The commutant of an instantaneous algebra need not be the redundancy of the process it generates

A finite exactly computed counterexample, with an operational witness

---

S. Tobias Croydon-McRae · Public working paper v1 · 9 September 2026

Scientific source: `bb3a22133cc6a12b6cb8ff026c2e1b2e56ed9526`. Non-author AI-agent review: `a419da15d17e4f1ddd933ba44a1e3ea0b0f17ca0`.

This publication edition incorporates the review's two editorial corrections and updates publication metadata, related-work context and review status. The original source and an exact record of the changes accompany this release. The finite-model claims received review within the research programme; external peer review and a comprehensive prior-art assessment remain open.

Abstract

In a quantum system with a distinguished subsystem ("register"), it is natural to describe a sector by the algebra of observables available at an instant, and to treat the commutant of that algebra as the redundancy — the directions in state space that carry no physical information. We exhibit a small, exactly computable model in which this inference fails.

For a six-spin chain with a two-state register restricted to the total singlet sector, the selected instantaneous algebra has a five-dimensional commutant. We show that this commutant is **not** the redundancy of the continuing process: under a declared class of measurement histories, the operationally invisible subspace is trivial. Concretely, an intervention invisible to the instantaneous algebra changes the statistics of a later permitted measurement.

We compute the span of realisable history effects exactly, in rational arithmetic, and find it is all of `Herm(C^10)`. We further find that repeated measurement of the **single binary register**, interleaved with free evolution, is informationally complete for the ten-dimensional system, and that this collapses to a two-dimensional span when the register coupling is removed.

Applied to a spatial cut, the same machinery shows that an interface which reproduces the correct Hilbert-space dimension can still discard information the process requires: charge-neutral matching across a `123|456` cut loses exactly 32 of 100 Hermitian directions, and we give two explicit density matrices with identical neutral data whose measurement statistics separate at first order in time with coefficient `1/40`.

Finally we report four **refuted** attempts to characterise which quotients are admissible, and identify the correct form of the question: the span is always an invariant direct summand, so admissibility is a reachability property and is not determined by the generated operator algebra. We give a two-state counterexample proving the algebra dimension is insufficient in general.

Separating two notions that are easily conflated then settles the matter for this model. Across 26 candidate gradings, **17** yield a modified law that remains informationally complete, while **none** admits an exact descent of the original dynamics; we prove that the kernel of a dephasing is invariant under `ad_H` precisely when `H` is block diagonal for that grading, which the model's Hamiltonian is for no declared cut. Independently, because the history effects already span `Herm(C^10)` and the trace form is nondegenerate, the operational equivalence kernel vanishes, so no quotient of the operator space with nonzero kernel preserves every declared history probability — on the normalised state domain the corresponding condition is that the kernel contain no traceless direction, which the dephasing kernels at issue do. A full-rank reconstruction of a dephased law is not a derived physical quotient.

We claim no new general mechanism. The framework is standard operational quantum theory; the contribution is a sharp, fully exact worked example and a set of negative results.

---

1. The point, informally

Suppose you have identified the observables accessible in some sector of a quantum system at a given instant, forming an algebra `A`. It is tempting to conclude that the commutant `A'` is physically redundant, and may be quotiented.

Two distinct, true statements have to be kept apart here.

- **Invisible differences at an instant** are the trace annihilator `A^⊥`, not the commutant. Two states are indistinguishable by `A` exactly when `ρ − σ ∈ A^⊥`.
- **Conjugation by a unitary in `A'`** leaves every instantaneous `A` expectation unchanged. This is the true statement about the commutant, and it is about a group action, not about additive state differences.

Identifying the two is a genuine error, not a slip of wording. In the very model below, with `P` the charge-rank-one projector on the five-dimensional DOWN space and `D = K12 + K23`,

```
P = (4/3) D (D + I) ∈ A,      Δ = P − I/5 ∈ A',      Tr Δ = 0,
ρ± = I/5 ± Δ/5      (both positive: eigenvalues {9,4,4,4,4}/25 and {1,6,6,6,6}/25)
Tr( P (ρ_+ − ρ_-) ) = 8/25 ≠ 0
```

so a state difference lying **inside** `A'` is already visible to `A` at the instant. Here `dim A = dim A' = 5` while `dim A^⊥ = 20`.

Our result concerns the second statement. An intervention that acts by conjugation within `A'` — invisible to `A` now — can be made visible later. If the system continues to evolve and further measurements are permitted, the commutant of the instantaneous algebra and the redundancy of the process are different objects, and they can differ maximally.

That statement is not surprising in the abstract. What we provide is a concrete finite model where it is exactly computable, where the gap is complete rather than marginal, and where the distinguishing measurement is explicitly constructed.

The practical relevance is to any programme that identifies a gauge or redundancy by inspecting an instantaneous structure — a stabiliser, a commutant, a symmetry of one Hamiltonian — and then quotients by it.

---

2. Setup

2.1 Model

Six spin-1/2 sites and a two-state register. Write `D_ij = S_i · S_j` and

```
D = D_12 + D_23
V = D_34 + D_45 + D_56
A = D_12

H(lambda) = [[ D + V , lambda A ],
             [ lambda A ,   D    ]]        in the (UP, DOWN) register basis
```

All terms commute with total matter spin. We work in the **total singlet sector**, which is five-dimensional; with the register the state space is `C^10`, and `dim_R Herm(C^10) = 100`. The main value is `lambda = 1`; `lambda = 0` is a conserved-register control.

Publication scope note. The register controls the extended right-region term `V`, so the interaction graph of this declared Hamiltonian includes a shared register hub. The `123|456` cut below is a partition of the matter sites; the finite sufficiency results do not establish a nearest-neighbour spatial Hamiltonian, a causal-cone bound or a geometric area law.

2.2 Declared process class

A *history* is any finite alternating sequence of free evolution and projective (Lüders) measurements drawn from a declared instrument set. For a history `h` with Kraus operator `K_h`, the realised effect is `E_h = K_h† K_h`. The observable content of the process class is

```
O = span_R { E_h : h an allowed history }
```

and the operationally invisible directions are its annihilator under the trace pairing,

```
N = { X : Tr(E X) = 0 for all E in O }.
```

Two states are operationally indistinguishable under the process class exactly when their difference lies in `N`.

`O` is the smallest real subspace containing `I` and closed under

```
E -> i[H, E]          (free evolution)
E -> P E P            (measurement backaction, one map per outcome projector)
```

This characterisation is standard and we claim nothing for it; see §7.

2.3 Instruments

The declared instrument set is the register measurement `{P_UP, P_DOWN}` together with the spectral projectors of `D_12` and `D_23` — the generators of the DOWN-block algebra.

---

3. Result 1 — the instantaneous commutant is not the process redundancy

The selected instantaneous DOWN algebra has a five-dimensional commutant. If that commutant were the redundancy of the process, `N` restricted to the DOWN domain would be correspondingly large.

It is not. Computing `O` exactly:

```
with register        dim O = 100 = dim_R Herm(C^10)      so  N = 0
on the DOWN domain   dim O =  25 = dim_R Herm(C^5)       so  N = 0
```

The invisible subspace is **trivial**. Two density matrices giving identical probabilities for every allowed history are equal.

3.1 The operational counterexample

The dimension count alone might be dismissed as an artifact of allowing too large a measurement class. It is not, and the following exhibits the failure directly.

Take the readout `O = D_23` (active) and the intervention `B = D_45` (idle). `B` commutes with every element of the compressed DOWN algebra, and both stay inside the total singlet sector, so the intervention is invisible to the instantaneous description. Expanding the restricted response in time:

```
order        0      1      2      3        4
response     0      0      0    nonzero  nonzero
```

The response vanishes identically at orders 0, 1, 2 and is **nonzero at order 3**. An intervention invisible to the compressed algebra changes the statistics of a later permitted readout. In the five-dimensional singlet basis the order-three restricted response matrix is

```
[    0      0      0   -33/16   27/16 ]
[    0      0   33/16      0   -21/16 ]
[    0   33/16      0       0    -3/4 ]
[-33/16     0      0       0     3/2  ]
[ 27/16 -21/16   -3/4    3/2     -3/2  ]
```

with witness expectation `-3/16` on a normalised valence-bond state.

A separate preregistered prediction for this model, fixed before computation, was that the singlet expectation of the `t^3` response coefficient would be `i·9/128`. It was reproduced exactly.

---

4. Result 2 — informational completeness from one binary register

With the same declared class, we find:

instrument set	coupling	`dim O` (full)	on DOWN
register + two pair measurements	`lambda = 1`	**100**	**25**
register measurement only	`lambda = 1`	**100**	**25**
register + two pair measurements	`lambda = 0`	30	5
register measurement only	`lambda = 0`	**2**	—

Repeated measurement of a **single binary register**, with free evolution between measurements, is informationally complete for the ten-dimensional system. Removing the register coupling collapses the same protocol from 100 to 2, so the completeness is a property of the coupled law and not of merely possessing a register.

We note the `lambda = 0` figure of 30 is **instrument-dependent**: across choices of the two pair observables it takes values in `{29, 30, 32, 34, 37, 38, 42, 50}`, with 30 attained by the DOWN generators `D_12` and `D_23`. Any statement of that control must name its instruments.

---

5. Result 3 — a dimension-correct interface can still lose the process

Apply the same machinery to the spatial cut `123|456`. Each three-spin region carries spin-3/2 and spin-1/2 sectors, and matching charges across the cut reproduces the familiar singlet multiplicity `1 + 2·2 = 5`. The Hilbert-space dimension is therefore correct.

The process is not.

	neutral charge-matched	charged reconstruction
with register	**68**	**100**
on DOWN	**17**	**25**

The neutral interface loses exactly **32** Hermitian directions, and **8** on the DOWN domain. These ceilings are structural: once the dephased dynamics and every instrument preserve the charge grading, the span cannot leave the block-diagonal subspace, of dimension $2^2 + 8^2 = 68$ and $1^2 + 4^2 = 17$. The substantive content is that the closure *attains* them.

5.1 An explicit operational witness

Let P project onto the matched spin-3/2 sector, $Q = I - P$, and let $E = Q_{23}$ be the pair-23 singlet outcome. With $E_1 = i[H,E]$ and

```
X = P E_1 Q + Q E_1 P
```

we find $\operatorname{rank} X = 2$ and $\operatorname{Tr} X^2 = 1/3$, and the exact operator identity $X^3 = (1/6)X$ holds, fixing the nonzero eigenvalues at $\pm 1/\sqrt{6}$. Hence

```
rho_± = I/10 ± (3/80) X
```

are positive normalised states. They have **identical neutral data** — $PXP + QXQ = 0$ exactly, so every charge-preserving observable assigns them the same value — yet measuring pair-23 after free evolution gives

```
p_+(t) − p_−(t) = t/40 + O(t²).
```

The lost information is not abstract: it names two states the permitted future process distinguishes at first order.

A caution on counting evidence. $\operatorname{Tr} X^2 = 1/3$ and the $t/40$ coefficient are **not independent** confirmations. Work throughout with the Hermitian

```
E₁ = i[H, E],       X = offblock(E₁),
```

so that X is the off-block part of a Hermitian operator and every trace below is real. Only the off-block part of E_1 contributes, giving $\operatorname{Tr}(X E_1) = \operatorname{Tr}(X^2) = 1/3$ and hence the derivative $(3/40) \operatorname{Tr}(X E_1) = (3/40)(1/3) = 1/40$. They are the same fact.

An earlier version of this paragraph wrote the same chain with the anti-Hermitian $C = [H,E]$ in place of E_1. That is a convention error: the exact value is $\operatorname{Tr}(XC) = -i/3$, not $\operatorname{Tr}(X^2)$. The two density matrices, the positivity argument and the first-order witness are unaffected — only this counting argument had to be repaired.

---

6. Result 4 — what admissibility is not, and what it is

6.0 Three questions that must not be merged

Everything in this section turns on keeping three different questions apart. The literature's word *reduction*, and our own earlier drafts, used one word for all three.

sense	question	where it is answered
exact descent	does the quotient's kernel stay invariant under the **original** dynamics and instruments?	§6.4, Theorem 3
informational completeness of a modified law	after replacing `H` by the dephased `R(H)`, do the declared instruments span everything?	§6.5, the `S` column
operational equivalence	do two distinct states ever give identical statistics for every declared history?	§6.6, Corollary 4

They are genuinely different: on this model the second has **17** positive answers among 26 candidate quotients while the first has **none**, and the third is settled outright by Result 2. A full-rank reconstruction of a dephased law is **not** a derived physical quotient, and neither is a sufficient boundary datum.

Which quotients may be taken without losing the process? We report four refuted characterisations, because the negative results constrain the answer more than any of the positive attempts did.

refuted claim	counterexample
admissible iff the interface carries charge	4 of 8 gradings lose nothing, including the interface bond `pair34`
admissible iff some declared instrument breaks the grading	`pair36` breaks it and still spans only 52
the slack confinement ceiling proves no syntactic criterion exists	refutes one criterion only; does not exclude an algebraic one
`dim O` = dimension of the generated `*`-algebra	compatible with seven cases of this model; **refuted in general** (§6.2)

6.1 What does hold

Proposition. `N` is invariant under the generating maps, so `Herm = O ⊕ N` with both summands invariant.

Proof. For `X ∈ N` and `E ∈ O`, `Tr(i[H,X]E) = −Tr(X·i[H,E]) = 0` since `i[H,E] ∈ O`; and `Tr(PXP·E) = Tr(X·PEP) = 0` since `PEP ∈ O`. ∎

So `O = A·I`, the orbit of the identity under the algebra of superoperators, and a shortfall is a **decomposition**, not a deficiency. Observational sufficiency takes the form:

> **(observational sufficiency)** a quotient `R` is observationally sufficient exactly when `ker R ⊆ N`

which is a reachability statement. The label matters: observational sufficiency is strictly weaker than **exact dynamical descent**, which additionally requires `ker R` to be invariant under the original process maps (§6.4). The word *admissible* has been used for both in earlier drafts; it is not used unqualified anywhere below. Every refuted criterion above attempted to read a decomposition off the generators, where it is not visible.

Two earlier statements of this condition are kept visible here rather than silently replaced, because both were wrong and the second was wrong in a subtler way than the first.

The **first** draft said "a union of invariant components not containing `I`". Ambiguous at best: if `V = ⊕ V_i` and `I` has a nonzero component in `V_i`, then `V_i` does not *contain* `I`, yet discarding it destroys history effects.

The **second** replaced that by identifying `N` with the sum of exactly those invariant components on which `I` has zero component. That is still too broad. It requires a decomposition adapted to `O ⊕ N`, and fails for an arbitrary invariant decomposition — even into irreducibles. Take a qubit with `H = 0` and sole projector `I`. The Hamiltonian commutator is the zero map, while compression by `I` is the identity map, so `O = span{I}` and `dim N = 3`. Decompose `Herm` into the invariant lines `E00, E11, σx, σy`. Then `I` has nonzero projection on **both** diagonal lines, yet the invisible vector `E00 − E11` does not lie in `span{σx, σy}`, the sum of the zero-projection lines. Equivalent invariant summands can carry correlated seed components.

What is correct, and all that is used below: `N = O^⊥`, and a quotient `R` is observationally sufficient exactly when `ker R ⊆ N`. No identification of `N` with a distinguished set of components of an arbitrary decomposition is claimed.

Verified by seeding: in the charge-dephased case a seed inside the off-block 32 generates exactly 32 and the identity generates 68, giving `68 + 32 = 100` with both components cyclic. The register-dephased case fragments further — a DOWN-block seed reaches only 12 — so component count measures how mixing the declared process class is.

6.2 A two-state counterexample to the algebraic criterion

Take a single qubit, `H = σ_z`, instrument the `σ_x` projectors. The generated `*`-algebra is all of `M_2` (dimension 4, irreducible), but

```
dim O = 3,     O = span{ I, σ_x, σ_y },     N = span{ σ_z }.
```

`σ_z` is never produced: `PEP` with a single projector can only destroy coherence in the `P` eigenbasis, never create it from `I`. So the generated algebra dimension does not determine `dim O`.

6.3 Two mechanisms of invisibility

The qubit's invisible direction is **conserved** and **killed by every compression**. Testing whether that characterises `N` in general:

```
S = { X : [H,X] = 0 and P_k X P_k = 0 for all k }

case                span    dim N   dim S
full, lambda = 1     100      0       0     (control)
qubit                  3      1       1     S = N
charge-dephased       68     32       0
register-dephased     30     70       0
```

`S = {0}` in both six-spin cases. Operational invisibility therefore has at least two distinct mechanisms — *static* (conserved and measurement-orthogonal) and *dynamical* (unreachable from the seed) — and the six-spin model is entirely the latter.

This refutes the specific **"conserved and killed by every compression"** proxy as a characterisation of `N`, and nothing more general. It does **not** show that no criterion phrased in terms of the generators can succeed: this manuscript already uses one that does, namely the generator-word orbit `O = span{gI}` followed by its trace annihilator. An earlier draft overstated this paragraph in exactly that way; the overstatement is retracted here and the claim restricted to the tested proxy.

6.4 Theorem 3 — exact descent, and why every declared quotient fails it

For a two-block grading `{Π, Q = 1 − Π}` write `R(E) = ΠEΠ + QEQ`, so `ker R` is the space of off-block coherences.

Theorem 3. `ker R` is invariant under `ad_H` **if and only if** `H` is block diagonal for the grading, i.e. `ΠHQ = 0`.

Proof. If `ΠHQ = 0` then `H = ΠHΠ + QHQ`, and the commutator of a block diagonal operator with an off-block operator is off-block, so `ad_H(ker R) ⊆ ker R`. Conversely suppose `ΠHQ ≠ 0` and take `X = i[Π, H]`, which is off-block, so `X ∈ ker R` and `R(X) = 0`. Then

```
R(i[H, X]) = 2 (Π H Q H Π − Q H Π H Q),
```

which is nonzero whenever `ΠHQ ≠ 0`, since its first term is `2 (ΠHQ)(ΠHQ)^†` restricted to the `Π` block and is positive semidefinite of the same rank as `ΠHQ`. So `ad_H` carries an element of `ker R` outside `ker R`. ∎

This is a theorem, not a pattern that happened to hold — a distinction that matters here, because four of the characterisations tabulated above also held on every case we tried before dying. It was additionally verified as an exact equivalence on all 26 candidate gradings.

For the charge grading the obstruction is exactly

```
rank R(i[H,X]) = 2,   Tr R(i[H,X]) = 0,   Tr [R(i[H,X])]^2 = 32/81,
```

computed independently along two coordinate paths.

Consequence. The original `H = [[D+V, A],[A, D]]` is block diagonal for **none** of the 26 declared gradings: the `A` blocks cross the register grading, and `K34` — the interface bond that carries the sufficiency result of §5 — is the only term crossing the charge grading. So *no declared quotient descends exactly*, and the reason is structural rather than subtle.

6.5 The two columns, side by side

Computed on the same 26 gradings, same model, same declared instruments. `D` is exact descent under the **original** `H`; `S` is the span under the **dephased** `R(H)`; `core` is the largest invariant subspace inside `ker R` under the original law.

```
gradings                      kerR  D:admissible  core  S:full  S:down  S:admissible
pair12..pair35 (11 rows)        48         False     0     100      25          True
pair36, 45, 46, 56              48         False     0      52      13         False
region123, region456            32         False     0      68      17         False
region12, 23, 34, 16            48         False     0     100      25          True
region45, region56              48         False     0      52      13         False
region124, region235            32         False     0     100      25          True
register                        50         False     0      30       5         False

totals                            0 of 26 descend      17 of 26 complete after dephasing
```

The columns disagree on exactly the 17 rows that the `S` column calls admissible. Reading an `S` entry as a statement about the original process is the error this section exists to prevent.

That the `core` column is uniformly zero is only meaningful because the same routine returns the **full 32-dimensional** kernel when applied to the dephased charge law, where a nonzero core is known to exist.

6.6 Corollary 4 — informational completeness already settles it

The lattice of invariant subspaces is **not** needed to rule out nontrivial quotients of the original process. Result 2 gives `O = Herm(C^10)`, and the trace form is nondegenerate on `Herm` (verified exactly: the Gram matrix of the trace form in our rational basis has rank 100). Hence

```
N = { X : Tr(X E) = 0 for all E in O } = 0,
```

so no two distinct states are operationally equivalent.

Scope — what "quotient" must mean here. The conclusion is a statement about a quotient of the **full Hermitian operator space** that must preserve every linear effect functional. Read that way, `ker Q ≠ 0` forces the loss of some declared history probability, for any candidate redundancy whatsoever and not merely the 26 tabulated ones. Read as a statement about *arbitrary linear representations of the normalised density states* it would be **false**, and the counterexample is one line:

```
T(X) = X − Tr(X) I / d      has   ker T = span{I} ≠ 0,
yet   ρ = T(ρ) + I/d        so T loses nothing about trace−one states.
```

The trace is already known, so discarding it costs nothing on the state domain. On that domain the correct condition is `ker Q ∩ {traceless Hermitian} = 0`, where the traceless subspace has dimension 99.

This does **not** weaken the charge-dephasing result. For a dephasing by orthogonal projectors, `Tr(ΠXQ) = Tr(QΠX) = 0` because `QΠ = 0`, so `ker R` lies entirely inside the traceless subspace — verified exactly, all 32 dimensions of the charge kernel are traceless. The qualification changes the wording, not the conclusion.

This corollary was pointed out by an external collaborator against an earlier, over-cautious statement of ours which held that the question could not be settled without computing the commutant. The same collaborator then supplied the scope qualification above, against our first, too-broad statement of the corollary itself. The commutant remains the right tool for a different question: classifying the dynamically invariant decompositions of the **modified** laws, where `N` is nonzero — 32-dimensional for the dephased charge law.

Note the division of labour between Theorem 3 and Corollary 4. Corollary 4 says no nontrivial quotient preserves all history probabilities. Theorem 3 says more: even the quotients one might hope to justify dynamically fail at the level of the generator, and identifies exactly which term obstructs each one.

---

7. Relation to existing work

This public working paper claims an exact worked example and bounded negative results, not a new general mechanism or priority over existing tomography and observability results. This section is publication editorial context; it is not a claim that the programme's literature review is complete.

The relevant established ideas include finite-dimensional quantum observability, sequential tomography, reachable operator spaces and quantum sufficient statistics. The following primary sources are concrete starting points:

- Xinhua Peng, Jiangfeng Du and Dieter Suter, *Measuring complete quantum states with a single observable*, Physical Review A 76, 042117 (2007). Their assistant-coupled protocol is related single-observable tomography; it is not the identical fixed binary-register history protocol tested here.
- Pengcheng Yang, Min Yu, Ralf Betzholz, Christian Arenz and Jianming Cai, *Complete Quantum-State Tomography with a Local Random Field*, Physical Review Letters 124, 010405 (2020). This provides a related route from local control and measurement to full-state tomography.
- *Operational Quantum Mereology and Minimal Scrambling* is related work on operational subsystem structure. The present computation does not establish a refutation of that paper.

These references and their stated connections were checked against the primary records for this edition. Full-text comparison across quantum reference frames, process tensors and combs, control and observability, and channel sufficiency remains necessary before claiming novelty or making a journal submission.

---

8. Reproducibility

All arithmetic is exact and rational. No floating point and no eigensolver appears in any reported identity; spectra are obtained from exact operator identities (e.g. `X³ = (1/6)X`, `χ² = (3/16)P_{j=1/2}`). Every reported rank is computed twice by two independently written elimination engines, one of which is order-independent by construction.

```
python3 verify_history_effect_span.py     # spans; 9/9
python3 verify_compose_reduce_claims.py   # cut, witness, t/40; 17/17
python3 verify_reduce_glue.py             # invariant subspaces; 13/13
python3 quotient_family_scan.py           # the eight-grading table
python3 criterion_stress_test.py          # 26 gradings; refutes a criterion
python3 independent_check_repair.py       # qubit counterexample; 8/8
python3 invisible_space_probe.py          # the two mechanisms
python3 seed_dependence_probe.py          # cyclicity by seed
python3 verify_invariant_decomposition.py # descent vs modified law; 14/14
python3 verify_descent_corollary.py       # equivalence kernel is zero; 7/7
python3 verify_review_repairs.py          # review counterexamples E1,E2; 7/7
```

The author commands above use the Python standard library. The review verified that committed preregistrations precede the committed implementations. The separate independent reconstruction uses SymPy and python-flint; its pinned requirements and step-by-step commands are in the public reproduction packet.

Errors made and corrected are preserved in the repository, including a run invalidated because a plain transpose was used where the non-orthogonal basis requires the Gram adjoint $X\# = G^{-1}X^{T}G$. Ranks and traces are basis-independent; the adjoint is not. A second invalidated run is preserved for §6.5: the routine computing the largest invariant subspace inside `ker R` initially kept only those *basis vectors* whose image stayed inside, which is basis-dependent and could not have detected a nonzero answer. Both the invalid run and the positive control added alongside the repair — the same routine returns the full 32-dimensional kernel on the dephased charge law — are in the repository.

---

9. What is not claimed

- **No new general mechanism.** §2.2 is textbook.
- **Nothing about quantum gravity, geometry, a continuum limit, or emergent spacetime.** The model is a six-spin chain.
- **The process class is declared, not derived.** Which histories are physically admissible is an input. Results state what *these* ingredients imply.
- **No claim that the charged interface is minimal**, or uniquely selected.
- **No complete classification of invariant subspaces.** §6.5 tests 26 declared gradings; §6.6 settles operational equivalence without the lattice, but the lattice of the modified laws, where `N` is nonzero, is not computed.
- **Nothing about POVMs or ancilla-assisted instruments.** Every statement is restricted to the declared sharp Lüders instruments.
- **The measurement-update rule is assumed, not derived.** Two instruments can share their outcome effects and differ in their post-measurement states; which one the relational principles select — if any — is open, and is the natural next study.
- **Review scope.** Claims A–E received non-author AI-agent review within the research programme at the frozen source revision above, including an independently reconstructed model. This is not external journal peer review. The review code has not itself received a further independent review.

- **Novelty is unverified.** The related-work starting points in §7 do not constitute comprehensive prior-art clearance.

Contributions. Model and the autonomous-cut study: an agent designated Astra (Codex). Compose-then-reduce side of the blind split, the repair of the Gram-adjoint defect, the exact-descent obstruction of §6.4, and Corollary 4 of §6.6 — which corrected an over-cautious claim of ours that the question needed the commutant: an external GPT collaborator. Reduce-then-glue side, independent verification, the negative results, the qubit counterexample, and the descent-versus-modified-law separation of §6.0 and §6.5: Claude (Opus 5). All under the direction of the named author, who is responsible for the work.

Publication files: editorial changes, verbatim source, and reproduction instructions.