

The Weight That Could Not Break It:

Machine-Checked *Endpoint* Reflection Positivity

for the Coupled $\mathbb{Z}_2$ Slice

real single-slice observables — two reflections, two hypotheses, and one of them sharp

Lluis Eriksson

Abstract

What is proved, stated before anything else. The reflected two-point form of the Gibbs measure, at the two *ends* of a path, for *real* observables of a *single* slice, is non-negative under the parity and coupling hypotheses stated below — at even separation for every β , and at every separation exactly for $\beta \geq 0$ — and the Gram matrix of a finite family of such observables is positive semidefinite under the same hypotheses. **This is not yet the Osterwalder–Schrader axiom**, which quantifies over observables of the whole past half-chain and over complex ones. The half-chain algebra, the reflection map and the sesquilinear form are not built here; §7 says what is missing and why it is a construction rather than a further inequality.

Eleven papers in this lane end with *reflection positivity is untouched*. This touches the endpoint form of it, and the interesting part is what the spatial source weight does — namely nothing. Earlier papers do prove statements about the coupled kernel; $\text{specGap} < \lambda$ holds there too [3]. What is new is that this result is *unchanged* by the weight: same hypothesis, same conclusion, and no constant that depends on w . Every earlier coupled-kernel statement contains a number that moves when w does. Positivity is the *sign* of a quadratic form, and conjugation by $\sqrt{w}$ is a congruence, so there is nothing for the weight to move.

Two reflections, two hypotheses. The reflected two-point sum is $\langle v, K^N v \rangle$ with v the dressed observable, and the two cases have genuinely different content. Through a *site* (N even) the sum is a square, hence non-negative for *every* β — negative coupling included — with nothing used but symmetry of the kernel. Through a *bond* (N odd) it is $\langle K^m v, K K^m v \rangle$, which needs K itself positive semidefinite; that holds exactly for $\beta \geq 0$, by an induction on the extent.

And the second hypothesis is active. At $\beta < 0$ and odd separation the single-site sign observable gives the reflected sum in closed form, $2(e^\beta - e^{-\beta})^N$, which is negative. Congruence is invertible, so the same witness divided by $\sqrt{w}$ shows the *coupled* kernel is indefinite below zero for *every* strictly positive weight, at every extent with at least one site. At $L = 0$ there is nothing to witness — one configuration, and there the decoupled kernel is 1 while the coupled one is the positive scalar $w(\emptyset)$ — and that exception is recorded rather than left to be found. The boundary $\beta = 0$ is sharp and exhibited, not inferred.

What this is not. Neither the full axiom (above) nor any reconstruction: the physical Hilbert space as the quotient of the past algebra by the null space of this form is not built. And the degeneration of $\text{specRatio}(L)$ under a ring weight is *measured*, not proved; no theorem here or elsewhere in the lane says the weight destroys uniformity, only that no uniform bound is proved for it. Nothing here concerns uniformity in the extent, $\text{SU}(N)$, the continuum limit, or the Yang–Mills mass gap [8].

On the pre-registration. Three *active* gates were committed before any Lean was written, and §6 reports what happened to each. A fourth is preserved there and not counted:

the original Gate B, which failed because of a design error of ours, and which two of the three active gates replaced.

1 Why this one survives the weight

The lane's obstruction has a shape. Every rate statement it has produced is a statement about `specGap` or `specRatio`, and those are spectral data: conjugating the kernel by $\sqrt{w}$ moves eigenvalues around, and the *measured* evidence is that with a ring weight it moves them to degeneracy outside the disordered region. (Measured: no theorem in this lane proves that, and none is claimed here.)

Positivity is not spectral data in that sense. It is the *sign* of a quadratic form, and

$$\langle u, \sqrt{w} T \sqrt{w} u \rangle = \langle \sqrt{w} u, T \sqrt{w} u \rangle$$

is the same form read on a relabelled vector. Sylvester's law is the general statement; here it is one line, and it is the whole reason the coupled case costs nothing once the decoupled one is done.

2 Through a site: free, at every coupling

Theorem 2.1 (site reflection). *Let K be any symmetric kernel. For every m and every observable u ,*

$$\sum_i (K^{2m} u)(i) u(i) \geq 0.$$

`sum_iterate_even_nonneg`. Move half the iterates across the pairing — which is what symmetry buys — and the sum becomes $\sum_i (K^m u)(i)^2$. No positivity of K is used, and no hypothesis on β : this holds at negative coupling, where the bond reflection below fails.

3 Through a bond: exactly $\beta \geq 0$

Theorem 3.1 (the decoupled kernel is positive semidefinite). *For $\beta \geq 0$, every extent L and every observable u , $\sum_{\sigma} u(\sigma) (Ku)(\sigma) \geq 0$.*

`spatialKernel_posSemidef`. By induction on the extent, splitting an observable at the first site into its even and odd parts. Each part inherits positivity from the inductive hypothesis, and the two are recombined with the two bond eigenvalues,

$$Z = e^{\beta} + e^{-\beta} > 0, \quad D = e^{\beta} - e^{-\beta} \geq 0 \iff \beta \geq 0,$$

which is precisely where the hypothesis enters. It is the same induction the companion uniformity paper runs [4], with different bookkeeping; the split lemma is shared and stated once (`act_spatialKernel_cons`).

Theorem 3.2 (congruence, and the coupled kernel). *If K is positive semidefinite then so is $c K c$ for any function c . Hence for $\beta \geq 0$ and any strictly positive source weight w , the symmetrised kernel $\sqrt{w} K \sqrt{w}$ is positive semidefinite.*

`posSemidef_congr` and `symWeighted_posSemidef`. This is the step §1 advertised, and it costs one reindexing.

Theorem 3.3 (bond reflection). *For $\beta \geq 0$, every m and every observable, $\sum_i (K^{2m+1}u)(i) u(i) \geq 0$.*

`sum_iterate_odd_nonneg`.

4 The same statements, about the measure

The bridge of the companion paper [2] identifies the reflected two-point sum of the Gibbs measure with $\langle v, K^N v \rangle$, so both theorems transfer verbatim.

Theorem 4.1 (endpoint reflection positivity in the measure). *Let A be a real observable of a single slice, evaluated at the two ends of the path. For any strictly positive source weight, at every extent:*

1. *at even separation the reflected two-point sum is non-negative for every β (`gibbsPathSum_nonneg_even`);*
2. *for $\beta \geq 0$ it is non-negative at every separation (`gibbsPathSum_nonneg`).*

The statement is about the *unnormalised* sum, which is the object reflection positivity is about; the partition function is positive, so the normalised expectation has the same sign, but that is a remark and not the theorem.

The whole matrix, not one diagonal entry. Reflection positivity is a statement about a Gram matrix being positive semidefinite. For these observables the matrix form follows from the diagonal one by bilinearity, and is stated so the two are not confused:

Theorem 4.2 (the Gram matrix). *For every finite family $A_1, \dots, A_k$ of endpoint observables and every real $c_1, \dots, c_k$,*

$$\sum_{i,j} c_i c_j \text{gibbsPathSum}(w, \beta, N, A_i, A_j) = \sum_{i,j} c_i c_j \langle \sqrt{w} A_i, K_w^N \sqrt{w} A_j \rangle \geq 0$$

under the hypotheses of Theorem 4.1, where $K_w = \sqrt{w} K \sqrt{w}$.

The factors $\sqrt{w}$ are not decoration. They are the dressing the bridge module produces [2], and they are what makes the operator self-adjoint; the same matrix written with undressed A_i and the bare K would be a true inequality about a *different* object.

`gibbsPathSum_gram_nonneg` and `gibbsPathSum_gram_nonneg_even`, with the bilinearity that makes the matrix visible at `gibbsPathSum_combo`.

5 Sharpness: the hypothesis is active

A hypothesis that is never in force is not worth stating.

Theorem 5.1 (the witness, in closed form). *At one site and constant weight, the single-site sign observable gives*

$$\sum_X A(X_0)A(X_N) \text{weight}(X) = 2(e^\beta - e^{-\beta})^N,$$

exactly. For $\beta < 0$ and N odd this is negative.

[gibbsPathSum_siteSign_eq](#) and [gibbsPathSum_neg_of_neg_beta](#). The observable is the one the extent paper already built [1]; it is an eigenvector of the decoupled kernel with eigenvalue D , so the iterate is D^N times it, and the norm counts the two configurations.

And for every weight, not only the constant one. That witness sits at constant weight, so on its own it refutes only the universally quantified statement. Congruence is *invertible* — $\sqrt{w}$ is nowhere zero — so inertia transports in both directions, and the same observable divided by $\sqrt{w}$ witnesses indefiniteness of the coupled kernel itself.

Theorem 5.2 (sharp for every positive weight). *For every extent with at least one site, every strictly positive w and every $\beta < 0$, the quadratic form of $\sqrt{w} K \sqrt{w}$ takes the value $(e^\beta - e^{-\beta}) Z^{L-1} 2^L < 0$.*

[symWeighted_quadForm_neg_of_neg_beta](#).

The exception, recorded — and it is two kernels, not one. At $L = 0$ the configuration space is a point, and there are two objects to keep apart. The *decoupled* kernel is the number 1 ([spatialKernel_posSemidef_zero](#)); the *weighted* kernel is the positive scalar $w(\emptyset)$, which need not be 1 ([symWeighted_posSemidef_zero](#)). Both quadratic forms are squares at every β — u^2 and $w(\emptyset)u^2$ — so Theorem 5.2 has nothing to witness there, and “exactly $\beta \geq 0$ ” is a statement from one site upwards.

6 The three gates, and the one that failed

The previous campaign in this lane committed a judge whose single joint gate read FAIL and was overridden; that is recorded as a process violation in the ledger and is not undoable [4]. The remedy applied here was structural: separate gates, each licensing one theorem.

- **Gate A** (site reflection, every β): PASS.
- **Gate B** (bond reflection): FAIL — and the reason is that we made the same design error one level down. Gate B bundled the $\beta \geq 0$ theorem with its sharpness witness behind “pass iff both”. The first half passed everywhere; the second missed one cell.

[scripts/judge_spatial_reflection.py](#). The autopsy measured why: at that cell the violating direction carries $|D|^N/Z^N \approx 10^{-3}$ of the weight, so a uniformly random observable lands in the violating cone with probability $\approx 2 \times 10^{-2}$, and sixty draws miss it about a third of the time. **An existence claim was being tested by sampling.** [scripts/judge_spatial_reflection_autopsy.py](#).

Gate B stays failed. It was replaced by two new gates, committed before being run and declared as a redesign rather than a reinterpretation: B1 keeps the $\beta \geq 0$ test, and B2 stops

sampling and predicts a *number* — the closed form of Theorem 5.1, to 10^{-12} — which is strictly harder than asking some draw to go negative. Both pass. [scripts/judge_spatial_reflection_v2.py](#).

7 Scope

An endpoint form, not the full axiom and not a reconstruction. Reflection positivity is one of the Osterwalder–Schrader axioms, and what is proved here is its endpoint form rather than the axiom itself. Building the physical Hilbert space — the quotient of the past algebra by the null space of this form, with the transfer operator descending to a self-adjoint contraction — is a further step again, and is neither done here nor claimed.

Endpoint observables only — and what the full axiom would need. The bridge module supplies correlations of observables at the two ends of the path. The Osterwalder–Schrader axiom is about a general observable F of the past half-chain, complex-valued, with the matrix $(\langle \Theta F_i, F_j \rangle)_{ij}$ positive semidefinite. Here the observables are real and single-slice, and no reflection map Θ is defined.

The missing step is identifiable and is a *construction*, not a further inequality: summing out the interior of each half collapses F to a boundary vector G_F , after which the site reflection reduces to $\sum_{\sigma} G_F(\sigma)^2 \geq 0$ and the bond reflection to $\langle G_F, K_w G_F \rangle \geq 0$ — which are Theorem 2.1 and Theorem 3.2, already proved. Extending from real to complex observables is then immediate for a real symmetric kernel, since $z^* K z = x^{\top} K x + y^{\top} K y$ with the cross terms cancelling by symmetry. Both layers are registered as the next step and neither is claimed here.

Nothing about uniformity. This paper says nothing about how any rate behaves as the extent grows. The coupled kernel’s uniformity is still the lane’s open analytic question.

No claim about the physical target. Nothing here concerns $SU(N)$, the continuum limit, or the Clay problem [8]. The distance to that problem is unchanged.

8 Reproducibility

All artifacts are at source checkpoint `854ff223396630a8a75b0a95a3ada6413f3a395d` on branch `main`. Every hyperlink was resolved against that commit, and checked to point at an actual declaration line, before compilation. Every counter below was measured in a window with no other elaborator running on the host, on commit `3ee8b1c7`, whose Lean sources, `oracle_check.lean`, `lakefile.lean` and `lake-manifest.json` are byte-identical to this checkpoint — the only change between them is the ledger entry recording the measurement, and `git diff 3ee8b1c7 854ff223396630a8a75b0a95a3ada6413f3a395d -- '*.lean' oracle_check.lean` is empty.

Between the previous version of this paper and this one, an unrelated branch was merged into `main` by a concurrent session. It adds no declaration to this module and changes none of its proofs, but it does enlarge the *repository-wide* oracle, which is why the totals below are larger than those printed in the previous version. The module’s own row is unchanged.

Quantity	Value
Full core build	8463 jobs, success
Oracle commands, and answers	2752
distinct declarations	2749
commands listed twice	3
distinct, with axiom dependencies	2726
distinct, axiom-free	23
Declarations in this module	19
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 module is [YangMills/OS/SpatialReflection.lean](#); verdicts are in [docs/VERIFICATION-LEDGER.md](#).

References

- [1] L. Eriksson, *Where the Elementary Reconstruction Stops: a Machine-Checked Obstruction at the Coupled $\mathbb{Z}_2$ Slice*, viXra:2607.0075, 2026.
- [2] L. Eriksson, *The Measure the Spectral Results Were About*, viXra, 2026.
- [3] L. Eriksson, *The Modulus: a Machine-Checked Operator Bound on the Fluctuation Sector of the Coupled $\mathbb{Z}_2$ Slice*, viXra, 2026.
- [4] L. Eriksson, *The Rate Without the Extent: a Machine-Checked Extent-Independent Spectral Rate for the Decoupled $\mathbb{Z}_2$ Slice*, viXra, 2026.
- [5] K. Osterwalder and R. Schrader, *Axioms for Euclidean Green’s functions*, Comm. Math. Phys. **31** (1973), 83–112.
- [6] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Ann. Physics **110** (1978), 440–471.
- [7] J. Glimm and A. Jaffe, *Quantum Physics: a Functional Integral Point of View*, 2nd ed., Springer, 1987.
- [8] A. Jaffe and E. Witten, *Quantum Yang–Mills theory*, Clay Mathematics Institute Millennium Prize problem description, 2000.