Region bounds, error margins and certified positivity in the numerical conformal bootstrap

or: the first mathematically rigorous bounds on the 3d Ising critical exponents

to appear

Marten Reehorst

Marten Reehorst

(King's College London)

The question of this talk

- We often call bootstrap results *rigorous bounds*.
- The existence of a positive functional proves an inconsistency in the assumptions.
- But as published, are bootstrap bounds a mathematically rigorous bound on CFT data?
- Not quite. There are four specific gaps.
- I'll discuss: what they are, how to close all four — and why closing them is useful even if you do not care about rigour.

The four loopholes

1. **Points vs. regions.** One SDP excludes *one point* $(\Delta_\sigma, \Delta_\epsilon)$. Bounds are read off by interpreting the results at sample points.
2. **Finitely many vs. all spins.** Positivity imposed for $\ell \leq \ell_{\text{max}}$; while crossing demands unbounded spin.
3. **Floating point vs. exact arithmetic.** Functionals are certified positive only up to some finite floating point precision.
4. **Approximate vs. true blocks.** The solver never sees a true conformal block; it sees a recursion truncated at order κ.

The four loopholes

- Despite these loopholes the claim that bounds are ‘rigorous’ are ubiquitous.
- I will just call out the worst offender which puts the false claim right in the title:
- ‘Rigorous bounds on irrelevant operators in the 3d Ising model CFT’

The four loopholes

- Despite these loopholes the claim that bounds are ‘rigorous’ are ubiquitous.
- I will just call out the worst offender which puts the false claim right in the title:
- ‘Rigorous bounds on irrelevant operators in the 3d Ising model CFT’ **by Marten Reehorst**.

Loophole 1 is not academic

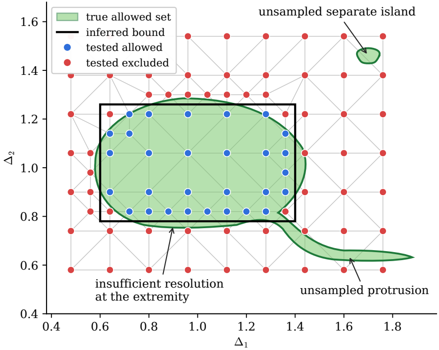

- Sampled points give a *picture*, not a bound.
- Three failure modes:
  - unsampled second island,
  - thin protrusion,
  - unclear resolution at the extremities.
- The last is a silent, (usually) un-estimated error on every quoted digit.

This has bitten the field

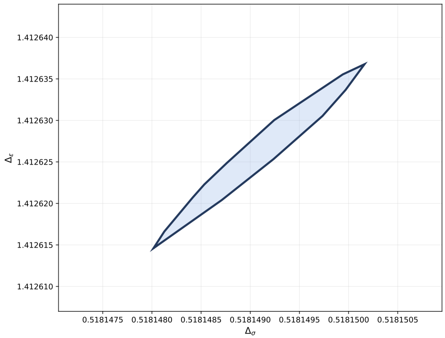

- [Kos, Poland, Simmons-Duffin, Vichi '16]: island charted by point sampling.

This has bitten the field

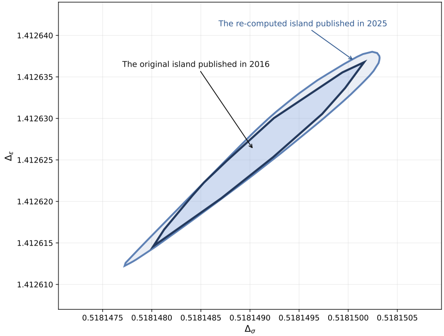

- [Kos, Poland, Simmons-Duffin, Vichi '16]: island charted by point sampling.
- Re-computed with better algorithms and more compute [Chang et al. '24]: *slightly larger* — the quoted errors were underestimated.
- Most of us have private anecdotes of this kind that did not make it to publication.

The fix: exclude *regions*, not points

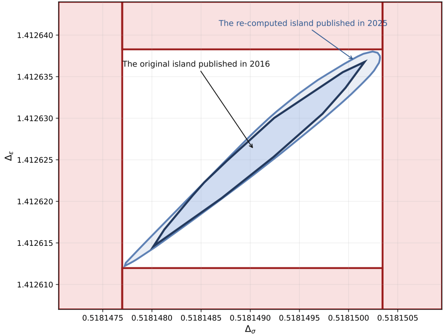

- A functional positive on a whole *box* excludes the whole box.
- A closing ring of boxes bounds the island **whatever its shape** — protrusions and satellites included.
- No interpolation, no interpretation.

KING'S College LONDON

Two reasons to care

- The 3d Ising exponents are among the most-measured numbers in physics. Nice if *one* determination were a theorem.
- **Practical.** The methods developed here are useful in ordinary workflows:
  - they catch failure modes that really occur;
  - they can make the numerics **faster**.
  - and more — sorry time constrained.

The Theorem

Theorem

Let $\mathcal{C}$ be a unitary 3d $\mathbb{Z}_2$-symmetric CFT with **(a)** one relevant $\mathbb{Z}_2$-odd scalar σ and $\Delta_\sigma \in [0.505, 0.5405]$; **(b)** at most one relevant $\mathbb{Z}_2$-even scalar; **(c)** a stress tensor at $\Delta = 3$; **(d)** next $\mathbb{Z}_2$-even spin-2 at $\Delta_{T'} \geq \frac{11}{2}$.

Then $\mathcal{C}$ contains *exactly one* relevant $\mathbb{Z}_2$-even scalar ϵ, and

$$(\Delta_\sigma, \Delta_\epsilon) \ \in \ (0.5055,\, 0.540) \times (1.23,\, 1.55).$$

And $\eta = 2\Delta_\sigma - 1$ and $\nu = 1/(3 - \Delta_\epsilon)$ obey

$$\eta \ \in \ (0.011,\, 0.080), \qquad \nu \ \in \ \left(\tfrac{100}{177},\, \tfrac{20}{29}\right) \ \approx \ (0.5650,\, 0.6897).$$

- Existence of ϵ is *proven*, not assumed.
- The rest of the talk explains how one gets to a statement like this.

The method

Everything that has to be proven is reduced to

- a short list of analytic inputs (standard CFT theorems, plus new scalar recursion relation coefficient bounds);
- a finite set of polynomial positivity conditions

$$P(\vec{x}) \ \geq \ 0 \quad \text{on } [0,1]^n, \qquad n \leq 3, \qquad (PSD)$$

  with P explicit and exact rational coefficients. Actually, pairs of rationals for reasons that will become apparent later.

The second kind is *decidable* — in exact arithmetic, by simple linear algebra.

1. First: how all four loopholes become statements of this form.
2. Then: how we certify such statements.

Setup: the plainest possible bootstrap problem

Single correlator $\langle\sigma\sigma\sigma\sigma\rangle$, unitary 3d $\mathbb{Z}_2$-symmetric CFT:

$$\sum_{\mathcal{O}\in\sigma\times\sigma} p_{\mathcal{O}}\, F_{\Delta,\ell}(u,v;\Delta_\sigma) = 0, \qquad p_{\mathcal{O}} = \lambda^2_{\sigma\sigma\mathcal{O}} \geq 0.$$

$\alpha\cdot\vec{\mathrm{Id}}(\Delta_\sigma) > 0, \qquad \alpha\cdot\vec{V}(\Delta_\sigma;\Delta,\ell) \geq 0$ on the assumed spectrum.

Assumptions: (a) exactly one relevant $\mathbb{Z}_2$-odd scalar σ; (b) *at most* one relevant $\mathbb{Z}_2$-even scalar; (c) stress tensor at $\Delta = 3$; (d) next $\mathbb{Z}_2$-even spin-2 at $\Delta_{T'} \geq \frac{11}{2}$.

$\Lambda = 9$ (15 components).

Region positivity: Extend the PSD condition in one variable.

The external dimension enters polynomially.

$$\vec{V}(\Delta_\sigma;\Delta,\ell)\Big|_i = \left(\tfrac{1}{4}\right)^{\Delta_\sigma} \times \left[\text{polynomial in } \Delta_\sigma,\ \deg \leq \Lambda\right]$$

- Blocks are Δ_σ-independent; Δ_σ enters only through $u^{\Delta_\sigma}, v^{\Delta_\sigma}$.
- The common factor $(1/4)^{\Delta_\sigma} > 0$ drops out of every sign condition.
- "Positive for all $\Delta_\sigma \in [\Delta_\sigma^-, \Delta_\sigma^+]$" is one more polynomial-positivity direction.

Region positivity applied to spin.

- L can be kept symbolic: The κ-truncated conformal block is a rational function of L and Δ.
- Positivity on real L imposes positivity on the desired subset. At worst makes bounds weaker. Never invalid.
- For example, every even $\ell \leq 100$ individually; $\ell \in [102, \infty)$ through one (PSD) condition.

Exact numbers, and compact domains

Limited irrationality. At the crossing-symmetric point $r_* = 3 - 2\sqrt{2} < \frac{1}{5}$.

- Everything lives in $\mathbb{Q}(\sqrt{2})$: pairs of rationals.
  Sign determination is one exact rational comparison: $a + b\sqrt{2}$ has the sign of a if $a^2 > 2b^2$, and the sign of b if $a^2 < 2b^2$. No decimal expansion of $\sqrt{2}$.

Compactification. Map every half-line onto $[0, 1]$ and clear denominators: $t = \tau_{\text{min}}/\tau$ for twists, $t = \ell_*/\ell$ for spins.

- $t = 0$ *is* the point at infinity: it becomes an ordinary *boundary face* of a closed box.
- Its value there is the sector's leading asymptotic coefficient data — finite, and what no finite sample constrains.

Loophole 4: the true blocks

We certify positivity on the true block by certifying positivity on the truncated block by sufficient margin:

$$V^{\text{true}} = V^{(\kappa)} + \left(V^{\text{true}} - V^{(\kappa)}\right) \;\geq\; V^{(\kappa)} - M \;\geq\; 0.$$

- M = a *certified* upper bound on our ignorance of the true block.
- Subtracting it keeps the polynomial positive by the required margin.
- The error is the ρ-truncation tail: levels $> \kappa$, evaluated at $r_* < \frac{1}{5}$, so each even pair is suppressed by $r_*^2 < \frac{1}{25}$.

Geometric smallness is *structural*; the work is turning it into an exact inequality with certified constants.

KING'S College LONDON

Bounding the truncation error: reduce to the diagonal

Diagonal ($z = \bar{z}$) stripped block: $h_{\Delta,\ell}(\rho, 1) = \sum_{m \geq 0} c_m(\Delta, \ell)\, \rho^m$, $c_0 = 1$. The recursion is organised by *radial level* m, and the derivative tables inherit that structure exactly:

$$\widehat{V}_i^{(\kappa)} = \sum_{m \leq \kappa} c_m(\Delta, \ell)\, \omega_{m,i}(\Delta, \ell).$$

$$i = 1, \ldots, N_\Lambda = 15$$

- $c_m \geq 0$: the *diagonal* radial coefficients
- Transverse (off-diagonal) derivatives are *reconstructed* from diagonal data by the Casimir recursion — through explicit rational weights.
- $\omega_{m,i}$: explicit *level maps* — rational in (Δ, ℓ), *independent of the block*, a finite chain of linear steps.

Truncation error $= \sum_{m > \kappa} c_m\, \omega_{m,i}$ — bound the c_m; linear algebra does the rest. (Majorants: same recursion with absolute weights $\Rightarrow \hat{\omega} \geq |\omega|$.)

The conceptual gem: a telescoping identity

- $\Delta_{12} = \Delta_{34} = 0$: Only even levels appear.
- This leaves only three terms (in $b_s = c_{2s}$);
- **whose middle coefficient is minus the sum of the outer two:**

$$P_0(s)\, b_s - \big(P_0(s) + P_2(s)\big)\, b_{s+1} + P_2(s)\, b_{s+2} = 0.$$

The conceptual gem: a telescoping identity

- $\Delta_{12} = \Delta_{34} = 0$: Only even levels appear.
- This leaves only three terms (in $b_s = c_{2s}$);
- **whose middle coefficient is minus the sum of the outer two:**

$$P_0(s)\, b_s - \big(P_0(s) + P_2(s)\big)\, b_{s+1} + P_2(s)\, b_{s+2} = 0.$$

$\Rightarrow$ First order recursion in the difference $D_s = b_{s+1} - b_s$:

$$D_{s+1} = \frac{P_0(s)}{P_2(s)} D_s \implies D_s = (c_2 - 1)\, \frac{(\frac{3}{2})_s (\Delta)_s (\frac{\tau+1}{2})_s (\frac{\Delta+\ell+2}{2})_s}{(2)_s (\Delta + \frac{1}{2})_s (\frac{\tau+2}{2})_s (\frac{\Delta+\ell+3}{2})_s} \sim \frac{\mathrm{const}}{s^2}$$
$$\implies 1 = c_0 \leq c_2 \leq c_4 \leq \cdots \nearrow c_\infty = 1 + \big(c_2 - 1\big)\, {}_5F_4[\cdots;1].$$

The conceptual gem: a telescoping identity

- $\Delta_{12} = \Delta_{34} = 0$: Only even levels appear.
- This leaves only three terms (in $b_s = c_{2s}$);
- **whose middle coefficient is minus the sum of the outer two:**

$$P_0(s)\, b_s - \big(P_0(s) + P_2(s)\big)\, b_{s+1} + P_2(s)\, b_{s+2} = 0.$$

$\Rightarrow$ First order recursion in the difference $D_s = b_{s+1} - b_s$:

$$D_{s+1} = \frac{P_0(s)}{P_2(s)} D_s \implies D_s = (c_2 - 1)\, \frac{(\frac{3}{2})_s (\Delta)_s (\frac{\tau+1}{2})_s (\frac{\Delta+\ell+2}{2})_s}{(2)_s (\Delta + \frac{1}{2})_s (\frac{\tau+2}{2})_s (\frac{\Delta+\ell+3}{2})_s} \sim \frac{\mathrm{const}}{s^2}$$

$$\implies \quad 1 = c_0 \leq c_2 \leq c_4 \leq \cdots \nearrow \; c_\infty = 1 + \big(c_2 - 1\big)\, {}_5F_4[\cdots; 1].$$

- Monotone and bounded.
- Summing gives a closed form for the diagonal block — a new derivation of [Rychkov, Yvernay '15].

KING'S College LONDON

All loopholes reduced to exact positivity of polynomials.

loophole	becomes
1. points $\to$ regions	one polynomial direction in Δ_σ
2. truncated $\to$ true blocks	subtract a certified envelope M
3. finite $\to$ all spins	one polynomial direction in real L
4. floats $\to$ exact	coefficients in $\mathbb{Q}(\sqrt{2})$, signs decided exactly
every half-line compactified: ∞ is an ordinary boundary face	

Per excluded box, what is left is a *finite* list of statements

$$P(\vec{x}) \;\geq\; 0 \quad \text{on } [0,1]^n, \qquad n \leq 3, \qquad P \in \mathbb{Q}(\sqrt{2})[\vec{x}].$$

A continuum of external dimensions, all twists, all spins and the true blocks —
replaced by finitely many machine-checkable sign conditions.

Certification architecture: search $\to$ certify

1. **Search.** Anything goes: e.g. floating-point LP/SDP. Output: a box B and an *exactly rational* $\hat{\alpha} \in \mathbb{Q}^N$.
   Nothing about it needs to be trusted.

2. **Certificate.** Verifies in *exact arithmetic* finitely many inequalities about explicit rational functions. Certifies or rejects.
   A rejected candidate triggers a new search, never results in a false theorem.

We discuss the certificate first — it is the only layer that has to be trusted.

Certifying positivity: the Bernstein basis

$$P(t) = \sum_{i=0}^{d} b_i\, B_i^d(t), \qquad B_i^d(t) = \binom{d}{i} t^i (1-t)^{d-i},$$

$$\text{with} \qquad B_i^d \;\geq\; 0 \quad \text{on } [0,1], \qquad \textstyle\sum_i B_i^d = 1.$$

$$\boxed{b_i \geq 0 \;\; \forall i \quad \Longrightarrow \quad P \geq 0 \;\; \text{on } [0,1].}$$

- Monomial $\to$ Bernstein is a *fixed, exact, rational linear map*. No rounding, no evaluation, no optimization.
- $P = 1 - t + t^2$: coefficients $(1, -1, 1)$, positivity invisible. Bernstein $(1, \frac{1}{2}, 1)$: $P = (1-t)^2 + \frac{1}{2} \cdot 2t(1-t) + t^2$.

Positivity has been made *syntactic*.

KING'S College LONDON

Bernstein procedure

- **Subdivision.** If some $b_i < 0$, bisect the interval and try again. Coefficients converge quadratically to values $\Rightarrow$ terminates on any strictly positive P.
- **Certifies counterexamples.** $b_0 = P(0)$, $b_d = P(1)$: a negative corner coefficient is an *exact certified counterexample*. Refutes positivity.
- Applies to $[0,1]^n$. We use $n \leq 3$ (τ, ℓ, Δ_σ.).

The search: Maximize positive margin!

Functionals found in conventional searches have *zero* slack — any perturbation flips a sign.
Instead maximize a **uniform relative margin**:

$$t^\star = \max_{\alpha} t \quad \text{s.t.} \quad \frac{\alpha \cdot \vec{V}(x)}{\|\vec{V}(x)\|_\infty} \geq t \quad \forall x, \qquad \alpha \cdot \vec{\mathrm{Id}} = 1.$$

$$\boxed{M < \; t^\star \implies \text{the verdict survives.}}$$

- In practice we find margins with ten orders to spare.
- Near the island boundary $t^\star \to 0$ — explicitly confirmed, as it should.

What the certificate taught the search

A Bernstein run either succeeds, or returns an exactly negative corner.

- Counterexample in a sector's *interior* = one violated inequality $\Rightarrow$ add one sample point.
- Counterexample *at* $t = 0$, i.e. at $\ell = \infty$: the asymptotic coefficient is negative $\Rightarrow \alpha \cdot \vec{V} < 0$ beyond some finite point: **infinitely many violated constraints at once**.
- Repair is structural: the sample set needs positivity conditions on the non-compact face!

With such coverage, all candidate tiles certified on the first pass.

The certified exclusion region

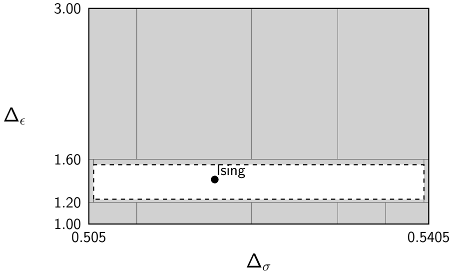

- $\Lambda = 9$, **27 boxes**.
- Boxes also certify positivity on *all* $\Delta_\epsilon \geq \frac{8}{5}$ — so they also exclude the **no-relevant-even-scalar** spectrum.
- Coverage checked exactly, in rational arithmetic.

The theorem

Theorem

Let $\mathcal{C}$ be a unitary 3d $\mathbb{Z}_2$-symmetric CFT with **(a)** exactly one relevant $\mathbb{Z}_2$-odd scalar σ and $\Delta_\sigma \in [0.505, 0.5405]$; **(b)** at most one relevant $\mathbb{Z}_2$-even scalar; **(c)** a stress tensor at $\Delta = 3$; **(d)** next $\mathbb{Z}_2$-even spin-2 at $\Delta_{T'} \geq \frac{11}{2}$.

Then $\mathcal{C}$ contains *exactly one* relevant $\mathbb{Z}_2$-even scalar ϵ, and

$$(\Delta_\sigma, \Delta_\epsilon) \ \in \ (0.5055,\, 0.540) \times (1.23,\, 1.55).$$

And $\eta = 2\Delta_\sigma - 1$ and $\nu = 1/(3 - \Delta_\epsilon)$ obey

$$\eta \ \in \ (0.011,\, 0.080), \qquad \nu \ \in \ \left(\tfrac{100}{177},\, \tfrac{20}{29}\right) \ \approx \ (0.5650,\, 0.6897).$$

What is, and is not, proved

- **Assumed**: the spectral assumptions including a gap in the spin-2 sector.
- The *uniqueness* of the $\mathbb{Z}_2$-even and odd scalar is never consumed by a solver — but is essential to read a paving of windows as a bound on the *dimensions*.
  With relevant even scalars at both 1.2 and 1.6, a theory evades each tile through a different operator.
- We don't show that the critical 3d Ising *is* a CFT within the window.
- Of course, experiment, perturbation theory, and lattice computations, and more all place it there.

Practical dividend: two-sided bounds get *cheaper*

Two-sided bound to resolution δ:

- **Standard:** map the island boundary $\Rightarrow O(\text{perimeter}/\delta)$ point solves — all spent *locating a boundary*, not *bounding a number*. And the conclusion is still an extrapolation.
- **Regions:** test a strip $[\Delta_\sigma^-, \Delta_\sigma^+] \times [e, e+h]$ and **bisect in** $e \Rightarrow O(\log 1/\delta)$ solves, and every affirmative answer *is* a bound.

$O(1/\delta) \longrightarrow O(\log 1/\delta)$: a change of optimality class.

Mixed correlators: the analytics are done

Target: $\sigma - \epsilon$ bootstrap, and *no spin-2 assumptions at all*.

Every analytic ingredient is already at certificate level:

- the seven-term recurrence with *symbolic* external parameters
- **external polynomiality:** at fixed truncation every entry of the mixed derivative tables is an exact *polynomial* in $(\Delta_{12}, \Delta_{34})$, degree $\leq 2\kappa + 2(N-1)$.
  A $(\Delta_\sigma, \Delta_\epsilon)$ box is just two more Bernstein directions.
- **certified coefficient majorants** $|c_m| \leq C \cdot 3^m$. A level costs $3r_* < \frac{3}{5}$, so tails still contract.

What remains is *compute*, not analysis: per-δ base-window certificates, of an already-certified type.

Mixed correlators: it also looks feasible numerically

- Numerics are in progress as we speak.
- Search side modification: PMP/SDPB rather than grid LPs (positivity on intervals by construction). Demanded by the certification!

KING'S College LONDON

Take-home

- Four concrete loopholes separate a bootstrap bound from a theorem. All four can be closed at once, with a small and classical trusted base.
- The 3d Ising exponents now have bounds that are theorems — modest ones, under explicit assumptions.
- Tools are generalizable to any bound.
- They are also practical:
  - *region bounds* turn island mapping into bisection;
  - *margins* replace convergence reruns with one inequality;
  - *audits* catch real failure modes.
- **A bound quoted with its error margin is more informative, and more credible, than the bound alone.**

Further directions

- **What is the best way to pave/bound?**
- Region bounds for *OPE* angles and magnitudes? — Required in $d = 4$ Non-Abelian Current Bootstrap.
- Higher Λ: very achievable. Scaling is fine.
- Verify the theorem in Lean?
- Spinning-block recursions error bounds? Blocks in other dimensions? Theorem grade bounds in larger and/or spinning mixed systems.
- Positive margin as a Navigator function for *region* bounds?

Three ways to make the continuum problem finite

We need $\alpha \cdot \vec{V}(x) \geq 0$ for a *continuum* of labels $x = (\text{sector}; \Delta_\sigma, \Delta, \ell)$.

1. **Sample everything ⇒ LP.** Positivity between samples not imposed — one hopes. What we ran (exact-rational LP support, unrestricted constraint types). Requires the audit loop.

2. **Keep Δ exact ⇒ PMP/SDPB.** The standard reduction; for region bounds one must still sample the externals.

3. **Keep Δ *and* externals exact ⇒ SOS-extended PMP.** Lukács: for *odd* deg $P = 2k+1$, $P \geq 0$ on $[\Delta_\sigma^-, \Delta_\sigma^+]$ iff
$$P(\Delta_\sigma) = (\Delta_\sigma - \Delta_\sigma^-)\, m^T G_0 m + (\Delta_\sigma^+ - \Delta_\sigma)\, m^T G_1 m, \qquad G_i \succeq 0,$$
Here $m = (1, \Delta_\sigma, \ldots, \Delta_\sigma^k)^T$ (even deg P has an analogous form). Promote the Gram entries to polynomials in Δ: matching powers of Δ_σ gives linear equations, and what remains is polynomial-*matrix* positivity in Δ — what SDPB does natively. Nothing sampled.

The 27 boxes

role	Δ_σ-range	Δ_ϵ-window
frame: left column	$[0.505, 0.5055]$	$[1.20, 1.60]$
frame: bottom row	$[0.5055, 0.5405]$	$[1.20, 1.23]$
frame: top row	$[0.5055, 0.5405]$	$[1.55, 1.60]$
frame: right column	$[0.540, 0.5405]$	$[1.20, 1.60]$
paving (5 tiles)	$[0.505, 0.5405]$	$[1.00, 1.20]$
gap (4 tiles)	$[0.505, 0.5405]$	$[\frac{8}{5}, \infty)$

- All endpoints exact rationals (denominators dividing 10^4). The frame rows/columns are subdivided into 18 tiles; the table shows their union.
- The gap tiles' window $[\frac{8}{5}, 3]$ merges with the $[3, \infty)$ continuum — one certificate covers $\Delta_\epsilon \in [\frac{8}{5}, 3)$ *and* the no-relevant-even-scalar spectrum.
- Not paved: $\Delta_\epsilon \in [\frac{1}{2}, 1)$, where scalar windows dip below twist one. Traditionally also absent from Ising island plots.

Objections to the all-spin certificate

- *Does the interpolant mean anything at non-integer spin?*
  It need not. No physics is claimed between integers; the only load-bearing property is that at integer L it coincides with the standard evaluator.
- *Is "$\ell = \infty$" used as if an operator sat there?*
  No. What is required is termwise positivity at all *finite* spins with no cutoff — exactly what the certificate gives. The face is instrumental: if the normalized W had a negative limit, positivity would already fail at large finite spin.
- *Is positivity at all real L too strong?*
  Genuinely stronger than physics requires — which can only *weaken* bounds, never invalidate them. Empirically it costs nothing: on every production tile the searched functional certified on the full real-L domain, far-tail relative margins $\sim 10\%$.

Analyticity in ℓ?

One of the three pole towers *lowers* spin, $\ell \to \ell - n$. A spin can only be lowered ℓ times, so standard evaluators contains a conditional:

$$\texttt{keep the term only if } n \leq \ell$$

An object built through conditionals is *not* obviously a rational function of continuous L.

- But the residue carries $\gamma(n, L) \propto \dfrac{(1+L-n)_n}{(1+\zeta+L-n)_n}$, $\zeta = \frac{d-2}{2} = \frac{1}{2}$.
- Numerator $= (1+L-n)\cdots L$ vanishes at integer $L = \ell < n$; the denominator, offset by a half-integer, never does.
- Exactly the terms the conditional would drop, drop themselves.

Only even levels are populated

The pole recursion builds $h^{(\kappa)}$ from residues of *three* pole towers. At equal externals, $\Delta_{12} = \Delta_{34} = 0$:

- towers I and III shift the radial level by n and carry the factor
$$\left(\frac{1-n}{2}\right)_n = \frac{1-n}{2} \cdot \frac{3-n}{2} \cdots \frac{n-1}{2},$$
whose n consecutive arguments run through 0 whenever n is *odd* $\Rightarrow$ every odd-shift residue vanishes identically;
- tower II shifts by $2k$ — always even.

Hence $c_{\mathrm{odd}} \equiv 0$: the diagonal expansion is a series in ρ^2, and we work with $b_s := c_{2s}$.

- Each retained level pair costs $r_*^2 < \frac{1}{25}$ — this is where the geometric smallness of the tail comes from.
- The c_m obey a *seven-term* recurrence (the Hogervorst–Osborn–Rychkov diagonal ODE, rewritten in ρ and cleared of denominators). On the even subsequence it collapses — next slide.

What has to be trusted

- Exact rational arithmetic; Sturm's theorem; the two-line Bernstein implication.
- Three analytic facts about conformal blocks:
  - convergence of the OPE [Pappadopulo, Rychkov, Espin, Rattazzi '12];
  - positivity + convergence of the radial expansion [Hogervorst, Rychkov '13];
  - the structure of the block recursion [Kos, Poland, Simmons-Duffin '13].
- One coefficient-boundedness statement — the telescoping identity.

Solvers, grids, heuristics and our own search code are outside the trusted base *by construction*.

Disciplines: fail-closed (any undecided sign aborts); negative controls (each certifier must *fail* on falsified claims); coverage checks.