

Congruence Rigidity and the Fusion Bound: what a positive weight can and cannot do to the spectrum of a transfer kernel

Lluis Eriksson

Draft — 1 August 2026

Abstract

A positive site weight acts on a transfer kernel by *congruence*, $K \mapsto DKD$ with D positive diagonal, not by similarity. Similarity preserves the entire spectrum; congruence preserves strictly less, and strictly more than nothing. We determine both halves for the kernel of L decoupled Ising bonds. *Rigid half*: the sign of every quadratic form value survives, so definiteness in either sign is a congruence invariant — the definite case of Sylvester’s law of inertia, recorded here in the elementary form that a machine can check without an eigenvalue count. *Fragile half*: the value of the subdominant ratio r — second largest eigenvalue modulus over the Perron root — does not survive at all, and we locate exactly how far it moves. Restricting the L -site kernel ($L \geq 1$) to the two antipodal configurations leaves *one* Ising bond of coupling βL ; for $\beta > 0$ its ratio is $\tanh(\beta L)$. A weight concentrating on the ferromagnetic configurations therefore *fuses* the L sites into a single effective site carrying L times the coupling. Since $\tanh(\beta L) \rightarrow 1$ for every fixed $\beta > 0$, there is no bound $r \leq \rho < 1$ holding *simultaneously* in L and over the whole positive-diagonal congruence orbit — which is not the same as ruling one out for a particular physical family of weights. The obstruction to volume-uniformity is named, exactly, and it is a property of the congruence and not of any spectral estimate. Consequently an argument that bounds r using only congruence invariants — rank, inertia, definiteness — cannot produce an L -uniform bound at all. We then ask which congruence is extremal. For *uniform* concentrations the answer is proved: keeping m sites of mutual correlation μ leaves $\mu J + (1 - \mu)I$, of ratio $(1 - \mu)/(1 + (m - 1)\mu)$, which is strictly decreasing in m — the more sites the weight keeps, the smaller the ratio it can reach, so the best uniform concentration is a *pair*. For *arbitrary* weights we then prove $\sup_{D>0} r(DMD) = (1 - \mu)/(1 + \mu)$ with $\mu = \min_{i \neq j} M_{ij}$: the *least correlated pair* determines the supremum, which is approached by concentrating on such a pair. (We evaluate the supremum; we do not classify its maximisers.) The proof is geometric rather than spectral — Hilbert’s projective diameter is itself a congruence invariant, and Birkhoff’s contraction theorem converts it into the bound — and it uses no definiteness at all, only positivity of the entries. The hypercube instance is the fusion bound. Every new algebraic lemma is machine-checked in Lean 4, and the supremum theorem exists in the kernel in *conditional* form, with its one remaining *non-elementary* input — the Birkhoff spectral bound — carried as an explicit hypothesis rather than an axiom. The lower bound uses no limiting argument about spectra: a two-supported fluctuation vector gives the estimate at each strictly positive ε separately, after which only scalar arithmetic remains, so no continuity of eigenvalues is imported anywhere.

1 The question

Let M be a real symmetric matrix and D an invertible diagonal matrix with positive entries. The map

$$M \longmapsto DMD$$

is a *congruence*. It arises whenever a statistical-mechanical transfer kernel is symmetrised against a site weight: if the unnormalised kernel is $w(\sigma)K(\sigma, \tau)$, the symmetric form with the same spectrum is $\sqrt{w(\sigma)}K(\sigma, \tau)\sqrt{w(\tau)}$, that is DKD with $D = \text{diag}(\sqrt{w})$.

Congruence is not similarity. Similarity $M \mapsto SMS^{-1}$ preserves the characteristic polynomial and therefore everything spectral. Congruence preserves the quadratic form only up to a change of variable, and the change of variable is not an isometry. The natural question — and, we will argue, the one that governs whether a spectral estimate can be uniform in the system size — is:

Which spectral data survive a positive-diagonal congruence, and how far can the rest move?

The quantity we track. Throughout, the kernels are entrywise positive, so by Perron–Frobenius the spectral radius $\rho(T)$ is itself a simple eigenvalue. We write

$$r(T) := \frac{\max\{|\lambda| : \lambda \in \sigma(T), \lambda \neq \rho(T)\}}{\rho(T)},$$

the *subdominant ratio*. Since we allow indefinite M , this must be said explicitly: ordering eigenvalues algebraically, or by modulus, or by index gives different conventions once negative eigenvalues are present, and it is the quantity above — second largest *modulus* over the Perron root — that Fact 6.1 controls.

Remark 1.1 (A collision of notation with the source). Our $r(T)$ is the quantity Eveson and Nussbaum call the *spectral clearance* $q(L)$. Their $r(L)$ denotes the *spectral radius*, i.e. our $\rho(T)$ — so the letter r names different objects in the two papers. We keep r for the ratio because that is what this paper is about, and flag the collision here so that Fact 6.1 can be checked against the source without a false start.

We answer both parts for the kernel of L decoupled Ising bonds. Section 2 gives the rigid half, Section 3 the fragile half, Section 4 the consequence, Section 5 the extremal congruence among uniform concentrations, and Section 6 the supremum over all positive weights — exactly, by way of Hilbert’s projective diameter.

What this paper does not do. It does not decide whether any particular coupled transfer operator has a spectral gap uniform in its extension. It *explains* why such a question is hard by exhibiting the exact mechanism that destroys uniformity, and explaining is not resolving. No statement here has a consequence for gauge theory in any dimension, and none is claimed.

2 The rigid half: congruence cannot move a sign

Write $q_M(x) = \sum_{i,j} x_i M_{ij} x_j$ for the quadratic form of M .

Lemma 2.1 (Change of variable). *For every M , every d , and every x ,*

$$q_{\text{diag}(d)M\text{diag}(d)}(x) = q_M(dx), \quad (dx)_i := d_i x_i.$$

Proof. $(\text{diag}(d)M\text{diag}(d))_{ij} = d_i M_{ij} d_j$; substitute and reassociate. $\square$

Theorem 2.2 (Definiteness is a congruence invariant). *Let $d_i \neq 0$ for every i . Then $\text{diag}(d)M\text{diag}(d)$ is positive definite if and only if M is, and negative definite if and only if M is.*

Proof. By Lemma 2.1 the two forms have the same value set, because $x \mapsto dx$ is a bijection of $\mathbb{R}^n$ (its inverse is $x \mapsto d^{-1}x$) and carries nonzero vectors to nonzero vectors: if $d_i x_i = 0$ with $d_i \neq 0$ then $x_i = 0$. Apply this in both directions. $\square$

Theorem 2.2 is the definite case of Sylvester’s law of inertia [1], which says more: the full signature (n_+, n_0, n_-) is a congruence invariant. We do not reprove the general law. We isolate the definite case because it admits a proof through the form alone — no diagonalisation, no eigenvalue count — and that is what makes it cheap to certify mechanically.

Remark 2.3. Invertibility of D is not decorative. A diagonal with k zero entries drops the rank by k and can turn a definite form into a degenerate one. Every claim below assumes D strictly positive.

3 The fragile half: the fusion identity

Let $s(\sigma_j, \tau_j) = +1$ if $\sigma_j = \tau_j$ and -1 otherwise, and let

$$K_\beta^{(L)}(\sigma, \tau) = \exp\left(\beta \sum_{j=1}^L s(\sigma_j, \tau_j)\right), \quad \sigma, \tau \in \{0, 1\}^L,$$

the kernel of L decoupled Ising bonds at coupling β . Write $\sigma^+ = (0, \dots, 0)$ and $\sigma^- = (1, \dots, 1)$ for the two antipodal (ferromagnetic) configurations, and let

$$B(a) = \begin{pmatrix} e^a & e^{-a} \\ e^{-a} & e^a \end{pmatrix}$$

be the transfer matrix of a *single* Ising bond of coupling a .

Lemma 3.1 (One bond). *Let $a \geq 0$. Then $B(a)$ has eigenvectors $(1, 1)$ and $(1, -1)$ with eigenvalues $2 \cosh a$ and $2 \sinh a$, and its subdominant ratio is*

$$r(B(a)) = \frac{2 \sinh a}{2 \cosh a} = \tanh a.$$

Proof. Direct: $B(a)(1, 1)^\top = (e^a + e^{-a})(1, 1)^\top$ and $B(a)(1, -1)^\top = (e^a - e^{-a})(1, -1)^\top$. Here $2 \cosh a > 0$ is the Perron root, and $a \geq 0$ makes $2 \sinh a \geq 0$, so the other eigenvalue is already its own modulus. $\square$

Remark 3.2 (The sign hypothesis is the content). For $a < 0$ the statement is *false*: $2 \sinh a < 0$, so the subdominant modulus is $|2 \sinh a|$ and $r(B(a)) = |\tanh a|$. An earlier version of this paper omitted the hypothesis, and the formalisation did not catch it — the Lean lemma proves the *algebraic identity* $2 \sinh a / 2 \cosh a = \tanh a$, which holds for every real a , while the English sentence read that identity as a statement about moduli, which it is not unless $a \geq 0$. A true machine-checked lemma can sit underneath a false sentence. Every subsequent *spectral-ratio* consequence uses $a = \beta L$ with $\beta \geq 0$; the block *identity* immediately below is unconditional in β .

That sentence is worth reading twice. In an earlier version it said “everything below uses $\beta > 0$, so the hypothesis costs nothing here” — and the theorem on the next line was then stated for all $\beta \in \mathbb{R}$. The remark documenting a failure to propagate a hypothesis had itself failed to propagate it. The fix is not to phrase the promise better but to make the two statements structurally different, which is what Theorem 3.3 and Corollary 3.4 now are.

Theorem 3.3 (Fusion identity). *For every $\beta \in \mathbb{R}$ and every $L \geq 1$, the 2×2 principal submatrix of $K_\beta^{(L)}$ on the antipodal pair $\{\sigma^+, \sigma^-\}$ is exactly one Ising bond of coupling βL :*

$$\begin{pmatrix} K_\beta^{(L)}(\sigma^+, \sigma^+) & K_\beta^{(L)}(\sigma^+, \sigma^-) \\ K_\beta^{(L)}(\sigma^-, \sigma^+) & K_\beta^{(L)}(\sigma^-, \sigma^-) \end{pmatrix} = B(\beta L).$$

Corollary 3.4 (Fusion ratio). *If moreover $\beta \geq 0$, that block has subdominant ratio $\tanh(\beta L)$. For $\beta < 0$ its ratio is $|\tanh(\beta L)|$.*

Proof. Lemma 3.1 with $a = \beta L \geq 0$; Remark 3.2 for the other sign. $\square$

The split is deliberate, and it mirrors what is and is not machine-checked. The *matrix identity* of Theorem 3.3 is unconditional and is the Lean theorem `antipodal_block_eq_bond`; the *spectral reading* needs the sign hypothesis and is a paper-level consumer. An earlier version merged them and asserted the ratio for all β — false for $\beta < 0$, where $r(B(\beta L)) = |\tanh(\beta L)| > 0$ while $\tanh(\beta L) < 0$. The hypothesis had been added to Lemma 3.1 one paragraph earlier and simply not carried to its consumer.

Proof. σ^+ agrees with itself at all L sites, so the sum of signs is $+L$ and the entry is $e^{\beta L}$; likewise for σ^- . The two antipodal configurations disagree at all L sites, so the sum is $-L$ and the entry is $e^{-\beta L}$. Compare with $B(\beta L)$ and apply Lemma 3.1. $\square$

Theorem 3.3 is the mechanism this paper is about, and it is worth stating in words — carefully, because the natural phrasing overstates. In the singular concentration limit a weight supported on a set S *restricts* the kernel to S ; at every strictly positive ε the principal block on S is already preserved *exactly*, and it is the rest of the matrix that shrinks (Proposition 4.3). So the weight does not perturb the block and then correct it; the block is there from the start. When S is the antipodal pair that block is a single bond whose coupling is L times the original. **The weight fuses the L sites into one effective site.**

4 Why that is a wall

Theorem 4.1 (No uniform ratio bound). *For every $\beta > 0$ and every $\rho < 1$ there is an L with $\tanh(\beta L) > \rho$. Consequently the fused ratio is not bounded away from 1, and no estimate of the form $r \leq \rho < 1$ can hold uniformly in the number of sites.*

Proof. $\tanh a = 1 - 2/(e^{2a} + 1)$ is strictly increasing with limit 1; choose L with $e^{2\beta L} + 1 > 2/(1 - \rho)$, which exists because $\beta > 0$ and $e^x > x$. $\square$

This is the honest content of the paper’s title. It is *not* a proof that any given coupled transfer operator fails to have a uniform gap; the concentration that realises the fusion is a limit, and the weight of a physical model may or may not approach it. What Theorem 4.1 establishes is that *the obstruction lives in the congruence*, not in whatever spectral estimate one happens to be attempting. An argument that bounds r using only data preserved by congruence — rank, inertia, definiteness — cannot possibly produce an L -uniform bound, because those data are identical for the unfused kernel (ratio $\tanh \beta$) and for every fused one (ratio up to $\tanh(\beta L)$).

Theorem 4.1 is a statement about the scalar $\tanh(\beta L)$; the consequence for the congruence orbit needs one more step, which we make explicit rather than leave to the reader to assemble — and make without a limit. Normalise to unit diagonal, $M = e^{-\beta L} K_\beta^{(L)}$, which changes no

ratio; then $M_{\sigma\tau} = e^{-2\beta h(\sigma,\tau)}$ with h the Hamming distance, the least correlated pair is the antipodal one, $\mu = e^{-2\beta L}$, and $n = 2^L$. The lower bound of §6 applies verbatim:

$$r(D_\varepsilon K_\beta^{(L)} D_\varepsilon) \geq \frac{1 - \mu}{1 + \mu + 2^L \varepsilon} > \rho \quad \text{for every } \varepsilon \in (0, 1] \text{ with } \varepsilon < \frac{(1 - \mu) - \rho(1 + \mu)}{\rho 2^L}$$

(any such ε if $\rho \leq 0$), the numerator on the right being positive exactly because $\tanh(\beta L) > \rho$. D_ε is strictly positive throughout. So the failure of uniformity is realised inside the orbit, at an explicitly named weight, not only in its closure.

Corollary 4.2. *Any bound on r that is uniform over both the full positive-diagonal congruence orbit and the number of sites must use a quantity that a positive-diagonal congruence can change.*

The qualification matters. Nothing here forbids a uniform bound for a *particular* physical family of weights; what is forbidden is a uniform bound derived from congruence invariants alone, since those are blind to the difference between the unfused kernel and every fused one.

The orbit, not its closure

The fusion of Theorem 3.3 is realised by a weight that *concentrates*, and the congruences considered here are by *strictly positive invertible* diagonals. Setting the remaining diagonal entries to zero would leave the orbit and land in its closure, so the interface deserves to be stated rather than assumed. It is stated, and the statement is stronger than a limiting argument.

Proposition 4.3 (Concentration stays inside the orbit). *Fix two configurations $p \neq q$ and $0 < \varepsilon \leq 1$, and let D_ε be the diagonal that is 1 at p and q and ε elsewhere. Then D_ε is strictly positive, and for every such ε ,*

$$(D_\varepsilon M D_\varepsilon)_{pq} = M_{pq}, \quad (D_\varepsilon M D_\varepsilon)_{qp} = M_{qp},$$

while $|(D_\varepsilon M D_\varepsilon)_{ij}| \leq \varepsilon |M_{ij}|$ whenever $i \notin \{p, q\}$.

Proof. $(D_\varepsilon M D_\varepsilon)_{ij} = d_i d_j M_{ij}$ with $d_p = d_q = 1$, giving the two equalities; and $d_i = \varepsilon$, $d_j \leq 1$ off the pair. $\square$

So the fused bond is *already present at every strictly positive ε* : the limit does not create the block, it deletes the rest of the matrix. An earlier version converted that deletion into a statement about eigenvalues by citing the continuity of the spectrum in finite dimensions. It no longer does. Because the block is exact at every ε , a vector supported on the pair already sees it exactly, and §6 turns that observation into a bound valid at each ε separately — which is the statement one actually wants here, since the whole point of Proposition 4.3 is that we never leave the orbit.

5 Which congruence is extremal: pairs beat exchangeable clusters

Theorem 3.3 identifies the pair value $\tanh(\beta L)$, approached by concentrating on the antipodal pair — for $L > 1$ it is a statement about a principal *block*, and the orbit of strictly positive congruences reaches $B(\beta L) \oplus 0$ only as $\varepsilon \rightarrow 0$ (Proposition 4.3). Is that value the best available? For the family of *uniform* concentrations the question has a complete and elementary answer, and the answer explains the mechanism.

Theorem 5.1 (Exchangeable concentration). *Fix $0 < \mu < 1$ and $m \geq 2$, and let $E_m = \mu J_m + (1 - \mu)I_m$ be the exchangeable correlation matrix on m sites. Its spectrum consists of exactly two points: $1 + (m - 1)\mu$ on the constant vector, and $1 - \mu$ on every vector summing to zero. Hence the ratio it reaches is*

$$\frac{1 - \mu}{1 + (m - 1)\mu},$$

which is strictly decreasing in m .

Proof. $E_m \mathbf{1} = (1 + (m - 1)\mu)\mathbf{1}$ by summing a row. If $\sum_i x_i = 0$ then $(E_m x)_i = \mu \sum_j x_j + (1 - \mu)x_i = (1 - \mu)x_i$. These two eigenspaces have dimensions 1 and $m - 1$, so they exhaust the spectrum. For monotonicity, $k \mapsto (1 - \mu)/(1 + k\mu)$ is strictly decreasing on $k \geq 0$ because the numerator is positive and the denominator strictly increasing. $\square$

The reading is the point. **The more sites an exchangeable concentration keeps, the smaller the ratio it can reach.** So among these the best is the smallest nondegenerate one — a *pair* — and at $m = 2$ the value is $(1 - \mu)/(1 + \mu)$. On the hypercube kernel, $\mu = e^{-2\beta L}$ at the antipodal pair, and $(1 - \mu)/(1 + \mu) = \tanh(\beta L)$: Theorem 3.3 is the $m = 2$ case of Theorem 5.1.

Remark 5.2 (Exactly how far this goes). Theorem 5.1 concerns subsets on which *all* pairwise correlations equal a common μ . On a general subset the entries M_{ij} differ and the spectrum has no such closed form, so the theorem does *not* say that a pair beats every uniformly weighted subset of an arbitrary M . The section title should be read with the word *exchangeable* attached. What the theorem does establish is the mechanism — spreading a concentration over more equally correlated sites lowers the reachable ratio — and it is that mechanism, not the general proof; the general proof is §6.

6 The supremum, exactly: Hilbert’s projective diameter

The remaining question is what an *arbitrary* positive weight can reach. It has a complete answer, and the answer comes from geometry rather than from spectral estimates: a positive weight cannot change the projective geometry it acts on.

For an entrywise positive matrix T , Hilbert’s projective diameter is

$$\Delta(T) = \max_{i,j,k,l} \log \frac{T_{ik}T_{jl}}{T_{jk}T_{il}}.$$

Birkhoff’s theorem [9] says the map induced by T contracts Hilbert’s projective metric with coefficient $\tanh(\Delta(T)/4)$, and the spectral consequence is classical.

Fact 6.1 ([10], Theorem 1.1 with Theorem 2.3, eq. (6)). *For an entrywise positive matrix T with Perron root $\rho(T)$,*

$$r(T) \leq \tanh(\Delta(T)/4).$$

The attribution is exact and worth unpacking, since this is the one non-elementary import. Eveson and Nussbaum call $q(L) = \sup\{|z|/r(L) : z \in \sigma(L), z \neq r(L)\}$ the *spectral clearance* (their eq. (5)); it is verbatim the $r(T)$ of §1. Their Theorem 2.3, eq. (6), gives $q(L) \leq \inf_{j \geq 1} N(L^j)^{1/j}$, and their Theorem 1.1 — the Birkhoff–Hopf theorem — gives $N(L) = k(L) = \tanh \frac{1}{4} \Delta(L)$. Taking $j = 1$ yields Fact 6.1. Their hypotheses hold here with $m = 1$: the standard cone in $\mathbb{R}^n$ is closed, normal and reproducing, T is bounded with $T(C) \subseteq C$, and T entrywise positive gives $\Delta(T; C, C) < \infty$ and $T^2 \neq 0$. Their introduction lists matrices with strictly positive entries as the motivating class.

This is the only non-elementary input imported by this section, and it is imported as a statement, not reproved. The two facts that make Fact 6.1 bite are elementary, and both are proved here.

An earlier version of this paper imported a *second* classical fact, the finite-dimensional continuity of eigenvalues, to run the lower bound as a limit. It is gone. The lower bound below is proved instead by exhibiting a fluctuation vector, which yields the estimate at *every* $\varepsilon > 0$ rather than in a limit, and uses only Perron–Frobenius and the variational characterisation of the second eigenvalue of a symmetric matrix. That is not a stylistic preference: a bound available at each ε is a bound about the orbit, whereas a limit is a statement about its closure, and §4 has already had to insist on the difference.

Remark 6.2 (An independent check on the value). The contraction coefficient is also written without logarithms or hyperbolic functions [7]: with $\varphi(A) = \min_{i,j,k,l} A_{ik}A_{jl}/(A_{jk}A_{il})$,

$$\tanh(\Delta(A)/4) = \frac{1 - \sqrt{\varphi(A)}}{1 + \sqrt{\varphi(A)}}.$$

For our M the cross-ratios come in reciprocal pairs, so $\varphi(M) = 1/e^{\Delta(M)} = \mu^2$ and the right-hand side is $(1 - \mu)/(1 + \mu)$ directly. The two closed forms agree, which is a useful check: the *upper-bound constant* of Theorem 6.6 can be read off the minimum cross-ratio with no transcendental function anywhere. It is a second derivation of the Birkhoff *coefficient*, not a second proof of the theorem — the equality of the supremum still needs the lower bound, hence the concentration and the fluctuation witness.

That the two closed forms really are the same number is not left to this paragraph. It is the Lean theorem `phi_form_eq_tanh`: $\tanh(\frac{1}{4}\log(1/\varphi)) = (1 - \sqrt{\varphi})/(1 + \sqrt{\varphi})$ for every $\varphi > 0$. The point is not the difficulty — there is none — but the placement: the hypothesis assumed in the machine-checked theorem of §9 is written in the $\sqrt{}$ form, while the fact cited from the literature is written with $\tanh$, and an unchecked sentence between the two is exactly the gap this paper has already published once (Remark 3.2).

Lemma 6.3 (Δ is a congruence invariant). *For every positive diagonal D and all indices,*

$$\frac{(DMD)_{ik}(DMD)_{jl}}{(DMD)_{jk}(DMD)_{il}} = \frac{M_{ik}M_{jl}}{M_{jk}M_{il}}, \quad \text{hence} \quad \Delta(DMD) = \Delta(M).$$

Proof. $(DMD)_{ab} = d_a d_b M_{ab}$, so numerator and denominator both acquire the factor $d_i d_j d_k d_l$, which cancels. $\square$

Lemma 6.4 (The diameter of a unit-diagonal kernel). *Let M be a symmetric $n \times n$ matrix with $n \geq 2$, $M_{ii} = 1$, $M_{ij} \leq 1$, and*

$$\mu := \min_{i \neq j} M_{ij} > 0.$$

Then $\Delta(M) = 2\log(1/\mu)$ exactly.

Proof. Every cross-ratio is at most $1 \cdot 1/(\mu \cdot \mu) = 1/\mu^2$. Taking $i = k$ and $j = l$ on a pair attaining the minimum gives $M_{ii}M_{jj}/(M_{ji}M_{ij}) = 1/\mu^2$, so the maximum is attained. $\square$

Remark 6.5 (μ must be the attained minimum). An earlier version stated this lemma with μ merely a *lower bound*, $\mu \leq M_{ij} \leq 1$. That is false: $M = J_n$ with $\mu = \frac{1}{2}$ satisfies every such hypothesis, yet all cross-ratios equal 1, so $\Delta(J_n) = 0$ while $2\log(1/\mu) = 2\log 2 > 0$. The proof always used attainment; only the statement did not.

The Lean module had it right and the prose did not, which is worth recording rather than quietly repairing. There, `crossRatio_le_of_bounds` assumes only $\mu \leq M_{ab}$ and concludes an *inequality*, while `crossRatio_attained` takes $M_{ij} = \mu$ as an explicit hypothesis. The two halves were kept apart in the kernel and merged in the English.

Theorem 6.6 (Least-correlated pair). *Let M be a symmetric $n \times n$ matrix with $n \geq 2$, $M_{ii} = 1$ and $0 < \mu \leq M_{ij} \leq 1$ for $i \neq j$, where $\mu = \min_{i \neq j} M_{ij}$. Then*

$$\sup_{D > 0 \text{ diagonal}} r(DMD) = \frac{1 - \mu}{1 + \mu}.$$

Proof. Upper bound. DMD is entrywise positive, so by Lemmas 6.3 and 6.4 and Fact 6.1,

$$r(DMD) \leq \tanh(\Delta(DMD)/4) = \tanh(\tfrac{1}{2} \log(1/\mu)) = \frac{1/\mu - 1}{1/\mu + 1} = \frac{1 - \mu}{1 + \mu},$$

the middle identity because $t = e^{\frac{1}{2} \log(1/\mu)}$ satisfies $t^2 = 1/\mu$ and $\tanh \log t = (t^2 - 1)/(t^2 + 1)$.

Lower bound, with no limit. Fix a pair $\{p, q\}$ attaining the minimum, let D_ε be as in Proposition 4.3 with $0 < \varepsilon \leq 1$, and put $T_\varepsilon := D_\varepsilon M D_\varepsilon$. It is symmetric and entrywise positive, and by Proposition 4.3 its block on $\{p, q\}$ is exactly $\begin{pmatrix} 1 & \mu \\ \mu & 1 \end{pmatrix}$.

By Perron–Frobenius the Perron root ρ_ε is simple with a strictly positive eigenvector Ω , so the second eigenvalue is

$$\lambda_1 = \max \left\{ \frac{\langle x, T_\varepsilon x \rangle}{\|x\|^2} : x \perp \Omega, x \neq 0 \right\}, \quad \lambda_1 < \rho_\varepsilon.$$

Take the two-supported *fluctuation witness* $x = \Omega_q e_p - \Omega_p e_q$, which is nonzero because $\Omega > 0$. Two things happen at once. It is orthogonal to Ω identically — $\langle \Omega, x \rangle = \Omega_p \Omega_q - \Omega_q \Omega_p = 0$, with no hypothesis on Ω whatever — and, being supported on $\{p, q\}$, it sees only the block the congruence preserves:

$$\langle x, T_\varepsilon x \rangle = \Omega_q^2 + \Omega_p^2 - 2\mu \Omega_p \Omega_q = (1 - \mu)(\Omega_q^2 + \Omega_p^2) + \mu(\Omega_p - \Omega_q)^2 \geq (1 - \mu) \|x\|^2.$$

The middle step is an identity, not an estimate, and the slack is a square; so $\lambda_1 \geq 1 - \mu > 0$ and the subdominant *modulus* is at least $1 - \mu$.

For the denominator, every row of T_ε sums to at most $1 + \mu + n\varepsilon$: row p contributes $M_{pp} = 1$ and $M_{pq} = \mu$ exactly and at most ε from each remaining entry, row q likewise, and a row outside the pair carries a factor ε throughout. A nonnegative matrix has Perron root at most its largest row sum, so $\rho_\varepsilon \leq 1 + \mu + n\varepsilon$. Therefore

$$r(T_\varepsilon) \geq \frac{1 - \mu}{1 + \mu + n\varepsilon} \quad \text{for every } \varepsilon \in (0, 1],$$

with D_ε strictly positive throughout. Given $t < (1 - \mu)/(1 + \mu)$, the right-hand side exceeds t for every $\varepsilon \in (0, 1]$ when $t \leq 0$, and as soon as $\varepsilon < [(1 - \mu) - t(1 + \mu)]/(tn)$ when $t > 0$ — the numerator being positive precisely by the choice of t . That is an explicit witness rather than a limiting argument, and it never leaves the orbit. Hence the supremum is at least $(1 - \mu)/(1 + \mu)$. $\square$

Remark 6.7 (Attainment, and a retracted claim). An earlier version of this paper asserted that the supremum is *not* attained. That is false. For $n = 2$ the matrix $\begin{pmatrix} 1 & \mu \\ \mu & 1 \end{pmatrix}$ is its own least correlated pair and already has eigenvalues $1 \pm \mu$ at $D = I$, so the value is reached with no concentration and no limit; for $\mu = 1$ the matrix is J , the ratio is 0, and the value is reached for every D . The stated reason was worse than unsupported: it claimed Δ is

“a strict maximum only in the limit”, which Lemma 6.3 directly contradicts — that lemma proves Δ is *constant* on the orbit. **We do not classify attainment here.** Doing so for $n > 2$ would need the equality cases of the Birkhoff bound, which the invariance of Δ alone cannot supply.

Remark 6.8 (Fact 6.1 is attained, at every diameter). The same 2×2 computation says something about the imported bound, and it is worth extracting. Given $\delta > 0$ put $\mu = e^{-\delta/2} \in (0, 1)$ and $M = \begin{pmatrix} 1 & \mu \\ \mu & 1 \end{pmatrix}$. By Lemma 6.4, $\Delta(M) = 2 \log(1/\mu) = \delta$; the eigenvalues are $1 \pm \mu$; so

$$r(M) = \frac{1 - \mu}{1 + \mu} = \tanh(\delta/4)$$

with equality. Hence at every admissible value of the diameter Fact 6.1 is met exactly, and **no bound on r that depends on Δ alone can be smaller.** The equality of the two numbers is the Lean theorem `birkhoff_attained_two`, and the eigenvalues $1 \pm \mu$ are `exch_two_eigen`. The argument is immediate and we claim nothing for it beyond having looked; we state it because it explains the shape of the classical literature.

The quantifier here is the whole content, and an earlier version of this remark got it wrong. It said that “the same is true along the whole congruence orbit” of that matrix, on the ground that Δ is invariant. That does *not* follow, and it is false: Lemma 6.3 says Δ is constant on the orbit, not that r is, and r certainly is not. For $M = \begin{pmatrix} 1 & \mu \\ \mu & 1 \end{pmatrix}$ with $\mu = 0.1$ and $D = \text{diag}(1, d)$, the diameter agrees to 10^{-11} at every d tested while r falls from 0.818 at $d = 1$ to 9.9×10^{-5} at $d = 0.01$; across three values of μ , twelve of fifteen orbit points sit strictly *below* the bound. The correct statement is the uniform one, and it is the one worth having:

Δ is constant on the orbit and at least one point of the orbit attains $\tanh(\Delta/4)$; therefore no bound valid over that whole orbit and depending only on Δ can be smaller.

The diameter pins the best *universal* bound, not the ratio of each individual member — which is, after all, exactly the paper’s thesis: the congruence moves the spectrum and cannot move the geometry. Bauer, Deutsch and Stoer [5] bound exactly this r — their $|\lambda(A)|/\pi(A)$ — and *improve* Hopf’s inequality by bringing in the eigenvector belonging to $\pi(A)$. They had to bring in something: a positive diagonal congruence moves the eigenvector and fixes the diameter, so any sharpening must be of that kind. Theorem 6.6 is the statement that makes this a theorem rather than an observation about proofs.

Three things about this proof are worth saying plainly.

It uses no definiteness whatever — only positivity of the entries. That is why the twelve indefinite witnesses in the certificate below all obeyed the bound: the pre-registered gate that predicted otherwise was pointing at this argument, and we did not see it at the time.

It explains the rigidity. The quantity that survives the weight is not a spectral one at all; it is the projective diameter. Congruence moves the spectrum but cannot move the geometry, and the reachable spectral ratio is pinned by the geometry.

Novelty, stated carefully. Both imported facts are classical, and Lemma 6.3 is immediate once written down. What we did not find in the searches we ran is this exact *evaluation* of the supremum of the spectral ratio over a diagonal congruence orbit. That is an absence of found prior art, not a claim of priority.

What the first attempt could not do

We record the route that failed, because it locates the difficulty. Since T is positive definite exactly when M is (Theorem 2.2),

$$r(T) \leq \frac{1-\mu}{1+\mu} \iff \lambda_0 - \lambda_1 \geq \mu(\lambda_0 + \lambda_1),$$

and writing $M = \mu J + (1-\mu)N$ — where N again has unit diagonal and entries in $[0, 1]$ — gives $T = \mu dd^\top + (1-\mu)DND$, a rank-one positive semidefinite piece plus an entrywise nonnegative one. Weyl’s inequality and the Perron root of the second piece bound $|\lambda_1|$ correctly but lose too much on λ_0 : testing against the Perron vector of DND reduces the claim to $\sum_{i \neq j} (1+\mu-2M_{ij})u_i u_j \geq (1-\mu) \sum_i u_i^2$, whose left-hand side is negative once some $M_{ij} > (1+\mu)/2$. The top eigenvector of T is not the Perron vector of DND , and the slack is exactly that gap.

On the hypercube, $K_\beta^{(L)}$ normalised to unit diagonal has $M_{\sigma\tau} = e^{-2\beta h(\sigma,\tau)}$ with h the Hamming distance; the least correlated pair is the antipodal one, $\mu = e^{-2\beta L}$, and $(1-\mu)/(1+\mu) = \tanh(\beta L)$. Theorem 6.6 therefore contains the *spectral-ratio consequence* of Theorem 3.3 (Corollary 3.4) as its hypercube instance; the identity of Theorem 3.3 itself supplies the explicit antipodal block that the lower bound concentrates on.

Numerical certificate. Random search plus Nelder–Mead ascent over positive diagonals, on families chosen to break it and run before the proof was found, never exceeded $(1-\mu)/(1+\mu)$ and in almost every case *matched* it to machine precision. We say “matched” rather than “attained”: a floating-point equality at the end of an ascent is not evidence that a maximiser exists, and this paper does not classify maximisers.

family	worst excess over target	cases
hypercube, $L \in \{2, 3, 4\}$, $\beta \in \{0.15, 0.4, 0.8\}$	1.0×10^{-15}	9
random positive Gram, $n \in \{4, 5, 6\}$	5.0×10^{-16}	12
full rank / near singular ($\kappa \sim 6 \times 10^6$)	3.3×10^{-16}	4
near-exchangeable, isolated argmin, entries near 1	4.4×10^{-16}	3
exactly exchangeable $\mu J + (1-\mu)I$	1.0×10^{-15}	2
<i>indefinite</i> , $\lambda_{\min} \in [-0.42, -0.06]$	1.6×10^{-15}	12
larger sizes, $n \in \{8, 10, 12\}$, κ up to 5×10^4	4.2×10^{-16}	6

The last row was added after an audit of our own procedure: the adversarial sweep as first written never actually reached $n = 8$. Its positive-definiteness *precondition* fired on that family, the generator was repaired, and the family was not re-run — so the largest size with evidence had been $n = 6$, not $n = 12$ as the sweep’s summary implied. A precondition that is never met carries no evidence either way, and the script now returns a distinct INCONCLUSIVE code rather than reporting such a case as a result.

The last row deserves comment, because it corrected us. A pre-registered gate predicted that positive definiteness would be load-bearing — that at least one of twelve indefinite witnesses would break the bound. None did; all twelve landed *on* it. The hypothesis was therefore dropped, and Theorem 6.6 as stated above assumes no definiteness. The gate, its prediction, and its failure are recorded in the repository rather than repaired in place.

Scope of the certificate. These are numerics, not proof. They cover $n \leq 12$ and $\mu > 0$, and they include indefinite symmetric matrices — but entrywise positivity still guarantees Perron dominance, so the Perron root is the spectral radius in every case and r always means

what §1 says it means. What is *not* tested is $\mu \leq 0$, or matrices with nonpositive entries, where both the Birkhoff argument and the definition of the class break down. (An earlier version claimed the sweep did not reach “a negative eigenvalue large enough to make $|\lambda_0|$ the negative extreme”. That situation cannot arise here at all: entrywise positivity forbids it. The sentence was a leftover from before r was defined through the Perron root.)

7 Relation to prior work

Sylvester (1852). The law of inertia is classical and Theorem 2.2 is its definite case. We claim no novelty on the rigid half; the contribution is that this form of it is checkable without diagonalisation.

Ostrowski. Ostrowski’s quantitative form of the inertia law gives, for Hermitian A and invertible S , $\lambda_k(SAS^*) = \theta_k \lambda_k(A)$ with $\theta_k \in [\lambda_{\min}(SS^*), \lambda_{\max}(SS^*)]$. For $S = D$ diagonal this yields

$$r(DAD) \leq \kappa(D^2) r(A), \quad \kappa(D^2) = \frac{\max_i d_i^2}{\min_i d_i^2},$$

which **diverges** exactly in the regime of interest, namely when D degenerates. Theorem 6.6 supplies a *finite* substitute in that regime — the exact supremum $(1 - \mu)/(1 + \mu)$, and $\tanh(\beta L)$ in the hypercube case, however badly conditioned D becomes — while Theorem 3.3 identifies the extremal antipodal block that realises it there. That contrast, divergent classical bound against finite exact value, is stated side by side so that it can be checked.

Onsager and Kaufman. The kernel $DK_\beta D$ with a ferromagnetic nearest-neighbour weight is the transfer matrix of the two-dimensional anisotropic Ising model on a strip, whose spectrum is known in closed form. Nothing in Sections 3–4 is new mathematics about that model, and the reader who knows the exact solution will find $\tanh(\beta L)$ unsurprising: it is the zero-temperature limit in the transverse direction, where the strip collapses to a single spin. The contribution is not the value. It is that the value is *exactly* the supremum over the whole diagonal congruence orbit (Theorem 6.6), that the orbit’s rigid invariants are simultaneously pinned, and that the pair of facts identifies which quantities an argument may and may not use. An earlier version of this paper could only conjecture the first of those; the phrasing here is now licensed by §6.

The DAD branch, audited. The previous version of this paper named this branch as an open obligation. It has now been read, and the outcome is more useful than a disclaimer: it *shrinks* the novelty claim in one place and sharpens it in another.

The scaling literature asks a different question. Idel’s survey [6] states its subject in its opening line — given a nonnegative A , find diagonal D_1, D_2 making $D_1 A D_2$ doubly stochastic — and its seventy years of material is about the existence, uniqueness and convergence of such a *normalisation*, together with the generalisation to positive maps between matrix algebras. Van der Sluis’s optimal scaling [11] is of the same type. Bushell [4] sits in the same circle from the other side: he computes the projective contraction ratio itself, which is an *ingredient* of Fact 6.1 rather than a statement about an orbit. None of this asks what a spectral quantity can *reach* as D varies.

The ratio itself is classical, and we say so plainly. Bauer, Deutsch and Stoer [5] bound $|\lambda(A)|/\pi(A)$ for positive A — our $r(A)$ verbatim — by cone-contraction arguments, and show Hopf’s inequality can be improved using the eigenvector of $\pi(A)$; Fact 6.1 is the descendant we cite. **The upper half of Theorem 6.6 is therefore not new in any part.** It is

a classical bound applied to DMD , plus the observation — immediate once written down — that Δ does not move. Remark 6.8 adds that the classical bound is exactly attained, so nothing was lost by importing the weakest form of it.

What the search did not find is the *evaluation*: the supremum of r over a positive diagonal congruence orbit, together with a matching lower bound realised at strictly positive weights inside the orbit. Its two non-classical ingredients are Lemma 6.3 and the fluctuation witness of §6. We record this as an absence of found prior art after a directed reading of the branch named above — Eveson and Nussbaum point to Borwein, Lewis and Nussbaum as the DAD entry point — and not as a claim of priority.

Formalisation. We are not aware of a mechanised treatment of congruence orbits of transfer kernels. This is an absence of found prior art in the searches we ran, not a claim of priority.

8 What is open

1. **Formalising Fact 6.1** — the one *mathematical* import that remains. The interface it must meet is already fixed and checked: the hypothesis is stated in the logarithm-free form of Remark 6.2, so a formalisation need carry neither `log` nor `tanh`, and `phi_form_eq_tanh` certifies that the form assumed is the form cited. Discharging it means proving the Birkhoff–Hopf contraction bound in Lean, which is a genuine project and not an afternoon.
2. **Bridging SpectralInterface to Mathlib** — the one *interface* import that remains, and a different kind of debt. Its content, *on symmetric entrywise positive matrices*, is Perron–Frobenius together with the variational characterisation of the second eigenvalue: standard, elementary, and not in question. Both qualifiers are load-bearing — dropping symmetry made an earlier version of this hypothesis false, and hence the theorem depending on it vacuous. What is missing is that the pinned Mathlib has neither fact in a form that applies to a `Matrix`-level definition of r , which this module deliberately does not introduce.

This one now looks reachable, and we say so with the reasons. The repository already contains, for entrywise positive kernels, the existence of a strictly positive eigenvector (`exists_pos_eigenvector`), a spectral orthonormal basis for real symmetric kernels, the subdominant modulus `specGap` with `specGap < lam`, the operator bound `norm_act_le_specGap` on the Perron complement, and `specRatio = specGap / lam` — which is r verbatim. The route is then short: apply the interface premise at the Perron vector, turn the Rayleigh bound into $b \leq \text{specGap}$ by Cauchy–Schwarz against the operator bound, get $\lambda \leq a$ from the row sums at a maximal coordinate of Ω , and divide. **The obstacle is not the analysis but the packaging:** r must be a *total* function `Matrix` $\rightarrow \mathbb{R}$, so the Perron data has to be chosen canonically and the definition split by cases on the symmetric-positive domain. We flag it as the next piece of work and do not claim it done. The lower bound that used to be assumed wholesale is now *derived* from this interface (`concentrated_lower_of_spec`), so the debt has moved from “assume my own conclusion” to “assume two textbook facts” — but it has not vanished, and we do not say it has.

These two are the whole gap, and they are of different kinds: the first is mathematics nobody here has done, the second is plumbing nobody here has laid. Only the first would be a discovery.

3. Whether the congruence-invariance of Δ extends usefully beyond diagonal congruences — for instance to congruences by positive matrices that are not diagonal, where the cross-ratio no longer cancels.
4. Whether the physical weight w_γ attains the antipodal restriction in the limit $\gamma \rightarrow \infty$ with the rate we measure. The deficit $\tanh(\beta L) - r(\gamma)$ was found to decay like $e^{-2\gamma}$ (fitted exponents -1.96 to -2.08 against a pre-registered prediction of -2), consistent with one-domain-wall configurations dominating the correction. The limit itself is not proved here.
5. The general signature-counting law as a mechanised statement.

Machine verification — status

The following are written in Lean 4: the rigid half (Theorem 2.2, Lemma 2.1); the *matrix identity* underlying Theorem 3.3; the *scalar tanh-approach lemma* underlying Theorem 4.1; Proposition 4.3; and the algebraic and eigenvector ingredients of Lemma 3.1, Theorem 5.1 and Theorem 6.6; the fluctuation witness and the scalar ε -step that replaced the limit in the lower bound; and the *conditional* form of Theorem 6.6 itself — against a pinned Mathlib in `YangMills/OS/CongruenceSpectrum.lean`, and **have been verified**, including the two lemmas that carry Theorem 6.6 and the $n = 2$ attainment witness of Remark 6.7.

What “ingredients” means, precisely. The module contains no Lean definition of r , no enumeration of a spectrum, and no theorem asserting that a second largest *modulus* equals any stated quantity. What it contains for Lemma 3.1 and Theorem 5.1 is the eigenvector equations, the algebraic ratio identity, the positivity and sign lemmas, and the scalar monotonicity. The same caution applies to Theorems 3.3 and 4.1: Lean proves the antipodal block equals $B(\beta L)$ as matrices, and proves $\forall \rho < 1 \exists L, \rho < \tanh(\beta L)$ as a scalar statement, but the reading of either as a claim about the spectrally defined r — and the passage from the block to a strictly positive member of the orbit — is completed at paper level. *Their interpretation in terms of the spectrally defined r of §1 is completed at paper level*, by elementary finite-dimensional linear algebra. We are explicit about this because it is exactly the gap Remark 3.2 is about: a true Lean lemma can sit underneath a sentence that says more than it does, and this paper has already published one such sentence.

Theorem 6.6 is *not* an unconditional Lean theorem, and we do not claim it is. What exists in Lean is the *conditional* theorem, which is a different and weaker object, stated exactly:

```
congruenceRatio_isLUB_of_birkhoff_of_spectralInterface :
  IsLUB (congrOrbit r M) ((1-μ)/(1+μ))
```

carrying **two** explicit hypotheses, and named for both. They are hypotheses, not **axiom** declarations: nothing is added to the ambient theory, the axiom oracle below is unaffected, and a reader who declines either one receives a true implication and no claim.

The name is part of the statement, and the previous version’s was not. It read `..._of_birkhoff` while the type carried a second hypothesis, and the surrounding prose then said that formalising Fact 6.1 was all that stood between this theorem and an unconditional one. That was true of the mathematics and false of the published type. It is the defect of Remark 3.2 in its mildest form — a label that says less than what sits under it — and the fix is not a better sentence but a name that cannot drift from the type.

The two imports are not of the same kind, and the difference is the whole remaining agenda:

- **hBirkhoff** is a **mathematical** import: Fact 6.1, a real theorem nobody here has proved.
- **hSpec** is an **interface** import. It says: for a symmetric, entrywise positive T , if *for every* strictly positive candidate Perron vector Ω *there exists* a nonzero $x \perp \Omega$ with $\langle x, Tx \rangle \geq b\|x\|^2$, and every row sum is at most a , then $r(T) \geq b/a$. That is Perron–Frobenius plus the variational characterisation of the second eigenvalue, packaged once, and stated so that *nothing from this paper appears in it* — no μ , no pair, no concentration, no congruence. It is carried because the pinned Mathlib supplies neither fact against a **Matrix**-level definition of r , not because either is in doubt.

Both hypotheses had the wrong domain, and one of them was false. This is the correction that cost a version, and it was found by reading the Lean types rather than the paragraph above them.

Symmetry. **hSpec** did not require $T^\top = T$. Without it the statement is not a packaging of textbook facts — it is *false*. Take

$$T = \begin{pmatrix} 3 & 5 & 2 \\ 5 & 4 & 1 \\ 2 & 5 & 3 \end{pmatrix},$$

entrywise positive, every row summing to 10, spectrum $\{10, 1, -1\}$, so $r(T) = 1/10$. But $\langle x, Tx \rangle$ sees only the symmetric part $(T + T^\top)/2$, whose second eigenvalue is 1.330519; by Courant–Fischer every positive Ω admits an $x \perp \Omega$ with $\langle x, Tx \rangle \geq 1.3\|x\|^2$ — measured over 4000 random positive Ω , worst attainable value 1.330518744. The premises hold with $b = 1.3$, $a = 10$; the conclusion demands $0.13 \leq 0.1$. **A conditional theorem whose hypothesis no ratio function can satisfy is vacuous**, which is a worse failure than one that merely assumes too much — and our own non-vacuity checks had audited the *conclusion* rather than the satisfiability of the premises, one level too low.

Positivity. **hBirkhoff** did not require entrywise positivity, so it quantified over every matrix obeying the cross-ratio inequality. Fact 6.1 is a theorem about *positive* operators, and §6 spends a paragraph checking exactly those hypotheses; so the Lean statement asked for strictly more than the source supplies and could not have been built from [10] at all. Assuming more than the source provides is the mirror image of claiming more than the lemma proves, and `phi_form_eq_tanh` does not help: it identifies the two *constants*, not the two *domains*.

Both are repaired by copying into the types the hypotheses the paper proof was already using: **BirkhoffInterface** and **SpectralInterface** now carry $T^\top = T$ and $0 < T_{ij}$, $M^\top = M$ is a hypothesis of the endpoint, and **congr_transpose** carries symmetry to every point of the orbit — which is the only reason either import applies at *DMD* rather than only at *M*. Positivity of *DMD* needs no new hypothesis: it follows from $0 < \mu \leq M_{ab}$.

What changed in this version is which of the two the lower bound is. It used to be its own hypothesis — the main theorem assumed the conclusion of §6’s lower bound outright. It is now *derived* (`concentrated_lower_of_spec`) from **hSpec** together with the block-preservation of Proposition 4.3, the fluctuation witness, and a new row-sum bound `concentrate_rowSum_le` giving $\rho \leq 1 + \mu + n\varepsilon$. The debt moved from “assume my own conclusion” to “assume two textbook facts”. It did not disappear, and we do not say it did.

Five endpoints carry it:

- `congrOrbit_upper_of_birkhoff` — Birkhoff at *DMD* reduces to Birkhoff at *M*. This is Lemma 6.3 doing its work: the cross-ratio hypothesis available at *M* transports to every point of the orbit.

- **leastPair_fluctuation_witness** — both halves of the lower bound’s input, with no eigenvalue and no limit in the statement: the witness is orthogonal to *any* Ω , and its quadratic form is at least $(1 - \mu)\|x\|^2$.
- **leastPair_specGap_lower** — the division step, and the only place a Perron root is mentioned at all.
- **concentrate_rowSum_le** — the denominator: every row of $D_\varepsilon M D_\varepsilon$ sums to at most $1 + \mu + n\varepsilon$, the two pair entries surviving untouched and everything else carrying a factor ε . This is what makes the limit $(1 - \mu)/(1 + \mu)$ and not something else.
- **concentrated_specRatio_approaches** — every $t < (1 - \mu)/(1 + \mu)$ is exceeded *inside* the orbit, at an ε the proof produces.
- **congr_transpose** — a positive diagonal congruence preserves symmetry, so the orbit stays inside the class both imports are about. Without it neither hypothesis is applicable at DMD , only at M .

The hypothesis **hBirkhoff** is stated in the logarithm-free form of Remark 6.2 — no **log** and no **tanh** occur in it, only **Real.sqrt** — and with φ merely a lower bound for the cross-ratios rather than their minimum. Both choices assume strictly less than the sharp statement, because $x \mapsto (1 - x)/(1 + x)$ decreases. That the form assumed is the form cited is itself checked, by **phi_form_eq_tanh**; a reader should not have to take a change of closed form on the author’s word, least of all in this paper.

Spectral continuity is no longer imported anywhere. It was the second classical input of the previous version, and it has been *eliminated* rather than formalised: **lower_bound_approaches** returns an explicit ε , and the Rayleigh estimate is the ring identity $a^2 + b^2 - 2\mu ab = (1 - \mu)(a^2 + b^2) + \mu(a - b)^2$, whose slack is a square. What remains unformalised is Fact 6.1 and the two textbook facts packaged in **hSpec** — one deep, two shallow, and none of them hidden.

Non-vacuity was audited rather than asserted: **congrOrbit_nonempty** (the supremum is not taken over $\emptyset$) and **isLUB_value_pos** (the pinned value is not the trivial 0), together with the $n = 2$ attainment of Remark 6.7, which is **exch_two_eigen**.

On provenance. Fact 6.1 is now the single *non-elementary* load-bearing import of Theorem 6.6 — the lower bound imports nothing beyond Perron–Frobenius and the variational characterisation of the second eigenvalue — so a reader must be able to check it at named statements rather than at a paper. It now is: Theorem 1.1 and Theorem 2.3, eq. (6) of [10], both read in the source rather than inferred from secondary literature. Three earlier attempts to locate it failed; the record is kept because a citation that took four tries is one a referee should re-check.

We cite only statements we have read. In particular we do *not* attribute Fact 6.1 to a corollary in §2 of [10]: the numbered results we located there are Propositions 2.1–2.2, Lemmas 2.1–2.3, Theorems 2.2–2.3 and Remarks 2.1–2.3. Bibliographic data for [9, 7] were checked against the publishers’ records; the internal numbering of those two was not, and neither is load-bearing.

No statement in this paper is conjectural.

commit (module, this version)	7d0be024
commit (aggregate import)	fdb0317b
module sha256	dd915eddd5ac2d81f91c1156
	c88bb83fdaa02cdca4aefa0804dde18f4cb519fc
toolchain	leanprover/lean4:v4.29.0-rc6
Mathlib pin	07642720480157414db592fa85b626dafb71355b
elaboration	exit 0; log 0 bytes — zero errors, zero warnings, zero sorry
compilation	lake build YangMills.OS.CongruenceSpectrum, exit 0
axiom oracle	exit 0; 82/82 declarations resolved
	80 on [propext, Classical.choice, Quot.sound]
	2 (cfgPlus, cfgMinus) on [propext] alone
sorryAx	none

The oracle was generated from the source, not hand-written: the declaration list is extracted by pattern and every extracted name is queried, so a declaration cannot be omitted by being forgotten. One of the 82 names carries a prime (`concentrate_block'`); the extractor was checked to capture it, because a name-blind parser has already cost this project a false all-clear elsewhere. The two declarations resting on `[propext]` alone are configuration constants, and `[propext]` is a subset of the standard triple, so the count of *nonstandard* axioms is zero. Every extracted name was also checked to *appear* in the oracle’s output, so a silently mistyped query cannot pass as a silent success.

And the counter was wrong once, in a way worth recording. On the run that produced these figures the first tally reported 78 declarations against 80 output lines, with $75 + 2$ accounted for: three lines unexplained. The cause was not a bad axiom. Renaming the main theorem to `congruenceRatio_isLUB_of_birkhoff_of_spectralInterface` made the name long enough that Lean wrapped its axiom list across three lines, and a check that matched `[propext, Classical.choice, Quot.sound]` as a substring of a *line* could no longer see it. The reconciliation is now done on whitespace-normalised text, per declaration rather than per line, and it balances exactly: 78 entries for 78 names, $76 + 2 + 0$, nothing unexplained. We report this because the failure mode is the interesting one — a verification script that becomes blind to a purely typographic change, in this case one caused by fixing a name — and because a tally that does not add up is a measurement not yet understood, not a rounding error.

Anchor versus branch head. Every figure in this section is measured at 7d0be024. The branch head when this version was written is 97c870cc, and the difference between them is confined to the module’s `/-! ... -/` section header: two paraphrases that had gone on describing the previous version’s interfaces — “a Perron vector for the lower half”, which understates `SpectralInterface`, and $\varepsilon \leq 1$ where the hypothesis is $\varepsilon \leq \frac{1}{2}$. The whole diff is comment text and is exhibited in that commit. We anchor at 7d0be024 rather than reprint measured numbers beside a hash they were not measured at, and we name the head rather than let a reader discover a hash mismatch on checkout.

The module hash was computed independently on the author’s machine and on the verification host and agrees, so the checked object is the source printed here. The figures above come from a *second* run, on a runtime created after the previous one was destroyed: a fresh clone at the same commit, a fresh toolchain install, and a fresh Mathlib cache, reproducing the same hash and the same verdicts. That is the independent-clone condition the project’s governance asks for, and it is met here rather than asserted. Compilation ran on a CPU (never GPU) high-memory cloud runtime, per the project’s environment rule of 1 August 2026; the desktop compiles nothing.

The registered aggregate check has now run, and passed. The pre-registration

also asked that the module be imported into the project’s aggregate library and that the build job count increment by exactly one. Both builds were run in the *same* tree, on the same host, toolchain and Mathlib pin, differing only in whether the import line is present:

build	exit	jobs
aggregate <i>without</i> the import (baseline)	0	8430
aggregate <i>with</i> the import (treatment)	0	8431
increment	+1, as predicted	

Both counters are printed rather than only the difference, and neither is a subtraction against a figure measured in some other tree.

These are the counts at the final commit, not inherited ones. An earlier version of this section carried figures measured at `fdb0317b`, when the module was smaller, and said so rather than reprinting them beside a newer hash. They have now been re-measured at `7d0be024`, on the fresh runtime described above: baseline **8430**, treatment **8431**, both exiting 0, in the same tree, the baseline produced by checking the parent’s `YangMillsCore.lean` into the treated tree — verified to contain zero occurrences of `CongruenceSpectrum` — rather than by a separate clone. The exit codes were read from sentinel files written atomically after termination, per the project’s sentinel rule; a sentinel’s existence is not success, and only the integer inside it is.

One difference from the earlier run, stated because it removes an argument we previously used. That run could point at `Built YangMills.OS.CongruenceSpectrum` appearing immediately before `Built YangMillsCore` in the treatment log. This run cannot: the module had already been compiled earlier in the same session by `lake build YangMills.OS.CongruenceSpectrum`, so Lake had nothing to rebuild and printed no such line. The reachability evidence here is the job count itself, which is better: the aggregate needs 8431 jobs with the import and 8430 without, and the one job of difference *is* the module. A dependency that were merely present on disk and not reached would contribute zero.

References

- [1] J. J. Sylvester, *A demonstration of the theorem that every homogeneous quadratic polynomial is reducible by real orthogonal substitutions to the form of a sum of positive and negative squares*, Philosophical Magazine **4** (1852), 138–142.
- [2] A. M. Ostrowski, *A quantitative formulation of Sylvester’s law of inertia*, Proc. Nat. Acad. Sci. U.S.A. **45** (1959), no. 5, 740–744.
- [3] A. M. Ostrowski, *A quantitative formulation of Sylvester’s law of inertia, II*, Proc. Nat. Acad. Sci. U.S.A. **46** (1960), no. 6, 859–862.
- [4] P. J. Bushell, *On the projective contraction ratio for positive linear mappings*, J. London Math. Soc. (2) **6** (1973), 256–258.
- [5] F. L. Bauer, E. Deutsch and J. Stoer, *Abschätzungen für die Eigenwerte positiver linearer Operatoren*, Linear Algebra Appl. **2** (1969), 275–301.
- [6] M. Idel, *A review of matrix scaling and Sinkhorn’s normal form for matrices and positive maps*, arXiv:1609.06349 (2016).

- [7] E. Seneta, *Non-negative Matrices and Markov Chains*, 2nd ed., 1981; revised paperback printing, Springer Series in Statistics, Springer, 2006. (Birkhoff’s contraction coefficient, §3.1–3.2.)
- [8] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, 2013. (Congruence and inertia, §4.5; Perron–Frobenius, §8.2; the variational characterisation of the eigenvalues of a Hermitian matrix, §4.2.)
- [9] G. Birkhoff, *Extensions of Jentzsch’s theorem*, Trans. Amer. Math. Soc. **85** (1957), 219–227.
- [10] S. P. Eveson and R. D. Nussbaum, *Applications of the Birkhoff–Hopf theorem to the spectral theory of positive linear operators*, Math. Proc. Cambridge Philos. Soc. **117** (1995), no. 3, 491–512.
- [11] A. van der Sluis, *Condition numbers and equilibration of matrices*, Numer. Math. **14** (1969), 14–23.
- [12] L. Onsager, *Crystal statistics. I. A two-dimensional model with an order-disorder transition*, Phys. Rev. **65** (1944), 117–149.
- [13] B. Kaufman, *Crystal statistics. II. Partition function evaluated by spinor analysis*, Phys. Rev. **76** (1949), 1232–1243.
- [14] The mathlib Community, *The Lean mathematical library*, Proc. CPP 2020, ACM, 367–381.