Mathematical Structures from a Distinguishability Axiom

Symplectic and Unitary Geometry as Consequences of Relational Topology

Abstract

We investigate the mathematical structures that emerge from a single axiom in modal logic: $\Diamond N \to \neg N$ ("if nothingness is possible, then nothingness does not obtain"). The axiom requires a conserved measure of distinguishability Ω and, through categorical exhaustion, a five-dimensional constraint structure representable as the Clifford algebra Cl(5). We show that the character of the resulting geometry depends entirely on relational topology — the pattern of connections among distinguishable elements.

At N = 2 (two elements, one pairwise relation, no closed loops), the single coupling matrix is always diagonalisable, decomposing the configuration space into independent two-dimensional sectors. A classical theorem of Arnold — that measure-preserving flow in two dimensions is equivalent to Hamiltonian flow — then yields symplectic structure, Liouville measure preservation, and canonical conjugate variables in each sector.

At $N \geqslant 3$ (three or more elements, closed relational loops), multiple coupling matrices are generically non-simultaneously-diagonalisable. A normalised trivector $\hat{T} = (a \wedge b \wedge c)/|a \wedge b \wedge c|$ satisfies $\hat{T}^2 = -1$, providing an intrinsic complex structure. We prove that this complex structure is preserved by the dynamics (using a symmetry argument specific to the relational framework), and that Ω conservation combined with $\hat{T}$ preservation implies the evolution preserves a Hermitian inner product — the defining property of unitary maps.

These mathematical results have structural parallels to Carcassi and Aidala's Reverse Physics programme, which identifies the minimal physical assumptions underlying classical and quantum mechanics. The "choice" between their starting assumptions — infinitesimal reducibility (yielding symplectic geometry) and irreducibility (yielding unitary geometry) — corresponds to the relational topology: open (N = 2) versus closed ($N \geqslant 3$). We develop this correspondence in detail, noting that two programmes designed without reference to each other converge on the same mathematical structures from opposite directions.

This paper explores the mathematical structures that emerge from the constraint framework's axiom. Connections to physical theories are developed as structural parallels — correspondences between mathematical forms — not as claims that physics is being derived from logic. The mathematical results stand independently of any physical interpretation.

---

1. Introduction

1.1 Motivation

Symplectic geometry and unitary groups are the mathematical foundations of classical and quantum mechanics respectively. These structures are typically introduced as mathematical frameworks chosen to model physical phenomena — symplectic manifolds because they accommodate Hamilton's equations, Hilbert spaces because they accommodate Schrödinger evolution. But the question of *why* these particular mathematical structures, and not others, is rarely addressed from a purely mathematical standpoint.

Carcassi and Aidala's Reverse Physics programme [1, 2, 3] provides one approach: working backward from the mathematical formalism of known theories, they identify the minimal assumptions from which each structure can be recovered. Their central finding is that the fork between symplectic and unitary geometry reduces to a single assumption about decomposability:

Infinitesimal reducibility (IR): the state of the whole is equivalent to the states of its parts. Combined with determinism and reversibility, this yields symplectic manifolds and Hamiltonian flow [3, 4].

Irreducibility (IRR): the state of the whole tells nothing about its parts. Combined with determinism and reversibility, this yields Hilbert spaces and unitary evolution [2, 3].

Reverse Physics requires a *choice* between IR and IRR as a starting assumption. It does not address when each holds or what determines the choice.

1.2 This Paper

We show that a constraint framework built from a single axiom in modal logic produces both symplectic and unitary structures as consequences of a single structural variable: the topology of the relational graph among distinguishable elements. When the graph has no closed loops ($N = 2$), the resulting geometry is symplectic. When it has closed loops ($N \geq 3$), the geometry is unitary. The "choice" between IR and IRR is determined by topology.

The axiom is $\Diamond N \to \neg N$: if nothingness is possible, then nothingness does not obtain. This is a statement in modal logic about distinguishability and existence. The mathematical structures that follow — Clifford algebra $Cl(5)$, conserved measures, symplectic forms, complex structures — are consequences of the axiom explored for their own interest. Where these structures parallel those of physical theories, we develop the correspondence in detail, following Carcassi's

vocabulary and results. But the mathematical content is independent of any physical interpretation.

1.3 Scope and Epistemological Position

This paper explores the mathematical structures that emerge from the constraint framework's axiom. Connections to physical concepts — Hamiltonian mechanics, quantum unitarity, the imaginary unit — are noted and developed as structural parallels: correspondences between mathematical forms arising in different contexts. They are not claims that physics is being derived from logic.

The distinction matters. Showing that a mathematical structure is *isomorphic* to the structure used in a physical theory is a mathematical result. Claiming that the physical theory *is* or *must be* described by that structure is a physical claim requiring additional bridge assumptions about what the mathematical objects represent in the physical world. This paper makes mathematical claims and identifies structural parallels, leaving physical interpretation open.

1.4 Relation to Other Work

The constraint framework [6] is related to but distinct from other programmes exploring connections between mathematical axiomatics and physical structure. Chiribella, D'Ariano, and Perinotti [7] derive the mathematical structure of quantum theory from five information-theoretic axioms. Hardy [8] derives it from five "reasonable" axioms. These programmes start from physical or operational axioms. Our axiom is ontological (about existence and distinguishability), and the resulting structures include both symplectic and unitary geometry as regime-dependent consequences, rather than one or the other.

---

2. The Framework

2.1 The Axiom and Its Mathematical Consequences

The axiom $\Diamond N \rightarrow \neg N$ states that nothingness is impossible. Any element in the framework must be distinguishable from nothingness and from other elements. This requires:

- A minimum separation $\varepsilon > 0$ in the configuration space (mutual distinguishability)
- A bounded configuration space (each constraint satisfies $\varepsilon < C\_i < 1 - \varepsilon$)
- A conserved measure Ω of total distinguishability (the measure cannot reach zero)

2.2 The Five Constraints and Cl(5)

For two elements to be robustly distinguishable — distinguishable under arbitrary observation conditions, not merely under specific measurements — their relational profile must differ along every independent axis of comparison. We identify these axes by categorical exhaustion:

systematically classifying what can differ between two distinguishable configurations and showing the classification is complete. The argument proceeds as follows.

Step 1: What must be specified. Any configuration of distinguishable elements has (a) spatial extent and boundaries between regions, (b) internal structure or pattern within regions, (c) the resources (energy, material, capacity) sustaining the configuration, and (d) the correlations linking distinct parts into a whole. These four categories — boundary (β), pattern (κ), resource (ρ), and integration (λ) — are pairwise properties: they characterise the relationship between two elements.

Step 2: Independence. The four categories are logically independent: specifying the boundary structure of a configuration does not determine its pattern complexity, resource density, or integration. This independence is confirmed empirically: principal component analysis across cellular automata (256 rules), Game of Life patterns (47 configurations), and Gray-Scott reaction-diffusion systems (202 parameter regimes) consistently requires five components, with no pair of constraints showing predictive redundancy ($R^2 < 0.02$ for all cross-predictions) [6].

Step 3: The fifth constraint. The four pairwise constraints are symmetric under exchange of relata (the β of the A–B relation equals the β of the B–A relation). A fifth category — ordering (τ) — captures asymmetric structure: the oriented circulation around a closed relational loop ($A \rightarrow B \rightarrow C \rightarrow A \neq A \rightarrow C \rightarrow B \rightarrow A$). This constraint is antisymmetric under exchange and requires at least three elements ($N \geqslant 3$) to be non-trivial. It is algebraically independent of the four symmetric constraints: no linear combination of symmetric quantities produces an antisymmetric one.

Step 4: Completeness. Any proposed sixth constraint must be either (a) a function of the existing five (hence redundant), (b) a higher-order combination (e.g., a product of two constraints, which belongs to a higher grade of the algebra rather than constituting a new independent axis), or (c) empirically distinguishable from all five (which the PCA analysis does not support — five components account for >96% of variance across all tested substrates). The five constraints are therefore both necessary and sufficient for robust distinguishability.

Step 5: Algebraic representation. Five independent constraint dimensions with the property that their pairwise products encode relational structure are naturally represented as the grade-1 basis vectors $e_1, ..., e_5$ of the Clifford algebra Cl(5), satisfying $e_i e_j + e_j e_i = 2\delta_{ij}$. This is not a choice of mathematical convenience but a consequence of the constraint structure: the geometric product $e_i e_j = e_i \cdot e_j + e_i \wedge e_j$ simultaneously encodes the correlation (scalar part) and oriented relation (bivector part) between constraint dimensions i and j.

The graded structure of Cl(5) provides a natural hierarchy:

- **Grade 1** (5 dimensions): Vectors — individual features (relational profiles)
- **Grade 2** (10 dimensions): Bivectors — pairwise relations
- **Grade 3** (10 dimensions): Trivectors — irreducible triple structures
- **Grade 4** (5 dimensions): Quadvectors — collective coherence

- **Grade 5** (1 dimension): Pseudoscalar — global orientation

A feature $\mathbf{a} = \Sigma\, a^i e\_i$ is a grade-1 element. The relation between features $\mathbf{a}$ and $\mathbf{b}$ is the bivector $B\_\{ab\} = \mathbf{a} \wedge \mathbf{b} \in \wedge^2 \mathbb{R}^5$.

Remark. The categorical exhaustion argument as presented here is a summary. The full treatment — including the formal independence proofs, the empirical validation across substrates, and the completeness argument — is developed in [6]. The results of the present paper depend on the Cl(5) structure but not on the specific physical interpretation of the five constraints; a reader who accepts the five-dimensional Clifford algebra as a mathematical starting point can follow the remainder without engaging the exhaustion argument.

2.3 What N Counts

A terminological clarification essential for the relational ontology. Throughout this paper, **N counts relational profiles (nodes)** — the grade-1 elements. A relational profile is defined by the totality of bivectors a feature participates in; it is not a substance with intrinsic properties. For N profiles, there are $C(N,2)$ pairwise relations (grade-2 bivectors), $C(N,3)$ triple structures (grade-3 trivectors), and so on.

Coupling matrices $M^{\wedge}(\alpha\beta)$, characterising how the constraint dimensions of profile α relate to those of profile β, are indexed by pairs of profiles (edges). Decomposability is a property of the set of coupling matrices.

2.4 Ω Conservation

Mathematical setting. The configuration space $M\_N = V^{\wedge}N$ is the N-fold product of the viable region $V \subset \mathbb{R}^5$, where $V = \{C \in \mathbb{R}^5 : \varepsilon < C\_i < 1 - \varepsilon$ for all $i\}$ is an open bounded subset of $\mathbb{R}^5$ with smooth boundary ∂V. The space M_N inherits the structure of a smooth Riemannian manifold with metric induced by the Fisher information metric $g\_\{ij\}$ on V. The measure μ is the Riemannian volume measure on M_N (equivalently, $\mu = \sqrt{\det(g)}\, dC_1 \wedge ... \wedge dC_{sn}$ in coordinates). The Ω density is a smooth positive function $\Omega: M\_N \to \mathbb{R}_+$, and the flow is generated by a smooth vector field X on M_N. These smoothness assumptions ensure the divergence theorem applies on any compact subregion of M_N with piecewise-smooth boundary.

Theorem 1 (Conservation of Total Distinguishability). The total distinguishability measure $\Omega\_\text{total} = \int\_V \Omega\, d\mu$ is conserved under any evolution consistent with the axiom.

Proof. The Ω density satisfies a continuity equation on the viable region V:

$$\frac{\partial \Omega}{\partial \lambda} + \mathrm{div}(\mathbf{j}_\Omega) = \sigma$$

where j_Ω is the Ω-current and σ is the source density. We establish $\sigma = 0$ and the boundary flux vanishes:

(i) **No sources or sinks (σ = 0).** The constraint field is all that exists — the axiom posits no external reservoir from which distinguishability could be created or into which it could be absorbed. Any σ ≠ 0 would require a mechanism for creating or destroying distinguishability ex nihilo, which contradicts the axiom's requirement that nothingness is impossible (creation from nothing) and that the viable region is bounded (unbounded creation). Therefore σ = 0 throughout V.

(ii) **No boundary flux.** The boundary ∂V corresponds to configurations where some constraint approaches 0 or 1 — the limits of distinguishability. The axiom forbids these limits (ε < C_i < 1 − ε), so the flow never reaches ∂V. Therefore j_Ω · n̂ = 0 on ∂V.

(iii) **Conservation.** Integrating the source-free continuity equation over V:

$$\frac{d}{d\lambda} \int_V \Omega \, d\mu = -\oint_{\partial V} \mathbf{j}_\Omega \cdot d\mathbf{A} + \int_V \sigma \, d\mu = 0 + 0 = 0$$

Therefore Ω_total is constant. Since Ω > 0 everywhere (the axiom forbids Ω = 0 at any point, not just globally), this is a non-trivial conservation law: the positive density is redistributed by the flow but neither created nor destroyed. ∎

Ω measures *relational* distinguishability — the aggregate of all bivector structure — not a count of features.

2.5 Coupling Matrices and Decomposability

For N features, the interaction structure is encoded in coupling matrices:

$$M_{ij}^{(\alpha\beta)} = \frac{\partial^2 \Phi_{\text{int}}}{\partial C_i^{(\alpha)} \partial C_j^{(\beta)}}$$

where Φ_int is the interaction potential and C_i^(α) is the i-th constraint value of feature α.

Definition 1 (Decomposability). An N-feature configuration is *decomposable* if there exists a basis in which all coupling matrices are simultaneously diagonal.

---

3. N = 2: Classical Hamiltonian Mechanics

3.1 Universal Decomposability

At N = 2, there is exactly one pair of features, hence one coupling matrix M^(12). This is a 5×5 real symmetric matrix.

Theorem 2 (Universal Decomposability at N = 2). For N = 2, the coupling matrix is always diagonalisable.

Proof. By the spectral theorem, every real symmetric matrix is orthogonally diagonalisable. Since there is only one matrix, "simultaneous" diagonalisation reduces to ordinary diagonalisation, which always succeeds. ∎

No genericity assumption is needed. Every N = 2 configuration is decomposable.

Since the ordering constraint τ is necessarily zero at N = 2 (τ requires closed loops, which cannot exist with only two features [6]), the diagonal basis decomposes the configuration space into four independent sectors, one for each symmetric constraint $\{\beta, \kappa, \rho, \lambda\}$. Each sector is a two-dimensional manifold with coordinates $(C_i^{(1)}, C_i^{(2)})$.

3.2 The 2D Theorem

Theorem 3 (Arnold [5]). On a two-dimensional manifold, a smooth flow is measure-preserving if and only if it is (locally) Hamiltonian.

Proof. Let $\mathbf{v} = (v_1, v_2)$ be a smooth vector field on a 2D manifold with coordinates (x_1, x_2). Measure preservation means the divergence vanishes: $\partial v_1 / \partial x_1 + \partial v_2 / \partial x_2 = 0$. By the Poincaré lemma, there exists a function H such that $v_1 = \partial H / \partial x_2$ and $v_2 = -\partial H / \partial x_1$. These are Hamilton's equations with symplectic form $\omega = dx_1 \wedge dx_2$. ∎

This equivalence holds *only in two dimensions*. In higher dimensions, measure preservation (divergence-free) is necessary but not sufficient for Hamiltonian flow. In 2D, every non-degenerate 2-form is automatically closed (there are no 3-forms on a 2-manifold), so every volume form is a symplectic form.

3.3 Symplectic Structure from the Axiom

Combining the three ingredients:

1. **Ω conservation** → the flow on the full configuration space preserves the Ω measure (Theorem 1, from the axiom).

2. **Decomposability** → the configuration space separates into independent 2D sectors (Theorem 2, from the spectral theorem).

3. **2D measure-preserving = Hamiltonian** → each sector carries Hamiltonian dynamics (Theorem 3, from Arnold).

Each sector has a Hamiltonian H_i, a symplectic form ω_i, and canonical Poisson brackets. H_i is conserved along flow lines (dH_i/dλ = {H_i, H_i} = 0 by antisymmetry). Darboux's theorem guarantees the local existence of canonical coordinates. Liouville's theorem — that the symplectic volume is preserved — is a restatement of Ω conservation in the decomposable regime.

Structural parallel to Carcassi: This result has the same mathematical form as Carcassi's theorem [3, 4] that determinism and reversibility (DR) plus infinitesimal reducibility (IR) yields Hamiltonian structure on symplectic manifolds. In the correspondence:

- DR $\leftrightarrow$ Ω conservation (from the axiom, Theorem 1)
- IR $\leftrightarrow$ decomposability (from the spectral theorem at N = 2, Theorem 2)

3.4 The Dynamical Hierarchy

Earlier formulations of the constraint framework describe the dynamics as "gradient flow on Φ" [6]. This requires clarification. Gradient flow increases Φ monotonically ($d\Phi/d\lambda = |\nabla\Phi|^2 \geqslant 0$), concentrating Ω into high-Φ regions, which is incompatible with Ω conservation.

The fundamental dynamics, as derived from the axiom, is Ω-preserving flow — which at N = 2 is Hamiltonian (this section).

Conjecture (Large-N effective dynamics). At large N, coarse-graining of many weakly-coupled Hamiltonian/unitary sectors produces effective dissipative dynamics on macroscopic variables, with gradient flow on Φ emerging as the leading-order effective description.

Motivation. This conjecture is structurally parallel to the standard result in statistical mechanics that coarse-graining of microscopic Hamiltonian dynamics produces macroscopic Langevin or gradient flow equations. The specific mechanism — many sectors with couplings of order ~1/N, producing effective friction coefficients and noise terms — is well-studied in the physics literature (see e.g. Zwanzig [15]). However, the formal derivation within the constraint framework (showing that Ω-preserving flow on M_N with large N reduces to gradient flow on coarse-grained variables) has not been carried out and remains an open problem (see Section 10.2).

3.5 Conjugate Variables

Define centre-of-mass and relative coordinates in each sector:

$$q_i = \frac{C_i^{(1)} + C_i^{(2)}}{2} \quad \text{(shared structure)}, \qquad p_i = C_i^{(1)} - C_i^{(2)} \quad \text{(distinguishing structure)}$$

In 2D, Ω conservation forces the flow to be tangent to level curves of Ω. The unique divergence-free vector field perpendicular to $\nabla\Omega$ is the Hamiltonian vector field. The canonical Poisson bracket {q_i, p_j} = δ_{ij} follows from the Hamiltonian structure.

The physical content of the conjugate pairing: q_i encodes what the two features share; p_i encodes what distinguishes them. The axiom requires both to be maintained. The symplectic structure is the mathematical expression of this dual requirement.

4. N ≥ 3: Generic Non-Decomposability

4.1 Loop Closure and Coupling

At N = 3, three features produce three coupling matrices: M^(12), M^(23), M^(31). Simultaneous diagonalisability requires all pairs to commute:

$$[M^{(12)}, M^{(23)}] = [M^{(23)}, M^{(31)}] = [M^{(31)}, M^{(12)}] = 0$$

Each commutator of 5×5 symmetric matrices is antisymmetric with at most 10 independent entries. The commutator conditions are polynomial equations in the matrix entries (each entry of [A, B] = AB − BA is a polynomial of degree 2 in the entries of A and B). The simultaneously-diagonalisable set is therefore a *real algebraic variety* — the zero set of a system of polynomial equations — within the 45-dimensional space of triples of 5×5 symmetric matrices.

To show this variety has measure zero, it suffices to show it is a *proper* subvariety — i.e., that the polynomial system is not identically zero. This is established by exhibiting a single triple (M^(12), M^(23), M^(31)) that does NOT simultaneously diagonalise. Any two generic 5×5 symmetric matrices fail to commute (their commutator is generically non-zero), so such triples are easy to construct. By the standard result that a proper real algebraic variety in $\mathbb{R}^n$ has Lebesgue measure zero [16], the simultaneously-diagonalisable set has measure zero in the parameter space.

(The three commutator conditions are not fully independent — they are related by identities of the form [A,[B,C]] + [B,[C,A]] + [C,[A,B]] = 0. This means the variety may have higher dimension than a naïve constraint count suggests. But the variety is still proper — it does not fill the ambient space — so the measure-zero conclusion holds.)

Theorem 4 (Generic Non-Decomposability at N ≥ 3). For N ≥ 3, the set of simultaneously-diagonalisable coupling matrix configurations is a proper algebraic subvariety of the parameter space, hence has measure zero. Generic N ≥ 3 configurations are non-decomposable.

The structural reason: three features form a closed loop (A → B → C → A). Each coupling matrix characterises an edge of this loop. The loop closure creates a consistency condition that generically prevents simultaneous diagonalisation. At N = 2, there is no loop, hence no obstruction.

4.2 The Mapping to Carcassi

Carcassi's Assumption	Constraint Framework	Mechanism
IR (reducible)	Decomposable	One coupling matrix; spectral theorem
IRR (irreducible)	Non-decomposable	Multiple coupling matrices; loop closure

Carcassi's Assumption	Constraint Framework	Mechanism
The "choice"	Determined by N	Whether relational topology has closed loops

The assumption that Carcassi treats as a starting point is here derived from the relational topology. The fork between classical and quantum is the fork between open (N = 2) and closed (N ⩾ 3) relational structures.

4.3 The Dimensional Obstruction

At N ⩾ 3, the configuration space does not decompose into independent 2D sectors. The flow lives in a higher-dimensional coupled space. Ω conservation still holds (the axiom doesn't depend on N), so the flow is measure-preserving. But in dimensions greater than 2, measure preservation does NOT imply Hamiltonian flow. There exist volume-preserving flows on 4D manifolds that preserve no symplectic form.

The symplectic structure of the N = 2 regime does not extend to N ⩾ 3. Something beyond Ω conservation is needed to determine the geometric character of the evolution. That additional structure is the trivector.

---

5. The Trivector Complex Structure

5.1 $T^2 = -1$

At N ⩾ 3, three features a, b, c ∈ $\mathbb{R}^5$ form a trivector:

$$T = \mathbf{a} \wedge \mathbf{b} \wedge \mathbf{c} \in \bigwedge^3 \mathbb{R}^5$$

This is generically non-zero — it vanishes only if the features are coplanar, which would collapse the system to effective N = 2 (collinearity of some pair violates mutual distinguishability [6]).

Theorem 5 (Trivector Squares to Minus One). For any unit trivector T̂ = T/|T| in Cl(5) with positive-definite signature:

$$\hat{T}^2 = -1$$

Proof. Choose an orthonormal frame where T̂ = $e_1 e_2 e_3$. Then:

$$\hat{T}^2 = (e_1 e_2 e_3)(e_1 e_2 e_3)$$

Using the anticommutation relations $e_a e_b = -e_b e_a$ for a ≠ b and $e_a^2 = 1$: move the second e_1 past e_3 and e_2 (two sign flips), giving $e_1^2 e_2 e_3 e_2 e_3 = e_2 e_3 e_2 e_3$. Then move the second e_2 past e_3 (one

sign flip), giving $-e_2{}^2 e_3{}^2 = -1$. For a general unit trivector, the result follows from the grade structure of Cl(5,0): the square of any simple grade-3 element has grade-0 component $-|T|^2$. ∎

5.2 The Complex Structure

The identity $T^2 = -1$ means $\hat{T}$ is algebraically identical to the imaginary unit. It defines a complex structure:

$$J(\mathbf{v}) = \hat{T}\mathbf{v}$$

satisfying $J^2 = -\mathrm{Id}$. This structure exists **only at N ≥ 3**. At N = 2, no trivector exists ($\tau = 0$), so there is no J, no complex structure, and no "i."

5.3 The Algebraic Hierarchy

Grade	Element	Squares to	Structure provided	N threshold
2	Bivector B = a ∧ b	−1	Rotation generator	N ≥ 2
3	Trivector T = a ∧ b ∧ c	−1	Complex structure	N ≥ 3

Both bivectors and trivectors square to −1, but they play different roles. Bivectors generate rotations within the decomposable (classical) regime. The trivector generates a complex structure in the irreducible (quantum) regime, coupling all sectors through a single irreducible action. The distinction: a bivector involves two features and acts within a single 2D sector; the trivector involves three features and couples sectors that would otherwise be independent.

5.4 Structural Parallel: The Complex Structure and "i"

In standard quantum mechanics, the imaginary unit appears in the Schrödinger equation, the commutation relations, and the inner product. Its presence is typically axiomatic.

The constraint framework produces a complex structure — $\hat{T}$ with $T^2 = -1$ — that is mathematically identical to the imaginary unit. It emerges at the precise structural threshold (N ≥ 3, closed loops) where the symplectic geometry of N = 2 is replaced by a different structure. Whether this mathematical correspondence reflects a deeper connection to physical quantum mechanics, or is merely a structural analogy, is an interpretive question beyond the scope of this paper. What is mathematical is the emergence of the complex structure itself from the axiom through Cl(5).

The framework also provides a structural reason for *complex* rather than real or quaternionic geometry: real structures correspond to N = 2 (no trivector); complex structures correspond to N ≥ 3 (one trivector providing $T^2 = -1$); quaternionic structure would require additional conditions not generically satisfied by the framework's geometry.

6. The T̂ Commutation Proof

6.1 The Question

Theorem 5 establishes $T^2 = -1$. To derive unitarity, we need the evolution to commute with the complex structure: $[U, J] = 0$ where $J(v) = T̂v$. This requires showing $L\_X(T̂) = 0$, where L_X is the Lie derivative along the flow.

6.2 The Symmetry Argument

Step 1: Relational dependence. Both the efficiency potential Φ and the measure Ω depend on the features only through their mutual inner products — the Gram matrix G:

$$\Phi(\mathbf{a}, \mathbf{b}, \mathbf{c}) = F(\mathbf{a} \cdot \mathbf{b}, \ \mathbf{b} \cdot \mathbf{c}, \ \mathbf{a} \cdot \mathbf{c}, \ |\mathbf{a}|^2, \ |\mathbf{b}|^2, \ |\mathbf{c}|^2)$$

This is because distinguishability is relational — it depends on how features relate to each other, not on absolute orientation. The Gram matrix is invariant under any orthogonal transformation R $\in$ SO(5) applied to all features simultaneously: (Ra)·(Rb) = a·b.

Step 2: The residual symmetry. Three generic features span a 3-dimensional subspace $V_3 \subset \mathbb{R}^5$. The orthogonal complement $V_3\perp$ has dimension 2. Both Φ and Ω are exactly invariant under $SO(V_3\perp) \cong SO(2)$ — rotations in the 2D complement — because these rotations preserve the Gram matrix.

Theorem 6 (Confinement). Any smooth flow determined by Φ and constrained by Ω conservation keeps the features in their spanned subspace: L_X $\mathbf{a}$, L_X $\mathbf{b}$, L_X $\mathbf{c} \in V_3$.

Proof. Suppose L_X $\mathbf{a}$ has a nonzero component v$\perp \in V_3\perp$. Consider R_$\pi \in$ SO($V_3\perp$) mapping v$\perp$ → −v$\perp$ (rotation by π in the plane containing v$\perp$, which exists because $V_3\perp$ is 2D). Under R_π: the configuration (a, b, c) is unchanged (R_π fixes V_3 pointwise), so Φ, Ω, and the full dynamical state are unchanged. By equivariance, the flow must also be unchanged. But R_π maps L_X $\mathbf{a}$ = v‖ + v$\perp$ to v‖ − v$\perp \neq$ v‖ + v$\perp$. Contradiction. Therefore v$\perp$ = 0. ∎

6.3 From Confinement to T̂ Preservation

Theorem 7 (T̂ Preservation). The normalised trivector is preserved by the dynamics: L_X(T̂) = 0.

Proof. By the Leibniz rule:

$$L_X(T) = (L_X \mathbf{a}) \wedge \mathbf{b} \wedge \mathbf{c} + \mathbf{a} \wedge (L_X \mathbf{b}) \wedge \mathbf{c} + \mathbf{a} \wedge \mathbf{b} \wedge (L_X \mathbf{c})$$

By Theorem 6, each velocity lies in V_3 = span(a, b, c). Therefore L_X(T) $\in \wedge^3 V_3$. Since dim($\wedge^3 V_3$) = 1 (the space of trivectors over a 3D subspace is one-dimensional), L_X(T) = αT for some scalar α.

For the normalised trivector:

$$L_X(\hat{T}) = \frac{L_X(T)}{|T|} - T \cdot \frac{L_X(|T|)}{|T|^2}$$

From L_X(T) = αT: L_X(|T|) = α|T|. Substituting: L_X(T̂) = αT̂ − αT̂ = 0. ∎

6.4 Commutation Follows

Since T̂ is constant along the flow, for any state ψ:

$$L_X(J\psi) = L_X(\hat{T}\psi) = (L_X\hat{T})\psi + \hat{T}(L_X\psi) = 0 + J(L_X\psi)$$

Therefore [L_X, J] = 0 at every point. For the finite evolution U = exp(λ L_X): [U, J] = 0.

6.5 Why This Argument Is Specific to the Relational Framework

The proof uses two properties specific to the constraint framework:

1. **Φ and Ω depend only on the Gram matrix** — a consequence of the axiom's requirement that all structure is defined through distinguishability relations. An external potential or preferred direction in $\mathbb{R}^5$ would break $SO(V_3\perp)$ invariance and invalidate Theorem 6.
2. **dim($\wedge^3 V_3$) = 1** — the specific dimensionality: three elements span a 3D subspace whose top exterior power is 1D. At N ⩾ 4, where elements span V_4 with dim($\wedge^3 V_4$) = 4, the argument requires extension (see Section 10).

---

7. Unitarity at N ⩾ 3

7.1 The Inner Product

In Cl(5), the reversal operation † (reversing the order of geometric products, with $v^\dagger = v$ for vectors) defines a natural inner product:

$$\langle\psi|\phi\rangle = \langle\psi^\dagger\phi\rangle_0$$

where $\langle\cdot\rangle_0$ denotes the grade-0 (scalar) projection. For vectors, this reduces to the dot product.

7.2 The Main Theorem

Theorem 8 (Unitarity from Ω Conservation + T). At N ⩾ 3, any evolution that preserves Ω and whose generator commutes with J(v) = T̂v preserves the full inner product $\langle\psi|\phi\rangle$.

Proof. Ω conservation requires norm preservation: $\langle\psi|\psi\rangle$ = const for all states ψ. The polarisation identity expresses the inner product in terms of norms:

$$\langle\psi|\phi\rangle = \tfrac{1}{4}\left[|\psi + \phi|^2 - |\psi - \phi|^2 + i|\psi + i\phi|^2 - i|\psi - i\phi|^2\right]$$

where "i" is T̂. If the evolution U preserves all norms (condition a) and commutes with multiplication by T̂ (condition b, established by Theorems 6–7), then each term is individually preserved:

- $|U\psi + U\varphi|^2 = |U(\psi + \varphi)|^2 = |\psi + \varphi|^2$ (norm preservation of ψ + φ)
- $|U\psi + T̂(U\varphi)|^2 = |U\psi + U(T̂\varphi)|^2 = |U(\psi + T̂\varphi)|^2 = |\psi + T̂\varphi|^2$ (commutativity + norm preservation)

Therefore $\langle U\psi|U\varphi\rangle = \langle\psi|\varphi\rangle$. An evolution preserving the inner product satisfies $U^{\dagger}U = I$: it is unitary. ∎

7.3 Why Ω Conservation Alone Is Insufficient

The proof requires BOTH Ω conservation and T̂ commutation. Neither alone suffices.

Ω conservation without T̂ gives measure-preserving but not necessarily inner-product-preserving evolution — the large group of volume-preserving diffeomorphisms rather than the small group of unitaries. T̂ without Ω conservation gives complex-linear maps that may change norms — the full GL(n,ℂ) rather than U(n). Only their combination selects the unitary group.

7.4 Structural Correspondence with Carcassi's Programme

The following table summarises the structural parallels between the mathematical content of the axiom and the assumptions Carcassi identifies as underlying classical and quantum mechanics.

Carcassi's Assumption	Axiom Framework Structure	Mathematical Mechanism
Determinism (D)	Ω conservation	Axiom forbids $\Omega = 0$; no sources or sinks
Reversibility (R)	Ω conservation	Measure preservation is bijective
Inf. reducibility (IR)	$N = 2$ decomposability	Spectral theorem (one symmetric matrix)
Irreducibility (IRR)	$N \geqslant 3$ non-decomposability	Loop closure (generically non-simultaneous-diagonalisability)
Symplectic structure	2D sectors with ω_i	Non-degenerate 2-form in 2D is automatically closed
Complex structure (i)	Trivector $T^2 = -1$	Emerges only in the irreducible regime (Theorem 5)
Hamiltonian flow	Ω + decomposability + 2D	Arnold's theorem (Theorem 3)
Unitary evolution	$\Omega + T\hat{} +$ commutation	Polarisation identity (Theorem 8)

The correspondence is structural: each row identifies the same mathematical form appearing in two different contexts. The axiom framework produces these forms through the mechanisms listed; Carcassi's programme identifies them as the minimal mathematical content of physical theories. That the same structures emerge from both directions — one working forward from an axiom, the other backward from established mathematics — is the central observation of this paper.

8. The Entropy–Geometry Equivalence

8.1 Two Independent Paths

The preceding sections derive the geometric structure of state space through *dynamics*: Ω-preserving flow in 2D sectors is Hamiltonian; Ω-preserving flow respecting $T\hat{}$ is unitary. Carcassi's recent work [4, 11] provides a second, *structural* path to the same conclusion — one that connects the geometric structure of state space to the information-theoretic requirement that entropy be well-defined, without invoking dynamics.

8.2 Carcassi's Entropy–Geometry Theorem

Shannon entropy (classical) and von Neumann entropy (quantum) are defined relative to a reference measure on state space. The value of entropy changes under transformations that do

not preserve this measure. Carcassi proves [4]:

Theorem (Carcassi, 2020). The group of transformations under which information entropy is invariant is exactly the symplectomorphism group. The symplectic form is the unique geometric structure whose preservation is equivalent to entropy invariance.

In the quantum case, the complex structure J defines a symplectic form $\omega(\psi, \varphi) = \text{Im}\langle\psi|\varphi\rangle$. The entropy-invariance group is the unitary group — the maps preserving both the inner product and the symplectic form.

8.3 The Structural Path in the Constraint Framework

The axiom requires Ω to be a property of the relational structure itself, not of any particular description. If Ω changed under redescription, "distinguishability" would not be well-defined. This gives a structural condition:

$$\Omega(\phi(C)) = \Omega(C) \quad \text{for all relational-content-preserving } \phi \qquad (*)$$

By Carcassi's theorem, the invariance group of an entropic quantity determines a symplectic structure. At N = 2, condition () *yields area-preserving diffeomorphisms of each 2D sector — the symplectomorphism group. At $N \geq 3$, condition* () combined with the complex structure T^yields the unitary group.

8.4 Convergence of the Two Paths

The paper now contains two independent derivations:

Path 1 (dynamical): Ω conservation in time + decomposability (N = 2) → Hamiltonian. Ω conservation in time + T^(N $\geq$ 3) → unitary.

Path 2 (structural): Ω description-independence + Carcassi's equivalence → symplectic structure. Ω description-independence + T^→ unitary group.

Path 1 uses Ω conservation as a dynamical constraint (time-invariance). Path 2 uses Ω description-independence as a structural constraint (coordinate-invariance). These are logically independent conditions, but the axiom requires both. That both paths produce the same geometric structure is a non-trivial consistency check on the framework's architecture.

---

9. The Regime Hierarchy

9.1 Summary

The axiom produces different geometric structures depending on the relational topology:

N Regime	Coupling	Additional Structure	Geometry	Status
N = 2	One matrix, diagonalisable	Independent 2D sectors	Symplectic (Thm 3)	**Proved**
N ⩾ 3, small	Multiple matrices, non-diag.	Trivector $T^2 = -1$	Unitary (Thm 8)	**Proved** (N = 3)
N → ∞	Statistical aggregate	Coarse-graining	Effective dissipative	**Conjectured**

9.2 The Large-N Regime (Conjectural)

At large N, individual coupling matrices are weak (~1/N per relation). The effective coupling for any pair of elements is approximately a single matrix, which is always diagonalisable. If this approximation holds rigorously, the decomposable (symplectic) structure of N = 2 re-emerges at the coarse-grained level, with loop corrections suppressed as $\sim 1/N^3$. This would provide a structural parallel to the standard observation that macroscopic systems are well-described by classical mechanics despite being composed of quantum-mechanical parts. The formal derivation of this large-N limit remains open (see Section 10.2).

9.3 Smooth Interpolation

The transition between geometric regimes is smooth. Per-relation constraint strength decreases continuously as ~1/N; cyclic contributions decrease as $\sim 1/N^3$. There is no critical N at which the geometry changes discontinuously.

10. Discussion

10.1 What Is Proved

#	Result	Method	Status
1	Ω conservation from axiom	No sources/sinks argument	Theorem 1
2	N = 2 always decomposable	Spectral theorem	Theorem 2
3	2D measure-preserving ↔ Hamiltonian	Poincaré lemma	Theorem 3
4	N ⩾ 3 generically non-decomposable	Dimension counting	Theorem 4
5	$T^2 = -1$	Cl(5) anticommutation	Theorem 5

#	Result	Method	Status
6	Features confined to V_3	$SO(V_3\perp)$ equivariance	Theorem 6
7	$\hat{T}$ preserved by evolution	Confinement + dim $\wedge^3 V_3 = 1$	Theorem 7
8	$\Omega + [U, \hat{T}] = 0 \to$ unitarity	Polarisation identity	Theorem 8

10.2 Scope Limitations

The extension to N > 3. The $\hat{T}$ commutation proof (Theorems 6–7) works at N = 3 where dim($\wedge^3 V_3$) = 1. At N = 4, the trivector space is 4-dimensional and individual trivectors can rotate into each other under the evolution. A companion paper [10] shows that Gram-preserving evolution induces an SO(3) rotation on the 3-dimensional K_4 cycle space — a gauge transformation that preserves unitarity up to conjugation. The isometry group is SO(3) because any positive-definite inner product on $\mathbb{R}^3$ has O(3) as its isometry group (a standard result from linear algebra, proved via Cholesky decomposition). Physical predictions are invariant under this gauge group. The SO(3) gauge freedom has the same dimension (3) as the spatial dimensions that emerge at the same N = 4 threshold [11], a structural parallel explored in [10]. The full extension to N ⩾ 5, where the cycle space is 6-dimensional and the Gram matrix is no longer of the form $\alpha I + \beta J$, remains open.

The specific form of evolution equations. We establish that the evolution preserves a Hermitian inner product (unitarity) and symplectic structure (Hamiltonian character). The specific *form* of evolution equations — which would correspond to the Schrödinger equation or Hamilton's equations in the physical parallel — requires additional structure beyond what is established here.

The large-N limit. The emergence of effective dissipative dynamics from underlying conservative dynamics through coarse-graining is a well-studied mathematical programme. The framework predicts specific scaling exponents but the systematic derivation of Langevin-type equations remains open.

10.3 The Convergence

The two programmes — Carcassi's Reverse Physics and the constraint framework — were developed independently and from opposite directions. Reverse Physics begins with established mathematical frameworks and identifies minimal assumptions. The constraint framework begins with a single axiom and constructs mathematical structures upward.

That both programmes arrive at the same structural distinction — decomposable (symplectic) versus irreducible (unitary) geometry, determined by the same mathematical properties (simultaneous diagonalisability of coupling matrices) — from opposite starting points is the central observation of this paper. We present this as a mathematical fact about the relationship between the two programmes, not as a claim about physics.

The convergence suggests that the symplectic and unitary structures are in some sense *necessary* — they are what you get whether you work forward from distinguishability or backward from the mathematical content of physical theories. But "necessary for what" is a question that depends on interpretation.

10.4 Structural Parallel: Co-Emergence

A companion paper [11] proves that three-dimensional Euclidean geometry (with round S^2 and cross product) emerges from the cycle space of K_4 — the complete graph on four vertices. That proof establishes all component results rigorously: the cycle rank formula gives dim = 3; the edge-sharing Gram matrix has the form $G_3 = 2I + J$; the central identity $PG_3P = 2P$ (roundness) is proved algebraically and verified computationally; the cross product, Jacobi identity, and $SO(3)$ symmetry are confirmed; and the uniqueness of K_4 among complete graphs is proved both by exhaustive computation (K_3 through K_6) and by arithmetic (the equation $(N-1)(N-2)/2 = 3$ has unique positive solution $N = 4$).

The mechanisms of this paper — non-decomposability, trivector complex structure, unitary geometry — activate at $N = 3$. Both results share the same structural cause: the formation of closed loops in the relational graph.

This co-emergence — three-dimensional geometry and unitary structure appearing at the same structural threshold — is a mathematical fact about the framework. Whether it has physical significance (relating the dimensionality of space to the character of quantum mechanics) is an interpretive question that we note but do not attempt to resolve.

Acknowledgments

References

[1] G. Carcassi and C.A. Aidala, "Reverse Physics: From Laws to Physical Assumptions," *Foundations of Physics* **52**, 40 (2022).

[2] G. Carcassi and C.A. Aidala, *Assumptions of Physics*, Version 2.0 (Michigan Publishing, 2023). assumptionsofphysics.org/book

[3] G. Carcassi, C.A. Aidala, D.J. Baker and L. Bieri, "From physical assumptions to classical and quantum Hamiltonian and Lagrangian particle mechanics," *J. Phys. Comm.* **2**, 045026 (2018).

[4] G. Carcassi and C.A. Aidala, "Hamiltonian mechanics is conservation of information entropy," *Stud. Hist. Phil. Sci. B* **70**, 19–27 (2020).

[5] V.I. Arnold, *Mathematical Methods of Classical Mechanics* (Springer, 1989).

[6] D. Neale, "Being from Nothingness: Deriving the Structure of Existence from a Single Axiom." Preprint available at goleudy.ai.

[7] G. Chiribella, G.M. D'Ariano, and P. Perinotti, "Informational derivation of quantum theory," *Phys. Rev. A* **84**, 012311 (2011).

[8] L. Hardy, "Quantum theory from five reasonable axioms," arXiv:quant-ph/0101012 (2001).

[9] G. Carcassi, L. Maccone, and C.A. Aidala, "Four postulates of quantum mechanics are three," *Phys. Rev. Lett.* **126**, 110402 (2021).

[10] D. Neale, "Unitarity at N = 4: Trivector Rotation and the Round Cycle Space." FC–Reverse Physics Bridge 3. Preprint available at goleudy.ai.

[11] D. Neale, "Three-Dimensional Geometry from a Distinguishability Axiom: The Cycle Space of K_4 and Frame-Invariant Directional Distributions." Preprint available at goleudy.ai.

[12] C. Doran and A. Lasenby, *Geometric Algebra for Physicists* (Cambridge University Press, 2003).

[13] D. Hestenes, "Spacetime physics with geometric algebra," *Am. J. Phys.* **71**, 691 (2003).

[14] J.A. Barandes, "The stochastic-quantum theorem," arXiv:2309.02691 (2023).

[15] R. Zwanzig, *Nonequilibrium Statistical Mechanics* (Oxford University Press, 2001).

[16] J. Bochnak, M. Coste, and M.-F. Roy, *Real Algebraic Geometry* (Springer, 1998).