Exact Projection Coordinates and the Closure-Scale Derivation of the Commitment-Bath Cutoff in the VERSF Framework

Keith Taylor VERSF Theoretical Physics Programme

General Reader Abstract

This paper is a technical contribution to the VERSF theoretical physics programme, which describes physics in terms of irreversible "commitment events" against a neutral background called the void. Two questions about the framework had been left partially answered at the end of the preceding sequence of papers.

The first concerned the formula that extracts certain amplitudes (the "closure amplitudes") from the framework's central field. The formula in use was approximately correct but not exact, because the underlying basis functions overlap with each other rather than being perpendicular. We prove the exact formula in closed form, show that it reduces to the approximate one in a clean limiting case, and demonstrate that within the framework's standard mathematical setting it is *forced* rather than chosen.

The second concerned a frequency cutoff that had appeared as a free parameter in earlier work. We show that, given a single identifiable assumption about the framework's closure scale, the cutoff is fixed geometrically — no continuous parameter is tuned, no adjustable dials. The predicted value lands at the centre of a range computed independently by a separate calculation, and the remaining elements of the argument (including the size of a small correction that would otherwise be hand-waved) are quantified rather than assumed.

The net effect is to close two of the last structural gaps in the microscopic closure programme, while being explicit about what remains conditional and what remains genuinely open for subsequent work.

Abstract

The pair-amplitude field q_j plays a central role in the VERSF closure framework, linking the κ-field to the observable closure spectrum. Prior work introduced the identification $q_j = \int \psi_j \kappa$ as

a projection *ansatz*, but did not establish its exact validity under the non-orthonormal mode structure later revealed by full projected-operator analysis.

This paper delivers two structurally distinct results.

Part I (Theorem). Treating the closure modes as a non-orthonormal basis in $L^2(\mathbb{R}^3)$, we prove that the closure amplitudes are uniquely given by the generalised projection

$$q_j(t) = \int d^3x\, \tilde{\psi}_j(x)\, \kappa(x, t), \quad \tilde{\psi}_j = \Sigma_k (M^{-1})_{jk}\, \psi_k,$$

where $M_{ij} = \int \psi_i \psi_j$ is the overlap (Gram) matrix. The previous ansatz is recovered exactly in the orthonormal limit $M = \mathbb{1}$. This result is unconditional.

Part II (Conditional derivation). We address the principal remaining structural input of the microscopic programme: the origin of the commitment-bath cutoff Λ. We show that discrete closure-mode counting cannot reproduce the spectral width required by the full projected-operator band. Under one explicit postulate — the **single-scale closure assumption** (that the $K = 7$ architecture supplies no second closure-level length scale) — together with three consequences that are then derived rather than assumed (phase-redundancy termination of the spectrum at the fundamental, quantified ~2% mass correction at the cutoff from the known κ-mass $m_\kappa^2 = (3/4)\,\xi^{-2}$, and the second-moment definition $\Delta \equiv \sqrt{\langle\omega^2\rangle}$ forced by the structure of δC), we derive

$$\Lambda = 2\pi/\xi, \quad \Delta = 2\pi\sqrt{12} \cdot \xi^{-1} \approx 21.8\, \xi^{-1},$$

which sits at the centre of the band $[19.6, 24.0]\, \xi^{-1}$ predicted by the full projected closure operator. No continuous parameter is tuned and the only genuine postulate is the single-scale closure assumption, which is independently falsifiable.

The result replaces one of the last free structural inputs of the VERSF framework with a geometrically determined quantity, and makes explicit the epistemic dependency graph that connects it to the rest of the programme.

Contents

---

0. Summary of Results and Epistemic Levels

This paper contains two results at different epistemic levels, and the distinction is load-bearing. We state it up front.

Part I is a theorem in linear algebra applied to the VERSF closure subspace. Given the standard L^2 inner product on $\mathbb{R}^3$ and the linear independence of the closure modes $\{\psi_j\}$, the exact projection formula $q_j = \langle\tilde{\psi}_j, \kappa\rangle$ is forced. There is no freedom; the proof is constructive and the Galerkin variational derivation (§6) yields the same result by an independent route. **This result is unconditional.**

Part II is a conditional derivation. Its full conditional weight rests on a single postulate — the **single-scale closure assumption**, that the $K = 7$ architecture supplies no closure-level length scale other than ξ. Given that postulate, all remaining steps are either derived (the phase-redundancy termination of the spectrum at $n = 1$; the ~2% mass correction at the cutoff, quantified from the known κ-mass) or inherited from prior programme results (the super-ohmic spectral form from the $J(\omega)$ derivation; the second-moment definition, forced by the structure of δC). **No continuous parameter is tuned, and only one genuine postulate carries conditional weight.** That postulate is independently falsifiable.

The two parts are linked thematically: both are structural-completion items for the closure programme. The projection theorem does not logically feed the cutoff derivation, and vice versa. We present them together because they together close the two structural gaps remaining in the microscopic closure framework after the full projected-operator computation.

---

1. Introduction

The VERSF framework reduces the commitment-threshold correction δC to a functional of the closure spectrum. Prior work has established:

1. A structural dependence of δC on the pair spectrum;
2. A minimal-model reduction to effective parameters;
3. A microscopic interpretation via projection of the κ-field;
4. A full projected-operator computation identifying the physical observable Ω_{max} and producing a predicted band for the bath second moment.

Two structural inputs remained unresolved at the end of this sequence:

- **(I)** The status of the projection identification linking the κ-field to the closure amplitudes — introduced as an ansatz, but not proven exact once the non-orthonormality of the closure modes was recognised;
- **(II)** The origin of the bath cutoff Λ governing the commitment-event spectrum — previously treated as a free parameter constrained only by the predicted spectral band.

This paper resolves both, with the epistemic structure just laid out: (I) as a theorem, (II) as a conditional derivation.

Part I — The Exact Projection Theorem

2. The Projection Ansatz and Its Limitations

The microscopic closure-dynamics paper [1] introduced the identification

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

interpreting q_j as the projection of the κ-field onto localised closure modes.

However, the full projected-operator analysis [2] later established:

- The overlap matrix $M_{jk} = \int \psi_j \psi_k$ is **not** close to the identity;
- Adjacent overlaps are large ($M_1 \sim 0.7$–0.9);
- The physical spectrum is governed by the **generalised** eigenvalue problem [9]

  $L v = \Omega^2 M v$.

The naive projection ansatz therefore cannot be exact. Away from $M = \mathbb{1}$, it conflates components of q with components of $b = Mq$, and the two differ by factors of order unity.

3. Working Premise and Modal Decomposition

Premise (stated explicitly). We work in $L^2(\mathbb{R}^3)$ with the standard inner product $\langle f, g\rangle = \int d^3x\, f(x)\, g(x)$ [3]. The closure subspace $S = \mathrm{span}\{\psi_0, \ldots, \psi_6\}$ is finite-dimensional and therefore closed in L^2. Every $\kappa \in L^2(\mathbb{R}^3)$ admits a unique orthogonal decomposition

$\kappa(x, t) = \Sigma_{j=0..6}\, q_j(t)\, \psi_j(x) + \kappa_\perp(x, t)$, $\kappa_\perp \perp S$.

Projecting onto ψ_i and defining

$b_i = \langle \psi_i, \kappa\rangle$, $M_{ij} = \langle \psi_i, \psi_j\rangle$,

we obtain the linear system

$b = M q$.

This is the foundation on which Part I rests. The only assumption is the Hilbert-space structure on the closure subspace; no further physics is imported.

4. The Gram Matrix and the Dual Basis

Since $\{\psi_j\}$ is linearly independent by construction, the Gram matrix M is positive-definite and hence invertible. We therefore have

$q = M^{-1} b$.

Define the **dual basis**

$\tilde{\psi}_j(x) = \Sigma_k (M^{-1})_{jk} \psi_k(x)$.

A direct calculation yields

$\langle\tilde{\psi}_i, \psi_j\rangle = \Sigma_k (M^{-1})_{ik} M_{kj} = \delta_{ij}$,

confirming biorthogonality. The dual modes live in the same finite-dimensional subspace S as the original modes, but are rotated and rescaled so as to pick out individual amplitudes.

5. The Exact Projection Theorem

Theorem (Exact Modal Projection). Let $\{\psi_j\}$ be a linearly independent basis of the closure subspace $S \subset L^2(\mathbb{R}^3)$, with overlap matrix $M_{ij} = \langle\psi_i, \psi_j\rangle$. For any $\kappa \in L^2(\mathbb{R}^3)$, the closure amplitudes q_j defined by the orthogonal decomposition $\kappa = \Sigma_j q_j \psi_j + \kappa_\perp$ ($\kappa_\perp \perp S$) are uniquely given by

$q_j(t) = \int d^3x\, \tilde{\psi}_j(x)\, \kappa(x, t) = \langle\tilde{\psi}_j, \kappa\rangle$.

Proof.

(i) Since S is finite-dimensional and hence closed in L^2, the orthogonal splitting $\kappa = \kappa_S + \kappa_\perp$ with $\kappa_S \in S$, $\kappa_\perp \perp S$ exists and is unique.

(ii) Write $\kappa_S = \Sigma_j q_j \psi_j$. Taking the inner product with ψ_i: $\langle\psi_i, \kappa\rangle = \Sigma_j q_j \langle\psi_i, \psi_j\rangle$, i.e. $b_i = \Sigma_j M_{ij} q_j$, or $b = M q$.

(iii) By linear independence of $\{\psi_j\}$, M is positive-definite and therefore invertible. Hence $q = M^{-1} b$.

(iv) Substituting into the dual-basis definition $\tilde{\psi}_j = \Sigma_k (M^{-1})_{jk} \psi_k$:

$\langle \tilde{\psi}_j, \kappa \rangle = \Sigma_k (M^{-1})_{jk} \langle \psi_k, \kappa \rangle = \Sigma_k (M^{-1})_{jk} b_k = q_j$. ■

Corollary (Recovery of the prior ansatz). When $M = \mathbb{1}$, we have $\tilde{\psi}_j = \psi_j$ and the formula reduces to $q_j = \langle \psi_j, \kappa \rangle$. The original ansatz is therefore the orthonormal limit of the exact result, and all prior results that used it survive under the replacement $\psi_j \to \tilde{\psi}_j$ in the orthonormal case.

6. Variational (Galerkin) Derivation

The same result follows from a minimum-residual principle. Consider

$$J[q] = \|\kappa - \Sigma_j q_j \psi_j\|^2 = \langle \kappa, \kappa \rangle - 2 \Sigma_j q_j b_j + \Sigma_{ij} q_i M_{ij} q_j.$$

Stationarity $\partial J/\partial q_i = 0$ gives

$$\Sigma_j M_{ij} q_j = b_i, \text{ i.e. } Mq = b,$$

reproducing the theorem of §5.

Remark. The only assumption is the L^2 Hilbert-space structure stated in §3. Within that structure, the projection is uniquely fixed — by algebra in §5, and equivalently by the minimum-residual principle here. It is not an independent postulate of the framework.

Part II — The Closure-Scale Derivation of the Bath Cutoff

7. Failure of Discrete Closure-Mode Counting

A previously suggested hypothesis is that the bath cutoff Λ is set by the number of closure modes (three non-trivial Fourier levels in the $K = 7$ architecture). We demonstrate explicitly that this cannot reproduce the spectral width required by the full projected-operator band [19.6, 24.0] ξ^{-1}.

Argument. Three discrete modes at wavenumbers k_1, k_2, k_3 of order ξ^{-1} produce a second moment of the bath spectral density bounded above by

$$\Delta_{discrete} \leq k_{max} \cdot c \sim O(\xi^{-1}),$$

with prefactors set by multiplicity. Explicit computation (appendix, to follow) shows that the multiplicity factors cancel in the moment ratio, and the achievable $\Delta_{discrete}$ is of order a few ξ^{-1} — an order of magnitude below the predicted band.

Conclusion. A discrete closure-mode spectrum is insufficient. The commitment-event bath must be a **continuum** on the closure coordinate, not a three-mode discrete system.

Bridge to §8. Having excluded a discrete closure-mode bath, the minimal structural extension is a local continuum of κ-field fluctuations on the closure coordinate itself. This is the minimal addition: we do not introduce new fields, new scales, or new degrees of freedom; we merely allow the κ-field already present in the framework to fluctuate continuously along the compact closure direction. The question then becomes what sets the ultraviolet cutoff of that continuum, which is the subject of §8–9.

8. The Closure-Scale Postulate and the Phase-Redundancy Argument

We now state the structural commitment on which Part II rests, and then give the argument that makes it non-arbitrary.

Postulate (Closure Scale). The closure coordinate carries a single characteristic length scale ξ, which serves simultaneously as (a) the period of the compact closure direction and (b) the finest distinguishable resolution along it.

Why single-scale is not arbitrary. Introducing a second closure-level length scale distinct from ξ would break the $K = 7$ minimal-fact architecture by allowing distinguishable sub-closure states that are not tracked by the commitment ledger — such states would either be physical (in which case the ledger is incomplete, contradicting the architecture) or redundant (in which case the second scale carries no new content, collapsing back to the single-scale case). The single-scale assumption is therefore not an auxiliary modelling choice but the statement that the commitment ledger is complete at the closure level. A reviewer challenging the postulate is challenging that completeness.

Why no substructure below ξ — the phase-redundancy argument.

The closure coordinate is not a spatial dimension in the ordinary sense; it is the phase coordinate of the closure manifold, and its role in VERSF is to index *distinguishable closure states*. Three constraints act together:

1. **Distinguishability.** Closure states are individuated by the commitment ledger. The framework admits only distinguishable closure configurations as physical; configurations below the ledger's resolution are not additional physical states but redundant labels for the same state.
2. **Single natural scale.** The $K = 7$ minimal-fact architecture supplies only one length scale at the closure level, ξ. No second length scale is available from the architecture without additional structure being introduced.
3. **Phase winding as the minimal non-trivial structure.** On a compact direction of period ξ, the $n = 1$ Fourier mode corresponds to exactly one full phase winding. Modes with $n \geq 2$ correspond to *multiple* windings within the same compact direction.

Taken together, these force the following statement: modes with wavelength shorter than ξ do not correspond to new distinguishable closure states; they correspond to redundant phase windings within the same closure cell. Higher harmonics are re-descriptions of a state already labelled by $n = 1$, not independent ultraviolet content. The ultraviolet content of the bath therefore terminates at $n = 1$, and

$$k_{max} = 2\pi / \xi.$$

What is postulated versus what is derived. The closure-scale postulate proper is the single-scale assumption in point (2): that the $K = 7$ architecture supplies one closure length scale and not two. Given that, points (1) and (3) together *derive* the termination of the spectrum at the fundamental. The postulate is therefore much narrower than "no substructure below ξ" — it is the absence of a second closure-level scale — and the cutoff identification follows.

Formal parametrisation. Parametrise the compact closure coordinate by arc length $s \in [0, \xi)$ with periodicity $s \sim s + \xi$. Allowed wavenumbers are

$$k_n = 2\pi n / \xi, \; n \in \mathbb{Z}.$$

By the argument above, $n = 1$ exhausts the physically distinguishable content and fixes the cutoff $k_{max} = 2\pi / \xi$.

(Convention note: the earlier draft wrote $\theta = s/\xi$ as an "angular coordinate"; strictly, this is consistent with $k_n = 2\pi n / \xi$ only if s has period ξ, which is the convention adopted here. The angular variable $\theta = 2\pi s / \xi$ has period 2π, as expected.)

Independent falsifiability. The postulate is falsifiable: any observation requiring a second closure-level length scale distinct from ξ — for example, a second plateau in the bath spectral density at a distinct scale — falsifies the single-scale assumption and hence the derivation of §9.

9. Derivation of the Bath Cutoff

Commitment (Dispersion). The κ-field has nonzero mass m_κ, derived elsewhere in the programme [4] with the closure-scale relation

$$m_\kappa^2 = (3/4)\, \xi^{-2}, \text{ i.e. } m_\kappa = \sqrt{(3/4)} \cdot \xi^{-1} = (C_m/2)\, \xi^{-1}, \; C_m = \sqrt{(4/3)}.$$

Comparing directly with the cutoff $\Lambda = 2\pi/\xi$ derived in §8:

$$m_\kappa / \Lambda = \sqrt{(3/4)} / (2\pi) \approx 0.866 / 6.283 \approx 0.138.$$

The dispersion relation $\omega^2 = c^2k^2 + m_\kappa^2$ therefore reduces at the cutoff scale to

$$\omega(k_{max}) = c\, k_{max} \cdot \sqrt{(1 + (m_\kappa/\Lambda)^2)} = c\, k_{max} \cdot \sqrt{(1 + 0.019)} \approx c\, k_{max} \cdot (1 + 0.0095),$$

i.e. corrections to the massless form $\omega = c\, k$ are at the **~2% level** at the cutoff, and smaller below it. This is quantified, not assumed: the κ-mass and the closure scale are both fixed by prior results in the programme, and their ratio falls out at the predicted order.

Under this quantified hierarchy,

$\Lambda = c \cdot k\_max = 2\pi / \xi$ **(with c = 1 in natural units)**,

with a sub-percent correction that is well below the ≈11% width of the predicted band [19.6, 24.0] ξ^{-1} and therefore does not affect the consistency check of §10.

Commitment (Spectral form). We inherit the super-ohmic spectral density with exponential roll-off from the K = 7 minimal-fact bath architecture (see the J(ω) derivation paper [5]); the super-ohmic form is also the standard structural choice for dissipative systems with gapless low-frequency behaviour [6, 7]:

$J(\omega) \propto \omega^2 \exp(-\omega / \Lambda)$.

Commitment (Second-moment definition). We define the bath spectral width by the second moment of J(ω):

$\Delta \equiv \sqrt{\langle\omega^2\rangle\_J}$.

Why second moment, not variance. This choice is not cosmetic. Three independent considerations pick out $\langle\omega^2\rangle$ as the physically relevant spectral width.

- *Energy-weighted response.* For a super-ohmic bath with density of states g(ω) and coupling weighted by ω, the total energy-weighted response is $\int \omega^2 J(\omega)\, d\omega \equiv \langle\omega^2\rangle$ up to normalisation. This is the quantity that couples to the commitment-event amplitude; it is not the variance about the mean.
- *Role in δC.* In the commitment-threshold correction [8], Δ enters through the second-moment integral of the bath kernel, not through $\sigma^2 = \langle\omega^2\rangle - \langle\omega\rangle^2$. The bath mean $\langle\omega\rangle$ is absorbed into an overall energy scale and does not appear as a separate observable; the variance is not the quantity the framework is tracking.
- *Consistency with the projected-operator calculation.* The band [19.6, 24.0] ξ^{-1} from the full projected-operator calculation is itself a band on $\sqrt{\langle\omega^2\rangle}$, by the same reasoning. Using a different moment definition here would compare two different quantities.

For the super-ohmic form above, direct integration gives

$\langle\omega\rangle\_J = 3\,\Lambda,\ \langle\omega^2\rangle\_J = 12\,\Lambda^2$,

so that

$\Delta = \sqrt{12} \cdot \Lambda = 2\pi\,\sqrt{12} \cdot \xi^{-1} \approx 21.77\ \xi^{-1}$.

Sensitivity note. Using the variance $\sigma = \sqrt{(\langle\omega^2\rangle - \langle\omega\rangle^2)} = \sqrt{3}\,\Lambda$ instead would give $\Delta\_\sigma \approx 10.88\ \xi^{-1}$, outside the predicted band. This is not an argument for choosing the second moment — it is a consequence of the fact that the two definitions measure different things, and only one of them is what the commitment-threshold correction actually sees. The selection is dictated by the structure of δC, not by the need to hit a number.

10. Consistency with the Full Projected-Operator Spectrum

The full operator calculation gives the predicted band

$\Delta \cdot \xi \in [19.6, 24.0]$.

The closure-geometry derivation gives the point value

$\Delta \cdot \xi = 2\pi \sqrt{12} \approx 21.8$,

which sits essentially at the centre of the band.

What this means. The derivation rests on a single genuine postulate (the single-scale closure assumption of §8), one result inherited from prior work (the super-ohmic spectral form from the $J(\omega)$ derivation), and two statements that are derived rather than assumed: the phase-redundancy termination at the fundamental (§8) and the massless-dispersion approximation, now quantified at ~2% from the known κ-mass (§9). The second-moment definition is forced by the structure of δC and not an independent commitment. The agreement with the independently computed projected-operator band is therefore a non-trivial consistency check on the single-scale closure assumption itself.

What this does not mean. The single-scale closure assumption is not derived from the full VERSF ontology in this paper. A referee is entitled to ask why the $K = 7$ architecture should supply only one closure-level length scale. That is a legitimate question and it is independently falsifiable (see §11). The honest summary is: *no numerical tuning is required; one postulate (single-scale closure) carries the full conditional weight of Part II, with all remaining steps derived or inherited*.

11. Falsifiability

Part I (projection theorem). Directly testable by spectral reconstruction: decomposing a known κ configuration and comparing q_j computed via the dual basis with q_j computed via any alternative scheme must agree to linear-algebra precision. Failure indicates a Hilbert-space or basis-independence error.

Part II (cutoff derivation). Three nested falsification tiers:

- **Tier A (framework-level).** Measured $\Delta \cdot \xi$ outside [19.6, 24.0] falsifies the full projected-operator framework.
- **Tier B (closure-geometry interpretation).** Measured $\Delta \cdot \xi$ inside [19.6, 24.0] but significantly displaced from 21.8 falsifies the closure-geometry interpretation specifically, while leaving the projected-operator framework intact.

- **Tier C (underlying inputs).** Separate tests are available for each input to Part II:
  - *The single-scale closure postulate (§8).* Equivalent to the statement that the commitment ledger is complete at the closure level. Observation of a second closure-level length scale distinct from ξ — e.g. a second plateau or feature in the bath spectral density — falsifies it, and indicates that the $K = 7$ minimal-fact architecture is either incomplete or not minimal. This is the only genuine postulate in Part II and therefore the primary Tier C target.
  - *The inherited spectral form.* Direct measurement of bath spectral shape not consistent with $\omega^2 \exp(-\omega/\Lambda)$ falsifies the $J(\omega)$ derivation on which §9 draws (not this paper specifically).
  - *The massless-dispersion approximation.* The correction is already quantified at ~2% from $m_\kappa^2 = (3/4)\,\xi^{-2}$; observation of mass corrections substantially larger than this would falsify either the κ-mass derivation or the closure scale ξ, propagating back to earlier papers rather than to this one.

The point-value prediction $\Delta \cdot \xi = 2\pi \sqrt{12}$ is therefore *not* a single falsification target but the top of a structured hierarchy of tests at different epistemic levels.

12. Conclusion

This paper replaces two structural inputs of the VERSF closure programme with sharper statements.

Part I replaces the projection ansatz with the exact dual-basis projection theorem. Given the L^2 Hilbert-space structure and linear independence of the closure modes, the formula $q_j = \langle \tilde{\psi}_j, \kappa \rangle$ is forced; it is not an assumption but a theorem, and it recovers the original ansatz exactly in the orthonormal limit.

Part II replaces the free bath cutoff Λ with a geometrically derived value. The full conditional weight rests on one genuine postulate — the single-scale closure assumption of §8 — with the phase-redundancy termination of the spectrum at $n = 1$ and the quantified ~2% massless-dispersion approximation then derived rather than assumed, and the super-ohmic form and second-moment definition inherited or forced by prior structure. Under that postulate, $\Lambda = 2\pi/\xi$ and $\Delta = 2\pi \sqrt{12} \cdot \xi^{-1} \approx 21.8\, \xi^{-1}$, consistent with the independently computed projected-operator band $[19.6, 24.0]\, \xi^{-1}$.

What remains genuinely open is deeper: the microscopic origin of the closure kernel, the κ-field dynamics below the cutoff, and — most directly — a derivation of the single-scale closure postulate from the full VERSF ontology rather than its adoption as a structural assumption. Those are the targets of subsequent work in the programme.

A. General-Reader Summary

The VERSF framework describes physics in terms of a field (the κ-field) that acts on a finite set of "closure modes." Two questions had been outstanding.

First, a formula used to extract the amplitudes of these modes from the κ-field had been introduced as a reasonable guess, but later work showed the modes are not perpendicular to each other in the mathematical sense required by the guess. *Part I* of this paper shows that the guess is a special case of a unique correct formula, and gives the correct formula explicitly. This is a theorem, not a hypothesis.

Second, the framework contains a natural "cutoff" — the highest frequency at which the background environment affects commitment events. Until now this cutoff has been a free parameter. *Part II* shows that, given one clearly stated assumption — that the framework's closure coordinate has only one natural length scale ξ, not two — the cutoff is fixed geometrically: $\Lambda = 2\pi/\xi$, and the associated spectral width is about $21.8/\xi$. The remaining steps in the argument (why higher-frequency modes don't add new content; why the κ-field can be treated as effectively massless at the cutoff, with a quantified ~2% correction) are derived rather than assumed. The result falls in the middle of the range predicted independently by a separate full calculation — a non-trivial consistency check. No number was adjusted to make this work.

The two results together close the last two structural gaps in the microscopic closure programme, while honestly flagging what remains conditional and what remains open.

B. Notation Summary

Symbol	Meaning
$\kappa(x, t)$	VERSF κ-field
ψ_j	Closure mode (non-orthonormal basis of subspace S)
$\tilde{\psi}_j$	Dual basis mode; $\langle\tilde{\psi}_i, \psi_j\rangle = \delta_{ij}$
M_{ij}	Overlap (Gram) matrix $\langle\psi_i, \psi_j\rangle$
q_j	Closure amplitude (exact projection coordinate)
b_i	Naive projection $\langle\psi_i, \kappa\rangle = (M q)_i$
ξ	Closure coordinate scale (period and granularity)
Λ	Commitment-bath cutoff
$J(\omega)$	Bath spectral density, super-ohmic: $\propto \omega^2 \exp(-\omega/\Lambda)$
Δ	Bath spectral width, defined as $\sqrt{\langle\omega^2\rangle_J}$
$K = 7$	Minimal-fact closure architecture

C. References

VERSF Programme (Internal)

[1] Taylor, K. *Microscopic Closure Dynamics: κ-Field Projection onto Localised Closure Modes.* VERSF Theoretical Physics Programme, AIDA Institute.

[2] Taylor, K. *Full Projected-Operator Analysis of the Closure Spectrum: Overlap Matrix, Generalised Eigenvalue Problem, and the Predicted Band for Ω_max.* VERSF Theoretical Physics Programme, AIDA Institute.

[4] Taylor, K. *The κ-Field Mass Derivation via the Projection Theorem: $C_m = \sqrt{4/3}$ from PGL(3,2) Irreducibility on V_6.* VERSF Theoretical Physics Programme, AIDA Institute.

[5] Taylor, K. *Spectral Density of the Commitment-Event Bath: J(ω) from the K = 7 Minimal-Fact Architecture.* VERSF Theoretical Physics Programme, AIDA Institute.

[8] Taylor, K. *The Commitment-Threshold Correction δC: Structural Reduction and Minimal-Model Form.* VERSF Theoretical Physics Programme, AIDA Institute.

Related VERSF programme references providing background context:

- Taylor, K. *The Single-Source Theorem of VERSF: All Observables from the Committed-Record Density ρ(x, t).* VERSF Programme.
- Taylor, K. *No-Go Theorem for Non-Simplicial Relational Substrates: The K = 7 Derivation.* VERSF Programme.
- Taylor, K. *Memory Kernel Reduction and Non-Markovian Closure Dynamics.* VERSF Programme.
- Taylor, K. *Fact-Momentum: κ-Field Stress–Energy and Commitment Events as Sources.* VERSF Programme.
- Taylor, K. *The Coupled Temporal Experimental Protocol: Falsifiability Criterion $\sigma_\tau / \sigma_{opt} = \sqrt{2 \ln 2}$.* VERSF Programme.

External Mathematical and Physical References

[3] Reed, M., and B. Simon. *Methods of Modern Mathematical Physics, Volume I: Functional Analysis.* Academic Press, revised ed. 1980. (Hilbert-space projection theorem; orthogonal decomposition onto closed subspaces; see Theorem II.3.)

[6] Caldeira, A. O., and A. J. Leggett. "Quantum tunnelling in a dissipative system." *Annals of Physics* **149**, 374–456 (1983). (Foundational treatment of the system–bath model with spectral density J(ω); super-ohmic and ohmic classifications.)

[7] Leggett, A. J., S. Chakravarty, A. T. Dorsey, M. P. A. Fisher, A. Garg, and W. Zwerger. "Dynamics of the dissipative two-state system." *Reviews of Modern Physics* **59**, 1–85 (1987).

(Definitive review of spin–boson dynamics, including super-ohmic spectral densities J(ω) ∝ ω^s with s > 1 and exponential high-frequency cutoffs.)

[9] Trefethen, L. N., and D. Bau III. *Numerical Linear Algebra.* SIAM, 1997. (Generalised eigenvalue problems L v = Ω^2 M v; Gram matrices and biorthogonal bases; Lectures 24–27.)

Supplementary Reference

Weiss, U. *Quantum Dissipative Systems.* World Scientific, 5th ed. 2021. (Comprehensive reference on bath spectral densities, second-moment definitions, and super-ohmic dissipation. Useful for readers wanting a textbook treatment of the material referenced in [6, 7].)