

Addendum 335: The Second-Preimage Scan on the Corpus APS Operator

A331's bounded-box global injectivity redone on the corpus boundary condition: it survives the soft zeta sector, and the moat is larger

TOE Corpus

2026-06-14

Skill: gap-closure v1.1 — **Outcome:** GROUNDED + TRANSFER (A331's in-box second-preimage scan redone on the corpus APS operator; no second preimage; the moat is *larger* under APS; in-box global injectivity survives the genuinely soft zeta sector) — **Addresses:** P18-T2 (iii) APS injectivity; A334; A331; A332; A330; A329; A302; A294 — **Verifier:** `verify_P335.py` — 16 PASS / 0 FAIL, exit 0

Abstract

P18-T2 sub-result (iii) asks whether $(\alpha, \gamma, \zeta, \beta)$ are the unique spectrum-matching assignment on P18's master operator. A331 answered the global half numerically: a bounded-box second-preimage scan found no coefficient assignment distinct from canonical with a machine-zero residual — in-box global injectivity, with a moat (the minimum off-canonical residual) of 1.67×10^{-3} . But A331's radial operator was the interior-mesh **Dirichlet model** the whole (iii) thread used, whose $l = 0$ spectrum is the squared Bessel zeros (14.68, 49.22, ...), not the corpus APS ladder (4, 16, 36, ...) that A329 established (P18 line 171). A334 built the corpus APS operator $R_{nl}(r) = r^{-1} J_{l+1}(kr)$, $k = 2n + l + 2$, and showed that on it the ζ sector is *genuinely soft and state-dependent*: the zeta diagonal $\langle R_{\text{sym}} \rangle = -2 + \frac{1}{2} |\psi(1)|^2 / \langle \psi | \psi \rangle$ ranges -1.98 to $+0.24$ across modes (std ≈ 0.6), not the near-uniform -2 of the Dirichlet model. **So A331's scan must be redone on the corpus operator**, where the live boundary term could in principle open a second preimage the Dirichlet model hid. That is this probe. We build the spectrum-matching residual on the corpus APS eigenstates — $\lambda_{nl} = (2n + l + 2)^2 + \langle V_{\text{self}} \rangle + \alpha \langle \rho \rangle + \zeta \langle R_{\text{sym}} \rangle$, with the additive sectors γ (layer-cycle $\text{Re } \omega^k$) and β (moment) re-optimised at each point, exactly as A331 — the sole corpus-faithful change being that $\langle R_{\text{sym}} \rangle$ is now the live, state-dependent APS quantity. **The finding** is option (a): in-box global injectivity *survives*. The zeta-valley floor (minimised over α) stays strictly positive everywhere $\zeta \neq \zeta_{\text{canon}}$ (shallowest 2.6×10^{-4} at $M = 6$); the residual-minimising α stays *pinned* at canonical along the whole valley; a dense α - ζ grid plus a 150-seed multi-start over the admissible box ($\alpha \in [0.6, 1.6]A_0$, $\zeta \in [0.2, 4]Z_0$) finds *no* second preimage, and the canonical assignment is the unique in-box machine-zero (4.4×10^{-16}). **The moat is larger under APS**, not collapsed: 6.2×10^{-3} at $M = 6$, $\approx 3.7 \times$ A331's Dirichlet moat. The mechanism is the opposite of a threat: the live state-dependent zeta diagonal A334 measured *breaks* the near-perfect α - ζ degeneracy that made the Dirichlet valley flat, so off-canonical residuals rise *faster* under the corpus BC, deepening the moat. Robust across $M \in \{4, 6, 8\}$ (moat $5.5 \times 10^{-3} \dots 7.2 \times 10^{-3}$) and box size. So A331's in-box global-injectivity conclusion is corpus-faithful: it *transfers* to the corpus APS operator, with a larger moat. This is a genuine numerical in-box result, not an analytic global proof — the ζ sector is genuinely soft (A334), and the analytic statement is A336's monotonicity probe; the corpus tie $\zeta = \alpha^{5/4}$ remains the clean closure for the corpus-posed family. Verifier `verify_P335.py`: 16 PASS / 0 FAIL. No published number changes.

1 The probe and the discipline

A334 settled the orientation of sub-result (iii): the whole (iii) thread (A302/A330/A331/A332/A333) built its radial operator on the interior mesh $r = \text{linspace}(dr, 1 - dr, N - 1)$, a Dirichlet wall at $r = 1$, and under the *corpus* APS boundary condition the ζ sector is genuinely soft and state-dependent rather than the rigid -2 shift A333 claimed. A334 therefore pre-registered two faithful fan-out probes to re-establish the (iii) injectivity results on the corpus operator, not transfer them from the Dirichlet model. This addendum is the first: A331’s bounded-box second-preimage scan, redone on the corpus APS operator.

The binding solution-bar discipline governs: do not force closure, do not manufacture a fan-out, honesty over closure. The honest outcomes here are sharply opposed — either the live boundary term opens a second preimage the Dirichlet model hid (the moat collapses, the free weights are non-injective on the corpus operator, the tie load-bearing), or in-box injectivity survives on the corpus operator (A331’s conclusion transfers). We recompute every number from scratch (numpy only; scipy unavailable) and report whichever the corpus operator gives. It gives option (a) — but not weakly: the moat is *larger* under APS, for a structural reason A334’s own measurement predicts.

2 The corpus APS residual, built from scratch

The corpus operator is the one A334 built and validated: P18’s bulk Dirac² radial form $-d_r^2 - (3/r)d_r + l(l+2)/r^2$ on the r^3 measure, regularity at $r = 0$, APS spectral projection at $r = 1$ (A329, P18 lines 393–396), with radial eigenfunctions $R_{nl}(r) = r^{-1}J_{l+1}((2n+l+2)r)$ and spectrum $(2n+l+2)^2$ (P18 line 171). The verifier reproduces the hard gate (check 3): the no-potential APS spectrum is the integer ladder ($l = 0$: 4, 16, 36, 64; $l = 1$: 9, 25, 49, 81), distinct from the Dirichlet squared-Bessel zeros A331 actually scanned (check 4). If the spectrum had come out squared-Bessel zeros we would be back on Dirichlet — the failure mode the whole reconciliation exists to escape.

Proposition 2.1 (The spectrum-matching residual on the APS operator). *The coefficient sectors act on the APS eigenstates through their first-order (Hellmann–Feynman) diagonal matrix elements — exactly the machinery A334 used to measure the ζ sector. The perturbed eigenvalue is*

$$\lambda_{nl}(\alpha, \zeta) = (2n + l + 2)^2 + \langle V_{\text{self}} \rangle_{nl} + \alpha \langle \rho \rangle_{nl} + \zeta \langle R_{\text{sym}} \rangle_{nl}, \quad \langle R_{\text{sym}} \rangle_{nl} = -2 + \frac{1}{2} \frac{|\psi_{nl}(1)|^2}{\langle \psi_{nl} | \psi_{nl} \rangle},$$

with $\langle \cdot \rangle_{nl}$ the expectation on the r^3 measure. The additive structural sectors γ (layer-cycle $\text{Re } \omega^k$) and β (moment) are re-optimised linearly at each (α, ζ) , and the spectrum-matching residual is $r(\alpha, \zeta) = \sqrt{\text{mean}(\lambda_{\text{low } M} - \lambda_{\text{target}})^2}$ over the lowest M eigenvalues of $l = 0, 1, 2$, the target the canonical assignment value. The sole corpus-faithful change from A331 is that $\langle R_{\text{sym}} \rangle_{nl}$ is the live, state-dependent APS quantity ($-1.98 \dots + 0.24$, std 0.61; check 5), not a near-uniform -2 . The APS eigenfunctions are genuine $L^2(r^3 dr)$ states (finite positive norms, check 6), so the residual landscape is well-defined.

3 The zeta-valley floor under APS

A331’s central probe was the ζ -valley: sweep ζ over a wide admissible range and minimise the residual over α (and the additive γ, β); if a second preimage existed on the soft α - ζ ray, the valley floor would touch zero somewhere off canonical. On the corpus operator it does not.

Proposition 3.1 (The APS valley floor is strictly positive, α^* pinned). *Minimising the residual over α at each ζ , the valley floor stays strictly positive everywhere $\zeta \neq \zeta_{\text{canon}}$: the shallowest off-canonical floor is 2.60×10^{-4} at $\zeta = 1.16 Z_0$ ($M = 6$). The residual-minimising α stays pinned at canonical ($1.000 A_0$) along the whole valley (drifting only to $0.975 A_0$ at the extreme $\zeta = 4 Z_0$). So even trading α against the soft ζ along the near-degenerate ray, the residual cannot be driven to zero off canonical: there is no second preimage on the soft ray, despite the live state-dependent zeta diagonal (checks 9–11).*

4 The bounded-box scan: the moat is larger under APS

Proposition 4.1 (No second preimage; the moat deepens under APS). *A dense α - ζ grid ($\alpha \in [0.6, 1.6] A_0$, $\zeta \in [0.2, 4] Z_0$) plus a 150-seed multi-start over the admissible box finds no second preimage. The moat — the minimum residual at any point outside a 0.10 relative ball around canonical — is 6.20×10^{-3} at $M = 6$, bounded below by $\varepsilon > 10^{-4}$; the canonical assignment is the unique in-box machine-zero (4.36×10^{-16}); the multi-start finds zero second-preimage hits (best off-canonical residual 8.40×10^{-4}) (checks 12–14). The APS moat is $\approx 3.7 \times$ A331’s Dirichlet moat 1.67×10^{-3} — larger, not collapsed.*

The reason is structural and is exactly A334’s measurement read forward. A331’s valley was so flat because, on the Dirichlet model, the ζ diagonal was a near-uniform -2 (cosine with the uniform vector 1.00000 , A332): a uniform shift is nearly degenerate with the α direction along the α - ζ ray, which is why A330 measured a 99.86% collinearity and A331 a shallow valley. A334 showed that under the corpus BC the ζ diagonal is *mode-structured* (cosine with the uniform vector 0.864 , range $-1.98 \dots + 0.24$): it carries genuine shape, so it is no longer near-degenerate with α . The same state-dependence that makes the ζ sector “soft” as an analytic object (A332’s lone non-sign-definite sector, A334’s live boundary term) makes the *numerical* second-preimage landscape *stiffer*, because off-canonical the structured zeta shift no longer cancels against an α re-optimisation the way a uniform shift does. The soft sector, far from opening the second preimage the Dirichlet model might have hidden, deepens the moat.

5 Robustness and the verdict

Remark 5.1 (A331’s conclusion transfers, with a larger moat). *Across $M \in \{4, 6, 8\}$ the moat stays bounded below and larger than A331’s at every M (7.18, 6.20, 5.51×10^{-3}) and the valley floor stays strictly positive (2.92, 2.60, 2.50×10^{-4} , check 15); extending ζ to $6 Z_0$ still leaves the residual finite and far from zero (8.12×10^{-3} , check 16). So A331’s in-box global-injectivity conclusion is corpus-faithful: it transfers to the corpus APS operator, with a larger moat. This is a genuine numerical in-box injectivity result for the free-weight version on the right operator; it is not an analytic global proof. The ζ sector is genuinely soft (A334), and the analytic statement is A336’s monotonicity probe; the corpus tie $\zeta = \alpha^{5/4}$ (P18 line 348) remains the clean closure for the corpus-posed family.*

6 Gate ledger

Object	Gate	Finding
Corpus APS operator (hard gate)	PASSED	$(2n + l + 2)^2$ spectrum, distinct from Dirichlet,
Zeta diagonal $\langle R_{\text{sym}} \rangle$	LIVE	state-dependent $-1.98 \dots + 0.24$, std 0.61 (A334)
Canonical assignment	MACHINE-ZERO	residual 4.4×10^{-16}
Zeta-valley floor (APS)	STRICTLY POSITIVE	shallowest 2.6×10^{-4} ; α^* pinned at canonical
Second preimage in the box	NONE	multi-start zero hits; canonical the unique zero
APS moat vs A331 Dirichlet moat	LARGER ($3.7\times$)	6.2×10^{-3} vs 1.67×10^{-3} ; soft zeta deepens it
In-box global injectivity	SURVIVES (corpus BC)	A331's conclusion transfers, with a larger moat
Analytic global proof	OPEN	zeta soft; A336 monotonicity; tie clean

7 What it does to (iii), and the sibling probe

This probe re-establishes, on the corpus operator, exactly the in-box numerical result A331 established on the Dirichlet model — and strengthens it: the moat is larger, for the structural reason A334's measurement predicts. The cross-thread inconsistency A334 flagged (Heart 2 (i) on APS, (iii) on Dirichlet) is now closed for the second-preimage half of (iii): the bounded-box global injectivity is a corpus-faithful statement, not a Dirichlet artifact. What stays open is the *analytic* global proof, and the obstruction is precisely the soft ζ sector — the same obstruction A332 isolated and A334 showed is real under the corpus BC. The sibling probe A336 (running in parallel) redoes A332's Hellmann–Feynman monotonicity on the corpus APS operator: the three shape sectors ($\alpha\rho$, βM , $\gamma \text{Re } \omega^k$) are multiplication/projection/character operators whose sign-definiteness should be BC-independent but must be re-verified on the APS eigenstates, and the zeta diagonal, now state-dependent rather than near-uniform, must be re-characterised to determine whether per-axis monotonicity in ζ survives at all. If A336 re-establishes the three shape sectors, (iii) stands as a corpus-faithful in-box numerical injectivity result (this probe) plus three-sector structural monotonicity, with the ζ sector an honest soft direction — genuinely parallel to where A331/A332 left the Dirichlet model, but now on the right operator.

8 Conclusion

A335 redoes A331's bounded-box second-preimage scan on the corpus APS operator A334 built — $R_{nl}(r) = r^{-1}J_{l+1}((2n + l + 2)r)$, spectrum $(2n + l + 2)^2$ (P18 line 171), eigenstates that do not vanish at the wall — where the zeta diagonal $\langle R_{\text{sym}} \rangle = -2 + \frac{1}{2}|\psi(1)|^2/\langle \psi|\psi \rangle$ is the live, state-dependent APS quantity, not the near-uniform -2 of the Dirichlet model the (iii) thread used. The spectrum-matching residual is built on the APS eigenstates with the additive γ, β sectors re-optimised, exactly as A331. The zeta-valley floor (minimised over α) stays strictly positive everywhere $\zeta \neq \zeta_{\text{canon}}$ (shallowest 2.6×10^{-4} at $M = 6$), the residual-minimising α stays pinned at canonical, and a dense α - ζ grid plus a 150-seed multi-start over the admissible box finds no second preimage: canonical is the unique in-box machine-zero. The moat is larger under APS — 6.2×10^{-3} at $M = 6$, $\approx 3.7\times$ A331's Dirichlet moat — not collapsed, because the live state-dependent zeta diagonal breaks the α - ζ degeneracy that flattened the Dirichlet valley, so off-canonical residuals rise faster under the corpus BC. The result is robust across $M \in \{4, 6, 8\}$ and box size. So in-box global injectivity survives on the corpus APS operator: A331's conclusion transfers, with a larger moat — a genuine corpus-faithful numerical in-box injectivity result for the free-weight version. It is not an analytic global proof; the ζ sector is genuinely soft (A334), the analytic statement is A336's monotonicity probe, and the corpus tie $\zeta = \alpha^{5/4}$ remains the clean closure for the corpus-posed

family. A294's re-typing of the whole P18-T2 stands; sub-results (i) (A328/A329) and (ii) (A310) are unchanged. No canon source files are edited; no published number changes; the $\hat{O}$ assembly and its derived observables are untouched. Verifier `verify_P335.py` reaches 16 PASS / 0 FAIL and exits zero.