

Addendum 339: The Gale–Nikaidô Route to ζ -Sector Global Injectivity

On the corpus APS operator the offset–difference functionals make the coefficient Jacobian a P-matrix; the lone non-floor gap sharpens to a single principal-minor inequality on the box

TOE Corpus

2026-06-14

Skill: gap-closure v1.1 — **Outcome:** PARTIAL (Gale–Nikaidô reduces free-weight ζ -sector global injectivity to a pointwise P-matrix condition; the offset–difference functionals achieve the P-matrix at canonical, min principal minor $+0.0305$; in the affine first-order model this is a P-matrix on the whole box, giving Gale–Nikaidô injectivity, but that model is rank-injective trivially, so the load-bearing exact claim reduces to one named principal-minor inequality on the box — true at canonical, numerical in-box, open in closed form) — **Addresses:** P18-T2 (iii) ζ -sector global injectivity; A336; A335; A334; A333; A332; A331; A330; A302; A294 — **Verifier:** `verify_P339.py` — 17 PASS / 0 FAIL, exit 0

Abstract

After A334 reconciled the whole sub-result (iii) thread to the corpus APS operator, two fan-outs settled both halves *on* that operator: A335 (numerical) found in-box global injectivity survives with a larger moat (6.20×10^{-3} vs A331’s Dirichlet 1.67×10^{-3}), and A336 (analytic, per-axis) found the three shape sectors $\alpha\rho, \beta M, \gamma \operatorname{Re} \omega^k$ stay sign-definite hence per-axis monotone/injective, but the renormalisation sector ζR , $R = r\partial_r$, is *non-monotone* — its Hellmann–Feynman diagonal $\langle \psi_n | R_{\text{sym}} | \psi_n \rangle$ ranges -1.98 to $+0.24$ and *sign-flips*. That sign-flip is the single obstruction to a clean analytic free-weight global-injectivity proof, and it is the lone live *non-floor* gap across both open hearts. This addendum attacks it via the **Gale–Nikaidô global univalence theorem**: a C^1 map on a closed box whose Jacobian is a *P-matrix* (all principal minors > 0) at every point is globally injective on the box. We reduce Φ to a square map $\mathbb{R}^4 \rightarrow \mathbb{R}^4$ by choosing four spectral functionals, whose 4×4 Jacobian has columns the sector Hellmann–Feynman diagonals on the corpus APS eigenstates. **The structural lever** is the exact integration-by-parts split (A333/A334): the ζ column is $-2\mathbf{1} + \mathbf{b}$ with $b_n = \frac{1}{2}|\psi_n(1)|^2 / \langle \psi_n | \psi_n \rangle \geq 0$; the uniform -2 bulk is the sign-flip culprit, but in eigenvalue *difference* functionals it cancels and ζ enters only through the *nonnegative* mode-structured $\mathbf{b}$ — the degeneracy-break A335 measured. **The witness**: with the offset plus three shape-differences (reference mode $(1, 2)$, partners $(0, 0), (2, 2), (1, 1)$) the 4×4 Jacobian *is* a P-matrix — all fifteen principal minors positive, min principal minor $+0.0305$ — where no raw eigenvalue-functional pick is one (the sign-indefinite ζ column blocks it; best raw min minor is negative). In the first-order (Hellmann–Feynman, frozen-eigenstate) model that the whole (iii) thread uses (A331/A335/A336) the coefficient→spectrum map is exactly *affine*, so the Jacobian is constant over the box and the P-matrix holds everywhere; all three Gale–Nikaidô hypotheses are met and the first-order model is globally injective on the box. **The honest catch** (why PARTIAL, not HIT): the affine map is injective by rank alone (the response has rank 4, smallest singular value 0.128, exactly A302), so there Gale–Nikaidô is not load-bearing. The non-trivial target is the *exact* nonlinear map, whose Jacobian is c -dependent through the shifting APS eigenstates and

whose ζ entries carry $b_n(c) = \frac{1}{2}|\psi_n(1;c)|^2/\langle\psi_n(c)|\psi_n(c)\rangle$, a Bessel value with no closed form. The whole proof then reduces to **one explicit scalar inequality** — the binding principal minor $M_{\min}(c)$ stays > 0 across the box — true at canonical (+0.0305), true numerically in-box (A335’s moat > 0 certifies no second preimage), open in closed form. So the lone non-floor gap *sharpens* from “a full analytic proof at the ζ sector” to this single principal-minor inequality. A pre-registered fan-out (Bessel-value bounds on $b_n(c)$) would upgrade PARTIAL→HIT. The corpus tie $\zeta = \alpha^{5/4}$ remains the clean closure for the corpus-posed family. Verifier `verify_P339.py`: 17 PASS / 0 FAIL. No published number changes.

1 The probe and the discipline

This addendum addresses the one live non-floor gap in P18-T2. The chain to here: A302 proved local rank-4 coefficient rigidity; A330/A331 established in-box global injectivity numerically and A332 the per-axis structural monotonicity, both on a Dirichlet model; A333 wrongly upgraded ζ to an exact rigid shift on that model; A334 reconciled the thread to the corpus APS operator and showed ζ is genuinely soft and sign-flipping there; A335 (numerical) and A336 (analytic per-axis) redid the two injectivity results on the corpus operator. What A335/A336 leave is a single analytic obstruction: the ζ sector’s sign-flipping Hellmann–Feynman diagonal blocks a per-axis monotonicity proof of *full* free-weight global injectivity, which A335 has only numerically and in-box. The discipline is the binding solution bar: do not force closure; a precise PARTIAL (one named inequality) or an honest floor is the acceptable and likely result, not a manufactured HIT.

The route is the Gale–Nikaidô global univalence theorem, the standard converter of global injectivity into a pointwise algebraic property.

Theorem 1.1 (Gale–Nikaidô, 1965). *Let $F : \overline{B} \rightarrow \mathbb{R}^n$ be C^1 on a closed rectangular region $\overline{B} \subset \mathbb{R}^n$, and suppose the Jacobian $JF(x)$ is a P-matrix (every principal minor strictly positive) at every $x \in \overline{B}$. Then F is injective on $\overline{B}$.*

2 The reduction and the corpus operator (the hard gate)

The operator is P18’s bulk Dirac² radial form on the unit ball, $-d_r^2 - (3/r)d_r + l(l+2)/r^2$ on the r^3 measure, with regularity at $r = 0$ and APS spectral projection at $r = 1$ (A329; P18 lines 393–396). Its APS radial eigenfunctions are $R_{nl}(r) = r^{-1}J_{l+1}(kr)$, $k = 2n + l + 2$, with spectrum $E_{nl} = (2n + l + 2)^2$ (P18 line 171) and boundary values $R_{nl}(1) = J_{l+1}(2n + l + 2) \neq 0$. The verifier’s hard gate (check 3) confirms the no-potential spectrum is the integer ladder $l = 0 : (4, 16, 36, 64)$, distinct from the Dirichlet squared-Bessel zeros (14.68, 49.22, ...); the eigenstates do not vanish at the wall (check 4); and the ζ diagonal recovers A334/A336 (range -1.98 to $+0.24$, sign-flipping at $l = 2, n = 0$, check 5).

To square the map we choose four spectral functionals $F_0, \dots, F_3$ of the spectrum, giving $\Phi = (F_0, \dots, F_3) : \mathbb{R}^4 \rightarrow \mathbb{R}^4$ in the coefficients $c = (\alpha, \gamma, \zeta, \beta)$. The 4×4 Jacobian has entries $\partial F_i / \partial c_j$, built from the sector Hellmann–Feynman diagonals $\langle \psi_m | S_j | \psi_m \rangle$ in the corpus r^3 inner product. The free degree of freedom Gale–Nikaidô grants is that it suffices for *some* choice of modes and functionals to give a P-matrix everywhere; we exploit this to control the sign-indefinite ζ column.

3 The structural lever: the exact IBP split

Proposition 3.1 (The ζ column splits into a uniform part and a nonnegative boundary part). *On the corpus APS eigenstates the renormalisation diagonal obeys the exact integration-by-parts*

identity (A333/A334)

$$\langle \psi_n | R_{\text{sym}} | \psi_n \rangle = -2 + b_n, \quad b_n = \frac{1}{2} \frac{|\psi_n(1)|^2}{\langle \psi_n | \psi_n \rangle} \geq 0 \quad \text{for every mode.}$$

The uniform bulk -2 is the source of sign-indefiniteness (it drives the small- b_n modes negative while the large- b_n mode $l=2, n=0$ goes positive). The boundary part $\mathbf{b} = (b_n)$ is nonnegative and mode-structured.

The lever is that in eigenvalue-difference functionals the uniform -2 cancels: a difference functional $F_i = \lambda_{m_i} - \lambda_{\text{ref}}$ reads, in its ζ entry, $b_{m_i} - b_{\text{ref}}$ (the -2 drops, check 7), so ζ enters the difference rows only through the nonnegative $\mathbf{b}$ — exactly the mode-structure that broke the α - ζ degeneracy in A335. The sign-flip that defeats a monotonicity proof in eigenvalue coordinates is therefore not an obstruction in difference coordinates, where the offending uniform term is gone.

4 The witness: the Jacobian is a P-matrix

Proposition 4.1 (A P-matrix from the offset-difference functionals). *With the functionals $F_0 = \lambda_{(1,2)}$ (offset) and $F_i = \lambda_{m_i} - \lambda_{(1,2)}$ for $m_i \in \{(0,0), (2,2), (1,1)\}$ (three shape-differences), the 4×4 Jacobian*

$$JF = \begin{pmatrix} 173.99 & -0.50 & -0.991 & 0.600 \\ 95.14 & 1.50 & 0.303 & 0.155 \\ 0.548 & 0 & 0.155 & 0.025 \\ -6.267 & 0 & -0.991 & 0.040 \end{pmatrix}$$

is a P-matrix: all 15 principal minors are strictly positive, with binding (smallest) principal minor $+0.0305$ (the zeta-carrying $\{F_2, F_3\}$ block; checks 8–9).

The difference trick is necessary, not incidental: the raw eigenvalue-functional Jacobian (functionals = four eigenvalues, sectors in natural order) is *never* a P-matrix for any ordered choice of four modes — the best achievable minimum principal minor is negative (check 10) — because the sign-indefinite ζ column forces a negative 2×2 minor against the always-positive α column (33 of the 66 α - ζ 2×2 minors are negative). Only the difference functionals, which excise the uniform -2 , rescue the P-matrix.

5 The first-order model: P-matrix everywhere, and the honest catch

The whole sub-result (iii) thread (A331/A335/A336) works in the first-order (Hellmann–Feynman, frozen-eigenstate) model, in which the eigenvalue is

$$\lambda_m(c) = \lambda_m^0 + \sum_i c_i \langle \psi_m | S_i | \psi_m \rangle,$$

the diagonals evaluated on the *unperturbed* eigenstates and hence c -independent. So λ is *affine* in c and the functional Jacobian JF is *constant* over the box (check 11): the P-matrix at the canonical point is a P-matrix at *every* point. The three Gale–Nikaidô hypotheses are then explicitly met (check 12): (a) the admissible region is the rectangular box $\alpha \in [0.6, 1.6]A_0$, $\zeta \in [0.2, 4]Z_0$ (A335’s box) with γ, β ; (b) the affine map is C^∞ ; (c) the constant JF is a P-matrix. Gale–Nikaidô then gives global injectivity of the first-order model on the box (check 13).

Remark 5.1 (Why this is PARTIAL, not HIT). *The first-order affine map is injective by rank alone: the 12×4 response matrix has rank 4 with smallest singular value 0.128 (exactly A302’s local rank-4, check 14). On the first-order model Gale–Nikaidô is therefore not load-bearing — it recertifies a fact linear algebra already gave. The non-trivial target is the exact nonlinear eigenvalue map, whose Jacobian is c -dependent because the eigenstates shift with c , and whose ζ entries carry the c -dependent APS boundary value $b_n(c) = \frac{1}{2}|\psi_n(1;c)|^2 / \langle \psi_n(c) | \psi_n(c) \rangle$, a Bessel value with no closed form (check 15). The constant-Jacobian argument does not extend to it.*

6 The single named inequality (the sharpened gap)

Proposition 6.1 (The gap reduces to one principal-minor inequality). *For the exact map, with the chosen offset–difference functionals, the Gale–Nikaidô P-matrix property on the box reduces to the single scalar inequality*

$$M_{\min}(c) > 0 \quad \text{for all } c \text{ in the box,}$$

where $M_{\min}(c)$ is the binding principal minor (the zeta-carrying difference block) of the c -dependent functional Jacobian. This inequality is true at the canonical point ($M_{\min} = +0.0305$), true numerically across the box (A335’s moat $6.20 \times 10^{-3} > 0$ certifies no second preimage in the box), and open in closed form because $b_n(c)$ is a Bessel-valued boundary quantity with no closed form (check 16).

This is sharper than the prior status “free-weight injectivity is numerical and in-box” (A335): the entire open content is now a single, explicit, scalar inequality on the box, parametrised by the APS boundary values. The verdict is a genuine PARTIAL — Gale–Nikaidô has done real work (it converted the global question to one pointwise inequality and the difference functionals made that inequality true at the canonical point and structurally tractable), and the residual is named precisely (check 17).

7 Gate ledger

Object	Gate	Finding
Corpus APS operator (hard gate)	PASSED	$(2n + l + 2)^2$, $\psi(1) \neq 0$, not squared-Bessel
ζ column split $-21 + \mathbf{b}$	DERIVED	$b_n = \frac{1}{2} \psi(1) ^2 / \ \psi\ ^2 \geq 0$ (A333/A334)
difference functionals kill the -2	DERIVED	ζ enters difference rows via $\mathbf{b} \geq 0$
4×4 Jacobian P-matrix (canonical)	ACHIEVED	all 15 minors > 0 , min $+0.0305$
raw eigenvalue-functional Jacobian	FAILS	best ordered-pick min minor < 0
first-order box P-matrix	PROVEN	affine $\Rightarrow$ constant $JF \Rightarrow$ GN injective
first-order Gale–Nikaidô	RANK-TRIVIAL	affine map injective by rank ($\sigma_{\min} = 0.128$)
exact-map Jacobian	c -DEPENDENT	ζ entries carry $b_n(c)$, Bessel, no closed form
the single named inequality	PARTIAL	$M_{\min}(c) > 0$ on the box; canonical/numerical yes, clo
(iii) free-weight global injectivity	SHARPENED	gap = one principal-minor inequality

8 Pre-registered fan-out

The PARTIAL has a real attack, so one probe is pre-registered (no manufacture beyond it): *bound $M_{\min}(c)$ analytically by Bessel-value bounds on $b_n(c)$* . The boundary value $|\psi_n(1;c)|^2$ and its c -

derivative are governed by J_{l+1} evaluated near the APS ladder $k = 2n + l + 2$, where J_{l+1} is $O(1)$ and bounded away from its zeros; standard two-sided Bessel envelopes (and the monotone dependence of the APS quantisation on c) would give a uniform lower bound on the difference-block minor over the box, upgrading PARTIAL→HIT. If that bound cannot be made uniform — if the boundary values approach a J_{l+1} zero somewhere in the admissible box — the result settles to a FLOOR parallel to G1/APS, the boundary value being an irreducible Bessel quantity. No second fan-out is opened: the alternative offset/shape decompositions reduce to the same boundary-value inequality.

9 Conclusion

The Gale–Nikaidô global univalence theorem converts free-weight ζ -sector global injectivity into a pointwise P-matrix condition on the box. The exact integration-by-parts split of the ζ column into a uniform -2 and a nonnegative boundary part b_n makes the offset–difference functionals a P-matrix at the canonical point (all 15 principal minors positive, min $+0.0305$), where no raw eigenvalue-functional choice is one. In the affine first-order model the whole (iii) thread uses, this is a P-matrix on the entire box, so Gale–Nikaidô gives global injectivity of that model — but the model is rank-injective trivially, so the load-bearing claim is the exact nonlinear map, where the Jacobian is c -dependent through the Bessel-valued APS boundary terms. The whole proof reduces to one explicit scalar inequality, the binding principal minor $M_{\min}(c) > 0$ across the box: true at canonical, true numerically in-box (A335), open in closed form. The lone non-floor gap in P18-T2 (iii) therefore sharpens from “a full analytic free-weight proof at the ζ sector” to this single principal-minor inequality, with a pre-registered Bessel-bound attack. The corpus tie $\zeta = \alpha^{5/4}$ (P18 line 348) 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; (iii)’s per-axis monotonicity in (α, γ, β) (A336) and in-box injectivity (A335) are unchanged. No canon source files are edited; no published number changes; the $\hat{O}$ assembly and its derived observables are untouched. Verifier `verify_P339.py` reaches 17 PASS / 0 FAIL and exits zero.