Matter Coupling and the Inertia Route in VERSF

Persistent U(1) Transport, Record-Current Coupling, and the Emergence of Source Dynamics

Keith Taylor *VERSF Theoretical Physics Programme*

General Reader Summary

The previous paper in this sequence established that the persistent transport sector of VERSF belongs to a fundamentally different dynamical class from the dissipative σ-sector. The σ-sector describes *admissibility restoration*: it smooths inconsistencies in the substrate much as diffusion smooths temperature gradients. The persistent sector, by contrast, preserves distinguishability and therefore belongs to the class of reversible dynamics. A separate admissibility theorem then showed that any reversible, composable, distinguishability-preserving evolution necessarily generates a Hamiltonian operator and unitary dynamics.

That result solved one structural problem but exposed another.

The persistent sector could now support reversible gauge-like transport, but it still lacked any *physical source structure*. The gauge field existed mathematically — as a cohomology class on the refinement-stable sector — but nothing generated it, nothing coupled to it, and nothing carried the inertia necessary for observable electromagnetic behaviour. The situation resembled a vacuum gauge field with no charged matter: structurally present, but physically inert.

This paper investigates the missing ingredient: **coupling**.

The central result is that once the persistent cohomological sector is coupled to the VERSF record-current sector — the substrate structure carrying irreversible commitment flow — the standard source structure of electromagnetism emerges naturally. The key interaction term turns out to be the familiar current–potential coupling

$$S_{\text{int}} = \int J^{\mu} A_{\mu}\, d^4x,$$

not because it is postulated from conventional field theory, but because it is the **unique** lowest-order admissible coupling compatible with locality, refinement persistence, gauge redundancy, distinguishability preservation, and Bit Conservation and Balance (BCB).

The paper shows that gauge redundancy on the persistent sector forces current conservation,

$$\partial_\mu J^\mu = 0,$$

and that the coupled persistent-sector dynamics therefore take Maxwell form,

$$\partial_\mu F^{\mu\nu} = J^\nu.$$

The physical reading is straightforward: the persistent cohomological sector supplies reversible transport structure; the record-current sector supplies the source structure to which the gauge field responds. Observable electromagnetic dynamics emerge not from the gauge sector alone, but from the *interaction* between reversible persistent transport and irreversible commitment flow. The gauge field itself remains massless, as it should — what coupling supplies is the dynamical response, not a mass term.

This paper does not claim to derive full quantum electrodynamics, fermion structure, charge quantization, or the wider Standard Model. It establishes something narrower but structurally important: persistent U(1) transport becomes physically dynamical only through coupling to the record-current sector, and this coupling reproduces the source structure of Maxwell electromagnetism with no remaining freedom at leading order.

The paper therefore closes the "inertia problem" identified previously. The persistent gauge sector does not require intrinsic mass-like terms inside the $K = 7$ constraint catalogue. Instead, observable dynamics arise through admissible coupling — structurally analogous to how gauge fields in conventional physics acquire physical content through interaction with matter.

Abstract

The previous VERSF paper on Hamiltonian persistent transport established that the refinement-persistent cohomological sector $H^1(\mathcal{G}(\Lambda))$ belongs to a reversible admissibility class distinct from the dissipative σ-sector. Combining the refinement-persistence theorems, the Wilson/cohomology identification, and the admissibility-generator theorem, the persistent sector was shown to support unitary Hamiltonian transport generated by a self-adjoint operator. However, the resulting gauge sector remained physically incomplete: it possessed reversible transport structure but no source dynamics, no inertia-carrying coupling, and no mechanism for observable electromagnetic behaviour.

This paper investigates the missing coupling structure and proves three theorems and one corollary.

Theorem 1 (Uniqueness of the admissible persistent-sector coupling, within the established substrate catalogue). Let A_μ denote the persistent cohomological transport potential on the refinement-stable sector $H^1(\mathcal{G}(\Lambda))$, and let J^μ denote the VERSF record-current associated with irreversible commitment flow. Under the *Catalogue Closure Theorem* (derived in §7.0 from the Single-Source Theorem and the admissibility principles; conditional only on the matter-sector characterisation of u^μ discussed in §15) — that J^μ is the unique admissible Lorentz-covariant

dimension-3 substrate vector, with no admissible dimension-2 antisymmetric tensor or primitive scalar at this order — and the admissibility principles of locality (P1), refinement persistence (P2), gauge redundancy (P3), distinguishability preservation (P4), and lowest-order closure consistency (P5), the unique admissible interaction Lagrangian density is

$$\mathscr{L}_{int} = -J^\mu A_\mu ,$$

modulo total derivatives. The result splits naturally into a classification (Theorem 1A: the admissible dimension-3 Lorentz vectors are $V_3 = \{J^\mu, \partial_\nu F^{\nu\mu}, \partial_\nu \tilde{F}^{\nu\mu}\}$, the latter two reducing to kinetic-renormalisation and identically-zero respectively) and a selection (Theorem 1B: only $J^\mu A_\mu$ survives as a genuinely new leading-order interaction). Higher-derivative variants are excluded by a refinement-persistence lemma giving suppression $(k \cdot a_n)^{2m}$ per $2m$ extra derivatives. Any future relaxation of the inputs to the Catalogue Closure Theorem — specifically the SST exhaustiveness principle or the u^μ matter-sector condition — requires re-examining the enumeration.

Theorem 2 (Gauge redundancy ⇔ current conservation). Invariance of the interaction action under $A_\mu \to A_\mu + \partial_\mu \chi$ for arbitrary smooth χ holds if and only if $\partial_\mu J^\mu = 0$. The VERSF Bit Conservation and Balance principle therefore coincides structurally with the current-conservation law of the emergent gauge sector.

Theorem 3 (Source-coupled Maxwell transport). Variation of the admissible persistent-sector action

$$S[A] = \int \left(-\tfrac{1}{4} F_{\mu\nu} F^{\mu\nu} - J^\mu A_\mu \right) d^4x$$

with respect to A_μ yields

$$\partial_\mu F^{\mu\nu} = J^\nu ,$$

while the Bianchi identity $\partial_{[\alpha} F_{\beta\gamma]} = 0$ holds automatically because $F = dA$. The persistent cohomological transport sector therefore reproduces the inhomogeneous and homogeneous Maxwell equations once coupled to the record-current sector.

Proposition 4 (Resolution of the inertia question, corollary of Theorem 1). The persistent gauge sector requires no intrinsic mass-like term within the $K = 7$ constraint catalogue. Observable electromagnetic dynamics emerge through coupling between the reversible persistent transport sector and the irreversible record-current sector. The observable dynamical response of the gauge field is therefore *interaction-generated* rather than intrinsic to the $K = 7$ harmonic sector itself. This is a re-reading of the mass-term exclusion in Theorem 1B in physical language, rather than an independent mathematical result.

The paper discusses the implications for the VERSF unification programme, including the emergence of electromagnetism from refinement persistence, the relation between reversible and irreversible admissibility classes, and the next open problems: charge quantization, coupling constants, matter-sector structure, and relativistic completion.

Contents

---

1. Introduction

The previous synthesis paper in this programme established that the persistent cohomological sector carries Maxwell-form U(1) transport as the unique admissible reversible continuum dynamics under BCB/TPB and the admissibility class (B1)–(B4). That construction, however, remained *source-free*: the gauge sector possessed reversible transport structure but no microscopic coupling to the substrate record-current sector. The present paper addresses precisely this missing layer, deriving the unique admissible leading-order coupling between the persistent gauge sector and the irreversible commitment-flow dynamics. In that sense it is the natural Part II of the Maxwell-form persistent transport programme.

To recall the broader setting: the persistent transport sector of VERSF belongs to a reversible admissibility class distinct from the dissipative σ-sector. The σ-sector is governed by gradient-flow admissibility restoration and therefore generates dissipative or invariant dynamics only. The persistent harmonic sector, by contrast, preserves distinguishability and supports reversible composable evolution. By the admissibility-generator theorem, such evolution necessarily generates a self-adjoint Hamiltonian operator and unitary transport dynamics:

$U(t) = \exp(-iHt)$.

The previous synthesis resolved one structural ambiguity but sharpened another. The persistent sector now possessed refinement persistence, cohomological transport structure, Wilson-loop observables, and unitary Hamiltonian evolution — yet it still lacked physical source coupling, observable field dynamics, and any interaction-carrying structure capable of producing the dynamical response associated with electromagnetic phenomena. The present paper addresses this missing layer.

The central claim is that the missing structure emerges naturally once the persistent cohomological sector is coupled to the VERSF record-current sector. The resulting coupled dynamics reproduce Maxwell-form U(1) gauge transport at leading order, with no remaining adjustable structure at that order. The physical reading is layered: the persistent sector carries reversible transport structure, the record-current sector carries irreversible commitment flow, and electromagnetism emerges from their interaction. Within VERSF this is the first explicit construction in which a reversible admissibility class and an irreversible admissibility class jointly generate an observable physical dynamics rather than acting in separate regimes.

The electromagnetic branch of the programme. The relationship between this paper and the surrounding sequence can be summarised compactly:

Stage	Achievement
Three-class / Maxwell-form transport synthesis	Identifies the reversible Maxwell-form gauge sector on $H^1(\mathcal{G}(\Lambda))$ under (B1)–(B4)
Present paper	Identifies the unique admissible source coupling $-J^\mu A_\mu$ and the resulting Maxwell-form dynamics
Next paper (matter sector)	Microscopic identification of the carriers of J^μ; species decomposition; charge quantisation via $H^1(\mathcal{G}(\Lambda); \mathbb{Z})$
Later paper (quantisation)	Quantum field structure on the persistent sector and full coupled QED

The present paper sits in the second row. Its job is narrower than full electromagnetism but structurally indispensable: it establishes that any matter sector subsequently constructed in row three must couple to the persistent gauge sector via the unique term derived here, with no remaining freedom at leading order. The rows below depend on this; the row above supplies its starting point.

---

2. Notation and Conventions

Spacetime is taken to be a 4-dimensional Lorentzian continuum emerging from the refinement-stable limit of the substrate lattice Λ, with metric signature $(+, -, -, -)$. Greek indices $\mu, \nu, \ldots$ run over 0, 1, 2, 3. Repeated indices are summed.

We work in natural Heaviside–Lorentz units ($\hbar = c = \varepsilon_0 = 1$) throughout. All factors of μ_0 and 4π are absorbed into the field normalisation; no factor appears in front of J^ν in the field equations.

The persistent cohomological transport potential is denoted A_μ. The field strength is

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

which in differential-form language is $F = dA$. The two notations are used interchangeably below; §10 makes explicit the structural consequence of the form-language version.

The record current is denoted J^μ. The Bit Conservation and Balance (BCB) principle is taken in its continuum form

$$\partial_\mu J^\mu = 0 ,$$

with the discrete-lattice version recovered under refinement.

By "lowest-order admissible interaction" we mean: local in the continuum limit, of total mass dimension 4 in $d = 4$, linear in A_μ, polynomial in the substrate fields with no higher derivatives than those already required for kinetic structure.

By "gauge redundancy" we mean the freedom $A_\mu \rightarrow A_\mu + \partial_\mu \chi$ for arbitrary smooth χ, with admissible falloff at the boundary of the integration region (so that boundary terms vanish under integration by parts).

---

3. Reversible vs Irreversible Admissibility Classes

The VERSF transport sequence established two distinct dynamical admissibility classes; their structural separation is what allows the present coupling theorem to be stated cleanly.

Dissipative admissibility (σ-sector). The σ-sector obeys gradient-flow dynamics, admissibility restoration, parabolic continuum limits, and entropy-reducing equilibration. Its continuum limit is diffusion-like,

$$\partial_\tau \varphi = D \nabla^2 \varphi ,$$

and its role is to smooth admissibility inconsistencies.

Reversible admissibility (persistent harmonic sector). The persistent harmonic sector satisfies composability, reversibility, and distinguishability preservation. By the admissibility-generator theorem, such sectors generate

$$U(t) = \exp(-iHt) ,$$

with H self-adjoint. The persistent sector therefore belongs to a fundamentally different dynamical class from the σ-sector.

These two classes are not in tension. They are *complementary*: dissipative dynamics restore admissibility, reversible dynamics propagate the persistent residue. The present paper concerns how the second of these becomes physically observable through coupling to a third structure — the record current — which itself sits in an irreversible class but at the substrate level rather than the persistent harmonic level.

---

4. The Persistent Cohomological Sector

The refinement papers established the persistent observable sector as the first Čech-type cohomology of the constraint graph,

$$H^1(\mathcal{G}(\Lambda)) = C^1(\Lambda) / \mathrm{Im}(d_0) ,$$

with cocycles representing closed 1-chains and coboundaries representing the lattice gradient of 0-chains. Wilson-loop observables on the persistent sector take the continuum form

$$W(\gamma) = \exp(i \oint_\gamma A_\mu dx^\mu) ,$$

with A_μ the transport potential associated with the cohomology class.

The persistent sector therefore already possesses, before any coupling is introduced:

- gauge redundancy $A_\mu \to A_\mu + \partial_\mu \chi$, by definition of the cohomology quotient;
- Wilson-loop observables, gauge-invariant by Stokes' theorem;
- refinement persistence under the lattice coarsening map;
- Hamiltonian transport structure, by the admissibility-generator theorem.

What it lacks is *source coupling*. There is, as yet, no object in the persistent sector that plays the role of a charge density or current to which A_μ responds.

---

5. The Record-Current Sector

The VERSF substrate independently carries an irreversible commitment current J^μ. This current arises from three structurally distinct conditions:

- Bit Conservation and Balance (BCB), which enforces local conservation of committed informational content;
- the commitment-flow architecture, which makes irreversible record formation directional;
- admissibility restoration at the σ-sector level, which provides the macroscopic continuum limit in which J^μ becomes a smooth vector field.

The continuum form of BCB is precisely

$$\partial_\mu J^\mu = 0 \,,$$

which is the local conservation law associated with admissible commitment transport.

Partial derivation: continuity from substrate commitment flow. The conservation law above is more than an inherited axiom; it admits a partial substrate-level derivation that is worth sketching, as it sharpens what J^μ is and what it is not.

The Single-Source Theorem requires every observable to be a functional of the committed-record density $\rho(x, t)$. For a vector observable the simplest SST-admissible form is

$$J^\mu(x) = \rho(x)\, u^\mu(x) \,,$$

where $u^\mu(x)$ is a local commitment-flow four-vector — the rate and direction at which records are being committed at x. This is the substrate-level analogue of a fluid four-current $\rho\, u^\mu$. Dimensional bookkeeping: with J^μ of dimension 3 (as established in §6) and u^μ taken as a dimensionless commitment-flow four-vector (the substrate analogue of a normalised four-velocity), ρ has dimension 3, consistent with a substrate number density. This is the dimension assignment used throughout the sketch below; it should be re-verified against the explicit substrate definition of ρ when the Single-Source Theorem's normalisation conventions are propagated downstream.

Now BCB at the substrate level is the statement that committed records are *not destroyed*: once a commitment event is registered at x, the resulting contribution to ρ persists and can only be transported, not annihilated. In differential form this is

$$\partial_t \rho + \nabla \cdot (\rho\, v) = 0 \,,$$

i.e., the time-derivative of ρ in any small region equals the flux of commitment-carrying flow across the region's boundary. Covariantly,

$$\partial_\mu (\rho\, u^\mu) = \partial_\mu J^\mu = 0 \,.$$

The continuity equation is therefore *not* an independent postulate but the local form of substrate-level commitment conservation, expressed as a constraint on the $J^\mu = \rho u^\mu$ decomposition. This converts the conservation law from an inherited axiom into a derivation from substrate dynamics — *given* the decomposition $J^\mu = \rho u^\mu$.

The decomposition itself is the remaining open question. $J^\mu = \rho u^\mu$ is the simplest SST-admissible form, but it is not yet established as the *unique* such form, and u^μ does not yet have a clean substrate-level definition independent of J^μ. The matter-sector programme (§15) must supply both. The derivation above is therefore a sketch of how continuity should fall out of substrate flow rather than a complete first-principles derivation; the gap is the same matter-sector gap discussed in §15, viewed from a different angle.

What J^μ is and is not. It is the substrate-level commitment-flow vector. It is not assumed to be the electromagnetic four-current of any particular matter species. The claim of this paper is that *under the admissibility principles below*, the coupling of J^μ to A_μ reproduces the structural role played by the electromagnetic four-current in Maxwell theory. Whether J^μ further admits a decomposition into species-resolved matter currents is a separate matter-sector question, left for future work.

The record-current sector therefore naturally supplies a physically conserved local vector field with a clear substrate origin. This makes it the natural — indeed, given the field content, the only — candidate for source-coupling the persistent gauge sector.

6. Admissible Coupling Principles

We impose five admissibility principles on candidate couplings between the persistent sector and the substrate. Each is motivated by an earlier VERSF result and is stated here as a constraint on the interaction Lagrangian density $\mathscr{L}_{\text{int}}$.

(P1) Locality. $\mathscr{L}_{\text{int}}(x)$ depends only on field values at x and finitely many derivatives, in the continuum refinement limit.

(P2) Refinement persistence. $\mathscr{L}_{\text{int}}$ must descend consistently under the lattice coarsening map and converge to a well-defined continuum limit; equivalently, it must lie in the refinement-stable image of the persistence functor.

(P3) Gauge redundancy compatibility. The action $S_{\text{int}} = \int \mathscr{L}_{\text{int}}\, d^4x$ must be invariant under $A_\mu \to A_\mu + \partial_\mu \chi$ modulo boundary terms, for arbitrary smooth χ with admissible falloff.

(P4) Distinguishability preservation. The coupling must not generate distinguishability-violating evolution on the persistent sector when J^μ is treated as a prescribed external source. (The full coupled-system question — whether unitarity persists once J^μ is itself dynamical and responds to A_μ — depends on matter-sector structure and is deferred; see the unitarity remark closing §7 and the matter question in §15.)

(P5) Lowest-order closure consistency. Among the admissible terms, we take the one of lowest mass dimension and lowest order in the fields that yields non-trivial source coupling.

These five principles are sufficient to fix $\mathscr{L}_{int}$ up to an overall constant, as the next section shows.

7. Theorem 1 — Uniqueness of the Interaction Term (within the Established Substrate Catalogue)

The uniqueness claim of this section depends on a structural input about the K = 7 substrate catalogue: that the leading-order admissible field content available for coupling to A_μ is exhausted, in a precise sense, by J^μ alone. We named this the Catalogue Closure Assumption (CCA) in earlier drafts and treated it as supplied by the K = 7 dimensionality and closure-trichotomy results. The present version *derives* CCA from the Single-Source Theorem and the admissibility principles introduced in §6, leaving only one residual condition — the substrate-level definition of u^μ — that is the same matter-sector gap §15 already identifies as the dominant remaining open question. Section §7.0 carries out the derivation; §7.1 then restates CCA as a theorem; Theorems 1A and 1B follow.

§7.0 — Derivation of the Catalogue Closure Assumption

The CCA derivation proceeds in three steps. Lemma 7.0.1 establishes the primitive substrate field content from the Single-Source Theorem. Lemma 7.0.2 establishes a structural property of admissible dimension-1 vectors. Proposition 7.0.3 then derives the three CCA clauses by dimensional analysis applied to the primitive content.

A preliminary distinction matters. By *primitive substrate field* we mean an independent dynamical degree of freedom in the substrate ontology — a field whose existence is not reducible to functional dependence on another substrate field. By *composite field* we mean a functional of primitive fields and their derivatives. The catalogue CCA refers to is the catalogue of primitive substrate fields. Composite fields are not catalogue entries; they are constructions from catalogue entries, and their admissibility in the coupling Lagrangian is settled separately by (P1)–(P5). This distinction is essential because the Single-Source Theorem permits unboundedly many composite scalars (e.g., $\Box \ln(\rho/\rho_0)$) while admitting only one primitive scalar (ρ); both facts are needed, and conflating them produces apparent contradictions.

A second preliminary clarifies the admissibility condition itself. Throughout this section, *non-fractional-power admissibility* means analytic functional dependence on ρ — excluding branch-cut structures from fractional powers such as $\rho^{(1/3)}$, but admitting analytic functions including polynomials, logarithms, exponentials, and their composites. So $\Box \ln(\rho/\rho_0)$ is admissible as a composite scalar; $\rho^{(1/3)}$ is not. This is the working definition used in Lemma 7.0.2 case (i) and in Proposition 7.0.3 case (iv).

Lemma 7.0.1 (Primitive Substrate Field Exhaustion). Under the Single-Source Theorem and the Constitutive Principle of Exhaustiveness, the primitive substrate fields in the VERSF catalogue are exhausted by:

- ρ: the committed-record density, a scalar of mass dimension 3, established by SST Definition 1 as the unique scalar density carrying irreducible distinguishability;
- u^μ: the commitment-flow four-vector, dimensionless, conditional on the matter-sector programme's substrate-level definition (see §15).

No other primitive substrate fields exist.

Argument. SST Definition 1 establishes ρ as the unique primitive scalar at any dimension: any alternative scalar field must either reduce to ρ or fail to represent committed physical records. Reducible scalars are by definition not primitive; non-admissible alternatives lie outside the framework. Hence ρ exhausts the scalar primitives.

The Constitutive Principle of Exhaustiveness (SST §8) requires any admissible field to be functionally reducible to ρ. A vector primitive cannot be reduced to ρ alone (a scalar cannot supply vector content without an independent directional structure), so the framework either admits no vector primitive at all or admits exactly one. The matter-coupling decomposition $J^\mu = \rho u^\mu$ used in §5 requires u^μ to exist as a substrate-level construct, so we list it as the conditional vector primitive. Its admissibility as an SST-derived primitive — rather than a residual postulate — depends on the matter-sector programme's substrate-level definition. The SST paper itself works with ρ and the commitment functional κ and does not enumerate u^μ; this gap is acknowledged in §15 and is the only residual condition of the CCA derivation.

No tensor primitives of rank ≥ 2 appear in the SST paper. By the Constitutive Principle of Exhaustiveness, any admissible field is functionally reducible to ρ; a tensor primitive (by definition independent of ρ) would therefore contradict exhaustiveness. Hence no tensor primitives exist. ■

Lemma 7.0.2 (Curl-freeness of admissible dimension-1 vectors). Under SST and the non-fractional-power admissibility condition (no branch cuts; only polynomial or analytic functional dependence on ρ), every admissible Lorentz vector V^μ of mass dimension 1 takes the form

$$V^\mu = \partial^\mu f(\rho/\rho_0)$$

for some dimensionless analytic function f, and therefore satisfies

$$\partial^{[\mu} V^{\nu]} = 0$$

identically, modulo curvature corrections that enter at higher mass dimension.

Argument. Under SST, V^μ must be a functional of ρ (and its derivatives, and u^μ if admissible). A dimension-1 vector built from ρ (dimension 3) and u^μ (dimension 0) by composition and differentiation has the following candidates:

(i) $\rho^a u^\mu$ for non-integer a: requires fractional powers, excluded by admissibility.

(ii) $\rho \times$ (dimensionless quantity) $\times$ (∂ acting on something): needs a vector index from somewhere; the only available sources are u^μ (dimensionless, picks up no dimension) or ∂^μ. With $\rho \times$ (dimensionless) $\times\, \partial^\mu$ (anything), the total dimension is $3 + 0 + 1 + [\geq 0] \geq 4$, exceeding 1. Excluded.

(iii) ∂^μ acting on a dimensionless scalar: gives total dimension 1. The dimensionless scalar must be analytic in ρ, hence a function of ρ/ρ_0 (since ρ has dimension 3, no other dimensional ratio is available without introducing additional scales, which would violate refinement persistence). Other dimensional scales such as the κ-field mass m_κ derive from ρ_0 via $m_\kappa^2 = (4/3)\rho_0^{(2/3)}$ per SST §2.2, so any putative additional dimensional ratio reduces to a function of ρ/ρ_0. Hence $V^\mu = \partial^\mu f(\rho/\rho_0)$.

Taking the antisymmetric derivative,

$$\partial^{[\mu} V^{\nu]} = \partial^{[\mu} \partial^{\nu]} f(\rho/\rho_0) = 0$$

identically, by commutativity of partial derivatives on the emergent Lorentzian continuum. On a curved manifold the commutator of covariant derivatives picks up a Riemann curvature term, but this enters at mass dimension ≥ 3 (since R has dimension 2) and is outside the leading order. ■

Proposition 7.0.3 (Derivation of CCA). Lemmas 7.0.1 and 7.0.2, together with dimensional analysis, non-fractional-power admissibility, and (A0) observer-invariant distinguishability, imply the three CCA clauses.

CCA-c (no admissible primitive scalar of dimension 2 in the substrate catalogue). By Lemma 7.0.1, the unique primitive scalar is ρ, with dimension 3. No primitive scalar of dimension 2 exists. ■

CCA-b (no admissible primitive antisymmetric rank-2 tensor of dimension 2 in the substrate catalogue). Candidates at this dimension:

(i) Primitive antisymmetric tensor of dimension 2: ruled out by Lemma 7.0.1 (no tensor primitives of rank ≥ 2 exist).

(ii) Composite of the form $\rho^a \times$ (antisymmetric dimensionless tensor): requires an antisymmetric dimensionless tensor built from primitives; the only candidates are $u^{[\mu} u^{\nu]} = 0$ (vanishes by antisymmetry) or contractions involving the metric (symmetric). Excluded.

(iii) Composite of the form $\partial^{[\mu} V^{\nu]}$ for an admissible V^μ of dimension 1: by Lemma 7.0.2, every such V^μ is curl-free, so $\partial^{[\mu} V^{\nu]} = 0$ identically. Excluded.

No surviving candidate. ■

CCA-a (J^μ is the unique Lorentz-covariant primitive substrate vector of dimension 3).
Candidates:

(i) ρ × dimensionless primitive vector = $\rho u^\mu = J^\mu$. Surviving candidate.

(ii) $\rho^a u^\mu$ for $a \neq 1$: requires fractional powers, excluded.

(iii) Constant Lorentz vector (no ρ dependence): picks out a preferred direction in spacetime, breaking Lorentz invariance at the emergent level, excluded by (A0).

(iv) Composite of the form $\partial^\mu \varphi$ for a scalar φ of dimension 2: φ could be a *composite* scalar of ρ (e.g., $\varphi = \Box \ln(\rho/\rho_0)$ is dimension 2 and SST-admissible as a composite). The candidate exists as a Lorentz vector of dimension 3, but it is not a primitive — it is a composite. As a coupling $A^\mu \partial_\mu \varphi$, it is excluded by (P3) gauge invariance: under $A_\mu \to A_\mu + \partial_\mu \chi$, the variation produces (after IBP) $\int \varphi \Box\chi \, d^4x$, which vanishes for arbitrary χ only if $\Box\varphi = 0$ off-shell, a non-trivial dynamical constraint. The composite is therefore not admissible as a leading-order coupling.

$J^\mu = \rho u^\mu$ is the unique surviving candidate. ■

Remark on the conditional clause. The derivation has one open condition: u^μ's status as an SST-admissible primitive. The SST paper establishes ρ as the unique primitive scalar but does not explicitly enumerate u^μ. §15 identifies this as part of the matter-sector gap (specifically, item 1 — "charged excitations" — requires a substrate-level characterisation of u^μ to be complete). Lemma 7.0.1's u^μ clause does not introduce a new gap; it makes visible that the matter-sector gap already acknowledged in §15 is doing structural work in the CCA derivation. Closing the matter-sector gap closes the CCA derivation completely.

Remark on primitive vs composite. Proposition 7.0.3's CCA-c says no *primitive* scalar of dimension 2 exists. *Composite* scalars of dimension 2 do exist — for example, $\Box \ln(\rho/\rho_0)$, or $|\nabla\rho|^2/\rho^{8/3}$ (the latter excluded by fractional powers, but the former perfectly admissible). What forbids these composites from generating new admissible couplings is not their absence from the catalogue but the gauge-invariance condition (P3): any candidate $A^\mu \partial_\mu \varphi$ with φ a composite scalar of dimension 2 fails gauge invariance for arbitrary χ. This is the structural role of (P3) in conjunction with CCA-c: CCA-c handles the catalogue side (no primitive at dimension 2), (P3) handles the composite side (no gauge-invariant scalar-gradient coupling). Together they exhaust the dimension-3 vector candidates from the scalar-gradient origin.

§7.1 — The Catalogue Closure Theorem

With Proposition 7.0.3 in hand, what was previously stated as an assumption is now a theorem. We retain the CCA name for reference and naming convenience.

Theorem (Catalogue Closure).[^cca-history] At leading admissible order in the current $K = 7$ substrate catalogue, conditional on the u^μ clause of Lemma 7.0.1:

- *(CCA-a)* The only admissible Lorentz-covariant primitive substrate vector of mass dimension 3 is the record current $J^\mu = \rho u^\mu$.
- *(CCA-b)* No admissible primitive antisymmetric rank-2 substrate tensor of dimension 2 exists.
- *(CCA-c)* No admissible primitive substrate scalar of dimension 2 exists.

[^cca-history]: Previously referred to as the Catalogue Closure Assumption in earlier drafts; §7.0 elevates it to theorem status. The CCA-a/b/c clause labels are retained for cross-reference continuity.

Proof. By Proposition 7.0.3. ■

Composite candidates (scalar-gradient couplings, kinetic-renormalising terms, parity-odd terms) are handled in the Theorem 1B proof below via (P3) gauge invariance and the Derivative Suppression Lemma. The uniqueness theorem of the remainder of §7 therefore splits naturally into two stages: Theorem 1A classifies the admissible dimension-3 Lorentz vectors available for contraction with A_μ at leading order, and Theorem 1B identifies the unique surviving interaction term from that classification, rules out non-vector candidates, and excludes higher-derivative variants via a refinement-persistence lemma.

Theorem 1A — Classification of admissible dimension-3 Lorentz vectors

Statement. Under CCA and (P1)–(P2), the set V_3 of admissible Lorentz vectors of mass dimension 3 available for contraction with A_μ at leading order, *enumerated by dimensional and tensorial source*, is

$$V_3 = \{ J^\mu, \partial_\nu F^{\nu\mu}, \partial_\nu \tilde{F}^{\nu\mu} \} ,$$

where $\tilde{F}^{\nu\mu} = \frac{1}{2} \varepsilon^{\nu\mu\rho\sigma} F_{\rho\sigma}$ is the dual field strength. The third element is identically zero by the Bianchi identity (as Theorem 1B's treatment of Element (iii) shows) and is listed here for completeness of the enumeration rather than as a non-trivial contributor; the non-trivial content of V_3 is the two-element set $\{ J^\mu, \partial_\nu F^{\nu\mu} \}$.

Proof. In natural units, A_μ has mass dimension 1, ∂_μ has dimension 1, and $F_{\mu\nu}$ has dimension 2. For a Lagrangian density of dimension 4 in $d = 4$ spacetime that is linear in A_μ, the contracted index of A_μ must be supplied by a Lorentz vector of dimension 3. Such a vector must be built from one of three sources: substrate fields directly, the field strength F and its derivatives, or gradients of admissible scalars.

Substrate origin. By CCA-a, the only admissible substrate vector of dimension 3 is J^μ.

Field-strength origin. At dimension 3, the only Lorentz vectors built from $F_{\mu\nu}$ and a single derivative are the divergence $\partial_\nu F^{\nu\mu}$ and the dual divergence $\partial_\nu \tilde{F}^{\nu\mu}$. Higher-derivative variants ($\partial^2 \partial_\nu F^{\nu\mu}$, etc.) carry additional derivatives and lie outside the leading order.

Scalar-gradient origin. A dimension-3 vector of the form $\partial^\mu \varphi$ would require φ to be a scalar of dimension 2. Two sub-cases arise. *(i)* If φ is a primitive substrate scalar, no such primitive exists at dimension 2 by §7.1 Theorem (c) (derived from Lemma 7.0.1). *(ii)* If φ is a composite scalar of dimension 2 (e.g., $\Box \ln(\rho/\rho_0)$), the candidate $A^\mu \partial_\mu \varphi$ exists at the Lagrangian level but is excluded by (P3) gauge invariance: under $A_\mu \to A_\mu + \partial_\mu \chi$, the variation produces (after IBP) $\int \varphi \, \Box\chi \, d^4x$, which vanishes for arbitrary χ only if $\Box\varphi = 0$ off-shell — a non-trivial dynamical constraint on φ, not an algebraic identity it can be expected to satisfy. Both sub-cases are excluded, so no admissible scalar-gradient candidate contributes to V_3.

Combining the three sources, $V_3 = \{ J^\mu, \partial_\nu F^{\nu\mu}, \partial_\nu \tilde{F}^{\nu\mu} \}$. ■

Theorem 1B — Unique leading-order interaction

Statement. Under CCA and (P1)–(P5), the unique admissible leading-order interaction Lagrangian density built from A_μ and the substrate and gauge field content available at this order is

$$\mathscr{L}_{\text{int}} = -J^\mu A_\mu \, ,$$

modulo total derivatives, with the overall sign absorbed into the definition of J^μ.

Proof. We treat each element of V_3 from Theorem 1A in turn as a candidate contraction with A_μ, then rule out non-vector candidates and higher-derivative variants.

Element (i): J^μ. The candidate is $c_1 J^\mu A_\mu$. Under $A_\mu \to A_\mu + \partial_\mu \chi$,

$$\mathscr{L}^{1)}{}_{\text{int}} \to c_1 J^\mu A_\mu + c_1 J^\mu \partial_\mu \chi \, ,$$

and the second term is, after IBP, $-c_1 (\partial_\mu J^\mu) \chi$ plus a boundary term. With BCB ($\partial_\mu J^\mu = 0$) the variation is a total derivative and the action is gauge-invariant; (P3) holds. This is the surviving candidate.

Element (ii): $\partial_\nu F^{\nu\mu}$. The candidate is $A_\mu \partial_\nu F^{\nu\mu}$. Integration by parts gives, up to boundary terms, $-F^{\nu\mu} \partial_\nu A_\mu$; using the antisymmetry of $F^{\nu\mu}$ this equals $-\frac{1}{2} F^{\nu\mu} F_{\nu\mu}$ — that is, twice the standard kinetic coefficient (since the conventional kinetic Lagrangian is $-\frac{1}{4} F^2$). Adding this term to the action therefore merely renormalises the kinetic coefficient and produces no independent structure. It is not a new leading-order interaction.

Element (iii): $\partial_\nu \tilde{F}^{\nu\mu}$. This vanishes identically before any IBP: $\partial_\nu \tilde{F}^{\nu\mu} = \frac{1}{2} \varepsilon^{\nu\mu\rho\sigma} \partial_\nu F_{\rho\sigma} = 0$ by the Bianchi identity, since the cyclic sum $\varepsilon^{\alpha\beta\gamma\delta} \partial_\alpha F_{\beta\gamma}$ vanishes. Equivalently, IBP converts the candidate to a multiple of $F_{\mu\nu} \tilde{F}^{\mu\nu}$, which is a classical total derivative (treated as a non-interaction below). Both routes agree: the candidate is classically inert. Quantum-mechanically the $F\tilde{F}$ total-derivative ambiguity has non-trivial topological content — θ-vacuum structure, the Witten effect, axion physics — but lies outside the leading-order classical enumeration; its admissibility depends on a substrate-level parity-admissibility constraint not addressed here.

Only Element (i) survives from V_3.

Non-vector candidate: mass term. The term $A_\mu A^\mu$ has dimension 2; a dimensionful prefactor of mass-squared dimension would bring it to dimension 4, so dimensional analysis alone does not exclude it. The decisive objection is gauge-theoretic: under $A_\mu \to A_\mu + \partial_\mu \chi$,

$$A_\mu A^\mu \to A_\mu A^\mu + 2 A^\mu \partial_\mu \chi + (\partial_\mu \chi)(\partial^\mu \chi) ,$$

and the variation is not a total derivative for arbitrary χ. The mass term is excluded by (P3) independent of any dimensional concern. This is the structural reason the persistent gauge sector cannot acquire an intrinsic mass — the topic of §11.

Non-vector candidates: pure gauge. The terms $F_{\mu\nu} F^{\mu\nu}$ and $F_{\mu\nu} \tilde{F}^{\mu\nu}$ are dimension 4, gauge-invariant, and involve no substrate field. They are therefore not *interactions* in the sense (P5) is selecting. The kinetic term F^2 is included separately in the full action with the $-\frac{1}{4}$ normalisation inherited from the Maxwell admissibility paper. The topological term $F\tilde{F}$ is a classical total derivative,

$$F_{\mu\nu} \tilde{F}^{\mu\nu} = \partial_\mu (\varepsilon^{\mu\nu\alpha\beta} A_\nu \partial_\alpha A_\beta) ,$$

contributing nothing to the classical field equations; its quantum content was discussed above under Element (iii).

Non-vector candidate: tensor-paired. A candidate of the form $K^{\mu\nu} F_{\mu\nu}$ would require an antisymmetric two-index substrate tensor $K^{\mu\nu}$. CCA-b rules this out at the relevant order.

Higher-derivative variants. Terms with additional derivatives — $J^\mu \Box A_\mu$, $(\partial_\nu J^\mu)(\partial^\nu A_\mu)$, $F^{\mu\nu} \partial_\mu J_\nu$, and similar — exist at the lattice level prior to refinement. Conventional EFT power-counting does *not* exclude them: the VERSF substrate possesses a natural scale (the lattice spacing a_n), so a candidate with $2m$ additional derivatives carries a prefactor of order $a_n^{(2m)}$, small but non-zero at finite refinement. Their exclusion at leading order is therefore not dimensional but structural, supplied by the following lemma.

Lemma (Derivative suppression under refinement). Let $\{a_n\}$ be a sequence of refinement scales with $a_n \to 0$, and let k denote a characteristic continuum momentum mode. Under the regularity assumption that the refinement coarse-graining form factor is smooth at zero momentum (see below), any local operator differing from a leading admissible operator by $2m$ additional derivatives scales under the refinement-stable persistence projection as

$$(k \cdot a_n)^{(2m)} ,$$

and therefore vanishes in the refinement-stable continuum limit $a_n \to 0$.

Sketch of argument. The persistence projection acts by coarse-graining the field content over cells of linear size a_n. A local operator with $2m$ additional derivatives picks up $2m$ extra factors of momentum k under Fourier transform. If the coarse-graining form factor $\hat{f}(k\, a_n)$ is smooth at zero momentum — that is, $\hat{f}(0) = 1$ and the leading correction is O((k·a_n)^2) (an even function of k · a_n, as is the case for any symmetric smooth filter such as a Gaussian or compactly-supported smooth bump) — then each pair of derivatives picks up suppression (k·a_n)^2 and the cumulative scaling is (k·a_n)^(2m). At fixed continuum momentum k, this vanishes as a_n → 0. The leading admissible operators are exactly those invariant under the persistence projection to leading order, with the higher-derivative variants suppressed by the scaling above. A complete proof requires the explicit refinement functor constructed in the persistence papers. ■

Regularity assumption: why it matters and what it requires. The (k·a_n)^(2m) scaling depends on the form factor being smooth at zero momentum, which gives O((k·a_n)^2) per derivative pair. A sharp-cutoff coarse-graining (Heaviside-type form factor) would give only O(k·a_n) per derivative pair, hence (k·a_n)^m rather than (k·a_n)^(2m), and the predictions of §14 would shift accordingly. The refinement functor used by the persistence papers is, to our knowledge, a smooth construction with the required regularity, so the (k·a_n)^(2m) scaling applies; this should be confirmed explicitly when the refinement functor is cited at point of use, and a future version of the persistence papers may wish to record this regularity property as a named lemma. If a sharp-cutoff coarse-graining is ever adopted, both the Lemma above and the §14 predictions need to be restated with the linear scaling.

By the Lemma, higher-derivative variants of the interaction term are eliminated by persistence in the continuum limit. They reappear only as the leading admissible *non-persistent* corrections at finite refinement scales, with magnitudes set by powers of (k · a_n). These corrections form the basis of the falsifiable predictions discussed in §14.

Normalisation. The constant c$_1$ in Element (i) is absorbed into the definition of J^μ. We take c$_1$ = −1, giving

ℒ_int = −J^μ A_μ ,

which yields ∂_μ F^μν = +J^ν under variation.[^sign]

[^sign]: Sign convention. Variation of S = ∫(−¼ F^2 − J·A) d^{4}x with respect to A_ν yields ∂_μ F^μν = +J^ν. The opposite sign in either term flips the sign of the source. This is standard particle-physics convention and is a choice, not a derivation.

This exhausts the candidate space under CCA and (P1)–(P5). ■

Remark on the scope of uniqueness. The claim is *conditional uniqueness*: uniqueness within the leading-order admissible field content currently established in the VERSF substrate catalogue, as encoded in CCA. It is not unconditional. Three observations sharpen this.

First, the Catalogue Closure Theorem itself rests on SST Definition 1, the Constitutive Principle of Exhaustiveness, Lemmas 7.0.1 and 7.0.2, and the u^μ matter-sector condition (see §7.0 and §16); any future revision of those inputs propagates to the uniqueness claim.

Second, the two most important hypothetical extensions of CCA — admission of an antisymmetric two-index substrate tensor $K^{\mu\nu}$ (relaxing CCA-b), or an admissible scalar φ of the relevant dimension (relaxing CCA-c) — are structurally non-trivial and would require independent admissibility justification. Neither has yet been established as admissible. Until one is, Theorem 1 stands.

Third, extension to non-abelian gauge structure requires precisely the kind of additional field content that the present catalogue does not contain. The narrowness of the uniqueness claim and the abelian character of the resulting gauge theory are not independent: both reflect the present scope of the substrate catalogue, and both will need to be revisited together if the catalogue is extended.

The right way to read the theorem is therefore not "the coupling is uniquely fixed forever" but: *given the structural commitments encoded in CCA and the admissibility principles (P1)–(P5), no other coupling at this order is admissible*. That is a strong claim, but a clearly delimited one.

Remark on unitarity under coupling. Theorem 1 establishes the unique admissible form of the interaction Lagrangian; it does not, by itself, establish that the *coupled dynamics* preserves distinguishability. (P4) is checked at the level of the coupling term with J^μ treated as a prescribed external source — confirming gauge invariance and admissibility of the interaction — not at the level of the full coupled evolution, which involves backreaction of J^μ on A_μ (Theorem 3) and of A_μ on whatever microscopic structure carries J^μ (§15).

The distinction is worth stating explicitly to avoid overclaiming. *Pre-coupling*, unitary Hamiltonian evolution on $H^1(\mathcal{G}(\Lambda))$ is established by the admissibility-generator theorem of the previous paper. *After coupling*, the present construction establishes only *linearised* unitarity — i.e., unitarity for J^μ treated as an external prescribed source rather than as the dynamical response of a matter sector. *Full coupled unitarity*, in which both A_μ and the carriers of J^μ are dynamical and respond to one another, depends on matter-sector structure and is deferred to the matter-sector programme. The present paper neither establishes nor relies on full coupled unitarity; the leading-order classical construction is consistent without it. ■

---

8. Theorem 2 — Gauge Redundancy ⇔ Current Conservation

Statement. The interaction action $S_{\text{int}} = \int (-J^\mu A_\mu)\, d^4x$ is invariant under $A_\mu \to A_\mu + \partial_\mu \chi$ for arbitrary smooth χ with admissible falloff *if and only if* $\partial_\mu J^\mu = 0$.

Proof. Under the gauge transformation,

$$\delta S_{\text{int}} = -\int J^\mu\, \partial_\mu \chi\, d^4x\ .$$

Integrating by parts and discarding the boundary term (admissible by assumption on the falloff of χ),

$$\delta S_{\text{int}} = \int (\partial_\mu J^\mu)\, \chi \, d^4x \, .$$

For arbitrary χ, $\delta S_{\text{int}} = 0$ if and only if $\partial_\mu J^\mu = 0$ as a distribution, and hence (by density of test functions) as a smooth field. ■

Structural reading. Two principles that arrived at the persistent sector through entirely different routes — gauge redundancy from the cohomology quotient on $H^1(\mathcal{G}(\Lambda))$, and current conservation from the BCB axiom on the substrate — turn out to be the same constraint. Either one implies the other in the coupled theory. This is the first explicit place in VERSF where a persistent-sector structural principle and a substrate-level conservation law become *identified* rather than merely consistent.

The statement is sharp. Off-shell, invariance of the action under arbitrary gauge transformations requires $\partial_\mu J^\mu = 0$ *exactly* — not approximately, not up to higher-order corrections, not modulo any tolerance. Conversely, exact current conservation guarantees gauge invariance of the action with no residual obstruction. Any weakening of one (approximate conservation, or invariance only for restricted χ) breaks the other in exactly the same measure.

It is worth emphasising what this does and does not say. It does not say BCB is *derived from* gauge invariance, nor that gauge invariance is derived from BCB; both arise independently from earlier structural arguments. It says that in the coupled theory they are dual expressions of a single admissibility constraint, locked together off-shell with no slack.

---

9. Theorem 3 — Source-Coupled Maxwell Dynamics

Statement. With the kinetic Lagrangian $-\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$ inherited from the Maxwell admissibility paper (which establishes its sign, signature, and normalisation by independent admissibility arguments on the persistent sector), the admissible persistent-sector action

$$S[A] = \int \left(-\tfrac{1}{4} F_{\mu\nu} F^{\mu\nu} - J^\mu A_\mu \right) d^4x$$

yields, under variation with respect to A_μ,

$$\partial_\mu F^{\mu\nu} = J^\nu \, .$$

Proof. Write $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ and compute the variation:

$$\delta S = \int \left(-\tfrac{1}{2} F^{\mu\nu} \delta F_{\mu\nu} - J^\mu \delta A_\mu \right) d^4x \, .$$

Using $\delta F_{\mu\nu} = \partial_\mu \delta A_\nu - \partial_\nu \delta A_\mu$ and the antisymmetry of $F^{\mu\nu}$,

$$-\tfrac{1}{2} F^{\mu\nu} \delta F_{\mu\nu} = -F^{\mu\nu} \partial_\mu \delta A_\nu .$$

Integrating by parts and discarding the boundary term,

$$\int -F^{\mu\nu} \partial_\mu \delta A_\nu \, d^4x = \int (\partial_\mu F^{\mu\nu}) \delta A_\nu \, d^4x .$$

Combining,

$$\delta S = \int (\partial_\mu F^{\mu\nu} - J^\nu) \delta A_\nu \, d^4x .$$

Setting $\delta S = 0$ for arbitrary δA_ν yields the field equation

$$\partial_\mu F^{\mu\nu} = J^\nu . \blacksquare$$

The persistent cohomological transport sector therefore reproduces the inhomogeneous Maxwell equation once coupled to the record current. The coupling is not introduced phenomenologically; it is forced by Theorem 1, and the source side of the equation is exactly the substrate conservation current with no rescaling required beyond the normalisation absorbed in Theorem 1B.

---

10. The Bianchi Identity from F = dA

A clean derivation should not stop at the inhomogeneous Maxwell equation. The homogeneous equation — equivalently, the Bianchi identity for F — follows automatically and deserves explicit statement.

Because A_μ is a connection 1-form on the persistent cohomological sector and $F = dA$, the exterior derivative of F vanishes identically:

$$dF = d(dA) = 0 .$$

In component form,

$$\partial_\alpha F_{\beta\gamma} + \partial_\beta F_{\gamma\alpha} + \partial_\gamma F_{\alpha\beta} = 0 ,$$

or equivalently $\partial_{[\alpha} F_{\beta\gamma]} = 0$. In familiar electromagnetic language this is the pair

$$\nabla \cdot B = 0 , \partial_t B + \nabla \times E = 0 ,$$

i.e. the absence of magnetic monopoles and Faraday's law.

No additional input is required. The Bianchi identity is therefore *not a dynamical equation but a cohomological identity*: it arises from the exactness of the curvature 2-form F as the exterior derivative of the connection 1-form A on the refinement-persistent sector. Equivalently, F is closed because $d^2 = 0$ on the de Rham complex carried by $H^1(\mathcal{G}(\Lambda))$. Within VERSF this is significant: it means *half* of the Maxwell equations emerge directly from the cohomological structure of the persistent sector, before any coupling to J^μ is introduced. Coupling to the record current is what generates the *other half* — the inhomogeneous source equation.

The split is structurally clean:

Equation	**Origin**
$\partial_{[\alpha} F_{\beta\gamma]} = 0$	$F = dA$ on $H^1(\mathcal{G}(\Lambda))$ — persistent sector alone
$\partial_\mu F^{\mu\nu} = J^\nu$	Coupling of persistent sector to record current

The structural content of this split deserves explicit statement: *the Bianchi identity is the only part of Maxwell theory that does not require coupling to matter*. It is a consequence of the persistent sector's cohomological structure alone, prior to any source. The inhomogeneous equation, by contrast, exists only in the presence of a coupled source — and that source is precisely what Theorem 3 supplies. The persistent sector therefore comes equipped with half of Maxwell's equations as a structural inevitability; the other half is the entire content of the coupling theorem.

---

11. Proposition 4 — The Inertia Question, Resolved (Corollary of Theorem 1)

Statement. The persistent gauge sector requires no intrinsic mass-like term within the $K = 7$ constraint catalogue. The effective dynamical response of the coupled gauge–matter system arises through interaction with the record-current sector rather than through any intrinsic harmonic-sector term.

This is presented as a proposition rather than as an independent theorem because its content is not a new mathematical result: it is the physical re-reading of the mass-term exclusion in Theorem 1B, where the candidate mass term $A_\mu A^\mu$ was excluded by gauge redundancy (P3). The point of §11 is therefore interpretive — to connect that exclusion to the inertia question posed by the previous paper — rather than to establish new structure.

Argument. The previous paper established that the $K = 7$ harmonic sector admits no intrinsic mass-like term consistent with admissibility — the mass-term exclusion in Theorem 1B makes this precise. This raised the question of whether the persistent gauge sector could generate any observable dynamics at all, since in conventional formulations of gauge theory it is the matter sector that supplies the dynamical content the gauge field responds to, and the gauge field itself remains massless.

The resolution is structurally clean. The persistent sector is not *meant* to carry intrinsic dynamical content. Its role in the layered architecture is to supply reversible transport structure — gauge redundancy, Wilson-loop observables, Hamiltonian evolution on $H^1(\mathcal{G}(\Lambda))$ — while the record-current sector supplies the source content to which the gauge field responds. The observable dynamics belong to the coupled system $A_\mu \oplus J^\mu$ and arise through the interaction term $-J^\mu A_\mu$ identified in Theorem 1.

This is precisely the situation in standard quantum electrodynamics: the photon field is massless, the electron field is massive, and observable electromagnetic phenomenology arises from their coupling. The VERSF construction recovers this layered structure from first principles. The persistent harmonic sector is the photon-like layer; the record-current sector is the matter-like layer; their interaction is the electromagnetic coupling. No $K = 7$ modification is required, and the absence of an admissible $A_\mu A^\mu$ term is no longer a problem — it is the structural reason the persistent sector remains massless, exactly as a gauge boson should. ■

The "inertia problem" of the previous paper is therefore not a problem in the VERSF framework but a *signature*: it is the structural fingerprint of an unbroken gauge sector.

12. Physical Interpretation

The dynamical architecture that has now emerged across the recent sequence of papers can be summarised compactly:

Sector	Dynamics	Physical role
σ-sector	dissipative gradient flow	admissibility restoration
persistent harmonic sector (pre-coupling)	unitary Hamiltonian transport	reversible gauge transport
record-current sector	irreversible commitment flow	source structure
coupled gauge theory	Maxwell-form source dynamics	observable electromagnetism

The "pre-coupling" qualifier in the second row is essential. Unitary Hamiltonian evolution on $H^1(\mathcal{G}(\Lambda))$ is established by the admissibility-generator theorem for the persistent sector taken in isolation. Once the coupling $-J^\mu A_\mu$ is turned on, full unitarity of the coupled dynamics requires the matter-sector structure that supplies J^μ to itself be unitary under the gauge backreaction. The present paper establishes only linearised unitarity (J^μ treated as a prescribed external source); the full coupled unitarity is deferred to the matter-sector programme. The fourth row therefore inherits unitarity from neither the persistent nor the matter sector alone but from their joint structure, and that joint structure has not yet been constructed in full.

Several features of conventional physics now receive a unified structural explanation.

The coexistence of *diffusion-like* and *wave-like* behaviour in physical systems is no longer a phenomenological observation but reflects the distinction between the dissipative σ-sector and the reversible persistent sector — two separate admissibility classes acting on the same substrate.

The *gauge redundancy* of the electromagnetic potential is no longer a postulate of field theory but an inherited feature of the cohomology quotient on $H^1(\mathcal{G}(\Lambda))$.

The *conservation of electric current* is no longer an independent input but the substrate-level BCB axiom expressed in continuum form.

And the *coupling structure* J^μ A_μ is no longer a minimal-coupling prescription chosen for convenience but the unique lowest-order admissible interaction given the field content.

Each of these conventionally-postulated features becomes, in the VERSF framework, a derived consequence of earlier structural commitments. This is the principal interest of the paper: not the recovery of Maxwell equations *per se* — which is a known structural inevitability of any U(1) gauge theory with a conserved source — but the demonstration that the gauge structure, the conservation law, and the coupling all emerge from independent substrate-level constructions and turn out to be mutually consistent at the level of admissibility.

---

13. Relation to Earlier VERSF Papers, and the Dependency Graph

This paper sits at the intersection of four earlier strands:

- **The refinement persistence and cohomology papers**, which established $H^1(\mathcal{G}(\Lambda))$ as the unique refinement-stable observable sector and identified A_μ as the cohomological transport potential.
- **The Maxwell admissibility paper (through v19)**, which derived U(1) gauge transport from BCB/TPB constraints and established Wilson-loop observables in the persistent sector.
- **The Hamiltonian admissibility paper**, which proved that any reversible, composable, distinguishability-preserving evolution generates a self-adjoint Hamiltonian and unitary dynamics.
- **The inertia problem paper**, which identified the absence of admissible mass-like terms in the $K = 7$ harmonic sector and posed the question this paper answers.

The present paper is the first to combine all four into a single coupled transport framework. It also incorporates the Single-Source Theorem as a structural input: §7.0 derives the Catalogue Closure Theorem from SST plus admissibility principles, so the uniqueness claim of Theorem 1 now rests directly on SST rather than on a separately-named catalogue assumption.

A separate consistency check, deferred to a companion paper, is the relation between this paper's interaction term and the fold/record-current structure developed in *The Fold and the Record* (through v31), particularly Theorem 12.2 on record-current uniqueness. The check has a clean falsifiable form: *the substrate-level record current J^μ of this paper must coincide, up to overall normalisation, with the unique current established in Theorem 12.2 of the Fold paper*. If that identification holds, the present construction inherits the uniqueness and refinement-persistence properties of Theorem 12.2 automatically, and the structural circle closes. If it does not, J^μ in this paper is a closely related but distinct object and the two constructions must be reconciled before either is built upon. Either outcome is informative.

Dependency graph. The construction of this paper rests on a non-trivial graph of earlier results. Making it explicit clarifies how the pieces interlock:

Result	Depends on
Persistent U(1) sector on $H^1(\mathcal{G}(\Lambda))$	Refinement persistence + Wilson identification
Gauge redundancy A_μ → A_μ + ∂_μ χ	Cohomology quotient on $H^1(\mathcal{G}(\Lambda))$
Unitary Hamiltonian transport (pre-coupling)	Admissibility-generator theorem
Conservation ∂_μ J^μ = 0	BCB axiom (with the J^μ = ρ u^μ sketch of §5)
Primitive Substrate Field Exhaustion (Lemma 7.0.1)	SST Definition 1 + Constitutive Principle of Exhaustiveness; conditional on u^μ substrate-level definition (matter-sector gap, §15)
Curl-freeness of admissible dim-1 vectors (Lemma 7.0.2)	SST + non-fractional-power admissibility + commutativity of partial derivatives
Catalogue Closure Theorem	Lemmas 7.0.1, 7.0.2 + dimensional analysis + (A0) observer-invariant distinguishability
Uniqueness of J^μ A_μ coupling	Catalogue Closure Theorem + (P1)–(P5), Theorem 1
Maxwell source equation ∂_μ F^μν = J^ν	Theorem 1 + variational principle (§9) + kinetic-term sign from Maxwell admissibility paper
Bianchi ∂_[α F_βγ] = 0	Cohomology of A on persistent sector ($d^2 = 0$)
Masslessness of A_μ	Theorem 1B mass-term exclusion (Proposition 4)

No row in this table is supplied by this paper alone; every result is either inherited from an earlier paper or derived by combining inherited results under the admissibility principles introduced here.

---

14. Admissible Corrections and Falsifiable Predictions

The construction so far has been purely structural. Before turning to the matter question, it is worth stating what the architecture predicts about *deviations* from Maxwell structure, since these are the closest the present paper comes to falsifiable content.

Within the persistent-sector interpretation, Maxwell structure is exact in the refinement-stable continuum limit. Deviations can therefore arise from precisely three sources, each structurally distinct and each in principle observable:

1. Finite-refinement corrections. At any finite refinement scale $a_n > 0$, the persistence projection has not yet eliminated higher-derivative operators. Under the regularity assumption of the §7 Lemma (smooth coarse-graining form factor), the leading correction to the $J^\mu A_\mu$ coupling enters as an operator of schematic form

$$a_n^2 \times (\partial^2 \times J^\mu A_\mu) ,$$

i.e., suppressed by two derivatives carrying powers of $(k \cdot a_n)^2$. For continuum momenta k well below the inverse lattice scale, this is small but non-zero. Any experiment sensitive to dispersion relations, photon propagation at extreme energies, or precision quantum-electrodynamic predictions provides, in principle, an upper bound on a_n and therefore a falsifiable handle on the refinement scale. *(If the refinement functor turns out to use a sharp-cutoff coarse-graining rather than a smooth one, the suppression weakens to $(k \cdot a_n)$ per derivative pair and the scaling becomes linear; the §7 Lemma's regularity discussion makes this dependency explicit.)*

2. Catalogue extensions. If the $K = 7$ substrate catalogue is extended — for example, if an admissible antisymmetric two-index substrate tensor $K^{\mu\nu}$ or an admissible scalar φ of dimension 2 is established by independent admissibility argument — then new coupling terms ($K^{\mu\nu} F_{\mu\nu}$, $A^\mu \partial_\mu \varphi$ in suitable gauges) become admissible. These would manifest as deviations from pure Maxwell structure of forms not reducible to higher-derivative corrections. Their absence in observation constrains the catalogue; their presence would force catalogue extension.

3. Breakdown of refinement persistence. If the persistence projection itself fails in some regime, the elimination of higher-derivative corrections fails non-perturbatively. The natural quantitative criterion would compare the σ-sector dissipation rate γ_σ against the characteristic persistent-sector frequency ω_{pers}; persistence is reliable in the regime $\gamma_\sigma \ll \omega_{pers}$ and fails as $\gamma_\sigma \to \omega_{pers}$. In a fully developed VERSF substrate calculation each of these would be a computable function of refinement scale and field content, and the failure criterion would be a sharp inequality with measurable consequences — most plausibly a coherence breakdown of the $F = dA$ closure on $H^1(\mathcal{G}(\Lambda))$, manifesting as momentum-dependent loss of Wilson-loop invariance.

This is the least sharp of the three channels: γ_σ and ω_{pers} are both schematic stand-ins, and a quantitative inequality requires substrate-level calculations not carried out here. Honestly stated, channel 3 is *qualitatively predicted* (the construction implies persistence breakdown is a logical possibility distinct from channels 1 and 2) but *not yet quantitatively constrained*. Channels 1 and 2 carry the falsifiable weight; channel 3 is included for structural completeness and will need a follow-up calculation to be made operationally testable.

The structural statement is therefore:

Within the persistent-sector interpretation, every admissible deviation from Maxwell structure enters as either a persistence-suppressed higher-derivative operator scaling as a power of (k · a_n), a catalogue-extension term of restricted form, or a coherence breakdown of the refinement-stable limit. No other admissible deviation is structurally available.

This is a constraint, not a prediction of a specific deviation. But it is a *falsifiable* constraint: any observation of a Maxwell deviation not reducible to one of these three structural categories would falsify the persistent-sector interpretation. Conversely, observed deviations consistent with the (k · a_n)^2 scaling at high energy would constitute positive evidence for the refinement-scale structure.

The predictive content is therefore narrower than that of a complete electromagnetic theory — which would predict specific numerical values for α, the deviation amplitudes, and so on — but it is non-trivial: it forecloses the open-ended landscape of "any deviation from Maxwell is admissible" and replaces it with a structurally constrained three-channel correction space.

---

15. The Matter Question: What Physically Carries J^μ?

The natural objection to this paper is sharp and should be stated directly:

You postulated a conserved current J^μ and recovered the Maxwell equations. Any U(1) gauge theory with a conserved source yields the same structure. What has actually been achieved?

The objection is fair. It identifies the dominant remaining gap in the construction, and any honest assessment of the paper must address it.

What has been achieved despite the objection. Two results are non-trivial and worth holding onto as preamble to the structural four-part claim below.

First, the coupling structure is forced, not chosen. In conventional field theory, the minimal-coupling prescription J^μ A_μ is selected by appeal to elegance, gauge invariance, and renormalisability. Here it is the *unique* admissible coupling at leading order, by Theorem 1. No tuning, no minimal-coupling postulate, no choice.

Second, gauge redundancy and current conservation are *identified*, not merely consistent. The cohomological structure of the persistent sector (which makes A_μ gauge-redundant) and the substrate-level BCB axiom (which makes J^μ conserved) arrive at the coupled theory through independent constructions, and Theorem 2 shows they are the same constraint. This is a unification, not a coincidence, and it is not visible from a Maxwell-theory starting point.

These two results, together with the structural masslessness statement of Proposition 4, establish how *rigid* the construction is once the framework is in place. None could be relaxed without breaking admissibility somewhere. That rigidity is the real claim of the paper.

The distinction worth stating aggressively is this. The novelty is **not** the recovery of Maxwell equations from gauge symmetry plus a conserved current; that recovery is a well-known structural inevitability of any U(1) gauge theory. The novelty **is** the four-part claim that, working from substrate primitives:

(a) the admissible gauge structure is *derived* (from cohomology on $H^1(\mathcal{G}(\Lambda))$), not postulated; (b) the conserved current is *derived independently* (from BCB on the substrate), not inserted; (c) the coupling between them is *uniquely forced* (by Theorem 1, with the Catalogue Closure Theorem derived from SST in §7.0), not chosen by minimal-coupling prescription; (d) the masslessness of the gauge field is *structurally derived* (Proposition 4, via the mass-term exclusion in Theorem 1B), not postulated.

A critic who responds "of course any conserved current plus U(1) gives Maxwell" is conceding (a)–(d) and pointing out that the *final assembly step* (variational derivation of ∂_μ F^μν = J^ν) is mechanical. That is correct: the final step is mechanical because (a)–(d) have done all the work. The interest is in those four substrate-level derivations, not in the variational step they enable.

What has not been achieved. The objection nevertheless lands on the genuine gap. The current J^μ is treated here as the conserved substrate commitment-flow vector — undecomposed, structureless, without microscopic carriers. This is sufficient for classical Maxwell structure, but it is insufficient for:

- actual matter, in any operational sense;
- charged excitations as identifiable substrate structures;
- fermionic species and their associated currents;
- spinorial structure, double-cover representations, and the connection to spin established in earlier VERSF papers;
- charge quantisation, which requires not just U(1) but ℤ-valued cohomology and discrete charge carriers;
- the species decomposition implicit in the CKM/PMNS flavour-mixing derivation;
- any quantum field theory of charged matter, and in particular QED.

Each of these is a separate structural problem, and none is addressed by the present construction. The paper establishes the *gauge-theoretic half* of an electromagnetic programme; the *matter-sector half* is still missing.

Why this is now the dominant remaining gap. Earlier papers in the sequence had distinct unresolved structures — the admissibility class of persistent transport, the existence of a Hamiltonian generator, the inertia question, and so on. With those resolved, the construction has converged to a single bottleneck: the microscopic origin of J^μ. Without an answer, the present paper sits as a structural waypoint — necessary, internally consistent, but physically incomplete.

The matter question is therefore not a deferred technicality. It is the next phase of the programme, and the present paper's value is partly defined by how completely the matter-sector phase eventually succeeds.

What the matter-sector programme must deliver. A complete VERSF account of charged matter will need at minimum:

1. *Charged excitations.* Identification of substrate configurations that carry conserved discrete charge under the persistent-sector gauge structure, with the integer cohomology $H^1(\mathcal{G}(\Lambda); \mathbb{Z})$ supplying the quantisation. The discreteness of the substrate makes this structurally available but it has not yet been constructed.
2. *Species decomposition.* A canonical decomposition of J^μ into species-resolved matter currents J^μ = Σ_a q_a J^μ_a, with the discreteness of the species set and the values q_a forced by admissibility rather than fit. The existing CKM/PMNS derivation suggests the closure-geometry machinery for this is in place; it has not yet been applied to the present construction.
3. *Spinorial structure.* Connection to the spin-as-double-cover representation established in earlier VERSF work, producing matter sectors that are not merely scalars carrying J^μ but spinorial fields with the correct fermionic statistics. This is the bridge to Dirac-type matter equations.
4. *Quantum field structure.* Promotion of the matter sector to operator-valued distributions on a Fock space, with the persistent gauge sector quantised in parallel, producing QED in the appropriate limit.

These are listed as a programme, not as commitments of the present paper. Their absence is the honest reading of what this paper does and does not do.

Structural guidance from the Single-Source Theorem. The matter-sector question is not entirely unconstrained. The Single-Source Theorem requires every observable to be a functional of the committed record density ρ(x, t). Under this constraint, J^μ is *forced* to take the form J^μ = J^μ[ρ] — for instance, as a commitment-flow vector ρ v^μ for some emergent velocity v^μ derived from the substrate dynamics, or as a gradient functional ∂^μ ρ in an alternative architecture, or as a more complex SST-admissible composite. Each candidate decomposition tightly constrains the microscopic content of the matter sector and converts the present construction's vague "what carries J^μ" into the sharp question "*which* SST-admissible functional dependence J^μ[ρ] is realised by the substrate." That is a much more tractable question than the unrestricted version, and identifying the answer is the natural first deliverable of the matter-sector programme. It also makes the eventual answer falsifiable: not every functional dependence is SST-admissible, so the constraint is genuinely informative.

The right framing. This paper establishes that *if* the substrate carries a conserved current, *then* the leading-order admissible coupling forces Maxwell structure with no remaining freedom. The next paper in this strand must establish that the substrate does in fact carry such a current with the right microscopic structure to be physically interpreted as electromagnetic — and the SST constraint above narrows what that current can look like. Both halves are needed. Without the second half, the present construction is mathematically tight but physically incomplete. With it, the construction becomes a derivation of electromagnetism from substrate primitives. The present paper does the first half cleanly; the second half is now the next priority of the programme.

16. Epistemic Status

For clarity, the results of this paper are labelled by status.

Proven within the framework.

- *Lemma 7.0.2 (Curl-freeness)* is a mathematical fact under SST and non-fractional-power admissibility, established directly from commutativity of partial derivatives in the emergent Lorentzian continuum.
- *Proposition 7.0.3 (Catalogue Closure Theorem)* follows from Lemma 7.0.1 and Lemma 7.0.2 plus dimensional analysis and (A0). Given the lemmas, the CCA clauses are immediate.
- *Theorem 1A (Classification)* is a mathematical consequence of the Catalogue Closure Theorem and dimensional analysis at this order: V_3 = {J^μ, ∂_ν F^νμ, ∂_ν F̃^νμ} exhausts the admissible dimension-3 Lorentz vectors.
- *Theorem 1B (Unique interaction)* is a mathematical consequence of the Catalogue Closure Theorem and (P1)–(P5): among candidates built from V_3 and the gauge field, plus non-vector candidates (mass, kinetic, topological, tensor-paired) and higher-derivative variants, only J^μ A_μ survives. The enumeration is exhaustive at this order.
- The *Derivative Suppression Lemma* establishes (k · a_n)^(2m) scaling of higher-derivative variants under the persistence projection, conditional on a smooth coarse-graining form factor (see the regularity remark in §7). The sketch is structural rather than exhaustively rigorous, and a complete proof requires the explicit refinement functor of the persistence papers. If the refinement functor is sharp-cutoff rather than smooth, the scaling becomes linear in (k · a_n) per derivative pair and both the Lemma and §14's channel-1 predictions need restating.
- *Theorem 2 (Gauge ⇔ conservation)* is a standard Noether-style result, made structurally meaningful by the independent origin of the two principles, and sharp off-shell.
- *Theorem 3 (Maxwell-form dynamics)* is a standard variational derivation; what is non-standard is that every ingredient was forced by earlier results rather than postulated.
- The Bianchi identity (§10) is structurally automatic from F = dA as a cohomological identity ($d^2 = 0$ on the de Rham complex of $H^1(\mathcal{G}(\Lambda))$).

Conditional on framework assumptions.

- The identification of the $H^1(\mathcal{G}(\Lambda))$ transport potential with the electromagnetic gauge potential A_μ rests on three independent structural inputs, each established in a separate paper: (i) the refinement-persistence results, which establish $H^1(\mathcal{G}(\Lambda))$ as the unique refinement-stable observable sector; (ii) the Wilson-loop identification in the Maxwell admissibility paper, which connects cocycle representatives to continuum gauge potentials and supplies the sign and normalisation of the kinetic term; and (iii) consistency with the α derivation series, which fixes the coupling strength and forecloses alternative U(1) identifications. The dependency graph is non-trivial: removing any one

input loosens the identification and opens room for the H^1 structure to be reinterpreted as some other (non-electromagnetic) U(1).
- The identification of the substrate record current J^μ with the physical electromagnetic four-current rests on the Single-Source Theorem and the fold/record-current architecture (§13's falsifiable check). The partial derivation of continuity from J^μ = ρ u^μ in §5 is conditional on this decomposition being SST-admissible and on u^μ admitting an independent substrate-level definition.
- *Lemma 7.0.1 (Primitive Substrate Field Exhaustion)* establishes ρ as the unique scalar primitive directly from SST Definition 1. The vector primitive u^μ is conditional on the matter-sector programme's substrate-level definition, which is the same gap §15 identifies as the dominant remaining open question. Closing the matter-sector gap closes this conditional clause completely. This is the only residual condition of the Catalogue Closure Theorem; the K = 7 catalogue and closure trichotomy are no longer load-bearing for the uniqueness claim of Theorem 1.
- The bidirectionality of Theorem 2 (gauge invariance ⇔ current conservation) assumes χ ranges over unrestricted smooth functions with admissible falloff at the boundary of the integration region. In topologically non-trivial geometries — compact spacetimes, non-trivial spatial topology, or restriction to compactly supported gauge transformations only — boundary effects can carry distinguishability information and the bidirectionality may weaken to a one-way implication. This is irrelevant for the leading-order construction on the emergent Lorentzian continuum but would need to be revisited in any topologically non-trivial extension.
- Pre-coupling unitary Hamiltonian evolution on $H^1(\mathcal{G}(\Lambda))$ is established by the admissibility-generator theorem. Post-coupling, only *linearised* unitarity (J^μ as a prescribed external source) is established here; full coupled unitarity awaits the matter-sector programme.

Conjectural / open.

- Full quantum electrodynamics, including the quantisation of the persistent-sector field and the resulting Fock structure.
- Charge quantisation, which requires the integer cohomology structure of $H^1(\mathcal{G}(\Lambda); \mathbb{Z})$ and is not addressed here.
- The numerical value of α (fine-structure constant); pursued separately in the α derivation series.
- Non-abelian generalisation, which requires additional substrate field content beyond the present construction.
- The species-resolved decomposition of J^μ into matter-sector currents, and full coupled unitarity that depends on it.

This labelling is conservative. Nothing in the "proven" category extends beyond what the proofs strictly establish; nothing in the "conditional" category is treated as derived; nothing in the "conjectural" category is assumed for the present results.

17. Open Problems

The construction of this paper opens, rather than closes, several lines of work. The first item below is the dominant remaining gap and is treated separately in §15; the others are subsidiary.

Matter-sector structure (dominant gap). As argued in §15, the microscopic origin of J^μ is the principal unresolved question. The matter-sector programme must deliver charged excitations as identifiable substrate structures, a species decomposition of J^μ into matter currents, spinorial structure connecting to the existing spin-as-double-cover work, and the eventual quantum field structure required for QED. Without this, the present paper is a structural waypoint rather than a physically complete derivation.

Charge quantisation. The persistent sector's cohomology has been treated here with real coefficients. The full charge-quantisation question requires the integer cohomology $H^1(\mathcal{G}(\Lambda); \mathbb{Z})$, and the discreteness of the substrate suggests this is structurally available. A clean derivation of the Dirac-style quantisation condition from the integer cohomology is the natural next paper in this strand.

Sign and signature of the kinetic term. This paper uses the $-\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$ kinetic Lagrangian without re-deriving it, inheriting it from the Maxwell admissibility paper. Verifying that the persistent-sector dynamics generate exactly this sign and normalisation under refinement — rather than, say, $+\frac{1}{4} F^2$ (which would invert the Hamiltonian and violate distinguishability) or a higher-derivative kinetic structure — is part of closing the structural circle. The Maxwell admissibility paper establishes this directly; downstream papers (this one included) should cite that result at every point of use, which §9 now does.

The numerical value of α. The fine-structure constant series ($\alpha^{-1} \approx 137.034$ to 15 ppm) sits in a separate strand of the programme. A clean connection between this paper's coupling structure and the closure-counting derivation of α is the natural next theoretical target.

Non-abelian generalisation. Extending the construction to non-abelian gauge structure requires additional substrate field content beyond what the current $K = 7$ catalogue admits. The structural shape of the required extension is suggested by the closure trichotomy theorem but has not been carried out explicitly. This is structurally tied to the catalogue scope of the Catalogue Closure Theorem discussed in §7.1; extending to non-abelian gauge structure would require introducing additional substrate-field primitives beyond ρ and u^μ, which would relax Lemma 7.0.1.

Relativistic completion and quantisation. The paper has worked at the level of classical field equations on the emergent Lorentzian continuum. Promoting A_μ and J^μ to operator-valued distributions, and constructing the resulting Fock space on the persistent sector, is the natural quantum-field-theoretic completion. This is downstream of the matter-sector programme.

Consistency with the fold/record-current uniqueness theorem. The companion check that this paper's J^μ coincides with the record current of *The Fold and the Record* (Theorem 12.2) should be carried out explicitly to close the structural circle.

---

18. Conclusion

The previous paper established that the persistent harmonic sector could not generate observable physical dynamics from the $K = 7$ gradient-flow catalogue alone, and identified this as an inertia problem requiring resolution.

The present paper provides the leading-order admissible resolution by showing that the missing physical structure emerges through admissible coupling to the record-current sector. The coupling is not postulated; by the Catalogue Closure Theorem of §7.0, it is the unique lowest-order term compatible with locality, refinement persistence, gauge redundancy, distinguishability preservation, and BCB current conservation. The same coupling forces the identification of gauge redundancy with current conservation, and yields Maxwell-form source dynamics under straightforward variation. The Bianchi identity comes free from the cohomological structure of the persistent sector itself. What is not achieved — and the dominant remaining gap, treated in §15 — is the microscopic origin of J^μ, without which the construction is a structural waypoint rather than a physically complete derivation.

The persistent cohomological sector supplies reversible Hamiltonian transport, gauge redundancy, refinement-stable Wilson observables, and the homogeneous Maxwell pair. The record-current sector supplies conserved source structure and the observable dynamical response of the coupled system. Their interaction reproduces the inhomogeneous Maxwell equation. The "inertia problem" is not a problem in this architecture; it is the structural signature of an unbroken gauge sector.

Within the VERSF programme this is the first explicit construction in which a reversible admissibility class and an irreversible admissibility class jointly generate a single observable physical dynamics. It also identifies the next priority of the programme without ambiguity: the matter-sector decomposition of J^μ, constrained by the Single-Source Theorem and tied to the integer cohomology of the persistent sector. Supplying that decomposition — and testing whether other Standard Model gauge structures admit analogous constructions traceable to differences in substrate field content — is the next stage of work.