Pair-Resolved Closure Spectrum and Commitment-Threshold Splitting in the VERSF Framework

Keith Taylor · versf-eos.com

Plain-language summary

Why this matters

A recurring frustration in modern physics is the number of *free parameters* our best theories contain — numbers like the mass of the electron or the strength of gravity that have to be measured by experiment rather than derived from deeper principles. The Standard Model of particle physics, for all its spectacular accuracy, has about nineteen such numbers. Every attempt over the last half-century to reduce that count has either stalled or failed outright. This is not a cosmetic worry: a theory with many free parameters is a theory that is, in a meaningful sense, *not yet complete*.

The VERSF research programme — of which this paper is one piece — takes the ambitious position that reality's fundamental constants shouldn't be free at all. It proposes that everything in physics, including the particles, forces, and constants we see, emerges from a single founding premise: that the universe must produce stable, locked-in facts. Each such "commitment event" costs a specific amount of energy, fixed by a universal constant that earlier VERSF work placed very close to a simple round number: **three-eighths of a natural unit**.

If VERSF is right, three-eighths isn't a measured parameter; it's a consequence of the framework. But three-eighths would be *exactly* the answer only if the universe were perfectly uniform at its smallest scale. Since no real universe is perfectly uniform, the true value should be three-eighths plus a small correction. Until now, that correction was a residual gap in the VERSF programme — a loose end that could have been almost anything.

What this paper does

This paper closes that gap at the level of *structure*. It doesn't compute the correction from absolute first principles — that remains the subject of a follow-up paper — but it shows that the correction cannot be anything at all. It is tightly constrained by several independent requirements, and the constraints turn out to be strong enough to reduce what looked like fourteen independent unknowns to just two or three.

What the calculation shows

- The correction is **never negative** in any physically sensible situation, and is exactly zero only if the universe's smallest-scale structure is perfectly symmetric. Any real-world departure from symmetry must produce a positive correction.
- The correction is **small** — under reasonable assumptions it is bounded at the sub-percent level, with the tightest form of the analysis giving an upper bound of about 0.2%.
- The fourteen independent numbers one might naively expect (seven for each of two structural features of the underlying manifold) collapse to just **two or three effective parameters**, thanks to an underlying cyclic symmetry. This is a dramatic reduction in the theory's residual freedom.
- Those same two or three parameters simultaneously control **a small correction to the strength of electromagnetism** and **a prediction about how the energy of empty space is distributed across the cosmos**. They are not independent quantities to be fitted sector by sector — they are consequences of one shared underlying structure. Measure one, and the others are predicted.
- Combined cross-sector consistency requires that the **bath frequency scale** — roughly speaking, the rate at which commitment events occur — must exceed the fundamental closure scale by at least an order of magnitude. This is itself a nontrivial testable prediction of the combined framework.

Why this is significant

First, it **reduces a theory's free-parameter count** — the currency in which theoretical physics measures progress. A framework that reduces fourteen unknowns to three is a framework that has become more predictive, more constrained, and therefore more falsifiable.

Second, it **produces a sharp, falsifiable prediction**. If a future measurement ever finds the commitment-barrier constant to be substantially larger than three-eighths — say, at the percent level or more — one of the structural assumptions in this paper must be wrong. Most modern theories do not produce predictions this specific.

Third, it **confirms that the three-eighths baseline is robust** rather than accidental. A worry with clean theoretical numbers is that they collapse under refinement. Here, three-eighths survives the inclusion of symmetry-breaking effects as the exact baseline, with corrections sitting on top of it rather than replacing it.

Fourth, it **shifts the remaining open question to a sharper target**. Instead of asking "what is the correction and how does it arise?" the question becomes "what are the three effective parameters of the minimal closure model?" That is a narrower, better-posed problem — which is what genuine progress in theoretical physics typically looks like.

What the paper does *not* do

It does not calculate the two or three parameters from VERSF's most basic equations. That is the subject of the next paper. What this paper does is take a previously undefined residual freedom and reduce it to a precisely located target — one a follow-up calculation can now aim at directly.

In short

A loose end of the VERSF programme — a small correction to a fundamental constant — has been pinned down to within a handful of numbers, with a concrete upper bound, a specific set of cross-sector consequences, and a sharper falsification criterion than the framework previously possessed. Those handful of numbers are now the next well-defined question to answer.

What this paper proves. This paper proves that the correction to the VERSF commitment-threshold coefficient is not an arbitrary residual term. Within the linear pair-resolved response formalism, it is fixed by the closure-pair spectrum through a constrained functional dependence: structurally by the theorems of Part I, and explicitly in a controlled near-symmetric regime by the formulas of Part II. Part IV then shows that the pair spectrum itself lies in a constrained Z_7 Fourier subspace, with splittings and participation anomalies correlated through shared Fourier structure. Under three further minimal-model assumptions — nearest-neighbour coupling, vacuum stability, and uniform leading bath coupling — the spectrum collapses further to just two effective parameters with a specific degeneracy structure, and the splitting-dominated sub-regime is selected automatically. The paper does not yet derive those two parameters microscopically, but it shows that once they are known, the barrier correction and its downstream consequences are determined without further freedom.

Abstract (technical)

The commitment barrier of VERSF has the structural form Φ_c = C* · ℏc/ξ, with C* = 3/8 + δC, where 3/8 is the symmetric-onset value and δC is a pair-resolved threshold splitting inherited from the K = 7 closure spectrum. **This paper reduces the threshold-splitting problem to an explicit functional dependence on the pair spectrum within the response formalism: it proves the structural properties of the correction for any admissible linear-onset spectrum, and computes that correction explicitly in a controlled near-symmetric regime.**

Part I (formal-structural). Four theorems on the pair-threshold spectrum — existence and uniqueness of pair thresholds, exact symmetric closure (δC = 0 iff spectrum is symmetric), conditional positivity with tight non-linear boundedness, and scale invariance. These are regime-independent within the linear-onset response formalism of §§2–3; they do not assume narrow-band forcing or near-degenerate splittings, but they do rely on the response formalism itself.

Part II (regime-conditional). Explicit closed-form computation of δC in the narrow-band, near-degenerate regime. Under three controlled assumptions we obtain

$\delta K_j / K_0 \approx \beta_j - r_j^2/8$, $r_j = \Omega_j / \Delta$,

$\delta_j \approx (3/32)\, r_j^2 - (3/8)\, \beta_j$,

$\delta C = \max_j \delta_j$.

In the splitting-dominated limit this collapses to the one-parameter prediction $\boldsymbol{\delta C = (3/32)\, r_{max}^2}$, with a falsifiability criterion tying it to the independently measurable ratio of pair splitting to bath bandwidth.

Part III. Consequences propagate to Φ_c ($\delta\Phi\_c/\Phi_0 = r_{max}^2/4$), Λ-spread, and α-corrections via a single shared input. Pair-splitting corrections to α enter at $O(r_{max}^4/64)$, preserving the symmetric-limit $\alpha^{-1} \approx 137.034$ result for $r_{max} \lesssim 0.15$ and supplying a cross-sector constraint on admissible closure-spectrum splitting.

Part IV (symmetry-level and minimal-model constraints on the spectrum). The pair spectrum $\{\Omega_j, \beta_j\}$ is not an unconstrained seven-parameter input. Cyclic Z_7 symmetry of the symmetric closure manifold restricts admissible perturbations to a six-dimensional Fourier subspace; stability of the symmetric configuration further ties pair splittings to eigenvalues of the linearized closure operator; and the same Fourier structure is inherited by the participation coefficients, correlating $\{\Omega_j\}$ and $\{\beta_j\}$. Under three further minimal-model assumptions — nearest-neighbour coupling, vacuum stability, and uniform leading bath coupling — the closure spectrum is uniquely determined up to two effective parameters (a, b) as $\Omega_n^2 = \xi^{-2}[a + 4b \sin^2(\pi n/7)]$ with three doubly-degenerate levels (n = 1,2,3 paired with n = 6,5,4), and participation anomalies vanish at leading order. The splitting-dominated sub-regime is selected automatically. The minimal-model prediction $\delta C = (3/32d^2)(a + 3.8019\, b)$, combined with the α-sector cross-sector constraint, further constrains the ratio of bath bandwidth to closure scale. The spectral problem thereby reduces to two effective parameters within the minimal model — the precisely located target of the next paper.

An explicit illustrative computation (Appendix D) confirms numerical stability: a representative 2% anisotropic seven-pair spectrum — recognizable as approximately the n = 1 Fourier harmonic of the Part IV decomposition — produces $\delta C \approx 8.4 \times 10^{-3}$ and exactly saturates the Part I tight bound. Multi-sector consistency with the α-bound identifies the toy spectrum as participation-dominated rather than splitting-dominated, demonstrating how cross-sector constraints sharpen interpretation. The remaining open problem is upstream — the microscopic derivation of the Fourier mode amplitudes — not downstream.

Abstract (general reader)

Reality in VERSF is held together by seven "closure pairs" at its finest grain. A universal constant — close to 3/8 of a natural unit — sets the energy needed to lock in a fact. This paper answers: what happens if the seven pairs are not perfectly identical?

Three kinds of answer are given. The first kind proves structural statements true for any pair spectrum: the correction is non-negative in the physically generic case, tightly bounded, and vanishes exactly when the pairs are symmetric. The second kind goes further — under controlled assumptions about the spectrum, we compute the correction in closed form. It scales as the square of a single measurable ratio: pair-splitting divided by bath-bandwidth. The third kind shows that the seven pairs are not truly seven free numbers: the cyclic symmetry of the closure manifold ties them together, and under a minimal locality assumption the entire spectrum collapses to just two effective parameters with a specific degeneracy pattern. An illustrative calculation at the end of the paper turns the crank on a specific seven-pair example and produces a correction of about 2%. One measurement fixes everything downstream.

Contents

Appendices

---

1. Introduction

1.1 The residual problem

Prior work established the commitment barrier in the form

$\Phi\_c = C^* \cdot \hbar c/\xi$, $C^* = 3/8 + \delta C$,

with 3/8 derived in the symmetric-onset limit of the K = 7 closure sector. The coefficient δC measures the correction induced by pair-level threshold asymmetry in the closure spectrum, and was reduced to a pair-resolved closure problem without explicit solution. The present paper reduces that problem to an explicit functional dependence on the pair spectrum within the response formalism.

1.2 Three-track strategy, and a demonstration

The problem admits three distinct modes of attack, and all three are needed.

- **Structural** (Part I): What can be proven about δC using only the formal definitions, without committing to a regime of the closure spectrum? Results here are universal over the admissible response class but shape-level: they fix sign conditions, bounds, and closure conditions.
- **Computational** (Part II): What is the value of δC in a tractable regime where explicit formulae are accessible? Results here are sharp but conditional on the regime holding.
- **Symmetry-level** (Part IV): Given the cyclic symmetry of the K = 7 closure manifold, what is the dimensional and functional structure of the admissible pair spectrum $\{\Omega_j, \beta_j\}$ that feeds Parts I–II? Results here narrow the spectrum from seven free parameters to a small Fourier-mode family, without attempting the microscopic derivation of the mode amplitudes themselves.

Part I answers "what shape must δC have?" Part II answers "what number does δC equal, here?" Part IV answers "how structured is the upstream spectrum that both parts depend on?" None alone suffices: without Part I, Part II has no universality guarantee; without Part II, Part I proves properties of an object whose magnitude is unknown; without Part IV, the upstream inputs look arbitrary. Together they bracket the answer from above (Part I bounds), from below (Part II explicit values), and from upstream (Part IV spectral constraints).

Appendix D supplies a **demonstration of the pipeline**: a specific seven-pair spectrum — recognizable as approximately the n = 1 Fourier harmonic of Part IV's decomposition — is plugged into the formalism end-to-end, producing a definite numerical δC. This is not a first-principles prediction, but it confirms that the formalism is computationally well-behaved and that Parts I, II, and IV are mutually consistent on a concrete case.

The key contribution of this paper is therefore not the first-principles derivation of the microscopic pair spectrum itself, but the proof that once the pair spectrum is specified within the Part IV symmetry-admissible subspace, the induced correction to the commitment threshold is no longer arbitrary. It is structurally constrained by the theorems of Part I, explicitly computable in the controlled regime of Part II, and sourced from a constrained Fourier-mode subspace per Part IV. The paper thus converts δC from a residual free correction into a well-defined functional of a small number of spectral mode amplitudes. This reduces the threshold-splitting problem to an explicit functional dependence on the pair spectrum within the response formalism, and further reduces the pair spectrum itself to a constrained Fourier-mode family, even though the microscopic mode amplitudes remain to be derived.

1.3 What is proven, what is assumed

Formal-structural scope of Part I. The results of Part I are regime-independent within the linear-onset response formalism defined in §§2–3. They do not assume narrow-band forcing, near-degenerate pair splittings, or any explicit spectral model, but they do rely on the monotonic linear scaling $R_j(C) = C \cdot K_j$, the dissipative definition of pair response, and the positivity of the resulting pair weights. In this precise sense, the Part I theorems are **universal over the admissible response class studied here, though not independent of the response formalism itself**. Their role is to constrain the allowed shape of the threshold correction δC prior to any regime-specific computation.

Part II scope. Part II results are conditional on three controlled assumptions (A1)–(A3) stated in §6, covering narrow-band bath forcing, near-degenerate pair spectrum, and smooth participation anomaly. Violations of these assumptions invalidate Part II but leave Part I intact.

Part IV scope. Part IV results follow from the cyclic Z_7 symmetry of the symmetric closure manifold together with stability of the symmetric configuration. They constrain the admissible form of the pair spectrum — restricting it to a Fourier-mode subspace and correlating splittings with participation anomalies — but do not derive the individual mode amplitudes. The explicit identification of pair splittings with stability eigenvalues assumes a kinetic–potential decomposition of the closure functional; the general case is reserved for the microscopic paper.

Appendix D scope. Appendix D is illustrative. The toy spectrum chosen there is representative of a generic ~2% anisotropic perturbation, recognizable as approximately the n = 1 Fourier harmonic of Part IV; it is not a derivation from the microscopic closure Hamiltonian, which remains the open problem of §12.2.

The epistemic status table in §10 tracks every named result.

2. κ-field response and the K = 7 channel basis

2.1 κ-field equation

The κ-field obeys the sourced Klein–Gordon equation

$(\square + m_\kappa^2)\, \kappa(x) = \rho_{src}(x; C)$, $m_\kappa^2 = \lambda_{eff} \cdot \xi^{-2}$, $\lambda_{eff} = 3/4$.

The value $\lambda_{eff} = 3/4$ is fixed by the K = 7 closure structure and inherited from prior work. Natural units $c = \hbar = 1$ throughout.

The dimensionless commitment-driving parameter is

$C = \Phi\xi / (\hbar c)$,

interpreted at onset as source strength at the coherence-cell scale, **not** as a tunnelling barrier height.

2.2 K = 7 mode basis

Let $\{\chi_a\}_{a=1..14}$ denote the closure-channel basis, grouped into seven forward/restore pairs labelled (j, ±), j = 1, …, 7. Decompose:

$\rho_{src}(x; C) = \Sigma_a\, s_a(\omega; C)\, \chi_a(x)$, $\kappa(x; C) = \Sigma_a\, \kappa_a(C)\, \chi_a(x)$.

Mode amplitudes follow from the retarded Green's function:

$\kappa_a(C) = \int d\omega\, G_a^{ret}(\omega)\, s_a(\omega; C)$, $G_a^{ret}(\omega) = 1 / (\omega_a^2 - \omega^2 - i0^+)$.

2.3 Source onset scaling

We assume throughout — both in Part I and Part II — that the source is linear in the onset driving parameter:

$s_a(\omega; C) = C \cdot \hat{f}(\omega) \cdot \eta_a$.

This is the defining assumption of the linear-response regime. η_a is the mode participation coefficient. Nonlinear onset corrections are outside the scope of this paper; their effect is to renormalize R_{crit} without changing the ordering of pair thresholds at leading order.

3. Pair-resolved response functions

3.1 Pair projectors and activation modes

For each pair j, the pair projector is

$P_j = |\chi_{j+}\rangle\langle\chi_{j+}| + |\chi_{j-}\rangle\langle\chi_{j-}|$.

The **pair activation mode** ψ_j is the unit vector in the pair-j sector whose growth first triggers the local closure condition B_phys = 1. In the symmetric limit of the pair (no intra-pair forward/restore anisotropy), the natural choice is

$\psi_j = (\chi_{j+} + \chi_{j-}) / \sqrt{2}$, $M_{ja} := \langle\psi_j, \chi_a\rangle$,

giving M_{j,(j,±)} = 1/√2 and M_{j,(k,±)} = 0 for k ≠ j. The antisymmetric alternative and intra-pair anisotropy corrections are treated in Appendix B; they do not affect δC at the order computed.

3.2 Pair response function

Define the pair response via the dissipative (absorptive) pairing (full convention in Appendix A):

$R_j(C) := \langle\psi_j, \kappa(C)\rangle$_diss = C · K_j,

where

$\mathbf{K_j = \Sigma_a M_{ja} \cdot \eta_a \cdot |\hat{f}(\omega_a)|^2 / (2\omega_a).}$ (master formula)

K_j is the pair spectral weight — the dissipative response of pair j per unit drive. It is positive definite and real.

Definition of the symmetric baseline K_0. Throughout this paper, K_0 denotes the exact common pair spectral weight in the fully symmetric K = 7 closure limit, i.e. the value for which $K_j = K_0$ for all seven pairs when all pair-level asymmetries vanish. Away from symmetry, K_0 is retained as a **fixed reference inherited from the symmetric limit**; it is *not*, unless explicitly stated, identified with the arithmetic mean of the perturbed $\{K_j\}$. This distinction is essential: the theorems of Part I are formulated relative to the symmetric baseline, not relative to an arbitrary average over an already-broken spectrum. Deviations are written $\delta K_j := K_j - K_0$, with K_0 fixed by the underlying symmetric configuration and not re-centred pair-by-pair.

3.3 Threshold condition

The pair threshold is the smallest C at which pair j crosses the closure condition:

$C_j = \min\{C : R_j(C) = R\_crit\} = R\_crit / K_j$.

The commitment barrier coefficient is

$C^* = \max_j C_j$, $\delta C = C^* - 3/8$.

3.4 Normalization of R_crit

Matching the symmetric limit (all $K_j = K_0$) to the derived baseline $C_0 = 3/8$ fixes

$R\_crit = (3/8) \cdot K_0$.

Epistemic status of this normalization. The identification $R\_crit = (3/8)\cdot K_0$ should be read carefully. It is *not* an independent first-principles derivation within the present paper. Rather, it is the **matching condition** that imports the previously established symmetric-limit barrier value $C_0 = 3/8$ into the pair-resolved response formalism. The present paper therefore does not derive the symmetric baseline anew; it derives how departures from pair symmetry shift the threshold away from that baseline. In this sense, the novelty of the paper lies in the structure and computation of δC, not in a re-derivation of the baseline value 3/8.

With this matching in place, R_crit is a derived constant of the theory relative to the symmetric limit, and every threshold in the theory is measured against it.

PART I — Structural Theorems

The results in this part follow from §§2–3 alone within the linear-onset response formalism. No regime-specific spectral assumption is invoked. Proofs are deferred to Appendix C; statements and consequences are given here.

4. Existence and symmetric closure

4.1 Existence of pair thresholds

Theorem 4.1 (Existence). *Under the linear-onset scaling of §2.3, each $R_j(C)$ is strictly monotonically increasing in C with $R_j(0) = 0$. The pair threshold $C_j = R\_crit/K_j$ is therefore unique and well-defined for every $j \in \{1, \ldots, 7\}$ with $K_j > 0$.*

Remark. Positivity $K_j > 0$ holds automatically under the dissipative definition of R_j (Appendix A). The theorem fails only in degenerate cases where a pair decouples entirely from the bath ($K_j = 0$), corresponding to an absent closure channel. Such cases are physically excluded in the $K = 7$ manifold, which is by construction fully closed.

4.2 Symmetric closure

Theorem 4.2 (Symmetric closure, iff). *$\delta C = 0$ if and only if $K_j = K_0$ for all $j \in \{1, \ldots, 7\}$.*

Scope of the iff. The iff character of Theorem 4.2 should be understood relative to the fixed symmetric normalization of §3.4. Once K_0 is defined as the exact common pair weight in the symmetric limit and R_crit is matched to the corresponding baseline $C_0 = 3/8$, vanishing threshold correction is equivalent to **exact flattening of the pair spectrum relative to that fixed reference**. The theorem therefore establishes not merely that spectral symmetry is sufficient for $\delta C = 0$, but that within the adopted normalization it is also necessary: any genuine pair-level asymmetry — any departure of some K_j from the K_0 inherited from the symmetric limit — must reappear as a nonzero barrier correction.

The theorem is therefore **structural within the normalized response framework, rather than invariant under arbitrary reparameterizations of the threshold condition**. Alternative normalizations that redefine either K_0 (e.g., as a spectrum-dependent mean) or R_crit would change the precise content of the iff statement, though not the qualitative fact that spectral asymmetry and barrier correction are in correspondence.

Consequence. Within the normalization of §3.4, the barrier correction and the spectral asymmetry are in one-to-one correspondence.

---

5. Positivity, boundedness, scale invariance

5.1 Positivity

Theorem 5.1 (Conditional positivity). *Let $K_{min} = \min_j K_j$. Then:*

- *$\delta C \geq 0$ iff $K_{min} \leq K_0$.*
- *$\delta C > 0$ iff $K_{min} < K_0$ strictly.*

Remark on physical positivity. Theorem 5.1 is algebraically exact, but its physical content depends on whether the perturbed spectrum contains at least one pair whose response is not stronger than the symmetric baseline. **This is satisfied in all the physically relevant regimes considered here.** In particular, any mean-preserving perturbation, any splitting-dominated perturbation, and any mixed perturbation with at least one participation deficit automatically satisfy $K_{min} \leq K_0$, implying $\delta C \geq 0$. The exceptional complement, in which every pair is uniformly enhanced above baseline, would correspond not to ordinary symmetry breaking but to a collective upward renormalization of the entire closure manifold. That case is logically allowed by the algebra but is not the regime physically associated with pair-threshold splitting.

Corollary 5.2. *In the splitting-dominated regime of Part II (β_j negligible, $\Omega_j \neq 0$ for some j), $\delta C > 0$ automatically.*

5.2 Non-linear boundedness

Theorem 5.3 (Tight boundedness). *Let $\varepsilon^- = max\_\{j: \delta K_j < 0\} (-\delta K_j/K_0)$ (the maximum fractional response deficit), and $\varepsilon^+ = max_j\ \delta K_j/K_0$. Then the pair splittings satisfy the exact bounds*

$\delta_j \in [-(3/8) \cdot \varepsilon^+/(1 + \varepsilon^+) , (3/8) \cdot \varepsilon^-/(1 - \varepsilon^-)]$

and the barrier correction is

$\boldsymbol{\delta C = (3/8) \cdot \varepsilon^- / (1 - \varepsilon^-)}$ **if** $\boldsymbol{\varepsilon^- > 0}$**,**

$\boldsymbol{\delta C = -(3/8) \cdot \varepsilon^+ / (1 + \varepsilon^+)}$ **if** $\boldsymbol{\varepsilon^- = 0}$**.**

Remark (versus linear bound). To first order in ε^-, this reduces to the familiar $\delta C \approx (3/8)\cdot\varepsilon^-$ bound. The full expression is exact (non-perturbative in ε^-) and tighter for $\varepsilon^- \gtrsim 0.3$. Its singularity at $\varepsilon^- \to 1$ marks the breakdown of linear response: at that limit $K_{min} \to 0$, a pair decouples entirely, and the barrier formally diverges — which is the correct physical behaviour (loss of a closure channel makes commitment infinitely costly).

Numerical verification. The illustrative computation of Appendix D saturates this bound with $\varepsilon^- = 0.022$ and produces $\delta C = 8.4 \times 10^{-3}$, matching the Theorem 5.3 value exactly.

5.3 Scale invariance

Theorem 5.4 (Scale invariance). *δC depends only on the dimensionless ratios $\{K_j/K_0\}$, not on their absolute values. Equivalently, the transformation $K_j \to \lambda K_j$ for all j (with $R\_crit \to \lambda R\_crit$ via the normalization of §3.4) leaves every C_j unchanged.*

This is a sanity check: the barrier correction is an intrinsic spectral quantity, independent of the overall normalization of the response. It also means measurement programmes targeting δC can work with any consistent normalization of the pair weights.

5.4 What Part I does not supply

Part I fixes the sign, bounds, and closure conditions of δC, but it does not provide:

- The value of δC for any specific spectrum.
- The functional form of δK_j in terms of physical pair parameters.
- A falsifiable numerical prediction.

These are the content of Part II.

PART II — Controlled-Regime Computation

This part computes K_j and δC in closed form under three controlled assumptions. All Part II statements are proven **conditional on (A1)–(A3)** below. Part I theorems continue to apply as universal constraints.

6. Controlled model

6.1 Assumptions

(A1) Narrow-band commitment bath.

$|\hat{f}(\omega)|^2 = 1/\sqrt{(2\pi\Delta^2)} \cdot \exp[-(\omega - \omega_0)^2 / (2\Delta^2)], \Delta / \omega_0 \equiv \varepsilon\_\Delta \ll 1.$

(A2) Near-degenerate pair spectrum.

$\omega\_\{j,\pm\} = \omega_0 \pm \Omega_j / 2, \max_j (\Omega_j / \omega_0) \equiv \varepsilon\_\Omega \ll 1.$

(A3) Smooth participation anomaly.

$\eta\_\{j,\pm\} = \eta_0 \cdot (1 + \beta_j), |\beta_j| \equiv \varepsilon\_\beta \ll 1.$

Intra-pair forward/restore asymmetry $\alpha_j = (\eta_{j+} - \eta_{j-})/\eta_0$ is assumed small compared to β_j (justification and treatment in Appendix B).

6.2 Why these assumptions are natural

The assumptions (A1)–(A3) are not arbitrary conveniences; they identify the **minimal tractable neighborhood of the symmetric K = 7 closure limit**.

- Assumption (A1) captures the regime in which the commitment-event bath is concentrated near the characteristic κ-sector scale, so that pair-to-pair differences arise primarily from displacement relative to a common spectral envelope rather than from broad multi-scale forcing. The commitment-event bath spectral density J(ω) of the commitment-barrier paper is already concentrated near the κ-sector mass scale; narrow-band forcing is the physical regime in which J(ω) is well-approximated by a Gaussian around its peak.
- Assumption (A2) encodes the fact that exact symmetry of the K = 7 manifold implies degeneracy, making small pair splittings the natural first symmetry-breaking deformation. In the K = 7 symmetric limit all pair frequencies coincide at ω_0; residual splittings $\{\Omega_j\}$ are the leading symmetry-breaking effect of the K = 7 group structure.
- Assumption (A3) similarly treats participation anomalies as weak smooth departures from the symmetric exposure pattern.

Taken together, these assumptions do not claim necessity for all realizations of the closure manifold; rather, they define the nearest analytically controllable regime around the symmetric point, which is **precisely the regime in which one expects the leading correction to the 3/8 baseline to be computable in closed form**.

6.3 Small parameters and ordering

Three independent small parameters: ε_Δ, ε_Ω, ε_β. The dominant spectral control parameter is the ratio

$\mathbf{r_j = \Omega_j / \Delta}$,

which is *not* required to be small — only $\varepsilon\_\Omega = \Omega_j/\omega_0$ is. Narrow-band (ε_Δ ≪ 1) is compatible with any value of r_j up to saturation of the Gaussian envelope; this is the regime where spectral resolution of the pair splittings becomes meaningful.

---

7. Pair weights in closed form

7.1 Direct computation

From the master formula of §3.2 with M_{j,(j,±)} = 1/√2 and the assumptions of §6.1:

$K_j = (1/\sqrt{2}) \cdot [\eta_{j+}|\hat{f}(\omega_{j+})|^2/(2\omega_{j+}) + \eta_{j-}|\hat{f}(\omega_{j-})|^2/(2\omega_{j-})]$.

Under (A1)–(A3):

$|\hat{f}(\omega\_\{j,\pm\})|^2 = 1/\sqrt{(2\pi\Delta^2)} \cdot \exp[-(\Omega_j/2)^2 / (2\Delta^2)]$ (identical for + and −),

$1/(\omega_{j+}) + 1/(\omega_{j-}) = (2/\omega_0) \cdot [1 - (\Omega_j/2\omega_0)^2]^{(-1)} \approx (2/\omega_0) \cdot [1 + \Omega_j^2/(4\omega_0^2)]$.

Collecting terms:

$\mathbf{K_j = K_0 \cdot (1 + \beta_j) \cdot \exp[-r_j^2/8] \cdot [1 + (\varepsilon\_\Omega\ \varepsilon\_\Delta)^2\ r_j^2 / 4 + O(\varepsilon^4)]}$,

with symmetric-limit normalization

$\mathbf{K_0 = (\eta_0 / (\sqrt{2} \cdot \omega_0)) \cdot 1/\sqrt{(2\pi\Delta^2)}}$.

This K_0 is the symmetric-limit value specified in §3.2 and is the fixed reference against which δK_j is measured.

7.2 Leading-order result

Expanding the Gaussian envelope to second order in r_j^2 and dropping ε_Δ²-suppressed terms:

⌈ $\boldsymbol{\delta K_j / K_0 \approx \beta_j - r_j^2 / 8}$ ⌉ (central quantitative result)

The $r_j^2/(8\omega_0^2)$ correction from the 1/ω Taylor expansion is $O(\varepsilon\_\Delta^2)$ smaller than the Gaussian-envelope contribution $r_j^2/8$ and is negligible in the narrow-band regime.

7.3 Retained-Gaussian form

For r_j not small but ε_Ω still small, the full envelope can be kept:

$\delta K_j / K_0 = (1 + \beta_j) \cdot \exp[-r_j^2/8] - 1$ (valid for $\varepsilon\_\Omega \ll 1$, any r_j up to envelope saturation).

This extends the validity of the computation beyond the r_j-perturbative regime while remaining closed-form.

---

8. Threshold splitting and falsifiability

8.1 Pair thresholds

Inserting §7.2 into $C_j = (3/8)/(1 + \delta K_j/K_0)$ and expanding:

⌈ $\boldsymbol{\delta_j = C_j - 3/8 \approx (3/32)\, r_j^2 - (3/8)\, \beta_j}$ ⌉

Two physical observations:

1. **Pair splitting always raises the threshold.** The r_j^2 term is strictly positive; spectral displacement from band-centre suppresses dissipative coupling.
2. **Participation anomaly can go either way.** $\beta_j > 0$ (enhancement) lowers C_j; $\beta_j < 0$ (deficit) raises it.

8.2 Barrier coefficient

The barrier is set by the hardest pair:

⌈ $\boldsymbol{\delta C = \max_j [\, (3/32)\, r_j^2 - (3/8)\, \beta_j \,]}$ ⌉

Consistency with Part I. Positivity (Theorem 5.1) requires at least one pair with $K_j \leq K_0$, i.e., at least one pair with $\beta_j - r_j^2/8 \leq 0$. In the splitting-dominated regime this is automatic: $r_j^2 \geq 0$ always, and if β_j is negligible, $K_j \leq K_0$ for all j with any splitting. Thus Part II lies safely within the Part I positivity region.

8.3 Splitting-dominated limit and one-parameter prediction

When β_j contributions are negligible (i.e., $|\beta_j| \ll r_j^2/8$ for all j):

$\delta C = (3/32) \cdot r_{max}^2$**,** $r_{max} = max_j (\Omega_j / \Delta)$**.**

This is a one-parameter prediction. r_{max} is in principle independently determinable from the K = 7 closure spectrum, giving a sharp falsifiability criterion:

Falsifiability statement. *Two independent measurements — (i) the maximum pair splitting Ω_{max} in the K = 7 closure manifold, and (ii) the bath spectral width Δ — determine r_{max}. Any inconsistency between (δC inferred from Φ_c measurements) and ($(3/32)\ r_{max}^2$) falsifies the splitting-dominated regime.*

8.4 Benchmark scaling

Illustrative values (parametric in r_{max}; full end-to-end numerical example in Appendix D):

$r_{max} = \Omega_{max}/\Delta$	δC	δC / (3/8)	δΦ_c / Φ_0
0.01	9.4×10^{-6}	2.5×10^{-5}	2.5×10^{-5}
0.1	9.4×10^{-4}	2.5×10^{-3}	2.5×10^{-3}
0.3	8.4×10^{-3}	2.25×10^{-2}	2.25×10^{-2}
1.0	9.4×10^{-2}	0.25	0.25

For $r_{max} \gtrsim 1$ the perturbative form breaks down; use the retained-Gaussian expression of §7.3. For $r_{max} \gtrsim 2$ the closed-form envelope itself saturates and full numerical integration is required.

PART III — Consequences

9. Downstream: Φ_c, Λ, α

9.1 Commitment barrier

From C* = 3/8 + δC and Φ_c = C* · ℏc/ξ:

$\delta\Phi\_c / \Phi_0 = (8/3)\ \delta C = r_{max}^2 / 4$ (splitting-dominated).

This is rigid: fractional correction to Φ_c equals one-quarter of r_{max}^2. Falsifiability inherited from §8.3.

9.2 Cosmological-constant spread

Conditional on the Λ-spread paper's identification of δC as the microphysical source of vacuum-closure variance:

$\delta\Lambda / \Lambda = f\_\Lambda \cdot r_{max}^2 / 4$,

where f_Λ is the transfer function fixed by the Λ-spread paper (not re-derived here). This paper supplies the microphysical input; the Λ-spread paper supplies the coupling.

9.3 α-sector

The symmetric-limit $\alpha^{-1} \approx 137.034$ result uses only K_0 and is therefore **unchanged at first order** by pair splitting. Corrections enter at second order via channel-competition terms:

$$\boldsymbol{\delta\alpha / \alpha = O((\delta C/C_0)^2) = O(r_{max}^4 / 64).}$$

Cross-sector constraint and empirical bound. For $r_{max} \lesssim 0.15$ the correction is below 2×10^{-6}, well below the 15 ppm accuracy of the symmetric-limit α derivation. The 15 ppm bound on δα/α therefore **constrains r_{max}**:

$r_{max} \lesssim 0.15$ (from α-sector cross-sector constraint, splitting-dominated regime).

This is not a loose consistency check: it is a structural constraint on the allowed spectral regime of the K = 7 manifold from an independent sector of the programme. Any measurement of r_{max} exceeding this bound would force reconsideration of either (i) the splitting-dominated assumption (requiring β_j contributions to cancel the r_{max}^4 term), or (ii) the α-sector derivation's inclusion of pair-splitting corrections.

Cross-check with Appendix D. The toy spectrum of Appendix D produces $\delta C \approx 8.4 \times 10^{-3}$, corresponding to $\delta C/C_0 \approx 0.022$. In a splitting-dominated interpretation this would require $r_{max}^2 \approx 0.18$, i.e. $r_{max} \approx 0.42$, which **exceeds** the α-sector constraint $r_{max} \lesssim 0.15$. The toy spectrum is therefore *not* in the splitting-dominated regime: its δC must be driven by participation-anomaly contributions β_j rather than pair splitting. This is precisely the regime flagged in §8.1 point (2) and does not violate the α-sector constraint, since the $O(r_{max}^4/64)$ correction vanishes for β-driven δC.

This is a useful illustration of how the α-bound refines interpretation: the same δC value can sit in different sub-regimes with different upstream physics, and cross-sector constraints identify which sub-regime is realized.

9.4 Unification statement

The three downstream quantities — Φ_c correction, Λ-spread, α-sector bound — depend on the K = 7 pair spectrum through a single dimensionless input in the splitting-dominated regime:

$\{\delta\Phi\_c/\Phi_0, \delta\Lambda/\Lambda$, bound on $\delta\alpha/\alpha\} \leftarrow r_{max}$ alone.

More generally (without the splitting-dominated assumption): $\{r_j, \beta_j\}$ per pair. This is the structural unification claim, now precise at the level of shared variables rather than asserted.

PART IV — Symmetry-Level and Minimal-Model Constraints on the Closure Spectrum

10. Structural constraints on the K = 7 closure spectrum

10.1 Scope and objective

Parts I–III reduce the commitment-threshold correction δC to a functional of the pair-resolved closure spectrum $\{\Omega_j, \beta_j\}$. The question now is whether that spectrum is genuinely free input or carries structure of its own.

The purpose of this part is to take a first step toward the microscopic derivation. We do *not* yet attempt a full first-principles calculation of $\{\Omega_j, \beta_j\}$ from a complete closure Hamiltonian. Instead, we show that the allowed form of the closure spectrum is already strongly constrained by **symmetry, stability, and coupling geometry**, and therefore cannot be treated as an arbitrary set of seven independent parameters.

The result is a reduction of the spectral freedom from an unconstrained seven-pair assignment to a structured, low-dimensional family of admissible spectra.

10.2 Closure symmetry and mode decomposition

In the symmetric limit, the K = 7 closure manifold is invariant under cyclic permutation of the seven pair indices j = 1, …, 7. This symmetry implies that the natural basis for perturbations of the pair spectrum is the discrete Fourier basis on the cyclic group Z_7.

Let deviations from the symmetric baseline be written as

$\delta K_j := K_j - K_0$.

Then any admissible perturbation decomposes as

$$\boldsymbol{\delta K_j = \Sigma_{n=1}^{3} [\, A_n \cos(2\pi nj/7) + B_n \sin(2\pi nj/7) \,],}$$

where only modes up to n = 3 are independent due to reality and periodicity constraints (the n = 0 mode is excluded by the definition of K_0 as the symmetric reference of §3.2, which fixes the zero-mean condition on $\{\delta K_j\}$).

This representation has two immediate consequences:

1. **The pair spectrum is not arbitrary**: it lies in a six-dimensional Fourier subspace (three amplitudes A_n and three phases B_n) rather than a seven-dimensional free space.
2. **Smooth, low-order perturbations are dominated by the lowest harmonic n = 1**, giving a natural leading-order form

   $\delta K_j \propto \cos(2\pi j/7 + \varphi)$.

This provides a symmetry-derived explanation for the structured toy spectra of Appendix D: the monotonic single-mode anisotropy used there is recognizable as the n = 1 Fourier harmonic of this symmetry-admissible family.

10.3 Closure stability and spectral splitting

The symmetric closure configuration corresponds to a stationary point of the underlying closure dynamics. Perturbations of the pair spectrum must therefore respect stability constraints.

Let $F[\{K_j\}]$ denote the effective closure functional governing the response amplitudes. Linearizing about the symmetric point $K_j = K_0$, define the stability operator

$L_{jk} := (\partial^2 F / \partial K_j \, \partial K_k)|\_\{K = K_0\}$.

L inherits the cyclic Z_7 symmetry of the closure manifold and is therefore **diagonal in the Fourier basis of §10.2**, with eigenvalues λ_n indexed by harmonic number n.

Identification of pair splittings with stability eigenvalues (conditional). When the closure functional F admits a kinetic–potential decomposition F = T + V with T normalized at the symmetric point, the pair splittings satisfy

$\Omega_n^2 = \lambda_n$,

where λ_n are the eigenvalues of $L = \partial^2 V/\partial K \partial K|_0$ in the Fourier basis. The general-F case reduces to a generalized eigenvalue problem involving a non-trivial kinetic metric and is reserved for the microscopic spectral paper; the leading-order identification used here is sufficient to establish the structural claims of Part IV.

Under this identification:

- Each harmonic mode n corresponds to a distinct splitting scale Ω_n.
- Degeneracies among harmonics are lifted only by symmetry-breaking perturbations beyond the cyclic level.

Thus the pair splittings $\{\Omega_j\}$ are not independent parameters but are determined by a small set of mode amplitudes $\{\Omega_n\}_{n=123}$ associated with the irreducible representations of the closure symmetry.

10.4 Scaling of spectral splittings

Dimensional analysis within VERSF fixes the overall scale of the splittings. The only intrinsic length scale in the closure sector is the coherence scale ξ, implying

$$\boldsymbol{\Omega_n \sim (c/\xi) \cdot \gamma_n,}$$

where γ_n are dimensionless coefficients associated with each harmonic mode.

The scale ξ^{-1} is fixed by the closure sector; the **pattern** $\{\gamma_n\}$ is exactly what the eigenvalue problem of §10.3 determines. Sections 10.3 and 10.4 are therefore complementary: §10.3 determines which harmonic modes exist and in what ratios, §10.4 fixes the overall dimensionful scale against which those ratios are measured.

Two structural consequences:

1. All splittings are parametrically of the same order, controlled by ξ^{-1}, preventing arbitrarily large hierarchies within the closure spectrum.
2. The dimensionless ratios $r_j = \Omega_j/\Delta$ of Part II are controlled by the ratio of the intrinsic closure scale to the bath bandwidth, directly linking the spectral structure of §10 to the controlled-regime parameter of §§6–8.

10.5 Participation structure from coupling geometry

The participation coefficients β_j encode deviations in how strongly each pair couples to the commitment-event bath. These are not independent of the closure geometry.

From §3.2, the pair weights depend on projections $\eta_a \sim \langle \chi_a, \text{bath} \rangle$. Projecting onto pair modes,

$$\beta_j \sim \langle \psi_j, G \rangle,$$

where G represents the geometric coupling profile of the bath.

Under the symmetry assumptions of §10.2, the coupling profile G itself decomposes into the same Fourier modes:

$$\mathbf{G = \Sigma_n \left[g_n^{(c)} \cos(2\pi nj/7) + g_n^{(s)} \sin(2\pi nj/7) \right],}$$

and therefore β_j inherits the harmonic structure of δK_j. The splittings and the participation anomalies live in the **same Fourier basis**, with amplitudes determined by related but distinct geometric inputs (the stability operator and the bath coupling profile, respectively).

Correlation prediction. This implies that in the $K = 7$ closure manifold, splittings and participation anomalies are not independent degrees of freedom but are correlated through shared mode structure. A spectrum dominated by the $n = 1$ harmonic in δK_j is accompanied by a matching $n = 1$ harmonic structure in β_j, with relative amplitude fixed by the ratio of stability and coupling geometries. This provides an in-principle testable structural prediction: the participation-vs-splitting sub-regime of §§8.1 and 9.3 is not a dichotomy selected by external choice but a consequence of the relative strengths of the stability spectrum and the bath-coupling geometry in the *same* harmonic channel.

10.6 Constraint on spectral ratios

Combining the scaling of §10.4 with the controlled-regime structure of Part II, the key dimensionless parameter $r_j = \Omega_j/\Delta$ is constrained by two considerations:

1. Closure stability prevents Ω_j from exceeding the intrinsic scale c/ξ (from §10.4).
2. The narrow-band regime of Part II requires Δ comparable to the dominant spectral support of the bath.

The first gives $r_j \lesssim (c/\xi)/\Delta$, which by itself does *not* fix the magnitude — it depends on the ratio of closure scale to bath bandwidth, neither of which is pinned by Part IV alone.

The binding constraint on r_{max} therefore remains the α-sector cross-sector constraint of §9.3:

$r_{max} \lesssim 0.15$ (from α-sector cross-sector constraint).

Closure stability does not yet independently produce a comparable bound on r_j; it supplies the shape and symmetry of the admissible spectrum but not the absolute ratio to the bath bandwidth. The α-sector cross-sector constraint remains the operational bound. When the constraint is exceeded, the system is forced into the participation-dominated sub-regime rather than indicating a failure of the framework — consistent with the correlated-harmonic picture of §10.5.

10.7 Reduction of spectral freedom

The results of §§10.2–10.6 imply a substantial reduction in the degrees of freedom of the closure spectrum:

- Instead of seven independent Ω_j, the spectrum is described by a small number of Fourier mode amplitudes $\{\Omega_n\}_{n=123}$.
- Instead of independent β_j, participation anomalies are tied to the same Fourier mode structure with amplitudes fixed by the bath coupling geometry.
- The overall scale is fixed by ξ^{-1}; admissible ratios r_j are bounded by the α-sector cross-sector constraint.

Thus the pair-resolved closure spectrum is **a constrained subspace within the closure manifold's symmetry-admissible perturbations**, not an unconstrained seven-pair input. Only a small number of effective mode amplitudes remain undetermined, and those are the targets of the microscopic paper.

This reduces the effective dimensionality of the spectral problem from 14 independent pair-level parameters (seven splittings Ω_j plus seven participation anomalies β_j) to a small

number of symmetry-constrained mode amplitudes. The gain is not merely notational: it is the difference between an underconstrained fit and a genuinely predictive upstream derivation.

10.8 Role within the VERSF programme

Part IV completes the logical bridge between the threshold-mapping results of Parts I–III and the spectral-closing objective of the next stage:

- **Parts I–III** show that δC is a functional of $\{\Omega_j, \beta_j\}$.
- **§§10.2–10.9** show that $\{\Omega_j, \beta_j\}$ are constrained by Z_7 symmetry, closure stability, and coupling geometry to a small Fourier-mode family.
- **§10.10** shows that under minimal-model locality assumptions, that family collapses further to two effective parameters (a, b) plus the bath-width parameter d, with $\beta_j = 0$ at leading order.
- The remaining task is the microscopic derivation of (a, b, d) from the closure sector of VERSF, together with an accounting of corrections beyond the minimal model.

In this sense the present paper reduces the entire commitment-threshold problem to a well-defined microscopic target: **the derivation of the three effective parameters of the minimal K = 7 closure model, plus corrections beyond it.**

10.9 Summary of Part IV

The K = 7 closure spectrum is not arbitrary. Its structure is constrained along two complementary levels:

Symmetry level (§§10.2–10.9):

- cyclic Z_7 symmetry of the closure manifold (Fourier subspace),
- stability of the symmetric configuration (mode-by-mode eigenvalue structure, conditional on kinetic–potential decomposition),
- geometric coupling to the commitment bath (shared Fourier basis correlating splittings and participation anomalies),
- consistency with downstream observables (α-sector cross-sector constraint binding r_{max}).

Minimal-model level (§10.10):

- nearest-neighbour pair coupling (M1),
- stable symmetric vacuum (M2),
- uniform leading bath coupling (M3),
- uniquely yield a two-parameter spectrum with three doubly-degenerate levels and $\beta_j = 0$.

Together these reduce the spectrum from an unconstrained seven-pair input to a small finite parameter set whose microscopic determination is the subject of the spectral-closing paper.

10.10 Minimal-model derivation of the K = 7 closure spectrum

§§10.2–10.9 show that the pair spectrum lies in a constrained Fourier subspace, without committing to specific mode amplitudes. We now take one step further — still short of a microscopic derivation — by asking what the closure spectrum looks like under a *minimal model* satisfying the Part IV symmetry assumptions plus two additional locality conditions. The result is a two-parameter spectrum, explicit closed form, and a genuine benchmark prediction for δC.

This is a minimal-model construction, not a first-principles derivation. The parameters μ^2 and κ below are effective; the microscopic derivation of their values from a complete closure Hamiltonian is the subject of the spectral-closing paper. The present subsection shows what the closure spectrum must look like *if* the minimal assumptions hold.

10.10.1 Closure amplitudes and symmetric vacuum

Let the seven closure pairs be indexed by $j = 0, \ldots, 6$ on a cyclic manifold. Define a real pair-amplitude field q_j measuring small deviations of pair j away from the symmetric closure configuration. The symmetric vacuum is

$q_j = q_0 \; \forall \; j,$

and fluctuations are written as $q_j = q_0 + \delta q_j$. The question is: *what quadratic fluctuation operator governs δq_j?*

10.10.2 Minimal-model assumptions

Beyond the Z_7 cyclic symmetry already invoked in §10.2, we adopt three further conditions:

(M1) Pair locality (nearest-neighbour coupling). The leading closure energy depends only on local pair amplitude and mismatch between adjacent pairs on the cyclic manifold.

(M2) Symmetric-vacuum stability. The quadratic form is positive semidefinite around the symmetric configuration.

(M3) Uniform leading bath coupling. In the unbroken vacuum the bath couples identically to all pairs.

(M1) is the substantive assumption. Further-neighbour couplings (next-nearest, etc.) are also Z_7-invariant and would give a more general circulant operator with more parameters; restricting to nearest-neighbour is a definite modelling choice that will have to be revisited in the microscopic paper. (M2) is physically required; (M3) is what will give $\beta_j = 0$ at leading order.

10.10.3 Unique local cyclic quadratic functional

Under (M1)–(M2) together with Z_7 symmetry and site-uniformity, the most general quadratic fluctuation functional is

$F^{(2)}[\delta q] = \frac{1}{2} \Sigma_{j=0}^{6} [\mu^2(\delta q_j)^2 + \kappa(\delta q_{j+1} - \delta q_j)^2]$,

with indices mod 7. This is the unique lowest-order cyclic, local, translation-invariant quadratic form up to parameter redefinitions. Expanding the neighbour term gives the Hessian

$\mathbf{L_{jk} = (\mu^2 + 2\kappa)\delta_{jk} - \kappa(\delta_{j,k+1} + \delta_{j,k-1})}$ (indices mod 7).

This is the minimal-model closure operator.

10.10.4 Fourier diagonalization

L is circulant, so its eigenvectors are the discrete Fourier modes on Z_7:

$v_j^{(n)} = (1/\sqrt{7}) \exp(2\pi inj/7)$, n = 0, …, 6,

with eigenvalues

$\mathbf{\lambda_n = \mu^2 + 2\kappa - 2\kappa \cos(2\pi n/7) = \mu^2 + 4\kappa \sin^2(\pi n/7).}$

This realizes the Fourier decomposition of §10.2 explicitly: the closure operator is diagonal in the Z_7 basis and its spectrum is determined by the two parameters μ^2 and κ.

10.10.5 Derived closure frequencies and degeneracy structure

Under the kinetic–potential decomposition assumption of §10.3, the normal-mode frequencies are

$\Omega_n^2 = \lambda_n = \mu^2 + 4\kappa \sin^2(\pi n/7)$.

Since $\sin^2(\pi n/7) = \sin^2(\pi(7-n)/7)$, the spectrum exhibits the symmetry

$\mathbf{\lambda_n = \lambda_{7-n}}$**,**

giving a twofold degeneracy for each nontrivial mode. The full spectrum therefore consists of:

- **One uniform mode** (n = 0): $\lambda_0 = \mu^2$. This is the collective breathing mode, not a pair-splitting mode.
- **Three doubly-degenerate nontrivial levels** (n = 1,2,3 paired with n = 6,5,4 respectively).

Using $\sin^2(\pi/7) \approx 0.1883$, $\sin^2(2\pi/7) \approx 0.6113$, $\sin^2(3\pi/7) \approx 0.9505$:

$\Omega_{(1)}^2 \approx \mu^2 + 0.7532\ \kappa$,

$\Omega_{(2)}^2 \approx \mu^2 + 2.4452\ \kappa$,

$\Omega_{(3)}^2 \approx \mu^2 + 3.8019\ \kappa$.

The maximum closure splitting is therefore

$$\boldsymbol{\Omega_{max} = \Omega_{(3)} = \sqrt{(\mu^2 + 3.8019\ \kappa)}.}$$

Degeneracy observation. The minimal-model prediction has degenerate pairs (n and 7−n); any spectrum lying strictly along a single harmonic (e.g., the toy of Appendix D, which is monotonic in j) must be read as a non-generic sub-family. Appendix D's toy serves its intended pipeline-demonstration role and is not claimed to match the minimal-model degeneracy structure.

10.10.6 Participation anomalies from bath-coupling structure

Under (M3) the bath-coupling field g_j on the seven-pair manifold is uniform in the leading vacuum, so its Fourier decomposition reduces to the n = 0 mode. Expanding general departures from uniform coupling in the Fourier basis of §10.2,

$$g_j = \overline{g} + \Sigma_{n=1}^{3} [\alpha_n \cos(2\pi nj/7) + \tilde{\alpha}_n \sin(2\pi nj/7)],$$

the participation anomaly is

$$\beta_j \propto g_j - \overline{g} = \Sigma_{n=1}^{3} [\alpha_n \cos(2\pi nj/7) + \tilde{\alpha}_n \sin(2\pi nj/7)].$$

Two consequences:

1. **$\boldsymbol{\beta_j}$ lives in the same Fourier sector as $\boldsymbol{\delta K_j}$**, realizing the correlated-harmonic picture of §10.5 explicitly.
2. **In the minimal (M3)-respecting limit, $\boldsymbol{\alpha_n = \tilde{\alpha}_n = 0}$ for all n, hence $\boldsymbol{\beta_j = 0}$.**

The second consequence is significant: it provides a derived reason why the splitting-dominated sub-regime of §§8.1 and 9.3 is the *default* starting point, while participation-dominated behaviour requires explicit symmetry breaking of (M3) in the bath-coupling geometry. The splitting vs participation dichotomy is not an external choice between two equally-likely options; the minimal model selects splitting.

10.10.7 Derived form of r_max and δC

To connect to VERSF scales, introduce dimensionless parameters

$$\mu^2 = a\ \xi^{-2},\ \kappa = b\ \xi^{-2},\ \Delta = d\ \xi^{-1},$$

with a, b, d > 0 dimensionless. Then

$$\Omega_n^2 = \xi^{-2} [a + 4b \sin^2(\pi n/7)],$$

and the maximum dimensionless spectral ratio is

$$\mathbf{r_{max}^2 = (a + 4b\,\sin^2(3\pi/7)) / d^2 = (a + 3.8019\,b) / d^2.}$$

In the splitting-dominated regime (which the minimal model selects automatically by §10.10.6), the threshold correction is

$$\boxed{\delta C = (3/32) \cdot (a + 3.8019\,b) / d^2} \text{ (minimal-model prediction, splitting-dominated)}$$

10.10.8 Benchmark scaling and α-sector consistency

For three effective parameters (a, b, d), the α-sector cross-sector constraint $r_{max} \lesssim 0.15$ of §9.3 becomes a constraint on the dimensionless ratio $\sqrt{(a + 3.80b)}/d$. Specifically:

$r_{max} \lesssim 0.15 \Leftrightarrow d^2 \gtrsim 44.4 \cdot (a + 3.80b)$,

i.e. **the bath bandwidth Δ must be much larger than the closure-stability scale ξ^{-1}** — specifically $d \gtrsim 6.7 \cdot \sqrt{(a + 3.80b)}$.

Benchmark cases (illustrative; conditional on the minimal-model assumptions and on specific choices of a, b, d):

(a, b, d)	$\mathbf{r_{max}^2}$	$\mathbf{r_{max}}$	**δC**	**α-consistent?**
(1, 1, 1)	4.80	2.19	0.450	No — exceeds bound by ~15×
(1, 0.5, 1)	2.90	1.70	0.272	No — exceeds bound
(1, 1, 5)	0.192	0.44	1.80×10^{-2}	No — exceeds bound by ~3×
(1, 1, 15)	0.0213	0.146	2.00×10^{-3}	Yes — at bound
(1, 1, 25)	0.0077	0.088	7.2×10^{-4}	Yes — well within bound

Interpretation. At "natural" parameter scales a ~ b ~ d ~ 1 the minimal model predicts δC of order 0.1–0.5, grossly violating the α-bound. Consistency therefore requires either $d \gg \sqrt{(a + 4b)}$ (bath bandwidth much larger than closure scale), or a complete reassessment of the minimal model's assumptions (M1)–(M3) for the physical manifold, or both. This is itself a predictive content of the combined Part IV + minimal-model + α-sector chain: **it constrains the bath-to-closure scale ratio rather than fitting it.**

Appendix E consolidates this α-consistency constraint with the closure-scale anchor (a, b ∼ 1) and a physical-relevance argument (b ∼ a conditional on δC being non-trivial) to produce sharper bounds: $d \gtrsim 14$ at the natural parameter scale, and a structural upper bound $\delta C \lesssim 2.1 \times 10^{-3}$ on the minimal-model threshold correction.

10.10.9 Theorem statement

Theorem 10.1 (Minimal K = 7 closure spectrum). *Consider small fluctuations of a seven-pair closure manifold about a symmetric vacuum, and assume: (i) cyclic Z_7 invariance, (ii) pair locality at quadratic order with nearest-neighbour coupling, (iii) stability of the symmetric vacuum, and (iv) uniform leading bath coupling. Then:*

(a) The quadratic closure operator is circulant and uniquely takes the form

$L_{jk} = (\mu^2 + 2\kappa)\delta_{jk} - \kappa(\delta_{j,k+1} + \delta_{j,k-1})$

up to parameter redefinition, with normal modes the discrete Fourier modes on Z_7 and eigenvalues

$\lambda_n = \mu^2 + 4\kappa \sin^2(\pi n/7)$, n = 0, …, 6.

(b) The nontrivial closure spectrum consists of three doubly-degenerate levels ($n \leftrightarrow 7-n$) with

$\Omega_{(n)}^2 = \mu^2 + 4\kappa \sin^2(\pi n/7)$, n = 1, 2, 3.

(c) In the uniform-bath-coupling limit, $\beta_j = 0$ at leading order; the system is automatically in the splitting-dominated sub-regime.

(d) The splitting-dominated threshold correction is

$\delta C = (3/32d^2)(a + 3.8019\ b)$

where $a = \mu^2\xi^2$, $b = \kappa\xi^2$, $d = \Delta\xi$.

10.10.10 Scope and what remains microscopic

Theorem 10.1 is a minimal-model result, not a full first-principles derivation. Specifically:

- **μ^2 and κ remain effective parameters.** Their values are not derived here from a microscopic closure Hamiltonian.
- **The bath-width coefficient $d = \Delta\xi$ is also effective.**
- **Nearest-neighbour coupling (M1) is an assumption.** Further-neighbour couplings would add more parameters while preserving Z_7 symmetry and locality in a weaker sense.
- **(M3) uniform-bath coupling is an assumption.** Participation-dominated regimes require its explicit violation, which the spectral-closing paper must address.

What the theorem *does* establish is that under the minimal assumptions, the K = 7 closure spectrum is **uniquely determined up to two dimensionless parameters (a, b)** — rather than fourteen freely assignable pair-level numbers — with a specific degeneracy structure, an explicit frequency formula, and a derived selection of the splitting-dominated sub-regime.

Status and Summary

11. Epistemic status and scope disclaimers

11.1 What this paper does not claim

A potential misunderstanding should be excluded explicitly. **This paper does not claim that the K = 7 pair splittings Ω_j or participation anomalies β_j have already been derived from the microscopic closure Hamiltonian.** Nor does it claim that Appendix D supplies empirical evidence for the physical spectrum. Nor does the minimal-model derivation of §10.10 claim to fix the two effective parameters μ^2 and κ. What is established here is narrower but precise:

1. The threshold correction δC has a tightly constrained formal structure.
2. It admits closed-form evaluation in a controlled neighborhood of the symmetric limit.
3. The upstream pair spectrum lies in a constrained Fourier-mode subspace determined by Z_7 symmetry, stability, and coupling geometry.
4. Under the additional minimal-model assumptions (M1)–(M3) of §10.10 — nearest-neighbour coupling, stable symmetric vacuum, uniform leading bath coupling — the spectrum is uniquely determined up to two effective parameters (a, b), with $\beta_j = 0$ at leading order and a derived degeneracy structure.
5. Once those parameters (and the bath scale d) are determined microscopically, downstream corrections to Φ_c, Λ, and the α-sector can be propagated without ambiguity.

The remaining open problem is upstream, not downstream — and is narrower than "derive the seven-pair spectrum," being specifically the microscopic determination of the minimal-model effective parameters, together with an accounting of corrections beyond (M1) nearest-neighbour coupling.

11.2 Status of each named result

§	Result	Status	Conditions
3.4	$R\_crit = (3/8)\ K_0$	**Matching condition**	Imports symmetric-limit baseline $C_0 = 3/8$
4.1	Existence and uniqueness of C_j	**Theorem** (structural)	Linear-onset formalism of §§2–3
4.2	Symmetric closure (iff)	**Theorem** (structural)	Normalization of §3.4
5.1	Conditional positivity	**Theorem** (structural)	—
5.3	Tight non-linear bound	**Theorem** (structural)	—
5.4	Scale invariance	**Theorem** (structural)	—
7.1	K_j closed form	**Derived**	(A1)–(A3)
7.2	$\delta K_j/K_0 \approx \beta_j - r_j^2/8$	**Derived**	(A1)–(A3), leading order
7.3	Retained-Gaussian form	**Derived**	(A2)–(A3), any r_j
8.1	$\delta_j = (3/32)\ r_j^2 - (3/8)\ \beta_j$	**Derived**	(A1)–(A3)

§	Result	Status	Conditions
8.2	δC formula	**Derived**	(A1)–(A3)
8.3	$\delta C = (3/32)\, r_{max}^2$	**Derived**	(A1)–(A3) + splitting-dominated
9.1	$\delta\Phi\_c/\Phi_0 = r_{max}^2/4$	**Derived**	§8.3 conditions
9.2	Λ-spread coupling	**Conditional**	Λ-spread paper
9.3	α bound $r_{max} \lesssim 0.15$	**Cross-sector constraint**	α-sector paper
9.4	Single-input unification	**Derived**	Splitting-dominated
10.2	Fourier subspace decomposition	**Symmetry-structural**	Z_7 cyclic symmetry
10.3	$\Omega_n^2 = \lambda_n$ identification	**Structural (conditional)**	Kinetic–potential decomposition of F
10.4	Scale $\Omega_n \sim c/\xi \cdot \gamma_n$	**Dimensional**	VERSF closure scale
10.5	β_j and δK_j share Fourier basis	**Symmetry-structural**	Z_7 symmetry of bath coupling
10.6	r_{max} constrained by α-sector	**Cross-sector** (binding)	α-sector paper
10.7	Spectrum → Fourier-mode subspace	**Derived** (symmetry level)	§§10.2–10.6
10.10.3	Unique nearest-neighbour quadratic form	**Minimal-model theorem**	(M1)–(M3)
10.10.4	$\lambda_n = \mu^2 + 4\kappa \sin^2(\pi n/7)$	**Derived** (minimal model)	Circulant diagonalization
10.10.5	Three doubly-degenerate levels ($n \leftrightarrow 7-n$)	**Derived** (minimal model)	Z_7 reflection symmetry
10.10.6	$\beta_j = 0$ at leading order	**Derived** (minimal model)	(M3) uniform bath coupling
10.10.7	$\delta C = (3/32d^2)(a + 3.8019b)$	**Derived** (minimal model)	§§10.10.2–10.10.6
10.10.8	$d \gtrsim 6.7\cdot\sqrt{(a + 3.80b)}$ from α-bound	**Cross-sector** (binding)	α-sector paper + minimal model
App. D	Toy $\delta C \approx 8.4 \times 10^{-3}$	**Illustrative**	Chosen toy ε_j values
App. E	a, b ~ O(1); b ~ a; $d \gtrsim 14$; $\delta C \lesssim 2.1\times10^{-3}$	**Conditional** (compiled)	Minimal model + α-bound + physical-relevance argument
—	Microscopic values of μ^2, κ, Δ	**Open**	Spectral-closing paper

11.3 The structural / computational / symmetry-level / illustrative split

Part I results (§§4–5) are theorems within the linear-onset response formalism: they require only the formal definitions of §§2–3 and hold for any K = 7 spectrum admitting monotonic onset. Part

II results (§§7–8) are closed-form derivations conditional on (A1)–(A3). Part IV results (§10) are symmetry-level constraints: they follow from the Z_7 symmetry of the symmetric closure manifold and from stability of that configuration, and narrow the upstream spectrum to a Fourier-mode subspace. Appendix D's numerical value is explicitly illustrative: it plugs a representative input — recognizable as the n = 1 Fourier harmonic of Part IV — into the Part I/II machinery and shows that the pipeline produces a specific, self-consistent number, but does not claim to predict δC from first principles.

---

12. Summary and open directions

12.1 What has been reduced

The key contribution of this paper is the proof that once the pair spectrum is specified within the Part IV symmetry-admissible subspace, the induced correction to the commitment threshold is no longer arbitrary. It is structurally constrained by the theorems of Part I, explicitly computable in the controlled regime of Part II, consistent across sectors by the α-sector constraint of Part III, sourced from a small Fourier-mode family by the symmetry constraints of §§10.2–10.9, and — under minimal-model locality assumptions — uniquely determined up to two effective parameters by §10.10. The paper converts δC from a residual free correction into a well-defined functional of a small number of spectral mode amplitudes.

Concretely, the threshold-splitting problem has been reduced along five axes:

- **Structurally.** Four theorems pin down the shape of δC for any K = 7 spectrum within the linear-onset response formalism. δC is non-negative in the physically relevant regimes considered here, bounded above by an exact non-linear function of the maximum fractional response deficit, vanishes iff the spectrum is symmetric (relative to the fixed symmetric normalization of §3.4), and is scale-invariant.
- **Computationally.** In the narrow-band near-degenerate regime, δC is given in closed form by two pair inputs per channel and by a single scalar r_{max} in the splitting-dominated sub-regime. The α-sector provides a cross-sector constraint $r_{max} \lesssim 0.15$ consistent with the 15 ppm accuracy of the α derivation.
- **Symmetry-structurally.** The pair spectrum is shown to lie in a constrained subspace of Z_7-admissible perturbations, with $\{\Omega_j\}$ and $\{\beta_j\}$ correlated through shared Fourier structure. The seven-pair spectrum reduces to a small number of mode amplitudes.
- **Minimal-model constructively.** Under (M1) nearest-neighbour coupling, (M2) vacuum stability, and (M3) uniform leading bath coupling, the closure spectrum collapses to two effective parameters (a, b) with three doubly-degenerate nontrivial levels $\Omega_n^2 = \xi^{-2}[a + 4b \sin^2(\pi n/7)]$. Participation anomalies vanish at leading order; the splitting-dominated sub-regime is selected automatically. The threshold correction is $\delta C = (3/32d^2)(a + 3.8019\ b)$, and α-sector consistency requires $d \gtrsim 6.7\cdot\sqrt{(a + 3.80b)}$. Appendix E compiles the scale anchor (a, b ∼ 1), the physical-relevance condition (b ∼ a), and the α-bound into a structural upper bound $\boldsymbol{\delta C \lesssim 2.1 \times 10^{-3}}$ — placing the minimal-model prediction firmly at the sub-percent level.

- **Illustratively.** Appendix D runs the formalism end-to-end on a specific toy seven-pair spectrum — recognizable as approximately the n = 1 harmonic of the Part IV decomposition — producing $\delta C \approx 8.4 \times 10^{-3}$. This value saturates the Part I tight bound of Theorem 5.3 exactly; the toy is monotonic across j and does not reproduce the minimal-model's (n ↔ 7−n) degeneracy structure, consistent with its purely illustrative role.

The brackets close: the Part II explicit value $(3/32)r_{max}^2$ is manifestly non-negative (Theorem 5.1), bounded (Theorem 5.3), vanishes iff the spectrum is symmetric (Theorem 4.2), scale-invariant (Theorem 5.4), Z_7-admissible (§10), and under minimal-model assumptions takes the specific two-parameter form of §10.10.7. The toy computation of Appendix D lives inside all four Part I theorem constraints.

At this stage, the commitment-threshold problem is no longer underdetermined in structure, but only in the microscopic values of a small number of parameters.

12.2 The remaining open problem

The bottleneck is no longer the commitment-threshold splitting itself, nor the question of whether the pair spectrum is free input. §§10.2–10.9 have narrowed the freedom to a small Fourier-mode family; §10.10 has further reduced it under minimal-model assumptions to two effective parameters (a, b) plus the bath-width parameter d. **The remaining open problem is twofold:**

1. **Microscopic determination of (a, b, d) from the closure Hamiltonian** — i.e. the values of the pair-locality coefficients and bath bandwidth in fundamental VERSF units.
2. **Accounting for corrections beyond the minimal model**, specifically further-neighbour contributions to (M1) and symmetry-breaking sources for (M3) that would generate non-zero β_j.

Once (a, b, d) are supplied, δC is fixed, the sub-regime is identified, and the downstream consequences in §9 become quantitative predictions rather than parametric relations.

The present paper therefore shows that all remaining freedom in the commitment barrier is reducible to a small finite set of effective parameters in the K = 7 closure manifold. The gap is not a weakness of the threshold framework but a precisely located target.

12.3 Immediate next paper

At this stage, the VERSF programme has reduced the commitment-threshold problem to a single remaining degree of freedom: the microscopic derivation of the K = 7 closure parameters. Under minimal-model assumptions that freedom reduces to just two dimensionless coefficients (a, b) plus the bath-width parameter d.

The immediate next task is therefore sharply defined. The next paper must:

1. Derive μ^2 and κ (equivalently a and b) from the closure sector of the underlying VERSF Lagrangian or its variational equivalent.
2. Derive Δ (equivalently d) from the commitment-event bath spectral structure.
3. Verify whether the minimal-model assumptions (M1)–(M3) hold for the physical manifold, and if not, compute the corrections.
4. Determine from the derived values whether the realized manifold is strictly in the minimal-model splitting-dominated regime, or whether (M3)-breaking corrections activate participation contributions.

Once these are done, the formal machinery developed here converts the microscopic result into quantitative predictions for the commitment barrier correction and its downstream propagation.

In that sense, **the present paper should be read as the threshold-mapping-plus-symmetry-constraint-plus-minimal-model paper, and the next one as the microscopic-closing paper.**

12.4 Medium-term programme

- Extension beyond linear onset into the nonlinear commitment regime.
- Coupling to the Λ-spread sector: explicit construction of f_Λ.
- Precision update to α^{-1} including the $O(r_{max}^4/64)$ correction once r_{max} is known.
- Full treatment of the generalized eigenvalue problem in §10.3 without the kinetic–potential decomposition assumption.
- Accounting for further-neighbour corrections beyond minimal-model (M1), and for (M3)-breaking bath-coupling asymmetries.

Appendix A: Green's function conventions and the dissipative normalization

Retarded Green's function.

$G_a^{ret}(\omega) = 1 / (\omega_a^2 - \omega^2 - i0^+)$.

Dissipative part.

$-(1/\pi) \text{ Im } G_a^{ret}(\omega) = (1/(2\omega_a)) \cdot [\delta(\omega - \omega_a) + \delta(\omega + \omega_a)] \equiv \rho_a(\omega)$.

Dissipative pairing. For states φ, χ in the closure-channel space,

$\langle\varphi, \kappa(C)\rangle$_diss := $C \cdot \Sigma_a \langle\varphi, \chi_a\rangle \cdot \eta_a \cdot \int_0$^∞ $d\omega\, |\hat{f}(\omega)|^2 \cdot \rho_a(\omega)$.

Using $\rho_a(\omega) = (2\omega_a)^{-1} \cdot \delta(\omega - \omega_a)$ on the positive-frequency support of $\hat{f}$:

$\langle\varphi, \kappa(C)\rangle$_diss $= C \cdot \Sigma_a \langle\varphi, \chi_a\rangle \cdot \eta_a \cdot |\hat{f}(\omega_a)|^2 / (2\omega_a)$.

For $\varphi = \psi_j$ this reduces to $R_j(C) = C \cdot K_j$ with K_j the master formula of §3.2. The dissipative normalization is physically the rate of energy absorbed by pair j per unit drive, which is the quantity that must cross R_crit to trigger commitment.

Why the dissipative pairing and not the full retarded response. The full $\langle\psi_j, \kappa(C)\rangle$ is complex: its real (reactive) part encodes frequency-shifted pair polarization, its imaginary (absorptive) part encodes pair excitation and commitment-event production. Only the absorptive part drives irreversible commitment. $R_j(C)$ is defined as the absorptive amplitude so that its crossing R_crit corresponds to a physical rate threshold, not a reactive phase.

Appendix B: Activation-mode choice and antisymmetric sector

B.1 The two natural choices

Within each pair-j sector, two natural unit-norm activation modes exist:

$\psi_j^{sym} = (\chi_{j+} + \chi_{j-}) / \sqrt{2}$, $\psi_j^{anti} = (\chi_{j+} - \chi_{j-}) / \sqrt{2}$.

§§3–8 use $\psi_j = \psi_j^{sym}$ throughout. The antisymmetric alternative probes intra-pair forward/restore imbalance.

B.2 Symmetric mode

$M_{\{j,(j,\pm)\}}^{sym} = 1/\sqrt{2}$ and $M_{\{j,(k,\pm)\}}^{sym} = 0$ for $k \neq j$. K_j^{sym} as computed in §7.

B.3 Antisymmetric mode

$M_{j(j+)}^{anti} = 1/\sqrt{2}$, $M_{j(j-)}^{anti} = -1/\sqrt{2}$, others zero. Then:

$K_j^{anti} = (1/\sqrt{2}) \cdot [\eta_{j+}|\hat{f}(\omega_{j+})|^2/(2\omega_{j+}) - \eta_{j-}|\hat{f}(\omega_{j-})|^2/(2\omega_{j-})]$.

Under (A3) with $\eta_{j+} = \eta_{j-}$, this vanishes identically at leading order. It becomes nonzero only when intra-pair asymmetry $\alpha_j \equiv (\eta_{j+} - \eta_{j-})/\eta_0$ is retained:

$K_j^{anti} \approx K_0 \cdot \alpha_j + O(\alpha_j \cdot \varepsilon\_\Omega)$.

B.4 Effective activation mode

Commitment triggers on the *first* crossing. The effective activation mode is the one minimizing C_j. Comparing:

$C_j^{sym} \propto 1 / K_j^{sym} \propto 1 / [1 + \beta_j - r_j^2/8]$,

$C_j^{anti} \propto 1 / K_j^{anti} \propto 1 / \alpha_j$.

For $\alpha_j \ll \beta_j$ and $\alpha_j \ll r_j^2$, the symmetric mode dominates (smaller C_j^{sym}). This is the generic case when α_j is parametrically smaller than the other pair anomalies — i.e., when forward/restore asymmetry is a higher-order effect in the K = 7 symmetry breaking than pair-level spectral splitting or participation anomaly.

B.5 Consequence for δC

The symmetric-mode analysis of §§7–8 captures the leading δC whenever α_j is of order ε_Ω^2 or smaller. For $\alpha_j \sim O(1)$, the antisymmetric mode can dominate and Part II must be redone in the antisymmetric sector. This regime is outside the Part II scope.

Appendix C: Proofs of Part I theorems

C.1 Theorem 4.1 (Existence)

Linear onset (§2.3) gives $R_j(C) = C \cdot K_j$ with K_j independent of C. $K_j > 0$ under the dissipative definition (Appendix A) because each term in the sum is non-negative and at least one is strictly positive (a fully decoupled pair having $K_j = 0$ is excluded). Therefore $R_j(C)$ is strictly monotonically increasing from $R_j(0) = 0$, and the equation $R_j(C) = R\_crit$ has the unique solution $C_j = R\_crit/K_j$. ■

C.2 Theorem 4.2 (Symmetric closure, iff)

(⇐) If $K_j = K_0$ for all j, then $C_j = R\_crit/K_0 = 3/8$ for all j, so $\max_j C_j = 3/8$ and $\delta C = 0$.

(⇒) Suppose $\delta C = 0$. Then $\max_j C_j = 3/8$, i.e., $C_j \leq 3/8$ for all j with equality for at least one. But $C_j = R\_crit/K_j = (3/8)\cdot K_0/K_j$. The inequality $C_j \leq 3/8$ gives $K_j \geq K_0$, and equality in $C_j = 3/8$ gives $K_j = K_0$.

If $K_j > K_0$ for any j, then $C_j < 3/8$, so $\max_j C_j = (3/8)\cdot K_0/K_{min}$ where $K_{min} = \min_j K_j$. For $\max_j C_j$ to equal exactly 3/8, we need $K_{min} = K_0$. Combined with $K_j \geq K_0$ for all j, this forces $K_j = K_0$ for all j. ■

C.3 Theorem 5.1 (Conditional positivity)

$\delta C = (3/8)\cdot K_0/K_{min} - 3/8 = (3/8)\cdot(K_0 - K_{min})/K_{min}$.

(⇒) $\delta C \geq 0 \Leftrightarrow (K_0 - K_{min})/K_{min} \geq 0 \Leftrightarrow K_{min} \leq K_0$ (using $K_{min} > 0$).

Strict form. $\delta C > 0 \Leftrightarrow K_{min} < K_0$ strictly. ■

C.4 Theorem 5.3 (Tight bound)

From $C_j = (3/8)\cdot K_0/K_j = (3/8)/(1 + \delta K_j/K_0)$ with $\delta K_j = K_j - K_0$:

- If $\delta K_j > 0$: $C_j < 3/8$, so $\delta_j < 0$ with $|\delta_j| = (3/8) \cdot (\delta K_j/K_0) / (1 + \delta K_j/K_0)$.
- If $\delta K_j < 0$: $C_j > 3/8$, so $\delta_j > 0$ with $\delta_j = (3/8) \cdot (|\delta K_j|/K_0) / (1 - |\delta K_j|/K_0)$.

Each is the exact non-linear expression of §5.2. The barrier picks $\delta C = \max_j \delta_j$; since only $\delta K_j < 0$ contributes positive δ_j, δC is determined by $\max_{\{j: \delta K_j < 0\}} |\delta K_j|/K_0 = \varepsilon^-$, giving $\delta C = (3/8) \cdot \varepsilon^-/(1 - \varepsilon^-)$. If no pair has $\delta K_j < 0$, then $\delta C = \max_j \delta_j = \max_j$ (negative) $= -(3/8) \cdot \varepsilon^+/(1 + \varepsilon^+)$. ■

C.5 Theorem 5.4 (Scale invariance)

Under $K_j \to \lambda K_j$ for all j: $K_0 \to \lambda K_0$ (since K_0 is a particular K_j). $R\_crit = (3/8)\cdot K_0 \to \lambda\cdot R\_crit$. Therefore $C_j = R\_crit/K_j \to (\lambda R\_crit)/(\lambda K_j) = C_j$. ■

Appendix D: Illustrative seven-pair computation

D.1 Purpose and status

Appendix D should be read as a pipeline-consistency demonstration rather than as a physical prediction. Its value is threefold:

1. It shows that the formal machinery is **numerically stable** under percent-level perturbations.
2. It verifies that the exact nonlinear bound of Theorem 5.3 is not merely qualitative but **operationally sharp** on a concrete example.
3. When cross-compared with the α-sector consistency bound, it illustrates how a given numerical value of δC can be assigned to different upstream sub-regimes, thereby showing that **multi-sector consistency genuinely adds interpretive power**.

The appendix therefore serves as a worked demonstration of method, **not as a substitute for the spectral paper still required by §12.2**.

We use the single per-pair input

$\varepsilon_j \equiv \delta K_j / K_0 = \beta_j - r_j^2 / 8$ (from §7.2)

and treat $\{\varepsilon_j\}$ as the primary variable, without committing to a specific decomposition into β_j and r_j^2 contributions. Physical interpretation of the toy result in terms of the β/r split is deferred to §D.6.

D.2 Narrow-band approximation (reminder from Part II)

Under (A1)–(A3) the master formula of §3.2 reduces to

$K_j \propto \Sigma_a M_{ja} \cdot \eta_a \cdot G_a(\omega_0) \cdot \hat{f}(\omega_0)$

at leading order (the common envelope factor absorbs into normalization). Writing

$K_j = K_0 (1 + \varepsilon_j)$,

with K_0 the fixed symmetric-limit reference of §3.2. The traceless constraint $(1/7) \Sigma_j \varepsilon_j = 0$ corresponds to mean-preserving perturbations; it is natural for an illustrative anisotropy but is not required by the formalism, and is adopted in the toy spectrum below for simplicity.

D.3 Pair thresholds in the toy model

From §3.3 with $R\_crit = (3/8) K_0$ (§3.4):

$C_j = R\_crit / K_j = (3/8) \cdot K_0 / K_j = (3/8) / (1 + \varepsilon_j)$.

For $|\varepsilon_j| \ll 1$:

$\delta_j = C_j - 3/8 = (3/8) \cdot [1/(1+\varepsilon_j) - 1] \approx -(3/8) \varepsilon_j + O(\varepsilon_j^2)$.

The barrier is

$\delta C = \max_j \delta_j$,

controlled by the most negative ε_j (the pair with smallest K_j — the hardest to activate).

D.4 Concrete seven-pair example

Take the weakly perturbed spectrum

$(\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5, \varepsilon_6, \varepsilon_7) = (+0.020, +0.012, +0.008, 0, -0.006, -0.012, -0.022)$.

This spectrum is:

- Small (all $|\varepsilon_j| \leq 0.022$).
- Traceless: $\Sigma_j \varepsilon_j = 0$ exactly.
- Monotonic across the seven pairs — **recognizable as approximately the n = 1 Fourier harmonic of the Part IV decomposition** (§10.2), i.e. $\varepsilon_j \approx \varepsilon \cos(2\pi j/7 + \varphi)$ with $\varepsilon \approx 0.022$. This is the lowest-order symmetry-admissible perturbation around the symmetric K = 7 baseline, and is the natural choice for an illustrative single-mode anisotropy.

Computing $C_j = (3/8)/(1 + \varepsilon_j)$ with $3/8 = 0.375$:

j	ε_j	$1 + \varepsilon_j$	C_j
1	+0.020	1.020	0.3676
2	+0.012	1.012	0.3706
3	+0.008	1.008	0.3720
4	0	1.000	0.3750
5	−0.006	0.994	0.3773
6	−0.012	0.988	0.3796
7	−0.022	0.978	0.3834

The maximum is

$C^* = \max_j C_j = C_7 \approx 0.3834$,

hence

$\boldsymbol{\delta C \approx 0.3834 - 0.375 = 8.4 \times 10^{-3}}$**.**

D.5 Cross-check with Theorem 5.3

Theorem 5.3 gives the exact non-linear bound

$\delta C = (3/8) \cdot \varepsilon^- / (1 - \varepsilon^-)$, $\varepsilon^- = \max\_\{j: \varepsilon_j < 0\} |\varepsilon_j|$.

For the toy spectrum, $\varepsilon^- = 0.022$ (attained at $j = 7$). Substituting:

$\delta C = 0.375 \cdot 0.022 / (1 - 0.022) = 0.375 \cdot 0.022 / 0.978 = 0.375 \cdot 0.02249\ldots = 8.43 \times 10^{-3}$.

This matches the direct computation of §D.4 to the fourth decimal. The toy spectrum **saturates the Theorem 5.3 bound**, which is expected: whenever the binding pair (the one setting δC) has $\varepsilon_j < 0$, the tight bound is not a loose envelope but the exact value. This confirms the consistency of the Part I and Part II machinery on the toy case.

Note on relation to the §10.10 minimal-model prediction. The toy spectrum used here is monotonic across $j = 1, \ldots, 7$. This is *not* the generic prediction of the minimal closure model of §10.10, which has the symmetry $\lambda_n = \lambda_{7-n}$ and therefore predicts three doubly-degenerate levels rather than a monotonic sequence. The toy is a deliberately simple single-harmonic illustration chosen to exercise the Part I / Part II pipeline end-to-end with an easily-interpretable ε^-; it does not claim to match the minimal-model degeneracy structure, and its role in the paper is pipeline demonstration rather than spectrum prediction. A spectrum compatible with the minimal-model structure would exhibit $\varepsilon_j = \varepsilon_{7-j}$, which the next paper (microscopic derivation of a, b) will produce quantitatively.

D.6 Interpretation and sub-regime identification

Three observations:

1. Stability of the 3/8 baseline. Weak pair asymmetry does not destroy 3/8; it perturbs it slightly. The symmetric value remains the correct structural baseline, with δC entering as a small controlled correction:

$\delta C / (3/8) = 8.4 \times 10^{-3} / 0.375 \approx 2.24\%$.

2. The barrier is set by the hardest pair. C* comes from pair 7, the pair with smallest K_j (largest $|\varepsilon^-|$). This confirms the §5 physical picture: commitment requires all pairs to activate, so the weakest pair sets the cost.

3. Sub-regime identification via the α-sector bound. The toy $\delta C = 8.4 \times 10^{-3}$ corresponds to $\delta C/C_0 \approx 0.022$. In a purely splitting-dominated interpretation (β_j negligible, all deviation from splitting alone), this would require $r_{max} \approx 0.42$ — **exceeding** the α-sector bound $r_{max} \lesssim 0.15$ established in §9.3. The toy spectrum is therefore *not* in the splitting-dominated regime.

A consistent reading of the toy example requires β_j contributions to dominate. Specifically, pair 7 must have $\beta_7 \approx -0.022$ with r_7^2 negligible, and the other pairs likewise dominated by β_j. The toy spectrum thus illustrates the **participation-dominated sub-regime** of §8.1, in which δC is driven by channel-strength anomalies rather than by spectral displacement.

This is a useful lesson: the same numerical δC can be realized in multiple upstream sub-regimes, and consistency with other sectors of VERSF (α, Λ, Φ_c) identifies which sub-regime is physically realized. The open problem of §12.2 — the microscopic determination of the minimal-model parameters (a, b, d) from the closure Hamiltonian — is precisely the question of which sub-regime the actual manifold occupies.

D.7 General structural bound

The result of §D.5 is an instance of the Part I Theorem 5.3, not an independent derivation. Restated for reference:

If $|\varepsilon_j| \leq \varepsilon_{max}$, then $\delta C \leq (3/8) \cdot \varepsilon_{max} / (1 - \varepsilon_{max})$.

For small ε_{max}: $\delta C \lesssim (3/8) \cdot \varepsilon_{max}$.

For $\varepsilon_{max} = 0.022$ this gives $\delta C \lesssim 8.4 \times 10^{-3}$, saturated by the toy spectrum. If the K = 7 closure manifold is approximately symmetric at the 1–3% level, δC is automatically constrained to the percent level — regardless of which sub-regime drives it.

D.8 Summary

The toy pair spectrum demonstrates that the formal pair-resolved framework is numerically stable and that weak asymmetry produces a small, controlled positive shift in the commitment barrier. In the specific seven-pair example considered here, the symmetric baseline C = 3/8 is shifted to C* ≈ 0.3834, corresponding to $\delta C \approx 8.4 \times 10^{-3}$, or a relative correction of approximately 2.2%. The toy value exactly saturates the Part I tight bound (Theorem 5.3), confirming the internal consistency of the structural and computational machinery. Cross-referencing with the α-sector bound identifies the toy example as participation-dominated, illustrating how multi-sector consistency sharpens interpretation of any given δC value.

This supports the interpretation of 3/8 as the exact symmetric-limit value and of δC as a computable spectral correction rather than a free parameter. What remains is the upstream question of what values (a, b, d) the minimal K = 7 closure model takes, and what corrections beyond the minimal model apply — the problem stated in §12.2 and reserved for the spectral-closing paper.

Appendix E: Conditional constraints on the minimal-model parameters

E.1 Scope and objective

Section 10.10 shows that under the minimal-model assumptions (M1)–(M3), the K = 7 closure spectrum is uniquely determined up to two effective parameters (μ^2, κ), together with the bath-width parameter Δ. In dimensionless form,

$\mu^2 = a\,\xi^{-2}$, $\kappa = b\,\xi^{-2}$, $\Delta = d\,\xi^{-1}$,

so the minimal model is specified by the three dimensionless parameters (a, b, d).

This appendix does *not* attempt to derive these parameters microscopically from a complete closure Hamiltonian — that remains the open problem of §12.2. Instead, we compile the scale and cross-sector constraints already scattered across §§10.4, 10.10.8, and 9.3 into a single consolidated statement, and add two structural arguments that further narrow the admissible parameter space.

The results of this appendix are conditional, not first-principles. They rely on the minimal-model assumptions, the kinetic–potential identification of §10.3, the α-sector constraint of §9.3, and the plausibility arguments introduced below.

E.2 Closure-scale normalization: a, b ~ O(1)

From §10.4, the only intrinsic length scale in the closure sector is the coherence scale ξ. The commitment barrier is set by $\hbar c/\xi$ and the κ-field mass satisfies $m_\kappa^2 \sim \xi^{-2}$ (§2.1). Characteristic closure-sector frequencies must therefore scale as

$\Omega \sim c/\xi$.

Comparing with the minimal-model spectrum $\Omega_n^2 = \xi^{-2}[a + 4b \sin^2(\pi n/7)]$, this means the dimensionless parameters (a, b) cannot be arbitrarily large or small. They are anchored to the same closure scale that governs the commitment barrier itself, giving

a ~ O(1), b ~ O(1).

This is not a precise numerical claim — it allows a, b anywhere in the range roughly 0.1 to 10 — but it excludes the extreme hierarchies $a \gg b$, $a \ll b$, or either parameter vastly larger or smaller than unity.

E.3 Relative magnitude of a and b: physical relevance argument

Within the O(1) range fixed by §E.2, the ratio b/a still varies. Two limiting cases illustrate what each regime implies:

- **Weak coupling, $b \ll a$**: the spectrum is nearly flat, $\Omega_n \approx \xi^{-1}\sqrt{a}$ independent of n, and the nontrivial pair splittings are vanishing. In this limit $\delta C \to 0$ automatically.
- **Strong coupling, $b \gg a$**: the spectrum is dominated by the sinusoidal term, producing large fractional separations between harmonic levels and generically driving r_n to values well above the α-bound of §9.3.

If δC is to be physically relevant in downstream sectors — contributing at all to Φ_c correction, to Λ-spread, or to the α-correction at its natural order — then neither limiting case is consistent with the VERSF programme's broader picture. The weak-coupling limit gives $\delta C = 0$, leaving no downstream effect. The strong-coupling limit violates the α-bound.

This is a physical-plausibility argument, not a structural theorem. **If δC is to be physically nontrivial**, this suggests the closure spectrum lies in the regime where both local stiffness a and inter-pair coupling b contribute non-negligibly:

b ~ a (conditional on δC being physically relevant).

The conditionality is real and should be stated explicitly: if the physical manifold happens to sit at $b \ll a$, the entire threshold-splitting programme becomes trivial but remains self-consistent. The present paper studies the nontrivial regime where threshold-splitting is a live quantity.

E.4 Cross-sector constraint: $d \gtrsim 14$

From §10.10.8, the α-sector constraint $r_{max} \lesssim 0.15$ becomes

$$d^2 \gtrsim 44.4 \cdot (a + 3.8019\, b).$$

Committing to the a ~ b ~ 1 range of §E.2–§E.3, the numerical coefficient evaluates to

$$d^2 \gtrsim 44.4 \cdot (1 + 3.8019 \cdot 1) \approx 213,$$

giving

$d \gtrsim 14$ (from α-consistency at the natural-parameter scale).

Equivalently, **$\Delta \gtrsim 14\,\xi^{-1}$**, so the commitment-event bath bandwidth must exceed the intrinsic closure scale by roughly an order of magnitude. Whether the κ-field dynamics naturally produce such a separation of scales is itself a question for the spectral-closing paper, which must derive Δ from the commitment-event bath structure; the present appendix only observes that consistency with the α-sector requires this hierarchy.

E.5 Constrained prediction for δC

Combining §§E.2–E.4, the minimal-model threshold correction

$\delta C = (3/32d^2) \cdot (a + 3.8019\,b)$

is bounded from above at the α-saturating edge by

$\delta C \lesssim (3/32) \cdot r_{max}^2 \lesssim (3/32) \cdot 0.0225 \approx 2.1 \times 10^{-3}$.

Lower bounds depend on how much room the actual parameters have inside the $a \sim b \sim 1$ envelope. Taking the realistic range $b/a \in [0.3, 3]$ and d at the α-bound:

$\delta C \in [\,{\sim}10^{-4}\,, {\sim}2 \times 10^{-3}\,]$ (conditional on §§E.2–E.4).

Thus, under the combined structural assumptions of the minimal closure model and α-sector consistency, **δC is confined to the sub-percent range within the minimal-model and α-consistent regime.**

This is consistent with the illustrative Appendix D result (2.2% for a non-minimal-model spectrum) but provides a tighter bound once the minimal-model constraints are imposed: a genuinely minimal-model-compliant spectrum cannot exceed $\delta C \approx 0.2\%$.

E.6 Interpretation: structural smallness of δC

Three interpretive consequences:

1. **The minimal-model parameters are not free.** §E.2 fixes their scale ($a, b \sim 1$), §E.3 fixes their relative magnitude ($b \sim a$, conditional on δC being physically relevant), and §E.4 fixes the bath parameter relative to them ($d \gtrsim 14$).
2. **The bath scale is constrained by cross-sector consistency.** The α-sector forces $\Delta \gtrsim 14\,\xi^{-1}$, a nontrivial hierarchy that the spectral-closing paper must independently reproduce.
3. **The smallness of δC is traceable to a scale hierarchy, not local fine-tuning.** The observation is: $\delta C \sim 1/d^2 \sim (\xi/\Delta)^2$. Given the α-imposed $\Delta \gg \xi^{-1}$ hierarchy, δC is automatically small. The "fine-tuning" is therefore pushed upstream to the Δ/ξ ratio — which itself is fixed by cross-sector consistency rather than by free-parameter adjustment.

Whether Δ/ξ is "naturally" large in the microscopic theory is a genuine question for the spectral-closing paper, not something resolved here.

Thus the smallness of δC is not a result of parameter tuning but of a hierarchy enforced by cross-sector consistency.

Taken together, the apparent freedom in (a, b, d) is already narrowed to a small corner of parameter space by internal VERSF consistency alone — without any appeal to microscopic first-principles derivation.

E.7 Status and limitations

The results of this appendix are conditional and rely on:

- the minimal-model assumptions (M1)–(M3) of §10.10;
- the identification $\Omega_n^2 = \lambda_n$ (kinetic–potential decomposition of §10.3);
- the α-sector constraint of §9.3;
- the physical-relevance argument of §E.3, which requires δC to be non-trivial.

They do not constitute a first-principles derivation of (a, b, d). In particular:

- μ^2 and κ are not derived from a microscopic closure Hamiltonian.
- Δ is not derived from the κ-field or commitment-event dynamics.
- Corrections beyond nearest-neighbour coupling (M1) are not included.
- (M3)-violating bath-coupling asymmetries that would reintroduce $\beta_j \neq 0$ are not analyzed.

The role of this appendix is to **narrow the admissible parameter space**, not to close it.

E.8 Outlook

With the constraints compiled here, the spectral-closing paper's task is sharpened to:

- derive μ^2, κ, and Δ from the microscopic closure dynamics;
- verify (or correct) the natural-scale estimates $a \sim b \sim 1$;
- independently produce the hierarchy $\Delta \gtrsim 14\ \xi^{-1}$ that the α-sector requires;
- account for corrections beyond (M1) and (M3).

Once these are done, the closure spectrum is fixed, δC is determined to within the sub-percent bound of §E.5, and all downstream consequences become quantitative predictions.