Half-Integer Spin as Closure Two-Cycle Structure in the VERSF Framework

K. Taylor *VERSF Theoretical Physics Programme*

Abstract for the General Reader

A spinning electron has a strange property: turning it once around does not return it to its starting state. It must be turned twice — through 720°, not 360° — before it looks the same as when it started. This is one of the genuine oddities of quantum mechanics, and it underlies all of ordinary matter: every electron, proton, and neutron in the universe behaves this way.

The standard explanation is mathematical. Particles like the electron are described by a special kind of object called a "spinor," whose mathematics happens to require two full turns rather than one. But this tells us *what* the rule is, not *why* the universe should have it.

This paper offers a *why*. It builds on a broader framework — the Void Energy-Regulated Space Framework (VERSF) — in which physical reality is constructed from irreversible distinctions and from the requirement that those distinctions remain consistent when they form loops. We show that a small set of foundational principles already established in that framework, together with one geometric fact about three-dimensional space (it has a kind of topological "twist" that takes two turns to undo), forces the 720° behaviour of spin-½ particles. The same construction predicts both kinds of fundamental particle — those that return on one turn (integer spin) and those that need two (half-integer spin) — from a single underlying structure. It also clarifies why exotic "anyon" particles can exist in two-dimensional systems but not in three.

No new mathematics is introduced. The standard mathematics of spin is shown to follow inevitably from a deeper structural principle, rather than being assumed.

Abstract

We derive the spin-½ transformation law from structural requirements within the Void Energy-Regulated Space Framework (VERSF), without postulating double-valued representations of the rotation group. Physical states are modelled as two-layer commitment objects $\Psi = (b, \varphi)$, where b is a distinguishability bit and $\varphi \in \mathbb{Z}_2$ is a closure phase tracking the cyclic consistency of the commitment network. We show that, given BC1 (distinguishability conservation), BC2 (finite information capacity), FIM (geometric distinguishability), CCC (Commitment Condition Criterion), and the binary structure of primitive commitment, three structural results follow: closure invariants factor through homotopy classes of admissible cycles (Theorem 3.1); continuous admissible transformations — including spatial rotations — induce admissible closure cycles, so the action of rotations on the closure register factors through $\pi_1(SO(3))$

(Propositions 3.2–3.4); and this factorisation is faithful and forces the non-trivial representation ρ_1 on the closure register (Theorem 3.5 and Corollary 3.6, both derived from BC1 and CCC). For three-dimensional space, $\pi_1(SO(3)) = \mathbb{Z}_2$, forcing a binary closure phase that inverts under 2π. Postulate L is reduced to the near-definitional statement that spatial rotations are continuous admissible transformations of physical states. The trivial action ρ_0 is excluded at the primitive level by Theorem 3.1 but emerges as the derived action on composite states in which constituent closure phases pair and cancel. The framework therefore predicts both half-integer spin (the primitive) and integer spin (composite), with the spinor 4π periodicity arising as the topological signature of the unique admissible primitive. The closure-complete transformation group is forced to be the universal cover SU(2). The standard $\mathbb{C}^2$ spinor is recovered as the minimal complex linear envelope of (b, φ), modulo the separately-treated VERSF derivation of complex Hilbert space from distinguishability. A composition rule generates the full integer/half-integer spectrum from tensor products of the primitive (b, φ) unit; a route to the Pauli exclusion principle via closure saturation is sketched and clearly labelled conjectural.

Table of Contents

1. Introduction

In standard quantum mechanics, spin enters through representation theory: particles are classified by irreducible representations of SU(2), the universal cover of the spatial rotation group SO(3). Half-integer spin corresponds to faithful representations of SU(2) that descend only projectively to SO(3), with the defining property

$R(2\pi)\ \psi = -\psi,\ R(4\pi)\ \psi = +\psi.$

This is mathematically clean but physically opaque. Why does the universe insist on tracking information beyond observable orientation? Why is the relevant group the *cover* rather than the geometric rotation group itself?

Within the VERSF programme, physical reality is constructed from irreversible distinctions ("facts") subject to admissibility and closure constraints. Particle properties should therefore be derivable from the structure of commitment and closure, not imposed by representation-theoretic fiat. This paper develops that perspective for spin-½.

The central claim is that half-integer spin is the signature of a state that carries two layers of structure: an **observable** layer (the distinguishability bit b) and a **closure** layer (the binary phase φ tracking cyclic consistency). The 4π periodicity follows once one notes that closure must respect the topology of the rotation path, and $\pi_1(SO(3)) = \mathbb{Z}_2$ admits exactly one non-trivial binary representation.

---

2. VERSF Preliminaries

We state only what is needed below.

2.1 Distinction and commitment

A physical state arises from an irreversible binary distinction:

$b \in \{0, 1\},\ \Delta S_{min} = k_B \ln 2.$

Such a distinction, once committed, is termed a **fact**.

2.2 Admissibility and closure

A distinction is physically admissible only if it participates in a closed commitment structure. Closure enforces:

- observer-invariant realisability,
- consistency under comparison and composition,

- finite distinguishability of the resulting fact.

Closure is therefore a property of *cycles* in the commitment network, not of isolated distinctions.

2.3 The two-layer state

We define a minimal admissible state as the pair

$\Psi = (b, \varphi),\ b \in \{0, 1\},\ \varphi \in \{+1, -1\}.$

Here b is the distinguishability bit and φ is the **closure phase**, an element of $\mathbb{Z}_2$ recording whether the commitment network has completed an even or odd number of closure cycles relative to a chosen reference. The reasons for the binary range of φ are addressed in §3.

2.4 BC1, BC2, FIM, and CCC

The derivations in §3.1 and §3.2 invoke four named results from the BCB foundations programme. We summarise them here so the present paper stands alone.

- **BC1 (distinguishability conservation).** Admissible evolution must preserve distinguishability and remain reversibly traceable. Trajectories cannot merge, collapse, or lose their histories without violating distinguishability conservation. Internal flows remain diffeomorphic.
- **BC2 (finite information capacity).** A physical state cannot retain redundant microscopic detail beyond its finite distinguishability. Closure invariants are restricted to information that survives admissible smooth deformation; arbitrary parametrisation-dependent data is excluded.
- **FIM (geometric distinguishability metric).** Distinguishability is geometric: only structure preserved under admissible smooth transformations may count as physical. Local comparison of states across fibres requires a connection, and curvature measures the path-dependence of distinguishability transport around loops. Holonomy is therefore the physical measure of path-dependent comparison, not an optional embellishment.
- **CCC (Commitment Condition Criterion).** A physical state is admissible only if its commitment structure is (i) closed, i.e. cycle-consistent; (ii) observer-invariant; and (iii) stable under admissible transformations of the distinguishability structure. In particular, any admissible transformation must map closure-complete states to closure-complete states.

Full statements and derivations are given in the BCB foundations programme [Taylor, *VERSF Constraint and Lagrangian Representations*; Taylor, *The Commitment Condition Criterion (CCC)*]. The present paper invokes BC1, BC2, FIM, and CCC as named results from that work.

CCC has multiple equivalent formulations across the VERSF programme — the structural-threshold form (the quartic inequality $\chi(L) = \rho L^4/\hbar c$), the action-principle form, and the operational form. The form invoked here, in §3.2 and §3.3, is the operational form stated above:

admissible transformations must preserve closure-completeness. This is the relevant face for the present derivation; the equivalence of formulations is treated in the CCC paper.

---

3. The Closure Phase as a $\mathbb{Z}_2$ Register

The closure phase φ is forced by one primitive fact about commitments, three derived results, and one near-definitional axiom. We name them in turn.

Postulate B (bit-level commitment). Every commitment is a bit. A register tracking commitments at the most primitive level therefore takes values in $\mathbb{Z}_2$.

Theorem 3.1 (path sensitivity of closure). Closure is a property of admissible commitment cycles. Since BC1 preserves distinguishability histories, BC2 excludes redundant microscopic path detail, and FIM restricts physical structure to deformation invariants, closure invariants factor through the homotopy class of closed admissible paths and are non-trivial whenever the homotopy class is non-trivial. Therefore closure-complete states distinguish paths that return to the same observable endpoint when those paths belong to different homotopy classes. The full statement and proof are given in §3.1.

Propositions 3.2–3.4 (rotation–closure coupling). Continuous admissible transformations induce closure cycles in the admissible state space (Proposition 3.2). Spatial rotations are continuous admissible transformations, and therefore define such cycles (Corollary 3.3). The action of rotations on closure invariants accordingly factors through $\pi_1(SO(3))$ (Proposition 3.4). The full statements and proofs are given in §3.2.

Theorem 3.5 (faithfulness of the rotation–closure mapping) and Corollary 3.6 (ρ_1 uniqueness). The induced mapping $\pi_1(SO(3)) \to \pi_1(A)$ is faithful: the non-trivial homotopy class of SO(3) does not collapse to the trivial class in the admissible commitment space (Theorem 3.5). The action of $\pi_1(SO(3))$ on the closure register must therefore be the non-trivial homomorphism ρ_1; the trivial action ρ_0 is incompatible with closure completeness (Corollary 3.6). Faithfulness is forced by BC1 (loss of distinguishability and reversibility) and CCC (incomplete closure encoding). The full statements and proofs are given in §3.3.

Postulate L (fully minimal). Spatial rotations act on physical states as continuous admissible transformations.

By Proposition 3.2 and Theorem 3.5, this induces a faithful action of $\pi_1(SO(3))$ on closure invariants. Postulate L is therefore reduced to the near-definitional content that rotations exist within the framework as continuous admissible transformations of admissible states; everything previously attributed to L is now derived.

Together, Postulate B, Theorem 3.1, Propositions 3.2–3.4, Theorem 3.5, Corollary 3.6, and Postulate L yield a $\mathbb{Z}_2$-valued closure register on which $\pi_1(SO(3))$ acts non-trivially, faithfully, and uniquely as ρ_1. The $2\pi/4\pi$ physical content of ρ_1 is identified in §4.2; the corresponding spinor structure follows in §5–§6.

The binary nature of φ is forced, not chosen. Since $\pi_1(SO(3))$ has only two homotopy classes, a closure register larger than $\mathbb{Z}_2$ cannot be irreducibly acted upon by $\pi_1(SO(3))$: it would split into a direct sum of $\mathbb{Z}_2$ subregisters, none more fundamental than the others. ($\mathbb{Z}_2$ has only two irreducible representations — trivial and sign — and the sign representation is one-dimensional.) The $\mathbb{Z}_2$ register is therefore uniquely forced by joint minimality and irreducibility, not merely the smallest sufficient choice.

3.1 Closure Invariants as Homotopy Invariants

The path sensitivity of closure need not be treated as an independent postulate. It follows from the existing BCB/VERSF structure.

BC1 requires distinguishability-preserving evolution to remain reversible and globally traceable. In the BCB foundations paper, this is expressed as the requirement that internal flows remain diffeomorphic and do not form caustics: trajectories cannot merge, collapse, or lose their histories without violating distinguishability conservation.

Similarly, local comparison of states across different fibres requires a connection; curvature then measures the path-dependence of distinguishability transport around loops. In BCB language, holonomy is therefore not an optional embellishment but the physical measure of path-dependent comparison.

Closure is precisely the requirement that a commitment structure remain consistent after completing a cycle. It is therefore not a property of an isolated endpoint, but of a closed admissible path. Two closed paths with the same endpoint may nevertheless differ in whether one can be continuously deformed into the other without leaving the admissible state space. If such deformation is impossible, then identifying the two paths would erase distinguishability history and violate the reversible traceability required by BC1.

At the same time, BC2 forbids retaining unnecessary path detail. Closure cannot encode every microscopic feature of a path, because such information would exceed finite distinguishability and introduce redundant internal structure. Only deformation-invariant path information may remain. Therefore, the natural object tracked by closure is the homotopy class of the closed path.

We can state this as a theorem.

Theorem 3.1 — Closure Invariants Are Homotopy Invariants. *Let A be the admissible state space of a commitment system satisfying BC1, BC2, and FIM. Let γ be a closed admissible path in A, representing a completed commitment cycle. Then any physically admissible closure invariant C(γ) factors through the homotopy class of γ — that is,*

$C(\gamma) = C(\gamma')$ whenever $\gamma \simeq \gamma'$ within A.

Moreover, for closure-complete states, C must be non-trivial whenever the homotopy class [γ] is non-trivial — i.e. $C(\gamma) \neq C(\gamma_0)$ whenever $[\gamma] \neq [\gamma_0]$, where γ_0 denotes a contractible loop — since closure-complete states must exhaust all admissible distinguishability permitted by BC1 and BC2.

Proof.

Closure is defined on completed cycles, not isolated endpoints. Hence C must be a function of closed paths.

By FIM, distinguishability is geometric: only structure preserved under admissible smooth transformations may count as physical. Therefore, closure cannot depend on arbitrary parametrisation, coordinate choice, or microscopic path deformation. Hence $C(\gamma) = C(\gamma')$ whenever γ and γ' are continuously deformable into one another within A. This establishes that C factors through the deformation classes of closed admissible paths — i.e. through the homotopy classes [γ].

By BC2, redundant internal information is excluded. Hence C cannot retain information beyond the deformation class. The minimal invariant distinguishing non-deformable cycles while remaining finite under BC2 is precisely the homotopy class. C is therefore a function of [γ] and no finer invariant.

By BC1, admissible evolution must preserve distinguishability and remain reversibly traceable. If two closed paths γ and γ_0 are not homotopic within A, then they represent genuinely distinct admissible commitment histories. Identifying them — i.e. assigning $C(\gamma) = C(\gamma_0)$ — would erase a real distinction between possible histories and violate reversible traceability. Hence for closure-complete states, $C(\gamma) \neq C(\gamma_0)$ whenever $[\gamma] \neq [\gamma_0]$.

Therefore physically admissible closure invariants factor through the homotopy class, and are non-trivial whenever the homotopy class is non-trivial. ■

Remark. Any invariant strictly finer than homotopy would violate BC2 by encoding redundant path detail, while any invariant strictly coarser would violate BC1 by identifying non-equivalent histories. The homotopy class is therefore not merely a sufficient choice but the unique invariant compatible with both bounds simultaneously.

The strength of this result is that it converts the logic chain from "assume closure is path-sensitive → derive spinor periodicity" to "closure + BC1 + BC2 + FIM → path sensitivity → $\mathbb{Z}_2$ closure register → spinor periodicity in §4."

3.2 From CCC and Admissibility to Rotation–Closure Coupling

We now show that the action of spatial rotations on the closure register is not an independent assumption, but follows — up to a minimal residual identification — from the Commitment Condition Criterion (CCC) and admissibility. The CCC is stated operationally in §2.4.

Proposition 3.2 — Continuous Physical Transformations Induce Closure Cycles. *Let A be the admissible state space satisfying BC1, BC2, FIM, and CCC. Let T(t), $t \in [0, 1]$, be a continuous family of admissible physical transformations such that*

T(0) = T(1) = identity on observable configuration.

Then T defines a closed admissible path γ in A, and therefore a closure cycle.

Proof.

By admissibility and CCC, each T(t) maps admissible states to admissible states; the family therefore lies entirely within A.

By continuity, the trajectory of any state under T(t) defines a continuous path in A.

Since T(0) = T(1) at the observable level, the path is closed in the observable sector.

Closure is defined on commitment cycles, not endpoints (§2.2), so this path defines a closure cycle γ in A. ■

Corollary 3.3 — Rotations Define Closure Cycles. *Any spatial rotation R(θ) defines a continuous admissible transformation. A rotation through 2π therefore defines a closed admissible path in A.*

Proof. Spatial rotations act continuously on the distinguishability structure of physical states and preserve admissibility (CCC). A rotation through 2π returns the observable configuration to its starting value. By Proposition 3.2, this defines a closed admissible path. ■

Interpretation. This is the key step. Rotations are not merely geometric operations on a separate spatial manifold: they are continuous transformations of the distinguishability structure of admissible states, and must therefore generate closure cycles in the admissible commitment space.

Proposition 3.4 — Closure Action Factors Through Homotopy. *Let R(θ) be a rotation path. Then the induced action on closure invariants depends only on the homotopy class of the induced path in A.*

Proof.

By Corollary 3.3, the rotation path induces a closure cycle in A.

By Theorem 3.1, closure invariants on A factor through the homotopy class of such cycles.

Therefore the action of rotations on closure invariants factors through

$\pi_1(SO(3))$. ■

Cumulative result. Combining §3.1–§3.2:

- Rotations induce admissible closure cycles (Proposition 3.2 and Corollary 3.3, derived from BC1, BC2, FIM, and CCC).
- Closure invariants factor through homotopy classes (Theorem 3.1, derived from BC1, BC2, and FIM).
- Therefore the action of rotations on the closure register factors through $\pi_1(SO(3))$ (Proposition 3.4).

The single remaining question is whether the induced mapping $\pi_1(SO(3)) \to \pi_1(A)$ is faithful — that is, whether the non-trivial homotopy class of SO(3) survives the passage to admissible commitment cycles, or collapses to the trivial class. We show in §3.3 that BC1 and CCC together force faithfulness, so this too is derived. Postulate L can then be reduced to its fully minimal form, given in §3.

3.3 Faithfulness of the Rotation–Closure Mapping

The faithfulness of the rotation-cycle / admissible-cycle correspondence — the last residual content previously assigned to Postulate L — is itself forced by BC1 and CCC. We state this as a theorem.

Theorem 3.5 — Faithfulness of the Rotation–Closure Mapping. *Let A be the admissible state space satisfying BC1, BC2, FIM, and CCC. Let spatial rotations induce closure cycles in A as in Proposition 3.2. Then the induced mapping*

$\pi_1(SO(3)) \to \pi_1(A)$

is faithful on admissible closure invariants.

Proof. Suppose, for contradiction, that the mapping is not faithful. Since $\pi_1(SO(3)) = \mathbb{Z}_2$, the only non-faithful possibility is that the non-trivial element is mapped to the trivial class in $\pi_1(A)$: the 2π rotation cycle is identified, at the closure level, with a contractible cycle.

The 2π rotation is topologically non-trivial in SO(3) — it is not continuously deformable to the identity within SO(3). It therefore represents a genuinely distinct admissible transformation history from any contractible loop.

Under the assumed non-faithful mapping, two non-equivalent admissible histories — the 2π rotation cycle and a contractible cycle — would be indistinguishable at the closure level. This is a loss of distinguishability between admissible histories, in violation of BC1. In particular, the identification of the 2π rotation with the trivial cycle removes a physically admissible loop without a corresponding inverse deformation, breaking the reversibility required by BC1.

Equivalently, under the non-faithful mapping, the closure register fails to encode the admissible distinction between the trivial and non-trivial rotation classes. The state is therefore not closure-complete in the sense of CCC.

Both consequences are forbidden. The assumption fails. The mapping must be faithful. ■

Corollary 3.6 — Necessity of Non-Trivial Representation. *The induced action of $\pi_1(SO(3))$ on the closure register must be non-trivial. The trivial representation is incompatible with closure completeness.*

Proof. Let $\rho: \pi_1(SO(3)) \to \mathrm{Sym}(\{+1, -1\})$ denote the induced action on the closure register. By Theorem 3.5, the mapping $\pi_1(SO(3)) \to \pi_1(A)$ is faithful, so the non-trivial element of $\pi_1(SO(3))$ corresponds to a non-trivial homotopy class in A. By Theorem 3.1 (closure-completeness clause, strengthened by the BC1/BC2 exhaustion requirement), closure invariants are non-trivial on non-trivial homotopy classes in A. Composing the two: the non-trivial element of $\pi_1(SO(3))$ acts non-trivially on the closure register, i.e. $\rho \neq \rho_0$. Since $\pi_1(SO(3)) = \mathbb{Z}_2$ admits only two homomorphisms into $\mathrm{Sym}(\{+1, -1\})$, $\rho = \rho_1$. ■

Interpretation. Thus the non-trivial representation is not a choice among representations, but the unique admissible encoding of topological distinguishability within closure-complete states.

Consequence. Faithfulness is now derived from BC1 and CCC, not assumed. The faithfulness clause that was the residual content of Postulate L in §3.2 is therefore itself a derived result, and Postulate L may be reduced to its fully minimal form — the near-definitional statement that spatial rotations are continuous admissible transformations of physical states. This form is given in §3. Furthermore, by Corollary 3.6 the non-trivial representation ρ_1 on the closure register is forced independently of the §4 development; §4.2 then identifies its physical content as the spinor 4π periodicity.

---

4. Topology of Closure: Why φ Must Invert Under 2π

By Theorem 3.1, closure invariants are homotopy invariants. When the admissible cycle is a spatial rotation, the relevant topology is that of SO(3). Since $\pi_1(SO(3)) = \mathbb{Z}_2$, the closure phase must distinguish exactly two rotation-cycle classes: trivial and non-trivial. The remainder of this section makes that distinction explicit and identifies the unique admissible primitive action.

4.1 The fundamental group of SO(3)

The rotation group SO(3) is doubly connected:

$$\pi_1(SO(3)) = \mathbb{Z}_2.$$

Loops in SO(3) fall into exactly two homotopy classes. A 2π rotation about any fixed axis represents the non-trivial class; a 4π rotation represents the trivial class and is contractible to the identity (the "Dirac belt trick").

4.2 Closure as a representation of $\pi_1(SO(3))$

Suppose $\Psi = (b, \varphi)$ is acted on by a continuous family of rotations, $R(\theta)$: SO(3) → Aut(state space), parameterised by a path γ from the identity to a target rotation $g \in SO(3)$. By Postulate L, rotations are continuous admissible transformations of physical states; by Corollary 3.3, the rotation path therefore defines an admissible closure cycle. By Theorem 3.1 and Proposition 3.4, the induced action on φ factors through $\pi_1(SO(3))$. By Theorem 3.5, this factorisation is faithful: the non-trivial homotopy class of SO(3) does not collapse to the trivial class in the admissible commitment space.

The closure phase is a two-valued register, $\varphi \in \{+1, -1\}$. The relevant action is therefore not an automorphism of $\mathbb{Z}_2$ as an abstract group, but a permutation action on the two possible closure states. Thus closure defines a homomorphism

$$\rho: \pi_1(SO(3)) \to \mathrm{Sym}(\{+1, -1\}) \cong \mathbb{Z}_2.$$

Since $\pi_1(SO(3)) \cong \mathbb{Z}_2$ and $\mathrm{Sym}(\{+1, -1\}) \cong \mathbb{Z}_2$, there are exactly two homomorphisms:

1. The **trivial** homomorphism ρ_0: every loop acts as identity on φ.
2. The **non-trivial** homomorphism ρ_1: the generator of $\pi_1(SO(3))$ acts by flip.

By Corollary 3.6, the trivial action ρ_0 is incompatible with closure completeness; the non-trivial homomorphism ρ_1 is the unique admissible primitive action. The $2\pi/4\pi$ content of ρ_1 is therefore:

$$\rho = \rho_1 \Longrightarrow R(2\pi): \varphi \mapsto -\varphi,\ R(4\pi): \varphi \mapsto \varphi.$$

This is the half-integer-spin transformation law, **derived** from Postulate B, Theorem 3.1 (and hence from BC1, BC2, FIM), Propositions 3.2–3.4 (and hence from CCC), Theorem 3.5 and Corollary 3.6 (faithfulness and ρ_1-uniqueness, derived from BC1 and CCC), and the now near-definitional Postulate L (rotations are continuous admissible transformations of physical states).

The exclusion of ρ_0 here applies only at the primitive level. Composite states built from multiple (b, φ) units can and do realise ρ_0 as a derived action, in which constituent closure phases pair and cancel under rotation; this is the integer-spin sector and is developed in §9. The ρ_0 action is therefore not absent from the framework — it appears as the *composite* action on tensor products of primitive (b, φ) units. Composite states realise integer-spin representations as a derived consequence of half-spinor composition (§9), rather than as a primitive sector in their own right. The framework is therefore not sector-agnostic at the primitive level: the half-spinor is the unique admissible primitive, and the integer-spin spectrum is generated by composition.

4.3 Action on the observable layer

Observable orientation is by definition a function of the endpoint $g \in SO(3)$ only. Hence

$$R(2\pi): b \mapsto b,\ R(4\pi): b \mapsto b.$$

4.4 Combined action

Combining §4.2 and §4.3:

$R(2\pi)$: $(b, \varphi) \mapsto (b, -\varphi)$,

$R(4\pi)$: $(b, \varphi) \mapsto (b, +\varphi)$.

Under the identification of φ with the global sign of a complex linear envelope (developed in §6), the discrete inversion on the closure register lifts to scalar multiplication by −1 on the complex state. A faithful linear realisation on a complex 2-dimensional state space (§6) then gives

$R(2\pi)\,\Psi = -\Psi$, $R(4\pi)\,\Psi = +\Psi$,

which is the defining transformation property of a spinor.

---

5. The Group Action: SO(3), SU(2), and Closure

Section 4 forces the closure phase, in the ρ_1 sector, to carry the non-trivial $\mathbb{Z}_2$ representation of $\pi_1(SO(3))$. Standard topology then identifies the group acting faithfully on the resulting (b, φ) structure.

5.1 The universal cover

The universal cover of SO(3) is SU(2), with covering map

$\pi: SU(2) \to SO(3)$, $\ker \pi = \mathbb{Z}_2 = \{+I, -I\}$.

Equivalently, SU(2) = Spin(3). Loops in SO(3) lift to paths in SU(2); a 2π loop in SO(3) lifts to a path connecting +I and −I, while a 4π loop lifts to a closed path.

5.2 Identification

The structure derived in §4 — a state space carrying both the orientation information of SO(3) and a $\mathbb{Z}_2$ register tracking the homotopy class of rotation paths — is, by construction, a representation of the universal cover SU(2). The kernel of π acts on Ψ by

$-I$: $(b, \varphi) \mapsto (b, -\varphi)$.

This is non-trivial in the closure layer but invisible to any quantity depending on b alone.

We may therefore restate the conventional double-cover relation in VERSF terms:

Group	What it tracks	VERSF interpretation
SO(3)	observable orientation (b)	endpoint of a rotation path
SU(2)	orientation + closure phase (b, φ)	full homotopy class of the path

The double cover is not a mathematical curiosity. It is the smallest group capable of representing both layers of an admissible state.

---

6. Recovery of the Standard Spinor — A Consistency Check

The two-layer object $\Psi = (b, \varphi)$ is discrete. The standard spinor is a complex two-component object carrying a faithful representation of all of SU(2), including the off-diagonal generators σ_x, σ_y that produce coherent superposition of distinguishability states. We must be careful about what the present construction does and does not establish.

What §4–§5 fix. The closure-phase argument fixes the action of $\ker \pi = \{\pm I\} \subset SU(2)$ on the closure-complete state — that is, the global sign behaviour under loops in SO(3). This is the topologically non-trivial content and is the subject of this paper.

What §4–§5 do not fix. The action of arbitrary SU(2) elements on superpositions of ($b = 0$, $b = 1$) — in particular the off-diagonal Pauli generators — is *not* derived here. It requires the VERSF derivation of complex Hilbert space from distinguishability structure, which is treated in a separate paper [Taylor, *Complex Hilbert Space from Distinguishability*, VERSF Programme].

What this section does. The remainder of §6 is a consistency check: we exhibit the standard $\mathbb{C}^2$ spinor as the minimal complex linear envelope carrying the closure-phase structure of §4 and verify that its $2\pi/4\pi$ behaviour matches §4.4 by construction. This is reconstruction at the level of the topologically essential action, not at the level of full SU(2) covariance.

The minimal faithful linear representation of SU(2) is the fundamental representation on $\mathbb{C}^2$, with basis states usually written $|\uparrow\rangle$, $|\downarrow\rangle$. Identifying:

- the basis label with the distinguishability bit b,
- the global sign with the closure phase φ,

a generic state is

$\psi = (\alpha, \beta)^T$, $\alpha, \beta \in \mathbb{C}$, $|\alpha|^2 + |\beta|^2 = 1$.

Under $R(2\pi) \in SU(2)$, the column inverts:

$R(2\pi)\, \psi = -\psi$,

reflecting closure-phase inversion. Under R(4π), ψ returns to itself. Quantities of the form ⟨ψ|A|ψ⟩ are insensitive to the closure phase, which explains why the global sign is *operationally* unobservable in single-state expectation values while remaining *structurally* essential — interference between branches with opposite closure phase remains accessible (e.g. neutron interferometry, Werner et al. 1975; Rauch et al. 1975), as expected.

The continuous Hilbert-space description is therefore not in tension with the discrete (b, φ) ontology. The closure-phase content is fully captured; the rest of the SU(2) structure follows from VERSF distinguishability complex-linearity, established elsewhere.

---

7. Fermionic Statistics: A Conjectural Extension

The argument above derives the spin-½ transformation law. It does not, by itself, derive Fermi–Dirac statistics. We outline a conjectural route and label it as such.

Conjecture (closure saturation). A closure-complete state occupies a unique configuration of the closure layer. Two distinct facts cannot share an identical closure configuration without violating the admissibility constraint that distinguishes them.

If this conjecture holds, then states with identical (b, φ) cannot coexist in the same admissibility class, yielding antisymmetric occupation. This is structurally suggestive of the Pauli principle but is **not** a substitute for the spin-statistics theorem. A full derivation requires:

1. a Lorentz-covariant extension of the closure-phase analysis (proto-time and emergent Lorentz invariance — addressed elsewhere in the VERSF programme);
2. a careful treatment of the exchange operation as a continuous path in the configuration space of identical particles, whose fundamental group provides the analogue of §4.2 in the multi-particle setting.

These are open problems within the present framework, not solved ones.

---

8. Scope, Status, and Limitations

To be explicit about what has and has not been established:

Derived (modulo stated postulates):

- Theorem 3.1 (closure invariants factor through homotopy class and are non-trivial whenever the homotopy class is non-trivial), proved in §3.1 from BC1, BC2, and FIM.

- Propositions 3.2, 3.3, and 3.4 (continuous admissible transformations induce closure cycles; rotations are such transformations; the action of rotations on closure invariants factors through $\pi_1(SO(3))$), proved in §3.2 from BC1, BC2, FIM, and CCC.
- Theorem 3.5 (faithfulness of the rotation–closure mapping), proved in §3.3 from BC1 and CCC. The faithfulness clause that previously constituted the residual content of Postulate L is now itself a derived result.
- Corollary 3.6 (ρ_1-uniqueness): the trivial representation ρ_0 is incompatible with closure completeness; the non-trivial representation ρ_1 is forced (§3.3, derived from Theorems 3.1 and 3.5).
- The reduction of Postulate L to its fully minimal, near-definitional form: spatial rotations are continuous admissible transformations of physical states (§3). Everything previously attributed to L beyond this is derived.
- The half-integer-spin transformation law itself, derived in §4.2 by combining Corollary 3.6 (ρ_1 uniqueness) with the topology of SO(3). The trivial action ρ_0 is excluded at the primitive level by Corollary 3.6.
- The 4π periodicity of the primitive admissible state (§4.2).
- The identification of the closure-complete transformation group with SU(2) (§5).
- The composition rule generating the full integer/half-integer spectrum from tensor products of the primitive (b, φ) unit (§9). The integer-spin sector arises as the derived ρ_0 action on composites in which constituent closure phases pair and cancel.

Established as consistency, not derivation:

- The recovery of the full $\mathbb{C}^2$ spinor with arbitrary SU(2) action (§6). The closure-phase argument fixes the action of $\ker \pi = \{\pm I\}$; the action of off-diagonal SU(2) generators on superpositions of $(b = 0, b = 1)$ requires the separate VERSF derivation of complex Hilbert space from distinguishability structure.

Structurally demonstrated, not derived:

- The interpretation of the SU(2) global sign as a closure phase rather than an unobservable artefact.
- The two-layer (b, φ) decomposition itself, which is motivated by VERSF principles but is here taken as a working ansatz for the smallest closure-complete state.

Conjectural:

- The connection between closure saturation and Pauli exclusion (§7).

Outside the scope of this paper:

- Spin-statistics, which requires Lorentz covariance and exchange-path topology.
- Spin in arbitrary spatial dimension $d \neq 3$, where $\pi_1(SO(d))$ differs (notably $\pi_1(SO(2)) = \mathbb{Z}$, which suggests the same machinery would predict anyonic structure in 2D — a feature, not a limitation, but one this paper does not develop).

9. Discussion

The construction has three notable features.

1. Topology does the work. The 4π periodicity of spinors is not a quirk of representation theory but a direct consequence of $\pi_1(SO(3)) = \mathbb{Z}_2$ acting on a closure register that is, by independent VERSF reasoning, also $\mathbb{Z}_2$. The match between these two $\mathbb{Z}_2$'s is what makes spin-½ structurally available.

2. The double cover acquires a physical reading. SO(3) tracks where you end up; SU(2) tracks how you got there, modulo continuous deformation. In a framework where physical states are commitment cycles rather than instantaneous configurations, the path-sensitive group is not an embellishment but the correct one.

3. The "unobservable" global sign is reinterpreted. It is operationally unobservable in single-state expectation values, but it is structurally necessary, and it becomes physically active in interference experiments. The framework therefore predicts no new phenomenology for spin-½ in isolation, but reframes existing phenomenology (notably the 4π neutron-rotation experiments) as direct evidence of the closure-phase layer.

The construction also shows explicitly how higher spins arise. Consider a composite state built from two (b, φ) units, with combined closure phase

$$\varphi\_total = \varphi_1 \cdot \varphi_2.$$

Under $R(2\pi)$, each factor inverts, so

$$\varphi\_total \mapsto (-\varphi_1)(-\varphi_2) = +\varphi\_total.$$

The composite therefore lies in the ρ_0 sector — it transforms as an integer-spin object. Its tensor product space is four-dimensional and decomposes under SO(3) into the standard singlet (spin-0) plus triplet (spin-1), exactly as in conventional angular-momentum addition. More generally, n-fold tensor products of the primitive (b, φ) unit yield states of integer spin if n is even and half-integer spin if n is odd. The full spin spectrum is therefore generated by composition from a single primitive — the closure-spinor unit derived in §4 — without need for further postulates. Both ρ_0 and ρ_1 sectors play a role: ρ_1 supplies the primitive, ρ_0 describes the composite states that emerge from even-fold combination.

10. Conclusion

Half-integer spin emerges in VERSF as a structural consequence of closure, distinguishability conservation, and finite information capacity. BC1 requires commitment histories to remain reversibly traceable; FIM makes state comparison path-dependent through transport; BC2 removes redundant microscopic path detail; and CCC requires admissible transformations to preserve closure. Closure invariants therefore reduce to homotopy invariants (Theorem 3.1), and continuous admissible transformations — including spatial rotations — induce admissible closure cycles (Propositions 3.2 and 3.3). The action of rotations on the closure register accordingly factors through $\pi_1(SO(3))$ (Proposition 3.4), and BC1 together with CCC forces this factorisation to be faithful (Theorem 3.5) and forces the action to be the non-trivial representation ρ_1 (Corollary 3.6). For three-dimensional space, the relevant invariant is $\pi_1(SO(3)) = \mathbb{Z}_2$. The closure-complete state must therefore carry a binary phase distinguishing the trivial and non-trivial rotation cycles, with rotation acting by inversion. Combined with Postulate B (commitments are bits) and Postulate L — now reduced to the near-definitional content that spatial rotations are continuous admissible transformations of physical states — this fixes the half-spinor as the unique admissible primitive state and gives $\varphi \mapsto -\varphi$ under 2π directly.

From the ρ_1 primitive, the 4π periodicity follows; the universal cover SU(2) is forced as the closure-complete transformation group; and the standard $\mathbb{C}^2$ spinor is recovered as the minimal complex linear envelope of (b, φ), once the VERSF derivation of complex Hilbert space (treated separately) is invoked. Composition of n primitive units yields integer spin for even n and half-integer spin for odd n: the integer-spin spectrum arises as the derived ρ_0 action on composites in which constituent closure phases have paired and cancelled. The full spin spectrum is therefore generated from a single primitive, the half-spinor, with no double-valued representations postulated.

The result is modest in what it changes (no new phenomenology for free spinors) but substantive in what it explains: the necessity, rather than the mere availability, of the spinor structure; the topological origin of the 4π periodicity; the demotion of path sensitivity from postulate to theorem proved from BC1, BC2, and FIM; and a structural unification of integer and half-integer spin within a single primitive. In 2D, $\pi_1(SO(2)) = \mathbb{Z}$ would predict a closure register with integer-valued (rather than $\mathbb{Z}_2$-valued) topological charge, naturally accommodating anyonic statistics; the framework therefore offers a structural account of why anyons are 2D-specific without requiring additional postulates. It also marks out a clear research programme: extending the analysis to multi-particle exchange topology, to Lorentz-covariant settings, and to spatial dimensions where $\pi_1(SO(d))$ admits richer structure.

References

Primary topology and spinor references:

- Cartan, É. (1913). *Les groupes projectifs qui ne laissent invariante aucune multiplicité plane.* Bulletin de la Société Mathématique de France 41, 53–96. [Original definition of spinors as objects of a double-valued representation; historical anchor for the topological framing.]
- Naber, G. L. (2011). *Topology, Geometry and Gauge Fields: Foundations* (2nd ed.). Springer. [Standard reference for $\pi_1(SO(3)) = \mathbb{Z}_2$ and the universal cover SU(2) → SO(3).]
- Penrose, R., & Rindler, W. (1984). *Spinors and Space-Time, Vol. 1: Two-Spinor Calculus and Relativistic Fields.* Cambridge University Press. [Geometric and topological treatment of spinors.]
- Sakurai, J. J., & Napolitano, J. (2020). *Modern Quantum Mechanics* (3rd ed.). Cambridge University Press. [Standard formalism of spin-½ and SU(2) representation theory.]

Empirical anchor:

- Werner, S. A., Colella, R., Overhauser, A. W., & Eagen, C. F. (1975). *Observation of the phase shift of a neutron due to precession in a magnetic field.* Phys. Rev. Lett. 35, 1053.
- Rauch, H., Zeilinger, A., Badurek, G., Wilfing, A., Bauspiess, W., & Bonse, U. (1975). *Verification of coherent spinor rotation of fermions.* Phys. Lett. A 54, 425.

VERSF programme (cross-references):

- Taylor, K. *Complex Hilbert Space from Distinguishability.* VERSF Programme [Required for the §6 reconstruction of arbitrary SU(2) action on (b, φ) superpositions.]
- Taylor, K. *VERSF Constraint and Lagrangian Representations.* VERSF Programme.
- Taylor, K. *The Commitment Condition Criterion (CCC).* VERSF Programme.
- Taylor, K. *No-Go Theorem for Non-Simplicial Relational Substrates: K = 7 as Forced.* VERSF Programme.

Appendix A: The Dirac Belt Trick

The topological fact $\pi_1(SO(3)) = \mathbb{Z}_2$, on which §4 turns, is unfamiliar to many readers. The classical demonstration is the Dirac belt trick (also called the plate trick or the cup trick). It is worth a brief pause because the entire spinor structure hangs on it.

The belt. Hold one end of a belt fixed and rotate the other end through 2π about a fixed axis, keeping the free end's orientation unchanged at the start and finish. The belt acquires a twist that cannot be removed by any continuous motion of the free end. The twist is the topological signature of the non-trivial element of $\pi_1(SO(3))$.

The 4π move. Now rotate the free end through a further 2π in the same direction, for a total of 4π. The belt now appears to have two twists. Remarkably, both can be removed by passing the belt around the free end, *without rotating either end*. The 4π loop is contractible to the identity; the 2π loop is not.

The cup or plate version. Hold a cup of water in one hand. By passing the arm under and over while keeping the cup upright, one can rotate the cup through 4π and return the arm to its starting configuration; a single 2π rotation leaves the arm in a knotted configuration that requires the second 2π to unwind. The arm is, in effect, a physical realisation of the closure register: it sees the difference between a 2π and a 4π rotation.

The connection to §4. The closure register φ is sensitive to exactly the same topological distinction. A 2π rotation places it in the non-trivial class (one twist, $\varphi \to -\varphi$); a 4π rotation places it in the trivial class (no net twist, $\varphi \to +\varphi$). The 4π periodicity of spinors is therefore not a quirk of complex linear algebra: it is the same fact that makes the belt trick work, expressed at the level of the closure phase.