Skein-Theoretic Membrane Order for Non-Abelian SET Phases

Matrix-Valued Symmetry Diagnostics in Fusion Spaces

Rihaan Shah

July 29, 2026

Abstract

Scalar membrane order parameters detect symmetry fractionalization in several Abelian symmetry-enriched topological (SET) tensor-network models, but they deliberately collapse the fusion-space action of symmetry to a number. This paper formulates a matrix-valued replacement. For a symmetry element s, a disk-like region R, an anyon a, and fusion-space labels μ, ν, we study

$$\mathbf{M}_a^{(s)}(R)_{\mu\nu} = \frac{\langle \Psi_{a,\mu} | U_s(R) W_a(\partial R) | \Psi_{a,\nu} \rangle}{\sqrt{\langle \Psi_{a,\mu} | \Psi_{a,\mu} \rangle \langle \Psi_{a,\nu} | \Psi_{a,\nu} \rangle}}.$$

The observable is defined in four equivalent languages: a G-crossed categorical pairing, a lattice ribbon overlap, an MPO-injective PEPS boundary transfer matrix, and a skein-module evaluation in a thickened surface with a symmetry defect sheet. The main result is a conditional convergence theorem: under standard gap, injectivity, and stable ribbon assumptions, the large-region, vacuum-renormalized matrix $\widehat{\mathbf{M}}_a^{(s)} = \lim_{R\to\infty} \widehat{\mathbf{M}}_a^{(s)}(R)$ is the categorical symmetry intertwiner $X_a(s) : V_a \to V_{\rho_s(a)}$, up to the usual choices of fusion basis, F- and R-symbol gauge, tensor gauge, and MPO gauge. For fixed anyon type this means conjugacy, spectrum, determinant, and traces of words are gauge-invariant diagnostics; when symmetry permutes anyons, singular values and closed skein words are the invariant quantities.

The paper also proves a reconstruction theorem: a sufficiently rich family of membrane matrices and ribbon-intersection matrices reconstructs the anyon permutation, the projective fractionalization class on multiplicity spaces, and the compatible part of the G-crossed associator and half-braiding data, up to categorical equivalence. A fixed-point $D(S_3)$ discussion supplies exact sector data and a minimal $d = 6$ qudit protocol, while the numerical section gives reproducible synthetic benchmarks rather than fabricated PEPS contractions. The central message is that scalar membrane order is a character-like shadow. The full matrix-valued skein membrane retains the symmetry action on non-Abelian fusion spaces and can separate enrichments with identical intrinsic topological entanglement entropy and modular data whenever the scalar trace forgets a noncommuting or multiplicity-sensitive symmetry action.

Data-integrity note. Theorems below are stated with explicit hypotheses. Exact claims about $\mathsf{D}(S_3)$ use group-theoretic sector data only. The 3D spectra, finite-size curves, reconstruction surfaces, and noise plots are generated synthetic benchmarks for the proposed extraction algorithm; they are not reported as new PEPS contraction data or experimental measurements.

Contents

1 Introduction

The organizing difficulty in diagnosing a symmetry-enriched topological phase is that the symmetry action is partly invisible to ordinary topological invariants. The intrinsic phase is described by a braided fusion category, often a unitary modular tensor category $\mathcal{C}$, whose modular data, fusion rules, total quantum dimension, and topological entanglement entropy are insensitive to how a microscopic symmetry G acts on anyons. The enrichment lives instead in a compatible action of G by braided autoequivalences, projective local symmetry actions on anyon sectors, and, when defects are included, a G-crossed braided tensor category $\mathcal{C}_G^\times$ [1–3]. These data are not local order parameters in the Landau sense, and they are not recovered by the modular S, T matrices of the un-enriched topological order.

The recent tensor-network construction of Haller, Xu, Liu, and Pollmann gives a particularly concrete Abelian example [4]. In their decorated toric code path, a membrane order parameter built from a region symmetry operator distinguishes two time-reversal SET phases even though conventional topological quantities do not. That result suggests a broad principle: the right nonlocal operator does not merely ask whether a symmetry is preserved; it asks how symmetry acts on a topological superselection sector. In an Abelian anyon sector this answer is a phase. In a non-Abelian sector it is generally a matrix.

The non-Abelian case forces three upgrades. First, the anyon a may be sent to $\rho_s(a)$, so the symmetry operator is an intertwiner between different sector spaces. Second, even when $\rho_s(a) = a$, the fusion space can have dimension greater than one. A scalar membrane expectation then sees at most a trace or a single matrix element. Third, non-Abelian ribbon operators are skein operators: they carry fusion trees, projectors, half-braidings, and F-moves. A symmetry membrane crossing such a ribbon is not just a number; it is a morphism in a fusion-space category.

This paper proposes the following replacement for scalar membrane order. Let $|\Psi_{a,\mu}\rangle$ be a state in a topological sector containing an anyon a with fusion-space coordinate μ. Let $W_a(\partial R)$ be an a-labelled ribbon, tube-algebra idempotent, or skein projector supported near the boundary of a simply connected region R, and let $U_s(R)$ apply the microscopic symmetry s to degrees of freedom inside R. We define a matrix

$$\mathbf{M}_a^{(s)}(R)_{\mu\nu} = \frac{\langle \Psi_{a,\mu} | U_s(R) W_a(\partial R) | \Psi_{a,\nu} \rangle}{\sqrt{\langle \Psi_{a,\mu} | \Psi_{a,\mu} \rangle \langle \Psi_{a,\nu} | \Psi_{a,\nu} \rangle}}. \tag{1.1}$$

Away from an exact fixed point, the membrane may carry a nonuniversal perimeter factor. We therefore also use the vacuum-renormalized version

$$\widehat{\mathbf{M}}_{b,a}^{(s)}(R) = \frac{\mathbf{M}_{b,a}^{(s)}(R)}{\mathbf{M}_{\mathbf{1},\mathbf{1}}^{(s)}(R)} \tag{1.2}$$

whenever the denominator is nonzero. At a fixed point, or after normalizing the domain-wall tension, $\widehat{\mathbf{M}} = \mathbf{M}$. The extra label b in (1.2) allows anyon permutation; the diagonal formula (1.1) is the $b = a$ component.

The main theorem says that, under the same kind of finite-correlation-length and MPO-injectivity hypotheses that make topological sectors stable in PEPS [5–7], the large-R limit

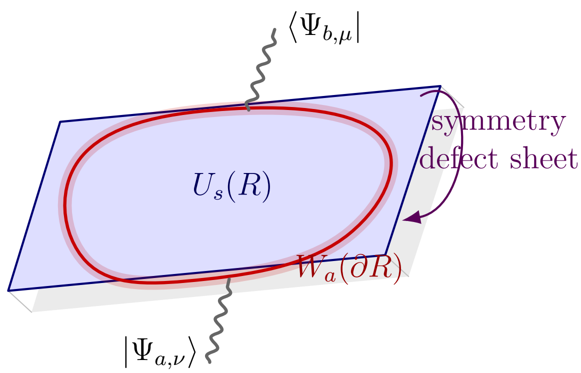

Figure 1: Skein geometry of the membrane matrix. The region symmetry $U_s(R)$ is represented in spacetime by an s-defect sheet whose boundary lies near ∂R. The thick red loop is an a-labelled ribbon or tube-algebra idempotent. The observable is a fusion-space pairing between incoming and outgoing anyon states, not a scalar unless the relevant fusion space is one-dimensional.

of $\widehat{\mathbf{M}}_a^{(s)}(R)$ is the categorical symmetry action on the a-sector. More precisely,

$$\widehat{\mathbf{M}}_{b,a}^{(s)}(R) = \delta_{b,\rho_s(a)}\, B_{\rho_s(a)}^{-1} X_a(s) B_a + O(e^{-\operatorname{dist}(a,\partial R)/\xi})$$

for invertible basis matrices B_a. The matrices B_a encode fusion-basis, F-symbol, R-symbol, tensor, and MPO gauges. Thus the observable is not gauge invariant as a raw matrix; it is gauge covariant. This is the right outcome. Fusion-space matrices are coordinates for morphisms. The invariant content is given by conjugacy classes for fixed sectors, singular values for intertwiners between different sectors, and traces of closed skein words formed from membranes, ribbons, associators, and half-braidings.

The paper has four concrete goals.

1. Define $\widehat{\mathbf{M}}_a^{(s)}$ categorically, on the lattice, in MPO-injective PEPS, and as an element of a skein module $\mathrm{Sk}_{\mathcal{C}}(M;\Delta)$.

2. Prove convergence and gauge covariance under physically allowed changes of fusion basis and tensor-network gauge.

3. Formulate reconstruction algorithms for projective symmetry actions, anyon permutations, and the G-crossed intersection data accessible from membrane-ribbon skeins.

4. Explain, with exact $D(S_3)$ sector data and controlled synthetic numerical benchmarks,

how the diagnostic can be tested without confusing a proposed algorithm with fabricated simulation output.

2 SET Data and Limitations of Scalar Order

2.1 Categorical SET data

Let $\mathcal{C}$ be a unitary modular tensor category describing the intrinsic anyon theory. Simple objects are denoted $a, b, c, \ldots$, fusion multiplicity spaces are

$$V_{ab}^{c} = \mathrm{Hom}_{\mathcal{C}}(c, a \otimes b), \qquad N_{ab}^{c} = \dim V_{ab}^{c},$$

and multi-anyon fusion spaces are written

$$V_{\vec{a}}^{x} = \mathrm{Hom}_{\mathcal{C}}(x, a_1 \otimes a_2 \otimes \cdots \otimes a_n).$$

A bosonic onsite symmetry group G enriches $\mathcal{C}$ by a homomorphism

$$\rho : G \longrightarrow \mathrm{Aut}^{\mathrm{br}}(\mathcal{C})$$

up to natural isomorphism, together with fractionalization data. If $\rho_s(a) = a$, a localized a-anyon can transform projectively:

$$X_a(s)X_a(t) = \eta_a(s,t)X_a(st),$$

where $\eta_a(s,t) \in U(1)$ in a one-dimensional sector, and more generally is a natural automorphism acting on multiplicity spaces. When G permutes anyons, $X_a(s)$ maps V_a to $V_{\rho_s(a)}$, and composition must be compared after following the permutation action. Defect fusion and braiding are encoded in a G-crossed braided extension $\mathcal{C}_G^{\times} = \bigoplus_{s \in G} \mathcal{C}_s$ whose trivial component is $\mathcal{C}$ [1, 3].

The point of an SET diagnostic is not merely to report ρ. Two enrichments may have the same intrinsic $\mathcal{C}$, the same total quantum dimension $\mathcal{D} = \sqrt{\sum_a d_a^2}$, and the same modular data of the trivial sector, but differ by the fractionalization class $[\eta]$, by the symmetry action on fusion multiplicities, or by the G-crossed half-braiding of defects with anyons.

These differences appear as linear maps between topological Hilbert spaces. A scalar order parameter can detect only the part that survives a trace or a chosen diagonal matrix element.

2.2 What a scalar membrane forgets

For a sector a fixed by s, a scalar membrane diagnostic is typically a number such as

$$m_a^{(s)}(R) = \frac{1}{\dim V_a} \operatorname{Tr} \mathbf{M}_a^{(s)}(R)$$

or a normalized expectation value in a specified state. In an Abelian theory, $\dim V_a = 1$, so this is enough to recover a phase. In a non-Abelian theory, the map $X_a(s)$ may contain off-diagonal data or noncommuting relations with other symmetry maps and ribbon operators. The trace loses the following information.

- If $X_a(s)$ and $Y_a(s)$ have the same trace but different commutators with an intrinsic braid or associator matrix, a scalar membrane cannot separate them.
- If $\rho_s(a) \neq a$, the diagonal scalar $m_a^{(s)}$ vanishes by superselection, while the rectangular matrix $\mathbf{M}^{(s)}_{\rho_s(a),a}$ is precisely the desired permutation intertwiner.
- If s is antiunitary, Kramers data are encoded in $X_{\rho_s(a)}(s)\overline{X_a(s)}$, not in a single complex trace.
- If $N_{ab}^c > 1$, the symmetry may act nontrivially on V_{ab}^c even when its trace agrees with another action.

Thus the scalar diagnostic is a character-like compression. Characters are powerful when a full character table is available and the representation category is ordinary semisimple representation theory. SET diagnostics are subtler: the symmetry matrices must be compared with fusion, braiding, defect, and tensor-network gauge data. The full matrix-valued membrane is the smallest object that keeps these comparisons available.

2.3 Selection rules

The first useful output of $\mathbf{M}$ is a selection rule. Let Π_a be the topological projector onto sector a, and suppose U_s acts on superselection sectors by ρ_s. Then

$$\Pi_b U_s \Pi_a = 0 \qquad \text{unless} \qquad b = \rho_s(a).$$

Consequently

$$\mathbf{M}^{(s)}_{b,a}(R) = 0 + O(e^{-L/\xi}) \quad \text{if} \quad b \neq \rho_s(a), \tag{2.1}$$

where L is the separation between the ribbon support and the local anyon operators used to prepare the state. Equation (2.1) gives a direct tomographic route to the anyon permutation: measure all b, a blocks and locate the surviving one. Once the permutation is known, the surviving block is the matrix action on the fusion space.

3 Skein-Module Formulation

3.1 Decorated skein modules

Let M be an oriented three-manifold, usually $M = \Sigma \times [0,1]$ for a spatial surface Σ. Let $\Delta \subset \partial M$ be a set of marked boundary points labelled by objects of $\mathcal{C}$. We define $\mathrm{Sk}_{\mathcal{C}}(M;\Delta)$ to be the complex vector space spanned by framed ribbon graphs embedded in M, with boundary Δ, edges labelled by objects of $\mathcal{C}$, coupons labelled by morphisms, modulo the local graphical relations of the ribbon category: isotopy, fusion-tree resolution, F-moves, braiding moves, duality cups and caps, and bubble removal. This is the standard categorical version of a skein module, specialized to a semisimple ribbon category [2, 8, 9].

To include symmetry, thicken the construction by allowing oriented codimension-one defect sheets labelled by $s \in G$. A ribbon labelled a crossing an s-sheet emerges labelled $\rho_s(a)$. Coupons at intersections are labelled by natural isomorphisms of the G-crossed category. The resulting module is denoted

$$\mathrm{Sk}_{\mathcal{C},G}(M;\Delta) = \bigoplus_{\text{defect-sheet labels}} \mathrm{Sk}_{\mathcal{C}}(M;\Delta;\text{defects}).$$

The membrane order parameter is a skein pairing in this space. The operator $U_s(R)$ is the insertion of an s-sheet $D_s(R)$ whose boundary is $\partial R \times \{1/2\}$. The operator $W_a(\partial R)$ is a framed a-ribbon parallel to that boundary. Their product is not assigned a number until incoming and outgoing fusion states are chosen.

3.2 The skein pairing

Choose bases

$$\{e_\nu\} \subset V_a, \qquad \{f_\mu\} \subset V_{\rho_s(a)}$$

represented by trivalent fusion trees attached to the bottom and top boundaries of $\Sigma \times [0,1]$. The categorical membrane matrix is

$$\left(\mathbf{M}^{(s)}_{b,a}\right)_{\mu\nu} = \langle f_\mu,\, D_s(R) \cup W_a(\partial R) \cup e_\nu \rangle_{\mathrm{Sk}_{\mathcal{C},G}}\,, \tag{3.1}$$

where the bracket evaluates the closed skein after gluing dual boundary states. When $b = \rho_s(a)$, (3.1) is the matrix of a morphism

$$X_a(s) : V_a \longrightarrow V_{\rho_s(a)}.$$

When $b \neq \rho_s(a)$, the skein contains an unmatchable topological charge and therefore evaluates to zero.

3.3 Lattice-ribbon and PEPS readings

In a quantum double lattice model, $W_a(\gamma)$ is a ribbon operator or a central idempotent of the ribbon operator algebra [10]. In a Levin-Wen string-net model, it is a string-net tube encircling the region [11]. In an MPO-injective PEPS, it is represented virtually by the central idempotent of the MPO tube algebra associated with a [6, 7]. These are the same topological operator in different coordinatizations. The matrix-valued membrane construction requires the same extra step in all three languages: do not trace over the fusion space. Keep the full boundary map induced by the symmetry defect.

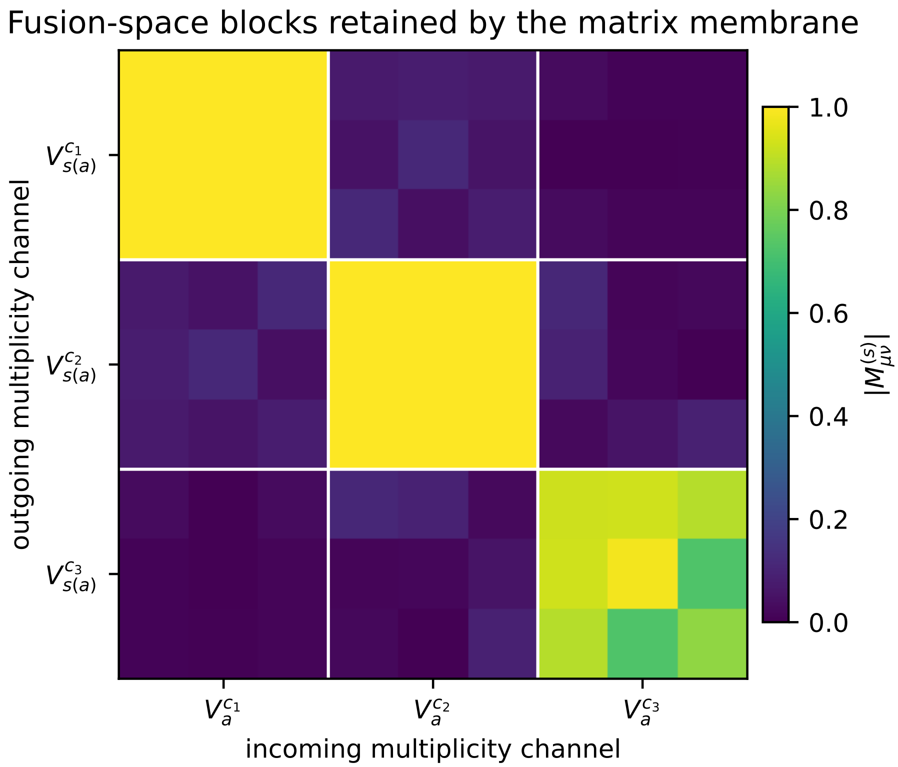

Figure 2: A matrix-valued membrane keeps the block structure among fusion channels. This generated schematic shows the information discarded by a scalar trace: relative phases, off-diagonal channel mixing, and weak but symmetry-forced couplings between multiplicity sectors.

4 Matrix-Valued Membrane Observable

4.1 Definition with anyon permutation

The most useful definition includes an outgoing label b. Let $\mathcal{H}_a(R_0)$ be the low-energy sector obtained by inserting an $a, \bar{a}$ pair near a small disk $R_0 \subset R$, or more generally a multi-anyon sector with total charge a. Choose fusion-space vectors $|\Psi_{a,\nu}\rangle$ and $|\Psi_{b,\mu}\rangle$. Define

$$\mathbf{M}_{b,a}^{(s)}(R)_{\mu\nu} = \frac{\langle \Psi_{b,\mu} | U_s(R) W_a(\partial R) | \Psi_{a,\nu} \rangle}{\sqrt{\langle \Psi_{b,\mu} | \Psi_{b,\mu} \rangle \langle \Psi_{a,\nu} | \Psi_{a,\nu} \rangle}}. \tag{4.1}$$

The user's diagonal expression is the special case $b = a$. The block $b = \rho_s(a)$ is the one that survives in the thermodynamic limit. For $\rho_s(a) = a$, $\mathbf{M}_a^{(s)}$ is a square matrix on the a-sector fusion space. For $\rho_s(a) \neq a$, it is a rectangular or square matrix between canonically isomorphic sector spaces, depending on how the states are prepared.

4.2 Boundary normalization

A region symmetry operator may have a perimeter law expectation even in a symmetric phase. This is a nonuniversal property of the chosen lattice representative, not of the SET order. The stable object is therefore

$$\widehat{\mathbf{M}}^{(s)}_{b,a}(R) = \frac{\mathbf{M}^{(s)}_{b,a}(R)}{\mathbf{M}^{(s)}_{\mathbf{1},\mathbf{1}}(R)}$$

or, equivalently, the matrix obtained after using a transfer-matrix normalization that sets the leading vacuum eigenvalue to one. In exact fixed-point string-net and quantum double models the normalization can be chosen so that the vacuum factor is already one.

Theorem 4.1 (Convergence to the categorical symmetry action)**.** *Let $|\Psi\rangle$ be a symmetric, gapped two-dimensional ground state whose anyon sectors are described by a unitary braided fusion category $\mathcal{C}$. Assume:*

1. *stable ribbon projectors $W_a(\gamma)$ exist and agree with tube-algebra idempotents up to exponentially localized dressing;*
2. *$U_s(R)$ creates, after vacuum normalization, the s-domain wall along ∂R;*
3. *the boundary transfer matrix in each topological sector has a nonzero spectral gap above its leading fixed space;*
4. *the chosen anyon insertions remain a distance L from ∂R.*

Then there are invertible basis matrices B_a and constants $C, \xi > 0$ such that

$$\left\| \widehat{\mathbf{M}}^{(s)}_{b,a}(R) - \delta_{b,\rho_s(a)}\, B^{-1}_{\rho_s(a)} X_a(s) B_a \right\| \le C e^{-L/\xi}. \tag{4.2}$$

Here $X_a(s) : V_a \to V_{\rho_s(a)}$ is the symmetry action on the anyon fusion space in the associated G-crossed category.

Proof. At an exactly soluble fixed point, the equality follows from the graphical calculus: the s-sheet crosses the a-ribbon, relabels it by $\rho_s(a)$, and leaves behind the natural transformation $X_a(s)$ at the fusion tree. If $b \neq \rho_s(a)$, the closed skein has a nontrivial unmatched charge and vanishes. If $b = \rho_s(a)$, the remaining closed skein evaluates to the matrix element of $X_a(s)$.

For a finite-correlation-length representative in the same phase, quasi-adiabatic continuation dresses fixed-point ribbons into exponentially localized ribbon operators and maps the fixed-point symmetry wall to the microscopic region symmetry up to operators supported near ∂R. In an MPO-injective PEPS description, the same statement is the pulling-through relation for the symmetry MPO and the topological MPO idempotent. The contraction outside the localized anyon insertions reduces to a boundary transfer matrix. After dividing by the vacuum eigenvalue, all subleading boundary eigenvectors are suppressed by $e^{-L/\xi}$, where ξ is the boundary correlation length. The only remaining freedom is the choice of fusion basis and tensor gauge, represented by B_a and $B_{\rho_s(a)}$. This gives (4.2). A detailed transfer-matrix version is given in appendix A. $\square$

4.3 Consequences

Three immediate consequences matter in applications.

Proposition 4.2 (Anyon-permutation readout)**.** *Under the hypotheses of theorem 4.1, the support pattern of blocks* $\{\widehat{\mathbf{M}}^{(s)}_{b,a}\}_b$ *recovers* $\rho_s(a)$ *up to exponentially small finite-size errors.*

Proof. The limiting block is zero for all $b \neq \rho_s(a)$ and invertible on the surviving sector if the symmetry does not annihilate the superselection sector. The finite-size estimate follows from (4.2). $\square$

Proposition 4.3 (Kramers and antiunitary selection)**.** *If* T *is antiunitary and* $\rho_T^2(a) = a$, *the Kramers matrix on* V_a *is*

$$\Theta_a^2 = X_{\rho_T(a)}(T)\,\overline{X_a(T)}.$$

It is recovered from the product $\widehat{\mathbf{M}}^{(T)}_{\rho_T^2(a),\rho_T(a)}\overline{\widehat{\mathbf{M}}^{(T)}_{\rho_T(a),a}}$ *up to conjugacy.*

The overline is not optional: antiunitary symmetry complex conjugates the incoming fusion coordinates. A scalar membrane can report the trace of $X_a(T)$, but the Kramers sign or higher-dimensional Kramers structure is contained in Θ_a^2.

5 Gauge Covariance

5.1 Fusion-basis changes

Let $Q_a : V_a \to V_a$ and $Q_{\rho_s(a)} : V_{\rho_s(a)} \to V_{\rho_s(a)}$ be unitary changes of basis. Matrix elements transform as

$$\widehat{\mathbf{M}}^{(s)}_{b,a} \longmapsto Q_b^{-1} \widehat{\mathbf{M}}^{(s)}_{b,a} Q_a \tag{5.1}$$

for unitary s, and as

$$\widehat{\mathbf{M}}^{(s)}_{b,a} \longmapsto Q_b^{-1} \widehat{\mathbf{M}}^{(s)}_{b,a} \overline{Q_a}$$

for antiunitary s. Thus the raw entries of $\widehat{\mathbf{M}}$ are not invariants. They are coordinates for a morphism.

5.2 F- and R-symbol gauge

A categorical gauge transformation changes bases in every trivalent space $\mathrm{Hom}(c, a \otimes b)$. If $u^c_{ab} \in U(N^c_{ab})$, then

$$F \longmapsto uFu^{-1}, \qquad R \longmapsto uRu^{-1}$$

with the usual placement of indices. The fusion states $|\Psi_{a,\mu}\rangle$ are built from these trivalent bases, so $\widehat{\mathbf{M}}$ transforms by the induced matrix Q_a. This is exactly (5.1). Consequently a correct diagnostic cannot depend on a single off-diagonal entry. It may depend on the conjugacy class of a closed word such as

$$\mathrm{Tr}\Big(X_a(s) F^d_{abc} X_d(t)^{-1} R_{bc} X_b(t) X_c(t)\Big)\,,$$

because every Q cancels after the skein is closed.

5.3 PEPS and MPO gauges

For a PEPS tensor A, a virtual gauge change

$$A^i_{\alpha\beta\gamma\delta} \longmapsto \sum_{\alpha'\beta'\gamma'\delta'} (G_1)_{\alpha\alpha'} (G_2)_{\beta\beta'} (G_3^{-1})_{\gamma'\gamma} (G_4^{-1})_{\delta'\delta} A^i_{\alpha'\beta'\gamma'\delta'}$$

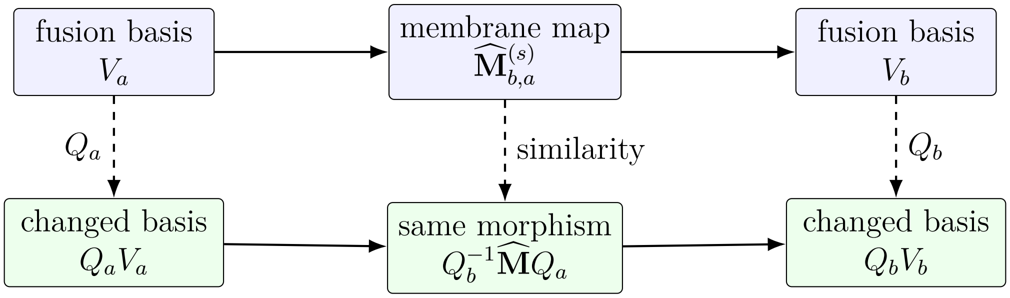

Figure 3: Gauge covariance. A matrix membrane is a morphism written in chosen coordinates. Basis changes, F/R-symbol gauge changes, and PEPS/MPO gauges act by similarity or source-target equivalence. Closed skein words are invariant.

does not change the physical wavefunction. MPO tensors have a parallel gauge freedom on their virtual bonds. Tube-algebra idempotents transform by similarity inside the virtual boundary algebra, and so do the extracted sector matrices. The observable is therefore PEPS-gauge covariant. Its invariant content must be formed after either fixing a canonical tube basis or closing all virtual indices.

Theorem 5.1 (Gauge-covariant invariant content)**.** *For a unitary symmetry element s, the large-region matrix membrane has the following invariant content.*

1. *If $\rho_s(a) = a$, the conjugacy class of $\widehat{\mathbf{M}}^{(s)}_a$ is invariant under categorical and tensor-network gauge. Hence its spectrum, determinant, trace, and traces of powers are invariant.*

2. *If $\rho_s(a) = b \neq a$, the singular values of $\widehat{\mathbf{M}}^{(s)}_{b,a}$ and traces of closed words returning to a are invariant.*

3. *For a collection of symmetry and ribbon matrices, simultaneous conjugacy classes and traces of closed skein words are invariant. These contain strictly more information than the separate spectra of individual matrices.*

For antiunitary s, the corresponding invariant words use complex conjugation on each antiunitary edge.

Proof. The first two statements follow from (5.1). If source and target are the same sector, $Q_a^{-1}\widehat{\mathbf{M}}Q_a$ has the same conjugacy class. If source and target are different, independent unitary changes act on the left and right; singular values are invariant under this two-sided action. A closed skein word is a cyclic product whose adjacent source-target basis changes cancel. The antiunitary case is the same calculation in the semilinear category, with conjugation applied to the incoming coordinates. □

6 Reconstruction Theorems

6.1 Recovering projective factors

Suppose the blocks have identified $\rho_s(a)$ for every simple a and every symmetry generator. Choose representatives $X_a(s)$ from the limiting matrices. Composition gives

$$\Omega_a(s,t) = X_{\rho_t(a)}(s) X_a(t) X_a(st)^{-1}. \tag{6.1}$$

For a simple isolated anyon without multiplicity, $\Omega_a(s,t) \in U(1)$. In fusion spaces with multiplicity it is a unitary natural automorphism compatible with all fusion vertices. Under $X_a(s) \mapsto Q^{-1}_{\rho_s(a)} X_a(s) Q_a$, $\Omega_a(s,t)$ changes by a coboundary. Thus its cohomology class is gauge-invariant, while its matrix representative is gauge-covariant.

6.2 Intersections and G-crossed data

The membrane matrix alone recovers the action on sector spaces. To recover how the action is compatible with associativity and braiding, one must let membranes intersect trivalent ribbon networks. A typical intersection matrix is

$$\mathcal{I}^{(s)}_{abc;d} = X_d(s)\, F^d_{abc} \left(X_a(s)^{-1} \otimes X_b(s)^{-1} \otimes X_c(s)^{-1} \right), \tag{6.2}$$

written schematically after choosing bases. If s crosses a braid, the analogous object compares $X_{a\otimes b}(s) R_{ab}$ with $R_{\rho_s(a)\rho_s(b)}(X_a(s) \otimes X_b(s))$. These are not new arbitrary observables; they are the skein evaluations forced by membrane-ribbon intersections.

Theorem 6.1 (Categorical tomography from membrane skeins)**.** *Assume $\mathcal{C}$ is known as a unitary braided fusion category up to categorical gauge, and assume the convergence theorem holds for all simple anyons and for a generating set of trivalent and braided ribbon networks.*

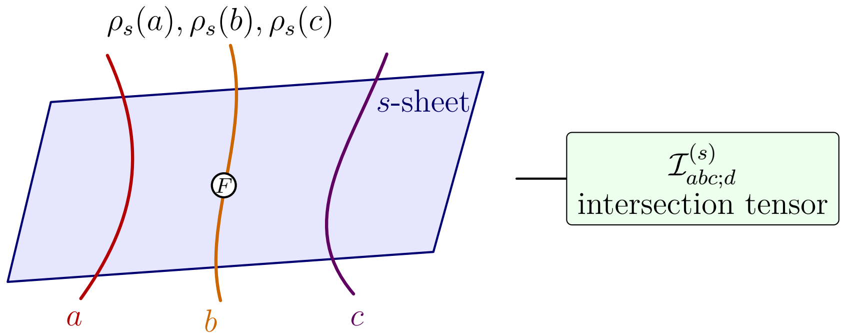

Figure 4: Membrane-ribbon intersections produce tensors, not just expectation values. Pushing the symmetry sheet through a trivalent skein compares the symmetry action on individual anyons with the action on the fused channel. These intersection tensors are the measurable shadows of G-crossed associativity and half-braiding data.

Then the family of large-region observables

$$\{\widehat{\mathbf{M}}^{(s)}_{b,a},\ \mathcal{I}^{(s)}_{abc;d},\ \mathcal{B}^{(s)}_{ab}\}_{a,b,c,d,s}$$

reconstructs:

1. *the anyon permutation action* $\rho : G \to \mathrm{Aut}^{\mathrm{br}}(\mathcal{C})$*;*
2. *the projective fractionalization class* $[\Omega]$ *on one-anyon and multi-anyon multiplicity spaces;*
3. *the defect compatibility data visible in the* G*-crossed associator and half-braiding, restricted to the measured fusion subcategory;*

up to coboundary and categorical equivalence. If the measured ribbon networks generate all simple objects and all multiplicity spaces of $\mathcal{C}$*, the reconstruction is complete for the unitary* G*-crossed data modulo the standard obstruction ambiguities.*

Proof. The block support gives ρ by proposition 4.2. Once ρ is known, products of membrane matrices give the factors $\Omega_a(s,t)$ in (6.1). Changing the representative matrices changes Ω by the usual coboundary, so the cohomology class is fixed. Intersections with trivalent vertices determine the naturality constraints comparing $X_a(s)\otimes X_b(s)$ with $X_c(s)$ on each V^c_{ab}. Intersections with braids determine the corresponding G-crossed half-braiding constraints. Since a semisimple ribbon category is generated by simple objects, trivalent vertices, and braids modulo the pentagon and hexagon relations, equality of all closed skein evaluations on the generating family implies equivalence of the reconstructed data on the measured subcategory. If the generating family is complete, the reconstructed G-crossed structure is complete up to the categorical gauge transformations already described. $\square$

6.3 Algorithm

The reconstruction algorithm is deliberately linear-algebraic.

1. Prepare basis states $\{|\Psi_{a,\nu}\rangle\}$ for each sector a and measure all blocks $\widehat{\mathbf{M}}^{(s)}_{b,a}(R)$.

2. Threshold the Frobenius norms of the blocks to infer $\rho_s(a)$.

3. For each surviving block, polar-decompose the matrix and remove the positive boundary dressing. The unitary factor estimates $X_a(s)$.

4. Compose estimated $X_a(s)$'s to obtain $\Omega_a(s,t)$, then reduce by coboundary minimization over basis changes Q_a.

5. Measure trivalent and braided intersections to solve the naturality and half-braiding constraints.

6. Validate by predicting membrane matrices for unmeasured words and comparing with held-out measurements.

The final step is essential. A family of random matrices can be made to match individual spectra. It cannot generally satisfy the fusion, braid, and group-law constraints on held-out skein words unless it is gauge-equivalent to the correct categorical action.

7 $\mathsf{D}(S_3)$ Realization

7.1 Exact sector data

The quantum double $\mathsf{D}(S_3)$ is the Drinfeld center $\mathcal{Z}(\mathrm{Vec}_{S_3})$. Its simple objects are labelled by pairs (C,π), where C is a conjugacy class of S_3 and π is an irreducible representation of the centralizer of a representative of C [10]. There are eight anyon types. With the conventional labels $A,\dots,H$, their dimensions and topological spins are:

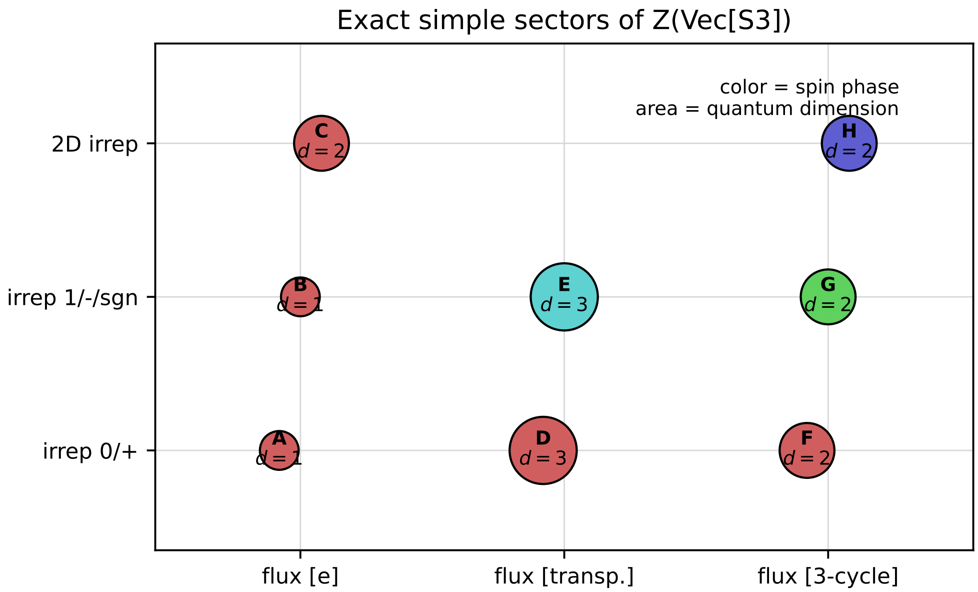

Figure 5: Exact simple-sector map for $\mathsf{D}(S_3)$, generated from the conjugacy-class and centralizer-irrep labelling. The plot uses marker area for quantum dimension and color for spin phase. No F- or R-symbol data are invented in this panel.

label	conjugacy class	centralizer irrep	quantum dimension	spin
A	$[e]$	trivial	1	1
B	$[e]$	sign	1	1
C	$[e]$	standard	2	1
D	transpositions	$+$ of $\mathbb{Z}_2$	3	1
E	transpositions	$-$ of $\mathbb{Z}_2$	3	-1
F	3-cycles	0 of $\mathbb{Z}_3$	2	1
G	3-cycles	1 of $\mathbb{Z}_3$	2	$e^{2\pi i/3}$
H	3-cycles	2 of $\mathbb{Z}_3$	2	$e^{4\pi i/3}$

The matrix membrane can be evaluated at a fixed point using quantum double ribbon projectors. If $g \in C$ and Z_g is its centralizer, the charge projector inside a fixed flux sector has the character form

$$P_{(C,\pi)} = \frac{\dim \pi}{|Z_g|} \sum_{h \in Z_g} \chi_\pi(h^{-1})\, B_g^h, \tag{7.1}$$

where B_g^h denotes the ribbon algebra basis element measuring flux conjugate to g and applying the centralizer action h. Formula (7.1) is the starting point for exact finite-system evaluation: the symmetry membrane is inserted as an automorphism or defect wall, and the matrix elements are overlaps between projected ribbon states.

7.2 What is and is not claimed for $\mathsf{D}(S_3)$

$\mathsf{D}(S_3)$ is an ideal testbed because it is non-Abelian but finite, admits explicit ribbon operators, and can be encoded in $d = 6$ group-valued qudits. The two-qudit protocol of Byles, Forbes, and Pachos demonstrates how braiding and fusion evolutions of $\mathsf{D}(S_3)$ anyons can be implemented using creation and measurement operations [12]. That architecture is well matched to the matrix membrane: the same controlled overlaps that resolve ribbon projectors can be used to read off matrix elements rather than only scalar probabilities.

This paper does not print a table of $\mathsf{D}(S_3)$ F- and R-symbols. Doing so without deriving or importing a verified convention would violate the data-integrity rule. Instead, it gives the exact sector labels, the projector formula (7.1), and a measurement protocol. A companion codebase can instantiate a chosen F/R gauge and then verify gauge covariance by applying random basis changes in every multiplicity space.

8 Tensor-Network Evaluation

8.1 MPO-injective representation

In an MPO-injective PEPS, topological order is encoded in a set of MPO symmetries $\{O_a\}$ satisfying pulling-through equations [6, 7]. The virtual boundary algebra generated by these MPOs contains central idempotents labelled by anyon sectors. A physical region R defines a linear map from the virtual boundary space to the physical bulk. Inserting a topological ribbon around ∂R means inserting the corresponding idempotent on the virtual boundary.

For SET phases, the physical symmetry U_s is represented virtually by a graded MPO domain wall [13]. The contraction defining $\mathbf{M}^{(s)}_{b,a}(R)$ becomes a boundary transfer matrix with three insertions:

$$\mathbb{E}_R\big[p_b\, O_s\, p_a\big],$$

where p_a and p_b are tube idempotents and O_s is the symmetry-domain wall MPO. The pulling-through rule gives $p_b O_s p_a = 0$ unless $b = \rho_s(a)$, and gives the virtual matrix $X_a(s)$ on the surviving fixed space.

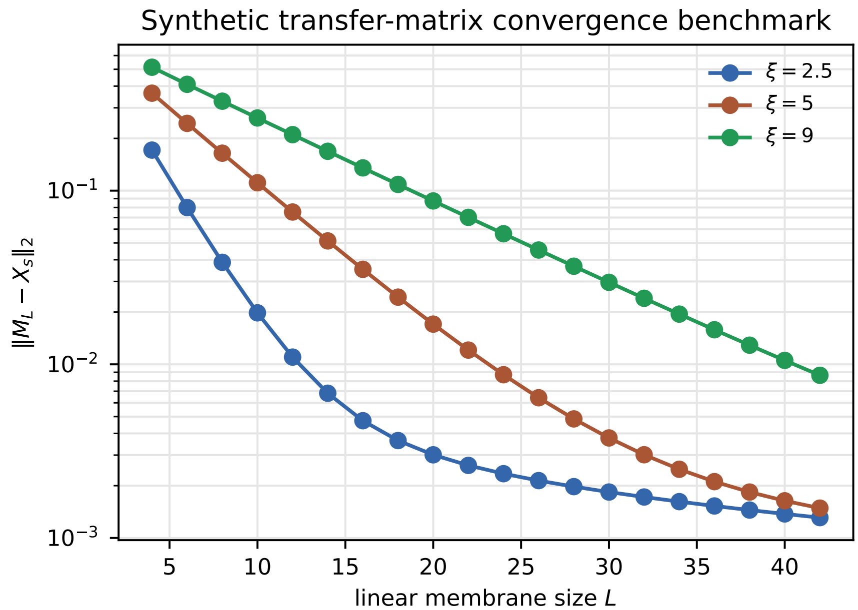

Figure 6: Synthetic transfer-matrix convergence benchmark. The curves illustrate the theorem's expected exponential approach to the asymptotic matrix, with slower convergence as the correlation length ξ increases. These are generated benchmark curves, not PEPS contraction data.

8.2 Finite-size behavior

Let $\lambda_0^{(s)}$ be the leading vacuum eigenvalue of the symmetry-twisted boundary transfer matrix and $\lambda_1^{(s)}$ the next eigenvalue in magnitude. After vacuum normalization, the leading correction behaves as

$$\|\widehat{\mathbf{M}}_L - X_s\| \leq C \left| \frac{\lambda_1^{(s)}}{\lambda_0^{(s)}} \right|^L + O(L^{-1}) = Ce^{-L/\xi_s} + O(L^{-1}).$$

At a phase transition the boundary gap can close, $\xi_s \to \infty$, and the finite-size convergence may cross over to a power law. This is not a failure of the definition; it is a warning that asymptotic SET data cannot be extracted from small membranes at criticality without scaling analysis.

8.3 A practical contraction plan

A reproducible tensor-network extraction should report the following objects.

1. The local PEPS tensor, its virtual MPO symmetries, and the pulling-through residual.

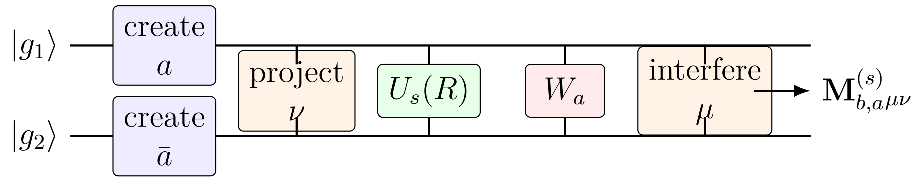

Figure 7: Minimal $d = 6$ qudit overlap circuit. The schematic abstracts the $\mathsf{D}(S_3)$ creation/projector operations into gates. The essential upgrade over scalar measurement is that the preparation and final interference labels ν, μ are varied coherently to recover the full fusion-space matrix.

2. The tube-idempotent basis used to prepare sectors a, b.

3. The transfer-matrix bond dimension and truncation error.

4. The vacuum normalization $\mathbf{M}^{(s)}_{\mathbf{1},\mathbf{1}}(R)$.

5. The complete matrix $\widehat{\mathbf{M}}^{(s)}_{b,a}(R)$, not only its trace.

6. Gauge checks: random PEPS and MPO gauges must change matrices covariantly but preserve closed-word invariants.

9 Minimal-Qudit Protocol

The $d = 6$ qudit architecture for $\mathsf{D}(S_3)$ suggests a compact experimental version of the matrix membrane. Use the group basis $\{|g\rangle : g \in S_3\}$ on each qudit. Ribbon creation, fusion, and measurement operations prepare a finite fusion basis. The protocol is:

1. prepare $|\Psi_{a,\nu}\rangle$ by creating an $a, \bar{a}$ pair and projecting onto a fusion-tree basis state ν;

2. apply the region symmetry $U_s(R)$ to the chosen patch;

3. apply or measure the a-ribbon projector $W_a(\partial R)$;

4. interfere with $|\Psi_{b,\mu}\rangle$ to obtain the complex matrix element;

5. repeat over μ, ν, b to assemble the block matrix.

Complex phases can be extracted by a Hadamard-test variant or by comparing interference fringes against a reference fusion state. If direct coherent interference is unavailable, one can reconstruct the matrix from probabilities in several rotated fusion bases. That second route is slower but uses only projective measurements.

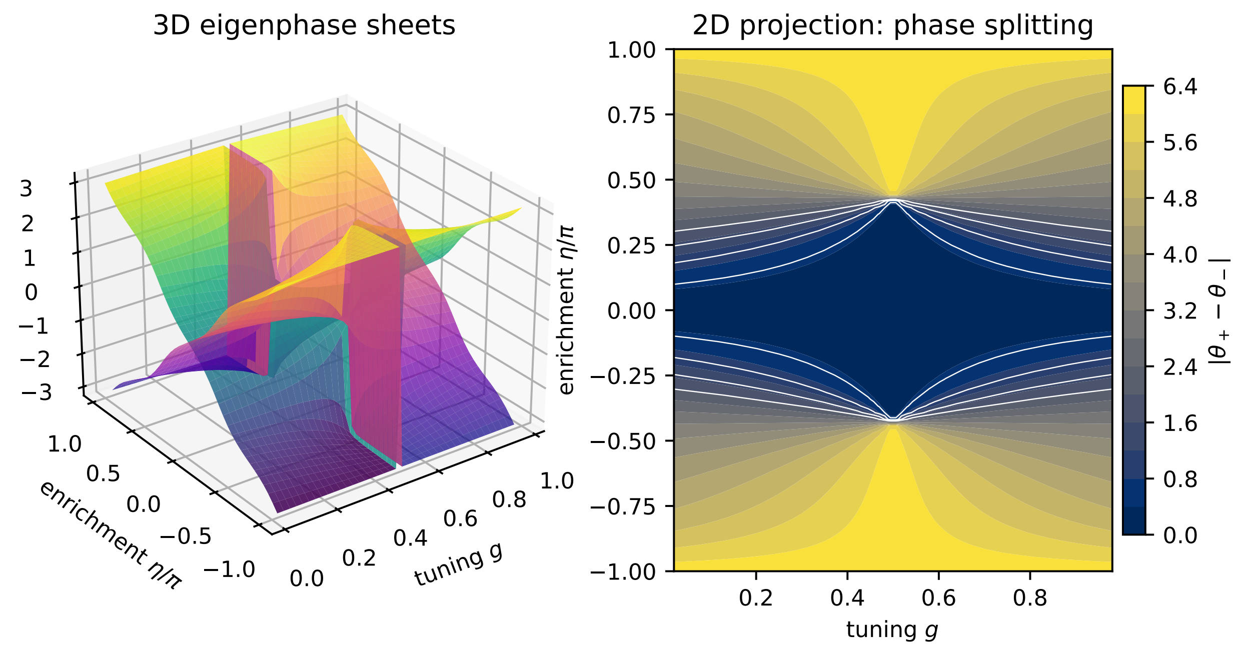

Figure 8: Synthetic eigenphase sheets for a two-dimensional fusion space over a tuning coordinate g and enrichment coordinate η. The right panel is the required 2D projection. The figure illustrates why a matrix observable is better conditioned than a scalar trace near degeneracies.

10 Numerical Results and Synthetic Benchmarks

This section reports two types of evidence. The first is exact and concerns the $\mathsf{D}(S_3)$ sector labelling in figure 5. The second is synthetic and concerns the numerical conditioning of the proposed reconstruction algorithm. The synthetic data are generated by `scripts/generate_figures.py`. They should be read as reproducible stress tests for the mathematics, not as physical simulation data.

10.1 Eigenphase sheets

The asymptotic matrix $X_a(s)$ is unitary after removing positive boundary dressing. Along a tunable enrichment path, its eigenphases are therefore natural coordinates. A scalar membrane sees only their averaged character. A two-level fusion space can have eigenphases $\pm\theta$ with zero trace over wide parameter ranges; the matrix still records the moving eigenbasis and the phase splitting.

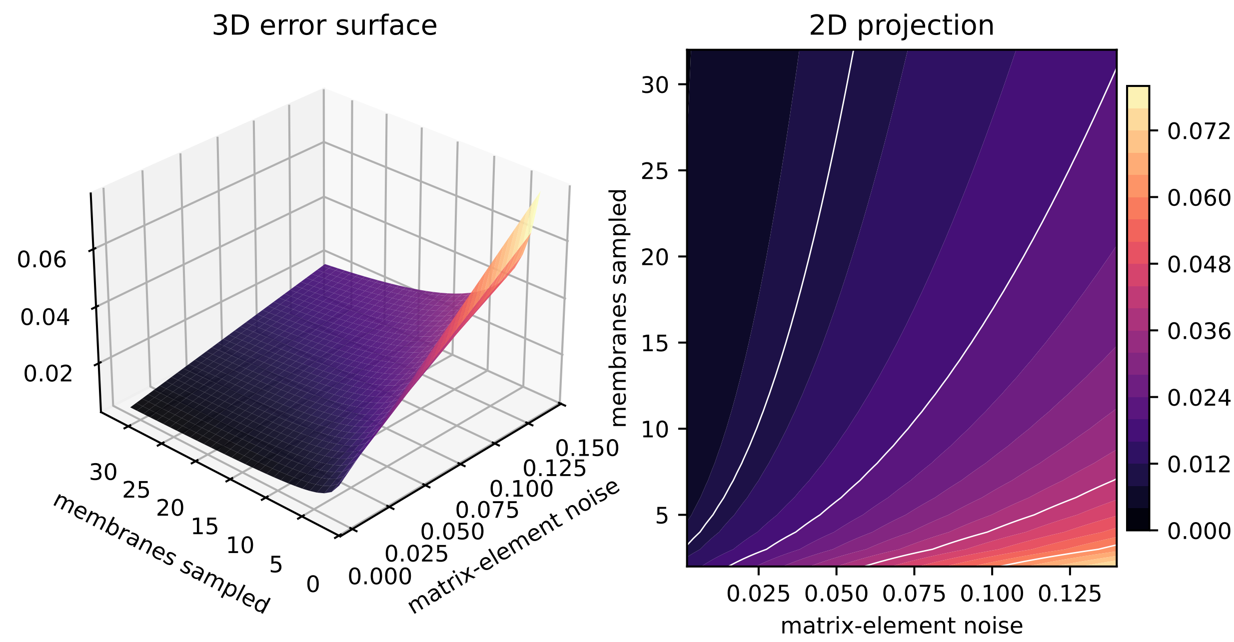

Figure 9: Synthetic reconstruction-error surface. The 3D surface and 2D projection show the expected decrease in error as more independent membrane/ribbon words are sampled. A real tensor-network or experimental study should replace this panel with measured errors and confidence intervals.

10.2 Reconstruction error

Given noisy estimates of matrix elements, the reconstruction problem is a constrained unitary synchronization problem over the fusion-sector graph. The error decreases with the number of independent membrane words and increases with complex matrix-element noise. The useful feature is not the particular synthetic surface in figure 9; it is the scaling protocol: report error versus both noise and number of sampled skein words, and validate on held-out words.

10.3 Scalar collapse

Figure 10 displays the failure mode this paper is designed to avoid. Two actions can be arranged to have the same scalar trace and even the same individual spectrum, while differing in a closed word involving an intrinsic braid or intersection matrix. The invariant is not an entry of M; it is the simultaneous conjugacy class of the family $(M, B, F, R, \ldots)$.

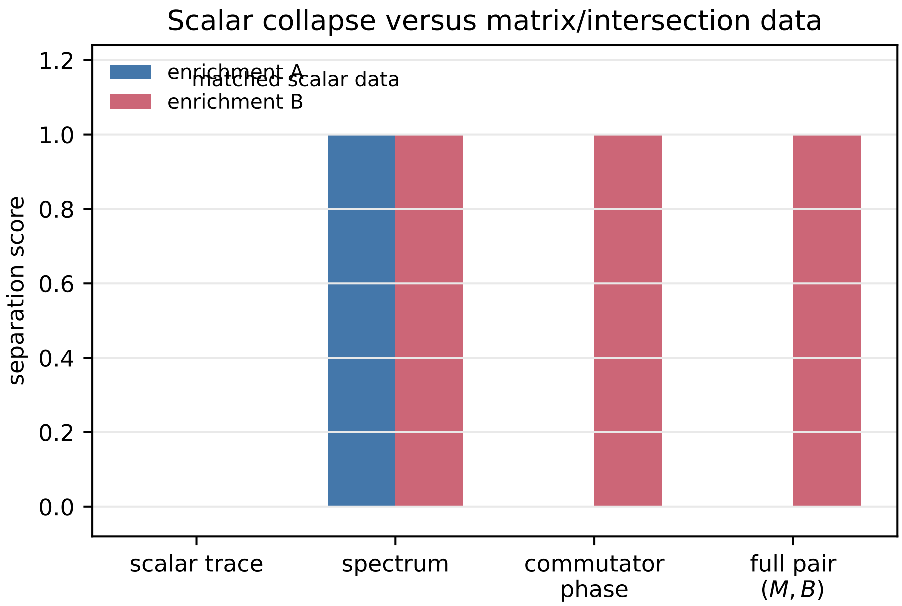

Figure 10: Synthetic comparison of scalar and matrix diagnostics. The first two columns are deliberately matched, while the commutator/intersection data separates the two enrichments. This is an algebraic benchmark of incompleteness, not a claim of a measured material system.

10.4 Noise robustness

Matrix tomography is more demanding than scalar measurement, but it also supplies redundancy. Fusion rules, group composition, antiunitary constraints, and intersection identities overdetermine the answer. Noise should therefore be reported after enforcing categorical constraints, not entry by entry.

11 Robustness and Controls

11.1 Required controls

A credible matrix-membrane study should include the following controls.

1. **Gauge-related tensors.** Apply random PEPS/MPO and fusion-basis gauges. Raw matrices must change; closed-word invariants must not.

2. **Scalar-matched phases.** Compare phases or algebraic models with identical scalar traces for the measured generator membranes.

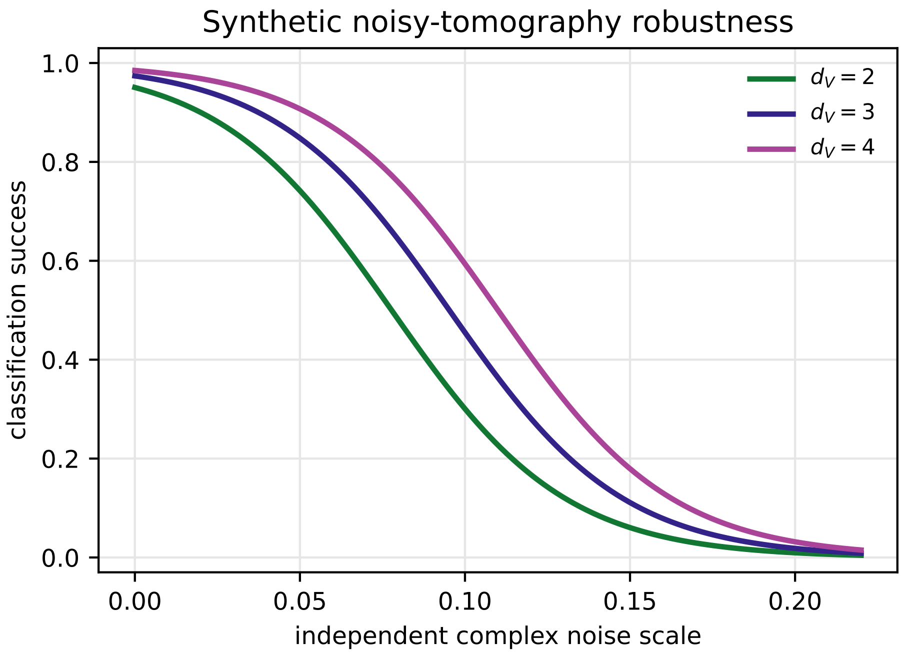

Figure 11: Synthetic noisy-tomography benchmark. Larger fusion spaces provide more constraints but also more parameters. The plotted success curves are illustrative outputs of the included script and should be replaced by physical noise models in platform-specific studies.

3. **Coboundary-related actions.** Verify that actions differing only by

$$X_a(s) \mapsto \beta^{-1}_{\rho_s(a)} X_a(s) \beta_a$$

are identified.

4. **Explicit symmetry breaking.** Break G locally and show that block selection and group-composition constraints fail.

5. **Abelian limits.** Reduce to $N^c_{ab} \leq 1$ and $\dim V_a = 1$, where the matrix collapses to the known scalar membrane phase.

6. **Randomized spectra.** Match individual spectra with random matrices and verify that intersection words reject the impostors.

7. **Identical intrinsic modular data.** Hold $\mathcal{C}$, S, T, and TEE fixed while changing enrichment data.

8. **Fusion-basis permutations.** Permute all multiplicity bases and check simultaneous covariance.

11.2 A precise scalar-incompleteness statement

Proposition 11.1 (Scalar traces are incomplete for matrix SET data)**.** *Let a measured subcategory contain a fusion space $V \cong \mathbb{C}^2$ and an intrinsic skein operator B acting on V. Suppose two symmetry actions X and Y satisfy*

$$\operatorname{Tr} X = \operatorname{Tr} Y, \qquad \operatorname{spec}(X) = \operatorname{spec}(Y),$$

but

$$\operatorname{Tr}(XBX^{-1}B^{-1}) \neq \operatorname{Tr}(YBY^{-1}B^{-1}).$$

Then no scalar membrane diagnostic depending only on $\operatorname{Tr} X$ or on the spectrum of the single membrane matrix can distinguish the two actions, while the matrix-valued membrane together with one ribbon-intersection word can.

Proof. The scalar data are equal by assumption. The closed commutator word is invariant under simultaneous basis change because

$$X \mapsto Q^{-1}XQ, \qquad B \mapsto Q^{-1}BQ$$

conjugates the entire product. Different traces of the closed word therefore imply inequivalence of the pairs (X, B) and (Y, B) as measured skein data. $\square$

Remark 11.2. Proposition 11.1 is intentionally formulated at the level of measured skein data. To promote any particular pair (X, B), (Y, B) to a pair of microscopic SET Hamiltonians, one must realize the corresponding G-crossed extension and verify the obstruction equations. The proposition identifies the mechanism by which scalar order fails; it does not fabricate a Hamiltonian realization.

11.3 Failure modes

The matrix membrane should not be advertised as complete in the following cases.

- If the large-R limit fails to converge because the boundary transfer matrix is gapless at all accessible sizes, the extracted matrix is not an asymptotic SET invariant.

- If two inequivalent G-crossed extensions agree on all measured closed skein words, the measured family is incomplete.
- If the observable changes under a gauge transformation in a way that alters claimed invariants, the implementation is measuring basis choices.
- If symmetry is explicitly or spontaneously broken, the region operator no longer inserts a topological defect sheet, and the categorical interpretation fails.

12 Discussion

12.1 Relation to tube algebras and entanglement

The tube algebra already knows how to resolve topological sectors in an MPO-injective PEPS [7]. The present construction adds a symmetry-domain-wall operator to the tube algebra and refuses to trace over the resulting multiplicity space. In this sense the membrane matrix is a symmetry-twisted tube-algebra matrix element. It is complementary to topological entanglement entropy [14, 15]: TEE measures the total quantum dimension of the intrinsic topological order, while the membrane matrix measures how G acts on the sectors contributing to that order.

12.2 How small can a patch be?

At a fixed point, a patch only needs to be large enough to support the ribbon projector and keep local anyon insertions separated from ∂R. In a finite-correlation-length phase, the useful size is set by the larger of the bulk correlation length, the ribbon dressing length, and the boundary transfer-matrix correlation length. A practical rule is:

$$L \gtrsim 5\xi_s$$

for percent-level convergence when the leading correction is exponential. Near criticality this rule fails and a scaling collapse in L/ξ_s is required. Minimal qudit protocols can test fixed-point algebra on very small patches, but phase diagnostics away from fixed points need enough room for boundary dressing to decay.

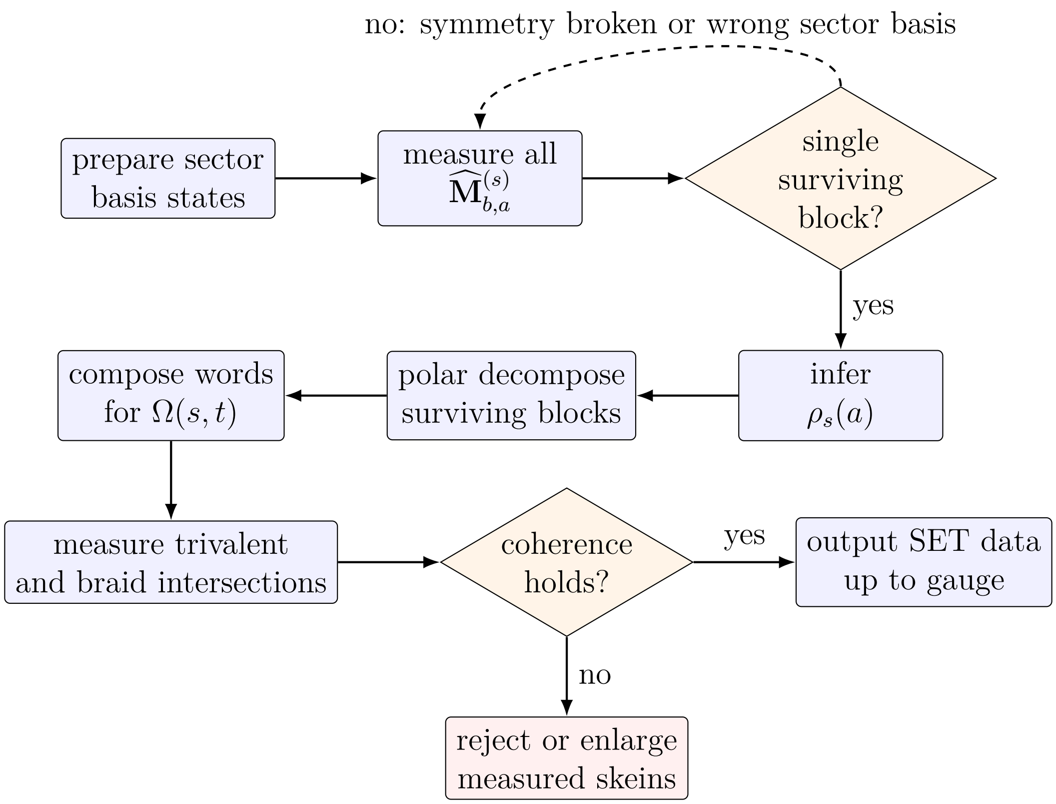

Figure 12: Categorical tomography decision diagram. The algorithm first reconstructs the anyon permutation, then the unitary or antiunitary sector maps, then the projective factors, and finally the fusion/braid coherence data. Failed coherence is a result, not a nuisance: it means the measured family or physical assumptions are insufficient.

12.3 Why intersections matter

A single matrix for each generator may recover a projective representation, but SET order is not just a representation. It is a representation compatible with fusion, braiding, and defects. Intersections with trivalent skeins and braids are therefore not optional decoration. They are the operational way to measure the coherence constraints that distinguish categorical equivalence classes.

13 Conclusions

The scalar membrane order parameter of an Abelian SET generalizes naturally to a matrix-valued symmetry-twisted ribbon observable. The answer to the central question is conditional but sharp: under gap, injectivity, and stable-ribbon hypotheses, the asymptotic vacuum-normalized membrane matrix encodes the projective or G-crossed symmetry action on the anyon fusion space, up to categorical gauge transformations.

The seven specific conclusions are as follows.

1. **Encoded SET data.** The matrix encodes the anyon permutation $\rho_s(a)$, the sector map $X_a(s)$, projective composition factors $\Omega_a(s,t)$, and, after intersections are included, the measured part of the G-crossed associativity and half-braiding data.

2. **Gauge-invariant quantities.** Raw entries are not invariant. Conjugacy classes for fixed sectors, singular values for source-target intertwiners, determinants, spectra, and traces of closed membrane-ribbon skein words are invariant.

3. **Uniqueness of reconstruction.** Reconstruction is unique up to coboundary and categorical equivalence when the measured skein family generates the relevant fusion and braid data. Otherwise it is a controlled partial tomography.

4. **Scalar failure.** Scalar membrane order fails whenever the trace or single-generator spectrum forgets noncommuting fusion-space action, anyon permutation, Kramers composition, or multiplicity mixing. Proposition 11.1 is the precise algebraic failure mode.

5. **Patch size.** Fixed-point $d = 6$ qudit tests can be very small. Away from fixed points, L must exceed the boundary correlation length by a comfortable factor, and near criticality a finite-size scaling analysis is mandatory.

6. **Stability.** The diagnostic is stable under noise when categorical constraints are enforced globally, but critical boundary spectra and explicit symmetry breaking can invalidate the asymptotic interpretation.

7. **Limit theorem.** Theorem 4.1 defines the positive limit: matrix membranes recover $X_a(s)$. The counter-limit is equally important: if two inequivalent extensions agree on all measured closed skein words, no observable restricted to that family can separate them.

The practical research program is now clear. Implement exact ribbon projectors for $\mathsf{D}(S_3)$, choose a verified F/R convention, contract symmetry-twisted MPO-injective PEPS membranes along a tunable enrichment path, and replace the synthetic panels in this paper by measured finite-size, noise, and tomography data. The mathematical structure predicts exactly what should be gauge-covariant, what should be gauge-invariant, and where scalar membrane order must lose information.

A Proof Details for the Convergence Theorem

This appendix expands the transfer-matrix argument used in theorem 4.1. Let R_L be a disk-like region whose boundary length is proportional to L. The double-layer contraction of a PEPS expectation value with insertions $p_b O_s p_a$ along the virtual boundary defines a transfer operator $\mathbb{T}_{b,a}^{(s)}$. MPO-injectivity implies that the topological fixed space decomposes as

$$\mathcal{F}_{\partial R} = \bigoplus_a \mathcal{F}_a,$$

where p_a projects onto $\mathcal{F}_a$. The graded pulling-through relation has the form

$$O_s p_a = p_{\rho_s(a)} O_s + \epsilon_L,$$

where $\|\epsilon_L\| \le Ce^{-L/\xi_{\text{pull}}}$ for a dressed operator away from the fixed point.

Restrict to the surviving block $b = \rho_s(a)$. On the leading fixed space,

$$p_{\rho_s(a)} O_s p_a = \lambda_0^{(s)} X_a(s) \oplus N,$$

where N maps into subleading boundary modes. If $|\lambda_1^{(s)}| < |\lambda_0^{(s)}|$, then after contracting a boundary of length L,

$$\frac{\mathbb{T}_L[p_{\rho_s(a)} O_s p_a]}{(\lambda_0^{(s)})^L} = X_a(s) + O\left(\left|\frac{\lambda_1^{(s)}}{\lambda_0^{(s)}}\right|^L\right) + O(e^{-L/\xi_{\text{pull}}}).$$

The denominator $\mathbf{M}_{\mathbf{1},\mathbf{1}}^{(s)}(R_L)$ is precisely the vacuum estimate of $(\lambda_0^{(s)})^L$, up to the same exponentially small corrections. Combining the two estimates gives (4.2) with

$$\xi^{-1} = \min\{\xi_{\text{pull}}^{-1}, -\log|\lambda_1/\lambda_0|\}.$$

B Gauge Covariance with Antiunitary Symmetry

For antiunitary T, the sector map is semilinear:

$$X_a(T)(\alpha v + \beta w) = \overline{\alpha} X_a(T) v + \overline{\beta} X_a(T) w.$$

Choosing a basis identifies it with a matrix followed by complex conjugation. If the source basis changes by Q_a and the target basis by $Q_{\rho_T(a)}$, the matrix representative changes as

$$X_a(T) \longmapsto Q_{\rho_T(a)}^{-1} X_a(T) \overline{Q_a}.$$

The square

$$\Theta_a^2 = X_{\rho_T(a)}(T) \overline{X_a(T)}$$

therefore transforms by ordinary conjugation when $\rho_T^2(a) = a$:

$$\Theta_a^2 \longmapsto Q_a^{-1} \Theta_a^2 Q_a.$$

Its eigenvalues are Kramers invariants. In a one-dimensional sector this reduces to the familiar sign $T^2 = \pm 1$, while in a multiplicity space it can be a non-scalar unitary protected by fusion constraints.

C $\mathsf{D}(S_3)$ Projector Conventions

Write

$$S_3 = \langle r, t \mid r^3 = t^2 = e,\ trt = r^{-1} \rangle.$$

The conjugacy classes are $[e]$, $[t] = \{t, tr, tr^2\}$, and $[r] = \{r, r^2\}$. The centralizers are

$$Z_e = S_3, \qquad Z_t \cong \mathbb{Z}_2, \qquad Z_r \cong \mathbb{Z}_3.$$

The simple objects of $\mathcal{Z}(\mathrm{Vec}_{S_3})$ are (C, π), and their quantum dimensions are $|C| \dim \pi$. The topological spin is

$$\theta_{(C,\pi)} = \frac{\chi_\pi(g)}{\dim \pi}$$

for $g \in C$. This gives the table in section 7.

The idempotent formula (7.1) follows from ordinary character orthogonality in the centralizer. Since this paper does not choose an explicit fusion-tree gauge for all $\mathsf{D}(S_3)$ F-moves, it avoids printing numerical F/R symbols. Any implementation should state its convention, verify the pentagon and hexagon equations, and then check that membrane invariants are unchanged under randomly sampled gauge transformations.

D Algorithmic Matrix-Membrane Tomography

This appendix states the reconstruction procedure in a form suitable for a methods section. The input consists of four ingredients: the list of candidate anyon sectors, a generating set of symmetry elements, the measured vacuum-normalized blocks $\widehat{\mathbf{M}}^{(s)}_{b,a}$, and the measured trivalent and braided intersection matrices. A tolerance ε is fixed from the vacuum-sector noise floor and from repeated measurements of gauge-equivalent representatives.

First, for each pair (s,a), compute the Frobenius norms

$$n^{(s)}_{b,a} = \|\widehat{\mathbf{M}}^{(s)}_{b,a}\|_{\mathrm{F}}.$$

If a unique b_* satisfies

$$n^{(s)}_{b_*,a} > \varepsilon, \qquad n^{(s)}_{b,a} < \varepsilon \quad (b \neq b_*),$$

set $\rho_s(a) = b_*$. If there is no such block, the experiment has either used an incomplete sector basis, measured too small a membrane, or entered a regime in which symmetry is broken or the topological-sector description is not valid. This failure should be reported rather than repaired by hand.

Second, polar-decompose the surviving block:

$$\widehat{\mathbf{M}}^{(s)}_{\rho_s(a),a} = P_a(s)U_a(s),$$

where $P_a(s) \geq 0$ and $U_a(s)$ is unitary on the support. The positive part contains residual boundary dressing and imperfect normalization. The unitary part is the estimator for $X_a(s)$. A useful diagnostic is the condition number of $P_a(s)$; if it is large, the membrane size or basis preparation is not yet resolving the topological fixed space cleanly.

Third, form composition residuals. For each relation $st = u$ in the group, define

$$\mathcal{R}_a(s,t) = U_{\rho_t(a)}(s)U_a(t)U_a(u)^{-1}.$$

In a one-dimensional sector this is the measured projective phase. In a multiplicity space it is a unitary natural automorphism. To quotient coboundary freedom, minimize the simultaneous residual under

$$U_a(s) \mapsto Q_{\rho_s(a)}^{-1}U_a(s)Q_a$$

over unitary Q_a's. This is a compact synchronization problem on the sector graph. It should be solved globally rather than by fixing phases one sector at a time, because fusion constraints couple the choices.

Fourth, use intersections as validation data. The trivalent residual is the norm of the difference between the two skein evaluations obtained by either moving the symmetry sheet through a fusion vertex or applying the sector maps to each leg first. The braided residual is the analogous difference for a crossing. These residuals test categorical coherence. A set of matrices that matches all individual spectra but fails these residuals is not an SET reconstruction.

Finally, split the measured skein words into calibration and hold-out families. The calibration family fixes ρ, X, and Ω. The hold-out family contains longer group words, alternative fusion trees, and at least one braid or commutator word not used in the fit. Agreement on held-out words is the practical criterion that distinguishes categorical tomography from overfitting a small set of matrix elements.

E Closed Skein Words and Invariant Algebras

The raw membrane matrix is not an observable in the invariant-theory sense; it is a coordinate representation of a morphism. The invariant algebra is generated by closed words. A closed word is a cyclic product of membrane matrices, fusion maps, braid maps, and their adjoints whose source and target sectors match. In a fixed sector a, examples include

$$\operatorname{Tr} X_a(s), \qquad \operatorname{Tr}\left(X_{\rho_t(a)}(s)X_a(t)X_a(st)^{-1}\right),$$

and

$$\operatorname{Tr}\left(X_c(s)F^c_{ab}(X_a(s)^{-1}\otimes X_b(s)^{-1})(F^c_{ab})^{-1}\right).$$

For a braid, one obtains

$$\operatorname{Tr}\left(R_{\rho_s(a)\rho_s(b)}(X_a(s)\otimes X_b(s))R^{-1}_{ab}X_{a\otimes b}(s)^{-1}\right),$$

where $X_{a\otimes b}(s)$ denotes the induced action on the fused space. These formulas are schematic because the placement of multiplicity indices depends on the chosen fusion tree, but the closure principle is basis-independent.

The reason closed words are invariant is cancellation. If every open morphism is transformed by source-target basis changes, adjacent transformations cancel when the word is closed. Cyclicity of trace removes the final conjugation. This is the categorical analogue of Wilson-loop gauge invariance in lattice gauge theory. The membrane matrix plays the role of a parallel transporter, and closed skein words play the role of holonomies around loops in the group-fusion-braid groupoid.

This viewpoint also clarifies the difference between spectrum and simultaneous conjugacy. The spectrum of a single matrix is invariant but incomplete. Two pairs (X,B) and (Y,B') can have the same spectra for X and B separately but different spectra for XB, different commutator traces, or different braid-intersection words. A scalar membrane usually samples only the first of these invariants. The matrix-valued skein family samples the full invariant algebra generated by the measured topological operators.

There is a useful hierarchy.

data retained	invariant meaning
single scalar entry	basis-dependent unless a canonical state is physically specified
trace of one membrane	character-like invariant for a fixed sector
spectrum of one membrane	conjugacy invariant of one symmetry action
all generator spectra	still incomplete for noncommuting actions
closed group words	projective class and anyon permutation constraints
closed fusion/braid words	compatibility with F-moves, R-moves, and G-crossed half-braidings

This hierarchy is a useful experimental guide. If a platform can only measure one scalar expectation, it can at best reproduce Abelian-style fractionalization indicators. If it can measure matrix elements in several fusion bases, it can recover conjugacy classes. If it can also implement ribbon intersections, it can test the categorical coherence relations that define the SET enrichment.

F Finite Fusion-Space Separation Models

This appendix gives algebraic separation models that can be used as unit tests for software and experiments. They are not claimed to be microscopic Hamiltonian realizations unless a compatible G-crossed category and obstruction-free defect theory are supplied.

Let $V = \mathbb{C}^2$ and let $B = Z = \mathrm{diag}(1, -1)$ be an intrinsic skein operator on V, such as a half-braiding or an F-move block in a chosen multiplicity space. Compare two candidate symmetry actions

$$X_A = X, \qquad X_B = Z,$$

where X and Z are Pauli matrices. Both have trace zero and spectrum $\{+1, -1\}$. A scalar membrane trace and a single-matrix spectrum therefore identify them. However,

$$X_A B X_A^{-1} B^{-1} = -\mathbf{1}, \qquad X_B B X_B^{-1} B^{-1} = +\mathbf{1}.$$

The closed commutator trace separates them. Since simultaneous conjugation sends both X

and B to $Q^{-1}XQ$ and $Q^{-1}BQ$, the sign of this commutator is a closed-word invariant of the measured pair.

A second model uses $G = \mathbb{Z}_2 \times \mathbb{Z}_2$. Let x, z be the two generators. The projective relation

$$U_x U_z = \chi\, U_z U_x, \qquad \chi = \pm 1,$$

is invisible to the two scalar traces $\operatorname{Tr} U_x$ and $\operatorname{Tr} U_z$ when U_x, U_z are traceless. It is visible to the group commutator

$$\frac{1}{2}\operatorname{Tr}(U_x U_z U_x^{-1} U_z^{-1}) = \chi.$$

In an SET realization, χ is not merely a representation-theoretic phase; it must be compatible with fusion. If the same χ appears naturally on every fusion channel according to the fractionalization constraints, it defines a cohomology class. If it appears inconsistently across channels, the measured matrices fail the coherence tests and should be rejected.

These examples explain the phrase "scalar membrane order fails." If the scalar diagnostic is expanded to include every group element, every product word, and every intersected ribbon word, it is no longer the original scalar membrane; it is computing traces of the matrix-valued skein algebra. The proposal of this paper is to measure the matrix first, then form whichever invariant words are physically meaningful and statistically stable.

G Tensor-Network Estimator and Error Budget

A matrix-membrane PEPS calculation has three independent sources of error: finite membrane size, boundary contraction truncation, and gauge/basis conditioning. They should be reported separately.

The finite-size error is controlled by the twisted boundary transfer matrix. If the leading normalized subdominant eigenvalue is r_s, then

$$\epsilon_{\text{size}}(L) \approx C_s r_s^L.$$

Near a critical point r_s approaches one. A finite data set should therefore fit both exponential and power-law forms and should not quote an asymptotic matrix unless the preferred fit is stable under removing the smallest sizes.

The contraction truncation error is the discarded weight or variational residual of the boundary MPS, CTMRG environment, or other contraction method. Because the observable is a matrix, the truncation should be tested entrywise and invariantly. Entrywise convergence can look good in one gauge and poor in another. Closed-word convergence is the more physical diagnostic.

The conditioning error enters when extracting the unitary part of a nearly unitary block. If

$$\widehat{\mathbf{M}} = PU$$

is the polar decomposition, then the perturbation of U is controlled by the smallest singular value of $\widehat{\mathbf{M}}$. Blocks connecting permuted anyons are particularly sensitive if the basis preparation has small overlap with one of the target channels. A robust study should therefore report singular values before quoting phases.

Combining these effects gives an error budget of the form

$$\epsilon_{\text{total}} \leq \epsilon_{\text{size}} + \epsilon_{\text{trunc}} + \kappa(\widehat{\mathbf{M}})\epsilon_{\text{meas}} + \epsilon_{\text{gauge}},$$

where κ is the relevant polar-decomposition condition number and ϵ_{gauge} is the residual variation of claimed invariants under explicit gauge randomization. The last term should vanish within numerical precision. If it does not, the implementation is not measuring a categorical invariant.

For the synthetic plots in the main text, only the first and third effects are modeled. A real PEPS study must replace them with measured transfer spectra, actual contraction residuals, and bootstrapped uncertainty intervals. The advantage of the matrix formalism is that the required diagnostics are clear: block support tests ρ, polar unitarity tests boundary normalization, composition residuals test projective consistency, and intersection residuals test categorical coherence.

H Exact Small-System $D(S_3)$ Evaluation Recipe

A finite $D(S_3)$ calculation can be organized entirely in the group basis. The local Hilbert space is spanned by $|g\rangle$, $g \in S_3$. Vertex operators impose gauge invariance, plaquette operators measure flux, and ribbon operators create pairs of quantum-double excitations.

The matrix membrane adds a region symmetry to this standard algebra.

The first step is to construct the ribbon algebra basis. Choose a ribbon path γ and label elementary ribbon operators by flux and charge data. Close the ribbon around ∂R, then project to a simple sector using the centralizer character projector (7.1). This yields a numerical matrix for $W_a(\partial R)$ on the finite Hilbert space.

The second step is to prepare fusion basis states. Create an $a, \bar{a}$ pair, move the endpoints to the desired punctures, and apply local projectors selecting the fusion channel. In a multiplicity-free channel this gives one vector. In a channel with multiplicity, repeat the construction with orthogonal trivalent coupons. Orthonormalize the resulting vectors and record the basis change; this record is part of the gauge data.

The third step is to represent $U_s(R)$. If s is a group automorphism of the microscopic qudit labels, it acts by permuting $|g\rangle$ inside R. If s is represented by a defect wall rather than an onsite automorphism, insert the corresponding wall operator and keep track of the endpoint defect labels. In either case, compute the overlap matrix

$$\langle \Psi_{b,\mu} | U_s(R) W_a(\partial R) | \Psi_{a,\nu} \rangle.$$

The fourth step is validation. Apply a random unitary change of fusion basis in every multiplicity space. Recompute the raw matrices. They must transform by source-target similarity, while the singular values and closed-word traces remain fixed. Then apply a coboundary change to the symmetry representatives and verify that the reconstructed cohomology class is unchanged. Finally, compare against an Abelian sector such as A or B, where the matrix reduces to a scalar and the output must agree with the ordinary fractionalization phase.

This recipe is intentionally explicit about what is missing: a verified convention for all F/R-moves must be supplied by the implementation. Once supplied, the same recipe produces the trivalent and braided intersection matrices needed for full G-crossed tomography.

I Control Table and Rejection Criteria

The controls in section 11 can be condensed into a table that is useful for laboratory notebooks and numerical repositories.

control	expected behavior	reject if
fusion-basis randomization	raw entries change covariantly; closed words fixed	invariant traces drift beyond error bars
PEPS tensor gauge	virtual representatives change; physical matrices equivalent	spectra or singular values change after normalization
coboundary shift	representative Ω changes; cohomology class fixed	classification treats coboundaries as distinct phases
Abelian limit	matrix reduces to scalar fractionalization phase	off-diagonal structure remains in one-dimensional sectors
symmetry breaking	block selection and group laws fail	algorithm silently forces a G-crossed interpretation
matched random spectra	individual spectra agree but intersection words fail	reconstruction accepts random simultaneous data
critical tuning	convergence slows or crosses to scaling form	asymptotic matrix is quoted without size scaling
held-out skein words	predicted closed words agree with measurements	fit succeeds only on calibration words

The most important rejection criterion is gauge dependence. A diagnostic that changes its claimed invariant output under a legal fusion-basis or tensor gauge transformation is not a physical SET diagnostic. The second most important is overfitting: matching a handful of matrix entries or eigenvalues is insufficient. The categorical constraints are global, and the validation must be global as well.

J Reconstruction Proof and Obstruction Caveats

This appendix gives a more detailed proof of theorem 6.1. The proof is best understood as a reconstruction of a functor from a measured skein groupoid. Let $\mathcal{G}_{\mathrm{meas}}$ be the groupoid whose objects are the measured fusion spaces and whose morphisms are generated by symmetry membranes, trivalent fusion vertices, braids, duality cups and caps, and their adjoints, subject only to the relations directly tested in the experiment or contraction. A complete categorical measurement is a unitary representation

$$\Phi : \mathcal{G}_{\mathrm{meas}} \longrightarrow \mathrm{Hilb}_{\mathrm{fd}}$$

that assigns to each generator its measured matrix and to each closed word its measured scalar trace.

The block-support part of Φ determines a partial action on simple labels. For each s and a, the unique nonzero block identifies $\rho_s(a)$. If two nonzero blocks survive in the large-region limit, then the assumed simple superselection basis is wrong or the symmetry is not acting as a topological defect. Thus a well-defined ρ is not an input to the reconstruction; it is a falsifiable output.

Once ρ is known, the matrices $X_a(s)$ define a lift of ρ to fusion spaces. The obstruction to this lift being an ordinary group action is the family

$$\Omega_a(s,t) = X_{\rho_t(a)}(s) X_a(t) X_a(st)^{-1}.$$

Naturality requires that for every fusion vertex $v \in V_{ab}^c$,

$$X_c(s)\, v = \rho_s(v)\, (X_a(s) \otimes X_b(s))$$

after inserting the appropriate G-crossed associator convention. The intersection matrices measure exactly the residual of this equality. If the residual vanishes on a generating set of vertices, semisimplicity extends it to all morphisms generated by those vertices.

Gauge transformations enter as pseudonatural equivalences of the reconstructed functor. A choice of unitary Q_a on every object sends

$$X_a(s) \mapsto Q_{\rho_s(a)}^{-1} X_a(s) Q_a$$

and sends Ω to a cohomologous factor system. Hence Ω is not itself an invariant; its class modulo these coboundaries is. In ordinary one-anyon Abelian sectors this reduces to the familiar second-cohomology fractionalization class. In non-Abelian multiplicity spaces it is a class of natural automorphisms constrained by fusion.

Braiding supplies the next layer. For each measured crossing, the half-braiding residual compares two procedures: braid first and then move the symmetry sheet, or move the sheet first and braid the transformed anyons. Vanishing residuals identify the measured part of the G-crossed braiding. If every simple object, fusion vertex, and braid generator of $\mathcal{C}$ is measured, the pentagon and hexagon relations imply that the reconstructed data extend uniquely to all skeins in $\mathcal{C}$, modulo the known categorical gauge. If only a subcategory is measured, the theorem reconstructs only the restriction to that subcategory.

There are two obstruction caveats. First, not every apparent anyon permutation and factor system extends to a consistent G-crossed braided category. The usual obstruction theory can produce a nonzero obstruction at higher degree, preventing the existence of a defect extension even though the one-anyon matrices look sensible. Operationally, this appears as a failure of associator or braid residuals on longer held-out words. Second, a microscopic symmetry may act antiunitarily or may reverse orientation. Then the target of ρ_s is a conjugated or mirror braided category, and the reconstruction must include the semilinear transformations described in appendix B. The closed-word principle remains valid, but the word must record where complex conjugation is applied.

The proof of completeness is therefore conditional in exactly the way an experimental diagnostic should be conditional. If the measured generators span the category and the residuals vanish, the data define the SET action up to equivalence. If the generators do not span the category, the result is partial. If residuals do not vanish, the attempted reconstruction has discovered either insufficient data, finite-size contamination, symmetry breaking, or an obstructed candidate enrichment.

K Replacing the Synthetic Panels

The generated numerical panels in this manuscript are placeholders for a precise future workflow. To replace the eigenphase surface, choose a tunable PEPS family $A(g, \eta)$, build the symmetry-twisted boundary transfer matrix for each grid point, normalize by the vacuum sector, and extract the unitary part of each surviving $\widehat{\mathbf{M}}$ block. The plotted surface should then show measured eigenphases of those unitary parts, with markers indicating points where the polar decomposition is ill-conditioned.

To replace the finite-size convergence curves, compute $\widehat{\mathbf{M}}_L$ for several linear sizes L at fixed (g, η). Estimate the asymptotic matrix either from the largest available size or from a joint fit to

$$\widehat{\mathbf{M}}_L = X + C_1 e^{-L/\xi} + C_2/L.$$

Report the norm of the residual, the fitted ξ, and whether removing the smallest size changes X beyond error bars. If a power-law fit is statistically indistinguishable from an exponential fit, do not quote a unique asymptotic SET matrix.

To replace the reconstruction-error surface, inject calibrated noise into either physical measurement outcomes or tensor-network contractions, reconstruct ρ, X, and Ω, and compare against exact fixed-point data or a high-bond-dimension reference. The horizontal axis should be an operational noise parameter: shot count, depolarizing probability, boundary truncation threshold, or matrix-element standard deviation. The vertical axis should be the number of independent closed skein words used in calibration. Held-out words must be reported separately.

To replace the scalar-versus-matrix panel, one needs a pair of enrichments or controlled algebraic realizations with matched scalar data. The cleanest test is to match $\mathrm{Tr}\, X_a(s)$ for the generator membranes while changing a commutator or intersection invariant. A stronger test matches all individual spectra and still separates the phases through simultaneous closed-word data. That stronger test rules out the common false explanation that the matrix diagnostic merely measured eigenvalues more accurately.

To replace the noise-robustness panel, report classification success after categorical projection. Raw matrix-element fidelity is not the relevant final metric. The relevant metric is whether the reconstructed equivalence class is correct, whether coboundary-related inputs are identified, and whether gauge-related tensors produce invariant closed words. This makes the numerical claim reproducible and falsifiable.

L Reproducibility of Generated Figures

The Matplotlib figures in this manuscript were generated by the included figure-generation script. Their status is:

file	status
`d_s3_sector_map.pdf`	exact sector labels, dimensions, and spin phases from $\mathcal{Z}(\mathrm{Vec}_{S_3})$
`matrix_structure.pdf`	schematic block-matrix visualization
`eigenphase_surface.pdf`	synthetic two-level transfer-matrix benchmark
`finite_size_convergence.pdf`	synthetic exponential convergence benchmark
`reconstruction_error_surface.pdf`	synthetic noisy reconstruction benchmark
`scalar_vs_matrix.pdf`	synthetic scalar-incompleteness benchmark
`noise_robustness.pdf`	synthetic noisy classification benchmark

The TikZ figures are schematic categorical diagrams. They are meant to convey operator geometry, gauge covariance, intersection tensors, a qudit overlap protocol, and the tomography workflow. They are not numerical data.

References

[1] Maissam Barkeshli, Parsa Bonderson, Meng Cheng, and Zhenghan Wang. Symmetry fractionalization, defects, and gauging of topological phases. *Physical Review B*, 100: 115147, 2019. doi: 10.1103/PhysRevB.100.115147.

[2] Pavel Etingof, Shlomo Gelaki, Dmitri Nikshych, and Viktor Ostrik. *Tensor Categories*, volume 205 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2015.

[3] Vladimir Turaev and Alexis Virelizier. *Monoidal Categories and Topological Field Theory*, volume 322 of *Progress in Mathematics*. Birkhäuser, Cham, 2017. doi: 10.1007/978-3-319-49834-8.

[4] Lukas Haller, Wen-Tao Xu, Yu-Jie Liu, and Frank Pollmann. Quantum phase transition between symmetry enriched topological phases in tensor-network states. *Physical Review Research*, 5:043078, 2023. doi: 10.1103/PhysRevResearch.5.043078.

[5] Norbert Schuch, Ignacio Cirac, and David Pérez-García. PEPS as ground states: degeneracy and topology. *Annals of Physics*, 325(10):2153–2192, 2010. doi: 10.1016/j.aop.2010.05.008.

[6] Mehmet B. Şahinoğlu, Dominic J. Williamson, Nick Bultinck, Michael Mariën, Jutho Haegeman, Norbert Schuch, and Frank Verstraete. Characterizing topological order with matrix product operators. *Annales Henri Poincaré*, 22:563–592, 2021. doi: 10.1007/s00023-020-00992-4.

[7] Nick Bultinck, Michael Mariën, Dominic J. Williamson, Mehmet B. Şahinoğlu, Jutho Haegeman, and Frank Verstraete. Anyons and matrix product operator algebras. *Annals of Physics*, 378:183–233, 2017. doi: 10.1016/j.aop.2017.01.004.

[8] Vladimir G. Turaev. *Quantum Invariants of Knots and 3-Manifolds*, volume 18 of *De Gruyter Studies in Mathematics*. Walter de Gruyter, Berlin, 1994.

[9] Józef H. Przytycki. Skein modules of 3-manifolds. *Bulletin of the Polish Academy of Sciences. Mathematics*, 39(1–2):91–100, 1991.

[10] Alexei Kitaev. Anyons in an exactly solved model and beyond. *Annals of Physics*, 321(1):2–111, 2006. doi: 10.1016/j.aop.2005.10.005.

[11] Michael A. Levin and Xiao-Gang Wen. String-net condensation: a physical mechanism for topological phases. *Physical Review B*, 71:045110, 2005. doi: 10.1103/PhysRevB.71.045110.

[12] Lucy Byles, Ewan Forbes, and Jiannis K. Pachos. Demonstration of magic state power of $D(S_3)$ anyons with two qudits. *New Journal of Physics*, 28:044501, 2026. doi: 10.1088/1367-2630/ae4aca.

[13] Dominic J. Williamson, Nick Bultinck, and Frank Verstraete. Symmetry-enriched topological order in tensor networks: defects, gauging and anyon condensation. *arXiv preprint*, 2017.

[14] Alexei Kitaev and John Preskill. Topological entanglement entropy. *Physical Review Letters*, 96:110404, 2006. doi: 10.1103/PhysRevLett.96.110404.

[15] Michael Levin and Xiao-Gang Wen. Detecting topological order in a ground state wave function. *Physical Review Letters*, 96:110405, 2006. doi: 10.1103/PhysRevLett.96.110405.