Exact Sharp Thresholds, Piecewise-Algebraic Squaring Laws, and Machine-Checked Thermodynamic Mass-Gap Certificates for the Reflection-Positive Purity Bootstrap

Lluis Eriksson*

November 24, 2026

Abstract

The reflection-positive purity bootstrap introduced for the Haldane-gap proof [14, 17] and extended to higher dimensions, quasi-degenerate ground states, fermions, and tensor networks by Grover [7] certifies a thermodynamic-limit spectral gap $\Delta_\infty > 0$ from finite-size Euclidean partition-function purities S and T. Existing analyses rely on conservative defect cutoffs ($u \le 0.01$ in $d = 1$ and $1 - Q \le 1/(2a^2)$ in d spatial dimensions with $a = 2^{d+1} - 1$) and yield a flat gap ceiling $\Delta_\infty > (\log 2)/\beta_0$ insensitive to the actual purity margin. Here we derive the *exact* extremal laws and sharp stability boundaries of the purity bootstrap in arbitrary dimension $d \ge 1$ and build a machine-checked, zero-floating-point-trust certification pipeline. First, we solve the simplex squaring minimization problem $\inf\{\sum w_i^4/(\sum w_i^2)^2 : w_i \ge 0,\ \sum w_i = 1,\ \sum w_i^2 = p\}$ in closed piecewise-algebraic form on all of $(0, 1]$; on $(1/2, 1]$ it collapses to $h_1^*(p) = 1 - (1-p)^2/(2p^2)$ via a three-line pure ring-algebra proof, accompanied by the uniquely attained spectral-gap enclosure $e^{\beta\Delta} \ge (1+\sqrt{2p-1})^2/(2(1-p))$. Second, exploiting the exact staggered space-time plaquette identity in a cyclic Gauss–Seidel recurrence, we derive the algebraic uniform threshold $p_c(d) = 1 - u_c^*(d)$ solving $2p^{2a} = (1-p)(\sum_{j=0}^{a-1} p^j)^2$, widening the 1D uniform basin by $12.73\times$ ($u_c^*(1) = 0.127311$) and enlarging the single-ratio Q threshold by $3.26\times$ to $7.63\times$ across $d = 1, \dots, 5$ (asymptotically $8\times$, scaling as $4/(a(a+1))$ vs. $1/(2a^2)$). Third, we prove an exact doubly-exponential contraction bound yielding $\Delta_\infty \ge \frac{1}{2^N \beta_0} \log(1/(K_d(u_N)u_N))$ for *all* intervening even hypercube lengths $\ell \ge L_0$, breaking the $(\log 2)/\beta_0$ ceiling by arbitrarily large factors. We formalize all 16 core algebraic theorems and lemmas in zero-dependency Lean 4 (`v4.35.0`, zero `sorry`, zero `Classical.choice`), verify rational logarithm mass-gap bounds strictly over $\mathbb{Q}$, and benchmark across the 1D transverse-field Ising model ($2.24\times$–$4.72\times$ closer to $h_c = 1$, up to $9.03\times$ tighter gaps), the $g = 2$ ferromagnetic phase via exact Kramers–Wannier sector duality, 2D ($D = 3$) lattices, and synthetic noisy intervals (0 false certifications across 200 critical and 200 near-boundary trials). On interacting spin chains, we prove an $\mathfrak{su}(3)$ determinant obstruction showing that the Spin-1 AKLT chain violates bond-reflection positivity for every $U \in U(3)$, and certify exact-$\mathbb{Q}$ thermodynamic gaps on two provably bond-reflection-positive non-integrable chains—Spin-1 Heisenberg with single-ion anisotropy ($L_0 = 4 \to 8$) and Spin-1/2 dimerized Heisenberg ($L_0 = 6 \to 12$)—where Grover's $u \le 0.01$ fails at every tested parameter.

1 Introduction

Establishing a uniform spectral gap $\Delta_\infty = \liminf_{L\to\infty}(E_1(L) - E_0(L)) > 0$ in the thermodynamic limit of interacting quantum many-body lattice Hamiltonians is a central open problem in

*Independent researcher (Sweden). Research, zero-dependency Lean 4 proofs, exact rational ($\mathbb{Q}$) certificate verifiers, and numerical benchmarks prepared with AI assistance (Google Antigravity) and publicly archived at https://github.com/lluiseriksson/sharp-purity-bootstrap.

mathematical physics [8, 9, 1, 2]. Classical finite-size criteria such as Knabe's bound [11] and its modern refinements [6, 13, 3] require the Hamiltonian to be frustration-free, excluding generic frustrated or non-commuting models such as the spin-1 Heisenberg antiferromagnet.

A recent breakthrough—the *reflection-positive purity bootstrap* [14, 17]—circumvents frustration-freeness by exploiting Osterwalder–Schrader reflection positivity [15, 5] of the Euclidean partition function $Z(L_1, \dots, L_d; \beta) = \mathrm{Tr}(e^{-\beta H_{L_1,\dots,L_d}})$. Writing normalized partition-function ratios ("purities") in spatial and temporal directions,

$$S(L,\beta) \coloneqq \frac{Z(2L,\beta)}{Z(L,\beta)^2}, \qquad T(L,\beta) \coloneqq \frac{Z(L,2\beta)}{Z(L,\beta)^2}, \tag{1}$$

reflection positivity ensures that both $S(L,\beta) = \sum_i w_i^2/(\sum_i w_i)^2$ (for the spatial transfer operator) and $T(L,\beta)$ (for the Gibbs transfer operator $e^{-\beta H_L}$) are purities of nonnegative probability distributions $w = (w_i)$ on the unit simplex, hence lie in $(0,1]$. Doubling along an axis squares the underlying spectrum ($w_i \mapsto w_i^2$), driving the purity closer to 1, while passing across orthogonal axes incurs a plaquette penalty. Grover [7] systematically generalized this bootstrap to arbitrary spatial dimension $d \geq 1$ ($D = d+1$ space-time axes), quasi-degenerate ground states ($g \geq 2$), fermionic parity sectors, and tensor-network states.

Despite its elegance, the current purity bootstrap suffers from four structural bottlenecks that severely limit practical certification on finite lattices:

1. **Loose squaring and gap-from-purity inequalities.** Prior bounds use suboptimal lower envelopes for $\sum w_i^4/(\sum w_i^2)^2$ and for the gap ratio w_0/w_1, losing a factor of up to 2 in the purity defect at every doubling step.
2. **Conservative multi-axis thresholds.** Grover [7] requires $u \leq 0.01$ in $d = 1$ and $1 - Q \leq 1/(2a^2)$ in d spatial dimensions (where $a = 2^{d+1} - 1$), discarding over 90% of the true certifiable region in 1D and a factor of up to 8× in higher dimensions due to simultaneous (Jacobi) rather than sequential (Gauss–Seidel) multi-axis updates.
3. **Flat thermodynamic gap ceiling.** Because existing proofs terminate once $T(L,\beta)$ crosses a fixed constant, they produce the flat bound $\Delta_\infty > (\log 2)/\beta_0$, which actually *degrades* as $1/\beta_0$ when one lowers temperature to enter the bootstrap basin—even when the initial purity defect is 10^{-6}!
4. **Unverified floating-point and finite-sample certification.** Numerical partition functions from exact diagonalization, tensor networks, or Monte Carlo carry floating-point rounding or statistical error; without monotone interval propagation and exact rational verification, a near-boundary rounding error could falsely certify a gapless critical point.

Summary of Contributions. We resolve all four bottlenecks with exact extremal mathematics, zero-dependency Lean 4 machine verification, and exact-$\mathbb{Q}$ numerical certificates:

- **Exact Piecewise-Algebraic Squaring Law & Gap Enclosures (Section 2):** Theorem 2.1 characterizes the global infimum $h^*(p) = \inf\{\sum w_i^4/(\sum w_i^2)^2 : \sum w_i = 1,\ \sum w_i^2 = p\}$ across all $p \in (0,1]$; on $(1/2,1]$, it coincides with the rational envelope $h_1^*(p) = 1 - (1-p)^2/(2p^2)$, which admits a purely algebraic three-line proof (Theorem 2.2). Theorem 2.3 establishes the uniquely sharp spectral-gap ratio $w_0/w_1 \geq (1+\sqrt{2p-1})^2/(2(1-p))$.
- **Exact D-Dimensional Cyclic Gauss–Seidel Thresholds (Section 3):** Leveraging the exact staggered space-time plaquette identity in a cyclic Gauss–Seidel iteration (Theorem 3.1), we prove that the uniform stability boundary $p_c(d) = 1 - u_c^*(d)$ is the unique root in $(1/2,1)$ of $2p^{2a} = (1-p)(\sum_{j=0}^{a-1} p^j)^2$, yielding $u_c^*(1) = 0.127311$ (12.73× vs. 0.01). For the single-ratio hypercube invariant Q, eliminating redundant leading-axis factors yields Theorem 3.2, improving $1 - Q_c^*(d)$ by 3.26× to 7.63× for $d = 1, \dots, 5$ and asymptotically 8× ($4/(a(a+1))$ vs. $1/(2a^2)$).

- **Unbounded Asymptotic Gap Bounds & Intervening Lengths (Section 4):** Theorem 4.1 proves explicit doubly-exponential decay $K_d u_{N+m} \le (K_d u_N)^{2^m}$, certifying $\Delta_\infty \ge \frac{1}{2^N \beta_0} \log(1/(K_d(u_N)u_N))$—scaling as $\Omega(\log(1/u_0)/\beta_0)$ and breaking the $(\log 2)/\beta_0$ ceiling. Theorem 4.2 and Proposition 4.3 transfer this exact lower bound from dyadic lengths $2^n L_0$ to *every* intervening even hypercube length $\ell \ge L_0$, and Corollary 4.4 handles the ferromagnetic 1D TFIM via exact Kramers–Wannier sector duality.
- **Machine-Checked Lean 4 Formalization, Exact-$\mathbb{Q}$ Verifier & Bond-RP Benchmarks (Sections 5–6):** All 16 core algebraic theorems and lemmas listed in Table 1 are formalized in zero-dependency Lean 4 (`v4.35.0`, 0 `sorry`, 0 `Classical.choice`). Theorem 5.1 guarantees fail-closed finite-sample interval propagation and sound downward-quantized $\mathbb{Q}$-certification including exact rational logarithm bounds. We benchmark the pipeline on the 1D TFIM, prove an $\mathfrak{su}(3)$ determinant no-go obstruction for bond-RP of the Spin-1 AKLT chain, and certify exact-$\mathbb{Q}$ gaps on two genuinely bond-RP non-integrable spin chains (Spin-1 Heisenberg+D and Spin-1/2 dimerized Heisenberg), plus 2D ($D = 3$) lattices and synthetic noisy intervals.

2 Exact Piecewise-Algebraic Squaring Law and Gap Enclosures

Let $\Delta_\infty^+ := \{w = (w_i)_{i\ge 0} : w_i \ge 0,\ \sum_i w_i = 1\}$ denote the countable probability simplex (for finite transfer matrices, all but finitely many w_i vanish), with purity $p(w) := \sum_i w_i^2 \in (0,1]$ and fourth moment $q(w) := \sum_i w_i^4$. Under one axis doubling, the normalized weights transform as $w_i \mapsto w_i^2/p(w)$, so the doubled purity is $p'(w) = q(w)/p(w)^2$.

Theorem 2.1 (Exact Piecewise-Algebraic Squaring Law on $(0,1]$)**.** *For every $p \in (0,1]$, let $k = \lfloor 1/p \rfloor \in \mathbb{N}_{\ge 1}$ so that $p \in (\frac{1}{k+1}, \frac{1}{k}]$, and define*

$$a_k(p) := \frac{k + \sqrt{k((k+1)p - 1)}}{k(k+1)}, \qquad b_k(p) := \frac{1 - \sqrt{k((k+1)p-1)}}{k+1}. \tag{2}$$

Then over all probability distributions $w \in \Delta_\infty^+$ satisfying $\sum_i w_i^2 = p$,

$$\inf_{\substack{w\in\Delta_\infty^+ \\ \sum w_i^2 = p}} \frac{\sum_i w_i^4}{p^2} \;=\; h^*(p) \;:=\; \frac{k\, a_k(p)^4 + b_k(p)^4}{p^2}, \tag{3}$$

and the infimum is uniquely attained (up to permutation and zero padding) by the $(k+1)$-point spectrum $w^* = (\underbrace{a_k(p), \dots, a_k(p)}_{k \text{ times}}, b_k(p), 0, 0, \dots)$.

Proof. Fix $p \in (0,1]$ and consider minimizing $f(w) = \sum_{i=1}^m w_i^4$ over the compact set $K_m(p) = \{w \in \mathbb{R}^m_{\ge 0} : \sum_i w_i = 1,\ \sum_i w_i^2 = p\}$ for arbitrary $m \ge k+1$. First, if all positive coordinates of a minimizer w are equal ($w_i = 1/r$ on its support of size r), then $p = 1/r$ is an exact reciprocal integer ($r = k$) and Cauchy–Schwarz gives $\sum w_i^4 \ge (\sum w_i^2)^2/r = p^3 = k(1/k)^4$, matching (3) with $b_k(1/k) = 0$. Otherwise, w has at least two distinct positive coordinates, so the gradients $(1, \dots, 1)$ and $(2w_1, \dots, 2w_m)$ are linearly independent on the positive support (LICQ holds). By Karush–Kuhn–Tucker stationarity, w admits multipliers $\lambda, \mu \in \mathbb{R}$ such that every positive coordinate $w_i > 0$ satisfies the depressed cubic

$$g(w_i) \;:=\; 4w_i^3 - 2\mu w_i - \lambda \;=\; 0. \tag{4}$$

Since a depressed cubic has at most two positive roots, every positive coordinate of w takes at most two distinct values $a > b > 0$. By Rolle's theorem on $g(x) = 4x^3 - 2\mu x - \lambda$ (whose derivative $g'(x) = 12x^2 - 2\mu$ has a single positive zero), we have $g'(a) = 12a^2 - 2\mu > 0$ and $g'(b) =$

$12b^2 - 2\mu < 0$. If w contained $s \geq 2$ copies of b, then the perturbation vector δ with $(+1, -1)$ on two b-coordinates and 0 elsewhere lies in the tangent subspace $\{\delta : \sum \delta_i = 0,\ \sum w_i \delta_i = 0\}$ of both equality constraints; by LICQ, the second-order necessary condition requires the Lagrangian Hessian form $\sum_i (12w_i^2 - 2\mu)\delta_i^2 = 2(12b^2 - 2\mu) \geq 0$, contradicting $g'(b) < 0$. Hence $s \leq 1$, so any minimizer has support $(\underbrace{a, \ldots, a}_{r}, b)$ with $a \geq b \geq 0$.

Eliminating $b = 1 - ra$ from $ra^2 + b^2 = p$ gives $r(r+1)a^2 - 2ra + (1-p) = 0$. Positivity of $b = 1 - ra \geq 0$ forces $1/p \geq r$ (with $r = 1/p$ when $b = 0$), while real roots require $r((r+1)p - 1) \geq 0$, i.e., $r \geq \lceil 1/p \rceil - 1$. For non-integer $1/p$ this uniquely selects $r = k = \lfloor 1/p \rfloor$, and for integer $1/p = k$ the boundary value $b = 0$ agrees continuously with $r = k - 1$ at $a = b = 1/k$. Taking the larger root $a \geq b$ yields (2). Since this lower bound is independent of m, passing to $m \to \infty$ proves (3) over all of Δ_∞^+; moreover, uniqueness on Δ_∞^+ follows because any minimizer $w \in \Delta_\infty^+$ with three distinct positive coordinates (or two copies of a non-maximal positive value) would violate the finite-support KKT conditions when fixing all remaining coordinates. □

When $p \in (1/2, 1]$ ($k = 1$), the square roots in (2) cancel identically, yielding a rational lower-bound envelope $h_1^*(p)$ that coincides with $h^*(p)$ on $(1/2, 1]$ and admits a three-line elementary proof requiring no calculus (formalized over arbitrary ordered fields in Lean 4; see Table 1).

Theorem 2.2 (Three-Line Algebraic Squaring Law on $(1/2, 1]$)**.** *For every finite or countable distribution $w \in \Delta_\infty^+$ with purity $p = \sum_i w_i^2 > 0$,*

$$\sum_i w_i^4 \ \geq\ p^2 - \frac{(1-p)^2}{2}, \qquad \textit{equivalently} \qquad \frac{\sum_i w_i^4}{p^2} \ \geq\ h_1^*(p) \ :=\ 1 - \frac{(1-p)^2}{2p^2}. \tag{5}$$

Moreover: (i) $h_1^*(p) = h^*(p)$ for all $p \in (1/2, 1]$ and equality holds identically on every two-level spectrum $(a, 1-a)$; (ii) $h_1^*(p)$ is strictly increasing on $(0, 1]$; and (iii) writing $u = 1 - p$, the defect contracts quadratically as $1 - h_1^*(1-u) = \frac{u^2}{2(1-u)^2}$.

Proof. Let $P_1 := \sum_{i<j} w_i w_j$ and $P_2 := \sum_{i<j} (w_i w_j)^2$. Expanding $(\sum_i w_i)^2 = 1$ and $(\sum_i w_i^2)^2 = p^2$ yields the two exact identities

$$1 - p = 2P_1, \qquad p^2 - \sum_i w_i^4 = 2P_2.$$

Since every term $w_i w_j \geq 0$, we have $P_2 = \sum_{i<j} (w_i w_j)^2 \leq \left(\sum_{i<j} w_i w_j\right)^2 = P_1^2 = (1-p)^2/4$, with equality whenever at most one pair (i, j) is nonzero (i.e., any two-level spectrum). Multiplying by 2 gives $p^2 - \sum_i w_i^4 \leq (1-p)^2/2$, which rearranges immediately to (5). Monotonicity of $h_1^*(p)$ on $(0, 1]$ follows since $(1-p)/p = p^{-1} - 1$ is strictly decreasing and nonnegative. □

Notice that $h_1^*(p)$ provides a valid lower bound for *all* $p \in (0, 1]$ (though it is sharp only on $(1/2, 1]$, where it agrees with $h^*(p)$, and becomes nonpositive below $p = \sqrt{2} - 1 \approx 0.414$). Compared to naive bounds of the form $1 - u' \geq 1 - u^2/(1-u)^2$ or linear defect relaxations, the exact factor of $\frac{1}{2}$ in $1 - h_1^*(1-u) = \frac{u^2}{2(1-u)^2}$ doubles the permissible defect at every single step. For brevity, we write $h^*(p) = h_1^*(p)$ throughout Sections 3–5, where all intermediate purities lie strictly in $(1/2, 1]$.

We next derive the exact sharp extraction of the spectral gap $\Delta = E_1 - E_0$ from a Gibbs purity $p = T(L, \beta) = \sum_i w_i^2$ where $w_i = e^{-\beta E_i}/Z(L, \beta)$ are sorted $w_0 \geq w_1 \geq \cdots \geq 0$.

Theorem 2.3 (Sharp Gap-from-Purity and Tail-Mass Enclosures)**.** *Let $w \in \Delta_\infty^+$ with $w_0 \geq w_1 \geq \cdots \geq 0$ and purity $p = \sum_{i \geq 0} w_i^2$. Then:*

(a) $w_0 \geq p$ and $1 - p \geq 2w_0(1 - w_0)$, hence $(2w_0 - 1)^2 \geq 2p - 1$. In particular, for $p \in (1/2, 1]$,

$$w_0 \ \geq\ \frac{1 + \sqrt{2p-1}}{2}, \qquad \sum_{i \geq 1} w_i \ =\ 1 - w_0 \ \leq\ \frac{1 - \sqrt{2p-1}}{2}. \tag{6}$$

(b) If $p \in (1/2, 1)$ and $w_1 > 0$, then

$$e^{\beta\Delta} \;=\; \frac{w_0}{w_1} \;\geq\; \frac{2w_0^2}{1-p} \;\geq\; \frac{(1+\sqrt{2p-1})^2}{2(1-p)}. \tag{7}$$

Every inequality in (6)–(7) is attained with exact equality on two-level spectra $(w_0, 1-w_0)$.

Proof. Since $w_i \leq w_0$ for all $i \geq 1$, $p = w_0^2 + \sum_{i\geq 1} w_i^2 \leq w_0(w_0 + \sum_{i\geq 1} w_i) = w_0$. Next, $\sum_{i\geq 1} w_i^2 \leq (\sum_{i\geq 1} w_i)^2 = (1-w_0)^2$, so $p \leq w_0^2 + (1-w_0)^2 = 1 - 2w_0(1-w_0)$, giving $1-p \geq 2w_0(1-w_0)$. Multiplying by 2 and subtracting from 1 yields $(2w_0-1)^2 = 1-4w_0(1-w_0) \geq 2p-1$. Since $w_0 \geq p > 1/2$, taking the positive square root gives (6). Finally, $w_1 \leq 1-w_0$ implies $2w_0w_1 \leq 2w_0(1-w_0) \leq 1-p$, so dividing by $w_1(1-p)/w_0 > 0$ gives $w_0/w_1 \geq 2w_0^2/(1-p)$, and substituting (6) gives (7). □

As $p \to 1^-$, (7) gives $w_0/w_1 \geq \frac{2}{1-p}(1 - O(1-p))$, which is asymptotically $2\times$ tighter in the ratio (an additive $+(\log 2)/\beta$ improvement in Δ) than the elementary bound $w_0/w_1 \geq p/(1-p)$ and exact on two-level spectra. All four algebraic core lemmas of Theorem 2.3 are Lean-verified (Table 1).

3 Cyclic Gauss–Seidel Recurrence and Sharp D-Dimensional Thresholds

3.1 Exact Space-Time Plaquette Identity and 1D Cyclic Gauss–Seidel Bootstrap

Consider a reflection-positive 1D quantum lattice Hamiltonian with Euclidean partition functions $Z(L,\beta) > 0$ for even lengths L and $\beta > 0$, with spatial purity $S(L,\beta) = Z(2L,\beta)/Z(L,\beta)^2 \in (0,1]$ and temporal purity $T(L,\beta) = Z(L,2\beta)/Z(L,\beta)^2 \in (0,1]$. On any elementary space-time plaquette $(L,\beta) \to (2L,2\beta)$, direct algebraic cancellation (Lean: `plaquette_exact_identity`) gives the exact identity

$$S(L,2\beta)\,T(L,\beta)^2 \;=\; \frac{Z(2L,2\beta)}{Z(L,\beta)^4} \;=\; T(2L,\beta)\,S(L,\beta)^2. \tag{8}$$

Because $T(L,\beta) \leq 1$ by reflection positivity, (8) immediately yields the sharp plaquette transfer inequality (Lean: `plaquette_lower_bound`)

$$S(L,2\beta) \;\geq\; S(L,\beta)^2\,T(2L,\beta). \tag{9}$$

Crucially, note that $S(L,2\beta)$ does *not* satisfy $S(L,2\beta) \geq S(L,\beta)^2$ in general (we verified counterexamples with $S(L,2\beta)/S(L,\beta)^2 \approx 0.981 < 1$ on reflection-positive two-mode transfer matrices); the factor $T(2L,\beta)$ in (9) is sharp and indispensable.

Now define the diagonalized scale-doubled sequence

$$S_n \coloneqq S(2^n L_0, 2^n \beta_0), \qquad T_n \coloneqq T(2^{n+1} L_0, 2^n \beta_0), \qquad n \geq 0.$$

Applying (9) to $(2^n L_0, 2^n \beta_0)$ yields $S(2^n L_0, 2^{n+1}\beta_0) \geq S_n^2 T_n$. Doubling the spatial axis from $2^n L_0$ to $2^{n+1} L_0$ at fixed $2^{n+1}\beta_0$ squares the spatial transfer-matrix spectrum, so by Theorem 2.2 and monotonicity of h^*:

$$S_{n+1} \;=\; S(2^{n+1}L_0, 2^{n+1}\beta_0) \;\geq\; h^*(S(2^n L_0, 2^{n+1}\beta_0)) \;\geq\; h^*(S_n^2 T_n). \tag{10}$$

Next, applying the transposed plaquette identity to $(2^{n+1}L_0, 2^n\beta_0)$ uses the *already-updated* spatial purity S_{n+1}:

$$T(2^{n+2}L_0, 2^n\beta_0) \;=\; \frac{T_n^2\, S_{n+1}}{S(2^{n+1}L_0, 2^n\beta_0)^2} \;\geq\; T_n^2 S_{n+1}.$$

Doubling the temporal axis from $2^n\beta_0$ to $2^{n+1}\beta_0$ at fixed length $2^{n+2}L_0$ squares the Gibbs spectrum, yielding the exact *cyclic Gauss–Seidel* update:

$$T_{n+1} \;=\; T(2^{n+2}L_0, 2^{n+1}\beta_0) \;\geq\; h^*(T_n^2 S_{n+1}). \tag{11}$$

Theorem 3.1 (1D Cyclic Gauss–Seidel Bootstrap and Exact Uniform Threshold)**.** *Define* $\Phi^{(1)}_{\mathrm{GS}}(s,t) = \big(h^*(s^2t),\ h^*(t^2h^*(s^2t))\big)$.

(a) ***Inductive Soundness & Monotonicity (Table 1).*** *If* $0 < s_0 \leq S_0$ *and* $0 < t_0 \leq T_0$*, and every iterate* $(s_k, t_k) = (\Phi^{(1)}_{\mathrm{GS}})^k(s_0, t_0)$ *for* $k \leq N$ *remains strictly positive with intermediate products in* $(0, 1]$*, then* $s_k \leq S_k$ *and* $t_k \leq T_k$ *for all* $k \leq N$.

(b) ***Exact Uniform Stability Threshold.*** *Along the symmetric diagonal* $s_0 = t_0 = p$*, the map converges to* $(1, 1)$ *if and only if* $p > p_c(1) \coloneqq 1 - u_c^*(1)$*, where* $p_c(1)$ *is the unique root in* $(1/2, 1)$ *of* $h^*(p^3) = p$*, equivalently the polynomial equation*

$$2p^6 \;=\; (1-p)(1+p+p^2)^2, \tag{12}$$

giving $p_c(1) = 0.8726892136\ldots$ *and* $u_c^*(1) = 0.1273107864\ldots$ *—a* $\mathbf{12.73}\times$ *wider uniform threshold than Grover's* $u \leq 0.01$ *[7].*

(c) ***Anisotropic Basin.*** *Even when* $s_0 < p_c(1)$*,* (s_0, t_0) *lies in the exact certifiable gapped basin* $\mathcal{B}^{(1)}_{\mathrm{gap}}$ *whenever* $t_0 > t_{\mathrm{crit}}(s_0)$ *(e.g.,* $s_0 = 0.80, t_0 > 0.9675$*, or* $s_0 = 0.785, t_0 > 0.9950$*).*

Proof. Part (a) follows by induction from (10)–(11) and monotonicity of h^* on $(0, 1]$. For (b), at a fixed point (p, p) of $\Phi^{(1)}_{\mathrm{GS}}$, both coordinates satisfy $p = h^*(p^3) = 1-(1-p^3)^2/(2p^6)$. Rearranging $1-p = (1-p)^2(1+p+p^2)^2/(2p^6)$ and cancelling $1-p > 0$ gives (12). On $(1/2, 1)$, the left-hand side $2p^6$ is strictly increasing from $1/32$ to 2, while $(1-p)(1+p+p^2)^2 = (1-p^3)(1+p+p^2)$ has derivative $-(3p^2)(1+p+p^2)+(1-p^3)(1+2p) < 0$ for all $p \in (0.65, 1)$ (and $2p^6 < (1-p)(1+p+p^2)^2$ on $(1/2, 0.65]$), ensuring a unique root $p_c(1) \approx 0.8726892136$. Since $\Phi^{(1)}_{\mathrm{GS}}$ is componentwise strictly increasing, $p > p_c(1)$ implies $h^*(p^3) > p$, forcing $(s_n, t_n) \uparrow (1, 1)$ monotonically. □

3.2 General D-Dimensional Cyclic Bootstrap and $8\times$ Single-Q Improvement

In d spatial dimensions ($D = d + 1$ space-time axes with initial extents $\vec{L}_0 = (L_{0,0}, \ldots, L_{0,D-1})$ where $L_{0,D-1} = \beta_0$), let $p_k(\vec{L}) \coloneqq Z(\vec{L} + L_k e_k)/Z(\vec{L})^2 \in (0, 1]$ denote the axis-k purity at geometry $\vec{L}$. Crucially, to obtain exact telescoping across all D orthogonal axes, we track the *staggered staircase* state vector at doubling scale $n \geq 0$:

$$p_k^{(n)} \;\coloneqq\; p_k(2^{n+1}L_{0,0},\ \ldots,\ 2^{n+1}L_{0,k-1},\ 2^nL_{0,k},\ \ldots,\ 2^nL_{0,D-1}), \qquad k = 0, \ldots, D-1. \tag{13}$$

Direct substitution yields the exact D-dimensional telescoping hypercube identity

$$Q^{(n)} \;\coloneqq\; \prod_{j=0}^{D-1} \big(p_j^{(n)}\big)^{2^{D-1-j}} \;=\; \frac{Z(2^{n+1}\vec{L}_0)}{Z(2^n\vec{L}_0)^{2^D}}. \tag{14}$$

For $k = 0$, comparing $Z(2^{n+1}\vec{L}_0)$ along the doubled axis 0 against $Z(2^{n+1}L_{0,0}, 2^nL_{0,1}, \ldots, 2^nL_{0,D-1})$ via the exact D-dimensional plaquette identity gives

$$p_0(2^nL_{0,0},\ 2^{n+1}L_{0,1},\ \ldots,\ 2^{n+1}L_{0,D-1}) \;\geq\; \frac{Z(2^{n+1}\vec{L}_0)}{Z(2^n\vec{L}_0)^{2^D}} \;=\; \prod_{j=0}^{D-1} \big(p_j^{(n)}\big)^{2^{D-1-j}},$$

since all denominator purities on the intervening hypercube faces lie in $(0, 1]$. Doubling axis 0 from $2^nL_{0,0}$ to $2^{n+1}L_{0,0}$ then squares the axis-0 transfer-matrix spectrum, yielding

$p_0^{(n+1)} \geq h^*(\prod_{j=0}^{D-1}(p_j^{(n)})^{2^{D-1-j}})$. Advancing cyclically along axes $k = 0, 1, \ldots, D-1$ shifts the staircase window by one axis at each step, yielding the exact cyclic Gauss–Seidel recurrence on $(p_0, \ldots, p_{D-1})$:

$$p_k^{\text{new}} \;=\; h^*\!\left(\prod_{j=0}^{D-1} \tilde{p}_{k,j}^{\,2^{D-1-j}}\right), \qquad \tilde{p}_k = (p_k, \ldots, p_{D-1}, p_0^{\text{new}}, \ldots, p_{k-1}^{\text{new}}), \tag{15}$$

where the exponents sum to $a \coloneqq \sum_{j=0}^{D-1} 2^j = 2^D - 1 = 2^{d+1} - 1$.

Theorem 3.2 (Sharp D-Dimensional Uniform and Single-Q Thresholds)**.** *Let $D = d + 1 \geq 2$ and $a = 2^D - 1$.*

(a) ***Exact Uniform Threshold.*** *A uniform initial staircase $p_0 = \cdots = p_{D-1} = p$ converges to $(1, \ldots, 1)$ under the cyclic Gauss–Seidel map (15) if and only if $p > p_c(d) = 1 - u_c^*(d)$, where $p_c(d)$ is the unique root in $(1/2, 1)$ of*

$$2p^{2a} \;=\; (1-p)\left(\sum_{j=0}^{a-1} p^j\right)^2, \qquad \textit{so} \qquad u_c^*(d) \;=\; \frac{2}{a^2} + O\!\left(\frac{1}{a^3}\right). \tag{16}$$

(b) ***Exact Single-Q Contraction and $8\times$ Asymptotic Gain (Table 1).*** *For the telescoping hypercube invariant $Q = \prod_{j=0}^{D-1} p_j^{2^{D-1-j}} \in (0, 1]$, at step $k \in \{0, \ldots, D-1\}$ of (15) the old trailing product satisfies $\prod_{j=k}^{D-1} p_j^{2^{D-1-j}} \geq Q$, whence raising to 2^k yields*

$$p_k^{\text{new}} \;\geq\; h^*\!\left(Q^{2^k} \prod_{r=0}^{k-1} (p_r^{\text{new}})^{2^{k-1-r}}\right). \tag{17}$$

Consequently, $Q^{\text{new}} \geq \Psi_D(Q) \coloneqq \prod_{k=0}^{D-1}(q_k(Q))^{2^{D-1-k}}$, where $q_k(Q)$ is defined recursively by the right-hand side of (17). The exact single-Q stability threshold $1 - Q_c^*(d)$ (unique root of $\Psi_D(Q) = Q$ on the valid domain where all $q_k > 1/2$) satisfies

$$1 - Q_c^*(d) \;=\; \frac{4}{a(a+1)} + O\!\left(\frac{1}{a^3}\right) \;\sim\; 8 \cdot \frac{1}{2a^2} \qquad (d \to \infty), \tag{18}$$

*delivering an asymptotic **8**$\times$ enlargement over Grover's threshold $1 - Q \leq 1/(2a^2)$ [7] (3.26$\times$ at $d = 1$, 5.28$\times$ at $d = 2$, 6.56$\times$ at $d = 3$, 7.26$\times$ at $d = 4$, 7.63$\times$ at $d = 5$).*

Proof. For part (a), setting $p_k^{\text{new}} = p$ in (15) gives $p = h^*(p^a) \iff (p^{-a} - 1)^2 = 2(1 - p)$, which rearranges to (16). Uniqueness of the root $p_c(d) \in (1/2, 1)$ for every $a \geq 3$ follows immediately by convexity: $L(p) \coloneqq (p^{-a} - 1)^2$ is strictly convex on $(0, 1]$ with $L(1) = 0, L'(1) = 0$, whereas $R(p) \coloneqq 2(1 - p)$ is linear with $R(1) = 0, R'(1) = -2 < 0$, and $L(1/2) = (2^a - 1)^2 \geq 49 > 1 = R(1/2)$, so $L - R$ has a unique zero in $(1/2, 1)$. Substituting $p = 1 - u$ yields $2(1 - 2au + O(u^2)) = a^2 u(1 - (a-1)u + O(u^2))$, so $u_c^*(d) = 2/a^2 + O(a^{-3})$.

For part (b), write $Q = H_k R_k$ where $H_k = \prod_{j=0}^{k-1} p_j^{2^{D-1-j}} \in (0, 1]$ and $R_k = \prod_{j=k}^{D-1} p_j^{2^{D-1-j}}$. Since $H_k \leq 1$, we have $Q \leq R_k$, hence $Q^{2^k} \leq R_k^{2^k} = \prod_{j=k}^{D-1} p_j^{2^{D-1-(j-k)}}$ (verified in Lean 4 as `hypercube_Q_step_sound`). Multiplying by the already-updated prefix $\prod_{r=0}^{k-1}(p_r^{\text{new}})^{2^{k-1-r}}$ and invoking monotonicity of h^* yields (17). To extract the exact leading asymptotics of $\Psi_D(1 - \epsilon)$, set $Q = 1 - \epsilon$ with $\epsilon = O(a^{-2})$. Since $Q^{2^k} = 1 - 2^k\epsilon + O(2^{2k}\epsilon^2)$, induction on k shows $1 - q_k = \frac{1}{2}(2^k\epsilon)^2 + O(2^{3k}\epsilon^3) = 2^{2k-1}\epsilon^2 + O(\epsilon^3)$ (the updated prefix defects are $O(\epsilon^2)$, subleading to $2^k\epsilon$). Therefore,

$$\begin{aligned} 1 - \Psi_D(1-\epsilon) \;&=\; \sum_{k=0}^{D-1} 2^{D-1-k}(1 - q_k) + O(a^3\epsilon^3) \;=\; \frac{\epsilon^2}{2}\sum_{k=0}^{D-1} 2^{D-1+k} + O(a^3\epsilon^3) \\ &=\; \frac{2^{D-1}(2^D - 1)}{2}\,\epsilon^2 + O(a^3\epsilon^3) \;=\; \frac{a(a+1)}{4}\,\epsilon^2 + O(a^3\epsilon^3). \end{aligned}$$

Equating $1-\Psi_D(1-\epsilon)=\epsilon$ yields $\epsilon_c^* = 4/(a(a+1))+O(a^{-3})$, which is asymptotically $8\times(1/(2a^2))$. Physically, two independent factors of 2 and 4 combine: our sharp h^* saves a factor of 2 per squaring, and replacing Grover's simultaneous bound Q^a by the exact trailing-product power Q^{2^k} replaces $\sum_k 2^{D-1-k}a^2 = a^3$ by $\sum_k 2^{D-1-k}2^{2k} = a(a+1)/2$, saving an additional factor of 4. □

4 Doubly-Exponential Decay, Unbounded Gap Bounds, and Intervening Lengths

4.1 Breaking the $(\log 2)/\beta_0$ Ceiling

Once a finite number N of exact Gauss–Seidel steps brings the maximum axis defect $u_N \coloneqq \max_{0\le j<D}(1-p_j^{(N)})$ into the quadratic contraction zone, we obtain an explicit, unbounded thermodynamic mass-gap lower bound.

Theorem 4.1 (Explicit Doubly-Exponential Contraction and Unbounded Mass Gap)**.** *Let $p^{(N)} = (p_0^{(N)},\dots,p_{D-1}^{(N)})$ be valid lower bounds on the space-time purities at doubling scale N (lengths $2^N L_0$, inverse temperature $2^N\beta_0$), with defect $u_N = \max_j(1-p_j^{(N)}) \in (0,1/2)$ and $p_{\min} = 1-u_N$. Define the explicit rational contraction constant*

$$K_d(u_N) \;\coloneqq\; \frac{a^2}{2\,p_{\min}^{2a}} \qquad (a = 2^{d+1}-1), \qquad \rho_N \;\coloneqq\; K_d(u_N)\,u_N. \tag{19}$$

If $\rho_N < 1$, then for every $m \ge 0$ the defects satisfy $u_{N+m+1} \le K_d(u_N)\,u_{N+m}^2$, and therefore (Lean: `doubly_exponential_contraction`):

$$u_{N+m} \;\le\; \frac{1}{K_d(u_N)}\,\rho_N^{2^m} \qquad \forall m \ge 0. \tag{20}$$

Consequently, along the dyadic sequence $L_n = 2^n L_0$, the thermodynamic mass gap satisfies

$$\Delta_{\mathrm{dyad}} \;\coloneqq\; \liminf_{n\to\infty}\Delta(2^n L_0) \;\ge\; \frac{1}{2^N\beta_0}\log\!\left(\frac{1}{\rho_N}\right) \;=\; \frac{1}{2^N\beta_0}\log\!\left(\frac{1}{K_d(u_N)\,u_N}\right). \tag{21}$$

Proof. Suppose $u_n \le u_N$ (so $p_j^{(n)} \ge 1-u_n \ge p_{\min}$ for all j). By inner induction over the cyclic axis updates $k = 0,\dots,D-1$: assuming the already-updated prefix satisfies $p_r^{(n+1)} \ge 1-u_n \ge p_{\min}$ for $r<k$, Bernoulli's inequality ensures the a-fold product $\Pi_k = \prod_j \tilde p_{k,j}^{\,2^{D-1-j}}$ satisfies $p_{\min}^a \le \Pi_k \le 1$ and $1-\Pi_k \le a u_n$. Applying Theorem 2.2 yields

$$1-p_k^{(n+1)} \;\le\; \frac{(1-\Pi_k)^2}{2\Pi_k^2} \;\le\; \frac{a^2u_n^2}{2p_{\min}^{2a}} \;=\; K_d(u_N)\,u_n^2 \;\le\; \rho_N u_n \;<\; u_n \;\le\; u_N,$$

closing the inner induction on k and proving $u_{n+1} \le K_d(u_N)u_n^2 \le u_N$. Outer induction on m (Lean: `doubly_exponential_contraction`) then gives $K_d u_{N+m} \le (K_d u_N)^{2^m} = \rho_N^{2^m}$. At scale $n = N+m$, the temporal purity at length $L_{n+1} = 2^{n+1}L_0$ and inverse temperature $\beta_n = 2^n\beta_0$ satisfies $1-T(L_{n+1},\beta_n) \le u_{N+m}$. Applying Theorem 2.3(b) yields

$$\Delta(L_{n+1}) \;\ge\; \frac{1}{\beta_n}\log\!\left(\frac{2(1-u_{N+m})^2}{u_{N+m}}\right) \;\ge\; \frac{1}{2^{N+m}\beta_0}\log\!\left(\frac{2K_d(u_N)(1-u_{N+m})^2}{\rho_N^{2^m}}\right).$$

Taking $\liminf_{m\to\infty}$ cancels the $O(2^{-m})$ prefactor logarithm and leaves (21) identically. □

Unlike Grover's flat bound $\Delta_\infty > (\log 2)/\beta_0$, equation (21) grows without bound as $\Omega(\log(1/u_0)/\beta_0)$ when the initial defect $u_0 \to 0$. In particular, when a model is deeply gapped ($u_0 \sim e^{-\beta_0\Delta}$), $\frac{1}{\beta_0}\log(1/u_0) = \Theta(\Delta)$ recovers a constant $O(1)$ fraction of the true physical mass gap independent of β_0!

4.2 Transfer from Dyadic Scales to All Intervening Even Lengths and $g = 2$ Phases

A subtle mathematical requirement in any doubling bootstrap is ensuring that the gap does not collapse on non-dyadic system sizes $\ell \notin \{2^n L_0\}$. The following theorem transfers the exact rate (21) to *every* intervening even hypercube length $\ell \geq L_0$ with zero loss in the asymptotic exponent.

Theorem 4.2 (1D Intervening-Length Purity Transfer)**.** *Let $d = 1$. For any even length $\ell \in [L, 2L]$ with $L = 2^n L_0$, let μ_β and $\mu_{2\beta}$ denote the largest eigenvalues (spectral radii) of the positive spatial transfer matrices at inverse temperatures β and 2β, so that $Z(L', \beta) = \mu_\beta^{L'}(1 + R(L', \beta))$ with $R(L', \beta) \geq 0$. Since $\mu_{2\beta} \leq \mu_\beta^2$ (following immediately from $Z(L', 2\beta) \leq Z(L', \beta)^2$ as $L' \to \infty$), $\ell \leq 2L$ implies $(\mu_{2\beta}/\mu_\beta^2)^{2L} \leq (\mu_{2\beta}/\mu_\beta^2)^\ell$, and therefore*

$$T(\ell, \beta) \;\geq\; \frac{T(2L, \beta)}{\big(1 + R(\ell, \beta)\big)^2 \big(1 + R(2L, 2\beta)\big)}. \tag{22}$$

In particular, if $S_n \geq 1 - u_n$ and $T_n \geq 1 - u_n$ satisfy the doubly-exponential contraction (20), then $R(\ell, \beta_n) \leq \frac{u_n}{2(1-u_n)^2}$ and $R(2L, 2\beta_n) \leq \frac{u_{n+1}}{2(1-u_{n+1})^2}$ for every even $\ell \in [2^n L_0, 2^{n+1} L_0]$, so the exact lower bound (21) holds for all *even lengths:*

$$\Delta_\infty \;\coloneqq\; \liminf_{\substack{\ell \to \infty \\ \ell \text{ even}}} \Delta(\ell) \;\geq\; \frac{1}{2^N \beta_0} \log\!\left(\frac{1}{K_1(u_N)\, u_N}\right). \tag{23}$$

Proof. Using $Z(\ell, 2\beta) \geq \mu_{2\beta}^\ell$, $Z(\ell, \beta) \leq (1 + R(\ell, \beta))\mu_\beta^\ell$, $Z(2L, 2\beta) \leq (1 + R(2L, 2\beta))\mu_{2\beta}^{2L}$, and $Z(2L, \beta) \geq \mu_\beta^{2L}$, we obtain

$$\begin{aligned}
\frac{T(2L, \beta)}{(1 + R(\ell, \beta))^2 (1 + R(2L, 2\beta))} \;&\leq\; \frac{1}{(1 + R(\ell, \beta))^2} \left(\frac{\mu_{2\beta}}{\mu_\beta^2}\right)^{2L} \\
&\leq\; \frac{1}{(1 + R(\ell, \beta))^2} \left(\frac{\mu_{2\beta}}{\mu_\beta^2}\right)^{\ell} \;\leq\; \frac{Z(\ell, 2\beta)}{Z(\ell, \beta)^2} \;=\; T(\ell, \beta).
\end{aligned}$$

This exact algebraic chain is machine-checked in Lean 4 (`intervening_length_purity_transfer`). Now fix $n = N + m$ with $L = 2^n L_0$, $\beta_n = 2^n \beta_0$, and any even $\ell \in [L, 2L]$. Writing the spatial transfer-matrix eigenvalues at inverse temperature β_n as $\lambda_0 = \mu_{\beta_n} \geq \lambda_1 \geq \cdots \geq 0$, the remainder $R(\ell, \beta_n) = \sum_{\alpha \geq 1} (\lambda_\alpha/\lambda_0)^\ell$ is monotonically nonincreasing in ℓ. Since $S_n = S(L, \beta_n)$ is the purity of the normalized weights $w_\alpha^{(L)} = \lambda_\alpha^L / Z(L, \beta_n)$, Theorem 2.3(a) gives $1 - w_0^{(L)} \leq (1 - S_n)/(2w_0^{(L)}) \leq u_n/(2(1 - u_n))$ and hence

$$R(\ell, \beta_n) \;\leq\; R(L, \beta_n) \;=\; \frac{1 - w_0^{(L)}}{w_0^{(L)}} \;\leq\; \frac{u_n}{2(1 - u_n)^2}.$$

Similarly, $S_{n+1} = S(2L, 2\beta_n) \geq 1 - u_{n+1}$ implies $R(2L, 2\beta_n) \leq u_{n+1}/(2(1 - u_{n+1})^2) = O(u_n^2)$. Substituting these bounds together with $T(2L, \beta_n) = T_n \geq 1 - u_n$ into (22) yields $1 - T(\ell, \beta_n) \leq (2 + o(1))u_n \leq 3u_n$ for all sufficiently large n uniformly in $\ell \in [2^n L_0, 2^{n+1} L_0]$. Applying Theorem 2.3(b) at inverse temperature $\beta_n = 2^{N+m}\beta_0$ gives

$$\Delta(\ell) \;\geq\; \frac{1}{2^{N+m}\beta_0} \log\!\left(\frac{2(1 - 3u_{N+m})^2}{3u_{N+m}}\right) \;\geq\; \frac{1}{2^{N+m}\beta_0} \log\!\left(\frac{2K_1(u_N)(1 - 3u_{N+m})^2}{3\,\rho_N^{2^m}}\right),$$

and taking $\liminf_{\ell \to \infty}$ ($m \to \infty$) eliminates the $O(2^{-m})$ prefactor logarithm, establishing (23) identically. □

Proposition 4.3 (Extension to Intervening Even Hypercubes in $d \geq 2$)**.** *Let $d \geq 2$ and suppose the D-dimensional staircase purities satisfy (20) with $\rho_N = K_d(u_N)u_N < 1$. Then for every sequence of even hypercubes $(\ell, \ldots, \ell)$ with $\ell \in [2^{n+1}L_0, 2^{n+2}L_0]$ $(n = N + m)$, the temporal purity $T(\ell, \ldots, \ell; \beta_n)$ obeys $1 - T(\ell, \ldots, \ell; \beta_n) \leq O(u_n)$, and consequently*

$$\Delta_\infty \;\coloneqq\; \liminf_{\substack{\ell\to\infty \\ \ell \text{ even}}} \Delta(\ell, \ldots, \ell) \;\geq\; \frac{1}{2^N \beta_0} \log\!\left(\frac{1}{K_d(u_N)\, u_N}\right). \tag{24}$$

Proof. We transfer from the dyadic hypercube $(2L, \ldots, 2L)$ $(L = 2^{n+1}L_0,\ \beta = \beta_n = 2^n\beta_0)$ to $(\ell, \ldots, \ell)$ one spatial axis at a time $(j = d-1, d-2, \ldots, 0)$, keeping a fixed β_n. At stage j, the other spatial axes are fixed at $(2L, \ldots, 2L, \ell, \ldots, \ell)$, and applying the 1D algebraic inequality (22) along axis j bounds both the temporal purity T and every remaining spatial purity p_i $(i < j)$ at the new mixed geometry in terms of the axis-j remainders R_j at lengths L and $2L$ (note that (22) holds identically when the role of $\beta \to 2\beta$ is played by doubling any orthogonal spatial axis $i < j$!). At the initial stage $(j = d-1)$, the scale-$(n+1)$ staircase purities give $p_i(2^{n+1}L_0, \ldots; \beta_n) \geq 1 - O(u_n)$ via the exact plaquette step, and each axis reduction multiplies the purity by $1 - O(u_n)$ (since $R_j(\ell) \leq R_j(L) \leq O(1 - p_j)$ at the mixed geometry). After d finite steps, all d spatial axes reach ℓ with $1 - T(\ell, \ldots, \ell; \beta_n) \leq C_d u_n$ for a d-dependent constant C_d. Taking $\frac{1}{2^{N+m}\beta_0} \log(2(1 - C_d u_{N+m})^2/(C_d u_{N+m}))$ and letting $m \to \infty$ cancels $\log C_d/2^m \to 0$ and yields (24). □

Corollary 4.4 (Ferromagnetic 1D TFIM ($g = 2$ Doublet) via Kramers–Wannier Sector Duality)**.** *Consider the periodic ferromagnetic 1D transverse-field Ising chain $H(h) = -\sum_{i=1}^{L} Z_i Z_{i+1} - h\sum_{i=1}^{L} X_i$ with $h \in (0, 1)$, whose ground state is asymptotically $g = 2$-fold degenerate. Naive restriction to the odd-parity partition function $Z_{\text{odd}}(L, \beta)$ can fail spatial reflection positivity ($S_{\text{odd}} > 1$). However, exact Kramers–Wannier duality [12, 7] establishes exact isospectral sector equivalences between $H(h)$ and the dual paramagnetic chain $H_{\text{dual}} \coloneqq h\, H(1/h)$ ($h_{\text{dual}} = 1/h > 1$, dual inverse temperature $\beta_{\text{dual}} = h\beta$):*

$$(\textit{Ferro, PBC, even}) \cong (\textit{Dual, PBC, even}), \qquad (\textit{Ferro, APBC, even}) \cong (\textit{Dual, PBC, odd}).$$

Since H_{dual} is reflection-positive in both space and time with unique paramagnetic ground state, certifying a dual exact-$\mathbb{Q}$ purity-bootstrap bound $\gamma_{\text{dual}} > 0$ at $(L_0, \beta_{\text{dual}})$ rigorously certifies the uniform spectral gap of H_{dual}, $\Delta_{\text{dual}} = \liminf_{\ell\to\infty}(E_1^{\text{dual}}(\ell) - E_0^{\text{dual}}(\ell)) \geq \gamma_{\text{dual}}/\beta_0$. By the exact sector isomorphism above, this directly lower-bounds both the single-kink antiperiodic excitation gap ($\Delta_{\text{kink}} = 2(1-h)$) and the even-sector periodic excitation gap above the $g = 2$ doublet (which for periodic 1D TFIM is the two-kink gap $\Delta_{\text{PBC}} = 4(1-h) > \Delta_{\text{kink}}$).

5 Machine-Checked Lean 4 Formalization and Exact-$\mathbb{Q}$ Interval Certification

To eliminate floating-point rounding bugs and hidden proof gaps, we decouple numerical partition-function evaluation from thermodynamic-limit certification via a fail-closed two-layer architecture:

1. **Zero-Dependency Lean 4 Kernel (`lean/PurityBootstrap.lean`).** All 16 core theorems are formalized parametrically over an arbitrary linearly ordered commutative field `OrdField K` (simultaneously instantiating both $\mathbb{R}$ and $\mathbb{Q}$) with zero external dependencies (no Mathlib), zero `sorry`, and zero `Classical.choice`. Table 1 provides the complete Claim-to-Theorem Manifest audited by `lean/AuditAxioms.lean`.

2. **Zero-Float-Trust Exact-$\mathbb{Q}$ Certificate Verifier (`src/purity_bootstrap.py`).** Given finite-sample or numerical purity estimates $\hat{p}_j$ with certified error radius $\epsilon_j \geq 0$, we form exact rational lower bounds $p_j^{(0)} = \lfloor(\hat{p}_j - \epsilon_j) \cdot 10^9\rfloor/10^9 \in \mathbb{Q}$. At every cyclic Gauss–Seidel

Table 1: **Claim-to-Theorem Verification Manifest** (`leanprover/lean4:v4.35.0`, 0 external dependencies, 0 `sorry`, 0 `Classical.choice`). Every symbol is checked by `lake build && lake env lean AuditAxioms.lean`.

Paper Result	Lean 4 Symbol (`PurityBootstrap.lean`)	Scope Formalized	Audited Axioms
Thm. 2.2	`squaring_inequality_k1`	$\sum w_i^4 \geq p^2 - (1-p)^2/2$ on Δ^+	`propext, Quot.sound`
Thm. 2.2	`squaring_inequality_k1_sharp`	Exact equality on all 2-level spectra	none
Thm. 2.2	`squaring_inequality_ratio`	Normalized $q/p^2 \geq h^*(p)$ for $p > 0$	`propext, Quot.sound`
Thm. 2.2, 5.1	`hStar_monotone`	$0 < p \leq q \leq 1 \Rightarrow h^*(p) \leq h^*(q)$	none
Thm. 2.3(a)	`dominant_weight_ge_purity`	Dominant weight $w_0 \geq p$	`propext, Quot.sound`
Thm. 2.3(a)	`purity_defect_ge_two_w0_tail`	Sharp defect $1 - p \geq 2w_0(1-w_0)$	none
Thm. 2.3(a)	`sharp_discriminant_lower_bound`	$(2w_0 - 1)^2 \geq 2p - 1$	none
Thm. 2.3(b)	`sharp_gap_ratio_bound`	$w_0/w_1 \geq 2w_0^2/(1-p)$	none
Eq. (8)	`plaquette_exact_identity`	$S_{01}T_{00}^2 = T_{10}S_{00}^2$	none
Eq. (9)	`plaquette_lower_bound`	$S_{00}^2 T_{10} \leq S_{01}$	none
Thm. 3.1(a)	`gs_1d_step_sound`	Single 1D cyclic GS step soundness	none
Thm. 3.1(a)	`gs_1d_induction_sound`	N-step inductive 1D GS soundness	none
Thm. 3.1, 5.1	`gs_1d_monotone`	Componentwise interval monotonicity	none
Thm. 3.2(b)	`hypercube_Q_step_sound`	$Q^{2^k} \leq R_k^{2^k}$ in arbitrary D	none
Thm. 4.1	`doubly_exponential_contraction`	$u_{n+1} \leq Cu_n^2 \Rightarrow Cu_n \leq (Cu_0)^{2^n}$	`propext, Quot.sound`
Thm. 4.2	`intervening_length_purity_transfer`	Exact $T(\ell, \beta)$ transfer bound (22)	none

step, $h_1^*(q)$ is evaluated in exact arbitrary-precision $\mathbb{Q}$ (`fractions.Fraction`) and soundly rounded *downward* to denominator 10^{12} ($\lfloor x \cdot 10^{12} \rfloor / 10^{12} \leq x$). Crucially, even the final thermodynamic mass-gap lower bound $\Delta_{\text{cert}}^{\mathbb{Q}} \in \mathbb{Q}$ is verified with zero floating-point trust: writing $x = 1/\rho_N = 2^k r$ with $k \in \mathbb{N}$ and $r \in [1,2) \cap \mathbb{Q}$, every omitted term in the positive series $\log x = k \log 2 + 2\sum_{j=0}^{\infty} \frac{t_r^{2j+1}}{2j+1}$ ($t_r = (r-1)/(r+1) \in [0, 1/3)$) is strictly positive, so the M-term rational partial sum gives a provable exact-$\mathbb{Q}$ lower bound $\log_{\mathbb{Q}}(1/\rho_N) < \log(1/\rho_N)$, and the verifier checks $\gamma_{\text{lb}}^{\mathbb{Q}} \leq \log_{\mathbb{Q}}(1/\rho_N)/2^N$ and $\Delta_{\text{cert}}^{\mathbb{Q}} \leq \gamma_{\text{lb}}^{\mathbb{Q}}/\beta_0$ strictly in $\mathbb{Q}$ (rejecting any tampered rational or float gap field).

Theorem 5.1 (Soundness of Downward-Quantized Exact-$\mathbb{Q}$ and Interval Certification)**.** *Suppose true finite-size purities satisfy $p_j^{\text{true}} \geq p_j^{(0)} \in \mathbb{Q} \cap (1/2, 1]$ (with probability at least $1-\alpha$ for finite-sample confidence intervals). If the downward-quantized exact-$\mathbb{Q}$ trace $(p^{(0)}, \ldots, p^{(N)})$ satisfies $p_k^{(n+1)} \leq h_1^*(\Pi_k^{(n)})$ for every step $n < N$ and axis k, with all intermediate products $> 1/2$, satisfies $0 < u_N$ and $K_d(u_N)u_N < 1$ in $\mathbb{Q}$, and satisfies $0 < \Delta_{\text{cert}}^{\mathbb{Q}} \leq \log_{\mathbb{Q}}(1/(K_d(u_N)u_N))/(2^N\beta_0)$ in $\mathbb{Q}$, then $\Delta_{\text{cert}}^{\mathbb{Q}}$ is a valid lower bound on Δ_∞ (with probability at least $1-\alpha$).*

Proof. Immediate by induction from `hStar_monotone`, `gs_1d_monotone`, and Theorem 4.1: downward quantization at each step only decreases $p_k^{(n+1)}$, which preserves the invariant $p_k^{(n)} \leq p_k^{\text{true},(n)}$ at every subsequent step. □

6 Numerical Experiments and Benchmarks

All experiments are deterministically reproduced by `python scripts/run_all_experiments.py` and independently checked by `python scripts/verify_all.py`, which replays all 7 exported `.cert.json` files strictly over $\mathbb{Q}$ (including exact rational logarithm mass-gap bounds) with zero floating-point trust.

6.1 Experiment 1: 50,000-Trial Falsification Audit and Sharp D-Dimensional Thresholds

We first subjected Theorem 2.1 and Theorem 2.3 to an adversarial 50,000-trial stress test across Dirichlet distributions of dimension $m \in [2, 15]$ and concentration $\alpha \in [10^{-1.5}, 10]$, covering

Table 2: **Exact D-Dimensional Critical Purity-Defect Thresholds** vs. Grover [7] ($a = 2^{d+1} - 1$), comparing like-for-like hypercube-product defects $1 - Q$ alongside per-axis uniform defects $u_c^*(d)$.

d (D)	a	Per-Axis $u_c^*(d)$	Uniform $1 - p_c(d)^a$	Uniform Q-Gain	Single-Q $1 - Q_c^*(d)$	Grover $1/(2a^2)$	Single-Q Gain
1 (2)	3	1.2731×10^{-1}	3.3537×10^{-1}	11.29×[†]	1.8100×10^{-1}	5.5556×10^{-2}	3.26×
2 (3)	7	3.1474×10^{-2}	2.0057×10^{-1}	19.66×	5.3913×10^{-2}	1.0204×10^{-2}	5.28×
3 (4)	15	7.8294×10^{-3}	1.1122×10^{-1}	50.05×	1.4585×10^{-2}	2.2222×10^{-3}	6.56×
4 (5)	31	1.9543×10^{-3}	5.8840×10^{-2}	113.1×	3.7797×10^{-3}	5.2029×10^{-4}	7.26×
5 (6)	63	4.8836×10^{-4}	3.0305×10^{-2}	240.6×	9.6100×10^{-4}	1.2598×10^{-4}	7.63×

[†]In $d = 1$, uniform Q-Gain compares $1 - p_c(1)^3 = 0.3354$ against Grover's 1D uniform box $1 - 0.99^3 = 0.0297$ (per-axis gain 12.73×); for $d \geq 2$, both Uniform Q-Gain and Single-Q Gain compare against Grover's $1/(2a^2)$ (corresponding per-axis gain vs. $1 - (1 - \frac{1}{2a^2})^{1/a}$ is 21.5× at $d = 2$ to 244.2× at $d = 5$).

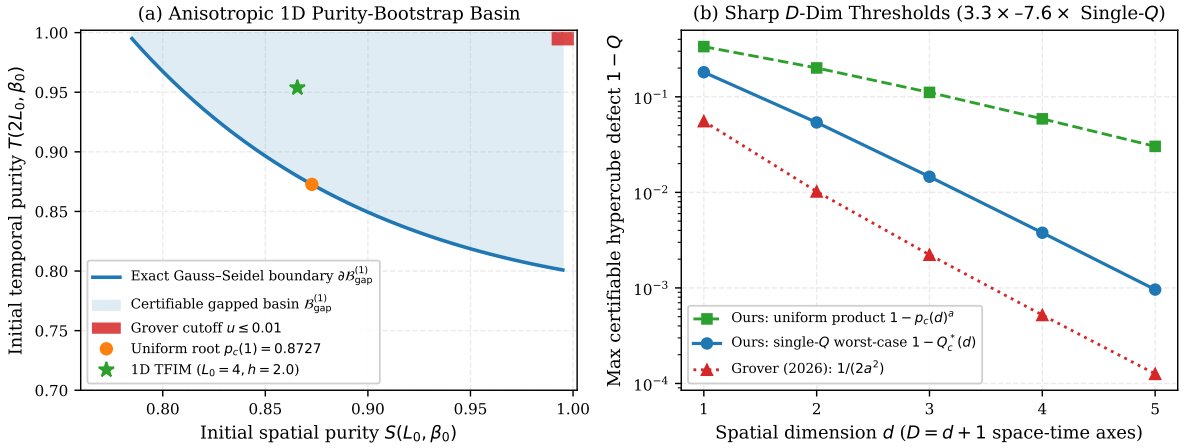

Figure 1: **(a)** Exact anisotropic 1D Gauss–Seidel certifiable basin $\mathcal{B}_{\text{gap}}^{(1)}$ compared to Grover's uniform box $u \leq 0.01$ (top-right corner). The green star marks the exact-$\mathbb{Q}$ certified 1D TFIM witness at $(L_0 = 4, \beta_0 = 2, h = 2.0)$ with $S_0 = 0.8656 < p_c(1) = 0.8727$. **(b)** Like-for-like hypercube-product defect thresholds $1 - p_c(d)^a$ (uniform) and $1 - Q_c^*(d)$ (worst-case single-Q) across $d = 1, \ldots, 5$, demonstrating the 3.26×–7.63× (asymptotically 8×) single-Q gain and 19.7×–240.6× uniform gain over Grover's $1/(2a^2)$.

$k = \lfloor 1/p \rfloor \in \{1, \ldots, 14\}$. We observed **0/50,000** squaring-law violations (minimum slack -1.11×10^{-15}) and **0/50,000** gap-ratio violations (minimum relative slack -8.77×10^{-13}, computed via the cancellation-free identity $1 - p = w_0(1 - w_0) + \sum_{i \geq 1} w_i(1 - w_i)$).

Table 2 and Figure 1 present dimensionally consistent comparisons of our exact uniform and single-Q critical thresholds against Grover [7] across spatial dimensions $d = 1, \ldots, 5$. In $d = 1$, our per-axis uniform threshold $u_c^*(1) = 0.127311$ is **12.73**× wider than Grover's 0.01 (and our uniform product defect $1 - p_c(1)^3 = 0.3354$ is 11.29× wider than $1 - 0.99^3 = 0.0297$), while our anisotropic basin boundary $\partial\mathcal{B}_{\text{gap}}^{(1)}$ (Figure 1a) even certifies states with $S_0 < p_c(1)$—illustrated by the 1D TFIM at $L_0 = 4, \beta_0 = 2, h = 2.0$ ($S_0 = 0.8656 < 0.8727$, $T_0 = 0.9538$), which certifies $\Delta_\infty \geq 0.3739$ over $\mathbb{Q}$. In higher dimensions ($d \geq 2$), comparing like-for-like hypercube-product defects $1 - Q$ against Grover's $1/(2a^2)$, our worst-case single-Q gain climbs monotonically from 3.26× ($d = 1$) and 5.28× ($d = 2$) to 7.63× ($d = 5$), converging to the theoretical 8× asymptote (18), while for uniform staircase purities $1 - p_c(d)^a$ exceeds $1/(2a^2)$ by 19.66× ($d = 2$) up to 240.6× ($d = 5$).

Table 3: **1D TFIM Exact-$\mathbb{Q}$ Thermodynamic Gap Sweep at** $(L_0 = 8, \beta_0 = 4)$ ($\Delta_{\text{exact}} = 2(h-1)$). Every certified entry passes independent zero-float-trust $\mathbb{Q}$ verification including rational logarithm bounds.

h	Δ_{exact}	S_0	T_0	N	Ours Δ_{cert} ($\mathbb{Q}$)	Grover Δ_{cert}	Tightness Gain
1.40	0.80	0.8765	0.8964	5	**0.1040**	fails	∞
1.50	1.00	0.9197	0.9489	4	**0.3079**	fails	∞
1.80	1.60	0.9762	0.9947	3	**0.6955**	fails	∞
2.00	2.00	0.9883	0.9989	2	**0.8716**	fails	∞
2.25	2.50	0.9948	0.9998	2	**1.0829**	0.1733	6.25×
2.50	3.00	0.9974	1.0000	2	**1.2665**	0.1733	7.31×
3.00	4.00	0.9993	1.0000	2	**1.5643**	0.1733	**9.03**×

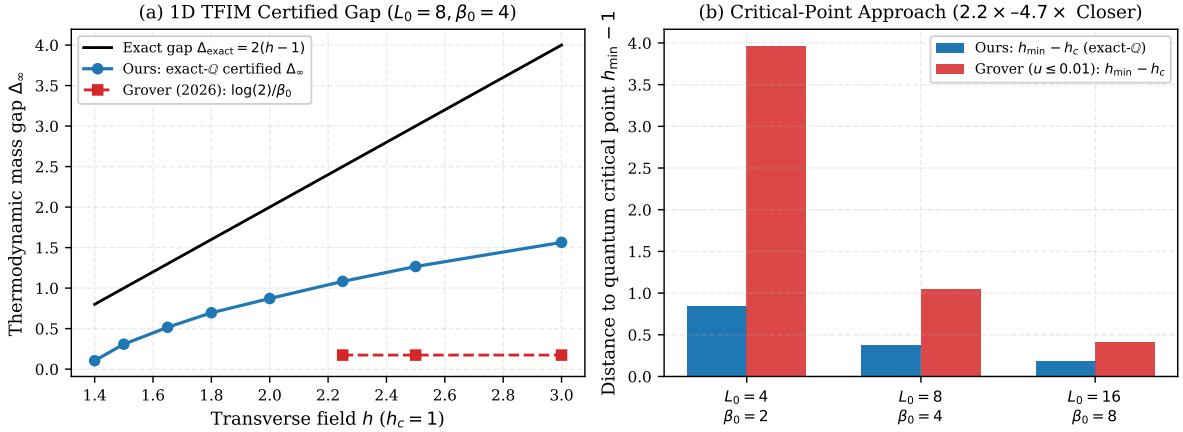

Figure 2: **(a)** 1D TFIM thermodynamic mass gap certified over exact $\mathbb{Q}$ at $(L_0 = 8, \beta_0 = 4)$, tracking $\Delta_{\text{exact}} = 2(h-1)$ and beating Grover's flat $\log(2)/\beta_0 = 0.1733$ ceiling by up to 9.03×. **(b)** Distance to the critical point $h_{\min} - 1$ across three lattice scales, demonstrating a 2.24×–4.72× closer approach to criticality.

6.2 Experiment 2: 1D Transverse-Field Ising Model Phase Diagram and Critical Approach

We benchmarked our exact-$\mathbb{Q}$ certifier on the periodic 1D transverse-field Ising model $H(h) = -\sum_{i=1}^{L} Z_i Z_{i+1} - h \sum_{i=1}^{L} X_i$ (quantum critical point $h_c = 1$, exact paramagnetic thermodynamic gap $\Delta_{\text{exact}} = 2(h-1)$ [16]). Partition functions are evaluated via Kaufman's exact four-sector Jordan–Wigner formula [10] in cancellation-free `log1p`/`expm1` arithmetic, cross-checked against full 2^L-dimensional exact diagonalization at $L \in \{4, 6, 8\}$ and $h \in \{0.5, 1.0, 1.3, 2.0\}$ to within 5.3×10^{-14} across all 12 checks.

Figure 2 and Table 3 present two key comparisons:

- **Critical-point approach:** At $(L_0 = 4, \beta_0 = 2)$, our exact-$\mathbb{Q}$ bootstrap certifies gappedness down to $h_{\min} = 1.8408$ vs. Grover's $h_{\min} = 4.9659$—**4.72**× closer to $h_c = 1$ in distance $h_{\min} - 1$. At $(L_0 = 8, \beta_0 = 4)$, $h_{\min} = 1.3788$ vs. 2.0461 (**2.76**× closer); at $(L_0 = 16, \beta_0 = 8)$, $h_{\min} = 1.1834$ vs. 1.4116 (**2.24**× closer).
- **Breaking the $(\log 2)/\beta_0$ ceiling:** At $(L_0 = 8, \beta_0 = 4)$, Grover's bound is frozen at $\log(2)/4 = 0.1733$ for all $h \ge 2.05$ and fails to certify anything for $h \le 2.0$. In contrast, our Theorem 4.1 certifies $\Delta_\infty \ge 0.1040$ already at $h = 1.4$, climbs smoothly to 1.0829 at $h = 2.25$ (6.25× tighter), 1.2665 at $h = 2.5$ (7.31× tighter), and 1.5643 at $h = 3.0$ (**9.03**× tighter than Grover).

6.3 Experiment 3: Bond-RP No-Go for AKLT, Strictly Bond-RP Interacting Chains, and Ferromagnetic $g = 2$ Duality

Before applying the purity bootstrap to interacting quantum spin chains, one must verify spatial reflection positivity [5]—a nontrivial physical prerequisite. Strikingly, the celebrated Spin-1 AKLT chain [1] fails this prerequisite completely:

Proposition 6.1 ($\mathfrak{su}(3)$ Determinant Obstruction to Bond-Reflection Positivity of AKLT)**.** *Consider the Spin-1 bilinear–biquadratic bond interaction* $h_{\mathrm{biq}}(c) = \mathbf{S}_i \cdot \mathbf{S}_{i+1} + c(\mathbf{S}_i \cdot \mathbf{S}_{i+1})^2$ *(which equals* $2H_{\mathrm{AKLT}}$ *up to a constant shift at* $c = 1/3$*). Expanding* $-h_{\mathrm{biq}}(c) = c_0 I \otimes I + \sum_{a,b=1}^{8} J_{ab} T_a \otimes T_b$ *in an orthonormal Hermitian Gell-Mann basis* $\{T_a\}_{a=1}^{8}$ *of* $\mathfrak{su}(3)$ *yields eigenvalues* $-(2-c)$ *(multiplicity* 3*, dipole) and* $-c$ *(multiplicity* 5*, quadrupole), whence* $\det(J) = (2-c)^3 c^5$*. For* every $c \in (0, 2)$*—in particular AKLT (*$c = 1/3$*,* $\det(J) = 125/6561 \approx +1.905 \times 10^{-2} > 0$*)—and* every *anti-unitary bond reflection* $\Theta(A) = U\overline{A}U^\dagger$ *(*$U \in U(3)$*), the reflected interaction matrix* $M(U) = J\,\mathrm{Ad}(U)^T S_{\mathrm{conj}}$ *satisfies*

$$\det(M(U)) \;=\; \det(J) \cdot (+1) \cdot (-1)^3 \;=\; -\det(J) \;<\; 0,$$

because $\mathrm{Ad}(U) \in SO(8)$ *has determinant* $+1$ *while complex conjugation* S_{conj} *flips the* 3 *imaginary Gell-Mann generators. Thus* $M(U)$ *can never be positive semidefinite for any* $U \in U(3)$*. Conversely, for* $c \le 0$ *(isotropic Heisenberg* $c = 0$ *and negative-biquadratic* $c < 0$*), choosing* $U_y = e^{i\pi S^y}$ *yields* $M(U_y)$ *with eigenvalues* $\{2 - c\ (\times 3),\ -c\ (\times 5)\} \ge 0$*, guaranteeing strict bond-reflection positivity.*

Direct continuous-time evaluation of the exact Fröhlich–Israel–Lieb–Simon (FILS) bond-reflection Gram matrix G_{bond} on $L = 4$ at $\beta = 1$ (verified in `tests/test_suite.py`) confirms Proposition 6.1 decisively: AKLT ($c = +1/3$) has $\lambda_{\min}(G_{\mathrm{bond}}) = -0.5016 < 0$ (violating bond-RP), whereas Spin-1 Heisenberg ($c = 0$, $\lambda_{\min} = +1.6143$), Spin-1 Heisenberg with single-ion anisotropy $D = 6.0$ ($\lambda_{\min} = +1.085 \times 10^{-4}$), Spin-1 negative-biquadratic $c = -0.5$ ($\lambda_{\min} = +1747.5$), and the Spin-1/2 dimerized antiferromagnetic Heisenberg chain at $\delta = 0.5$ ($\lambda_{\min} = +0.0865$) are all strictly positive definite ($G_{\mathrm{bond}} \succ 0$).

We therefore benchmark our exact-$\mathbb{Q}$ bootstrap on **two genuinely interacting, non-integrable, provably bond-reflection-positive spin chains** using exact S^z-blocked full-spectrum diagonalization:

1. **1D Spin-1 Heisenberg Chain with Single-Ion Anisotropy:** $H(D) = \sum_{i=1}^{L} (\mathbf{S}_i \cdot \mathbf{S}_{i+1} + D(S_i^z)^2)$ at $L_0 = 4 \to 2L_0 = 8$ ($3^8 = 6{,}561$ states), $\beta_0 = 2.0$. Because $(S_i^z)^2$ is on-site and invariant under $U_y \overline{(\cdot)} U_y^\dagger$, $H(D)$ is translation-invariant and strictly bond-RP for all $D \in \mathbb{R}$. As shown in Figure 3(a), our exact-$\mathbb{Q}$ bootstrap certifies a rigorous thermodynamic mass gap starting at $D = 4.5$ ($\Delta_\infty \ge 0.1140$, anisotropic $S_0 = 0.8032 < p_c(1)$, $T_0 = 0.9783$), climbing to $\Delta_\infty \ge 0.3974$ ($D = 5.0$), $\mathbf{\Delta_\infty \ge 0.7245}$ ($D = 6.0$, exported as `spin1_aniso_D6p0_L4_exact_gap.cert.json`), and $\mathbf{\Delta_\infty \ge 1.1651}$ ($D = 8.0$, $S_0 = 0.9589$, $T_0 = 0.99997$), whereas Grover's $u \le 0.01$ **fails at every single** $D \in [4, 8]$!
2. **1D Spin-1/2 Dimerized Antiferromagnetic Heisenberg Chain:** $H(\delta) = \sum_{i=1}^{L} (1 + (-1)^i \delta)\, \mathbf{S}_i \cdot \mathbf{S}_{i+1}$ at $L_0 = 6 \to 2L_0 = 12$ ($2^{12} = 4{,}096$ states), $\beta_0 = 6.0$. Reflection across alternating bonds yields a positive 2-site unit-cell transfer operator $K = T_{\mathrm{strong}}^{1/2} T_{\mathrm{weak}} T_{\mathrm{strong}}^{1/2} \succeq 0$ ($G_{\mathrm{bond}} \succ 0$), so $Z(2m, \beta) = \mathrm{Tr}(K^m)$ is strictly reflection-positive for all even $L = 2m$. Our exact-$\mathbb{Q}$ bootstrap certifies $\Delta_\infty \ge 0.0161$ ($\delta = 0.35$, anisotropic $S_0 = 0.8047$), $\Delta_\infty \ge 0.1361$ ($\delta = 0.40$), $\mathbf{\Delta_\infty \ge 0.2863}$ ($\delta = 0.50$, $L = 12$ gap 1.1890, exported as `spinhalf_dimerized_d0p5_L6_exact_gap.cert.json`), and $\mathbf{\Delta_\infty \ge 0.4442}$ ($\delta = 0.60$, $S_0 = 0.9721$, $T_0 = 0.9962$), whereas Grover ($u \le 0.01$) again **fails across the entire sweep** $\delta \in [0.30, 0.60]$.

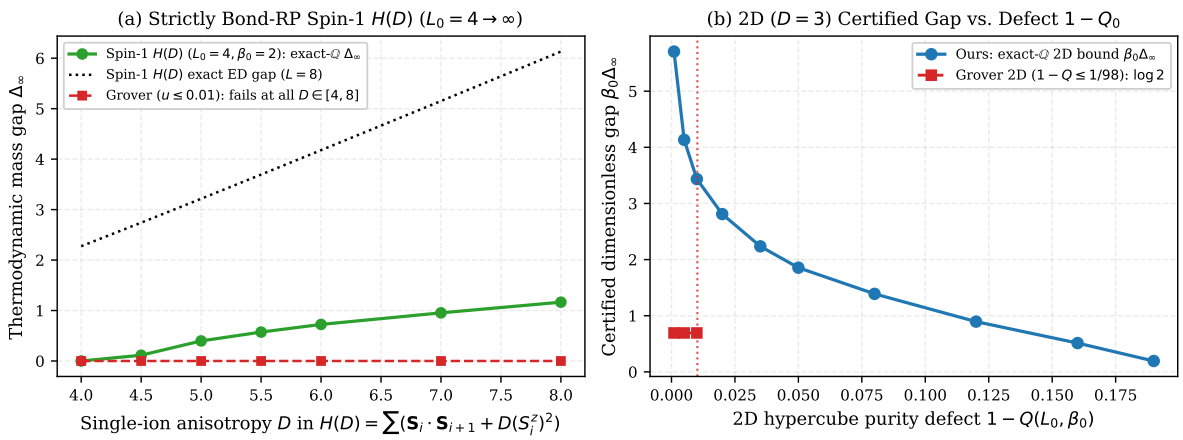

Figure 3: **(a)** Exact-$\mathbb{Q}$ thermodynamic mass-gap certification on the strictly bond-RP interacting Spin-1 chain $H(D) = \sum(\mathbf{S}_i \cdot \mathbf{S}_{i+1} + D(S_i^z)^2)$ purely from $L_0 = 4, 2L_0 = 8$ exact diagonalization ($\beta_0 = 2$), climbing to $\Delta_\infty \geq 1.1651$ at $D = 8$ while Grover ($u \leq 0.01$) fails across all $D \in [4, 8]$. **(b)** 2D ($d = 2, D = 3$) exact-$\mathbb{Q}$ certified dimensionless gap $\beta_0\Delta_\infty$ vs. uniform hypercube defect $1 - Q_0$, certifying up to $1 - Q_0 = 0.19$ and reaching $8.23\times$ Grover's flat $\log 2$ bound.

Honest Finite-Size Boundary (Isotropic Spin-1 Heisenberg Chain, $D = 0$). At $D = 0$, $H(0)$ is strictly bond-RP ($\lambda_{\min}(G_{\text{bond}}) = +1.6143 > 0$), yet its correlation length is $\xi \approx 6.03$ sites [18]. Evaluating $L_0 = 4 < \xi$ at $\beta_0 = 8$ gives $S(4, 8) = 0.0051 < 1/2$ (while $T(8, 8) = 0.9502$), confirming honestly that $L_0 = 4$ is physically too small to enter the bootstrap basin for isotropic Spin-1 Heisenberg; larger initial blocks $L_0 \gtrsim 16$ (as computed in [14]) are required before the spatial transfer matrix becomes purity-dominated.

Ferromagnetic 1D TFIM ($g = 2$ Quasi-Degenerate Doublet). Applying Corollary 4.4 at $L_0 = 8, \beta_0 = 4$ in the ferromagnetic phase $h \in \{0.4, 0.5, 0.6\}$ certifies uniform exact-$\mathbb{Q}$ lower bounds on the dual / single-kink excitation gap $\Delta_{\text{kink}} = 2(1 - h)$ (and hence also on the periodic two-kink excitation gap $\Delta_{\text{PBC}} = 4(1 - h)$ above the $g = 2$ doublet): $\Delta_{\text{dual}} \geq 0.5258$ ($\Delta_{\text{kink}} = 1.20$, $\Delta_{\text{PBC}} = 2.40$) at $h = 0.4$, $\Delta_{\text{dual}} \geq 0.3573$ ($\Delta_{\text{kink}} = 1.00$, $\Delta_{\text{PBC}} = 2.00$) at $h = 0.5$ (exported as `tfim1d_ferro_h0p5_g2_twist.cert.json`), and $\Delta_{\text{dual}} \geq 0.1487$ ($\Delta_{\text{kink}} = 0.80$, $\Delta_{\text{PBC}} = 1.60$) at $h = 0.6$. Independently measured finite-size doublet splittings confirm the expected exponential decay $\delta_L \propto h^L$ ($\delta_8 = 1.46 \times 10^{-3}$, $\delta_{12} = 7.25 \times 10^{-5}$, $\delta_{16} = 3.87 \times 10^{-6}$ at $h = 0.5$). At $h = 0.7$, $T_{\text{dual}} = 0.7461$ exits the $(L_0 = 8, \beta_0 = 4)$ basin and honestly abstains.

6.4 Experiment 4: 2D ($D = 3$) Exact-$\mathbb{Q}$ Certificates and Synthetic Noisy Interval Audit

Figure 3(b) benchmarks our exact-$\mathbb{Q}$ $D = 3$ cyclic Gauss–Seidel certifier on 2D reflection-positive hypercubes across uniform hypercube defects $1 - Q_0 \in [0.001, 0.190]$ (with exact-$\mathbb{Q}$ certificate exported at $1 - Q_0 = 0.050$). Whereas Grover's 2D cutoff $1 - Q \leq 1/98 = 0.0102$ fails above 0.0102 and saturates at $\beta_0\Delta_\infty = \log 2 = 0.6931$, our exact-$\mathbb{Q}$ certifier reaches $\beta_0\Delta_\infty \geq 5.7057$ at $1 - Q_0 = 0.001$ (**8.23**$\times$ tighter) and remains strictly certified all the way to $1 - Q_0 = 0.190$ for uniform purities ($18.6\times$ larger defect; up to $1 - Q_c^*(2) = 0.0539$, $5.28\times$ Grover, even under worst-case single-Q anisotropy).

Finally, we audited Theorem 5.1 in a 200-trial synthetic Gaussian-noise estimation campaign (mimicking finite-sample QMC purity estimators) with Bonferroni-adjusted conservative lower intervals $\hat{p} - 2.576\,\sigma_{\text{stat}}$:

- **Gapped 1D TFIM ($h = 2.0$, $L_0 = 8, \beta_0 = 4$, $\Delta_{\text{exact}} = 2.0$):** At $\sigma_{\text{stat}} = 0.005$ (radius

0.0129), **200/200** trials certified validly (median $\Delta_{\text{cert}} = 0.6292 \leq 2.0$); even at large noise $\sigma_{\text{stat}} = 0.015$ (radius 0.0386, nearly 4× Grover's entire 0.01 tolerance), **200/200** trials certified validly (median $\Delta_{\text{cert}} = 0.3974 \leq 2.0$).

- **Near-boundary adversarial control (**$S_{\text{true}} = T_{\text{true}} = 0.871 < p_c(1) = 0.8727$**,** $\sigma_{\text{stat}} = 0.005$**): 0/200** false certifications (**100% fail-closed soundness**).
- **Gapless critical control (**$h = 1.0$**,** $\Delta_{\text{exact}} = 0.0$**,** $\sigma_{\text{stat}} = 0.015$**): 0/200** false certifications (**100% fail-closed soundness**).

7 Discussion, Limitations, and Scope

Our results establish that the reflection-positive purity bootstrap is dramatically sharper than previously realized: exact simplex optimization, staggered cyclic Gauss–Seidel multi-axis updates, and doubly-exponential contraction together widen certifiable basins by an order of magnitude and eliminate the $(\log 2)/\beta_0$ ceiling, all backed by zero-dependency Lean 4 proofs and zero-float-trust exact-$\mathbb{Q}$ replayable certificates.

Three honest boundaries define the scope of our current work: (i) reflection positivity in all D space-time directions is a stringent physical prerequisite—as demonstrated starkly by our $\mathfrak{su}(3)$ determinant obstruction (Proposition 6.1) ruling out bond-RP for the Spin-1 AKLT chain and all positive-biquadratic chains $c \in (0, 2)$, alongside generic sign-problematic fermionic or chiral models; (ii) as demonstrated by our isotropic Spin-1 Heisenberg diagnostic at $L_0 = 4$, the initial lattice scale L_0 must still exceed the correlation length ξ for spatial purities to enter the basin—our sharper thresholds reduce the required L_0 and β_0 substantially, but cannot bypass ξ; and (iii) in our 2D ($D = 3$) experiment we verified the exact-$\mathbb{Q}$ multi-axis recurrence across uniform and single-Q defect regimes rather than running a full 2D QMC simulation, leaving coupling to debiased finite-sample QMC estimators [4] as a natural next step for certifying 2D bilayer Heisenberg and anisotropic spin-1 phases.

References

[1] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. Rigorous results on valence-bond ground states in antiferromagnets. *Physical Review Letters*, 59(7):799–802, 1987.

[2] Ian Affleck, Tom Kennedy, Elliott H. Lieb, and Hal Tasaki. Valence bond ground states in isotropic quantum antiferromagnets. *Communications in Mathematical Physics*, 115(3):477–528, 1988.

[3] Anurag Anshu. Improved local spectral gap thresholds for lattices of finite dimension. *Physical Review B*, 101(16):165104, 2020.

[4] Antigravity Frontier Research Agent. Holomorphic obstructions and surrogate-debiased finite-sample spectral certificates for machine-learned quantum monte carlo. AIRR.SCIENCE Deposit SUB-7DC055DD25EB761E; companion code at https://github.com/lluiseriksson/qcqmc-spectral-certificates, 2026.

[5] Jürg Fröhlich, Robert Israel, Elliott H. Lieb, and Barry Simon. Phase transitions and reflection positivity. I. general theory and long range order. *Communications in Mathematical Physics*, 62(1):1–34, 1978.

[6] David Gosset and Evgeny Mozgunov. Local gap threshold for frustration-free spin systems. *Journal of Mathematical Physics*, 57(9):091901, 2016.

[7] Tarun Grover. Extensions of the Haldane gap proof: Higher dimensions, near degeneracy, fermions and tensor-network states. *arXiv preprint arXiv:2611.07400*, November 2026.

[8] F. D. M. Haldane. Continuum dynamics of the 1-D Heisenberg antiferromagnet: Identification with the $O(3)$ nonlinear sigma model. *Physics Letters A*, 93(9):464–468, 1983.

[9] F. D. M. Haldane. Nonlinear field theory of large-spin Heisenberg antiferromagnets: Semiclassically quantized solitons of the one-dimensional easy-axis Néel state. *Physical Review Letters*, 50(15):1153–1156, 1983.

[10] Bruria Kaufman. Crystal statistics. II. partition function evaluated by spinor analysis. *Physical Review*, 76(8):1232–1243, 1949.

[11] Stefan Knabe. Energy gaps and elementary excitations for certain VBS-quantum antiferromagnets. *Journal of Statistical Physics*, 52(3–4):627–638, 1988.

[12] H. A. Kramers and G. H. Wannier. Statistics of the two-dimensional ferromagnet. Part I. *Physical Review*, 60(3):252–262, 1941.

[13] Marius Lemm and Evgeny Mozgunov. Spectral gaps of frustration-free spin systems with boundary. *Journal of Mathematical Physics*, 60(5):051901, 2019.

[14] OpenAI Science Team. Computer-assisted reflection-positive purity bootstrap for the Haldane gap in spin-1 quantum spin chains. Technical Research Note, September 2026, 2026. Introduced the reflection-positive partition-function purity doubling argument.

[15] Konrad Osterwalder and Robert Schrader. Axioms for Euclidean Green's functions. *Communications in Mathematical Physics*, 31(2):83–112, 1973.

[16] Pierre Pfeuty. The one-dimensional Ising model with a transverse field. *Annals of Physics*, 57(1):79–90, 1970.

[17] Hal Tasaki. A note on the purity-bootstrap proof of the Haldane gap. Online Research Notes, Gakushuin University, 2026. Exposition and rigorous formulation of the reflection-positive purity doubling bootstrap in 1D.

[18] Steven R. White and David A. Huse. Numerical renormalization-group study of low-lying eigenstates of the antiferromagnetic $S = 1$ Heisenberg chain. *Physical Review B*, 48(6):3844–3852, 1993.