

Addendum 303: The Degree Factor

An eighth excluded class, and why a column weight cannot localize

TOE Corpus

2026-06-13

Skill: gap-closure v1.1 — **Outcome:** Case D (mechanism class excluded; OI-287-1 lead unchanged) — **Addresses:** OI-287-1 — **Verifier:** `verify_P303.py` — all PASS

Abstract

A300 bracketed the factor missing from the A299 base $B_{nk} = \langle \psi_n | c_k \pi^k x^k | \psi_n \rangle / \lambda_n$, whose diagonal normalizes to (1, 1.12, 7.50) against the target (1, 1.5, 8), and excluded the node count as a uniform mechanism: it lifts the μ and τ rows by different relative amounts, so the factor is not a function of the family index alone. A300's stated next step was to test the complementary hypothesis, that the missing factor is a fixed function of the monomial degree k , not of the family n . This addendum runs that test. Four degree-dependent candidates are declared in advance and applied as a per-column weight $w(k)$ on the A299 base: the Taylor factorial $1/k!$, the inverse degree $1/k$, a degree-resolved strength $\pi^{-k} \kappa \Omega$, and a linear strength $\kappa \Omega k$, with $\kappa = \alpha^{5/4}$ and $\Omega = \alpha^{-1}$. None peaks on the diagonal under the column-standardized criterion, and all four carry the normalized diagonal away from (1, 1.5, 8): the closest, inverse degree, reaches (1, 0.56, 2.50) at $L^2 = 5.58$, far worse than the A299 control at $L^2 = 0.63$, and its μ entry moves to 0.56, away from 1.5 rather than toward it. A degree-dependent column weight is therefore excluded as an eighth mechanism class. The structural reason is simple and is the content of the result: a weight that depends only on k scales an entire column uniformly, so it cannot change the position of the peak within a column and cannot lift one diagonal entry relative to the column it sits in. The missing factor is not a fixed function of the monomial degree alone.

1 The Probe

Proposition 1.1 (A degree-only column weight cannot localize). *Let $B_{nk} = \langle \psi_n | c_k \pi^k x^k | \psi_n \rangle / \lambda_n$ be the A299 base, diagonal (1, 1.12, 7.50). For a weight $w(k)$ depending only on the degree k , the weighted matrix is $M_{nk} = w(k) B_{nk}$, a uniform rescaling of each column. Such a rescaling leaves $\arg \max_n$ of every column fixed, so M peaks on the diagonal exactly when B does, namely not at all. Evaluated on the four declared weights, the normalized diagonals are (1, 0.56, 1.25) for $w = 1/k!$, (1, 0.56, 2.50) for $w = 1/k$, (1, 0.36, 0.76) for $w = \pi^{-k} \kappa \Omega$, and (1, 2.24, 22.50) for $w = \kappa \Omega k$. None peaks, none lands within 20% of (1, 1.5, 8), and the closest by L^2 is the inverse degree at $L^2 = 5.58$, against the unweighted control's $L^2 = 0.63$. The pre-registered HIT criterion, a diagonal peak together with all three entries within 20%, is met by no candidate. A degree-only factor is excluded as an eighth mechanism class.*

Remark 1.2 (Why the column weight fails, and what it implies). *The failure is not numerical accident but a consequence of where the target information lives. The diagonal entry of column k is row k ; a weight $w(k)$ multiplies that entry and the two off-diagonal entries of the column by the*

same number, so the ratio of the diagonal entry to its column is invariant. Localization requires lifting the diagonal entry relative to the others in its column, which a function of k alone cannot do. The four candidates merely re-weight the already-correct relative magnitudes between columns, and since the A299 diagonal is already close in the cross-column sense, any nontrivial $w(k)$ can only spoil it; that is why every candidate increases the L^2 distance. Read together with A300, the two complementary hypotheses are now both excluded: the factor is neither a function of the family index n alone (A300) nor of the monomial degree k alone (here). The honest narrowing is that the missing factor must couple n and k jointly, acting within a column to raise the $n = k$ entry above its neighbors. The selection rule's localization is an off-diagonal-suppression effect, not a row or column rescaling.

Status. OI-287-1 remains open. Its form is unchanged by this addendum: the missing factor lifts the μ entry by roughly 34% and the τ entry by roughly 7%, and is now known to be neither a function of the family index alone (A300) nor of the monomial degree alone (this addendum). The next probe must test a jointly indexed kernel that suppresses entries with $n \neq k$. No published number changes.

Verifier. `verify_P303.py`: recomputes the eigenstructure (eigenvalues ordered, node counts (0,1,2)), confirms the A299 base diagonal reproduces (1,1.12,7.50), confirms each declared candidate's diagonal and that none yields a valid diagonal peak, and confirms the NULL conditions, that no candidate lands within 20% and the closest candidate's μ entry moves away from 1.5. All PASS.