

Addendum 300: The Node-Count Bracket

A seventh excluded class, and a two-sided bound on the localizing factor

TOE Corpus

2026-06-13

Skill: gap-closure v1.1 — **Outcome:** Case D (mechanism class excluded; OI-287-1 lead bounded from both sides) — **Addresses:** OI-287-1 — **Verifier:** `verify_P300.py` — all PASS

Abstract

A299 left OI-287-1 in a sharp state: the density-Taylor overlap divided by λ_n reproduces the target diagonal (1, 1.5, 8) to (1, 1.12, 7.50), missing only by a factor that lifts the μ entry and sharpens the diagonal. A297 established that the three families are distinguished by their node count, 0, 1, 2 interior zeros. This addendum asks the obvious next question: is that node count the missing localizing factor? It is not, but the way it fails is informative. The node count brackets the target. The A299 base, with no node weighting, undershoots ($\mu = 1.12 < 1.5$, $\tau = 7.50 < 8$); a linear node gain ($1 + \nu_n$) overshoots ($\mu = 2.24$, $\tau = 22.50$). The target lies strictly between. But no single power $(1 + \nu_n)^p$ reconciles the two entries: hitting $\mu = 1.5$ requires $p \approx 0.42$, while hitting $\tau = 8$ requires $p \approx 0.06$. A uniform node-power gain is therefore excluded as a seventh mechanism class. The localizing factor lifts the μ and τ entries by different relative amounts, so it is not a function of the node count alone.

1 The Probe

Proposition 1.1 (The node count brackets but does not reconcile). *Let $B_{nk} = \langle \psi_n | c_k \pi^k x^k | \psi_n \rangle / \lambda_n$ be the A299 base, whose diagonal normalizes to (1, 1.12, 7.50). Let ν_n be the node count of ψ_n , equal to (0, 1, 2). Weighting row n by $(1 + \nu_n)$ gives the diagonal (1, 2.24, 22.50); by $(1 + \nu_n)^2$, (1, 4.48, 67.50). Neither peaks on the diagonal under the column-standardized criterion, so a seventh mechanism class is excluded. The target (1, 1.5, 8) is bracketed: each of its nontrivial entries lies between the unweighted value ($p = 0$) and the linear value ($p = 1$). But the single power that places μ at 1.5, $p = \log(1.5/1.12)/\log 2 \approx 0.42$, is not the single power that places τ at 8, $p = \log(8/7.50)/\log 3 \approx 0.06$. No uniform $(1 + \nu_n)^p$ satisfies both.*

Remark 1.2 (What the bracket says, and does not say). *The result is two-sided, which is more than the prior nulls gave. The selection rule's missing factor is now bounded: it must lift the μ row by more than nothing and less than a full factor of the node count, and the τ row by far less in relative terms. The disparity is the content. A factor that depended only on ν_n would lift both rows by the same power and could be tuned to either entry but not both; the data refuse that. So the missing factor is either jointly indexed by family and degree, not by family alone, or it is a function of a quantity that happens to coincide with the node count for the μ row but not the τ row. The honest narrowing: the localizing factor is not a power of the node count. It distinguishes the second and third families by more than their zero counts differ.*

Status. OI-287-1 remains open. Its form after this addendum: the missing factor in $\langle \psi_n | c_k \pi^k x^k | \psi_n \rangle / \lambda_n$ lifts the μ entry by roughly 34% and the τ entry by roughly 7%, lies between the identity and a linear node gain, and is not a uniform power of the node count. No published number changes.

Verifier. `verify_P300.py`: recomputes the node counts $(0, 1, 2)$, the A299 base and its linear and quadratic node gains, confirms none peaks on the diagonal, confirms the target is bracketed entry by entry, and confirms the μ and τ entries demand different node powers (0.42 versus 0.06). All PASS.