Temporal adjoints and injectivity lengths in stationary quantum processes

Lluis Eriksson

29 October 2026
Version 2: revised proof and source conventions

Abstract

Taking local matrix adjoints along a forward random process does not reverse its temporal order. We give finite, ergodic qubit examples in which first and eventual Kraus-span injectivity lengths change under this operation. The failure persists for bistochastic, two-sided faithful local channels; a three-phase example also changes the onset distribution. These examples contradict the pointwise adjoint-length identity in Lemma A.1 of arXiv:2607.01446v1, without providing a counterexample to that manuscript's main parent-Hamiltonian or spectral-gap theorems. We state a corrected inverse-base identity. It preserves fixed-block distributions, and, under local and dual faithfulness, preserves first and eventual onset laws. We also replace an invalid common-kernel step in the source appendix by a Choi-rank convergence argument, recovering its eventual-strict-positivity/injectivity implication under the cited ergodic convergence theorem. The required normalization and endpoint shift are explicit. Fixed-block, first-occurrence, eventual, and pathwise assertions remain distinct.

1 Definitions and the assertion being tested

Let $(\Omega, \mathcal{F}, \mu, T)$ be an invertible probability-preserving system. A measurable tuple $S(\omega) = (A_1(\omega), \dots, A_g(\omega))$ defines

$$\Phi_\omega(a) = \sum_{i=1}^{g} A_i(\omega) a A_i(\omega)^*.$$

For $\sigma = (i_1, \dots, i_n)$, put

$$\Pi_n(\omega; \sigma) = A_{i_n}(T^{n-1}\omega) \cdots A_{i_1}(\omega), \qquad W_n(\omega) = \operatorname{span}_{\mathbb{C}}\{\Pi_n(\omega; \sigma) : \sigma \in [g]^n\}.$$

Thus later matrices multiply on the left. Define, with $\inf \varnothing = \infty$,

$$\begin{aligned} \ell_0(\omega) &= \inf\{n \geq 1 : W_n(\omega) = \mathbb{M}_d(\mathbb{C})\}, \\ \ell(\omega) &= \inf\{n \geq 1 : W_m(\omega) = \mathbb{M}_d(\mathbb{C}) \text{ for every } m \geq n\}. \end{aligned}$$

These are the site-count conventions in Appendix A of [1]; an interval-endpoint convention can differ by one.

Write $S^*(\omega) = (A_i(\omega)^*)_{i=1}^g$. In this notation S^* is still driven by the *forward* base T. Lemma A.1 of [1], version v1, PDF page 34, asserts equality almost surely of both $\ell_0(S)$ and $\ell(S)$ with their S^* counterparts. Its proposed reversal of Kraus labels does not reverse the time arguments:

$$\begin{aligned} \Pi_n^{S^*}(\omega; \sigma) &= A_{i_n}(T^{n-1}\omega)^* \cdots A_{i_1}(\omega)^*, \\ \Pi_n^{S}(\omega; \operatorname{rev}\sigma)^* &= A_{i_n}(\omega)^* \cdots A_{i_1}(T^{n-1}\omega)^*. \end{aligned}$$

The autonomous case is different: when the matrices do not depend on the site, label reversal does give the desired identity.

For calculations we use the unnormalized Choi matrix

$$J(\Phi) = \sum_{a,b=1}^{d} \Phi(|a\rangle\langle b|) \otimes |a\rangle\langle b|.$$

Its rank is the dimension of the Kraus span. In particular $W_n = \mathbb{M}_d(\mathbb{C})$ precisely when the Choi matrix of the corresponding block channel is positive definite.

2 Bistochastic counterexamples

Let $P_0 = |0\rangle\langle 0|$, $P_1 = |1\rangle\langle 1|$ and

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad D(a) = P_0 a P_0 + P_1 a P_1, \qquad U = \mathrm{Ad}_H\,.$$

Write $L(a) = \mathrm{Tr}(a) I_2/2$. Direct multiplication gives

$$DUD = L, \qquad D^2 = D, \qquad U^2 = \mathrm{id}\,. \tag{1}$$

All the channels in this section are trace preserving and unital, as are their Hilbert–Schmidt adjoints. A positive map Ψ is called faithful if $\Psi(a) = 0$ with $a \geq 0$ implies $a = 0$. These channels and their adjoints are faithful: for a positive map Φ in finite dimensions, faithfulness of Φ is equivalent to $\Phi^*(I) > 0$. Also $\Phi L = L$ for a unital Φ, so an exact depolarizing block persists under further forward steps.

Theorem 1. *There is an invertible ergodic two-point system of bistochastic qubit channels, with both local and dual faithfulness and finite eventual injectivity, for which*

$$(\ell_0(S), \ell(S)) = ((3,2),(3,2)), \qquad (\ell_0(S^*), \ell(S^*)) = ((2,3),(2,3)).$$

The displayed pairs list the values at phases $0, 1$.

Proof. Use the uniform measure on $\{0,1\}$ and $T(p) = p + 1 \bmod 2$. This transformation is ergodic, although it is not mixing. At phase 0 choose the Kraus tuple (P_0, P_1), and at phase 1 choose (HP_0, HP_1). The forward channels are D, UD; the pointwise adjoints on the same base are D, DU.

For the original process, phase 0 has block channels D, UD, L at lengths $1, 2, 3$. Phase 1 has UD, L at lengths $1, 2$. In the adjoint process, phase 0 has D, L, while phase 1 has DU, DU, L. A dephasing channel preceded or followed by a unitary has Choi rank 2, whereas $J(L) = I_4/2$ has rank 4. Equation (1) proves these identities exactly. Persistence of L proves the assertions for all later lengths. Both Kraus normalizations are I_2 in each phase, so all stated normalization and faithfulness hypotheses hold. □

Theorem 1 refutes the pointwise equalities, but its two onset distributions happen to agree. That observation cannot be promoted to a distributional repair on the unchanged forward base.

Theorem 2. *A uniform three-cycle of bistochastic qubit channels has*

$$\ell_0(S) = \ell(S) = (4,3,2), \qquad \ell_0(S^*) = \ell(S^*) = (2,2,2).$$

Consequently the first and eventual onset distributions can also differ.

Proof. On $\Omega = \{0, 1, 2\}$, with uniform measure and $T(p) = p + 1 \bmod 3$, take

$$\Phi_0 = D, \qquad \Phi_1 = UD, \qquad \Phi_2 = D_X := UDU.$$

The last map has Kraus operators HP_0H, HP_1H. All channels are bistochastic. Before the first depolarizing block, the original channel sequences are

$$\begin{array}{c|l} 0 & D,\ UD,\ UD,\ L \\ 1 & UD,\ UD,\ L \\ 2 & D_X,\ L. \end{array}$$

These identities follow from (1). The pointwise-adjoint local channels are D, DU, D_X. Every consecutive two-step product in this forward three-cycle is L. Each one-step channel has Choi rank 2, and L persists. Thus the stated lengths hold at every phase. In particular

$$\mu\{\ell_0(S) \le 2\} = \tfrac{1}{3}, \qquad \mu\{\ell_0(S^*) \le 2\} = 1.$$

□

3 The correctly reversed process

Let $\widehat{T} = T^{-1}$ and

$$\widehat{A}_i(\omega) = A_i(T^{-1}\omega)^*.$$

Both the inverse base and the one-step shift in this definition matter. This is an algebraic inverse-base adjoint construction; outside the bistochastic case its local maps need not be trace preserving, so it is not automatically a physical reverse channel.

Proposition 3. *For every fixed $n \ge 1$,*

$$\widehat{\Pi}_n(\omega; \sigma) = \Pi_n(T^{-n}\omega; \operatorname{rev}\sigma)^*, \qquad \widehat{W_n}(\omega) = W_n(T^{-n}\omega)^*.$$

Therefore $\dim \widehat{W_n}$ *and* $\dim W_n$ *have the same distribution. This fixed-block assertion does not require ergodicity, trace preservation, unitality, or faithfulness.*

Proof. The reversed word is

$$A_{i_n}(T^{-n}\omega)^* \cdots A_{i_1}(T^{-1}\omega)^*,$$

which is the adjoint of the indicated original word. Label reversal is a bijection on $[g]^n$, and matrix adjunction preserves the complex dimension of a span. Finally T^{-n} preserves μ. □

The shift in Proposition 3 depends on n. Put $E_n = \{W_n = \mathbb{M}_d(\mathbb{C})\}$. Then $\widehat{E}_n = T^n E_n$, but

$$\{\ell_0 \le N\} = \bigcup_{n \le N} E_n, \qquad \{\widehat{\ell}_0 \le N\} = \bigcup_{n \le N} T^n E_n.$$

Equality of the probabilities of individual events does not identify these unions, or the analogous eventual-event intersections.

Theorem 4. *Suppose Φ_ω and Φ_ω^* are faithful almost surely. Then, on a common full-measure set, $\ell = \ell_0$ for the original and genuinely reversed processes. Their first-onset laws agree, as do their eventual-onset laws, including any mass at infinity. Equality at the same base point is not asserted.*

Proof. Intersect the full-measure faithfulness set over all integer shifts. For a local Kraus tuple write

$$F = \sum_i A_i A_i^* = \Phi(I), \qquad G = \sum_i A_i^* A_i = \Phi^*(I).$$

Both are positive definite. Thus

$$\operatorname{span}_i A_i \mathbb{M}_d = \mathbb{M}_d, \qquad \operatorname{span}_i \mathbb{M}_d A_i = \mathbb{M}_d.$$

Indeed, for any Y,

$$Y = \sum_i A_i (A_i^* F^{-1} Y) = \sum_i (Y G^{-1} A_i^*) A_i.$$

This full-span persistence is the established ingredient in [1, Lemma A.3], presented here in direct Kraus algebra. In particular a full Kraus span stays full when a later map with $\Phi(I) > 0$ is added on the left. The same property applies to the reversed process, whose local maps are shifted original adjoints. Hence the events E_n and $\widehat{E}_n$ are increasing in n. First and eventual full-span times coincide, and

$$\mu\{\ell_0 \le n\} = \mu(E_n) = \mu(\widehat{E}_n) = \mu\{\widehat{\ell}_0 \le n\}.$$

These equalities determine the laws on $\mathbb{N} \cup \{\infty\}$. □

In the three-cycle of Theorem 2, genuine reversal has first and eventual lengths $(4, 2, 3)$. Its law agrees with $(4, 3, 2)$, while the values differ at phases $1, 2$. This also separates Theorem 4 from a pathwise assertion.

4 Why persistence cannot be omitted

The following example supplies an explicit obstruction to inferring onset laws from fixed-block laws alone. On a uniform two-cycle use

$$R = (E_{10}, E_{11}), \qquad B = \left(\begin{pmatrix} 1 & 0 \\ 0 & 1/2 \end{pmatrix}, \frac{\sqrt{3}}{2} E_{01} \right), \qquad E_{ab} = |a\rangle\langle b|.$$

Both local channels are trace preserving. The reset channel is $\Phi_R(a) = \operatorname{Tr}(a) E_{11}$ and its adjoint is not faithful.

Let $V = \operatorname{span}\{E_{10}, E_{11}\}$ and $Q = \operatorname{span}\{\operatorname{diag}(1, 1/2), E_{01}\}$. Direct subspace multiplication yields $QV = \mathbb{M}_2$, $V\mathbb{M}_2 = V$, and $VQ = V$. Starting in phase 0, forward spans alternate between dimension 2 at odd lengths and 4 at even lengths. Starting in phase 1, the first span has dimension 2, even lengths have dimension 2, and odd lengths at least 3 have dimension 4. Thus

$$\ell_0(S) = (2, 3), \qquad \ell(S) = (\infty, \infty).$$

For genuine reversal, the local tuples are B^*, R^* at phases $0, 1$. With $C = \operatorname{span}\{E_{01}, E_{11}\}$ one has $CQ^* = \mathbb{M}_2$, $Q^*\mathbb{M}_2 = \mathbb{M}_2$, $Q^*C = C$, and $CC = C$. It follows that

$$\widehat{\ell}_0 = \widehat{\ell} = (2, \infty).$$

At each fixed length $n \ge 2$, the original and reversed full-span probabilities are both $1/2$. Nevertheless their first-onset laws differ, and their eventual-finiteness probabilities are respectively 0 and $1/2$. This does not contradict Proposition 3: the joint behavior over lengths matters.

5 A Choi-rank convergence lemma

In this section only, the base need be neither invertible nor ergodic. All spaces have fixed finite dimension. For a complex-linear map Ψ write

$$\|\Psi\|_{1\to 1} = \sup_{\|a\|_1=1} \|\Psi(a)\|_1, \qquad \Delta_z(a) = \mathrm{Tr}(a)z.$$

The norm is taken over all complex matrices. A stationary family means that its local map at time j is $\Phi_{T^j\omega}$.

Theorem 5. *Let T preserve a probability measure and let $\omega \mapsto \Phi_\omega$ be a measurable family of completely positive maps on $\mathbb{M}_d(\mathbb{C})$. Suppose $z_\omega > 0$ is measurable, $\mathrm{Tr}\, z_\omega = 1$, and*

$$e_n(\omega) := \|\Phi_\omega^{(n)} - \Delta_{z_{T^n\omega}}\|_{1\to 1} \longrightarrow 0 \quad \textit{almost surely}, \qquad \Phi_\omega^{(n)} = \Phi_{T^{n-1}\omega}\cdots\Phi_\omega. \tag{2}$$

Then there is a measurable $N(\omega) < \infty$ almost surely such that $W_n(\omega) = \mathbb{M}_d(\mathbb{C})$ for every $n \ge N(\omega)$. No exponential rate or deterministic positive lower bound for z_ω is required.

Proof. Put $f(\omega) = \lambda_{\min}(z_\omega) > 0$. With the output-first Choi convention fixed above,

$$J(\Delta_z) = z \otimes I_d, \qquad \|J(\Psi)\|_\infty \le d^2\|\Psi\|_{1\to 1}.$$

The second inequality follows by expanding in matrix units, using $\|E_{ab}\|_1 = 1$, and applying the triangle inequality; sharpness of the constant is irrelevant. Thus

$$a_n(\omega) := \|J(\Phi_\omega^{(n)}) - z_{T^n\omega} \otimes I_d\|_\infty \le d^2 e_n(\omega) \longrightarrow 0.$$

We first derive the local condition needed for persistence. Let

$$B = \{\omega : \Phi_\omega(I_d) \text{ is singular}\}, \qquad g(\omega) = 1_B(\omega) f(T\omega).$$

If $\Phi_\theta(I_d)$ is singular, every local Kraus matrix at θ has range in the same proper output subspace: a vector in the kernel of $\sum_i A_i A_i^*$ is annihilated by every A_i^*. Consequently any composition ending in this map has a singular Choi matrix, regardless of its earlier prefix. A unit vector in the Choi kernel gives

$$0 \le g(T^{n-1}\omega) \le a_n(\omega).$$

This pointwise norm bound requires no measurable selection of kernel vectors. Its left-hand side converges to zero almost surely and hence in probability. By stationarity, for each $\varepsilon > 0$,

$$\mu\{g > \varepsilon\} = \mu\{g \circ T^{n-1} > \varepsilon\} \longrightarrow 0.$$

Therefore $g = 0$ almost surely. Since $f \circ T > 0$ almost surely, $\mu(B) = 0$, so $\Phi_\omega(I_d) > 0$ almost surely. A countable intersection of forward pullbacks makes this true at every later site of almost every trajectory.

For each $k \ge 1$ set $A_k = \{f \ge 1/k\}$. These sets cover a full-measure set. Poincaré recurrence says that almost every point of each A_k returns to it infinitely often. Removing the countable union of exceptional sets, each remaining ω therefore has some $k = k(\omega)$ and arbitrarily large n for which $f(T^n\omega) \ge 1/k$. Choose such an n after $a_n < 1/(2k)$. Then

$$J(\Phi_\omega^{(n)}) \ge \big(f(T^n\omega) - a_n(\omega)\big) I_{d^2} > 0.$$

The random recurrent floor is sufficient; it need not be uniform in ω or in all times.

Finally, if $J(\Phi_\omega^{(n)}) \ge \eta I_{d^2} > 0$, complete positivity gives

$$J(\Phi_\omega^{(n+1)}) = (\Phi_{T^n\omega} \otimes \mathrm{id}) J(\Phi_\omega^{(n)}) \ge \eta \Phi_{T^n\omega}(I_d) \otimes I_d > 0.$$

Induction proves persistence. The first index with positive smallest Choi eigenvalue is a measurable finite N, because every finite-time matrix is measurable. All full-measure intersections used are countable. The same proof applies when $d = 1$. □

Remark 6. *The CP wording in Theorem 5 is not an extension to arbitrary normalized, non-trace-preserving cocycles. Its strong* unnormalized *hypothesis already forces local trace preservation. On a common set where* (2) *holds at* ω *and* $T\omega$*, the identity* $\Phi_\omega^{(n+1)} = \Phi_{T\omega}^{(n)}\Phi_\omega$ *and* $\operatorname{Tr} z = 1$ *imply, for every* a*,*

$$|\operatorname{Tr}\Phi_\omega(a) - \operatorname{Tr} a| \le \big(e_n(T\omega)\|\Phi_\omega\|_{1\to1} + e_{n+1}(\omega)\big)\|a\|_1.$$

Letting $n \to \infty$ *proves the claim. In particular, normalization of a general CP product cannot be discarded when invoking* (2)*.*

The finite-dimensional faithful-limit mechanism has an autonomous antecedent in [3, Proposition 3(c)⇒(b)]. The role of stationarity and recurrence here is to handle a moving positive state without a deterministic spectral floor. This argument supplies a scoped proof repair, not a claim of a new general mixing theorem.

6 Repairing the eventual-strict-positivity implication

Return to an invertible ergodic probability-preserving base and suppose the local maps are trace preserving. Define

$$q(\omega) = \inf\{n \ge 1 : \Phi_\omega^{(m)}(a) > 0 \text{ for every } m \ge n \text{ and every } 0 \ne a \ge 0\}.$$

The following is the statement of [1, Proposition A.9], with a replacement proof of the nontrivial implication.

Corollary 7. *Under these hypotheses,* $q < \infty$ *almost surely if and only if* $\ell < \infty$ *almost surely.*

Proof. If $J(\Phi) > 0$, then Φ is strictly positive: testing the partial-trace formula against rank-one input and output vectors gives a strictly positive quadratic form. Thus $q \le \ell$ pointwise, and one direction follows.

For the other direction assume $q < \infty$ almost surely. This is the eventual strict positivity assumption used in [2]. We invoke its Theorems 1 and 2 under their standing invertible, ergodic hypotheses; we do not reprove those theorems. The endpoint convention and normalization are as follows. Write the local map at site j as $\phi_j = \Phi_{T^j\omega}$. Then our length-$n$ block is their $\Psi_{n-1,0}$. For trace-preserving maps,

$$\operatorname{Tr}[\Psi_{n-1,0}^*(I_d)] = d, \qquad z_0' = I_d/d, \qquad P_{n-1,0}(a) = \frac{\operatorname{Tr} a}{d}\, z^{\mathrm{MS}}_{T^{n-1}\omega}.$$

Here $z^{\mathrm{MS}} > 0$ is their stationary density field. Put $\widetilde{z}_\omega = z^{\mathrm{MS}}_{T^{-1}\omega}$. Multiplying their normalized estimate by d, for $n \ge 2$ with source endpoints $a = 0$, $b = n-1$, $x = 0$, yields

$$\|\Phi_\omega^{(n)} - \Delta_{\widetilde{z}_{T^n\omega}}\|_{1\to1} \le dC_{r,0}(\omega)r^{n-1}, \qquad 0 < r < 1, \tag{3}$$

with $C_{r,0} < \infty$ almost surely. The initial length does not affect convergence. Theorem 5 applies and gives $\ell < \infty$ almost surely. □

The original forward proof on PDF page 37 of [1] infers a nonzero common right-kernel projection from failure of the Kraus span to fill $\mathbb{M}_d(\mathbb{C})$. That inference is invalid: the single Kraus matrix I_2 has a proper span but zero kernel. In fact a trace-preserving tuple satisfies $\sum_i K_i^*K_i = I_d$, so it cannot have a nonzero common right-kernel vector. The proof above instead uses singularity of the *Choi matrix* and convergence to a faithful replacement. It avoids both that kernel inference and the route through Lemmas A.1, A.8 and A.6 of [1]. Their downstream uses of A.1 are therefore not relied upon in this proof. It does not assert that the statement of A.9 was false.

7 Scope, reproducibility and review context

The counterexamples invalidate the specific local-adjoint identity in [1, Lemma A.1]. Corollary 7 repairs a separate appendix proof using the cited convergence theorem with its actual assumptions. Neither correction refutes or certifies the main parent-Hamiltonian construction, spectral-gap criteria, or a classification of gapped ergodic matrix product states. Inspection of selected primary statements and proof passages is not a whole-paper audit of either cited source.

The all-length conclusions rest on algebraic proofs, not extrapolation from finite ranks. The earlier scripts are:

`check_bistochastic1221.py`
`check_three_phase1222.py`
`check_reverse1223.py`.

They are fixed in PhysMath Knowledge Tree at the immutable commit and directory

`de24e4e461fb35bda8336cf248fe85bb5c3e2f61`
`integrations/research-radar/evidence/2026-10-29-temporal-audit-context/`.

Copy them to a writable scratch directory and run each with Python; the recorded environment is Python 3.12.14 with SymPy 1.14.0. The version-specific source bundle also supplies `check_repair1237_pure.py`, a self-contained extraction of the original repair-control prefix. It checks singular-last-output support and the nonstationary-target/nonfaithful-limit boundary examples. The earlier combined script's unrelated prior-task-reading suffix was excluded by the reviewer; that packaging defect and the original frozen script remain recorded, rather than retrospectively altered. These controls do not prove the general convergence theorem. No Lean formalization or certified worldwide originality is claimed.

Model assessments and review context. This revision adds the scoped proof repair in Sections 5–6; it is a different PDF from the published first version. Two unscored component audits examined, respectively, the counterexamples/reversal and the conditional Choi-rank repair. They are not assessments of this complete revised manuscript. The latter audit retained the dimension/index matching, the trace-preservation consequence of the convergence hypothesis, the control-script packaging defect, known autonomous antecedents and provenance limitations. Original reports and all objections are preserved. Requested configuration was GPT-6 Astra with xhigh effort in separately initiated contexts; effective model identity, effort, runtime isolation, prior model-level participation and independence were unattested. The filesystem was shared. Any full assessment of this frozen revised PDF must have its own exact-hash public record and full report; no score from a component audit or an earlier PDF transfers to it.

AI and responsibility disclosure. Lluis Eriksson directed OpenAI Codex assistance in mathematical development, writing and exact computation and is responsible for the manuscript. The audits were AI-assisted under the limitations stated above. No independent human review, full formal verification, guaranteed priority, predicted assessment score or editorial admission is claimed.

References

[1] O. Ekblad, E. Moreno-Nadales, E. B. Roon and J. H. Schenker, *Parent Hamiltonians of Ergodic Matrix Product States*, arXiv:2607.01446v1, 1 July 2026. `https://arxiv.org/abs/2607.01446v1`.

[2] R. Movassagh and J. Schenker, *An ergodic theorem for quantum processes with applications to matrix product states*, Communications in Mathematical Physics **395** (2022), 1175–1196. Version consulted: `https://arxiv.org/abs/1909.11769v4`.

[3] M. Sanz, D. Pérez-García, M. M. Wolf and J. I. Cirac, *A quantum version of Wielandt's inequality*, IEEE Transactions on Information Theory **56** (2010), 4668–4673. Version consulted: `https://arxiv.org/abs/0909.5347v2`.