An Open Letter to CERN: On a Continuity Law for Photonic Propagation Across Dimensional Extensions of Spacetime

Peter De Ceuster[*]

[*]Dimensional Electrodynamics Working Note: Toward a Plausible Unobserved Photonic Law Coupling 4D Electromagnetism to Higher-Dimensional Structure, (peterdeceuster5@gmail.com)

October 27, 2025

Abstract

This document is written as an open letter to CERN physicists and associated experimental and theoretical groups. Its purpose is to request focused attention on what appears to be a physically admissible, symmetry-compatible, and experimentally testable extension of Maxwell electrodynamics. In brief: if photons in 3+1 dimensions (X) are allowed to couple to, or partially propagate through, compact extra directions in a higher-dimensional ambient manifold (Y), then standard Maxwell theory on X should be deformed in a specific, derivable way. That deformation implies a conserved cross-dimensional charge and predicts a nonzero floor in interferometric visibility. The claim here is that such a law is plausible, not exotic, and that it has not yet been explicitly instrumented for.

What follows is a technical statement of that mechanism, together with a direct request that CERN evaluate its theoretical consistency and experimental accessibility.

The discussion is organized in five points.

(1) Distinct dimensional sectors of nature need not share geometric or dynamical structure. Any program that treats observed 3+1-dimensional spacetime as an effective slice of a higher-dimensional geometry must accept that different sectors of the full theory can be stabilized by backgrounds of radically different dimensional profile. We denote by X the observed Lorentzian 3+1 manifold, and by Y a $(4+k)$-dimensional bulk with $k \geq 1$ compact (or effectively compact) directions.

Empirically and theoretically, X exhibits locality, Lorentz symmetry, and $U(1)$ gauge structure that are extremely rigid and have been measured in detail. By contrast, generic candidates for Y — e.g. warped throats, internal fibered components, nontrivial compactification manifolds, higher-dimensional strata with nontrivial holonomy, etc. — are anisotropic, topologically nontrivial, and typically support additional field content that does not descend to light degrees of freedom in X. The local excitation spectrum, available homology, and causal texture in Y are not expected to be isomorphic to those of X.

It is therefore natural to model X not as “the whole theory,” but as a boundary-type or brane-type sector: comparatively spectrally clean, aggressively renormalized, and symmetry-constrained. Y should be regarded as a higher-capacity reservoir with additional directions, additional allowed modes, and looser local structure. In formal terms, we assume an asymmetry of structure between X and Y.

A convenient language for making that asymmetry mathematically explicit is the language of Grothendieck-style algebraic geometry: sheaves, derived functors, and pushforwards. One can regard physically relevant fields in Y (gauge fields, scalar sectors, Higgs-like condensates) as living in appropriate sheaf-theoretic objects on Y. The projection

$$\pi : Y \longrightarrow X$$

then induces pullbacks and pushforwards between the corresponding derived categories. In particular:

- The pullback π^* takes a sheaf/bundle on X and promotes it to Y. Physically: lift data from the observed 4D sector into the higher-dimensional bulk.
- The derived pushforward $R^\bullet \pi_*$ takes structured data on Y and induces effective data back on X. Physically: integrate out, or compress, bulk structure back down to the observed sector.

In this formalism, fields defined globally or locally on Y do not simply restrict to X in the naive sense. Instead, their effective imprint on X is given by an object like $R^\bullet \pi_*(\cdot)$, which in general carries the memory of topological obstructions, cohomology classes, compactification scales, and so on. This already encodes the geometric asymmetry: X and Y are not treated as equivalent stages of one manifold, but as related by a functorial interface that remembers higher-dimensional structure in compressed form. The essential point is that X may inherit structured, conserved data from Y without itself housing the full higher-dimensional degrees of freedom.

That interface — sheaf-theoretic lifting via π^* and induced data via $R^\bullet \pi_*$ — is the backbone of what follows.

(2) The photon, as observed in X, appears not to undergo species redefinition across environments. Within X, the photon is an exceptionally stable gauge excitation: it is massless to current bounds, it carries a $U(1)$ gauge charge, and it satisfies Maxwell equations with extremely high precision. No terrestrial observation suggests that the photon bifurcates into distinct environment-dependent species, nor that it acquires different effective identities in different regions of ordinary spacetime.

This motivates the following inference. If the photon can, in any sense, probe or partially occupy directions in Y, then there must exist a continuity principle constraining that process. Equivalently: even if X and Y are not dynamically equivalent sectors, electromagnetic flux must still satisfy a transport law that is respected across their interface. The observed absence of photon "relabeling" in X is empirical pressure for such a law. Standard 3+1 Maxwell theory, in its usual form, does not explicitly encode any such cross-dimensional continuity requirement.

In categorical terms: P, the photon sheaf/bundle on X, appears to behave as a single object under evolution, not as a direct sum of context-specific variants. This strongly suggests a functorial constraint on how P may be lifted to Y and then pushed back down to X.

(3) A concrete formalization: an extended Maxwell system with a bulk-induced current. Let P denote the sheaf (or bundle) encoding photon gauge data on X, and let H denote a Higgs-like sector defined (possibly nontrivially) on Y. We assume the existence of a morphism in the bulk,

$$\eta : \quad \pi^* P \otimes H \longrightarrow \mathcal{O}_Y[\ell], \tag{1}$$

where $\mathcal{O}_Y[\ell]$ denotes an appropriate twisting of the structure sheaf on Y (e.g. with degree/weight ℓ). Physically: η is a permissible channel through which lifted photon data $\pi^* P$ can couple to bulk scalar structure H.

We then define an induced 3-form current on X via the derived pushforward:

$$J_s \equiv R^\bullet \pi_* \Big[\eta(\pi^* P \otimes H) \Big] \in \Omega^3(X). \tag{2}$$

Interpretation:

- $J_{\rm s}$ is *not* introduced arbitrarily. It is the shadow, on X, of allowed interactions in Y between the photon sector and a Higgs-like sector.
- If Y introduces no additional effective channels (for example, $k = 0$, or η is trivial in cohomology), then $J_{\rm s} = 0$.
- If Y does provide a nontrivial channel, then $J_{\rm s}$ can be nonzero and encodes a leakage / throughput / transport correction that is invisible to purely 3+1 Maxwell theory.

On X we now propose the following deformation of Maxwell's equations:

$$\mathrm{d}F \;=\; 0, \qquad \mathrm{d}\big(\star F \;-\; J_{\rm s}\big) \;=\; 0, \tag{3}$$

where A is the $U(1)$ gauge potential on X, $F = \mathrm{d}A$ is the field strength, and $\star$ denotes the Hodge dual on X.

Key remarks:

- $\mathrm{d}F = 0$ is unchanged. We preserve the Bianchi identity.
- If $J_{\rm s} = 0$, we recover the standard source-free Maxwell equation $\mathrm{d}(\star F) = 0$.
- If $J_{\rm s} \neq 0$, the conserved object is not simply $\star F$, but rather $\star F - J_{\rm s}$. This states that electromagnetic flux in X is not closed by itself, but becomes closed once the contribution inherited from Y is included.

The system (3) is not imposed by hand. It can be obtained from an action functional on X of the schematic form

$$S[A; J_{\rm s}] \;=\; \int_X \left[\frac{1}{2}F \wedge \star F \;-\; A \wedge \mathrm{d}J_{\rm s}\right] \;+\; S_{\rm bulk}\big[H, g_Y, \mathcal{M}_Y\big]\,, \tag{4}$$

which yields $\mathrm{d}(\star F - J_{\rm s}) = 0$ upon variation with respect to A. Gauge invariance under $A \mapsto A + \mathrm{d}\lambda$ is maintained provided $J_{\rm s}$ is itself gauge-invariant. That in turn is enforced because $J_{\rm s}$ is defined functorially, through (2), from bulk data, not inserted arbitrarily into X.

Two immediate theoretical consequences follow.

(i) A conserved cross-dimensional charge. Define, for any closed 3-surface $\Sigma_3 \subset X$,

$$Q_{\rm s} \;=\; \int_{\Sigma_3} (\star F \;-\; J_{\rm s})\,. \tag{5}$$

Because $\mathrm{d}(\star F - J_{\rm s}) = 0$, the quantity $Q_{\rm s}$ is conserved. This is a genuine continuity law. $Q_{\rm s}$ measures a corrected electromagnetic flux: ordinary $\star F$ minus the bulk-induced contribution $J_{\rm s}$. In physical terms, $Q_{\rm s}$ can be read as the invariant that remains constant when a photon sector that "touches" higher-dimensional structure is projected back down and observed purely in X.

This provides a candidate for a cross-dimensional conservation principle: electromagnetic transport is not purely confined to X, but it is still globally conserved once the $Y \to X$ interface is taken into account.

(ii) A lower bound on optical interference visibility. If $J_{\rm s} \neq 0$, then not all phase information encoded in F remains localized in X. Even in the limit of arbitrarily well-controlled instrumentation, interference fringe visibility V in a sufficiently sensitive photon interferometer should saturate at

$$1 - V \;\sim\; \big\|\omega(\eta)\big\|^2\,\Lambda^{-2k}, \tag{6}$$

where $\omega(\eta)$ is the obstruction class associated with the bulk morphism η, and Λ denotes the characteristic compactification scale of the extra k dimensions. The scaling $1 - V \propto \Lambda^{-2k}$ captures the expectation that when the extra dimensions are more highly compactified (larger effective energy scale Λ), the leakage channel into Y is further suppressed, and the deviation from perfect visibility becomes correspondingly smaller but does not vanish identically for any finite Λ.

This is important experimentally: it asserts the existence of a nonzero, geometry-driven floor on interferometric visibility, not attributable to conventional environmental noise, optics tolerances, or detector inefficiency. In principle, this is falsifiable.

(4) Request to CERN. This open letter is making two requests.

- *Theoretical analysis.* We ask that the deformation in Eq. (3) be examined for consistency with Lorentz covariance on X, gauge invariance, unitarity, and anomaly cancellation; that possible quantization of Q_s be assessed (in analogy to Dirac-type arguments); and that the UV admissibility of J_s within known high-energy frameworks be evaluated.
- *Experimental feasibility.* We ask whether interferometric platforms under CERN-associated or partner instrumentation groups can, in principle, be operated in a regime where an irreducible floor in fringe visibility at order Λ^{-2k} would be observable, or else tightly bounded.

The mechanism above is specific enough to be interrogated with established field-theoretic tools. It also directly suggests an optical observable. The request is that both lines of inquiry be treated as active physics questions rather than deferred curiosities.

(5) Conclusion and forward work. Three premises are advanced:

1. The effective geometries of X and Y differ in topology, causal microstructure, excitation content, and renormalization behavior. They should not be treated as trivially equivalent sectors of a single homogeneous spacetime.
2. The photon in X behaves empirically as a single stable gauge species. This suggests the existence of a cross-dimensional continuity principle for electromagnetic flux, rather than ad hoc environment-specific photon identities.
3. A mathematically well-defined deformation of Maxwell theory, summarized in Eq. (3), produces such a continuity principle. It predicts a conserved quantity Q_s and implies a small but nonzero, geometry-controlled lower bound on interferometric visibility.

If these premises hold simultaneously, then the proposed law should be regarded not as speculation but as an unmeasured element of electrodynamics at the interface between X and Y. The sheaf-theoretic / derived-functor machinery is not used here as ornament; it is the natural language for transporting field data along $\pi : Y \to X$ and for exhibiting how Y can leave a conserved trace in X without fully revealing its own degrees of freedom.

A more extended treatment — including explicit toy compactification models for Y, construction of $\omega(\eta)$ in concrete examples, quantitative discussion of anomaly freedom and unitarity, and interferometer sensitivity estimates required to probe $1 - V \sim \Lambda^{-2k}$ — is intended for submission to the arXiv or Zenodo in 2026, potentially in collaboration with a co-author. That longer document will include full derivations in the Grothendieck-style language described above: explicit definitions of P and H as objects in the appropriate derived category on Y, explicit computation of $R^\bullet\pi_*$, and explicit recovery of J_s in X.

We ask CERN to recognize this continuity law as an experimentally addressable question in photon physics, and to treat the associated interferometric floor not as a curiosity but as a target. The analysis presented here reflects the limits of what can be advanced by a single investigator operating under severe time constraints; a complete treatment now requires institutional theoretical review and dedicated interferometric effort, which this letter formally requests CERN to initiate.