Supplementary Information

Quantized composition penalties alter state selection in phase-field Ising models

Zhiyi Liu[1,2], Anping Liu[1], Zicong Zhu[1,*]

[1]Tianli Qiming AI Research Institute, Sichuan Qiming Daren Technology Co., Ltd., Chengdu, China

[2]International Research Center for Complexity Sciences, Hangzhou International Innovation Institute, Beihang University, Hangzhou, China

[*]Correspondence: zhuzc0907@outlook.com

This supplement presents the encoding derivations, finite-precision benchmark, state-level mechanism analysis, sensitivity analyses and continuum references. All numerical models are dimensionless. The study protocol and subsequent diagnostic amendments are documented with the reproducibility materials.

Contents

Study question, units of inference and scope 2

Exact encoding of the constrained four-state lattice 2

Penalty thresholds and ground-set certificates 3

Fixed-grid penalty splitting and its physical meaning 4

Conventional shared-range pure-penalty splitting 5

Constraint-aligned integer compilation 6

Independent verification of the encoding identities 7

Compilation grid and allocation selection 8

Full encoded ground sets and outcome definitions 8

State-level mechanism reanalysis 9

Statistical analysis and interpretation 12

Held-out results and numerical diagnostics 12

Results of the state-level mechanism reanalysis 14

Coefficient perturbations and matched-grid control 21

Independent continuum and representation audit 22

Reproducibility, independent checks and retained records 25

Interpretation of the comparisons 26

Study question, units of inference and scope

The study asks whether finite-precision coefficients preserve the objective and composition constraint of a specified four-state phase-field model. Exact enumeration identifies every encoded ground state, meaning a configuration with the lowest encoded energy, allowing its composition and energy to be evaluated in the original finite-state objective. This isolates changes in the encoded objective from solver error.

There are 16 physical-parameter strata: two site counts, four gradient coefficients and two forcing amplitudes. Each stratum has one training instance, one validation instance and six held-out test instances, yielding $16 + 16 + 96 = 128$ independent random-phase instances. Training screens penalty allocations and validation selects the final allocation; testing evaluates these fixed choices. Different bit budgets, allocations, ground states and coefficient-noise draws within one instance are repeated evaluations rather than independent sampling units. Averaging strata weights the designed random-phase distribution equally.

The dimensionless domain length is $L = 8$, the site counts are $n \in \{6, 8\}$, and the spacing is $h = L/n$. Gradient coefficients are $\kappa \in \{0.1, 0.5, 1.0, 1.8\}$ and forcing amplitudes are $A \in \{0.02, 0.08\}$. The sampled forcing is

$$v_i = \frac{A}{1.5}\left[\cos(2\pi i/n + \theta_1) + \frac{1}{2}\cos(4\pi i/n + \theta_2)\right], \qquad i = 0, \ldots, n-1. \tag{S1}$$

Here the grid index starts at zero, as in the implementation; the main text uses the equivalent one-based coordinates $x_i = (i-1)h$. The two angles shift the spatial bias waves. Both angles are independent uniform draws on $[0, 2\pi)$. The normalization ensures $|v_i| \le A$, rather than forcing each realized field to attain amplitude A. The seed namespace is $[42, 20261014, 3]$; split identifiers 1, 2, and 3 distinguish training, validation and test. Each instance uses a NumPy seed sequence from this namespace followed by split identifier, instance identifier and stream identifier 100.

Exact encoding of the constrained four-state lattice

Consider an even number $n \ge 4$ of sites on a periodic ring of dimensionless length L, with $h = L/n$. At each grid point, ϕ_i represents local composition; the two pure components correspond to $\phi_i = \pm 1$. The discrete (nodal) free energy is the total energy evaluated from these grid-point values, using a double-well local potential and a gradient cost [1, 2], together with a prescribed spatial bias:

$$F(\boldsymbol{\phi}) = h\sum_{i=1}^{n}\left[\frac{(\phi_i^2 - 1)^2}{4} + \frac{\kappa}{2h^2}(\phi_{i+1} - \phi_i)^2 + v_i\phi_i\right], \qquad \kappa \ge 0. \tag{S2}$$

Periodic indexing sets $\phi_{n+1} = \phi_1$. The local potential favors ± 1, the gradient term penalizes neighboring differences, and positive v_i favors negative ϕ_i. The factors h and $1/h$ follow from approximating integration by a grid sum and the derivative by $(\phi_{i+1} - \phi_i)/h$. Equal total composition imposes $\sum_i \phi_i = 0$. The restriction $\phi_i \in \{-1, -1/3, 1/3, 1\}$ admits the exact two-bit map

$$\phi_i = -1 + \frac{2}{3}b_{i0} + \frac{4}{3}b_{i1} = \frac{s_{i0} + 2s_{i1}}{3}, \qquad s_{ia} = 2b_{ia} - 1. \tag{S3}$$

The map represents the four admissible values exactly; the continuum comparison below assesses the separate restriction to those values. The integer-valued composition error is

$$C = \sum_i (b_{i0} + 2b_{i1}) - \frac{3n}{2} = \frac{1}{2}\sum_i (s_{i0} + 2s_{i1}), \qquad \sum_i \phi_i = \frac{2}{3}C. \tag{S4}$$

Because n is even, every infeasible configuration satisfies $|C| \geq 1$.

The local quartic term reduces to a quadratic binary expression on this alphabet:

$$\frac{(\phi_i^2 - 1)^2}{4} = \frac{16}{81}(b_{i0} + b_{i1} - 2b_{i0}b_{i1}) = \frac{8}{81}(1 - s_{i0}s_{i1}). \tag{S5}$$

This is a quadratic unconstrained binary optimization (QUBO) representation without auxiliary spins. With the convention that each pair coefficient is counted once, the physical Ising polynomial is

$$\begin{aligned} F = {} & \frac{8hn}{81} + \frac{5\kappa n}{9h} + \sum_i \left(\frac{4\kappa}{9h} - \frac{8h}{81}\right) s_{i0}s_{i1} \\ & - \frac{\kappa}{9h}\sum_i (s_{i0} + 2s_{i1})(s_{i+1,0} + 2s_{i+1,1}) + \frac{h}{3}\sum_i v_i(s_{i0} + 2s_{i1}). \end{aligned} \tag{S6}$$

The gradient term follows by expanding $\sum_i (\phi_{i+1} - \phi_i)^2 = 2\sum_i \phi_i^2 - 2\sum_i \phi_i\phi_{i+1}$. The on-site identity and this expansion establish exactness directly.

Penalty thresholds and ground-set certificates

The exact penalized Hamiltonian is $H_P = F + PC^2$, $P \geq 0$. If $w = (1, 2, 1, 2, \ldots, 1, 2)$, then

$$PC^2 = \frac{5Pn}{4} + \frac{P}{2}\sum_{a<b} w_a w_b s_a s_b. \tag{S7}$$

The composition penalty therefore adds dense pair couplings but no fields. Let $F_* = \min_{C=0} F$ and define

$$P_{\mathrm{crit}} = \max_{C \neq 0} \frac{F_* - F}{C^2}. \tag{S8}$$

For $P > P_{\mathrm{crit}}$, every infeasible state satisfies $F + PC^2 > F_*$, so all exact ground states are feasible. For $P < P_{\mathrm{crit}}$, some infeasible state has lower energy than F_*. Equality permits an infeasible ground-state tie. The threshold is evaluated by exhaustive enumeration in this study.

The following sufficient bound uses a feasible reference with energy F_{ref} and a lower bound $F_{\mathrm{lb}} \leq \min F$:

$$P > F_{\mathrm{ref}} - F_{\mathrm{lb}}. \tag{S9}$$

Indeed, $C^2 \geq 1$ gives $F + PC^2 \geq F_{\mathrm{lb}} + P > F_{\mathrm{ref}}$ for every infeasible state. In Eq. (S2), $F_{\mathrm{lb}} = -h\sum_i |v_i|$ is valid. A configuration with one contiguous half-ring at $+1$ and the other at -1 is feasible and has gradient energy $4\kappa/h$. This bound can be conservative.

These bounds specialize standard penalty and coefficient-perturbation arguments to this encoding [3–5]. Let θ contain the nonconstant fields and pair coefficients of H_P, $R = \max_a |\theta_a|$, and $q = 2^{B-1} - 1$. Conventional compilation uses

$$\Delta = R/q, \qquad \widetilde{\theta}_a = \Delta\,\mathrm{rint}(\theta_a/\Delta), \qquad |\widetilde{\theta}_a - \theta_a| \leq \Delta/2, \tag{S10}$$

where rint is nearest-integer rounding with ties to even. The nominal signed bit budget has $2q+1$ represented levels, including zero. An all-zero Hamiltonian is handled separately. Constants are omitted from compilation because they cannot change the ordering of configurations.

For two configurations differing in k of the $p = 2n$ spins, only k field features and $k(p-k)$ pair features change sign. Thus

$$\begin{aligned}\left|[\widetilde{H}(s) - \widetilde{H}(t)] - [H(s) - H(t)]\right| &\le \frac{R}{q}k(p-k+1) \\ &\le \frac{R}{q}n(n+1) \equiv B_{\mathrm{ord}}. \end{aligned} \tag{S11}$$

The first inequality sums the changed feature errors, each at most R/q. The maximum of $k(p+1-k)$ over integer k is $\lfloor (p+1)^2/4 \rfloor = n(n+1)$. A coefficient-dependent alternative, $2\sum_a |\widetilde{\theta}_a - \theta_a|$, can tighten the global bound.

Define the strict physical feasible gap and constraint gap by

$$\Delta_f = \min_{C=0,\, F>F_*} (F - F_*), \tag{S12}$$

$$\Delta_c(P) = \min_{C\neq 0} (F + PC^2 - F_*). \tag{S13}$$

If every feasible state is a physical ground state, take $\Delta_f = +\infty$. If $\Delta_f > B_{\mathrm{ord}}$ and $\Delta_c(P) > B_{\mathrm{ord}}$, every quantized ground state is a feasible physical ground state. To prove the claim, compare each infeasible or feasible suboptimal state with any feasible physical ground state. Its exact energy difference exceeds the maximal ordering perturbation, so it cannot become a quantized ground state. The certificate is sufficient, not necessary; failure to certify does not establish failure of compilation.

This result separates insufficient penalties from lost physical resolution. Let R_0 be the largest unpenalized physical coefficient and set $\eta_B = n(n+1)/q$. Since $R(P) \le R_0 + 2P$, the conservative conditions

$$P > F_{\mathrm{ref}} - F_{\mathrm{lb}} + \eta_B(R_0 + 2P), \qquad \Delta_f > \eta_B(R_0 + 2P) \tag{S14}$$

give the explicit sufficient interval, when $2\eta_B < 1$,

$$\frac{F_{\mathrm{ref}} - F_{\mathrm{lb}} + \eta_B R_0}{1 - 2\eta_B} < P < \frac{\Delta_f/\eta_B - R_0}{2}. \tag{S15}$$

The interval can be empty even when successful compilations exist.

Fixed-grid penalty splitting and its physical meaning

Let $\Delta > 0$ be a fixed coefficient step and define the unclipped nearest-integer quantizer

$$Q_\Delta(u) = \Delta\, \mathrm{rint}(u/\Delta). \tag{S16}$$

Analytical statements use exact real arguments and the declared nearest-even tie rule. Implementation diagnostics retain the actual floating-point arguments; the stabilized-half control is described separately below. The same step is used for every pair coefficient. For the pure penalty PC^2, define its three rounded pair classes by

$$a_\Delta = Q_\Delta(P/2), \qquad b_\Delta = Q_\Delta(P), \qquad c_\Delta = Q_\Delta(2P). \tag{S17}$$

Here a_Δ couples two low-weight spins, b_Δ a low- and a high-weight spin, and c_Δ two high-weight spins. The counts are summarized by $X = \sum_i s_{i0}$ and $Y = \sum_i s_{i1}$. Expanding the three pair sums gives

$$\sum_{i<j} s_{i0}s_{j0} = (X^2 - n)/2, \tag{S18}$$

$$\sum_{i,j} s_{i0}s_{j1} = XY, \tag{S19}$$

$$\sum_{i<j} s_{i1}s_{j1} = (Y^2 - n)/2. \tag{S20}$$

The middle sum includes the within-site pair. Therefore the rounded nonconstant penalty is

$$\widetilde{H}_{P,\Delta} = \frac{a_\Delta}{2}(X^2 - n) + b_\Delta XY + \frac{c_\Delta}{2}(Y^2 - n). \tag{S21}$$

On $C = 0$, the identity $X = -2Y$ yields the exact reduction

$$\widetilde{H}_{P,\Delta}\Big|_{C=0} = -\frac{n}{2}(a_\Delta + c_\Delta) + \beta_\Delta Y^2, \qquad \beta_\Delta = 2a_\Delta - 2b_\Delta + \frac{c_\Delta}{2}. \tag{S22}$$

Adding or retaining the omitted constant $5Pn/4$ changes only the state-independent term. Thus $\beta_\Delta = 0$ means that the rounded penalty remains constant on the feasible sector; a negative value favors larger $|Y|$, while a positive value favors smaller $|Y|$. For even $n \geq 4$ the feasible set contains more than one value of Y^2, so this condition is also necessary for constancy. Feasible values obey $Y \in 2\mathbb{Z}$ and $|Y| \leq n/2$; when n is divisible by four, $|Y| = n/2$ is attainable.

Let N_r count sites at each of the four field levels r. The encoding gives

$$Y = -N_{-1} - N_{-1/3} + N_{+1/3} + N_{+1}. \tag{S23}$$

Zero composition is $-3N_{-1} - N_{-1/3} + N_{+1/3} + 3N_{+1} = 0$. Eliminating the middle-level population difference proves

$$Y = 2(N_{-1} - N_{+1}), \qquad \beta_\Delta Y^2 = 4\beta_\Delta(N_{-1} - N_{+1})^2. \tag{S24}$$

The penalty therefore biases the populations of the four physical levels, even when their weighted mean is exactly zero. It does not distinguish spatial rearrangements with the same four populations: their local gradient and bias energies are determined by the physical objective.

For the ideal conventional pure-penalty grid, $\Delta = 2P/q$ and $q = 2^{B-1}-1$, Eq. (S22) reduces to the expression proved in the following section. In floating-point computations, a value intended to be a half-integer can lie just below or above it. The reanalysis therefore computes β_Δ from the three rounded classes rather than substituting $-P/q$ when the recorded range is merely close to $2P$. The fixed-grid identity applies to other steps without renormalizing the three coefficient classes separately. It describes an isolated rounded penalty. Joint rounding of the physical and penalty coefficients requires the residual term defined in the Supplementary Note "Paired-state energy account"; the isolated expression alone does not give the ordering of the full compiled model.

Conventional shared-range pure-penalty splitting

Consider only PC^2, with $P > 0$. Its largest coefficient is $R = 2P$, and its normalized pair coefficients have three classes: 1/4, 1/2, and 1. For $B \geq 3$, $q = 2^{B-1} - 1$ is congruent to three modulo

four, giving

$$a = \frac{1}{4} + \frac{1}{4q}, \qquad b = \frac{1}{2} + \frac{1}{2q}, \qquad c = 1 \tag{S25}$$

after the declared rounding. With the encoding-spin sums defined above, the rounded pair energy, divided by $2P$, is

$$\frac{a}{2}(X^2 - n) + bXY + \frac{1}{2}(Y^2 - n). \tag{S26}$$

On the feasible subspace $X = -2Y$, multiplication by $2P$ yields

$$\widetilde{H}_{\text{penalty}}\Big|_{C=0} = \text{constant} - \frac{P}{q}Y^2. \tag{S27}$$

For n divisible by four, feasible states include $Y = 0$ and $|Y| = n/2$, so the spurious feasible energy range is $Pn^2/(4q)$. The constraint is exactly satisfied throughout this comparison. The distortion therefore concerns physical ordering within the feasible set, rather than composition error alone.

The coefficient $-P/q$ follows from the specified grid and nearest-even rounding. The full-model analysis below retains joint rounding, including branch changes when a physical coefficient shifts a half-grid penalty value.

Independent enumeration of all 65,536 states for $n = 8$ finds 8,092 feasible states before quantization. At each tested budget of 4, 5, 6, 8, 10 and 12 bits, the pure compiled penalty has 56 ground states, all feasible and all with $Y^2 = 16$. The unweighted mean of Y^2 over the full exact feasible set is 1.660899654. This comparison counts the complete ground-state sets. A separate three-bit stress case has 16 compiled ground states, all infeasible with $|C| = 2$; three bits is outside the main benchmark grid.

Constraint-aligned integer compilation

Let $M > 0$ be the maximum absolute nonconstant coefficient of the physical energy. Choose an integer $\ell \geq 1$ satisfying $4\ell < q$, and define

$$\delta = \frac{M}{q - 4\ell}, \qquad P_{\text{eff}} = 2\ell\delta. \tag{S28}$$

Round the physical fields and pairs to integers, $k_F = \text{rint}(\theta_F/\delta)$, and add the penalty pair integers $\ell w_a w_b$. Since $|k_F| \leq q - 4\ell$ and $w_a w_b \leq 4$, all stored fields and pair coefficients obey $|k| \leq q$. The compiled polynomial is

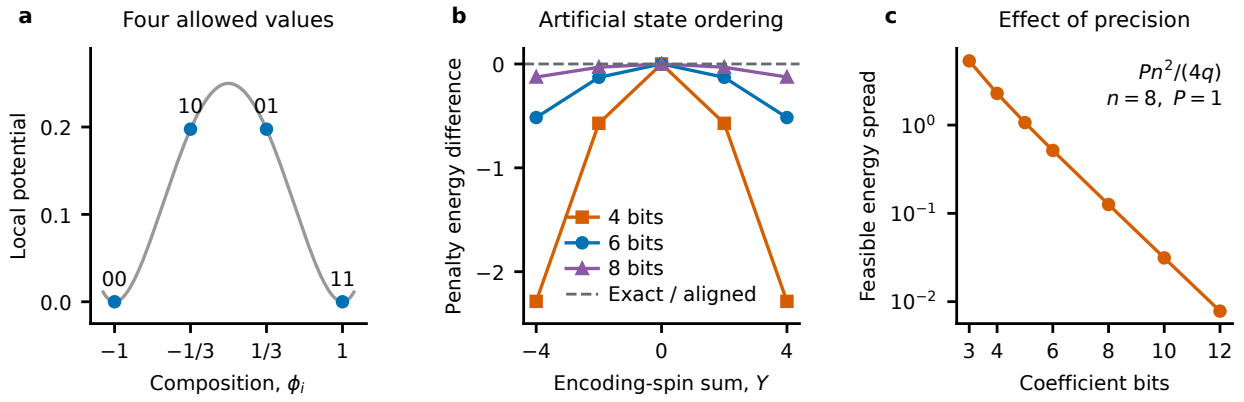
$$\delta \sum_a k_a z_a(s) = F_{\text{rounded}}(s) + P_{\text{eff}} C(s)^2 + \text{constant}. \tag{S29}$$

Thus the penalty is exactly constant on the feasible subspace. The physical objective remains rounded, with errors that grow as the penalty allocation increases. If a compiled ground state is feasible, its physical energy excess $F - F_*$ is at most $\delta n(n+1)$ by Eq. (S11) applied to the physical coefficient errors alone. This conditional bound does not apply to an infeasible returned state. Exact constraint preservation does not guarantee that the penalty is sufficiently large or that the physical optimum survives.

The comparison with conventional compilation must match P_{eff} and report both coefficient steps. Conventional normalization generally gives $R/q \leq \delta$, so it can have a finer step while failing to preserve the constraint's coefficient ratios. At a matched step, direct full-objective rounding is a useful negative control: integer translation commutes with nearest rounding except at half-grid ties with odd shifts. Specifically, $\text{rint}(0.5 + 1) = 2$ but $\text{rint}(0.5) + 1 = 1$. Actual coefficient differences must therefore be checked rather than assumed to vanish.

Independent verification of the encoding identities

The independent script `verify_encoding_theory.py` does not import the benchmark implementation. Eleven checks cover exact physical/Ising and binary/QUBO agreement, the composition identity, strict penalty thresholds, state-pair ordering bounds, a sufficient ground-set certificate, an oracle-free feasibility bound, the optional field-gauge minimax identity, integer-aligned constraint preservation, pure-penalty leakage, and the nearest-even translation caveat. Direct physical and Ising energies agree to 7.11×10^{-15} over 24 exhaustive parameter cases. These checks establish the stated algebraic and numerical identities.



Supplementary Figure S1: **Exact finite-state encoding and spurious feasible ordering from a rounded composition penalty. a**, Four admissible values of the order parameter and their two-bit labels (b_{i0}, b_{i1}) on the local double-well potential. The continuous curve is a guide to the specified local function; only the four marked levels enter the optimization model. **b**, The rounded pure penalty among exactly feasible states, relative to $Y = 0$, with $Y = \sum_i s_{i1}$, $n = 8$ and $P = 1$. Markers are exhaustive enumeration results and follow $-PY^2/q$ for the declared signed grid with nearest-even rounding. The exact or constraint-aligned penalty is constant on this subspace. **c**, The induced feasible energy range follows $Pn^2/(4q)$. All values are dimensionless. The calculation exhaustively enumerates the pure penalty, with no sampling uncertainty. Joint rounding is treated separately in the state-level analysis.

Compilation grid and allocation selection

The nominal signed coefficient budgets are $B \in \{4, 5, 6, 8, 10, 12\}$. For $q = 2^{B-1} - 1$, the six prespecified fractions $\alpha \in \{0.04, 0.08, 0.16, 0.32, 0.64, 0.9\}$ are mapped to integer allocations by

$$\ell = \mathrm{clip}\left(\mathrm{rint}(\alpha q/4), 1, \lfloor (q-1)/4 \rfloor\right). \tag{S30}$$

Duplicate integer allocations are removed. The resulting candidate counts are 1, 3, 4, 6, 6 and 6, respectively. A complete instance thus has $2(1 + 3 + 4 + 6 + 6 + 6) = 52$ compiled Hamiltonians. There are 832 training, 832 validation and 4,992 test compilations in the full grid.

Each grid comparison at a shared (B, ℓ) uses the same physical model and effective penalty $P = 2\ell\delta$ for conventional and aligned compilation. The coefficient steps generally differ. The conventional compiler uses the entire available shared range, with step R/q, while the aligned compiler uses $\delta = M/(q - 4\ell)$. The full-grid comparison therefore matches the penalty, not the physical coefficient step. The conventional compiler can use a finer step because $R/q \leq \delta$.

The selected policy is determined separately for every precision and compiler. Candidates are ordered by mean training target-energy fidelity loss, then smaller ℓ. Up to three candidates advance to validation. The validation winner is ordered by validation loss, then training loss, then smaller ℓ. All validation candidates were computed and archived; only the training-selected finalists were eligible. The test set did not determine allocations or tolerance.

Supplementary Table S1: Locked integer allocations selected before held-out evaluation. The conventional and aligned allocations differ at eight bits, so the primary selected-policy contrast at that precision does not match penalty strength. The full-grid same-allocation comparison remains available separately.

B	q	Unique candidates	Conventional ℓ	Aligned ℓ
4	7	1	1	1
5	15	3	1	1
6	31	4	2	2
8	127	6	20	10
10	511	6	41	41
12	2047	6	164	164

Full encoded ground sets and outcome definitions

State identifiers enumerate the $2n$ site-major bits, with low-index bits in the least significant positions. Both global spin reversals and every other distinct configuration are retained; the forcing generally removes spin-reversal symmetry. There are $4^6 = 4{,}096$ states at six sites and $4^8 = 65{,}536$ states at eight sites. A feature matrix contains each spin and each unordered pair product exactly once. Integer coefficients are multiplied by these integer-valued features. The resulting baseline energy sums are exactly representable in double precision at the tested ranges. The implementation retains all states within 10^{-9} of the lowest integer energy. This tolerance does not merge distinct baseline integer energies, which differ by at least one unit.

For an encoded ground-state set G, the primary counting endpoints are

$$f_{\text{feasible}} = \frac{1}{|G|} \sum_{s \in G} \mathbf{1}\{C(s) = 0\}, \tag{S31}$$

$$f_{\text{faithful}} = \frac{1}{|G|} \sum_{s \in G} \mathbf{1}\{C(s) = 0,\ F(s) \le F_* + 0.001L\}. \tag{S32}$$

These fractions are denoted by f_C and f_E, respectively, in the main text. The tolerance is fixed on the dimensionless nodal-energy scale before the benchmark. Fractions use equal counting over the complete encoded ground set, including infeasible states. They describe the composition of that set, not a device sampling distribution.

Secondary endpoints include the ground-set multiplicity, whether every ground state is feasible, whether any ground state is faithful, and the mean and worst feasible physical regret $(F - F_*)/L$. Conditional regret is undefined, and stored as missing, if no encoded ground state is feasible; it is never replaced by zero. The feasible-sector distortion is the range over all $C = 0$ states of $\widetilde{H}(s) - F(s)$, after irrelevant constant offsets. It measures coefficient-induced ordering distortion independently of which ground states happen to be selected. It can be nonzero even when the ground set remains faithful. The archived `pcrit` field stores $\max(0, P_{\text{crit}})$; the algebraic threshold in Eq. (S8) itself can be negative. These are distinct conventions, and the clipped diagnostic is not used to choose test-instance allocations.

Additional morphology summaries describe feasible returned states only: the root-mean-square order parameter, the number of sign-changing periodic edges, and a mean squared distance to a stated exact physical minimizer. Composition error follows the stored definition: $2|C|/3$ is the absolute sum of nodal values; its integrated counterpart is $2h|C|/3$. Exact feasibility is the same under either convention.

State-level mechanism reanalysis

Analyzed configurations

The original benchmark fixes the physical instances, both compilation rules, all allocation candidates, and the locked policies. The mechanism reanalysis covers all 52 compiled columns in each of the 16 validation and 96 test instances, giving 5,824 configurations. It uses the archived coefficients and complete state identifiers. It adds physical population counts, state-selection events and energy-difference accounts without changing the original fidelity threshold or selecting new compiler parameters. Configuration-level summaries refer to the full grid; instance-level summaries retain the original random-phase instance as their unit. The results below identify the analyzed configurations and the states used for illustration.

Exact minimizers and distinct selection events

Let $\mathcal{S} = \{-1, +1\}^{2n}$ and $\mathcal{S}_0 = \{s \in \mathcal{S} : C(s) = 0\}$. Define the exact physical target and the two compiled minimizer sets by

$$G_* = \arg\min_{s \in \mathcal{S}_0} F(s), \qquad \widetilde{G}_0 = \arg\min_{s \in \mathcal{S}_0} \widetilde{H}(s), \qquad \widetilde{G} = \arg\min_{s \in \mathcal{S}} \widetilde{H}(s). \tag{S33}$$

Thus $\widetilde{G}_0$ describes ordering within the feasible sector even when an infeasible state lies below it in the full encoded problem. For a physical-energy tolerance τ, let

$$T_\tau = \{s \in \mathcal{S}_0 : F(s) \le F_* + \tau\}. \tag{S34}$$

Three mutually exclusive classes describe the feasible sector:

$$\text{preserved:} \quad \widetilde{G}_0 \subseteq T_\tau, \tag{S35}$$

$$\text{tie expansion:} \quad \widetilde{G}_0 \cap T_\tau \neq \varnothing, \quad \widetilde{G}_0 \setminus T_\tau \neq \varnothing, \tag{S36}$$

$$\text{strict reversal:} \quad \widetilde{G}_0 \cap T_\tau = \varnothing. \tag{S37}$$

In the last case, the minimum feasible compiled energy is strictly below that of every state in T_τ. A tie expansion instead admits both a target state and a physically worse state at the same compiled minimum. The exact-target classification uses the archived physical-optimum comparison tolerance 10^{-10}. The near-optimal classification uses $\tau = 0.001L$ with the original 10^{-12} numerical comparison guard. These two classifications distinguish exact ordering from the original near-optimal fidelity endpoint.

Feasibility is assessed separately on the global set $\widetilde{G}$: "all feasible" means every member is in $\mathcal{S}_0$; "mixed" means some but not all members are feasible; and "none feasible" means no member is feasible. The reanalysis reports the cross-classification, so a feasible ordering event is never mistaken for a guarantee of global feasibility. Original integer coefficients and integer-valued features give exactly representable integer energy sums at the tested sizes. The reanalysis uses exact equality for their minima and verifies agreement with the archived 10^{-9} grid-unit tie rule.

Paired-state energy account

Separate the exact physical energy into

$$F_{\text{local}}(s) = \frac{h}{4} \sum_i (\phi_i(s)^2 - 1)^2, \tag{S38}$$

$$F_{\text{gradient}}(s) = \frac{\kappa}{2h} \sum_i (\phi_{i+1}(s) - \phi_i(s))^2, \tag{S39}$$

$$F_{\text{bias}}(s) = h \sum_i v_i \phi_i(s). \tag{S40}$$

Thus $F = F_{\text{local}} + F_{\text{gradient}} + F_{\text{bias}}$. For the actual compiled coefficient vector, let

$$e(s) = \widetilde{H}(s) - F(s) - PC(s)^2, \tag{S41}$$

with any consistent constant convention. For a returned state s and physical reference $t \in G_*$, write $\mathrm{d}A = A(s) - A(t)$. Direct subtraction gives

$$\mathrm{d}\widetilde{H} = \mathrm{d}F_{\text{local}} + \mathrm{d}F_{\text{gradient}} + \mathrm{d}F_{\text{bias}} + P\mathrm{d}(C^2) + \mathrm{d}e. \tag{S42}$$

For two feasible states, the penalty difference vanishes. A state with $\mathrm{d}F > 0$ ties or outranks the reference exactly when the error difference is at most $-\mathrm{d}F$. Equality produces a pairwise tie; a strict inequality reverses the ordering relative to that reference. Testing the best compiled member of G_* establishes whether every physical optimum has been displaced.

The isolated fixed-grid penalty can also be separated from joint-rounding effects. Let θ_F and θ_P contain physical and penalty coefficients, let $\widetilde{\theta}$ contain the actual compiled coefficients, and let $\boldsymbol{\psi}(s)$ contain all spin and pair features. At the compiler's actual coefficient step, define

$$\boldsymbol{e}_{F,\Delta} = Q_\Delta(\theta_F) - \theta_F, \tag{S43}$$

$$\boldsymbol{I}_\Delta = \widetilde{\theta} - Q_\Delta(\theta_F) - Q_\Delta(\theta_P). \tag{S44}$$

The first vector contains physical-coefficient rounding errors; the second is the residual relative to rounding the two components separately. For conventional compilation, $\widetilde{\theta} = Q_\Delta(\theta_F + \theta_P)$, so this residual is the interaction introduced by joint rounding. For integer-aligned compilation, $\widetilde{\theta} = Q_\Delta(\theta_F) + \theta_P$ and $Q_\Delta(\theta_P) = \theta_P$, giving $\boldsymbol{I}_\Delta = 0$. On the feasible sector, Eq. (S22) then gives

$$\begin{aligned} \mathrm{d}\widetilde{H} &= \mathrm{d}F + \mathrm{d}e_F + \beta_\Delta\, \mathrm{d}(Y^2) + \mathrm{d}I, \\ \mathrm{d}e_F &= \boldsymbol{e}_{F,\Delta}^{\mathsf{T}}[\boldsymbol{\psi}(s) - \boldsymbol{\psi}(t)], \\ \mathrm{d}I &= \boldsymbol{I}_\Delta^{\mathsf{T}}[\boldsymbol{\psi}(s) - \boldsymbol{\psi}(t)]. \end{aligned} \tag{S45}$$

These evaluations use the implemented rounding rule and the stored compiled coefficients, retaining floating-point branch decisions. The terms can reinforce or cancel one another. The isolated penalty term need not explain a full-model reversal. For constraint-aligned compilation, $\beta_\Delta = 0$ and $\boldsymbol{I}_\Delta = 0$, so the corresponding difference is $\mathrm{d}F + \mathrm{d}e_F$ evaluated at that compiler's step.

To summarize a component u across the feasible sector, the analysis uses the centered root mean square

$$\mathrm{RMS}_0(u) = \left[\frac{1}{|\mathcal{S}_0|}\sum_{s\in\mathcal{S}_0}(u(s) - \overline{u})^2\right]^{1/2}, \qquad \overline{u} = \frac{1}{|\mathcal{S}_0|}\sum_{s\in\mathcal{S}_0} u(s). \tag{S46}$$

Centering removes energy offsets. Component root mean squares are not additive, because the components can be correlated.

Selection of illustrative states

Examples are drawn from the validation set only. Candidates are ordered by increasing bit budget, instance identifier, then integer allocation. The first two examples are the first conventional-compiler cases classified as a near-optimal strict reversal and tie expansion, respectively, with at least one feasible global ground state. The third example is the first same-instance, same-bit, same-allocation pair for which aligned fidelity is lower and both global ground sets are entirely feasible. If a category is absent, it is reported absent without substituting a test-set example.

Within a selected case, the displayed witness is the feasible-sector compiled minimizer with the highest physical energy; the lowest state identifier breaks a numerical tie. The reference is the exact physical minimizer with the lowest compiled integer energy, again resolving ties by state identifier. This rule displays the worst selected feasible state alongside the physical optimum with the best compiled ranking. The full ground-set sizes and event counts accompany each example; numerical outcomes always use all tied states rather than the displayed pair alone.

The state identifiers reconstruct the complete vector $\boldsymbol{\phi}$ and its four population counts, C and Y. The corresponding records retain $\mathrm{d}F$, $\mathrm{d}e_F$, $\beta_\Delta\mathrm{d}(Y^2)$, $\mathrm{d}I$ and $\mathrm{d}\widetilde{H}$. These fields expose the physical meaning and the exact energy balance behind the displayed state change.

Statistical analysis and interpretation

For each of the 96 held-out instances, outcomes are averaged equally over the six precisions after applying the locked method-specific allocations. The two primary contrasts are aligned minus conventional target-energy fidelity and aligned minus conventional composition feasibility. The independent unit is the forcing-phase instance, not the precision or a member of a ground set.

Uncertainty is estimated by 10,000 paired stratified bootstrap draws. Within each of the 16 strata, six instances are sampled with replacement, preserving the paired compiler outcomes. Percentile intervals use the 2.5th and 97.5th percentiles. Two-sided random-sign tests use 99,999 draws of independent signs applied to the instance-level paired contrasts, with a plus-one correction. Their null assumes exchangeability of paired labels or symmetric signed differences. Holm adjustment is applied to the two primary tests. Root seed 42 fixes analysis randomness.

Precision-specific results, the full allocation grids, coefficient-noise sensitivity, certificates and morphology are descriptive diagnostics. They are not additional independent confirmatory tests. Intervals summarize the designed random-phase distribution.

Held-out results and numerical diagnostics

The selected aligned policy has lower aggregate fidelity than the selected conventional policy. Both have nearly complete composition feasibility. The precision-specific results below show where the fidelity difference changes direction.

Supplementary Table S2: Primary paired contrasts over 96 held-out instances. Outcomes are averaged equally over the six precision budgets within each instance. Differences are aligned minus conventional. Confidence intervals (CIs) are paired stratified bootstrap intervals; p values use the prespecified random-sign test and Holm correction across two outcomes.

Outcome	Difference	95% CI	Raw p	Holm p
Target-energy fidelity	-0.021420	[-0.032742, -0.010147]	0.01046	0.02092
Composition feasibility	-0.000926	[-0.002315, 0.000000]	0.49858	0.49858

Supplementary Table S3: Selected-policy means over held-out instances, after averaging precision budgets. Intervals are 95% stratified instance-bootstrap confidence intervals (CIs) across the 96 test instances.

Compiler	Outcome	Mean [95% CI]
Conventional	Target-energy fidelity	0.7964 [0.7862, 0.8064]
Conventional	Composition feasibility	1.0000 [1.0000, 1.0000]
Conventional	Ground-set size	2.3698 [2.3299, 2.4097]
Aligned	Target-energy fidelity	0.7750 [0.7640, 0.7855]
Aligned	Composition feasibility	0.9991 [0.9977, 1.0000]
Aligned	Ground-set size	5.7812 [5.7014, 5.8733]

Supplementary Table S4: Precision-specific descriptive results for the locked policies. Each row contains 96 independent instances; intervals are 95% stratified instance-bootstrap confidence intervals (CIs). Ground-state fractions are counting conventions, not device sampling probabilities.

Bits	Compiler	Target-energy fidelity [95% CI]	Composition feasibility	Mean $\|G\|$
4	Conventional	0.3838 [0.3591, 0.4100]	1.0000	6.271
4	Aligned	0.2525 [0.2306, 0.2752]	1.0000	24.729
5	Conventional	0.6997 [0.6632, 0.7326]	1.0000	2.573
5	Aligned	0.6139 [0.5743, 0.6536]	0.9944	4.198
6	Conventional	0.7418 [0.7088, 0.7726]	1.0000	1.917
6	Aligned	0.8060 [0.7661, 0.8429]	1.0000	2.302
8	Conventional	0.9531 [0.9323, 0.9722]	1.0000	1.406
8	Aligned	0.9774 [0.9601, 0.9931]	1.0000	1.323
10	Conventional	1.0000 [1.0000, 1.0000]	1.0000	1.042
10	Aligned	1.0000 [1.0000, 1.0000]	1.0000	1.094
12	Conventional	1.0000 [1.0000, 1.0000]	1.0000	1.010
12	Aligned	1.0000 [1.0000, 1.0000]	1.0000	1.042

Both selected compilers reach fidelity 1.0 at 10 and 12 bits for every held-out instance at the prespecified near-optimal tolerance. This does not imply that every encoded ground state is an exact physical minimizer. At four and five bits, the aligned allocation yields lower descriptive fidelity, while six-bit fidelity is higher. The eight-bit policy comparison also changes the selected penalty allocation. The same-allocation comparisons separate this allocation choice from the compilation rule.

Supplementary Table S5: Coefficient-noise sensitivity with locked allocations. Four draws are averaged within each instance before bootstrapping 96 instances; intervals are 95% confidence intervals (CIs). Standard deviations are in each compiler's own coefficient grid units. These diagnostic comparisons are not additional confirmatory tests.

Bits	σ	Compiler	Target-energy fidelity [95% CI]	Feasibility
6	0.05	Conventional	0.7500 [0.7135, 0.7839]	1.0000
6	0.05	Aligned	0.8203 [0.7760, 0.8620]	1.0000
6	0.25	Conventional	0.6823 [0.6406, 0.7240]	0.9922
6	0.25	Aligned	0.7031 [0.6615, 0.7448]	0.9844
8	0.05	Conventional	0.9557 [0.9245, 0.9818]	1.0000
8	0.05	Aligned	0.9844 [0.9714, 0.9948]	1.0000
8	0.25	Conventional	0.8698 [0.8411, 0.8958]	1.0000
8	0.25	Aligned	0.8984 [0.8672, 0.9271]	1.0000

Across the full held-out grid, the sufficient strict-ground-set certificate holds for 306 of 2,496 conventional compilations and 249 of 2,496 aligned compilations. The certificate protects exact ordering against every admissible half-bin error; actual errors can be smaller or cancel.

Direct full-objective rounding on the aligned grid differs from the constructed aligned integer vector in 372 of 2,496 allocation-instance combinations. The largest observed integer coefficient difference is 1. Every differing combination has at least one physical coefficient within the recorded numerical tolerance of a half-grid location. Exact half-tie parity and floating-point placement around those ties must therefore be distinguished from an algebraic failure of the aligned constraint identity.

Supplementary Table S6: Exploratory sensitivity to the physical-energy tolerance, reconstructed from the archived locked-policy ground sets without selecting new allocations. The zero-tolerance row retains the stated numerical energy tolerance.

Tolerance per L	Conventional fidelity	Aligned fidelity
0	0.710503	0.678695
0.0001	0.724103	0.697016
0.001	0.796393	0.774972
0.01	0.928241	0.896776

Results of the state-level mechanism reanalysis

Joint rounding changes the route to physical error

The fixed-grid analysis reconstructs every saved coefficient vector and every encoded ground-state tie for the 16 validation and 96 test instances: 5,824 method–precision–allocation conditions in total. The maximum residual in the feasible-state decomposition is 2.2×10^{-14}. The analysis separates the ordering of feasible states from whether the global encoded minima are feasible.

Under the locked policies, conventional rounding produces 48 strict reversals and 110 tie expansions at the original near-optimal tolerance, compared with 10 and 164 for aligned compilation (Table S7). The aligned method therefore has fewer strict reversals but more enlarged minimum sets containing physically poorer states. This distinction explains why counting only whether a good minimum survives would miss part of the fidelity loss. Mean target-energy fidelity is 0.796393 for conventional rounding and 0.774972 for aligned compilation.

The pure-penalty formula does not by itself explain the observed full-objective reversals. Across all 154 conventional strict-reversal conditions in the full test grid, every feasible compiled minimum and every physical optimum has $Y = 0$; these conditions contain 1,196 feasible compiled minima when counted separately within each condition. Thus $\beta \Delta Y^2 = 0$ in all of these comparisons. The same holds for the 48 locked reversals and their 260 feasible compiled minima. Their ordering changes arise through physical rounding error and the joint-rounding interaction. For the first validation reversal selected by the fixed rule, $\Delta F = 0.169023$, $\Delta e_F = 0.000289$, and $\Delta I = -0.507937$, giving $\Delta \widetilde{H} = -0.338624$ (Fig. S2). The companion validation examples show a different route: a coarse physical grid removes an energy distinction and enlarges the minimum set.

Fixed-grid decomposition and complete-set event classification

For each saved coefficient vector, all spin configurations were enumerated independently of the original benchmark implementation. Integer compiled energies were evaluated exactly at the tested sizes, and every minimizing state was retained. Physical energies were independently recomputed from the nodal functional. No solver samples, optimization training, new noise experiments or parameter retuning were introduced.

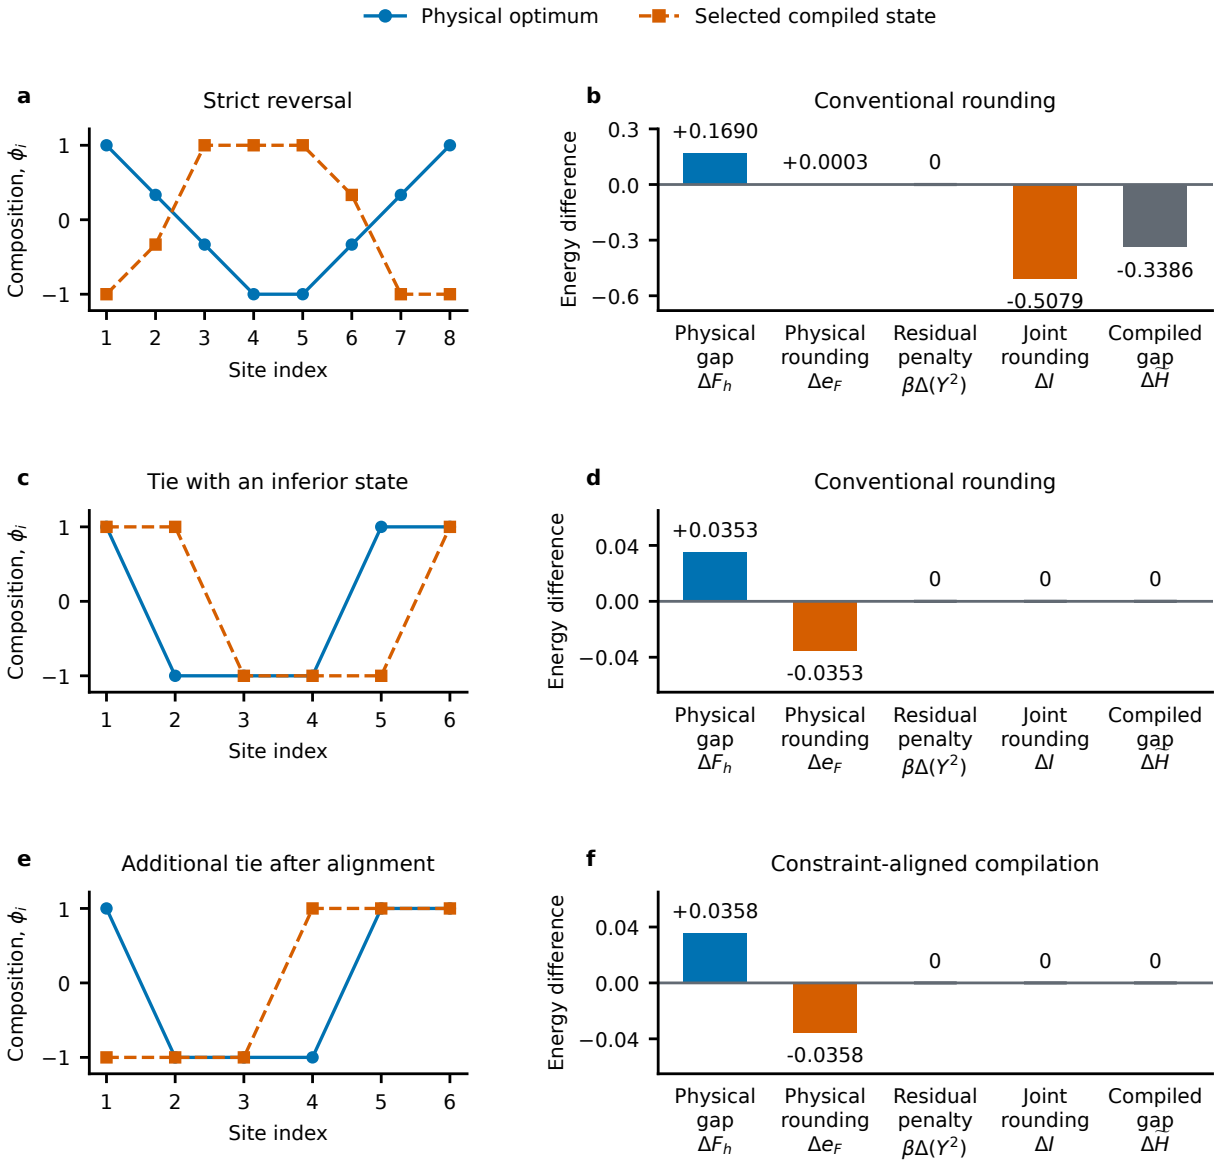
Let G_C be the minimum set of the compiled energy restricted to exactly composition-feasible states. A preserved ordering means that every member of G_C belongs to the acceptable physical set. Tie expansion means that G_C contains both acceptable and unacceptable physical states. A

Supplementary Table S7: Ordering events and contributions to target-energy fidelity loss. Counts refer to 576 instance–precision conditions per method, from 96 independent instances. Exact and near-optimal classifications use tolerances 10^{-10} and $0.001L$, respectively. A strict reversal excludes acceptable states from the feasible-sector minima; a tie expansion retains both acceptable and inferior states. Loss contributions weight each condition equally and sum to $100(1-\overline{f}_E)$. The final column is aligned minus conventional, calculated before rounding; it reports condition-count differences in the upper sections and percentage-point differences in the loss section.

Event or loss component	Conventional	Aligned	Difference
Exact optimum: condition counts			
Preserved	353	332	-21
Tie expansion	163	221	+58
Strict reversal	60	23	-37
Near optimum: condition counts			
Preserved	418	402	-16
Tie expansion	110	164	+54
Strict reversal	48	10	-38
Global feasibility: condition counts			
All global minima feasible	576	574	-2
Mixed feasible/infeasible minima	0	2	+2
No global minimum feasible	0	0	0
Near-optimal fidelity loss: percentage points			
Inferior feasible states: strict reversal	8.33	1.74	-6.60
Inferior feasible states: tie expansion	12.03	20.67	+8.65
Infeasible states	0.00	0.09	+0.09
Total loss	**20.36**	**22.50**	**+2.14**

Supplementary Table S8: State-pair energy decomposition for the selected validation state pairs. Differences are selected state minus reference physical optimum, with all values in the dimensionless physical energy units of the nodal model. Every state has zero composition; the residual penalty contribution is zero in each displayed comparison.

Validation example	ΔF_h	Δe_F	$\beta\Delta Y^2$	ΔI	$\Delta\widetilde{H}$
Strict reversal	0.169023	0.000289	0	-0.507937	-0.338624
Tie expansion	0.035346	-0.035346	0	0	0
Aligned adverse case	0.035764	-0.035764	0	0	0



Supplementary Figure S2: Validation examples of full-objective state selection. Panels a,c,e compare field values ϕ_i of the physical optimum and the selected feasible compiled minimum. Panels b,d,f separate their energy difference into the physical gap ΔF_h, physical rounding Δe_F, residual penalty $\beta\Delta(Y^2)$ and compilation interaction ΔI, with the resulting total $\Delta\widetilde{H}$. Panels b,d,f have individually labelled vertical scales. The first pair shows strict reversal, the second a conventional tie, and the third an additional inferior tie on the aligned grid. All examples have four-bit coefficients and allocation $\ell = 1$. State selection and complete-set classifications are specified in the text.

strict reversal means that G_C contains no acceptable state and has strictly lower compiled energy than every acceptable state. These events were classified separately for the exact physical optimum (numerical tolerance 10^{-10}) and for the original near-optimal set $F \leq F_* + 0.001L$ (comparison guard 10^{-12}). Global ground-state feasibility was classified separately as all feasible, mixed feasible/infeasible, or none feasible. Physical regret was averaged only over feasible global minima and remained missing when none existed. Full-grid averages describe all 26 allocation–precision choices per method; locked-policy results retain the original six selected precision-specific settings. Only the latter reproduce the primary policy comparison.

Examples were selected from the existing validation set before examining their explanatory profiles. Configurations were ordered by coefficient precision, instance identifier and integer allocation. The first conventional strict reversal and first conventional tie expansion at the near-optimal tolerance were retained, provided the global ground set contained a feasible state. The adverse aligned example was the first same-instance, same-precision, same-allocation pair with all global minima feasible and lower aligned fidelity. For each example the displayed witness has the largest physical energy among feasible-sector compiled minima; the smallest state identifier breaks a numerical tie. The reference is a true physical optimum with the lowest compiled energy, again breaking ties by state identifier. Outcome counts always use all states. The resulting examples are validation instance 10 at four bits and allocation one (strict reversal), and instance 0 at four bits and allocation one (conventional tie expansion and the adverse aligned comparison).

For the implementation audit, $Q_d(v)$ denotes the declared floating-point division followed by nearest-even rounding, multiplied by the fixed grid step d. The three penalty classes must be evaluated with this same operator. The exact-arithmetic specialization $\beta = -P/q$ at $d = 2P/q$ has half-integer arguments; floating-point division can place an argument on one side of that boundary. For validation instance 0, $P/d = 3.4999999999999996$, so the original implementation rounds it to three rather than the ideal nearest-even value four. Consequently the implemented β is 0.0281011, whereas half-boundary stabilization gives -0.00936704. In the full test grid, 108 of the 960 conventional conditions with $R \simeq 2P$ have this raw-boundary discrepancy; the locked subset has 48 discrepancies among 144 such conditions. None remains after the original 10^{-12} half-boundary stabilization. This changes the algebraic components and their interaction, not their exact sum.

The previously saved 1,152 stabilized locked-policy coefficient vectors were reclassified without changing the selected allocations. Their complete ground sets, fidelity values and both exact/near-optimal event counts agree with the original locked analysis. Thus the locked 48-versus-10 strict-reversal counts, 110-versus-164 tie-expansion counts and negative average fidelity comparison are unaffected by this existing implementation sensitivity. These are descriptive repeated-condition counts, not new independent statistical samples.

Physical content of the selected validation states

The physical reference states (r) and displayed compiled-minimum witnesses (w) have the following nodal values. A denotes the strict-reversal example; B and C denote conventional and aligned

compilation on the same six-site instance. Vectors follow the site order in Fig. S2.

$$\begin{aligned}
\boldsymbol{\phi}_{A,r} &= \left(1, \frac{1}{3}, -\frac{1}{3}, -1, -1, -\frac{1}{3}, \frac{1}{3}, 1\right), \\
\boldsymbol{\phi}_{A,w} &= \left(-1, -\frac{1}{3}, 1, 1, 1, \frac{1}{3}, -1, -1\right), \\
\boldsymbol{\phi}_{B,r} &= (1, -1, -1, -1, 1, 1), \\
\boldsymbol{\phi}_{B,w} &= (1, 1, -1, -1, -1, 1), \\
\boldsymbol{\phi}_{C,r} &= (1, -1, -1, -1, 1, 1), \\
\boldsymbol{\phi}_{C,w} &= (-1, -1, -1, 1, 1, 1),
\end{aligned}$$

The local, gradient and bias contributions are evaluated directly from the archived bias field v_i:

$$F_{\text{loc}} = \frac{h}{4}\sum_i (\phi_i^2 - 1)^2, \qquad F_{\text{grad}} = \frac{\kappa}{2h}\sum_i (\phi_{i+1} - \phi_i)^2, \qquad F_{\text{bias}} = h\sum_i v_i \phi_i,$$

where $\phi_{n+1} = \phi_1$ and $h = L/n$, with $L = 8$. In A, $n = 8$ and $\kappa = 0.5$; in B and C, $n = 6$ and $\kappa = 0.1$. The bias amplitude is 0.02 in all three examples. The occupancy tuple is $(N_{-1}, N_{-1/3}, N_{1/3}, N_1)$; $Y = 2(N_{-1} - N_1)$ on the feasible set. All six states have $\sum_i \phi_i = 0$ and $Y = 0$.

Supplementary Table S9: Physical energies and field-level occupancies of the selected validation states. Labels A, B and C refer to the reversal, conventional-tie and aligned-tie examples; subscripts r and w denote physical reference and displayed witness. The occupancies are ordered as $(N_{-1}, N_{-1/3}, N_{1/3}, N_1)$, the site counts at the four field levels; F_{loc}, F_{grad} and F_{bias} are the local-potential, gradient and bias energies. All quantities are dimensionless, and every state satisfies zero composition.

State	Occupancies	F_{loc}	F_{grad}	F_{bias}	F	$F - F_*$
A_r	(2,2,2,2)	0.790123	0.666667	-0.056004	1.400786	0
A_w	(3,1,1,3)	0.395062	1.111111	0.063636	1.569809	0.169023
B_r	(3,0,0,3)	0	0.300000	-0.071111	0.228889	0
B_w	(3,0,0,3)	0	0.300000	-0.035764	0.264236	0.035346
C_r	(3,0,0,3)	0	0.300000	-0.071111	0.228889	0
C_w	(3,0,0,3)	0	0.300000	-0.035346	0.264654	0.035764

In A, the witness lowers the local potential contribution but raises the gradient contribution; the bias term also becomes less favorable. Its net physical disadvantage is 0.169023, despite its lower compiled energy. In B and C, all occupied values are ± 1, so the local potential contribution vanishes and the gradient contribution is unchanged. Their physical energy differences come entirely from the archived bias field, which the low-precision compilation no longer resolves.

Coupling-level origin of the representative reversal

The original validation selection fixes instance 10, $n = 8$, $\kappa = 1/2$, bias amplitude $A = 0.02$, four bits and allocation $\ell = 1$. With $h = 1$, the largest physical coefficient has magnitude $M = 2/9$,

the penalty is $P = 4/27$, and conventional compilation uses $\delta = 8/189$. Let a_j and b_j denote physical and penalty coefficients, respectively, and define the integer interaction residual

$$r_j = \mathrm{rint}((a_j + b_j)/\delta) - \mathrm{rint}(a_j/\delta) - \mathrm{rint}(b_j/\delta). \tag{S47}$$

The coefficients in Table S10 account for all 16 fields and 120 distinct pairs. Rational nearest-even rounding gives the same residuals as the archived compiled vector. For example, the within-site arguments are $a_j/\delta = 35/12$, $b_j/\delta = 7/2$ and $(a_j + b_j)/\delta = 77/12$, so $r_j = 6 - 3 - 4 = -1$. The neighboring first-group pair gives $r_j = 0 - (-1) - 2 = -1$. All other classes give zero.

Supplementary Table S10: All coefficient classes in the fixed strict-reversal example. The indices 0 and 1 denote the first and second encoding spins at a site. Each distinct pair is counted once; neighboring sites include the periodic boundary. Fields have no penalty contribution, so their interaction vanishes regardless of the sampled bias. Here r_j is the interaction coefficient in units of $\delta = 8/189$.

Coefficient class	Count	a_j	b_j	r_j
Fields	16	$v_i w_\alpha/3$	0	0
Within-site pair	8	10/81	4/27	−1
Neighboring 00	8	−1/18	2/27	−1
Neighboring 01 or 10	16	−1/9	4/27	0
Neighboring 11	8	−2/9	8/27	0
Nonneighboring 00	20	0	2/27	0
Nonneighboring 01 or 10	40	0	4/27	0
Nonneighboring 11	20	0	8/27	0

Consequently, with periodic site indices,

$$I(s) = -\delta \left[\sum_i s_{i0} s_{i1} + \sum_i s_{i0} s_{i+1,0} \right]. \tag{S48}$$

The physical optimum and the selected inferior state have within-site sums 0 and 4, and neighboring first-group sums −4 and 4, respectively. Their interaction difference is therefore $-12\delta = -32/63$. It overcomes the positive physical gap 0.169023; adding the physical-rounding difference 0.000289 gives the compiled gap −0.338624. Both states have $Y = 0$, so the residual penalty contributes zero. The within-site product is +1 precisely at the saturated levels $\phi_i = \pm 1$, so this interaction favors the more saturated witness. The neighboring term favors alignment of the first encoding spins. These spins do not vary monotonically with ϕ_i; in this witness their alignment and the physical gradient energy both increase.

Exact accounting of the mean fidelity loss

For a saved condition c, let G_c be its complete global ground set, b_c the fraction of G_c that is feasible but outside the near-optimal tolerance, and u_c its infeasible fraction. Every state belongs to exactly one of the acceptable, inferior-feasible or infeasible classes, giving

$$1 - f_{E,c} = b_c + u_c. \tag{S49}$$

Let R_c and T_c indicate strict reversal and tie expansion of the feasible-sector ordering. If that ordering is preserved, $b_c = 0$. Otherwise the inferior feasible global minima belong to the corresponding reversal or tie-expansion condition. Thus

$$1 - \overline{f}_E = \overline{R_c b_c} + \overline{T_c b_c} + \overline{u_c}, \tag{S50}$$

where the bar weights the six precisions equally within each of the 96 test instances. Each identity was verified directly from the saved ground-state identifiers and independently recomputed physical energies. Table 2 of the main text reports the three contributions. Alignment reduces the reversal contribution by 6.597 percentage points, increases the tie-expansion contribution by 8.647 points and adds 0.093 points from infeasible global ties. Their sum is a 2.142-point increase in loss. This is descriptive accounting of the original policy comparison; the terms are not effects of independently manipulated interventions.

Eight-bit policy and same-allocation comparisons

The original development-stage calibration selected $\ell = 20$ for conventional rounding and $\ell = 10$ for alignment at eight bits. Table S11 reports that locked-policy comparison alongside the already saved comparisons at common allocations 10 and 20. No allocation was retuned. Every row averages the same 96 test instances and uses the original near-optimal tolerance. Alignment has higher fidelity in the locked-policy comparison, but lower fidelity at either common allocation. The policy comparison consequently measures the combined consequences of compilation and its separately calibrated allocation, rather than an isolated benefit of alignment. Even at common allocation, conventional and aligned physical grid steps can differ because their normalization rules differ.

Supplementary Table S11: Eight-bit fidelity under locked policies and common allocations. The allocations ℓ_C and ℓ_A belong to conventional and aligned compilation. The difference is aligned minus conventional mean f_E across 96 test instances. These are descriptive summaries of existing conditions.

Comparison	ℓ_C	ℓ_A	Conventional	Aligned	Difference
Common allocation 10	10	10	0.984375	0.977431	-0.006944
Common allocation 20	20	20	0.953125	0.904514	-0.048611
Locked policies	20	10	0.953125	0.977431	+0.024306

Physical size of the ground-state errors

Target-energy fidelity counts states at a fixed tolerance; physical regret describes the energy size of those errors. For each condition, the mean regret is the uniform mean of $(F_h - F_*)/L$ over its feasible global minima. Table S12 averages these condition means across the 96 instances at each precision. Its worst value is the maximum over all feasible global minima at that precision. All locked conditions contain at least one feasible global minimum, including the two aligned conditions with additional infeasible ties; hence no condition is omitted from these regret means. The infeasible states remain fully included in the separate fidelity loss above. These dimensionless quantities are calculated from saved states and do not represent device sampling frequencies.

The reproducibility package contains per-condition class counts, coefficient residuals and summary tables for the fixed instances and allocations.

Supplementary Table S12: Physical regret under the original locked policies. The mean gives equal weight to each physical instance after uniformly averaging its feasible global minima. The worst column is the largest individual-state regret across those same conditions. The values describe F_h on the four-level alphabet, independently of the continuum-reference comparison.

	Mean regret		Worst regret	
Bits	Conventional	Aligned	Conventional	Aligned
4	6.43×10^{-3}	4.39×10^{-2}	4.57×10^{-2}	3.12×10^{-1}
5	1.34×10^{-2}	2.18×10^{-3}	1.11×10^{-1}	2.23×10^{-2}
6	1.51×10^{-3}	9.19×10^{-4}	1.28×10^{-2}	1.50×10^{-2}
8	1.68×10^{-4}	1.15×10^{-4}	2.91×10^{-3}	2.47×10^{-3}
10	3.39×10^{-6}	8.41×10^{-6}	2.87×10^{-4}	4.24×10^{-4}
12	8.15×10^{-9}	1.24×10^{-6}	1.56×10^{-6}	1.13×10^{-4}

Coefficient perturbations and matched-grid control

The prespecified sensitivity experiment uses the locked allocations at six and eight bits. For each test instance, bit budget, and standard deviation $\sigma \in \{0.05, 0.25\}$, four coefficient vectors are drawn. The same standard-normal vector is used for both compilers in a paired draw. If k is the integer coefficient vector, the perturbed vector in grid units is $k + \sigma z$, with independent standard-normal coordinates in z. Both fields and all pair slots are perturbed, including slots whose intended coefficient is zero. Perturbations are continuous and are not clipped to the nominal signed range.

There are $96 \times 2 \times 2 \times 4 \times 2 = 3{,}072$ noisy compilations. Their complete ground sets are enumerated with the same 10^{-9} energy tolerance, now in continuous grid-unit energies. Common grid-unit noise does not imply common physical-unit noise: the latter has standard deviation $\sigma\delta$ or $\sigma R/q$, and the steps can differ between the selected policies. This sensitivity therefore probes resistance to additive coefficient errors measured in grid units. Gaussian errors are unbounded, so deterministic half-bin certificates are disabled for these rows.

The matched-grid control rounds the full penalized coefficient vector at the aligned step. Its coefficient vector is compared directly with the aligned construction, retaining any nearest-even half-tie discrepancies. This control distinguishes preserving the integer constraint block from an unreported change of coefficient step. The precise exception is proved in the Supplementary Note "Constraint-aligned integer compilation". The ideal algebraic identity of the aligned penalty is unaffected by those direct-rounding exceptions.

Post hoc half-grid implementation sensitivity

After the matched-grid differences were observed, a deterministic rounding sensitivity was added without changing the primary analysis or retuning allocations. For every scaled coefficient argument z, values within 10^{-12} of $\lfloor z \rfloor + 1/2$ were snapped to that exact half-grid value before nearest-even rounding. All $96 \times 6 \times 2 = 1{,}152$ locked-policy ground sets were re-enumerated. Integer coefficients changed in 73 configurations, but no complete ground set changed. Fidelity and feasibility therefore remained unchanged; the pooled fidelity difference and its interval were unchanged up to floating-point output precision.

The same stable half-grid convention removed all 372 original matched-grid exceptions across the

2,496 allocation–instance cases. All original exceptions lay near a half-grid location. No residual discrepancy remained in this benchmark after stabilizing the half-grid arithmetic. Thus the aligned construction is equivalent to joint rounding on its chosen grid in the stable matched-grid control. Its comparison with full-range conventional compilation concerns grid and penalty allocation, rather than an additional independent improvement at an identical stable grid. The original and stabilized records are both retained.

Independent continuum and representation audit

Why three comparisons are needed

Compilation fidelity compares the encoded objective with its exact four-state nodal target. A separate alphabet comparison allows continuous nodal values in the same nodal objective. A continuum comparison evaluates the spatially continuous functional and the exact continuum energy of a specified coarse interpolation. These comparisons involve different admissible sets and, for the nodal and continuum objectives, different energy evaluations. Their gaps should not be added or interpreted as a single universal error decomposition.

The periodic continuum functional is

$$\mathcal{F}[\phi] = \int_0^L \left[\frac{(\phi^2-1)^2}{4} + \frac{\kappa}{2}(\partial_x \phi)^2 + v(x)\phi \right] \mathrm{d}x, \qquad \int_0^L \phi \, \mathrm{d}x = 0. \tag{S51}$$

Eight diagnostic cases use $\kappa \in \{0.1, 0.5, 1.0, 1.8\}$ and $A \in \{0, 0.08\}$. One pair of phases, drawn using NumPy `default_rng(42)`, is fixed across these cases: $\theta_1 = 4.862909272689599$ and $\theta_2 = 2.757554564287996$. The forcing formula is the continuous counterpart of the benchmark's two-mode field. These cases are distinct from the independently randomized training, validation and held-out instances. Their prescribed settings were recorded before the continuum solver ran.

Fourier–Galerkin stationary references

The mean-zero field is represented as

$$\phi(x) = \sum_{k=1}^{K} \left[a_k \cos(2\pi k x/L) + b_k \sin(2\pi k x/L) \right]. \tag{S52}$$

Omitting the constant mode enforces mass exactly. Mode counts $K = 15, 31, 63$ use $N = 64, 128, 256$ periodic trapezoidal quadrature nodes. Since $4K < N$, the quartic potential of the represented field has no unresolved aliased mode in its integral. The forcing product is also resolved. The gradient contribution is evaluated analytically as

$$\frac{\kappa L}{4} \sum_{k=1}^{K} (2\pi k/L)^2 (a_k^2 + b_k^2). \tag{S53}$$

Thus the specified finite Fourier functional is evaluated exactly up to floating-point arithmetic; this statement is about quadrature, not global minimization.

The limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm with bounds (L-BFGS-B) [6] uses an analytic coefficient gradient, gradient tolerance 10^{-10}, relative objective tolerance 10^{-15}, at most 12,000 iterations, 50 line-search steps and 30 stored correction pairs. There are eight starts at each resolution: four first harmonics with amplitude 1.1 and phases $0, \pi/2, \pi, 3\pi/2$; two second harmonics with amplitude 0.9 and phases $0, \pi/2$; and two smooth random Fourier fields with coefficient scales decaying as $0.8/k^{3/2}$. Their random streams derive from $[42, 20261014, 3, 701, \text{start identifier}]$. At each finer resolution, the preceding best field is embedded unchanged as a ninth start. In total, $8(8 + 9 + 9) = 208$ optimizations are retained.

The lowest returned value is a numerical stationary reference. Apart from the explicit convexity case below, it is not a certified global continuum minimum. Every optimizer reported success, but an objective stopping criterion can trigger before the gradient criterion. Achieved residuals are therefore recorded independently. The chemical-potential residual is the root mean square of $\mu - \overline{\mu}$ on 4,096 nodes, where $\mu = \phi^3 - \phi - \kappa\phi_{xx} + v$. It includes cubic modes beyond the optimized Fourier space.

Across all starts, the largest mean error was 6.1583×10^{-17}; no increase between accepted objective iterates was recorded. The largest coefficient-gradient infinity norm was 7.4652×10^{-6}. Among the best $K = 63$ fields, the largest dense residual was 1.4681×10^{-7}. The largest absolute energy change from $K = 31$ to $K = 63$ was 9.6589×10^{-14}. These checks establish numerical resolution of the reported stationary references. Optimizer iterations are optimization steps rather than physical time, and unforced fields have translation degeneracy.

Exact continuum energy of a piecewise-linear coarse field

For adjacent values l and r, the exact potential integral on a cell is

$$h\left[\frac{l^4 + l^3 r + l^2 r^2 + l r^3 + r^4}{20} - \frac{l^2 + lr + r^2}{6} + \frac{1}{4}\right]. \tag{S54}$$

The gradient energy is $\kappa(r - l)^2/(2h)$. The product of the linear interpolant and each cosine forcing mode is integrated analytically. Together these expressions evaluate $\mathcal{F}$ of the periodic piecewise-linear interpolation without a quadrature approximation. They do not solve the continuum minimization problem.

At $n = 8$, every four-level configuration is enumerated and the 8,092 composition-feasible states are retained. Both the nodal objective and exact piecewise-linear continuum functional are evaluated on this same finite set. Ties within 10^{-11} are retained. Nodal and piecewise-linear minimizer sets intersect in all eight audit cases. The minimum piecewise-linear regret among nodal minimizers is zero in all eight cases. The data therefore do not demonstrate a quadrature-induced minimizer reversal. They show energy-value differences: the nodal value minus the exact interpolant energy ranges from -0.2802937252 to -0.0098765432.

Continuous nodal reference

The same $n = 8$ nodal objective is minimized under $-1 \le \phi_i \le 1$ and $\sum_i \phi_i = 0$, using sequential least squares programming (SLSQP) with exact gradients, objective tolerance 10^{-13} and at most 5,000 iterations. The initial eight starts are the zero field, four phase-shifted first harmonics of amplitude 0.9, two phase-shifted second harmonics of amplitude 0.9, and one zero-mean scaled random field.

The initial unforced $\kappa = 0.1$ calculation returned a best local value of 0.680082141, above the known feasible four-level value 0.4. A ninth start at the exact four-level optimum was therefore added uniformly to all eight cases. This post-initial-diagnostic amendment is recorded in `continuum/nodal_reference_amendment.json`. Initial results remain recoverable from the first eight starts in the raw table. The amendment does not change the primary compilation benchmark or the Fourier-reference settings. It prevents a missed local minimum from being misinterpreted as an alphabet advantage; it does not turn the nodal reference into a global certificate.

All 72 final nodal runs reported success. Their largest mean error was 5.5511×10^{-17} and largest box/mass Karush–Kuhn–Tucker (KKT) residual was 2.8236×10^{-7}. The residual fits the equality multiplier, uses absolute residuals on free nodes, and applies the appropriate one-sided conditions at active bounds. The Fourier field has no pointwise box bound, whereas this nodal reference does. Accordingly, a gap between the piecewise-linear and Fourier energies cannot be attributed solely to the finite alphabet or solely to spatial discretization.

Supplementary Table S13: Eight fixed cases for the representation comparison. Entries are dimensionless objective values. "Nodal cont." is the best returned box-constrained local reference; "P1" is the exact continuum energy of the piecewise-linear interpolation of the selected four-level nodal minimizer; "Fourier" is the best $K = 63$ stationary reference. These columns do not all minimize the same objective over the same admissible set.

κ	A	Four-level nodal	Nodal cont.	P1	Fourier
0.1	0	0.400000	0.362004	0.666667	0.596285
0.1	0.08	0.129401	0.098420	0.409694	0.333820
0.5	0	1.456790	1.250000	1.466667	1.327909
0.5	0.08	1.239029	1.028326	1.259871	1.111014
1.0	0	2.024691	1.763298	2.093827	1.801015
1.0	0.08	1.905695	1.595156	1.926538	1.639164
1.8	0	2.380247	2.000000	2.449383	2.000000
1.8	0.08	2.290047	1.949334	2.363725	1.961527

A certified uniform-state representation limit

For any mean-zero perturbation η, the continuum second variation obeys

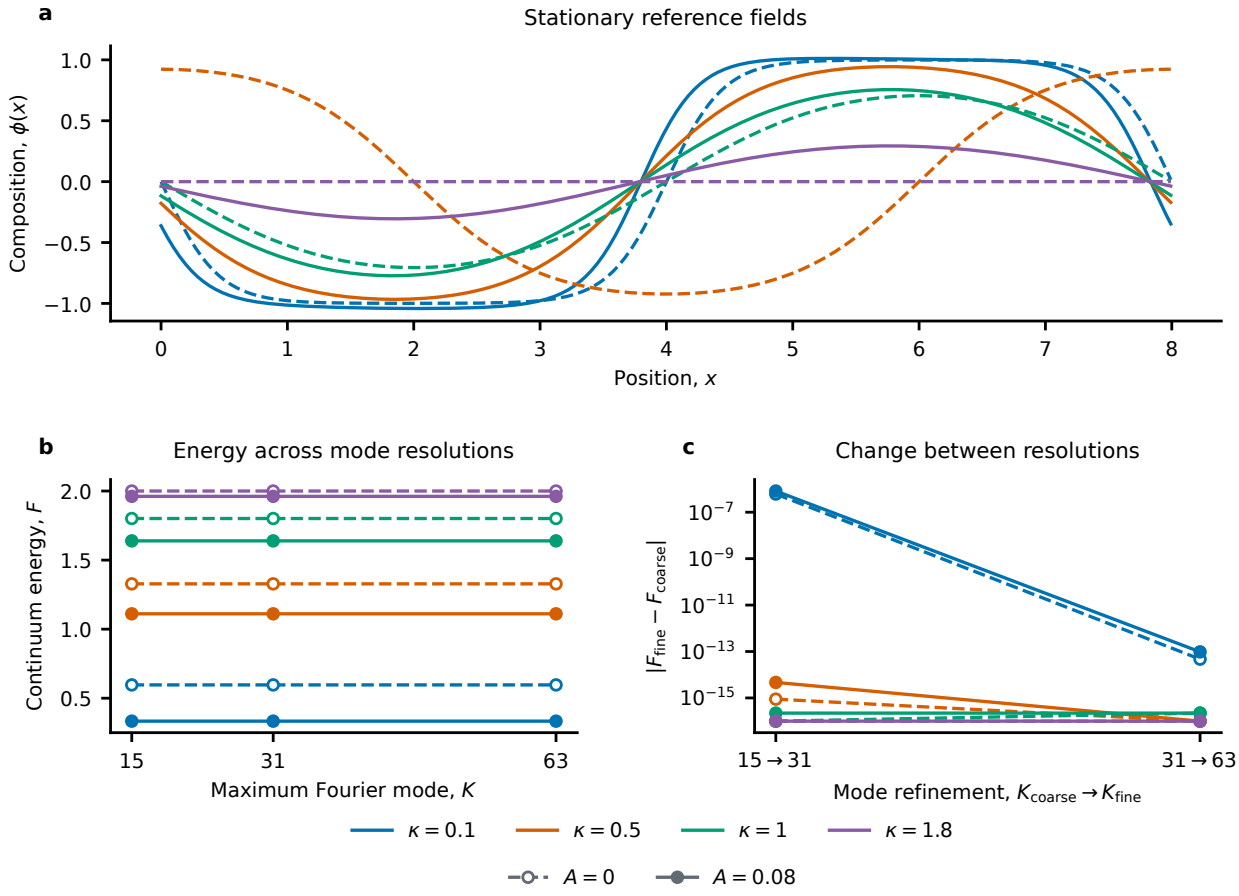
$$\delta^2\mathcal{F}[\phi](\eta,\eta) = \int_0^L [(3\phi^2-1)\eta^2 + \kappa\eta_x^2]\mathrm{d}x \tag{S55}$$

$$\geq [\kappa(2\pi/L)^2 - 1]\int_0^L \eta^2\mathrm{d}x. \tag{S56}$$

Poincare's inequality therefore gives strict convexity at fixed mean for $\kappa > (L/2\pi)^2 = 1.621138938$. With $A = 0$, the unique minimum is $\phi = 0$ and $\mathcal{F} = L/4 = 2$. The corresponding nodal threshold at $n = 8$ is $h^2/[4\sin^2(\pi/n)] = 1.707106781$, so the continuous nodal objective has the same unique minimum at $\kappa = 1.8$.

The four-level alphabet excludes zero. Its exact nodal minimum is 2.380246914, and the exact continuum energy of its periodic piecewise-linear interpolation is 2.449382716. Thus faithful encoding and exact optimization of the finite-state model still cannot represent this certified uni-

form solution. This is an alphabet limitation in a finite periodic model, not a thermodynamic phase-transition result.



Supplementary Figure S3: Independent stationary references and spectral resolution. **a**, Best returned $K = 63$ fields for the eight fixed audit cases. Colours identify gradient strength κ; dashed lines and open markers indicate bias amplitude $A = 0$, and solid lines and filled markers indicate $A = 0.08$. Unforced translations are arbitrary. **b**, Lowest returned functional at maximum Fourier modes $K = 15, 31, 63$. **c**, Absolute energy changes for mode refinements $15 \to 31$ and $31 \to 63$; values below 10^{-16} are displayed at 10^{-16} only for the logarithmic plot. The corresponding $N = 64, 128, 256$ quadrature nodes exceed $4K$ and exactly integrate the represented quartic. Each level uses eight initial fields, with an additional prolonged preceding best at the two finer levels. The reference is numerical and is not a general global-optimum certificate.

Reproducibility, independent checks and retained records

The source, model protocols, coefficient vectors, complete ground-state identifier arrays and comma-separated values (CSV) summaries are provided in the accompanying reproducibility material. The manuscript source package additionally contains the current figure-rendering code and compact display inputs, which reuse these archived values. The separate archive `Reproducibility_Package.zip` retains the original numerical research material under `src/`, the

mechanism-analysis code and results under `scripts/` and `analysis/`, and a file manifest. Its current guide and package metadata supersede historical administrative records. Its reproduction guide distinguishes regenerating analysis from compiling the lightweight manuscript source package. Each primary instance retains its forcing phases and samples, unpenalized Ising coefficients, constant, exact physical ground-state identifiers, all compiled coefficient columns, and every encoded ground-state set. No compiled instance is omitted on the basis of performance. Continuous-reference records include every local-optimization result, coefficients of selected fields, and accepted objective histories.

The independent theory script checks the identities without importing the benchmark implementation. The additional training archive audit reconstructs all 16 training instances, 832 compiled ground sets and 3,167 ground-state entries. It finds zero integer-coefficient differences, a maximum physical mapping discrepancy of 7.11×10^{-15}, and a maximum primary-metric discrepancy of 4.44×10^{-16}. All 105 training compilations marked certified have feasible, physically optimal ground sets. The locked choices are recovered from training and validation outcomes only. All 128 instance seed streams are distinct.

A separate independent audit verifies all 96 held-out archives and 4,992 compiled ground sets, including 19,010 ground-state entries counted across configurations. All 555 certified configurations satisfy their claimed strict physical-ground-set property, with no false positive found. The largest physical mapping error is 7.11×10^{-15} and the largest metric discrepancy is 4.44×10^{-16}. All 96 test phase pairs are distinct. This audit checks the complete finite-state archive.

Six independent continuum checks cover analytic Fourier gradients, nested quadrature invariance, exact piecewise-linear energies against 30-point Gauss–Legendre cell integration, the uniform energy, exact mean fixing, and forcing normalization. The largest centered finite- difference Fourier-gradient discrepancy is 2.4927×10^{-10}; the largest independent piecewise-linear integration discrepancy is 8.8818×10^{-16}. Quadrature refinement has zero discrepancy at the tested floating-point precision.

The original numerical environment is Python 3.12.14, NumPy 2.3.5, SciPy 1.18.1, pandas 3.0.1, matplotlib 3.11.2, and threadpoolctl 3.7.0. Basic Linear Algebra Subprograms (BLAS) are limited to one thread within each benchmark worker. Independent instances are evaluated in separate processes; process scheduling does not determine their seeds. The reproducibility package records the software environment and executed analysis versions. The current figures use Matplotlib 3.10.7 (`https://matplotlib.org/`); their rendering code and saved display inputs are provided in the separate manuscript source package.

Interpretation of the comparisons

The benchmark concerns a one-dimensional periodic model with four local values and complete finite-state enumeration. Ground-set fractions summarize logical minimizers. The prescribed coefficient grid and Gaussian perturbations isolate numerical effects rather than a particular hardware calibration.

The selected-policy comparison measures the combined effect of allocation selection and compilation. Same-allocation comparisons hold the effective penalty fixed; matched-grid comparisons additionally identify whether the integer construction differs from direct rounding. The state-level account then separates composition errors, strict feasible-state reversals and tie expansion. To-

gether these comparisons explain how enforcing the constraint and resolving physical energy differences can compete for coefficient range.

The continuum references show that nodal and piecewise-linear energies share minimizers in the eight cases even though their values differ. The certified uniform-state example establishes that an exactly solved four-level target can exclude the continuum minimum.

References

[1] Cahn, J. W. & Hilliard, J. E. Free Energy of a Nonuniform System. I. Interfacial Free Energy. *J. Chem. Phys.* **28**, 258–267 (1958). URL `https://doi.org/10.1063/1.1744102`.

[2] Chen, L. Q. & Shen, J. Applications of semi-implicit Fourier-spectral method to phase field equations. *Comput. Phys. Commun.* **108**, 147–158 (1998). URL `https://doi.org/10.1016/s0010-4655(97)00115-x`.

[3] Verma, A. & Lewis, M. Penalty and partitioning techniques to improve performance of QUBO solvers. *Discrete Optim.* **44**, 100594 (2022). URL `https://doi.org/10.1016/j.disopt.2020.100594`.

[4] Mücke, S., Gerlach, T. & Piatkowski, N. Optimum-preserving QUBO parameter compression. *Quantum Mach. Intell.* **7**, 1 (2025). URL `https://doi.org/10.1007/s42484-024-00219-3`.

[5] Alessandroni, E., Ramos-Calderer, S., Roth, I., Traversi, E. & Aolita, L. Alleviating the quantum Big-M problem. *npj Quantum Inf.* **11**, 125 (2025). URL `https://doi.org/10.1038/s41534-025-01067-0`.

[6] Byrd, R. H., Lu, P., Nocedal, J. & Zhu, C. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM J. Sci. Comput.* **16**, 1190–1208 (1995). URL `https://doi.org/10.1137/0916069`.