

Where the Elementary Reconstruction Stops: Spatial Coupling Breaks the Uniform Vacuum, Machine-Checked

A negative result, and the positive one that isolates it

Lluis Eriksson

Abstract

Two companion developments verified an Osterwalder–Seiler reconstruction end to end for a lattice gauge chain whose spatial slice is a single point. Every step of that chain begins by knowing the vacuum, and in the one-dimensional case the vacuum is free: the normalised transfer kernel has constant row sums, so the uniform vector is fixed, and $T\Omega = \Omega$ follows from normalisation alone. This paper asks what survives when the slice acquires spatial extent, and answers in Lean 4 with `mathlib`.

The algebraic half survives untouched. With time bonds only, the row sums of the transfer kernel are constant for every spatial extent L , so the uniform vacuum persists on a space of dimension 2^L ; and the single-site sign observable is an eigenvector whose normalised eigenvalue is exactly $\tanh \beta$, with L free in the statement.

The uniform vacuum does not survive. Switching on a coupling between sites inside a slice makes the spatial weight depend only on the source configuration, so it factors out of the sum over the target and the row sums become configuration-dependent. We exhibit two explicit configurations of a two-site slice with different row sums, and conclude that no constant row sum exists: the uniform vector is not fixed, so $T\Omega = \Omega$ is *false* for it. The vacuum becomes a Perron vector that row-sum normalisation no longer supplies in closed form, and every later step of the reconstruction loses its starting point.

We state plainly what the positive half is and is not. The decoupled system is L non-interacting copies of a two-state system, and the rate it yields — the eigenvalue $\tanh \beta$ of the single-site sign mode — is independent of L for trivial reasons, so it is physically empty and is recorded only because it isolates which half of the construction survives. **No gap for the coupled system is proved here, and none is claimed.** Nothing in this paper is a claim about $SU(N)$, the continuum limit, or the Yang–Mills mass gap.

1 Introduction

1.1 A prediction, and a test of it

The companion papers [3, 4] verified, for the $\mathbb{Z}_2$ chain, a complete passage from a Gibbs weight to exponential decay of a correlation function: reflection positivity, a Gelfand–Naimark–Segal space, a transfer operator, an identification theorem, a spectral bound, and an exact two-point function. Both papers state in print that the next step is spatial extent, and both predict that it is where the construction stops being easy.

This paper tests that prediction rather than repeating it. The result is that the prediction was right, and that it is possible to say exactly *which* step fails and why.

1.2 Why the vacuum is the load-bearing step

Every theorem of the one-dimensional chain begins by knowing the vacuum. The identification is a statement about $\langle A\Omega, T^n B\Omega \rangle$; the spectral bound is a bound on $T - |\Omega\rangle\langle\Omega|$; the endpoint is an expectation evaluated against Ω . And in one dimension the vacuum is free: the normalised kernel has constant row sums, so the uniform vector is fixed, and $T\Omega = \Omega$ is a consequence of normalisation — the companion paper makes a point of the fact that it is a consequence and not a hypothesis.

If that step fails, nothing downstream has an input. That is what makes it the right thing to test first, and it is what this paper tests.

1.3 What is proved

Slices carry L spins, so configurations are maps $\mathbb{Z}_2^L$ and the transfer operator acts on a space of dimension 2^L . Write k_β for the Boltzmann bond weight and $Z_\beta = e^\beta + e^{-\beta}$.

Decoupled — kernel $K(\sigma, \tau) = \prod_j k_\beta(\sigma_j, \tau_j)$:

- $\sum_\tau K(\sigma, \tau) = Z_\beta^L$, independent of σ , for arbitrary L ;
- the single-site sign observable is an eigenvector, with normalised eigenvalue exactly $\tanh \beta$ — the same number as in one dimension, with L free.

Coupled — the same, times a weight $w_\gamma(\sigma)$ linking sites inside the slice:

- $\sum_\tau K_\gamma(\sigma, \tau) = w_\gamma(\sigma) Z_\beta^L$;
- two explicit configurations of a two-site slice have different row sums at every $\gamma > 0$;
- hence no constant row sum exists, and $T\Omega = \Omega$ fails.

1.4 What is deliberately not claimed

This subsection is not boilerplate; the first item is the one that matters, and it was fixed in the governing charter before any of this was written.

1. **The volume-independent rate is physically empty.** The decoupled system is L copies of a two-state system with no interaction between them, and the eigenvalue of the single-site sign mode does not depend on L . We do not call this a volume-uniform gap, we do not headline it, and it is *not* evidence about the coupled case. It is recorded because it isolates which half of the construction survives.
2. **No gap for the coupled system.** None is proved and none is claimed. That is the open problem.
3. **Still $\mathbb{Z}_2$,** one variable per site, fixed finite size, and the spatial coupling is exhibited on a two-site slice.
4. **Nothing about $SU(N)$, the continuum limit, or the Clay problem [7].** No reduction is claimed and none is known to the author.

2 The factorisation that carries both halves

Lemma 2.1. For $f : \{1, \dots, L\} \times \mathbb{Z}_2 \rightarrow \mathbb{R}$, $\sum_{\tau} \prod_j f_j(\tau_j) = \prod_j \sum_t f_j(t)$.

[sum_prod_config](#). Both the positive and the negative half below are corollaries of this one identity applied to different f .

3 The decoupled system: the algebraic half ports

Theorem 3.1 (Constant row sums at every spatial extent). $\sum_{\tau} K(\sigma, \tau) = Z_{\beta}^L$ for every σ and every L .

Proof. Lemma 2.1 with $f_j(t) = k_{\beta}(\sigma_j, t)$; each factor is $\sum_t k_{\beta}(s, t) = Z_{\beta}$ whatever s is. $\square$

[sum_spatialKernel](#), and as an existence statement [decoupled_constant_rowSum](#). This is what carries the uniform vacuum to arbitrary spatial extent.

Theorem 3.2 (A mode whose rate does not depend on the volume). Let $s_i(\sigma)$ be the sign of site i . Then $\sum_{\tau} K(\sigma, \tau) s_i(\tau) = (e^{\beta} - e^{-\beta}) Z_{\beta}^{L-1} s_i(\sigma)$, so after dividing by the row sum the eigenvalue is $(e^{\beta} - e^{-\beta}) / Z_{\beta} = \tanh \beta$.

[spatialKernel_siteSign](#) and [spatial_rate_eq_tanh](#). In the formal statement L is a free variable, so the independence is visible in the statement rather than inferred from a family of instances.

Remark 3.3 (What Theorem 3.2 is worth). Almost nothing on its own, and it is important to say so. The decoupled system is a tensor power; its spectrum consists of products of one-site eigenvalues, so the non-vacuum spectral radius is $|\tanh \beta|$ at every L , and the mode of Theorem 3.2 is volume-independent for the same reason that $1^{L-1} \lambda = \lambda$. Two boundaries belong here. First, the superlative: for $\beta \geq 0$ one may say instead that $\tanh \beta$ is the largest non-vacuum eigenvalue, but for $\beta < 0$ that is false, since $(\tanh \beta)^2 > \tanh \beta$, and nothing in this paper needs it — which is why the statements above name the mode rather than a supremum. Second, the spectral description of a tensor power is standard, and it is prose here; the machine-checked statement is Theorem 3.2, and it holds for every real β . A reader who takes this as progress towards a volume-uniform statement about an interacting theory has been misled, and we would rather forestall that here than be quoted out of context later. The value of the theorem is comparative: it fixes what the coupled case has to be measured against.

4 The coupled system: the vacuum breaks

Definition 4.1. $K_{\gamma}(\sigma, \tau) = w_{\gamma}(\sigma) K(\sigma, \tau)$, where w_{γ} is a bond weight between two sites of the slice.

The essential feature is that w_{γ} depends on the *source* only: [spatialWeight](#), [coupledKernel](#).

Theorem 4.2 (The row sum becomes configuration-dependent). $\sum_{\tau} K_{\gamma}(\sigma, \tau) = w_{\gamma}(\sigma) Z_{\beta}^L$.

`sum_coupledKernel`. The spatial weight is constant in τ , so it leaves the sum untouched and survives into the answer.

Theorem 4.3 (The obstruction). *For every $\gamma > 0$ there are two configurations of a two-site slice with different row sums. Consequently there is no constant c with $\sum_{\tau} K_{\gamma}(\sigma, \tau) = c$ for all σ , and the uniform vector is not fixed: $T\Omega = \Omega$ is false.*

Proof. Take σ constant and σ' alternating. By Theorem 4.2 the row sums are $e^{\gamma} Z_{\beta}^2$ and $e^{-\gamma} Z_{\beta}^2$, and $e^{-\gamma} < e^{\gamma}$ for $\gamma > 0$. $\square$

Artifacts: `cfgFlat`, `cfgAlt`, `coupled_rowSums_not_constant`, `coupled_no_constant_rowSum`.

Remark 4.4 (Why this ends the elementary route). *In one dimension $T\Omega = \Omega$ was free. It is not a small convenience: the identification, the spectral bound and the endpoint of the companion papers all take the vacuum as an input, and all three are stated in terms of it. Once the row sums vary with the configuration, the top eigenvector is a Perron vector that row-sum normalisation no longer supplies, and none of those statements can be written down without first solving for it. The obstruction is therefore not that a proof got harder; it is that the first object in the chain stopped being available.*

Remark 4.5 (The symmetric kernel, and what is and is not formalised). *The coupled kernel above carries the spatial weight on the source only, which makes the row sum transparent but is not symmetric. The form usual in the physics literature is the symmetrised*

$$\tilde{K}_{\gamma}(\sigma, \tau) = w_{\gamma}(\sigma)^{1/2} K(\sigma, \tau) w_{\gamma}(\tau)^{1/2},$$

which is $D^{-1}K_{\gamma}D$ for $D = \text{diag}(w_{\gamma}^{1/2})$ and therefore has the same spectrum. Three things should be kept apart.

- **What is proved:** *for the oriented kernel, the row sums are not constant and the uniform vector is not fixed. That is the statement the elementary route actually consumes, because that route obtains the vacuum from constant row sums.*
- **What is asserted and not formalised:** *that the symmetrised kernel does not fix the uniform vector either. The diagonal similarity above is standard, but neither it nor that consequence is machine-checked in this development, and we mark it as prose rather than let it pass for a theorem.*
- **What is not claimed at all:** *that the transfer operator ceases to exist, or that no vacuum exists. A positive kernel on a finite set has a Perron vector. The obstruction is that the elementary normalisation stops handing it to us, not that the object is absent.*

A reader who suspects the obstruction is an artifact of a kernel convention is asking the right question, and the honest answer has a boundary in it. The convention does change the coordinate representative of the vacuum: a diagonal similarity changes the row sums and it changes the entries of the Perron vector, so no invariance should be asserted by fiat here. What is proved is narrower, and is exactly what the chain of [3] consumed: for the oriented kernel, the constant-row-sum certificate by which the elementary route produces the vacuum fails. The corresponding statement for the symmetrised kernel is recorded above as prose.

Remark 4.6 (What would have to happen next). *The open problem is a spectral bound for a transfer operator whose vacuum is not known in closed form. That is the genuine analytic content of the subject, it is where the constructive literature is difficult, and nothing in the present development touches it. Two honest continuations exist: a Perron–Frobenius route that produces the vacuum at spatial extent without a closed form, or a statement that the elementary methods of this lane end here. We do not know which, and we decline to estimate.*

5 Formalisation notes

One lemma, two opposite conclusions. Lemma 2.1 proves both that the decoupled row sums are constant and that the coupled ones are not. The difference is entirely in which function is fed to it, which is a compact way of exhibiting that the obstruction is structural rather than incidental.

A shared case split. Three separate proofs originally repeated the same two-case analysis of a single site, each slightly differently, and three of the module’s failures had that single cause. Extracting the case split and the one-site row sum as shared lemmas removed all three at once: [fin_two_cases](#), [sum_bond_row](#).

A linter warning obeyed without checking. A simplification step flagged as unused was in fact load-bearing; removing it broke a proof whose error message pointed elsewhere. The reduction is now performed explicitly so that it does not depend on a linter’s opinion. We record this because it is the second occurrence in this development.

6 Reproducibility

All artifacts are at source checkpoint `d8187d7f9caa9ce93a58a70a382e98a239942afe` 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	8424 jobs, success
Oracle commands	2481
with axiom dependencies	2459
axiom-free	22
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 criteria this work had to meet — including the requirement that the obstruction be a proved negative with explicit witnesses, and the requirement that the decoupled result not be presented as evidence about the coupled case — were fixed in [docs/O-BRIDGE-CHARTER.md](#), Amendment 10, in a commit preceding the existence of the module they govern. Their verdicts

are in [docs/VERIFICATION-LEDGER.md](#), Addendum 520. The oracle script is [oracle_check.lean](#).

References

- [1] K. Osterwalder and E. Seiler, *Gauge field theories on a lattice*, Annals of Physics **110** (1978), 440–471.
- [2] J. Glimm and A. Jaffe, *Quantum Physics: A Functional Integral Point of View*, 2nd ed., Springer, 1987.
- [3] L. Eriksson, *From the Gibbs Weight to the Spectral Gap: A Complete Machine-Checked Osterwalder–Seiler Chain for the $\mathbb{Z}_2$ Lattice Gauge Chain*, viXra, 2026.
- [4] L. Eriksson, *The Quotient That Is Not the Identity: A Machine-Checked Degenerate Reflection Pairing and Its Gelfand–Naimark–Segal Quotient*, viXra, 2026.
- [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.