer to real contrasts. Section 6 gives the three-arm diagnostic and the support-specific conclusions. Section 7 discusses design implications, limitations, and future work, and Sections A to C collect the auxiliary lemmas, the verification record, and the proofs of the main results.

2 Related work

Randomization-based causal inference treats the assignment mechanism as the source of uncertainty and potential outcomes as fixed features of the finite population. This design-based perspective goes back to Fisher [1935], Splawa-Neyman [1990], and the potential-outcomes formulation of Rubin [1974, 1978]; see also Holland [1986], Imbens and Rubin [2015], Imbens and Wooldridge [2009], Athey and Imbens [2017], and Ding [2024]. Recent econometric treatments emphasize how the choice of randomization affects precision, robustness, and finite-sample behavior, including regression adjustment and covariate balance [Freedman, 2008, Lin, 2013, Dasgupta et al., 2015, Morgan and Rubin, 2012, Li et al., 2018, Bai et al., 2024]. The present paper works in the same finite-population tradition, with a fixed number of treatment arms K, a fixed number of labeled units n, binary potential outcomes, and a nonzero zero-sum treatment contrast c. Its target is the exact design-estimator minimax risk and its first- and second-order behavior under the randomization distribution.

Finite-population sampling theory provides a parallel design-based language for choosing sampling laws and estimators under adversarial finite populations. Classical unequal-probability sampling and Horvitz–Thompson estimation [Horvitz and Thompson, 1952, Hansen et al., 1953, Särndal et al., 1992] connect naturally to randomized experiments because treatment assignment can be read as a sampling design over exposure cells. Gabler [1990] and Aronow and Lopatto [2026] study minimax ideas in finite-population sampling, with the latter giving a close comparison through finite-population minimax sampling under unbiasedness. The analysis here is aligned with that decision-theoretic finite-population viewpoint and specializes it to multi-arm binary randomization, where both the assignment law and the clipped estimator enter the minimax criterion.

The paper is also connected to optimum experimental design. The classical design literature studies allocation rules through criteria derived from information matrices or risk functionals [Smith, 1918, Kiefer, 1959, 1974, Atkinson et al., 2007]. In causal experiments, minimax and optimality criteria have been used to evaluate and construct randomization designs under finite-population or model-assisted objectives. Kallus [2018] develops optimal a priori balance and Kallus [2020] studies optimal randomization designs; related modern work includes Basse et al. [2023], Harshaw et al. [2024], Kandiros et al. [2024, 2026], Eckles et al. [2014], Dereziński et al. [2019], and Hu et al. [2024]. Relative to this literature, the contribution is an exact finite binary minimax analysis for treatment contrasts, including the orbit reduction and finite-program certificates used to evaluate sharp finite-n quantities.

Information inequalities and local asymptotic arguments supply the lower-bound tools behind many minimax results. The decision-theoretic tradition begins with Wald [1950]; asymptotic minimax and local experiment methods are developed in Le Cam [1986], Ibragimov and Has'minskii [1981], and related work. Bayesian Cramér–Rao and van Trees inequalities [Van Trees, 1968, Gill and Levit, 1995], together with classical information identities [Rao, 1945], provide lower bounds for estimation risk. Second-order minimax refinements have a long statistical lineage, including Levit [1981], and related shrinkage and nonparametric minimax phenomena appear in Bickel [1981], Pinsker [1980], Feldman [1991], and Johnstone and MacGibbon [1992]. The results below use this tradition to identify the first-order minimax constant and to analyze the size of the second-order improvement in the finite-population binary randomization game.

The closest econometric comparisons are recent sharp analyses of two-arm designs. Sudijono et al. [2026] obtains a sharp two-arm binary expansion, and Hull [2026] studies bounded two-arm potential outcomes. Rosenman and Hunter [2026] is also related through shrinkage ideas for experimental estimation. The present paper places those two-arm themes inside a multi-arm contrast framework: it gives exact first-order risk for general fixed contrasts, embeds the two-arm converse where it is sharp, and constructs finite linear-program certificates for rational and real contrasts. A three-arm example then shows, at $n = 3$ and for the single contrast $c^\dagger = (1, -1/2, -1/2)$ under one fixed allocation, that estimators using the full observed arm label and outcome achieve a strictly smaller worst-case risk than estimators using only a scalar signed score. That comparison isolates what the data reduction discards for this fixed three-arm support, fixed allocation, and sample size.

3 Setup and the finite orbit game

Throughout the paper, n denotes the fixed number of labeled units, K the fixed number of treatment arms, and asymptotic notation is taken with K and the contrast fixed. Sums over arms run over the treatment-arm set and sums over response types run over the corresponding binary type space. Generic constants may depend on the fixed contrast and on K, but not on n.

The finite-population model starts with a labeled binary potential-outcome schedule. The treatment-arm set $\mathcal{A}_K$ indexes the arms, and the response-type space records the vector of binary potential outcomes attached to each unit.

Definition 1. [`synth_1`] (Treatment-arm label set) *For each natural number K, the treatment-arm set is the finite set of its K arm labels:*

$$\mathcal{A}_K := \{1, \ldots, K\}.$$

Definition 2. [`synth_2`] (Binary response types) *For a fixed number K of treatment arms, let $\mathcal{A}_K$ be the treatment-arm label set. The binary response-type space is*

$$\mathcal{T}_K := \{0,1\}^{\mathcal{A}_K}.$$

Thus each $t \in \mathcal{T}_K$ assigns one binary potential-outcome value $t_a \in \{0,1\}$ to every arm $a \in \mathcal{A}_K$.

This notation is the multi-arm version of the binary potential-outcome framework used in design-based causal inference [Splawa-Neyman, 1990, Rubin, 1974, Holland, 1986, Imbens and Rubin, 2015]. A complete schedule assigns one response type to each labeled unit; an assignment design then randomizes only the observed arm, while the schedule itself remains fixed.

Three objects enter the next definition and are fixed here once. The contrast c is a nonzero zero-sum vector, that is $c \in \mathbb{R}^K \setminus \{0\}$ with $\sum_{a \in \mathcal{A}_K} c_a = 0$ (restated as Definition 10 below, where the rational case is also named); we write

$$L_c = \sum_{a \in \mathcal{A}_K} |c_a|, \qquad I_c = \left[-\frac{L_c}{2}, \frac{L_c}{2}\right], \qquad \tau_c(z) = \frac{1}{n} \sum_{i=1}^{n} \sum_{a \in \mathcal{A}_K} c_a \, t_{i,a}$$

for its ℓ_1 norm, the interval it generates, and the finite-population estimand at the complete schedule $z = (t_1, \ldots, t_n)$. Because c sums to zero and each $t_{i,a}$ is 0 or 1, the estimand always lies in I_c. The infimum in the definition below runs over assignment laws $\mathcal{D} \in \Delta(\mathcal{A}_K^n)$ and over estimators

$$\widehat{\tau} \colon (\mathcal{A}_K \times \{0,1\})^n \to I_c.$$

Confining the estimator to I_c costs nothing: projecting a real-valued rule onto I_c never increases squared error at a target that lies in I_c, so the value of the game is the same as it would be with arbitrary real-valued estimators. Finally, $\mathcal{M}_{n,K}$ denotes the set of nonnegative integer vectors $(m_t)_{t \in \mathcal{T}_K}$ with $\sum_t m_t = n$, and $\mathcal{R}_{n,K}$ the set of nonnegative integer vectors $(r_a)_{a \in \mathcal{A}_K}$ with $\sum_a r_a = n$; these are finite sets of lattice points, and the shorter word "simplex" is occasionally used for them below.

Definition 3. [`def:labeled-schedule-game`] (Labeled-schedule risk $\rho_n(c)$) *The unrestricted labeled-schedule minimax risk for contrast c is*

$$\rho_n(c) = \inf_{\mathcal{D} \in \Delta(\mathcal{A}_K^n), \, \widehat{\tau}} \max_{z \in \mathcal{T}_K^n} \mathbb{E}_{\mathcal{D}} \left[\left\{ \widehat{\tau}((A_i, Y_i^{\mathrm{obs}})_{i=1}^n) - \tau_c(z) \right\}^2 \right].$$

The risk $\rho_n(c)$ is the benchmark for joint design and estimation; by the previous paragraph nothing is lost by the clipping, so we drop the qualifier "unrestricted" from here on. It optimizes over an assignment law $\mathcal{D}$ and a clipped estimator $\widehat{\tau}$, then evaluates the largest design-based squared error over complete binary schedules. The contrast c is zero-sum, so the estimand is a treatment comparison rather than a level.

The next definitions pass from labeled schedules to response-type counts. This is the statistic that remains after averaging over relabelings of the units.

Definition 4. [`synth_4`] (Response-type count vector) *For a complete labeled schedule $z = (t_1, \dots, t_n) \in \mathcal{T}_K^n$, define*

$$\mathrm{counts}(z) = \big(\mathrm{counts}(z)_t\big)_{t\in\mathcal{T}_K} \in \mathcal{M}_{n,K}, \qquad \mathrm{counts}(z)_t = \#\{i \in \{1,\dots,n\} : t_i = t\}.$$

Thus $\mathrm{counts}(z)$ *is the unit-permutation orbit of z.*

Definition 5. [`synth_19`] (Notation for response-type counts) *For a complete labeled response schedule $z = (t_1, \dots, t_n) \in \mathcal{T}_K^n$ we write $m(z) = \mathrm{counts}(z)$ for the response-type count vector of Definition 4, with coordinates*

$$m_t(z) = \#\{i \in \{1,\dots,n\} : t_i = t\}, \qquad t \in \mathcal{T}_K,$$

so that $\sum_{t\in\mathcal{T}_K} m_t(z) = n$. A generic element of $\mathcal{M}_{n,K}$ is written m, and m_t is its coordinate at t.

The count vector m belongs to the response-count simplex $\mathcal{M}_{n,K}$. It preserves the finite-population target because the contrast averages unit-level potential outcomes symmetrically across labels.

Given a response-count vector and an allocation-count vector, the observable data reduce to arm-specific success counts. The likelihood below counts all response-type-by-arm contingency tables compatible with those margins.

Definition 6. [`def:orbit-likelihood`] (Orbit likelihood $P_m(x \mid r)$) *The orbit likelihood of arm-success counts x given allocation counts r and response-type counts m is*

$$P_m(x \mid r) = \frac{\prod_{a\in\mathcal{A}_K} r_a!}{n!} \sum_{h\in\mathcal{H}(m,r,x)} \prod_{t\in\mathcal{T}_K} \frac{m_t!}{\prod_{a\in\mathcal{A}_K} h_{t,a}!}.$$

Here $\mathcal{H}(m, r, x)$ is the set of nonnegative integer tables $h = (h_{t,a})$ with response-type margins m, allocation margins r, and arm-success totals x.

The orbit likelihood $P_m(x \mid r)$ is purely finite-population: it is induced by random assignment among labeled units with fixed response types. The allocation-count vector r lies in the allocation-count simplex $\mathcal{R}_{n,K}$, and the observed-success vector x records one count for each arm.

Definition 7. [`def:finite-orbit-game`] (Finite orbit game $G_n(c)$) *The finite response-type orbit minimax value is*

$$G_n(c) = \inf_{\pi\in\Delta(\mathcal{R}_{n,K}),\,\delta}\ \max_{m\in\mathcal{M}_{n,K}} \sum_{r\in\mathcal{R}_{n,K}} \pi_r \sum_{x\in\prod_{a\in\mathcal{A}_K}\{0,\dots,r_a\}} P_m(x \mid r)\{\delta(r,x) - \tau_c(m)\}^2.$$

Here $\mathcal{M}_{n,K}$ is the simplex of response-type count vectors and $\mathcal{R}_{n,K}$ is the simplex of allocation-count vectors.

The finite value $G_n(c)$ optimizes over an invariant design distribution π on allocation-count orbits and an orbit estimator δ. Its target $\tau_c(m)$ is the count-level version of the labeled finite-population contrast.

For later reference, it is useful to name both the labeled and orbit risks explicitly.

Definition 8. [`synth_3`] (Design-based squared-error risk) *For a complete labeled response schedule $z \in \mathcal{T}_K^n$, an assignment design $\mathcal{D} \in \Delta(\mathcal{A}_K^n)$, and an estimator $\widehat{\tau} : (\mathcal{A}_K \times \{0,1\})^n \to [-L_c/2, L_c/2]$, define*

$$R_n(\mathcal{D}, \widehat{\tau}; z) = \mathbb{E}_{\mathcal{D}}\left[\left\{\widehat{\tau}((A_i, Y_i^{\mathrm{obs}})_{i=1}^n) - \tau_c(z)\right\}^2\right],$$

where $A \sim \mathcal{D}$, $Y_i^{\mathrm{obs}} = t_{i,A_i}$, and the expectation is only over the assignment randomization. If $p = (\mathcal{D}, \widehat{\tau})$ is a labeled procedure, write $R_n(p; z)$ for the same quantity.

Definition 9. [`synth_10`] (Orbit-procedure risk) *For a contrast c, an orbit procedure is a pair $q = (\pi, \delta)$ with $\pi \in \Delta(\mathcal{R}_{n,K})$ and $\delta(r, x) \in [-L_c/2, L_c/2]$. Its risk at response-type count vector $m \in \mathcal{M}_{n,K}$ is*

$$R_n(q; m) = \sum_{r \in \mathcal{R}_{n,K}} \pi_r \sum_{x \in \prod_{a \in \mathcal{A}_K}\{0,\dots,r_a\}} P_m(x \mid r)\{\delta(r, x) - \tau_c(m)\}^2.$$

In the two-arm contrast $c = (1, -1)$, write

$$R^{\mathrm{orb}}_{(1,-1)}(q; m) = \sum_{r \in \mathcal{R}_{n,2}} \pi_r \sum_{x \in \prod_{a \in \mathcal{A}_2}\{0,\dots,r_a\}} P_m(x \mid r)\{\delta(r, x) - \tau_{(1,-1)}(m)\}^2.$$

Definition 10. [`synth_20`] (Nonzero zero-sum contrast) *For $K \geq 2$, a zero-sum contrast is a vector $c = (c_a)_{a \in \mathcal{A}_K} \in \mathbb{R}^K$ satisfying $\sum_{a \in \mathcal{A}_K} c_a = 0$. It is nonzero when $c \neq 0$. A rational nonzero zero-sum contrast is such a vector with $c \in \mathbb{Q}^K$.*

Definition 11. [`synth_15`] (Compatible observation-success counts) *For an allocation-count vector $r \in \mathcal{R}_{n,K}$, an observed-success vector $x = (x_a)_{a \in \mathcal{A}_K}$ is compatible with r when $x_a \in \{0, \dots, r_a\}$ for every arm a. For $m \in \mathcal{M}_{n,K}$, the vector x is compatible with (m, r) when there exists a contingency table $h = (h_{t,a})_{t \in \mathcal{T}_K, a \in \mathcal{A}_K} \in \mathbb{Z}_+^{\mathcal{T}_K \times \mathcal{A}_K}$ such that $\sum_a h_{t,a} = m_t$ for every t, $\sum_t h_{t,a} = r_a$ for every a, and $\sum_{t : t_a = 1} h_{t,a} = x_a$ for every a; equivalently, $\mathcal{H}(m, r, x)$ is nonempty.*

The compatibility condition spells out the feasible fibers in Definition 6. A table h records how many units of each response type are assigned to each arm; its margins determine the allocation counts and the observed successes.

The exact reduction uses the unit-permutation orbit representation of the experiment. Invariance requires the design and estimator to ignore the arbitrary labels of the units.

Definition 12. [`synth_14`] (Invariant labeled procedure) *A labeled procedure is a pair $p = (\mathcal{D}, \widehat{\tau})$, where $\mathcal{D} \in \Delta(\mathcal{A}_K^n)$ is an assignment law and $\widehat{\tau} : (\mathcal{A}_K \times \{0,1\})^n \to [-L_c/2, L_c/2]$ is an estimator. For $\sigma \in \mathfrak{S}_n$, define $(\sigma A)_i = A_{\sigma^{-1}(i)}$, $(\sigma z)_i = z_{\sigma^{-1}(i)}$, and $\sigma((A_i, Y_i^{\mathrm{obs}})_{i=1}^n) = ((A_{\sigma^{-1}(i)}, Y^{\mathrm{obs}}_{\sigma^{-1}(i)}))_{i=1}^n$. The procedure p is invariant when, for every $\sigma \in \mathfrak{S}_n$, the law of σA under $\mathcal{D}$ equals $\mathcal{D}$ and $\widehat{\tau}(\sigma o) = \widehat{\tau}(o)$ for every observed-data vector $o \in (\mathcal{A}_K \times \{0,1\})^n$.*

The following theorem is the main reduction in this section. It establishes that symmetrizing over unit labels preserves the minimax value and identifies the resulting finite game exactly.

Theorem 1. [`thm:exact-response-type-game`] (Exact orbit game reduction) *Fix $n \in \mathbb{N}$, $K \geq 2$, and a nonzero zero-sum contrast c in the labeled-schedule and finite-orbit games of Definitions 3 and 7. Then:*

- ***(Lossless symmetrization.)*** *For every labeled procedure p, there are an invariant labeled procedure $\bar{p}$ and an orbit procedure q such that $\bar{p}$ is the common-unit-permutation average of p, and for every complete labeled schedule $z \in \mathcal{T}_K^n$,*

$$R_n(\bar{p}; z) \leq \frac{1}{|\mathfrak{S}_n|} \sum_{\sigma \in \mathfrak{S}_n} R_n(p; \sigma z).$$

  The orbit procedure q is dominated by this averaged invariant representative in the corresponding orbit risk.

- ***(Exact invariant correspondence.)*** *Invariant labeled procedures are in bijection with orbit procedures (π, δ). Under this bijection, each invariant labeled procedure and its orbit representative realize the same allocation-count and observation-count rule, and for every $z \in \mathcal{T}_K^n$,*

$$R_n(q; \operatorname{counts}(z)) = R_n(\bar{p}; z).$$

- ***(Target preservation.)*** *If $m = \operatorname{counts}(z)$, with $m_t = \#\{i : t_i = t\}$, then*

$$\tau_c(z) = \frac{1}{n} \sum_i \sum_{a \in \mathcal{A}_K} c_a \, t_{i,a} = \frac{1}{n} \sum_t m_t \sum_{a \in \mathcal{A}_K} c_a \, t_a = \tau_c(m).$$

- ***(Value equality.)*** *The unrestricted labeled-schedule minimax risk equals the finite response-type orbit minimax value:*

$$\rho_n(c) = G_n(c).$$

- ***(Orbit saddle point.)*** *There exist an orbit procedure $q^\star$ and a prior ν on $\mathcal{M}_{n,K}$ such that, for every response-count vector m and every orbit procedure q,*

$$R_n(q^\star; m) \leq G_n(c) \qquad \text{and} \qquad G_n(c) \leq \mathbb{E}_\nu[R_n(q; m)] \,.$$

The intuition is that the labels of units carry no information about the contrast target once the complete schedule is fixed. Averaging a procedure over common relabelings cannot increase worst-case risk, and the averaged procedure depends on the assignment and outcomes only through allocation counts and observed-success counts. Theorem 1 therefore turns the unrestricted labeled problem into a finite decision problem over response-count vectors, while preserving both the target and the minimax value.

The final part of the setup records the two-arm benchmark and the fixed-K extension. The two-arm bounded schedule domain is included because recent two-arm results provide an important comparison point for binary minimax analysis [Sudijono et al., 2026, Hull, 2026].

Definition 13. [synth_9] (Hull bounded-schedule domain and risk) *For bounds $L < U$, let $\mathcal{Y}^{\mathrm{Hull}}_{n,L,U}$ be the set of two-arm bounded potential-outcome schedules*

$$Y = \big(Y_i(1), Y_i(2)\big)_{i=1}^{n} \qquad \text{with} \qquad Y_i(a) \in [L, U]$$

for every unit i and arm $a \in \{1, 2\}$. For a design $\mathcal{D}$, estimator $\widehat{\tau}$, and schedule $Y \in \mathcal{Y}^{\mathrm{Hull}}_{n,L,U}$, define the Hull squared-error risk

$$R_{\mathrm{Hull}}(\mathcal{D}, \widehat{\tau}; Y) = \mathbb{E}_{\mathcal{D}}\left[\left\{\widehat{\tau}((A_i, Y_i^{\mathrm{obs}})_{i=1}^{n}) - \frac{1}{n}\sum_{i=1}^{n}\big(Y_i(1) - Y_i(2)\big)\right\}^2\right],$$

where $Y_i^{\mathrm{obs}} = Y_i(A_i)$ and the expectation is over assignment from $\mathcal{D}$.

Theorem 2. [thm:multiarm-strict-extension] (Two-arm specialization and the fixed-K orbit game) *Throughout, $\rho_n(c)$ is the labeled-schedule minimax risk of Definition 3 and $G_n(c)$ the finite orbit-game value of Definition 7. The following statements hold.*

- ***(Two-arm orbit value.)*** *For every $n > 0$, with $c = (1, -1)$,*

$$\rho_n(c) = G_n(c) \qquad \text{and} \qquad C_0(c) = 1.$$

- ***(Two-arm target.)*** *For every $n > 0$ and every $m \in \mathcal{M}_{n,2}$, with $p_+(m) = m_{(1,0)}$ and $p_-(m) = m_{(0,1)}$,*

$$\tau_c(m) = \frac{p_+(m) - p_-(m)}{n}.$$

- ***(Four two-arm response types and likelihood.)*** *For every n, every $m \in \mathcal{M}_{n,2}$, every $r \in \mathcal{R}_{n,2}$, and every compatible observed-success vector x, let*

$$t_{00}(a) = 0, \qquad t_{01} = (0, 1), \qquad t_{10} = (1, 0), \qquad t_{11}(a) = 1.$$

  These four response types are pairwise distinct and exhaust $\mathcal{T}_2$. Moreover, with the likelihood notation of Definition 6,

$$P_m(x \mid r) = \frac{\prod_{a\in\mathcal{A}_2} r_a!}{n!} \sum_{h\in\mathcal{H}(m,r,x)} \prod_{t\in\mathcal{T}_2} \frac{m_t!}{\prod_{a\in\mathcal{A}_2} h_{t,a}!}.$$

- ***(Two-arm second-order scale.)*** *For every $n \geq 8$,*

$$d_n(c) \leq 43\, n^{-4/3}.$$

  There are constants $\kappa > 0$ and N such that, for every $n \geq N$,

$$\kappa n^{-4/3} \leq d_n(c).$$

- ***(Exact multiarm orbit game.)*** *For every $K \geq 3$, every n, and every contrast c,*

$$\rho_n(c) = G_n(c) \qquad \text{and} \qquad |\mathcal{T}_K| = 2^K.$$

- ***(Hull risk-value inclusion.)*** *For every N, L, U, every finite design $\mathcal{D}$ on two-arm assignments, and every measurable Hull estimator $\widehat{\tau}$,*

$$\sup_Y R_{\mathrm{Hull}}(\mathcal{D}, \widehat{\tau}; Y) \in \left\{ v : \exists\, \mathcal{D}', \widehat{\tau}' \text{ with } \widehat{\tau}' \text{ measurable and } v = \sup_Y R_{\mathrm{Hull}}(\mathcal{D}', \widehat{\tau}'; Y) \right\},$$

  where the suprema range over Hull schedules with parameters N, L, U.

- ***(Strict treatment-domain separation.)*** *For every $K \geq 3$, every $n > 0$, and every $L < U$, the Hull bounded-schedule domain with parameters n, L, U is not a subset of the binary K-arm schedule domain:*

$$\mathcal{Y}^{\mathrm{Hull}}_{n,L,U} \not\subseteq \mathcal{T}^n_K.$$

Write

$$d_n(c) = \frac{C_0(c)}{n} - \rho_n(c), \qquad C_0(c) = \frac{L_c^2}{4},$$

for the improvement of the minimax risk on the first-order envelope; Section 4 is devoted to it. For $K = 2$, Theorem 2 identifies the orbit value for $c = (1, -1)$, expresses the estimand through the effect-class counts p_+ and p_-, reproduces the four-type likelihood, and places $d_n(c)$ on the $n^{-4/3}$ scale. For every fixed $K \geq 3$, it gives the exact multi-arm orbit equality and the binary type count 2^K, which supplies the finite state space used by the linear programs of Section 5.

Two clauses of that theorem fix the scope of the bounded-outcome comparison. They record that the worst-case Hull risk of a design–estimator pair is itself an attainable Hull risk value, and that the bounded real-valued two-arm schedule domain and the binary K-arm schedule domain are distinct as sets. Both are elementary, and their role here is bookkeeping: they make precise that the two problems are posed over different schedule domains, so the route between them is an explicit construction. That is the route this paper takes, through the sign-pair embedding used in the converse of Section 4.

4 First-order risk and the second-order rate

The finite orbit reduction in Theorem 1 turns the unrestricted labeled-schedule risk into a finite response-type decision problem. This section first identifies the first-order benchmark in the original labeled game and then refines that benchmark by an explicit second-order construction. Throughout the section, the contrast c is fixed, nonzero, and zero-sum on the treatment-arm set from Definition 3.

The first-order procedure allocates experimental effort in proportion to the absolute contrast weights and estimates the centered contrast by a clipped Horvitz–Thompson rule. The clipping interval is the natural range of the finite-population contrast target.

Definition 14. [`def:first-order-procedure`] (First-order rule $\widehat{\tau}^\star_{\mathrm{cHT}}$) *Let*

$$S_c = \{a \in \mathcal{A}_K : c_a \neq 0\}, \qquad L_c = \sum_{a \in \mathcal{A}_K} |c_a|, \qquad q^\star_a = \frac{|c_a|}{L_c} \quad (a \in S_c).$$

Assign each A_i independently on S_c with $\Pr(A_i = a) = q^\star_a$. The projected contrast-weighted Horvitz–Thompson rule is

$$\widehat{\tau}^\star_{\mathrm{cHT}} = \mathrm{proj}_{[-L_c/2, L_c/2]} \left[n^{-1} \sum_{i=1}^n \sum_{a \in S_c} c_a \mathbf{1}\{A_i = a\} \frac{Y_i^{\mathrm{obs}} - 1/2}{q^\star_a} \right].$$

The active support S_c selects the arms that enter the contrast, while the contrast norm L_c fixes both the target range and the first-order scale. The assignment probabilities $q^\star$ place larger sampling mass on arms with larger absolute contrast weights, matching the variance contribution of the centered Horvitz–Thompson summands.

Definition 15. [`synth_5`] (Contrast-weighted independent design) *Let* $S_c = \{a \in \mathcal{A}_K : c_a \neq 0\}$, $L_c = \sum_{a\in\mathcal{A}_K} |c_a|$, *and*

$$q_a^\star = \begin{cases} |c_a|/L_c, & a \in S_c, \\ 0, & a \notin S_c. \end{cases}$$

The design $\mathcal{D}^\star$ is the product assignment law under which $A_1, \ldots, A_n$ are independent and $\Pr_{\mathcal{D}^\star}(A_i = a) = q_a^\star$ for every unit i and arm a.

Definition 15 records the same allocation as a design object. This formulation will also be used for the shrinkage rule below, so that the first- and second-order procedures share the same contrast-weighted randomization scheme.

Theorem 3. [`thm:first-order-saddle`] (First-order minimax constant and attaining procedure) *Fix a natural number K and a nonzero zero-sum contrast c on the treatment-arm set $\mathcal{A}_K$, in the labeled-schedule game of Definition 3. Let L_c, $q_a^\star$, and $\widehat{\tau}^\star_{\mathrm{cHT}}$ be the contrast norm, contrast-weighted assignment probabilities, and projected contrast-weighted Horvitz–Thompson rule from Definition 14, and set*

$$C_0(c) = \frac{L_c^2}{4}.$$

Then the unrestricted labeled-schedule minimax risk satisfies

$$\lim_{n\to\infty} n\,\rho_n(c) = C_0(c).$$

Moreover, for every $n \geq 1$, the independent contrast-weighted design together with $\widehat{\tau}^\star_{\mathrm{cHT}}$ attains the finite-sample worst-case bound

$$\max_{z\in\mathcal{T}_K^n} R_n(\mathcal{D}^\star, \widehat{\tau}^\star_{\mathrm{cHT}}; z) \leq \frac{C_0(c)}{n}.$$

The theorem gives the first-order robustness constant and a finite-sample certificate attaining its envelope. The constant $C_0(c) = L_c^2/4$ depends on the contrast through its ℓ_1 norm, so the allocation benchmark is determined by the signed estimand rather than by a default equal-allocation rule. The bound is stated in the unrestricted labeled-schedule game of Definition 3, which places it on the same decision scale as the finite orbit value characterized in Theorem 1.

The converse input comes from embedding the sharp two-arm difficulty into the multi-arm contrast. This comparison links every fixed contrast to the two-arm contrast $(1, -1)$, and it also gives an exact specialization when the contrast has two active arms.

Theorem 4. [`thm:embedded-two-arm-converse`] (Embedded two-arm converse) *Fix integers K and n, and a contrast c on $\mathcal{A}_K$, with labeled-schedule minimax risk $\rho_n(c)$ as in Definition 3. Suppose that*

- ***(Arm count.)*** $K \geq 2$.
- ***(Population size.)*** $n \geq 1$.

- ***(Contrast.)*** *c is nonzero and zero-sum.*

Let

$$C_0(c) = \frac{L_c^2}{4}, \qquad \rho_{n,2} = \rho_n((1,-1)), \qquad d_n(c) = \frac{C_0(c)}{n} - \rho_n(c),$$

where $\rho_{n,2}$ is the unrestricted two-arm minimax risk at population size n. Then

$$C_0(c)\rho_{n,2} \le \rho_n(c) \qquad \text{and} \qquad \rho_n(c) \le \frac{C_0(c)}{n}.$$

Moreover, if $n \ge 8$, then

$$0 \le d_n(c), \qquad d_n(c) \le C_0(c)\{n^{-1} - \rho_{n,2}\}, \qquad d_n(c) \le 43\, C_0(c) n^{-4/3}.$$

Finally, writing $S_c = \{a \in \mathcal{A}_K : c_a \ne 0\}$ for the active support, if $|S_c| = 2$, then

$$\rho_n(c) = C_0(c)\rho_{n,2} \qquad \text{and} \qquad d_n(c) = C_0(c)\{n^{-1} - \rho_{n,2}\}.$$

The embedded comparison supplies the lower side of the second-order scale. It says that any improvement over the first-order envelope for a fixed multi-arm contrast is bounded by the corresponding two-arm improvement, up to the scale factor $C_0(c)$. When the active support has size two, the comparison becomes an equality, so the two-arm subproblem fully determines the multi-arm risk for that contrast.

The second-order construction keeps the contrast-weighted design and modifies the estimator. The rule first normalizes the centered Horvitz–Thompson score to the interval scale $h_c = L_c/2$, averages the normalized scores, and then applies a clipped shrinkage toward zero over a bandwidth of order $n^{-1/3}$.

Definition 16. [`synth_16`] (Regular variation of a normalizer) *A positive sequence $(a_n)_{n\ge1}$ is regularly varying with index β when, for every $x > 0$,*

$$\frac{a_{\lfloor xn \rfloor}}{a_n} \longrightarrow x^\beta \qquad \text{as } n \to \infty.$$

Definition 17. [`synth_6`] (Clipped-shrinkage estimator) *Put $h_c = L_c/2$, and let*

$$\lambda_c = \min\left\{ \frac{\left|\sum_{a\in\mathcal{A}_K} c_a t_a\right|}{h_c} : t \in \mathcal{T}_K,\ \sum_{a\in\mathcal{A}_K} c_a t_a \ne 0 \right\}$$

be the smallest nonzero normalized response-type contrast magnitude. Under $\mathcal{D}^\star$ of Definition 15, define

$$W_i = \sum_{a\in S_c} \frac{c_a \mathbf{1}\{A_i = a\}}{q_a^\star}\left(Y_i^{\mathrm{obs}} - \frac{1}{2}\right), \qquad V_i = \frac{W_i}{h_c}, \qquad X_n = \frac{1}{n}\sum_{i=1}^n V_i.$$

For

$$b_n = n^{-1/3}, \qquad \varepsilon_n = \frac{\sqrt{\lambda_c}}{4} n^{-1/3}, \qquad g_n(x) = \max\{-b_n, \min(x, b_n)\},$$

define the clipped-shrinkage rule

$$\widehat{\tau}_n^{\mathrm{sh}} = \mathrm{proj}_{[-h_c,h_c]}\big(h_c\{X_n - \varepsilon_n g_n(X_n)\}\big).$$

The projection is never active for $n \ge 1$, because $|X_n| \le 1$, $0 \le b_n \le 1$ and $0 \le \varepsilon_n \le 1/4$; it is written explicitly only so that the rule takes values in the target interval by construction. This estimator is evaluated as a function of the realized assignment and observed outcomes.

The regular-variation definition in Definition 16 is used to state the uniqueness of the second-order normalizing rate. In Definition 17, the centered arm score W_i, its normalized version V_i, and their average X_n are all computed under the same design as in Definition 15. The constants b_n and ε_n set the clipping width and shrinkage amplitude; their common $n^{-1/3}$ scale produces an $n^{-4/3}$ risk correction after averaging.

Theorem 5. [`thm:universal-second-order-rate`] (Universal second-order rate) *For every fixed $K \geq 2$ and every nonzero zero-sum contrast c, let λ_c, the smallest nonzero normalized response-type contrast magnitude, be*

$$\lambda_c = \min\left\{ \frac{\left|\sum_{a\in\mathcal{A}_K} c_a t_a\right|}{L_c/2} : t \in \mathcal{T}_K,\ \sum_{a\in\mathcal{A}_K} c_a t_a \neq 0 \right\}.$$

Define

$$C_0(c) = \frac{L_c^2}{4}, \qquad A_c = \frac{\sqrt{\lambda_c}}{2} - \frac{\lambda_c}{16}, \qquad B_c = \frac{7\lambda_c}{16}, \qquad \kappa_c = \frac{1}{2}\min\{A_c, B_c\},$$

and

$$d_n(c) = \frac{C_0(c)}{n} - \rho_n(c), \qquad s_n = n^{4/3}.$$

The following assertions hold:

- ***(Spacing.)*** *The constants satisfy*

  $$0 < \lambda_c \leq 1, \qquad 0 < \kappa_c.$$

- ***(First-order envelope.)*** *For every n,*

  $$\rho_n(c) \leq \frac{C_0(c)}{n}.$$

- ***(Uniform shrinkage bound.)*** *There is a finite integer N_c such that, for every integer $m \geq N_c$,*

  $$4\sqrt{\lambda_c}\, m \exp\{-m^{1/3}/8\} \leq \frac{\sqrt{\lambda_c}}{2} - \frac{\lambda_c}{16},$$

  and, for every $n \geq N_c$, the explicit clipped-shrinkage procedure $\widehat{\tau}_n^{\mathrm{sh}}$ satisfies

  $$\sup_{z\in\mathcal{T}_K^n} R_n(\mathcal{D}^\star, \widehat{\tau}_n^{\mathrm{sh}}; z) \leq C_0(c)\left\{n^{-1} - \kappa_c n^{-4/3}\right\}.$$

- ***(Second-order envelope.)*** *The normalized second-order improvement satisfies*

  $$C_0(c)\kappa_c \leq \liminf_{n\to\infty} n^{4/3} d_n(c) \leq \limsup_{n\to\infty} n^{4/3} d_n(c) \leq 43 C_0(c).$$

- ***(Regular-variation index.)*** *For every positive sequence a_n, every β, and every $C > 0$, if a_n is regularly varying with index β and*

  $$a_n d_n(c) \to C,$$

  then $\beta = 4/3$.

Theorem 5 establishes the $n^{-4/3}$ order of the improvement for every fixed nonzero zero-sum contrast. The positive spacing constant λ_c is defined over the finite response-type space, so the shrinkage amount can be chosen directly from the contrast. The lower side of the displayed envelope is witnessed by the explicit clipped-shrinkage procedure of Definition 17, which guarantees an improvement of size $C_0(c)\kappa_c n^{-4/3}$; the upper side follows from the embedded two-arm comparison in Theorem 4. The two sides pin the order of the improvement, not its constant. For contrasts with active support of size at least three, the same sandwich applies after the orbit reduction to the genuinely multi-arm response-type count game; support-two contrasts follow the exact two-arm reduction. The regular-variation clause characterizes $n^{4/3}$ as the unique regularly varying normalization that can yield a positive finite limit for the improvement.

It remains to record the two-arm lower bound used inside the embedded comparison. The result is stated for the unrestricted labeled-schedule minimax risk with the two-arm contrast $(1, -1)$, matching the notation of Theorem 4.

Theorem 6. [`thm:coarse-two-arm-minimax-lower`] (Coarse two-arm lower bound) *For every integer $n \geq 8$, let $\rho_{n,2} := \rho_n((1, -1))$ denote the unrestricted labeled-schedule minimax risk for the two-arm contrast $(1, -1)$. Then*

$$\frac{(1 - n^{-1/3})^2}{n/(1 - n^{-2/3}/4) + 40n^{2/3}} \leq \rho_{n,2}$$

and

$$n^{-1} - 43n^{-4/3} \leq \rho_{n,2}.$$

The bound supplies a fully internal rate input for the two-arm risk. Combined with Theorem 4, it transfers the $43C_0(c)n^{-4/3}$ upper bound on the second-order improvement to every fixed multi-arm contrast. Together with the shrinkage construction in Theorem 5, the section therefore pins down the second-order minimax scale while retaining the contrast-weighted design identified in Theorem 3.

5 Finite-program certificates for rational and real contrasts

The preceding section identifies $n^{4/3}$ as the scale on which the second-order improvement can have a nondegenerate limit. This section represents that scale by finite linear programs. The programs are finite and all their data are exactly rational, so a bracket at a given (n, K) gives an exact rational certificate for that finite orbit game. The number of response-count rows is $\binom{n+2^K-1}{2^K-1}$, which grows quickly in both arguments. A finite bracket at a chosen (n, K) compares a proposed asymptotic design against the exact finite orbit game and identifies which response-count states drive the lower bound. The construction replaces the continuum of estimator values by a finite rational grid, solves a primal linear program over allocation-count orbits and grid actions, and reads a matching lower bound from the dual response-count multipliers. The first version is the three-arm diagnostic contrast $c^\dagger$; the general version covers rational contrasts and then transfers to real contrasts by continuity.

We begin with the three-arm program because it is the smallest setting in which the multi-arm response-type geometry is visible. The grid Γ_M records candidate estimator values, the variable $w_{r,x,g}$ couples an allocation count and observation with a grid action g, and the epigraph variable u bounds the risk uniformly over response-count vectors.

Definition 18. [def:k3-rational-lp] (Three-arm LP value $\Lambda^{\dagger}_{n,M}$) *For $K = 3$, $c^{\dagger} = (1, -1/2, -1/2)$, and*

$$\Gamma_M = \{-1 + j/M : 0 \leq j \leq 2M\},$$

the three-arm rational grid-action LP value $\Lambda^{\dagger}_{n,M}$ is the minimum of u over rational variables $\pi_r \geq 0$ and $w_{r,x,g} \geq 0$ satisfying

$$\sum_{r \in \mathcal{R}_{n,3}} \pi_r = 1,$$

$$\sum_{g \in \Gamma_M} w_{r,x,g} = \pi_r \quad \text{for every } r \in \mathcal{R}_{n,3} \text{ and } x \in \prod_{a \in \mathcal{A}_3} \{0, \ldots, r_a\},$$

and

$$\sum_{r,x,g} P_m(x \mid r) w_{r,x,g} \{g - \tau_{c^{\dagger}}(m)\}^2 \leq u \quad \text{for every } m \in \mathcal{M}_{n,3}.$$

The same construction applies to any rational nonzero zero-sum contrast. Its only additional bookkeeping is the contrast-dependent target range: the grid $\Gamma_{M,c}$ spans $[-L_c/2, L_c/2]$, so each feasible barycenter remains in the admissible range for the orbit game of Definition 7. The resulting value $\Lambda_{n,M}(c)$ is the rational finite program used throughout the rest of the section.

Definition 19. [def:rational-contrast-grid-lp] (Rational grid-action value $\Lambda_{n,M}(c)$) *For $K \geq 2$, $n, M \geq 1$, and a rational nonzero zero-sum contrast $c \in \mathbb{Q}^K$, define the rational contrast grid-action value $\Lambda_{n,M}(c)$ as the optimal value of the following finite linear program. Let*

$$h_c = L_c/2, \qquad \Gamma_{M,c} = \{-h_c + jh_c/M : j = 0, \ldots, 2M\}.$$

The program minimizes u over variables $\pi_r \geq 0$ and $w_{r,x,g} \geq 0$, subject to

$$\sum_{r \in \mathcal{R}_{n,K}} \pi_r = 1,$$

$$\sum_{g \in \Gamma_{M,c}} w_{r,x,g} = \pi_r \quad \text{for every } (r, x),$$

and

$$\sum_{r \in \mathcal{R}_{n,K}} \sum_{x} \sum_{g \in \Gamma_{M,c}} P_m(x \mid r) w_{r,x,g} \{g - \tau_c(m)\}^2 \leq u \quad \text{for every } m \in \mathcal{M}_{n,K}.$$

For later lower bounds, the dual variables are interpreted as a response-count prior. The next definition fixes the objective convention used to turn dual feasibility into a numerical certificate.

Definition 20. [synth_8] (Dual objective for the grid linear program) *For the grid-action linear program with primal constraints written in standard form $A\xi \leq b$ and objective u, let y denote the nonnegative vector of dual row multipliers satisfying the stationarity equations for the variables (u, π, w). Define*

$$\mathrm{obj}_{\mathrm{dual}}(y) = -\sum_i y_i b_i.$$

In the grid program, the restriction of y to the response-count risk rows is the response-count multiplier ν, and dual feasibility gives a lower certificate for the corresponding primal epigraph value.

The three-arm diagnostic program has an exact rational certificate at every sample size and grid resolution. The statement also records the mesh size $M = n^2$, which is fine enough for the grid error to vanish at the $n^{4/3}$ scale.

Proposition 1. [`prop:k3-lp-certificate`] (Three-arm LP certificate) *For the three-arm diagnostic contrast $c^\dagger = (1, -1/2, -1/2)$ and every pair of positive integers n and M, let $\rho_n^\dagger$ be the unrestricted three-arm minimax risk and let $\Lambda_{n,M}^\dagger$ be the three-arm rational grid-action LP value from Definition 18. Then*

$$\rho_n^\dagger \le \Lambda_{n,M}^\dagger \le \rho_n^\dagger + \frac{2}{M} + \frac{1}{M^2}, \qquad 0 \le \Lambda_{n,M}^\dagger \le 1.$$

At the resolution $M = n^2$, there exist a rational allocation mass π on allocation-count vectors, a joint design-action mass w, a rational epigraph value u, and rational response-count multipliers ν on count vectors $m : \mathcal{T}_3 \to \{0, \dots, n\}$ with $\sum_{t \in \mathcal{T}_3} m_t = n$, such that there are a rational grid weight $w^{\mathbb{Q}}$ and dual multipliers y satisfying

$$w = w^{\mathbb{Q}} \text{ after real embedding}, \qquad (\pi, w^{\mathbb{Q}}, u) \text{ is primal feasible}, \qquad y \text{ is dual feasible}, \qquad \nu = y_{\mathrm{risk}}, \qquad u = \mathrm{obj}_{\mathrm{d}}$$

and

$$u = \Lambda_{n,n^2}^\dagger.$$

Moreover, every binary response type $t \in \mathcal{T}_3$ is represented by some response-count vector m with total population size n and $m_t > 0$, and the grid-discretization error at the certificate scale satisfies

$$\lim_{k \to \infty} k^{4/3} \left(\frac{2}{k^2} + \frac{1}{k^4} \right) = 0.$$

Proposition 1 supplies a concrete finite certificate for the three-arm contrast. Its primal side gives a randomized invariant procedure whose risk is bounded by the grid value; its dual side gives response-count multipliers that certify a matching lower value up to the displayed mesh error. The representation clause records that every response type occurs in at least one feasible response-count vector, so the finite state space of the program covers the whole type space; which types carry positive dual mass is decided by the solution, not by this clause.

The lower endpoint in the general certificate is a Bayes risk computed at a fixed allocation count. The formula below is the posterior-variance form of that Bayes risk, expressed entirely with the orbit likelihood from Definition 6.

Definition 21. [`synth_7`] (Orbit Bayes risk at an allocation count) *For a response-count prior $\nu \in \Delta(\mathcal{M}_{n,K})$, a contrast c, and an allocation count $r \in \mathcal{R}_{n,K}$, define*

$$\mathrm{BayesRisk}_{n,c}(r, \nu) = \inf_{\delta_r} \sum_{m \in \mathcal{M}_{n,K}} \nu_m \sum_x P_m(x \mid r) \{\delta_r(x) - \tau_c(m)\}^2,$$

where $x \in \prod_{a \in \mathcal{A}_K} \{0, \dots, r_a\}$ and $\delta_r(x) \in [-L_c/2, L_c/2]$. Equivalently, with

$$D_{r,x,c}(\nu) = \sum_m \nu_m P_m(x \mid r), \qquad N_{r,x,c}(\nu) = \sum_m \nu_m P_m(x \mid r) \tau_c(m),$$

$$\mathrm{BayesRisk}_{n,c}(r, \nu) = \sum_m \nu_m \tau_c(m)^2 - \sum_{x : D_{r,x,c}(\nu) > 0} \frac{N_{r,x,c}(\nu)^2}{D_{r,x,c}(\nu)}.$$

The rational certificate theorem combines the primal grid solution, its barycenter estimator, and the dual Bayes-risk endpoint. The statement also gives the orbit counts, making the computational size of the certificate explicit.

Theorem 7. [`thm:rational-contrast-grid-certificate-sandwich`] (Rational grid certificate sandwich) *Fix natural numbers K, n, M, a rational nonzero zero-sum contrast $c \in \mathbb{Q}^K$, and suppose $n \geq 1$ and $M \geq 1$; here $\rho_n(c)$ is the labeled-schedule minimax risk of Definition 3. In the rational grid-action program of Definition 19, there exist a rational allocation certificate $\pi^{n,M,c}$, a joint design-action weight $w^{n,M,c}$, a rational epigraph value u, a rational response-count multiplier $\nu_{n,M,c} : \mathcal{M}_{n,K} \to \mathbb{Q}$, and an invariant estimator $\delta_{n,M,c}$ on allocation and observation counts such that:*

- ***(Exact rational certificate.)*** *There are rational grid weights $w^{\mathbb{Q}}$ and rational dual multipliers y such that $w^{n,M,c}$ is the real embedding of $w^{\mathbb{Q}}$, the tuple $(\pi^{n,M,c}, w^{\mathbb{Q}}, u)$ is primal feasible for the grid program, y is dual feasible, $\nu_{n,M,c}$ is the risk component of y, and u is the dual objective value of y.*
- ***(Barycenter estimator.)*** *For every allocation count $r \in \mathcal{R}_{n,K}$ and compatible observation x,*
$$\delta_{n,M,c}(r,x) = \begin{cases} \sum_{g \in \Gamma_{M,c}} g\, w^{n,M,c}_{r,x,g} / \pi^{n,M,c}_r, & \pi^{n,M,c}_r > 0, \\ 0, & \pi^{n,M,c}_r = 0. \end{cases}$$
- ***(Optimal value.)*** *The real value of u is the grid LP value:*
$$u = \Lambda_{n,M}(c).$$
- ***(Certificate sandwich.)*** *With*
$$B_{n,c}(\nu_{n,M,c}) = \inf\Big\{ v \in \mathbb{R} : \exists r \in \mathcal{R}_{n,K} \text{ with } v = \mathrm{BayesRisk}_{n,c}(r, \nu_{n,M,c}) \Big\},$$
and with $U_{n,M}(c)$ denoting the unrestricted labeled-schedule risk of $\delta_{n,M,c}$ under $\pi^{n,M,c}$, the certificates satisfy
$$B_{n,c}(\nu_{n,M,c}) \leq \rho_n(c) \leq U_{n,M}(c) \leq \Lambda_{n,M}(c) \leq B_{n,c}(\nu_{n,M,c}) + \frac{C_0(c)}{4M^2}.$$
- ***(Orbit counts.)*** *The finite sets in the certificate have cardinalities*
$$|\mathcal{M}_{n,K}| = \binom{n + 2^K - 1}{2^K - 1}, \qquad |\mathcal{R}_{n,K}| = \binom{n + K - 1}{K - 1},$$
and
$$\sum_{r \in \mathcal{R}_{n,K}} |\{x : x \text{ is compatible with } r\}| = \binom{n + 2K - 1}{2K - 1}.$$
- ***(Mesh-rate scaling.)*** *For every positive sequence a_k and every integer sequence $M_k \geq 1$, if*
$$\frac{a_k}{M_k^2} \longrightarrow 0,$$
then
$$a_k \frac{C_0(c)}{4M_k^2} \longrightarrow 0.$$

- ***(Asymptotic certificates.)*** *For every positive normalizing sequence a_k and every grid-resolution sequence $M_k \geq 1$ satisfying $a_k/M_k^2 \to 0$, the rational certificates at mesh M_k admit lower and upper endpoints B_k and U_k with*

  $$B_k \leq \rho_k(c) \leq U_k, \qquad a_k\{U_k - B_k\} \longrightarrow 0.$$

  Consequently, whenever $a_k d_k(c)$ has a finite limit ℓ, the scaled lower and upper certificate improvements both converge to that limit:

  $$a_k \left\{ \frac{C_0(c)}{k} - U_k \right\} \longrightarrow \ell, \qquad a_k \left\{ \frac{C_0(c)}{k} - B_k \right\} \longrightarrow \ell.$$

  At the second-order scale $a_k = k^{4/3}$, the specialization $M_k = k$ satisfies the mesh condition and yields

  $$k^{4/3}\{U_k - B_k\} \longrightarrow 0.$$

Theorem 7 turns each rational contrast into two computable endpoints. The upper endpoint is obtained by replacing the randomized grid action with its conditional barycenter, which preserves the relevant risk inequality by convexity. The lower endpoint is the best allocation-count Bayes risk under the dual response-count multiplier. Strong duality for the finite rational program then leaves only the grid mesh gap, and that gap is $o(k^{-4/3})$ when $M_k = k$.

The transfer from rational to real contrasts uses a direct Lipschitz comparison for the square root of the minimax risk. The distance is the half-ℓ_1 contrast distance, so the comparison is expressed in the same scale as the target interval.

Theorem 8. [`thm:contrast-risk-continuity`] (Contrast risk continuity) *Let K and n be integers, and let $c, c' \in \mathbb{R}^K$ be nonzero zero-sum treatment contrasts for the labeled-schedule game in Definition 3. Suppose that*

$$K \geq 2 \qquad \textit{and} \qquad n \geq 1.$$

Define the contrast distance

$$\eta(c, c') = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c_a - c'_a|.$$

Then the unrestricted labeled-schedule minimax risk satisfies

$$\left| \sqrt{\rho_n(c)} - \sqrt{\rho_n(c')} \right| \leq \eta(c, c').$$

The square-root form of Theorem 8 matches the geometry of squared-error risk: changing the contrast shifts every target by at most $\eta(c, c')$, and the minimax risk changes accordingly in root-risk units. This gives a stable way to carry finite rational certificates to nearby real contrasts.

It remains to name the exact rational primal-dual certificate as a reusable object. The definition records the feasibility, dual stationarity, and zero-gap conditions that are invoked in the transfer theorem.

Definition 22. [`synth_12`] (Exact rational primal-dual grid certificate) *Fix rational nonzero zero-sum $c \in \mathbb{Q}^K$ and integers $n, M \geq 1$. An exact rational primal-dual grid certificate for the program defining $\Lambda_{n,M}(c)$ consists of rational values $(\pi_r)_{r \in \mathcal{R}_{n,K}}$, rational values $(w_{r,x,g})$,* *a rational epigraph value u, and rational dual multipliers for the linear constraints, such that $\pi_r \geq 0$, $w_{r,x,g} \geq 0$, $\sum_r \pi_r = 1$, $\sum_{g\in\Gamma_{M,c}} w_{r,x,g} = \pi_r$ for every (r,x), and*

$$\sum_{r\in\mathcal{R}_{n,K}} \sum_{x} \sum_{g\in\Gamma_{M,c}} P_m(x \mid r) w_{r,x,g} \{g - \tau_c(m)\}^2 \leq u \quad \text{for every } m \in \mathcal{M}_{n,K}.$$

The dual multipliers are rational, nonnegative on inequality rows, satisfy the stationarity equations for the same rational linear program, and have dual objective equal to the primal objective u. The common value is $\Lambda_{n,M}(c)$.

The real-contrast result chooses rational approximants close enough that the continuity loss is dominated by the second-order scale. For an already rational contrast, the same statement specializes to a bracket with the n^{-2} mesh width.

Theorem 9. [`thm:real-contrast-grid-certificate-transfer`] (Real contrast certificate transfer) *Fix $K \geq 2$ and a nonzero zero-sum real contrast c on $\mathcal{A}_K$. For real contrasts c_1, c_2, write*

$$\eta(c_1, c_2) = \frac{1}{2} \sum_{a\in\mathcal{A}_K} |c_1(a) - c_2(a)|, \qquad C_0(c_1) = \frac{L_{c_1}^2}{4}.$$

A tuple (B, U, R^-, R^+) is a transferred certificate for (n, M, c, q), where $q \in \mathbb{Q}^K \setminus \{0\}$ is zero-sum, if there exist an invariant allocation-count distribution π, a rational grid-action weight w, a rational epigraph value u, a rational response-count prior ν, an invariant barycenter rule δ, a projected upper procedure, and a schedule prior such that the exact rational primal-dual grid certificate and barycenter identities hold for $q, \pi, w, u, \nu, \delta$, and

$$B = B_{n,q}(\nu), \qquad U = U_{n,M}(q; \pi, \delta), \qquad u = \Lambda_{n,M}(q),$$

$$B \leq \rho_n(q) \leq U \leq u \leq B + \frac{C_0(q)}{4M^2},$$

$$R^- = \left\{\max\left(0, \sqrt{B} - \eta(c, q)\right)\right\}^2, \qquad R^+ = \left(\sqrt{U} + \eta(c, q)\right)^2,$$

$$R^- \leq \rho_n(c) \leq R^+,$$

and

$$R^+ - R^- \leq \frac{5C_0(q)}{4M^2} + 4\sqrt{C_0(q)}\, \eta(c, q)\, n^{-1/2} + 3\eta(c, q)^2.$$

The projected upper procedure certifies worst-case risk at most R^+, and the schedule prior certifies Bayes risk at least R^-.

Then the following hold.

- ***(Finite-grid transfer.)*** *For every rational zero-sum contrast $q \in \mathbb{Q}^K \setminus \{0\}$ and every pair of integers $n, M \geq 1$, there exist $B, U, R^-, R^+ \in \mathbb{R}$ forming a transferred certificate for (n, M, c, q).*

- ***(Real-contrast approximation.)*** *There exist rational zero-sum contrasts $c^{(n)} \in \mathbb{Q}^K \setminus \{0\}$ and real sequences B_n, U_n, R_n^-, R_n^+ such that, for every arm $a \in \mathcal{A}_K$,*

$$c^{(n)}(a) \longrightarrow c(a),$$

$$n^{5/6}\eta(c, c^{(n)}) \longrightarrow 0,$$

for every $n \geq 1$, (B_n, U_n, R_n^-, R_n^+) *is a transferred certificate for* $(n, n, c, c^{(n)})$, *and*

$$n^{4/3}\big(R_n^+ - R_n^-\big) \longrightarrow 0.$$

- ***(Exact rational case.)*** *If a rational zero-sum contrast* $q \in \mathbb{Q}^K \setminus \{0\}$ *embeds coefficientwise as* c, *then for every* $n \geq 1$ *there exist* $B, U, R^-, R^+ \in \mathbb{R}$ *forming a transferred certificate for* (n, n, c, q) *and satisfying*

$$U - B \leq \frac{C_0(c)}{4n^2}.$$

Theorem 9 gives the certificate construction for every real contrast. The rational program supplies the inner bracket for the approximating contrast q, while Theorem 8 expands the bracket just enough to cover the target contrast c. The displayed bound separates the mesh contribution, the contrast-approximation contribution, and the quadratic approximation term, so choosing $M = n$ and $n^{5/6}\eta(c, c^{(n)}) \to 0$ produces a bracket whose width is negligible at the $n^{4/3}$ scale.

We close the section by recording the specialized three-arm sandwich. It packages the diagnostic program of Definition 18 in the same lower-upper form as the general rational theorem and gives the certificate comparison used in the three-arm analysis that follows.

Theorem 10. [`thm:k3-grid-certificate-sandwich`] (Three-arm grid sandwich) *For integers* $n \geq 1$ *and* $M \geq 1$, *in the three-arm rational grid program of Definition 18, there exist a rational allocation certificate* $\pi^{n,M}$, *a joint design–action weight* $w^{n,M}$, *a rational epigraph value* u, *a rational response-count multiplier* $\nu_{n,M}$, *and an estimator* $\delta_{n,M}$ *such that:*

- ***(Exact certificate.)*** *The tuple* $(\pi^{n,M}, w^{n,M}, u, \nu_{n,M})$ *is represented by a rational grid weight whose real embedding is* $w^{n,M}$, *a primal feasible point of the rational grid LP, and a dual feasible point whose risk component is* $\nu_{n,M}$ *and whose dual objective is* u.

- ***(Barycenter rule.)*** *For every allocation count* $r \in \mathcal{R}_{n,3}$ *and compatible observation* x, *the estimator is*

$$\delta_{n,M}(r, x) = \begin{cases} \displaystyle\sum_{g \in \Gamma_M} g \, \frac{w^{n,M}_{r,x,g}}{\pi^{n,M}_r}, & \pi^{n,M}_r > 0, \\ 0, & \pi^{n,M}_r = 0. \end{cases}$$

- ***(Grid value.)*** *The real value of* u *is* $\Lambda^\dagger_{n,M}$.

Define

$$B_n(\nu_{n,M}) = \inf_{r \in \mathcal{R}_{n,3}} \mathrm{BayesRisk}_{n,c^\dagger}(r, \nu_{n,M})$$

and

$$U_{n,M} = \sup_{m \in \mathcal{M}_{n,3}} \sum_{r \in \mathcal{R}_{n,3}} \pi^{n,M}_r \sum_x P_m(x \mid r)\{\delta_{n,M}(r, x) - \tau_{c^\dagger}(m)\}^2.$$

Then

$$B_n(\nu_{n,M}) \leq \rho^\dagger_n \leq U_{n,M} \leq \Lambda^\dagger_{n,M} \leq B_n(\nu_{n,M}) + \frac{1}{4M^2}.$$

Moreover, for every positive normalizing sequence a_k and every grid-resolution sequence M_k with $M_k \geq 1$ for all k, if

$$\frac{a_k}{M_k^2} \longrightarrow 0,$$

then

$$a_k \frac{1}{4M_k^2} \longrightarrow 0.$$

Finally, for every positive normalizing sequence a_k, every grid-resolution sequence M_k, every pair of real sequences L_k, U_k, and every real number C, if for every $k \geq 1$,

$$M_k \geq 1, \qquad L_k \leq \rho_k^\dagger \leq U_k, \qquad U_k - L_k \leq \frac{1}{4M_k^2},$$

and if

$$\frac{a_k}{M_k^2} \longrightarrow 0, \qquad a_k d_k(c^\dagger) \longrightarrow C,$$

then

$$a_k \left\{ \frac{C_0(c^\dagger)}{k} - L_k \right\} \longrightarrow C \quad \text{and} \quad a_k \left\{ \frac{C_0(c^\dagger)}{k} - U_k \right\} \longrightarrow C.$$

Theorem 10 specializes the rational programs to $c^\dagger$, where the mesh gap is $1/(4M^2)$. Thus any independently computed bracket with that width inherits the same second-order limit for the scaled improvement. The next section uses this certificate form to diagnose the three-arm geometry and to compare support-specific conclusions at the $n^{4/3}$ scale.

6 Three-arm diagnostics and support-specific conclusions

The finite certificates of Theorem 10 give a computable way to audit three-arm risks. We first use them to separate the full labeled experiment from a natural scalar signed-score reduction. For the diagnostic contrast $c^\dagger = (1, -1/2, -1/2)$, the scalar statistic records the sign-adjusted observed outcome, which retains the one-dimensional target direction and nothing else. The full experiment keeps a richer summary of the assignment and the outcomes: the rule exhibited below reads the number of units sent to the positively weighted arm, the signed outcome total within that arm, and the signed outcome total within the negatively weighted group. That summary already suffices, under the same independent allocation rule, to attain a strictly smaller worst-case risk than any rule based on the scalar score. The comparison is therefore between the scalar score and this richer assignment–outcome summary, which the exhibited rule reads while aggregating the two negatively weighted arms.

Proposition 2. [`prop:k3-scalar-score-not-minimax-preserving`] (Scalar score separation) *For $K = 3$, let $c^\dagger = (1, -1/2, -1/2)$, and let $\mathcal{D}^\star$ be the independent design assigning each unit with probabilities $q^\star = (1/2, 1/4, 1/4)$. For any $n \geq 1$, any schedule $z \in \mathcal{T}_3^n$, and any unit i, define*

$$Z_i = \operatorname{sign}(c^\dagger_{A_i})(2Y_i^{\text{obs}} - 1), \qquad \mu_i = t_{i,1} - \frac{t_{i,2} + t_{i,3}}{2}.$$

Then the following statements hold.

- ***(Score moments.)*** *For every $n \geq 1$, every $z \in \mathcal{T}_3^n$, and every unit i,*

$$\mathbb{E}_{\mathcal{D}^\star}(Z_i \mid z) = \mu_i, \qquad \operatorname{Var}_{\mathcal{D}^\star}(Z_i \mid z) = 1 - \mu_i^2.$$

- ***(Average-score risk.)*** *For every $n \ge 1$ and every $z \in \mathcal{T}_3^n$,*

$$\mathbb{E}_{\mathcal{D}^\star}\left[\left\{\frac{1}{n}\sum_i Z_i - \tau_{c^\dagger}(z)\right\}^2 \middle| z\right] = \frac{1}{n} - \frac{1}{n^2}\sum_i \mu_i^2.$$

- ***(Scalar minimax value.)*** *For the scalar experiment at $n = 3$ that retains (Z_1, Z_2, Z_3) and lets each μ_i range over the five feasible values, its minimax value is*

$$V_3^Z = 1 - \frac{\sqrt{3}}{2}.$$

- ***(Full-data certificate.)*** *There exists a clipped full-data estimator using the observed arm labels and outcomes under $\mathcal{D}^\star$ whose worst-case risk over $z \in \mathcal{T}_3^3$ is*

$$\frac{511653}{4000000}.$$

  Moreover,

$$\frac{511653}{4000000} < \frac{18213}{136000} < 1 - \frac{\sqrt{3}}{2}.$$

- ***(Scalar moment envelope.)*** *For every n, every real M, and every feasible three-arm scalar total $M = |\sum_i \mu_i|$, the minimum attainable value of $\sum_i \mu_i^2$ among scalar mean vectors with $|\sum_i \mu_i| = M$ is*

$$\begin{cases} M/2, & M \le n/2, \\ (3M - n)/2, & M > n/2. \end{cases}$$

Thus the scalar signed-score experiment has exact value V_3^Z at $n = 3$, while the full observed arm-label-and-outcome experiment under the same independent $q^\star$ design attains the smaller certified worst-case risk displayed above.

The diagnostic has two consequences for the rest of the section. First, scalar moment calculations remain useful for locating difficult schedules, because the displayed envelope identifies how the average-score risk depends on the feasible means. Second, the minimax audit is conducted in the full labeled experiment: Proposition 2 gives an explicit $n = 3$ comparison in which arm labels strictly lower the worst-case risk.

We now package the certificate objects used to compare rational finite programs and real contrasts. The cluster terminology names the primal weights, barycenter estimator, and dual response-count prior that move together in the finite orbit calculation.

Definition 23. [`synth_18`] (Exact rational grid-certificate cluster) *Fix a rational nonzero zero-sum contrast $c \in \mathbb{Q}^K$ and integers $n, M \ge 1$. An* exact rational grid-certificate cluster *at (n, M, c) consists of an exact rational primal-dual grid certificate for the grid-action program of Definition 19, in the sense of Definition 22 — rational primal weights $(\pi^{n,M,c}, w^{n,M,c})$ and the rational epigraph value $\Lambda_{n,M}(c)$ — together with the rational nonnegative response-count multiplier*

$$\nu_{n,M,c} \in \Delta(\mathcal{M}_{n,K})$$

read off its dual solution. Its derived objects are the barycenter estimator $\delta_{n,M,c}$ and the upper risk certificate $U_{n,M}(c)$, both built from $w^{n,M,c}$, and the posterior-mean lower Bayes-risk certificate $B_{n,c}(\nu_{n,M,c})$, built from $\nu_{n,M,c}$.

For real contrasts, the transfer argument uses the same half-ℓ_1 metric as Theorem 8. We record the notation before stating the support-three-and-larger result.

Definition 24. [synth_11] (Contrast distance) *For two real contrasts c_1, c_2 on $\mathcal{A}_K$, define*

$$\operatorname{dist}(c_1, c_2) = \eta(c_1, c_2) = \frac{L_{c_1 - c_2}}{2} = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c_1(a) - c_2(a)|.$$

Thus $\operatorname{dist}(c, c^{(n)}) = \eta(c, c^{(n)})$.

The next result collects the resolved second-order scale and the finite-certificate approximations for contrasts with at least three active arms. It uses the improvement $d_n(c) = C_0(c)/n - \rho_n(c)$ and the normalized quantity $x_n(c) = n^{4/3} d_n(c)$, matching the rate established in Theorem 5 and the lower comparison in Theorem 6.

Theorem 11. [thm:second-order-rate-and-certificate-frontier] (Second-order certificate frontier) *Fix $K \in \mathbb{N}$ and a contrast c on $\mathcal{A}_K$ as in Definition 3. Suppose that*

- ***(Arm count.)*** $2 \le K$.
- ***(Active support.)*** *With $S_c = \{a \in \mathcal{A}_K : c_a \neq 0\}$, one has $|S_c| \ge 3$.*

Define

$$x_n(c) = n^{4/3} d_n(c), \qquad d_n(c) = \frac{C_0(c)}{n} - \rho_n(c).$$

Then

$$0 < \liminf_{n \to \infty} x_n(c) \le \limsup_{n \to \infty} x_n(c) \le 43 C_0(c).$$

Moreover, $x_n(c)$ has at least one subsequential limit, and its subsequential-limit set

$$\{a \in \mathbb{R} : \text{there is a strictly increasing } \phi : \mathbb{N} \to \mathbb{N} \text{ with } x_{\phi(n)}(c) \to a\}$$

is compact.

For every positive sequence a_n, every $\beta \in \mathbb{R}$, and every $C > 0$, if

$$\frac{a_{\lfloor tn \rfloor}}{a_n} \to t^{\beta} \quad \text{for every } t > 0$$

and

$$a_n d_n(c) \to C,$$

then $\beta = 4/3$. In addition,

$$\exists\, a_n > 0,\ \exists\, C > 0 \text{ such that } \frac{a_{\lfloor tn \rfloor}}{a_n} \to t^{4/3} \text{ for every } t > 0 \text{ and } a_n d_n(c) \to C$$

$$\Longleftrightarrow$$

$$x_n(c) > 0 \text{ eventually and } \frac{x_{\lfloor tn \rfloor}(c)}{x_n(c)} \to 1 \text{ for every } t > 0.$$

For every rational contrast q on $\mathcal{A}_K$ whose real-valued contrast is c, there exist sequences G_n, P_n, and L_n such that G_n is an exact rational grid-certificate cluster for q, and

$$G_n - x_n(c) \to 0, \qquad P_n - G_n \to 0, \qquad L_n - G_n \to 0.$$

There also exist rational contrasts $c^{(n)}$*, certificate values* B_n, U_n*, and endpoints* R_n^-, R_n^+ *such that, for every arm* a*,*

$$c_a^{(n)} \to c_a,$$

and

$$n^{5/6}\,\mathrm{dist}(c, c^{(n)}) \to 0.$$

For every $n > 0$*, there are an exact rational primal-dual certificate, a grid barycenter estimator, and a rational prior for* $c^{(n)}$ *such that*

$B_n =$ *the lower certificate from the prior,* $\qquad U_n =$ *the upper certificate from the grid barycenter,*

$$R_n^- = \left(\max\left\{0, \sqrt{B_n} - \mathrm{dist}(c, c^{(n)})\right\}\right)^2, \qquad R_n^+ = \left(\sqrt{U_n} + \mathrm{dist}(c, c^{(n)})\right)^2,$$

and

$$R_n^- \le \rho_n(c) \le R_n^+.$$

The transferred normalized endpoint improvements are asymptotically coincident with $x_n(c)$*:*

$$n^{4/3}\left(\frac{C_0(c)}{n} - R_n^-\right) - x_n(c) \to 0, \qquad n^{4/3}\left(\frac{C_0(c)}{n} - R_n^+\right) - x_n(c) \to 0.$$

Finally, for the three-arm contrast $c^\dagger = (1, -1/2, -1/2)$*, there exists an estimator rule* $\widehat{\tau}$ *in the* $K = 3, n = 3$ *labeled game such that*

$$\sup_{z \in \mathcal{T}_3^3} R_3(\mathcal{D}^\star, \widehat{\tau}; z) < V_3^Z,$$

where $\mathcal{D}^\star$ *is the product design using the* $q^\star$ *allocation rule for* $c^\dagger$ *and* V_3^Z *is the three-observation scalar sign-score minimax value.*

The theorem gives the $n^{-4/3}$ scale an operational interpretation. The bounded positive liminf and finite limsup locate the improvement between two constants on that scale, while the regular-variation clause says that any convergent power normalization must use exponent 4/3. The finite-program clauses connect the asymptotic quantity to exactly solvable rational programs in principle: rational clusters approximate the normalized improvement, and rational approximants transfer the same bracketing statement to real contrasts through the distance defined in the setup. The three-arm scalar diagnostic is stated separately as a finite $n = 3$ illustration and is then used as the running example for the bracket construction.

It remains to isolate the support-two case. When the contrast has exactly one positive and one negative active arm, the multi-arm problem embeds the two-arm labeled experiment by assigning mass only to those active arms and rescaling the target by the half-range h_c.

Definition 25. [`synth_17`] (Sign-pair embedding and lift) *Assume* $|S_c| = 2$*, and write* $S_c = \{a_+, a_-\}$ *with* $c_{a_+} = h_c = L_c/2$ *and* $c_{a_-} = -h_c$*. The sign-schedule embedding* $\iota_c : \mathcal{T}_2^n \to \mathcal{T}_K^n$ *sends a two-arm schedule* $s = ((s_{i,+}, s_{i,-}))_{i=1}^n$ *to the* K*-arm schedule* $z = \iota_c(s)$ *with* $z_{i,a_+} = s_{i,+}$*,* $z_{i,a_-} = s_{i,-}$*, and fixed zero coordinates on arms outside* $\{a_+, a_-\}$*. The sign-pair lift of a two-arm labeled procedure* $p_2 = (\mathcal{D}_2, \widehat{\tau}_2)$ *is the* K*-arm procedure that maps the two labels* $(+, -)$ *to* (a_+, a_-)*, assigns no mass to other arms, and reports* $h_c\widehat{\tau}_2$ *after translating the observed active-arm data back to two-arm notation. Under this embedding,* $\tau_c(\iota_c(s)) = h_c\tau_{(1,-1)}(s)$*, and pushed-forward priors on* $\mathcal{T}_2^n$ *become priors on* $\mathcal{T}_K^n$*.*

The concluding theorem combines the support-two exact value, the order-attaining shrinkage rule, and the certificate statements used for finite audits. In applications with several randomized arms, such as multi-arm binary endpoint analyses in the style of ACTG 175 [Hammer et al., 1996], these clauses identify which support geometry determines the minimax comparison.

Theorem 12. [`thm:attainment-and-k3-certified-converse`] (Two-arm value transfer and finite brackets) *Fix $K \geq 2$ and a nonzero zero-sum contrast c on $\mathcal{A}_K$. Here $\rho_n(c)$ is the labeled-schedule minimax risk of Definition 3 and $G_n(c)$ the finite orbit-game value of Definition 7. The following statements hold.*

- ***(Two-active-arm value and saddle transfer.)*** *For every $n \geq 1$, if $|S_c| = 2$, then*

$$\rho_n(c) = C_0(c)\rho_{n,2}.$$

Moreover, there exist a two-arm orbit procedure q_2, a prior ν_2 on $\mathcal{M}_{n,2}$, a two-arm labeled procedure p_2, a prior prior_2 on $\mathcal{T}_2^n$, a K-arm labeled procedure p_K, and a prior prior_K on $\mathcal{T}_K^n$ such that

$$R^{\mathrm{orb}}_{(1,-1)}(q_2; m) \leq G_n((1,-1)) \quad \textit{for every } m \in \mathcal{M}_{n,2},$$

and, for every two-arm orbit procedure q_2',

$$G_n((1,-1)) \leq \mathbb{E}_{\nu_2}\left[R^{\mathrm{orb}}_{(1,-1)}(q_2'; m)\right].$$

The labeled procedure p_2 is the invariant labeled procedure induced by q_2, and the labeled prior satisfies

$$\mathrm{prior}_2(z) = \frac{\nu_2(\mathrm{counts}(z))}{|\{z' \in \mathcal{T}_2^n : \mathrm{counts}(z') = \mathrm{counts}(z)\}|}.$$

The K-arm procedure p_K is the sign-pair lift of p_2 along c, and prior_K is the pushforward of prior_2 under the sign-schedule embedding. These objects satisfy

$$R_n(p_2; z) \leq \rho_{n,2} \quad \textit{for every } z \in \mathcal{T}_2^n,$$

$$\rho_{n,2} \leq \mathbb{E}_{\mathrm{prior}_2}\left[R_n(p_2'; z)\right] \quad \textit{for every two-arm labeled procedure } p_2',$$

$$R_n(p_K; z) \leq \rho_n(c) \quad \textit{for every } z \in \mathcal{T}_K^n,$$

and

$$\rho_n(c) \leq \mathbb{E}_{\mathrm{prior}_K}\left[R_n(p_K'; z)\right] \quad \textit{for every } K\textit{-arm labeled procedure } p_K'.$$

- ***(Second-order scale.)*** *The normalized improvement $d_n(c) = C_0(c)/n - \rho_n(c)$ satisfies*

$$C_0(c)\kappa_c \leq \liminf_{n\to\infty} n^{4/3} d_n(c) \leq \limsup_{n\to\infty} n^{4/3} d_n(c) \leq 43 C_0(c).$$

- ***(Shrinkage attainment.)*** *There is an integer N such that, for every $n \geq N$, the explicit clipped-shrinkage procedure satisfies*

$$\sup_{z \in \mathcal{T}_K^n} R_n(\mathcal{D}^\star, \widehat{\tau}_n^{\mathrm{sh}}; z) \leq C_0(c)\left\{n^{-1} - \kappa_c n^{-4/3}\right\}.$$

- ***(Real-contrast transfer certificates.)*** *For every rational nonzero zero-sum contrast q, every $n \geq 1$, and every $M \geq 1$, there exist real numbers B, U, R^-, R^+ satisfying the real-contrast transfer-certificate conditions of Theorem 9 for (K, n, M, c, q). There also exist rational contrasts $c^{(n)}$ and real sequences B_n, U_n, R_n^-, R_n^+ such that*

$$c_a^{(n)} \to c_a \quad \text{for every } a \in \mathcal{A}_K,$$

$$n^{5/6}\, \eta(c, c^{(n)}) \to 0,$$

the tuple (B_n, U_n, R_n^-, R_n^+) satisfies the real-contrast transfer-certificate conditions of Theorem 9 for $(K, n, n, c, c^{(n)})$ whenever $n \geq 1$, and

$$n^{4/3}(R_n^+ - R_n^-) \to 0.$$

If a rational contrast q represents c exactly, then for every $n \geq 1$ there exist real numbers B, U, R^-, R^+ satisfying the same transfer-certificate conditions for (K, n, n, c, q) and

$$U - B \leq \frac{C_0(c)}{4n^2}.$$

- ***(Three-arm rational certificates.)*** *For every $n \geq 1$ and $M \geq 1$, there exist a three-arm grid allocation certificate π, a grid weight w, a rational epigraph value u, a rational response-count prior ν, and a barycenter estimator δ satisfying the exact primal-dual certificate and barycenter conditions of Theorem 10 for $c^\dagger$. These objects satisfy*

$$B_n(\nu) \leq \rho_n^\dagger \leq U_{n,M} \leq B_n(\nu) + \frac{1}{4M^2}.$$

The support-two clause is an exact reduction: the sign-pair embedding in Definition 25 transports both the upper procedure and the least favorable prior, so the K-arm value is the two-arm value scaled by $C_0(c)$. For larger supports, the shrinkage and certificate clauses give a common conclusion on the same second-order scale. The explicit procedure attains a $C_0(c)\{n^{-1} - \kappa_c n^{-4/3}\}$ bound eventually, while the rational and real-certificate brackets provide finite lower and upper checks around the unrestricted minimax risk. The specialized three-arm certificate closes the loop with Proposition 2: the same orbit-program machinery that audits $\rho_n^\dagger$ also records where the scalar signed-score reduction separates from the full observed-data experiment.

7 Discussion, limitations, and future work

7.1 Design implications for multi-arm binary experiments

The results give minimax benchmarks and finite-program tools for prespecified finite-population contrasts in multi-arm binary randomization. For any fixed nonzero zero-sum contrast, Theorem 3 identifies the first-order minimax constant and shows that the contrast-weighted assignment rule, paired with the projected contrast-weighted Horvitz–Thompson estimator, attains the first-order bound. Theorem 5 then shows that the same design can be paired with a clipped shrinkage estimator to improve the worst-case risk on the $n^{-4/3}$ scale. The model fixes a finite population and fixed K, encodes each unit by binary potential outcomes for the arms, uses zero-sum contrasts to target treatment differences, clips estimators to the feasible contrast range, and lets the adversary range over all labeled binary schedules; the active support records which arms enter the contrast. The finite orbit reduction in Theorem 1 and the finite LP constructions in Theorems 7 and 9 turn the minimax problem into finite optimization problems that give exact rational certificates and transferred finite bounds at fixed n.

A concrete setting is ACTG 175 [Hammer et al., 1996], which randomized HIV-1 infected patients among four regimens — zidovudine monotherapy, didanosine monotherapy, zidovudine–didanosine, and zidovudine–zalcitabine — and whose primary endpoint was a binary composite event: a decline of at least 50% in CD4 cell count, progression to AIDS, or death. Taking $K = 4$, $Y_i(a) \in \{0, 1\}$ for the indicator of that event under regimen a, and the enrolled cohort as the finite population, the model of Section 3 applies verbatim once a contrast c is fixed in the protocol, before outcomes are seen. With the contrast fixed, its active support determines which of the support-two or support-at-least-three conclusions applies. A comparison that places mass on one regimen against one other regimen falls under the support-two reduction in Theorem 12. A contrast comparing one regimen against an average of two or more alternatives uses the multi-arm conclusions of Theorems 11 and 12, with finite lower and upper certificates available through the rational or real-contrast programs.

The three-arm diagnostic case clarifies why the full observed-data experiment matters. Proposition 2 separates the signed-score reduction from the orbit experiment for the contrast $c^\dagger = (1, -1/2, -1/2)$, while Theorem 10 supplies matching finite-program certificates for the corresponding minimax risk. Together these statements say that collapsing observations to a scalar sign score can change the minimax value, and that the response-type orbit game retains the information needed for finite-n certification.

The logical spine is short, and it is worth stating once. Everything rests on the exact orbit equality of Theorem 1. Given that equality, Theorem 3 fixes the first-order constant, Theorem 5 supplies the attaining shrinkage rule, and Theorems 4 and 6 supply the matching converse; Theorems 7 to 9 then make the finite value computable, first for rational contrasts and then for all real ones. The remaining statements — Theorems 10 to 12 and Proposition 1 — assemble those ingredients for specific supports and for the three-arm example, and are read as consequences rather than as new mechanisms.

For design practice, the results suggest three complementary checks. First, Definition 14 and Theorem 3 identify the contrast-weighted randomization probabilities as the first-order benchmark. Second, Theorem 5 gives an explicit estimator achieving a uniform second-order improvement over that benchmark. Third, Theorems 7, 9 and 10 provide computable finite-sample brackets around the minimax risk, so the asymptotic design criterion can be checked against the exact finite orbit game.

7.2 Limitations and future work

Four limitations bound the reach of what is proved here, and the ACTG 175 reading above carries two more: the realized trial used a fixed equal-allocation randomization rather than the contrast-weighted $q^\star$, and its published analyses use time-to-event information that the binary composite endpoint discards. The potential outcomes are binary, so the results do not cover bounded or continuous responses. The number of arms K and the contrast c are fixed before the asymptotics, so nothing here speaks to regimes in which K grows with n. The second-order result is an order statement: the exponent $4/3$ is pinned down, but the constant in front of $n^{-4/3}$ is bracketed rather than identified, and the scaled improvement $n^{4/3} d_n(c)$ is not shown to converge. And the finite linear programs are established as objects, with their primal–dual structure proved; the paper reports no solved instance beyond the three-arm example at $n = 3$, and the number of response-count rows grows as $\binom{n+2^K-1}{2^K-1}$, so the size of a program at useful (n, K) is an open practical question.

Beyond those, the theory leaves two analytic programs that would sharpen the present finite-bracket and rate conclusions. The first concerns the limiting constant and the local structure of the hardest response-type configurations. The second concerns attainment: the form of procedures and priors that realize, or approximate, the limiting second-order value. Both questions are naturally tied to the orbit representation and to the contrast-weighted scores used in Theorem 5.

Remark 1. [`def:boundary-layer-handle`] (Boundary-layer problem) *A natural next question is whether the exact orbit game admits a boundary-layer limit obtained by enumerating binding response-type and allocation faces, expanding the discrete risk by a vector Stein identity within each stratum, using the $c^\dagger$ LP sequence to choose the normalizer, and passing the primal and dual games to a stratified diffusion-control or spectral value problem with face and corner compatibility conditions.*

The boundary-layer program would convert the finite orbit certificates into an analytic limiting problem. Its role is to identify the faces that bind at the second-order scale and to determine whether the normalized finite-n values converge to a sharp constant. The proposed Stein expansion and stratified limiting game parallel classical links between minimax estimation and differential-operator problems [Levit, 1986, Johnstone and MacGibbon, 1992], adapted here to the discrete response-type and allocation faces generated by multi-arm binary randomization.

Remark 2. [`def:attainment-handle`] (Attainment program) *A natural direction for future work is to start from the vector of centered arm scores induced by $q^\star$, derive an orbit-saddle posterior mean or feedback correction from the limiting value function, lift that correction to finite n, and discretize the squared ground state or dual occupation measure into a least-favorable prior on response-type counts.*

The attainment program would connect the constructive shrinkage estimator in Theorem 5 with the finite dual certificates in Theorems 7 and 10. It asks for a posterior-mean or feedback correction whose finite-n version matches the orbit game, together with a least-favorable prior on response-type counts. Related shrinkage ideas appear in recent work on randomized experiments [Rosenman and Hunter, 2026]; the orbit formulation here gives a finite-population route for studying their minimax constants and support-specific behavior.

Appendices

A Proofs and auxiliary lemmas

This appendix collects the auxiliary statements used by the paper's reduction, first-order, finite-program, continuity, and two-arm lower-bound arguments. The organizing principle is the finite orbit representation from Definition 7: permutation symmetry turns the original labeled experiment into a finite response-count game, while the two-arm lower bound uses a mixture over labeled binary schedules whose induced experiment depends on the assignment only through a one-dimensional score, so that the remaining assignment randomness can be averaged out without changing the Bayes risk. The scalar construction below supplies that interface before the Bayesian information inequality is applied to the resulting one-dimensional family.

The first auxiliary object is the canonical schedule law for the two-arm reduced scalar experiment. It lifts a distribution on effect-class triples to a distribution on complete response schedules, while recording the scalar observation induced by any assignment vector.

Definition 26. [`synth_13`] (Canonical two-arm scalar schedule kernel) *For a triple $\vartheta = (p_+, p_-, r_0)$ with $p_+ + p_- + r_0 = n$, the canonical schedule kernel K_ϑ is the probability law on complete two-arm schedules obtained by choosing uniformly an ordered partition (P, N, R) of $\{1, \dots, n\}$ with sizes (p_+, p_-, r_0), assigning response type $(1, 0)$ on P and $(0, 1)$ on N, and, independently for $i \in R$, drawing $H_i \sim \mathrm{Bernoulli}(1/2)$ and assigning type (H_i, H_i). Equivalently,*

$$K_\vartheta(z) = \mathbf{1}\{m_{(1,0)}(z) = p_+,\ m_{(0,1)}(z) = p_-,\ m_{(0,0)}(z) + m_{(1,1)}(z) = r_0\} \frac{p_+! p_-! r_0!}{n!} 2^{-r_0}.$$

For a schedule z and assignment vector $A \in \mathcal{A}_2^n$, define the transformed observed-score vector $S(z, A) \in \{0, 1\}^n$ by $S_i(z, A) = Y_i^{\mathrm{obs}}$ when $A_i = 1$ and $S_i(z, A) = 1 - Y_i^{\mathrm{obs}}$ when $A_i = 2$, and define the scalar observation

$$X(z, A) = \sum_{i=1}^{n} S_i(z, A).$$

Under K_ϑ, conditionally on ϑ, $X(z, A) = p_+ + B$ with $B \sim \mathrm{Binomial}(r_0, 1/2)$, whose probability mass is

$$\mathrm{binom}_{1/2}(r_0, k) = \binom{r_0}{k} 2^{-r_0}, \qquad k = 0, \dots, r_0.$$

The construction in Definition 26 is useful because it keeps the prior on complete schedules inside the original design problem while exposing a scalar sufficient statistic for the two-arm comparison. The next result states the exact mixture identity and the associated Rao–Blackwell comparison.

Lemma 1. [`lem:two-arm-scalar-prior-schedule-kernel`] (Scalar two-arm lift) *Fix $n \geq 0$, and let ν be any probability distribution on triples $\vartheta = (p_+, p_-, r_0)$ of nonnegative integers with $p_+ + p_- + r_0 = n$. There exists a probability distribution $\widetilde{\nu}$ on complete two-arm schedules $z \in \mathcal{T}_2^n$ such that:*

- ***(Canonical mixture.)*** *For every schedule z,*

  $$\widetilde{\nu}(z) = \sum_{\vartheta} \nu_\vartheta \, K_\vartheta(z),$$

  where K_ϑ is the canonical schedule kernel for the triple ϑ.

- ***(Scalar kernel.)*** *For every triple $\vartheta = (p_+, p_-, r_0)$, every assignment vector $A \in \mathcal{A}_2^n$, and every $x \in \{0, \dots, n\}$,*

  $$\sum_{z \in \mathcal{T}_2^n} \mathbf{1}\{X(z, A) = x\} K_\vartheta(z) = \sum_{k=0}^{r_0} \mathbf{1}\{x = p_+ + k\} \, \mathrm{binom}_{1/2}(r_0, k).$$

  Moreover, for every $A \in \mathcal{A}_2^n$,

  $$\sum_{z \in \mathcal{T}_2^n} K_\vartheta(z) = 1,$$

and, for every binary vector s with $\sum_i s_i = x$,

$$\sum_{z\in\mathcal{T}_2^n} \mathbf{1}\{X(z,A)=x,\ S(z,A)=s\}K_\vartheta(z) = \frac{\sum_{z\in\mathcal{T}_2^n}\mathbf{1}\{X(z,A)=x\}K_\vartheta(z)}{\binom{n}{x}}.$$

- ***(Rao–Blackwell domination.)*** *For every assignment law $\mathcal{D}$ on $\mathcal{A}_2^n$ and every estimator $\widehat{\tau}$, there is a function $f:\mathbb{N}\to\mathbb{R}$ such that*

$$\sum_{\vartheta=(p_+,p_-,r_0)} \nu_\vartheta \sum_{k=0}^{r_0} \mathrm{binom}_{1/2}(r_0,k)\left\{f(p_+ + k) - \frac{p_+ - p_-}{n}\right\}^2 \le \mathbb{E}_{\widetilde{\nu}}[R_n(\mathcal{D},\widehat{\tau};z)]\,.$$

- ***(Bayes-risk identity.)*** *For every assignment law $\mathcal{D}$ on $\mathcal{A}_2^n$,*

$$\inf_{\widehat{\tau}} \mathbb{E}_{\widetilde{\nu}}[R_n(\mathcal{D},\widehat{\tau};z)] = \inf_{f:\mathbb{N}\to\mathbb{R}} \sum_{\vartheta=(p_+,p_-,r_0)} \nu_\vartheta \sum_{k=0}^{r_0} \mathrm{binom}_{1/2}(r_0,k)\left\{f(p_+ + k) - \frac{p_+ - p_-}{n}\right\}^2.$$

Proof of Lemma 1. Fix a triple $\vartheta = (p_+, p_-, r_0)$. For a complete two-arm schedule z, write

$$P(z) = \{i : z_{i,1} = 1,\ z_{i,2} = 0\}, \qquad N(z) = \{i : z_{i,1} = 0,\ z_{i,2} = 1\}, \qquad R(z) = \{i : z_{i,1} = z_{i,2}\}.$$

Let

$$\Omega_\vartheta = \frac{n!}{p_+!\,p_-!\,r_0!}\,2^{r_0} = \binom{n}{p_+}\binom{n-p_+}{p_-}2^{r_0},$$

where the second equality uses $p_+ + p_- + r_0 = n$. Define

$$K_\vartheta(z) = \begin{cases} \Omega_\vartheta^{-1}, & |P(z)| = p_+,\ |N(z)| = p_-,\ |R(z)| = r_0, \\ 0, & \text{otherwise.} \end{cases}$$

When $n = 0$, the only triple is $(0,0,0)$, $\Omega_\vartheta = 1$, and all sums below are the corresponding one-point or empty sums.

1. Define the lifted prior by

$$\widetilde{\nu}(z) = \sum_\vartheta \nu_\vartheta K_\vartheta(z), \qquad z \in \mathcal{T}_2^n.$$

   Each summand is nonnegative. For fixed ϑ, there are $\binom{n}{p_+}\binom{n-p_+}{p_-}$ choices of the sets P, N, after which R is forced, and the r_0 equal-potential-outcome units have 2^{r_0} binary choices. Thus exactly Ω_ϑ schedules have positive K_ϑ-mass, each with mass Ω_ϑ^{-1}, so

$$\sum_{z\in\mathcal{T}_2^n} K_\vartheta(z) = 1.$$

   It follows that

$$\sum_{z\in\mathcal{T}_2^n} \widetilde{\nu}(z) = \sum_\vartheta \nu_\vartheta \sum_{z\in\mathcal{T}_2^n} K_\vartheta(z) = \sum_\vartheta \nu_\vartheta = 1,$$

   and the displayed definition is exactly the canonical mixture formula.

2. For $A \in \mathcal{A}_2^n$, define the transformed score vector and scalar count by

$$S_i(z,A) = \begin{cases} z_{i,1}, & A_i = 1, \\ 1 - z_{i,2}, & A_i = 2, \end{cases} \qquad X(z,A) = \sum_{i=1}^n S_i(z,A).$$

Fix a binary vector $s \in \{0,1\}^n$, and put $x_s = \sum_i s_i$. A schedule in the ϑ-fiber can produce $S(z,A) = s$ exactly by choosing $P(z)$ as a p_+-element subset of $\{i : s_i = 1\}$ and $N(z)$ as a p_--element subset of $\{i : s_i = 0\}$; after these choices the schedule is determined. Hence

$$\sum_{z \in \mathcal{T}_2^n} \mathbf{1}\{S(z,A) = s\} K_\vartheta(z) = \frac{\binom{x_s}{p_+}\binom{n-x_s}{p_-}}{\binom{n}{p_+}\binom{n-p_+}{p_-} 2^{r_0}}.$$

If $p_+ \le x_s \le p_+ + r_0$, the factorial identity

$$\binom{n}{x_s}\binom{x_s}{p_+}\binom{n-x_s}{p_-} = \binom{n}{p_+}\binom{n-p_+}{p_-}\binom{r_0}{x_s - p_+}$$

gives

$$\sum_{z \in \mathcal{T}_2^n} \mathbf{1}\{S(z,A) = s\} K_\vartheta(z) = \frac{\mathrm{binom}_{1/2}(r_0, x_s - p_+)}{\binom{n}{x_s}}.$$

If $x_s < p_+$ or $p_+ + r_0 < x_s$, the same mass is 0, since respectively $\binom{x_s}{p_+} = 0$ or $\binom{n-x_s}{p_-} = 0$. Summing over the $\binom{n}{x}$ vectors s with $x_s = x$ yields

$$\sum_{z \in \mathcal{T}_2^n} \mathbf{1}\{X(z,A) = x\} K_\vartheta(z) = \sum_{k=0}^{r_0} \mathbf{1}\{x = p_+ + k\}\, \mathrm{binom}_{1/2}(r_0, k).$$

Summing this identity over $x = 0, \ldots, n$ gives

$$\sum_{z \in \mathcal{T}_2^n} K_\vartheta(z) = \sum_{k=0}^{r_0} \mathrm{binom}_{1/2}(r_0, k) = 1.$$

Finally, if $\sum_i s_i = x$, then the event $\{S(z,A) = s\}$ already implies $\{X(z,A) = x\}$, and the preceding display gives

$$\sum_{z \in \mathcal{T}_2^n} \mathbf{1}\{X(z,A) = x,\ S(z,A) = s\} K_\vartheta(z) = \frac{\sum_{z \in \mathcal{T}_2^n} \mathbf{1}\{X(z,A) = x\} K_\vartheta(z)}{\binom{n}{x}}.$$

These identities are independent of A.

3. Fix an assignment law $\mathcal{D}$ and an estimator $\widehat{\tau}$. For a binary score vector s, let $y(A,s)$ be the observed-outcome vector obtained by inverting the preceding score map:

$$y_i(A,s) = \begin{cases} s_i, & A_i = 1, \\ 1 - s_i, & A_i = 2. \end{cases}$$

Set

$$T(A,s) = \widehat{\tau}(A, y(A,s)), \qquad \psi(\vartheta) = \begin{cases} (p_+ - p_-)/n, & n > 0, \\ 0, & n = 0. \end{cases}$$

Throughout the displays below, the expression $(p_+ - p_-)/n$ is read with this same total convention when $n = 0$. The risk under the lifted schedule prior is the prior risk in the induced score experiment:

$$\mathbb{E}_{\widetilde{\nu}}[R_n(\mathcal{D}, \widehat{\tau}; z)] = \sum_{\vartheta} \nu_\vartheta \sum_{A} \mathcal{D}(A) \sum_{z} K_\vartheta(z) \{T(A, S(z, A)) - \psi(\vartheta)\}^2.$$

For $x = 0, \dots, n$, define

$$\beta_\vartheta(x) = \sum_{k=0}^{r_0} \mathbf{1}\{x = p_+ + k\} \operatorname{binom}_{1/2}(r_0, k), \qquad \omega(A, s) = \frac{\mathcal{D}(A)}{\binom{n}{\sum_i s_i}}.$$

The scalar-kernel calculation gives the factorization

$$\mathcal{D}(A) \sum_{z} \mathbf{1}\{S(z, A) = s\} K_\vartheta(z) = \omega(A, s)\, \beta_\vartheta\left(\sum_i s_i\right).$$

Moreover, for each x,

$$\sum_{A} \sum_{s : \sum_i s_i = x} \omega(A, s) = \sum_{A} \mathcal{D}(A) \frac{\binom{n}{x}}{\binom{n}{x}} = 1.$$

Since the displayed weights form a probability distribution on each fiber $\{(A, s) : \sum_i s_i = x\}$, define the scalar average

$$g(x) = \sum_{A} \sum_{s : \sum_i s_i = x} \omega(A, s) T(A, s), \qquad x = 0, \dots, n.$$

For each fixed ϑ and x, Jensen's inequality on this finite fiber gives

$$\{g(x) - \psi(\vartheta)\}^2 \le \sum_{A} \sum_{s : \sum_i s_i = x} \omega(A, s) \{T(A, s) - \psi(\vartheta)\}^2.$$

Multiplying by $\nu_\vartheta \beta_\vartheta(x)$, summing over ϑ and x, and using the preceding factorization gives

$$\sum_{\vartheta} \nu_\vartheta \sum_{x=0}^{n} \beta_\vartheta(x) \{g(x) - \psi(\vartheta)\}^2 \le \mathbb{E}_{\widetilde{\nu}}[R_n(\mathcal{D}, \widehat{\tau}; z)].$$

Extend g to $f : \mathbb{N} \to \mathbb{R}$ by setting $f(x) = g(x)$ for $0 \le x \le n$ and $f(x) = 0$ for $x > n$. Expanding β_ϑ gives

$$\sum_{\vartheta = (p_+, p_-, r_0)} \nu_\vartheta \sum_{k=0}^{r_0} \operatorname{binom}_{1/2}(r_0, k) \left\{ f(p_+ + k) - \frac{p_+ - p_-}{n} \right\}^2 \le \mathbb{E}_{\widetilde{\nu}}[R_n(\mathcal{D}, \widehat{\tau}; z)].$$

4. It remains to identify the two infima. Let

$$\operatorname{clip}(u) = \max\{-1, \min\{1, u\}\}.$$

For every ϑ, the target $\psi(\vartheta)$ lies in $[-1, 1]$: this is immediate when $n = 0$, and for $n > 0$ the identities $p_+ + p_- + r_0 = n$ and nonnegativity give $0 \le p_+, p_- \le n$, so $-n \le p_+ - p_- \le n$, and division by the positive number n gives $-1 \le (p_+ - p_-)/n \le 1$. Therefore, for every scalar rule $f : \mathbb{N} \to \mathbb{R}$,

$$\{\mathrm{clip}(f(p_+ + k)) - \psi(\vartheta)\}^2 \le \{f(p_+ + k) - \psi(\vartheta)\}^2 .$$

The full-data estimator

$$\widehat{\tau}_f(A, y) = \mathrm{clip}\left(f\left(\sum_{i=1}^{n} \begin{cases} y_i, & A_i = 1, \\ 1 - y_i, & A_i = 2 \end{cases} \right) \right)$$

therefore has lifted Bayes risk at most the scalar risk of f. Taking the infimum over f gives

$$\inf_{\widehat{\tau}} \mathbb{E}_{\widetilde{\nu}}[R_n(\mathcal{D}, \widehat{\tau}; z)] \le \inf_{f:\mathbb{N}\to\mathbb{R}} \sum_{\vartheta=(p_+,p_-,r_0)} \nu_\vartheta \sum_{k=0}^{r_0} \mathrm{binom}_{1/2}(r_0, k) \left\{ f(p_+ + k) - \frac{p_+ - p_-}{n} \right\}^2 .$$

Conversely, the Rao–Blackwell domination from the previous step says that every full-data estimator has a scalar rule whose scalar risk is no larger than its lifted Bayes risk. Taking the infimum over $\widehat{\tau}$ gives the reverse inequality. Both risk classes are nonempty and bounded below by 0, since they contain the zero rule and all losses are squared losses, so the two infimum comparisons combine to the claimed identity.

□

Lemma 1 provides the paper-local Rao–Blackwell step behind the coarse two-arm converse in Theorem 6. For every prior on effect-class triples, it constructs a complete-schedule prior, identifies the binomial scalar experiment, and equates the induced Bayes risk with the best scalar decision rule. In this form the lower-bound calculation can be carried out on X while remaining a lower bound for the original assignment problem.

The final auxiliary ingredient is a Bayesian information inequality in the spirit of the van Trees bound [Van Trees, 1968, Gill and Levit, 1995]. The statement is written for an observation-dependent target $g(\theta, x)$, because the two-arm lower bound applies the inequality after conditioning and then evaluates the derivative of the induced conditional mean. The regularity clauses impose the smoothness, integrability, and boundary cancellation needed for the integration-by-parts argument underlying the information bound.

Lemma 2. [`lem:bayesian-information-inequality`] (Bayesian information inequality) *Let $\Theta = (\ell, u)$ be a scalar parameter interval with $\ell < u$, let $(\Omega, \mathcal{F})$ be a measurable observation space, and let μ be a σ-finite dominating measure. Let $w, dw : \mathbb{R} \to \mathbb{R}$, let $p_\theta(x)$, $dp_\theta(x)$, $g(\theta, x)$, and $dg(\theta, x)$ be real-valued measurable fields, and let $T : \Omega \to \mathbb{R}$. Define*

$$J(w) = \int_\ell^u \begin{cases} dw(\theta)^2/w(\theta), & w(\theta) > 0, \\ 0, & w(\theta) \le 0, \end{cases} d\theta$$

and

$$I(\theta) = \int \begin{cases} dp_\theta(x)^2/p_\theta(x), & p_\theta(x) > 0, \\ 0, & p_\theta(x) \le 0, \end{cases} d\mu(x).$$

Assume the following regularity conditions.

- ***(Prior.)*** *The density w is nonnegative, satisfies $\int_\ell^u w(\theta)\,d\theta = 1$, is C^1, has compact topological support contained in (ℓ, u), is positive on the interior of that support, and satisfies $dw(\theta) = w'(\theta)$ for every θ.*

- ***(Likelihood.)*** *For every $\theta \in (\ell, u)$, $p_\theta(x) \geq 0$ for μ-almost every x, $\int p_\theta(x)\,d\mu(x) = 1$, the fields p_θ and dp_θ are μ-integrable, and*

$$\left.\frac{d}{dt}\right|_{t=\theta} \int p_t(x)\,d\mu(x) = \int dp_\theta(x)\,d\mu(x) = 0.$$

  For μ-almost every x, $\theta \mapsto p_\theta(x)$ is absolutely continuous on every compact subinterval of (ℓ, u), and $dp_\theta(x) = \partial_\theta p_\theta(x)$ for Lebesgue-almost every $\theta \in (\ell, u)$.

- ***(Target.)*** *For μ-almost every x, $\theta \mapsto g(\theta, x)$ is absolutely continuous on every compact subinterval of (ℓ, u), and $dg(\theta, x) = \partial_\theta g(\theta, x)$ for Lebesgue-almost every $\theta \in (\ell, u)$.*

- ***(Pointwise derivatives.)*** *For (θ, x)-almost every point under Lebesgue measure on (ℓ, u) times μ, the derivatives $\partial_\theta p_\theta(x) = dp_\theta(x)$ and $\partial_\theta g(\theta, x) = dg(\theta, x)$ exist in the ordinary pointwise sense.*

- ***(Information and joint integrability.)*** *The prior information $J(w)$ is finite, $I(\theta)$ is finite for every $\theta \in (\ell, u)$, and*

$$0 < J(w) + \int_\ell^u I(\theta) w(\theta)\,d\theta.$$

  Moreover, all joint fields entering the risk, target derivative, squared prior score, squared likelihood score, score cross-product, squared total score, derivative-balance identity, and error-score identity are integrable with respect to $w(\theta)p_\theta(x)\,d\theta\,d\mu(x)$ on $[\ell, u] \times \Omega$.

- ***(Boundary product.)*** *For μ-almost every x, the map*

$$\theta \mapsto w(\theta)p_\theta(x)\{T(x) - g(\theta, x)\}$$

  is absolutely continuous on $[\ell, u]$, and its endpoint values vanish:

$$w(\ell)p_\ell(x)\{T(x) - g(\ell, x)\} = 0, \qquad w(u)p_u(x)\{T(x) - g(u, x)\} = 0.$$

Then

$$\int_\ell^u \int \{T(x) - g(\theta, x)\}^2 p_\theta(x) w(\theta)\,d\mu(x)\,d\theta \geq \frac{\left(\int_\ell^u \int dg(\theta, x) p_\theta(x) w(\theta)\,d\mu(x)\,d\theta\right)^2}{J(w) + \int_\ell^u I(\theta) w(\theta)\,d\theta}.$$

Proof of Lemma 2. 1. Let $\lambda_{\ell,u}$ denote Lebesgue measure restricted to $[\ell, u]$, and write

$$\varpi(\theta, x) = w(\theta)p_\theta(x), \qquad e(\theta, x) = T(x) - g(\theta, x)$$

for $(\theta, x) \in [\ell, u] \times \Omega$. Since $\ell < u$, the interval integral over $[\ell, u]$ is the integral with respect to $\lambda_{\ell,u}$. The stated joint integrability hypotheses justify Fubini and the following identifications:

$$\int_\ell^u \int e(\theta, x)^2 p_\theta(x) w(\theta)\,d\mu(x)\,d\theta = \iint e(\theta, x)^2 \varpi(\theta, x)\,d(\lambda_{\ell,u} \otimes \mu)(\theta, x),$$

and

$$\int_\ell^u \int dg(\theta,x)p_\theta(x)w(\theta)\,d\mu(x)\,d\theta = \iint dg(\theta,x)\varpi(\theta,x)\,d(\lambda_{\ell,u}\otimes\mu)(\theta,x).$$

The same change of measure rewrites $J(w)$ and the prior-averaged Fisher information in the guarded-score notation introduced below.

2. Define the guarded prior, likelihood, and joint scores on their full domains by

$$s_w(\theta) = \begin{cases} dw(\theta)/w(\theta), & w(\theta) > 0, \\ 0, & w(\theta) \le 0, \end{cases}$$

$$s_p(\theta,x) = \begin{cases} dp_\theta(x)/p_\theta(x), & p_\theta(x) > 0, \\ 0, & p_\theta(x) \le 0, \end{cases}$$

and

$$s(\theta,x) = \begin{cases} \{dw(\theta)p_\theta(x) + w(\theta)dp_\theta(x)\}/\varpi(\theta,x), & \varpi(\theta,x) > 0, \\ 0, & \varpi(\theta,x) \le 0. \end{cases}$$

For μ-almost every x, the section $\theta \mapsto p_\theta(x)$ is absolutely continuous, hence continuous on compact subintervals of (ℓ,u). Intersecting the full-measure nonnegativity sets over rational θ's and passing to limits gives

$$p_\theta(x) \ge 0 \quad \text{for every } \theta \in (\ell,u)$$

for μ-almost every x. Since the endpoints have $\lambda_{\ell,u}$-measure zero, $\varpi(\theta,x) \ge 0$ for $(\lambda_{\ell,u}\otimes\mu)$-almost every (θ,x).

3. The zero-density cases are controlled by local-minimum arguments. Because $w \ge 0$ and $dw = w'$, whenever $w(\theta) = 0$ one has $dw(\theta) = 0$. Similarly, at $(\lambda_{\ell,u}\otimes\mu)$-almost every (θ,x), if $p_\theta(x) = 0$, then θ is a local minimum of the nonnegative section $t \mapsto p_t(x)$, and therefore $dp_\theta(x) = 0$. Consequently the guarded likelihood score satisfies

$$s_p(\theta,x)p_\theta(x) = dp_\theta(x)$$

for almost every relevant point. The identity is immediate on the positive-density set and follows from the displayed zero-derivative fact on the zero-density set. The same case analysis will also be used below for the guarded total score: on the set $w(\theta)p_\theta(x) > 0$ it is the quotient defining the logarithmic derivative of $w(\theta)p_\theta(x)$, while if $w(\theta)p_\theta(x) = 0$, the preceding zero-derivative facts make the corresponding numerator vanish.

4. Define the derivative-balance field

$$B(\theta,x) = \{dw(\theta)p_\theta(x) + w(\theta)dp_\theta(x)\}e(\theta,x) - \varpi(\theta,x)dg(\theta,x).$$

For μ-almost every x, the map

$$\theta \longmapsto w(\theta)p_\theta(x)e(\theta,x)$$

is absolutely continuous on $[\ell,u]$ and has endpoint values equal to zero. At $\lambda_{\ell,u}$-almost every θ, the product rule gives its derivative as

$$\{dw(\theta)p_\theta(x) + w(\theta)dp_\theta(x)\}e(\theta,x) + w(\theta)p_\theta(x)(-dg(\theta,x)),$$

because $T(x)$ is constant in θ and $\partial_\theta e(\theta, x) = -dg(\theta, x)$. This displayed derivative is exactly $B(\theta, x)$. Thus the fundamental theorem for absolutely continuous functions gives

$$\int_\ell^u B(\theta, x)\, d\theta = w(u)p_u(x)e(u, x) - w(\ell)p_\ell(x)e(\ell, x) = 0.$$

Fubini and the assumed integrability of B yield

$$\iint B(\theta, x)\, d(\lambda_{\ell,u} \otimes \mu)(\theta, x) = 0.$$

Using the guarded total-score case analysis from the previous step,

$$e(\theta, x)s(\theta, x)\varpi(\theta, x) = e(\theta, x)\{dw(\theta)p_\theta(x) + w(\theta)dp_\theta(x)\}$$

almost everywhere: on $w(\theta)p_\theta(x) > 0$ this is the defining quotient multiplied by $w(\theta)p_\theta(x)$, and on $w(\theta)p_\theta(x) = 0$ the numerator vanishes by the zero-density derivative facts for w and p_θ. Hence

$$B(\theta, x) = e(\theta, x)s(\theta, x)\varpi(\theta, x) - dg(\theta, x)\varpi(\theta, x)$$

almost everywhere. Therefore

$$\iint e(\theta, x)s(\theta, x)\varpi(\theta, x)\, d(\lambda_{\ell,u} \otimes \mu) = \iint dg(\theta, x)\varpi(\theta, x)\, d(\lambda_{\ell,u} \otimes \mu).$$

5. For $\lambda_{\ell,u}$-almost every θ, the point lies in (ℓ, u), and the likelihood normalization and differentiation-under-the-integral assumptions give

$$\int dp_\theta(x)\, d\mu(x) = 0.$$

Combining this with $s_p(\theta, x)p_\theta(x) = dp_\theta(x)$ gives the conditional centering identity

$$\int s_p(\theta, x)p_\theta(x)\, d\mu(x) = 0.$$

Now expand the squared joint score. If $w(\theta) > 0$ and $p_\theta(x) > 0$, then $s(\theta, x) = s_w(\theta) + s_p(\theta, x)$. If $w(\theta) \leq 0$, the prior nonnegativity hypothesis gives $w(\theta) = 0$, so every term below contains the factor $w(\theta)$ or the guarded joint density and is zero. If $w(\theta) > 0$ but $p_\theta(x) \leq 0$, the almost-everywhere likelihood nonnegativity established above gives $p_\theta(x) = 0$, and again every term below vanishes by the guarded conventions. Hence, almost everywhere,

$$s(\theta, x)^2\varpi(\theta, x) = w(\theta)p_\theta(x)s_w(\theta)^2 + w(\theta)p_\theta(x)s_p(\theta, x)^2 + 2w(\theta)p_\theta(x)s_w(\theta)s_p(\theta, x).$$

Integrating, the cross term vanishes by the centering identity:

$$\iint w(\theta)p_\theta(x)s_w(\theta)s_p(\theta, x)\, d(\lambda_{\ell,u} \otimes \mu) = 0.$$

The prior term evaluates to

$$\iint w(\theta)p_\theta(x)s_w(\theta)^2\, d(\lambda_{\ell,u} \otimes \mu) = \int_\ell^u \begin{cases} dw(\theta)^2/w(\theta), & w(\theta) > 0, \\ 0, & w(\theta) \leq 0, \end{cases} d\theta = J(w),$$

using $\int p_\theta \, d\mu = 1$ for almost every interior θ. The likelihood term evaluates to

$$\iint w(\theta) p_\theta(x) s_p(\theta, x)^2 \, d(\lambda_{\ell,u} \otimes \mu) = \int_\ell^u I(\theta) w(\theta) \, d\theta.$$

Thus

$$\iint s(\theta, x)^2 \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes \mu) = J(w) + \int_\ell^u I(\theta) w(\theta) \, d\theta.$$

6. Finally, apply the weighted L^2 inequality with weight ϖ, functions e and s. Its proof is the standard quadratic argument: for every $\alpha \in \mathbb{R}$,

$$0 \le \iint \{e(\theta, x) - \alpha s(\theta, x)\}^2 \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes \mu),$$

and expanding the square, using the assumed integrability of the two square terms and the cross term, gives a nonnegative quadratic in α. Its discriminant is nonpositive, so

$$\left(\iint e(\theta, x) s(\theta, x) \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes \mu) \right)^2 \le \left(\iint e(\theta, x)^2 \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes \mu) \right) \left(\iint s(\theta, x)^2 \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes$$

Substituting the error–sensitivity identity and the information decomposition obtained above gives

$$\left(\iint dg(\theta, x) \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes \mu) \right)^2 \le \left(\iint e(\theta, x)^2 \varpi(\theta, x) \, d(\lambda_{\ell,u} \otimes \mu) \right) \left(J(w) + \int_\ell^u I(\theta) w(\theta) \, d\theta \right).$$

The denominator is strictly positive by assumption, so division yields the claimed inequality, and the identifications made at the outset translate it back to the statement's iterated-integral notation.

□

Lemma 2 is the reusable information step for the two-arm Bayes lower bound. The numerator measures the averaged derivative of the target, and the denominator combines the prior information with the likelihood information, matching the standard Bayesian Cramér–Rao structure associated with Van Trees [1968]. In the application to Theorem 6, Lemma 1 first reduces the complete-schedule prior to the scalar experiment, and Lemma 2 then lower-bounds the scalar Bayes risk used in the converse.

B Verification note

This appendix records what is machine-checked, in what environment, and what is not.

Every displayed definition, proposition, lemma and theorem in the paper falls into exactly one of two classes. A *matched* environment renders a Lean declaration in reader-facing mathematics: the displayed statement unfolds the declaration's structures and composite predicates into ordinary notation, with the same hypotheses and the same conclusion, and the proof reproduced in Section C follows the machine-checked proof of that declaration. The rendering is faithful to the declaration's content, not a transcription of Lean source. A *presentation-only* environment introduces notation or packages several matched objects for the reader; it carries no separate mathematical content and is not itself checked. The two lists below are exhaustive.

Environment. The development is written in Lean 4 against the toolchain `leanprover/lean4:v4.33.0`, on top of Mathlib and the project's own `Causalean` library. The results reported here were checked at commit `715add8e9475f5193526bce0c26baef559a47b68` of the source repository. The source of this development contains no `sorry` and no `axiom` declaration, and an axiom audit of all 26 declarations listed below reports exactly the three standard Lean axioms — `propext`, `Classical.choice` and `Quot.sound` — and nothing else. In particular, the paper contains no statement that is conditional on an unproved input, and no result is imported from a published source as a black box.

Matched environments. Each entry gives the paper environment, its Lean declaration, and the file that contains it.

- Definition 1: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.Arm` (`World.lean`).
- Definition 2: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.RespType` (`World.lean`).
- Definition 3: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.rhoN` (`Basic.lean`).
- Definition 6: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.orbitLik` (`Helpers/OrbitLikelihood.lean`).
- Definition 7: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.orbitGameValue` (`Basic.lean`).
- Theorem 1: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.exact_response_` (`T_exact_response_type_game.lean`).
- Theorem 2: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.multiarm_strict_` (`T_multiarm_strict_extension.lean`).
- Definition 14: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.contrastWeight` (`Basic.lean`).
- Theorem 3: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.first_order_sad` (`T_first_order_saddle.lean`).
- Theorem 4: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.embedded_two_ar` (`T_embedded_two_arm_converse.lean`).
- Theorem 5: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.universal_secon` (`T_universal_second_order_rate.lean`).
- Theorem 6: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.coarse_two_arm_` (`T_coarse_two_arm_minimax_lower.lean`).
- Definition 18: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.gridLPValueK3` (`Basic.lean`).
- Definition 19: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.gridLPValue` (`Basic.lean`).
- Proposition 1: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.k3_lp_certifi` (`T_k3_lp_certificate.lean`).

- Theorem 7: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.rational_contra` (`T_rational_contrast_grid_certificate_sandwich.lean`).
- Theorem 8: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.contrast_risk_c` (`T_contrast_risk_continuity.lean`).
- Theorem 9: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.real_contrast_g` (`T_real_contrast_grid_certificate_transfer.lean`).
- Theorem 10: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.k3_grid_certif` (`T_k3_grid_certificate_sandwich.lean`).
- Proposition 2: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.k3_scalar_sco` (`T_k3_scalar_score_not_minimax_preserving.lean`).
- Theorem 11: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.second_order_r` (`T_second_order_rate_and_certificate_frontier.lean`).
- Theorem 12: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.attainment_and` (`T_attainment_and_k3_certified_converse.lean`).
- Remark 1: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.activeFaceBounda` (`Basic.lean`).
- Remark 2: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.contrastScoreFee` (`Basic.lean`).
- Lemma 1: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.twoArmScalarPrio` (`Helpers/TwoArmScheduleKernel.lean`).
- Lemma 2: `CausalSmith.Experimentation.MultiarmSecondorderMinimaxFrontier.bayesianInformat` (`Helpers/BayesInformation.lean`).

Presentation-only environments. These introduce notation used by the matched statements and are not separately checked.

- Definition 4: Response-type count vector.
- Definition 5: Notation for response-type counts.
- Definition 8: Design-based squared-error risk.
- Definition 9: Orbit-procedure risk.
- Definition 10: Nonzero zero-sum contrast.
- Definition 11: Compatible observation-success counts.
- Definition 12: Invariant labeled procedure.
- Definition 13: Hull bounded-schedule domain and risk.
- Definition 15: Contrast-weighted independent design.
- Definition 16: Regular variation of a normalizer.
- Definition 17: Clipped-shrinkage estimator.
- Definition 20: Dual objective for the grid linear program.
- Definition 21: Orbit Bayes risk at an allocation count.
- Definition 22: Exact rational primal-dual grid certificate.

- Definition 23: Exact rational grid-certificate cluster.
- Definition 24: Contrast distance.
- Definition 25: Sign-pair embedding and lift.
- Definition 26: Canonical two-arm scalar schedule kernel.

What is not checked. The bibliographic comparisons in Section 2, including those to Sudijono et al. [2026] and Hull [2026], are positioning claims about the published literature and are not part of the formal development; no theorem in this paper takes a published result as a hypothesis. The applied reading of the ACTG 175 design in Section 7 is likewise interpretive. The numerical values reported for the three-arm example — the table T_F and the exact rational risk 511653/4000000 — are checked, but the search that produced the table is not: only the risk it attains is claimed.

C Proofs of the main results

Proof of Theorem 1. Write $\mathfrak{S}_n$ for the group of unit permutations. The proof packages five finite-sum facts.

1. Fix a labeled procedure $p = (\mathcal{D}, \widehat{\tau})$. For an assignment A and observed vector y, define

$$\mathsf{den}(A) = \sum_{\sigma \in \mathfrak{S}_n} \mathcal{D}(\sigma A), \qquad \mathsf{num}(A, y) = \sum_{\sigma \in \mathfrak{S}_n} \mathcal{D}(\sigma A)\, \widehat{\tau}(\sigma A, \sigma y),$$

where $(\sigma A)_i = A_{\sigma^{-1}(i)}$ and $(\sigma y)_i = y_{\sigma^{-1}(i)}$. Let

$$\bar{\mathcal{D}}(A) = \frac{1}{|\mathfrak{S}_n|} \sum_{\sigma \in \mathfrak{S}_n} \mathcal{D}(\sigma A),$$

and let

$$\bar{\widehat{\tau}}(A, y) = \begin{cases} \mathrm{clip}_{[-L_c/2, L_c/2]}\left(\mathsf{num}(A, y)/\mathsf{den}(A)\right), & \mathsf{den}(A) > 0, \\ 0, & \mathsf{den}(A) = 0. \end{cases}$$

Since $L_c = \sum_a |c_a| \geq 0$, the second value lies in the estimator interval. Reindexing the finite sum by left composition with any fixed unit permutation gives

$$\bar{\mathcal{D}}(\sigma_0 A) = \bar{\mathcal{D}}(A), \qquad \bar{\widehat{\tau}}(\sigma_0 A, \sigma_0 y) = \bar{\widehat{\tau}}(A, y),$$

so $\bar{p} = (\bar{\mathcal{D}}, \bar{\widehat{\tau}})$ is invariant.

For fixed A with $\mathsf{den}(A) > 0$, the weighted mean

$$\mathsf{num}(A, y)/\mathsf{den}(A)$$

is a convex combination of values in $[-L_c/2, L_c/2]$, hence the clipping is inactive. Therefore, for every real target value θ,

$$\mathsf{den}(A)\{\bar{\widehat{\tau}}(A, y) - \theta\}^2 \leq \sum_{\sigma \in \mathfrak{S}_n} \mathcal{D}(\sigma A)\{\widehat{\tau}(\sigma A, \sigma y) - \theta\}^2$$

by Jensen's inequality for the convex map $t \mapsto (t-\theta)^2$, applied to the weights $\mathcal{D}(\sigma A)$. If $\mathsf{den}(A) = 0$, the left side is zero and the right side is nonnegative, so the same inequality holds.

Apply this with $y = Y^{\mathrm{obs}}(z, A)$ and $\theta = \tau_c(z)$, then sum over A. The identities

$$\tau_c(\sigma z) = \tau_c(z), \qquad Y^{\mathrm{obs}}(\sigma z, \sigma A) = \sigma Y^{\mathrm{obs}}(z, A)$$

convert the right-hand side into the permutation average of labeled risks:

$$R_n(\bar{p}; z) \le \frac{1}{|\mathfrak{S}_n|} \sum_{\sigma \in \mathfrak{S}_n} R_n(p; \sigma z).$$

Let q be the orbit procedure obtained by collapsing the invariant procedure $\bar{p}$ to allocation and observation counts. For this collapsed procedure, the labeled and orbit descriptions put the same total probability on each allocation-count and observation-count fiber, so summing the squared loss over those fibers gives

$$R_n(q; \mathrm{counts}(z)) = R_n(\bar{p}; z).$$

Hence q is dominated by the averaged invariant representative.

2. An invariant labeled procedure is determined exactly by allocation-count and observation-count orbits. For an assignment $A \in \mathcal{A}_K^n$ and an observed binary vector y, define

$$\mathrm{alloc}(A)_a = \#\{i : A_i = a\}, \qquad \mathrm{succ}(A, y)_a = \#\{i : A_i = a,\ y_i = 1\}.$$

Thus $\mathrm{alloc}(A) \in \mathcal{R}_{n,K}$, and, when $\mathrm{alloc}(A) = r$, the vector $\mathrm{succ}(A, y)$ is an element of $\prod_{a \in \mathcal{A}_K} \{0, \ldots, r_a\}$. Write

$$\mathcal{O}_{\mathrm{alloc}}(r) = \{A' \in \mathcal{A}_K^n : \mathrm{alloc}(A') = r\}.$$

For an orbit procedure $q = (\pi, \delta)$, inflate it to a labeled procedure by

$$\mathcal{D}_q(A) = \frac{\pi_{\mathrm{alloc}(A)}}{|\mathcal{O}_{\mathrm{alloc}}(\mathrm{alloc}(A))|}, \qquad \widehat{\tau}_q(A, y) = \delta(\mathrm{alloc}(A), \mathrm{succ}(A, y)).$$

Conversely, if $p = (\mathcal{D}, \widehat{\tau})$ is invariant, choose one representative $A_r^\circ \in \mathcal{O}_{\mathrm{alloc}}(r)$ for each allocation count r, and one representative pair $(A_{r,x}^\circ, y_{r,x}^\circ)$ with $\mathrm{alloc}(A_{r,x}^\circ) = r$ and $\mathrm{succ}(A_{r,x}^\circ, y_{r,x}^\circ) = x$ for each observation count $x \in \prod_{a \in \mathcal{A}_K} \{0, \ldots, r_a\}$. Set

$$\pi_p(r) = |\mathcal{O}_{\mathrm{alloc}}(r)|\, \mathcal{D}(A_r^\circ), \qquad \delta_p(r, x) = \widehat{\tau}(A_{r,x}^\circ, y_{r,x}^\circ).$$

These definitions are independent of the representatives: equal allocation counts are exactly equality up to a unit permutation, and equal allocation and observation counts are exactly equality up to a simultaneous unit permutation of assignment and observed outcomes.

The two constructions are inverse to each other. Under either construction,

$$\mathcal{D}(A) = \frac{\pi_{\mathrm{alloc}(A)}}{|\mathcal{O}_{\mathrm{alloc}}(\mathrm{alloc}(A))|}, \qquad \widehat{\tau}(A, y) = \delta(\mathrm{alloc}(A), \mathrm{succ}(A, y)).$$

It remains to identify the regrouped labeled risk with the orbit risk. For a labeled schedule z, put $m_z = \mathrm{counts}(z)$ and define the observed labeled fiber

$$\mathcal{L}_z(r,x) = \{A \in \mathcal{A}_K^n : \mathrm{alloc}(A) = r, \mathrm{succ}(A, Y^{\mathrm{obs}}(z,A)) = x\}.$$

For a table $h \in \mathcal{H}(m_z, r, x)$, the number of assignments in $\mathcal{L}_z(r,x)$ having $h_{t,a}$ units of response type t assigned to arm a is

$$\prod_{t\in\mathcal{T}_K} \frac{m_{z,t}!}{\prod_{a\in\mathcal{A}_K} h_{t,a}!}.$$

Summing over the feasible contingency-table fiber and dividing by the allocation-orbit cardinality

$$|\mathcal{O}_{\mathrm{alloc}}(r)| = \frac{n!}{\prod_{a\in\mathcal{A}_K} r_a!}$$

gives

$$\frac{|\mathcal{L}_z(r,x)|}{|\mathcal{O}_{\mathrm{alloc}}(r)|} = \frac{\prod_{a\in\mathcal{A}_K} r_a!}{n!} \sum_{h\in\mathcal{H}(m_z,r,x)} \prod_{t\in\mathcal{T}_K} \frac{m_{z,t}!}{\prod_{a\in\mathcal{A}_K} h_{t,a}!} = P_{m_z}(x \mid r),$$

where the last equality is the likelihood formula in Definition 6. The target term is also constant on the response-count orbit,

$$\tau_c(z) = \frac{1}{n}\sum_i \sum_{a\in\mathcal{A}_K} c_a t_{i,a} = \frac{1}{n}\sum_{t\in\mathcal{T}_K} m_{z,t} \sum_{a\in\mathcal{A}_K} c_a t_a = \tau_c(m_z).$$

Therefore, regrouping first by allocation counts and then by observed fibers yields

$$\begin{aligned} R_n(p;z) &= \sum_r \sum_{A\in\mathcal{O}_{\mathrm{alloc}}(r)} \frac{\pi_r}{|\mathcal{O}_{\mathrm{alloc}}(r)|}\{\delta(r, \mathrm{succ}(A, Y^{\mathrm{obs}}(z,A))) - \tau_c(m_z)\}^2 \\ &= \sum_r \pi_r \sum_x \frac{|\mathcal{L}_z(r,x)|}{|\mathcal{O}_{\mathrm{alloc}}(r)|}\{\delta(r,x) - \tau_c(m_z)\}^2 \\ &= \sum_r \pi_r \sum_x P_{m_z}(x \mid r)\{\delta(r,x) - \tau_c(m_z)\}^2 = R_n(q; m_z). \end{aligned}$$

This proves the stated bijection and the exact equality of allocation-count, observation-count, and risk rules.

3. For every schedule z, the count vector satisfies

$$\mathrm{counts}(z)_t = \#\{i : t_i = t\}.$$

Thus the finite sum over labeled units may be regrouped by its response-type fibers:

$$\sum_i \sum_a c_a t_{i,a} = \sum_t \#\{i : t_i = t\} \sum_a c_a t_a.$$

Here and in the displayed target formula, $1/n$ denotes the real inverse $(n : \mathbb{R})^{-1}$, with the convention $0^{-1} = 0$. Scaling the regrouped identity by this factor gives

$$\tau_c(z) = \frac{1}{n}\sum_i \sum_a c_a t_{i,a} = \frac{1}{n}\sum_t m_t \sum_a c_a t_a = \tau_c(m), \qquad m = \mathrm{counts}(z).$$

4. The arm-count condition $K \geq 2$ supplies nonempty assignment and orbit spaces, while the zero clipped estimator supplies nonempty procedure classes. Risks are nonnegative, so both minimax values in Definitions 3 and 7 are formed over bounded-below finite risk ranges.

   First fix an orbit procedure q, and inflate it to its invariant labeled representative p_q. By the risk identity above, for every labeled schedule z,

   $$R_n(p_q; z) = R_n(q; \mathrm{counts}(z)) \leq \max_m R_n(q; m).$$

   Taking the worst case over z and then the infimum over q gives

   $$\rho_n(c) \leq G_n(c).$$

   Conversely, fix a labeled procedure p. The lossless symmetrization step gives an orbit procedure q_p such that, for every z,

   $$R_n(q_p; \mathrm{counts}(z)) \leq \frac{1}{|\mathfrak{S}_n|} \sum_{\sigma \in \mathfrak{S}_n} R_n(p; \sigma z) \leq \max_{z'} R_n(p; z').$$

   Every response-count vector is $\mathrm{counts}(z)$ for some schedule z, so

   $$\max_m R_n(q_p; m) \leq \max_{z'} R_n(p; z').$$

   Taking the infimum over p gives $G_n(c) \leq \rho_n(c)$. Hence

   $$\rho_n(c) = G_n(c).$$

5. Put the orbit game in finite squared-loss form with state space $\mathcal{M}_{n,K}$, design space $\mathcal{R}_{n,K}$, observation space

   $$\mathcal{X}(r) = \prod_{a \in \mathcal{A}_K} \{0, \ldots, r_a\},$$

   likelihood coefficient

   $$\mathsf{L}_{\mathrm{orb}}(m, r, x) = P_m(x \mid r),$$

   target $\tau_c(m)$, and action interval $[-L_c/2, L_c/2]$. The interval is nonempty because $L_c \geq 0$, and $\mathsf{L}_{\mathrm{orb}}(m, r, x) = P_m(x \mid r) \geq 0$ follows term by term from the factorial formula in Definition 6.

   The state set $\mathcal{M}_{n,K}$ is finite, the action set is the compact interval $[-L_c/2, L_c/2]$, and squared loss is bounded and convex in the action, so the game has a saddle point [Wald, 1950]. Applied to these data it gives an orbit procedure q_{sad} and a probability law ν on $\mathcal{M}_{n,K}$ such that

   $$R_n(q_{\mathrm{sad}}; m) \leq \inf_q \max_{m'} R_n(q; m') \quad \text{for every } m,$$

   and

   $$\inf_q \max_{m'} R_n(q; m') \leq \mathbb{E}_\nu[R_n(q; m)] \quad \text{for every orbit procedure } q.$$

   By Definition 7, the middle value is $G_n(c)$, giving

   $$R_n(q_{\mathrm{sad}}; m) \leq G_n(c), \qquad G_n(c) \leq \mathbb{E}_\nu[R_n(q; m)]\,.$$

□

Proof of Theorem 2. 1. For $K = 2$ and $c = (1, -1)$, the arm-count condition required in Theorem 1 is immediate. Its value-equality conclusion, applied to this two-arm contrast, gives

$$\rho_n(c) = G_n(c)$$

for every $n > 0$.

2. For the same contrast,

$$L_c = |1| + |-1| = 2, \qquad C_0(c) = \frac{L_c^2}{4} = 1.$$

This proves the asserted first-order constant.

3. Fix $n > 0$ and $m \in \mathcal{M}_{n,2}$. The four two-arm response types are

$$t_{00} = (0,0), \qquad t_{01} = (0,1), \qquad t_{10} = (1,0), \qquad t_{11} = (1,1).$$

They are pairwise distinct and every $t \in \mathcal{T}_2$ is exactly one of these four types: the pair (t_1, t_2) is one of

$$(0,0), \qquad (0,1), \qquad (1,0), \qquad (1,1).$$

Thus the orbit-target sum may be evaluated over this four-element list:

$$\tau_c(m) = \frac{1}{n} \sum_{t \in \{t_{00}, t_{01}, t_{10}, t_{11}\}} m_t \sum_{a \in \mathcal{A}_2} c_a t_a.$$

For $c = (1, -1)$, their contrast scores are

$$\sum_a c_a t_{00,a} = 0, \qquad \sum_a c_a t_{01,a} = -1, \qquad \sum_a c_a t_{10,a} = 1, \qquad \sum_a c_a t_{11,a} = 0.$$

Therefore

$$\tau_c(m) = \frac{m_{t_{10}} - m_{t_{01}}}{n} = \frac{p_+(m) - p_-(m)}{n},$$

which is the claimed two-arm target formula.

4. The four functions above differ by evaluating them at one of the two arms: t_{00} differs from t_{01} at the second arm, from t_{10} at the first arm, and from t_{11} at either arm; t_{01} and t_{10} differ at the first arm; t_{01} and t_{11} differ at the first arm; and t_{10} and t_{11} differ at the second arm. Hence they are pairwise distinct.

5. If $t \in \mathcal{T}_2$, then the pair (t_1, t_2) is one of

$$(0,0), \qquad (0,1), \qquad (1,0), \qquad (1,1),$$

so t is respectively t_{00}, t_{01}, t_{10}, or t_{11}. Thus these four types exhaust $\mathcal{T}_2$.

6. With these four response types kept as the elements of $\mathcal{T}_2$, the likelihood identity is exactly the specialization of Definition 6 to $K = 2$:

$$P_m(x \mid r) = \frac{\prod_{a\in\mathcal{A}_2} r_a!}{n!} \sum_{h\in\mathcal{H}(m,r,x)} \prod_{t\in\mathcal{T}_2} \frac{m_t!}{\prod_{a\in\mathcal{A}_2} h_{t,a}!}.$$

Here the fiber $\mathcal{H}(m, r, x)$ imposes the response-type margins m, allocation margins r, and arm-success totals x, as in Definition 6.

7. Let $n \geq 8$. Applying Theorem 4 with $K = 2$ and $c = (1, -1)$ gives

$$d_n(c) \leq 43\, C_0(c) n^{-4/3}.$$

Using $C_0(c) = 1$ from the calculation above yields

$$d_n(c) \leq 43\, n^{-4/3}.$$

8. For the lower second-order bound, apply Theorem 5 with $K = 2$ and $c = (1, -1)$. Let $\kappa = \kappa_c > 0$, and choose $N_\star$ large enough that the shrinkage procedure in that result satisfies, for every $n \geq N_\star$,

$$\sup_{z\in\mathcal{T}_2^n} R_n(\mathcal{D}^\star, \widehat{\tau}_n^{\mathrm{sh}}; z) \leq C_0(c)\{n^{-1} - \kappa n^{-4/3}\}.$$

Since $\rho_n(c)$ is the infimum of the same worst-case risk over all admissible procedures,

$$\rho_n(c) \leq C_0(c)\{n^{-1} - \kappa n^{-4/3}\}.$$

With $C_0(c) = 1$ and $d_n(c) = n^{-1} - \rho_n(c)$, this rearranges to

$$\kappa n^{-4/3} \leq d_n(c)$$

for every $n \geq N_\star$.

9. Now fix $K \geq 3$, n, and a contrast c. Since $K \geq 3$ implies $K \geq 2$, Theorem 1 applies and gives

$$\rho_n(c) = G_n(c).$$

The response-type set is the set of binary functions on $\mathcal{A}_K$. Because $|\mathcal{A}_K| = K$, its cardinality is

$$|\mathcal{T}_K| = |\{0,1\}^{\mathcal{A}_K}| = 2^K.$$

10. Fix N, L, U, a finite design $\mathcal{D}$ on two-arm assignments, and a measurable Hull estimator $\widehat{\tau}$. Set

$$v = \sup_Y R_{\mathrm{Hull}}(\mathcal{D}, \widehat{\tau}; Y),$$

where Y ranges over Hull schedules with parameters N, L, U. Choosing

$$\mathcal{D}' = \mathcal{D}, \qquad \widehat{\tau}' = \widehat{\tau}$$

preserves measurability and gives

$$v = \sup_Y R_{\mathrm{Hull}}(\mathcal{D}', \widehat{\tau}'; Y).$$

Thus v belongs to the displayed risk-value class.

11. Finally fix $K \geq 3$, $n > 0$, and $L < U$. First choose a number

$$v_{\text{sep}} \in [L, U], \qquad v_{\text{sep}} \neq 0, \qquad v_{\text{sep}} \neq 1.$$

The choice is by endpoint cases. If $L = 0$ and $U = 1$, take $v_{\text{sep}} = 1/2$. If $L = 0$ and $U \neq 1$, then the subcase $U = 0$ contradicts $L < U$, and in the remaining subcase take $v_{\text{sep}} = U$. If $L \neq 0$ and $L = 1$, then the subcases $U = 0$ and $U = 1$ contradict $L < U$, and in the remaining subcase take $v_{\text{sep}} = U$. If $L \neq 0$ and $L \neq 1$, take $v_{\text{sep}} = L$. In every surviving case this gives the displayed membership and exclusions.

Define the Hull schedule

$$Y_i(a) = v_{\text{sep}}, \qquad i = 1, \ldots, n, \quad a \in \mathcal{A}_2.$$

This schedule lies in $\mathcal{Y}^{\text{Hull}}_{n,L,U}$. Since $n > 0$, choose one unit i_0. If this embedded Hull schedule belonged to the binary K-arm schedule domain, then at the first arm its value would be the value of a binary potential outcome, hence either 0 or 1. But the same coordinate equals v_{sep}, contradicting $v_{\text{sep}} \notin \{0, 1\}$. Therefore

$$\mathcal{Y}^{\text{Hull}}_{n,L,U} \not\subseteq \mathcal{T}^n_K.$$

□

Proof of Theorem 3. 1. We first prove the limiting assertion. The hypotheses on c imply $K \geq 2$: if $K = 0$, there are no arm coefficients and the contrast is the zero function; if $K = 1$, the zero-sum identity forces the unique coefficient to be zero. Both alternatives contradict the assumed nonzero contrast.

For $n \geq 1$, set

$$d_n(c) = \frac{C_0(c)}{n} - \rho_n(c).$$

With $K \geq 2$, $n \geq 1$, and c nonzero and zero-sum, Theorem 4 applies. In particular, for every $n \geq 8$,

$$0 \leq d_n(c), \qquad d_n(c) \leq 43\, C_0(c) n^{-4/3}.$$

Since $C_0(c) = L_c^2/4 \geq 0$, for $n \geq 8$ we have

$$n\rho_n(c) - C_0(c) = n \left\{ \frac{C_0(c)}{n} - d_n(c) \right\} - C_0(c) = -n d_n(c),$$

and therefore

$$|n\rho_n(c) - C_0(c)| = n d_n(c) \leq 43\, C_0(c) n^{-1/3}.$$

The upper bound tends to 0, so the squeeze theorem gives

$$\lim_{n \to \infty} n\rho_n(c) = C_0(c).$$

2. We next prove the finite-sample risk bound. Fix $n \geq 1$ and a labeled schedule $z \in \mathcal{T}^n_K$. Since c is nonzero,

$$L_c = \sum_{a \in \mathcal{A}_K} |c_a| > 0.$$

Thus $q^\star_a = |c_a|/L_c$ is a probability vector, and $q^\star_a > 0$ exactly on the active support $S_c = \{a : c_a \neq 0\}$.

For each unit i and arm a, define the centered one-unit score on all of $\mathcal{A}_K$ by

$$\psi_i(a; z) = \begin{cases} \dfrac{c_a(t_{i,a} - 1/2)}{q^\star_a}, & q^\star_a > 0, \\ 0, & q^\star_a = 0. \end{cases}$$

Under the independent $q^\star$ design in Definition 14, the unprojected estimator is

$$\widetilde{\tau}_{\text{raw}}(A, z) = \frac{1}{n}\sum_{i=1}^{n} \psi_i(A_i; z).$$

For each i,

$$\mathbb{E}\{\psi_i(A_i; z)\} = \sum_{a \in S_c} q^\star_a \frac{c_a(t_{i,a} - 1/2)}{q^\star_a} = \sum_{a \in \mathcal{A}_K} c_a(t_{i,a} - 1/2) = \sum_{a \in \mathcal{A}_K} c_a t_{i,a},$$

where the last equality uses $\sum_a c_a = 0$. Hence

$$\mathbb{E}\{\widetilde{\tau}_{\text{raw}}(A, z)\} = \frac{1}{n}\sum_{i=1}^{n}\sum_{a \in \mathcal{A}_K} c_a t_{i,a} = \tau_c(z).$$

Also, since $t_{i,a} \in \{0, 1\}$,

$$(t_{i,a} - 1/2)^2 = \frac{1}{4},$$

and therefore

$$\mathbb{E}\{\psi_i(A_i; z)^2\} = \frac{1}{4}\sum_{a \in S_c} \frac{c_a^2}{q^\star_a} = \frac{1}{4}\sum_{a \in S_c} L_c |c_a| = \frac{L_c^2}{4} = C_0(c).$$

Thus

$$\operatorname{Var}\{\psi_i(A_i; z)\} = C_0(c) - \{\mathbb{E}\psi_i(A_i; z)\}^2 \le C_0(c).$$

Independence across units gives

$$\operatorname{MSE}\{\widetilde{\tau}_{\text{raw}}; z\} = \operatorname{Var}\{\widetilde{\tau}_{\text{raw}}(A, z)\} = \frac{1}{n^2}\sum_{i=1}^{n} \operatorname{Var}\{\psi_i(A_i; z)\} \le \frac{C_0(c)}{n}.$$

It remains to pass from the raw rule to the projected rule. For any two zero-sum contrasts $\mathfrak{c}_1, \mathfrak{c}_2$, put

$$\operatorname{dist}(\mathfrak{c}_1, \mathfrak{c}_2) = \frac{1}{2}\sum_{a \in \mathcal{A}_K} |(\mathfrak{c}_1)_a - (\mathfrak{c}_2)_a|.$$

For any binary response type, splitting the arms according to $t_a = 1$ and $t_a = 0$, the zero-sum identity for $\mathfrak{c}_1 - \mathfrak{c}_2$ yields

$$\left| \sum_{a \in \mathcal{A}_K} \{(\mathfrak{c}_1)_a - (\mathfrak{c}_2)_a\} t_a \right| \le \operatorname{dist}(\mathfrak{c}_1, \mathfrak{c}_2).$$

Averaging over the n units gives

$$|\tau_{\mathfrak{c}_1}(z) - \tau_{\mathfrak{c}_2}(z)| \le \operatorname{dist}(\mathfrak{c}_1, \mathfrak{c}_2).$$

Taking $\mathfrak{c}_1 = c$ and $\mathfrak{c}_2 = -c$, we have

$$\tau_{-c}(z) = -\tau_c(z), \qquad \operatorname{dist}(c, -c) = L_c,$$

so

$$|\tau_c(z)| \le \frac{L_c}{2}, \qquad \text{that is,} \qquad \tau_c(z) \in [-L_c/2, L_c/2].$$

For every $y \in [-L_c/2, L_c/2]$ and every real x,

$$\Big\{\operatorname{proj}_{[-L_c/2, L_c/2]}(x) - y\Big\}^2 \le (x-y)^2.$$

Indeed, if x lies in the interval there is equality; if $x < -L_c/2$, then $0 \le y + L_c/2 \le y - x$; and if $x > L_c/2$, then $0 \le L_c/2 - y \le x - y$. Squaring gives the displayed inequality in the two boundary cases. Applying it with $y = \tau_c(z)$ shows, pointwise in the assignment,

$$\Big\{\widehat{\tau}^{\star}_{\mathrm{cHT}}(A, Y^{\mathrm{obs}}) - \tau_c(z)\Big\}^2 \le \{\widetilde{\tau}_{\mathrm{raw}}(A, z) - \tau_c(z)\}^2.$$

Consequently,

$$R_n(\mathcal{D}^\star, \widehat{\tau}^{\star}_{\mathrm{cHT}}; z) \le \frac{C_0(c)}{n}.$$

Since z was arbitrary,

$$\max_{z \in \mathcal{T}_K^n} R_n(\mathcal{D}^\star, \widehat{\tau}^{\star}_{\mathrm{cHT}}; z) \le \frac{C_0(c)}{n}.$$

3. The two preceding steps are exactly the two components of the asserted conjunction: the first gives the asymptotic value of the unrestricted labeled-schedule minimax risk, and the second gives the finite-sample worst-case bound for the independent contrast-weighted design together with the projected contrast-weighted Horvitz–Thompson rule.

□

Proof of Theorem 4. Write

$$R_n(\mathcal{D}, \widehat{\tau}; z) = \mathbb{E}_{\mathcal{D}}\Big[\{\widehat{\tau}((A_i, Y_i^{\mathrm{obs}})_{i=1}^n) - \tau_c(z)\}^2\Big]$$

for the schedulewise risk appearing in Definition 3. For the two-arm contrast $(1, -1)$, write

$$\tau_2(y) = \frac{1}{n}\sum_{i=1}^n (y_{i,+} - y_{i,-}), \qquad R_{n,2}(\mathcal{D}_2, \widehat{v}; y) = \mathbb{E}_{\mathcal{D}_2}\left[\{\widehat{v}((B_i, V_i)_{i=1}^n) - \tau_2(y)\}^2\right],$$

for $y \in \mathcal{T}_2^n$, with $y_{i,+}, y_{i,-} \in \{0, 1\}$.

1. Define the sign groups and scale

$$S_c^+ = \{a \in \mathcal{A}_K : c_a > 0\}, \qquad S_c^- = \{a \in \mathcal{A}_K : c_a < 0\}, \qquad S_c^0 = \{a \in \mathcal{A}_K : c_a = 0\}, \qquad h_c = \frac{L_c}{2}.$$

Since c is nonzero and zero-sum, both S_c^+ and S_c^- are nonempty, and

$$\sum_{a \in S_c^+} c_a = -\sum_{a \in S_c^-} c_a = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c_a| = h_c, \qquad h_c^2 = C_0(c).$$

For $y \in \mathcal{T}_2^n$, embed it into a K-arm schedule $E_c y \in \mathcal{T}_K^n$ by

$$(E_c y)_{i,a} = \begin{cases} y_{i,+}, & a \in S_c^+, \\ y_{i,-}, & a \in S_c^-, \\ 0, & a \in S_c^0. \end{cases}$$

Then

$$\tau_c(E_c y) = \frac{1}{n} \sum_{i=1}^{n} \left\{ y_{i,+} \sum_{a \in S_c^+} c_a + y_{i,-} \sum_{a \in S_c^-} c_a \right\} = h_c \, \tau_2(y).$$

Fix an arbitrary K-arm procedure $(\mathcal{D}, \widehat{\tau})$. Let $\Omega = \{0,1\}^{\{1,\ldots,n\}}$, put the product law $\mathcal{L}$ on $(A, \omega) \in \mathcal{A}_K^n \times \Omega$ by drawing $A \sim \mathcal{D}$ and independent fair bits ω_i, and define the induced two-arm assignment

$$\sigma_{c,i}(A, \omega) = \begin{cases} +, & A_i \in S_c^+, \\ -, & A_i \in S_c^-, \\ +, & A_i \in S_c^0 \text{ and } \omega_i = 1, \\ -, & A_i \in S_c^0 \text{ and } \omega_i = 0. \end{cases}$$

For $v \in \{0,1\}^n$, define the inactive-zero observation vector

$$\widetilde{v}_i(A, v) = \begin{cases} v_i, & c_{A_i} \neq 0, \\ 0, & c_{A_i} = 0. \end{cases}$$

Let $\mathcal{D}_2^\circ$ be the pushforward of $\mathcal{L}$ under σ_c, so that

$$\mathcal{D}_2^\circ\{b\} = \mathcal{L}\{(A, \omega) : \sigma_c(A, \omega) = b\}, \qquad b \in \{+,-\}^n.$$

All conditional means below use the following totalization convention: on a cell with $\mathcal{D}_2^\circ\{b\} = 0$, the conditional mean is set to 0; this value is immaterial in risks because the cell has zero $\mathcal{D}_2^\circ$-mass. Define the induced two-arm estimator, on every $b \in \{+,-\}^n$ and $v \in \{0,1\}^n$, by

$$\widehat{v}^\circ(b, v) = \begin{cases} \mathbb{E}_{\mathcal{L}}\left[h_c^{-1} \widehat{\tau}(A, \widetilde{v}(A, v)) \,\middle|\, \sigma_c(A, \omega) = b \right], & \mathcal{D}_2^\circ\{b\} > 0, \\ 0, & \mathcal{D}_2^\circ\{b\} = 0. \end{cases}$$

Because $\widehat{\tau}$ is clipped to $[-h_c, h_c]$, this conditional mean lies in $[-1, 1]$. For each $y \in \mathcal{T}_2^n$, expanding the risk under the pushed-forward law and applying conditional Jensen's inequality gives

$$\begin{aligned} R_{n,2}(\mathcal{D}_2^\circ, \widehat{v}^\circ; y) &= \mathbb{E}_{\mathcal{D}_2^\circ}\big[\{\widehat{v}^\circ(B, (y_{i,B_i})_{i=1}^n) - \tau_2(y)\}^2\big] \\ &\le \mathbb{E}_{\mathcal{L}}\Big[\{h_c^{-1}\widehat{\tau}(A, \widetilde{v}(A, (y_{i,\sigma_{c,i}(A,\omega)})_{i=1}^n)) - \tau_2(y)\}^2\Big] \\ &= h_c^{-2} R_n(\mathcal{D}, \widehat{\tau}; E_c y). \end{aligned}$$

The last equality follows from the observation identity

$$\widetilde{v}_i(A, (y_{j,\sigma_{c,j}(A,\omega)})_{j=1}^n) = (E_c y)_{i,A_i}, \qquad i = 1, \ldots, n,$$

and the target identity $\tau_c(E_c y) = h_c \tau_2(y)$. Thus the construction supplies the statewise comparison

$$R_{n,2}(\mathcal{D}_2^\circ, \widehat{v}^\circ; y) \le h_c^{-2} R_n(\mathcal{D}, \widehat{\tau}; E_c y), \qquad y \in \mathcal{T}_2^n.$$

Taking the supremum over y, and using $E_c y \in \mathcal{T}_K^n$, yields the induced-procedure worst-case bound

$$\sup_{y \in \mathcal{T}_2^n} R_{n,2}(\mathcal{D}_2^\circ, \widehat{v}^\circ; y) \le h_c^{-2} \sup_{z \in \mathcal{T}_K^n} R_n(\mathcal{D}, \widehat{\tau}; z).$$

Since $\rho_{n,2}$ is the infimum of the left-hand worst-case risk over all two-arm procedures,

$$\rho_{n,2} \le h_c^{-2} \sup_{z \in \mathcal{T}_K^n} R_n(\mathcal{D}, \widehat{\tau}; z).$$

Multiplying by $h_c^2 = C_0(c)$ and taking the infimum over the arbitrary K-arm procedure yields

$$C_0(c)\rho_{n,2} \le \rho_n(c).$$

2. The identity $C_0(c) = L_c^2/4$ gives $C_0(c) \ge 0$. Consider the product design and projected estimator in Definition 14. For every schedule $z \in \mathcal{T}_K^n$, introduce the auxiliary totalized weights

$$\bar{q}_a^\star = \frac{|c_a|}{L_c}, \qquad a \in \mathcal{A}_K.$$

These agree with the probabilities $q_a^\star$ from Definition 14 on S_c, and $\bar{q}_a^\star = 0$ exactly when $c_a = 0$. For $a \in \mathcal{A}_K$, define

$$\psi_i(a; z) = \begin{cases} \dfrac{c_a\{z_{i,a} - 1/2\}}{\bar{q}_a^\star}, & \bar{q}_a^\star > 0, \\ 0, & \bar{q}_a^\star = 0. \end{cases}$$

Under the independent assignment law in Definition 14, equivalently with totalized weights $\bar{q}^\star$ on $\mathcal{A}_K$,

$$\mathbb{E}\{\psi_i(A_i; z)\} = \sum_{a \in S_c} \bar{q}_a^\star \frac{c_a(z_{i,a} - 1/2)}{\bar{q}_a^\star} = \sum_a c_a(z_{i,a} - 1/2) = \sum_a c_a z_{i,a},$$

because $\sum_a c_a = 0$. Also,

$$\mathbb{E}\{\psi_i(A_i; z)^2\} = \sum_{a \in S_c} \bar{q}_a^\star \frac{c_a^2(z_{i,a} - 1/2)^2}{(\bar{q}_a^\star)^2} = \frac{1}{4}\sum_{a \in S_c} \frac{c_a^2}{\bar{q}_a^\star} = \frac{1}{4}\sum_{a \in S_c} L_c|c_a| = C_0(c).$$

Thus, for

$$\overline{\psi}(A;z) = \frac{1}{n}\sum_{i=1}^{n}\psi_i(A_i;z),$$

the product structure gives

$$\mathbb{E}\{\overline{\psi}(A;z)\} = \tau_c(z), \qquad \mathrm{Var}\{\overline{\psi}(A;z)\} = \frac{1}{n^2}\sum_{i=1}^{n}\mathrm{Var}\{\psi_i(A_i;z)\} \le \frac{C_0(c)}{n}.$$

For each unit,

$$-\frac{L_c}{2} \le \sum_a c_a z_{i,a} \le \frac{L_c}{2},$$

so $\tau_c(z) \in [-L_c/2, L_c/2]$. Projection onto this interval cannot increase squared distance from $\tau_c(z)$, hence

$$R_n(\mathcal{D}^\star, \widehat{\tau}^\star_{\mathrm{cHT}}; z) \le \mathrm{Var}\{\overline{\psi}(A;z)\} \le \frac{C_0(c)}{n}.$$

Taking the maximum over z and then the infimum over procedures gives

$$\rho_n(c) \le \frac{C_0(c)}{n}.$$

3. Suppose $|S_c| = 2$. Let a_c^+ and a_c^- be the unique active arms with positive and negative coefficients. The preceding sign-sum identities give

$$c_{a_c^+} = h_c, \qquad c_{a_c^-} = -h_c, \qquad c_a = 0 \quad (a \notin \{a_c^+, a_c^-\}).$$

For any two-arm procedure $(\mathcal{D}_2, \widehat{v})$, lift a two-arm assignment $b \in \{+,-\}^n$ to the K-arm assignment A^b given by

$$A_i^b = \begin{cases} a_c^+, & b_i = +, \\ a_c^-, & b_i = -. \end{cases}$$

Let $\mathcal{D}_2^\uparrow$ be the pushforward of $\mathcal{D}_2$ under $b \mapsto A^b$. For a general K-arm assignment A, define the associated two-arm assignment

$$b_i^A = \begin{cases} +, & A_i = a_c^+, \\ -, & A_i \neq a_c^+, \end{cases}$$

and define the lifted estimator on its full domain by

$$\widehat{\tau}^\uparrow(A, v) = h_c \widehat{v}(b^A, v), \qquad A \in \mathcal{A}_K^n, \quad v \in \{0,1\}^n.$$

For a K-arm schedule z, define the active two-arm schedule

$$z_{i,+}^\downarrow = z_{i,a_c^+}, \qquad z_{i,-}^\downarrow = z_{i,a_c^-}.$$

The coefficient identities above give

$$\tau_c(z) = h_c \tau_2(z^\downarrow).$$

Moreover $b^{A^b} = b$ and the lifted observations satisfy

$$(z^{\downarrow})_{i,b_i} = z_{i,A_i^b}, \qquad i = 1, \dots, n.$$

Therefore, for every $z \in \mathcal{T}_K^n$, the lifted procedure has the statewise risk identity

$$R_n(\mathcal{D}_2^{\uparrow}, \widehat{\tau}^{\uparrow}; z) = h_c^2 R_{n,2}(\mathcal{D}_2, \widehat{\upsilon}; z^{\downarrow}).$$

It follows that

$$\sup_{z \in \mathcal{T}_K^n} R_n(\mathcal{D}_2^{\uparrow}, \widehat{\tau}^{\uparrow}; z) \le h_c^2 \sup_{y \in \mathcal{T}_2^n} R_{n,2}(\mathcal{D}_2, \widehat{\upsilon}; y).$$

Taking the infimum over two-arm procedures and using $h_c^2 > 0$ to pull the fixed scale through the infimum gives the minimax comparison

$$\rho_n(c) \le h_c^2 \rho_{n,2} = C_0(c)\rho_{n,2}.$$

Together with the lower bound obtained above,

$$\rho_n(c) = C_0(c)\rho_{n,2} \qquad \text{whenever } |S_c| = 2.$$

4. Now assume $n \ge 8$. From the upper bound,

$$d_n(c) = \frac{C_0(c)}{n} - \rho_n(c) \ge 0.$$

From the lower bound,

$$d_n(c) = C_0(c)n^{-1} - \rho_n(c) \le C_0(c)n^{-1} - C_0(c)\rho_{n,2} = C_0(c)\{n^{-1} - \rho_{n,2}\}.$$

Finally, Theorem 6 gives

$$n^{-1} - 43n^{-4/3} \le \rho_{n,2},$$

and therefore

$$n^{-1} - \rho_{n,2} \le 43n^{-4/3}.$$

Since $C_0(c) \ge 0$,

$$d_n(c) \le C_0(c)\{n^{-1} - \rho_{n,2}\} \le 43\, C_0(c) n^{-4/3}.$$

5. In the support-two case, the equality established above gives

$$d_n(c) = \frac{C_0(c)}{n} - \rho_n(c) = C_0(c)n^{-1} - C_0(c)\rho_{n,2} = C_0(c)\{n^{-1} - \rho_{n,2}\}.$$

This proves both asserted identities for $|S_c| = 2$.

□

Proof of Theorem 5. Fix $K \ge 2$ and a nonzero zero-sum contrast c. We prove the five asserted conclusions in the order displayed.

1. Let

$$\sigma_c(t) = \frac{\sum_{a\in\mathcal{A}_K} c_a t_a}{L_c/2}, \qquad t \in \mathcal{T}_K.$$

Since $c \neq 0$, choose a_0 with $c_{a_0} \neq 0$. For the response type t^{a_0} with $t^{a_0}_{a_0} = 1$ and $t^{a_0}_a = 0$ for $a \neq a_0$, the numerator is c_{a_0}, so the finite set of nonzero values $|\sigma_c(t)|$ is nonempty and has a positive minimum. This is exactly λ_c, hence $0 < \lambda_c$. Also, zero-summability gives

$$c_{a_0} = -\sum_{a\neq a_0} c_a, \qquad |c_{a_0}| \leq \sum_{a\neq a_0} |c_a|,$$

and therefore $2|c_{a_0}| \leq L_c$. Thus one admissible normalized nonzero value is at most one, and the minimum satisfies $\lambda_c \leq 1$. Finally,

$$A_c = \frac{\sqrt{\lambda_c}}{2} - \frac{\lambda_c}{16} > 0, \qquad B_c = \frac{7\lambda_c}{16} > 0,$$

because $0 < \lambda_c \leq 1$. Hence

$$\kappa_c = \frac{1}{2}\min\{A_c, B_c\} > 0.$$

2. The first-order rule $\widehat{\tau}^{\star}_{\mathrm{cHT}}$ from Definition 14 is controlled first before projection. For the degenerate case $n = 0$, there is a single empty schedule, its target is

$$\tau_c(\varnothing) = 0,$$

and the centered contrast-weighted sum is the empty sum, hence equals 0. The projection fixes 0, so the first-order procedure has risk

$$R_0(\mathcal{D}^\star, \widehat{\tau}^{\star}_{\mathrm{cHT}}; \varnothing) = 0.$$

Since squared-error risks are nonnegative, this gives $\rho_0(c) = 0$. The reciprocal estimate now begins with positive population size, so assume $n \geq 1$. For $z = (t_i)_{i=1}^n$, define, for an assignment vector $A \in \mathcal{A}_K^n$,

$$\xi_i(A_i, z) = \sum_{a\in S_c} c_a \mathbf{1}\{A_i = a\}\frac{t_{i,a} - 1/2}{q^\star_a}, \qquad \widetilde{\tau}_n(A, z) = \frac{1}{n}\sum_{i=1}^n \xi_i(A_i, z).$$

Under the independent $q^\star$-design, inverse-probability cancellation and $\sum_a c_a = 0$ give

$$\mathbb{E}\{\xi_i(A_i, z) \mid z\} = \sum_{a\in\mathcal{A}_K} c_a t_{i,a}, \qquad \mathbb{E}\{\xi_i(A_i, z)^2 \mid z\} = C_0(c),$$

and hence

$$\mathbb{E}\{\widetilde{\tau}_n(A, z) \mid z\} = \tau_c(z), \qquad \mathrm{Var}\{\widetilde{\tau}_n(A, z) \mid z\} \leq \frac{C_0(c)}{n}.$$

The target lies in the projection interval: for each response type t, the positive and negative coefficient masses are both $L_c/2$, so

$$-\frac{L_c}{2} \leq \sum_{a\in\mathcal{A}_K} c_a t_a \leq \frac{L_c}{2}, \qquad -\frac{L_c}{2} \leq \tau_c(z) \leq \frac{L_c}{2}.$$

For every $y \in [-L_c/2, L_c/2]$, interval projection satisfies

$$\left|\mathrm{proj}_{[-L_c/2,L_c/2]}(x) - y\right| \le |x - y|,$$

because the projection fixes x inside the interval and otherwise moves x to the nearer endpoint between x and y. Therefore

$$\left\{\mathrm{proj}_{[-L_c/2,L_c/2]}(\widetilde{\tau}_n(A, z)) - \tau_c(z)\right\}^2 \le \{\widetilde{\tau}_n(A, z) - \tau_c(z)\}^2 .$$

Taking expectations gives

$$R_n(\mathcal{D}^\star, \widehat{\tau}^\star_{\mathrm{cHT}}; z) \le \frac{C_0(c)}{n} \qquad \text{for every } z \in \mathcal{T}_K^n,$$

and taking the infimum over procedures gives $\rho_n(c) \le C_0(c)/n$ for $n \ge 1$. Together with the zero-unit branch, this proves the first-order envelope for every n.

3. Since

$$j_{\mathrm{tail}} \exp\{-j_{\mathrm{tail}}^{1/3}/8\} \longrightarrow 0 \qquad (j_{\mathrm{tail}} \to \infty),$$

multiplying by $4\sqrt{\lambda_c}$ gives

$$4\sqrt{\lambda_c}\, j_{\mathrm{tail}} \exp\{-j_{\mathrm{tail}}^{1/3}/8\} \longrightarrow 0.$$

Because $A_c > 0$, there is N_0 such that, for all integers $j_{\mathrm{tail}} \ge N_0$,

$$4\sqrt{\lambda_c}\, j_{\mathrm{tail}} \exp\{-j_{\mathrm{tail}}^{1/3}/8\} \le A_c.$$

Set $N_c = \max\{N_0, 1\}$. Then the stated tail side condition holds for every integer threshold variable at least N_c.

For the risk bound, fix $n \ge N_c$ and a schedule $z = (t_i)_{i=1}^n$. Under the independent $q^\star$-design, define

$$V_i = \frac{1}{L_c/2} \sum_{a \in S_c} c_a \mathbf{1}\{A_i = a\} \frac{t_{i,a} - 1/2}{q_a^\star}, \qquad X_n = \frac{1}{n}\sum_{i=1}^n V_i,$$

$$\Theta_n(z) = \frac{1}{n}\sum_{i=1}^n \sigma_c(t_i), \qquad Q_n(z) = \sum_{i=1}^n \sigma_c(t_i)^2.$$

Then

$$\mathbb{E}(X_n \mid z) = \Theta_n(z), \qquad \mathrm{Var}(X_n \mid z) = \frac{1}{n} - \frac{Q_n(z)}{n^2},$$

and the definition of λ_c gives

$$Q_n(z) \ge \lambda_c \sum_{i=1}^n |\sigma_c(t_i)| \ge \lambda_c n |\Theta_n(z)|.$$

For the tail estimate, write

$$\zeta_i(A) = V_i(A) - \sigma_c(t_i), \qquad X_n(A) - \Theta_n(z) = \frac{1}{n}\sum_{i=1}^n \zeta_i(A).$$

The product design makes the variables ζ_i independent conditional on z. Also $V_i(A) \in [-1,1]$ and $\mathbb{E}(V_i \mid z) = \sigma_c(t_i)$, so ζ_i has conditional mean zero and lies in the interval $[-1-\sigma_c(t_i), 1-\sigma_c(t_i)]$, whose length is 2. Hoeffding's lemma therefore gives, for every real α_{mgf},

$$\mathbb{E}[\exp\{\alpha_{\mathrm{mgf}}\zeta_i\} \mid z] \le \exp\{\alpha_{\mathrm{mgf}}^2/2\}.$$

Multiplying the moment-generating-function bounds across the independent units yields

$$\mathbb{E}\left[\exp\left\{\alpha_{\mathrm{mgf}} \sum_{i=1}^{n} \zeta_i\right\} \mid z\right] \le \exp\{n\alpha_{\mathrm{mgf}}^2/2\}.$$

Thus Chernoff's bound, with $\alpha_{\mathrm{mgf}} = \omega_{\mathrm{tail}}$ for $\omega_{\mathrm{tail}} > 0$ and with the endpoint $\omega_{\mathrm{tail}} = 0$ immediate, gives

$$\Pr\left\{\sum_{i=1}^{n} \zeta_i \ge n\omega_{\mathrm{tail}} \mid z\right\} \le \exp\{-n\omega_{\mathrm{tail}}^2/2\} \qquad (\omega_{\mathrm{tail}} \ge 0).$$

Applying the same argument to $-\sum_i \zeta_i$ and taking the union of the two one-sided events gives, for every $\omega_{\mathrm{tail}} \ge 0$,

$$\Pr\{|X_n - \Theta_n(z)| \ge \omega_{\mathrm{tail}} \mid z\} \le 2\exp\{-n\omega_{\mathrm{tail}}^2/2\}.$$

Put

$$b_n = n^{-1/3}, \qquad \varepsilon_n = \frac{\sqrt{\lambda_c}}{4} b_n, \qquad g_n(x) = \max\{-b_n, \min\{x, b_n\}\}.$$

The clipped-shrinkage estimator is

$$\widehat{\tau}_n^{\mathrm{sh}} = \frac{L_c}{2}\{X_n - \varepsilon_n g_n(X_n)\},$$

followed by projection onto $[-L_c/2, L_c/2]$; that projection fixes the displayed value. Indeed, fix a realized assignment and write $x = X_n$. The bounds

$$|x| \le 1, \qquad 0 \le b_n \le 1, \qquad 0 \le \varepsilon_n \le \frac{1}{4} \le 1$$

imply

$$|x - \varepsilon_n g_n(x)| \le 1.$$

To see this, use the three regions in the definition of g_n. If $x < -b_n$, then $g_n(x) = -b_n$ and

$$-1 \le x \le x + \varepsilon_n b_n \le x + b_n < 0 \le 1.$$

If $b_n < x$, then $g_n(x) = b_n$ and

$$-1 \le 0 < x - b_n \le x - \varepsilon_n b_n \le x \le 1.$$

In the remaining case, $-b_n \le x \le b_n$, so $g_n(x) = x$ and

$$|x - \varepsilon_n g_n(x)| = (1-\varepsilon_n)|x| \le |x| \le 1.$$

Thus, for every realized assignment,

$$-\frac{L_c}{2} \le \frac{L_c}{2}\{X_n - \varepsilon_n g_n(X_n)\} \le \frac{L_c}{2},$$

because $L_c/2 \ge 0$, and the final interval projection is inactive. Let

$$\mathcal{C}_n(z) = \mathrm{Cov}\{X_n, g_n(X_n) \mid z\}.$$

Since g_n is monotone,

$$\mathcal{C}_n(z) \ge 0,$$

and

$$\mathbb{E}\big[\{X_n - \varepsilon_n g_n(X_n) - \Theta_n(z)\}^2 \mid z\big] \le \mathrm{Var}(X_n \mid z) - 2\varepsilon_n \mathcal{C}_n(z) + \varepsilon_n^2 b_n^2.$$

If $|\Theta_n(z)| \le b_n/2$, define the far-tail event

$$\mathscr{E}_{\mathrm{tail}}(z) = \{A \in \mathcal{A}_K^n : |X_n(A) - \Theta_n(z)| \ge b_n/2\}.$$

The tail bound gives

$$\Pr\{\mathscr{E}_{\mathrm{tail}}(z) \mid z\} \le 2\exp\{-n^{1/3}/8\}.$$

On $\mathscr{E}_{\mathrm{tail}}(z)^c$, the local condition gives

$$|X_n(A)| \le |X_n(A) - \Theta_n(z)| + |\Theta_n(z)| \le b_n,$$

so $g_n(X_n(A)) = X_n(A)$. Since $|X_n(A)| \le 1$ and $b_n \le 1$, clipping also gives the uniform perturbation bound

$$|g_n(X_n(A)) - X_n(A)| \le 1.$$

Combining the two displays,

$$|g_n(X_n(A)) - X_n(A)| \le \mathbf{1}\{A \in \mathscr{E}_{\mathrm{tail}}(z)\}.$$

Moreover $|X_n(A) - \Theta_n(z)| \le |X_n(A)| + |\Theta_n(z)| \le 2$. Since $\mathbb{E}(X_n \mid z) = \Theta_n(z)$,

$$\mathcal{C}_n(z) = \mathrm{Var}(X_n \mid z) + \mathbb{E}[(X_n - \Theta_n(z))\{g_n(X_n) - X_n\} \mid z],$$

and therefore

$$\begin{aligned}\mathcal{C}_n(z) &\ge \mathrm{Var}(X_n \mid z) - \mathbb{E}[|X_n - \Theta_n(z)|\,|g_n(X_n) - X_n| \mid z] \\ &\ge \mathrm{Var}(X_n \mid z) - 2\Pr\{\mathscr{E}_{\mathrm{tail}}(z) \mid z\} \\ &\ge \mathrm{Var}(X_n \mid z) - 4\exp\{-n^{1/3}/8\}.\end{aligned}$$

Since $0 \le \varepsilon_n \le 1/4$,

$$1 - 2\varepsilon_n \ge 0, \qquad Q_n(z) \ge 0.$$

Substituting $\mathrm{Var}(X_n \mid z) = n^{-1} - Q_n(z)/n^2$ and the covariance lower bound into the preceding MSE inequality gives

$$\begin{aligned}&\mathbb{E}\big[\{X_n - \varepsilon_n g_n(X_n) - \Theta_n(z)\}^2 \mid z\big] \\ &\quad\le (1 - 2\varepsilon_n)\left\{\frac{1}{n} - \frac{Q_n(z)}{n^2}\right\} + 8\varepsilon_n \exp\{-n^{1/3}/8\} + \varepsilon_n^2 b_n^2 \\ &\quad\le \frac{1}{n} - \frac{2\varepsilon_n}{n} + 8\varepsilon_n \exp\{-n^{1/3}/8\} + \varepsilon_n^2 b_n^2,\end{aligned}$$

where the last line drops the nonpositive term

$$-(1-2\varepsilon_n)Q_n(z)/n^2$$

. Using

$$\frac{b_n}{n}=n^{-4/3}, \qquad \varepsilon_n=\frac{\sqrt{\lambda_c}}{4}b_n, \qquad \varepsilon_n^2 b_n^2=\frac{\lambda_c}{16}n^{-4/3},$$

we obtain

$$\mathbb{E}\big[\{X_n-\varepsilon_n g_n(X_n)-\Theta_n(z)\}^2 \mid z\big] \le \frac{1}{n}-A_c n^{-4/3}+2\sqrt{\lambda_c}\, b_n \exp\{-n^{1/3}/8\}.$$

The choice of N_c implies

$$2\sqrt{\lambda_c}\, b_n \exp\{-n^{1/3}/8\} \le \frac{A_c}{2}n^{-4/3},$$

so the local case gives

$$\mathbb{E}\big[\{X_n-\varepsilon_n g_n(X_n)-\Theta_n(z)\}^2 \mid z\big] \le \frac{1}{n}-\frac{A_c}{2}n^{-4/3} \le \frac{1}{n}-\kappa_c n^{-4/3}.$$

If $|\Theta_n(z)|>b_n/2$, then the lower bound on $Q_n(z)$ yields

$$\frac{Q_n(z)}{n^2} \ge \frac{\lambda_c}{2}n^{-4/3}.$$

Using $\mathcal{C}_n(z)\ge 0$ and the same identity for $\varepsilon_n^2 b_n^2$,

$$\mathbb{E}\big[\{X_n-\varepsilon_n g_n(X_n)-\Theta_n(z)\}^2 \mid z\big] \le \frac{1}{n}-\frac{7\lambda_c}{16}n^{-4/3} \le \frac{1}{n}-\kappa_c n^{-4/3}.$$

The two cases cover every schedule. Multiplying by $(L_c/2)^2=C_0(c)$ gives, uniformly in z,

$$R_n(\mathcal{D}^\star,\widehat{\tau}_n^{\mathrm{sh}};z) \le C_0(c)\{n^{-1}-\kappa_c n^{-4/3}\}.$$

Taking the supremum over schedules proves the asserted clipped-shrinkage risk bound for all $n\ge N_c$.

4. Since $\rho_n(c)$ is the minimax infimum, the shrinkage bound implies, for all $n\ge N_c$,

$$\rho_n(c) \le C_0(c)\{n^{-1}-\kappa_c n^{-4/3}\}.$$

Thus

$$d_n(c)=\frac{C_0(c)}{n}-\rho_n(c) \ge C_0(c)\kappa_c n^{-4/3},$$

and therefore

$$n^{4/3}d_n(c) \ge C_0(c)\kappa_c \qquad (n\ge N_c).$$

On the other hand, Theorem 4 gives, for all $n\ge 8$,

$$d_n(c) \le 43C_0(c)n^{-4/3},$$

so

$$n^{4/3}d_n(c) \le 43C_0(c) \qquad (n\ge 8).$$

The eventual lower and upper bounds imply

$$C_0(c)\kappa_c \le \liminf_{n\to\infty} n^{4/3}d_n(c) \le \limsup_{n\to\infty} n^{4/3}d_n(c) \le 43C_0(c).$$

5. Let a_n be a positive regularly varying sequence with index β, and suppose

$$a_n d_n(c) \longrightarrow C \qquad \text{with } C > 0.$$

Define the auxiliary positive sequence

$$\upsilon_n = \begin{cases} 1, & n = 0, \\ a_n n^{-4/3}, & n \geq 1. \end{cases}$$

For $n \geq 1$,

$$\upsilon_n \, n^{4/3} d_n(c) = a_n d_n(c).$$

The bounds from the previous step and the convergence to C imply that, eventually,

$$\frac{C}{2 \cdot 43 C_0(c)} \leq \upsilon_n \leq \frac{2C}{C_0(c)\kappa_c}.$$

Regular variation at the scale factor 2 gives

$$\frac{a_{2n}}{a_n} \longrightarrow 2^{\beta},$$

and hence

$$\frac{\upsilon_{2n}}{\upsilon_n} = \frac{a_{2n}}{a_n} 2^{-4/3} \longrightarrow 2^{\beta - 4/3}.$$

A positive sequence that is eventually bounded above and away from zero can have a dyadic ratio limit only equal to 1: indeed, for the dyadic logarithms

$$\chi_k = \log \upsilon_{2^k},$$

the increments $\chi_{k+1} - \chi_k$ converge to the logarithm of the dyadic ratio, while the eventual boundedness of υ_{2^k} makes χ_k eventually bounded; Cesaro averaging of

$$\frac{\chi_J - \chi_0}{J} = \frac{1}{J} \sum_{k=0}^{J-1} (\chi_{k+1} - \chi_k)$$

therefore forces that logarithm to be 0. Applying this to υ_n gives

$$2^{\beta - 4/3} = 1.$$

Since the map $r \mapsto 2^r$ is strictly increasing,

$$\beta = \frac{4}{3}.$$

$\square$

Proof of Theorem 6. Fix $n \geq 8$ and put

$$\alpha := n^{-1/3}.$$

1. Since $n > 0$, also $\alpha > 0$, and the real-power identities give

$$\alpha^3 = n^{-1}, \qquad n^{-2/3} = \alpha^2, \qquad n^{2/3} = n\alpha.$$

Moreover $\alpha \le 1/2$. Indeed, $n^{-1} \le 8^{-1}$. If $1/2 < \alpha$, then

$$\frac{1}{4} < \alpha^2$$

and, multiplying by the positive inequality $1/2 < \alpha$,

$$\frac{1}{8} < \alpha^3 = n^{-1},$$

contradicting $n^{-1} \le 8^{-1}$. The same identities also give

$$\frac{1}{\alpha^2} = n^{2/3}.$$

2. We first prove the sharper fractional lower bound in the auxiliary parameter α. Let

$$w_\alpha(\theta) = \frac{15}{8\alpha}\left(1 - \frac{4\theta^2}{\alpha^2}\right)^2 \mathbf{1}\{|\theta| \le \alpha/2\}.$$

This is a probability density on $(-1/2, 1/2)$, is positive on its support interior, and vanishes at the support endpoints. Direct differentiation on $|\theta| < \alpha/2$ gives

$$w'_\alpha(\theta) = -\frac{30\theta}{\alpha^3}\left(1 - \frac{4\theta^2}{\alpha^2}\right),$$

so

$$J(w_\alpha) = \int \frac{w'_\alpha(\theta)^2}{w_\alpha(\theta)}\, d\theta = \int_{-\alpha/2}^{\alpha/2} \frac{480\theta^2}{\alpha^5}\, d\theta = \frac{40}{\alpha^2}.$$

For fixed θ with $|\theta| \le \alpha/2$, draw independent two-arm response types $r_i = (r_{i,1}, r_{i,2}) \in \{0,1\}^2$ with

$$\Pr_\theta\{(1,0)\} = \frac{\alpha+\theta}{2}, \qquad \Pr_\theta\{(0,1)\} = \frac{\alpha-\theta}{2}, \qquad \Pr_\theta\{(0,0)\} = \Pr_\theta\{(1,1)\} = \frac{1-\alpha}{2}.$$

These masses are nonnegative and sum to one because $0 < \alpha \le 1/2$ and $|\theta| \le \alpha/2$. Let

$$\sigma_i := r_{i,1}, \qquad X(\sigma) := \sum_{i=1}^n \mathbf{1}\{\sigma_i = 1\}, \qquad \bar{R}(\sigma) := \frac{1}{n}\sum_{i=1}^n (2\mathbf{1}\{\sigma_i = 1\} - 1).$$

Then the marginal law of σ is the product Bernoulli law

$$p_\theta(\sigma) = \prod_{i=1}^n \left(\frac{1+\theta}{2}\right)^{\mathbf{1}\{\sigma_i=1\}} \left(\frac{1-\theta}{2}\right)^{\mathbf{1}\{\sigma_i=0\}},$$

with Fisher information

$$I_n(\theta) = \frac{n}{1-\theta^2} \le \frac{n}{1-\alpha^2/4}.$$

The finite-population contrast target for a response vector r is

$$\tau(r) = \frac{1}{n}\sum_{i=1}^{n}(r_{i,1} - r_{i,2}).$$

Conditional on $\sigma_i = 1$,

$$\mathbb{E}_\theta(r_{i,1} - r_{i,2} \mid \sigma_i = 1) = \frac{\alpha + \theta}{1 + \theta},$$

and conditional on $\sigma_i = 0$,

$$\mathbb{E}_\theta(r_{i,1} - r_{i,2} \mid \sigma_i = 0) = -\frac{\alpha - \theta}{1 - \theta}.$$

Thus the posterior target is

$$\gamma_\alpha(\theta, \sigma) := \mathbb{E}_\theta(\tau(r) \mid \sigma) = \theta + \omega_\alpha(\theta)\{\bar{R}(\sigma) - \theta\}, \qquad \omega_\alpha(\theta) := \frac{\alpha - \theta^2}{1 - \theta^2}.$$

For $|\theta| \leq \alpha/2$,

$$0 \leq \omega_\alpha(\theta) \leq \alpha.$$

At fixed σ,

$$\partial_\theta \gamma_\alpha(\theta, \sigma) = 1 - \omega_\alpha(\theta) + \frac{2\theta(\alpha - 1)}{(1 - \theta^2)^2}\{\bar{R}(\sigma) - \theta\}.$$

Since $\mathbb{E}_\theta \bar{R}(\sigma) = \theta$,

$$\mathbb{E}_\theta\{\partial_\theta \gamma_\alpha(\theta, \sigma)\} = 1 - \omega_\alpha(\theta) \geq 1 - \alpha.$$

For any scalar estimator $t : \{0, \ldots, n\} \to \mathbb{R}$, let t° be its clipping to $[-1, 1]$. Since $\tau(r) \in [-1, 1]$,

$$\{t^\circ(X(\sigma)) - \tau(r)\}^2 \leq \{t(X(\sigma)) - \tau(r)\}^2.$$

The conditional square completion gives, for each θ and σ,

$$p_\theta(\sigma)\{t^\circ(X(\sigma)) - \gamma_\alpha(\theta, \sigma)\}^2 \leq \sum_r \Pr_\theta(r)\mathbf{1}\{r_{i,1} = \sigma_i \ \forall i\}\{t^\circ(X(\sigma)) - \tau(r)\}^2.$$

Applying Lemma 2 to the finite observation space $\{0, 1\}^n$, the density w_α, likelihood p_θ, target γ_α, and estimator $t^\circ(X(\sigma))$, with the regularity hypotheses discharged by the displayed compact support and finite product likelihood, yields

$$\int w_\alpha(\theta) \sum_\sigma p_\theta(\sigma)\{t^\circ(X(\sigma)) - \gamma_\alpha(\theta, \sigma)\}^2 \, d\theta \geq \frac{(1 - \alpha)^2}{40/\alpha^2 + n/(1 - \alpha^2/4)}.$$

Combining this estimate with the square-completion and clipping inequalities shows that every scalar estimator has Bayes risk at least

$$\frac{(1 - \alpha)^2}{40/\alpha^2 + n/(1 - \alpha^2/4)}.$$

Let ν_α be the induced distribution on triples

$$\vartheta = (p_+, p_-, r_0),$$

where p_+ counts $(1,0)$ units, p_- counts $(0,1)$ units, and r_0 counts equal-potential-outcome units. Applied to ν_α, Lemma 1 gives a complete-schedule prior whose Rao–Blackwell scalar risk is dominated by every design-estimator Bayes risk. Hence, by Definition 3,

$$\frac{(1-\alpha)^2}{40/\alpha^2 + n/(1-\alpha^2/4)} \le \rho_{n,2}.$$

Using $n^{-2/3} = \alpha^2$, $\alpha = n^{-1/3}$, and $1/\alpha^2 = n^{2/3}$, this is exactly

$$\frac{(1-n^{-1/3})^2}{n/(1-n^{-2/3}/4) + 40n^{2/3}} \le \rho_{n,2}.$$

3. It remains to convert the fractional bound to the advertised coarse form. Since $0 < \alpha \le 1/2$,

$$0 < 1 - \frac{\alpha^2}{4}$$

and

$$\left(1 - \frac{\alpha^2}{4}\right)^{-1} \le 1 + \alpha.$$

Therefore

$$40\alpha + \left(1 - \frac{\alpha^2}{4}\right)^{-1} \le 1 + 41\alpha.$$

We claim

$$1 - 43\alpha \le \frac{(1-\alpha)^2}{40\alpha + (1-\alpha^2/4)^{-1}}.$$

The denominator is positive. If $1 - 43\alpha \ge 0$, then

$$(1-43\alpha)\left\{40\alpha + \left(1 - \frac{\alpha^2}{4}\right)^{-1}\right\} \le (1-43\alpha)(1+41\alpha) \le (1-\alpha)^2.$$

If $1 - 43\alpha < 0$, then the left-hand side after multiplication by the positive denominator is at most 0, while $(1-\alpha)^2 \ge 0$. This proves the claim in both cases.

Finally,

$$n^{-4/3} = n^{-1}\alpha$$

and

$$\frac{n}{1-\alpha^2/4} + 40n\alpha = n\left\{40\alpha + \left(1 - \frac{\alpha^2}{4}\right)^{-1}\right\}.$$

Thus

$$\begin{aligned} n^{-1} - 43n^{-4/3} &= n^{-1}(1-43\alpha) \\ &\le n^{-1}\frac{(1-\alpha)^2}{40\alpha + (1-\alpha^2/4)^{-1}} \\ &= \frac{(1-\alpha)^2}{n/(1-\alpha^2/4) + 40n\alpha}. \end{aligned}$$

The last expression is the fractional lower bound proved in the preceding step, so

$$n^{-1} - 43n^{-4/3} \le \rho_{n,2}.$$

Together with the first displayed bound, this proves both conclusions.

$\square$

Proof of Proposition 1. Let $n, M \geq 1$. We write $\rho_n^\dagger = \rho_n(c^\dagger)$ and use the identification of $\Lambda_{n,M}^\dagger$ with the three-arm specialization of the rational grid value from Definition 18.

1. Applying Theorem 7 with $K = 3$, contrast $c^\dagger$, and mesh M gives certificates $\pi^{n,M,c^\dagger}$, $w^{n,M,c^\dagger}$, u, $\nu_{n,M,c^\dagger}$, and a barycenter estimator $\delta_{n,M,c^\dagger}$ satisfying

$$B_{n,c^\dagger}(\nu_{n,M,c^\dagger}) \leq \rho_n^\dagger \leq U_{n,M}(c^\dagger) \leq \Lambda_{n,M}^\dagger \leq B_{n,c^\dagger}(\nu_{n,M,c^\dagger}) + \frac{C_0(c^\dagger)}{4M^2}.$$

   For $c^\dagger = (1, -1/2, -1/2)$,

$$L_{c^\dagger} = |1| + |-1/2| + |-1/2| = 2, \qquad C_0(c^\dagger) = \frac{L_{c^\dagger}^2}{4} = 1.$$

   The displayed sandwich first gives
$$\rho_n^\dagger \leq \Lambda_{n,M}^\dagger.$$
   It also gives
$$\Lambda_{n,M}^\dagger \leq B_{n,c^\dagger}(\nu_{n,M,c^\dagger}) + \frac{1}{4M^2} \leq \rho_n^\dagger + \frac{1}{4M^2}.$$
   Since $M \geq 1$,
$$\frac{1}{4M^2} \leq \frac{2}{M} + \frac{1}{M^2},$$
   and hence
$$\Lambda_{n,M}^\dagger \leq \rho_n^\dagger + \frac{2}{M} + \frac{1}{M^2}.$$

2. The unrestricted minimax risk is an infimum of suprema of squared-error risks, so

$$0 \leq \rho_n^\dagger.$$

   Combining this with $\rho_n^\dagger \leq \Lambda_{n,M}^\dagger$ yields

$$0 \leq \Lambda_{n,M}^\dagger.$$

3. It remains to record the elementary feasible point giving the upper bound 1. Let $r^\star = (n, 0, 0) \in \mathcal{R}_{n,3}$, let $0 \in \Gamma_M$ be the grid action corresponding to the index $j = M$, and define
$$\pi_r = \mathbf{1}\{r = r^\star\}, \qquad w_{r,x,g} = \mathbf{1}\{r = r^\star,\ g = 0\}, \qquad u_\star = 1.$$
   The normalization constraints in Definition 18 are then immediate. For a response-count vector $m \in \mathcal{M}_{n,3}$, set
$$\alpha(\zeta) = \zeta_1 - \frac{1}{2}\zeta_2 - \frac{1}{2}\zeta_3, \qquad \zeta \in \mathcal{T}_3.$$
   Because each $\zeta_a \in \{0, 1\}$, one has $-1 \leq \alpha(\zeta) \leq 1$. Therefore

$$\tau_{c^\dagger}(m) = \frac{1}{n} \sum_{\zeta \in \mathcal{T}_3} m_\zeta\, \alpha(\zeta) \in [-1, 1], \qquad \tau_{c^\dagger}(m)^2 \leq 1.$$

Using the orbit likelihood of Definition 6, whose finite probabilities in x are nonnegative and sum to 1, the risk row for the displayed feasible point is

$$\sum_{r,x,g} P_m(x \mid r) w_{r,x,g}\{g - \tau_{c^\dagger}(m)\}^2 = \sum_x P_m(x \mid r^\star)\tau_{c^\dagger}(m)^2 = \tau_{c^\dagger}(m)^2 \leq 1.$$

Thus $u_\star = 1$ is feasible, and the minimizing value satisfies

$$\Lambda^\dagger_{n,M} \leq 1.$$

4. Apply Theorem 7 again, now with mesh $M = n^2$. Since $n \geq 1$, also $n^2 \geq 1$. The exact rational certificate clause gives a rational allocation mass π, real joint design-action mass w, rational epigraph value u, and rational response-count multipliers ν, packaged as an exact primal-dual certificate. Unpacking that certificate supplies rational grid weights $w^{\mathbb{Q}}$ and dual multipliers y such that

$$w = w^{\mathbb{Q}} \text{ after real embedding}, \qquad (\pi, w^{\mathbb{Q}}, u) \text{ is primal feasible}, \qquad y \text{ is dual feasible},$$

$$\nu = y_{\mathrm{risk}}, \qquad u = \mathrm{obj}_{\mathrm{dual}}(y).$$

The optimal-value clause gives
$$u = \Lambda^\dagger_{n,n^2}.$$

5. Fix $t \in \mathcal{T}_3$. Choose any labeling of the n units by $\{1, \ldots, n\}$, and let $z^{(t)} \in \mathcal{T}_3^n$ be the constant labeled schedule
$$z_i^{(t)} = t, \qquad i = 1, \ldots, n,$$
and let $m^{(t)} \in \mathcal{M}_{n,3}$ be its response-count vector,

$$m_\zeta^{(t)} = \#\{i : z_i^{(t)} = \zeta\}, \qquad \zeta \in \mathcal{T}_3.$$

Then $m_t^{(t)} = n$. Since $n \geq 1$,
$$m_t^{(t)} > 0,$$
so the response type t is represented by a response-count vector with total population size n.

6. Finally, work at the second-order scale $k^{4/3}$. The asserted limit follows from the following direct calculation. For all $k \geq 1$,

$$k^{4/3}\left(\frac{2}{k^2} + \frac{1}{k^4}\right) = 2k^{-2/3} + k^{-8/3}.$$

Both powers on the right tend to zero as $k \to \infty$, because $2/3 > 0$ and $8/3 > 0$. Hence

$$\lim_{k\to\infty} k^{4/3}\left(\frac{2}{k^2} + \frac{1}{k^4}\right) = 0.$$

□

Proof of Theorem 7. 1. Since a zero-sum contrast on $K = 0$ or $K = 1$ is identically zero, the hypotheses imply $K \geq 2$. The rational grid program of Definition 19 is a finite rational linear program. A feasible point is obtained by placing all allocation mass on one allocation orbit and all grid-action mass at 0, with an epigraph coordinate large enough to dominate the finitely many risk rows. The risk rows also force every feasible epigraph coordinate to be nonnegative. Finite-dimensional rational LP strong duality therefore gives rational primal and dual optimizers. Decoding the primal optimizer gives

$$\pi^{n,M,c}, \qquad w^{\mathbb{Q}}, \qquad u \in \mathbb{Q},$$

and decoding the dual optimizer gives multipliers y, whose risk-row component is

$$\nu_{n,M,c}(m) = y_m^{\text{risk}}, \qquad m \in \mathcal{M}_{n,K}.$$

The decoded certificate satisfies

$$w_{r,x,g}^{n,M,c} = w_{r,x,g}^{\mathbb{Q}}, \qquad u = y^{\text{norm},-} - y^{\text{norm},+}, \qquad u = \Lambda_{n,M}(c),$$

with primal feasibility and dual feasibility in the grid program. In particular,

$$y_m^{\text{risk}} \geq 0 \quad \text{for all } m, \qquad \sum_{m \in \mathcal{M}_{n,K}} y_m^{\text{risk}} = 1,$$

so $\nu_{n,M,c}$ is a rational prior on response-count vectors.

2. Define the estimator, for $r \in \mathcal{R}_{n,K}$ and compatible x, by

$$\delta_{n,M,c}(r,x) = \begin{cases} \displaystyle\sum_{g \in \Gamma_{M,c}} g\, w_{r,x,g}^{n,M,c} / \pi_r^{n,M,c}, & \pi_r^{n,M,c} > 0, \\ 0, & \pi_r^{n,M,c} = 0. \end{cases}$$

This is exactly the barycenter condition in the certificate statement; the zero convention applies precisely on allocation orbits with zero certificate mass.

3. For a rational prior ν on $\mathcal{M}_{n,K}$, set

$$D_{r,x,c}(\nu) = \sum_{m \in \mathcal{M}_{n,K}} \nu_m P_m(x \mid r), \qquad N_{r,x,c}(\nu) = \sum_{m \in \mathcal{M}_{n,K}} \nu_m P_m(x \mid r) \tau_c(m),$$

and use the posterior-mean convention

$$\bar{\tau}_{r,x,c}(\nu) = \begin{cases} N_{r,x,c}(\nu)/D_{r,x,c}(\nu), & D_{r,x,c}(\nu) > 0, \\ 0, & D_{r,x,c}(\nu) = 0. \end{cases}$$

Then the Bayes risk at allocation orbit r is

$$\text{BayesRisk}_{n,c}(r,\nu) = \sum_m \nu_m \tau_c(m)^2 - \sum_x \begin{cases} N_{r,x,c}(\nu)^2/D_{r,x,c}(\nu), & D_{r,x,c}(\nu) > 0, \\ 0, & D_{r,x,c}(\nu) = 0. \end{cases}$$

Completing the square gives, for every invariant orbit procedure (π, δ),

$$\sum_m \nu_m R_n((\pi,\delta); m) \geq \sum_r \pi_r \ \text{BayesRisk}_{n,c}(r,\nu) \geq B_{n,c}(\nu).$$

Thus $B_{n,c}(\nu)$ is a Bayes lower bound for the finite orbit game. By the value equality in Theorem 1,

$$B_{n,c}(\nu) \le \rho_n(c).$$

Applying this to $\nu = \nu_{n,M,c}$ gives the stated lower certificate inequality.

4. The barycenter rule $\delta_{n,M,c}$, together with $\pi^{n,M,c}$, is an admissible orbit procedure. Through the invariant correspondence and value equality in Theorem 1,

$$\rho_n(c) \le \sup_{m \in \mathcal{M}_{n,K}} \sum_r \pi_r^{n,M,c} \sum_x P_m(x \mid r)\{\delta_{n,M,c}(r,x) - \tau_c(m)\}^2 = U_{n,M}(c).$$

This proves $\rho_n(c) \le U_{n,M}(c)$.

5. For $\pi_r^{n,M,c} > 0$, let

$$\kappa_{r,x}^{n,M,c}(g) = w_{r,x,g}^{n,M,c} / \pi_r^{n,M,c}, \qquad g \in \Gamma_{M,c}.$$

Primal feasibility gives $\sum_g \kappa_{r,x}^{n,M,c}(g) = 1$, and the barycenter definition gives

$$\delta_{n,M,c}(r,x) = \sum_g g\, \kappa_{r,x}^{n,M,c}(g).$$

Hence, for every m, r, x with $\pi_r^{n,M,c} > 0$,

$$\{\delta_{n,M,c}(r,x) - \tau_c(m)\}^2 \le \sum_g \kappa_{r,x}^{n,M,c}(g)\{g - \tau_c(m)\}^2.$$

When $\pi_r^{n,M,c} = 0$, the corresponding contribution is multiplied by zero. Multiplying by $\pi_r^{n,M,c} P_m(x \mid r)$, summing over (r,x), and using the primal risk row gives

$$U_{n,M}(c) \le u = \Lambda_{n,M}(c).$$

6. It remains to compare $\Lambda_{n,M}(c)$ with the lower certificate. Fix an allocation orbit r. Since each $\tau_c(m)$ lies in $[-h_c, h_c]$, the posterior mean $\bar{\tau}_{r,x,c}(\nu_{n,M,c})$ also lies in this interval when $D_{r,x,c}(\nu_{n,M,c}) > 0$, and the convention value 0 lies there when $D_{r,x,c}(\nu_{n,M,c}) = 0$. The grid $\Gamma_{M,c}$ has mesh h_c/M, so choose $\gamma_{r,x}^\star \in \Gamma_{M,c}$ with

$$\left|\gamma_{r,x}^\star - \bar{\tau}_{r,x,c}(\nu_{n,M,c})\right| \le \frac{h_c}{2M}.$$

Write the relevant dual occupancy difference and nonnegativity multipliers as

$$\alpha_{r,x} = y_{r,x}^{\mathrm{occ},+} - y_{r,x}^{\mathrm{occ},-}, \qquad \zeta_r = y_r^\pi \ge 0, \qquad \omega_{r,x}^\star = y_{r,x,\gamma_{r,x}^\star}^w \ge 0.$$

For the fixed allocation orbit r and the selected grid points $\gamma_{r,x}^\star$, dual feasibility gives the stationarity rows

$$y^{\mathrm{norm},+} - y^{\mathrm{norm},-} - \sum_x \alpha_{r,x} - \zeta_r = 0$$

and, for every compatible observation x,

$$\alpha_{r,x} - \omega_{r,x}^\star + \sum_m \nu_{n,M,c}(m) P_m(x \mid r)\{\gamma_{r,x}^\star - \tau_c(m)\}^2 = 0.$$

Since the dual objective equals $u = y^{\mathrm{norm},-} - y^{\mathrm{norm},+}$, these rows imply

$$
\begin{aligned}
u &= -\sum_x \alpha_{r,x} - \zeta_r \\
&\le \sum_x (-\alpha_{r,x}) \\
&\le \sum_x \sum_m \nu_{n,M,c}(m) P_m(x \mid r)\{\gamma^\star_{r,x} - \tau_c(m)\}^2 \\
&= \sum_m \nu_{n,M,c}(m) \sum_x P_m(x \mid r)\{\gamma^\star_{r,x} - \tau_c(m)\}^2.
\end{aligned}
$$

The posterior square decomposition gives

$$
\sum_m \nu_{n,M,c}(m) \sum_x P_m(x \mid r)\{\gamma^\star_{r,x} - \tau_c(m)\}^2 = \mathrm{BayesRisk}_{n,c}(r, \nu_{n,M,c}) + \sum_x D_{r,x,c}(\nu_{n,M,c})\{\gamma^\star_{r,x} - \bar{\tau}_{r,x,c}(\nu_{n,M,c})\}^2.
$$

Indeed, for each x with positive predictive mass this is completion of the square around the posterior mean $\bar{\tau}_{r,x,c}(\nu_{n,M,c})$; when the predictive mass is zero, all terms carrying $P_m(x \mid r)\nu_{n,M,c}(m)$ vanish and the displayed convention contributes zero. Moreover

$$
\sum_x D_{r,x,c}(\nu_{n,M,c}) = \sum_m \nu_{n,M,c}(m) \sum_x P_m(x \mid r) = 1.
$$

Using the grid approximation bound for every x,

$$
\sum_x D_{r,x,c}(\nu_{n,M,c})\{\gamma^\star_{r,x} - \bar{\tau}_{r,x,c}(\nu_{n,M,c})\}^2 \le \frac{h_c^2}{4M^2} = \frac{C_0(c)}{4M^2}.
$$

Thus, for every allocation orbit r,

$$
u \le \mathrm{BayesRisk}_{n,c}(r, \nu_{n,M,c}) + \frac{C_0(c)}{4M^2}.
$$

The allocation-orbit set is nonempty, so taking the infimum over r and using $u = \Lambda_{n,M}(c)$ gives

$$
\Lambda_{n,M}(c) \le B_{n,c}(\nu_{n,M,c}) + \frac{C_0(c)}{4M^2}.
$$

Together with the preceding steps, this proves the full certificate sandwich.

7. The orbit counts are stars-and-bars counts. Response-count vectors are nonnegative counts over 2^K binary response types with total n, so

$$
|\mathcal{M}_{n,K}| = \binom{n + 2^K - 1}{2^K - 1}.
$$

Allocation-count vectors are nonnegative counts over K arms with total n, so

$$
|\mathcal{R}_{n,K}| = \binom{n + K - 1}{K - 1}.
$$

Finally, a pair (r, x), with x compatible with r, is equivalently a nonnegative count vector on arm-outcome cells $\mathcal{A}_K \times \{0, 1\}$ with total n. There are $2K$ such cells, hence

$$
\sum_{r \in \mathcal{R}_{n,K}} |\{x : x \text{ is compatible with } r\}| = \binom{n + 2K - 1}{2K - 1}.
$$

8. Let (a_k) be any positive sequence and let $M_k \geq 1$. If

$$\frac{a_k}{M_k^2} \longrightarrow 0,$$

then multiplication by the fixed constant $C_0(c)/4$ gives

$$a_k \frac{C_0(c)}{4M_k^2} = \frac{C_0(c)}{4} \frac{a_k}{M_k^2} \longrightarrow 0.$$

For certificate sequences at mesh M_k, the finite-sample sandwich gives, for all positive k,

$$B_k \leq \rho_k(c) \leq U_k, \qquad 0 \leq U_k - B_k \leq \frac{C_0(c)}{4M_k^2}.$$

Finite initial values do not affect limits, and the preceding display therefore implies

$$a_k(U_k - B_k) \longrightarrow 0.$$

Writing $d_k(c) = C_0(c)/k - \rho_k(c)$, if $a_k d_k(c) \to \ell$, then

$$a_k \left\{ \frac{C_0(c)}{k} - U_k \right\} = a_k d_k(c) - a_k\{U_k - \rho_k(c)\} \longrightarrow \ell$$

and

$$a_k \left\{ \frac{C_0(c)}{k} - B_k \right\} = a_k d_k(c) + a_k\{\rho_k(c) - B_k\} \longrightarrow \ell,$$

because both error terms are squeezed by $a_k(U_k - B_k)$. At the second-order scale $a_k = k^{4/3}$ and with $M_k = k$,

$$\frac{k^{4/3}}{k^2} = k^{-2/3} \longrightarrow 0,$$

so

$$k^{4/3}(U_k - B_k) \longrightarrow 0.$$

□

Proof of Theorem 8. 1. First prove the one-sided estimate

$$\sqrt{\rho_n(c)} \leq \sqrt{\rho_n(c')} + \eta(c, c').$$

Put

$$\xi_a = c_a - c'_a, \qquad \eta(c, c') = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |\xi_a|.$$

For a binary response type t, let

$$\mathcal{I}(t) = \{a \in \mathcal{A}_K : t_a = 1\}.$$

Since both contrasts have zero sum, $\sum_a \xi_a = 0$, and therefore

$$\sum_{a \in \mathcal{I}(t)} \xi_a = - \sum_{a \notin \mathcal{I}(t)} \xi_a.$$

Thus

$$2\left|\sum_{a\in\mathcal{I}(t)}\xi_a\right| \le \sum_{a\in\mathcal{I}(t)}|\xi_a| + \sum_{a\notin\mathcal{I}(t)}|\xi_a| = \sum_{a\in\mathcal{A}_K}|\xi_a|,$$

so

$$\left|\sum_{a\in\mathcal{A}_K}\xi_a t_a\right| \le \eta(c,c').$$

For any labeled schedule $z \in \mathcal{T}_K^n$, using $n \ge 1$,

$$\begin{aligned}|\tau_c(z) - \tau_{c'}(z)| &= \left|\frac{1}{n}\sum_{i=1}^n\sum_{a\in\mathcal{A}_K}\xi_a z_{i,a}\right| \\ &\le \frac{1}{n}\sum_{i=1}^n\left|\sum_{a\in\mathcal{A}_K}\xi_a z_{i,a}\right| \le \eta(c,c').\end{aligned}$$

Also set

$$L_c = \sum_{a\in\mathcal{A}_K}|c_a|, \qquad I_c = \left[-\frac{L_c}{2}, \frac{L_c}{2}\right].$$

Applying the preceding target bound to the pair $c, -c$ gives

$$|2\tau_c(z)| = |\tau_c(z) - \tau_{-c}(z)| \le L_c,$$

hence $\tau_c(z) \in I_c$.

Fix any procedure $(\mathcal{D}, \widehat{\tau}')$ for contrast c'. Define the transferred procedure for c by keeping the same assignment law and projecting the estimator output onto I_c:

$$\widehat{\tau}^c(A,Y) = \mathrm{proj}_{I_c}\{\widehat{\tau}'(A,Y)\}.$$

For $y \in I_c$, projection onto the closed interval I_c satisfies

$$\left|\mathrm{proj}_{I_c}(s) - y\right|^2 \le |s-y|^2 \qquad (s\in\mathbb{R}),$$

by the three cases $s < -L_c/2$, $s \in I_c$, and $s > L_c/2$. For fixed z and assignment vector A, define

$$e_z(A) = \widehat{\tau}'(A, Y^{\mathrm{obs}}(z,A)) - \tau_{c'}(z).$$

Using $\tau_c(z) \in I_c$ and the target bound above,

$$\begin{aligned}\left|\widehat{\tau}^c(A,Y^{\mathrm{obs}}(z,A)) - \tau_c(z)\right| &\le \left|\widehat{\tau}'(A,Y^{\mathrm{obs}}(z,A)) - \tau_c(z)\right| \\ &\le |e_z(A)| + |\tau_{c'}(z) - \tau_c(z)| \\ &\le |e_z(A)| + \eta(c,c').\end{aligned}$$

For each target contrast $\bar{c} \in \{c, c'\}$ and any estimator $\widetilde{\tau}$, write its squared-error risk at the fixed schedule z as

$$R_n^{\bar{c}}(\mathcal{D}, \widetilde{\tau}; z) = \mathbb{E}_{\mathcal{D}}\left[\left\{\widetilde{\tau}(A, Y^{\mathrm{obs}}(z,A)) - \tau_{\bar{c}}(z)\right\}^2\right].$$

Taking $\mathcal{D}$-expectations gives

$$R_n^c(\mathcal{D}, \widehat{\tau}^c; z) \le \mathbb{E}_{\mathcal{D}}\left[(|e_z(A)| + \eta(c, c'))^2\right].$$

Since $\eta(c, c') \ge 0$, Cauchy–Schwarz yields

$$\mathbb{E}_{\mathcal{D}}|e_z(A)| \le \sqrt{\mathbb{E}_{\mathcal{D}} e_z(A)^2},$$

and therefore

$$\sqrt{\mathbb{E}_{\mathcal{D}}\left[(|e_z(A)| + \eta(c, c'))^2\right]} \le \sqrt{\mathbb{E}_{\mathcal{D}} e_z(A)^2} + \eta(c, c').$$

Here $\mathbb{E}_{\mathcal{D}} e_z(A)^2 = R_n^{c'}(\mathcal{D}, \widehat{\tau}'; z)$, so

$$\sqrt{R_n^c(\mathcal{D}, \widehat{\tau}^c; z)} \le \sqrt{R_n^{c'}(\mathcal{D}, \widehat{\tau}'; z)} + \eta(c, c').$$

It remains to pass from this pointwise comparison to the two minimax values. For each $\bar{c} \in \{c, c'\}$, set

$$\Psi_{\bar{c}}(\mathcal{D}, \widetilde{\tau}) = \max_{z \in \mathcal{T}_K^n} R_n^{\bar{c}}(\mathcal{D}, \widetilde{\tau}; z),$$

so Definition 3 says

$$\rho_n(\bar{c}) = \inf_{\mathcal{D}, \widetilde{\tau}} \Psi_{\bar{c}}(\mathcal{D}, \widetilde{\tau}).$$

The schedule set is finite and nonempty, and $\sqrt{\cdot}$ is increasing on $[0, \infty)$, hence

$$\max_{z \in \mathcal{T}_K^n} \sqrt{R_n^{\bar{c}}(\mathcal{D}, \widetilde{\tau}; z)} = \sqrt{\Psi_{\bar{c}}(\mathcal{D}, \widetilde{\tau})}.$$

We also have

$$\sqrt{\rho_n(\bar{c})} = \inf_{\mathcal{D}, \widetilde{\tau}} \sqrt{\Psi_{\bar{c}}(\mathcal{D}, \widetilde{\tau})}.$$

Indeed, every $\Psi_{\bar{c}}(\mathcal{D}, \widetilde{\tau})$ is at least $\rho_n(\bar{c})$, giving the $\ge$ inequality. Conversely, the procedure class is nonempty, and for every $\varepsilon > 0$ there is a procedure with

$$\Psi_{\bar{c}}(\mathcal{D}, \widetilde{\tau}) < \rho_n(\bar{c}) + \varepsilon;$$

therefore the infimum of the square roots is at most $\sqrt{\rho_n(\bar{c}) + \varepsilon}$, and continuity at $\rho_n(\bar{c})$ gives the reverse inequality as $\varepsilon \downarrow 0$.

Taking the maximum over z in the previous pointwise display gives

$$\sqrt{\Psi_c(\mathcal{D}, \widehat{\tau}^c)} \le \sqrt{\Psi_{c'}(\mathcal{D}, \widehat{\tau}')} + \eta(c, c').$$

Since $(\mathcal{D}, \widehat{\tau}^c)$ is one admissible procedure for contrast c, taking the infimum over all procedures $(\mathcal{D}, \widehat{\tau}')$ for c' gives

$$\sqrt{\rho_n(c)} \le \sqrt{\rho_n(c')} + \eta(c, c').$$

2. The contrast distance is symmetric:

$$\eta(c', c) = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c'_a - c_a| = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c_a - c'_a| = \eta(c, c').$$

Applying the one-sided estimate to the ordered pair (c', c) gives

$$\sqrt{\rho_n(c')} \le \sqrt{\rho_n(c)} + \eta(c', c) = \sqrt{\rho_n(c)} + \eta(c, c'),$$

hence

$$-\eta(c, c') \le \sqrt{\rho_n(c)} - \sqrt{\rho_n(c')}.$$

3. Applying the same one-sided estimate to the ordered pair (c, c') gives

$$\sqrt{\rho_n(c)} - \sqrt{\rho_n(c')} \le \eta(c, c').$$

Together with the previous displayed lower bound, this is exactly

$$\left|\sqrt{\rho_n(c)} - \sqrt{\rho_n(c')}\right| \le \eta(c, c').$$

□

Proof of Theorem 9. 1. Fix a rational zero-sum contrast $q \in \mathbb{Q}^K \setminus \{0\}$ and integers $n, M \ge 1$. Apply Theorem 7 to q. It gives an invariant allocation certificate π, a rational grid weight w, a rational epigraph value u, a rational response-count multiplier ν, and a barycenter rule δ, with

$$B := B_{n,q}(\nu), \qquad U := U_{n,M}(q; \pi, \delta), \qquad u = \Lambda_{n,M}(q),$$

and

$$B \le \rho_n(q) \le U \le u \le B + \frac{C_0(q)}{4M^2}.$$

The multiplier ν is a response-count probability prior. Indeed, by Theorem 7, ν is the risk component of a dual-feasible multiplier for the grid-action program of Definition 19. In that program the response-count constraints are the inequality rows

$$\sum_{r,x,g} P_m(x \mid r)\, w_{r,x,g}\{g - \tau_q(m)\}^2 - u \le 0, \qquad m \in \mathcal{M}_{n,K},$$

so dual feasibility gives $\nu(m) \ge 0$ for every m. The epigraph variable u has objective coefficient 1, coefficient -1 in each of those rows, and no other occurrence, so its dual stationarity equation is

$$1 - \sum_{m \in \mathcal{M}_{n,K}} \nu(m) = 0.$$

Thus

$$\nu(m) \ge 0 \quad \text{for every } m \in \mathcal{M}_{n,K}, \qquad \sum_{m \in \mathcal{M}_{n,K}} \nu(m) = 1.$$

With ν now a probability prior, the lower certificate is an infimum of Bayes risks over the finite allocation-count set, hence $B \ge 0$; the displayed sandwich gives $U \ge 0$. Put

$$\eta_q := \eta(c, q), \qquad R^- := \left\{\max(0, \sqrt{B} - \eta_q)\right\}^2, \qquad R^+ := (\sqrt{U} + \eta_q)^2.$$

The pointwise transfer estimate used for both certificate directions is the following. If a procedure for one contrast is clipped to the natural interval of another contrast, then for every schedule z,

$$\sqrt{R_n(\mathcal{D}, \widehat{\tau}^{\,c}; z; c)} \le \sqrt{R_n(\mathcal{D}, \widehat{\tau}^{\,q}; z; q)} + \eta_q.$$

Indeed, for each assignment A,

$$|\tau_c(z) - \tau_q(z)| \le \eta_q$$

because the target difference is an average of binary response-type scores and

$$\eta_q = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c(a) - q(a)|.$$

Projection onto $[-L_c/2, L_c/2]$ can only reduce squared distance to $\tau_c(z)$, and Cauchy–Schwarz gives

$$\mathbb{E}_{\mathcal{D}}(|E_A| + \eta_q)^2 \le \left\{ \sqrt{\mathbb{E}_{\mathcal{D}} E_A^2} + \eta_q \right\}^2$$

for the original residual $E_A = \widehat{\tau}^{\,q}(A, Y^{\mathrm{obs}}) - \tau_q(z)$.

The upper procedure keeps the invariant allocation law induced by π and uses

$$\widehat{\tau}^{\,c}(A, Y^{\mathrm{obs}}) = \mathrm{proj}_{[-L_c/2, L_c/2]}\{\delta(r(A), x(A, Y^{\mathrm{obs}}))\}.$$

The rational certificate bounds the q-risk of the barycenter rule by U, so the transfer estimate yields worst-case c-risk at most R^+, and therefore

$$\rho_n(c) \le R^+.$$

For the lower endpoint, spread this response-count prior uniformly over each labeled schedule orbit, and write the resulting expectation as

$$\mathbb{E}_{\nu}^{\mathrm{sch}}[\varphi] := \sum_{z \in \mathcal{T}_K^n} \frac{\nu(m(z))}{|\{z' : m(z') = m(z)\}|} \, \varphi(z).$$

The denominator is positive because each response-count vector is represented by at least one labeled schedule, and the displayed nonnegativity and mass identity make $\mathbb{E}_{\nu}^{\mathrm{sch}}$ a probability expectation. We first record the all-procedure lower bound supplied by this lifted prior. Let $\mathfrak{p} = (\mathcal{D}_{\mathfrak{p}}, \widehat{\tau}_{\mathfrak{p}})$ be any labeled procedure for contrast q, and average $\mathfrak{p}$ over the unit-permutation group $\mathfrak{S}_n$. By the lossless symmetrization and exact invariant correspondence in Theorem 1, the averaged procedure is invariant and corresponds to an orbit procedure $\mathfrak{o}$. If $\alpha_{\mathfrak{o}}$ is its allocation-count law, define its orbit risk at response-count vector m by

$$\mathrm{Risk}_{n,q}^{\mathrm{orb}}(\mathfrak{o}; m) := \sum_{r \in \mathcal{R}_{n,K}} \alpha_{\mathfrak{o}}(r) \sum_x P_m(x \mid r) \{\widehat{\tau}_{\mathfrak{o}}(r, x) - \tau_q(m)\}^2.$$

For a fixed allocation count r, the conditional Bayes risk under ν is bounded below by $B = B_{n,q}(\nu)$; on zero prior-predictive fibers the corresponding summand is zero, and on positive prior-predictive fibers the posterior-mean residual is the conditional minimum. Hence, for every orbit procedure,

$$B \le \sum_m \nu(m) \, \mathrm{Risk}_{n,q}^{\mathrm{orb}}(\mathfrak{o}; m).$$

Because the lifted prior is uniform within each response-count orbit,

$$\sum_m \nu(m) \, \mathrm{Risk}_{n,q}^{\mathrm{orb}}(\mathfrak{o}; m) = \mathbb{E}_{\nu}^{\mathrm{sch}} \left[\mathrm{Risk}_{n,q}^{\mathrm{orb}}(\mathfrak{o}; m(z)) \right].$$

Lossless symmetrization gives, for every labeled schedule z,

$$\mathrm{Risk}^{\mathrm{orb}}_{n,q}(\mathfrak{o}; m(z)) \le \frac{1}{|\mathfrak{S}_n|} \sum_{\sigma \in \mathfrak{S}_n} R_n(\mathcal{D}_{\mathfrak{p}}, \widehat{\tau}_{\mathfrak{p}}; \sigma z; q).$$

The lifted prior is invariant under unit permutations, so averaging the last display over z yields

$$B \le \mathbb{E}^{\mathrm{sch}}_{\nu} \left[R_n(\mathcal{D}_{\mathfrak{p}}, \widehat{\tau}_{\mathfrak{p}}; z; q) \right]$$

for every labeled q-procedure $\mathfrak{p}$.

Now let $\mathfrak{r} = (\mathcal{D}_{\mathfrak{r}}, \widehat{\tau}_{\mathfrak{r}})$ be any procedure for the real contrast c, and let $\mathfrak{r}^q$ be the procedure obtained by clipping its action to $[-L_q/2, L_q/2]$. Applying the pointwise transfer estimate in the direction from c to q,

$$\sqrt{R_n(\mathfrak{r}^q; z; q)} \le \sqrt{R_n(\mathcal{D}_{\mathfrak{r}}, \widehat{\tau}_{\mathfrak{r}}; z; c)} + \eta_q.$$

Set

$$\mathscr{V}_c(\mathfrak{r}) := \mathbb{E}^{\mathrm{sch}}_{\nu} \left[R_n(\mathcal{D}_{\mathfrak{r}}, \widehat{\tau}_{\mathfrak{r}}; z; c) \right].$$

Squaring the preceding inequality, integrating with respect to $\mathbb{E}^{\mathrm{sch}}_{\nu}$, and using Cauchy–Schwarz in the finite schedule space gives

$$\mathbb{E}^{\mathrm{sch}}_{\nu} \left[R_n(\mathfrak{r}^q; z; q) \right] \le \mathscr{V}_c(\mathfrak{r}) + 2\eta_q \sqrt{\mathscr{V}_c(\mathfrak{r})} + \eta_q^2 = \{ \sqrt{\mathscr{V}_c(\mathfrak{r})} + \eta_q \}^2.$$

Combining this with the all-procedure q-lower bound applied to $\mathfrak{r}^q$ gives

$$B \le \{ \sqrt{\mathscr{V}_c(\mathfrak{r})} + \eta_q \}^2.$$

If $B \le 0$, then $\sqrt{B} = 0$ and $R^- = 0 \le \mathscr{V}_c(\mathfrak{r})$. If $B > 0$, the displayed inequality gives $\sqrt{B} \le \sqrt{\mathscr{V}_c(\mathfrak{r})} + \eta_q$. When $\sqrt{B} \le \eta_q$ this again gives $R^- = 0 \le \mathscr{V}_c(\mathfrak{r})$, and otherwise

$$R^- = (\sqrt{B} - \eta_q)^2 \le \mathscr{V}_c(\mathfrak{r}).$$

Thus the lifted schedule prior has Bayes risk at least R^- against every real-contrast procedure. Since prior risk is bounded above by worst-case risk for each procedure, taking the minimax value gives

$$R^- \le \rho_n(c).$$

It remains to verify the stated width. From the rational sandwich,

$$U - B \le \frac{C_0(q)}{4M^2}.$$

Also Theorem 4 gives

$$\rho_n(q) \le \frac{C_0(q)}{n},$$

so

$$U \le \frac{C_0(q)}{n} + \frac{C_0(q)}{4M^2}.$$

Consequently,

$$\sqrt{U} \le \sqrt{C_0(q)}\, n^{-1/2} + \frac{\sqrt{C_0(q)}}{2M},$$

because the square of the right-hand side contains the displayed upper bound plus a nonnegative cross term. Finally,

$$B - \{\max(0, \sqrt{B} - \eta_q)\}^2 \le 2\eta_q \sqrt{B} + \eta_q^2,$$

and $\sqrt{B} \le \sqrt{U}$, whence

$$R^+ - R^- \le \frac{C_0(q)}{4M^2} + 4\eta_q \sqrt{U} + 2\eta_q^2.$$

Substituting the bound on $\sqrt{U}$ and using

$$\frac{2\sqrt{C_0(q)}\,\eta_q}{M} \le \frac{C_0(q)}{M^2} + \eta_q^2$$

gives

$$R^+ - R^- \le \frac{5C_0(q)}{4M^2} + 4\sqrt{C_0(q)}\,\eta_q\, n^{-1/2} + 3\eta_q^2.$$

Thus B, U, R^-, R^+ form a transferred certificate for (n, M, c, q), proving the finite-grid transfer assertion.

2. We next construct the rational approximating sequence. Since c is nonzero, $L_c > 0$. For $s \in \mathbb{N}$, define the approximation radius

$$\operatorname{rad}_s = \begin{cases} L_c/4, & s = 0, \\ \min\{L_c/4, s^{-3}\}, & s \ge 1. \end{cases}$$

Then $0 < \operatorname{rad}_s < L_c/2$. Choose a fixed arm $a_\circ$. By density of $\mathbb{Q}$ in $\mathbb{R}$, choose rational numbers $\widetilde{c}_s(a)$ for $a \ne a_\circ$ with

$$|c(a) - \widetilde{c}_s(a)| < \frac{\operatorname{rad}_s}{2K},$$

and set

$$\widetilde{c}_s(a_\circ) = -\sum_{a \ne a_\circ} \widetilde{c}_s(a).$$

Because c is zero-sum,

$$c(a_\circ) = -\sum_{a \ne a_\circ} c(a),$$

and hence

$$|c(a_\circ) - \widetilde{c}_s(a_\circ)| \le \sum_{a \ne a_\circ} |c(a) - \widetilde{c}_s(a)|.$$

Therefore

$$\sum_{a \in \mathcal{A}_K} |c(a) - \widetilde{c}_s(a)| \le \operatorname{rad}_s .$$

Since $\operatorname{rad}_s < L_c/2$, this rational zero-sum vector is nonzero; otherwise $L_c \le \operatorname{rad}_s$. Denote the resulting rational contrast by $c^{(s)}$. Then

$$\eta(c, c^{(s)}) \le \operatorname{rad}_s,$$

and for $s \geq 1$,

$$\eta(c, c^{(s)}) \leq s^{-3}.$$

It follows that

$$\eta(c, c^{(s)}) \to 0, \qquad s^{5/6}\eta(c, c^{(s)}) \leq s^{5/6}s^{-3} = s^{-13/6} \to 0.$$

For each arm a,

$$|c^{(s)}(a) - c(a)| \leq \sum_{b \in \mathcal{A}_K} |c^{(s)}(b) - c(b)| = 2\eta(c, c^{(s)}),$$

so $c^{(s)}(a) \to c(a)$.

3. For each $s \geq 1$, apply the fixed construction from Item 1 with $q = c^{(s)}$, $n = s$, and $M = s$, obtaining B_s, U_s, R_s^-, R_s^+. The construction gives a transferred certificate for every $s \geq 1$. Write

$$\eta_s := \eta(c, c^{(s)}), \qquad C_s := C_0(c^{(s)}).$$

Because $c^{(s)}(a) \to c(a)$ for every arm and the arm set is finite,

$$L_{c^{(s)}} = \sum_{a \in \mathcal{A}_K} |c^{(s)}(a)| \longrightarrow \sum_{a \in \mathcal{A}_K} |c(a)| = L_c,$$

and hence

$$C_s = \frac{L_{c^{(s)}}^2}{4} \longrightarrow C_0(c).$$

The certificate width bound gives, for $s \geq 1$,

$$0 \leq s^{4/3}(R_s^+ - R_s^-) \leq s^{4/3}\left\{\frac{5C_s}{4s^2} + 4\sqrt{C_s}\,\eta_s\, s^{-1/2} + 3\eta_s^2\right\}.$$

The three terms on the right converge to zero:

$$s^{4/3}\frac{5C_s}{4s^2} = \frac{5C_s}{4}s^{-2/3} \to 0,$$

$$s^{4/3}\,4\sqrt{C_s}\,\eta_s\, s^{-1/2} = 4\sqrt{C_s}\,\{s^{5/6}\eta_s\} \to 0,$$

and, using $\eta_s \leq s^{-3}$,

$$s^{4/3}\,3\eta_s^2 \leq 3s^{4/3}s^{-6} = 3s^{-14/3} \to 0.$$

By squeezing,

$$s^{4/3}(R_s^+ - R_s^-) \to 0.$$

This proves the real-contrast approximation assertion.

4. Finally suppose a rational zero-sum contrast q embeds coefficientwise as c. For any $n \geq 1$, apply the fixed construction with $M = n$. The rational sandwich gives

$$U - B \leq \frac{C_0(q)}{4n^2}.$$

Since the real embedding of q is c, $C_0(q) = C_0(c)$, and hence

$$U - B \leq \frac{C_0(c)}{4n^2}.$$

Together with the transferred certificate constructed in Item 1, this proves the exact rational case.

□

Proof of Theorem 10. Fix $n \geq 1$ and $M \geq 1$.

1. Let
$$c^\dagger_{\mathbb{Q}} = (1, -1/2, -1/2) \in \mathbb{Q}^3, \qquad c^\dagger = (1, -1/2, -1/2) \in \mathbb{R}^3.$$
This contrast is nonzero and zero-sum, and
$$L_{c^\dagger} = |1| + |-1/2| + |-1/2| = 2, \qquad C_0(c^\dagger) = \left(\frac{L_{c^\dagger}}{2}\right)^2 = 1.$$
Hence the real embedding of $c^\dagger_{\mathbb{Q}}$ has $C_0 = 1$, and the grid in Definition 19 becomes
$$\Gamma_{M,c^\dagger} = \{-1 + j/M : 0 \leq j \leq 2M\} = \Gamma_M.$$
Therefore, after specializing Theorem 7 to $K = 3$ and $c = c^\dagger_{\mathbb{Q}}$, the rational contrast grid program of Definition 19 is the three-arm rational grid program of Definition 18, with
$$\Lambda_{n,M}(c^\dagger) = \Lambda^\dagger_{n,M}, \qquad \rho_n(c^\dagger) = \rho^\dagger_n.$$

2. The specialized certificate theorem supplies a rational allocation certificate $\pi^{n,M}$, a joint design–action weight $w^{n,M}$, a rational epigraph value u, a rational response-count multiplier $\nu_{n,M}$, and an estimator $\delta_{n,M}$. Its exact-certificate conclusion gives the rational grid weight whose real embedding is $w^{n,M}$, primal feasibility, dual feasibility, the identification of the risk component with $\nu_{n,M}$, and dual objective value u. Its barycenter conclusion gives, for every $r \in \mathcal{R}_{n,3}$ and compatible x,
$$\delta_{n,M}(r,x) = \begin{cases} \displaystyle\sum_{g \in \Gamma_M} g \, \frac{w^{n,M}_{r,x,g}}{\pi^{n,M}_r}, & \pi^{n,M}_r > 0, \\ 0, & \pi^{n,M}_r = 0. \end{cases}$$
The second branch fixes the estimator at allocation counts that carry no certificate mass; those counts are never charged by the certificate, so the choice is immaterial.

   The same specialization gives
$$u = \Lambda^\dagger_{n,M}.$$
For a response-count multiplier ν, the specialized lower endpoint is
$$B_{n,c^\dagger}(\nu) = \inf_{r \in \mathcal{R}_{n,3}} \operatorname{BayesRisk}_{n,c^\dagger}(r, \nu) = B_n(\nu),$$
and the specialized upper endpoint is precisely
$$U_{n,M} = \sup_{m \in \mathcal{M}_{n,3}} \sum_{r \in \mathcal{R}_{n,3}} \pi^{n,M}_r \sum_x P_m(x \mid r)\{\delta_{n,M}(r,x) - \tau_{c^\dagger}(m)\}^2.$$

Thus the certificate sandwich in Theorem 7, together with $C_0(c^\dagger) = 1$, yields

$$B_n(\nu_{n,M}) \leq \rho_n^\dagger \leq U_{n,M} \leq \Lambda_{n,M}^\dagger \leq B_n(\nu_{n,M}) + \frac{1}{4M^2}.$$

3. Let $a_k > 0$ be any positive normalizing sequence and let $M_k \geq 1$ for all k. If

$$\frac{a_k}{M_k^2} \longrightarrow 0,$$

then

$$a_k \frac{1}{4M_k^2} = \frac{1}{4} \frac{a_k}{M_k^2} \longrightarrow 0.$$

This is the mesh-rate conclusion at $C_0(c^\dagger) = 1$.

4. Now let $a_k > 0$, M_k, L_k, U_k, and C satisfy the hypotheses of the final assertion. All inequalities below are for $k \geq 1$, which is an eventual range for limits along $k \to \infty$. From

$$L_k \leq \rho_k^\dagger \leq U_k, \qquad U_k - L_k \leq \frac{1}{4M_k^2},$$

we have

$$0 \leq a_k(\rho_k^\dagger - L_k) \leq a_k(U_k - L_k) \leq a_k \frac{1}{4M_k^2},$$

and also

$$0 \leq a_k(U_k - \rho_k^\dagger) \leq a_k(U_k - L_k) \leq a_k \frac{1}{4M_k^2}.$$

The mesh assumption and the preceding step give

$$a_k(\rho_k^\dagger - L_k) \longrightarrow 0, \qquad a_k(U_k - \rho_k^\dagger) \longrightarrow 0.$$

Finally,

$$d_k(c^\dagger) = \frac{C_0(c^\dagger)}{k} - \rho_k^\dagger.$$

Therefore

$$a_k \left\{ \frac{C_0(c^\dagger)}{k} - L_k \right\} = a_k d_k(c^\dagger) + a_k(\rho_k^\dagger - L_k),$$

and

$$a_k \left\{ \frac{C_0(c^\dagger)}{k} - U_k \right\} = a_k d_k(c^\dagger) - a_k(U_k - \rho_k^\dagger).$$

Since $a_k d_k(c^\dagger) \to C$ and both residual terms converge to 0, the two displayed sequences converge to C.

□

Proof of Proposition 2. 1. For $c^\dagger = (1, -1/2, -1/2)$, the contrast-weighted one-unit law is

$$\Pr(A_i = 1) = \frac{|1|}{|1| + |-1/2| + |-1/2|} = \frac{1}{2}, \qquad \Pr(A_i = 2) = \Pr(A_i = 3) = \frac{1}{4}.$$

Since $c^\dagger$ is positive on arm 1 and negative on arms $2, 3$,

$$Z_i = \begin{cases} 2t_{i,1} - 1, & A_i = 1, \\ 1 - 2t_{i,2}, & A_i = 2, \\ 1 - 2t_{i,3}, & A_i = 3. \end{cases}$$

Thus, conditionally on z,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}^\star}(Z_i \mid z) &= \frac{1}{2}(2t_{i,1} - 1) + \frac{1}{4}(1 - 2t_{i,2}) + \frac{1}{4}(1 - 2t_{i,3}) \\ &= t_{i,1} - \frac{t_{i,2} + t_{i,3}}{2} = \mu_i. \end{aligned}$$

This proves the asserted conditional mean for every $n \geq 1$, schedule z, and unit i.

2. The score is always sign-valued, so $Z_i^2 = 1$. Combining this with the preceding mean identity gives

$$\mathrm{Var}_{\mathcal{D}^\star}(Z_i \mid z) = \mathbb{E}_{\mathcal{D}^\star}(Z_i^2 \mid z) - \{\mathbb{E}_{\mathcal{D}^\star}(Z_i \mid z)\}^2 = 1 - \mu_i^2.$$

3. Write

$$\overline{Z} = \frac{1}{n} \sum_i Z_i.$$

The assignment law is a product law, so the Z_i's are conditionally independent given z. Moreover

$$\tau_{c^\dagger}(z) = \frac{1}{n} \sum_i \left(t_{i,1} - \frac{t_{i,2} + t_{i,3}}{2} \right) = \frac{1}{n} \sum_i \mu_i,$$

and therefore $\overline{Z}$ is conditionally unbiased for $\tau_{c^\dagger}(z)$. Hence, for $n \geq 1$,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}^\star}\left[\{\overline{Z} - \tau_{c^\dagger}(z)\}^2 \,\middle|\, z \right] &= \mathrm{Var}_{\mathcal{D}^\star}(\overline{Z} \mid z) \\ &= \frac{1}{n^2} \sum_i \mathrm{Var}_{\mathcal{D}^\star}(Z_i \mid z) \\ &= \frac{1}{n^2} \sum_i (1 - \mu_i^2) = \frac{1}{n} - \frac{1}{n^2} \sum_i \mu_i^2. \end{aligned}$$

4. For the scalar experiment at $n = 3$, let

$$\mathfrak{L}_Z = \{-1, -1/2, 0, 1/2, 1\}$$

be the feasible set of scalar means. Given $\boldsymbol{\mu} = (\mu_1, \mu_2, \mu_3) \in \mathfrak{L}_Z^3$, the retained observation has independent coordinates with

$$\Pr_\mu(Z_i = 1) = \frac{1 + \mu_i}{2}, \qquad \Pr_\mu(Z_i = -1) = \frac{1 - \mu_i}{2},$$

and the scalar target is $M_\mu / 3$, where

$$M_\mu = \sum_{i=1}^{3} \mu_i, \qquad Q_\mu = \sum_{i=1}^{3} \mu_i^2.$$

Set

$$\alpha = \frac{3-\sqrt{3}}{6}, \qquad \widehat{\theta}_Z(Z_1, Z_2, Z_3) = \alpha \sum_{i=1}^{3} Z_i.$$

Then

$$\mathbb{E}_\mu\left[\left\{\widehat{\theta}_Z - \frac{M_\mu}{3}\right\}^2\right] = \alpha^2(3 - Q_\mu) + \left(\alpha - \frac{1}{3}\right)^2 M_\mu^2$$
$$= \frac{2-\sqrt{3}}{18}\{9 - 3Q_\mu + M_\mu^2\}.$$

By Cauchy's inequality, $M_\mu^2 \le 3Q_\mu$, and hence

$$\mathbb{E}_\mu\left[\left\{\widehat{\theta}_Z - \frac{M_\mu}{3}\right\}^2\right] \le \frac{2-\sqrt{3}}{2} = 1 - \frac{\sqrt{3}}{2}.$$

Taking the supremum over $\boldsymbol{\mu} \in \mathfrak{L}_Z^3$ and then the infimum over scalar rules gives

$$V_3^Z \le 1 - \frac{\sqrt{3}}{2}.$$

5. For the reverse inequality, put the following weights on the five homogeneous states

$$\lambda_j = -1 + \frac{j}{2}, \qquad j = 0, 1, 2, 3, 4:$$

$$\varpi_0 = \varpi_4 = \frac{-11 + 7\sqrt{3}}{12}, \qquad \varpi_1 = \varpi_3 = \frac{8 - 4\sqrt{3}}{3}, \qquad \varpi_2 = \frac{-5 + 3\sqrt{3}}{2}.$$

They are nonnegative and satisfy $\sum_{j=0}^{4} \varpi_j = 1$. For any scalar rule $\widehat{\theta}$, its Bayes risk under this prior is

$$\mathcal{B}(\widehat{\theta}) = \sum_{j=0}^{4} \varpi_j \, \mathbb{E}_{\lambda_j}\left[(\widehat{\theta}(Z) - \lambda_j)^2\right],$$

where $\mathbb{E}_{\lambda_j}$ denotes the product law with all three scalar means equal to λ_j. If $N_+ = \#\{i : Z_i = 1\}$, then for $b = 0, 1, 2, 3$ the prior predictive mass of the event $\{N_+ = b\}$ and its target numerator are

$$\mathfrak{P}_b = \binom{3}{b} \sum_{j=0}^{4} \varpi_j \left(\frac{1+\lambda_j}{2}\right)^b \left(\frac{1-\lambda_j}{2}\right)^{3-b}, \qquad \mathfrak{H}_b = \binom{3}{b} \sum_{j=0}^{4} \varpi_j \lambda_j \left(\frac{1+\lambda_j}{2}\right)^b \left(\frac{1-\lambda_j}{2}\right)^{3-b}.$$

A direct expansion gives

b	$\mathfrak{P}_b$	$\mathfrak{H}_b$	$\mathfrak{H}_b/\mathfrak{P}_b$
0	$(-1+3\sqrt{3})/16$	$(6-5\sqrt{3})/16$	-3α
1	$(9-3\sqrt{3})/16$	$(-6+3\sqrt{3})/16$	$-\alpha$
2	$(9-3\sqrt{3})/16$	$(6-3\sqrt{3})/16$	α
3	$(-1+3\sqrt{3})/16$	$(-6+5\sqrt{3})/16$	$3\alpha.$

Thus the posterior mean is exactly $\alpha(2N_+ - 3) = \widehat{\theta}_Z(Z)$, and completing the square yields

$$\mathcal{B}(\widehat{\theta}) = \mathcal{B}(\widehat{\theta}_Z) + \sum_{Z\in\{-1,1\}^3} \Pr_{\text{prior}}(Z)\{\widehat{\theta}(Z) - \widehat{\theta}_Z(Z)\}^2 \geq \mathcal{B}(\widehat{\theta}_Z).$$

For a homogeneous level λ_j,

$$\mathbb{E}_{\lambda_j}\Big[(\widehat{\theta}_Z(Z) - \lambda_j)^2\Big] = \alpha^2\, 3(1-\lambda_j^2) + (3\alpha\lambda_j - \lambda_j)^2,$$

and summing this identity with the weights ϖ_j gives

$$\mathcal{B}(\widehat{\theta}_Z) = 1 - \frac{\sqrt{3}}{2}.$$

Every supremum risk dominates this Bayes risk, so

$$1 - \frac{\sqrt{3}}{2} \leq V_3^Z.$$

Together with the preceding item, this proves $V_3^Z = 1 - \sqrt{3}/2$.

6. The full-data rule is defined from the observed arm labels and outcomes by

$$r_1 = \sum_i \mathbf{1}\{A_i = 1\}, \qquad s_1 = \sum_{i:A_i=1} (2Y_i^{\text{obs}} - 1), \qquad s_- = \sum_{i:A_i\in\{2,3\}} (1 - 2Y_i^{\text{obs}}).$$

The rule is given by a lookup table. Let $\mathrm{T}_F(r_1, s_1, s_-)$ be the integer in the table

$(0,0,-3)$	-623	$(0,0,-1)$	-183
$(0,0,1)$	183	$(0,0,3)$	623
$(1,-1,-2)$	-638	$(1,-1,0)$	-199
$(1,-1,2)$	148	$(1,1,-2)$	-148
$(1,1,0)$	199	$(1,1,2)$	638
$(2,-2,-1)$	-648	$(2,-2,1)$	-210
$(2,0,-1)$	-167	$(2,0,1)$	167
$(2,2,-1)$	210	$(2,2,1)$	648
$(3,-3,0)$	-660	$(3,-1,0)$	-167
$(3,1,0)$	167	$(3,3,0)$	$660.$

The displayed triples are exactly the triples (r_1, s_1, s_-) that assignments and binary outcomes can generate, so the rule is defined wherever it is used; set $\mathrm{T}_F = 0$ elsewhere. Define

$$\widehat{\tau}_F(A, Y^{\text{obs}}) = \operatorname{proj}_{[-1,1]}\left(\frac{\mathrm{T}_F(r_1, s_1, s_-)}{1000}\right).$$

All displayed values lie in $[-1000, 1000]$, so the projection leaves the table value divided by 1000 unchanged. For any schedule z, its risk under $\mathcal{D}^\star$ is the finite rational sum

$$R_F(z) = \sum_{a\in\mathcal{A}_3^3} \prod_{i=1}^3 q_{a_i}^\star \left\{\frac{\mathrm{T}_F(r_1(a), s_1(a,z), s_-(a,z))}{1000} - \tau_{c^\dagger}(z)\right\}^2.$$

Identifying an assignment $a \in \mathcal{A}_3^3$ with its ordered triple gives a 27-term rational sum. To certify the uniform bound, write each schedule as

$$z = (\mathbf{t}^{(1)}, \mathbf{t}^{(2)}, \mathbf{t}^{(3)}), \qquad \mathbf{t}^{(j)} \in \{0,1\}^3,$$

and order the eight possible response types as

$$\mathfrak{O}_3 = (000, 001, 010, 011, 100, 101, 110, 111).$$

For $\mathbf{t}^{(1)}, \mathbf{t}^{(2)}, \mathbf{t}^{(3)} \in \mathfrak{O}_3$, define the cleared slack

$$\Xi(\mathbf{t}^{(1)}, \mathbf{t}^{(2)}, \mathbf{t}^{(3)}) = 144000000 \left\{ \frac{511653}{4000000} - R_F(\mathbf{t}^{(1)}, \mathbf{t}^{(2)}, \mathbf{t}^{(3)}) \right\}.$$

Substituting the 27-term expression for R_F and clearing denominators makes Ξ an integer-valued function of the triple, and the check is a finite one: there are $8^3 = 512$ triples, and $\Xi \geq 0$ at every one of them. The evaluation is an exact rational computation: after substituting the displayed formula for R_F, the 512 integer values of Ξ are all nonnegative. The bound is attained: for instance

$$\Xi(000, 111, 111) = 0,$$

so no smaller constant than 511653/4000000 works for this estimator. Since $\Xi \geq 0$ on all of $(\{0,1\}^3)^3$,

$$R_F(z) \leq \frac{511653}{4000000} \qquad \text{for every } z \in \mathcal{T}_3^3.$$

For the schedule with response types $(0,0,0), (1,1,1), (1,1,1)$, the target is 0, and the same finite sum reduces to

$$\begin{aligned} R_F(z) = {} & \frac{1}{8}\left(\frac{183}{1000}\right)^2 + \frac{1}{8}\left(\frac{638}{1000}\right)^2 + \frac{1}{4}\left(\frac{199}{1000}\right)^2 + \frac{1}{4}\left(\frac{167}{1000}\right)^2 \\ & + \frac{1}{8}\left(\frac{648}{1000}\right)^2 + \frac{1}{8}\left(\frac{167}{1000}\right)^2 = \frac{511653}{4000000}. \end{aligned}$$

Therefore the worst-case risk of this full-data rule is exactly 511653/4000000.

7. The rational comparison in the full-data certificate is

$$\frac{511653}{4000000} < \frac{18213}{136000}.$$

Also,

$$\sqrt{3} < \frac{117787}{68000},$$

because squaring both positive sides gives $3 < (117787/68000)^2$. Hence

$$1 - \frac{\sqrt{3}}{2} > 1 - \frac{1}{2}\frac{117787}{68000} = \frac{18213}{136000}.$$

Thus

$$\frac{511653}{4000000} < \frac{18213}{136000} < 1 - \frac{\sqrt{3}}{2}.$$

8. It remains to compute the scalar moment envelope. Let

$$\mathfrak{L}_Z = \{-1, -1/2, 0, 1/2, 1\}, \qquad M = \left|\sum_i \mu_i\right|, \qquad Q_\mu = \sum_i \mu_i^2,$$

with each $\mu_i \in \mathfrak{L}_Z$. Pointwise on $\mathfrak{L}_Z$,

$$\frac{|\mu_i|}{2} \le \mu_i^2, \qquad \frac{3|\mu_i| - 1}{2} \le \mu_i^2.$$

Since $\sum_i |\mu_i| \ge M$, summing these two inequalities gives

$$Q_\mu \ge \frac{M}{2}, \qquad Q_\mu \ge \frac{3M - n}{2}.$$

Consequently

$$Q_\mu \ge \begin{cases} M/2, & M \le n/2, \\ (3M - n)/2, & M > n/2. \end{cases}$$

For attainment, feasibility implies $M = k/2$ for some integer $0 \le k \le 2n$. If $k \le n$, take k coordinates equal to $1/2$ and the remaining coordinates equal to 0. Then

$$\left|\sum_i \mu_i\right| = \frac{k}{2} = M, \qquad \sum_i \mu_i^2 = \frac{k}{4} = \frac{M}{2}.$$

If $k > n$, write $\ell = 2n - k$, so $0 \le \ell < n$, take ℓ coordinates equal to $1/2$, and take the remaining $n - \ell$ coordinates equal to 1. Then

$$\left|\sum_i \mu_i\right| = n - \frac{\ell}{2} = \frac{k}{2} = M,$$

and

$$\sum_i \mu_i^2 = n - \frac{3\ell}{4} = \frac{3M - n}{2}.$$

For $n = 0$, feasibility forces $M = 0$, and the empty vector gives value 0, matching the first branch. Taking the infimum over feasible scalar mean vectors yields

$$\min\left\{Q_\mu : \ \mu \in \mathfrak{L}_Z^n, \ \left|\sum_i \mu_i\right| = M\right\} = \begin{cases} M/2, & M \le n/2, \\ (3M - n)/2, & M > n/2. \end{cases}$$

Equivalently, in the notation of the statement, this is the asserted formula for the minimum attainable $\sum_i \mu_i^2$.

9. Combining the score identities, the average-score risk formula, the exact scalar minimax value, the full-data witness with its rational separation, and the scalar moment envelope proves all asserted clauses.

□

Proof of Theorem 11. Let

$$s_n = n^{4/3}, \qquad x_n = x_n(c) = s_n d_n(c), \qquad d_n(c) = \frac{C_0(c)}{n} - \rho_n(c),$$

Since $C_0(c) = L_c^2/4$ and the contrast is nonzero, $C_0(c) > 0$.

1. **Eventual two-sided bounds.** Theorem 5 supplies $\kappa_c > 0$, the liminf–limsup chain

$$C_0(c)\kappa_c \le \liminf_{n\to\infty} x_n \le \limsup_{n\to\infty} x_n \le 43C_0(c),$$

and, for all sufficiently large n, a clipped-shrinkage rule satisfying

$$\sup_{z\in\mathcal{T}_K^n} R_n(\mathcal{D}^\star, \widehat{\tau}_n^{\mathrm{sh}}; z) \le C_0(c)\{n^{-1} - \kappa_c n^{-4/3}\}.$$

By the definition of $\rho_n(c)$ in Definition 3,

$$\rho_n(c) \le C_0(c)\{n^{-1} - \kappa_c n^{-4/3}\},$$

and hence

$$d_n(c) \ge C_0(c)\kappa_c n^{-4/3}, \qquad x_n = s_n d_n(c) \ge C_0(c)\kappa_c,$$

for all sufficiently large n, using $s_n n^{-4/3} = 1$. On the other side, Theorem 4 gives, for every $n \ge 8$,

$$d_n(c) \le 43C_0(c)n^{-4/3}, \qquad x_n \le 43C_0(c).$$

Thus, eventually,

$$x_n \in [C_0(c)\kappa_c,\, 43C_0(c)]\,, \qquad 0 < x_n.$$

The displayed liminf–limsup inequalities above give the first assertion of the theorem.

2. **Subsequential limits.** Put

$$\mathfrak{L}_c = \left\{y \in \mathbb{R} : \text{there is a strictly increasing } \phi : \mathbb{N} \to \mathbb{N} \text{ with } x_{\phi(n)} \to y\right\}.$$

The eventual containment from the previous step places every tail of (x_n) in the compact interval

$$I_c = [C_0(c)\kappa_c,\, 43C_0(c)]\,.$$

Bolzano–Weierstrass applied to a tail subsequence gives $\mathfrak{L}_c \ne \varnothing$. Also $\mathfrak{L}_c \subseteq I_c$. To see that $\mathfrak{L}_c$ is closed, take $y_j \in \mathfrak{L}_c$ with $y_j \to y$. For each j, choose an index $n_j > n_{j-1}$ such that

$$|x_{n_j} - y_j| < \frac{1}{j},$$

possible because y_j is a subsequential limit. Then $x_{n_j} \to y$, so $y \in \mathfrak{L}_c$. Hence $\mathfrak{L}_c$ is a closed subset of the compact interval I_c, and is compact.

3. **Regular-variation index and the slowly varying criterion.** The index assertion is exactly the regular-variation conclusion in Theorem 5: if $a_n > 0$ is regularly varying with index β, $C > 0$, and

$$a_n d_n(c) \to C,$$

then

$$\beta = \frac{4}{3}.$$

It remains to identify when a normalizer of index 4/3 exists. Suppose first that $a_n > 0$, a_n is regularly varying with index 4/3, $C > 0$, and $a_n d_n(c) \to C$. For $t > 0$, set

$$p_n = a_n d_n(c).$$

Since $\lfloor tn \rfloor \to \infty$,

$$\frac{p_{\lfloor tn \rfloor}}{p_n} \to 1.$$

Moreover,

$$\frac{s_{\lfloor tn \rfloor}}{s_n} \to t^{4/3}, \qquad \frac{a_{\lfloor tn \rfloor}}{a_n} \to t^{4/3}.$$

For all sufficiently large n, the denominators are nonzero, and

$$\frac{x_{\lfloor tn \rfloor}}{x_n} = \frac{s_{\lfloor tn \rfloor} d_{\lfloor tn \rfloor}(c)}{s_n d_n(c)} = \frac{s_{\lfloor tn \rfloor}}{s_n} \frac{p_{\lfloor tn \rfloor}}{p_n} \left(\frac{a_{\lfloor tn \rfloor}}{a_n} \right)^{-1} \longrightarrow t^{4/3} \cdot 1 \cdot t^{-4/3} = 1.$$

Together with the eventual positivity established above, this is the asserted eventual positivity and slow variation of x_n.

Conversely, suppose $x_n > 0$ eventually and

$$\frac{x_{\lfloor tn \rfloor}}{x_n} \to 1 \qquad (t > 0).$$

Choose an integer $N_+ \geq 1$ such that $x_n > 0$ for all $n \geq N_+$, and define a positive sequence b_n on all of $\mathbb{N}$ by assigning arbitrary positive values for $n < N_+$ and setting

$$b_n = x_n \qquad (n \geq N_+).$$

Since $\lfloor tn \rfloor \to \infty$, the finite initial assignments are invisible in the ratio, and hence

$$\frac{b_{\lfloor tn \rfloor}}{b_n} \to 1,$$

so b_n is regularly varying with index 0. Now define a positive sequence a_n on all of $\mathbb{N}$ by assigning arbitrary positive values for $n < N_+$ and setting

$$a_n = \frac{s_n}{b_n} \qquad (n \geq N_+).$$

This totalization is positive because $N_+ \geq 1$, so $s_n = n^{4/3} > 0$ on the displayed tail and $b_n > 0$ everywhere. Since finite initial values do not affect regular variation, and since s_n is regularly varying with index 4/3, the quotient $a_n = s_n / b_n$ on the tail is regularly varying with index 4/3. For all $n \geq N_+$,

$$a_n d_n(c) = \frac{s_n}{x_n} d_n(c) = 1.$$

Thus the desired normalizer exists, with $C = 1$.

4. **Exact rational grid-certificate cluster.** Let q be a rational contrast whose real embedding is c. For each $n \geq 1$, apply Theorem 7 with mesh $M = n$. This gives rational certificate data

$$\pi_n, \quad w_n, \quad u_n, \quad \nu_n, \quad \delta_n$$

such that the exact primal-dual certificate and barycenter identities hold, and such that, writing

$$\mathsf{B}_n^q = B_{n,q}(\nu_n), \qquad \mathsf{U}_n^q = U_{n,n}(q; \pi_n, \delta_n), \qquad \mathsf{T}_n^q = u_n = \Lambda_{n,n}(q),$$

one has

$$\mathsf{B}_n^q \leq \rho_n(c) \leq \mathsf{U}_n^q \leq \mathsf{T}_n^q \leq \mathsf{B}_n^q + \frac{C_0(c)}{4n^2}.$$

Define, with arbitrary zero-th values,

$$G_n = s_n \left\{ \frac{C_0(c)}{n} - \mathsf{T}_n^q \right\}, \quad P_n = s_n \left\{ \frac{C_0(c)}{n} - \mathsf{U}_n^q \right\}, \quad L_n = s_n \left\{ \frac{C_0(c)}{n} - \mathsf{B}_n^q \right\}.$$

The same certificate data make G_n an exact rational grid-certificate cluster for q. The sandwich width satisfies

$$0 \leq s_n(\mathsf{T}_n^q - \mathsf{B}_n^q) \leq s_n \frac{C_0(c)}{4n^2} = \frac{C_0(c)}{4} n^{-2/3} \longrightarrow 0.$$

Since

$$\mathsf{B}_n^q \leq \rho_n(c) \leq \mathsf{T}_n^q,$$

we get

$$|G_n - x_n| \leq s_n(\mathsf{T}_n^q - \mathsf{B}_n^q) \to 0.$$

Similarly, because $\mathsf{U}_n^q \leq \mathsf{T}_n^q$ and $\mathsf{B}_n^q \leq \mathsf{U}_n^q$,

$$|P_n - G_n| = s_n(\mathsf{T}_n^q - \mathsf{U}_n^q) \leq s_n(\mathsf{T}_n^q - \mathsf{B}_n^q) \to 0,$$

and

$$|L_n - G_n| = s_n(\mathsf{T}_n^q - \mathsf{B}_n^q) \to 0.$$

These are precisely

$$G_n - x_n(c) \to 0, \qquad P_n - G_n \to 0, \qquad L_n - G_n \to 0.$$

5. **Transferred certificates for real contrasts.** Apply Theorem 9 to the real contrast c. It gives rational contrasts $c^{(n)}$, certificate values B_n, U_n, and transferred endpoints R_n^-, R_n^+ such that, for every arm a,

$$c_a^{(n)} \to c_a, \qquad n^{5/6}\eta(c, c^{(n)}) \to 0,$$

where

$$\eta(c, c^{(n)}) = \frac{1}{2} \sum_{a \in \mathcal{A}_K} |c_a - c_a^{(n)}|.$$

For every $n \geq 1$, the transferred certificate supplies an exact rational primal-dual certificate, a grid barycenter estimator, and a rational prior for $c^{(n)}$, with

$B_n =$ the lower certificate from the prior, $\qquad U_n =$ the upper certificate from the grid barycenter, and

$$R_n^- = \left(\max\left\{0, \sqrt{B_n} - \eta(c, c^{(n)})\right\}\right)^2, \qquad R_n^+ = \left(\sqrt{U_n} + \eta(c, c^{(n)})\right)^2.$$

The same transferred certificate gives

$$R_n^- \le \rho_n(c) \le R_n^+$$

and, at mesh $M = n$,

$$s_n(R_n^+ - R_n^-) \to 0.$$

Therefore

$$0 \le s_n \left\{\frac{C_0(c)}{n} - R_n^-\right\} - x_n(c) = s_n\{\rho_n(c) - R_n^-\} \le s_n(R_n^+ - R_n^-) \to 0,$$

and

$$0 \le x_n(c) - s_n \left\{\frac{C_0(c)}{n} - R_n^+\right\} = s_n\{R_n^+ - \rho_n(c)\} \le s_n(R_n^+ - R_n^-) \to 0.$$

Thus

$$n^{4/3}\left(\frac{C_0(c)}{n} - R_n^-\right) - x_n(c) \to 0, \qquad n^{4/3}\left(\frac{C_0(c)}{n} - R_n^+\right) - x_n(c) \to 0.$$

6. **The three-arm scalar-score separation.** For $K = 3$ and $c^\dagger = (1, -1/2, -1/2)$, Proposition 2 supplies a full observed arm-label-and-outcome estimator $\widehat{\tau}$ under the product $q^\star$ design $\mathcal{D}^\star$ whose worst-case risk over $\mathcal{T}_3^3$ is

$$\frac{511653}{4000000}.$$

The same result identifies the three-observation scalar sign-score minimax value as

$$V_3^Z = 1 - \frac{\sqrt{3}}{2}$$

and gives the strict numerical chain

$$\frac{511653}{4000000} < \frac{18213}{136000} < 1 - \frac{\sqrt{3}}{2}.$$

Consequently

$$\sup_{z \in \mathcal{T}_3^3} R_3(\mathcal{D}^\star, \widehat{\tau}; z) < V_3^Z,$$

as claimed.

□

Proof of Theorem 12. 1. Suppose first that $n \ge 1$ and $|S_c| = 2$. Choose the two active arms $a_+(c)$ and $a_-(c)$ so that

$$c_{a_+(c)} = \alpha_c, \qquad c_{a_-(c)} = -\alpha_c, \qquad \alpha_c = \frac{L_c}{2},$$

and hence

$$\alpha_c^2 = C_0(c).$$

For $z_K = (t_i)_{i=1}^n \in \mathcal{T}_K^n$, define the active two-arm schedule

$$\mathrm{act}_c(z_K)_i = \big(t_{i,a_+(c)}, t_{i,a_-(c)}\big) \in \mathcal{T}_2.$$

For $z_2 = (s_i)_{i=1}^n \in \mathcal{T}_2^n$, define the sign-schedule embedding $E_c(z_2) \in \mathcal{T}_K^n$ by

$$E_c(z_2)_{i,a_+(c)} = s_{i,1}, \qquad E_c(z_2)_{i,a_-(c)} = s_{i,2}, \qquad E_c(z_2)_{i,a} = 0 \quad (a \notin S_c).$$

Then

$$\tau_c(z_K) = \alpha_c\, \tau_{(1,-1)}\{\mathrm{act}_c(z_K)\}, \qquad \tau_c\{E_c(z_2)\} = \alpha_c\, \tau_{(1,-1)}(z_2).$$

By Theorem 4,

$$\rho_n(c) = C_0(c)\rho_{n,2}.$$

2. Apply Theorem 1 with $K = 2$ and contrast $(1, -1)$. It gives a two-arm orbit procedure q_2 and a prior ν_2 on $\mathcal{M}_{n,2}$ such that

$$R^{\mathrm{orb}}_{(1,-1)}(q_2; m) \le G_n((1, -1)) \quad (m \in \mathcal{M}_{n,2}),$$

and, for every two-arm orbit procedure q_2',

$$G_n((1, -1)) \le \mathbb{E}_{\nu_2}\left[R^{\mathrm{orb}}_{(1,-1)}(q_2'; m)\right].$$

The same result identifies $G_n((1, -1)) = \rho_{n,2}$. Let p_2 be the invariant labeled procedure induced by q_2. For $m \in \mathcal{M}_{n,2}$, set

$$\mathrm{Orb}_2(m) = \{z_2 \in \mathcal{T}_2^n : \mathrm{counts}(z_2) = m\},$$

and define the labeled prior

$$\mathrm{prior}_2(z_2) = \frac{\nu_2(\mathrm{counts}(z_2))}{|\mathrm{Orb}_2(\mathrm{counts}(z_2))|}.$$

Since $z_2 \in \mathrm{Orb}_2(\mathrm{counts}(z_2))$, the denominator is positive for every schedule in the displayed formula. The invariant correspondence gives

$$R_n^{(1,-1)}(p_2; z_2) = R^{\mathrm{orb}}_{(1,-1)}\big(q_2; \mathrm{counts}(z_2)\big),$$

so

$$R_n^{(1,-1)}(p_2; z_2) \le \rho_{n,2} \quad (z_2 \in \mathcal{T}_2^n).$$

For any two-arm labeled procedure p_2', symmetrizing p_2' over unit permutations and passing through the invariant-orbit correspondence yields an orbit procedure whose orbit risk, averaged under ν_2, is bounded above by the prior_2-Bayes risk of p_2'. Therefore

$$\rho_{n,2} \le \mathbb{E}_{\mathrm{prior}_2}\left[R_n^{(1,-1)}(p_2'; z_2)\right].$$

3. Let p_K be the sign-pair lift of p_2, and let

$$\mathrm{prior}_K = (E_c)_{\#}\mathrm{prior}_2, \qquad \mathbb{E}_{\mathrm{prior}_K} F = \mathbb{E}_{\mathrm{prior}_2}[F\{E_c(z_2)\}]\,.$$

The lifted experiment observes exactly the two active coordinates and multiplies the two-arm estimate by α_c. Hence, for every $z_K \in \mathcal{T}_K^n$,

$$R_n^c(p_K; z_K) = \alpha_c^2 R_n^{(1,-1)}\big(p_2; \mathrm{act}_c(z_K)\big) \leq C_0(c)\rho_{n,2} = \rho_n(c).$$

Now fix an arbitrary K-arm labeled procedure $p'_K = (\mathcal{D}'_K, \widehat{\tau}'_K)$. Define the induced two-arm procedure $p'_{2\leftarrow K}$ by the following Rao–Blackwellization. Let $\widetilde{\mathcal{D}}'_K$ be the product experiment in which $\mathfrak{a} \in \mathcal{A}_K^n$ is drawn from $\mathcal{D}'_K$ and, independently, $\xi_i \in \{0,1\}$ is fair for each unit. Define the sign-group assignment map $\phi_c : \mathcal{A}_K^n \times \{0,1\}^n \to \mathcal{A}_2^n$ by

$$\phi_c(\mathfrak{a}, \xi)_i = \begin{cases} 1, & c_{\mathfrak{a}_i} > 0, \\ 2, & c_{\mathfrak{a}_i} < 0, \\ 1, & c_{\mathfrak{a}_i} = 0 \text{ and } \xi_i = 1, \\ 2, & c_{\mathfrak{a}_i} = 0 \text{ and } \xi_i = 0. \end{cases}$$

Thus the fair bits are used exactly on zero-coefficient assignments. Define the inactive-zero pseudo-observation by

$$\widetilde{y}_c(\mathfrak{a}, y)_i = \begin{cases} y_i, & c_{\mathfrak{a}_i} \neq 0, \\ 0, & c_{\mathfrak{a}_i} = 0, \end{cases}$$

where 0 is the binary response value assigned to inactive coordinates. For a two-arm observed vector y, set

$$\zeta'_c(\mathfrak{a}, \xi; y) = \alpha_c^{-1}\widehat{\tau}'_K\{\mathfrak{a}, \widetilde{y}_c(\mathfrak{a}, y)\}.$$

Here $\alpha_c = L_c/2$ and $\alpha_c > 0$, while the K-arm estimator is clipped to $[-L_c/2, L_c/2] = [-\alpha_c, \alpha_c]$. Hence

$$\zeta'_c(\mathfrak{a}, \xi; y) \in [-1, 1].$$

The induced two-arm procedure $p'_{2\leftarrow K}$ has assignment law $(\phi_c)_{\#}\widetilde{\mathcal{D}}'_K$ and estimator

$$\widehat{\tau}'_{2\leftarrow K}(\mathfrak{b}, y) = \begin{cases} \mathbb{E}_{\widetilde{\mathcal{D}}'_K}[\zeta'_c(\mathfrak{a}, \xi; y) \mid \phi_c(\mathfrak{a}, \xi) = \mathfrak{b}]\,, & \mathbb{P}_{\widetilde{\mathcal{D}}'_K}\{\phi_c = \mathfrak{b}\} > 0, \\ 0, & \mathbb{P}_{\widetilde{\mathcal{D}}'_K}\{\phi_c = \mathfrak{b}\} = 0. \end{cases}$$

The displayed interval containment makes this a legitimate two-arm procedure. For a fixed $z_2 \in \mathcal{T}_2^n$, write $Y_{z_2}(\mathfrak{b})$ for the observed outcome vector under two-arm assignment $\mathfrak{b}$. Conditional Jensen, applied on each fiber of ϕ_c, gives

$$R_n^{(1,-1)}(p'_{2\leftarrow K}; z_2) \leq \mathbb{E}_{\widetilde{\mathcal{D}}'_K}\Big[\big\{\zeta'_c\big(\mathfrak{a}, \xi; Y_{z_2}(\phi_c(\mathfrak{a}, \xi))\big) - \tau_{(1,-1)}(z_2)\big\}^2\Big]\,.$$

For every $(\mathfrak{a}, \xi)$, the sign-schedule embedding and inactive-zero convention give

$$\widetilde{y}_c\{\mathfrak{a}, Y_{z_2}(\phi_c(\mathfrak{a}, \xi))\} = Y_{E_c(z_2)}(\mathfrak{a}).$$

Indeed, a positive-coefficient assigned arm reads the first two-arm coordinate, a negative-coefficient assigned arm reads the second, and an inactive assigned arm has binary response value 0 on both sides: $E_c(z_2)$ fixes inactive coordinates at 0, while $\widetilde{y}_c$ replaces inactive observations by 0. Also,

$$\begin{aligned}\tau_c\{E_c(z_2)\} &= \frac{1}{n}\sum_{i=1}^{n}\left(\sum_{a:c_a>0} c_a(z_2)_{i,1} + \sum_{a:c_a<0} c_a(z_2)_{i,2}\right)\\ &= \frac{1}{n}\sum_{i=1}^{n}\left(\alpha_c(z_2)_{i,1} - \alpha_c(z_2)_{i,2}\right) = \alpha_c\tau_{(1,-1)}(z_2),\end{aligned}$$

where the middle equality uses $\sum_{a:c_a>0} c_a = \alpha_c$ and $\sum_{a:c_a<0} c_a = -\alpha_c$. Substituting these two identities into the preceding Jensen bound yields

$$\begin{aligned}R_n^{(1,-1)}(p'_{2\leftarrow K}; z_2) &\le \mathbb{E}_{\widetilde{\mathcal{D}}'_K}\left[\alpha_c^{-2}\left\{\widehat{\tau}'_K\big(\mathfrak{a}, Y_{E_c(z_2)}(\mathfrak{a})\big) - \tau_c\{E_c(z_2)\}\right\}^2\right]\\ &= \alpha_c^{-2}R_n^c\big(p'_K; E_c(z_2)\big),\end{aligned}$$

since $\widetilde{\mathcal{D}}'_K$ has $\mathcal{D}'_K$ as its $\mathfrak{a}$-marginal. Integrating this inequality under prior_2 and using the two-arm Bayes lower bound from the preceding step gives

$$\rho_{n,2} \le \mathbb{E}_{\mathrm{prior}_2}\left[R_n^{(1,-1)}(p'_{2\leftarrow K}; z_2)\right] \le \alpha_c^{-2}\mathbb{E}_{\mathrm{prior}_K}\left[R_n^c(p'_K; z_K)\right].$$

Multiplying by $\alpha_c^2 = C_0(c)$ and using $\rho_n(c) = C_0(c)\rho_{n,2}$ yields

$$\rho_n(c) \le \mathbb{E}_{\mathrm{prior}_K}\left[R_n^c(p'_K; z_K)\right].$$

Together with the definitions of $p_2, \mathrm{prior}_2, p_K, \mathrm{prior}_K$, this proves the two-active-arm value and saddle-transfer assertion.

4. For arbitrary active support, Theorem 5 supplies

$$C_0(c)\kappa_c \le \liminf_{n\to\infty} n^{4/3}d_n(c) \le \limsup_{n\to\infty} n^{4/3}d_n(c) \le 43C_0(c).$$

It also supplies an integer N such that, for every $n \ge N$, the clipped-shrinkage procedure satisfies the contrast-c risk bound

$$\sup_{z\in\mathcal{T}_K^n} R_n^c(\mathcal{D}^\star, \widehat{\tau}_n^{\mathrm{sh}}; z) \le C_0(c)\left\{n^{-1} - \kappa_c n^{-4/3}\right\}.$$

These are exactly the second-order scale and shrinkage-attainment assertions.

5. By Theorem 9, for every rational nonzero zero-sum contrast q and every $n, M \ge 1$, there exist real numbers B, U, R^-, R^+ forming a transferred certificate for (n, M, c, q). The same result supplies rational contrasts $c^{(n)}$ and real sequences B_n, U_n, R_n^-, R_n^+ such that

$$c_a^{(n)} \to c_a \quad (a \in \mathcal{A}_K), \qquad n^{5/6}\eta(c, c^{(n)}) \to 0,$$

for every $n \ge 1$, the tuple (B_n, U_n, R_n^-, R_n^+) is a transferred certificate for $(n, n, c, c^{(n)})$, and

$$n^{4/3}(R_n^+ - R_n^-) \to 0.$$

Finally, if q represents c coefficientwise, the exact rational part of the same result gives, for every $n \ge 1$, a transferred certificate for (n, n, c, q) satisfying

$$U - B \le \frac{C_0(c)}{4n^2}.$$

This proves all real-contrast transfer certificate assertions.

6. Apply Theorem 10 with the given $n, M \geq 1$, and write the objects it supplies as π, w, u, ν, δ. For these, $B_n(\nu)$ is the lower certificate formed from the response-count prior ν, and $U_{n,M}$ is the upper certificate formed from π and δ. The cited result gives the exact primal-dual certificate and barycenter conditions for $c^\dagger$, and its returned inequalities include

$$B_n(\nu) \leq \rho_n^\dagger, \qquad \rho_n^\dagger \leq U_{n,M}, \qquad U_{n,M} \leq B_n(\nu) + \frac{1}{4M^2}.$$

These are exactly the displayed sandwich in the three-arm rational certificate assertion.

□

References

P. M. Aronow and Patrick Lopatto. Minimax unbiased estimation for finite populations with bounded outcomes, 2026. URL https://arxiv.org/abs/2605.20572.

Susan Athey and Guido W. Imbens. The econometrics of randomized experiments. In Abhijit Vinayak Banerjee and Esther Duflo, editors, *Handbook of Economic Field Experiments*, volume 1, pages 73–140. Elsevier, 2017. doi: 10.1016/bs.hefe.2016.10.003. URL https://doi.org/10.1016/bs.hefe.2016.10.003.

Anthony C. Atkinson, Alexander N. Donev, and Randall D. Tobias. *Optimum Experimental Designs, with SAS*. Oxford University Press, Oxford, 2007. ISBN 9780199296590.

Yuehao Bai. Why randomize? minimax optimality under permutation invariance. *Journal of Econometrics*, 232(2):565–575, 2023. doi: 10.1016/j.jeconom.2021.10.009. URL https://doi.org/10.1016/j.jeconom.2021.10.009.

Yuehao Bai, Azeem M. Shaikh, and Max Tabord-Meehan. A primer on the analysis of randomized experiments and a survey of some recent advances, 2024. URL https://arxiv.org/abs/2405.03910.

Guillaume W. Basse, Yi Ding, and Panos Toulis. Minimax designs for causal effects in temporal experiments with treatment habituation. *Biometrika*, 110(1):155–168, 2023. URL https://arxiv.org/abs/1908.03531.

Peter J. Bickel. Minimax estimation of the mean of a normal distribution when the parameter space is restricted. *The Annals of Statistics*, 9(6):1301–1309, 1981. doi: 10.1214/aos/1176345646. URL https://doi.org/10.1214/aos/1176345646.

Tirthankar Dasgupta, Natesh S. Pillai, and Donald B. Rubin. Causal inference from 2^K factorial designs by using potential outcomes. *Journal of the Royal Statistical Society: Series B*, 77(4):727–753, 2015. doi: 10.1111/rssb.12085. URL https://doi.org/10.1111/rssb.12085.

Michał Dereziński, Kenneth L. Clarkson, Michael W. Mahoney, and Manfred K. Warmuth. Minimax experimental design: Bridging the gap between statistical and worst-case approaches to least squares regression, 2019. URL https://arxiv.org/abs/1902.00995.

Peng Ding. *A First Course in Causal Inference*. Chapman and Hall/CRC, Boca Raton, 2024. ISBN 9781032758626. URL https://www.routledge.com/A-First-Course-in-Causal-Inference/Ding/p/book/9781032758626.

Dean Eckles, Brian Karrer, and Johan Ugander. Design and analysis of experiments in networks: Reducing bias from interference, 2014. URL https://arxiv.org/abs/1404.7530.

Israel Feldman. Constrained minimax estimation of the mean of the normal distribution with known variance. *The Annals of Statistics*, 19(4):2259–2265, 1991. doi: 10.1214/aos/1176348398. URL https://doi.org/10.1214/aos/1176348398.

Ronald A. Fisher. *The Design of Experiments*. Oliver and Boyd, Edinburgh, 1935. URL https://archive.org/details/in.ernet.dli.2015.502684.

David A. Freedman. On regression adjustments in experiments with several treatments. *The Annals of Applied Statistics*, 2(1):176–196, 2008. doi: 10.1214/07-AOAS143. URL https://doi.org/10.1214/07-AOAS143.

Siegfried Gabler. *Minimax Solutions in Sampling from Finite Populations*, volume 64 of *Lecture Notes in Statistics*. Springer-Verlag, Berlin, 1990. ISBN 9783540973584. URL https://books.google.com/books?id=dFwZAQAAIAAJ.

Richard D. Gill and Boris Y. Levit. Applications of the van trees inequality: A bayesian cramér-rao bound. *Bernoulli*, 1(1/2):59–79, 1995. doi: 10.2307/3318681. URL https://doi.org/10.2307/3318681.

Scott M. Hammer, David A. Katzenstein, Michael D. Hughes, Heather Gundacker, Robert T. Schooley, Richard H. Haubrich, W. Keith Henry, Michael M. Lederman, John P. Phair, Marek Niu, Martin S. Hirsch, and Thomas C. Merigan. A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with CD4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine*, 335(15):1081–1090, 1996. doi: 10.1056/NEJM199610103351501. URL https://doi.org/10.1056/NEJM199610103351501.

Morris H. Hansen, William N. Hurwitz, and William G. Madow. *Sample Survey Methods and Theory*. John Wiley and Sons, New York, 1953.

Christopher Harshaw, Fredrik Sävje, Daniel Spielman, and Peng Zhang. Balancing covariates in randomized experiments with the gram-schmidt walk design, 2024. URL https://arxiv.org/abs/1911.03071. Journal of the American Statistical Association, 119(548), 2934–2946.

Paul W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945–960, 1986. doi: 10.1080/01621459.1986.10478354. URL https://doi.org/10.1080/01621459.1986.10478354.

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. URL https://doi.org/10.1080/01621459.1952.10483446.

Yuchen Hu, Henry Zhu, Emma Brunskill, and Stefan Wager. Minimax-regret sample selection in randomized experiments, 2024. URL https://arxiv.org/abs/2403.01386.

Peter Hull. Optimal experimental design and estimation when potential outcomes are bounded, 2026. URL https://arxiv.org/abs/2608.09812.

I. A. Ibragimov and R. Z. Has'minskii. *Statistical Estimation: Asymptotic Theory*, volume 16 of *Applications of Mathematics*. Springer-Verlag, New York, 1981. ISBN 9780387905235.

Guido W. Imbens and Donald B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge University Press, New York, 2015. doi: 10.1017/CBO9781139025751. URL https://doi.org/10.1017/CBO9781139025751.

Guido W. Imbens and Jeffrey M. Wooldridge. Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86, 2009. doi: 10.1257/jel.47.1.5. URL https://doi.org/10.1257/jel.47.1.5.

Iain M. Johnstone and Brenda MacGibbon. Minimax estimation of a constrained poisson vector. *The Annals of Statistics*, 20(2):807–831, 1992.

Nathan Kallus. Optimal a priori balance in the design of controlled experiments. *Journal of the Royal Statistical Society: Series B*, 80(1):85–112, 2018. doi: 10.1111/rssb.12240. URL https://doi.org/10.1111/rssb.12240.

Nathan Kallus. On the optimality of randomization in experimental design: How to randomize for minimax variance and design-based inference, 2020. URL https://arxiv.org/abs/2005.03151.

Vardis Kandiros, Charilaos Pipis, Constantinos Daskalakis, and Christopher Harshaw. The conflict graph design: Estimating causal effects under arbitrary neighborhood interference, 2024. URL https://arxiv.org/abs/2411.10908.

Vardis Kandiros, Christopher Harshaw, and Fredrik Sävje. A design-based minimax theory for network experiments, 2026. URL https://arxiv.org/abs/2608.04909.

Jack Kiefer. Optimum experimental designs. *Journal of the Royal Statistical Society: Series B*, 21(2):272–319, 1959.

Jack Kiefer. General equivalence theory for optimum designs (approximate theory). *The Annals of Statistics*, 2(5):849–879, 1974. doi: 10.1214/aos/1176342810. URL https://doi.org/10.1214/aos/1176342810.

Lucien Le Cam. *Asymptotic Methods in Statistical Decision Theory*. Springer-Verlag, New York, 1986. ISBN 9780387963075.

B. Ya. Levit. On asymptotic minimax estimates of the second order. *Theory of Probability and Its Applications*, 25(3):552–568, 1981. doi: 10.1137/1125066. URL https://doi.org/10.1137/1125066.

B. Ya. Levit. Second-order asymptotic optimality and positive solutions of schrödinger's equation. *Theory of Probability and Its Applications*, 30(2):333–363, 1986.

Xinran Li, Peng Ding, and Donald B. Rubin. Asymptotic theory of rerandomization in treatment-control experiments. *Proceedings of the National Academy of Sciences*, 115(37):9157–9162, 2018. doi: 10.1073/pnas.1808191115. URL https://doi.org/10.1073/pnas.1808191115.

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

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. URL https://doi.org/10.1214/12-AOS1008.

M. S. Pinsker. Optimal filtering of square-integrable signals in gaussian noise. *Problems of Information Transmission*, 16(2):120–133, 1980.

C. Radhakrishna Rao. Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society*, 37:81–91, 1945.

Evan T. R. Rosenman and Kristen B. Hunter. Adaptive experimental design using shrinkage estimators, 2026. URL https://arxiv.org/abs/2602.07404.

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. URL https://doi.org/10.1037/h0037350.

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

Kirstine Smith. On the standard deviations of adjusted and interpolated values of an observed polynomial function and its constants and the guidance they give towards a proper choice of the distribution of observations. *Biometrika*, 12(1/2):1–85, 1918. doi: 10.1093/biomet/12.1-2.1. URL https://doi.org/10.1093/biomet/12.1-2.1.

Jerzy Splawa-Neyman. On the application of probability theory to agricultural experiments. essay on principles. section 9. *Statistical Science*, 5(4):465–472, 1990. doi: 10.1214/ss/1177012031. URL https://doi.org/10.1214/ss/1177012031. Translated by D. M. Dabrowska and T. P. Speed.

Timothy Sudijono, Edgar Dobriban, and Eric Tchetgen Tchetgen. Sharp minimax theory for randomized experiments, 2026. URL https://arxiv.org/abs/2608.13822.

Carl-Erik Särndal, Bengt Swensson, and Jan Wretman. *Model Assisted Survey Sampling*. Springer, New York, 1992. ISBN 9780387975283.

Harry L. Van Trees. *Detection, Estimation, and Modulation Theory, Part I*. John Wiley and Sons, New York, 1968.

Abraham Wald. *Statistical Decision Functions*. John Wiley and Sons, New York, 1950. URL https://archive.org/details/statisticaldecis0000abra_t1e7.