CERTIFYING PAIRWISE READOUTS UNDER RANK-TWO REPRESENTATION UNCERTAINTY

Anonymous authors
Paper under double-blind review

ABSTRACT

In graph-based decision systems, some node representations may be missing when a score must be compared with a threshold. Certifying the decision then requires a bound over all admissible completions. We study this problem for nonnegative weighted sums of pairwise inner products. Unknown vectors vary independently in the unit cube, with coordinatewise-bounded orthogonal residuals from a supplied rational rank-two subspace; observed vectors lie in that subspace. In a computable small-radius regime, we construct the exact projected domain and ordered sets of projected candidates, reducing their coupled optimization to a polynomial-size minimum-cut problem. Feasible lifting provides a witness and a quadratic-width interval for the original maximum, whose first-order perturbation coefficient is exactly recoverable. Thresholding the second-order coefficient is strongly NP-complete in a fixed sixteen-dimensional geometry, even with a unique zero-radius optimizer. At every fixed sufficiently small positive rational radius in that geometry, maximum-value thresholding remains strongly NP-complete and maximization admits no fully polynomial-time approximation scheme unless $\mathrm{P} = \mathrm{NP}$. Exact comparisons on constructed instances quantify certificate quality, while coupled multi-label examples demonstrate the role of joint optimization. As an application, a monotone deviation certificate supports selecting vectors to fix at reference values. Together, these results separate efficient certification from exact residual optimization in the rank-two uncertainty model.

1 INTRODUCTION

Learning from incomplete data may require downstream decisions before missing quantities are uniquely determined. Matrix completion reconstructs missing entries using structural assumptions and related data (Fan et al., 2024; Jalan et al., 2025); domain constraints can improve inference for unobserved outcomes (Balachandar et al., 2024). Threshold decisions require assessing ambiguity across all completions compatible with the assumed structure.

Consider a graph score summing inner products between node representations with nonnegative weights. With some vectors missing, one completion gives only one score. Certifying that no allowed completion exceeds the threshold requires an upper bound; a feasible completion above the threshold supplies a violation witness. The task is joint worst-case optimization over missing vectors.

Reasoning over compatible sets underlies robust reward learning (Lazzati & Metelli, 2026) and causal sensitivity analysis (Frauen et al., 2024). We study how shared representation geometry affects computation. Even with low-dimensional vectors, the coupled problem grows as their number increases. Fixed-global-rank box algorithms (Hladík et al., 2021) and PSD rank-collapse methods (Soltanalian & Mousavi, 2026) address a different structure: here the full quadratic interaction matrix can be indefinite and grow in rank. We ask when shared geometry permits efficient worst-case certification with arbitrarily many coupled vectors, and which parts of that response remain hard.

We study this question in a rank-two model. Unknown vectors vary independently in the unit cube, with coordinatewise-bounded orthogonal residuals from a supplied rational two-dimensional subspace. Observed vectors lie in that subspace and induce nonnegative linear terms. The subspace and radius specify the admissible completions. At zero radius, nominal optimization maximizes the score over all exact-subspace completions; positive radius introduces additional residual directions.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

An explicit planar projection and an optimum-preserving cone reduce the projected problem to ordered candidate sets and minimum cut. Lifting the optimizer back to the original domain yields a feasible witness and a quadratic-width interval for the maximum. This interval retains the full first-order response, while the remaining residual optimization marks a complexity boundary at second order.

Contributions.

- **Exact projected optimization and a quadratic-width certificate.** Within a computable small-radius regime, we construct the projected feasible set explicitly and solve the coupled projected problem in polynomial bit time for an unrestricted number of unknown vectors. A feasible lift and certified upper bound enclose the original maximum with width at most $2nW\varepsilon^2$, where n is the ambient dimension and W the total unknown-pair weight (Theorem 1, Section 3). An explicit instance attains width exactly $2nW\varepsilon^2$ for this certificate (Appendix F.5).
- **A separation between first-order recovery and residual optimization.** The first-order perturbation coefficient is exactly computable, whereas thresholding the second-order coefficient is strongly NP-complete in a fixed sixteen-dimensional geometry with a unique nominal optimizer (Theorem 2). At every fixed sufficiently small positive rational radius in that geometry, maximum-value thresholding is strongly NP-complete and maximization has no FPTAS unless $\mathrm{P} = \mathrm{NP}$ (Theorem 3, Section 4).

Application to disclosure. The certificate also bounds the score deviation when selected vectors are fixed to reference values. We derive an upper bound that is monotone in the unknown set, with an explicit additive error. Standard deletion and hitting-set arguments then yield certified disclosure sets and cost comparisons across tolerances (Section 5); the new ingredient is the computable deviation certificate.

2 Model and main results

Fixing observed vectors turns their interactions with unknown vectors into linear terms; unknown–unknown interactions remain coupled. The input specifies a rational two-dimensional subspace $S \subset \mathbb{R}^n$ and rational residual radius $\varepsilon \geq 0$. Let P be the Euclidean orthogonal projector onto S, and assume $C = S \cap \mathbb{R}^n_+$ spans S. The admissible representations form the residual tube:

$$K_\varepsilon = \{b \in [0,1]^n : \|(I-P)b\|_\infty \leq \varepsilon\}, \qquad Q_\varepsilon = PK_\varepsilon. \tag{1}$$

The budget bounds the orthogonal residual in infinity norm. The g unknown vectors form the columns of $B = (b_1, \ldots, b_g)$ and vary independently over K_ε, with score:

$$J(B) = \sum_{i<j} w_{ij}\langle b_i, b_j\rangle + \sum_i \langle h_i, b_i\rangle, \qquad w_{ij} \geq 0, \quad h_i \in C. \tag{2}$$

Each cross-column pair is counted once, with no self-interactions. All data are rational; n, g, and all binary encoding lengths contribute to input size. For projected columns $s = (s_1, \ldots, s_g)$, write:

$$U_\varepsilon = \max_{B \in K_\varepsilon^g} J(B), \quad V_\varepsilon = \max_{s \in Q_\varepsilon^g} J(s), \quad W = \sum_{i<j} w_{ij}. \tag{3}$$

Native fields and threshold decisions. We call the nonnegative linear coefficients lying in S *native fields*; the model permits every $h_i \in C$. For a readout $\beta + \sum_{i<j} w_{ij}\langle x_i, x_j\rangle$, fixing public representations $a_j \in S \cap [0,1]^n$ produces a public constant and fields $h_i = \sum_{j \text{ observed}} w_{ij}a_j \in C$. Subtracting the bias and public–public terms from the readout threshold gives the threshold for J. For rational θ and $0 < \varepsilon \leq \varepsilon_B(S)$, the oracle returns a feasible lift and an interval $[\underline{U}, \overline{U}]$ containing the maximum U_ε. The tests $\overline{U} \leq \theta$ and $\underline{U} > \theta$ certify safety and a feasible violation, respectively; a straddling interval returns Unknown. These certificates apply to completions in (1). An admissible encoder preimage of a violating representation also supplies an input-space witness.

Table 1: Computational guarantees within the native rank-two tube model. The number of unknown vectors is unrestricted. Hardness statements use binary rational inputs.

Quantity	Guarantee	Geometry needed
U_0	Exact polynomial time	Rational rank two
V_ε	Exact polynomial time	$\varepsilon \le \varepsilon_B(S)$
Λ	Exact polynomial time	Rational rank two
U_ε enclosure	Width $\le 2nW\varepsilon^2$	$\varepsilon \le \varepsilon_B(S)$
Γ threshold	Strongly NP-complete	Constructed fixed rational $S \subset \mathbb{R}^{16}$
U_ε threshold	Strongly NP-complete	Same S; fixed rational $0 < \varepsilon \le \varepsilon_*$

Theorem 1 (Projected certificate)**.** *There is a positive rational $\varepsilon_B(S)$, computable in polynomial time and with polynomial encoding length, such that for every rational $0 < \varepsilon \le \varepsilon_B(S)$, V_ε is exactly computable in polynomial time. One can also return $B^{\mathrm{lift}} \in K_\varepsilon^g$ satisfying:*

$$J(B^{\mathrm{lift}}) \le U_\varepsilon \le V_\varepsilon + nW\varepsilon^2, \tag{4}$$

$$V_\varepsilon + nW\varepsilon^2 - J(B^{\mathrm{lift}}) \le 2nW\varepsilon^2. \tag{5}$$

In particular, $|U_\varepsilon - V_\varepsilon| \le nW\varepsilon^2$.

The feasible lift and the certified upper bound bracket the maximum within a quadratic-width interval.

Theorem 2 (Perturbation coefficients)**.** *For fixed inputs other than ε, there is a positive rational ε_0 of polynomial encoding length, computable in polynomial time, such that for $0 \le \varepsilon \le \varepsilon_0$, we have:*

$$U_\varepsilon = U_0 + \Lambda\varepsilon + \Gamma\varepsilon^2. \tag{6}$$

The coefficient $\Lambda \ge 0$ is exactly computable in polynomial time. Deciding $\Gamma \ge \tau$ for a rational threshold is NP-complete. Its restriction to a fixed rational $S \subset \mathbb{R}^{16}$ is strongly NP-complete, even with strictly positive native fields in four fixed directions and a unique nominal maximizer.

The linear response is exactly recoverable, while residual interactions carry the second-order combinatorial difficulty.

Theorem 3 (Fixed positive radius)**.** *There exist a fixed rational $S \subset \mathbb{R}^{16}$ and a fixed rational $\varepsilon_* > 0$ such that, for every fixed rational $0 < \varepsilon \le \varepsilon_*$, deciding $U_\varepsilon \ge \tau$ is strongly NP-complete. Unless* $\mathrm{P} = \mathrm{NP}$*, maximizing* (2) *over this fixed-radius domain has no fully polynomial-time approximation scheme (FPTAS).*

The difficulty persists with geometry and radius both fixed.

The box constraints and orthogonal-residual budget generate three residual labels for the hardness proofs. In the stated hard classes, thresholding is NP-complete and value computation NP-hard.

3 An exact projected oracle

The certification problem couples unknown vectors through their inner products. Projection is useful because orthogonality and the native fields leave only residual–residual interactions outside the projected score. We construct the projected domain, optimize on ordered response chains, and lift the selected points to the original domain. Figure 1 illustrates these steps and the possible nonuniqueness of a feasible lift.

3.1 An explicit description of the projected domain

To optimize in the projected space, we first need an exact test for whether a projected point has an original feasible preimage. Fix a rational basis matrix $R \in \mathbb{Q}^{n\times 2}$, put $G = R^\top R$, and write $s = Rx$. The box and the infinity-norm budget on the orthogonal residual give:

$$R^\top b = Gx, \qquad L_j(x) := \max\{0, (Rx)_j - \varepsilon\} \le b_j \le T_j(x) := \min\{1, (Rx)_j + \varepsilon\}. \tag{7}$$

The intervals are nonempty exactly on $\Omega_\varepsilon = \{x : -\varepsilon \le Rx \le 1 + \varepsilon\}$. Projected-domain membership is therefore a two-dimensional moment problem.

Let $r_j = R_j^\top$ and let $\mathcal{N}$ contain both perpendicular orientations to every nonzero r_j. The attainable moments form the zonotope $Z(x) = \sum_j r_j[L_j(x), T_j(x)]$, whose support inequalities give the exact membership test:

$$\begin{aligned} Gx \in Z(x) \iff \nu^\top Gx \le F_\nu(x) \quad (\nu \in \mathcal{N}), \\ F_\nu(x) = \sum_{\nu^\top r_j \ge 0} (\nu^\top r_j) T_j(x) + \sum_{\nu^\top r_j < 0} (\nu^\top r_j) L_j(x). \end{aligned} \tag{8}$$

The clipping lines $(Rx)_j = \varepsilon$ and $(Rx)_j = 1 - \varepsilon$ divide Ω_ε into $O(n^2)$ cells. On each cell A, the clipped support function F_ν has an affine expression $F_{\nu,A}$. Each branch is a global affine upper bound, and some branch is attained at every point. Continuity includes cell boundaries, so we have:

$$F_\nu(x) = \min_A F_{\nu,A}(x), \qquad \nu^\top Gx \le F_\nu(x) \iff \nu^\top Gx \le F_{\nu,A}(x) \quad \text{for every } A. \tag{9}$$

Intersecting these halfplanes with Ω_ε gives the exact projected domain in basis coordinates.

Lemma 4 (Explicit projection)**.** *For rational $\varepsilon > 0$, Q_ε has a description by $O(n^3)$ rational halfplanes, and hence $O(n^3)$ vertices, constructible in polynomial bit time.*

At $\varepsilon = 0$, use $0 \le Rx \le 1$ directly.

3.2 An optimum-preserving cone and ordered responses

Projection alone does not provide the ordered interactions needed for a cut formulation: Q_ε can contain negative coordinates and obtuse pairs. An optimum-preserving restriction to a common nonobtuse cone supplies the required geometry. The starting point is the dual-cone inclusion within S:

$$Q_\varepsilon \subset C_S^* := \{s \in S : \langle s, c\rangle \ge 0 \ \forall c \in C\}, \qquad \langle Pb, c\rangle = \langle b, c\rangle \ge 0. \tag{10}$$

Consequently every projected conditional field $t_i = h_i + \sum_{j \ne i} w_{ij} s_j$ lies in C_S^*.

Let u, v be the extreme rays of C. When $\langle u, v\rangle > 0$, define:

$$q = \frac{\langle u, v\rangle}{4(\|u\|^2 + \|v\|^2)}, \qquad D = \operatorname{cone}(u - qv, v - qu). \tag{11}$$

Then $0 < q \le 1/8$, $C \subset D$, and $\langle u - qv, v - qu\rangle = \langle u, v\rangle(3/4 + q^2) > 0$. Choose ℓ, k with $v_\ell = 0 < u_\ell$ and $u_k = 0 < v_k$, and set:

$$\gamma = \max_j (u_j + v_j), \quad M = \max\{u_\ell^{-1}, v_k^{-1}\}, \quad \varepsilon_B = \frac{q}{2\gamma M}, \quad c_0 = \frac{u + v}{\gamma} \in K_0. \tag{12}$$

For $0 < \varepsilon \le \varepsilon_B$, the residual budget makes $c_0 - s$ a strictly positive combination of u and v whenever $s \in Q_\varepsilon \setminus D$. Thus $\langle t, c_0 - s\rangle > 0$ for every nonzero $t \in C_S^*$. Appendix A gives coefficient bounds establishing this replacement property. For the instances in this paper, ε_B equals $1/64$ and $1/72$ for the two closed-form examples (Appendices C and F.5) and is at least $1/128$ for every member of the growing family in Appendix F.2.

At a global optimum, an outside column cannot have $t_i \ne 0$, since replacement would strictly improve the value. If $t_i = 0$, replacing it by c_0 preserves optimality. Earlier replacements remain in D, so at most g replacements give:

$$V_\varepsilon = \max_{s \in (Q_\varepsilon \cap D)^g} J(s). \tag{13}$$

If $u \perp v$, their nonnegative supports are disjoint, P is entrywise nonnegative, and $Q_\varepsilon \subset C$ for all radii; take $D = C$ and $\varepsilon_B = 1$. Both cases preserve an optimum for arbitrary nonnegative weights, with optimizer ties allowed.

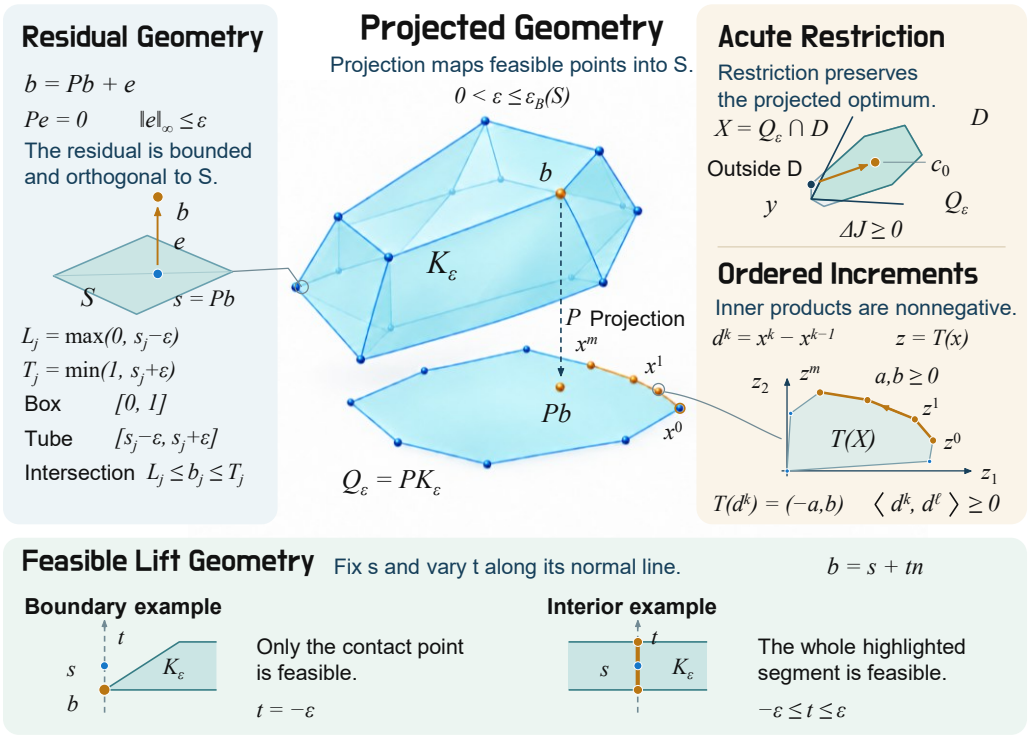

Figure 1: **Projected geometry and feasible lifts.** The schematic pair illustrates an arbitrary feasible point b, its projection Pb, and their bounded orthogonal residual. Under Theorem 1's assumptions and $0 < \varepsilon \leq \varepsilon_B(S)$, restriction to $Q_\varepsilon \cap D$ preserves the projected optimum; D is the common nonobtuse cone of Section 3.2, and ordered response increments have nonnegative pairwise inner products. The bottom panels show singleton and segment feasible lift sets for two exact points in $S = \{z = x + y\}$ at $\varepsilon = 1/64$: blue marks projections, dashed lines normal fibres, and amber feasible intersections. Upper panels are schematic; Appendix C gives exact coordinates and local-section conventions.

The Gram inverse orders interactions. Write $D = \mathrm{cone}(r, s)$ and $X = Q_\varepsilon \cap D$. A field in D scores x by a nonnegative combination of the response coordinates $T(x) = (\langle r, x\rangle, \langle s, x\rangle)$. Retain $T(X)$'s upper Pareto vertex chain, ordered by decreasing first and increasing second coordinate. Its preimages in X form the response chain. For each nonzero field in D, a chain vertex maximizes this score; a zero field permits any chain vertex. Conditional fields lie in D, so successive optimal replacements preserve an optimum on the chains.

Points $x_i^0, \ldots, x_i^{m_i}$ in the projected feasible set $X \subset S$ have responses $T(x_i^k)$ ordered along the vertex chain of $T(X)$ and increments $d_i^k = x_i^k - x_i^{k-1} \in S$. For increments d, d', write $T(d) = (-a, b)$ and $T(d') = (-a', b')$, with $a, b, a', b' \geq 0$. With $A = \|r\|^2$, $B = \|s\|^2$, $c = \langle r, s\rangle \geq 0$, the inverse Gram matrix gives their inner product in S:

$$\langle d, d'\rangle = \frac{Baa' + c(ab' + ba') + Abb'}{AB - c^2} \geq 0. \tag{14}$$

Prefix indicators for the chain indices turn these nonnegative inner products into nonnegative pair coefficients, so the negative score is a cut objective under prefix constraints: a polynomial-size minimum-cut formulation for unrestricted g (Kolmogorov & Zabih, 2004) (Appendix A).

3.3 Certificate construction and exact first-order recovery

The two-moment construction in Appendix H.3 lifts a projected optimizer to rational original-domain witnesses. An LP backend provides an alternative through rational reconstruction. Algorithm 1 checks the lift constraints and cut–flow equality exactly before issuing a certificate. Every

Algorithm 1 Projected maximum certificate and threshold query

Require: Rational input of Section 2, $0 < \varepsilon \le \varepsilon_B(S)$, threshold θ
1: Construct Q_ε by (8)–(9).
2: Construct D, the response chains of $Q_\varepsilon \cap D$, and the cut network.
3: Compute an exact minimum cut and an equal-valued maximum flow.
4: Decode the projected optimizer s^* and compute $V_\varepsilon = J(s^*)$.
5: Lift each s_i^* by the exact two-moment construction.
6: Verify cut/flow equality and every lift certificate in exact arithmetic.
7: Set $B^{\mathrm{lift}} = (b_i)_i$, $\underline{U} = J(B^{\mathrm{lift}})$, $\overline{U} = V_\varepsilon + nW\varepsilon^2$.
8: Set $\mathcal{C} = (B^{\mathrm{lift}}, [\underline{U}, \overline{U}], \text{verified certificates})$.
9: **if** $\overline{U} \le \theta$, **return** $(\mathsf{VALIDATED}, \mathsf{Safe}, \mathcal{C})$.
10: **if** $\underline{U} > \theta$, **return** $(\mathsf{VALIDATED}, \mathsf{Conflict}, \mathcal{C})$.
11: **return** $(\mathsf{VALIDATED}, \mathsf{Unknown}, \mathcal{C})$.

issued certificate is therefore exactly verified, and Unknown denotes a verified interval that straddles the threshold. The threshold outcomes are shown in Figure 2 (Appendix A).

The enclosure concerns the maximum. A lift supplies an attainable score, while the projected optimum and residual bound supply an upper endpoint. Writing $b_i = s_i + e_i$, orthogonality and the native fields give:

$$J(B) = J(s) + \sum_{i<j} w_{ij}\langle e_i, e_j\rangle, \qquad |J(B) - J(PB)| \le nW\varepsilon^2. \tag{15}$$

The enclosure brackets U_ε, while the zero completion and continuity give the full feasible score range $[0, U_\varepsilon]$. A public constant shifts both intervals equally.

Rational reconstruction recovers the first coefficient. Affine vertex branches give finitely many quadratic objectives. Their coefficients have a height bound H_{coef} of polynomial bit length (denoted H in Appendix B). Thus $\mathrm{den}(\Lambda) \le H_{\mathrm{coef}}$ and $|\Gamma| \le H_{\mathrm{coef}}$. Feasibility and ordering stabilize below ε_F and ε_J. At the query radius in (29), $z = (V_\varepsilon - U_0)/\varepsilon$ satisfies $|z - \Lambda| < 1/(8H_{\mathrm{coef}}^2)$. Continued fractions recover the unique such rational without enumerating the branches. Nestedness gives $\Lambda \ge 0$; Γ need not be nonnegative. The roles of ε_B, the reconstruction radius, and the hardness radius ε_* are compared in Appendix H.4.

3.4 The retained first-order response changes the decision

Consider a four-node similarity score. Nodes 3 and 4 are fixed anchors $u = (1, 0, 1)$ and $v = (0, 1, 1)$, with $S = \mathrm{span}(u, v)$ and unit-weight edges $(1, 2), (1, 3), (2, 4)$. Updating nodes 1 and 2 gives $g = 2$, $w_{12} = 1$, $(h_1, h_2) = (u, v)$, and reference score $J(u, v) = 5$. Direct optimization gives $U_\varepsilon = 5 + 6\varepsilon$ for $0 \le \varepsilon \le 1/3$ (Appendix C). For the analytic threshold $t_\varepsilon = 5 + 3\varepsilon$ and $0 < \varepsilon \le \varepsilon_B = 1/64$, the nominal maximum and feasible lift satisfy:

$$U_0 = 5 < t_\varepsilon, \qquad J(B^{\mathrm{lift}}) - t_\varepsilon \ge 3\varepsilon - 6\varepsilon^2 > 0. \tag{16}$$

The lift therefore supplies a feasible threshold-violation witness.

Certification beyond the closed-form example. A prespecified grid varies geometry, graph density, field strength, $n \le 18$, and $g \le 32$, with $W = 1$ and degree-scaled native fields (Appendix H.5). Projection returns verified certificates on all 98 instances. Table 2 compares it with an original-coordinate McCormick LP (McCormick, 1976) that retains every tube constraint, and with the cube bound; McCormick coverage combines 56 prespecified runs with 42 completed afterwards under the same protocol. Exact LP-feasible witnesses establish that the projected upper bound is strictly below the McCormick relaxation optimum in every pair. Original-domain enumeration validates all compared intervals on twelve small cases. On these twelve cases, projection resolves 44 of 60 near-optimum queries; an additional separable upper-bound check raises this to 54.

Table 2: Paired threshold outcomes on the 98 grid instances. S/C/U denote Safe/Conflict/Unknown over the 392 fixed-tolerance and 294 first-order-scale queries; the last column reports the median interval width normalized by the input-dependent score scale defined in Appendix H.5.

Method	Fixed tolerance S/C/U	First-order scale S/C/U	Median width
Projection	186/206/0	135/159/0	8.2×10^{-7}
McCormick	55/68/269	20/25/249	1.4×10^{-2}
Cube	0/0/392	0/0/294	1.4×10^{-1}

A three-node graph with six-dimensional representations has a unique optimizer over the full original domain, with three distinct projections (Appendix H.7). Exact restrictions to two projected values certify the need for at least three values in all ten grid instances for which the computed projected optimizer used at least three labels; the residual bound transfers this conclusion to eight corresponding original-domain problems (Appendix H.6). We next characterize the complexity of exact residual optimization.

4 THE PERTURBATION BOUNDARY

The residual contribution encodes a hard combinatorial problem even when the nominal optimizer is unique and the first-order response is fixed.

4.1 A UNIQUE NOMINAL OPTIMUM LEAVES A GENUINE RESIDUAL FIBRE

To realize this obstruction within the original domain, we reduce unweighted three-terminal edge cut (Dahlhaus et al., 1994) using the fixed rational $R \in \mathbb{Q}^{16\times 2}$ in Appendix D. Its four consecutive blocks G_A, G_B, G_C, G_D each contain four coordinates and have row sum $(7/2, 0)$. Writing $R = (a, x)$ gives independent public cube columns $u = (a + x)/2$, $v = (a - x)/2$. For type $\tau \in \{0, B, C, D\}$, let A_τ contain the rows $I_\tau = (a_\tau, b_\tau)$ in Figure 3 of Appendix D, and define:

$$p^\tau = RA_\tau^{-1}\begin{pmatrix}1\\1\end{pmatrix}, \qquad q^\tau = P(e_{a_\tau} + e_{b_\tau}), \qquad h_i = Hq^{\tau(i)}. \tag{17}$$

Each p^τ has exactly the two specified active upper coordinates. Each q^τ is strictly positive and a positive combination of u, v, making every field native.

A source graph with N vertices and m edges gives $g = 2N$ independently variable columns, two per vertex. Nonterminals have type 0 and terminal t_X has type X. Pair each vertex's copies with weight L and join copy-zero columns by unit-weight source edges. Set

$$D_0 = 3 \cdot 10^{20}, \quad L = 4D_0^4(32m + 1), \quad d_{\max} = L + m, \quad H \geq 1 + 10910 d_{\max}, \tag{18}$$

with integer H, unchanged domains K_ε, and no self-interactions. Each neighbor $b \in K_0$ has expansion coefficients along Pe_{a_τ}, Pe_{b_τ} with absolute value at most the fixed $\kappa < 10910$ (Appendix D, (39)). Thus each nominal conditional field is:

$$Hq^\tau + \sum_j w_{ij} b_j = \mu_a Pe_{a_\tau} + \mu_b Pe_{b_\tau}, \qquad \mu_a, \mu_b \geq H - \kappa d_{\max} \geq 1. \tag{19}$$

Maximizing this conditional field forces both active coordinates to one. Invertibility of A_τ therefore gives the unique nominal optimizer $b_i = p_i := p^{\tau(i)}$, with $U_0 = \mathcal{B} + 2Hg$ and $\mathcal{B} = \sum_{i<j} w_{ij}\langle p_i, p_j\rangle$.

First-order equality produces the fibre. For a right-feasible path $p_i + \varepsilon d_i$, set $z_i = (I - P)d_i$. The original constraints imply $z_i \in S^\perp \cap [-1, 1]^{16}$ and $(d_i)_{a_\tau}, (d_i)_{b_\tau} \leq 0$, so each projected active coordinate satisfies $(Pd_i)_a = (d_i)_a - (z_i)_a \leq 1$. For $G_i = h_i + \sum_j w_{ij} p_j$, positivity of the multipliers gives $\langle G_i, d_i\rangle \leq \langle G_i, p_i\rangle$; equality holds exactly when both projected active coordinates equal one. At the active coordinates, $d_a \leq 0$, $z_a \geq -1$, and $d_a - z_a = 1$ force $(d_i)_{a_\tau} = (d_i)_{b_\tau} = 0$ and $(z_i)_{a_\tau} = (z_i)_{b_\tau} = -1$. Inverting the active-row matrix then gives:

$$Pd_i = p_i, \qquad d_i = p_i + z_i, \qquad z_i \in F_{\tau(i)} := \{z \in S^\perp \cap [-1, 1]^{16} : z_{a_\tau} = z_{b_\tau} = -1\}. \tag{20}$$

Conversely, for every such direction, the fixed slack bounds give $b_i(\varepsilon) = (1+\varepsilon)p_i + \varepsilon z_i \in K_\varepsilon$ for $0 \le \varepsilon \le 1/40000$, with orthogonal residual exactly εz_i. Thus the original box and residual budget generate the entire residual fibre.

4.2 Vertex exclusion and a graph-independent positive radius

With the unique nominal optimizer fixed, lexicographic selection of affine vertex branches restricts the optimal first-order directions to the fibres $F_{\tau(i)}$. Proposition 8 proves the matching reverse bound using feasible fibre paths. The resulting coefficient is $\Lambda = 2\mathcal{B} + 2Hg$; orthogonality yields the exact residual objective:

$$\Gamma = \mathcal{B} + \max_{z_i \in F_{\tau(i)}} \sum_{i<j} w_{ij}\langle z_i, z_j\rangle. \tag{21}$$

Enforcing discrete labels at a vertex optimum. Separate affinity gives a product-vertex optimizer of this continuous problem. Lemma 9 uses the fixed moment matrix to bound vertex denominators: a pair of residual-fibre vertices z, z' attains inner product 16 only at identical cube corners, and otherwise satisfies:

$$16 - \langle z, z'\rangle \ge \frac{1}{4D_0^4}, \qquad L(16 - \langle z, z'\rangle) \ge 32m + 1 > 32m. \tag{22}$$

The source-edge sum varies by at most $32m$, while all pairing edges can be saturated simultaneously. Thus maximizing vertex configurations use synchronized legal corners. Separate affinity transfers their value to the continuous problem, without requiring a gap at nonvertex points.

The three legal corner labels have Gram matrix $16I$ (Figure 3, Appendix D), rewarding equal labels. Removing edges whose endpoints have different labels separates the terminals. A terminal-separating cut admits terminal-consistent component labels. Thus its minimum size c_* gives:

$$U_0 = \mathcal{B} + 2Hg, \quad \Lambda = 2\mathcal{B} + 2Hg, \quad \Gamma = \mathcal{B} + 16LN + 16(m - c_*). \tag{23}$$

Replacing c_* by the cut threshold k gives gap 16, with polynomially bounded data after clearing fixed denominators.

Two replacements establish a common radius. The coefficient reduction must also exclude competing vertex branches at a fixed positive radius. Let Q be the least common multiple of the denominators of P, set $T = 16!Q^{16}$, and use $H = (64QT + 10911)(d_{\max} + 1)$ and ε_* from (53). Lemma 11 controls each right-feasible branch $a + \varepsilon d$ on $0 < \varepsilon \le \varepsilon_*$. Its denominator and coordinate bounds enable two strict improvements: first fix the nominal intercepts, then remove any deficient first-order direction.

Write $\deg_w(i) = \sum_j w_{ij}$. At a product-vertex optimum, replacing a column with an incorrect intercept by a feasible column on a fibre path at the current radius gains at least $H/(2QT) > 16\deg_w(i)$ in field contribution, exceeding all interaction loss. Hence each intercept is p_i. Replacing a deficient direction then gains at least $\varepsilon(\deg_w(i) + 1)/(QT)$ at first order, with second-order loss bounded as follows:

$$16\varepsilon^2 T(T+2)\deg_w(i) \le 32\varepsilon^2 T^2 \deg_w(i) \le \frac{\varepsilon \deg_w(i)}{2QT}. \tag{24}$$

Both improvements are strict, so for $0 < \varepsilon \le \varepsilon_*$ each column of an optimal vertex configuration is $b_i = (1+\varepsilon)p_i + \varepsilon z_i$ with $z_i \in F_{\tau(i)}$. The nominal optimizer gives the value at $\varepsilon = 0$; (23) gives the exact quadratic value on $[0, \varepsilon_*]$. Here Q, T, ε_* are graph-independent; field strength scales with the graph.

Fixed radius and approximation. For every fixed rational $0 < \varepsilon \le \varepsilon_*$, thresholding U_ε is strongly NP-complete, with gap $16\varepsilon^2$ and product-vertex NP witnesses. The polynomial upper bound on the objective $M_{\text{box}} = 16\sum_{i<j} w_{ij} + \sum_{i,\ell}(h_i)_\ell$ makes the relative approximation-error parameter $4\varepsilon^2/M_{\text{box}}$ inverse-polynomial. An FPTAS would then achieve absolute error at most $4\varepsilon^2$ and distinguish the cut instances. Thus these fixed-radius classes admit no FPTAS unless $\mathrm{P} = \mathrm{NP}$. For disclosure, the computable enclosure provides conservative guarantees without requiring the exact maximum.

5 APPLICATION: CERTIFIED DISCLOSURE

The maximum-value certificate also supports selecting columns to reveal. We call this *disclosure*: selected columns are fixed to reference values, and the remaining columns stay uncertain. A monotone upper bound on score deviation provides a common certificate across these choices.

A monotone certificate for joint uncertainty. Fix reference matrix A with columns in K_0, universe $\mathcal{U}$, and shared inputs S, P, ε with $0 \le \varepsilon \le \varepsilon_B(S)$. For $E \subseteq \mathcal{U}$, columns outside E retain their reference values. Let $M(E)$ denote the largest full score, $c(E)$ the public–public constant, and $r_A = \sum_{i<j} w_{ij}\langle a_i, a_j\rangle$. Zeroing unknown columns attains the minimum score $c(E)$, giving anchored deviation $f(E) = \max\{r_A - c(E), M(E) - r_A\}$. Let $V(E)$ be the maximum full score with columns in E ranging over Q_ε, those outside E fixed at reference values, and $c(E)$ included. Set $\delta_E = n\varepsilon^2 \sum_{i<j\in E} w_{ij}$. Define:

$$u(E) = \max\{r_A - c(E), V(E) + \delta_E - r_A\}, \qquad f(E) \le u(E) \le f(E) + 2\delta_E. \tag{25}$$

As E grows, projected completion sets are nested, $c(E)$ is nonincreasing, and δ_E is nondecreasing. Hence u is monotone. A set with $u(E) > \theta$ is called a *potential conflict*; a feasible completion exceeding the deviation tolerance certifies a true conflict.

Proposition 5 (Certified disclosure)**.** *For $\theta \ge 0$, let $\mathcal{H}_u$ be the inclusion-minimal sets with $u(E) > \theta$. A disclosure Q hitting all edges of $\mathcal{H}_u$ (for every $E \in \mathcal{H}_u$, $Q \cap E \neq \varnothing$) guarantees $f(\mathcal{U} \setminus Q) \le \theta$. Given E with $u(E) > \theta$, an inclusion-minimal potential conflict $C \subseteq E$ can be extracted using at most $|E| + 1$ certificate queries, including the initial query. Its proper subsets are truly safe; a feasible violating completion also certifies true inclusion minimality. Otherwise C remains unresolved.*

The lift and zero completion provide lower certificates. Together with $u(E)$, they decide every query whose true deviation is more than $2\delta_E$ from the threshold.

For nonnegative disclosure costs and $\Delta = 2\delta_{\mathcal{U}}$, we have:

$$\mathrm{OPT}_{\mathrm{true}}(\theta) \le \mathrm{OPT}_{\mathrm{cons}}(\theta) \le \mathrm{OPT}_{\mathrm{true}}(\theta - \Delta), \tag{26}$$

Here the two optima are minimum disclosure costs for fixed reference values, with infeasible problems assigned $+\infty$; true feasibility requires $f(\mathcal{U} \setminus Q) \le \theta$ and conservative feasibility requires $u(\mathcal{U} \setminus Q) \le \theta$ (Appendix E). Appendix I extends the enclosure to reference residuals under an explicit error budget and projected-cone condition.

6 RELATED WORK

Ordered interactions and projected geometry. Luo & Yang (2026) study cone-compatible Monge geometry for fixed-marginal optimal transport; Kolmogorov & Zabih (2004) characterize graph-representable binary pairwise energies. Darbon (2009) reduces ordered-label submodular energies to minimum cut via level sets. Our geometric contribution is the explicit projection and common cone that make these ordered-label tools applicable to the original residual tube.

Minimal sets and disclosure. Minimal sets over monotone predicates provide a standard extraction framework (Marques-Silva et al., 2013); conflict–hitting-set relations are classical in diagnosis (Reiter, 1987). Our deviation bound supplies a computable monotone predicate for this framework, with an explicit additive error inherited from the geometric certificate.

Value sensitivity. Cifuentes et al. (2026) study finite right-hand-side sensitivity of mixed-binary quadratic programs; Hladík (2024) proves hardness for worst-case first derivatives of degenerate linear programs. Our first-/second-order separation arises from residual choices within one continuous domain.

7 CONCLUSION

Within the computable small-radius regime, exact projected optimization and feasible lifting give a quadratic-width certificate for the maximum. The first-order response remains exactly recoverable, while second-order and fixed-radius optimization encode strongly NP-complete problems. The same upper-bound construction yields monotone certificates for fixed-reference disclosure.

REPRODUCIBILITY STATEMENT

Section 3 describes projected optimization, feasible lifting, and certificate validation. Appendices A–B provide the geometric proofs and bit bounds; Appendices D, E, and I give the hardness, disclosure, and residual-reference arguments. Appendices F and H specify the analytic constructions, experiment grid, thresholds, resource limits, and independent checks. Executable code, dependencies, exact inputs, and complete experimental records accompany the submission.

AI USE STATEMENT

The core research ideas originated with the authors. Generative AI assisted with literature discovery and analysis, development and checking of mathematical formulations and proofs, experiment design, construction of synthetic test instances, verification-code implementation, interpretation of exact computations, figure preparation, translation, and manuscript drafting. Mathematical additions were checked through separate proof reviews and exact rational tests. The authors are responsible for the scientific claims and final submission.

REFERENCES

Sidhika Balachandar, Nikhil Garg, and Emma Pierson. Domain constraints improve risk prediction when outcome data is missing. In *International Conference on Learning Representations*, volume 2024, pp. 13578–13594, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/3a91841d2bcc0b13a3d0d5d60c9f0581-Paper-Conference.pdf.

Diego Cifuentes, Santanu S. Dey, and Jingye Xu. Sensitivity analysis for mixed binary quadratic programming. *Mathematical Programming*, 216(1-2):379–424, March 2026. doi: 10.1007/s10107-025-02265-6. URL http://dx.doi.org/10.1007/s10107-025-02265-6.

E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, August 1994. doi: 10.1137/s0097539792225297. URL http://dx.doi.org/10.1137/S0097539792225297.

Jérôme Darbon. Global optimization for first order Markov Random Fields with submodular priors. *Discrete Applied Mathematics*, 157(16):3412–3423, 2009. doi: 10.1016/j.dam.2009.02.026. URL https://doi.org/10.1016/j.dam.2009.02.026.

Jicong Fan, Rui Chen, Zhao Zhang, and Chris Ding. Neuron-enhanced AutoEncoder matrix completion and collaborative filtering: Theory and practice. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/6cd3ac24cdb789beeaa9f7145670fcae-Abstract-Conference.html.

Dennis Frauen, Fergus Imrie, Alicia Curth, Valentyn Melnychuk, Stefan Feuerriegel, and Mihaela van der Schaar. A neural framework for generalized causal sensitivity analysis. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://openreview.net/forum?id=ikX6D1oM1c.

Milan Hladík. Linear programming sensitivity measured by the optimal value worst-case analysis. *Optimization Methods and Software*, 39(5):1168–1184, September 2024. doi: 10.1080/10556788.2024.2329590. URL http://dx.doi.org/10.1080/10556788.2024.2329590.

Milan Hladík, Michal Černý, and Miroslav Rada. A new polynomially solvable class of quadratic optimization problems with box constraints. *Optimization Letters*, 15(6):2331–2341, September 2021. doi: 10.1007/s11590-021-01711-6. URL http://dx.doi.org/10.1007/s11590-021-01711-6.

Akhil Jalan, Yassir Jedra, Arya Mazumdar, Soumendu Sundar Mukherjee, and Purnamrita Sarkar. Optimal transfer learning for missing not-at-random matrix completion. In *Proceedings of*

the 42nd International Conference on Machine Learning, volume 267 of *Proceedings of Machine Learning Research*, pp. 26797–26829. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/jalan25a.html.

V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(2):147–159, February 2004. doi: 10.1109/tpami.2004.1262177. URL http://dx.doi.org/10.1109/TPAMI.2004.1262177.

Filippo Lazzati and Alberto Maria Metelli. Robustness in the face of partial identifiability in reward learning. In *International Conference on Learning Representations*, volume 2026, pp. 107353–107390, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/af179b3066cae40d9e57241a75b9db2c-Paper-Conference.pdf.

Lei Luo and Jian Yang. From cone geometry to monge structure: Local defects and global regret in high-dimensional optimal transport, 2026. URL https://arxiv.org/abs/2606.04695v2. arXiv:2606.04695v2, version of 23 August 2026.

Joao Marques-Silva, Mikoláš Janota, and Anton Belov. Minimal sets over monotone predicates in boolean formulae. In *Computer Aided Verification*, volume 8044 of *Lecture Notes in Computer Science*, pp. 592–607. Springer, 2013. doi: 10.1007/978-3-642-39799-8_39. URL https://doi.org/10.1007/978-3-642-39799-8_39.

Garth P. McCormick. Computability of global solutions to factorable nonconvex programs: Part i — convex underestimating problems. *Mathematical Programming*, 10(1):147–175, December 1976. ISSN 1436-4646. doi: 10.1007/bf01580665. URL http://dx.doi.org/10.1007/BF01580665.

Panagiotis Patrinos and Haralambos Sarimveis. Convex parametric piecewise quadratic optimization: Theory and algorithms. *Automatica*, 47(8):1770–1777, August 2011. doi: 10.1016/j.automatica.2011.04.003. URL http://dx.doi.org/10.1016/j.automatica.2011.04.003.

Raymond Reiter. A theory of diagnosis from first principles. *Artificial Intelligence*, 32(1):57–95, 1987. doi: 10.1016/0004-3702(87)90062-2. URL https://doi.org/10.1016/0004-3702(87)90062-2.

Joseph K. Scott, Davide M. Raimondo, Giuseppe Roberto Marseglia, and Richard D. Braatz. Constrained zonotopes: A new tool for set-based estimation and fault detection. *Automatica*, 69:126–136, July 2016. doi: 10.1016/j.automatica.2016.02.036. URL http://dx.doi.org/10.1016/j.automatica.2016.02.036.

Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018. URL https://iclr.cc/virtual/2018/oral/378.

Mojtaba Soltanalian and Ahmad Mousavi. The rank-collapse principle for quadratic optimization, 2026. URL https://arxiv.org/abs/2608.07828v1. arXiv:2608.07828v1.

Johan Thapper and Stanislav Živný. The complexity of finite-valued CSPs. *Journal of the ACM*, 63(4):1–33, November 2016. doi: 10.1145/2974019. URL http://dx.doi.org/10.1145/2974019. Article 37.

Peihao Wang, Shenghao Yang, Yunyu Liu, Zhangyang Wang, and Pan Li. Equivariant hypergraph diffusion neural operators. In *International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/forum?id=RiTjKoscnNd.

Yanbang Wang and Jon Kleinberg. From graphs to hypergraphs: Hypergraph projection and its reconstruction. In *International Conference on Learning Representations*, volume 2024, pp. 52745–52775, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/file/e92b8441b702d142435f243806f9ec5a-Paper-Conference.pdf.

ORGANIZATION OF THE APPENDICES

Geometry and sensitivity

Appendix A. Projection and exact optimization. An explicit description of the projected domain, an optimum-preserving cone, the response-chain reduction to minimum cut, and the maximum-value enclosure.

Appendix B. Local coefficients and bit bounds. Stable vertex branches, coefficient-height bounds, and polynomial-time recovery of the first-order response.

Appendix C. Analytic mechanisms. The geometry and feasible fibres in Figure 1, a closed-form threshold example, and finite comparisons for the hardness construction.

Hardness and disclosure

Appendix D. Fixed-geometry hardness. The sixteen-dimensional construction, residual corner labels, the second-order reduction, and hardness at a uniform positive radius.

Appendix E. Disclosure certificates. The fixed-reference model, monotone upper bounds, minimal potential conflict sets, disclosure cost bounds, and exact end-to-end examples.

Appendix I. References with residual components. Linear fields and fixed references with residual components, their error budgets, monotonicity, and exact linear-programming cases.

Computational evidence

Appendix F. Analytic mechanisms and geometric scaling. First-order sensitivity, response-chain size as dimension and radius vary, attainment of the enclosure-width bound, and threshold reversal.

Appendix H. Same-domain comparisons and exact lifting. Geometric lifting, McCormick comparisons, a 98-instance certificate study, separable refinement, continuous label constraints, and a unique original-domain optimizer.

Related work

Appendix G. Supplementary comparisons. Connections to distributional robustness, hypergraph representations, fixed-rank optimization, and ordered-label graph cuts.

A PROJECTION, CONE REDUCTION, AND EXACT OPTIMIZATION

The projected certificate follows from an explicit feasible set, a common cone, ordered response chains, and feasible lifting. The argument below proves these steps and their polynomial bit bounds. All inputs are rational, and arithmetic complexity is measured in bit operations. The shared subspace has dimension two; the ambient dimension and number of unknown columns remain part of the input.

A.1 EXACT PROJECTION BY SUPPORT INEQUALITIES

Let R have rank two and put $G = R^\top R$. The orthogonal projection of b is Rx exactly when $R^\top b = Gx$. The residual constraint becomes $|b_j - (Rx)_j| \le \varepsilon$ in each coordinate. Intersecting these bounds with the box gives (7), whose coordinate intervals are nonempty precisely on the polygon:

$$\Omega_\varepsilon = \{x : -\varepsilon \le (Rx)_j \le 1 + \varepsilon \text{ for every } j\}.$$

This polygon is bounded because R is injective. For $\varepsilon > 0$ it has nonempty interior: a sufficiently small neighborhood of the origin satisfies its inequalities strictly.

Write $r_j = R_j^\top \in \mathbb{Q}^2$. At a fixed $x \in \Omega_\varepsilon$, the attainable moments form the zonotope:

$$Z(x) = \sum_j r_j[L_j(x), T_j(x)].$$

Let $\mathcal{N}$ contain both perpendicular orientations to every nonzero r_j. Then we have:

$$Gx \in Z(x) \quad \Longleftrightarrow \quad \nu^\top Gx \le F_\nu(x) := \sum_{\nu^\top r_j \ge 0} (\nu^\top r_j) T_j(x) + \sum_{\nu^\top r_j < 0} (\nu^\top r_j) L_j(x) \quad (\nu \in \mathcal{N}). \tag{27}$$

The listed normals describe every case. In full dimension, each edge is parallel to a nonzero generator, so the list contains every facet normal. For a segment, perpendicular normals fix its supporting line, and other listed normals bound its endpoints through exact support values. For a point, the rows of R span $\mathbb{R}^2$, so two independent normals and their opposites force that point. Retaining normals from zero-width intervals makes the description valid throughout.

Each T_j is a minimum of two affine functions, and each L_j a maximum. The function F_ν is therefore concave and piecewise affine: coefficients of T_j are nonnegative, and those of L_j are negative. The $O(n)$ lines $(Rx)_j = \varepsilon$ and $(Rx)_j = 1 - \varepsilon$, together with the boundary of Ω_ε, form $O(n^2)$ full-dimensional cells. On a cell A, denote the affine expression of F_ν by $F_{\nu,A}$. Each branch choice supplies a global affine upper bound, giving the following identity:

$$F_\nu(x) = \min_A F_{\nu,A}(x), \qquad x \in \Omega_\varepsilon.$$

Equality on cell boundaries follows by continuity from an adjacent cell. Consequently (27) is equivalent to all the linear inequalities:

$$\nu^\top Gx \le F_{\nu,A}(x), \qquad \nu \in \mathcal{N},\ A \text{ a cell.}$$

The $O(n^3)$ inequalities, intersected with Ω_ε, describe exactly the projected domain in basis coordinates. A bounded planar intersection of m halfplanes has at most m vertices. Rational line intersections, sign tests, and arrangement construction use polynomially many operations on polynomial-length rationals. All formulas contain polynomial-length sums and products of input rationals. This proves the polynomial bit-complexity bound in Lemma 4.

At $\varepsilon = 0$, use $K_0 = S \cap [0,1]^n$ directly, with basis-coordinate description $0 \le Rx \le 1$. Since C spans S, the section is two-dimensional.

A.2 AN EXPLICIT OPTIMUM-PRESERVING NONOBTUSE CONE

The pointed two-dimensional cone C has two distinct rational extreme rays u, v. Obtain them from the planar inequalities $Rx \ge 0$ and normalize each by its maximum coordinate. Along the boundary ray v, some coordinate vanishes but is positive on u; otherwise v would not bound the nonnegative section. This gives ℓ in (12), and exchanging the rays gives k.

Set $A = \|u\|^2$, $B_0 = \|v\|^2$, and $c = \langle u, v\rangle$. If $c > 0$, the parameter in (11) satisfies $0 < q \le 1/8$, by $2c \le A + B_0$. The new boundary rays have inner product:

$$\langle u - qv, v - qu\rangle = c - q(A + B_0) + q^2 c = c(3/4 + q^2) > 0.$$

The new rays are independent. Expressing u, v in this basis through the inverse matrix, whose denominator is $1 - q^2$, gives $C \subset D$.

For $s = Pb \in Q_\varepsilon$ and $c' \in C$, orthogonality gives $\langle s, c'\rangle = \langle b, c'\rangle \ge 0$, proving (10). Write $s = \alpha u + \beta v$. Since $s_j \ge -\varepsilon$, the two selected coordinates imply $\alpha, \beta \ge -M\varepsilon$. Membership in D is equivalent to the following conditions:

$$\alpha + q\beta \ge 0, \qquad \beta + q\alpha \ge 0.$$

If the first inequality is violated, $\alpha < qM\varepsilon$ and $\beta < M\varepsilon/q$; violating the second reverses these roles. Both coefficients are therefore strictly below $M\varepsilon/q$, which for $\varepsilon \le \varepsilon_B$ is at most $1/(2\gamma)$.

Let $c_0 = (u + v)/\gamma \in K_0$. Then $c_0 - s$ is a strictly positive combination of both extreme rays whenever $s \notin D$. Every nonzero $t \in C_S^*$ has nonnegative inner products with the rays and a positive inner product with at least one, since the rays span S. Hence we have:

$$\langle t, c_0 - s\rangle > 0 \qquad (t \in C_S^* \setminus \{0\},\ s \in Q_\varepsilon \setminus D).$$

At a projected global optimum, column i has conditional field $h_i + \sum_{j \ne i} w_{ij} s_j \in C_S^*$. A nonzero field rules out an outside column by strict improvement. With a zero field, replacement by c_0 preserves the objective. Replacing all outside columns proves (13). Each intermediate configuration remains globally optimal, and earlier replacements stay in D, so at most g replacements suffice.

If $c = 0$, the nonnegative rays have disjoint supports and

$$P = uu^\top/\|u\|^2 + vv^\top/\|v\|^2$$

is entrywise nonnegative. Hence $Pb \in C$ for every $b \ge 0$. Take $D = C$ at any radius, or $\varepsilon_B = 1$ for a uniform statement. In the other case, constructing the rational cone and radius takes polynomially many rational operations, with polynomial encoding length. The construction allows coordinates that vanish throughout S.

A.3 Response chains and the cut reduction

To isolate the geometric condition, allow different rational polygons $X_i \subset D$. Suppose $D = \mathrm{cone}(r, s)$ is pointed, two-dimensional, and nonobtuse, and $h_i \in D$. Use the invertible map:

$$T(x) = (\langle r, x\rangle, \langle s, x\rangle).$$

The upper Pareto boundary of $T(X_i)$ gives a finite vertex chain ordered by increasing second and decreasing first coordinate. On a horizontal or vertical extreme face, retain its Pareto-maximal endpoint. For every nonzero $ar + bs \in D$ with $a, b \ge 0$, at least one retained vertex maximizes the corresponding linear functional over X_i. For a segment or point, retain the appropriate endpoint or endpoints.

At a global optimum, the conditional field $h_i + \sum_{j \ne i} w_{ij} x_j$ lies in D. For a nonzero field, replace x_i by a maximizing chain vertex; for a zero field, replace x_i by any chain vertex. Each replacement preserves optimality. Since all other vectors remain in D, at most g replacements produce an optimum entirely on the chains.

Let the chain be $x_i^0, \dots, x_i^{m_i}$, with increments $d_i^k = x_i^k - x_i^{k-1}$. Its image increments have the form $(-a, b)$ with $a, b \ge 0$. Write:

$$G_D = \begin{pmatrix} A & c \\ c & B \end{pmatrix}, \qquad A, B > 0, \quad c \ge 0, \quad AB - c^2 > 0.$$

For increments with image coordinates $(-a, b)$ and $(-a', b')$, we have:

$$\langle d, d'\rangle = \frac{Baa' + c(ab' + ba') + Abb'}{AB - c^2} \ge 0. \tag{28}$$

The variables $y_i^k = \mathbf{1}\{\text{chosen index} \geq k\}$ satisfy $y_i^k \geq y_i^{k+1}$ and represent $x_i = x_i^0 + \sum_k d_i^k y_i^k$. Expanding the objective gives linear terms and pair coefficients:

$$\beta_{ik,j\ell} = w_{ij}\langle d_i^k, d_j^\ell \rangle \geq 0.$$

Extend the coefficients symmetrically across columns and set all same-column coefficients to zero. For binary y, z, the identity $yz = (y + z - |y - z|)/2$ expresses the objective as linear rewards minus nonnegative disagreement costs. Its negative is minimized by an s–t cut (Kolmogorov & Zabih, 2004). Encode a unary reward ay with a source arc of capacity a if $a \geq 0$, or a sink arc of capacity $-a$ otherwise, taking $y = 1$ on the source side. Disagreement edges have capacity $\beta/2$. Prefix constraints use directed arcs with capacity greater than the sum of all finite costs. The rational capacities have polynomial encoding length, and the network size is polynomial in the total chain size: $O(gn^3)$ binary nodes, at most $O(g^2n^6)$ cross-column terms, and unary and prefix arcs. Rational minimum cut has polynomial bit complexity, so the problem is exactly solvable for unrestricted g. For fixed g, product-vertex enumeration would suffice after projection; the cut formulation removes that restriction.

Apply this argument to $X_i = Q_\varepsilon \cap D$, using the previous two subsections. At zero radius use $X_i = K_0 \subset C$ and $D = C$. This proves exact polynomial-time computation of both V_ε and U_0 in their stated ranges.

A.4 Lifting and the maximum-value enclosure

Each projected column admits an independent lift through (7), found by rational LP in polynomial bit time. For every feasible completion, orthogonality eliminates field–residual and projected–residual cross terms, yielding (15). Moreover, we have:

$$\left| \sum_{i<j} w_{ij}\langle e_i, e_j \rangle \right| \leq \sum_{i<j} w_{ij} \sum_{k=1}^{n} |(e_i)_k (e_j)_k| \leq nW\varepsilon^2.$$

Projecting an optimizer of U_ε shows $U_\varepsilon \leq V_\varepsilon + nW\varepsilon^2$. Lifting an optimizer of V_ε shows $J(B^{\text{lift}}) \geq V_\varepsilon - nW\varepsilon^2$, while feasibility gives $J(B^{\text{lift}}) \leq U_\varepsilon$. These statements prove (4)–(5).

The enclosure brackets the *maximum*. Setting all unknown vectors to zero attains the minimum of the original nonnegative score. Continuity on the connected feasible domain then gives the full score interval $[0, U_\varepsilon]$. A public additive constant shifts this interval and every maximum-value certificate equally. A fixed off-subspace component of h_i contributes a field–residual term of order ε.

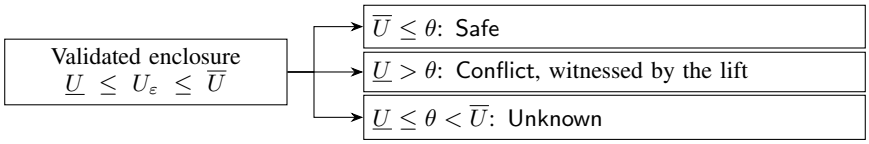

Figure 2: The enclosure brackets the maximum. An upper endpoint at or below the threshold proves safety; a feasible lower endpoint above it proves a violation. A straddling interval remains a valid certificate with an unresolved threshold.

B Exact local coefficients and bit bounds

Exact recovery of the first-order response requires a radius at which the maximizing vertex branches have stabilized. We bound that radius and the coefficient denominators, then apply rational reconstruction to the projected value.

Write $K_\varepsilon = \{x : Ax \leq b_0 + \varepsilon b_1\}$ and clear a common denominator so that A, b_0, b_1 are integer arrays. Let $K \geq 1$ bound their absolute entries; this scaling has polynomial bit length. For $\varepsilon > 0$, the tube has nonempty ambient interior. Scale a relative-interior point of the nonnegative section of S sufficiently far into the box, then assign sufficiently small positive values to its common zero coordinates. The resulting point satisfies all box and residual inequalities strictly, so every vertex admits n independent active rows.

For a nonsingular basis I, put $a_I = A_I^{-1}(b_0)_I$ and $d_I = A_I^{-1}(b_1)_I$, giving candidate vertex $a_I + \varepsilon d_I$. Cramer's rule bounds each vector's common denominator by $H_v = n!K^n$ and its absolute numerators by H_v. The basis is right-feasible at zero exactly when every slack $\sigma_0 + \varepsilon\sigma_1$ satisfies either $\sigma_0 > 0$ or both $\sigma_0 = 0$ and $\sigma_1 \geq 0$. A negative constant, or a zero constant with negative slope, rules out right-feasibility. A common height bound for the slack coefficients is:

$$H_F = (n+1)(KH_v)^{n+1} + 2$$

This bounds both numerator and denominator heights of the slack coefficients. For $0 < \varepsilon \leq \varepsilon_F := 1/(2H_F^2)$, the signs of all affine slacks are stable. The same finite family of right-feasible bases therefore supplies all vertices on this interval.

Because the score is affine in each column, successive conditionally maximizing vertex replacements turn any global optimizer into a product-vertex optimizer. Substituting the affine vertex branches yields finitely many polynomials $c_0 + c_1\varepsilon + c_2\varepsilon^2$. Let $H_c \geq 2$ bound the numerator and denominator heights of all weights and field entries, and set:

$$M = 2ng^2 + ng + 1, \qquad H = M(H_c H_v^2)^M + 2.$$

Each coefficient is a sum of fewer than M products with height at most $H_c H_v^2$, so its reduced numerator and denominator are bounded by H. Although the number of candidate products can be exponential, $\log H$ is polynomial in the input size.

Two distinct rationals with denominators at most H differ by at least H^{-2}, while linear and quadratic coefficient differences have magnitude at most $2H$. For $0 < \varepsilon \leq \varepsilon_J := 1/(8H^3)$, polynomial comparisons therefore follow the lexicographic order of (c_0, c_1, c_2), and one fixed polynomial gives the maximum. Compactness and nestedness imply $U_\varepsilon \to U_0$: as radii tend to zero, optimizing completions have a convergent subsequence with vanishing residual, and continuity gives the upper limiting bound; $K_0 \subset K_\varepsilon$ gives the lower bound. The maximal constant coefficient is thus U_0. This proves the local identity and polynomial encoding lengths for the coefficients and a computable interval endpoint. Nestedness also gives $\Lambda \geq 0$, without requiring $\Gamma \geq 0$ when $\Lambda > 0$.

To recover Λ, compute U_0 by the zero-radius chain oracle and choose:

$$\varepsilon = \min\left\{\frac{\varepsilon_B}{2}, \varepsilon_F, \varepsilon_J, \frac{1}{8H^2(H + nW + 1)}\right\}. \tag{29}$$

All these rational numbers have polynomial encoding length. Query the exact projected oracle and set $z = (V_\varepsilon - U_0)/\varepsilon$. The certificate and local identity imply:

$$|z - \Lambda| \leq (|\Gamma| + nW)\varepsilon < 1/(8H^2).$$

Since Λ has denominator at most H, it is the unique rational with that denominator bound this close to z. Continued-fraction reconstruction finds it in time polynomial in $\log H$ and the encoding length of z. The error bound is stronger than the convergent criterion $|z - p/q| < 1/(2q^2)$, giving exact recovery with polynomial bit complexity.

Membership of the coefficient decision problem in NP. A certificate gives one right-feasible nonsingular basis per column. The verifier checks the slacks, expands the score into (c_0, c_1, c_2), computes U_0 and Λ by the polynomial algorithms above, and checks $c_0 = U_0$, $c_1 = \Lambda$, and $c_2 \geq \tau$. Feasibility on an initial positive interval rules out $c_2 > \Gamma$. Conversely, an optimal product branch supplies a witness whenever $\Gamma \geq \tau$. The basis list and arithmetic have polynomial size. Together with Appendix D, this proves NP-completeness.

C Analytic mechanisms and finite constructions

Geometry and exact fibres in Figure 1. The paired points illustrate a generic feasible projection. The residual $e = (I - P)b$ satisfies $Pe = 0$ and $\|e\|_\infty \leq \varepsilon$; coordinatewise box and tube constraints give $[L_j, T_j] = [0, 1] \cap [s_j - \varepsilon, s_j + \varepsilon]$ for $s = Pb$. For $0 < \varepsilon \leq \varepsilon_B(S)$, $w_{ij} \geq 0$, and $h_i \in C$, the constructed D, c_0 permit an outside-point replacement with conditional gain $\Delta J \geq 0$, preserving the projected optimum. In response coordinates $T(x) = (\langle r, x\rangle, \langle s, x\rangle)$ for $D = \mathrm{cone}(r, s)$, successive increments $d^k = x^k - x^{k-1}$ satisfy $T(d^k) = (-a, b)$ with $a, b \geq 0$; the nonobtuse-cone Gram argument gives $\langle d^k, d^\ell\rangle \geq 0$.

The bottom diagrams show local feasible sections for two projected points at $S = \{z = x + y\}$ and $\varepsilon = 1/64$. Writing $b = s + t(-1, -1, 1)$, the boundary point $s = (63, 2, 65)/64$ admits only $t = -1/64$, whereas the interior point $s = (1, 1, 2)/4$ admits $-1/64 \le t \le 1/64$. The shaded sections use the in-plane direction $w = (-1, 0, -1)$ and normal $n = (-1, -1, 1)$ in $b = s+uw+tn$ and are cropped to local viewing windows. The dashed line fixes $u = 0$, blue points mark s, and amber marks its feasible intersection with K_ε. These intersections give a singleton and a segment, respectively.

Exact maximum of the four-node example. The threshold example in Section 3.4 compares the nominal maximum with the maximum over the residual tube. To compute the latter, we reduce the original-domain optimization to two line segments. For this three-dimensional representation geometry, the normal to S is $(-1, -1, 1)$, so the tube is the box intersected with $|b_3 - b_1 - b_2| \le 3\varepsilon$. Put $t = 3\varepsilon \in [0, 1]$. Any feasible (x, y, z) can be increased coordinatewise to $(x', y', 1)$ satisfying $x' + y' = 1 + t$: the original inequalities imply $x + y \le 1 + t$, and $1 + t \le 2$ ensures sufficient remaining box capacity in the first two coordinates. The new point is feasible. Coordinate monotonicity therefore restricts both columns to the segment with endpoints $A = (1, t, 1)$ and $B = (t, 1, 1)$. Separate affinity leaves four objective values:

$$J(A, B) = 5 + 2t, \quad J(A, A) = J(B, B) = 5 + t + t^2, \quad J(B, A) = 3 + 4t.$$

For $0 \le t \le 1$, the first value is maximal, proving $U_\varepsilon = 5 + 6\varepsilon$. Equation (12) gives $\varepsilon_B = 1/64$ for the stated rays. For $0 < \varepsilon \le 1/64$, the projected point-value error relative to the nominal error satisfies:

$$\frac{|U_\varepsilon - V_\varepsilon|}{U_\varepsilon - U_0} \le \frac{3\varepsilon^2}{6\varepsilon} = \frac{\varepsilon}{2}.$$

The enclosure-width ratio is bounded by ε, since the width bound is $6\varepsilon^2$. Both ratios are defined for $\varepsilon > 0$.

Finite checks of the hardness construction. The reduction in Appendix D identifies residual corner labels and relates their scores to terminal cuts. Exhaustive enumeration checks these identities on finite instances: the two moment constraints are evaluated at all $65{,}536$ cube corners, and terminal-consistent objectives are evaluated on all $1{,}096$ simple graphs on three through five vertices across $9{,}416$ labelings. Eighteen original-objective expansions are also evaluated. The original-domain comparison additionally uses the uniform-radius field and rational radius from Appendix D.5.

On these graph instances, enumeration compares the equality-label objective with the terminal-cut objective. Exhaustive edge-removal comparisons cover all graphs on three and four vertices and the complete graph on five vertices. The continuous-domain mechanism and its scaling behavior are analyzed in Appendix F.

D HARDNESS IN A FIXED SIXTEEN-DIMENSIONAL GEOMETRY

The hardness proof realizes terminal-cut labels as residual choices inside the original continuous tube. Native fields first fix the nominal optimizer; residual pairing then enforces the labels, and a uniform-radius argument extends the reduction to positive radii. The reduction fixes the subspace, projector, and four field directions independently of the source graph. Only the number of free columns, their field types, and polynomially bounded positive coefficients vary. Every unordered interaction is counted once, with no self-interactions.

Our source is unweighted three-terminal edge cut, which is NP-complete by Theorem 3 and Section 4.2 of Dahlhaus et al. (1994). The source graph has three specified terminals. In that proof, replacing each weight-four edge with four internally disjoint length-two unit-weight paths gives the unit-weight case. Thus the source graph can be simple and unweighted.

D.1 THE FIXED SUBSPACE AND FOUR EXPOSED VERTICES

The geometry needs distinct nominal types for nonterminals and for each terminal. The matrix below supplies those exposed points while its block structure determines which residual corners remain available.

Partition the sixteen coordinates into consecutive blocks $G_A = \{1,2,3,4\}$, $G_B = \{5,6,7,8\}$, $G_C = \{9,10,11,12\}$, and $G_D = \{13,14,15,16\}$. Define the rational matrix $R^{(0)}$ as follows:

j	$R^{(0)}_{j1}$	$R^{(0)}_{j2}$
1	$9991/10000$	$-3/100$
2	$9999/10000$	$-1/100$
3	$9991/10000$	$3/100$
4	$5019/10000$	$1/100$
5	$399/400$	$-1/20$
6	$1001/1200$	$3/5$
7	$1001/1200$	$-3/5$
8	$1001/1200$	$1/20$
9	$9999/10000$	$1/100$
10	$25001/30000$	$3/5$
11	$25001/30000$	$-3/5$
12	$25001/30000$	$-1/100$
13	$399/400$	$1/20$
14	$1001/1200$	$3/5$
15	$1001/1200$	$-3/5$
16	$1001/1200$	$-1/20$

(30)

For $t \in \{0,1,2,3\}$ and $s \in \{1,2,3\}$, set $\alpha_{t,s} = 10^{-(8+3t+s)}$ and perturb only the first column:

$$R_{4t+s,1} = R^{(0)}_{4t+s,1} + \alpha_{t,s}, \qquad R_{4t+4,1} = R^{(0)}_{4t+4,1} - \sum_{s=1}^{3} \alpha_{t,s}, \tag{31}$$

$$R_{j2} = R^{(0)}_{j2}.$$

Every block still has row sum $(7/2, 0)$. Write the two columns of R as a, x, and define:

$$S = \mathrm{col}(R), \qquad M = R^\top R, \qquad P = RM^{-1}R^\top, \qquad u = (a+x)/2, \quad v = (a-x)/2. \tag{32}$$

The public columns u, v lie strictly inside the cube because $R_{j1} > |R_{j2}|$ in every row and $R_{j1} + |R_{j2}| < 2$. They are independent: a has nonzero coordinate sum, whereas x is nonzero with zero coordinate sum. Hence $C = S \cap \mathbb{R}^{16}_+$ spans the fixed two-dimensional space S.

For $\tau \in \{0, B, C, D\}$, use the index pairs:

$$I_0 = (1,2), \qquad I_B = (5,1), \qquad I_C = (2,9), \qquad I_D = (3,13). \tag{33}$$

If $I_\tau = (i_\tau, j_\tau)$, let A_τ consist of those two rows of R, and put:

$$p^\tau = RA_\tau^{-1}\begin{pmatrix}1\\1\end{pmatrix}, \qquad q^\tau = P(e_{i_\tau} + e_{j_\tau}), \qquad K_0 = S \cap [0,1]^{16}. \tag{34}$$

Lemma 6 (Fixed geometric bounds)**.** *Each A_τ is invertible. The point p^τ has exactly its two specified coordinates equal to one, and*

$$\min_j p^\tau_j > \frac{2}{5}, \qquad \min_{j \notin I_\tau}(1 - p^\tau_j) > \frac{1}{10000}. \tag{35}$$

Writing $q^\tau = R(\eta_\tau, \xi_\tau)^\top$, we have $\eta_\tau > |\xi_\tau|$ and $q^\tau_j > 1/20$ for every coordinate. In particular, $q^\tau \in C$ and $\langle q^\tau, p^\tau \rangle = 2$. Moreover, for

$$\widehat{M}_\tau = A_\tau^{-\top} M A_\tau^{-1}, \qquad \kappa = \max_{\tau, r} \sum_s |(\widehat{M}_\tau)_{rs}|, \tag{36}$$

we have $\kappa < 10910$.

Proof. The following table provides the unperturbed solutions and strict bounds for the perturbed normal matrices:

τ	$(A^{(0)}_\tau)^{-1}(1,1)^\top$	$\|\widehat{M}_\tau\|_\infty$ lies in
0	$(10000, -400)/10003$	$(10875, 10876)$
B	$(2000, -160)/2003$	$(10906, 10907)$
C	$(10000, 0)/9999$	$(10861, 10862)$
D	$(2000, 160)/2003$	$(10909, 10910)$

(37)

Substitution in (30) gives unperturbed coordinates above $1/2$ and inactive upper slacks at least $8/10015 > 1/2000$. Each unperturbed inverse has infinity norm below 102, while the largest row perturbation in (31) is below $2\cdot 10^{-9}$. The inverse perturbation identity bounds the parameter change strictly below 10^{-6} and the resulting sixteen-dimensional point change strictly below 10^{-5}. This proves (35) with strict slack.

For the field claims, the row table gives:

$$12 < M_{11} < 13, \qquad M_{22} = 10861/5000 > 17/8, \qquad |M_{12}| < 1/100. \tag{38}$$

The unperturbed cross entry is $-1243/375000$, and its perturbation has magnitude below $2 \cdot 10^{-9}$. For each selected row pair, the first component of the row sum lies in $(199/100, 2)$, and the second has absolute value at most $2/25$. Applying the two-by-two inverse formula to $M^{-1}(R_{i_\tau} + R_{j_\tau})$ gives:

$$\eta_\tau > 3/20, \qquad |\xi_\tau| < 10600/254999 < 1/24.$$

As $R_{j1} > 1/2$ and $|R_{j2}| \le 3/5$, this implies $q_j^\tau > (3/20)(1/2) - (1/24)(3/5) = 1/20$. Also $q^\tau = (\eta_\tau + \xi_\tau)u + (\eta_\tau - \xi_\tau)v$, with both coefficients positive. Finally, $p^\tau \in S$ gives $\langle q^\tau, p^\tau \rangle = p^\tau_{i_\tau} + p^\tau_{j_\tau} = 2$.

The final column of (37) consists of exact rational inequalities. For the rows of A_τ given by $(\alpha_1, \beta_1), (\alpha_2, \beta_2)$, write $\Delta = \alpha_1\beta_2 - \alpha_2\beta_1$. The normal-matrix entries satisfy:

$$\begin{aligned}
\Delta^2(\widehat{M}_\tau)_{11} &= M_{11}\beta_2^2 - 2M_{12}\alpha_2\beta_2 + M_{22}\alpha_2^2, \\
\Delta^2(\widehat{M}_\tau)_{22} &= M_{11}\beta_1^2 - 2M_{12}\alpha_1\beta_1 + M_{22}\alpha_1^2, \\
\Delta^2(\widehat{M}_\tau)_{12} &= -M_{11}\beta_1\beta_2 + M_{12}(\alpha_1\beta_2 + \alpha_2\beta_1) - M_{22}\alpha_1\alpha_2.
\end{aligned}$$

Substitution and clearing positive denominators give the table; its largest row sum, for type D, is given exactly by:

$$\kappa = \frac{6565129627480406321409371014562512859296150 7}{60180135003004499819767499995500000013500} < 10910.$$

□

The matrix $\widehat{M}_\tau$ bounds the normal coordinates of every neighbor $b \in K_0$:

$$b = \lambda_1 P e_{i_\tau} + \lambda_2 P e_{j_\tau}, \qquad \begin{pmatrix}\lambda_1 \\ \lambda_2\end{pmatrix} = \widehat{M}_\tau \begin{pmatrix} b_{i_\tau} \\ b_{j_\tau}\end{pmatrix}, \qquad |\lambda_1|, |\lambda_2| \le \kappa. \tag{39}$$

The identity follows by writing $b = R\theta$ and solving for the coefficients in the independent basis Pe_{i_τ}, Pe_{j_τ}.

D.2 Residual corners and native-field pinning

The residual choices will encode terminal-cut labels after native fields fix each column's nominal type. The relevant residual sets are:

$$F = S^\perp \cap [-1, 1]^{16}, \qquad F_\tau = \{z \in F : z_{i_\tau} = z_{j_\tau} = -1\}. \tag{40}$$

These sets arise from first-order equality in the original constraints.

Lemma 7 (The six corners)**.** *The cube corners in F are exactly the vectors that are constant on each block, negative on two blocks, and positive on the other two. Let $z^{AX} = \mathbf{1} - 2\mathbf{1}_{G_A \cup G_X}$ for $X \in \{B, C, D\}$. Then we have:*

$$\begin{aligned}
F_0 \cap \{-1, 1\}^{16} &= \{z^{AB}, z^{AC}, z^{AD}\}, \\
F_X \cap \{-1, 1\}^{16} &= \{z^{AX}\} \quad (X = B, C, D),
\end{aligned} \qquad \langle z^{AX}, z^{AY} \rangle = 16\mathbf{1}\{X = Y\}. \tag{41}$$

Proof. For a sign vector z satisfying $R^\top z = 0$, the first moment equation is the sum of an element of $(1/30000)\mathbb{Z}$ and

$$\sum_{t=0}^{3} \sum_{s=1}^{3} 10^{-(8+3t+s)}(z_{4t+s} - z_{4t+4}).$$

The perturbation sum has absolute value below $2\sum_{j=1}^{12} 10^{-(8+j)} < 1/30000$, so it and the unperturbed moment both vanish. Its coefficients belong to $\{-2, 0, 2\}$. A first nonzero term would dominate all later terms, so every coefficient vanishes and each block is constant. The preserved block sums $(7/2, 0)$ then force exactly two negative blocks. Conversely, these six vectors satisfy both moments. The active coordinates select the lists in (41), and counting agreeing and disagreeing blocks gives the Gram matrix. □

nominal	active upper pair	fibre	cube-corner labels
p^0	$(1, 2)$	F_0	$1, 2, 3$
p^B	$(5, 1)$	F_B	1
p^C	$(2, 9)$	F_C	2
p^D	$(3, 13)$	F_D	3

	G_A	G_B	G_C	G_D
z^1	$-$	$-$	$+$	$+$
z^2	$-$	$+$	$-$	$+$
z^3	$-$	$+$	$+$	$-$

$\langle z^r, z^s \rangle = 16\mathbb{1}\{r = s\}$
Each sign repeats on four coordinates.

Figure 3: Four nominal types produce three orthogonal residual corner labels: $z^1 = z^{AB}$, $z^2 = z^{AC}$, $z^3 = z^{AD}$. The active coordinates select the cube corners of each residual fibre, not its entire point set.

Let the source graph have N vertices, m edges, and terminals t_B, t_C, t_D. Replace each vertex v by two independent free columns $b_{v,0}, b_{v,1}$, giving $g = 2N$ columns. Nonterminal copies have type 0; the copies of terminal t_X have type X. Set

$$D_0 = 3 \cdot 10^{20}, \qquad L = 4D_0^4(32m + 1), \qquad d_{\max} = L + m, \tag{42}$$

give each pair $(v, 0), (v, 1)$ weight L, and give $(u, 0), (v, 0)$ weight one for each source edge. All other free-free weights are zero. Choose any integer

$$H \geq 1 + 10910 d_{\max}, \qquad h_i = Hq^{\tau(i)}. \tag{43}$$

Each field is a positive combination of the two fixed public columns in (32), so it is native. In a public-free edge representation, the additional weights are fixed positive rational multiples of H. The restriction to $\{1, L\}$ applies only to free-free weights.

For any choice of the other columns in K_0, a type-τ column has conditional field:

$$Hq^\tau + \sum_j w_{ij} b_j = \mu_1 P e_{i_\tau} + \mu_2 P e_{j_\tau}, \qquad \mu_1, \mu_2 \geq H - \kappa d_{\max} \geq 1. \tag{44}$$

The conditional objective is uniquely maximized by setting both active coordinates to one, identifying p^τ. Hence the full nominal optimizer is unique: $b_i = p_i := p^{\tau(i)}$. With

$$\mathcal{B} = \sum_{i<j} w_{ij} \langle p_i, p_j \rangle, \tag{45}$$

the nominal value is $U_0 = \mathcal{B} + 2Hg$.

D.3 The complete first-order fibre

Fixing a unique nominal optimizer still permits several feasible first-order directions. We identify all directions that attain the optimal first-order response, since their residual interactions determine the second-order coefficient.

At the nominal point, put $G_i = h_i + \sum_j w_{ij} p_j$ and consider a right-feasible affine path $p_i + \varepsilon d_i$ in the original domain:

$$K_\varepsilon = \{b \in [0, 1]^{16} : \|(I - P)b\|_\infty \leq \varepsilon\}.$$

Writing $z_i = (I - P)d_i$, feasibility gives $z_i \in F$ and $(d_i)_{i_\tau}, (d_i)_{j_\tau} \leq 0$. Hence both corresponding coordinates of Pd_i are at most one. By the positive multipliers in (44), we have:

$$\langle G_i, d_i \rangle \leq \langle G_i, p_i \rangle. \tag{46}$$

Equality holds precisely when both projected active coordinates equal one. Since $d_a \le 0$ and $z_a \ge -1$, the equality $d_a - z_a = 1$ forces $d_a = 0, z_a = -1$. The active-row matrix is invertible, so equality in (46) is equivalent to the following conditions:

$$Pd_i = p_i, \qquad d_i = p_i + z_i, \qquad z_i \in F_{\tau(i)}. \tag{47}$$

Conversely, every such direction yields a feasible path:

$$b_i(\varepsilon) = (1+\varepsilon)p_i + \varepsilon z_i \quad \text{for } 0 \le \varepsilon \le 1/40000. \tag{48}$$

The active coordinates stay exactly one; the other coordinates remain in the cube by (35) and $\|p_i + z_i\|_\infty \le 2$. The orthogonal residual is exactly εz_i.

Proposition 8 (Exact second-order reduction)**.** *For the instance above, we have:*

$$\Lambda = 2\mathcal{B} + 2Hg, \qquad \Gamma = \mathcal{B} + \max_{z_i \in F_{\tau(i)}} \sum_{i<j} w_{ij}\langle z_i, z_j\rangle. \tag{49}$$

Proof. The fixed rational constraint matrix and affine right-hand side give a fixed finite family of affine vertex branches on an initial positive-radius interval. Columnwise affinity gives an optimum at a product of vertices. The resulting finite quadratic envelope selects coefficients lexicographically. Since the nominal optimum is unique, every winning intercept equals p_i.

The bounds (46) are attained simultaneously by (48), and their sum is $\sum_i \langle G_i, p_i\rangle = 2\mathcal{B} + 2Hg$. Nonnegativity of each deficit means that attaining the optimal total first-order coefficient requires all equalities in (47).

For a winning vertex branch, orthogonality of S and $S^\perp$ gives quadratic coefficient $\mathcal{B} + \sum_{i<j} w_{ij}\langle z_i, z_j\rangle$, establishing the upper bound in (49). Conversely, choose a maximizer of the compact residual product and follow the feasible paths (48). Their objective has the optimal constant and linear coefficients and the displayed quadratic coefficient. Compare it with U_ε, divide the remaining difference by ε^2, and take the right limit to obtain the reverse bound. □

D.4 A FIXED DENOMINATOR GAP ENFORCES THE THREE LABELS

The residual faces contain continuous choices as well as cube corners. To force the desired corner labels at a vertex optimum, we bound the loss on a pairing edge whenever its endpoints are not identical corners.

Lemma 9 (Residual vertex gap)**.** *If z, z' are vertices of the type residual faces and are not the same cube corner, then we have:*

$$16 - \langle z, z'\rangle \ge \frac{1}{4D_0^4}. \tag{50}$$

Proof. The integer matrix $D_0 R$ has entries of absolute value at most D_0. Fixing the two prescribed residual-face coordinates leaves two linear moment equations. At a vertex, at most two other coordinates lie strictly between -1 and 1; more would admit a nonzero two-sided perturbation preserving both equations. With two interior coordinates, their coefficient matrix must be invertible for the same reason, and its determinant has magnitude at most $2D_0^2$. With one, a denominator at most D_0 suffices. All vertex coordinates therefore share a denominator at most $2D_0^2$.

The rational deficit $16 - \langle z, z'\rangle$ therefore has denominator at most $4D_0^4$. Cauchy–Schwarz bounds the inner product by sixteen; equality requires both norms to equal four and the vectors to coincide. The deficit is thus positive unless the points are the same cube corner. □

By separate affinity, the residual objective in (49) has a vertex optimizer. A pairing edge that does not join identical legal corners loses at least $L/(4D_0^4) = 32m + 1$ from its maximum $16L$. The source-edge contribution ranges from $-16m$ to $16m$, so it cannot offset this loss. Every type has a legal corner, making simultaneous saturation of all pairing edges feasible. Thus every maximizing vertex configuration synchronizes the copies at a legal corner. Separate affinity transfers this vertex optimum to the full continuous residual product.

The remaining choices are precisely the three labels for nonterminals and the three distinct pinned labels for the terminals. Their Gram matrix is (41). If c_* is the minimum three-terminal cut, the residual maximum is therefore:

$$16LN + 16(m - c_*). \tag{51}$$

Unequal-label edges separate the terminals. Conversely, label each component of a terminal-separating cut consistently with its terminal, assigning labels arbitrarily to components without terminals. Thus the minimum unequal-label edge count equals the minimum cut cardinality.

Theorem 10 (Fixed-geometry second-order hardness)**.** *The construction satisfies:*

$$U_0 = \mathcal{B} + 2Hg, \qquad \Lambda = 2\mathcal{B} + 2Hg, \qquad \Gamma = \mathcal{B} + 16LN + 16(m - c_*). \tag{52}$$

Consequently, thresholding Γ is strongly NP-hard with fixed $n = 16$ and fixed rational S.

Proof. To prove the final claim, for a source threshold k set $T_\Gamma = \mathcal{B} + 16LN + 16(m - k)$. Then $\Gamma \geq T_\Gamma$ holds exactly when $c_* \leq k$, with a gap of at least sixteen for a NO instance. The constants $R, P, p^\tau, q^\tau, D_0$ are fixed. Both L and the choice $H = 1+10910(L+m)$ are polynomially bounded integers. Fields and $\mathcal{B}$ have fixed common denominators and polynomially bounded numerators. Clearing these denominators gives polynomially bounded integer data, proving a strong reduction from the unweighted source problem with fixed dimension and fixed denominators. □

D.5 A UNIFORM POSITIVE RADIUS

Increase the fixed multiplier in H so that the reduction is exact on a graph-independent positive-radius interval. The second-order formula stays unchanged; the constant and linear coefficients use the new field strength.

Let the scalar Q be the least common multiple of the coordinate denominators of P. Clear denominators in the tube system:

$$A = Q[I; -I; I - P; -(I - P)], \quad b^0 = Q[\mathbf{1}; 0; 0; 0], \quad b^1 = Q[0; 0; \mathbf{1}; \mathbf{1}].$$

Every entry has absolute value at most Q. Put:

$$\begin{aligned} T &= 16!Q^{16}, \qquad H = (64QT + 10911)(d_{\max} + 1), \\ \varepsilon_* &= \min\left\{\frac{1}{40000}, \frac{1}{34QT^2}, \frac{1}{68QT^2}, \frac{1}{64QT^3}\right\}. \end{aligned} \tag{53}$$

The constants Q, T, ε_* depend only on the fixed geometry; H scales with the weighted degree bound.

Lemma 11 (Uniform vertex control)**.** *For $0 < \varepsilon \leq \varepsilon_*$, every vertex of K_ε is given by a right-feasible affine branch $a+\varepsilon d$, with a common denominator at most T for a, d and $\|a\|_\infty, \|d\|_\infty \leq T$. For each type, if $c = 2 - \langle q^\tau, a\rangle > 0$, then $c \geq 1/(QT)$. If $a = p^\tau$, the first-order deficiency $l = 2 - \langle q^\tau, d\rangle$ is nonnegative, and if positive is at least $1/(QT)$.*

Proof. Cramer's rule bounds basis determinants and numerators by T. Every nonzero constant slack has magnitude at least $1/T$, while its slope has magnitude at most $17QT$. Thus radius $1/(34QT^2)$ fixes the feasibility status of all bases. Feasible branch limits belong to K_0. Since q^τ has denominator dividing Q, the stated positive gaps have denominator at most QT. For $a = p^\tau$, the active-coordinate argument of (46) gives $l \geq 0$, with equality exactly for the fibre directions. Finally, $|l| \leq 2 + 32T \leq 34T$. □

Proposition 12 (Exact formula on the common radius interval)**.** *With* (53)*, for every source graph and every $0 \leq \varepsilon \leq \varepsilon_*$, we have:*

$$U_\varepsilon = \mathcal{B} + 2Hg + (2\mathcal{B} + 2Hg)\varepsilon + [\mathcal{B} + 16LN + 16(m - c_*)]\varepsilon^2. \tag{54}$$

Proof. Let $\deg_w(i)$ denote the total free-edge weight incident to column i. Fix positive $\varepsilon \leq \varepsilon_*$ and a vertex optimizer. If a branch intercept differs from its prescribed p_i, replace the column by a feasible point at the current radius on a fibre path in (48). Its linear score in direction $q^{\tau(i)}$ increases by $c + \varepsilon l \geq 1/(2QT)$. This gives field gain at least $H/(2QT)$, while the incident free-edge terms can lose at most $16\deg_w(i)$: old and new columns and all neighbors lie in the cube. The choice of H makes the net gain strict, forcing every intercept to equal its p_i.

Each neighbor now has form $p_j + \varepsilon d_j$, with $\|d_j\|_\infty \le T$. For a column with positive deficiency l_i, replace its direction by the fibre direction $d_i^* = p_i + z_i$. Both nominal normal coefficients are at least $H - \kappa \deg_w(i) \ge \deg_w(i) + 1$, giving first-order gain at least $\varepsilon(\deg_w(i) + 1)/(QT)$. The second-order loss is bounded as follows:

$$\varepsilon^2 \, 16T(T+2) \deg_w(i) \le 32\varepsilon^2 T^2 \deg_w(i) \le \frac{\varepsilon \deg_w(i)}{2QT}.$$

The improvement is again strict, so every optimal vertex configuration lies in the first-order-optimal fibres. Applying the pairing argument to the exact objective expansion proves (54). At zero radius, the identity is the nominal optimum already established. □

Corollary 13 (Fixed-radius hardness)**.** *For every fixed positive rational $\varepsilon \le \varepsilon_*$, thresholding the original maximum U_ε is strongly NP-complete for the fixed geometry above. This class admits no maximum-value FPTAS unless* $\mathrm{P} = \mathrm{NP}$.

Proof. Substitute k for c_* in (54). The YES/NO gap is $16\varepsilon^2$. With dimension, projector, and radius fixed, the polytope has a fixed finite rational vertex set, so columnwise vertex assignments give polynomial-size NP certificates. All varying numerical data remain polynomially bounded after clearing fixed denominators.

The cube bound $M_{\mathrm{box}} = 16 \sum_{i<j} w_{ij} + \sum_{i,\ell} (h_i)_\ell$ is polynomially bounded in the source size. An FPTAS at relative tolerance $4\varepsilon^2 / M_{\mathrm{box}}$ would have absolute error at most $4\varepsilon^2$, enough to distinguish the cut instances in polynomial time. □

E Certified disclosure and potential conflict hypergraphs

Disclosure fixes selected representations to specified reference values and leaves the others uncertain. The maximum-value certificate bounds the resulting score deviation. Its monotonicity makes it possible to identify minimal potential conflicts and certify which disclosures remove them.

E.1 Fixed-reference completion model

Fix a rational two-dimensional space $S \subset \mathbb{R}^n$, its orthogonal projector P, and $\varepsilon \in (0, \varepsilon_B(S)]$. Write:

$$K_\varepsilon = \{b \in [0,1]^n : \|(I - P)b\|_\infty \le \varepsilon\}, \qquad Q_\varepsilon = PK_\varepsilon.$$

Assume $C = S \cap \mathbb{R}^n_+$ is full dimensional relative to S. Let $A_j \in S \cap [0,1]^n$ be fixed reference columns and let the rational unordered-pair weights w_{ij} be nonnegative, with no self terms. For a candidate unknown-column universe $\mathcal{U}$ and $E \subseteq \mathcal{U}$, set:

$$\mathcal{F}(E) = \{B : B_j = A_j \ (j \notin E),\ b_i \in K_\varepsilon \ (i \in E)\}, \qquad R(B) = \sum_{i<j} w_{ij} \langle b_i, b_j \rangle.$$

As E varies, keep the space, radius, reference, weights, and threshold fixed. Disclosing $Q \subseteq \mathcal{U}$ fixes those columns to A_Q.

To measure deviation from the reference, separate the fixed score contribution from the maximum over unknown columns:

$$r_A = R(A), \qquad c(E) = \sum_{i<j,\ i,j \notin E} w_{ij} \langle A_i, A_j \rangle,$$

$$D(E) = r_A - c(E), \qquad M(E) = \max_{B \in \mathcal{F}(E)} R(B),$$

$$f(E) = \max\{D(E), M(E) - r_A\}.$$

Zeroing all unknown columns is feasible and minimizes the score, so $c(E) = \min_{\mathcal{F}(E)} R$. Hence $f(E)$ is the anchored deviation $\max_{B \in \mathcal{F}(E)} |R(B) - R(A)|$. For $\theta \ge 0$, a true conflict has $f(E) > \theta$; equality is safe.

E.2 A MONOTONE UPPER BOUND

Let $V(E)$ be the maximum of the complete projected score: unknown columns range over Q_ε, public columns remain A_j, and the public constant $c(E)$ is included. The native field for an unknown column is $h_i(E) = \sum_{j \notin E} w_{ij} A_j \in C$. Set:

$$W(E) = \sum_{i<j,\ i,j \in E} w_{ij}, \qquad \delta_E = nW(E)\varepsilon^2.$$

The small-tube certificate computes $V(E)$ and a feasible original completion $B_{\text{lift}}(E)$ in polynomial time. Writing $L(E) = R(B_{\text{lift}}(E))$, it gives:

$$L(E) \le M(E) \le V(E) + \delta_E, \qquad V(E) + \delta_E - L(E) \le 2\delta_E.$$

Proposition 14 (Uniform conservative predicate)**.** *Define:*

$$\ell(E) = \max\{D(E), L(E) - r_A\}, \qquad u(E) = \max\{D(E), V(E) + \delta_E - r_A\}.$$

Then $\ell(E) \le f(E) \le u(E)$ *and*

$$u(E) - \ell(E) \le 2\delta_E, \qquad f(E) \le u(E) \le f(E) + 2\delta_E. \tag{55}$$

Moreover, u *is monotone under inclusion and* $u(\varnothing) = 0$*. The same geometric radius applies simultaneously to every* $E \subseteq \mathcal{U}$*.*

Proof. For fixed D, the map $x \mapsto \max\{D, x - r_A\}$ is nondecreasing and 1-Lipschitz. Applied to the maximum-value certificate, it gives the first inequality in (55); $\ell \le f$ gives the second. For monotonicity, $A_j \in K_0 \subseteq Q_\varepsilon$ makes projected completion domains nested as unknown columns are added. Thus $V(E)$ is nondecreasing, including when projected coordinates are negative. Both $D(E)$ and $W(E)$ are nondecreasing: removing public columns removes nonnegative public-public terms, while adding unknown columns adds nonnegative free-free weights. Each branch defining u is therefore nondecreasing. With no unknown columns, $V = r_A$ and $D = \delta_E = 0$. All queries have native fields in the same cone C, so one geometric radius applies to every subset. □

The nesting argument establishes monotonicity for exact projected values uniformly over all subsets, including adaptively chosen queries. Put $\bar{\delta} = \delta_{\mathcal{U}}$ and $\Delta = 2\bar{\delta}$. When $W(E) = 0$, the certificate is exact, including for empty and singleton unknown sets.

A three-way threshold oracle returns SAFE if $u(E) \le \theta$, CONFLICT if $\ell(E) > \theta$, and UNKNOWN otherwise. A conflict has an original feasible witness: use the zero completion if $D(E) > \theta$, and otherwise use $B_{\text{lift}}(E)$. Equation (55) implies:

$$f(E) \le \theta - 2\delta_E \implies \text{SAFE}, \qquad f(E) > \theta + 2\delta_E \implies \text{CONFLICT}.$$

The two-sided margin $|f(E) - \theta| > \Delta$ therefore suffices for a definite answer, though an exact branch can decide a query without it.

Validity of a disclosure certificate. A certificate is tied to its exact reference, weights, projector, radius, and unknown set. The verifier checks agreement with public columns, feasibility under the original cube and orthogonal-residual constraints, and the reported rational score. Zeroing unknown columns supplies the downward witness. An upper answer additionally requires the projected optimum and its error budget, checked by rerunning the polynomial projected oracle or verifying its exact optimization certificates.

Diameter queries. For score-interval diameter, use $d(E) = M(E) - c(E)$, bounded by $L(E) - c(E)$ and $V(E) + \delta_E - c(E)$. The gap is at most $2\delta_E$, and the upper bound is monotone because V is nondecreasing and c nonincreasing. The arguments below therefore also apply to diameter thresholds. A current-state diameter query uses the public values; comparing future disclosures additionally requires their specified reference values. For absolute maximum-score thresholds, feasibility of the fully revealed reference is necessary. The anchored formulation with $\theta \ge 0$ avoids this empty-conflict obstruction.

E.3 INCLUSION-MINIMAL POTENTIAL CONFLICTS

A potential conflict identifies a set of uncertain columns whose certified upper deviation exceeds the tolerance. Removing redundant columns gives an inclusion-minimal set. For a monotone predicate value q, define:

$$\mathcal{H}_q = \{C \subseteq \mathcal{U} : q(C) > \theta,\ q(C \setminus \{e\}) \leq \theta \text{ for every } e \in C\}.$$

This is minimal-set extraction over a monotone predicate (Marques-Silva et al., 2013). By monotonicity, the single-deletion tests characterize inclusion minimality; minimum cardinality and cost are separate optimization objectives. The families $\mathcal{H}_f$ and $\mathcal{H}_u$ are the true and potential conflict hypergraphs. Every true conflict is potential, but this inclusion holds for the upward-closed families, not necessarily for their minimal edge lists.

Proposition 15 (Certified deletion-based extraction)**.** *Given E with $u(E) > \theta$, an element-deletion pass extracts a member $C \in \mathcal{H}_u$ using at most $|E|$ additional certificate queries. Every proper subset of C is truly safe. If $\ell(C) > \theta$, the output is a true inclusion-minimal conflict with a feasible violating completion. Otherwise it is an unresolved potential conflict, with uncertainty interval $[\ell(C), u(C)]$ of length at most $2\delta_C$.*

Proof. Start with $C = E$ and test each original element e once by querying $T = C \setminus \{e\}$. If $u(T) > \theta$, replace C by T; otherwise retain e and store the safe certificate for $T_e = T$. This preserves $u(C) > \theta$. If e survives, then $C_{\text{final}} \setminus \{e\} \subseteq T_e$, so its deletion has upper value at most $u(T_e) \leq \theta$. Every proper subset lies within such a deletion set. This proves inclusion-minimal potential conflict and, by $f \leq u$, true safety of all proper subsets. The final certificate is already available from the initial query or last accepted deletion. A lower value above θ proves a true conflict and hence true inclusion minimality. □

Including the initial query gives at most $|E| + 1$ calls. Deletion order can affect the output. Stored safe-deletion certificates remain valid by containment, whereas a violation witness for a larger set need not remain a witness after deletion. A two-sided margin in the true deviations of the queried sets is sufficient for a true inclusion-minimal output. A feasible violation for the final set, together with its safe-deletion certificates, establishes true inclusion minimality even when the margin is unknown.

E.4 DISCLOSURE CORRECTNESS AND COST CONTROL

A disclosure removes every potential conflict precisely when it intersects every minimal potential conflict set. This conflict–hitting-set relationship is standard in diagnosis (Reiter, 1987). Here it is applied to the conservative predicate $u(E) > \theta$.

Corollary 16 (Conservative hitting sets)**.** *For $Q \subseteq \mathcal{U}$, we have:*

$$Q \text{ hits every edge of } \mathcal{H}_u \quad \Longleftrightarrow \quad u(\mathcal{U} \setminus Q) \leq \theta \quad \Longrightarrow \quad f(\mathcal{U} \setminus Q) \leq \theta.$$

Proof. Every set with upper value above θ contains an inclusion-minimal such set, by finiteness. Thus a remaining potential conflict exists exactly when an edge of $\mathcal{H}_u$ is disjoint from Q. The last implication is the soundness of the upper certificate. □

For nonnegative disclosure costs c_j, let $\mathrm{OPT}_{\mathrm{true}}(t)$ minimize $\sum_{j \in Q} c_j$ subject to $f(\mathcal{U} \setminus Q) \leq t$, and let $\mathrm{OPT}_{\mathrm{cons}}(\theta)$ use u instead.

Proposition 17 (Threshold-based cost sandwich)**.** *With infeasible minima interpreted as $+\infty$, we have:*

$$\mathrm{OPT}_{\mathrm{true}}(\theta) \leq \mathrm{OPT}_{\mathrm{cons}}(\theta) \leq \mathrm{OPT}_{\mathrm{true}}(\theta - \Delta).$$

Proof. Every conservative-feasible choice is truly feasible at θ. Conversely, true feasibility at $\theta - \Delta$ implies $u \leq f + \Delta \leq \theta$. The two feasible-set inclusions give the claimed inequalities. □

The sandwich compares optimal disclosure costs across deviation tolerances. For an integer formulation of disclosure selection, an upper query checks a proposed set Q. If the remaining columns have upper deviation above the tolerance, a potential conflict $C \subseteq \mathcal{U} \setminus Q$ supplies the violated inequality $\sum_{j \in C} z_j \geq 1$ for the conservative integer master problem; validity for the true master requires a true violation certificate for C. Identical actionable-column restrictions may be imposed on all three problems, retaining explicit infeasibility.

E.5 Exact conflict membership at positive radius

The conservative predicate supports efficient certificates. Deciding the true conflict predicate also includes the exact residual optimization from Appendix D.5, as the following reduction shows.

Proposition 18 (Strong hardness of the true predicate)**.** *There are a fixed rational $S \subset \mathbb{R}^{16}$ and a fixed positive rational ρ such that, for every fixed rational $0 < \varepsilon \leq \rho$, deciding $f(E) > \theta$ is strongly NP-complete. This holds with two permanent public anchors and zero reference values for every candidate unknown column; the number of candidate columns is part of the input.*

Proof. Take the fixed-geometry construction of Appendix D.5 and fix $\varepsilon \leq \min\{\varepsilon_*, \varepsilon_B(S)\}$. Realize its fields through public anchors $u = (a+x)/2$ and $v = (a-x)/2$, with respective incident weights $H(\eta_\tau + \xi_\tau)$ and $H(\eta_\tau - \xi_\tau)$, which are nonnegative. Keep the candidate-candidate weights, set every candidate reference column to zero, and set the public-public pair weight to zero. For $E = \mathcal{U}$, we then have $r_A = c(E) = 0$ and hence $f(E) = U_\varepsilon$ for the original tube objective.

Restrict the source cut bound to $0 \leq k \leq m$, mapping trivial cases separately. The uniform reduction supplies a computable threshold τ with $U_\varepsilon \geq \tau$ in YES instances and $U_\varepsilon \leq \tau - 16\varepsilon^2$ in NO instances. Set:

$$\theta = \tau - 8\varepsilon^2.$$

Then YES implies $f(E) > \theta$ and NO implies $f(E) < \theta$. The threshold is nonnegative: the reduction has $\tau \geq 2Hg$, $H \geq 1$, $g \geq 1$, and $\varepsilon_* \leq 1/40000$. Fixed geometry and radius leave all variable numerical data polynomially bounded, preserving strong hardness. Finally, multilinearity attains the maximum at product vertices of the rational tube. Such a violating completion has polynomial encoding length and is checked by exact feasibility and score evaluation, proving NP membership. □

The hardness concerns the true predicate; the upper predicate u remains polynomially computable. A sound polynomial-time interface that always resolved the true threshold would imply $\mathrm{P} = \mathrm{NP}$; the complementary safety condition is $f(E) \leq \theta$. Margins or additional structure can resolve particular instances.

E.6 Exact examples and interpretation

The first example shows how pairwise score terms create a minimal conflict involving three columns. The second verifies extraction and minimum-cost disclosure when the representations have nonzero residuals.

Take $n = 2$, $S = \mathbb{R}^2$, and $P = I$. The two permanent public columns e_1, e_2 fix the space and have zero incident weights. Three candidate columns have zero reference values and unit weights on all three pairs. Thus $K_\varepsilon = [0,1]^2$ at every radius, giving:

$$f(E) = V(E) = 2\binom{|E|}{2}.$$

The residual is identically zero, so an exact upper predicate is:

$$\widetilde{u}(E) = V(E) = f(E).$$

For either $\theta = 3$ or $\theta = 129/64$, every pair is safe with gain 2, while the triple has gain 6. Thus $\mathcal{H}_{\widetilde{u}} = \mathcal{H}_f = \{\mathcal{U}\}$. Four certificate queries extract this true inclusion-minimal conflict, and disclosing any one candidate is sufficient and has minimum unit cost 1. This three-element hyperedge records a threshold dependency of the pairwise score.

For a nonzero-residual example, use the geometry, radius, and score of Appendix H.7. Realize its fields with the two permanent anchors u, v, and give the three candidate columns zero reference values. Choose the threshold from the certificates as follows:

$$\theta = \frac{1}{2}\left(\max_{E \subsetneq \mathcal{U}} u(E) + \ell(\mathcal{U})\right).$$

The maximum proper-subset upper bound is strictly below the full-set feasible lower bound. All seven proper subsets are certified safe, and the full set has a feasible violation. Exact original-domain enumeration for all eight subsets confirms these decisions. Each of the six deletion orders uses four queries and extracts the triple. With candidate disclosure costs $(1, 2, 4)$, disclosing the first candidate is optimal: all three optima in Proposition 17 equal 1, with $\Delta = 3/65536$.

Appendix I treats fixed references with residual components, and Appendix G discusses representation equivalence.

F Analytic mechanisms and geometric scaling

The certificate raises distinct questions about first-order accuracy, the size of the response chain, and the sharpness of its width bound. The constructions below address each question with exact examples. Computational comparisons on a prespecified finite grid appear in Appendix H.

F.1 The five analytic mechanisms

Table 3 compares first-order sensitivity across a parametrized family of planes, an orthogonal-ray geometry, and a native-field example with tied nominal optima. For M1–M3, set $S = \operatorname{col} R(a)$, where $R(a)$ has rows $(1, 1)$, $(0, 1/2)$, and $(a, 0)$. These cases use $g = 2$, $w_{12} = 1$, and $h_1 = h_2 = 0$, with the three values of a listed in the table. Exact rational recovery yields the reported slopes. Appendix C derives the three-dimensional case M5; complete input specifications and query radii accompany the supplementary results.

Table 3: Five analytic mechanisms. M1–M3 use the same oracle on the plane family defined by $R(a)$. M4 covers $u \perp v$, where the projector is entrywise nonnegative and the common cone reduces to C. M5 is the three-dimensional mechanism of Section 3.4: its projected-value error equals $nW\varepsilon^2$, while its enclosure has zero width. Bit counts give the denominator length of the slope-query radius; vertex counts and enclosure evaluations use the separately specified geometric radii.

Case	Geometry	Exact slope	Slope-query denom. (bits)	Projected vertices
M1	parametrized plane, $a = 1/2$	$9/4$	2,268	5
M2	parametrized plane, $a = 1/10$	$21/4$	4,246	5
M3	parametrized plane, $a = 1/100$	$2001/40$	7,914	5
M4	orthogonal extreme rays	$215/108$	4,416	4
M5	native fields, tied nominal optima	6	1,384	9

F.2 A native family with growing exposed chains

The next native rank-two family proves that the number of available exposed labels has no uniform constant bound as dimension grows and radius decreases.

For an integer $m \geq 2$ put $n = m + 2$,

$$d = \frac{1}{8m^2}, \qquad t_j = \frac{2j - m - 1}{16m^2} \quad (j = 1, \ldots, m),$$

so that $t_{j+1} - t_j = d$ and $|t_j| < 1/(16m)$. Let $R_m \in \mathbb{Q}^{n \times 2}$ have rows

$$(1, 0), \quad (0, 1), \quad r(t_j) := \left(1 + t_j - t_j^2,\ 1 - t_j - t_j^2\right),$$

let $S_m = \operatorname{col}(R_m)$ and set $\varepsilon_m = d^2/64 = 1/(4096m^4)$. The first two rows fix the rank and imply $S_m \cap \mathbb{R}^n_+ = R_m \mathbb{R}^2_+$. Let $\ell_D(X)$ count the vertices of a polygon $X \subset S_m$ exposed by nonzero directions in the cone D from (11).

Lemma 19 (Growing exposed chains). *For every $m \geq 2$, $0 < \varepsilon_m < \varepsilon_B(S_m)$ and $\ell_D(Q_{\varepsilon_m} \cap D) \geq m - 1 = n - 3$.*

Proof. Step 1: nominal vertices. For adjacent parameters $a = t_j$, $b = t_{j+1}$ put:

$$x^{(j)} = \Big(\frac{1 + a + b}{2(1 + ab)},\ \frac{1 - a - b}{2(1 + ab)}\Big).$$

A direct expansion gives $r(t)\, x^{(j)} = 1 - (t-a)(t-b)/(1+ab)$ for every t, so $r(a)x^{(j)} = r(b)x^{(j)} = 1$ and $r(t_k)x^{(j)} \leq 1$ for every grid parameter, because no grid parameter lies strictly between a and b. Both coordinates of $x^{(j)}$ lie in $(0, 1)$ and R_m is nonnegative, so $R_m x^{(j)} \in K_0$. Adjacent rows are independent, hence the $m - 1$ points $x^{(j)}$ are distinct vertices of the nominal domain.

Step 2: native exposing directions. Let $A = \sum_k (1 - t_k^2)^2$ and $B = \sum_k t_k^2$. Symmetry of the grid gives:

$$G = R_m^\top R_m = \begin{pmatrix} 1 + A + B & A - B \\ A - B & 1 + A + B \end{pmatrix},$$

with eigenvalues $1 + 2A$ and $1 + 2B$ on the sum and difference directions. Put $w_j = r(a) + r(b) = (c + s, c - s)$ with $c = 2 - a^2 - b^2$ and $s = a + b$. Since $|s|/c \leq 1/(8m) < 1/(2m + 1) \leq (1 + 2B)/(1 + 2A)$, the vector $z_j = G^{-1} w_j$ has two strictly positive coordinates, so $h^{(j)} = R_m z_j \in C$ is a native direction with $\langle h^{(j)}, R_m x\rangle = w_j^\top x$ for all x.

Step 3: the radius lies in the proved regime. By symmetry, the two raw columns have the same maximum coordinate κ. After maximum-coordinate normalization, the first two rows give $M = \kappa$, while $\gamma = \max_j (u_{0j} + v_{0j})/\kappa$ for the raw columns u_0, v_0. Hence $\gamma M \leq 2$; the common normalization also leaves q unchanged. From $|t_j| \leq 1/32$ one has $A - B \geq m/2$ and $1 + A + B \leq 2m$, so $q = (A-B)/(8(1+A+B)) \geq 1/32$ and $\varepsilon_B(S_m) = q/(2\gamma M) \geq 1/128 > 1/65536 \geq \varepsilon_m$.

Step 4: distinct maximizing faces. Maximize $w_j^\top x$ over $\{x : R_m x \in Q_\varepsilon \cap D\}$. The value is at least 2 at $x^{(j)}$, while $r(a)x \leq 1 + \varepsilon$ and $r(b)x \leq 1 + \varepsilon$ on the domain, so every maximizer satisfies $|r(a)x - 1| \leq \varepsilon$ and $|r(b)x - 1| \leq \varepsilon$. The adjacent-row matrix has determinant $-2d(1 + ab)$ and inverse of infinity norm at most $2/d$, hence $\|x - x^{(j)}\|_\infty \leq 2\varepsilon/d$. The coordinate differences $x_1^{(j)} - x_2^{(j)}$ increase strictly in j, with consecutive gap $2d(1 - b^2)/((1 + ab)(1 + bc)) \geq d$ for consecutive parameters a, b, c; distinct centres are therefore at least $d/2$ apart. At $\varepsilon = \varepsilon_m = d^2/64$ each maximizing face lies in a neighbourhood of radius $d/32$ of its own centre, and these neighbourhoods are disjoint. Each face contains a vertex. Since $h^{(j)}$ lies in the relative interior of C, a sufficiently small rational perturbation within C uniquely selects a vertex of that face. Vertices outside the face have a positive objective gap under the original direction, so small enough perturbations do not select them. The disjoint faces therefore give at least $m - 1$ distinct strictly natively exposed vertices, also exposed by directions in D. □

The family establishes unbounded available chain length at decreasing radii $\varepsilon_m = \Theta(m^{-4})$. The distinction between available labels and labels required by an optimal assignment is examined in Appendix H.6.

F.3 Fixed-radius response counts for the growing family

Fixing the radius separates the effect of growing dimension from that of shrinking the uncertainty set. For the same row family, consider the fixed radius $\varepsilon = 1/512$. Define $R_m \in \mathbb{Q}^{(m+2)\times 2}$ by the rows:

$$(1, 0), \quad (0, 1), \quad (a_j + t_j, a_j - t_j), \qquad t_j = \frac{2j - m - 1}{16m^2}, \quad a_j = 1 - t_j^2 \quad (1 \leq j \leq m).$$

Let $S_m = \operatorname{col} R_m$, let P be its orthogonal projector, and set $K_\varepsilon = \{b \in [0, 1]^{m+2} : \|(I - P)b\|_\infty \leq \varepsilon\}$ and $Q_\varepsilon = PK_\varepsilon$. The first two rows give $C_m = S_m \cap \mathbb{R}_+^{m+2} = R_m \mathbb{R}_+^2$. For the raw columns u_0, v_0, put $T = (m - 1)/(16m^2)$ and $\kappa = 1 + T - T^2$. Their maximum-coordinate-normalized rays are $u = u_0/\kappa$ and $v = v_0/\kappa$. Use the cone $D_m = \operatorname{cone}(u - q_m v, v - q_m u)$ from (11). Let ℓ_C count vertices uniquely maximizing a nonzero field in C_m; ℓ_D counts the vertices of the D_m response chain.

Proposition 20 (Fixed-radius boundary)**.** *Fix $\varepsilon = 1/512$. For every integer $m \geq 32$, $\varepsilon < \varepsilon_B(S_m)$ and:*

$$\ell_C(Q_\varepsilon \cap D_m) = 3, \qquad \ell_D(Q_\varepsilon \cap D_m) = 5.$$

These are response counts, not the total number of projected vertices.

Proof. The symmetry $t_{m+1-j} = -t_j$ holds for both parities. The resulting moments and Gram eigenvalues are:

$$B = \sum_j t_j^2 = \frac{m^2-1}{768m^3}, \qquad E = \sum_j t_j^4 = \frac{(m^2-1)(3m^2-7)}{983040m^7},$$
$$A = \sum_j a_j^2 = m - 2B + E, \quad H = \sum_j a_j = m - B, \qquad L_+ = 1 + 2A, \quad L_- = 1 + 2B.$$

Indeed, $R_m^\top R_m$ has diagonal $1 + A + B$ and off-diagonal $A - B$. Define:

$$d_m = \min_j t_j^2 = \begin{cases} 0, & m \text{ odd}, \\ 1/(256m^4), & m \text{ even}. \end{cases}$$

The normalized rays satisfy $\gamma = 2(1 - d_m)/\kappa$, and give:

$$q_m = \frac{A - B}{8(1 + A + B)}, \qquad M = \kappa, \qquad \gamma M = 2(1 - d_m).$$

The bounds $A - B \geq m/2$ and $1 + A + B \leq 2m$ imply $q_m \geq 1/32$ and $\varepsilon_B = q_m/(2\gamma M) \geq 1/128 > \varepsilon$.

Five feasible projected points. Write $s = R_m x$, $p = x_1 + x_2$, and $y = x_1 - x_2$. Its coordinates are:

$$s_1 = (p + y)/2, \qquad s_2 = (p - y)/2, \qquad s_{j+2} = a_j p + t_j y.$$

Define:

$$p_0 = \frac{H + \varepsilon}{A}, \qquad y_R = \frac{2A(1-\varepsilon) - (H + \varepsilon)}{A - B}, \qquad p_R = 2(1 - \varepsilon) - y_R,$$
$$\beta = 1 - L_- y_R, \qquad p_E = \frac{1 + 2H}{L_+}, \qquad y_E = \frac{1}{L_-}.$$

The names below denote projected points through their (p, y) coordinates:

Point	(p, y)	Original-domain lift b
O	$(p_0, 0)$	$(p_0/2 + \varepsilon, p_0/2 + \varepsilon, \mathbf{1}_m)$
R	(p_R, y_R)	$(1, \beta, \mathbf{1}_m)$
L	$(p_R, -y_R)$	$(\beta, 1, \mathbf{1}_m)$
E_+	(p_E, y_E)	$(1, 0, \mathbf{1}_m)$
E_-	$(p_E, -y_E)$	$(0, 1, \mathbf{1}_m)$

When the last m lift coordinates equal one, the two moment equations for $Pb = s$ are exactly:

$$L_+ p = b_1 + b_2 + 2H, \qquad L_- y = b_1 - b_2.$$

Each listed lift satisfies these moment identities.

For $m \geq 32$, $B < \varepsilon/2$, $A > 3m/4$, and $0 < p_0 - 1 = (B - E + \varepsilon)/A < 2\varepsilon/m$. For a lift $(1, \xi, \mathbf{1}_m)$, its first residual is:

$$f(\xi) = 1 - \frac{1}{2}\left(\frac{1 + \xi + 2H}{L_+} + \frac{1 - \xi}{L_-}\right).$$

Here $f' > 0$, $f(0) < B < \varepsilon$, $f(1) > 1/6 > \varepsilon$, and $f(\beta) = \varepsilon$. Thus $0 < \beta < 1$, giving $1 < p_E < p_R < p_0$ and $0 < y_R < y_E < 1$. The cube constraint for O follows from:

$$p_0/2 + \varepsilon < 1/2 + \varepsilon(1 + 1/m) < 1.$$

The first two residuals of O are $(\varepsilon, \varepsilon)$; those of R are $(\varepsilon, \varepsilon - 2By_R)$, with L symmetric. For $E_\pm$, both absolute residuals are less than $\varepsilon/m + B < 5\varepsilon/(3m) < \varepsilon$. All five points satisfy $1 < p \le p_0$ and $|y| < 1$, hence lie in the relative interior of C_m. For every integer $m \ge 35$, their tail residuals obey:

$$|1 - p + pt_j^2 - t_j y| < \frac{2\varepsilon}{m} + \frac{1}{128m^2} + \frac{1}{16m} = \varepsilon\left(\frac{34}{m} + \frac{4}{m^2}\right) \le \frac{1194}{1225}\varepsilon < \varepsilon.$$

For the remaining cases $m \in \{32, 33, 34\}$, direct substitution into the same rational formulas gives, for each of the five lifts,

$$\varepsilon - \max_{1 \le j \le m} |e_{j+2}| > \frac{\varepsilon}{256} > 0.$$

These finite inequalities, together with the preceding cube and first-two-coordinate bounds, complete feasibility of all five lifts.

Support over the full domain. For arbitrary $b \in K_\varepsilon$, write $s = Pb$, $e = b - s$. Orthogonality gives:

$$Ap + By = e_1 + \sum_j (a_j + t_j) b_{j+2}, \qquad Ap - By = e_2 + \sum_j (a_j - t_j) b_{j+2}.$$

Since all tail coefficients are positive, the residual and box bounds give four global support inequalities:

$$Ap \pm By \le H + \varepsilon, \qquad L_+ p \pm L_- y \le 2(H + 1). \tag{56}$$

The second pair also follows from $u_0, v_0 \ge 0$ and $\langle u_0, Pb\rangle = \langle u_0, b\rangle$, with the analogous identity for v_0.

Put $\eta = B/A$, $\rho = L_-/L_+$, and $\rho_D = \rho(1 + q_m)/(1 - q_m)$. Since $AL_- - BL_+ = A - B > 0$, $0 < \eta < \rho < \rho_D$. For $h = R_m z$, the objective has the form:

$$\langle h, s\rangle = \frac{L_+(z_1 + z_2)}{2}(p + \lambda y), \qquad \lambda = \rho\frac{z_1 - z_2}{z_1 + z_2}.$$

The prefactor is positive for nonzero $h \in D_m$. Native fields cover $[-\rho, \rho]$, and D_m fields cover $[-\rho_D, \rho_D]$. For $|\lambda| < \eta$, the following is a strictly positive combination:

$$p + \lambda y = \frac{1 + \lambda/\eta}{2A}(Ap + By) + \frac{1 - \lambda/\eta}{2A}(Ap - By).$$

Both bounds must be tight, uniquely selecting O. For $\eta < \lambda < \rho$, write $p + \lambda y = \alpha(p + \eta y) + (1 - \alpha)(p + \rho y)$ with $\alpha = (\rho - \lambda)/(\rho - \eta) \in (0, 1)$. Equality in the two corresponding bounds in (56) uniquely selects R; their determinant is $A - B > 0$. Reflection selects L on the negative interval.

At $\lambda = \eta$, equality forces every tail lift coordinate to one and $e_1 = \varepsilon$. Then $p = p_0 - \eta y$ and $e_2 = \varepsilon - 2By$; $e_2 \le \varepsilon$ and $b_1 \le 1$ give $0 \le y \le y_R$. The maximizing face is therefore exactly $[O, R]$, since its endpoints are feasible and the domain is convex. At $\lambda = \rho$, equality forces $b_1 = 1$ and every tail coordinate to one. The remaining coordinate $b_2 = \xi$ satisfies $0 \le \xi \le \beta$ by the cube bound and $f(\xi) \le \varepsilon$. The face is exactly $[E_+, R]$; reflection gives the negative faces.

The tail coordinates of both generating fields of D_m are positive, being proportional to $(1 - q_m)a_j \pm (1 + q_m)t_j$. For $\rho < \lambda \le \rho_D$, the field has signs $(+, -, +, \ldots, +)$, so its unique cube maximizer is the already feasible lift $(1, 0, \mathbf{1}_m)$. It uniquely selects E_+ over the full Q_ε; the negative interval selects E_-. This exhausts the directions and their tie faces. All resulting faces lie in $C_m \subset D_m$, so intersecting with D_m changes none of them. Only L, O, R admit strict native exposure; $E_\pm$ are tied in the native endpoint directions but strictly exposed by the extreme D_m fields.

For $r = u - q_m v$, $s = v - q_m u$, and response coordinates $(\langle r, x\rangle, \langle s, x\rangle)$, decreasing the first coordinate and increasing the second orders the five response vertices as:

$$E_+ \longrightarrow R \longrightarrow O \longrightarrow L \longrightarrow E_-.$$

Explicit strict native fields for O, R, L are respectively $R_m(1, 1)^\top$, $R_m(3A + B + 8AB, A - B)^\top$, and its coefficient-swapped counterpart; their slopes are $0, (\eta + \rho)/2, -(\eta + \rho)/2$. This proves both counts. All listed vector coordinates are fixed-degree rational expressions in m, j, giving $O(m \log m)$ total encoding length for these vector certificates. □

F.4 Error orders in the closed-form example

On the mechanism of Section 3.4, the certificate was evaluated at $\varepsilon = 2^{-k}$ for $k = 6, \dots, 14$, with threshold $t_\varepsilon = 5 + 3\varepsilon$. At each tested radius, the projected-value error equals $3\varepsilon^2$, the nominal error equals 6ε, their ratio is $\varepsilon/2$, and the enclosure width is 0. The nominal value 5 is below t_ε, yet the certificate returns Conflict with both witness columns in the cube. These exact identities give the linear and quadratic error orders.

F.5 Enclosure width and threshold reversal

To examine enclosure width and threshold reversal beyond one example, we evaluate 107 exact instances. The comparison covers the closed-form example at three radii, four elementary families at $g \in \{2, 4, 8\}$ and two radii, and forty random native instances with $n \in \{3, 4, 5\}$ and $g \in \{2, 3\}$ at two radii within $\varepsilon_B(S)$. All 107 enclosures are valid. Their widths are at most half of $2nW\varepsilon^2$; the maximum ratio is exactly $1/2$, and 24 of 96 positive-width cases exceed one quarter of the bound. The margin $J(B^{\mathrm{lift}}) - U_0$ is positive in 105 cases, giving an instance-specific reversal interval $[U_0, J(B^{\mathrm{lift}}))$; the other two have zero first-order response. The separate extremal construction below attains the full width bound.

An exact extremal instance. The prescribed maximum-value enclosure attains the constant 2 in (5), even with a unique projected optimizer and unique columnwise lifts. Take

$$R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 2 & 1 \\ 1 & 2 \end{pmatrix}, \qquad S = \operatorname{col} R, \qquad u = \tfrac{1}{2}(1, 0, 2, 1)^\top, \quad v = \tfrac{1}{2}(0, 1, 1, 2)^\top,$$

with $g = 2$, $w_{12} = 1$, $h_1 = 6u$, and $h_2 = 6v$. The independent public columns u, v lie in the cube, and the first two rows of R give $C = \operatorname{cone}(u, v)$. Thus the fields are native. The projector is:

$$P = \frac{1}{10} \begin{pmatrix} 3 & -2 & 4 & -1 \\ -2 & 3 & -1 & 4 \\ 4 & -1 & 7 & 2 \\ -1 & 4 & 2 & 7 \end{pmatrix}.$$

For the normalized rays u, v, (11)–(12) give $q = 1/12$, $\gamma = 3/2$, $M = 2$, and $\varepsilon_B = 1/72$. For any rational $0 < \varepsilon \le 1/72$, put:

$$\begin{aligned} p &= (1/3, 1/3, 1, 1)^\top, & d &= (1, -1, 1, -1)^\top, \\ z &= (1, -1, -1, 1)^\top, & s_\pm &= p \pm \varepsilon d, \\ b_\pm &= p \pm 2\varepsilon(1, -1, 0, 0)^\top. \end{aligned}$$

Direct multiplication gives $Pb_\pm = s_\pm$ and $(I - P)b_\pm = \pm\varepsilon z$. Since $\varepsilon \le 1/72 < 1/6$, $b_\pm \in K_\varepsilon$.

Projected optimality and uniqueness. Write $N = I - P$ and let f_k be the kth standard basis vector. For an arbitrary opponent lift $b \in [0, 1]^4$, define:

$$\begin{aligned} \beta &= (6 - b_1 + 2b_2 + 3b_4)/3 \ge 5/3, \\ \eta &= (12 + 4b_1 - 5b_2 + 3b_3 - 6b_4)/3 \ge 1/3, \\ \alpha &= 2\beta + \eta > 0. \end{aligned}$$

The displayed projector gives the identity:

$$h_1 + Pb = \alpha f_3 + \beta f_4 + \beta N f_1 - \eta N f_3.$$

Consequently every $x \in K_\varepsilon$ satisfies:

$$\begin{aligned} \langle h_1 + Pb, x \rangle &= \alpha x_3 + \beta x_4 + \beta(Nx)_1 - \eta(Nx)_3 \\ &\le \alpha + \beta + (\beta + \eta)\varepsilon. \end{aligned}$$

The point b_+ attains the bound. Since α, β, η are positive, equality forces $x_3 = x_4 = 1$, $(Nx)_1 = \varepsilon$, and $(Nx)_3 = -\varepsilon$. With $e = Nx$, the equations $R^\top e = 0$ then give $e_4 = \varepsilon$ and $e_2 = -\varepsilon$. The last

two coordinates of Px are therefore $1+\varepsilon$ and $1-\varepsilon$. Independence of the final rows of R determines $Px = s_+$ and then $x = b_+$. Thus b_+ is the unique maximizer of this conditional projected objective for *every* opponent lift in the cube.

The permutation swapping coordinates $1, 2$ and $3, 4$ commutes with P and swaps h_1, h_2 and b_+, b_-. Symmetry gives the unique conditional maximizer b_- for the second column, hence the unique projected optimizer (s_+, s_-). Every lift of s_+ has the same value under the first conditional functional. Uniqueness therefore forces its lift to be b_+; likewise, the unique lift of s_- is b_-. □

The identities $\langle p, d\rangle = 0$, $\|p\|^2 = 20/9$, $\|d\|^2 = \|z\|^2 = 4$, $\langle u, p\rangle = \langle v, p\rangle = 5/3$, and $\langle u, d\rangle = 1 = -\langle v, d\rangle$ now give:

$$\begin{aligned} V_\varepsilon &= 200/9 + 12\varepsilon - 4\varepsilon^2, \\ J(b_+, b_-) &= V_\varepsilon - 4\varepsilon^2, \\ \overline{U} &= V_\varepsilon + 4\varepsilon^2 = 200/9 + 12\varepsilon, \\ \overline{U} - J(b_+, b_-) &= 8\varepsilon^2 = 2nW\varepsilon^2. \end{aligned}$$

This equality establishes sharpness of the prescribed enclosure width. It compares the feasible lower endpoint with the certified upper endpoint. At the concrete radius $\varepsilon = 1/144$, all outputs are rational:

$$\begin{gathered} b_+ = (25/72, 23/72, 1, 1)^\top, \qquad b_- = (23/72, 25/72, 1, 1)^\top, \\ V_\varepsilon = 115631/5184, \quad \underline{U} = 57815/2592, \quad \overline{U} = 803/36, \\ \overline{U} - \underline{U} = 1/2592 > 1/5184 = nW\varepsilon^2. \end{gathered}$$

Exact evaluation at the sixteen cube vertices confirms the coefficient bounds used in the conditional optimality argument. Evaluations at several radii agree with the closed forms above; at $\varepsilon = 1/144$, projected optimization and lifting yield the displayed enclosure.

F.6 Linear and quadratic error scales

Figure 4 illustrates the two error orders on the closed-form example.

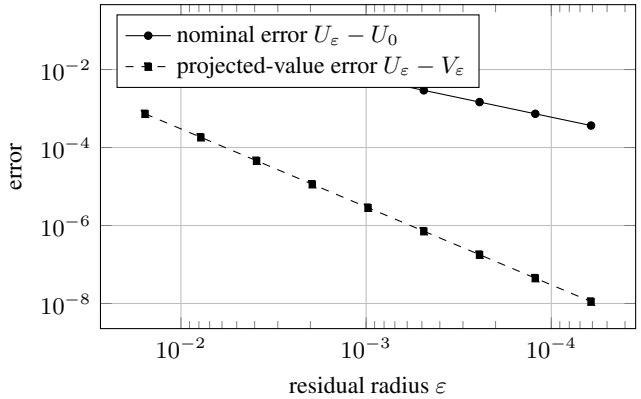

Figure 4: For the closed-form example of Section 3.4 at $\varepsilon = 2^{-6}, \ldots, 2^{-14}$, the nominal error is 6ε (slope one) and the projected-value error is $3\varepsilon^2$ (slope two). These identities hold at every tested radius and show that projection retains the full first-order response. The enclosure has zero width at each radius, giving an exact maximum-value certificate on this instance.

G Supplementary related work

The comparisons below distinguish the uncertainty model, the representation of joint conflicts, and the geometric structure that makes projected optimization tractable.

Table 4: Scope of the closest ingredients. The table compares the assumptions and structural ingredients of the cited formulations.

Ingredient	Supplied structure	Additional step here
Fixed-rank optimization (Hladík et al., 2021)	Fixed rank of the full quadratic form over a box	Shared rank two per vector; the global interaction rank can grow with the number of columns.
Ordered-label cuts (Kolmogorov & Zabih, 2004; Darbon, 2009)	Discrete labels and graph-representable pairwise terms	Derive the finite ordered chain and required signs from the continuous residual tube.
Bilinear envelopes (McCormick, 1976)	Coordinate bounds give a valid convex relaxation	Solve the projected geometric problem exactly, then bound the loss from original-domain lifting.
Multiterminal cut (Dahlhaus et al., 1994)	Discrete labels and prescribed terminal separation	Realize the labels in the original fixed continuous tube using native fields, with a uniform-radius argument.

Sinha et al. (2018) study Wasserstein distributional robustness through penalized adversarial training, obtaining population guarantees under smoothness conditions. Here uncertainty is specified deterministically by a coordinatewise residual budget around a supplied shared subspace.

Hypergraph reconstruction from clique projections addresses information lost by the observation map (Wang & Kleinberg, 2024); equivariant hypergraph diffusion admits message passing on a star expansion (Wang et al., 2023). Matrix, vector, and incidence-graph implementations give the same predicates when they preserve the feasible set and score function. The conflict hypergraph in Appendix E records joint threshold dependencies and their disclosure certificates.

Multiterminal cuts and finite-valued constraint languages provide the classical hardness mechanisms (Dahlhaus et al., 1994; Thapper & Živný, 2016). Appendix D contributes the realization of the required label choices inside the original residual tube using native fields.

Implicit set representations and value functions. Constrained zonotopes represent linear images and intersections implicitly (Scott et al., 2016); the explicit $O(n^3)$ description of Lemma 4 instead uses two moments and coordinatewise clipping. Convex piecewise-quadratic value functions (Patrinos & Sarimveis, 2011) give a further comparison for the local expansion of Theorem 2, which is quadratic on an initial interval but arises from residual choices in one continuous domain.

Global rank and residual certificates. Soltanalian & Mousavi (2026, Theorem 3.2 and Corollary 4.4) connect PSD quadratic optimization to linear exposure and enumeration in global rank space. Writing $s_i = Rx_i$ and $G = R^\top R$, our projected Hessian is $A_w \otimes G$, where A_w has off-diagonal entries w_{ij} and zero diagonal. Its rank is $2\,\mathrm{rank}(A_w)$, which can grow with g, and it is indefinite whenever $A_w \neq 0$. Their Theorem 7.1 permits an arbitrary symmetric residual around a PSD core whose exact optimizer is available; the residual may make the full objective indefinite. Our natural projected core is itself generally indefinite. Both certificates bound a residual around a solvable objective. Here common-cone response chains solve the projected problem, and the coordinatewise tube supplies the residual budget.

H Same-domain certification and exact lifting

Certificate quality depends on how tightly an interval encloses the original maximum and how often it resolves a threshold query. We compare both quantities on explicitly specified feasible domains. Six exactly solvable cases connect the bounds to analytic mechanisms. A geometric lifting construction then supports a broader comparison across geometry, graph size, and field strength, including the distinction between available and necessary response labels.

H.1 A ONE-SIDED SIMILARITY THRESHOLD EXAMPLE

Consider four nodes with fixed anchor vectors $A_3 = u = (1, 0, 1)$ and $A_4 = v = (0, 1, 1)$. Set $w_{12} = w_{13} = w_{24} = 1$ and all other pair weights to zero. Nodes 1 and 2 have unknown updated vectors $b_1, b_2 \in K_\varepsilon$, where $S = \{z = x + y\}$. Their score has the following form:

$$J(b_1, b_2) = \langle b_1, b_2 \rangle + \langle u, b_1 \rangle + \langle v, b_2 \rangle.$$

The reference assignment $A_1 = u, A_2 = v$ gives $r_A = 5$. This value fixes the one-sided thresholds $r_A + \delta$ before evaluation. The question is whether every admissible score is at most that threshold, or a feasible exceedance can be returned. The domain is the full tube K_ε, equivalently the cube intersected with $|(b_i)_3 - (b_i)_1 - (b_i)_2| \le 3\varepsilon$ for each column b_i. The full tube includes residual directions beyond nonnegative mixtures of the anchors.

H.2 A FIXED SAME-DOMAIN COMPARISON

Instances and thresholds. The two families test different certificate properties: exactness in the three-dimensional example and attainment of the width bound in the four-dimensional construction. Each certificate is compared against a known original-domain maximum. Both families have $g = 2$, $w_{12} = 1$, and the common original domain K_ε^2 for every method. The three-dimensional family (3D) is the four-node construction above, with $U_0 = 5$ and $\varepsilon \in \{1/256, 1/128, 1/64\}$. The four-dimensional family (4D) is the enclosure-width construction of Section F.5:

$$R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 2 & 1 \\ 1 & 2 \end{pmatrix}, \qquad u = (1, 0, 2, 1)/2, \quad v = (0, 1, 1, 2)/2, \quad h_1 = 6u, \quad h_2 = 6v.$$

Here $U_0 = 200/9$, and we use $\varepsilon \in \{1/288, 1/144, 1/72\}$, namely one quarter, one half, and all of its specified geometric radius. The public feasible reference is (u, v) for 3D and (p, p) for 4D, where $p = (1/3, 1/3, 1, 1)$. Exact enumeration also verifies both nominal values. For each of the six positive-radius instances, the three thresholds are $\theta = U_0 + \delta$ with $\delta \in \{1/64, 1/16, 1/8\}$. These choices were fixed before evaluating positive-radius answers; the margins are absolute, not normalized across families.

Compared intervals. The cube baseline uses the feasible reference score as its lower endpoint and the valid bound $nW + \sum_{i,j}(h_i)_j$ as its upper endpoint. The original-coordinate McCormick relaxation (McCormick, 1976) retains every cube and residual inequality of K_ε for both vectors. For each coordinate it introduces z_j for $(b_1)_j(b_2)_j$, subject to the following inequalities:

$$z_j \ge 0, \qquad z_j \ge (b_1)_j + (b_2)_j - 1, \qquad z_j \le (b_1)_j, \qquad z_j \le (b_2)_j.$$

Its objective is $\langle h_1, b_1 \rangle + \langle h_2, b_2 \rangle + \sum_j z_j$. Numerical LP proposals are accepted only after exact rational checks of primal feasibility, nonnegative dual multipliers, dual stationarity, and equality of primal and dual objectives. The lower endpoint is the larger of the reference score and the true quadratic score at the returned feasible (b_1, b_2).

The projected method combines the optimization in Section 3 with numerical lifting and exact rational certificate validation. Reference maxima are obtained independently by enumerating all full-rank active-constraint intersections of the original tube, retaining exactly feasible vertices, and maximizing over their Cartesian square. Separate affinity guarantees a product-vertex optimum. Rational arithmetic makes both vertex feasibility and objective comparisons exact; the resulting maximum determines the excess of each upper bound.

Threshold decisions. All six instances return validated intervals for all methods. McCormick and projection–lifting tie on all nine 3D queries. On 4D, projection–lifting resolves five additional queries: three as safe and two with feasible violations. Over the full grid, the projected method returns eight safe decisions and ten conflicts; McCormick returns five safe decisions, eight conflicts, and five unknowns; the cube bound leaves all 18 unknown. All definite decisions agree with enumeration. Table 5 and the complete per-query supplementary results show that the additional decisions arise from the four-dimensional family, while the projected enclosure and McCormick relaxation coincide on the three-dimensional family.

Table 5: Exact upper-endpoint excess over the enumerated maximum. The projected lower endpoint equals U_ε in all six cases. For 3D, McCormick also returns the exact interval; for 4D, the projected excess is exactly $8\varepsilon^2$. Entries are exact rational differences.

Family	ε	U_ε	Cube gap	MC gap	Projected gap
3D	$1/256$	$643/128$	$253/128$	0	0
3D	$1/128$	$323/64$	$125/64$	0	0
3D	$1/64$	$163/32$	$61/32$	0	0
4D	$1/288$	$230831/10368$	$59473/10368$	$1499/3456$	$1/10368$
4D	$1/144$	$57815/2592$	$14761/2592$	$1097/2592$	$1/2592$
4D	$1/72$	$14507/648$	$3637/648$	$29/72$	$1/648$

The 4D construction attains the enclosure-width constant. The 18 threshold queries share six feasible sets. The broader comparison in Section H.5 tests additional geometries, column counts, and field strengths; all individual outcomes of both studies remain in the supplementary results.

H.3 Structural exact lifting

A projected optimizer must be converted into an original feasible completion before its score can serve as a lower bound. The two moment constraints permit a direct geometric construction of that lift.

A two-moment construction. The prepared basis R_+ consists of the two normalized nonnegative extreme rays. For a projected point $s \in S$, put $L_j = \max\{0, s_j - \varepsilon\}$ and $T_j = \min\{1, s_j + \varepsilon\}$. A lift must satisfy $R_+^\top b = R_+^\top s$ with $L \le b \le T$. Its two-dimensional box image has the following form:

$$R_+^\top L + \sum_{j=1}^{n} [0, (T_j - L_j)(R_+)_{j,:}^\top].$$

The geometric lifting method sorts these nonnegative planar generators by direction and walks the boundary, retaining an original box corner for each boundary vertex. Adding or subtracting a generator switches its associated coordinate between L_j and T_j. A triangle containing the target $R_+^\top s$ provides nonnegative barycentric coefficients summing to one. Applying them to the corresponding original box corners returns a box-feasible vector with the required moments, hence $Pb = s$. Point and segment images are handled separately; parallel generators and zero-width coordinates are allowed. Every returned vector is checked again in the original cube and residual constraints using rational arithmetic. The normalized nonnegative basis supplies the generator signs required by this construction.

H.4 The roles of the radius bounds

The geometric radius ε_B guarantees the common-cone reduction used by the projected algorithm. In Appendix B, ε_F stabilizes feasible affine vertex branches and ε_J stabilizes objective ordering; the local identity uses both restrictions. Exact recovery of Λ uses the still smaller radius in (29), which also controls rational separation and enclosure error. These sufficient bounds have polynomial encoding length. Ordinary certificate queries use the geometric range, whereas exact coefficient recovery additionally uses the local stability and rational-separation bounds.

The distinct ε_* in Appendix D.5 gives a graph-independent interval for the fixed 16-dimensional hardness construction; its field magnitude still scales with the source graph. This hardness radius is specific to that construction. An instance-specific identity can hold beyond a sufficient general algorithmic guarantee: for 3D, the closed form $U_\varepsilon = 5 + 6\varepsilon$ holds on $0 \le \varepsilon \le 1/3$, whereas the supplied common-cone guarantee is $\varepsilon_B = 1/64$. The comparisons above stay within the corresponding geometric ranges.

H.5 A BROADER FINITE COMPARISON

The six exact cases above isolate particular mechanisms. The broader comparison asks how certificate width and threshold resolution vary with geometry, graph density, field strength, dimension, and the number of unknown columns. The fixed grid uses the same inputs and threshold formulas for all applicable methods. Its 98 optimization instances span nine distinct subspaces and twelve subspace–radius pairs.

Domains, graphs, and fields. Besides the family R_m of Appendix F.2, two families start with rows $(1,0),(0,1)$. Their remaining $m = n-2$ rows are either $(1+t_j-t_j^2, 1-t_j-t_j^2)$ with $t_j = (2j-m-1)/(8m)$, or $(a_j/8, b_j/8)$ with independent integers uniform on $\{2,\dots,16\}$. Cycle and complete graphs share the same prespecified random streams. Raw edge weights are uniform on $\{1,2,4,8\}$ for these two families and one for R_m; normalization gives $W = 1$. With normalized extreme rays u, v and $d_i = \sum_j w_{ij}$, fields are:

$$h_i = \frac{\kappa d_i}{2}\big((1-\lambda_i)u + \lambda_i v\big), \qquad S_{\text{score}} = n(1+\kappa).$$

This construction supplies native fields without restricting the theorem's general condition $h_i \in C$. The mixed-geometry block uses two fixed seeds; random fields use λ_i uniform on $\{0, 1/8, \dots, 1\}$. For the R_m field mixtures, exact native exposure intervals select up to four evenly spaced eligible chain vertices; their interval midpoints are assigned cyclically to nodes. Field-strength comparisons use complete graphs at both radii. The common-field control uses $\lambda_i = 1/2$ and $\kappa = 1$.

Table 6: Prespecified instance grid. Both graphs means cycle and complete. Matched diagonal uses the pairs $(n,g) = (6,4),(10,16),(18,32)$. All three bounds cover every row and target the same full original tube; the cube method relaxes its residual constraints. McCormick coverage combines its initial 56-case allocation with a complete supplementary evaluation of the other 42 instances.

Purpose	Cases	n	g	Radius and field scale
Mixed geometry	32	$4, 8$	$4, 12$	$\varepsilon_B/2$, $\kappa = 1$, both graphs
Joint scale	36	$6, 10, 18$	$4, 16, 32$	$\varepsilon_m, 1/512$, $\kappa = 1$, both graphs
Field strength	12	$6, 10, 18$	$4, 16, 32$	Matched diagonal, $\kappa = 1/4, 4$
Common field	6	$6, 10, 18$	$4, 16, 32$	Matched diagonal, $1/512$, both graphs
Exact truth	12	4	$2, 3$	Three geometries, $\varepsilon_B/2$, $\kappa = 1/4, 1$

An exact nominal solve provides U_0 and the common feasible reference before each positive-radius query. We report two threshold groups separately:

$$\theta = U_0 + S_{\text{score}}\delta, \quad \delta \in \{0, 1/1024, 1/256, 1/64\}; \qquad \theta = U_0 + S_{\text{score}} a\varepsilon, \quad a \in \{1/4, 1, 4\}.$$

Every instance in the planned grid is evaluated, and each threshold decision is checked against a validated interval.

Exact checks for the compared bounds. The projected method uses geometric lifting and the larger of its lifted score and the common reference as its lower endpoint. The cube bound uses that reference and the cube upper bound. McCormick introduces one product variable per nonzero edge and coordinate, retaining every original tube constraint. For an LP written as $\max c^\top z$ subject to $A_{\text{LP}} z \le d$, $0 \le z \le 1$, a rational nonnegative multiplier vector y gives the valid bound:

$$\widehat{U}_{\text{MC}} = y^\top d + \sum_k (c - A_{\text{LP}}^\top y)_k^+.$$

The residual correction is evaluated exactly. An LP-feasible point attaining the validated upper bound additionally certifies LP optimality. Original-domain lower witnesses are scored with the quadratic objective at the feasible vectors. For exact maximum flow, all capacities of the completed network are uniformly scaled to integers and solved by highest-label preflow-push; values and flows are then restored to rational form. All geometric, flow, lifting, and bound checks use exact arithmetic, including exact replay of every projected optimum and independent checks of the saved flow and lift certificates.

Certificate quality across the full grid. Projection, McCormick, and the cube bound produce valid certificates on all 98 instances. The 56 prespecified McCormick comparisons are supplemented by a post-hoc evaluation of every remaining grid instance, using the same mathematical inputs, solver, and resource limits. The two execution batches remain separately identified in the supplementary records. All nominal preparations and all twelve exact-truth computations also complete with exact validation.

On all 98 pairs, the projected upper endpoint is strictly below the true McCormick relaxation optimum. To certify this comparison, take a verified original-domain witness B and set each product variable to $z_{ij\ell} = \min\{b_{i\ell}, b_{j\ell}\}$. Since both coordinates lie in $[0, 1]$, these values satisfy all four McCormick inequalities while retaining every original tube constraint. We apply this construction to the nominal reference and the two methods' feasible witnesses, and retain the largest exactly feasible LP objective L_{MC}. Writing V_{MC} for the relaxation optimum, every pair satisfies:

$$\overline{U}_{\mathrm{proj}} < L_{\mathrm{MC}} \leq V_{\mathrm{MC}} \leq \widehat{U}_{\mathrm{MC}}.$$

Thus the separation holds against the relaxation optimum itself, even when the numerical LP solve returns only a validated upper bound. The score intervals in Table 7 retain their original-domain lower witnesses.

Table 7: Paired threshold outcomes on all 98 grid instances, combining the initial 56 McCormick comparisons and the complete 42-instance coverage extension. S/C/U denote Safe/Conflict/Unknown. The rightmost column is the median interval width divided by S_{score}, measured once per instance. Thresholds on one instance are dependent queries.

Method	Fixed tolerance: S/C/U (392)	First-order scale: S/C/U (294)	Median width
Projection	$186/206/0$	$135/159/0$	8.206×10^{-7}
McCormick	$55/68/269$	$20/25/249$	1.386×10^{-2}
Cube	$0/0/392$	$0/0/294$	1.376×10^{-1}

A common feasible lower witness separates the effect of the upper bound from that of the chosen completion. Sharing the best validated original-domain witness among the three methods changes 138 fixed-tolerance and 134 first-order McCormick queries from Unknown to Conflict. This auxiliary comparison isolates the contribution of feasible lower witnesses while preserving every upper bound; it uses no truth-oracle witness.

Exact truth and near-threshold behavior. For twelve four-dimensional cases, independent enumeration covers all original-domain vertices and their products. The three distinct tubes have 43, 33, and 41 vertices. All 36 projection, McCormick, and cube intervals contain the enumerated maximum; the projected lower endpoint equals it in all twelve cases. To examine resolution near the optimum, a separately prespecified diagnostic uses $\theta = U_\varepsilon + c\, nW\varepsilon^2$ with $c \in \{-2, -1/2, 0, 1/2, 2\}$. Across its 60 queries per method, projection yields 20 Safe, 24 Conflict, and 16 Unknown; McCormick yields $0/9/51$ and the cube bound $0/0/60$. These truth-centered queries measure near-optimum resolution separately from the nominal-margin queries in Table 7.

A separable upper-bound refinement. The unresolved near-optimum queries motivate a second upper bound on the same feasible domain. Bounding each pairwise interaction by the two squared norms separates the remaining maximization across columns. The inequality $2\langle b_i, b_j\rangle \leq \|b_i\|^2 + \|b_j\|^2$ gives an additional bound on the same original domain. With $d_i = \sum_j w_{ij}$, define:

$$U_{\mathrm{sep}} = \sum_i \max_{v \in \mathrm{vert}(K_\varepsilon)} \left\{ \frac{d_i}{2}\|v\|^2 + \langle h_i, v\rangle \right\}.$$

Each single-column objective is convex, so its maximum over K_ε is attained at a vertex, and $U_\varepsilon \leq U_{\mathrm{sep}}$. On all twelve small cases, we enumerate these single-column vertices anew and replace the upper endpoint by $\min\{\overline{U}, U_{\mathrm{sep}}\}$. This additional post-hoc check improves eight upper endpoints and changes ten Unknown queries to Safe, yielding $30/24/6$ on the same sixty queries. All 48 queries with $\theta \neq U_\varepsilon$ are decided; the six remaining Unknowns occur at equality. The lower witnesses

and original projection results are unchanged. This finite-instance refinement uses explicit vertex enumeration.

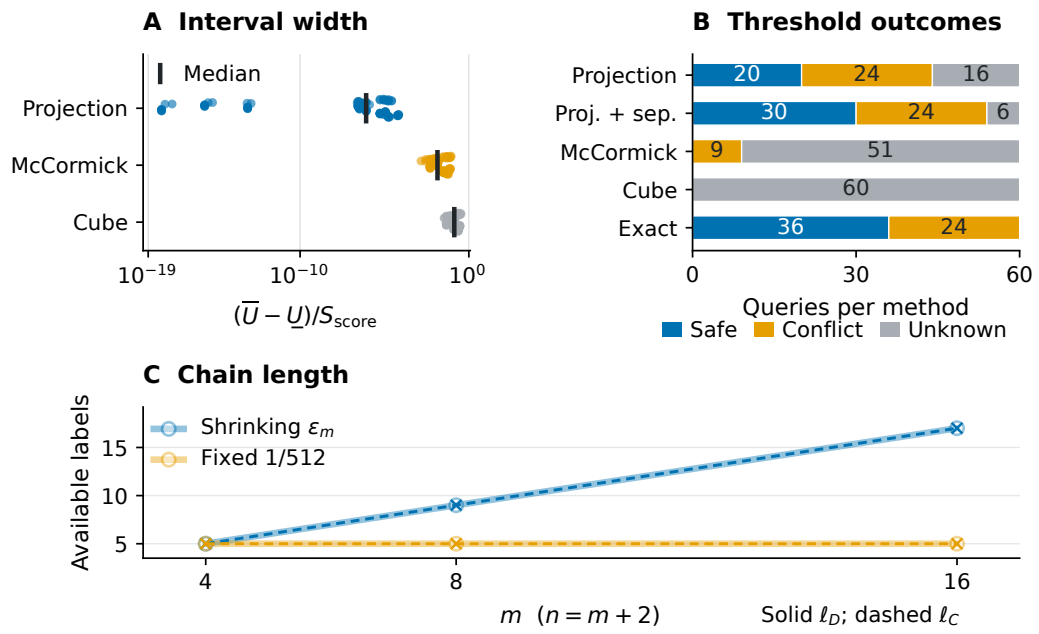

Figure 5: **Certificate quality and available response labels.** (A) Normalized interval widths on all 98 grid instances. Points are instances; black marks are medians. Each method uses its own lower witness and the common nominal completion; all 56 initial and 42 supplementary McCormick evaluations are included. (B) All 60 truth-centered diagnostic queries per method, separate from the nominal-margin threshold groups. "Proj. + sep." adds the single-column upper-bound refinement on the same full domain; the original projection outcomes remain shown. Exact enumeration supplies reference decisions. (C) Response-chain length ℓ_D and strictly native-exposable count ℓ_C for the original family. The solid and dashed curves coincide. Shrinking radius is $\varepsilon_m = 1/(4096m^4)$, fixed radius is $1/512$. These counts describe the candidate-set geometry.

Computation time. We measure the complete cost of producing and checking a certificate. Each method is charged input-geometry and field preparation, nominal optimization, its own construction, exact verification, and serialization. The original input-generation times are reused, while nominal optimization and method execution are timed anew, including the complete McCormick coverage extension. All methods have a common 3600-second full-cost limit, with at most 600 seconds inside the McCormick LP solve, and a 2 GiB process-memory limit. At most two instances run concurrently within each batch.

Across all 98 projection instances, the median charged cost is 2.69 seconds and the maximum 57.02 seconds; the corresponding McCormick costs over 98 instances are 2.35 and 30.17 seconds. These are descriptive costs under the stated resource policy, with both McCormick batches retained. The additional single-column refinement over all twelve small cases takes 0.66 seconds, including its fresh vertex enumerations. Full per-instance wall and CPU times, stages, sampled memory, and shared-cost attribution accompany the supplement.

H.6 Available labels and necessary label use

We distinguish available chain length, labels selected by one optimum, and whether fewer labels can attain the same value. Across all 98 projected solves, 77 select one label, eleven select two, and ten select three to five. The original family at $m = 4, 8, 16$ has chain lengths $5, 9, 17$ at ε_m, but $5, 5, 5$ at fixed $1/512$. The fixed-radius counts are constant across these three sizes. To determine whether

multiple selected labels are necessary for optimality, we compare the unrestricted value with the best value under a bound on the number of distinct projections.

Proposition 21 (At most k projected values)**.** *Under the native model and geometric radius conditions of Theorem 1, let $\mathcal{L}$ be the complete response chain in the common domain $Q_\varepsilon \cap D$. For $1 \le k \le g$, the continuous and discrete optima coincide:*

$$\max_{\substack{s \in Q_\varepsilon^g \\ |\{s_1,\dots,s_g\}| \le k}} J(s) = \max_{\substack{s \in \mathcal{L}^g \\ |\{s_1,\dots,s_g\}| \le k}} J(s) =: V_{\le k}.$$

For each fixed k, this value is computable in polynomial bit time.

Proof. Fix the initial partition of columns into $r \le k$ groups sharing a projected value. Keep this partition throughout. Conditional on the other groups, the objective in one group's value x has the form $f(x) = \alpha\|x\|^2 + \langle t, x\rangle + c$, where $\alpha \ge 0$ sums its internal pair weights and t aggregates its fields and interactions with other groups. Before cone restriction, $x, t \in C_S^*$. For $x \notin D$, the domination point c_0 of Appendix A has $d = c_0 - x \in C$, giving:

$$f(c_0) - f(x) = \langle t, d\rangle + \alpha\big(2\langle x, d\rangle + \|d\|^2\big) \ge 0.$$

Replacing each whole group puts all group values in D; when $Q_\varepsilon \subset D$, this step is unnecessary. The cone D is convex and nonobtuse, and now $p = 2\alpha x + t \in D$. Choose $y \in \mathcal{L}$ maximizing $\langle p, y\rangle$ over $Q_\varepsilon \cap D$, taking any chain point if $p = 0$. The resulting change is nonnegative:

$$f(y) - f(x) = \langle p, y - x\rangle + \alpha\|y - x\|^2 \ge 0.$$

Sequential whole-group replacements preserve the bound of $r \le k$ distinct values and do not decrease the objective. The reverse inequality follows from $\mathcal{L} \subset Q_\varepsilon$. Finally, enumerate subsets of at most k chain labels and solve each by the cut construction. Subsequence increments are sums of the original increments, so their pairwise inner products remain nonnegative. There are $O(|\mathcal{L}|^k)$ such subsets for fixed k, with the existing polynomial encoding bounds. □

The same value also provides an original-domain exclusion certificate. If B has at most k distinct projections, then $J(B) \le V_{\le k} + nW\varepsilon^2$. Thus every original-domain maximizer has more than k distinct projections whenever a feasible witness satisfies:

$$J(B^{\mathrm{lift}}) > V_{\le k} + nW\varepsilon^2$$

Four prespecified $m = 4, g = 4$ checks enumerate all $5^4 = 625$ chain assignments and reproduce the cut optimum. In the two prespecified $m = 8, g = 16$ fixed-radius checks, one label already attains the optimum. Both have $V - V_{\le 2} = 0$, where $V_{\le 2}$ restricts assignments to at most two labels from the complete response chain.

After observing the selected labels, we separately evaluate this restriction for *all ten* cases selecting at least three labels. This post-hoc diagnostic includes every such case. Every gap is strictly positive; $(V - V_{\le 2})/S_{\mathrm{score}}$ ranges from 2.95×10^{-5} to 2.35×10^{-3}. Two $g = 4$ cases also admit independent exhaustive checks of all 625 assignments, including all 145 at-most-two-label assignments. For the other 88 cases, the saved one- or two-label optimum directly witnesses $V_{\le 2} = V$.

To test whether the selected labels reflect joint interactions, vary the neighboring labels and compare a node's conditional scores. For a node assigned label k, the exact preference range against another label ℓ is:

$$A_i^{\pm}(k, \ell) = \langle h_i, x^k - x^\ell\rangle + d_i \min_r / \max \langle x^k - x^\ell, x^r\rangle.$$

In each of the ten cases, every node has an alternative with $A_i^- < 0 < A_i^+$, so its conditional preference depends on its neighbors. By Proposition 21, every continuous projected maximizer in these ten cases has at least three distinct values. In eight of them, the saved feasible original-domain witness also exceeds $V_{\le 2} + nW\varepsilon^2$, certifying at least three distinct projections in every original-domain maximizer. All ten restricted values and original-domain comparison margins are supplied with the exact checks.

H.7 A COUPLED EXAMPLE WITH A UNIQUE ORIGINAL-DOMAIN OPTIMIZER

This example shows that coupling can require distinct projected values even when the original-domain optimizer is unique. A three-node graph with six-dimensional representations gives an explicit construction. We establish its projected optimum, compare it with independent column choices, and then compute the original-domain maximum.

Consider the full residual tube at $\varepsilon = 1/512$ for

$$R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 11/10 & 9/10 \\ 21/20 & 1 \\ 1 & 21/20 \\ 9/10 & 11/10 \end{pmatrix}, \qquad S = \mathrm{col}(R), \qquad G = R^\top R = \frac{1}{400}\begin{pmatrix} 2049 & 1632 \\ 1632 & 2049 \end{pmatrix}.$$

Its normalized nonnegative extreme rays are $u = (10/11)Re_1$ and $v = (10/11)Re_2$. The cone construction gives $q = 68/683$ and $\varepsilon_B = 680/28003 > \varepsilon$. Three unknown columns form a triangle with $w_{12} = w_{13} = w_{23} = 1/3$. Thus $W = 1$ and $d_i = 2/3$. Taking $\kappa = 2$ and $(\lambda_1, \lambda_2, \lambda_3) = (0, 1/2, 1)$ in the degree-scaled anchor construction gives:

$$h_1 = \frac{2}{3}u, \qquad h_2 = \frac{1}{3}(u + v), \qquad h_3 = \frac{2}{3}v.$$

In particular, $\|h_i\|_\infty \le d_i$. Every unknown column varies independently over the same complete K_ε.

Projected candidates and optimum. The full retained chain $\mathcal{L}$ consists of $x^k = Ra^k$, where

$$\begin{aligned} a^0 &= (5139/5632, -1/512), & a^1 &= (2565/3968, 2565/7936), \\ a^2 &= (2565/5248, 2565/5248), & a^3 &= (2565/7936, 2565/3968), \\ a^4 &= (-1/512, 5139/5632). \end{aligned}$$

Explicit feasible lifts are $b^k = x^k + \varepsilon z^k$, with:

$$\begin{aligned} z^0 = z^4 &= (1, 1, -1, 20/41, 20/41, -1), \\ z^1 &= (1, 43/80, -1, -1, 1/4, 1), \\ z^2 &= (1, 1, 21/40, -1, -1, 21/40), \\ z^3 &= (43/80, 1, 1, 1/4, -1, -1). \end{aligned}$$

Direct substitution gives $R^\top z^k = 0$, $\|z^k\|_\infty = 1$, and $b^k \in [0, 1]^6$.

To establish completeness of the chain, every projected coordinate a lies in the outer hexagon

$$\mathcal{H} = \{a : a_1, a_2 \ge -\varepsilon,\ R_j a \le 1 + \varepsilon \text{ for } j = 3, 4, 5, 6\}.$$

Its vertices are $(-\varepsilon, -\varepsilon)$ and the five a^k above. Every nonzero D-field has two strictly positive objective coefficients in these coordinates. The five feasible vertices and their connecting edges therefore attain all the relevant upper supports of $\mathcal{H}$. The directions u, $RG^{-1}(R_3 + R_4)^\top$, $RG^{-1}(R_4 + R_5)^\top$, $RG^{-1}(R_5 + R_6)^\top$, and v strictly expose the respective five vertices. The three middle directions have positive R-basis coefficients, so all five fields are native. For the present coupled objective, the coefficients of every conditional objective are also strictly positive: $R^\top h_i > 0$ and $R^\top P b_j = R^\top b_j \ge 0$. Successive conditional replacements consequently preserve a continuous projected optimum on this chain.

Exact enumeration of all 125 chain configurations gives the unique chain optimum $(1, 2, 3)$, with:

$$V_\varepsilon = \frac{65429319429}{5164367872}.$$

Enumerating all 65 configurations using at most two labels from this same full chain gives

$$V_{\le 2}^{\mathcal{L}} = \frac{1898007635889}{150266445824}, \qquad V_\varepsilon - V_{\le 2}^{\mathcal{L}} = \frac{179009412399}{4658259820544} > 0.$$

Each selected label occurs once, a fraction $1/3$ of the columns. The same two-label bound holds on the continuous projected domain by Proposition 21. Continuous projected uniqueness can also be checked directly: the $6^3 = 216$ vertex configurations of $\mathcal{H}^3$ have the same unique maximizing tuple. Separate affinity expresses every point in $\mathcal{H}^3$ as a product convex combination of these configurations. Equality forces concentration on that single tuple. Since its three points are feasible in Q_ε, it is the unique maximizer over Q_ε^3 as well.

How coupling changes the choices. The distinct optimal labels alone do not identify the role of the pairwise terms. We therefore compare conditional preferences and the optimum obtained by using only the linear fields. For a selected label k and alternative l, consider:

$$A_i^{\pm}(k,l) = \langle h_i, x^k - x^l \rangle + \frac{2}{3} \min_r \max \langle x^k - x^l, x^r \rangle.$$

The alternatives $(k,l) = (1,0), (2,1), (3,4)$ satisfy $A_i^- < 0 < A_i^+$ for nodes $1,2,3$, respectively. Thus every node is interaction-sensitive and none is uniformly forced by its field. Maximizing the linear fields independently gives the unique labels $(0,2,4)$; coupling changes two choices and increases the complete objective by

$$J(x^1, x^2, x^3) - J(x^0, x^2, x^4) = \frac{107752892484003}{249955405004800} > 0.$$

The original-domain optimum. The 24 cube and residual inequalities defining K_ε have exactly 380 vertices. Exact enumeration of all $\binom{24}{6} = 134596$ six-constraint sets establishes this vertex list. Evaluation of the $380^3 = 54872000$ vertex configurations gives the unique maximizing tuple (b^1, b^2, b^3) specified above. Separate affinity makes this vertex optimum the maximum over the original continuous domain. Its value is:

$$U_\varepsilon = J(b^1, b^2, b^3) = V_\varepsilon + \frac{21}{20}\varepsilon^2 = \frac{2617173604581}{206574714880}.$$

The residual contribution is $\frac{1}{3}\sum_{i<j}\langle z^i, z^j \rangle \varepsilon^2 = \frac{21}{20}\varepsilon^2$. Every original feasible tuple is a product convex combination of vertex tuples, and J is separately affine. A unique maximizing vertex tuple therefore implies a unique continuous original-domain optimizer. Its projections are precisely (x^1, x^2, x^3) and are pairwise distinct. The vertex enumeration, complete product search, and separate exact verification accompany the construction.

I Linear fields and fixed references with residual components

Observed representations may have components outside the shared subspace. Those components create additional linear error terms in the score. We first extend the projected certificate to such fields, then derive disclosure bounds for fixed references with residuals and identify a case that can be solved exactly.

Retain the rational rank-two space S, projector P, cone $C = S \cap \mathbb{R}^n_+$ spanning S, and domains $K_\varepsilon, Q_\varepsilon$ of (1). Fix the nonobtuse cone D and radius $\varepsilon_B(S)$ constructed in (11)–(12); for orthogonal extreme rays use $D = C$ and $\varepsilon_B = 1$. All data are rational, $0 \le \varepsilon \le \varepsilon_B(S)$, and pair weights are nonnegative with no self-interactions. The cone D is kept fixed even at $\varepsilon = 0$. Write $w_{ji} = w_{ij}$ and $w_{ii} = 0$ when summing over ordered indices.

Proposition 22 (Linear fields with residual components)**.** *Let $h_i \in \mathbb{Q}^n_+$ satisfy the additional condition $\bar{h}_i = Ph_i \in D$. Put $h_i^\perp = (I-P)h_i$, $\eta = \sum_i \|h_i^\perp\|_1$, and $W = \sum_{i<j} w_{ij}$. Define:*

$$J(B) = \sum_{i<j} w_{ij}\langle b_i, b_j \rangle + \sum_i \langle h_i, b_i \rangle, \qquad U = \max_{K_\varepsilon^g} J,$$

$$F(s) = \sum_{i<j} w_{ij}\langle s_i, s_j \rangle + \sum_i \langle \bar{h}_i, s_i \rangle, \qquad V = \max_{Q_\varepsilon^g} F.$$

One can compute V, a projected optimizer $s^\star$, and a rational feasible lift B^{lift} with $Pb_i^{\mathrm{lift}} = s_i^\star$ in polynomial bit time. With $a_\varepsilon = nW\varepsilon^2 + \varepsilon\eta$,

$$J(B^{\mathrm{lift}}) \le U \le V + a_\varepsilon, \qquad V + a_\varepsilon - J(B^{\mathrm{lift}}) \le 2a_\varepsilon, \qquad |U - V| \le a_\varepsilon. \tag{57}$$

The bounds hold for any feasible lift of any projected global optimizer.

Proof. The projection in Lemma 4 is independent of the fields; at zero radius use $Q_0 = K_0$. Since D is nonobtuse and $C \subseteq D$, $D \subseteq C_S^*$. Before restriction, (10) therefore gives $t_i = \bar{h}_i + \sum_{j \neq i} w_{ij} s_j \in C_S^*$. For positive radius, the outside-point argument in Appendix A gives $\langle t, c_0 - s, > \rangle 0$ whenever $t \in C_S^* \setminus \{0\}$ and $s \in Q_\varepsilon \setminus D$. At a global optimum an outside column thus has zero conditional field; replacing it by c_0 preserves its value. At most g replacements yield an optimum on $(Q_\varepsilon \cap D)^g$. At zero radius, or with orthogonal extreme rays, the inclusion in D already holds.

After restriction all conditional fields belong to D. The response-chain argument and increment sign (28) in Appendix A therefore apply with fields $\bar{h}_i$. Pareto maximality is relative to the whole response polygon, retaining the maximal endpoints of horizontal or vertical faces. In the binary expansion, extend $\beta_{ik,j\ell} = w_{ij}\langle d_i^k, d_j^\ell \rangle$ for $i < j$ by $\beta_{j\ell,ik} = \beta_{ik,j\ell}$, and set $\beta_{ik,i\ell} = 0$, including the diagonal. Use its exact cut construction and rational lifting through (7). At $\varepsilon = 0$, retain the response coordinates of the frozen D, including when $\bar{h}_i \notin C$. The existing polynomial bit bounds remain valid, since computing Ph_i and the new unary coefficients requires only polynomially many rational operations. The chain replacement also covers zero fields and tied maximizers.

For each original feasible completion, write $b_i = Pb_i + e_i$. Orthogonality and the tube constraints give:

$$J(B) - F(PB) = \sum_{i<j} w_{ij}\langle e_i, e_j \rangle + \sum_i \langle h_i^\perp, e_i \rangle, \qquad |J(B) - F(PB)| \leq nW\varepsilon^2 + \varepsilon\eta.$$

Project an original optimizer for the upper bound and lift a projected optimizer for the lower bound. Feasibility gives $J(B^{\text{lift}}) \leq U$, proving (57). □

Fixed original references. When the revealed columns have residual components, their interactions with unknown columns contribute to the enclosure error. Their mutual interactions remain a known constant. Keeping these two contributions separate gives the following disclosure certificate. Let $\mathcal{U} = \{1, \ldots, N\}$ index all columns, including any permanently known columns. Fix $A_i \in K_\varepsilon$ with the separate premise $s_i = PA_i \in D$, and write $r_i = (I - P)A_i$. For $E \subseteq \mathcal{U}$, let $\mathcal{F}(E)$ allow $B_i \in K_\varepsilon$ for $i \in E$ and fix $B_j = A_j$ for $j \notin E$. Define the full original score and its extrema:

$$R(B) = \sum_{i<j} w_{ij}\langle B_i, B_j \rangle, \qquad r_A = R(A),$$

$$c(E) = \sum_{\substack{i<j \\ i,j \notin E}} w_{ij}\langle A_i, A_j \rangle, \qquad M(E) = \max_{B \in \mathcal{F}(E)} R(B).$$

Set $f(E) = \max\{r_A - c(E), M(E) - r_A\}$. The complete projected value and residual terms are:

$$V(E) = \max_{\substack{s_i' \in Q_\varepsilon \ (i \in E) \\ s_j' = s_j \ (j \notin E)}} \sum_{i<j} w_{ij}\langle s_i', s_j' \rangle,$$

$$C_\perp(E) = \sum_{\substack{i<j \\ i,j \notin E}} w_{ij}\langle r_i, r_j \rangle,$$

$$\sigma(E) = n\varepsilon^2 \sum_{\substack{i<j \\ i,j \in E}} w_{ij} + \varepsilon \sum_{\substack{i \in E \\ j \notin E}} w_{ij}\|r_j\|_1, \tag{58}$$

$$\overline{M}(E) = V(E) + C_\perp(E) + \sigma(E), \qquad \overline{f}(E) = \max\{r_A - c(E), \overline{M}(E) - r_A\}. \tag{59}$$

The mixed sum counts each unknown–known edge once, and $C_\perp(E)$ retains its sign.

Proposition 23 (Fixed references outside the subspace)**.** *For each E, $V(E)$ and a feasible completion B_E lifting a projected optimizer are computable in polynomial bit time. Set $L(E) = R(B_E)$ and $\ell(E) = \max\{r_A - c(E), L(E) - r_A\}$. Then*

$$L(E) \leq M(E) \leq \overline{M}(E), \qquad \overline{M}(E) - L(E) \leq 2\sigma(E),$$

$$\ell(E) \leq f(E) \leq \overline{f}(E), \qquad \overline{f}(E) - \ell(E) \leq 2\sigma(E). \tag{60}$$

Thus $f(E) \leq \overline{f}(E) \leq f(E) + 2\sigma(E)$. The functions $\overline{M}$, $\overline{f}$, and σ are nondecreasing under inclusion, and $\overline{M}(\varnothing) = r_A$, $\ell(\varnothing) = f(\varnothing) = \overline{f}(\varnothing) = 0$.

Proof. The fields for unknown columns are $h_i(E) = \sum_{j \notin E} w_{ij} A_j \geq 0$, with $Ph_i(E) = \sum_{j \notin E} w_{ij} s_j \in D$. Proposition 22 computes their variable projected maximum. Adding the projected known–known constant $c_S(E)$ gives $V(E)$; the identity $c(E) = c_S(E) + C_\perp(E)$ converts between the two constants. Lift only the unknown columns and retain A_j elsewhere.

For any such completion, writing $e_i = (I - P)B_i$ for $i \in E$ gives:

$$R(B) - R(PB) - C_\perp(E) = \sum_{\substack{i<j \\ i,j \in E}} w_{ij} \langle e_i, e_j \rangle + \sum_{\substack{i \in E \\ j \notin E}} w_{ij} \langle e_i, r_j \rangle.$$

Its absolute value is at most $\sigma(E)$, proving the first line of (60). Zeroing unknown columns gives the exact minimum $c(E)$, and $A \in \mathcal{F}(E)$; hence f is the maximum absolute deviation from r_A. Applying the nondecreasing, 1-Lipschitz map $x \mapsto \max\{r_A - c(E), x - r_A\}$ proves the second line.

For monotonicity, the following per-edge contributions omit the common nonnegative weight:

Endpoints	$C_\perp(E) + \sigma(E)$	$\sigma(E)$
Both known	$\langle r_i, r_j \rangle$	0
i unknown, j known	$\varepsilon \|r_j\|_1$	$\varepsilon \|r_j\|_1$
Both unknown	$n\varepsilon^2$	$n\varepsilon^2$

Each column of contributions is nondecreasing as endpoints become unknown: $\langle r_i, r_j \rangle \leq \varepsilon \|r_j\|_1 \leq n\varepsilon^2$ and $0 \leq \varepsilon \|r_j\|_1$. This includes negative residual inner products. The projected domains are nested since $s_j \in Q_\varepsilon$, so V is nondecreasing. Thus $\overline{M}$ is nondecreasing; so is $\overline{f}$, because c is nonincreasing. At the empty set, $V = c_S$, $\sigma = 0$, and $c_S + C_\perp = r_A$, giving the stated identities. □

Queries within an initial unknown set. During deletion, the unknown set shrinks while the reference values remain fixed. One error budget therefore applies to every query: for fixed $E_0 \subseteq \mathcal{U}$ and every $E \subseteq E_0$, we have:

$$\overline{f}(E) - \ell(E) \leq 2\sigma(E) \leq \Delta_{E_0} := 2\sigma(E_0),$$
$$\Delta_{E_0} = 2n\varepsilon^2 \sum_{\substack{i<j \\ i,j \in E_0}} w_{ij} + 2\varepsilon \sum_{\substack{i \in E_0 \\ j \notin E_0}} w_{ij} \|r_j\|_1.$$

The second term includes permanently known columns outside E_0. With the reference, space, radius and weights fixed, $\theta \geq 0$, and nonnegative disclosure costs, the standard arguments of Proposition 15, Corollary 16, and Proposition 17 apply on E_0 with upper predicate $\overline{f}(E) > \theta$ and tolerance shift Δ_{E_0}; deletion starts from a set above threshold, and a true conflict still requires a feasible lower witness. When $E_0 = \mathcal{U}$, the mixed budget vanishes.

Exact queries without unknown–unknown edges. If $W(E) = \sum_{i<j,\, i,j \in E} w_{ij} = 0$, nonnegative weights give $w_{ij} = 0$ for all $i, j \in E$. The original objective then separates across unknown columns:

$$M(E) = c(E) + \sum_{i \in E} \max_{b \in K_\varepsilon} \langle h_i(E), b \rangle, \qquad h_i(E) = \sum_{j \notin E} w_{ij} A_j.$$

Each maximum is a rational linear program on the original tube, so $M(E)$ and $f(E) = \max\{r_A - c(E), M(E) - r_A\}$ are exactly computable in polynomial bit time. This case retains fixed nonnegative references $A_j \in K_\varepsilon$ and requires no condition $PA_j \in D$ on the known columns. Combining the individual maximizers gives a feasible completion attaining $M(E)$, while zeroing the unknown columns attains $c(E)$. Using these exact values removes the slack that the mixed-edge budget $\sigma(E)$ can introduce for such queries.

Recovery of the native model. The projected certificates require cone membership $Ph_i \in D$ or $PA_i \in D$ in addition to nonnegativity and the residual budget. Native fields recover $a_\varepsilon = nW\varepsilon^2$, and references in K_0 recover $C_\perp = 0$, $\sigma = \delta_E$, and the original upper bound $u(E)$. Zero radius gives exact certificates with the same fixed D. Empty sums are zero, covering $g = 0$ and reference problems with no edges. For nonnative references, mixed edges can contribute to $\sigma(E)$ even when $W(E) = 0$; the independent linear programs above compute such queries exactly.

Radius and first-order accuracy. The size of the reference residual determines the radii for which the fixed reference remains feasible. For fixed nonnative references, the projected certificate uses the range:

$$\max_i \|(I - P)A_i\|_\infty \le \varepsilon \le \varepsilon_B(S).$$

The reference values remain fixed throughout this range. For fixed $\eta > 0$, field residuals contribute the linear term $\varepsilon\eta$ to the error budget. Dividing the bound by ε gives $nW\varepsilon+\eta$; exact native first-order recovery uses the case $\eta = 0$.