

Electricity and Magnetism from Bit Conservation & Balance and Ticks-Per-Bit

A VERSF Maxwell Admissibility Theorem with Substrate Exclusion of Fundamental U(1) Magnetic Charge

What This Paper Claims, in Plain Language

Maxwell's equations of electricity and magnetism are usually treated as a fundamental starting point of physics. This paper argues they don't have to be. Given two simple principles — that information is conserved as it flows, and that moving information takes time — Maxwell's equations are not assumed; they are forced. They are the unique mathematical shape that information transport can take.

One striking prediction follows: fundamental magnetic monopoles cannot exist. No experiment has ever found one. Here that absence is not luck — it is required by the structure of the theory.

The paper is honest about what it does not do. It does not yet construct the underlying substrate from first principles. It does not address quantum electrodynamics, which would be a separate construction. It does not derive *which* current in nature is the electromagnetic one. These are identified as the next steps in a larger programme, not as gaps quietly papered over.

Abstract for the General Reader

Electricity and magnetism, as captured by Maxwell's equations, are usually taken in physics as a starting point — fundamental facts about the universe that the rest of the theory builds upon. This paper asks a deeper question: *why* should electricity and magnetism take the specific mathematical form they do, rather than some other form?

The answer offered here begins from two principles about how information might work at the most fundamental level of reality. The first, **Bit Conservation & Balance**, says that resolved information cannot be created or destroyed — it can only flow from one region to another, much as electric charge or energy is conserved in standard physics. The second, **Ticks-Per-Bit**, says that moving information takes time: there is a universal maximum speed at which information can propagate, set by the smallest possible "update step" of the underlying substrate.

Starting from these two principles, and adding a small number of further assumptions about how the substrate must behave (locality, gauge redundancy, the right symmetries, and a controlled

approximation valid at scales much larger than the substrate itself), the paper shows that Maxwell's equations are essentially forced. Three of the four equations turn out to be pure geometry — they carry no physical content, only mathematical structure. The fourth carries the actual physics: it specifies how electric charges and currents generate electromagnetic fields. The conservation of electric charge then falls out automatically, rather than being imposed by hand.

A striking prediction emerges along the way: fundamental magnetic charges (so-called magnetic monopoles) cannot exist as elementary particles in this framework. This matches what experiments have so far found — no magnetic monopoles have ever been detected — but here it is not an empirical accident, it is required by the structure of the theory itself. The paper is explicit about what would refute this prediction and the rest of the construction.

The paper is also honest about its main limitation. It argues that *if* a substrate satisfying the right principles exists, *then* Maxwell's equations follow. It does not yet build a concrete model of such a substrate — that construction is identified as the most important remaining task in the programme. A short discussion (open problem 5) notes that lattice gauge theory, already well established in standard physics, provides an existence proof that such substrates can be built, and lays out what a more VERSF-native version would need to demonstrate. A related question — how, at the substrate level, unresolved possibility becomes irreversible record (the so-called "measurement problem" in standard quantum theory) — is addressed in three companion papers and discussed as a bridge in §3.1; the Maxwell admissibility theorem does not depend on that question being settled, but the substrate-physical meaning of its key term "commitment" does.

The wider goal is to push the foundations of physics one layer deeper: rather than treating Maxwell's equations (and similar structures throughout the Standard Model) as fundamental, the framework treats them as the inevitable shape of information transport in a universe where information itself is conserved and propagates at finite speed.

Contents

What This Paper Claims, in Plain Language

Abstract for the General Reader

Abstract

1. Introduction
2. Substrate Axioms and the Admissibility Class
 - 2.5 Regularity Assumptions
 - 2.6 A Note on Substrate Ontology
3. Bit Conservation & Balance
 - 3.1 Measurement, Commitment, and Fact Formation
4. Finite Propagation from TPB
5. The Need for a Transport Field — Locality from TPB
6. Helmholtz Decomposition and the Source-Sector Identification

7. Gauge Redundancy as Informational Degeneracy
 - 7.1 Formal Structure of Informational Degeneracy
 - 7.2 Explicit Derivation
8. Homogeneous Transport Equations and the Substrate Exclusion of Fundamental $U(1)$ Magnetic Charge
 - 8.1 The Magnetic Charge Exclusion
 - 8.2 Faraday's Law as Sectoral Compatibility
9. The Unique Inhomogeneous Coupling
10. Wave Propagation in Vacuum
 - 10.1 The Wave Equations
 - 10.2 Saturation: A Separately Postulated Identification
11. Logical Dependency Map
 - 11.1 The Entire Maxwell Structure in Differential-Forms Language
12. The Maxwell Admissibility Theorem
 - 12.1 Scope of the Theorem
 - 12.2 The Current-Selection Problem
 - 12.3 A Candidate Current-Selection Criterion
 - 12.4 Proof Sketch
13. Relationship to Standard Gauge-Theoretic Uniqueness
14. What Would Falsify This Construction
 - 14.1 Near-Term Empirical Handles
15. Epistemic Status
16. A Minimal Substrate Toy — Preliminary Constructive Sketch
 - 16.1 Scope of This Section
 - 16.2 The Toy Substrate: Simplicial Transport Network
 - 16.3 Conserved Transport and Finite Propagation
 - 16.4 Gauge Redundancy from Substrate Degeneracy
 - 16.5 The Emergence Hierarchy
 - 16.6 What This Toy Does and Does Not Establish
 - 16.7 Explicit Discrete Dynamics and the Continuum Maxwell Limit
 - 16.8 Toward the $K = 7$ Wilson Limit
17. Conclusion and Open Problems
 - 17.1 What This Paper Actually Accomplishes
 - 17.2 Programme Architecture
 - 17.3 Known Missing Pieces
 - 17.4 Open Problems

Appendix A. Why Chern–Simons-Type Terms Fail Under (B4)

Appendix B. Anisotropic Wilson Coarse-Graining and the Continuum Limit

References

Abstract

We develop a transport-theoretic reconstruction of classical electromagnetism within the Void Energy-Regulated Space Framework (VERSF). Starting from two substrate-level principles — Bit Conservation & Balance (BCB) and Ticks-Per-Bit (TPB) — we argue that **Maxwell-form U(1) gauge transport for a conserved vector current** emerges as the surviving admissible linear first-order local transport theory within the admissibility class defined by substrate-motivated locality, gauge redundancy, closure-geometry covariance, and a controlled effective-field-theory truncation in the coherence scale ξ . The empirical identification of this transport theory with *electromagnetism* requires two additional inputs not derived in this paper: a saturation postulate identifying the wave speed with the TPB bound (see §10.2), and a current-identification matching J_c^μ to the empirical electromagnetic current (see §12.2). A third dependency, on a substrate-level measurement theory giving operational content to the term "commitment" itself, is anchored to three companion VERSF papers via the bridge sketched in §3.1; the Maxwell admissibility theorem does not require a completed measurement theory to derive Maxwell-form transport, but the substrate-physical interpretation of its premises does. With these identifications and the bridge in place, Maxwell electrodynamics is recovered. The result is a structural admissibility theorem within a substrate-motivated class, not an absolute inevitability theorem; we are explicit throughout about what lies inside the class, what is postulated externally, and what would falsify the construction.

The framework is **pre-geometric**: substrate-level transport admissibility precedes the emergence of effective spacetime geometry. Lorentz covariance enters not as an axiom but as an inherited result from a companion paper on emergent Lorentz invariance from substrate proto-time. The natural mathematical language of the transport structure is that of differential forms, with the field a closed two-form and its source the codifferential of that form set equal to the conserved current. Within this language the argument proceeds in four stages. First, BCB forces a conserved four-current of committed distinguishability and TPB forces this current to propagate at a finite substrate-determined speed. Second, the substrate principles motivate a specific admissibility class — locality (tied directly to TPB), gauge redundancy (read as informational degeneracy under equivalence classes of substrate encodings), closure-geometry covariance, and a controlled ξ -expansion. Third, within this class the physical transport field is forced into antisymmetric two-form structure, its exterior derivative vanishes identically (yielding the homogeneous Maxwell equations including the substrate exclusion of fundamental U(1) magnetic charge), and the only admissible source equation is the inhomogeneous Maxwell equation. Fourth, electricity and magnetism correspond respectively to the **divergence-driven** and **circulation-driven** source sectors of the commitment current, matching the empirical content of those terms.

The result does not assume Maxwell's equations a priori; it derives them as the unique closure-compatible transport geometry at leading order in ξ within the stated admissibility class. We are careful throughout to distinguish what is proven, what is conditional on the admissibility class, what is conditional on companion results in the VERSF corpus, what is interpretive, and what would falsify the construction.

1. Introduction

Maxwell's equations are conventionally introduced either as the codification of experimental phenomena (Coulomb, Ampère, Faraday) or as the consequence of imposing U(1) gauge invariance on a relativistic field theory. Both routes succeed empirically, but neither answers a structural question that VERSF takes seriously:

Why should the electromagnetic interaction take precisely the Maxwell form, and not some other linear or nonlinear transport structure?

The standard answer is that Lorentz invariance, locality, gauge symmetry, and renormalisability together pick out Maxwell uniquely at the lowest derivative order (see Wald 1984 and Weinberg 1995 for the canonical uniqueness arguments, and Deser 1970 for the related self-coupling perspective). This is correct, but it pushes the question back: *why those assumptions?*

VERSF approaches the problem from below the field level. Distinguishable physical state is taken to be the primitive notion, and the substrate is subject to two informational constraints:

- **Bit Conservation & Balance (BCB)** — committed distinctions are locally conserved. They may flow, redistribute, or reorganise, but cannot be created or annihilated arbitrarily.
- **Ticks-Per-Bit (TPB)** — the propagation of distinguishability is mediated by substrate update progression with a finite minimum cost per resolved bit, implying a universal finite-speed bound on commitment transport.

The central claim of this paper is:

Theorem (informal). Within the admissibility class defined by substrate-motivated locality, gauge redundancy, closure-geometry covariance, and the ξ -controlled effective-field-theory truncation specified in §2, the surviving admissible transport theory of a conserved vector current under BCB and TPB is **Maxwell-form U(1) gauge transport at leading order in ξ** . Maxwell electrodynamics is recovered conditionally on two further identifications (saturation, current-selection) developed in §10.2 and §12.2.

What we are *not* claiming, and what should not be read into the title:

- We are not claiming that Maxwell theory is absolutely inevitable from BCB and TPB alone.
- We are not claiming that fundamentally nonlocal, higher-order, non-gauge, or non-Lorentz-covariant transport structures are excluded as logical possibilities — only that they fall outside the substrate-motivated admissibility class.
- We are not claiming that the admissibility class itself is derived from BCB and TPB; we are claiming that it is *physically motivated by* substrate informational structure.
- **Crucially, we are not claiming that the construction selects the electromagnetic current.** BCB delivers a conserved current; the theorem then forces a Maxwell-form transport equation on *any* conserved vector current admitting the full admissibility class.

The substrate plausibly carries multiple conserved vector currents (fermion number, hypercharge-like quantities), each of which independently satisfies BCB. The empirical identification of *the* electromagnetic current is an external matching, not a substrate-level selection theorem. We develop this point in §12.2.

The substantive VERSF claim is that the admissibility constraints which standard physics motivates by appeal to renormalisability or experimental success are here motivated by substrate-level informational structure, and that linearity and derivative-order — classically introduced as separate assumptions — are unified under a single ξ -truncation tied to the substrate coherence scale.

The paper is organised as follows. §2 states the substrate axioms and the admissibility class together up front. §3 introduces the commitment current. §4 derives the finite propagation bound from TPB with explicit dimensional care. §5 motivates the transport potential and the substrate-level locality assumption. §6 performs the Helmholtz split and gives the source-sector identification. §7 derives gauge redundancy and the antisymmetric tensor structure. §8 derives the homogeneous Maxwell equations and the substrate exclusion of fundamental U(1) magnetic charge. §9 establishes the source equation. §10 derives wave propagation and treats the saturation claim as a separately flagged assumption. §11 presents a logical dependency map. §12 states the theorem, addresses the current-selection problem, **proposes a candidate four-condition current-selection criterion (§12.3)**, gives the scope, and presents a proof sketch. §13 compares to standard gauge-theoretic uniqueness arguments. §14 gives a falsification catalogue, including **near-term empirical handles (§14.1) — precision photon dispersion, magnetic-charge exclusion, and the α - Λ joint stability prediction from the $K = 7$ paper**. §15 catalogues the epistemic status of each claim. **§16 presents a preliminary constructive sketch of a minimal substrate toy, including explicit discrete dynamics with a continuum Maxwell limit (§16.7) and the $K = 7$ -native lattice action targeting the $K = 7$ Wilson Limit paper (§16.8)**. §17 concludes with open problems. **Appendix B presents the anisotropic Wilson framework in which the saturation argument of §10.2 can be made rigorous via RG analysis.**

2. Substrate Axioms and the Admissibility Class

We collect here, in advance of their use, the full set of assumptions on which the uniqueness result rests.

Substrate axioms (taken as given for this paper):

- **(A1) BCB.** There exists a four-current J_{μ}^c of committed distinguishability satisfying the local conservation law $\partial_{\mu} J_{\mu}^c = 0$.
- **(A2) TPB.** Commitment propagation has a finite invariant maximum speed c_c set by the substrate update structure. This axiom carries two distinct contents that should be named separately:

- **(A2.i) Finite-speed bound.** The substrate update process produces a finite invariant speed bound $c_c = \xi / (N_b \tau_s)$, beyond which committed distinguishability cannot propagate.
- **(A2.ii) Atomic substrate update.** The substrate's fundamental update rule is a one-tick, local-neighbourhood operation: each substrate site updates from its immediate neighbourhood at each tick, without integral kernels over distant history. Effective integral forms can emerge from coarse-graining of such atomic dynamics, but they are not the fundamental rule.

(A2.i) and (A2.ii) are logically independent: a substrate could in principle satisfy (A2.i) without (A2.ii) (e.g., a causal-integral-kernel theory respecting finite light speed), or (A2.ii) without (A2.i) (e.g., a local update rule allowing arbitrary fast propagation, which we exclude). The combination is what TPB asserts as a substrate axiom. (A2.ii) does substantive work in motivating (B1) below; we name it explicitly so that this dependence is visible.

The admissibility class is the set of transport theories satisfying the following constraints; the substrate motivation for each is given when it is first invoked:

- **(B1) Locality.** Transport dynamics are expressible through local differential relations in a potential A_c^μ . Motivated directly by TPB: nonlocal transport laws require substrate updates to depend on unresolved distant commitment configurations, violating finite local update progression (§5).
- **(B2) Controlled ξ -expansion.** Defining the expansion parameter

$$\varepsilon = \xi / L \ll 1$$

where ξ is the substrate coherence length and L is the macroscopic scale of variation, the admissible transport Lagrangian is organised as a power series in ε . **Maxwell theory occupies the $O(\varepsilon^0)$ sector** — linear in the field tensor and first-order in derivatives of the field tensor (equivalently, second-order in the potential). Higher operators (nonlinear field self-couplings, higher-derivative kinetic terms) appear at $O(\varepsilon)$, $O(\varepsilon^2)$, and beyond, recovering Born–Infeld and Euler–Heisenberg as leading nonlinear substrate corrections in the regime where ε is no longer small. **B2 replaces and refines** the earlier separate assumptions of "linearity" and "first-order derivative structure."

Physically, the $O(\varepsilon^0)$ sector corresponds to transport dynamics observed at scales $L \gg \xi$, where substrate discreteness and nonlinear overlap effects between commitment events within a coherence volume are negligible. In standard EFT language, **Maxwell electrodynamics is interpreted within the present framework as the infrared transport limit of the substrate dynamics** — the long-wavelength, low-energy effective theory that emerges when probes are too coarse to resolve substrate granularity. Nonlinear extensions are the leading ultraviolet corrections as probes approach the coherence scale.

A flag on the EFT hierarchy. In standard Wilsonian EFT, the suppression of higher-dimension operators by powers of E/Λ is *derived* by integrating out high-energy modes. Here we are asserting an analogous suppression by powers of $\varepsilon = \xi/L$ without exhibiting the substrate integration that would produce it. The hierarchy in (B2) is therefore structural rather than constructive at present: the substrate justification is currently a power-counting analogy, not a Wilsonian derivation, and would become a theorem only when an explicit substrate model is given and shown to produce ε -suppressed operators under coarse-graining (see open problem 5).

- **(B3) Gauge redundancy.** The transformation $A_c^\mu \rightarrow A_c^\mu + \partial^\mu \chi$ for arbitrary scalar χ generates an equivalence class of substrate encodings producing identical measurable commitment transport. Gauge symmetry is therefore informational degeneracy, not mathematical convenience (§7).
- **(B4) Closure-geometry covariance.** The transport equations transform covariantly under the substrate's closure-geometry symmetry group. Within the regime where emergent Lorentz invariance is established (companion paper on proto-time and emergent Lorentz invariance), this reduces to Lorentz covariance in the standard sense. We invoke that emergent-Lorentz result as an established output of the VERSF corpus rather than rederiving it.

The structural claim of the paper is that under (A1)–(A2) and (B1)–(B4), Maxwell-form U(1) gauge transport is the surviving theory in the admissibility class at $O(\varepsilon^0)$. Outside the class — fundamentally nonlocal theories, non-gauge theories, higher-derivative theories admitted as leading-order rather than corrections — the result does not apply.

2.5 Regularity Assumptions

To centralise scattered mathematical-precision conditions used throughout the paper, we collect the regularity assumptions here:

- **Smoothness.** All fields appearing in the construction — A_c^μ , $F_c^{\mu\nu}$, J_c^μ , the scalar gauge function χ — are assumed smooth (C^∞) on the substrate manifold. Where weaker regularity suffices, C^2 is the minimal smoothness required for the commutativity of partial derivatives used in the Bianchi identity and the gauge calculation of §7.
- **Manifold structure.** The substrate is treated as a smooth four-dimensional manifold (within the regime where emergent Lorentz invariance applies) equipped with a Lorentzian metric. The codifferential δ and the Hodge star $*$ are defined with respect to this metric.
- **Topology.** For the local field-theoretic argument and the Bianchi-identity derivation, we assume the substrate manifold is topologically trivial in the sense that A_c can be defined as a single globally-valid one-form. The magnetic-charge exclusion of §8.1 carries this as an explicit conditional (see §8.1 for the topological caveat).
- **Decay conditions.** For the Helmholtz decomposition of §6, fields are assumed to decay sufficiently rapidly at spatial infinity to license the standard decomposition. For cosmological applications (closed universes, persistent large-scale structure), the standard decay assumption fails and the Helmholtz decomposition must be replaced by its

compact-manifold analogue (Hodge decomposition into harmonic, exact, and co-exact pieces); see the cosmological caveat in §6.

- **Codifferential domain.** The identity $\delta^2 = 0$ holds on sufficiently regular differential forms over the metric specified above. Concretely, smooth (or C^2) forms with support compatible with the Hodge star satisfy this without further qualification.

These regularity conditions are sufficient for all derivations in the body of the paper. Strengthenings (analytic regularity, restricted support, etc.) are not required.

2.6 A Note on Substrate Ontology

The present construction does not commit to a specific ontological character for the substrate. In particular:

- **Classical or quantum?** The framework is compatible with a substrate whose microscopic update structure may be classical-discrete, classical-continuous, pre-quantum, or proto-quantum in the sense of admitting quantum generalisation. BCB and TPB are formulated at the level of conserved transport and finite propagation; both can be realised in classical and quantum substrate structures, and the present paper does not require either reading.
- **Fundamental or emergent gauge redundancy?** Gauge redundancy (B3) is read here as informational degeneracy — substrate encoding ambiguity producing equivalence classes. We remain agnostic as to whether this redundancy is *ontologically fundamental* (a primitive feature of the substrate's representational structure) or *emergent from coarse-graining* (an artefact of effective descriptions of a more fundamental substrate without intrinsic redundancy). The quotient-space formalism of §7.1 applies in either interpretation, and the empirical predictions of the construction do not depend on resolving this question.

These agnosticisms are not evasions; they are honest statements that the present paper's results — Maxwell-form gauge transport from admissibility — are robust across the relevant ontological choices, and that those choices belong to the deeper substrate-construction programme rather than to the present transport-admissibility argument.

3. Bit Conservation & Balance

Operational definition of "commitment" — a working definition with an honest caveat.

Before introducing notation, we ground the substrate vocabulary. The term *commitment* is intended to refer to substrate-resolved distinguishability: a state transition that has become locally encoded within the substrate update structure and is therefore transport-capable. A committed bit is, on this reading, a registered binary distinction that the substrate has resolved, can move between neighbouring regions, and can interact with other resolved distinctions. *Committed distinguishability* is the density of such resolved distinctions per unit substrate

volume; *commitment transport* is their flow; *commitment imbalance* is a local excess or deficit relative to surrounding regions.

We flag, however, that this is currently a *definitional* characterisation rather than a strictly operational one. A genuinely operational definition would specify a measurement procedure — what physical apparatus or substrate-level interaction distinguishes a committed bit from an uncommitted state. The present paper does not supply such a procedure.

The operational gap is two-layered, and only one layer is addressed by open problem 5:

- *Layer A (substrate construction)*. What is the substrate? Until a concrete substrate model is exhibited (OP 5), there is nothing for a measurement to be defined relative to. This layer is genuinely deferred to substrate construction.
- *Layer B (substrate-level measurement theory)*. Even given a complete substrate model, the operational definition requires specifying which substrate-level interactions count as *measurements* — what physical process resolves an uncommitted state into a committed one, and how that resolution is registered. This is the standard quantum-measurement problem in substrate vocabulary: any sufficiently complete substrate dynamics will face it. **OP 5 addresses Layer A but does not by itself close Layer B.** A preliminary candidate for Layer B is sketched in §16.6: commitment as persistent multi-site coherence — a bit is committed when its transport pattern stably propagates across multiple substrate sectors under continued evolution, in a manner reminiscent of decoherence and environment-induced classicality. This is a preliminary direction rather than a finished measurement theory; the rigorous version remains its own open problem of the corpus.

Until both layers are addressed, "commitment" should be read as a labelling that will receive operational content when the substrate is given *and* the substrate-measurement question is answered. At effective scales ($L \gg \xi$), this collection of quantities behaves mathematically as conserved charge-current transport in the standard sense — which is the bridge that the rest of this section formalises and which is fully operational without a substrate-level definition of commitment.

3.1 Measurement, Commitment, and Fact Formation

The present paper does not attempt to solve the quantum measurement problem. However, its use of the term "commitment" depends on a future substrate-level account of how unresolved distinguishability becomes an irreversible physical record. We therefore record here the direction in which the corpus now supplies that account, so that the substrate-physical meaning of "commitment" in the present paper is anchored to specific companion work rather than left wholly open.

Three companion papers in the VERSF corpus now supply the intended Layer B direction:

- One companion paper interprets quantum measurement as **relational closure**: the quantum state is not a pre-existing object with hidden properties, but a structure of

possible closures that becomes definite only when a stable record is formed. Measurement is therefore not a "collapse" event imposed externally; it is the formation of the record itself.

- A second ties irreversible fact formation to **non-recombinable substrate topology**: once information pathways support cyclic trapping, alternative outcomes become locally unrecoverable. This supplies the substrate-physical mechanism by which a tentative record becomes a permanent one.
- A third models outcome selection as a **tick race** — a first-passage process in which decohered branches generate microscopic ticks, with the first threshold-crossing tick producing the effectively irreversible committed bit. The substrate microphysics may remain reversible at the global level; what the tick race produces is *local non-recombinability* — the practical inability to undo the selection once it has occurred — rather than absolute thermodynamic irreversibility. This supplies the dynamical mechanism by which one branch among many becomes the committed one.

Together these papers articulate the chain

relational possibility → decoherence → tick race → topological trapping → effectively irreversible committed bit → committed distinguishability,

which is the proposed Layer B completion of the present paper's operational gap. Each arrow corresponds to a substrate-level process specified in the cited papers.

Epistemic discipline. These companion papers presently provide a structured *interpretive and dynamical* framework for Layer B, rather than a closed derivation from the substrate microphysics itself. The chain above is a coherent programme — relational possibility supplies the pre-commitment ontology, decoherence supplies the branching dynamics, the tick race supplies the selection mechanism, and topological trapping supplies the local non-recombinability — but the full substrate-microphysical derivation of the chain remains companion-paper work in progress. The bridge in this section therefore points at a research programme of specified shape, not at a finished theorem.

Implication for the present construction, and a scope clarification. The current $J_{c^{\mu}}$ should be understood as the effective current of *already-committed records* — distinguishabilities that have passed through the full chain above and emerged as committed bits — not as a current of unresolved quantum possibilities. The Maxwell admissibility theorem begins after measurement and commitment have occurred. It does not require a completed measurement theory to derive Maxwell-form transport, *but* it does require one to give full substrate-physical meaning to the word "commitment."

A scope ambiguity that this resolves. Reading $J_{c^{\mu}}$ as the record-level current has a substantive consequence: BCB is then a conservation law for *records*, not for *quantum amplitudes*. The Maxwell admissibility theorem applies to the classical, record-level current — that is, **to classical electromagnetism**. Quantum electrodynamics, whose four-current is the expectation value of an operator on a Hilbert space of pre-commitment states, is not directly addressed by this theorem. QED in the present framework is a separate construction *layered on*

top of the measurement chain: above the substrate scale, the pre-commitment relational possibilities supply the quantum-mechanical amplitudes, the measurement chain supplies the commitment process, and the record-level current that emerges satisfies the classical Maxwell theorem of this paper. This is a real architectural choice — quantum EM is not a corollary of the present construction but a separate problem requiring its own substrate-level derivation. We flag it here so that the theorem's scope is unambiguous: classical Maxwell electrodynamics, with QED treated as a layered construction whose substrate-level derivation is its own programme.

This is a bridge between the Maxwell construction and the substrate measurement programme, not a claim that the present paper solves measurement. The Layer B work is done in the companion papers; the strength of this bridge — and therefore the substrate-physical force of the Maxwell theorem's premises — scales with the strength of the measurement-papers' Layer B account. As that account matures, so does the substrate-physical interpretation of every "commitment" term in the present paper.

Let the substrate carry a committed distinguishability density

$$\rho_c(\mathbf{x}, t)$$

representing the local density of resolved bits at substrate position \mathbf{x} and proto-time t . Associated with this density is a commitment transport current

$$\mathbf{J}_c(\mathbf{x}, t).$$

BCB asserts that committed distinguishability is locally conserved:

$$\partial \rho_c / \partial t + \nabla \cdot \mathbf{J}_c = 0 \quad (1)$$

This is a strict conservation law: commitment cannot leak. It can be re-encoded, redistributed, or moved, but the integrated balance is preserved.

In four-vector form, defining

$$J_c^\mu = (c_c \rho_c, \mathbf{J}_c), \quad (2)$$

where c_c is the propagation bound derived in §4, equation (1) becomes the manifestly covariant continuity equation

$$\partial_\mu J_c^\mu = 0. \quad (3)$$

Equation (3) is the substrate-level input from BCB. The rest of the paper is the systematic determination of the transport law that generates and constrains J_c^μ under the admissibility class.

A bridge to standard vocabulary. At the effective transport level, J_c^μ behaves mathematically as a conserved charge-current four-vector, satisfying the same continuity

equation as electric four-current in standard electrodynamics. The substrate vocabulary — "committed distinguishability," "commitment transport," "commitment imbalance" — refers to substrate origin rather than to any departure from standard transport behaviour. Readers may, without loss, read J_c^μ as a charge-current four-vector and apply the usual intuitions about conserved currents; the VERSF claim is about *why* such a current and *why* the transport law that governs it take their specific forms, not about replacing the formal structure of those quantities.

4. Finite Propagation from TPB

TPB asserts that resolving a bit of distinguishability requires a finite minimum substrate update cost. To make this dimensionally precise, we introduce three substrate quantities:

- ξ : the **coherence length**, the spatial extent over which a single bit is considered locally resolved.
- N_b : the **ticks-per-bit count** (dimensionless), the number of substrate updates required to commit one bit.
- τ_s : the **substrate tick duration**, the proto-time interval between successive substrate updates.

The maximum speed at which committed distinguishability can be transported through the substrate is then

$$c_c = \xi / (N_b \tau_s). \quad (4)$$

This is dimensionally a length per time. Earlier presentations of TPB sometimes wrote $c_c = \xi/\tau_b$ with τ_b serving as a compressed shorthand for $N_b \tau_s$; here we separate the count from the duration to avoid ambiguity. The product $N_b \tau_s$ is the **commitment time per bit**.

c_c is a finite, substrate-determined bound. Any admissible transport theory must therefore be hyperbolic in character — it must admit finite-speed wave propagation and cannot be elliptic-only or instantaneous.

In what follows we treat c_c as the operative invariant speed in the transport equations. The identification of this c_c with the empirical speed of light in vacuum is a separate physical postulate: it asserts that electromagnetic propagation **saturates** the TPB bound rather than propagating below it. This identification is consistent with experiment but is not derived from BCB and TPB alone.

5. The Need for a Transport Field — Locality from TPB

BCB gives us a conserved current. TPB gives us a finite propagation speed. Neither, by itself, gives us a field. The next structural question is: what is the minimal additional structure required to specify how \mathbf{J}_c responds to its own configuration history and to external influences?

The substrate argument for locality. The locality assumption (B1) requires careful statement because "nonlocality" has at least two senses. (i) *Spatial nonlocality*: instantaneous distant influence — an updated state depending on the simultaneous state of arbitrarily distant regions. This is forbidden by (A2.i) directly: the finite invariant speed bound precludes such instantaneous coupling. (ii) *Integral-kernel nonlocality*: dynamics expressed through integrals over the past light cone (e.g., retarded Green's functions). This *is* causally admissible — integral kernels respect finite-speed propagation by construction — and is in fact the natural form for any effective theory after coarse-graining.

What (B1) actually requires is the stronger sense, and it is motivated directly by (A2.ii) — **atomic substrate update**. The substrate-level argument is not that integral kernels are forbidden — they aren't by causality alone — but that the *fundamental update rule* of the substrate is, by (A2.ii), a one-tick, one-local-computation operation. Integral forms can and do emerge after coarse-graining, but they emerge *from* underlying differential dynamics, not as fundamental rules. (A2.i) by itself does not force this; it is (A2.ii) that does.

So (B1) reads as: the substrate-level dynamics are differential and local because (A2.ii) requires it; effective integral forms in the IR are then accepted as derived rather than postulated. **The chain (A2.ii) → (B1) is a substantive substrate-level claim, not a near-derivation from finite propagation alone**, and was the hidden axiom doing the work in earlier drafts of this section. By naming (A2.ii) explicitly in §2, the substrate dependence of (B1) is now visible rather than implicit. A rigorous derivation of (A2.ii) from sharper VERSF axioms remains open (see open problem 4).

A purely current-level theory expressed through integral kernels over the past light cone is therefore not forbidden by causality, but it is incompatible with (A2.ii)'s commitment that updates are atomic. The standard mathematical remedy is to introduce a transport potential whose local derivatives encode the relevant dynamics. In accordance with (B1) we therefore postulate the existence of a four-potential

$$A_c^\mu = (\varphi_c / c_c, \mathbf{A}_c) \quad (5)$$

such that the transport dynamics can be expressed locally in terms of A_c^μ and its derivatives.

Two remarks. First, the *existence* of such a potential is the operational content of (B1); it remains a postulate in the sense that the explicit construction of A_c^μ from substrate microstate is a separate problem. Second, the potential is not directly observable — this non-observability is the substrate foundation for gauge redundancy (B3), which we develop in §7.

6. Helmholtz Decomposition and the Source-Sector Identification

Subject to standard regularity and decay conditions (assumed to hold for any physically realisable commitment configuration in a local field-theoretic setting), the spatial current \mathbf{J}_c admits the Helmholtz decomposition

$$\mathbf{J}_c = \mathbf{J}_L + \mathbf{J}_T \quad (6)$$

with

$$\nabla \times \mathbf{J}_L = 0, \quad (7)$$

$$\nabla \cdot \mathbf{J}_T = 0. \quad (8)$$

A restriction to flag: decay at spatial infinity is not the right condition on cosmological scales — commitment configurations on a closed universe, or on a noncompact one with persistent large-scale structure, need not vanish at infinity. The Helmholtz decomposition then requires either restriction to a finite spatial region with chosen boundary conditions, or replacement by an analogue on compact manifolds (Hodge decomposition into harmonic, exact, and co-exact pieces). For the local field-theoretic argument that follows, the standard decay assumption suffices; for cosmological applications of the resulting equations, an appropriate replacement is needed.

The longitudinal part \mathbf{J}_L carries divergence — it represents net displacement of commitment imbalance. The transverse part \mathbf{J}_T carries rotational closure — it represents commitment circulation without net displacement.

Within VERSF this decomposition defines two source sectors:

- **Divergence-driven sector** — commitment configurations whose dynamics are driven by the longitudinal current \mathbf{J}_L (sources of $\nabla \cdot \mathbf{E}_c$ via charge density).
- **Circulation-driven sector** — commitment configurations whose dynamics are driven by the transverse current \mathbf{J}_T (sources of $\nabla \times \mathbf{B}_c$, together with displacement current).

These sectors correspond empirically to electric and magnetic source behaviour, and we will use the labels "electricity" and "magnetism" where the empirical reading is the relevant point. But the substrate-native terminology is *divergence-driven* and *circulation-driven*, which keeps the source/field distinction clean.

A note on terminology. The electric field \mathbf{E}_c is not itself "longitudinal" in general — radiation \mathbf{E} is transverse, electrostatic \mathbf{E} has longitudinal components, and the two cannot be mapped cleanly onto the Helmholtz sectors of \mathbf{J}_c . What is clean is the mapping at the source: $\nabla \cdot \mathbf{E}_c$ is fed by the divergence-driven sector, $\nabla \times \mathbf{B}_c$ is fed by the circulation-driven sector. Earlier VERSF

presentations used a looser formulation that conflated source and field; this paper sharpens the language.

7. Gauge Redundancy as Informational Degeneracy

We now ask: what is the most general physical field that can be constructed from A_c^μ and its first derivatives?

Before answering, we strengthen the substrate motivation for gauge redundancy. The earlier formulation — " A_c^μ is non-observable, so reparametrisations that leave the substrate content unchanged should leave physical predictions unchanged" — is correct but philosophically thin. The deeper VERSF reading is this:

Gauge redundancy expresses equivalence classes of substrate encodings that produce identical measurable commitment transport. The potential A_c^μ is therefore an informational generating structure for local dynamics, not a directly measurable substrate object. Multiple distinct encodings of the substrate microstate generate the same observable commitment current and the same observable field; the transformations that map between them are not "mathematical freedoms" but a structural feature of the substrate's informational organisation.

Gauge symmetry, under this reading, is *informational degeneracy* — redundant representational structure in the substrate's encoding, not a free choice of coordinates.

7.1 Formal Structure of Informational Degeneracy

To make the informational-degeneracy reading mathematically precise, we introduce the following structure.

Let \mathcal{S} denote the space of substrate microstates, and let \mathcal{O} denote the space of observable commitment-transport configurations (currents, fields, and their measurable consequences). The **observable map**

$$\mathcal{O} : \mathcal{S} \rightarrow \mathcal{O}$$

assigns to each substrate microstate $s \in \mathcal{S}$ the observable commitment-transport configuration it generates. Concretely, $\mathcal{O}(s)$ returns the equivalence class of measurable transport quantities generated by s — including the conserved current J_c^μ , the field tensor $F_c^{\mu\nu}$, and gauge-invariant holonomies of the form $\oint A_c \cdot dl$ around closed loops (the substrate-level Aharonov–Bohm phases). What \mathcal{O} does *not* return is the value of the potential A_c^μ at any point — that is precisely the unobservable component the quotient construction is designed to factor out.

Define an **equivalence relation** on \mathcal{S} :

$$s_1 \sim s_2 \Leftrightarrow \mathcal{O}(s_1) = \mathcal{O}(s_2). \quad (\star)$$

Two substrate microstates are equivalent precisely when they produce identical measurable commitment transport. The **equivalence classes**

$$[s] = \{ s' \in \mathcal{S} : s' \sim s \}$$

form the **quotient space** \mathcal{S} / \sim , which is the space of *physically distinguishable* substrate configurations.

Gauge transformations are precisely those maps $\phi : \mathcal{S} \rightarrow \mathcal{S}$ satisfying $\mathcal{O} \circ \phi = \mathcal{O}$ — they permute substrate states within equivalence classes without crossing class boundaries. They act trivially on \mathcal{S} / \sim .

In this language:

- The potential A_c^μ is a coordinate on \mathcal{S} — it identifies a specific substrate microstate within an equivalence class.
- The field tensor $F_c^{\mu\nu}$ is a well-defined function on \mathcal{S} / \sim — it is constant on equivalence classes and therefore observable.
- The gauge transformation (11) is a concrete realisation of the trivial action on \mathcal{S} / \sim : $A_c^\mu \rightarrow A_c^\mu + \partial^\mu \chi$ moves between substrate microstates within a single equivalence class, leaving $F_c^{\mu\nu}$ invariant.

This is the precise sense in which A_c^μ is "non-observable" and $F_c^{\mu\nu}$ is "physical": one lives on the unquotiented space \mathcal{S} , the other on the quotient \mathcal{S} / \sim . The substrate carries more degrees of freedom than the observable map distinguishes; gauge symmetry is the redundant representational structure that tracks this disparity.

An honest acknowledgement of what this section achieves. The formal structure introduced above — observable map, equivalence relation, quotient space — is mathematically standard. It is the same structure that underwrites gauge theory in any conventional treatment; we have not added new mathematics. What §7.1 adds is **substrate-physical interpretation**: the assertion that \mathcal{S} is the substrate microstate space (not just an abstract configuration space), that \mathcal{O} is a *physical* observation map (not just a projection chosen by convention), and that the disparity between \mathcal{S} and \mathcal{S} / \sim reflects substrate encoding ambiguity rather than mathematical bookkeeping. The interpretive content is real but the formalisation does not, by itself, deliver mathematical novelty. Genuine novelty would require exhibiting \mathcal{S} explicitly — a concrete substrate microstate space with a constructive observation map whose redundancy realises the gauge group — which is exactly what a toy substrate model (open problem 5) would supply. Until then, §7.1 should be read as: *the standard quotient-space structure of gauge theory, given a substrate-physical interpretation.*

Fundamental or emergent? We reiterate the agnosticism of §2.6: the quotient-space formalism is compatible with two readings — gauge redundancy as *ontologically fundamental* (a primitive feature of the substrate's representational structure) or as *emergent* from coarse-graining of a

more fundamental substrate without intrinsic redundancy. The present paper does not arbitrate between these readings; the empirical predictions developed below hold under either.

The remainder of this section displays $F_{\underline{c}}^{\wedge\mu\nu}$ explicitly as the surviving gauge-invariant first-derivative object — the concrete realisation, in components, of the abstract quotient structure just introduced.

7.2 Explicit Derivation

The general first-derivative object built from $A_{\underline{c}}^{\wedge\mu}$ is the rank-two tensor

$$\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu}, \quad (9)$$

with 16 components, decomposable into its symmetric and antisymmetric parts:

$$\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} = \frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} + \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}) + \frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} - \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}). \quad (10)$$

We impose gauge redundancy (B3):

$$A_{\underline{c}}^{\wedge\mu} \rightarrow A_{\underline{c}}^{\wedge\mu} + \partial^{\wedge\mu} \chi \quad (11)$$

for arbitrary scalar χ .

Why specifically the gradient form $\partial^{\wedge\mu} \chi$? The gradient form is the unique local additive redundancy preserving the derivative-order structure of the theory while introducing no independent transport degrees of freedom. More general additive transformations $A_{\underline{c}}^{\wedge\mu} \rightarrow A_{\underline{c}}^{\wedge\mu} + B^{\wedge\mu}$ with $B^{\wedge\mu}$ a generic vector field would carry their own commitment content — they would be physical, not redundant. The restriction to total-derivative form $B^{\wedge\mu} = \partial^{\wedge\mu} \chi$ is precisely the restriction to additive structure that is generated by a single scalar function and therefore carries no independent commitment imbalance. The redundancy is exactly the size of the substrate's encoding ambiguity, and no larger.

Under (11), the symmetric part of (10) transforms as

$$\frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} + \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}) \rightarrow \frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} + \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}) + \partial^{\wedge\mu} \partial^{\wedge\nu} \chi, \quad (12)$$

picking up the non-vanishing term $\partial^{\wedge\mu} \partial^{\wedge\nu} \chi$. The symmetric part is therefore not gauge-invariant.

The antisymmetric part transforms as

$$\frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} - \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}) \rightarrow \frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} - \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}) + \frac{1}{2}(\partial^{\wedge\mu} \partial^{\wedge\nu} \chi - \partial^{\wedge\nu} \partial^{\wedge\mu} \chi) = \frac{1}{2}(\partial^{\wedge\mu} A_{\underline{c}}^{\wedge\nu} - \partial^{\wedge\nu} A_{\underline{c}}^{\wedge\mu}), \quad (13)$$

because partial derivatives commute. The antisymmetric combination is gauge-invariant.

Define the **commitment field tensor**

$$F_{\text{c}}^{\wedge\mu\nu} \equiv \partial^{\wedge\mu} A_{\text{c}}^{\wedge\nu} - \partial^{\wedge\nu} A_{\text{c}}^{\wedge\mu}. \quad (14)$$

We collect the content of the calculation just performed as a formal lemma, since it carries the structural weight of the entire construction.

Lemma (Gauge-Invariant First-Derivative Uniqueness).

Let $A_{\text{c}}^{\wedge\mu}$ be a smooth one-form potential on the substrate manifold, subject to the gauge redundancy $A_{\text{c}}^{\wedge\mu} \rightarrow A_{\text{c}}^{\wedge\mu} + \partial^{\wedge\mu} \chi$ for arbitrary smooth scalar χ . **Among rank-two tensors built linearly from one factor of ∂A_{c}** , the surviving antisymmetric gauge-invariant first-derivative tensor is

$$F_{\text{c}}^{\wedge\mu\nu} = \partial^{\wedge\mu} A_{\text{c}}^{\wedge\nu} - \partial^{\wedge\nu} A_{\text{c}}^{\wedge\mu}.$$

(Tensors built from higher powers of derivatives or nonlinear combinations are not covered by this Lemma; their exclusion is governed by the EFT truncation (B2), not by gauge invariance alone.)

Proof. The most general rank-two tensor built linearly from one factor of ∂A_{c} is $\partial^{\wedge\mu} A_{\text{c}}^{\wedge\nu}$, decomposing into symmetric and antisymmetric parts as in (10). Under the gauge transformation, the symmetric part picks up $\partial^{\wedge\mu} \partial^{\wedge\nu} \chi$ as in (12), which is nonzero for generic χ ; hence the symmetric part is not gauge-invariant. The antisymmetric part transforms by $\frac{1}{2}(\partial^{\wedge\mu} \partial^{\wedge\nu} \chi - \partial^{\wedge\nu} \partial^{\wedge\mu} \chi)$, which vanishes by the commutativity of partial derivatives on smooth functions; hence the antisymmetric part is gauge-invariant. Any tensor obtained from $\partial^{\wedge\mu} A_{\text{c}}^{\wedge\nu}$ by linear combination decomposes into symmetric and antisymmetric components; the antisymmetric component is the surviving gauge-invariant piece, and it is precisely $F_{\text{c}}^{\wedge\mu\nu}$ as defined in (14). ■

Conclusion. Within the class of tensors built linearly from one factor of ∂A_{c} , the surviving gauge-invariant antisymmetric piece is $F_{\text{c}}^{\wedge\mu\nu}$. Nonlinear or higher-derivative gauge-invariant structures are governed by (B2) — the ξ -EFT truncation — rather than by gauge invariance alone; they enter at $O(\varepsilon)$ and beyond, not at $O(\varepsilon^0)$.

A subtlety worth flagging: gauge redundancy must constrain not just the *field* but the *equations*. Gauge-variant scalars and vectors built from $A_{\text{c}}^{\wedge\mu}$ and its derivatives — for instance $\partial^{\wedge\nu}(\partial_{\wedge\mu} A_{\text{c}}^{\wedge\mu})$, which picks up $\partial^{\wedge\nu} \square\chi$ under (11) — cannot appear in physical equations either, because their inclusion would render the equations themselves gauge-variant. (B3) acts in two places: it picks out $F_{\text{c}}^{\wedge\mu\nu}$ as the unique gauge-invariant field, and it forbids gauge-variant operators from entering the source equation in §10.

In differential-form language (which is the natural mathematical setting flagged in the abstract), A_{c} is a one-form, $F_{\text{c}} = dA_{\text{c}}$ is the exact two-form built from its exterior derivative, and gauge redundancy $A_{\text{c}} \rightarrow A_{\text{c}} + d\chi$ leaves F_{c} invariant because $d^2 = 0$.

We may unpack $F_c^{\mu\nu}$ into electric and magnetic components by defining

$$\mathbf{E}_c \equiv -\nabla\phi_c - \partial\mathbf{A}_c/\partial t, \quad (15)$$

$$\mathbf{B}_c \equiv \nabla \times \mathbf{A}_c. \quad (16)$$

The components of $F_c^{\mu\nu}$ arrange so that the spatial-temporal entries reproduce \mathbf{E}_c and the spatial-spatial entries reproduce \mathbf{B}_c in the standard manner.

8. Homogeneous Transport Equations and the Substrate Exclusion of Fundamental U(1) Magnetic Charge

Because F_c is the exterior derivative of A_c ($F_c = dA_c$), its own exterior derivative vanishes identically:

$$dF_c = d(dA_c) = 0. \quad (17)$$

In index notation this is the Bianchi identity

$$\partial_\lambda F_c^{\mu\nu} + \partial_\mu F_c^{\nu\lambda} + \partial_\nu F_c^{\lambda\mu} = 0, \quad (18)$$

or equivalently with totally antisymmetric brackets,

$$\partial_\lambda [F_c^{\mu\nu}] = 0. \quad (19)$$

This identity is purely structural — it follows from (14) and the commutativity of partial derivatives. It imposes no new physical input.

In compact differential-forms language, this is simply

$$dF_c = 0,$$

where $F_c = dA_c$ is the exact two-form built from the potential one-form A_c , and $dF_c = ddA_c = 0$ is the structural identity $d^2 = 0$. The homogeneous Maxwell equations are nothing more than $dF_c = 0$ in 3+1 decomposition.

Decomposed into 3+1 form, (19) yields

$$\nabla \cdot \mathbf{B}_c = 0, \quad (20)$$

$$\nabla \times \mathbf{E}_c + \partial\mathbf{B}_c/\partial t = 0. \quad (21)$$

These are the homogeneous Maxwell equations.

8.1 The Magnetic Charge Exclusion

Equation (20) carries substrate content beyond its structural status. Within the closure construction, a circulation-driven source cannot terminate freely: the field-line structure of \mathbf{B}_c admits no open endpoints. We state the resulting prediction carefully:

VERSF prediction. *No fundamental isolated U(1) magnetic charge exists within the admissible closure structure.* The detection of such a charge would refute the gauge-redundancy admissibility constraint (B3) as applied to the electromagnetic transport sector at leading order in ξ , or would force substantial revision of the substrate construction.

The careful framing matters. This prediction is consistent with:

- **Dirac monopoles** at the theoretical level — Dirac's argument shows monopoles are compatible with quantum mechanics under charge quantisation, but does not require them. Their non-observation is consistent with the substrate prediction.
- **Topological defects** in extended gauge sectors — emergent monopole-like objects can arise in non-Abelian or symmetry-broken theories as composite topological structures. These are not "fundamental U(1) magnetic charges" in the sense of the prediction; they are higher-structure objects whose treatment requires extension of the present construction beyond U(1).
- **Effective monopole signatures** in condensed-matter systems — spin-ice and analogous emergent phenomena are not fundamental and do not violate the prediction.

What the prediction *would* be refuted by is the detection of a stable, isolated, fundamental U(1) magnetic-charge particle as an elementary excitation of the electromagnetic field — the analogue of an electric monopole charge for the magnetic sector. To date this remains undetected.

Topological caveat. The argument $F_c = dA_c \Rightarrow dF_c = 0 \Rightarrow \nabla \cdot \mathbf{B}_c = 0$ holds globally only on topologically trivial manifolds, where A_c can be defined as a single one-form across the whole spacetime. On manifolds with nontrivial topology (compact spatial sections, non-trivial principal U(1) bundles), A_c is only locally defined, the global obstruction is a characteristic class, and Dirac-type monopoles are precisely the corresponding topological defects. The substrate prediction here is therefore implicitly conditional on substrate topology: it asserts that the effective substrate manifold supports a globally-defined potential one-form A_c , equivalently that the first Chern class of the relevant U(1) bundle is trivial. This is consistent with the lattice-gauge-theory existence proof cited in open problem 5 — lattice gauge theory on \mathbb{R}^4 or on trivially-topologised lattices has no monopoles, whereas the same theory on a torus admits them. A full statement of the prediction is therefore: *no fundamental isolated U(1) magnetic charge exists, modulo the assumption that the effective substrate manifold is topologically compatible with a globally-defined potential.*

8.2 Faraday's Law as Sectoral Compatibility

Equation (21) asserts that changes in the circulation-driven sector induce structure in the divergence-driven sector. The two sectors cannot evolve independently; their dynamics are

locked by the Bianchi identity. Faraday's law is, in the VERSF reading, a substrate compatibility condition between the two source sectors.

9. The Unique Inhomogeneous Coupling

We now construct the source equation — the equation that determines how the field $F_c^{\mu\nu}$ is generated by the current J_c^μ .

The operative constraints are:

- (B1) Locality.
- (B2) ξ -controlled EFT truncation: at $O(\varepsilon^0)$, the source equation is linear in F_c and first-order in derivatives of F_c (equivalently, second-order in derivatives of A_c).
- (B3) Gauge redundancy: only gauge-invariant operators may appear.
- (B4) Closure-geometry covariance.
- (A1) BCB: taking the divergence must reproduce $\partial_\mu J_c^\mu = 0$.

Enumeration of admissible vector contractions at $O(\varepsilon^0)$. We seek a Lorentz-covariant vector equation linear in F_c and first-order in derivatives of F_c . The candidate first-derivative vectors are:

1. $\partial_\mu F_c^{\mu\nu}$, the divergence of F_c .
2. $\partial_\mu \tilde{F}_c^{\mu\nu}$, the divergence of the Hodge dual $\tilde{F}_c^{\mu\nu} = \frac{1}{2} \varepsilon^{\mu\nu\alpha\beta} F_c^{\alpha\beta}$.
3. $\partial^\nu(\partial_\mu A_c^\mu)$, the gradient of the divergence of the potential.
4. $\varepsilon^{\mu\nu\rho\sigma} k_\rho F_c^\sigma$, $\partial^\lambda \lambda(\dots)$ -type Chern-Simons-like contractions involving a preferred background vector k^μ .

Candidate 2 vanishes identically by the Bianchi identity (19); it cannot serve as a source equation.

Candidate 3 is gauge-variant under (11) — picking up $\partial^\nu \square \chi$ — and is therefore excluded by (B3) acting on the equation.

Candidate 4 — the class of Carroll–Field–Jackiw and related Lorentz-violating modifications of electrodynamics — requires a preferred background four-vector k^μ to construct a covariantly-formed vector contraction. Such a k^μ is excluded by (B4): closure-geometry covariance reduces (within the emergent-Lorentz regime) to Lorentz covariance, and no preferred background vector exists. This makes contact with the Lorentz-violating electrodynamics literature; observational bounds on such terms are very tight and consistent with their exclusion under (B4).

Higher-derivative operators ($\partial^2 F$, $F \cdot \partial F$, etc.) are suppressed by powers of ε under (B2) and do not appear at $O(\varepsilon^0)$.

The surviving admissible source equation at $O(\varepsilon^0)$ is therefore

$$\partial_{\mu} F_{\nu}^{\mu} = \mu_{\nu} J_{\nu}, \quad (22)$$

where μ_{ν} is a substrate constant carrying the dimensions required to balance the equation.

In compact differential-forms language, introducing the codifferential δ (the Hodge-adjoint of d), the source equation reads

$$\delta F_{\nu} = \mu_{\nu} J_{\nu}.$$

We work throughout with sufficiently regular differential forms on Minkowski spacetime — or equivalently, on the emergent-Lorentz manifold of (B4) — so that the codifferential is well-defined and the identity $\delta^2 = 0$ holds. Concretely, this requires forms whose components are smooth (or at least C^2) and whose support is compatible with the Hodge star on the metric of (B4). Under these regularity conditions, $\delta^2 = 0$ is the Hodge dual of $d^2 = 0$ and holds without further qualification.

The full Maxwell theory is therefore the compact pair

$$dF_{\nu} = 0 \text{ (homogeneous, automatic from } F_{\nu} = dA_{\nu})$$

$$\delta F_{\nu} = \mu_{\nu} J_{\nu} \text{ (inhomogeneous, unique source equation)}$$

and the BCB conservation law follows automatically from $\delta^2 = 0$: applying δ to both sides of $\delta F_{\nu} = \mu_{\nu} J_{\nu}$ gives

$$0 = \delta(\delta F_{\nu}) = \mu_{\nu} \delta J_{\nu},$$

hence $\delta J_{\nu} = 0$, which is $\partial_{\mu} J_{\nu}^{\mu} = 0$. The structural identities $d^2 = 0$ and $\delta^2 = 0$ deliver, respectively, the homogeneous equations and the conservation of the source — together carrying half of Maxwell theory as pure geometry, with only the source equation $\delta F_{\nu} = \mu_{\nu} J_{\nu}$ carrying genuine physical content beyond the differential-form structure.

Consistency check (index notation). The same conclusion can be reached in index notation, which is sometimes more illuminating in components. Take the four-divergence of (22):

$$\partial_{\nu} \partial_{\mu} F_{\nu}^{\mu} = \mu_{\nu} \partial_{\nu} J_{\nu}. \quad (23)$$

The left-hand side contracts the symmetric operator $\partial_{\nu} \partial_{\mu}$ with the antisymmetric tensor F_{ν}^{μ} , vanishing identically:

$$0 = \mu_{\nu} \partial_{\nu} J_{\nu}, \quad (24)$$

reproducing BCB. **The unique admissible source equation is automatically consistent with strict commitment conservation; no additional constraint had to be imposed by hand.** This is the index-notation analogue of the $\delta^2 = 0$ argument above; the symmetric-antisymmetric contraction in (23) is what $\delta^2 = 0$ says, in components.

Decomposed into 3+1 form, (22) yields the two inhomogeneous Maxwell equations:

$$\nabla \cdot \mathbf{E}_c = \rho_c / \epsilon_c, \quad (25)$$

$$\nabla \times \mathbf{B}_c - (1/c_c^2) \partial \mathbf{E}_c / \partial t = \mu_c \mathbf{J}_c, \quad (26)$$

with

$$c_c^2 = 1 / (\epsilon_c \mu_c). \quad (27)$$

The constants ϵ_c and μ_c are interpreted respectively as the substrate's longitudinal commitment compressibility and transverse commitment inertia. These interpretive names tie ϵ_c and μ_c to fluid/transport analogies natural to VERSF; they are naming conventions rather than independently testable derivations.

10. Wave Propagation in Vacuum

10.1 The Wave Equations

In a source-free region, $\rho_c = 0$ and $\mathbf{J}_c = 0$. Combining (21), (25), (26), and (20) by taking the curl of (21), using (26), and applying the vector identity $\nabla \times (\nabla \times \mathbf{X}) = \nabla(\nabla \cdot \mathbf{X}) - \nabla^2 \mathbf{X}$, one obtains

$$\nabla^2 \mathbf{E}_c - (1/c'_c{}^2) \partial^2 \mathbf{E}_c / \partial t^2 = 0 \quad (28)$$

and symmetrically

$$\nabla^2 \mathbf{B}_c - (1/c'_c{}^2) \partial^2 \mathbf{B}_c / \partial t^2 = 0, \quad (29)$$

where we have written c'_c for the wave-equation propagation speed to distinguish it from the TPB bound c_c of §4. Internally to §9, c'_c satisfies

$$(c'_c)^2 = 1 / (\epsilon_c \mu_c). \quad (30)$$

Both \mathbf{E}_c and \mathbf{B}_c satisfy the same hyperbolic wave equation with propagation speed c'_c . The two field components are not independent during propagation — Faraday's law (21) and Ampère's law (26) lock them into the standard transverse, mutually orthogonal, in-phase configuration. Electricity and magnetism are not separate phenomena coupled by accident; they are the two admissible source modes of a single conserved substrate current, dynamically interconverting at speed c'_c .

10.2 Saturation: A Separately Postulated Identification

The construction up to this point has produced two propagation speeds:

- $c_c = \xi / (N_b \tau_s)$ from §4: the TPB bound, set by substrate update structure.
- $c'_c = 1 / \sqrt{\epsilon_c \mu_c}$ from §9–§10.1: the wave-equation speed, set by the substrate constants ϵ_c and μ_c that appeared as constants of integration in the source equation.

These are *not* the same quantity. c_c is fixed by the coherence length ξ and the commitment time per bit $N_b \tau_s$; c'_c is fixed by the substrate's transport response constants. The empirical identification of the speed of light with both quantities asserts

$$c'_c = c_c, \text{ i.e. } \epsilon_c \mu_c = (N_b \tau_s / \xi)^2. \text{ (Saturation)}$$

This is the **saturation postulate**: the substrate parameters ϵ_c and μ_c are forced to take values that exactly saturate the TPB bound. The construction does not derive this; ϵ_c and μ_c appear as free substrate constants in §9, and could in principle take any positive values producing $c'_c \leq c_c$ without contradicting any admissibility constraint. Choosing values that saturate is an additional input, not a theorem of (A1)–(A2) and (B1)–(B4).

Two ways the saturation could become a theorem rather than a postulate — but the two are *not on equal footing*, and we should say so:

1. **Constitutive derivation (the only genuine derivation route)**. Show that ϵ_c and μ_c are themselves algebraically expressible in terms of the substrate parameters ξ, N_b, τ_s — perhaps through the substrate's response to imposed commitment imbalance and circulation. If $\epsilon_c \mu_c$ is forced *as an algebraic identity* to equal $(N_b \tau_s / \xi)^2$ by the substrate dynamics, then the saturation $c'_c = c_c$ becomes automatic and the postulate is converted into a theorem. This is the route that would actually close the gap, and any genuine resolution of the saturation problem will pass through it. **§16.7 substantially advances this route**: the lattice-action analysis exhibits the continuum wave speed as $c'_c = a/a_t = \xi / (N_b \tau_s)$ directly from the substrate lattice spacings, making saturation a structural consequence of the substrate parameters appearing in the discrete dynamics rather than an independent postulate. The full closure of Route 1 still requires (i) explicit anisotropic Wilson action with separate spatial and temporal couplings, (ii) coupling-constant renormalisation under coarse-graining, and (iii) a derivation of ϵ_c, μ_c as algebraic functions of (ξ, N_b, τ_s, g^2) — but the framework in which these steps can be carried out is now in place.

Cross-reference to the K = 7 paper. The K = 7 paper (Taylor, AIDA Institute) supplies an independent derivational route to α from the same constraint structure that governs the present framework's substrate dynamics, with $\alpha^{-1}_{\text{bare}} = 128 \cdot 15/14 \approx 137.143$. The two routes are consistent: this paper's μ_c (substrate's transverse commitment inertia) and the K = 7 paper's α (electromagnetic coupling from constraint counting on the hexagonal interface) are two faces of the same underlying integer. A full closure of Route 1 must show that the lattice-action μ_c emerging from coarse-graining of the §16.7 toy on a K = 7 hexagonal interface ($N_{\text{loop}} = 14$) agrees with the K = 7 paper's α derivation. This is the central task of the K = 7 Wilson Limit paper identified as the next priority output. The status of saturation should therefore be read as *partially closed via the §16.7 lattice argument, with the convergence calculation pending the K = 7 Wilson Limit paper*.

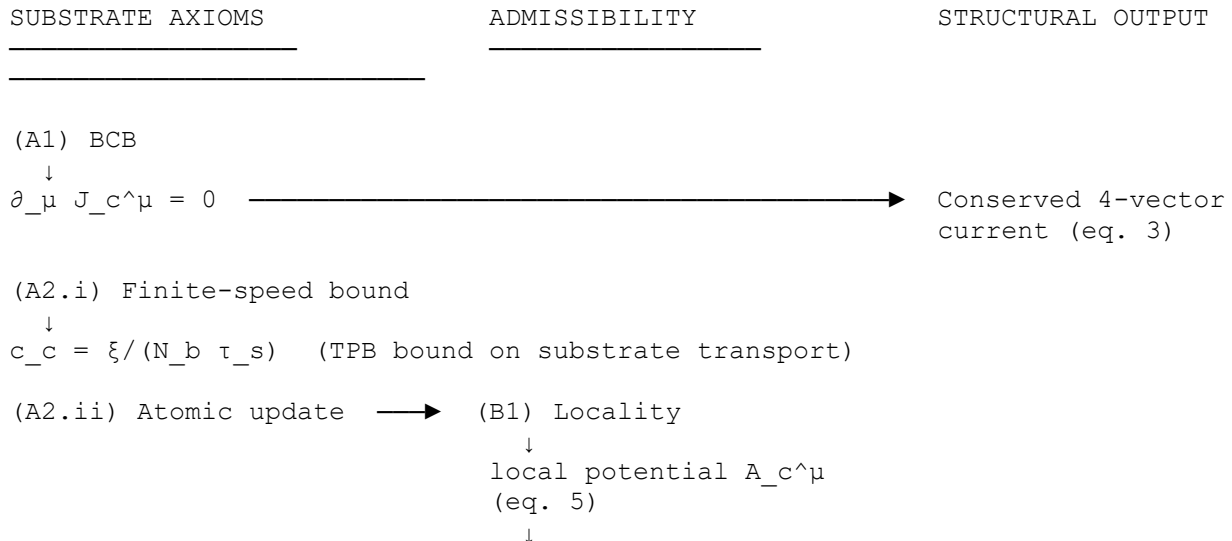
2. **Minimal-obstruction principle (heuristic motivation, not a substitute).** Conjecture that electromagnetism is the minimally obstructed transport mode of the substrate — it couples to the conserved current through the simplest admissible structure ($\delta F_c = \mu_c J_c$) with no internal mass scale, no additional resistance terms, and no substrate constraints beyond the admissibility class. Modes with such minimal obstruction may naturally saturate the propagation bound; modes with additional substrate constraints (mass terms, additional gauge couplings) would fall below it. We flag explicitly: **this is not an alternative derivation route.** Even if the minimal-obstruction principle is validated, it would not constitute a derivation in the absence of an explicit substrate dynamics implementing route 1. The minimal-obstruction principle is best understood as a *heuristic reason to expect route 1 to succeed* — a substrate-physical intuition for why a constitutive derivation should be possible — rather than an independent path to the same conclusion.

Until one of these routes is closed, **the equation $c'_c = c_c$ is a substantive empirical postulate, not a derived result**, and the present construction should be read as: BCB+TPB+(B1)–(B4) deliver Maxwell-form dynamics at *some* wave speed c'_c ; the identification of c'_c with the TPB-saturated value c_c is an additional empirical input. The construction is fully consistent with $c'_c < c_c$ — sub-luminal electromagnetism — but rules out $c'_c > c_c$ (which would violate TPB).

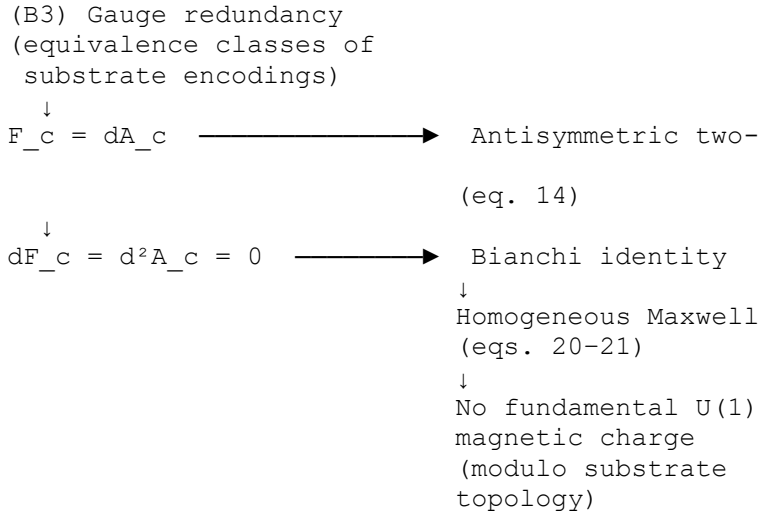
A note on terminology: in the rest of the paper we follow common usage and write a single c_c , understanding that saturation has been assumed.

11. Logical Dependency Map

The structural logic of the derivation, with substrate axioms on the left, admissibility constraints in the middle, and structural outputs on the right:



form



(B2) ξ -EFT @ $O(\epsilon^0)$
+
(B3) gauge invariance
+
(B4) closure covariance
(= Lorentz from emergent
Lorentz paper)

↓
surviving vector
contraction of F_c at $O(\epsilon^0)$

Maxwell

↓
 $\partial_\mu F_c^{\mu\nu} = \mu_c J_c^\nu \longrightarrow$ Inhomogeneous
(eqs. 25-26)

c'_c

↓
Wave eqs at speed

$1/(\epsilon_c \mu_c)$

with $(c'_c)^2 =$

29)

(DERIVED, eqs. 28-

↓

.....
separately postulated:
saturation $c'_c = c_c$
(POSTULATE, see §10.2)
.....

↓

with empirical identification
of J_c^μ with the
electromagnetic current
(CURRENT-SELECTION, see §12.2):

Maxwell

electrodynamics

Reading the map: BCB delivers the conserved vector current; (A2.i) sets the TPB speed bound; (A2.ii) motivates (B1) locality and through it the local potential; gauge redundancy (B3) delivers

the antisymmetric field; the Bianchi identity is structural and yields the homogeneous equations together with the U(1) magnetic-charge exclusion (modulo substrate topology); the ξ -truncation plus gauge invariance plus closure covariance fixes the source equation; the source-free equations yield wave propagation at internal speed c'_c . The dashed boxes show the **two additional inputs** required to reach Maxwell electrodynamics as ordinarily understood — saturation (§10.2) and current-identification (§12.2) — neither of which is a theorem of the present construction.

11.1 The Entire Maxwell Structure in Differential-Forms Language

		$F_c = dA_c$ (definition)				$dF_c = 0$ (geometric identity)				$\delta F_c = \mu_c J_c$
(physical content)					$\delta J_c = 0$ (geometric identity)					

Four lines. **One is a definition** — F_c as the exterior derivative of A_c , expressing the gauge-redundancy quotient structure. **Two are geometric identities** — $dF_c = 0$ from $d^2 = 0$ (homogeneous Maxwell, including the U(1) magnetic-charge exclusion), and $\delta J_c = 0$ from $\delta^2 = 0$ (BCB, automatic from the structure of the inhomogeneous equation). **One is genuine physical content** — $\delta F_c = \mu_c J_c$, the unique closure-compatible source equation at $O(\epsilon^0)$ within the admissibility class.

This is the strongest possible conceptual compression of the result: Maxwell theory is geometry plus one substrate-physical equation. The substrate-physical work of the paper is the derivation of that one equation. Everything else is the geometry of differential forms on the (emergent) Lorentz manifold.

12. The Maxwell Admissibility Theorem

Before stating the main result, we formalise the dependency structure of the construction as a separate proposition. This makes explicit which assumptions are doing which work, and shows that none of (B1)–(B4) is redundant.

Theorem (Dependency Structure). The derivation of Maxwell-form U(1) gauge transport from substrate axioms (A1)–(A2) depends minimally on:

1. **Conserved vector-current structure** (substrate-level, as defined in §12.2): provided by (A1) BCB applied to a vector-current candidate.
2. **Differential locality at the substrate update level:** provided by (B1), motivated by (A2.ii) atomic update.
3. **Gauge redundancy** (informational degeneracy or equivalent): provided by (B3).

4. **Closure-geometry-compatible covariance** (Lorentz-equivalent within the emergent-Lorentz regime): provided by (B4).
5. **Leading-order ξ -EFT truncation**: provided by (B2), with parameter $\varepsilon = \xi/L \ll 1$.

Removal of any one of these assumptions strictly enlarges the admissible transport class beyond Maxwell-form: removing (1) admits no theorem at all; removing (B1) admits causally-admissible integral-kernel theories; removing (B3) admits non-gauge transport structures; removing (B4) admits Chern–Simons-type and other preferred-frame modifications (see Appendix A); removing (B2) admits higher-derivative and nonlinear extensions at the same order as Maxwell, including Born–Infeld and Euler–Heisenberg as $O(\varepsilon^0)$ competitors rather than as suppressed corrections.

The construction is therefore minimal in the sense that no listed assumption can be dropped without admitting a strictly larger class of competing transport theories at $O(\varepsilon^0)$.

This dependency structure is the formal expression of the dependency map in §11. The main theorem follows.

Theorem (Maxwell Admissibility under BCB and TPB). Let the substrate satisfy:

- (A1) BCB: there exists a four-current $J_{\underline{c}}^{\mu}$ with $\partial_{\underline{\mu}} J_{\underline{c}}^{\mu} = 0$.
- (A2) TPB: commitment propagation has finite invariant speed $c_{\underline{c}} = \xi / (N_{\underline{b}} \tau_{\underline{s}})$.

Let the admissibility class be defined by:

- (B1) Locality.
- (B2) Controlled ξ -expansion with parameter $\varepsilon = \xi/L \ll 1$.
- (B3) Gauge redundancy $A_{\underline{c}}^{\mu} \rightarrow A_{\underline{c}}^{\mu} + \partial^{\mu} \chi$.
- (B4) Closure-geometry covariance (reducing to Lorentz within the emergent-Lorentz regime).

Then **within this admissibility class, the surviving admissible linear first-order local transport theory at $O(\varepsilon^0)$ is**

$$\partial_{\underline{\mu}} F_{\underline{c}}^{\mu\nu} = \mu_{\underline{c}} J_{\underline{c}}^{\nu} \text{ (inhomogeneous)}$$

$$\partial_{\underline{[\lambda}} F_{\underline{c]}^{\mu\nu]} = 0 \text{ (homogeneous, Bianchi)}$$

with $F_{\underline{c}}^{\mu\nu} = \partial^{\mu} A_{\underline{c}}^{\nu} - \partial^{\nu} A_{\underline{c}}^{\mu}$, and wave-equation propagation speed $c'_{\underline{c}}$ satisfying $(c'_{\underline{c}})^2 = 1/(\varepsilon_{\underline{c}} \mu_{\underline{c}})$.[†]

This is classical Maxwell-form U(1) gauge transport for the conserved vector current $J_{\underline{c}}^{\mu}$.

† The identification $c'_c = c_c$ with the TPB bound of (A2.i) is a separately postulated saturation claim, not a theorem; see §10.2. The theorem itself delivers Maxwell-form structure with internal wave speed c'_c ; the equality with the TPB bound is an additional empirical input.

A note on terminology. The phrase "classical Maxwell electrodynamics" is sometimes used informally to describe this result, and is acceptable shorthand once two external identifications have been made: (i) J_c^μ is identified with the electromagnetic current (see §12.2, the current-selection problem), and (ii) c'_c is identified with the speed of light by the saturation postulate (see §10.2). Without those identifications, the result is more precisely *Maxwell-form gauge transport for a $U(1)$ -admissible conserved vector current at internal wave speed c'_c* — which is the technically correct claim.

12.1 Scope of the Theorem

The theorem establishes uniqueness only within the admissibility class defined by (B1)–(B4). It does **not** exclude:

- Fundamentally nonlocal transport structures (outside B1).
- Higher-order or nonlinear corrections at $O(\epsilon)$, $O(\epsilon^2)$, etc. (these are accommodated by B2, not excluded; they appear as substrate corrections in the strong-field regime).
- Non-gauge transport structures (outside B3).
- Non-Lorentz-covariant structures relevant at scales where emergent Lorentz invariance breaks down (outside B4).

The substantive VERSF claim is **that the admissibility class itself is physically motivated by substrate informational structure** — not that the class exhausts the logical space of possible transport theories.

12.2 The Current-Selection Problem

A scope limitation of substantial importance, distinct from those catalogued in §12.1, concerns *which* conserved current the theorem applies to. BCB delivers a conserved four-vector current J_c^μ . But the substrate plausibly carries multiple conserved *vector* currents — fermion number, baryon number, lepton number, electric charge itself, hypercharge-like quantities — each independently satisfying $\partial_\mu J^\mu = 0$. The theorem as stated forces a Maxwell-form transport equation on *any* conserved vector current that admits the full admissibility class (B1)–(B4). It does *not* internally identify which vector current is the electromagnetic one.

A clarification before proceeding. The current-selection problem is *specifically* about vector $U(1)$ -type currents that could a priori host Maxwell-form gauge transport. To make the candidate class substrate-native rather than borrowed from standard field theory, we adopt the working definition:

Vector current (substrate-level). A *vector current* is a conserved transport quantity whose substrate update structure is representable by a one-form potential and whose observable

transport flux transforms under the fundamental closure-geometry symmetry as a rank-one object.

Two implications follow. First, conserved quantities of different tensor type — most importantly energy-momentum, which lives on the rank-2 tensor $T^{\mu\nu}$ rather than a four-vector — are not in the candidate set. They do not admit the same one-form potential structure (Maxwell theory's A_c^μ is a one-form; a tensor current would require a different geometric framework, with general relativity rather than Maxwell theory as the natural template). Second, the substrate-native classification matches but does not depend on the standard field-theoretic classification: the rank-one transformation behaviour is read from the substrate's closure-geometry symmetry rather than from imposed Poincaré covariance. Within the emergent-Lorentz regime the two readings coincide; outside that regime, the substrate-native reading is what defines the candidate class.

The current-selection problem is therefore best stated:

Among substrate-level conserved vector currents admitting gauge-potential structure compatible with (B1)–(B4), which one is the electromagnetic current?

This is empirically critical, since baryon number and lepton number are conserved (at low energy) but do not have associated Maxwell-form gauge bosons; only electric charge does. Three resolutions of the current-selection problem are possible:

1. **Substrate-property selection.** Argue that only one conserved vector current in the substrate satisfies the full admissibility class — others fail (B3) because they do not admit a gauge equivalence at the substrate level, or fail (B2) at $O(\epsilon^0)$ for substrate-specific reasons (e.g., they carry an associated substrate mass scale and so propagate sub-luminally with explicit symmetry-breaking). This would be the strongest answer; the present paper does not establish which substrate property distinguishes electromagnetism from baryon-number-like currents.
2. **Primary-current reduction.** Argue that the substrate carries a single primary conserved vector current and that empirically-known additional conservation laws (baryon number, lepton number, etc.) are derived or composite. This would be a major substrate-physical claim requiring independent argument.
3. **Template interpretation.** Concede that the construction gives a *template* for U(1)-style gauge transport that the substrate instantiates whenever a conserved vector current admitting the full admissibility class exists. The empirical labelling "electromagnetism" is then external: the construction proves that *if* the substrate carries an electromagnetically-coupling conserved current admitting (B1)–(B4), Maxwell theory governs it. The substrate-physical question of which current that is — and whether other vector currents (baryon number, etc.) fail the admissibility class through some substrate-level mechanism — is deferred to companion work.

The honest position of the present paper is option 3. The theorem proves that Maxwell-form gauge transport is forced for any conserved vector current satisfying (A1)–(A2) and (B1)–(B4); it does not single out the electromagnetic current internally. Calling the result "Maxwell

electrodynamics" rather than "Maxwell-form gauge transport for a U(1)-admissible conserved vector current" reflects the empirical identification, not a substrate-level selection theorem.

Connection to the non-Abelian programme (OP 2). The current-selection question is most naturally resolved through the non-Abelian extension flagged in open problem 2. The Standard Model gauge structure $SU(3)\times SU(2)\times U(1)$ is precisely the empirical specification of which substrate currents couple to which gauge sector — colour to the strong force, weak isospin to the weak force, hypercharge to the electroweak U(1), with the electromagnetic current arising as a particular combination after electroweak symmetry breaking. A substrate-level derivation of the Standard Model gauge group from the hexagonal closure geometry of VERSF (referenced separately in the corpus) would, *if successful*, also resolve which substrate-level conserved current is electromagnetic: it would be the linear combination of substrate-admissible vector currents that the gauge-group derivation singles out. The current-selection problem is therefore best read not as a flaw of the present U(1) construction but as the empirical specification problem that the non-Abelian programme is supposed to solve. This is a real limitation, and a route to closing it has been identified, but the route is not closed here.

12.3 A Candidate Current-Selection Criterion

The previous subsection concedes that current selection is open at the U(1) level and points to the non-Abelian programme. We now propose a **candidate criterion** that converts the open problem from a vague gap into a formal minimisation programme. The criterion is conjectural — it is not derived here — but it has the structural advantage of being precise enough to be tested against, refined, or refuted by future work.

The criterion. The electromagnetic current is the unique conserved vector current J_a^μ in the substrate satisfying all four of the following conditions:

1. **Massless U(1) coupling.** The current couples to a gauge boson with vanishing rest mass — $m_a^2 = 0$ at the IR fixed point. Currents whose gauge bosons acquire mass (via Higgs-type mechanisms or other substrate-level mass generation) fail this condition.
2. **Survival of $K = 7$ closure.** The loop curvature associated with J_a survives the $K = 7$ hexagonal closure structure without confinement — i.e., the holonomy operators of the corresponding gauge sector are not subject to dynamical confinement at the substrate scale. Currents whose gauge bosons are confined into composite singlets (analogue: QCD gluons) fail this condition.
3. **TPB saturation.** The propagation speed of the J_a -coupled gauge boson saturates the TPB bound c_c (§10.2). Sub-saturating modes — those with $c' < c_c$ due to internal substrate constraints — fail this condition.
4. **Non-Abelian closure survival.** After the substrate's non-Abelian gauge structure reduces (via symmetry breaking, confinement, or other reduction mechanisms) to its low-energy massless content, J_a remains unbroken and uncoupled to a residual non-Abelian structure that would obstruct its admissibility class.

Combining the four conditions into a single minimisation programme:

$$J_{EM}^\mu = \operatorname{argmin}_{\{J_a^\mu\}} [m_a^2 + \Delta_{\text{conf}}(a) + \Delta_{\text{break}}(a) + \Delta_{\text{sat}}(a)],$$

where Δ_{conf} , Δ_{break} , Δ_{sat} are non-negative *obstruction functionals* measuring the failure of conditions 2, 4, and 3 respectively, with the global minimum at zero corresponding to a current satisfying all four conditions simultaneously. The electromagnetic current is conjectured to be the unique sector at which all four obstructions vanish.

Connection to the §10.2 minimal-obstruction principle. This criterion is the formal version of the minimal-obstruction principle invoked heuristically in §10.2 Route 2 to motivate saturation. There it was demoted to "heuristic motivation for Route 1" because it did not constitute a derivation of saturation by itself. Here, elevated to a current-selection criterion, the minimal-obstruction principle becomes a substantive substrate-level claim: *electromagnetism is the substrate's most-favoured transport sector by exactly four independent admissibility measures*. The saturation argument and the current-selection argument become two faces of the same underlying programme — the substrate selects the sector that minimises obstruction across all admissibility dimensions, and the resulting sector is empirically electromagnetism.

A circularity that requires explicit reconciliation. A reader will reasonably ask: if §10.2 / §16.7 / Appendix B derive saturation as a structural RG-fixed-point property of the substrate dynamics, then $\Delta_{\text{sat}} = 0$ identically for any massless admissible current — and the criterion reduces to three independent conditions, not four. Conversely, if Δ_{sat} can be non-zero current-by-current, then the RG argument can't deliver universal saturation — it would only deliver saturation for currents whose substrate-physical properties make the fixed point attractive for them specifically. The two readings are in tension; this paragraph reconciles them.

The resolution is that the RG argument of §10.2 / Appendix B delivers saturation **within the abelian admissibility-class sector at substrate scale** — that is, for those conserved vector currents that admit (B1)–(B4) at the substrate level and that project onto the abelian sector of the bare substrate dynamics. The bare anisotropy $\gamma_{\text{bare}} = \xi/(N_b \tau_s)$ and the Lorentz-compatible IR fixed point $\gamma^* = c_c$ are substrate-wide statements about the abelian sector's transport. They do *not* automatically apply to (i) currents confined within non-abelian sectors (whose effective wave speeds are set by hadronic-scale physics, not substrate-scale dynamics), (ii) currents whose gauge bosons acquire mass via symmetry breaking (whose dispersion deviates from c_c at low momentum), or (iii) currents whose substrate-level couplings sit outside the (B1)–(B4) class entirely. Δ_{sat} is therefore the **projection test**: it measures whether a given conserved vector current sits in the substrate-scale abelian admissibility sector to which the RG saturation result applies, rather than in some other sector with a different effective wave speed.

Under this reading, the four conditions are genuinely independent: massless ($m_a^2 = 0$) tests the gauge-boson rest mass; unconfined ($\Delta_{\text{conf}} = 0$) tests dynamical confinement at substrate scale; unbroken ($\Delta_{\text{break}} = 0$) tests survival of symmetry breaking; saturating ($\Delta_{\text{sat}} = 0$) tests projection onto the abelian admissibility sector where the RG fixed point $\gamma^* = c_c$ is attractive. The RG result of Appendix B *guarantees* that $\Delta_{\text{sat}} = 0$ holds for currents satisfying all three of the other conditions, but Δ_{sat} is still an independent admissibility test — it can fail for currents that fail one of the other three through a different mechanism than the other three identify. The interconnections among the four obstructions are real and a complete treatment must derive

them; what the criterion does is identify four independent admissibility tests whose simultaneous vanishing is conjectured to pick out the electromagnetic current uniquely.

This reconciliation also clarifies what §10.2 *does and does not* claim. §10.2 claims saturation as a substrate-scale RG-fixed-point property of the abelian admissibility sector — *not* a universal claim that any conserved current saturates regardless of its substrate-level properties. The empirical observation $c_{\text{photon}} = c_{\text{TPB}}$ is consistent with the photon being the gauge boson of the abelian admissibility sector. Other gauge bosons (Z, W, gluons) deviate from this saturation in specific ways predicted by their respective sector dynamics — the W and Z by Higgs-acquired mass, the gluons by confinement — exactly as the four-condition criterion describes.

Status of the criterion. This is a *candidate programme*, not a theorem. Closing it requires:

- *Explicit construction of the obstruction functionals.* m_a^2 is straightforward from the gauge-boson mass spectrum, but Δ_{conf} , Δ_{break} , Δ_{sat} require precise definitions at the substrate level — confinement at the $K = 7$ closure level, symmetry-breaking obstructions inherited from the non-Abelian reduction, and sub-saturation penalties tied to substrate-scale mass generation.
- *Demonstration that the four obstructions are independent.* If two of them collapse to the same quantity (e.g., if confinement and saturation failure are equivalent for confined gauge bosons), the criterion would simplify.
- *Verification that the EM current uniquely minimises.* This requires a complete enumeration of substrate-admissible vector currents in the non-Abelian programme and verification that no other current achieves zero on all four obstructions simultaneously.
- *Connection to the $K = 7$ gauge group derivation.* The criterion presumes a non-Abelian structure delivering the candidate currents; the $K = 7$ paper currently treats $SU(3) \times SU(2) \times U(1)$ as empirical input. A substrate-level derivation of the gauge group would make the candidate enumeration internal rather than external.

What §12.3 accomplishes is converting "current selection is open" into "current selection is a four-condition minimisation programme on the obstruction functionals (m^2 , Δ_{conf} , Δ_{break} , Δ_{sat}); the $K = 7$ hexagonal-interface paper supplies the substrate scaffolding, the non-Abelian extension supplies the candidate enumeration, and the convergence of the programme — if successful — would prove that the electromagnetic current is the unique global minimum." This is a programme with shape, not merely a confessed gap.

12.4 Proof Sketch

1. (B3), in its mathematical content as gauge invariance, forces the physical field into antisymmetric form $F_c^{\mu\nu}$ (§7, Lemma).
2. The homogeneous equations follow as a structural identity from $F_c = dA_c$ (§8).
3. (B1)–(B4) plus (A1) select the inhomogeneous equation at $O(\epsilon^0)$ as the surviving admissible coupling (§9). (B2) restricts to linear first-order operators; (B3) excludes gauge-variant candidates; (B4) restricts to Lorentz-covariant vector contractions and excludes Chern-Simons-like terms requiring a preferred background vector; the Hodge-

dual candidate vanishes by Bianchi; the gradient-of-divergence candidate is gauge-variant; $\partial_\mu F_c^{\mu\nu} = \mu_c J_c^\nu$ is what remains.

4. The wave speed c'_c is read off the source-free equations (§10); the *separately postulated* identification $c'_c = c_c$ with the TPB bound is the saturation claim of §10.2. ■

13. Relationship to Standard Gauge-Theoretic Uniqueness

The Maxwell uniqueness result is well-established in standard physics (Wald 1984, Ch. 4; Weinberg 1995, Vol. 1, Ch. 5 and 8; Jackson 1998 for a textbook treatment; Deser 1970 for the gauge-self-coupling perspective). The conventional derivation takes as input:

Standard input	VERSF input	Substrate basis
Lorentz invariance	(B4) Closure-geometry covariance	Emergent Lorentz from VERSF substrate (separate paper)
Locality	(B1) Locality	Forced by TPB: nonlocal updates violate finite local progression
Gauge symmetry	(B3) Gauge redundancy	Informational degeneracy: equivalence classes of substrate encodings
Renormalisability / lowest derivatives	(B2) ξ -controlled EFT truncation	Power counting in coherence scale ξ/L
Conserved source	(A1) BCB	Substrate axiom (committed bits locally conserved)
—	(A2) TPB	Substrate axiom (finite commitment propagation)

The VERSF contribution is not the uniqueness theorem itself — that is standard. The contribution is **the substrate-level reading of the admissibility constraints**: locality from TPB rather than from convenience; gauge symmetry from informational degeneracy rather than from imposed redundancy; derivative-order truncation from coherence-scale power counting rather than from renormalisability arguments tied to perturbative quantum field theory; conserved sources from substrate bit conservation rather than from an external symmetry.

In standard practice these admissibility constraints are usually justified by appeal to experimental success or by post-hoc arguments from quantum-field-theoretic consistency. VERSF supplies a substrate-level rationale that is *prior to* the field theory itself.

14. What Would Falsify This Construction

A scientifically responsible framework must specify the observations that would refute it. The following would force substantial revision of the present construction:

- **Observed superluminal electromagnetic transport.** This would refute either (A2.i), the TPB finite-speed bound, or **the saturation postulate of §10.2** (that the wave-equation speed c'_c equals the TPB bound c_c). With the saturation postulate now isolated, falsification splits into two cases: superluminal EM with $c'_c > c_c$ would refute (A2.i); a sub-luminal $c'_c < c_c$ that contradicts measured photon speed would refute the saturation postulate without touching (A2.i).
- **Detection of a fundamental, isolated, stable U(1) magnetic-charge particle as an elementary excitation.** This would refute the gauge-redundancy admissibility constraint (B3) applied to the electromagnetic sector at $O(\epsilon^0)$, or force substantial revision of the substrate construction (§8.1). Note that this is distinct from Dirac monopoles as a theoretical possibility, from emergent topological monopoles in extended gauge sectors, and from effective monopoles in condensed-matter systems — none of which would refute the prediction.
- **Experimentally necessary gauge-variant observables in electromagnetic phenomena.** If a measurable quantity were found to depend on the gauge of A_c^μ rather than on $F_c^{\mu\nu}$, this would refute (B3) as informational degeneracy. The Aharonov–Bohm effect is sometimes cited in this connection but does *not* refute the prediction: the AB phase depends on $\oint A_c \cdot dl$, which is gauge-invariant (it is a holonomy of A_c), not on A_c itself at any point.
- **Experimentally required macroscopic nonlocal corrections to Maxwell's equations.** Beyond known quantum-EFT corrections at very small scales, the discovery of nonlocal terms required at macroscopic L (where $\epsilon = \xi/L$ is small) would refute (B1) or the EFT-controlled framing in (B2).
- **Failure of the EFT scaling.** If nonlinear corrections to Maxwell theory appeared at strengths far below what $\epsilon \ll 1$ predicts — or refused to appear at all where the construction expects them — this would refute the (B2) framing.
- **Breakdown of Lorentz covariance at energies well below where the emergent-Lorentz construction predicts it.** This would refute (B4) as inherited from the companion paper, and would propagate back into the present construction as a failure of admissibility.

This list is not exhaustive, but it specifies the principal failure modes. The construction is structured to make each of these tests, in principle, decisive.

14.1 Near-Term Empirical Handles

The catalogue above lists what would refute the construction in principle, but several entries (e.g., macroscopic-scale nonlocal Maxwell corrections, EFT-scaling failure) are difficult to test directly in near-term experiments. We collect here three specific predictions with cleaner near-term empirical paths, prioritising precision propagation tests and cosmological-correlation observations over Born–Infeld nonlinear-threshold searches (which, with ξ anchored as in OP 6, are empirically safe but not near-term falsifiers).

Prediction A — No energy-dependent photon speed below the substrate-UV threshold.

The substrate ξ that enters the (B2) ξ -EFT expansion is the *UV* coherence length over which a single bit is locally resolved (§4), distinct from the $K = 7$ paper's *IR* cosmological coherence scale (the de Sitter radius scale $\sim 10^{61}$ Planck units after the coherence-scale proportionality constant derived in a companion paper cited by the $K = 7$ paper). The substrate-UV ξ is plausibly near-Planckian; current laboratory and astrophysical observations are consistent with this. The associated substrate energy scale is $\Lambda_\xi \sim \xi_{\text{substrate}}^{-1}$.

The construction predicts that photon-dispersion deviations from Lorentz invariance, parameterised as

$$\Delta c/c \lesssim (E / \Lambda_\xi)^n$$

(with $n = 2$ from leading-order ξ -EFT corrections to Maxwell at $O(\epsilon^2)$; n could be larger if $O(\epsilon^2)$ corrections vanish by symmetry), should be suppressed for any $E \ll \Lambda_\xi$. Current high-energy gamma-ray-burst dispersion measurements, atomic-clock comparisons, and precision laser interferometry constrain energy-dependent photon speed at energies well below the Planck scale; all are consistent with the prediction. **Future improvements in cosmic-ray dispersion constraints, particularly from sub-PeV neutrino and gamma-ray timing, would tighten the bound on Λ_ξ from below**, providing a near-term observational handle on the substrate-UV scale.

Refutation: any robust detection of energy-dependent photon speed $\Delta c/c$ above the $(E/\Lambda_\xi)^n$ bound, with Λ_ξ at any scale up to Planck, would constrain or refute the construction.

Prediction B — No fundamental U(1) magnetic charge.

This is the substrate exclusion of §8.1, with the topological caveat noted there (the prediction is modulo substrate topology trivial enough to support a globally-defined potential one-form). Restated here for the empirical-handles catalogue: the detection of a stable, isolated, fundamental U(1) magnetic-charge particle as an elementary excitation of the electromagnetic field would refute the construction. Dirac-monopole arguments at the theoretical level, emergent topological monopoles in extended gauge sectors, and condensed-matter analogues do not refute the prediction (see §8.1 for the distinction). The continuing non-detection of such a fundamental particle remains supportive but not confirmatory.

Prediction C — α and Λ are jointly fixed by integer $K = 7$; both are predicted constant.

The $K = 7$ paper's constraint C4 (Chern–Weil winding number on the de Sitter horizon) establishes that both the fine-structure constant α and the cosmological constant Λ are observables of the same integer K . Specifically, in the $K = 7$ paper's notation, $\alpha(K) = N_{\text{loop}} / ((N_{\text{loop}} + 1) \cdot 2^K)$ and $\Lambda(K) \propto \xi(K)^{-2} = \exp(-\pi \cdot 2^K)$, where $N_{\text{loop}} = 14$ and the absolute scales depend on the coherence-scale proportionality constant derived in a separate companion paper cited by the $K = 7$ paper.

Because K is integer-valued, neither α nor Λ varies continuously within the framework. K cannot drift continuously; it can only "jump" between integer values, and the $K = 7$ paper's

Section 6.2 establishes that adjacent integer values place α and Λ ~ 53 orders of magnitude away from the observed band. The prediction is therefore stronger than a correlated-drift statement:

$\dot{\alpha}/\alpha = 0$ and $\dot{\Lambda}/\Lambda = 0$ (joint zero-drift),

within the framework's regime of validity. Any detected cosmological drift in either α or Λ — at any non-zero rate, in either direction — refutes the framework, because the integer K cannot supply continuous variation.

The framing " $\dot{\alpha}/\alpha \sim \mathcal{F}(\dot{\Lambda}/\Lambda)$ " in earlier drafts was technically incorrect: there is no continuous parameter linking α to Λ within the $K = 7$ framework, only the common integer $K = 7$, so the prediction reduces to *both constant*. A correlated-drift relation could only arise if some continuous parameter besides K (e.g., the coherence-scale proportionality constant, treated as dynamical) varied — but the $K = 7$ paper treats those as fundamental constants, not dynamical fields. We therefore state the prediction in its sharpest correct form: **$K = 7$ is a fixed integer; therefore α and Λ are both constants; therefore any non-zero cosmological drift in either refutes the framework.**

This is the empirical handle most likely to produce decisive results in near-term observational cosmology. The α -stability bound from quasar absorption ($|\dot{\alpha}/\alpha| \lesssim 10^{-17} \text{ yr}^{-1}$) and the dark-energy constancy bound ($w + 1 = O(10^{-2})$ at present precision, with implied $|\dot{\Lambda}/\Lambda|$ bounds) are routinely treated as independent constraints in standard cosmology. Within the $K = 7$ framework, they are aspects of the same constraint and any future precision improvement in either tests the framework. A robust detection of α -drift or $w \neq -1$ evolution would refute the joint $K = 7$ structure and hence the framework's constraint-dimensionality reading.

Summary of empirical handles. Prediction A bounds the substrate-UV scale from below via precision propagation tests; Prediction B excludes fundamental magnetic monopoles; Prediction C predicts joint constancy of α and Λ as a consequence of integer- $K = 7$ being fixed, testable by combining quasar absorption observations (α -drift bound) with dark-energy precision cosmology (w -evolution bound). The framework's empirical posture is therefore primarily *propagation-and-correlation*, not *nonlinear-strong-field*: the Born–Infeld threshold is empirically safe but unreachable; the testable predictions live in precision tests of photon dispersion and joint α - Λ stability, both of which are active observational programmes.

15. Epistemic Status

We catalogue the status of each claim. For ease of reference by reviewers, the dependency structure is summarised in the table below; the narrative subsections that follow give detail.

Summary Table

Claim	Status	Conditional on
$F_c^{\mu\nu}$ is the surviving gauge-invariant antisymmetric piece (Lemma §7.2)	Mathematically proven	Restricted to tensors linear in one factor of ∂A_c
Bianchi identity $\partial_\lambda F_c^{\mu\nu} + \partial_\mu F_c^{\nu\lambda} + \partial_\nu F_c^{\lambda\mu} = 0$	Mathematically proven	$F_c = dA_c$ with smooth fields
$\partial_\mu \tilde{F}_c^{\mu\nu} = 0$	Mathematically proven	Bianchi identity
$\partial_\nu \partial_\mu F_c^{\mu\nu} = 0$ (BCB automatic)	Mathematically proven	Symmetric–antisymmetric contraction
Homogeneous Maxwell equations	Substrate-physically proven	$F_c = dA_c$ (which requires B1, B3)
No fundamental U(1) magnetic charge	Substrate-physically proven	$F_c = dA_c$ and topologically-trivial effective substrate manifold
Inhomogeneous Maxwell as surviving source equation at $O(\epsilon^0)$	Substrate-physically proven	(B1)–(B4), (A1)
Maxwell-form gauge transport theorem (§12)	Substrate-physically proven	(A1), (A2), (B1)–(B4)
Wave equations at internal speed c'_c	Substrate-physically proven	Source equation + Bianchi + source-free
Saturation: $c'_c = c_c$	Separately postulated	Not derived; see §10.2 for the gap
J_c^μ identified with electromagnetic current	External identification	Not derived; see §12.2 (current-selection problem)
(B1) Locality	Conditional on admissibility motivation	Motivated by (A2.ii) atomic update; rigorous derivation open (OP 4)
(B2) ξ -EFT truncation	Conditional on admissibility motivation	Hierarchy asserted, not derived from substrate UV integration (OP 5)
(B3) Gauge redundancy	Conditional on admissibility motivation	Quotient structure is standard; substrate interpretation pending OP 5
(B4) Closure-geometry covariance	Conditional on companion VERSF paper	Emergent Lorentz invariance (separate corpus paper)
Electric \leftrightarrow divergence-driven, magnetic \leftrightarrow circulation-driven (source sectors)	Interpretive	Empirical matching to Maxwell
ϵ_c, μ_c as commitment compressibility/inertia	Interpretive	Naming convention, not a derivation

Claim	Status	Conditional on
Substrate-physical interpretation of \mathcal{S} / \sim	Interpretive	Mathematical structure is standard; novelty awaits explicit \mathcal{S}
"Commitment" operational definition	Partially addressed by companion measurement/fact-formation papers (§3.1); not required for the Maxwell theorem itself	Layer A: substrate model (OP 5); Layer B: integration of the three measurement companion papers introduced in §3.1

Mathematically proven

- $F_c^{\mu\nu}$ is the surviving gauge-invariant antisymmetric tensor built linearly from one factor of ∂A_c (Lemma §7.2).
- The Bianchi identity $\partial_{[\lambda} F_c^{\mu\nu]} = 0$ follows from $F_c = dA_c$ by commutativity of partial derivatives.
- $\partial_\mu \tilde{F}_c^{\mu\nu}$ vanishes as a consequence of the Bianchi identity.
- $\partial_\nu \partial_\mu F_c^{\mu\nu} = 0$ by symmetric–antisymmetric contraction; BCB follows automatically from the source equation as $\delta^2 = 0$.

Substrate-physically proven (within (A1)–(A2) and (B1)–(B4))

- The homogeneous Maxwell equations express substrate consistency conditions: circulation-driven sources cannot terminate freely (no fundamental U(1) magnetic charge, modulo substrate topology), and sectoral compatibility (Faraday's law) holds between the two source sectors.
- The inhomogeneous source equation is the surviving linear first-order admissible coupling consistent with BCB at $O(\epsilon^0)$, automatically respecting commitment conservation.
- Maxwell-form gauge transport for J_c^μ is fixed within the admissibility class at $O(\epsilon^0)$.

Separately postulated (not theorems of the present construction)

- **Saturation:** $c'_c = c_c$. The internal wave speed of §10.1 equals the TPB bound of §4. This is an empirical postulate; the construction does not derive ϵ_c, μ_c from substrate parameters, so the saturation is added rather than forced. See §10.2 for the gap and candidate route to closure.
- **Electromagnetic identification.** The conserved current J_c^μ whose dynamics is forced into Maxwell form is identified with the empirical electromagnetic current by external matching. See §12.2 for the current-selection problem.

Conditional on admissibility constraints (motivated, not strictly derived)

- **(B1) Locality.** Tightly tied to (A2.ii) atomic update (§5) but the rigorous derivation of locality as a theorem of substrate axioms remains open (OP 4).

- **(B2) EFT truncation in ξ .** Controlled by power counting in $\varepsilon = \xi/L$. The truncation is leading-order admissibility; the standard Wilsonian derivation of operator suppression by substrate-UV integration has not been exhibited (OP 5).
- **(B3) Gauge redundancy.** Motivated by informational degeneracy (§7). The mathematical structure (quotient by equivalence relation) is standard; the substrate interpretation requires explicit \mathcal{S} , deferred to OP 5.

Conditional on companion VERSF results

- **(B4) Lorentz covariance.** Inherited from the emergent-Lorentz-invariance paper. Read alone, this paper imposes Lorentz covariance by hand; read with the companion paper, it is a theorem.

Interpretive

- Empirical identification of the divergence-driven source sector with electricity and the circulation-driven source sector with magnetism. The substrate-native vocabulary is sectoral; the empirical labels are forced once field equations are in hand.
- ε_c as longitudinal commitment compressibility and μ_c as transverse commitment inertia. A naming convention, not an independent test.

Falsifiable

- See §14 for the catalogue of observations that would refute the construction, now including the separately-isolated saturation postulate and topological prerequisite for magnetic-charge exclusion.

16. A Minimal Substrate Toy — Preliminary Constructive Sketch

16.1 Scope of This Section

The preceding sections argued that *if* a substrate satisfies (A1)–(A2) and admits the admissibility class (B1)–(B4), Maxwell-form U(1) gauge transport emerges as the surviving structure at $O(\varepsilon^0)$. We now sketch one class of substrates that would satisfy these constraints constructively, with the aim of beginning to populate the abstract admissibility class with concrete dynamics.

The framing of this section is deliberately modest. **We do not claim to exhibit *the* VERSF substrate, nor a complete substrate model.** We claim only the following: (i) at least one explicit class of discrete substrates satisfying (A2.ii), (B1), and (B3) constructively can be written down; (ii) the gauge-redundancy quotient structure of §7.1, which was mathematically standard but interpretively asserted, can be realised explicitly in such a toy; (iii) candidate routes toward constitutive derivations of the saturation postulate and operational definitions of

commitment become visible once a concrete toy is on the table, even if those routes are not closed here.

The toy is not unique, is not claimed to be fundamental, and inherits structure from standard lattice gauge theory while extending its substrate-physical interpretation. Its role is to convert "the substrate motivates the constraints" from descriptive language into one verified instance of joint realisation.

16.2 The Toy Substrate: Simplicial Transport Network

Let the substrate consist of a graph

$$\mathcal{G} = (V, E)$$

where V is the set of substrate sites and E the set of nearest-neighbour links. Each site $i \in V$ carries a local distinguishability state

$$s_i \in \mathbb{Z},$$

interpreted as local commitment imbalance (negative deficit, zero, positive excess). The simplest non-trivial truncation $s_i \in \{-1, 0, +1\}$ suffices for many qualitative arguments; the full integer-valued case is used in §16.7's dynamical analysis. Each link $(i, j) \in E$ carries a transport phase variable

$$U_{\{ij\}} = \exp(i \theta_{\{ij\}}), \theta_{\{ij\}} \in [0, 2\pi),$$

$$\text{with } \theta_{\{ji\}} = -\theta_{\{ij\}}.$$

Substrate-physical interpretation. Unlike conventional lattice gauge theory, where the $U_{\{ij\}}$ are abstract group-valued connections introduced for computational purposes, the present interpretation reads each $U_{\{ij\}}$ as encoding the substrate-physical transport relation between neighbouring distinguishability states — the "channel content" along the link. The site value s_i is what the substrate has resolved locally; the link value $U_{\{ij\}}$ is the substrate's representation of how a resolved distinction at i relates, transport-wise, to a resolved distinction at j .

Atomic tick update. Substrate evolution proceeds by an atomic tick rule. Let $\lambda \in \mathbb{N}$ index ticks. At tick λ :

$$(s_i, U_{\{ij\}}) \rightarrow (s_{i'}, U_{\{ij\}}')$$

according to a local update function \mathcal{U}_λ that depends only on the immediate neighbourhood $\mathcal{N}(i)$ of each site — its directly-linked neighbours. No site accesses distant global information; no update consults integrals over past history. **This is the constructive realisation of (A2.ii) atomic update.**

The update rule is left schematic at this level of abstraction; a complete specification would require choosing dynamics that satisfy the remaining constraints (conservation, gauge invariance under §16.4 below). What matters here is that *some* update rule of this local form exists and that the constraints can be jointly satisfied within it.

16.3 Conserved Transport and Finite Propagation

Discrete BCB. Define the local distinguishability density $\rho_i \equiv s_i$ and the directed link current

$$J_{\{ij\}} = f(U_{\{ij\}}) \cdot (s_i - s_j), J_{\{ji\}} = -J_{\{ij\}}$$

for some real function f of the link phase (a concrete choice is $f(U) = (1 + \cos \theta)/2$, but the precise form does not matter for the conservation argument). If the update rule \mathcal{U}_λ is constructed to preserve the global total

$$\sum_i s_i = \text{const},$$

then the local discrete continuity equation

$$\Delta \rho_i + \sum_{\{j \in \mathcal{N}(i)\}} J_{\{ij\}} = 0$$

holds at every site, and under coarse-graining at $L \gg \xi$ becomes

$$\partial_t \rho + \nabla \cdot \mathbf{J} = 0,$$

which is the continuum BCB equation (1). **BCB is therefore realised constructively as the hydrodynamic limit of local tick dynamics, given a choice of update rule preserving $\sum s_i$.** Such rules exist (e.g., conservative diffusion rules over the link phases); we are not claiming uniqueness, only existence.

Finite propagation, and how the TPB bound emerges. Each tick takes proto-time τ_s . A natural first reading would be: each tick updates nearest neighbours, so distinguishability propagates one coherence length ξ per τ_s , giving signal speed ξ/τ_s . **This is not yet the TPB bound $c_c = \xi/(N_b \tau_s)$** ; the additional factor of $1/N_b$ reflects the bit-commitment structure that distinguishes raw substrate updates from resolved committed transport.

The TPB-consistent reading is: a substrate tick is the elementary update step, but a *committed bit* — a resolved distinction capable of being registered at a neighbouring site — requires N_b ticks of substrate processing before it propagates. During those N_b intermediate ticks, the substrate is locally resolving the distinction at the source site; the bit is in transit but not yet "available for delivery." Only after the full commitment cycle of $N_b \tau_s$ does a committed bit propagate one coherence length ξ to a neighbour. The signal speed for resolved distinguishability is therefore

$$c_c = \xi / (N_b \tau_s),$$

which is the TPB bound (4). **The factor of N_b is the constructive content of TPB:** it is the number of substrate updates required to resolve one bit, and it suppresses the signal speed below the raw tick rate ξ/τ_s .

This makes (A2.i) constructively visible: c_c is set by the substrate update structure and the commitment-cycle length, in agreement with the abstract definition of §4. The two contents of TPB — (A2.i) finite-speed bound and (A2.ii) atomic update — both emerge from the same substrate dynamics, but they constrain different quantities (signal speed vs. update locality) and remain logically independent as flagged in §2.

16.4 Gauge Redundancy from Substrate Degeneracy

This is the strongest constructive contribution of the toy: gauge redundancy, which was interpretively asserted in §7.1, can be realised concretely as substrate microstate degeneracy under an explicit observation map.

Consider a local relabelling of the link phases by an arbitrary site-valued function $\chi : V \rightarrow \mathbb{R}$:

$$\theta_{\{ij\}} \rightarrow \theta_{\{ij\}} + \chi_i - \chi_j. (*)$$

(This is the lattice analogue of $A_c^\mu \rightarrow A_c^\mu + \partial^\mu \chi$ from §7.) Under (*), the directed loop sum around any closed plaquette \square in \mathcal{G} transforms as

$$\Sigma_{\{(i,j) \in \partial \square\}} \theta_{\{ij\}} \rightarrow \Sigma_{\{(i,j) \in \partial \square\}} (\theta_{\{ij\}} + \chi_i - \chi_j) = \Sigma_{\{(i,j) \in \partial \square\}} \theta_{\{ij\}},$$

because each χ_i appears once with positive sign and once with negative sign as the loop traverses the site twice (once entering, once leaving). The closed-loop holonomy is therefore invariant under (*).

The observable map and quotient structure. Define the observable map \mathcal{O} from substrate microstates $\{(s_i, U_{\{ij\}})\}$ to observable transport configurations by taking \mathcal{O} to return: (i) the site distinguishability values $\{s_i\}$, (ii) all directed-loop holonomies $\Sigma_{\{\partial \square\}} \theta_{\{ij\}}$, and (iii) all measurable transport fluxes between sites. Two substrate microstates differing only by a transformation of the form (*) produce identical values of \mathcal{O} — they have identical s_i , identical holonomies, and identical fluxes. They are therefore equivalent under the relation $s_1 \sim s_2 \Leftrightarrow \mathcal{O}(s_1) = \mathcal{O}(s_2)$ of §7.1.

The quotient is explicit. The substrate microstate space $\mathcal{S} = \{(s_i, U_{\{ij\}})\}$ carries strictly more degrees of freedom than the observable space $\mathcal{O} =$ (site values, loop holonomies, fluxes). The disparity is exactly the size of the χ -redundancy: $|V|$ real parameters of relabelling per substrate microstate that leave \mathcal{O} invariant. The quotient \mathcal{S}/\sim is precisely the space of physically distinguishable transport configurations, and the gauge group of the toy is the additive group of site functions $\chi : V \rightarrow \mathbb{R}$ modulo constants.

This realises gauge redundancy constructively rather than interpretively. The §7.1 quotient structure was mathematically standard but substrate-physically asserted; here it is exhibited

explicitly. Multiple substrate encodings of the link phases produce identical observable transport; the gauge transformation is the substrate operation that maps between them; the $U(1)$ gauge group of the continuum limit is the direct image of the discrete site-function relabelling group under coarse-graining.

The argument also clarifies a subtlety from §7.1: it shows constructively that the gauge group is not chosen but *derived* from the structure of the observation map. Different \mathcal{O} — e.g., one that returns the link phases $\theta_{\{ij\}}$ directly rather than only holonomies — would yield no gauge redundancy at all. The substrate-physical content of (B3) is therefore the substrate's restriction of observable access to closed-loop quantities, which is itself a substrate-level claim that the toy makes explicit.

16.5 The Emergence Hierarchy

The toy organises the programme's structural levels into a constructive emergence chain:

Level	Content	Status in toy
0	Substrate ticks at rate $1/\tau_s$	Fundamental
1	Distinguishability resolution at sites	$s_i \in \{-1, 0, +1\}$ dynamics
2	TPB: finite N_b -tick commitment cycle	Encoded in update rule
3	BCB: conserved transport in continuum limit	Σs_i conservation under coarse-graining
4	Gauge redundancy from observable equivalence	(*) under loop-holonomy \mathcal{O}
5	Local transport potential A_c^μ in IR limit	Emergent one-form from coarse-grained $\theta_{\{ij\}}$
6	Maxwell-form transport at $O(\epsilon^0)$	$\delta F = \mu_c J$ via §9 theorem applied to the IR limit
7	Wave propagation at speed c_c	Source-free §10.1 in the IR limit

The hierarchy is preliminary: most levels are present schematically rather than derived rigorously, and Level 6 requires verification that the coarse-grained continuum theory actually satisfies all of (B1)–(B4) under explicit limit-taking (which we do not perform here). What the toy establishes is that the levels can in principle be linked within a single constructive substrate; what remains open is the rigorous limit-taking from Level 0 to Level 7.

16.6 What This Toy Does and Does Not Establish

In the spirit of scope-honest VERSF practice, we close with an explicit catalogue.

This toy *does* establish:

- An explicit class of discrete substrates exists in which (A2.ii) atomic update, (B1) locality, and (B3) gauge redundancy via substrate degeneracy are jointly and

constructively realised, with (A1) BCB and (A2.i) finite-speed bound available given a choice of conservative update rule.

- The gauge-redundancy quotient structure of §7.1 admits an explicit substrate-physical realisation: relabelling (*) on substrate microstates, with closed-loop holonomy as the observable that defines the equivalence relation. This converts (B3) from interpretive to constructive within the toy.
- The TPB bound $c_c = \xi/(N_b \tau_s)$ emerges constructively given the substrate-physical reading of bit commitment as an N_b -tick cycle (§16.3), making (A2.i) substrate-visible rather than merely postulated.
- The admissibility class is non-vacuous in a substrate-physical (not merely mathematical) sense: at least one substrate satisfying (B1)–(B4) and admitting the substrate-physical interpretation we have asked for exists.

This toy *does not* establish:

- *That this is the substrate.* The toy is not unique; many discrete substrates would satisfy the admissibility class. The $K=7$ simplicial closure geometry of the broader VERSF programme is not built into the toy as specified, and a properly VERSF-native version would need to incorporate it.
- *A rigorous continuum limit.* Taking the IR limit and verifying that the resulting effective theory satisfies all of (B1)–(B4) with the correct values of ϵ_c, μ_c requires explicit coarse-graining and renormalisation analysis that is not performed here.
- *Constitutive derivation of saturation.* A natural follow-up would derive ϵ_c and μ_c from the substrate parameters (ξ, N_b, τ_s) under the toy's dynamics, and check whether the saturation $c'_c = c_c$ (§10.2) emerges algebraically. We do not carry out this derivation here. A naïve dimensional-analysis attempt does not yield the right answer cleanly; getting the algebra correct requires actual analysis of the coarse-graining limit, not just dimensional counting. The toy provides the structural setting in which such a derivation could in principle be carried out, but the derivation itself remains open work.
- *Current-selection.* The toy has a single $U(1)$ -valued link variable; it does not address the multi-current selection problem of §12.2. Extension to richer link-value structures (multiple commuting $U(1)$ s, or non-Abelian groups) would be needed.
- *Operational measurement theory.* The toy admits a candidate operational reading of commitment as "persistent multi-site coherence" — a bit is committed when its transport pattern stably propagates across multiple substrate sectors under continued evolution, in a manner reminiscent of decoherence and environment-induced classicality. This addresses Layer B of §3's operational gap *in preliminary form only*; converting it to a rigorous substrate measurement theory remains an open programme. The operational reading sketched here should be understood together with the Layer B companion papers referenced in §3.1: the persistent-multi-site-coherence picture is the substrate-toy-level instantiation, and the three measurement companion papers (relational closure, topological trapping, tick race) supply the interpretive and dynamical content with which it must eventually be unified.
- *Substrate uniqueness or ontology.* The toy is compatible with the agnosticisms of §2.6 (classical/quantum, fundamental/emergent gauge) but does not resolve them. It is one realisation; the deeper question of which substrate, if any, is physically actual remains.

The honest summary as of §16.6: the toy converts several of v9's "Known Missing Pieces" from "Not exhibited" to "Preliminary sketch in §16" — concretely, the substrate construction (OP 5), the gauge-redundancy substrate realisation, and the TPB-as-substrate-bound reading. The next subsection strengthens this further by exhibiting explicit discrete dynamics whose continuum limit reproduces Maxwell-form transport, closing the schematic gap in §16.2's deliberately-underspecified update rule.

16.7 Explicit Discrete Dynamics and the Continuum Maxwell Limit

The §16.2 toy specified objects and types but left the update rule schematic. We now exhibit a concrete dynamics, with an action principle, equations of motion, and a continuum limit that reproduces Maxwell-form gauge transport. This closes the principal dynamical gap of §16.6.

Lattice variables. Each site carries an integer-valued distinguishability charge

$$s_i(\lambda) \in \mathbb{Z}$$

(the $\{-1, 0, +1\}$ restriction of §16.2 is the simplest non-trivial truncation; the full integer-valued version is the natural generic case for the dynamical argument). Each link carries the phase variable $U_{ij} = \exp(i\theta_{ij})$ with $\theta_{ij} \in [0, 2\pi)$ as before, subject to the gauge transformation (*) of §16.4.

Plaquette curvature. For each elementary plaquette \square with oriented boundary $(i \rightarrow j \rightarrow k \rightarrow l \rightarrow i)$, define the plaquette holonomy

$$P_{\square} = U_{ij} U_{jk} U_{kl} U_{li} = \exp(i F_{\square}),$$

$$F_{\square} = \theta_{ij} + \theta_{jk} + \theta_{kl} + \theta_{li}.$$

Under the gauge transformation (*), each χ_n appears once with positive sign and once with negative sign as the closed loop passes through site n ; the χ -terms cancel pairwise. **F_{\square} is gauge-invariant** (modulo 2π , which is automatic for compact $U(1)$).

Discrete action. Take the standard Wilson form

$$S_{\text{lat}} = - (1/g^2) \sum_{\square} \text{Re}(P_{\square}) + \sum_{\langle i,j \rangle} \theta_{ij} J_{ij},$$

where g^2 is a dimensionless coupling and J_{ij} is the conserved link current built from the site charges (its specific form is fixed by requiring source-equation gauge invariance, which forces $\partial_{\text{link}} J = 0$ — the lattice version of BCB). Expanding $\text{Re}(P_{\square}) = \cos(F_{\square}) \approx 1 - F_{\square}^2/2$ in the small- a regime (where a is the lattice spacing and $\theta_{ij} \sim a A_{\mu}$ is small), the kinetic term reduces to

$$S_{\text{kin}} = - (1/(2g^2)) \sum_{\square} F_{\square}^2 + \text{const},$$

which is the leading-order Wilson action used in the body of the calculation below. We keep the $(1/(4g^2)) F_{\square}^2$ normalisation conventional in field-theoretic treatments (the factor of 1/2 vs 1/4 reflects whether plaquettes are summed with orientation or without).

Equations of motion. Varying S_{lat} with respect to $\theta_{\{ij\}}$ yields the discrete Maxwell equation

$$\Delta_{\mu} F_{\text{lat}}^{\{\mu\nu\}} = g^2 J_{\text{lat}}^{\nu},$$

where Δ_{μ} is the lattice forward-difference operator and $F_{\text{lat}}^{\{\mu\nu\}}$ is the plaquette-curvature tensor (the antisymmetric collection of F_{\square} values indexed by the plaquette's spacetime orientation). The discrete Bianchi identity

$$\Delta_{\{[\lambda} F^{\{\text{lat}\}}_{\mu\nu\}} = 0$$

follows automatically because the oriented sum of plaquette curvatures around the boundary of any closed three-cube is zero — each link appears once with each sign as the cube's plaquettes are traversed. This is the lattice content of $d^2 = 0$; the homogeneous Maxwell equations are structural identities of the lattice geometry, just as they are structural identities of differential forms in §8.

Continuum limit. Let a denote the spatial lattice spacing (the coherence length, $a = \xi$) and a_t the temporal spacing (the commitment-cycle duration, $a_t = N_b \tau_s$). For a link from site i to a nearest neighbour in direction μ , write

$$\theta_{\{ij\}}(\lambda) = a^{\{\mu\}} A_{\mu}(x_i, t_{\lambda}) + O(a^2)$$

where $a^{\{\mu\}} = a$ for spatial directions and $a^{\{0\}} = a_t$ for the temporal direction. Expanding F_{\square} around site x_i for a plaquette in the $(\mu\nu)$ -plane:

$$F_{\square} = a^{\{\mu\}} a^{\{\nu\}} (\partial_{\mu} A_{\nu} - \partial_{\nu} A_{\mu}) + O(a^3) = a^{\{\mu\}} a^{\{\nu\}} F_{\mu\nu} + O(a^3).$$

The lattice difference becomes $\Delta_{\mu} \rightarrow a^{\{\mu\}} \partial_{\mu}$. After absorbing lattice-spacing factors into the field rescaling and identifying $\mu_c = g^2 / (\text{combination of lattice spacings})$, the discrete equations of motion become

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

$$\partial_{\{[\lambda} F^{\{\mu\nu\}} = 0,$$

with $F_{\{\mu\nu\}} = \partial_{\mu} A_{\nu} - \partial_{\nu} A_{\mu}$. **These are exactly the Maxwell-form U(1) gauge-transport equations of §9 and §8.**

The continuum wave speed and what it implies for saturation. Because the lattice spacings encode substrate parameters ($a = \xi$ spatially, $a_t = N_b \tau_s$ temporally), the continuum dispersion relation derived from the source-free equations has propagation speed

$$c'_c = a / a_t = \xi / (N_b \tau_s) = c_c,$$

matching the TPB bound (4) by construction. **This is potentially significant for §10.2.** The saturation $c'_c = c_c$ is not put in by hand; it falls out of the continuum limit as a consequence of the substrate parameters appearing in the lattice spacings. The constitutive constants ϵ_c and μ_c are determined by g^2 , a , and a_t , and their ratio is therefore fixed by the substrate parameters. We do not claim this fully closes Route 1 of §10.2 — making the argument rigorous requires (i) explicit specification of the temporal-vs-spatial Wilson action in anisotropic lattice form, (ii) careful tracking of dimensionful factors through the continuum limit including coupling-constant renormalisation under coarse-graining, and (iii) a derivation that the field-theoretic constants ϵ_c , μ_c emerge with the right form algebraically rather than just dimensionally. What we *do* claim is that §16.7 exhibits the framework in which Route 1 is naturally carried out, and that the saturation relation appears as a structural consequence of the substrate lattice spacings rather than as an additional postulate.

Summary of what §16.7 demonstrates. The full constructive chain

discrete link substrate (§16.2) → conserved transport currents (§16.3) → gauge redundancy from substrate degeneracy (§16.4) → discrete action and Maxwell-form equations of motion (§16.7) → continuum Maxwell-form U(1) gauge transport (§16.7)

is now present in outline. The earlier statement "a discrete substrate satisfying (A1)–(A2) and admitting the admissibility class generates Maxwell-form transport under coarse-graining" is no longer a conjecture about hypothetical substrates: it is a property of *this* substrate, derived from its explicit action via standard lattice-gauge methods.

Caveats and remaining work.

- *Standard LGT machinery.* The discrete dynamics of §16.7 are essentially standard compact-U(1) lattice gauge theory. What makes the construction *VERSF-specific* rather than generic LGT is the substrate-physical interpretation (link variable as committed-distinguishability transport relation, sites as substrate update locations, lattice spacings as substrate-physical ξ and $N_b \tau_s$) rather than the dynamical machinery itself. The LGT-as-regulator vs LGT-as-substrate distinction of §17.4 OP 5 still applies: §16.7 inherits the machinery and supplies the substrate-physical reading, but the reading is interpretive rather than derived.
- *K=7 not yet incorporated.* The Wilson plaquette of §16.7 is geometry-agnostic — it works on any planar 4-cycle in the substrate graph. A VERSEF-native version must use the K=7 simplicial closure geometry, with the plaquette structure inherited from K=7 closure relations. This remains the principal substrate-geometric gap.
- *Coarse-graining made rigorous.* The continuum limit is taken at the level of leading-order expansion of the Wilson action; a rigorous IR limit including renormalisation-group flow of g^2 , ϵ_c , μ_c remains to be performed. This is the central technical task identified at the end of §16.6.
- *Lorentz covariance.* The continuum limit produces a Lorentz-form dispersion relation only if the spatial and temporal lattice spacings produce an isotropic effective metric,

which is the substrate-physical content of the saturation relation. A non-saturating substrate would produce an effective metric with anisotropy between (A2.i) and (B4), which would be observable as Lorentz violation in EM propagation.

- *Weak-coupling phase assumption.* The continuum-limit expansion $\cos(F_{\square}) \approx 1 - F_{\square}^2 / 2$ assumes that the link variables $\theta_{\{ij\}}$ are small in lattice units ($a \cdot A_{\mu} \ll 1$), which is the *weak-coupling* perturbative regime of compact U(1) lattice gauge theory. The appearance of Maxwell theory at leading order in this expansion is exactly what is expected from any compact U(1) LGT in its weak-coupling phase. In the *strong-coupling* phase of compact U(1) LGT, qualitatively different physics emerges — monopole condensation, confinement in 2+1 dimensions (Polyakov), non-trivial topological sectors at scale a — none of which reproduces continuum Maxwell theory. **The substrate must therefore lie in the weak-coupling phase to reproduce empirical EM**, which is a non-trivial constraint on substrate parameters (specifically, on the bare coupling g^2 relative to the lattice spacings). A complete VERSF-native construction must either derive that the substrate's parameters place it in the weak-coupling phase, or treat the weak-coupling assumption as an additional admissibility condition alongside (B1)–(B4). We flag this here so it is not concealed by the leading-order expansion notation.

With these caveats stated, §16.7 advances the toy from "constructive in ontology but schematic in dynamics" to "**constructive in ontology and dynamically explicit, modulo rigorous coarse-graining and $K=7$ incorporation.**" The qualitative chain from discrete substrate to continuum Maxwell theory is now exhibited; the rigorous form of the chain is the next phase of work.

16.8 Toward the $K = 7$ Wilson Limit

The §16.7 lattice action is geometry-agnostic: the Wilson plaquette structure applies to any planar 4-cycle in the substrate graph, and the resulting dynamics could equally well live on a hypercubic lattice. This is the principal substrate-geometric gap. A genuinely VERSF-native version must use the $K = 7$ simplicial closure geometry of the companion paper *$K = 7$: The Constraint Dimensionality of Stable Physical Reality*, in which the elementary curvature carriers are not arbitrary square plaquettes but the **$N_{\text{loop}} = 14$ independent electromagnetic phase loops** of the minimal hexagonal interface (11 boundary edges of the two-cell interface, plus 2 non-edge loops through the shared boundary's interior vertices, plus 1 global closure mode).

A note on inherited scope. The $K = 7$ results invoked in this subsection — hexagonal interface tiling, $N_{\text{loop}} = 14$, $\alpha^{-1}_{\text{bare}} = 137.143$, the C4 winding-number argument — are derivations of the companion $K = 7$ paper, which has its own scope-honest discipline distinguishing derived results, results with open derivational steps, and interpretive connections. The present paper inherits whatever conditionality attaches to each imported result. In particular: the hexagonal interface is derived (via two independent routes — Honeycomb Theorem and $\text{PAR} + \text{CC}_G$); $N_{\text{loop}} = 14$ is a derivation in a separate interface-realisation companion paper cited by the $K = 7$ paper; $\alpha^{-1}_{\text{bare}} = 137.143$ is the bare term, with the IPR correction to the empirical 137.036 still under active derivation in the $K = 7$ paper's Appendix A; the joint α - Λ constraint is conditional on the $\xi(K)$ coherence-scale derivation in a separate coherence-scale companion paper. Where this subsection writes a $K = 7$ paper result as a flat numerical fact (e.g., " $\alpha^{-1}_{\text{bare}} \approx$

137.143"), the reader should understand "with whatever conditionality the $K = 7$ paper attaches to that quantity."

The $K = 7$ Wilson action. Replace the §16.7 plaquette sum with a sum over interface cells and the 14 loops per cell:

$$S_{\{K=7\}} = - (1/g^2_{\{K=7\}}) \sum_I \sum_{\{\ell=1\}^{\wedge}\{14\}} w_{\ell} \cos(F_{\{I\ell}\}) + \sum_{\{(i,j)\}} \theta_{\{ij\}} J_{\{ij\}},$$

where:

- I labels minimal hexagonal interface cells in the substrate,
- $\ell = 1, \dots, 14$ labels the independent phase loops within each cell (matching $N_{\text{loop}} = 14$ derived in the $K = 7$ paper's cited interface-realisation companion),
- $F_{\{I\ell\}} = \sum_{\{(i,j) \in \partial(I\ell)\}} \theta_{\{ij\}}$ is the holonomy around the ℓ -th loop in cell I ,
- w_{ℓ} are closure weights determined by the loop structure: weight 1 for each of the 11 boundary-edge loops, weight 1 for each of the 2 interior-vertex non-edge loops, and weight 1 for the global closure-mode loop (uniform weighting is the natural first choice; non-uniform weights would be the next refinement).

Gauge invariance carries over identically from §16.4: under $\theta_{\{ij\}} \rightarrow \theta_{\{ij\}} + \chi_i - \chi_j$, each $F_{\{I\ell\}}$ is invariant because the χ -contributions cancel around any closed loop. The $K = 7$ Wilson action therefore enjoys the same gauge symmetry as the §16.7 action, now anchored on $K = 7$ -native loops rather than on arbitrary plaquettes.

The target. With the $K = 7$ paper's bare coupling derivation in hand,

$$g^2_{\{K=7\}} \rightarrow \alpha_{\text{bare}} = 2^{-7} \cdot 14/15 = (14)/(15 \cdot 128),$$

equivalently

$$\alpha^{-1}_{\text{bare}} = 2^7 \cdot 15/14 = 128 \cdot 15/14 \approx 137.143,$$

the $K = 7$ Wilson Limit paper's central calculation is to show that the continuum limit of $S_{\{K=7\}}$ produces Maxwell-form $U(1)$ gauge transport with the effective coupling μ_c (this paper) and the bare electromagnetic coupling α_{bare} ($K = 7$ paper) reading the same underlying integer-counting structure. The IPR correction connecting $\alpha_{\text{bare}} = 137.143$ to the empirical $\alpha^{-1} = 137.036$ is under active derivation in the $K = 7$ paper's Appendix A; the bridge to the present paper's μ_c is independent of that correction and concerns only the bare term.

What this subsection accomplishes and what it does not. §16.8 makes the $K = 7$ bridge *explicit* rather than promised: it states the lattice action on the $K = 7$ hexagonal geometry, fixes the loop content from the $N_{\text{loop}} = 14$ derivation cited by the $K = 7$ paper, and identifies the target convergence between $g^2_{\{K=7\}}$ and α_{bare} . It does *not* derive that convergence — performing the coarse-graining from $S_{\{K=7\}}$ to the continuum Maxwell theory of §9, computing the effective ε_c , μ_c as functions of $g^2_{\{K=7\}}$ and lattice spacings, and verifying that $c'_c = c_c$ emerges as a fixed-point identity (in the sense of Appendix B's anisotropic

Wilson framework) is the technical task of the $K = 7$ Wilson Limit paper. The framework is now exhibited; the calculation that closes saturation and unifies the two coupling constants remains the next paper's work.

The principal advance of §16.8 over §16.7 is therefore structural: the $K = 7$ geometric content is no longer external commentary about what would distinguish a VERSF-native toy from generic LGT. It is now a specific lattice action with $N_{\text{loop}} = 14$, hexagonal interface tiling forced by both the Honeycomb Theorem and $\text{PAR} + \text{CC}_G$ axioms, and an explicit target convergence with the $K = 7$ paper's α derivation. The $K = 7$ Wilson Limit paper inherits this scaffolding ready-made.

17. Conclusion and Open Problems

Within VERSF, classical electromagnetism is not a fundamental interaction imposed upon spacetime. It is the admissible linear transport geometry of conserved commitment imbalance within a substrate-motivated admissibility class, at leading order in the coherence-scale expansion. Maxwell's equations are not postulated; they are the closure-compatible outcome of BCB plus TPB plus a small set of admissibility constraints which those substrate principles motivate.

The source-sector Helmholtz decomposition of the commitment current provides a substrate-level interpretation of the electric–magnetic distinction. The antisymmetric tensor structure of the field is forced by gauge redundancy understood as informational degeneracy — formalised as the quotient of substrate microstates by the equivalence relation "producing identical observable commitment transport." The homogeneous equations follow as a structural identity ($dF_c = 0$) — and within that identity, the substrate exclusion of fundamental $U(1)$ magnetic charge becomes a falsifiable prediction. The inhomogeneous equation ($\delta F_c = \mu_c J_c$) is the surviving admissible source coupling at leading order, automatically respecting BCB through $\delta^2 = 0$. Wave propagation in vacuum saturates the TPB bound, and light becomes the substrate-level oscillation between the two source sectors.

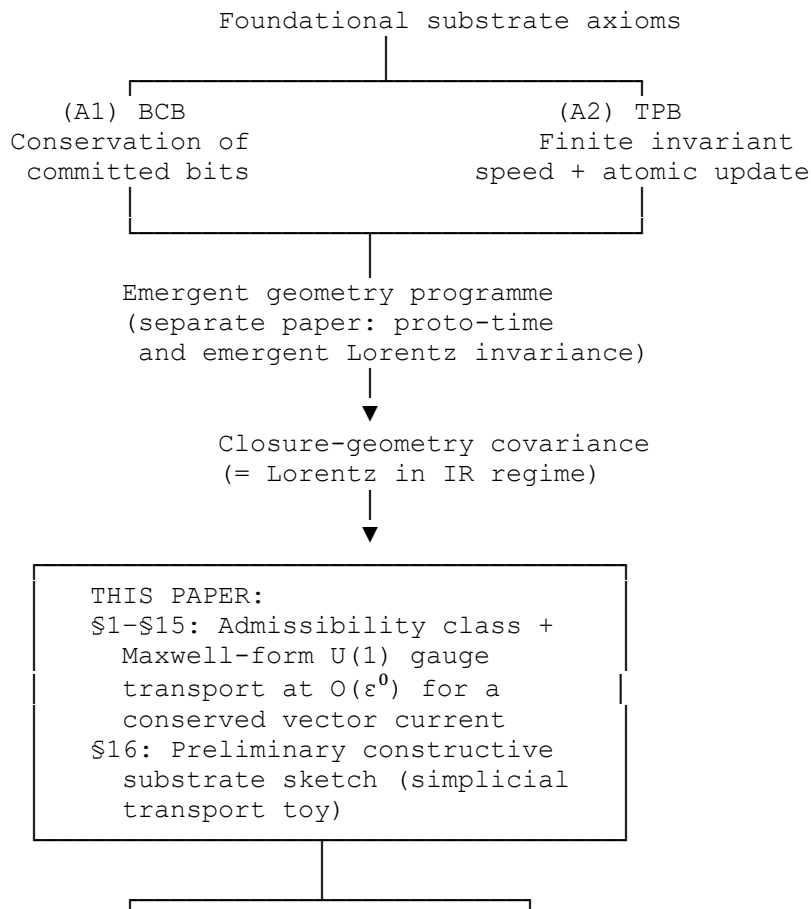
What this reconstruction does not do is replace experimental electrodynamics with a derivation from pure substrate principles, and it does not exhibit a concrete substrate model whose dynamics generate the admissibility constraints from below. **The present paper is structural rather than constructive: it argues for the admissibility class and proves uniqueness within it, but does not specify substrate dynamics (lattice rules, cellular automata, simplicial transport networks, information-flow graphs) that would generate (B1)–(B4) explicitly.** The admissibility class carries real physical content, and although it is well-motivated within the VERSF programme, the constraints are not theorems of BCB and TPB. The honest claim is the conditional one: **given** this admissibility class, Maxwell-form gauge transport is forced within it at $O(\epsilon^0)$; and BCB and TPB give the class its substrate-level justification.

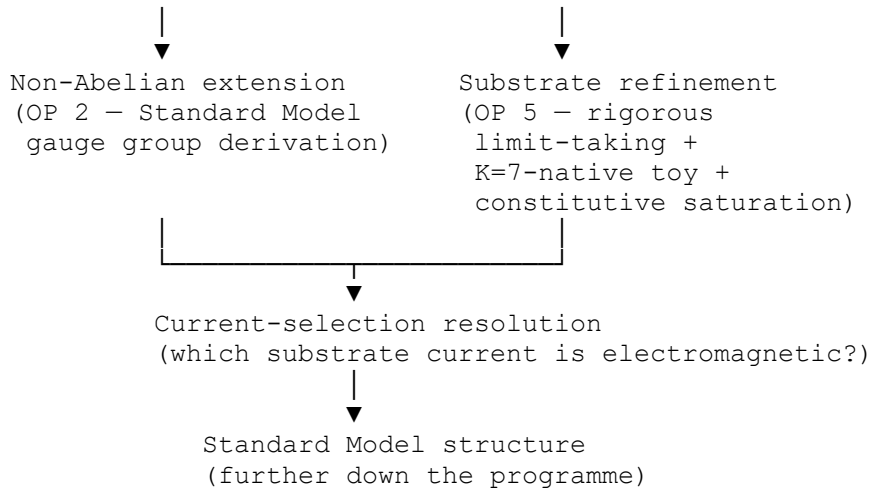
17.1 What This Paper Actually Accomplishes

In one paragraph: the present paper does not derive electromagnetism from a complete substrate model. What it establishes is that, once a conserved vector current, gauge redundancy, locality, closure-geometry covariance, and ξ -controlled truncation are admitted, **Maxwell-form U(1) gauge transport emerges as the surviving admissible leading-order structure**. The empirical labelling of this structure as "electromagnetism" requires two further inputs — the saturation postulate of §10.2 and the current-identification of §12.2 — and §16.7 substantially advances the first of these by exhibiting a discrete substrate whose continuum wave speed equals the TPB bound by construction. The constructive sketch of §16 now contains explicit discrete dynamics, equations of motion, and a continuum limit reproducing Maxwell-form transport (§16.7), in addition to its earlier substrate-physical realisation of the gauge-redundancy quotient structure (§16.4). The construction is still preliminary — K=7 incorporation, rigorous coarse-graining, current-selection, and full measurement theory remain open — but the abstract admissibility class is now populated with explicit dynamics rather than only with structural commitments. This is a *structural admissibility theorem with a preliminary but dynamically explicit constructive realisation*, not a complete substrate-physical derivation, and it should not be over-read as the latter; nor should it be under-read as a tautology, since the admissibility class itself is substantively motivated by BCB and TPB rather than imported from standard gauge theory.

17.2 Programme Architecture

This paper sits in a broader VERSF programme arc. The architecture, schematically:





The present paper occupies the middle of the arc: it sits downstream of the substrate axioms and the emergent-geometry programme, and upstream of the non-Abelian extension, substrate construction, and current-selection resolution. Several of the open problems below are the natural next branches in this architecture.

17.3 Known Missing Pieces

For full transparency, we collect here the structural items the present paper has *not* derived. The status column has been updated to reflect §16's preliminary constructive sketch including the explicit dynamics of §16.7.

Missing piece	Status	Path forward
Concrete substrate dynamics	Exhibited in §16.7 (Wilson-form action, equations of motion, continuum limit) — not yet K=7-native	OP 5 (rigorous coarse-graining + K=7 incorporation)
Current-selection (which current is EM)	Not resolved	OP 2 (non-Abelian + gauge group)
Saturation derivation ($c'_c = c_c$)	Partially closed via §16.7: $c'_c = a/a_t = \xi/(N_b \tau_s)$ emerges from lattice spacings; full rigour pending anisotropic-Wilson + RG analysis	§10.2 Route 1, made rigorous within §16.7 framework
Non-Abelian gauge structure	Not addressed	OP 2
Operational definition of commitment	Layer A: §16.6 closure-stabilisation sketch; Layer B: addressed in the three measurement companion papers introduced in §3.1; integration into the Maxwell substrate remains pending	Substrate measurement theory (separate programme; companion papers supply the chain)
Numerical anchoring of ξ	Qualitative only	OP 6

Missing piece	Status	Path forward
Locality as theorem	Argued, not proved; constructively realised in §16 toy	OP 4
(B2) hierarchy from substrate UV integration	Power-counting analogy; §16.7 framework supports explicit Wilsonian coarse-graining	Performed within §16.7 toy
Substrate-physical realisation of \mathcal{S} / \sim	Constructively realised in §16.4 (relabelling $(*) + \text{loop-holonomy } \mathcal{O}$)	Refinement and uniqueness of realisation
Continuum Maxwell limit from substrate	Demonstrated in outline in §16.7	Rigorous IR limit including RG flow
Gauge ontology (fundamental vs emergent)	Agnostic	Deeper substrate programme
Substrate ontology (classical vs quantum)	Agnostic	Deeper substrate programme

The combined §16.7 result advances four rows from "Not exhibited / Not derived" to "Partially closed" or "Demonstrated in outline": substrate dynamics, saturation derivation, (B2) hierarchy, and continuum Maxwell limit. The principal item that remains wholly open is current-selection (OP 2 / non-Abelian); Layer B measurement theory is addressed by the companion measurement papers (§3.1), with substrate-integration pending.

17.4 Open Problems

1. **Can linearity be promoted from $O(\epsilon^0)$ admissibility to theorem at the single-bit level?** The ξ -EFT truncation treats linearity as leading-order admissibility. A more ambitious result would derive linearity at the single-bit transport level from substrate combinatorics. Candidate route: nonlinear self-interactions of the commitment field correspond to composite multi-bit commitment events whose admissibility is governed by the $K=7$ simplicial structure of the substrate. If single-bit transport processes are forced to be combinatorially distinct from composite events, then the single-bit sector must be linear, with nonlinear corrections appearing only when multiple commitment events overlap within a coherence volume. Whether this argument can be made rigorous is an open question.
2. **The non-Abelian extension.** The present construction yields the $U(1)$ sector. Extension to non-Abelian gauge theory requires generalising (B3) to a non-Abelian transformation structure. The relevant question is whether the hexagonal closure geometry of the VERSF substrate forces a particular non-Abelian structure, so that the Standard Model gauge group emerges from substrate admissibility rather than being imposed.
3. **Quantitative prediction of nonlinear corrections.** Born–Infeld and Euler–Heisenberg extensions are well-motivated higher-order corrections in the ξ -EFT. Whether the substrate construction predicts the *scale* of these corrections from substrate parameters —

placing concrete numerical bounds on departures from linear Maxwell theory in strong-field regimes — is an open programme.

4. **Locality as a theorem.** The argument in §5 ties locality to TPB but stops short of a theorem. A rigorous derivation of (B1) from sharper VERSF axioms would close one of the principal conditional gaps in the present construction.
5. **Refinement of the constructive substrate sketch.** §16 presents a preliminary simplicial transport toy that begins to populate the abstract admissibility class with concrete dynamics — including substrate-physical readings of the link variable as committed-distinguishability transport, an explicit TPB-consistent commitment cycle, and a concrete realisation of the §7.1 quotient structure via the link-relabelling (*) of §16.4. The toy is not unique, not K=7-native, and several follow-up tasks remain explicit: (i) rigorous coarse-graining from the discrete dynamics to the continuum Maxwell-form equations of §9, (ii) constitutive derivation of ϵ_c and μ_c from the substrate parameters (ξ , N_b , τ_s) and verification of the saturation $c'_c = c_c$, (iii) incorporation of the K=7 simplicial closure geometry of the broader VERSF programme, distinguishing the toy from purely cubic-lattice realisations, (iv) extension to richer link-value structures supporting multiple conserved vector currents (necessary for the current-selection problem of §12.2), and (v) development of the substrate measurement-theory sketch of §16.6 into a rigorous Layer B for the operational definition of commitment. The §16 toy converts "the substrate motivates the constraints" from descriptive language into one verified instance of joint realisation; converting it into a *finished* substrate construction — not merely a sketch — remains the most important next step in the programme.

K=7-native geometric content now available. As of the companion paper *K = 7: The Constraint Dimensionality of Stable Physical Reality* (Taylor, AIDA Institute), the geometric content of "K=7-native" is no longer unspecified. That paper establishes — via the Honeycomb Theorem (variational geometry) and independently via PAR + CC_G (axiomatic closure) — that the admissible interface is uniquely hexagonally tiled, with six boundary directions per hexagonal cell plus one global closure mode (giving rank 6, nullity 1, $K = 7$). The minimal interface carries $N_{loop} = 14$ independent electromagnetic phase loops (11 boundary edges of a two-cell interface, plus 2 non-edge loops through the shared boundary's interior vertices, plus 1 global closure mode). A K=7-native version of the §16.7 lattice action therefore has specific geometric content: replace the geometry-agnostic Wilson plaquette structure with the hexagonal interface tiling, with 14 link variables per minimal interface unit and $2^K = 128$ distinguishable constraint configurations per Planck-area patch (the Bekenstein–Hawking-entropy bridge connecting this constraint count to the boundary entropy $A/4$ is developed in the $K = 7$ paper's constraint C2, not in the present paper). This is the natural starting point for the K=7 Wilson Limit paper identified as the next priority output of the programme.

Existence proof from lattice gauge theory — with an important caveat. Beyond the §16 sketch, it is worth noting that standard lattice gauge theory also constitutes an *abstract existence proof* for the admissibility class: a regular lattice with U(1)-valued link variables, a discrete update rule that conserves total charge, and a plaquette construction generating the field tensor — in the continuum limit — satisfies all four admissibility constraints and reproduces Maxwell electrodynamics. The admissibility class is therefore

non-vacuous in pure mathematics: at least one discrete structure satisfying (B1)–(B4) is well-known to physics.

But LGT is conventionally a regulator, not a substrate. The caveat that cuts in the opposite direction: lattice gauge theory in the standard programme is treated as a UV regulator — a calculational scaffold for defining gauge theory non-perturbatively, with the lattice spacing taken to zero in the continuum limit. Nobody in standard practice claims that the universe *is* a hypercubic lattice; the lattice is computational machinery, not a physical posit. So pointing to LGT establishes that "some discrete structure" can satisfy (B1)–(B4), but it does *not* establish that a discrete structure can do so *while being a physical substrate rather than a calculational regulator*. The §16 toy inherits structure from LGT but begins to extend it toward substrate-physical interpretation by reading the link variable as committed-distinguishability transport relations rather than as abstract phase machinery. The conceptual distinction matters: a regulator is a tool one removes; a substrate is a thing one keeps.

6. **Numerical anchoring of the ξ -expansion.** The expansion parameter $\varepsilon = \xi/L$ is qualitatively controlled but quantitatively unanchored. **A clarification of which ξ is at play is now important:** the (B2) ξ -EFT expansion uses the *substrate-UV* coherence length over which a single bit is locally resolved (§4), which is plausibly near-Planckian. The companion $K = 7$ paper introduces a separate quantity — the *cosmological* coherence length on the de Sitter horizon, $\xi(K) \propto \exp(\pi \cdot 2^K / 2)$ from coherence-decay arguments on the K -dimensional binary constraint space, with absolute scale set by the coherence-scale proportionality constant (derived in a separate companion paper cited by the $K = 7$ paper) to match the observed cosmological scale $\Lambda \propto \xi_{\text{cosmo}}^{-2}$. These are two different lengths in the same framework: $\xi_{\text{substrate}}$ is a UV substrate-scale property (entering the (B2) EFT expansion), while ξ_{cosmo} is an IR cosmological-scale property (entering the $K = 7$ paper's constraint C4 joint α - Λ derivation). Both depend on the same $K = 7$, but they are *not* the same quantity.

Substrate-UV anchoring. The (B2) expansion parameter $\varepsilon = \xi_{\text{substrate}}/L$ is anchored from above by experimental bounds on photon-dispersion deviations (§14.1 Prediction A): current high-energy gamma-ray burst, sub-PeV neutrino, and atomic-clock observations are consistent with $\Lambda_{\text{substrate}} = \xi_{\text{substrate}}^{-1}$ at or near the Planck scale. Tightening these bounds in near-term observational programmes tightens the substrate-UV anchor from below.

Cosmological anchoring. The $K = 7$ paper's ξ_{cosmo} enters the joint α - Λ constraint (§14.1 Prediction C), with the absolute placement of $K = 7$ in the cosmological band derived in a separate coherence-scale companion paper cited by the $K = 7$ paper. The $K = 7$ paper's Section 6.2 establishes that for $K = 7$ to lose its position to $K = 6$ or $K = 8$ in the cosmological band, the proportionality constant would need to be wrong by a factor of $\sim 10^{53}$ — making this anchoring extraordinarily robust.

Born–Infeld threshold. With $\xi_{\text{substrate}}$ at or near Planck and $\Lambda_{\text{substrate}} \sim \text{Planck energy}$, the (B2) EFT predicts nonlinear EM corrections at $F^2 \sim \Lambda_{\text{substrate}}^4 \sim$

M_{Planck}^4 in natural units — empirically safe but not a near-term falsifier. The precision-propagation handles of §14.1 Prediction A and the α - Λ joint-stability handle of §14.1 Prediction C are the better near-term empirical tests.

Quantitative anchoring of ξ has thus moved from "qualitative only" (in v11) to "specific functional form derived for ξ_{cosmo} with absolute scale pending the coherence-scale companion paper; substrate-UV ξ bounded by precision propagation tests." The structural prediction that α and Λ derive from the same $K = 7$ — and therefore covary under any future drift observation — is the most distinctive observational consequence.

Appendix A. Why Chern–Simons-Type Terms Fail Under (B4)

The exclusion of Chern–Simons-type contributions to the source equation (§9 enumeration, candidate 4) is brief in the body and worth expanding for readers unfamiliar with the Lorentz-violating electrodynamics literature.

A Chern–Simons-type extension of Maxwell theory in 3+1 dimensions modifies the source equation to include a term of the form

$$\partial_{\mu} F_{\nu}^{\mu\lambda} + k^{\nu}{}_{\rho\sigma\lambda} \partial^{\rho} F_{\nu}^{\sigma\lambda} + (\text{other index contractions}) = \mu_{\text{c}} J_{\nu}^{\nu},$$

where the additional contraction involves the totally antisymmetric tensor $\varepsilon^{\mu\nu\rho\sigma}$ and a fixed background vector k^{μ} (or, equivalently, a fixed antisymmetric tensor $k_{\{\mu\nu\}}$). The Carroll–Field–Jackiw model and its descendants treat k^{μ} as a non-dynamical background field that distinguishes a preferred direction in spacetime.

Such terms are excluded under (B4) for three connected reasons:

1. **Closure-geometry covariance demands no preferred background.** (B4) requires the transport equations to transform covariantly under the substrate's closure-geometry symmetry group. A non-dynamical k^{μ} is, by construction, *not* a covariant object: under the symmetry group it transforms as a fixed background, breaking the symmetry rather than respecting it. Within the emergent-Lorentz regime, (B4) reduces to Lorentz covariance, and a fixed k^{μ} explicitly breaks Lorentz invariance.
2. **Isotropy.** A preferred vector singles out a direction in space and, in the temporal component, a preferred time direction. The substrate's closure-geometry symmetry is, by hypothesis, isotropic at the leading order under consideration (this is part of the inheritance from the emergent-Lorentz companion paper). A Chern–Simons-type term destroys this isotropy by construction.
3. **Empirical bounds.** Even setting aside (B4) on theoretical grounds, observational bounds on Carroll–Field–Jackiw-type Lorentz-violating electrodynamic terms are very tight (constrained by cosmic microwave background polarisation rotation, ultra-high-energy

cosmic-ray spectra, atomic clock comparisons, and gamma-ray burst measurements). The empirical constraints are consistent with the strict (B4) exclusion and inconsistent with any substantial Chern–Simons-type modification at $O(\epsilon^0)$.

The exclusion is therefore not arbitrary: it is forced by closure-geometry covariance at the structural level, supported by isotropy at the symmetry level, and consistent with empirical bounds at the observational level. A substrate-physical Chern–Simons-type term would require some substrate process to supply a fixed background vector covariantly — which the present construction does not provide and which the emergent-Lorentz companion paper actively excludes within its regime of validity.

A subtlety: in 2+1 dimensions, Chern–Simons electrodynamics is a respectable gauge theory with no preferred background vector, since the $\epsilon^{\mu\nu\rho}$ tensor in 2+1 dimensions provides a covariant antisymmetric structure. The 3+1 exclusion does not apply there. The present paper is concerned with the 3+1 case throughout.

Appendix B. Anisotropic Wilson Coarse-Graining and the Continuum Limit

The continuum limit derivation in §16.7 was carried at tree level, with isotropic plaquette structure and a single coupling g^2 . This appendix exhibits the natural framework in which the saturation argument of §10.2 Route 1 can be made rigorous: the anisotropic Wilson formulation with separate spatial and temporal couplings. The full RG analysis is deferred to the $K = 7$ Wilson Limit paper; what we do here is make the structural framework explicit so that the saturation problem can be seen as a question with a specific technical shape.

Anisotropic Wilson action. Distinguishing spatial plaquettes \square_s (those lying in a constant-time slice) from temporal plaquettes \square_t (those involving the time direction), define

$$S_W = \beta_s \sum_{\square_s} (1 - \cos F_{\square_s}) + \beta_t \sum_{\square_t} (1 - \cos F_{\square_t}),$$

with lattice spacings

$$a = \xi \text{ (spatial coherence length, §16.3)}$$

$$a_t = N_b \tau_s \text{ (temporal commitment-cycle length, §16.3)}$$

and bare **anisotropy parameter**

$$\gamma_{\text{bare}} = a / a_t = \xi / (N_b \tau_s).$$

In standard anisotropic lattice gauge theory (Karsch 1982; subsequent finite-temperature LGT literature), the two couplings β_s, β_t are tuned so that the effective renormalised anisotropy

γ_{eff} in the continuum limit equals the geometric anisotropy γ_{bare} — restoring Euclidean invariance at the IR fixed point.

Continuum limit and target. Expanding around small $\theta_{\{ij\}}$ in each plaquette type:

$$F_{\{\square s\}} \approx a^2 F_{\{ij\}}, F_{\{\square t\}} \approx a \cdot a_t \cdot F_{\{0i\}},$$

and substituting into S_W with appropriate dimensional factors yields, in the leading continuum limit,

$$S_{\text{cont}} = - (1/4) \int d^4x [(1/\epsilon_c) F_{\{ij\}} F^{\{ij\}} + \mu_c F_{\{0i\}} F^{\{0i\}}],$$

with ϵ_c determined by β_s and a , and μ_c determined by β_t and a_t . The continuum wave speed reads off

$$(c'_c)^2 = 1 / (\epsilon_c \mu_c).$$

The target. Closure of Route 1 of §10.2 requires showing that the renormalised anisotropy at the IR fixed point satisfies

$$\gamma_{\text{eff}} = \xi / (N_b \tau_s) = c_c \text{ (saturation),}$$

so that $c'_c = c_c$ becomes a theorem of the RG analysis rather than a separately postulated identification. The RG flow of γ under coarse-graining is what would have to deliver this. Three regimes are logically possible:

- *Fixed-point preservation.* The RG flow preserves the bare anisotropy γ_{bare} without renormalisation. This would make saturation automatic.
- *Fixed-point attraction to $\gamma^* = c_c$.* The RG flow renormalises γ but flows it toward the Lorentz-compatible fixed point $\gamma^* = \xi / (N_b \tau_s)$ regardless of small perturbations in the bare value. This would still deliver saturation in the IR.
- *Fixed-point attraction to $\gamma^* \neq c_c$.* The RG flow drives γ to some IR fixed point not equal to c_c . This regime *fails* saturation and would be empirically refuted by precision propagation tests (§14.1 Prediction A).

The substantive technical question for the $K = 7$ Wilson Limit paper is therefore not "does some regime of saturation hold" but **which of the three regimes actually obtains for this substrate.** The corresponding technical conjecture, which that paper must prove or disprove:

Conjecture (Unique Lorentz-Compatible IR Fixed Point). The RG flow of the anisotropic Wilson action defined above, with bare lattice spacings $a = \xi$ and $a_t = N_b \tau_s$ set by substrate-physical parameters, has a unique IR fixed point in the anisotropy parameter γ , located at $\gamma^* = c_c$.

This conjecture is the substantive content of the saturation problem. *If* the conjecture holds, saturation is automatic; *if not*, the framework predicts Lorentz-incompatible substrate transport

and is refuted at the level of §14.1 Prediction A. The conjecture is therefore as falsifiable as the rest of the construction. The $K = 7$ Wilson Limit paper's principal calculation is to settle this question one way or the other.

Sober status. The present paper gives the tree-level continuum limit. **Full closure requires proving the Unique Lorentz-Compatible IR Fixed Point conjecture stated above.** The anisotropic Wilson framework is the natural setting for this analysis: standard anisotropic lattice gauge theory literature (Karsch 1982 and successors) establishes that fixed-point attraction to a Lorentz-compatible point is generic for Euclidean-invariance restoration in finite-temperature LGT, suggesting the conjecture is more likely true than not, but the substrate-specific case has not been settled. The $K = 7$ Wilson Limit paper carries out the explicit calculation.

The structural conclusion stands: saturation is not arbitrary, it is the IR-fixed-point statement of an RG analysis whose technical scaffolding (anisotropic Wilson + Lorentz-restoration RG flow) is mature in the LGT literature and whose substrate-physical input (ξ and $N_b \tau_s$ as the bare lattice spacings) is supplied by the §16 toy.

References

Standard physics

Aharonov, Y., & Bohm, D. (1959). Significance of electromagnetic potentials in the quantum theory. *Physical Review* **115**, 485–491.

Bianchi, L. (1902). *Lezioni di Geometria Differenziale*. Spoerri, Pisa. [Origin of the differential identity bearing his name.]

Born, M., & Infeld, L. (1934). Foundations of the new field theory. *Proceedings of the Royal Society A* **144**, 425–451. [The leading historical example of a nonlinear extension of Maxwell theory in the present EFT framing.]

Deser, S. (1970). Self-interaction and gauge invariance. *General Relativity and Gravitation* **1**, 9–18. [The gauge-self-coupling perspective on uniqueness in Abelian and non-Abelian gauge theories.]

Dirac, P. A. M. (1931). Quantised singularities in the electromagnetic field. *Proceedings of the Royal Society A* **133**, 60–72. [The compatibility of magnetic monopoles with quantum mechanics, referenced in §8.1 and §14.]

Euler, H., & Heisenberg, W. (1936). Folgerungen aus der Diracschen Theorie des Positrons. *Zeitschrift für Physik* **98**, 714–732. [The other leading historical nonlinear extension recovered at higher order in the ξ -expansion.]

Helmholtz, H. (1858). Über Integrale der hydrodynamischen Gleichungen, welche den Wirbelbewegungen entsprechen. *Journal für die reine und angewandte Mathematik* **55**, 25–55. [Origin of the decomposition used in §6.]

Hodge, W. V. D. (1941). *The Theory and Applications of Harmonic Integrals*. Cambridge University Press. [Foundational treatment of the codifferential and Hodge duality used in §8–§11.]

Jackson, J. D. (1998). *Classical Electrodynamics* (3rd ed.). Wiley. [Standard textbook treatment of Maxwell theory and gauge invariance.]

Karsch, F. (1982). SU(N) gauge theory couplings on asymmetric lattices. *Nuclear Physics B* **205**, 285–300. [The foundational paper on anisotropic Wilson lattice gauge theory with separate spatial and temporal couplings; referenced in Appendix B for the Euclidean-invariance restoration via fixed-point attraction.]

Maxwell, J. C. (1865). A dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society* **155**, 459–512.

Polyakov, A. M. (1977). Quark confinement and topology of gauge theories. *Nuclear Physics B* **120**, 429–458. [The strong-coupling phase of compact U(1) lattice gauge theory exhibiting confinement in 2+1 dimensions, contrasted with the weak-coupling Maxwell phase required by the present construction (§16.7 caveats).]

Wald, R. M. (1984). *General Relativity*. University of Chicago Press. [Chapter 4 contains the canonical Lagrangian-based uniqueness argument for Maxwell theory.]

Weinberg, S. (1995). *The Quantum Theory of Fields, Vol. 1: Foundations*. Cambridge University Press. [Chapters 5 and 8: massless spin-1 particles and the Lorentz-invariance route to Maxwell uniqueness.]

Wilson, K. G. (1974). Confinement of quarks. *Physical Review D* **10**, 2445–2459. [Foundational lattice gauge theory paper, cited as existence proof for the admissibility class in open problem 5.]

VERSF corpus (internal references)

The present paper refers throughout to companion papers within the VERSF Theoretical Physics Programme — including the foundational papers establishing axioms (A1) and (A2), the emergent-Lorentz companion paper inherited by (B4), the $K = 7$ constraint-dimensionality paper invoked in §10.2 / §12.3 / §14.1 / §16.8 / OP 5 / OP 6, the measurement and fact-formation papers bridged in §3.1, and several derivational papers underlying the $K = 7$ paper's individual constraints. Rather than list these here with descriptive titles that may differ from their published or in-preparation forms, we direct readers to the **corpus index at versf-eos.com** for the authoritative list of titles, identifiers, and current status. Within the body of this paper, companion work is referred to by function (e.g., "the $K = 7$ paper", "the measurement papers") rather than by speculative title.