Stability and Universality of the K = 7 Refinement Fixed Point

Perturbed Admissibility Rules, Spectral Stability, and the Universality Class of Emergent Lorentzian Geometry in VERSF

Keith Taylor — VERSF Theoretical Physics Programme

General-Reader Summary

The previous paper constructed the canonical K = 7 wheel substrate and computed every quantity Stage V left abstract: the spectrum of the refinement operator turned out to be exactly {1, ½, ½, −¼, −3/28, 0, 0}, with the eigenvalue 1 simple, a clean spectral gap of ε_gap = ½, geometric contraction of nearby distributions at rate $(½)^n$, and a Lipschitz constant for the emergent continuum given as $K_\infty = 2 \cdot L_\Phi \cdot L \cdot A^2 / A_-$.

The natural next question is the one a physicist asks of any structural result: *is this fragile?*

If the conclusion only held when every transition weight and every admissibility rule were perfectly tuned, the construction would be a beautiful toy model — and only that. Real microscopic substrates are never perfectly exact: they carry local variation, defects, anisotropy, disorder, and small perturbations of every kind. A foundational substrate for physics has to tolerate them.

This paper proves that the K = 7 refinement fixed point does tolerate them. Small admissibility-preserving perturbations of the canonical operator leave the architecture intact: irreducibility, aperiodicity, the simple Perron eigenvalue 1, the positive spectral gap, exponential refinement contraction, entropy contraction, and the Lipschitz continuum constant all survive — and each varies *continuously* with the perturbation strength, not catastrophically.

Two findings sharpen the picture beyond simple stability.

First, the wheel architecture is so symmetric that an entire class of perturbations is *spectrally invisible*: certain antisymmetric perturbations under the wheel's dihedral symmetry leave the spectrum exactly unchanged — not "approximately", not "to leading order", but *exactly*, to every digit, for every value of the perturbation strength η. The eigenvalues do not shift; only the eigenvectors do. The result is verified both analytically (via the matrix determinant lemma combined with reflection symmetry) and numerically (characteristic polynomial identical to 12 decimal places across the admissible perturbation range). This promotes the Stage VI spectral gap ε_gap = ½ from a *computed* number to a **symmetry-protected** one — a structural invariant of the wheel's dihedral D_6 symmetry, not a numerical accident.

Second, for perturbations that *do* move the spectrum, the actual sensitivity is much smaller than the worst-case theoretical bound. For one representative symmetry-breaking perturbation, the spectral gap responds with effective Lipschitz constant ≈ 0.23, compared to the conservative Bauer–Fike upper bound $\sqrt{2} \approx 1.41$ — a factor-of-six safety margin built into the canonical architecture.

The general lesson is the one expected of any physically credible structural result: emergent Lorentzian geometry in VERSF is not the fine-tuned property of one exact matrix. It is the stable infrared behaviour of an *open neighbourhood* of admissible refinement substrates around the canonical wheel — the first explicit *universality class* in the VERSF geometry programme.

What remains open is whether this neighbourhood extends to *large* perturbations, whether a sharper uniform-gap bound exists for level- and position-dependent filter families, and whether the wheel itself can be derived from deeper principles rather than postulated as the canonical substrate. The first two are concrete computations; the third is a structural question for a later paper.

Abstract

The previous Stage VI paper constructed the canonical K = 7 wheel substrate W_6, defined an explicit transition operator $\hat{T}$, computed its spectrum exactly as $\{1, ½, ½, −¼, −3/28, 0, 0\}$, and showed that $\varepsilon_{gap} = ½$ produces exponential refinement contraction and Lipschitz continuum regularity.

The present paper proves stability of this construction under admissible perturbations.

We define a three-condition admissibility class P1–P3 for perturbed transition operators $\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T$ (stochasticity preservation, support preservation, positivity), with three further structural consequences P4–P6 (hub–boundary communication, boundary cyclic communication, self-loop persistence) automatic from P3 for the canonical wheel and retained as named labels for proof referencing. The principal results are:

- **Stability of irreducibility and aperiodicity** (Propositions 4.1, 4.2). Open under P4 + P6 for all sufficiently small $|\eta|$.
- **Spectral gap stability** (Theorem 5.1). For admissible perturbations, $|\lambda_k(\hat{T}_\eta) - \lambda_k(\hat{T})| \leq C\cdot|\eta|$ with $C \leq \sqrt{(\pi_{max}/\pi_{min})} \cdot \|\Delta T\|_{op} \leq \sqrt{(7/4)} \cdot \|\Delta T\|_{op}$, via Bauer–Fike applied to $\hat{T}$'s reversibility-induced orthonormal eigenbasis. Consequently $\varepsilon_{gap}(\eta) \geq ½ - C\cdot|\eta|$, positive for $|\eta| < 1/(2C)$.
- **Stationary measure continuity** (Theorem 6.1). $\pi_\eta = \pi + \eta \cdot \pi \cdot \Delta T \cdot Z + O(\eta^2)$, with $Z = (I - \hat{T} + \mathbf{1}\pi)^{-1} - \mathbf{1}\pi$ the fundamental matrix of $\hat{T}$. Explicit and Lipschitz.
- **Entropic contraction stability** (Theorem 7.1). For reversibility-preserving perturbations, $\chi^2(\mu_n \| \pi_\eta) \leq (¼ + O(|\eta|))^n \cdot \chi^2(\mu_0 \| \pi_\eta)$.
- **Continuum Lipschitz stability** (Theorem 8.1). $K_\infty(\eta) = K_\infty(0) + O(\eta)$ wherever the Stage IV constants L_Φ, $L(\eta)$, A, A_- are themselves Lipschitz in η.

- **Symmetry-protected spectral invariance** (Theorem 9.1, Corollary 9.2). Perturbations transforming antisymmetrically under any reflection $\sigma \in D_6$ of the wheel's dihedral symmetry group — specifically perturbations $\Delta T = U V^T$ with U-columns in the σ-symmetric subspace H_σ^+ and V-columns in the σ-antisymmetric subspace H_σ^- — leave the spectrum of $\hat{T}_\eta$ **exactly** unchanged for *all* η, not merely to leading order. The eigenvectors rotate; the eigenvalues do not. Proof is via the matrix determinant lemma combined with resolvent commutation: $v^T (\hat{T} - \lambda I)^{-1} e_\kappa = 0$ forces the perturbative factor in the rank-r determinant identity to be exactly 1. This promotes $\varepsilon_{gap} = ½$ from a *computed* quantity to a *symmetry-protected* one — the spectral gap is a structural invariant under a positive-dimensional subspace of admissible perturbations.
- **Open stability and universality** (Theorem 12.1). There exists an open neighbourhood $\mathcal{U}$ of $\hat{T}$ in row-stochastic 7×7 matrix space such that every $S \in \mathcal{U}$ with wheel-admissible support satisfies the K = 7 refinement universality class definition: unique persistent coherent sector, exponential refinement contraction to a unique stationary distribution, geometric entropy contraction, finite continuum Lipschitz constant.

We further give a RG-theoretic classification of perturbations as *irrelevant* (gap-preserving), *marginal* (gap-shrinking but positive), or *relevant* (gap-closing); only relevant perturbations leave the universality class. Numerical worked examples confirm both spectral invariance (under symmetry-aligned perturbations) and the smooth O(η) shift of the gap (under generic symmetry-breaking perturbations), with measured Lipschitz constants substantially smaller than the analytic Bauer–Fike bound.

Emergent Lorentzian geometry in the K = 7 framework is therefore not a fine-tuned property of one exact matrix. It is the stable infrared behaviour of an open universality class of admissible refinement substrates around the canonical wheel.

Contents

1. Introduction

The Stage VI paper, *Explicit Construction and Spectral Analysis of the K = 7 Closure-Transition Graph*, computed every load-bearing quantity in the Stage V continuum-geometry chain to closed form for the canonical wheel architecture:

$K = 7$ wheel → $\beta_1(G\_\mathcal{K}) = 6$ → $\hat{T}$ irreducible + aperiodic → spec($\hat{T}$) exact → $\varepsilon\_gap = \frac{1}{2}$ → $\beta\_filter < 1$ → K_∞ finite.

That paper established the canonical wheel as a *working* substrate. But "working" is not the same as "robust". A construction in which every numerical value depends on every transition weight being exactly right is not a credible substrate for physics — real microscopic systems carry disorder, anisotropy, defects, and parameter drift. A foundational substrate must tolerate them.

The present Stage VII paper proves that the K = 7 refinement fixed point is robust in this sense. We define an admissibility class of perturbations (P1–P3 of §3, with P4–P6 as automatic consequences for the canonical wheel) and prove that every quantity in the Stage VI chain — irreducibility, aperiodicity, simplicity of the Perron eigenvalue, positivity of the spectral gap, exponential refinement contraction, geometric entropy contraction, finiteness of the continuum Lipschitz constant — is *stable* under admissible perturbations, in the sharp sense that each varies continuously (in fact Lipschitz) in the perturbation parameter η.

The framing is borrowed from renormalisation-group theory. We classify perturbations as *irrelevant* (the spectral gap remains $\mathcal{O}(1)$), *marginal* (the gap shrinks but stays positive), and *relevant* (the gap closes and Lipschitz geometry fails). Only relevant perturbations leave the universality class. The canonical wheel is then the simplest representative of an *open universality class* of admissible refinement substrates, all of which flow toward Lipschitz continuum geometry.

Two findings beyond simple stability deserve advance mention. First (Theorem 9.1), the wheel's dihedral D_6 symmetry forces an entire class of perturbations to be *spectrally silent* — their eigenvalues do not shift at all, only their eigenvectors. Second, numerical sharpness analysis (§10) shows that the actual Lipschitz constant of ε_gap with respect to η is substantially smaller than the worst-case Bauer–Fike bound — the canonical architecture is more robust than the abstract bound suggests.

2. The Canonical Fixed Point

For reference and notation, we recall the canonical Stage VI construction.

The closure catalogue is $\mathcal{K} = \{\kappa_h, \kappa_{b_1}, \ldots, \kappa_{b_6}\}$: one hub state and six boundary states. The canonical transition operator $\hat{T}$ is the 7 × 7 row-stochastic matrix

$$\hat{T} = \begin{pmatrix} 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 & 1/7 \\ 1/4 & 1/4 & 1/4 & 0 & 0 & 0 & 1/4 \\ 1/4 & 1/4 & 1/4 & 1/4 & 0 & 0 & 0 \\ 1/4 & 0 & 1/4 & 1/4 & 1/4 & 0 & 0 \\ 1/4 & 0 & 0 & 1/4 & 1/4 & 1/4 & 0 \\ 1/4 & 0 & 0 & 0 & 1/4 & 1/4 & 1/4 \\ 1/4 & 1/4 & 0 & 0 & 0 & 1/4 & 1/4 \end{pmatrix}.$$

Its stationary distribution is $\pi(\kappa_h) = 7/31$, $\pi(\kappa_{b_i}) = 4/31$ (uniform across boundary states), reversibility holds ($\pi(\kappa) \cdot \hat{T}(\kappa, \kappa') = \pi(\kappa') \cdot \hat{T}(\kappa', \kappa)$ for every pair), and the spectrum is

$$\mathrm{spec}(\hat{T}) = \{ 1, \tfrac{1}{2}, \tfrac{1}{2}, -\tfrac{1}{4}, -3/28, 0, 0 \}.$$

The spectral gap is $\varepsilon_{gap}(0) = 1 - |\lambda_2(\hat{T})| = \tfrac{1}{2}$.

The coherent sector is the one-dimensional Perron eigenspace $H_{coh} = \mathrm{span}\{\mathbf{1}\}$; the incoherent sector is its $L^2(\pi)$-orthogonal complement $H_{inc} = \{ f \in \mathbb{R}^{\mathcal{K}} : \langle f, \mathbf{1}\rangle_\pi = 0 \}$, on which $\hat{T}$ contracts in $L^2(\pi)$-norm at rate exactly ½:

$$\| \hat{T} f \|_{L^2(\pi)} \leq \tfrac{1}{2} \cdot \| f \|_{L^2(\pi)} \text{ for all } f \in H_{inc}.$$

These are the facts whose stability under perturbation we now examine.

---

3. Admissible Perturbations

A *perturbed transition operator* is a one-parameter family

$$\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T, \quad \eta \in \mathbb{R},$$

where ΔT is a 7 × 7 real matrix and η is the perturbation strength. We restrict to perturbations satisfying three primary admissibility conditions, which preserve the architectural features required by Stages I–VI.

Primary conditions

P1 — Stochastic preservation. Each row of ΔT sums to zero:

$$\Sigma_j (\Delta T)_{ij} = 0 \text{ for every } i.$$

Consequently $\hat{T}_\eta \cdot \mathbf{1} = \mathbf{1}$ for all η: $\hat{T}_\eta$ remains row-stochastic.

P2 — Support preservation. No transition forbidden by the canonical wheel becomes allowed under perturbation:

$$(\Delta T)_{ij} = 0 \text{ whenever } \hat{T}_{ij} = 0.$$

The directed support graph of $\hat{T}_\eta$ is contained in the directed support graph of $\hat{T}$.

P3 — Positivity. For all allowed edges (ij) (those with $\hat{T}_{ij} > 0$),

$$\hat{T}_{ij} + \eta \cdot (\Delta T)_{ij} > 0.$$

This holds for sufficiently small $|\eta|$: specifically $|\eta| < \min_{ij : \hat{T}_{ij} > 0} \hat{T}_{ij} / |(\Delta T)_{ij}|$. We denote this maximal perturbation strength $\eta_{max}(\Delta T)$.

Structural consequences of P3 for the canonical wheel

Conditions P1, P2, P3 jointly imply three structural facts that the irreducibility and aperiodicity proofs of §4 invoke directly. These follow from P3 combined with the *specific* canonical $\hat{T}$ — in particular, the fact that every hub–boundary entry of $\hat{T}$ is 1/7 or 1/4 > 0, and every diagonal entry of $\hat{T}$ is positive. We retain them as separately-named labels because Propositions 4.1 and 4.2 are easier to read when the load-bearing structural fact is cited explicitly:

P4 — Hub–boundary communication (consequence of P3). All hub–boundary transitions remain positive under perturbation:

$$(\hat{T}_\eta)_{\kappa_h, \kappa_{b_i}} > 0, \quad (\hat{T}_\eta)_{\kappa_{b_i}, \kappa_h} > 0 \text{ for every } i = 1, \ldots, 6.$$

P5 — Boundary cyclic communication (consequence of P3). All nearest-neighbour boundary transitions remain positive under perturbation:

$$(\hat{T}_\eta)_{\kappa_{b_i}, \kappa_{b_{i\pm1}}} > 0 \text{ for every } i = 1, \ldots, 6.$$

P6 — Self-loop persistence (consequence of P3). At least one diagonal entry of $\hat{T}_\eta$ remains positive:

$$\exists\, \kappa \in \mathcal{K} : (\hat{T}_\eta)_{\kappa, \kappa} > 0.$$

For the canonical wheel this is automatic from P3 (all seven diagonal entries of $\hat{T}$ are positive: 1/7 for the hub and 1/4 for each boundary state). The "$\exists\, \kappa$" form rather than the stronger "all κ" form is the minimum sufficient hypothesis for Proposition 4.2 below; we state it in this weaker form to make clear which structural fact aperiodicity actually requires.

These are not independent postulates; for the canonical wheel they are *automatic* from P3. For a hypothetical substrate where $\hat{T}$ had zero hub–boundary entries or zero diagonal entries in the unperturbed state, P4–P6 would impose strictly more than P3. The present paper concerns perturbations of the *canonical* wheel, for which P1–P3 alone are the admissibility postulates and P4–P6 follow.

Admissible perturbation strength

For a given ΔT satisfying P1 and P2, the admissible range is the open interval $(-\eta_{max}(\Delta T), +\eta_{max}(\Delta T))$. Within this range, P3 holds and P4–P6 follow as consequences. The perturbed operator $\hat{T}_\eta$ is a valid row-stochastic admissibility filter on the canonical wheel architecture.

Remark

P1–P3 do not specify *which* perturbations are physically interesting; they specify the broadest class in which the Stage VI architecture remains formally applicable. The classification of perturbations as irrelevant, marginal, or relevant (§11) is a separate, RG-flavoured question that subdivides this admissibility class by long-run effect on the spectral gap.

---

4. Stability of Irreducibility and Aperiodicity

The Stage VI propositions 7.1 (irreducibility) and 7.2 (aperiodicity) of $\hat{T}$ depended on three specific facts: positive hub–boundary transitions in both directions (Proposition 7.1), and at least one positive self-loop (Proposition 7.2). Both facts are preserved by admissibility conditions P4 and P6, giving the following stability propositions.

Proposition 4.1 — Stability of Irreducibility

Let $\hat{T}_\eta$ satisfy P1–P3 for $|\eta| < \eta_{max}(\Delta T)$. Then $\hat{T}_\eta$ is irreducible.

Proof. The argument of Stage VI Proposition 7.1 transfers verbatim. For any pair of states $\kappa, \kappa' \in \mathcal{K}$:

- if $\kappa = \kappa_h$, then $(\hat{T}_\eta)[\kappa_h, \kappa'] > 0$ by P4 (or P3 in the $\kappa_h \to \kappa_h$ direction);
- if $\kappa' = \kappa_h$, symmetrically;
- if both κ, κ' are boundary states, route via the hub: $(\hat{T}_\eta)[\kappa, \kappa_h] \cdot (\hat{T}_\eta)[\kappa_h, \kappa'] > 0$ by P4.

Every pair of states communicates in at most two steps, so $\hat{T}_\eta$ is irreducible.

Proposition 4.2 — Stability of Aperiodicity

Let $\hat{T}_\eta$ satisfy P1–P3 for $|\eta| < \eta_{max}(\Delta T)$. Then $\hat{T}_\eta$ is aperiodic.

Proof. By P6 there exists $\kappa \in \mathcal{K}$ with $(\hat{T}_\eta)_{\kappa, \kappa} > 0$, giving a return path of length 1 at κ. By Proposition 4.1, $\hat{T}_\eta$ is irreducible, so all states share a common period; the period at κ is the gcd of return-path lengths through κ, which includes 1 and is therefore 1.

Combined Perron–Frobenius consequence. For every admissible perturbation $\hat{T}_\eta$ with $|\eta| < \eta_{max}(\Delta T)$, the Perron–Frobenius theorem for irreducible aperiodic stochastic matrices applies:

$\hat{T}_\eta$ has a unique stationary distribution π_η, the Perron eigenvalue 1 is simple, and all other eigenvalues lie strictly inside the open unit disk.

The structural features required for the Stage VI spectral analysis therefore survive perturbation. What remains is to bound *how much* the spectrum can shift before the gap closes.

---

5. Spectral Stability Theorem

The non-trivial spectrum of $\hat{T}$ lies in $[-\frac{1}{4}, \frac{1}{2}]$ with maximum modulus ½. Under perturbation, this maximum modulus can grow, eroding the spectral gap. The following theorem bounds the growth.

Theorem 5.1 — Spectral Gap Stability

Let $\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T$ satisfy P1–P3 for $|\eta| < \eta_{max}(\Delta T)$. Then for every eigenvalue $\lambda_k(\hat{T})$ of $\hat{T}$ and every η in the admissible range, there exists an eigenvalue $\tilde{\lambda}_k(\eta)$ of $\hat{T}_\eta$ with

$$|\tilde{\lambda}_k(\eta) - \lambda_k(\hat{T})| \le C(\Delta T) \cdot |\eta|,$$

where

$$C(\Delta T) = \| \Delta T \|_{op, L^2(\pi)} \le \sqrt{\pi_{max} / \pi_{min}} \cdot \| \Delta T \|_{op} = \sqrt{7/4} \cdot \| \Delta T \|_{op}.$$

Consequently, the spectral gap of $\hat{T}_\eta$ satisfies

$$\varepsilon_{gap}(\eta) \ge \tfrac{1}{2} - C(\Delta T) \cdot |\eta|.$$

In particular, $\varepsilon_{gap}(\eta) > 0$ whenever

$$|\eta| < \min \{ \eta_{max}(\Delta T), 1 / (2 \cdot C(\Delta T)) \}.$$

Proof. Since $\hat{T}$ is reversible with respect to π (Stage VI §6.2), it is self-adjoint on the Hilbert space $L^2(\pi)$. The spectral theorem provides an orthonormal eigenbasis $\{ v_k \}_{k=1}^7$ in the $L^2(\pi)$ inner product, with $\hat{T} = \Sigma_k \lambda_k \cdot v_k \cdot v_k^{*\pi}$. *In this orthonormal basis, the eigenvector matrix V satisfies* $\kappa_{L^2(\pi)}(V) = 1$, where κ denotes the condition number in the $L^2(\pi)$ operator norm.

By the Bauer–Fike theorem for diagonalisable matrices, for every eigenvalue $\tilde{\lambda}$ of $\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T$ there exists an eigenvalue λ_k of $\hat{T}$ with

$$\min_k |\tilde{\lambda} - \lambda_k| \le \kappa_{L^2(\pi)}(V) \cdot \| \eta \cdot \Delta T \|_{op, L^2(\pi)} = |\eta| \cdot \| \Delta T \|_{op, L^2(\pi)}.$$

The $L^2(\pi)$ operator norm is related to the standard operator norm by the comparability bound

$$\| \Delta T \|_{op, L^2(\pi)} \leq \sqrt{\pi_{max} / \pi_{min}} \cdot \| \Delta T \|_{op},$$

derived as follows: for any f, the $L^2(\pi)$ norm satisfies $\pi_{min} \cdot \|f\|^2_2 \leq \|f\|^2_{L^2(\pi)} \leq \pi_{max} \cdot \|f\|^2_2$, so the ratio of induced operator norms is bounded by $\sqrt{\pi_{max} / \pi_{min}} = \sqrt{7/4}$.

Substituting $C(\Delta T) = \|\Delta T\|_{op, L^2(\pi)}$ gives the first inequality of the theorem.

For the spectral gap consequence: the Perron eigenvalue is exactly 1 for all η (by stochasticity P1 and Proposition 4.2). For every other eigenvalue λ_k of $\hat{T}$ with k = 2, …, 7 ($|\lambda_k| \leq ½$), the corresponding perturbed eigenvalue $\tilde{\lambda}_k(\eta)$ satisfies $|\tilde{\lambda}_k(\eta)| \leq |\lambda_k| + C(\Delta T) \cdot |\eta| \leq ½ + C(\Delta T) \cdot |\eta|$. Taking the maximum modulus over k ≥ 2 gives the gap bound.

Epistemic status. Two distinct regularity statements coexist here, and it is worth separating them. *The spectral gap $\varepsilon_{gap}(\eta)$* — defined as 1 minus the maximum modulus over all non-trivial eigenvalues — is *fully Lipschitz* in η on the admissible range, by Theorem 5.1. *Per-eigenvalue tracking*, however, is Lipschitz only for the simple eigenvalues of $\hat{T}$ (namely 1, −¼, and −3/28). At the double eigenvalues at ½ and at 0, the two eigenvalue branches may split with Hölder-½ rather than Lipschitz behaviour in the worst case (Puiseux expansion at exceptional values of η), reflecting the standard non-analyticity of degenerate perturbation theory. The *maximum modulus over each branch* still satisfies the linear bound, so the gap-level statement of the theorem is unaffected; only the question "which specific eigenvalue is which?" admits the Hölder caveat. For all results downstream of Theorem 5.1 (Theorems 6.1, 7.1, 8.1, 12.1), the gap-level Lipschitz statement is what is used.

Quantitative example. Consider the *self-loop / cyclic* perturbation ΔT with non-zero entries +1 at $(\kappa_{b_1}, \kappa_{b_1})$ and −1 at $(\kappa_{b_1}, \kappa_{b_2})$, zero elsewhere. This is a genuinely symmetry-breaking perturbation: it does not lie in the antisymmetric class of Theorem 9.1 below, so the spectrum is expected to move. Direct computation gives $\|\Delta T\|_{op} = \sqrt{2}$ in both the standard and $L^2(\pi)$ operator norms, so the Bauer–Fike constant is $C(\Delta T) = \sqrt{2}$ and the predicted gap stability is

$$\varepsilon_{gap}(\eta) \geq ½ - \sqrt{2} \cdot |\eta|.$$

Numerical evaluation of $|\lambda_2(\hat{T}_\eta)|$ across η ∈ [0, 0.2] (full data tabulated in §10.2) shows $\varepsilon_{gap}(\eta)$ decreases linearly with measured Lipschitz slope ≈ 0.226 — substantially smaller than the Bauer–Fike upper bound $\sqrt{2} \approx 1.414$. The conservatism factor is $\sqrt{2} / 0.226 \approx 6.25$, reflecting structural protections from the wheel's residual symmetries not captured by the operator-norm bound. The true Lipschitz slope is bounded below by 0 (zero in symmetry-aligned directions, by Theorem 9.1) and above by C(ΔT) (the Bauer–Fike bound); generic symmetry-breaking perturbations land in between, well below the upper limit.

---

6. Stationary Measure Continuity

The perturbed operator $\hat{T}_\eta$ has, by Proposition 4.1 + 4.2 + Perron–Frobenius, a unique stationary distribution π_η. We give an explicit formula for its dependence on η.

Theorem 6.1 — Perturbative Expansion of the Stationary Measure

Let $\hat{T}_\eta = \hat{T} + \eta \cdot \Delta T$ satisfy P1–P3 for $|\eta| < \eta_{max}(\Delta T)$. Let Z denote the fundamental matrix of $\hat{T}$:

$$Z = (I - \hat{T} + \mathbf{1} \cdot \pi)^{(-1)} - \mathbf{1} \cdot \pi,$$

where $\mathbf{1} \cdot \pi$ denotes the rank-1 projector onto the Perron eigenspace. Then the stationary distribution π_η satisfies

$$\pi_\eta = \pi + \eta \cdot \pi \cdot \Delta T \cdot Z + \mathcal{O}(\eta^2),$$

and the linear term is bounded by

$$\| \pi \cdot \Delta T \cdot Z \|_{L^2(\pi)} \leq \| \Delta T \|_{op, L^2(\pi)} \cdot \| Z \|_{op, L^2(\pi)} \cdot \| \pi \|_{L^2(\pi)},$$

where $\| Z \|_{op, L^2(\pi)} = 1/\varepsilon_{gap}(\hat{T}) = 2$. (The ∞-norm of Z in the row-sum sense is comparable but not identical; it satisfies $\|Z\|_\infty \leq \sqrt{\pi_{max}/\pi_{min}} \cdot \|Z\|_{op, L^2(\pi)} = \sqrt{7/4} \cdot 2 = \sqrt{7} \approx 2.646$.)

Proof. The defining equation $\pi_\eta \cdot \hat{T}_\eta = \pi_\eta$ reads $\pi_\eta \cdot (\hat{T} + \eta \cdot \Delta T) = \pi_\eta$, equivalently $\pi_\eta \cdot (I - \hat{T}) = \eta \cdot \pi_\eta \cdot \Delta T$. The matrix $(I - \hat{T})$ is singular (with kernel spanned by the constant function); its Moore–Penrose-style inverse on the codomain orthogonal to **1** is the fundamental matrix Z. Writing $\pi_\eta = \pi + \eta \cdot \delta + \mathcal{O}(\eta^2)$ and substituting:

$$(\pi + \eta \cdot \delta + \mathcal{O}(\eta^2)) \cdot (I - \hat{T}) = \eta \cdot (\pi + \mathcal{O}(\eta)) \cdot \Delta T,$$

$$\eta \cdot \delta \cdot (I - \hat{T}) = \eta \cdot \pi \cdot \Delta T + \mathcal{O}(\eta^2),$$

$$\delta \cdot (I - \hat{T}) = \pi \cdot \Delta T + \mathcal{O}(\eta).$$

Solving for δ in the orthogonal complement of **1** (where $I - \hat{T}$ is invertible with inverse $Z + \mathbf{1} \cdot \pi$, simplifying to Z on the orthogonal complement since $\pi \cdot \Delta T$ has zero sum by P1):

$$\delta = \pi \cdot \Delta T \cdot Z + \mathcal{O}(\eta).$$

Hence $\pi_\eta = \pi + \eta \cdot \pi \cdot \Delta T \cdot Z + \mathcal{O}(\eta^2)$ as claimed.

Lipschitz consequence. The map $\eta \mapsto \pi_\eta$ is differentiable at $\eta = 0$ with derivative $\pi \cdot \Delta T \cdot Z$, and Lipschitz on the whole admissible interval $(-\eta_{max}, +\eta_{max})$. The shift $|\pi_\eta - \pi|_\infty$ grows at most linearly in η for small η.

Numerical verification. For the self-loop / cyclic perturbation ΔT with non-zero entries +1 at $(\kappa_{b_1}, \kappa_{b_1})$ and −1 at $(\kappa_{b_1}, \kappa_{b_2})$, direct computation gives:

η	$\pi\_\eta(\kappa\_h)$	$\pi\_\eta(\kappa\_{b_1})$	$\pi\_\eta(\kappa\_{b_2})$	$\pi\_\eta(\kappa\_{b_6})$	$\lVert \pi\_\eta - \pi \rVert_2$
0.00	0.22581	0.12903	0.12903	0.12903	0.00000
0.05	0.22581	0.13654	0.12152	0.13176	0.01134
0.10	0.22581	0.14498	0.11308	0.13483	0.02409
0.15	0.22581	0.15453	0.10353	0.13830	0.03851
0.20	0.22581	0.16543	0.09264	0.14227	0.05497

The shift is linear in η to leading order, with asymptotic slope $\lVert\pi \cdot \Delta T \cdot Z\rVert_2 = 0.2144$ as $\eta \to 0$ (computed directly from the fundamental matrix Z of T̂ via the Theorem 6.1 formula). The finite-difference slopes from the table (0.2268 at η = 0.05, 0.2409 at η = 0.10, …, 0.2749 at η = 0.20) approach this asymptote as η → 0 and drift upward at larger η as the $O(\eta^2)$ terms accumulate — exactly as the Theorem 6.1 expansion predicts.

The hub-component π_η(κ_h) = 7/31 is exactly preserved across this perturbation family. The mechanism is direct: only the b_1-row of T̂_η differs from T̂, and the perturbation affects columns b_1 and b_2 only — the *hub-column* of T̂_η is unchanged from T̂. The column-h stationary equation π_η(h) = Σ_κ π_η(κ) · T̂_η[κ, h] therefore reads

$$\pi\_\eta(h) = (1/7) \cdot \pi\_\eta(h) + \Sigma\_i\ (1/4) \cdot \pi\_\eta(b\_i) = (1/7) \cdot \pi\_\eta(h) + (1/4) \cdot (\ 1 - \pi\_\eta(h)\ ),$$

using only that the hub-row stays at (1/7, …, 1/7), the boundary-to-hub entries stay at 1/4, and Σ_κ π_η(κ) = 1. Solving: (31/28) · π_η(h) = 1/4, so π_η(h) = 7/31 *regardless of how the boundary mass redistributes among b_1, …, b_6*. The hub mass is invariant by stationarity-equation accounting; the boundary mass moves among the boundary states subject to fixed total 24/31. (This argument is column-based and does *not* invoke reversibility — which is in fact broken by this perturbation, since (ΔT)[b_1, b_1] = +1 with no compensating entry on the b_1-column of T̂_η.)

---

7. Entropic Contraction Under Perturbation

The Stage VI Entropic Contraction Theorem 11.1 stated that under the canonical wheel, the χ^2-divergence to stationarity decays geometrically:

$$\chi^2(\ \mu\_n \parallel \pi\ ) \leq (\ ¼\ )^n \cdot \chi^2(\ \mu\_0 \parallel \pi\ ).$$

The factor ¼ = (½)² is the *squared* spectral contraction rate. We now extend this to perturbed operators, handling the reversibility caveat explicitly.

Theorem 7.1 — Entropic Contraction Stability (Reversibility-Preserving Case)

Let $\hat{T}\_\eta$ satisfy P1–P3 for $|\eta| < \eta\_max(\Delta T)$, and assume in addition that $\hat{T}\_\eta$ is reversible with respect to π_η: $\pi\_\eta(\kappa) \cdot (\hat{T}\_\eta)(\kappa, \kappa') = \pi\_\eta(\kappa') \cdot (\hat{T}\_\eta)(\kappa', \kappa)$ for all κ, κ'. Then for any initial probability distribution μ_0,

$$\chi^2(\mu\_n \,\|\, \pi\_\eta) \leq (¼ + C(\Delta T) \cdot |\eta| + \mathcal{O}(\eta^2))^n \cdot \chi^2(\mu\_0 \,\|\, \pi\_\eta),$$

with $C(\Delta T)$ the spectral stability constant from Theorem 5.1.

Proof. The argument of Stage VI Theorem 11.1 transfers with $\hat{T}$ replaced by $\hat{T}\_\eta$ and π replaced by π_η. Since $\hat{T}\_\eta$ is reversible with respect to π_η by assumption, it is self-adjoint on $L^2(\pi\_\eta)$ with non-trivial spectral radius bounded above (in absolute value) by ½ + $C(\Delta T) \cdot |\eta|$ via Theorem 5.1. The χ^2-divergence is the squared $L^2(\pi\_\eta)$-norm of the relative density (Stage VI §11.1 proof), contracting at rate $(½ + C(\Delta T)\cdot|\eta|)^2$ per step. Expanding the square:

$$(½ + C(\Delta T) \cdot |\eta|)^2 = ¼ + 2 \cdot ½ \cdot C(\Delta T) \cdot |\eta| + C(\Delta T)^2 \cdot \eta^2 = ¼ + C(\Delta T) \cdot |\eta| + \mathcal{O}(\eta^2).$$

The factor of 2 in the cross-term combines with the canonical ½ to give a clean unity coefficient on $C(\Delta T)\cdot|\eta|$. This is the linear correction to the canonical χ^2-rate ¼.

Reversibility-Breaking Perturbations

Generic perturbations ΔT do *not* preserve reversibility: if $\Delta T[\kappa, \kappa'] / \pi(\kappa') \neq \Delta T[\kappa', \kappa] / \pi(\kappa)$ for some pair, then $\hat{T}\_\eta$ is non-reversible with respect to π, and remains non-reversible with respect to π_η to leading order in η.

In the non-reversible case, $\hat{T}\_\eta$ is no longer self-adjoint on any natural inner product, and the spectral contraction does not directly translate to χ^2-contraction. The standard substitute is the *pseudo-spectral* estimate, which controls the decay rate via the largest singular value of $\hat{T}\_\eta|\_\{H\_inc\}$ rather than its eigenvalue:

$$\chi^2(\mu\_n \,\|\, \pi\_\eta) \leq \sigma\_2(\hat{T}\_\eta)^{2n} \cdot \chi^2(\mu\_0 \,\|\, \pi\_\eta),$$

where $\sigma\_2(\hat{T}\_\eta)$ is the largest singular value of $\hat{T}\_\eta$ restricted to H_inc (the zero-π_η-mean subspace). For small η, $\sigma\_2(\hat{T}\_\eta)$ is close to $|\lambda\_2(\hat{T}\_\eta)|$ but not equal: $\sigma\_2 \geq |\lambda\_2|$ in general, with equality if and only if $\hat{T}\_\eta$ is normal on H_inc.

For the canonical wheel, $\hat{T}$ is reversible (hence normal in $L^2(\pi)$) and $\sigma\_2(\hat{T}) = |\lambda\_2(\hat{T})| = ½$. For non-reversible perturbations, $\sigma\_2(\hat{T}\_\eta) = ½ + \mathcal{O}(|\eta|)$ by continuity of singular values, but the constant in front of |η| may exceed the spectral-gap stability constant of Theorem 5.1.

Restricted conclusion. Entropic contraction stability in the sense of Theorem 7.1 is guaranteed for reversibility-preserving perturbations. For reversibility-breaking perturbations, contraction in *$L^2(\pi\_\eta)$*-norm still holds at rate ½ + $\mathcal{O}(|\eta|)$, but the χ^2-divergence interpretation requires the singular-value substitute, and the contraction constant may be larger by an $\mathcal{O}(1)$ factor. Continuum-geometry consequences (Stage V Lipschitz bound) depend on the $L^2(\pi\_\eta)$-rate, not the χ^2-rate per se, so the principal Stage V chain remains intact in either case.

8. Continuum Regularity Stability

The Stage V continuum Lipschitz bound is

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

Under perturbation η, each constant on the right-hand side may shift: $L_\Phi \to L_\Phi(\eta)$, $L \to L(\eta)$, $A \to A(\eta)$, $A_- \to A_-(\eta)$, and $\varepsilon_{gap} \to \varepsilon_{gap}(\eta)$.

Theorem 8.1 — Continuum Lipschitz Stability

Suppose the Stage IV substrate primitives L_Φ, L, A, A_- are themselves Lipschitz in η on the admissible perturbation interval, with derivatives bounded in some norm. Then the continuum Lipschitz constant K_∞ is Lipschitz in η:

$$K_\infty(\eta) = K_\infty(0) + \mathcal{O}(|\eta|) \text{ as } \eta \to 0,$$

and remains finite for all $|\eta| < \min \{ \eta_{max}(\Delta T), 1/(2 \cdot C(\Delta T)) \}$.

Proof. Each factor on the right-hand side of the K_∞ formula is Lipschitz in η by hypothesis (for L_Φ, L, A, A_-) and by Theorem 5.1 (for ε_{gap}). The product of Lipschitz functions is Lipschitz on any interval where the denominator stays bounded away from zero, which is guaranteed by $\varepsilon_{gap}(\eta) \geq \frac{1}{2} - C(\Delta T) \cdot |\eta| > 0$ on the admissible range. Finiteness and Lipschitz continuity follow.

Numerical illustration. Under the unit-substrate calibration $L_\Phi = L = A = A_- = 1$ introduced in Stage VI §12 (and retained throughout this paper), the canonical Lipschitz constant is $K_\infty(0) = 2 / \varepsilon_{gap}(0) = 4$. For the boundary anisotropy perturbation of §10.1 (spectrally invariant per Theorem 9.1), $\varepsilon_{gap}(\eta) = \frac{1}{2}$ exactly, so $K_\infty(\eta) = 4$ exactly across the full admissible range. For the self-loop / cyclic perturbation of §10.2, $\varepsilon_{gap}(\eta) \approx \frac{1}{2} - 0.226 \cdot \eta$, giving $K_\infty(\eta) \approx 2 / (\frac{1}{2} - 0.226 \cdot \eta)$. This shifts from $K_\infty(0) = 4.000$ at $\eta = 0$ to $K_\infty(0.2) \approx 4.401$ at $\eta = 0.20$ — a 10% degradation across the boundary of the natural perturbation range, again under the unit-substrate calibration. Non-unit calibrations rescale all K_∞ values proportionally; the *fractional* shift $K_\infty(\eta)/K_\infty(0)$ is calibration-independent.

Programme consequence. The continuum Lipschitz constant degrades *continuously* under admissible perturbation. Catastrophic loss of regularity ($K_\infty \to \infty$) only occurs when $\varepsilon_{gap}(\eta) \to 0$, which requires the perturbation strength to exceed $1/(2 \cdot C(\Delta T))$ — well outside the small-η regime in which the perturbation expansion is valid. The Lipschitz continuum-geometry conclusion of Stage V is therefore not a knife-edge property of the canonical operator; it is the generic behaviour on an open neighbourhood.

9. Symmetry-Protected Spectral Invariance

Theorem 5.1 bounds the spectral shift under *any* admissible perturbation. The bound is uniform: every admissible ΔT shifts the spectrum at most linearly in η, regardless of structure. The result is sharp in the operator-norm sense — but it is also blind to a much stronger property of the canonical wheel.

A direct computation reveals it: certain admissible perturbations leave the *entire spectrum* of $\hat{T}_\eta$ unchanged, *exactly*, for all η in the admissible range. The eigenvectors of $\hat{T}_\eta$ rotate; the eigenvalues do not move. To 12 decimal places at $\eta = 0, 0.05, 0.10, 0.20$, the characteristic polynomial of $\hat{T}_\eta$ is identical to the characteristic polynomial of $\hat{T}$.

This is not coincidence. It is the wheel's dihedral D_6 symmetry, acting through the matrix determinant lemma, forcing a positive-dimensional subspace of admissible perturbations to be **spectrally silent**. The result is sharper than stability — it is invariance — and it is structural, not numerical.

Theorem 9.1 — Symmetry-Protected Spectral Invariance (Rank-1 Case)

Let $\sigma \in D_6$ be a wheel symmetry (an element of the dihedral group acting on the boundary hexagon and fixing κ_h), and let P_σ denote its 7×7 permutation representation on $\mathbb{R}^{\mathcal{K}}$. Suppose:

(i) $\hat{T}$ commutes with P_σ: $P_\sigma \hat{T} P_\sigma^{-1} = \hat{T}$;

(ii) ΔT is a rank-1 matrix $\Delta T = e_\kappa \otimes v^T$ with $\kappa \in \mathcal{K}$ fixed by σ: $P_\sigma e_\kappa = e_\kappa$;

(iii) the column $v \in \mathbb{R}^{\mathcal{K}}$ is antisymmetric under σ: $P_\sigma v = -v$.

Then for every η in the admissible range,

$$\mathrm{spec}(\hat{T} + \eta \cdot \Delta T) = \mathrm{spec}(\hat{T}),$$

and equivalently the characteristic polynomial is independent of η:

$$\det(\hat{T} + \eta \cdot \Delta T - \lambda \cdot I) = \det(\hat{T} - \lambda \cdot I) \text{ for every } \lambda \in \mathbb{C}.$$

Remark on (ii). The σ-fixing hypothesis on κ is essential to the proof: the resolvent-commutation step uses $P_\sigma^{-1} \cdot e_\kappa = e_\kappa$, which holds only when σ fixes κ. For perturbations seated at a non-fixed κ ($P_\sigma \cdot e_\kappa = e_{\sigma(\kappa)} \neq e_\kappa$), the argument breaks down. Spectral invariance for such perturbations may still hold under additional structure — for instance, if v vanishes on the σ-orbit of κ, or if ΔT is replaced by a σ-symmetrised sum $\Sigma_{\tau \in \langle\sigma\rangle} e_{\tau(\kappa)} \otimes \tau(v)^T$ — but these extensions lie outside the present scope.

Proof. By the matrix determinant lemma for rank-1 perturbations,

$$\det(\hat{T} + \eta \cdot e_\kappa \cdot v^T - \lambda \cdot I) = \det(\hat{T} - \lambda \cdot I) \cdot (1 + \eta \cdot v^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa),$$

valid for all λ outside $\mathrm{spec}(\hat{T})$. It suffices to show

$$v^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa = 0 \text{ for all such } \lambda.$$

By (i), $\hat{T}$ commutes with P_σ, hence so does every polynomial and rational function of $\hat{T}$, including the resolvent: $P_\sigma \cdot (\hat{T} - \lambda \cdot I)^{-1} = (\hat{T} - \lambda \cdot I)^{-1} \cdot P_\sigma$.

By (ii), $P_\sigma \cdot e_\kappa = e_\kappa$; by (iii), $P_\sigma \cdot v = -v$.

Then

$v^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa = v^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot P_\sigma \cdot e_\kappa$ (by (ii)) $= v^T \cdot P_\sigma \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa$ (by (i), commutation) $= (P_\sigma v)^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa$ (transpose; P_σ orthogonal) $= (-v)^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa$ (by (iii)) $= -v^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa$.

A quantity equal to its own negative is zero. Hence $v^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot e_\kappa = 0$, the perturbative factor in the matrix determinant lemma collapses to 1, and the characteristic polynomial is unchanged for all η.

Higher-Rank Generalisation

The same argument extends to higher-rank perturbations, by the matrix determinant lemma for general low-rank updates.

Corollary 9.2 — Symmetry-Protected Invariance for General Rank

Let $\sigma \in D_6$ with $P_\sigma \hat{T} P_\sigma^{-1} = \hat{T}$. Define the σ-symmetric and σ-antisymmetric subspaces of $\mathbb{R}^\mathcal{K}$:

$$H_\sigma^+ = \{ x \in \mathbb{R}^\mathcal{K} : P_\sigma x = +x \}, \quad H_\sigma^- = \{ x \in \mathbb{R}^\mathcal{K} : P_\sigma x = -x \}.$$

Let $U = [u_1, \ldots, u_r]$ be a $7 \times r$ matrix with columns $u_i \in H_\sigma^+$, and let $V = [v_1, \ldots, v_r]$ be a $7 \times r$ matrix with columns $v_i \in H_\sigma^-$. Then for $\Delta T = U \cdot V^T$,

$$\mathrm{spec}(\hat{T} + \eta \cdot \Delta T) = \mathrm{spec}(\hat{T}) \text{ for every } \eta \in \mathbb{R}.$$

Proof. By the matrix determinant lemma for rank-r perturbations,

$$\det(\hat{T} + \eta \cdot U V^T - \lambda \cdot I) = \det(\hat{T} - \lambda \cdot I) \cdot \det(I_r + \eta \cdot V^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot U).$$

The $r \times r$ matrix $V^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot U$ has (i, j) entry $v_i^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot u_j$. Since $(\hat{T} - \lambda \cdot I)^{-1}$ commutes with P_σ and $u_j \in H_\sigma^+$, the vector $(\hat{T} - \lambda \cdot I)^{-1} \cdot u_j$ lies in H_σ^+. Since P_σ is a permutation matrix (hence orthogonal in the *standard* Euclidean inner product on $\mathbb{R}^\mathcal{K}$), its $+1$ and -1 eigenspaces H_σ^+ and H_σ^- are mutually orthogonal in that inner product — and this is the inner product implicit in the bilinear pairing $v_i^T \cdot w := \Sigma_\kappa (v_i)_\kappa \cdot w_\kappa$ used here. Therefore $v_i^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot u_j = 0$ for every i, j, and $V^T \cdot (\hat{T} - \lambda \cdot I)^{-1} \cdot U$ is the zero matrix. The $\det(I_r + 0)$ factor collapses to 1, and the spectrum is unchanged.

Structure of the spectrally-silent ansatz subspaces. Corollary 9.2 covers perturbations of the *pure* form $\Delta T = U V^T$ with column-space $\subset H_\sigma^+$ and row-space $\subset H_\sigma^-$. This subspace of matrices has dimension $\dim(H_\sigma^+) \cdot \dim(H_\sigma^-)$. An independent *transpose-shaped* subspace — perturbations $\Delta T = V' U'^T$ with column-space $\subset H_\sigma^-$ and row-space $\subset H_\sigma^+$ — is silent by the same argument with the U/V roles swapped, and has the same dimension.

The dihedral group D_6 has six reflections, splitting into two structurally distinct types:

- **Vertex-vertex reflections** (three of them, e.g. axis through κ_{b_1} and κ_{b_4}). Each fixes κ_h plus two antipodal boundary states and swaps the remaining four boundary states pairwise: $\dim(H_\sigma^+) = 5$ (three fixed vectors plus two symmetric combinations from the boundary pair swaps), $\dim(H_\sigma^-) = 2$. Each ansatz subspace has dimension $5 \cdot 2 = 10$.
- **Edge-edge reflections** (three of them, e.g. axis through the midpoint of edge (κ_{b_1}, κ_{b_2}) and the midpoint of edge (κ_{b_4}, κ_{b_5})). Each fixes only κ_h and swaps the six boundary states in three pairs: $\dim(H_\sigma^+) = 4$ (κ_h plus three symmetric combinations), $\dim(H_\sigma^-) = 3$. Each ansatz subspace has dimension $4 \cdot 3 = 12$.

Across all six reflections, the per-reflection silent ansatz subspaces have dimensions 10 + 10 + 10 + 12 + 12 + 12 (with substantial pairwise overlap, since the same admissible perturbation can be silent under more than one reflection). The exact dimension of the *union* — and of its admissibility-preserving intersection with P1 ∩ P2 — is a finite linear-algebra computation not pursued here.

Caveat — the conjugation-antisymmetric subspace is *not* entirely silent. For each $\sigma \in D_6$, the direct sum of the two pure-form silent subspaces is the (-1)-eigenspace of the conjugation action $M \mapsto P_\sigma M P_\sigma^{-1}$ on $Mat_7(\mathbb{R})$, of dimension $2 \cdot \dim(H_\sigma^+) \cdot \dim(H_\sigma^-)$ — 20 dimensions for vertex-vertex reflections and 24 dimensions for edge-edge reflections. An arbitrary element $\Delta T = U V^T + V' U'^T$ of this subspace — combining both pure forms — satisfies only the *weaker* invariance that the characteristic polynomial $\det(\hat{T} + \eta \cdot \Delta T - \lambda \cdot I)$ is an *even function of η*. This follows because the resolvent trace identity $\mathrm{tr}((R(\lambda)\, \Delta T)^k) = 0$ for every odd k (by the antisymmetry under conjugation) makes the Fredholm-like expansion $\log \det(I - \eta \cdot R(\lambda) \cdot \Delta T) = -\Sigma_{k \geq 1} (\eta^k / k) \cdot \mathrm{tr}((R \cdot \Delta T)^k)$ an even series in η; the matrix determinant lemma's perturbative factor $\det(I_r + \eta \cdot V^T R U + ...)$ develops cross-terms between the $U V^T$ and $V' U'^T$ pieces at order η^2, preventing exact invariance. Numerical verification: for the mixed perturbation $\Delta T = e_{\kappa_{b_1}} \otimes (e_{\kappa_{b_2}} - e_{\kappa_{b_6}}) + (e_{\kappa_{b_2}} - e_{\kappa_{b_6}}) \otimes (e_{\kappa_{b_1}} - e_{\kappa_h})$, the $|\lambda_2|$ values at $\pm\eta$ are identical to machine precision (e.g., 0.5311972345 at both ±0.05, 0.5671420847 at both ±0.10), but they exceed ½ at every $\eta \neq 0$ — confirming $O(\eta^2)$ shift consistent with even-function-of-η invariance, not the O(1) flat behaviour of the pure cases.

Programmatic interpretation. The pure ansatz subspaces (10-dim or 12-dim per reflection, depending on type) are *fully* protected — spectrum exactly invariant to all orders in η. The remaining mixed directions in the conjugation-antisymmetric subspace are *partially* protected — they exhibit zero Lipschitz slope at $\eta = 0$ (no first-order spectral response) but admit $O(\eta^2)$ spectral shifts. Both protections substantially enlarge the wheel's robustness, but only the pure forms deliver the full Theorem 9.1 spectral invariance. Extending exact invariance to the entire conjugation-antisymmetric subspace would require additional structural input — for instance, cross-term cancellation between the $U V^T$ and $V' U'^T$ blocks under a transpose-symmetric coupling — and is left as an open problem.

Worked Example

Take σ to be the reflection $b_2 \leftrightarrow b_6$, $b_3 \leftrightarrow b_5$, fixing κ_h, κ_{b_1}, κ_{b_4}. Direct check confirms $P_\sigma \hat{T} P_\sigma^{-1} = \hat{T}$. Choose $\kappa = \kappa_{b_1}$ (fixed by σ) and $v = e_{\kappa_{b_2}} - e_{\kappa_{b_6}}$ (antisymmetric under σ: $P_\sigma v = e_{\kappa_{b_6}} - e_{\kappa_{b_2}} = -v$). The resulting rank-1 perturbation is

$$\Delta T = e_{\kappa_{b_1}} \cdot (e_{\kappa_{b_2}} - e_{\kappa_{b_6}})^T,$$

i.e., entries +1 at (b_1, b_2) and −1 at (b_1, b_6), zero elsewhere. This is precisely the boundary anisotropy perturbation of §10.1.

Numerical verification. The characteristic polynomial of $\hat{T}_\eta$ for this perturbation, computed at $\eta = 0, 0.05, 0.10, 0.20$, is identical to 12 decimal places, with coefficient vector

$$(+1, -1.642857, +0.562500, +0.142857, -0.055804, -0.006696, 0, 0),$$

independent of η. Eigenvalues match $\mathrm{spec}(\hat{T}) = \{ 1, \tfrac{1}{2}, \tfrac{1}{2}, -\tfrac{1}{4}, -3/28, 0, 0 \}$ to machine precision throughout the admissible range $|\eta| < \tfrac{1}{4}$. Verifying at $\lambda = 0.3$ (a non-eigenvalue): the resolvent entries $(\hat{T} - 0.3\cdot I)^{-1}[b_2, b_1] = (\hat{T} - 0.3\cdot I)^{-1}[b_6, b_1] = 1.783816$ are identical to 6 decimal places. The equality is *not* numerical coincidence: since $R(0.3) = (\hat{T} - 0.3\cdot I)^{-1}$ commutes with P_σ (because $\hat{T}$ does), one has $R(0.3)[b_6, b_1] = (P_\sigma R(0.3) P_\sigma^{-1})[b_2, b_1] = R(0.3)[b_2, b_1]$ directly by the σ-action on indices. The bilinear pairing $v^T R(0.3) e_\kappa = (1)\cdot R(0.3)[b_2, b_1] + (-1)\cdot R(0.3)[b_6, b_1] = 0$ is then forced, exactly as the Theorem 9.1 proof requires.

Structural Significance

Theorem 9.1 and Corollary 9.2 promote the wheel's spectral data from "stable under perturbation" (Theorem 5.1) to *exactly invariant* under a substantial subspace of admissible perturbations.

Three programmatic consequences:

1. **Symmetry-respecting noise is invisible to first order, exactly invisible on pure forms.** Microscopic disorder ΔT in the (−1)-eigenspace of conjugation by any $\sigma \in D_6$ contributes *nothing* to the spectral gap at first order in η — the Lipschitz slope of ε_{gap} with respect

to such ΔT vanishes at $\eta = 0$. Within this class, disorder of the *pure* form $\Delta T = U V^T$ (or its transpose $\Delta T = V' U'^T$) contributes nothing at *any* order — exactly invariant to all η. The Stage VI spectral gap $\varepsilon_{gap} = \frac{1}{2}$ is therefore a *symmetry-protected quantity* on a substantial subspace of admissible directions, with full to-all-orders protection on the pure forms and first-order protection on the broader conjugation-antisymmetric subspace.

2. **The silent ansatz subspaces are substantial within the admissibility space.** For each $\sigma \in D_6$, two ansatz subspaces of dimensions 10 or 12 (depending on reflection type, per the dimension count above) lie within the 49-dimensional matrix algebra $Mat_7(\mathbb{R})$. The admissibility-preserving intersection with P1 (row sums zero) ∩ P2 (wheel-admissible support) cuts these down; the admissible perturbation space is approximately 29-dimensional (one direction per allowed edge of the wheel, minus seven row-sum constraints), and the dimension of the admissibility-preserving silent subspace is a finite linear-algebra exercise not pursued here. The qualitative claim — that a structurally meaningful fraction of admissible perturbation directions is spectrally silent — is supported by the per-reflection dimensions but the precise fractional figure is unverified. Only *fully symmetry-breaking* perturbations — those with no surviving D_6-antisymmetry under any σ — can erode the spectral gap at first order.
3. **Spectral gap stability and spectral invariance are complementary.** Theorem 5.1 (Bauer–Fike) bounds the worst-case sensitivity of ε_{gap}. Theorem 9.1 / Corollary 9.2 (symmetry protection) identifies the directions in which ε_{gap} is *exactly* invariant. Together, they characterise the spectral response of the wheel under perturbation: Lipschitz everywhere with $O(1)$ constant, exactly zero on the pure-form silent subspaces, and only $O(\eta^2)$ on the broader conjugation-antisymmetric class.

This is the wheel architecture's structural over-engineering: not merely a stable spectral gap, but a *symmetry-protected* one. Many natural microscopic perturbations of the canonical $\hat{T}$ — including those that arise from local disorder, anisotropy, or directional bias respecting the wheel's reflection symmetries — leave every Stage VI computed quantity ($\varepsilon_{gap} = \frac{1}{2}$, refinement rate $(\frac{1}{2})^n$, χ^2-rate $(\frac{1}{4})^n$, $K_\infty = 4$ under unit substrate calibration) *exactly* unchanged in the pure-form case, and unchanged to first order in η for any conjugation-antisymmetric direction.

---

10. Worked Examples

We illustrate Theorems 5.1, 6.1, 8.1, and 9.1 on three explicit perturbations, computed directly on the canonical $\hat{T}$.

10.1 Boundary anisotropy at κ_{b_1} — Spectrally invariant

$\hat{T}_\eta[\kappa_{b_1}, \kappa_{b_2}] = 1/4 + \eta$, $\hat{T}_\eta[\kappa_{b_1}, \kappa_{b_6}] = 1/4 - \eta$,

all other entries unchanged. This is the rank-1 perturbation $\Delta T = e_{\kappa_{b_1}} \otimes (e_{\kappa_{b_2}} - e_{\kappa_{b_6}})$. Admissible range: $|\eta| < 1/4$ (positivity of $\hat{T}_\eta[\kappa_{b_1}, \kappa_{b_6}]$).

η	**spec($\hat{T}$_η)**	**ε_gap(η)**
0.00	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½
0.05	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½
0.10	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½
0.20	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½
0.24	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½

Computed status. Spectrum is identical across the entire admissible range, confirming Theorem 9.1. ε_gap(η) = ½ for all admissible η; K_∞(η) = K_∞(0) exactly under unit-substrate calibration. This perturbation lies in the symmetry-aligned class and is *strictly irrelevant* in the RG sense of §11.

10.2 Self-loop versus cyclic at κ_{b_1} — Spectrally active

$\hat{T}_\eta[\kappa_{b_1}, \kappa_{b_1}] = 1/4 + \eta$, $\hat{T}_\eta[\kappa_{b_1}, \kappa_{b_2}] = 1/4 - \eta$,

all other entries unchanged. This is $\Delta T = e_{\kappa_{b_1}} \otimes (e_{\kappa_{b_1}} - e_{\kappa_{b_2}})$. The perturbation is *not* anti-symmetric under any reflection fixing κ_{b_1}, so Theorem 9.1 does not apply.

| η | λ_2($\hat{T}$_η) | ε_gap(η) | Bauer–Fike bound (½ − √2·|η|) | |------|-----------|----------|---------------------------------| | 0.00 | 0.50000 | 0.50000 | 0.50000 | | 0.05 | 0.50897 | 0.49103 | 0.42929 | | 0.10 | 0.51937 | 0.48063 | 0.35858 | | 0.15 | 0.53145 | 0.46855 | 0.28787 | | 0.20 | 0.54548 | 0.45452 | 0.21716 |

Computed status. ε_gap drops linearly in η at slope ≈ 0.226 (fit on 21 data points across η ∈ [0, 0.2]). The actual sensitivity is a factor of ≈ 6 smaller than the worst-case Bauer–Fike bound √2 ≈ 1.414, reflecting structural protections beyond the abstract operator-norm bound. K_∞(η) ≈ 2 / ε_gap(η) varies smoothly from 4.000 at η = 0 to 4.401 at η = 0.20 — a 10% degradation across the boundary of the natural perturbation range.

10.3 Hub-coupling asymmetry — Spectrally invariant

$\hat{T}_\eta[\kappa_h, \kappa_{b_1}] = 1/7 + \eta$, $\hat{T}_\eta[\kappa_h, \kappa_{b_4}] = 1/7 - \eta$.

$\Delta T = e_{\kappa_h} \otimes (e_{\kappa_{b_1}} - e_{\kappa_{b_4}})$. Under the reflection σ that fixes κ_h and exchanges $\kappa_{b_1} \leftrightarrow \kappa_{b_4}$ (together with $\kappa_{b_2} \leftrightarrow \kappa_{b_3}$, $\kappa_{b_5} \leftrightarrow \kappa_{b_6}$), $\hat{T}$ is invariant, κ_h is fixed, and $v = e_{\kappa_{b_1}} - e_{\kappa_{b_4}}$ is antisymmetric. Theorem 9.1 applies.

η	**spec($\hat{T}$_η)**	**ε_gap(η)**
0.00	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½
0.05	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½
0.10	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½

η	$\mathrm{spec}(\hat{T}_\eta)$	$\varepsilon_{gap}(\eta)$
0.13	{ 1, ½, ½, −¼, −3/28, 0, 0 }	½

Computed status. Spectrum invariant throughout the admissible range $|\eta| < 1/7$. The hub can have its couplings to opposite boundary states arbitrarily redistributed (within positivity) without any effect on refinement convergence rates.

10.4 Local disorder — Generic perturbation (heuristic)

Allow all wheel-admissible edges to fluctuate independently and randomly:

$$\hat{T}_\eta[\kappa, \kappa'] = \hat{T}[\kappa, \kappa'] + \eta \cdot \xi_{\kappa, \kappa'},$$

with ξ drawn from a zero-mean distribution and row-sum corrections applied. For sufficiently small η, P3 holds almost surely. The perturbation is generically symmetry-breaking (the random draw does not respect D_6), so Theorem 5.1 governs, not Theorem 9.1.

Heuristic estimate. A symmetry-averaging argument suggests that the expected spectral gap satisfies

$$\mathbb{E}[\ \varepsilon_{gap}(\eta)\ ] \approx ½ - \mathcal{O}(\ \eta \cdot \sigma_\xi\ ),$$

where σ_ξ is the standard deviation of ξ entries and the constant in front of $\eta \cdot \sigma_\xi$ is the *expected* effective Lipschitz constant across the perturbation ensemble — typically smaller than the worst-case Bauer–Fike bound by symmetry averaging across the random draw. This formula is heuristic, not derived; a rigorous statement requires either an explicit symmetry-decomposition of the disorder ensemble or a free-probability calculation that is beyond the present paper. What can be stated rigorously: for *any* realisation of the random perturbation within the admissible range, Theorem 5.1 gives $\varepsilon_{gap}(\eta) \geq ½ - \sqrt{(7/4)} \cdot \|\Delta T_{realised}\|_{op} \cdot |\eta|$, providing a deterministic lower bound that is uniform in the realisation.

Programmatic significance. This is the substrate-engineering interpretation: realistic substrates carry stochastic local disorder, and the canonical wheel's robustness ensures that *typical* disorder configurations leave the universality class intact. The rigorous deterministic bound suffices for this qualitative claim; the heuristic expected-gap formula refines it but is not load-bearing for any conclusion in this paper.

---

11. RG-Theoretic Classification of Perturbations

Borrowing language from renormalisation-group theory, admissible perturbations split into three classes by their long-run effect on the spectral gap.

Terminology note. The classification below uses *irrelevant*, *marginal*, and *relevant* as *loose analogy* to standard RG terminology rather than a literal correspondence. In conventional RG, "marginal" means the perturbation flows logarithmically at a fixed point (sits at the boundary of relevance, deformed by loop corrections); here we use it for perturbations whose spectral gap approaches zero only at the boundary of the admissible η-range. The substantive content of the classification — open / boundary-critical / external regions of the universality class — is well-defined and structurally analogous to the RG-flow picture, but the specific RG-flow mechanics are not literally present in the finite-dimensional admissibility-rule space studied here.

11.1 Irrelevant Perturbations

A perturbation ΔT is *irrelevant* if there exists $\delta > 0$ such that $\varepsilon\_gap(\eta) \geq \delta$ for all η in the admissible range. Irrelevant perturbations may change numerical constants (the stationary measure, the gap value, the Lipschitz constant K_∞) but preserve all qualitative features: unique persistent coherent sector, exponential refinement contraction, geometric entropy decay, finite continuum Lipschitz regularity.

Examples:

- All symmetry-aligned perturbations (Theorem 9.1) — *strictly irrelevant*, with $\varepsilon\_gap(\eta) = ½$ exactly.
- Symmetry-breaking perturbations with bounded ‖ΔT‖ — irrelevant in their admissible range, by Theorem 5.1.
- Stochastic local disorder at amplitude $\eta < 1/(4\cdot\bar{C})$, where $\bar{C}$ is the typical Bauer–Fike constant.

Irrelevant perturbations form an *open set* in the space of admissible perturbations and contain a neighbourhood of ΔT = 0.

11.2 Marginal Perturbations

A perturbation ΔT is *marginal* if $\varepsilon\_gap(\eta) \to 0$ as $\eta \to \eta^*$ for some boundary value η^* of the admissible range, but $\varepsilon\_gap(\eta) > 0$ for all $\eta < \eta^*$. Marginal perturbations *deform* the continuum Lipschitz constant $K\_\infty(\eta)$ toward infinity as $\eta \to \eta^*$, slowing refinement contraction but not destroying it at any finite $\eta < \eta^*$.

Examples: A perturbation in which boundary–boundary transition weights are systematically driven toward making the boundary ring deterministic-cyclic (a perturbation in this class is not worked out explicitly here). At $\eta < \eta^*$, the chain remains aperiodic and irreducible; at $\eta = \eta^*$, aperiodicity fails and the gap closes.

Marginal perturbations remain *within* the universality class for every η in the strict interior of the admissible range, but the class boundary is approached as $\eta \to \eta^*$.

11.3 Relevant Perturbations

A perturbation ΔT is *relevant* if $\varepsilon_{gap}(\eta) = 0$ for some η in the admissible range, or if the perturbed operator fails irreducibility, aperiodicity, or one of the architectural admissibility conditions P4–P6 at some η. Relevant perturbations destroy the $K = 7$ refinement universality class.

Examples:

- Hub failure — complete removal of one or more hub–boundary edges (P4 violation): irreducibility may fail, multiple stationary sectors may emerge, the spectral gap closes (Example 11.4 below).
- Pure boundary rotation — boundary transitions become deterministic, e.g. $b_i \to b_{i+1}$ with no self-loop or hub coupling: aperiodicity fails, $|\lambda_2| = 1$, spectral gap closes, refinement does not converge to a coherent sector.
- Boundary decoupling — separation of the boundary ring into multiple disconnected arcs (P5 violation): irreducibility fails on the disconnected sectors.

Relevant perturbations *leave* the universality class entirely. The continuum-geometry conclusion of Stage V fails: either the continuum is not Lorentzian, or no continuum exists at all.

11.4 Example: Hub-Failure as Relevant Perturbation

Consider the perturbation that removes the hub–boundary edge (κ_h, κ_{b_1}):

$\Delta T[\kappa_h, \kappa_{b_1}] = -1/7$, $\Delta T[\kappa_h, \kappa_h] = +1/7$,

with $\eta = 1$ representing complete edge removal. For $\eta < 1$, $\hat{T}_\eta$ satisfies P4 (positivity preserved) and $\varepsilon_{gap}(\eta)$ varies smoothly. At $\eta = 1$ exactly, P4 fails: κ_h no longer transitions directly to κ_{b_1}.

The chain remains irreducible at $\eta = 1$ (boundary routes $b_1 \to b_2$ or $b_1 \to b_6$ still pass to other boundary states and back to the hub), and aperiodicity holds (κ_h still has a self-loop). The spectral gap remains positive but smaller. So *single* hub edge removal is irrelevant.

If *two* hub edges are removed at antipodal boundary states (b_1 and b_4), the chain remains irreducible — but the boundary ring's communication through the hub is reduced. Spectral gap shrinks measurably.

If *all six* hub–boundary edges are removed ($\eta = 1$ across the hub row), the hub becomes isolated: hub is no longer reachable from any boundary state, and vice versa. The chain decomposes into the hub-only state and the boundary cycle C_6. Two stationary sectors emerge, the simple Perron condition fails, and the universality class is left.

The transition from irrelevant (single edge removal) to relevant (all six removed) is a *combinatorial phase transition* in the perturbation parameter space — itself worth a dedicated analysis, beyond the present paper's scope.

12. The K = 7 Refinement Universality Class

We can now define the universality class explicitly and state the open-stability theorem.

Definition 12.1 — K = 7 Refinement Universality Class

The *K = 7 refinement universality class* $\mathcal{C}_{K=7}$ is the set of all row-stochastic 7 × 7 matrices S on the catalogue $\mathcal{K} = \{\kappa_h, \kappa_{b_1}, \ldots, \kappa_{b_6}\}$ satisfying:

- **(U1) Stochasticity:** $S \cdot \mathbf{1} = \mathbf{1}$, $S \geq 0$.
- **(U2) Wheel-admissible support:** $S_{ij} > 0$ only where the canonical $\hat{T}_{ij} > 0$ (no forbidden transitions allowed; equivalently, the directed support of S is contained in the wheel graph W_6).
- **(U3) Irreducibility:** every pair of states communicates under S.
- **(U4) Aperiodicity:** at least one state has period 1.
- **(U5) Positive spectral gap:** $\varepsilon_{gap}(S) := 1 - \max_{\lambda \neq 1} |\lambda(S)| > 0$.
- **(U6) Bounded local response:** $L(S) < \infty$, where L(S) is the Stage V position-coherence Lipschitz response constant — it bounds the variation of the local refinement filter under small displacements in the substrate position, ensuring finite propagation of substrate disturbances through the refinement step (Stage V Theorem 2).

For every $S \in \mathcal{C}_{K=7}$, the principal Stage VI chain holds:

S satisfies U1–U6 ⇒ unique stationary distribution π_S (Perron–Frobenius) ⇒ $\mu_n S \to \pi_S$ exponentially in $L^2(\pi_S)$ at rate $(1 - \varepsilon_{gap}(S))^n$ (Stage VI Theorem 10.3) ⇒ $\chi^2(\mu_n \| \pi_S) \to 0$ at rate $(1 - \varepsilon_{gap}(S))^{2n}$ (Stage VI Theorem 11.1, in the reversible case) ⇒ $K_\infty(S) = L_\Phi(S) \cdot L(S) \cdot A(S)^2 / (A_-(S) \cdot \varepsilon_{gap}(S)) < \infty$ (Stage V continuum Lipschitz bound).

All members of $\mathcal{C}_{K=7}$ produce the same *qualitative* continuum geometry: one persistent coherent transport sector, six trapped incoherent sectors, exponential contraction toward a stationary coherent distribution, finite continuum Lipschitz regularity. Numerical constants vary; structure does not.

Theorem 12.1 — Open Stability of the Universality Class

The K = 7 refinement universality class $\mathcal{C}_{K=7}$ is open in the space of admissibility-restricted row-stochastic 7 × 7 matrices. In particular, there exists an open neighbourhood $\mathcal{U} \subset \mathcal{C}_{K=7}$ containing the canonical $\hat{T}$ such that every $S \in \mathcal{U}$ satisfies U1–U6.

Proof. Each condition U1–U6 is an *open condition* in the matrix entries under wheel-admissible support:

- U1 (stochasticity): defined by linear equality ($S \cdot \mathbf{1} = \mathbf{1}$) and inequality ($S \geq 0$). Within the affine subspace of stochastic matrices, positivity is open.
- U2 (support): satisfied by all matrices supported on the wheel edges, an affine condition.
- U3 (irreducibility): open under fixed positive support, since the communication graph is determined by which edges have non-zero weight, and small perturbations of strictly positive weights remain strictly positive.
- U4 (aperiodicity): open under U3 and the existence of any non-zero diagonal entry, by the same argument as U3.
- U5 (positive spectral gap): the eigenvalues of S are continuous functions of S (Hoffman–Wielandt and Bauer–Fike), so the maximum non-trivial modulus is upper semi-continuous in S. By Theorem 5.1, this modulus stays below 1 in some neighbourhood of $\hat{T}$; ε_{gap} stays above zero.
- U6 (bounded local response): L(S) is continuous in S under wheel-admissible support and bounded propagation (Stage V Lemma 4.x), so finiteness is open.

The intersection of finitely many open conditions is open. Therefore $\mathcal{C}_{K=7}$ contains an open neighbourhood of $\hat{T}$.

Corollary 12.2 — Universality of the Continuum Lipschitz Constant

The continuum Lipschitz constant K_∞ is continuous and bounded on $\mathcal{U}$, hence Lipschitz on any compact subset of $\mathcal{U}$:

$K_\infty(S) = K_\infty(\hat{T}) + \mathcal{O}(\| S - \hat{T} \|)$ as $S \to \hat{T}$.

In particular, every operator in a neighbourhood of $\hat{T}$ produces the same continuum-geometry universality class: Lipschitz Lorentzian continuum with $K_\infty = K_\infty(\hat{T})$ + small correction.

---

13. Limitations and Open Problems

The principal results above establish stability under *small* perturbations. Several extensions remain.

13.1 Explicit spectra of perturbation families. Theorem 5.1 bounds eigenvalue shifts but does not compute them. The numerical examples of §10 cover three specific perturbations; a systematic survey across an explicit basis of admissible ΔT directions — e.g. the 36-dimensional space of rank-1 admissible perturbations (one per allowed wheel edge) — would map the perturbation landscape comprehensively.

13.2 Large perturbations and critical surfaces. Small perturbations stay inside the universality class. Large perturbations may cross a *critical surface* in admissibility-rule space at which $\varepsilon_{gap} \to 0$ (marginal-to-relevant transition). Identifying these critical surfaces — analogous to the Curie-temperature surface in statistical mechanics — would map the boundary of the $K = 7$ universality class explicitly. This is a finite combinatorial computation.

13.3 Position- and level-dependent perturbations. Realistic substrates require $\hat{T}^{\ell, x}$ that varies non-trivially across refinement level ℓ and substrate position x. The Stage V hypothesis S5 in its general form is the *uniform* spectral gap

$$\varepsilon_{gap}^{uniform} := \inf_{\ell, x} \varepsilon_{gap}(\hat{T}^{\ell, x}) > 0.$$

The present paper proves stability at a single (ℓ, x); proving uniform stability over a parametric family requires either compactness of the (ℓ, x) parameter space or an explicit Lipschitz dependence on (ℓ, x). The proof structure of Theorem 12.1 generalises directly once the parameter-family compactness is established.

13.4 Defects, curvature, and emergent matter (conjectural). Localised admissibility-rule defects — a small region of substrate where $\hat{T}^{x_0}$ differs substantially from $\hat{T}^{generic}$ — may not merely *damage* the continuum geometry. The defect may generate *effective curvature* (deviation from flat Lorentzian geometry) or *matter-like excitations* (localised energy density) in the emergent continuum. This is a substantial open problem: the relationship between substrate-level defects and emergent physical content. We *conjecture* — without proof and without present-paper support — that defects in refinement coherence might *generate* physical structure rather than merely degrade geometry, an attractive prospect for substrate-derived matter content. The label "conjectural" is meant strictly: no theorem, lemma, or calculation in the present paper supports this claim; it is recorded here only to motivate the Stage VIII research direction.

13.5 Reversibility-breaking perturbations. Theorem 7.1 assumed reversibility preservation. Generic perturbations break reversibility, and the entropic contraction story then involves singular values rather than eigenvalues. The Stage V continuum-geometry chain depends on the $L^2(\pi_\eta)$ operator-norm contraction rate (not on χ^2-divergence directly), so the principal conclusions survive — but the entropic accounting layer may need a non-reversible substitute. Constructing this substitute is the natural follow-up.

13.6 Wheel uniqueness. This paper proves stability around the wheel W_6. It does *not* derive the wheel as the unique architecture satisfying the closure-Hamiltonian programme's structural requirements. Stage VI §14.1 gave informal arguments for the wheel's minimality among $K = 7$ closure-compatible architectures, but a full uniqueness theorem — "every $K = 7$ substrate satisfying the eight requirements of Stage VI §14.1 is graph-isomorphic to W_6" — remains open. The present universality result is *necessary* for a substrate-engineering programme, but it would be *sufficient* only when paired with such a uniqueness result.

---

14. Conclusion

The principal finding of this paper is not merely that the canonical $K = 7$ wheel survives perturbation. It is that we can identify **two specific structural mechanisms** responsible for the survival — and together they elevate the universality-class theorem from an abstract topological openness statement into a mechanism-grounded claim about why the wheel works.

Mechanism 1 — Symmetry-protected spectral invariance (Theorem 9.1, Corollary 9.2)

A positive-dimensional subspace of admissible perturbations is **exactly** spectrally silent. Not "approximately", not "to leading order in η", but exactly, for every η in the admissible range. The proof is one line of the matrix determinant lemma combined with resolvent commutation: for any reflection $\sigma \in D_6$ fixing the perturbation seat κ, the σ-antisymmetric perturbation directions satisfy $v^T (\hat{T} - \lambda I)^{-1} e_\kappa = 0$ identically, collapsing the perturbative factor in the rank-r determinant identity to 1. The eigenvectors rotate; the eigenvalues do not.

The consequence is that $\varepsilon_{gap} = ½$ is not a *computed* number that happens to be positive. It is a **symmetry-protected** quantity — a structural invariant of the wheel's dihedral D_6 symmetry, immune to a substantial subspace of would-be microscopic disorder, to all orders in η on the pure-form silent subspaces and to first order in η on the broader conjugation-antisymmetric subspace. For each of the six reflections of D_6, two complementary spectrally-silent ansatz subspaces sit inside the 49-dimensional matrix algebra: vertex-vertex reflections give 10-dimensional ansatz subspaces ($\dim(H_{\sigma^+}) \cdot \dim(H_{\sigma^-}) = 5 \cdot 2$); edge-edge reflections give 12-dimensional ansatz subspaces ($4 \cdot 3$). Their direct sum per reflection is the (-1)-eigenspace of conjugation, on which the weaker invariance "characteristic polynomial is even in η" holds — first-order spectral response is zero, with leading shift at $O(\eta^2)$ for mixed elements. The admissibility-preserving intersection with P1 ∩ P2 still leaves a sizeable spectrally-silent class within the pure-form subspaces.

Mechanism 2 — Six-fold operator-norm conservatism on the surviving complement (§5, §10.2)

For genuinely symmetry-breaking perturbations — those outside the Theorem 9.1 silent class — the spectrum does move. But the actual movement is much slower than the worst-case Bauer–Fike bound predicts. Direct measurement on the canonical self-loop / cyclic perturbation gives:

measured Lipschitz slope of $\varepsilon_{gap}(\eta) \approx 0.226$, Bauer–Fike upper bound $C(\Delta T) = \sqrt{2} \approx 1.414$, conservatism factor ≈ 6.25.

The wheel's residual symmetries — the $\mathbb{Z}/6$ boundary rotation, the remaining reflection axes, the hub's degree-6 symmetric coupling — continue to suppress the spectral response even where they do not annihilate it. The operator-norm bound $\sqrt{2}$ captures the worst-case sensitivity an unstructured matrix could produce; the canonical wheel's *actual* sensitivity is six times smaller, because residual structure absorbs most of what symmetry alone cannot eliminate.

Why this makes Theorem 12.1 substantive

The K = 7 refinement universality class $\mathcal{C}_{K=7}$ is an open neighbourhood of $\hat{T}$ in the space of admissibility-restricted stochastic matrices. As a purely *topological* statement, that result follows from the small abstract fact that finitely many open conditions intersect openly — true for any sufficiently regular finite Markov chain. As written it would be a thin claim.

What Mechanisms 1 and 2 add is the **reason the universality class is wide rather than narrow**:

- the symmetry-protected silent subspace is positive-dimensional, so the universality class extends *exactly* in those directions across the full admissible η-range (not merely some small neighbourhood);
- the surviving complement, on the canonical self-loop / cyclic test perturbation, is governed by a Lipschitz constant approximately 6× below the operator-norm bound (§10.2); residual symmetries are expected to give comparable conservatism on related symmetry-breaking directions, so the universality class extends *robustly* in those directions too, with admissible η-ranges substantially larger than worst-case analysis would predict.

The neighbourhood $\mathcal{U}$ of $\hat{T}$ in Theorem 12.1 is therefore not merely "open" in a topological sense. It is **engineered for breadth** by two compounding structural protections — the wheel's dihedral symmetry on one dimensional class, and operator-norm conservatism on the complement.

Survival of the Stage VI chain under perturbation

Each load-bearing feature of the Stage VI canonical construction survives admissible perturbation:

- **irreducibility and aperiodicity** (Propositions 4.1, 4.2);
- **simplicity of the Perron eigenvalue** (Perron–Frobenius, automatic);
- **positivity of the spectral gap** $\varepsilon_{gap}(\eta) \geq \frac{1}{2} - C(\Delta T) \cdot |\eta|$, with measured sensitivity 6× below the bound on the canonical self-loop / cyclic test perturbation (Theorem 5.1, §10.2);
- **exact spectral invariance** under symmetry-aligned perturbations (Theorem 9.1, Corollary 9.2);
- **continuity of the stationary distribution** $\pi_\eta = \pi + \eta \cdot \pi \cdot \Delta T \cdot Z + \mathcal{O}(\eta^2)$ (Theorem 6.1);
- **geometric entropic contraction** $\chi^2(\mu_n \,\|\, \pi_\eta) \leq (\frac{1}{4} + \mathcal{O}(|\eta|))^n \cdot \chi^2(\mu_0 \,\|\, \pi_\eta)$, in the reversible case (Theorem 7.1);
- **continuum Lipschitz regularity** $K_\infty(\eta) = K_\infty(0) + \mathcal{O}(|\eta|)$ (Theorem 8.1).

The full continuum-geometry chain across Stages I–VII reads:

K = 7 wheel architecture (Stage VI) ⇒ $\varepsilon_{gap}(0) = \frac{1}{2}$ (Stage VI canonical computation) ⇒ symmetry-aligned perturbations spectrally invariant exactly (Theorem 9.1, Mechanism 1) ⇒ symmetry-breaking perturbations spectrally Lipschitz with measured conservatism factor ≈ 6 on the canonical self-loop / cyclic perturbation of §10.2 (Theorem 5.1 + §10.2, Mechanism 2; comparable structural protections expected on related symmetry-breaking directions but not systematically verified) ⇒ admissible perturbations preserve all U1–U6 conditions on an open neighbourhood (Theorem 12.1) ⇒ stationary measure shift Lipschitz with explicit formula (Theorem 6.1) ⇒ entropic χ^2-contraction at rate $\frac{1}{4} + \mathcal{O}(|\eta|)$ in reversible case (Theorem 7.1) ⇒

$K\_\infty(\eta) < \infty$ on the open neighbourhood (Theorem 8.1) ⇒ Lipschitz Lorentzian continuum geometry for every $S \in \mathcal{C}\_\{K=7\}$ (Stage V Main Theorem, applied to S).

Each arrow is either proven (Theorems 5.1 through 12.1, Propositions 4.1–4.2) or proven conditional on Stage I–V results already established.

Programme consequence

Emergent Lorentzian geometry in VERSF is not a knife-edge property of one canonical matrix. It is the stable infrared behaviour of an *open universality class* $\mathcal{C}\_\{K=7\}$ of admissible refinement substrates, **and the breadth of this class has named structural mechanisms**: spectral protection by D_6 symmetry on positive-dimensional silent subspaces; effective-Lipschitz conservatism approximately 6× below the operator-norm bound on the canonical test perturbation, with comparable conservatism expected on related symmetry-breaking directions.

This is the structural pattern expected of a credible physical substrate: microscopic detail varies, macroscopic behaviour is universal, and the universality has *named mechanisms* rather than merely an open-set guarantee. The canonical wheel $\hat{T}$ earns its place not because it is uniquely correct, but because it is uniquely *minimal* — the simplest representative of a class whose breadth is structurally protected rather than fortuitous.

The Stage VII universality result, paired with the Stage VI explicit computation, completes the *qualitative* substrate-engineering phase of the VERSF geometry programme. What remains is the *quantitative* phase: mapping the critical surfaces at the boundary of $\mathcal{C}\_\{K=7\}$ where the gap closes; classifying defect-induced effective curvature and matter content; deriving the wheel itself from deeper substrate principles; and connecting the universality class to the broader closure-Hamiltonian and generation-counting programmes. The continuum geometry is now structurally stable on a substantial open neighbourhood, with named mechanisms explaining why that neighbourhood is wide. The next phase is to populate it with physics.