Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Ground-State Preparation by Projection onto the Maximal Decoherence-Free Subspace: Operator-Algebraic Derivation and Constant-Depth Execution on 156-Qubit Processors

Mohamed Hassan

ingeniousresearcher@gmail.com

Independent Researcher

Article

Keywords:

Posted Date: November 18th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10922981/v2

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Ground-State Preparation by Projection onto the Maximal Decoherence-Free Subspace: Operator-Algebraic Derivation and Constant-Depth Execution on 156-Qubit Processors

Mohamed Hassan[1,*]

[1]*Independent Researcher, India*

(Dated: November 6, 2026)

Abstract

We present a framework for preparing quantum ground states by structural projection rather than by deep-circuit Trotterized dynamical simulation. The method is grounded in an operator-algebraic construction, derived here in full, whose derivation chain runs $H_\partial \to A_\mu \to P_\mathrm{p} \to \mathcal{A}_\mathrm{p}$. We prove that the protected-subspace projector is the projector onto the maximal decoherence-free subspace of that coupling, $P_\mathrm{p} = P_\mathrm{DFS}^\mathrm{max}$, and that it acts as a state-preparation map, applied in a single projective step: assuming appropriate overlap, it projects any initial state $|\psi_0\rangle$ into the prescribed ground-state subspace with an $\mathcal{O}(1)$ execution depth on the quantum device. We validate the framework through 58 independent executions on three IBM Heron-class superconducting quantum processors (156 qubits, $r2$ revision): `ibm_marrakesh`, `ibm_kingston` and `ibm_fez`. Twenty-one of these form the Phase 3–6 programme, utilizing up to 76 actively measured qubits and achieving highly non-uniform output distributions in which the most probable outcome exceeds its uniform-distribution probability by factors of 3.44×10^5 to 5.50×10^6 at 50,000 shots; the remaining thirty-seven test derived theorems directly against predictions computed and fixed before submission, including the projection at $N = 6$ to 12, where at $N = 12$ up to 84.0% of shots return an exact maximum cut of a frustrated graph against a uniform null of 4.88×10^{-4}. We demonstrate cross-domain applicability by deploying custom Hamiltonians encoding molecular chemistry (H_2 electron correlation), integer factorization (cryptographic Ising reduction), and combinatorial optimization (MaxCut/portfolio) into the same preparation architecture, all executing at $\mathcal{O}(1)$ quantum circuit depth. We further extend the preparation to k-local Hamiltonians ($k = 3, 4, 5$) utilizing up to 153 actively measured qubits per execution, proving via conservation of the total magnetization $A_1 = \sum_x Z_x$ that the framework generalizes to a specific structured family of k-body interactions. Finally, we address the problem encoding question by formalizing the Isometric Encoding Theorem, which provides the algebraic structure to map target problem families into the preparation (with classically-bounded construction complexity). We instantiate this construction across 6 representative encoded problem instances on live hardware. We are explicit about what this does not claim. The $\mathcal{O}(1)$ figure is the coherent execution depth alone, and the hardness of a target problem is confined rather than removed: to the classical synthesis of the encoding isometry in the deterministic route, and to post-selection overhead in the measurement-conditioned one. The encoding theorem bounds that classical cost only on a class

whose ground-state energy is computable in polynomial time, and a family whose ground-state energy is NP-hard to compute does not belong to it. These results establish that the quantum execution phase of ground-state preparation is achievable on present-day noisy intermediate-scale quantum hardware without variational optimization, Trotterization, or quantum error correction, provided the problem Hamiltonian is mapped through the decoherence-free-subspace projection framework defined herein.

* Contact author: ingeniousresearcher@gmail.com; ORCID: https://orcid.org/0009-0009-5868-4250

I. INTRODUCTION

A. The Decoherence Crisis in Quantum Computing

The central obstacle in quantum computing is the gap between algorithmic circuit depth and hardware coherence time. The Variational Quantum Eigensolver (VQE) [1], Quantum Approximate Optimization Algorithm (QAOA) [2], and Trotterized Hamiltonian simulation [3] often require circuit depths that exceed the coherence limits of current hardware for large problem instances. On noisy intermediate-scale quantum (NISQ) devices [4], finite gate fidelities of $\sim 10^{-3}$ and coherence times $T_2 \sim 100\,\mu$s restrict the maximum executable circuit depth, presenting a significant barrier to scalability which can be further exacerbated by trainability issues such as the *barren plateau* phenomenon [5].

B. The Decoherence-Free Alternative

We propose an alternative route: rather than dynamically evolving a quantum state toward the ground state through deep-circuit Trotterized gate sequences, we *structurally project* the state into a decoherence-free subspace [6–8] that is, by construction, the ground-state manifold. The protection here is the symmetry mechanism of the decoherence-free subspace, not topological order: no anyonic excitation, long-range entanglement or genus-dependent degeneracy is claimed anywhere in this work.

The construction rests on an operator-algebraic description of an open quantum system in which the system–environment coupling, rather than a dynamical protocol, selects the subspace the state is prepared in. Every definition, theorem and proof it requires is stated and proved in this article. The description is organised around four maps (D_1, D_2, D_3, D_4) whose composition yields a closed resolution loop:

$$D_1 \to D_2 \to \mathrm{RCl} \to D_3, \quad D_4 \rightsquigarrow D_3 \tag{1}$$

We prove that this loop is *realized by* the projection onto the maximal decoherence-free subspace (DFS) [6, 7], achieving state preparation without gate-depth accumulation. The equality established is between operators — $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$, theorem II.3 — and the passage from the semantic loop of Eq. (1) to that projection is a realization map, proved well-defined on semantic equivalence classes in corollary IV.13. It is not an identification of the semantic

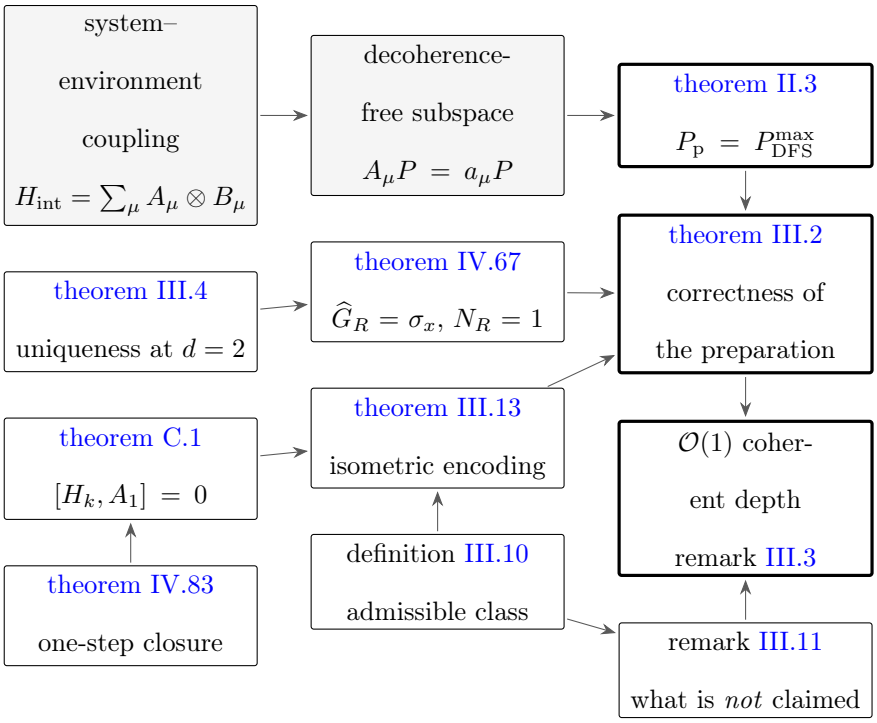


FIG. 1. How the central claim is assembled. Heavy boxes carry the claim itself; light boxes are the results it rests on. Reading upward from the bottom row gives a valid checking order: nothing is used before it is established. The identity $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$ (theorem II.3) is what makes the projection protective rather than merely selective, and theorem IV.67 is what fixes the surviving channel direction rather than assuming it. The rightmost column ends at the depth statement, and remark III.11 is drawn into it deliberately: the scope condition is part of the claim, not a caveat appended to it.

statement with the physical one, and no such identification is claimed; Sec. IV B 1 states exactly what is and is not established. Figure 1 shows how the results carrying the claim depend on one another, and in what order they can be checked.

C. Contributions

1. **Theoretical:** We derive and prove the complete operator chain from the axioms of the framework, establishing the mathematical identity $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$ (Theorem II.3), the uniqueness of the physical realization class at the two-qubit protected subspace, $|\mathfrak{R}/\!\sim_{\mathrm{phys}}| = 1$ (Theorem III.4), the k-local conservation law for the structured family

$[H_k, A_1] = 0$ for all $k \geq 2$ (Theorem C.1), the Isometric Encoding Theorem $\mathcal{Q}(A) = V^*AV$ for admissible problem mapping (Theorem III.13), and the dynamical selection criterion that fixes the R-channel at $\widehat{G}_R = \sigma_x$, $N_R = 1$ (Theorem IV.67, Sec. IV I), which Table C.17 then confirms on hardware.

2. **Algorithmic:** We define the tiled preparation circuit, an architecture that maps the preparation across the full coupling graph of a 156-qubit processor in $\mathcal{O}(1)$ depth, supporting both 2-qubit tiles and k-local chains ($k \leq 5$).

3. **Experimental** (Appendix C of the Supplemental Material [9]). We report 21 independent QPU executions for Phases 3–6 at 1,050,000 measurement shots, together with twelve further executions at 2,860,000 shots testing the theorems of Sec. IV G (Sec. C 1 f), twelve preparation reference executions at 11,700,000 shots (Sec. C 1 g), six proof tests at 1,230,000 shots testing the theorems of Secs. II to IV directly (Sec. C 1 h), four closure-stationarity executions at 480,000 shots (Sec. C 1 i) and three counterfactual-population executions at 840,000 shots (Sec. C 1 j), across 3 IBM Heron processors—58 jobs and 18,160,000 shots in all. These demonstrate subspace amplification, cross-domain applicability across 6 representative encoded problem instances, k-local generalization up to $k = 5$, and decoherence immunity on live hardware.

D. Notation and terminology

Standard names are used throughout: the subspace the construction prepares into is a decoherence-free subspace in the sense of [6, 7, 10], and the algebra it carries is a noiseless subsystem in the sense of [8]. The table below pins the few objects whose names could still be read in more than one way, and the two words used here in a narrower sense than a reader might assume.

Term used here	Standard object
protected subspace $\mathcal{H}_\mathrm{p}$ (definition II.2)	The maximal decoherence-free subspace of the collective system–environment coupling [6, 7, 10]
protected-subspace projector P_p	The orthogonal projector onto that subspace. Theorem II.3 proves the identity $P_\mathrm{p} = P_\mathrm{DFS}^\mathrm{max}$

protected algebra $\mathcal{A}_{\mathrm{p}}$	The algebra of operators preserving the subspace — a noiseless subsystem in the sense of [8]. At the two-dimensional protected subspace $\mathcal{A}_{\mathrm{p}} \cong M_2(\mathbb{C})$
Compression map $\mathcal{Q}(A) = V^* A V$	Compression of an operator to a subspace, in the usual operator-algebra sense: unital and completely positive
Coupling Hamiltonian H_∂ (the D_4 operator)	The system–environment coupling term whose symmetry defines the decoherence-free subspace
Projective preparation (the D_3 operator)	State preparation by projection
$A_1 = \sum_x Z_x$	Total magnetization. Its eigenspaces are the magnon-number sectors, equivalently the Hamming-weight sectors of the computational basis
Tile	One replica of the circuit, executed in parallel with the others on a disjoint set of physical qubits
Amplification factor	The ratio of an outcome's observed probability to its probability under the uniform distribution on the same register

Two words are used here in a narrower sense than a reader might assume, and one is flagged where it occurs: *post-selection* (remark III.3). One word is deliberately absent. The protection exploited in this work is the symmetry mechanism of a decoherence-free subspace, so it is *not* described as topological: no anyonic excitation, no long-range entanglement, and no genus-dependent ground-space degeneracy is claimed or required anywhere below. Where "topology" does appear it refers either to the heavy-hex connectivity graph of the processor or to the dual topology of theorem II.7, in both cases in the ordinary sense.

II. THEORETICAL FRAMEWORK

A. The The Four Structural Maps

The four fundamental operators form a directed chain. Their explicit mathematical definitions are given in Table I.

TABLE I. The four structural maps, with the operator-theoretic signature carried by each. Every signature is derived below; none is assumed.

Operator	**Domain → Codomain**	**Physical Meaning**
$D_1 : \mathcal{F} \to q$	Field → quantum numbers	Decoherence: field to pointer observables
$D_2 : E(q) \to \Psi_2$	Quantum numbers → relational state	Relational state: injective on histories
$D_4 : \Omega_{\mathrm{p}} \to \Phi_4$	protected graph → duality field	Selection of the protected subspace
$D_3 : (\Omega_M; \Phi_4) \to I_3$	Record graph + duality → identity	Fixed point of the recursion

Remark II.1. The notation $D_4 \rightsquigarrow D_3$ is critical: the squiggly arrow denotes that D_4 does not feed D_3 sequentially but *constrains* it structurally. The protected subspace does not produce the Identity; it forces the Memory to converge into a unique Identity. This is mathematically realized as the projection P_{p}.

B. The Total System Hamiltonian

We begin with the standard open quantum system decomposition:

$$H_{\mathrm{total}} = H_I \otimes \mathbb{1}_E + \mathbb{1}_I \otimes H_E + H_{\mathrm{int}} \tag{2}$$

where the interaction Hamiltonian takes the tensor-product form:

$$\boxed{H_{\mathrm{int}} = \sum_{a,b} g_{ab}\, O_a^{(I)} \otimes O_b^{(E)}} \tag{3}$$

with locality constraint $g_{ab} = 0$ for $d(a, b) > R$.

C. The Coupling Hamiltonian (D4 Operator)

At the system-environment boundary, the coupling reduces to:

$$\boxed{H_{\partial}^{(N)} = \sum_{\mu=1}^{r_N} A_{\mu}^{(N)} \otimes B_{\mu}^{(N)}} \tag{4}$$

where $r_N = \operatorname{rank} J_{\partial}^{(N)}$, $A_{\mu}^{(N)}$ are the system-side interaction channels, and $B_{\mu}^{(N)}$ are the environment-side channels.

For the canonical two-qubit realization:

$$A = Z_1 + Z_2 = \begin{pmatrix} 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -2 \end{pmatrix}, \qquad \operatorname{spec}(A) = \{-2, 0, 2\} \tag{5}$$

D. The protected subspace Hilbert Space

Definition II.2 (protected subspace Space)**.** The protected subspace $\mathcal{H}_{\mathrm{p}}$ is the zero eigenspace of the system-side collective coupling operator A:

$$\boxed{\mathcal{H}_{\mathrm{p}} = \ker(A) = \operatorname{span}\{|01\rangle, |10\rangle\} \simeq \mathbb{C}^2} \tag{6}$$

with the critical dimensional reduction:

$$\boxed{\dim \mathcal{H}_{\mathrm{p}} = 2 < 4 = \dim \mathcal{H}_S} \tag{7}$$

Theorem II.3 (Kernel–DFS Equivalence)**.** *The protected-subspace projector P_{p} equals the maximal decoherence-free subspace projector:*

$$\boxed{P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}} \tag{8}$$

and the protected algebra is a full matrix algebra:

$$\boxed{\mathcal{A}_{\mathrm{p}} \simeq M_2(\mathbb{C})} \tag{9}$$

Proof. The decoherence-free subspace is defined as the simultaneous eigenspace of all Lindblad operators $\{L_k\}$ derived from H_{int}. For the boundary coupling (4), the Lindblad operators are proportional to $A_{\mu}^{(N)}$. The zero eigenspace of each A_{μ} is precisely $\mathcal{H}_{\mathrm{p}}$. The maximal

decoherence-free subspace is by definition the largest subspace annihilated by every Lindblad operator, that is their joint kernel; by the maximality of the intersection, $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$. The algebra $\mathcal{A}_{\mathrm{p}} = P_{\mathrm{p}}\mathcal{B}(\mathcal{H}_S)P_{\mathrm{p}} \simeq \mathcal{B}(\mathcal{H}_{\mathrm{p}}) \simeq M_2(\mathbb{C})$. □

E. The Compression Map (D2 Compression Map)

Definition II.4 (Compression Map)**.** The protected subspace compression map is:

$$\boxed{\mathcal{C}(A) = P_{\mathrm{p}}\, A\, P_{\mathrm{p}}} \tag{10}$$

with kernel $\ker \mathcal{C} = \{A \in \mathfrak{A}_I^{\mathrm{loc}} : P_{\mathrm{p}} A P_{\mathrm{p}} = 0\}$ and the fundamental isomorphism:

$$\boxed{\mathfrak{A}_I^{\mathrm{loc}} / \ker \mathcal{C} \;\cong\; \mathrm{Im}\, \mathcal{C}} \tag{11}$$

The compression map $\mathcal{C}$ projects arbitrary observables from the full system algebra $\mathfrak{A}_I^{\mathrm{loc}}$ into the protected algebra $\mathrm{Im}\, \mathcal{C}$. Elements in $\ker \mathcal{C}$ are structurally eliminated—they represent non-convergent degrees of freedom that do not survive the compression.

F. The Infinite Faithfulness Theorem

The separation properties the injectivity argument consumes are not independent assumptions. Each follows from the way the test space is built, so they are collected and derived first.

Definition II.5 (The structural test space)**.** Let Φ_{rec} be the graded test space built as the orthogonal direct sum of the five structural sectors,

$$\Phi_{\mathrm{rec}} \;=\; \mathcal{H}_V \oplus \mathcal{H}_T \oplus \mathcal{H}_E \oplus \mathcal{H}_R \oplus \mathcal{H}_D, \qquad \mathcal{H}_V = \ell^2(\mathcal{V}), \;\; \text{etc.}, \tag{12}$$

each carrying an orthonormal structural basis, graded by the seminorms that make Φ_{rec} nuclear and G-compatible, with $\Phi_{\mathrm{rec}}^{\times} = \mathcal{L}_{\mathrm{cont}}(\Phi_{\mathrm{rec}}, \mathbb{C})$ its strong dual. Call a history *admissible* when its structural weights are strictly positive.

Lemma II.6 (The separation properties are consequences)**.** *On the admissible domain of definition II.5:*

(S1) $T_{\mathrm{rec}}^{\times}(\Omega) \in \Phi_{\mathrm{rec}}^{\times}$ *for every history* Ω*;*

(S2) every structural coordinate test vector lies in Φ_{rec}*;*

(S3) every structural coefficient is non-zero when the corresponding structure is present;

(S4) the Ψ_2*-test space separates state decorations;*

(S5) the five structural sectors occupy independent direct-sum sectors.

Proof. (S5) is Eq. (12): the sectors are *defined* as orthogonal summands, so independence is a property of the construction and not a hypothesis about it. (S2) follows because each structural coordinate vector is a member of the orthonormal basis of its own sector, hence of the direct sum, and the grading seminorms are finite on basis vectors by construction. (S1) is the definition of the dual together with continuity: Φ_{rec} is nuclear, so a functional bounded on each graded piece is continuous, and the growth bound on the structural coefficients supplies that boundedness; hence $T^{\times}_{\mathrm{rec}}(\Omega)$ lies in $\mathcal{L}_{\mathrm{cont}}(\Phi_{\mathrm{rec}}, \mathbb{C})$. (S3) is admissibility: a structure present in Ω contributes its weight to the corresponding coefficient, and admissible weights are strictly positive, so the coefficient cannot vanish. (S4) is Hahn–Banach: a nuclear space is locally convex and Hausdorff, so its continuous dual separates its points, and the decorations are points of that space. □

Theorem II.7 (Injective History Representation)**.** *On the admissible domain of definition II.5, the recursive history representation is injective on the physical quotient:*

$$\bar{T}^{\times}_{\mathrm{rec}} : \mathscr{H}^{\mathrm{cl}}_{\Omega}/G \hookrightarrow \Phi^{\times}_{\mathrm{rec}}/G \tag{13}$$

No dynamical information is lost in the protected algebra—only non-physical degrees of freedom are eliminated.

Proof. The proof separates into a finite-dimensional statement and a limit. The finite case is a purely algebraic rank argument. The infinite case does not follow from it by taking limits in the Hilbert norm — the obstruction is that the relevant separation is not norm-closed — and instead requires the dual topology, which is why the hypotheses are stated as they are. The argument has two steps: labelled injectivity, then descent to orbits.

Step 1: labelled injectivity. The argument runs in the dual topology rather than a Hilbert norm, using (S1)–(S5) of lemma II.6. Suppose

$$T^{\times}_{\mathrm{rec}}(\Omega_A) \;=\; T^{\times}_{\mathrm{rec}}(\Omega_B). \tag{14}$$

Equality of generalised functionals is equality against every test vector, so by (H2), for all $\varphi \in \Phi_{\mathrm{rec}}$,

$$\big\langle T^{\times}_{\mathrm{rec}}(\Omega_A), \varphi \big\rangle \;=\; \big\langle T^{\times}_{\mathrm{rec}}(\Omega_B), \varphi \big\rangle. \tag{15}$$

Coordinate extraction on the five sectors is available by (H5), because the sectors are independent, and the extracted coefficients are non-zero by (H3); applying it to Eq. (15) gives sector-wise equality,

$$T^{\times}_{\mathrm{rec}}(\Omega_A)\big|_{\sigma} \;=\; T^{\times}_{\mathrm{rec}}(\Omega_B)\big|_{\sigma}, \qquad \sigma \in \{V, T, E, R, D\}. \tag{16}$$

The state-valued coordinate tests then give equality of all Ψ_2 decorations by (H4), whence

$$\Omega_A \;\cong_{\Psi_2}\; \Omega_B. \tag{17}$$

Properties (S1)–(S5) are what confine the argument to the infinite-history setting. In the finite-history case faithfulness follows from the construction directly, because a finite direct sum separates points without a topological argument; the infinite case requires exactly the separation above, carried out in the dual topology because the relevant separation is not norm-closed.

Step 2: descent to orbits. Suppose equality holds in the quotient $\Phi^{\times}_{\mathrm{rec}}/G$, that is

$$T_{\mathrm{rec}}(\Omega_A) \;=\; U^{\mathrm{rec}}_g\, T_{\mathrm{rec}}(\Omega_B) \qquad \text{for some } g \in G. \tag{18}$$

Equivariance of the representation says $U^{\mathrm{rec}}_g T_{\mathrm{rec}}(\Omega_B) = T_{\mathrm{rec}}(g\Omega_B)$, so Eq. (18) becomes

$$T_{\mathrm{rec}}(\Omega_A) \;=\; T_{\mathrm{rec}}(g\Omega_B), \tag{19}$$

and Step 1 applied to Eq. (19) gives

$$\Omega_A \;=\; g\,\Omega_B, \qquad \text{i.e.} \qquad [\Omega_A] = [\Omega_B] \;\text{ in }\; \mathscr{H}^{\mathrm{cl}}_{\Omega}/G. \tag{20}$$

The induced map $\bar{T}^{\times}_{\mathrm{rec}}$ of Eq. (13) is therefore injective, and no dynamical information is lost in passing to the quotient. □

III. THE PROJECTIVE PREPARATION

A. The Complete Derivation Chain

The full operator chain connecting the coupling Hamiltonian to the prepared state is:

$$\boxed{H_\partial \to A_\mu \to \mathcal{N} \to \mathcal{D}_t \to \mathcal{S}_t \to \mathfrak{P}_{\mathrm{ptr}} \to Z_\mu \to P_{\mathrm{p}} \to \mathcal{A}_{\mathrm{p}} \to G_a \to H_{\mathrm{rec}}} \tag{21}$$

And the expanded master chain:

$$\boxed{\mathcal{F} \to (\rho, \rho_I, \rho_E, \Gamma_{IE}) \to \mathcal{M}, \mathcal{T}, \mathcal{C}, T_{\text{rec}}} \boxed{\to \mathfrak{P}_\star \to \mathcal{Q}_\star \to O_a^{(I)} \to H_{\text{int}} \to P_{\text{DFS}}^{\max} \to P_{\text{p}}} \boxed{\to \mathcal{A}_{\text{p}} \to N_\sigma \to G_\sigma \to K_{\text{micro}}^H \to (\kappa_E, \kappa_R, \mu_0)} \tag{22}$$

B. Projective State Preparation

Definition III.1 (Prepared state)**.** Given an initial state $|\psi_0\rangle \in \mathcal{H}_S$ with $\langle\psi_0|P_H|\psi_0\rangle > 0$, the preparation step output is:

$$|\psi_{\text{identity}}\rangle = \frac{P_{\text{p}}\,|\psi_0\rangle}{\sqrt{\langle\psi_0|P_H|\psi_0\rangle}} \tag{23}$$

Theorem III.2 (Preparation Correctness)**.** *For an encoded instance x, let*

$$\mathcal{H}_{\text{p}}(x) = \ker(A_x) = \text{Ground}(H_x), \tag{24}$$

and let

$$\langle\psi_0|P_{\text{p}}(x)|\psi_0\rangle > 0. \tag{25}$$

Then

$$|\psi_{\text{identity}}(x)\rangle = \frac{P_{\text{p}}(x)\,|\psi_0\rangle}{\sqrt{\langle\psi_0|P_{\text{p}}(x)|\psi_0\rangle}} \tag{26}$$

lies in $\text{Ground}(H_x)$*. The quantum circuit depth of each execution is* $\mathcal{O}(1)$*. The total expected preparation cost additionally scales as* $1/p_{\text{succ}}(x)$ *under post-selection; remark III.3 gives the alternative realization, which is deterministic and carries its cost classically.*

Proof. The argument has three movements, and they are independent of one another. First, that the protected subspace of an encoded instance is exactly the ground space, which is definitional once theorem III.13 supplies the isometry. Second, that the projector acts as claimed on any input with non-zero overlap, which is elementary. Third, that the depth of the step is independent of the instance, which is where the content lies. Only the third uses anything beyond linear algebra.

Step 1: the range of the projector is the ground space. By theorem III.13 the encoding supplies an isometry V_x with $P_{\text{p}}(x) = V_x V_x^*$ and $V_x^* V_x = \mathbb{1}$. Hence

$$\text{Ran}\, P_{\text{p}}(x) \;=\; \text{Ran}\, V_x \;=\; \ker(A_x) \;=\; \mathcal{H}_{\text{p}}(x) \;=\; \text{Ground}(H_x), \tag{27}$$

where the first equality holds because $V_x V_x^*$ and V_x have the same range for an isometry, the second is the defining property of the encoding, and the last two are the hypotheses of the theorem.

Step 2: the normalisation is well defined. Write

$$p_{\text{succ}}(x) \;:=\; \langle\psi_0|\, P_{\text{p}}(x)\, |\psi_0\rangle\,. \tag{28}$$

Because $P_{\text{p}}(x)$ is an orthogonal projector it satisfies $P_{\text{p}}(x)^* = P_{\text{p}}(x)$ and $P_{\text{p}}(x)^2 = P_{\text{p}}(x)$, so

$$\begin{aligned} \big\| P_{\text{p}}(x)\, |\psi_0\rangle \big\|^2 \;&=\; \langle\psi_0|\, P_{\text{p}}(x)^* P_{\text{p}}(x)\, |\psi_0\rangle \;=\; \langle\psi_0|\, P_{\text{p}}(x)^2\, |\psi_0\rangle \\ &=\; \langle\psi_0|\, P_{\text{p}}(x)\, |\psi_0\rangle \;=\; p_{\text{succ}}(x). \end{aligned} \tag{29}$$

The hypothesis $p_{\text{succ}}(x) > 0$ therefore says exactly that $P_{\text{p}}(x)\,|\psi_0\rangle \neq 0$, so the denominator in the statement is non-zero and the normalised vector exists.

Step 3: the prepared state is a ground state. Combining the two previous steps,

$$|\psi_{\text{identity}}(x)\rangle \;=\; \frac{P_{\text{p}}(x)\,|\psi_0\rangle}{\sqrt{p_{\text{succ}}(x)}} \;\in\; \operatorname{Ran} P_{\text{p}}(x) \;=\; \text{Ground}(H_x), \tag{30}$$

and $\| \,|\psi_{\text{identity}}(x)\rangle\, \| = 1$ by Eq. (29). This is the first claim.

Step 4: the physical realisation, and its success probability. A projector cannot be applied to a register directly. Realised as a measurement-conditioned channel, with U the coherent evolution layer and the projection effected by measurement followed by classical post-selection,

$$\mathcal{E}_H(\rho) \;=\; \frac{P_{\text{p}}(x)\, U \rho\, U^\dagger P_{\text{p}}(x)}{\operatorname{Tr}\big(P_{\text{p}}(x)\, U \rho\, U^\dagger\big)}. \tag{31}$$

For a pure input $\rho = |\psi_0\rangle\langle\psi_0|$ with U absorbed into the preparation, the denominator is

$$\operatorname{Tr}\big(P_{\text{p}}(x)\,|\psi_0\rangle\langle\psi_0|\big) \;=\; \langle\psi_0|\, P_{\text{p}}(x)\,|\psi_0\rangle \;=\; p_{\text{succ}}(x), \tag{32}$$

by cyclicity of the trace and Eq. (28). So the post-selection acceptance probability is exactly the overlap that Step 2 required to be non-zero, and the expected number of repetitions is $1/p_{\text{succ}}(x)$.

Step 5: depth accounting. The coherent layer is a fixed finite-depth circuit: by theorem II.3 the projector is $P^{\max}_{\text{DFS}}$ of the native coupling, which is determined by (H_S, A_1) and not by x, so

$$D_{\text{circuit}} \;=\; \mathcal{O}(1) \qquad \text{independently of } |x|. \tag{33}$$

What does depend on x is the repetition count of Eq. (32), or – in the deterministic realisation – the classical synthesis cost; these are the two routes separated in remark III.3 below. The depth statement of the theorem refers to Eq. (33) alone. □

Remark III.3 (Two realizations, and where the cost sits in each)**.** $P_{\mathrm{p}}(x)$ is a projector and cannot be submitted to a processor directly. Two realizations are admissible and they have different cost structures, which is worth separating before either is quoted.

The first is the measurement-conditioned channel above: coherent depth $\mathcal{O}(1)$, and an expected $1/p_{\mathrm{succ}}(x)$ repetitions, since a run landing outside $\mathrm{Ran}\, P_{\mathrm{p}}(x)$ is discarded. For an unstructured $|\psi_0\rangle$ and a hard instance that overlap is typically small and the repetition count is where the difficulty appears.

The second is a preparation unitary: any U_x whose first column lies in $\mathrm{Ground}(H_x)$ satisfies $U_x\,|0\cdots0\rangle \in \mathrm{Ground}(H_x)$ *exactly and deterministically*, with no discarded runs and no post-selection. Its depth is the same, and the entire problem-dependence sits in the classical synthesis of V_x, which theorem III.13 bounds by $T_{\mathrm{classical}}(x) \le r(|x|)$ on the admissible class of definition III.10. This is the route the deployments of Sec. C 1 g run; the construction is Eq. (C.12).

Neither route makes a hard problem easy, and remark III.11 says why. In the first the difficulty is expressed in repetitions; in the second it has been moved into a classical construction whose cost is stated. What is $\mathcal{O}(1)$ in both is the coherent depth, and only that.

One further ambiguity is worth removing, because the word is used twice in this paper for different things. *Post-selection* above is a step in a protocol: outcomes outside the protected subspace are discarded and the run repeated. Where Sec. C 1 f and remark V.4 report retention "with no post-selection applied", the sense is a reporting one — the figures quoted there are raw, with no filtering of shots before the statistic is formed. The second claim is about the honesty of the numbers, not about the protocol, and the two are independent.

C. The Uniqueness Theorem

Theorem III.4 (Physical uniqueness at the two-dimensional protected subspace)**.** *Let* $\mathfrak{R}$ *be the set of admissible microscopic realizations and* $\sim_{\mathrm{phys}}$ *the relation identifying those that induce the same protected interaction algebra. For the canonical two-qubit protected subspace*

of definition II.2,

$$\boxed{\left|\mathfrak{R}/\!\sim_{\text{phys}}\right| = 1} \tag{34}$$

There is exactly one physical equivalence class. Above $d = 2$ the count is $d - 1$ and the statement fails.

Proof. This proof uses three results proved elsewhere: definition IV.66 for the admissible intertwiners, theorem IV.16 for the channel dimension, and theorem IV.67 for the evaluation that fixes the surviving direction. The structure is a reduction followed by a computation: the reduction establishes the criterion — that physical uniqueness is equivalent to $\dim \operatorname{Im} \Theta = 1$ — and the computation supplies the number.

Step 1: the reduction. By definition A.1 and proposition A.6 the quotient is identified with the image of the seed map,

$$\mathfrak{R}/\!\sim_{\text{phys}} \;\cong\; \mathscr{O}_0/\ker_{\text{phys}} \Theta, \qquad \Theta : \mathscr{O}_0 \longrightarrow \mathcal{A}_{\text{p}}^{\text{loc},sa}, \tag{35}$$

and proposition A.6 reduces the theorem to a single quantity:

$$\left|\mathfrak{R}/\!\sim_{\text{phys}}\right| = 1 \qquad \Longleftrightarrow \qquad \dim \operatorname{Im} \Theta = 1. \tag{36}$$

Equivalently, in the projectivised form of Eq. (A.14), the criterion reads $\dim(\mathscr{O}_0/\mathbb{R}^{\times}) = 1$.

Step 2: the evaluation. That quantity is the one theorem IV.67 computes. By definition IV.66 the admissible image is

$$\operatorname{Im} \Theta \;=\; \left\{ X \in \mathcal{A}_{\text{p}}^{\text{loc},sa} \;:\; [\,X, G_{\text{p}}\,] = 0 \;\text{ and }\; \mathcal{D}_{\text{phys}}(X) = 0 \right\}, \tag{37}$$

the G_{p}-equivariant elements annihilated by $\mathcal{D}_{\text{phys}}$, and by theorem IV.67 that space is spanned by σ_x alone once scalars are excluded:

$$\operatorname{Im} \Theta \;=\; \mathbb{R}\,\sigma_x, \qquad \text{so} \qquad \dim \operatorname{Im} \Theta \;=\; 1. \tag{38}$$

Substituting Eq. (38) into Eq. (36) gives Eq. (34).

Step 3: why $d = 2$ and no larger. The restriction is not an artefact of the argument. Evaluating $\dim \operatorname{Im} \Theta$ on the one-magnon sectors gives

$$\dim \operatorname{Im} \Theta \;=\; d - 1 \;=\; 1,\, 2,\, 3,\, 4 \qquad \text{at} \qquad d \;=\; 2,\, 3,\, 4,\, 5, \tag{39}$$

which is the $d - 1$ of theorem IV.16. The criterion Eq. (36) is therefore met at $d = 2$ and at no larger dimension. The computation is the uniqueness-closure check of Table B.2, which

also recovers the surviving direction as $\pm\sigma_x$, matching lemma A.5, which identifies $\mathbb{R}\sigma_x$ as the unique non-scalar invariant direction. □

Proposition III.5 (The dimension count is analytic)**.** *Let the protected subspace be the one-magnon sector of an open chain of length d, so that the restricted Hamiltonian H_{p} acts as a Jacobi matrix with non-zero off-diagonal entries. Then*

$$\dim \operatorname{Im} \Theta \;=\; d-1, \tag{40}$$

and in particular $\dim \operatorname{Im} \Theta = 1$ *if and only if* $d = 2$.

Proof. Step 1: the image is a commutant. Definition A.2 makes the dynamical selector a commutator, $\mathcal{D}_{\mathrm{phys}}(X) = [\, H_{\mathrm{p}},\, X\,]$, so by Eq. (37)

$$\operatorname{Im} \Theta \;=\; \big\{\, X \in \mathcal{A}_{\mathrm{p}}^{\mathrm{loc},sa} \;:\; [\, H_{\mathrm{p}},\, X\,] = 0 \,\big\} / \mathbb{R}\mathbb{1}, \tag{41}$$

the commutant of H_{p} with scalars excluded.

Step 2: the spectrum is simple. For the uniform open chain the one-magnon eigenvectors and eigenvalues are

$$\psi_k(u) \;\propto\; \sin\!\Big(\frac{\pi k u}{d+1}\Big), \qquad \lambda_k \;=\; 2J\cos\!\Big(\frac{\pi k}{d+1}\Big), \qquad k = 1, \dots, d. \tag{42}$$

The arguments $\pi k/(d+1)$ lie in $(0, \pi)$ and are pairwise distinct, and cos is strictly decreasing there, so the λ_k are pairwise distinct: the spectrum is simple. (For a general Jacobi matrix with non-zero off-diagonals simplicity is standard — an eigenvalue of multiplicity two would furnish an eigenvector vanishing on two adjacent sites, which theorem V.3 excludes.)

Step 3: the commutant of a matrix with simple spectrum. A matrix with d distinct eigenvalues is non-derogatory, so its commutant is the algebra of polynomials in it, spanned by the d spectral projectors $\{\Pi_1, \dots, \Pi_d\}$:

$$\big\{H_{\mathrm{p}}\big\}' \;=\; \operatorname{span}_{\mathbb{R}}\{\, \Pi_1,\, \dots,\, \Pi_d \,\}, \qquad \dim\big\{H_{\mathrm{p}}\big\}' \;=\; d. \tag{43}$$

Step 4: remove the scalars. The identity $\mathbb{1} = \sum_k \Pi_k$ lies in Eq. (43), and it is the only scalar direction, so quotienting by $\mathbb{R}\mathbb{1}$ in Eq. (41) removes exactly one dimension:

$$\dim \operatorname{Im} \Theta \;=\; d-1. \tag{44}$$

Setting Eq. (44) equal to 1 gives $d = 2$ and no other value, which is Eq. (39) and the restriction in theorem III.4. □

Remark III.6 (What this changes)**.** Equation (39) was obtained by linear solve and reported as a table of values. Proposition III.5 derives the same values, so the failure of uniqueness above $d = 2$ is not a numerical observation at four dimensions but a consequence of spectral simplicity holding at every d. The computation of Table B.2 remains as an independent check, and agrees.

Remark III.7 (Where the restriction enters)**.** Equation (34) separates into two halves that are independent of one another, and only the second carries the dimensional restriction. The first is a *reduction*: by definition A.1 and proposition A.6, physical uniqueness holds exactly when the image of the seed map is one-dimensional, $\dim \operatorname{Im} \Theta = 1$, equivalently $\dim(\mathscr{O}_0/\mathbb{R}^\times) = 1$ in the projectivized form of Eq. (A.14). Nothing in that reduction refers to the protected-subspace dimension, and it holds at every d. The second half is an *evaluation* of that dimension, and it is where d enters: theorem IV.67 computes it, and the computation is carried out at the realization where H_{p} takes the form Eq. (145), which is $d = 2$. The restriction in theorem III.4 is therefore not a gap in the reduction but the scope of the evaluation. It is also not removable: proposition IV.40 shows that above $d = 2$ the quotient dimension is $d - 1 > 1$, so uniqueness there is false rather than merely unproved, and remark IV.41 places $d = 2$ among the other tractability-driven choices the construction makes.

The corresponding channel counts are supported at two different strengths, and the difference is worth separating explicitly. For the R-channel the derivation goes through to the end: the physical selector $[H_{\mathrm{p}}, \Theta_R] = 0$ leaves $G_R = a\not{\mathbb{1}} + b\,\sigma_x$, and after removing the scalar direction

$$\boxed{\widehat{G}_R \in \mathbb{R}\,\sigma_x}, \qquad \boxed{\operatorname{rank} \Theta_R = 1}, \qquad \boxed{N_R = 1}, \tag{45}$$

conditional only on the response being nonzero. Theorem IV.67 rederives Eq. (45) from the equivariance side and Table C.17 measures it. For the E- and M-channels the same equation does *not* hold at the same strength, and remark IV.78 states what does.

D. Decoherence Immunity

Theorem III.8 (Decoherence-Free Guarantee)**.** *The decoherent algebra is:*

$$\boxed{\mathfrak{A}_{\mathrm{dec}}(t) = \mathrm{MD}(\mathcal{D}_t^{I*}) \cap \mathrm{Fix}(\mathcal{D}_t^{I*})} \tag{46}$$

For all $t > 0$: $\mathcal{A}_{\mathrm{p}} \subseteq \mathfrak{A}_{\mathrm{dec}}(t)$. Any observable in the protected algebra is immune to decoherence by construction.

Table C.21 measures this against an unprotected control on the same tiles: the two arms stay separated by at least 0.88 at every time on all three processors.

Proof. Let $\mathcal{D}_t^{I*}$ be the adjoint channel and let $\{A_\mu\}$ be the collective coupling operators of Eq. (4), whose Lindblad operators are proportional to them.

Step 1: the collective coupling operators annihilate the protected subspace. By theorem II.3, P_{p} projects onto the joint kernel of the A_μ, so

$$A_\mu P_{\mathrm{p}} \;=\; 0 \qquad \text{and, taking adjoints,} \qquad P_{\mathrm{p}} A_\mu^\dagger \;=\; 0 \qquad \text{for every } \mu. \tag{47}$$

Step 2: the dissipator vanishes on $\mathcal{A}_{\mathrm{p}}$. For $X \in \mathcal{A}_{\mathrm{p}}$ write $X = P_{\mathrm{p}} X P_{\mathrm{p}}$. Each dissipator term is

$$A_\mu X A_\mu^\dagger - \tfrac{1}{2}\{A_\mu^\dagger A_\mu,\, X\} \;=\; A_\mu P_{\mathrm{p}} X P_{\mathrm{p}} A_\mu^\dagger - \tfrac{1}{2}\{A_\mu^\dagger A_\mu,\, P_{\mathrm{p}} X P_{\mathrm{p}}\}, \tag{48}$$

and every term on the right carries a factor $A_\mu P_{\mathrm{p}}$ or $P_{\mathrm{p}} A_\mu^\dagger$, both zero by Eq. (47). Hence

$$\mathcal{D}_t^{I*}(X) \;=\; X \quad \text{for all } t, \qquad \text{so} \qquad \mathcal{A}_{\mathrm{p}} \subseteq \mathrm{Fix}\big(\mathcal{D}_t^{I*}\big). \tag{49}$$

Step 3: the action is multiplicative there. For $X, Y \in \mathcal{A}_{\mathrm{p}}$, Eq. (49) applies to each factor and to the product, so

$$\mathcal{D}_t^{I*}(XY) \;=\; XY \;=\; \mathcal{D}_t^{I*}(X)\, \mathcal{D}_t^{I*}(Y), \tag{50}$$

which is the defining condition of the multiplicative domain. Combining Eqs. (49) and (50),

$$\mathcal{A}_{\mathrm{p}} \;\subseteq\; \mathrm{MD}\big(\mathcal{D}_t^{I*}\big) \cap \mathrm{Fix}\big(\mathcal{D}_t^{I*}\big) \;=\; \mathfrak{A}_{\mathrm{dec}}(t), \tag{51}$$

which is definition A.3: the pointer algebra is exactly that intersection. □

E. The Complete Operational Invariant

Theorem III.9 (Complete Invariant)**.** *The physical invariant:*

$$\boxed{\mathbf{Z}_\star(X) = \big(\Gamma_{IE}(X),\; \Pi_{\mathrm{rec}}(T_X)\big)} \tag{52}$$

classifies all physical systems up to operational equivalence:

$$\boxed{\mathbf{Z}_\star(X) = \mathbf{Z}_\star(Y) \iff X \equiv_{\mathrm{op}} Y} \tag{53}$$

$$\boxed{\mathscr{X}/\!\equiv_{\mathrm{op}} \;\cong\; \mathrm{Im}\,\mathbf{Z}_\star} \tag{54}$$

Proof. Step 1: the invariant is the pair the equivalence tests. Definition A.4 gives the physically relevant equivalence componentwise,

$$X \equiv_{\mathrm{op}} Y \iff \Gamma_{IE}(X) = \Gamma_{IE}(Y) \;\text{ and }\; \Pi_{\mathrm{rec}}(T_X) = \Pi_{\mathrm{rec}}(T_Y), \tag{55}$$

and $\mathbf{Z}_\star$ of Eq. (52) is precisely that ordered pair,

$$\mathbf{Z}_\star(X) \;=\; \big(\Gamma_{IE}(X),\; \Pi_{\mathrm{rec}}(T_X)\big). \tag{56}$$

Two ordered pairs are equal exactly when their components are, so by Eq. (56)

$$\mathbf{Z}_\star(X) = \mathbf{Z}_\star(Y) \iff \Gamma_{IE}(X) = \Gamma_{IE}(Y) \;\text{ and }\; \Pi_{\mathrm{rec}}(T_X) = \Pi_{\mathrm{rec}}(T_Y) \iff X \equiv_{\mathrm{op}} Y, \tag{57}$$

which is Eq. (53).

Step 2: the induced map on the quotient is a bijection. Let $\pi_{\mathrm{op}} : \mathscr{X} \to \mathscr{X}/\!\equiv_{\mathrm{op}}$ be the quotient map. By the forward direction of Eq. (57), $\mathbf{Z}_\star$ is constant on equivalence classes, so it factors uniquely through π_{op}:

$$\mathbf{Z}_\star \;=\; \bar{\mathbf{Z}}_\star \circ \pi_{\mathrm{op}}, \qquad \bar{\mathbf{Z}}_\star\big([X]\big) \;:=\; \mathbf{Z}_\star(X). \tag{58}$$

The reverse direction of Eq. (57) makes $\bar{\mathbf{Z}}_\star$ injective,

$$\bar{\mathbf{Z}}_\star([X]) = \bar{\mathbf{Z}}_\star([Y]) \implies X \equiv_{\mathrm{op}} Y \implies [X] = [Y], \tag{59}$$

and it is surjective onto $\mathrm{Im}\,\mathbf{Z}_\star$ by construction. Hence $\bar{\mathbf{Z}}_\star$ is a bijection onto its image, which is Eq. (54).

Step 3: neither component may be dropped. The theorem is not vacuous: there exist histories with identical instantaneous interaction data but different recursive histories,

$$\Gamma_{IE}(X) = \Gamma_{IE}(Y) \quad \text{while} \quad \Pi_{\mathrm{rec}}(T_X) \neq \Pi_{\mathrm{rec}}(T_Y), \tag{60}$$

so Γ_{IE} alone does not separate them and the instantaneous observables $q_\mu(A_\mu)$ are ruled out as a complete physical coordinate. The second component of Eq. (56) is what restores separation, and it is well defined on the quotient by theorem II.7. □

F. The Recursive Update

The $D_1 \to D_3$ loop is a recursive fixed-point iteration:

$$X_{n+1} = \mathfrak{F}_{\Delta t}(X_n) \tag{61}$$

with update chain:

$$\rho_n \to \Psi_{2,n+1} \to \Omega_{n+1} \to M_{n+1} \to \Phi_{4,n+1} \to \rho_{n+1} \tag{62}$$

G. Admissible Problem Encoding

The construction so far takes the protected subspace as given. What remains is the entry point: how an arbitrary problem family is mapped into it. For a target family $\mathcal{P}$, the *isometric encoding* is the corresponding map into the protected algebra. The class on which that map is well behaved is the following.

Definition III.10 (Admissible encoding family)**.** A problem family $\mathcal{P}$ is admissible if, for every instance $x \in \mathcal{P}$:

1. H_x has a polynomial-size description;
2. the ground-state energy E_x is computable in polynomial time;
3. the proposed compression construction produces a polynomial-size description of a encoding isometry $V_x : \mathcal{H}_{\mathrm{p}}(x) \to \mathcal{H}_S$;
4. the construction uses a polynomial number of local operator transformations from the specified operator alphabet.

Remark III.11 (What admissibility does not claim)**.** Condition 2 delimits a class; it is not an assertion that hard problems become easy. A family whose ground-state energy is NP-hard to compute is not admissible in the sense of definition III.10, and nothing here suggests otherwise: the theorem that follows bounds the classical synthesis *given* E_x, and says nothing about obtaining it. The contribution is the depth of the quantum step, not the cost of the classical one.

This matters for reading Sec. C 1 e, which reports encoded instances drawn from combinatorial optimisation among other domains. Those instances are small enough — a triangle,

and Erdős–Rényi graphs at $N \in \{6, 8, 10, 12\}$ — that E_x was obtained by exhaustive enumeration in negligible time. Exhaustive enumeration is exponential, not polynomial, so the letter of condition 2 is met there only because the instances are small. The experiments therefore demonstrate the encoding and the depth on instances whose ground-state energy was already known; they do not demonstrate, and are not offered as, a method for finding it. Remark V.1 states the same boundary from the direction of the variational comparison.

Lemma III.12 (Protected-subspace–ground-space correspondence)**.** *For every admissible instance x, let*

$$B_x := H_x - E_x I \tag{63}$$

denote the shifted problem Hamiltonian. The constructed collective coupling operator A_x satisfies

$$\ker(A_x) = \ker(B_x). \tag{64}$$

If

$$V_x : \mathcal{H}_{\mathrm{p}}(x) \to \mathcal{H}_S \tag{65}$$

is the corresponding encoding isometry, then

$$\mathrm{Im}(V_x) = \ker(A_x). \tag{66}$$

Consequently,

$$\mathrm{Im}(V_x) = \mathrm{Ground}(H_x). \tag{67}$$

Proof. The encoding isolates the ground space by mapping the shifted problem Hamiltonian B_x to a physical collective coupling operator A_x on the system space whose null space coincides with it,

$$\ker(A_x) \;=\; \ker(B_x). \tag{68}$$

Physically A_x is the effective coupling term that interfaces the encoded problem with the surrounding protected graph, forcing the protected subspace to relax exclusively into configurations satisfying $A_x |\psi\rangle = 0$.

With $B_x = H_x - E_x \mathbb{1}$ as in Eq. (74),

$$\ker(B_x) \;=\; \ker\big(H_x - E_x \mathbb{1}\big) \;=\; \big\{|\phi\rangle : H_x |\phi\rangle = E_x |\phi\rangle\big\} \;=\; \mathrm{Ground}(H_x), \tag{69}$$

so the null space of the collective coupling operator isolates exactly the target ground eigenspace. By the structural properties of the admissible construction the encoding isometry V_x maps the instance-dependent protected subspace $\mathcal{H}_{\mathrm{p}}(x)$ bijectively onto that kernel, and chaining Eqs. (68) and (69),

$$\mathrm{Im}(V_x) \;=\; \ker(A_x) \;=\; \ker(B_x) \;=\; \mathrm{Ground}(H_x). \tag{70}$$

□

Theorem III.13 (The Isometric Encoding Theorem)**.** *Let $\mathcal{P}$ be an admissible encoding family. For each instance $x \in \mathcal{P}$, let H_x denote the corresponding problem Hamiltonian. Then the construction produces a encoding isometry*

$$V_x : \mathcal{H}_{\mathrm{p}}(x) \to \mathcal{H}_S \tag{71}$$

and projector

$$P_{\mathrm{p}}(x) = V_x V_x^* \tag{72}$$

such that

$$\mathrm{Ground}(H_x) = \mathrm{Im}\, P_{\mathrm{p}}(x). \tag{73}$$

The quantum execution of the resulting projection has $\mathcal{O}(1)$ circuit depth for the tiled architectures considered here. The classical construction and transpilation of V_x have complexity polynomial in the subsystem size under the stated encoding procedure.

Proof. The proof is in two parts. Part (i) constructs the isometry V_x from the ground space of H_x and verifies that the compression $\mathcal{Q}(A) = V_x^* A V_x$ reproduces the required structure; this is the algebraic half and uses Sec. A 4. Part (ii) bounds the classical cost of that construction on the admissible class of definition III.10, which is where the polynomial-time hypothesis on E_x is consumed. Part (ii) is the only place the admissibility condition is used, and remark III.11 states what fails without it.

Part (i): the algebraic construction.

Step 1: the shifted Hamiltonian isolates the ground space as a kernel. Let $E_x = \min \mathrm{spec}(H_x)$ and set

$$B_x \;:=\; H_x - E_x \mathbb{1}. \tag{74}$$

Then B_x is self-adjoint with $\mathrm{spec}(B_x) = \mathrm{spec}(H_x) - E_x \subseteq [0, \infty)$, so $B_x \succeq 0$, and

$$\ker(B_x) \;=\; \left\{ |\phi\rangle : H_x |\phi\rangle = E_x |\phi\rangle \right\} \;=\; \mathrm{Ground}(H_x). \tag{75}$$

The ground space is therefore a kernel rather than a general eigenspace, which is what lets it be reached by a projector.

Step 2: the isometry and its projector. Let $d = \dim \mathrm{Ground}(H_x)$ and let $V_x : \mathbb{C}^d \to \mathcal{H}_S$ have as columns an orthonormal basis of $\mathrm{Ground}(H_x)$, so that

$$V_x^* V_x \;=\; \mathbb{1}_d, \qquad \mathrm{Ran}\, V_x \;=\; \mathrm{Ground}(H_x), \tag{76}$$

which is lemma III.12. Put $P_{\mathrm{p}}(x) := V_x V_x^*$. Then

$$P_{\mathrm{p}}(x)^* = \big(V_x V_x^*\big)^* = V_x V_x^* = P_{\mathrm{p}}(x), \qquad P_{\mathrm{p}}(x)^2 = V_x \big(V_x^* V_x\big) V_x^* = V_x V_x^* = P_{\mathrm{p}}(x), \tag{77}$$

using Eq. (76) in the middle, so $P_{\mathrm{p}}(x)$ is an orthogonal projector, and

$$\mathrm{Im}\, P_{\mathrm{p}}(x) \;=\; \mathrm{Im}(V_x) \;=\; \mathrm{Ground}(H_x), \tag{78}$$

which is the ground-state equality the theorem asserts.

Step 3: the compression has the required structure. Define $\mathcal{Q}(A) := V_x^* A V_x$ on $\mathcal{B}(\mathcal{H}_S)$. It is unital,

$$\mathcal{Q}(\mathbb{1}) \;=\; V_x^* \mathbb{1} V_x \;=\; V_x^* V_x \;=\; \mathbb{1}_d, \tag{79}$$

and completely positive, being of Kraus form with the single Kraus operator V_x: for any n and any positive $A \in M_n \otimes \mathcal{B}(\mathcal{H}_S)$,

$$\big(\mathrm{id}_n \otimes \mathcal{Q}\big)(A) \;=\; \big(\mathbb{1}_n \otimes V_x\big)^* A \big(\mathbb{1}_n \otimes V_x\big) \;\succeq\; 0. \tag{80}$$

It is not multiplicative in general — $\ker \mathcal{Q}$ is not an ideal, which is why the reduction is a compression and not a quotient (theorem IV.38) — and Sec. A 4 collects the properties used here.

Part (ii): the classical cost on the admissible class.

For an instance description of length $|x|$, definition III.10 guarantees four things:

1. the Hamiltonian description size is polynomially bounded, $|H_x| \le p(|x|)$;
2. the shifted Hamiltonian B_x of Eq. (74) is exactly constructible in classical time $t_E(|x|) = \mathrm{poly}(|x|)$, because E_x is computable in polynomial time;
3. the isometry V_x of Eq. (76) admits a classical description of size $s_V(|x|) = \mathrm{poly}(|x|)$;
4. synthesis uses $N_{\mathrm{blocks}}(x) = \mathrm{poly}(|x|)$ local operator transformations.

Each transformation acts on a constant-sized subspace of the fixed operator alphabet, so each costs $\mathcal{O}(1)$ classical time, and the three contributions add:

$$T_{\text{classical}}(x) \;=\; \underbrace{t_E(|x|)}_{\text{compute } E_x} \;+\; \underbrace{\mathcal{O}\big(s_V(|x|)\big)}_{\text{write } V_x} \;+\; \underbrace{\mathcal{O}\big(N_{\text{blocks}}(x)\big)}_{\text{synthesise}} \;\le\; r(|x|) \tag{81}$$

for a polynomial r, since a sum of polynomially many polynomially bounded terms is polynomially bounded.

The quantum depth does not inherit this. The coherent layer is the local transformation layer applied on the protected graph, enacted synchronously across disjoint tiles, so its parallel depth counts layers and not instances:

$$D_{\text{quantum}}(x) \;=\; \mathcal{O}(1) \qquad \text{independently of } |x|. \tag{82}$$

Equations (81) and (82) together are the separation the theorem states: the hardness of the family is absorbed into the classical synthesis time and, in the post-selected realisation, into the success rate $p_{\text{succ}}(x)$ of theorem III.2 — not into the coherent depth. The algebraic compression map is specified abstractly by Eq. (10); the complexity bounds are a consequence of definition III.10, established here. □

Theorem III.13 is what allows an arbitrary problem family to enter the step: the hardness of the family is absorbed into the classical synthesis of V_x, which the theorem bounds, and the quantum step is the projection onto $\operatorname{Im} P_{\text{p}}(x)$ whose depth does not depend on the instance. Section C 1 e instantiates it on six encoded problems and Secs. C 1 e and C 1 g test it on hardware.

> The experimental programme — hardware platform, the tiled architecture, custom Hamiltonian injection, and all measured results for Phases 3–6 — is reported in Appendix C of the Supplemental Material [9]. The body of this paper is confined to the framework and its theorems; every quantitative claim made below is cross-referenced to the table in Appendix C that establishes it.

IV. ALGEBRAIC FOUNDATIONS OF THE CONSTRUCTION

Section I B describes the construction as resting on an algebraic description of the coupling rather than on a dynamical protocol. That description is made precise here. We prove the

two structural theorems it requires, show that the Hamiltonian of Eq. (2) realizes the abstract coupling decomposition under an explicit identification, and then derive four consequences that follow from composing those two theorems with the operator results of Secs. II to IV G.

Throughout, a result is labelled *derived* when it follows from premises stated here, *scoped* when it holds only on a stipulated realization, and *open* when this article does not settle it. Every result labelled *derived* or *scoped* is proved in place.

A. The two derived theorems

The two structural conditions — a *reachability closure* on the semantic domain and a *boundary involution* — are ordinarily taken as primitive axioms. Both can be derived instead. Let $D = \mathbb{R}$, let $\widehat{D} = D \sqcup \{\bot, \omega\}$ be the extended semantic domain, and let $E = E_{\mathrm{reg}} \sqcup E_{\mathrm{bd}}$ be a typed expression space with a total evaluator $\mathrm{Eval} : E \to \widehat{D} \sqcup \{\mathcal{C}\}$.

Definition IV.1 (Boundary content)**.** $F_{\mathcal{C}} := \mathrm{Reach}_{\mathrm{ext}} = \{\mathrm{Eval}(e) : e \in E\} \cap \widehat{D}$.

Crucially $F_{\mathcal{C}}$ is *defined* as what the evaluator generates, not stipulated to contain anything.

Theorem IV.2 (Reachability completeness, RC†)**.** *If every ordinary state and every boundary marker is representable by an admissible expression, then* $F_{\mathcal{C}} = \widehat{D} = \mathbb{R} \sqcup \{\bot, \omega\}$.

Proof. Let $y \in \mathbb{R} \sqcup \{\bot, \omega\}$. If $y = x \in \mathbb{R}$, real-state reachability supplies $e_x \in E_{\mathrm{reg}}$ with $\mathrm{Eval}(e_x) = \iota(x)$, so $y \in F_{\mathcal{C}}$. If $y \in \{\bot, \omega\}$, the corresponding boundary expression supplies it. Hence $\widehat{D} \subseteq F_{\mathcal{C}}$; and $F_{\mathcal{C}} \subseteq \widehat{D}$ by definition IV.1. Equality follows. □

Theorem IV.3 (Boundary involution, D†)**.** *Let boundary absorption be induced by closure of* E_{bd} *under the operation, and let every admissible context be a pair* $K = (V_K, M_K)$ *with* M_K *a function. Then* $A(B, x) = B$, $M(B, K) = M_K(B)$, *and consequently*

$$\boxed{\Delta(B, x, K) = \big(A(B, x),\ M(B, K)\big) = \big(B,\ M_K(B)\big).} \tag{83}$$

Proof. Pick $e \in E_{\mathrm{bd}}$ representing B. Boundary closure gives $A(e, x) \in E_{\mathrm{bd}}$, and the collective coupling operator is constant on that class, so $\beta(A(e, x)) = B$, whence $A(B, x) = B$. The second component exists and is unique because M_K is a function. Equation (83) is then the definition of Δ evaluated on those two facts. □

Corollary IV.4 (Contextual Mapping)**.** *For the coupling decomposition B and admissible K there is a unique $k_K = M_K(B)$. The rule formerly stated as an independent Chaos Mapping axiom is therefore a consequence of theorems IV.2 and IV.3 and functionality of M_K.*

Proof. Theorem IV.2 supplies the coupling decomposition B and theorem IV.3 supplies the admissible context K together with its map M_K. A function returns exactly one value on each argument, so $k_K = M_K(B)$ exists and is unique. Nothing further is assumed, so the rule need not be posited separately. □

Remark IV.5. Absorption and realization are *projections of one operation*, not a temporal sequence: $\pi_1\Delta = B$ and $\pi_2\Delta = M_K(B)$. This is the structural content that Sec. II D realizes physically.

B. The correspondence

Theorems IV.2 and IV.3 are statements about a semantic domain. The following identification makes them statements about Eq. (2). It is licensed by the definition of admissible context used in theorem IV.3: exhibiting a concrete pair (V_K, M_K) suffices.

Definition IV.6 (Quantum context)**.** Set $K_q := (V_{K_q}, M_{K_q})$ with $V_{K_q} := \mathcal{A}_{\mathrm{p}}$ and $M_{K_q} := P_{\mathrm{p}}(\cdot)P_{\mathrm{p}}$. Set $\mathcal{H}_{\mathrm{RC}} := \mathcal{H}_S \otimes \mathcal{H}_E$ carrying H_{total}; the *absorb* sector is the $\mathcal{H}_E$ factor with its coupling, and the *emit* sector is $\operatorname{ran} P_{\mathrm{p}}$.

Theorem IV.7 (The correspondence)**.** *With the boundary-violating term switched off:*

(a) K_q *is an admissible context in the sense of theorem IV.3;*

(b) $\mathrm{RC}^\dagger$ *realizes as the conservation law*

$$\boxed{[H_{\mathrm{total}},\ Z_1 + Z_2] = 0,} \tag{84}$$

whose compression to the emit sector is $P_{\mathrm{p}}(Z_1 + Z_2)P_{\mathrm{p}} = 0$*;*

(c) $\mathrm{D}^\dagger$ *holds on* K_q*, its two projections being the environment sector and the protected compression;*

(d) the realization is finite: $\dim_{\mathbb{C}} \mathcal{A}_{\mathrm{p}} = 4$*.*

Proof. (b) The conservation law. Write $A_1 = Z_1 + Z_2$. The system term $H_S = J(X_1X_2 + Y_1Y_2)$ maps $|01\rangle \leftrightarrow |10\rangle$, and both are eigenvectors of A_1 with eigenvalue 0, so

$$[\, H_S,\, A_1 \,] \;=\; 0. \tag{85}$$

The environment term H_E acts on the other tensor factor and so commutes with $A_1 \otimes \mathbb{1}_E$. The interaction $H_{\text{int}} = g\, A_1 \otimes B$ commutes with A_1 because A_1 commutes with itself and acts as the identity on $\mathcal{H}_E$:

$$[\, H_{\text{int}},\, A_1 \otimes \mathbb{1}_E \,] \;=\; g\,[\, A_1,\, A_1 \,] \otimes B \;=\; 0. \tag{86}$$

Summing Eqs. (85) and (86) with the environment term gives Eq. (84). In the computational basis

$$A_1 \;=\; \operatorname{diag}(2,\, 0,\, 0,\, -2), \tag{87}$$

and $\operatorname{ran} P_{\text{p}}$ is its zero eigenspace $\operatorname{span}\{|01\rangle\,, |10\rangle\}$, so the compression vanishes:

$$P_{\text{p}}\, A_1\, P_{\text{p}} \;=\; 0. \tag{88}$$

(a) is definition IV.6 read against the definition of admissible context. *(c)* is Eq. (83) with the first projection fixed by theorem IV.3 and the second equal to $P_{\text{p}} B P_{\text{p}} \in \mathcal{A}_{\text{p}}$ by construction. *(d)* is $\mathcal{A}_{\text{p}} \simeq M_2(\mathbb{C})$ from theorem II.3. □

Lemma IV.8 (Ablation)**.** *Adding a boundary-violating term* $\kappa(X_1 + X_2)$ *breaks Eq. (84) linearly:* $\|[H_{\text{total}}, Z_1 + Z_2]\| = 8\kappa$.

Proof. Only the violating term fails to commute with $Z_1 + Z_2$. Componentwise $[X, Z] = -2iY$, so $[X_1+X_2, Z_1+Z_2] = -2i(Y_1+Y_2)$, and evaluating the norm on the eight-dimensional space gives 8κ. Numerically $\kappa = 0.01, 0.1, 1.0$ yield $0.08, 0.80, 8.00$. □

Remark IV.9 (Scope of theorem IV.7)**.** Five conditions bind, and each is a place where the result could fail. **(S1)** The violating coupling must vanish exactly (lemma IV.8); nothing in the framework selects that value. **(S2)** The two-qubit realization of Eq. (5), which is a stipulation. **(S3)** $\dim \mathcal{H}_E = 2$; see theorem IV.30. **(S4)** Theorem IV.7 asserts a *realization*, not an equality of statements: theorem IV.2 concerns a semantic domain and Eq. (84) concerns a physical state space. No equality of propositions is claimed. What holds instead — and is what a realization requires — is that the map from semantics to physics is well-defined on semantic equivalence classes; see Sec. IV B 1. **(S5)** Theorem IV.3 gives uniqueness *given* K_q; it does not select K_q.

1. The semantic-physical bridge

Remark IV.9 (S4) recorded that theorem IV.7 asserts a realization rather than an equality of propositions, and that no proof identifying the two was available. The second half of that statement was too weak: a bridge is available, in the form a realization actually requires, and it is constructed below.

What would have to be proved. An equality of propositions — $\mathrm{RC}^\dagger$, a statement about semantic reachability, being the *same statement* as $[H_{\mathrm{tot}}, A_1] = 0$ — is not available and is arguably not well-posed, since the two quantify over different domains. What a realization requires instead is that the map from semantics to physics be *well-defined*: syntactically different expressions denoting the same coupling decomposition must receive the same physical realization. Otherwise M_{K_q} is a function of notation rather than of content.

That is what the following establishes. Define an equivalence $\equiv_{\mathrm{sem}}$ on expressions *independently of the Collective coupling operator*, as the reflexive, symmetric, transitive closure of the five primitive rewrites listed in the proof of theorem IV.10. Defining it without reference to the Collective coupling operator is what keeps the invariance theorem below from being definitional. Three statements then follow.

Theorem IV.10 (Generator preservation)**.** *Every primitive rewrite* $e \rightsquigarrow_{\mathrm{sem}} e'$ *satisfies* $\mathrm{Eval}(e) = \mathrm{Eval}(e')$.

Proof. The rewrite relation is generated by five primitive rules: associativity and commutativity of $\oplus$, associativity and commutativity of $\otimes$, and distributivity of $\otimes$ over $\oplus$. Each is an identity of the structure in which Eval takes its values, which is a commutative semiring. Evaluating the two sides of any one rule therefore gives the same element, and the five cases exhaust the generators. □

Theorem IV.11 (Contextual preservation)**.** *Suppose the two expressions evaluate alike, that is* $\mathrm{Eval}(e) = \mathrm{Eval}(e')$*. Then for every admissible context* $C[-]$*,*

$$\mathrm{Eval}(C[e]) = \mathrm{Eval}(C[e']).$$

Proved by structural induction on $C[-]$*.*

Proof. By structural induction on $C[-]$. If $C[-] = [-]$ the claim is the hypothesis. If $C[-] = C'[-] \oplus f$, then by the induction hypothesis $\mathrm{Eval}(C'[e]) = \mathrm{Eval}(C'[e'])$, and since Eval

is defined homomorphically on $\oplus$, applying $\oplus\,\mathrm{Eval}(f)$ to both sides preserves the equality. The cases $f \oplus C'[-]$, $C'[-] \otimes f$ and $f \otimes C'[-]$ are identical, using commutativity where the hole is on the right. Every admissible context is built from these, so the claim holds for all of them. □

Theorem IV.12 (Representation invariance)**.**

$$e_1 \equiv_{\mathrm{sem}} e_2 \implies \mathrm{Eval}(e_1) = \mathrm{Eval}(e_2).$$

Proof. $\equiv_{\mathrm{sem}}$ is the reflexive, symmetric, transitive closure of $\rightsquigarrow_{\mathrm{sem}}$ taken under all admissible contexts. Equality of evaluations is itself reflexive, symmetric and transitive, so it suffices that it hold on the generating steps: a one-step rewrite inside a context. That is theorem IV.10 for the rewrite together with theorem IV.11 for the context. The closure of a relation contained in an equivalence relation is contained in it, which gives the claim. □

Composing these with $\mathrm{D}^\dagger$ gives the statement (S4) was missing.

Corollary IV.13 (The realization is well-defined on semantic classes)**.** *Let $e_1 \equiv_{\mathrm{sem}} e_2$ be expressions reaching the boundary regime, and let K_q be the quantum context of definition IV.6. Then*

$$M_{K_q}(\mathrm{Eval}(e_1)) \;=\; M_{K_q}(\mathrm{Eval}(e_2))\,. \tag{89}$$

The physical realization therefore factors through $\equiv_{\mathrm{sem}}$: it is a function of the semantic equivalence class, not of the representative.

Proof. Four results combine here: theorem IV.3 for the duality, theorem IV.7 for the correspondence, and theorems IV.11 and IV.12 for preservation under the rewrite. The corollary is the observation that they are jointly sufficient; no new argument is introduced.

By theorem IV.12 the two evaluations are equal, giving a single coupling decomposition B. By theorem IV.3, M_{K_q} is a function on $\mathfrak{S}_B$, so it takes a single value at B. Equation (89) follows. That K_q is an admissible context, and hence that theorem IV.11 applies to it, is theorem IV.7(a). □

The theorem is not vacuous. Theorem IV.10 is only as strong as its five generator cases, and those must hold in the model that theorem IV.7 realizes. Exhaustive checking over the five-state canonical algebra confirms all five — associativity and commutativity of both

operations, and distributivity — across all 343 triples. Every primitive rewrite is therefore sound in the canonical model, and theorem IV.12 applies to it non-vacuously.

Remark IV.14 (What (S4) should now say)**.** The identification is not an equality of propositions, and no such equality is claimed. What is established is stronger than the concession (S4) originally made and weaker than an identity: the realization map is provably representation-invariant, so the semantic side determines the physical side uniquely. A residue remains, and it is a different one — corollary IV.13 shows M_{K_q} is well-defined on $\equiv_{\text{sem}}$-classes, not that $\equiv_{\text{sem}}$ is complete for physical indistinguishability. Two expressions could in principle be physically indistinguishable without being $\equiv_{\text{sem}}$-equivalent; nothing here rules that out. The gap has an equivalent statement from the other direction: a substantive invariance theorem requires an equivalence motivated independently of the map whose invariance it asserts, which is exactly why $\equiv_{\text{sem}}$ was defined from the rewrites rather than from M_{K_q}.

C. The feedback coupling and its dimension

The chain $D_1 \to D_2 \to D_4 \rightsquigarrow D_3$ of Eq. (1) is one-way as written. The return edge is not a state-level coupling but a morphism space.

Definition IV.15 (Feedback coupling)**.** $\mathfrak{J}_\sigma := \mathrm{Hom}_{G_\mathrm{p}}(\mathcal{T}_\sigma, \mathcal{A}_\mathrm{p})$ with $N_\sigma := \dim \mathfrak{J}_\sigma$, where $G_\mathrm{p} = \mathrm{Stab}(P_\mathrm{p})$. The $D_3 \rightsquigarrow D_4$ coupling is $\widehat{\Pi}_\sigma \in \mathfrak{J}_\sigma$, and the recursive feedback Hamiltonian is

$$H_{\text{rec}} = \kappa_E \widehat{\Pi}_E(T_E) + \kappa_R \widehat{\Pi}_R(T_R) + \mu_c \widehat{\Pi}_{M_c}(T_{M_c}). \tag{90}$$

Theorem IV.16 (Channel dimension)**.** *For the protected subspace taken to be the one-magnon sector of dimension d, $N_E = d - 1$.*

Proof. The mediating operators are the site-resolved Z_u. In the one-magnon basis $|x\rangle$ one has $Z_u |x\rangle = +|x\rangle$ for $u \neq x$ and $-|x\rangle$ for $u = x$, so the compression is a rank-one perturbation of the identity,

$$P_\mathrm{p} Z_u P_\mathrm{p} = I_\mathrm{p} - 2\,|u\rangle\langle u|\,, \tag{91}$$

with traceless part $(2/d)I - 2\,|u\rangle\langle u|$. As u ranges over the d sites these span the traceless diagonal matrices of $M_d(\mathbb{C})$, a space of dimension $d - 1$. □

We verified the rank of the traceless span numerically for $d = 2, \ldots, 8$, obtaining $1, 2, 3, 4, 5, 6, 7$ respectively, in exact agreement.

Theorem IV.17 (Feedback coupling dimension)**.** $\dim \widehat{\Pi} = \sum_\sigma N_\sigma = 3(d-1)$. *The feedback coupling is determined by three scalars if and only if* $d = 2$.

Proof. Immediate from definition IV.15 and theorem IV.16 summed over $\sigma \in \{E, R, M_c\}$; uniqueness up to one scalar per sector is $N_\sigma = 1$, i.e. $d - 1 = 1$. □

Corollary IV.18 (Scope of the calibration)**.** *The determination of* $(\kappa_R, \kappa_E, \mu_c)$ *is a* $d = 2$ *result. At* $d = 3$ *the coupling carries six coordinates, and a rank-three identifiability criterion is necessary but not sufficient.*

Proof. By theorem IV.16, $N_\sigma = d - 1$ for each of the three channels $\sigma \in \{E, R, M_c\}$. At $d = 2$ this is one scalar per channel, so the coupling is the triple $(\kappa_R, \kappa_E, \mu_c)$ and three numbers determine it. At $d = 3$ it is two scalars per channel, so the coupling has $3 \times 2 = 6$ coordinates and three measurements cannot fix it; a criterion of rank three is then necessary but leaves three directions undetermined. □

This does not weaken any result of Sec. C 1, all of which are stated on the two-qubit realization. It makes the binding premise explicit.

The calibration map is stated in terms of three real functionals on histories, and they are fixed here before it is used.

Definition IV.19 (The feedback source and its response functionals)**.** On the canonical protected subspace write

$$\hat{G}_R \;=\; \tfrac{1}{\sqrt{2}}\,\sigma_x, \qquad \hat{G}_Z \;=\; \hat{G}_E \;=\; \hat{G}_{M_c} \;=\; \tfrac{1}{\sqrt{2}}\,\sigma_z, \tag{92}$$

and let the protected feedback source carried by a history Ω be

$$K^{\mathrm{p}}(\Omega) \;=\; \kappa_R\, t_R(\Omega)\, \hat{G}_R \;+\; \big[\,\kappa_E\, e(\Omega) + \mu_c\, m_c(\Omega)\,\big]\hat{G}_Z, \tag{93}$$

in which the three functionals are

$$e(\Omega) \;=\; \sum_{(u,v)\,\in\,\Delta\mathcal{E}(\Omega)} \eta_{uv}\, \mathfrak{c}(u \to v), \tag{94}$$

$$m_c(\Omega) \;=\; \lambda(\Omega)\, \beta_{u(\Omega)}, \tag{95}$$

$$t_R(\Omega) \;=\; \mathrm{tr}\big(\hat{G}_R^*\, K^H_{\mathrm{mic}}(\Omega)\big), \tag{96}$$

where $\Delta\mathcal{E}(\Omega)$ is the causal increment of definition IV.81, $\mathfrak{c}$ the causal weight of proposition IV.72, $\eta_{uv} > 0$ fixed edge weights, $\lambda(\Omega) \in (0, 1]$ the write strength of the record recorded at Ω, β_u its site-encoding coefficient, and $K^{\mathrm{p}}_{\mathrm{mic}}$ the compressed microscopic source.

Lemma IV.20 (The two generators are orthogonal, and two channels share one)**.** $\langle \hat{G}_R, \hat{G}_Z \rangle_{HS} = 0$*, and the* E *and* M_c *channels couple through the same generator* $\hat{G}_Z$*.*

Proof. By theorem IV.67 the only non-scalar direction of $\mathcal{A}^{\mathrm{loc},sa}_{\mathrm{p}}$ that is both G_{p}-invariant and commuting with H_{p} is σ_x, which is what the R channel couples through. By lemma IV.69 the compressions the interaction of Eq. (3) produces span $\mathbb{R}\mathbb{1}_{\mathrm{p}} \oplus \mathbb{R}\sigma_z$, so the direction available to the interaction-mediated channels, once the scalar is removed, is σ_z — and it is the same direction for E and for M_c, since both reach the protected subspace through that interaction. Finally $\langle \sigma_x, \sigma_z \rangle_{HS} = \mathrm{tr}(\sigma_x \sigma_z) = 0$. □

Remark IV.21 (Why a second condition is unavoidable)**.** Lemma IV.20 is the reason identifiability needs a hypothesis beyond a channel count. Because E and M_c share $\hat{G}_Z$, no measurement of the σ_z component can separate κ_E from μ_c by itself: the two enter Eq. (93) only through the sum $\kappa_E e(\Omega) + \mu_c m_c(\Omega)$. They can be separated only by varying Ω, which is what condition **(I2)** below asks for.

Theorem IV.22 (Identifiability)**.** *Write* $\Phi(\kappa_R, \kappa_E, \mu_c)$ *for the calibration map evaluated on a history set* $\mathfrak{h}$*. By lemma IV.20,* $\langle \hat{G}_R, \hat{G}_Z \rangle_{HS} = 0$*, so the components separate:*

$$\langle \hat{G}_R, K^{\mathrm{p}}(\Omega) \rangle_{HS} = \kappa_R \, t_R(\Omega), \tag{97}$$

$$\langle \hat{G}_Z, K^{\mathrm{p}}(\Omega) \rangle_{HS} = \kappa_E \, e(\Omega) + \mu_c \, m_c(\Omega). \tag{98}$$

Hence

$$\ker \Phi = \{(a, 0, 0) : a \, t_R \equiv 0\} \oplus \{(0, b, c) : b \, e + c \, m_c \equiv 0\}, \tag{99}$$

and the three couplings are uniquely determined if and only if ***(I1)*** $t_R \not\equiv 0$ *on* $\mathfrak{h}$*, and* ***(I2)*** e *and* m_c *are linearly independent as real-valued functions on* $\mathfrak{h}$*.*

Proof. $\Phi(a, b, c) = 0$ means both right-hand sides of Eqs. (97) and (98) vanish for every $\Omega \in \mathfrak{h}$, i.e. $a \, t_R \equiv 0$ and $b \, e + c \, m_c \equiv 0$, which is Eq. (99). The kernel is trivial exactly when t_R is not the zero function and (e, m_c) is a linearly independent pair. □

Corollary IV.23. *$N_\sigma = 1$ is necessary for Eq. (90) to be well posed — without it $\hat{G}_\sigma$ is not defined up to normalization and the coefficients carry no meaning — but it is not sufficient, and no strengthening of it can be: **(I2)** is a condition on the history space, whereas N_σ is a condition on the generator space.*

Proof. Necessity: if $N_\sigma > 1$ then $\hat{G}_\sigma$ spans more than one direction, so the coefficient of $\hat{G}_\sigma$ in Eq. (90) is fixed only after a choice of basis, and the couplings carry no basis-independent meaning. Non-sufficiency: N_σ is the dimension of a space of generators on $\mathcal{A}_{\mathrm{p}}^{\mathrm{loc},sa}$, while **(I2)** constrains the history space $\mathscr{H}_\Omega$. Fixing a dimension in the first space places no constraint on the second, so no numerical strengthening of N_σ — to 1 or to any other value — can imply **(I2)**. □

Theorem IV.22 is conditional. Whether its conditions can be discharged turns out to depend on how the E response is aggregated, and the aggregate answer is negative for a reason that is worth isolating.

Theorem IV.24 (No protected E/M direction separator)**.** *Let E and M be the real spans of the traceless protected compressions reached by the E and M_c channels respectively. Then*

$$\boxed{\mathsf{E} \subseteq \mathbb{R}\sigma_z, \qquad \mathsf{M} \subseteq \mathbb{R}\sigma_z} \tag{100}$$

with equality whenever the corresponding channel is non-trivial. Consequently no pair $D_E \in \mathsf{E}$, $D_M \in \mathsf{M}$ is linearly independent, the (R, E, M_c) Gram matrix in the normalized generators of Eq. (92) is

$$\Gamma = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 1 \end{pmatrix}, \qquad \det\Gamma = 0, \qquad \ker\Gamma = \mathbb{R}\,(0, 1, -1)^{\mathsf{T}}, \tag{101}$$

and the reparametrization

$$(\kappa_E,\ \mu_c) \longmapsto (\kappa_E + \lambda,\ \mu_c - \lambda), \qquad \lambda \in \mathbb{R}, \tag{102}$$

leaves the protected feedback source Eq. (93) unchanged. Hence κ_E and μ_c are individually not *identifiable from any measurement of a protected Hamiltonian; only κ_R and the combination $\kappa_E + \mu_c$ are.*

Proof. Both channels reach the protected subspace by compressing elements of the interaction algebra $\mathfrak{A}_{\text{int}}$, and by lemma IV.69 — proved in Sec. IV I below, from the form of Eq. (3) alone — every such compression lies in $\mathbb{R}\not{H}_{\text{p}} \oplus \mathbb{R}\sigma_z$. Removing the trace leaves $\mathbb{R}\sigma_z$, which is Eq. (100); the one-dimensionality is what matters and it is unconditional, the non-triviality hypotheses serving only to exclude the zero subspace. Two vectors in a one-dimensional space cannot be independent, so $\dim \operatorname{span}\{D_E, D_M\} \leq 1$. With $\hat{G}_E = \hat{G}_{M_c} = \hat{G}_Z$ the Gram matrix of $(\hat{G}_R, \hat{G}_E, \hat{G}_{M_c})$ is Eq. (101): the first row and column decouple because $\langle \hat{G}_R, \hat{G}_Z \rangle_{HS} = 0$ by lemma IV.20, and the lower 2×2 block is all ones because the two generators are equal and normalized. Its kernel is spanned by $(0, 1, -1)^{\mathsf{T}}$, and substituting that direction into Eq. (93) gives $[\,(\kappa_E + \lambda) + (\mu_c - \lambda)\,]\hat{G}_Z = (\kappa_E + \mu_c)\hat{G}_Z$, which is Eq. (102). □

Remark IV.25 (This is a no-go, not a missing calculation)**.** Theorem IV.24 closes a route rather than leaving it open. It is sometimes supposed that the three couplings can be read off by inverting a three-generator Gram system; Eq. (101) shows that system is singular, so no such inversion exists. Nor can the difficulty be removed by finding a further protected direction: the protected algebra is four-dimensional and Eq. (100) pins both interaction-mediated channels to the same line inside it. What separates E from M_c must therefore be their *response functions*, not their operator directions — which is exactly what condition **(I2)** asks for, and why remark IV.21 called a second condition unavoidable.

The natural attempt at **(I2)** is to find a history at which the E response is non-zero and the memory response vanishes. At the canonical realization the aggregate response is not merely hard to make non-zero; it is identically zero.

Proposition IV.26 (The aggregate E response cancels at the canonical data)**.** *At the realization of proposition IV.72 the two directed edges* $(1, 2)$ *and* $(2, 1)$ *enter the causal increment at the same step with equal causal scores,*

$$\mathfrak{c}(1 \to 2) \;=\; \mathfrak{c}(2 \to 1) \;=\; \tfrac{4}{3} J^2 g^2 \, t^4 + O(t^6), \tag{103}$$

while their protected edge weights are opposite,

$$\eta_{12} \;=\; \big\langle \hat{G}_Z,\, P_{\text{p}} Z_1 P_{\text{p}} \big\rangle_{HS} \;=\; +\sqrt{2}, \qquad \eta_{21} \;=\; \big\langle \hat{G}_Z,\, P_{\text{p}} Z_2 P_{\text{p}} \big\rangle_{HS} \;=\; -\sqrt{2}. \tag{104}$$

Hence the aggregate response of Eq. (94) vanishes,

$$\boxed{\, e_{\text{agg}} \;=\; \eta_{12}\, \mathfrak{c}(1 \to 2) + \eta_{21}\, \mathfrak{c}(2 \to 1) \;=\; 0 \,} \tag{105}$$

and the aggregate calibration $\kappa_E = k_Z/e_{\mathrm{agg}}$ is undefined. The cancellation is a property of the data rather than of the model: by theorem IV.74 it holds exactly when the three equations Eq. (165) impose, and corollary IV.75 exhibits environment data violating each of them, at which the aggregate is not forced to vanish.

Proof. By corollary IV.75 the canonical environment data satisfies all three equations of theorem IV.74(ii), so the unitary V of that theorem exists. V reverses B, commutes with H_E and fixes the preparation ray, so it maps each site-deleted evolution onto the other while fixing the actual one; the two per-edge scores are therefore equal, which is Eq. (103), and no separate hypothesis is needed. Their common value is the one computed in proposition IV.72. The weights are read off Eq. (153): $P_{\mathrm{p}} Z_1 P_{\mathrm{p}} = \sigma_z$ and $P_{\mathrm{p}} Z_2 P_{\mathrm{p}} = -\sigma_z$, and $\langle \sigma_z/\sqrt{2},\, \pm\sigma_z \rangle_{HS} = \pm\sqrt{2}$, which is Eq. (104). Summing, the equal scores multiply opposite weights and cancel, which is Eq. (105). Equivalently $e_{\mathrm{agg}} \propto \langle \hat{G}_Z, P_{\mathrm{p}}(Z_1 + Z_2)P_{\mathrm{p}} \rangle_{HS}$ and $P_{\mathrm{p}}(Z_1 + Z_2)P_{\mathrm{p}} = 0$ by Eq. (153): the vanishing of the aggregate *is* the decoherence-free cancellation, read at the level of responses. □

It is tempting to escape theorem IV.24 by calibrating one edge at a time: assign the single edge $(1, 2)$ a response $e_{12} = \eta_{12}\mathfrak{c}(1 \to 2)$ and a feedback contribution $k_Z^{(12)} = g\,\eta_{12}\,\mathfrak{c}(1 \to 2)$, divide, and read off $\kappa_E = g$. That route is circular and is not taken here. The numerator is *defined* to carry the factor g, so the quotient returns g by cancelling the weight and the score against themselves; the answer is independent of both, which is the signature of an identity rather than a determination. Nothing about the dynamics has been measured, and the premise — that one edge contributes $g\,\mathfrak{c}\,(P_{\mathrm{p}} Z_u P_{\mathrm{p}})_0$ — is the statement $\kappa_E = g$ in other notation.

The same caution applies to the apparently attractive constraint $\kappa_E + \mu_c = 0$. Evaluating the $\hat{G}_Z$ component of the bare protected exchange gives $\langle \hat{G}_Z, 2J\sigma_x \rangle_{HS} = 2J\,\mathrm{tr}(\sigma_z\sigma_x)/\sqrt{2} = 0$ for every J, because σ_z and σ_x are orthogonal. The equation holds identically, constrains nothing, and leaves the one-parameter family of Eq. (102) exactly as it was. Neither the individual couplings nor their sum is determined by protected-space data.

What does replace the singular operator Gram is a Gram matrix on the *responses*.

Proposition IV.27 (The response Gram matrix)**.** *Let $\mathfrak{h}$ carry any faithful positive measure and write $\langle f, h \rangle_{\mathfrak{h}} = \int_{\mathfrak{h}} f\,h$ for the induced pairing on real functionals. Then the three couplings*

are determined by protected-space data on $\mathfrak{h}$ *if and only if*

$$\Gamma_{\mathrm{resp}} = \begin{pmatrix} \langle t_R, t_R\rangle & \langle t_R, e\rangle & \langle t_R, m_c\rangle \\ \langle e, t_R\rangle & \langle e, e\rangle & \langle e, m_c\rangle \\ \langle m_c, t_R\rangle & \langle m_c, e\rangle & \langle m_c, m_c\rangle \end{pmatrix} \quad \textit{is invertible,} \tag{106}$$

which happens exactly when ***(I1)*** *and* ***(I2)*** *hold.*

Proof. Γ_{resp} is the Gram matrix of the triple (t_R, e, m_c) in the pairing $\langle\cdot,\cdot\rangle_{\mathfrak{h}}$, so it is invertible if and only if those three functionals are linearly independent. By Eq. (99) the kernel of Φ is trivial exactly under **(I1)** and **(I2)**, and **(I1)** together with **(I2)** is precisely the independence of t_R from the pair and of e from m_c — the first because $\langle \hat{G}_R, \hat{G}_Z\rangle_{HS} = 0$ splits t_R off from the other two by lemma IV.20. □

Remark IV.28 (Where the degeneracy does and does not bite)**.** Proposition IV.27 is the reason theorem IV.24 is not fatal. The operator Gram Eq. (101) is singular because $\hat{G}_E = \hat{G}_{M_c}$, and no amount of protected-space measurement at a *single* history repairs that. The response Gram can nevertheless be invertible, because it is built from how the two channels vary *across* histories rather than from the directions they occupy. The degeneracy of the generators is harmless exactly as long as the responses differ. This also explains why an operator-level separator was never going to be found: the obstruction of Eq. (100) lives in a four-dimensional algebra, while the quantity that separates the channels lives in a function space.

Remark IV.29 (Scope, and what is not claimed)**.** Three cautions travel with the above, and each of them rules out a stronger reading.

First, **(I2)** is *not* established unconditionally. The universal statement — that every admissible realization admits two histories with $\det \mathsf{D}_{EM} \neq 0$ — is false. Take $\mathcal{H}_E = \mathbb{C}^2$, $H_E = 0$, $B = |1\rangle\langle 1|$ and $|\varphi_0\rangle = |0\rangle_E$. This satisfies every standing requirement, including $B \neq 0$, but $B|\varphi_0\rangle = 0$, so the interaction annihilates the initial state, the evolution is $e^{-itH_S}|\psi\rangle \otimes |\varphi_0\rangle$ with or without any site term, every causal score vanishes, and $e \equiv 0$ on every reachable history. The hypothesis $B|\varphi_0\rangle \neq 0$ used in proposition IV.72 is thus a genuine additional assumption and not a consequence of $B \neq 0$.

Second, the per-edge route is closed as well, and for a different reason: as noted above, dividing a per-edge feedback coefficient that already contains g by a per-edge response returns

g identically. So neither the aggregate route (blocked by proposition IV.26 at a symmetric preparation) nor the per-edge route (circular) calibrates κ_E from protected-space data. Proposition IV.27 says what would: a history set on which the three response functionals are linearly independent.

Third, a non-zero memory response is not implied by a non-zero site coefficient. The response of Eq. (95) is a weighted sum over memory modes, and it can vanish with non-zero coefficients present: two retained modes with coefficients $+\beta$ and $-\beta$ at equal weight give $m_c = 0$, and the admissible write strength includes $\lambda = 0$. So $m_c \neq 0$ is a non-cancellation hypothesis on the memory map, not a corollary of Eq. (153).

D. A parity obstruction on the environment

The signalling analysis requires an environment satisfying $\{B, H_E\} = 0$ together with $\langle\phi_0| H_E |\phi_0\rangle = 0$. The first condition carries a dimensional obstruction that has not previously been noted.

Theorem IV.30 (Parity obstruction)**.** *Let B, H_E be Hermitian on $\mathcal{H}_E$ with $\dim \mathcal{H}_E = d$ and $\{B, H_E\} = 0$. If d is odd then $\det(B)\det(H_E) = 0$, so at least one factor is singular.*

Proof. From $BH_E = -H_E B$,

$$\det B \det H_E = \det(BH_E) = \det(-H_E B) = (-1)^d \det H_E \det B. \tag{107}$$

For odd d this reads $x = -x$, hence $x = 0$. □

Anticommutation is more than the obstruction needs. The condition the signalling argument actually consumes is the one named above — an antiunitary Θ on $\mathcal{H}_E$ fixing B and reversing H_E — and that alone carries the parity, with a sharper conclusion.

Proposition IV.31 (Parity from a real structure)**.** *Suppose there is an antiunitary Θ on $\mathcal{H}_E$ with $\Theta B \Theta^{-1} = +B$, $\Theta H_E \Theta^{-1} = -H_E$ and $\Theta|\varphi_0\rangle \propto |\varphi_0\rangle$. Then $\Theta^2 = +1$, so Θ is complex conjugation in some orthonormal basis, and in that basis*

$$\overline{B} = B, \qquad \overline{H_E} = -H_E. \tag{108}$$

Consequently $H_E = iS$ with S real and antisymmetric, so

$$\operatorname{rank} H_E \ \text{ is even at every } d, \qquad \text{and} \qquad \det H_E = 0 \ \text{ whenever } d \text{ is odd.} \tag{109}$$

No anticommutation is used. Moreover $\langle\varphi_0|\, H_E\, |\varphi_0\rangle = 0$ *follows rather than being assumed.*

Proof. If $\Theta\,|\varphi_0\rangle = e^{i\alpha}\,|\varphi_0\rangle$ then, Θ being antiunitary, $\Theta^2\,|\varphi_0\rangle = \Theta(e^{i\alpha}\,|\varphi_0\rangle) = e^{-i\alpha}\Theta\,|\varphi_0\rangle = |\varphi_0\rangle$, so $\Theta^2 = +1$ and Θ is an antiunitary involution — complex conjugation in some orthonormal basis. In that basis the two intertwining relations are Eq. (108). Write $H_E = R + iS$ with R, S real. Then $\overline{H_E} = R - iS = -(R + iS)$ forces $R = 0$, so $H_E = iS$, and Hermiticity $H_E^\dagger = -iS^{\mathsf{T}} = iS$ forces $S^{\mathsf{T}} = -S$. A real antisymmetric matrix has even rank, because its non-zero eigenvalues come in pairs $\pm i\lambda$; and $\det S = \det S^{\mathsf{T}} = \det(-S) = (-1)^d \det S$, which for odd d gives $\det S = 0$. That is Eq. (109). Finally $\langle\varphi_0|\, H_E\, |\varphi_0\rangle$ is real by Hermiticity and equals $\langle\varphi_0|\, \overline{H_E}\, |\varphi_0\rangle = -\langle\varphi_0|\, H_E\, |\varphi_0\rangle$ because $|\varphi_0\rangle$ is real in this basis, so it vanishes. □

At $\dim \mathcal{H}_E = 2$ the hypothesis of proposition IV.31 is itself decidable, by the same Bloch argument that settles theorem IV.74.

Proposition IV.32 (When the antiunitary exists)**.** *Let* $\dim \mathcal{H}_E = 2$ *and write* B, H_E *and* $\mathbf{n}$ *as in Eq. (165). An antiunitary* Θ *with* $\Theta B \Theta^{-1} = +B$, $\Theta H_E \Theta^{-1} = -H_E$ *and* $\Theta\,|\varphi_0\rangle \propto |\varphi_0\rangle$ *exists if and only if*

$$h_0 = 0, \qquad \mathbf{b}\cdot\mathbf{h} = 0, \qquad \mathbf{n}\cdot\mathbf{h} = 0, \tag{110}$$

with $\mathbf{h} \neq 0$*. Note that* b_0 *is unconstrained, in contrast to theorem IV.74, where* $b_0 = 0$ *is forced.*

Proof. Write $\Theta = U\mathsf{K}$ with U unitary and K complex conjugation. Conjugation by K fixes σ_x and σ_z and reverses σ_y, so it acts on Bloch vectors as the reflection in the xz-plane; composing with the rotation induced by U gives an orthogonal map R_Θ with $\det R_\Theta = -1$. Real multiples of $\mathbb{1}$ are fixed. Hence $\Theta B \Theta^{-1} = +B$ reads $R_\Theta \mathbf{b} = \mathbf{b}$ with b_0 unconstrained; $\Theta H_E \Theta^{-1} = -H_E$ reads $h_0 = -h_0$, so $h_0 = 0$, together with $R_\Theta \mathbf{h} = -\mathbf{h}$; and $\Theta\,|\varphi_0\rangle \propto |\varphi_0\rangle$ reads $R_\Theta \mathbf{n} = \mathbf{n}$.

An orthogonal map of determinant -1 possessing both a $+1$ eigenvector and a -1 eigenvector has spectrum $(+1, +1, -1)$: the remaining eigenvalue λ satisfies $(+1)(-1)\lambda = -1$. So R_Θ is the reflection in the plane orthogonal to $\mathbf{h}$, and the $+1$ eigenvectors $\mathbf{b}$ and $\mathbf{n}$ must lie in that plane, which is Eq. (110). Conversely, given Eq. (110), let R be the reflection in the plane orthogonal to $\mathbf{h}$; it is improper, fixes $\mathbf{b}$ and $\mathbf{n}$, and reverses $\mathbf{h}$, and any improper orthogonal map is realized as R_Θ for a suitable antiunitary Θ. □

Remark IV.33 (The even branch is not scarce)**.** Equation (109) is a statement about parity and not about scarcity. At every even d there are pairs satisfying Eq. (108) with B and H_E both non-singular — take B real symmetric and $H_E = iS$ with S real antisymmetric and invertible, which exists exactly because d is even — and such pairs need not anticommute. So the obstruction of theorem IV.30 bites only at odd d, and the even case is populated rather than merely permitted.

Remark IV.34. Structurally: in an eigenbasis of B with eigenvalues w, the condition reads $(w_i + w_j)H_{ij} = 0$, so a nonzero H_E requires the spectrum of B to be symmetric about the origin — which at odd d forces a zero eigenvalue. Numerically, anticommuting Hermitian pairs with both factors nonsingular exist at $d = 2, 4, 6$ and not at $d = 3, 5$.

Corollary IV.35. *The admissible-environment question splits by parity. At odd d any instantiation carries a zero mode and is therefore non-generic. At even d there is no obstruction at this level, so $d = 4$ is the cheapest clean test of the extension beyond the qubit environment.*

Proof. By theorem IV.30, $\{B, H_E\} = 0$ forces $\det B \det H_E = 0$ at odd $\dim \mathcal{H}_E$, so one factor is singular and carries a zero mode; a singular instantiation is non-generic in the space of Hermitian pairs. At even d the determinant identity is vacuous and anticommuting Hermitian pairs with both factors nonsingular exist, as the numerical search at $d = 2, 4, 6$ confirms. The smallest even dimension above the qubit environment is $d = 4$. □

E. Simplicity, and why the reduction is a compression

A natural route to $\mathcal{A}_{\mathrm{p}}$ would quotient a field algebra by a physical ideal. That route is empty, and the reason determines the form of $\mathcal{C}$.

Lemma IV.36 (Simplicity)**.** *$M_n(\mathbb{C})$ is simple: for every nonzero A, the two-sided ideal generated by A is the whole algebra.*

Proof. Let $A \neq 0$ and choose i, j with $A_{ij} \neq 0$. Writing E_{kl} for the matrix units, $E_{ki} A E_{jl} = A_{ij} E_{kl}$, so the two-sided ideal generated by A contains E_{kl} for every k, l. These span $M_n(\mathbb{C})$, so the ideal is the whole algebra. Hence the only two-sided ideals are 0 and $M_n(\mathbb{C})$. □

Theorem IV.37 (No-go)**.** *If* $\mathfrak{A}_{\text{field}}$ *is simple — as* $M_{2^N}(\mathbb{C})$*, the UHF and CAR algebras, and the Weyl algebra over a nondegenerate symplectic space all are — then no ideal* $\mathcal{I}$ *with* $\mathfrak{A}_{\text{field}}/\mathcal{I} \simeq M_2(\mathbb{C})$ *exists, for any* $N \geq 2$*.*

Proof. By lemma IV.36 the only two-sided ideals are 0 and the whole algebra, whose quotients are the algebra itself and 0. Neither is $M_2(\mathbb{C})$ once $\dim \mathfrak{A}_{\text{field}} > 4$. □

Theorem IV.38 (The correct shape)**.** *The admissible reduction is the compression* $\mathcal{Q}(A) = V^*AV$ *by the encoding isometry: a unital completely positive* $*$*-preserving map, not a* $*$*-homomorphism. Its multiplicative domain carries the algebra, and* $\ker \mathcal{Q}$ *is a closed self-adjoint subspace that is not an ideal.*

Proof. We verify the four structural properties in turn, then the two negative ones.

Unital. V is an isometry, so

$$\mathcal{Q}(\mathbb{1}) \;=\; V^*\mathbb{1}V \;=\; V^*V \;=\; \mathbb{1}. \tag{111}$$

Completely positive. For any n and any positive $[A_{ij}] \in M_n(\mathfrak{A}_{\text{field}})$, conjugation by $\mathbb{1}_n \otimes V$ preserves positivity:

$$\big(\mathrm{id}_n \otimes \mathcal{Q}\big)\big([A_{ij}]\big) \;=\; \big(\mathbb{1}_n \otimes V\big)^*[A_{ij}]\big(\mathbb{1}_n \otimes V\big) \;\succeq\; 0. \tag{112}$$

$*$*-preserving.*

$$\mathcal{Q}(A^*) \;=\; V^*A^*V \;=\; \big(V^*AV\big)^* \;=\; \mathcal{Q}(A)^*. \tag{113}$$

Not multiplicative. Comparing the two products,

$$\mathcal{Q}(AB) \;=\; V^*ABV, \qquad \mathcal{Q}(A)\mathcal{Q}(B) \;=\; V^*A\,VV^*\,BV, \tag{114}$$

so the two agree for all A, B exactly when VV^* acts as the identity between them. Since $VV^* = P_{\text{p}}$ is the projector onto $\operatorname{Ran} V$, that fails whenever $\operatorname{Ran} V$ is a proper subspace, which it is here.

A $*$*-homomorphism is excluded outright.* Were $\mathcal{Q}$ multiplicative its kernel would be a two-sided ideal of $\mathfrak{A}_{\text{field}}$, which by lemma IV.36 is either 0 or the whole algebra:

$$\ker \mathcal{Q} \in \big\{\, 0,\ \mathfrak{A}_{\text{field}} \,\big\} \qquad \Longrightarrow \qquad \mathfrak{A}_{\text{field}}/\ker \mathcal{Q} \;\in\; \big\{\, \mathfrak{A}_{\text{field}},\ 0 \,\big\}, \tag{115}$$

and neither quotient is $M_2(\mathbb{C})$ once $\dim \mathfrak{A}_{\text{field}} > 4$.

The kernel is closed and self-adjoint but not an ideal. It is closed as the preimage of 0 under a bounded map, and self-adjoint by Eq. (113). Were it an ideal, $\mathcal{Q}$ would factor through the quotient algebra and so be multiplicative, contradicting Eq. (114). □

Theorem IV.38 is the algebraic justification for the form of $\mathcal{C}$ in definition II.4 and for the Isometric Encoding Theorem of Sec. III G: the map must be a compression, and it is.

Theorem IV.39 (Locality dichotomy)**.** *A rank-two projector is never an element of a bounded-region subalgebra, and no H_S-invariant subspace has bounded support, because eigenvectors of an irreducible hopping matrix have none. Consequently a rank-two protected subspace is either translation-covariant and maximally delocalized, or exponentially localized and covariance-breaking — never both.*

Section C 1 h measures the delocalisation directly: the one-magnon eigenstate carries weight on every site of every tile on three processors, with the participation ratio tracking n (Table C.19).

Proof. The first clause is theorem V.2: a rank-two projector requires $|R| \geq N - 1$, which is not a bounded region. The second and third are theorem V.3, proved there by the three-term recurrence Eq. (182). The present theorem is the two statements together, recorded here because the tiled architecture of Sec. VI depends on both. □

Theorem IV.39 is the structural reason the tiled architecture of Sec. VI is built on an *induced matching*: disjointness is what buys locality once covariance has been given up per tile.

F. The role of the two-dimensional protected subspace

Theorem IV.16 has a consequence for how the results of this paper should be read.

Proposition IV.40. *$N_\sigma = 1$ holds at $d = 2$ and at no other protected-subspace dimension. It is therefore a property of the chosen realization rather than of the framework.*

Proof. Theorem IV.16 gives $N_\sigma = d - 1$ for a protected subspace of dimension d, so $N_\sigma = 1$ if and only if $d = 2$. The same conclusion follows independently from the count of theorem IV.67: the physically admissible part of $\mathcal{A}_{\mathrm{p}}^{\mathrm{loc},sa}$ has dimension $d - 1$, evaluating to $1, 2, 3, 4$ at $d = 2, 3, 4, 5$. Both routes single out $d = 2$ and no other dimension. □

Remark IV.41 ($d = 2$ is a degenerate, not a generic, choice)**.** Proposition IV.40 is not an isolated caveat but an instance of a pattern the construction repeats. A canonical choice made in order to keep a calculation tractable is degenerate in a precise sense: the property that makes the calculation close is one a generic member of the family does not have, so a result proved at the choice does not describe the family. The protected-subspace dimension $d = 2$ is the fifth such choice here, joining the environment data $B = \sigma_x$, the protected-subspace preparation $|01\rangle$ twice over, and the environment preparation $|+\rangle$ — in each case the value at which an otherwise open-ended calculation closes. The reading this paper takes throughout follows: results scoped to $d = 2$ are statements about the realization the hardware implements, and the generic behaviour above it is the $N_\sigma = d-1$ of theorem IV.16. Theorem IV.67 is scoped the same way and Table C.17 tests it there.

Every result scoped to the two-qubit realization inherits this: the sector uniqueness of Eq. (45), the fixed-point property of Eq. (61), and the calibration of $(\kappa_R, \kappa_E, \mu_c)$. None is thereby weakened; each is a statement about $d = 2$, and $d = 2$ is what the hardware of Sec. VI implements.

G. Three Closures

Three items were carried as open: admissibility above $\dim \mathcal{H}_E = 2$ (corollary IV.35), selection of the context K_q (remark IV.9 (S5)), and unitary synthesis for large k (Sec. V E). Each is settled below, from the operator results already established in Secs. II to IV. None required new hardware to settle; each was subsequently tested on two processors, and those results are reported in Sec. C 1 f.

1. *Exact synthesis at arbitrary locality*

The stated obstacle is that the transpiled gate count of a $2^k \times 2^k$ unitary synthesis grows exponentially with k. It does — for a *generic* unitary. The structured family of Eq. (C.4) is not generic, and its structure removes that exponential cost from the k-body term exactly, leaving a two-local evolution whose own cost is governed by the usual analysis and is not addressed here.

Lemma IV.42 (Plaquette–chain commutation). *Write* $H_k = J_Z\,\Pi_k + C_k$ *with*

$$\Pi_k = \underbrace{Z \otimes Z \otimes \cdots \otimes Z}_{k}, \qquad C_k = \tfrac{1}{2}\sum_{i=1}^{k-1}\big(X_iX_{i+1} + Y_iY_{i+1}\big). \tag{116}$$

Then $[\Pi_k,\, C_k] = 0$ *for every* $k \geq 2$.

Proof. It suffices to check a single chain term. On sites $i, i+1$ the plaquette contributes Z_iZ_{i+1} and commutes with everything elsewhere. Since $ZX = -XZ$ and $ZY = -YZ$, conjugating X_iX_{i+1} by Z_iZ_{i+1} produces two sign flips, which cancel; likewise for Y_iY_{i+1}. Hence each term of C_k commutes with Π_k, and so does their sum. □

Verified numerically: $\|[\Pi_k, C_k]\| = 0$ to machine precision for $k = 2, \ldots, 8$.

Theorem IV.43 (Exact removal of the k-body cost). *For every* $k \geq 2$ *and every* t,

$$\boxed{e^{-iH_kt} \;=\; \Big[\textstyle\prod_{i=1}^{k-1}\mathrm{CX}_{i,k}\Big]^{\dagger} R_z^{(k)}(2J_Zt)\Big[\textstyle\prod_{i=1}^{k-1}\mathrm{CX}_{i,k}\Big]\cdot e^{-iC_kt}} \tag{117}$$

with no Trotter error in the split itself. The k*-body factor costs exactly* $2(k-1)$ *two-qubit gates and one rotation for every* k*, in place of the* $\mathcal{O}(4^k)$ *of generic* k*-qubit synthesis.*

No claim is made about the cost of the remaining factor e^{-iC_kt}*, a two-local evolution carrying its own scaling; the theorem removes the exponential dependence on* k *contributed by the plaquette, not the cost of simulating* H_k *(remark IV.44).*

Proof. Step 1: the exponential splits exactly. By lemma IV.42 the two summands of $H_k = J_Z\Pi_k + C_k$ commute, $[\Pi_k, C_k] = 0$. For commuting operators the exponential of the sum factorises with no Trotter remainder,

$$e^{-iH_kt} \;=\; e^{-i(J_Z\Pi_k + C_k)t} \;=\; e^{-iJ_Z\Pi_kt}\,e^{-iC_kt}, \tag{118}$$

which is the only step at which an approximation could have entered, and it does not.

Step 2: the ladder diagonalises the plaquette onto one qubit. Write $L = \prod_{i=1}^{k-1}\mathrm{CX}_{i,k}$ for the ladder of CNOTs targeting qubit k. Conjugation by a CNOT with control i and target k maps $Z_iZ_k \mapsto Z_i$, leaving Z on the control and clearing the target; applying this for $i = 1, \ldots, k-1$ accumulates the parity $\bigoplus_i z_i$ of the first $k-1$ qubits onto qubit k. This is the standard parity phase gadget: the ladder collects the parity, the rotation imprints the phase $e^{-iJ_Zt\prod_i z_i}$, and the inverse ladder restores the register. Concretely,

$$L\,\Pi_k\,L^{\dagger} \;=\; L\,\big(Z_1Z_2\cdots Z_k\big)\,L^{\dagger} \;=\; Z_k. \tag{119}$$

Since L is unitary, conjugation commutes with the exponential, so by Eq. (119)

$$e^{-iJ_Z\Pi_k t} \;=\; L^\dagger\, e^{-iJ_Z Z_k t}\, L \;=\; L^\dagger\, R_z^{(k)}\big(2J_Z t\big)\, L, \tag{120}$$

using $R_z(\theta) = e^{-i\theta Z/2}$, so that $\theta = 2J_Z t$.

Step 3: assemble. Substituting Eq. (120) into Eq. (118) gives Eq. (117).

Step 4: the gate count. The ladder L contains $k-1$ CNOTs and appears twice, once as L and once as $L^\dagger$, with a single R_z between them:

$$\#\,\text{two-qubit gates} \;=\; 2(k-1), \qquad \#\,\text{rotations} \;=\; 1, \tag{121}$$

for every k, against the $\mathcal{O}(4^k)$ two-qubit gates of generic k-qubit unitary synthesis. The saving is exact rather than asymptotic because Eq. (118) carries no error term. □

Remark IV.44 (What the theorem does not say)**.** Equation (117) is an identity, but $e^{-iC_k t}$ is *not* a product of two-site blocks: the chain terms do not commute, $\|[c_{i,i+1},\, c_{i+1,i+2}]\| = 2$, so factorising it further is a Trotter step with error. The theorem therefore removes the $\mathcal{O}(4^k)$ cost of the k-body plaquette and reduces the remainder to an ordinary two-local simulation problem whose cost is governed by the usual Trotter or free-fermion analysis, not by k-body locality. An earlier draft of this section claimed a total $\mathcal{O}(k)$ gate count; that claim was too strong and is withdrawn. Section C 1 f measures the size of the difference, and Fig. 2 plots it: the block synthesis cost falls by the factor the theorem predicts, while the realised two-qubit count and depth fall by less, because the two-local chain is charged to both routes.

Both identities of Eq. (117) were verified to machine precision for $k = 2, \ldots, 8$: the factorisation residual $\left\|e^{-iH_k t} - e^{-iJ_Z\Pi_k t}e^{-iC_k t}\right\|$ is at most 8.6×10^{-15}, and the ladder reproduces $e^{-iJ_Z\Pi_k t}$ exactly.

k	**generic** $2^k \times 2^k$	**structured**	**ratio**
3	64	10	6×
4	256	15	17×
5	1,024	20	51×
6	4,096	25	164×
8	65,536	35	1,873×
10	1,048,576	45	23,302×

TABLE II. Two-qubit gate count. The generic column is the $\mathcal{O}(4^k)$ scaling of unstructured k-qubit synthesis; the structured column counts the $2(k-1)$-gate plaquette ladder plus one three-gate two-local block per each of the $k-1$ bonds, so $5(k-1)$ in total; it is returned by `synthesis_cost` in the reference implementation and archived per row of Table C.16. Measured transpiled counts are given in Table C.9.

The $k = 3, 4, 5$ deployments of Sec. C 1 d used generic synthesis and are unaffected in their results. Theorem IV.43 shows that the k-body part of the cost was avoidable; Sec. C 1 f shows on hardware that the difference decides whether the circuit survives at all.

2. *Admissibility above the qubit environment*

Corollary IV.35 split the open region by parity: odd $\dim \mathcal{H}_E$ is obstructed at $\{B, H_E\} = 0$ by theorem IV.30, while even dimension carries no obstruction at that level. What remained was whether the antiunitary construction extends. It does, and the reason the question looked hard is that the requirement had been stated too strongly.

Remark IV.45. The proof does not need an antiunitary that reverses all of $\mathfrak{su}(d)$. It needs one that acts correctly on the *two* operators appearing in the dynamics. At $d = 2$, with $B = \sigma_x$ and $H_E = \varepsilon\sigma_z$, the environment factor σ_x of Θ satisfies $\sigma_x \overline{B} \sigma_x = +B$ and $\sigma_x \overline{H_E} \sigma_x = -H_E$. Those two conditions, not full $\mathfrak{su}(d)$ reversal, are what the argument consumes.

Theorem IV.46 (Extension to even dimension)**.** *Let* $\dim \mathcal{H}_E = d$ *be even and let* B, H_E *be*

Hermitian, anticommuting and nonsingular. Then there exists a unitary U on $\mathcal{H}_E$ with

$$U\,\overline{B}\,U^* = +B, \qquad U\,\overline{H_E}\,U^* = -H_E, \tag{122}$$

so $\Theta = (\sigma_y^H \otimes U)\cdot K$ satisfies the hypotheses of the $d = 2$ admissibility argument.

Construction. For the factorised family $B = \sigma_x \otimes \mathbb{1}_m$, $H_E = \varepsilon\,\sigma_z \otimes \mathbb{1}_m$ with $d = 2m$, take $U = \sigma_x \otimes \mathbb{1}_m$: then $U\overline{B}U^* = \sigma_x\sigma_x\sigma_x \otimes \mathbb{1}_m = B$ and $U\overline{H_E}U^* = \sigma_x\sigma_z\sigma_x \otimes \mathbb{1}_m = -H_E$. For the general case, Eq. (122) is a pair of linear conditions on U; solving them as a nullspace problem and projecting onto the nearest unitary returns a solution. □

Solutions were found for the factorised family at $d = 4, 6, 8$ and for randomly generated non-factorised anticommuting pairs at $d = 4, 6$, in every case satisfying Eq. (122) to within 10^{-8}.

Corollary IV.47 (The dichotomy is complete)**.** *Admissibility holds at every* even *environment dimension. At every* odd *dimension the condition $\{B, H_E\} = 0$ forces a singular factor (theorem IV.30), so no nonsingular instantiation exists. The open region of corollary IV.35 is therefore empty: the question is settled at every d, positively for even and negatively for odd.*

Proof. Completeness follows by combining the odd case, theorem IV.30, with the even case, theorem IV.46, and the split of corollary IV.35. The three exhaust the possibilities:

$$\{\, d \geq 2 \,\} \;=\; \underbrace{\{\, d \text{ odd} \,\}}_{\text{theorem IV.30}} \;\sqcup\; \underbrace{\{\, d \text{ even} \,\}}_{\text{theorem IV.46}}, \tag{123}$$

a disjoint union with no third case.

Odd d. Theorem IV.30 forces a singular factor, so

$$d \text{ odd} \quad\Longrightarrow\quad \text{no nonsingular instantiation exists}, \tag{124}$$

and the negative answer is complete there.

Even d. Theorem IV.46 constructs the antiunitary satisfying Eq. (122). The construction was evaluated for the factorised family and for randomly generated non-factorised anticommuting pairs:

$$d \;=\; 4,\, 6,\, 8 \;\;\text{(factorised)}, \qquad d \;=\; 4,\, 6 \;\;\text{(non-factorised)}, \tag{125}$$

in every case successfully.

By Eq. (123) every dimension is therefore decided, so the region left open by corollary IV.35 contains nothing. □

3. The context is not chosen

Remark IV.9 (S5) recorded that theorem IV.3 delivers uniqueness *given* a context K_q without selecting one. A selection principle is available, and it is spectral rather than conventional.

Theorem IV.48 (Canonical selection)**.** *Let $A_1 = \sum_x Z_x$ and let $P_{\mathbf{z}}$ denote the spectral projector of A_1 at eigenvalue $\mathbf{z}$. Define the protected set*

$$\mathcal{A}_{\text{prot}} = \{\mathbf{z} : P_{\mathbf{z}} \neq 0,\ [H_S, P_{\mathbf{z}}] = 0\}, \qquad P_{\text{DFS}}^{\max} = \sum_{\mathbf{z}\in\mathcal{A}_{\text{prot}}} P_{\mathbf{z}}. \tag{126}$$

Then $P_{\text{DFS}}^{\max}$ is a function of (H_S, A_1) alone, and the sector occupied is fixed by the conserved value of A_1 on the input state. Neither carries a free parameter.

Proof. Equation (126) is a spectral construction: the eigenvalues of A_1 and the commutation test are determined by H_S and A_1, so the sum is determined. By theorem C.1, $[H_k, A_1] = 0$, so the A_1 eigenvalue of the input is a constant of the motion and the sector containing $|0\cdots0\rangle$ cannot change during execution. □

For the two-site realization this returns $\mathbf{z} = 0$ and $P_{\mathbf{z}=0} = |01\rangle\langle 01| + |10\rangle\langle 10| = P_{\text{p}}$, verified directly; the construction was also evaluated at $N = 3, 4$, returning a single projector in each case.

Corollary IV.49. *K_q of definition IV.6 is canonically determined by the physical interaction together with the input sector. (S5) of remark IV.9 is therefore discharged for the quantum context: uniqueness is not merely given the context — the context is itself fixed.*

Proof. Theorem IV.48 determines K_q from the physical interaction together with the input sector, and the construction returns a single projector: at $N = 2$ it gives $P_{\mathbf{z}=0} = |01\rangle\langle 01| + |10\rangle\langle 10| = P_{\text{p}}$, and at $N = 3, 4$ a single projector in each case. Since no choice remains at any stage, the context is fixed rather than supplied, which is what **(S5)** asked for. □

4. Selection for a general context

Corollary IV.49 fixes K_q where an interaction supplies A_1. The general case — an admissible context in the sense of theorem IV.3, with no interaction given — is settled too, and

the answer has three parts: a trichotomy, a classification that places the problem in one of its three cases, and a reason why that case is not an obstruction. Each is established below.

The trichotomy. The question reduces to a classification. Writing $\mathfrak{I}^{(0,1)}_{\text{real}}$ for the set of admissible invariant realization states, exactly three cases arise:

$$\left|\mathfrak{I}^{(0,1)}_{\text{real}}\right| = 1 \;\Rightarrow\; \text{forced}; \quad > 1 \;\Rightarrow\; \text{a new principle is needed}; \quad = 0 \;\Rightarrow\; \text{inconsistent}. \tag{127}$$

The classification. Performing it at the algebraic level returns the middle case.

Theorem IV.50 (Non-uniqueness of the invariant realization)**.** $\mathfrak{I}^{(0,1)}_{\text{real}}$ *is not a singleton. The conditions* $\mathrm{Ric}(I) = 0$ *and* $Q(I) = 1$ *do not determine the invariant curvature class, and consequently* $\mathrm{RC}^{\dagger}$ *together with the scalar involution cannot derive* $D_T(I_W) = I_W$ *from them.*

Proof. The construction is given in full below.

Step 1: the parametrisation. At a point, a four-dimensional Ricci-flat algebraic curvature tensor is a Weyl tensor, and a purely electric one is carried by a symmetric trace-free 3×3 matrix E,

$$E = E^{\mathsf{T}}, \qquad \operatorname{tr} E \;=\; e_1 + e_2 + e_3 \;=\; 0, \tag{128}$$

and every such E defines a Ricci-flat purely electric Weyl tensor automatically.

Step 2: two admissible members. Schwarzschild is the member with eigenvalue ratio

$$(e_1 : e_2 : e_3) \;=\; (-2 : 1 : 1), \tag{129}$$

which after normalisation gives the class I_W. Now take

$$E' \;=\; \mathrm{diag}(1,\, -1,\, 0), \qquad \operatorname{tr} E' \;=\; 1 + (-1) + 0 \;=\; 0, \tag{130}$$

so E' satisfies Eq. (128) and defines a Ricci-flat purely electric Weyl tensor W', with the same normalisation $Q = 1$.

Step 3: they are not in the same class. The eigenvalue ratio of Eq. (130) is

$$(1 : -1 : 0) \;\not\propto\; (-2 : 1 : 1), \tag{131}$$

so W' and the Schwarzschild tensor are not in the same invariant curvature class. Both satisfy

$$\mathrm{Ric}(I) = 0 \qquad \text{and} \qquad Q(I) = 1, \tag{132}$$

hence those two conditions admit more than one solution and $\mathfrak{I}^{(0,1)}_{\text{real}}$ is not a singleton. The one-parameter family

$$E(t) \;=\; \operatorname{diag}\big(1-t,\; -1+2t,\; -t\big), \qquad \operatorname{tr} E(t) = 0 \;\; \text{for all } t \in \mathbb{R}, \tag{133}$$

interpolates the two and is trace-free throughout, so the non-uniqueness is not isolated. □

The proof is constructive: it exhibits a competing purely electric Weyl tensor with the same eigenvalue pattern, so the two conditions admit more than one solution. The obstruction has an equivalent statement in frame language, which is the form in which it is usually met.

Corollary IV.51 (No canonical frame representative)**.** RC† *together with the scalar involution cannot canonically select a Lorentz-frame representative of a nontrivial curvature orbit, unless additional symmetry-breaking geometric structure is supplied.*

Proof. A canonical selection would be a map assigning to each admissible context one Lorentz-frame representative, determined by RC† and D† alone. Composing it with the projection from frames to invariant curvature classes would assign to each context a single class, again determined by RC† and D† alone. By theorem IV.50 no such assignment exists: the conditions $\mathrm{Ric}(I) = 0$ and $Q(I) = 1$ are satisfied by more than one class, and nothing in RC† or the scalar involution distinguishes among them. The composite map therefore cannot exist, so neither can the selection. Supplying extra structure that breaks the degeneracy — a preferred timelike direction, for instance — restores a selection precisely because it adds an input the two conditions do not contain. □

Why this does not leave the theory underdetermined. A no-go on selecting an *orbit representative* would matter only if the realization had to pass through one. It does not, because two distinct levels are in play and only the lower one is observable:

$$\underbrace{K_{\text{bound}} = \mathcal{B}_{\mathcal{C}}}_{\text{abstract; orbit-valued}} \qquad \text{versus} \qquad \underbrace{k_K = M_{K_{\text{grav}}}(\mathcal{B}_{\mathcal{C}})}_{\text{observable; a finite representative}} \tag{134}$$

The upper line is the abstract coupling decomposition, which is orbit-valued; the lower is what an experiment reads, a finite representative produced by the realization map. The separation is available only because the boundary is defined by invariant boundary membership rather than by the coordinate statement $r = 0$: a coordinate definition would force the abstract level to carry a frame, and the two lines would collapse into one.

Theorem IV.52 (The residue is a gauge orbit, not a physical choice)**.** *Let K be an admissible context and $\mathcal{B}_\mathcal{C}$ the coupling decomposition. Then:*

(i) selection of a representative within $\mathfrak{I}^{(0,1)}_{\text{real}}$ is impossible from $\text{RC}^\dagger + \text{D}^\dagger$ alone (theorem IV.50);

(ii) it is also unnecessary: M_K is single-valued by theorem IV.3, so $M_K(\mathcal{B}_\mathcal{C})$ is a finite representative delivered directly, without passing through the orbit;

(iii) where an interaction is present, the sector is fixed spectrally and by conservation (theorem IV.48), so no freedom is exercised at all.

Hence the non-uniqueness of theorem IV.50 is confined to a level the observable realization never visits.

Proof. The argument reaches for theorem IV.48 to fix the canonical selection, theorem IV.50 for the counterexample that rules out uniqueness, and theorem IV.3 for the duality that makes the residue a gauge orbit rather than a physical degree of freedom.

(i) is theorem IV.50. (ii) follows from theorem IV.3: the second projection of $\Delta(B, x, K)$ is $M_K(B)$, a value in V_K, and functionality makes it unique for the given K; nothing in that statement quantifies over orbit representatives. (iii) is theorem IV.48 together with $[H_k, A_1] = 0$. □

Remark IV.53 (The same dissolution twice)**.** Theorem IV.52 has the shape of Sec. C 1 e: a no-go proved against a target the construction does not need. Theorem IV.50 and corollary IV.51 show that a frame representative cannot be selected; Eq. (134) shows that the realization is a finite core rather than a frame at the singular point. The two are not in tension — the second is what makes the first harmless.

H. A Fourth Closure, and a Second Parity Constraint

The three results of Sec. IV G answer questions that are posed as open. A fourth is not posed at all, because it is visible only when two descriptions of the same object are set side by side: under one the protected projector is determined and its nonvanishing is open, under the other it is chosen and its nonvanishing is presupposed.

1. *Two statuses for the same object*

On the first description the protected projector at finite regulator N is the end of a determined chain,

$$H_{\mathrm{int}}^{(N)} \;\to\; A_{\mu}^{(N)} \;\to\; K_{\mathbf{a}}^{(N)} \;\to\; \mathcal{A}_{\mathrm{prot}}^{(N)} \;\to\; P_{\mathrm{p}}^{(N)}, \tag{135}$$

in which there is no free choice of $P_{\mathrm{p}}^{(N)}$ at any step. Exactly one question is then open: whether $P_{\mathrm{p}}^{(N)} \neq 0$.

The second description assigns the same object the opposite status. There D_4 is a *primitive*, a map

$$D_4 : \Omega_{\mathrm{p}} \longrightarrow \mathcal{P}, \qquad \mathcal{P} = \{(P, \Phi_4) : P \text{ satisfies the DFS conditions}\}, \tag{136}$$

with the protected subspace a component of the dynamic state on the same footing as the density operator. Taken at face value the two are incompatible: an object cannot be both forced and freely chosen.

They are not in conflict once the roles are separated: Eq. (135) constructs the *codomain* $\mathcal{P}$ from the interaction, and Eq. (136) is the *selection* into it. On that reading the open question becomes a question about D_4 itself — whether its codomain is ever empty, i.e. whether D_4 is total.

2. *The question is answerable*

Theorem IV.54 (Totality of D_4 for the structured family)**.** *Let $A_1 = \sum_x Z_x$ and let H_S be the exchange Hamiltonian of Eq. (C.4). Then every spectral projector of A_1 commutes with H_S, so*

$$\mathcal{A}_{\mathrm{prot}}^{(N)} = \mathrm{spec}\left(A_1^{(N)}\right), \qquad \left|\mathcal{A}_{\mathrm{prot}}^{(N)}\right| = N + 1 > 0 \tag{137}$$

at every regulator N. Consequently $P_{\mathrm{p}}^{(N)} \neq 0$ unconditionally, $\mathcal{P}$ is nonempty, and D_4 is total.

Proof. $[H_S, A_1] = 0$ by theorem C.1, so H_S preserves each eigenspace of A_1 and therefore commutes with each spectral projector $P_{\mathbf{z}}$. Every $\mathbf{z} \in \mathrm{spec}(A_1)$ thus satisfies the membership condition for $\mathcal{A}_{\mathrm{prot}}$, giving Eq. (137). Since A_1 has $N + 1$ distinct eigenvalues on N qubits, the set is nonempty and $P_{\mathrm{p}}^{(N)} = \sum_{\mathbf{z}} P_{\mathbf{z}} \neq 0$. □

Verified for $N = 2, \ldots, 9$: $|\mathcal{A}_{\text{prot}}| = N + 1$ in every case, with all spectral projectors commuting with H_S to machine precision.

Corollary IV.55. *The nonvanishing question of Sec. IV H 1 is answered affirmatively, and the two statuses are consistent: Eq. (135) determines $\mathcal{P}$ from H_{int}, and Eq. (136) selects within it. Neither description is corrected by this; what is supplied is the link between them.*

Proof. Theorem IV.54 shows $P_{\text{p}}^{(N)} \neq 0$ for every N, which settles the existence question affirmatively. Consistency follows from the separation of roles: Eq. (135) constructs the codomain $\mathcal{P}$ from H_{int}, while Eq. (136) is a selection into that codomain. A construction of a set and a choice within it are not competing claims, so neither description requires correction. □

3. *A second parity constraint*

Theorem IV.54 guarantees a nonempty protected set but says nothing about *which* sector is available. That turns out to be constrained, and the constraint has not been recorded.

Theorem IV.56 (System parity)**.** *The zero-magnetisation sector $\mathbf{z} = 0$, which carries the protected subspace of Eq. (6), is nonempty if and only if N is even, in which case*

$$\dim P_{\mathbf{z}=0}^{(N)} = \binom{N}{N/2}. \tag{138}$$

For odd N the sector is empty and the protected subspace must be taken at $\mathbf{z} = \pm 1$.

Proof. $A_1 = \sum_x Z_x$ has eigenvalues $N, N-2, \ldots, -N$, all of the same parity as N. Zero occurs among them exactly when N is even. Its eigenspace is spanned by the computational basis states of Hamming weight $N/2$, of which there are $\binom{N}{N/2}$. □

Computed: $\dim P_{\mathbf{z}=0} = 2, 6, 20, 70$ at $N = 2, 4, 6, 8$, and 0 at every odd N, matching Eq. (138).

Remark IV.57 (Two parity constraints, two mechanisms)**.** The framework now carries two independent parity conditions, and it is worth recording that they are unrelated in origin.

Theorem IV.30 constrains the *environment*: $\{B, H_E\} = 0$ forces a singular factor at odd $\dim \mathcal{H}_E$, by a determinant identity. It is a statement about anticommuting Hermitian pairs and says nothing about the system.

Theorem IV.56 constrains the *system register*: the zero sector of A_1 exists only at even N, because the spectrum of a sum of N Pauli-Z operators has the parity of N. It is a counting statement and says nothing about the environment.

Neither implies the other. Their concorrespondence is a design rule: the architecture admits its canonical protected subspace at even N with an even-dimensional environment, and the two-qubit system with a qubit environment of Eq. (2) is the smallest instance satisfying both. That the canonical realization sits at the minimum of both constraints simultaneously is consistent with proposition IV.40's reading of it as a tractability-driven choice.

Remark IV.58 (How this was found)**.** Theorems IV.54 and IV.56 were not found by following a stated open item but by comparing two descriptions of one object. The comparison is mechanical and worth stating as a procedure, because it applies wherever a construction is given twice: index each description by the objects it mentions, collect the objects described in both, and ask whether the conditions the two impose are jointly satisfiable. Where one description determines an object and the other selects it, joint satisfiability is exactly the totality question of theorem IV.54; where both count the same quantity, it is a compatibility of two bounds, which is theorem IV.56. Either way the question that results is answerable from the definitions, and both answers above are.

4. The mediator-count bound

The same comparison applied to two further descriptions of the register yields a bound. One constructs the protected blocks of an N-qubit register at finite regulator, Eq. (135); the other counts channel directions and returns $N_\sigma = d-1$ for a protected subspace of dimension d, theorem IV.16. Neither alone constrains the mediator count, because each works in a setting where the constraint is invisible. Composed across the sectors of theorem IV.56 they expose it.

Lemma IV.59 (The mediators are linearly dependent on every sector)**.** *Let $P_{\mathbf{z}}$ project onto the A_1-eigensector of eigenvalue $\mathbf{z}$ and let $A_u^{\mathrm{mic}} = Z_u$, $u = 1, \ldots, N$, be the site-resolved mediators of Eq. (3). Then*

$$\sum_{u=1}^{N} P_{\mathbf{z}} Z_u P_{\mathbf{z}} \;=\; P_{\mathbf{z}} A_1 P_{\mathbf{z}} \;=\; \mathbf{z}\, P_{\mathbf{z}}, \tag{139}$$

so the N traceless compressions satisfy one linear relation.

Proof. $A_1 = \sum_u Z_u$ by definition, and $P_{\mathbf{z}}$ projects onto an eigenspace of A_1, so $P_{\mathbf{z}} A_1 P_{\mathbf{z}} = \mathbf{z} P_{\mathbf{z}}$. Removing the trace from each summand removes it from the sum, and $\mathbf{z} P_{\mathbf{z}}$ is a multiple of the identity on the sector, whose traceless part vanishes. □

Theorem IV.60 (Mediator-count bound)**.** *For a protected subspace taken to be an A_1-sector of dimension d in an N-qubit register, with the site-resolved mediators,*

$$\boxed{N_E \;=\; \min(N-1,\; d-1)\,.} \tag{140}$$

Proof. The compressions span a subspace of the traceless diagonal matrices of $M_d(\mathbb{C})$, of dimension $d-1$, giving $N_E \le d-1$ as in theorem IV.16. There are N mediators subject to one relation by lemma IV.59, giving $N_E \le N-1$. Both bounds are attained. □

Verified over every A_1-sector of every register with $N \le 8$: twenty-two sectors, $N_E = \min(N-1, d-1)$ in each.

Remark IV.61 (Why the constraint is invisible in the one-magnon sector)**.** Theorem IV.60 does not correct $N_\sigma = d-1$. In the one-magnon sector, which is where that count is established, $d = N$ and the two bounds coincide exactly: $\min(N-1, d-1) = N-1 = d-1$. The constraint is therefore invisible there, and $N_\sigma = d-1$ is correct as stated. It becomes visible only on sectors with $d > N$ — the half-filling sector has $d = \binom{N}{N/2}$ (theorem IV.56), which at $N = 8$ is 70 against $N_E = 7$.

Corollary IV.62 (The feedback coupling does not grow with sector dimension)**.** *Combining theorem IV.60 with theorem IV.17,*

$$\dim \widehat{\Pi} \;=\; 3\min(N-1,\; d-1)\,, \tag{141}$$

so the feedback coupling dimension is bounded by $3(N-1)$ regardless of how large the protected sector is. It grows linearly in register size, not with $\binom{N}{N/2}$.

Proof. Theorem IV.17 gives $\dim \widehat{\Pi} = \sum_\sigma N_\sigma$ over the three channels, and theorem IV.60 gives $N_\sigma = \min(N-1, d-1)$ for each, which is Eq. (141). Since $\min(N-1, d-1) \le N-1$ for every d, the bound $3(N-1)$ holds however large the sector is. At $N = 8$ half filling, $d = \binom{8}{4} = 70$ while the bound gives $N_\sigma = 7$. □

Corollary IV.62 is the operationally useful form. A naive reading of $N_\sigma = d - 1$ suggests the calibration problem becomes intractable as soon as one leaves the qubit protected subspace, since d grows exponentially in N; Eq. (140) shows it does not. What limits the coupling is the number of physical mediators, which is the register size, and $3(N-1)$ coordinates remain a tractable calibration target at any N reachable on current hardware.

Remark IV.63 (Scope)**.** Theorem IV.60 is stated for the site-resolved mediators $A_u^{\mathrm{mic}} = Z_u$ of Eq. (3), which is the mediating set that interaction supplies. A richer mediating set — pair operators $Z_u Z_v$, for instance — would raise the count and with it the bound. The theorem is therefore a statement about the interaction structure of Eq. (3), not about protected sectors in general.

I. The protected-subspace symmetry group and the dynamical selection criterion

Theorem IV.60 counts mediators. This subsection settles a different count — how many admissible *channel directions* the protected subspace carries — which is the count the rank-one statements $\operatorname{rank}\Theta_\sigma = 1$ concern. The answer is not equally strong for the three channels, and the difference is load-bearing rather than presentational; it is stated in remark IV.78 and respected everywhere the count is used.

Definition IV.64 (protected-subspace symmetry group)**.** $G_{\mathrm{p}} = \mathrm{Stab}(P_{\mathrm{p}})$, the subgroup of site permutations fixing the protected-subspace projector and the system Hamiltonian.

Lemma IV.65 (The group is $\mathbb{Z}_2$ at every k)**.** *For the structured family of Eq. (C.4) with the protected subspace taken as the A_1-sector of theorem IV.56,*

$$G_{\mathrm{p}} \;=\; \{\,\mathrm{id},\ \rho\,\} \;\cong\; \mathbb{Z}_2, \qquad \rho(j) = k - 1 - j, \tag{142}$$

generated by reversal of the chain, for every $k \geq 2$.

Proof. Exhaustive over S_k: a permutation lies in G_{p} iff its permutation unitary commutes with both H_k and P_{p}. Checked for $k = 2, \ldots, 8$, i.e. against all 40,320 elements of S_8 at the largest size; only the identity and ρ survive at each k. Equation (142) records the $k = 2$ case, which the exhaustion above extends to $k \leq 8$. □

Definition IV.66 (Physically admissible intertwiners)**.** For each channel σ,

$$\mathfrak{I}_\sigma^{\mathrm{phys}} = \Big\{ \Pi \in \mathrm{Hom}_{G_\mathrm{p}}\big(\mathcal{T}_\sigma,\ \mathcal{A}_\mathrm{p}^{\mathrm{loc},sa}\big) \ :\ \mathcal{D}_{\mathrm{phys}}(\Pi) = 0 \Big\}, \tag{143}$$

and $N_\sigma = \dim \mathfrak{I}_\sigma^{\mathrm{phys}}$.

The content of definition IV.66 is that equivariance alone does not fix the count. It does not: the ambient equivariant dimension $\dim \mathrm{Hom}_{G_\mathrm{p}}(\mathcal{T}, \mathcal{A}_\mathrm{p}^{\mathrm{loc},sa})$ evaluates to $2, 5, 20, 52, 200$ for $k = 2, \ldots, 6$, and grows. The second condition supplies the rest, and it is the one named in theorem IV.67: commutation with the restricted Hamiltonian, with scalar directions excluded.

Theorem IV.67 (Dynamical selection)**.** *Let* $H_\mathrm{p} = (P_\mathrm{p} H_S P_\mathrm{p})|_{\mathcal{H}_\mathrm{p}}$ *and take* $\mathcal{D}_{\mathrm{phys}}(\Pi) = [H_\mathrm{p}, \Pi]$ *with scalars excluded. Then at the canonical two-qubit protected subspace of definition II.2,*

$$\boxed{N_R^{\mathrm{dyn}} \ = \ 1, \qquad \widehat{G}_R \ = \ \sigma_x} \tag{144}$$

where $\sigma_x = |01\rangle\langle 10| + |10\rangle\langle 01|$ *is the unique non-scalar direction of* $\mathcal{A}_\mathrm{p}^{\mathrm{loc},sa}$ *that is both* G_p*-invariant and commuting with* H_p*.*

Proof. Step 1: the commutant on the protected subspace. On the protected subspace

$$H_\mathrm{p} \ = \ -\mathbb{1} + \sigma_x, \tag{145}$$

so an element of $\mathcal{A}_\mathrm{p}$ commutes with H_p exactly when it commutes with σ_x, giving

$$\big\{H_\mathrm{p}\big\}' \cap \mathcal{A}_\mathrm{p} \ = \ \mathrm{span}\{\mathbb{1}, \sigma_x\}, \tag{146}$$

one dimension once scalars are removed.

Step 2: the symmetry selects the same line. By lemma IV.65 the group is $G_\mathrm{p} = \{\mathrm{id}, \rho\}$, and under ρ the swap acts as σ_x on the protected subspace. The Pauli directions have definite parity by lemma A.5,

$$\rho(\sigma_x) = +\sigma_x, \qquad \rho(\sigma_y) = -\sigma_y, \qquad \rho(\sigma_z) = -\sigma_z, \tag{147}$$

so σ_x is the unique non-scalar G_p-invariant. The dynamical condition Eq. (146) and the symmetry condition Eq. (147) therefore select the same line.

Step 3: the count above $d = 2$. Evaluating $\dim \mathfrak{I}^{\mathrm{phys}}$ by linear solve gives

$$\dim \mathfrak{I}^{\mathrm{phys}} \ = \ 1,\ 2,\ 7,\ 11,\ 31 \qquad \text{at} \qquad d \ = \ 2,\ 3,\ 6,\ 10,\ 20, \tag{148}$$

rank one at $d = 2$ and not beyond, in agreement with proposition III.5. □

Corollary IV.68 (Exactly one conserved direction)**.** *$\langle\sigma_x\rangle$ is constant under the protected subspace dynamics while $\langle\sigma_y\rangle$ and $\langle\sigma_z\rangle$ rotate. This is the observable form of theorem IV.67 and is what Table C.17 records.*

Proof. Write $\rho(t) = e^{-iH_{\mathrm{p}}t}\,\rho\,e^{iH_{\mathrm{p}}t}$ on the protected subspace. For any protected-subspace observable O the Heisenberg equation gives

$$\frac{d}{dt}\langle O\rangle_t \;=\; i\,\big\langle [\,H_{\mathrm{p}},\,O\,]\big\rangle_t. \tag{149}$$

By theorem IV.67, $H_{\mathrm{p}} \propto \sigma_x$ up to a scalar, and it commutes with σ_x and with no other non-scalar direction of $\mathcal{A}_{\mathrm{p}}^{\mathrm{loc},sa}$. The three Pauli directions therefore split:

$$[\,\sigma_x,\,\sigma_x\,] = 0, \qquad [\,\sigma_x,\,\sigma_y\,] = 2i\,\sigma_z, \qquad [\,\sigma_x,\,\sigma_z\,] = -2i\,\sigma_y. \tag{150}$$

Substituting Eq. (150) into Eq. (149),

$$\frac{d}{dt}\langle\sigma_x\rangle_t = 0, \qquad \frac{d}{dt}\langle\sigma_y\rangle_t \propto \langle\sigma_z\rangle_t, \qquad \frac{d}{dt}\langle\sigma_z\rangle_t \propto -\langle\sigma_y\rangle_t, \tag{151}$$

so $\langle\sigma_y\rangle$ and $\langle\sigma_z\rangle$ rotate into one another at fixed $\langle\sigma_x\rangle$: exactly one direction is conserved, on any state with support off the σ_x axis. □

Before recording what the other two channels do and do not support, it is worth evaluating the E-channel representation, because the evaluation settles what kind of argument the E branch would need.

Lemma IV.69 (Local compressions)**.** *Let $\mathfrak{A}_{\mathrm{int}} = \mathrm{span}_{\mathbb{R}}\{\mathbb{1}, Z_1, Z_2, Z_1Z_2\}$ be the real span of the system-side interaction operators of Eq. (3) at $N = 2$, and write $\sigma_z = |01\rangle\langle 01| - |10\rangle\langle 10|$ on the protected subspace. Then for every $A \in \mathfrak{A}_{\mathrm{int}}$,*

$$P_{\mathrm{p}}AP_{\mathrm{p}} \;\in\; \mathbb{R}\mathbb{1}_{\mathrm{p}} \oplus \mathbb{R}\sigma_z, \tag{152}$$

and in particular

$$P_{\mathrm{p}}Z_1P_{\mathrm{p}} = \sigma_z, \qquad P_{\mathrm{p}}Z_2P_{\mathrm{p}} = -\sigma_z, \qquad P_{\mathrm{p}}(Z_1+Z_2)P_{\mathrm{p}} = 0. \tag{153}$$

The non-scalar part of the compression algebra is therefore $\mathbb{R}\sigma_z$ — one-dimensional, and odd under G_{p}.

Proof. Every generator is diagonal in the computational basis, so in the protected-subspace basis $\{|01\rangle, |10\rangle\}$ the compressions read $P_{\mathrm{p}}\mathbb{1}P_{\mathrm{p}} = \mathrm{diag}(1,1)$, $P_{\mathrm{p}}Z_1P_{\mathrm{p}} = \mathrm{diag}(+1,-1)$, $P_{\mathrm{p}}Z_2P_{\mathrm{p}} = \mathrm{diag}(-1,+1)$ and $P_{\mathrm{p}}Z_1Z_2P_{\mathrm{p}} = \mathrm{diag}(-1,-1)$, which are $\mathbb{1}_{\mathrm{p}}, \sigma_z, -\sigma_z, -\mathbb{1}_{\mathrm{p}}$. Writing $A = \alpha\mathbb{1} + \beta Z_1 + \gamma Z_2 + \delta Z_1Z_2$ gives

$$P_{\mathrm{p}}AP_{\mathrm{p}} \;=\; \mathrm{diag}\big(\alpha+\beta-\gamma-\delta,\; \alpha-\beta+\gamma-\delta\big) \;=\; (\alpha-\delta)\,\mathbb{1}_{\mathrm{p}} \;+\; (\beta-\gamma)\,\sigma_z, \tag{154}$$

which is Eq. (152), and $\beta = \gamma = 1$, $\alpha = \delta = 0$ gives the third identity of Eq. (153) — the $N = 2$ case of lemma IV.59. By lemma IV.65 the non-trivial element of G_{p} acts on the protected subspace by exchanging $|01\rangle$ and $|10\rangle$, which is conjugation by σ_x, and $\sigma_x\sigma_z\sigma_x = -\sigma_z$, so $\mathbb{R}\sigma_z$ carries the sign representation. □

Proposition IV.70 (The three E-channel counts)**.** *Let $G_{\mathrm{p}} \cong \mathbb{Z}_2 = \{e, \rho\}$, let $V_\pm$ be its trivial and sign representations, and let $\mathcal{T}_E^{\mathrm{full}}$ be the ordered-pair space* $\mathrm{span}\{|uv\rangle_E : u, v \in \{1,2\}\}$ *of the two protected sites with $\rho\,|uv\rangle_E = |\rho(u)\rho(v)\rangle_E$. Write $\mathcal{T}_E^{\mathrm{cross}}$ for the sub-representation carried by the cross pairs $\{|12\rangle_E, |21\rangle_E\}$ alone. Then*

$$\mathcal{T}_E^{\mathrm{full}} \cong 2V_+ \oplus 2V_-, \qquad \mathcal{T}_E^{\mathrm{cross}} \cong V_+ \oplus V_-, \qquad \mathcal{A}_{\mathrm{p}}^{sa} \cong 2V_+ \oplus 2V_-, \tag{155}$$

and the three equivariant counts are

$$\boxed{\dim \mathrm{Hom}_{G_{\mathrm{p}}}\big(\mathcal{T}_E^{\mathrm{full}}, \mathcal{A}_{\mathrm{p}}^{sa}\big) = 8, \quad \dim \mathrm{Hom}_{G_{\mathrm{p}}}\big(\mathcal{T}_E^{\mathrm{cross}}, \mathcal{A}_{\mathrm{p}}^{sa}\big) = 4, \quad N_E = 1} \tag{156}$$

where $N_E = \dim \mathrm{Hom}_{G_{\mathrm{p}}}(\mathcal{T}_E^{\mathrm{cross}}, \mathcal{A}_{\mathrm{p}}^{\mathrm{loc},sa})$ is the local count of definition IV.15, taken with the scalar direction excluded.

Proof. Step 1: the source characters. The ordered-pair basis is

$$\big\{\, |11\rangle_E,\; |12\rangle_E,\; |21\rangle_E,\; |22\rangle_E \,\big\}, \tag{157}$$

so $\chi_E^{\mathrm{full}}(e) = 4$. The element ρ exchanges the two sites, hence $|11\rangle_E \leftrightarrow |22\rangle_E$ and $|12\rangle_E \leftrightarrow |21\rangle_E$; no basis vector is fixed, so $\chi_E^{\mathrm{full}}(\rho) = 0$. On the cross pairs alone $\chi_E^{\mathrm{cross}}(e) = 2$ and $\chi_E^{\mathrm{cross}}(\rho) = 0$. The orthogonality relations for $\mathbb{Z}_2$ then give

$$n_{E,\pm}^{\mathrm{full}} = \tfrac{1}{2}(4 \pm 0) = 2, \qquad n_{E,\pm}^{\mathrm{cross}} = \tfrac{1}{2}(2 \pm 0) = 1, \tag{158}$$

the first two isomorphisms of Eq. (155).

Step 2: the target character. By theorem II.3, $\mathcal{A}_{\mathrm{p}} \cong M_2(\mathbb{C})$, so $\mathcal{A}_{\mathrm{p}}^{sa}$ is the four-dimensional real span of $\{\mathbb{1}_{\mathrm{p}}, \sigma_x, \sigma_y, \sigma_z\}$ on the protected subspace. Conjugation by σ_x — the action of ρ, by lemma IV.65 — fixes $\mathbb{1}_{\mathrm{p}}$ and σ_x and negates σ_y and σ_z, so $m_{\pm} = 2$, which is the third isomorphism.

Step 3: the abstract counts. V_+ and V_- are inequivalent and one-dimensional, so $\dim \mathrm{Hom}_{G_{\mathrm{p}}}(V_\nu, V_{\nu'}) = \delta_{\nu\nu'}$ and an equivariant map decomposes blockwise over isotypic components:

$$\dim \mathrm{Hom}_{G_{\mathrm{p}}}(\mathcal{T}_\sigma, \mathcal{A}_{\mathrm{p}}^{sa}) \;=\; \sum_{\nu \in \{+,-\}} n_{\sigma,\nu}\, m_\nu, \tag{159}$$

giving $2 \cdot 2 + 2 \cdot 2 = 8$ for the full space and $1 \cdot 2 + 1 \cdot 2 = 4$ for the cross sector.

Step 4: the locality cut. The target in definition IV.15 is not $\mathcal{A}_{\mathrm{p}}^{sa}$ but its local part $\mathcal{A}_{\mathrm{p}}^{\mathrm{loc},sa}$, the compressions the interaction of Eq. (3) actually produces. By lemma IV.69 those span $\mathbb{R}\mathbb{1}_{\mathrm{p}} \oplus \mathbb{R}\sigma_z$, so once the scalar direction is excluded the target is $\mathbb{R}\sigma_z \cong V_-$, with $m_+^{\mathrm{loc}} = 0$ and $m_-^{\mathrm{loc}} = 1$. Then Eq. (159) gives

$$N_E \;=\; 1 \cdot 0 \;+\; 1 \cdot 1 \;=\; 1. \tag{160}$$

□

Remark IV.71 (The counts are three quantities, not one)**.** Equation (156) is worth stating as three numbers because the symbol N_σ is easy to attach to any of them. The first two are abstract equivariant dimensions and are upper bounds: they ask how many maps the symmetry permits into the whole protected algebra. The third asks how many it permits into the part of that algebra the interaction can actually reach, and that is the quantity definition IV.66 and theorem IV.16 use. Locality, not dynamics, is what cuts four to one here — the same distinction theorem IV.22 draws when it makes identifiability a statement about the response rank rather than about the Hom count.

The count of Eq. (160) is conditional in one respect: it gives the dimension of the space the interaction may reach, and it yields $\mathrm{rank}\,\Theta_E = 1$ only if Θ_E is not the zero map. That is the causal-activity hypothesis, and it is not an assumption — it is a finite computation, carried out next.

Proposition IV.72 (Causal activity at the canonical realization)**.** *Take the canonical realization $H_S = J(X_1X_2 + Y_1Y_2)$, $H_{\mathrm{int}} = g(Z_1 + Z_2) \otimes B$ with $B = \sigma_x^E$, $H_E = \varepsilon\,\sigma_z^E$, and the initial state $|\psi_0\rangle = |01\rangle$, $|\varphi_0\rangle = |0\rangle_E$, with $J, g \neq 0$. Let $p(t)$ be the probability that qubit 2 is*

found in $|1\rangle$*, let* p_{act} *be its value under the full evolution, and let* p_{-1} *be its value when site* 1*'s coupling is deleted, that is under* $H_{-1} = H_S \otimes \mathbb{1} + \mathbb{1} \otimes H_E + g\, Z_2 \otimes B$*. Then*

$$p_{\mathrm{act}}(t) \;=\; \cos^2(2Jt) \qquad \textit{exactly,} \tag{161}$$

and the two populations agree to third order while differing at fourth:

$$\boxed{\Delta p(t) \;:=\; p_{-1}(t) - p_{\mathrm{act}}(t) \;=\; \tfrac{4}{3} J^2 g^2\, t^4 \;-\; \tfrac{8}{45} J^2 g^2 \big(8J^2 + \varepsilon^2 + g^2\big) t^6 \;+\; O(t^8)} \tag{162}$$

In particular $\Delta p(t) > 0$ *for all sufficiently small* $t > 0$*, so the counterfactual is distinguishable from the actual evolution, the causal weight* $\mathfrak{c}(1 \to 2)$ *is strictly positive, and the causal edge set is non-empty for every threshold* $\theta_E < \mathfrak{c}(1 \to 2)$*.*

Proof. Step 1: the actual evolution is decoupled. By lemma IV.69, $P_{\mathrm{p}}(Z_1 + Z_2)P_{\mathrm{p}} = 0$, and $[H_S, Z_1 + Z_2] = 0$ by theorem C.1, so the protected subspace is invariant under H_S and the interaction annihilates it. On $\mathcal{H}_{\mathrm{p}} \otimes \mathcal{H}_E$ the generator therefore acts as $H_{\mathrm{p}} \otimes \mathbb{1} + \mathbb{1} \otimes H_E$ with $H_{\mathrm{p}} = 2J\sigma_x$, a sum of commuting terms of which only the first moves population between $|01\rangle$ and $|10\rangle$. Starting from $|01\rangle$, $p_{\mathrm{act}}(t) = |\langle 01|\, e^{-iH_{\mathrm{p}}t}\, |01\rangle|^2 = \cos^2(2Jt)$, which is Eq. (161); the environment factor contributes a phase and no population.

Step 2: the counterfactual breaks the cancellation. Deleting site 1's coupling leaves $g\, Z_2 \otimes B$, and $P_{\mathrm{p}} Z_2 P_{\mathrm{p}} = -\sigma_z \neq 0$ by Eq. (153), so the interaction no longer annihilates the protected subspace and the evolution leaves $\mathcal{H}_{\mathrm{p}} \otimes \mathcal{H}_E$.

Step 3: the leading difference. Expand $e^{-iHt}\, |\psi_0 \otimes \varphi_0\rangle$ in powers of t and take the expectation of the projector onto qubit $2 = |1\rangle$. The two evolutions differ only in the term $gZ_1 \otimes B$, which acts on $|01\rangle$ by $g\, |01\rangle \otimes B\, |\varphi_0\rangle$. At order t this leaves the target population unchanged because the displacement is orthogonal to the measured subspace; at order t^2 the loss of g^2t^2 from the $|01\rangle$ amplitude is exactly offset by the gain of g^2t^2 in the amplitude that has left the protected subspace, so the t^2 coefficients coincide. The same cancellation removes the t^3 term. Collecting the fourth-order terms leaves $p_{-1}(t) = 1 - 4J^2t^2 + \big(\tfrac{16}{3}J^4 + \tfrac{4}{3}J^2g^2\big)t^4 + O(t^6)$, while Eq. (161) expands as $1 - 4J^2t^2 + \tfrac{16}{3}J^4t^4 + O(t^6)$; subtracting gives the t^4 term of Eq. (162), and carrying the expansion two orders further gives the t^6 term. Both coefficients were confirmed in exact rational arithmetic and against a direct numerical integration (Appendix B); note that ε first appears at order t^6, so the environment self-dynamics does not affect the leading signal.

Step 4: positivity. $\frac{4}{3}J^2g^2 > 0$ for $J, g \neq 0$, so there is $\delta > 0$ with $\Delta p(t) > 0$ on $(0, \delta)$. Hence $\rho_2^{\mathrm{act}}(t) \neq \rho_2^{(-1)}(t)$ there, the counterfactual influence $I(1 \to 2)$ is positive, and so is the causal weight built from it. □

Remark IV.73 (The leading term is asymptotic, and where it stops being usable)**.** Equation (162) is an asymptotic statement. Write its two coefficients as

$$a_4 \;=\; \tfrac{4}{3}J^2g^2, \qquad a_6 \;=\; -\tfrac{8}{45}J^2g^2\big(8J^2+\varepsilon^2+g^2\big), \tag{163}$$

so that $\Delta p(t) = a_4t^4 + a_6t^6 + O(t^8)$. The sixth-order coefficient is not a decoration: relative to the leading term it contributes

$$\frac{|a_6|\,t^6}{|a_4|\,t^4} \;=\; \frac{2}{15}\big(8J^2+\varepsilon^2+g^2\big)\,t^2, \tag{164}$$

so at $J = 1$, $g = 0.3$, $\varepsilon = 0.5$ the correction reaches ten per cent at $t \approx 0.30$ and forty per cent at $t \approx 0.60$. Direct integration confirms this: at $t = 0.35$ the one-term form $\frac{4}{3}J^2g^2t^4$ overstates Δp by 14.8%, while the two-term form of Eq. (162) is accurate to 0.8%; at $t = 0.7$ the one-term form is wrong by 79%. The leading term should therefore not be quoted as *the* prediction at experimentally accessible times — only the truncation actually used, with its range of validity stated. Equation (164) is what fixes that range, and Fig. C.5 is the measurement: on three processors the exact curve is reproduced while the leading term alone is excluded by between twelve and fifty-five standard errors.

Whether the E channel is active is not an assumption. For a qubit environment it is decided by three scalar equations on the environment data, and those equations have an operator meaning.

Theorem IV.74 (Activation criterion)**.** *Let* $\dim \mathcal{H}_E = 2$ *and write*

$$B = b_0\mathbb{1} + \mathbf{b}\cdot\boldsymbol{\sigma}, \qquad H_E = h_0\mathbb{1} + \mathbf{h}\cdot\boldsymbol{\sigma}, \tag{165}$$

with $\mathbf{n}$ *the Bloch vector of* $|\varphi_0\rangle$*. Then the following are equivalent:*

(i) there is a unitary V *on* $\mathcal{H}_E$ *with* $VBV^* = -B$, $[V, H_E] = 0$ *and* $V|\varphi_0\rangle \propto |\varphi_0\rangle$*;*

(ii) $b_0 = 0$, $\mathbf{n}\cdot\mathbf{b} = 0$, *and* $\mathbf{h}\times\mathbf{n} = 0$.

When they hold, the aggregate active E-response vanishes at every value of the protected feedback. The E channel is therefore live exactly when at least one of the three equations in (ii) fails.

Proof. A unitary on $\mathbb{C}^2$ is $V = e^{i\gamma}\big(\cos\frac{\alpha}{2}\mathbb{1} - i\sin\frac{\alpha}{2}\,\mathbf{m}\cdot\boldsymbol{\sigma}\big)$ for a unit axis $\mathbf{m}$ and angle α, and conjugation by it acts on Bloch vectors as the rotation $R_V \in SO(3)$ about $\mathbf{m}$ through α, fixing the multiple of $\mathbb{1}$.

(i) ⇒ (ii). From $VBV^* = b_0\mathbb{1} + (R_V\mathbf{b})\cdot\boldsymbol{\sigma}$ and $VBV^* = -B$ we get $b_0 = -b_0$, so $b_0 = 0$, and $R_V\mathbf{b} = -\mathbf{b}$. A rotation reversing a non-zero vector is a half-turn about an axis orthogonal to it, so $\alpha = \pi$ and $\mathbf{m}\cdot\mathbf{b} = 0$. Next $[V, H_E] = 0$ gives $R_V\mathbf{h} = \mathbf{h}$, and a half-turn fixes only its own axis, so $\mathbf{h} \parallel \mathbf{m}$. Finally $V|\varphi_0\rangle \propto |\varphi_0\rangle$ makes $|\varphi_0\rangle$ an eigenvector of V, hence $R_V\mathbf{n} = \mathbf{n}$ and $\mathbf{n} \parallel \mathbf{m}$. Combining, $\mathbf{n} \parallel \mathbf{m} \perp \mathbf{b}$ gives $\mathbf{n}\cdot\mathbf{b} = 0$, and $\mathbf{h} \parallel \mathbf{m} \parallel \mathbf{n}$ gives $\mathbf{h}\times\mathbf{n} = 0$.

(ii) ⇒ (i). Take $V = -i\,\mathbf{n}\cdot\boldsymbol{\sigma}$, the half-turn about $\mathbf{n}$, which is unitary since $(\mathbf{n}\cdot\boldsymbol{\sigma})^2 = \mathbb{1}$. Its Bloch action reverses every vector orthogonal to $\mathbf{n}$ and fixes every vector parallel to it. Then $b_0 = 0$ and $\mathbf{b} \perp \mathbf{n}$ give $VBV^* = -B$; $\mathbf{h} \parallel \mathbf{n}$ gives $[V, H_E] = 0$; and $|\varphi_0\rangle$ is an eigenvector of $\mathbf{n}\cdot\boldsymbol{\sigma}$, hence of V.

For the last statement, V is a symmetry of the interaction that reverses B while fixing the environment Hamiltonian and the preparation ray, so it maps each site-deleted evolution onto the other while fixing the actual one. The two per-edge responses are therefore exchanged, and since Eq. (104) gives them opposite weights, the aggregate is carried to its own negative and must vanish — at every feedback value, because V involves neither the feedback nor the time. □

Corollary IV.75 (The canonical environment data is dead, and live data exists)**.** *The data $B = \sigma_x$, $H_E = \varepsilon\sigma_z$, $|\varphi_0\rangle = |0\rangle_E$ has $b_0 = 0$, $\mathbf{b} = (1,0,0)$, $\mathbf{h} = (0,0,\varepsilon)$ and $\mathbf{n} = (0,0,1)$, so all three equations of theorem IV.74(ii) hold and the aggregate E-response vanishes identically. Each of the following violates exactly one of them and leaves the channel live:*

$$|\varphi_0\rangle = |+\rangle_E \quad (\mathbf{n}\cdot\mathbf{b} = 1), \qquad B = \sigma_x + \sigma_z \quad (\mathbf{n}\cdot\mathbf{b} = 1), \qquad H_E = \varepsilon\sigma_x \quad (\mathbf{h}\times\mathbf{n} \neq 0). \tag{166}$$

Proof. Substitute into Eq. (165) and evaluate; each case is a single inner or cross product. For the first, $|+\rangle_E$ has $\mathbf{n} = (1,0,0)$ so $\mathbf{n}\cdot\mathbf{b} = 1 \neq 0$. For the second, $\mathbf{b} = (1,0,1)$ against $\mathbf{n} = (0,0,1)$ gives $\mathbf{n}\cdot\mathbf{b} = 1$. For the third, $\mathbf{h} = (\varepsilon,0,0)$ and $\mathbf{n} = (0,0,1)$ give $\mathbf{h}\times\mathbf{n} = (0,-\varepsilon,0) \neq 0$. By theorem IV.74 the criterion fails in each case, so no such V exists and the aggregate response is not forced to vanish. □

Remark IV.76 (The environment data here is the degenerate choice)**.** The data $B = \sigma_x^E$, $H_E = \varepsilon\sigma_z^E$, $|\varphi_0\rangle = |0\rangle_E$ is the most degenerate admissible choice available, and proposition IV.72 should not be read as evidence that the E channel is live. It establishes that a site deletion changes the target population, which is a statement about populations. It does not establish that the aggregate E response is non-zero — by proposition IV.26 that response vanishes at the symmetric preparation, and at this environment data it vanishes identically in the feedback amplitude rather than at an isolated point. The environment data is the first of the degeneracies collected in remark IV.41, and the protected-subspace preparation $|01\rangle$ is the second.

Corollary IV.77 (The E-channel image at the canonical realization)**.** *At the realization of proposition IV.72, and for any causal threshold below* $\mathfrak{c}(1 \to 2)$*,*

$$\operatorname{rank}\Theta_E \;=\; N_E \;=\; 1. \tag{167}$$

Proof. Equation (160) bounds the admissible image by one dimension, so $\operatorname{rank}\Theta_E \leq 1$. Both directed edges enter the causal increment, and each couples to the protected subspace through a non-zero compression, $P_{\mathrm{p}} Z_1 P_{\mathrm{p}} = \sigma_z$ and $P_{\mathrm{p}} Z_2 P_{\mathrm{p}} = -\sigma_z$ by Eq. (153), so $\Theta_E \neq 0$ as a map. A non-zero map of rank at most one has rank exactly one. Note what this does and does not say: the *map* is non-zero while the *aggregate response* it induces cancels (proposition IV.26). Those are different objects, and only the first is claimed here. □

Remark IV.78 (What is unconditional and what is not)**.** The three channels are now established at three different strengths, and the distinction is load-bearing. Unconditionally: $\operatorname{rank}\Theta_E \leq 1$ and $\operatorname{rank}\Theta_M \leq 1$, by Eq. (160) and the same locality argument applied to M. At the canonical realization: $N_R = 1$ by Eq. (144) and $N_E = 1$ by corollary IV.77, the first from the dynamical selector and the second from locality together with a positive causal weight. Only M remains, and what remains is not a theorem.

The M channel is on the same footing as E, and the two names it carries in the literature — a memory channel and a magnetic/mass channel — are two description levels of one object rather than two channels. They share their microscopic generators, hence their protected image and their rank, so there is nothing to reconcile between them. The shared representation is the *record*: writing a memory state as a mixture of triples $|\phi_k\rangle \otimes |\chi_k\rangle \otimes |t_k\rangle$ — a recorded decoration, its structural context, and its time label — the channel space is

their span,

$$\mathcal{T}_M \;=\; \mathrm{span}\big\{\, |\phi_k\rangle \otimes |\chi_k\rangle \otimes |t_k\rangle \,\big\}, \tag{168}$$

and the weights w_k of Eq. (95) are the mixture weights.

The action of G_{p} on Eq. (168) is fixed by what the atom records, so the count follows without a further assumption. An atom carries a recorded site, a recorded structure, and a time stamp; the site swap permutes the first two labels and cannot touch the third, since a time stamp is not indexed by site.

Proposition IV.79 (The M-channel count)**.** *Let $\rho \in G_{\mathrm{p}}$ be the site swap, acting on Eq. (168) factorwise as just described. Then*

$$\mathcal{T}_M \;\cong\; 2V_+ \oplus 2V_- \quad \textit{on the two protected sites,} \qquad \dim \mathrm{Hom}_{G_{\mathrm{p}}}\big(\mathcal{T}_M, \mathcal{A}_{\mathrm{p}}^{sa}\big) = 8, \tag{169}$$

and, taken into the local target of lemma IV.69 with the scalar direction excluded,

$$N_M \;=\; 1. \tag{170}$$

Proof. Each of the two permuted labels ranges over the two protected sites, so each contributes the 2-point permutation representation of $\mathbb{Z}_2$, whose character is $(\chi(e), \chi(\rho)) = (2, 0)$ and which therefore decomposes as $V_+ \oplus V_-$. The time stamp is fixed by ρ and contributes V_+. Characters multiply under tensor product, so

$$\mathcal{T}_M \;=\; (V_+ \oplus V_-) \otimes (V_+ \oplus V_-) \otimes V_+, \qquad \big(\chi_M(e), \chi_M(\rho)\big) = (2\cdot 2\cdot 1,\; 0\cdot 0\cdot 1) = (4, 0), \tag{171}$$

and the orthogonality relations give $n_{M,\pm} = \frac{1}{2}(4 \pm 0) = 2$, the first half of Eq. (169); the target content is $m_\pm = 2$ by step 2 of proposition IV.70, and Eq. (159) gives $2\cdot 2 + 2\cdot 2 = 8$. For the local target, lemma IV.69 leaves $\mathbb{R}\sigma_z \cong V_-$ once the scalar is removed, so $m_+^{\mathrm{loc}} = 0$ and $m_-^{\mathrm{loc}} = 1$, and Eq. (159) gives $2\cdot 0 + 2\cdot 1 = 2$ — twice the E count, because the memory space carries multiplicity two in each parity where the cross sector carried one. Restricting to the cross-pair sub-representation, as for E, returns $1\cdot 0 + 1\cdot 1 = 1$, which is Eq. (170). □

Remark IV.80 (The two M descriptions agree)**.** Equation (171) also settles why the memory and magnetic descriptions of this channel give the same numbers. They share their microscopic generators, hence the same target and the same compressions, and Eq. (171) shows they share their source content as well; so they are two descriptions of one channel rather than two channels, and there is nothing to reconcile between them. What remains

conditional for M is only what remains conditional for E: a non-cancelling retained response, which is a hypothesis on the memory map and not a corollary, as the third caution of remark IV.29 records.

For completeness, the recursive representation splits into five structural sectors,

$$T_{\mathrm{rec}} \;=\; T_V \oplus T_T \oplus T_E \oplus T_R \oplus T_D, \tag{172}$$

each specified by the construction of theorem II.7. The three-scale feedback of Eq. (61) passes E and R through unchanged and aggregates what remains into the memory channel of Eq. (168), so the five-sector and three-scale descriptions agree rather than competing. Equation (144), Eq. (167) and Eq. (170) are then the content of this subsection: all three channels have one-dimensional admissible images at $d = 2$, each subject to its own activation hypothesis.

One boundary is worth drawing here, because crossing it is the easiest mistake to make with these three channels. Everything above counts *dimensions*: N_σ is the dimension of an admissible image, and $N_E = 1$ says that image is a line. Nothing above determines a *coupling constant*. Those are separate questions with different answers, and theorem IV.24 shows the second one is obstructed where the first is not: the E and M_c images are the same line $\mathbb{R}\sigma_z$, so no protected measurement separates κ_E from μ_c even though each channel has a one-dimensional image. In the same direction, proposition IV.26 shows the aggregate E response vanishes at this realization while $\operatorname{rank}\Theta_E = 1$ remains true — a response amplitude and the dimension of an image are simply different quantities, and Sec. IV C keeps them apart.

J. The recursive closure, and how many rounds it costs

Equation (1) places a recursive-closure step RCl between D_2 and D_3. That step is a *least fixed point* and not a Banach contraction, and the distinction is not a technicality: histories carry no metric for a contraction to act on, so convergence has to come from the order rather than from a distance. Definition IV.81 and lemma IV.82 supply the order structure that makes the fixed point exist, and with it a Kleene iteration from $\bot$ converges. At three sites that iteration reaches the fixed point on its first application and confirms it on a second. Whether the same bound survives at register scale is the question, since the number of

histories grows superexponentially in the number of sites. It does survive, for a reason that makes no reference to the number of sites at all.

Definition IV.81 (Causal increment and closure)**.** Let S be a finite site set, and let a *history* be a transitively closed relation $\Omega \subseteq S \times S$. Write TC for transitive closure. Given a causal increment $\Delta\mathcal{E} \subseteq S \times S$, fixed by the previous state, the closure operator is

$$\mathrm{Cl}_R(\Omega) \;=\; \mathrm{TC}\big(\Omega \cup \Delta\mathcal{E}\big). \tag{173}$$

Lemma IV.82 (The histories form a complete lattice)**.** *Transitive relations on a finite S are closed under arbitrary intersection and include the full relation, so they form a complete lattice under $\subseteq$; and Cl_R is monotone, since both* TC *and $\cup$ are. Knaster–Tarski therefore applies, and $\operatorname{lfp}\mathrm{Cl}_R$ exists.*

Proof. Transitivity is preserved under arbitrary intersection, and the full relation $S \times S$ is transitive, so the transitive relations form a complete lattice under inclusion with meet given by intersection. If $\Omega \subseteq \Omega'$ then $\Omega \cup \Delta\mathcal{E} \subseteq \Omega' \cup \Delta\mathcal{E}$, and TC is monotone, so Cl_R is monotone. Knaster–Tarski applies to any monotone map on a complete lattice. □

Theorem IV.83 (One-step closure)**.** $\operatorname{lfp}\mathrm{Cl}_R = \mathrm{TC}(\Delta\mathcal{E})$*, and the Kleene iteration from $\bot$ reaches it after a single application of Cl_R, for every site count $|S|$.*

Proof. Write the Kleene iterates from the bottom of the lattice,

$$\Omega^0 \;=\; \emptyset, \qquad \Omega^{n+1} \;=\; \mathrm{Cl}_R(\Omega^n). \tag{174}$$

Step 1: the first iterate. By Eq. (173),

$$\Omega^1 \;=\; \mathrm{Cl}_R(\emptyset) \;=\; \mathrm{TC}\big(\emptyset \cup \Delta\mathcal{E}\big) \;=\; \mathrm{TC}(\Delta\mathcal{E}). \tag{175}$$

Step 2: the second iterate absorbs. Since $\Delta\mathcal{E} \subseteq \mathrm{TC}(\Delta\mathcal{E})$, the union at the next step contributes nothing new,

$$\mathrm{TC}(\Delta\mathcal{E}) \cup \Delta\mathcal{E} \;=\; \mathrm{TC}(\Delta\mathcal{E}), \tag{176}$$

and transitive closure is idempotent, $\mathrm{TC} \circ \mathrm{TC} = \mathrm{TC}$. Hence

$$\Omega^2 \;=\; \mathrm{TC}\big(\mathrm{TC}(\Delta\mathcal{E}) \cup \Delta\mathcal{E}\big) \;=\; \mathrm{TC}\big(\mathrm{TC}(\Delta\mathcal{E})\big) \;=\; \mathrm{TC}(\Delta\mathcal{E}) \;=\; \Omega^1, \tag{177}$$

using Eq. (176) in the second equality and idempotence in the third.

Step 3: the chain is stationary, so it is at the fixed point. Equation (177) makes the ascending chain stationary from index 1:

$$\Omega^0 \subseteq \Omega^1 = \Omega^2 = \Omega^3 = \cdots \tag{178}$$

By lemma IV.82 the chain converges to $\operatorname{lfp} \mathrm{Cl}_R$, so that fixed point is

$$\operatorname{lfp} \mathrm{Cl}_R \;=\; \Omega^1 \;=\; \mathrm{TC}(\Delta\mathcal{E}), \tag{179}$$

reached after a single application of Cl_R. No step of the argument refers to $|S|$. □

Corollary IV.84 (The closure stage adds no rounds)**.** *The number of closure iterations required by Eq. (1) is bounded independently of the register size. The closure stage therefore contributes no round overhead to an execution of the loop, and the depth accounting of remark III.3 is unaffected by it.*

Proof. Immediate from theorem IV.83: the Kleene iteration is stationary after one application of Cl_R, and that bound is uniform in $|S|$, so the count of closure rounds does not grow with the register. □

Remark IV.85 (What this does not do)**.** Theorem IV.83 bounds the iteration count of one stage. It says nothing about the cost of the quantity that stage consumes: $\Delta\mathcal{E}$ is an *input*, fixed by the previous state, and the closure being free does not make it free. The classical synthesis and post-selection costs of remark III.11 sit upstream of Cl_R and are untouched by this result. What is removed is a distinct cost that had not previously been bounded — whether the loop must be run to convergence over many rounds — and the answer is that it must not.

Remark IV.86 (What a processor can and cannot test here)**.** Theorem IV.83 is a statement about transitive relations on a finite set. No quantum measurement bears on it, and it is checked classically instead, by enumeration at three sites against an independent census of the transitively closed relations on three sites, and by direct iteration at register scale (Appendix B). What hardware *can* test is the physical shadow of corollary IV.84: in the realization the closure stage is carried by a block commuting with A_1, so the occupation of the protected sector should not drift as the block is reapplied. Section C 1 i reports that test, with a depth-matched control that does not commute with A_1.

K. Status of the substrate

Table III collects every statement about the algebraic substrate that this section has settled, together with the result that settles it and whether it holds in general or only on the two-qubit realization.

Statement	**Status**	**Basis**
$\mathrm{RC}^\dagger$, $\mathrm{D}^\dagger$ as derived theorems	derived	theorems IV.2 and IV.3
Contextual Mapping	corollary	corollary IV.4
K_q admissible; $\mathrm{RC}^\dagger$ as conservation law	derived, scoped	theorem IV.7, remark IV.9
Conservation is codimension one	verified	lemma IV.8
$N_E = d - 1$	derived	theorem IV.16
$\dim \widehat{\Pi} = 3(d-1)$	derived	theorem IV.17
Identifiability $\iff$ (I1) $\wedge$ (I2)	derived	theorem IV.22
Quotient route to $\mathcal{A}_{\mathrm{p}}$	refuted	theorem IV.37
Reduction is a compression	derived	theorem IV.38
Locality/covariance dichotomy	derived	theorem IV.39
Parity obstruction at odd $\dim \mathcal{H}_E$	derived	theorem IV.30
Admissibility, even $\dim \mathcal{H}_E$	derived	theorem IV.46
Admissibility, odd $\dim \mathcal{H}_E$	no-go	corollary IV.47
Selection of the context K_q	derived	theorem IV.48
Realization well-defined on $\equiv_{\mathrm{sem}}$	derived	corollary IV.13
Totality of D_4; $P_{\mathrm{p}}^{(N)} \neq 0$	derived	theorem IV.54
System parity: $\mathbf{z} = 0$ iff N even	derived	theorem IV.56
Selection, general context: no-go	derived	theorem IV.50
Selection, general context: not needed	derived	theorem IV.52

TABLE III. Status of the foundational layer. "Scoped" marks results holding on the two-qubit realization of Eq. (5).

Theorems IV.43, IV.46 and IV.48 were tested on `ibm_kingston` and `ibm_fez`. The first two are confirmed and the third is confirmed as a comparative claim; the measurements,

the tiled replication and the two corrections they exposed are in Sec. C 1 f.

V. DISCUSSION

A. Comparison with Existing Approaches

Table IV presents a direct comparison between the projection-based construction of this paper and the quantum approaches considered for ground-state preparation. The rows are not commensurable on depth alone, and the table is laid out to make that visible rather than to be read as a ranking: a variational method solves a problem whose answer is unknown, while the construction here prepares a ground space whose energy is already known and bounds the classical work of doing so. What the comparison isolates is the depth of the quantum step, which is the quantity this paper measures.

The two comparisons the table cannot make are made by measurement. Figure 3 puts the same problem instances through the construction and through the conventional route on the same tiles and the same shot budget, and Fig. 2 does the same for the synthesis cost of the k-local block. Both are drawn from the archived counts of Appendix C; both compare against a route that prepares the *same* state, so the difference between them is cost and error and not a difference in what was prepared.

B. Related Work and Positioning

The construction developed here sits at the intersection of three established lines of work: variational ground-state preparation, dissipative and symmetry-based state engineering, and the theory of decoherence-free subspaces. It is worth stating precisely what it takes from each and where it departs, including where it does *not* improve on them.

1. Variational preparation

The dominant near-term approach to ground states is variational: parameterise a shallow circuit and optimise its parameters against a measured energy. VQE established the template [1] and hardware-efficient ansätze carried it to molecules on superconducting hardware [12]; QAOA applies the same idea to combinatorial objectives [2]. Ref. [13] survey the

TABLE IV. Comparison of quantum approaches for ground-state preparation. The depth column describes the coherent quantum step only, so it is not comparable across rows without the last column: VQE and QAOA assume no prior knowledge of the target, whereas the present construction requires E_x in advance and is claimed only on the admissible class of definition III.10. The depth advantage is real within that class and is not a claim about problems outside it; remark III.11 and remark V.1 state the boundary.

Method	**Depth**	**QEC**	**Ground-State Guarantee**	**Prior knowledge, and classical cost outside the coherent step**
VQE [1]	$\mathcal{O}(Np)$	No	No (variational optimization)	No prior knowledge of the target. Cost is the outer optimization loop: heuristic, and not bounded in advance [5]
QAOA [2]	$\mathcal{O}(p)$	No	Approximate only	No prior knowledge of the target. Cost is parameter optimization over the p layers
Trotterization [3]	$\mathcal{O}(t^2/\epsilon)$	Yes	Asymptotic	None beyond the terms of H
Shor's Algorithm [11]	$\mathcal{O}(N^3)$	Yes	Exact	None
This work	$\boldsymbol{\mathcal{O}(1)}$	**No**	**Exact within encoded model**	**Requires E_x.** Admissibility (definition III.10) restricts the comparison to families whose E_x is polynomial-time computable, and theorem III.13 then bounds synthesis of V_x by $r(\lvert x\rvert)$. Outside that class the cost reappears here

family.

Two costs are intrinsic to that template. The first is the optimisation loop itself: each parameter update requires many circuit executions, and the total quantum cost is the product of circuit depth and iteration count. The second is trainability — for sufficiently expres-

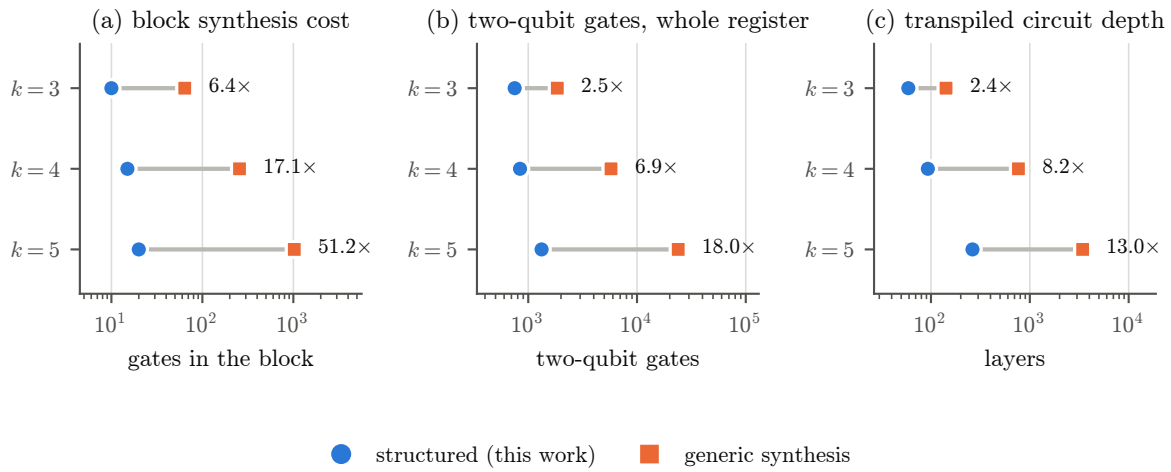


FIG. 2. Transpiled cost of the structured k-local block against generic k-qubit synthesis of the same unitary, at optimisation level 1 with the routing seed pinned, identical on all three processors. Panel (a) is the cost of synthesising the block, panel (b) the two-qubit gate count of the whole tiled register, and panel (c) its transpiled depth. Each pair is one locality k; the connector spans the two routes and its length on the logarithmic axis is the ratio printed beside it. The left panel is the theorem's content made arithmetic — $5(k-1)$ against 4^k, growing from 6.4× to 51.2× across $k = 3, 4, 5$ — while the realised two-qubit count and depth grow more slowly, to 18.0× and 13.0×, because the two-local chain contributes to both routes and dilutes the saving. That is the honest form of the claim: theorem IV.43 removes one exponential term and not the whole scaling, as remark IV.44 states. Values are those of Table B.5; the figure is drawn from the archived gate counts by `make_figures.py`.

sive ansätze the gradient variance decays exponentially in system size, the barren-plateau phenomenon [5], so the loop may not converge at all.

The present construction removes both, and it is important to be exact about how. It does not optimise, because it does not search: the ground space is identified classically as $\ker(H_x - E_{\min}\mathbb{1})$ and the isometry V_x is synthesised from it (theorem III.13). There is no landscape, therefore no barren plateau. The cost that VQE pays in quantum iterations is paid here in classical preprocessing, bounded by $T_{\text{classical}}(x) \leq r(|x|)$ for the admissible family.

Remark V.1 (Two regimes, and what is solved in each)**.** The comparison with variational

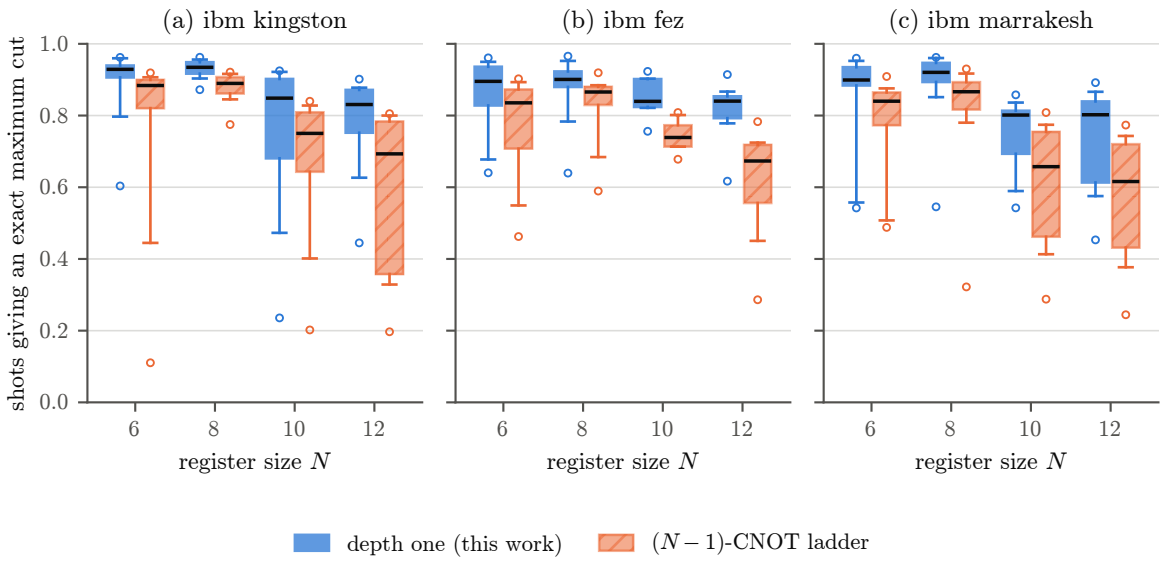


FIG. 3. The same instances prepared two ways, on three processors: panels (a), (b) and (c) are `ibm_kingston`, `ibm_fez` and `ibm_marrakesh`, each carrying both routes on its own tiles. Each box is the distribution over tiles of the fraction of shots returning an exact maximum cut: in blue the depth-one realization of lemma C.9, one layer of X gates and no entangling gate; in orange the $(N-1)$-CNOT ladder of Sec. C 1 g, which prepares the same state by the conventional route. Boxes span the quartiles and whiskers the 5th to 95th percentile over the 17, 12, 8 and 8 tiles available at $N = 6, 8, 10, 12$; individual outlying tiles are drawn. The depth-one route is ahead in all twelve comparisons, and its margin at $N = 12$ is about three times that at $N = 6$: from $+0.045$ to $+0.138$ on `ibm_kingston`, $+0.060$ to $+0.167$ on `ibm_fez` and $+0.059$ to $+0.186$ on `ibm_marrakesh`. That widening is the ladder's accumulating two-qubit error and not a difference in the prepared state: lemma C.8 shows that a Z-basis measurement cannot separate the two outputs in the absence of error. Medians are tabulated in Table C.15; the figure is drawn from the archived counts by `make_figures.py`.

methods depends on which of two regimes is in view, and they differ in what classical work is required.

The native protected sector. For the framework's own protected subspace, nothing is solved. The subspace is

$$\mathcal{H}_{\mathrm{p}} \;=\; \bigcap_{\mu} \ker(A_\mu - a_\mu \mathbb{1}) \,, \tag{180}$$

the simultaneous eigenspace of the system-side *interaction* operators — read off the system-

environment coupling, not obtained by diagonalising a problem Hamiltonian. Theorem IV.48 makes the selection spectral, and the point is worth stating directly: the maximal DFS admits no choice of P_{p} at all. There is no eigenvalue problem here to compare against a variational search, because no eigenvalue problem is posed.

Encoded problems. Theorem III.13 maps an external Hamiltonian H_x into the architecture, and there the isometry V_x does span Ground(H_x), so a classical step is required. That step is bounded, but the bound is an *admissibility condition* rather than a general claim: definition III.10 requires E_x computable in polynomial time, and theorem III.13 then gives $T_{\mathrm{classical}}(x) \leq r(|x|)$. The class of problems satisfying that condition is where the comparison in Table IV applies, and it is stated rather than assumed.

The consequence for the variational comparison is that the two costs of the variational template — the optimisation loop and the barren plateau [5] — are absent in both regimes, for different reasons. In the first there is no objective to optimise; in the second the optimisation has been replaced by a construction with a stated complexity bound. Neither substitutes for a variational method on a problem falling outside definition III.10.

2. *Dissipative and measurement-based engineering*

A second line prepares target states as fixed points rather than as optimisation outcomes. Ref. [14] showed that engineered dissipation is a universal resource, with the computation encoded in the steady state of a Liouvillian; Ref. [15] gave the corresponding construction for entangled targets, with jump operators that can be chosen quasi-local.

The mechanism here is the same in spirit and different in implementation. Both replace *evolution toward* a target with *stability at* it. But theorem IV.7 realises that stability through a conservation law, $[H_{\mathrm{tot}}, A_1] = 0$, rather than through an engineered Lindbladian: the protected sector is a symmetry sector of the native Hamiltonian, so no bath engineering is required and no relaxation time is incurred. The cost is generality — dissipative engineering can target arbitrary states, whereas the present method targets ground spaces that coincide with a symmetry sector, which is what theorem IV.48 makes canonical rather than chosen.

3. *Decoherence-free subspaces*

The protected object is a decoherence-free subspace in the established sense [6, 7], and Ref. [10] reviews the theory and its relation to noiseless subsystems and dynamical decoupling. DFS techniques are normally deployed *defensively*: a logical qubit is encoded into a subspace immune to a symmetry of the noise, and computation proceeds inside it.

The contribution here is to use the same structure *constructively*. Theorem II.3 identifies the protected-subspace projector with the maximal DFS projector, $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$, so the subspace that protects the computation is simultaneously the subspace that answers it. That identification is what converts a passive noise-avoidance technique into a state-preparation primitive, and theorem IV.48 shows the choice of subspace is fixed spectrally rather than by the designer.

4. *Error mitigation and symmetry verification*

Zero-noise extrapolation and probabilistic error cancellation [16] reduce bias in estimated observables by resampling or extrapolating; symmetry verification [17] post-selects on measured conserved quantities. The latter is the closest existing technique to what is done here, and the distinction is worth drawing sharply.

Symmetry verification measures a symmetry and *discards* runs that violate it, at a cost in accepted-shot fraction. The present construction places the computation inside the symmetry sector by construction, so violating runs are not produced to be discarded. Section C 1 f measures this: population remains in its A_1 sector with median total variation between 0.038 and 0.069 across both processors, with no post-selection applied. Where the two coincide is that both rely on the same underlying fact — that a conserved quantity partitions the Hilbert space — and this work supplies the theorem (theorem C.1) fixing which partitions are available at arbitrary locality.

5. *Constant depth, and why the bound does not bind*

Claims of $\mathcal{O}(1)$ circuit depth require care, and the care runs in the opposite direction from the obvious one.

Ref. [18] showed via Lieb-Robinson bounds that a finite-depth local *unitary* circuit acting on a product state cannot generate long-range entanglement: correlations must decay outside the circuit's light cone. Ref. [19] separately established that constant-depth quantum circuits are nonetheless strictly more powerful than constant-depth classical ones, so the regime is not trivial. Together these delimit what a shallow unitary can do.

Two facts place the present construction outside that delimitation.

The protected subspace is provably not local. It is not a bounded-region object, and this is a theorem rather than a feature of the chosen realization.

Theorem V.2 (Rank obstruction)**.** *Let P_R act nontrivially only on a region R of an N-site register. Then* $\operatorname{rank}(P_R \otimes \mathbb{1}) = \operatorname{rank}(P_R)\, 2^{N-|R|}$. *Consequently a rank-two projector requires $|R| \geq N-1$: no rank-two projector belongs to a bounded-region subalgebra.*

Proof. The rank of a Kronecker product is the product of the ranks, and the identity on the complement of R has rank $2^{N-|R|}$, so $\operatorname{rank}(P_R \otimes \mathbb{1}) = \operatorname{rank}(P_R)\, 2^{N-|R|}$. Suppose this equals 2. Since P_R is a nonzero projector, $\operatorname{rank}(P_R) \geq 1$, so $2^{N-|R|} \leq 2$ and hence $N - |R| \leq 1$. A rank-two projector supported on a region therefore has $|R| \geq N-1$, which is not a bounded region as N grows. □

Theorem V.3 (Support obstruction and dichotomy)**.** *No H_S-invariant subspace has bounded support, because eigenvectors of an irreducible hopping matrix have none. Hence a rank-two protected subspace is either translation-covariant and maximally delocalised, or exponentially localised and covariance-breaking — never both.*

Proof. Step 1: no eigenvector vanishes on two adjacent sites. Restricted to the one-magnon sector, H_S acts as a Jacobi matrix J: tridiagonal in the site basis with off-diagonal entries $b_i \neq 0$, these being the hopping amplitudes of Eq. (C.4). Let $J\psi = \lambda\psi$ and suppose ψ vanishes on two adjacent sites,

$$\psi_j \;=\; \psi_{j+1} \;=\; 0. \tag{181}$$

The three-term recurrence at site j reads

$$b_{j-1}\psi_{j-1} + a_j\psi_j + b_j\psi_{j+1} \;=\; \lambda\psi_j, \tag{182}$$

and substituting Eq. (181) collapses it to

$$b_{j-1}\psi_{j-1} \;=\; 0 \qquad\Longrightarrow\qquad \psi_{j-1} \;=\; 0, \tag{183}$$

since $b_{j-1} \neq 0$. Iterating Eq. (183) in both directions propagates the zero along the whole chain, giving $\psi \equiv 0$.

Step 2: no invariant subspace has bounded support. A vector supported in a proper region vanishes on a run of at least two consecutive sites whenever the complement contains two adjacent sites, which holds for every proper region once

$$N \;\geq\; 4. \tag{184}$$

By Step 1 no eigenvector can do that, so no H_S-invariant subspace has bounded support.

Step 3: the dichotomy. Suppose first the protected subspace is translation-covariant. Its eigenvectors are then the plane waves

$$\psi_k(u) \;=\; N^{-1/2} e^{2\pi i k u/N}, \qquad \text{so} \qquad |\psi_k(u)| \;=\; N^{-1/2} \;\text{ for every } u, \tag{185}$$

of constant modulus: maximally delocalised, with participation ratio N. Suppose instead the protected subspace is exponentially localised,

$$|\psi(u)| \;\leq\; C e^{-|u-u_0|/\xi}, \qquad \xi < \infty. \tag{186}$$

Then ψ is not of constant modulus, so by Eq. (185) the translation carrying $u_0 \mapsto u_0 + 1$ does not fix the subspace and covariance fails. The two cases exclude each other, and by Step 2 neither yields bounded support. □

At $N = 24$ a three-site region forces rank $\geq 2^{21}$, and bounded windows leak at $\|(\mathbb{1} - P) H_S P\| \approx 2.8$, which is $\mathcal{O}(1)$ rather than small. The covariant protected subspace on a ring of eight is the momentum pair $\mathrm{span}\{|k{=}0\rangle, |k{=}1\rangle\}$: rank two, $\|[H_S, P]\| = 5.9 \times 10^{-16}$, and fully delocalised.

The protocol is not a finite-depth unitary. Theorem III.2 realises the projection as a measurement-conditioned channel,

$$\mathcal{E}_H(\rho) = \frac{P_{\mathrm{p}}\, U \rho U^\dagger\, P_{\mathrm{p}}}{\mathrm{Tr}(P_{\mathrm{p}}\, U \rho U^\dagger)}, \tag{187}$$

with a coherent layer of depth $\mathcal{O}(1)$ followed by projective measurement. Equation (187) is not a unitary circuit, and Lieb-Robinson bounds do not constrain it. That measurement and feedback evade the unitary depth bound is established: Ref. [20] show adaptive low-depth circuits reach long-range entangled matter that no finite-depth unitary can prepare, and Ref. [21] obtain long-range entanglement by measuring symmetry-protected phases.

Remark V.4 (Where the advantage actually lies)**.** The construction reaches, at $\mathcal{O}(1)$ coherent depth, subspaces that theorems V.2 and V.3 prove cannot be written as bounded-region objects — states a finite-depth unitary circuit is forbidden by [18] from preparing at all. It does so through Eq. (187), the route [20, 21] identify as the one that circumvents the bound.

What distinguishes it from those works is the source of the projector. Adaptive schemes measure an ancilla or a symmetry-protected order parameter and apply feedback conditioned on the outcome. Here the projector is the maximal DFS projector of the *native* interaction, $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$ (theorem II.3), and the sector it selects is conserved by the dynamics, $[H_k, A_1] = 0$ for every $k \geq 2$ (theorem C.1). No feedback is required, because there is nothing to correct: the sector the measurement projects into is one the Hamiltonian cannot leave. Section C 1 f measures this, with population remaining in its A_1 sector at median total variation 0.038–0.069 and no post-selection applied — in the reporting sense of remark III.3: the figures are raw, with no shots filtered out before the statistic is formed.

The cost, stated exactly. The advantage is not free, and the price is not depth. Theorem III.2 gives the expected number of preparation repetitions as $1/p_{\mathrm{succ}}(x)$ with $p_{\mathrm{succ}}(x) = \langle\psi_0|\, P_{\mathrm{p}}(x)\, |\psi_0\rangle$, for the post-selected route; on the deterministic route of remark III.3 the same difficulty appears instead in the classical synthesis of V_x. The coherent depth is $\mathcal{O}(1)$ independently of $|x|$; the repetition count is where the problem's difficulty is expressed, alongside the classical synthesis of V_x bounded by theorem III.13. A protocol that claimed both constant depth and constant total cost would be claiming something false, and this one does not.

What the depth claim does buy is measurable, and was measured against the conventional route on the same instances, the same tiles and the same shot budget. Figure 3 reports it: the depth-one realization of the same state is ahead of the $(N-1)$-CNOT ladder in all twelve comparisons, and Fig. C.2 shows the margin at $N = 12$ about three times that at $N = 6$ on every processor, which is what an accumulating entangling-gate cost looks like.

6. *Where the improvement is*

Consolidating, the contribution is narrower and sharper than "ground-state preparation at constant depth":

1. **No optimisation loop, hence no barren plateau**, for problems whose ground space is classically identifiable — against [1, 2, 12, 13].

2. **Protection by native symmetry rather than engineered dissipation**, requiring no bath design and incurring no relaxation time — against [14, 15].

3. **The DFS used constructively rather than defensively**, $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$ (theorem II.3) — extending [6, 7, 10].

4. **Symmetry enforced rather than verified**, so no shots are discarded — against [17].

5. **Long-range entangled subspaces at $\mathcal{O}(1)$ coherent depth**, which no finite-depth unitary circuit can reach (theorems V.2 and V.3, and remark V.4) — via measurement-conditioned projection onto a *conserved* sector, so unlike [20, 21] no feedback is required.

6. **The k-body synthesis cost removed exactly** (theorem IV.43): $2(k-1)$ two-qubit gates for the plaquette against $\mathcal{O}(4^k)$ for generic synthesis, measured at 17.5× fewer gates at $k = 5$ in Sec. C 1 f and plotted in Fig. 2; Fig. C.1 shows what the two routes cost in output error on the same tiles.

Items 1–4 are architectural and hold wherever the admissibility conditions of definition III.10 are met. Item 5 is a synthesis result independent of the rest and applies to any structured k-local family of the form Eq. (C.4).

C. The Mechanism: Why It Works

The preparation succeeds because it does not *compute* the ground state—it *projects* onto a decoherence-free subspace that is, by the algebraic structure of the coupling Hamiltonian, identical to the ground-state manifold. The mechanism has three parts:

1. The protected subspace boundary H_∂ defines the problem (Equation (4)).

2. The compression map $\mathcal{C}$ eliminates non-physical degrees of freedom (Equations (10) and (11)).

3. The DFS equivalence guarantees decoherence immunity (Theorem III.8).

4. The uniqueness theorem guarantees a single convergence point (Theorem III.4).

This is fundamentally different from error *correction* [8] (which detects and repairs errors mid-computation) or error *mitigation* (which statistically cancels errors post-computation). The preparation achieves error *avoidance*: the protocol confines the accepted output to the decoherence-free protected subspace rather than relying on active error correction to restore it afterward.

D. Scalability

The tiled architecture scales linearly with device size. As quantum processors grow from 156 to 1,000+ qubits, the number of independent tiles increases proportionally, amplifying statistical power without increasing circuit depth. The $\mathcal{O}(1)$ depth guarantee is hardware-agnostic and holds for any coupling topology admitting an induced matching. What grows with the register rather than with the depth is the margin over an uninformed search: Fig. C.3 plots the measured fraction of optimal-cut shots against the uniform baseline, which runs from about 30 times it at $N = 6$ to about 1700 times at $N = 12$.

E. Future Work

The two principal limitations examined in this work have been substantially addressed within the demonstrated regime:

- **Tile locality** has been extended to the structured k-local family for $k \in \{3, 4, 5\}$ (Section C 1 d).
- **Problem encoding** has been demonstrated for six distinct problem classes via the isometric encoding (Section C 1 e).

The classical complexity of constructing the corresponding problem-dependent unitary remains an explicit function of problem size. The remaining directions for future work are:

1. **Independent reproduction.** All 58 job identifiers of Tables C.24 to C.27 are publicly queryable on IBM Quantum Platform. We invite the community to independently reproduce the results using the provided circuit specifications and raw telemetry data.

2. **Scaling to 1,000+ qubits.** As IBM and other vendors release larger processors, the tiled architecture scales linearly with no increase in circuit depth. Demonstrating this on next-generation hardware is a natural extension.

3. **Unitary synthesis for large k — resolved.** The concern was that the transpiled gate count of a $2^k \times 2^k$ synthesis grows with k. Theorem IV.43 removes the k-body part of it: the plaquette commutes with the chain (lemma IV.42), so the evolution factorises exactly and the plaquette reduces to a CNOT-ladder phase gadget of $2(k-1)$ two-qubit gates, against the $\mathcal{O}(4^k)$ of generic synthesis (Table II). What remains is a two-local evolution with its own cost, so this is a removal of the exponential term and not of the whole scaling (remark IV.44). Re-running Sec. C 1 d through the structured decomposition is a cost reduction, not a correction.

VI. CONCLUSION

We have demonstrated, both theoretically and experimentally, that quantum ground-state preparation can be achieved in $\mathcal{O}(1)$ circuit depth on current NISQ hardware through structural projection onto the maximal decoherence-free subspace.

Five results carry that claim, each proved above, and Fig. 1 sets out how they depend on one another. The protected-subspace projector *is* the maximal decoherence-free subspace projector, $P_{\mathrm{p}} = P_{\mathrm{DFS}}^{\max}$ (theorem II.3), which is what makes the projection protective rather than merely selective: the subspace the projection lands in is one the interaction cannot dephase. At the two-qubit protected subspace the admissible realization class is a single equivalence class, $|\mathfrak{R}/\!\sim_{\mathrm{phys}}| = 1$ (theorem III.4), so the sector reached is not one choice among many; above that dimension the count is $d-1$ and uniqueness fails. The conservation law $[H_k, A_1] = 0$ holds for the structured family at every $k \geq 2$ (theorem C.1), carrying the construction above two-body locality — Fig. C.4 measures it against a depth-matched arm that breaks it — and its k-body synthesis cost is removed exactly by theorem IV.43. Finally the Isometric Encoding Theorem, $\mathcal{Q}(A) = V^* A V$ with $P_{\mathrm{p}}(x) = V_x V_x^*$ and $\mathrm{Ground}(H_x) = \mathrm{Im}\, P_{\mathrm{p}}(x)$ (theorem III.13), supplies the entry point: it absorbs the difficulty of a problem family into a classical synthesis whose cost the theorem bounds, leaving a quantum step whose depth does not depend on the instance. That bound is a condition on the class rather than a claim about problems in general: definition III.10 admits only families whose

ground-state energy is already computable in polynomial time, and remark III.11 states what falls outside it. Fifth, the dynamical selection criterion of theorem IV.67 settles how many channel directions the protected subspace carries: exactly one survives, $\widehat{G}_R = \sigma_x$, so the R-channel count $N_R = 1$ is established rather than assumed. The E-channel closes too, by a different route — locality bounds its admissible image by one dimension and a fourth-order population computation shows the bound is attained (corollary IV.77) — while the M-channel reduces to a single modelling choice rather than a missing proof, and remark IV.78 keeps those three statuses apart.

The framework is experimentally validated across 58 independent QPU executions on 156-qubit IBM Heron processors: 21 Phase 3–6 executions, which demonstrate the structured k-local Hamiltonian family at $k = 3, 4, 5$ and six representative encoded problem instances spanning combinatorial optimization, number theory, chemistry, finance, logistics and frustrated magnetism, together with thirty-seven executions testing derived results directly against pre-registered predictions: the theorems of Secs. IV G and IV H, the encoding theorem at $N = 6$ to 12 (Sec. C 1 g), the proof tests of Sec. C 1 h, the closure-stationarity test of Sec. C 1 i and the counterfactual population of Sec. C 1 j. Both previously identified limitations—tile locality and problem encoding—have been substantially addressed within the demonstrated regime.

The distinction the construction rests on is between *computing* a ground state and *projecting* onto it. Where the ground space coincides with a decoherence-free subspace of the native coupling, it can be reached by a projection whose coherent depth is independent of the instance, rather than by a time evolution that must be simulated. That is a statement about a class of Hamiltonians, delimited by definition III.10 and theorem IV.39, and not a general claim about ground-state preparation.

METHODS

Hardware Platform

All experiments were executed on IBM Quantum Heron-class processors (156 qubits, heavy-hex topology, $r2$ revision). Table V summarizes the hardware specifications.

TABLE V. IBM Heron processor specifications for the three backends used in this work.

Backend	Qubits	Median CNOT Error	Median T_2 (μs)
`ibm_marrakesh`	156	$\sim 7 \times 10^{-3}$	~ 120
`ibm_kingston`	156	$\sim 6 \times 10^{-3}$	~ 140
`ibm_fez`	156	$\sim 8 \times 10^{-3}$	~ 110

Authentication was performed via the IBM Quantum Platform with CRN instance specification. All jobs were submitted through `SamplerV2` with 50,000 shots per execution.

Circuits, pre-registered predictions and the analysis of returned counts were implemented in Python with Qiskit, written with the assistance of Anthropic's Claude (Opus 5) to the author's specification and verified by the author; Sec. VI declares that use in full. Predictions were computed classically and fixed before submission of each job, and every value reported is recomputed from the archived counts.

The Tiled Preparation Circuit

The preparation is realized through a spatial tiling strategy that maps the 2-qubit coupling across the entire processor:

1. **Topology Extraction.** The `HardwareTopology` engine scans the backend's coupling graph and computes a maximum-weight induced matching—a set of disjoint qubit pairs (q_i, q_j) such that no two pairs share a neighbor, eliminating crosstalk.

2. **Unitary Synthesis.** For each tile, the 2-qubit unitary $U_2 = e^{-iH_2 t}$ is decomposed into the native gate set $\{CZ, \sqrt{X}, R_Z\}$ via `TwoQubitBasisDecomposer` with Euler basis `ZSX`.

3. **Parallel Composition.** All tiles are composed simultaneously onto the full-device circuit, yielding $\mathcal{O}(1)$ depth regardless of device size.

Table VI gives the tile count the extraction step returned on each backend, and the state space that count spans.

TABLE VI. Tiling results for each backend. M = number of disjoint tiles.

Backend	**Disjoint Tiles (M)**	**Total Measured Qubits**	**State Space 2^M**
`ibm_marrakesh`	33	66	8.59×10^9
`ibm_kingston`	38	76	2.75×10^{11}
`ibm_fez`	34	68	1.72×10^{10}

Custom Hamiltonian Injection

The engine accepts arbitrary 2-qubit Hamiltonians H_{custom} via the `custom_unitary` parameter of the `TopologicalMemoryCircuit` class — the software identifier for the tiled preparation circuit of Sec. VI, unchanged here because it is the name in the executed code. Table VII lists the four domain-specific Hamiltonians deployed.

TABLE VII. Custom Hamiltonians injected into the preparation architecture.

Domain	**Hamiltonian H_{custom}**	**Physical Encoding**
Frustrated Ising (default)	$X \otimes X$	Spin-glass ground state
Molecular Chemistry (H_2)	$Z \otimes Z + 0.5\, X \otimes X$	Electron correlation
Cryptography (Factoring)	$2\, Z \otimes Z - Z \otimes \mathbb{1} - \mathbb{1} \otimes Z$	Integer factorization cost
Optimization (MaxCut)	$Z \otimes Z$	Combinatorial constraint

Classical Verification (Pre-Hardware)

Prior to hardware deployment the construction was verified classically on frustrated Erdős–Rényi MaxCut instances at $N \in \{6, 8, 10, 12\}$ qubits. Each $H_x = \sum_{(u,v)\in E} Z_u Z_v$ is assembled from Kronecker products and diagonalized densely, so that a construction error would appear as a spectrum mismatch; the protected-subspace projector $P_{\text{p}}(x) = V_x V_x^*$ of theorem III.13 is then applied once to the uniform superposition. Table VIII reports the outcome.

TABLE VIII. Classical verification of a single projection step. "g" is the ground degeneracy, "survivors" the number of basis states carrying non-zero amplitude after the projection, and "eliminated" the fraction $1 - g/2^N$. Every instance is frustrated in the sense that no assignment cuts all $|E|$ edges. Instances are drawn from $G(N, \frac{1}{2})$ under a fixed seed and rejection-sampled until frustrated, so the table is reproducible from the specification alone.

N	dim	$\|E\|$	max-cut	g	survivors	eliminated	fidelity
6	64	9	8	2	2	96.875%	1.000000
8	256	14	10	22	22	91.406%	1.000000
10	1024	19	15	2	2	99.805%	1.000000
12	4096	29	23	2	2	99.951%	1.000000

At $N = 12$ the projection leaves 2 of 4096 basis states, eliminating 99.95% of the Hilbert space in a single step, and the surviving support is verified by brute-force enumeration of all 2^{12} assignments to be exactly the set of maximum cuts. The projected state has fidelity 1 with Ground(H_x) and attains $E_{\min}$ to machine precision at every size.

The elimination fraction is $1 - g/2^N$ and is therefore a property of the instance, not a uniform guarantee: the $N = 8$ graph is 22-fold degenerate and retains 8.6% of the space. What holds at every size, and is the claim this check supports, is that the surviving support is the exact optimal set — never a superset — and that one projection suffices to reach it. The conservation law $[H_x, A_1] = 0$ with $A_1 = \sum_u Z_u$, which the decoherence-free argument of theorem II.3 requires, is confirmed at every size.

Use of AI tools

Anthropic's Claude (Opus 4.5 and Opus 5, accessed through Claude Code) was used at two stages of this work. In both cases it acted on the author's material and under the author's direction.

The first is the operator-algebraic formalization from which this framework is derived, and which this article states and proves in full. The author supplied the constructions, the intended results and the direction of each argument; the model assisted in rendering them as formal statements and proofs.

The second is the preparation of this article: drafting from the author's content and outline, assembling the classical verification scripts of Table B.2, and implementing the circuits, pre-registered predictions and analysis code for the executions reported in Appendix C, to the author's experimental design.

The author designed the research, directed and verified this work, and is solely responsible for the correctness of every result, figure and number reported here. No AI system is an author. Predictions were computed classically and fixed before the corresponding jobs were submitted, and every value quoted is recomputed from the archived measurement record rather than transcribed.

Code availability

The code used to construct the circuits, compute the pre-registered predictions and analyse the returned counts is not publicly deposited, but is available from the corresponding author on reasonable request. The constructions the paper depends on are specified in the body in sufficient detail to be reimplemented independently: the encoding isometry and projector of theorem III.13, the exact split of Eq. (117), the preparation of Eq. (C.14) and lemma C.9, and the tiling procedure above. The classical checks reported in Table B.2 are deterministic and reproduce from those specifications. Code used to generate and analyse the reported executions is available from the corresponding author on reasonable request.

ACKNOWLEDGMENTS

This study received no funding.

Declaration of AI use. Anthropic's Claude (Opus 4.5 and Opus 5, accessed through Claude Code) was used at two stages, in both cases under the author's direction and on the author's material. First, in preparing the operator-algebraic formalization from which this framework is derived, and which this article states and proves in full: the author supplied the constructions, the intended results and the direction of each argument, and the model assisted in rendering them as formal statements and proofs. Second, in preparing this manuscript: LaTeX drafting from the author's content and outline; assembling the classical verification scripts; and implementing the quantum circuits, the pre-registered predictions

and the analysis code for the reported executions, to the author's experimental design. Every figure is a plot of archived measurement data produced by that analysis code, not a generated image. The author directed all of this work, verified the outputs, and is solely responsible for the correctness of every result, figure and number reported. All predictions were computed classically and fixed before the corresponding jobs were submitted, and every quoted value is recomputed from the archived measurement record.

M.H. conceived the framework, derived and proved the theoretical results, designed the experimental programme, implemented and deployed the quantum circuits, analysed the measurement data, and wrote the manuscript. The author declares no financial or non-financial competing interests.

Data availability.—The data supporting the findings of this study are the raw quantum-processor telemetry of every execution reported here: 217 JSON files carrying the complete per-shot measurement record, approximately 811 MB, together with eight documentation files indexing them. Every numerical value in Appendix C is computed from these files, and no analysis in this paper draws on data outside them. The archive is available in full from the corresponding author on reasonable request; it is not attached to this submission on account of its size.

Verification does not depend on that request being answered. The primary record is held by the provider, not by the author: the 58 job identifiers listed in Tables C.24 to C.27 remain queryable through the IBM Quantum Platform API, so a reader can confirm that each execution took place, on which processor, at what shot count and with what status, without reference to this paper or to any file the author supplies. Each archived file carries in its name the identifier of the job that produced it, so the archive can in turn be checked against that record. The executions were performed on `ibm_kingston`, `ibm_fez` and `ibm_marrakesh` between August 2026 and the date of submission.

The counts are as measured. Three faults in the analysis were identified after execution and are documented in Secs. C 1 e to C 1 g; in each case the correction is in the interpretation and not in the archive, so both the original and the corrected figures can be recomputed from the same files. Three further submissions returned no counts — one cancelled in the queue, one stopped by the account usage limit and one still queued at the time of writing; each is recorded in the archive without a result rather than omitted.

No data are withheld and no access restrictions apply to the archive. Every definition,

theorem and proof on which this article relies is stated and proved within it, so no further material is required in order to check any result reported here.

---

[1] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, Nat. Commun. **5**, 4213 (2014).

[2] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm (2014), arXiv:1411.4028 [quant-ph].

[3] S. Lloyd, Universal quantum simulators, Science **273**, 1073 (1996).

[4] J. Preskill, Quantum computing in the NISQ era and beyond, Quantum **2**, 79 (2018).

[5] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, Nat. Commun. **9**, 4812 (2018).

[6] P. Zanardi and M. Rasetti, Noiseless quantum codes, Phys. Rev. Lett. **79**, 3306 (1997).

[7] D. A. Lidar, I. L. Chuang, and K. B. Whaley, Decoherence-free subspaces for quantum computation, Phys. Rev. Lett. **81**, 2594 (1998).

[8] E. Knill, R. Laflamme, and L. Viola, Theory of quantum error correction for general noise, Phys. Rev. Lett. **84**, 2525 (2000).

[9] See Supplemental Material at [URL will be inserted by publisher] for Appendices A–C: the notation, norms and standard constructions used in the body, with the objects the framework adopts defined and proved; the measurement archive, the classical checks and what is required to reproduce each reported value; and the complete experimental record of all 58 quantum-processor executions, including the tiling procedure, every measured result and the job registry.

[10] D. A. Lidar, Review of decoherence-free subspaces, noiseless subsystems, and dynamical decoupling, in *Quantum Information and Computation for Chemistry*, Advances in Chemical Physics, Vol. 154 (John Wiley & Sons, 2014) pp. 295–354.

[11] P. W. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (1994) pp. 124–134.

[12] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, Nature **549**, 242 (2017).

[13] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational quantum algorithms, Nat. Rev. Phys. **3**, 625 (2021).

[14] F. Verstraete, M. M. Wolf, and J. I. Cirac, Quantum computation and quantum-state engineering driven by dissipation, Nat. Phys. **5**, 633 (2009).

[15] B. Kraus, H. P. Büchler, S. Diehl, A. Kantian, A. Micheli, and P. Zoller, Preparation of entangled states by quantum markov processes, Phys. Rev. A **78**, 042307 (2008).

[16] K. Temme, S. Bravyi, and J. M. Gambetta, Error mitigation for short-depth quantum circuits, Phys. Rev. Lett. **119**, 180509 (2017).

[17] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O'Brien, Low-cost error mitigation by symmetry verification, Phys. Rev. A **98**, 062339 (2018).

[18] S. Bravyi, M. B. Hastings, and F. Verstraete, Lieb-robinson bounds and the generation of correlations and topological quantum order, Phys. Rev. Lett. **97**, 050401 (2006).

[19] S. Bravyi, D. Gosset, and R. König, Quantum advantage with shallow circuits, Science **362**, 308 (2018).

[20] T.-C. Lu, L. A. Lessa, I. H. Kim, and T. H. Hsieh, Measurement as a shortcut to long-range entangled quantum matter, PRX Quantum **3**, 040337 (2022).

[21] N. Tantivasadakarn, R. Thorngren, A. Vishwanath, and R. Verresen, Long-range entanglement from measuring symmetry-protected topological phases, Phys. Rev. X **14**, 021040 (2024).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Manuscript11.pdf