

Typed Pipeline for Recoverability–Rate–Power Links: A Contract Paper with a Closed Recoverability Lane and Falsification Criteria

Lluis Eriksson

January 2026 (v1); July 2026 (v2)

Abstract

We present an audit-friendly *logical contract* for a multi-layer program connecting (i) static locality/Markovness, (ii) recoverability bounds, (iii) separation-dependent dissipation rates, and (iv) thermodynamic maintenance power. Each interface is typed with explicit quantifiers, tagged as **[PROVED]**/**[IMPORTED]**/**[ASSUMED]**/**[CONJECTURED]**, and paired with falsification routes. We do *not* claim a proof of the Clay Yang–Mills mass gap; we separate a Clay (closed-Hamiltonian) track from an operational (open-system/maintenance) track. As a fully closed lane inside this paper, we prove that an exponential conditional mutual information (CMI) decay hypothesis implies exponential recoverability via an imported Fawzi–Renner inequality, with fidelity conventions fixed explicitly. *v2 (no v1 number is changed)*: the cross-reference labels are repaired (in v1 every assumption, lemma, theorem, remark and corollary was typeset as “Definition $x.y$ ”) and Tables 1 and 2 are re-typeset; the constant-alignment step of Appendix A is executed rather than prescribed — in the locked conventions (squared fidelity, $E_{\text{rec}} = -\log F$) the Fawzi–Renner import yields $c_{\text{FR}} = 1$ exactly (Remark 4.10); the deferred interfaces of Appendix B are now concrete series papers and are cited as such (the Type III interface is ai.viXra:2601.0065, the Davies/RIP-U dynamics note is ai.viXra:2601.0064, the benchmark infrastructure lives in the series repository); the BATO-LAW upgrade path of Appendix C records the partial progress of ai.viXra:2601.0031; and a verification suite instantiates everything instantiable: the closed recoverability lane end to end on a gapped transverse-field Ising collar ($I \approx 0.20 e^{-1.06w}$, $E_{\text{rec}} \leq I$ at every width in the generated dataset, full $1 - F \leq E_{\text{rec}} \leq c_{\text{FR}} K e^{-\alpha\epsilon}$ chain), the exact-Markov example at machine precision, and the dephasing example with an explicit $\kappa^\uparrow/\kappa^\downarrow$ spread of two orders of magnitude illustrating the directionality golden rule.

1 Positioning and related work (minimal)

Recoverability from small conditional mutual information (CMI) was initiated in the finite-dimensional setting by Fawzi–Renner [1], with subsequent developments clarifying strengthened and universal recovery maps (see e.g. [2] and related work). Exponential CMI decay in Gibbs or quasi-local settings has been studied in several regimes in the many-body/QIT literature; in this paper we keep that input as an explicit falsifiable hypothesis (A-CMI) rather than embedding it as a black-box “Gibbs theorem” claim. The present contribution is the *typed contract* structure and the explicit separation between proved implications and bridge assumptions (especially toward operational power costs).

2 Executive summary (audit view)

Scope and non-claims

- **Not a Clay proof:** no claim is made about existence of continuum Yang–Mills theory or its Hamiltonian spectral gap.
- **This is a contract paper:** the value is explicit typing, status tagging, and falsification routes.
- **Closed lane included:** we include one fully self-contained implication: CMI decay (assumption) + FR-type recoverability (imported) $\Rightarrow$ recoverability decay (proved).

Status table (what is actually proved inside this PDF)

Item	Status	Depends on	Falsification route
L1: $1 - F \leq -\log F$	[PROVED]	elementary calculus	direct check
T2: FR-type inequality (constant-fixed)	[IMPORTED]	finite-dimensional QIT	disproof or convention mismatch
A-CMI: CMI decay hypothesis	[ASSUMED]	chosen model/geometry	exhibit family violating stated decay
T3: CMI decay $\Rightarrow$ recoverability decay	[PROVED]	A-CMI + T2 + L1	refute A-CMI or mismatch hypotheses
BATO-LAW: maintenance inequality	[ASSUMED]	operational definition	violate hypothesis or give counter-protocol

Table 1: Audit status for this paper. Anything tagged **[ASSUMED]** is a typed hypothesis (not a hidden dependency). (Re-typeset in v2; v1’s table had overlapping column text.)

3 Typing, notation, and objects

Remark 3.1 (Parameter convention). Throughout: ϵ is a **geometric** separation/collar parameter, while δ denotes **tolerances** (regularization thresholds, error budgets, floors). They are never interchanged.

Definition 3.2 (D1: Δ -track coherence). Fix a dephasing (pinching) CPTP map Δ . Define $C_\Delta(\rho) := S(\rho \| \Delta[\rho])$, where $S(\cdot \| \cdot)$ is quantum relative entropy (with the usual support convention).

Definition 3.3 (D2: instantaneous coherence loss under fixed dynamics). Let $T_t = e^{t\mathcal{L}}$ be a fixed CPTP Markov semigroup on states. Define

$$\dot{C}_{\Delta, \text{loss}}(\rho) := - \left. \frac{d}{dt} \right|_{t=0} C_\Delta(T_t(\rho)),$$

whenever the derivative exists.

Remark 3.4 (A minimal sufficient condition for existence). A sufficient (not necessary) condition for $\dot{C}_{\Delta, \text{loss}}(\rho)$ to exist is differentiability at $t = 0$ of $t \mapsto C_\Delta(T_t(\rho))$. In finite dimension this holds under mild spectral regularity (e.g. full-rank conditions), but this paper does not fix a universal regularity class; we treat existence as part of the typing of any application.

Definition 3.5 (D3: physically implementable local operations $\text{Ops}_A^{\text{phys}}$). Fix a reference state ρ_0 , a tripartition $A-B-C$, and a separation notion $\text{sep}(A, C)$. Let $\text{Ops}_A^{\text{phys}}$ denote a declared class of *physically implementable* operations on A (CPTP maps with local ancillas allowed and discarded), realizable within stated resource tolerances δ (time, energy, ancilla size, etc.).

Definition 3.6 (D4: the family $\mathcal{F}_\epsilon$ (minimal; typed by $\text{Ops}_A^{\text{phys}}$)). Define $\mathcal{F}_\epsilon := \{\rho = \Lambda_A(\rho_0) : \Lambda_A \in \text{Ops}_A^{\text{phys}}, \text{sep}(A, C) \geq \epsilon, C_\Delta(\rho) > 0\}$.

Assumption 3.7 (REG: regularity restrictions (invoked only when needed)). *When needed to exclude trivial degeneracies (e.g. near-fixed points or nearly singular diagonals) we impose explicit regularity restrictions on $\mathcal{F}_\epsilon$, such as $\text{dist}(\rho, \text{Fix}(T)) \geq \delta_{\text{fix}}$ and/or $\Delta[\rho] \succeq p_{\min} \mathbf{1}$. These restrictions are not part of the minimal definition of $\mathcal{F}_\epsilon$; they are invoked only where required and always stated at the point of use.*

Definition 3.8 (D5: envelopes (quantifiers explicit)). Define the domain of the ratio at scale ϵ as $\mathcal{D}_\epsilon := \{\rho \in \mathcal{F}_\epsilon : C_\Delta(\rho) > 0, \dot{C}_{\Delta, \text{loss}}(\rho) \text{ exists}\}$. Then define

$$\kappa^\uparrow(\epsilon) := \sup \left\{ \frac{\dot{C}_{\Delta, \text{loss}}(\rho)}{C_\Delta(\rho)} : \rho \in \mathcal{D}_\epsilon \right\}, \quad \kappa^\downarrow(\epsilon) := \inf \left\{ \frac{\dot{C}_{\Delta, \text{loss}}(\rho)}{C_\Delta(\rho)} : \rho \in \mathcal{D}_\epsilon \right\}.$$

Remark 3.9 (Empty-domain convention). In any concrete application we assume $\mathcal{D}_\epsilon \neq \emptyset$ in the regime of interest. If $\mathcal{D}_\epsilon = \emptyset$, we adopt the convention $\kappa^\uparrow(\epsilon) := 0, \kappa^\downarrow(\epsilon) := +\infty$, so that envelope statements become vacuous rather than ambiguous. (Any alternative convention should be stated explicitly if used.)

Remark 3.10 (Directionality (golden rule)). A **minimum-power** statement requires a **lower envelope** bound (involving $\kappa^\downarrow$), *not* an upper envelope bound (involving $\kappa^\uparrow$).

4 Closed recoverability lane: CMI decay $\Rightarrow$ exponential recovery

Definition 4.1 (Fidelity convention). For density matrices ρ, σ on the same finite-dimensional Hilbert space, we use the *squared* Uhlmann fidelity $F(\rho, \sigma) := \|\sqrt{\rho}\sqrt{\sigma}\|_1^2 \in [0, 1]$.

Definition 4.2 (Recovery error). For a tripartite finite-dimensional state ρ_{ABC} and a recovery map $\mathcal{R}_{B \rightarrow BC}$, define $E_{\text{rec}}(\rho_{ABC}; \mathcal{R}) := -\log F(\rho_{ABC}, (\text{id}_A \otimes \mathcal{R}_{B \rightarrow BC})(\rho_{AB}))$.

Lemma 4.3 (L1: elementary inequality). **[PROVED]** For $F \in (0, 1]$ one has $1 - F \leq -\log F$.

Proof. Define $g(F) := -\log F - (1 - F)$ on $(0, 1]$. Then $g'(F) = -\frac{1}{F} + 1 \leq 0$ for $F \in (0, 1]$ and $g(1) = 0$, so $g(F) \geq 0$ for all $F \in (0, 1]$. $\square$

Theorem 4.4 (T2: FR-type recoverability inequality (constant explicit)). **[IMPORTED]** (finite-dimensional recoverability literature; see e.g. [1, 2]). *There exists a universal constant $c_{\text{FR}} > 0$ (depending only on the choice of fidelity/error convention) such that for any finite-dimensional ρ_{ABC} , there exists a CPTP recovery map $\mathcal{R}_{B \rightarrow BC}$ with*

$$E_{\text{rec}}(\rho_{ABC}; \mathcal{R}) \leq c_{\text{FR}} I(A : C|B)_\rho.$$

Remark 4.5 (Convention lock-in). This paper fixes: (i) squared fidelity (Definition 4.1), and (ii) error functional $E_{\text{rec}} = -\log F$ (Definition 4.2). The value of c_{FR} should be aligned to the imported theorem under these conventions (Appendix A gives the alignment step; *executed in v2*, Remark 4.10). Using c_{FR} avoids silent factor drift across conventions.

Assumption 4.6 (A-CMI: exponential CMI decay in a shielded geometry). *[ASSUMED] (model/geometry hypothesis). Fix a family of tripartitions $A-B(\epsilon)-C$ (“collar” of width/separation ϵ) and a family of finite-dimensional states $\rho_{ABC}(\epsilon)$. Assume there exist constants $K > 0$ and $\alpha > 0$ (independent of ϵ in the regime of interest) such that $I(A : C|B)_{\rho(\epsilon)} \leq K e^{-\alpha\epsilon}$.*

Remark 4.7 (Where A-CMI is expected/known). A-CMI is intended to capture “approximate quantum Markov” behavior in geometries where $B(\epsilon)$ shields A from C . In concrete applications one either: (i) imports a theorem establishing such decay in a stated regime (e.g. a temperature/interaction range regime), or (ii) tests it numerically/empirically in a benchmark family. This paper keeps the hypothesis explicit so that failures are interpretable as falsification rather than hidden assumption.

Theorem 4.8 (T3: exponential recoverability under exponential CMI decay). *[PROVED] (given Theorem 4.4 and Assumption 4.6). Under Assumption 4.6, for each ϵ there exists a CPTP recovery map $\mathcal{R}_{B \rightarrow BC}$ such that*

$$E_{\text{rec}}(\rho(\epsilon); \mathcal{R}) \leq c_{\text{FR}} K e^{-\alpha\epsilon}.$$

Moreover, $1 - F(\rho_{ABC}(\epsilon), (\text{id}_A \otimes \mathcal{R}_{B \rightarrow BC})(\rho_{AB}(\epsilon))) \leq c_{\text{FR}} K e^{-\alpha\epsilon}$.

Proof. Combine Theorem 4.4 with Assumption 4.6. Then apply Lemma 4.3. □

Remark 4.9 (Why T3 is still worth stating). T3 is logically simple, but it is *audit-critical*: it fixes conventions, isolates the single falsifiable hypothesis A-CMI, and cleanly separates the imported recoverability mechanism from downstream “rate–power” bridges.

Remark 4.10 (v2: the alignment step executed — $c_{\text{FR}} = 1$). *[PROVED] (v2).* In the locked conventions of Definitions 4.1 and 4.2, the Fawzi–Renner theorem [1] states $I(A : C|B) \geq -\log F$ with F the squared fidelity of the recovered state for a rotated Petz map, i.e. FR imports as Theorem 4.4 with $c_{\text{FR}} = 1$ exactly. With the *unsquared* (root) fidelity the same theorem reads $I \geq -2 \log f$, so silently mixing conventions produces a factor-2 drift — the recurring factor corrected across this series (2512.0101, 2601.0020, 2601.0034, 2601.0035, all v2). The verification suite checks the identity $-\log F = -2 \log \sqrt{f}$ and instantiates $c_{\text{FR}} = 1$ on a benchmark family (gapped transverse-field Ising collar): $E_{\text{rec}} \leq I(A : C|B)$ at every collar width in the generated dataset, with $I \approx 0.20 e^{-1.06w}$, closing the T3 chain end to end including the $1 - F$ form via Lemma 4.3. This settles the Type I anchor; the same alignment in the split-regularized Type III setting is Lemma 10.2 of the companion interface note ai.viXra:2601.0065 (v2).

5 Operational power interface (typed, not proved here)

Definition 5.1 (D6: maintenance strategies and average power). Fix an operational control model (battery, allowed controls, coupling to noise), and an admissible strategy class **Strat**. Fix an error budget δ and a noise semigroup T_t . For a target state ρ , let $P(\rho; \mathbf{S}) \in [0, \infty]$ denote the average battery power required by a strategy $\mathbf{S} \in \mathbf{Strat}$ to maintain ρ under T_t within tolerance δ (as defined by that model).

Definition 5.2 (D7: extra maintenance power). Define the extra power as

$$P_{\text{extra}}(\rho) := \max \left\{ 0, \inf_{\mathbf{S} \in \mathbf{Strat}} P(\rho; \mathbf{S}) - \inf_{\mathbf{S} \in \mathbf{Strat}} P(\Delta[\rho]; \mathbf{S}) \right\}.$$

Assumption 5.3 (BATO-LAW: maintenance inequality (operational postulate)). *[ASSUMED] (typed law to be proved or imported in a companion note). Under the hypotheses of the chosen control model, for all target states ρ one has $P_{\text{extra}}(\rho) \geq k_B T \dot{C}_{\Delta, \text{loss}}(\rho)$, whenever $\dot{C}_{\Delta, \text{loss}}(\rho)$ exists.*

Corollary 5.4 (C1: envelope consequence (directionality explicit)). *Assume Assumption 5.3. Assume Assumption 3.7 whenever needed to exclude degeneracies. Then for $\rho \in \mathcal{F}_\epsilon$,*

$$P_{\text{extra}}(\rho) \geq k_B T \kappa^\downarrow(\epsilon) C_\Delta(\rho).$$

Proof. By Assumption 5.3, $P_{\text{extra}}(\rho) \geq k_B T \dot{C}_{\Delta, \text{loss}}(\rho)$. By Definition 3.8, $\dot{C}_{\Delta, \text{loss}}(\rho) \geq \kappa^\downarrow(\epsilon) C_\Delta(\rho)$ for $\rho \in \mathcal{F}_\epsilon$. $\square$

6 Two small examples (toy but honest)

Example 6.1 (Exact Markov case ($I(A : C|B) = 0$ implies perfect recovery)). If a state satisfies $I(A : C|B)_\rho = 0$ (i.e. it is a quantum Markov chain in the exact sense), then recoverability can be perfect: there exists a CPTP recovery map $\mathcal{R}$ such that $F(\rho_{ABC}, (\text{id}_A \otimes \mathcal{R})(\rho_{AB})) = 1$, hence $E_{\text{rec}} = 0$. This anchors the interpretation of T2/T3: CMI is the static “Markovness defect” that controls recoverability; the implication $I(A : C|B)_\rho = 0 \Rightarrow$ exact recovery is standard (see e.g. [3]).

Example 6.2 (A semigroup where $\dot{C}_{\Delta, \text{loss}}$ is directly computable (envelope intuition)). Let Δ be dephasing in a fixed basis and let T_t be pure dephasing in the same basis (so T_t commutes with Δ and damps off-diagonal terms exponentially). In such cases, $C_\Delta(T_t(\rho))$ is typically monotone decreasing in t on broad families, and $\dot{C}_{\Delta, \text{loss}}(\rho)$ exists and is computable from the initial damping rate. This illustrates what $\kappa^\uparrow(\epsilon)$ and $\kappa^\downarrow(\epsilon)$ measure: worst-case and best-case *initial relative* coherence-loss rates over a typed family $\mathcal{F}_\epsilon$. (We do not claim a universal closed form for $\dot{C}_{\Delta, \text{loss}}$ in this paper; the point is operational meaning.)

Remark 6.3 (v2: both examples instantiated numerically). The verification suite instantiates Example 6.1 with a random state of Hayden–Jozsa–Petz–Winter structure $\rho = \rho_{AB_L} \otimes \rho_{B_R C}$: $I(A : C|B) = 4 \times 10^{-16}$ and Petz-type reconstruction with $1 - F = 1.3 \times 10^{-15}$. It instantiates Example 6.2 with a $d = 3$ pure-dephasing semigroup with strongly unequal pair rates: $C_\Delta(T_t \rho)$ is monotone on all sampled trajectories, $\dot{C}_{\Delta, \text{loss}} > 0$ on the family, and the empirical envelopes spread by $\kappa^\uparrow/\kappa^\downarrow \approx 100$ — a concrete reminder of the golden rule (Remark 3.10): an upper envelope never certifies a floor. The same directionality is exhibited dynamically in the Davies setting by the companion ai.viXra:2601.0064 (v2).

7 Falsification matrix (minimal)

Item	Status	Verification proxy	Concrete falsifier
A-CMI (CMI decay)	[ASSUMED]	compute/estimate $I(A : C B)$ vs ϵ	show non-decay or slower decay than claimed
T2 (FR-type)	[IMPORTED]	align c_{FR} to cited convention (done in v2: $c_{\text{FR}} = 1$)	detect convention mismatch / wrong constant
BATO-LAW	[ASSUMED]	explicit model calculation	explicit protocol violating inequality
Lower-envelope need	[PROVED]	logic in Remark 3.10 and Corollary 5.4	N/A (directionality statement)

Table 2: Minimal falsification matrix for this contract paper (re-typeset in v2).

8 Clay vs operational track (explicit separation)

Clay track (closed Hamiltonian). The Clay mass gap concerns a *closed* Yang–Mills QFT in the continuum and its spectral gap in an axiomatic setting.

Operational track (this paper). This paper treats an open-system/maintenance track: typed inequalities relating coherence loss and power. Any bridge from operational rate floors to closed Hamiltonian gaps would require additional hypotheses not provided here.

Note added in v2

No v1 definition, assumption, theorem statement or numbering is changed; new v2 material is appended at the end of the affected sections. v2: (1) repairs the cross-reference labels — in v1 every reference was typeset as “Definition $x.y$ ” regardless of environment, so e.g. “given Definitions 4.4 and 4.6” in v1 denotes Theorem 4.4 and Assumption 4.6, “apply Definition 4.3” denotes Lemma 4.3, and “Assume Definition 5.3” denotes Assumption 5.3; (2) re-typesets Tables 1 and 2 (overlapping column text in v1); (3) executes the c_{FR} alignment of Appendix A: in the locked conventions the Fawzi–Renner import gives $c_{\text{FR}} = 1$ exactly (Remark 4.10); (4) updates Appendix B: the three deferred interfaces announced in v1 now exist as series papers and are cited concretely; (5) records in Appendix C the partial progress toward BATO-LAW in ai.viXra:2601.0031 (v2); (6) adds series positioning (v1 cited no companion papers) and a verification suite (`verification/2601-0066/` in the series repository) instantiating the closed recoverability lane end to end and both examples (Remarks 4.10 and 6.3).

A Imported theorem conventions (how to set c_{FR})

To make Theorem 4.4 fully audit-aligned, pick a single source theorem statement and restate it verbatim under the conventions: squared fidelity F (Definition 4.1), error functional $E_{\text{rec}} = -\log F$ (Definition 4.2). Then set c_{FR} to the numerical constant appearing in that statement after translating conventions. Using c_{FR} prevents silent factor drift across papers. v2: *this step is executed in Remark 4.10 with source [1]: $c_{\text{FR}} = 1$ in the locked conventions (equivalently, 2 if the root-fidelity convention were used).*

B Deferred interfaces (updated in v2: now concrete)

To keep this paper self-contained and citable, we do not develop here:

- Type III/AQFT split-regularized CMI definitions and conditional expectations: *now* the interface note ai.viXra:2601.0065 (v2), with the CE-pullback direction discipline of ai.viXra:2601.0046 (v2);
- Davies/Dirichlet-form identities and RIP-U envelope bounds: *now* the dynamics note ai.viXra:2601.0064 (v2), with the Bohr-channel decomposition of ai.viXra:2601.0023 (v2);
- non-abelian benchmark implementation contracts and reproducibility harness: *now* the benchmark notes ai.viXra:2601.0043, 2601.0044, 2601.0047, 2601.0050, 2601.0051 (all v2) with suites in the series repository.

Nothing in Sections 4–6 depends on these components; the citations record where each deferred interface is developed.

C Upgrade path for BATO-LAW (audit guidance)

In this paper, BATO-LAW is treated as **[ASSUMED]**, since its validity is operational-model dependent. To upgrade Assumption 5.3 to **[IMPORTED]** or **[PROVED]** within a specified control model, it suffices to provide:

- an axiomatized specification of the control model (battery, admissible controls, error metric, and the maintenance criterion within tolerance δ);
- either a complete proof within that model or an imported theorem with hypotheses restated verbatim under the present conventions;
- an explicit identification (or inequality) relating the model’s dissipation functional to $\dot{C}_{\Delta, \text{loss}}(\rho)$ as defined in Definition 3.3.

v2 note: partial progress exists in the series: ai.viXra:2601.0031 (v2) proves a work-cost inequality in a GNS/KMS battery model (with the v1 sign error corrected there to a battery free-energy *decrease*), which supplies the third bullet’s dissipation-functional identification in that model class; a full BATO-LAW derivation remains open.

References

- [1] O. Fawzi and R. Renner, *Quantum conditional mutual information and approximate Markov chains*, Commun. Math. Phys. **340**, 575–611 (2015).
- [2] D. Sutter, M. Tomamichel, and P. Harremoës, *Strengthened monotonicity of relative entropy via pinched Petz recovery map*, IEEE Trans. Inf. Theory **62**, 2907–2913 (2016).
- [3] P. Hayden, R. Jozsa, D. Petz, and A. Winter, *Structure of states which satisfy strong subadditivity of quantum entropy with equality*, Commun. Math. Phys. **246**, 359–374 (2004).
- [4] L. Eriksson, ai.viXra:2601.0065 (2026; v2 2026).
- [5] L. Eriksson, ai.viXra:2601.0064 (2026; v2 2026).
- [6] L. Eriksson, ai.viXra:2601.0046 (2026; v2 2026).
- [7] L. Eriksson, ai.viXra:2601.0023 (2026; v2 2026).
- [8] L. Eriksson, ai.viXra:2601.0031 (2026; v2 2026).
- [9] L. Eriksson, ai.viXra:2601.0035 (2026; v2 2026).
- [10] L. Eriksson, ai.viXra:2601.0043 (2026; v2 2026).
- [11] L. Eriksson, ai.viXra:2601.0044 (2026; v2 2026).