# The spectral gap is the wrong object: observable-projected gradient signal-to-noise and a reversibility–mixing obstruction for thermodynamic energy-based models

Prasanna*

2026

**Abstract**

Training an energy-based model (EBM), and in particular a Denoising Thermodynamic Model (DTM) on a thermodynamic sampler, hinges on the recoverability of the negative-phase gradient from a finite Markov-chain Monte Carlo (MCMC) estimate. We study the squared gradient signal-to-noise ratio $Q_{\rm op} := \|g\|^2/\mathbb{E}\|\hat{g} - g\|^2$, whose collapse below $O(1)$ is the training plateau observed in practice. We make three contributions, each grounded in controlled exact-diagonalization and at-scale experiments. First, the textbook predictor built on the spectral gap $\gamma = 1 - \sigma_2$ is the *wrong object*: under a $\mathbb{Z}_2$ spin-flip symmetry the slowest mode is observable-orthogonal to the even gradient observables (measured overlap $\leq 3.5 \times 10^{-17}$ for an RBM, $\sim 10^{-24}$ to $10^{-29}$ for Ising), so a single-$\gamma$ ratio over-predicts $Q_{\rm op}$ by $10^{26}$ to $10^{30}$ — a divide-by-symmetry-zero, repaired by an $L^2(\pi_\theta)$ function-subspace projection rather than state-space conditioning. Second, we define the observable-projected, multi-mode predictor $Q^{\perp}_{\rm struct}$ via an observable-relevant mode set, an aggregate timescale, and a harmonic-mean effective gap; it tracks $Q_{\rm op}$ in 45/48 controlled cells and in 92–99% of cells across model sizes $m = 4$ to 16 under two kernels. Third, we expose an A2↔A6 obstruction: the reversibility the factorization requires antagonizes the mixing it needs, manifesting as an effective-gap collapse ($\sim 100\times$), a dead-linear $\tau \propto L$ autocorrelation growth at scale, and a thermodynamic-length cost wall (136 rungs needed against a 96-rung budget). We preserve a strict claim-status discipline: a *conditional* factorization is [solid], the *operational* claim stays [conjectured], and the obstruction's fundamentality remains open.

## 1 Introduction and setup

A Denoising Thermodynamic Model (DTM) [7] stacks reverse-process energy-based layers, each sampled by a thermodynamic Gibbs kernel. Training one layer reduces to the classical restricted-Boltzmann-machine (RBM) maximum-likelihood problem [1]: the gradient of the log-likelihood is a difference of a data-phase and a model-phase expectation, and the model phase must be estimated by MCMC. Both contrastive divergence [1] and its persistent variant [2] replace the intractable equilibrium model phase by a short-chain surrogate, and the practical pain — the part that controls whether training descends at all — lives in how well that finite chain has mixed. The DTM paper diagnoses this with the negative-phase autocorrelation and a plateau in generation quality [7]. The question this paper formalizes and measures is: *which spectral quantity predicts that plateau, and can it ever be reached on the reversible sampler the analysis requires?*

*Correspondence: `prasa.sanna9080@gmail.com`.

### 1.1 The model and the gradient

Take one reverse-process EBM layer with energy $E_\theta$ on configurations $x \in \{-1,+1\}^N$, model law

$$\pi_\theta(x) = \frac{e^{-E_\theta(x)}}{Z_\theta}, \tag{1}$$

sampled by a $\pi_\theta$-reversible Gibbs kernel $P_\theta$. For a parameter $a$ write the per-parameter observable $f_a(x) := \partial_a E_\theta(x)$. The training-gradient component is the data−model difference

$$g_a := \mathbb{E}_{\text{data}}[f_a] - \mathbb{E}_{\pi_\theta}[f_a], \tag{2}$$

the descendant of the RBM identity $\langle s_i s_j\rangle_{\text{data}} - \langle s_i s_j\rangle_{\text{model}}$ of Hinton [1], here with $f_a = \partial_{J_{ij}} E_\theta = -x_i x_j$ for a pairwise coupling. The data phase is cheap; *all the mixing cost is in the model (negative) phase.*

### 1.2 The canonical estimator and the quantity $Q$

We estimate the model phase with a burn-in $B$ plus a time-average window $K$: run $P_\theta$ from a start law $\psi_0$, discard the first $B$ steps, and average $f_a$ over the next $K$,

$$\bar{f}_a := \frac{1}{K}\sum_{j=B+1}^{B+K} f_a(x[j]), \qquad \hat{g}_a := \mathbb{E}_{\text{data}}[f_a] - \bar{f}_a. \tag{3}$$

Pinning $B$ and $K$ *separately* keeps the bias and variance legs describing the same estimator. The object we predict is the squared gradient signal-to-noise ratio:

$$\boxed{Q_{\text{op}} := \frac{\|g\|^2}{\mathbb{E}\|\hat{g} - g\|^2}} \qquad \text{(true-gradient power/estimator MSE).} \tag{4}$$

When $Q_{\text{op}} \gg 1$ the estimated direction is reliable and SGD descends; when $Q_{\text{op}} \lesssim 1$ the estimator MSE swamps the signal — the training plateau. **This is an estimation plateau** (the recoverability of $g$ collapses), *not* a signal-extinction plateau (the true gradient itself does not vanish). The distinction from the quantum barren plateau is taken up in Section 5.

### 1.3 Spectral ingredients

For the reversible kernel $P_\theta$, standard reversible-MCMC theory [4] gives a real $L^2(\pi_\theta)$-orthonormal eigenbasis $\varphi_1 = 1, \varphi_2, \dots$ with eigenvalues $1 = \sigma_1 > \sigma_2 \geq \cdots \geq \sigma_d \geq 0$ and spectral gap $\gamma := 1 - \sigma_2$. Expand each observable $f_a = \sum_j \hat{f}_{a,j}\varphi_j$ with $\hat{f}_{a,j} := \langle f_a, \varphi_j\rangle_\pi$, and define the slow-mode weight $w_a := \hat{f}_{a,2}^2 / \operatorname{Var}_\pi[f_a] \in [0,1]$. Then for the $B+K$ estimator the bias decays geometrically in $B$ and the window-$K$ variance scales as [4, 3]

$$\mathbb{E}[\bar{f}_a] - \mathbb{E}_\pi[f_a] \asymp (1-\gamma)^B\, \hat{m}_2\, \hat{f}_{a,2}, \qquad \operatorname{Var}[\hat{g}_a] \approx \frac{2w_a}{\gamma K}\operatorname{Var}_\pi[f_a], \tag{5}$$

where $\hat{m}_2$ is the overlap of the initial error with $\varphi_2$. The effective sample size is $\text{ESS}_a \approx \gamma K/(2w_a)$, scaling as $\gamma K$. The variance leg's constant $2w_a$ is a single-isolated-slow-mode refinement of the worst-case Levin–Peres constant; it presumes one observable-relevant slow mode. Section 2 shows that presumption fails — both mis-anchored and insufficient.

In the mixed regime (burn-in adequate, $K$ large) the historical conjecture was that $Q_{\rm op}$ factorizes as

$$Q_{\rm struct}(\theta, K) := \frac{\gamma_\theta K}{2}\, R(\theta), \qquad R(\theta) := \frac{\|g\|^2}{\sum_a w_a \,{\rm Var}_\pi[f_a]}, \tag{6}$$

i.e. $Q_{\rm op} \approx Q_{\rm struct} = {\rm ESS} \times$ (signal-to-slow-fluctuation ratio). This is computable from the kernel spectrum without training to convergence, and differentiable in the couplings $J$. We now show that Eq. (6), as written, is mis-specified.

# 2 Result 1: the spectral gap is the wrong object

## 2.1 The divide-by-symmetry-zero

For the canonical $b = 0$ pairwise-Ising and RBM energies the model possesses a global $\mathbb{Z}_2$ spin-flip symmetry $x \mapsto -x$. The Hilbert space $L^2(\pi_\theta)$ then splits into parity sectors $H = H^+ \oplus H^-$, with even functions $\{f : f(-x) = f(x)\}$ and odd functions. The pairwise gradient observables $f_{ij} = -x_i x_j$ are *even*, since $f_{ij}(-x) = -(-x_i)(-x_j) = -x_i x_j = f_{ij}(x)$. The slowest non-trivial mode $\varphi_2$, by contrast, is *odd*.

**Proposition 1** (Observable orthogonality)**.** *If $\pi_\theta$ and $P_\theta$ are $\mathbb{Z}_2$-invariant and $f_a \in H^+$, then $\hat{f}_{a,j} = \langle f_a, \varphi_j\rangle_\pi = 0$ for every odd eigenfunction $\varphi_j \in H^-$. In particular the odd slowest mode satisfies $\hat{f}_{a,2} = 0$ exactly.*

*Proof.* The integrand $f_a \cdot \varphi_j$ is odd under $x \mapsto -x$ (even times odd), and $\pi_\theta$ is even, so the sum over the $\pm x$ pairs cancels identically. The proof is the isotypic decomposition of $L^2(\pi_\theta)$ under the $\mathbb{Z}_2$ unitary $(U_s f)(x) = f(-x)$, which commutes with $P_\theta$ by equivariance; see [14]. □

The consequence is fatal for Eq. (6). The denominator of $R(\theta)$ is $\sum_a w_a \,{\rm Var}_\pi[f_a] = \sum_a \hat{f}_{a,2}^2$, which is *exactly zero* when $\varphi_2$ is observable-orthogonal: $R$ is ill-defined, a *divide-by-symmetry-zero* rather than a finite mis-estimate. Empirically, the slowest-mode observable overlap $\sum_a \hat{f}_{a,2}^2$ sits at machine zero — measured at $\leq 3.5 \times 10^{-17}$ for an RBM under both single-site and block-Gibbs kernels [9], and at $\sim 2 \times 10^{-29}$ (weak coupling) to $\sim 6 \times 10^{-24}$ (strong coupling) for the Ising family [8], with an independent exact-diagonalization check measuring $\sim 10^{-20}$ to $10^{-23}$ on the bipartite RBM [10]. With this floor the single-$\gamma$ predictor over-predicts $Q_{\rm op}$ by

$$Q_{\rm struct}/Q_{\rm op} \;\sim\; 10^{26} \text{ to } 10^{30} \tag{7}$$

in every cluster cell tested (exact-diagonalization, $N \leq 14$ controlled families) [8]. The naive single-$\gamma$ predictor anchored to $\sigma_2$ tracks $Q_{\rm op}$ (median ratio in $[\frac{1}{3}, 3]$) in 0 **of** 48 symmetric ($b = 0$) cells [8]. When the symmetry is broken ($b \neq 0$) the naive predictor becomes well-defined and over-predicts by a finite $\sim 8$ to $9\times$ in cluster cells — the clean, finite mechanism the original conjecture envisioned, recovered only once the divide-by-zero is removed.

## 2.2 The fix is an $L^2(\pi)$-projection, not state-space conditioning

The repair is a function-subspace projection, not a restriction of the chain to a sub-region of state space. Let $\mathcal{C}^\star(\mathcal{O}) := \{j \geq 2 : \sum_a \hat{f}_{a,j}^2 > \varepsilon \sum_a {\rm Var}_\pi[f_a]\}$ be the observable-relevant mode set, and $\Pi_C$ the $L^2(\pi_\theta)$-orthogonal projection onto $V_C := \bigoplus_{j\in\mathcal{C}^\star(\mathcal{O})} {\rm span}(\varphi_j)$. We have proved the following, in the regime A1–A4 plus a symmetry-compatibility assumption [14].

**Theorem 1** (Projection $\neq$ conditioning; [proven-here] for O1.c)**.** *$\Pi_C$ is an orthogonal projection in $H = L^2(\pi_\theta)$, with $\sum_a \|\Pi_C f_a\|_\pi^2 = \sum_a \sum_{j\in\mathcal{C}^\star(\mathcal{O})} \hat{f}_{a,j}^2$. It is* not *the operator of conditioning $\pi_\theta$ on any state-space event: if $V_C$ equalled a support subspace $H_S = \{f : \operatorname{supp} f \subseteq S\}$, then $S$ and $S^c$ would both be absorbing, contradicting irreducibility for any nontrivial $S$. Moreover the* stationary *signal-to-noise objects ($T_\mathcal{O}, \mathrm{Var}_\pi, \tau_{\mathrm{int}}, \|g\|^2$, hence $Q^\perp_{\mathrm{struct}}$) are invariant to the choice between the $H$-projection form and a fundamental-domain quotient, and the $H$-form is canonical.*

A $\mathbb{Z}_2$ "even sector" is doubly malformed as a conditioning event: $H^+$ is a function-parity sector, not a support subspace, and the free flip $x \mapsto -x$ has no fixed point, so "even configurations" is not even a well-defined set. The slow odd mode is therefore annihilated by $L^2(\pi_\theta)$-orthogonality, not by any conditioning. The O1.c invariance — that the stationary predictor is unchanged whether one projects in function space or runs a fundamental-domain quotient chain — is the load-bearing new lemma, and is the first [proven-here] result in this program [14]. The empirical overlaps at machine zero corroborate the projection *mechanism*, but the fundamental-domain comparison was never run as a chain, so the invariance rests on the written proof, never on a [not validated] flip.

# 3 Result 2: the observable-projected predictor $Q^\perp_{\mathrm{struct}}$

## 3.1 Construction

We replace each ingredient of Eq. (6) by its observable-projected, multi-mode form.

**Definition 1** (Observable-relevant ingredients)**.** With $\mu_C := \sum_a \sum_{j\in\mathcal{C}^\star(\mathcal{O})} \hat{f}_{a,j}^2$ the cluster overlap mass, define

$$\text{(1) mode set:} \quad \mathcal{C}^\star(\mathcal{O}) = \Big\{ j \geq 2 : \textstyle\sum_a \hat{f}_{a,j}^2 > \varepsilon \sum_a \mathrm{Var}_\pi[f_a] \Big\}, \tag{8}$$

$$\text{(2) aggregate timescale:} \quad T_\mathcal{O} := \sum_a \tau_{\mathrm{int}}[f_a]\ \mathrm{Var}_\pi[f_a] = \sum_a \sum_{j\geq 2} \hat{f}_{a,j}^2 \frac{1+\sigma_j}{2(1-\sigma_j)}, \tag{9}$$

$$\text{(3) effective gap:} \quad \gamma_{\mathrm{eff}} := \frac{\mu_C}{\sum_a \sum_{j\in\mathcal{C}^\star(\mathcal{O})} \hat{f}_{a,j}^2/(1-\sigma_j)}, \tag{10}$$

where $\tau_{\mathrm{int}}$ uses the half-Sokal convention and $\gamma_{\mathrm{eff}}$ is the overlap-weighted harmonic mean of the cluster gaps $\{1-\sigma_j\}$. The corrected predictor and claim are

$$\boxed{Q^\perp_{\mathrm{struct}}(\theta, K) := \frac{K}{2}\frac{\|g\|^2}{T_\mathcal{O}} = \frac{\gamma_{\mathrm{eff}} K}{2} R_{\mathrm{eff}}}, \qquad R_{\mathrm{eff}} := \frac{\|g\|^2}{\mu_C}, \qquad \text{claim: } Q_{\mathrm{op}} \approx Q^\perp_{\mathrm{struct}}. \tag{11}$$

The aggregate timescale $T_\mathcal{O}$ is exactly the half-aggregate MCMC-CLT asymptotic variance of the gradient observables in the sense of Younes [3], and it is directly measurable at scale from the autocorrelation of $f_a$ (which, being even, already projects onto $\mathcal{C}^\star(\mathcal{O})$). The harmonic-mean reduction is exact: using $(1+\sigma_j)/(2(1-\sigma_j)) = 1/(1-\sigma_j) - \frac{1}{2}$, the cluster term satisfies $T^C_\mathcal{O} = (\mu_C/\gamma_{\mathrm{eff}})(1-\gamma_{\mathrm{eff}}/2)$, which in the plateau regime $\gamma_{\mathrm{eff}} \to 0$ collapses to $\mu_C/\gamma_{\mathrm{eff}} \cdot (1+o(1))$; a single isolated relevant mode recovers $\gamma_{\mathrm{eff}} = 1-\sigma_{j^*}$ [15]. The predictor remains computable without training to convergence and differentiable in $J$ (given a smooth surrogate for the $\mathcal{C}^\star(\mathcal{O})$-selection), and $\gamma_{\mathrm{eff}}$ is smoother than $\sigma_2$ across within-cluster eigenvalue crossings [15].

### 3.2 Tracking evidence

**Exact diagonalization (exp1).** [8] On $N \leq 14$ controlled Ising/Hopfield families with the chain estimator validated against exact bias-plus-window MSE to within 0.7 to 3.3%, a *single* observable-relevant gap tracks $Q_{\mathrm{op}}$ in 23/48 symmetric cells, while the *multi-mode cluster correction* of Eq. (11) tracks in 45/48. Both fixes — the projection and the multi-mode aggregation — are therefore necessary; one without the other fails.

**At scale, two kernels (exp2).** [9] On bipartite RBMs the observable-relevant $Q^{\perp}_{\mathrm{struct}}$ tracks $Q_{\mathrm{op}}$ in 92 to 99% of cells at every size $m \in \{4, 6, 8, 12, 16\}$, including $m \geq 8$ beyond the spectrum-overlap regime, under both single-site and 2-block-Gibbs kernels — it extrapolates. The MC pipeline reproduces the exact joint-eigendecomposition timescale to a median 2.2% relative error, and an independent install of the Extropic `thrml` library [7] reproduces the self-contained sampler's integrated autocorrelation to 1.4%. By contrast the *naive* odd-scalar diagnostic (the un-projected autocorrelation $r_{yy}$ of a symmetry-odd scalar such as the magnetization) tracks 90% of cells when $\tau_{\mathrm{int}}[\mathrm{odd}]/\tau_{\mathrm{int}}[f] < 1.5$ but only 14% when that ratio exceeds 2.5 — it fails precisely in the symmetry-separated regime, the at-scale analog of the Result-1 wrong-mode failure. This sharpens a practical warning: the DTM autocorrelation diagnostic predicts the gradient SNR only when measured on the *even* gradient observables.

### 3.3 Claim status: conditional [solid] versus operational [conjectured]

We are careful to state two distinct claims with two distinct statuses.

**Theorem 2** (Conditional factorization; [solid])**.** *In the regime A1–A8 plus the plateau condition $\gamma_{\mathrm{eff}} \to 0$ plus the positive-phase-subdominance condition F4,*

$$Q_{\mathrm{op}} = Q^{\perp}_{\mathrm{struct}}\,(1 + o(1)). \tag{12}$$

*This is a written assembly over six obligations: O1.c (projection invariance, [proven-here]) and O2–O6 (each [solid], textbook reversible-MCMC and spectral algebra closing to [4, 3]).*

The assembly is: O2 grounds $T_{\mathcal{O}}$ as the MCMC-CLT asymptotic variance; O3 shows the finite-$B$ bias is $O(1/\mathrm{ESS})$-subdominant (so the binding gate is $K \gg \tau_{\mathrm{int}}$, not burn-in); O4 proves the projection commutes with the dynamics, $[\Pi_C, P_\theta] = 0$; O5 reduces $T_{\mathcal{O}}$ to the harmonic-mean $\gamma_{\mathrm{eff}}$ [15]; O6 closes the Euclidean aggregation $\mathbb{E}\|\hat{g} - g\|^2 = \mathrm{tr}(\hat{\Sigma}) + \|b\|^2$, the off-diagonal covariance being trace-invisible.

*Remark* 1 (The operational claim stays [conjectured]). The *unconditional* statement "$Q_{\mathrm{op}} \approx Q^{\perp}_{\mathrm{struct}}$ on a real DTM at the $K$ one runs" is [conjectured]. It is gated on two regime preconditions that are open at scale: A7 (the observable-overlapping bulk relaxes at an $\Omega(1)$ rate) and $K \gg \tau_{\mathrm{int}}$. The conditional theorem is vacuous on the plateau of a real DTM until those gates are met. Supporting evidence for the operational direction is therefore *construction-confirmation on small and moderate controlled models* (exp1, exp2), not [not validated] validation; neither tier is validated.

## 4 Result 3: the reversibility–mixing obstruction

The factorization needs $\pi$-reversibility (assumption A2) for the spectral machinery and $K \gg \tau_{\mathrm{int}}$ (assumption A6) for the window. These are logically independent — nothing forbids a reversible kernel that also mixes fast — yet we find them *operationally antagonistic* exactly in the multimodal plateau regime the theorem is about. We document the antagonism at three scales.

### 4.1 Mechanism: effective-gap collapse on a controlled family (exp5)

On a controlled bipartite-RBM family with $\gamma_{\text{eff}}$ and $\tau_{\text{int}}$ computed *exactly from the spectrum and resolvent* (no chain, hence no possible burn-in artifact), the reversible kernel's observable-relevant gap collapses 1 to 2 orders of magnitude as planted multimodality $M$ turns on [10]. Representative values at $\beta = 3$:

$$\gamma_{\text{eff}}(M{=}1) = 0.648 \;\longrightarrow\; \gamma_{\text{eff}}(M{=}2) = 0.006, \qquad \text{a} \sim 100\times \text{ collapse}, \tag{13}$$

with the raw observable gap $\gamma_{\perp}$ collapsing identically ($0.577 \to 0.006$), confirming the effect is physical and not a clustering artifact. Being exact-from-spectrum, the collapse *cannot* be a burn-in artifact; it favors the reading that the plateau is a genuine $\gamma_{\text{eff}} \to 0$ effect. Separately, a cross-kernel test shows that *reversibilization itself* adds a mixing penalty that grows with multimodality: the reversible symmetrization $P_{\text{sym}} = \frac{1}{2}(P_{\text{fwd}} + P^{*}_{\text{fwd}})$ is slower than the non-reversible $P_{\text{fwd}}$ in 16/16 cells (integrated-autocorrelation ratio 1.001 to 1.284, monotone increasing in $M$, resolvent-derived with a $6 \times 10^{-14}$ exact-diagonalization cross-check). The honest ceiling: at this small scale the antagonism is *bounded* — $\gamma_{\text{eff}}$ bottoms at $\sim 0.006$ (not 0), the window sweet spot stays marginally open (ESS at $K = 1000$ is $\approx 3$), and the reversibility penalty is 1 to 28% (not unbounded). So exp5 is mechanism-confirming, not a proof that the trained DTM sits at the same depth.

### 4.2 At scale: $\tau \propto L$ never equilibrates (exp4, exp6)

On the real 60×12 MNIST DTM single-input conditional, with the A2-required reversible kernel (forward/reverse symmetrized 4-block Gibbs, self-adjointness-gated), the integrated autocorrelation time does not stabilize: a non-circular doubling-stability probe finds $\tau_{\max}$ growing *dead-linearly in the trajectory length* $L$ (Table 1). The ratio $\tau/L \approx 0.16$ is essentially constant over six doublings, so the autocorrelation function never decays within the window: the integrated $\tau_{\text{int}}$ accumulates as fast as data is added, signalling a near-zero gap or an effectively non-equilibrating chain [11]. The doubling-stability criterion is never met, $\hat{\tau}$ is UNRESOLVED, and the $\gtrsim 50\,\hat{\tau}$ windows the analysis needs are uninstantiable: **A6 ($K \gg \tau_{\text{int}}$) is unreachable at this checkpoint.**

| $L$ (sweeps) | $\tau_{\max}$ | $\tau/L$ |
|---:|---:|---:|
| 1,000 | 166.7 | 0.167 |
| 2,000 | 333.2 | 0.167 |
| 4,000 | 662.6 | 0.166 |
| 8,000 | 1,325 | 0.166 |
| 16,000 | 2,431 | 0.152 |
| 32,000 | **5,280** | 0.165 |

Table 1: Reversible-kernel autocorrelation does not stabilize: $\tau_{\max} \propto L$ on the trained 60×12 MNIST DTM ($t = 200$ checkpoint), $\tau/L \approx 0.16$ over six doublings [11].

An earlier-checkpoint sweep [12] over $t \in \{25, 50, 100, 200\}$ (one cumulative trajectory, provenance proven by opt_count $= t \times 61$ and probe-RNG isolation) finds the same $\tau \propto L$ ($\tau/L \approx 0.13$ to 0.17 out to $L = 64{,}000$) at $t \in \{50, 100, 200\}$, with $t = 200$ reproducing exp4. A single sub-tolerance step at $t = 25$ rule-resolved a finite $\hat{\tau} \approx 2094$, but the audit finds this a fragile artifact (probed one doubling shorter than the others), so the conservative reading is "slow from the start": the antagonism is present from very early training, not only at convergence. Even taken at face value, the only finite-$\hat{\tau}$ checkpoint is the least-trained, barely-useful model, and its $50\,\hat{\tau} \approx 1.05 \times 10^5$ window underscores an A6-versus-utility tension.

### 4.3 The thermodynamic-length cost wall (exp19)

A natural escape is to accelerate the reversible negative phase with parallel tempering. We tested an equal-acceptance reversible-PT ladder on the same trained $t = 200$ DTM [13]. A hotter top *does* cure the decorrelation — a temperature effect, not a local-kernel non-ergodicity — with the single-replica reversible kernel reaching $\tau \le 25$ only at the deep top temperature $\alpha^\star_{\text{top}} = 0.02$ ($\tau = 15.8$; at $\alpha = 0.03$ already $\tau = 79$). But the cost of a single-span equal-acceptance ladder is set by the *thermodynamic length* between the free-mixing hot regime and the cold target,

$$R^\star(\alpha_{\text{top}}) = 1 + \text{round}\left(\frac{\beta \int_{\alpha_{\text{top}}}^{1} \sqrt{C}\, d\alpha}{\delta^\star}\right), \qquad \delta^\star = 1.683. \tag{14}$$

At the cheapest decorrelating span $[\,0.02, 1.0\,]$ the length is $\beta \int \sqrt{C} \approx 227.5$, requiring $R^\star = \mathbf{136}$ reversible rungs against a budget of $R_{\max} = 96$. The two frontiers do not overlap: decorrelation requires $\alpha_{\text{top}} \le 0.02$, tractability requires $\alpha_{\text{top}} \gtrsim 0.25$ ($R^\star(0.25) = 94$), and no top temperature satisfies both. The gap between the temperature you need to decorrelate and the temperature at which a tractable ladder exists *is* the obstruction. This is a feasibility stop, config-scoped to this kernel and landscape; because $R^\star$ is a Gaussian-overlap lower bound for a multimodal target, the true rung count could be higher still.

### 4.4 The A2↔A6 antagonism, framed

The theory leg is one-step algebra over the [solid] identity $T^C_{\mathcal{O}} = (\mu_C/\gamma_{\text{eff}})(1 - \gamma_{\text{eff}}/2) \to \infty$ as $\gamma_{\text{eff}} \to 0$: as the gap collapses (deep plateau), the window A6 demands diverges as $K \gg 1/\gamma_{\text{eff}}$. The mechanism leg (exp5) makes the collapse and the reversibility penalty a measured fact. The scale leg (exp4, exp6, exp19) shows the consequence at DTM scale is *unreachable* A6. Net: satisfying the theorem's reversibility assumption A2 tends to push the sampler toward the slow-mixing plateau where the operational gate A6 is hardest. **The fundamentality of this obstruction — whether it is a deep impossibility for multimodal reversible kernels or a deep-checkpoint, scale-dependent effect — is OPEN.** The small-scale evidence (bounded antagonism, escapable sweet spot) and the at-scale evidence ($\tau \propto L$, 136-rung wall) are complementary, not contradictory; whether small-family escapability survives to DTM scale is exactly the open question.

## 5 Discussion

**What is and is not established.** Result 1 is established: the single-$\gamma$ predictor is mis-anchored under symmetry, and the repair is an $L^2(\pi_\theta)$-projection, with O1.c [proven-here] and the orthogonality empirically at machine zero. Result 2's *conditional* factorization is [solid] (a written O1–O6 assembly), and the predictor tracks $Q_{\text{op}}$ on small and moderate controlled models; but the *operational* claim stays [conjectured], gated on A7 and $K \gg \tau_{\text{int}}$, and is [not validated]. Result 3's antagonism is a measured mechanism at small scale and an unreachable A6 at DTM scale, but its fundamentality is OPEN. We have flipped nothing to validated; terminal status is conferred only on the basis of a written proof (O1.c) or a pre-registered at-scale run that meets its preconditions (not yet achieved).

**Relation to quantum barren plateaus.** The phenomenology — a squared-SNR collapse that stalls training — is shared with the quantum barren-plateau literature [5, 6], and the three-factor structure of trainability ($\gamma \to 0$, $R \to 0$, budget starvation) maps factor-by-factor onto the quantum picture ($\gamma \leftrightarrow 1/\dim \mathfrak{g}$, $R \leftrightarrow$ purity products). *But the mechanism differs.* In the barren plateau the true gradient variance itself vanishes exponentially — signal extinction. Here the true gradient $g$ does not vanish; what collapses is its *recoverability* from a finite chain — an estimation plateau. The mapping is a shape analogy, not a literal transfer: one extinguishes the signal, the other defeats the estimator.

**Implications for thermodynamic computing.** The reversible kernel is precisely the one the analysis requires and the one that mixes slowest on the trained DTM. Two routes remain open. A *sampler* route — switching to hierarchical PT, simulated tempering, or population annealing to amortize the 136-rung length differently — risks breaking the reversibility A2 the theorem needs, trading a mixing problem for a theorem problem. A *measurement* route — a $\tau_{\text{int}}$-robust estimator of $Q_{\text{op}}$ that validates the operational claim without a fast-mixing kernel — fixes the measurement rather than the sampler. Either way, the diagnostic these results most strongly support is using the projected predictor $Q^{\perp}_{\text{struct}}$ on the gradient observables *before* training, rather than diagnosing the autocorrelation after the fact.

**Scope.** All positive statements are scoped: "on controlled small-$N$ planted or random families" (exp1, exp5), "on small-to-moderate RBMs under block-Gibbs" (exp2), and "on this $60 \times 12$ MNIST DTM, seed 0, single-input conditional, with this reversible kernel" (exp4, exp6, exp19). None is "validated for DTM training."

## Acknowledgments

This work was carried out as part of a thermodynamic-machine-learning research program. The author thanks the maintainers of the Extropic `thrml` library for a clean, reproducible reference implementation.

## References

[1] G. E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800, 2002.

[2] T. Tieleman. Training restricted Boltzmann machines using approximations to the likelihood gradient. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 1064–1071, 2008.

[3] L. Younes. On the convergence of Markovian stochastic algorithms with rapidly decreasing ergodicity rates. *Stochastics and Stochastic Reports*, 65(3–4):177–228, 1999.

[4] D. A. Levin and Y. Peres. *Markov Chains and Mixing Times.* American Mathematical Society, 2nd edition, 2017.

[5] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven. Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9:4812, 2018.

[6] M. Ragone, B. N. Bakalov, F. Sauvage, A. F. Kemper, C. Ortiz Marrero, M. Larocca, and M. Cerezo. A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits. *Nature Communications*, 15:7172, 2024.

[7] F. Jelinčič et al. Denoising Thermodynamic Models. *arXiv preprint* arXiv:2510.23972, 2025.

[8] Thermodynamic-trainability experiment exp1 (exact-diagonalization, $N \leq 14$ controlled Ising/Hopfield families): report on symmetry mis-anchoring and the observable-relevant cluster correction, 2026.

[9] Thermodynamic-trainability experiment exp2 (bipartite-RBM block-Gibbs smoke test, `thrml` fidelity check): report on at-scale tracking of $Q^{\perp}_{\mathrm{struct}}$ and the odd-scalar diagnostic failure, 2026.

[10] Thermodynamic-trainability experiment exp5 (gap-versus-multimodality, exact-from-spectrum bipartite RBM): report on $\gamma_{\mathrm{eff}}$ collapse and the reversibilization penalty, 2026.

[11] Thermodynamic-trainability experiment exp4 (reversible-kernel $\gtrsim 50\,\tau$ re-freeze, $60{\times}12$ MNIST DTM, H200): report on the $\tau \propto L$ doubling probe and the P0-HALT, 2026.

[12] Thermodynamic-trainability experiment exp6 (checkpoint $\tau$-sweep, reversible kernel, $t \in \{25, 50, 100, 200\}$): report on $\tau \propto L$ from very early training, 2026.

[13] Thermodynamic-trainability experiment exp19 (hotter-top equal-acceptance reversible-PT feasibility probe): report on the thermodynamic-length cost wall ($R^{\star} = 136$ vs $R_{\max} = 96$), 2026.

[14] Thermodynamic-trainability synthesis, proof O1: projection in $L^2(\pi_\theta)$, not state-space conditioning (O1.a, O1.b [solid]; O1.c [proven-here]), 2026.

[15] Thermodynamic-trainability synthesis, proof O5: off-cluster subdominance and the harmonic-mean $\gamma_{\mathrm{eff}}$ identity ([solid]), 2026.