Closure Algebra and the Origin of Spin, Gauge, and Gravity in the VERSF Framework

A Representation-Theoretic Unification of Relational Closure Dynamics

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Abstract

Modern physics rests on three pillars that, in the standard textbook story, are introduced as separate ideas: *matter* (described by quantum states), *forces* (described by gauge fields like electromagnetism and the strong and weak nuclear forces), and *gravity* (described by Einstein's equation for the geometry of spacetime). Each pillar works beautifully on its own, but each is normally postulated as an independent starting point — three different "first ideas" that physicists have learned to fit together, but never derived from a single underlying source.

This paper proposes — and partly proves — that all three pillars are structurally derived consequences of a single organizing principle. The principle is the *closure ontology* developed in earlier papers in the VERSF programme: physical reality is built up from facts that distinguish one state of affairs from another, and the structure of physical fields is determined by the geometry of those distinctions. From this single starting point, the paper constructs a mathematical object called the *closure algebra* with three natural layers, each corresponding to one of the three pillars:

- The **first layer** asks how physical states can be identified at all. The answer forces them to be quantum states with the spin-½ structure of fermionic matter (electrons, quarks, neutrinos).
- The **second layer** asks how to compare states at neighbouring points in space. The answer forces gauge fields with spin-1 structure — the kind of fields that mediate electromagnetism and the other forces.
- The **third layer** asks how the substrate responds when distinguishing events accumulate and flow. The answer forces a symmetric tensor field with spin-2 structure, whose dynamics take the form of Einstein's equation for gravity.

A central technical contribution is a **No-Alternative Theorem** (§7.2): under stated structural assumptions, no fourth layer can exist. Any operation that might seem to require a new layer turns out to be expressible in terms of derivatives or combinations of objects already present in the three layers. The three pillars of physics are therefore *complete* — under the assumptions —

and there is no room for an unsuspected fourth fundamental field type. The most consequential of the structural assumptions is what the paper calls the *trichotomy*: the claim that any admissible operation on closure states must do one of three things — identify a state, compare neighbouring states, or register accumulated response — and that no fourth role is possible. The trichotomy follows naturally from the closure ontology, but a fully primitive derivation of it is identified in the paper as the leading open structural problem. Until that derivation is supplied, the No-Alternative Theorem rests on the trichotomy as a structural commitment, not as a derived fact.

The paper does not claim mathematical novelty in any individual sector. The equations recovered (Dirac, Yang–Mills, Einstein) are the standard equations of physics. The contribution is the *consolidation*: where standard physics requires three separate postulational starting points, the closure-algebra picture requires one. This is a substantial reduction in the postulate count of physics, even when the equations themselves remain familiar — a referee would correctly say we have not changed the mathematics, but we have changed where it comes from.

Several substantial questions remain open. How does the smooth-manifold setting used throughout the paper emerge from the underlying discrete substrate? How do the non-Abelian gauge groups of the Standard Model arise from the Abelian case derived here? Can the structural assumptions of the No-Alternative Theorem themselves be derived from yet more primitive principles? These are flagged as the central open problems of the broader VERSF programme.

Table of Contents

Abstract

We propose a unifying organizational principle for the origin of physical fields within the **Void Energy-Regulated Space Framework (VERSF)**: all fundamental field structures arise as representations of constraints imposed by relational closure under distinguishability-preserving symmetry. Building on prior reconstruction results — in which (i) complex Hilbert space $\mathscr{H}$ and unitary dynamics are derived from distinguishability geometry [VERSF–CHS], (ii) half-integer spin emerges from the projective action of spatial isometries [VERSF–SPIN], and (iii) Abelian gauge fields arise as connections required for closure-consistent local comparison of ray representatives [VERSF–GAUGE] — we introduce a **closure algebra** $\mathfrak{c}$ with a graded differential structure whose admissible representations contain these structures simultaneously. The algebra has three graded layers:

$\mathfrak{c} = \mathfrak{c}^{(0)} \oplus \mathfrak{c}^{(1)} \oplus \mathfrak{c}^{(2)}$,

corresponding to *state closure*, *transport closure*, and *back-reaction closure*. The $k = 2$ layer carries both an antisymmetric sector (housing the gauge curvature F generated by the covariant differential D acting on $\mathfrak{c}^{(1)}$) and a symmetric sector (housing the metric response $h_{\mu\nu}$ of closure back-reaction). Closure-density conservation $\nabla^{\mu} T^{cl}_{\mu\nu} = 0$ combined with the imported Lovelock-type uniqueness theorem then forces a field equation of Einstein form

$G_{\mu\nu} + \Lambda_{void} g_{\mu\nu} = \kappa_{cl} T^{cl}_{\mu\nu}$.

The contribution is **interpretive consolidation under a single algebraic object**: technical content at each layer overlaps with standard treatments (projective representation theory, Abelian gauge theory, Lovelock-type uniqueness for the Einstein tensor), but all three sectors are now organized by a single graded algebra of closure operations. Under five stated structural assumptions (ray admissibility, local comparability, closure conservation, locality in the smooth-manifold continuum limit, and minimal non-redundancy), we prove a **No-Alternative Theorem** (§7.2): no independent grade $k > 2$ exists, the {½, 1, 2} hierarchy is closed, and any apparent $k > 2$ operation reduces to derivatives or polynomial composites of lower-grade objects via the cohomological closures of [VERSF–GAUGE] §7 (Henneaux–Teitelboim 1992) and Anderson 1989. Three residual structural commitments — the trichotomy of admissible roles, the locality assumption's dependence on the unconstructed continuum limit, and postulate (M) at layer 0 — are owned in §7.3 and identified as open problems. The strongest unconditional version of the theorem, with these residual commitments themselves derived, remains open. We are explicit throughout about what is *derived*, what is *imported* (the differential-geometric apparatus on a smooth metric manifold), what is *postulated* (the minimal-nontrivial-representation selection principle at layer 0), and what is *deferred* (the Lorentz lift, the non-Abelian extension, coupling constants, dynamics, and the continuum limit from the $K = 7$ simplicial substrate to smooth manifold structure).

---

1. Introduction

Modern physics rests on three structural layers traditionally introduced as independent postulates:

1. **Quantum states** — vectors in a complex Hilbert space $\mathscr{H}$, identified up to global phase.
2. **Gauge fields** — connections on principal bundles, introduced by demanding local symmetry.
3. **Gravitational fields** — a metric g_μν on spacetime, governed by Einstein's equation.

Each is mathematically clean. None of them, in standard treatments, is *derived* — they are postulated alongside one another and connected by additional couplings.

The central question of this paper is:

Do these three structures arise from a single underlying organization?

The VERSF programme has previously addressed parts of this question in isolation. Hilbert space and complex amplitudes were derived from distinguishability geometry in [VERSF–CHS]. Half-integer spin was derived from the projective action of SO(3) on ray space in [VERSF–SPIN]. The Abelian gauge connection was derived from the local comparison of ray representatives in [VERSF–GAUGE]. What was missing — and what this paper supplies — is a **single algebraic object** that contains all three structures as its graded layers.

The object is the **closure algebra** $\mathfrak{c}$: a graded differential algebra of operations on closure-state ray fields over the distinguishability substrate $\mathcal{S}$. Its three graded components $\mathfrak{c}^{(0)}$, $\mathfrak{c}^{(1)}$, $\mathfrak{c}^{(2)}$ correspond to operations of three structural types — *state*, *transport*, and *back-reaction* — and their admissible representations contain spins ½, 1, and 2 respectively.

1.1 The contribution: interpretive consolidation, not new mathematics

The technical content of the spin-½ layer overlaps with standard projective representation theory (Wigner–Bargmann). The technical content of the spin-1 layer overlaps with standard Abelian gauge theory (Yang–Mills). The technical content of the spin-2 layer overlaps with the Lovelock-type uniqueness arguments selecting the Einstein tensor as the unique divergence-free symmetric-tensor curvature on a metric manifold. None of this technical mathematics is new.

What is new is the **algebraic organization**. Three independent structural postulates of standard physics — Hilbert space + Born rule + spinor representation; the gauge principle; the equivalence principle + general covariance — are replaced by a single graded algebraic structure, with the three sectors appearing as its three forced grades. Within a reconstruction programme aiming to derive physics from a small ontological base, this consolidation reduces the postulate count of physics by combining three apparently independent structural inputs into one. We use this framing once here and refer back to it where relevant; we do not repeat it as a defense at every section.

1.2 Notation: substrate vs. cosmological constant

We denote the distinguishability substrate by $\mathcal{S}$ throughout this paper, reserving the symbol Λ (or Λ_void) for the cosmological-constant-like term that appears in the closure back-reaction equation in §8. This differs from the convention in earlier VERSF papers, where Λ was used for the substrate; we make the change to avoid collision with the standard physics symbol for the cosmological constant.

1.3 Structure of this paper

§2 fixes the input/output ledger. §3 introduces the closure bundle and the closure algebra $\mathfrak{c}$, with a separate §3.5 developing a structural argument for the three-grade decomposition. §§4–6 derive the three layers — spin-½, spin-1, spin-2 — as forced or contained representations of $\mathfrak{c}^{(0)}$, $\mathfrak{c}^{(1)}$, $\mathfrak{c}^{(2)}$, summarising results from [VERSF–SPIN] and [VERSF–GAUGE] for layers 0 and 1 (with self-contained statements of the imported facts) and giving the new derivation for layer 2. §6 includes a side-by-side ledger of where each closure-structural input replaces a corresponding standard motivation for the Einstein equation. §7 states the **Hierarchy Theorem** as a consolidation of §§4–6 and the **No-Alternative Theorem** establishing closure of the algebra at grade 2 under five stated structural assumptions, with §7.3 owning the residual commitments that prevent the unconditional version. §8 derives the Einstein-form field equation. §9 discusses relations to standard treatments and other reconstruction programmes. §§10–11 state open problems and what is not claimed. §12 concludes.

---

2. What Is Assumed and What Is Derived

To avoid importing physical structure through the back door, we state the input/output ledger explicitly.

2.1 Imported from prior VERSF reconstruction

Imported result	Source
Distinguishability substrate $\mathcal{S}$ with metric d preserving an interference-compatible algebra	VERSF–CHS
$\mathbb{C}$-valued amplitudes; complex Hilbert space $\mathscr{H}$	VERSF–CHS
Reversible dynamics are unitary on $\mathscr{H}$	VERSF–CHS
Relational closure ontology (states are closure patterns)	VERSF–FSN
Physical states are rays in $\mathbb{P}(\mathscr{H})$; global phase is unobservable	VERSF–FSN, VERSF–SPIN §5
Continuous spatial isometry group is SO(3); states transform projectively	VERSF–SPIN
Three-dimensional simplicial K = 7 closure substrate	VERSF–KSEVEN

Imported result	Source
Closure events produce conserved distinguishability density	VERSF–CCB, VERSF–FSN

2.2 Imported from standard mathematics

The derivation in §6 imports the apparatus of differential geometry on a smooth metric manifold — covariant derivatives, the Riemann tensor, the contracted Bianchi identity $\nabla^\mu G_{\mu\nu} = 0$ — and the standard **Lovelock-type uniqueness theorem**: among local symmetric rank-2 tensors built from a metric $g_{\mu\nu}$ and its first and second derivatives, with vanishing covariant divergence, the unique family up to overall scale and an additive constant is

$$G_{\mu\nu} + \Lambda_{\text{void}}\, g_{\mu\nu},$$

where $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} R\, g_{\mu\nu}$ and Λ_{void} is a constant (Lovelock 1971).

These imports are not derived from VERSF principles. The closure-algebra contribution at the spin-2 level is to *force the input shape* of the Lovelock argument from a single ontological source (see §6 and the ledger in §6.8), not to re-derive the differential-geometric machinery itself.

2.3 Continuum-limit assumption

A serious gap in the present formulation is the bridge between the discrete substrate of [VERSF–KSEVEN] (a $K = 7$ simplicial 2-complex with discrete closure events) and the smooth differential-geometric apparatus invoked in §§5–8. The smooth-manifold setting is taken as an emergent description that the discrete closure dynamics must reproduce in an appropriate scaling limit. We do not construct that limit here; it is the second residual commitment of the No-Alternative Theorem and the canonical statement of the gap is at §10.5(ii). A reader concerned with the emergent-vs-fundamental status of the smooth structure used here is correct to note the gap; we identify it as one of the central open problems of the programme rather than claim it is solved.

2.4 Derived in this paper

Derived result	Section
Closure bundle $\mathcal{E} \to \mathcal{S}$ as the natural object on closure-state ray fields	§3
Closure algebra $\mathfrak{c} = \mathfrak{c}^{(0)} \oplus \mathfrak{c}^{(1)} \oplus \mathfrak{c}^{(2)}$ as graded differential algebra	§3
Layer 0 (state closure): contains projective SO(3) representations including $j = \frac{1}{2}$	§4
Layer 1 (transport closure): minimal admissible representation lies in $j = 1$	§5
Layer 2 (back-reaction closure): admissible response contains $j = 2$ sector	§6
Closure source tensor $T^{\text{cl}}_{\mu\nu}$ must be conserved rank-2	§6
Hierarchy Theorem: $\mathfrak{c}$ contains the $\{\frac{1}{2}, 1, 2\}$ spin sectors as its three graded layers	§7.1

Derived result	Section
No-Alternative Theorem: under (A1)–(A5), $\mathfrak{c}$ is closed at grade 2 (no independent k > 2)	§7.2
Einstein-form field equation as minimal closure back-reaction	§8

2.5 Postulated (selection principles)

The argument as it stands requires one selection principle, listed here for transparency.

Postulate	Where invoked
(M) The minimal nontrivial admissible representation is the one realized in physics	§4 (layer 0); §6 (layer 2 sub-cases)

Discussion. Postulate (M) is doing real work at layer 0 and at parts of layer 2 (§6.6); it does *not* enter at layer 1, where the connection is uniquely determined up to scale by the five-step argument of [VERSF–GAUGE] §5.4. We discuss the status of (M) in §4.4 and §10.7. We adopt it explicitly rather than smuggle it in via "minimal nontrivial" language, since the three layers are doing structurally different work.

2.6 Deferred to companion papers

Deferred question	Status
Lorentz-spin assignment of all three layers	VERSF–PROTO
Non-Abelian extension of the spin-1 layer	Open problems §10.1
Specific gauge groups (SU(3) × SU(2) × U(1))	Open problems §10.1
Coupling constants κ_cl, gauge couplings	Open problems §10.2
Full nonlinear gravitational dynamics	VERSF–GRAV-DYN
Yang–Mills + Einstein–Hilbert action principle	VERSF–CCB
Continuum limit from K = 7 simplicial substrate	Open problems §10.5(ii) (canonical), §10.6 (pointer); also VERSF–PROTO and VERSF–CCB
Unconditional No-Alternative Theorem (without residual commitments of §7.3)	Open problems §10.5
Quantum-gravity sector	Future work
Mass generation / Higgs mechanism	Future work

---

3. The Closure Bundle and the Closure Algebra

3.1 The closure bundle

Let $\mathcal{S}$ be the distinguishability substrate and $\mathcal{H}$ the complex Hilbert space of [VERSF–CHS]. Closure states at point $x \in \mathcal{S}$ are rays $[\psi(x)] = \{ e^{i\theta} |\psi(x)\rangle : \theta \in \mathbb{R} \} \subset \mathcal{H}$. Vector representatives $|\psi(x)\rangle$ are obtained by an arbitrary choice of phase at each point, and the relational closure ontology imposes no relation between phase choices at different points.

Definition 3.1 (Closure Bundle). *The closure bundle $\mathcal{E} \to \mathcal{S}$ is the bundle whose fibre over $x \in \mathcal{S}$ is the ray space $\mathbb{P}(\mathcal{H}_x)$ and whose local sections are vector representatives $|\psi(x)\rangle \in \mathcal{H}_x$ defined modulo local U(1) re-phasings.*

Structurally, $\mathcal{E}$ is a U(1) principal bundle in the Abelian case derived in [VERSF–GAUGE]; larger structure groups arise when closure states carry internal indices (see §10.1). What VERSF adds to the standard fibre-bundle picture is a *physical motivation* for why $\mathcal{E}$ exists: it is forced by the relational closure ontology — physical states are rays, vector representatives are auxiliary, and any local comparison structure must respect this ambiguity.

3.2 The closure algebra: definition

We now define the algebraic object that organises operations on $\mathcal{E}$. The definition fixes the structural problem flagged in earlier internal review: the curvature $F = dA$ generated by acting with the covariant differential on layer 1 is antisymmetric, while the gravitational response field $h_{\mu\nu}$ is symmetric; both must therefore live in layer 2.

Definition 3.2 (Closure Algebra). *The closure algebra $\mathfrak{c}$ over $\mathcal{S}$ is the graded $\mathbb{R}$-algebra*

$$\mathfrak{c} = \mathfrak{c}^{(0)} \oplus \mathfrak{c}^{(1)} \oplus \mathfrak{c}^{(2)},$$

with grades:

- $\mathfrak{c}^{(0)} = \Omega^0(\mathcal{S}) \otimes \mathrm{End\_ray}(\mathcal{H}_x)$ — *operations on closure states (state closure), where $\mathrm{End\_ray}(\mathcal{H}_x)$ denotes those endomorphisms of $\mathcal{H}_x$ that descend to well-defined operations on the ray space $\mathbb{P}(\mathcal{H}_x)$;*
- $\mathfrak{c}^{(1)} = \Omega^1(\mathcal{S}; i\mathbb{R})$ — *one-form-valued operations on closure-state comparison (transport closure), in the Abelian case;*
- $\mathfrak{c}^{(2)} = \Lambda^2(\Omega^1(\mathcal{S})) \oplus \mathrm{Sym}^2(\Omega^1(\mathcal{S}))$ — *two-form-valued operations on closure response (back-reaction closure), with antisymmetric and symmetric sectors hosting respectively the gauge curvature and the metric response,*

equipped with the exterior derivative $d : \Omega^k \to \Omega^{k+1}$ on the form sectors and the natural product structure on $\mathrm{End\_ray}(\mathcal{H}_x)$.

The $k = 0$ grade is restricted to ray-respecting endomorphisms because the physically meaningful operations at this layer act on ray space $\mathbb{P}(\mathcal{H}_x)$, not on $\mathcal{H}_x$ itself; this tightens the definition relative to a naïve $\mathrm{End}(\mathcal{H}_x)$. The $k = 2$ grade is decomposed into antisymmetric and symmetric sectors because both are forced: $F = dA$ is antisymmetric (Λ^2), $h_{\mu\nu}$ is symmetric (Sym^2), and the algebra must contain both for the unification to hold.

3.3 The covariant differential

The graded structure carries a covariant differential D extending the exterior derivative d. On the $k = 0$ sector, D is the covariant derivative $D_\mu = \partial_\mu + iA_\mu$ derived in [VERSF–GAUGE]. On the $k = 1$ sector, the curvature

$$F = dA + A \wedge A \ (= dA \text{ in the Abelian case})$$

is an element of the antisymmetric sector $\Lambda^2(\Omega^1(\mathcal{S})) \subset \mathfrak{c}^{(2)}$. The symmetric sector $\mathrm{Sym}^2(\Omega^1(\mathcal{S})) \subset \mathfrak{c}^{(2)}$ houses the metric perturbation $h_{\mu\nu}$ of layer 2; its gauge-like redundancy and curvature structure are developed in §6. The algebra $\mathfrak{c}$ is therefore graded-differential, with each layer's "field strength" living in the layer above it, and with the $k = 2$ layer carrying *both* the gauge curvature (antisymmetric sector) and the metric response (symmetric sector).

3.4 Closure invariance

Throughout, *closure-invariant* means invariant under the closure-equivalence transformations of all three layers:

$$|\psi(x)\rangle \longmapsto e^{i\theta(x)}|\psi(x)\rangle,$$

$$A_\mu(x) \longrightarrow A_\mu(x) - \partial_\mu\theta(x),$$

$$h_{\mu\nu}(x) \longrightarrow h_{\mu\nu}(x) + \partial_\mu\xi_\nu(x) + \partial_\nu\xi_\mu(x),$$

for any smooth real $\theta : \mathcal{S} \longrightarrow \mathbb{R}$ and any smooth vector field ξ^μ on $\mathcal{S}$. The first transformation is the U(1) phase ambiguity of layer 0; the second is the gauge redundancy of layer 1; the third is the diffeomorphism-like redundancy of layer 2. Each is a structural feature of the corresponding layer's representation — an ambiguity in how we represent the physical content, not a postulated symmetry.

3.5 Why three grades, and why these specific tensor types

A skeptical reading will object that the three-grade decomposition is reverse-engineered to match the $\{½, 1, 2\}$ answer. We address that worry directly here, while being honest about what is and is not established by the argument.

The structural problem at each grade. Closure operations on $\boldsymbol{\mathcal{E}}$ fall into three structurally distinct types:

1. **Within a single fibre ($k = 0$).** Operations that compare or combine closure states at a single substrate point. These act on rays in $\mathbb{P}(\mathcal{H}_x)$ and require no differential structure on $\mathcal{S}$; they are point-local in x. The natural form-degree is zero (scalar functions over $\mathcal{S}$).

2. **Between neighbouring fibres ($k = 1$).** Operations that compare closure states at infinitesimally separated points x and $x + dx$. Such a comparison must contract with dx^μ to produce an x-independent quantity, which forces a one-form ($k = 1$) index structure on $\mathcal{S}$.
3. **Bilinear in differential displacements ($k = 2$).** Operations that register the *response* of the substrate to closure flow — either as the path-dependence of transport (the antisymmetric commutator $[D_\mu, D_\nu]$, generating curvature) or as the response of the substrate metric to closure-density flow (the symmetric coupling $h_{\mu\nu} T^{\mu\nu}_{cl}$). Both involve two differential displacements simultaneously, hence two-form structure ($k = 2$).

Why no $c^{(3)}$. A $k = 3$ grade would correspond to operations involving three independent differential displacements. We claim these decompose into iterated applications of the existing grades. The argument has two parts. *First*, the cohomological theorem of [VERSF–GAUGE] §7 (Henneaux–Teitelboim 1992, Ch. 12) establishes that all local closure-invariants built from A_μ are polynomial in $F_{\mu\nu}$ and its derivatives — so higher-form invariants in the gauge sector are derivatives of layer-2 elements, not independent degrees of freedom. *Second*, the corresponding statement for the symmetric sector of $c^{(2)}$ is that all local diffeomorphism-invariants of $h_{\mu\nu}$ are polynomial in the Riemann tensor and its covariant derivatives (Anderson 1989; this is the symmetric analogue of the same cohomological structure). Higher-grade closure operations therefore reduce to derivatives of existing-grade quantities and do not constitute an independent fourth layer.

Status of this argument. The argument is structural rather than fully formal at this point in the paper. The structural commitments motivating Definition 3.2 are made theorem-rigorous in §7.2 (the No-Alternative Theorem) under five stated assumptions, with residual commitments owned in §7.3. What §3.5 does is motivate the form-degree choices in Definition 3.2 in a way consistent with the cohomological closure results imported in §2.2 and the [VERSF–GAUGE] §7 result; what §7.2 does is establish that these choices are forced under stated assumptions on admissible closure operations. The two sections are complementary: §3.5 is the structural narrative, §7.2 is the theorem.

We therefore present Definition 3.2 as a **proposal** justified by the structural narrative above. The proposal is consistent with the cohomological closures imported in §2.2 and is made theorem-rigorous, under stated assumptions, in §7.2.

A logical dependence to flag. The form-degree organization presupposes a smooth-manifold structure on $\mathcal{S}$; on the underlying $K = 7$ simplicial substrate, the natural analogue is simplicial cochain degree, and the relationship between simplicial and smooth gradings is itself a continuum-limit question (§10.5(ii)). The structural argument of this section is therefore logically downstream of the continuum-limit construction: it shows that *if* the smooth-manifold setting is the correct continuum description, then form-degree organizes admissible local closure operations as described. The continuum-limit gap is not closed by §3.5; it is presupposed by it.

---

4. Layer 0: State Closure and Spin-½

We summarise the spin-½ derivation of [VERSF–SPIN]; full details are given there.

4.1 The state-closure problem

Closure states at a single substrate point are rays $[\psi] \in \mathbb{P}(\mathscr{H})$. Spatial symmetries $g \in SO(3)$ act on $\mathscr{H}$ by unitary operators $U(g)$, but the action on rays is only *projective* — the relation $U(g_1) U(g_2) = U(g_1 g_2)$ holds only up to a phase factor, since global phases are unobservable.

4.2 The lift to SU(2)

By the projective representation theorem (Wigner 1931, Bargmann 1954), every projective unitary representation of SO(3) lifts to a true unitary representation of its universal cover SU(2). The representations of SU(2) are labelled by $j \in \{0, \tfrac{1}{2}, 1, 3/2, 2, \ldots\}$, with integer j descending to genuine SO(3) representations and half-integer j being faithfully projective on SO(3).

4.3 Theorem 4.1: spin-½ in layer 0

Theorem 4.1 (Layer 0 Contains Spin-½). *Layer $\mathcal{C}^{(0)}$ admits the $j = \tfrac{1}{2}$ projective representation of SO(3), realized as a faithful (non-projective) representation of SU(2). It is the smallest-dimensional SU(2) representation that is projective rather than linear when restricted to SO(3) — i.e., the smallest representation in $\mathcal{C}^{(0)}$ that does not descend to a genuine linear SO(3) representation.*

Proof. See [VERSF–SPIN]. The argument is: state closure on rays forces projective SO(3) action; projective SO(3) lifts to SU(2); the $j = \tfrac{1}{2}$ representation is the smallest SU(2) representation that is projective rather than linear when restricted to SO(3). ■

4.4 Layer 0 vs layers 1–2: the role of postulate (M)

Theorem 4.1 establishes that $j = \tfrac{1}{2}$ *appears* in $\mathcal{C}^{(0)}$ — it does not establish that physics realizes the $j = \tfrac{1}{2}$ representation rather than $j = 0$ or higher half-integer j. To bridge this gap we invoke postulate (M) of §2.5: physics realizes the minimal nontrivial admissible representation. This is a selection principle, not a uniqueness result, and it is doing genuinely different work than the uniqueness arguments of layers 1 and 2.

We flag this asymmetry explicitly. Layer 1's connection is uniquely determined up to constant rescaling by the five-step argument of [VERSF–GAUGE] §5.4 (locality, linearity, Hermiticity, single covariant index, coefficient fixing). Layer 2's response field, given the closure-structural inputs of §6, is uniquely determined up to scale and a cosmological term by the imported Lovelock theorem. Layer 0, however, has no analogous uniqueness result; it has admissibility plus minimality.

A more satisfying replacement for (M) would be a stability or no-fine-tuning argument: that higher-j representations are unstable under closure dynamics, or are excluded by the fact-production constraints of [VERSF–FSN]. We do not have such an argument and identify it as an open problem (§10.7).

Layer 0 therefore generates the spin-½ sector — the fermionic matter fields, in standard terminology — *given (M)*. The Lorentz lift relating SO(3) $j = \frac{1}{2}$ to the 4-component Dirac spinor under the Lorentz group is deferred to [VERSF–PROTO]; the present theorem establishes only the SO(3) $j = \frac{1}{2}$ assignment.

---

5. Layer 1: Transport Closure and Spin-1

We summarise the spin-1 derivation of [VERSF–GAUGE]; the full derivation is given there. We include here a self-contained statement of two facts from that paper that are invoked below: the existence theorem for the connection (Theorem 5.1 of [VERSF–GAUGE]) and the curvature uniqueness theorem (Theorem 7.1 of [VERSF–GAUGE]).

5.1 The transport-closure problem

The partial derivative $\partial_\mu |\psi(x)\rangle$ is not closure-invariant: under $|\psi\rangle \longmapsto e^{i\theta(x)} |\psi\rangle$, it acquires the anomalous term $i(\partial_\mu \theta)\, e^{i\theta} |\psi\rangle$, which depends on the gradient of the unphysical phase choice and not just on the physical ray.

5.2 The forced connection

Closure-consistent local comparison forces a unique compensating structure: a real one-form A_μ on $\mathcal{S}$ with covariant derivative

$$D_\mu = \partial_\mu + i A_\mu$$

and gauge transformation rule

$$A_\mu \longrightarrow A_\mu - \partial_\mu \theta$$

for arbitrary smooth real $\theta : \mathcal{S} \longrightarrow \mathbb{R}$. Uniqueness in the Abelian case follows from a five-step argument in [VERSF–GAUGE] §5.4: locality, linearity in $|\psi\rangle$ inherited from unitarity, Hermiticity forced by reality of θ, single covariant index from contraction with dx^μ, and coefficient fixed up to constant rescaling.

5.3 Theorem 5.1: spin-1 in layer 1

Theorem 5.1 (Layer 1 Representation). *Under spatial isometries $g \in SO(3)$, the connection A_μ transforms as a one-form on $\mathcal{S}$. The minimal admissible representation of $\mathcal{C}^{(1)}$ therefore lies* *in the j = 1 (vector) representation of SO(3), and any quantum carrying its degrees of freedom carries SO(3) spin 1.*

Proof. See [VERSF–GAUGE] §8. Briefly: A_μ is a real one-form (Theorem 5.1 of that paper); one-forms transform as vectors under SO(3); the vector representation has dimension 3 = 2j + 1 with j = 1. ■

The minimality framing here is benign — the connection is genuinely uniquely determined up to constant rescaling, so there is no postulate (M) to invoke. To make the asymmetry across the three layers transparent at the point of statement: **postulate (M) of §2.5 is *not* invoked at layer 1, in contrast to layer 0 (§4.4)**, because the uniqueness of A_μ up to constant rescaling is a genuine theorem (the five-step argument of [VERSF–GAUGE] §5.4), not a selection principle. Layer 1 therefore generates the spin-1 sector — the Abelian gauge field, in standard terminology — without needing any selection postulate. The non-Abelian extension is deferred (§10.1).

5.4 Curvature lives in the antisymmetric sector of layer 2

The curvature F_μν = ∂_μ A_ν − ∂_ν A_μ is the natural element of the antisymmetric sector $\Lambda^2(\Omega^1(\mathcal{S}))$ of $c^{(2)}$ generated by transport closure: the covariant differential D maps $c^{(1)}$ to $c^{(2)}$, and the antisymmetric part dA = F lives in this sector.

The result we invoke from [VERSF–GAUGE] is that F_μν *generates* the local closure-invariants of the gauge sector. Specifically:

Imported Result ([VERSF–GAUGE] Theorem 7.1). *Let P[A, ∂A, $\partial^2 A$, ...] be a local polynomial functional of A_μ and its derivatives that is closure-invariant in the sense of §3.4. Then P is expressible as a polynomial in F_μν and its ordinary derivatives ∂_λ F_μν, ∂_λ ∂_κ F_μν, ...*

This is the algebraic Poincaré lemma for the Abelian gauge complex (see Henneaux–Teitelboim 1992, Ch. 12; Anderson 1989). It tells us that the antisymmetric sector of $c^{(2)}$ contains all local closure-invariants of the gauge-field content of layer 1 — specifically, F_μν F^μν is the simplest scalar density built from these, and (as developed in [VERSF–CCB]) becomes the kinetic term of the Yang–Mills Lagrangian under the Lorentz lift.

The *symmetric* sector of $c^{(2)}$ — the part not generated by curvature of A_μ — is what carries the spin-2 gravitational response. This is the content of the next section.

---

6. Layer 2: Back-Reaction Closure and Spin-2

This is the new derivation of the present paper. It establishes that closure events accumulate into a conserved rank-2 distinguishability source whose substrate response is a symmetric rank-2 field.

6.1 The back-reaction problem

Layers 0 and 1 describe how closure states *exist* and *are compared*. Layer 2's symmetric sector asks the further question:

What happens when closure events accumulate and flow across $\mathcal{S}$?

A closure event registers a fact — it is, by [VERSF–FSN], the production of an irreversible distinguishability increment. Accumulated closure produces accumulated distinguishability density, and this density must be carried somewhere on the substrate. The question is: what is the structural form of the response?

6.2 The closure source must be rank-2 — a sharper argument

The naïve dimension-counting argument ("a scalar can't encode flow, a vector can't encode stress, so we need rank-2") is too quick: a scalar density with a separate vector flow field is logically possible. The real argument runs through closure-density conservation and the fact that closure transfer is *bilinear* in directional structure.

Let $\rho_{cl}(x)$ denote the closure density (the rate at which closure events are produced at x). By the closure ontology of [VERSF–FSN], ρ_{cl} is conserved in the sense that the integrated closure on a region changes only via flux through its boundary. There is therefore a conserved current j^μ_{cl} with

$\partial_\mu j^\mu_{cl} = 0$ (or $\nabla_\mu j^\mu_{cl} = 0$ in the curved case).

This rank-1 object captures the *amount* of closure flowing in direction μ but not the *content* being transferred. Closure events carry content because they register distinctions between substrate states; on a smooth substrate, distinctions between neighbouring states have tangent-space directional structure inherited from the distinguishability metric d.

The cleanest way to see why this forces a rank-2 source is via the Noether parallel. In standard field theory, spacetime translation invariance generates one Noether current per translation direction, and the collection combines into $T_{\mu\nu}$ with one index for the current and one for the translation direction. At layer 2, the analogous role is played by isometries of the substrate distinguishability metric ($\mathcal{S}$, d): in the smooth-manifold limit, isometries of ($\mathcal{S}$, d) act as substrate translations and generate per-direction conserved currents, whose combination is rank-2. The second index of $T^{cl}_{\mu\nu}$ is the Noether-conjugate index of substrate isometries — a tangent-space index in the standard sense, and the closure analogue of the translation-direction index in the standard stress-energy construction.

The parallel is explicit but not exact. Standard $T_{\mu\nu}$ arises from Noether's first theorem applied to a Lagrangian symmetry; $T^{cl}_{\mu\nu}$ here arises from the combination of closure-conservation and substrate isometry, without committing to a particular Lagrangian. The two paths converge once the action principle of [VERSF–CCB] is supplied; without it, the present argument establishes the *structural type* of $T^{cl}_{\mu\nu}$ (rank-2, conserved, with one flux index and one isometry/content index) but not its precise relation to a canonical Noether current. The Lovelock-type uniqueness argument of §6.6 only requires the structural type, so the present argument suffices for that purpose; readers wanting a fully canonical Noether derivation should consult [VERSF–CCB].

Conservation then takes the rank-2 form

$$\nabla^\mu T^{cl}_{\mu\nu} = 0,$$

generalizing the rank-1 closure-density conservation to its content-resolved form. This is the structural input we need, and it is genuinely forced — given that closure events carry directional content and that the source must register both flux and content direction simultaneously, no lower-rank object suffices.

Symmetry of $T^{cl}_{\mu\nu}$. A separate question is why $T^{cl}_{\mu\nu}$ is symmetric in $(\mu\nu)$. The argument parallels the Belinfante–Rosenfeld construction in standard field theory: the antisymmetric part of a conserved rank-2 tensor decouples from the response field by a divergence-redefinition (a "Belinfante improvement"), so the physically coupled part of $T^{cl}_{\mu\nu}$ is its symmetric part. Equivalently, the antisymmetric part of $T^{cl}_{\mu\nu}$ would couple only to the antisymmetric sector of $c^{(2)}$ — i.e., to the gauge curvature $F_{\mu\nu}$ of layer 1, which is already accounted for there. The symmetric part is what couples to the symmetric sector $\mathrm{Sym}^2(\Omega^1)$ of $c^{(2)}$, which is the layer-2 response. We adopt $T^{cl}_{\mu\nu}$ symmetric on this basis.

6.3 The response field must be symmetric rank-2

Let $h_{\mu\nu}$ be the substrate response field in the symmetric sector of $c^{(2)}$. The minimal coupling consistent with locality and linearity in the source is

$$S_{int} = \tfrac{1}{2} \int d^4x\, h_{\mu\nu} T^{\mu\nu}_{cl}.$$

Since $T^{cl}_{\mu\nu}$ is symmetric (§6.2), only the symmetric part of $h_{\mu\nu}$ couples. The minimal admissible response field is therefore a symmetric rank-2 tensor:

$$h_{\mu\nu} = h_{\nu\mu}.$$

A scalar response cannot couple to $T^{cl}_{\mu\nu}$ directionally; a vector response cannot couple to the stress structure. The response *is* the symmetric sector of $c^{(2)}$.

6.4 Gauge-like redundancy of the response field

The response field $h_{\mu\nu}$ inherits a structural redundancy analogous to layer 1. Just as A_μ is defined only up to $A_\mu \longrightarrow A_\mu - \partial_\mu \theta$, the response field $h_{\mu\nu}$ is defined only up to

$$h_{\mu\nu} \longrightarrow h_{\mu\nu} + \partial_\mu \xi_\nu + \partial_\nu \xi_\mu,$$

for arbitrary smooth vector field ξ^μ on $\mathcal{S}$. This is the spin-2 analogue of gauge redundancy: not all components of $h_{\mu\nu}$ carry physical content; only the part invariant under this redundancy registers physical facts.

The structural origin of this redundancy is the same as for layer 1: closure events register *invariants* of the closure structure, and any quantity that changes under closure-equivalent re-labellings cannot be a physical fact. In the spin-2 case, the redundancy corresponds to infinitesimal substrate diffeomorphisms — relabellings of substrate coordinates that do not change the underlying closure pattern.

6.5 Theorem 6.1: layer 2 contains spin-2

Theorem 6.1 (Layer 2 Symmetric-Sector Representation). *The minimal admissible response field for closure back-reaction is a symmetric rank-2 tensor* $h_{\mu\nu}$ *with gauge-like redundancy* $h_{\mu\nu} \longrightarrow h_{\mu\nu} + \partial_\mu \xi_\nu + \partial_\nu \xi_\mu$. *Under spatial isometries* $g \in SO(3)$, *the symmetric-traceless part of* $h_{\mu\nu}$ *transforms in the* $j = 2$ *representation, and any quantum carrying the corresponding degrees of freedom carries SO(3) spin 2.*

Proof. By §§6.2–6.4, the response field is a symmetric rank-2 tensor. Under SO(3), a symmetric rank-2 tensor decomposes as

$$h_{\mu\nu} = (\text{trace part}) \oplus (\text{symmetric traceless part}),$$

with the trace part transforming as $j = 0$ and the symmetric traceless part as $j = 2$ (dimension $2j + 1 = 5$, matching the 5 independent components of a 3×3 symmetric traceless tensor). The gauge-like redundancy of §6.4 projects out the trace and the longitudinal modes; the surviving physical content lies in the $j = 2$ sector. ■

Layer 2 (symmetric sector) therefore generates the spin-2 sector — the gravitational field, in standard terminology. The minimality framing here is partial: the symmetric-rank-2 structure is forced by §§6.2–6.4, but the selection of the $j = 2$ sector over the $j = 0$ trace mode is delivered by the gauge-like redundancy projection rather than by postulate (M) directly. Postulate (M) re-enters when asking why the gravitational sector realizes only the $j = 2$ traceless mode and not also independent scalar gravitational degrees of freedom, which is the standard scalar-tensor-theory question; we leave that question to the dynamical companion paper [VERSF–GRAV-DYN].

6.6 The unique closure-invariant curvature: Einstein tensor

We now invoke the imported Lovelock-type uniqueness theorem of §2.2.

The closure-algebra contribution at this step is to show that the *input shape* of the Lovelock theorem is forced by closure structure: §6.3 forces a symmetric rank-2 response, §6.2 forces divergence-freeness, §6.4 forces the gauge-like redundancy that makes curvature (rather than $h_{\mu\nu}$ itself) the physical observable. Given these inputs, the Lovelock-type theorem then selects $G_{\mu\nu} + \Lambda_{\text{void}}\, g_{\mu\nu}$ uniquely.

This is the layer-2 analogue of the role played by [VERSF–GAUGE] Theorem 7.1 in the layer-1 sector: there, $F_{\mu\nu}$ was selected as the unique local closure-invariant generated by the connection; here, $G_{\mu\nu}$ is selected as the unique local divergence-free symmetric rank-2 invariant of the response field's curvature.

6.7 Forward reference to the action principle

The structural result of §6 is kinematic: it identifies the field type, its redundancy structure, and the unique invariant curvature object. Two distinct dynamical questions are deferred to [VERSF–CCB]: (i) the *existence* of an action principle whose Euler–Lagrange equations reproduce the field equation of §8, and (ii) the *selection* of the Einstein–Hilbert form $L = R/(2\kappa_{cl}) - \Lambda_{void}/\kappa_{cl}$ over higher-curvature variants such as $L = R + \alpha R^2 + \beta R_{\mu\nu} R^{\mu\nu} + \dots$. The present paper neither constructs nor selects an action; the reader noticing the absence of dynamical content in §§6–8 is correctly identifying that the structural form of the field equation is established here but its action-principle origin is deferred.

6.8 Ledger: closure-structural inputs vs standard motivations for Einstein's equation

The reviewer's third concern is that the closure-algebra contribution at layer 2 is thinner than the framing suggests, since the structural inputs we provide for the Lovelock argument (symmetric, rank-2, conserved, gauge-redundant) are nearly identical to the standard motivations for Einstein's equation. We agree with the technical observation and respond to it as follows.

The contribution at layer 2 is *not* the production of new structural inputs for the Lovelock argument. The contribution is the **derivation of these inputs from a single ontological source**, where standard treatments derive them from three separate sources. The ledger:

Standard input to Einstein equation	Standard origin	Closure-structural origin
$T_{\mu\nu}$ is a rank-2 tensor	postulate / convention from matter Lagrangian	closure events carry both flux direction (from conservation) and content direction (from substrate metric); both indices are forced (§6.2)
$\nabla^\mu T_{\mu\nu} = 0$	Noether's theorem applied to spacetime translation symmetry	closure-density conservation in [VERSF–FSN], extended to the rank-2 form by the bilinear directional structure of closure transfer (§6.2)
$T_{\mu\nu}$ is symmetric	Belinfante–Rosenfeld improvement applied to canonical $T_{\mu\nu}$	closure-algebra grading: the antisymmetric part couples to $F_{\mu\nu}$ in $\mathfrak{c}^{(2)}$ and is accounted for in layer 1; only the symmetric part is layer-2 content (§6.2)
$h_{\mu\nu}$ is symmetric rank-2	postulate of metric gravity ($g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$)	minimal admissible response in $\mathrm{Sym}^2(\Omega^1) \subset \mathfrak{c}^{(2)}$ that couples to $T^{cl}_{\mu\nu}$ (§6.3)

Standard input to Einstein equation	Standard origin	Closure-structural origin
$h_{\mu\nu}$ is diffeomorphism-invariant	postulate of general covariance	closure-equivalence redundancy of substrate relabellings (§6.4)

The **standard origin** column has five distinct postulational sources (the matter sector's stress-energy structure, spacetime translation invariance, the Belinfante–Rosenfeld procedure, metric gravity, and general covariance). The **closure-structural origin** column has one: the closure ontology of [VERSF–FSN] applied to a distinguishability substrate. This is the substantive content of the unification at layer 2 — not the invention of new structural inputs for Lovelock, but the consolidation of five standard postulates into a single ontological source. Whether this consolidation rises to a publishable contribution in foundations of physics is a judgement we leave to the reader; we believe it does, on the grounds that reducing five postulates to one is genuinely useful work in a reconstruction programme.

---

7. The Hierarchy Theorem and the No-Alternative Theorem

We can now state the consolidation result that organises §§4–6.

7.1 Hierarchy Theorem

Theorem 7.1 (Closure Representation Hierarchy). *Let* $\mathcal{E} \to \mathcal{S}$ *be a closure bundle over a distinguishability substrate, with local phase redundancy and admissible transformations preserving closure equivalence and transition probability. Let* $c = c^{(0)} \oplus c^{(1)} \oplus c^{(2)}$ *be the closure algebra of Definition 3.2. Adopt postulate (M) of §2.5 where invoked.*

Then the three graded layers contain the following SO(3) spin sectors:

Layer	Closure problem	Required structure	SO(3) spin sector
$c^{(0)}$	state identity	projective ray	$j = \frac{1}{2}$ (under (M))
$c^{(1)}$	local comparison	one-form connection A_μ	$j = 1$ (uniquely)
$c^{(2)}$ (sym.)	accumulated response	symmetric tensor $h_{\mu\nu}$	$j = 2$ (after redundancy projection)
$c^{(2)}$ (antisym.)	gauge curvature	two-form $F_{\mu\nu}$	(already in layer 1's content)

Proof. Theorem 4.1 establishes the $j = \frac{1}{2}$ result for layer 0 under (M) (citing [VERSF–SPIN]). Theorem 5.1 establishes the $j = 1$ result for layer 1 (citing [VERSF–GAUGE]). Theorem 6.1 establishes the $j = 2$ result for the symmetric sector of layer 2 (this paper §6). The antisymmetric sector of layer 2 contains the curvature F_μν of layer 1, by §5.4 and the imported result of [VERSF–GAUGE] Theorem 7.1. ■

The theorem is a consolidation: it organises the three layer results plus the cohomological closures of [VERSF–GAUGE] Theorem 7.1 and Anderson 1989 into a single statement about the algebra $\mathfrak{c}$. We do not claim it is more than that.

7.2 No-Alternative Theorem

The Hierarchy Theorem of §7.1 organizes the three layer results into a single statement. We now ask the further question: are there other independent grades that admissible closure operations could populate? Could there be a $\mathfrak{c}^{(3)}$, $\mathfrak{c}^{(4)}$, …, carrying additional spin sectors not contained in the three layers?

Under stated assumptions on what counts as an admissible local closure operation, the answer is no. Every higher-grade closure operation is either a derivative or polynomial composite of lower-grade objects, generated by F_μν, R_μνρσ, and h_μν via the cohomological closures of [VERSF–GAUGE] §7 and Anderson 1989. We state the result precisely, prove it, and own the residual structural commitments in §7.3.

Theorem 7.2 (No-Alternative Closure Algebra). *Let ℰ → 𝒮 be a closure bundle over a distinguishability substrate. Let an admissible operation on ℰ satisfy the following five assumptions:*

(A1) ***Ray admissibility*** *— the operation preserves the closure-equivalence |ψ⟩ ~ e^(iθ(x))|ψ⟩ for arbitrary smooth real θ : 𝒮 → ℝ;*

(A2) ***Local comparability*** *— comparison of closure states at neighbouring substrate points produces a well-defined object that is closure-invariant in the sense of §3.4;*

(A3) ***Closure conservation*** *— accumulated closure content is conserved on 𝒮 in the sense of [VERSF–FSN] (no closure appears or disappears without corresponding transfer), with the directional structure of closure content registered by the substrate distinguishability metric d. The rank-2 form of the resulting source tensor T^cl_μν is derived in §6.2 from this primitive conservation plus the directional content carried by the substrate metric, not assumed here;*

(A4) ***Locality (smooth limit)*** *— the operation is local or generated by local differential data on the smooth-manifold continuum limit of 𝒮 (see §2.3, §10.5(ii));*

(A5) ***Minimal non-redundancy*** *— no independent structure is admitted if it can be expressed as a derivative, polynomial product, or curvature/response object built from objects of lower grade.*

Then:

(a) The admissible graded algebra of closure operations on ℰ satisfying (A1)–(A5) is

$\mathfrak{c} = \mathfrak{c}^{(0)} \oplus \mathfrak{c}^{(1)} \oplus \mathfrak{c}^{(2)}$

with the form-degrees and sector decomposition of Definition 3.2.

(b) Under postulate (M) at layer 0, the graded layers contain the {½, 1, 2} SO(3) spin sectors as their physical content, after the redundancy projections of §§4.4, 5.3, 6.5. Other spin sectors that arise within the algebra (notably the $j = 0$ trace mode of the symmetric rank-2 sector at grade 2) are projected out by these redundancies and do not appear in the physical content of the corresponding grade.

(c) No independent grade $k > 2$ admissible operation exists: every such operation is expressible as $\nabla^m F_{\mu\nu}$, $\nabla^m R_{\mu\nu\rho\sigma}$, $\nabla^m h_{\mu\nu}$, or a polynomial composite of these, by the cohomological closures of [VERSF–GAUGE] Theorem 7.1 (Henneaux–Teitelboim 1992 Ch. 12) and Anderson 1989.

Proof. The proof proceeds in five steps: forcing each of the three grades from the corresponding admissibility assumption, establishing that the three structural roles of closure operations are exhaustive, and closing off higher grades via the imported cohomological theorems.

Step 1 — k = 0 is forced by ray admissibility (A1). The first admissible operation is identifying physical closure states modulo the global phase ambiguity. By [VERSF–SPIN] (summarised in §4), this gives projective state structure on rays; the projective representation theorem lifts it to SU(2), and postulate (M) selects $j = \tfrac{1}{2}$ as the smallest projective-rather-than-linear sector. So $\mathfrak{c}^{(0)}$ is the layer of state-closure operations and contains the spin-½ sector.

Step 2 — k = 1 is forced by local comparability (A2). Comparing closure states at neighbouring substrate points requires a connection. By [VERSF–GAUGE] §5.4 (summarised in §5), the unique structure (up to constant rescaling) that defines a closure-invariant local derivative on closure-state vector representatives is the one-form connection $D_\mu = \partial_\mu + i A_\mu$ with gauge transformation $A_\mu \longrightarrow A_\mu - \partial_\mu \theta$. The connection transforms as a one-form under SO(3) and carries $j = 1$. So $\mathfrak{c}^{(1)}$ is the layer of transport-closure operations and contains the spin-1 sector. (Postulate (M) is *not* invoked here; the result is uniqueness-driven.)

Step 3 — k = 2 is forced by curvature obstruction and conserved back-reaction (A3). Once a connection exists, the commutator $[D_\mu, D_\nu] = i F_{\mu\nu}$ is the natural two-index obstruction to path-independent comparison; $F_{\mu\nu}$ lives in the antisymmetric sector of $\mathfrak{c}^{(2)}$. Independently, the primitive closure-conservation of (A3), combined with the directional structure of closure content, is shown in §6.2 to force a conserved rank-2 source $T^{cl}_{\mu\nu}$ (the derivation runs through the Noether-parallel argument with substrate isometries playing the role of standard translations). The minimal symmetric rank-2 response $h_{\mu\nu}$ coupling to $T^{cl}_{\mu\nu}$ (§6.3) lives in the symmetric sector of $\mathfrak{c}^{(2)}$, with $j = 2$ content selected by the redundancy projection of §6.5. Both sectors are forced, and together they span $\mathfrak{c}^{(2)} = \Lambda^2(\Omega^1(\mathcal{S})) \oplus \mathrm{Sym}^2(\Omega^1(\mathcal{S}))$ as in Definition 3.2.

Step 4 — the trichotomy of admissible roles is exhaustive. Any admissible closure operation must fall into one of three structural roles: it must either (i) identify physical states, (ii) compare states at neighbouring points, or (iii) register accumulated response or obstruction. These three roles correspond precisely to the assumptions (A1), (A2), (A3) and to grades 0, 1, 2 by Steps 1–3.

We pre-empt three candidate operations that might appear to require a fourth role and show that each falls into one of the three or is reducible to them:

- *Finite-path comparison (parallel transport along non-infinitesimal paths, holonomies).* Such operations are not new closure primitives; they reduce to iterated infinitesimal comparisons (grade 1 applied along a path) plus path-ordered composition. They are therefore composites of grade-1 objects in the sense of (A5), not independent grade-k operations. Their gauge-invariant content is captured by the holonomy $\oint_\gamma A_\mu dx^\mu$, which is non-local and falls outside (A4); see [VERSF–GAUGE] §9.
- *Bundle automorphisms beyond gauge action.* Operations that act on the closure-bundle structure itself — changes of fibre, large bundle automorphisms — either act on closure states (role (i), grade 0) or on the connection one-form (role (ii), grade 1, since they enter the gauge transformation rule of A_μ). They do not constitute an independent role; they are decompositions of role-(i) and role-(ii) operations restricted to particular substructures.
- *Multi-state operations (entanglement-type couplings across substrate points).* Operations that act on tensor-product closure states $|\psi_1(x)\rangle \otimes |\psi_2(y)\rangle$ are role (i) operations on a composite state-space; the comparison structure between x and y is role (ii) content. For neighbouring x and y, this is direct grade-1. For finite (but local) separation, the comparison reduces to iterated infinitesimal grade-1, exactly as in the finite-path bullet above. For genuinely nonlocal entanglement structures (Bell-scenario-type correlations between distant points, with no path-integrated reduction), the comparison falls under (A4)'s locality exclusion and is outside the scope of Theorem 7.2 — the same status as the Wilson-loop / non-contractible-holonomy case noted in §7.3 below. Multi-state operations therefore introduce no new grade in any of the three sub-cases.

An operation that does *none* of (i), (ii), (iii) is structurally idle on $\mathcal{E}$ and is excluded by minimal non-redundancy (A5). *The trichotomy is the principal structural commitment of Theorem 7.2, and we own it as such — we have made the candidate-dismissal explicit here, but a future theorem deriving the trichotomy from a more primitive characterization of admissible closure operations would strengthen the result; see §7.3(i).*

Step 5 — no independent $k > 2$ (cohomological closure). Suppose some grade $k > 2$ contained an independent admissible operation $X^{(k)}$. By Step 4, $X^{(k)}$ must fall into one of the three structural roles and therefore reduce to grade ≤ 2 in its primitive content, *unless* $X^{(k)}$ is genuinely a higher-derivative or polynomial composite expression built from existing grade-2 objects. We exclude even this possibility by invoking the imported cohomological theorems:

- *Gauge sector.* [VERSF–GAUGE] Theorem 7.1 (the algebraic Poincaré lemma for the Abelian gauge complex; Henneaux–Teitelboim 1992 Ch. 12) establishes that every local closure-invariant polynomial in A_μ and its derivatives is expressible as a polynomial in $F_{\mu\nu}$ and $\nabla^m F_{\mu\nu}$. Higher-grade local closure-invariants in the gauge sector are therefore generated by the existing F_μν and its covariant derivatives — i.e., by composites and derivatives of grade-2 objects.

- *Gravitational sector.* By Anderson 1989 and the variational-bicomplex literature, every local diffeomorphism-invariant polynomial in g_μν and its derivatives is expressible as a polynomial in the Riemann tensor R_μνρσ and its covariant derivatives. Higher-grade local diffeomorphism-invariants in the symmetric sector are therefore generated by the existing curvature and its covariant derivatives — i.e., again by composites and derivatives of grade-2 objects.

By minimal non-redundancy (A5), composites and derivatives of grade-2 objects do not constitute independent grades.

Cross-sector composites. The two cohomological theorems above cover within-sector polynomial invariants (gauge-only and gravitational-only). A polynomial mixing F_μν and R_μνρσ — e.g., F_μν F^μν R, R_μνρσ F^μν F^ρσ — is a cross-sector composite. Such polynomials are gauge-invariant on the gauge factor (by construction from F_μν and ∇F) and diffeomorphism-invariant on the gravitational factor (by construction from R_μνρσ and ∇R), hence closure-invariant in the sense of §3.4. They are therefore polynomial composites of grade-2 objects from both sectors, and (A5) excludes them as independent grades on the same basis as within-sector composites. The cross-sector closure follows from the within-sector results plus closure of polynomial composition under (A5); we note it explicitly here so that statement (c)'s reference to "polynomial composites" is unambiguous.

There is therefore no independent k > 2 admissible operation, and the closure algebra is closed at grade 2: $\mathfrak{c} = \mathfrak{c}^{(0)} \oplus \mathfrak{c}^{(1)} \oplus \mathfrak{c}^{(2)}$.

This establishes (a) and (c). Statement (b) follows from Theorems 4.1, 5.1, and 6.1 with postulate (M) invoked at layer 0. ■

7.3 Remarks on residual structural commitments of Theorem 7.2

Theorem 7.2 establishes that no independent k > 2 grade exists *under the stated assumptions (A1)–(A5)*. We are explicit about the three residual structural commitments that remain part of the framework rather than being derived from more primitive principles. Owning these honestly is the difference between a defensible theorem and an overclaim.

(i) The trichotomy of admissible roles (Step 4). The three-way decomposition of closure operations into state-identification, comparison, and response/obstruction is the key structural commitment. It is presented as following from the closure ontology — every admissible operation must do something with closure states, and the three roles enumerate what "something" can mean — but a future theorem deriving the trichotomy from a more primitive characterization of "admissible local closure operation" (perhaps drawing only on [VERSF–FSN]'s definition of fact-production) would strengthen Theorem 7.2 from "no fourth grade given the three-role decomposition" to "no fourth grade simpliciter." We identify this as the leading possible upgrade in §10.5.

(ii) Locality presupposes the continuum limit. Assumption (A4) presupposes that $\mathcal{S}$ admits a smooth-manifold continuum description in which "local differential data" is well-defined. On the underlying K = 7 simplicial substrate, the natural locality notion is simplicial-local (involving cochains on simplices and their faces), and the bridge from simplicial to smooth locality is itself an open construction (§10.5(ii)). Theorem 7.2 should therefore be read as: *no independent k > 2 grade in the smooth-manifold setting that the closure dynamics are assumed to produce in the continuum limit*. The unconditional version, freed from the continuum-limit assumption, is open until that limit is constructed.

(iii) Postulate (M) at layer 0. Theorem 7.2's statement (b) — that the {½, 1, 2} sectors are the spin content of the three grades — invokes postulate (M) at layer 0 (§4.4). Layer 1 is uniqueness-driven and does not require (M); layer 2's symmetric sector is forced by §§6.2–6.4 with the j = 2 content selected by the redundancy projection rather than by (M) directly. The asymmetry is honest: (M) is the residual selection principle at layer 0, not a uniqueness theorem, and replacing it with a derived statement is identified as an open problem (§10.7).

What the theorem does not exclude. Theorem 7.2 rules out independent local-polynomial closure-invariants at grade k > 2 under (A1)–(A5). It does *not* exclude:

- **Nonlocal topological sectors** such as Wilson loops, Chern–Simons-like global structures, or holonomies on non-contractible loops. These are excluded from (A4)'s locality assumption by construction; their relationship to the closure-algebra picture is part of the broader VERSF programme but lies outside the present theorem.
- **Higher-spin extensions** in which closure states carry additional internal indices generating higher-spin representations beyond the {½, 1, 2} sectors. These would correspond to extending $\mathfrak{c}$ with new content rather than adding a fourth grade, and are addressed by the non-Abelian-extension open problem (§10.1) and the higher-spin open problem (§10.3).
- **Internal gauge-group structures** beyond the Abelian U(1) of layer 1. The non-Abelian extension is open (§10.1).

With these residual commitments owned, Theorem 7.2 establishes a substantive result: under the stated assumptions, the closure algebra $\mathfrak{c} = \mathfrak{c}^{(0)} \oplus \mathfrak{c}^{(1)} \oplus \mathfrak{c}^{(2)}$ is closed at grade 2, and the {½, 1, 2} spin hierarchy is the complete structural content of layer-by-layer closure operations. The strongest unconditional form of the theorem — with all three residual commitments derived from yet more primitive principles — remains open, and is the central structural target of the closure-algebra programme.

---

8. The Einstein-Form Field Equation

Combining §6 with the imported Lovelock-type result of §6.6 gives the structural form of the closure back-reaction equation.

Theorem 8.1 (Closure Back-Reaction). *Let T^cl_μν be a conserved symmetric rank-2 closure-density source on a metric manifold (𝒮, g_μν), with ∇^μ T^cl_μν = 0. Let the substrate response h_μν (entering through g_μν = η_μν + h_μν or its nonlinear completion) be required to be local, universal, symmetric, divergence-compatible with the source, and closure-invariant in the sense of §6.4. Then the minimal nonlinear field equation has Einstein form:*

G_μν + Λ_void g_μν = κ_cl T^cl_μν,

where G_μν = R_μν − ½ R g_μν is the Einstein tensor, Λ_void is a constant (the void closure pressure / cosmological term), and κ_cl is a coupling constant.

Proof sketch. By §§6.2–6.4, T^cl_μν is symmetric, conserved, and rank-2, and h_μν has the gauge-like redundancy of §6.4. The unique local symmetric rank-2 tensor with vanishing covariant divergence built from g_μν and at most its second derivatives is, up to the cosmological term, the Einstein tensor (Lovelock 1971, imported per §2.2). Matching divergences on both sides of the equation and absorbing the freedom into a single coupling constant κ_cl yields the stated form. ■

8.1 What this equation is and is not

What it is. A *structural* equation of motion: the minimal closure-back-reaction equation consistent with the structural inputs derived in §6 (symmetric rank-2, conservation, gauge-like redundancy) plus the standard differential-geometric machinery imported in §2.2.

What it is not. A derivation of the *coupling constant* κ_cl (which would require connection to Newton's G or to closure microphysics — see §10.2), a derivation of the *value* of the cosmological constant Λ_void (which is the subject of [VERSF–TWO-PLANCK]), a derivation of the full *nonlinear* gravitational dynamics including matter coupling (which requires [VERSF–CCB] + [VERSF–GRAV-DYN]), or a derivation that the smooth-manifold setting in which the equation lives is itself the correct continuum limit of the discrete substrate (open problem §10.5(ii)).

On the equivalence principle. Layer 2 derives the universal-coupling half of the equivalence principle: all closure-density sources couple to the same response field, since they live in the same closure ontology and couple through the symmetric sector of $c^{(2)}$. The geometric-equivalence half (local flatness) is imported, not derived: it comes packaged with the differential-geometric apparatus of §2.2. The standard equivalence principle is the conjunction of the two; the closure-algebra picture derives the easy half and imports the hard half. We claim no more than that.

8.2 Identification with the standard stress-energy tensor

The closure source tensor T^cl_μν can be identified, in the appropriate limit, with the standard stress-energy tensor T_μν of matter and field content. The identification is structural: matter fields and gauge fields are themselves (by Theorems 4.1 and 5.1) representations of layers 0 and 1 of the closure algebra, and their associated stress-energy is the closure-density content of those representations. A complete derivation of this identification — i.e., of why the stress-energy tensor of matter+gauge fields equals their closure-density tensor up to convention — requires the action-principle framework of [VERSF–CCB] and is deferred to that companion paper.

9. Discussion

9.1 Relation to standard treatments

In standard treatments, the three field types are introduced through three independent postulates: Hilbert space + Born rule + spinor representation for matter, the local-symmetry / gauge principle for gauge fields, and the equivalence principle + general covariance for gravity. In the VERSF account given here, all three flow from a single graded algebra:

closure ontology + distinguishability substrate ⟹ closure bundle $\mathcal{E} \to \mathcal{S}$ ⟹ closure algebra $\mathfrak{c}$ ⟹ {spin-½, spin-1, spin-2}.

The technical mathematics in each layer overlaps with standard treatments — projective representation theory in layer 0, Abelian gauge theory in layer 1, Lovelock-type uniqueness in layer 2 — but the three layers are organized by a single algebraic structure. The reduction in postulate count is the substantive contribution; the §6.8 ledger makes the layer-2 reduction explicit.

9.2 Relation to other unification programmes

The closure-algebra unification differs from other unification approaches in instructive ways.

Grand unified theories (GUTs) unify the gauge sector by embedding SU(3) × SU(2) × U(1) into a larger simple group. They unify the gauge layer with itself, not the gauge layer with matter or gravity, and they leave the spin hierarchy untouched.

Supersymmetry unifies fermions (spin-½, layer 0) with bosons (spin-1, layer 1; spin-2, layer 2 in supergravity). It unifies layers across the grading rather than within. The closure-algebra picture is structurally compatible with supersymmetry but does not require it: the {½, 1, 2} hierarchy is contained in $\mathfrak{c}$ regardless of whether the layers are related by additional symmetry.

Algebraic quantum field theory (AQFT) derives gauge structure from local algebras and superselection sectors via the Doplicher–Haag–Roberts construction. AQFT and the closure-algebra picture differ in which structural inputs they take as primitive: AQFT imports Lorentz covariance and microcausality; VERSF defers both to [VERSF–PROTO] (proto-time) and [VERSF–CCB] (closure-event ordering). Whether AQFT-with-emergent-relativity reduces to the closure-algebra picture, or vice versa, is open methodological work we do not attempt to settle here.

9.3 Relation to operational reconstructions

Operational reconstructions of quantum theory (Hardy 2001, Masanes–Müller 2011, Chiribella–D'Ariano–Perinotti 2011) derive single-system Hilbert space and Born rule from operational primitives but do not address gauge or gravitational structure. The closure-algebra picture extends the reconstruction programme to all three sectors by adding a single structural input — the distinguishability substrate *S* over which closure states vary — to the operational machinery. Whether other structural augmentations of the operational programmes could deliver the same {½, 1, 2} hierarchy, and with what minimal inputs, is an open question for the foundations community.

9.4 Relation to the BCB Lagrangian formulation

The closure-invariant Lagrangian densities at each layer — the Dirac kinetic term at layer 0, F_μν F^μν at layer 1, and the Einstein–Hilbert term R at layer 2 — form the kinetic content of the BCB action of [VERSF–CCB]. The present paper establishes the *kinematic* origin of each layer; [VERSF–CCB] establishes the *dynamic* action principle that selects these specific Lagrangian densities. The two papers are complementary: this one shows that the field types and their closure-invariance structures are forced; the BCB paper shows how their dynamics are constrained by a unified action principle.

---

10. Open Problems

The closure-algebra unification raises a structured set of open problems.

10.1 Non-Abelian extension of layer 1

The spin-1 derivation of [VERSF–GAUGE] is rigorously Abelian. The non-Abelian extension requires (i) internal closure indices on closure states, (ii) Lie-group structure on those indices rather than discrete or other algebraic structure, (iii) derivation of the specific Standard Model gauge group SU(3) × SU(2) × U(1). Whether the K = 7 simplicial closure structure of [VERSF–KSEVEN] generates these is open. A non-Abelian analogue of [VERSF–GAUGE] Theorem 7.1 (gauge-invariant local polynomials in A_μ^A and its derivatives generated by the non-Abelian curvature F_μν^A) is also required and would import from Henneaux–Teitelboim 1992 Ch. 12 if not derived structurally.

10.2 Coupling constants

The three layers have associated coupling constants — gauge couplings g_s, g, g' at layer 1, and the gravitational coupling κ_cl ↔ Newton's G at layer 2. None are derived in the present paper. Connection to the VERSF derivation of $\alpha^{-1} \approx 137.143$ [VERSF–ALPHA] is suggestive but open.

10.3 Higher-derivative and higher-spin extensions

The hierarchy theorem identifies the {½, 1, 2} spin sectors. It does not address higher-spin content (j = 3/2 gravitinos, j ≥ 3) or higher-derivative gravitational theories (Gauss–Bonnet, Lovelock with higher-curvature terms). These correspond to extensions of c or to higher-grade content; their VERSF status is open.

10.4 Quantum gravity

The present paper derives the structural form of the gravitational field equation but does not address quantum gravity — i.e., what happens when the response field h_μν itself becomes a closure-state field carrying its own spin-½/spin-1 closure content. The framework suggests this question has a natural formulation but no derivation is attempted.

10.5 Strengthening the No-Alternative Theorem

Theorem 7.2 establishes that no independent k > 2 grade exists under five stated assumptions, with three residual structural commitments owned in §7.3. Strengthening the theorem to its unconditional form requires resolving each of those commitments, in the order in which §7.3 enumerates them.

(i) Deriving the trichotomy of admissible roles (§7.3(i)). Step 4 of the Theorem 7.2 proof — that admissible closure operations fall into exactly three structural roles (state, comparison, response) — is presented as following from the closure ontology rather than derived from a precise prior characterization. A future theorem deriving the trichotomy from [VERSF–FSN]'s definition of fact-production (or from a similarly primitive structural source) would upgrade Theorem 7.2 from "no fourth grade given the three-role decomposition" to "no fourth grade simpliciter." This is the leading structural open problem of the programme.

(ii) Constructing the continuum limit (§7.3(ii)). Assumption (A4)'s locality presupposes the smooth-manifold continuum description of S. On the underlying K = 7 simplicial 2-complex of [VERSF–KSEVEN], the natural locality notion is simplicial-local (involving cochains on simplices and their faces); the smooth-manifold differential geometry invoked in §§3, 5, 6, 8 must be derived as an emergent description of closure dynamics on this discrete substrate. The construction is plausibly addressed in [VERSF–PROTO] (where smooth Lorentz invariance is itself derived as emergent from proto-time) and [VERSF–CCB] (where the action principle is constructed in the continuum), but the explicit limit construction is open. Until it is supplied, Theorem 7.2 holds only in the smooth setting. The unconditional version requires either constructing this limit explicitly or reformulating Theorem 7.2 directly in the simplicial setting (with simplicial cochain degree replacing form degree). This is one of the most consequential gaps in the programme.

(iii) Replacing postulate (M) (§7.3(iii)). Theorem 7.2(b) — that the {½, 1, 2} sectors are the spin content of the three grades — invokes postulate (M) at layer 0 (§10.7). Replacing (M) with a derived selection principle (perhaps via a stability or no-fine-tuning argument) would close the asymmetry across the three layers.

All three are open structural problems. The unconditional version of Theorem 7.2 — no independent k > 2 grade simpliciter, with the {½, 1, 2} hierarchy forced rather than postulated — depends on resolving all three.

10.6 Continuum limit from K = 7 simplicial substrate

See §10.5(ii) for the canonical statement of this gap. The continuum-limit construction is logically prior to the unconditional version of Theorem 7.2 and is identified there as residual commitment (ii) of the No-Alternative Theorem.

10.7 Status of postulate (M)

Postulate (M) — that physics realizes the minimal nontrivial admissible representation — is invoked at layer 0 (and partially at layer 2). Replacing it with a stability or no-fine-tuning argument that excludes higher-j representations from physical realization is a desirable upgrade. Possible routes: stability under closure dynamics, exclusion via fact-production constraints from [VERSF–FSN], or a representation-theoretic uniqueness result we have not yet identified.

10.8 Mass generation

The Higgs mechanism, electroweak symmetry breaking, and gauge-boson masses are dynamical questions outside the scope of the present structural framework.

10.9 Lorentz lift

The spin assignments are SO(3) statements. The Lorentz lift — relating SO(3) j = ½ to the 4-component Dirac spinor, j = 1 to the 4-vector A^μ, and j = 2 to the symmetric tensor g^μν under the Poincaré group — is deferred to [VERSF–PROTO].

---

11. What This Paper Does Not Claim

To restate the scope explicitly:

- **Mathematical novelty in any sector.** The technical content of each layer overlaps with standard treatments. The contribution is interpretive consolidation under a single algebraic object (§1.1).
- **Dynamical equations.** The {½, 1, 2} hierarchy is *kinematic*. Equations of motion (Dirac, Yang–Mills, Einstein) require [VERSF–CCB].
- **Coupling constants.** Numerical strengths of gauge and gravitational couplings are not derived here.
- **Specific gauge groups.** Only the Abelian U(1) case is rigorously derived for layer 1.
- **Lorentz-spin assignment.** Theorems 4.1, 5.1, 6.1 are SO(3) statements; the Lorentz lift is deferred.

- **Continuum limit.** The smooth-manifold setting is taken as an emergent description; its derivation from the discrete $K = 7$ substrate is open.
- **The unconditional No-Alternative Theorem.** Theorem 7.2 establishes closure of the algebra at grade 2 under five stated assumptions (A1)–(A5), with three residual structural commitments owned in §7.3 (the trichotomy of admissible roles, the continuum-limit-presupposing locality, and postulate (M) at layer 0). The unconditional version, with these commitments derived rather than assumed, is open (§10.5).
- **The full equivalence principle.** Layer 2 produces universal coupling to closure sources; the local-flatness content of the equivalence principle's strong form requires the imported differential-geometric apparatus.
- **Quantum gravity.** Quantisation of the layer-2 response field is not addressed.
- **New empirical predictions.** None are claimed beyond those of the underlying standard sectors.

12. Conclusion

We have shown that the spin hierarchy $\{\tfrac{1}{2}, 1, 2\}$ of fundamental physics — fermionic matter, gauge bosons, gravity — appears as the three graded layers of a single closure algebra $\mathfrak{c}$ over the distinguishability substrate of VERSF, with structure:

Layer	Closure problem	Required structure	Spin (SO(3))
$\mathfrak{c}^{(0)}$	state identity (rays)	projective representation	$j = \tfrac{1}{2}$ (under postulate (M))
$\mathfrak{c}^{(1)}$	local comparison	one-form connection A_μ	$j = 1$ (uniquely)
$\mathfrak{c}^{(2)}$ symmetric	accumulated back-reaction	symmetric tensor $h_{\mu\nu}$	$j = 2$ (after redundancy projection)
$\mathfrak{c}^{(2)}$ antisymmetric	gauge curvature	two-form $F_{\mu\nu}$	(layer 1 content)

Closure-density conservation $\nabla^\mu T^{cl}_{\mu\nu} = 0$, combined with the imported Lovelock-type uniqueness theorem for divergence-free symmetric rank-2 curvature objects, forces the structural form of the gravitational field equation:

$$G_{\mu\nu} + \Lambda_{void}\, g_{\mu\nu} = \kappa_{cl}\, T^{cl}_{\mu\nu}.$$

The contribution is interpretive consolidation under a single algebraic object: three apparently independent structural postulates of physics — Hilbert space + Born rule + spinor representation, the gauge principle, the equivalence principle + general covariance — are organized by a single graded algebra of closure operations. The §6.8 ledger makes the layer-2 reduction explicit by showing five standard postulational sources for the inputs to the Einstein equation collapsing to one ontological source: the closure ontology of [VERSF–FSN] applied to a distinguishability substrate.

What we have done in this paper is establish two theorems organizing the closure-algebra picture: the *Hierarchy Theorem* (§7.1), consolidating the three layer results into a single statement, and the *No-Alternative Theorem* (§7.2), establishing under five stated assumptions (ray admissibility, local comparability, closure conservation, locality in the smooth-manifold limit, and minimal non-redundancy) that no independent grade $k > 2$ exists and that the {½, 1, 2} hierarchy is closed. The latter is the strongest unification result the present paper supports; its unconditional form (with the three residual commitments of §7.3 — the trichotomy of admissible roles, the continuum-limit-presupposing locality, and postulate (M) at layer 0 — derived rather than assumed) remains open (§10.5).

We have also not constructed the continuum limit from the discrete $K = 7$ substrate to the smooth-manifold setting in which §§3, 5, 6, 8 operate; that bridge is the second consequential open problem and is logically prior to the unconditional version of Theorem 7.2.

The conceptual content of the result, stated honestly: the spin-½, spin-1, and spin-2 sectors of physics are *contained as the three graded layers* of a single closure algebra on a distinguishability substrate, and *no independent fourth layer exists* under the structural assumptions (A1)–(A5) of Theorem 7.2. Whether the underlying assumptions can themselves be derived from more primitive principles is open work.

The next steps are: the Lorentz lift [VERSF–PROTO]; the non-Abelian extension and gauge-group selection (§10.1); the dynamical action principle [VERSF–CCB]; the dynamical gravitational sector [VERSF–GRAV-DYN]; the continuum-limit construction (§10.5(ii)); and strengthening Theorem 7.2 to its unconditional form (§10.5).

References (selected)

VERSF programme

- VERSF–CHS: *Why a Fact-Producing Universe Must Satisfy Interference: Deriving Complex Hilbert Space from Admissibility Axioms.*
- VERSF–FSN: *Facts as Structural Necessities: The Joint Necessity Theorem.*
- VERSF–SPIN: *Spin as a Double-Cover Representation of Distinguishability Isometries.*
- VERSF–GAUGE: *Gauge Fields as Closure Connections: A Reconstruction-Level Derivation of the Abelian Connection and Spin-1.*
- VERSF–KSEVEN: *No-Go Theorem on Non-Simplicial Relational Substrates: K = 7 and Triangular 2-Complex Geometry as Structural Necessities.*
- VERSF–CCB: *The VERSF Constraint and Lagrangian: BCB as the Action-Principle Face of VERSF.*
- VERSF–PROTO: *Proto-Time and Emergent Lorentz Invariance.*
- VERSF–GRAV-DYN: *Gravity from Fold-Density Gradients: The Spin-2 Closure Back-Reaction Dynamics.*
- VERSF–TWO-PLANCK: *Cosmological Constant from the Two-Planck Principle.*

- VERSF–ALPHA: *Derivation of the Fine-Structure Constant from K = 7 Hexagonal Closure Constraints.*

Standard literature

- Y. Aharonov and D. Bohm, *Significance of Electromagnetic Potentials in the Quantum Theory*, Phys. Rev. **115** (1959) 485–491.
- I. M. Anderson, *The Variational Bicomplex*, Utah State University (1989).
- V. Bargmann, *On Unitary Ray Representations of Continuous Groups*, Annals of Mathematics **59** (1954) 1–46.
- F. J. Belinfante, *On the Spin Angular Momentum of Mesons*, Physica **6** (1939) 887–898; L. Rosenfeld, *Sur le tenseur d'impulsion-énergie*, Mém. Acad. Roy. Belg. **18** (1940) 1–30.
- G. Chiribella, G. M. D'Ariano, P. Perinotti, *Informational Derivation of Quantum Theory*, Phys. Rev. A **84** (2011) 012311.
- S. Doplicher, R. Haag, J. E. Roberts, *Local Observables and Particle Statistics I, II*, Comm. Math. Phys. **23** (1971) 199; **35** (1974) 49.
- N. Dragon, *BRS Symmetry and Cohomology*, hep-th/9602163 (1996).
- L. Hardy, *Quantum Theory From Five Reasonable Axioms*, quant-ph/0101012 (2001).
- M. Henneaux and C. Teitelboim, *Quantization of Gauge Systems*, Princeton University Press (1992).
- D. Lovelock, *The Einstein Tensor and Its Generalizations*, J. Math. Phys. **12** (1971) 498–501.
- L. Masanes and M. Müller, *A Derivation of Quantum Theory from Physical Requirements*, New J. Phys. **13** (2011) 063001.
- C. N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96** (1954) 191–195.
- E. Wigner, *Group Theory and Its Application to the Quantum Mechanics of Atomic Spectra* (1931 / 1959).