Gauge Fields as Closure Connections in the VERSF Framework

A Reconstruction-Level Derivation of the Abelian Connection and Spin-1

Keith Taylor *VERSF Theoretical Physics Programme*

Plain-Language Summary

Modern physics describes electromagnetism, the strong nuclear force, and the weak nuclear force as **gauge fields** — structures whose mathematical form is dictated by certain symmetry requirements. The standard textbook story tells you how to build a gauge theory once you decide you want one, but it leaves a basic question unanswered: *why does nature have these structures at all?*

This paper addresses the simplest case (electromagnetism, mathematically a "U(1) gauge field") within the VERSF reconstruction programme, which derives the structure of quantum theory from a small set of ontological starting points. The result: given that physical states are what earlier VERSF papers establish them to be — rays in a complex Hilbert space, where the choice of "phase" of a quantum state at any point in space is physically arbitrary — the existence of an electromagnetic-like field is *forced* as a structural consequence of comparing states at different points. The field is precisely what makes such comparison consistent. The quantum that carries the field is automatically what physicists call "spin-1."

The technical mathematics is identical to the standard gauge-theory derivation given by Yang and Mills (1954) and found in every quantum field theory textbook. What changes is the *route to it*: the equations no longer follow from a separately postulated symmetry whose origin remains mysterious; they follow from a structural feature of how quantum states relate to one another. The contribution is therefore best understood as **interpretive consolidation** — showing that what previously required two postulates (the structure of quantum states, plus the existence of gauge fields) can be reduced to one (the structure of quantum states alone).

The Aharonov–Bohm effect — a well-known experiment confirming that electromagnetic potentials carry physical content beyond the local field strength — is reproduced as a direct consequence. The extensions to the strong and weak nuclear forces, and the connection to gravity, are flagged explicitly as open problems for companion work, not claims of this paper.

Abstract

We give a reconstruction-level derivation of the Abelian gauge connection and its spin-1 quantum within the **Void Energy-Regulated Space Framework (VERSF)**. Building on prior results — in which complex Hilbert space, the Born rule, and unitary dynamics are derived from the geometry of distinguishable configurations, and SU(2) is shown to be the universal cover under which physical states transform — we show that (i) the local comparison of closure states across the distinguishability substrate is ill-defined without additional structure, (ii) given a single global phase ambiguity, closure-consistent comparison forces a unique connection structure A_μ with covariant derivative $D_\mu = \partial_\mu + iA_\mu$ and gauge transformation rule $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$, (iii) the curvature $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the unique elementary local closure-invariant observable built from A_μ, and (iv) under spatial isometries the connection transforms as a one-form and therefore corresponds to a quantum carrying $j = 1$ in the SO(3)/SU(2) representation theory of VERSF–SPIN. We give an explicit worked example of the U(1) case, with the Aharonov–Bohm effect as empirical anchor. The technical mathematics overlaps with the standard gauge-principle derivation of Yang and Mills; the present contribution is *interpretive consolidation* — the same equations now flow from a structural feature of the closure ontology rather than from a separately postulated local invariance principle. We are careful to distinguish what is *derived* in this paper (the Abelian connection, its uniqueness given a single phase ambiguity, the gauge transformation rule, the closure-invariance and uniqueness of curvature, the SO(3) $j = 1$ character of the connection) from what is *imported* ($\mathscr{H}$, ray-space ontology, SO(3)/SU(2) representation theory) and from what is *deferred* (the Lorentz-spin assignment, which depends on the emergent-Lorentz reconstruction of VERSF–PROTO; the non-Abelian extension; specific gauge groups; coupling constants; and full Yang–Mills dynamics, which is the subject of VERSF–CCB).

Contents

---

1. Introduction

In standard quantum field theory, gauge fields are introduced via the **gauge principle**: one *postulates* invariance under a local symmetry group, observes that a partial derivative ∂_μψ no longer transforms covariantly, and *introduces* a compensating field A_μ whose transformation rule restores covariance. The procedure is mathematically clean but conceptually unsatisfying: it explains how to write down a gauge theory once one decides one wants one, but it does not explain why nature should have gauge fields at all.

Three questions remain unanswered in the standard treatment:

1. **Why is local symmetry "required"?** The principle is usually motivated by analogy with general covariance, or by the wish to make a global symmetry local. Neither is a derivation.
2. **Why does the gauge field have spin-1?** This is read off from the Lorentz transformation properties of A_μ once it has been introduced, but it is not derived from anything more fundamental.
3. **Why is the curvature F_μν physically observable while A_μ is not?** The standard answer — gauge invariance — restates the question rather than answering it.

The VERSF programme has previously addressed the prior structural questions: ℋ and complex amplitudes are derived in [VERSF–CHS], and the SO(3)/SU(2) representation structure governing rotational symmetry is derived in [VERSF–SPIN]. This paper addresses the gauge sector in the **Abelian** case. We show that for a single global phase ambiguity, gauge structure is a *forced consequence* of comparing closure states across the distinguishability substrate when those states are defined only up to local phase — and that the field carrying this comparison structure is necessarily a one-form under SO(3), corresponding to the $j = 1$ representation.

1.1 What is and is not new

Before proceeding, we state the contribution of this paper precisely. The technical mathematics derived below — the covariant derivative $D_\mu = \partial_\mu + iA_\mu$, the gauge transformation rule $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$, the curvature $F_{\mu\nu}$, and the spin-1 assignment of the connection's quantum — is well-established and overlaps with the standard derivation of Abelian gauge theory dating to Yang and Mills (1954) and earlier work on electromagnetic gauge invariance.

What is new is the *motivational route*. In the standard derivation, the procedure begins with a *postulated* local symmetry and derives the connection as a compensating field. In the VERSF derivation, the procedure begins with a *structural feature of the relational closure ontology* — the fact that physical states are rays, and vector representatives are independently chosen at each substrate point — and derives the connection as the unique structure that renders local comparison closure-consistent. The two paths converge on the same equations.

The contribution of this paper is therefore **interpretive consolidation**, not new mathematics. Within a reconstruction programme that aims to derive physics from a small ontological base, this consolidation is genuinely valuable: it removes the postulated local-symmetry principle from the list of inputs to physics, replacing it with a structural feature of the closure ontology that is already imported. But the paper does not claim mathematical novelty in the gauge sector, and the technical content should be expected to coincide with standard treatments. This framing is what licenses the rest of the paper.

The argument proceeds as a chain of forced consequences from a small set of clearly-stated assumptions, summarised in §2.

2. What Is Assumed and What Is Derived

To avoid the common pitfall of importing classical field-theoretic structure through the back door, we state the input/output ledger explicitly.

2.1 Imported from prior VERSF reconstruction

Imported result	Source
A distinguishability substrate Λ with metric d preserving an interference-compatible algebra	VERSF–CHS
$\mathbb{C}$-valued amplitudes; complex Hilbert space $\mathscr{H}$ as the state space	VERSF–CHS
Reversible dynamics are unitary on $\mathscr{H}$	VERSF–CHS
Relational closure ontology (states are closure patterns)	VERSF–FSN

Imported result	Source
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 under SO(3) and linearly under SU(2)	VERSF–SPIN
Three-dimensional simplicial $K = 7$ closure substrate	VERSF–KSEVEN

2.2 Derived in this paper

Derived result	Section
Local comparison of closure states is ill-defined without additional structure	§4
Closure-consistent comparison forces a connection $D_\mu = \partial_\mu + iA_\mu$ (Abelian case)	§5
Gauge transformation $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$ emerges as structural identity	§5
Gauge "redundancy" is closure equivalence, not postulated symmetry	§6
Curvature $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the unique elementary local closure-invariant	§7
The connection transforms as a one-form ⇒ field quantum carries $j = 1$ (under SO(3))	§8
Aharonov–Bohm effect predicted as direct consequence	§9

2.3 Deferred to companion papers

Deferred question	Status
Lorentz-spin assignment of gauge bosons	Deferred to VERSF–PROTO (emergent Lorentz from proto-time)
Non-Abelian extension and which gauge groups appear in nature	Open problems flagged in §10
Yang–Mills dynamics and the action principle	Subject of VERSF–CCB
Mass generation, electroweak symmetry breaking	Outside scope of present programme stage
Spin–statistics for gauge bosons	Requires VERSF–PROTO (relativistic causal structure) plus additional locality input not yet developed in the programme

The seven imported items in §2.1 are *all* of the structural inputs. No gauge field, no covariant derivative, no curvature, no gauge transformation rule, no spin-1 assignment, and no gauge group is imported. Each appears as output (the Abelian case rigorously; the Lorentz lift and non-Abelian generalisation explicitly deferred).

---

3. Setup and Definitions

3.1 Closure states across the substrate

Let Λ be the distinguishability substrate of [VERSF–CHS] and let

$$|\psi(x)\rangle \in \mathscr{H}, \; x \in \Lambda ,$$

denote the closure state at substrate point x. By the imported reconstruction (§2.1), the physical content of $|\psi(x)\rangle$ is its ray:

$$[\psi(x)] = \{ e^{i\theta} \, |\psi(x)\rangle : \theta \in \mathbb{R} \} \subset \mathscr{H} .$$

Two vectors in the same ray correspond to the same physical state — global phase is, by the relational closure ontology of [VERSF–FSN], a quantity that no closure event can register and is therefore not part of the physical content of the state.

3.2 Definition: closure invariance

The phrase *closure-invariant* is used throughout the paper in a single technical sense, which we now fix.

Definition (Closure Invariance). A quantity $Q[\psi, A, \ldots]$ constructed from closure-state vector representatives $|\psi(x)\rangle$, the connection $A_\mu(x)$, and their derivatives is *closure-invariant* if it is unchanged under arbitrary local re-phasing

$$|\psi(x)\rangle \mapsto e^{i\theta(x)}|\psi(x)\rangle , \; A_\mu(x) \longrightarrow A_\mu(x) - \partial_\mu\theta(x) ,$$

for any smooth real-valued function $\theta : \Lambda \longrightarrow \mathbb{R}$.

Closure invariance in this sense coincides with what the standard literature calls *gauge invariance*. We adopt the term "closure-invariant" because it makes explicit the structural reason for the requirement: the invariance reflects the fact that re-phasing changes only the vector representative, not the physical ray, and so any quantity registering a physical fact must be invariant under it. There is no separate "ontological closure-invariance" carrying additional content beyond this.

3.3 Notational conventions

Two arrow conventions are used throughout. The mapsto symbol $\mapsto$ denotes the action of a transformation on a state (specifically, a vector representative of a closure-state ray); the long arrow $\longrightarrow$ denotes the induced transformation rule on a field (specifically, the gauge field A_μ). For example, the joint re-phasing of Definition 3.2 takes the form

$$|\psi\rangle \mapsto e^{i\theta}|\psi\rangle \text{ paired with } A_\mu \longrightarrow A_\mu - \partial_\mu\theta .$$

Both arrows describe parts of the same closure-equivalence operation; the typographic distinction tracks whether the object transformed is a state or a field.

---

4. The Local Comparison Problem

We now make precise the sense in which "comparing closure states at different substrate points" is ill-defined without additional structure. This is the result on which the rest of the paper rests.

The setting is straightforward. We have closure states $|\psi(x)\rangle$ varying across the substrate Λ, and we wish to define a notion of how they change from point to point — a local derivative compatible with the closure ontology. In a flat physical theory, change is captured by the partial derivative $\partial_\mu|\psi(x)\rangle$. The central observation is that this expression, taken at face value, is not closure-invariant.

4.1 Independent local phase choices

At each substrate point x, the physical state is the ray $[\psi(x)]$. A *vector representative* $|\psi(x)\rangle \in [\psi(x)]$ is obtained by an arbitrary choice of phase. The relational closure ontology imposes no relation between phase choices at different points: the choice at x and the choice at $x + dx$ are independent.

Formally, let

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

denote a *re-phasing* of the state, where $\theta : \Lambda \longrightarrow \mathbb{R}$ is an arbitrary smooth local phase function. Both $|\psi\rangle$ and $|\tilde{\psi}\rangle$ encode the same ray-valued physical state. Any closure-respecting structure must give the same physical predictions under either choice.

4.2 The partial derivative is not closure-invariant

Consider the partial derivative

$$\partial_\mu|\psi(x)\rangle = \lim_{dx \to 0} \left[\, |\psi(x + dx)\rangle - |\psi(x)\rangle \,\right] / dx_\mu \,.$$

Under the local phase transformation $|\psi(x)\rangle \longmapsto e^{i\theta(x)}\, |\psi(x)\rangle$, this expression transforms as

$$\partial_\mu|\psi\rangle \longmapsto e^{i\theta}\, \partial_\mu|\psi\rangle + i\, (\partial_\mu\theta)\, e^{i\theta}\, |\psi\rangle \,.$$

The second term — the *anomalous derivative term* — depends on the gradient of the phase choice and is not present in the original expression. Two observers who chose different phase representatives at x and $x + dx$ would compute different "rates of change" for the same physical state. The partial derivative is therefore *not closure-invariant*: it depends on the unphysical phase choice as well as the physical ray.

4.3 Theorem: closure-invariant comparison requires additional structure

Theorem 4.1 (Local Comparison Problem). *Let* $|\psi(x)\rangle$ *be a vector representative of the closure-state ray* $[\psi(x)]$*. There is no operator that takes* $|\psi(x)\rangle$ *and* $|\psi(x + dx)\rangle$ *to a closure-invariant comparison without supplying additional data beyond the ray-valued state itself.*

Proof. Any operator constructed from $|\psi(x)\rangle$ and $|\psi(x + dx)\rangle$ alone — including $\partial_\mu|\psi\rangle$ and the inner product $\langle\psi(x) \mid \psi(x + dx)\rangle$ — depends on the chosen vector representatives. Under the independent re-phasings $|\psi(x)\rangle \longmapsto e^{i\theta(x)}|\psi(x)\rangle$ and $|\psi(x + dx)\rangle \longmapsto e^{i\theta(x + dx)}|\psi(x + dx)\rangle$, any such operator acquires a phase factor $e^{i[\theta(x + dx) - \theta(x)]} \neq 1$ in general. Closure invariance — invariance under arbitrary local re-phasing in the sense of Definition 3.2 — is therefore unattainable from the ray data alone. ■

This theorem is the structural pivot of the paper. It states that the local comparison of closure states is *not just undetermined but structurally underspecified* — there is no closure-invariant comparison of nearby states using only the ray-valued data. To recover a closure-respecting derivative, additional structure must be supplied.

The remainder of the paper analyses what that additional structure must be, **in the case of a single global phase ambiguity**.

5. The Forced Introduction of a Connection (Abelian Case)

The derivation that follows reaches the standard equations of Abelian gauge theory: the covariant derivative, the gauge transformation rule, and the curvature. Per §1.1, the contribution of this paper lies in the structural motivation for these equations, not in the equations themselves. The reader will recognise the technical mathematics; what is new is that it now flows from the closure ontology imported in §2.1 rather than from a separately postulated local-invariance principle.

5.1 Scope of this section

We derive the unique compensating structure that resolves the local comparison problem of §4, **under the restriction that the local phase ambiguity is described by a single real-valued function** $\boldsymbol{\theta : \Lambda \longrightarrow \mathbb{R}}$. This corresponds physically to a single global phase symmetry — the Abelian U(1) case. The non-Abelian extension, in which θ is replaced by a Lie-algebra-valued matrix function, is a structurally distinct theorem requiring additional input on internal closure indices; see §10.

5.2 What the additional structure must do

The additional structure must accomplish two things simultaneously:

1. Define a closure-invariant rate of change of $|\psi(x)\rangle$ across the substrate.
2. Transform under local re-phasing in a way that exactly compensates the anomalous derivative term identified in §4.2.

We denote this structure by $A_\mu(x)$ — a real-valued vector field on the substrate — and define the *covariant derivative*

$$D_\mu|\psi\rangle := (\partial_\mu + iA_\mu)\,|\psi\rangle\ .$$

The role of A_μ is to absorb the anomalous gradient term so that $D_\mu|\psi\rangle$ transforms as a closure-invariant object.

5.3 The gauge transformation rule is forced

We require that under $|\psi\rangle \longmapsto e^{i\theta(x)}\,|\psi\rangle$, the covariant derivative transforms covariantly:

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

Expanding the left-hand side with the transformed connection A'_μ:

$$(\partial_\mu + iA'_\mu)(e^{i\theta}\,|\psi\rangle) = e^{i\theta}\,[\ \partial_\mu|\psi\rangle + i(\partial_\mu\theta)|\psi\rangle + iA'_\mu\,|\psi\rangle\ ]\ .$$

Equating with the right-hand side $e^{i\theta}(\partial_\mu + iA_\mu)|\psi\rangle$ and comparing coefficients of $|\psi\rangle$:

$$i(\partial_\mu\theta) + iA'_\mu = iA_\mu\ ,$$

which rearranges to

$$A_\mu \longrightarrow A'_\mu = A_\mu - \partial_\mu\theta\ .$$

This is the **gauge transformation rule**. In the standard treatment it is *postulated*; in the present derivation it is the unique transformation rule consistent with closure-invariant comparison under a single-valued local phase ambiguity.

5.4 Theorem: closure-consistent comparison forces a unique connection (Abelian case)

Theorem 5.1 (Forced Connection — Abelian Case). *Let $|\psi(x)\rangle$ be a vector representative of a closure-state ray field on the distinguishability substrate Λ, with local phase ambiguity described by a smooth real-valued function $\theta : \Lambda \longrightarrow \mathbb{R}$. Then the unique structure (up to constant rescaling) that (i) defines a closure-invariant local derivative on closure-state vector representatives and (ii) transforms locally to absorb arbitrary smooth re-phasings is a real-*

valued one-form A_μ(x) with covariant derivative D_μ = ∂_μ + iA_μ and gauge transformation rule A_μ ⟶ A_μ − ∂_μθ.

Proof.

Existence. The construction of §§5.2–5.3 satisfies (i) and (ii).

Uniqueness. We establish uniqueness through a sequence of forced reductions, each justified explicitly.

Step 1 — Locality. We restrict attention to compensating structures that are local in x — i.e., that depend on data in an arbitrarily small neighbourhood of x. This is a structural requirement of any derivative-like operator, and we adopt it as part of the definition of "compensating structure for ∂_μ". A non-local compensating structure would correspond to a different mathematical object (e.g., a Wilson line) and is treated separately in §9.2.

Step 2 — Linearity in |ψ⟩. The covariant derivative D_μ enters physical dynamics through the kinetic term of the Hamiltonian. By the imported reconstruction (§2.1), reversible dynamics are unitary on *ℋ*, generated by a Hamiltonian H that must act linearly on *ℋ* in order for time evolution e^(−iHt) to preserve superpositions. Since D_μ enters H through polynomial combinations such as D_μ D^μ (we restrict to this universal case), and H must act linearly on |ψ⟩, the covariant derivative D_μ must also act linearly on |ψ⟩; a non-linear D_μ would propagate non-linearity into H and break the unitarity of dynamics. The compensating structure A_μ acting via D_μ = ∂_μ + iA_μ must therefore act linearly on |ψ⟩, ruling out structures that depend on |ψ⟩ nonlinearly (e.g., self-interaction terms involving ⟨ψ | ψ⟩).

Step 3 — Hermiticity. Under closure-invariant re-phasing |ψ⟩ ⟼ e^(iθ)|ψ⟩ with θ ∈ ℝ, the anomalous derivative term takes the form i(∂_μθ)|ψ⟩, which is purely imaginary on the right-hand side relative to |ψ⟩. The compensating structure must absorb a term of this exact form. If A_μ had an anti-Hermitian part A_μ = A_μ^H + iA_μ^A (with A_μ^H, A_μ^A both Hermitian), the absorption equation iA_μ ⟶ iA_μ + i(∂_μθ) would require ∂_μθ to be complex, contradicting the reality of θ established in §3.1. Hermiticity of A_μ is therefore *forced* by the reality of the phase parameter, independently of any appeal to dynamics.

Step 4 — Index structure. The compensating term must contract with dx^μ to produce a derivative-like quantity, so it carries one covariant spacetime index. By Steps 2 and 3 and the Abelian restriction (single real-valued θ), the structure has no internal index; it is simply a real one-form A_μ(x).

Step 5 — Coefficient. The coefficient of A_μ in the covariant derivative D_μ = ∂_μ + iA_μ is fixed up to constant rescaling by requiring that the transformation rule A_μ ⟶ A_μ − ∂_μθ exactly cancel the anomalous i(∂_μθ)|ψ⟩ term. Any other coefficient produces residual anomalous terms or absorbs more than required.

The five steps determine A_μ uniquely up to constant rescaling. ■

The theorem establishes that A_μ is not a postulate but a structural identity *given the Abelian restriction*: under a single global phase ambiguity, the only object that can render $\partial_\mu|\psi\rangle$ closure-invariant is a real Hermitian one-form. The non-Abelian generalisation, in which the phase ambiguity is matrix-valued, is a different theorem; we discuss it in §10.

6. Gauge Redundancy as Closure Equivalence

We now address a conceptual point that is often obscured in standard treatments of gauge theory.

6.1 What gauge "symmetry" actually is

In the standard formulation, gauge symmetry is presented as a *symmetry* — a transformation of the fields that leaves the action invariant. This framing invites the question: why do these symmetries exist? In the VERSF derivation, the answer is structurally different.

The transformation

$$|\psi\rangle \mapsto e^{(i\theta(x))} |\psi\rangle , \quad A_\mu \longrightarrow A_\mu - \partial_\mu\theta$$

is not a symmetry of physical states. It is the **identity transformation on physical states**, expressed in the language of vector representatives. The rays $[\psi(x)]$ before and after the transformation are identical; the transformation merely changes the arbitrary phase representative chosen at each substrate point.

What standard treatments call "gauge symmetry" is therefore better described as **gauge redundancy** — the mathematical multi-valuedness of vector representatives over the same physical ray field.

6.2 Why this matters

This re-framing has three structural consequences:

1. **There is no question "why does the symmetry exist?"** It is not a symmetry of physical states; it is a structural feature of the vector-to-ray map. The ambiguity is built into the relational closure ontology.
2. **The "gauge group" is the ambiguity group of phase choices.** In the Abelian case derived here it is U(1) (or $\mathbb{R}$ if one drops periodicity), corresponding to the freedom of choosing a phase at each point. Non-Abelian gauge groups arise when there are larger internal closure ambiguities; see §10.
3. **Gauge-fixing is not symmetry-breaking.** Choosing a representative phase function at every substrate point does not break any physical symmetry — it merely fixes a specific labelling of the rays. Physical predictions are invariant under any such choice.

This re-framing is the substantive content of the consolidation flagged in §1.1: the same equations are forced by the standard treatment and by VERSF, but the *motivational structure* is different. In the standard treatment, gauge invariance is a deep principle requiring justification; in the VERSF account, it is a structural feature of the relational closure ontology imported in §2.1, with no further justification needed.

---

7. Curvature as the Unique Local Closure-Invariant Observable

We have established that A_μ is closure-equivalent under $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$. The connection itself is therefore *not* a closure-invariant observable — it changes under closure-equivalent re-phasings. We now ask: what local quantity built from A_μ *is* closure-invariant?

7.1 The curvature

Consider the commutator of covariant derivatives:

$$[D_\mu, D_\nu]\,|\psi\rangle = (\partial_\mu + iA_\mu)(\partial_\nu + iA_\nu)|\psi\rangle - (\mu \leftrightarrow \nu)\ .$$

Expanding and using $\partial_\mu\partial_\nu = \partial_\nu\partial_\mu$:

$$[D_\mu, D_\nu]\,|\psi\rangle = i(\partial_\mu A_\nu - \partial_\nu A_\mu)\,|\psi\rangle\ .$$

Define the **curvature** (or *field strength*):

$$F_{\mu\nu} := \partial_\mu A_\nu - \partial_\nu A_\mu\ .$$

Direct computation under $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$ gives $F_{\mu\nu} \longrightarrow F_{\mu\nu}$, since the additional terms $\partial_\mu\partial_\nu\theta$ and $\partial_\nu\partial_\mu\theta$ are equal and cancel. The curvature is therefore **closure-invariant**: it does not depend on the phase choice.

7.2 Theorem: $F_{\mu\nu}$ generates the local closure-invariants

Theorem 7.1 (Local Closure-Invariants). *Let $P[A, \partial A, \partial^2 A, \ldots]$ be a polynomial functional of A_μ and its derivatives that is local at the point x (depends only on the values of A and finitely many derivatives at x) and is closure-invariant in the sense of Definition 3.2. Then P is expressible as a polynomial in $F_{\mu\nu}$ and its ordinary derivatives $\partial_\lambda F_{\mu\nu}$, $\partial_\lambda\partial_\kappa F_{\mu\nu}$, ...*

Proof. The result is the algebraic Poincaré lemma for Abelian gauge theories: the local cohomology of A_μ in the variational bicomplex, restricted to closure-invariant polynomials, is generated by $F_{\mu\nu}$ and its derivatives. We sketch the inductive structure for motivation and then cite the standard cohomological closure.

Inductive sketch. For polynomials of derivative order N = 0 (no derivatives of A_μ), closure invariance under A_μ ⟶ A_μ − ∂_μθ requires P[A_μ − ∂_μθ] = P[A_μ] for arbitrary smooth θ. Choosing θ such that ∂_μθ at x takes any prescribed value forces P to be independent of A_μ at x — i.e., P must be a constant.

For polynomials with derivatives of A_μ, the variation under A_μ ⟶ A_μ − ∂_μθ produces, at each derivative order, a constraint that can only be satisfied if the corresponding derivative of A_μ enters P only through the antisymmetric combination ∂_μA_ν − ∂_νA_μ = F_μν, possibly with further derivatives applied. The symmetric part of any derivative of A_μ transforms by a non-vanishing total derivative of θ and cannot be made closure-invariant.

Cohomological closure. The full theorem — that *every* polynomial closure-invariant local functional, including cross-terms between derivatives of different orders, is expressible in F_μν and its ordinary derivatives — is the algebraic Poincaré lemma for the Abelian gauge complex. A complete treatment is given in Henneaux–Teitelboim, *Quantization of Gauge Systems* (1992), Chapter 12, and in Anderson, *The Variational Bicomplex* (1989), with the underlying mathematical structure being BRST cohomology in the trivial sector. The cohomological result handles cross-terms between derivative orders that the inductive sketch does not cleanly resolve, and we adopt it here without re-deriving. ■

7.3 Physical content

Theorem 7.1 has two structural consequences. First, all *local* polynomial closure-invariants built from the gauge field must be expressible as functionals of F_μν and its ordinary derivatives — not of A_μ directly. Second, the simplest non-trivial closure-invariant scalar density built from F_μν is F_μν F^μν, which under the Lorentz lift of [VERSF–PROTO] becomes the Lorentz-invariant kinetic term in the Yang–Mills Lagrangian. The dynamical content of the gauge field — the action principle that selects F_μν F^μν specifically — is the subject of [VERSF–CCB] and lies outside the scope of this paper.

The Aharonov–Bohm effect (§9) shows that the *complete* closure-invariant content of the gauge sector is *not* exhausted by Theorem 7.1: non-local closure-invariants exist (in particular, holonomies around non-contractible loops), and these carry information not reducible to local functionals of F_μν. We discuss this distinction precisely in §9.

---

8. Spin-1 of the Gauge Quantum (under SO(3))

We now derive the spin of the field quantum associated with A_μ *under spatial isometries*. The Lorentz-spin assignment is a structurally distinct claim that we address separately in §8.5.

8.1 The connection is a one-form

Under spatial transformations $g \in SO(3)$ acting on the substrate, the closure state transforms as $|\psi(x)\rangle \longmapsto U(g)\,|\psi(g^{-1}x)\rangle$, where $U(g)$ is the relevant unitary representative on $\mathscr{H}$.

The covariant derivative $D_\mu|\psi\rangle$ must transform consistently, which requires that the index μ on A_μ transform as a vector index under SO(3):

$$A_\mu(x) \longrightarrow (g^{-1})_\mu{}^\nu A_\nu(g^{-1} x) .$$

This is the transformation rule of a **one-form** (covector field). The connection is therefore not a scalar (which would carry $j = 0$) and not a higher-rank tensor.

8.2 SO(3)/SU(2) representation theory

By the result of [VERSF–SPIN], the projective unitary representations of SO(3) lift to true unitary representations of its universal cover SU(2). The integer-j representations are exactly those that descend to genuine linear representations of SO(3); the half-integer-j representations are faithful on SU(2) and only projective on SO(3). Dimensions are $2j + 1$.

A one-form has three components in three-dimensional space, corresponding to $\dim = 3 = 2j + 1$ with $j = 1$:

$j = 1$: vector representation, dimension 3, descending to a genuine SO(3) representation as an integer-spin case.

The connection A_μ therefore lies in the $j = 1$ representation under SO(3), and any quantum carrying its degrees of freedom carries spin 1 in the spatial-isometry sense.

8.3 Theorem: SO(3) spin assignment

Theorem 8.1 (SO(3) Spin of the Gauge Quantum). *Within VERSF, the field quantum associated with the closure connection A_μ derived in Theorem 5.1 transforms in the $j = 1$ (vector) representation of SO(3) under spatial isometries.*

Proof. By Theorem 5.1, A_μ is a real-valued one-form on the distinguishability substrate. By §8.1, one-forms transform under SO(3) as vectors. By §8.2 and [VERSF–SPIN], the vector representation is the $j = 1$ irreducible representation of SO(3). The associated quantum therefore carries SO(3) spin 1. ■

8.4 Caveat: kinematic only, not dynamical

The theorem above establishes the *kinematic* spin of the gauge quantum but does not address whether it is *massive* or *massless*. In standard quantum field theory, unbroken gauge invariance forces the gauge quantum to be massless, with two physical polarisations rather than three — the Higgs mechanism is the standard way to generate masses while preserving gauge invariance at the Lagrangian level. The mass question is a dynamical one that depends on the action and the symmetry-breaking content of the theory, both of which lie outside the scope of the present paper.

8.5 The Lorentz-spin assignment is deferred

The spin assignment of physical gauge bosons is a Lorentz-covariant claim, not just a spatial-isometry claim. The two are related but not identical. Under the Lorentz group, A_μ transforms as a four-vector with four components, of which only two are physical polarisations for a massless gauge boson — the longitudinal and timelike modes are pure-gauge artefacts and decouple from physical observables. The spin-1 assignment of physical massless gauge bosons therefore depends on the structure of the *Poincaré* representation theory and the gauge-fixing procedure, not just on SO(3).

VERSF treats Lorentz invariance as itself emergent from proto-time, with the Poincaré group arising as a structural consequence of irreversible commitment events on the closure substrate. The companion paper [VERSF–PROTO] derives this emergence and the associated representation theory. The Lorentz-spin lift of Theorem 8.1 — i.e., the demonstration that the SO(3) $j = 1$ assignment derived here is the kinematic prerequisite for the four-vector Lorentz character of the physical gauge boson — is a structural step that must be carried out within that companion framework rather than the present one.

We therefore explicitly defer the Lorentz-spin claim to [VERSF–PROTO] and limit the result of this paper to the SO(3) $j = 1$ statement of Theorem 8.1. The Theorem 8.1 result is the kinematic prerequisite for the Lorentz lift but does not constitute it.

9. Worked Example: U(1) and the Aharonov–Bohm Effect

We make the abstract result concrete by working through the simplest case — a single phase ambiguity $\theta \in U(1)$ — and showing that the framework predicts the Aharonov–Bohm effect. The example also makes precise the statement at the end of §7.3 about local versus non-local closure-invariants.

9.1 The U(1) closure connection

In the simplest case, the local phase choice is parameterised by a single real number $\theta(x) \in \mathbb{R}$, with θ and $\theta + 2\pi$ identified. The phase ambiguity group is U(1). The connection A_μ is a real-valued one-form, the covariant derivative is

$$D_\mu = \partial_\mu + iA_\mu ,$$

and the curvature

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$$

corresponds to the electromagnetic field strength tensor in the standard identification.

9.2 Holonomy: a non-local closure-invariant

Consider parallel transport of a closure state around a closed loop $\gamma \subset \Lambda$. The state acquires a phase

$$\Phi_\gamma = \oint_\gamma A_\mu \, dx^\mu ,$$

known as the **holonomy** of the connection around γ. Under the gauge transformation $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$, the holonomy transforms as

$$\Phi_\gamma \longrightarrow \Phi_\gamma - \oint_\gamma \partial_\mu\theta \, dx^\mu = \Phi_\gamma ,$$

where the integral of the total derivative around the closed loop vanishes for any single-valued θ.[1] The holonomy is therefore closure-invariant, but it is a *non-local* invariant — it depends on the values of A_μ along the entire loop γ, not just at a single point.

By Stokes' theorem, if γ is the boundary of a simply-connected surface S:

$$\Phi_\gamma = \oint_\gamma A_\mu \, dx^\mu = \int_S F_{\mu\nu} \, dS^{\mu\nu} .$$

When γ bounds a region of zero curvature, $\Phi_\gamma = 0$. But when γ encloses a region of non-zero curvature *without passing through it* — for example, when the topology forbids a simply-connected bounding surface — Φ_γ can be non-zero even though $F_{\mu\nu} = 0$ along γ itself.

This is the precise sense in which the closure-invariant content of the gauge sector exceeds Theorem 7.1: the local invariants are exhausted by $F_{\mu\nu}$ and its derivatives, but *non-local* invariants (holonomies around non-contractible loops) carry independent information.

[1] *Footnote on large gauge transformations.* For multi-valued θ — large gauge transformations winding non-trivially around the U(1) target — the integral $\oint_\gamma \partial_\mu\theta \, dx^\mu$ shifts by 2π integer multiples rather than vanishing. This is the topological content of the U(1) periodicity: the holonomy Φ_γ is well-defined modulo 2π, with shifts by 2π corresponding to closure-equivalent re-labellings of the phase representatives. For the single-valued re-phasings considered in the main argument, the holonomy is fully invariant; the modulo-2π subtlety does not affect the closure-invariance of the holonomy as a physical quantity, only the choice of branch.

9.3 The Aharonov–Bohm prediction

Consider a closure state passing through a two-slit interferometer with a region of non-zero $F_{\mu\nu}$ (a magnetic flux tube) confined between the slits. The state at the screen is a superposition of the two paths γ_1 (above the flux) and γ_2 (below the flux). The relative phase between the two paths is

$$\Delta\Phi = \oint_{(\gamma_1 - \gamma_2)} A_\mu \, dx^\mu = \int_S F_{\mu\nu} \, dS^{\mu\nu} = \Phi_{\text{enclosed}} ,$$

where the surface S is bounded by the difference of the two paths and intersects the flux tube. This is the **Aharonov–Bohm phase**. The interference pattern at the screen shifts by an amount proportional to the enclosed flux $\Phi_{enclosed}$, even though the closure state never passes through any region of non-zero $F_{\mu\nu}$ along its actual trajectory.

9.4 Empirical anchor

The Aharonov–Bohm effect was predicted by Aharonov and Bohm (1959) and observed in a series of increasingly stringent experiments, beginning with Chambers (1960) and culminating in the Tonomura electron-holography experiments (1986), which confirmed the effect with magnetic flux fully enclosed inside a superconducting torus that the electrons could not penetrate.

The empirical content is precise. The observed phase shift confirms that the closure-invariant content of the gauge sector includes non-local quantities — holonomies around non-contractible loops — that are not exhausted by local functionals of $F_{\mu\nu}$ alone. The connection A_μ has physical content beyond what can be reconstructed point-by-point from the local curvature.

The VERSF derivation reproduces this prediction by construction: the connection arises as the structure required to render local comparison closure-consistent, and its holonomy is a non-local closure-invariant. Theorem 7.1 limits *local* polynomial closure-invariants to functionals of $F_{\mu\nu}$, but does not exhaust the closure-invariant content of the connection — the Aharonov–Bohm effect is the canonical empirical demonstration of this distinction.

---

10. Open Problems for the Non-Abelian Extension

The Abelian derivation above is the result of this paper. We now state — explicitly as **open problems**, not as derived results — what would be required to extend the framework to non-Abelian gauge structure. This section is descriptive of the research programme, not contributory to the present derivation.

10.1 Internal closure degrees of freedom

The natural generalisation is to allow closure states to carry **internal indices**:

$$|\psi(x)\rangle \longrightarrow |\psi^a(x)\rangle\ ,\ a = 1, \ldots, N\ ,$$

where the index a labels distinguishable internal closure modes. Local re-phasings would then be matrix-valued:

$$|\psi^a(x)\rangle \longmapsto U(x)^a{}_b\, |\psi^b(x)\rangle\ ,\ U(x) \in G\ ,$$

with G a Lie group acting unitarily on the internal index. The Abelian case is the special case $G = U(1)$ with $N = 1$.

If this structure is granted, the local-comparison argument of §§4–5 generalises to a matrix-valued connection $A_\mu^A T_A$ with covariant derivative $D_\mu = \partial_\mu + iA_\mu^A T_A$ and the standard non-Abelian curvature

$$F_{\mu\nu}^A = \partial_\mu A_\nu^A - \partial_\nu A_\mu^A - f^{(ABC)} A_\mu^B A_\nu^C .$$

This is Yang–Mills in the standard form. **However, this construction imports the matrix structure rather than deriving it**, and the present paper does not claim to derive any of the following:

10.2 What is not derived

The non-Abelian extension as sketched above takes for granted four things that VERSF would need to derive separately:

1. **That internal closure indices exist at all.** The Abelian derivation requires only the global phase ambiguity built into the ray-to-vector lift. For non-Abelian structure, additional internal degrees of freedom must be present in the closure ontology. Whether the $K = 7$ simplicial closure structure of [VERSF–KSEVEN] generates such internal indices, and how, is an open question.
2. **That internal indices carry Lie-group structure rather than discrete or other algebraic structure.** A priori, internal closure modes could carry any group structure — discrete (finite group), continuous (Lie group), or more general (groupoid, semigroup). The continuous-Lie-group case is what produces standard Yang–Mills theory; deriving this restriction from VERSF principles is a separate problem.
3. **Which specific gauge groups appear in nature.** The Standard Model gauge group $SU(3) \times SU(2) \times U(1)$ has a specific structure. Whether VERSF derives this exact group, or only constrains the space of allowed groups, is an open question. The $K = 7$ hexagonal closure structure of [VERSF–KSEVEN] is a candidate substrate for fixing the gauge group content; this is an active research direction.
4. **The non-Abelian analogue of Theorem 7.1.** The corresponding result in the non-Abelian case — that gauge-invariant local polynomials in A_μ^A and its derivatives are generated by the non-Abelian curvature $F_{\mu\nu}^A$ (transforming covariantly as $F_{\mu\nu} \longrightarrow U F_{\mu\nu} U^{(-1)}$ under $U(x) \in G$) and its covariant derivatives — is also a known cohomological theorem [Henneaux–Teitelboim 1992, Ch. 12; Dragon 1996]. A VERSF derivation of this result, as opposed to its importation from the standard cohomological literature, is a fourth open question for the non-Abelian programme. The §7.3 and §12.4 references to $F_{\mu\nu} F^{\mu\nu}$ as the simplest kinetic term tacitly invoke the analogous non-Abelian result for the trace $\mathrm{tr}(F_{\mu\nu} F^{\mu\nu})$; for the present Abelian-only paper this is unproblematic, but it is part of what a complete non-Abelian VERSF derivation would have to establish.

These four questions are stated here as a roadmap for the non-Abelian programme. The present paper does not contribute to their resolution; it establishes only the Abelian connection and notes the structural pathway.

---

11. What This Paper Does Not Claim

To restate the scope explicitly, the following are *not* claimed as results of this paper:

- **Specific gauge groups.** The Abelian U(1) case is derived; non-Abelian groups are sketched in §10 as open problems. Which specific groups appear in nature is an open question.
- **Coupling constants.** Numerical strengths of gauge interactions are not derived here.
- **Full Yang–Mills dynamics.** The Yang–Mills Lagrangian $L = -\frac{1}{4} F_{\mu\nu}^A F^{\mu\nu}_A$ is the simplest closure-invariant action constructible from $F_{\mu\nu}$, but the variational principle that selects it is the subject of [VERSF–CCB], not the present paper.
- **Lorentz-spin assignment of physical gauge bosons.** The result of this paper is the SO(3) $j = 1$ assignment of Theorem 8.1. The Lorentz lift, which determines physical polarisations and the massless-gauge-boson structure, is deferred to [VERSF–PROTO].
- **Spin–statistics for gauge bosons.** That gauge bosons obey Bose–Einstein statistics is consistent with their integer spin, but the spin–statistics theorem requires relativistic + locality input not provided here. [VERSF–PROTO] supplies the relativistic causal structure; the locality input beyond that is a separate question for future work.
- **Mass generation.** The Higgs mechanism, electroweak symmetry breaking, and gauge boson masses are dynamical questions outside the scope of this paper (§8.4).
- **Connection to gravitation.** The spin-2 sector [VERSF–GRAV] is an analogous but distinct construction involving back-reaction of closure events on the substrate metric, not a connection on the closure bundle.
- **Mathematical novelty in the gauge sector.** The technical content of §§5–8 overlaps with standard Abelian gauge theory; the contribution is interpretive consolidation, not new mathematics. See §1.1 and §12.2.

---

12. Discussion

12.1 Relation to the fiber bundle picture

The VERSF derivation can be re-cast in the standard fiber-bundle language of differential geometry: the rays $[\psi(x)]$ form a U(1) principal bundle over the substrate, the connection A_μ is a connection one-form on this bundle, and $F_{\mu\nu}$ is its curvature. What VERSF adds is a *physical motivation* for why this bundle structure exists: it is forced by the local-comparison problem of §4, which is itself forced by the relational closure ontology imported in §2.1. The fiber bundle is not assumed; it is derived as the natural structure on closure-state ray fields.

12.2 Relation to standard gauge-principle derivations

The standard pedagogical derivation of gauge fields proceeds as follows: postulate a global symmetry, "make it local", observe that the kinetic term is no longer invariant, introduce a connection to restore invariance. This procedure is mathematically identical to the VERSF derivation in the Abelian case — the same equations are written down, the same transformation rule is forced, and the same curvature is identified.

The contribution of this paper is *interpretive consolidation*, not new mathematics. The standard derivation begins with a *postulated* symmetry; the VERSF derivation begins with the *closure-equivalence ambiguity* (a structural feature of the ontology, not a postulated symmetry) and arrives at the same equations. The two paths converge mathematically. Within a reconstruction programme that aims to derive physics from a small ontological base, this consolidation is genuinely valuable — it removes the gauge principle from the list of inputs to physics, replacing it with the relational closure ontology that is already required for $\mathscr{H}$ and the Born rule. But it is not a claim of mathematical novelty in the gauge sector.

The appropriate frame for the result is: "given the relational closure ontology, the standard gauge-principle equations are no longer separately postulated but are forced by the structure already present." This is a strict reduction in the postulate count of physics, not a different theory.

12.3 Relation to other reconstruction programmes

Operational reconstructions of quantum theory — Hardy's five reasonable axioms, the Masanes–Müller information-theoretic reconstruction, and the Chiribella et al. categorical-quantum programme — derive the structure of single-system quantum mechanics (Hilbert space, Born rule, composition rules) from operational primitives. They do not derive gauge structure, but neither do they aim to: their target is the kinematic structure of single-system quantum theory, and the gauge sector lies outside that target.

A more substantive comparison is therefore the question of *what would have to be added* to those programmes to derive gauge structure. The VERSF answer is specific: the derivation in §§4–5 above requires (i) a substrate Λ over which states vary, and (ii) a relational closure ontology that makes ray representation physically primary and vector representation a derived (and ambiguous) auxiliary. The operational reconstructions take the single-system Hilbert space as primary and do not commit to either a substrate or a ray-versus-vector ontology in the structural sense VERSF requires; the Abelian gauge derivation is therefore not naturally available within them.

A different and more developed comparison is with **algebraic quantum field theory** (AQFT), in which gauge structure is *recovered* — not postulated — via the Doplicher–Haag–Roberts (DHR) analysis of superselection sectors over a Lorentz-covariant net of local von Neumann algebras. The DHR construction is itself non-trivial: it shows that under suitable assumptions on the algebraic net and its sectors, a compact gauge group and a field algebra implementing it can be reconstructed from the observables alone, without being put in by hand. AQFT and VERSF therefore both derive gauge structure rather than postulate it; the comparison concerns the structural level at which their respective inputs sit. AQFT takes Lorentz covariance and microcausality as primitive operational requirements on the algebraic net; VERSF treats both as derived — Lorentz invariance from proto-time [VERSF–PROTO], microcausality from closure-event ordering. AQFT's inputs are therefore *kinematically richer* than VERSF's at the operational level (it imports relativistic structure that VERSF derives) but *ontologically thinner* (no substrate, no closure ontology, no ray-versus-vector distinction). Whether AQFT-with-emergent-relativity reduces to VERSF, or VERSF-with-completed-derivations reduces to an AQFT-style local-algebra picture, or the two are genuinely complementary routes that converge only at the level of empirical predictions, is itself an open methodological question. We note it here as a comparison worth developing, not as one this paper resolves.

12.4 Connection to the BCB Lagrangian formulation

The closure-invariant Lagrangian density built from $F_{\mu\nu}$ and its contractions provides the action principle for the gauge sector. In the [VERSF–CCB] formulation, this density appears as one term in the Bit Conservation and Balance (BCB) action, alongside the matter and gravitational contributions. The present paper establishes the *kinematic* origin of the gauge field; the [VERSF–CCB] paper establishes the *dynamic* origin of its action principle. The two papers are complementary: this one shows that A_μ exists and what its symmetries are, the BCB paper shows how its dynamics are constrained.

---

13. Conclusion

We have shown that the Abelian gauge connection and its kinematic spin-1 character under SO(3) are *forced consequences* of comparing closure states across the distinguishability substrate, given the structural inputs of §2:

1. Physical states are rays in $\mathscr{H}$ (closure ontology — imported from VERSF–FSN).
2. Local phase representatives at different substrate points are independently chosen — a structural feature of the ray-to-vector lift.
3. Closure-consistent comparison of nearby states requires absorbing the local-phase ambiguity, which forces the introduction of a connection A_μ with the unique transformation rule $A_\mu \longrightarrow A_\mu - \partial_\mu\theta$ (Theorem 5.1, conditional on a single global phase ambiguity).
4. The curvature $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ generates the local closure-invariants built from A_μ (Theorem 7.1, with cohomological closure cited from Henneaux–Teitelboim).
5. The connection transforms as a one-form under SO(3) and therefore carries the $j = 1$ representation (Theorem 8.1). The Lorentz-spin lift is deferred to [VERSF–PROTO].

The U(1) case is derived rigorously and confirmed empirically by the Aharonov–Bohm effect, which also shows that the closure-invariant content of the gauge sector includes non-local quantities (holonomies on non-contractible loops) beyond the local invariants of Theorem 7.1.

The non-Abelian extension is flagged in §10 as a structured set of open problems, not a derived result.

The technical content of §§5–8 overlaps with the standard Abelian gauge-principle derivation; the contribution of this paper is *interpretive consolidation* — the same equations now flow from a structural feature of the closure ontology rather than from a separately postulated local-invariance principle. Within a reconstruction programme that aims to derive physics from a small ontological base, this consolidation removes the gauge principle from the list of inputs to physics.

Gauge structure thus joins $\mathscr{H}$, the Born rule, unitary dynamics, and the half-integer spin sector on the list of features that earlier appeared as quantum-theoretic postulates and now appear as derived consequences (in the Abelian case) of distinguishability physics. The conceptual content of the result is this: **Abelian gauge fields are, structurally, the geometry of comparing relational closure states.**

The next steps in the VERSF reconstruction sequence are: the Lorentz lift [VERSF–PROTO]; the non-Abelian extension and gauge-group selection (open problems §10); the dynamical origin of the Yang–Mills action [VERSF–CCB]; and the spin-2 sector [VERSF–GRAV].

References (selected)

- 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–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: *Gravity from Fold-Density Gradients: The Spin-2 Closure Back-Reaction.*
- I. M. Anderson, *The Variational Bicomplex*, technical report, Department of Mathematics, Utah State University, Logan, Utah (1989). Available online at https://digitalcommons.usu.edu/dg_pubs/.
- N. Dragon, *BRS Symmetry and Cohomology*, arXiv:hep-th/9602163 (1996).
- M. Henneaux and C. Teitelboim, *Quantization of Gauge Systems*, Princeton University Press, Princeton, NJ (1992), ISBN 978-0-691-08775-0. Chapter 12 covers the local cohomology of gauge theories.
- Y. Aharonov and D. Bohm, *Significance of Electromagnetic Potentials in the Quantum Theory*, Phys. Rev. **115** (1959) 485–491.
- R. G. Chambers, *Shift of an Electron Interference Pattern by Enclosed Magnetic Flux*, Phys. Rev. Lett. **5** (1960) 3–5.

- A. Tonomura, N. Osakabe, T. Matsuda, T. Kawasaki, J. Endo, S. Yano, and H. Yamada, *Evidence for Aharonov–Bohm Effect with Magnetic Field Completely Shielded from Electron Wave*, Phys. Rev. Lett. **56** (1986) 792–795.
- C. N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96** (1954) 191–195.
- E. P. Wigner, *Group Theory and Its Application to the Quantum Mechanics of Atomic Spectra*, Academic Press, New York (1959).