

Strict but Not Uniform: a Machine-Checked Spectral Gap at Every Finite Extent of the Coupled Slice

peripheral separation without an equality case,
and the modulus the theorem does not provide

Lluis Eriksson

Abstract

A companion paper supplied the vacuum of the coupled $\mathbb{Z}_2$ slice at every spatial extent: a strictly positive eigenvector, unique up to scale, carrying the spectral radius [1]. It listed *peripheral separation* as out of scope, and that omission is not cosmetic — without it $|\mu| \leq \lambda$ leaves $\mu = -\lambda$ open, and no gap follows at all.

This paper closes it, and then draws the consequence that matters, which is negative.

Proved. For a strictly positive kernel on a finite nonempty type, $-\lambda$ is not an eigenvalue, hence every real eigenvalue other than the Perron eigenvalue is *strictly* smaller in absolute value. Specialised to the coupled slice: at every extent L , every β , and every strictly positive source weight, the transfer operator has a strict spectral gap. The proof of peripheral separation avoids the equality case of the triangle inequality, which is where the classical argument spends its effort: writing $u = |w|$, $p = u - w$ and $q = u + w$, one gets $Ap = \lambda q$ and $Aq = \lambda p$, so a nonzero p would make Ap strictly positive, hence q strictly positive, hence w nonnegative, hence $p = 0$.

The separation is then extended from the real eigenvalues to *all* of them. The coupled kernel is conjugate by a positive diagonal to its symmetrised form, which is symmetric; and a real symmetric kernel has real eigenvalues, by a computation that pairs the eigenvector against its image twice and is two rearrangements of a double sum. So there are no complex peripheral eigenvalues left to exclude, and the strict gap is a statement about the whole spectrum. We also deliver the vacuum in Euclidean normalisation, $\|\Omega\| = 1$ with $T\Omega = \Omega$.

Not proved, and this is the title. The gap is *strict*, not *quantitative*: the theorem provides no modulus of separation, and in particular nothing uniform in L . That is not a limitation we expect to remove by working harder in this direction. Direct numerical diagonalisation shows the subdominant ratio running 0.9205, 0.9829, 0.9964, 0.9992 at $L = 2, 3, 4, 5$ for $\beta = 0.8$, $\gamma = 1.2$ — collapsing towards 1. That computation is reported as measured and unproved, and no theorem here depends on it; its role is to say that a paper reporting only the positive half would be reporting the half that does not matter.

Nothing in this paper is a claim about $SU(N)$, the continuum limit, or the Yang–Mills mass gap.

1 Introduction

1.1 What was missing, and why it was not cosmetic

[1] proved that the coupled slice has a vacuum at every finite extent — strictly positive, unique up to scale, carrying the spectral radius — and listed among its exclusions that *peripheral*

separation was not proved. Read quickly, that looks like a technical refinement. It is not. The domination statement is $|\mu| \leq \lambda$; the case $\mu = -\lambda$ saturates it. Until that case is excluded, the second-largest modulus may equal the largest, and *no gap statement exists at all*, at any extent.

So the first half of this paper is not an improvement of the previous one. It is the step without which the previous one supports no spectral conclusion.

1.2 What is proved

Throughout A is a strictly positive kernel on a finite nonempty type, $v > 0$ with $Av = \lambda v$.

- $-\lambda$ is not an eigenvalue of A ;
- hence for every real eigenpair (μ, w) with $w \neq 0$ and $\mu \neq \lambda$, $|\mu| < \lambda$ — a strict gap;
- specialised: at every extent L , every β , and every strictly positive source weight, the coupled transfer operator has a strict gap;
- the vacuum in Euclidean normalisation: $\|\Omega\| = 1$ and $T\Omega = \Omega$;
- the symmetrised kernel is symmetric, and strictly positive;
- *every* eigenvalue of the coupled kernel is real — the kernel is conjugate by a positive diagonal to the symmetrised one — so the strict separation above is a statement about the whole spectrum, not only about its real part.

1.3 What is deliberately not claimed

1. **No modulus of separation, and nothing uniform in L .** The gap theorem is a strict inequality with no quantitative content. This is the paper’s own headline and Section 7 says what the numbers show.
2. **No proof that a volume-uniform gap fails.** The collapse in Section 7 is measured on a finite grid, is labelled as such, and carries no theorem. What *is* established is that this paper does not provide a uniform bound — not that none exists.
3. **Algebraic simplicity** is not proved, and neither is Perron–Frobenius for merely irreducible kernels.
4. **Nothing about $SU(N)$, the continuum limit, or the Clay problem [7].**

2 The positive left eigenvector, and what it buys

Lemma 2.1 (Left eigenvector). $A^\top$ is strictly positive, so it too has a strictly positive eigenvector φ , and its eigenvalue equals λ .

[exists_pos_left_eigenvector](#), [left_eigenvalue_eq](#). Equality of the two eigenvalues follows from pairing: $\langle \varphi, Av \rangle$ evaluates to $\lambda \langle \varphi, v \rangle$ one way and $\nu \langle \varphi, v \rangle$ the other, and $\langle \varphi, v \rangle > 0$.

Lemma 2.2 (A nonnegative sub-eigenvector is an eigenvector). If $u \geq 0$ and $Au \geq \lambda u$ then $Au = \lambda u$.

Proof. Let $z = Au - \lambda u \geq 0$. Then $\langle \varphi, z \rangle = \langle A^\top \varphi, u \rangle - \lambda \langle \varphi, u \rangle = 0$, and a nonnegative vector with zero positive-weighted sum vanishes. $\square$

[subeigen_eq](#). This is the standard device, and it is what makes the next section cheap.

3 Peripheral separation, without an equality case

Theorem 3.1 ($-\lambda$ is not an eigenvalue). *There is no $w \neq 0$ with $Aw = -\lambda w$.*

Proof. Let $u = |w|$ entrywise. From $\lambda u_i = |(Aw)_i| \leq (Au)_i$ we get $Au \geq \lambda u$, so $Au = \lambda u$ by Lemma 2.2. Put $p = u - w \geq 0$ and $q = u + w \geq 0$. Then

$$Ap = Au - Aw = \lambda u + \lambda w = \lambda q, \quad Aq = Au + Aw = \lambda u - \lambda w = \lambda p.$$

Suppose $p \neq 0$. Then $Ap > 0$ entrywise, so $\lambda q > 0$, so $q_i > 0$ for every i ; but $q_i = |w_i| + w_i$, which vanishes wherever $w_i < 0$, so $w \geq 0$ — and then $u = w$, i.e. $p = 0$, a contradiction. Hence $p = 0$, so $w = u \geq 0$, so $Aw = \lambda w$; with $Aw = -\lambda w$ and $\lambda > 0$ this forces $w = 0$. $\square$

[neg_not_eigenvalue](#).

Remark 3.2 (Why this route was taken). *The classical proof of peripheral separation goes through the equality case of the triangle inequality: from $|\sum_j A_{ij} w_j| = \sum_j A_{ij} |w_j|$ one concludes that the w_j share a phase. Formalising an equality case is markedly more work than formalising an inequality, and it is avoidable here. The pair (p, q) carries the same information in a form that only ever needs strict positivity of the image of a nonnegative nonzero vector — a lemma the companion development already had.*

4 The gap

Theorem 4.1 (Strict gap). *If $Aw = \mu w$ with $w \neq 0$ real and $\mu \neq \lambda$, then $|\mu| < \lambda$.*

Proof. $|\mu| \leq \lambda$ by the domination theorem of [2]. If $|\mu| = \lambda$ then $\mu = \lambda$, excluded by hypothesis, or $\mu = -\lambda$, excluded by Theorem 3.1. $\square$

[abs_lt_of_ne_perron](#), and for the coupled slice at every extent [coupled_gap_of_sourceWeight](#).

Remark 4.2 (What the statement does and does not contain). *Theorem 4.1 is a strict inequality between two real numbers. It contains no ε , no dependence on L , and therefore no information whatever about how the separation behaves as the extent grows. Every use of it that sounds quantitative is a misuse.*

5 The two interfaces the next module needs

Theorem 5.1 (Euclidean-normalised vacuum). *With $T = \lambda^{-1}A$ there is $\Omega > 0$ with $\sum_i \Omega_i^2 = 1$ and $T\Omega = \Omega$.*

[eucNorm](#), [unitVacuum](#), [unitVacuum_norm](#), [unitVacuum_fixed](#). The companion paper normalised on the simplex, which is what its existence proof produced; the Osterwalder–Seiler side asks for unit Euclidean norm, and the passage between them is a rescaling.

Theorem 5.2 (The symmetrised kernel is symmetric). $\tilde{K}_w(\sigma, \tau) = w(\sigma)^{1/2} K(\sigma, \tau) w(\tau)^{1/2}$ satisfies $\tilde{K}_w(\sigma, \tau) = \tilde{K}_w(\tau, \sigma)$, and is strictly positive when w is.

`symWeighted_symm`, `symWeighted_pos`. It is symmetric because the decoupled kernel is.

6 Every eigenvalue is real, so the gap is about the whole spectrum

Theorem 4.1 separates the *real* eigenvalues. For a real non-symmetric kernel that is not yet a statement about the spectrum: nothing so far excludes a peripheral pair $\mu = \lambda e^{\pm i\theta}$. The symmetry just established closes exactly that gap in the argument.

Theorem 6.1 (A real symmetric kernel has real eigenvalues). *If S is real and symmetric and $Sy = \mu y$ with $y \neq 0$ complex, then $\bar{\mu} = \mu$.*

Proof. Pair y against Sy : the result is μN with $N = \sum_i |y_i|^2 > 0$. Conjugating and using that S is real and symmetric returns the same quantity, so $\bar{\mu}N = \mu N$. $\square$

`eigenvalue_real_of_symm`. No spectral theory is invoked; the computation is two rearrangements of a double sum.

Theorem 6.2 (Every eigenvalue of the coupled kernel is real). *K_w is conjugate to $\tilde{K}_w$ by the positive diagonal $\text{diag}(w^{1/2})$, so every complex eigenvalue of K_w is an eigenvalue of $\tilde{K}_w$, hence real.*

`sourceWeighted_eq_sym`, `coupled_eigenvalue_real`.

Corollary 6.3. *For the coupled slice, Theorem 4.1 separates every eigenvalue from λ , not merely the real ones.*

`coupled_gap_all_eigenvalues`.

Remark 6.4 (What this does not turn into). *Making the separation a statement about the whole spectrum changes nothing about its quantitative content, which remains empty. Self-adjointness on an inner-product space, and the variational characterisation that would supply a modulus, still belong to the module that consumes this one.*

7 What is measured and not proved

Direct dense numerical diagonalisation of the $2^L \times 2^L$ coupled kernel gives, for $\beta = 0.8$ and $\gamma = 1.2$, subdominant ratios

$$0.9205, \quad 0.9829, \quad 0.9964, \quad 0.9992 \quad \text{at } L = 2, 3, 4, 5,$$

and for $\beta = 0.3$, $\gamma = 0.4$,

$$0.4665, \quad 0.5518, \quad 0.5938, \quad 0.6159.$$

The first family collapses towards 1; the second does not, and the parameters of the second satisfy $\sinh 2\beta \sinh 2\gamma < 1$. All of this is **verified** in the numerical sense and **not proved**; the identification of that threshold with the Kramers–Wannier line of the two-dimensional Ising

transfer matrix is standard literature and is not established here. The probe and its transcript are in the repository.

No theorem of this paper depends on this section. Its function is to prevent a misreading: Theorem 4.1 holds at every finite extent, and the numbers show that this is compatible with the separation vanishing in the limit. A strict gap at each L and a uniform gap in L are different statements, and only the first is proved here.

8 Formalisation notes

An inequality instead of an equality case. The one design decision that shaped the module is described in the remark after Theorem 3.1: the (p, q) pair replaces the equality case of the triangle inequality by two applications of a positivity lemma that already existed. The saving is not marginal — an equality case would have required either complex-phase reasoning or a sign decomposition of the index set.

The left eigenvector was free. Existence for $A^\top$ is the companion theorem applied to a transposed kernel, and equality of the two eigenvalues is one pairing. That single object converts the sub-eigenvector inequality into an equality, which is the hinge of the whole argument.

Normalisation conventions are interfaces, not mathematics. The companion paper produced $\sum \Omega = 1$ because its existence proof lives on the simplex; this one adds $\|\Omega\| = 1$ because that is what the consuming side asks. Both are one rescaling from each other, and neither is a theorem about the model. They are recorded as separate endpoints so that no later module has to guess which convention is in force.

9 Reproducibility

All artifacts are at source checkpoint `ac8979631c1541b059b3c10f590ca174b0b5f6be` on branch `main` of <https://github.com/lluiseiriksson/THE-ERIKSSON-PROGRAMME>. Every link was resolved against that commit, and checked to point at an actual declaration line, before compilation.

Quantity	Value
Full core build	8428 jobs, success
Oracle commands	2583
with axiom dependencies	2560
axiom-free	23
Distinct axioms across all outputs	<code>propext</code> , <code>Quot.sound</code> , <code>Classical.choice</code>
Occurrences of <code>sorryAx</code>	0
Project axioms	0
Errors	0

The numerical probe is [scripts/probe_spatial_birkhoff.py](#) with transcript [docs/O-LANE-PROBE-TRANSCRIPT-20260728.txt](#). Verdicts are in [docs/VERIFICATION-LEDGER.md](#).

References

- [1] L. Eriksson, *The Vacuum Was Never Absent: A Machine-Checked Perron Theorem for Strictly Positive Kernels, and the Coupled-Slice Vacuum at Every Spatial Extent*, viXra, 2026.
- [2] L. Eriksson, *Blind to the Coupling: a Second Machine-Checked Obstruction at Spatial Extent*, viXra, 2026.
- [3] L. Eriksson, *Where the Elementary Reconstruction Stops: Spatial Coupling Breaks the Uniform Vacuum, Machine-Checked*, viXra, 2026.
- [4] E. Seneta, *Non-negative Matrices and Markov Chains*, 2nd ed., Springer, 2006.
- [5] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, CADE-28, Lecture Notes in Computer Science **12699** (2021), 625–635.
- [6] The mathlib Community, *The Lean mathematical library*, CPP 2020, 367–381.
- [7] A. Jaffe and E. Witten, *Quantum Yang–Mills theory*, Clay Mathematics Institute Millennium Prize Problem description, 2000.