Supplemental Numerical and Reproducibility Information for "Few-Outcome Readout of Quantum-Fisher-Optimal Measurements: Sharp Bounds and Coherent-Synthesis Hardness"

Maestro v0.1
Pudim AI

August 20, 2026

Abstract

Proofs of results first established in the main article are integrated into its Appendices; the general benign projective-landscape theorems imported from the companion reference are not reproduced. This document records the numerical ensembles, deterministic generation rules, physical benchmark, heuristic product-search protocol, validation thresholds, and file-level reproducibility map. The archived calculations use dense exact linear algebra to test mathematical identities and finite-dimensional approximation guarantees. They do not emulate the logical circuit or establish the compiler's operational score-intertwining contract.

1 Reproducibility contract

The package is designed to rebuild without network access after the pinned Python and LaTeX dependencies have been installed. The master seed is

$$s_0 = 20260818. \tag{1}$$

Every stochastic experiment derives a separate deterministic generator from this value. The command

```
python reproduce.py --mode complete --run-json run.json
```

regenerates the numerical data, both main-text and supplemental figures, validation report, article, supplemental document, and combined PDF in a staging directory, validates them, promotes them only on success, and records hashes and timings in `run.json`. A faster document-and-validation check is

```
python reproduce.py --mode smoke --run-json run-smoke.json
```

which recompiles the archived data and figures without rerunning the optimizations. The container contract is declared separately; the candidate attestation identifies whether a container and offline isolation were actually observed.

The numerical code uses NumPy and SciPy complex double precision. Matrix products are explicitly Hermitized whenever roundoff could otherwise spoil a formal Hermiticity assumption. Eigenproblems use dense Hermitian solvers; Sylvester equations use a Bartels–Stewart implementation. All quantities displayed in the article can be reconstructed from the CSV, NPY, NPZ, and JSON files under `data/`.

2 Random-state ensembles

2.1 Faithful states and admissible tangents

For dimension d and requested condition number κ, a random faithful state is generated as

$$\rho = U \frac{\operatorname{diag}(w_1, \ldots, w_d)}{\sum_i w_i} U^\dagger, \qquad w_i \in [1/\kappa, 1], \tag{2}$$

where U is Haar distributed and the positive weights are sampled before normalization. A random traceless Hermitian matrix H is transformed to

$$\dot\rho = \tau\, \rho^{1/2} H \rho^{1/2}, \qquad \operatorname{Tr} \dot\rho = 0, \tag{3}$$

with τ chosen so that $\rho \pm \dot\rho$ remain positive with a fixed safety margin. The SLD is then obtained by solving

$$\rho L + L\rho = 2\dot\rho. \tag{4}$$

This construction prevents the numerical tests from being dominated by boundary singularities, which lie outside the faithful-state theorem.

2.2 Gap-identity ensemble

The identity test uses $d = 16$, $\kappa = 8$, and 600 PVMs. Four hundred eighty bases are geodesic perturbations of the SLD eigenbasis with perturbation amplitudes logarithmically spaced from 2×10^{-5} to 1.2; the remaining bases are Haar random. This mix resolves the identity over approximately eight decades of normalized suboptimality while retaining far-from-optimal measurements.

For each basis, the archived quantities are

$$x_P = \frac{F_Q - J(P)}{F_Q}, \qquad y_P = \frac{\mathrm{Tr}[\rho(L_P - L)^2]}{F_Q}, \qquad e_P = |x_P - y_P| F_Q. \tag{5}$$

The largest absolute discrepancy is 8.59×10^{-18}. Twenty-four records close to exact equality round to zero error in IEEE double precision. Figure S1(a) shows the nonzero error distribution.

2.3 Score-ascent ensemble

For every $\kappa \in \{2, 4, 8, 16\}$, twenty independent $d = 32$ instances are initialized at $X_0 = 0$ and iterated according to

$$X_{t+1} = X_t + \frac{2\dot\rho - \rho X_t - X_t\rho}{\lambda_{\min}(\rho) + \lambda_{\max}(\rho)}. \tag{6}$$

The stored observable is the normalized potential gap

$$g_t = \frac{\mathrm{Tr}[\rho(X_t - L)^2]}{\mathrm{Tr}(\rho L^2)}. \tag{7}$$

The dashed curves in Fig. S1(c) are the worst-case PL envelopes q^{2t} with $q = (\kappa - 1)/(\kappa + 1)$. Solid curves are medians over the twenty instances. For $\kappa = 16$, the median gap after ninety iterations is 7.815×10^{-13}.

2.4 Finite-resolution ensemble

For forty independent $d = 32$, $\kappa = 8$ pairs, the SLD spectrum is divided into fixed-width intervals with requested widths

$$\delta / \mathrm{diam}(L) \in [1/128, 1/4]. \tag{8}$$

Because an interval may contain no eigenvalues, the code records both the requested width δ and the largest *actual* occupied-bin diameter $\Delta_B \le \delta$. Every coarse projector is built by summing the SLD eigenprojectors in its bin and its Fisher information is evaluated directly from Born probabilities. Across all 520 trials,

$$\max \frac{F_Q - J_\delta}{\delta^2/4} = 0.366541 < 1. \tag{9}$$

The largest value of $F_Q - J_\delta - \Delta_B^2/4$, after the 5×10^{-13} floating-point allowance used by the validator, is negative. Figure S1(b) shows the complete tight-resolution sweep.

2.5 Optimal score-compression ensemble

For each $d \in \{8, 16, 32, 64\}$, twenty independently generated faithful models are diagonalized and their state-weighted SLD spectra are optimally partitioned by the $O(Kd^2)$ Bellman dynamic program in the article. The archived table records both the exact Lloyd–Max optimum and the sufficient bin count obtained by evaluating Theorem V.3 on each measured spectral tail. At 99% retention, the exact medians are $7, 10, 12, 14$ and the theorem medians are $25.5, 28, 30, 31.5$. Over the four tested dimensions,

$$K_{99}^\star = 7.30 + 2.30 \log_2(d/8), \qquad R^2 = 0.989, \tag{10}$$

$$K_{99}^{\mathrm{tail}} = 25.75 + 2.00 \log_2(d/8), \qquad R^2 = 0.988. \tag{11}$$

These are finite-range diagnostics, not asymptotic laws.

The separate constructive tail gate uses the four-spin sensor at $\beta = 0.30$, radii $R \in \{0.5, 1, 1.5, 2\}$, and central-bin counts $K \in \{4, 8\}$. For each pair the code constructs the K equal central intervals plus two tail bins exactly as in the proof. The maximum observed loss-to-bound ratio over the 8 cases is 0.323037. Raw tail moments, partitions, and losses are archived in `data/tail_bound.csv`.

2.6 Dense-spectrum high-rate and atomic-crossover ensemble

To test the dense-spectrum corollary in a setting where its hypotheses are nonvacuous, the code constructs commuting models

$$\rho_m = I/m, \qquad L_m = \operatorname{diag}\left(-1 + \frac{2j-1}{m}\right)_{j=1}^{m}, \qquad \dot\rho_m = L_m/m. \tag{12}$$

The score measure is the uniform midpoint rule for the limiting density $f = 1/2$ on $[-1, 1]$, and $W_2(\nu_m, f\,\mathrm{d}\lambda) = 1/(\sqrt{3}m)$. When K divides m, exact contiguous quantization gives

$$D_K(\nu_m) = \frac{1}{3K^2} - \frac{1}{3m^2}, \qquad D_m(\nu_m) = 0. \tag{13}$$

The limiting PDZ constant is 1/3. Across $m \in \{64, 128, 256\}$, the largest formula error is 0, the largest relative deviation of $K^2 D_K$ from 1/3 over the archived intermediate points is 0.015625, and the terminal loss is 0. Figure S4(a) shows the plateau and exact crossover. Raw records are archived in `data/thermodynamic_quantization.csv`.

3 Explicit strict-saddle instance

The three-dimensional landscape uses

$$\rho = \frac{I_3}{3}, \qquad \dot\rho = \begin{pmatrix} 0.10 & 0.20 & 0 \\ 0.20 & 0.10 & 0 \\ 0 & 0 & -0.20 \end{pmatrix}. \tag{14}$$

At the computational basis, the first two outcomes have the same score but a nonzero off-diagonal score-block residual. Rotating basis vectors 1 and 2 gives the analytically predicted positive curvature

$$\left.\frac{\mathrm{d}^2 J}{\mathrm{d}x^2}\right|_0 = 1.920000, \tag{15}$$

while the finite-difference value is 1.920000. An inter-block rotation has curvature -1.080001. The two signs establish that the point is a genuine saddle rather than a minimum or a flat nonglobal stationary point. The landscape grid archived in `saddle_landscape.npz` is sampled at 181×151 points.

4 Frustrated thermal sensor

4.1 Hamiltonian and parameter derivative

The physical benchmark is a four-spin open chain with

$$\begin{aligned} H_0 = {} & \sum_{i=1}^{3} (0.90 X_i X_{i+1} + 0.60 Y_i Y_{i+1} + 1.10 Z_i Z_{i+1}) \\ & + 0.35 \sum_{i=1}^{2} Z_i Z_{i+2} + 0.30 \sum_{i=1}^{4} (-1)^{i-1} X_i \\ & + 0.12 \sum_{i=1}^{4} \cos(i\sqrt{2}) Z_i, \end{aligned} \tag{16}$$

whose estimated perturbation is

$$V = \sum_{i=1}^{3} \left(X_i Z_{i+1} + 0.70 Z_i X_{i+1}\right). \tag{17}$$

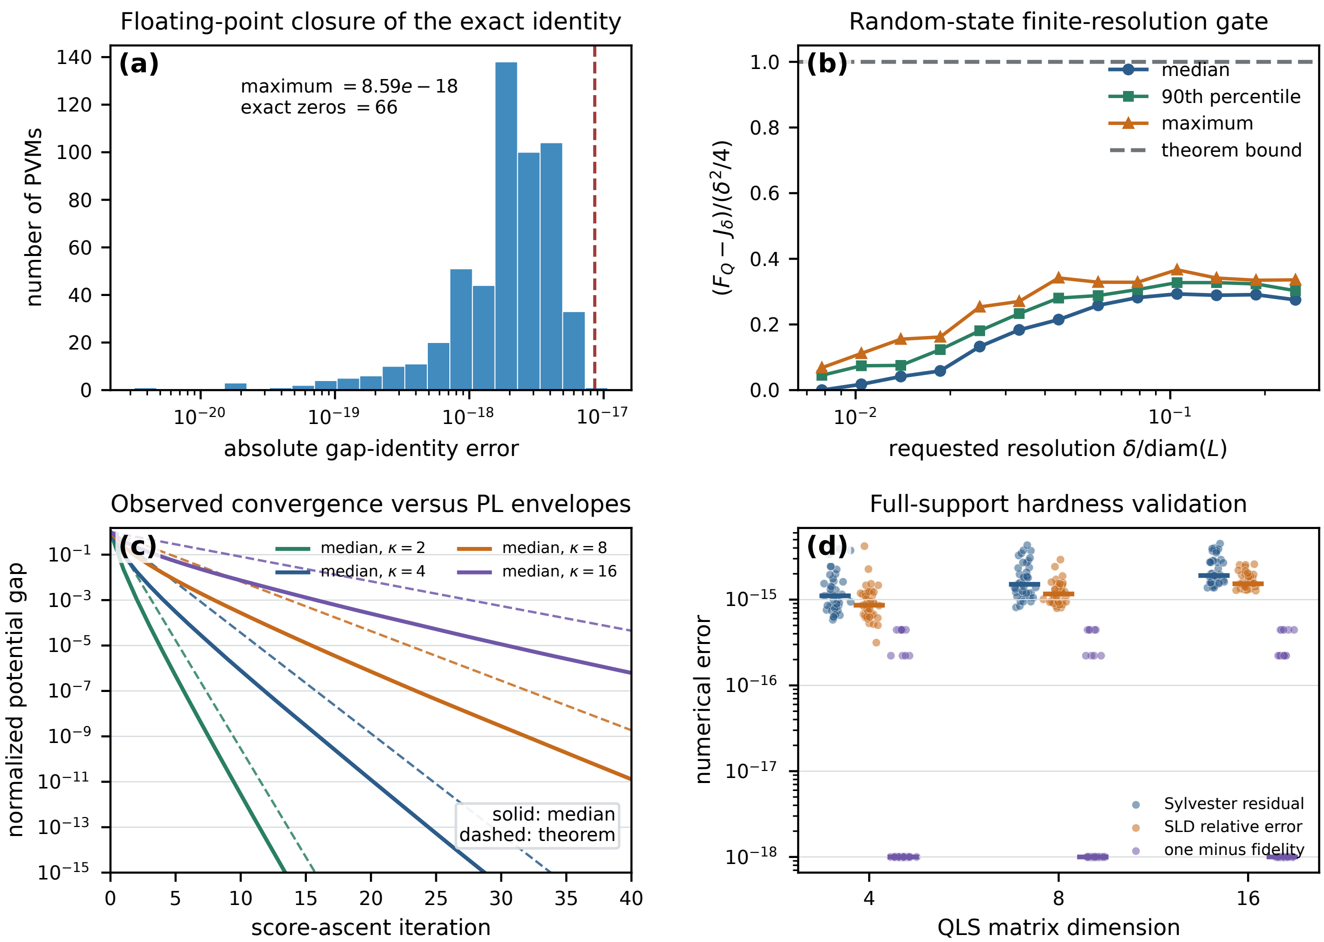

Figure S1: Independent numerical gates. (a) Floating-point residuals for the exact global gap identity. Exact zeros are omitted from the logarithmic histogram and counted in the annotation. (b) Median, 90th percentile, and maximum Fisher loss normalized by the sharp budget $\delta^2/4$; the dashed line is the theorem bound. (c) Median Richardson gaps (solid) and exact contraction envelopes (dashed). (d) Numerical Sylvester residual, SLD error, and QLS-state infidelity for all 123 full-support hardness-family instances; short horizontal segments mark medians.

The state family is

$$\rho_\lambda(\beta) = \frac{e^{-\beta(H_0+\lambda V)}}{\operatorname{Tr} e^{-\beta(H_0+\lambda V)}} \tag{18}$$

and the derivative is evaluated at $\lambda = 0$ by the exact divided-difference formula in the eigenbasis of H_0. The scan contains twelve equally spaced inverse temperatures from 0.18 to 1.50.

4.2 Product-measurement baseline

A rank-one local projective basis on qubit i is the eigenbasis of

$$\boldsymbol{n}_i\cdot\boldsymbol{\sigma}, \qquad \boldsymbol{n}_i = (\sin\theta_i\cos\phi_i, \sin\theta_i\sin\phi_i, \cos\theta_i). \tag{19}$$

The product PVM therefore has eight angular parameters. For each temperature, 128 scrambled Sobol points screen the domain, the five best starts are refined by bounded L-BFGS-B, and the best solution is polished with Powell's method using periodic angle wrapping. The scan used 29438 objective evaluations in total. This is a deterministic, strong heuristic baseline; no claim is made that it proves the global optimum within the product-PVM class. At $\beta = 0.30$, the best value found retains only 48.55% of F_Q.

4.3 Coarse SLD readout

At the same operating point, the exact SLD has sixteen eigenvalues. The exact optimal seven-bin partition retains 98.34% of F_Q, while the exact optimal eight-bin partition retains 99.26%. Figure S2(a) colors equal-bin eigenvalues

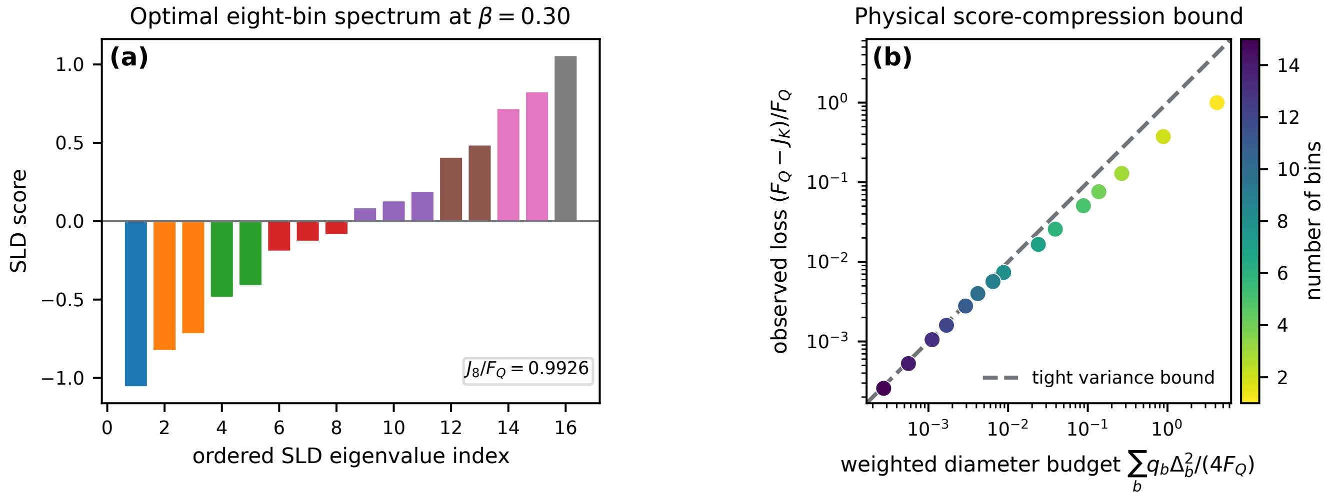

Figure S2: Score compression for the four-spin sensor at $\beta = 0.30$. (a) Ordered SLD eigenvalues; common colors indicate the eight intervals of the exact optimal eight-bin readout. (b) Observed Fisher loss versus the sharp occupied-diameter budget $\Delta_B^2/4$. Color records the number of occupied bins. The dashed identity is the sharp rigorous boundary.

together. Panel (b) tests the physical-model version of the conditional-variance identity using the actual occupied-bin diameter rather than the requested phase-estimation width. Every point lies below the identity line.

4.4 Gibbs access scaling and the inhabited $\beta = c/n$ window

For sizes $n = 2, 4, 6, 8, 10$, the code constructs the same local Hamiltonian pattern and records its full spectral span ΔE_n. Since every positive rescaling of the Gibbs state has eigenvalue ratio $e^{\beta \Delta E_n}$, the condition number is computed both by diagonalization and by that closed form. The maximum relative discrepancy is 1.36×10^{-16}. At $n = 10$, $\Delta E_n/n = 2.866$, giving $\kappa = 5416.7$ at fixed $\beta = 0.30$ and 4.66×10^{18} at fixed $\beta = 1.50$.

The same diagonalizations evaluate $\beta = c/n$ for $c = 1, 3$. At $c = 1$, the largest tested condition number is 17.56, while the exact $n = 10$ QFI is 0.1202 and $nF_Q = 1.2023$. Over $n = 4, 6, 8, 10$, nF_Q lies in $[0.809, 1.202]$. The exact infinite-temperature variance $2^{-n} \operatorname{Tr} V_n^2 = 1.49(n-1)$ supplies the leading predictor $F_Q \simeq 1.49c^2(n-1)/n^2$. Figure S4(b) compares exact conditioning and sensitivity with this predictor. The archived column `beta_for_kappa_n2` remains the largest inverse temperature consistent with the illustrative polynomial target $\kappa = n^2$.

5 Partial readout-resource worksheet

At $\beta = 0.30$, the code shifts H_0 by its minimum eigenvalue and constructs the balanced unnormalized operators

$$S = e^{-\beta(H_0 - E_{\min} I)}, \qquad D = \partial_\lambda S - S\,\partial_\lambda \log Z, \tag{20}$$

followed by the exact SLD L. Every operator is expanded in the four-qubit Pauli basis and coefficients below 10^{-12} are discarded. The resulting nonzero term counts are 136, 135, and 136.

The worksheet uses a deliberately transparent, conservative convention. A unary M-term LCU query is charged $4(M-1)$ synthesized rotations. Target block error 10^{-6} is divided evenly among those rotations and the single-qubit synthesis cost is estimated by $3\log_2(1/\epsilon_{\mathrm{rot}}) + 10$ T gates. This gives 52,920 T gates per L-LCU query. The exact optimal eight-bin boundaries require 123 controlled score queries, computed as $\lceil 2\pi\alpha_L/\delta \rceil$ from the physical score normalization and target width. The 7-bit register merely has enough capacity because $2^7 = 128 > 123$; the query count is not the textbook $2^7 - 1$ schedule. This gives 6,509,160 T gates for the readout stage under this convention. The binary-index and unary logical footprints are 22 and 150 qubits, respectively.

The Sylvester entry reports only the dimensionless proxy $\sqrt{\kappa(S)}\log_2[\kappa(S)2^4/10^{-3}] = 82.4$; the hidden constant in the cited theorem is left unresolved. This worksheet excludes state preparation, complete positive-square-root and Sylvester constants, a composed allocation of the matrix and circuit precision budgets, routing, surface-code cycles, and magic-state factories. It is a readout-side diagnostic, not an end-to-end implementation estimate. The exact coefficients and all assumptions are archived in `data/resource_estimate.json` and `data/resource_operators.csv`.

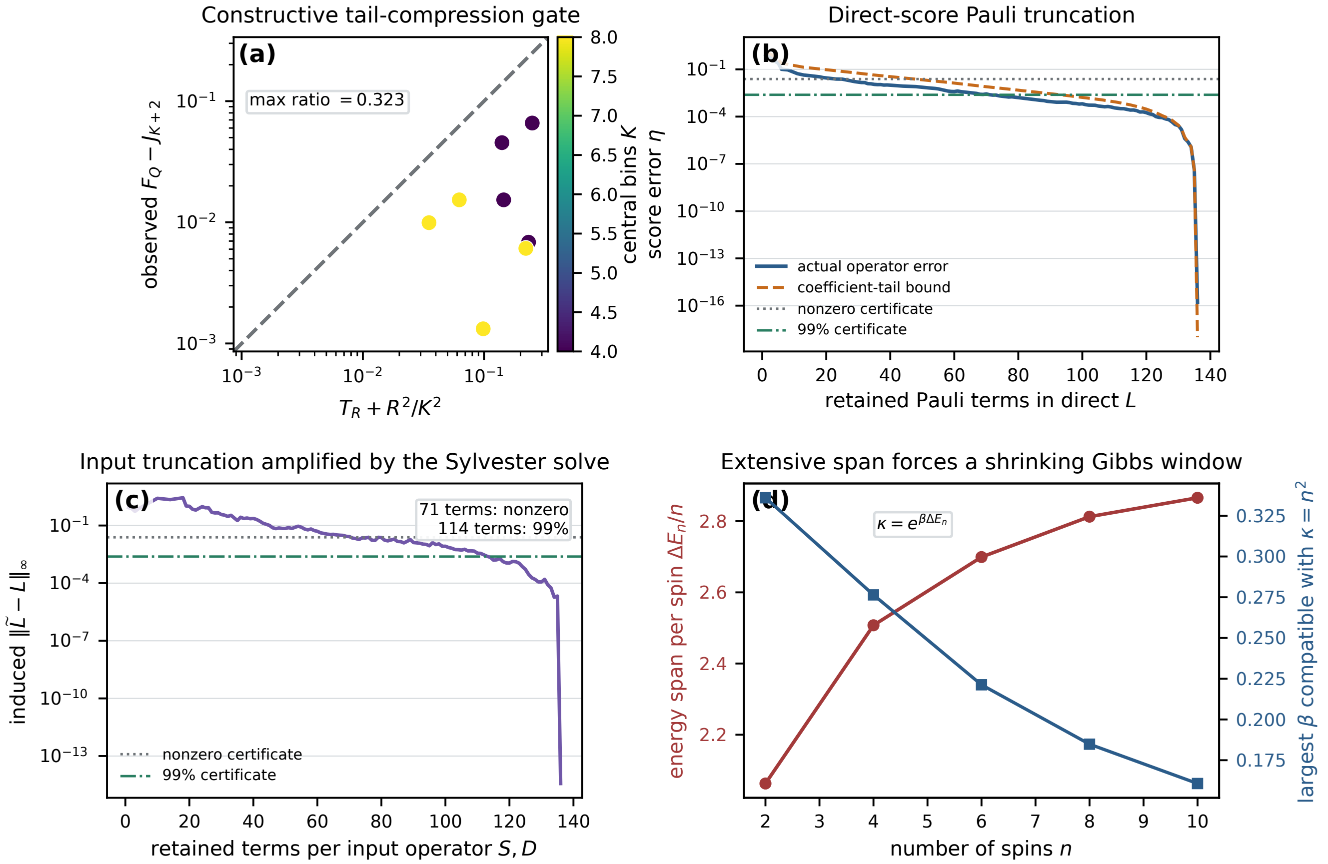

Figure S3: Tail compression, access scaling, and truncation diagnostics. (a) Eight constructive tests of $F_Q - J_{K+2} \le T_R + R^2/K^2$; color gives the number of central bins. (b) Direct truncation of the exact Pauli expansion of L: actual operator error and the coefficient-tail sufficient bound. (c) SLD error induced by truncating both physical input expansions S, D and resolving the Sylvester equation. Horizontal lines in (b,c) are the nonzero- and 99%-information thresholds from the robust compiler theorem. (d) Measured energy span per spin and the inverse temperature compatible with the illustrative polynomial target $\kappa(S) = n^2$.

The robust rank-one certificate is nonzero only when $\eta < \sqrt{F_Q}/(\kappa+1) = 0.02425$. Direct operator-norm truncation of the exact L reaches the nonzero and 99% thresholds after 23 and 72 retained strings. The coefficient-tail sufficient certificate is more conservative, requiring 48 and 94. Truncating the physical inputs S, D and resolving the perturbed Sylvester equation requires 71 and 114 terms per input. The last diagnostic is the relevant one for the compiler input model; it shows directly that the κ-amplified score accuracy, not the nominal Pauli count alone, controls certification.

6 Promise-compatible sensor and full-support hardness ensembles

The parity-coherence sensor is evaluated analytically for $n \in \{2, 4, 8, 16, 32, 64, 128\}$ at $\epsilon = 1/2$. The archived rows verify $\kappa = 3$, $F_Q = n/4$, $n+1$ distinct exact scores, and $\lceil \log_2(n+1) \rceil$ score bits. This family supplies a scalable sensitivity and readout benchmark; it is not used for hardness.

Positive linear-system matrices are generated in dimensions $d \in \{4, 8, 16\}$ and condition numbers $\kappa \in \{2, 4, 8, 16, 32\}$. Each spectrum is logarithmically spaced between 1 and $1/\kappa$ and conjugated by a Haar unitary. A normalized complex Gaussian vector defines $|b\rangle$, and a deterministic unitary completion gives $U_b|0\rangle = |b\rangle$. Eight repetitions at each (d, κ) produce 123 instances.

For each instance the code constructs the full-support block matrices

$$S = \mathrm{diag}(A - aI, aI), \qquad T = \frac{a}{2} \begin{pmatrix} 0 & U_b \\ U_b^\dagger & 0 \end{pmatrix}, \tag{21}$$

solves $SL + LS = 2T$, and compares with the closed-form SLD. It applies L to $|2, 0\rangle$ and compares the normalized

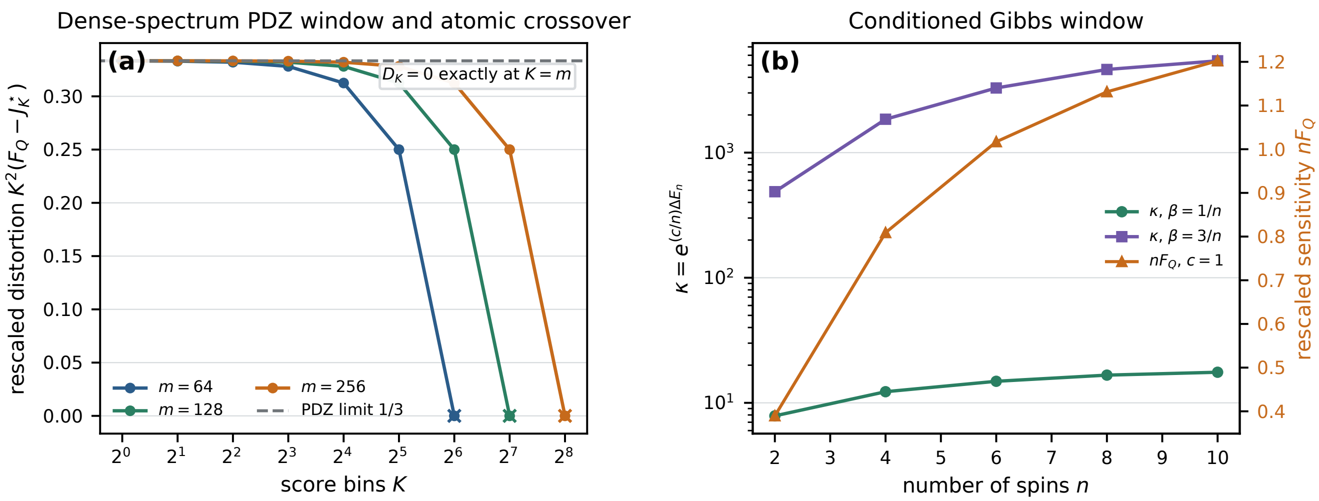

Figure S4: Dense-spectrum and shrinking-temperature gates. (a) Rescaled optimal distortion $K^2(F_Q - J_K^\star)$ for uniform atomic score measures with $m = 64, 128, 256$. The dashed line is the limiting PDZ constant $1/3$; crosses mark the exact zero-loss crossover at $K = m$. (b) Exact Gibbs condition numbers for $\beta = 1/n$ and $3/n$ together with exact nF_Q for $c = 1$. The condition number remains bounded over the tested sequence while the leading high-temperature predictor approaches an $O(1)$ rescaled sensitivity.

output with $A^{-1}|b\rangle$. The archived gates also verify the QFI formula, full nonzero SLD support, inverse-polynomial application success, zero-score margin, and positivity margin.

7 File map and independent audit path

An independent audit can avoid the manuscript entirely: run `reproduce.py`, which regenerates the scientific data before applying the declared tolerances and records the resulting hashes and timings. The quantum circuit is not numerically simulated. Its hardness statement uses the analytic reduction and the explicit score-intertwining interface, while primitive-level implementation cost remains conditional. The four-spin example tests measurement structure and does not establish asymptotic quantum speedup.

Table S1: Archived numerical ensembles. "Instances" counts independent states or reductions; "records" counts rows in the corresponding CSV file.

Gate	Instances	Records	Purpose
Three-level saddle	1	181×151 grid	Resolve one positive and one negative Hessian direction at a nonglobal critical PVM.
Exact gap identity	600	600	Compare $F_Q - J(P)$ against $\mathrm{Tr}[\rho(L_P - L)^2]$ over near-optimal and Haar-random PVMs.
Score ascent	4×20	4×91 quantile rows	Test global linear convergence for $\kappa \in \{2, 4, 8, 16\}$.
Finite resolution	40	40×13	Test the sharp $\delta^2/4$ score-diameter theorem.
Score compression	4×20	$4 \times 20 \times 4$	Find exact optimal contiguous partitions for four target QFI fractions through $d = 64$.
Tail theorem	1 sensor	8	Construct the proof partition for four radii and two central-bin counts.
Dense-spectrum limit	$m = 64, 128, 256$	24	Test the PDZ intermediate window, exact finite-m formula, and zero-loss crossover.
Thermal sensor	12 temperatures	$12 + 16$	Compare SLD and product PVMs; compute exact optimal score partitions.
Resource worksheet	1 operating point	3 operators	Archive exact Pauli-LCU decompositions and stated logical-resource assumptions.
Gibbs conditioning	5 sizes	5	Verify rescaling-invariant $\kappa = e^{\beta \Delta E_n}$, fixed-temperature cost, and the $\beta = c/n$ window with exact QFI.
Pauli truncation	1 operating point	$136 + 136$	Separate direct-score, coefficient-tail, and input-induced SLD errors.
Parity-coherence sensor	7 sizes	7	Verify constant conditioning, extensive QFI, and the exact $n+1$-score alphabet through $n = 128$.
Hardness family	123	123	Verify the full-support faithful embedding, QFI formula, application success, and QLS-state recovery.

Table S2: Twelve validation families comprising 34 executable checks. Every check is recomputed by `validate.py`, which exits nonzero on failure.

Validation family	Archived value	Required condition
G1: Exact landscape geometry	gap error 8.59×10^{-18}; curvatures $1.920000, -1.080001$	identity $< 10^{-14}$; ascent positive; transverse negative
G2: Richardson diagnostic	$\kappa = 16$ final median 7.815×10^{-13}	$< 10^{-10}$
G3: Tight finite resolution	maximum ratio 0.366541	$\leq 1 + 10^{-10}$
G4: Random score quantization	$K_{99} = 14$ at $d = 64$; fit $R^2 = 0.989$	$K_{99} \leq 16$; finite-range fit $R^2 > 0.95$
G5: Tail theorem	maximum loss/bound 0.323037	$\leq 1 + 10^{-10}$
G6: Dense-spectrum limit	PDZ relative error 0.015625; terminal loss 0	relative error < 0.02; crossover exact
G7: Full-support hardness construction	residual 4.494×10^{-15}; fidelity 0.9999999999999996; $\kappa^2 p_{\mathrm{app}} = 1.047458$; $4\kappa^2 F_Q = 1.060359$	errors $< 10^{-12}$; fidelity $> 1 - 10^{-12}$; scaled quantities ≥ 1; full support
G8: Physical measurement contrast	product 48.55%; seven/eight bins 98.34%/99.26%	product $< 60\%$; bins $> 98\%/99\%$
G9: Resource and truncation audit	123 queries; input thresholds 71/114 terms	files generated; arithmetic internally consistent
G10: Fixed-temperature Gibbs scaling	identity error 1.36×10^{-16}; $\kappa_{n=10,\beta=.3} = 5416.7$	identity $< 10^{-12}$; exponential cost exposed
G11: Shrinking Gibbs window	$\kappa_{c=1} \leq 17.56$; $nF_Q \in [0.809, 1.202]$	$\kappa < 20$; $0.75 < nF_Q < 1.30$
G12: Parity-coherence sensor	$\kappa = 3.0$; $F_Q = n/4$; $n+1$ scores through $n = 128$	conditioning, QFI slope, and score count exact

Table S3: Principal reproducibility assets.

Path	Role
`src/qsa_core.py`	Linear algebra, SLD, Fisher objectives, score ascent, spin Hamiltonian, and hardness embedding.
`src/simulate.py`	Deterministic data generation and product-PVM optimization.
`src/plot_pudim_research_quantum.py`	Main-text Figs. 1–3.
`src/plot_supplement.py`	Supplemental Figs. S1–S4.
`src/generate_numbers.py`	Converts `data/summary.json` into LaTeX macros used here.
`validate.py`	Threshold checks with nonzero exit status on failure, run only after regeneration by `reproduce.py`.
`data/*.csv`	Row-level numerical observations.
`data/saddle_landscape.npz`	Full two-dimensional strict-saddle grid.
`data/physical_product_parameters.npy`	Eight optimized local-basis angles at every temperature.
`data/score_compression.csv`	Exact optimal bin counts for all random compression models.
`data/tail_bound.csv`	Constructive tail-theorem partitions, moments, losses, and bounds.
`data/thermodynamic_quantization.csv`	Dense-spectrum PDZ window, exact finite-support formula, and zero-loss crossover.
`data/gibbs_conditioning.csv`	Energy spans, fixed-temperature and c/n condition numbers, and exact shrinking-window QFI through $n = 10$.
`data/parity_sensor.csv`	Constant-conditioned extensive-QFI family and exact score-register size through $n = 128$.
`data/hardness_reduction.csv`	Full-support QLS-to-SLD residuals, fidelity, application success, QFI, zero-score and positivity margins.
`data/pauli_truncation.csv`	Direct-score operator errors and coefficient-tail certificates.
`data/input_pauli_truncation.csv`	Input truncation and induced SLD errors after the Sylvester solve.
`data/resource_estimate.json`	Physical Pauli-LCU and logical-resource assumptions and totals.
`data/resource_operators.csv`	Exact term counts, LCU norms, and largest coefficients for S, D, L.
`data/summary.json`	Machine-readable headline values and software metadata.
`metadata/oracle-interface.json`	Machine-readable hardness input, oracle, normalization, signal-gadget, error-composition, success, and output-boundary audit.
`requirements.lock`	Exact Python dependency lock with package hashes.
`tolerances.json`	Machine-readable scientific validation thresholds.
`attestations/*.json`	Observed commands, environments, resources, validation outcomes, content hashes, and explicit isolation limitations.