Defects, Local Spectral Distortion, and Candidate Curvature Indicators in the K = 7 Refinement Universality Class

Localized Admissibility Perturbations, Spatially Varying Spectral Gap, and Trapped Incoherent Modes in VERSF

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

This paper is part of a research programme proposing that physical reality — space, time, fields, perhaps eventually particles — emerges from a more fundamental layer of discrete mathematical structure called the *refinement substrate*. The substrate is built from a simple primitive: at every point of an underlying lattice, there is a small set of 7 internal states (one "hub" state and six "boundary" states arranged around it like the spokes of a wheel), together with a fixed canonical rule for how probability flows between them — the precise rule whose consequences when *locally modified* are the subject of this paper. Earlier papers in this programme have shown that, at large scales, this discrete machinery produces *smooth* space-and-time-like structure, with the smoothness controlled by a single technical number called the *spectral gap* — roughly, a measure of how quickly the system settles toward equilibrium. A large spectral gap means rapid settling, which translates into smooth large-scale structure; a small gap means slow settling and rougher structure.

The Stage VII paper, which preceded this one, established that the substrate is *robust*: small adjustments to the wheel's transition rules (within mathematical limits that keep the rules valid as probability distributions) leave the smooth large-scale structure essentially intact, and the spectral gap remains positive. The wheel sits inside a wide family of similar structures that all behave the same way at large scales, rather than at a fragile, finely-tuned point. The technical reasons *why* this family is wide were the substance of Stage VII; the only fact we need from it here is that the family exists and is robust.

This naturally raises the next question: what happens when the substrate's mathematical structure fails *in a small local region* rather than globally?

In ordinary materials, local imperfections produce real and well-understood physical effects. A missing atom in a crystal traps sound waves; an impurity in a semiconductor binds electrons into discrete localized states; a defect in a magnetic chain hosts a localized magnetic excitation. The underlying physical mechanism is the same in all these cases: a region where the local rules differ from the bulk creates a *bound state* — a self-sustaining pattern that lives inside the defect and decays away outside it.

This paper investigates whether the same kind of phenomenon occurs in the refinement substrate. We consider *localized defects* — small regions where the wheel's transition rules are modified, with the standard rules holding everywhere else — and ask what such defects do to the emergent smooth structure.

We find three primary structural phenomena (and a fourth conjectural extension), with the epistemic status of each marked explicitly so the reader can see which claims are theorems and which are speculation:

- **Proven (rigorous mathematical theorems).** A localized defect produces a measurable "coherence weakness" in the surrounding region, captured by a number we call the *local spectral gap* ε_gap(x) that varies from point to point on the substrate. This local number varies smoothly with position whenever the defect itself varies smoothly. The local roughness of the emergent smooth structure depends inversely on the local spectral gap — where coherence is weak, the emergent geometry is rougher.
- **Conditional (depends on a hypothesis we cannot directly verify with the present mathematical tools).** If a defect is strong enough to *trap* a refinement mode — a self-sustaining pattern of internal-state activity that does not propagate away from the defect — then that trapped mode persists indefinitely, behaving like a stable localized excitation in the emergent geometry. We give the mathematical conditions for trapping and a finite-dimensional test to detect part of them; verifying whether any specific defect actually satisfies the full conditions requires further machinery (sketched in §12.2) which is left to a future paper.
- **Heuristic (a structural ansatz, motivated but not derived).** The spatial variation of the coherence weakness behaves *like* a curvature in the emergent geometry. We propose a specific scalar quantity — the discrete Laplacian of ε_gap(x) — as a candidate "curvature indicator" R(x). The right way to read R(x) is as analogous to an *order parameter* in condensed-matter physics (like magnetisation, or the BCS superconducting gap): a single scalar diagnosing local departures from a uniform coherent state, without committing to a full tensorial geometric theory. We do *not* derive curvature in the differential-geometric sense; that is an open problem (§12.1).
- **Conjectural (a structural analogy with no derivation, recorded only to motivate future investigation).** Trapped refinement modes might be the substrate-level precursors of *matter* in the emergent geometry. This is a suggestive structural analogy — bound states play this role in condensed-matter physics — not a derivation of any specific particle content. No phenomenology is claimed. The reader should treat this strictly as a conjecture flagged for future programme work, not as a result of the present paper.

The overall picture is: smooth emergent geometry corresponds to *uniform* coherence in the substrate; candidate curvature corresponds to *spatial variation* of coherence; trapped modes correspond to *isolated regions of coherence failure*. A further consequence of locally reduced coherence is slower local relaxation to equilibrium, which sits as a quantitative corollary of the first phenomenon rather than as a separate primary phenomenon (it is one of several mathematical *functionals* of the local coherence). All these effects are different ways of seeing the same underlying quantity: the local spectral gap ε_gap(x).

What this paper does *not* do: it does not derive Einstein's equations, produce particle phenomenology, recover any specific known curvature formula, or quantize the defect spectrum. The contribution is structural — identifying localized defects as the natural next layer of the substrate programme, providing the mathematical machinery for studying them, and clearly separating what has been proven from what remains conjectural or open.

Abstract

The previous Stage VII paper established stability and universality of the $K = 7$ refinement fixed point under admissibility-preserving perturbations. The present paper extends the analysis to *spatially localized* perturbations.

We define a defect operator as a Stage VII admissibility perturbation supported in a finite substrate region:

$\hat{T}_x = \hat{T} + \Delta T_x$, $\Delta T_x = 0$ outside $B_r(x)$,

with ΔT_x satisfying the Stage VII admissibility conditions P1–P3. The principal results are:

- **Local spectral gap (Proposition 3.1, Theorem 3.2).** For sufficiently weak defects, the local refinement operator $\hat{T}_x$ has a well-defined local spectral gap $\varepsilon_{\text{gap}}(x) := 1 - \max_{\{\lambda \neq 1\}} |\lambda(\hat{T}_x)|$, satisfying the pointwise bound $\varepsilon_{\text{gap}}(x) \geq \frac{1}{2} - C_0 \cdot \|\Delta T_x\|_{\text{op}}$ with $C_0 := \sqrt{7/4}$ the Stage VII Bauer–Fike constant, and the Lipschitz regularity bound $|\varepsilon_{\text{gap}}(x) - \varepsilon_{\text{gap}}(y)| \leq C_0 \cdot K \cdot d_X(x, y)$ when the defect profile is Lipschitz in the substrate metric with constant K.
- **Local continuum Lipschitz constant (Proposition 4.1, Corollary 4.2).** The Stage V bound applied to the local operator gives $K_\infty(x) = L_\Phi \cdot L(x) \cdot A^2 / (A_- \cdot \varepsilon_{\text{gap}}(x))$, spatially varying with the gap, with gradient $\nabla K_\infty / K_\infty = -\nabla\varepsilon_{\text{gap}} / \varepsilon_{\text{gap}}$ (modulo ∇L corrections).
- **Candidate curvature indicator (heuristic Definition 5.1).** Motivated by the structure of the local Lipschitz bound, we propose the scalar field $R(x) := \nabla^2\varepsilon_{\text{gap}}(x)$ (analyst convention for the discrete Laplacian) as a candidate curvature-like indicator. The "candidate" label is meant strictly: we do not derive this from a tensorial geometric theory, we observe that its sign and spatial profile track expected curvature behaviour in worked examples, and we list its limitations explicitly.
- **Trapped-mode existence (Theorem 6.1, conditional and global; Proposition 6.2, variational test).** If the global position-indexed refinement operator on the full substrate develops an eigenvector spatially localized around the defect region with eigenvalue λ_{trap} satisfying $\frac{1}{2} < |\lambda_{\text{trap}}| < 1$, then this trapped mode persists under refinement and dominates the local late-time dynamics. The condition is on the *global* operator, not merely the local operator $\hat{T}_{x_0}$; for reversibility-preserving defects, the necessary eigenvalue condition $\lambda_2(\mathbf{T}) > \frac{1}{2}$ admits a finite-dimensional variational test via Rayleigh quotients of localized test functions. The full existence question (including eigenvector localization) is governed by a Birman–Schwinger-type criterion we do not develop here.

- **Defect classification (Section 7).** We classify defects by the *spatial localization of their perturbed eigenvectors* into *bounded* (eigenvectors remain delocalised), *gap-suppressing* (an eigenvector becomes localized, satisfying Theorem 6.1's hypothesis), and *gap-closing* (admissibility lost) types. §7.5 defines four critical-strength thresholds (α_admiss, α_smallpert, α_trap, α_c) and gives the phase diagram for the §9 examples; the classification differs from the Stage VII RG-style trichotomy in being based on eigenvector behaviour rather than eigenvalue magnitude.
- **Local entropy persistence (Corollary 8.1).** For reversibility-preserving defects, the local χ^2-divergence at position x satisfies s(x, n) ≤ (1 − ε_gap(x))^{2n} · s(x, 0); a 10% local-gap suppression increases entropy retention by a factor $(1.21)^n$ per refinement step, with the doubled exponent reflecting the χ^2-vs-L^2 relationship.
- **Worked examples (Section 9).** A localized boundary self-loop / cyclic defect produces explicit and numerically computed local-gap suppression matching the Stage VII §10.2 numerics (slope ≈ 0.226, conservatism factor ≈ 6.25); the candidate-curvature spatial profile is positive at the defect core and negative on the immediate neighbour ring. A contrasting hub-coupling defect (§9.6) leaves the local gap *exactly* invariant across the entire admissible range, supplying the localized analogue of Stage VII §10.1's symmetry-protected silent perturbations.
- **Defect-Coherence Principle (§13.1).** Gap reduction, candidate curvature, and trapped-mode persistence are three primary phenomena, jointly producing four functionals of ε_gap(x) (continuum roughening, candidate curvature, entropy retention, trapped-mode persistence) — different mathematical extractions from the same underlying spectral-gap field. This is a calibrated structural unification, not a derivational one; entropy retention sits as a corollary of gap reduction rather than as a fourth primary phenomenon.

We do not derive Einstein-style field equations, particle phenomenology, or any specific physical content. The contribution is the operator-theoretic machinery for treating local refinement failure within the K = 7 universality class, the explicit identification of trapped incoherent modes as candidate matter-like substrate excitations, and the conditional / heuristic / conjectural separation of claims that further stages must address.

Contents

---

1. Introduction

The Stage VII paper, *Stability and Universality of the K = 7 Refinement Fixed Point*, established that the canonical wheel operator $\hat{T}$ sits inside an open universality class $\mathcal{C}_{K=7}$ of admissibility-restricted refinement operators. The class is characterised by the structural mechanism of symmetry-protected spectral invariance on a positive-dimensional subspace of conjugation-antisymmetric perturbations, and by operator-norm conservatism approximately six-fold below the worst-case Bauer–Fike bound on the symmetry-breaking complement.

That paper considered only *globally uniform* perturbations:

$$\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T,$$

where ΔT acts everywhere on the substrate. The present paper considers *localized* perturbations:

$$\hat{T}_x = \hat{T} + \Delta T_x,$$

where ΔT_x has finite spatial support in a neighbourhood of substrate location x and vanishes elsewhere.

The motivation is structural. The Stage V continuum Lipschitz bound

$$K_\infty = L_\Phi \cdot L \cdot A^2 / (A_- \cdot \varepsilon_{gap})$$

involves a single global spectral gap ε_gap. In a uniform substrate this is one number; in a substrate with local defects, every quantity in the bound becomes potentially spatially varying. We therefore ask: what is the structural content of a *local* version of the Stage V bound, and what local phenomena does it predict?

Three structural phenomena emerge naturally:

1. **Local gap reduction.** The Stage VII spectral stability theorem applied to the local operator $\hat{T}$_x bounds the local gap reduction in terms of the local defect strength. The gap becomes a *field* ε_gap(x) rather than a constant.
2. **Spatial gap variation as a curvature precursor.** The local Lipschitz constant K_∞(x) varies inversely with ε_gap(x). Spatial derivatives of ε_gap(x) thus naturally enter any local-geometry expression. We propose, *heuristically*, that ∇^2ε_gap(x) functions as a scalar curvature-like indicator. This is a structural ansatz, not a derived tensor; we discuss its limitations in §5.
3. **Trapped incoherent modes.** When a defect is sufficiently strong, the position-indexed global refinement operator may develop an eigenvector spatially localized inside the defect region with eigenvalue larger in modulus than the bulk subdominant eigenvalue ½. Such modes do not dissipate globally and constitute a structural candidate for substrate-level matter-like excitations.

The three phenomena are connected. A defect that reduces the local gap weakly is a candidate for smooth curvature variation; a defect that reduces it strongly enough to create a trapped mode is a candidate for matter-like localization. The transition between these regimes is governed by an explicit spectral criterion (Theorem 6.1).

What this paper does *not* do:

- It does not derive Einstein's equations or any specific gravitational dynamics.
- It does not produce particle phenomenology, gauge structure, or quantum statistics.
- It does not quantize the defect spectrum.
- It does not derive the candidate curvature indicator R(x) from a tensorial geometric theory.

These are deferred to later stages. The present paper's contribution is the operator-theoretic machinery for studying refinement defects within the K = 7 universality class, the conditional and heuristic results that machinery makes available, and the explicit identification of which further structural input is needed for each downstream claim.

2. Localized Defect Operators

2.1 Substrate Notation

We work with the Stage VI canonical $K = 7$ substrate. Following Stage VI §3, the refinement substrate is a regular lattice X of *substrate positions*, each carrying a copy of the closure catalogue $\mathcal{K} = \{\kappa_h, \kappa_{b_1}, \ldots, \kappa_{b_6}\}$. The full state space is $\mathcal{K} \times X$ (closure label, substrate position).

At each position $x \in X$, the canonical refinement operator is the wheel operator $\hat{T}$ defined in Stage VI:

$$\hat{T}_{\text{canonical}}(x) = \hat{T} \text{ for every } x \in X.$$

This is the *uniform-substrate* assumption: the same $K = 7$ wheel operator acts identically at every substrate position.

2.2 Defect Operators

A *localized defect at position $x_0 \in X$ with radius r* is a modification of the refinement operator in a neighbourhood $B_r(x_0) := \{ x \in X : d_X(x, x_0) \leq r \}$:

$$\hat{T}_x = \hat{T} + \Delta T_x,$$

where ΔT_x is a one-parameter family of 7×7 matrices satisfying:

- **D1 (Stage VII admissibility).** For every $x \in X$, ΔT_x satisfies the Stage VII conditions P1 (stochasticity preservation), P2 (wheel-admissible support), and P3 (positivity within an admissible range).
- **D2 (finite support).** $\Delta T_x = 0$ for $x \notin B_r(x_0)$.
- **D3 (bulk recovery).** At positions x outside the defect region, $\hat{T}_x = \hat{T}$, recovering the canonical bulk substrate.

A defect is *weak* if $\|\Delta T_x\|_{op} \leq 1/(2 \cdot C_0)$ uniformly across $B_r(x_0)$, where $C_0 := \sqrt{7/4}$ is the Stage VII Bauer–Fike constant for the canonical wheel. It is *strong* otherwise.

2.3 Notation Disclaimer

Throughout this paper we write the local operator as $\hat{T}_x$ to emphasise its position dependence. This is *not* the same notation as Stage VII's $\hat{T}_\eta$ (perturbed by uniform strength η); the subscript here is a substrate position, not a perturbation parameter. When discussing the strength of a localized defect, we use a separate parameter α and write $\hat{T}_x = \hat{T} + \alpha \cdot \Delta T_{x_0,r}$ for a defect of strength α at position x_0 with profile $\Delta T_{x_0,r}$.

2.4 The Local-vs-Global Distinction

Stage VII considered the case where every position x has the same perturbed operator $\hat{T} + \eta \cdot \Delta T$. The present paper considers the case where positions inside the defect region have $\hat{T} + \Delta T_x$ with $\Delta T_x \neq 0$, while positions outside have $\hat{T}$ exactly. The two situations have different spectral content:

- In Stage VII, every local operator is identical, so there is one global spectrum and one global gap.
- Here, each position has its own local operator $\hat{T}_x$ and its own local spectrum $\mathrm{spec}(\hat{T}_x)$, giving a *field* of spectra over X.

The local field $\varepsilon_{gap}(x)$ introduced in §3 is well-defined precisely because each $\hat{T}_x$ is a perturbation of the same bulk $\hat{T}$, so Stage VII spectral stability applies position by position. The *global* refinement operator on the full state space $\mathcal{K} \times X$ (which couples local operators via substrate spatial structure) is distinct from any individual $\hat{T}_x$ and is discussed separately in §6.

---

3. The Local Spectral Gap Field

3.1 Definition

Given the local operator $\hat{T}_x$ for each position $x \in X$, the *local spectral gap* at position x is

$$\varepsilon_{gap}(x) := 1 - \max_{\{\lambda \in \mathrm{spec}(\hat{T}_x),\ \lambda \neq 1\}} |\lambda|.$$

In the uniform-substrate vacuum ($\Delta T_x = 0$ everywhere), this reduces to the Stage VI canonical value:

$$\varepsilon_{gap}(x) = \varepsilon_{gap}^{(0)} = ½ \text{ for all } x \in X.$$

In a substrate with defects, $\varepsilon_{gap}(x)$ varies spatially: equal to ½ outside the defect support, and generically perturbed inside, with $\varepsilon_{gap}(x) = ½$ in the symmetry-protected case of §9.6. Proposition 3.1 gives a lower bound on $\varepsilon_{gap}(x)$; whether a corresponding upper bound $\varepsilon_{gap}(x) \leq ½$ holds in general (i.e., whether localized defects always depress rather than enhance the local gap) is not established here and depends on the spectral direction of the perturbation.

Terminological note. Throughout the paper we refer to $\varepsilon_{gap}(x)$ as a "field" on the substrate. This is a *discrete scalar field* on the substrate lattice X, with derivatives $\nabla\varepsilon_{gap}$ and $\nabla^2\varepsilon_{gap}$ understood as finite differences in the analyst convention (§5.2). It is not a continuum-differentiable field; the discrete lattice X is the underlying ontology, and any continuum interpretation of ε_{gap} as a smooth function is a *limiting* operation that requires either taking a substrate-spacing limit $a \to 0$ (Stage IV would specify whether and how) or interpolating between lattice values. Statements about "the gradient" or "the Laplacian" of ε_{gap} throughout the paper refer to the discrete versions; statements about "continuous variation" mean the discrete-Lipschitz continuity of Theorem 3.2, not differentiability in the continuum sense. Readers should keep the substrate ontology in mind when translating any of the present claims into continuum-geometric language.

3.2 Stage VII Spectral Stability Applied Locally

The Stage VII Theorem 5.1 (spectral gap stability) was stated for the uniform perturbation $\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T$. The same argument applies *verbatim* to each local operator $\hat{T}_x$, since $\hat{T}_x$ is itself a Stage VII admissible perturbation of $\hat{T}$ at the single position x:

Proposition 3.1 — Local Spectral Gap Bound

Let $C_0 := \sqrt{(7/4)}$ denote the Stage VII Bauer–Fike constant for the canonical wheel. Let $\hat{T}_x = \hat{T} + \Delta T_x$ satisfy D1–D3 with

$\|\Delta T_x\|_{op} \leq 1/(2 \cdot C_0)$ for every $x \in X$.

Then for every $x \in X$,

$\varepsilon_{gap}(x) \geq ½ - C_0 \cdot \|\Delta T_x\|_{op} \geq ½ - ½ > 0$.

In particular, the local spectral gap field $\varepsilon_{gap} : X \to (0, 1]$ is strictly positive everywhere.

Proof. Direct application of Stage VII Theorem 5.1 to the operator $\hat{T}_x$ at fixed x. The Bauer–Fike bound on the $L^2(\pi)$ operator norm controls the maximum non-trivial eigenvalue modulus, with $C(\Delta T_x) = \|\Delta T_x\|_{\{op, L^2(\pi)\}} \leq C_0 \cdot \|\Delta T_x\|_{op}$ by the comparability bound. The smallness condition $\|\Delta T_x\|_{op} \leq 1/(2 \cdot C_0)$ ensures $\varepsilon_{gap}(x) \geq ½ - ½ > 0$.

Consequence. The local spectral gap is a well-defined positive scalar field $\varepsilon_{gap} : X \to (0, 1]$ on the substrate, equal to ½ in the canonical vacuum and reduced inside defect regions by an amount controlled by the local defect strength.

3.3 Lipschitz Continuity of the Local Spectral Gap Field

The Bauer–Fike spectral perturbation bound transfers cleanly between *two* perturbed operators only under additional structural hypotheses, since individual-eigenvalue tracking on non-normal matrices is delicate. We give the sharp result under reversibility preservation (Theorem 3.2), then the weaker but more general Hausdorff-distance bound for arbitrary admissibility-preserving defects (Remark following).

Theorem 3.2 — Lipschitz Continuity (Reversibility-Preserving Case)

Suppose ΔT_x preserves reversibility for every $x \in X$ (so each $\hat{T} + \Delta T_x$ is self-adjoint with respect to its own stationary measure π_x in $L^2(\pi_x)$), with $\|\Delta T_x\|_{op} \leq 1/(2 \cdot C_0)$ uniformly across X. Suppose the defect profile is Lipschitz in the substrate metric:

$\|\Delta T_x - \Delta T_y\|_{op} \leq K \cdot d_X(x, y)$ for all $x, y \in X$.

Then the local spectral gap field is Lipschitz-continuous with constant $C_0 \cdot K$:

$$|\varepsilon_{gap}(x) - \varepsilon_{gap}(y)| \le C_0 \cdot K \cdot d_X(x, y).$$

Proof. Under reversibility preservation, $\hat{T} + \Delta T_x$ is self-adjoint in $L^2(\pi_x)$, and the operators $(\hat{T} + \Delta T_x)$ and $(\hat{T} + \Delta T_y)$ lie in a common self-adjoint setting after passing to a uniform reference measure π (justified for small perturbations since $\pi_x \to \pi$ as $\Delta T_x \to 0$). The Weyl interlacing inequality for self-adjoint perturbations gives, for each index i,

$$|\lambda_i(\hat{T} + \Delta T_x) - \lambda_i(\hat{T} + \Delta T_y)| \le \|\Delta T_x - \Delta T_y\|_{op, L^2(\pi)} \le C_0 \cdot K \cdot d_X(x, y).$$

Taking $i = 2$ (the subdominant index, well-defined under reversibility's spectral simplicity) and noting $\varepsilon_{gap}(z) = 1 - |\lambda_2(\hat{T}_z)|$ transfers the bound to the gap field.

Remark (general admissibility-preserving defects). When ΔT_x breaks reversibility, $\hat{T} + \Delta T_x$ is non-normal and individual-eigenvalue tracking between perturbed operators is delicate (no Weyl interlacing). The cleaner statement is in terms of Hausdorff spectral distance:

$$d_H(\mathrm{spec}(\hat{T} + \Delta T_x), \mathrm{spec}(\hat{T} + \Delta T_y)) \le \kappa \cdot \|\Delta T_x - \Delta T_y\|_{op} \le \kappa \cdot K \cdot d_X(x, y),$$

where κ is the uniform Bauer–Fike condition-number bound on the perturbed eigenbasis. Inside the small-perturbation regime $\|\Delta T_x\|_{op} \le 1/(2 \cdot C_0)$ where the entire paper operates, κ inherits a concrete bound from the canonical wheel: by standard analytic-eigenvector perturbation theory, the eigenbasis of $\hat{T} + \Delta T_x$ differs from the eigenbasis of $\hat{T}$ by an amount of order $\|\Delta T_x\|_{op} / \Delta_0$, where $\Delta_0 = ½$ is the canonical gap separation; this gives $\kappa \le C_0 \cdot (1 + O(\|\Delta T_x\|_{op} / \Delta_0)) \le C_0 \cdot (1 + O(1)) = O(C_0)$ *explicitly, with the constant absorbing the canonical eigenbasis condition number* ($\le C_0$ *by the* $L^2(\pi)$*-comparability already used in Proposition 3.1). Outside the small-perturbation regime,* κ *can grow without bound and the Hausdorff statement becomes only qualitatively useful. The gap-edge Lipschitz constant in this concrete regime is therefore* $C_0 \cdot (1 + O(1)) \cdot K = O(C_0) \cdot K$*, generally larger than the reversibility-preserving* $C_0 \cdot K$ *but of the same order. For the purposes of §5 (the second-difference quantity* $\nabla^2\varepsilon_{gap}$*), the Hausdorff bound is sufficient — it guarantees that the maximum non-trivial eigenvalue modulus varies Lipschitz-continuously, which is what controls* ε_{gap}*. The general modulus-Lipschitz bound on the non-Perron spectral radius (i.e., that* $\max_{\{\lambda \ne 1\}} |\lambda(\hat{T} + \Delta T_x)|$ is Lipschitz in x even without individual-eigenvalue tracking) follows from Kato §IV.3.5 (Kato, *Perturbation Theory for Linear Operators*, 2nd ed., 1995), applied to the family of perturbed operators with the Perron eigenvalue separated out.

Consequences for §5. In either case (reversibility-preserving with constant $C_0 \cdot K$, or general with constant $\kappa \cdot K$), Theorem 3.2 elevates the curvature ansatz $R(x) := \nabla^2\varepsilon_{gap}(x)$ from a symbolic to a controlled object: second differences of ε_{gap} exist as bona fide finite-difference quantities on the lattice X, bounded by the lattice spacing and the relevant Lipschitz constant. Geometric roughness, as defined by Proposition 4.1's $K_\infty(x)$, varies continuously with x rather than jumping discontinuously across defect interfaces.

For finite-support defects (D2), $\varepsilon_{gap}(x)$ is exactly ½ outside $B_r(x_0)$ and varies in a controlled fashion across $B_r(x_0)$ provided ΔT_x varies Lipschitz-continuously in x; the boundary x =

$\partial B_r(x_0)$ is the locus where the Lipschitz interpolation between bulk and defect must be smooth enough for Theorem 3.2 to apply uniformly.

---

4. Local Continuum Lipschitz Field

The Stage V continuum bound

$$K_\infty = L_\Phi \cdot L \cdot A^2 / (A_- \cdot \varepsilon_{gap})$$

was derived under the assumption of a single global spectral gap. When the gap becomes spatially varying, we must verify that the bound can be applied position-by-position.

Proposition 4.1 — Local Lipschitz Bound (Conditional, Slowly-Varying Substrate)

Let $\hat{T}_x$ *satisfy D1–D3 with* $\|\Delta T_x\|_{op} \leq 1/(2 \cdot C_0)$*. Assume the following conditions:*

- *(S2-local) The Stage IV substrate primitives* L_Φ, A, A_- *are uniform in x;* $L(x)$ *is Lipschitz in x with constant* K_L*.*
- *(smooth defect interface) The defect profile varies Lipschitz-continuously across* $\partial B_r(x_0)$*: there exists* $K_{int} < \infty$ *such that* $\|\Delta T_x - \Delta T_y\|_{op} \leq K_{int} \cdot d_X(x, y)$ *for all x, y in a neighbourhood of the interface. Equivalently, the defect has effective radius* $r \gg 1$ *in substrate units, allowing the bulk-to-defect transition to be resolved across multiple substrate sites.*

Then the local continuum Lipschitz constant

$$K_\infty(x) := L_\Phi \cdot L(x) \cdot A^2 / (A_- \cdot \varepsilon_{gap}(x))$$

is finite at every x and satisfies the Stage V continuum regularity bound restricted to the neighbourhood of x.

Proof sketch. The Stage V bound is derived by composing the local refinement filter (controlled by $L(x)$) with the spectral contraction rate (controlled by $\varepsilon_{gap}(x)$) over a single refinement step. The S2-local hypothesis ensures interactions between refinement steps at adjacent positions can be treated as effectively independent (Stage V's bounded-branching hypothesis applied locally). The smooth-interface hypothesis ensures that the local-bound application does not pick up additional interface contributions at $\partial B_r(x_0)$. Under both hypotheses, the Stage V derivation transfers position-by-position.

Note (sharp interfaces). Defects with $r \lesssim 1$ in substrate units violate the smooth-interface hypothesis. For such defects, Proposition 4.1 holds inside $B_r(x_0)$ and outside (where each region is a homogeneous Stage V application), but additional interface contributions may appear at

$\partial B_r(x_0)$. We do not characterise these in the present paper; this is the sharp-interface contribution flagged in §12.7.

Consequence — Geometric Roughness Field

The continuum Lipschitz constant $K_\infty(x)$ is *inversely* proportional to $\varepsilon_{gap}(x)$. Therefore:

- In the canonical vacuum ($\varepsilon_{gap} = \frac{1}{2}$), $K_\infty(x) = 2 \cdot L_\Phi \cdot L \cdot A^2 / A_-$ is uniform — the continuum has globally constant Lipschitz regularity.
- Inside a weak defect (ε_{gap} reduced by $\Delta\varepsilon > 0$), $K_\infty(x)$ increases by approximately $\Delta\varepsilon / \varepsilon_{gap}^2$ of its bulk value — the continuum is *rougher* in the defect region.
- As $\varepsilon_{gap}(x) \to 0$ (defect strength approaching the Stage VII admissibility limit), $K_\infty(x) \to \infty$ — the continuum Lipschitz bound fails inside the defect.

This is the structural input to §5: spatial variation of $\varepsilon_{gap}(x)$ produces spatial variation of $K_\infty(x)$, and the latter is what an emergent continuum geometry would see as variation in local regularity.

Corollary 4.2 — Continuum-Roughness Gradient

Under the hypotheses of Proposition 4.1, with substrate primitives L_Φ, A, A_- uniform in x and $L(x)$ Lipschitz in x, the spatial gradient of the local continuum Lipschitz constant satisfies

$$\nabla K_\infty(x) = -\,(L_\Phi \cdot L(x) \cdot A^2 / A_-) \cdot \nabla\varepsilon_{gap}(x) / \varepsilon_{gap}(x)^2 + L_\Phi \cdot A^2 \cdot \nabla L(x) / (A_- \cdot \varepsilon_{gap}(x)).$$

In the special case $\nabla L(x) = 0$ (uniform refinement-filter primitive), this simplifies to the chain-rule identity

$$\nabla K_\infty(x) / K_\infty(x) = -\,\nabla\varepsilon_{gap}(x) / \varepsilon_{gap}(x):$$

the relative continuum-roughness gradient is the negative of the relative spectral-gap gradient.

Note. This is the chain rule applied to K_∞ as a function of ε_{gap}, not a deep theorem; the content is recorded as a corollary to make explicit that the spectral-gap *field gradient* — not only the gap value — controls continuum-regularity variation. Combined with Theorem 3.2, this gives a quantitative control sequence:

$$K \text{ bound on } \nabla\Delta T_x \Longrightarrow C_0\cdot K \text{ bound on } \nabla\varepsilon_{gap} \Longrightarrow (C_0\cdot K/\varepsilon_{gap}(x)) \times K_\infty(x) \text{ bound on } \nabla K_\infty.$$

The continuum-roughness field is therefore as regular as the defect profile, with explicit Lipschitz constants traceable to the input data.

---

5. Candidate Curvature Indicators (Heuristic)

This section proposes a scalar candidate curvature indicator built from the local spectral gap field. The proposal is *heuristic* in the strict sense: we do not derive it from a tensorial geometric theory, we do not show it equals any specific known curvature scalar, and we list its limitations explicitly. The contribution is the identification of a structurally natural ansatz that further stages can test against a derived geometry.

5.1 Motivation

In Riemannian geometry, the scalar curvature $R(x)$ measures local deviation from flatness. Several structural features of $R(x)$ in standard geometry are useful to bear in mind:

- $R(x)$ is a *second-order* differential invariant of the metric — second derivatives of the metric field.
- $R(x)$ is *signed*: positive for "spherical" deviation, negative for "saddle-like" deviation, zero in flat regions.
- $R(x)$ vanishes in flat space.

We seek an indicator on the $K = 7$ substrate that shares these structural features and is constructed from the spectral data made available by Propositions 3.1 and 4.1.

5.2 The Heuristic Ansatz

The natural candidate is the *Laplacian of the local spectral gap*:

$R(x) := \nabla^2\varepsilon\_gap(x)$, (heuristic curvature indicator)

where ∇^2 is the substrate Laplacian on the lattice X in the standard analyst convention:

$\nabla^2 f(x) := (1/a^2) \cdot \Sigma\_\{y \sim x\} [f(y) - f(x)]$,

with lattice spacing a and sum running over nearest neighbours y of x. In this convention, a function with a local minimum at x has $\nabla^2 f(x) > 0$; therefore a *suppressed-gap region* (where ε_gap is locally minimal) gives *positive* $R(x)$, recovering the "well = positive curvature" intuition.

Structural features matching standard scalar curvature:

- Second-order in the gap field, hence second-order in the local refinement data.
- Signed, with the sign convention that "wells" in ε_gap (minima) produce positive $R(x)$.
- Vanishes in the canonical vacuum ($\varepsilon\_gap \equiv ½$ everywhere), recovering flat behaviour.

5.3 Limitations Listed Explicitly

This ansatz is *not* a derived tensorial curvature. The following limitations are intrinsic to the present formulation:

1. **Scalar, not tensorial.** A genuine curvature theory requires a metric tensor, a connection, and a Riemann tensor whose contractions yield scalar curvature. The present ansatz produces a scalar field by hand; it is silent on tensorial structure.
2. **No coordinate-invariance proof.** The substrate Laplacian ∇^2 is defined with respect to the substrate lattice X; its behaviour under continuum diffeomorphisms is not established by the present construction.
3. **No connection to a metric.** In standard geometry, R(x) is derived from the metric g_{μν}(x). The present construction does not produce a metric; it produces ε_gap(x) and operates on it directly. The relationship between ε_gap and any putative emergent metric is not derived.
4. **No Einstein-style dynamics.** Even granting the candidate-curvature interpretation, no field equation of the form G_{μν} ∝ T_{μν} is derived or proposed. The ansatz is *kinematic*, not dynamical.
5. **No conformal-vs-Ricci distinction.** Standard geometry distinguishes the Ricci scalar, the conformal factor, the Weyl tensor, and many other curvature-like objects. The present scalar makes no such distinctions.
6. **The sign of R(x) = ±∇^2ε_gap(x) is a structural choice.** The "+" convention is adopted throughout in conjunction with the analyst Laplacian to give gap-wells → positive R; a tensorial derivation may confirm this sign or invert it. Until such derivation exists, the sign should be regarded as fixed by the present convention rather than as an intrinsic feature of the ansatz.

5.4 What the Ansatz *Does* Provide

Given these limitations, the ansatz remains useful as a structural probe:

- **Sign-tracking.** It correctly tracks where the continuum is locally rougher (ε_gap depressed → K_∞ enhanced → larger R(x) at the defect core) and where it relaxes back (negative R(x) on neighbouring sites; see §9.4).
- **Vacuum recovery.** R(x) = 0 in the canonical vacuum, as required.
- **Defect localization.** R(x) is supported within and near the defect region, as required for a local curvature-like indicator.

The intended use is as a *placeholder*: any genuine tensorial curvature theory derived from the refinement substrate should reproduce, in its scalar contraction, qualitatively similar behaviour to R(x). The ansatz provides a target structure for future derivations to hit.

5.5 Why ε_gap and Not K_∞, 1/ε_gap, or log ε_gap?

A natural critic-style objection to Definition 5.1 is: given that the structural picture involves several inter-related quantities — ε_gap(x), K_∞(x) ∝ 1/ε_gap(x), the log-gap log ε_gap(x), and various other functionals — why specifically pick ∇^2ε_gap as the curvature indicator? The answer is provisional, but several structural considerations favour ε_gap:

- **Regularity at gap closure.** A genuine curvature scalar should be regular wherever the underlying field is regular and should diverge only where the underlying field becomes singular. The functions 1/ε_gap and log ε_gap both diverge as ε_gap → 0, which is the §7.3 gap-closing transition where admissibility is *lost* — that boundary should appear as a singularity of the geometry, mirroring metric singularities at curvature blow-ups in standard geometry. ε_gap itself remains regular up to and including the closure point ε_gap = 0; its second differences are finite-difference quantities on the lattice and do not diverge at the transition. So $\nabla^2\varepsilon\_gap$ behaves like a "field strength" that goes to a finite limit, while $\nabla^2(1/\varepsilon\_gap)$ and $\nabla^2(\log \varepsilon\_gap)$ blow up there. We want the former.
- **Direction-invariance in the spectrum.** ε_gap is a *property of the spectrum* (the distance from the unit circle of the second-largest eigenvalue modulus), invariant under any change of eigenbasis. The candidate-curvature ansatz inherits this invariance: $R(x) = \nabla^2\varepsilon\_gap(x)$ is well-defined without choosing a specific eigenbasis at any position. K_∞(x) shares this property (it depends only on ε_gap and substrate primitives), so this consideration does not distinguish ε_gap from K_∞.
- **Unification across phenomena (§13.1 Defect-Coherence Principle).** Continuum roughening, entropy retention, and trapped-mode persistence all depend on ε_gap(x) directly (not on K_∞ or log ε_gap). Using a different functional in the curvature ansatz would break the §13.1 unified-functional structure, requiring separate spectral primitives for each phenomenon.
- **Affine transforms vs nonlinear functionals.** The vacuum value ε_gap = ½ is a finite positive number; $\nabla^2(½) = 0$ recovers flat curvature trivially. The functions log(½) and 1/(½) also give zero second derivative in the bulk, so bulk-flatness alone does not distinguish ε_gap from K_∞ or log ε_gap. The substantive structural point is rather: since the Laplacian annihilates constants, any *affine* transform $a\cdot\varepsilon\_gap + b$ gives the identical curvature ansatz $a\cdot\nabla^2\varepsilon\_gap$ (modulo overall scale and the sign issue noted below). The interesting distinctions arise from *nonlinear* functionals. Exponentiation gives $\nabla^2(e^{\varepsilon\_gap}) = e^{\varepsilon\_gap} \cdot (\nabla^2\varepsilon\_gap + |\nabla\varepsilon\_gap|^2)$, adding a $|\nabla\varepsilon\_gap|^2$ correction; logarithm gives $\nabla^2(\log \varepsilon\_gap) = \nabla^2\varepsilon\_gap / \varepsilon\_gap - |\nabla\varepsilon\_gap|^2 / \varepsilon\_gap^2$, adding the analogous correction with opposite sign and additional 1/ε_gap factors. These nonlinear shifts can be either suppressed (in the small-gap-gradient regime) or structurally significant; the §12.1 tensorial derivation would resolve which functional class is correct.
- **K_∞ vs ε_gap.** Since $K\_\infty \propto 1/\varepsilon\_gap$ (Corollary 4.2), a curvature ansatz built on K_∞ rather than ε_gap would give $R'(x) := \nabla^2 K\_\infty(x) \propto \nabla^2(1/\varepsilon\_gap(x))$, differing from R(x) by the identity $\nabla^2(1/f) = -\nabla^2 f / f^2 + 2\,|\nabla f|^2 / f^3$ in the continuum limit. The first term recovers ε_gap-based curvature with prefactor $1/\varepsilon\_gap^2$ and *opposite sign* (note the leading minus): a K_∞-based ansatz would therefore invert the §5.2 sign convention, recovering "gap-wells → negative R'" in the analyst convention. The second term adds a $|\nabla\varepsilon\_gap|^2 / \varepsilon\_gap^3$ contribution that vanishes in linear order in $\nabla\varepsilon\_gap$. So R and R' carry the same leading-order structural information up to a sign flip and prefactor, differing at second order in the gap gradient. The choice between them — including the sign convention — is a calibration question that a tensorial derivation would resolve.

These considerations are *not* a derivation; they are provisional structural reasons for the present choice. A tensorial geometric theory derived from the substrate may favour a different scalar functional of ε_gap, may show that the difference between R, R', and other candidates is conformal-rescaling-class equivalent, or may show that the right primitive is neither ε_{gap} nor K_∞ but a third quantity related to the resolvent structure of $\hat{T}$. Each of these outcomes is consistent with the present heuristic; ε_{gap} is recorded as the *placeholder choice* on the grounds above, not as a uniquely correct primitive.

Reading guidance. The scalar $R(x) = \nabla^2 \varepsilon_{gap}(x)$ should presently be viewed as analogous to an *order parameter* or *coherence-density indicator* — a substrate-level scalar field whose spatial variation diagnoses local departures from the uniform coherent phase — not yet as geometric curvature in the differential-geometric sense. Order parameters in condensed-matter physics (magnetisation, superfluid density, BCS gap) play the same structural role: they indicate phase variation without committing to specific dynamics or to a particular tensorial geometric interpretation. The promotion from order-parameter scalar to a contraction of a derived curvature tensor is the §12.1 problem, not a step already taken in this paper.

---

6. Trapped Incoherent Modes

The most structurally interesting consequence of localized defects is the possibility of *trapped modes*: eigenvectors of the global position-indexed refinement operator that are spatially confined to the defect region and have eigenvalue larger in modulus than the bulk second eigenvalue.

This phenomenon is familiar from quantum-mechanical bound states (an attractive potential creates bound eigenstates below the continuum) and from defect modes in periodic systems (an impurity creates an eigenfrequency outside the bulk band). We adapt the structural argument to the refinement-substrate setting.

6.1 Local Operators vs the Global Position-Indexed Operator

We must distinguish two operators:

- The *local operator at the defect core*, $\hat{T}_{x_0}$, which is a 7×7 matrix acting on the closure catalogue at position x_0 alone.
- The *global position-indexed refinement operator* $\mathbf{T}$ acting on the full state space $\mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^X$, defined by $\mathbf{T} |\kappa, y\rangle = \hat{T}_y |\kappa\rangle \otimes |y\rangle$ plus the substrate spatial-coupling structure inherited from Stage IV.

The spectrum spec($\hat{T}_{x_0}$) describes what would happen if $\hat{T}_{x_0}$ acted everywhere; spec($\mathbf{T}$) describes what actually happens with $\hat{T}_{x_0}$ acting only at x_0 and $\hat{T}$ acting at every other position. These spectra coincide only in the uniform vacuum ($\hat{T}_x = \hat{T}$ for all x).

For a localized defect, the relationship between spec($\hat{T}_{x_0}$) and spec($\mathbf{T}$) is governed by a Birman–Schwinger / Krein-resolvent type problem analogous to impurity-state theory in periodic tight-binding models. The standard picture is: the local operator's discrete eigenvalues survive globally as bound states when they lie *outside the band* generated by the bulk operator acting through the substrate spatial coupling. Inside the band, local eigenvalues become resonances rather than true bound states.

In the present finite-state $K = 7$ setting, the wheel spectrum $\{1, \tfrac{1}{2}, \tfrac{1}{2}, -\tfrac{1}{4}, -3/28, 0, 0\}$ is itself discrete (it is the spectrum of a 7×7 row-stochastic matrix). Band structure for the global operator $\mathbf{T}_{bulk}$ emerges from tensoring this 7-element set with the spectrum of the substrate spatial-coupling operator inherited from Stage IV; the "bulk bands" of $\mathbf{T}_{bulk}$ are thus images of spec($\hat{T}$) under the spatial coupling, with widths and locations determined by Stage IV substrate structure. The trapped-mode regime corresponds to local eigenvalues that survive embedding into $\mathbf{T}$ as discrete eigenvalues outside these bands.

6.2 The Trapped-Mode Theorem (Global, Conditional)

Theorem 6.1 — Trapped-Mode Persistence (Conditional, Global)

Let $\hat{T}_x$ be a defect operator satisfying D1–D3, and let $\boldsymbol{T}$ be the corresponding global position-indexed refinement operator. Define $\lambda_{bulk} := \max\{|\lambda| : \lambda \in spec(\hat{T}), \lambda \neq 1\} = \tfrac{1}{2}$ as the canonical bulk subdominant modulus. Suppose there exists an eigenvalue λ_{trap} of $\boldsymbol{T}$ with

$$\lambda_{bulk} < |\lambda_{trap}| < 1,$$

and a corresponding eigenvector $\Psi_{trap} \in \mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^{X}$ spatially extended over $B_r(x_0)$ and exponentially decaying outside it:

$$\|\Psi_{trap}(\cdot, y)\|_{\mathcal{K}} \leq C \cdot e^{-d_X(y, x_0) / \xi} \text{ for some } \xi > 0 \text{ and all } y \in X.$$

Then under iterated refinement, Ψ_{trap} is a slowly-decaying component of the global dynamics, with decay rate $|\lambda_{trap}|^n$ per refinement step, exceeding the bulk decay rate $(\tfrac{1}{2})^n$. Locally inside $B_r(x_0)$, Ψ_{trap} dominates the late-time evolution.

Epistemic status. This theorem is *doubly conditional*. The hypothesis is on the global operator $\mathbf{T}$, not the local operator $\hat{T}_{x_0}$; one cannot infer trapped-mode existence from spec($\hat{T}_{x_0}$) alone. The existence of a localized eigenvector with eigenvalue above the bulk subdominant modulus is itself a non-trivial spectral fact that requires either (a) explicit construction for a given defect, or (b) a Birman–Schwinger criterion translating local-operator conditions into global eigenvalue persistence. We do not develop (b) here.

Proof sketch. Under the stated hypothesis, Ψ_{trap} is by definition an eigenvector of $\mathbf{T}$ with eigenvalue λ_{trap}, so $\mathbf{T}^n \Psi_{trap} = \lambda_{trap}^n \cdot \Psi_{trap}$. The exponential spatial decay then ensures the mode remains concentrated near x_0 at every refinement step. The bulk modes orthogonal to Ψ_{trap} (with eigenvalue moduli $\leq \tfrac{1}{2}$) decay at the bulk rate, so for sufficiently large n the trapped mode dominates locally inside $B_r(x_0)$.

Remark on the Birman–Schwinger structure. A full existence criterion would proceed as follows. The bulk operator $\mathbf{T}_{bulk}$ (with $\hat{T}$ at every position) has an essential spectrum determined by the substrate spatial structure tensored with spec($\hat{T}$) — typically a union of bands on an infinite regular substrate, or a discrete set on finite X. A localized perturbation $V := \mathbf{T} - \mathbf{T}_{\text{bulk}}$ (supported on $B_r(x_0)$) shifts this spectrum; the shifted operator $\mathbf{T} = \mathbf{T}_{\text{bulk}} + V$ has additional discrete eigenvalues exactly when the Birman–Schwinger operator $V \cdot (\mathbf{T}_{\text{bulk}} - \lambda)^{-1}$ has an eigenvalue at 1, in the region outside $\mathrm{spec}_{\mathrm{ess}}(\mathbf{T}_{\text{bulk}})$. The detailed criterion depends on the substrate spatial-coupling Stage IV structure and the substrate cardinality, and is not developed in this paper; it is the natural follow-up problem identified in §12.

Illustrative Remark — A Toy Form for T. The Stage IV substrate determines the precise form of **T**; the present paper does not commit to a specific form. As an illustration, a particularly tractable *toy* form (used here only to indicate the operator-theoretic shape of the Birman–Schwinger problem above) is

$$\mathbf{T}_{\text{toy}} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X,$$

where A_X is the substrate adjacency operator on $\mathbb{R}^X$ and $\gamma > 0$ is a spatial-coupling strength. The bulk operator (with $\hat{T}_x = \hat{T}$ for all x) has spectrum

$$\mathrm{spec}(\mathbf{T}_{\text{toy, bulk}}) = \{ \lambda + \gamma \cdot \mu : \lambda \in \mathrm{spec}(\hat{T}),\ \mu \in \mathrm{spec}(A_X) \},$$

which on an infinite regular lattice X gives bulk *bands* spanning $\gamma \cdot \mathrm{spec}(A_X)$ centred at each eigenvalue of $\hat{T}$ (with width $2\gamma \cdot \|A_X\|$ in the self-adjoint case where $\mathrm{spec}(A_X) \subset [-\|A_X\|, \|A_X\|]$; the width formula depends on the actual spectrum for general non-self-adjoint or unbounded A_X). A localized defect $V = (\hat{T}_x - \hat{T}) \otimes 1_{B_r(x_0)}$ is a finite-rank perturbation, and the Birman–Schwinger criterion becomes the concrete eigenvalue problem $V \cdot (\mathbf{T}_{\text{toy, bulk}} - \lambda)^{-1}$ has eigenvalue 1, with resolvent computable via the spectral decomposition of A_X (Fourier transform on regular substrates). We do not pursue this calculation here; the toy form is offered solely to indicate the concrete operator-theoretic shape that the Birman–Schwinger machinery would take, once Stage IV specifies the spatial-coupling structure.

6.3 Variational Detection of Spectral Excursion (Reversibility-Preserving Case)

The Theorem 6.1 hypothesis is global and difficult to verify directly. Under the additional assumption of reversibility preservation, the *eigenvalue* condition $|\lambda_{\text{trap}}| > \frac{1}{2}$ admits a finite-dimensional variational test via Rayleigh quotients, partially closing the gap between the local-operator computation (easy) and the global-operator hypothesis (hard).

Proposition 6.2 — Variational Detection of Spectral Excursion Above the Bulk Gap

*Suppose the defect operator $\hat{T}_x$ preserves reversibility for every $x \in X$, so that the global position-indexed operator **T** is self-adjoint with respect to the substrate-weighted inner product $\langle \cdot, \cdot \rangle_{\pi \otimes \mu_X}$, where μ_X is the substrate measure with respect to which the spatial-coupling operator of Stage IV is self-adjoint. (Concretely, for regular substrates with adjacency-type coupling as in the §6.2 Illustrative Remark, μ_X is the substrate counting measure $\mu_X(x) \equiv 1$; for inhomogeneous substrates, Stage IV's spatial-coupling construction induces μ_X as the measure making the spatial coupling self-adjoint in the corresponding L^2.) Suppose further that* *there exists* $f \in \mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^{X}$ *orthogonal to the Perron eigenvector* Ψ_1 *of* $\mathbf{T}$*, with support predominantly inside* $B_r(x_0)$*, such that*

$$\langle f, \mathbf{T} f\rangle_{\pi\otimes\mu_X} / \langle f, f\rangle_{\pi\otimes\mu_X} > \tfrac{1}{2}.$$

Then $\mathbf{T}$ *has at least one eigenvalue* $\lambda > \tfrac{1}{2}$*, with the corresponding eigenvector having non-zero overlap* $\langle f, \psi_\lambda\rangle \neq 0$*.*

Proof. The Courant–Fischer min–max principle for self-adjoint operators gives $\lambda_2(\mathbf{T}) = \sup_{f \perp \Psi_1, f \neq 0} \langle f, \mathbf{T} f\rangle / \langle f, f\rangle$. The hypothesis exhibits a test function $f \perp \Psi_1$ with Rayleigh quotient strictly greater than ½, hence $\lambda_2(\mathbf{T}) > \tfrac{1}{2}$. The supremum is attained by the second eigenvector ψ_2; if f had zero overlap with ψ_2 and all other eigenvectors with eigenvalue > ½, the Rayleigh quotient would be bounded by the next eigenvalue below, contradicting the strict inequality. Hence $\langle f, \psi_\lambda\rangle \neq 0$ for at least one eigenvector ψ_λ with eigenvalue $\lambda > \tfrac{1}{2}$.

What this gives. Proposition 6.2 detects the eigenvalue condition $\tfrac{1}{2} < |\lambda_{\text{trap}}| < 1$ of Theorem 6.1 while leaving the exponential-localization condition open. Concretely: one constructs a specific localized test function f, computes a single Rayleigh quotient, and if the value exceeds ½, the spectral-excursion condition is verified by a finite-dimensional variational calculation.

What this does not give. *Eigenvector localization* in the sense of Theorem 6.1's exponential-decay hypothesis. The min–max criterion shows that some eigenvector with eigenvalue > ½ has non-zero overlap with the localized test function f, but it does not directly imply that the eigenvector itself is exponentially localized — a delocalised eigenvector can have substantial overlap with a localized test function. Verifying exponential decay requires additional input, e.g. Combes–Thomas-type resolvent estimates or explicit construction of the global eigenvector, which we do not develop here.

Caveat on non-reversible defects. When ΔT_x breaks reversibility, $\hat{T}_x$ is not self-adjoint in $L^2(\pi)$ and the Rayleigh quotient $\langle f, \mathbf{T} f\rangle / \langle f, f\rangle$ bounds the *numerical range* $W(\mathbf{T}) := \{ \langle f, \mathbf{T} f\rangle / \langle f, f\rangle : f \neq 0 \}$ rather than the spectrum directly. The numerical range contains the spectrum but generally exceeds it, so the criterion above becomes a statement about $W(\mathbf{T})$ only. A symmetrized variant using $(\mathbf{T} + \mathbf{T}^*) / 2$ (which is self-adjoint) gives bounds on $\mathrm{Re}(\mathrm{spec}(\mathbf{T}))$ and is the appropriate substitute for non-reversible defects, at the cost of potentially missing complex eigenvalues with $|\lambda| > \tfrac{1}{2}$ but $\mathrm{Re}(\lambda) \leq \tfrac{1}{2}$.

6.4 Why Trapped Modes Matter Structurally

A trapped mode is structurally distinguished from bulk incoherent modes by three features:

1. **Spatial localization.** It does not propagate; its support remains confined near x_0.
2. **Slow decay.** Its decay rate $|\lambda_{\text{trap}}|^n$ is slower than the bulk rate $(\tfrac{1}{2})^n$, so at sufficiently late refinement steps the trapped mode dominates locally.
3. **Persistent incoherence.** It carries refinement incoherence (it is orthogonal to the Perron eigenspace) but does not dissipate to the same Perron equilibrium as bulk modes.

These three properties are exactly the structural features one would demand of a candidate matter-like excitation in the emergent continuum: localized, persistent, distinguishable from the vacuum.

6.5 What This Does Not Establish

Theorem 6.1 establishes that *if* the global operator **T** develops a localized eigenvector with eigenvalue above the bulk subdominant modulus, *then* the resulting mode has matter-like structural features. It does not establish:

- That such global eigenvectors actually exist for any specific defect (this is the open Birman–Schwinger problem of §6.2 Remark).
- That trapped modes have any specific particle content (no mass, charge, spin assignments).
- That trapped modes interact with each other in any specific way.
- That trapped modes obey quantum statistics or possess gauge structure.
- That trapped modes form a discrete spectrum corresponding to a particle table.

The structural analogy to matter is *suggestive*, not derivational. We label this clearly as conjectural in §13.

---

7. Defect Classification

The Stage VII paper classified *uniform* perturbations by their long-run effect on the global spectral gap (irrelevant / marginal / relevant). The present paper classifies *localized* defects by a different criterion based on the *spatial localization of the perturbed eigenvectors*. The classifications are independent: a defect can be irrelevant in the Stage VII sense yet gap-suppressing in the present sense, or vice versa.

7.1 Bounded Defects

A defect is *bounded* if the perturbed eigenvectors of the global operator **T** near the bulk spectrum remain *delocalised* — they are small perturbations of the bulk eigenvectors and do not concentrate in $B_r(x_0)$. The defect produces local-gap reduction (Proposition 3.1) but does *not* satisfy the Theorem 6.1 hypothesis: no trapped mode forms.

Sufficient condition. In the Stage VII Theorem 5.1 small-perturbation regime, $\|\Delta T_x\|_{op}$ is small enough that the eigenvectors of $\hat{T}_x$ near each bulk eigenvalue are approximately the bulk eigenvectors (which are inherently delocalised over $\mathcal{K} \times X$). This approximation is expected to persist in **T** under mild substrate-coupling regularity, though making this precise — and deriving the precise threshold below which trapped-mode formation is precluded — requires the §12.2 Birman–Schwinger machinery.

Consequences. The local continuum Lipschitz constant $K_\infty(x)$ remains uniformly bounded inside the defect, varying smoothly between the bulk value $2 \cdot L_\Phi \cdot L \cdot A^2 / A_-$ and a slightly enhanced value at the defect core. The continuum geometry is locally rougher inside the defect but remains globally Lipschitz with a (worse) constant.

Example. The §9 worked example at small defect strength $\alpha \leq 0.2$ falls in this class (verified in §9.2 numerics).

Special sub-class — spectrally invisible defects. A bounded defect can be sharper than weak-perturbation theory predicts: certain defect operators leave the local spectral gap *exactly* invariant, $\varepsilon_{gap}(x_0) = ½$, despite genuinely perturbing $\hat{T}_{x_0}$. These are the localized analogue of Stage VII §10.1 / Theorem 9.1's symmetry-protected silent perturbations and are exhibited in §9.6.

7.2 Gap-Suppressing Defects

A defect is *gap-suppressing* if at least one eigenvector of the global operator **T** becomes *spatially localized* in $B_r(x_0)$ with eigenvalue satisfying $½ < |\lambda_{trap}| < 1$, i.e. it satisfies the Theorem 6.1 hypothesis.

Equivalent characterisation. The defect is strong enough to push at least one eigenvalue meaningfully outside the bulk continuum, with the corresponding eigenvector concentrating in the defect region (exponential decay away from x_0).

Consequences. A trapped mode exists. The local continuum Lipschitz constant $K_\infty(x_0)$ is substantially enhanced. The defect supports persistent localized incoherent dynamics and is the candidate matter-like substrate structure of §§6.4–6.5 (with §6.4 giving the positive structural characterisation and §6.5 listing what is *not* established).

Existence remark. Constructing a wheel-admissible defect that demonstrably satisfies the gap-suppressing condition (in the *global* spectral sense of Theorem 6.1, not merely the local-spectrum sense) is an explicit open problem; see §12.2 and §12.3. The §9 worked example does not reach this regime within the admissible perturbation range.

7.3 Gap-Closing Defects

A defect is *gap-closing* if $\varepsilon_{gap}(x) \to 0$ at some point inside the defect, or if the defect causes $\hat{T}_x$ to violate one of the Stage VII admissibility conditions P1–P3 (loss of irreducibility, aperiodicity, or positivity).

Consequences. The Stage V continuum Lipschitz bound $K_\infty(x)$ diverges at gap-closing points; the continuum geometry is not Lipschitz across the defect. Multiple persistent sectors may emerge if the defect disconnects the local communication graph.

Example. A defect that completely severs the hub–boundary coupling at position x_0 (analogous to the Stage VII Example 11.4) is gap-closing.

7.4 Relation to Stage VII Classification

The Stage VII classification (irrelevant / marginal / relevant) is for *uniform* perturbations and concerns whether the *global* gap closes as the perturbation strength is varied. The present classification (bounded / gap-suppressing / gap-closing) is for *localized* defects and concerns *eigenvector localization* in the global position-indexed operator. The two classifications are independent:

- A bounded localized defect has delocalised perturbed eigenvectors — no trapped modes form. (Includes the spectrally-invisible sub-class.)
- A gap-suppressing localized defect produces at least one spatially-localized eigenvector with eigenvalue above the bulk subdominant modulus.
- A gap-closing localized defect is the analogue of Stage VII's "relevant" — outside the universality class locally.

The gap-suppressing category is the novel structural feature of localized defects, with no clean analogue in the Stage VII uniform-perturbation analysis: it is precisely the regime in which localization (a *spatial* property) and gap-violation (a *spectral* property) coincide.

7.5 Phase Structure: Critical Defect Strengths

Given a one-parameter family of defects $\Delta T_x = \alpha \cdot \Delta T_{x_0,r}$ of fixed profile and varying strength $\alpha \geq 0$, four characteristic α-values structure the phase diagram:

- **$\alpha_{\mathrm{admiss}}(\Delta T_{x_0,r})$** — the positivity threshold beyond which $\hat{T} + \alpha \cdot \Delta T_{x_0,r}$ ceases to be a stochastic matrix (P1 or P3 fails).
- **$\alpha_{\mathrm{smallpert}} := 1/(2 \cdot C_0 \cdot \|\Delta T_{x_0,r}\|_{\mathrm{op}})$** — the Stage VII Theorem 5.1 small-perturbation threshold, below which Proposition 3.1 guarantees $\varepsilon_{\mathrm{gap}}(x_0) > 0$ and the §7.1 bounded class applies pointwise.
- **$\alpha_{\mathrm{trap}}(\Delta T_{x_0,r})$** — the (defect- and substrate-dependent) threshold above which a *global* trapped eigenmode forms in the sense of Theorem 6.1.
- **$\alpha_c(\Delta T_{x_0,r}) := \inf\{\alpha \geq 0 : \varepsilon_{\mathrm{gap}}(x_0, \alpha) = 0\}$** — the gap-closing threshold beyond which §7.3 applies.

The §7 classification can be re-stated in terms of these thresholds: the bounded class corresponds to $\alpha < \min(\alpha_{\mathrm{trap}}, \alpha_c, \alpha_{\mathrm{admiss}})$; the gap-suppressing class corresponds to $\alpha_{\mathrm{trap}} \leq \alpha < \min(\alpha_c, \alpha_{\mathrm{admiss}})$; the gap-closing class corresponds to $\alpha \geq \alpha_c$.

Phase structure for the §9 examples. For the §9.1 boundary defect, $\|\Delta T_{x_0,r}\|_{\mathrm{op}} = \sqrt{2}$, so

$$\alpha_{\mathrm{smallpert}} = 1 / (2 \cdot \sqrt{(7/4)} \cdot \sqrt{2}) = 1/\sqrt{14} \approx 0.267,$$

while $\alpha_{\mathrm{admiss}} = 1/4 = 0.250$ (the positivity bound on $\hat{T}[\kappa_{b_1}, \kappa_{b_2}] = \tfrac{1}{4} - \alpha$). Therefore $\alpha_{\mathrm{admiss}} < \alpha_{\mathrm{smallpert}}$: the entire admissible range $\alpha \in [0, 0.25)$ sits inside the small-perturbation regime, with Proposition 3.1 active throughout. The §9.2 numerics show $\varepsilon_{\mathrm{gap}} > 0.45$ across this range, so α_c is *not reached before admissibility fails*; this defect family does not realise the gap-closing transition. The threshold α_{trap} is undetermined pending the §12.2 Birman–Schwinger machinery.

For the §9.6 hub-coupling defect, $\varepsilon_{\mathrm{gap}}(x_0, \alpha) = \frac{1}{2}$ identically across $\alpha \in [0, 1/7)$, so $\alpha_c = \infty$ and $\alpha_{\mathrm{trap}} = \infty$ within the admissible range: the entire family sits in the spectrally-invisible sub-class of §7.1, with no phase transition.

Structural note: why α_c typically coincides with α_{admiss}. The bounded → gap-suppressing → gap-closing trichotomy is a *possible* phase sequence, not a generic one. Realising all three phases in succession requires $\alpha_{\mathrm{smallpert}} < \alpha_{\mathrm{trap}} < \alpha_c < \alpha_{\mathrm{admiss}}$, and the most consequential of these inequalities is the last: $\alpha_c < \alpha_{\mathrm{admiss}}$. The structural reason this is non-trivial is spectral. Gap closure means $\varepsilon_{\mathrm{gap}}(x_0, \alpha) \to 0$, equivalently $\max_{\lambda \neq 1} |\lambda(\hat{T}_{x_0}(\alpha))| \to 1$. By Perron–Frobenius theory for stochastic matrices, an eigenvalue of modulus 1 distinct from the Perron eigenvalue 1 itself is precisely the spectral signature of a *reducible* chain (multiplicity of the Perron eigenvalue exceeds 1) or a *periodic* one (additional eigenvalues at roots of unity). Both are violations of P2 — irreducibility and aperiodicity — which are admissibility conditions in the Stage VII sense. So for generic defect families, gap closure (the eigenvalue modulus reaching 1) coincides with admissibility failure (loss of irreducibility or aperiodicity), giving $\alpha_c = \alpha_{\mathrm{admiss}}$ not by coincidence but because the two criteria collapse onto each other in the limit. Realising α_c strictly less than α_{admiss} requires *engineered* defects that push a non-Perron eigenvalue to near-modulus-1 while preserving the spectral-multiplicity and root-of-unity-avoidance criteria — a deliberate fine-tuning rather than a generic outcome.

The actual structural obstruction to populating all three phases of the §7 classification is therefore not the absence of an example but the spectral conflation: any natural family driving $\varepsilon_{\mathrm{gap}} \to 0$ will typically drive irreducibility/aperiodicity to fail at the same parameter value. The §12.3 open construction problem is therefore equivalent to constructing a defect with eigenvalue *approaching* the unit circle through the interior — staying off any root of unity, and approaching 1 transversely rather than radially — within an otherwise admissible family. This is the precise structural target.

Note on α_{trap}. The threshold α_{trap} is not directly computable from the present results: it is a *global*-operator quantity (requiring the Birman–Schwinger machinery of §6.2), not a *local*-operator one. The phase diagram for the §9.1 example (and for any specific defect family) is therefore necessarily incomplete in the α_{trap} direction by exactly the same conditionality that affects Theorem 6.1: until §12.2's Birman–Schwinger criterion is developed, α_{trap} remains symbolic in every concrete case, including ours.

---

8. Local Entropy Concentration

The Stage VII Theorem 7.1 established global entropic contraction $\chi^2(\mu_n \,\|\, \pi_\eta) \leq (\frac{1}{4} + O(|\eta|))^n \cdot \chi^2(\mu_0 \,\|\, \pi_\eta)$ under reversibility-preserving uniform perturbations. We adapt this to the local-defect setting.

8.1 Local χ^2-Divergence

Define the local χ^2-divergence at refinement step n and substrate position x as

$s(x, n) := \chi^2(\mu_n |_{\{B_r(x)\}} \| \pi_x |_{\{B_r(x)\}})$,

where $\mu_n |_{\{B_r(x)\}}$ is the restriction of the global state distribution to a neighbourhood of x and π_x is the local stationary distribution of $\hat{T}_x$. This is well-defined whenever the local operator $\hat{T}_x$ admits a unique stationary distribution (Stage VII Proposition 4.1 + Perron–Frobenius applied locally).

8.2 Local Contraction Rate

By direct application of Stage VII Theorem 7.1 to the local operator (assuming reversibility preservation), the local χ^2-divergence contracts at rate

$s(x, n) \leq (\tfrac{1}{4} + s_\alpha \cdot |\alpha| + \mathcal{O}(\alpha^2))^n \cdot s(x, 0)$,

with a slope coefficient $s_\alpha \in [0, C_0 \cdot \|\Delta T_{\{x_0,r\}}\|_{op}]$ *inheriting the Stage VII directional behaviour (writing the local defect as* $\Delta T_x = \alpha \cdot \Delta T_{\{x_0,r\}}$ for unit-profile $\Delta T_{\{x_0,r\}}$ and using the universal Bauer–Fike constant $C_0 = \sqrt{(7/4)}$). The worst case $s_\alpha = C_0 \cdot \|\Delta T_{\{x_0,r\}}\|_{op}$ is the Bauer–Fike upper bound, attained by symmetry-breaking directions such as the §9.1 boundary defect. The opposite extreme $s_\alpha = 0$ occurs for symmetry-protected perturbations such as the §9.6 hub-coupling defect, which leave the local gap (and hence the contraction rate) exactly invariant at the local level. Outside the defect region, the bulk Stage VI rate $(\tfrac{1}{4})^n$ applies exactly.

8.3 Entropy Trapping at Gap-Suppressing Defects

When the local gap $\varepsilon_{gap}(x)$ is suppressed at position x, the local contraction rate $(1 - \varepsilon_{gap}(x))^2$ is *larger* than the bulk rate $\tfrac{1}{4} = (1 - \tfrac{1}{2})^2$. Therefore:

s(x, n) decays slower at gap-suppressed positions than in the bulk.

In the limit of trapped-mode formation (Theorem 6.1), $|\lambda_{trap}| > \tfrac{1}{2}$ means $(1 - \varepsilon_{gap}(x_0))^2 > \tfrac{1}{4}$ locally, and entropy concentration persists at the defect core indefinitely under refinement.

Corollary 8.1 — Local Entropy Persistence Bound

Under the hypotheses of §8.2 with reversibility-preserving defects, the local χ^2-divergence at position x satisfies the sharp local bound

$s(x, n) \leq (1 - \varepsilon_{gap}(x))^{2n} \cdot s(x, 0)$,

expressed directly in terms of the local gap (not the bulk leading-order rate ¼ plus a correction). For bulk positions $\varepsilon_{gap}(x) = \tfrac{1}{2}$, this reproduces the $(\tfrac{1}{4})^n$ bulk rate exactly. For gap-suppressed *positions with* $\varepsilon_{gap}(x) = \frac{1}{2} - \delta$*, the local rate becomes* $(\frac{1}{2} + \delta)^{2n} = (\frac{1}{4})^{n} \cdot (1 + 2\delta)^{2n}$*, so the ratio of local-to-bulk entropy retention is*

$$s(x, n) / s(x_{bulk}, n) \le (1 + 2\delta)^{2n},$$

exceeding 1 exponentially fast whenever $\delta > 0$*.*

Proof. Standard $L^2(\pi)$-contraction for reversible Markov operators: the operator $\hat{T}_x$ restricted to mean-zero functions has spectral radius $(1 - \varepsilon_{gap}(x))$, so $\|\hat{T}_x^n f\|_{L^2(\pi)} \le (1 - \varepsilon_{gap}(x))^n \cdot \|f\|_{L^2(\pi)}$ for any f orthogonal to the Perron eigenvector. The χ^2-divergence is the squared $L^2(\pi)$ norm of the density ratio minus 1, hence inherits the squared rate $(1 - \varepsilon_{gap}(x))^{2n}$.

Quantitative consequence. The exponent is *doubly* sensitive to the local gap: a 10% suppression of ε_{gap} (from ½ to 0.45, i.e. $\delta = 0.05$) shifts the local contraction rate from $(\frac{1}{4})^n$ to $(0.55)^{2n} = (0.3025)^n$, increasing local entropy retention relative to the bulk by a factor of $(1.21)^n$ per refinement step. After 10 refinement steps, the local entropy is approximately 6.7× greater than the bulk value at the same step. This is the precise quantitative form of the qualitative §8.3 statement and is the load-bearing input to the §8.4 conjectural interpretation.

8.4 Conjectural Interpretation

The Stage V continuum bound $K_\infty \propto 1/\varepsilon_{gap}$ and the local entropy concentration $s(x, n) \to \text{const}$ at trapped-mode defects suggest a structural analogy: regions of persistent refinement incoherence carry persistent local entropy, which would translate at the continuum level into a localized energy-density-like quantity.

We *conjecture*, without proof, that trapped-mode entropy concentration is the substrate-level precursor to localized continuum energy density. The label "conjectural" is meant strictly: no theorem in this paper supports this identification, no Einstein-style equation is derived, and the analogy is structural only. It is recorded here to motivate the next stage of investigation.

---

9. Worked Examples

We exhibit two explicit localized defects: a *boundary self-loop / cyclic* defect that produces actual gap suppression (§9.1–9.5), and a contrasting *hub-coupling* defect that leaves the local gap exactly invariant (§9.6). The two together illustrate the bounded class and its spectrally-invisible sub-class.

9.1 The Boundary Defect

At a distinguished substrate position $x_0 \in X$, consider the local operator

$$\hat{T}_{x_0} = \hat{T} + \alpha \cdot \Delta T, \quad \Delta T = e_{\kappa_{b_1}} \otimes (e_{\kappa_{b_1}} - e_{\kappa_{b_2}})^T,$$

which has the rank-1 entry pattern

$\Delta T[\kappa_{b_1}, \kappa_{b_1}] = +1$, $\Delta T[\kappa_{b_1}, \kappa_{b_2}] = -1$, all other entries zero.

This is the Stage VII §10.2 self-loop / cyclic perturbation, applied here only at substrate position x_0 (and zero at every other substrate position). The admissible range is $\alpha \in [0, 1/4)$ (positivity of $\hat{T}_{x_0}[\kappa_{b_1}, \kappa_{b_2}] = 1/4 - \alpha$). The perturbation is genuinely symmetry-breaking under D_6: not antisymmetric under any reflection fixing κ_{b_1}.

9.2 Computed Local Spectrum

By direct numerical computation of $\mathrm{spec}(\hat{T} + \alpha \cdot \Delta T)$ across $\alpha \in [0, 0.20]$:

α	**$\lambda_2(\hat{T}_{x_0})$**	**$\varepsilon_{gap}(x_0)$**	**$K_\infty(x_0)$ (unit calibration)**	**Bauer–Fike bound ($\frac{1}{2} - \sqrt{2} \cdot \alpha$)**
0.00	0.50000	0.50000	4.000	0.50000
0.02	0.50343	0.49657	4.028	0.47172
0.05	0.50897	0.49103	4.073	0.42929
0.08	0.51502	0.48498	4.124	0.38686
0.10	0.51937	0.48063	4.161	0.35858
0.15	0.53145	0.46855	4.268	0.28787
0.20	0.54548	0.45452	4.400	0.21716

The local gap $\varepsilon_{gap}(x_0)$ decreases linearly with measured Lipschitz slope −0.226 across $\alpha \in [0, 0.20]$. The Bauer–Fike bound predicts slope $-\|\Delta T\|_{op, L^2(\pi)} = -\sqrt{2} \approx -1.414$, where $\sqrt{2}$ is the $L^2(\pi)$ operator norm of this *specific* perturbation ΔT, distinct from the universal Bauer–Fike comparability constant $C_0 = \sqrt{7/4}$ of Proposition 3.1. The conservatism factor is

$\|\Delta T\|_{op, L^2(\pi)} / |\mathrm{slope}| \approx \sqrt{2} / 0.226 \approx 6.25.$

This reproduces the Stage VII §10.2 numerics exactly, since the local analysis at a single substrate position is the Stage VII single-operator analysis. The matching values confirm that the localization framework correctly inherits Stage VII conservatism.

Direction-dependence of the conservatism factor. The value 6.25 here is specific to the symmetry-breaking direction of Stage VII §10.2; it is *not* a universal substrate constant. The §9.6 example below exhibits the opposite limiting case, where the conservatism factor is effectively infinite on the symmetry-protected subspace (the local gap is preserved exactly across the entire admissible range). Stage VII's conservatism is in this sense a *function* of spectral direction, ranging from 6.25 (worst observed admissibility-preserving direction) to ∞ (silent directions of Stage VII Theorem 9.1 and its localized analogue in §9.6).

9.3 Local Lipschitz Field

Under unit-substrate calibration $L\_\Phi = L = A = A_- = 1$ (Stage VII §8), the local continuum Lipschitz constant is

$$K\_\infty(x_0) = 2 / \varepsilon\_gap(x_0),$$

shifting from $K\_\infty = 4.000$ at $\alpha = 0$ to $K\_\infty = 4.400$ at $\alpha = 0.20$ — a 10% increase across the boundary of the natural perturbation range. The defect produces local geometric roughening directly proportional to the local gap suppression, consistent with the $K\_\infty \propto 1/\varepsilon\_gap$ dependence.

9.4 Spatial Profile and Candidate Curvature

Suppose the defect is localized at a single substrate position x_0 on a regular lattice with spacing a and coordination number z (number of nearest neighbours of each lattice site), with $\varepsilon\_gap(x) = ½$ for all $x \neq x_0$ and $\varepsilon\_gap(x_0) = ½ − \Delta$ where $\Delta = ½ − \varepsilon\_gap(x_0) > 0$ is the gap depression at the defect core. The discrete Laplacian (analyst convention) at x_0 is

$$\nabla^2\varepsilon\_gap(x_0) = (1/a^2) \cdot \Sigma\_\{y \sim x_0\} [\varepsilon\_gap(y) − \varepsilon\_gap(x_0)] = (z/a^2) \cdot \Delta,$$

so the candidate curvature indicator at the defect core is

$$R(x_0) = \nabla^2\varepsilon\_gap(x_0) = (z/a^2) \cdot \Delta > 0,$$

positive as required by the §5.2 sign convention.

Spatial profile beyond the core. Consider a neighbour y of x_0 (with $y \neq x_0$ but $d\_X(y, x_0) = 1$). The site y sees one neighbour (namely x_0) at the depressed value $\varepsilon\_gap(x_0) = ½ − \Delta$, and $z − 1$ neighbours at the bulk value $\varepsilon\_gap = ½$. The discrete Laplacian at y is therefore

$$\nabla^2\varepsilon\_gap(y) = (1/a^2) \cdot [(½ − \Delta − ½) + (z − 1) \cdot (½ − ½)] = −\Delta/a^2,$$

giving $R(y) = −\Delta/a^2 < 0$. At every neighbour of x_0, the candidate curvature is *negative*.

Structural interpretation. A point defect generates a *positive core surrounded by a negative ring* in the candidate-curvature field. The total contribution $\Sigma\_y R(y)$ summed over x_0 and its z nearest neighbours satisfies

$$R(x_0) + z \cdot R(y\_neighbour) = (z/a^2) \cdot \Delta + z \cdot (−\Delta/a^2) = 0,$$

so the candidate curvature has zero spatial integral over the defect cluster (defect core plus immediate neighbour ring). This is structurally reminiscent of the trace pattern of a stress concentrated at a point — the source is balanced by the response in its immediate vicinity. Whether this cancellation reflects a genuine conservation law of the substrate (a candidate divergence-free property of the curvature indicator) or is a coincidence of the lattice discretisation is an open question; resolving it would require the full derived tensorial geometry of §12.1.

Generalisation to finite-radius defects. The zero-sum identity $R(x_0) + z \cdot R(y_{\text{neighbour}}) = 0$ above is specific to a single-site depression with bulk recovery at every neighbour. For a finite-radius defect ($r > 1$), the depressed region spans multiple sites and the "neighbour ring" is one shell beyond the support. The cleaner general statement is a discrete-divergence theorem: $\Sigma_x \nabla^2 f(x)$ telescopes to boundary terms, which vanish identically when f returns to its bulk value outside finite support. Therefore, for *any* compactly-supported gap depression $\varepsilon_{\text{gap}}(x) - \frac{1}{2}$ supported on a finite subset $S \subset X$, the candidate curvature satisfies

$$\Sigma_{\{x \in X\}} R(x) = 0,$$

with the sum identically zero by discrete summation by parts on the bounded support. The point-defect cancellation above is the single-site case of this general identity, and the conjectural "conservation law" of §12.10 is therefore a generic discrete-Laplacian property rather than a coincidence of the point-defect example — strengthening the conjecture, since the property is genuinely structural rather than situational.

9.5 Trapped-Mode Check

For $\alpha \leq 0.20$ (within the admissible range), the maximum non-Perron eigenvalue is $\lambda_2(\hat{T}_{\{x_0\}}) \leq 0.5455$ — above ½ but only by about 9% at the boundary of the admissible range. To trigger Theorem 6.1's trapped-mode mechanism robustly, one would need:

1. A *global* operator eigenvalue (not merely a local-operator eigenvalue) with $|\lambda_{\text{trap}}|$ meaningfully greater than ½.
2. A corresponding *globally-localized* eigenvector — exponentially decaying away from x_0 on the full substrate, not merely the single-position eigenvector that the local-operator analysis yields.

Neither condition can be verified from the present local-operator computation alone. The numerics show that at $\alpha = 0.20$, the local operator $\hat{T}_{\{x_0\}}$ has $\lambda_2 \approx 0.55$, but whether this survives embedding into the global $\mathbf{T}$ as a localized eigenvector (rather than smearing into the bulk continuum) is the open Birman–Schwinger question of §6.2.

In the present worked example, the defect is therefore *bounded* in the §7.1 sense: local-gap reduction is real, but the conditions for trapped-mode formation are not demonstrably met within the admissible range. A stronger example exhibiting demonstrable trapped-mode behaviour requires either (a) a defect operator pushed outside the small-perturbation regime, with explicit construction of the global localized eigenvector, or (b) development of the Birman–Schwinger criterion to predict when local eigenvalues survive globally. Both are flagged in §12 as immediate open problems.

9.6 Contrasting Example — Hub-Coupling Defect (Spectrally Invisible)

For contrast, consider a different localized defect at substrate position x_0:

$$\hat{T}_{\{x_0\}} = \hat{T} + \alpha \cdot \Delta T_{\{hub\}}, \quad \Delta T_{\{hub\}} = e_{\{\kappa_h\}} \otimes (e_{\{\kappa_h\}} - e_{\{\kappa_{\{b_1\}}\}})^T,$$

with non-zero entries +1 at (κ_h, κ_h) and -1 at (κ_h, κ_{b_1}). The admissible range is $\alpha \in [0, 1/7)$ (positivity of $\hat{T}_{x_0}[\kappa_h, \kappa_{b_1}] = 1/7 - \alpha$).

By direct numerical computation, the *full spectrum* of $\hat{T}_{x_0}$ as a function of α is:

α	$\mathrm{spec}(\hat{T}_{x_0})$ (sorted by $\lvert\cdot\rvert$ descending)	$\varepsilon_{gap}(x_0)$
0.00	$\{ 1, ½, ½, -¼, -3/28, 0, 0 \}$	½
0.05	$\{ 1, ½, ½, -¼, -0.0571, 0, 0 \}$	½
0.10	$\{ 1, ½, ½, -¼, -0.0071, 0, 0 \}$	½
0.14	$\{ 1, ½, ½, -¼, +0.0329, 0, 0 \}$	½

Across the entire admissible range, *six of the seven eigenvalues are exactly invariant*; only the eigenvalue at $-3/28$ drifts (toward 0 and past it). Crucially, the moving eigenvalue's modulus stays well below ½ throughout the admissible range, so the local spectral gap $\varepsilon_{gap}(x_0) = ½$ is *exactly* invariant.

Structural mechanism. Both vectors in the outer-product decomposition $\Delta T_{hub} = e_{\kappa_h} \otimes (e_{\kappa_h} - e_{\kappa_{b_1}})^T$ lie in $H_{\sigma^*}^+$ (since σ^* fixes both κ_h and κ_{b_1}; σ^* additionally fixes κ_{b_4}, the antipodal vertex on the b_1–b_4 reflection axis, but κ_{b_4} does not appear in this perturbation). Applying the block-diagonalisation procedure of Stage VII §10.1 with respect to the σ^*-eigenspace decomposition (here adapted to a different perturbation direction), the perturbation acts as the *zero* operator on the 2-dimensional $H_{\sigma^*}^-$ block (since ΔT_{hub} maps $H_{\sigma^*}^-$ to $\{0\}$), and as a rank-1 operator on the 5-dimensional $H_{\sigma^*}^+$ block.

The $H_{\sigma^*}^-$ eigenvalues of $\hat{T}$ are $\{½, 0\}$, computed as follows. In the σ^*-antisymmetric basis $u_1 := (e_{\kappa_{b_2}} - e_{\kappa_{b_6}})/\sqrt{2}$, $u_2 := (e_{\kappa_{b_3}} - e_{\kappa_{b_5}})/\sqrt{2}$ of $H_{\sigma^*}^-$, the Stage VI boundary-row structure of $\hat{T}$ (each boundary row has three non-zero entries of value ¼ at the hub and the two neighbouring boundary states, with the b_2/b_6 and b_3/b_5 swaps acting transitively under σ^*) gives directly

$$\hat{T}\,(e_{\kappa_{b_2}} - e_{\kappa_{b_6}}) = (¼) \cdot (0, 0, 1, 1, 0, -1, -1)^T = (\sqrt{2}/4) \cdot (u_1 + u_2),$$

and identically $\hat{T}\,(e_{\kappa_{b_3}} - e_{\kappa_{b_5}}) = (\sqrt{2}/4) \cdot (u_1 + u_2)$. Since $u_1 = (e_{\kappa_{b_2}} - e_{\kappa_{b_6}})/\sqrt{2}$, normalising by $1/\sqrt{2}$ gives

$$\hat{T}\,u_1 = (1/\sqrt{2}) \cdot \hat{T}\,(e_{\kappa_{b_2}} - e_{\kappa_{b_6}}) = (1/\sqrt{2}) \cdot (\sqrt{2}/4) \cdot (u_1 + u_2) = (¼) \cdot (u_1 + u_2),$$

and similarly $\hat{T}\,u_2 = (¼) \cdot (u_1 + u_2)$. Hence $\hat{T}|_{H_{\sigma^*}^-}$ acts in the orthonormal (u_1, u_2) basis as the 2 × 2 matrix

$$\hat{T}|_{H_{\sigma^*}^-} = [\, [¼, ¼], [¼, ¼] \,],$$

with trace ½, determinant 0, and eigenvalues $\{½, 0\}$. Since ΔT_{hub} maps $H_{\sigma^*}^- \to \{0\}$ (verified above), these block eigenvalues are preserved exactly under arbitrary α.

Within $H_{\sigma^*}^+$, the 5 eigenvalues $\{1, \tfrac{1}{2}, -\tfrac{1}{4}, -3/28, 0\}$ of $\hat{T}|_{H_{\sigma^*}^+}$ respond to the rank-1 perturbation according to the matrix-determinant lemma:

$$\det(\hat{T}|_{H_{\sigma^*}^+} + \alpha \cdot u v^T - \lambda I) = \det(\hat{T}|_{H_{\sigma^*}^+} - \lambda I) \cdot (1 + \alpha \cdot v^T (\hat{T}|_{H_{\sigma^*}^+} - \lambda I)^{-1} u),$$

with u, v the $H_{\sigma^*}^+$ projections of e_{κ_h} and $e_{\kappa_h} - e_{\kappa_{b_1}}$ respectively. An eigenvalue λ_k is preserved under perturbation if and only if the resolvent factor $v^T (\hat{T}|_{H_{\sigma^*}^+} - \lambda I)^{-1} u$ has *no pole* at λ_k — equivalently, either the right eigenvector at λ_k satisfies $v^T \psi_R^{(k)} = 0$ or the left eigenvector satisfies $u^T \psi_L^{(k)} = 0$. Direct computation gives:

- $\lambda_k = 1$: $v^T \psi_R^{(1)} = 0$ → preserved.
- $\lambda_k = \tfrac{1}{2}$: $u^T \psi_L^{(\frac{1}{2})} = 0$ → preserved.
- $\lambda_k = -\tfrac{1}{4}$: $u^T \psi_L^{(-\frac{1}{4})} = 0$ → preserved.
- $\lambda_k = 0$: $u^T \psi_L^{(0)} = 0$ → preserved.
- $\lambda_k = -3/28$: both products nonzero → moves under α.

Only the $-3/28$ eigenvalue lacks the orthogonality condition that suppresses the pole, and is therefore the only $H_{\sigma^*}^+$ eigenvalue that responds to the perturbation.

The net effect is that 6 of 7 spectrum eigenvalues — the entire $H_{\sigma^*}^-$ block plus four of the five $H_{\sigma^*}^+$ eigenvalues — are invariant under arbitrary α in the admissible range, with the moving $-3/28$ eigenvalue bounded well below ½ in modulus throughout.

This is *not* the Stage VII Theorem 9.1 silent class (which requires conjugation-*antisymmetric* perturbations). It is a distinct invariance phenomenon driven by ΔT_{hub} having its entire image in $H_{\sigma^*}^+$ plus the within-block orthogonality structure above. We do not attempt to characterise the general class of such "block-aligned + resolvent-orthogonal" perturbations; the §9.6 example demonstrates the phenomenon's existence and structural origin.

In plain terms. The hub-coupling defect changes a hidden part of the operator that does not touch the gap-controlling modes. The wheel's spectrum has two parts: the modes that determine the gap (the eigenvalues at ½, which sit at the gap edge, and the Perron eigenvalue 1, which provides the reference from which the gap is measured), and the modes deep inside (the eigenvalues at $-\tfrac{1}{4}$, $-3/28$, and 0). The hub-coupling defect lives in a direction of the operator space that is orthogonal (in the right and left eigenvector sense) to all of the gap-controlling modes; it can only affect the one deep-inside eigenvalue at $-3/28$, and even when that one is moved, it stays bounded well away from the gap edge throughout the admissible range. So although the defect is a real, non-trivial perturbation of $\hat{T}_{x_0}$, the *quantity that matters for emergent geometry* (the local spectral gap) does not see it at all.

This is the localized analogue of Stage VII §10.1 (boundary anisotropy as spectrally invariant) and Theorem 9.1 (symmetry-protected silent perturbations). The hub-coupling defect is a *spectrally invisible defect*: a genuine perturbation of $\hat{T}_{x_0}$ that leaves the local-gap structure exactly unchanged.

Candidate-curvature consequence. Since $\varepsilon_{gap}(x_0) = \frac{1}{2} = \varepsilon_{gap}(y)$ for all y, the gap field is spatially uniform, and $R(x) = \nabla^2\varepsilon_{gap}(x) = 0$ everywhere. The hub-coupling defect produces *no* candidate curvature in the emergent geometry — it is invisible at every level of the present formulation. Whether such defects would be invisible in a derived tensorial geometry, or whether tensorial structure would reveal something the scalar ansatz misses, is an open question.

Programmatic significance. The two §9 examples together show that the bounded class of §7.1 contains genuinely different sub-types: defects that produce smooth local-gap variation (§9.1–9.5, ordinary bounded behaviour) and defects that produce no variation at all (§9.6, spectrally invisible). The wheel's residual D_6 symmetry — already shown by Stage VII §10.1 and Theorem 9.1 to protect uniform perturbations — also protects localized ones, when the defect direction lies in the σ^*-symmetric subspace $H_{\sigma^*}^+ \otimes H_{\sigma^*}^+$ of admissible perturbations at x_0 (i.e., when both outer-product vectors of the rank-1 ΔT_x are σ^*-symmetric) *and* the within-block right/left orthogonality conditions of the mechanism above are satisfied.

---

10. Defect Interactions

When multiple defects coexist on the substrate, their effects compose. The structural behaviour depends on whether the defects overlap and on the spatial coupling between substrate positions.

10.1 Non-Overlapping Local Operators, Possibly Coupled Global Dynamics

If two defects at positions x_1, x_2 have non-overlapping support ($B_r(x_1) \cap B_r(x_2) = \emptyset$), then the local operators $\hat{T}_x = \hat{T}$ for x outside both regions and $\hat{T}_x = \hat{T}_{x_i}$ for x inside region i. The local spectral gap field is the union:

$\varepsilon_{gap}(x) = \varepsilon_{gap}^{(i)}(x)$ for $x \in B_r(x_i)$, $i = 1$ or 2, $\varepsilon_{gap}(x) = \frac{1}{2}$ elsewhere.

Each defect contributes its own local Lipschitz enhancement and (potentially) its own candidate trapped mode, with no interaction *at the level of the local-operator spectra*.

Hedge on global dynamics. This non-interaction is at the *local-operator* level. The *global* position-indexed operator $\mathbf{T}$ still couples positions x_1 and x_2 through the substrate spatial structure (Stage IV propagation), so the corresponding global trapped modes — if they exist — can in principle hybridise across the defect pair at finite separation. The strength of this hybridisation falls off with $d_X(x_1, x_2)$, governed by the spatial-coupling decay rate of $\mathbf{T}$. For small separations, two nearby gap-suppressing defects may produce a *coupled* trapped-mode pair rather than two independent ones; for large separations, the modes decouple. A precise crossover scale would emerge from the Birman–Schwinger criterion deferred in §6.2. We do not develop this here.

10.2 Overlapping Defects

When defect supports overlap, the local operators must be summed:

$\Delta T_x = \Delta T_x^{(1)} + \Delta T_x^{(2)}$ for $x \in B_r(x_1) \cap B_r(x_2)$.

The local spectral gap at overlap points is not simply the minimum of the individual gaps; it is computed from the *sum* of perturbations, which can produce nonlinear combinations of the individual effects.

Heuristic prediction. If both defects produce gap depressions of magnitude Δ_i individually, the overlap region may produce gap depression up to $\Delta_1 + \Delta_2$ (additive) or qualitatively different behaviour (constructive or destructive interference, depending on whether $\Delta T^{(1)}$ and $\Delta T^{(2)}$ are in the same or opposite spectral directions).

We do not develop the interaction theory in detail. The qualitative structure — that nearby defects can interact constructively or destructively, with possible composite-defect structures emerging at strong overlap — is recorded for later development.

10.3 Multi-Defect Curvature Field

Generalising the candidate curvature indicator of §5:

$R(x) = \nabla^2 \varepsilon_{gap}(x)$ with $\varepsilon_{gap}(x)$ computed from the full multi-defect operator $\hat{T}_x$.

This is automatic from the definition and inherits all the limitations of §5.

11. Relation to the Stage VII Universality Class

The Stage VII paper established that the canonical wheel sits inside an open universality class $\mathcal{C}_{K=7}$. The present paper studies what happens when admissibility-preserving perturbations are *localized* rather than uniform.

11.1 Bounded Defects: Pointwise in the Universality Class

A bounded defect (§7.1) produces local-operator $\hat{T}_x$ at each x with $\varepsilon_{gap}(x) > 0$ uniformly. By Stage VII Theorem 12.1, each local $\hat{T}_x$ lies in $\mathcal{C}_{K=7}$. Therefore the entire defected substrate is *pointwise* in the universality class:

$\hat{T}_x \in \mathcal{C}_{K=7}$ for every $x \in X$.

This is the *pointwise* / *substrate-section* version of universality: at every substrate position, the local operator is a representative of the Stage VII class.

Note on substrate-level universality. Stage VII's $\mathcal{C}_{K=7}$ is a class of *single* row-stochastic matrices, not of position-indexed families. The "pointwise in the class" statement above is a natural extension, but a substrate-level universality class covering position-dependent operator families $(\hat{T}_x : x \in X)$ and incorporating the substrate spatial coupling has not been formalised in the present paper. Properly defining such a class — and proving its openness, robustness, etc., as substrate-level analogues of the Stage VII single-operator results — is a definitional extension whose development is left for later stages.

11.2 Gap-Suppressing Defects: At the Edge of the Pointwise Class

Gap-suppressing defects (§7.2) require defect strength outside the Proposition 3.1 small-perturbation regime $\|\Delta T_x\|_{op} < 1/(2 \cdot C_0)$. *At each individual position* $x \in B_r(x_0)$*, the local operator* $\hat{T}_x$ *may still lie in* $\mathcal{C}\{K=7\}$ (irreducibility and aperiodicity hold by Stage VII Propositions 4.1–4.2 transferring without smallness restriction; positive spectral gap holds whenever $|\lambda_2(\hat{T}_x)| < 1$, which is automatic if stochasticity is preserved).

So pointwise membership $\hat{T}_x \in \mathcal{C}_{K=7}$ for all x can still hold for gap-suppressing defects; the difference is that the small-perturbation Stage VII bounds no longer apply position-by-position. The trapped-mode phenomenon of §6 is a *substrate-level* feature requiring the substrate-level extension flagged in §11.1; it is not visible at the pointwise level, because the pointwise universality class concerns single operators at individual positions, whereas trapped-mode existence is intrinsically a global property of the position-indexed operator **T** acting on the full $\mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^X$ state space.

11.3 Gap-Closing Defects: Outside the Pointwise Class

Gap-closing defects (§7.3) produce $\varepsilon_{gap}(x) \to 0$ or violate P1–P3, hence U5 or earlier U-conditions of Stage VII Definition 12.1. These exit the pointwise universality class at the defect core. The continuum geometry breaks down across the defect.

11.4 Summary

The Stage VII universality class $\mathcal{C}_{K=7}$ is *pointwise robust* against all bounded localized defects and contains gap-suppressing defects pointwise. The trapped-mode and curvature phenomena of §§5–6, however, are *substrate-level* features that require an explicit substrate-level universality class extension not formalised here. Gap-closing defects exit the pointwise class. The picture is structurally consistent with the Stage VII open-neighbourhood result: localized defects redistribute the substrate's local representative within $\mathcal{C}_{K=7}$ without leaving it (in the bounded and many gap-suppressing cases) or fracture the class (gap-closing). What remains is the substrate-level extension that lifts pointwise statements to position-indexed-operator statements.

---

12. Limitations and Open Problems

The present paper provides the operator-theoretic machinery for studying localized defects in the $K = 7$ universality class and identifies three structural phenomena (local gap reduction, candidate curvature indicators, trapped modes). The following extensions remain open.

12.1 Tensorial curvature theory. The §5 candidate indicator $R(x) = \nabla^2\varepsilon_{gap}(x)$ is a scalar, not a tensor. A genuine emergent geometry requires a metric tensor $g_{\mu\nu}(x)$, a connection, a Riemann tensor, and the standard contractions. Deriving any of these from the substrate is the principal open problem of the geometry programme. The Stage V Lipschitz bound and the present local-gap field are scalar precursors only.

12.2 Birman–Schwinger criterion for trapped-mode existence. Theorem 6.1 is conditional on the existence of a *globally* localized eigenvector with $|\lambda_{trap}| > \frac{1}{2}$ in the spectrum of the position-indexed operator $\mathbf{T}$. The local-operator condition $|\lambda| > \frac{1}{2}$ in $\mathrm{spec}(\hat{T}_{x_0})$ is necessary but not sufficient for global trapped-mode existence — local eigenvalues inside the bulk band become resonances rather than true bound states. Developing the explicit Birman–Schwinger / Krein-resolvent criterion for the $K = 7$ substrate is a concrete operator-theoretic problem requiring the Stage IV spatial-coupling structure as input.

The natural next paper in the programme would attack this directly, with the following engineering pieces required:

- *Substrate spatial coupling.* Fix the explicit form of the spatial-coupling operator (the $I_{\mathcal{K}} \otimes A_X$ part of $\mathbf{T}$, in the §6.2 Illustrative Remark's notation) that Stage IV induces — the toy form $\mathbf{T}_{toy} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X$ is the simplest case but Stage IV may require a more general structure (e.g. non-uniform γ, or coupling terms mixing closure labels with substrate positions).
- *Resolvent structure of $\mathbf{T}_{bulk}$.* Compute the resolvent $(\mathbf{T}_{bulk} - \lambda)^{-1}$ explicitly. For the toy form on regular substrates, this reduces to Fourier analysis on A_X together with the spectral decomposition of $\hat{T}$. For general Stage IV substrates, more work is needed.
- *Localization criteria.* Derive explicit conditions on a defect operator $V = \Delta T_x \otimes 1_{B_r(x_0)}$ under which the Birman–Schwinger operator $V \cdot (\mathbf{T}_{bulk} - \lambda)^{-1}$ has eigenvalue 1 in the gap region $|\lambda| \in (\frac{1}{2}, 1)$.
- *Band-edge analysis.* Determine the band edges of $\mathrm{spec}(\mathbf{T}_{bulk})$, since trapped modes exist only outside the bands. For finite-$K = 7$ substrates the band structure has at most 7 bands, with widths and gaps determined by $\mathrm{spec}(\hat{T})$ and the spatial-coupling strength γ.
- *Combes–Thomas-type decay estimates.* Prove exponential decay of trapped-mode eigenvectors away from the defect core, completing the Theorem 6.1 hypothesis (which requires not just eigenvalue existence but eigenvector localization with rate ξ).

These pieces are largely sequential: (a) substrate spatial coupling is prerequisite to all others; (b) resolvent structure builds on (a) and is prerequisite to (c) and (d); (c) localization criteria and (d) band-edge analysis can be developed in parallel after (b); and (e) Combes–Thomas decay estimates attach after (c), since they presuppose that a trapped eigenvalue exists before its eigenvector can be shown to decay. A future-paper plan would naturally proceed (a) → (b) → (c, d in parallel) → (e).

This is the authentic next step for the geometry programme. With these pieces in place, Theorem 6.1's hypothesis becomes verifiable for explicit defect families, the §7.5 phase-diagram threshold α_trap becomes computable, and §12.3's strong-defect construction problem can be attempted on a concrete operator-theoretic footing.

12.3 Explicit strong-defect example with verified global trapped mode. The §9.1 worked example exhibits local gap reduction but stays within the bounded class. An explicit wheel-admissible defect with a verified *global* trapped mode (not merely a local-spectrum bump above ½) is the most direct extension of the present paper and remains to be constructed.

12.4 Substrate-level universality class. §11.1 notes that Stage VII's 𝒞_{K=7} is a class of single operators, not of position-indexed families. A substrate-level universality class — incorporating spatial coupling and supporting substrate-level theorems analogous to Stage VII Theorem 12.1 — is a definitional extension whose proper development would underpin the trapped-mode and curvature claims at the substrate level rather than pointwise.

12.5 Matter interpretation. §6.5 lists what trapped-mode existence does *not* establish: no mass, charge, spin, gauge structure, statistics, or particle table. Whether trapped modes can be classified, quantized, or matched to known particle content is the central question of any emergent-matter programme. Nothing in the present paper bears on this beyond the structural analogy.

12.6 Defect interaction theory. §10 sketches the qualitative structure of defect interactions without developing the calculation. A controlled overlapping-defect analysis (perturbative in defect separation distance) is a finite computational exercise; the hybridisation of trapped modes across non-overlapping defects (§10.1 hedge) requires the §12.2 Birman–Schwinger machinery.

12.7 Sharp-interface contributions. Proposition 4.1 assumes slowly-varying substrate primitives. Sharp defect interfaces (radius r comparable to substrate spacing) may carry contributions not captured by the local-bound argument. A boundary-layer analysis at defect interfaces is a deferred problem.

12.8 Quantization. The framework is classical / refinement-theoretic throughout. Whether trapped modes admit a natural quantization (perhaps from the underlying refinement event statistics) is open and requires input from the quantum-foundations branch of the programme.

12.9 Connection to known emergent-matter approaches. The picture proposed here — localized refinement defects as matter precursors — has structural analogies to topological-defect approaches in condensed-matter physics (vortices, skyrmions) and to certain loop-quantum-gravity approaches (spin-network defects). A systematic comparison is deferred.

12.10 Conservation properties of the candidate curvature field. The §9.4 observation that $R(x_0) + z \cdot R(y\_neighbour) = 0$ over a point defect plus its neighbour ring suggests a possible divergence-free structure in the candidate curvature. Whether this reflects a genuine conservation law of the substrate or is a lattice-discretisation artefact, and whether such a property survives in a derived tensorial geometry, are open questions requiring §12.1's tensorial input.

13. Conclusion

Programme map. To orient: Stage V established continuum regularity from spectral coherence — the canonical $K = 7$ wheel operator's spectral gap of $\tfrac{1}{2}$ gives the substrate a Lipschitz continuum limit. Stage VII established robustness of that coherence substrate — small admissibility-preserving perturbations leave the gap positive and the universality class open. The present paper studies *localized failures* of substrate coherence — defects that lower the local spectral gap in finite regions, with consequent effects on continuum regularity, candidate curvature, entropy retention, and trapped-mode persistence. The trajectory Stage V → VII → VIII therefore moves from *establishing coherence* to *establishing the robustness of coherence* to *characterising local breakdowns of coherence*, with each stage contained operator-theoretically in the previous.

The Stage VII universality result established that the canonical $K = 7$ wheel operator sits inside an open class of admissible refinement substrates, with two named structural mechanisms (symmetry-protected spectral invariance, operator-norm conservatism) explaining the breadth of the class. The present paper extends the operator-theoretic machinery to *spatially localized* admissibility perturbations and identifies three structural phenomena:

- **Local spectral gap as a field** (Propositions 3.1, 4.1). The Stage VII spectral stability theorem applied position-by-position gives a well-defined positive scalar field $\varepsilon_{gap} : X \to (0, 1]$, lower-bounded by $\tfrac{1}{2} - C_0 \cdot \|\Delta T_x\|_{op} > 0$ *in the small-perturbation regime, with the local Stage V Lipschitz constant* $K_\infty(x) \propto 1/\varepsilon_{gap}(x)$ varying inversely. This is *proven*, by direct local application of Stage VII results.
- **Candidate curvature indicators** (§5, heuristic). The Laplacian $R(x) := \nabla^2 \varepsilon_{gap}(x)$ (analyst convention) is a structurally natural scalar built from the local gap field, tracking spatial variation of refinement coherence. We label this *heuristic* throughout and list its limitations explicitly: it is scalar, not tensorial; not derived from a metric; not connected to Einstein dynamics; not coordinate-invariant in any established sense. Its intended use is as a *target structure* that a future derived tensorial geometry should reproduce in scalar contraction. The §9.4 observation of a positive core + negative neighbour-ring structure, with zero spatial integral over the defect cluster, is a structural feature the candidate ansatz exhibits and that a derived geometry should ideally explain or refine.
- **Trapped incoherent modes** (Theorem 6.1, conditional and global). If the global position-indexed refinement operator $\mathbf{T}$ develops an eigenvector with $|\lambda_{trap}| > \tfrac{1}{2}$ exponentially localized in the defect region, that mode persists under refinement at decay rate slower than the bulk. The existence condition is on the *global* operator, not the local operator $\hat{T}_{x_0}$; the relationship is structurally analogous to impurity-state existence in periodic tight-binding models and requires a Birman–Schwinger criterion (§6.2 Remark) we have flagged but not developed. The §9.1 worked example produces local-spectrum behaviour consistent with the bounded class; the gap-suppressing regime with verified global trapped modes remains an open construction.

The defect classification of §7 — bounded / gap-suppressing / gap-closing — is the localized analogue of the Stage VII RG-style classification (irrelevant / marginal / relevant), but based on *eigenvector localization* rather than eigenvalue magnitude. The two classifications are independent. The §9.6 contrasting example — the hub-coupling defect with exactly invariant local spectral gap — exhibits the localized analogue of Stage VII §10.1's symmetry-protected silent perturbations, identifying a *spectrally invisible* sub-class of bounded defects.

The conjectural claim recorded in §8.4 is that local entropy concentration at trapped-mode defects is the substrate-level precursor to localized continuum energy density. This is structural analogy, not derivation. No phenomenology is claimed, no field equation is proposed, and no comparison to specific matter content is established.

What the paper *does* establish is the operator-theoretic groundwork for treating local refinement failure within the $K = 7$ universality class. The Stage V → VI → VII chain — explicit substrate, exact spectrum, universality class — is now extended to:

Stage VIII: localized defects within the universality class → spatially varying gap field → candidate curvature indicator → conditional trapped-mode structure → conjectural matter-like precursors.

Each arrow carries its epistemic label. The first two are *proven*. The third is *heuristic*. The fourth is *conditional* on Theorem 6.1's hypothesis (which is itself *global* and requires the §12.2 Birman–Schwinger machinery to verify in any specific case). The fifth is *conjectural*.

The honest summary is that this paper opens a structurally natural direction without claiming to have traversed it. Smooth continuum geometry corresponds to uniform refinement coherence in a way that is now rigorously stated and locally bounded; spatial variation of refinement coherence is a candidate curvature precursor in a way that is now structurally well-formulated but not yet tensorially derived; localized refinement failure carries trapped incoherent modes in a way that is now spectrally *diagnosed* at the local level (via the necessary condition $|\lambda_2(\hat{T}_{x_0})| > \frac{1}{2}$), but *verification* of trapped-mode existence requires global-operator analysis.

The next stages of the geometry programme must address each open item in §12. The principal targets are: a tensorial curvature theory derived from the substrate (12.1), the Birman–Schwinger existence criterion for global trapped modes (12.2), an explicit strong-defect construction exhibiting verified trapped-mode persistence (12.3), and a substrate-level universality class extension (12.4). With those in place, the structural analogy to matter-like content can be tested against derivation rather than left as a label.

The picture proposed across Stages V–VIII is therefore consistent and incomplete in identifiable, addressable ways. Geometry, candidate curvature, and candidate matter-like structure all emerge as *manifestations of the same refinement substrate* — coherent uniformity, spatially varying coherence, and localized incoherence respectively. Whether this picture survives derivation in the tensorial direction, and whether trapped modes recover any specific physical content, are explicit programme problems carried forward. The present paper's contribution is to make those problems precisely statable in operator-theoretic terms within the established universality class.

13.1 Defect-Coherence Principle (Structural Summary)

The three primary structural phenomena identified in this paper — local spectral gap reduction, candidate curvature, and trapped incoherent modes — yield *four distinct mathematical functionals* of the local spectral gap field $\varepsilon_{gap}(x)$, since gap reduction has two consequences (continuum roughening and entropy retention) that are each ε_{gap}-functionals in their own right. All four functionals are listed below; $\varepsilon_{gap}(x)$ is a *discrete scalar field on the substrate lattice X* (per the §3.1 terminological note; derivatives below are finite differences and prefactors involving substrate primitives L_Φ, L, A, A_- are suppressed for readability):

- **Continuum roughening** (Proposition 4.1, Corollary 4.2): $K_\infty(x) \propto 1 / \varepsilon_{gap}(x)$; $\nabla K_\infty \propto -\nabla\varepsilon_{gap} / \varepsilon_{gap}^2$ (modulo substrate primitives).
- **Candidate curvature** (Definition 5.1, §5.5): $R(x) = \nabla^2\varepsilon_{gap}(x)$, the second-order functional of the same field — the natural second-derivative counterpart to the first-derivative roughness-gradient relation above.
- **Entropy retention** (Corollary 8.1): $s(x, n) \leq (1 - \varepsilon_{gap}(x))^{2n} \cdot s(x, 0)$. This sits as a corollary of gap reduction rather than as a separate primary phenomenon.
- **Trapped-mode persistence** (Theorem 6.1, Proposition 6.2): requires $\lambda_2(\mathbf{T}) > \frac{1}{2}$, equivalently effective local gap $< \frac{1}{2}$ in the global operator.

We record this shared dependence as a structural principle:

Defect-Coherence Principle. *Localized suppression of refinement coherence — i.e. reduction of $\varepsilon_{gap}(x)$ in a finite region — simultaneously increases local continuum roughness (Corollary 4.2), produces candidate curvature $R(x) = \nabla^2\varepsilon_{gap}(x) \neq 0$ (Definition 5.1), slows local entropy contraction (Corollary 8.1), and increases the candidate persistence of localized incoherent modes (Theorem 6.1, Proposition 6.2). The four functionals are different mathematical extractions from the same underlying spectral-gap field — some pointwise, some involving spatial derivatives, some involving the global-operator spectrum — but all controlled by $\varepsilon_{gap}(x)$.*

This is a *structural unification* statement, not a derivational one: it asserts that the four functionals share a common functional dependence on $\varepsilon_{gap}(x)$, not that they share a deeper physical origin. Whether the unification can be lifted to a genuine common derivation — perhaps in the tensorial geometry of §12.1, where $\varepsilon_{gap}(x)$ itself would emerge as a scalar contraction of more primitive substrate tensors — is an open programme question. The §9.6 hub-coupling example is the principle's limiting case: $\varepsilon_{gap}(x)$ is unperturbed, so none of the four ε_{gap}-mediated functionals respond, although $\hat{T}_x$ is genuinely changed at the operator level.

The Defect-Coherence Principle replaces, in a calibrated form, the conjectural metaphysical statement that "curvature, entropy concentration, and trapped excitations are different manifestations of one underlying phenomenon": the underlying common thing is the gap field, the unification is functional rather than physical, and the principle delineates exactly what structural symmetry the four functionals share and what they do not.

13.2 Scope of the Candidate Curvature Claim

To prevent misreading of the candidate curvature claim:

The present paper does not claim that curvature has been derived. It identifies the first scalar diagnostic of departure from uniform continuum regularity produced by localized refinement defects — the spatial variation of the local spectral gap field ε_gap(x), with the ansatz R(x) := ∇^2ε_gap(x) recording this diagnostic in a form that mirrors the standard geometric scalar-curvature signature. The tensorial problem — deriving a metric tensor, connection, and Riemann tensor from the substrate, and showing that R(x) is a scalar contraction of these — remains open and is the §12.1 principal target.

Equivalently: this paper establishes that something *capable of carrying* curvature content exists at the substrate level (a positive scalar field whose second differences are finite, signed, and localized to defect regions). Whether that scalar is in fact the contraction of a derived curvature tensor, what dynamics it would obey, and how it relates to standard gravitational variables are subsequent questions. The contribution is the *substrate-level scalar precursor*, calibrated to be hit by a future derivation, not the derivation itself.

The order-parameter / coherence-density framing — developed at length as reading guidance in §5.5 — fits this end-of-paper summary perfectly: across the Stage V → VII → VIII programme trajectory, what the present paper produces is the substrate-level scalar diagnostic that any future tensorial geometry must reproduce as a scalar contraction. The order-parameter analogy is thus not just a reading aid for R(x) in isolation; it is the programme-positional statement of where Stage VIII sits in the larger derivation chain — at the layer that produces the diagnostic, not yet at the layer that derives the tensors whose contraction the diagnostic should match.