SUPPORT SELECTION BEYOND SMOOTH DAG EXACTNESS: COMPLETION GEOMETRY, SCORE MARGINS, AND SELECTIVE CERTIFICATES

Anonymous authors
Paper under double-blind review

ABSTRACT

Continuous DAG learning often relies on a smooth exact constraint to enforce acyclicity. Exactness specifies the zero set, but leaves two questions unanswered. How informative is a small nonzero residual? Which edges should be removed when several repairs restore acyclicity? For the first question, we prove sharp local limits for sufficiently smooth exact representations on signed candidate neighborhoods of a fixed DAG. The bounds reveal that the smallest minimal cycle completion controls the earliest possible response, while the largest controls the best uniform distance-error exponent. They also distinguish signed or vector residuals from nonnegative scalars. For the second question, we show how a score margin breaks a feasibility tie in an isolated-cycle flow. We then use simultaneous parent-set score intervals to account for sampling uncertainty and certify features shared by all population optima of a frozen score, without requiring a unique graph. Controlled experiments recover the predicted orders and verify certification under nonunique optima, while comparisons with standard optimizers show the flow model's limitations. The results clarify what changing a smooth constraint can improve and what requires evidence from scores and data.

1 INTRODUCTION

Smooth exact constraints made continuous optimization possible for weighted DAG learning (Zheng et al., 2018). Trace polynomials, DAGMA's log determinant, and spectral or analytic alternatives provide different optimization landscapes (Yu et al., 2019; Bello et al., 2022; Nazaret et al., 2024; Zhang et al., 2025). Their shared guarantee is that the residual vanishes precisely on acyclic supports. This does not yet explain how informative a small nonzero residual is, or how to choose between repairs that restore acyclicity.

Consider a path $1 \to 2 \to 3$ and candidate edges $3 \to 4$ and $4 \to 1$, with weights x and y. Either edge alone preserves acyclicity; together they close a cycle. When both weights are small, a smooth residual can be tiny while the graph remains cyclic. Deleting either candidate breaks the cycle, but leaves a different dependency structure. Detecting the cycle and comparing the repairs are separate tasks. These roles appear separately in the usual objective,

$$\min_{\mathbf{W}} \ \mathcal{L}(\mathbf{W}; X) + \lambda \mathcal{R}(\mathbf{W}) + \Psi(h(\mathbf{W})). \tag{1}$$

where $\mathcal{L}$ scores the fit to data, $\mathcal{R}$ regularizes the weights, and Ψ penalizes the acyclicity residual h. We ask what residual size can tell us about distance to the acyclic set, and what score evidence can support a particular repair. Figure 1 follows these two questions from local sensitivity to support selection and its uncertainty.

Our analysis fixes a DAG and allows selected absent edges to vary with either sign. For the candidates $3 \to 4$ and $4 \to 1$, the residual is zero on both $x = 0$ and $y = 0$. Every nonzero term in its Taylor expansion must therefore involve both weights. More generally, a *cycle completion* is a set of candidates that creates a cycle when added to the base DAG. It is minimal if no proper subset does so. An exact residual vanishes on every coordinate subspace whose active candidates preserve acyclicity. Its Taylor expansion can therefore contain only products whose active candidates include a minimal completion. Using standard combinatorial and algebraic tools (Hultman, 2004; Miller &

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.

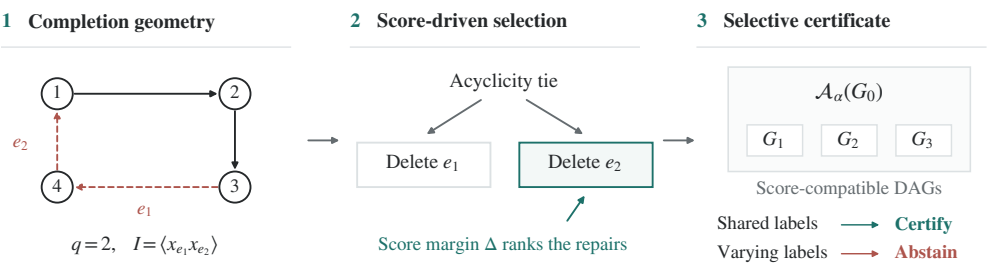


Figure 1: From feasibility to a support decision. Completion geometry identifies competing repairs and constrains the local response of a smooth exact residual. The score ranks repairs. Finite-sample intervals retain the score-compatible DAGs and certify only the labels they share.

Sturmfels, 2005; Reznick, 1978), we turn these simultaneous zero conditions into local bounds that depend on the candidate support rather than the chosen formula.

These bounds distinguish two properties of the completion structure. The smallest minimal completion controls the earliest possible local response, while the largest controls the best uniform distance-error exponent (Theorems 2.1 and 2.2). We prove sharp bounds for signed or vector residuals and nonnegative scalar residuals, and construct exact representations that attain them. A constraint can therefore respond at a low order in some directions and still give weak distance control across the candidate section. Controlled experiments on constructed and randomly generated sections recover the predicted orders. The bounds apply to sufficiently smooth exact maps on signed candidate neighborhoods, not to every way of enforcing acyclicity.

The score and regularizer can distinguish deleting $3 \to 4$ from deleting $4 \to 1$ even when the feasibility gradient treats the edges alike. GOLEM and later work study these roles in the full objective (Ng et al., 2020; Deng et al., 2024; Jin et al., 2026). To isolate this distinction, Section 3 studies an isolated-cycle gradient flow with two nearly equal weights. Under the feasibility penalty alone, separation time diverges as the initial imbalance vanishes. A fixed positive pathwise margin bounds this time under the stated continuation assumptions. Controlled experiments recover the predicted flow law; official NOTEARS and DAGMA trajectories test what the separation diagnostics describe under different update rules.

An empirical score can favor one repair even when the data do not resolve the population ordering. To account for this uncertainty, we use set-valued model comparison (Hansen et al., 2011; Park et al., 2026). For a frozen candidate family and score, simultaneous intervals from an independent holdout retain all population-optimal DAGs at the stated confidence level. Features shared by this set can be certified without a unique optimum (Theorem 4.1). The queries cover adjacencies, directed edges, and unshielded colliders. Exact-population calibration checks certification under nonunique optima. On learned graphs, certified labels are more stable than uncertified labels under the tested score perturbations. Interpreting these labels causally requires additional assumptions connecting the population score optima to the data-generating graph.

1.1 Related Work

Due to space constraints, we discuss the closest work here; Appendix D.6 gives a fuller comparison of assumptions and guarantees.

Constraints and degeneracy. NOTEARS and DAGMA provide smooth exact acyclicity constraints (Zheng et al., 2018; Bello et al., 2022). No Fears identifies gradient and KKT degeneracy (Wei et al., 2020); Ng et al. (2022) study the resulting convergence issues. Higher-order local search uses Taylor derivatives for a specified constraint family (Shridharan & Iyengar, 2025). Our response and distance bounds instead depend on minimal cycle completions and apply across sufficiently smooth exact representations on signed candidate neighborhoods.

Support selection. GOLEM and subsequent analysis examine how scores, sparsity, and thresholding affect the selected graph (Ng et al., 2020; 2024). Direct acyclic parameterizations, continuous–

discrete hybrids, and nonnegative domains use other assumptions (Yu et al., 2021; Massidda et al., 2024; Gillot & Parviainen, 2022; Rey et al., 2026). Our cycle-flow analysis isolates when a score preference separates repairs that the feasibility term treats alike.

Uncertainty in model selection. Model confidence sets and split-sample relative-fit tests retain multiple models when data do not resolve their ranking (Hansen et al., 2011; Park et al., 2026). Wang et al. (2026) give asymptotically valid ordering confidence sets for identifiable additive-error structural equation models. We target a frozen score and candidate family: simultaneous parent-set intervals yield a robust DAG objective and finite-sample certificates for features shared by all population score optima. Parent-set integer programming is an established computational tool (Cussens, 2011; 2020).

2 How Cycle Completions Limit Local Sensitivity

For the path $1 \to 2 \to 3$, neither candidate $3 \to 4$ nor $4 \to 1$ alone creates a cycle. Any Taylor response must involve both weights. We now derive the corresponding restrictions on local sensitivity and distance from the acyclic set.

Candidate coordinates. Fix a zero-diagonal weighted directed acyclic graph (DAG) $\mathbf{W} \in \mathbb{R}^{d\times d}$, where W_{ij} weights the edge $i \to j$, and absent off-diagonal edges $F = \{e_1, \dots, e_m\}$. The coordinate matrix $\mathbf{E}_s$ has a single one at e_s and zeros elsewhere. For candidate weights $x \in \mathbb{R}^m$, write $G_F(x) = G(\mathbf{W} + \sum_s x_s \mathbf{E}_s)$: the base weights stay fixed and only the candidate edges vary. The domain is an open neighborhood of $x = 0$, allowing either sign. Here G denotes either a finite-dimensional signed/vector residual H or a nonnegative scalar residual h. Both are *exact*: they vanish precisely on DAG supports in this domain. Derivatives below act on the candidate coordinates; vector norms and distances are Euclidean, and Jacobian norms are induced operator norms.

A completion is a candidate subset that forms a cycle with the base DAG. It is *minimal* if removing any candidate removes all cycles: inclusion-minimality, not minimum size. Denote their collection by $\mathcal{C}_{\min}$. Let $q = q_{\mathbf{W}}(F)$ be the fewest candidates needed to form a cycle, so $q = \min_{C\in\mathcal{C}_{\min}} |C|$, with $\min \varnothing = +\infty$, and set $x^C = \prod_{s\in C} x_s$. In the four-node example, $q = 2$, although the cycle has length four: we count missing edges, not the full cycle. At the empty graph, q is the shortest cycle length in F. If F includes every missing directed edge of a nonempty DAG, then $q = 1$, since a reverse edge closes a two-cycle. Higher values concern restricted candidate groups.

Every coordinate subspace with fewer than q active candidates is acyclic. The next theorem turns these zero conditions into a limit on how early a smooth residual can respond. Write C^k for continuous derivatives through order k, and D for the derivative. For a multi-index $\beta \in \{0, 1, \dots\}^m$, set $|\beta| = \sum_s \beta_s$, $x^\beta = \prod_s x_s^{\beta_s}$, $D^\beta = \partial_1^{\beta_1} \cdots \partial_m^{\beta_m}$, and $\operatorname{supp}(\beta) = \{s : \beta_s > 0\}$.

Theorem 2.1 (Completion structure and sharp local response orders)**.** *Fix $(\mathbf{W}, F)$ with $q_{\mathbf{W}}(F) = q < \infty$. For any integer $k \geq 0$ and C^k exact residual G_F, each nonzero coefficient of x^β in a component of its order-k Taylor polynomial satisfies*

$$\exists C \in \mathcal{C}_{\min} \quad C \subseteq \operatorname{supp}(\beta).$$

The polynomial vector $(x^C)_{C\in\mathcal{C}_{\min}}$ is itself exact. If $H_F \in C^q$ is signed/vector exact and $h_F \in C^{2q}$ is nonnegative exact, then, for every fixed $u \in \mathbb{R}^m$, as $t \to 0$,

$$\begin{aligned} \|H_F(tu)\|_2 &= O(|t|^q), & \|DH_F(tu)\|_{\mathrm{op}} &= O(|t|^{q-1}), \\ h_F(tu) &= O(|t|^{2q}), & \|Dh_F(tu)\|_2 &= O(|t|^{2q-1}). \end{aligned} \tag{2}$$

Moreover, $D^\beta H_F(0) = 0$ for $|\beta| < q$, and $D^\beta h_F(0) = 0$ for $|\beta| < 2q$. Both value exponents are attained by exact representations along directions whose support is a smallest completion.

Beyond gradient degeneracy. Earlier analyses establish polynomial-constraint degeneracy and failures of standard ALM regularity (Wei et al., 2020; Ng et al., 2022). The monomial condition in Theorem 2.1 identifies which derivatives must vanish for *every* sufficiently smooth exact representation on the same signed section. The two rows of Eq. (2) separate the order imposed by completion size from the additional order imposed by nonnegativity. Changing a smooth nonnegative formula cannot improve on $2q$; signed or vector outputs can attain q, but no earlier order.

The proof uses two facts. Each surviving monomial contains a minimal completion. A nonnegative scalar's first nonzero homogeneous Taylor term, if one exists, is nonnegative. Its Newton polytope, the convex hull of its exponent vectors, has even vertex exponents (Lemma A.7): at least q coordinates have exponent at least two. Appendix A develops the *cycle-completion ideal*, unequal-rate paths, factorization, and weaker regularity conditions (Theorems A.5 and A.8).

Residual size and distance. A stopping tolerance needs a different guarantee: how much must the candidate weights change to remove every cycle? Let $\mathcal{D}_F$ be the candidate vectors whose support, together with the base DAG, is acyclic, and let $r = r_{\mathbf{W}}(F) = \max_{C \in \mathcal{C}_{\min}} |C|$. The smallest size q determines the earliest possible response. A distance bound must also cover the least sensitive directions, which depend on the largest size r.

For example, the minimal products x_1 and $x_2x_3x_4$ give $h(x) = x_1^2 + x_2^2x_3^2x_4^2$, with $q = 1$ but $r = 3$. Along $(0, t, t, t)$, the residual is t^6 while the distance is $|t|$. A residual of this size does not mean that a change of the same size will remove the cycle. Such a section is realized by disjoint base paths, one closed by a single candidate and another closed by three candidates. The exponent θ below quantifies this conversion from residual to distance. A larger exponent gives a tighter bound as the residual approaches zero; the example rules out any exponent above $1/6$. The sharp vector below has one component per simple directed cycle, equal to the product of its edge weights.

Theorem 2.2 (Sharp topological error-bound exponents)**.** *Suppose $q_{\mathbf{W}}(F) < \infty$ and let $r = r_{\mathbf{W}}(F)$. If an exact residual G_F admits constants $K, \rho, \theta > 0$, with the ball $B_\rho(0)$ contained in its domain, such that*

$$\operatorname{dist}(x, \mathcal{D}_F) \le K\|G_F(x)\|^\theta \qquad \text{for every } \|x\|_2 < \rho,$$

then $\theta \le 1/r$ when $G_F = H_F \in C^r$ is signed/vector exact, and $\theta \le 1/(2r)$ when $G_F = h_F \in C^{2r}$ is nonnegative exact. Both bounds are attained: the simple-cycle product vector has optimal exponent $1/r$, while the positive walk constraints in Corollary A.10 have optimal exponent $1/(2r)$. Consequently, squaring these sharp residuals gives optimal exponents $1/(2r)$ and $1/(4r)$.

Consequences for stopping rules. The upper bounds on θ in Theorem 2.2 rule out a tighter power-law conversion from residual to distance on the whole section. Its attainment statement shows that these limits are not a loose estimate. Unlike the response order q, this conversion depends on the largest minimal completion r. In the Hölder error-bound framework (Li & Mordukhovich, 2012), necessity follows from that completion; zeroing a smallest-magnitude coordinate in each completion proves attainment (Proposition C.4). An arbitrary smooth exact residual need not admit a power-law bound.

Multiplying a residual by a positive constant preserves its zero set and response order but changes the constant in a distance bound. For a nonnegative residual attaining the sharp exponent, $h_F(x) \le \epsilon$ yields $\operatorname{dist}(x, \mathcal{D}_F) \le K\epsilon^{1/(2r)}$ in the neighborhood where the bound holds. A stopping rule based on this exponent therefore also needs control of K and that neighborhood.

Implications for NOTEARS and DAGMA. For NOTEARS,

$$h_{\exp}(\mathbf{W}) = \operatorname{tr}(\exp(\mathbf{W} \circ \mathbf{W})) - d. \tag{3}$$

Here $\circ$ is entrywise multiplication, $\exp$ the matrix exponential, and tr the trace; we abbreviate this residual as EXP. A cycle contributes the product of its squared weights to the positive walk expansion. A minimal q-edge completion thus gives a term of order $2q$. The same argument applies to positive trace polynomials and to DAGMA's convergent log-determinant expansion inside its M-matrix domain. On a minimal completion with nonzero equal-scale candidate coefficients, all three satisfy

$$h_F(tu) = \Theta(|t|^{2q}), \qquad \|Dh_F(tu)\|_2 = \Theta(|t|^{2q-1}). \tag{4}$$

Corollary A.10 states the common positive-walk conditions. DAGMA can improve coefficient sizes and conditioning away from the boundary while retaining this local exponent.

Learning also depends on the score, regularizer, and support operations. Absolute-value reformulations and support-based local search use information beyond the residual's response (Wei et al., 2020). Sparsity and post-hoc thresholding also affect the selected graph (Ng et al., 2024).

Local-response diagnostics. The settings in Table 4 test two consequences of Theorem 2.1. On prescribed completions with $q = 1, \ldots, 5$, changing the representation or squaring its residual gives the predicted $q, 2q, 4q$ hierarchy; value and derivative slope errors are at most 0.046. Random sections test whether completion size also predicts response outside these prescribed paths. Across 80 instances (20 seeds per $q \in \{1, 2, 3, 4\}$), the maximum error over 160 instance–method fits is 0.0033 (Appendix F). These checks distinguish changes in response order from changes in coefficients. They do not rank the feasible repairs, which is the next question.

3 How a Score Breaks Feasibility Ties

Deleting $3 \to 4$ or $4 \to 1$ breaks the four-node cycle, but leaves different supports. Reducing the residual need not distinguish these repairs. We first study how a score breaks such a tie in a tractable flow, then test what carries over to standard optimizers.

For relabeling-invariant constraints built from squared weights, exchangeable edge pairs have equal feasibility gradients on a tied manifold. Disjoint blocks give exponentially many such manifolds at nonzero weights (Theorem D.1). Near such a manifold, a locally Lipschitz gradient has a tangential component at most proportional to the distance from it (Corollary D.2).

Separation under the feasibility term. Consider the continuous flow $\dot{z} = -\nabla_z[\Psi(h(W(z)))]$ on an isolated directed cycle. The matrix $W(z)$ places the positive entries of z on the cycle and is zero elsewhere. Every cycle edge evolves. Two tracked weights start at $(x_0, y_0) = (a + \varepsilon, a - \varepsilon)$, with fixed $a > 0$ and small $\varepsilon > 0$; the others start at fixed $b_j > a$. Let $T_0(\varepsilon)$ be the first time $y/x = \zeta$, for fixed $\zeta \in (0, 1)$. This measures separation, not arrival at zero.

Proposition 3.1 (Separation under a feasibility penalty)**.** *Under this initialization, suppose $h(W(z)) = \phi(p)$, where $p = \prod_i z_i^2$, $\phi(0) = 0$, and $\phi'(0) > 0$. Let ϕ and Ψ be C^1 on their relevant nonnegative ranges, including zero, with $\phi' > 0$ and $\Psi'(s) > 0$ for $s > 0$. If $\Psi'(s) \sim cs^\nu$ as $s \downarrow 0$, for $c > 0$ and $\nu \geq 0$, then, as $\varepsilon \downarrow 0$,*

$$T_0(\varepsilon) = \Theta(\varepsilon^{-(2\nu+1)}). \tag{5}$$

Equation (5) follows from conservation of $z_i^2 - z_j^2$: all squared weights have the same derivative, so the initial imbalance remains fixed as they shrink. Integrating the resulting one-dimensional flow gives the leading constant in Proposition D.3(i). Both EXP and DAGMA have this cycle-product form. A linear feasibility penalty gives exponent one; a pure quadratic penalty gives three. Thus squaring the residual increases the cost of separating a near tie, not just its constant. Related balancing invariants occur in homogeneous models (Du et al., 2018).

For the fixed ALM scalarization $\Psi(h) = \mu h + \rho h^2/2$, $\mu > 0$ and $\rho \geq 0$ give $\nu = 0$, whereas $\mu = 0$ and $\rho > 0$ give $\nu = 1$. A positive multiplier changes the separation exponent, not just the time scale. This comparison concerns the fixed inner flow: outer multiplier updates change the vector field and require a separate analysis.

Adding a score. A score can favor one deletion. To control separation time, that preference must persist along the path, not just at initialization.

Proposition 3.2 (Separation under a persistent score margin)**.** *Add a smooth score S and let $\dot{z} = -\nabla_z[S + \Psi(h)]$. For positive x_0, y_0 with $y_0/x_0 > \zeta$, let T_ζ be the first time $y/x = \zeta$. Suppose the trajectory can be continued in its positive domain through the bound below unless it reaches that ratio earlier. If $\partial_y(S + \Psi(h))/y - \partial_x(S + \Psi(h))/x \geq \gamma_{\text{flow}} > 0$ along the path until the target ratio is reached, then*

$$T_\zeta \leq \gamma_{\text{flow}}^{-1} \log\left(\frac{y_0}{\zeta x_0}\right). \tag{6}$$

Along the flow, $\frac{d}{dt}\log(y/x) \leq -\gamma_{\text{flow}}$; integration proves Eq. (6), as in Proposition D.3(ii). For a fixed positive margin, its upper bound stays finite as $y_0/x_0 \to 1$, unlike Eq. (5). A weak feasibility response therefore need not imply slow separation under the full objective. Part (iii) of Proposition D.3 treats a shrinking quadratic-score preference, including the logarithmic correction for a linear penalty.

Table 1: The initial ratio accounts for most of the prediction gain. Entries are mean absolute errors across 160 trajectories per method on the normalized separation time. Rows add features cumulatively. All fits use development data; the added comparisons are post-hoc. Paired intervals and censoring checks are in Appendix F.10.

Prediction features	NOTEARS	DAGMA
Constraint-and-scale baseline	0.04112	0.10780
+ initial ratio	0.04039	0.08250
+ log distance	0.04024	0.08313
+ completion depth	0.04054	0.08165
+ constraint separation	0.04125	0.08228
+ score separation	0.04105	0.08758

Controlled cycle flows. Across 72 isolated-cycle settings, the largest absolute error in the exponent predicted by Eq. (5) is 0.044. Direct matrix gradients agree with the cycle-product calculation to 1.1×10^{-15}. A 480-run experiment checks the score crossover; Appendix D.4 reports the full trajectories.

Matched-mask intervention. To separate completion size from initial imbalance, we keep the data, initial matrix, marked-edge ratio, and number of optimized coordinates fixed while varying $q \in \{2, 3, 4\}$. Across 24 fresh seeds, a larger completion increases median flow separation time at the two separating scales. Under the same masks, Adam's response changes sign across scales; L-BFGS-B has no consistent trend. The flow's ordering does not carry across solvers. Appendix F.11 gives the frozen protocol, paired contrasts, and all nonseparating runs.

Official optimizer trajectories. We record 320 official NOTEARS and DAGMA trajectories on ER2/ER4/SF2/SF4 graphs at $d = 20$, with 500 samples and formal seeds 0–9. The outcome is the first recorded event where the smallest tracked cycle weight is at most one tenth of the second-smallest. Elapsed steps are divided by the remaining steps from the early snapshot; unresolved runs receive one. Early constraint-force differences are associated with faster separation (Appendix F.9).

The predictor, fitted on development data and held fixed for evaluation, improved on a baseline using constraint value, maximum weight, and the relative critical threshold: the smallest acyclic cutoff divided by the maximum weight. Table 1 tests whether the initial edge-weight ratio explains the gain. These post-hoc fits add the ratio and its logarithmic distance from the target, keeping the development split and ridge penalty fixed. The ratio accounts for most of the DAGMA improvement. The statistic q_2 is the minimum number of edges in the critical-to-double-critical weight band needed to complete a cycle with the edges above that band (Appendix F.9). For the two smallest cycle magnitudes $a_{(1)} \le a_{(2)}$, on a two-cycle it is exactly $q_2 = 1 + \mathbf{1}\{a_{(1)}/a_{(2)} \ge 1/2\}$. On these cycles, adding it to the predictor encodes the same ratio nonlinearly; it adds no independent structural information. Constraint and score separation do not consistently improve the stronger baseline. With 319 of the 320 critical cycles being two-cycles, the audit primarily tests pairwise separation rather than higher-order completions.

When the score is estimated from data, its preference between repairs is also subject to sampling uncertainty.

4 Certifying Score-Based Support Decisions

An empirical ranking need not resolve the population ordering. If the only two population optima differ in one edge's direction, their shared adjacency may be certified without choosing an orientation.

The completion structure specifies the valid repairs: each must remove an edge from every minimal completion (Lemma G.1). Here we compare discrete alternatives under sampling uncertainty, rather than follow a fixed score's gradient as in Section 3. The margins for these two questions are different.

Frozen family and score. We score each node's parents jointly, so the contribution of one incoming edge can depend on the others. For each node j, training data determine a finite family $\mathcal{P}_j$ of parent sets, a fitted predictor for each set, a bounded per-observation loss $\ell_{j,S}$, and a penalty $a_{j,S}$. The allowed DAG family is nonempty. These objects are frozen before an independent holdout of n observations. Write $\mathrm{pa}_G(j)$ for node j's parents in G, $\mathcal{P} = (\mathcal{P}_j)_j$, and $\mathcal{G}(\mathcal{P})$ for the allowed DAGs, with $Q(G) = \sum_j (R_{j,\mathrm{pa}_G(j)} + a_{j,\mathrm{pa}_G(j)})$ for their population score. Here $R_{j,S} = \mathbb{E}[\ell_{j,S}(X) \mid \text{training}]$ averages over a fresh observation from the holdout distribution; $a_{j,S}$ is the fixed penalty. The audit uses these objects and holdout losses to certify features of the population score optima. Connecting those optima to the generating graph requires additional assumptions.

Confidence family. Fix a failure probability $0 < \alpha < 1$ and choose a reference DAG G_0, possibly using the holdout. Intervals covering all pairs of allowed parent sets supply a simultaneous lower bound $g_{j,S}$ on each local population score difference from the reference parent set. Set the reference differences to zero and define

$$\Gamma(G) = \sum_j g_{j,\mathrm{pa}_G(j)}, \qquad \mathcal{A}_\alpha(G_0) = \{G \in \mathcal{G}(\mathcal{P}) : \Gamma(G) \le 0\}. \tag{7}$$

On the coverage event, a positive sum rules out a graph as worse than the reference. The set $\mathcal{A}_\alpha(G_0)$ retains the graphs that these comparisons do not rule out.

An adjacency ignores edge direction; an unshielded collider is $u \to v \leftarrow w$ with no edge between u and w. For any such binary feature $\varphi : \mathcal{G}(\mathcal{P}) \to \{0, 1\}$, ask whether a retained graph has the opposite label:

$$L_\varphi^{\mathrm{opp}} = \min_{\substack{G \in \mathcal{G}(\mathcal{P}) \\ \varphi(G) \neq \varphi(G_0)}} \Gamma(G). \tag{8}$$

Use $\min \varnothing = +\infty$. If the minimum is positive, certify the reference label; otherwise abstain. Here *selective* refers to feature-wise decisions with abstention, not conditioning on a selection event.

Theorem 4.1 (Interaction-aware selective feature certificate)**.** *Fix $0 < \alpha < 1$ and condition on training, fixing $\mathcal{P}$, ℓ, and a. Let $\mathcal{M} = \arg\min_{G \in \mathcal{G}(\mathcal{P})} Q(G)$, and let Φ be the queried binary structural features. Suppose the independent holdout observations are i.i.d. and all paired local score intervals cover simultaneously with conditional probability at least $1 - \alpha$. For any holdout-dependent reference $G_0 \in \mathcal{G}(\mathcal{P})$, on that event,*

$$\mathcal{M} \subseteq \mathcal{A}_\alpha(G_0),$$
$$\forall \varphi \in \Phi : \quad L_\varphi^{\mathrm{opp}} > 0 \implies \forall G^* \in \mathcal{M}, \quad \varphi(G^*) = \varphi(G_0).$$

These conclusions therefore hold simultaneously with conditional probability at least $1 - \alpha$.

What can be resolved without a unique graph. Model confidence sets and split-sample tests retain plausible optima (Hansen et al., 2011; Park et al., 2026). The first inclusion in Theorem 4.1 retains *all* population score optima. Its second statement permits a feature decision without selecting one of them: different orientations may share an adjacency. The extension here is from simultaneous local parent-set comparisons to global structural labels under acyclicity, with the reference itself allowed to depend on the holdout.

The proof sums the local bounds to give $\Gamma(G) \le Q(G) - Q(G_0)$ on the joint coverage event. Every population minimizer has a nonpositive difference. A positive opposite-label minimum therefore excludes that label from all optima; all-pairs coverage permits holdout selection of G_0. Appendix H gives the intervals and proof.

Feature-resolution margin. The test decides if every opposite-label score gap exceeds the uncertainty in its changed parent sets. Write $m(G, G_0)$ for the number of nodes whose parent sets differ from the reference. For a feature with feasible opposite labels, define

$$\kappa_\varphi(G_0) = \min_{\varphi(G) \neq \varphi(G_0)} \frac{Q(G) - Q(G_0)}{m(G, G_0)}, \qquad G \in \mathcal{G}(\mathcal{P}). \tag{9}$$

On the simultaneous coverage event, if every paired radius is at most ε_n, then $\Gamma(G) \ge Q(G) - Q(G_0) - 2\varepsilon_n m(G, G_0)$. Thus $\kappa_\varphi(G_0) > 2\varepsilon_n$ is sufficient for $L_\varphi^{\mathrm{opp}} > 0$, even when several optimal DAGs remain indistinguishable. The implementation uses a solver lower bound, reduced by a

numerical guard, rather than a feasible solution's objective to certify positivity. Appendix H.2 gives the derivation; Appendix G.4 describes the numerical checks.

The margin counts changed parent sets, not changed edges, because each local comparison evaluates a joint regression. When G_0 is a population optimum, tied orientations can leave the graph-level gap at zero while all graphs lacking their shared adjacency have a positive gap. The feature margin compares against those opposite-label graphs. It can support an adjacency decision once uncertainty is small enough, without selecting an orientation.

Frontend-to-feature certification. ER2/ER4/SF2/SF4 graphs, four noise/weight regimes, and seeds 0–4 give 80 datasets at $d = 20$. NOTEARS, DAGMA, SDCD, and a GOLEM port share these datasets, for 320 runs. Each uses 500 training, 5,000 holdout, and 20,000 independent proxy observations, used only afterwards to estimate Q and its optimum. Table 2 reports the family recall and certified fraction after screening and the four-candidate-parent cap. Family recall limits recoverable support. Appendix H.7 gives the implementations, screening rules, and metrics.

Table 2: Frontend-to-feature certification (320 runs). Family recall is the mean fraction of generating edges retained after screening and the four-parent cap, over 80 datasets per frontend. Certified percentages pool all candidate-family queries of each feature type within a frontend. All decided labels are score-certified; the remainder are statistical abstentions, not family-forced conclusions or solver timeouts.

Frontend	Family recall (%)	Certified adj. (%)	Certified coll. (%)
NOTEARS	57.4	19.5	16.8
DAGMA	77.6	44.5	34.6
SDCD	53.4	10.5	4.4
GOLEM	77.2	44.4	36.3

Regret is $Q(G_0) - \min_{G \in \mathcal{G}(\mathcal{P})} Q(G)$; Appendix H gives its bound. All bounds cover the proxy regrets, and all certified labels agree with the proxy optimum, an estimate rather than the population optimum. Against the generating graph, 4.4% of certified skeleton labels and 5.5% of certified collider labels disagree, including 23 of the 80 certified absent adjacencies (Table 23). The holdout and proxy parent-set optima have mean structural Hamming distance (SHD) 14.63 and 14.56; refitting does not uniformly improve the frontend's estimate.

The certificate resolves 2,962 present adjacencies and 1,382 present colliders, compared with 80 absent adjacencies and 1,014 absent colliders. Among queries present in G_0, the resolution rates are 45.3% and 31.9%, respectively. Changing only the proxy loss's clipping range $[-B, B]$, from $B = 0.02$ to $B = 0.01$, leaves 99.45% of certified labels unchanged, compared with 87.60% of uncertified labels. This post-hoc comparison matches the same queries between proxy optima, keeping nominal certification groups fixed. Coverage applies to the nominal score (Appendix H.5).

Exact-population calibration. We enumerate all 25 three-variable DAGs under four finite-support laws with rational population scores. Two laws have two optimal orientations of one edge and a shared adjacency. In a collider-generating third law, clipping cancels every local expected loss, leaving $Q(G) = \lambda_{\mathrm{sp}}|E(G)|$, where $E(G)$ is the edge set and λ_{sp} the per-edge penalty; the unique score optimum is empty. A fourth law is a population-constructed positive control with unique optimum $0 \to 2 \leftarrow 1$. At $n = 5000$, its positive collider is certified in all 200 draws. Across four nested sample sizes per law, all 3,200 coverage checks retain every population optimum, with no observed feature error. Figure 2 separates positive-label resolution from unique-graph (singleton) certification. Appendix H.4 gives the population construction, the retained negative control, and pointwise sampling uncertainty.

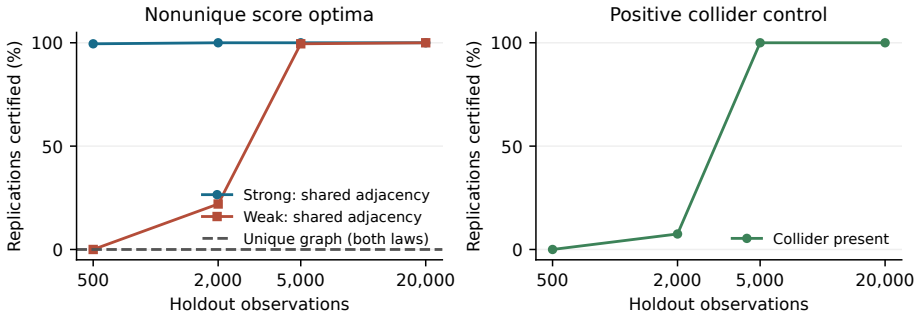


Figure 2: Exact-population calibration: feature resolution need not imply a unique graph. Left: a shared adjacency is resolved although two orientations remain optimal. Right: the constructed collider control has a unique score optimum; positive-collider and singleton certification rates coincide here. Each point uses 200 independent draws, with nested sample sizes. The control law was selected using its known population score before sampling; these are calibration rates, not benchmark accuracy.

The score and candidate family affect both resolution and computational cost. The nominal clipping level $B = 0.02$ changes 55.1% of candidate-loss entries; the certificate targets this bounded score. Exact interaction-aware queries are combinatorial. The core runs to $d = 200$ with at most two candidate parents per node, a restriction that limits the scaling result to sparse candidate families. For additive edge scores, the cycle-cover linear program (LP) provides a polynomial-time conservative specialization. It does not retain general parent interactions. Appendices G and H report the SHD comparisons, score and screen sensitivity, coverage diagnostics, and search costs.

5 Discussion

What connects the three analyses. Completions determine where an exact residual vanishes and which deletions restore acyclicity. The flow studies how a score selects among these repairs; the certificate assesses uncertainty in the population ranking. A discrete score gap does not supply the pathwise derivative margin required by the flow bound.

What constraint design can change. At a fixed score and update rule, residuals of the same local order can differ in coefficients and behavior away from the boundary. Distance control must cover all nearby directions; its constants depend on normalization.

Where further work is needed. Practical use needs computable distance constants and neighborhoods, optimizer-specific analyses, and tighter intervals and search bounds for richer parent families. The implementation separates screening limitations, score-target disagreement, statistical abstention, and solver limits. The scaling results concern bounded-parent families.

6 Conclusion

Minimal cycle completions determine two limits of smooth DAG representations: the smallest completion bounds the earliest local response, while the largest determines the best uniform distance-error exponent. Matching constructions attain both bounds within the stated representation classes.

Choosing a support requires another source of information. A persistent score margin resolves a feasibility tie in the analyzed flow. With an estimated score, simultaneous parent-set comparisons can instead certify features shared by all population optima, even when no unique graph is resolved. The numerical studies separate these guarantees from optimizer behavior and agreement with the generating DAG. For continuous DAG learning, the resulting reporting principle is concrete: state what the residual controls, how the score selects a support, and which features the data resolve under that score.

Reproducibility statement

The appendix gives the assumptions and proofs. The Lean 4 supplement formalizes mathematical cores of completion order, Newton parity, symmetry, cycle-flow implications, and certificate logic; its claim map lists the formal hypotheses and coverage limits. It does not verify global ODE existence, the LP separation oracle, or numerical performance. Frozen configurations, result tables, tests, and figure scripts accompany the code. The official NOTEARS/DAGMA trajectory instrumentation and all baseline modifications are identified in the supplement.

Statement on generative AI

Generative AI tools assisted with manuscript organization, drafting and editing, literature discovery, and the development and review of experimental code. We also used these tools to discuss theoretical claims, counterexample constructions, and proof steps, and to audit assumptions and derivations. For the experiments, assistance included designing diagnostic and sensitivity checks, preparing data-generation and evaluation scripts, and interpreting results. Synthetic observations were sampled from the structural models specified in the paper using numerical code. The authors independently checked and revised AI-assisted text and suggestions, verified references against the cited sources, examined the mathematical arguments, and tested the code by rerunning the relevant computations. The authors take full responsibility for the mathematical claims, experiments, code, and final text.

References

Kevin Bello, Bryon Aragam, and Pradeep Ravikumar. DAGMA: Learning DAGs via M-matrices and a log-determinant acyclicity characterization. In *Advances in Neural Information Processing Systems*, volume 35, pp. 8226–8239, 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/36e2967f87c3362e37cf988781a887ad-Abstract-Conference.html.

Lijie Chen, Anupam Gupta, Jian Li, Mingda Qiao, and Ruosong Wang. Nearly optimal sampling algorithms for combinatorial pure exploration. In *Proceedings of the 30th Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pp. 482–534, 2017. URL https://proceedings.mlr.press/v65/chen17a.html.

David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, 2002. URL https://jmlr.org/papers/v3/chickering02b.html.

James Cussens. Bayesian network learning with cutting planes. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence*, pp. 153–160, 2011. URL https://arxiv.org/abs/1202.3713.

James Cussens. GOBNILP: Learning bayesian network structure with integer programming. In *Proceedings of the 10th International Conference on Probabilistic Graphical Models*, volume 138 of *Proceedings of Machine Learning Research*, pp. 605–608, 2020. URL https://proceedings.mlr.press/v138/cussens20a.html.

Chang Deng, Kevin Bello, Pradeep Ravikumar, and Bryon Aragam. Markov equivalence and consistency in differentiable structure learning. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2912.

Mathias Drton and Han Xiao. Wald tests of singular hypotheses. *Bernoulli*, 22(1):38–59, 2016. doi: 10.3150/14-BEJ620.

Simon S. Du, Wei Hu, and Jason D. Lee. Algorithmic regularization in learning deep homogeneous models: Layers are automatically balanced. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/fe131d7f5a6b38b23cc967316c13dae2-Abstract.html.

Richard Ehrenborg and Gábor Hetyei. The topology of the independence complex. *European Journal of Combinatorics*, 27(6):906–923, 2006. doi: 10.1016/j.ejc.2005.04.010.

Pierre Gillot and Pekka Parviainen. Learning large DAGs by combining continuous optimization and feedback arc set heuristics. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 6713–6720, 2022. doi: 10.1609/aaai.v36i6.20626.

Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981. doi: 10.1007/BF02579273.

Peter R. Hansen, Asger Lunde, and James M. Nason. The model confidence set. *Econometrica*, 79 (2):453–497, 2011. doi: 10.3982/ECTA5771.

Jason A. Howald. Multiplier ideals of monomial ideals. *Transactions of the American Mathematical Society*, 353(7):2665–2671, 2001. doi: 10.1090/S0002-9947-01-02720-9. Preprint version: `https://arxiv.org/abs/math/0003232v1`.

Axel Hultman. Directed subgraph complexes. *The Electronic Journal of Combinatorics*, 11(1):R75, 2004. doi: 10.37236/1828.

Kaifeng Jin, Ignavier Ng, Kun Zhang, and Biwei Huang. Revisiting differentiable structure learning: Inconsistency of L_1 penalty and beyond. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, number 27, pp. 22399–22407, 2026. doi: 10.1609/aaai.v40i27.39398.

Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller, James W. Thatcher, and Jean D. Bohlinger, editors, *Complexity of Computer Computations*, pp. 85–103. Plenum Press, 1972. doi: 10.1007/978-1-4684-2001-2_9.

Guoyin Li and Boris S. Mordukhovich. Hölder metric subregularity with applications to proximal point method. *SIAM Journal on Optimization*, 22(4):1655–1684, 2012. doi: 10.1137/120864660.

Riccardo Massidda, Francesco Landolfi, Martina Cinquini, and Davide Bacciu. Constraint-free structure learning with smooth acyclic orientations. In *International Conference on Learning Representations*, 2024.

Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the 22nd Annual Conference on Learning Theory*, 2009. URL `https://arxiv.org/abs/0907.3740`.

Ezra Miller and Bernd Sturmfels. *Combinatorial Commutative Algebra*, volume 227 of *Graduate Texts in Mathematics*. Springer, 2005. doi: 10.1007/b138602.

Achille Nazaret, Justin Hong, Elham Azizi, and David Blei. Stable differentiable causal discovery. In *International Conference on Machine Learning*, volume 235, pp. 37413–37445, 2024. URL `https://proceedings.mlr.press/v235/nazaret24a.html`.

Ignavier Ng, AmirEmad Ghassami, and Kun Zhang. On the role of sparsity and DAG constraints for learning linear DAGs. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17943–17954, 2020. URL `https://proceedings.neurips.cc/paper/2020/hash/d04d42cdf14579cd294e5079e0745411-Abstract.html`.

Ignavier Ng, Sébastien Lachapelle, Nan Rosemary Ke, Simon Lacoste-Julien, and Kun Zhang. On the convergence of continuous constrained optimization for structure learning. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 8176–8198, 2022. URL `https://proceedings.mlr.press/v151/ng22b.html`.

Ignavier Ng, Biwei Huang, and Kun Zhang. Structure learning with continuous optimization: A sober look and beyond. In *Conference on Causal Learning and Reasoning*, volume 236, pp. 71–105, 2024.

Richard S. Palais. The principle of symmetric criticality. *Communications in Mathematical Physics*, 69(1):19–30, 1979. doi: 10.1007/BF01941322.

Beomjo Park, Sivaraman Balakrishnan, and Larry Wasserman. Robust universal inference for misspecified models. *Biometrika*, 113(2):asaf070, 2026. doi: 10.1093/biomet/asaf070.

Samuel Rey, Madeline Navarro, and Gonzalo Mateos. Exploiting non-negativity in DAG structure learning. *arXiv preprint arXiv:2605.19947*, 2026. URL https://arxiv.org/abs/2605.19947v1.

Bruce Reznick. Extremal PSD forms with few terms. *Duke Mathematical Journal*, 45(2):363–374, 1978. doi: 10.1215/S0012-7094-78-04519-2.

Madhumitha Shridharan and Garud Iyengar. β-th order acyclicity derivatives for DAG learning. In *International Conference on Artificial Intelligence and Statistics*, volume 258, pp. 208–216, 2025. URL https://proceedings.mlr.press/v258/shridharan25a.html.

Nils Sturma, Mathias Drton, and Dennis Leung. Testing many constraints in possibly irregular models using incomplete U-statistics. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(4):987–1012, 2024. doi: 10.1093/jrsssb/qkae022.

Y. Samuel Wang, Mladen Kolar, and Mathias Drton. Confidence sets for causal orderings. *Journal of the American Statistical Association*, 121(553):690–703, 2026. doi: 10.1080/01621459.2025.2542552.

Dennis Wei, Tian Gao, and Yue Yu. DAGs with no fears: A closer look at continuous optimization for learning bayesian networks. In *Advances in Neural Information Processing Systems*, volume 33, pp. 3895–3906, 2020.

Yue Yu, Jie Chen, Tian Gao, and Mo Yu. DAG-GNN: DAG structure learning with graph neural networks. In *International Conference on Machine Learning*, volume 97, pp. 7154–7163, 2019. URL https://proceedings.mlr.press/v97/yu19a.html.

Yue Yu, Tian Gao, Naiyu Yin, and Qiang Ji. DAGs with no curl: An efficient DAG structure learning approach. In *International Conference on Machine Learning*, volume 139, pp. 12156–12166, 2021.

Zhen Zhang, Ignavier Ng, Dong Gong, Yuhang Liu, Mingming Gong, Biwei Huang, Kun Zhang, Anton van den Hengel, and Javen Qinfeng Shi. Analytic DAG constraints for differentiable DAG learning. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=oCdIo9757e.

Xun Zheng, Bryon Aragam, Pradeep K. Ravikumar, and Eric P. Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper_files/paper/2018/hash/e347c51419ffb23ca3fd5050202f9c3d-Abstract.html.

Appendix guide. Appendices A and C give the definitions and proofs for the response and distance bounds. Appendix B develops their geometric consequences. Appendix D treats symmetry, cycle-flow dynamics, and score perturbations, with numerical tests in Appendix F. The main feature certificate is developed in Appendix H, including its solver formulation and experimental protocols. Appendix G covers the additive-edge specialization and its LP relaxation.

A Formal Order Results and Extensions

Notation. The graph has d nodes and $m = |F|$ candidate coordinates. The collection $\mathcal{C}_{\min}$ consists of minimal cycle completions, not all cycles in the graph. Its smallest and largest member sizes are q and r; weighted paths use τ instead of q. We use the positive integer N_H for a residual's output dimension. The matrix direction $\mathbf{U}$ and its coordinate vector u satisfy $\mathbf{U} = \sum_s u_s \mathbf{E}_s$. The notation D denotes a derivative, D^β a multi-index derivative, and C^k continuity of all derivatives through order k. Norms are specified where a result requires a choice other than the Euclidean norm. We write $[m] = \{1, \dots, m\}$, identify a candidate edge e_s with its index s, and use supp for nonzero coordinates. The space $\mathbb{R}^S \subseteq \mathbb{R}^m$ consists of vectors zero outside S. The notation $\mathbf{1}_S$ denotes an indicator vector, while $\mathbf{1}\{A\}$ is one when statement A holds and zero otherwise. For a scalar v, $[v]_+ = \max(v, 0)$.

A.1 EXACT REPRESENTATIONS

The representation bounds compare residuals with the same acyclic zero set but different values and derivatives away from that set.

Let $\mathcal{Z} = \{\mathbf{W} \in \mathbb{R}^{d\times d} : \operatorname{diag}(\mathbf{W}) = 0\}$, and let $\mathcal{D} \subset \mathcal{Z}$ be the matrices whose nonzero support is a DAG. We work on an open domain $\Omega \subseteq \mathcal{Z}$ containing the candidate subspaces of interest; a proper domain accommodates log-determinant constraints.

Definition A.1 (Exact representations)**.** *A smooth map $H : \Omega \to \mathbb{R}^{N_H}$ is a* signed/vector exact representation *if, for every $\mathbf{W} \in \Omega$,*

$$H(\mathbf{W}) = 0 \quad \Longleftrightarrow \quad \mathbf{W} \in \mathcal{D}. \tag{10}$$

A smooth scalar $h : \Omega \to \mathbb{R}$ is nonnegative exact *if, for every $\mathbf{W} \in \Omega$,*

$$h(\mathbf{W}) \geq 0, \qquad h(\mathbf{W}) = 0 \quad \Longleftrightarrow \quad \mathbf{W} \in \mathcal{D}. \tag{11}$$

A pathwise quantity has *exact order* p if $R(t) = t^p v + o(t^p)$ as $t \downarrow 0$ for nonzero v. The results below bound the first possible order; the positive-walk constraints attain it. Each theorem states the finite differentiability it needs; analyticity is not assumed.

A.2 CYCLE-COMPLETION COMPLEXITY

Fix a DAG $\mathbf{W} \in \mathcal{D}$ and a finite set $F = \{e_1, \ldots, e_m\}$ of absent off-diagonal coordinates. Here we also allow a signed coordinate matrix $\mathbf{E}_s$: its only nonzero entry is a fixed $+1$ or -1 at e_s. For $G \in \{H, h\}$ define the local restriction

$$V_F = \left\{ x \in \mathbb{R}^m : \mathbf{W} + \sum_{s=1}^{m} x_s \mathbf{E}_s \in \Omega \right\}, \qquad G_F(x) = G\left(\mathbf{W} + \sum_{s=1}^{m} x_s \mathbf{E}_s \right), \quad x \in V_F. \tag{12}$$

Because Ω is open, V_F is an open neighborhood of the origin.

Definition A.2 (Completion number)**.** *The candidate-subspace completion number is*

$$q_{\mathbf{W}}(F) = \min \{|S| : S \subseteq F,\ \operatorname{supp}(\mathbf{W}) \cup S \text{ is cyclic}\}, \tag{13}$$

with $q_{\mathbf{W}}(F) = \infty$ if no subset completes a cycle.

Candidate edges need not approach zero at the same rate. For positive weights $a = (a_1, \ldots, a_m) \in \mathbb{R}^m_{>0}$, define the weighted analogue

$$\tau_{\mathbf{W}}(F; a) = \min_{S \subseteq F:\ \operatorname{supp}(\mathbf{W}) \cup S \text{ cyclic}} \sum_{e_s \in S} a_s. \tag{14}$$

Along $x_s = c_s t^{a_s}$ as $t \downarrow 0$, $\tau_{\mathbf{W}}(F; a)$ is the cheapest cycle-completion exponent. At the empty graph it is weighted directed girth. Equivalently, $q_{\mathbf{W}}(F)$ is a minimum directed-cycle cost when base edges have cost zero and candidate edges unit cost; $\tau_{\mathbf{W}}(F; a)$ assigns candidate edge e_s the cost a_s. Both are computable by weighted shortest-cycle routines on the finite candidate section.

Definition A.3 (Cycle-completion ideal)**.** *The acyclic candidate supports form a restricted directed-subgraph complex (Hultman, 2004),*

$$\Delta_{\mathbf{W},F} = \{S \subseteq [m] : \operatorname{supp}(\mathbf{W}) \cup \{e_s : s \in S\} \text{ is acyclic}\}. \tag{15}$$

Let $\mathcal{C}_{\min}(\mathbf{W}, F)$ collect its inclusion-minimal nonfaces. Its Stanley–Reisner ideal is

$$I_{\mathbf{W},F} = \left\langle x^C : C \in \mathcal{C}_{\min}(\mathbf{W}, F) \right\rangle \subset \mathbb{R}[x_1, \ldots, x_m], \tag{16}$$

where $x^C = \prod_{s\in C} x_s$.

Thus $q_{\mathbf{W}}(F)$ and $\tau_{\mathbf{W}}(F; a)$ are the least ordinary and weighted degrees in $I_{\mathbf{W},F}$. When this ideal is nonzero, call $r_{\mathbf{W}}(F) = \max_{C \in \mathcal{C}_{\min}(\mathbf{W},F)} |C|$ its *completion width*. The two unweighted degrees can differ: $q_{\mathbf{W}}(F)$ controls the earliest Taylor signal, while $r_{\mathbf{W}}(F)$ controls the worst local error bound.

Theorem A.4 (Cycle-completion jet ideal)**.** *Let G_F be the restriction of either exact representation in Eqs.* (10)–(11)*, and suppose G_F is C^k. Every component of its order-k Taylor polynomial at zero belongs to $I_{\mathbf{W},F}$. Conversely, the vector of minimal generators*

$$M_{\mathbf{W},F}(x) = (x^C)_{C\in\mathcal{C}_{\min}(\mathbf{W},F)} \tag{17}$$

vanishes exactly on the acyclic candidate supports.

The ideal-membership statement in Theorem A.4 identifies which monomials any exact residual can use, rather than only its first possible nonzero order. Its converse makes the same minimal completions an exact polynomial representation. The support restrictions are therefore both necessary for Taylor coefficients and sufficient to construct a zero set; no particular trace or determinant formula is needed.

Proof. If $\operatorname{supp}(\alpha) \in \Delta_{\mathbf{W},F}$, exactness makes G_F identically zero near the origin on that coordinate subspace, so $D^\alpha G_F(0) = 0$. Every surviving Taylor monomial therefore contains a minimal nonface and belongs to $I_{\mathbf{W},F}$. The generator vector is zero exactly when the active support contains no minimal nonface, which is precisely membership in $\Delta_{\mathbf{W},F}$. □

A.3 Signed and Vector Representations

Exactness forces certain Taylor coefficients to vanish even without a nonnegativity assumption. We establish that restriction first.

Theorem A.5 (Vector cycle-completion barrier)**.** *Assume Eq.* (10)*, let $q = q_{\mathbf{W}}(F) < \infty$, and suppose H is C^{q-1} near $\mathbf{W}$. Then*

$$D^\alpha H_F(0) = 0 \qquad \textit{for every multi-index } |\alpha| \le q-1. \tag{18}$$

If H is C^q, then for every $\mathbf{U}$ supported on F,

$$\|H(\mathbf{W}+t\mathbf{U})\|_2 = O(|t|^q), \qquad \|DH_F(tu)\|_{\mathrm{op}} = O(|t|^{q-1}). \tag{19}$$

Equation (18) shows that removing nonnegativity does not remove all flatness: fewer than q candidate coordinates cannot detect a cycle. Equation (19) converts that exact derivative statement into value and Jacobian bounds under one additional derivative. The proof restricts to acyclic coordinate subspaces and then uses Taylor's remainder (Appendix C). The restriction concerns F; other directions in the full Jacobian may close a path at lower order.

Corollary A.6 (Weighted vector barrier)**.** *Let $a \in \mathbb{N}^m$, $\tau = \tau_{\mathbf{W}}(F;a) < \infty$, $c_s \neq 0$, and $\mathbf{W}_a(t) = \mathbf{W} + \sum_s c_s t^{a_s}\mathbf{E}_s$. If H is $C^{\tau-1}$, every derivative of $H(\mathbf{W}_a(t))$ below order τ vanishes at zero. If H is C^τ, then*

$$\|H(\mathbf{W}_a(t))\|_2 = O(|t|^\tau). \tag{20}$$

Here the integer path exponents are strictly positive, as in the weighted completion definition. The statement concerns this prescribed path, not a uniform bound over paths with vanishing coefficients. The proof lifts edge e_s to a_s auxiliary coordinates whose product activates the edge, then applies the local coordinate-subspace jet lemma underlying Theorem A.5; Appendix C.2 gives details.

The bound is sharp. To see this directly, let $\mathcal{C}_d$ be the set of all simple directed cycles and define

$$H_{\mathrm{cyc}}(\mathbf{W}) = \left(\prod_{(i,j)\in C} W_{ij}\right)_{C\in\mathcal{C}_d}. \tag{21}$$

This polynomial vector is zero exactly on DAG support. A minimal q-edge completion produces a component of order exactly q, and a minimum-weight completion produces order τ. Its representation size is

$$|\mathcal{C}_d| = \sum_{k=2}^{d}\binom{d}{k}(k-1)!. \tag{22}$$

Appendix B studies the output dimension needed for uniform directional sharpness.

A.4 NONNEGATIVE SCALAR REPRESENTATIONS

Nonnegativity restricts which of the surviving monomials can lead the expansion. The Newton polytope is the convex hull of the exponent vectors of a polynomial's nonzero monomials. For background on sparse nonnegative forms, see Reznick (1978). The elementary parity property below turns the support restriction into a degree bound; Appendix C.5 gives a direct proof for the polynomial statement used here.

Lemma A.7 (Newton vertices)**.** *Every vertex exponent of the Newton polytope of a globally nonnegative real polynomial is coordinatewise even, and its coefficient is positive.*

Theorem A.8 (Nonnegative scalar barrier)**.** *Assume Eq.* (11)*, let $q = q_{\mathbf{W}}(F) < \infty$, and suppose h is C^{2q-1} near $\mathbf{W}$. Then*

$$D^\alpha h_F(0) = 0 \qquad \textit{for every } |\alpha| \le 2q-1. \tag{23}$$

If h is C^{2q}, then

$$h(\mathbf{W}+t\mathbf{U}) = O(|t|^{2q}), \qquad \|Dh_F(tu)\| = O(|t|^{2q-1}). \tag{24}$$

Equation (23) adds a full block of vanishing orders, $q, \dots, 2q-1$, to the vector result. This additional restriction comes from nonnegativity on a signed neighborhood, not specifically from squaring weights in NOTEARS. Equation (24) gives its quantitative effect on the residual and derivative. Lemma A.7 supplies the parity argument (Appendix C.5); a smooth nonanalytic residual can vanish even faster, with no finite leading Taylor form.

Corollary A.9 (Newton-square geometry)**.** *Let P_r be the first nonzero homogeneous Taylor form of h_F. Every vertex exponent α of its Newton polytope satisfies $x^\alpha \in I^2_{\mathbf{W},F}$. Hence, for every positive weight vector a, the minimum a-weight at a Newton vertex is at least $2\tau_{\mathbf{W}}(F;a)$.*

Proof. Theorem A.4 gives $P_r \in I_{\mathbf{W},F}$. Lemma A.7 makes α even, while $\operatorname{supp}(\alpha/2) = \operatorname{supp}(\alpha)$ is a nonface. Thus $x^{\alpha/2} \in I_{\mathbf{W},F}$ and $x^\alpha \in I^2_{\mathbf{W},F}$. □

For $a \in \mathbb{N}^m$, the same auxiliary-coordinate argument gives a weighted $2\tau_{\mathbf{W}}(F;a)$ barrier with the required regularity made explicit. If h is $C^{2\tau-1}$ near $\mathbf{W}$, every derivative of $h(\mathbf{W}_a(t))$ below order 2τ vanishes; if h is $C^{2\tau}$, then $h(\mathbf{W}_a(t)) = O(|t|^{2\tau})$. For one minimal completion, a stronger local factorization holds:

$$h_F(x) = \left(\prod_{s=1}^{q} x_s^2\right)\Psi(x), \qquad \Psi(x) \ge 0 \quad \text{locally}, \tag{25}$$

under C^{2q} regularity. The positive walk families used by standard continuous methods attain this lower bound.

Write $h_f(W) = \sum_{k\ge1} c_k \operatorname{tr}((W\circ W)^k)$ for a positive walk sum, with $c_k \ge 0$ and positive coefficients at every simple-cycle length. In DAGMA, $s > 0$ is the log-determinant domain parameter and $\rho(A)$ is the spectral radius of $A = W \circ W$. Its M-matrix domain here is precisely $\rho(A) < s$.

Corollary A.10 (Local order of positive walk constraints)**.** *Let F be a minimal q-edge completion at a DAG W and let U activate every edge in F with a nonzero equal-scale coefficient. Then NOTEARS EXP, the usual positive trace polynomial, and DAGMA's log determinant while $\rho(W\circ W) < s$ satisfy*

$$h(W+tU) = \Theta(|t|^{2q}), \qquad \|Dh_F(tu)\|_2 = \Theta(|t|^{2q-1}). \tag{26}$$

The same conclusion holds for every positive walk sum with a positive coefficient at each simple-cycle length that converges locally to a C^{2q} function. It also holds with an entrywise squash $\sigma(W)$ in place of $W\circ W$, provided the same local regularity holds and σ is smooth with $\sigma(0) = 0$, $\sigma(z) > 0$ for $z \ne 0$, and $\sigma(z) = \gamma z^2 + o(z^2)$ as $z \to 0$, where $\gamma > 0$.

The proof is Proposition C.3 in Appendix C.6. For DAGMA, $-\log\det(sI - A) + d\log s = \sum_{k\ge1} \operatorname{tr}(A^k)/(ks^k)$ inside its M-matrix domain. The log-determinant can therefore improve coefficients and the landscape away from the boundary without changing the topology-controlled order at the boundary itself.

B Additional Geometry and Consequences

The main results concern response order and distance control. The results in this appendix answer four further questions: how many residual outputs uniform sharpness requires, which local topologies candidate sections can realize, how large small-residual regions are, and how a residual scales at a noisy estimator. They describe further consequences of the same completion structure; the feature certificate uses only its feasible repair family.

B.1 Uniform sharpness and residual dimension

Write $J_q(U) = D^q H(0)[U, \dots, U]/q!$ for the degree-q Taylor jet. Call H *directionwise q-sharp* on a family $U(z)$ if $J_q(U(z)) \neq 0$ for every $z \neq 0$ in that family. Its *q-sharp residual complexity* $s_q(U)$ is the minimum output dimension of a C^q map $H : \mathcal{Z} \to \mathbb{R}^{N_H}$ that is globally exact and has this property.

Theorem B.1 (Near-tight sensitivity–succinctness law)**.** *For every $3 \leq q \leq d$, there is an explicit channel family U_q with*

$$M_q(d) = \left\lfloor \frac{(d-q+2)^2}{4} \right\rfloor \tag{27}$$

independent q-cycle channels satisfying

$$M_q(d) \leq s_q(U_q) \leq M_q(d) + 1. \tag{28}$$

For $q = 2$, there is a family with one channel per unordered node pair and

$$\binom{d}{2} \leq s_2(U_2) \leq \binom{d}{2} + 1. \tag{29}$$

Consequently, $s_q(U_q) = \Theta(d^2)$ for every fixed q. Below either lower bound, some cyclic mixture has zero degree-q jet; the upper bound is globally exact and directionwise sharp.

The lower bounds in Eqs. (28) and (29) concern sensitivity in *every* nonzero channel mixture, not exactness alone. A scalar can have the correct zero set while losing the degree-q response in some mixture. The upper bounds show that retaining this response costs at most one output beyond the rank lower bound. Adding an exact scalar to the selected cycle products gives that construction (Section C.4).

B.2 Coordinate topology

A simplicial complex is a family of finite sets closed under taking subsets; its members are faces and the other sets are nonfaces. Write $|K|$ for its geometric realization and $\mathrm{sd}\, K$ for its barycentric subdivision, whose faces are chains of nonempty faces of K. Below, $\cong$ denotes an isomorphism or homeomorphism as appropriate, and $\simeq$ denotes homotopy equivalence. The symbol $\bigvee_N S^k$ means N spheres joined at one point; $\widetilde{\beta}_j$ is the dimension of reduced homology over $\mathbb{Q}$.

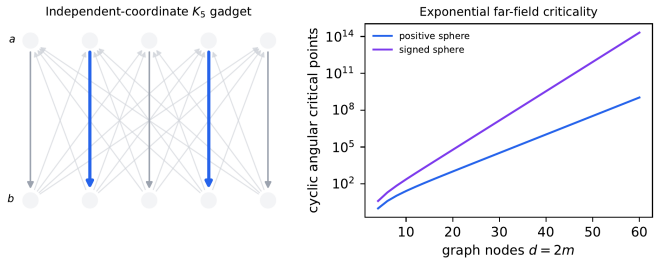


Figure 3: Ordinary-coordinate construction. A K_5 gadget realizes all pairwise cyclic candidate supports (left), producing exponentially many cyclic angular critical points under disjoint union (right).

The distance bound concerns the hardest direction near a support. To describe how those directions are arranged, we now consider the topology of the positive normalized candidate section,

$$\mathcal{A}^{+}_{\mathbf{W},F} = \left\{ x \geq 0 : \|x\|_2 = 1,\ \mathbf{W} + \sum_{e \in F} x_e \mathbf{E}_e \in \mathcal{D} \right\} \cong |\Delta_{\mathbf{W},F}|. \tag{30}$$

Theorem B.2 (Coordinate DAG-germ universality)**.** *For every finite simplicial complex K, there are a base DAG $\mathbf{W}_0$ and independently parameterized absent adjacency entries F such that*

$$\Delta_{\mathbf{W}_0,F} \cong \operatorname{sd} K. \tag{31}$$

Consequently $\mathcal{A}^{+}_{\mathbf{W}_0,F}$ is homeomorphic to $|K|$.

The isomorphism in Eq. (31) realizes any finite complex after subdivision, using one ordinary adjacency coordinate per candidate. Its homeomorphism consequence transfers that complex's topology to the positive acyclic section. In contrast, the full acyclic-subgraph complex has a restricted homotopy type (Hultman, 2004). The result identifies what changes after fixing base edges and restricting the remaining candidates: a simple full complex need not leave a simple conditional search section.

Corollary B.3 (Exponential local homology)**.** *There is a family with $d = 10t$ graph nodes and $5t$ independent candidates such that*

$$\Delta_{\mathbf{W}_t,F_t} \simeq \bigvee_{4^t} S^{t-1}, \qquad \sum_j \widetilde{\beta}_j(\Delta_{\mathbf{W}_t,F_t}) = 4^t = 2^{d/5}. \tag{32}$$

For the next corollary, use a differentiable square-symmetric residual (Eq. (87)) on a domain containing the points considered. An angular critical point is a stationary point of its candidate restriction on the indicated unit sphere. The K_m gadget comes from the complete undirected graph on m vertices, as constructed in Appendix D.

Corollary B.4 (Exponential angular criticality)**.** *On one K_m gadget, the positive and signed candidate unit spheres contain at least $2^m - m - 1$ and $3^m - 2m - 1$ distinct cyclic critical points, respectively. All corresponding DAGMA points lie in its $s = 1$ M-matrix domain.*

The proof uses the fact that every finite complex becomes an independence complex after barycentric subdivision (Ehrenborg & Hetyei, 2006, p. 922); the DAG contribution is an ordinary-adjacency realization of that complex.

B.3 Metric consequences

Derivative order also controls how reliably a residual measures distance to the DAG set. Write

$$\mathcal{D}_F = \left\{ x \in V_F : \mathbf{W} + \sum_s x_s \mathbf{E}_s \in \mathcal{D} \right\}. \tag{33}$$

For an exact residual R, write R_F for its candidate restriction. A local Hölder error bound with exponent $\theta > 0$ is an inequality

$$\operatorname{dist}(x, \mathcal{D}_F) \leq \kappa \, \|R_F(x)\|^{\theta} \tag{34}$$

for all x near zero and some $\kappa > 0$; this is also called Hölder metric subregularity (Li & Mordukhovich, 2012).

Theorem 2.2 states the sharp exponents in the main text. Proposition C.4 proves the generator bound and its sharpness; positive walk comparison transfers it to the standard scalar constraints.

Two further invariants describe the size of neighborhoods rather than their worst direction. Let

$$b_{\mathbf{W}}(F) = m - \max_{S \in \Delta_{\mathbf{W},F}} |S|,$$

$$\lambda_{\mathbf{W}}(F) = \min_{\substack{z \geq 0 \\ \sum_{s \in C} z_s \geq 1,\ C \in \mathcal{C}_{\min}}} \sum_{s=1}^{m} z_s. \tag{35}$$

Thus b is the minimum integral completion cover (the height of $I_{\mathbf{W},F}$), while λ is its fractional relaxation. In the volume statements, residuals are evaluated on the candidate section, with the restriction subscript F suppressed, and vol denotes Lebesgue volume in $\mathbb{R}^m$.

Theorem B.5 (Completion-volume geometry)**.** *For a sufficiently small box $B_\rho = [-\rho, \rho]^m$ and $0 \le \varepsilon \le \rho$, the normalized ℓ_∞ tube volume is exactly*

$$\sum_{S \in \Delta_{\mathbf{W},F}} (1 - \varepsilon/\rho)^{|S|} (\varepsilon/\rho)^{m-|S|}; \tag{36}$$

its small-ε exponent is b. If $\mathcal{V}_R(\eta) = \mathrm{vol}\{x \in B_\rho : \|R(x)\| \le \eta\}$, then

$$\lim_{\eta \downarrow 0} \frac{\log \mathcal{V}_{H_{\mathrm{cyc}}}(\eta)}{\log \eta} = \lambda, \qquad \lim_{\eta \downarrow 0} \frac{\log \mathcal{V}_{h_f}(\eta)}{\log \eta} = \frac{\lambda}{2} \tag{37}$$

for the simple-cycle vector and every locally convergent positive-walk scalar h_f. Moreover, $b/r \le \lambda \le b$.

Equation (36) counts points geometrically close to the acyclic set and gives the integral-cover exponent b. Equation (37) counts points accepted by a residual tolerance and instead gives λ or $\lambda/2$. The two statements separate geometric proximity from residual acceptance: the size of an accepted region is not determined by the worst-direction exponent r. They identify the classical monomial-ideal threshold (Howald, 2001) with the fractional completion cover. The substitution $|x_s| = e^{-u_s}$ turns product bounds into cover inequalities; LP duality gives the matching exponents (Section C.8).

The error-bound exponents rule out Lipschitz metric subregularity for a smooth nonnegative exact scalar at such a support boundary, and for a vector residual when $r \ge 2$. The response orders also prevent finite smooth penalties from being exact against first-order score descent (Corollary D.9).

Smooth scalarization adds a second loss of sensitivity. On paths where a vector residual is q-sharp and a nonnegative scalar residual has exact order $2q$, the common squared-penalty constructions obey

$$\underbrace{H}_{q} \quad \longrightarrow \quad \underbrace{\tfrac{1}{2}\|H\|^2}_{2q} \quad \text{and} \quad \underbrace{h}_{2q} \quad \longrightarrow \quad \underbrace{\tfrac{1}{2}h^2}_{4q}. \tag{38}$$

The subscripts are exact orders under sharpness and lower bounds otherwise. Increasing an augmented-Lagrangian penalty changes coefficients, not these missing derivatives; a residual-sized multiplier also leaves a cold start at the squared order. Sections D.7–D.8 give the full sensitivity–smoothness frontier, ALM statements, and finite-precision calculations.

B.4 Statistical consequences

The derivative orders in the main paper also determine the null scale of a residual evaluated at a noisy estimator. For fixed d, let $W, \widehat{W}_n \in \mathcal{Z}$, where $\mathcal{Z}$ is the zero-diagonal matrix space, and let F contain all absent off-diagonal coordinates of the fixed DAG W. Suppose

$$\sqrt{n}(\widehat{W}_n - W) \;\Rightarrow\; Z \tag{39}$$

for a Gaussian matrix $Z \in \mathcal{Z}$ whose off-diagonal coordinates have a positive-definite covariance matrix. Nondegeneracy is on $\mathcal{Z}$, not on the full matrix space. The sampling assumption is separate from the representation assumptions. It can arise from randomized interventions, temporal information, or another model that identifies the directed parameter. Here $\Rightarrow$ denotes convergence in distribution, and a subscript F extracts the candidate coordinates. The notation $o_p(1)$ used in the proof denotes convergence to zero in probability.

Theorem B.6 (Topology-indexed delta law)**.** *Under Eq.* (39) *and the zero-diagonal Gaussian setting above, let $q = q_W(F) < \infty$ and consider a locally convergent positive-walk constraint*

$$h_f(A) = \sum_{k \ge 1} c_k \,\mathrm{tr}((A \circ A)^k), \qquad c_k > 0. \tag{40}$$

Then

$$n^q h_f(\widehat{W}_n) \Rightarrow P_{W,q}(Z_F), \qquad P_{W,q}(z) = \sum_{\substack{C \in \mathcal{C}_{\min}(W,F) \\ |C| = q}} a_C \prod_{e \in C} z_e^2, \tag{41}$$

where every a_C is positive. For the local-alternative extension, suppose $U \in \mathcal{Z}$ and $W_n = W + U/\sqrt{n}$ belongs to a specified statistical model with laws P_{W_n}, and assume separately that

$$\sqrt{n}(\widehat{W}_n - W_n) \xRightarrow[P_{W_n}]{} Z. \tag{42}$$

Then, under P_{W_n}, the limit is $P_{W,q}(Z_F + U_F)$. The fixed-parameter convergence in Eq. (39) *alone does not imply this extension.*

Equation (41) specifies both the normalization n^q and the limiting polynomial from the smallest completions. Thus equal root-n estimation accuracy need not give equal residual scales at different supports; fixed active-edge fluctuations enter later. The local-alternative statement shifts this same polynomial only under the separate stability condition in Eq. (42). The higher-order delta method is the established tool (Drton & Xiao, 2016; Sturma et al., 2024); the completion structure identifies the terms and powers required for this DAG residual. Section C.9 gives the proof.

Fix a positive threshold constant c and a real scaling exponent a for the following acceptance rule.

Corollary B.7 (No universal scalar tolerance)**.** *For a rule that accepts when $h_f(\widehat{W}_n) \le cn^{-a}$,*

$$\Pr(\text{accept}) \longrightarrow \begin{cases} 0, & a > q, \\ \Pr\{P_{W,q}(Z_F) \le c\}, & a = q, \\ 1, & a < q. \end{cases} \tag{43}$$

Thus one exponent a cannot be nondegenerately calibrated at two DAG strata with different completion numbers. Moreover, if one completion complex has minimal generators of sizes $q < r$, the limit in Eq. (41) *retains only size-q generators. A local shift on a disjoint size-r generator is invisible to the shortest-order scalar limit.*

A residual cutoff calibrated on one support need not retain its nominal behavior on another. Even within one support, calibration at the shortest order can miss shifts along larger disjoint completions. The consequence is a need for support-dependent calibration under the stated sampling model.

C Proofs of Completion-Order Results

The response proof starts with vanishing on acyclic coordinate subspaces. Newton parity supplies the extra factor of two for nonnegative scalars. The distance proof then uses a largest minimal completion for necessity and a coordinate-deletion construction for sufficiency. These are separate steps: the smallest completion alone does not determine a uniform error bound.

All derivatives below are taken in the Euclidean space $\mathcal{Z} = \{W \in \mathbb{R}^{d\times d} : \operatorname{diag}(W) = 0\}$. For a multi-index α, $\operatorname{supp}(\alpha) = \{s : \alpha_s > 0\}$.

C.1 Coordinate-subspace vanishing

We first record the elementary mechanism behind the vector result.

Lemma C.1 (Local coordinate-subspace jet)**.** *Let $V \subseteq \mathbb{R}^m$ be an open neighborhood of zero and let $G : V \to \mathbb{R}^{N_H}$ be C^k. Suppose that, for every coordinate set S with $|S| \le k$, G is zero on a neighborhood of zero in $V \cap \mathbb{R}^S$. Then*

$$D^\alpha G(0) = 0 \qquad \textit{for every } |\alpha| \le k. \tag{44}$$

Proof. Fix α with $|\alpha| \le k$ and set $S = \operatorname{supp}(\alpha)$. Since $|S| \le |\alpha| \le k$, the local restriction G_S obtained by setting all coordinates outside S to zero is identically zero near the origin. The derivative $D^\alpha G(0)$ is the corresponding derivative of G_S and is therefore zero, component by component. □

Proof of Theorem A.5. For every $S \subseteq F$ with $|S| < q_W(F)$, the support $\operatorname{supp}(W) \cup S$ is acyclic. Exactness therefore makes H_F zero on a neighborhood of the origin in the coordinate subspace indexed by S. Lemma C.1, with $k = q - 1$, proves Eq. (18). If H is C^q, the multivariate Taylor formula with integral remainder gives, uniformly for u in a compact set,

$$\|H_F(tu)\|_2 \le C|t|^q \|u\|_2^q. \tag{45}$$

Apply the same formula to DH_F. Its derivatives through order $q - 2$ are derivatives of H_F through order $q - 1$ and vanish at zero, giving

$$\|DH_F(tu)\|_{\text{op}} \le C'|t|^{q-1}. \tag{46}$$

This is the Jacobian of the restriction to F. A coordinate outside F may close an existing path with fewer edges, so the same order need not hold for the full Jacobian. □

The theorem gives a lower bound on the first possible nonzero order, not an equality. Smooth nonanalytic exact maps may be flat to all orders, and signed components can cancel on a particular path.

C.2 WEIGHTED PATHS

We prove both weighted statements by a clone-variable lift, using the local coordinate-subspace argument directly rather than treating clones as graph edges. Let $a_s \in \mathbb{N}$, introduce variables $z_{s,1}, \dots, z_{s,a_s}$, and define

$$\Gamma(z) = W + \sum_{s=1}^{m} c_s \left(\prod_{j=1}^{a_s} z_{s,j} \right) E_s. \tag{47}$$

Original edge e_s is active exactly when all of its clones are nonzero. If $\mathrm{supp}(W)$ together with the active original edges is cyclic, the support of z contains at least

$$\min_{S:\mathrm{supp}(W)\cup S \text{ cyclic}} \sum_{e_s \in S} a_s = \tau_W(F; a) \tag{48}$$

clone coordinates. Consequently, every clone-coordinate subspace of dimension below τ maps to DAG support. Conversely, activating all clones of a minimizing set S uses exactly τ clone coordinates and maps to cyclic support. Thus τ is exactly the coordinate-subspace completion number of the lifted composite.

Let $V_\Gamma = \{z : \Gamma(z) \in \Omega\}$, an open neighborhood of zero, and set $G = H \circ \Gamma$. If H is $C^{\tau-1}$, then so is G, and G is zero near the origin on every clone-coordinate subspace of dimension below τ. Lemma C.1 makes all derivatives of G below order τ vanish. If H is C^τ, Taylor's formula also gives $G(z) = O(\|z\|^\tau)$. On the diagonal clone path $z_{s,j} = t$, Eq. (47) is precisely $W_a(t)$, proving Corollary A.6.

For the scalar statement set $g = h \circ \Gamma$. Assume first that h is $C^{2\tau-1}$. The function g is locally nonnegative and is zero on every clone-coordinate subspace of dimension below τ. If its first nonzero homogeneous Taylor term had degree $r < 2\tau$, the proof of Theorem A.8 would apply verbatim: every monomial would use at least τ clones, while a Newton vertex of a nonnegative polynomial has even exponents, forcing $r \geq 2\tau$. Hence all derivatives below 2τ vanish. Under $C^{2\tau}$ regularity, Taylor's formula yields

$$h(W_a(t)) = O(|t|^{2\tau}). \tag{49}$$

The same one-sided asymptotic bounds for positive rational exponents follow after multiplying all exponents by a common denominator and reparameterizing $t \downarrow 0$; derivative statements at zero require the resulting path to possess the corresponding regularity.

C.3 A SHARP VECTOR REPRESENTATION

We verify exactness and count the components of Eq. (21).

Proposition C.2 (Simple-cycle vector)**.** $H_{\mathrm{cyc}}(W) = 0$ *if and only if* $\mathrm{supp}(W)$ *is acyclic. For a minimal* q*-edge completion, its restricted residual and Jacobian orders are exactly* q *and* $q - 1$*, respectively. Along a weighted path the residual order is exactly* τ*, provided the coefficients on a cheapest completion are nonzero.*

Proof. Every finite directed graph is cyclic if and only if it contains a simple directed cycle. The component indexed by such a cycle is nonzero exactly when all its edges are present, proving exactness. A cycle created by a minimal completion must use all q completion edges; otherwise a proper subset would already be cyclic. Its component is a nonzero fixed-edge coefficient times the product of those q candidate coordinates. This has equal-scale order q, and differentiating in a completion coordinate has order $q - 1$. No component has lower order by Theorem A.5. The same argument chooses a cycle whose completion cost is τ on a weighted path. □

There are $\binom{d}{k}$ choices of nodes for a length-k cycle. Fixing the smallest selected node as the first node leaves $(k-1)!$ directed orders, each representing one cyclic-rotation class. Summing over $k = 2, \dots, d$ proves Eq. (22).

C.4 Uniform sharpness and residual dimension

We prove Theorem B.1. Let

$$P_q(U) = \frac{1}{q!} D^q H(0)[U, \dots, U] \tag{50}$$

be the homogeneous degree-q Taylor jet of H at the empty graph.

Fix $3 \le q \le d$ and set $n = d - q + 2$. Partition n endpoint nodes into disjoint sets S, T of sizes $a = \lfloor n/2 \rfloor$ and $b = \lceil n/2 \rceil$. The remaining $q-2$ nodes form singleton layers $v_1, \dots, v_{q-2}$. Let $U(z)$ contain all forward funnel edges

$$s \to v_1, \quad v_1 \to \cdots \to v_{q-2}, \quad v_{q-2} \to t \tag{51}$$

with weight one for $s \in S, t \in T$, and let the reverse edge $t \to s$ have weight z_{st}. The forward graph is acyclic. Every nonzero reverse channel completes the simple q-cycle $s \to v_1 \to \cdots \to v_{q-2} \to t \to s$; hence $tU(z)$ is cyclic for every $t \neq 0$ and $z \neq 0$.

Any cyclic support in this funnel must contain one reverse edge and all $q-1$ edges of its associated forward path. A degree-q Taylor monomial with any other support is therefore acyclic and its coefficient vanishes by exactness on that coordinate subspace. Total degree q also forces exponent one on each cycle edge. It follows that

$$P_q(U(z)) = \sum_{s \in S, t \in T} v_{st} z_{st} = Lz, \qquad L : \mathbb{R}^{ab} \to \mathbb{R}^{N_H}. \tag{52}$$

Directionwise q-sharpness requires $Lz \neq 0$ for every $z \neq 0$, so L is injective and

$$N_H \ge ab = \left\lfloor \frac{(d-q+2)^2}{4} \right\rfloor. \tag{53}$$

Conversely, if $N_H < ab$, rank–nullity gives $z \neq 0$ with $P_q(U(z)) = 0$. Theorem A.5 removes every lower jet, so $H(tU(z)) = o(|t|^q)$ even though exactness requires it to be nonzero for all sufficiently small $t \neq 0$.

It remains to prove the upper bound. Let $C_q(W) \in \mathbb{R}^{ab}$ collect the cycle products associated with the selected funnel channels, and define

$$H_q^+(W) = (C_q(W), h_{\exp}(W)) \in \mathbb{R}^{ab+1}. \tag{54}$$

The EXP component is zero exactly on DAGs, so H_q^+ is globally exact. On the funnel family, $C_q(tU(z)) = t^q z$ up to fixed nonzero forward-edge coefficients. Hence its degree-q jet is nonzero for every $z \neq 0$, proving $s_q(U_q) \le ab + 1$.

For $q = 2$, label nodes by a total order. Let $U(z)_{ij} = 1$ for every $i < j$, and let $U(z)_{ji} = z_{ij}$. Any two directed edges form a cycle if and only if they are the two opposite orientations of one unordered pair. Thus exactness eliminates every quadratic Taylor monomial except $W_{ij}W_{ji}$. Consequently

$$P_2(U(z)) = \sum_{i<j} v_{ij} z_{ij} = Lz. \tag{55}$$

Directionwise second-order sharpness makes $L : \mathbb{R}^{\binom{d}{2}} \to \mathbb{R}^{N_H}$ injective, proving $N_H \ge \binom{d}{2}$. Conversely, if $N_H < \binom{d}{2}$, rank–nullity supplies a nonzero z with $P_2(U(z)) = 0$. The support of $tU(z)$ contains a two-cycle for every $t \neq 0$, so exactness still requires $H(tU(z)) \neq 0$; the first two Taylor jets vanish and $H(tU(z)) = o(t^2)$. For the upper bound, collect the $\binom{d}{2}$ products $W_{ij}W_{ji}$ and append $h_{\exp}(W)$ as in Eq. (54). This map is globally exact and its quadratic jet on the pair-channel family is the coefficient vector z.

C.5 NONNEGATIVE SCALAR REPRESENTATIONS

We give a proof of the Newton-polytope fact used in the main text and then prove the scalar barrier in full.

Proof of Lemma A.7. Let $P(x) = \sum_\beta c_\beta x^\beta$ be a nonzero globally nonnegative polynomial and let α be a vertex of its Newton polytope. A vector $w \in \mathbb{R}^m$ uniquely exposes α; changing its sign if necessary, take $\langle w, \alpha\rangle < \langle w, \beta\rangle$ for every other exponent. For a sign vector $s \in \{-1, 1\}^m$, evaluate P on the monomial curve $x_i = s_i t^{w_i}$. Although coordinates with $w_i < 0$ diverge as $t \downarrow 0$, the polynomial remains nonnegative. Unique exposure gives

$$t^{-\langle w,\alpha\rangle} P(s_1 t^{w_1}, \ldots, s_m t^{w_m}) \longrightarrow c_\alpha s^\alpha. \tag{56}$$

The limit is nonnegative for every s. If some α_i were odd, flipping s_i would flip the nonzero limit. Thus α is coordinatewise even, and taking $s = \mathbf{1}$ gives $c_\alpha > 0$. □

Proof of Theorem A.8. Suppose, for contradiction, that the Taylor jet of h_F has a first nonzero homogeneous term P_r with $r < 2q$. Local nonnegativity of h_F is sufficient to make P_r globally nonnegative: for fixed x and $t \downarrow 0$,

$$t^{-r} h_F(tx) \longrightarrow P_r(x) \geq 0. \tag{57}$$

Every coordinate subspace supported on fewer than q candidates is acyclic, so the restriction of h_F, and hence of P_r, to that subspace is zero. Therefore every monomial of P_r uses at least q distinct coordinates. Choose any vertex exponent α of the Newton polytope of P_r. Lemma A.7 makes every nonzero α_s at least two. Hence

$$r = |\alpha| \geq 2|\operatorname{supp}(\alpha)| \geq 2q, \tag{58}$$

a contradiction. Thus all derivatives through order $2q - 1$ vanish. Under C^{2q} regularity, Taylor remainder bounds give $h_F(tu) = O(|t|^{2q})$. Applying the same argument to Dh_F gives the restricted gradient order $O(|t|^{2q-1})$. □

Nonnegativity is essential for the factor of two. For $d = 2$, the signed scalar $W_{12}W_{21}$ is exact and has order two at the empty graph, whereas every smooth nonnegative scalar exact representation has order at least four there.

C.6 MINIMAL-COMPLETION FACTORIZATION

Let $F = \{e_1, \ldots, e_q\}$ itself be a minimal completion and write $\Phi(x) = h(W + \sum_s x_s E_s)$ on a sufficiently small box contained in V_F. If $x_s = 0$, only a proper subset of F is active and exactness gives $\Phi(x) = 0$. For fixed x_{-s}, the nonnegative differentiable function $u \mapsto \Phi(x_{-s}, u)$ has a minimum at zero, so

$$\Phi(x_{-s}, 0) = 0, \qquad \partial_s \Phi(x_{-s}, 0) = 0. \tag{59}$$

The integral form of the second-order Hadamard lemma is

$$f(u, z) = u^2 \int_0^1 (1 - r)\partial_{uu} f(ru, z)\, dr \tag{60}$$

whenever $f(0, z) = \partial_u f(0, z) = 0$. We apply it inductively. After $k - 1$ steps, write

$$\Phi(x) = \left(\prod_{s<k} x_s^2 \right) \Phi_{k-1}(x), \tag{61}$$

where Φ_{k-1} is nonnegative and $C^{2q-2(k-1)}$. On the dense set where the preceding coordinates are nonzero, $\Phi_{k-1}(x) = 0$ whenever $x_k = 0$; continuity extends this identity to the divided hyperplanes. Because Φ_{k-1} is nonnegative and differentiable, its x_k derivative also vanishes there. Equation (60) extracts another factor x_k^2 and leaves a nonnegative C^{2q-2k} quotient. The C^{2q} assumption permits all q steps and yields

$$\Phi(x) = \left(\prod_{s=1}^{q} x_s^2 \right) \Psi(x) \tag{62}$$

with continuous Ψ. Off the coordinate hyperplanes the squared product is positive and exactness gives $\Phi > 0$, so $\Psi > 0$ there; continuity gives $\Psi \geq 0$ everywhere locally.

Proposition C.3 (Sharpness for nonnegative walk sums)**.** *Let $A = W \circ W$ and $h_f(W) = \sum_{k\geq 1} c_k \operatorname{tr}(A^k)$, where $c_k \geq 0$ and the coefficient for every simple-cycle length is positive. Suppose the series converges in a neighborhood of a minimal q-edge completion to a C^{2q} function. Then*

$$h_f(W + tU) = C|t|^{2q} + o(|t|^{2q}), \qquad \|D(h_f)_F(tu)\|_2 = \Theta(|t|^{2q-1}) \tag{63}$$

for some $C > 0$.

Proof. Choose a simple cycle in the completed support. Minimality forces it to use all q candidate edges. Its cyclic rotations contribute a positive constant times the product of the q squared candidate coordinates to the appropriate trace power. All other walk-sum terms are nonnegative, and the same expansion is zero exactly on DAG support. Thus $h_f(W + tU) \geq c|t|^{2q}$ for some $c > 0$ and all sufficiently small nonzero t. Theorem A.8 and C^{2q} regularity give $h_f(W+tU) = C|t|^{2q} + o(|t|^{2q})$; the lower bound forces $C > 0$. Let P_{2q} be the first homogeneous Taylor term, so $P_{2q}(u) = C > 0$. The restricted gradient has leading term $t^{2q-1}\nabla P_{2q}(u)$. Euler's identity gives $\langle \nabla P_{2q}(u), u\rangle = 2qP_{2q}(u) = 2qC \neq 0$, so the gradient norm is bounded below by a positive multiple of $|t|^{2q-1}$. Its Taylor expansion gives the matching upper bound. □

This proposition covers EXP and the usual positive trace polynomial. Under the squash conditions in Corollary A.10, fixed active edges retain positive factors, and each candidate contributes $\sigma(tu_e) = \gamma t^2 u_e^2 + o(t^2)$. The cycle product therefore still has a positive coefficient at degree $2q$; the same Taylor argument gives the gradient order. The log-determinant series has the same positive walk expansion inside its convergence domain.

C.7 General completion complexes and exact error bounds

Proposition C.4 (Generator domination and optimal exponents)**.** *Let $\mathcal{C}_{\min} = \mathcal{C}_{\min}(W, F)$ be nonempty, $N = |\mathcal{C}_{\min}|$, and $r = \max_{C\in\mathcal{C}_{\min}} |C|$. On a sufficiently small candidate box,*

$$\operatorname{dist}(x, \mathcal{D}_F) \leq \sqrt{N}\, \|M_{W,F}(x)\|_2^{1/r}. \tag{64}$$

Moreover, there are constants $c_1, c_2 > 0$ such that

$$\|(H_{\mathrm{cyc}})_F(x)\|_2 \geq c_1 \|M_{W,F}(x)\|_2, \tag{65}$$

$$h_f\left(W + \sum_s x_s E_s\right) \geq c_2 \|M_{W,F}(x)\|_2^2 \tag{66}$$

for every positive walk scalar in Proposition C.3. Consequently, the two families have optimal error-bound exponents $1/r$ and $1/(2r)$, respectively.

Proof. For each $C \in \mathcal{C}_{\min}$, choose $i(C) \in \arg\min_{s\in C} |x_s|$ and set every selected coordinate to zero. This deletion set meets every minimal nonface, so the remaining candidate support lies in $\Delta_{W,F}$. If $\|x\|_\infty \leq 1$, the distance to this particular acyclic point satisfies

$$\operatorname{dist}(x, \mathcal{D}_F)^2 \leq \sum_{C\in\mathcal{C}_{\min}} \min_{s\in C} |x_s|^2 \tag{67}$$

$$\leq \sum_{C\in\mathcal{C}_{\min}} |(x^C)|^{2/|C|} \leq N \|M_{W,F}(x)\|_2^{2/r}. \tag{68}$$

The middle inequality is the minimum–geometric-mean inequality. The last uses $|C| \leq r$ and $|x^C| \leq 1$, proving Eq. (64).

For every minimal completion C, choose a simple directed cycle in $\operatorname{supp}(W) \cup C$. Minimality forces this cycle to use every candidate in C; its other edges have fixed nonzero weights in W. The corresponding component of H_{cyc} is therefore $b_C x^C$ with $b_C \neq 0$. Taking the minimum of finitely many $|b_C|$ proves Eq. (65). The same selected cycles occur with positive coefficients and squared edge weights in h_f, and all other walk terms are nonnegative. Their finite minimum gives Eq. (66).

Finally choose $C_\star \in \mathcal{C}_{\min}$ with $|C_\star| = r$ and set $x = t\mathbf{1}_{C_\star}$. Every other minimal generator contains a coordinate outside $C_\star$, while acyclicity requires at least one coordinate of $C_\star$ to be zero. Hence

$$\operatorname{dist}(x, \mathcal{D}_F) = |t|, \qquad \|M_{W,F}(x)\|_2 = |t|^r. \tag{69}$$

Proposition C.3, restricted to $C_\star$, gives scalar order $2r$. Thus neither exponent can be increased. □

C.8 Coordinate tubes and residual sublevel volume

The completion ideal also controls two volume notions that are not determined by q or r. Let $m = |F|$ and define

$$b_W(F) = m - \max_{S \in \Delta_{W,F}} |S| \\ = \min\{|T| : T \cap C \neq \varnothing \text{ for every } C \in \mathcal{C}_{\min}\}, \tag{70}$$

$$\lambda_W(F) = \min_{\substack{z \in \mathbb{R}^m_+ \\ \sum_{s \in C} z_s \geq 1,\ C \in \mathcal{C}_{\min}}} \sum_{s=1}^{m} z_s. \tag{71}$$

The first quantity is the height of $I_{W,F}$ and the minimum number of candidate edges whose removal from the full candidate graph makes it acyclic. The second is the fractional cover number of the minimal-completion hypergraph. In the standard algebraic terminology, Eq. (71) is the Newton-polyhedron formula for the log-canonical threshold of this monomial ideal (Howald, 2001, preprint, Main Theorem and Example 5).

Proposition C.5 (Exact coordinate-tube law)**.** *Choose $\rho > 0$ such that $B_\rho = [-\rho, \rho]^m \subset V_F$. For $0 \leq \varepsilon \leq \rho$,*

$$\frac{\mathrm{vol}\{x \in B_\rho : \mathrm{dist}_\infty(x, \mathcal{D}_F) \leq \varepsilon\}}{(2\rho)^m} \\ = \sum_{S \in \Delta_{W,F}} \left(1 - \frac{\varepsilon}{\rho}\right)^{|S|} \left(\frac{\varepsilon}{\rho}\right)^{m-|S|}. \tag{72}$$

In particular, its small-ε exponent is $b_W(F)$, and the leading coefficient after normalization by $(\varepsilon/\rho)^{b_W(F)}$ is the number of maximum-cardinality faces of $\Delta_{W,F}$.

Proof. For $x \in B_\rho$, set $S_\varepsilon(x) = \{s : |x_s| > \varepsilon\}$. If some $y \in \mathcal{D}_F$ satisfies $\|x - y\|_\infty \leq \varepsilon$, then every coordinate in $S_\varepsilon(x)$ is nonzero in y. Hence $S_\varepsilon(x) \subseteq \mathrm{supp}(y) \in \Delta_{W,F}$ and, by downward closure, $S_\varepsilon(x) \in \Delta_{W,F}$. Conversely, if $S_\varepsilon(x)$ is a face, retaining those coordinates and setting all others to zero produces an element of $\mathcal{D}_F$ within distance ε.

For a uniform point in B_ρ, the probability that $S_\varepsilon(x) = S$ is $(1-\varepsilon/\rho)^{|S|}(\varepsilon/\rho)^{m-|S|}$. The events are disjoint, so summing over the faces proves Eq. (72). Its smallest power of ε is $m - \max_{S \in \Delta_{W,F}} |S| = b_W(F)$; all coefficients at that power are positive. □

Theorem C.6 (Residual tolerance-volume exponent)**.** *Let $M = M_{W,F}$ be the minimal-generator vector and*

$$V_M(\eta) = \mathrm{vol}\{x \in B_\rho : \|M(x)\|_\infty \leq \eta\}.$$

Then

$$\lim_{\eta \downarrow 0} \frac{\log V_M(\eta)}{\log \eta} = \lambda_W(F). \tag{73}$$

Moreover, on a sufficiently small candidate box there are positive constants c_i, C_i such that

$$c_1 \|M(x)\|_\infty \leq \|(H_{\mathrm{cyc}})_F(x)\|_2 \leq C_1 \|M(x)\|_\infty, \tag{74}$$

$$c_2 \|M(x)\|_\infty^2 \leq h_f\left(W + \sum_s x_s E_s\right) \leq C_2 \|M(x)\|_\infty^2 \tag{75}$$

for every locally uniformly convergent positive-walk scalar in Proposition C.3. Consequently, the sublevel-volume exponents of the cycle vector and positive-walk scalar are respectively $\lambda_W(F)$ and $\lambda_W(F)/2$.

Equation (73) first determines the exponent for the minimal monomials. The two-sided comparisons in Eqs. (74)–(75) transfer it to the actual cycle-vector and positive-walk residuals. Constants and higher walk terms cannot change this logarithmic exponent; squaring the generator scale divides it by two. This supplies the residual-volume statement of Theorem B.5, beyond the monomial construction itself.

Proof. Fixed coordinate rescaling changes only constants, so first take $\rho = 1$. By symmetry it suffices to work on $[0,1]^m$. Put $x_s = e^{-u_s}$ and $L = \log(1/\eta)$. The coordinates u_s have independent unit-exponential density, and $\|M(x)\|_\infty \le \eta$ is equivalent to

$$u \ge 0, \qquad \sum_{s \in C} u_s \ge L \quad \text{for every } C \in \mathcal{C}_{\min}. \tag{76}$$

Let $z^\star$ solve Eq. (71). The orthant event $u_s \ge L z_s^\star$ for all s is contained in Eq. (76) and has probability $e^{-L\sum_s z_s^\star} = e^{-\lambda L}$.

For the reverse bound, linear-programming duality gives $y_C \ge 0$ satisfying

$$\sum_C y_C = \lambda, \qquad \sum_{C \ni s} y_C \le 1 \quad \text{for every } s. \tag{77}$$

On the event in Eq. (76),

$$\sum_s u_s \ge \sum_s u_s \sum_{C \ni s} y_C = \sum_C y_C \sum_{s \in C} u_s \ge \lambda L.$$

(The first inequality uses $u_s \ge 0$.) Since $\sum_s u_s$ has a Gamma$(m,1)$ distribution,

$$e^{-\lambda L} \le \Pr\{\|M(x)\|_\infty \le e^{-L}\} \le e^{-\lambda L} \sum_{j=0}^{m-1} \frac{(\lambda L)^j}{j!}. \tag{78}$$

Orthant factors and the fixed radius ρ do not affect the logarithmic limit. Dividing logarithms by $\log \eta = -L$ proves Eq. (73).

The lower inequalities in Eqs. (74) and (75) were proved in Proposition C.4. For the cycle-vector upper bound, every nonzero restricted simple-cycle product contains a minimal completion C among its candidate edges. On a box with $\|x\|_\infty \le 1$, its absolute value is at most a fixed base-edge coefficient times $|x^C|$. There are finitely many simple cycles, which gives C_1.

For the scalar upper bound, every nonzero closed-walk monomial likewise contains a minimal completion and is bounded by a fixed coefficient times $|x^C|^2$. More explicitly, choose a slightly larger candidate polydisc whose closure remains in the convergence domain and assign each walk to one minimal completion dividing its monomial. Absolute convergence on the larger polydisc implies absolute convergence of the finitely many termwise-divided series on the smaller box. Their suprema sum to a finite C_2. The two-sided comparisons sandwich each residual sublevel set between generator sublevel sets with constant-rescaled tolerances. Equation (73) then yields λ for the vector and $\lambda/2$ for the scalar. □

Remark C.7 (A neighborhood invariant, not one origin ray)**.** *An optimal fractional cover may have zero coordinates. The associated valuation then approaches a higher-dimensional acyclic support stratum inside B_ρ rather than the all-zero candidate point. This is intended: λ describes the volume of a neighborhood of the full local coordinate arrangement $\mathcal{D}_F \cap B_\rho$, whereas q and r are witnessed by paths into the distinguished base point.*

Corollary C.8 (Four scales of the completion ideal)**.** *The invariants satisfy*

$$\frac{b_W(F)}{r_W(F)} \le \lambda_W(F) \le b_W(F). \tag{79}$$

Here q is the first possible jet degree, r controls the worst Hölder error-bound exponent, b is the coordinate-tube codimension, and λ is the residual tolerance-volume exponent.

Proof. The indicator of an integral completion cover is feasible in Eq. (71), so $\lambda \le b$. Conversely, given any feasible z, the set $T = \{s : z_s \ge 1/r\}$ meets every minimal completion: otherwise some C of size at most r would satisfy $\sum_{s \in C} z_s < 1$. Thus $b \le |T| \le r \sum_s z_s$. Minimize over z. □

C.9 Statistical desingularization

Proof of Theorem B.6. Write $A = W \circ W$. Every term of $\mathrm{tr}(A^k)$ is a positive closed-walk monomial. After restricting to the absent coordinates F, its candidate degree is twice the number of candidate edges traversed with multiplicity. Theorem A.4 rules out candidate degree below $2q$. At degree $2q$, the candidate support of a closed walk contains a cycle completion of size at least q, so it has exactly q distinct candidate edges and traverses each once. Removing any one of them breaks all cycles in that support; otherwise a smaller completion would occur. Its candidate factor is therefore $\prod_{e\in C} x_e^2$ for a size-q minimal completion C.

Conversely, every minimal completion C contains a simple directed cycle in $\mathrm{supp}(W) \cup C$. The cycle uses every edge of C, since otherwise a proper subset would complete a cycle. Its fixed-edge product is nonzero, and the coefficient c_k at its length is positive. Summing all degree-$2q$ closed walks therefore gives

$$h_{f,F}(x) = \sum_{\substack{C\in\mathcal{C}_{\min}(W,F)\\ |C|=q}} a_C \prod_{e\in C} x_e^2 + o(\|x\|^{2q}), \qquad a_C > 0. \tag{80}$$

Let $Z_n = \sqrt{n}(\widehat{W}_n - W) \in \mathcal{Z}$. The diagonal remains zero, so no self-loop term enters. Substitution into Eq. (80) yields

$$n^q h_f(\widehat{W}_n) = P_{W,q}((Z_n)_F) + o_p(1). \tag{81}$$

The continuous mapping theorem gives Eq. (41). The fluctuations on fixed active edges perturb the coefficients a_C by $o_p(1)$ and hence disappear by Slutsky's theorem. For the local sequence, write $Z_n^{\mathrm{loc}} = \sqrt{n}(\widehat{W}_n - W_n)$. Under Eq. (42), $\sqrt{n}(\widehat{W}_n - W) = Z_n^{\mathrm{loc}} + U \Rightarrow Z + U$ under the changing laws P_{W_n}. Apply the same local expansion under those laws to obtain the shifted limit. An arbitrary cyclic perturbation is not an admissible local alternative in a model defined only on DAGs; the stated model-membership assumption is necessary.

Proof of Corollary B.7. The acceptance event is

$$n^q h_f(\widehat{W}_n) \le c n^{q-a}. \tag{82}$$

The right side tends to zero, c, or infinity according as $a > q$, $a = q$, or $a < q$. The limiting polynomial is positive almost surely because its coefficients are positive and each candidate coordinate of the Gaussian Z_F has positive variance, hence is nonzero almost surely. This proves Eq. (43). The degree-$2q$ initial form contains only size-q generators. A local shift supported on a disjoint size-r generator with $r > q$ leaves this initial form unchanged, which proves the final statement. Disjointness is sufficient, not necessary; without it, a shifted coordinate may also belong to a shortest generator.

D Symmetry and the Score–Topology Crossover

D.1 Coordinate realization and homology

Proof of Theorem B.2. Let $L = \mathrm{sd}\, K$. Its vertices are the nonempty faces of K, and its simplices are chains under inclusion. In particular, L is flag: a set of vertices is a face exactly when every pair is comparable. Let H be the graph whose vertices are those of L, with an edge between two incomparable faces. Write $\mathrm{Ind}(H)$ for the complex of vertex sets containing no edge of H. Then

$$\mathrm{Ind}(H) = L. \tag{83}$$

For each $i \in V(H)$, create graph nodes a_i, b_i and candidate edge $e_i : a_i \to b_i$. For every $\{i,j\} \in E(H)$, add fixed edges $b_i \to a_j$ and $b_j \to a_i$. Every fixed edge goes from the b-layer to the a-layer, so the fixed graph W_0 is acyclic.

If e_i, e_j are active for $\{i,j\} \in E(H)$, then

$$a_i \to b_i \to a_j \to b_j \to a_i \tag{84}$$

is a directed cycle. Conversely, every directed cycle must alternate between a candidate edge $a_i \to b_i$ and a fixed edge $b_i \to a_j$. Any such fixed edge certifies $\{i,j\} \in E(H)$, while the cycle certifies

that both candidates are active. Thus an active candidate set is acyclic if and only if it is independent in H. Equation (83) proves $\Delta_{W_0,F} \cong \operatorname{sd} K$. Radial projection identifies the positive normalized coordinate faces with the usual geometric realization, and barycentric subdivision preserves that realization.

Proof of Corollary B.3. Let $H_t = \bigsqcup_{b=1}^{t} K_5$ and apply the preceding construction. There are $5t$ candidates and $10t$ graph nodes. Independence complexes turn disjoint graph unions into simplicial joins: $A * B$ consists of unions of one face of A and one face of B, on disjoint vertex sets; A^{*t} is its t-fold join.

$$\operatorname{Ind}(H_t) = \operatorname{Ind}(K_5)^{*t}. \tag{85}$$

The complex $\operatorname{Ind}(K_5)$ is a set of five points, so its only nonzero reduced homology has rank four in degree zero. The reduced join formula

$$\widetilde{H}_r(A * B) \cong \bigoplus_{i+j=r-1} \widetilde{H}_i(A) \otimes \widetilde{H}_j(B) \tag{86}$$

shows inductively that the t-fold join has reduced homology of rank 4^t in degree $t-1$ and zero in all other degrees. Equivalently it is homotopy equivalent to a wedge of 4^t copies of S^{t-1}.

D.2 Far-field symmetry and invariant flows

Support-selection ties can also occur at nonzero edge weights. We construct them using a small ordinary-coordinate gadget. For each $i \in [m]$, introduce nodes a_i, b_i, a candidate edge $a_i \to b_i$, and unit-weight fixed cross-edges $b_i \to a_j$ for every $i \neq j$. The fixed graph is acyclic, while any two active candidates form a directed four-cycle. Taking t disjoint copies produces the family below. The same reduction also shows that local candidate sections can realize any finite homotopy type; Appendix B develops that broader topological consequence.

Call h *square-symmetric* if

$$h(\mathbf{W}) = \Phi(\mathbf{W} \circ \mathbf{W}), \qquad h(P\mathbf{W}P^\top) = h(\mathbf{W}) \tag{87}$$

for every permutation matrix P. NOTEARS, positive trace polynomials, and DAGMA are square-symmetric on their domains.

Write $x = (x_b)_{b=1}^{t}$, with $x_b \in \mathbb{R}^m$ the candidate weights in block b; $\mathbf{W}_0 + x$ abbreviates their coordinate embedding. For a residual $h : \Omega \to \mathbb{R}$, write $g(x) = h(\mathbf{W}_0 + x)$ on $V = \{x \in \mathbb{R}^{tm} : \mathbf{W}_0 + x \in \Omega\}$. A within-block redistribution direction varies the two active weights by opposite amounts and leaves all other coordinates fixed.

Theorem D.1 (Far-field support-selection ties)**.** *For every integer $m \geq 2$ and integer $t \geq 1$, there is a base DAG on $d = 2mt$ nodes with mt independent candidate entries and*

$$\binom{m}{2}^t \tag{88}$$

disjoint positive cyclic support manifolds. Let $h : \Omega \to \mathbb{R}$ be a C^1 square-symmetric exact scalar on an open domain $\Omega \subseteq \mathcal{Z}$. At every point of each manifold that lies in V,

$$\nabla_{x_b} g(x) \in \operatorname{span}\{x_b\}, \qquad b = 1, \ldots, t, \tag{89}$$

so the constraint gives zero first-order signal in every within-block edge-redistribution direction.

If ∇g is locally Lipschitz on V, every scalarized flow $\dot{x} = -\omega(x)\nabla g(x)$ with locally Lipschitz scalar $\omega : V \to \mathbb{R}$, when initialized on one of these manifolds in V, stays on it throughout the solution's maximal interval of existence in V. Positive active coefficients cannot reach zero at a finite time in that interval, so the support remains cyclic whenever the solution exists. With $m = 5$ there are $10^t = 2^{0.332\ldots d}$ such manifolds. Each contains a point with every active candidate coefficient equal to $1/2$. These points lie in the domains of NOTEARS and DAGMA with $s = 1$.

The radial-gradient identity in Eq. (89) allows a nonzero shrinking force while ruling out first-order preference between the two deletions. The flow statement strengthens this pointwise tie to an invariant trajectory when the vector field is locally Lipschitz. These are support-selection ties at positive

weights, distinct from the boundary derivative zeros in Theorem 2.1. An automorphism equates the active derivatives, evenness kills the inactive ones, and uniqueness preserves the manifold in the stated domain. When ∇g is L-Lipschitz, Corollary D.2 bounds the redistribution signal by $L\|x - \bar{x}\|_2$ for a nearby manifold point $\bar{x}$, so small perturbations yield a proportionally small signal rather than an immediate uniform ranking.

Proof of Theorem D.1. Use t disjoint coordinate gadgets for K_m. Write the candidate restriction as

$$g(x) = h\left(W_0 + \sum_{b=1}^{t}\sum_{i=1}^{m} x_{bi}E_{bi}\right). \tag{90}$$

For each block choose a pair $S_b = \{i_b, j_b\}$ and define

$$\mathcal{M}_S = \{x : \; x_{bi_b} = x_{bj_b} = a_b > 0, \; x_{b\ell} = 0 \text{ for } \ell \notin S_b, \; b \in [t]\}. \tag{91}$$

Each block contains the four-cycle in Eq. (84), hence every point of $\mathcal{M}_S$ is cyclic. The open set V contains a neighborhood of each point of $\mathcal{M}_S \cap V$. All derivative and flow statements below are restricted to this intersection.

Because $h(W) = \Phi(W \circ W)$, g is even in every candidate coordinate. Therefore

$$\partial_{b\ell} g(x) = 0 \qquad (\ell \notin S_b, \; x \in \mathcal{M}_S \cap V). \tag{92}$$

The node permutation exchanging simultaneously $a_{i_b} \leftrightarrow a_{j_b}$ and $b_{i_b} \leftrightarrow b_{j_b}$ is an automorphism of the K_m gadget. It fixes every point of $\mathcal{M}_S$. Relabeling invariance of h therefore gives

$$\partial_{bi_b} g(x) = \partial_{bj_b} g(x). \tag{93}$$

Equations (92)–(93) say exactly that the block gradient is proportional to x_b. There are $\binom{m}{2}$ choices in each block, and distinct choices have different algebraic supports, proving Eq. (88).

Now suppose ∇g and the scalar function ω are locally Lipschitz. The vector field $-\omega\nabla g$ satisfies the same inactive-zero and active-tie identities, hence is tangent to $\mathcal{M}_S \cap V$. The common amplitudes solve a locally Lipschitz system; on a boundary face lying in V, evenness gives

$$\dot{a}_b = F_b(a), \qquad F_b(0, a_{-b}) = 0. \tag{94}$$

Local uniqueness prevents a positive solution from reaching such an invariant face at a finite time within V. On any compact subinterval of its maximal interval of existence there is $L < \infty$ with $|F_b(a)| \le L|a_b|$. Gronwall's inequality yields

$$a_b(t) \ge a_b(0)e^{-Lt} > 0 \tag{95}$$

whenever $a_b(0) > 0$, on every compact subinterval of the solution's maximal interval of existence. No active coefficient reaches zero at a finite time in that interval, and every block retains its four-cycle while the solution exists.

For $m = 5$, $d = 10t$ and the number of manifolds is 10^t. For $A = W \circ W$, the nonzero eigenvalues of A^2 in block b are $\pm a_b^2$, so that block has spectral radius a_b. The blocks are disjoint, hence $\rho(A) = \max_b a_b$. DAGMA's $s = 1$ domain therefore intersects each $\mathcal{M}_S$ in $0 < a_b < 1$ for all b. In particular, the point with every $a_b = 1/2$ lies strictly inside the domain. The flow conclusion applies on the maximal interval in that domain, not outside it.

Corollary D.2 (Near-symmetry ranking bound)**.** *Let $\bar{x} \in \mathcal{M}_S \cap V$ for a manifold from Theorem D.1. Let $v \in \mathbb{R}^{tm}$ have $\|v\|_2 = 1$, $v_{bi_b} + v_{bj_b} = 0$ in each block, and all other coordinates zero. Suppose ∇g is L-Lipschitz on a convex neighborhood $N \subseteq V$ of $\bar{x}$. Then*

$$\left|\langle \nabla g(x), v\rangle\right| \le L\,\|x - \bar{x}\|_2$$

for every $x \in N$. If $\bar{x}$ attains $\operatorname{dist}(x, \mathcal{M}_S \cap V)$, the right side is $L \operatorname{dist}(x, \mathcal{M}_S \cap V)$.

Proof. Theorem D.1 gives $\langle \nabla g(\bar{x}), v\rangle = 0$. Lipschitz continuity and Cauchy–Schwarz yield

$$\begin{aligned}\left|\langle \nabla g(x), v\rangle\right| &= \left|\langle \nabla g(x) - \nabla g(\bar{x}), v\rangle\right| \\ &\le \|\nabla g(x) - \nabla g(\bar{x})\|_2 \, \|v\|_2 \\ &\le L\,\|x - \bar{x}\|_2 .\end{aligned}$$

The distance form follows when a nearest point lies in N. The vector v redistributes active weights; it is not a tangent to the equal-weight manifold itself. □

D.3 FULL-CYCLE SEPARATION AND SCORE PERTURBATIONS

The next proposition gives the separation time under feasibility-only flow, a bound under a positive pathwise score margin, and the effect of a shrinking score perturbation. The margin bound requires trajectory continuation. The perturbation result uses the specified quadratic score throughout the path; its initial gradient alone would not suffice.

Proposition D.3 (Full-cycle scalarization and score perturbation)**.** *Let $W(z)$ be supported on an isolated directed L-cycle with positive edge weights $z = (z_1, \dots, z_L)$. Suppose*

$$h(W(z)) = \phi(p), \qquad p = \prod_{i=1}^{L} z_i^2, \qquad \phi(0) = 0, \quad \phi'(0) > 0,$$

where ϕ is C^1 with $\phi' > 0$ on the relevant nonnegative range, including zero. Let Ψ be C^1 on the corresponding range, including zero, with $\Psi'(s) > 0$ for positive s and $\Psi'(s) = cs^\nu(1 + o(1))$ as $s \downarrow 0$, where $c > 0$ and $\nu \geq 0$. Under $\dot{z} = -\nabla_z(\Psi \circ h)$, initialize two tracked edges as $(x_0, y_0) = (a + \varepsilon, a - \varepsilon)$ and initialize the remaining edges at $b_j > a$. Put

$$C = \prod_{j=1}^{L-2} (b_j^2 - a^2), \qquad I_\nu(r) = \int_{r^2/(1-r^2)}^{\infty} [u(u+1)]^{-(\nu+1)}\, du.$$

For fixed $0 < r < 1$, let $T_r(\varepsilon)$ be the first time $y/x = r$. Here r denotes the ratio target, written ζ in the main text.

(i) Feasibility-only flow. *Every difference $z_i^2 - z_j^2$ is conserved and*

$$T_r(\varepsilon) \sim \frac{I_\nu(r)}{4c[\phi'(0)C]^{\nu+1}(4a\varepsilon)^{2\nu+1}}. \tag{96}$$

(ii) A pathwise score margin. *For a flow with a smooth score S, allow arbitrary positive tracked initial values with $y_0/x_0 > r$. For $\gamma_{\text{flow}} > 0$, set $T_\gamma = \gamma_{\text{flow}}^{-1} \log(y_0/(rx_0))$. Assume the trajectory can be continued in its positive domain through T_γ, unless $y/x = r$ is reached earlier. Define the total logarithmic deletion margin*

$$m_{\text{sc}}(z) = \frac{\partial_y[S + \Psi(h)]}{y} - \frac{\partial_x[S + \Psi(h)]}{x}.$$

If $m_{\text{sc}}(z) \geq \gamma_{\text{flow}} > 0$ until $y/x = r$, then

$$T_r \leq \frac{1}{\gamma_{\text{flow}}} \log \frac{y_0}{rx_0}, \tag{97}$$

(iii) A specified score perturbation. *Add the smooth edge score $S_\gamma(z) = \gamma y^2/2$, let $T_r(\varepsilon, \gamma)$ be the new hitting time, and write $T_0(\varepsilon) = T_r(\varepsilon, 0)$. For any positive schedule $\gamma = \gamma(\varepsilon)$,*

$$\left.\begin{array}{ll} \gamma T_0(\varepsilon) \to 0, & \nu > 0, \\ \gamma T_0(\varepsilon) \log(1/\varepsilon) \to 0, & \nu = 0 \end{array}\right\} \quad \Longrightarrow \quad \frac{T_r(\varepsilon, \gamma)}{T_0(\varepsilon)} \to 1, \tag{98}$$

whereas, for every $\nu \geq 0$,

$$\gamma T_0(\varepsilon) \to \infty \quad \Longrightarrow \quad \frac{T_r(\varepsilon, \gamma)}{T_0(\varepsilon)} \to 0. \tag{99}$$

Thus, when $\nu > 0$, the two regimes meet at the power scale

$$\gamma_c(\varepsilon) = T_0(\varepsilon)^{-1} = \Theta(\varepsilon^{2\nu+1}). \tag{100}$$

At $\nu = 0$, the scale $[T_0(\varepsilon) \log(1/\varepsilon)]^{-1} = \Theta(\varepsilon/\log(1/\varepsilon))$ is sufficient for asymptotically negligible score perturbations, while $\gamma T_0 \to \infty$ is sufficient for score domination. In general a logarithmic transition layer lies between these statements.

Part (i) gives the delay caused by a vanishing initial imbalance; part (ii) bounds that delay when the full objective maintains a positive logarithmic margin. Part (iii) asks when a shrinking quadratic score can still change the leading time scale. Its $\nu = 0$ condition needs the extra logarithm to

control entry into the limiting regime, so the pure-flow time alone does not determine whether that perturbation is negligible.

The need for the endpoint logarithm can be seen in the critical limiting system. Let $\Delta = x_0^2 - y_0^2$, $Q = x^2/\Delta$, and $V = y^2/\Delta$. At $\nu = 0$, use scaled time $c\phi'(0)C\Delta t$ and coefficient $\eta = \gamma/(c\phi'(0)C\Delta)$. This local normalization absorbs the leading flow coefficient; the proof below retains it explicitly. Primes denote derivatives in this scaled time. For the limiting system

$$Q' = -4QV, \qquad V' = -4QV - 2\eta V,$$

the quantity $V - Q - (\eta/2)\log Q$ is conserved. There are schedules with $\eta \to 0$ but $\eta \log(1/\Delta) \to \infty$ for which the normalized hitting time tends to zero. Thus the implication $\gamma T_0 \to 0 \Rightarrow T_r/T_0 \to 1$ is false at $\nu = 0$ without additional entrance control.

The pathwise quantity γ_{flow} and the completion-exchange margin $\kappa_{\mathcal{C}}$ introduced below measure different objects. The former lower-bounds a gradient ratio along one trajectory; the latter is a discrete score gap between repairs. Neither bounds the other without additional assumptions on the score and path.

Conserved squared differences also occur in homogeneous models (Du et al., 2018). On the exact DAG-cycle restriction, this invariant yields the scalarization-dependent hitting-time law and its score-margin counterpart. For an L-cycle, $\phi_{\text{EXP}}(p) = L\sum_{\ell \geq 1} p^\ell/(\ell L)!$, so $\phi'_{\text{EXP}}(0) = 1/(L-1)!$, whereas $\phi_{\text{DAGMA}}(p) = -\log(1-p)$ and $\phi'_{\text{DAGMA}}(0) = 1$. Thus the proposition applies exactly to both constraints. Linear feasibility and seeded ALM have $\nu = 0$, while a cold quadratic penalty has $\nu = 1$, yielding the feasibility-only exponents $1, 3, 1$. The logarithmic endpoint qualification applies to the first and third scalarizations.

Proof of Proposition D.3. Put $p = \prod_{i=1}^{L} z_i^2$. On the positive orthant,

$$\partial_{z_i} h = 2p\phi'(p)/z_i, \qquad \dot{z}_i = -2\Psi'(\phi(p))p\phi'(p)/z_i. \tag{101}$$

It follows immediately that

$$\frac{d}{dt} z_i^2 = -4\Psi'(\phi(p))p\phi'(p) \tag{102}$$

is independent of i. Hence every $z_i^2 - z_j^2$ is conserved. In particular, $\Delta = x^2 - y^2 = 4a\varepsilon$. Conservation of squared-norm differences is a familiar balancing mechanism in homogeneous gradient flows (Du et al., 2018); the calculation below uses the exact cycle product to obtain the sharp fixed-ratio edge-separation time.

Write $u = y^2$ and

$$d_{j,\varepsilon} = b_j^2 - (a-\varepsilon)^2, \qquad p_\varepsilon(u) = u(u+\Delta)\prod_{j=1}^{L-2}(u + d_{j,\varepsilon}).$$

Equation (102) becomes

$$\dot{u} = -4\Psi'(\phi(p_\varepsilon(u)))p_\varepsilon(u)\phi'(p_\varepsilon(u)). \tag{103}$$

The event $y/x = r$ occurs at $u_r = r^2\Delta/(1-r^2)$. With $u = \Delta v$, separation of variables gives

$$T_r(\varepsilon) = \frac{\Delta}{4}\int_{v_r}^{(a-\varepsilon)^2/\Delta} \frac{dv}{\Psi'(\phi(p_\varepsilon(\Delta v)))p_\varepsilon(\Delta v)\phi'(p_\varepsilon(\Delta v))}, \qquad v_r = \frac{r^2}{1-r^2}. \tag{104}$$

For fixed v,

$$p_\varepsilon(\Delta v) = \Delta^2 v(v+1)\prod_{j=1}^{L-2}(\Delta v + d_{j,\varepsilon}), \qquad d_{j,\varepsilon} \longrightarrow b_j^2 - a^2.$$

Let $\alpha = \phi'(0)$ and $C = \prod_j (b_j^2 - a^2)$. The assumptions on Ψ' and ϕ' imply

$$\Psi'(\phi(p))p\phi'(p) = c\alpha^{\nu+1}p^{\nu+1}(1 + o(1)).$$

After multiplying Eq. (104) by $\Delta^{2\nu+1}$, its integrand therefore converges pointwise to

$$\frac{1}{4c(\alpha C)^{\nu+1}[v(v+1)]^{\nu+1}}.$$

To justify taking the limit with a moving integration endpoint, split the integral at a small fixed $u_0 > 0$. On $u \in [0, u_0]$, the asymptotic relations above hold with two-sided uniform bounds, and the normalized integrand is dominated by a constant multiple of $[v(v+1)]^{-(\nu+1)}$, which is integrable on $[v_r, \infty)$. On $u \in [u_0, (a-\varepsilon)^2]$, all factors in Eq. (103) are bounded away from zero, so that part of the unnormalized hitting time is $O(1)$ and vanishes after multiplication by $\Delta^{2\nu+1}$. Dominated convergence now yields

$$\lim_{\varepsilon \downarrow 0} \Delta^{2\nu+1} T_r(\varepsilon) = \frac{I_\nu(r)}{4c[\alpha C]^{\nu+1}}.$$

Substitution of $\Delta = 4a\varepsilon$ proves Eq. (96).

For the score statement, the full flow satisfies

$$\frac{d}{dt} \log \frac{y}{x} = -\left(\frac{\partial_y [S + \Psi(h)]}{y} - \frac{\partial_x [S + \Psi(h)]}{x} \right) = -m_{\mathrm{sc}}(z). \tag{105}$$

If $m_{\mathrm{sc}}(z) \geq \gamma_{\mathrm{flow}}$ until the stopping event, integration gives $\log(y(t)/x(t)) \leq \log(y_0/x_0) - \gamma_{\mathrm{flow}} t$ before that event. If no crossing occurred by T_γ, the assumed continuation would make the ratio defined and continuous on $[0, T_\gamma]$, while the integrated bound would give $y(T_\gamma)/x(T_\gamma) \leq r$. The intermediate value theorem then supplies a crossing, a contradiction. This proves Eq. (97) without assuming that the preferred edge is initially smaller. Smoothness of S alone does not supply continuation. With $h = x^2 y^2$, $\Psi(s) = s$, and $S = J - h$, the total objective $J = -(x^2+y^2)^2/4 + y^2/2$ has margin one, but from $x_0 = y_0 = 1$ its solution satisfies $x(t)^{-2} = e^{-2t} - 2t$ and $y(t)/x(t) = e^{-t}$. It explodes before $t = 1/2$, hence before reaching ratio $1/2$.

For the canonical score $S_\gamma = \gamma y^2/2$, continuation follows from the coordinate equations. Write $f(p) = \Psi'(\phi(p))\phi'(p)$. Its continuous extension to $[0, p_0]$, where $p_0 = \prod_i z_i(0)^2$, is bounded and nonnegative. Each coordinate obeys

$$\dot{z}_i = -a_i(z) z_i, \qquad a_i(z) = 2f(p) \prod_{j \neq i} z_j^2 + \gamma \mathbf{1}\{i = y\}.$$

The coordinates are nonincreasing, so throughout the positive solution $0 \leq a_i(z) \leq K$ for a finite constant depending on its initial box and γ. Thus $z_i(0) e^{-Kt} \leq z_i(t) \leq z_i(0)$. On every finite time interval the trajectory stays in a compact subset of the positive domain; its bounded vector field gives a limit at any finite endpoint, and local existence there extends the solution. The continuation condition therefore holds for this score. The corresponding Lean check uses a path defined on the full comparison interval; it verifies the crossing argument, not a general ODE-existence theorem.

It remains to prove the crossover claim for $S_\gamma(z) = \gamma y^2/2$. Put

$$A(p) = \Psi'(\phi(p)) p \phi'(p).$$

The squared-weight equations become

$$\frac{d}{dt} x^2 = -4A(p), \qquad \frac{d}{dt} y^2 = -4A(p) - 2\gamma y^2, \tag{106}$$

while every unpenalized cycle edge has the same squared speed as x. The gap $x^2 - y^2$ starts positive and has derivative $2\gamma y^2$, so it stays positive. Consequently,

$$\frac{d}{dt} \log \frac{y}{x} = -\gamma - 2A(p) \left(\frac{1}{y^2} - \frac{1}{x^2} \right) \leq -\gamma. \tag{107}$$

This proves

$$T_r(\varepsilon, \gamma) \leq \gamma^{-1} \log \frac{y_0}{r x_0},$$

and therefore Eq. (99) whenever $\gamma T_0 \to \infty$.

For the other regime, let $\Delta = x_0^2 - y_0^2 = 4a\varepsilon$, write $q = x^2$, $v = y^2$, and note that every other squared cycle weight equals $q + d_{j,\varepsilon}$ for a constant $d_{j,\varepsilon} \to b_j^2 - a^2 > 0$. On the bottleneck scale

$$q = \Delta Q, \qquad v = \Delta V, \qquad \tau = \Delta^{2\nu+1} t,$$

Eq. (106) has, on compact subsets of the positive (Q, V) quadrant, the uniform limit

$$\frac{dQ}{d\tau} = -4c[\phi'(0)C]^{\nu+1}(QV)^{\nu+1}, \tag{108}$$

$$\frac{dV}{d\tau} = -4c[\phi'(0)C]^{\nu+1}(QV)^{\nu+1} - 2\eta V, \tag{109}$$

where $\eta = \gamma/\Delta^{2\nu+1}$. The pure-flow calculation above shows that $T_0\Delta^{2\nu+1}$ converges to a finite positive constant. Thus $\gamma T_0 \to 0$ is equivalent, up to a convergent positive factor, to $\eta \to 0$.

The entrance from $Q, V \to \infty$ is where the cases $\nu > 0$ and $\nu = 0$ differ. Fix M, and let τ_M be the first scaled time at which $Q = M$. Before the target ratio is reached, write the exact scaled system as

$$Q' = -4K_\varepsilon Q^{\nu+1}V^{\nu+1}, \qquad V' = -4K_\varepsilon Q^{\nu+1}V^{\nu+1} - 2\eta V,$$

where the scalar-speed assumptions give $K_\varepsilon \geq k > 0$ for all sufficiently small ε. With $D = Q - V$ and $V/Q \geq r^2$, a change of variables from scaled time to the decreasing coordinate Q yields, for $\nu > 0$,

$$0 \leq D(\tau_M) - 1 = \int_M^{Q(0)} \frac{\eta}{2K_\varepsilon Q^{\nu+1}V^\nu}\, dQ \leq \frac{\eta}{4\nu k r^{2\nu}} M^{-2\nu}, \tag{110}$$

$$\tau_M \leq \frac{1}{4kr^{2\nu+2}(2\nu+1)} M^{-(2\nu+1)}. \tag{111}$$

Thus the entrance state converges to $(M, M-1)$ for every fixed M, and the time omitted above that section vanishes uniformly as $M \to \infty$.

At $\nu = 0$, write the exact scaled equations before the section as

$$Q' = -4K_\varepsilon QV, \qquad V' = -4K_\varepsilon QV - 2\eta V,$$

where the scalar-speed assumptions give $K_\varepsilon \geq k > 0$ for all sufficiently small ε. With $D = Q - V$,

$$D' = 2\eta V, \qquad -(\log Q)' = 4K_\varepsilon V \geq 4kV.$$

Since $D(0) = 1$, integration to τ_M gives the entrance estimate

$$0 \leq D(\tau_M) - 1 \leq \frac{\eta}{2k} \log \frac{Q(0)}{M}. \tag{112}$$

Here $Q(0) = x_0^2/\Delta = \Theta(\Delta^{-1})$. The condition $\gamma T_0 \log(1/\varepsilon) \to 0$ is therefore equivalent at this scale to $\eta \log(1/\Delta) \to 0$, and Eq. (112) again sends the entrance state to $(M, M-1)$. Moreover, while $V/Q \geq r^2$, $-Q' \geq 4kr^2Q^2$, so $\tau_M \leq (4kr^2M)^{-1}$. This entrance envelope vanishes as $M \to \infty$.

On the compact corridor below $Q = M$, uniform convergence of the vector fields, continuous dependence of solutions, and transversality of the target ratio give convergence to the $\eta = 0$ hitting time. Taking first $\varepsilon \downarrow 0$ and then $M \to \infty$ proves Eq. (98). For $\nu > 0$, combining this result with Eq. (96) gives the power scale in Eq. (100).

The additional logarithm at $\nu = 0$ cannot be removed. In the constant critical system $K_\varepsilon \equiv 1$, direct differentiation gives the first integral

$$V - Q - \frac{\eta}{2} \log Q = \text{constant}.$$

For example, take $\Delta_n = \exp(-(n+1)^2)$ and $\eta_n = (n+1)^{-1}$. Then $\eta_n \to 0$, but $\eta_n \log(1/\Delta_n) \to \infty$. The first integral forces the target section to be reached at $Q_{\mathrm{hit,n}} \to \infty$; since $T_n \leq (4r^2 Q_{\mathrm{hit,n}})^{-1}$, the normalized hitting time tends to zero rather than to the positive feasibility-only limit. This proves the claimed endpoint failure under the weaker condition $\gamma T_0 \to 0$.

Finally, on an isolated L-cycle, the only closed walks have lengths ℓL, and $\mathrm{tr}((W \circ W)^{\ell L}) = Lp^{\ell}$. Hence

$$\phi_{\mathrm{EXP}}(p) = L \sum_{\ell \geq 1} \frac{p^{\ell}}{(\ell L)!}, \qquad \phi'_{\mathrm{EXP}}(0) = \frac{1}{(L-1)!}.$$

The same cycle block satisfies $\det(I - W \circ W) = 1 - p$, so $\phi_{\mathrm{DAGMA}}(p) = -\log(1-p)$ and $\phi'_{\mathrm{DAGMA}}(0) = 1$. The linear scalarization has $\Psi'(s) = 1$, the cold quadratic penalty has $\Psi'(s) = s$, and a seeded ALM term $\Psi(s) = \mu s + \rho s^2/2$ with $\mu > 0$ has $\Psi'(s) \sim \mu$. □

Proof of Corollary B.4. Use one K_m gadget. For every $S \subseteq [m]$ of size $k \geq 2$, set

$$x_i = \frac{1}{\sqrt{k}} \mathbf{1}\{i \in S\}. \tag{113}$$

All active coordinates are related by automorphisms and all inactive derivatives vanish by evenness. This is the elementary finite-group case of the symmetric-criticality principle (Palais, 1979). Thus $\nabla g(x) = \lambda x$, the Lagrange equation on the unit sphere. Every such support is cyclic, and there are $\sum_{k=2}^{m} \binom{m}{k} = 2^m - m - 1$ of them.

Giving each active coordinate an arbitrary sign preserves the Lagrange equation because g is coordinatewise even. This gives

$$\sum_{k=2}^{m} \binom{m}{k} 2^k = 3^m - 2m - 1 \tag{114}$$

signed critical points. For DAGMA, with $A = W \circ W$, the active block of A^2 is $k^{-1}(J_k - I_k)$. Thus $\rho(A) = \sqrt{(k-1)/k} < 1$.

D.4 Controlled cycle flows: score–topology audit

The controlled cycle-flow setting compares EXP and DAGMA residuals under specified gradient flows, rather than their learning algorithms. The deterministic crossover grid uses cycle lengths $4, 6, 8$, $\nu \in \{0, 1\}$, $a = 0.4$, and target ratio $1/4$. It crosses four imbalances $\varepsilon \in \{3, 6, 12, 24\} \times 10^{-4}$ with ten scaled margins $\gamma T_0 \in \{0, 0.01, 0.03, 0.1, 0.3, 1, 3, 10, 30, 100\}$, giving 480 runs. No sampling seeds enter this grid. The full-matrix and SEM-score variants use the protocols in Appendix F.8; the latter compares population covariances with sample covariances at three sample sizes and seed indices 0–19.

Figure 5 reports the full-matrix checks that complement the normalized crossover experiment in Figure 4. The population score has a positive pathwise margin in every audited setting and selects the false back edge even when that edge starts larger. Sampling can reverse this margin, which is why the finite-data experiment does not attain perfect selection. The flow experiment fixes the score to isolate its dynamical effect. The confidence set in Section 4 addresses the separate uncertainty in an estimated score's preference.

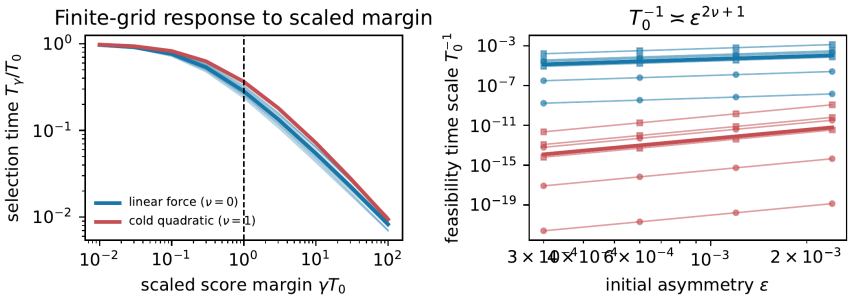


Figure 4: Finite-range score perturbation audit. The normalized hitting time separates around $\gamma T_0 = 1$ on the plotted grid (left), and the fitted power scale follows $T_0^{-1} \asymp \varepsilon^{2\nu+1}$ (right). For $\nu = 0$, the proposition gives the additional logarithmic entrance condition.

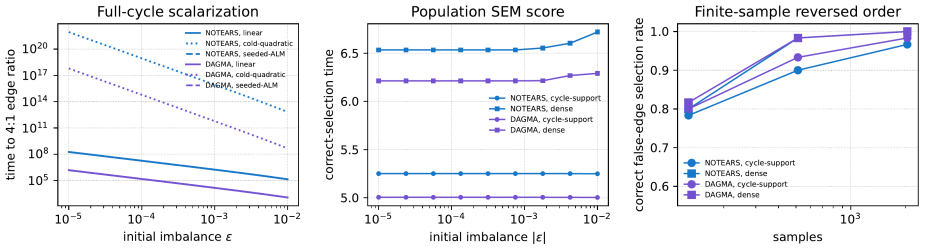


Figure 5: Full-matrix checks of Proposition D.3. Feasibility-only linear and seeded-ALM flows scale as ε^{-1}, while the cold quadratic scales as ε^{-3} (left). A population linear-SEM score removes the divergence for reversed initial rankings and dense flows (center). The same event is not guaranteed from an empirical covariance, although its frequency increases with sample size (right).

D.5 Hard-threshold support stability

In this subsection, the matrix norm $\|\Delta\|_\infty$ is the largest absolute entry, not the maximum row sum.

Proposition D.4 (Stability of hard-threshold support)**.** *For $\theta > 0$, define*

$$\mathcal{S}_\theta(W) = \{(i,j) : i \neq j,\ |W_{ij}| > \theta\}, \qquad m_\theta(W) = \min_{i \neq j} \big||W_{ij}| - \theta\big|. \tag{115}$$

The support $\mathcal{S}_\theta(W)$ is unchanged for every off-diagonal perturbation $\|\Delta\|_\infty < m_\theta(W)$; conversely, for every $\eta > 0$ some Δ with $\|\Delta\|_\infty \leq m_\theta(W) + \eta$ changes it. Moreover, a positive diagonal rescaling preserves algebraic support but need not preserve $\mathcal{S}_\theta(W)$.

For a data matrix $X \in \mathbb{R}^{n \times d}$, X_j is response column j and X_P contains the columns in a parent set P. Let Π_P be orthogonal projection onto their span and $\mathrm{RSS}_j(P) = \|(I - \Pi_P)X_j\|_2^2$ the least-squares residual sum of squares. A complete cycle oracle returns a directed cycle whenever one exists, and reports acyclicity otherwise.

Proposition D.5 (Scale invariance of profile-score event paths)**.** *For centered data X, define*

$$\Delta_{i \to j}(P) = \tfrac{1}{2} \log \frac{\mathrm{RSS}_j(P \setminus \{i\})}{\mathrm{RSS}_j(P)}. \tag{116}$$

Under $X' = XD$ with positive diagonal D, every positive-RSS deletion score is unchanged. Start from the same algebraic support and use the same deterministic complete cycle oracle, depending only on the current support, with a fixed edge order to break score ties. If all queried residual sums of squares are positive, repeatedly deleting a minimum-score edge of the returned cycle produces the same event path in both units.

We prove Proposition D.4. For every off-diagonal coordinate, the reverse triangle inequality gives

$$\big||W_{ij} + \Delta_{ij}| - |W_{ij}|\big| \leq |\Delta_{ij}| < m_\theta(W) \leq \big||W_{ij}| - \theta\big|. \tag{117}$$

Thus $|W_{ij} + \Delta_{ij}| - \theta$ has the same sign as $|W_{ij}| - \theta$ for every coordinate, and the strict threshold support is unchanged.

Conversely, choose a coordinate attaining the finite minimum $m = m_\theta(W)$. If $|W_{ij}| > \theta$, perturb opposite to its sign by $m + \eta'$; if $|W_{ij}| \leq \theta$, perturb along its sign (choosing either sign at zero) by $m + \eta'$, where $0 < \eta' \leq \eta$ is small enough to avoid crossing zero in the first case. The chosen coordinate crosses the strict threshold and the perturbation norm is at most $m + \eta$. Hence m is the supremum of open ℓ_∞ balls on which support is constant. Taking one coefficient equal to θ shows $\inf_W m_\theta(W) = 0$.

For the scaling claim, let $D = \mathrm{diag}(d_1, \ldots, d_d)$ with $d_i > 0$. Then

$$(D^{-1}WD)_{ij} = \frac{d_j}{d_i} W_{ij}. \tag{118}$$

Every multiplier is positive, so exact zero support, directed cycles, and acyclicity are preserved. For any selected nonzero coefficient, however, d_j/d_i can be any positive number. Choosing it above or

below $\theta/|W_{ij}|$ moves that coefficient across the fixed threshold. This proves non-invariance. If standardization is prescribed and repeated after rescaling, this particular source of variation disappears within that fixed pipeline. The proposition makes the narrower statement that a numerical cutoff depends on the chosen coordinate convention, and that its robustness within any such convention is governed by $m_\theta(W)$.

We next prove Proposition D.5. Let P be a parent set for node j, and let Π_P denote orthogonal projection onto the column span of X_P. Under $X' = XD$, the response becomes $X'_j = d_j X_j$, while the design becomes $X'_P = X_P D_P$. Because D_P is invertible, $\mathrm{col}(X'_P) = \mathrm{col}(X_P)$ and hence the two designs have the same projection operator. Therefore

$$\mathrm{RSS}'_j(P) = \|(I - \Pi_P) d_j X_j\|_2^2 = d_j^2 \mathrm{RSS}_j(P). \tag{119}$$

The same factor d_j^2 multiplies the numerator and denominator of Eq. (116), so every positive-RSS deletion score is unchanged. At the initial support, the same deterministic support-only oracle returns the same cycle in both runs. Equality of its deletion scores and the fixed edge order imply the same first deletion. If the supports agree after any number of deletions, the oracle again returns the same cycle, and the scores and tie breaking give the same next deletion. Induction therefore gives the same event path and final DAG. Completeness detects acyclicity, and deletion terminates because the support is finite. This proves the proposition. Its starting point is a fixed algebraic support; rescaling invariance of a statistical frontend is a separate property of that frontend.

D.6 Relation to prior work and alternative assumptions

Table 3 compares the assumptions and conclusions of the closest results. The middle column identifies the tools and guarantees already available. The right column locates the present results in completion geometry, cycle dynamics, and structural score certification.

Table 3: Closest prior results and the boundary of the present claims.

Work	Established there	Distinction here
Hultman (2004)	The simplicial complex of acyclic edge sets and its homotopy type.	We condition on a fixed base DAG and restrict to candidate coordinates. These local sections realize arbitrary finite homotopy types, despite the restricted homotopy type of the full complex.
Ehrenborg & Hetyei (2006)	Every finite complex becomes an independence complex after barycentric subdivision.	This fact is an ingredient in an ordinary-adjacency reduction from graph independence complexes to local DAG sections; the reduction yields the homology and gradient consequences used here.
Howald (2001)	Multiplier ideals and log-canonical thresholds of monomial ideals from their Newton polyhedra.	The classical fractional-cover formula, together with an elementary sublevel-volume sandwich, identifies the continuous DAG tolerance exponent.
Drton & Xiao (2016); Sturma et al. (2024)	Nonstandard inference for singular polynomial restrictions and bootstrap tests for many irregular constraints.	Theorem B.6 supplies the problem-specific part: the exact singular degree and initial polynomial determined by DAG completion topology.
Wei et al. (2020)	Gradient degeneracy for positive polynomial constraints; absolute-value KKT conditions and local search.	We cover arbitrary smooth exact maps on candidate subspaces, identify the topology-dependent q versus $2q$ orders, and bound vector output dimension. Their local search is prior art for the hybrid category.

Continued on next page

Table 3: Closest prior results and the boundary of the present claims (continued).

Work	Established there	Distinction here
Ng et al. (2022)	Failure of standard ALM regularity conditions; feasibility of QPM limit points under stated assumptions.	Our sharp response and distance bounds concern exact representations on candidate sections, independently of a solver's convergence conditions.
Zhang et al. (2025)	Positive-coefficient analytic functions that yield exact DAG constraints, closure rules for constructing new constraints, and efficient evaluation.	These families furnish examples to which the representation-level lower orders and edge-ranking obstructions apply on the candidate restrictions considered here.
Bello et al. (2022)	An exact M-matrix log-determinant constraint, invexity, a reachability-dependent Hessian, and central-path optimization.	We prove exact $2q$ local order and exponentially many nonzero cyclic strata with zero tangential edge-ranking signal. These are not Euclidean stationary points, so the conclusion is compatible with invexity.
Du et al. (2018)	Squared-norm differences are conserved by gradient flow in homogeneous multi-layer models.	On an exact DAG-cycle restriction, this invariant reduces the flow to a one-dimensional product dynamics and yields the scalarization-dependent $\varepsilon^{-(2\nu+1)}$ edge-separation law and the score-margin escape condition.
Shridharan & Iyengar (2025)	Higher-order Taylor approximations for $h(\lvert W\rvert^{\beta})$, used to construct score-improving local search directions.	We derive the completion-dependent order $q_W(F)$ for arbitrary smooth exact maps, rather than a specified function family, and separate vector from nonnegative scalar maps with a dimension lower bound.
Ng et al. (2024)	Empirical and optimization consequences of sparsity, nonconvexity, and thresholding.	We prove that no positive continuous tolerance certifies support and quantify the local detection order before a threshold is applied.
Yu et al. (2021); Massidda et al. (2024); Gillot & Parviainen (2022)	Direct acyclic parameterization/orientation and continuous–feedback-arc-set hybrid optimization.	These methods lie outside a fixed smooth exact scalar. The invariant-flow result explains why a nonsmooth, discrete, or score-asymmetric mechanism can supply support-selection information absent from the feasibility gradient.
Rey et al. (2026)	A constraint without the Hadamard square; population landscape guarantees under the stated nonnegative-SEM, identifiability, and multiplier assumptions.	Sparse DAGs lie on the boundary of the nonnegative orthant, so the open signed-neighborhood assumption used here does not apply.
Hansen et al. (2011)	Confidence sets retain the best models without assuming a true candidate model.	We specialize score uncertainty to a frozen parent-set family and query shared graph features; set-valued model comparison is established prior work.

Continued on next page

Table 3: Closest prior results and the boundary of the present claims (continued).

Work	Established there	Distinction here
Park et al. (2026)	Split-sample relative-fit tests for projections under misspecification, with construction-specific targets and coverage conditions.	Simultaneous local intervals yield a robust DAG objective, permit holdout selection of the reference, and certify features shared by all score optima.

The bounds apply to local Taylor information in an exact representation near a DAG support boundary. Algorithms that use higher derivatives, change the parameterization, modify support discretely, or stop before exact support fall outside the corresponding assumptions.

D.7 Scalarization and augmented Lagrangians

For $p \geq 1$, define $P_p(W) = \|H(W)\|_2^p$. In an augmented Lagrangian (ALM), λ and μ below are respectively vector and scalar multipliers, and $\rho > 0$ is the penalty coefficient. A cold start sets the multiplier to zero; a seeded start uses a fixed nonzero multiplier. These multipliers are distinct from the fractional-cover invariant $\lambda_{\mathbf{W}}(F)$.

Proposition D.6 (Sensitivity–smoothness frontier)**.** *If $H(W + tU) = t^q v + o(t^q)$ with $v \neq 0$, then $P_p(W + tU) = \|v\|_2^p |t|^{pq} + o(|t|^{pq})$. As a function of the residual, $z \mapsto \|z\|_2^p$ is nonsmooth at zero for $p = 1$, is C^1 but has non-Lipschitz gradient for $1 < p < 2$, and has locally Lipschitz gradient for $p \geq 2$.*

Corollary D.7 (Augmented-Lagrangian hierarchy)**.** *For the vector ALM $Q_{\lambda,\rho} = \lambda^\top H + \rho \|H\|^2 /2$, the zero-multiplier initialization vanishes to order at least $2q$ along an equal-scale completion; on a q-sharp path this order is exact, while a generic nonzero multiplier exposes order q. For a nonnegative scalar ALM $Q_{\mu,\rho} = \mu h + \rho h^2/2$, the corresponding lower bounds are $4q$ for $\mu = 0$ and $2q$ for $\mu \neq 0$, with equality when h has exact order $2q$. In these sharp cases, derivative orders are one lower.*

Corollary D.8 (Order under residual-scaled multipliers)**.** *Let $R(t) = O(t^p)$ and $DR(t) = O(t^{p-1})$. If $\lambda(t) = O(t^p)$ and $\rho = O(1)$, then the primal gradient of $\lambda^\top R + \rho \|R\|^2 /2$, treating the multiplier as fixed within the primal step, is $O(t^{2p-1})$. Any fixed number of updates $\lambda \leftarrow \lambda + \rho R$ from zero, with bounded penalties and residual-comparable iterates, preserves this lower bound. Recovering order $p - 1$ requires a multiplier that does not vanish with the local residual and is not orthogonal to its leading coefficient.*

Residual-power frontier. If $H(t) = t^q v + o(t^q)$, continuity of the norm and homogeneity give

$$\|H(t)\|_2^p = |t|^{pq} \|v + o(1)\|_2^p = \|v\|_2^p |t|^{pq} + o(|t|^{pq}). \tag{120}$$

For $z \neq 0$, $\nabla \|z\|_2^p = p \|z\|_2^{p-2} z$. At zero, the norm is nonsmooth for $p = 1$; for $1 < p < 2$ this gradient extends continuously but its derivative scales as $\|z\|^{p-2}$ and is unbounded; for $p \geq 2$ it is locally Lipschitz. These are residual-space statements. Composition with a particular high-order H can add pathwise regularity, but it cannot provide a uniform Lipschitz guarantee over arbitrary residual directions.

Assume a sharp path $H(t) = t^q v + o(t^q)$ with $v \neq 0$. Then

$$\frac{1}{2} \|H(t)\|_2^2 = \frac{1}{2} \|v\|_2^2 t^{2q} + o(t^{2q}). \tag{121}$$

More generally, Theorem A.5 always yields the lower-order bound $O(t^{2q})$. For $Q_{\lambda,\rho} = \lambda^\top H + \rho \|H\|_2^2/2$, zero initialization $\lambda = 0$ therefore removes the order-q term. A multiplier with $\lambda^\top v \neq 0$ restores it; this is the meaning of "generic" in Corollary D.7.

For a sharp nonnegative scalar $h(t) = a t^{2q} + o(t^{2q})$, $a > 0$,

$$Q_{\mu,\rho}(t) = \mu a t^{2q} + \frac{\rho}{2} a^2 t^{4q} + \mu\, o(t^{2q}) + o(t^{4q}). \tag{122}$$

Thus $\mu = 0$ has order $4q$, while every nonzero μ has order $2q$. Differentiation lowers each sharp order by one. If the representation has higher-than-minimal order, these statements remain lower bounds on the first possible nonzero term.

Residual-generated multipliers. Let $R(t) = O(t^p)$ and $DR(t) = O(t^{p-1})$. During a primal update the multiplier is fixed, so

$$\nabla_t Q_{\lambda,\rho} = DR(t)^\top \lambda + \rho DR(t)^\top R(t). \tag{123}$$

If $\lambda = O(t^p)$ and $\rho = O(1)$, both terms are $O(t^{2p-1})$. Starting from zero, a fixed number of updates $\lambda_{s+1} = \lambda_s + \rho_s R(t_s)$ with $t_s = \Theta(t)$ and bounded ρ_s leaves $\lambda_s = O(t^p)$ by induction. An order-$p-1$ term can appear only if the multiplier has a nonvanishing leading scale and nonzero inner product with the leading residual coefficient. Substituting $p = q$ for a vector representation and $p = 2q$ for a nonnegative scalar proves Corollary D.8. A penalty schedule that diverges as $t \to 0$ can counteract the coefficient, but then does so by ill-conditioning, not by restoring a missing Taylor term.

D.8 Finite precision and restricted correction time

If $|R(t)| \le C|t|^p$, then $|R(t)| < \varepsilon$ whenever

$$|t| < r_\varepsilon := \left(\frac{\varepsilon}{C}\right)^{1/p}. \tag{124}$$

This interval follows directly from the stated bound. Its numerical value depends on the formulation-specific constant C and on the comparison rule. For fixed C and $0 < \varepsilon < C$, the radius increases to one as $p \to \infty$.

Here ϕ is a one-dimensional penalty as a function of the positive candidate scale t. Let $T(\varepsilon)$ be the time for $\dot{t} = -\phi'(t)$ to move from fixed small $t_0 > 0$ to $0 < \varepsilon < t_0$. For this calculation, assume explicitly that

$$\phi'(t) = apt^{p-1} + o(t^{p-1}), \qquad a > 0, p \ge 2, t \downarrow 0. \tag{125}$$

For sufficiently small t, $\phi'(t)$ is bounded above and below by positive constants times t^{p-1}. Separating variables in $\dot{t} = -\phi'(t)$ gives

$$T(\varepsilon) \asymp \int_\varepsilon^{t_0} t^{-(p-1)}\, dt, \tag{126}$$

and therefore

$$T(\varepsilon) = \begin{cases} \Theta(\log(t_0/\varepsilon)), & p = 2, \\ \Theta(\varepsilon^{-(p-2)}), & p > 2. \end{cases} \tag{127}$$

Thus the sharp vector least-squares and scalar zero-multiplier ALM paths have one-dimensional correction exponents $2q - 2$ and $4q - 2$, respectively.

Corollary D.9 (Finite smooth penalties are not locally exact)**.** *Let F be a minimal cycle completion at a DAG W, let U activate every edge in F, and let $\mathcal{L}$ be C^1 with $D\mathcal{L}(W)[U] < 0$. For any finite $\gamma \ge 0$, W is not a local minimizer of*

$$\mathcal{L}(\widetilde{W}) + \gamma P(\widetilde{W}) \tag{128}$$

when $P = h$ is a smooth nonnegative exact scalar or $P = \frac{1}{2}\|H\|_2^2$ is a squared smooth exact vector residual.

Proof. Along $W + tU$ with $t \downarrow 0$,

$$\mathcal{L}(W + tU) = \mathcal{L}(W) + tD\mathcal{L}(W)[U] + o(t). \tag{129}$$

The scalar and squared-vector barriers give $P(W + tU) = O(t^{2q}) = o(t)$. The negative linear score term therefore dominates every finite multiple of the penalty for all sufficiently small $t > 0$. □

D.9 Universal first- and second-order consequences

The notation $j \rightsquigarrow i$ means that a directed path runs from j to i; E_{ij} has a single one at entry (i, j).

Proposition D.10 (Gradient and Hessian degeneracy)**.** *Let h be differentiable, locally nonnegative, and exact. At every DAG W, $\nabla h(W) = 0$. If h is twice differentiable, define*

$$\mathcal{S}_W = \operatorname{span}\{E_{ij} : i \neq j,\ j \not\rightsquigarrow i \text{ in } \operatorname{supp}(W)\}. \tag{130}$$

Then $\mathcal{S}_W \subseteq \ker \nabla^2 h(W)$ and

$$\dim \ker \nabla^2 h(W) \geq d(d-1) - r(W) \geq \frac{d(d-1)}{2}, \tag{131}$$

where $r(W)$ is the number of reachable ordered pairs. At $W = 0$ the Hessian is zero.

Proof. A feasible DAG is a local minimizer of h, so Fermat's condition gives the zero gradient. If $j \not\rightsquigarrow i$, adding only edge $i \to j$ cannot create a cycle. Hence $t \mapsto h(W + tE_{ij})$ is identically zero and $\langle E_{ij}, \nabla^2 h(W) E_{ij} \rangle = 0$. The Hessian at a local minimum is positive semidefinite; for such an operator, zero quadratic form implies membership in its kernel. There are $d(d-1) - r(W)$ safe coordinates, and a DAG has at most one reachable orientation for each unordered pair, so $r(W) \leq d(d-1)/2$. □

The proposition explains why the linear independence and Mangasarian–Fromovitz constraint qualifications (LICQ and MFCQ) fail for a scalar exact equality at every feasible DAG. It is weaker than the completion theorem when $q > 1$, but makes the topology dependence of the Hessian explicit.

Proof that tolerance is not a support certificate. Let a DAG W_0 contain a directed path $j \rightsquigarrow i$ and add the back edge $i \to j$ with weight t. Every $t \neq 0$ gives cyclic support, whereas $W_t \to W_0$. Exactness gives $h(W_t) > 0$ and continuity gives $h(W_t) \to h(W_0) = 0$. Therefore every positive tolerance accepts some cyclic support. Starting from an empty graph and shrinking all edges of a cycle gives the same conclusion without requiring a nonempty base DAG. □

E Sharpness and Exact Constructions

This construction separates overflow control from low-order sensitivity. Let $s(x) = x^2/(1 + x^2)$, define $S(W)_{ij} = s(W_{ij})$ for $i \neq j$ and $S(W)_{ii} = 0$, and set

$$C(W) = \frac{\kappa}{d} S(W), \qquad h_{\mathrm{BLD}}(W) = -\left(\frac{d}{\kappa}\right)^2 \log\det(I - C(W)), \quad 0 < \kappa < 1. \tag{132}$$

Proposition E.1 (Bounded log-determinant representation)**.** *For fixed d and $0 < \kappa < 1$, h_{BLD} is C^∞, nonnegative, exact, and globally bounded. Its gradient and Hessian are globally bounded, so its gradient is globally Lipschitz. Nevertheless, it obeys the same $2q$ and 2τ barriers as every nonnegative scalar representation.*

Proof. Every entry of S lies in $[0, 1)$, so the maximum row and column sums of the nonnegative matrix C are below κ. Thus $\rho(C) < \kappa < 1$ and

$$-\log\det(I - C) = \sum_{r \geq 1} \frac{\operatorname{tr}(C^r)}{r}. \tag{133}$$

Every term is nonnegative. A DAG makes C nilpotent and every trace zero; a directed cycle contributes a positive product to an appropriate trace. This proves exactness. Moreover $\operatorname{tr}(C^r) \leq d\kappa^r$, so the unscaled series is at most $-d\log(1 - \kappa)$.

Let $R = (I - C)^{-1}$. In this proof, matrix norms with subscripts 1 and ∞ are the induced maximum column and row sums, and the 2-norm is the spectral norm. The Neumann series gives

$\|R\|_1, \|R\|_\infty, \|R\|_2 \le (1-\kappa)^{-1}$. The first two derivatives of s are globally bounded. For perturbations U, V,

$$Dh[U] = p\,\mathrm{tr}(R\,DC[U]), \tag{134}$$

$$D^2h[U,V] = p\,\mathrm{tr}(R\,DC[V]R\,DC[U]) + p\,\mathrm{tr}(R\,D^2C[U,V]), \tag{135}$$

where $p = (d/\kappa)^2$. Standard Frobenius inequalities therefore give global bounds for the gradient and Hessian. The final claim follows from Theorem A.8 and Eq. (49); bounded derivatives change coefficients but not the missing Taylor orders. □

F Experimental Settings and Numerical Verification

Shared conventions and setting overview

Table 4 defines the six settings used in the main text. Local-response diagnostics evaluate constraint formulas; controlled flows and the matched-mask intervention test specified update rules. Official trajectories record released learners, while the last two settings evaluate statistical certificates. Their data, repetition units, and targets differ. The detailed protocols below state preprocessing, optimizer budgets, and aggregation rules before their results.

Evaluation conventions. A slope is a least-squares fit in log–log coordinates, with its grid and fitting window specified below. A seed generates one graph, coefficient draw, or projection as stated in each protocol; points along its scale grid are repeated evaluations of that instance. Slopes are compared with the predicted exponent, while separation experiments report a hitting time or a budget-capped statistic. The latter two are distinguished explicitly. Constraint formulas evaluated under our controlled solvers are not called NOTEARS or DAGMA learning runs; that designation is reserved for the instrumented released implementations in Appendix F.9.

For data $X \in \mathbb{R}^{n\times d}$, our least-squares score is $\mathcal{L}_{\mathrm{LS}}(W;X) = \|X - XW\|_F^2/(2n)$, where $\|\cdot\|_F$ is the Frobenius norm and $\|W\|_1 = \sum_{ij}|W_{ij}|$. The common regularization scale is $\lambda_{\max}(X) = \max_{i\neq j}|(X^TX/n)_{ij}|$, computed after the preprocessing specified for each experiment. ER and SF denote Erdős–Rényi and scale-free graph generators; their suffixes 2 and 4 request mean total degrees 2 and 4, respectively. SEM abbreviates structural equation model.

Table 4: Experimental settings. Seed ranges refer to data or instance generation, not to points on a scale grid. Deterministic calculations have no sampling repetitions. The named protocols give the remaining settings.

Setting	Data and comparisons	Repetitions and evaluation
Local-response diagnostics	Prescribed completions, $q = 1,\dots,5$; random $d = 8$ sections, $q = 1,\dots,4$. Cycle-vector and scalar residuals and their penalties.	Deterministic hierarchy; census seed indices 0–19 per q. Value/derivative slope error (§F.1, §F.2).
Controlled cycle flows	Isolated cycles of length 4, 6, or 8; EXP and DAGMA residuals with linear, quadratic, or seeded-ALM penalties. Population and empirical Gaussian SEM scores.	Deterministic scale grids; SEM samples $n = 128, 512, 2048$, seed indices 0–19. Separation time and false-edge selection (§F.8).
Matched-mask intervention	Seven-node SEM, $n = 500$; common data and initial weights, $q = 2, 3, 4$. Masked flow, Adam, and L-BFGS-B.	Formal seeds 7100–7123; smoke tests 6090–6091. Paired budget-capped separation times; failures reported (§F.11).
Official optimizer trajectories	ER2/ER4/SF2/SF4, $d = 20, n = 500$. Official NOTEARS and DAGMA; four formal noise/weight regimes.	Development seeds 100–104 on two regimes; formal seeds 0–9 on four. Normalized separation time, rank correlation, and prediction MAE (§F.9).

Table 4: Experimental settings (continued).

Setting	Data and comparisons	Repetitions and evaluation
Frontend-to-feature certification	The same four graph types and four regimes, $d = 20$. NOTEARS, DAGMA, SDCD, GOLEM port; 500 training, 5,000 holdout, 20,000 proxy samples.	Seeds 0–4, shared across frontends. Family recall, SHD, certified-label fraction and agreement (§H.7).
Exact-population calibration	Four finite-support laws; all 25 three-node DAGs. Frozen rational regressions, no training sample; holdout sizes 500, 2,000, 5,000, and 20,000.	200 draws per law, master seed 20260906; sizes nested within a draw. Coverage, feature resolution, singleton rate (§H.4).

The four formal SEM regimes are standard Gaussian, exponential-noise, 40%-weak-edge Gaussian, and heteroscedastic Gaussian. Trajectory development uses only the standard and weak-edge Gaussian regimes. The certificate setting shares samples across frontends within each cell; it is not a second analysis of the trajectory dataset. The separate Lasso-screen audit uses 160 datasets and four holdout sizes, giving 640 records (Appendix H). Screen, score, and computational-budget variants are specified before their tables in Appendices G and H.

F.1 Local-response diagnostics: equal-scale hierarchy

We hold the base DAG and candidate path fixed and compare four formulations. Value and derivative slopes measure whether changing the residual or its scalarization changes the order by the predicted amount.

For each $q \in \{1, \dots, 5\}$, the diagnostic uses $d = \max(3, q + 1)$ nodes. It fixes the first $d - q$ edges of the path $0 \to 1 \to \cdots \to d - 1$ at weight 0.7. The remaining path edges and the closing edge $d-1 \to 0$ are the q completion coordinates; the closing edge has negative sign to ensure that signed weights are exercised. Every proper subset is acyclic.

The scale grid is $\mathrm{geomspace}(0.08, 0.30, 8)$. We fit the least-squares slope of $\log R(t)$ against $\log t$ on the four smallest positive finite values. The exact simple-cycle vector is enumerated for $d \le 6$. Its Jacobian is computed by forward automatic differentiation and then restricted to the q completion coordinates before taking the Frobenius norm. EXP and all gradients use JAX float64 automatic differentiation. The 320 hierarchy rows contain eight scales for every (q, method) pair. A separate 160-row check applies the same base DAGs to EXP, the trace polynomial, DAGMA log determinant, and BLD on $\mathrm{geomspace}(0.10, 0.30, 8)$. EXP and the log determinant are the NOTEARS and DAGMA constraints (Zheng et al., 2018; Bello et al., 2022); the polynomial is $\mathrm{tr}[(I + (W \circ W)/d)^d] - d$. BLD is our bounded construction from Appendix E, evaluated at $\kappa = 0.9$ to separate boundedness from local sensitivity.

Table 5: Local-response diagnostics: topology-dependent order hierarchy. The last column is the maximum absolute log–log slope error over $q = 1, \dots, 5$.

Formulation	Value order	Derivative order	Max. error
Signed/vector residual H	q	$q - 1$	0.000
Vector least squares $\|H\|_2^2/2$	$2q$	$2q - 1$	0.046
Nonnegative scalar h	$2q$	$2q - 1$	0.046
Scalar ALM at $\mu = 0$	$4q$	$4q - 1$	0.046

Table 6: Observed scalar orders for $q = 1, \ldots, 5$. Every family is predicted to have order $2q$.

Constraint	Observed orders ($q = 1, \ldots, 5$)	Max. error
EXP	2.00, 4.00, 6.00, 8.00, 10.01	0.007
Trace polynomial	2.00, 4.00, 6.00, 8.00, 10.00	0.005
DAGMA log-det	2.00, 4.00, 6.00, 8.00, 10.00	0.004
BLD (bounded)	1.97, 3.93, 5.90, 7.87, 9.98	0.130

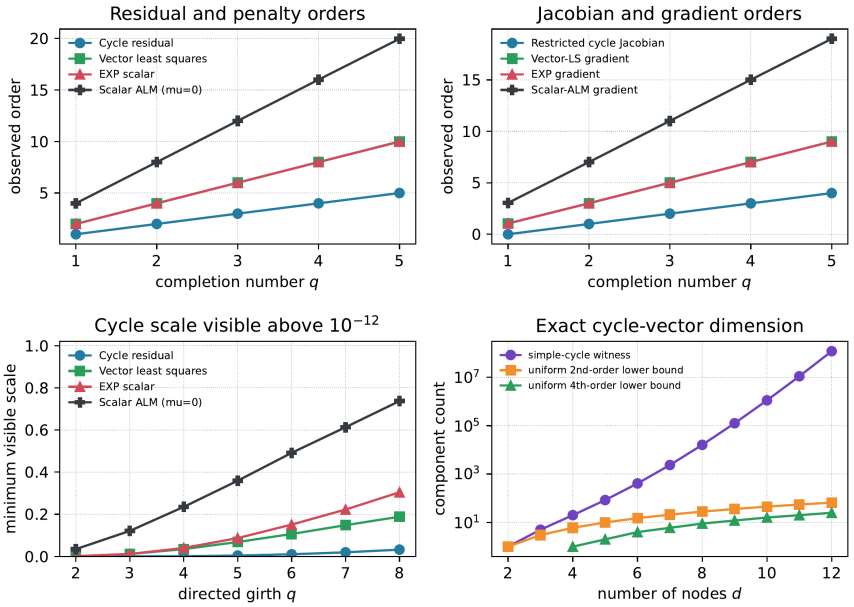


Figure 6: Empirical order hierarchy and its consequences. Top: residual/penalty and restricted Jacobian/gradient orders. Bottom left: smallest isolated-cycle scale visible above 10^{-12} in float64. Bottom right: exact cycle-vector size and uniform-sharpness lower bounds.

F.2 Local-response diagnostics: random-subspace census

We test whether the computed completion number predicts local order away from the controlled families. For each $q \in \{1, 2, 3, 4\}$ and seed index $s \in \{0, \ldots, 19\}$, the random seed is $70000 + 100q + s$. The generator chooses a random directed cycle on $d = 8$ nodes, assigns exactly q of its edges to the candidate set, and places the remainder in the base DAG. It then proposes additional random base edges, retaining one only if the base remains acyclic and $q_{\mathbf{W}}(F)$ is unchanged. It similarly adds two distractor candidate edges only when $q_{\mathbf{W}}(F)$ remains unchanged. Base weights are drawn uniformly in magnitude from $[0.45, 0.90]$, candidate coefficients from $[0.70, 1.30]$, and signs are independent.

As an independent check, the implementation exhaustively enumerates all candidate subsets and recomputes $q_{\mathbf{W}}(F)$ before evaluation. The resulting instances contain 5–10 base edges, 3–6 candidate edges, and 1–5 active simple cycles. For eight scales in $\mathrm{geomspace}(0.006, 0.05, 8)$, we evaluate the norm of all active simple-cycle products and the positive trace series for EXP. Slopes use the four smallest scales, exactly as in the controlled diagnostics. The 1,280 raw rows yield 160 instance/method fits. The predicted orders q and $2q$ hold with maximum absolute error 0.0033 and median error 4.9×10^{-15}.

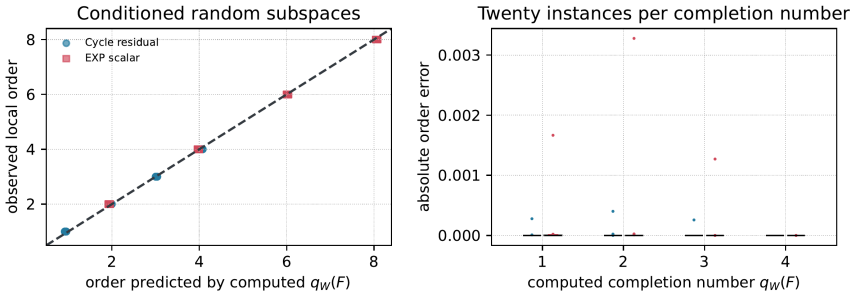


Figure 7: Conditioned random-subspace census. Left: every observed order lies on the prediction from the independently recomputed $q_{\mathbf{W}}(F)$. Right: absolute errors over 20 instances per completion number.

F.3 Completion-volume audit

The script `paper_support/explore_completion_invariants.py` tests the formulas in Theorem B.5 on finite examples. For every random candidate complex it enumerates all supports, extracts the minimal completions, solves both the integral and fractional cover problems, and evaluates Eq. (72) on 30 values of ε. It also normalizes an optimal fractional-cover vector $z^\star$ so that $\min_C \sum_{s\in C} z_s^\star = 1$ and evaluates the logarithmic path $|x_s| = t^{z_s^\star}$. Along this path, the volume-density scale is $t^{\sum_s z_s^\star} = t^\lambda$, while the generator and squared-generator scales are t and t^2.

Table 7: Completion-volume checks. "Ray" lists the fitted vector and scalar exponents. The tournaments are genuine fractional feedback-cover gaps, not arbitrary hypergraphs.

Suite	Cases	$q < r$	q	r	(b, λ)	Ray
Random DAG subspaces	597	440	–	–	23 distinct	exact
$d = 7$ tournament	1	1	3	7	$(5, 4.5)$	$(4.5, 2.25)$
$d = 8$ tournament	1	1	3	8	$(7, 6.667)$	$(6.667, 3.333)$

For the random suite, the maximum error between the fitted tube slope and b is 1.7×10^{-4}; the valuation-ray errors for λ and $\lambda/2$ are below 10^{-8}. The $d = 7$ tournament has 21 arcs and 64 simple cycles. Its minimum feedback-arc cover has size five, whereas the LP relaxation is 4.5. The $d = 8$ example has 28 arcs and 288 cycles. On the same rays, direct positive-series evaluations give orders $(1.000000, 1.000001, 2.000001, 2.000003)$ for the generator, cycle vector, EXP, and DAGMA in the $d = 7$ case; the maximum error is 1.5×10^{-4} across both tournaments.

We also draw scrambled Sobol points in $[0, 1]^m$ and directly count generator sublevel events. At practical thresholds this gives slopes 3.219 and 4.594 for the two tournaments, below the asymptotic values 4.5 and 6.667. Equation (78) permits polynomial factors in $\log(1/\eta)$, and the high-dimensional events become too rare before those factors are negligible. At the tested tolerances, the Sobol estimates do not recover the asymptotic exponent. They were not used to select parameters; the logarithmic limit is established by the proof.

F.4 Weighted overlapping cycles

To test the weighted order, we let candidate edges shrink at different rates on two overlapping cycles,

$$C_A : 0 \to 1 \to 2 \to 3 \to 0, \tag{136}$$

$$C_B : 0 \to 1 \to 4 \to 3 \to 0. \tag{137}$$

They share edges $0 \to 1$ and $3 \to 0$. The four exponent patterns give cycle costs A/B of $4/4$, $5/8$, $6/8$, and $8/5$. The final pattern switches the cheapest cycle. Eight scales follow $\mathrm{geomspace}(0.12, 0.30, 8)$. A generic weighted shortest-directed-cycle routine computes τ from the graph rather than reading either displayed cycle cost. There are 128 raw weighted rows.

Table 8: Weighted overlapping cycles. Costs A/B are the two cycle costs and τ is their minimum. The final row switches the cheapest cycle.

A/B	τ	$\|H\|$	$\|H\|^2/2$	EXP	$\text{EXP}^2/2$
4/4	4	4.00	8.00	8.00	16.00
5/8	5	5.00	10.00	10.00	20.00
6/8	6	6.00	12.00	12.00	24.00
8/5	5	5.00	10.00	10.00	20.00

F.5 Subcritical compression diagnostic

We test the loss of sharpness under dimension reduction by comparing generic directions with directions in the projection's null space. Both retain the same exact fallback component.

For $d = 10$ and $q = 2, \ldots, 6$, let $C_q(W) \in \mathbb{R}^{M_q(10)}$ collect the selected cycle products from Theorem B.1. For a matrix $A \in \mathbb{R}^{k_{\text{proj}} \times M_q(10)}$ with orthonormal rows, we evaluate the globally exact residual

$$H_A(W) = (AC_q(W),\ h_{\text{exp}}(W)). \tag{138}$$

Here $M_2(d) = \binom{d}{2}$; for $q \geq 3$, $M_q(d)$ is given by Eq. (27). Because the EXP component is globally exact, $H_A(W) = 0$ if and only if W is a DAG, regardless of the projection. Along the channel family, $C_q(tU(z)) = t^q z$ and $h_{\text{exp}}(tU(z)) = \Theta(t^{2q})$. A generic unit z therefore gives $\|H_A(tU(z))\|_2 = \Theta(t^q)$, whereas any nonzero $z \in \ker A$ gives $\Theta(t^{2q})$. Rank–nullity supplies such a direction whenever $k_{\text{proj}} < M_q(10)$.

We draw 20 random orthogonal maps at each of ranks $\lfloor M_q/4 \rfloor$, $\lfloor M_q/2 \rfloor$, and $M_q - 2$. After adding the fallback, every total residual dimension remains strictly below M_q. An SVD gives the null direction and an independent normalized Gaussian vector gives the generic direction. Eight local scales follow $\text{geomspace}(0.02, 0.12, 8)$, producing 4,800 raw rows. EXP is evaluated by its positive trace series rather than subtracting d, and the validator checks that it is positive on every cyclic mixture. The dense EXP fallback preserves exactness in this diagnostic; we do not assess its large-graph computational cost here. The largest absolute fitted-order error is 0.0031.

Table 9: Subcritical random compression on the $d = 10$ channel families of Theorem B.1. Observed orders are medians over three ranks and 20 seeds.

q	Channels $M_q(10)$	Generic	Null direction
2	45	2.000	4.002
3	20	3.000	6.000
4	16	4.000	8.000
5	12	5.000	10.000
6	9	6.000	12.000

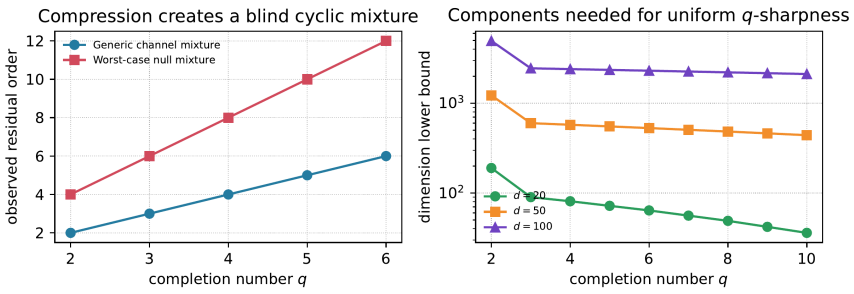


Figure 8: Subcritical compression creates a worst-case cyclic mixture with doubled residual order (left). The funnel lower bound is quadratic in d for each fixed q (right).

F.6 VISIBILITY AND CYCLE-VECTOR SIZE

We measure how response order affects visibility at a fixed numerical tolerance and record the output size of the sharp vector construction. For each isolated cycle length $q = 2, \dots, 8$, the diagnostic evaluates 500 scales from 10^{-4} to 0.99 and records the first residual at least 10^{-8}, 10^{-12}, and 10^{-16}. We report the grid-based value so that EXP cancellation and float64 evaluation are included. Cycle-vector dimensions use the closed count in Eq. (22) for $d = 2, \dots, 12$.

Table 10: Smallest isolated-cycle scale whose residual exceeds 10^{-12} in float64. Lower is better.

q	$\|H\|$	$\|H\|^2/2$	EXP	EXP$^2/2$
2	0.0001	0.0012	0.0010	0.0345
4	0.0010	0.0345	0.0400	0.2350
6	0.0100	0.1064	0.1510	0.4913
8	0.0321	0.1884	0.3042	0.7371

F.7 RESTRICTED OPTIMIZATION DYNAMICS

We compare the remaining candidate-edge scale after a fixed step budget. The p rows minimize powers $|R(t)|^p$; cold and seeded ALM differ in their initial multiplier. All objectives use the same projected-gradient solver, so the comparison varies the objective while keeping the update rule fixed.

The family has one fixed edge $c \sim \mathrm{Uniform}[0.55, 0.90]$ and q completion edges of scale t, giving $R(t) = ct^q$. With $L = q + 1$ and $P = c^2 t^{2q}$, NOTEARS equals

$$h_{\exp}(t) = \sum_{m=1}^{\infty} \frac{LP^m}{(mL)!}, \qquad \mathrm{tr}(A^k) = 0 \text{ unless } L \mid k. \tag{139}$$

Twelve terms match the JAX matrix exponential through $q = 6$ within 2×10^{-14}. Every method uses 1,000 projected-gradient steps with Armijo parameter 10^{-4} and maximum step 0.25. ALM uses 40 updates, 25 primal steps per update, $\rho_0 = 1$, and growth 1.2 capped at 10^6; only the seeded variant sets $\mu_0 = 1$. Every trajectory starts at $t = 0.8$. These trajectories follow the restricted correction path into the local regime; all noncandidate weights remain fixed.

The coefficient-controlled protocol rescales each objective to unit initial gradient. At $q = 6$, its medians are $0.134/0.266/0.357/0.357/0.376/0.357$ for $p = 1/p = 1.5/p = 2$/EXP/cold/seeded; full records are in the released CSV.

Table 11: Median completion-edge scale after 1,000 raw-objective projected-gradient steps over 20 seeds. Lower is better; all runs start at 0.8.

q	$p = 1$	$p = 1.5$	$p = 2$	EXP	cold ALM	seeded ALM
2	0.000	0.002	0.044	0.044	0.094	0.044
4	0.025	0.125	0.256	0.387	0.469	0.386
6	0.122	0.268	0.406	0.707	0.790	0.706
8	0.219	0.372	0.501	0.798	0.800	0.798

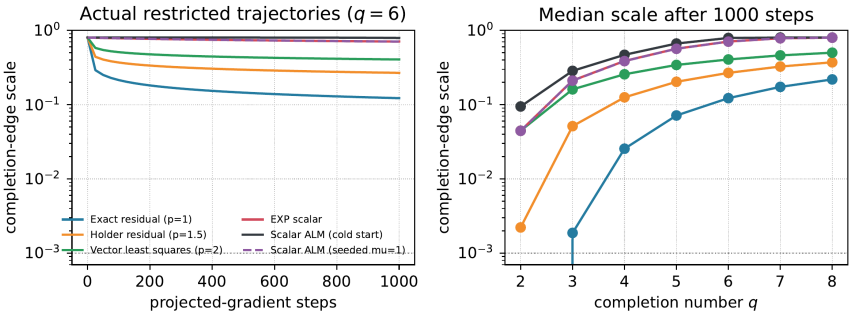


Figure 9: Restricted feasibility-correction trajectories under the fixed protocol. Lower is better. The exact $p = 1$ residual retains the strongest signal, while smooth scalarization and cold-start ALM become stiffer as the completion number grows.

F.8 Controlled cycle flows: full-matrix and SEM-score protocol

We evaluate the isolated-cycle flow before adding a population SEM score to measure its effect on edge selection. The flow tests use directed cycles of lengths $L \in \{4, 6, 8\}$. Every cycle edge is optimized. The two tracked edges start at $a \pm \varepsilon$, the remaining $L - 2$ edges start at 0.8, and we vary $a \in \{0.3, 0.5\}$, $\zeta \in \{0.1, 0.25\}$, and $\varepsilon \in [10^{-5}, 10^{-2}]$. We evaluate the exact full-matrix NOTEARS and DAGMA constraints under

$$\Psi(h) = h, \qquad \Psi(h) = \frac{1}{2}h^2, \qquad \Psi(h) = h + \frac{1}{2}h^2.$$

These are the linear feasibility force, a cold quadratic penalty, and an ALM inner objective with unit multiplier and penalty. Here ζ is the smaller-to-larger ratio defining the event. Selection times are computed by quadrature of the exact matrix-flow reduction. We also compare the closed-form values and every cycle-edge gradient with direct matrix exponential and inverse evaluations at 120 random cycle matrices.

For the score experiment, the population data distribution is the linear Gaussian SEM whose true graph is the path

$$0 \longrightarrow 1 \longrightarrow \cdots \longrightarrow L - 1$$

with every coefficient equal to 0.5 and unit noise variance. If $W_\star$ is this path matrix, its covariance under the row-vector convention is

$$\Sigma = (I - W_\star)^{-T}(I - W_\star)^{-1}.$$

Initialization adds the false back edge $(L - 1) \to 0$, closing the cycle. The optimized smooth objective is

$$F_\Sigma(W) = \frac{0.2}{2}\operatorname{tr}\big((I - W)^T\Sigma(I - W)\big) + h(W) + \frac{1}{2}h(W)^2. \tag{140}$$

We run both a cycle-support flow and a dense flow in which every off-diagonal entry may change. "Aligned" initialization makes the false edge smaller by 2ε; the adversarial "reversed" initialization makes it larger. Correct selection is declared when the false-to-true tracked-edge ratio first reaches $1/4$. The opposite $1/4$ event is recorded as an incorrect selection. No ground-truth quantity enters the gradient or stopping rule other than this post-run correctness label.

For the finite-sample check, we replace Σ by the empirical covariance from $n \in \{128, 512, 2048\}$ observations. Seed indices $s = 0, \ldots, 19$ use random seed $s + 10000L + n$. All 720 finite-sample runs use reversed initialization.

All 216 population flows reach the false-edge ratio event before its opposite, as do explicit gradient-descent runs with step 0.01 on the predeclared subset of ε values. The fitted edge-separation-time slopes range from -0.0041 to 0.0039, and the minimum total deletion margin in every population trajectory is positive. Dense trajectories are genuinely different from the cycle restriction: their off-cycle Frobenius norm reaches 0.384. We enforce the DAGMA domain $\rho(W \circ W) < 1$ at every gradient evaluation; the largest value over all population and finite-sample DAGMA trajectories is 0.453.

Finite-sample correctness is 91.25%; the per-setting range is 65% to 100%, increasing toward the population result with sample size. The positive-margin condition in Proposition D.3 applies to a fixed objective and need not hold for every empirical covariance.

F.9 Official optimizer trajectories

We test whether separation diagnostics measured at an early cyclic iterate predict subsequent edge separation beyond the constraint value and weight scale. All graphs have $d = 20$ and 500 observations. Development seeds 100–104 cover ER2/ER4/SF2/SF4 graphs under standard Gaussian and 40%-weak-edge Gaussian SEMs. Formal seeds 0–9 use these regimes plus exponential and heteroscedastic noise, giving 160 formal trajectories per optimizer. The prediction model is fitted on development trajectories and evaluated on formal trajectories.

We record the unthresholded iterates of the official linear NOTEARS and DAGMA updates (Zheng et al., 2018; Bello et al., 2022). The update equations, continuation schedules, stopping rules, and regularization parameters are unchanged. Read-only callbacks store every fifth L-BFGS-B iterate for NOTEARS and every 250th Adam iterate for DAGMA. On small deterministic tests, the instrumented and uninstrumented implementations agree to relative tolerances 10^{-9} and 10^{-12}, respectively.

F.9.1 Truth-free early diagnostics

For an unthresholded matrix W, define its critical DAG threshold

$$\tau(W) = \min\{t \geq 0 : \mathrm{supp}(|W| > t) \text{ is a DAG}\}. \tag{141}$$

This quantity is used as a diagnostic, not as the reported graph. At the first completed continuation stage whose nonzero support is cyclic, let $B = \mathrm{supp}(|W| > 2\tau)$, and treat entries with magnitude in $[\tau, 2\tau]$ as completion coordinates. The base B is acyclic. We compute the smallest number q_2 of band coordinates needed to complete a cycle and choose a deterministic shortest witnessing cycle C. Thus q_2 is a two-scale operational version of the completion statistic, not a claim that the random trajectory lies in the exact local model of Section 2.

For a two-cycle witness, the smaller weight is exactly τ: both weights are at least τ, and if both were larger, the strict-threshold graph in Eq. (141) would still contain a cycle. Writing $u = a_{(1)}/a_{(2)}$, the larger edge is a base edge precisely when $a_{(2)} > 2\tau$. Hence

$$q_2 = 1 + \mathbf{1}\{u \geq 1/2\}.$$

All 80 development witnesses and 319 of the 320 formal witnesses are two-cycles; the identity holds in every such record. The remaining formal witness is a NOTEARS three-cycle (ER4, exponential noise, seed 8), with $q_2 = 2$ and $u = 0.6140$. It happens to have the same encoding, so replacing q_2 by the ratio step leaves all 400 recorded feature vectors unchanged. The two-cycle identity alone does not extend to general cycles.

Three dimensionless quantities describe the early ranking on C. Write

$$a_e = |W_e|, \qquad b_e = |\partial_e h(W)|, \qquad c_e = \mathcal{L}(W \text{ with } W_e = 0) - \mathcal{L}(W).$$

Here $\mathcal{L} = \mathcal{L}_{\mathrm{LS}}(\cdot; X)$ is evaluated without refitting the other weights. For the two selected values v_1, v_2, the relative gap is $|v_1 - v_2|/\max(|v_1|, |v_2|, \varepsilon_{\mathrm{mach}})$, with $\varepsilon_{\mathrm{mach}} = 2^{-52}$ for float64. For the weights and deletion costs, the separation is the relative gap between the smallest and second-smallest values. For the constraint force it is the relative gap between the largest and second-largest values. A zero gap denotes a first-order tie. The outcome is the first later snapshot at which the smallest cycle weight is at most one tenth of the second-smallest. Time is divided by the remaining optimizer steps, separately for each method; an unresolved trajectory is right-censored at one. No generating edge, coefficient, or SHD enters the anchor, cycle, predictor, or outcome.

The development split fixes one ridge model per optimizer with penalty one. The baseline uses $\log h(W)$, $\log(\tau(W)/\|W\|_{\max})$, and $\log \|W\|_{\max}$. The geometry model adds q_2 and the three separation statistics. Spearman correlations, bootstrap intervals, and permutation tests are computed only after this freeze. Here $\|W\|_{\max} = \max_{i \neq j} |W_{ij}|$. In the tables, ρ_s is Spearman's rank correlation, p_{perm} its permutation-test p-value, and CI a confidence interval.

F.9.2 FORMAL RESULTS

Table 12: Official optimizer trajectories: 320 formal runs. The correlation is between early constraint separation and normalized edge-separation time. Models are fitted only on disjoint development seeds.

Method	Early predictor	ρ_s [95% CI]	p_{perm}	Base MAE	Geom. MAE
NOTEARS	constraint separation	-0.52 [-0.62,-0.40]	0.0001	0.041	0.041
DAGMA	constraint separation	-0.66 [-0.75,-0.56]	0.0001	0.108	0.088

Constraint separation is associated with faster edge separation for both optimizers: $\rho_s = -0.521$ for NOTEARS and -0.664 for DAGMA. The 95% bootstrap intervals exclude zero, and both permutation tests give $p < 10^{-4}$. Ranking observations within each graph–noise stratum gives -0.564 and -0.649, respectively, so the association persists after separating sparse and dense regimes. Completion depth has the opposite sign, with correlations 0.435 for NOTEARS and 0.386 for DAGMA. The geometry model, fitted on development data and held fixed for evaluation, lowers DAGMA's formal MAE by 18.5%, from 0.1078 to 0.0879, and raises its prediction rank correlation from 0.316 to 0.648. For NOTEARS, the richer model leaves MAE essentially unchanged (0.04112 to 0.04100), despite the local association.

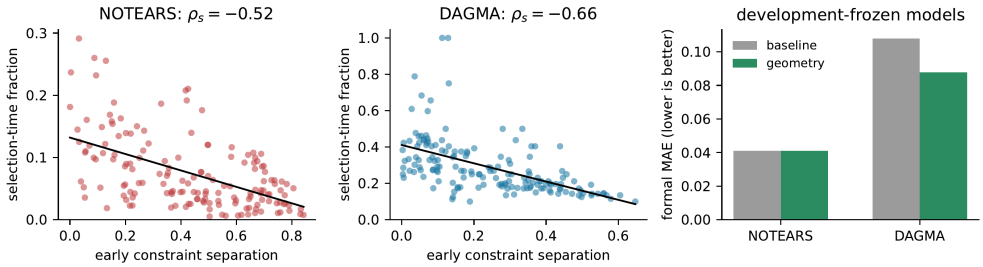


Figure 10: Predicting a later support event from early optimizer states, without using the generating graph. Left and center: formal trajectories and pooled rank correlations. Right: MAE of baseline and geometry models fitted on development data and held fixed for evaluation. Prediction improves materially for DAGMA and is essentially unchanged for NOTEARS.

Table 13: Censoring and one-factor sensitivity. The range varies the critical band multiplier over $1.5, 2, 3$, the endpoint ratio over $1/5, 1/10, 1/20$, or doubles the snapshot interval.

Method	Censored	Baseline ρ_s	Uncensored ρ_s	Audit range
NOTEARS	0/160	-0.521	-0.521	[-0.644,-0.450]
DAGMA	2/160	-0.664	-0.664	[-0.703,-0.577]

Only two baseline trajectories are censored, and the uncensored estimate is unchanged at the reported precision. All sensitivity settings preserve the sign and a moderate-to-strong rank association. The complete configuration grid and trajectory-level outcomes are included in the supplement.

The random-graph experiment also shows which local configurations the diagnostic encounters. Of the 320 critical cycles, 319 are directed two-cycles and one is a three-cycle; $q_2 = 1$ in 78 runs and $q_2 = 2$ in 242. Dense signed frontends therefore expose mostly pairwise orientation conflicts at their first continuation stage, whereas the controlled experiments deliberately span higher completion orders. Score separation is method-dependent as well: its formal correlation is -0.304 for NOTEARS and 0.118 for DAGMA. In this setting, an early feasibility tie and a larger two-scale completion depth are associated with a slower support decision. The nested comparison below tests whether these associations add information beyond the initial edge-weight ratio.

F.10 NESTED PREDICTIVE ABLATION

We specified this post-hoc comparison after inspecting the original formal results. We retain the original development/formal split and fit all six nested ridge models with penalty one. No model or penalty is selected by formal MAE. Starting from the original baseline, the features added in order are the anchor-cycle ratio $a_{(1)}/a_{(2)}$, $\log[(a_{(1)}/a_{(2)})/0.1]$, q_2, constraint separation, and score separation. Features are standardized using development observations only. The logarithm supplies a nonlinear encoding of the same initial ratio, not an independent geometric measurement. The same distinction applies to q_2, which is the ratio step above on every two-cycle.

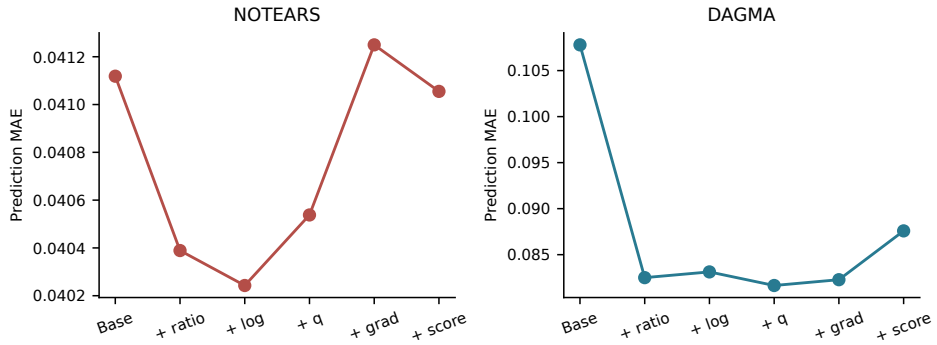


Figure 11: Post-hoc nested prediction audit. All models use the original development split. The direct initial-ratio feature explains most of the DAGMA gain over the original baseline. Later additions have smaller and method-dependent effects.

Table 14: Incremental reduction in formal MAE, with pointwise 95% paired bootstrap intervals. Positive values favor the added feature. Each bootstrap resamples trajectories within graph–noise strata; 10,000 draws use seed 20260906. These intervals do not adjust for the five comparisons.

Feature added	NOTEARS	DAGMA
Initial ratio	$0.00073\ [0.00034, 0.00111]$	$0.02529\ [0.01682, 0.03407]$
Log distance	$0.00015\ [-0.00024, 0.00056]$	$-0.00062\ [-0.00104, -0.00027]$
Completion depth	$-0.00030\ [-0.00093, 0.00036]$	$0.00148\ [0.00020, 0.00291]$
Constraint separation	$-0.00071\ [-0.00126, -0.00023]$	$-0.00063\ [-0.00448, 0.00309]$
Score separation	$0.00019\ [-0.00038, 0.00077]$	$-0.00530\ [-0.01020, -0.00030]$

For DAGMA, adding the ratio lowers MAE from 0.10780 to 0.08250, below the geometry model's MAE in Table 12. Completion depth gives a smaller improvement after the log-distance baseline. The final feature bundle has MAE 0.08758, compared with 0.08313 for that stronger baseline; the paired gain is -0.00446, with interval $[-0.01026, 0.00139]$. NOTEARS shows no improvement from the full bundle over the stronger baseline either: the gain is -0.00081, with interval $[-0.00179, 0.00012]$. Removing the two censored DAGMA trajectories leaves the same pattern: MAE is 0.07557 for the stronger baseline, 0.07408 after completion depth, and 0.07979 for the full bundle. In this dataset, q_2 adds a nonlinear basis function to a linear predictor. The improvement does not establish information beyond the initial ratio. Higher-order completion effects must be studied in candidate sections where this encoding identity does not apply.

F.11 MATCHED-MASK INTERVENTION

The official trajectories suggest that initial edge imbalance explains much of the prediction gain. To separate this effect from the allowed coordinate updates, we hold the imbalance fixed and change only the update mask. On seven nodes, the initial matrix has a four-cycle $0 \to 1 \to 2 \to 3 \to 0$ with weights $(0.9a, a, 2a, 2a)$, and a disjoint chain $4 \to 5 \to 6$ with weights $(2a, 2a)$. The first q cycle

coordinates are variable; the remaining cycle coordinates are fixed. We make $4-q$ chain coordinates variable so that every mask has four optimized parameters. Zeroing the variable coordinates gives an acyclic base, and precisely q candidates are needed to complete its only cycle. We use $q = 2, 3, 4$ and $a = 0.2, 0.4, 0.6$.

Within a seed–scale block, the full initial matrix and data are identical across masks. The initial ratio of the two marked edges is 0.9, and their two objective-gradient components, the initial residual, and the initial objective also agree exactly in the recorded computations. The norm of the whole masked gradient is not matched: the intervention changes which other cycle edges can contract. It therefore tests a mask effect at a common initial state, not the isolated effect of an abstract completion number with all subsequent dynamics held fixed.

For each of 24 fresh seeds (7100–7123), we generate 500 centered Gaussian SEM observations. The generating matrix is the common initial matrix with edge $0 \to 1$ removed. This generating support is not used to select the mask, penalty, stopping event, or optimizer settings. The objective is

$$\mu\{\mathcal{L}_{\mathrm{LS}}(W; X) + 0.01\|W\|_1\} + h(W)/h(W_0), \qquad \mu \in \{0, 1\},$$

restricted to the mask and the box $[-1.25, 1.25]^4$. Here W_0 is the common initial matrix and the four-dimensional box constrains only the optimized coordinates. We use the exponential residual and the log determinant with $s = 2$. Both are evaluated by their four-cycle product formulas, checked against the matrix expressions and finite-difference gradients. The full objective is the same across masks. The score-free controls have no random data dependence and are run only once per scale.

The three solvers are numerically integrated projected gradient flow, PyTorch Adam, and SciPy L-BFGS-B. Each uses the prescribed four-coordinate mask. The event is a marked-edge smaller-to-larger magnitude ratio at most 0.1. The budgets are flow time 20, 3,000 Adam steps, and 300 L-BFGS-B iterations. Adam uses learning rate 0.002, moments $(0.9, 0.999)$, and denominator offset 10^{-8}. L-BFGS-B uses gradient tolerance 10^{-10}, zero relative objective tolerance, and at most 50 line-search steps. The protocol and source hash were recorded before the formal outcomes; two separate seeds (6090–6091) were used for smoke checks. There are 1,296 score-on runs and 54 deterministic controls.

The prespecified contrast is the mean paired difference in $\log(1 + T_{\mathrm{cap}})$ between $q = 4$ and $q = 2$, where T_{cap} is event time or iteration count, assigned the budget if the solver never reaches the event. Stationary termination without an event is included in this rule. The statistic is not an uncensored mean hitting time. Bootstrap draws resample whole seeds, preserving all three scales within each seed. The six solver–constraint comparisons use percentile tails $0.05/12$ and $1 - 0.05/12$; these are approximate Bonferroni-adjusted bootstrap intervals, not finite-sample coverage bounds.

Table 15: Matched-mask contrast, score on. Positive values indicate a larger capped separation statistic for $q = 4$. Each row uses 72 paired seed–scale blocks and 24 independent seeds; 10,000 seed-clustered bootstrap draws use seed 17803. All non-event runs remain in the comparison.

Constraint	Solver	Mean paired difference	Adjusted interval
EXP	Flow	0.04581	[0.03559, 0.05903]
EXP	Adam	0.02241	[0.00114, 0.04401]
EXP	L-BFGS-B	−0.04897	[−0.15828, 0.06712]
Log det	Flow	0.04581	[0.03546, 0.05917]
Log det	Adam	0.02231	[0.00090, 0.04342]
Log det	L-BFGS-B	0.06204	[−0.21834, 0.27734]

Scale-specific results qualify the pooled contrasts. For EXP at $a = 0.2$, median flow times for $q = 2$ and $q = 4$ are 0.1425 and 0.1867; at $a = 0.4$, they are 0.6970 and 0.8104, although individual seeds can reverse this difference. Adam takes 366 versus 409 steps at $a = 0.2$, but 759 versus 695.5 steps at $a = 0.4$: the direction reverses. At $a = 0.6$, neither flow nor Adam reaches the event in any of the 24 seeds, for any mask or constraint. L-BFGS-B reaches it in one seed per configuration at this scale and otherwise terminates without separation. The log-determinant results have the same scale-dependent pattern. Overall, 433 of 1,350 runs do not reach the event; none has a recorded numerical failure. Thus a common initial ratio does not determine the subsequent separation, but

neither does q impose a solver-independent ordering. The experiment measures feedback from permitted coordinate updates, rather than the Taylor exponent itself.

For numerical validation, 54 selected flow runs were repeated with relative and absolute tolerances reduced from $(10^{-8}, 10^{-10})$ to $(10^{-10}, 10^{-12})$. Every event/status label agrees, and the largest change in capped event time is 2.03×10^{-8}. The original projected-flow integration exceeds the box by at most 3.39×10^{-6}; the tighter checks show no box excess. The source, protocol hashes, all per-run outcomes, and scale-specific summaries are included in the supplement.

G Modular Certificate: Proof and Scaling Audit

This appendix studies an additive deletion score: the cost of removing a set of edges is the sum of their fixed individual costs. That target permits the cycle-cover LP relaxation below. General refitted parent-set losses are not additive in edges; the main certificate retains those interactions and is developed in Appendix H.

G.1 The algebraic bridge

Here every edge of the active support E may be deleted: the completion complex has empty base and candidate set E. Its minimal completions are the simple directed cycles. A transversal is a subset of E meeting every such cycle; $\mathcal{T}$ contains all transversals, not only the inclusion-minimal ones. The superscript $\vee$ denotes Alexander duality, given explicitly by the intersection below. See Miller & Sturmfels (2005, Chapter 5) for background; the proof spells out the squarefree definition used here.

Lemma G.1 (Completion-repair duality)**.** *Let*

$$I_E = \langle x^C : C \in \mathcal{C}_{\min} \rangle$$

be the squarefree cycle-completion ideal on an active directed support E. If $\mathcal{T}_{\min}$ is the family of inclusion-minimal deletion sets that leave a DAG, then

$$I_E^\vee = \bigcap_{C \in \mathcal{C}_{\min}} \langle x_e : e \in C \rangle = \langle x^D : D \in \mathcal{T}_{\min} \rangle. \tag{142}$$

Proof. Let $\Delta_E = \{S \subseteq E : C \not\subseteq S \text{ for every } C \in \mathcal{C}_{\min}\}$, so $I_E = I_{\Delta_E}$. The dual complex is $\Delta_E^\vee = \{D \subseteq E : E \setminus D \notin \Delta_E\}$, and $I_E^\vee = I_{\Delta_E^\vee}$ by definition. For $D \subseteq E$, the monomial x^D belongs to this ideal exactly when $D \notin \Delta_E^\vee$, or equivalently $E \setminus D \in \Delta_E$. This is the condition $D \cap C \neq \varnothing$ for every $C \in \mathcal{C}_{\min}$, which also characterizes membership in the intersection of coordinate primes. Both ideals are squarefree, so

$$I_E^\vee = \bigcap_{C \in \mathcal{C}_{\min}} \langle x_e : e \in C \rangle.$$

Divisibility-minimal squarefree monomials correspond to inclusion-minimal transversals. These are precisely the minimal edge deletions that make E acyclic. □

We record the edge-modular specialization used below. After a frozen training stage, let each active edge have bounded holdout deletion loss $Z_e \in [-B, B]$ and population mean $\mu_e = \mathbb{E}Z_e$. For n independent holdout observations, let $\widehat{\mu}_e = n^{-1} \sum_{i=1}^n Z_{i,e}$ and $\widehat{\mu}(D) = \sum_{e \in D} \widehat{\mu}_e$. Write $D \triangle D^*$ for edges in exactly one of the two sets. Write $\mu(D) = \sum_{e \in D} \mu_e$ for a feasible transversal $D \in \mathcal{T}$, and let D^* be a minimum-cost transversal. When this minimizer is unique, its normalized completion-exchange margin is

$$\kappa_{\mathcal{C}}(\mu) = \min_{D \in \mathcal{T},\, D \neq D^*} \frac{\mu(D) - \mu(D^*)}{|D \triangle D^*|}. \tag{143}$$

Given simultaneous intervals $[\ell_e, u_e]$ and an empirical minimizer $\widehat{D}$, define

$$\underline{\Delta}(D; \widehat{D}) = \sum_{e \in D \setminus \widehat{D}} \ell_e - \sum_{e \in \widehat{D} \setminus D} u_e, \qquad \mathcal{A}_\alpha(\widehat{D}) = \{D \in \mathcal{T} : \underline{\Delta}(D; \widehat{D}) \leq 0\}. \tag{144}$$

Let $E_{\mathrm{del}} = \cap_{D\in\mathcal{A}_\alpha} D$, $E_{\mathrm{keep}} = E \setminus \cup_{D\in\mathcal{A}_\alpha} D$, $L = \min_{D\in\mathcal{T}} \underline{\Delta}(D;\widehat{D})$, and $L_- = \min_{D\neq\widehat{D}} \underline{\Delta}(D;\widehat{D})$. Common-radius intervals below mean $[\ell_e, u_e] = [\widehat{\mu}_e - r, \widehat{\mu}_e + r]$; $\widehat{D}$ minimizes $\widehat{\mu}(D)$ over $\mathcal{T}$.

Theorem G.2 (Selective completion-repair confidence set)**.** *If all intervals cover simultaneously, every population-optimal repair belongs to $\mathcal{A}_\alpha(\widehat{D})$. Consequently, every population optimum deletes E_{del} and retains E_{keep}. On the same event,*

$$\mu(\widehat{D}) - \min_{D\in\mathcal{T}} \mu(D) \le [-L]_+. \tag{145}$$

Moreover, $\mathcal{A}_\alpha = \{\widehat{D}\}$ exactly when $L_- > 0$. If the intervals have common radius r and $\kappa_{\mathcal{C}}(\mu) > 0$, then $r < \kappa_{\mathcal{C}}(\mu)$ implies $\widehat{D} = D^*$, while $r < \kappa_{\mathcal{C}}(\mu)/2$ makes the confidence set a singleton.

The coverage statement in Theorem G.2 justifies the shared deletion and retention labels even when several optima remain. Equation (145) instead controls the reference repair's excess score. The final two radius conditions give separate sufficient precision levels: $r < \kappa_{\mathcal{C}}$ guarantees correct point selection, while $r < \kappa_{\mathcal{C}}/2$ also guarantees a singleton confidence set. Correct point selection alone is not the certificate's test for a unique optimum.

The following relaxation uses the same simultaneous interval-coverage event. Its guarantee about population optima is conditional on that event.

Proposition G.3 (Conservative cycle-cover certificate)**.** *For any feasible reference D_0, relax repair indicators to $P = \{z \in [0,1]^E : \sum_{e\in C} z_e \ge 1 \ \forall C\}$. If the robust forced-opposite LP for edge e has lower bound above the robust score of D_0, every population optimum has the D_0 label at e. These labels are a subset of those returned by exact forced-opposite optimization for the same reference. Minimum-weight cycle separation gives a polynomial-time separation oracle for the LP.*

The relaxation trades resolution for computation, not the coverage event: fractional alternatives can prevent a decision but cannot create a false one on that event. The cycle-cover relaxation is classical; its role here is to preserve the certificate's direction of implication for modular scores.

G.2 Proof of Theorem G.2

A deletion set leaves an acyclic support if and only if it intersects every simple directed cycle. The simple cycles are precisely the inclusion-minimal nonfaces of the acyclic edge complex, so the feasible deletion sets are the transversals of $\mathcal{C}_{\min}$.

Let $m = |E|$. For the common-radius statement, Hoeffding's inequality and a union bound give the event

$$\mathcal{E}_n = \left\{ \max_{e\in E} |\widehat{\mu}_e - \mu_e| \le r_n \right\}, \qquad r_n = B\sqrt{2\log(2m/\alpha)/n}, \qquad \Pr(\mathcal{E}_n) \ge 1-\alpha. \tag{146}$$

The first part of the theorem only needs the more general simultaneous event $\mathcal{E} = \{\ell_e \le \mu_e \le u_e \ \forall e\}$. Dependence among the coordinates of one holdout vector is allowed. On $\mathcal{E}$, every $D \in \mathcal{T}$ satisfies

$$\mu(D) - \mu(\widehat{D}) = \sum_{e\in D\setminus\widehat{D}} \mu_e - \sum_{e\in\widehat{D}\setminus D} \mu_e \tag{147}$$

$$\ge \underline{\Delta}(D;\widehat{D}). \tag{148}$$

Let D^* be any population-optimal repair. Since $\mu(D^*) - \mu(\widehat{D}) \le 0$, Eq. (148) implies $\underline{\Delta}(D^*;\widehat{D}) \le 0$, and hence $D^* \in \mathcal{A}_\alpha(\widehat{D})$. This holds for every population optimum. An edge in the intersection of the confidence set is therefore deleted by every optimum; an edge outside its union is retained by every optimum.

Minimizing Eq. (148) over D proves Eq. (145). By definition, $\mathcal{A}_\alpha = \{\widehat{D}\}$ if and only if $\underline{\Delta}(D;\widehat{D}) > 0$ for every $D \neq \widehat{D}$, which is equivalent to $L_- > 0$. Equation (148) then makes $\widehat{D}$ the unique population optimum.

Suppose D^* is unique. For every other transversal,

$$\widehat{\mu}(D) - \widehat{\mu}(D^*) \geq (\kappa_{\mathcal{C}}(\mu) - r_n)\,|D \triangle D^*|. \tag{149}$$

Hence $r_n < \kappa_{\mathcal{C}}(\mu)$ makes D^* the unique empirical minimizer. Applying the confidence subtraction once more gives

$$\underline{\Delta}(D; D^*) \geq (\kappa_{\mathcal{C}}(\mu) - 2r_n)\,|D \triangle D^*|, \tag{150}$$

which is positive for all competitors when $r_n < \kappa_{\mathcal{C}}(\mu)/2$. This proves the theorem.

G.3 Computing the selective labels

The confidence set need not be enumerated. Define

$$c_e = \begin{cases} u_e, & e \in \widehat{D}, \\ \ell_e, & e \notin \widehat{D}. \end{cases} \qquad \underline{\Delta}(D; \widehat{D}) = \sum_{e \in D} c_e - \sum_{e \in \widehat{D}} u_e. \tag{151}$$

Membership in $\mathcal{A}_\alpha$ is therefore a threshold test on a weighted feedback-edge objective. For each edge, one solve with its label forced opposite to $\widehat{D}$ determines whether any repair in the set can reverse that label. The whole selective output takes one unconstrained and at most $|E|$ forced-opposite solves. Weighted feedback-edge optimization is NP-hard in the worst case (Karp, 1972). The formulation avoids enumerating the confidence set but retains this computational cost. The modular protocol below allows one second per forced query and records timeouts as computationally unresolved.

G.4 Proof of Proposition G.3

The reference D_0 need not minimize either the empirical or population score. Define

$$\mathcal{A}_\alpha(D_0) = \{D \in \mathcal{T} : \underline{\Delta}(D; D_0) \leq 0\}.$$

Let

$$P = \left\{ z \in [0,1]^E : \sum_{e \in C} z_e \geq 1 \text{ for every directed cycle } C \right\},$$

$$c_e = \begin{cases} u_e, & e \in D_0, \\ \ell_e, & e \notin D_0, \end{cases} \qquad \tau_0 = \sum_{e \in D_0} u_e.$$

For an active edge e, let P_e^{opp} impose $z_e = 0$ when $e \in D_0$, and $z_e = 1$ otherwise. The relaxation reports the D_0 label precisely when

$$\min_{z \in P_e^{\mathrm{opp}}} c^\top z > \tau_0. \tag{152}$$

On the simultaneous coverage event, Eq. (148) holds with D_0 in place of $\widehat{D}$. If D^* is any population optimum, then $\mu(D^*) - \mu(D_0) \leq 0$, and hence $\underline{\Delta}(D^*; D_0) \leq 0$. Thus every population optimum belongs to $\mathcal{A}_\alpha(D_0)$, even when D_0 is obtained by a greedy repair.

For the costs in Proposition G.3, direct expansion gives

$$\underline{\Delta}(D; D_0) = \sum_{e \in D} c_e - \tau_0. \tag{153}$$

Every incidence vector $\mathbf{1}_D$ of a feasible repair intersects every directed cycle, so $\mathbf{1}_D \in P$. Suppose that $D \in \mathcal{A}_\alpha(D_0)$ gives edge e the label opposite to D_0. Then $\mathbf{1}_D \in P_e^{\mathrm{opp}}$, while Eq. (153) gives $c^\top \mathbf{1}_D \leq \tau_0$. This contradicts Eq. (152). The same argument covers an empty forced face by assigning it optimum $+\infty$. Therefore every repair in the confidence set, and in particular every population optimum, shares the reported label.

The exact forced-opposite problem minimizes the same objective over the integer points of P_e^{opp}. Its LP optimum is no larger than its integer optimum. Any label certified by the relaxation is consequently also certified by an exact forced-opposite query using the same reference D_0; this need not be the confidence set centered at a different empirical minimizer. The relaxation can lose labels only by abstaining.

It remains to justify the computational claim. For a candidate point $z \in [0,1]^E$, a cycle inequality is violated exactly when the directed graph contains a cycle of total z-weight below one. Since the weights are nonnegative, a minimum-weight directed cycle can be found by shortest paths, which supplies a polynomial-time separation oracle. The separation–optimization equivalence of Grötschel et al. (1981) gives a polynomial-time algorithm for the rational LP. Our implementation uses repeated shortest-cycle separation with a standard LP solver. This practical cutting-plane loop is not itself claimed to have a strongly polynomial iteration bound.

For each LP, the implementation reconstructs a weak-duality bound from the row multipliers and box constraints; integer queries use the reported mixed-integer linear program (MILP) dual bound, never a primal objective in its place. Missing, nonfinite, or inconsistent bounds stop the query with an error, rather than a statistical abstention or a family-forced label. A floating-point slack is subtracted before a bound can certify a label. These numerical checks do not constitute an exact-arithmetic verification of the solver.

G.5 Two-point rate calibration

A standard two-point calculation gives the sample scale associated with this score margin. It concerns selection between the two specified score distributions; it is not a bound on general DAG learning.

Take a directed two-cycle with edge set $E = \{e_1, e_2\}$. Its minimum transversals are $\{e_1\}$ and $\{e_2\}$. Fix $a > \kappa$ and consider independent Gaussian score coordinates with mean vectors

$$\mu_+ = (a-\kappa, a+\kappa), \qquad \mu_- = (a+\kappa, a-\kappa) \tag{154}$$

and covariance $\sigma^2 I_2$. The optimal deletion is different under the two models, and Eq. (143) equals κ under both. Their n-sample Kullback–Leibler divergence is

$$\mathrm{KL}(P_+^n, P_-^n) = \frac{4n\kappa^2}{\sigma^2}. \tag{155}$$

For $\delta \in (0, 1/4)$, any selector that succeeds under both models induces a test with both error probabilities at most δ. We include the short argument for the testing bound. Write $P = P_+^n$, $Q = P_-^n$, and let p, q be their densities. Define the affinity $A = \int \sqrt{pq}$ and total variation $V = \frac{1}{2}\int |p-q|$. Jensen's inequality gives $\log A = \log \mathbb{E}_P \sqrt{q/p} \geq -\mathrm{KL}(P,Q)/2$. Applying Cauchy–Schwarz to $|p-q| = |\sqrt{p}-\sqrt{q}|(\sqrt{p}+\sqrt{q})$ gives $V \leq \sqrt{1-A^2}$. The sum of the two testing errors is at least $1-V$, so

$$2\delta \geq 1 - V \geq 1 - \sqrt{1-A^2} \geq \tfrac{1}{2}A^2 \geq \tfrac{1}{2}e^{-\mathrm{KL}(P,Q)}.$$

Substituting the Gaussian divergence yields

$$n \geq \frac{\sigma^2}{4\kappa^2} \log \frac{1}{4\delta}. \tag{156}$$

General instance-dependent best-set bounds are substantially sharper (Chen et al., 2017). We use this familiar rate only to calibrate the score margin.

G.6 Controlled and random-graph audits

The two-cycle calculation has an exact error curve. Under P_+, the empirical score difference is Gaussian with mean 2κ and variance $2\sigma^2/n$. With Φ denoting the standard normal cumulative distribution function, the wrong-deletion probability is

$$\Phi\Big(-\sqrt{2n}\,\kappa/\sigma\Big). \tag{157}$$

We use four margins, six values of $n\kappa^2/\sigma^2$, and 200,000 replications per cell. The simulated curve differs from the formula by at most 0.0016; the fitted log–log slope of the sample size required for error 0.05 is -2.000.

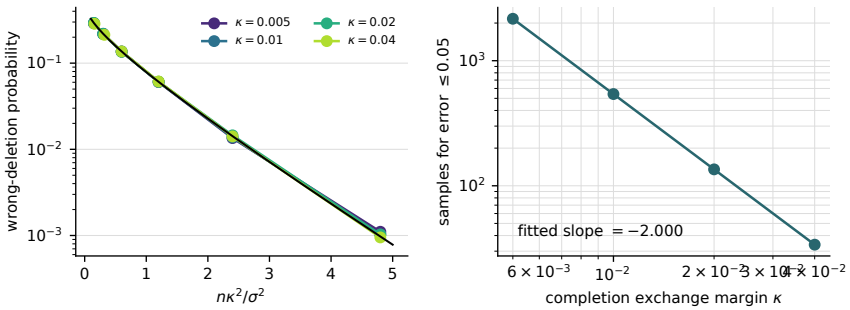


Figure 12: Statistical selection on the directed two-cycle. Error curves collapse under the scaled budget $n\kappa^2/\sigma^2$ (left), and the sample size for fixed error has slope -2 in κ (right).

The random-graph audit uses the 160 candidate supports obtained by all-node Lasso in the modular score experiment. We hold these supports fixed and evaluate the score with an independent 20,000-sample proxy. For each screen, we minimize the normalized gap over all alternative transversals, rather than normalizing only the unnormalized runner-up gap. The distinction matters because a more costly competitor can differ on more edges and have a smaller normalized gap. The proxy is used only for the audit, never for selection or certification.

To compute the ratio, let $d(D) = |D \triangle D^*|$ and solve

$$\Phi(t) = \min_{D \neq D^*} \{\mu(D) - \mu(D^*) - t d(D)\}.$$

With binary deletion variables, $d(D) = |D^*| + \sum_{e \notin D^*} z_e - \sum_{e \in D^*} z_e$, so each subproblem is another weighted feedback-edge solve with modified costs and one no-good constraint. Start at the feasible runner-up ratio. A negative minimum supplies a smaller ratio; at zero, no smaller ratio exists. There are only finitely many feasible ratios, giving finite termination in exact arithmetic. The numerical routine also records a solver lower bound $b \leq \Phi(t)$. Since $d(D) \geq 1$,

$$t + \min(0, b) \leq \kappa_{\mathcal{C}} \leq t.$$

The implementation retains a roundoff allowance and reports this interval, not a claim of exact arithmetic for the integer solver.

For a concrete counterexample to the runner-up shortcut, take the edges $1 \to 0$ of cost 9, $0 \to i$ of cost 4 and $i \to 1$ of cost 100 for $i \in \{2, 3, 4\}$, and a disjoint two-cycle with costs 1 and 3. The best deletion set costs 10. Switching the deletion in the two-cycle costs 12 and changes two edges, giving ratio 1. Deleting the three branches instead of the common back edge costs 13 and changes four edges, giving ratio $3/4$. Exhaustive enumeration of all 2^9 deletion subsets confirms that $3/4$ is the minimum. This case and random small supports are included as regression tests.

Table 16: Completion-margin audit over 160 formal datasets per holdout size. The margin minimizes the ratio under an independent 20,000-sample score proxy. Here r_n denotes the paired confidence radius, not the maximal completion size. The sufficient-condition column uses the numerical margin lower bound.

n	median κ	median $2r_n/\kappa$	$2r_n < \kappa$	unique cert.
500	1.16e-05	809.8	0.6%	2.5%
1,000	1.16e-05	522.3	1.9%	2.5%
5,000	1.16e-05	204.0	2.5%	4.4%
20,000	1.16e-05	95.9	4.4%	7.5%

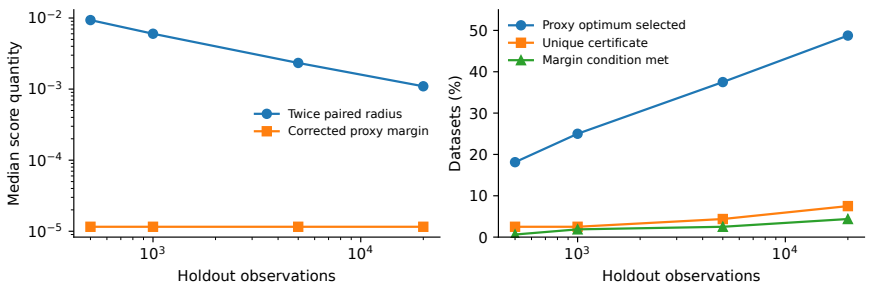


Figure 13: Why unique certificates remain rare. The median confidence radius is far larger than the oracle-proxy exchange margin (left), even as graph selection and certification improve with holdout size (right).

All 640 regret estimates from the independent high-sample proxy lie at or below their reported bounds. The median exchange margin is 1.16×10^{-5} at every nested holdout size. At $n = 20{,}000$, the empirical optimizer selects the proxy deletion set in 48.75% of runs, but only 7.5% are uniquely certified. The median resolution ratio remains 95.9, so the gap between selection and proof is consistent with near-tied transversals. Because the oracle margin is itself estimated, the column $2r_n < \kappa$ is a diagnostic rather than a population certificate. Minimizing over all competitors gives a normalized gap smaller than the runner-up shortcut by more than 10^{-10} in 24 of the 160 clipped-score screens, although the medians agree. The direct regret and forced-opposite certificates do not use the shortcut.

G.7 Selective output and score dependence

To measure resolution before unique-graph certification, we record shared labels across $\mathcal{A}_\alpha$, singleton rates, and proxy agreement, using the regret experiment's simultaneous empirical-Bernstein intervals. Varying the time cap separates statistical abstention from unfinished optimization.

Table 17: Selective repair over 160 formal datasets per holdout size under the one-second main protocol. "Statistical" means that an optimally solved forced-opposite query remains inside the confidence set; "computational" means that the query timed out. "Proxy retained" checks the independent oracle-proxy optimum.

n	decided	statistical	computational	singleton	proxy retained
500	3.8%	96.2%	0.0%	2.5%	100.0%
1,000	6.0%	94.0%	0.0%	2.5%	100.0%
5,000	16.5%	83.5%	0.0%	4.4%	100.0%
20,000	36.8%	63.2%	0.0%	7.5%	100.0%

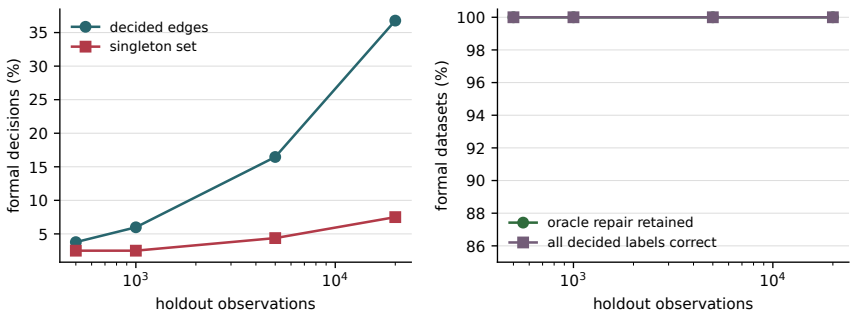


Figure 14: Selective rather than forced output. More edges become decidable as the holdout grows (left); the independent high-sample audit checks set and label coverage (right).

Table 18: Compute-budget sensitivity at $n = 20{,}000$ over the same 160 datasets. Solved queries are reused as the cap increases; only previous timeouts are retried.

cap	decided	statistical	timeout	singleton
0.01s	34.6%	48.0%	17.3%	7.5%
0.1s	36.8%	62.9%	0.3%	7.5%
1s	36.8%	63.2%	0.0%	7.5%
10s	36.8%	63.2%	0.0%	7.5%

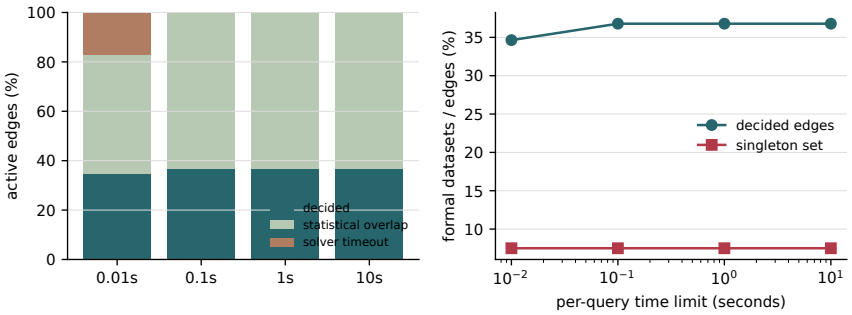


Figure 15: Statistical and computational abstention are reported separately. A 0.01-second cap leaves a visible timeout component; one second eliminates it on the formal suite.

Under the one-second protocol, every query finishes. As holdout size grows, the mean decided fraction rises from 3.8% to 36.8%, while the singleton rate rises from 2.5% to 7.5%. All 640 proxy-optimal repairs remain in the reported sets, and all reported labels agree with the proxy. At $n = 20{,}000$, a 0.01-second cap leaves 17.3% of queries unfinished; this falls to 0.3% at 0.1 seconds and zero at one second. The 63.2% statistical abstention at that budget persists after timeouts have been removed.

G.8 Scaling audit for exact and relaxed queries

The $d = 20$ main certificate suite was chosen to keep the statistical audit from being obscured by a failed combinatorial solve. We separately test the computational boundary on $d \in \{20, 40, 60, 80, 100\}$. The frozen grid has ER2 and SF2 graphs, standard and mixed-strength Gaussian SEMs, three seeds, 500 training observations, and 20,000 independent holdout observations: 60 datasets in total. A Lasso screen fixes the active directed support. A deterministic greedy repair supplies the feasible D_0, which is permitted by Proposition G.3; the experiment does not require an exact solve to enter the confidence set.

The uncapped LP relaxation audits every active edge. For comparison, each dataset has 20 evenly spaced active-edge indices fixed before any score is evaluated. Their forced-opposite integer queries receive one second each. Ground truth is not used by either path.

Table 19: Scaling audit. Percentages in the upper panel are edge- or query-weighted within dimension. The lower panel separates solved statistical overlap from exact-solver timeout by candidate-support size. LP sec. is the median time to audit every active edge.

d	sets	edges	LP decided	LP sec.	exact decided	timeout
20	12	85.7	6.7%	0.32	15.9%	0.0%
40	12	101.5	5.9%	0.38	9.2%	0.0%
60	12	203.8	1.3%	0.68	7.3%	16.7%
80	12	122.2	5.5%	0.53	22.7%	0.0%
100	12	199.1	0.7%	0.74	0.8%	0.0%

candidate edges	sets	LP sec.	exact decided	statistical	timeout
≤ 100	23	0.17	26.7%	73.3%	0.0%
101–200	26	0.58	2.7%	97.3%	0.0%
201–400	9	2.03	0.0%	100.0%	0.0%
> 400	2	16.13	0.0%	2.5%	97.5%

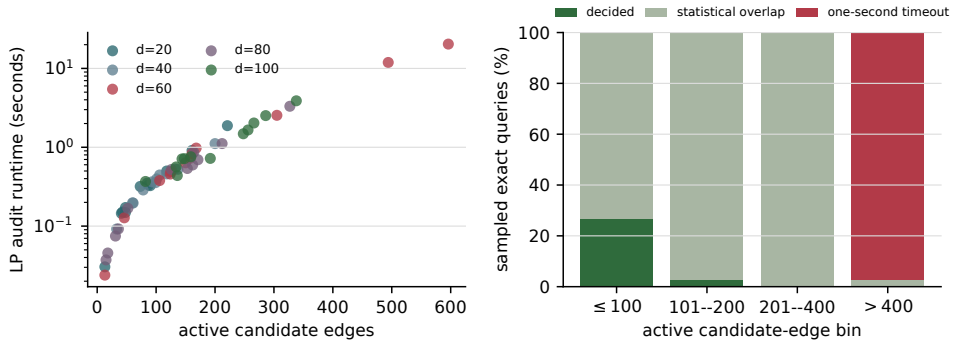


Figure 16: Computational audit. The relaxed path audits every active edge (left). Exact-query outcomes (right) show statistical overlap up to 400 candidate edges; the 39 one-second timeouts all occur in the two supports above 490 edges.

Of 1,179 sampled integer queries, 131 are certified, 1,009 are statistically unresolved, and 39 time out. On completed queries, the LP labels coincide with the exact labels. The uncapped LP completes every all-edge audit; median runtime ranges from 0.32 to 0.74 seconds by dimension, though the two supports above 490 edges take 11.90 and 20.35 seconds. At $d = 100$, it certifies 0.7% of active edges. The relaxation avoids query-time failures in this suite but cannot remove weak score separation. Exact-solver timeouts concentrate in the densest candidate supports rather than the largest ambient dimension.

G.9 Frontend-agnostic certificate audit

We compare NOTEARS, DAGMA, SDCD, and GOLEM as screens for the same modular certificate, using the implementations described in Appendix H.7. Each contributes its top $2d$ off-diagonal magnitudes with deterministic ties at $d = 20$. The downstream procedure is otherwise identical: regressors and clipping are frozen on $n = 500$ training samples, simultaneous empirical-Bernstein intervals use an independent holdout of 5,000 samples, and a separate 20,000-sample draw is consulted only after all labels have been fixed. The grid contains ER2/ER4/SF2/SF4, Gaussian, non-Gaussian, mixed-strength, and heteroscedastic SEMs, and five seeds, for 320 frontend–dataset pairs.

Table 20: The same selective certificate after four frontend screens. Coverage refers to simultaneous coverage of all screened deletion costs; errors compare LP-certified labels with the independent score proxy.

Frontend	Screen recall	Labels (%)	Coverage (%)	Errors
NOTEARS	67.3	29.5	100.0	0
DAGMA	95.3	51.2	98.8	0
SDCD	63.1	13.2	100.0	0
GOLEM	94.8	46.1	98.8	0

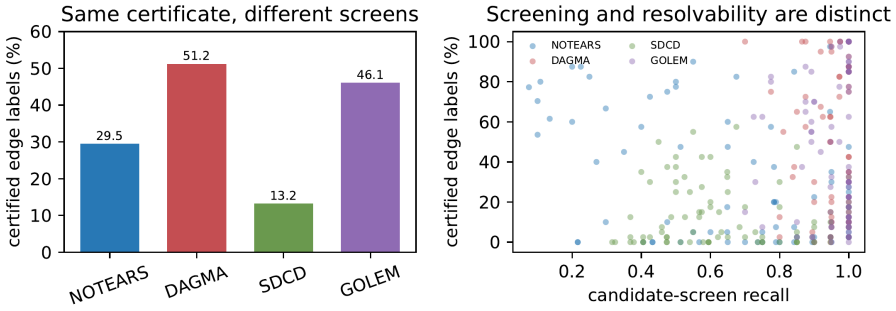


Figure 17: Frontend-agnostic audit. The common certificate resolves different fractions of the four screened supports (left). Candidate recall and selective resolvability are distinct: a broad screen may retain true edges without creating a large holdout margin between feasible repairs (right).

Regret coverage is 100% for every frontend, and no LP-certified label disagrees with the score proxy. Simultaneous interval coverage is 98.75% for DAGMA and GOLEM and 100% for NOTEARS and SDCD. Table 20 reports resolution together with screen recall: retaining more generating edges and resolving more score decisions are distinct outcomes.

The exchange margin belongs to the candidate support, deletion score, and fitting protocol jointly. We therefore repeat the 160-dataset high-sample audit with four dimensionless scores: the frozen clipped score used by the certificate, the same frozen regressors without clipping, oracle-refitted normalized quadratic loss, and oracle-refitted Gaussian profile loss. The two oracle-refitted variants serve as score diagnostics; selection continues to use the frozen clipped score.

Table 21: Score sensitivity on the same 160 frozen Lasso supports. Relative κ divides by the median absolute edge cost. "Unresolved" means that the numerical margin interval includes zero, not a proven exact tie. SHD evaluates the score-optimal repair, whose selection does not use the margin estimate.

Score	median κ	relative κ	unresolved	SHD
Frozen clipped	1.16e-05	4.31e-03	0.0%	34.90
Frozen raw	2.09e-05	8.38e-04	0.0%	37.33
Refit quadratic	2.94e-07	2.07e-05	15.6%	42.49
Refit profile	7.88e-08	1.92e-06	43.8%	38.79

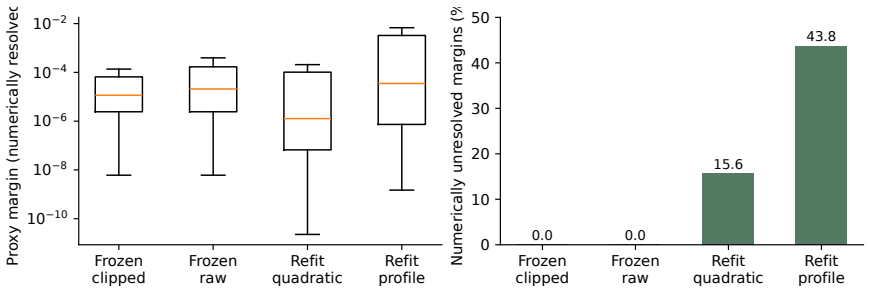


Figure 18: The completion margin is score-conditional. Replacing clipping, refitting regressors, or changing the loss moves the margin distribution and changes the fraction of margins indistinguishable from zero at the solver's numerical precision.

The median absolute margin ranges from 7.88×10^{-8} for refitted profile loss to 2.09×10^{-5} for the frozen unclipped score. The numerical lower bound is nonpositive in 43.75% of profile-score datasets and 15.63% of refitted-quadratic datasets, but not for either frozen score. The maximum numerical interval width is 1.72×10^{-12} across all 640 score audits. Relative to the runner-up shortcut, the ratio decreases by more than 10^{-10} in 10 unclipped, 9 refitted-quadratic, and 4 profile-score screens, in addition to the 24 clipped-score screens. The associated SHD varies nonmonotonically with the margin. The primary audit's value of 1.16×10^{-5} characterizes its frozen score and frontend; changing either can change the margin.

H Interaction-Aware Certificate: Proof and Unified Audit

The main certificate scores each parent set jointly, retaining interactions that separate edge costs omit. We derive the local score intervals and the global guarantee, then express adjacency and collider queries in the parent-set integer program. The experiments measure coverage, feature resolution, and computational cost separately. The unified frontend audit is the experiment summarized in the main text. The Lasso validation-budget experiment and bounded-parent scaling test use different screens and parent caps, specified in their own subsections.

For each compatible DAG G, define

$$Q(G) = \sum_j (R_{j,\mathrm{pa}_G(j)} + a_{j,\mathrm{pa}_G(j)}), \qquad G_0 \in \mathcal{G}(\mathcal{P}). \tag{158}$$

The reference G_0 is arbitrary and may be data-dependent. The benchmark chooses the empirical holdout minimizer, but the theorem also covers a compatible graph supplied by a continuous, constraint-based, or score-based frontend.

Condition throughout on the training data: the candidate sets, predictors, loss bound $B > 0$, and penalties are fixed before observing the independent holdout. Each observation $X_i \in \mathbb{R}^d$ is drawn i.i.d. from its population law, and $\ell_{j,S}(X_i) \in [-B, B]$. All expectations and coverage probabilities below are conditional on training and refer to this holdout law. For $n \geq 2$ observations, write $R_{j,S} = \mathbb{E}\ell_{j,S}(X)$, $\widehat{R}_{j,S} = n^{-1}\sum_i \ell_{j,S}(X_i)$, and $M = \sum_j |\mathcal{P}_j|(|\mathcal{P}_j| - 1)$. When $M = 0$, there are no alternative local pairs and no intervals to estimate. For $M > 0$ and each distinct ordered pair $S, T \in \mathcal{P}_j$, let $\widehat{V}_{j,S,T}$ be the sample variance, with denominator $n - 1$, of $\ell_{j,S}(X_i) - \ell_{j,T}(X_i)$, and set

$$r_{j,S,T} = \sqrt{\frac{2\widehat{V}_{j,S,T}\log(2M/\alpha)}{n}} + \frac{28B\log(2M/\alpha)}{3(n-1)}. \tag{159}$$

For $S_{0j} = \mathrm{pa}_{G_0}(j)$, define

$$g_{j,S_{0j}} = 0, \qquad g_{j,S} = \widehat{R}_{j,S} - \widehat{R}_{j,S_{0j}} + a_{j,S} - a_{j,S_{0j}} - r_{j,S,S_{0j}} \quad (S \neq S_{0j}). \tag{160}$$

For the graph-level conclusions, define

$$L = \min_G \Gamma(G), \qquad L_- = \min_{G \neq G_0} \Gamma(G). \tag{161}$$

On the simultaneous interval-coverage event,

$$Q(G_0) - \min_{G \in \mathcal{G}(\mathcal{P})} Q(G) \le -L, \tag{162}$$

and $L_- > 0$ makes G_0 the unique population minimizer. We now prove these regret and confidence-family statements together with the feature claim in Theorem 4.1.

H.1 Proof of Theorem 4.1

Condition on the training sample. The parent-set families, fitted predictors, losses, and penalties are then fixed, while the validation observations remain i.i.d. For an ordered pair $S, T \in \mathcal{P}_j$, put

$$D_i^{j,S,T} = \ell_{j,S}(X_i) - \ell_{j,T}(X_i).$$

Because each loss lies in $[-B, B]$, the paired difference lies in $[-2B, 2B]$, an interval of length $4B$. The empirical Bernstein inequality of Maurer & Pontil (2009, Theorem 4), with failure probability α/M, gives

$$R_{j,S} - R_{j,T} \ge \widehat{R}_{j,S} - \widehat{R}_{j,T} - r_{j,S,T} \tag{163}$$

with the radius in Eq. (159). A union bound over all M ordered local pairs makes Eq. (163) simultaneous with probability at least $1 - \alpha$. Covering the entire ordered-pair family is essential when G_0 depends on the same holdout.

Work on this simultaneous event and set $T = S_{0j}$. Adding the deterministic penalty difference to Eq. (163) yields

$$R_{j,S} + a_{j,S} - R_{j,S_{0j}} - a_{j,S_{0j}} \ge g_{j,S}$$

for every local alternative, including equality at $S = S_{0j}$. Hence, for every candidate DAG G,

$$Q(G) - Q(G_0) \ge \sum_j g_{j,\mathrm{pa}_G(j)}. \tag{164}$$

Minimizing both sides over $G \in \mathcal{G}(\mathcal{P})$ gives $\min_G Q(G) - Q(G_0) \ge L$. Rearrangement proves Eq. (162). Notice that $L \le 0$ because G_0 is feasible and has robust relative cost zero.

Let G^* be any population minimizer. Since $Q(G^*) - Q(G_0) \le 0$, Eq. (164) gives $\sum_j g_{j,\mathrm{pa}_{G^*}(j)} \le 0$. Hence every such G^* lies in Eq. (7). If $L_- > 0$, Eq. (164) is strictly positive for every $G \ne G_0$; the reference graph is therefore the unique minimizer of Q. Finally, if a population minimizer had the feature label opposite to G_0, its robust relative cost would be at least $L_\varphi^{\mathrm{opp}} > 0$, contradicting its membership in $\mathcal{A}_\alpha(G_0)$. This proves the feature statement.

H.2 A sufficient feature-resolution margin

Assume both labels of φ are feasible in $\mathcal{G}(\mathcal{P})$, and use $\kappa_\varphi(G_0)$ from Eq. (9). On the simultaneous ordered-pair event, both signs of each error are bounded: writing $\delta_{j,S} = R_{j,S} - R_{j,S_{0j}}$,

$$|\widehat{R}_{j,S} - \widehat{R}_{j,S_{0j}} - \delta_{j,S}| \le r_{j,S,S_{0j}}.$$

The reverse ordered pair has the same sample variance and radius. Subtracting the radius once more in the definition of $g_{j,S}$ gives

$$g_{j,S} \ge \delta_{j,S} + a_{j,S} - a_{j,S_{0j}} - 2r_{j,S,S_{0j}}.$$

At an unchanged parent set, both the true difference and g are exactly zero. If all used radii are at most ε_n, summing only over changed nodes yields

$$\Gamma(G) \ge Q(G) - Q(G_0) - 2\varepsilon_n m(G, G_0).$$

Every opposite-feature graph has $m(G, G_0) \ge 1$. Its right-hand side is at least $(\kappa_\varphi(G_0) - 2\varepsilon_n)m(G, G_0)$, strictly positive when $\kappa_\varphi(G_0) > 2\varepsilon_n$. Minimization proves the sufficient condition in the main text. It concerns only opposite-feature competitors; same-feature ties need not be resolved. When $0 < \kappa_\varphi(G_0) \le 2\varepsilon_n$, certification may still succeed because the uniform radius can be conservative. In contrast, if $\kappa_\varphi(G_0) \le 0$, the finite family contains an opposite-label witness $\bar{G}$ with $Q(\bar{G}) - Q(G_0) \le 0$. On the coverage event,

$$L_\varphi^{\mathrm{opp}} \le \Gamma(\bar{G}) \le Q(\bar{G}) - Q(G_0) \le 0.$$

The present rule therefore cannot certify that reference label on this event. Positivity of the margin is necessary; the stronger uniform-radius condition is sufficient.

The arithmetic implications for this margin and for the parametric exchange ratio are checked in the supplementary Lean module `RevisionAudit`. Those checks assume their stated error or solver bounds. They do not verify floating-point MILP output or derive the statistical assumptions of a particular data model.

H.3 Optimization and target of the certificate

For each child, the implementation caps the training-screened candidate parents and fits every subset by least squares. The unified audit uses four candidate parents; the separate validation-budget experiment uses eight. Centering and regression use training data only. In the loss below, x is centered by the training mean, $\widehat{f}_{j,S}$ is the frozen linear predictor from parent set S, and $\widehat{v}_j$ is the training mean squared centered response, floored at the smallest positive normal floating-point number. The validation loss is

$$\ell_{j,S}(x) = \mathrm{clip}_{[-B,B]}\left\{ \frac{(x_j - \widehat{f}_{j,S}(x))^2 - (x_j - \widehat{f}_{j,F_j}(x))^2}{2\widehat{v}_j} \right\}, \tag{165}$$

where F_j is the full screened parent set. We use $a_{j,S} = \lambda_{\mathrm{sp}}|S|$. The clipping operator is $\mathrm{clip}_{[-B,B]}(u) = \min(B, \max(-B, u))$. Before clipping, the common full-model baseline cancels in every paired comparison and therefore changes neither the score difference nor its variance. After clipping, the baseline and B jointly define the bounded target; they need not preserve the unclipped minimizer or reduce paired variance. Division by the frozen training response variance makes the score dimensionless.

Both minimizations in Eq. (161) are standard exact Bayesian-network structure-learning problems over a restricted parent-set family. We use the parent-set integer program and add cluster inequalities until the selected support is acyclic (Cussens, 2011). A no-good constraint excludes G_0 when computing L_-. These optimization devices, including parent-set screening and cluster cuts, are established tools (Cussens, 2020). We use them to optimize the robust local gaps in Eq. (160); the resulting objective converts simultaneous score intervals into the global guarantee of Theorem 4.1.

The feature queries use binary parent-set variables $x_{j,S} \in \{0, 1\}$, with $\sum_{S \in \mathcal{P}_j} x_{j,S} = 1$ for each child j. If $x_{j,S} = 1$ selects parent set S for child j, write

$$a_{uv} = \sum_{S \in \mathcal{P}_v : u \in S} x_{v,S}, \qquad s_{uv} = a_{uv} + a_{vu} \quad (u < v). \tag{166}$$

Use $s_{vu} = s_{uv}$ when the pair is written in reverse order. On a DAG, s_{uv} is the skeleton-adjacency indicator. Thus an opposite skeleton or directed-edge label is one linear equality. An unshielded collider $u \to v \leftarrow w$ is present exactly when $a_{uv} = a_{wv} = 1$ and $s_{uw} = 0$. Its presence is imposed by these three equalities; its absence is the single inequality

$$a_{uv} + a_{wv} - s_{uw} \le 1. \tag{167}$$

One forced-opposite solve therefore computes Eq. (8). Skeletons and unshielded colliders characterize a Markov equivalence class by the result of Verma and Pearl reproduced in Chickering (2002, Theorem 1). This projection does not require the data to resolve arbitrary orientations inside the class.

The theorem uses exact minima. In the implementation, a guarded MILP dual bound must be strictly positive to certify an opposite-label query; a positive feasible objective alone is insufficient. Missing or invalid bounds raise an error. Infeasible opposite labels are recorded separately as family-forced. The parent-set audit has no query time limit and aborts on a non-optimal status; it does not count unfinished searches as statistical abstentions. The time-limited modular audit records these separately as computationally unresolved.

Screening and the score define the target of the certificate. Screening can omit a true parent, and the population minimizer of the clipped regularized score can differ from the causal DAG. The population-regret guarantee is restricted to $\mathcal{G}(\mathcal{P})$ and does not imply causal support recovery.

H.4 EXACT-POPULATION CALIBRATION

We use four explicitly specified finite distributions to check the implemented certificate against an exactly computable population target. The predictors use rational least-squares coefficients computed from these laws, with no training-sample estimation. Each node has two candidate parents, and every subset is allowed. We enumerate all 25 compatible three-node DAGs. The loss in Eq. (165) uses the fixed population response variance, $B = 1/50$, and penalty $\lambda_{\mathrm{sp}} = 1/1000$. Exact rational evaluation over all eight atoms identifies every population minimizer, including ties.

For each law, 200 independent draws use seed sequences indexed by the law and replication, with master seed 20260906. Holdout sizes $500, 2000, 5000, 20{,}000$ are nested within each draw. The certificate receives only the frozen loss family and holdout values; the exact target is used afterwards. We query all six directed edges, three adjacencies, and three unshielded colliders. Both labels of every query are feasible in the family, so none is family-forced. For each draw, enumeration also checks the reported solver lower bounds against the minima of Γ over all opposite-label graphs.

For the strong and weak exchangeable laws,

$$\Pr(X = x) = \frac{1 + \rho x_0 x_1}{8}, \qquad x \in \{-1, 1\}^3, \qquad \rho \in \{3/5, 3/20\}.$$

Each has two population-optimal graphs, containing respectively $0 \to 1$ or $1 \to 0$, and no other edges. The third law takes independent Rademacher variables Z_0, Z_1, Z_2 and sets $X_0 = Z_0$, $X_1 = Z_1$, $X_2 = (Z_0 + Z_1)/2 + Z_2$. Its generating graph is a collider, but the specified clipped regularized score has the empty graph as its unique population minimizer. We retain this case to distinguish calibration of the score guarantee from agreement with a generating DAG. Each Rademacher variable is $+1$ or -1 with equal probability.

The cancellation in this third law can be evaluated explicitly. For the empty parent set of node 2, the eight unclipped losses contain two values of 1, two of $-1/3$, and four zeros. Clipping at $1/50$ makes their expectation zero. Rational evaluation gives $R_{j,S} = 0$ for all twelve local models, hence

$$Q(G) = |E(G)|/1000.$$

The empty optimum is selected entirely by the penalty, not by a retained average prediction-risk difference.

For a positive control, let independent Y_i be Bernoulli with probabilities $(1/4, 1/4, 7/8)$, set $Z_i = Y_i - \mathbb{E}Y_i$, and use

$$X_0 = Z_0, \qquad X_1 = Z_1, \qquad X_2 = 2Z_0 + Z_1/8 + Z_2.$$

The score, parent family, and penalty are unchanged. Its twelve local risks, listed exactly in the supplement, yield the unique optimum $0 \to 2 \leftarrow 1$, with score $43/4000$ and gap $1/1000$ to the next graph. This law was constructed using population scores: 147 symmetric-root candidates gave no positive collider optimum; an asymmetric-root search with a fixed random seed found this law on its 66th draw. All 213 examined laws are logged. The selected law was frozen before drawing any holdout sample. It serves as a constructed positive control; its success rate does not estimate performance on random SEMs.

Table 22: Exact-population calibration: 200 replications per cell. All population optima were retained and no feature error occurred in any cell. Adjacency and collider percentages count all three queries of that type, including negative labels. Singleton means a certified unique graph.

Law	Holdout n	Adj. decided (%)	Coll. decided (%)	Singleton (%)
Strong	500	33.2	33.3	0.0
	2,000	33.3	66.7	0.0
	5,000	33.3	66.7	0.0
	20,000	33.3	89.3	0.0
Weak	500	0.0	0.0	0.0
	2,000	7.3	65.8	0.0
	5,000	33.2	66.7	0.0
	20,000	33.3	66.7	0.0
Collider law	500	0.0	0.0	0.0
	2,000	0.0	28.3	0.0
	5,000	9.2	98.3	0.0
	20,000	81.0	100.0	65.5
Positive control	500	33.3	33.3	0.0
	2,000	39.2	39.3	7.5
	5,000	100.0	100.0	100.0
	20,000	100.0	100.0	100.0

At $n = 20{,}000$, both exchangeable laws certify the common positive adjacency in all 200 draws and no negative adjacency. They certify, respectively, 536 and 400 negative collider labels out of 600 queries; neither has a population-optimal collider to certify as present. For the collider-generating law, all 486 adjacency and 600 collider certifications are negative, agreeing with the empty score optimum. The sign counts distinguish certified absence from positive discovery. The strong law already certifies its positive adjacency in 199 of 200 draws at $n = 500$; the weak law reaches that count at $n = 5000$.

The positive control certifies its collider as present in $0, 15, 200, 200$ of the 200 draws at the four sample sizes. At $n = 5000$ and 20,000, all six directed-edge, three adjacency, and three collider labels are resolved in every draw. No label is family-forced: all 25 DAGs remain eligible in this calibration.

All 3,200 regret checks and all-optima coverage checks pass. Zero observed feature errors in 200 independent draws gives a one-sided 95% binomial upper limit $1 - 0.05^{1/200} = 0.01487$ on the per-draw probability of any feature error, separately for each fixed law and sample size. The nested sizes are not independent additional replications, and these pointwise limits are not simultaneous across the 16 cells. This experiment calibrates a bounded, fixed-score implementation; it does not establish a coverage rate for unrestricted frontends or causal-truth labels.

H.5 Signed features and sensitivity to the frozen score

We examine signed feature labels in the 320 unified frontend runs without refitting any frontend. The training and holdout streams, top-$2d$ screens, four-parent cap, penalty, and nominal $B = 0.02$ are held fixed. For each run, we retain every queried label and its forced-opposite bound. Proxy optimization at $B = 0.01, 0.02, 0.05$ gives 960 selections for the cross-score comparison. The released records allow each label to be matched across these scores.

Table 23 separates presence from absence relative to the holdout-selected reference graph. The certified set is held fixed at the nominal score. For each alternative clipping level we compare its labels with one solver-selected optimum of the independent 20,000-sample proxy. Agreement with this finite-sample optimizer is not a simultaneous guarantee across scores or a check over all possible proxy ties. The rates are descriptive pooled-label summaries; queries within a graph are dependent.

Table 23: Signed resolution and cross-score agreement on the 320 frozen runs. The denominator of the certified percentage is the number of queries with the indicated reference label. At nominal $B = 0.02$, every certificate agrees with the proxy. The last column counts disagreements with generating graph labels, a different target. No decision is family-forced.

Feature	Label	Queries	Certified, n (%)	Agree $B = .01$	Agree $B = .05$	Truth diff.
Adjacency	Present	6,536	2,962 (45.32)	99.43%	100%	112
	Absent	3,389	80 (2.36)	100%	100%	23
Collider	Present	4,328	1,382 (31.93)	99.28%	100%	80
	Absent	5,918	1,014 (17.13)	99.70%	100%	51

At $B = 0.01$, 17 positive adjacency, 10 positive collider, and three negative collider labels change relative to their nominal certification. For a same-unit comparison, we partition all 20,171 queries by certification at $B = 0.02$ and compare each proxy label with its counterpart at that nominal score. Of the 5,438 certified labels, 30 change at $B = 0.01$ and none at $B = 0.05$, giving agreement rates of 99.45% and 100%. Of the 14,733 uncertified labels, 1,827 and 971 change, giving 87.60% and 93.41% agreement. Certification groups are not reselected at the alternative scores. These pooled comparisons show greater stability within the nominal certified group, but a label certified for one score need not hold for another.

H.6 Frozen protocol and results

This experiment varies the validation budget while keeping the training family fixed. Its Lasso screen and eight-parent cap differ from the four-frontend, four-parent protocol in Appendix H.7; their counts are reported separately.

The formal suite uses $d = 20$ and 500 training observations. A standardized all-node Lasso frontend uses $0.002\lambda_{\max}$; at most eight incoming parents per node are retained, and all of their subsets are enumerated. We set $B = 0.02$ and $\lambda_{\mathrm{sp}} = 0.001$ using disjoint development seeds 100–102, then freeze them. Formal seeds are 0–9. We test ER2, ER4, SF2, and SF4 graphs under standard Gaussian, exponential-noise, 40%-weak-edge Gaussian, and heteroscedastic Gaussian SEMs. The holdout is independent of training; its nested prefixes contain $500, 1000, 5000$, or $20{,}000$ observations. A further independent 20,000-sample proxy is used only after selection and certification. This gives 160 datasets at each holdout size and 640 records in total, not 640 independent datasets.

The "sequential local" row below is our stepwise repair control: each round uses a fresh validation batch to choose an edge deletion from a selected cycle. Its regret sums the score gaps against the best candidate deletion at each visited support. Its exact-path indicator requires every step to be certified. The global rows instead compare complete repairs or DAGs; their exact-graph indicator requires a certified unique optimum. Budget is the mean number of validation observations used, summed across rounds for the sequential control. Bound and proxy-regret entries are dataset means.

Table 24: Matched validation-budget audit over 160 formal datasets. The three rows certify different targets; their regret values are not losses on a common scale.

Method	Budget	Bound	Proxy reg.	Exact path/graph	SHD
Sequential local	21,275	0.0492	0.00048	0.0%	34.59
Global modular	20,000	0.0037	0.00007	7.5%	34.88
Global parent-set	20,000	0.0065	0.00010	0.0%	26.22

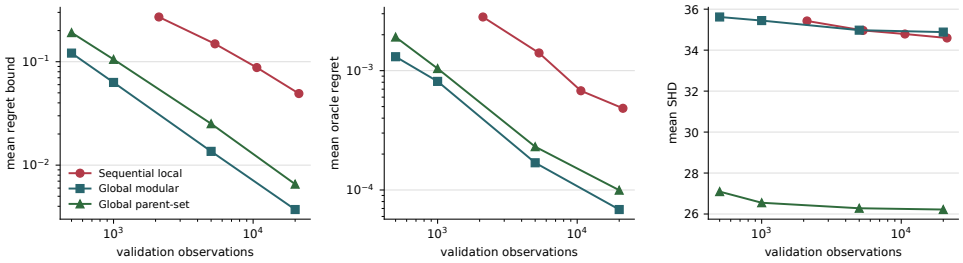


Figure 19: Matched-budget comparison of the original sequential certificate, a globally optimized modular deletion score, and the parent-set certificate. The methods certify different population losses, so regret values should be compared as tightening curves rather than as a common performance metric. Every proxy regret estimate lay at or below its reported bound.

For the parent-set target, all 640 proxy regret estimates lie at or below their reported bounds. The mean bound decreases from 0.1911 at $n = 500$ to 0.00655 at $n = 20{,}000$, while the mean proxy regret decreases from 0.00191 to 0.00010. The selected graph matches the proxy graph in 3.75% and 47.5% of runs, respectively, but L_- is never positive. Thus finite samples often select the proxy graph before the simultaneous confidence box is narrow enough to prove uniqueness.

At the largest matched budget, parent-set refitting lowers mean SHD from 34.88 to 26.22 relative to modular repair. Mean family recall is 0.692, and ER4/SF4 cells account for most remaining errors. This comparison concerns two scores on the same Lasso family; the next experiment compares them with four learning frontends on matched datasets.

The modular companion fixes one bounded loss per active-edge deletion and solves an exact weighted feedback-edge problem. At $n = 20{,}000$, its mean bound is 0.00372, its proxy regret is 0.000069, and 7.5% of graphs are uniquely certified. Its empirical modular objective is lower than that of a same-cost greedy cycle repair in every formal run. Global optimization improves this additive objective; changing to jointly refitted parent sets also changes the statistical target.

H.7 Frontend-to-feature certification

Each dataset passes through screening, point estimation, and feature certification. We measure the support retained by screening, the labels resolved by the holdout, and their agreement with the generating graph. ER2/ER4/SF2/SF4 graphs at $d = 20$, the four SEM regimes in Table 4, and seeds 0–4 give 80 datasets. Each frontend uses the same 500 training, 5,000 holdout, and 20,000 proxy observations within a dataset. The nominal certificate uses $\alpha = 0.05$, clipping level $B = 0.02$, and sparsity penalty $\lambda_{\mathrm{sp}} = 0.001$, fixed across cells.

Frontends and adaptations. We use the released linear NOTEARS and DAGMA implementations (Zheng et al., 2018; Bello et al., 2022) and the SDCD package (Nazaret et al., 2024). GOLEM uses our JAX port of the authors' equal-variance objective, with zero initialization, rather than their TensorFlow implementation (Ng et al., 2020). NOTEARS uses least squares, $\lambda = 0.01\lambda_{\max}$, and at most 30 outer iterations. DAGMA uses least squares, $\lambda = 0.03$, four continuation stages, and warm/main budgets of 5,000/10,000 iterations. SDCD uses 500 epochs per stage without internal standard scaling; its variance option is fixed by the declared equal- or unequal-noise regime. The GOLEM port uses 30,000 Adam steps at learning rate 10^{-3}, with sparsity weight 0.02 and acyclicity weight 5. For reported frontend DAGs, NOTEARS, DAGMA, and GOLEM use cutoff 0.3 followed by cycle pruning if needed; SDCD uses its package thresholded output with the same acyclicity check. Screening instead uses the unthresholded weights. The released adapters record these choices per run.

Training and evaluation samples. The same training, holdout, and proxy samples are reused across frontends within each graph–regime–seed cell. Node means and all parent-set regressions are fitted on training observations; the stored training means are also used to center holdout and proxy observations. Each frontend supplies its top-$2d$ nonzero off-diagonal entries, ordered by absolute

weight with deterministic index tie-breaking. A four-parent-per-child cap fixes the parent-set family before holdout scores are evaluated; every subset is refitted on 500 training observations. The interaction-aware score uses the holdout; the independent proxy is reserved for post-selection evaluation. The four frontends give 320 runs, with the same parent cap in every cell. In this generator, the ER2/ER4 and SF2/SF4 suffixes specify target mean total degree: the requested edge budgets are d and $2d$, respectively. The four regimes are standard Gaussian, exponential noise, 40%-weak-edge Gaussian, and heteroscedastic Gaussian SEMs. Generating supports and proxy scores are used for evaluation, not to choose a screen or a certified label.

The generator uses the NOTEARS graph and linear-SEM routines. Standard edge magnitudes are uniform on $[0.5, 2.0]$, with independent random signs. In the weak-edge regime, each edge independently uses magnitude $\mathrm{Uniform}[0.1, 0.3]$ with probability 0.4, and the standard range otherwise. Gaussian noises have unit standard deviation except in the heteroscedastic regime, which draws one standard deviation per node from $[0.5, 2.0]$; exponential noises have unit scale. Graph, coefficient, and sample draws are fixed by the reported seed.

Metrics and aggregation. Family recall is the fraction of generating directed edges retained after both screening steps. SHD counts an extra adjacency, a missing adjacency, or a reversed edge as one edit, following the NOTEARS evaluation routine. Both are averaged over datasets within each frontend. SHD evaluates the selected complete DAG, not a graph assembled from the certified labels. Adjacency queries range over unordered pairs with at least one candidate edge; collider queries range over pairs of candidate parents of one child. Each query has a present or absent reference label. Certified percentages pool query counts within each frontend and feature type; generating-graph disagreement uses certified queries as its denominator. This differs from averaging per-dataset percentages. Queries within a graph and frontends using the same dataset are dependent. Runtime medians summarize runs. "Coverage" in the learned-graph tables compares a bound with a finite-sample proxy regret; only the exact-population calibration uses known population optima.

Table 25: Frontend-to-feature certification. "Front." is the frontend DAG under the reporting rules above; "Modular" and "Parent" apply the two repair scores to the same top-$2d$ screen. "Proxy parent" uses 20,000 independent observations; label columns are certified fractions of the candidate-family queries.

Frontend	Family recall	Front. SHD	Modular SHD	Parent SHD	Proxy parent SHD	Cert. adj.	Cert. coll.
NOTEARS	57.4	15.6	19.9	20.4	20.1	19.5	16.8
DAGMA	77.6	4.5	3.6	8.5	8.6	44.5	34.6
SDCD	53.4	39.3	24.1	21.1	21.0	10.5	4.4
GOLEM	77.2	8.1	4.0	8.5	8.5	44.4	36.3

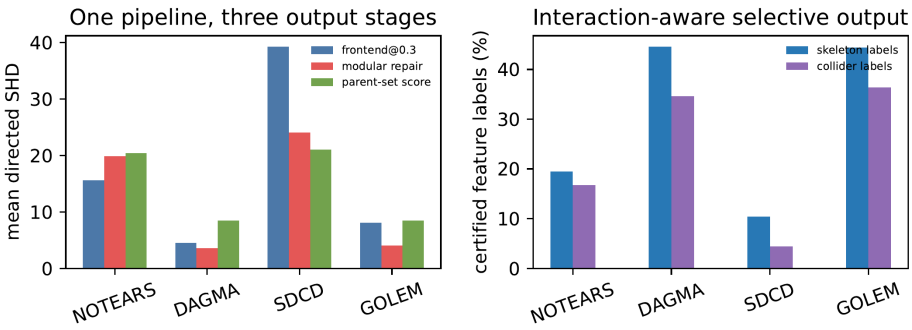


Figure 20: One frozen pipeline, reported end to end. Joint parent-set refitting changes both SHD relative to modular repair (left) and the statistical target, allowing selective skeleton and collider labels (right).

Where decisions come from. The candidate parent-set family is downward closed and contains the empty set. For any queried adjacency, the empty graph realizes absence and a graph with one candidate edge realizes presence. For any queried collider $u \to v \leftarrow w$, the empty graph realizes absence, while choosing $\{u, w\}$ for v and empty parent sets elsewhere realizes presence. Both constructions are DAGs. Thus a forced-opposite problem is feasible for every query in this benchmark; family feasibility alone certifies none of them. The exact solver is run without a time limit, and a non-optimal solver status aborts the run rather than being counted as an abstention.

Table 26: Source of feature decisions in the unified audit, as a percentage of all candidate-family queries. Family-forced and solver-timeout rates are zero for every row.

Frontend	Feature	Family-forced	Score-certified	Statistical abstention
NOTEARS	Adjacency	0.0	19.5	80.5
NOTEARS	Collider	0.0	16.8	83.2
DAGMA	Adjacency	0.0	44.5	55.5
DAGMA	Collider	0.0	34.6	65.4
SDCD	Adjacency	0.0	10.5	89.5
SDCD	Collider	0.0	4.4	95.6
GOLEM	Adjacency	0.0	44.4	55.6
GOLEM	Collider	0.0	36.3	63.7

All 320 proxy regret estimates lie at or below their reported bounds. None of the 3,042 certified skeleton labels or 2,396 certified collider labels disagrees with the corresponding proxy-optimal graph. Parent-set refitting lowers mean SHD for SDCD but trails modular repair for DAGMA and GOLEM. Its mean SHD is 14.63, close to the proxy's 14.56; modular and frontend SHDs average 12.91 and 16.87. Parent-set SHD and family recall have Pearson correlation -0.886 across the 320 runs. The small holdout–proxy difference suggests that more validation data alone would not remove the observed structural error; the candidate family and score target also matter.

Against the generating graph, 4.4% of certified skeleton labels and 5.5% of certified collider labels disagree. Proxy agreement checks the score ranking; generating-graph agreement checks its causal interpretation. Screen omissions, finite-proxy error, and a score optimum different from the generator can all contribute to the latter discrepancy.

H.8 Clipping and confidence-level sensitivity

Clipping changes the score target; the confidence level changes interval width. We compare these effects by re-evaluating all 320 frozen families at $B \in \{0.01, 0.02, 0.05\}$ and varying $\alpha \in \{0.01, 0.05, 0.10\}$ at the nominal $B = 0.02$. No frontend, candidate family, regression fit, or data split changes.

Table 27: Parent-set score sensitivity. "Clip" is the fraction of holdout candidate-loss entries changed by clipping. "H–P" compares the holdout point estimate with its same-score 20,000-sample proxy; "P–nom." compares that proxy with the nominal-(B) proxy.

(B)	α	Clip (%)	Coverage (%)	H–P (%)	P–nom. (%)
0.01	0.05	62.6	100	59.4	22.5
0.02	0.01	55.1	100	58.4	100.0
0.02	0.05	55.1	100	58.4	100.0
0.02	0.10	55.1	100	58.4	100.0
0.05	0.05	42.7	100	69.4	47.2

Proxy-regret coverage is unchanged. Changing B changes the high-sample optimum; changing α changes interval width but not the point estimate. The two controls distinguish score sensitivity from uncertainty about a fixed score, the target of Theorem 4.1.

H.9 SCREEN-BUDGET DIAGNOSIS

We compare candidate recall and selected-graph error as the screening budget grows. Starting from the primary top-$2d$ screen, we rerun all 320 frontend–dataset pairs with predeclared top-$2d$, top-$4d$, and top-$6d$ budgets. Every frontend fit, data split, four-parent cap, loss, and penalty is unchanged; no budget is selected using truth. "Raw recall" is measured before the per-child cap, whereas "family recall" describes the parent sets actually passed to the certificate.

Table 28: Frozen screen-budget audit over 960 cells. Coverage compares the holdout regret bound with the independent 20,000-sample score optimum.

Budget	Raw recall	Family recall	Parent SHD	Proxy SHD	Coverage	Time (s)
top-$2d$	80.1	66.4	14.63	14.56	100.0	0.009
top-$4d$	84.5	68.4	15.40	15.41	100.0	0.028
top-$6d$	85.4	68.5	15.73	15.67	100.0	0.058

The extra candidates raise raw recall by 5.3 percentage points, while the local cap passes only 2.1 points to the parent-set family. Mean proxy-score SHD rises from 14.56 to 15.67. The effect is frontend-dependent. At top-$4d$, proxy-parent SHD is 8.54 for DAGMA and 8.18 for GOLEM, but 21.65 for NOTEARS and 23.28 for SDCD. Enlarging the family therefore recovers some omitted edges while also admitting alternatives favored by the frozen score over the generating DAG.

H.10 BOUNDED-PARENT SCALING

The unified audit uses $d = 20$ to examine statistical resolution without unfinished exact solves. For the scaling test, we increase dimension under a fixed candidate-parent cap. A truth-blind absolute-correlation screen supplies at most two candidate parents per node. For each instance we solve the global reference certificate and 20 feature queries fixed by index before scores are evaluated: ten skeleton adjacencies and ten collider candidates. The grid contains ER2 and SF4, standard and mixed-strength Gaussian SEMs, and five seeds, giving 20 cases at each dimension. In the runtime table, "med." and "p95" denote the median and the 95th percentile across those cases.

Table 29: Interaction-aware computational scaling. Times are seconds for the global certificate (core) and the batch of 20 forced-opposite queries.

d	Cases	Parent sets	Queries	Core med./p95	Query med./p95
20	20	80	20	0.051/0.069	0.343/0.472
50	20	200	20	0.175/0.241	1.139/1.674
100	20	400	20	0.533/0.742	3.675/5.226
200	20	800	20	1.800/2.433	12.798/17.487

All 80 core solves and all 1,600 feature queries terminate. At $d = 200$, the median core time is 1.80 seconds and the median time for 20 queries is 12.80 seconds. This experiment measures the bounded-parent regime: runtimes grow superlinearly, two-parent families capture only local pair interactions, and denser candidate families retain the worst-case complexity of exact Bayesian-network structure learning. The edge-modular LP in Appendix G provides the corresponding scalable, conservative certificate.