Reproducible Spectral Diagnostics for Modular Phase Transitions

Nesen Oleg

Abstract

We develop a reproducible diagnostic framework for modular phase transitions in RG-proxy flows of quantum states. The diagnostics are derived from the modular generator $K_\rho = -\log\rho$ and include (i) spectral quantile coordinates $k_q = -\log\lambda_q(\rho)$, (ii) a normalized commutator probe $L(\rho; O) = \|[K_\rho, O]\|_F/(\|K_\rho\|_F\|O\|_F)$, and (iii) a running exponent ν extracted from sliding log–log fits of $L(t)$. We introduce pre-registered stability rules using bootstrap over seeds and model-selection criteria (AICc/BIC/LOOCV), converting qualitative phase narratives into PASS/FAIL statements. Using a reference experiment (dim=64; $N_{\mathrm{vis}} = 5$, $N_{\mathrm{hid}} = 1$; Gibbs-chaotic family with mixing parameter p), we demonstrate robustness across observables (near vs far), normalization choices, and mild noise channels, and we identify explicit domain limitations (failure on extreme random-pure families). Finally, we explain how these diagnostics interface with phase/locality criteria through spectral-flow markers and provide a practical route to numerical phase mapping under modular RG-proxy dynamics.

1 Introduction

Foundational programs that aim to relate quantum structure, dynamics, and emergent effective descriptions often face a methodological bottleneck: phase claims are typically stated conceptually but lack diagnostics that are simultaneously (i) representation-invariant, (ii) reproducible under protocol variations, and (iii) equipped with explicit acceptance criteria. Without such a layer, numerical evidence is hard to compare across models and parameter choices, and "phase transitions" risk becoming narrative-dependent.

This paper provides a reproducible diagnostic framework for modular phase transitions in the A→B→C series. Paper A yields the canonical state object ρ from normalized probabilistic structure. Paper B yields the canonical state-derived spectral coordinate $K = -\log\rho$ and CPTP-compatible RG-proxy evolution (continuous or discrete). Paper C defines phases and emergent locality via access structures and information distances. Here we supply the diagnostic layer linking these components: spectral quantile coordinates k_q, commutator probes $L(\rho; O)$, and running exponents ν, together with pre-registered stability rules (bootstrap and model selection) that convert qualitative phase narratives into PASS/FAIL statements. The focus is conservative: we report domain limitations explicitly and treat closure relations as validated diagnostics rather than universal laws.

2 Setup: RG-proxy flow and observed objects

We assume an RG-proxy evolution producing $\rho(t)$ (or a discrete sequence ρ_n) of density operators. The parameter t is interpreted as an informational scale, consistent with Paper B; discrete iterative CPTP implementations are primary, and continuous-time generators serve as effective limits.

For faithful states, the modular generator is $K_\rho = -\log\rho$. When numerical stability requires it, we use the faithful regularization

$$\rho_\epsilon = (1-\epsilon)\rho + \epsilon\,\mathbb{1}/d,$$

ensuring $\log \rho_\epsilon$ is well-defined. Diagnostics are computed from $\rho(t)$, $K(t)$, and commutators with selected observables O.

3 Spectral coordinates and quantiles

Let $\lambda_1 \geq \cdots \geq \lambda_d$ be eigenvalues of ρ. Fix quantiles $q \in (0,1)$ and define $\lambda_q(\rho)$ under a fixed convention, and the spectral quantile coordinates

$$k_q(\rho) = -\log \lambda_q(\rho).$$

These coordinates are representation-invariant and serve as compact tail-scale variables for modular spectral deformation along dissipative RG-proxy flows.

4 Commutator probe and normalized diagnostic

Definition 1 (Normalized commutator probe)**.** Given an observable O, define

$$L(\rho; O) = \frac{\|[K_\rho, O]\|_F}{\|K_\rho\|_F \, \|O\|_F}, \qquad K_\rho = -\log \rho,$$

where $\| \cdot \|_F$ is the Frobenius norm.

Proposition 1 (Representation invariance)**.** *For a unitary change of representation $\rho \mapsto U\rho U^\dagger$ and $O \mapsto UOU^\dagger$,*

$$L(U\rho U^\dagger; UOU^\dagger) = L(\rho; O).$$

Proof sketch. Functional calculus gives $K_{U\rho U^\dagger} = UK_\rho U^\dagger$; Frobenius norms are unitary invariant, and commutators transform covariantly. □

This diagnostic is designed to detect changes in modular spectral organization that are operationally visible through noncommutativity with O.

5 Running exponent and transition markers

Given a discretized scale parameter t (or mixing parameter p), we compute $L(t)$ and define a running exponent $\nu(t_c)$ by sliding log–log fits on windows centered at t_c:

$$\log L(t) \approx a + \nu \log t, \qquad t \in [t_c - w/2,\ t_c + w/2].$$

Windows are accepted only when fit quality exceeds a pre-registered threshold (e.g., $R^2 \geq R^2_{\min}$). A crossover marker t_x (or p_x) is then defined procedurally via the maximal per-step drop $\Delta\nu$ between consecutive window centers.

6 Reference Experiment

All robustness and stability claims are benchmarked against a fixed baseline.

Table 1. Reference Experiment: baseline parameters for reproducible diagnostics

System / dimension	$d = 64$, $N_{\text{vis}} = 5$, $N_{\text{hid}} = 1$
Family / path	$\rho(p) = (1-p)\sigma_{\text{geo}} + p\,\rho_{\text{ng}}$
Mixing range	$p \in [0.01, 0.20]$
Grid step	$\Delta p = 0.01$
Running exponent	window width $w = 0.05$ (in p units), center step 0.01
Fit model	$\log L(p) \approx a + \nu \log p$ (per window)
Observed fit quality	$R^2_{\text{min}} \approx 0.9971$–$0.9972$ (reported runs)
Seeds (near–far)	$\{501, 503, 509, 521, 541, 557\}$ (6 seeds)
Seeds (normalization)	$\{501, 503, 509, 521\}$ (4 seeds)
Primary marker	$p_x \approx 0.04$; $\Delta\nu_{\text{max}} \approx -0.051$ to -0.052

7 Results

7.1 Observable robustness: near vs far

In the reference experiment (Table 1), running-exponent diagnostics are essentially invariant under switching between a near and a far observable (e.g., Z_0Z_1 vs Z_0Z_4). The crossover estimate is stable at

$$p_x \approx 0.04$$

for both observables. The maximal drop in ν occurs between window centers $p_c = 0.035 \to 0.045$, with $\Delta\nu_{\text{max}} \approx -0.051$ (near and far). Across accepted windows, fit quality remains uniformly high ($R^2_{\text{min}} \approx 0.9971$–$0.9972$).

7.2 Normalization robustness

Using the normalized probe $L(\rho; O)$ (Table 1), the crossover location remains unchanged: $p_x = 0.04$ for both near and far observables. Normalization shifts the overall ν-range slightly (as expected for scale normalization) but preserves the profile and marker; $R^2_{\text{min}} \approx 0.9971$ throughout accepted windows.

7.3 Spectral link: ν tracks quantile spectral scale

Across high-quality windows, ν is strongly tied to a modular spectral tail scale (illustrated via a 50% commutator-energy cutoff k_{50}): $\text{corr}(\nu, k_{50}) \approx 0.968$ (far) and 0.980 (near), while $\text{corr}(\nu, \lambda_{50}) \approx -0.949$ (far) and -0.968 (near). Near the crossover region ($p_c = 0.035 \to 0.045$), the drop in ν coincides with a consistent shift of the tail scale, supporting the interpretation that the observed crossover is a modular spectral-tail phenomenon rather than an observable-specific artifact.

7.4 Closure laws $\nu \approx f(k_q, q)$ as predictive diagnostics

We treat $\nu(k_q, q)$ as a predictive diagnostic relation and evaluate it out-of-sample. In validated domains: (i) with sparse quantiles $q \in \{0.5, 0.8, 0.9\}$, a quadratic-in-k and quadratic-in-q closure achieves mean holdout $R^2 \approx 0.98$; (ii) for extended quantiles $q = 0.5, \ldots, 0.9$, a quadratic-in-k closure with cubic dependence in q restores strong generalization, with holdout $R^2 \approx 0.91$–0.92 across additional observable classes.

8 Stability rules (PASS/FAIL)

We define reproducibility as a property of the pre-registered pipeline.

8.1 Bootstrap stability

Key quantities (e.g., p_x, fitted coefficients, holdout errors) are evaluated under seed bootstrap. In the reference experiment, window fits are uniformly high ($R^2_{\min} \gtrsim 0.997$), yielding stable ν-profiles and a robust crossover marker.

8.2 Model selection and acceptance

Diagnostic closures are accepted only if they improve out-of-sample performance under explicit criteria (AICc/BIC thresholds and cross-validation error comparisons) and achieve a target bootstrap win-rate. In the Gibbs-like/mixed domain, quad_cubic consistently outperforms linear baselines; representative performance under mild channels is mean holdout $R^2 \approx 0.95$ with RMSE ≈ 0.011–0.012.

8.3 Invariance checks

A diagnostic claim must persist under reasonable variations of windowing, observable class, and mild channel perturbations. Portability across observable classes improves materially for quad_cubic (e.g., mean holdout $R^2 \approx 0.880$ vs 0.723 for a linear baseline; RMSE ≈ 0.0114 vs 0.0183 in tested scenarios).

9 Domain limitations

The diagnostic closures and robustness properties are domain-specific. They perform strongly for Gibbs-like/mixed families and under mild noise channels (e.g., depolarization $\mu \sim 0.02$, dephasing $\kappa \sim 0.10$), but can fail for extreme random-pure ensembles. In that failure regime, predictive performance collapses (e.g., quad_cubic holdout $R^2 \approx 0.14$; linear baseline ≈ 0.43), and adding depolarization does not restore generalization; dephasing improves performance but remains subtarget (quad_cubic mean holdout $R^2 \approx 0.69$). We therefore treat the closures as stable diagnostics for a "thermal/mixed-like" universality class of modular flows rather than as universal laws across arbitrary pure-state ensembles.

10 Bridge to phases and locality (Paper C)

Paper C defines phase-local descriptions via an optimal partition P and quantifies locality by the block total correlation

$$\mathcal{T}_{P(\rho)=D\left(\rho \middle\| \bigotimes_{X\in P}\rho_X\right)}.$$

The diagnostics developed here provide a spectral interface to this phase structure. In practice, changes in quantile flows $k_q(t)$, commutator probes $L(t)$, and running exponents $\nu(t)$ supply reproducible transition markers that can be correlated with (i) changes in the optimal partition P and (ii) growth of $\mathcal{T}_{P(\rho)}$. This yields an implementable protocol for phase mapping under modular RG-proxy dynamics.

11 Reproducibility package

We treat reproducibility as part of the scientific claim. A minimal package includes: (i) per-seed raw diagnostics (CSV), (ii) aggregated summaries with bootstrap confidence intervals (CSV/JSON), and (iii) a manifest specifying parameter values, windowing rules, and version information. This structure supports independent verification and cross-model comparison.

12 Conclusion

Scientific value

This paper provides the diagnostic layer for a modular foundations program: spectral quantile coordinates k_q, a normalized commutator probe $L(\rho; O)$, and running exponents ν, together with explicit stability and model-selection rules. By coupling invariance (unitary covariance), reproducibility (bootstrap), and pre-registered acceptance criteria (information criteria and cross-validation), the framework turns modular phase-transition discussions into quantitative PASS/FAIL statements and links the canonical state representation (Paper A), modular/CPTP dynamics (Paper B), and phase/locality criteria (Paper C) into a testable pipeline.

Applied value and future directions

Practically, the framework serves as a reproducible numerical instrument for (i) detecting transition markers (e.g., p_x) in modular RG-proxy flows, (ii) selecting between competing diagnostic relations $\nu \approx f(k_q, q)$ under explicit model-selection rules, and (iii) supporting automated phase mapping in open-system simulations where state spectra deform nontrivially. It can also guide principled decisions about when phase-local (block-factorized) descriptions are justified, by pairing spectral transition markers with growth in $\mathcal{T}_{P(\rho)}$.

Scientifically, the results point toward identifying universality classes of modular spectral flow and isolating robust spectral signatures of locality breakdown. The quantile-based coordinates k_q provide a compact description of tail deformation, suggesting testable targets for subsequent work: scaling of transition markers under spectator-large-N extensions, robustness under broader channel classes, and the characterization of critical regimes where stable partitions fail to persist. These directions set up the next steps of the series: full numerical phase diagrams under RG-proxy dynamics and systematic exploration of modular criticality using reproducible, pre-registered diagnostics.

REFERENCES

References

[1] J. von Neumann, *Mathematical Foundations of Quantum Mechanics*, Princeton University Press, 1955.

[2] E. T. Jaynes, Information theory and statistical mechanics, *Physical Review* **106** (1957), 620–630; **108** (1957), 171–190.

[3] G. Lindblad, On the generators of quantum dynamical semigroups, *Communications in Mathematical Physics* **48** (1976), 119–130.

[4] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, Completely positive dynamical semigroups of N-level systems, *Journal of Mathematical Physics* **17** (1976), 821–825.

[5] H. Spohn, Entropy production for quantum dynamical semigroups, *Journal of Mathematical Physics* **19** (1978), 1227–1230.

[6] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems*, Oxford University Press, 2002.

[7] D. Petz, Monotone metrics on matrix spaces, *Linear Algebra and its Applications* **244** (1996), 81–96.

[8] S. Amari and H. Nagaoka, *Methods of Information Geometry*, AMS & Oxford University Press, 2000.

[9] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary Edition, Cambridge University Press, 2010.