

Addendum 305: The Joint Kernel

A ninth excluded class, and why off-diagonal suppression cannot lift a diagonal entry

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_P305.py` — all PASS

Abstract

A300 and A303 together excluded the two separable hypotheses for 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). A300 ruled out a factor depending only on the family index n ; A303 ruled out a factor depending only on the monomial degree k . Both pointed to the same next hypothesis: the missing factor must couple n and k jointly, acting within a column to raise the $n = k$ entry above its neighbors, an off-diagonal-suppression effect. This addendum runs that test. Four non-separable joint kernels K_{nk} are declared in advance and applied as $M_{nk} = B_{nk} K_{nk}$: a Gaussian off-diagonal suppression $\exp(-(k - (n+1))^2/2)$ with fixed width, a Kronecker band $(1/2)^{|k-(n+1)|}$ with fixed decay, the overlap's own structure as a self-suppressing kernel, and a hard projection onto the diagonal. A declared separable control $f(n)g(k)$ is run alongside. None of the joint kernels yields a valid diagonal peak: the Gaussian moves only the third column's argmax one row toward the diagonal and stalls; the hard projection collapses every column's variance and so peaks trivially on the matched sine-box control as well, failing the control clause. Every kernel that equals one on the matched pair leaves the normalized diagonal exactly at (1, 1.12, 7.50), because off-diagonal suppression by construction does not touch the diagonal entry it sits beside. The separable control lands at the same L^2 as the joint kernels and does not beat them. A non-separable joint kernel acting on the existing overlap matrix is therefore excluded as a ninth mechanism class. The localizing factor is not a kernel that reweights the overlaps the operator already produces; it must change the overlaps themselves.

1 The Probe

Proposition 1.1 (A joint kernel on the overlap cannot localize the diagonal). *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), in which every column's largest entry sits in row 0. For a joint kernel K_{nk} with $K_{nn} = 1$, the weighted matrix $M_{nk} = B_{nk} K_{nk}$ has $M_{nn} = B_{nn}$, so its normalized diagonal is identical to that of B . Evaluated on the declared kernels, each of the Gaussian $\exp(-(k - (n+1))^2/2)$, the band $(1/2)^{|k-(n+1)|}$, and the self-structure overlap returns the diagonal (1, 1.12, 7.50) unchanged, and none produces a valid diagonal peak under the column-standardized criterion paired with the sine-box control. The hard projection $K = I$ peaks on the diagonal but also peaks on the control, so it fails the control clause. The closest candidate by L^2 sits at 0.63, the A299 control distance, and its μ entry remains at 1.12, not moved toward 1.5. The pre-registered HIT criterion, a valid diagonal peak together with all three entries within 20%, is met by no candidate. A non-separable joint kernel is excluded as a ninth mechanism class. The separable control $f(n)g(k)$ lands at the same L^2 and does not beat the joint kernels.*

Remark 1.2 (Why the kernel fails, and what it implies). *The failure has two parts, and both are structural rather than numerical. The first part concerns the diagonal entries. A kernel that suppresses off-diagonal entries by definition equals one on the matched pair, so it multiplies B_{nn} by one and leaves it fixed; the diagonal ratio is invariant under any such kernel. The lift the target demands, μ from 1.12 toward 1.5, raises a diagonal entry, and no off-diagonal suppression can raise a diagonal entry. The second part concerns the peak. Suppressing the off-diagonal does move a column's argmax toward the diagonal, and the Gaussian shows this in the third column, where the argmax shifts from row 0 to row 1; but the dominance of the first family in the raw overlaps is large enough that finite suppression does not carry every column all the way to its diagonal row, and the only kernel that forces it, the hard projection, destroys the column variance that distinguishes the operator's eigenfunctions from a generic basis and so peaks on the sine-box control too. Read together with A300 and A303, the three exclusions close the space of post-hoc reweightings: the factor is not a function of the family alone, not a function of the degree alone, and not a kernel on the overlaps that depends on both. The honest narrowing is that the localizing information is not present in the overlap matrix B to be recovered by any reweighting of it. It must enter earlier, in the overlaps themselves, which means in the operator or the states, not in a correction applied after the overlaps are formed.*

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 (A303), nor a kernel applied to the A299 overlap matrix (this addendum). The next probe must look inside the overlap, at the operator or the eigenstates, not at a post-overlap correction. No published number changes.

Verifier. `verify_P305.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 joint kernel that equals one on the matched pair leaves the diagonal at the A299 value and that none yields a valid diagonal peak, and confirms the NULL conditions, that no candidate lands within 20%, the closest sits at the control $L^2 = 0.63$ with μ unmoved, and the separable control does not beat the joint kernels. All PASS.