

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

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

Lluís Eriksson

Independent Researcher

lluiseriksson@gmail.com

January 2026

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 (buffer) of width w separating A from C , we show that 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 then obtain a proved interacting lane (in a shielded small-region geometry) at 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 an integrable regime ($h_z = 0$) and a non-integrable regime ($h_z = 0.5$), and we evaluate the explicit Petz recovery map. We adopt a censored log-plotting and fit protocol that avoids numerical-floor artifacts and report decay-rate estimates only when at least three pre-floor points are available. Finally, motivated by the quantum error-correction interpretation of subregion duality, we state a conditional application to entanglement wedge reconstruction, cleanly separating proved information-theoretic content from bulk–boundary interface assumptions.

Contents

1	Introduction	2
2	Setup: tripartitions, CMI, and recovery	3
2.1	Geometric tripartition	3
2.2	Fidelity and conditional mutual information	3
3	From CMI to recoverability (general, non-Gaussian)	3
4	Proved interacting lane: exponential CMI decay for Gibbs states	4
4.1	Imported local Markovness input	4
4.2	Main theorem (clustering–recovery bridge)	4
5	Numerical evidence (exact diagonalization)	4
5.1	Model and geometry	4
5.2	Quantities and plotting policy	4
5.3	Numerical stability and reporting policy	5
5.4	Results	5
6	Conditional application: entanglement wedge reconstruction	5
A	Finite-size spot-check at $N = 12$ (non-integrable)	6

1 Introduction

A basic locality intuition in quantum many-body physics is that geometric separation suppresses influence between distant regions. A distinct, operational question is whether such suppression implies *recoverability*: can the global state on ABC be approximately reconstructed from the reduced state on AB when B is a buffer separating A from C ?

In quasi-free (Gaussian) regimes, recovery can often be analyzed through covariance-block techniques. In interacting systems, such reductions fail in general. This motivates using *intrinsic* information-theoretic quantities, in particular the conditional mutual information $I_\rho(A : C|B)$, which measures approximate quantum Markovness and directly controls recoverability via universal theorems.

This paper has three goals:

- Establish a general (non-Gaussian) implication $I(A : C|B) \Rightarrow$ recoverability.
- Provide a proved interacting lane (in a shielded small-region geometry) by invoking the Gibbs local Markovness theorem of Chen–Rouzé.
- Benchmark the story numerically in a canonical interacting spin chain, using an explicit recovery candidate (Petz).

We additionally discuss a conditional holographic corollary in the language of quantum error correction (QEC): collar-controlled Markovness is a natural sufficient condition for approximate recovery from erasure, and hence for quantitative subregion duality once bulk–boundary interface assumptions are fixed.

2 Setup: tripartitions, CMI, and recovery

2.1 Geometric tripartition

Definition 2.1 (Shielded small-region geometry). Fix a lattice system (e.g. a 1D chain) and a tripartition A – B – C such that B geometrically separates A from C with collar width w (graph distance). We call this a *shielded small-region geometry* if $|A|$ is fixed (independent of total system size) and B is a contiguous buffer of width w between A and C .

Remark 2.2 (1D geometry used in numerics). In the numerics we use a contiguous tripartition (open chain):

$$A = \{1, \dots, L_A\}, \quad B = \{L_A + 1, \dots, L_A + w\}, \quad C = \{L_A + w + 1, \dots, N\}.$$

2.2 Fidelity and conditional mutual information

Definition 2.3 (Fidelity). For density matrices ρ, σ on the same Hilbert space, define the (Uhlmann) fidelity

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

Definition 2.4 (Conditional mutual information). For a tripartite state ρ_{ABC} , define

$$I_\rho(A : C|B) := S(\rho_{AB}) + S(\rho_{BC}) - S(\rho_B) - S(\rho_{ABC}),$$

where $S(\rho) := -\text{Tr}(\rho \log \rho)$ is the von Neumann entropy (with natural logarithm).

Definition 2.5 (Recovery map). A recovery map is a completely positive trace-preserving (CPTP) map $R_{B \rightarrow BC}$ such that $(\text{id}_A \otimes R_{B \rightarrow BC})(\rho_{AB})$ approximates ρ_{ABC} .

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

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

Theorem 3.2 (Fawzi–Renner [1]). For any state ρ_{ABC} , there exists a CPTP map $R_{B \rightarrow BC}$ such that

$$-\log F(\rho_{ABC}, (\text{id}_A \otimes R_{B \rightarrow BC})(\rho_{AB})) \leq \frac{1}{2} I_\rho(A : C|B).$$

Corollary 3.3 (CMI controls fidelity error). For any ρ_{ABC} there exists a CPTP map $R_{B \rightarrow BC}$ such that

$$1 - F(\rho_{ABC}, (\text{id}_A \otimes R_{B \rightarrow BC})(\rho_{AB})) \leq \frac{1}{2} I_\rho(A : C|B).$$

In particular, if $I_\rho(A : C|B) \leq \delta$, then the right-hand side is at most $\delta/2$.

Proof. Combine Theorem 3.2 with Theorem 3.1. □

Remark 3.4 (Existential versus explicit recovery). Theorem 3.2 is existential: it constrains the *optimal* recovery. Explicit candidates (e.g. Petz) need not satisfy the FR inequality pointwise, especially in interacting regimes.

4 Proved interacting lane: exponential CMI decay for Gibbs states

4.1 Imported local Markovness input

We rely on a recent theorem establishing local Markovness for Gibbs states of finite-range Hamiltonians in shielded geometries at arbitrary temperature, in a regime where A is of fixed (small) size and shielded from the complement by a collar.

Theorem 4.1 (Local Markovness of Gibbs states (imported form) [6]). *Let $\rho_\beta = e^{-\beta H} / \text{Tr}(e^{-\beta H})$ be a Gibbs state of a finite-range lattice Hamiltonian with bounded interaction degree. In a shielded geometry where a fixed-size region A is separated from the complement C by a collar B of width w , there exist constants $K(\beta), \alpha(\beta) > 0$ (depending on β and local interaction parameters, but independent of total system size) such that*

$$I_{\rho_\beta}(A : C|B) \leq K(\beta) e^{-\alpha(\beta)w}.$$

4.2 Main theorem (clustering–recovery bridge)

Theorem 4.2 (Quantitative collar \Rightarrow recovery). *Under the hypotheses of Theorem 4.1, fix a tripartition A – B – C and set $\rho_{ABC} := (\rho_\beta)_{ABC}$ and $\rho_{AB} := (\rho_\beta)_{AB}$. Then there exists a CPTP recovery map $R_{B \rightarrow BC}$ such that*

$$-\log F(\rho_{ABC}, (\text{id}_A \otimes R_{B \rightarrow BC})(\rho_{AB})) \leq \frac{1}{2} K(\beta) e^{-\alpha(\beta)w}.$$

Proof. Combine Theorem 3.2 with Theorem 4.1. □

Remark 4.3. By Theorem 3.1, the same hypothesis also implies

$$1 - F(\rho_{ABC}, (\text{id}_A \otimes R_{B \rightarrow BC})(\rho_{AB})) \leq \frac{1}{2} K(\beta) e^{-\alpha(\beta)w}.$$

5 Numerical evidence (exact diagonalization)

5.1 Model and geometry

We benchmark recoverability in the transverse-field Ising chain with longitudinal field

$$H = -J \sum_{i=1}^{N-1} Z_i Z_{i+1} - h_x \sum_{i=1}^N X_i - h_z \sum_{i=1}^N Z_i,$$

with open boundary conditions. We compare an integrable setting ($h_z = 0$) and a non-integrable setting ($h_z = 0.5$). For each inverse temperature β we form the Gibbs state ρ_β by exact diagonalization.

We use a shielded tripartition A – B – C with A consisting of the leftmost L_A sites, B the contiguous collar of width w adjacent to A , and C the remaining sites. In the main text we fix $N = 11$ and $L_A = 2$, and sweep $\beta \in \{0.5, 1, 2, 3, 4, 5\}$ and $w \in \{1, 2, 3, 4, 5\}$.

5.2 Quantities and plotting policy

For each (β, w) we compute $I_{\rho_\beta}(A : C|B)$ and benchmark an explicit recovery candidate via the Petz map. We denote the Petz-reconstructed state by $\tilde{\rho}_{ABC}^{\text{Petz}} := (\text{id}_A \otimes R_{B \rightarrow BC}^{\text{Petz}})(\rho_{AB})$, and report recovery performance in the theorem-facing metric $-\log F$.

On log plots, recovery errors can hit machine precision quickly; we adopt a censored policy with threshold $y_{\min} = 10^{-12}$. Decay rates are fitted from $\log y \approx a - \mu w$ using only pre-floor points $y \geq y_{\min}$. If fewer than three pre-floor points are available for a given observable, we report “—” rather than an unstable slope.

5.3 Numerical stability and reporting policy

Petz recovery involves inverses of reduced states (e.g. $\rho_B^{-1/2}$). In finite precision we regularize by eigenvalue clipping: eigenvalues $\lambda < \epsilon_{\text{inv}}$ are replaced by $\epsilon_{\text{inv}} = 10^{-12}$ before taking matrix powers (e.g. $\lambda^{-1/2}$). We also clip the computed fidelity to $[0, 1]$ to remove floating-point overshoots. To compute $-\log F$ stably when $F \approx 1$, we set $\epsilon := 1 - F$ and evaluate

$$-\log F = -\log(1 - \epsilon),$$

implemented in code as `-log1p(-epsilon)` to avoid catastrophic cancellation when $\epsilon \ll 1$. To avoid artifacts on log scales when the error reaches machine precision, we adopt the censoring threshold $y_{\text{min}} = 10^{-12}$ and exclude censored points from slope fits. Slopes are reported only when at least three pre-floor points are available (otherwise shown as “—”).

5.4 Results

Figure 1 compares integrable and non-integrable regimes at $\beta = 3$. We observe clear exponential decay in $I(A : C|B)$ over the available pre-floor window and rapid saturation of Petz recovery to the numerical floor for modest collar widths.

Table 1 summarizes decay-rate fits. CMI slopes are robust across parameters, while Petz error slopes in the non-integrable case often cannot be reliably fitted due to rapid floor saturation (insufficient pre-floor points), and are therefore marked “—” by policy.

FR-scale benchmark (explicit recovery). Figure 2 benchmarks Petz recovery against the Fawzi–Renner recoverability scale by plotting $-\log F(\rho_{ABC}, \tilde{\rho}_{ABC}^{\text{Petz}})$ versus $\frac{1}{2}I_\rho(A : C|B)$, together with the reference line $y = x$. Since Theorem 3.2 is existential (optimal recovery), an explicit candidate need not satisfy the bound pointwise. Empirically, Petz lies below the reference line throughout the integrable dataset (0 overshoots in the pre-floor window). In the non-integrable case we observe a mild overshoot at low temperature and minimal collar width; in a spot-check at $N = 12$ the maximum ratio is $r_{\text{max}} \approx 1.23$ at $(h_z, \beta, w) = (0.5, 5, 1)$.

Rotated and twirled Petz controls (spot-check at $N = 12$). At the unique overshoot point $(N, h_z, \beta, w) = (12, 0.5, 5, 1)$ we tested explicit Petz-type variants. First, varying the eigenvalue clipping threshold over $\epsilon_{\text{inv}} \in [10^{-14}, 10^{-3}]$ leaves the recovery error and FR ratio unchanged to all displayed digits; in this regime $\epsilon_{\text{inv}} < \lambda_{\text{min}}(\rho_B) \approx 1.77 \times 10^{-3}$ and clipping is inactive. Second, within the rotated Petz family $R^{(t)}$ (tested at $t \in \{0, 0.5, 1, 2, 4\}$), the recovery error $-\log F$ increases relative to $t = 0$; the FR ratio grows from $r \approx 1.23$ at $t = 0$ to larger values as t increases. Finally, a twirled (averaged) rotated-Petz reconstruction was implemented by discretizing $t \in [-8, 8]$ with $n = 81$ equally spaced points and averaging the recovered states with weights proportional to $\frac{\pi^2}{4} \left(\frac{\pi t}{2}\right)$, followed by a final trace renormalization. This twirled channel also does not improve the overshoot: we obtain $-\log F \approx 9.84 \times 10^{-4}$, corresponding to an FR ratio $r \approx 1.26$. Thus the mild overshoot reflects a genuine gap between explicit Petz-type constructions and the existential optimal-recovery scale in this challenging regime.

6 Conditional application: entanglement wedge reconstruction

Assumption 6.1 (AdS/CFT QEC interface). *We assume a standard bulk–boundary QEC dictionary in which entanglement wedge reconstruction can be phrased as recovery from erasure on an appropriate code subspace, with a notion of reconstruction error controlled by boundary recovery error.*




Figure 1: Integrable ($h_z = 0$) vs non-integrable ($h_z = 0.5$) comparison at $\beta = 3$ for $I(A : C|B)$ and Petz recovery error as functions of collar width w , using the censored plotting convention.

Table 1: Exponential decay-rate fits $\log y \approx a - \mu w$ using only pre-floor points $y \geq y_{\min}$ with $y_{\min} = 10^{-12}$; entries are marked “—” when fewer than three pre-floor points are available.

h_z	β	μ_{CMI}	$\mu_{-\log F, \text{Petz}}$
0.0	0.5	5.1829 (4)	—
0.0	1.0	3.0362 (5)	6.7161 (3)
0.0	2.0	1.6140 (5)	3.4888 (3)
0.0	3.0	1.1034 (5)	2.4346 (5)
0.0	4.0	0.8623 (5)	1.6526 (5)
0.0	5.0	0.7290 (5)	1.3275 (5)
0.5	0.5	5.3022 (4)	—
0.5	1.0	3.5241 (5)	—
0.5	2.0	2.8260 (5)	—
0.5	3.0	3.0681 (5)	—
0.5	4.0	3.3958 (5)	—
0.5	5.0	3.5147 (5)	—

Theorem 6.2 (Conditional quantitative wedge reconstruction). *Assume Theorem 6.1. If a boundary state family satisfies a collar CMI bound of the form $I(A : C|B) \leq K e^{-\alpha w}$ for tripartitions associated to a boundary region and its buffer, then there exists a corresponding recovery procedure with boundary error bounded by $O(e^{-\alpha w})$ (e.g. in $-\log F$), and the bulk reconstruction error inherits an exponentially suppressed bound under the interface assumptions.*

Remark 6.3. This section isolates the interface assumptions; all CMI \Rightarrow recovery statements earlier are purely information-theoretic.

A Finite-size spot-check at $N = 12$ (non-integrable)

We provide a minimal finite-size robustness check at $N = 12$ for the non-integrable case ($h_z = 0.5$) and representative inverse temperatures $\beta \in \{1, 3, 5\}$. We keep the same shielded geometry ($L_A = 2$) and collar widths $w \in \{1, 2, 3, 4, 5\}$ for CMI. For $-\log F$ (Petz) we evaluate only a subset of collar




Figure 2: Benchmark against the FR recoverability scale using Petz recovery. We plot $-\log F(\rho_{ABC}, \tilde{\rho}_{ABC}^{\text{Petz}})$ versus $\frac{1}{2}I_{\rho}(A : C|B)$ (pre-floor points), with the reference line $y = x$. FR bounds the *optimal* recovery; Petz is an explicit candidate and can mildly overshoot the reference scale in the non-integrable, low-temperature, minimal-collar regime.

Table 2: Spot-check decay-rate fits at $N = 12$ using only pre-floor points $y \geq y_{\min}$ with $y_{\min} = 10^{-12}$; entries are marked “—” when fewer than three pre-floor points are available.

β	μ_{CMI}	$\mu_{-\log F, \text{Petz}}$
1.0	3.5343 (5)	—
3.0	3.0729 (5)	—
5.0	3.5275 (5)	—

widths at $N = 12$ due to computational cost. As a result, decay-rate fits for $-\log F$ at $N = 12$ may fail the $n \geq 3$ reporting threshold even when CMI fits remain robust.




Figure 3: Spot-check at $N = 12$ for the non-integrable case ($h_z = 0.5$) and $\beta \in \{1, 3, 5\}$. Left: $I(A : C|B)$ versus collar width w . Right: Petz recovery error $-\log F$ computed on a subset of w values (higher N cost), displayed with the same censored convention as in the main text.

References

- [1] O. Fawzi and R. Renner. Quantum Conditional Mutual Information and Approximate Markov Chains. *Communications in Mathematical Physics*, 340:575–611, 2015.
- [2] D. Petz. Sufficient subalgebras and the relative entropy of states of a von Neumann algebra. *Communications in Mathematical Physics*, 105:123–131, 1986.
- [3] A. Almheiri, X. Dong, and D. Harlow. Bulk locality and quantum error correction in AdS/CFT. *JHEP*, 04:163, 2015.
- [4] X. Dong, D. Harlow, and A. C. Wall. Reconstruction of Bulk Operators within the Entanglement Wedge in Gauge-Gravity Duality. *Phys. Rev. Lett.*, 117:021601, 2016.
- [5] J. Cotler, P. Hayden, G. Penington, G. Salton, B. Swingle, and M. Walter. Entanglement Wedge Reconstruction via Universal Recovery Channels. *arXiv:1704.05839*, 2017.
- [6] C.-F. Chen and C. Rouzé. Quantum Gibbs states are locally Markovian. *arXiv:2504.02208*, 2025.