

A Non-Gaussian Clustering–Recovery Bridge via Conditional Mutual Information

Interacting Gibbs States, Explicit Petz Benchmarks, and a Conditional Link to Entanglement Wedge Reconstruction

Version 2

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

January 2026 (v1); July 2026 (v2)

Abstract

We present a quantitative clustering–recovery bridge for interacting quantum many-body systems that is intrinsically non-Gaussian, organized around conditional mutual information (CMI). For a geometric tripartition A – B – C in which B is a collar of width w separating A from C , an exponential geometric Markov bound $I_\rho(A : C|B) \leq K e^{-\alpha w}$ implies exponentially accurate recovery of ρ_{ABC} from ρ_{AB} in the theorem-facing metric $-\log F$, by combining the Fawzi–Renner inequality with an elementary conversion to fidelity error bounds. We obtain a proved interacting lane (shielded small-region geometry, arbitrary temperature) by invoking recent local Markovness results for finite-range lattice Gibbs states. Numerically, we benchmark the mechanism in the transverse-field Ising chain with longitudinal field, comparing integrable ($h_z = 0$) and non-integrable ($h_z = 0.5$) regimes, and evaluate the explicit Petz recovery map with a censored log-plotting and fit protocol. *v2 corrects the Fawzi–Renner factor* under the squared-fidelity convention used throughout ($-\log F \leq I$, not $\frac{1}{2}I$; fourth occurrence of this correction in the series), with a substantive empirical consequence: v1’s headline “mild overshoot” of Petz over the FR scale ($r \approx 1.23$ – 1.26 in the non-integrable, low-temperature, minimal-collar regime, including rotated and twirled controls) was measured against the incorrect half scale; against the corrected scale the overshoot *disappears entirely* ($r \approx 0.62$), and the v1 conclusion of “a genuine gap between explicit Petz-type constructions and the existential optimal-recovery scale” is withdrawn. The corrected conclusion is stronger: explicit Petz satisfies the FR scale throughout the dataset, with at least $\sim 40\%$ margin even at the hardest point. A fully regenerable verification suite reproduces the pipeline from scratch. Finally, we state a conditional application to entanglement wedge reconstruction, separating proved information-theoretic content from bulk–boundary interface assumptions.

1 Introduction

A basic locality intuition is that geometric separation suppresses influence between distant regions; a distinct, operational question is whether such suppression implies recoverability of ρ_{ABC} from ρ_{AB} . In quasi-free regimes recovery can be analyzed through covariance blocks; in interacting systems such reductions fail, motivating intrinsic information-theoretic control via $I_\rho(A : C|B)$. Goals: (i) the general non-Gaussian implication $\text{CMI} \Rightarrow \text{recoverability}$; (ii) a proved interacting

lane via the Gibbs local-Markovness theorem of Chen–Rouzé; (iii) numerical benchmarks with an explicit recovery candidate (Petz); and a conditional holographic corollary.

Series positioning (added in v2). This note is the empirical partner of the Type III capstone 2601.0034 [11] within the clustering–recovery program: 2512.0101 [8] formulates the uniform-in- Λ non-Gaussian bridge as a conjecture (its v2 fixes the same FR factor corrected here); the present note trades uniformity for a *proved* shielded small-region lane by importing Chen–Rouzé [6] – the same input recorded in 2601.0020 [10] – and adds integrable/non-integrable Petz benchmarks. The Gaussian core is 2601.0007 [9] (with 2512.0060 [7] upstream).

2 Setup: tripartitions, CMI, and recovery

Definition 2.1 (Shielded small-region geometry). *A tripartition A – B – C where B separates A from C with collar width w , $|A|$ fixed, B contiguous. Numerics use the open-chain tripartition $A = \{1, \dots, L_A\}$, $B = \{L_A+1, \dots, L_A+w\}$, C the rest.*

Definition 2.2 (Fidelity). $F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1^2 \in [0, 1]$ (squared convention).

Definition 2.3 (CMI). $I_\rho(A : C|B) := S(\rho_{AB}) + S(\rho_{BC}) - S(\rho_B) - S(\rho_{ABC})$, natural logarithm.

3 From CMI to recoverability (general, non-Gaussian)

Lemma 3.1. For $F \in (0, 1]$: $1 - F \leq -\log F$.

Theorem 3.2 (Fawzi–Renner [1]; factor corrected in v2). *For any ρ_{ABC} there exists a CPTP map $\mathcal{R}_{B \rightarrow BC}$ such that*

$$-\log F(\rho_{ABC}, (\text{id}_A \otimes \mathcal{R}_{B \rightarrow BC})(\rho_{AB})) \leq I_\rho(A : C|B).$$

Remark 3.3 (Correction note). v1 stated Theorem 3.2 with $\frac{1}{2}I$ on the right-hand side. With F in the squared convention the correct constant is I ; the $\frac{1}{2}$ belongs to the root-fidelity form $-2\log f \leq I$. Same correction as 2512.0101 v2, 2601.0020 v2 and 2601.0034 v2. The change propagates to Corollary 3.4, Theorem 4.2, and – substantively – to the empirical FR benchmark of Section 5.

Corollary 3.4 (CMI controls fidelity error). *For any ρ_{ABC} there is a CPTP $\mathcal{R}_{B \rightarrow BC}$ with $1 - F \leq I_\rho(A : C|B)$; if $I_\rho(A : C|B) \leq \delta$ then $1 - F \leq \delta$.*

Remark 3.5 (Existential versus explicit). Theorem 3.2 constrains the *optimal* recovery; explicit candidates (Petz) need not satisfy it pointwise a priori.

4 Proved interacting lane

Theorem 4.1 (Local Markovness of Gibbs states (imported) [6]). *For Gibbs states of finite-range lattice Hamiltonians with bounded interaction degree, in shielded geometries with fixed-size A : $I_{\rho_\beta}(A : C|B) \leq K(\beta)e^{-\alpha(\beta)w}$, constants independent of total system size.*

Theorem 4.2 (Quantitative collar $\Rightarrow$ recovery; constant corrected in v2). *Under Theorem 4.1, there exists a CPTP recovery map with $-\log F \leq K(\beta)e^{-\alpha(\beta)w}$, and by Lemma 3.1 also $1 - F \leq K(\beta)e^{-\alpha(\beta)w}$.*

5 Numerical evidence (exact diagonalization)

Model: $H = -J \sum Z_i Z_{i+1} - h_x \sum X_i - h_z \sum Z_i$, open chain, $J = h_x = 1$; integrable ($h_z = 0$) vs non-integrable ($h_z = 0.5$); Gibbs states by ED; $N = 11$, $L_A = 2$, $\beta \in \{0.5, \dots, 5\}$, $w \in \{1, \dots, 5\}$. Plotting/fit policy as in v1 (unchanged and correct): censoring threshold $y_{\min} = 10^{-12}$, slopes only with ≥ 3 pre-floor points, eigenvalue clipping $\varepsilon_{\text{inv}} = 10^{-12}$ for Petz inverses, $-\log_{10}(-\text{eps})$ for $-\log F$ near $F \approx 1$.

Results. Figure 1 compares regimes at $\beta = 3$: clear exponential CMI decay over the pre-floor window; rapid saturation of Petz recovery to the numerical floor. Table 1 of v1 (decay fits; unchanged data) shows robust CMI slopes; Petz slopes in the non-integrable case often fail the $n \geq 3$ policy.

FR-scale benchmark (rewritten in v2). Figure 2 benchmarks explicit Petz recovery against the FR scale. v1 plotted $-\log F$ versus $\frac{1}{2}I_\rho(A : C|B)$ and reported a “mild overshoot” ($r_{\max} \approx 1.23$ at $(h_z, \beta, w) = (0.5, 5, 1)$, $N = 12$), robust under rotated ($t \in \{0, 0.5, 1, 2, 4\}$) and twirled (β_0 -weighted, $n = 81$ nodes) Petz variants and under clipping sweeps – and concluded a “genuine gap between explicit Petz-type constructions and the existential optimal-recovery scale.” *That conclusion is withdrawn in v2:* against the corrected scale $I_\rho(A : C|B)$, the same data give $r_{\max} \approx 0.62$, and Petz satisfies the FR bound *everywhere* in the dataset. The verification suite reproduces this from scratch (at $N = 9$: $r_{\max} = 0.615$ at the analogous hardest point, i.e. $2r_{\max} = 1.231$ against the v1 half scale – numerically matching v1’s reported overshoot and identifying it as exactly the factor-2 artifact). The rotated/twirled spot-checks are retained with reversed interpretation: they confirm that plain Petz is within the corrected FR scale with $\sim 40\%$ margin even in the non-integrable, low-temperature, minimal-collar regime, and that the rotated variants (which degrade monotonically in t) remain so.




Figure 1: Integrable ($h_z = 0$) vs non-integrable ($h_z = 0.5$) at $\beta = 3$: $I(A : C|B)$ and Petz recovery error vs collar width w (censored convention).




Figure 2: FR-scale benchmark (v1 figure, plotted against the v1 half scale $\frac{1}{2}I$; retained for transparency). Under the corrected scale I (Remark 3.3) all points, including the apparent overshoots, fall strictly below the reference line: $r_{\max} \approx 0.62$.

6 Conditional application: entanglement wedge reconstruction

Assumption 6.1 (AdS/CFT QEC interface). *A standard bulk-boundary QEC dictionary in which entanglement wedge reconstruction is recovery from erasure on a code subspace, with bulk reconstruction error controlled by boundary recovery error.*

Theorem 6.2 (Conditional quantitative wedge reconstruction). *Assume Assumption 6.1 (v1 mislabeled this cross-reference as a theorem). If a boundary family satisfies $I(A : C|B) \leq K e^{-\alpha w}$ for tripartitions associated to a boundary region and its buffer, then there is a recovery procedure with boundary error $O(e^{-\alpha w})$ in $-\log F$, and bulk reconstruction error inherits an exponentially suppressed bound under the interface assumptions. The Type III form of this dictionary is developed in 2601.0034 [11].*

7 Verification suite (added in v2)

`verification/2601-0035/verify_2601_0035.py` regenerates the pipeline from scratch (no archived data): ED Gibbs states of the declared model, CMI decay with the censored-fit policy, explicit Petz recovery, and the FR-scale ratio. Checks: corrected ratio $r = -\log F/I < 1$ on the whole grid (max 0.615); the v1 half-scale reproduces the spurious overshoot ($2r = 1.231$ at the hardest point); CMI slopes positive and consistent with Table 1 (e.g. $\mu_{\text{CMI}} = 1.15$ at $h_z = 0$, $\beta = 3$, $N = 9$ vs 1.10 at $N = 11$); and Lemma 3.1. Appendix A of v1 ($N = 12$ spot-check; figures unchanged) is retained.




Figure 3: $N = 12$ spot-check, non-integrable case (v1 figure, unchanged).

Note added in v2

One correction with empirical consequences, plus integration. (1) *FR factor* (Remark 3.3): Theorem 3.2, Corollary 3.4 and Theorem 4.2 restated with the squared-fidelity constant I (not $\frac{1}{2}I$); fourth occurrence of this correction in the series. (2) *Withdrawn finding*: v1’s “mild overshoot”/“genuine gap” of explicit Petz over the FR scale was an artifact of the half scale; corrected, explicit Petz satisfies the FR scale throughout the regenerated ED dataset (not claimed as a general theorem) with $r_{\max} \approx 0.62$, quantitatively confirmed from scratch by the suite ($2r_{\max} = 1.231$ matches v1’s reported 1.23). The rotated/twirled controls are retained as margin confirmations. (3) Cross-reference labels fixed (“Assume Theorem 6.1” $\rightarrow$ Assumption; “Theorem 3.1” $\rightarrow$ Lemma). (4) Series positioning and verification suite added; v1 cited no companion papers. All v1 data, figures, tables and the numerical policy are unchanged.

References

- [1] O. Fawzi and R. Renner, Commun. Math. Phys. **340** (2015) 575–611.
- [2] D. Petz, Commun. Math. Phys. **105** (1986) 123–131.
- [3] A. Almheiri, X. Dong, and D. Harlow, JHEP **04** (2015) 163.
- [4] X. Dong, D. Harlow, and A. C. Wall, Phys. Rev. Lett. **117** (2016) 021601.
- [5] J. Cotler et al., arXiv:1704.05839 (2017).
- [6] C.-F. Chen and C. Rouzé, arXiv:2504.02208 (2025).
- [7] L. Eriksson, ai.viXra:2512.0060 (2025; v2 2026).
- [8] L. Eriksson, ai.viXra:2512.0101 (2025; v2 2026).
- [9] L. Eriksson, ai.viXra:2601.0007 (2026; v2 2026).
- [10] L. Eriksson, ai.viXra:2601.0020 (2026; v2 2026).
- [11] L. Eriksson, ai.viXra:2601.0034 (2026; v2 2026).