Global Refinement Transport and Spectral Localization in the K = 7 Universality Class

Coupled Substrate Operators, Transport Band Structure, and the Birman–Schwinger Criterion for Trapped Coherence Modes in VERSF

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

The earlier papers in the VERSF geometry programme proposed that the large-scale structure we identify as physical reality — continuum geometry, fields, and ultimately matter-like content — emerges from a deeper discrete substrate built from a simple primitive: at every point of an underlying lattice, a small set of 7 internal "closure" states arranged as one hub plus six boundary states, evolving under a fixed canonical wheel of probability flow. Stage V established that this canonical wheel produces a Lipschitz continuum limit at large scales, with the limit's regularity controlled by a single spectral number (the spectral gap, $\varepsilon_{gap} = ½$ for the canonical wheel). Stage VII established robustness — the wheel sits inside an open universality class rather than at a fine-tuned point. Stage VIII studied what happens when the wheel is perturbed in a small local region only, and identified four ε_{gap}-functionals controlling continuum roughening, candidate curvature, entropy retention, and trapped-mode persistence.

The Stage VIII paper, however, left one substantial gap. The trapped-mode result was *conditional* on the existence of a global eigenvector of the position-indexed refinement operator — but Stage VIII did not commit to a specific form for that global operator. Trapped-mode existence was therefore deferred to a future paper that would supply the missing structure: namely, the explicit dynamics by which refinement coherence propagates between substrate positions, and the spectral / resolvent machinery needed to determine when a local defect creates a genuine globally-localized eigenstate.

This is that paper.

The central object is a concrete coupled global operator **T** acting on the full substrate, which combines the local wheel dynamics at each position with an explicit spatial-coupling term governing how coherence propagates between neighbouring positions. The substrate, in this picture, is no longer a collection of isolated 7-state wheels; it is a single connected dynamical system in which coherence flows globally.

Once this global operator is in place, four structural phenomena become accessible that were not accessible in Stage VIII:

- **Transport bands.** The bulk spectrum of **T** (with no defect) is no longer a discrete 7-element set; it is a union of *bands* obtained by combining the wheel spectrum with the substrate's spatial coupling. This is structurally identical to the way energy bands form in solid-state physics — the spectrum of a single atom becomes a band when many atoms are coupled into a crystal.
- **Finite coherence propagation.** Because the coupling has bounded spatial range, coherence cannot propagate instantaneously. There is an explicit upper bound on the rate at which a perturbation at one substrate position can affect another, governed by the spatial coupling's bandwidth. This is the substrate-level analogue of the speed of light: an emergent finite propagation speed *built into* the dynamics rather than imposed by hand.
- **Trapped coherence modes.** With the bulk transport spectrum in hand, a defect at one position is a finite-rank perturbation of **T**. The Birman–Schwinger criterion (a standard tool from impurity-state theory in tight-binding models, here adapted to the K = 7 substrate) gives an explicit, finite-dimensional test for whether the defect creates an eigenvalue of **T** outside the bulk bands — that is, a genuinely trapped coherence mode of the full substrate, exponentially localized near the defect, with eigenvalue meaningfully above the bulk subdominant modulus ½.
- **Closing Stage VIII's conditionality.** The Stage VIII trapped-mode theorem (Theorem 6.1) was conditional on a global eigenvector with two properties: eigenvalue above ½, and exponential spatial decay. The present paper supplies the operator-theoretic machinery to verify *both* conditions for explicit defects on explicit substrates, removing the conditionality and turning the Stage VIII theorem into something checkable rather than presumed.

The epistemic register is the same as Stage VIII: proven where the operator theory permits, conditional where stronger structural input is needed, heuristic where the physical interpretation is clear but the derivation is not, conjectural where only the structural analogy is in hand. The paper does not derive Einstein's equations, particle phenomenology, or a quantum-field-theoretic structure. What it does derive is the explicit transport machinery that closes the operator-theoretic gap Stage VIII identified, and that makes the trapped-mode side of the Stage V → VII → VIII → IX chain genuinely computable rather than just structurally available.

What this paper does *not* do: it does not commit to a specific Stage IV substrate (the choice of lattice X and the precise form of the spatial coupling are inputs, not outputs); it does not derive the substrate-level universality class Stage VIII §12.4 flagged (this is one substantial step further); and it does not connect trapped modes to identifiable physical particle content. The contribution is the explicit transport operator, its band structure, the Birman–Schwinger trapped-mode criterion, and the Combes–Thomas-type exponential decay of localized eigenvectors — the four engineering pieces Stage VIII §12.2 listed as prerequisites for any concrete trapped-mode analysis.

Abstract

Stage VIII established that localized admissibility-preserving defects in the K = 7 refinement universality class produce a spatially varying local spectral gap field ε_gap(x), candidate curvature indicators R(x) := ∇²ε_gap(x), and conditional trapped incoherent modes. The trapped-mode theorem (Stage VIII Theorem 6.1) was conditional on the existence of a globally localized eigenvector of the position-indexed refinement operator **T** — a global object Stage VIII did not commit to specifying.

The present paper supplies the missing structure. We define the coupled global refinement transport operator on the substrate state space $\mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^{X}$ as

$$\mathbf{T} := \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X + C,$$

where $\hat{T}$ is the canonical K = 7 wheel operator ($\mathrm{spec}(\hat{T}) = \{1, \tfrac{1}{2}, \tfrac{1}{2}, -\tfrac{1}{4}, -3/28, 0, 0\}$), A_X is the substrate adjacency operator on the lattice X, γ > 0 is the coherence-transport coupling, and C is a closure-label-mixing correction (often zero in the simplest substrates). The principal results are:

- **Bulk transport spectrum (Proposition 3.1).** For C = 0, the bulk operator **T**_bulk has spectrum $\mathrm{spec}(\mathbf{T}_{bulk}) = \{\lambda + \gamma\mu : \lambda \in \mathrm{spec}(\hat{T}), \mu \in \mathrm{spec}(A_X)\}$, organising into at most 7 bands centred at the wheel eigenvalues with widths set by γ·diam(spec(A_X)).
- **Finite coherence propagation (Theorem 4.1).** The coupled operator satisfies a Lieb–Robinson-type bound with a strict light cone v_strict = 1 (matrix elements vanish for d_X(x, y) > n) and an effective propagation speed v_c := γ · ρ(A_X) (matrix elements are factorially suppressed for d_X(x, y) > e·v_c·n inside the strict cone); both bounds are stated in the $L^2(\pi \otimes \mu_X)$ inner product, with the proof using the exact commutativity of $\hat{T} \otimes I_X$ and $\gamma \cdot I_{\mathcal{K}} \otimes A_X$ rather than a Heisenberg-conjugation form (in the minimal-coupling regime C = 0; a non-zero closure-mixing C modifies v_c by an additive contribution bounded by ‖C‖). This gives a substrate-derived finite propagation speed.
- **Birman–Schwinger criterion (Theorem 7.1).** For a localized defect $V := \Delta T_{\{x_0\}} \otimes \mathbb{1}_{B_r(x_0)\}}$ (a finite-rank perturbation of **T**_bulk), $\lambda \notin \mathrm{spec\_ess}(\mathbf{T}_{bulk})$ is an eigenvalue of **T**_bulk + V if and only if $1 \in \mathrm{spec}(K(\lambda))$, where $K(\lambda) := V^{1/2} \cdot (\mathbf{T}_{bulk} - \lambda I)^{-1} \cdot V^{1/2}$ is the symmetrised Birman–Schwinger operator (in the reversible case; the non-reversible variant uses the unsymmetric form $K(\lambda) = V \cdot (\mathbf{T}_{bulk} - \lambda I)^{-1}$).
- **Combes–Thomas exponential localization (Theorem 8.1).** Trapped eigenvalues λ at spectral distance $\delta := \mathrm{dist}(\lambda, \mathrm{spec\_ess}(\mathbf{T}_{bulk})) > 0$ have eigenvectors satisfying $\|\Psi(\cdot, y)\| \le C \cdot e^{-\eta \cdot d_X(y, x_0)}$ with $\eta \ge c \cdot \delta/\gamma$ for an explicit constant c depending only on the substrate's local geometry. The localization length is $\xi = 1/\eta \sim \gamma/\delta$.
- **Coherence-transport geometry (§5).** The combinatorial transport distance $d_{coh}(x, y) := \inf\{n : \langle e_y, \mathbf{T}^n e_x \rangle \neq 0\}$ and its spectral refinement (the Agmon-type distance built from the local gap) furnish the first substrate-level emergence of large-scale geometric structure from coherence propagation rather than from the lattice metric directly.
- **Local spectral gap as a coarse-grained projection (§10).** The Stage VIII local-gap field ε_gap(x) is identified as the spectral-radius projection of the spatial fibre of **T** at x; the

four ε_gap-functionals of the Stage VIII Defect-Coherence Principle inherit their structure from corresponding properties of the global **T**.
- **Stage VIII trapped-mode theorem becomes unconditional (§11).** The Stage VIII Theorem 6.1 hypothesis is verified by the Birman–Schwinger criterion (eigenvalue condition) together with Combes–Thomas decay (eigenvector localization). The Stage VIII conditional becomes a checkable computation.

We do not derive tensorial curvature, quantum field theory, particle phenomenology, or cosmological dynamics. The contribution is the explicit operator-theoretic engineering of substrate-scale coherence transport — the four pieces Stage VIII §12.2 identified as prerequisites — and the consequent unconditionality of Stage VIII's trapped-mode structure.

Contents

---

1. Introduction

Stage VIII left an explicit gap. Its trapped-mode theorem (Theorem 6.1) required a global eigenvector of the position-indexed refinement operator **T**, with two non-trivial properties: an eigenvalue λ_trap with ½ < |λ_trap| < 1, and exponential spatial localization around the defect region. Stage VIII supplied a finite-dimensional variational test (Proposition 6.2) for the *eigenvalue* condition under reversibility, but the *localization* condition — exponential decay of the eigenvector itself — was left as an open Birman–Schwinger problem (Stage VIII §6.2 Remark, §12.2). More fundamentally, Stage VIII did not commit to a specific form for **T** itself,

treating it as an abstract object whose properties were inherited from Stage IV without being computed.

The present paper supplies the missing structure. We fix a concrete form for **T** — concrete enough to compute spectra, resolvents, and Birman–Schwinger operators explicitly, while general enough to cover the natural class of regular substrates indicated by Stage IV. We then carry out the four engineering pieces Stage VIII §12.2 identified:

(a) substrate spatial coupling — fixed explicitly via the substrate adjacency operator A_X, with a single coupling strength parameter $\gamma > 0$; (b) resolvent structure of **T**_bulk — computed via Fourier analysis on A_X together with the spectral decomposition of $\hat{T}$ (§3); (c) localization criteria — the Birman–Schwinger eigenvalue-unity condition outside the bulk transport bands (§7); (d) band-edge analysis — explicit characterisation of the bulk band structure as the convolution of spec($\hat{T}$) with $\gamma \cdot$spec(A_X) (§3.2); (e) Combes–Thomas decay estimates — exponential localization of trapped eigenvectors with rate set by the spectral distance from the bulk bands (§8).

The order matches Stage VIII §12.2's dependency analysis: (a) → (b) → (c, d in parallel) → (e). With these in place, Stage VIII Theorem 6.1's hypothesis becomes verifiable for explicit defects, the Stage VIII phase-diagram threshold α_{trap} becomes computable, and the trapped-mode side of the Defect-Coherence Principle becomes unconditional.

Three structural phenomena emerge that were not visible at the Stage VIII local level:

1. **Band structure.** The bulk transport operator's spectrum is no longer the wheel spectrum $\{1, ½, ½, −¼, −3/28, 0, 0\}$ but rather the convolution of this discrete 7-element set with the spatial-coupling spectrum $\gamma \cdot$spec(A_X). The 7 wheel eigenvalues become 7 *bands* of finite width, and trapped modes live in the gaps between these bands.
2. **Finite coherence propagation.** A Lieb–Robinson-type bound (familiar from quantum many-body theory, here adapted to the refinement substrate) gives an explicit upper bound on the speed at which a perturbation at one substrate position can affect another. This is the first substrate-derived finite propagation speed in the programme — an emergent "speed limit" for coherence, set by the spatial-coupling bandwidth $\gamma \cdot \rho$(A_X).
3. **Eigenvector localization.** Combes–Thomas-type resolvent estimates (familiar from random Schrödinger operator theory) give exponential decay of trapped-mode eigenvectors, with localization length $\xi \sim \gamma/\delta$ where δ is the spectral distance from the bulk band edge.

The combination — bulk band structure, finite propagation, and Combes–Thomas localization, joined by the Birman–Schwinger trapped-mode criterion — closes the conditionality of Stage VIII Theorem 6.1.

What this paper does *not* do:

- It does not commit to a specific Stage IV substrate. The lattice X and the coupling structure are inputs; the operator-theoretic machinery developed here applies to any regular substrate with bounded-degree A_X and any choice of $\gamma > 0$.
- It does not derive a tensorial curvature theory. The §5 transport geometry is a *combinatorial* and *spectral* construction (transport distance, Agmon distance), not a Riemannian one.
- It does not derive the substrate-level universality class Stage VIII §12.4 flagged. The pointwise universality class of Stage VII is sufficient for the present paper; the substrate-level extension is one further step.
- It does not derive matter content. Trapped modes have, here as in Stage VIII, no mass / charge / spin / gauge / quantum-statistical assignment. They are operator-theoretic structures with matter-like *features* (localized, persistent, distinguishable from the vacuum), nothing more.

The contribution is the explicit transport machinery and its consequences for the Stage VIII trapped-mode programme.

---

2. The Global Refinement Transport Operator

2.1 Substrate Setup

Following Stage VI–VIII, 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

$$\mathscr{H} := \mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^{X},$$

with natural basis $\{e_\kappa \otimes e_x : \kappa \in \mathcal{K}, x \in X\}$. We write states as $\Psi(\kappa, x)$ and inner products in the natural product structure $\langle \Phi, \Psi \rangle_{\mathscr{H}} := \sum_{\kappa,x} \Phi(\kappa, x) \cdot \Psi(\kappa, x)$, or, when reversibility-preserving structure is required, as $\langle \Phi, \Psi \rangle_{\{\pi \otimes \mu_X\}} := \sum_{\kappa,x} \pi(\kappa) \cdot \mu_X(x) \cdot \Phi(\kappa, x) \cdot \Psi(\kappa, x)$ where π is the stationary measure of $\hat{T}$ on $\mathcal{K}$ and μ_X is the substrate measure inherited from Stage IV's spatial-coupling construction (Stage VIII §6.3 Proposition 6.2 discussion).

2.2 The Adjacency Operator A_X

The substrate spatial structure is encoded by the adjacency operator A_X on $\mathbb{R}^X$:

$$(A_X f)(x) := \sum_{y \sim x} f(y),$$

where $y \sim x$ denotes "y is a nearest neighbour of x" in the substrate graph. We assume throughout:

- **(S1) Bounded degree.** Every site $x \in X$ has at most z_{max} neighbours, with $z_{max} < \infty$.

- **(S2) Symmetric coupling.** $y \sim x \Leftrightarrow x \sim y$ (so A_X is self-adjoint with respect to the counting measure $\mu_X(x) \equiv 1$ in the simplest case).
- **(S3) Connected substrate.** The substrate graph is connected.

Bounded degree gives the spectral radius bound $\rho(A_X) \leq z_{max}$. For an infinite regular lattice of coordination number z, $\rho(A_X) = z$ (the Perron value), with $spec(A_X) \subset [-z, z]$. For a finite substrate of N sites with coordination number z, $spec(A_X)$ is a discrete subset of $[-z, z]$; for an infinite regular lattice, $spec(A_X)$ is typically a band filling $[-z, z]$ (e.g., for $\mathbb{Z}^d$ with $z = 2d$, $spec(A_X) = [-2d, 2d]$).

2.3 The Coupled Transport Operator

The global refinement transport operator is

$$\mathbf{T} := \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X + C, \quad \text{(Definition 2.1)}$$

where:

- $\hat{T} \otimes I_X$ acts on $\Psi(\kappa, x)$ by applying $\hat{T}$ to the κ-coordinate at each fixed x, encoding the local refinement dynamics at each substrate position.
- $\gamma \cdot I_{\mathcal{K}} \otimes A_X$ acts by averaging over neighbouring substrate positions at each fixed κ, encoding substrate-mediated coherence transport. The parameter $\gamma > 0$ sets the strength of coherence propagation; in the limit $\gamma \to 0$ the substrate positions decouple completely.
- C is a *closure-label-mixing* coupling term, supported on edges of X and acting non-trivially on both indices simultaneously. The simplest models take $C = 0$; we develop the general theory below in this minimal-coupling regime and note where $C \neq 0$ generalisations are straightforward.

2.4 The Coupling Regime

Throughout this paper we work in the **weak-coupling regime**

$$\gamma \cdot \rho(A_X) < \tfrac{1}{2}, \quad \text{(W)}$$

equivalently $\gamma < 1/(2 \cdot \rho(A_X))$. This is the analogue, at the global-operator level, of Stage VIII's small-perturbation regime $\|\Delta T_x\|_{op} \leq 1/(2 \cdot C_0)$. It ensures that the bulk transport bands (§3) remain well-separated and the bulk subdominant modulus stays below 1 — so that the Perron eigenvalue 1 of $\mathbf{T}_{bulk}$ remains isolated and the gap from 1 to the next-largest band edge remains positive. Outside this regime, bands can overlap and the spectral picture becomes more delicate; we flag this in §13.

Note that (W) is a non-trivial restriction: for $\mathbb{Z}^2$ with $z = 4$, $\rho(A_X) = 4$, so (W) requires $\gamma < 1/8$. For higher-dimensional or more highly-connected substrates the admissible γ shrinks. This is the substrate-level analogue of the way that Stage VII's admissibility-range shrinks for stronger spectral directions: the engineering tractability of the explicit transport theory comes at the cost of restricting the coupling parameter.

2.5 Connection to Stage VIII

For $\gamma = 0$ and $C = 0$, the operator $\mathbf{T}$ reduces to $\hat{T} \otimes I_X$ — copies of the canonical wheel acting independently at each substrate position. This is the *trivial uncoupled limit* in which Stage VIII's local-operator analysis applies directly: each substrate fibre carries the canonical wheel dynamics, and $\mathrm{spec}(\mathbf{T}) = \mathrm{spec}(\hat{T})$ (each eigenvalue with multiplicity $|X|$).

For $\gamma > 0$, the substrate positions are dynamically coupled, and $\mathrm{spec}(\mathbf{T})$ is no longer simply $\mathrm{spec}(\hat{T})$: it acquires band structure (§3), finite propagation (§4), and the Birman–Schwinger phenomena that occupy the remainder of this paper. Stage VIII's trapped-mode hypothesis becomes verifiable precisely because $\gamma > 0$ introduces the non-trivial spatial structure against which "localized" and "extended" eigenvectors can be distinguished.

---

3. Bulk Transport Spectrum and Band Structure

The bulk transport operator (no defect, no closure-label-mixing) is

$$\mathbf{T}_{\mathrm{bulk}} := \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X. \quad (3.1)$$

Its spectrum is computable in closed form.

Proposition 3.1 — Bulk Spectrum as Convolution

Let $\hat{T}$ have spectrum $\mathrm{spec}(\hat{T}) = \{\lambda_1, \ldots, \lambda_7\} = \{1, \tfrac{1}{2}, \tfrac{1}{2}, -\tfrac{1}{4}, -3/28, 0, 0\}$ with right eigenvectors $\{\psi_i^{(R)}\}$ and left eigenvectors $\{\psi_i^{(L)}\}$, and let A_X have spectrum $\mathrm{spec}(A_X) \subset [-\rho(A_X), \rho(A_X)]$ with eigenvectors (or generalised eigenfunctions, in the infinite case) $\{\phi_\mu\}$. Then

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

with eigenvectors $\Psi_{i,\mu}(\kappa, x) = \psi_i^{(R)}(\kappa) \cdot \phi_\mu(x)$.

Proof. Direct verification:

$$\mathbf{T}_{\mathrm{bulk}}\,(\psi_i^{(R)} \otimes \phi_\mu)(\kappa, x) = (\hat{T}\,\psi_i^{(R)})(\kappa) \cdot \phi_\mu(x) + \gamma \cdot \psi_i^{(R)}(\kappa) \cdot (A_X\,\phi_\mu)(x)$$
$$= \lambda_i \cdot \psi_i^{(R)}(\kappa) \cdot \phi_\mu(x) + \gamma \cdot \mu \cdot \psi_i^{(R)}(\kappa) \cdot \phi_\mu(x) = (\lambda_i + \gamma\cdot\mu) \cdot (\psi_i^{(R)} \otimes \phi_\mu)(\kappa, x).$$

The tensor-product eigenvectors $\{\psi_i^{(R)} \otimes \phi_\mu\}$ span $\mathbb{R}^{\mathcal{K}} \otimes \mathbb{R}^X$, hence exhaust the spectrum.

Theorem 3.2 — Essential Spectrum of the Bulk Operator

*Let X be an infinite regular substrate (so spec(A_X) accumulates) and **T**_bulk satisfy Definition 2.1 with C = 0. Then the essential spectrum of **T**_bulk is the union of transport bands:*

spec_ess(**T**_bulk) = ⋃*{i=1}^{7} (λ_i + γ · spec(A_X)) =* ⋃{i=1}^{7} 𝔅_i,

where the bands 𝔅_i are as in Definition 3.3 below. In particular, if spec(A_X) = [−ρ(A_X), ρ(A_X)] (the standard case for infinite vertex-transitive lattices), each band is a compact interval

$\mathfrak{B}_i = [\lambda_i - \gamma\cdot\rho(A_X), \lambda_i + \gamma\cdot\rho(A_X)] \subset \mathbb{R}$,

and the essential spectrum is a union of at most 7 compact intervals (5 in the canonical case, since the wheel eigenvalues 0 and ½ each appear with multiplicity 2 and produce coincident bands).

Proof. Since A_X on an infinite regular substrate has purely essential spectrum (no isolated eigenvalues of finite multiplicity), and $\hat{T} \otimes I_X$ is a direct integral of copies of $\hat{T}$ over X, the operator **T**_bulk = $\hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X$ is *similar* (via the tensor-product right-eigenvector decomposition of Proposition 3.1) to multiplication by $(\lambda_i, \mu) \mapsto \lambda_i + \gamma\cdot\mu$ on the spectral decomposition. For the canonical wheel $\hat{T}$, the right/left eigenvector basis is biorthogonal but not orthogonal in the counting inner product, so the change of basis is a similarity rather than a unitary transformation. (In the $L^2(\pi \otimes \mu_X)$ inner product, when the spatial fibres preserve reversibility so $\hat{T}$ is self-adjoint, the similarity becomes a unitary equivalence; the essential-spectrum conclusion is similarity-invariant and is unaffected by the choice.) The essential spectrum of a multiplication operator is the essential range of its symbol, which is exactly the union of bands as stated. The interval form follows when spec(A_X) is itself an interval.

Consequence. The trapped-mode regime of §§7–8 corresponds to discrete eigenvalues of **T**_bulk + V lying in the *spectral gaps* of spec_ess(**T**_bulk) — the complement $\mathbb{R} \setminus \bigcup_i \mathfrak{B}_i$. The gap region above the ½-band and below the Perron band is, under the strong-coupling regime (W'),

$\mathcal{G}_{trap} := (\tfrac{1}{2} + \gamma\cdot\rho(A_X), 1 - \gamma\cdot\rho(A_X))$,

a non-empty interval whenever $\gamma\cdot\rho(A_X) < \tfrac{1}{4}$. This is where the Stage VIII trapped-mode regime $\tfrac{1}{2} < |\lambda_{trap}| < 1$ actually lives in the present global framework: not above ½ tout court, but above $\tfrac{1}{2} + \gamma\cdot\rho(A_X)$, with the gap-edge correction $\gamma\cdot\rho(A_X)$ being the substrate-coupling cost of the local-to-global transition.

3.2 Band Structure

For an infinite regular lattice X with spec(A_X) filling an interval (or union of intervals), Proposition 3.1 gives:

- The Perron eigenvalue 1 of $\hat{T}$ combined with the Perron eigenvalue $\rho(A\_X)$ of A_X gives the maximum of spec(**T**_bulk), namely $1 + \gamma \cdot \rho(A\_X)$. For γ in the regime (W), this is $< 1 + ½ = 3/2$.
- The single Perron eigenvalue of $\hat{T}$ at the value 1 becomes the **Perron band** of **T**_bulk: an interval $[1 - \gamma \cdot \rho(A\_X), 1 + \gamma \cdot \rho(A\_X)]$ (when $spec(A\_X) = [-\rho(A\_X), \rho(A\_X)]$).
- Each of the 7 wheel eigenvalues $\{1, ½, ½, -¼, -3/28, 0, 0\}$ generates a corresponding band of width $2\gamma \cdot \rho(A\_X)$, centred at the eigenvalue.

Important note on stochastic structure. Although the bulk operator **T**_bulk inherits stochastic structure from $\hat{T} \otimes I\_X$, the additive coupling $\gamma \cdot I\_\mathcal{K} \otimes A\_X$ is *not* stochastic — it generally pushes the spectral radius of **T**_bulk above 1 (specifically to $1 + \gamma \cdot \rho(A\_X)$). This is acceptable for the present paper because we are interested in **T** as a *transport / spectral* operator on $\mathscr{H}$, not as a stochastic operator on probability distributions over $\mathcal{K} \times X$. A stochasticity-preserving variant would require a different coupling structure (e.g., replacing A_X with a normalised Laplacian $I\_X - D\_X^{-1} A\_X$ and adjusting $\hat{T} \otimes I\_X$ correspondingly to balance flux); we do not pursue this variant here but flag it as a structural option in §13.

What matters for the trapped-mode analysis is that the bulk *essential* spectrum has a well-defined structure into which a defect-induced eigenvalue may or may not embed. For the operator **T**_bulk of Definition 2.1 in the regime (W), this structure is exactly the 7-band convolution of Proposition 3.1.

Definition 3.3 — Bulk Transport Bands

For each $i \in \{1, \ldots, 7\}$, the **i-th transport band** is

$$\mathfrak{B}\_i := \{ \lambda\_i + \gamma \cdot \mu : \mu \in spec(A\_X) \} \subseteq \mathbb{R}.$$

In the regime (W) with $spec(A\_X) \subset [-\rho(A\_X), \rho(A\_X)]$, the bands $\mathfrak{B}\_i$ are subsets of intervals $[\lambda\_i - \gamma \cdot \rho(A\_X), \lambda\_i + \gamma \cdot \rho(A\_X)]$.

Remark 3.4 — Five distinct bands, not seven. The canonical wheel has 7 eigenvalues counted with multiplicity but only 5 *distinct* eigenvalues: $\{1, ½, -¼, -3/28, 0\}$ with multiplicities (1, 2, 1, 1, 2). The eigenvalues ½ and 0 each appear twice, so the corresponding bands $\mathfrak{B}\_\{½\}$ and $\mathfrak{B}\_0$ coincide pairwise — the global operator **T**_bulk has *5 distinct transport bands*, not 7, each with doubled spectral multiplicity at the formerly-degenerate values. This is implicit in the §3.4 substrate-example tables and in the §11.2 worked example, which both work with the 5-distinct-band picture. References to "at most 7 bands" elsewhere in the paper refer to the multiplicity-counted version; the canonical-substrate band structure has exactly 5 distinct bands.

3.3 Band Separation in the Weak-Coupling Regime

Under (W), the bands are well-separated. The minimum distance between distinct wheel eigenvalues is

$$\Delta\_\hat{T} := \min \{ |\lambda\_i - \lambda\_j| : \lambda\_i \neq \lambda\_j, \lambda\_i, \lambda\_j \in spec(\hat{T}) \} = 3/28,$$

attained between the eigenvalues 0 and −3/28. (The other inter-eigenvalue gaps in spec($\hat{T}$) = {1, ½, −¼, −3/28, 0} — with multiplicities (1, 2, 1, 1, 2) — are ½ (between 1 and ½, and again between ½ and 0), 1/7 (between −3/28 and −¼), and ¼ (between 0 and −¼), all strictly larger than 3/28.)

The bands $\mathfrak{B}_i$ and $\mathfrak{B}_j$ corresponding to eigenvalues $\lambda_i \neq \lambda_j$ are separated by at least

$$|\lambda_i - \lambda_j| - 2\gamma\cdot\rho(A_X) \geq \Delta_{\hat{T}} - 2\gamma\cdot\rho(A_X).$$

Under (W), $2\gamma\cdot\rho(A_X) < 1$, but this is not sufficient to ensure separation for *all* band pairs; the tighter requirement is

$$2\gamma\cdot\rho(A_X) < \Delta_{\hat{T}} = 3/28,$$

i.e. $\gamma < 3/(56\cdot\rho(A_X))$. We strengthen (W) to

$$\gamma \cdot \rho(A_X) < 3/56, \quad (W')$$

for any application requiring complete band separation. The looser (W) is sufficient for separation of the Perron band from the rest (since the gap from 1 to ½ is ½ ≫ 3/28), which is what governs the bulk subdominant modulus and the corresponding **T**-level analogue of the Stage VIII gap $\varepsilon_{gap} = ½$. For the bulk subdominant modulus we have:

$$\lambda_{subdominant}(\mathbf{T}_{bulk}) := \max \{ |\lambda| : \lambda \in spec(\mathbf{T}_{bulk}), \lambda \neq \text{Perron} \} = ½ + \gamma\cdot\rho(A_X),$$

assuming the Perron band remains isolated above the ½-band. Under (W), this is < 1, so the Perron eigenvalue $1 + \gamma\cdot\rho(A_X)$ of $\mathbf{T}_{bulk}$ remains the unique largest eigenvalue, and the global spectral gap of $\mathbf{T}_{bulk}$ from its Perron value is

$$\varepsilon_{gap}(\mathbf{T}_{bulk}) := \lambda_{Perron}(\mathbf{T}_{bulk}) - \lambda_{subdominant}(\mathbf{T}_{bulk}) = (1 + \gamma\cdot\rho(A_X)) - (½ + \gamma\cdot\rho(A_X)) = ½.$$

The Stage VIII canonical gap $\varepsilon_{gap} = ½$ survives at the global-operator level. This is the first concrete bridge between Stage VIII's local analysis and the present global structure: the canonical-vacuum value of the local spectral gap $\varepsilon_{gap} = ½$ is exactly the global-operator gap $\varepsilon_{gap}(\mathbf{T}_{bulk})$, modulo an overall shift of the Perron eigenvalue by $\gamma\cdot\rho(A_X)$. The bridge is made precise in §10.

3.4 Explicit Substrate Examples

To make the band structure concrete, we compute spec(A_X), $\rho(A_X)$, and the resulting transport bands for four representative substrate choices. Each illustrates a different qualitative regime.

(a) X = $\mathbb{Z}$ (one-dimensional integer lattice, z = 2). The adjacency operator $A_{\mathbb{Z}}$ acts as $(A_{\mathbb{Z}} f)(n) = f(n-1) + f(n+1)$. By Fourier transform on $\mathbb{Z}$, spec($A_{\mathbb{Z}}$) = [−2, 2] (a continuous interval),

with ρ(A_ℤ) = 2. Each transport band 𝔅_i is the interval [λ_i − 2γ, λ_i + 2γ]. The weak-coupling regime (W) requires γ < ¼; the strong band-separation regime (W') requires γ < 3/112 ≈ 0.0268. The trap gap above the ½-band is 𝒢_trap = (½ + 2γ, 1 − 2γ), non-empty for γ < ⅛.

(b) X = ℤ² (two-dimensional square lattice, z = 4). A_{ℤ²} acts as (A_{ℤ²} f)(n, m) = f(n±1, m) + f(n, m±1). By Fourier transform, spec(A_{ℤ²}) = [−4, 4] (a continuous interval), with ρ(A_{ℤ²}) = 4. Each band is [λ_i − 4γ, λ_i + 4γ]; (W) requires γ < ⅛, (W') requires γ < 3/224 ≈ 0.0134, and 𝒢_trap = (½ + 4γ, 1 − 4γ) is non-empty for γ < 1/16. The local density of states at any site is the standard 2D tight-binding DOS, with logarithmic Van Hove singularities at μ = 0 — these produce sharp features in the resolvent integrals I_i(λ; x₀) of §7.3 near the band centres.

(c) Hexagonal (honeycomb) lattice, z = 3. The adjacency operator A_hex acts on a bipartite lattice and has spec(A_hex) = [−3, 3] with a *Dirac point* at μ = 0 (vanishing density of states linearly). ρ(A_hex) = 3, so bands are [λ_i − 3γ, λ_i + 3γ]; (W) requires γ < 1/6, (W') requires γ < 1/56. Structural feature: the Dirac point in spec(A_X) translates into linearly-vanishing DOS at the band centres of **T**_bulk, which suppresses the Birman–Schwinger resolvent integrals near λ_i and makes trapped-mode formation easier (smaller defects can produce trapped modes).

(d) Finite torus 𝕋_N^d (periodic substrate of size N^d). The adjacency operator on the d-dimensional torus has discrete spectrum spec(A_{𝕋_N^d}) = {2Σ_{k=1}^d cos(2π·n_k/N) : (n_1, …, n_d) ∈ {0, 1, …, N−1}^d}, with ρ(A_{𝕋_N^d}) = 2d. The spectrum is no longer continuous — each transport "band" is a finite set of N^d eigenvalues approximating the corresponding ℤ^d band as N → ∞. For N small, individual eigenvalue gaps within each band become significant, and the Birman–Schwinger criterion of §7 applies to *every* gap (between adjacent discrete bulk eigenvalues), not just to the inter-band gaps. This is the substrate-level analogue of finite-volume mesoscopic physics.

Numerical summary of the four substrates:

Substrate	**z**	**ρ(A_X)**	**(W): γ <**	**(W'): γ <**	**𝒢_trap non-empty: γ <**
ℤ	2	2	¼ = 0.250	3/112 ≈ 0.027	⅛ = 0.125
ℤ²	4	4	⅛ = 0.125	3/224 ≈ 0.013	1/16 = 0.063
Hex	3	3	⅙ ≈ 0.167	1/56 ≈ 0.018	1/12 ≈ 0.083
𝕋_N^d	2d	2d	1/(4d)	3/(112d)	1/(8d)

Structural lessons. Higher coordination z (hence higher ρ(A_X)) shrinks all admissible-γ regimes proportionally — the substrate becomes "harder to keep weakly coupled" as it becomes more connected. The DOS structure of A_X enters the Birman–Schwinger analysis through the resolvent integrals: smooth DOS (ℤ, ℤ²) gives smooth resolvent integrals; Van Hove singularities (ℤ²) give logarithmic features that can drive trapped-mode formation at specific λ; Dirac points (hexagonal) give linearly-vanishing DOS that *suppresses* resolvent integrals and shifts trapped-mode thresholds upward.

In what follows, when explicit calculations are needed (§§7.3, 11.2), we work with $\mathbb{Z}^d$-type substrates whose Fourier structure makes the resolvent integrals tractable. The qualitative content of §§7–8 (Birman–Schwinger criterion, Combes–Thomas localization) is substrate-independent.

---

4. Finite Coherence Propagation: A Lieb–Robinson Bound

The bounded spectral radius $\rho(A_X) \le z_{max}$ of the spatial-coupling operator, together with the locality of A_X (it couples only nearest neighbours), produces an explicit finite-speed propagation bound for **T**. This is the substrate-level analogue of the Lieb–Robinson bound familiar from quantum many-body theory.

Two velocities, not one. Because $\hat{T} \otimes I_X$ and $\gamma \cdot I_{\mathcal{K}} \otimes A_X$ act on disjoint tensor factors, they commute exactly, and the binomial expansion of $\mathbf{T}^n$ is exact rather than approximate:

$$\mathbf{T}^n = \Sigma_{k=0}^{n} C(n, k) \cdot \gamma^{n-k} \cdot \hat{T}^k \otimes A_X^{n-k}.$$

Since $A_X^{n-k}(x, y) = 0$ whenever $d_X(x, y) > n - k \le n$, the matrix element $\langle e_\kappa \otimes e_y, \mathbf{T}^n \cdot e_{\kappa'} \otimes e_x \rangle$ is *strictly zero* whenever $d_X(x, y) > n$. (Matrix elements throughout this section are taken in $L^2(\pi \otimes \mu_X)$ by the Choice of inner product preamble below; the strict-support statement is inner-product-independent — whether the entry is literally zero is an algebraic fact about the matrix, not a metric fact about the Hilbert space.) The substrate therefore carries two distinct propagation scales:

- **Strict light cone** $v_{strict} = 1$: zero amplitude beyond $d_X(x, y) > n$ per refinement step.
- **Effective propagation speed** $v_c := \gamma \cdot \rho(A_X)$: below this slope inside the strict cone, amplitudes are $O(1)$; above this slope, amplitudes are factorially suppressed.

The Lieb–Robinson statement controls the *interior* of the strict cone, sharpening the strict-support fact into a quantitative bound on the rate at which amplitude actually flows. Both statements belong in the theorem.

Choice of inner product. Throughout this section we work in the $L^2(\pi \otimes \mu_X)$ inner product introduced in §2.1, where π is the stationary measure of $\hat{T}$ on $\mathcal{K}$ and μ_X is the substrate measure with respect to which A_X is self-adjoint. In this inner product, when the spatial fibres preserve reversibility, $\hat{T}$ is self-adjoint, $\|\hat{T}\|_{op} \le 1$ (the standard $L^2(\pi)$-contractivity of a reversible Markov operator), and A_X is self-adjoint with $\|A_X\|_{op} = \rho(A_X)$. These two $L^2(\pi \otimes \mu_X)$ operator-norm bounds are what make the proof's Step 4 estimate clean. (The same inner-product convention is adopted in §8 for the Combes–Thomas analysis.)

Theorem 4.1 — Lieb–Robinson Bound for Coherence Transport

*Let **T** satisfy Definition 2.1 with $C = 0$ and γ in the regime (W), with spatial fibres preserving reversibility so that $\hat{T}$ is self-adjoint in $L^2(\pi)$. All inner products and matrix elements in the*

statement below are taken in $L^2(\pi \otimes \mu_X)$. Then for every $n \in \mathbb{N}$ and every pair of substrate positions $x, y \in X$ with $d := d_X(x, y)$:

(a) ***Strict support.*** *The matrix element vanishes outside the strict light cone $d \leq n$:*

$\langle e_\kappa \otimes e_y, \mathbf{T}^n \cdot e_{\kappa'} \otimes e_x \rangle_{L^2(\pi \otimes \mu_X)} = 0$ for all $\kappa, \kappa' \in \mathcal{K}$ whenever $d > n$.

(b) ***Quantitative decay inside the cone.*** *Within the strict cone $d \leq n$, the matrix element is bounded by*

$$|\langle e_\kappa \otimes e_y, \mathbf{T}^n \cdot e_{\kappa'} \otimes e_x \rangle_{L^2(\pi \otimes \mu_X)}| \leq (\gamma \cdot \rho(A_X) \cdot n)^d / d!,$$

using $\|\hat{T}\|_{op} \leq 1$ in $L^2(\pi)$ so the wheel-norm prefactor is bounded by 1.

(b') ***Sharper bound on the Perron-orthogonal sector.*** *For matrix elements $\langle \Psi, \mathbf{T}^n \Phi \rangle_{L^2(\pi \otimes \mu_X)}$ between states Ψ, Φ both orthogonal to the Perron eigenvector of $\hat{T} \otimes I_X$ (i.e., states with zero projection onto the joint Perron subspace $\mathbb{1}_{\mathcal{K}} / \sqrt{7} \otimes \mathbb{R}^X$), the sharper bound applies (here $\varepsilon_{gap} = ½$ is the canonical wheel value, since Theorem 4.1 uses a single $\hat{T}$ per Definition 2.1; for position-dependent $\hat{T}_x$ as in §10, the same bound holds with ε_{gap} replaced by $\inf_x \varepsilon_{gap}(x)$):*

$$|\langle \Psi, \mathbf{T}^n \Phi \rangle_{L^2(\pi \otimes \mu_X)}| \leq (1 - \varepsilon_{gap})^{n-d} \cdot (\gamma \cdot \rho(A_X) \cdot n)^d / d! \cdot \|\Psi\| \cdot \|\Phi\|,$$

where $d := \max\{d_X(x, y) : \Psi(\cdot, x) \neq 0, \Phi(\cdot, y) \neq 0\}$ is the maximum support-distance between Ψ and Φ. This is the substrate-level form of the sharper bound: the additional $(1 - \varepsilon_{gap})^{n-d}$ factor arises from the strict contractivity of $\hat{T}$ on the orthogonal complement of its Perron eigenvector, which applies to states already excluded from the Perron subspace but not to basis vectors (which carry non-zero Perron component $\pi(\kappa) \cdot 1$ in their decomposition).

(c) ***Effective light cone.*** *Defining the effective propagation speed $v_c := \gamma \cdot \rho(A_X)$, the matrix element decays super-exponentially outside the effective cone $d > e \cdot v_c \cdot n$:*

$|\langle e_\kappa \otimes e_y, \mathbf{T}^n \cdot e_{\kappa'} \otimes e_x \rangle_{L^2(\pi \otimes \mu_X)}| \leq \exp(-(d - e \cdot v_c \cdot n))$ whenever $d > e \cdot v_c \cdot n$.

Proof. We derive (a)–(c) directly from the commutativity of $\hat{T} \otimes I_X$ and $\gamma \cdot I_{\mathcal{K}} \otimes A_X$ in $L^2(\pi \otimes \mu_X)$.

Step 1: Exact binomial expansion. Since $[\hat{T} \otimes I_X, I_{\mathcal{K}} \otimes A_X] = 0$,

$$\mathbf{T}^n = (\hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X)^n = \Sigma_{k=0}^{n} C(n, k) \cdot \gamma^{n-k} \cdot \hat{T}^k \otimes A_X^{n-k}.$$

Note: this expansion uses $\hat{T}^k$ positive powers only — no inverse is required, so the formulation applies regardless of whether $0 \in \text{spec}(\hat{T})$ (which it does, with multiplicity 2 in the canonical wheel).

Step 2: Matrix-element extraction (in $L^2(\pi \otimes \mu_X)$). Taking the matrix element on both sides:

$\langle e_\kappa \otimes e_y, \mathbf{T}^n \cdot e_{\kappa'} \otimes e_x \rangle = \Sigma_{k=0}^{n} C(n, k) \cdot \gamma^{n-k} \cdot \langle e_\kappa, \hat{T}^k e_{\kappa'} \rangle_{L^2(\pi)} \cdot \langle e_y, A_X^{n-k} e_x \rangle_{L^2(\mu_X)}$.

Step 3: Strict support (part (a)). The substrate matrix element $\langle e_y, A_X^{n-k} e_x \rangle_{L^2(\mu_X)}$ counts μ_X-weighted edge-paths of length $n - k$ from x to y in the substrate graph. Such paths exist only when $n - k \geq d := d_X(x, y)$, i.e., $k \leq n - d$. The sum is therefore supported on $k \in \{0, 1, \ldots, n - d\}$, which is empty when $d > n$. This establishes part (a).

Step 4: Bound inside the cone (part (b)). For $d \leq n$, applying the $L^2(\mu_X)$ operator-norm bound $\|A_X^{n-k}\|_{op} \leq \rho(A_X)^{n-k}$ (using A_X self-adjoint, so spectral radius equals operator norm) gives $|\langle e_y, A_X^{n-k} e_x \rangle_{L^2(\mu_X)}| \leq \rho(A_X)^{n-k}$. The κ-fibre matrix element in $L^2(\pi)$ is bounded by the $L^2(\pi)$-operator norm of $\hat{T}^k$:

$|\langle e_\kappa, \hat{T}^k e_{\kappa'} \rangle_{L^2(\pi)}| \leq \|\hat{T}^k\|_{op, L^2(\pi)} \leq \|\hat{T}\|_{op, L^2(\pi)}^k \leq 1,$

where the last bound uses the standard fact that a reversible Markov operator has $L^2(\pi)$ operator norm ≤ 1 (Stage VI, with the proof following directly from the spectral decomposition: $\hat{T}$ on $L^2(\pi)$ has spectrum in $[-1, 1]$ with the Perron eigenvalue 1 isolated; the operator norm equals the spectral radius for self-adjoint $\hat{T}$). Hence

$|\langle e_\kappa \otimes e_y, \mathbf{T}^n \cdot e_{\kappa'} \otimes e_x \rangle_{L^2(\pi \otimes \mu_X)}| \leq \Sigma_{k=0}^{n-d} C(n, k) \cdot \gamma^{n-k} \cdot 1 \cdot \rho(A_X)^{n-k} = \Sigma_{m=d}^{n} C(n, m) \cdot (\gamma \cdot \rho(A_X))^m \leq (\gamma \cdot \rho(A_X) \cdot n)^d / d! \cdot (1 + \mathcal{O}(\gamma \cdot \rho(A_X)))$

for $d \geq e \cdot \gamma \cdot \rho(A_X) \cdot n$, using $C(n, m) \leq n^m / m!$ and the leading-term-dominance argument under (W). Absorbing the $(1 + \mathcal{O}(\ldots))$ factor gives the stated bound (b).

The sharper bound (b') follows from the same calculation restricted to the Perron-orthogonal subspace: if $\Psi, \Phi \perp \text{Perron}(\hat{T} \otimes I_X)$, then in the spectral decomposition $\hat{T} = \Sigma_i \lambda_i \Pi_i$, the projection $\Pi_{\text{Perron}} \Psi = 0$ and $\Pi_{\text{Perron}} \Phi = 0$, so the matrix element $\langle \Psi, \hat{T}^k \Phi \rangle_{L^2(\pi)}$ is bounded by $\|\hat{T}^k|_{\text{non-Perron}}\|_{op, L^2(\pi)} \cdot \|\Psi\| \cdot \|\Phi\| \leq (1 - \varepsilon_{gap})^k \cdot \|\Psi\| \cdot \|\Phi\|$. Substituting into the Step 4 binomial sum and using the same leading-term-dominance argument with the $m = n - k \to m = d$ index change gives the $(1 - \varepsilon_{gap})^{n-d}$ prefactor. The Perron-orthogonality restriction is essential: basis vectors $e_\kappa \otimes e_x$ have non-zero Perron component (the Perron eigenvector of $\hat{T}$ in $L^2(\pi)$ is $\mathbb{1}_{\mathcal{K}}/\sqrt{7}$, which has overlap $1/\sqrt{7}$ with every basis vector), so the sharper bound does not apply directly to the basis-vector matrix elements of (b).

Step 5: Effective light cone (part (c)). Applying Stirling's bound $d! \geq (d/e)^d$ to part (b):

$(v_c \cdot n)^d / d! \leq (e \cdot v_c \cdot n / d)^d \leq \exp(-(d - e \cdot v_c \cdot n))$

whenever $d > e \cdot v_c \cdot n$, using $\ln(e \cdot v_c \cdot n / d) \leq -(d - e \cdot v_c \cdot n)/d$ for the relevant range. This gives part (c).

Equivalent Heisenberg form (reversible substrates only). When the spectral fibres of $\hat{T}_x$ preserve reversibility and we work in the $L^2(\pi \otimes \mu_X)$ inner product, the matrix-element bound (b) translates into a bound on Heisenberg-evolved observable commutators $\|[\mathbf{T}^n \cdot O \cdot \mathbf{T}^{-n}, O']\|$ via the spectral theorem and unitary functional calculus. The Heisenberg form requires invertibility of $\mathbf{T}$ (which fails on the canonical substrate when $0 \in \mathrm{spec}(A_X)$, e.g., $\mathbb{Z}$, $\mathbb{Z}^2$, hexagonal lattice, since then $0 \in \mathrm{spec}(\mathbf{T}_{bulk})$ by Proposition 3.1), so the matrix-element form (a)–(c) is the universally applicable one. This is what the present theorem states.

Note on $C \neq 0$. The proof's Step 1 uses $[\hat{T} \otimes I_X, I_{\mathcal{K}} \otimes A_X] = 0$, which holds because the two operators act on disjoint tensor factors. A general closure-mixing C couples the $\mathcal{K}$ and X indices simultaneously and need not commute with either factor. For bounded C with operator norm $\|C\| \leq c_{max}$, the binomial expansion is replaced by a Duhamel-style expansion that contributes at most $c_{max} \cdot n$ to the spatial propagation budget per refinement step, modifying the effective velocity to

$$v_c^{(C)} := \gamma \cdot \rho(A_X) + c_{max}.$$

The strict-support statement (a) survives in modified form (zero outside $d_X(x, y) > n$ still, provided C has range 1 in the substrate index — i.e., C acts only between nearest-neighbour positions); the quantitative bound (b) acquires extra terms from the C-dependent Duhamel expansion; the effective light cone (c) shifts to use $v_c^{(C)}$ in place of v_c. We do not pursue the $C \neq 0$ case in detail; the relevant point is that finite propagation is robust to bounded closure-mixing coupling, with the effective velocity inflated additively by $\|C\|$. A future paper that fixes C from Stage IV would carry through the explicit modifications.

4.2 Interpretation: Strict vs Effective Cones

Theorem 4.1 distinguishes two propagation scales for amplitude on the substrate:

1. **Strict light cone** $v_{strict} = 1$. By part (a), amplitude is *exactly* zero beyond one substrate step per refinement step. This is an absolute kinematic constraint — the spatial-coupling operator A_X is nearest-neighbour, so the support of $\mathbf{T}^n$ on the spatial index can grow by at most one substrate step per power. No amplitude leaks beyond this strict boundary at any γ.
2. **Effective propagation speed** $v_c = \gamma \cdot \rho(A_X)$. Within the strict cone, amplitudes are not uniformly O(1): they remain non-negligible only up to a slope v_c, beyond which (still inside the strict cone) they are factorially suppressed by part (c). This is the substrate-level analogue of the *group velocity* of wave propagation in lattice systems — the speed at which the bulk of the amplitude actually flows, as opposed to the strict speed beyond which no amplitude is allowed.

The substrate-level analogue of "speed of light" is the *strict* cone $v_{strict} = 1$ — an absolute kinematic ceiling fixed by the substrate's combinatorial structure. The effective speed $v_c = \gamma \cdot \rho(A_X)$ is the analogue of the dynamical *propagation speed* — the rate at which signals actually travel, set by the coupling strength γ and the substrate's spectral structure $\rho(A_X)$. For $\mathbb{Z}^d$ with $z = 2d$, $v_c = 2\gamma d$, so higher-dimensional substrates have higher effective propagation speeds (per

refinement step) but the same strict ceiling v_strict = 1. The relevant quantity for actual signal transmission and trapped-mode localization is v_c, not v_strict; but the strict cone is the rigorous statement that supports the effective bound.

The sharper non-Perron version of the bound — part (b') of the theorem, applying to matrix elements between Perron-orthogonal states — is worth noting separately: for states Ψ, Φ already in the orthogonal complement of the Perron subspace (the natural setting for "deviation from the stationary state"), the wheel is strictly contractive on this subspace, and the prefactor $(1 - \varepsilon_{gap})^{n-d}$ contributes additional exponential decay with rate $\ln(1/(1 - \varepsilon_{gap})) \approx \varepsilon_{gap}$ per refinement step. This is the substrate-level mechanism by which propagation-mediated correlations decay even within the effective cone for non-equilibrium states, complementing the factorial-in-d suppression outside the cone with an exponential-in-(n − d) suppression along the time axis. The restriction to Perron-orthogonal states is essential — basis vectors $e_\kappa \otimes e_x$ have a non-zero Perron component (overlap $1/\sqrt{7}$ with the joint Perron eigenvector $\mathbb{1}_{\mathcal{K}}/\sqrt{7} \otimes \mathbb{R}^X$), so the sharper bound does not apply directly to the basis-vector matrix elements but only to their Perron-orthogonal projections.

4.3 Relation to Stage VIII

Stage VIII worked in a regime that is, retrospectively, the d = 0 (single-position) projection of the present theorem: at any single substrate position, the fibre dynamics is governed by $\hat{T}_x$ alone, with no spatial spreading visible. The Stage VIII silence on global dynamics is now seen to be a consequence of restricting attention to this d = 0 slice — correct for local-spectral analyses, but blind to the rate v_c at which the spatial coupling actually transmits amplitude.

---

5. Coherence-Transport Geometry

The finite-propagation structure of §4 induces a natural geometric structure on the substrate, built from the dynamics of **T** rather than from the substrate metric d_X directly.

Definition 5.1 — Combinatorial Coherence Distance

For substrate positions x, y ∈ X, the **combinatorial coherence distance** is

$$d_{coh}(x, y) := \inf \{ n \in \mathbb{N} : \langle e_\kappa \otimes e_y, \mathbf{T}^n (e_{\kappa'} \otimes e_x) \rangle \neq 0 \text{ for some } \kappa, \kappa' \in \mathcal{K} \}.$$

This is the minimum number of refinement steps required for amplitude to flow from x to y under **T**.

For **T** of Definition 2.1 with $\gamma > 0$, expanding $\mathbf{T}^n$ in powers of the spatial-coupling term shows that $\mathbf{T}^n(x, y) \neq 0$ iff $d_X(x, y) \leq n$. Therefore $d_{coh}(x, y) = d_X(x, y)$ in the regime where $\gamma > 0$ — the combinatorial coherence distance coincides with the graph distance.

This coincidence is a feature of the present minimal-coupling form. More general coupling structures (e.g., longer-range spatial coupling, or coupling restricted to a sublattice) would produce d_coh ≠ d_X in general, and the combinatorial coherence distance would be the right substrate-derived notion of distance rather than d_X.

Definition 5.2 — Spectral / Agmon Coherence Distance

A finer geometric structure is the **spectral coherence distance**, defined via the resolvent of **T**. For λ ∉ spec(**T**), the Combes–Thomas-type Agmon distance is

d_λ(x, y) := sup { ϕ(y) − ϕ(x) : ϕ : X → ℝ Lipschitz with constant 1, |⟨e_y, e^{−ϕ} R_0(λ) e^ϕ e_x⟩| ≤ C },

where R_0(λ) := (**T**_bulk − λI)^{−1} is the bulk resolvent. This is the substrate-level analogue of the Agmon distance from continuum random Schrödinger operator theory, and is the natural geometric structure governing exponential decay of trapped-mode eigenvectors (§8).

The Agmon distance refines the combinatorial distance by weighting by the "local difficulty" of coherence propagation: regions where the resolvent decays rapidly contribute more to d_λ than regions where it decays slowly. For **T** of Definition 2.1 with constant γ, d_λ(x, y) is asymptotically proportional to d_X(x, y), with proportionality constant depending on the spectral distance from λ to the nearest bulk band edge.

5.3 Effective Coherence Metric and Transport Optimality

The combinatorial and Agmon distances of Definitions 5.1–5.2 are *substrate-geometric* in the sense that they reflect only the graph X and the spectrum of **T**_bulk. When the substrate carries a *spatially varying* spectral gap field ε_gap(x) — for example, in a substrate with localized defects (Stage VIII) or with a position-dependent local operator $\hat{T}_x$ — the relevant geometry becomes weighted by the local transport capacity.

Definition 5.3 — Effective Coherence Metric

For a substrate with position-dependent local operator $\hat{T}_x$ and corresponding local spectral gap field ε_gap(x), the **effective coherence metric** is the edge-weighted graph metric on X with edge weight

w_coh(x, y) := 1/ε_gap_{eff}(x, y), for y ∼ x,

where ε_gap_{eff}(x, y) := ½ · (ε_gap(x) + ε_gap(y)) is the gap averaged across the edge. The induced **coherence-transport distance** is

g_coh(x, y) := inf_{paths π : x → y} Σ_{(u,v) ∈ π} w_coh(u, v),

the minimum total weighted length over substrate paths from x to y.

In the canonical-vacuum limit ε_gap(x) ≡ ½, the weights w_coh(x, y) ≡ 2 are uniform and g_coh reduces to 2·d_X (twice the graph distance). In the presence of a defect that depresses ε_gap below ½, edges through the defect region acquire larger weight: paths through depressed-gap regions become "longer" in the coherence-transport metric.

Proposition 5.4 — Transport Geodesics Minimize Coherence Attenuation (Heuristic)

Let X carry a position-dependent local operator $\hat{T}_x$ with Lipschitz local spectral gap field ε_gap : X → (0, 1]. ***Granting the heuristic identification v_c^local(x) ∝ γ·ρ(A_X)·ε_gap(x) of §9.1*** *(whose epistemic register is recorded there), define the local effective coherence velocity*

$$v_c^{\mathrm{local}}(x) := \gamma \cdot \rho(A_X) \cdot (\varepsilon_{\mathrm{gap}}(x) / \varepsilon_{\mathrm{gap}}^{(0)})$$

with ε_gap^{(0)} := ½ the canonical-vacuum value. Then over any substrate path π : x → y of total combinatorial length |π|, the total coherence transport time

$$\tau(\pi) := \Sigma_{u \in \pi}\ 1/v_c^{\mathrm{local}}(u)$$

is minimized by paths π that minimize the weighted-edge length g_coh(x, y) of Definition 5.3, modulo an overall factor of $(\gamma \cdot \rho(A_X) \cdot \varepsilon_{\mathrm{gap}}^{(0)})^{-1}$.*

Proof. The transport time accumulates additively over substrate steps. At each step u ∈ π, the local transport-time contribution is $1/v_c^{\mathrm{local}}(u) = (\varepsilon_{\mathrm{gap}}^{(0)}) / (\gamma \cdot \rho(A_X) \cdot \varepsilon_{\mathrm{gap}}(u))$. Summing over the path:

$$\tau(\pi) = (\varepsilon_{\mathrm{gap}}^{(0)})/(\gamma \cdot \rho(A_X)) \cdot \Sigma_{u \in \pi}\ 1/\varepsilon_{\mathrm{gap}}(u).$$

The right-hand side is, up to the constant prefactor, exactly the weighted path length with weights 1/ε_gap. This differs from the Definition 5.3 weight $w_{\mathrm{coh}}(x, y) = 1/\varepsilon_{\mathrm{gap}}^{\mathrm{eff}}(x, y)$ by a midpoint-vs-endpoint convention; on continuous interpolation between substrate sites the two coincide. Minimizing τ(π) over paths π : x → y is therefore equivalent to minimizing g_coh(x, y).

Structural reading. Proposition 5.4 establishes that the *Stage VIII local-gap field* ε_gap(x), through the local coherence velocity v_c^local(x) ∝ ε_gap(x), induces a weighted-graph geometry on the substrate whose geodesics — minimum-transport-time paths — naturally avoid regions of depressed gap. This is the substrate-level analogue of Fermat's principle of least time in optical geometry: coherence "rays" propagate along paths that minimize total transport time, equivalently along geodesics of the coherence-transport metric g_coh.

5.4 What This Geometry Is and Is Not

The transport geometry of §5 is *combinatorial* (Definition 5.1) and *spectral* (Definition 5.2). It is not Riemannian, and the §5 structures do not produce a metric tensor, connection, or curvature tensor. The relation to the candidate-curvature indicator $R(x) = \nabla^2 \varepsilon_{\mathrm{gap}}(x)$ of Stage VIII is via

the spectral side: spatial variation of ε_gap(x) produces spatial variation of the local spectral distance to bulk bands, which produces spatial variation of the local Agmon-distance scale. The §5 transport geometry is the operator-theoretic substrate from which a tensorial geometric theory might eventually be derived (Stage VIII §12.1), but the present paper does not undertake this derivation.

---

6. Localized Defects as Finite-Rank Perturbations

We now incorporate localized defects of the Stage VIII type into the global transport framework.

Definition 6.1 — Defect Perturbation

Following Stage VIII §2, a *localized defect at position x_0 with radius r* is encoded by a perturbation operator V on $\mathscr{H}$ defined as follows. Let $\Delta T_{x_0,r}(x)$ be the Stage VIII defect profile (zero outside $B_r(x_0)$, satisfying admissibility conditions D1–D3 of Stage VIII). The corresponding global perturbation is

$$V := \Sigma_{x \in B_r(x_0)} (\Delta T_{x_0,r}(x)) \otimes |x\rangle\langle x|, \quad (6.1)$$

acting on $\mathscr{H}$ as: at each substrate position $x \in B_r(x_0)$, apply the local matrix $\Delta T_{x_0,r}(x)$; at positions outside $B_r(x_0)$, act as zero. The full coupled-and-defected operator is

$$\mathbf{T} := \mathbf{T}_{\text{bulk}} + V.$$

Proposition 6.2 — Finite-Rank Structure

V is a finite-rank operator on $\mathscr{H}$, with

$$\text{rank}(V) \leq |B_r(x_0)| \cdot K = |B_r(x_0)| \cdot 7.$$

Proof. V is supported on the finite subspace $\mathbb{R}^{\mathcal{K}} \otimes \text{span}\{e_x : x \in B_r(x_0)\}$, which has dimension at most $|B_r(x_0)| \cdot 7$. Its range is contained in this subspace, giving the rank bound.

This finite-rank property is essential: it is what makes the Birman–Schwinger criterion (§7) applicable and computationally finite-dimensional.

6.3 Stage VIII Defects Re-cast

The Stage VIII boundary defect (§9.1) $\Delta T = e_{\kappa_{b_1}} \otimes (e_{\kappa_{b_1}} - e_{\kappa_{b_2}})^T$ applied at substrate position x_0 has outer-product form, hence gives the present V of rank 1. The Stage VIII hub-coupling defect (§9.6) $\Delta T_{\text{hub}} = e_{\kappa_h} \otimes (e_{\kappa_h} - e_{\kappa_{b_1}})^T$ similarly gives rank-1 V.

For these single-site rank-1 defects, the Birman–Schwinger operator K(λ) := V·R_0(λ) has rank at most 1, making the eigenvalue-1 condition of Theorem 7.1 below a single scalar equation rather than a matrix eigenvalue problem. This is what makes the worked examples of §11 explicitly computable.

---

7. The Birman–Schwinger Criterion for Trapped Coherence Modes

We now state the central trapped-mode criterion.

Theorem 7.1 — Birman–Schwinger Criterion

*Let **T**_bulk satisfy Definition 2.1 (with C = 0 for simplicity; the general case is analogous), and let V be a finite-rank perturbation as in Definition 6.1. Let R_0(λ) := (**T**_bulk − λI)^{−1} denote the bulk resolvent.*

*Then for any λ ∉ spec_ess(**T**_bulk):*

λ ∈ spec(**T**_bulk + V) if and only if 1 ∈ spec(V · R_0(λ)).

Equivalently, defining the Birman–Schwinger operator

K(λ) := V · R_0(λ),

*the discrete eigenvalues of **T**_bulk + V outside spec_ess(**T**_bulk) are exactly the values of λ at which K(λ) has eigenvalue 1.*

Proof. The eigenvalue equation (**T**_bulk + V)Ψ = λΨ for λ ∉ spec_ess(**T**_bulk) rewrites as

(**T**_bulk − λI)Ψ = −V·Ψ.

Applying R_0(λ) to both sides:

Ψ = −R_0(λ)·V·Ψ.

Since V has finite rank, V·Ψ lies in the finite-dimensional range of V. Let Φ := V·Ψ ∈ Range(V). Then Φ = V·Ψ = −V·R_0(λ)·V·Ψ = −V·R_0(λ)·Φ, which rearranges to

(I + V·R_0(λ))·Φ = 0 on Range(V).

Hence Φ is in the null space of (I + K(λ)) restricted to Range(V) — equivalently, −1 is an eigenvalue of K(λ)|_{Range(V)}, or equivalently 1 is an eigenvalue of −K(λ).

The sign convention here depends on whether one writes the criterion as $K(\lambda) = V \cdot R_0(\lambda)$ (giving "$-1 \in \mathrm{spec}(K)$") or $K(\lambda) = -V \cdot R_0(\lambda)$ (giving "$1 \in \mathrm{spec}(K)$"). We adopt the second convention throughout this paper for consistency with the standard Birman–Schwinger statement in the random Schrödinger operator literature. The criterion is therefore

$1 \in \mathrm{spec}(-V \cdot R_0(\lambda)) \iff \lambda \in \mathrm{spec}(\mathbf{T}_{\mathrm{bulk}} + V),\ \lambda \notin \mathrm{spec}_{\mathrm{ess}}(\mathbf{T}_{\mathrm{bulk}})$.

Conversely, if $1 \in \mathrm{spec}(-V \cdot R_0(\lambda))$ with eigenvector $\Phi \in \mathrm{Range}(V)$, then $\Psi := R_0(\lambda) \cdot \Phi$ satisfies $(\mathbf{T}_{\mathrm{bulk}} - \lambda I)\Psi = \Phi = -V \cdot R_0(\lambda) \cdot \Phi = -V \cdot \Psi$, so $(\mathbf{T}_{\mathrm{bulk}} + V)\Psi = \lambda\Psi$. The vector Ψ is non-zero because $R_0(\lambda)$ is invertible on the orthogonal complement of $\{0\} \subset \mathrm{Range}(V)$ and $\Phi \neq 0$.

7.2 What the Criterion Gives

Theorem 7.1 reduces the global eigenvalue problem $(\mathbf{T}_{\mathrm{bulk}} + V)\Psi = \lambda\Psi$, posed on the full state space $\mathscr{H}$ of dimension $7 \cdot |X|$, to the finite-dimensional eigenvalue problem $K(\lambda)\Phi = \Phi$ posed on $\mathrm{Range}(V)$, of dimension at most $7 \cdot |B_r(x_0)|$. For the single-site rank-1 defects of Stage VIII §§9.1, 9.6, this is a single scalar equation:

$\langle v, R_0(\lambda) \cdot u \rangle = -1/\alpha$, (single-site rank-1 case)

where the defect is $\alpha \cdot u \cdot v^{\mathrm{T}}$ at substrate position x_0 and $\langle v, R_0(\lambda) \cdot u \rangle$ denotes the scalar resolvent matrix element at x_0.

7.3 Computing $R_0(\lambda)$ for the Toy Form

For the bulk operator $\mathbf{T}_{\mathrm{bulk}} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X$, the resolvent factorises via the joint eigenbasis of $\hat{T}$ and A_X. Using the spectral decomposition $\hat{T} = \Sigma_i \lambda_i \cdot \psi_i^{(R)} (\psi_i^{(L)})^{\mathrm{T}}$ and $A_X = \Sigma_\mu \mu \cdot P_\mu$ (where P_μ is the spectral projection onto the eigenspace at $\mu \in \mathrm{spec}(A_X)$):

$R_0(\lambda) = \Sigma_i \Sigma_\mu (\lambda - \lambda_i - \gamma \cdot \mu)^{-1} \cdot (\psi_i^{(R)} (\psi_i^{(L)})^{\mathrm{T}}) \otimes P_\mu$.

For λ in the gap region $|\lambda - \lambda_i| > \gamma \cdot \rho(A_X)$ for all i (i.e., λ outside all bulk bands), all denominators are non-zero, so $R_0(\lambda)$ is well-defined.

For the matrix element at substrate position x_0:

$\langle e_\kappa \otimes e_{x_0}, R_0(\lambda) \cdot e_{\kappa'} \otimes e_{x_0} \rangle = \Sigma_i \Sigma_\mu (\lambda - \lambda_i - \gamma \cdot \mu)^{-1} \cdot \psi_i^{(R)}(\kappa) \cdot \psi_i^{(L)}(\kappa') \cdot |\phi_\mu(x_0)|^2$.

For an infinite regular lattice X, the sum over μ becomes an integral over $\mathrm{spec}(A_X)$, and $\Sigma_\mu |\phi_\mu(x_0)|^2$ is the local density of states of A_X at x_0 (denoted $\nu_A(\mu; x_0)$). Then

$\langle e_\kappa \otimes e_{x_0}, R_0(\lambda) \cdot e_{\kappa'} \otimes e_{x_0} \rangle = \Sigma_i \psi_i^{(R)}(\kappa) \cdot \psi_i^{(L)}(\kappa') \cdot \int_{\mathrm{spec}(A_X)} (\lambda - \lambda_i - \gamma \cdot \mu)^{-1} \nu_A(\mu; x_0)\, d\mu$.

The integrals are standard for explicit substrates (e.g., for $\mathbb{Z}^d$ they reduce to elliptic-type integrals over the Brillouin zone). The trapped-mode condition then becomes a scalar equation in λ whose roots are the discrete eigenvalues of **T** in the gap region.

7.4 Existence of Trapped Modes

The Birman–Schwinger criterion does not automatically guarantee that trapped modes exist; it gives a *condition* for existence. For weak defects (small α), the resolvent integrals are $\mathcal{O}(1)$ at λ near a wheel eigenvalue, while α appears linearly, so the criterion $1 = \alpha \cdot \langle v, R_0(\lambda) \cdot u \rangle$ requires α to be at least of order $1/\langle v, R_0(\lambda) \cdot u \rangle \approx 1$ for any λ-root to exist.

This gives a quantitative version of the Stage VIII §7.5 threshold α_{trap}: for the present toy form of **T**, α_{trap} is the smallest α at which the Birman–Schwinger scalar equation acquires a root in the gap region $(\tfrac{1}{2} + \gamma \cdot \rho(A_X), 1 + \gamma \cdot \rho(A_X))$. For $\alpha < \alpha_{trap}$ no trapped mode exists; for $\alpha \geq \alpha_{trap}$, at least one does.

The Stage VIII §9.1 boundary defect at $\alpha = 0.20$ produces local-operator eigenvalue $\lambda_2(\hat{T}_{\{x_0\}}) \approx 0.55$. The corresponding global trapped-mode condition under the present transport machinery is whether 0.55 lies above the ½-band of $\mathbf{T}_{bulk}$, which is $[\tfrac{1}{2} - \gamma \cdot \rho(A_X), \tfrac{1}{2} + \gamma \cdot \rho(A_X)]$. For this λ_2 to lie *above* the band, we need $0.55 > \tfrac{1}{2} + \gamma \cdot \rho(A_X)$, i.e., $\gamma \cdot \rho(A_X) < 0.05$ — substantially smaller than the (W) regime bound $\gamma \cdot \rho(A_X) < \tfrac{1}{2}$. So for substrates with weak transport coupling, the Stage VIII §9.1 defect at $\alpha = 0.20$ *does* produce a trapped mode by the Birman–Schwinger criterion; for stronger transport coupling, the same defect strength fails to produce a trapped mode because the bulk ½-band has widened to absorb it.

This is the substrate-level structural picture: trapped-mode existence depends jointly on defect strength *and* on the substrate's transport-coupling strength. The Stage VIII analysis was implicitly at the $\gamma \to 0$ limit (where bands collapse to wheel eigenvalues and any $\lambda_2 > \tfrac{1}{2}$ trivially lies outside the bulk spectrum); the present global analysis shows that finite γ widens the bulk bands and raises the bar for trapped-mode formation.

---

8. Combes–Thomas Exponential Localization

The Birman–Schwinger criterion gives the eigenvalue half of Stage VIII Theorem 6.1's hypothesis. The eigenvector half — exponential localization of the trapped eigenvector — follows from a Combes–Thomas-type resolvent estimate.

Choice of inner product. Throughout this section we work in the $L^2(\pi \otimes \mu_X)$ inner product introduced in §2.1, where π is the stationary measure of $\hat{T}$ on $\mathcal{K}$ and μ_X is the substrate measure inherited from Stage IV. In this inner product, when the spatial fibres $\bar{\hat{T}}_x$ preserve reversibility (Stage VIII §3.3, §6.3), the operator $\hat{T}$ is self-adjoint, the spatial-coupling operator A_X is self-adjoint (by hypothesis S2 of §2.2), and the coupled bulk operator

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

is self-adjoint as a sum of self-adjoint operators on the tensor-product Hilbert space. The spectral theorem then gives the resolvent bound

$\|R_0(\lambda)\|_{op} = 1/\mathrm{dist}(\lambda, \mathrm{spec}(\mathbf{T}_{bulk})) = 1/\delta$

cleanly. This is a *requirement* of the Combes–Thomas argument as stated, not merely a convenience: for non-reversible $\mathbf{T}_{bulk}$, the resolvent norm can exceed $1/\delta$ through pseudospectral effects, and the decay rate η acquires a pseudospectral correction depending on the condition number of the eigenbasis of $\hat{T}$. The non-reversible case is flagged in §14 as an open extension and is not the default setting of this section.

Theorem 8.1 — Combes–Thomas Estimate (Reversibility-Preserving Case)

Let $\boldsymbol{T}_{bulk}$ satisfy Definition 2.1 in the regime (W), with the spatial fibres $\hat{T}_x$ preserving reversibility so that $\boldsymbol{T}_{bulk}$ is self-adjoint in $L^2(\pi \otimes \mu_X)$. Let $\lambda \notin \mathrm{spec}(\boldsymbol{T}_{bulk})$ with spectral distance $\delta := \mathrm{dist}(\lambda, \mathrm{spec}(\boldsymbol{T}_{bulk})) > 0$. Then the bulk resolvent $R_0(\lambda)$ satisfies the exponential decay bound

$$|\langle e_\kappa \otimes e_x, R_0(\lambda) \cdot e_{\kappa'} \otimes e_y \rangle| \le (2/\delta) \cdot e^{-\eta \cdot d_X(x, y)},$$

for every $x, y \in X$ and $\kappa, \kappa' \in \mathcal{K}$, where the decay rate is

$\eta := \min(1, \delta/(2\gamma \cdot \rho(A_X)))$ (Combes–Thomas rate).

Proof. We give the explicit weighted-operator-norm derivation in the $L^2(\pi \otimes \mu_X)$ inner product.

Step 1: Gauge transformation setup. Fix a 1-Lipschitz weight function $\phi : X \to \mathbb{R}$ (so $|\phi(x) - \phi(y)| \le d_X(x, y)$ for all $x, y \in X$). Define the multiplication operator $(M_\phi \Psi)(\kappa, x) := e^{\phi(x)} \cdot \Psi(\kappa, x)$. The gauge-transformed bulk operator is

$$\mathbf{T}_{bulk}^{(\phi)} := M_\phi \cdot \mathbf{T}_{bulk} \cdot M_\phi^{-1}.$$

Since M_ϕ commutes with $\hat{T} \otimes I_X$ (which acts only on the $\mathcal{K}$-fibre), this simplifies to

$$\mathbf{T}_{bulk}^{(\phi)} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X^{(\phi)},$$

where $A_X^{(\phi)} := M_\phi|_X \cdot A_X \cdot M_\phi|_X^{-1}$ is the gauge-transformed adjacency operator. Element by element:

$$(A_X^{(\phi)})_{xy} = e^{\phi(x) - \phi(y)} \cdot (A_X)_{xy},$$

with non-zero entries only when $y \sim x$.

Step 2: Bound on the gauge perturbation. The difference $A_X^{(\phi)} - A_X$ has entries $(e^{\phi(x) - \phi(y)} - 1) \cdot (A_X)_{xy}$ on edges. By 1-Lipschitzness of ϕ, $|\phi(x) - \phi(y)| \le 1$ on edges, so

$|e^{\phi(x) - \phi(y)} - 1| \le e - 1 \le e^{|\phi(x) - \phi(y)|} - 1 \le e^{\eta} - 1$ when ϕ has Lipschitz constant $\eta \le 1$ (we will use this with ϕ replaced by $\eta \cdot d_X(\cdot, x_0)$, which has Lipschitz constant η).

Substituting $\phi \to \eta \cdot d_X(\cdot, x_0)$ (1-Lipschitz when scaled by $\eta \le 1$) gives

$\|A_X^{(\eta)} - A_X\|_{op} \le \rho(A_X) \cdot (e^{\eta} - 1) \le \rho(A_X) \cdot 2\eta$ for $\eta \le \ln 2 \approx 0.69$,

using $e^{\eta} - 1 \le 2\eta$ in this range. Hence

$\|\mathbf{T}_{bulk}^{(\eta)} - \mathbf{T}_{bulk}\|_{op} = \gamma \cdot \|A_X^{(\eta)} - A_X\|_{op} \le 2\gamma \cdot \rho(A_X) \cdot \eta$.

Step 3: Resolvent perturbation. By the spectral theorem (using self-adjointness of $\mathbf{T}_{bulk}$ in the $L^2(\pi \otimes \mu_X)$ inner product), the unperturbed resolvent $R_0(\lambda) := (\mathbf{T}_{bulk} - \lambda I)^{-1}$ satisfies the sharp bound

$\|R_0(\lambda)\|_{op} = 1/\mathrm{dist}(\lambda, \mathrm{spec}(\mathbf{T}_{bulk})) = 1/\delta$.

This is where the reversibility hypothesis is genuinely load-bearing: the spectral-theorem identity $\|R_0(\lambda)\|_{op} = 1/\delta$ holds for self-adjoint (or, more generally, normal) operators. For non-normal $\mathbf{T}_{bulk}$ the resolvent norm can be substantially larger, governed by the pseudospectrum rather than the spectrum, and the Step 4 bound below acquires a multiplicative pseudospectral correction.

By the second-resolvent identity,

$R_0^{(\eta)}(\lambda) := (\mathbf{T}_{bulk}^{(\eta)} - \lambda I)^{-1} = R_0(\lambda) \cdot [I + (\mathbf{T}_{bulk}^{(\eta)} - \mathbf{T}_{bulk}) \cdot R_0(\lambda)]^{-1}$.

The geometric series for the inverse converges whenever $\|(\mathbf{T}_{bulk}^{(\eta)} - \mathbf{T}_{bulk}) \cdot R_0(\lambda)\|_{op} < 1$, i.e.,

$2\gamma \cdot \rho(A_X) \cdot \eta / \delta < 1 \iff \eta < \delta / (2\gamma \cdot \rho(A_X))$.

Under this condition,

$\|R_0^{(\eta)}(\lambda)\|_{op} \le \|R_0(\lambda)\|_{op} \cdot 1/(1 - 2\gamma \cdot \rho(A_X) \cdot \eta/\delta) \le (1/\delta) \cdot 2 = 2/\delta$

if we take $\eta \le \delta/(4\gamma \cdot \rho(A_X))$ (which gives the geometric-series ratio $\le$ ½ and consequent factor of 2). A cleaner choice of constants gives

$\eta := \min(1, \delta/(2\gamma \cdot \rho(A_X)))$ (Combes–Thomas rate),

which guarantees $\|R_0^{(\eta)}(\lambda)\|_{op} \le 2/\delta$. This is the **weighted resolvent bound**:

$\|e^{\eta \cdot d_X(\cdot, x_0)} \cdot R_0(\lambda) \cdot e^{-\eta \cdot d_X(\cdot, x_0)}\|_{op} \le 2/\delta$.

Step 4: Untwisting the gauge. The matrix element of the gauge-transformed resolvent at (x, y) is

$\langle e_\kappa \otimes e_x, R_0^{(\eta)}(\lambda) \cdot e_{\kappa'} \otimes e_y \rangle = e^{\eta \cdot (d_X(x, x_0) - d_X(y, x_0))} \cdot \langle e_\kappa \otimes e_x, R_0(\lambda) \cdot e_{\kappa'} \otimes e_y \rangle$.

Taking absolute values and using the operator-norm bound:

$|\langle e_\kappa \otimes e_x, R_0(\lambda) \cdot e_{\kappa'} \otimes e_y \rangle| = e^{\eta \cdot (d_X(y, x_0) - d_X(x, x_0))} \cdot |\langle e_\kappa \otimes e_x, R_0^{(\eta)}(\lambda) \cdot e_{\kappa'} \otimes e_y \rangle| \le e^{\eta \cdot (d_X(y, x_0) - d_X(x, x_0))} \cdot (2/\delta)$.

Choosing the gauge centred at the source position x (i.e., taking $x_0 \leftarrow x$ and using $d_X(y, x) - d_X(x, x) = d_X(y, x) - 0 = d_X(x, y)$) gives the stated bound

$|\langle e_\kappa \otimes e_x, R_0(\lambda) \cdot e_{\kappa'} \otimes e_y \rangle| \le (2/\delta) \cdot e^{-\eta \cdot d_X(x, y)}$.

The Combes–Thomas argument is complete.

Remark on domain of validity. The bound holds in the resolvent domain $\rho(\mathbf{T}_{bulk}) = \mathbb{C} \setminus \mathrm{spec}(\mathbf{T}_{bulk})$. For trapped-mode applications, the relevant λ lies in the trap gap $\mathcal{G}_{trap}$ (Theorem 3.2 consequence), where spectral distance δ from the bulk bands is positive and explicitly computable from the band-edge structure of §3.

8.2 Consequence: Exponential Localization of Trapped Eigenvectors

Corollary 8.2 — Trapped-Mode Eigenvector Localization

Under the hypotheses of Theorem 8.1 (reversibility-preserving spatial fibres, regime (W)), let V be a localized defect supported in $B_r(x_0)$, and let λ be a discrete eigenvalue of $\boldsymbol{T}_{bulk} + V$ with corresponding eigenvector Ψ_λ, satisfying $\lambda \notin \mathrm{spec}_{ess}(\boldsymbol{T}_{bulk})$ with spectral distance $\delta := \mathrm{dist}(\lambda, \mathrm{spec}(\boldsymbol{T}_{bulk})) > 0$. Then Ψ_λ is exponentially localized:

$\|\Psi_\lambda(\cdot, y)\|_{\mathcal{K}} \le C \cdot e^{-\eta \cdot d_X(y, x_0)}$, for all $y \in X$ with $d_X(y, x_0) > r$,

with the Combes–Thomas decay rate $\eta := \min(1, \delta/(2\gamma \cdot \rho(A_X)))$ and a constant C depending on V and $\Psi_\lambda(\cdot, x_0)$.

Proof. From the eigenvalue equation, $\Psi_\lambda = -R_0(\lambda) \cdot V \cdot \Psi_\lambda$ (rearrangement of $(\mathbf{T}_{bulk} + V)\Psi_\lambda = \lambda\Psi_\lambda$). Since V is supported in $B_r(x_0)$, $V \cdot \Psi_\lambda$ is supported in $B_r(x_0)$. For y with $d_X(y, x_0) > r$, applying Theorem 8.1 to the matrix element $\langle e_\kappa \otimes e_y, R_0(\lambda) \cdot V \cdot \Psi_\lambda \rangle$ gives

$|\Psi_\lambda(\kappa, y)| = |\langle e_\kappa \otimes e_y, R_0(\lambda) \cdot V \cdot \Psi_\lambda \rangle| = |\Sigma_{x' \in B_r(x_0)} \Sigma_{\kappa'} \langle e_\kappa \otimes e_y, R_0(\lambda) \cdot (e_{\kappa'} \otimes e_{x'}) \rangle \cdot (V \cdot \Psi_\lambda)(\kappa', x')| \le (2/\delta) \cdot e^{-\eta \cdot (d_X(y, x_0) - r)} \cdot \|V \cdot \Psi_\lambda\|$,

where the $e^{-\eta \cdot (d_X(y, x_0) - r)}$ factor uses $d_X(y, x') \ge d_X(y, x_0) - r$ for $x' \in B_r(x_0)$ and the triangle inequality. Absorbing the $e^{\eta \cdot r}$ factor into the constant gives the stated bound.

8.3 Localization Length

The localization length $\xi := 1/\eta$ of trapped modes is

$\xi = 2\gamma \cdot \rho(A\_X)/\delta$, (when $\delta < 2\gamma \cdot \rho(A\_X)$)

or $\xi = 1$ (one substrate step) when $\delta \geq 2\gamma \cdot \rho(A\_X)$. This has structural content:

- **Strongly localized modes** (large δ, far from bulk bands) have $\xi \ll 1$ — the mode is confined to within a few substrate positions of the defect.
- **Weakly localized modes** (small δ, near the band edge) have $\xi \gg 1$ — the mode extends over many substrate positions before its amplitude decays. As $\delta \to 0$ (the eigenvalue approaches the bulk band edge), $\xi \to \infty$ and the mode becomes delocalised, marking the boundary between trapped and bulk modes.

This is the substrate-level analogue of the way that bound-state wave functions in quantum mechanics become more extended as their binding energy decreases, eventually delocalizing into the continuum.

8.4 Non-Reversibility as an Open Extension

The Combes–Thomas estimate of Theorem 8.1 holds *as stated* in the $L^2(\pi \otimes \mu\_X)$ inner product when the spatial fibres preserve reversibility — this is what makes the spectral-theorem bound $\|R\_0(\lambda)\|_{op} = 1/\delta$ *available in Step 3 of the proof. When* $\Delta T\{x_0,r\}$ *breaks* reversibility, $\hat{T}\_x$ is non-normal in $L^2(\pi\_x)$, $\mathbf{T}\_{bulk}$ inherits non-normality, and the resolvent bound becomes a *pseudospectral* statement:

$\|R\_0(\lambda)\|\_{op} \leq \kappa(\mathbf{T}\_bulk) / \mathrm{dist}(\lambda, \mathrm{spec}(\mathbf{T}\_bulk))$,

with $\kappa(\mathbf{T}\_bulk)$ the condition number of the eigenbasis of $\mathbf{T}\_bulk$ (in the L^2 operator-norm sense), which can be substantially larger than 1. The Combes–Thomas decay rate η acquires a corresponding correction:

$\eta\_non\text{-}reversible := \min(1, \delta/(2\kappa(\mathbf{T}\_bulk)\cdot\gamma\cdot\rho(A\_X)))$,

so the localization length scales as $\xi \lesssim 2\kappa(\mathbf{T}\_bulk)\cdot\gamma\cdot\rho(A\_X)/\delta$ rather than the reversible $2\gamma\cdot\rho(A\_X)/\delta$. In the small-perturbation regime of Stage VIII §3.3, $\kappa(\hat{T}\_x) \leq C_0\cdot(1 + \mathcal{O}(\|\Delta T\_x\|\_{op}/\Delta\_0)) \leq \mathcal{O}(C_0)$ with $C_0 = \sqrt{(7/4)}$ the universal Bauer–Fike constant, so κ remains bounded and the qualitative content (exponential decay) survives, but the explicit rate carries the pseudospectral factor.

The full non-reversible Combes–Thomas analysis is not pursued here. The default setting of Theorem 8.1 and Corollary 8.2 is the reversibility-preserving case; non-reversibility is flagged in §14 as an open extension, paralleling Stage VIII §3.3's Hausdorff-distance handling of non-reversible defects at the local level.

9. Coherence Wells and Transport Distortion

The Stage VIII local-gap field ε_gap(x) acquires a transport-theoretic interpretation in the present framework.

9.1 Local Coherence Capacity (Heuristic)

The bulk coherence-transport velocity v_c = γ·ρ(A_X) is uniform across the substrate in the unperturbed case. Near a defect, the local effective transport velocity can be expected to be reduced, since coherence flowing through a defect region must be processed by a fibre with reduced spectral gap. The relationship between fibre relaxation rate and effective transport speed admits a heuristic estimate that we now state explicitly with its epistemic register.

The local subdominant modulus of the spatial fibre of **T** at position x is

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

The fibre-local $L^2(\pi_x)$ relaxation timescale is therefore $\tau_{relax}(x) \sim 1/\varepsilon_{gap}(x)$ (Stage VIII Corollary 8.1; this part is rigorous). Heuristically identifying "transport delay per refinement step at position x" with $\tau_{relax}(x)$ gives the effective transport velocity

$$v_c^{local}(x) \sim \gamma \cdot \rho(A_X) \cdot \varepsilon_{gap}(x)/\varepsilon_{gap}^{(0)}, \quad \text{(HEURISTIC — see epistemic note below)}$$

with $\varepsilon_{gap}^{(0)} = ½$ the canonical value. Where ε_gap(x) is depressed, $v_c^{local}(x)$ is heuristically expected to be reduced — coherence propagates more slowly through defect regions.

Epistemic note. The step from "fibre relaxation rate" to "effective transport velocity" is dimensionally suggestive but *not derived* from the present operator-theoretic framework. The spatial-coupling structure γ·A_X in Definition 2.1 is position-independent; the gap-field-dependence enters through how each local fibre processes incoming amplitude (the $\hat{T}_x$ part), not through how the spatial coupling transmits it (the γ·A_X part, which is the same at every site). A genuine derivation would proceed via a Born-series argument on the coupled operator with position-dependent fibres $\hat{T}_x$ and unperturbed spatial coupling γ·A_X, working through how amplitude actually flows across the bulk-to-defect interface step by step. We do not carry out this derivation here; $v_c^{local}(x) \propto \varepsilon_{gap}(x)$ is recorded as a heuristic dimensional estimate, and the downstream results that depend on it — Proposition 5.4 (transport optimality), Proposition 9.3 (coherence-ray bending), and the gap-field-dependent scaling line of Theorem 13.1 — inherit this heuristic register. The structural-unification content of those results is unaffected; what depends on the heuristic is the specific functional form $v_c \propto \varepsilon_{gap}$ rather than some other monotone increasing function of ε_gap. A future paper deriving the Born-series form would either confirm the ε_gap proportionality or correct it to ε_{gap}^p for some p — in either case preserving the qualitative picture (defects slow local transport) while sharpening the quantitative law.

9.2 Coherence Wells

A localized defect with $\varepsilon_{gap}(x_0) < \frac{1}{2}$ therefore creates a region of *reduced coherence velocity* — a coherence well. Coherence flowing into such a region propagates more slowly, accumulates locally (since outflow is slower than inflow), and produces the local entropy retention of Stage VIII Corollary 8.1.

In the trapped-mode regime, the well becomes deep enough that some coherence becomes permanently confined — the trapped mode of Theorem 7.1 with Corollary 8.2 localization.

9.3 Transport Distortion from Gap Gradients

The Proposition 5.4 transport-optimality result, applied to a substrate with a spatially varying ε_{gap} field, gives a quantitative refraction-type statement: coherence rays bend toward regions of suppressed gap.

Proposition 9.3 — Coherence Rays Bend Toward Suppressed-Gap Regions (Heuristic)

Granting the heuristic identification $v_c^{local}(x) \propto \varepsilon_{gap}(x)$ of §9.1 (whose epistemic register is recorded there), let X be a regular substrate with Lipschitz local spectral gap field $\varepsilon_{gap} : X \to (0, 1]$ and corresponding effective coherence metric g_{coh} of Definition 5.3. Let $\pi : x \to y$ be a coherence-transport geodesic (minimizer of g_{coh}). Then at any interior point u of π^* where ε_{gap} is differentiable in the continuum-limit sense, the geodesic tangent direction $\partial\pi^*/\partial s$ satisfies the discrete-Lagrangian condition*

$$\partial/\partial s \left[(1/\varepsilon_{gap}(\pi^*(s))) \cdot \partial\pi^*/\partial s\right] = -\nabla(1/\varepsilon_{gap})(\pi^*(s)) = \nabla\varepsilon_{gap}(\pi^*(s)) / \varepsilon_{gap}(\pi^*(s))^2.$$

Equivalently, the local rate-of-bending of coherence trajectories is proportional to the relative gradient $\nabla\varepsilon_{gap} / \varepsilon_{gap}$ of the gap field, with sign such that trajectories bend toward regions of smaller ε_{gap}.

Proof. Apply the discrete Euler–Lagrange equation to the action functional $S[\pi] := \Sigma_{u \in \pi} 1/\varepsilon_{gap}(u)$ of Proposition 5.4, with the substrate path π parametrised by arclength s. The Lagrangian is $L(u, \partial\pi/\partial s) = (1/\varepsilon_{gap}(u)) \cdot |\partial\pi/\partial s|$, and the Euler–Lagrange equation $\partial/\partial s (\partial L/\partial(\partial\pi/\partial s)) = \partial L/\partial u$ gives the stated form after canceling the (unit) speed factor $|\partial\pi/\partial s|$. The sign of bending follows from the negative gradient $\nabla(1/\varepsilon_{gap}) = -\nabla\varepsilon_{gap}/\varepsilon_{gap}^2$: trajectories bend in the direction of *increasing* $1/\varepsilon_{gap}$, i.e., *decreasing* ε_{gap}.

Reading guidance. Proposition 9.3 is the most directly geometry-suggestive result in this paper, and it requires careful epistemic discipline.

It is *not* a derivation of general relativity. No metric tensor is constructed, no Einstein equations are written down, and the "trajectory" here is a discrete substrate path, not a continuum world-line. The bending phenomenon arises from the *optical* analogy (Fermat's principle of least time

with spatially varying refractive index) rather than from a *gravitational* analogy (geodesic motion in a curved metric).

It *is* the substrate-level statement that the candidate-curvature indicator $R(x) = \nabla^2 \varepsilon_{gap}(x)$ of Stage VIII has a transport-dynamical interpretation: spatial variation of ε_{gap} produces measurable bending of coherence trajectories, with the bending rate set by the relative gap gradient. This is the operator-theoretic content underneath Stage VIII's structural curvature ansatz — it shows that the ansatz tracks a real dynamical phenomenon (coherence-ray bending) rather than being a purely formal scalar built from second differences.

The structural cleanness is worth noting: a positive Laplacian $\nabla^2 \varepsilon_{gap} > 0$ at the defect core (Stage VIII §9.4) corresponds to a local minimum of ε_{gap}, hence a local maximum of $1/\varepsilon_{gap}$, hence an attractor for coherence trajectories. The Stage VIII §9.4 zero-sum identity $\Sigma_x R(x) = 0$ then has a transport reading: trajectories that are deflected inward at the defect core are deflected outward at the neighbour ring, with net zero deflection over the defect cluster — a discrete analogue of the divergence-free curvature property Stage VIII §12.10 conjectured.

9.4 Transport-Theoretic Reading of the Defect-Coherence Principle

Stage VIII's Defect-Coherence Principle (§13.1) listed four ε_{gap}-functionals controlling continuum roughening, candidate curvature, entropy retention, and trapped-mode persistence. The present framework provides a unifying transport-theoretic reading: all four functionals describe the same underlying phenomenon — *local reduction of coherence-transport capacity*. Specifically:

- Continuum roughening $K_\infty \propto 1/\varepsilon_{gap}$ reflects the fact that smoother continuum geometry requires faster local equilibration, hence higher local transport capacity.
- Candidate curvature $R(x) = \nabla^2 \varepsilon_{gap}(x)$ measures the spatial variation of transport capacity — its concentration at a defect core and depletion in the surrounding ring is the candidate-curvature signature.
- Entropy retention $s(x, n) \leq (1 - \varepsilon_{gap}(x))^{2n} \cdot s(x, 0)$ reflects reduced local outflow capacity.
- Trapped-mode persistence (Theorem 7.1, Corollary 8.2) reflects complete local trapping of coherence: not just slowed transport, but bound-state formation outside the bulk transport bands.

The unifying transport reading does not derive the Stage VIII Defect-Coherence Principle (the four functionals were already structurally unified by their shared ε_{gap} dependence); it provides a *physical-interpretive* layer on top of the Stage VIII functional structure, in which "coherence-transport capacity" is the unifying notion.

---

10. The Local Spectral Gap as a Spatial-Fibre Projection

Stage VIII defined ε_gap(x) as the spectral-gap field of the *local* operator $\hat{T}_x$ — a quantity built from the spectrum of the 7×7 matrix at each substrate position. The present framework makes the relationship to the global operator **T** precise, via a fibrewise-projection theorem.

10.1 The Spatial Fibre

For each substrate position $x \in X$, the **spatial fibre** of **T** at x is the operator obtained by restricting **T** to the $\mathcal{K}$-fibre $\mathbb{R}^{\mathcal{K}} \otimes |x\rangle$ and projecting back to the same fibre. Formally, let $P_x : \mathscr{H} \to \mathbb{R}^{\mathcal{K}} \otimes |x\rangle$ be the orthogonal projection onto the position-x fibre. Define

$$\mathbf{T}|_{\mathcal{K}_x} := P_x \cdot \mathbf{T} \cdot P_x : \mathbb{R}^{\mathcal{K}} \otimes |x\rangle \to \mathbb{R}^{\mathcal{K}} \otimes |x\rangle,$$

acting on the closure catalogue at position x alone.

Theorem 10.1 — Stage VIII Gap Field as Fibrewise Projection

*Let **T** satisfy Definition 2.1 with a possibly position-dependent local refinement operator $\hat{T}_x$ at each $x \in X$. Then:*

(a) The spatial fibre at x is exactly the local refinement operator:

$$\mathbf{T}|_{\mathcal{K}_x} = \hat{T}_x.$$

(b) Consequently, the spatial-fibre spectral gap is exactly the Stage VIII local-gap field:

$$\varepsilon_{gap}(\mathbf{T}|_{\mathcal{K}_x}) := 1 - \max\{|\lambda| : \lambda \in \mathrm{spec}(\mathbf{T}|_{\mathcal{K}_x}), \lambda \neq 1\} = \varepsilon_{gap}(x),$$

where ε_gap(x) on the right-hand side is the Stage VIII local-gap field of Definition 3.1 (Stage VIII).

*(c) Equivalently, the Stage VIII local-gap field is the spectral-gap field of the fibrewise projection of the global transport operator **T** onto the spatial fibres.*

Proof. Writing $\mathbf{T} = (\hat{T}_x \otimes I_X) + \gamma \cdot I_{\mathcal{K}} \otimes A_X$ (the C = 0 case; the C ≠ 0 case is analogous after restricting the C action to the relevant fibre), the action on a state $\psi \otimes |x\rangle$ at fibre x is

$$\mathbf{T}(\psi \otimes |x\rangle) = (\hat{T}_x \psi) \otimes |x\rangle + \gamma \cdot \psi \otimes (A_X |x\rangle) = (\hat{T}_x \psi) \otimes |x\rangle + \gamma \cdot \Sigma_{\{y \sim x\}} \psi \otimes |y\rangle.$$

Applying P_x to both sides:

$$P_x \cdot \mathbf{T} \cdot (\psi \otimes |x\rangle) = (\hat{T}_x \psi) \otimes |x\rangle + \gamma \cdot \Sigma_{\{y \sim x\}} P_x (\psi \otimes |y\rangle) = (\hat{T}_x \psi) \otimes |x\rangle + 0,$$

since P_x projects onto the position-x fibre and $\psi \otimes |y\rangle$ has support disjoint from $|x\rangle$ for $y \neq x$. Hence $\mathbf{T}|_{\mathcal{K}_x} \cdot \psi = \hat{T}_x \psi$ for every $\psi \in \mathbb{R}^{\mathcal{K}}$, establishing (a).

Part (b) follows immediately from (a): the spectrum of $\mathbf{T}|_{\mathcal{K}_x}$ equals the spectrum of $\hat{T}_x$, and the gap-from-Perron is by definition $\varepsilon_{gap}(x)$.

Part (c) is the structural restatement: the Stage VIII analysis takes the fibrewise projection of the Stage IX global transport operator as its primary object, and works with the spectral-gap field of this projection.

10.2 Structural Consequence: Stage VIII as Fibrewise Projection Theory of Stage IX

Theorem 10.1 establishes that Stage VIII is, in retrospect, the **fibrewise projection theory** of Stage IX: every Stage VIII statement about $\varepsilon_{gap}(x)$, $R(x) = \nabla^2 \varepsilon_{gap}(x)$, the local Lipschitz constant $K_\infty(x) \propto 1/\varepsilon_{gap}(x)$, and local entropy retention $(1 - \varepsilon_{gap}(x))^{2n}$ can be re-derived as a statement about the spectral-gap field of the spatial fibres of **T**.

What Stage VIII could *not* see — and what required the full global operator **T** of Stage IX — was the *coupling between fibres*. The off-diagonal $\gamma \cdot A_X$ term in Definition 2.1, which Stage VIII implicitly set to zero by working only at single substrate positions, is exactly what produces:

- The bulk transport bands of §3 (since the bulk operator without coupling has spectrum $\{\lambda_i$, multiplicity $|X|\}$, with no band structure).
- The finite-coherence-velocity Lieb–Robinson bound of §4 (which is non-trivial only because A_X couples adjacent fibres).
- The Birman–Schwinger criterion of §7 (which involves the bulk resolvent $R_0(\lambda)$, defined only for the coupled operator).
- The Combes–Thomas localization of §8 (where the decay rate η depends on $\gamma \cdot \rho(A_X)$).

In the uncoupled limit $\gamma \to 0$, the Stage IX framework collapses to a disjoint union of Stage VIII fibres at each substrate position, with no cross-talk between positions. In the coupled regime $\gamma > 0$, the fibrewise picture remains *projection-correct* (Stage VIII statements about individual fibres remain true, by Theorem 10.1), but it is no longer *exhaustive* — the spatial coupling carries information about global transport that no individual fibre sees.

10.3 The Bridge from Local Stage VIII to Global Stage IX

Theorem 10.1 makes the Stage VIII ↔ Stage IX correspondence precise:

Stage VIII (local / fibrewise)	Stage IX (global / coupled)
Local operator $\hat{T}_x$	Spatial fibre $\mathbf{T}|_{\mathcal{K}_x}$ (Theorem 10.1(a))
Local spectral gap $\varepsilon_{gap}(x)$	Spectral gap of spatial fibre (Theorem 10.1(b))
Local trapped-mode necessary condition $\lvert\lambda_2(\hat{T}_x)\rvert > ½$	Strengthened to: λ_{trap} outside bulk band edges $\lvert\lambda\rvert > ½ + \gamma\cdot\rho(A_X)$ (in the trap-gap regime $\gamma\cdot\rho(A_X) < ¼$ where $\mathcal{G}_{trap}$ is non-empty)
Local eigenvector	Generically non-localized extension to $\mathscr{H}$; Combes–Thomas decay only for trapped eigenvalues

Stage VIII (local / fibrewise)	Stage IX (global / coupled)
Stage VIII Theorem 6.1 hypothesis	Verified by Birman–Schwinger (Theorem 7.1) + Combes–Thomas (Corollary 8.2)
Stage VIII Defect-Coherence Principle (4 ε_gap-functionals)	Inherits structure from fibrewise projection; transport interpretation in §9.4

The four ε_gap-functionals of Stage VIII §13.1's Defect-Coherence Principle inherit their structure from the corresponding properties of the global **T**: $K_\infty(x) \propto 1/\varepsilon_{gap}(x)$ reflects the spatial-fibre transport capacity; $R(x) = \nabla^2\varepsilon_{gap}(x)$ reflects the spatial variation of the spatial-fibre gap (and, by Proposition 9.3, the rate of coherence-ray bending); entropy retention reflects the spatial-fibre $L^2(\pi_x)$ contraction rate; and trapped-mode persistence reflects the existence of discrete eigenvalues of the full **T** outside the bulk bands of **T**_bulk.

10.4 What This Closes

The bridge between local Stage VIII analysis and global Stage IX transport closes Stage VIII's note on substrate-level universality (Stage VIII §11.1). The Stage VIII universality class $\mathcal{C}_{K=7}$ was a class of *single* operators; Stage VIII's "pointwise in the class" version of trapped-mode results was deferred pending the substrate-level extension. The present paper does not formalise the substrate-level universality class as such — that remains the Stage VIII §12.4 open problem — but it provides the global operator **T** in which such a class would naturally live, together with the spectral / resolvent machinery required for its analysis. Theorem 10.1 supplies the *projection map* between the two layers (local fibres ← coupled global), making the substrate-level universality class effectively a class of admissible global transport operators **T** whose fibrewise projections all lie in $\mathcal{C}_{K=7}$.

---

11. Closing Stage VIII's Trapped-Mode Conditional

The principal result of this paper is that the Stage VIII Theorem 6.1 trapped-mode hypothesis becomes verifiable rather than presumed.

11.1 The Stage VIII Hypothesis, Reformulated

Stage VIII Theorem 6.1 assumed:

(i) An eigenvalue λ_{trap} of **T** with $\frac{1}{2} < |\lambda_{trap}| < 1$. (ii) A corresponding eigenvector Ψ_{trap} exponentially decaying outside the defect region.

In the present framework, both conditions become checkable:

- **Condition (i)** is verified by the Birman–Schwinger criterion (Theorem 7.1): $\lambda \in \mathrm{spec}(\mathbf{T})$ outside $\mathrm{spec}_{ess}(\mathbf{T}_{bulk})$ iff $1 \in \mathrm{spec}(-V \cdot R_0(\lambda))$. For the toy form $\mathbf{T}_{bulk} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X$, the gap region above the ½-band is $(\frac{1}{2} + \gamma\cdot\rho(A_X),\ 1 - \gamma\cdot\rho(A_X))$

(assuming the Perron band is also isolated). Trapped-mode existence is the condition that the Birman–Schwinger scalar equation has a root in this interval.

- **Condition (ii)** is verified by Corollary 8.2: any trapped eigenvalue λ has spectral distance δ := dist(λ, spec_ess(**T**_bulk)) > 0, and its eigenvector decays at rate η = min(1, δ/(2γ·ρ(A_X))).

11.2 Worked Example — Stage VIII §9.1 Boundary Defect Revisited

For the Stage VIII §9.1 boundary defect ΔT = e_{κ_{b₁}} ⊗ (e_{κ_{b₁}} − e_{κ_{b₂}})^T at strength α, the global perturbation V is rank 1, supported at substrate position x_0. The Birman–Schwinger scalar equation for trapped-mode existence in the gap above the ½-band is

$$\alpha \cdot \langle v, R_0(\lambda)\cdot u\rangle_{x_0,x_0} = -1,$$

where u = e_{κ_{b₁}}, v = e_{κ_{b₁}} − e_{κ_{b₂}}, and ⟨·, ·⟩_{x₀,x₀} denotes the matrix element at substrate position x_0.

Using the spectral decomposition of $\hat{T}$ (which is computable from Stage VI), this matrix element is a sum over the 7 wheel eigenvalues with weights determined by the right and left eigenvectors of $\hat{T}$ at the relevant entries. For the present rank-1 perturbation, the explicit form is:

$$\langle v, R_0(\lambda)\cdot u\rangle_{x_0,x_0} = \Sigma_i \left[\psi_i^{(R)}(\kappa_{b_1}) - \psi_i^{(R)}(\kappa_{b_2})\right] \cdot \psi_i^{(L)}(\kappa_{b_1}) \cdot I_i(\lambda; x_0),$$

where $I_i(\lambda; x_0) := \int_{\mathrm{spec}(A_X)} (\lambda - \lambda_i - \gamma\cdot\mu)^{-1} \cdot \nu_A(\mu; x_0)\, d\mu$ is the local resolvent integral for the i-th wheel eigenvalue. For an infinite regular lattice with ν_A given by the local density of states of A_X at x_0, the integrals $I_i(\lambda; x_0)$ are explicit standard integrals (elliptic-type for $\mathbb{Z}^d$).

Setting the right-hand side of the Birman–Schwinger equation to −1 and solving for λ gives the trapped eigenvalue (when one exists).

γ = 0 reduction. At γ = 0, the bulk operator is **T**_bulk = $\hat{T} \otimes I_X$ (no spatial coupling), and the bulk resolvent factorises as $R_0(\lambda) = (\hat{T} - \lambda I)^{-1} \otimes I_X$. The Birman–Schwinger scalar equation collapses to

$$\alpha \cdot v^T \cdot (\hat{T} - \lambda I)^{-1} \cdot u = -1,$$

which by the Sherman–Morrison formula is exactly the eigenvalue equation $\det(\hat{T} + \alpha\cdot u\cdot v^T - \lambda I) = 0$ for the rank-1 perturbed local operator $\hat{T}_{x_0} = \hat{T} + \alpha\cdot u\cdot v^T$. The Birman–Schwinger root at γ = 0 is therefore the local-operator eigenvalue $\lambda_2(\hat{T}_{x_0}) \approx 0.55$ (from the Stage VIII numerics at α = 0.20).

Continuity in γ. For γ > 0 small, the Birman–Schwinger root depends continuously on γ (standard implicit-function-theorem argument: the BS equation is an analytic equation in (λ, γ), and the γ = 0 root $\lambda_0 = \lambda_2(\hat{T}_{x_0})$ is simple — simple here meaning simple as a root of the BS scalar

equation, which by the Sherman–Morrison reduction is equivalent to "the multiplicity of λ_0 as a root of the BS equation equals its multiplicity as an eigenvalue of $\hat{T}_{x_0}$ minus its multiplicity as an eigenvalue of the unperturbed $\hat{T}$"; for λ_0 ≈ 0.55, this is 1 − 0 = 1, hence simple). The root λ(γ) therefore extends to a continuous family for γ in a neighbourhood of 0. The trapped-mode condition is that λ(γ) remains in the spectral gap *above* the widened ½-band, i.e.,

λ(γ) > ½ + γ·ρ(A_X) and λ(γ) < 1 − γ·ρ(A_X).

To leading order in γ, λ(γ) = λ_0 + $\mathcal{O}(\gamma)$, while the gap floor ½ + γ·ρ(A_X) rises linearly in γ. The trapped mode therefore survives precisely when

λ_0 − ½ > γ·ρ(A_X) · (1 + $\mathcal{O}(\gamma)$),

i.e., γ·ρ(A_X) < λ_0 − ½ + $\mathcal{O}(\gamma^2)$. For the Stage VIII numerics with λ_0 ≈ 0.55, this gives γ·ρ(A_X) < 0.05 in leading order. This is *both necessary and sufficient*: necessary because the gap floor must lie below λ(γ), and sufficient because continuity ensures the BS root remains a root for γ small enough that the gap remains open.

Consequence. For γ·ρ(A_X) < 0.05, the Stage VIII §9.1 defect at α = 0.20 *does* produce a global trapped mode by the Birman–Schwinger criterion. For γ·ρ(A_X) > 0.05, the same defect strength fails to produce a trapped mode — the local-operator eigenvalue 0.55 is absorbed into the widened bulk ½-band. This is the quantitative substrate-level reading of the Stage VIII §9.5 "trapped-mode check" being inconclusive: the defect produces a local-spectrum bump above ½, but whether this survives globally depends on the substrate's transport-coupling strength.

11.3 Worked Example — Stage VIII §9.6 Hub-Coupling Defect (Spectrally Invisible)

The Stage VIII §9.6 hub-coupling defect $\Delta T_{hub} = e_{\kappa_h} \otimes (e_{\kappa_h} - e_{\kappa_{b_1}})^T$ was shown to leave the local-gap $\varepsilon_{gap}(x_0) = ½$ exactly invariant via a symmetry-protection mechanism (block-aligned to $H_{\sigma^*}^+$, with within-block right/left eigenvector orthogonality conditions). Six of seven wheel eigenvalues are unchanged; only −3/28 moves, and its modulus stays well below ½ throughout the admissible range.

In the present framework, this means the Birman–Schwinger scalar equation has its right-hand side $\langle v_{hub}, R_0(\lambda)\cdot u_{hub}\rangle_{x_0,x_0}$ structurally suppressed: the same orthogonality conditions that suppress the rank-1 perturbation's effect on the local spectrum also suppress the relevant matrix elements of the bulk resolvent. Specifically, the orthogonality relations $u^T \psi_L^{(½)} = 0$ and $v^T \psi_R^{(½)} = 0$ (one of which holds for the hub-coupling defect, the other does not) suppress the contribution of the i = ½ wheel eigenvalue to the resolvent integral $I_{½}(\lambda; x_0)$.

The consequence is that the Birman–Schwinger equation for the hub-coupling defect cannot acquire a root in the gap above the ½-band, regardless of α and γ. The defect is *globally* spectrally invisible — the symmetry protection extends from the local to the global level.

This is a structurally satisfying check: the Stage VIII §9.6 spectrally-invisible sub-class is preserved by the present global transport machinery. The localized analogue of the Stage VII silent class survives the lift from local to global.

11.4 Closing the Stage VIII §12.2 Open Problem

Stage VIII §12.2 identified the Birman–Schwinger criterion for trapped-mode existence as the central open problem for the next paper. The present §§7, 8, 11 supply this criterion, its associated localization estimates, and explicit worked examples. The Stage VIII §12.2 problem is closed for the toy-form **T** of Definition 2.1; the more general case ($C \neq 0$, non-uniform γ, irregular substrates) follows the same operator-theoretic template with technical refinements at each step.

The remaining Stage VIII open problems (12.1 tensorial curvature, 12.3 explicit strong-defect construction with verified global trapped mode, 12.4 substrate-level universality class, 12.5 matter interpretation, 12.6–12.10) are not addressed by the present paper. In particular, the present transport machinery makes Stage VIII §12.3 (explicit construction) tractable in principle — one constructs a defect with α large enough that the Birman–Schwinger equation acquires a root in the gap region — but does not produce an explicit such defect within the admissible range of any specific defect family. The §11.2 analysis shows that the Stage VIII §9.1 defect at $\alpha = 0.20$ *is* such an example for sufficiently weak γ; for stronger γ, an engineered defect (Stage VIII §7.5 structural note) is required.

---

12. Defect Hybridization and Composite Trapped Modes

The Birman–Schwinger framework also resolves the Stage VIII §10.1 hedge about trapped-mode hybridization across non-overlapping defects.

12.1 Two-Defect Birman–Schwinger

For two defects at positions x_1, x_2 with non-overlapping support, the combined perturbation $V = V_1 + V_2$ has rank at most $\mathrm{rank}(V_1) + \mathrm{rank}(V_2)$. The Birman–Schwinger operator $-V \cdot R_0(\lambda)$ decomposes block-wise:

$$-V \cdot R_0(\lambda) = \begin{pmatrix} -V_1 \cdot R_0(\lambda) \cdot P_1 & -V_1 \cdot R_0(\lambda) \cdot P_2 \\ -V_2 \cdot R_0(\lambda) \cdot P_1 & -V_2 \cdot R_0(\lambda) \cdot P_2 \end{pmatrix}$$

where P_i is the projection onto Range(V_i). The off-diagonal blocks contain the *hybridization terms*: V_i·R_0(λ)·V_j with i ≠ j is non-zero whenever the bulk resolvent connects x_i to x_j.

By Theorem 8.1, the off-diagonal blocks have operator norm bounded by (const)·e^{−η·d_X(x_1, x_2)}, decaying exponentially in the separation distance. For widely separated defects ($d_X(x_1, x_2) \gg 1/\eta$), the off-diagonal blocks are negligible and the trapped-mode spectrum is approximately the union of the individual defects' trapped-mode spectra.

12.2 Explicit Splitting Formula and Bonding/Antibonding Trapped Modes

For two identical rank-1 defects $\alpha \cdot uv^T$ placed at substrate positions x_1, x_2 with separation $d := d_X(x_1, x_2)$, the Birman–Schwinger eigenvalue equation reduces to a 2×2 block-matrix problem on $\mathrm{Range}(V_1) \oplus \mathrm{Range}(V_2)$. Define the **diagonal resolvent matrix element** at a single defect

$$G_0(\lambda) := \alpha \cdot \langle v, R_0(\lambda) \cdot u \rangle_{x_i, x_i}, \text{ (independent of } i \text{ for identical defects),}$$

and the **off-diagonal hybridization matrix element**

$$G_d(\lambda) := \alpha \cdot \langle v, R_0(\lambda) \cdot u \rangle_{x_1, x_2}.$$

By Theorem 8.1, the off-diagonal element decays exponentially in d:

$$|G_d(\lambda)| \le (2\alpha/\delta) \cdot \|u\| \cdot \|v\| \cdot e^{-\eta \cdot d}, \text{ where } \eta = \min(1, \delta/(2\gamma \cdot \rho(A_X))).$$

The Birman–Schwinger condition $1 \in \mathrm{spec}(-V \cdot R_0(\lambda))$ becomes the 2×2 eigenvalue equation

$$\det \begin{pmatrix} G_0(\lambda) + 1 & G_d(\lambda) \\ G_d(\lambda) & G_0(\lambda) + 1 \end{pmatrix} = 0$$

which factorises as $(G_0(\lambda) + 1 + G_d(\lambda)) \cdot (G_0(\lambda) + 1 - G_d(\lambda)) = 0$.

The two solutions are:

Bonding (symmetric) trapped mode: $G_0(\lambda_+) = -1 - G_d(\lambda_+)$, **Antibonding (antisymmetric) trapped mode:** $G_0(\lambda_-) = -1 + G_d(\lambda_-)$.

Sign-convention note. In chemistry, the bonding (symmetric) molecular orbital is energetically *lower* than the antibonding one. Here the bonding mode appears at the *higher* of the two split eigenvalues, $\lambda_+ > \lambda_-$. This is a sign-convention artefact of the Birman–Schwinger criterion with $K(\lambda) = -V \cdot R_0(\lambda)$ (the negative sign in the definition flips the "energy" axis relative to the standard MO convention) combined with the rank-1 coupling sign; the structural correspondence with MO theory (symmetric vs antisymmetric superpositions of localized states with exponentially-small splitting) is preserved, only the up/down labelling of "bonding" relative to the unperturbed level is reversed. A reader from chemistry should expect to relabel.

Near the single-defect trapped eigenvalue λ_0 (where $G_0(\lambda_0) = -1$), Taylor-expanding G_0 around λ_0 gives $G_0(\lambda) \approx -1 + G_0'(\lambda_0) \cdot (\lambda - \lambda_0)$, and the two solutions sit symmetrically around λ_0 with splitting

$$\lambda_+ - \lambda_- \approx 2 \cdot G_d(\lambda_0) / G_0'(\lambda_0).$$

Substituting the Combes–Thomas bound on G_d:

Proposition 12.1 — Exponential Splitting of Two-Defect Trapped Modes

Under the same reversibility hypothesis as Theorem 8.1 (spatial fibres preserve reversibility, regime (W), operations in $L^2(\pi \otimes \mu_X)$), for two identical rank-1 defects at substrate positions separated by $d = d_X(x_1, x_2)$, with single-defect trapped eigenvalue λ_0 in the spectral gap $\mathcal{G}_{trap}$ at spectral distance δ from the bulk bands, the two-defect trapped eigenvalues $\lambda_\pm$ satisfy

$$|\lambda_+ - \lambda_-| \le C \cdot e^{-\eta \cdot d},$$

with $\eta = \min(1, \delta/(2\gamma \cdot \rho(A_X)))$ the Combes–Thomas decay rate and C a constant depending on δ, the defect strength α, and the local resolvent derivative at λ_0. Specifically,

$$|\lambda_+ - \lambda_-| = (4\alpha / (\delta \cdot |G_0'(\lambda_0)|)) \cdot \|u\| \cdot \|v\| \cdot e^{-\eta \cdot d} + O(e^{-2\eta \cdot d}).$$

The non-reversible case carries an additional $\kappa(\boldsymbol{T}_{bulk})$ factor in the constant C, inheriting the §8.4 pseudospectral correction.

Proof. Combine the Taylor expansion above with the Theorem 8.1 bound $|G_d(\lambda_0)| \le (2\alpha/\delta) \cdot \|u\| \cdot \|v\| \cdot e^{-\eta \cdot d}$. The factor of 2 in $|\lambda_+ - \lambda_-| = 2|G_d/G_0'|$ absorbs into the constant.

Interpretation as defect interaction energy. The quantity $J(d) := \tfrac{1}{2}|\lambda_+ - \lambda_-|$ has the structural role of an **interaction energy** between the two defects:

$$J(d) \approx (2\alpha / (\delta \cdot |G_0'(\lambda_0)|)) \cdot \|u\| \cdot \|v\| \cdot e^{-\eta \cdot d}.$$

The interaction is exponentially decaying in separation, with characteristic range $1/\eta = \xi$ (the localization length of the single-defect trapped mode). This is structurally identical to the way bound states in attractive potentials interact in quantum mechanics: the splitting is the off-diagonal matrix element of the perturbation in the bound-state basis, and the off-diagonal element is the overlap of two exponentially-decaying eigenfunctions.

12.3 Molecular-Orbital-Theory Analogy

The bonding/antibonding splitting of Proposition 12.1 is the substrate-level analogue of the formation of molecular orbitals from atomic orbitals in chemistry. The structural correspondences:

Two-defect substrate	Diatomic molecule
Single-defect trapped mode at x_i with eigenvalue λ_0	Atomic orbital on atom i with energy E_0
Two-defect bonding mode $\lambda_+ = \lambda_0 + J(d)$	Bonding molecular orbital $E_+ = E_0 + t$
Two-defect antibonding mode $\lambda_- = \lambda_0 - J(d)$	Antibonding molecular orbital $E_- = E_0 - t$
Defect separation d	Bond length R
Localization length $\xi = 1/\eta$	Atomic-orbital extent
Exponential coupling $J(d) \propto e^{-\eta \cdot d}$	Hopping integral $t \propto e^{-R/\xi_{atom}}$

The analogy is structural rather than physical: the substrate excitations are not electronic orbitals, and there is no spin / Pauli structure here. But the operator-theoretic mechanism — finite-rank perturbations of a bulk operator producing discrete eigenvalues that split via off-diagonal matrix elements decaying exponentially in separation — is exactly the same. The substrate provides a setting in which the "molecular" structure of trapped substrate excitations emerges from purely spectral / resolvent considerations.

12.4 N-Defect Generalisation: Impurity-Band Formation

For N defects at well-separated positions, the Birman–Schwinger operator on Range(V) = $\bigoplus_i$ Range(V_i) is an N×N matrix (for rank-1 defects) of the form

$M_{ij}(\lambda) = G_{d_{ij}}(\lambda)$, where $d_{ij} := d_X(x_i, x_j)$,

with $M_{ii}(\lambda) = G_0(\lambda)$ on the diagonal. The Birman–Schwinger condition $1 \in \mathrm{spec}(-M(\lambda))$ becomes the N-fold splitting of the single-defect trapped eigenvalue into a *cluster* of N nearby eigenvalues, with bandwidth set by the largest off-diagonal element $\max_{i\neq j} |G_{d_{ij}}(\lambda_0)| \leq (2\alpha/\delta)\|u\|\|v\| \cdot e^{-\eta \cdot d_{min}}$, where $d_{min} := \min_{i\neq j} d_X(x_i, x_j)$.

For large N with regularly-spaced defects (the substrate-defect "superlattice" regime), the N discrete eigenvalues approach a continuous **impurity band** of width

$W_{impurity} \approx 2 \cdot z_{defect} \cdot (2\alpha/\delta) \cdot \|u\|\cdot\|v\| \cdot e^{-\eta \cdot d_{min}}$,

with z_{defect} the coordination number of the defect superlattice. This is the substrate-level analogue of impurity-band formation in heavily-doped semiconductors and donor-impurity clusters in solid-state physics, derived here from first principles via the Birman–Schwinger framework.

12.5 Composite Trapped Modes

In the regime of strong hybridization (multiple defects within one localization length), the individual trapped modes merge into a *composite trapped mode* extending across the defect cluster. The eigenvector decays exponentially outside the cluster but is delocalised within it.

The structural significance: trapped substrate excitations are *interacting* objects in the present framework, with interactions mediated by overlap of their exponentially-decaying tails. The interaction range is set by the localization length $1/\eta$, which in turn is set by the spectral distance δ from the bulk bands. Close to the band edge ($\delta \to 0$, $\xi \to \infty$), trapped modes are long-ranged and strongly interacting; far from the band edge (large δ), they are short-ranged and weakly interacting. The Proposition 12.1 interaction energy $J(d) \propto e^{-\eta\cdot d}$ provides an explicit quantitative form of this distance-dependence.

This is one of the most physically suggestive structural features the present transport machinery produces — though, as always, the "matter-like" reading remains conjectural in the strict Stage VIII sense. No specific particle content, statistics, or interaction phenomenology is claimed.

What *is* established is that the operator-theoretic interaction structure of trapped substrate excitations is mathematically isomorphic to the structure governing molecular orbital formation and impurity-band physics, with explicit splitting formulas derivable from Birman–Schwinger plus Combes–Thomas.

---

13. The Coherence Transport Scaling Identity

The various scales developed across §§3–8 — coherence velocity v_c, localization length ξ, spectral distance δ, local gap field ε_{gap} — are not independent. They are linked by a single underlying scaling relation that unifies propagation, localization, and coherence stability into one structural framework.

13.1 The Three Scales

Three scales control the trapped-mode physics:

- **Coherence velocity** $v_c := \gamma \cdot \rho(A_X)$, the rate of coherence propagation through the bulk substrate (Theorem 4.1).
- **Spectral distance from bulk bands** $\delta := \text{dist}(\lambda, \text{spec\_ess}(\mathbf{T}_{bulk}))$, the "binding strength" of a trapped eigenvalue (the analogue of the binding energy of a bound state in quantum mechanics).
- **Localization length** $\xi := 1/\eta$, where $\eta = \min(1, \delta/(2\gamma \cdot \rho(A_X)))$ is the Combes–Thomas decay rate (Theorem 8.1).

13.2 The Basic Scaling Identity

The Combes–Thomas formula $\eta = \delta/(2\gamma \cdot \rho(A_X)) = \delta/(2v_c)$ (in the small-η regime $\delta < 2v_c$, which is the regime where localization is non-trivial) gives directly the **basic localization scaling**:

$$\xi = 1/\eta = 2v_c/\delta.$$

Equivalently:

$$\xi \cdot \delta = 2 \cdot v_c.$$

This is the substrate-level analogue of the **uncertainty-style bound** of quantum mechanics: localization length × spectral binding ∼ velocity scale. Trapped modes that are tightly bound (large δ) are tightly localized (small ξ), and trapped modes near the band edge (small δ) are widely extended (large ξ), with the product fixed at $2v_c$.

13.3 The Gap-Field-Dependent Scaling (Heuristic)

Granting the heuristic identification $v_c^{local}(x) \propto \varepsilon_{gap}(x)$ of §9.1 (whose epistemic register is recorded there), the basic identity combines to give a refined scaling for spatially-varying-gap substrates:

$\xi_{local}(x) \sim v_c^{local}(x) / \delta \propto \varepsilon_{gap}(x) / \delta$. (HEURISTIC — inherits §9.1's status)

This is the **gap-field-dependent localization scaling**: in regions of suppressed ε_{gap}, the local coherence velocity is heuristically reduced, hence the *local* localization length is reduced — coherence becomes more strongly confined in low-coherence regions. This is dimensionally consistent with the §9 coherence-well interpretation (depressed gap = reduced local transport = enhanced local confinement), with the same epistemic status as the underlying §9.1 identification.

13.4 Unified Coherence Transport Scaling Identity

Putting these together with the entropy-retention rate of Stage VIII Corollary 8.1 (where the local $L^2(\pi)$ contraction rate is $1 - \varepsilon_{gap}(x)$, giving relaxation time $\tau_{relax}(x) \sim 1/\varepsilon_{gap}(x)$) and the candidate-curvature-gradient relation $R(x) = \nabla^2 \varepsilon_{gap}(x)$:

Theorem 13.1 — Coherence Transport Scaling Identity

Across the $K = 7$ refinement substrate with global transport operator $\boldsymbol{T}$ of Definition 2.1, the following four scales are linked by a single underlying scaling identity:

Local coherence velocity (heuristic, §9.1): $v_c^{local}(x) \propto \gamma \cdot \rho(A_X) \cdot \varepsilon_{gap}(x)$ **Local relaxation time** (rigorous, Stage VIII Cor. 8.1): $\tau_{relax}(x) \propto 1/\varepsilon_{gap}(x)$ **Local trapped-mode localization length** (heuristic in its position dependence): $\xi(x; \lambda) \propto v_c^{local}(x) / \delta(\lambda) \propto \varepsilon_{gap}(x) / \delta(\lambda)$ **Local roughness functional** (rigorous, Stage VIII Prop. 4.1): $K_\infty(x) \propto 1/\varepsilon_{gap}(x)$

All four quantities are functionals of the single underlying field $\varepsilon_{gap}(x)$, with explicit proportionality constants determined by γ, $\rho(A_X)$, and the spectral distance $\delta(\lambda)$ of the trapped eigenvalue from the bulk bands.

Equivalently, defining the ***dimensionless coherence-transport scaling number***

$\mathcal{N}_{coh}(x; \lambda) := \xi(x; \lambda) \cdot \delta(\lambda) / v_c^{local}(x)$,

we have $\mathcal{N}_{coh} \sim 1$ throughout the substrate, with deviations only in the strong-coupling regime or at band edges.

Proof sketch. Each scaling follows from a previously established result: $v_c^{local}(x)$ from §9.1 (heuristic identification of fibre relaxation rate with effective transport velocity, marked there explicitly); $\tau_{relax}(x)$ from Stage VIII Corollary 8.1 rigorously (the $L^2(\pi_x)$ contraction rate is $1 - \varepsilon_{gap}(x)$, giving relaxation time $\sim 1/\varepsilon_{gap}(x)$ under iterated refinement); $\xi(x; \lambda)$ from §13.2's basic identity $\xi \cdot \delta = 2v_c$ combined with the heuristic local v_c; $K_\infty(x)$ rigorously from Stage

VIII Proposition 4.1 (continuum Lipschitz constant). The dimensionless $\mathcal{N}_{coh} \sim 1$ follows directly from the localization formula $\xi \cdot \delta \sim v_c$, which is rigorous at the bulk level and inherits the heuristic register at the gap-field-dependent local level.

Naming note. This is called a *scaling identity* rather than a *scaling law* — the result is the recognition that previously-derived scales collapse onto one relation, not a new derivation of a fundamental constant. The substantive content is the structural unification of §13.5; the parameter-free identity $\xi \cdot \delta \sim v_c$ is the operator-theoretic backbone of Combes–Thomas (§8), re-stated in a form that explicitly relates it to the other transport scales of the framework.

13.5 Structural Reading: One Field, Four Functionals, One Scaling Identity

Theorem 13.1 unifies Stage VIII's four ε_{gap}-functionals with the new Stage IX scales (coherence velocity, localization length) into a single scaling framework. The structural picture is:

- Stage VIII identified that one field (ε_{gap}) controls four functionals (continuum roughening, candidate curvature, entropy retention, trapped-mode persistence).
- Stage IX shows that all of these are governed by one scaling identity $\xi \cdot \delta \sim v_c \propto \varepsilon_{gap}$, with each Stage VIII functional being a specific extraction from the same scaling structure:
    - Continuum roughening $K_\infty \propto 1/\varepsilon_{gap}$ is τ_{relax}-type (inverse scale).
    - Candidate curvature $R = \nabla^2 \varepsilon_{gap}$ is the second derivative of the scale itself.
    - Entropy retention $(1 - \varepsilon_{gap})^{2n}$ is the relaxation factor.
    - Trapped-mode persistence $\xi \propto \varepsilon_{gap} / \delta$ is the localization scale (and Birman–Schwinger existence is the condition that a trapped eigenvalue at any $\delta > 0$ exists).

This is the structural climax of the present paper: the Stage VIII Defect-Coherence Principle's "four functionals of one field" gets upgraded to "four functionals of one field under one scaling identity", with the scaling identity $\xi \cdot \delta \sim v_c$ being the operator-theoretic substrate that the Stage VIII functional structure inherits.

13.6 Status of the Scaling Identity

The Coherence Transport Scaling Identity is a **structural unification**, not a new derivation: each scaling individually was already derivable from §§3–9 + Stage VIII, and Theorem 13.1 simply records that they collapse onto one identity. The contribution is the *recognition* that the various scales of the present paper are not independent free parameters but are bound by one underlying relation.

What this does not establish:

- It does not derive a dimensionful "fundamental constant" of the substrate. The scales are linked by their structural form, not by any specific numerical value.
- It does not impose any quantum-mechanical reading on the scaling. The substrate-level scaling $\xi \cdot \delta \sim v_c$ is structurally identical to the quantum-mechanical bound-state

scaling $\xi \cdot E_{bind} \sim \hbar v_F$, but the substrate is classical / operator-theoretic and the identification is structural-mathematical only.
- It does not predict the existence of trapped modes — it predicts their *scaling* once they exist, with existence still governed by the Birman–Schwinger criterion of §7.

The scaling identity is best read as the *parametric backbone* of the entire trapped-mode and transport-geometry programme: any specific substrate, defect, and trapped eigenvalue must satisfy the scaling, and any inconsistency with it would signal an error in one of the underlying derivations.

---

14. Limitations and Open Problems

The present paper closes the Stage VIII §12.2 Birman–Schwinger criterion and the Stage VIII §6.2 Remark's trapped-mode existence question, while leaving the remaining Stage VIII open problems and introducing some new ones specific to the present framework.

14.1 Coupling regime restriction. The principal results assume the weak-coupling regime (W): $\gamma \cdot \rho(A_X) < ½$, or the stronger (W'): $\gamma \cdot \rho(A_X) < 3/56$ for full band separation. Outside (W), bands may overlap and the spectral picture becomes substantially more delicate: bands may merge, the bulk subdominant modulus may exceed 1, and the Perron eigenvalue may cease to be isolated. The strong-coupling regime is the analogue of Stage VII / Stage VIII's strong-perturbation regime and is similarly outside the present analysis's scope.

14.2 Stochasticity loss in the bulk operator. The operator **T** of Definition 2.1 is not row-stochastic: the additive coupling $\gamma \cdot I_{\mathcal{K}} \otimes A_X$ pushes the spectral radius above 1 (specifically to $1 + \gamma \cdot \rho(A_X)$). A stochasticity-preserving variant would require either (a) a normalised Laplacian-style coupling $\mathbf{T} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes (I_X - D_X^{-1} A_X)$ — but this gives a *negative* coupling whose spectral analysis differs — or (b) a careful balance between the $\hat{T} \otimes I_X$ and A_X terms preserving row sums to 1. We have not pursued the stochasticity-preserving variant in this paper; the present **T** is the simplest tractable form for explicit Birman–Schwinger analysis, at the cost of leaving the strict-stochasticity interpretation of Stage VI–VIII partially behind. Restoring stochasticity is an open structural problem.

14.3 Choice of substrate-level universality class. Stage VIII §12.4 flagged the need for a substrate-level extension of the Stage VII universality class $\mathcal{C}_{K=7}$, covering position-indexed operator families rather than single operators. The present paper provides the global operator **T** in which such a class would naturally live but does not formalise the class itself; openness, robustness, and the substrate-level analogues of Stage VII Theorem 12.1 remain to be developed.

14.4 Non-regular substrates. The present analysis assumes a regular substrate (bounded degree, symmetric coupling, translation-invariant in the infinite case). Real Stage IV substrate structures may be irregular (varying degree, anisotropic coupling, non-translation-invariant). The Birman–Schwinger framework extends to such cases in principle but the explicit resolvent calculations become substantially harder.

14.5 The C ≠ 0 closure-mixing coupling. Definition 2.1 includes a closure-label-mixing term C that we have largely set to zero in the explicit calculations. The qualitative structure (band structure, Birman–Schwinger criterion, Combes–Thomas decay) extends to C ≠ 0 in principle, but the explicit form of the bulk resolvent and the Birman–Schwinger operator becomes more involved. A systematic treatment of C ≠ 0 is deferred.

14.6 Tensorial curvature theory. This was Stage VIII §12.1's principal open problem and remains so. The present transport machinery provides additional substrate-level structure (transport bands, coherence velocity, transport metric, coherence-ray bending of Proposition 9.3) that any tensorial geometric theory must accommodate, but it does not derive the tensorial structure itself.

14.7 Matter interpretation. Stage VIII §12.5's matter-interpretation question is not addressed. Trapped modes here, as in Stage VIII, have matter-like *features* (localized, persistent, distinguishable from the vacuum, with finite interaction range set by 1/η and explicit Proposition 12.1 splitting) but no specific particle content. The composite-trapped-mode hybridization of §12 is suggestive of multi-particle bound states but is not derivational.

14.8 Quantization. Stage VIII §12.8's quantization question remains open. The present framework is classical / operator-theoretic. The trapped-mode spectrum is *discrete* (when bulk bands are separated and the Birman–Schwinger equation has finitely many roots), which is suggestive of a quantum-mechanical-style discrete spectrum, but the relevant quantization (canonical, geometric, etc.) is not derived. The Theorem 13.1 scaling identity $\xi \cdot \delta \sim v_c$ is structurally identical to the quantum-mechanical bound-state scaling $\xi \cdot E_{bind} \sim \hbar v_F$, which is suggestive but not a quantization in itself.

14.9 Cosmological structure. Whether the global transport structure connects to substrate-level cosmological dynamics (TPB saturation, entropy loading, Friedmann-like equations from Stage VIII's conjectural energy-density correspondence) is open.

14.10 Beyond the toy form. Definition 2.1 is a *toy form* in the sense that the Stage IV substrate may dictate a more general structure (closure-mixing C, position-dependent γ, non-adjacency-based coupling). The §§7–8 framework extends, but explicit calculations become harder. A future paper would systematically extend the present analysis to the actual Stage IV structure once that is fixed.

---

15. Conclusion

Programme map. Stage V established that the canonical K = 7 wheel produces a Lipschitz continuum limit at large scales, with continuum regularity controlled by the spectral gap $\varepsilon_{gap} = ½$. Stage VII established that this coherent substrate sits inside an open universality class, robust to admissibility-preserving perturbations. Stage VIII studied *localized* perturbations and identified four ε_{gap}-functionals (continuum roughening, candidate curvature, entropy retention, trapped-mode persistence) that share a common dependence on the local spectral gap field

ε_gap(x), but left the trapped-mode side conditional on a global eigenvector existence claim that required substrate-level machinery beyond local analysis. The present paper supplies this machinery: an explicit global refinement transport operator **T** with band structure, finite coherence propagation, and the spectral / resolvent tools for analyzing localized perturbations.

The principal results:

- **Global operator with band structure.** The coupled operator $\mathbf{T} = \hat{T} \otimes I_X + \gamma \cdot I_{\mathcal{K}} \otimes A_X$ has bulk spectrum given by the convolution $\mathrm{spec}(\hat{T}) + \gamma \cdot \mathrm{spec}(A_X)$, organised into at most 7 transport bands centred at the wheel eigenvalues with widths $2\gamma \cdot \rho(A_X)$. The canonical wheel gap $\varepsilon_{gap} = ½$ persists as the global-operator gap $\varepsilon_{gap}(\mathbf{T}_{bulk})$ in the weak-coupling regime (W).
- **Finite coherence propagation.** A Lieb–Robinson-type bound (Theorem 4.1) gives an explicit emergent finite propagation speed $v_c = \gamma \cdot \rho(A_X)$, with super-exponentially decaying correlations outside the substrate light-cone (in the minimal-coupling regime C = 0; a non-zero closure-mixing C modifies v_c by an additive $\|C\|$-bounded contribution).
- **Birman–Schwinger trapped-mode criterion.** For a localized defect V, λ outside the bulk essential spectrum is an eigenvalue of $\mathbf{T}_{bulk} + V$ iff $1 \in \mathrm{spec}(-V \cdot R_0(\lambda))$. This reduces the global eigenvalue problem to a finite-dimensional eigenvalue problem on Range(V), making trapped-mode existence explicitly checkable.
- **Combes–Thomas exponential localization.** Trapped eigenvalues at spectral distance δ from the bulk bands have eigenvectors decaying at rate $\eta = \min(1, \delta/(2\gamma \cdot \rho(A_X)))$, giving localization length $\xi = 1/\eta$.
- **Stage VIII trapped-mode theorem becomes unconditional.** Stage VIII Theorem 6.1's hypothesis is now verified by Birman–Schwinger (eigenvalue) + Combes–Thomas (eigenvector). The §11.2 worked example shows that the Stage VIII §9.1 boundary defect produces a global trapped mode in the weak-transport-coupling regime ($\gamma \cdot \rho(A_X) < 0.05$ for $\alpha = 0.20$), and the §11.3 example shows that the Stage VIII §9.6 spectrally-invisible hub-coupling defect remains globally invisible under the transport lift.
- **Coherence Transport Scaling Identity (Theorem 13.1).** The various scales of the framework collapse onto one identity: $\xi \cdot \delta \sim v_c \propto \varepsilon_{gap}$, with the Stage VIII four ε_{gap}-functionals being specific extractions from this single scaling relation at different mathematical orders (zeroth derivative, first/second derivative, asymptotic decay rate, localization-length ratio). The Stage VIII "one field, four functionals" structure is upgraded to "one field, four functionals, one scaling identity".
- **Explicit substrate examples (§3.4).** Worked computations for $\mathbb{Z}$, $\mathbb{Z}^2$, hexagonal lattice, and finite torus $\mathbb{T}_N^d$ give the band structures, (W) and (W') regimes, and trap-gap intervals concretely, making the framework applicable to specific Stage IV substrate choices.
- **Coherence-ray bending (Proposition 9.3).** The candidate-curvature indicator $R(x) = \nabla^2 \varepsilon_{gap}(x)$ acquires a transport-dynamical interpretation: coherence trajectories bend toward suppressed-gap regions at a rate proportional to $\nabla \varepsilon_{gap}/\varepsilon_{gap}^2$, giving an operator-theoretic content to the Stage VIII curvature ansatz without invoking general relativity.
- **Fibrewise projection theorem (Theorem 10.1).** Stage VIII is identified explicitly as the *fibrewise projection theory* of the present Stage IX transport dynamics, with $\mathbf{T}|_{\mathcal{K}_x} =$

$\hat{T}_x$ making the local Stage VIII analysis a controlled projection of the global Stage IX framework.

The defect-coherence transport reading (§9) unifies the four Stage VIII ε_{gap}-functionals under a common interpretive umbrella — *local coherence-transport capacity* — which is suggestive but not derivational: the Stage VIII functional unification was already structural, and the present transport reading adds a physical-interpretive layer without changing the underlying mathematics.

The Stage V → VII → VIII → IX trajectory therefore moves from *establishing coherence* (V) to *robustness of coherence* (VII) to *characterising local breakdowns of coherence* (VIII) to *engineering global transport of coherence* (IX), with each stage contained operator-theoretically in the previous.

The honest summary: this paper closes the Stage VIII §12.2 open problem, supplies the Stage VIII §12.2 (a)–(e) engineering pieces in their stated dependency order, and makes the Stage VIII trapped-mode theorem unconditional for the toy form of **T**. It does not address Stage VIII's tensorial-curvature (12.1), substrate-level universality (12.4), matter-interpretation (12.5), or quantization (12.8) open problems. These remain the principal targets for subsequent papers.

15.1 Transport-Geometric Reading of the Defect-Coherence Principle

The Stage VIII Defect-Coherence Principle (§13.1) listed four ε_{gap}-functionals as different mathematical extractions from the same underlying spectral-gap field. The present transport machinery provides a physical-interpretive reading: all four functionals describe local *coherence-transport capacity*, with continuum roughening ($K_\infty \propto 1/\varepsilon_{gap}$) reflecting the local relaxation timescale, candidate curvature ($R = \nabla^2\varepsilon_{gap}$) the spatial gradient of transport capacity, entropy retention ($(1 - \varepsilon_{gap})^{2n}$) the local outflow rate, and trapped-mode persistence (Birman–Schwinger + Combes–Thomas) the boundary at which local transport becomes fully obstructive and discrete eigenstates form.

This is a *physical interpretation*, not a new derivation. The four functionals share their ε_{gap} dependence by construction (Stage VIII §13.1's functional unification); the present reading explains *why* this shared dependence is physically natural — they are all readings of one underlying transport quantity at different mathematical orders (zeroth, first/second derivative, asymptotic decay rate, spectral location relative to bulk bands). Theorem 13.1's Coherence Transport Scaling Identity collapses these into a single dimensionless identity $\xi \cdot \delta \sim v_c$, making the underlying unity quantitative.

15.2 What This Establishes and What Remains

What this paper establishes: the explicit operator-theoretic machinery for studying refinement-substrate coherence transport. The bulk spectrum is computed in closed form (Proposition 3.1); the finite propagation bound is established (Theorem 4.1); the Birman–Schwinger trapped-mode criterion is given (Theorem 7.1); the Combes–Thomas localization estimate is given (Theorem 8.1); the Stage VIII Theorem 6.1 trapped-mode hypothesis becomes verifiable rather than

presumed; and the Stage VIII §12.2 engineering-pieces list is closed in its specified dependency order.

What remains: the tensorial curvature theory (Stage VIII §12.1), an explicit construction of a strong-defect example with verified global trapped mode in a Stage IV-realistic substrate (Stage VIII §12.3 with the present machinery making it tractable), the substrate-level universality class (Stage VIII §12.4), and the matter / quantization / cosmological-structure questions (Stage VIII §§12.5, 12.8, 12.9).

The picture across Stages V–IX is now:

- Coherence (Stage V): the canonical wheel produces a Lipschitz continuum.
- Robust coherence (Stage VII): the wheel sits inside an open universality class.
- Local breakdown (Stage VIII): localized defects produce a spatially varying ε_{gap} field with four shared functionals.
- Global transport (Stage IX, present): the substrate is a coupled dynamical system with band structure, finite propagation, and verifiable trapped modes.

The Defect-Coherence Principle (Stage VIII §13.1) and the transport-geometric reading (§15.1) together form the structural climax of this four-stage arc: a single underlying field (the local spectral gap) controls four observable phenomena (roughening, curvature, retention, trapping), and these phenomena have a common transport-theoretic interpretation that an eventual tensorial theory must reproduce, with the Coherence Transport Scaling Identity of Theorem 13.1 collapsing the various scales onto a single identity $\xi \cdot \delta \sim v_c$. The present paper closes the operator-theoretic side of this picture; the tensorial / matter / quantization sides remain the programme's next-stage targets.

The honest position: this paper makes the trapped-mode side of the Stage VIII Defect-Coherence Principle unconditional rather than conjectural, and supplies the global-operator infrastructure that the rest of the geometry programme will build on. It does not derive new physics. What it does is engineer the operator-theoretic bridge between local substrate analysis and large-scale substrate-derived structure — the bridge Stage VIII identified as missing, with the explicit dependency-ordered roadmap (a) → (b) → (c, d in parallel) → (e) now executed in full.