

The Row Sums Were the Method, Not the Theorem: a Machine-Checked Chain from a Positive Weight to Exponential Decay of Correlations, and a Misattributed Uniformity Wall

Lluis Eriksson

Abstract

A formal development of a two-dimensional spatial transfer kernel proves its uniform spectral statements by Schur’s test on *constant* row sums, and carries a theorem, labelled THE OBSTRUCTION in its own source, stating that those row sums stop being constant the moment the spatial coupling is switched on. That theorem obstructs the *method*. It was read as obstructing the *conclusion*, and the reading outlived its own retraction.

We prove, in Lean 4 with no `sorry` and no project axiom, twelve theorems: a chain of six headline theorems, the star-cell and rectangle-family specialisations they compose to, the three transport theorems of the final contribution below, and the corollary they compose to. **First**, the one-bond influence envelope: flipping one neighbour across a bond of strength $J \geq 0$ moves a two-state site’s conditional by at most $\tanh J$, and by exactly $\tanh J$ at one field, so $\tanh J$ is the *least field-uniform* bound. It is a field-uniform interdependence majorant and need not equal the intrinsic coefficient of a given finite model, whose reachable fields form a discrete set — a distinction Section 2 makes precisely, and both of whose sides occur: equality is proved below for the one-bond cell. **Second**, those envelopes assembled into a matrix discharge every *matrix-side* hypothesis of the third (the two conditions on the ambient distance remain hypotheses), and turn a purely *local* covering condition into the remaining row-sum bound holding at every site of every volume — which is the coupling window, and which we compute: at a site carrying two bonds of strength β and two of strength γ it is exactly $2 \tanh |\beta| + 2 \tanh |\gamma|$, so the window is the hypothesis of a theorem rather than a sentence. **Third**, the resolvent bound: for a nonnegative matrix C over a finite index type, supported on pairs at distance at most one and with row sums *bounded* by $\alpha < 1$, without any constancy assumption,

$$\sum_{n < N} (C^n)_{ij} \leq \frac{\alpha^{\text{dist}(i,j)}}{1 - \alpha} \quad \text{for every } N,$$

with no reference of any kind to the cardinality of the index type. **Fourth**, *Dobrushin’s comparison estimate itself*, which an earlier version of this paper named as its one open analytic step: for a probability weight μ on a finite product space S^ι , invariant under a family of single-site conditionals whose influence is dominated by C ,

$$|\text{Cov}_\mu(f, g)| \leq \frac{1}{4} \sum_{i,j} \delta_i(f) \left(\sum_n (C^n)_{ij} \right) \delta_j(g),$$

where δ_i is oscillation in the single coordinate i ; composing with the third theorem replaces the inner series by $\alpha^{\text{dist}(i,j)} / (1 - \alpha)$ and yields exponential decay of connected correlations at a rate read off single-site data. The constant $\frac{1}{4}$ is *attained*, on a witness inside the kernel.

Fifth, the instantiation at an arbitrary weight: from a strictly positive weight w on S^i *alone* we construct the measure $\mu_w = w/Z$ and its heat-bath kernel, *prove* that μ_w is invariant — by an involution on $S^i \times S$, so that step is an identity rather than an estimate — and *prove* that the intrinsic Dobrushin matrix c^w has vanishing diagonal and dominates the influence. Every hypothesis of the fourth theorem is thereby discharged except one: $\sum_k c_{ik}^w \leq \alpha < 1$, Dobrushin’s condition itself, which is a condition on the interaction and not on the machinery. (The *geometric* form of the decay consumes one further hypothesis, stated where it is used: that c^w has range one for a distance obeying the two metric conditions. The resolvent form needs nothing else.) **Sixth**, the interaction itself: for the *Ising* weight of a symmetric zero-diagonal coupling J we prove that the heat-bath conditional is the sigmoid $(1 + \tanh h)/2$ at the local field $h_i = \sum_k J_{ik} \sigma_k$, and that the intrinsic matrix is dominated by the envelope of the first theorem, $c_{ik}^w \leq \tanh |J_{ik}|$ — so every hypothesis of the fifth theorem’s endpoints becomes a computation on J . On the one-bond system the domination is an *equality*: the optimal field is reachable there, and the intrinsic coefficient *is* $\tanh \beta$; that system moreover satisfies Dobrushin’s condition at every coupling, so its decay bound carries no smallness hypothesis at all. And on the star cell the window $2 \tanh |\beta| + 2 \tanh |\gamma| < 1$ — the second theorem’s row-sum computation — becomes the hypothesis of an exponential-decay theorem about the anisotropic Ising weight itself. **And the family**: on the $L \times T$ rectangle with β on horizontal bonds and γ on vertical ones, the envelope row sums respect the window at every site of every rectangle, so the decay theorem holds with the couplings, the rate and the prefactor fixed *before* the quantifier over volumes — the volume-uniformity of the measure-level statement is the quantifier order of one machine-checked declaration. At zero coupling its bound degenerates to 0^d and forces exact independence of distinct single-site observables at every volume, which is the family’s witness. **And the transport**: the pairing $\langle v, T^n v \rangle$ of a normalised kernel with Perron boundary is *exactly* a finite band measure whose mass is independent of the horizon; uniform decay of those band covariances at a common rate — a hypothesis on *path sums*, carrying no operator, no norm and no spectrum — forces one strictly positive mass for the whole family of projected transfer operators, with the identification measure-to-operator a proved identity rather than an assumption; and for Perron data of a symmetrised kernel the band measure is the free strip measure tilted at its two end slices by exactly $\Omega/\sqrt{w}$, so free-covariance decay transports to band-covariance decay at the same rate with the whole boundary cost absorbed into a per-extent constant. Each rung carries a non-vacuity witness, and every claimed endpoint is oracled with exactly Lean’s three standard axioms.

Separately, and by measurement rather than proof, we remove the wall from where it was being placed. For the coupled kernel the second-eigenvalue ratio $\text{specRatio}(L)$ saturates strictly below one in each of six cells tested inside Onsager’s disordered region — with the spatial weight switched on — and reaches one to six decimal places in the two ordered cells tested. The control at zero spatial coupling reproduces the *proved* identity $\text{specRatio} = \tanh \beta$ to six decimals, so the harness is calibrated against a theorem rather than against itself. These data are eight cells at $L \leq 12$ with an extrapolation. On the tested grid they *contradict* the blanket attribution of the degeneracy to the mere presence of the spatial weight; they are not, and are not offered as, a refutation of an asymptotic claim, since a sequence can appear to saturate at $L \leq 12$ and turn afterwards. They do *not* locate the transition, which would need a dense sweep and finite-size scaling that we do not attempt.

And the corollary the chain was for: inside the window there is ONE $m > 0$ bounding the projected transfer operator of the coupled kernel’s normalised Perron data by e^{-m} at *every* extent — the currying identification of the strip sums with the rectangle covariances, and the feed of the family theorem into the abstract transport theorem, are both exact and machine-checked. **What is not proved**. No infinite-volume state is constructed; the window is sufficient, not sharp; the identification of the operator-norm bound with a spectral statement is the standard symmetric-operator reading and is stated as such where it is used. The intrinsic Dobrushin condition is now a computation on the coupling: it holds

unconditionally for the one-bond system, reduces to the window for the star cell and for the whole rectangle family, and for other lattice geometries is the same computation, not repeated here. Dobrushin’s uniqueness theorem and his comparison estimate are classical [1, 2, 5]; nothing in this paper is new mathematics. What is offered is a mechanised, non-vacuous chain from a single bond to exponential decay of correlations, and measured counterevidence against an attribution. No consequence for Yang–Mills theory is stated, suggested, or implied.

1 What this paper claims, and what it does not

Claimed, machine-checked. Theorems 3.1, 3.2, 3.3, 4.3, 5.2 and 5.5, with Theorems 5.7 and 5.8, their witnesses and the corollaries that compose them, in [DobrushinCoefficient.lean](#), [DobrushinRowSum.lean](#), [DobrushinMatrix.lean](#), [DobrushinOscillation.lean](#), [DobrushinGruss.lean](#), [DobrushinConditional.lean](#) and [DobrushinComparison.lean](#), [DobrushinGibbs.lean](#), [DobrushinIsing.lean](#) and [DobrushinLattice.lean](#); and the transport rungs of Section 6 — Theorems 6.1, 6.2, 6.4 and 6.5 with Proposition 6.3 — in [DobrushinBridge.lean](#), [DobrushinTransport.lean](#), [DobrushinTilt.lean](#) and [DobrushinCorollary.lean](#). Every claimed endpoint is oracled and depends on exactly Lean’s three standard axioms; the counts and the runs are itemised in the Reproducibility section.

Claimed, measured. The saturation table of Section 7. These are floating-point computations at finite extent with a documented extrapolation; they are evidence, not proof, and they are labelled as such wherever they appear.

Claimed, about the record. Section 2 states precisely what the source theorem `coupled_rowSums_not_constant` does and does not license. This is a statement about a formal development, and it is verifiable by reading that development.

Not claimed. An infinite-volume state — none is constructed, and the word *clustering* is not used for one. Sharpness of the window — it is sufficient, and Section 9 says by how much it undershoots the exact answer where one exists. Anything beyond the symmetric-operator reading of the norm bound of Theorem 6.5, which is stated where it is used. That Dobrushin’s condition holds beyond the geometries treated — the one-bond and star cells and the rectangle family of Theorem 5.8; for a general weight it is the hypothesis Theorem 5.2 leaves to its user, together with, for the geometric form of the decay, the range-one support of the intrinsic matrix. That the window derived from Dobrushin’s condition is optimal — it is not, and Section 9 says by how much. That any of this bears on the Yang–Mills mass gap.

2 The setting, and the theorem that was read too widely

Fix an extent L and consider configurations $\sigma \in \{0, 1\}^L$ of one slice. The development under discussion carries a time-bond kernel that factorises over sites,

$$K_\beta(\sigma, \tau) = \prod_{j < L} b_\beta(\sigma_j, \tau_j), \quad b_\beta(s, s') = e^{\pm\beta},$$

and an arbitrary strictly positive slice weight w , the Gibbs weight of a space-time configuration being the product of the slice weights and the time bonds. When w is itself a product over spatial bonds with coupling γ , the system is the anisotropic two-dimensional Ising model.

Two facts from that development matter here.

Proposition 2.1 (proved elsewhere in the development). *At $\gamma = 0$ and for every $L \geq 1$, $\text{specRatio} = \tanh \beta$ exactly.*

The proof is Schur’s test [8]: the decoupled kernel has *constant* row sums, and the loss on the odd part of an observable is exactly compensated by the odd bond eigenvalue at every step.

Proposition 2.2 (`coupled_rowSums_not_constant`). *At $\gamma > 0$ the row sums of the coupled kernel are not constant: two explicit configurations of a two-site slice have different row sums.*

Proposition 2.2 is labelled THE OBSTRUCTION in the source. *It obstructs Schur’s test.* It says nothing whatever about `specRatio`, it is proved on a two-site witness, and no statement quantified over L appears in it.

A stronger reading — that the spatial weight destroys uniformity — was entered into the development and then *retracted of record* as unproved, leaving behind only a measurement of degeneracy.

What is documented, and can be checked. The retraction of the sentence from the manuscript is recorded at Addendum 561 of the development’s verification ledger; its removal from the module docstring, which the first retraction had missed, at Addendum 563; and its survival in a hand-typed submission form, caught only because the form was read in full before delivery, at Addendum 568. Three removals were required for one sentence, and the third was the one that would have been published.

What is inference, and is labelled as such. We observe that no statement quantified over L for the coupled kernel was attempted in the interval, and we attribute this to the retracted reading continuing to govern target selection. That attribution is ours; it is not established by a diff, and no artefact in the development records a decision not to attempt something. The verifiable half of the observation is the narrower one, and it is enough for this paper: the guards that caught the sentence three times compare identifiers and substrings inside files, so whatever else was carrying the reading was not a file.

3 The chain

Three theorems, each machine-checked, composing into one statement. The first is about a single bond, the second assembles bonds into a matrix, the third is about that matrix and never sees the model.

3.1 The one-bond influence envelope, and that it is least among field-uniform bounds

A two-state site in a real field h has conditional $P(+1) = (1 + \tanh h)/2$, written from the Boltzmann weights and nothing else. Flipping one neighbour across a bond of strength J moves that field to $h - 2J$. Write $\text{tv}(h, h') = |P_h(+1) - P_{h'}(+1)|$.

Theorem 3.1 (machine-checked). *For every $J \geq 0$,*

$$\text{tv}(h, h - 2J) \leq \tanh J \quad \text{for all } h, \quad \text{and} \quad \text{tv}(J, -J) = \tanh J.$$

Hence $\tanh J$ is the least upper bound of $\{\text{tv}(h, h - 2J) : h \in \mathbb{R}\}$.

Proof. $\tanh a - \tanh b = \sinh(a - b)/(\cosh a \cosh b)$, and $\cosh a \cosh b = (\cosh(a + b) + \cosh(a - b))/2 \geq (1 + \cosh(a - b))/2$ since $\cosh \geq 1$. With $a = h$ and $b = h - 2J$ the difference of the fields is $2J$, so the denominator is at least $(1 + \cosh 2J)/2 = \cosh^2 J$ while the numerator is $\sinh 2J = 2 \sinh J \cosh J$. Dividing by two, $\text{tv} \leq \sinh J \cosh J / \cosh^2 J = \tanh J$. At $h = J$ the two fields are $+J$ and $-J$, their sum is zero, $\cosh 0 = 1$, and every inequality above is an equality. $\square$

What this is, stated carefully, because the obvious reading of it is wrong. The supremum is taken over *all* $h \in \mathbb{R}$. In a concrete finite system with the remaining couplings and any external field fixed, the effective fields actually reachable at a site form a discrete subset of $\mathbb{R}$, and the value $h = J$ at which equality holds need not correspond to any admissible boundary condition of that system. So the true Dobrushin coefficient

$$c_{ij} = \sup\{\|\mu_i(\cdot | \eta) - \mu_i(\cdot | \eta')\|_{\text{TV}} : \eta, \eta' \text{ differ only at } j\}$$

of a particular model may satisfy $c_{ij} < \tanh |J_{ij}|$ strictly.

Theorem 3.1 is therefore a **sharp field-uniform one-bond influence envelope**: sharp as a bound over arbitrary real fields, and a valid interdependence majorant, $c_{ij} \leq \tanh |J_{ij}|$, for every model with that bond. That is all Dobrushin's condition needs and it invalidates nothing downstream — but it is not the claim that $\tanh |J|$ is the minimal interdependence matrix of *every* system with that bond: on multi-bond cells the reachable fields form a discrete set and the intrinsic coefficient can sit strictly below the envelope. Where the two coincide is itself settled below: on the one-bond cell of Section 5 the coincidence is exact, and it is a theorem (Proposition 5.6). The window below is a sufficient condition, explicit and not optimal.

The attainment still earns its place: without it the envelope would be some number above the influence, and one could not say the envelope is the best field-uniform one. In Lean the three statements are `tvField_le`, `tvField_attained` and `tvField_isLUB`.

3.2 The matrix, and the window as a hypothesis

Given sites i , bond strengths J and a distance d , put

$$C_{ij} = \begin{cases} \tanh |J_{ij}| & i \neq j \text{ and } d(i, j) \leq 1, \\ 0 & \text{otherwise.} \end{cases}$$

Theorem 3.2 (machine-checked). *C is entrywise nonnegative, vanishes on the diagonal, and vanishes whenever $d(i, j) > 1$. Consequently every matrix-side hypothesis of Theorem 3.3 — (i) and (ii) — holds by construction. Furthermore, if at every site the influence support is contained in a finite set whose bond coefficients sum to at most α , then the row-sum hypothesis holds at every site, with neither the statement nor α mentioning the cardinality of the index type. Finally, at a site carrying two bonds of strength β and two of strength γ ,*

$$\sum_j C_{ij} = 2 \tanh |\beta| + 2 \tanh |\gamma|.$$

Three remarks, in decreasing order of how easily the statement is over-read.

What is not discharged. The construction settles the matrix-side hypotheses. It does *not* discharge the two conditions on the ambient distance — $d(i, i) = 0$ and the triangle inequality — which Corollary 3.7 still takes as hypotheses, as does the Lean endpoint. For a graph distance

they are immediate; for an arbitrary d they are assumptions, and the paper does not pretend otherwise.

One site versus a family. The last display is arithmetic at one site of one five-point graph, and by itself it would say nothing about a family of volumes. The middle clause is what makes it say something: it converts a *local* geometric fact — each site’s influence support is covered by a set whose coefficients sum below one — into the row-sum hypothesis holding simultaneously at every site of every volume with that local geometry. In Lean these are `rowSum_bound_of_local_cover` and `dobrushin_bound_of_local_cover`.

The window. It has been carried in this lane’s prose as $2 \tanh \beta + 2 \tanh \gamma < 1$; it is now computed rather than asserted and is consumed by Corollary 3.7. It remains a sufficient condition and, by Section 9, a conservative one. The distance used in the computation is the graph distance on a star, proved reflexive at zero and triangular by exhaustive checking over the 125 triples.

3.3 The resolvent bound

Theorem 3.3 (machine-checked). *Let ι be a finite type, $d : \iota \times \iota \rightarrow \mathbb{N}$ satisfy $d(i, i) = 0$ and the triangle inequality, and let $C : \iota \times \iota \rightarrow \mathbb{R}$ satisfy*

- (i) $C_{ij} \geq 0$ for all i, j ;
- (ii) $C_{ij} = 0$ whenever $d(i, j) > 1$;
- (iii) $\sum_j C_{ij} \leq \alpha$ for all i , with $0 \leq \alpha < 1$.

Then for all i, j and every $N \in \mathbb{N}$,

$$\sum_{n \in [0, N)} (C^n)_{ij} \leq \frac{\alpha^{d(i, j)}}{1 - \alpha},$$

and consequently $\sum_n (C^n)_{ij} \leq \alpha^{d(i, j)} / (1 - \alpha)$.

Proof. Three steps, each an induction on n . Nonnegativity is preserved by powers. Row sums multiply: expanding (C^{n+1}) and exchanging the two finite sums, $\sum_j (C^{n+1})_{ij} = \sum_k (C^n)_{ik} \sum_j C_{kj} \leq \alpha \sum_k (C^n)_{ik} \leq \alpha^{n+1}$, so each entry, being at most its row sum, is at most α^n . Support: if $n < d(i, j)$ then $(C^n)_{ij} = 0$, since in $(C^{n+1})_{ij} = \sum_k (C^n)_{ik} C_{kj}$ each k has either $d(i, k) > n$, killing the first factor by induction, or $d(i, k) \leq n$, whence $d(i, j) \leq d(i, k) + d(k, j) \leq n + d(k, j)$ forces $d(k, j) > 1$ and kills the second. Combining, $(C^n)_{ij} \leq \alpha^n$ when $n \geq d(i, j)$ and vanishes otherwise; splitting $[0, N)$ at $d(i, j)$ and reindexing the surviving block by $n = d(i, j) + m$ gives $\alpha^{d(i, j)} \sum_{m < N - d(i, j)} \alpha^m$, and the geometric sum is at most $(1 - \alpha)^{-1}$. $\square$

Remark 3.4 (where the volume-freeness is). *The cardinality of ι appears nowhere in the hypotheses or the conclusion. That is the entire content of the theorem: the index type may be as large as one likes and neither α nor the exponent notices. Every constant is read off single-site data.*

Remark 3.5 (the hypothesis that matters here). *Hypothesis (iii) asks for row sums bounded by α . It never asks for them to be constant. Proposition 2.2 refutes constancy and leaves (iii) untouched: it is precisely the crack through which the coupled kernel passes after Schur’s test has died.*

Proposition 3.6 (non-vacuity, machine-checked). *On the three-point chain with d the graph distance and $C_{ij} = 1/4$ for $d(i, j) \leq 1$ and 0 otherwise, hypotheses (i)–(iii) hold with $\alpha = 3/4 < 1$, the matrix is not the zero matrix, the two endpoints are at distance exactly two, and the conclusion reads: the whole resolvent series between the endpoints is at most $9/4$.*

The witness is not decoration. Hypotheses (ii) and (iii) are jointly satisfiable by the zero matrix, for which the conclusion is trivial, and by data with $d \equiv 0$, for which the exponent is trivial. Proposition 3.6 excludes both.

3.4 The composition

Corollary 3.7 (machine-checked). *Let J be bond strengths on a finite set of sites with a distance d satisfying $d(i, i) = 0$ and the triangle inequality, and let C be the matrix of Theorem 3.2. If every row sum of C is at most $\alpha < 1$, then for all i, j and every N ,*

$$\sum_{n < N} (C^n)_{ij} \leq \frac{\alpha^{d(i,j)}}{1 - \alpha},$$

and the same holds for the full series.

This is the whole chain in one line: the envelope of Theorem 3.1 enters the matrix of Theorem 3.2, which satisfies the hypotheses of Theorem 3.3, whose conclusion mentions no cardinality. In Lean it is `dobrushin_resolvent_bound`, and its proof is the application of one theorem to three others.

4 The comparison estimate

Theorems 3.1–3.3 bound a resolvent. Nothing in them mentions a measure, and so nothing in them says that correlations are controlled by that resolvent. That implication is Dobrushin’s comparison estimate, and the previous version of this paper named it as the single open analytic step. This section proves it.

4.1 The objects

Let ι be a finite index type and S a finite nonempty state space, so that the configuration space S^ι is finite. For $f : S^\iota \rightarrow \mathbb{R}$ put

$$\delta_i(f) = \max\{|f(\eta) - f(\eta')| : \eta, \eta' \text{ agree off } i\},$$

the oscillation of f in the single coordinate i . A *single-site kernel* is a family $p_i(\eta, \cdot)$ of probability weights on S that is *local*: $p_i(\eta, \cdot)$ does not depend on η_i . It acts on observables by

$$(E_i f)(\eta) = \sum_{s \in S} p_i(\eta, s) f(\eta^{i \rightarrow s}),$$

where $\eta^{i \rightarrow s}$ is η with the value at i replaced by s . We say a matrix C *dominates the influence* of the kernel if, for $k \neq i$,

$$\|p_i(\eta, \cdot) - p_i(\eta', \cdot)\|_{\text{TV}} \leq C_{ik} \quad \text{whenever } \eta, \eta' \text{ agree off } k.$$

The orientation is the one fixed before any of this Lean was written: the *first* index of C is the site being conditioned, the *second* is the site whose value moved. On a symmetric cell the two conventions are indistinguishable, which is why fixing it in advance mattered; the gate that discriminates is described at the end of this section.

4.2 The key lemma

Lemma 4.1 (machine-checked). *Let p be a local single-site kernel and C dominate its influence. Then for every observable f :*

$$\delta_i(E_i f) = 0, \quad \delta_k(E_i f) \leq \delta_k(f) + C_{ik} \delta_i(f) \quad (k \neq i).$$

Proof. The first clause is semantics, not estimation: by locality $E_i f$ does not depend on the coordinate i at all. For the second, split the difference of $E_i f$ at two configurations agreeing off k into a *transport* term, where the kernel is held fixed and f is evaluated at two configurations that still agree off k — here $k \neq i$ is load-bearing — and a *kernel variation* term $\sum_s (p_i(\eta, s) - p_i(\eta', s)) f(\eta'^{i \rightarrow s})$. The first is at most $\delta_k(f)$ because $p_i(\eta, \cdot)$ is a probability weight. The second is a signed mass of total zero tested against a section of f along coordinate i ; a zero-sum mass cannot see a constant, so subtracting the midpoint of the range gives $|\sum_s a_s g_s| \leq \frac{1}{2} (\sum_s |a_s|) \text{osc}(g)$, which is total variation times oscillation, at most $C_{ik} \delta_i(f)$. $\square$

The three statements are kept apart in the source on purpose. The vanishing clause is `deltaAt_condExp_self` and mentions no matrix; the transport clause is `deltaAt_condExp_le` and asks only that C *dominate*, never that it be minimal — the same distinction that cost this lane two versions over $\tanh|J|$; and their assembly into a single vector inequality, `deltaAt_condExp_le_matrix`, is the *only* statement that assumes $C_{ii} = 0$. That the zero diagonal is a declared hypothesis of the matrix form and not an accident of the physical cell was decided, and written into the lane charter, before this module existed.

4.3 The one-step quarter

Lemma 4.2 (machine-checked). *For a probability weight p on a finite set and any f, g with ranges in $[m_1, M_1]$ and $[m_2, M_2]$,*

$$|\mathbb{E}_p[fg] - \mathbb{E}_p[f] \mathbb{E}_p[g]| \leq \frac{1}{4} (M_1 - m_1)(M_2 - m_2),$$

and the constant $\frac{1}{4}$ is attained.

Proof. Not Cauchy–Schwarz. The centred mass $a_x = p_x(f_x - \mathbb{E}_p f)$ has total zero, so the signed-mass inequality above gives $|\text{Cov}| \leq \frac{1}{2} (\sum_x |a_x|) \text{osc}(g)$, and $\sum_x |a_x|$ is the mean absolute deviation of f about its mean, which is at most $(M_1 - m_1)/2$: split it into its positive and negative parts, observe they are *equal* because the mass is centred, bound one by $(M_1 - \mathbb{E}f)T$ and the other by $(\mathbb{E}f - m_1)(1 - T)$ with T the mass above the mean, and multiply. The fair two-point weight with $f = g$ an indicator attains it. $\square$

Each factor $\frac{1}{2}$ is born in exactly one lemma and the quarter is their product. Popoviciu’s variance bound, also machine-checked in the same module, is *not* used: it is the L^2 companion, kept separate so that a difficulty in either could not block the other.

4.4 The estimate

Theorem 4.3 (machine-checked). *Let μ be a probability weight on S^ι , invariant under every E_i (that is, $\mathbb{E}_\mu[E_i F] = \mathbb{E}_\mu[F]$ for all F and i), let C dominate the influence of the kernel with $C_{ii} = 0$ and row sums at most $\alpha < 1$. Then for all observables f, g ,*

$$|\text{Cov}_\mu(f, g)| \leq \frac{1}{4} \sum_{i,j} \delta_i(f) \left(\sum_n (C^n)_{ij} \right) \delta_j(g),$$

and, composing with Theorem 3.3 when C is supported at distance at most one,

$$|\text{Cov}_\mu(f, g)| \leq \frac{1}{4(1-\alpha)} \sum_{i,j} \delta_i(f) \alpha^{\text{dist}(i,j)} \delta_j(g).$$

For observables depending on single sites i_0 and j_0 this reads

$$|\text{Cov}_\mu(f, g)| \leq \frac{\delta_{i_0}(f) \delta_{j_0}(g)}{4(1-\alpha)} \alpha^{\text{dist}(i_0, j_0)} :$$

exponential decay in the distance, at a rate free of the volume.

Proof. Let $P = \frac{1}{m} \sum_i E_i$ be the random-scan average, $m = \#\iota$. It preserves μ -expectations, so telescoping over N steps,

$$\text{Cov}_\mu(f, g) = \sum_{n < N} \left(\text{Cov}_\mu(P^n f, g) - \text{Cov}_\mu(P^{n+1} f, g) \right) + \text{Cov}_\mu(P^N f, g).$$

Each increment is $\mathbb{E}_\mu[g \cdot (h - Ph)]$ with $h = P^n f$; averaging over the site chosen and using invariance once more, the i -th contribution equals $\mathbb{E}_\mu[E_i(g(h - E_i h))]$, whose integrand is *pointwise* the covariance of the two sections of g and h along coordinate i under $p_i(\eta, \cdot)$. By Lemma 4.2 and the fact that a section oscillates at most δ_i , each increment is at most $\frac{1}{4m} \sum_i \delta_i(g) \delta_i(P^n f)$.

By Lemma 4.1 the oscillation vector of the iterates is transported by $v \mapsto (1 - \frac{1}{m})v + \frac{1}{m} C^\top v$, so $\delta(P^n f)$ is dominated coordinatewise by the n -th iterate of that map applied to $\delta(f)$. Two things follow, both by finite induction and neither taking a limit. Its ℓ^1 norm contracts by $1 - \frac{1-\alpha}{m}$ per step, so the telescoping remainder $|\text{Cov}_\mu(P^N f, g)|$, bounded crudely by $\frac{1}{4} \text{osc}(g) \sum_i \delta_i(P^N f)$, dies geometrically. And the *accumulated* transport $\sum_{n < N}$ of those iterates is dominated entrywise by m times the N -th partial sum of $\sum_n (C^\top)^n$ — which is the step where the factor m introduced by the averaging is exactly cancelled, leaving a bound with no m in it. Assembling gives the finite- N statement; letting $N \rightarrow \infty$ in the *bound* gives the series form. $\square$

Every statement up to the last sentence of that proof is a finite inequality over a finite configuration space. The single limit is taken at the level of the bound, which is a design decision this lane recorded before the module existed, so that no infinite-volume object is constructed anywhere.

Proposition 4.4 (non-vacuity and attainment, machine-checked). *One site, two states, the fair kernel, the uniform weight, $C = 0$, $\alpha = 0$, $d \equiv 0$. Every hypothesis of Theorem 4.3 holds by computation, and on the indicator observable the two-point conclusion is an equality: $|\text{Cov}| = \frac{1}{4}$ against a bound whose value is $\frac{1}{4}$.*

The witness does two jobs. It shows the hypotheses are jointly satisfiable — invariance of μ under all E_i is a real constraint, not a formality. And because the bound is *attained* there, the constant $\frac{1}{4}$ cannot be lowered by any argument, in this or any other proof of the same inequality.

4.5 The gates, and the one that discriminated

Five gates were registered before the corresponding Lean and are in [judge_dobrushin_d3.py](#) and [judge_dobrushin_d3b.py](#). Two are worth naming.

The gate that could have killed the rung. Lemma 4.1 was checked by brute force on small Ising cells over *all* $2^{|S^c|}$ Boolean observables — 256 of them at three sites — computing both sides exactly for every site pair. The gate reports the worst value of lhs – rhs, and a positive value would have refuted the lemma before a line of it was written. Measured: $+1.11 \times 10^{-16}$, one unit in the last place of a float64 comparison at scale one.

The gate that discriminated the orientation. On an Ising cell the minimal influence coefficients are symmetric, so $C = C^\top$ and a transposed index would pass every other gate and surface much later. A three-site cell built from eight distinct positive weights has $\max_{ij} |C_{ij} - C_{ji}| > 0$, and there the declared orientation passes with zero violations while the transposed one **fails with** 24. The convention in Lemma 4.1 is therefore the one the lemma needs, and that is established by measurement rather than by assertion.

Both certifiers were run in `normal` and in `optimized (python -O)` mode on the Linux plane, all four runs exiting zero with their check counters matching.

5 The instantiation: from a weight, not from data

Theorem 4.3 takes *data* — a weight, a kernel, a majorant — and six hypotheses. A reader is entitled to ask whether that data exists for anything, and whether the hypotheses are met by the object one actually has, which is an interaction, not a kernel. This section removes the question. Its input is a single function.

Definition 5.1. Let $w : S^c \rightarrow \mathbb{R}$ be strictly positive. Its *Gibbs weight* is $\mu_w = w/Z$ with $Z = \sum_{\eta} w(\eta)$, and its *heat-bath kernel* is

$$p_i^w(\eta, s) = \frac{w(\eta^{i \rightarrow s})}{\sum_{t \in S} w(\eta^{i \rightarrow t})}.$$

Its *intrinsic Dobrushin matrix* is

$$c_{ik}^w = \max \left\{ \|p_i^w(\eta, \cdot) - p_i^w(\eta', \cdot)\|_{\text{TV}} : \eta, \eta' \text{ agree off } k \right\},$$

the maximum over the finite configuration space — the coefficient itself, not a majorant chosen from outside.

Theorem 5.2 (machine-checked). *For every strictly positive w : μ_w is a probability weight; p^w is nonnegative, normalised and local; μ_w is invariant under every E_i^w ; $c_{ii}^w = 0$; and c^w dominates the influence of p^w . Consequently, if $\sum_k c_{ik}^w \leq \alpha < 1$ at every site, then for all observables*

$$|\text{Cov}_{\mu_w}(f, g)| \leq \frac{1}{4} \sum_{i,j} \delta_i(f) \left(\sum_n (c^w)_{ij}^n \right) \delta_j(g),$$

and if moreover c^w vanishes beyond distance one for a distance d , then for single-site observables at i_0 and j_0 ,

$$|\text{Cov}_{\mu_w}(f, g)| \leq \frac{\delta_{i_0}(f) \delta_{j_0}(g)}{4(1 - \alpha)} \alpha^{d(i_0, j_0)}.$$

Proof. Normalisation and locality are computations: the local partition sum $\sum_t w(\eta^{i \rightarrow t})$ is positive, and it does not see η_i because $\eta^{i \rightarrow t}$ does not.

Invariance is the only step with content, and it is an identity rather than an estimate. Write $\mathbb{E}_{\mu_w}[E_i^w F]$ as a sum over $S^t \times S$ and apply the involution

$$\tau_i(\eta, s) = (\eta^{i \rightarrow s}, \eta_i), \quad \tau_i \circ \tau_i = \text{id}.$$

Under τ_i the integrand $\mu_w(\eta) p_i^w(\eta, s) F(\eta^{i \rightarrow s})$ becomes $\mu_w(\eta^{i \rightarrow s}) \cdot \frac{w(\eta)}{\sum_t w(\eta^{i \rightarrow t})} \cdot F(\eta)$, using that the local partition sum is unchanged along the fibre. Summing that over s first reconstitutes $\sum_s \mu_w(\eta^{i \rightarrow s})$, which is the local partition sum over Z ; it cancels the denominator exactly, leaving $\mu_w(\eta) F(\eta)$. Reindexing a finite sum by an involution is free, and that is the whole proof.

The zero diagonal is a theorem, not a hypothesis: c_{ii}^w is a maximum of total variations between $p_i^w(\eta, \cdot)$ and $p_i^w(\eta^{i \rightarrow s}, \cdot)$, which are equal by locality. Domination is immediate because c^w is defined as the maximum. The conclusions are Theorem 4.3 and its corollaries applied to this data. $\square$

Three remarks, in the order in which the statement is most easily over-read.

What is left as a hypothesis. Dobrushin’s condition $\sum_k c_{ik}^w \leq \alpha < 1$ — plus, for the distance form, that c^w has range one. Both are conditions on the *interaction*, not on the machinery: nothing about kernels, invariance or majorants survives for a user to discharge. In Lean the endpoints are `gibbs_covar_le_resolvent`, `gibbs_covar_exp_decay` and `gibbs_covar_two_point`.

Why the heat-bath kernel and not another. The comparison estimate holds for *any* invariant local kernel; the heat-bath one is chosen because its invariance is an involution rather than a computation, and because its influence coefficients are the ones Dobrushin’s condition is usually stated for. Nothing downstream depends on the choice.

The intrinsic matrix versus the envelope. c^w is the coefficient, so for a bonded interaction it is bounded by the envelope $\tanh |J_{ij}|$ of Theorem 3.1 and may be strictly smaller. That domination is no longer only measured: for the Ising weight it is Theorem 5.5 below. What stays measured-only is the *strictness* of the gap on multi-bond cells (gate G14) — and the one-bond witness shows the gap is not universal, because there the two coincide exactly.

Proposition 5.3 (non-vacuity of the instantiation, machine-checked). *For the uniform weight on one site with two states, every hypothesis of Theorem 5.2 holds by computation, its Gibbs weight is the witness weight of Proposition 4.4, and the instantiated two-point bound is again an equality at $\frac{1}{4}$.*

So the instantiated hypotheses are jointly satisfiable and the instantiated bound is still the sharp one. This matters more than the usual non-vacuity check, because invariance is a genuine constraint on the pair (weight, kernel) and a construction that failed it would make Theorem 5.2 say nothing.

Gates. `judge_dobrushin_d4.py`, registered before this Lean. G12 checks invariance as an *identity* to 10^{-12} on deliberately asymmetric weights, exhaustively over all Boolean observables and all sites — a residual above roundoff would have killed the involution argument before it was written. G13 checks that the heat-bath conditional of an Ising weight is $(1 + \tanh h)/2$ at the local field. G14 checks that the intrinsic coefficients are dominated by $\tanh |J|$ entrywise, which is the measured content of “the envelope is a majorant, not the minimum” at the level of the measure.

The interaction, instantiated

Theorem 5.2 starts at an arbitrary positive weight. A reader's weight is not arbitrary: it is built from a coupling. The last rung instantiates that too.

Definition 5.4. Let $J : \iota \times \iota \rightarrow \mathbb{R}$ be symmetric with zero diagonal. The *Ising weight* of J is $w_J(\eta) = \exp(\frac{1}{2} \sum_{a,b} J_{ab} \sigma_a \sigma_b)$ with $\sigma_a = \pm 1$ the spin of η_a , and the *local field* at i is $h_i(\eta) = \sum_k J_{ik} \sigma_k$.

Theorem 5.5 (machine-checked). *For every such J : the heat-bath conditional of w_J at site i is the sigmoid of the local field, $p_i^{w_J}(\eta, +1) = \frac{1 + \tanh h_i(\eta)}{2}$; the intrinsic matrix is dominated by the envelope,*

$$c_{ik}^{w_J} \leq \tanh |J_{ik}|,$$

vanishing wherever the coupling does; and consequently, if the envelope row sums obey $\sum_k \tanh |J_{ik}| \leq \alpha < 1$ at every site, every conclusion of Theorem 5.2 holds for w_J with the same α — the comparison estimate, and, when J has range one for a distance, exponential decay in that distance. In Lean the endpoints are `heatBath-ising-pPlus`, `dobCoeff-ising-le` and `ising-covar-{le-resolvent, exp-decay, two-point}`.

Proof. The quadratic energy of a configuration updated at i splits as an i -free block plus $\sigma_i h_i(\eta)$ (`energy_update`: split the double sum at $a = i$ and $b = i$, kill the diagonal with $J_{ii} = 0$, merge the two cross blocks by symmetry). The local partition sum is then $e^h + e^{-h}$ times an i -free factor, which cancels in the heat-bath quotient: the sigmoid identity. Updating site k moves the field by $J_{ik}(\sigma'_k - \sigma_k)$, and $\sigma'_k - \sigma_k \in \{0, \pm 2\}$, so the total variation of the two conditionals is the envelope's $\text{tv}(h, h - 2J')$ at $|J'| = |J_{ik}|$ — Theorem 3.1 bounds it by $\tanh |J_{ik}|$. Taking the supremum gives the entrywise domination; the endpoints compose with Theorem 5.2. $\square$

Proposition 5.6 (the one-bond witness, and the domination is an equality there). *For two sites joined by one bond $\beta \geq 0$: the intrinsic coefficient equals the envelope, $c_{01}^w = \tanh \beta$ (`dobCoeff-bond-attained`), because the all-up configuration realises the optimal field $h = \beta$ and flipping the neighbour realises $-\beta$. Moreover the envelope row sum is $\tanh |\beta| < 1$ at every coupling, so this system's two-point decay bound (`bond-covar-two-point`) carries no smallness hypothesis at all.*

The equality matters beyond non-vacuity: Section 2's distinction between the envelope and the minimal coefficient is a genuine gap on multi-bond cells (G14 measures it), and the one-bond system shows the gap can also close — the reachable fields decide, exactly as Theorem 3.1's attainment said they would.

Theorem 5.7 (the window, instantiated at a weight; machine-checked). *On the star of Section 3 — a centre carrying two bonds β and two bonds γ — mask the bond function to the star's edges, so the coupling is symmetric, zero-diagonal and of range one for the star distance. If*

$$2 \tanh |\beta| + 2 \tanh |\gamma| < 1,$$

then for single-site observables f at i_0 and g at j_0 ,

$$|\text{Cov}_\mu(f, g)| \leq \frac{\delta_{i_0}(f) \delta_{j_0}(g)}{4(1 - (2 \tanh |\beta| + 2 \tanh |\gamma|))} (2 \tanh |\beta| + 2 \tanh |\gamma|)^{d_\star(i_0, j_0)},$$

where μ is the Gibbs weight of the star’s Ising weight and $d_\star$ the star distance. In Lean: `star_covar_two_point`. The row-sum hypothesis is discharged by computation: the centre’s envelope row sum is exactly the window, each leaf’s is $\tanh |\beta|$ or $\tanh |\gamma|$.

The window of Section 3 was “the hypothesis of a theorem rather than a sentence”; it is now the hypothesis of a theorem *about the anisotropic Ising weight itself*, which is the composition the abstract promises. The star is a cell, not a lattice, and the theorem does not evaluate the window for any numerical (β, γ) ; no interval arithmetic is used anywhere in this paper.

The family, and the quantifier order

A cell is not a family of volumes. The last rung supplies the family: the $L \times T$ rectangle with coupling β on horizontal nearest-neighbour bonds, γ on vertical ones, free boundary, and the Manhattan distance $d_\square((a, b), (c, d)) = |a - c| + |b - d|$.

Theorem 5.8 (the volume family; machine-checked). *Fix β, γ and α with $2 \tanh |\beta| + 2 \tanh |\gamma| \leq \alpha < 1$. Then for every $L \geq 1$ and $T \geq 1$, for single-site observables f at p_0 and g at q_0 on the $L \times T$ rectangle,*

$$|\text{Cov}_\mu(f, g)| \leq \frac{\delta_{p_0}(f) \delta_{q_0}(g)}{4(1 - \alpha)} \alpha^{d_\square(p_0, q_0)},$$

where μ is the Gibbs weight of the rectangle’s Ising coupling. In Lean: `rect_ising_uniform_two_point`, with the general-observable form `rect_ising_uniform_decay`. The quantifier order is the theorem: β, γ, α and the prefactor are bound before (L, T) , so a reader asking “uniform in what?” is answered by the statement’s own binder list.

Proof. Each site of each rectangle carries at most two horizontal and two vertical bonds: the horizontal neighbour set injects, through the first coordinate’s value, into the two-element set $\{a-1, a+1\} \subset \mathbb{N}$, and likewise vertically. So the envelope row sum is at most $2 \tanh |\beta| + 2 \tanh |\gamma|$ at every site — interior sites attain it, boundary sites sit strictly below — and Theorem 5.5’s endpoints apply with the given α and the range-one support of the coupling for $d_\square$. $\square$

Proposition 5.9 (non-vacuity across the family, machine-checked). *At $\beta = \gamma = 0$ the window degenerates to $\alpha = 0$ and Theorem 5.8 forces, for distinct sites, $\text{Cov}_\mu(f, g) = 0$ exactly, at every volume — true, since the zero-coupling weight is a product measure, and proof that the family statement carries content at every (L, T) simultaneously. In Lean: `rect_zero_coupling_indep`.*

Gates. `judge_dobrushin_d5.py`, registered before this rung’s Lean. G15 measures the row sums against the window on four lattices with interior attainment; G16 checks the bound itself against *exhaustive* Gibbs covariances — full enumeration, no sampling — over every ordered pair of distinct sites; G17 computes the bound once and holds four growing volumes under it.

6 The transport into the operator, and how much of it is now closed

The frontier section of an earlier version of this paper named one open step: the transport of the measure-level decay into the operator formulation. All four of its planks are now machine-checked theorems, in `DobrushinBridge.lean`, `DobrushinTransport.lean`, `DobrushinTilt.lean` and `DobrushinCorollary.lean`; this section states them at their exact size, ending at the corollary the chain was for.

6.1 The band identity

Theorem 6.1 (the finite band identity; machine-checked). *Let X be a finite nonempty type, M a symmetric real $X \times X$ matrix, and $\Omega : X \rightarrow \mathbb{R}$ strictly positive with $M\Omega = \Omega$ (the normalised eigen-relation; `normalise_eig` shows how a raw Perron pair (T, λ, Ω) produces it, so the eigenvalue is visible rather than hidden in a convention). Define the band weight on paths $v : \{0, \dots, n\} \rightarrow X$ by*

$$W_n(v) = \Omega(v_0) \left(\prod_{k < n} M_{v_k v_{k+1}} \right) \Omega(v_n).$$

Then, exactly and for every n : the total mass is $\sum_v W_n(v) = \sum_x \Omega(x)^2$, independent of n (`bandZ`); two-endpoint sums against f, g are the matrix elements $\sum_{ij} (f\Omega)_i (M^n)_{ij} (g\Omega)_j$ (`band_pair`); and the band covariance of $f(v_0)$ and $g(v_n)$ equals

$$\frac{\langle (f - \mathbb{E}f)\Omega, M^n (g - \mathbb{E}g)\Omega \rangle}{\sum_x \Omega(x)^2}, \quad \mathbb{E}f = \frac{\sum_x f(x)\Omega(x)^2}{\sum_x \Omega(x)^2}$$

(`band_covariance_eq`). *Nothing assumes $0 \leq M$ as a form: positive entries do not give spectral positivity, subdominant eigenvalues may be negative, and the statement never notices — writing it so that it would was a failure mode named in the rung’s charter before any Lean existed.*

6.2 The abstract transport theorem

Theorem 6.2 (uniform band decay transports to a uniform gap; machine-checked). *Let ι index a family of finite nonempty slice types X_i , symmetric kernels M_i and strictly positive Ω_i with $M_i\Omega_i = \Omega_i$. Suppose there is one rate $0 < r < 1$ such that for every i and every observable $f : X_i \rightarrow \mathbb{R}$ there is a constant C with*

$$|\text{bandCov}_i(f, f, n)| \leq C r^n \quad \text{for all } n,$$

where `bandCov` is the measure-level band covariance of Theorem 6.1 — a finite path sum; no operator, no norm and no spectrum occurs in the hypothesis. Then there is one $m > 0$ with

$$\|\text{projectedTransfer } T_i \hat{\Omega}_i\| \leq e^{-m} \quad \text{for every } i$$

for the packaged operators T_i on $\ell^2(X_i)$ with unit vacua $\hat{\Omega}_i = \Omega_i / \|\Omega_i\|$. In Lean: `abstract_uniform_gap`.

Two sentences of the design are load-bearing. *First*, the anti-circularity of the hypothesis: an “abstract” statement of the shape “positive weight + symmetric matrix $\Rightarrow$ gap” could come out true by assuming the spectral conclusion in other notation. The hypothesis above quantifies over path sums of the band weight; its identification with operator matrix elements is the exact identity

$$\text{connCorr } T \hat{\Omega} v n = \left(\sum_x \Omega(x)^2 \right) \cdot \text{bandCov}(v/\Omega, v/\Omega, n)$$

(`connCorr_eq_bandCov`), proved for *every* vector v and every n through Theorem 6.1 — a theorem, not an assumption, which is what makes the interface reusable rather than circular. *Second*, the quantifier order: only r is shared; the constant may depend on the family member and the observable, and the representation $f_v = v/\Omega$ through Perron positivity is precisely where that freedom is spent. **No Ising, no Dobrushin and no rectangle occurs in the statement.**

Proposition 6.3 (non-vacuity of the transport; machine-checked). *For every rate $0 < r < 1$ the two-state kernel with diagonal $(1 + r)/2$ and off-diagonal $(1 - r)/2$, with constant boundary vector, satisfies every hypothesis of Theorem 6.2 — its band covariance is exactly $((f_0 - f_1)/2)^2 r^n$ (`wKernel_bandCov`) — and its projected operator is not the zero operator (`wKernel_fluctuation_ne, transport_witness`). The theorem is inhabited at every rate by an instance with a nonzero fluctuation sector; it is not satisfied only by degenerate data.*

6.3 The tilt layer: the boundary cost, absorbed

The physical kernel’s band measure and the free strip measure of the rectangle family differ at exactly two slices. The next theorem makes the difference an identity rather than an estimate.

Theorem 6.4 (the tilt layer; machine-checked). *Let w be a strictly positive slice weight, β a coupling, and (λ, Ω) raw Perron data for the symmetrised kernel $T = \sqrt{w} K_\beta \sqrt{w}$, normalised as in Theorem 6.1. Put $\psi = \Omega/\sqrt{w}$. Then, exactly:*

$$W_n(v) = \frac{\psi(v_0) \mu_\beta^{\text{free}}(v) \psi(v_n)}{\lambda^n}$$

where μ_β^{free} is the free strip Gibbs weight of the development’s spatial system (`bandW_eq_tilt`); the band mass and both band marginals are the corresponding tilted free path sums (`mass_mul_pow, marginal_left_mul_pow, marginal_right_mul_pow`); the band covariance satisfies the division-free identity

$$\text{bandCov} \cdot \text{PS}(\psi, \psi)^2 = \text{PS}(f\psi, g\psi) \text{PS}(\psi, \psi) - \text{PS}(f\psi, \psi) \text{PS}(\psi, g\psi)$$

with PS the two-endpoint free path sum (`bandCov_mul_sq`); and a pointwise floor $c \leq \psi$ gives $c^2 Z \leq \text{PS}(\psi, \psi)$ uniformly in the horizon, with no decay input (`floor_PS_tilt`). Consequently (`bandCov_decay_of_free_decay`): if the free two-endpoint covariances of the strip decay at rate α with per-observable constants of product form, then the band covariances decay at the same rate, with a constant depending on the extent and the observables only — never on n .

The factor is $\psi = \Omega/\sqrt{w}$ and not $\Omega\sqrt{w}$; the two dressings differ by exactly the square root the transfer kernel absorbs, conflating them silently was the second failure mode the charter named in advance, and the fabricating desk’s own first scratch derivation committed it — caught by the derivation itself and then measured by the gates, which is what pre-registered gates are for.

6.4 The last plank, closed

An earlier version of this paper carried, at this spot, the one discharge separating Theorems 6.2 and 6.4 from a volume-uniform operator statement for the coupled kernel: supplying the free-decay hypothesis of Theorem 6.4, for the physical kernel’s strip, from Theorem 5.8. It is now machine-checked, in [DobrushinCorollary.lean](#), exactly along the two steps that version named. *The currying identity* (`curry_weight`): the free strip Gibbs weight IS the rectangle Ising weight under the currying bijection between $\{0, \dots, n\} \times \{\text{sites}\}$ configurations and slice paths — an exact weight identity, proved through a one-dimensional nearest-neighbour pairing (`adj_pairing`: the ordered double sum over adjacent pairs is twice the sum over steps) and the energy split (`rectJ_energy_split`) in which the half carried by the development’s Ising weight compensates ordered double counting exactly, the couplings landing as β on the time axis and γ

on the space axis. Hence the tilt layer’s free covariance IS the rectangle covariance of end-slice observables (`freeCov_eq_rect_covar`). *The feed* (`rect_feed`): end-slice observables oscillate only on their slices (`row_deltaAt_zero`, `row_deltaAt_le`), every site pair the covariance sees is at rectangle distance at least the horizon, so Theorem 5.8’s $\alpha^{d_\square}$ becomes the needed α^n with slice-level oscillation-sum constants — the growth with the extent exactly where the consumer’s quantifier order permits it.

Theorem 6.5 (the Dobrushin–Ising corollary; machine-checked). *Fix β, γ and α with $0 < \alpha < 1$ and $2 \tanh |\beta| + 2 \tanh |\gamma| \leq \alpha$. Then there is one $m > 0$ such that for every extent L there are normalised Perron data (λ_L, Ω_L) for the coupled kernel $T_L = \sqrt{w_\gamma} K_\beta \sqrt{w_\gamma}$ — strictly positive Ω_L fixed by the normalised kernel — with*

$$\|\text{projectedTransfer } \hat{T}_L \hat{\Omega}_L\| \leq e^{-m}.$$

In Lean: `dobrushin_ising_uniform_gap`, through Theorem 6.2; the constant is $m = -\log \alpha$. For the symmetric normalised kernel this operator norm is the modulus of the fluctuation sector, so the statement is the volume-uniform separation the development’s frontier had carried as open — obtained from the measure side, with no spectral input anywhere in the chain.

Gate G23 had assembled this exact bound and held the measured band covariance under it before any Lean of the layer existed; the theorem is the gate’s shadow made checkable. Ising certifies the abstract bridge of Theorem 6.2 is inhabited by the physical family; the bridge itself never mentions it, which is what makes it reusable.

Gates. Two scripts, both registered and passed *before* the Lean they license. `judge_dobrushin_d6.py` (G18–G20, 60 of 60 in both interpreter modes): the strip identity measured exactly at 10^{-10} ; the target’s numeric shadow — `specRatio` under the window on every in-window cell at $L = 2..8$, the out-of-window control reported unjudged; per-extent prefactor convergence. `judge_dobrushin_d6b.py` (G21–G23, 142 of 142 in both modes): the identity of `connCorr_eq_bandCov` on registered matrices *including one with a negative subdominant eigenvalue*, so spectral positivity cannot be smuggled and pass unnoticed; the tilt algebra exact with $\psi = \Omega/\sqrt{w}$; and the assembled per-extent bound of the previous paragraph dominating the measured band covariance.

7 The measurement, and where the wall is

This section is numerical. Nothing in it is proved, and nothing in it is used by Section 3.

Let $T_L = W_\gamma^{1/2} K_\beta W_\gamma^{1/2}$ be the symmetrised coupled kernel on $\{\pm 1\}^L$ with free spatial boundary, and $r(L) = |\lambda_1|/|\lambda_0|$ its second-eigenvalue ratio, which is `specRatio` of the development. We compute $r(L)$ for $L = 2, \dots, 12$ by dense symmetric eigendecomposition and extrapolate by Aitken.

β	γ	$\sinh 2\beta \sinh 2\gamma$	$r(2)$	$r(12)$	$r(\infty)$
<i>controls at zero spatial coupling</i>					
0.20	0.00	0.0000	0.197375	0.197375	$\tanh \beta = 0.197375$
0.40	0.00	0.0000	0.379949	0.379949	$\tanh \beta = 0.379949$
<i>disordered</i>					
0.10	0.10	0.0405	0.1096	0.1209	0.1214
0.20	0.20	0.1687	0.2360	0.2895	0.2922
0.30	0.30	0.4053	0.3733	0.5127	0.5230
0.40	0.40	0.7887	0.5122	0.7813	0.8162
0.10	0.80	0.4783	0.1651	0.4013	0.4570
0.80	0.10	0.4783	0.7246	0.8031	0.8075
<i>ordered</i>					
0.60	0.60	2.2785	0.7489	0.9992	1.000003
0.80	0.80	5.6433	0.8893	0.999999	1.000000

$r(\infty)$ is an Aitken extrapolation of the last three extents. Values above one in the ordered rows are extrapolation overshoot at the sixth decimal and are reported as measured.

Each of the six cells tested inside Onsager’s disordered region, $\sinh 2\beta \sinh 2\gamma < 1$, saturates strictly below one, with increment ratios near 0.6; both ordered cells tested reach one to six decimal places. The control at $\gamma = 0$ reproduces $\tanh \beta$ to six decimals, which is Proposition 2.1 — the harness is calibrated against a theorem of the development rather than against itself.

What these data support, stated at their strength and no further. They are consistent with the degeneracy being a property of the *ordered* phase, and they are *not* consistent with the degeneracy being caused by switching on the spatial weight: six disordered cells with the weight switched on fail to degenerate. That second, negative reading is the one this section is for, and it is the one that reopened the question.

What they do not support. They do not locate the transition. Eight cells at $L \leq 12$, with a three-term Aitken extrapolation whose overshoot is visible in the ordered rows above, cannot place a phase boundary; doing so would require a dense sweep near $\sinh 2\beta \sinh 2\gamma = 1$, finite-size scaling, and control of the extrapolator, none of which is attempted here. A reader should take from this section that the wall is *not* where the development had been treating it as being, and nothing about where it is.

Gates. Three gates were registered in [scripts/judge_dobrushin.py](#) together with the lane charter, before a line of the Lean of Section 3 was written. Each exits non-zero on failure; none bundles a theorem with its witness; the bound gate predicts a number rather than sampling. One of them was later found not to exercise the hypothesis of Remark 3.5 — it normalised every row to *exactly* α — and was redesigned to carry unequal row sums, strictly harder, declared in its own docstring and committed before being run. The gates emit a verdict only after an explicit counter confirms that the expected number of checks actually executed, because an empty failure list is also what zero checks produce. No acceptance decision anywhere in this lane depends on a Python `assert`.

8 What is now closed, and what remains

Theorems 3.1–5.2 are an engine, not a conclusion. Turning them into a volume-uniform statement about specRatio requires, in order:

1. a valid single-site interdependence majorant, and the sufficient window it produces. **Done:** Theorems 3.1 and 3.2, machine-checked, with the window computed rather than asserted, and holding at every site of every volume with the given local geometry. Not needed, and done only where the one-bond witness makes it coincide with the envelope: the *minimal* interdependence matrix of a general system;
2. Dobrushin’s comparison estimate [2, 5] — that connected correlations are controlled by the resolvent of C , hence decay exponentially at a rate free of the extent. **Done:** Theorem 4.3;
3. construction of the single-site conditionals of a Gibbs weight, and discharge of every hypothesis that estimate consumes — invariance above all, which is a genuine constraint on the pair and not a formality. **Done:** Theorem 5.2, from a strictly positive weight alone, leaving Dobrushin’s condition on the intrinsic matrix as the hypothesis a user supplies — plus, for the geometric form of the decay, the range-one support of that matrix; Theorem 5.5 instantiates the interaction itself, turning that hypothesis into a computation on the coupling, discharged on the one-bond and star cells; and Theorem 5.8 carries it across the rectangle family, constants first, volumes after;
4. transport of that decay into the operator formulation: from correlations of the Gibbs weight to specRatio of the transfer kernel, through the finite-time interface described below. **Three of four planks done:** the finite band identity (Theorem 6.1), the abstract transport theorem with its inhabited witness (Theorem 6.2, Proposition 6.3), and the tilt layer that absorbs the whole boundary cost into a per-extent constant (Theorem 6.4) — fabricated under the committed charter and its two amendments ([DOBRUSHIN-D6-CHARTER.md](#), [DOBRUSHIN-D6-B2-DESIGN.md](#)), with gates G18–G23 passed before each Lean layer existed. **And now closed entirely:** the currying identification of the strip sums with Theorem 5.8’s rectangle covariances and the resulting instantiation of Theorem 6.2 for the coupled family are machine-checked (Section 6.4, Theorem 6.5). **Done**, all four planks.

What changed, stated at its exact size. An earlier version of this paper named step 2 as the only currently identified new analytic estimate and the bottleneck of the chain; a later one proved it, for arbitrary invariant data, and left step 3 open; the version before this one closed both and described step 4 as untouched; its successor closed three of step 4’s four planks and named the fourth. Step 4 is now closed. The finite band identity, the abstract transport theorem, the Perron-boundary tilt, the currying identification and the rectangle-family feed compose to Theorem 6.5, a volume-uniform projected-operator gap throughout the Dobrushin window. The translation *did* run through finite time extent with explicit boundary control, exactly as the design decision below prescribed: the band identity is finite and exact at every horizon, the boundary control is the tilt layer’s floor and sup bounds, and the limit is taken nowhere — the conclusion is an operator-norm bound obtained from finite covariances. What this paper claims is that between a positive weight and exponential decay of its correlations there is no gap, that between uniform band decay and a uniform operator gap there is no gap either, and that for the coupled family the two chains are connected. What remains lies beyond this finite-volume

transport: no infinite-volume state, thermodynamic limit or boundary-condition independence is constructed here.

The consumer, named. It is `volumeUniform_gap` in [YangMills/OS/TransferGap.lean](#), and its signature, read off the elaborated declaration rather than from memory, is: for a family of real inner-product spaces H_i indexed by ι , transfer operators T_i and vectors Ω_i with `VacuumTransfer` $T_i \Omega_i$ for every i , and a rate $0 < r < 1$, the hypothesis

$$\forall i, \forall v \in H_i, \exists C, \forall n, \quad |\text{connCorr } T_i \Omega_i v n| \leq C r^n$$

yields $\exists m > 0$ with $\|\text{projectedTransfer } T_i \Omega_i\| \leq e^{-m}$ for every i . It takes, over a family of finite volumes, decay at a *common* rate and returns a *common* mass. Its quantifier order is worth stating because it is load-bearing: the prefactor C may depend on the volume and on the observable, and only r is shared. The development’s own standing objection — that a bound $C(L)\rho^N$ with $C(L) \rightarrow \infty$ says nothing — therefore does not apply to it, the conclusion being an operator-norm bound.

A design decision is recorded here rather than discovered later: the vacuum pairing is an infinite-time object, so step 3 must proceed through finite time extent with explicit boundary control, taking the limit only at the level of the bound. No infinite-volume state is constructed anywhere in this paper, and the word *clustering* is accordingly not used for one.

9 Prior art, and the exact delta

Dobrushin’s uniqueness theorem is classical [1, 2, 5] and nothing in this paper is new mathematics. We locate the contribution precisely.

Not explicit constants. The cell here is the anisotropic two-dimensional Ising model, whose transfer spectrum has been available in closed form since Onsager [6] and Kaufman [7]. A constant obtained from Dobrushin’s condition is *worse* than the exact answer — on the cells we measured, by a factor between roughly two and thirty-six. Any claim to have improved a constant for this model would be false, and we make none.

Generality, instead. Dobrushin’s coefficient is read off the single-site conditionals of an *arbitrary* positive finite-range slice weight, which is what the development’s Gibbs weight quantifies over, and where the exact solution is silent.

Mechanisation. A search conducted on 1 August 2026 over the public proof-assistant repositories and the arXiv listings for Lean, Coq and Isabelle returned no mechanisation of the finite-lattice Dobrushin interdependence matrix or of the comparison estimate. We state it that narrowly on purpose: the name *Dobrushin* attaches to several distinct results, and mechanised work exists that uses others of them, so a blanket claim of first formalisation would be both unsupported by the search we ran and ambiguous about which theorem it referred to. The same search does return a Lean 4 formalisation of the transverse-field Ising chain [11], with momentum-mode decomposition and a Jordan–Wigner transformation; it concerns the quantum chain rather than the classical two-dimensional transfer matrix, and displaces nothing here, but it means a blanket claim that the Ising model has not been formalised would not survive review, and we make none. A formalisation of the Perron–Frobenius theorem for irreducible nonnegative matrices also exists [10], as does a mechanisation of the Osterwalder–Schrader axioms for the Gaussian free field [12]. Reflection positivity and finite-lattice spectra for related models have been treated without mechanisation [13].

10 Limitations

All twelve theorems are elementary. Theorem 3.3 is three inductions; Theorem 3.1 is a hyperbolic identity and a bound on a denominator; Theorem 3.2 is a construction and a finite computation; Theorem 4.3 is a telescoping sum, one pointwise application of a ninety-year-old covariance bound, and two finite inductions; Theorem 5.2 is a reindexing by an involution; Theorem 5.5 is a split of a quadratic sum at one index and the envelope applied to a field move of 0 or ± 2 ; Theorem 5.8 is a neighbour count by injection; Theorem 6.1 is one path-splitting induction and the eigen-relation applied twice; Theorem 6.2 is an exact change of description fed to the development’s operator interface, whose uniformisation is Banach–Steinhaus and was proved there, not here; and Theorem 6.4 is division-free algebra under a pointwise floor and sup bounds; and Theorem 6.5 is a reindexing by currying, a nearest-neighbour pairing count, and the feed of one theorem into another. The family is rectangles with free boundary, not tori. The transport is closed: four machine-checked planks ending at the corollary, the only spectral content being the symmetric-operator reading of an operator-norm bound. Their value is that they are machine-checked, non-vacuous, composed, and carry the hypothesis the dead method violated — not that any of them is deep.

Section 7 is numerics at $L \leq 12$ with an Aitken extrapolation. It provides measured counterevidence against an attribution and contradicts it on the tested grid; it does not prove where the wall is, and an extrapolation can mislead near a transition. The window is not optimal and is never called sharp.

The substantive limitation is now the infinite volume. The volume-uniform operator bound of Theorem 6.5 holds at every finite extent with one constant; no infinite-volume state, no thermodynamic limit and no boundary-condition independence is constructed anywhere, and within-volume decay does not build them — that is a separate programme with its own charter, and nothing here licenses it. Two further boundaries are worth naming. Theorem 5.2 discharges every hypothesis *except* Dobrushin’s condition; Theorem 5.8 discharges it uniformly for the rectangle family inside the window, and for other geometries — tori, higher dimension — it is the same neighbour count, to be repeated rather than presumed. The envelope of Theorem 3.1 is what a user would apply, and it is a sufficient condition, not a characterisation. And the decay proved is decay of correlations of a measure on a *finite* product space: no infinite-volume state is constructed anywhere, and the word *clustering* is accordingly not used for one.

11 Reproducibility

Anchor. All sources cited above are read at commit `c7b870b05bd6d1bec270d499e1358a64070ce0e1` of <https://github.com/lluiseiriksson/THE-ERIKSSON-PROGRAMME> — the anchor at which the corollary of Section 6.4 was verified; the family rung’s record at `3bd8a463d647d2ef4ea8afd44c19855244729705` and the transport rungs’ own verification anchors (`b4469a797a41ed3c053ca30f4e6a5ff7edddbb85` for the band identity, `9088eea6cdbbbe67a37dc37496581801925154f5` for the abstract transport theorem, `082e89b1a6bc37136b44037086d9620adee3abc2` for the tilt layer) stand beneath it, every lane module byte-identical across them except where a rung’s own passes are itemised below. Toolchain `leanprover/lean4:v4.29.0-rc6`; Mathlib pinned at `07642720480157414db592fa85b626dafb71355b`. A runner verifies all three of these before compiling anything and aborts on any mismatch, so that a stale cache cannot masquerade

as a reproduction.

Job counts, and how each delta was obtained. The development machine measured the core at 8465 jobs with Theorem 3.3’s module removed from the root and 8466 with it restored — that delta by rebuilding the same tree twice, not by subtracting against a baseline someone else measured. A fresh Linux clone reproduced 8466 at that commit. With the modules of Theorems 3.1 and 3.2 added, the Linux plane measures 8468 with zero errors: the +2 is the difference between two measurements on the same plane, one job per new module, and is not an extrapolation.

The modules of Sections 4 and 5, and what is measured about them. All seven — oscillation, Grüss, conditional, comparison, Gibbs, Ising, lattice — are imported by `YangMillsCore.lean` at this anchor, and their endpoints are listed by name in `oracle_check.lean`. *What has been measured now includes the full core, in a fresh clone.* The job accounting is stated in full because part of it was wrong and was corrected *before* the number was taken: the wiring commit registered 8468 $\rightarrow$ 8473, using this lane’s own D-2 count as the base, but the branch base already carries a sibling paper’s module measured at 8469, so the corrected expectation with six new modules — registered before any measurement — was **8475**. A *fresh* Linux clone (new runtime, `git clone` from scratch, toolchain installed by `elan`, Mathlib artefacts from `cache get` and nothing else reused) then built this paper’s two lane modules and their entire chain green, and the core build completed with lake’s own closing line — *Build completed successfully (8475 jobs)* — and zero errors: the predicted number, taken. One honesty note travels with it. At the fresh clone’s first commit a module belonging to a *different* paper’s campaign (`SpatialReconstruction.lean`, material committed without compilation) still failed; its seven defects were repaired across two commits — five dead closing tactics deleted, one pin-absent lemma name re-derived through `add_right_cancel`, one `simp` replaced by the definitional term it was fighting for — with no statement changed anywhere, and the mathematical audit of that lane stays with its author. The green count was then taken in the same clone after the final one-line repair, with every module of this paper byte-identical throughout, as the hash table below witnesses. The family rung then repeated the discipline in a *second* fresh clone: the prediction 8475 $\rightarrow$ 8476 registered in its wiring commit, two elaboration passes — the only defects two lambdas whose domain the elaborator could not yet see — and lake’s closing line at this anchor, *Build completed successfully (8476 jobs)*: the number again exactly as registered before the run. A previous version of this paper carried, at this spot, the opposite sentence — that the modules were not yet imported — left over from the draft that preceded the wiring; its correction had been appended past the end of the document environment and so existed in the source while never reaching the PDF. The paragraph now standing is the one the anchor’s bytes support, and the defect is recorded because a contradiction a reader finds is worth more than one an author hides. The four transport rungs then extended the run of registered-then-taken counts to seven: 8476 $\rightarrow$ 8477 for the band identity, 8477 $\rightarrow$ 8478 for the abstract transport theorem, 8478 $\rightarrow$ 8479 for the tilt layer, and 8479 $\rightarrow$ **8480** for the corollary — each prediction in the rung’s wiring commit, each fresh Linux clone printing lake’s closing line with exactly the registered number (*Build completed successfully (8477/8478/8479/8480 jobs)*). One operational honesty note travels with the last: the corollary’s first verification queue died with a client session mid-flight (a queued Colab cell is submitted by the browser, not the server), so its pass sequence includes one rerun of an identical commit on a fresh runtime; the counts and hashes above are from the completed run.

Oracle. For Theorems 3.1–3.3, at an earlier anchor of this lane: the repository-wide or-

acle ran to completion with exit code zero over 2817 axiom reports and **zero** occurrences of **sorryAx**, thirty-seven of them this chain's. For Theorems 4.3, 5.2 and 5.5: three *targeted* oracle runs at their lane anchors (the third at this anchor), exit code zero, printing the axioms of seventeen, eleven and twelve endpoints respectively — among them the three clauses of Lemma 4.1, both Grüss forms with the Grüss attainment, the transport domination, the finite- N and series estimates, the exponential-decay composition and the two-point corollary; on the Gibbs side invariance, the zero diagonal, domination, the three instantiated endpoints and both witnesses; and on the Ising side the energy split, the sigmoid identity, the envelope domination, the three instantiated endpoints, the one-bond attainment and unconditional decay, and the star window theorem. The union of every axiom list printed, over the three runs, is exactly **{propext, Classical.choice, Quot.sound}**: Lean's three standard axioms, no project axiom, no **sorry** in any build log. And the *repository-wide* oracle then ran to completion in a fresh clone at each successive rung: 2985 axiom reports at the family anchor, 2993 at the band identity's (+8, its endpoints), 2999 at the abstract transport theorem's (+6), 3004 at the tilt layer's (+5), and **3010** at this paper's anchor (+6, the corollary's endpoints) — each count the one registered in the rung's wiring commit before the run, each with exit code zero, **zero** occurrences of **sorryAx**, and the union of every axiom list printed exactly the standard triple, the wrapped log lines read to their continuation lines rather than truncated at the bracket. (The earlier runs, 2980 and 2985 reports with the same union, stand in the record.) The union is computed as a *union* rather than by matching each printed line against the standard triple, because a declaration depending on *fewer* axioms is perfectly clean and would fail that naive test.

One gap in the runner's own coverage, stated because a reader counting modules would otherwise not find it. The runner builds two lane modules by name; [DobrushinRowSum.lean](#) is not one of them. It is nevertheless imported by [DobrushinIsing.lean](#) at this anchor, so the zero exit of the *lane* build covers it — and the core build's zero exit at this anchor now covers it a second time.

Independent hash agreement. A second desk, working from its own check-out on a different operating system and without coordination on the value, reported 12c1fcdb12e39e46381fa372a102503f5d63522ecbae27c0b3911671a7c04eb7 at 11467 bytes for Theorem 3.3's module. The Linux run computed the same digest and the same size. The oleans, the axiom reports and the elaboration therefore refer to the same bytes, which is the only reason any of the numbers above can be compared at all.

The bytes of Sections 4 and 5. Printed in the same executions that built and oracled them:

module	sha256	bytes
DobrushinOscillation.lean	ca7c6842b8e08d13c405517994df86f868047b48c53fb9ab2e24fcea10777d89	7777
DobrushinGruss.lean	8e00f418a42fb30d6f6b367d5e353b4f34149c2e70d3625ae3307ce8bceb9ace	16462
DobrushinConditional.lean	eb52b816f2c401d17edb72c5a7d0f8d413c33bbf9292f770e92a787fc8f5833a	18270
DobrushinComparison.lean	a489fc2823e564685f3ffbf586f1a1f63ce4b78a955a9ec88edbcd70afae4133	54738
DobrushinGibbs.lean	b525731d274bf0ce7511e0e0f796d5892f22cee1601daa5a8fa7fea6bd507ba1	16787
DobrushinIsing.lean	4e880ff013508c55bbbe43db0a511bcecfb51f509a8ebc78f9d49201f577bf23	29158
DobrushinLattice.lean	78b57a5bed869378557c2a1a9c4442be1ad41207e08a6607999365926422d1e7	15135
DobrushinBridge.lean	d2db5202110baade1663edf5cc3baab5c748c68a3215cbbc9f057e6f092bec38	16594
DobrushinTransport.lean	82a783544cf60a7672c49bfeb32102eb85a3587a9d7fc2f60cb44b125505abb0	18467
DobrushinTilt.lean	a14bb43eee5503e93b2e55ae71321df13dcb863014b6db98d23341fd84aa4d05	31280
DobrushinCorollary.lean	ffdc5a6a910211ec34678007625d6eb69a2bc50dd8fa3539ee262deec71f5a0	34006

Certifiers. Eight gate scripts, each run in `normal` and in `optimized` (`python -O`) mode on the Linux plane; all sixteen runs exited zero. `judge_dobrushin.py` reported 29 of 29 checks in each mode and `judge_dobrushin_d2.py` 18 of 18, at the previous anchor. At this anchor `judge_dobrushin_d3.py` reported 11 of 11, `judge_dobrushin_d3b.py` 6 of 6, and `judge_dobrushin_d4.py` — the gates G12–G14 of Section 5 — 14 of 14, and `judge_dobrushin_d5.py` — G15–G17 of the family rung — 37 of 37, `judge_dobrushin_d6.py` — G18–G20 of the transport — 60 of 60, and `judge_dobrushin_d6b.py` — G21–G23 of the tilt layer — 142 of 142, in both modes, all running *before* the build in the same runner invocation and, for every rung from D-4 onward, before *every* fabrication pass of the rung they license — the record shows the gates passing before the theorems they license were compiled. No acceptance decision in any of the eight depends on a Python `assert`; each counts the checks it performed and refuses to print a verdict when fewer ran than expected, because an empty failure list is also what zero checks produce.

What the passes cost, and the one that was not a repair. Seven elaboration rounds on the Linux plane closed Section 4’s four modules; two closed the Gibbs module; three closed the Ising module, whose first pass elaborated the entire mathematical chain — energy split, sigmoid identity, envelope domination, endpoints, attainment witness — at the first attempt, failing only on a missing import, two dead closing tactics, one no-progress `simp`, and four rewrites that could not see through `fin_cases`’ literal representation (an `exact` bridges it by definitional equality; a rewrite does not); and two closed the family module, whose only defects were two bare lambdas whose projections the elaborator met before their domain. The transport rungs kept the shape: two passes closed the band identity, whose single error was a higher-order rewrite unable to match a β -redex; two closed the abstract transport module, whose sections 1–4 — the identity and the theorem — elaborated clean on the first pass, the thirteen errors all in the two-state witness (one under-expanded `Fin-sum` script, ten universe-inference failures from a single unpinned `PUnit`); and three closed the tilt layer, whose thirteen errors across two passes were two closing tactics dead behind `field_simp`, one associativity mismatch, one normaliser choking on sum-sized atoms (cured by proving the five-variable identity abstractly first), two product-order slips, one under-determined `le_sup`’, and eight `positivity` calls against atoms the extension cannot see (cured by explicit `mul_nonneg` chains). Six passes closed the corollary, its error ladder $20 \rightarrow 9 \rightarrow 5 \rightarrow 4 \rightarrow 2 \rightarrow 0$ and every rung machinery: library names absent under their classical spelling in the pinned Mathlib; a `have` that made the currying equivalence *opaque*, killing a definitional `rfl` (data must be `let`-bound or inlined); a typeclass search stuck on the coupling’s metavariable extents until they were pinned by name; four rewrites blind to β -redexes until a `show` exposed the reduced form; a summation lemma used in the wrong direction; and five closers that died as the chains beneath them became exact. The mechanism — the pairing count, the halved double energy, the currying identity, the feed — is the one written in the first draft and was never altered. Not one of the failures was mathematical: they were library names absent under their classical spelling in the pinned Mathlib, section variables that a statement never mentioned and so never bound, a global rewrite that consumed the literal 1 inside an indicator, and goals equal only after β -reduction. The exception is worth its sentence. The first round failed in `DobrushinOscillation.lean`, a module committed at an *earlier* anchor: the file had been verified in its pre-split working form and split afterwards, and the committed bytes had never been elaborated. The lane had declared that debt in advance and the runner was built to re-verify those bytes; it found the defect on its first run. *A module that was green before an edit is not green after it, and the hash is what says which one was compiled.*

Child processes. Every child of the runner writes a sentinel holding exactly one line, its real decimal exit code, captured after termination and written atomically with validation. Four states are distinguished on read: absent, malformed, non-zero, zero. A run in which every sentinel reads zero is reported as a *pass candidate*, never as a pass, because the logs still have to be read.

Every module in the anchor compiles — this paper’s and, after the cross-desk repair the job-count paragraph records, the sibling’s too. An earlier version of this paper was anchored to a tree in which `DobrushinCoefficient.lean` did not, and said so in this section rather than leaving it to be discovered. It took five passes to close, and the record of those passes is worth one sentence because it is the honest description of what mechanising a classical argument costs: none of the fifteen errors across them was mathematical. The mechanism — the identity of Theorem 3.1 and the denominator bound under it — is the one written in the first draft and was never altered. Every failure was a library lemma whose name or form had been recalled wrongly, or a tactic emitting a different number of goals than the script assumed.

References

- [1] R. L. Dobrushin, *The description of a random field by means of conditional probabilities and conditions of its regularity*, Theory of Probability and its Applications **13** (1968), 197–224. The uniqueness condition and the comparison estimate whose mechanised form is Theorem 4.3. The theorems of this paper are classical and due here.
- [2] R. L. Dobrushin, *Prescribing a system of random variables by conditional distributions*, Theory of Probability and its Applications **15** (1970), 458–486. Together with [5], the source of the argument mechanised in Section 4: the single-site update acting on oscillation vectors, and the resolvent of the interdependence matrix as the resulting comparison.
- [3] G. Grüss, *Über das Maximum des absoluten Betrages von $\frac{1}{b-a} \int_a^b f(x)g(x) dx - \frac{1}{(b-a)^2} \int_a^b f(x) dx \int_a^b g(x) dx$* , Mathematische Zeitschrift **39** (1935), 215–226. The covariance bound with the quarter constant, applied pointwise to the conditional in Lemma 4.2. Our proof is the discrete signed-mass one rather than Grüss’s, so that each factor $\frac{1}{2}$ is created in exactly one lemma.
- [4] T. Popoviciu, *Sur les équations algébriques ayant toutes leurs racines réelles*, Mathematica (Cluj) **9** (1935), 129–145. The variance bound $(M - m)^2/4$, machine-checked in the same module as Lemma 4.2 and deliberately *not* used by it.
- [5] B. Simon, *The Statistical Mechanics of Lattice Gases, Volume I*, Princeton University Press, 1993. Chapter V gives the Dobrushin uniqueness condition, the interdependence matrix, and the exponential decay of correlations it yields; the one-bond envelope $\tanh |J|$ for the Ising cell and the resulting window $2 \tanh \beta + 2 \tanh \gamma < 1$ are the standard specialisation.
- [6] L. Onsager, *Crystal statistics I: a two-dimensional model with an order-disorder transition*, Physical Review **65** (1944), 117–149. The exact solution, and the source of the line $\sinh 2\beta \sinh 2\gamma = 1$ used as the reference boundary in Section 7.
- [7] B. Kaufman, *Crystal statistics II: partition function evaluated by spinor analysis*, Physical Review **76** (1949), 1232–1243. The transfer-matrix spectrum in closed form. This is the

reason Section 9 claims generality rather than constants: for this cell the exact second eigenvalue is available and the Dobrushin constant is worse than it.

- [8] I. Schur, *Bemerkungen zur Theorie der beschränkten Bilinearformen mit unendlich vielen Veränderlichen*, Journal für die reine und angewandte Mathematik **140** (1911), 1–28. The test whose constant-row-sum form is the method discussed in Section 2.
- [9] E. Seneta, *Non-negative Matrices and Markov Chains*, revised printing, Springer, 2006. Standard reference for nonnegative matrices, ergodicity coefficients and the resolvent series $\sum_n C^n = (I - C)^{-1}$ of which Theorem 3.3 is a locally-supported, entrywise refinement.
- [10] The mathlib Community, *The Lean mathematical library*, Proceedings of CPP 2020, 367–381. The library against which the development is pinned; it carries a formalisation of the Perron–Frobenius theorem for irreducible nonnegative matrices.
- [11] *A Machine-Verified Proof of a Quantum-Optimization Conjecture*, arXiv:2606.29687 (2026). A Lean 4 formalisation of the Ising chain with momentum-mode decomposition and the Jordan–Wigner transformation. It concerns the transverse-field quantum chain rather than the classical two-dimensional transfer matrix, and displaces nothing here; it is cited because a blanket claim that the Ising model has not been formalised would be false.
- [12] *Formalization of Quantum Field Theory*, arXiv:2603.15770 (2026). Osterwalder–Schrader axioms for the Gaussian free field in Lean 4; the nearest mechanised neighbour on the constructive side, and free rather than interacting.
- [13] *Spectrum and reflection positivity on a finite lattice*, arXiv:2606.07922 (2026). Finite-lattice spectral and reflection-positivity results for a related model, not mechanised.