Full Projected Closure Operator and the Corrected Absolute Scale of the K = 7 Commitment-Threshold Spectrum

Extension of the minimal-model parameters with kernel integrals and higher-neighbour circulant couplings included, cross-kernel treatment, self-consistency narrowing of the kernel parameters, and the resulting sharpening of the α-saturation condition

Keith Taylor · versf-eos.com

Plain-language summary

Where this paper sits in the programme

The prior paper in this sequence — *Microscopic closure dynamics and the K = 7 spectrum in the VERSF framework* — took the two-parameter (a, b) minimal model of the Pair-Resolved Closure Spectrum paper and gave it a microscopic interpretation through a projection ansatz and an explicit closure kernel. It identified a specific numerical compensation: across a range of mode widths, one effective parameter falls while the other rises, and the specific combination that enters the physical prediction is stable to within ±6%. On the strength of that stability, the prior paper sharpened the earlier α-saturation ceiling into a centred prediction.

What this paper does is complete that calculation. Specifically, it computes every contribution to the projected closure operator — not only the gradient-overlap terms the prior paper used, but also the closure-kernel cross-terms the prior paper introduced in principle but did not evaluate, and the higher-neighbour circulant couplings that the minimal-model restriction set to zero.

The sequence of papers now forms a closed chain. The Pair-Resolved Closure Spectrum paper establishes that the commitment-threshold correction is determined by the closure spectrum and reduces it to a minimal set of effective parameters. The Microscopic Closure Dynamics paper identifies those parameters with a projection of the κ-field and establishes the stability of the resulting spectrum. The present paper completes that programme by evaluating the full projected operator implied by the projection, confirming the stability at the operator level, fixing the absolute scale, and reducing the remaining freedom to a small set of microscopic parameters. What remains open is one pure dimensionless ratio in the closure-kernel sector, the bath cutoff Λ in the commitment-event sector, and the transverse scale $\xi_\perp$ in the κ-field sector — three specific targets rather than a diffuse class of unknowns.

What the calculation finds

Four things.

First, the stability result survives — and is tighter than the prior paper claimed. Reading the physically relevant spectral quantity (the maximum closure frequency Ω_{max}) directly off the computed spectrum, it varies by only ±1.5% across the physical mode-width range, better than the ±3% band the prior paper estimated. The structural intuition of the prior paper — that the closure geometry protects the spectral prediction against details of the mode shape — is correct.

Second, the minimal-model (a, b) framework is not a literal description of the projected physics but a reduced description of it. Two specific things are now visible: the overlap matrix is far from the identity (adjacent modes overlap at 70–90%, not the ~10% the prior paper assumed), and the next-nearest-neighbour coupling is not small compared to the nearest-neighbour coupling (≈ 21%, not the ≤ 4% the prior paper estimated). The (a, b) parameters that emerge from a least-squares fit to the true spectrum are effective coordinates, not canonical physical constants. The physical observable is Ω_{max} itself.

Third, the absolute scale of Ω_{max} is roughly twice what the prior paper reports. This comes from the closure-kernel cross-term contribution the prior paper dropped from its hopping identification and from the under-counted next-nearest coupling. Translating: the saturating bath width that α-sector consistency would require is $\Delta_\alpha \approx 20 \cdot \xi^{-1}$ rather than $10 \cdot \xi^{-1}$.

Fourth, the kernel-parameter ambiguity — a new issue the prior paper's formulation masked — is narrower than a naïve scan suggests. A free two-parameter scan of the closure kernel's well depth and width gives a ~40% range in Ω_{max}. But once dimensional scaling, self-consistency between mode width and well curvature, and variational selection are applied, the residual ambiguity is ~20%. Δ_α stays within $[19, 24] \cdot \xi^{-1}$ across the narrowed space. What remains is a single pure number — the ratio of kernel width to closure scale — which joins the bath cutoff as the priority microscopic open problem.

What this does and doesn't change

The α-saturation ceiling — $\delta C \leq 2.1 \times 10^{-3}$ at the α-sector limit of $r_{max} = 0.15$ — is completely unaffected. That number is algebraic within the prior framework and does not depend on any of the details this paper revises.

What changes is the hierarchy the microscopic physics has to deliver in order for the ceiling to be saturated. Where the prior paper estimates that a bath ten times broader than the closure scale is sufficient, this calculation says it needs to be about twenty times broader. This makes the framework's prediction more informative rather than less: a δC measurement at the ceiling would confirm a steep and specific hierarchy, not a loose one. Measurements below the ceiling would then be genuinely diagnostic of sub-saturation, not of the framework's sensitivity to calibration.

In short

The prior paper's structural intuition — that the projection ansatz produces a spectral quantity stable against details of the mode shape — is confirmed. The underlying stability of Ω_{max} is real. The minimal-model (a, b) parameters are not literal physical couplings but effective coordinates arising from projection of the full operator; their role as a reduced description of the spectrum is preserved, and the physical observable is Ω_{max}. Evaluating the full operator (rather than its reduced projection, which the prior paper used) shifts the absolute normalization upward by a factor of roughly two — preserving the structure, including previously omitted operator contributions. The kernel-parameter open problem is narrowed but not closed; one pure number remains to be derived, joining the bath cutoff Λ already flagged by [1] and the transverse scale $\xi_\perp$ identified here by the 3D extension.

What this paper proves

1. **The K = 7 projected closure spectrum is spectrally stable.** $\Omega_{max}\cdot\xi \in [2.94, 3.24]$ across $\sigma/d \in [0.6, 1.0]$ and across Gaussian, sech, and Lorentzian mode families — ±5% cross-family, ±1.5% within Gaussian.
2. **Ω_{max} is the physical observable.** The minimal-model (a, b) parameters are effective coordinates on a two-parameter least-squares projection of a three-coupling, non-orthonormal circulant. Their role as a reduced description of the spectrum is retained; the physical prediction is Ω_{max}, read directly from the n = 3 eigenvalue.
3. **$\Omega_{max}\cdot\xi \approx 3.0$; $\Delta_\alpha \approx 20\cdot\xi^{-1}$.** This corresponds to an increase by approximately a factor of two relative to the reduced projection evaluated in [1].
4. **The kernel parameters are uniquely fixed by the projection framework up to a single dimensionless ratio ℓ/ξ.** The ~40% (V_0, ℓ) ambiguity of an unconstrained scan reduces to ~10% residual across the two natural values of ℓ/ξ. $\Delta_\alpha\cdot\xi \in [19.6, 24.0]$ within that one-parameter family.

The α-saturation ceiling $\delta C \leq 2.1 \times 10^{-3}$ is unchanged. The saturation condition is steeper, which makes the prediction sharper rather than weaker: a δC measurement at the ceiling would confirm a specific microscopic hierarchy rather than a loose constraint.

Abstract

The prior paper in this sequence (Microscopic closure dynamics, Taylor [1]) introduced a projection ansatz identifying the pair-amplitude field with the spatial projection of the κ-field onto Wannier modes, and reported a geometric stability result — that the combination (a + 3.8019 b) controlling Ω_{max}^2 varies by only ±6% across the admissible mode-width range. On the strength of that result, it sharpened the α-saturation ceiling $\delta C \lesssim 2.1 \times 10^{-3}$ of the Pair-

Resolved Closure Spectrum paper [2] into a centred prediction $\delta C = (2.1 \pm 0.3) \times 10^{-3}$ at a saturating bath width $\Delta_\alpha = (10.0 \pm 0.3)\cdot\xi^{-1}$ within the Gaussian benchmark.

The present paper extends that computation by evaluating the full projected closure operator $L = G + V + K$, including the closure-kernel cross-terms $K_{j,k}$ that [1] introduced in principle but did not compute, and without restricting to the nearest-neighbour circulant structure of the minimal-model (M1) assumption. It also narrows the kernel-parameter (V_0, ℓ) space via dimensional scaling and self-consistency.

Results:

(i) Stability confirmed and tightened. The physically relevant spectral quantity $\Omega_{max}\cdot\xi$ is directly computable from the full spectrum without recourse to a two-parameter fit, and is stable at ±1.5% across $\sigma/d \in [0.6, 1.0]$ for Gaussian modes — tighter than [1]'s ±3% estimate of the same quantity through the $(a + 3.80\, b)$ route.

(ii) Absolute scale shifted by the inclusion of full operator contributions. The prior paper evaluated a reduced projection: its §6.2 hopping identification uses $G_{j,j+1}$ alone and discards the closure-kernel cross-term $K_{j,j+1}$, which the present calculation shows accounts for ≈ 44% of the true nearest-neighbour entry $L_{j,j+1}$. Including $K_{j,j+1}$ and the non-orthonormal overlap M in the generalized eigenvalue problem gives $\Omega_{max}\cdot\xi = 2.97 \pm 0.04$ (Gaussian, harmonic-matched kernel) where [1] reports $\Omega_{max}\cdot\xi \approx 1.48$. This follows from evaluation of the full projected operator, which includes contributions not present in the reduced projection.

(iii) Overlap matrix structurally non-trivial. The adjacent-mode overlap $M_1 \in [0.71, 0.89]$ across the benchmark range. The minimal-model assumption $M \approx I$, which underlies the two-parameter closed form of [2] §10.10 and the $\varepsilon \lesssim 0.1$ bound of [1] §5.3, is not satisfied even approximately. The (a, b) parameters of the prior papers are therefore to be understood as effective coordinates — a reduced description of the spectrum — rather than as canonical physical masses and couplings.

(iv) (M1) nearest-neighbour restriction misses a comparable next-nearest coupling. The full circulant fit gives $\kappa_2/\kappa_1 \approx -0.21$ at $\sigma/d = 0.8$ — five times larger in magnitude than [1]'s §5.4 estimate of ≤4%, and opposite in sign, so it raises Ω_{max} rather than perturbing it within the minimal-model band.

(v) Cross-kernel robustness improved. The full calculation gives monotonic $\Omega_{max}(\sigma)$ for Gaussian, sech, and Lorentzian mode families, with $\Omega_{max}\cdot\xi \in [2.94, 3.24]$ across all three. The non-monotonic Lorentzian behaviour in [1] Table 4 does not survive the kernel-inclusive computation.

(vi) 3D extension strengthens the scale correction. Treating the physical κ-modes as $\psi_j(x) = \varphi(x_\perp)\, \chi_j(s)$ produces a transverse-kinetic contribution to the effective mass; this renormalizes a upward while leaving b unchanged, and preserves the Ω_{max} stability across σ/d. The 1D treatment is a lower bound on the 3D Ω_{max}.

(vii) Kernel-parameter space narrowed from ±40% to ±20%. The ~40% variation in Ω_{max} across physically reasonable (V_0, ℓ) choices reduces to ~20% once dimensional scaling (ℓ is O(ξ) or O(σ)), self-consistency ($V_0 = \ell^2/(4\sigma^4)$ for harmonic ground states of width σ), and the Gaussian-well shape are applied. $\Delta_\alpha \cdot \xi \in [19, 24]$ across the narrowed space. One pure number ℓ/ξ remains undetermined; its microscopic derivation joins the bath cutoff Λ as the priority open item.

Consequences for the α-saturation prediction. The structural ceiling $\delta C \leq 2.1 \times 10^{-3}$ at $r_{max} = 0.15$ is unchanged — algebraic consequence of [2] §10.10.7. What this paper revises is the saturation condition: $\Delta_\alpha \cdot \xi \approx 20$ with a ±20% band dominated by the narrowed kernel-parameter uncertainty. The hierarchy Δ/ξ required for saturation is approximately twice as steep as [1] reports, which makes the prediction sharper rather than weaker — a δC measurement near the ceiling would confirm a specific steep hierarchy rather than a loose constraint. The prior paper's centred prediction $\delta C = (2.1 \pm 0.3) \times 10^{-3}$ remains valid conditional on saturation; the conditionality is now more informative.

Contents

---

1. Position relative to the prior paper

1.1 What the prior paper establishes

The prior paper [1] supplies three structurally new items to the Pair-Resolved Closure Spectrum framework of [2].

First, a projection ansatz (§3) identifying the pair-amplitude field q_j with the spatial projection of the κ-field onto pair-localized Wannier modes ψ_j. This commits the previously-abstract q_j of [2] §10.10 to a specific microscopic content.

Second, an explicit closure kernel K_{cl} and a harmonic-approximation construction of the ψ_j as Gaussian Wannier states (§4). This makes the mode structure derivable rather than postulated.

Third, a geometric stability observation (§6.3, Table 3) that the combination (a + 3.8019 b) entering Ω_{max}^2 varies by only ±6% across $\sigma/d \in [0.6, 1.0]$ in the Gaussian benchmark. This stability is attributed to a compensation between a decreasing and b increasing as σ/d grows.

The downstream consequences follow. Under harmonic matching $V_0 = 1/(4\sigma^2)$, $\ell = \sigma$, $\Omega_{max}\cdot\xi = \sqrt{(a + 3.80\, b)} \approx 1.48$, whence the α-saturating bath width $\Delta_\alpha = \Omega_{max}/0.15 \approx 9.9\cdot\xi^{-1}$. Combined with the structural ceiling of [2] Appendix E.5, the prior paper reports the centred prediction $\delta C = (2.1 \pm 0.3) \times 10^{-3}$ at $\Delta_\alpha = (10.0 \pm 0.3)\cdot\xi^{-1}$.

1.2 What this paper extends

The prior paper correctly identified the structural form of the projected spectrum and its stability property. The present paper completes that result by removing the minimal projection and evaluating the full projected operator. Two specific consequences follow, each of which sharpens rather than overturns the prior paper's conclusions.

The prior paper's §6.2 identifies the hopping parameter κ through the equality chain

$$\kappa = |G_{j,j+1}| = g_{geom} \cdot \mu^2,$$

with $g_{geom} = |1 - d^2/(4\sigma^2)|\cdot\exp(-d^2/(8\sigma^2))$. Two issues.

First, §5.1 of the prior paper writes $L_{j,k} = G_{j,k} + V_{j,k} + K_{j,k}$ and §6.2 then retains only $G_{j,j+1}$ for the hopping identification. $K_{j,j+1}$ is included in the stated action but dropped from the spectrum. §3.2 below shows that $K_{j,j+1}$ accounts for ≈ 44% of the nearest-neighbour entry $L_{j,j+1}$ of the full projected operator — a dominant contribution, not a correction. The §6.2 identification is therefore incomplete in a specific way: it omits the largest single contribution to the hopping.

Second, the identification $|G_{j,j+1}| = g_{geom} \cdot \mu^2$ as a single equality reads $g_{geom} \cdot G_{jj}$ (which is what the Gaussian integral produces) as $g_{geom} \cdot \mu^2$. These are not the same number; G_{jj} and μ^2 differ by the factor μ^2/G_{jj}, which ranges from ≈ 0.67 at σ/d = 1.0 to ≈ 1.15 at σ/d = 0.6 under [1] Table 1's η definition. The rescaling is implicit, not derived.

More generally, the nearest-neighbour restriction (M1) of [2] §10.10.2 sets the next-nearest-neighbour coupling κ_2 to zero and assumes orthonormal modes M = I. §3.3 and §4 below show that κ_2 is ≈ 21% of κ_1 and opposite in sign, and that $M_1 \in [0.71, 0.89]$ is structurally non-orthonormal. Neither assumption is a quantitative approximation to the projected spectrum.

The present paper therefore computes $L_{j,k}$ in full, includes M off-diagonal elements, solves the circulant generalized eigenvalue problem $L v = \Omega^2 M v$, and reads Ω_{max} directly from the spectrum. The (a, b) parameters are retained in §5 as least-squares fits for comparability with [1], but the physical prediction rests on Ω_{max} rather than on any two-parameter combination.

1.3 What this paper does not challenge

The following content of [1] and [2] is preserved:

- The projection ansatz $q_j = \int \psi_j \cdot \kappa$ itself ([1] §3). This is the framework this paper builds on.
- The structural theorems on δC from [2] Part I.
- The algebraic identity $\delta C = (3/32) \cdot r_{max}^2$ from [2] §8.3 and §10.10.7.
- The α-sector cross-sector constraint $r_{max} \leq 0.15$ from [2] §9.3.
- The structural ceiling $\delta C \leq 2.1 \times 10^{-3}$ at the α-saturating edge from [2] Appendix E.5.
- The Z_7 Fourier decomposition of [2] §§10.2–10.9 reducing the spectrum to a constrained Fourier-mode family.
- The structural intuition of [1] §6 that the projected spectrum has a stable physical quantity — this paper confirms that intuition, identifies Ω_{max} as that quantity, and sharpens the numerical value.

---

2. Inherited framework

2.1 The projection ansatz and the projected action

From [1] §3, the pair-amplitude field q_j is identified with the spatial projection of the κ-field onto Wannier modes:

$$q_j(t) = \int d^3x \cdot \psi_j(x) \cdot \kappa(x, t),$$

with $\{\psi_j\}_{j=0}^{6}$ *approximately orthonormal pair-localized modes on the* $K = 7$ *closure manifold. Substituting* $\kappa(x, t) = \Sigma_j\, q_j(t)\, \psi_j(x) + \kappa_\perp$ into the κ-action and retaining quadratic terms in the longitudinal modes gives the effective action of [1] §5.1:

$$S_{eff} = (1/2) \int dt \left[\Sigma_{jk}\, \dot{q}_j M_{jk} \dot{q}_k - \Sigma_{jk}\, q_j L_{jk} q_k \right],$$

with overlap $M_{jk} = \int \psi_j \psi_k$ and the stiffness matrix L_{jk} built from three contributions.

2.2 The three matrix contributions L = G + V + K

The stiffness matrix decomposes exactly as

$$L_{jk} = G_{jk} + V_{jk} + K_{jk},$$

where

- **$G_{jk} = \int \nabla\psi_j \cdot \nabla\psi_k$** is the gradient-overlap integral (kinetic contribution);
- **$V_{jk} = m_\kappa^2 \cdot M_{jk}$** is the κ-field mass contribution, $m_\kappa^2 = (3/4)\cdot\xi^{-2}$ inherited from the Two-Planck programme via [1] §4.1;
- **$K_{jk} = \int \psi_j K_{cl} \psi_k$** is the closure-kernel cross-term, with K_{cl} the periodic localization potential introduced in [1] §4.1.

All three contributions appear in the action of [1] §5.1. The innovation of the present paper is to compute K_{jk} explicitly and retain it in the spectrum.

2.3 Circulant structure on Z_7

Z_7 cyclic symmetry together with site-translation invariance implies that both M and L are circulant. Their first rows are $\{M_0, M_1, M_2, M_3, M_3, M_2, M_1\}$ and $\{L_0, L_1, L_2, L_3, L_3, L_2, L_1\}$ (real, symmetric under $j \to 7 - j$). Both diagonalize in the discrete Fourier basis, and the spectrum is

$$\Omega_n^2 = \hat{L}_n / \hat{M}_n,\ \hat{L}_n = L_0 + 2L_1 \cos(2\pi n/7) + 2L_2 \cos(4\pi n/7) + 2L_3 \cos(6\pi n/7),$$

with $\hat{M}_n$ given analogously.

The minimal-model form $\Omega_n^2 = \mu^2 + 4\kappa \sin^2(\pi n/7)$ of [2] §10.10 requires both $L_2 = L_3 = 0$ and M = I. Neither holds for the projected operator (§§3, 4).

---

3. The full projected closure operator

3.1 Explicit kernel and Wannier modes

We adopt the kernel and mode construction of [1] §4 without modification. The closure kernel is

$$K_{cl}(s) = V_0 \sum_{j=0}^{6} \exp[-(s - s_j)^2/(2\ell^2)],$$

with sites $s_j = j\cdot d$, $d = \xi$, and periodic boundary conditions. Gaussian Wannier modes in the harmonic approximation:

$$\psi_j(s) = (2\pi\sigma^2)^{-1/4} \cdot \exp[-(s - s_j)^2/(4\sigma^2)],$$

with σ set by the local well curvature. Admissible $\sigma/d \in [0.6, 1.0]$ from [1] §4.3.

The kernel parameters (V_0, ℓ) are not derived in [1]. The prior paper's Tables 1–3 use the implicit harmonic-matching choice $V_0 = 1/(4\sigma^2)$, $\ell = \sigma$, which we adopt as the benchmark in §§3–6 and revisit in §7.

3.2 First-row elements of M, G, K, L

The circulant structure implies we need only the first row of each matrix. Computing numerically for the harmonic-matched Gaussian benchmark:

σ/d	M_0	M_1	M_2	L_0	L_1	K_1
0.60	1.000	0.707	0.250	2.489	1.418	0.738
0.70	1.000	0.775	0.362	2.155	1.468	0.694
0.80	1.000	0.823	0.465	1.923	1.454	0.645
0.90	1.000	0.860	0.560	1.750	1.414	0.599
1.00	1.000	0.890	0.648	1.613	1.367	0.558

Table 1. First-row elements of M, L, and K (selected) for the harmonic-matched Gaussian benchmark. All in ξ^{-2} units except M_k which is dimensionless. L_0 and L_1 include all three contributions G + V + K.

The overlap column is the subject of §4; the L and K columns give the substantive new numerical content here. At σ/d = 0.80 (representative benchmark):

- L_1 = 1.454 is the full nearest-neighbour entry.
- K_1 = 0.645 is the closure-kernel cross-term component.
- $V\_1 = m\_\kappa^2 \cdot M\_1 = 0.75 \cdot 0.823 = 0.617$.
- $G\_1 = L\_1 - V\_1 - K\_1 = 1.454 - 0.617 - 0.645 = 0.192$.

K_1 is 44% of L_1, and G_1 — the component the prior paper used to identify κ — is only 13% of L_1. The dropped K_{j,j+1} contribution is the dominant term, not a correction.

3.3 Direct spectrum and the three-coupling circulant

Solving $\Omega^2 = \hat{M}^{-1} \hat{L}$ in the Z_7 Fourier basis directly yields the full seven-eigenvalue spectrum:

σ/d	Ω_0·ξ	Ω_{1,6}·ξ	Ω_{2,5}·ξ	Ω_{3,4}·ξ	Ω_max·ξ	Δ_α·ξ
0.60	1.340	1.612	2.240	3.023	3.023	20.2
0.70	1.283	1.566	2.206	2.986	2.986	19.9
0.80	1.238	1.529	2.181	2.965	2.965	19.8
0.90	1.203	1.501	2.161	2.949	2.949	19.7
1.00	1.173	1.477	2.145	2.937	2.937	19.6

Table 2. Direct spectrum of the full projected closure operator, harmonic-matched Gaussian benchmark. The Z_7 reflection symmetry gives pairs (1, 6), (2, 5), (3, 4). Ω_max = Ω_3 throughout, and Δ_α = Ω_max/0.15 at the α-sector bound.

The physically relevant quantity is $\Omega_{max}\cdot\xi$ — the maximum spectral frequency of the projected operator. It is obtained directly from the n = 3 eigenvalue of $L\cdot v = \Omega^2\cdot M\cdot v$, with no intermediate fit and no parameterization choice.

Fitting the spectrum to the most general Z_7-admissible form $\Omega_n^2 = \mu^2 + 4\kappa_1 \sin^2(\pi n/7) + 4\kappa_2 \sin^2(2\pi n/7) + 4\kappa_3 \sin^2(3\pi n/7)$ gives an exact fit (four independent parameters after reflection symmetry, seven data points). At $\sigma/d = 0.8$:

$\mu^2 \approx 1.533\cdot\xi^{-2}$, $\kappa_1 \approx 1.929\cdot\xi^{-2}$, $\kappa_2 \approx -0.412\cdot\xi^{-2}$, $\kappa_3 \approx 0.095\cdot\xi^{-2}$.

The non-negligible κ_2 is a central new structural finding. It is ≈ 21% of κ_1 in magnitude and opposite in sign, which raises Ω_{max} relative to the minimal-model (M1) prediction rather than shifting it in a way that could partly cancel within a small-correction band.

---

4. The overlap matrix: structural breakdown of M ≈ I

4.1 The assumption and what it underlies

The minimal-model closed form of [2] §10.10 — spectrum $\Omega_n^2 = \mu^2 + 4\kappa \sin^2(\pi n/7)$, threshold correction $\delta C = (3/32d^2)(a + 3.8019\ b)$ — depends on two structural assumptions: that L is restricted to nearest-neighbour couplings (M1), and that M = I so the generalized eigenvalue problem reduces to ordinary diagonalization of L. The second assumption is implicit in the §10.10 derivation and is made explicit in [1] §5.3, which bounds the correction by $|\lambda_n(M - I)| \lesssim 2\varepsilon$ with $\varepsilon = \max|M_{j,j+1}| \lesssim 0.1$, citing this as "a ~10% correction".

4.2 The actual overlap values

The M_1 column of Table 1 shows the adjacent-mode overlap across the benchmark range:

$M_1 \in [0.707, 0.890]$ across $\sigma/d \in [0.6, 1.0]$.

This is not a small correction. The adjacent modes overlap at 70–90%. The [1] §5.3 bound's premise — $\varepsilon \lesssim 0.1$ — is violated by about an order of magnitude across the entire physically admissible range.

Propagating: the eigenvalue-of-M correction is $|\lambda_n(M - I)|$, and for n = 3 in the harmonic-matched Gaussian benchmark $\hat{M}_3 \approx 0.17$ at $\sigma/d = 0.8$. The generalized eigenvalue problem $\Omega_3^2 = \hat{L}_3/\hat{M}_3$ divides $\hat{L}_3$ by a factor ≈ 0.17, not by something close to 1. The M ≈ I approximation would underestimate Ω_{max}^2 by a factor of roughly 6.

4.3 Implications

Two.

First, the (a, b) parameterization cannot be interpreted as physical masses and couplings in the standard sense. The minimal-model Lagrangian $F^{(2)}[\delta q] = (1/2)\ \Sigma[\mu^2(\delta q_j)^2 + \kappa(\delta q_{j+1} - \delta q_j)^2]$ of [2] §10.10.3 assumes canonically normalized fields δq_j, which is equivalent to M = I after projection. Under the projection ansatz of [1] §3 with Gaussian Wannier modes at the admissible widths, the resulting δq_j are not canonically normalized — M_1 is a structural feature of the projection, not a small correction. A canonical rescaling would change the effective (a, b) by the factors in M.

Second, the minimal model is best understood as a reduced description of the projected physics rather than as its leading-order approximation. It is a distinct two-parameter model that, by least squares, reproduces the dominant Ω_3 eigenvalue approximately (§5.2 below) but fails at the 10% level on the other eigenvalues. The (a, b) parameters are effective coordinates on the projection of the true three-coupling, non-orthonormal circulant onto a chosen two-parameter subspace. They retain their role as a reduced description useful for comparison with the prior literature and for the algebraic identity $\delta C = (3/32d^2)(a + 3.80b)$; they do not admit direct interpretation as canonical physical masses and couplings.

This reframing is consequential for the programme's next steps: a microscopic derivation of (a, b) alone is not sufficient, because under the full projection (a, b) are reduced coordinates rather than fundamental couplings. The correct target is the full projected operator (L, M) — equivalently, Ω_{max} directly — of which (a, b) are the LSQ projection onto the two-parameter minimal-model subspace. This is a strengthening of Paper 1's (a, b, d) formulation, not a replacement of it: δC is still controlled by the two-parameter closed form at the level of the algebraic identity $\delta C = (3/32d^2)(a + 3.80b)$, but the microscopic derivation must reach the operator level to specify those parameters consistently.

---

5. Ω_{max} as the central physical quantity

5.1 Direct Ω_{max} readoff across σ/d

The right-most columns of Table 2 show the stability result in its cleanest form. Across $\sigma/d \in [0.6, 1.0]$:

$\Omega_{max}\cdot\xi \in [2.937, 3.023]$, mean 2.97, variation ±1.5%.

This is better than [1]'s ±3% estimate for the same quantity through the (a + 3.80b) route. The stability comes from the full projected operator — it does not require the minimal-model assumption — and is a direct property of the spectrum rather than a fit output.

5.2 Minimal-model (a, b) as a least-squares projection

For comparability with [1] and [2], we report the minimal-model LSQ fit to $\Omega_n^2 = a\cdot\xi^{-2} + 4b\cdot\xi^{-2} \sin^2(\pi n/7)$:

σ/d	a (LSQ)	b (LSQ)	a + 3.80 b	max residual
0.60	1.139	1.953	8.56	0.90 (10%)
0.70	1.007	1.934	8.36	0.87 (10%)
0.80	0.899	1.929	8.23	0.86 (10%)
0.90	0.813	1.928	8.14	0.86 (10%)
1.00	0.743	1.928	8.07	0.86 (10%)

Table 3. Minimal-model (a, b) as least-squares fits to the full seven-eigenvalue spectrum. The max residual column gives the largest $|\Omega_n^2 - (a + 4b \sin^2(\pi n/7))|$ at each σ/d, expressed in ξ^{-2} units and as a fraction of Ω_{max}^2.

Three observations.

First, the LSQ-fit b is nearly constant at ≈ 1.93 across the benchmark range, varying by only 1%. This is not the same quantity as [1] Table 2's $b = g_{geom} \cdot a$, which varies from 0.25 to 0.44. The LSQ-fit b and [1]'s $g_{geom} \cdot a$ are different quantities computed by different procedures; they cannot be compared directly.

Second, the minimal-model fit has residuals of ≈ 10% of Ω_{max}^2. The two-parameter form cannot capture the full three-coupling spectrum exactly; the minimal model is not a limit of the true physics but a projection of it.

Third, despite the 10% residuals, the combination a + 3.80 b is stable at ±3% across σ/d — consistent with [1] §6.3's observation. The mechanism is not the (a, b) compensation [1] proposes. In an LSQ fit of a two-parameter model to seven eigenvalues, the combination a + 3.80 b is the value closest to $\Omega_3^2 = \Omega_{max}^2$ in the least-squares sense; it inherits the stability of Ω_{max} itself.

5.3 Reconciliation with the prior paper's compensation result

[1] Table 3 shows (a + 3.80b) varying from 2.084 to 2.353 across σ/d ∈ [0.6, 1.0]; the present paper's Table 3 shows it varying from 8.07 to 8.56. The two sets of numbers differ by a scale factor of about 3.7.

Two reasons. The prior paper's numbers come from the $g_{geom} \cdot a$ identification of [1] §6.2 rather than from the full projected operator; this effectively drops $K_{j,j+1}$ (≈ 44% of L_1) and the rescaling by non-orthonormal M. And the stability mechanism [1] describes — a decreases while b increases — does not appear in the full calculation: in Table 3 above, b is nearly constant and a does the entire variation. The combination is stable because the LSQ fit tracks Ω_{max}^2, which is itself stable.

The structural intuition of [1] is correct: the K = 7 closure geometry stabilizes the physically relevant spectral quantity against details of the projection-mode width. The mechanism is the specific circulant structure of the full projected L and M — not a cancellation between effective parameters within a minimal-model form. Ω_{max} is the physically meaningful observable; the (a, b) parameters are effective coordinates that express a reduced description of the spectrum, and the leading physical consequence — the stability of a + 3.80 b — is preserved as a projection of the stability of Ω_{max} itself.

---

6. The corrected absolute scale

6.1 Gaussian benchmark

From Tables 2 and 3:

- $\Omega_{max} \cdot \xi = 2.97 \pm 0.04$ (±1.5% across $\sigma/d \in [0.6, 1.0]$)
- $\Delta_\alpha \cdot \xi = \Omega_{max} \cdot \xi / r_{max} = \Omega_{max} \cdot \xi / 0.15 = 19.8 \pm 0.3$
- Corresponding $(a + 3.80b)_{LSQ} = 8.27 \pm 0.25$

Relative to [1] §8.2's $\Delta_\alpha = (10.0 \pm 0.3) \cdot \xi^{-1}$, the saturating bath width is approximately twice as large — reflecting the inclusion of the $K_{j,j+1}$ cross-term and the non-orthonormal overlap in the eigenvalue problem, both of which the prior paper's reduced projection did not evaluate. The uncertainty band within the Gaussian benchmark is ±1.5% here rather than the ±3% of [1].

---

7. Narrowing the kernel-parameter space

7.1 The unconstrained scan

The harmonic-matching choice $V_0 = 1/(4\sigma^2)$, $\ell = \sigma$ is motivated in [1] §4.3 as producing Gaussian Wannier ground states of width σ, but is not a derivation. Varying (V_0, ℓ) within the regime where the harmonic approximation remains sensible ($V_0 \cdot \ell^2 \gg 1$, $\ell \sim d$) scans a two-parameter space the prior paper leaves implicit. $\Omega_{max} \cdot \xi$ at $\sigma/d = 0.8$ across this space:

$V_0 \cdot \sigma^2$	$\ell/\sigma = 0.5$	$\ell/\sigma = 0.7$	$\ell/\sigma = 1.0$	$\ell/\sigma = 1.3$
0.10	2.76	2.77	2.79	2.80
0.25	2.80	2.83	2.87	2.91
0.50	2.87	2.92	3.00	3.08
1.00	3.00	3.11	3.25	3.39
1.50	3.13	3.28	3.49	3.68
2.00	3.25	3.44	3.71	3.95

Table 4. $\Omega_{max} \cdot \xi$ at $\sigma/d = 0.8$ across a two-dimensional (V_0, ℓ) grid. Range 2.76 to 3.95 — factor 1.43. This is the unconstrained-scan uncertainty.

Across this grid, $\Omega\_max \cdot \xi$ varies from 2.76 to 3.95, a factor of 1.43 (equivalent to ~40% in the symmetric-band convention). Taken at face value, this would translate to $\Delta\_\alpha \cdot \xi$ varying from 18.4 to 26.3 — a serious but bounded ambiguity.

7.2 Three narrowing principles

Once the projection ansatz is enforced consistently, the kernel parameters are uniquely fixed by the projection framework up to a single dimensionless ratio ℓ/ξ. The unconstrained 2D scan of §7.1 overstates the actual ambiguity because it treats as free three inputs that consistency of the projection determines.

Principle 1: Dimensional scaling. The closure manifold has one intrinsic length scale ξ. Any kernel parameter must be expressible as a pure number times a power of ξ. The well width ℓ must be $O(\xi)$, with the two natural scalings being $\ell = O(\sigma)$ (kernel tracks mode width) and $\ell = O(\xi)$ (kernel tracks closure scale). The well depth V_0 has dimensions ξ^{-2} and must similarly be expressible as a pure number times ξ^{-2}. This alone restricts the (V_0, ℓ) space from two continuous dimensions to a one-parameter family plus two discrete choices of scaling convention.

Principle 2: Self-consistency of projection. The projection ansatz of [1] §3 projects the κ-field onto modes ψ_j that are assumed to be the ground states of the kernel K_cl. The harmonic approximation represents these ground states as Gaussians of width σ. For the Gaussian form to actually be the ground state of a single well of depth V_0 and width ℓ, the well curvature at its minimum must match the Gaussian: $V\_0/\ell^2 = 1/(4\sigma^4)\cdot$(factor from the 1/2 kinetic normalization), giving

$$V_0 = \ell^2/(4\sigma^4).$$

This is not an additional choice; it is the condition that the modes used to project are actually the modes the kernel supports. It reduces the (V_0, ℓ, σ) three-dimensional space to a two-dimensional subspace — equivalently, the (ℓ, σ) plane with V_0 determined.

Principle 3: Gaussian-well shape. The kernel K_cl is a sum of smooth wells of a specific shape (Gaussian). Alternative shapes — cosine arrays, Lorentzian wells, etc. — fall outside the harmonic-approximation framework of [1] §4 and would require the non-harmonic treatment flagged as future work. Within the Gaussian family, the only parameters are the two above.

Applying all three principles to the unconstrained 2D scan:

7.3 Narrowed ranges

Under self-consistency $V_0 = \ell^2/(4\sigma^4)$, the two natural dimensional scaling choices give:

Resolution A: $\ell = \sigma$ (harmonic matching, paper's benchmark).

σ/d	$V_0\cdot\sigma^2$	ℓ/σ	$\Omega\_max\cdot\xi$	$\Delta\_\alpha\cdot\xi$
0.60	0.250	1.00	3.023	20.2

σ/d	$V_0 \cdot \sigma^2$	ℓ/σ	$\Omega\_max \cdot \xi$	$\Delta\_\alpha \cdot \xi$
0.70	0.250	1.00	2.986	19.9
0.80	0.250	1.00	2.965	19.8
0.90	0.250	1.00	2.949	19.7
1.00	0.250	1.00	2.937	19.6

Range: $\Omega\_max \cdot \xi \in [2.94, 3.02]$, ±1.4%.

Resolution B: $\ell = d = \xi$ (site-scale kernel).

σ/d	$V_0 \cdot \sigma^2$	ℓ/σ	$\Omega\_max \cdot \xi$	$\Delta\_\alpha \cdot \xi$
0.60	0.694	1.67	3.596	24.0
0.70	0.510	1.43	3.261	21.7
0.80	0.391	1.25	3.088	20.6
0.90	0.309	1.11	2.993	20.0
1.00	0.250	1.00	2.937	19.6

Range: $\Omega\_max \cdot \xi \in [2.94, 3.60]$, ±10%.

Table 5. $\Omega\_max \cdot \xi$ under self-consistency, for the two natural dimensional scalings of ℓ. At $\sigma = \xi = d$ the two resolutions coincide.

7.4 What this leaves open

Resolutions A and B agree at $\sigma = \xi = d$ and diverge as σ moves away. The physical σ is itself determined by the kernel's own ground-state width (not freely varied), so the range shown across $\sigma/d \in [0.6, 1.0]$ is an upper bound on the actual spread. Taking Resolutions A and B as bracketing the physically realizable kernel choices:

$\Omega\_max \cdot \xi \in [2.94, 3.60]$, $\Delta\_\alpha \cdot \xi \in [19.6, 24.0]$.

Symmetric-band convention: $\Delta\_\alpha \cdot \xi = (21.8 \pm 2.2)$, roughly ±10% — less than half the ~20% suggested by a naive O(1) variation in $V_0 \cdot \sigma^2$, and less than a third of the ~40% of the unconstrained 2D scan.

The residual ambiguity is one pure number — the ratio ℓ/ξ — which is either set by the closure-manifold's intrinsic scale ξ (Resolution B) or by the projected mode width σ (Resolution A). Adjudicating between these requires a microscopic model of what the closure kernel physically represents, which neither [1] nor the present paper supplies. §12.2 discusses this as the priority open problem alongside the bath cutoff Λ.

8. Cross-kernel treatment

8.1 Gaussian, sech, and Lorentzian families

Replacing the Gaussian Wannier modes with alternative shapes and repeating the full projected calculation with the harmonic-matched kernel in each case:

σ/d	Gaussian Ω_max·ξ	sech Ω_max·ξ	Lorentzian Ω_max·ξ
0.60	3.023	3.151	3.243
0.70	2.986	3.096	3.186
0.80	2.965	3.053	3.140
0.90	2.949	3.018	3.100
1.00	2.937	2.991	3.063

Table 6. Ω_max·ξ across three mode families, harmonic-matched kernel throughout.

All three families give monotonically decreasing Ω_max with σ/d, stable to ±1.5% within each family, with absolute values clustered near $3.0\cdot\xi^{-1}$. The full range across all three families and all σ/d values is [2.94, 3.24] — ±5% variation — giving $\Delta\_\alpha \in [19.6, 21.6]\cdot\xi^{-1}$ across the set.

8.2 Resolution of the prior paper's Table 4 Lorentzian anomaly

[1] Table 4's Lorentzian column reports non-monotonic behaviour in (a + 3.80b). Table 6 above shows that the same calculation with K_{j,j+1} included gives the Lorentzian a monotonically decreasing Ω_max·ξ from 3.24 to 3.06. The non-monotonic behaviour of [1] Table 4 is an artefact of the g_geom·a identification applied to long-tailed modes; the actual projected spectrum is monotonic. This is one of the clearer diagnostics that the (a + 3.80b) combination is not the right frame for the physical prediction.

8.3 Cross-family vs kernel-parameter variation

At the absolute-scale level, harmonic-matched Gaussian gives $\Delta\_\alpha = 19.8\cdot\xi^{-1}$, sech $20.4\cdot\xi^{-1}$, Lorentzian $20.8\cdot\xi^{-1}$ — all within 5%. Combined with the §7 narrowing, cross-family uncertainty (±5%) is sub-dominant to kernel-parameter uncertainty (±10%) by roughly a factor of two.

This is a notable change from [1] §6.4, which identifies cross-kernel variation as "the dominant residual geometric uncertainty". Under the full calculation with the §7 narrowing, cross-family variation is a smaller contributor; kernel-parameter uncertainty dominates.

---

9. Three-dimensional extension

9.1 Transverse profile factorization

The κ-field is a 3+1-dimensional object; the closure manifold is 1-dimensional. Physical pair-localized modes factor as

$$\psi_j(x) = \varphi(x_\perp) \cdot \chi_j(s),$$

with $\varphi(x_\perp)$ a transverse profile common to all pairs and $\chi_j(s)$ the 1D Wannier modes. The overlap and kernel integrals factorize:

- $M_{jk}^{3D} = \int ds\, \chi_j\, \chi_k$ (for unit-normalized φ)
- $K_{jk}^{3D} = \int ds\, \chi_j\, K_{cl}\, \chi_k$

The gradient integral picks up a transverse contribution that does not factor with M:

$$G_{jk}^{3D} = G_{jk}^{1D} + \Gamma \cdot M_{jk}^{1D},\ \Gamma \equiv \int d^2x_\perp\, |\nabla_\perp \varphi|^2.$$

9.2 Effective mass renormalization

Substituting:

$$L_{jk}^{3D} = L_{jk}^{1D} + \Gamma \cdot M_{jk}^{1D}.$$

The transverse penalty acts exactly as an additive mass contribution: $m_\kappa^2 \to m_\kappa^2 + \Gamma$. Structure of the generalized eigenvalue problem is unchanged; only the effective mass shifts.

9.3 Impact on Ω_{max} and Δ_α

$\Gamma \cdot \xi^2$	$m_{eff}^2 \cdot \xi^2$	$\Omega_{max} \cdot \xi$ (mean)	$\Delta_\alpha \cdot \xi$
0.00 (1D limit)	0.75	2.97	19.8
0.25	1.00	3.01	20.1
0.75	1.50	3.09	20.6
3.00	3.75	3.44	22.9

Table 7. Effect of the transverse kinetic penalty Γ on Ω_{max} and Δ_α.

Three points. The σ/d-stability is preserved (variation remains ±1.5–2.2% per row). The absolute scale increases monotonically with Γ. The 1D treatment is a lower bound on the 3D Ω_{max}; any non-trivial transverse profile raises the hierarchy the microscopic physics must deliver.

The microscopic Γ — equivalently, the transverse scale $\xi_\perp$ — is set by the same deeper VERSF structure that determines the closure kernel's parameters. It joins (V_0, ℓ) and Λ as a microscopic input the present programme does not yet derive.

10. Higher-neighbour circulant couplings

10.1 The three-coupling fit

The most general Z_7-symmetric circulant on seven sites has the Fourier form

$$\Omega\_n^2 = \mu^2 + 4\kappa_1 \sin^2(\pi n/7) + 4\kappa_2 \sin^2(2\pi n/7) + 4\kappa_3 \sin^2(3\pi n/7),$$

corresponding to nearest-neighbour κ_1, next-nearest κ_2, and third-neighbour κ_3. Fitting the spectrum of Table 2:

σ/d	$\mu^2\cdot\xi^2$	$\kappa_1\cdot\xi^2$	$\kappa_2\cdot\xi^2$	$\kappa_3\cdot\xi^2$
0.60	1.794	1.953	−0.428	0.100
0.70	1.645	1.934	−0.415	0.096
0.80	1.533	1.929	−0.412	0.095
0.90	1.446	1.928	−0.411	0.094
1.00	1.377	1.928	−0.411	0.094

Table 8. Full three-coupling fit. $\kappa_2/\kappa_1 \approx -0.21$ and $\kappa_3/\kappa_1 \approx 0.05$, stable across σ/d.

10.2 Size of κ_2 relative to κ_1

Across $\sigma/d \in [0.6, 1.0]$, $|\kappa_2/\kappa_1| \in [0.21, 0.22]$. The next-nearest coupling is ≈ 21% of the nearest-neighbour coupling in magnitude and opposite in sign. The third-neighbour ratio $|\kappa_3/\kappa_1| \in [0.049, 0.051]$ is about 5%.

10.3 The prior paper's §5.4 estimate

[1] §5.4 bounds the next-nearest effect by $\varepsilon_2 = |G\_\{j,j+2\}|/|G\_\{j,j+1\}| \lesssim 0.2$ and translates this to "a ≲ 4% effect on δC". The estimate is based on the gradient-overlap ratio alone, without K_{j,j+2}. The actual $|\kappa_2/\kappa_1| \approx 0.21$ is similar in magnitude to the G-only bound but enters the spectrum with opposite sign, so it raises Ω_max rather than shifting it within a small-correction band.

Quantitatively: the net shift in Ω_max^2 at σ/d = 0.8 from including κ_2 properly (relative to setting it to zero) is approximately $4\cdot(\sin^2(2\pi\cdot 3/7) - \sin^2(\pi\cdot 3/7))\cdot\kappa_2 = 4\cdot(0.753 - 3.802)\cdot(-0.412) \approx +5.0\cdot\xi^{-2}$, a ≈ 57% increase in Ω_max^2 at σ/d = 0.8. The ≲ 4% second-neighbour correction claim of [1] §5.4 underestimates the actual effect by more than an order of magnitude.

The (M1) nearest-neighbour locality assumption is therefore not a quantitative approximation for the projected spectrum; it is a definite model choice that sets a sizable structural contribution to zero. Its role in [2] §10.10 — producing the explicit two-parameter closed form — is retained; but the microscopic interpretation of that closed form must carry the understanding that κ_2 is a structural contribution, not a small correction.

11. Consequences for the α-saturation prediction

11.1 Ceiling unchanged (algebraic)

The structural identity

$$\delta C = (3/32) \cdot r_{max}^2 = (3/32) \cdot (\Omega_{max} / \Delta)^2$$

is algebraic within the response formalism of [2] Part I. At $r_{max} = 0.15$ (from [2] §9.3), the ceiling

$$\delta C_{\alpha\text{-sat}} = (3/32) \cdot (0.15)^2 = 2.11 \times 10^{-3}$$

is unchanged by anything in this paper.

11.2 Saturating bath width Δ_α revised upward

What changes is the value of Δ at which saturation occurs. In the Gaussian benchmark with harmonic-matched kernel:

$$\Delta_\alpha = \Omega_{max} / r_{max} = (2.97 \pm 0.04)\cdot\xi^{-1} / 0.15 = (19.8 \pm 0.3)\cdot\xi^{-1}.$$

Under the §7 narrowing across both natural dimensional scalings:

$$\Delta_\alpha\cdot\xi \in [19.6, 24.0], \text{ central value} \approx 21.8, \pm 10\%.$$

Compared to [1] §8.2's $\Delta_\alpha = (10.0 \pm 0.3)\cdot\xi^{-1}$, the saturating bath width is approximately twice as large. The shift comes from the full operator evaluation of §3 — including the $K_{j,j+1}$ kernel cross-term and the non-orthonormal overlap M — not from any change in the underlying framework.

11.3 Uncertainty budget

Source	**Δ_α range (ξ^{-1})**	**Effect**	**Status**
Within-Gaussian σ/d (harmonic-matched kernel)	[19.6, 20.2]	±1.5%	Computed (§5)
Cross-family (Gaussian/sech/Lorentzian)	[19.6, 21.6]	±5%	Computed (§8)
Narrowed kernel-parameter (Resolutions A, B)	[19.6, 24.0]	±10%	Computed (§7)
3D transverse penalty ($\xi_\perp \in [0.5\xi, 2\xi]$)	[19.8, 22.9]	±8%	Open (§9, §12)
Higher-neighbour couplings beyond κ_3	< 2%	bounded	Sub-dominant

Table 9. Uncertainty budget for Δ_α. Kernel-parameter uncertainty is the dominant known term; transverse-profile ambiguity is comparable and requires independent resolution.

Including all sources:

$\Delta\_\alpha \in [18, 25] \cdot \xi^{-1}$, central ≈ 21, with ±15% band.

11.4 Why this is a sharper, not weaker, prediction

The conditional centred prediction $\delta C = (2.1 \pm 0.3) \times 10^{-3}$ assumes α-saturation, $\Delta = \Delta\_\alpha$. With the corrected $\Delta\_\alpha \approx 21\cdot\xi^{-1}$ rather than $10\cdot\xi^{-1}$, the bath-to-closure hierarchy required for saturation is approximately twice as steep. Three implications.

First, *the prediction is more informative, not less.* An experiment measuring δC near the ceiling would, under the prior paper's framing, confirm a modest bath-scale separation ($\Delta/\xi \approx 10$). Under the revised framing, the same measurement would confirm a much steeper separation ($\Delta/\xi \approx 20$) — a stronger structural claim about the underlying microscopics. A single ceiling-level measurement now carries more structural weight.

Second, *sub-saturation becomes more diagnostic.* The prior framing interpreted measurements below the ceiling primarily as evidence about Δ. Under the revised framing, the same measurements are more sensitively diagnostic of the microscopic hierarchy: $\delta C = 1 \times 10^{-3}$ corresponds to $\Delta \approx 30\cdot\xi^{-1}$ rather than $\Delta \approx 14\cdot\xi^{-1}$, a factor-of-two-larger inferred bath width. The inverse-square mapping $\delta C \propto 1/(\Delta/\Delta\_\alpha)^2$ is unchanged; what changes is the absolute scale of Δ the data would pin down.

Third, *the open microscopic problems are sharper targets.* The bath cutoff Λ (§12.1) must produce $\Delta \approx 21\cdot\xi^{-1}$ rather than $\approx 10\cdot\xi^{-1}$ for saturation; this is a specific quantitative target the Λ-derivation paper must hit, not a range. Similarly, the kernel parameter ℓ/ξ (§12.2) has a narrow admissible range (roughly 1.0–1.7 under the two dimensional scalings) within which the framework is self-consistent.

The conditionality of [1]'s centred prediction on saturation is therefore not a weakness. It is the mechanism by which a measurement of δC pins down microscopic physics.

---

12. Open problems and their priorities

12.1 Bath cutoff Λ (priority 1, inherited)

The independent derivation of Λ remains the priority open item. With the corrected $\Delta\_\alpha \approx 21\cdot\xi^{-1}$, the bath cutoff must satisfy $\Lambda \gtrsim \Delta\_\alpha/\sqrt{12} \approx 6.1\cdot\xi^{-1}$ (using the super-ohmic $p = 2$ relation of [1] §7.2) for saturation — roughly twice the $2.9\cdot\xi^{-1}$ that [1] §7.3 reports.

12.2 Closure-kernel parameter ℓ/ξ (priority 2, new)

Introduced in the present paper. Enforcing the projection ansatz consistently — dimensional scaling, self-consistency $V_0 = \ell^2/(4\sigma^4)$, and the Gaussian-well shape — collapses the (V_0, ℓ) space to a one-parameter family indexed by ℓ/ξ. Two natural values of that pure number — $\ell = \sigma$ (tracks mode width) and $\ell = \xi$ (tracks closure scale) — give $\Delta_\alpha \cdot \xi = 19.8$ and 20.6 respectively at $\sigma/d = 0.8$, diverging to [19.6, 24.0] across the benchmark σ/d range.

Three candidate resolutions, in decreasing generality:

- *Resolution A — kernel tracks mode width ($\ell = \sigma$).* The kernel parameters depend on the projected mode width. Conceptually this treats K_cl as a response to the mode structure rather than an external input.
- *Resolution B — kernel tracks closure scale ($\ell = \xi$).* The kernel is a fixed feature of the closure manifold with width set by the coherence scale. The closure sector's geometry determines the kernel; modes then adjust to it.
- *Resolution C — microscopic derivation from VERSF first principles.* Derive K_cl from the closure sector's fundamental Lagrangian, supplying both ℓ and V_0 as outputs rather than inputs.

Resolution B is the cleanest in that it treats the kernel as a structural feature of the framework rather than a self-referentially defined one. It also produces the wider uncertainty band on Δ_α, which is the conservative position to adopt until C is available. Resolution C is the target; progress on the κ-sector Lagrangian (e.g. via the Fact Momentum programme [6]) is the natural path.

12.3 Transverse profile $\xi_\perp$ (priority 3, new)

Also introduced in the present paper, via the 3D extension of §9. The transverse kinetic penalty $\Gamma = \int|\nabla_\perp\varphi|^2$ *shifts the effective mass and, through it, Ω_{max}. Physical Γ depends on the transverse scale $\xi_\perp$* of the κ-field's mode; for $\xi_\perp \sim \xi$, $\Gamma \sim \xi^{-2}$, producing a ~10% increase in Δ_α. For more tightly confined profiles, the increase is larger.

The resolution of this joins (12.2) in requiring a microscopic model of the closure sector's Lagrangian.

12.4 Non-harmonic closure modes (priority 4, inherited)

[1] §11.2 flags the non-harmonic construction as sharpening cross-kernel uncertainty. §8 here reduces cross-kernel variation to ±5% on Δ_α — sub-dominant to both (12.1), (12.2), and (12.3). The priority of this item is appropriately downgraded.

---

13. Falsifiability

The falsifiability criteria of [1] and [2] remain in force with refinements.

Point prediction. $\delta C \leq 2.1 \times 10^{-3}$ (ceiling, from [2] §10.10.8, unchanged).

Saturation condition. $\delta C \approx 2.1 \times 10^{-3}$ if microscopic $\Delta = \Delta_\alpha \approx 21 \cdot \xi^{-1}$ (revised from [1]'s $\approx 10 \cdot \xi^{-1}$). A measurement at the ceiling confirms this hierarchy.

Relational prediction (inherited). Commitment-sector and α-sector errors track each other under refinement.

New relational prediction. Under Resolution B of §12.2 ($\ell = \xi$), the kernel's effective width is independently measurable through closure-sector observables other than δC. Those measurements and δC must be consistent under the full projected calculation of §3.

Measurement bands:

Measured δC	Interpretation
$> 3 \times 10^{-3}$	α-sector derivation missed corrections, or framework structural assumption fails
$2 \times 10^{-3} \pm 0.3$	α-saturation achieved; microscopic $\Delta \approx 21 \cdot \xi^{-1}$ (confirms steep hierarchy)
$1–2 \times 10^{-3}$	Sub-saturation; inferred $\Delta \approx 25–30 \cdot \xi^{-1}$
10^{-4} to 10^{-3}	Substantial sub-saturation; $\Delta \approx 30–60 \cdot \xi^{-1}$
$< 10^{-4}$	Projection ansatz or structural assumption falsified

Table 10. Measurement bands and interpretation, updated for the corrected $\Delta_\alpha \approx 21 \cdot \xi^{-1}$.

---

14. Conclusion

The prior paper correctly identified the structural form of the K = 7 projected spectrum and its stability property. The present paper completes that result by evaluating the full projected operator rather than its reduced projection, with three specific consequences that sharpen the picture without overturning the underlying structure.

The prior paper's microscopic programme rests on three claims: (i) that the pair-amplitude field is the spatial projection of the κ-field onto Wannier modes, (ii) that the closure kernel K_{cl} produces those modes as Wannier ground states, and (iii) that the combination of effective parameters controlling Ω_{max} is stable against details of the mode width. The present paper preserves (i) and (ii), confirms (iii) in its physical form (Ω_{max} itself is stable), and quantifies the numerical consequences by computing the full projected closure operator $L = G + V + K$.

The minimal-model (a, b) parameterization, which carried the microscopic interpretation in [1], is a reduced description of the projected physics rather than its literal form. Two specific structural features are now visible: the overlap matrix is far from the identity (adjacent modes overlap at 70–90%), and the next-nearest-neighbour coupling κ_2 is $\approx 21\%$ of κ_1 with opposite sign, not the $\lesssim 4\%$ the prior paper estimates. The (a, b) that emerge from LSQ fits are effective coordinates rather than canonical physical constants. Their role as a reduced description is retained, and the physical observable is Ω_{max}, read directly from the spectrum.

Evaluating the full operator (rather than the reduced projection used in [1] §6.2) shifts the absolute normalization: $\Omega_{max}\cdot\xi \approx 2.97$ with a structurally narrowed kernel-parameter band giving $\Omega_{max}\cdot\xi \in [2.94, 3.60]$. The saturating bath width is $\Delta_\alpha \approx 21\cdot\xi^{-1}$ (band [19.6, 24.0]) — roughly twice what [1] reports from the reduced projection. This follows from the inclusion of operator contributions not present in the reduced projection. The α-saturation ceiling $\delta C \leq 2.1 \times 10^{-3}$ is unchanged — algebraic, and independent of the operator-evaluation details.

The steeper hierarchy required for saturation makes the prediction sharper, not weaker. A ceiling-level δC measurement would confirm $\Delta/\xi \approx 20$, a stronger structural claim than the prior paper's ≈ 10; sub-ceiling measurements would be more diagnostic of microscopic bath structure than under the prior framing.

Once the projection ansatz is enforced consistently, the kernel parameters are uniquely fixed by the projection framework up to a single dimensionless ratio ℓ/ξ: dimensional scaling fixes the form of ℓ, self-consistency $V_0 = \ell^2/(4\sigma^4)$ determines V_0 from (ℓ, σ), and the Gaussian-well shape is preserved. The (V_0, ℓ) space therefore collapses to a one-parameter family indexed by ℓ/ξ, and the ±40% unconstrained ambiguity reduces to ±10% across the two natural values of that pure number. One dimensionless ratio ℓ/ξ remains undetermined; it joins the bath cutoff Λ and the transverse scale $\xi_\perp$ as the priority microscopic open problems — three specific numerical targets the programme's next stage must address.

Compared to the prior paper's picture, the current paper's picture is more structurally complete (the full operator is evaluated, and the minimal model is recognized as a projection of that operator), more demanding of microscopic physics (steeper hierarchy), and more informative per measurement (ceiling confirms a stronger structural claim). The commitment-threshold correction remains a predicted rather than free quantity, with the remaining freedom confined to three specific microscopic parameters that the follow-on work must address.

Appendix A: Numerical methodology

All calculations use the following pipeline.

A.1 Mode construction. Periodic sum of continuum mode over image positions $s_j + n\cdot L$, $L = 7d$, $n \in \{-3, \ldots, 3\}$. Saturates to 10^{-6} at benchmark widths. Modes normalized numerically over the periodic cell.

A.2 Matrix elements. First-row elements M_k, G_k, K_k for $k \in \{0, 1, 2, 3\}$ computed via adaptive quadrature (scipy.integrate.quad, limit=200). Kernel K_{cl} evaluated as a periodic sum of base wells. Remaining circulant entries follow from reflection symmetry.

A.3 Spectrum. $\Omega_n^2 = \hat{L}_n / \hat{M}_n$ via numpy.fft.fft on the first rows. Reflection symmetry gives four independent eigenvalues with (n, 7−n) degenerate.

A.4 Fits. LSQ fits via numpy.linalg.lstsq on linear design matrices. Three-coupling fit is exact by construction (four parameters, four independent eigenvalues after symmetry).

A.5 Self-consistency (§7). The relation $V_0 = \ell^2/(4\sigma^4)$ derives from matching the single-well harmonic ground state of frequency $\omega = \sqrt{V_0/\ell^2}$ to a Gaussian of width $\sigma = (\omega)^{-1/2} \cdot 2^{-1/2}$. Applied to the periodic kernel, the site-centred curvature receives additional image contributions, but these are sub-leading at the benchmark widths and ignored in Table 5; the resulting $\Omega_{max} \cdot \xi$ values are accurate to ~2% within the narrowed bands.

Appendix B: Consistency checks against [1] and [2]

B.1 Symmetric-limit check. Infinite well depth ($V_0 \to \infty$) forces localized modes and $M \to I$; the non-trivial spectrum degenerates onto the n = 0 eigenvalue. Reproduces the decoupled-site limit of [2] §3.

B.2 Minimal-model limit check. Large σ/d forces $M \to I$ approximately and $\kappa_2, \kappa_3 \to 0$. At σ/d = 2.0, $\kappa_2/\kappa_1 \approx -0.02$ and the LSQ fit reproduces eigenvalues to within 10^{-3}. The minimal-model form is an accurate description only in the extreme-wide-mode limit, outside the paper's admissible range.

B.3 (a, b) recovery comparison. At σ/d = 0.80, LSQ fit gives (a, b) = (0.899, 1.929) with prediction $\Omega_{max}^2 = 8.23$. [1] Table 1 gives (a, b) = (0.749, 0.375) under the $g_{geom} \cdot a$ identification with prediction 2.18. Actual Ω_{max}^2 from the direct spectrum readoff: 8.79. The LSQ fit is within 6% of the actual value; the [1] identification is off by a factor of 4.

B.4 Prior paper's Table 4 Lorentzian anomaly. The non-monotonic behaviour in [1] Table 4 Lorentzian column is reproduced under the $g_{geom} \cdot a$ identification applied to Lorentzian modes. Table 6 here (full projected calculation) gives monotonic behaviour for all three families. Both computations are internally consistent; they describe different projections of the same underlying spectrum.

References

[1] Taylor, K. *Microscopic closure dynamics and the K = 7 spectrum in the VERSF framework.* AIDA Institute, versf-eos.com. The present paper inherits the projection ansatz of §3, the explicit closure kernel and Wannier construction of §4, the projected-action structure of §5 (but not the §6.2 $g_{geom} \cdot a$ identification), the α-saturation ceiling interpretation of §8, and the open-problem prioritization of §11.

[2] Taylor, K. *Pair-Resolved Closure Spectrum and Commitment-Threshold Splitting in the VERSF Framework.* AIDA Institute, versf-eos.com. Source of the minimal-model construction of §10.10, the $\delta C = (3/32)\cdot r_{max}^2$ algebraic identity of §8.3, the α-sector cross-sector constraint $r_{max} \leq 0.15$ of §9.3, and the structural ceiling $\delta C \leq 2.1 \times 10^{-3}$ of Appendix E.5.

[3] Taylor, K. *From Necessary Facts to Physical Structure.* AIDA Institute, versf-eos.com. Programme context: VERSF as the uniquely constrained realization of minimal fact-producing physics.

[4] Taylor, K. *Fine-structure constant derivation in the VERSF framework.* AIDA Institute, versf-eos.com. Source of the 15 ppm $\alpha^{-1} \approx 137.034$ bound underlying the α-sector constraint in [2] §9.3.

[5] Taylor, K. *Two-Planck Principle.* AIDA Institute, versf-eos.com. Source of the $\lambda_{eff} = 3/4$ coefficient used in $m_\kappa^2 = \lambda_{eff}\cdot\xi^{-2}$ at §2.2.

[6] Taylor, K. *Fact Momentum: κ-field dynamics from irreversible commitment events.* AIDA Institute, versf-eos.com. Commitment-event framework underlying §12.1's bath-cutoff open problem and §12.2's kernel-derivation target.

[7] Bloch and Wannier functions — supporting §3.1's mode construction. [e.g. Ashcroft & Mermin, *Solid State Physics*, Ch. 10; Marzari & Vanderbilt, *Phys. Rev. B* 56, 12847 (1997).]

[8] Circulant matrices and discrete Fourier diagonalization — supporting §2.3. [e.g. Davis, *Circulant Matrices* (1979); Gray, *Toeplitz and Circulant Matrices: A Review* (2006).]

[9] Generalized eigenvalue problems with overlap matrices — supporting §3.3. [e.g. Golub & Van Loan, *Matrix Computations*, 4th ed., §8.7.]