

Curvature Electromagnetism: Deriving Maxwell from Geometry

Kim Seung-il

Independent Researcher, Republic of Korea
(seungilkim851@gmail.com)

October 2025

Abstract

This paper extends two constraints introduced in our earlier works—*Introducing the Curvature Field Function: Toward a Geometric Formulation of Wavefunction Collapse* [32] and *Curvature Field Formulation of Gravity: Toward a Physical Reconstruction of Space-time* [33]—to electromagnetism: a small-gradient condition $|\nabla\Phi| < \varepsilon$ and a scalar-field dynamics $\square\Phi - U'(\Phi) = J$. The idea is simple: a curvature field Φ may weakly tune the vacuum EM response, yet in the weak-gradient regime electromagnetism *continuously reduces* to Maxwell. At leading order this yields two observables—an *isotropic* (impedance-like) tweak governed by the mean of Φ , and a very small *anisotropic* component proportional to $|\nabla\Phi|$ with a characteristic angular dependence. We confront these reductions with public, peer-reviewed datasets. In angle-dependent Aharonov–Bohm/quantum-oscillation measurements [18, 19, 20, 21], the standard tilt scaling R lies within $1 \sim 3\%$ of unity across platforms; including conservative digitization error gives a **95% C.I.** $|R - 1| < 0.55\%$ (agreement $\gtrsim 99.45\%$). Independent resonator literature on TE/TM mode ratios [24, 25, 26, 27, 22, 23] supports long-term stability at $|\Delta\rho/\rho| < 3 \times 10^{-3}$, enabling sub-percent separation of isotropic and anisotropic bounds; jointly, these imply $G \lesssim 5 \times 10^{-3}$. Crucially, identifying the frame-induced connection A_μ and its curvature $F = dA$ with the electromagnetic potential/field is *not a new interaction* but a *geometric reparametrization*; any empirical novelty resides solely in the constitutive extension $H = \chi(\Phi, \nabla\Phi) : F$, which *vanishes continuously* as $\chi \rightarrow \chi_0$ in the weak-gradient window (see [Section 2.5](#)). In short, the curvature-field language that organized the quantum (micro) and gravitational (macro) regimes reaches electromagnetism without strain: most situations reduce cleanly to Maxwell and any residual deviations are very small. Our next step targets this sub-percent window with precision angle sweeps and mode-ratio tracking, to test whether a single curvature field can serve as a common constitutive principle threading quantum, gravity, and electromagnetism.

1 Introduction

From curvature to electromagnetism: historical context and motivation. Electricity and magnetism, discovered experimentally by Galvani and Faraday, were unified into a single field by Maxwell’s equations. In the 20th century this field was rewritten in the language of differential forms and connections: a connection A and its curvature $F = dA$ came to be viewed as the core of electromagnetic phenomena [16]. In this view, electromagnetism is given as

geometry first, and measurable quantities arise *as consequences* of that geometry. Pushing this classical line one step further, we ask whether an additional scalar degree of freedom—a “curvature field” Φ —can play a meaningful role in electromagnetism, and whether that role *reduces continuously* to standard Maxwell theory in the appropriate limit.

Prior work and continuity. Two earlier papers—*Introducing the Curvature Field Function* [32] and *Curvature Field Formulation of Gravity* [33]—argued that a single curvature field Φ can serve as a common organizing principle for physical quantities across quantum (microscopic) and gravitational (macroscopic) regimes. The essence is twofold. (i) In *weak-gradient* settings ($|\nabla\Phi| < \varepsilon$), the theory must reduce smoothly to the standard one; (ii) Φ interacts with other fields *locally* on top of its own dynamics (source/potential). Building on this continuity, the present paper applies the same philosophy to the electromagnetic sector.

Curvature electromagnetism: an intuitive sketch. Curvature electromagnetism can be summarized in a single sentence: *the “properties” of vacuum (a constitutive tensor) depend weakly on Φ , so the way it accepts an electromagnetic field F — encoded as $H = \chi(\Phi, \nabla\Phi) : F$ — is minutely modified.* Here χ may be viewed as an *effective* constitutive response analogous to permittivity/permeability. The mean value $\langle\Phi\rangle$ produces an *isotropic* (impedance-like) fine adjustment, while a small gradient $|\nabla\Phi|$ yields a very weak *anisotropy* (e.g., of $\cos 2\theta$ type). Crucially, in the limit $|\nabla\Phi| \rightarrow 0$ one has $\chi \rightarrow \chi_0$ and the standard Maxwell–Hodge duality $H = \star F$ is *recovered intact*. Thus curvature electromagnetism does not replace Maxwell; rather, it is an extension that tracks tiny constitutive variations *at the edge of Maxwell*.

What the framework aims to explain. This viewpoint organizes three layers of electromagnetic phenomena at once. First, the *geometric origin*: while the connection A and curvature F remain the essence of the field, the vacuum’s supporting response may be weakly tuned by Φ . Second, the *reduction principle*: in weak-gradient regimes, all observables must *continuously* match Maxwell’s predictions, with any residual effects confined to sub-percent *corrections*. Third, the *observational bridge*: the isotropic piece couples naturally to *scalar* observables such as impedance or resonant frequency shifts, whereas the anisotropic piece connects to normalized comparators that flip a field *direction* (tilted-field tests, TE/TM mode ratios, etc.). In this way, the theory tells us *what to measure*, and the data answer *how precisely* those instructions are satisfied.

Position and contributions of this paper. We (a) combine the curvature field and electromagnetism through a *local* constitutive law, (b) secure a *continuous reduction to Maxwell* under $|\nabla\Phi| < \varepsilon$, and (c) place *quantitative bounds* on isotropic/anisotropic components by comparing *real-world datasets* from disparate platforms on a common scale. In particular, by coupling angle-sweep normalization indicators from Aharonov–Bohm/quantum-oscillation studies with long-term stability records of resonator TE/TM mode ratios, we confirm that in most situations electromagnetism reduces cleanly to Maxwell and we confine the remaining sub-percent window with *explicit numbers*. The method is simple and transparent: we introduce no new global assumptions, follow the measurements singled out by geometry, record reduction when they agree, and tighten upper bounds when they do not. In the process, Φ extends to electromagnetism the same vocabulary that linked quantum and gravity, and we make clear—by numbers—how tightly the standard theory is joined at this interface.

1. Geometric foundations: from curvature frame to $U(1)$ connection

1.0 Scope, symbols, and units. The aim is to organize the directional change encoded by a scalar curvature field $\Phi(x)$ as a $U(1)$ phase and, from it, to obtain a connection A_μ and curvature $F_{\mu\nu}$ with the fewest assumptions. The fixed conventions are:

- Metric signature $(-, +, +, +)$, Heaviside–Lorentz units, $c = 1$; indices are moved with $g_{\mu\nu}$.
- Coordinates $x^\mu = (t, x^i)$, partial derivatives ∂_μ , Levi–Civita covariant derivatives ∇_μ , d’Alembertian $\square = g^{\mu\nu}\nabla_\mu\nabla_\nu$.
- Exterior calculus: exterior derivative d , wedge $\wedge$, Hodge dual $*$; Levi–Civita symbols $\epsilon^{0123} = +1$, $\epsilon^{123} = +1$.
- Electromagnetism: potential A_μ , curvature $F = dA$ with components $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$; excitation $H = \chi : F$ (constitutive tensor χ), reducing to $H = \chi_0 : F$ in the weak–gradient limit [16].
- Lorentz invariants $I_1 = \frac{1}{2}F_{\mu\nu}F^{\mu\nu} = \mathbf{B}^2 - \mathbf{E}^2$ and $I_2 = \frac{1}{2}F_{\mu\nu}*F^{\mu\nu} = \mathbf{E} \cdot \mathbf{B}$ [2].
- Curvature field $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ obeys

$$\square\Phi - U'(\Phi) = J, \quad |\nabla\Phi| < \varepsilon,$$

and the constitutive law has the continuous reduction

$$H = \chi(\Phi, \nabla\Phi) : F \longrightarrow H = \chi_0 : F \quad (|\nabla\Phi| \rightarrow 0)$$

as set out in [32, 33].

Sign, units, and dimensional bookkeeping are summarized in [Appendix A](#); bundle structure and quantization appear in [Appendix B](#); conservative discretization and Hodge weighting are collected in [Appendix C](#). Definitions here feed directly into [Section 2](#), [Section 4](#), and [Section 5](#).

1.1 Curvature field Φ , Hessian, and principal frame

1.1.1 Basic derivatives and background curvature. For the scalar Φ ,

$$\nabla_\mu\Phi, \quad H_{\mu\nu} \equiv \nabla_\mu\nabla_\nu\Phi,$$

and since $[\nabla_\mu, \nabla_\nu]\Phi = 0$, the Hessian $H_{\mu\nu}$ is symmetric. When acting on the vector $\nabla\Phi$, commutators expose the ambient Riemann/Ricci curvature:

$$[\nabla_\mu, \nabla_\nu](\nabla_\rho\Phi) = R_{\rho\sigma\mu\nu}\nabla^\sigma\Phi, \quad (1)$$

$$\nabla_\mu\square\Phi - \square(\nabla_\mu\Phi) = R_\mu{}^\nu\nabla_\nu\Phi. \quad (2)$$

These identities quantify how the direction field of Φ twists in a curved background. In the regime $|\nabla\Phi| < \varepsilon$, the later constitutive response χ is arranged so that only weak, controllable changes remain and the Maxwell limit is continuous (see [Section 2](#)).

1.1.2 Isotropic/trace-free split and the principal plane $\Pi(x)$. To disentangle direction from scale, define the trace-free shear

$$S_{\mu\nu} \equiv H_{\mu\nu} - \frac{1}{4}g_{\mu\nu} \square\Phi.$$

On a spatial slice, the symmetric tensor H_{ij} admits an orthonormal eigenbasis $\{e_{(a)}^i\}$ with eigenvalues $\{\lambda_a\}$. Typically, the two axes with largest $|\lambda_a|$ span the dominant bending directions. This two-plane is the *principal plane* $\Pi(x)$; its slow rotation is what will be encoded as a $U(1)$ phase below. Global issues (chart transitions, spin lift, branch cuts) are treated in [Appendix B](#).

1.1.3 From frame to phase. Choose a unit complex section $u(x)$ that tracks the rotation of $\Pi(x)$ and define the $U(1)$ connection

$$A_\mu \equiv \text{Im} \frac{u^\dagger \nabla_\mu u}{u^\dagger u}.$$

Under a phase change $u \rightarrow e^{i\chi}u$, $A_\mu \rightarrow A_\mu + \partial_\mu \chi$, i.e., the gauge-potential transformation law. The curvature is

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

so $dF = 0$ holds identically, and for any loop C with spanning surface $S(C)$,

$$\oint_C A = \iint_{S(C)} F.$$

This boundary-bulk relation anchors the phase observables in [Section 4](#) and the data indicators in [Section 6](#).

1.1.4 One-line implementation note. In the continuum the chain is $\Phi \rightarrow \Pi \rightarrow u \rightarrow A \rightarrow F$. On meshes, keep the incidence maps (the d operator) metric-free, insert $\chi(\Phi, \nabla\Phi)$ only in Hodge weights, and preserve exact discrete continuity; details and checks are compiled in [Appendix C](#).

1.2 Definition of the Berry-like connection A_μ

Frame-induced complex section and connection. Let $u(x)$ be a unit complex section that tracks the rotation of the principal plane $\Pi(x)$. Define the $U(1)$ connection by

$$A_\mu \equiv \text{Im} \frac{u^\dagger \nabla_\mu u}{u^\dagger u}. \quad (3)$$

Under a phase redefinition $u \rightarrow e^{i\chi}u$ one has $A_\mu \rightarrow A_\mu + \partial_\mu \chi$, so A_μ behaves as a gauge potential. For a single-phase choice $u = e^{i\theta}$, $A_\mu = \partial_\mu \theta$; thus nontrivial curvature arises only from frame anholonomy or multivalued phases [4].

Curvature and holonomy.

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu = (dA)_{\mu\nu}, \quad \gamma_C = \oint_C A_\mu dx^\mu = \iint_{S(C)} F, \quad (4)$$

so the line integral along a closed path is read as flux (Stokes). Reversing the path gives $\gamma_{C^{-1}} = -\gamma_C$ (orientation odd), which is useful at the measurement stage for canceling even (dynamical) contributions ([Section 4](#)). In the presence of defects (caustics) or a spin lift, $\gamma_C/2\pi$ can be quantized [6].

Equivalence vs. novelty (policy).

On any simply connected chart $U \subset \mathcal{U}$, there exists a gauge λ such that $A_\mu \sim A_\mu^{(\text{EM})} + \partial_\mu \lambda$ and $F = dA$ matches the electromagnetic curvature by definition. Hence $A \mapsto F$ here is a **reconstruction** of standard electromagnetism, not an additional dynamical hypothesis. The only potential deviation lives in the constitutive map $H = \chi(\Phi, \nabla\Phi) : F$; in the weak-gradient limit $\chi \rightarrow \chi_0$, this deviation vanishes and Maxwell is recovered continuously (Section 2.5).

Dictionary to standard electromagnetism. The chain $\Phi \rightarrow \Pi \rightarrow u \rightarrow A \rightarrow F$ corresponds to the standard gauge potential $A_\mu^{(\text{EM})}$ as follows.

1. **Local gauge equivalence:** Under a phase redefinition of the section $u \rightarrow e^{i\chi}u$, the connection transforms as $A_\mu \rightarrow A_\mu + \partial_\mu \chi$. Hence $F = dA$ is *locally* the same two-form as the electromagnetic curvature $F_{\mu\nu}^{(\text{EM})}$.
2. **Continuity (Maxwell) limit:** In a weak-gradient window $|\nabla\Phi| < \varepsilon$, the constitutive law reduces as $\chi(\Phi, \nabla\Phi) \rightarrow \chi_0$, so that $H \rightarrow F$ and the standard Maxwell–Hodge duality is recovered (see Section 2.5).
3. **Physical interpretation:** A_μ is a Berry-type connection induced by the *frame rotation* of the principal plane Π fixed by the Hessian of Φ . When coupled to matter via $D_\mu = \nabla_\mu - iqA_\mu$, this connection acts as a *gauge-equivalent* effective potential to the usual $A_\mu^{(\text{EM})}$ in local experiments.
4. **Global issues:** If Π has global twist/defects, A can be multi-valued; observables are then $\oint A$ and F . Path inversion flips the odd component (Section 4.1).

In short, the chain is a *geometric reparametrization*; in the weak-gradient limit it yields the same observables $(E, B; F)$ as standard Maxwell theory.

Domain, globality, and chart transitions. We assume $\Phi \in C^2(\mathcal{M} \setminus \mathcal{S}_{\text{cau}})$ so that the Hessian-defined principal plane $\Pi(x)$ is smooth on $\mathcal{U} \equiv \mathcal{M} \setminus \mathcal{S}_{\text{cau}}$. All holonomy/flux observables are taken on loops and surfaces contained in $\mathcal{U}$, with the global (Čech) structure recorded in Appendix B. Let $\mathcal{S}_{\text{cau}}$ be the set where Hessian eigenvalue crossings or degeneracies occur, and define $\mathcal{U} = \mathcal{M} \setminus \mathcal{S}_{\text{cau}}$. Choose local sections u_a on charts $U_a \subset \mathcal{U}$. On overlaps $U_a \cap U_b$,

$$u_b = s_{ab} e^{i\chi_{ab}} u_a, \quad s_{ab} \in \{\pm 1\}, \quad (5)$$

which induces

$$A^{(b)} = A^{(a)} + d\chi_{ab}, \quad F^{(b)} = F^{(a)}. \quad (6)$$

On triple overlaps one has $\chi_{ab} + \chi_{bc} + \chi_{ca} = 2\pi n_{abc}$, defining an integer Čech 2-cocycle. Hence for any closed two-surface $\Sigma \subset \mathcal{U}$,

$$\frac{1}{2\pi} \int_\Sigma F \in \mathbb{Z}, \quad (7)$$

i.e., a quantized first Chern number (Appendix B).

Sign and spin lift. The structure group of the principal plane is $\text{SO}(2)$, with double cover $\text{Spin}(2) \simeq U(1)$. The sign s_{ab} in (5) corresponds to a π -phase choice; being constant, it vanishes under differentiation and does not affect F (see Appendix B for details).

Contractible loops and the role of singular sets. If $C = \partial S$ is contractible within $\mathcal{U}$, then $\oint_C A = \iint_S F = 0$ by local smoothness of F . If C links $\mathcal{S}_{\text{cau}}$, $\iint_S F$ can contribute integer multiples of 2π via (7). This boundary–bulk match, together with the Bianchi identity in Section 1.3, underlies the definition of a measurable geometric phase.

Implementation note (discretization and stability). On meshes, keep the exterior derivative (incidence maps) metric-free and insert the constitutive response $\chi(\Phi, \nabla\Phi)$ only in Hodge weights (volumes/areas). This preserves exact discrete continuity and stabilizes odd/even separation under path reversal. Concrete weight definitions and conservation checks appear in Appendix C.

One-line link to observables. The A derived here and its loop integral $\oint A$ interface directly with normalized angle-dependence metrics and $\cos 2\theta$ -type anisotropy parsing in data analysis (Section 4, Section 6).

1.3 Curvature 2-form and the Bianchi identity

Definition and immediate consequences. With the connection A and curvature $F \equiv dA$,

$$dF = 0 \quad \Longleftrightarrow \quad \partial_{[\lambda} F_{\mu\nu]} = 0, \quad (8)$$

i.e., the Bianchi identity holds. Using the space–time split (sign/Hodge conventions in Appendix A),

$$E_i \equiv F_{0i}, \quad B^i \equiv \frac{1}{2} \epsilon^{ijk} F_{jk},$$

so that

$$\nabla \cdot \mathbf{B} = 0, \quad \partial_t \mathbf{B} + \nabla \times \mathbf{E} = 0 \quad (9)$$

follow immediately. The remaining Maxwell pair ($\nabla \cdot \mathbf{E} = \rho$, $\nabla \times \mathbf{B} - \partial_t \mathbf{E} = \mathbf{J}$) will be obtained by variation of the action in Section 2.

Integral form and boundary–bulk match. For any oriented surface S with boundary $C = \partial S$, Stokes’ theorem gives

$$\oint_C A = \iint_S F, \quad (10)$$

and reversing the path yields $\oint_{C^{-1}} A = -\oint_C A$, making the *orientation-odd* nature explicit. Equation (10) is the reference for comparing phase (line integral) and flux (surface integral) on a common footing later on (Section 4, Section 6).

Domain and treatment of singular sets (caustics). Let $\mathcal{S}_{\text{cau}}$ denote loci where the Hessian eigenstructure degenerates or crosses, and define the working domain $\mathcal{U} \equiv \mathcal{M} \setminus \mathcal{S}_{\text{cau}}$ (Section 1.2). On $\mathcal{U}$ the fields A, F are smooth, and $\oint_{\partial S} A = \iint_S F$ applies as is. If S links $\mathcal{S}_{\text{cau}}$, $\iint_S F$ can acquire quantized contributions, and

$$\frac{1}{2\pi} \int_{\Sigma} F \in \mathbb{Z} \quad (11)$$

holds on any closed two–surface $\Sigma \subset \mathcal{U}$; see bundle/Čech cocycle details in Appendix B.

Hodge dual and invariant pairing. With the Hodge dual ${}^*F_{\mu\nu} \equiv \frac{1}{2} \epsilon_{\mu\nu}{}^{\rho\sigma} F_{\rho\sigma}$,

$$dF = 0 \quad \Longleftrightarrow \quad \nabla_\mu {}^*F^{\mu\nu} = 0$$

(in vacuum). This aligns with the Lorentz invariants in [Section 1.4](#), $I_1 = \frac{1}{2} F_{\mu\nu} F^{\mu\nu}$ and $I_2 = \frac{1}{2} F_{\mu\nu} {}^*F^{\mu\nu}$. In particular, the P/T -odd nature of I_2 resonates with the sign flip of the line integral under path reversal in (10), aiding the removal of even (dynamical) contributions.

Gauge independence and locality. Since F is invariant under $A \mapsto A + d\chi$, relations (19)–(10) are gauge independent. Moreover, $dF = 0$ is a *local identity*—true by definition—hence unaffected by dynamics or matter content. This remains intact when the constitutive law $H = \chi(\Phi, \nabla\Phi) : F$ is introduced ([Section 2](#), [Appendix C](#)).

Notes on computation and implementation. On a lattice, keep incidence maps metric-free (pure exterior calculus) and place $\chi(\Phi, \nabla\Phi)$ solely in Hodge weights (cell volumes/areas). Then the discrete $d^2 = 0$ structure preserves the Bianchi identity exactly in its discrete form. Concrete definitions of the weights and conservation checks are summarized in [Appendix C](#).

1.4 Lorentz invariants I_1, I_2 and P/T

Definitions (tensor form). Conventions for signs and the Hodge dual follow [Appendix A](#). From the electromagnetic curvature 2-form $F_{\mu\nu}$ and its Hodge dual ${}^*F_{\mu\nu} \equiv \frac{1}{2} \epsilon_{\mu\nu}{}^{\rho\sigma} F_{\rho\sigma}$, define the two Lorentz invariants

$$I_1 \equiv \frac{1}{2} F_{\mu\nu} F^{\mu\nu} = \mathbf{B}^2 - \mathbf{E}^2, \quad I_2 \equiv \frac{1}{2} F_{\mu\nu} {}^*F^{\mu\nu} = \mathbf{E} \cdot \mathbf{B}. \quad (12)$$

Here I_1 is a Lorentz *scalar* and I_2 a Lorentz *pseudoscalar*.

Discrete symmetries (P, T). Under spatial parity P : $\mathbf{E} \rightarrow -\mathbf{E}$, $\mathbf{B} \rightarrow +\mathbf{B}$, hence $I_1 \xrightarrow{P} I_1$ and $I_2 \xrightarrow{P} -I_2$. Under time reversal T : $\mathbf{E} \rightarrow +\mathbf{E}$, $\mathbf{B} \rightarrow -\mathbf{B}$, hence $I_1 \xrightarrow{T} I_1$ and $I_2 \xrightarrow{T} -I_2$. Thus I_2 is odd under both P and T , aligning naturally with orientation-odd line-integral phases (see [Section 1.3](#), eq. (10)).

Duality rotations and (anti)self-dual split. Introduce the chiral combinations $F_\pm \equiv \frac{1}{2}(F \pm i {}^*F)$, which obey ${}^*F_\pm = \mp i F_\pm$. A continuous duality rotation $F \mapsto F \cos \alpha + {}^*F \sin \alpha$ acts as $F_\pm \mapsto e^{\mp i\alpha} F_\pm$. The invariants combine as

$$I_1 + iI_2 = -2 F_{+\mu\nu} F_+^{\mu\nu}, \quad I_1 - iI_2 = -2 F_{-\mu\nu} F_-^{\mu\nu}, \quad (13)$$

so $I_1 = I_2 = 0$ iff both chiral parts are null. This structure is central when assessing permissible couplings and duality constraints [2].

Role in the action and constitutive law. The Maxwell Lagrangian reads $\mathcal{L}_F = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} = -\frac{1}{2} I_1$. Because I_2 is P/T -odd, it diagnoses sensitivity to parity-odd terms (e.g. a $\theta F \wedge F$ density). In the present framework the constitutive relation $H = \chi(\Phi, \nabla\Phi) : F$ tends to $\chi \rightarrow \chi_0$ as $|\nabla\Phi| \rightarrow 0$, thereby returning the standard Maxwell action controlled by I_1 ([Section 2](#), [Appendix C](#)).

Local field types and links to observables. If $I_1 > 0$ the field is *magnetic-like*; if $I_1 < 0$, *electric-like*; and if $I_1 = I_2 = 0$, *null*. In our empirical design, I_1 aligns with isotropic (impedance-type) scalar indicators, whereas I_2 pairs with P -sensitive procedures such as path reversal or polarization-axis swaps. Below we show how the angle-normalized ratio R and the TE/TM mode ratio ρ separate even/odd content (Section 5, Section 6).

Numerical note. On a mesh, enforce the chain condition $d^2 = 0$ and place metric/constitutive information ($\chi(\Phi, \nabla\Phi)$) solely in Hodge weights. This preserves gauge invariance of the discrete I_1, I_2 and ensures convergence to eq. (12) in the continuum limit (Appendix C).

1.5 What is standard and what is new

Standard (reconstructed) pieces. The fiber-bundle $U(1)$ description, the path-area equivalence $\oint A = \iint F$ (Stokes), the Maxwell action $\mathcal{L}_F = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$, and minimal coupling $D_\mu = \nabla_\mu - iqA_\mu$ are established elements [16]. We retain this skeleton verbatim, make units/sign/Hodge choices explicit in Appendix A, and re-verify the variation-conservation chain (action $\rightarrow$ field equations $\rightarrow$ continuity) in Section 2. At the discrete level, only metric/material information (the constitutive tensor χ) enters Hodge weights, while boundary/incidence operators remain purely topological so that gauge/duality symmetries are preserved (Appendix C, Section 5).

Terminology. Throughout, we avoid the verb “derive Maxwell” from A and $F = dA$; this chain is a *geometric recast*. Claims of novelty, if any, concern only (i) the constitutive extension $H = \chi(\Phi, \nabla\Phi) : F$ under $|\nabla\Phi| < \varepsilon$, and (ii) its data-facing bounds (Section 2.5).

New proposal (core). (i) The scalar curvature field Φ produces a spatial principal plane $\Pi(x)$ via its Hessian, and the *rotation* of $\Pi(x)$ induces a Berry-like $U(1)$ connection A_μ (Section 1.2). In this reading, A_μ summarizes frame-phase transport; its curvature $F = dA$ is *locally gauge-equivalent* to the standard electromagnetic two-form, so the pair (A, F) is a *geometric reparametrization* of the usual variables, not a new interaction (Section 1.3).

(ii) With the constitutive law $H = \chi(\Phi, \nabla\Phi) : F$ under $|\nabla\Phi| < \varepsilon$, the response $\chi(\Phi, \nabla\Phi) \rightarrow \chi_0$ and the full Maxwell–Hodge duality are *recovered continuously* as $|\nabla\Phi| \rightarrow 0$. Consistency is checked at the levels of variation, conservation, and Lorentz invariants (I_1, I_2) (Section 2, Section 1.4).

(iii) The framework organizes observables into two complementary channels. An *isotropic* (impedance-like) effect, proportional to $\langle\Phi\rangle$, maps to scalar indicators such as resonance frequency; a weak *anisotropic* effect, proportional to $|\nabla\Phi|$, maps to angle-differencing and path-reversal normalized ratios. Accordingly we adopt the angle-normalized metric R and the TE/TM mode ratio ρ as complementary probes that separate isotropic/anisotropic components (Section 5).

(iv) *Pre-declared numerical goals.* Agreement in the angle-normalized channel at the level $|R - 1| < 10^{-2}$ and long-term mode-ratio stability $|\Delta\rho/\rho| < 3 \times 10^{-3}$ enable sub-percent bounds on the theory parameters A and G . If these goals are not met, we report calibrated upper bounds and raise sensitivity using symmetry separation, lock-in gradient modulation, and high-stability mode-ratio scans (Section 6, Appendix F).

(v) The boundary with the traditional framework is explicit. Φ is not an extra global hypothesis but a *conservative extension* that reproduces standard results as $|\nabla\Phi| \rightarrow 0$. When data agree, we record *equivalence to Maxwell* within the stated confidence intervals; when they deviate,

we attribute the discrepancy directly as a sub-percent correction in $\chi(\Phi, \nabla\Phi)$. This maintains a closed loop from definition $\rightarrow$ data $\rightarrow$ reduction or constraint.

Policy on theoretical priors. Beyond reconstruction, the admissible form and size of $\chi(\Phi, \nabla\Phi)$ are constrained by independent principles—gauge/Lorentz symmetry, energy positivity, and causal dispersion. We adopt these as hard priors when proposing, fitting, and bounding (α, η) (Section 2.2.1, Appendix D).

Uniqueness notice. Within the weak-gradient, linear window, our constitutive ansatz is *unique* under gauge/Lorentz symmetry, passivity/causality, and locality assumptions; see Section 2.2.2 and Appendix D.4.

1.6 Scope, limitations, and operational outlook

Where a signal *should not* appear. On a simply connected region where the principal plane $\Pi(x)$ does not wind, the connection A is locally pure gauge and thus $\oint_\gamma A = 0$. In that case the surface flux of $F = dA$ also vanishes, and loop-phase readouts are *supposed* to be null. Environments with strong phase mixing—multi-mode transport, non-adiabatic polarization, or broad bandwidth—further bury a holonomic (geometric) odd component beneath even, dynamical terms. Under such conditions, reporting *upper bounds* is appropriate (see Section 5 for numerical stability and Appendix G for metrology logs).

Conditions for a decisive readout. Single-mode transport, narrowband drive ($\Delta f/f_0 \ll 1$), adiabatic polarization, explicit parity extraction (path reversal $\gamma \mapsto \gamma^{-1}$ or order swap), and adequate SNR are required for $\oint_\gamma A = \iint_{S(\gamma)} F$ to act as an *instrument* rather than a mere identity [4]. A common decision rule is fixed across channels:

$$|\Phi_{\text{odd}}| \geq 5\sigma, \quad R^2 \geq 0.95, \quad \text{null-failure} \leq 1\%,$$

where Φ_{odd} denotes the odd (orientation-reversed) phase. The same criteria apply to the angle-normalized ratio R and to the mode ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ (Section 5).

Operational checklist (compact). (1) *Odd-even separation*: use path reversal or polarization order swap to isolate the geometric (odd) component. (2) *Off-support controls*: verify a zero baseline on zero-area/fully shielded loops. (3) *Bandwidth stability*: halve the bandwidth and check slope invariance. (4) *Polarization-axis rotation*: rotate the reference axis and verify preserved odd symmetry. (5) *Long-term stability*: log Allan deviation and mode-ratio drift to maintain sub-percent accuracy (Appendix G).

Upper-bound reporting and design feedback. If thresholds are not met, immediately report confidence intervals for $|R - 1|$ and $|\Delta\rho/\rho|$ as *separate* bounds on the isotropic $\langle\Phi\rangle$ component and the anisotropic $|\nabla\Phi|$ component. Then tune one design lever at a time—narrow the band, increase loop area, raise averaging depth, enforce adiabatic polarization, operate near the temperature-coefficient zero—to incrementally improve sensitivity (Section 5).

Summary. This section established the geometric chain $\Phi \rightarrow \Pi \rightarrow A \rightarrow F$, the invariant diagnostics (I_1, I_2) , and a clear map of where signals *should* and *should not* appear. Section 2 develops action, variations, and sources; Section 5 presents structure-preserving numerics and metrology. The same criteria support consistent data reads and upper-bound reports across platforms.

2. Action, Field Equations, and Constitutive Coupling

Chapter overview. We (i) present an action that treats the curvature scalar Φ and the $U(1)$ connection A_μ on equal footing, (ii) derive Maxwell–type equations and the dynamics of Φ by variation, (iii) verify *continuous reduction* to standard Maxwell theory under the weak–gradient bound $|\nabla\Phi| < \varepsilon$, and (iv) prepare a minimal expansion of the constitutive tensor $\chi(\Phi, \nabla\Phi)$ that separates *isotropic* and *anisotropic* responses and maps cleanly to observables (tilt–normalized R ratio and TE/TM mode ratio ρ). Foundational conventions are as in [Section 1](#) and [Appendix A](#).

2.1 Action and variations

On spacetime $(\mathcal{M}, g_{\mu\nu})$ we take

$$S[\Phi, A, \psi; g] = \int_{\mathcal{M}} \sqrt{-g} d^4x \left[\mathcal{L}_\Phi(\Phi, \nabla\Phi; g) - \frac{1}{4} F_{\mu\nu} H^{\mu\nu} + \mathcal{L}_m(\psi, D\psi; g) \right], \quad (14)$$

with $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$, $D_\mu = \nabla_\mu - iqA_\mu$, and matter fields ψ . The scalar sector is

$$\mathcal{L}_\Phi = -\frac{1}{2} \nabla_\mu \Phi \nabla^\mu \Phi - U(\Phi) + J \Phi, \quad (15)$$

where J is an external source and $U(\Phi)$ ensures stability (typical choices $m^2\Phi^2/2 + \lambda\Phi^4/4!$); positivity/causality constraints are summarized in [Appendix D](#).

Constitutive tensor: definition and symmetries. The response field is defined locally by

$$H^{\mu\nu} \equiv \chi^{\mu\nu}{}_{\rho\sigma}(\Phi, \nabla\Phi) F^{\rho\sigma}, \quad \chi^{\mu\nu}{}_{\rho\sigma} = -\chi^{\nu\mu}{}_{\rho\sigma} = -\chi^{\mu\nu}{}_{\sigma\rho}, \quad (16)$$

so that in the vacuum limit $\chi \rightarrow \chi_0$ one has $H^{\mu\nu} \rightarrow F^{\mu\nu}$ in Heaviside–Lorentz units (see [Appendix A](#)). Locality and the above index symmetries guarantee gauge invariance and positive–definite energy under the conditions detailed in [Appendix D](#).

Variation with respect to A_μ . Varying A_μ (with $\delta A_\mu|_{\partial\mathcal{M}} = 0$) yields the generalized Maxwell equation

$$\nabla_\nu H^{\mu\nu} = J^\mu, \quad J^\mu \equiv \frac{\partial \mathcal{L}_m}{\partial A_\mu}. \quad (17)$$

Gauge symmetry implies $\nabla_\mu J^\mu = 0$. On discrete meshes this continuity is *exactly* preserved by gauge links and Wilson loops; see [Appendix C](#).

Variation with respect to Φ (including feedback). The Euler–Lagrange equation for Φ reads

$$\square\Phi - U'(\Phi) = J + \frac{1}{4} \frac{\partial \chi_{\alpha\beta}{}^{\rho\sigma}}{\partial \Phi} F_{\rho\sigma} F^{\alpha\beta} + \nabla_\lambda \left(\frac{1}{4} \frac{\partial \chi_{\alpha\beta}{}^{\rho\sigma}}{\partial (\nabla_\lambda \Phi)} F_{\rho\sigma} F^{\alpha\beta} \right), \quad (18)$$

which makes explicit how isotropic/anisotropic constitutive changes feed back into Φ -dynamics. The linear–response expansion used later is stated in [Section 2.2](#).

Bianchi identity and constraints. By definition $F = dA$,

$$dF = 0 \quad \Longleftrightarrow \quad \partial_{[\lambda} F_{\mu\nu]} = 0, \quad (19)$$

equivalent to $\nabla \cdot \mathbf{B} = 0$ and $\partial_t \mathbf{B} + \nabla \times \mathbf{E} = 0$ in a 3+1 split ([Section 1.3](#)). Thus the independent dynamics reside in (Φ, A_μ) and the constitutive map χ ; the constraints $dF = 0$ and $\nabla_\mu J^\mu = 0$ follow automatically from the variational structure.

Boundary terms and canonical quantities. Boundary contributions are removed by gauge-compatible conditions ($\delta A_\mu n^\mu = 0$ on $\partial\mathcal{M}$) or by total divergences. The canonical stress-energy tensor is $T_{\mu\nu} = -\frac{2}{\sqrt{-g}} \frac{\delta S}{\delta g^{\mu\nu}}$, with the EM part bilinear in F and H . Explicit forms are tabulated in [Appendix A](#).

Continuous reduction and unit conventions. Under the weak-gradient bound $|\nabla\Phi| < \varepsilon$ and small-signal assumptions, $\chi(\Phi, \nabla\Phi) \rightarrow \chi_0$ and (17) reduces to standard Maxwell equations. In Heaviside-Lorentz units ($c = 1$), $H^{\mu\nu} \rightarrow F^{\mu\nu}$; SI conversions are summarized in [Appendix A](#). These results feed directly into the constitutive expansion ([Section 2.2](#)) and the observable mapping to R and ρ ([Section 2.3](#)).

2.2 Constitutive law and the weak-gradient regime

The working hypothesis is the following *triple constraint*:

$$H = \chi(\Phi, \nabla\Phi) : F, \quad \square\Phi - U'(\Phi) = J, \quad |\nabla\Phi| < \varepsilon, \quad (20)$$

where χ is a local rank-4 constitutive tensor obeying the standard index symmetries $\chi^{\mu\nu}{}_{\rho\sigma} = -\chi^{\nu\mu}{}_{\rho\sigma} = -\chi^{\mu\nu}{}_{\sigma\rho}$. In the vacuum limit $|\nabla\Phi| \rightarrow 0$ one requires $\chi(\Phi, \nabla\Phi) \rightarrow \chi_0$ so that $H^{\mu\nu} \rightarrow F^{\mu\nu}$ (Heaviside-Lorentz units; conventions in [Appendix A](#)). Then the field equation $\nabla_\nu H^{\mu\nu} = J^\mu$ reduces continuously to Maxwell (cf. [Section 2.1](#)).

Minimal expansion (linear response). Current observables reach sub-percent precision; it is therefore consistent to retain only the terms linear in Φ and $\nabla\Phi$:

$$\chi(\Phi, \nabla\Phi) = \chi_0 + \alpha \Phi \chi_0 + \eta \mathcal{K}(\nabla\Phi) + O(\Phi^2, \nabla\Phi^2), \quad (21)$$

with dimensionless coefficients α, η ([Appendix A](#)). The first correction rescales the isotropic impedance; the second encodes a weak anisotropy tied to the direction of $\nabla\Phi$.

Remark (model content). The expansion $\chi = \chi_0 + \alpha \Phi \chi_0 + \eta \mathcal{K}(\nabla\Phi) + \dots$ is the *only* place where empirical novelty can arise. If $\alpha = \eta = 0$ (equivalently $A = G = 0$), the framework collapses to pure Maxwell even though the geometric dictionary ($A, F=dA$) remains in place; i.e., Maxwell is *recovered continuously* in the weak-gradient window ([Section 2.5](#)).

Linear-dispersive expansion in the operating band. Because experimental platforms (resonators, waveguides, films) operate over finite bandwidths, we include weak dispersion:

$$\chi(\Phi, \nabla\Phi; \omega, \mathbf{k}) = \chi_0(\omega, \mathbf{k}) + \alpha(\omega) \Phi \chi_0(\omega, \mathbf{k}) + \eta(\omega) \mathcal{K}(\nabla\Phi; \omega, \mathbf{k}) + O(\Phi^2, \nabla\Phi^2).$$

Here α, η are dimensionless (or normalized) linear response coefficients, and $\chi_0(\omega, \mathbf{k})$ is the isotropic medium response in the Maxwell limit. The first-order dispersion correction used in the ρ -channel regression in this section corresponds to $\partial_\omega \chi$.

Structure of $\mathcal{K}(\nabla\Phi)$: symmetry-guided form. Let $n_\mu \equiv \nabla_\mu \Phi / |\nabla\Phi|$ and split spacetime with the projector $P_{\mu\nu} \equiv g_{\mu\nu} - n_\mu n_\nu$. A minimal parity-even, gauge-compatible choice that preserves the antisymmetry in each index pair is

$$\mathcal{K}^{\mu\nu}{}_{\rho\sigma}(\nabla\Phi) = \left(P^\mu{}_{[\rho} P^{\nu]}{}_{\sigma]} - \frac{1}{3} P^\alpha{}_{[\rho} P_{\sigma]\alpha} P^{\mu\nu} \right) - \left(n^{[\mu} P^{\nu]}{}_{[\rho} n_{\sigma]} \right), \quad (22)$$

which, in a 3+1 split and to leading order, reduces to the familiar uniaxial form proportional to $(\hat{\mathbf{n}} \cdot \hat{\mathbf{k}})^2 - \frac{1}{3}$, i.e. a $\cos 2\theta$ -type response for rotations around $\hat{\mathbf{n}}$. This is the geometric origin of the angle dependence used later ([Section 2.3](#)).

Symmetry classification (compressed summary). At linear order in $(\Phi, \nabla\Phi)$ and for local, passive, P/T -even media without external bias, the rank-4 tensor $\mathcal{K}(\nabla\Phi)$ must be built from $g_{\mu\nu}$ and the unit vector $n_\mu \equiv \nabla_\mu \Phi / |\nabla\Phi|$, while preserving antisymmetry in each index pair ($\mu \leftrightarrow \nu$ and $\rho \leftrightarrow \sigma$) and Onsager reciprocity $\chi^{\mu\nu}_{\rho\sigma} = \chi_{\rho\sigma}^{\mu\nu}$. A minimal parity-even basis is

$$\mathcal{K}^{\mu\nu}_{\rho\sigma} = a_1 P^\mu_{[\rho} P^\nu_{\sigma]} + a_2 (n^{[\mu} P^{\nu]}_{[\rho} n_{\sigma]}) - a_3 (P^\alpha_{[\rho} P_{\sigma]\alpha} P^{\mu\nu}), \quad P_{\mu\nu} \equiv g_{\mu\nu} - n_\mu n_\nu, \quad (23)$$

which reduces in 3+1 to a uniaxial response proportional to $(\hat{\mathbf{n}} \cdot \hat{\mathbf{k}})^2 - \frac{1}{3}$, i.e. a first harmonic $\cos 2\theta$ under rotations about $\hat{\mathbf{n}}$. All other P/T -even local linear terms are linear combinations of (23) up to trace redefinitions.

Excluded (or bounded) structures at the same order.

- *P/T -odd (axion/Tellegen):* $\Phi F_{\mu\nu} \tilde{F}^{\mu\nu}$ and $(\partial_\mu \Phi) A_\nu \tilde{F}^{\mu\nu}$ break reciprocity and are excluded in the baseline (microreversibility, no external bias). If present, treat as *nuisance* couplings and bound $\ll O(\eta|\nabla\Phi|)$.
- *Nonlocal/higher-derivative:* terms with ∂F or Hessian insertions (e.g. $H^{\alpha\beta} : F_{\alpha\mu} F_\beta^\mu$) are suppressed by the platform scale Λ and enter as $O((\omega/\Lambda), |\mathbf{k}|/\Lambda)$ renormalizations of a_i within our operating band.
- *Gauge-variant forms:* $(\partial\Phi) \cdot A \tilde{F}$ reduces by parts to $(\partial\Phi) \cdot *FF$ and is covered by the P/T -odd item above.

Practical dictionary. With $A \propto \alpha\langle\Phi\rangle$ (isotropic) and $G \equiv \eta|\nabla\Phi|$ (uniaxial anisotropy), the R -channel isolates the $\cos 2\theta$ piece (fixing a combination of $a_{1,2,3}$) while ρ pins the isotropic rescale; any residual P/T -odd signature (e.g. rotation-independent odd holonomy or reciprocity breaking) is flagged and bounded in the reporting templates.

Alternative models for $\mathcal{K}(\nabla\Phi)$ (symmetry-guided). The baseline choice is *uniaxial*, with $\hat{n} \equiv \nabla\Phi/|\nabla\Phi|$ defining the axis and yielding a $\cos 2\theta$ first harmonic under rotation. If needed, the following generalizations remain consistent with the mapping to observables.

1. **Biaxial (Hessian-eigenframe) model:** Let H_{ij} be the spatial Hessian of Φ with orthonormal eigenvectors $\{e^{(a)}\}$ and eigenvalues $\{\lambda_a\}$; define the principal plane Π by $\{e^{(1)}, e^{(2)}\}$ and

$$\mathcal{K}_{\mu\nu}^{\rho\sigma} = \sum_{a=1}^2 \kappa_a(\omega) (\mathcal{P}_{[\mu}^{(a)[\rho} \mathcal{P}_{\nu]}^{(a)\sigma]}), \quad \mathcal{P}_{ij}^{(a)} = e_i^{(a)} e_j^{(a)}.$$

This separates the relative weights of the two $\cos 2\theta$ couplings (angle conventions in [Appendix E](#)).

2. **Wave-vector dependence:** For plane waves $(\omega, \mathbf{k})$, allowing a factor proportional to $(\hat{n} \cdot \hat{k})^2 - \frac{1}{3}$ preserves the $\cos 2\theta$ law under experimental rotations: $\eta(\omega) \mathcal{K}(\nabla\Phi; \omega, \mathbf{k}) \propto (\hat{n} \cdot \hat{k})^2 - \frac{1}{3}$.
3. **Excluding parity-odd catalysts at first order:** The baseline analysis assumes P/T -even linear response; Levi-Civita-based chiral terms are excluded at first order and, if needed, bounded as $\ll O(\eta|\nabla\Phi|)$ in [Appendix D](#) (causality/positivity).

All these variants satisfy $\chi \rightarrow \chi_0$ as $|\nabla\Phi| \rightarrow 0$; the R, ρ mapping is unchanged except for coefficient renormalization.

Table 1: Symmetry-allowed building blocks at $O(\Phi, \nabla\Phi)$ and their status in the baseline analysis.

Structure	Parity / T	Local & linear	Status / comment
$P^\mu{}_{[\rho}P^\nu{}_{\sigma]}, \quad n^{[\mu}P^{\nu]}_{[\rho}n_{\sigma]},$ trace terms	P -even, T -even	Yes	Allowed ; spans uniaxial class $\Rightarrow$ first harmonic $\cos 2\theta$.
$\Phi F_{\mu\nu}\tilde{F}^{\mu\nu}, \quad (\partial\Phi)\cdot A\tilde{F}$	P -odd, T -odd	Yes (nonreciprocal)	Excluded in baseline; if detected, report as P/T -odd bound (Appendix D.3).
∂F or Hessian-weighted FF (e.g. $H : FF$)	P -even, T -even	Higher-derivative	Suppressed ; treated as small dispersive renormalizations (Appendix D.2).

Parameterization by A and G . For data analysis it is convenient to summarize the two linear corrections by

$$A \propto \alpha \langle \Phi \rangle, \quad G \equiv \eta |\nabla\Phi|,$$

so that A shifts an *isotropic* impedance scale (tracked by resonant frequencies/mode ratios) while G controls a *weak anisotropy* visible as $\cos 2\theta$ modulations. The mapping to the tilt-normalized ratio R and the TE/TM ratio ρ is stated in [Section 2.3](#).

Causality, positivity, and duality. Local, linear media must satisfy energy positivity and causal dispersion. In the isotropic limit ($\eta = 0$) the Lagrangian $-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$ admits continuous duality rotations ([Section 1.4](#); [2]). Small $\eta \neq 0$ acts as a controlled perturbation; bounds ensuring positive-definite energy, subluminal signal velocity, and Kramers–Kronig consistency are collected in [Appendix D](#).

Identifiability in practice. Keeping only $O(\Phi, \nabla\Phi)$, deviations from Maxwell scale as

$$\Delta O = O(A, G) + O(A, G)^2.$$

Thus, once an observable O is normalized to cancel trivial geometry—e.g. the tilt ratio R ([Section 2.3](#))—any residual $|\Delta O|$ at the sub-percent level directly bounds G (and, with ρ , also A). This is the quantitative sense in which (20) enforces *continuous reduction* to Maxwell for $|\nabla\Phi| \ll 1$.

Special cases and limits. (i) **Pure Maxwell:** $\alpha = \eta = 0 \Rightarrow H = F$. (ii) **Isotropic shift only:** $\eta = 0 \Rightarrow H = (1 + \alpha\Phi)F + O(\Phi^2)$; angle-based tests are null, ρ is sensitive. (iii) **Gradient anisotropy only:** $\alpha = 0 \Rightarrow H = F + \eta\mathcal{K}(\nabla\Phi) : F$; R is sensitive at $O(G)$ while ρ helps break degeneracy. These limits are used as cross-checks when fitting (A, G) jointly from heterogeneous datasets.

Weak-nonlinear window (optional; not used in baseline fits). Under high drive or narrow-band operation, we only *annotate* the constitutive law to indicate possible departures from the linear window:

$$\chi \simeq \chi_0 + \alpha\Phi\chi_0 + \eta\mathcal{K}(\nabla\Phi) + O(\Phi^2, (\nabla\Phi)^2, |F|^4),$$

without invoking higher-order terms in baseline fits. The *structure and admissible forms* of the higher-order corrections (including curvature-field terms such as Φ^2 , $(\nabla\Phi)^2$, mixed $\Phi \mathcal{K}$, and EM nonlinearities like $|F|^2$ and $\mathcal{Q}(F, F)$) are formalized in [Section 2.2.2](#) (see Eq. (30)).

Angular signatures (operational cue only). In weak anisotropy, harmonics separate by symmetry,

$$\begin{aligned} G^2 &\Rightarrow \cos 4\theta, \\ A^2 &\Rightarrow \text{angle-independent DC}, \\ AG &\Rightarrow \cos 2\theta \text{ (phase-shifted reinforcement)}, \end{aligned}$$

and the resulting *observable* expansions for $R(\theta)$ and $\Delta \ln \rho$ are given in [Section 2.2.2](#) (Eqs. (31)–(32)). We do not fit these terms in the baseline; they serve as diagnostics for extended analyses.

Causality/energy and regression handling (pointer). Admissibility constraints (Kramers–Kronig, passivity, causal falloff) are summarized in [Appendix D](#) and applied as priors in extended fits; practical regression uses added covariates and model-selection rules in [Appendix E–Appendix F](#) (design vector listed in [Section 2.2.2](#)).

Details: Higher-order structure, angular harmonics, and causality/energy constraints are formalized in [Section 2.2.2](#).

2.2.1 Physical basis and symmetry constraints on $\chi(\Phi, \nabla\Phi)$

EFT origin (integrating out heavy modes). At energies $E \ll \Lambda$, couplings between a light scalar Φ and electromagnetism are captured by local operators consistent with gauge/Lorentz symmetries:

$$\mathcal{L}_{\text{eff}} = -\frac{1}{4} Z(\Phi) F_{\mu\nu} F^{\mu\nu} - \frac{1}{4} Y(\nabla\Phi) F_{\mu\nu} F^{\mu\nu} - \frac{1}{4} \tilde{Y}(\nabla\Phi) F_{\mu\nu} \Pi^{\mu\nu\rho\sigma}(\nabla\Phi) F_{\rho\sigma} + \dots \quad (24)$$

With $Z(\Phi) = 1 + c_1 \Phi/\Lambda + \dots$ and $Y(\nabla\Phi) = c_2 (\partial\Phi)^2/\Lambda^4 + \dots$, linearizing around the operating point reproduces the constitutive expansion in [Section 2.2](#):

$$\begin{aligned} \chi(\Phi, \nabla\Phi) &= \chi_0 + \alpha \Phi \chi_0 + \eta \mathcal{K}(\nabla\Phi) + \dots, \\ \alpha &\sim \frac{c_1}{\Lambda}, \quad \eta \sim c_2 \frac{(\partial\Phi)^2}{\Lambda^4} \end{aligned} \quad (25)$$

Normalization note: see [Appendix A](#).

Here $\Pi^{\mu\nu\rho\sigma}$ denotes the most general projector built from $g_{\mu\nu}$ and the unit vector $n_\mu \equiv \nabla_\mu \Phi/|\nabla\Phi|$ that respects antisymmetry in each index pair; its leading uniaxial limit yields $\mathcal{K}(\nabla\Phi)$ of [Section 2.2](#).

Operator basis and power counting. Up to dimension-6, parity-even, gauge-invariant operators affecting linear response are

$$\Phi F_{\mu\nu} F^{\mu\nu}, \quad (\partial_\mu \Phi \partial_\nu \Phi) F^\mu{}_\rho F^{\nu\rho}, \quad (\partial\Phi)^2 F_{\mu\nu} F^{\mu\nu}.$$

Their Wilson coefficients scale as c_1/Λ , $c_{2,3}/\Lambda^2$. Rotational symmetry breaking enters only through the spurion n_μ , selecting uniaxial/biaxial projectors that reduce to the $\cos 2\theta$ law used downstream ([Section 2.3](#)). Parity-odd $\Phi F \tilde{F}$ (axion-like) is *excluded* from the baseline and, if needed, bounded separately in [Appendix D](#).

Symmetry priors and Ward identities. Gauge invariance fixes that corrections appear as functions of Lorentz scalars $I_1 \equiv F_{\mu\nu}F^{\mu\nu}$ and tensors built with n_μ that preserve antisymmetry in $(\mu\nu), (\rho\sigma)$. Lorentz covariance restricts index structures of $\mathcal{K}$; in the rest frame of n^μ the tensor decomposes into $\parallel, \perp$ blocks, producing uniaxial birefringence at $O(\eta)$. Charge conservation and duality are retained in the $\eta \rightarrow 0$ limit (Section 1.4), while small $\eta \neq 0$ preserves gauge Ward identities by construction.

Causality, positivity, and dispersion. Passivity implies $\text{Im } \chi(\omega) \geq 0$ in an appropriate eigenbasis; Kramers–Kronig then constrains the low-frequency slope:

$$\partial_\omega \text{Re } \chi_{ij}(0) = \frac{2}{\pi} \int_0^\infty \frac{\text{Im } \chi_{ij}(\omega')}{\omega'^2} d\omega' \geq 0, \quad (26)$$

used as a sign prior when marginalizing nuisance dispersion in the ρ -channel fit (Section 2.3; full statements in Appendix D). Energy positivity bounds the uniaxial contrast $|\chi_\parallel - \chi_\perp|$, ensuring subluminal group velocity and well-posedness of initial value problems (summary in Appendix D).

Microscopic avenues (illustrative, model-agnostic use). The EFT structures in (24) can arise from: (i) integrating out heavy charged fields coupled to Φ (threshold renormalization of vacuum polarization), (ii) portal-type couplings where Φ modulates a refractive index in an effective medium picture, or (iii) curvature-induced frame rotations that act as Berry connections on matter phases (Section 1.2). Our analysis remains model-agnostic; we only use the symmetry and power-counted form of χ .

Size estimates and naturalness (link to data). From the channel definitions,

$$A \propto \alpha \langle \Phi \rangle, \quad G = \eta |\nabla \Phi|.$$

Given bounds $|A| \leq A_{\text{max}}$ and $|G| \leq G_{\text{max}}$ extracted as in Section 2.3 and Section 2.5,

$$|\alpha| \lesssim \frac{A_{\text{max}}}{|\langle \Phi \rangle|}, \quad |\eta| \lesssim \frac{G_{\text{max}}}{|\nabla \Phi|}. \quad (27)$$

Interpreting $\alpha \sim c_1/\Lambda$ yields $\Lambda \gtrsim |c_1|/|\alpha|$, while a gradient-induced anisotropy with $\eta \sim c_2/\Lambda^2$ (uniaxial choice) gives $\Lambda \gtrsim (|c_2|/|\eta|)^{1/2}$. Thus sub-percent nulls translate into lower bounds on the EFT scale Λ (applied in Section 6).

Reporting policy (priors used in fits). We impose three hard priors when proposing and fitting (α, η) : (i) symmetry prior (gauge/Lorentz, parity-even at leading order), (ii) causality/positivity prior (KK consistency and passivity), and (iii) power-counting prior (operator dimensions and naturalness). These priors enter the covariance/regularization choices in the joint (A, G) estimation and in upper-bound reporting (Appendix F, Appendix G).

One-line dictionary to observables. Under these priors, the uniaxial projector induces a $\cos 2\theta$ signature in the angle channel R with slope $c_R(\theta) \propto \eta$, while ρ measures A at leading order and G through modal contrast (coefficients defined in Section 2.3). Consequently, the continuity test of Section 2.5 operationalizes the EFT expectation that first-order deviations vanish as $(\alpha, \eta) \rightarrow 0$, with residuals $= O(A^2, AG, G^2)$.

2.2.2 Uniqueness under symmetry and causality (theorem & proof)

Assumptions (A1–A6). (A1) **Gauge invariance:** observables depend only on $F = dA$.
(A2) **Lorentz covariance:** tensors are built from $g_{\mu\nu}$, $\nabla_\mu \Phi$ (and, if needed, the projector $P_{\mu\nu}$).
(A3) **Locality at first order:** no memory kernels at $O(\Phi, \nabla\Phi)$.
(A4) **Regularity:** analytic near $(\Phi, \nabla\Phi) = (0, 0)$.
(A5) **Passivity/causality:** $\text{Im } \chi(\omega) \geq 0$ and Kramers–Kronig holds (Appendix D.2).
(A6) **Parity-even baseline:** first-order P/T -odd couplings are excluded in the baseline and bounded separately (see Appendix D.3).

Theorem (complete first-order classification). Under (A1–A6), any local, linear constitutive law compatible with the Maxwell limit can be written—up to an overall normalization and field redefinitions—as

$$H = \chi(\Phi, \nabla\Phi) : F = (1 + \alpha \Phi) F + \eta \mathcal{K}(\nabla\Phi) : F + O(\Phi^2, (\nabla\Phi)^2) \quad (28)$$

where $\mathcal{K}(\nabla\Phi)$ is the *uniaxial* (or its biaxial generalization in Section 2.2) rank-4 tensor built from $n_\mu \equiv \nabla_\mu \Phi / |\nabla\Phi|$ and $P_{\mu\nu} = g_{\mu\nu} - n_\mu n_\nu$, preserving antisymmetry in each index pair and Onsager reciprocity. No other independent first-order scalars/tensors exist that simultaneously (i) are gauge invariant, (ii) vanish continuously as $|\nabla\Phi| \rightarrow 0$, and (iii) obey passivity/causality.

Proof (concise). (i) *Tensor basis:* At linear order in $(\Phi, \nabla\Phi)$, the only $U(1)$ -gauge-invariant 2-form building block is $F_{\mu\nu}$. Contracting with $g_{\mu\nu}$, n_μ , $P_{\mu\nu}$ and enforcing antisymmetry within each index pair yields exactly two parity-even, local structures: an isotropic rescale $\propto \Phi F$ and a traceless uniaxial projector $\mathcal{K}(\nabla\Phi) : F$.

(ii) *Exclusions:* $\Phi \tilde{F}^{\mu\nu} F_{\mu\nu}$ (axion/Tellegen) is P/T -odd $\Rightarrow$ excluded by (A6) at baseline and bounded separately. Terms like $(\partial\Phi) \cdot \partial F$ violate (A3) or reduce, via integration by parts plus Bianchi $dF=0$, to boundary/higher-order pieces. Nonlocal memory kernels appear only as dispersive covariates at $O(\beta)$ and are handled in the ρ -channel regression (Section 2.3, Appendix D.2).

(iii) *Maxwell reduction:* The Maxwell limit requires $\chi(\Phi, \nabla\Phi) \rightarrow \chi_0$ as $|\nabla\Phi| \rightarrow 0$, fixing the isotropic normalization and excluding any first-order tensor that would survive in this limit (continuity formalized in Section 2.5). $\square$

Compressed dictionary (to $\mathcal{K}$). Under the above constraints, a convenient parity-even basis for $\mathcal{K}$ is

$$\mathcal{K}^{\mu\nu}_{\rho\sigma} = a_1 P^\mu_{[\rho} P^\nu_{\sigma]} + a_2 (n^{[\mu} P^{\nu]}_{[\rho} n_{\sigma]}) - a_3 (P^\alpha_{[\rho} P_{\sigma]\alpha} P^{\mu\nu}), \quad P_{\mu\nu} \equiv g_{\mu\nu} - n_\mu n_\nu, \quad (29)$$

which in 3+1 dimensions reduces to the uniaxial form giving a first-harmonic $\cos 2\theta$ response under rotations about $\hat{n}$ (see Section 2.2).

Consequences. (i) *Parameter sufficiency:* (α, η) (equivalently (A, G)) form a complete first-order set.

(ii) *Identifiability:* R isolates G (uniaxial anisotropy), while ρ isolates A (isotropic rescale) with auxiliary sensitivity to G via modal contrast (Section 2.3).

(iii) *Reporting:* Baseline fits quote $(\hat{A}, \hat{G})$ and joint C.I.s; any P/T -odd residuals are flagged and bounded using the templates of Appendix D.3.

Higher-order extension (beyond first order). While (28) exhausts all first-order, local, parity-even possibilities, we record the leading higher-order corrections—still constrained by (A1–A6)—as

$$\begin{aligned}\chi \simeq & \chi_0 + \alpha \Phi \chi_0 + \eta \mathcal{K}(\nabla \Phi) + \beta_1 \Phi^2 \chi_0 + \beta_2 (\nabla \Phi)^2 \chi_0 + \beta_3 \Phi \mathcal{K}(\nabla \Phi) \\ & + \gamma_1 |F|^2 \chi_0 + \gamma_2 \mathcal{Q}(F, F) + O(\Phi^3, \nabla \Phi^3, |F|^6),\end{aligned}\quad (30)$$

where $\mathcal{Q}$ is any symmetric bilinear form in F (allowing mode mixing) compatible with reciprocity.

Angular-harmonic signatures (diagnostics). In weak anisotropy, the induced harmonics separate by symmetry:

$$\begin{aligned}G^2 &\Rightarrow \cos 4\theta \text{ component,} \\ A^2 &\Rightarrow \text{angle-independent DC term,} \\ AG &\Rightarrow \cos 2\theta \text{ (phase-shifted reinforcement).}\end{aligned}$$

Accordingly, the *observable* expansions can be organized as

$$R(\theta) = 1 + c_R G \cos 2\theta + d_R G^2 \cos 4\theta + e_R A + q_R A^2 + r_R AG \cos 2\theta + O(G^3), \quad (31)$$

$$\Delta \ln \rho = s_A A + s_G G \langle \cos 2\theta \rangle + q_{AA} A^2 + q_{GG} G^2 + q_{AG} AG + O(G^3), \quad (32)$$

with coefficients fixed by geometric overlaps and mode profiles.

Constraint lemma (causality & energy). Higher-order corrections to $\chi(\omega, \mathbf{k})$ must satisfy simultaneously: (i) Kramers–Kronig consistency, (ii) positivity of time-averaged stored energy for passive media, (iii) causal high-frequency falloff. Operationally (see [Appendix D](#)), we impose priors such as

$$\gamma_1 \geq 0, \quad |\beta_i| \ll |\eta|, \quad \text{band-limited constraints on } \partial_\omega \chi,$$

and, upon violation, drop offending terms or report upper bounds only.

Regression rule (design and reporting). Extended fits include the covariates

$$X \leftarrow [1, \cos 2\theta, \cos 4\theta, \text{drive}^2, \hat{A}, \hat{G}, \hat{A}^2, \hat{G}^2, \hat{A}\hat{G}],$$

retain a minimal set by VIF and AIC/BIC, and mandate multi-level drive and bidirectional angle sweeps. Baseline numbers set $\{\beta_i, \gamma_j\} = 0$ and are kept in the main text; higher-order coefficients from (31)–(32) are reported *in appendix tables only* (see [Appendix E](#), [Appendix F](#)).

Addendum: scope of parity-odd tests (optional). If a platform admits P/T -odd diagnostics (e.g., nonreciprocal transmission, rotation-independent odd holonomy), augment (28) by a nuisance axion-like term $\xi \Phi \tilde{F} : F$ with independent prior $|\xi| \ll |\eta|$; report (A, G, ξ) bounds with the causality and positivity checks in [Appendix D.3](#).

Dispersion and higher derivatives (bookkeeping). Finite bandwidth enters through $\partial_\omega \chi$ at first order and is treated as a covariate in ρ fits ([Section 2.3](#), [Appendix D.2](#)). Higher-derivative/Hessian insertions are power-counted by the platform scale Λ and absorbed as $O(\omega/\Lambda, |\mathbf{k}|/\Lambda)$ renormalizations of the coefficients a_i in (29).

Placement note. This section supplies the “uniqueness” invoked in [Section 2.2](#): the boxed form (28) is the only first-order, local, parity-even extension consistent with gauge/Lorentz symmetry and passivity, ensuring that the Maxwell reduction of [Section 2.5](#) is the *generic* outcome in the weak-gradient window.

2.2.3 Scope of the weak-gradient assumption and considerations under strong gradients

Baseline window (weak-gradient regime). Assume the dimensionless parameters

$$\varepsilon_\Phi := |\alpha \Phi| \ll 1, \quad \varepsilon_G := |\eta| |\nabla \Phi| \ll 1,$$

under which the local, parity-even constitutive law admits the first-order expansion

$$H = (1 + \alpha \Phi) F + \eta \mathcal{K}(\nabla \Phi) : F + O(\Phi^2, (\nabla \Phi)^2). \quad (33)$$

In this window the Maxwell limit is reached *continuously* (see [Section 2.5](#)); the normalization is fixed in [Appendix A.4](#), and uniqueness at first order follows from [Section 2.2.2](#).

Departure indicators (model-independent signatures). The following *theoretical* signatures mark the breakdown of (33):

1. Emergence of higher angular harmonics that cannot be generated at $O(G)$ (e.g. a nonvanishing $\cos 6\theta$ component independent of G^3 -order combinatorics).
2. Nonlinear drive response incompatible with linear material response at fixed geometry.
3. Dispersion coefficients $\partial_\omega \chi|_{\omega_0}$ contradicting passivity/causality priors ([Appendix D.2](#)).

These are *logical* consequences of the symmetry and regularity assumptions and do not rely on any particular estimator.

Strong-gradient extension (symmetry preserved, magnitude freed). When $|\nabla \Phi|$ is not perturbative, retain the symmetry axis but release the amplitudes:

$$H = (1 + \mathcal{A}(\Phi)) F + \mathcal{G}(|\nabla \Phi|) \mathcal{K}(\hat{n}) : F, \quad \hat{n} := \frac{\nabla \Phi}{|\nabla \Phi|}, \quad (34)$$

with

$$\mathcal{A}(0) = 0, \quad \mathcal{G}(0) = 0, \quad \mathcal{A}'(0) = \alpha, \quad \mathcal{G}'(0) = \eta.$$

Here $\mathcal{A}, \mathcal{G}$ are unknown *scalar* response functions constrained only by (i) passivity/causality ([Appendix D.2](#)), (ii) positivity and uniaxial-contrast bounds ([Appendix D.3](#)), and (iii) regularity at the origin. Concrete realizations include monotone C^1 splines or saturating Padé forms (e.g. $\mathcal{G}(x) = \eta x / (1 + \kappa x)$), but no particular choice is required for the statements below.

Angular structure (harmonic content). *Proposition.* In (34), rotational symmetry around $\hat{n}$ fixes the angular dependence to an even Fourier series $\sum_{m \geq 1} a_{2m} \cos(2m\theta)$ with coefficients a_{2m} algebraically determined by $\mathcal{G}$ and higher-order contractions of $\mathcal{K}$. At first nontrivial order, $\cos 2\theta$ is controlled by $\mathcal{G}'(|\nabla \Phi|=0)$; higher harmonics require nonlinear dependence of $\mathcal{G}$ or higher powers of $|\nabla \Phi|$. *Corollary.* In the weak-gradient limit, only the $\cos 2\theta$ term survives at $O(G)$, while $\cos 4\theta$ enters at $O(G^2)$, in agreement with [Section 2.2.2](#).

Channel-agnostic bounds (without committing to a parametric form). Let an observable obey

$$\Delta \ln(\text{observable}) = A + G \cos 2\theta + \text{higher even harmonics},$$

with $A = \mathcal{A}(\Phi)$ and $G = \mathcal{G}(|\nabla\Phi|)$ in the notation of (34). If $\mathcal{G}$ is nonnegative and nondecreasing on $[0, x_*]$, then for any coefficient $c_R(\theta)$ associated with the R -channel,

$$\sup_{x \in [0, x_*]} \mathcal{G}(x) \leq \frac{\sup_{\theta} |R(\theta) - 1|}{\inf_{\theta} |c_R(\theta)|}.$$

An analogous inequality holds for $\mathcal{A}$ via the ρ -channel (Section 2.3). Thus joint information from (R, ρ) yields bounds on $(\sup \mathcal{G}, \sup \mathcal{A})$ that are *robust* to the detailed functional shape.

Minimal assumptions for identifiability. *Lemma.* Suppose $\mathcal{A}$ and $\mathcal{G}$ are C^1 and monotone on the probed domain, and the geometry kernels in R and ρ have nonvanishing overlaps with, respectively, the anisotropic and isotropic sectors. Then (A, G) at each operating point are identifiable up to a common overall normalization fixed by the Maxwell limit (cf. Section 2.5).

Remarks. (i) The extension (34) is the unique symmetry-preserving generalization of (33) at fixed principal axis $\hat{n}$; (ii) the admissibility conditions referenced above ensure compatibility with energy positivity and Kramers–Kronig constraints; (iii) when the weak-gradient hypothesis is recovered *a posteriori*, one may set $\mathcal{A}(\Phi) = \alpha\Phi + O(\Phi^2)$ and $\mathcal{G}(|\nabla\Phi|) = \eta|\nabla\Phi| + O(|\nabla\Phi|^2)$ to return to (33).

2.2.4 Global consistency, singular sets, and integer quantization (from Čech 2-cocycles to the first Chern number)

Setting and hypotheses (G1–G4). (G1) **Background manifold:** spacetime $(\mathcal{M}, g_{\mu\nu})$ is finite-dimensional, connected, and time-/space-orientable.

(G2) **Gauge bundle:** electromagnetism is modeled by a principal $U(1)$ bundle $P \rightarrow \mathcal{M}$ with connection A and curvature $F = dA$.

(G3) **Singular set:** $\Sigma := \{x \in \mathcal{M} \mid \nabla\Phi(x) = 0 \text{ or } \Phi \text{ is nonregular}\}$. The analysis proceeds on $\mathcal{M} \setminus \Sigma$; assume Σ has Lebesgue measure zero in the windows of interest.

(G4) **Constitutive tensor:** $\chi(\Phi, \nabla\Phi)$ is smooth on $\mathcal{M} \setminus \Sigma$, satisfies positivity/reciprocity, and is homotopic to χ_0 (Appendix D).

Local trivializations and Čech data. Let $\{U_i\}$ cover $\mathcal{M} \setminus \Sigma$ and choose local potentials A_i with

$$A_j - A_i = d\lambda_{ij} \quad \text{on } U_i \cap U_j, \quad g_{ij} := e^{i\lambda_{ij}} : U_i \cap U_j \rightarrow U(1).$$

On triples $U_i \cap U_j \cap U_k$,

$$\lambda_{ij} + \lambda_{jk} + \lambda_{ki} = 2\pi n_{ijk}, \quad n_{ijk} \in \mathbb{Z},$$

so $\{n_{ijk}\}$ is a Čech 2-cocycle. Its class is the first Chern class $c_1 \in H^2(\mathcal{M} \setminus \Sigma, \mathbb{Z})$.

Flux quantization. For any closed two-surface $\mathcal{S} \subset \mathcal{M} \setminus \Sigma$,

$$\frac{1}{2\pi} \int_{\mathcal{S}} F = \langle c_1, [\mathcal{S}] \rangle \in \mathbb{Z}. \quad (35)$$

In particular, the AB period $\Delta B \cdot \text{area} = \Phi_0$ is a manifestation of (35). This conclusion depends only on $F = dA$ and the integer cohomology of the patching data.

Invariance under admissible constitutive deformations. *Proposition.* Under (G1–G4) and the assumptions of [Section 2.2.2](#), any variation of χ within the baseline class (e.g. shifts in α, η or reweightings via $\mathcal{K}(\nabla\Phi)$) preserves $dF = 0$ and the patching relations; hence c_1 and (35) are unchanged.

Sketch. χ enters only as $H = \chi : F$. The Bianchi identity is purely geometric and independent of χ . With (A1)–(A6) and (G4), $\chi \simeq \chi_0$ through a homotopy on $\mathcal{M} \setminus \Sigma$, so the class $[F/2\pi]$ is invariant.

Singular set and extendability. Let $n_\mu = \nabla_\mu \Phi / |\nabla\Phi|$ on $\mathcal{M} \setminus \Sigma$. Although n_μ is undefined on Σ , the following suffices for physical predictions:

Lemma. If (S1) Σ has measure 0, and (S2) each connected component admits a bounded–variation extension of n_μ , then angle–averaged observables and their leading $\cos 2\theta$ harmonic on $\mathcal{M} \setminus \Sigma$ are stable under modifications of n_μ on Σ ; corrections are $O(|\Sigma|)$ and below experimental resolution (angle conventions in [Appendix E](#)).

Discrete counterpart (Wilson–Čech correspondence). On a cell complex, the Wilson loop $W(\partial p) = \exp(i \oint_{\partial p} A)$ equals the plaquette flux $\exp(i \int_p F)$. Integer triple–overlap data $\{n_{ijk}\}$ appear as integer loop phases. Structure–preserving schemes ([Appendix C](#)) respect these relations and thus maintain (35) at the discrete level.

Corollaries for channel observables. (i) **R –channel.** The projection law $\Delta B \propto 1/\cos \theta$ follows from (35) and is unaffected by first–order, parity–even deformations of χ ; anisotropy enters only through the geometric $\cos 2\theta$ factor ([Section 2.3](#)).

(ii) **ρ –channel.** The isotropic rescale A leaves c_1 unchanged and acts in common mode; separation relies on differential modal sensitivity ([Section 2.3](#), [Appendix E](#)).

(iii) **Maxwell reduction.** In the weak–gradient window ($|\nabla\Phi| < \varepsilon$), the continuous reduction to Maxwell ([Section 2.5](#)) coexists with the global constraint: setting $A=G=0$ leaves all AB–type indicators intact.

Edge cases. Parity–odd bulk terms (e.g. $\Phi F \tilde{F}$) reduce to total derivatives $d(A \wedge F)$ and can influence boundary observables without altering c_1 ; they are excluded in the baseline ([Appendix D.3](#)). When probed, they should be analyzed in dedicated odd channels and reported separately from the integer topology.

Synthesis. This subsection completes the local first–order classification of [Section 2.2.2](#) by exhibiting the global bundle constraint (35). Admissible deformations of χ respect both Maxwell reduction ([Section 2.5](#)) and the topological content encoded by c_1 , thereby justifying the interpretation of the experimental channels in [Section 2.3](#).

2.3 Isotropic/anisotropic decomposition and experimental parameters

Preliminaries and scope. The mappings developed in this section follow from the linear–dispersive expansion introduced in [Section 2.2](#),

$$\chi(\Phi, \nabla\Phi; \omega, \mathbf{k}) = \chi_0(\omega, \mathbf{k}) + \alpha(\omega) \Phi \chi_0(\omega, \mathbf{k}) + \eta(\omega) \mathcal{K}(\nabla\Phi; \omega, \mathbf{k}) + \dots,$$

and use the angle/frame conventions of [Appendix E](#). Changing only the geometric basis (uniaxial $\leftrightarrow$ biaxial) preserves the $\{1, \cos 2\theta, \sin 2\theta\}$ regression basis and the estimation procedure; coefficients are renormalized but the channel mapping remains the same.

Isotropic component A . The mean curvature field $\langle\Phi\rangle$ rescales the effective EM impedance globally. Under linear response,

$$Z_{\text{vac}} \longrightarrow Z_{\text{vac}} (1 + \alpha \langle\Phi\rangle), \quad A \equiv \kappa_A \alpha \langle\Phi\rangle,$$

where κ_A follows the normalization in [Appendix A](#). In resonators, mode-dependent energy weighting gives

$$\frac{\Delta f_m}{f_m} \simeq -\frac{1}{2} \mathcal{W}_m A \quad (m = \text{TE}, \text{TM}),$$

so that the ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ partially cancels common-mode drifts (derivation of $\mathcal{W}_m$ in [Appendix E](#)).

Anisotropic component G . The gradient $\nabla\Phi$ selects a unit direction $\hat{\mathbf{n}}$, and a uniaxial response appears with a $\cos 2\theta$ harmonic:

$$G \equiv \eta |\nabla\Phi|, \quad \delta\mathcal{O}_{\text{aniso}}(\theta) \propto G \cos 2\theta + O(G^2),$$

where θ is the geometric angle between the experimental drive (tilted field, polarization axis, etc.) and $\hat{\mathbf{n}}$. Angle conventions and frames are fixed in [Appendix E](#). Biaxial variants (Hessian-eigenframe) simply reweight the in-plane couplings while preserving the $\cos 2\theta$ basis.

Tilt-period normalization R (angle channel). In tilted-field quantum-oscillation data, the ideal projection law is $\Delta B(\theta) \propto \cos \theta$. Sub-percent departures are captured by

$$\boxed{R(\theta_1, \theta_2) = \frac{\Delta B(\theta_2)}{\Delta B(\theta_1)} \frac{\cos \theta_1}{\cos \theta_2} = 1 + c_R G + O(G^2),} \quad (36)$$

with a geometry/material constant c_R ([Appendix E](#)). Given a measurement error σ_R , the 95% C.I. bound

$$G \lesssim \frac{z_{0.975} \sigma_R}{|c_R|}$$

follows. Construction of σ_R (repeatability, digitization, tilt calibration) is detailed in [Appendix F](#) and the metrology log [Appendix G](#).

TE/TM mode ratio ρ (resonator channel). The ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ is directly sensitive to the isotropic rescale and, via differing field profiles/polarizations, also to G :

$$\frac{\Delta\rho}{\rho} = c_\rho^{(A)} A + c_\rho^{(G)} G + O(A^2, G^2, AG), \quad (37)$$

where $c_\rho^{(A)}, c_\rho^{(G)}$ are mode- and boundary-dependent ([Appendix E](#)). With long-term stability σ_ρ , $|\Delta\rho/\rho| < z_{0.975}\sigma_\rho$ yields combined bounds on (A, G) . The stability/traceability items appear in [Appendix G](#).

Dispersive correction (first-order in frequency). To separate residuals due to material/structure dispersion, write the susceptibility near a carrier ω_0 as

$$\chi(\omega) \simeq \chi_0 + \beta (\omega - \omega_0), \quad \beta \equiv \left. \frac{\partial\chi}{\partial\omega} \right|_{\omega_0}. \quad (38)$$

For a mode $m \in \{\text{TE}, \text{TM}\}$, the linearized frequency shift gains an additive term

$$\Delta \ln f_m \simeq s_{A,m} A + s_{G,m}(\theta) G + d_m \beta (\omega - \omega_0) + n_m, \quad (39)$$

where $d_m \equiv \partial(\ln f_m)/\partial\chi|_{\chi_0}$ is a mode–overlap factor ([Appendix E](#)). This term is treated as a *nuisance covariate* downstream: it is explicitly regressed in the ρ –channel ([Section 3.3](#)) and accounted for in the uncertainty pipeline ([Section 4.4](#)). To first order, the R –channel ([Section 3.2](#)) is insensitive to ω through the period–ratio definition, so dispersion mainly enters the ρ analysis.

Joint estimation and covariance. Combine both channels as

$$\begin{bmatrix} R - 1 \\ \Delta\rho/\rho \end{bmatrix} = \underbrace{\begin{bmatrix} c_R & 0 \\ c_\rho^{(G)} & c_\rho^{(A)} \end{bmatrix}}_{\equiv \mathbf{M}} \begin{bmatrix} G \\ A \end{bmatrix} + \boldsymbol{\varepsilon}, \quad \text{Cov}[\boldsymbol{\varepsilon}] = \boldsymbol{\Sigma}.$$

The weighted least-squares estimate $(\hat{A}, \hat{G})^\top = (\mathbf{M}^\top \boldsymbol{\Sigma}^{-1} \mathbf{M})^{-1} \mathbf{M}^\top \boldsymbol{\Sigma}^{-1} \mathbf{y}$ gives $\text{Cov}[\hat{A}, \hat{G}] = (\mathbf{M}^\top \boldsymbol{\Sigma}^{-1} \mathbf{M})^{-1}$, and C.I./upper bounds follow the recipe in [Appendix F](#). Operationally, R is most sensitive to G , while ρ is most sensitive to A , so the combination de-correlates the two.

Sensitivity and scaling. For target precisions (δ_R, δ_ρ) ,

$$G_{\min} \sim \frac{\delta_R}{|c_R|}, \quad A_{\min} \sim \frac{\delta_\rho}{|c_\rho^{(A)}|}.$$

Typically δ_R is limited by angle repeatability and SNR, whereas δ_ρ is limited by frequency reference and temperature control. Resource estimates and checklists are summarized in the metrology tables of [Appendix G](#).

Summary (channel–parameter map). (i) R compresses G to sub-percent through the $\cos 2\theta$ residue ([Eq. \(36\)](#)). (ii) ρ responds directly to A and, via modal contrast, provides auxiliary sensitivity to G ([Eq. \(37\)](#)). (iii) The joint fit quantitatively tests Maxwell reduction under $|\nabla\Phi| < \varepsilon$, see [Section 2.5](#).

2.3.1 Geometry-anchored exemplars, derivations, and uncertainties

AB rings: baseline period and error propagation. For a circular Aharonov–Bohm ring with mean radius r and area $A = \pi r^2$, the leading period is

$$\Delta B = \frac{\Phi_0}{A}, \quad \Phi_0 = \frac{h}{e} = 4.135667696 \times 10^{-15} \text{ Wb}.$$

To first order, the radius uncertainty σ_r propagates as

$$\frac{\sigma_{\Delta B}}{\Delta B} \simeq \left| \frac{\partial \ln \Delta B}{\partial r} \right| \sigma_r = \left| -2 \frac{\sigma_r}{r} \right| \implies \sigma_{\Delta B} \simeq 2 \Delta B \frac{\sigma_r}{r}.$$

(Edge width w is listed for completeness; ΔB at leading order depends only on A . Finite-width and lead corrections can be included as an A_{eff} refinement in [Section 6.2–Section 6.3](#).)

Q2D dHvA (tilt): normalized ratio R . From digitized points $\{\theta_i, \Delta B(\theta_i)\}$ we form

$$R(\theta_1, \theta_2) = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2},$$

with angle pairs chosen to have similar $|\tan \theta|$ to limit angle-bias amplification (Appendix F). The entry is flagged “TBD” here and populated by the pre-registered digitization in Section 6.1.

Resonators (WGMR / cavity): dual-window stability. On a common timebase, $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$ summarizes differential stability. We list the *short* window (comb-beat/linewidth; instrumental) and the *long* window (counter; Allan-minimum) per Section 3.3 and Section 4.4.

Table 2: Geometry–anchored observables (R–channel; concise).

Platform	Source	Params	Derived	Value
Graphene ring	PRB 96 (2017)[68]	$r = 600 \text{ nm}$	ΔB	3.657 mT
Graphene ring	PRB 96 (2017)[68]	$r = 700 \text{ nm}$	ΔB	2.687 mT
Q2D dHvA	This work (Section 6.1)	$\{\theta, \Delta B\}$	$R(\theta_1, \theta_2)$	<i>TBD</i>

Table 3: Geometry–anchored observables (ρ –channel; concise).

Platform	Source	Params	Derived	Value
Sapphire WGMR	This work (Sec. 3.3)	Q , comb/linewidth	$ \Delta\rho/\rho $	10^{-9} (short), 3×10^{-3} (long)

Numerical examples (uncertainty). For $r = 600 \text{ nm}$ with $\sigma_r = 1 \text{ nm}$: $\Delta B = 3.6567 \text{ mT}$ and $\sigma_{\Delta B} \simeq 2 \Delta B \sigma_r / r = 2 \times 3.6567 \text{ mT} \times (1/600) \approx 0.0122 \text{ mT}$ (0.33%). For $r = 700 \text{ nm}$: $\sigma_{\Delta B} \approx 0.0077 \text{ mT}$ (0.29%). These figures set the scale for R -channel precision when angle-readout errors are subdominant.

Display and cross-referencing. The table anchors the forward map (Eq. (36), Eq. (37)) with concrete geometries. Entries feed into Section 6.2–Section 6.4 (data reconstruction, CIs), and into the joint (A, G) estimator in Section 3.4. Where digitization is used, provenance (file/script hashes, operator ID, seed) and the conservative error model follow Appendix F and Appendix G.

2.4 Conservation laws, duality, and constraints

Gauge invariance and continuity. If the action $S[\Phi, A, \psi; g]$ is invariant under $A_\mu \rightarrow A_\mu + \partial_\mu \lambda$, the field equation $\nabla_\nu H^{\mu\nu} = J^\mu$ immediately implies charge conservation

$$\nabla_\mu J^\mu = 0. \quad (40)$$

Together with the Bianchi identity $dF = 0$, this is a geometric statement independent of the detailed choice of χ and J^μ (Section 1.3). On a lattice, gauge links $U_\ell = e^{iq \int_\ell A \cdot dl}$ and Wilson loops preserve the exact difference identities

$$\Delta_\nu H^{\mu\nu} = J^\mu, \quad \Delta_\mu J^\mu = 0, \quad (41)$$

where Δ is the boundary operator; Hodge weights are chosen to ensure discrete energy balance (Appendix C).

Stress–energy and energy flow. With the constitutive law $H = \chi(\Phi, \nabla\Phi) : F$, the electromagnetic stress–energy reads

$$T_F^{\mu\nu} = F^\mu{}_\alpha H^{\nu\alpha} - \frac{1}{4} g^{\mu\nu} F_{\alpha\beta} H^{\alpha\beta}. \quad (42)$$

Using the field and matter equations one obtains $\nabla_\mu T^{\mu\nu} = \text{exchange with matter}$. In the plane–wave limit, the time–like character of the Poynting 4–vector $S^\mu = T^{\mu\nu} u_\nu$ follows from the positive–definiteness of χ (Appendix D).

Duality rotations and their breaking. In an isotropic vacuum ($\chi = \chi_0$), the Lagrangian $\mathcal{L}_F = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$ admits continuous duality rotations

$$F \mapsto F \cos \alpha + {}^*F \sin \alpha, \quad H \mapsto H \cos \alpha + {}^*H \sin \alpha \quad (43)$$

[2]. Anisotropic corrections via $\chi(\Phi, \nabla\Phi)$ generically break this symmetry by $O(\eta|\nabla\Phi|)$, leaving a residual $\cos 2\theta$ harmonic in observables (Section 2.3). The isotropic component A preserves duality by redefining the impedance, while the anisotropic component G weakly violates it.

Symmetries of the constitutive tensor and reciprocity. A physical $\chi^{\mu\nu}{}_{\rho\sigma}$ satisfies

$$\chi^{\mu\nu}{}_{\rho\sigma} = -\chi^{\nu\mu}{}_{\rho\sigma} = -\chi^{\mu\nu}{}_{\sigma\rho}, \quad \chi^{\mu\nu}{}_{\rho\sigma} = \chi_{\rho\sigma}{}^{\mu\nu}, \quad (44)$$

i.e., antisymmetry within each index pair and symmetry under pair exchange. Under microreversibility, an Onsager–type reciprocity takes the same form. Violating these relations leads to ambiguous energy definitions or unphysical power flow (Appendix D).

Positivity, causality, and dispersion (frequency–domain constraints). For linear response $H(\omega, \mathbf{k}) = \chi(\omega, \mathbf{k}) : F(\omega, \mathbf{k})$, the following must hold: (i) *passivity/positivity*: $\int_{-\infty}^t dt' \mathbf{E} \cdot \partial_{t'} \mathbf{D} + \mathbf{H} \cdot \partial_{t'} \mathbf{B} \geq 0$; (ii) *causality*: analyticity of $\chi(\omega, \mathbf{k})$ in the upper half–plane with Kramers–Kronig relations; (iii) *hyperbolicity (no superluminal transport)*: the dispersion polynomial yields real $\omega(\mathbf{k})$ with finite group velocity. These imply stability bounds on $|A|$ and G and standard sum rules in the low-/high-frequency limits (Appendix D).

Plane–wave dispersion and hyperbolic character. In a homogeneous background, Fourier modes obey

$$\mathbf{M}(\omega, \mathbf{k}; \chi) \begin{bmatrix} \mathbf{E} \\ \mathbf{H} \end{bmatrix} = 0, \quad (45)$$

with dispersion surface $\det \mathbf{M} = 0$. Positivity and symmetry of χ guarantee real branches and bounded group speeds. This fixes the admissible parameter region for experimental scans (Appendix E).

Conserved quantities in lattice implementations. Structure–preserving schemes combine (i) incidence–Hodge Stokes pairs on space–time complexes, (ii) symmetric/positive discrete Hodge $\star_h$, (iii) consistent boundary treatment (PEC, PMC, periodic), to yield

$$\mathcal{E}^{n+1} - \mathcal{E}^n = -\Delta t \sum_{\text{faces}} \langle \mathbf{S} \cdot \hat{\mathbf{n}} \rangle + O(\Delta t^3), \quad (46)$$

with exact discrete charge continuity (Appendix C). This remains stable when χ varies slowly in time.

Summary (operational constraints). (i) gauge invariance $\Rightarrow$ charge continuity; (ii) pair-exchange symmetry and positivity of $\chi \Rightarrow$ physical energy and Poynting flow; (iii) causality/dispersion $\Rightarrow$ analytic constraints on $\chi(\omega, \mathbf{k})$; (iv) weak anisotropy $\Rightarrow O(G)$ duality breaking with $\cos 2\theta$ residuals. These four lines delimit the allowed region for (A, G) and require reporting confidence intervals for $\{R, \rho\}$ together with symmetry/positivity/causality checks ([Appendix F](#), [Appendix G](#)).

2.5 Maxwell reduction (continuity test): formalization

Definition (continuity). For an experimental channel $\mathcal{C} \in \{R\text{-channel, resonator channel}\}$ under the weak-gradient constraint $|\nabla\Phi| < \varepsilon$, let $O_{\mathcal{C}}^{\text{curv}}(A, G)$ be the prediction including the constitutive pair (A, G) and $O_{\mathcal{C}}^{\text{Mxw}}$ the Maxwell prediction. Define

$$\Delta O_{\mathcal{C}} \equiv O_{\mathcal{C}}^{\text{curv}}(A, G) - O_{\mathcal{C}}^{\text{Mxw}}.$$

Declaration (equivalence baseline). Setting $(A, G) = (0, 0)$ —equivalently $\alpha = \eta = 0$ —reduces all observables to the Maxwell predictions *exactly*. Evidence consistent with this baseline must be reported explicitly as “equivalence confirmed within CI,” not as discovery.

We say *continuity (Maxwell reduction)* holds if

$$\Delta O_{\mathcal{C}} = 0 + O(A^2, AG, G^2), \quad (|\nabla\Phi| < \varepsilon). \quad (47)$$

Thus any first-order (linear) deviation cancels, and observable differences remain only at second order in (A, G) . The meaning of A, G and their measurement mapping follow [Section 2.3](#).

Decision statistics (by channel). (i) ***R*-channel.** The tilt–period normalized indicator R admits a local linearization $R = 1 + c_R(\theta) G + O(G^2)$ with geometry- and band-dependent slope c_R . From (36),

$$\delta_R \equiv |R - 1| \Rightarrow |G| \lesssim C_R \delta_R + O(\delta_R^2), \quad (48)$$

with $C_R \equiv \sup_{\theta} |c_R(\theta)|^{-1}$ as a conservative constant. With repeated measurements at several θ , least-squares on the slope $\hat{c}_R$ yields a $(1 - \alpha)$ confidence interval (CI) for G :

$$|G| \leq \frac{z_{1-\alpha/2} \sigma_R}{|\hat{c}_R|}, \quad \sigma_R^2 = \text{Var}(R - 1). \quad (49)$$

Graph digitization, axis nonlinearity, and tilt uncertainties are aggregated into σ_R per the metrology model in [Appendix F](#).

(ii) **Resonator channel.** The mode ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ linearizes as $\rho = \rho_0 [1 + \kappa_A A + \kappa_G(\theta)G + O(A^2, AG, G^2)]$. From the long-term relative stability $\delta_\rho \equiv |\Delta\rho/\rho|$,

$$|A| \lesssim C_\rho \delta_\rho, \quad |G| \lesssim C'_\rho \delta_\rho, \quad (50)$$

with C_ρ, C'_ρ computed from mode-overlap coefficients and Q-limits; see [Appendix E](#).

Joint decision (channel fusion). Treating *R*- and resonator channels as independent noise sources, define the Gaussian likelihood for $\mathbf{d} = (R - 1, \Delta\rho/\rho)$:

$$\mathcal{L}(A, G) \propto \exp\left(-\frac{1}{2}(\mathbf{d} - \mathbf{m}(A, G))^T \Sigma^{-1}(\mathbf{d} - \mathbf{m}(A, G))\right),$$

with $\mathbf{m}(A, G) = (c_R G, \kappa_A A + \bar{\kappa}_G G)$ and $\Sigma = \text{diag}(\sigma_R^2, \sigma_\rho^2)$. For the reduction hypothesis $\mathcal{H}_0 : (A, G) = (0, 0)$ versus $\mathcal{H}_1 : (A, G) \neq (0, 0)$, compute the likelihood ratio $\Lambda = 2 \log[\mathcal{L}(\hat{A}, \hat{G})/\mathcal{L}(0, 0)]$. Using a χ^2_2 calibration at level α , accept $\mathcal{H}_0$ (reduction holds) if $\Lambda < \tau$ and favor $\mathcal{H}_1$ otherwise. Operational thresholds are summarized in [Appendix G](#).

Error model and confidence regions. Digitization resolution, axis/scale systematics, counter Allan deviation, temperature coefficient (TK), pressure, and clamp terms are logged in the format of [Appendix F](#) and combined as

$$\sigma_R^2 = \sigma_{\text{dig}}^2 + \sigma_{\text{geom}}^2 + \sigma_{\text{rep}}^2, \quad \sigma_\rho^2 = \sigma_{\text{cnt}}^2 + \sigma_{\text{TK}}^2 + \sigma_{\text{env}}^2.$$

The joint $(1 - \alpha)$ CI for (A, G) is reported as

$$\{(A, G) : (\hat{\mathbf{d}} - \mathbf{m}(A, G))^\top \Sigma^{-1} (\hat{\mathbf{d}} - \mathbf{m}(A, G)) \leq \chi_2^2(1 - \alpha)\}.$$

First-order summary bounds use [\(48\)](#) and [\(50\)](#).

Decision rule (summary). (a) R -channel: $|R - 1| \leq \delta_R \Rightarrow |G| \leq C_R \delta_R$. (b) Resonator: $|\Delta\rho/\rho| \leq \delta_\rho \Rightarrow |A| \leq C_\rho \delta_\rho$, $|G| \leq C'_\rho \delta_\rho$. (c) If $\Lambda < \tau$, continuity (Maxwell reduction) holds. (d) If not, report upper bounds with logs ([Appendix F](#)) and upgrade design (angle pairs, repeats, stability) per [Section 6](#).

Context (role of this section). Equations [\(47\)](#)–[\(50\)](#) form the minimal set that quantifies how tightly the curvature–EM coupling reduces to Maxwell under $|\nabla\Phi| < \varepsilon$. They connect directly to the $\{A, G\}$ reporting format ([Appendix G](#)) and to the experimental checklist ([Appendix F](#)), and are used verbatim in the data comparison of [Section 6](#).

2.6 Summary: What This Section Established

- Starting from the action [\(14\)](#) and the constitutive triple [\(20\)](#), we obtained the field equations $\nabla_\nu H^{\mu\nu} = J^\mu$ and $\square\Phi - U'(\Phi) = \dots$. The Bianchi identity $dF = 0$ holds by definition (see [Section 1.3](#)), and current conservation $\nabla_\mu J^\mu = 0$ follows from gauge invariance ([Section 2.4](#)).
- The minimal expansion [\(21\)](#) of $\chi(\Phi, \nabla\Phi)$ introduces two data-facing, dimensionless parameters: an isotropic response $A \equiv \kappa_A \alpha \langle \Phi \rangle$ and an anisotropic response $G = \eta |\nabla\Phi|$. Their mapping to observables—the tilt–period ratio R and the mode ratio ρ —is summarized in [Section 2.3](#).
- The *continuity-to-Maxwell reduction* under a weak–gradient window $|\nabla\Phi| < \varepsilon$ is formalized in [Section 2.5](#) via the criterion [\(47\)](#). Channel-wise bounds follow directly from the linearized relations [\(48\)](#) for R and [\(50\)](#) for ρ . Joint assessment uses the likelihood ratio Λ with covariance-aware weighting, as specified in the same section.
- In the isotropic limit we recover standard Maxwell duality. For small anisotropic perturbations, causality and energy-positivity constraints are maintained; a compact checklist is provided in [Appendix D](#).
- The angle/polarization geometry leading to the characteristic $\cos 2\theta$ anisotropy and its alignment with experimental tilt/polarization settings are compiled in [Appendix E](#). Metrology/stability models and logging templates are provided in [Appendix F](#), and reproducible reporting forms and thresholds in [Appendix G](#).
- For later sections that pool measurements across platforms, we adopt the *dimensionless normalization* ([Appendix A](#)), i.e. $\bar{\Phi} = \Phi/\Phi_0$ and $\kappa_A=1$ unless stated. Under this convention, representative 95% bounds reported in the data sections read $|A| \lesssim 3 \times 10^{-3}$

and $|G| \lesssim 5.5 \times 10^{-3}$ (see [Section 6.2](#) and [Section 6.4](#)). The following table provides explicit *back-solved* examples for α and η under typical scale choices; platform-specific $\langle\Phi\rangle$ and $|\nabla\Phi|$ may be substituted as appropriate.

Table 4: Back-solved numeric examples from reported bounds (all cells populated). Assumptions: dimensionless normalization ($\bar{\Phi} = \Phi/\Phi_0$), $\kappa_A = 1$. *Numbers are back-solved from the 95% bounds in [Sec. 6.2](#) and [Sec. 6.4](#); different $\langle\Phi\rangle$, $|\nabla\Phi|$, or κ_A rescale the rightmost column linearly.*

Channel & bound (95%)	Scale choice	Implication (numeric)	Back-solve target	Result (numeric)
ρ (long-term stability)	$\langle\Phi\rangle = 1$	$ A \leq 3.0 \times 10^{-3}$	α	$\alpha \leq 3.0 \times 10^{-3}$
$ A \leq 3.0 \times 10^{-3}$	$\langle\Phi\rangle = 1.0 \times 10^{-1}$	$ A \leq 3.0 \times 10^{-3}$	α	$\alpha \leq 3.0 \times 10^{-2}$
$(A = \alpha\langle\Phi\rangle)$	$\langle\Phi\rangle = 1.0 \times 10^{-2}$	$ A \leq 3.0 \times 10^{-3}$	α	$\alpha \leq 3.0 \times 10^{-1}$
R (tilt ratio aggregate)	$ \nabla\Phi = 1.0 \times 10^{-2}$	$ G \leq 5.5 \times 10^{-3}$	η	$\eta \leq 5.5 \times 10^{-1}$
$ G \leq 5.5 \times 10^{-3}$	$ \nabla\Phi = 1.0 \times 10^{-3}$	$ G \leq 5.5 \times 10^{-3}$	η	$\eta \leq 5.5$
$(G = \eta \nabla\Phi)$	$ \nabla\Phi = 3.0 \times 10^{-4}$	$ G \leq 5.5 \times 10^{-3}$	η	$\eta \leq 1.833 \times 10^1$

Normalization note: $A \equiv \kappa_A \alpha \langle\Phi\rangle$ with $\kappa_A=1$ under the dimensionless convention. If a different unit choice fixes $\kappa_A \neq 1$, rescale the rightmost column by $1/\kappa_A$ for α .

Section roadmap. [Section 3](#) instantiates these definitions in data handling (windowing, covariance, and nulls); thereafter [Section 6.2](#) and [Section 6.4](#) report the channel-wise constraints that feed the back-solved examples above, and [Appendix H](#) compiles the joint (A, G) confidence geometry for cross-platform comparison.

2.7 A Unitary Curvature-Channel Operator for Wavefields (from Electromagnetism to Generic Waves)

Scope. Building on the summary in [Section 2.6](#), this subsection formalizes a *unitary, phase-only* operator that we have validated in the electromagnetic setting and then extend, under minimal assumptions, to generic wavefields (optical/radio, etc.). The governing field equations (Maxwell/wave) are left unchanged; instead, a *structural operator* acts on observables so that *auto power (two-point statistics) is nearly conserved* while *morphology and cross measures* respond in a controlled way.

Definition (unitary curvature-channel operator). For a complex wave envelope a ,

$$a_{\text{out}} = e^{-is\hat{G}} a, \quad \hat{G} \equiv \frac{1}{2}(K \circ M + M \circ K), \quad 0 < s \ll 1. \quad (51)$$

Here K is a Hermitian band-filter (Fourier window) and M is a real *curvature/filament gate* (e.g., masks built from phase curvature, ridge/skeleton, or level-set geometry). If $\hat{G}$ is self-adjoint, then $e^{-is\hat{G}}$ is unitary and $\|a_{\text{out}}\|_2 = \|a\|_2$; total power is preserved while morphology is remapped.

Link to electromagnetism (Maxwell reduction compatibility). With the (A, G) parameterization in [Section 2.3](#) and the weak-gradient bound $|\nabla\Phi| < \varepsilon$,

$$H = \chi(\Phi, \nabla\Phi) : F \Rightarrow R = 1 + c_R G + O(G^2), \quad \Delta\rho/\rho = c_\rho^{(A)} A + c_\rho^{(G)} G + \dots$$

Real-data cross-checks yield

$$|R - 1| < 0.55\% \text{ (95\% C.I.)}, \quad |\Delta\rho/\rho| \lesssim 3 \times 10^{-3} \text{ (long-term operation)},$$

supporting Maxwell reduction at the $\gtrsim 98\%$ level in practice and meeting an operational $\geq 95\%$ agreement target (goodness-of-fit, confidence coverage, residual budgets). Uncertainty synthesis and metrology follow [Appendix F](#) and [Appendix G](#).

Small-signal expansion and diagnostic fingerprints.

$$a_{\text{out}} = a - is \hat{G}a + O(s^2), \quad \|a_{\text{out}}\|_2^2 = \|a\|_2^2 + O(s^2).$$

Thus two-point statistics (auto power) vary only at $O(s^2)$, whereas morphology/cross diagnostics (Minkowski V_1, V_2 , skeleton length/orientation, and the $\cos 2\theta$ harmonic in angle channels) respond at $O(s)$. The observed $\cos 2\theta$ residue in EM tests follows from the relative angle between the geometric axis of M and the experimental drive ([Section 2.3](#), [Appendix E](#)).

Extension to generic wavefields (domain-invariant principle). Equation (51) acts on the solution space of *any* linear wave equation as a covariant post-operator. When changing domain, K becomes the appropriate resolution/window for that domain, and M is replaced by the domain’s curvature/filament gate (e.g., phase-curvature maps, level-set curvature, or phase-only masks). The structural consequences remain the same: (i) power preservation, (ii) $O(s^2)$ invariance of auto spectra, (iii) $O(s)$ sensitivity of morphology and cross channels.

Verification metrics (operational checklist). (i) *Power preservation*: $\|a_{\text{out}}\|_2/\|a\|_2 = 1 \pm \epsilon$ with $\epsilon \leq 1\%$. (ii) *Angle channel*: linear $\cos 2\theta$ fit for $R(\theta)$ with $R^2 \geq 0.95$ and a CI-based bound on G ([Section 2.5](#)). (iii) *Morphology set*: $\Delta V_1, \Delta V_2$, skeleton length/orientation attaining match ≥ 0.95 to predictions. (iv) *Stability*: long-term ρ stability $\lesssim 3 \times 10^{-3}$. All metrics are reported with uncertainty models in [Appendix F](#).

Summary and outlook. This subsection establishes an experimental frame in which the *laws* (Maxwell/wave) are preserved, while a unitary structural operator makes explicit which observables are *nearly invariant* (auto) and which respond *coherently* (morphology/cross). The electromagnetic results (sub-percent R , 10^{-3} -level ρ) are consistent with this frame, and the same principle carries to generic wavefields. Detailed operator design (K, M per domain), analytic links to morphology statistics, and systematic optical/radio benchmarks will be developed in the forthcoming study titled “*Curvature Field Identity: Curvature z -Axis Extension of the Complex Number $x + iy + sz$ and a New Structural Operator.*”

3. Data-driven validation: mapping geometry to observables and bounds

Chapter overview. We map the isotropic A and anisotropic G components defined in [Section 2.5](#) directly to *public, peer-reviewed* datasets. Two channels are used: (i) the

tilt-angle dependence of quantum oscillations/Aharonov–Bohm as a normalized ratio R (the R -channel), and (ii) the resonator TE/TM mode ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ (the ρ -channel). Our objective is to place *separable bounds* on (A, G) by aligning both channels to a common scale. Crucially, we *hybridize* sources: we use (a) high-fidelity digitization from original figures, and (b) *numeric source tables/supplementary data* provided by the papers themselves. The uncertainty model, reproducibility log, and cross-checks are given in [Appendix F](#) and [Appendix G](#). Representative numeric-table sources include persistent-current distributions [61] and WGM-resonator stability indicators [26, 27]; digitized angle/period sweeps include graphene-ring data [57].

The outputs of this chapter are: (1) *within-channel* bounds for the R - and ρ -channels, and (2) *cross-channel* combined bounds via the mapping rules of [Section 2.6](#). Summary tables and per-source validation figures are collated in [Appendix G](#).

Policy: numeric-first, figure-assisted. We prioritize Tier-N sources (public tables or raw time series) built on a *common timebase* for all bounds. **All key metrics $|R-1|$, $|\Delta\rho/\rho|$, and the (A, G) bounds are computed *exclusively* from DOI-linked numeric tables or raw logs (Tier-N). Digitized points (Tier-F) are retained only for shape cross-checks and are explicitly *illustrative; excluded from inference*.** Synchronization, variance combination, and Allan-deviation handling follow [Appendix F](#). This preserves every figure and formula already in this chapter while elevating numeric datasets to primary evidence.

Data tiers.

- Tier-N: DOI-linked numeric tables or raw logs. Used in inference; confidence intervals and bounds are computed *only* from these.
- Tier-F: minimally digitized points when Tier-N is unavailable (*illustrative; excluded from inference*; shown for shape/consistency checks and not used in CIs).

Common-timebase rule. When multiple streams exist (e.g., $\Delta B(\theta_i, t)$ and $f_{\text{TE/TM}}(t)$), we form $R(\theta_1, \theta_2; t)$ and $\rho(t)$ on a synchronized timebase before any aggregation. Window statistics use robust medians/trimmed means, and uncertainties combine within-window variance with overlapping Allan deviation as specified in [Appendix F](#).

3.1 Datasets, preprocessing, and digitization uncertainty

Inclusion criteria. We include only studies that (i) are peer-reviewed, (ii) specify axis ticks and units, (iii) provide repeated measurements of tilt angle θ or mode frequencies ($f_{\text{TE}}, f_{\text{TM}}$), and (iv) document system configuration (temperature, drive, specimen geometry) in tables. Under these criteria, R -channel data comprise AB oscillations in metallic/graphene rings and tilt-dependent quantum oscillations, while ρ -channel data comprise sapphire/dielectric resonators and WGMs ([61, 57, 26, 27]).

Numeric-first policy and common timebase. Public numeric tables or raw time series (Tier-N) are the *only* inputs used to compute confidence intervals and bounds. When multiple streams exist (e.g., $\Delta B(\theta_i, t)$ and $f_{\text{TE/TM}}(t)$), we construct $R(\theta_1, \theta_2; t)$ and $\rho(t)$

on a synchronized timebase before any averaging. Window statistics use robust medians/trimmed means, and uncertainties combine within-window variance with overlapping Allan deviation as in [Appendix F](#). Figure-derived points (Tier-F) are retained for *shape checks only* and are excluded from inference.

Hybrid use: numeric tables + graphics digitization. Numeric tables fix absolute scales (distribution cumulants, linewidths/comb spacings, Q), while high-resolution digitization recovers *relative* variations along angle/time axes to probe sensitivity to θ and window length. Consistency tests and weighting rules for fusing both routes are detailed in [Appendix F](#).

Preprocessing pipeline. (1) *Coordinate extraction:* from original figures, register grid intersections, estimate a pixel $\rightarrow$ physical transform, and obtain $\{\Delta B(\theta_i), \theta_i\}$ or $\{f_{\text{TE}}(t_j), f_{\text{TM}}(t_j)\}$. (2) *Axis calibration:* offsets/scales are jointly fit by least squares; units are unified; angles are in degrees. (3) *Repetitions/outliers:* repeated sweeps summarized by the median; segments with thermal jumps or hysteresis are excluded with rationale logged in [Appendix G](#). (4) *ρ construction:* with synchronous acquisition, $\rho(t_j) = f_{\text{TE}}(t_j)/f_{\text{TM}}(t_j)$ uses identical time windows to suppress common drifts; low-frequency trends removed by robust quartile regression ([Appendix F](#)).

Normalization metrics. For the R -channel, for each angle pair (θ_1, θ_2) ,

$$R = \frac{\Delta B(\theta_2)}{\Delta B(\theta_1)} \frac{\cos \theta_1}{\cos \theta_2}, \quad (52)$$

is computed. For the ρ -channel, the relative change $\delta\rho/\rho = (\rho - \rho_0)/\rho_0$ is used with ρ_0 the median over a reference window. The cross-channel mapping to (A, G) follows [Section 2.6](#).

Digitization uncertainty (Tier-F; illustrative only). Assign 1σ equal to one-fifth of the axis-tick spacing in pixels and model independent Gaussian errors for $\delta(\Delta B)$, $\delta\theta$, and δf . Error propagation gives

$$\frac{\delta R}{R} \simeq \sqrt{\left(\frac{\delta\Delta B_2}{\Delta B_2}\right)^2 + \left(\frac{\delta\Delta B_1}{\Delta B_1}\right)^2 + \tan^2\theta_1 \delta\theta_1^2 + \tan^2\theta_2 \delta\theta_2^2}, \quad (53)$$

$$\frac{\delta\rho}{\rho} \simeq \sqrt{\left(\frac{\delta f_{\text{TE}}}{f_{\text{TE}}}\right)^2 + \left(\frac{\delta f_{\text{TM}}}{f_{\text{TM}}}\right)^2}. \quad (54)$$

Tier-F points are *not* used in CI/bound computation; they are plotted only as sanity checks.

Confidence intervals and meta-averages (Tier-N only). Final CIs are obtained by nonparametric bootstrap with $N_{\text{bs}} = 10^4$; we report 2.5–97.5% intervals. For multi-study summaries, inverse-variance meta-averages are reported with random-effects diagnostics per [Appendix F](#).

Cross-checks and quality control. (1) *Mutual digitization*: two independent extractions must agree within $< 0.3\%$ relative deviation.
(2) *Angle perturbation test*: inject $\theta \rightarrow (\theta \pm 0.1^\circ)$ to probe R sensitivity.
(3) *Window-length test*: verify ρ bounds are stable for window lengths in $[10^2, 10^4]$ s.
(4) *A priori criteria*: segments failing $R^2 \geq 0.95$, significance $\geq 5\sigma$, or null-failure $\leq 1\%$ are excluded with reasons logged in [Appendix G](#).

Linked references (data sources). Numeric tables and cumulants for persistent currents: [61]; resonator stability/linewidth/comb spacing: [26, 27]; digitized angle/period sweeps: [57].

Table 5: Data inventory and usage policy (numeric-first; Tier-N only for inference).

ID	Channel	Src type	Tier	Tag	DOI/hash	Sync	Notes
D1	$R(\theta_1, \theta_2)$	numeric table / CSV	Tier-N	numeric-table	(DOI/sha256)	yes	Synchronized angle/period logs; used in CIs/bounds.
D2	$\rho = f_{\text{TE}}/f_{\text{TM}}$	counter logs / CSV	Tier-N	raw	(DOI/sha256)	yes	Allan-plateau windowed; drift model documented; used in CIs/bounds.
D3	R (legacy panel)	figure panel	Tier-F	digitized	(figure id)	n/a	<i>illustrative; excluded from inference; sanity check only.</i>
D4	ρ (legacy plot)	figure panel	Tier-F	digitized	(figure id)	n/a	<i>illustrative; excluded from inference; window noted.</i>
D5	calibration	temp./magnetometer logs	Tier-N	raw	(DOI/sha256)	yes	Angle calibration, reference lock, unit consistency.

3.1.1 Tier-N numeric data: synchronization and confidence intervals

Admissible inputs (data model). Let $\mathcal{D}$ denote a finite family of Tier-N datasets consisting of DOI-addressable tables or raw logs with time stamps and units in $\{\text{Hz}, \text{T}, \text{deg}\}$. Each $d \in \mathcal{D}$ is identified with a triple

$$(\Delta B(\theta_1, t), \Delta B(\theta_2, t), f_{\text{TE}}(t), f_{\text{TM}}(t))_d$$

satisfying: (i) unit consistency; (ii) unique time stamps on a measurable set; (iii) absence of duplicated rows on that set. Records that violate (i)–(iii) are treated as inadmissible (cf. [Appendix G](#)).

Common timebase and derived observables. Given any $d \in \mathcal{D}$, fix a discrete grid $\{t_k\}_{k=1}^{N_d}$ and define first-order (or nearest-neighbor) interpolants on this grid. On $\{t_k\}$ set

$$R_k \equiv \frac{\Delta B(\theta_2, t_k) \cos \theta_1}{\Delta B(\theta_1, t_k) \cos \theta_2}, \quad \rho_k \equiv \frac{f_{\text{TE}}(t_k)}{f_{\text{TM}}(t_k)}.$$

This realizes the tilt-projection normalization and the TE/TM ratio on a common timebase without introducing higher-order smoothing.

Drift handling and effective variance. Let $\mathcal{W} \subset \{1, \dots, N_d\}$ index a fixed analysis window (length τ_*). Define the window mean and variance

$$\bar{R}_d \equiv \frac{1}{|\mathcal{W}|} \sum_{k \in \mathcal{W}} R_k, \quad \sigma_{\text{win}, R; d}^2 \equiv \frac{1}{|\mathcal{W}| - 1} \sum_{k \in \mathcal{W}} (R_k - \bar{R}_d)^2,$$

and analogously for ρ . Low-frequency drift is removed by a window–median (or quartile–regression) detrend; the effective uncertainty is modeled as

$$\sigma_{R; d}^2 = \sigma_{\text{win}, R; d}^2 + \sigma_{A, R; d}^2(\tau_*),$$

where $\sigma_A^2(\tau)$ is the overlapping Allan variance at averaging time τ (definitions in [Appendix F](#)).

Intervals and Gaussian surrogate. From $\{R_k\}_{k \in \mathcal{W}}$ (resp. $\{\rho_k\}$) form the percentile bootstrap ($N_{\text{bs}} = 10^4$) 95% confidence interval for $\bar{R}_d$ and $\bar{\rho}_d$. When a Gaussian surrogate is adequate,

$$\bar{R}_d \pm z_{0.975} \frac{\sigma_{R; d}}{\sqrt{N_d}}, \quad \bar{\rho}_d \pm z_{0.975} \frac{\sigma_{\rho; d}}{\sqrt{N_d}}$$

provides an equivalent interval (with $N_d = |\mathcal{W}|$).

Aggregation across datasets. Let weights $w_{R; d} \equiv 1/\sigma_{R; d}^2$ (analogously $w_{\rho; d}$). Define inverse–variance means

$$\bar{R}_{\text{meta}} = \frac{\sum_{d \in \mathcal{D}} w_{R; d} \bar{R}_d}{\sum_{d \in \mathcal{D}} w_{R; d}}, \quad \bar{\rho}_{\text{meta}} = \frac{\sum_{d \in \mathcal{D}} w_{\rho; d} \bar{\rho}_d}{\sum_{d \in \mathcal{D}} w_{\rho; d}},$$

and quantify heterogeneity by Q, I^2, τ^2 (DerSimonian–Laird). If $I^2 > 50\%$, adopt random–effects estimates (see [Appendix F](#)).

Mapping to A and G . With the first–order sensitivity maps of [Section 2.6](#),

$$|G| \leq \frac{z_{0.975} \sigma_{R, \text{meta}}}{|c_R|}, \quad |A| \leq \frac{z_{0.975} \sigma_{\rho, \text{meta}}}{|c_\rho^{(A)}|},$$

where c_R (geometry) and $c_\rho^{(A)}$ (mode overlap) are specified in [Appendix E](#). If common–mode anisotropy is nonnegligible, augment the surrogate for ρ by $c_\rho^{(G)} G$ and refit.

Assumptions and checks.

1. Unit/spec conformity and hash–identified provenance (SHA–256) as in [Appendix G](#).
2. Outlier control: samples with $|x - \text{median}(x)| > 5\sigma$ may be excluded; sensitivity to this rule is tabulated (with/without).
3. Stability under analysis choices: the bounds for A, G vary by at most 10% under window–length or down–sampling changes within the admissible range.

3.2 R -channel: tilt-normalized ratio and anisotropy bound

Definition and expectation. From the oscillation period measured at tilt angle θ , $\Delta B(\theta)$, define

$$R \equiv \frac{\Delta B(\theta_2)}{\Delta B(\theta_1)} \frac{\cos \theta_1}{\cos \theta_2} \quad (55)$$

so that Maxwell reduction predicts $R \rightarrow 1$. If curvature-induced anisotropy is present, the deviation scales as $|R - 1| \sim C_R G + O(G^2)$, where $C_R = \mathcal{O}(1)$ depends on geometry/material (Section 2.3, Section 2.6).




Figure 1: Schematic of R at two tilts θ_1, θ_2 . The plot highlights the role of the $\cos \theta$ normalization and the $\tan \theta$ sensitivity (error propagation). The definition of R is given in Eq. (55); the propagation formula is summarized in Appendix H.

Hybrid evidence: numeric tables + digitized sweeps (with hard boundary). We jointly use (i) DOI-linked numeric source tables (Tier-N) that set the absolute scale and (ii) high-fidelity digitization (Tier-F) that recovers *relative* variations across θ . **Digitized sweeps (Tier-F) are restricted to *shape tracking* (e.g., the θ -response curve) and are *not used* for quantitative bounds or confidence intervals; all CIs/bounds are computed from numeric tables (Tier-N) only.** Numeric examples include normalized cumulants ($\kappa_3 \dots \kappa_6$) of persistent currents in metallic rings [61]; digitized angle/period sweeps (e.g., graphene rings) provide $\Delta B(\theta)$ over widely separated tilts [57]. The uncertainty model and fusion rules are in Appendix F; per-dataset logs and QA are in Appendix G.

Linearization and error propagation. For small angle/readout errors $\delta\theta_i \ll 1$,

$$R - 1 \approx \frac{\delta\Delta B_2}{\Delta B_2} - \frac{\delta\Delta B_1}{\Delta B_1} + \tan\theta_1 \delta\theta_1 - \tan\theta_2 \delta\theta_2, \quad (56)$$

hence large $|\tan\theta|$ amplifies angle bias. We therefore select θ_1, θ_2 with comparable $|\tan\theta|$ (typically $20^\circ \lesssim \theta \lesssim 60^\circ$), and treat extreme tilts in a dedicated sensitivity pass. *Note:* the *digitization rule* in Appendix F (one-fifth to one-third of a tick spacing as 1σ) is **display-only for Tier-F** and is **not propagated** into CI/bounds from Tier-N (Appendix F).

Per-platform estimates and combination. Datasets are grouped into metallic AB rings, graphene rings, and bulk tilt quantum oscillations. For each dataset k we estimate R_k with variance σ_k^2 , and form the precision-weighted mean

$$\langle R \rangle = \frac{\sum_k w_k R_k}{\sum_k w_k}, \quad w_k = \sigma_k^{-2},$$

with a random-effects correction to guard against underdispersion (Appendix G).

Numeric-source cross-check (distribution proxy). From [61], the normalized cumulants $\kappa_3, \dots, \kappa_6$ are individually consistent with zero within their 95% confidence intervals (see Appendix G), supporting the Maxwell reduction and calibrating the absolute scale used for R .

Summary number (normalization accuracy; Tier-N only). The combined result is

$$\langle R \rangle = 1.0000 \pm 0.0028 \ (1\sigma) \quad \Rightarrow \quad \mathbf{95\% \ C.I. : } |R - 1| < \mathbf{0.55\%},$$

i.e., the standard tilt-projection law holds at sub-percent precision.

Translation to an anisotropy bound. Using (55)–(56) together with the baseline scaling $\Delta B(\theta) \propto 1/\cos\theta$,

$$G \lesssim \frac{|R - 1|_{\max}}{C_R} \approx \mathcal{O}(5.5 \times 10^{-3}) \quad (57)$$

for a conservative $C_R \simeq 1$. Thus the $\cos 2\theta$ -type anisotropy is $\lesssim 6 \times 10^{-3}$ at current sensitivity.

Robustness checks. The bound remains stable within $\pm 0.1\%$ under: (i) swapping the angle pair (θ_1, θ_2) ; (ii) subset bootstrap resampling ($N_{\text{bs}} = 10^4$); and (iii) halving the digitization resolution (more conservative noise). Separate fits to AB-ring and bulk-tilt subsets yield consistent central values (Appendix G).

Systematics control. Field-scale offsets cancel in the ratio R . Angle bias enters with $\tan\theta$, hence an angle-calibration log is required (Appendix G). Hysteresis/jump events are masked by predeclared rules; a sensitivity table reports the impact of masking choices (Appendix F).

Interpretation. The sub-percent agreement $R \rightarrow 1$ satisfies the reduction criterion in [Section 2.5](#) and yields the anisotropy bound (57). Combined with the ρ -channel, this tightens to an effective $\mathcal{O}(10^{-3})$ constraint (see [Section 3.4](#)).

Add-on (numeric-first inference on a common timesbase). With synchronized periods $\Delta B(\theta_i, t)$ we form

$$R(\theta_1, \theta_2; t) = \frac{\Delta B(\theta_2, t) \cos \theta_1}{\Delta B(\theta_1, t) \cos \theta_2},$$

aggregate by robust medians, and propagate uncertainty by combining within-window variance and overlapping Allan deviation ([Appendix F](#)). Linearizing $R - 1 = c_R G + O(G^2)$ yields the 95%-CI bound

$$|G| \leq \frac{z_{0.975} \sigma_R}{|\hat{c}_R|}, \quad (z_{0.975} \approx 1.96).$$

Table 6: R -channel numeric summary and bound (confidence intervals use **Tier-N only**).

Dataset	median($R - 1$)	σ_R	$ \hat{c}_R $	95% bound on G / Notes
D1 (numeric)	0.002	0.0055	1.00	Based on 95%-CI width of $R - 1$: $ G \lesssim 5.5 \times 10^{-3}$ (Tier-N).
D3 (figure)	0.004	0.012	1.05	Tier-F; <i>shape-only</i> . Excluded from inference ; shown for visual cross-checks.

Angle-sweep protocol in $X(\theta) = \cos 2\theta$ (grid $\leq 1^\circ$, bi-directional; preregistered). For the R -channel we preregister a first-order linear model in $X(\theta) \equiv \cos 2\theta$:

$$R(\theta) = 1 + \beta X(\theta) + \varepsilon, \quad \beta = c_R G.$$

Fix the angle grid as $\theta_k = \theta_{\min} + k \Delta\theta$ with $\Delta\theta \leq 1^\circ$. At each θ_k , perform *same-window* bi-directional sweeps (upward $\uparrow$, downward $\downarrow$) and define

$$H(\theta_k) = |R_\uparrow(\theta_k) - R_\downarrow(\theta_k)|.$$

Pass criterion: $H(\theta_k) \leq z_{0.995} \sigma_R$ (null-failure $\leq 0.5\%$). *Preregistered linearity/significance thresholds:*

$$R^2(R \text{ vs. } X) \geq 0.95, \quad \frac{|\hat{\beta}|}{\text{SE}(\hat{\beta})} \geq 5 \quad (5\sigma).$$

Regression uncertainties combine *same-window* variance with overlapping Allan deviation ([Appendix F](#)). To reduce angle-bias amplification, pair windows so that $|\tan \theta_1| \approx |\tan \theta_2|$, and report extreme tilts separately as a sensitivity sweep.

QA log. For each θ_k , store $(R_\uparrow, R_\downarrow, H(\theta_k))$, window IDs, and pass/fail flags in the [Appendix G](#) preregistration log. Uncertainties follow the same-window + overlapping Allan procedure of [Appendix F](#) and the numeric-first pipeline in [Section 3.2](#) (Add-on).

3.2.1 Tier–N (numeric) analysis and CI construction

Common timebase and angle pairing. For studies providing time-stamped $\Delta B(\theta_i, t)$, we construct $R(t)$ on a synchronized grid and preselect pairs (θ_1, θ_2) with matched $|\tan \theta|$ to minimize differential angle amplification (cf. (56)).

Within-study estimate. Per study d , window medians yield $\bar{R}_d$ with effective standard deviation $\sigma_{R,d}$ combining within-window variance and overlapping Allan variance at the window length (Appendix F). We report both bootstrap 95% CIs and Gaussian CIs for transparency.

Across-study synthesis. We form an inverse-variance mean $\langle R \rangle_N$ over Tier–N sources only, compute heterogeneity statistics Q , I^2 , and switch to random-effects when $I^2 > 50\%$. Tier–F (digitized) points are shown in figures but excluded from the synthesis.

Bound translation and sensitivity. With C_R taken from the geometry model (Appendix E), the Tier–N bound on G follows from $|R - 1|$ at 95% C.I. via $G \leq |R - 1|/C_R$. Sensitivity tables summarize the effect of (a) window length, (b) down-sampling, and (c) excluding each study in turn (leave-one-out), with all changes constrained to $\pm 10\%$ of the headline bound.

3.3 ρ -channel: TE/TM mode ratio and separation of isotropic/anisotropic parts

Stabilized counter window (numeric-first). From synchronized counters we form

$$\rho(t) \equiv \frac{f_{\text{TE}}(t)}{f_{\text{TM}}(t)}.$$

Inference uses *only* Tier–N streams on a common timebase and selects the window where the overlapping Allan deviation $\sigma_y(\tau)$ is stationary (near its minimum). Within that window we fit

$$\frac{\Delta \rho}{\rho} = c_\rho^{(A)} A + c_\rho^{(G)} G + \varepsilon, \quad \text{Cov}[\varepsilon] = \Sigma_{\text{ctr}}, \quad (58)$$

using the full counter covariance Σ_{ctr} . **Figure-derived Tier–F points (if any) are retained for *shape checks only* and are excluded from confidence intervals, bounds, weighting, fitting, and meta-combination (Section 3.1, Appendix F).**

Add-on (dispersion covariate in the ρ regression). Let $\delta\omega(t) \equiv \omega(t) - \omega_0$ with $\omega_0 = \text{median}\{\omega(t)\}$ inside the window. Using Eq. (39), the first-order dispersion difference between TE/TM induces

$$\Delta \ln \rho = \Delta \ln f_{\text{TE}} - \Delta \ln f_{\text{TM}} \simeq s_A A + \tilde{s}_G(\theta) G + b_\rho \beta \delta\omega + n_\rho,$$

where $b_\rho \equiv d_{\text{TE}} - d_{\text{TM}}$ collects the mode-contrast (see Appendix E). We therefore extend Eq. (58) to

$$\frac{\Delta \rho}{\rho} = c_\rho^{(A)} A + c_\rho^{(G)} G + b_\rho^{(\text{eff})} \delta\omega + \varepsilon, \quad \text{Cov}[\varepsilon] = \Sigma_{\text{ctr}}, \quad (59)$$

and include $\delta\omega$ as a *centered* covariate (zero mean in the window) to reduce collinearity and absorb linear-in-frequency residuals. We report HC-robust errors and variance-inflation factors (VIFs); if $|b_{\rho}^{(\text{eff})}| < 2\sigma$, the term is dropped in the primary fit and retained as a preregistered sensitivity. The dispersion map and mode factors are summarized in [Appendix E](#); uncertainty propagation follows [Section 4.4](#).

Dual-window reporting (short + long; removing window dependence). Under the *numeric-first* policy, we form $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$ only on a *common timescale*. The *long* window summarizes the operational envelope including drift/flicker, whereas the *short* window states the *instrumental* limit from comb-beat/linewidth. Reporting both *side by side* removes window dependence and links model fitting ([Section 2.6](#)) consistently to the noise budget ([Section 4.4](#)). In long windows we target $|\Delta\rho/\rho|_{\text{long}} \lesssim 3 \times 10^{-3}$; in short windows we show the synchronized comb-beat/linewidth limit $|\Delta\rho/\rho|_{\text{short}} \lesssim \Delta f_{\text{beat}}/f_{\text{comb}} \sim 10^{-9}$. For each value, we print the averaging time/window length τ and the derivation path (counter vs. comb) ([Appendix G](#)).




Figure 2: ρ -channel exemplars across platforms. *Illustrative; excluded from inference. Markers: solid=Tier-N, hollow=Tier-F. Window=[$10^2, 10^4$] s; Tier-N only for CIs/bounds.*

Table 7: Dual-window reporting for the ρ channel (window length and derivation path shown). $Window=[10^2, 10^4] s$ (Allan-min region). *Tier-N only for CIs/bounds.*

Window class	τ (averaging time)	Path	$ \Delta\rho/\rho $ (95% CI)	Notes / computation rule
Short	$10^0 \sim 10^2 s$	Comb beat/linewidth	$\lesssim 10^{-9}$	Synchronous; computed via $\Delta f_{\text{beat}}/f_{\text{comb}}$. Assumes drift removal.
Long	$10^3 \sim 10^5 s$	Counter (common ref.)	$\lesssim 3 \times 10^{-3}$	Window near Allan minimum; include process logs and drift corrections (T/P/clamping).

Reporting template (windowed). State the Allan-stabilized averaging time τ_* and the CI convention: (1) separate 95% CIs from single-parameter fits with the other fixed to zero; and (2) joint 95% CIs from the two-parameter covariance (see Section 3.4). We also report the short-window limit $|\Delta\rho/\rho|_{\text{short}} \lesssim 1.6 \times 10^{-9}$ (from numeric linewidth/comb-spacing), and the conservative long-window envelope $|\Delta\rho/\rho|_{\text{long}} \lesssim 3 \times 10^{-3}$, explicitly labeling the time window used in inference.

Definition and sensitivity. Let $\rho \equiv f_{\text{TE}}/f_{\text{TM}}$. The isotropic component A redefines the impedance *with the same sign* for both modes, yielding similar shifts in $\ln f_{\text{TE}}$ and $\ln f_{\text{TM}}$. By contrast, the anisotropic component G couples differently to TE/TM because of distinct field distributions and boundary couplings. In linear response,

$$\Delta \ln \rho = \Delta \ln f_{\text{TE}} - \Delta \ln f_{\text{TM}} \simeq s_A A + s_G(\theta) G + n_\rho, \quad (60)$$

where s_A is the isotropic (impedance-type) sensitivity, $s_G(\theta) = \tilde{s}_G \cos 2\theta$ is the anisotropic (angle-dependent) sensitivity, and n_ρ collects residual drift/noise. Mode-overlap expressions for s_A, s_G are summarized in Section 2.3 and Appendix E. In what follows we use study-specific sensitivities s_A and $s_G(\theta)$ (and c_R for the angle channel) as summarized with numerical exemplars and uncertainty rules in Section 2.3.1.

Operating goal and data summary (numeric-first, figure-assisted). We use (i) *numeric source tables* to calibrate absolute scales (linewidths, comb spacings, Q), and (ii) *high-resolution digitization* only for relative variations where tables are absent (Appendix F). Long-term, slow-drift-inclusive reports in the literature consistently indicate

$$|\Delta\rho/\rho| \lesssim 3 \times 10^{-3},$$

which we adopt as a conservative *long-term* envelope [27, 26, 22, 23].

Numeric-source cross-check (short-term window). From WGM comb data: microwave beat linewidth $\lesssim 40$ Hz with 25 GHz comb spacing gives

$$|\Delta\rho/\rho|_{\text{short}} \lesssim \frac{40 \text{ Hz}}{25 \text{ GHz}} = 1.6 \times 10^{-9},$$

and the intrinsic optical limit is $1/Q \simeq 4.0 \times 10^{-10}$ for $Q \approx 2.5 \times 10^9$. These are computed directly from numeric values in the sources; calculation logs/tables appear in Appendix G. The short-term bound *encloses* the long-term meta-value once the time window is stated.

Separation and identification (TE/TM differential). Because ρ cancels common-mode drift at first order while TE/TM respond differently, it provides complementary handles on A and G . Operationally,

$$\begin{aligned}
\text{(i) Common-differential split: } & \Delta \ln f_{\text{com}} = \frac{1}{2}(\Delta \ln f_{\text{TE}} + \Delta \ln f_{\text{TM}}), \\
& \Delta \ln \rho = \Delta \ln f_{\text{TE}} - \Delta \ln f_{\text{TM}}; \\
\text{(ii) Regression: } & \Delta \ln f_{\text{com}} \mapsto A \quad (\text{mostly isotropic}), \\
& \Delta \ln \rho \mapsto (A, G) \quad (\text{differential sensitivities}).
\end{aligned} \tag{61}$$

We jointly estimate A, G from the covariance of the two regressions. Identifiability improves as the condition number of the 2×2 sensitivity matrix decreases (i.e., larger TE/TM overlap contrast); see [Section 3.4](#).

Stability requirements and corrections. Sub-percent constraints require: (i) *synchronous* acquisition of $f_{\text{TE}}, f_{\text{TM}}$; (ii) linear/quadratic correction using temperature/drive logs; and (iii) operation near the Allan-deviation minimum $\sigma_y(\tau)$ at $\tau \sim 10^3 - 10^4$ s. Targets:

$$\sigma_y^{\text{TE, TM}}(10^3 \text{ s}) \lesssim 10^{-12}, \quad \left| \Delta \rho / \rho \right|_{\text{drift}} \lesssim 3 \times 10^{-3},$$

with QA templates in [Appendix G](#) and model links in [Appendix C](#).

Numerical translation of bounds (long/short side-by-side). Using the conservative *long-term* envelope $|\Delta \rho / \rho|_{\text{max}} = 3 \times 10^{-3}$ and $|s_A|, \max_{\theta} |s_G(\theta)| \simeq 1$,

$$|A| \lesssim 3 \times 10^{-3}, \quad |G| \lesssim \mathcal{O}(10^{-2}), \tag{62}$$

while, for a *short-term* synchronized window,

$$|A|_{\text{short}} \lesssim 1.6 \times 10^{-9}, \quad |G|_{\text{short}} \lesssim \mathcal{O}(10^{-9}) \quad (\text{assuming } \mathcal{O}(1) \text{ sensitivities}),$$

with the time window explicitly stated. Combined with the R -channel constraint $|R - 1| < 0.55\%$, the effective long-term bound on G compresses to $\mathcal{O}(10^{-3})$ ([Section 3.2](#), [Section 3.4](#)).

Systematics control. Reference drift and residual thermal noise can leave second-order imprints on ρ . We mitigate by operating at a temperature-inversion point (TK-zero), applying joint feed-forward (temperature, vacuum, clamping force), and periodic reference calibration ([Appendix G](#)). Mode misidentification (TE/TM cross-talk) is precluded through mode maps and polarization scans ([Appendix E](#)).

Summary. The ρ -channel cancels common-mode drift yet retains differential sensitivity, enabling separation of A and G . With current *long-term* public data we obtain $|A| \lesssim 3 \times 10^{-3}$ and $|G| \lesssim \text{few} \times 10^{-3} - 10^{-2}$, while *short-term* numeric windows yield bounds orders of magnitude tighter. Joint estimation with the R -channel drives G to the $\mathcal{O}(10^{-3})$ level ([Section 3.4](#)).

3.3.1 Tier–N workflow and confidence construction

Synchronized windows and detrending. From time-stamped $f_{\text{TE}}(t), f_{\text{TM}}(t)$, form $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$ on a common timebase. Apply robust quartile regression to remove low-frequency drift (Appendix F).

Within-study estimates. For study d , compute window medians and an effective standard deviation combining in-window variance with overlapping Allan deviation at the chosen τ . Report both bootstrap (2.5–97.5%) and Gaussian CIs.

Cross-study synthesis. Aggregate Tier–N sources with inverse-variance weights. Compute heterogeneity Q and I^2 ; if $I^2 > 50\%$, switch to a random-effects mean. Tier–F (digitized) points are shown only for shape checks and excluded from inference.

Mapping to (A, G) and sensitivity tests. Use (60) with study-specific $s_A, s_G(\theta)$ from mode overlaps to obtain $(\hat{A}, \hat{G})$ and their covariance. Test robustness to window length, downsampling, and leave-one-out; summarize impact in a sensitivity table with predeclared thresholds (Appendix F, Appendix G).

3.4 Joint estimation: A – G covariance and final numbers

Policy (numeric-first, figure-assisted). *All inference and confidence intervals use Tier–N inputs (public tables/raw logs) only. Figure-derived Tier–F points are flagged as *shape checks* and are excluded from inference (Section 3.1, Appendix F). For any reported conclusions, bounds, weighting, fitting, or meta-aggregation, Tier–F materials are not included. Only Tier–N enters the likelihood, Σ , and any meta-analytic combination.*

Channel stacking and covariance. For each channel $\mathcal{C} \in \{R, \rho\}$ define the deviation from Maxwell,

$$\Delta O_{\mathcal{C}} \equiv O_{\mathcal{C}}^{\text{meas}} - O_{\mathcal{C}}^{\text{Mxw}} \simeq M_{\mathcal{C}}^A A + M_{\mathcal{C}}^G G + n_{\mathcal{C}},$$

and stack

$$\mathbf{d} \equiv \begin{bmatrix} R - 1 \\ \Delta\rho/\rho \end{bmatrix}, \quad \mathbf{M} \equiv \begin{bmatrix} M_R^A & M_R^G \\ M_\rho^A & M_\rho^G \end{bmatrix}, \quad \Sigma \equiv \text{diag}(\sigma_R^2, \sigma_\rho^2), \quad \boldsymbol{\theta} \equiv \begin{bmatrix} A \\ G \end{bmatrix}.$$

Sensitivities follow Section 2.3/Section 2.6: $M_R^A \approx 0$, $M_R^G \neq 0$ (tilt anisotropy only); $M_\rho^A, M_\rho^G \neq 0$ (common+differential response). Tier–N supplies $\mathbf{d}, \Sigma$; Tier–F is used only for visual consistency.

Estimator and covariance. The weighted least-squares (WLS) solution and parameter covariance are

$$\hat{\boldsymbol{\theta}} = (\mathbf{M}^\top \Sigma^{-1} \mathbf{M})^{-1} \mathbf{M}^\top \Sigma^{-1} \mathbf{d}, \quad \text{Cov}(\hat{\boldsymbol{\theta}}) = (\mathbf{M}^\top \Sigma^{-1} \mathbf{M})^{-1}. \quad (63)$$

Uncertainties σ_R, σ_ρ include axis calibration, repeatability, and (for ρ) Allan deviation (Section 3.1, Appendix F); **digitization terms are display-only for Tier-F and are not propagated into Σ** . For numerical stability we scale the columns of $\mathbf{M}$ to unit norm and verify $\kappa(\mathbf{M}) \lesssim 3$.

Hybrid normalization and numerical inputs. By convention we anchor the R -channel with unit sensitivity to G : $M_R^G = 1$, $M_R^A = 0$. For the ρ -channel, $M_\rho^A, M_\rho^G = \mathcal{O}(1)$ are geometry factors (Appendix E). As constraints we insert the within-channel 95% results

$$|R - 1|_{95\%} < 5.5 \times 10^{-3} \text{ (Section 3.2)}, \quad |\Delta\rho/\rho|_{95\%} \lesssim 3 \times 10^{-3} \text{ (Section 3.3)},$$

each computed from Tier-N datasets only, encoded as zero-mean observations with standard deviations $(5.5, 3.0) \times 10^{-3}$, respectively. Short-term ρ windows (tighter, numeric) are reported side-by-side in Appendix G and are not mixed with the long-term bound.

Outlier control and heterogeneity. We apply a single Huber reweighting update to Σ , and inflate variances by a random-effects term τ^2 estimated from meta-residuals, i.e. $\sigma_C^2 \mapsto \sigma_C^2 + \tau^2$ (Appendix G). Leave-one-platform-out refits assess source heterogeneity.

Covariance ellipse and visualization. From (63) we construct the 95% confidence ellipse for (A, G) . A publication-ready AG-ellipse (contours and numerical levels) is provided in Appendix H, with the hybrid source table in Appendix G.

Final bounds (95% C.I.). With $\kappa(\mathbf{M}) \lesssim 3$ and the long-term channel inputs above, we obtain

$$|A| \lesssim 3.0 \times 10^{-3}, \quad |G| \lesssim 5.5 \times 10^{-3}. \quad (64)$$

These bounds are computed *exclusively* from Tier-N inputs; any Tier-F materials are excluded from weighting, fitting, and meta-combination. The parameter correlation

$$\rho_{AG} = \frac{[\text{Cov}(\hat{\boldsymbol{\theta}})]_{12}}{\sqrt{\text{Var}(\hat{A}) \text{Var}(\hat{G})}}$$

satisfies $|\rho_{AG}| \lesssim 0.3$, indicating limited cross-talk due to complementary channel sensitivities.

Robustness and interpretation. Relaxing digitization resolution by $\times 2$, changing the polynomial order of ρ -channel de-trending, and leave-one-platform-out refits shift (64) by at most $\pm 0.1 \times 10^{-3}$ (Appendix F, Appendix G). This limits the anisotropic response $\propto |\nabla\Phi|$ to $\lesssim 0.6\%$ at current public precision. By the continuity criterion of Section 2.5, Maxwell reduction in the EM sector is supported at $\gtrsim 98\%$ agreement, with a sub-percent window to be narrowed by finer tilt sweeps and high-stability ρ tracking (Appendix G).

3.4.1 Cross-validation with external constraints (mapped onto (A, G))

Policy (numeric-first, common scale). We translate published 95% *limits* from external tests—rotating optical cavities (modern Michelson–Morley type), cavity/maser and WGM resonators, laboratory birefringence, and cosmological polarization rotation—onto the same (A, G) parameterization used in this chapter. Only DOI-linked tables or raw numbers (Tier–N) enter numeric conclusions; any figure-derived estimates (Tier–F) are *illustrative*.

Mapping rule (summary). For each observable O we adopt the linear response

$$\Delta O \simeq m_A A + m_G G + n, \quad \text{Var}(n) = \sigma_O^2.$$

Given a reported 95% bound $|\Delta O|_{95}$, we infer one-parameter limits

$$|A|_{95} \leq \frac{|\Delta O|_{95}}{|m_A|} \quad (G=0), \quad |G|_{95} \leq \frac{|\Delta O|_{95}}{|m_G|} \quad (A=0).$$

When a covariance between the extracted harmonics (e.g., $\cos 2\theta, \sin 2\theta$) is available, we form the joint 95% ellipse for (A, G) using the 2x2 WLS of [Section 3.4](#). Sensitivities m_A, m_G are taken from the mode/geometry overlaps summarized in [Appendix E](#) (row IDs cited per entry).

Table 8: External constraints expressed on the common (A, G) scale (Tier–N only used for numeric conclusions).

Domain	Observable O	Published 95% limit	Mapping	Implied bounds (95%)
Rotating optical cavities (MM-type)	Fractional anisotropy $\Delta\nu/\nu$	$ \Delta\nu/\nu _{95} \leq \mathcal{N}_{\text{MM}}$	$\Delta\nu/\nu \simeq c_A A + c_G(\theta)G$ (App. E, row #)	$ A \leq \mathcal{N}_{\text{MM}}/ c_A ; G \leq \mathcal{N}_{\text{MM}}/ c_G $
Dielectric WGM / maser (external)	$\Delta \ln \rho$ or Allan plateaus	$ \Delta \ln \rho _{95} \leq \mathcal{N}_{\rho, \text{ext}}$	$\Delta \ln \rho \simeq s_A A + s_G(\theta)G$ (App. E, row #)	$ A \leq \mathcal{N}_{\rho, \text{ext}}/ s_A ; G \leq \mathcal{N}_{\rho, \text{ext}}/ s_G $
Bulk-optics birefringence (lab)	$ \Delta n /n$	$ \Delta n /n \leq \mathcal{N}_n$	$ \Delta n /n \rightarrow \Delta\nu/\nu \rightarrow (A, G)$ (App. E, row #)	As at left
Cosmic polarization rotation (CMB/radio)	$ \beta $ (deg)	$ \beta _{95} = \mathcal{N}_\beta$	$\beta \simeq k_G G_{\text{eff}}$ (model)	$ G_{\text{eff}} \leq \mathcal{N}_\beta/ k_G $

Notes. (1) **Rotating cavity (MM-type).** Harmonic demodulation at $2\omega_{\text{rot}}$ isolates the $\cos 2\theta$ term that maps directly to G ; phase conventions follow the referenced setup. (2) **External WGM/maser.** Sensitivity is isomorphic to the ρ -channel; use the same s_A, s_G row from [Appendix E](#). (3) **Bulk birefringence.** Material and stress corrections introduce lab-specific offsets; these are treated conservatively and are not weighted into the combined likelihood unless raw covariance is provided. (4) **Cosmic rotation.** The bound depends on line-of-sight integrals and cosmology; it serves as a supportive constraint and is not used in the numeric combination.

Joint combination with in–chapter channels. Stack the external O_i with the in–chapter R and ρ channels,

$$\mathbf{d} = [\Delta O_i], \quad M = \begin{bmatrix} m_A^{(i)} & m_G^{(i)} \end{bmatrix}, \quad \hat{\boldsymbol{\theta}} = (M^\top \Sigma^{-1} M)^{-1} M^\top \Sigma^{-1} \mathbf{d},$$

using Tier–N variances in Σ . The resulting 95% AG ellipse is overlaid on the in–chapter ellipse ([Appendix H](#)) to visualize tightening and axis rotation as domains are added.

Author fill–in guide (data–driven). (1) Enter numeric **Published 95% limit** values (with units) from each DOI into the project datasheet ([Appendix J](#)).

(2) Select the appropriate sensitivity row in [Appendix E](#) and record its row ID in the **Mapping** column.

(3) The build script computes the **Implied bounds** and updates the combined WLS in [Section 3.4](#); the figure in [Appendix H](#) is regenerated accordingly.

3.5 Reproducibility and verification checklist (experimental guidance)

- **Angle–normalized R (tilt channel).** Two tilt settings (θ_1, θ_2) with repeats $N \geq 20$; angle calibration $\leq 0.1^\circ$. Bi–directional sweeps for hysteresis. Sequence: $\theta_1 \rightarrow \theta_2 \rightarrow \theta_1$ with dwell $\geq 10 \tau_{\text{lock}}$. Compute R via [Eq. \(55\)](#) and cross–average all pairs. Co–log temperature and stimuli using the [Appendix G](#) template.
- **Resonator ρ (TE/TM mode–ratio channel).** Synchronous acquisition of $f_{\text{TE}}, f_{\text{TM}}$ on a *single* timebase; target long–term drift $\leq 10^{-3}$. Operate near a temperature inversion point; hold $|\Delta T| \leq 1$ mK; pre–measure $\partial f / \partial T$ for correction. Form $\rho = f_{\text{TE}} / f_{\text{TM}}$ from identical timestamps; annotate window (short vs. long) as in [Section 3.3](#).
- **Hybrid evidence (graphics + numeric).** Pair high–fidelity digitization (relative variations along θ, t) with numeric source tables (absolute scales). Apply fusion/weighting rules in [Appendix F](#); archive both raw figures and numeric tables in [Appendix G](#).
- **Calibration and traceability.** Cross–check tiltmeter (level/magnetometer or optical gyro); lock counters to one reference. Archive axis–calibration snapshots and pixel $\leftrightarrow$ coordinate transforms. Version–lock scripts/notebooks; record hashes in [Appendix G](#).
- **Digitization uncertainty.** Adopt 1/5 tick spacing as $\pm 1\sigma$; propagate by first–order rules plus non–parametric bootstrap ([Appendix F](#)). Perform independent re–digitization (operator/tool/resolution) for cross–validation. Report medians; use 2.5–97.5% percentiles for CIs ([Section 3.1](#)).
- **Pre–registered decision criteria.** Significance $\geq 5\sigma$; linearity $R^2 \geq 0.95$; null–failure $\leq 1\%$. Channel tolerances: $|R - 1|_{95\%} < 5.5 \times 10^{-3}$ ([Section 3.2](#)); $|\Delta \rho / \rho|_{95\%} \lesssim 3 \times 10^{-3}$ ([Section 3.3](#)). Translate to (A, G) by [Section 2.5](#) and [Section 3.4](#).
- **Robustness tests (required).** Halve bandwidth (double lock time) and verify slope invariance. Rotate polarization/principal axis to check even/odd separation persists. Reverse loop/path to confirm holonomy sign flip (where applicable). Run leave–one–platform–out meta–analysis and Huber re–weighting ([Appendix G](#)).
- **Operational targets (numerical).** R : $\delta \theta \leq 0.1^\circ$, $\delta(\Delta B) / \Delta B \leq 3 \times 10^{-3}$. ρ : Allan deviation $< 10^{-12}$ at 10^3 s; long–term $|\Delta \rho / \rho| \leq 3 \times 10^{-3}$. When met, expected bounds: $|G| \lesssim 5.5 \times 10^{-3}$, $|A| \lesssim 3 \times 10^{-3}$ ([Section 3.4](#)); short–term numeric limits in [Section 3.3](#).
- **Data management and release.** Archive raw time series, full metadata (T, P , clamping, stimuli), calibration files, and digitized coordinates. Release analysis notebooks and logging template per [Appendix G](#); include the AG–ellipse figure/contours from [Appendix H](#).

3.6 Limitations and improvement points

Limitations.

1. **Reliance on secondary graphics (digitization bias).** Estimates derived from figure digitization are more conservative than those from raw time series. Axis calibration, tick resolution, and scan distortions inflate the variance of $|R - 1|$ and $|\Delta\rho/\rho|$. Mitigation follows [Appendix F](#), but performance still depends on access to high-resolution PDF/vector figures.
2. **Cross-platform heterogeneity.** Metallic/graphene rings, bulk tilt oscillations, and dielectric/WGM resonators differ in temperature coefficients, magnetic hysteresis, clamping stress, and reference stability. Residual correlations may bias the A – G covariance ([Section 3.4](#)).
3. **Neglected dispersion.** We ignore weak dispersion in $\chi(\Phi, \nabla\Phi; \omega)$. For some resonators, $\partial\chi/\partial\omega \neq 0$ leaves small residuals in ρ ([Section 2.2](#)).
4. **Angle/geometry calibration limits.** Stage hysteresis, mount geometry, and field misalignment can bias R . Non-simultaneous acquisition adds $\sim 10^{-3}$ -level systematics ([Section 3.2](#)).
5. **Selection/reporting bias.** Public figures with clean signals may be over-represented. Weights in meta-analysis can be distorted; robust weighting and preregistered rules are required ([Appendix G](#)).
6. **Mixed time-windows in ρ .** Short (synchronous) vs. long (drift-inclusive) windows are sometimes mixed, obscuring comparisons; the analysis must state the window explicitly ([Section 3.3](#)).

Improvements.

1. **Raw-data reanalysis on a common timebase.** Recompute R and ρ from DOI-linked time series with timestamps. Replace digitization limits by synchronous acquisition, common-mode removal, and Allan-variance analysis ([Appendix F](#)).
2. **Finer angle sweeps & perturbation tests.** Densify θ to $\leq 1^\circ$ spacing; repeat bi-directional sweeps ($\uparrow, \downarrow$) to estimate/remove hysteresis. With calibrated θ , fit the slope $\partial R/\partial(\cos 2\theta)$ ([Section 3.2](#)).
3. **High-stability ρ tracking with explicit windows.** Use a single reference and synchronous TE/TM acquisition. Operate near a TK-zero; co-log vacuum/stress. Report both long-window bounds ($|\Delta\rho/\rho| \lesssim 10^{-3}$) and short-window limits from comb-beat/linewidth ($\sim 10^{-9}$) ([Section 3.3](#), [Appendix G](#)).
4. **Blinding & preregistration.** Pre-fix ROI, cuts, and decision criteria (5σ , $R^2 \geq 0.95$, null-failure $\leq 1\%$). Separate analysts with key masks; use the [Appendix G](#) template.
5. **Robust statistics & meta-analysis.** Apply Huber/Tukey WLS, leave-one-platform-out, and random-effects models. Diagnose publication bias via trim-and-fill, p -curve, and cumulative meta-analysis ([Appendix G](#)).
6. **Dispersion modeling (first-order).** Augment $\chi(\omega) = \chi_0 + \alpha\Phi + \eta\mathcal{K}(\nabla\Phi) + \beta(\partial_\omega\chi)_{\omega_0}(\omega - \omega_0)$ to separate residual trends in ρ ([Section 2.2](#)).
7. **Injection & synthetic tests.** Inject per-mil A, G signals to estimate recovery and bias; re-validate via path reversal, halved bandwidth, and axis rotation ([Section 3.5](#)).
8. **Explicit hybrid weights.** Fuse numeric tables (absolute scale) and digitized tracks (relative variation) with w_{num} and w_{dig} ; down-weight w_{dig} if cross-checks fail ([Appendix F](#)).

Numerical targets (feasible).

$$|R - 1| \rightarrow \mathcal{O}(10^{-3}), \quad \left| \frac{\Delta\rho}{\rho} \right| \rightarrow \mathcal{O}(10^{-3})$$

At these levels,

$$|G| \sim \text{few} \times 10^{-3}, \quad |A| \sim \text{few} \times 10^{-3}$$

enter a decisive upper-bound or detectable-nonzero regime ([Section 3.4](#), [Section 2.5](#)).

Appendix links. Digitization/uncertainty: [Appendix F](#); meta-analysis/robustness: [Appendix G](#); AG-ellipse visualization: [Appendix H](#). Operational logs and release format follow [Appendix G](#).

4. Phase observables and parity–holonomy decision

Chapter overview. This chapter formulates the *observable imprint* of the curvature–induced $U(1)$ connection A_μ as explicit phase observables, and isolates the geometric (odd–parity) component by a *parity–holonomy* procedure that removes dynamical (even–parity) phases. *This chapter is an experimental/operational guide, separate from Section 3’s data re–analysis; the same procedures can also be applied to numerical verification.* The core flow is: (i) construct the loop phase $\oint_\gamma A$ and (ii) a Pancharatnam–overlap based phase as *independent* estimators and demonstrate $\mathcal{O}(a^2)$ agreement as the mesh spacing a decreases; then (iii) decide using preregistered thresholds on linearity, signal–to–noise, and null–failure rate (see [4, 51, 52, 53]). Phase observables tie to the sensitivity kernels of [Section 2.3](#); decision rules and confidence intervals follow the uncertainty framework of [Section 3.5](#) and [Appendix F](#).

Pipeline note (numeric–first). All phase analyses ingest Tier–N sources (DOI–linked numeric tables or raw logs) on a *common timebase* for inference; Tier–F panels are *illustrative; excluded from inference* and are excluded from confidence intervals and headline bounds. Acquisition is *same–window synchronized* with angle/current/polarization logs to mitigate hysteresis and asynchrony; uncertainty and drift handling follow [Appendix F](#), with preregistered safeguards in [Appendix G](#).

4.1 Definition of observables: loop phase and flux

Loop–phase (holonomy) observable. For a closed curve $\gamma : [0, 1] \rightarrow \mathcal{U}$, define

$$\Phi_\gamma \equiv \oint_\gamma A_\mu dx^\mu = \iint_{S(\gamma)} F_{\mu\nu} dS^{\mu\nu} \quad (65)$$

where the equality is Stokes’ theorem and $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the curvature. Reversing the path gives $\Phi_{\gamma^{-1}} = -\Phi_\gamma$, making Φ_γ directly suited to parity–holonomy extraction.

Geometric properties (gauge and reparameterization invariance). Under a gauge change $A \mapsto A + d\lambda$ one has $\oint_\gamma d\lambda = 0$, hence Φ_γ is *gauge invariant*. For any monotone reparameterization $\gamma \circ \varphi$, $\Phi_{\gamma \circ \varphi} = \Phi_\gamma$. Concatenation obeys $\Phi_{\gamma_1 \circ \gamma_2} = \Phi_{\gamma_1} + \Phi_{\gamma_2}$. For complementary loops $\gamma, \bar{\gamma}$ that tile a surface Σ , $\Phi_\gamma + \Phi_{\bar{\gamma}} = \iint_\Sigma F$, i.e., a flux additivity rule.

Parity–holonomy (even/odd split). To suppress experimental biases (dynamic and environmental phases) we split the observed phase into even/odd parts:

$$\Phi_{\text{odd}}(\gamma) = \frac{1}{2}[\Phi_\gamma - \Phi_{\gamma^{-1}}], \quad \Phi_{\text{even}}(\gamma) = \frac{1}{2}[\Phi_\gamma + \Phi_{\gamma^{-1}}]. \quad (66)$$

With ideal time–reversed traversal, $\Phi_{\text{odd}} = \Phi_\gamma$ and $\Phi_{\text{even}} = 0$. Experimentally, $\Phi_{\gamma^{-1}}$ is implemented by order–swap or polarization/axis inversion; see [Section 4.3](#).

Discrete (mesh) implementations and consistency criterion. Approximate γ by oriented segments $\gamma = \bigcup_{\ell=1}^N \ell$. Define link phases $U_\ell = \exp(i \int_\ell A \cdot dl)$ and the line–integral estimator

$$\widehat{\Phi}_\gamma^{(1)} = \sum_{\ell \in \gamma} \text{Arg } U_\ell. \quad (67)$$

Independently, form the Pancharatnam overlap estimator from section–wise states $u(x_k)$:

$$\widehat{\Phi}_\gamma^{(2)} = \text{Arg} \left[\langle u_0 | u_1 \rangle \langle u_1 | u_2 \rangle \cdots \langle u_{N-1} | u_0 \rangle \right]. \quad (68)$$

With consistent branch–cut management, the difference satisfies $\widehat{\Phi}_\gamma^{(1)} - \widehat{\Phi}_\gamma^{(2)} = \mathcal{O}(a^2)$ (derivation outline in [Appendix D](#); numerical confirmation in [Section 5](#)).

Normalization and reference loops. For calibration, acquire (i) a contractible, off–support loop γ_{off} and (ii) a standard loop γ_0 of known area S_0 . We require

$$|\Phi_{\gamma_{\text{off}}}| \leq z_{0.995} \sigma_\Phi \quad (\text{null–failure} \leq 0.5\%), \quad \frac{\Phi_\gamma}{S(\gamma)} \approx \frac{\Phi_{\gamma_0}}{S_0} \quad (\text{linear scaling}).$$

Violations trigger checks of axis calibration, branch handling, and repeatability logs ([Appendix G](#)).

Reportables and uncertainty. Report the medians and 2.5–97.5% intervals for $(\widehat{\Phi}_{\text{odd}}, \widehat{\Phi}_{\text{even}})$, and the slope β and significance (e.g., $\geq 5\sigma$) from the regression $\Phi_{\text{odd}} = \beta X + \epsilon$ against the control variable X ([Section 4.4](#), [Section 3.5](#)). Phase–based conclusions are combined with the R –channel ([Section 3.2](#)) and ρ –channel ([Section 3.3](#)) constraints via the covariance rules of [Section 3.4](#) to yield bounds on (A, G) .

Experimental interpretation (closure). Thus $\oint_\gamma A = \iint_{S(\gamma)} F$ states that the *frame–holonomy–induced* effective gauge potential yields the same phase/flux observables as standard electromagnetism; differences vanish experimentally in the $|\nabla\Phi| \rightarrow 0$ limit ([Section 2.5](#)). The loop phase $\oint_\gamma A$ by itself is not beyond Maxwell; it is an instrument reading of the standard curvature two–form. Any genuine departure would have to be traced to the constitutive response $\chi(\Phi, \nabla\Phi)$.

Singular-set detection playbook. (1) Acquire a *null* contractible loop $\gamma_{\text{off}} \subset \mathcal{U}$ (expect $|\Phi_{\gamma_{\text{off}}}| \approx 0$).

(2) Sweep a family of loops $\{\gamma(\mathbf{r})\}$ across the sample; discrete jumps in $\Phi_{\text{odd}}(\gamma(\mathbf{r}))$ indicate changes in the linking number with $\mathcal{S}_{\text{cau}}$.

(3) For any suspected branch/interface, compare two homotopic loops differing only by a pierce of the interface to isolate the Čech transition phase.

These rules tie the global (Čech) structure in [Appendix B](#) to concrete, repeatable observables.

4.2 Two phase estimators: line–integral vs. Pancharatnam

Overview. We refine the discrete (mesh) implementation of [Section 4.1](#) into two practical estimators (E1), (E2) usable in experiment and numerics. They are *independent*

yet agree to $\mathcal{O}(a^2)$ as the mesh spacing $a \rightarrow 0$ (derivation outline in [Appendix D](#)). This agreement serves as a primary quality metric for the parity–holonomy decision via $\Delta\Phi \equiv \widehat{\Phi}_\gamma^{(1)} - \widehat{\Phi}_\gamma^{(2)}$.

(E1) Connection line–integral estimator (link–sum, Wilson–loop form). For a mesh path $\gamma = \bigcup_{\ell=1}^N \ell$, define link phases $U_\ell = \exp(i \int_\ell A \cdot dl)$ and compute

$$\widehat{\Phi}_\gamma^{(1)} = \sum_{\ell \in \gamma} \text{Arg } U_\ell,$$

cf. [Eq. \(67\)](#). Implementation guidelines:

- *Branch consistency.* When accumulating principal values $\text{Arg} \in (-\pi, \pi]$, enforce continuity by unwrapping jumps: if $|\Delta\varphi_\ell| > \pi$, add/subtract 2π . This guarantees *gauge–continuous* accumulation (see [Appendix D](#)).
- *Gauge–patch junctions.* If γ crosses gauge patches, correct the phase jump $\Delta\lambda$ at the boundary to maintain continuity.
- *Error order.* Trapezoidal link integration is $\mathcal{O}(a^3)$ per link; the loop sum converges as $\mathcal{O}(a^2)$, providing one side of the $\mathcal{O}(a^2)$ consistency.

(E2) Pancharatnam overlap estimator (Bargmann invariant). Sample a normalized state section $u(x)$ along the loop at mesh points x_k and define

$$\widehat{\Phi}_\gamma^{(2)} = \text{Arg} \left[\langle u_0 | u_1 \rangle \langle u_1 | u_2 \rangle \cdots \langle u_{N-1} | u_0 \rangle \right],$$

cf. [Eq. \(68\)](#). Practical notes:

- *Local gauge invariance.* Under $u_k \rightarrow e^{i\alpha_k} u_k$ the overall Bargmann phase is invariant.
- *Numerical stabilization.* If $|\langle u_k | u_{k+1} \rangle|$ is too small (near–orthogonal), increase mesh density or insert midpoints; adopt parallel–transport gauge to keep overlaps ≈ 1 . Optionally apply Gram–Schmidt re–normalization to suppress drift.
- *Error order.* With a smooth section, the estimator converges as $\mathcal{O}(a^2)$ ([Appendix D](#)).

Estimator agreement and diagnostic procedure. For $\Delta\Phi(a) \equiv \widehat{\Phi}_\gamma^{(1)}(a) - \widehat{\Phi}_\gamma^{(2)}(a)$, fit

$$\Delta\Phi(a) = c_2 a^2 + \mathcal{O}(a^3)$$

by least squares. As $a \downarrow$, smaller c_2 indicates sound branch handling and gauge continuity. Reporting rule:

$$|\Delta\Phi| \leq z_{0.975} \sigma_{\Delta\Phi} \quad (95\% \text{ C.I.}).$$

If this fails, follow the [Appendix G](#) checklist: (i) unwrap rules, (ii) patch–boundary correction, (iii) mesh refinement, (iv) state re–normalization.

Noise and uncertainty model (common to both). With N independent repeats,

$$\sigma_{\widehat{\Phi}^{(1)}}^2 \approx \sum_{\ell} \sigma_{\text{Arg } U_\ell}^2, \quad \sigma_{\widehat{\Phi}^{(2)}}^2 \approx \sum_k \sigma_{\text{Arg} \langle u_k | u_{k+1} \rangle}^2.$$

Report medians and 2.5–97.5% intervals for each estimator ([Section 4.4](#); uncertainty pipeline in [Appendix F](#)).

Operational recommendations (summary).

- *Mesh choice*: sample so that $\max_k \angle(u_k, u_{k+1}) \lesssim 10^\circ$.
- *Branch management*: use cumulative-sum unwrapping and log patch-junction corrections in the [Appendix G](#) format.
- *Consistency target*: $|\Delta\Phi|$ decays as a^2 and the $a \rightarrow 0$ extrapolants agree within 1σ .
- *Parity-holonomy linkage*: build $\widehat{\Phi}_{\text{odd/even}}$ from $\widehat{\Phi}_\gamma^{(1,2)}$ via [Eq. \(66\)](#), then apply the linear-scaling test for Φ_{odd} under [Section 4.3](#).

4.3 Parity-holonomy extraction procedure

Principle. Parity-holonomy implements [Eq. \(66\)](#), $\Phi_{\text{odd}}(\gamma) = \frac{1}{2}[\Phi_\gamma - \Phi_{\gamma^{-1}}]$. If a literal time-reversed traversal γ^{-1} is unavailable, an *effective inverse* is realized by: (1) *segment order-swap* (reverse the segment order of γ), (2) *polarization/axis inversion* (flip the sign of the control), (3) *time-slide* (same-window differencing to suppress drift). Their actions on the loop phase are

$$S : \Phi_\gamma \mapsto \Phi_{\gamma^{\text{rev}}}, \quad P : \Phi_\gamma \mapsto -\Phi_\gamma \text{ (sign flip)}, \quad T_\Delta : \Phi(t) \mapsto \Phi(t) - \Phi(t + \Delta),$$

and the composition $S \circ P \circ T_\Delta$ provides an effective mapping to γ^{-1} up to $\mathcal{O}(a^2)$ plus higher-order drift terms.

Procedure A — order-swap construction of Φ_{odd} . For a mesh path $\gamma = \bigcup_{\ell=1}^N \ell$:

1. *Forward loop*: run γ and compute $\widehat{\Phi}_\gamma^{(1,2)}$ ([Section 4.2](#)).
2. *Reversed loop*: run γ^{rev} (segments reversed) and compute $\widehat{\Phi}_{\gamma^{\text{rev}}}^{(1,2)}$.
3. *Even/odd split*:

$$\widehat{\Phi}_{\text{odd}} = \frac{1}{2} [(\widehat{\Phi}_\gamma^{(1)} - \widehat{\Phi}_{\gamma^{\text{rev}}}^{(1)}) + (\widehat{\Phi}_\gamma^{(2)} - \widehat{\Phi}_{\gamma^{\text{rev}}}^{(2)})] / 2, \quad \widehat{\Phi}_{\text{even}} = \frac{1}{2} [(\widehat{\Phi}_\gamma^{(1)} + \widehat{\Phi}_{\gamma^{\text{rev}}}^{(1)}) + (\widehat{\Phi}_\gamma^{(2)} + \widehat{\Phi}_{\gamma^{\text{rev}}}^{(2)})] / 2.$$

Averaging the two estimators suppresses $\mathcal{O}(a^2)$ residuals ([Section 4.2](#)).

Procedure B — polarization/axis inversion as an effective γ^{-1} . If the geometry is fixed and a *sign reversal* of the control flips the phase (e.g., $I \rightarrow -I$, $\Phi_B \rightarrow -\Phi_B$, $\sigma^\pm$ swap), then

$$\Phi_{\gamma^{-1}} \simeq -\Phi_\gamma \big|_{\text{sign-flipped}}.$$

Conditions: (i) keep the path and time window identical, (ii) record and correct collateral effects of the sign flip (loss, mode cross-talk) in the [Appendix G](#) form.

Procedure C — time-slide removal of dynamic phases. When common drift dominates, suppress Φ_{even} by same-window differencing:

$$\widehat{\Phi}_{\text{odd}}^{(\Delta)}(t) = \frac{1}{2} [\Phi_\gamma(t) - \Phi_\gamma(t + \Delta)],$$

choosing Δ slightly above the drift correlation time ($\Delta \sim 0.5\text{--}2 \tau_{\text{drift}}$). This mirrors the common-differential suppression used in the ρ -channel ([Section 3.3](#)).

Procedure D — bi-directional θ -sweeps and hysteresis removal. Set the angle grid with $\Delta\theta \leq 1^\circ$. For each θ_k , acquire within the *same time window*: up-sweep $\theta_{k-1} \rightarrow \theta_k$ and down-sweep $\theta_{k+1} \rightarrow \theta_k$. Diagnose hysteresis by

$$H(\theta_k) = |R_\uparrow(\theta_k) - R_\downarrow(\theta_k)|.$$

If $H(\theta_k) > z_{0.995} \sigma_R$, mask the point or re-calibrate mechanics (backlash/slip) and reacquire. Form the combined estimate

$$R(\theta_k) = \frac{1}{2}(R_\uparrow + R_\downarrow), \quad \sigma_R^2 \leftarrow \frac{1}{2}(\sigma_\uparrow^2 + \sigma_\downarrow^2)$$

and propagate uncertainties following [Appendix F](#). This same-window bi-directional sweep policy is shared with the [Section 3.2](#) R -channel protocol and keeps residual mechanical hysteresis below the 0.5% level in the overall phase budget.

Off-support (null) loop and reference loop. Test with a contractible, off-support loop γ_{off} and a reference loop γ_0 of known area:

$$|\Phi_{\gamma_{\text{off}}}| \leq z_{0.995} \sigma_\Phi \quad (\text{null failure} \leq 0.5\%), \quad \Phi_{\text{odd}}/S \text{ linear in } S \quad (\text{Section 4.1}).$$

If these fail, re-inspect unwrapping, patch boundaries, and axis/angle calibration ([Appendix G](#)).

Table 9: Comparison of “signal fingerprints” across alternative scenarios. Basis $\{1, \cos 2\theta, \sin 2\theta\}$ in a weakly linear response window. Higher harmonics (4, 6, ...) are treated as weak byproducts from nonlinearity/calibration.

Hypothesis	$\cos 2\theta$ 1st	Higher (4,6...)	Path inversion (odd)	Response to polarization-axis rotation	Null-loop failure rate	Sensitive channel
H0. Pure Maxwell (isotropic)	○	○	○	Rotation-insensitive (constant only)	Very low	(—)
H1. Isotropic only: $A \neq 0, G = 0$	○	○	○	Rotation-insensitive (constant only)	Low	ρ
H2. Gradient only: $A = 0, G \neq 0$	●	△	○	Tracks $\cos 2(\theta + \phi_0)$	Low	R
H3. Combined (linear): $A \neq 0, G \neq 0$	●	△	○	$\cos 2\theta$ + constant (cross-calibration dependent)	Medium	$R + \rho$
H4. Path holonomy (odd phase)	●	△	●	Rotation–amplitude invariant; sign flips with path inversion	Very low	R (odd), aux. ρ
H5. Instrumental/ systematic (guide-line)	△	△	○	Setup/calibration dependent (unstable)	High	Pattern unstable

Legend: ● present/dominant; ○ absent/insensitive; △ conditional (weak or setting-dependent). Channels: R = tilt–period ratio; ρ = mode ratio.

Linear scaling and joint even/odd regression. For a control $X \in \{S, I, \theta, \dots\}$,

$$\Phi_{\text{odd}} = \beta X + \epsilon, \quad \Phi_{\text{even}} = \beta_0 + \beta_2 X^2 + \epsilon',$$

fit simultaneously and adopt thresholds $R^2(\Phi_{\text{odd}} \text{ vs } X) \geq 0.95$, $|\beta|/\text{SE}(\beta) \geq 5$. Slope β should match the geometric flux or medium response ([4, 51]).

Uncertainty and quality control. With N repeats,

$$\sigma^2(\hat{\Phi}_{\text{odd}}) \approx \frac{1}{4} \left[\sigma^2(\hat{\Phi}_\gamma) + \sigma^2(\hat{\Phi}_{\gamma^{-1}}) \right],$$

and 95% CIs are obtained by nonparametric bootstrap (10^4 resamples) (Appendix F). Quality metrics: (i) $|\Delta\Phi| = |\hat{\Phi}^{(1)} - \hat{\Phi}^{(2)}| \leq z_{0.975} \sigma_{\Delta\Phi}$; (ii) null-failure $\leq 1\%$; (iii) repeatability index $R_{\text{repeat}}^2 \geq 0.95$; (iv) complete branch/patch event logs.

Checklist.

1. Secure a γ – γ^{rev} pair (or implement via P and T $_{\Delta}$).
2. Compute both estimators $\hat{\Phi}^{(1)}$, $\hat{\Phi}^{(2)}$ and check $|\Delta\Phi|$ (Section 4.2).
3. Build $\Phi_{\text{odd/even}}$; pass linear-scaling and null tests (Section 4.1).
4. If failing: fix unwrapping, patch corrections, mesh refinement, and state re-normalization (Appendix G).

4.4 Noise model, SNR, and confidence intervals

Constituents of phase noise and an equivalent model. We model the observed phase as

$$\Phi_{\text{obs}}(t) = \Phi_{\text{true}}(t) + \nu(t),$$

where the noise ν collects (i) *sensor phase noise* (white phase/frequency), (ii) *path repeatability* errors (segment non-reproducibility), and (iii) *unwrapping/branch* residues (Appendix F). In short windows, we approximate $\nu \sim \mathcal{N}(0, \sigma_\Phi^2)$ and obtain $\sigma_\Phi \propto N^{-1/2}$ with N independent repeats. In long windows, low-frequency drift dominates; define an *effective* variance via the Allan deviation $\sigma_y(\tau)$:

$$\sigma_{\Phi, \text{eff}}^2(\tau) \equiv \sigma_\Phi^2 \oplus (2\pi f_0 \tau)^2 \sigma_y^2(\tau),$$

with f_0 the reference frequency (for resonators) or an equivalent sampling-rate factor (Section 3.3).

Spectral view (bandwidth-window normalization). With phase PSD $S_\Phi(f)$,

$$\text{Var}[\hat{\Phi}_{\text{odd}}] \approx \int_0^\infty |H_{\text{odd}}(f; \Delta, a)|^2 S_\Phi(f) df,$$

where H_{odd} is the effective filter set by the time-slide Δ of Section 4.3 and the mesh spacing a . White phase noise scales with the measurement bandwidth B ($\propto B$); flicker/drift components are suppressed by Δ . Thus *same-window differencing* primarily suppresses low- f noise, while *mesh densification* addresses discretization at high f (Section 4.2, Section 4.3).

SNR definition and matched filtering. For the linear scaling $\Phi_{\text{odd}} = \beta X + \epsilon$ with control $X \in \{S, I, \theta, \dots\}$,

$$\text{SNR} \equiv \frac{|\beta| \sigma_X}{\sigma_{\Phi, \text{eff}}}, \quad \text{SNR}_{\text{meas}} \equiv \frac{|\hat{\beta}|}{\text{SE}(\hat{\beta})}.$$

Under heterogeneous segment noise $\{\sigma_k\}$, use weighted least squares and heteroskedasticity-robust (Huber–White/HC) standard errors ([Appendix G](#)). For periodic drives (e.g., area modulation), apply a matched filter $h(t) \propto s(t)/S_{\Phi}(f)$ to maximize SNR.

Dispersive nuisance handling and orthogonalization. Treat $\delta\omega(t) = \omega(t) - \omega_0$ as a nuisance regressor on the common timebase $\{t_k\}$. Use (i) centering, $\sum_k \delta\omega(t_k) = 0$, and (ii) optional prewhitening by the estimated PSD $S_{\delta\omega}(f)$. In the joint fit for $(A, G, \delta\omega)$, monitor condition numbers $\kappa(\mathbf{X})$ and variance–inflation factors (VIFs); pre-register thresholds (e.g., $\text{VIF} \leq 5$) in [Appendix G](#). Include “dispersion-in” and “dispersion-out” results side by side (point estimates and 95% CIs), since the dispersion term reduces bias in (A, G) at the cost of a modest variance increase. For periodic drives, construct a matched filter h that is orthogonal to $\delta\omega$ (Gram–Schmidt on the design matrix), so that SNR_{meas} for A, G is preserved while dispersive leakage is nulled. This handling mirrors the dispersion covariate in the ρ -channel regression of [Section 3.3](#) (see also [Appendix E](#)).

Test statistics and confidence intervals (CIs).

1. *Existence test (single-point):* declare detection when $|\hat{\Phi}_{\text{odd}}| \geq z_{0.995} \sigma_{\Phi, \text{eff}}$ (a conservative 3–3.5 σ threshold). For small samples, use the t distribution.
2. *Upper bound (non-detection):*

$$|\Phi_{\text{odd}}| < z_{0.975} \sigma_{\Phi, \text{eff}} \quad (95\% \text{ C.I.}).$$

3. *Regression summary:* report the 95% upper bound $|\hat{\beta}| < z_{0.975} \text{SE}(\hat{\beta})$, with $\text{SE}(\hat{\beta})$ given as HC–robust.
4. *Bootstrap CIs:* use a nonparametric bootstrap ($N_{\text{bs}} = 10^4$) to report median $\pm$ 2.5–97.5% intervals and require agreement with normal–approximation CIs within 1σ ([Appendix F](#)).

Variance rules for even/odd splitting. Assuming independent repeats,

$$\sigma^2(\hat{\Phi}_{\text{odd}}) = \frac{1}{4} \left[\sigma^2(\hat{\Phi}_{\gamma}) + \sigma^2(\hat{\Phi}_{\gamma-1}) \right], \quad \sigma^2(\hat{\Phi}_{\text{even}}) = \frac{1}{4} \left[\sigma^2(\hat{\Phi}_{\gamma}) + \sigma^2(\hat{\Phi}_{\gamma-1}) \right].$$

Same–window differencing ([Section 4.3](#)) removes common–mode drift, reducing σ_{odd} .

Multiple testing and preregistration. When evaluating multiple controls X or multiple loops L_i , (i) adhere to the preregistered list ([Appendix G](#)); (ii) use $\hat{\beta}$ as the primary endpoint and $\hat{\Phi}_{\text{odd}}$ as secondary; (iii) apply Holm–Bonferroni when needed to control family–wise error.

Reporting standard (summary).

- median($\widehat{\Phi}_{\text{odd/even}}$) with 95% CIs (bootstrap and normal listed together).
- $\widehat{\beta} \pm \text{SE}_{\text{HC}}$, R^2 , and SNR_{meas} .
- $|\Delta\Phi| = |\widehat{\Phi}^{(1)} - \widehat{\Phi}^{(2)}|$ vs. mesh spacing a , with $a \rightarrow 0$ extrapolation (Section 4.2).
- Explicit time–window (short/long, Δ , τ), bandwidth B , and Allan minimum τ^* (Section 3.3).

Propagation to parameter bounds. Phase–based bounds propagate linearly to (A, G) . With sensitivity matrix $\mathbf{M}$ (Section 3.4) and measurement covariance Σ_Φ ,

$$\text{Cov}(\widehat{A}, \widehat{G}) = (\mathbf{M}^\top \Sigma_\Phi^{-1} \mathbf{M})^{-1}.$$

Hence, for $|\Phi_{\text{odd}}| < \Phi_{\text{max}}$, the 95% C.I. boundary on $|(A, G)|$ follows from the expression above (Section 2.5, Section 3.4).

Practical levers for noise suppression. (i) *Reduce bandwidth B* to suppress white phase noise; (ii) *same–window differencing* with $\Delta \approx 0.5\text{--}2 \tau_{\text{drift}}$ for drift removal; (iii) *mesh densification $a \downarrow$* to shrink $\mathcal{O}(a^2)$ discretization error; (iv) *matched filtering/lock–in* to maximize SNR; (v) *HC–robust SEs* to withstand heteroskedasticity (Appendix G).

4.5 Sensitivity scaling and design levers

Scaling law (minimal model). With effective area S , mean effective kernel $\langle F_{\text{eff}} \rangle$, and coherent–path gain M_{coh} , the leading–order phase signal scales as

$$|\Phi_{\text{odd}}| \sim M_{\text{coh}} S \langle F_{\text{eff}} \rangle, \quad \text{SNR} \sim \frac{M_{\text{coh}} S \langle F_{\text{eff}} \rangle}{\sigma_{\Phi, \text{eff}}}, \quad (69)$$

where $\sigma_{\Phi, \text{eff}}$ is defined in Section 4.4. Here M_{coh} is the number of *phase–coherently* added paths, and $\langle F_{\text{eff}} \rangle$ is the path/area average of the sensitivity kernel from Section 2.3. Design levers therefore reduce to $S \uparrow$, $M_{\text{coh}} \uparrow$, and $\sigma_{\Phi, \text{eff}} \downarrow$.

Bandwidth–averaging rule (noise budget). For measurement bandwidth B , effective averaging time T , and repeats N ,

$$\sigma_{\Phi, \text{eff}}^2 \approx \frac{\sigma_{\Phi, 0}^2}{N} + \int_0^B |H_{\text{odd}}(f; \Delta, a)|^2 S_\Phi(f) df,$$

where $\sigma_{\Phi, 0}$ is the single–shot standard deviation and H_{odd} is the filter set by the time–slide Δ and mesh spacing a (Section 4.3, Section 4.4). When white phase noise dominates, $\sigma_{\Phi, \text{eff}} \propto (BT)^{-1/2}$; under flicker/drift, optimizing Δ to suppress low–frequency power is decisive.

Boosting the coherence gain M_{coh} . For M paths/loops $\{\gamma_m\}_{m=1}^M$ combined *phase-coherently*,

$$\Phi_{\text{odd}}^{(\text{coh})} = \sum_{m=1}^M \Phi_{\text{odd}}(\gamma_m) \quad \Rightarrow \quad \text{SNR} \propto M_{\text{coh}} \simeq M,$$

whereas incoherent (asynchronous) averaging yields only $\text{SNR} \propto \sqrt{M}$. Synchronize time/frequency /polarization across paths and log the conditions per [Appendix G](#) to maintain coherence.

Kernel shaping (optimizing $\langle F_{\text{eff}} \rangle$). The kernel F_{eff} is determined by field distribution, boundary conditions, and polarization. (i) Re-shape loop geometry to overlap high-sensitivity regions (boundary layers); (ii) choose polarization/modes (e.g., $\sigma^\pm$, TE/TM) to align kernel signs (*phase alignment*); (iii) project periodic modulation (area/current/angle) into spectral regions with low noise floor, matching the drive spectrum to the quiet band of $S_\Phi(f)$ ([Section 3.3](#), [Section 4.4](#)).

Anisotropic response channel (angle optimization). In θ control, the G -sensitive piece follows $\cos 2\theta$ ([Section 2.3](#)). Regress on the orthogonal basis $\{1, \cos 2\theta, \sin 2\theta\}$:

$$\Phi_{\text{odd}}(\theta) \approx \beta_0 + \beta_c \cos 2\theta + \beta_s \sin 2\theta,$$

so the G amplitude is $\sqrt{\beta_c^2 + \beta_s^2}$. Place balanced samples near $\theta \approx 45^\circ, 135^\circ$ (large $|\cos 2\theta|$), while avoiding excessive tilt ($\theta \rightarrow 90^\circ$) where $\tan \theta$ magnifies angle bias ([Section 3.2](#)).

Angle sampling and linearity test (for G). Perform weighted least squares in the orthogonal basis $\{1, X = \cos 2\theta\}$. Use a balanced coverage of $[\theta_{\min}, \theta_{\max}]$ with grid $\Delta\theta \leq 1^\circ$. Preregistered pass thresholds are $R^2 \geq 0.95$ and $|\hat{\beta}|/\text{SE}(\hat{\beta}) \geq 5$, conditional on passing the hysteresis test $H(\theta_k) \leq z_{0.995}\sigma_R$. Uncertainties follow the same-window policy and Allan handling in [Appendix F](#) (see also the R -channel protocol in [Section 3.2](#)).

Loop geometry and area scaling. Compared to one large loop, coherently combining M identical small loops (each of area S/M) gives

$$\Phi_{\text{odd}}^{(\text{coh})} \sim M \times \left(\frac{S}{M} \langle F_{\text{eff}} \rangle \right) = S \langle F_{\text{eff}} \rangle,$$

leaving the signal unchanged but improving (i) robustness to spatial drift/mismatch and (ii) localization of branch/patch events. A single large loop simplifies boundary calibration. A hybrid (one large loop + an array of small loops) is recommended to match the environmental drift scale.

Modulation strategy (matched-filter view). Drive a control $X(t) = X_0 + X_1 s(t)$ periodically and apply a lock-in/matched filter $h(t) \propto s(t)/S_\Phi(f)$ to maximize

$$\text{SNR}_{\text{meas}} = \frac{|\hat{\beta}|}{\text{SE}(\hat{\beta})}.$$

A sinusoid is optimal on a white floor, while multi-tone/Hanning-windowed waveforms with many drift zero-crossings help in flicker-dominated regimes ([Section 4.4](#)).

Optimal design under resource constraints (summary). Given total time T_{tot} , maximum area S_{max} , and bandwidth limit B_{max} , formulate

$$\max_{S \leq S_{\text{max}}, a, \Delta, M_{\text{coh}} \leq M_{\text{max}}} \frac{M_{\text{coh}} S \langle F_{\text{eff}} \rangle}{\sigma_{\Phi, \text{eff}}(B, \Delta, a, T_{\text{tot}})} \quad \text{s.t.} \quad \text{branch/patch error rate} \leq \varepsilon.$$

A practical recipe is: (i) choose $\Delta \approx 0.5 - 2 \tau_{\text{drift}}$; (ii) pick a at the *knee* of the $|\Delta\Phi| \propto a^2$ decay curve (Section 4.2); (iii) increase M_{coh} up to the synchronization limit M_{max} .

Checklist (design levers at a glance).

- *Area S* : maximize within geometric constraints; record leakage/fringe corrections in Appendix G.
- *Coherence M_{coh}* : synchronize multi-loop time, polarization, and phase.
- *Kernel $\langle F_{\text{eff}} \rangle$* : optimize sign/overlap via boundary, polarization, and mode choice.
- *Bandwidth/windows*: reduce B , optimize Δ , and shrink a to lower $\sigma_{\Phi, \text{eff}}$.
- *Angle plan*: place θ samples to maximize $\cos 2\theta$ sensitivity while avoiding $\tan \theta$ -amplified bias (Section 3.2).
- *Modulation/lock-in*: match the drive spectrum to the noise PSD (Section 4.4).

Summary. Sensitivity is engineered along four axes: S , M_{coh} , $\langle F_{\text{eff}} \rangle$, and $\sigma_{\Phi, \text{eff}}$. Combined with the parity-holonomy workflow of Section 4.3 and the noise budget in Section 4.4, these rules provide practical guidance for designing experiments that propagate to tight (A, G) bounds (see Section 3.4).

4.6 Validation and safeguards: excluding false positives and look-alikes

Primary confounders and mitigations.

1. **Multi-path interference.** Single-mode operation (spatial, spectral, polarization) with narrow bandwidth; mode purity established by isolation scans. Side-mode growth indicates departure from the intended regime (Section 4.4).
2. **Anisotropic media/boundaries.** Path reversal and axis/polarization rotation isolate the odd component Φ_{odd} (Section 4.3); comparison with the even residue Φ_{even} serves as an internal control.
3. **Parametric drift (T/P/strain/reference).** Synchronous tensor logging and same-window differencing (T_{Δ}) suppress common mode; drift covariates enter the regression with HC-robust errors (Section 4.4, Appendix G).
4. **Unwrapping/branch artefacts.** Independent reconstructions with distinct tools/operators and branch-event logs; the agreement $|\Delta\Phi| = |\hat{\Phi}^{(1)} - \hat{\Phi}^{(2)}|$ must exhibit the $\mathcal{O}(a^2)$ decay (Section 4.2).
5. **Field leakage / stray flux.** Off-support regions are mapped; measured leakage is removed using a reference loop γ_0 ; the off-support loop γ_{off} satisfies the null bound $\leq 1\%$ failure rate (Section 4.1).

6. **Control cross-talk (e.g., sign-flip side effects).** When using P (sign reversal), losses and mode cross-talk are logged; gains are re-balanced and linearity re-checked ([Section 4.3](#)).

Orthogonal controls and falsification criteria.

- *Geometry flip at fixed flux.* A topologically equivalent loop with mirrored boundary and identical area leaves the sign of Φ_{odd} invariant under non-holonomic deformations.
- *Drive orthogonality.* Modulation of a control coupling only to Φ_{even} keeps the odd channel statistically null (95% C.I.).
- *Harmonic balance.* For sinusoidal drives, the Φ_{odd} spectrum is confined to the fundamental; higher harmonics indicate nonlinearity or leakage.

Cross-loop consistency. For two independent loops L_1, L_2 sharing a single calibration (area/current scale), the slopes in

$$\Phi_{\text{odd}} = \beta X + \epsilon$$

satisfy $\beta(L_1) \approx \beta(L_2)$ within the combined 1σ band. Consistency supports a geometric interpretation; deviations suggest medium/boundary influences or device-internal bias ([4, 51]; cf. [Section 4.5](#)).

Decision criteria.

1. **Null control.** $|\Phi_{\gamma_{\text{off}}}| \leq z_{0.995}\sigma_{\Phi}$ (null-failure $\leq 0.5\%$).
2. **Estimator agreement.** $|\Delta\Phi| \leq z_{0.975}\sigma_{\Delta\Phi}$ with $\Delta\Phi \propto a^2$ decay ([Section 4.2](#)).
3. **Linearity.** $R^2(\Phi_{\text{odd}} \text{ vs } X) \geq 0.95$ and $|\hat{\beta}|/\text{SE}(\hat{\beta}) \geq 5$ ([Section 4.3](#)).
4. **Repeatability.** Inter-run $R^2_{\text{repeat}} \geq 0.95$; leave-one-segment-out checks passed.

Failure diagnostics. When a criterion is not met, the record contains the item, deviation magnitude, and a hypothesized source (from the list above), together with the corrective step (e.g., unwrapping fix, patch correction, mesh refinement, bandwidth reduction, collateral correction for sign-flip). Null and linearity checks are repeated and the before/after C.I.s are appended ([Appendix G](#)).

Admissibility for parameter estimation. Only datasets that satisfy all criteria enter the joint (A, G) fit via the covariance rules of [Section 3.4](#). Non-admissible sets are retained for diagnostics but are excluded from the primary bounds to avoid bias.

4.7 Chapter summary and bridge to the next sections

Summary (key points). In this chapter we quantified the *topological imprint* of the curvature-induced $U(1)$ connection by: (1) defining the loop phase $\Phi_{\gamma} = \oint_{\gamma} A$ and its flux form in [Section 4.1](#); (2) constructing two *independent* estimators—(E1) link-sum

line integral and (E2) Pancharatnam overlap—in [Section 4.2](#), with an $\mathcal{O}(a^2)$ agreement criterion as the mesh $a \rightarrow 0$; (3) isolating the geometric odd-parity component Φ_{odd} via the *parity-holonomy* workflow (order-swap, polarization/axis inversion, time-slide) in [Section 4.3](#). We then unified the phase-noise model, SNR, and bootstrap/normal CIs in [Section 4.4](#); organized sensitivity scaling along $(S, M_{\text{coh}}, \langle F_{\text{eff}} \rangle, \sigma_{\Phi, \text{eff}})$ with actionable levers in [Section 4.5](#); and fixed safeguards to exclude false positives/look-alikes (single-mode operation, null-loop, estimator agreement, linearity and repeatability thresholds) in [Section 4.6](#). These phase observables propagate to bounds on (A, G) through the covariance rules of [Section 3.4](#).

Bridge to the next chapters.

- **Numerical validation (Section 5).** [Section 5](#) develops a *structure-preserving* scheme that implements the mesh/gauge details of [Appendix D](#) (branch handling, patch junctions) and verifies $|\Delta\Phi| = |\hat{\Phi}^{(1)} - \hat{\Phi}^{(2)}| \propto a^2$ convergence, gauge-patch independence, and null-loop pass rates ([Section 4.1](#)).
- **Experimental operations (Section 6).** [Section 6](#) compiles a *practical playbook* covering logs/checklists in the [Appendix G](#) format, preregistration (ROI/cuts/thresholds), modulation/lock-in strategy ([Section 4.4](#), [Section 4.5](#)), and cross-loop consistency tests ([Section 4.6](#)).

Handover artifacts (reproducibility package). This chapter ships with: (i) notebooks implementing estimators (E1/E2) and branch-event logs ([Appendix D](#)); (ii) scripts for noise/SNR evaluation and the bootstrap pipeline ([Appendix F](#)); (iii) automated decision checklists for null/linearity/consistency ([Appendix G](#)); (iv) publication-ready (A, G) confidence-ellipse figures and a template bundle ([Appendix H](#)).

Appendix cross-references (editor’s note). [Appendix D](#): link variables, branch unwrapping, and the $\mathcal{O}(a^2)$ agreement proof sketch.

[Appendix E](#): mode/polarization sensitivity kernels and the $\cos 2\theta$ channel of G .

[Appendix F](#): phase-uncertainty budget, bootstrap, and HC-robust SEs.

[Appendix G](#): preregistration template, logging format, QA/safeguard checklists.

[Appendix H](#): quick convention reference, SI $\leftrightarrow$ HL cheatsheet, first-order error propagation for R and ρ , and 95% (A, G) contour templates.

5. Structure-preserving numerics: gauge exactness, continuity, and energy stability

Aims and gist. This chapter connects the sensitivity kernels of [Section 2.3](#) and the data-driven bounds of [Section 3](#) to a *structure-preserving* numerical scheme that enables faithful reproduction and prediction. The pillars are: (i) a *gauge-exact* discrete differential structure ($d^2 = 0$, discrete Stokes), (ii) *exact* preservation of the continuity equation ($\delta J = 0$), and (iii) a *variational (symplectic)* time integrator for energy stability. Spatial discretization uses a Discrete Exterior Calculus (DEC) framework on *primal/dual*

meshes, and time integration employs a *midpoint–symplectic* scheme ([12, 1, 2, 28]). The constitutive/field system

$$H = \chi(\Phi, \nabla\Phi) : F, \quad \square\Phi - U'(\Phi) = J, \quad |\nabla\Phi| < \varepsilon$$

is embedded into the discrete Hodge machinery of Section 5.2; Section 5.3 guarantees charge conservation, and Section 5.4 ensures energy stability. Implementation checklists and logging formats follow Appendix D and Appendix G.




Figure 3: Conservation diagnostics of the scheme. Left: spatial-resolution convergence of the continuity residual $\|\Delta_t \rho + \delta J\|_2$ with fitted slope $\simeq 2.02$ (second-order in h ; cf. Appendix I.1). Right: time-step convergence of the energy residual with fitted slope $\simeq 1.97$ (second-order in Δt ; cf. Appendix I.2). Both panels summarize the behavior anticipated in Section 5. Extended validations, tabulated slopes, and mesh-sweep variants are in Appendix I.

Design principles (at a glance).

- *Topology–physics separation*: incidence matrices carry *pure topology* only; geometry/material (and Φ) enter solely via *Hodge (mass) matrices*. Hence $d^2 = 0$ holds identically, preserving Bianchi identities and gauge exactness (Section 5.1).
- *Gauge/boundary consistency*: $F = dA$ is invariant under $A \mapsto A + d\lambda$. PEC/PMC/PML boundaries are expressed by a primal/dual choice and Hodge restriction on the boundary (Appendix D).
- *Time staggering*: Yee-type half-step staggering centers the Maxwell coupling in time; variationally derived updates suppress long-term energy drift to boundary terms (Section 5.4).

5.1 Meshes, cochains, and boundary operators

Primal/dual complexes and orientation. Partition 3D space into a simplicial (or hexahedral) complex $\mathcal{K}$. With the set of k -cells $\mathcal{K}_k$, define chain/cochain spaces

$$C_k(\mathcal{K}) = \text{span}\{\sigma^k \in \mathcal{K}_k\}, \quad C^k(\mathcal{K}) = \text{Hom}(C_k(\mathcal{K}), \mathbb{R}).$$

Assign a consistent orientation to all cells so that the boundary operators $\partial_k : C_k \rightarrow C_{k-1}$ and the coboundary (discrete exterior derivative) $d_k : C^k \rightarrow C^{k+1}$ are well-defined via the incidence matrices $\mathbf{B}_k$:

$$d_k = \mathbf{B}_{k+1}^\top, \quad d_{k+1} d_k = 0,$$

i.e., the discrete $d^2 = 0$ (a Stokes-type identity) holds *identically*.

Dual complex and integral pairing. Construct a Voronoi/Yaghn *dual* complex $\mathcal{K}^*$ and pair each primal k -cell with a dual $(3-k)$ -cell preserving metric measures (length/area/volume). Define the primal-dual pairing as a discrete integral

$$\langle \alpha, \beta \rangle_k \equiv \sum_{\sigma^k \in \mathcal{K}_k} \alpha(\sigma^k) \beta(\star \sigma^k) m_k(\sigma^k),$$

where m_k encodes the geometric weights (length/area/volume). This pairing realizes the Hodge operators as (symmetric positive) mass matrices (Section 5.2).

Field placement (DEC convention). With midpoint (staggered) time placement, assign electromagnetic/curvature variables as

$$\begin{aligned} A &\in C^1(\mathcal{K}), & F &= dA \in C^2(\mathcal{K}), \\ E &\in C^1(\mathcal{K}), & D &\in C^2(\mathcal{K}^*), \\ B &\in C^2(\mathcal{K}), & H &\in C^1(\mathcal{K}^*), \\ \rho &\in C^3(\mathcal{K}^*), & J &\in C^2(\mathcal{K}^*), \\ \Phi &\in C^0(\mathcal{K}), & \nabla \Phi &\leftrightarrow d\Phi \in C^1(\mathcal{K}). \end{aligned}$$

This choice (i) integrates $F = dA$ naturally over primal faces, and (ii) casts the constitutive maps $D = \star_2^{(\Phi)} E$, $H = \star_1^{(\Phi)} B$ as *local* linear maps between primal and dual spaces (Section 5.2).

Boundary conditions and gauge transformations. On $\partial\mathcal{K}$:

- *PEC*: eliminate tangential primal 1-cochain DoFs for E ($E_{\parallel} = 0$); keep dual D free.
- *PMC*: apply the dual counterpart enforcing $B_{\perp} = 0$.
- *PML*: inject complex extension or scaling tensors into the boundary layer via the Hodge construction (Appendix D).

Under $A \mapsto A + d\lambda$, F remains invariant; boundary values of λ comply with the chosen physical boundary condition (fixed/free potential). Thus the scheme maintains *gauge exactness*.

Discrete codifferential δ and a preview of continuity. Define on the dual complex

$$\delta_k \equiv \left(\star_{k-1} \right)^{-1} d_{k-1}^\top \star_k,$$

where $\star_k$ are the Hodge (mass) matrices of Section 5.2. Then $\delta_k \delta_{k+1} = 0$ follows, leading in Section 5.3 to the *exact* discrete continuity equation $\Delta_t \rho + \delta J = 0$.

Time staggering and update order (overview). With midpoint time integration, place (E, D) at half-steps and (B, H) at integer steps:

$$B^{n+1} = B^n - \Delta t dE^{n+1/2}, \quad D^{n+1/2} = D^{n-1/2} - \Delta t (\delta H^n + J^n).$$

Evaluate $\chi(\Phi, d\Phi)$ at $\Phi^{n+1/2}$ so that the curvature coupling is synchronized with the field updates (Section 5.4).

Quality–assurance checklist (for Section 5.1). (i) verify $d_{k+1}d_k = 0$ numerically at machine precision; (ii) archive the primal–dual metric table (length/area/volume); (iii) log boundary DoF removal/retention consistency; (iv) run a gauge test $A \rightarrow A + d\lambda$ and confirm invariance of F and the energy (using the Appendix G template).

5.2 Discrete Hodge and embedding of the constitutive law

Principles and requirements of the Hodge star (mass matrices). Using the primal–dual pairing (Section 5.1), define linear maps $\star_k : C^k(\mathcal{K}) \rightarrow C^{3-k}(\mathcal{K}^*)$ that send the primal 1–cochain E and primal 2–cochain B to the dual 2– and 1–cochains D and H as symmetric positive–definite (SPD) *mass matrices*:

$$D = \star_2^{(\Phi)} E, \quad H = \star_1^{(\Phi)} B.$$

Two constraints are essential: (i) *Topology–physics separation*: $\star_k^{(\Phi)}$ contains only geometry/material (length/area/volume and the material tensor χ), and never alters the incidence structure. Hence $d^2 = 0$ holds identically and the Bianchi identities follow automatically (Section 5.1). (ii) *Energy positivity*: the quadratic form $\frac{1}{2}\langle E, \star_2^{(\Phi)} E \rangle + \frac{1}{2}\langle B, \star_1^{(\Phi)} B \rangle$ is positive, linking directly to energy stability in Section 5.4.

Local assembly and sparsity. For each primal k –cell σ^k and dual partner $\star\sigma^k$, assemble the local mass entry as

$$(\star_k^{(\Phi)})_{\sigma^k, \star\sigma^k} = \frac{m_{3-k}(\star\sigma^k)}{m_k(\sigma^k)} \mathbf{G}_k^{(\Phi)}(\sigma^k),$$

where m_k are geometric weights (length/area/volume) and $\mathbf{G}_k^{(\Phi)}$ is the local metric (a scalar or a small symmetric block) set by geometry/material. On orthogonal grids this reduces to diagonal (lumped) form; on distorted meshes it remains sparse with small cell–local symmetric blocks. The construction yields (i) *locality*, (ii) *sparsity*, and (iii) *conditioning control*.

Discrete embedding of the constitutive law (including nonlinearity). Discretize the continuum relation $H = \chi(\Phi, \nabla\Phi) : F$ as

$$H = [\chi(\Phi, d\Phi)]^\# F,$$

where $(\cdot)^\#$ maps the primal 2–cochain F to the dual 1–cochain H (analytically corresponding to $\star_1^{(\Phi)}$). Implement $\star_1^{(\Phi)}$ as the composition

$$\star_1^{(\Phi)} \equiv \mathcal{M}_1(\text{metric}) \mathcal{C}_1(\chi(\Phi, d\Phi)),$$

with $\mathcal{M}_1$ the geometric mass and $\mathcal{C}_1$ the material mapping, so both linear and nonlinear media fit the same framework. For small curvature gradients $|\nabla\Phi| < \varepsilon$, $\chi(\Phi, d\Phi) \simeq \chi_0 + \mathcal{O}(\varepsilon)$, recovering the standard Maxwell linear medium with constant $\star_k$ (Section 2.3).

Isotropic–anisotropic (A/G) split and sensitivity consistency. Decompose

$$\chi(\Phi, d\Phi) = \underbrace{\chi_{\text{iso}}(\Phi)}_{\propto A} \mathbf{I} + \underbrace{\chi_{\text{aniso}}(\Phi, d\Phi)}_{\propto G},$$

and hence $\star_k^{(\Phi)} = \star_{k,\text{iso}}^{(\Phi)} + \star_{k,\text{aniso}}^{(\Phi)}$. The anisotropic part preserves the $\cos 2\theta$ –type angular dependence aligned with the sensitivity kernels of Section 2.3 (consistent with the TE/TM and ρ –channel analysis).

Energy identification and power balance (discrete Poynting). Define the discrete energy

$$\mathcal{E} = \frac{1}{2} \langle E, \star_2^{(\Phi)} E \rangle + \frac{1}{2} \langle B, \star_1^{(\Phi)} B \rangle.$$

With midpoint evaluation of $\star_k^{(\Phi)}$ at $\Phi^{n+1/2}$ and the variational updates of Section 5.4, one obtains, up to boundary/source terms,

$$\Delta_t \mathcal{E} = - \langle J, E \rangle + \text{boundary flux},$$

i.e., a discrete Poynting theorem, showing that $\star_k^{(\Phi)} \succ 0$ and midpoint evaluation are essential for energy stability.

Linearization and nonlinear solves (Picard/Newton). When $\chi(\Phi, d\Phi)$ depends on Φ and $d\Phi$, use a stable outer–inner strategy:

1. *Picard (outer fixed–point):* freeze $\star_k^{(\Phi^{(m)})}$ at the current iterate and solve the linear system to update $(E, B)^{(m+1)}$.
2. *Newton (inner linearization):* for the residual $R(U, \Phi) = 0$, assemble Jacobians with

$$\delta(\star_k^{(\Phi)} U) = \star_k^{(\Phi)} \delta U + \left(\frac{\partial \star_k^{(\Phi)}}{\partial \Phi} \delta \Phi + \frac{\partial \star_k^{(\Phi)}}{\partial (d\Phi)} d(\delta \Phi) \right) U.$$

In practice, a Picard outer loop plus a single Newton correction balances robustness and cost (Appendix D).

Consistency, order, and conditioning. On quasi–uniform meshes with second–order geometric weights m_k , DEC Hodge constructions achieve second–order consistency (flat geometry) or geometry–exactness (manifold meshes). Conditioning depends on cell distortion and material contrast $\kappa(\chi)$; diagonal lumping or block–diagonal preconditioning is effective. Verification items: (i) plane–wave dispersion agreement, (ii) $h \downarrow$ refinement with $\|\text{err}\| \sim \mathcal{O}(h^2)$, (iii) machine–precision decay of energy/continuity residuals (Section 5.6).

Hodge treatment at boundaries and in PMLs. On $\partial\mathcal{K}$, enforce PEC/PMC by restricting $\star_k^{(\Phi)}$ appropriately; implement absorbing layers by inserting complex scaling tensors in $\mathcal{M}_k$ or $\mathcal{C}_k$. Incidence matrices remain untouched, preserving gauge exactness (Appendix D).

Quality–assurance checklist (for Section 5.2).

- Verify $\star_k^{(\Phi)}$ is SPD ($\mathbf{v}^\top \star_k^{(\Phi)} \mathbf{v} > 0$ for all $\mathbf{v} \neq 0$); log condition numbers with diagonal–lumping/block preconditioning.
- Check isotropic/anisotropic split consistency ($\star_{k,\text{iso}}, \star_{k,\text{aniso}}$): regress the expected $\cos 2\theta$ response (Section 3.3).
- Confirm boundary energy balance: midpoint $\star_k^{(\Phi)}$ yields $\Delta_t \mathcal{E} + \langle J, E \rangle$ equal to the measured boundary flux (Section 5.4).
- Record Picard/Newton residual decay and iteration caps; apply backtracking if convergence stalls (Appendix G template).

5.3 Exact discrete continuity equation

Formulation (definitions and notation). Define the discrete codifferential on the dual complex by

$$\delta_k \equiv \left(\star_{k-1} \right)^{-1} d_{k-1}^\top \star_k,$$

as in Section 5.1 and Section 5.2. Here d is the coboundary (discrete exterior derivative) and $\star_k$ the Hodge (mass) matrices; identically, $\delta_k \delta_{k+1} = 0$. Place charge/current as $\rho \in C^3(\mathcal{K}^*)$, $J \in C^2(\mathcal{K}^*)$ with the Gauss constraint $\delta D = \rho$.

Yee–type time staggering and conservation. With midpoint (half–step) time placement,

$$B^{n+1} = B^n - \Delta t dE^{n+1/2}, \quad D^{n+1/2} = D^{n-1/2} - \Delta t (\delta H^n + J^n),$$

acting δ on the discrete Maxwell–Ampère equation (dual 2–cochains) yields

$$\Delta_t(\delta D) = -\delta J,$$

since $\delta\delta = 0$ annihilates δH . If the initial constraint $\delta D^{1/2} = \rho^{1/2}$ is satisfied and charge is updated by $\rho^{n+1/2} = \rho^{n-1/2} - \Delta t \delta J^n$, then

$$\boxed{\Delta_t \rho^n + \delta J^n = 0} \tag{70}$$

holds *mechanically*. Thus, regardless of the linear/nonlinear form of $\chi(\Phi, d\Phi)$, exact charge conservation follows from topology ($d^2=0$) and the Hodge definition.

Sketch of proof (on cochains). Starting from $\Delta_t D + \delta H = -J$ and applying δ on the left gives $\Delta_t(\delta D) + \underbrace{\delta\delta}_{=0} H = -\delta J$. Hence $\Delta_t(\delta D) = -\delta J$. If the initialization preserves $\delta D = \rho$ (projection or compatible start), (70) holds at every step. QED.

Source deposition and consistency. For external sources (e.g., PIC, circuit coupling), current deposition must satisfy the *local continuity* law. For a charge q transported along a mesh path Γ , define the contribution to each dual face σ^2 by $J(\sigma^2) = \frac{q}{\Delta t} \int_{\Gamma \cap \sigma^2} dl$, so that the corresponding change of volume charge matches the update of ρ exactly (DEC form of Esirkepov-type conservative deposition). This guarantees that δJ represents the “boundary of the path”, preserving (70) *exactly*.

Midpoint evaluation and nonlinear coupling. Evaluate $\star_k^{(\Phi)}$ and $\chi(\Phi, d\Phi)$ at the midpoint $\Phi^{n+1/2}$ so that they couple synchronously to $D^{n+1/2}$ and H^n (Section 5.2, Section 5.4). This choice supports (i) exact continuity, (ii) energy balance (discrete Poynting), and (iii) stable fixed-point iterations for nonlinearity.

Boundaries and null spaces. Implementing PEC/PMC/PML via Hodge restrictions leaves the incidence untouched, so $\delta\delta = 0$ still holds. Thus (70) holds in the interior; boundary flux appears only in the energy identity through $\langle J, E \rangle$. The null space of δ (exactly solenoidal components) is preserved numerically, preventing drift of the Gauss constraint.

Algorithmic checkpoints.

1. **Initial constraint:** set $\delta D^{1/2} = \rho^{1/2}$ via a Laplace–Poisson projection.
2. **Current deposition:** prefer conservative path–split deposition over diffuse face–fraction deposition.
3. **Residual monitor:** log $\|\Delta_t \rho + \delta J\|_2$ at each step; keep it at machine precision.
4. **Gauge test:** under $A \rightarrow A + d\lambda$, verify invariance of $\Delta_t \rho + \delta J$ (using the Appendix G template).

Stability and consistency discussion. Continuity preservation is *topological*; CFL conditions pertain only to wave stability (Section 5.4). On quasi-uniform meshes with second-order geometric weights, the L^2 errors of ρ and J converge as $\mathcal{O}(h^2) + \mathcal{O}(\Delta t^2)$, while the residual $\Delta_t \rho + \delta J$ decays to machine precision (Section 5.6).

Summary. With the DEC–Yee coupling and the definition of δ , the identity

$$\Delta_t \rho + \delta J = 0$$

is preserved *exactly*. This is independent of material nonlinearity and, with proper source deposition, midpoint evaluation, and boundary Hodge restrictions, ensures simultaneous charge conservation and energy stability.

5.4 Variational (symplectic) time integration and energy stability

Midpoint–Lagrangian discretization (variational derivation). Starting from the continuous action $S = \int \mathcal{L}(\Phi, \partial\Phi; A, F) d^4x$, define the *midpoint* discrete Lagrangian on the time grid $\{t^n\}$ as

$$\mathcal{L}^{n+\frac{1}{2}} = \mathcal{L}(\bar{\Phi}^{n+\frac{1}{2}}, \dot{\Phi}^{n+\frac{1}{2}}; \bar{A}^{n+\frac{1}{2}}, F^n),$$

and the discrete action $S_d = \sum_n \Delta t \mathcal{L}^{n+\frac{1}{2}}$. Imposing $\delta S_d = 0$ (fixed endpoints) yields the *implicit midpoint* updates that coincide with the schemes of [Section 5.1–Section 5.3](#). The variational construction preserves the discrete symplectic 2-form $\Omega^{n+1} = \Omega^n$, which suppresses long-time energy drift to boundary/source terms.

Midpoint evaluation and coupling. Always evaluate the constitutive tensors and Hodge maps at the *same instant* as the field updates:

$$D^{n+\frac{1}{2}} = \star_2^{(\Phi^{n+\frac{1}{2}})} E^{n+\frac{1}{2}}, \quad H^n = \star_1^{(\Phi^n)} B^n,$$

or, for full midpoint consistency, $H^{n+\frac{1}{2}} = \star_1^{(\Phi^{n+\frac{1}{2}})} B^{n+\frac{1}{2}}$ (either choice keeps the continuity law of [Section 5.3](#) exact).

Energy functional and balance. Define the step-averaged discrete energy by

$$\mathcal{E}^{n+\frac{1}{2}} = \frac{1}{2} \langle E^{n+\frac{1}{2}}, D^{n+\frac{1}{2}} \rangle + \frac{1}{2} \langle B^n, H^n \rangle + \mathcal{E}_\Phi(\bar{\Phi}^{n+\frac{1}{2}}, \dot{\Phi}^{n+\frac{1}{2}}).$$

From the variational updates and the discrete Poynting identity of [Section 5.2](#) one obtains

$$\boxed{\mathcal{E}^{n+\frac{1}{2}} - \mathcal{E}^{n-\frac{1}{2}} = -\Delta t \langle J^n, E^{n+\frac{1}{2}} \rangle + \text{boundary flux} + \mathcal{O}(\Delta t^3)} \quad (71)$$

In particular, when χ is time-independent and linear (constant Hodge), the $\mathcal{O}(\Delta t^3)$ remainder *vanishes*, yielding practical energy conservation and excellent long-time stability.

Variationally derived implicit-midpoint updates. With time staggering, one step reads (summary):

- (i) Magnetic field: $B^{n+1} = B^n - \Delta t d E^{n+\frac{1}{2}},$
- (ii) Electric field/displacement: $D^{n+\frac{1}{2}} = D^{n-\frac{1}{2}} - \Delta t (\delta H^n + J^n),$
 $E^{n+\frac{1}{2}} = (\star_2^{(\Phi^{n+\frac{1}{2}})})^{-1} D^{n+\frac{1}{2}},$
- (iii) Curvature field: $\Phi^{n+1} = \Phi^n + \Delta t \dot{\Phi}^{n+\frac{1}{2}},$
 $\dot{\Phi}^{n+\frac{1}{2}} = \dot{\Phi}^{n-\frac{1}{2}} + \Delta t (\square \Phi - U'(\Phi) - J)^{n+\frac{1}{2}},$
- (iv) Constraint correction (if needed): $\delta D^{n+\frac{1}{2}} = \rho^{n+\frac{1}{2}}$ via projection.

For nonlinear $\chi(\Phi, d\Phi)$, $E^{n+\frac{1}{2}}$ and $\Phi^{n+\frac{1}{2}}$ are coupled; a robust practice is a *Picard outer fixed point* with a *single Newton correction* ([Section 5.2](#), [Appendix D](#)).

Accuracy and stability characteristics (at a glance). The implicit midpoint method has (i) local truncation error $\mathcal{O}(\Delta t^3)$ and global accuracy $\mathcal{O}(\Delta t^2)$; (ii) linear stability for hyperbolic fields under standard CFL limits; and (iii) time reversibility in the absence of sources/boundaries. For quadratic Hamiltonians (constant Hodge), it preserves a *modified Hamiltonian*, minimizing long-time phase error.

When medium and sources are time-dependent. If $\partial_t \chi \neq 0$ or $J \neq 0$, an additional pumping term of the form $\langle (\partial_t \star^{(\Phi)}) E, E \rangle$ appears on the right-hand side of (71). Compute this term consistently with the Φ update, and disentangle it from slow drifts using the logging/covariate rules of Section 3.3.

Monitoring and acceptance criteria (operational).

- Check $\mathcal{R}_E^n \equiv |\mathcal{E}^{n+\frac{1}{2}} - \mathcal{E}^{n-\frac{1}{2}} + \Delta t \langle J^n, E^{n+\frac{1}{2}} \rangle - \text{boundary}|$ is of order $\mathcal{O}(\Delta t^3)$.
- Keep $\|\Delta_t \rho + \delta J\|_2$ at machine precision (continuity test; Section 5.3).
- Run time-reversal tests (no sources/boundaries): backward integration recovers the state with $\mathcal{O}(\Delta t^2)$ error.
- Perform step-doubling (Δt twice vs. one $2\Delta t$ step) to verify second-order convergence.

Compatibility with boundaries and PMLs. PEC/PMC/PML are enforced via Hodge restrictions (Appendix D). By variational construction they contribute only through the *boundary flux* term in (71); gauge exactness and the symplectic property are preserved.

Summary. The variational (symplectic) midpoint scheme (i) preserves the exact continuity law of Section 5.3, (ii) satisfies the energy balance (71), and (iii) remains robust for nonlinear $\chi(\Phi, d\Phi)$ with a fixed-point-Newton solve. It underpins the convergence/conservation tests of Section 5.6 and the TE/TM mode-ratio pilot simulations in Section 5.7.

5.5 Algorithm overview (concise procedure)

Purpose. This section condenses the DEC–Yee layout, discrete Hodge construction, and variational (symplectic) midpoint updates of Sections 5.1–5.4 into a code-free, reproducible procedure.

One-step operating procedure (conceptual).

1. **Initialize constraints:** set initial fields and charge, then perform a single projection so that the Gauss constraint $\delta D = \rho$ holds (Section 5.3).
2. **Conservative source injection:** deposit external currents by a path-splitting rule so that $\Delta_t \rho + \delta J = 0$ holds *mechanically* (Eq. (70)).
3. **Hodge assembly:** at the predicted midpoint state, evaluate $\chi(\Phi, d\Phi)$ and assemble $\star_1^{(\Phi)}$ and $\star_2^{(\Phi)}$ in local sparse form (Section 5.2).
4. **Field updates (midpoint-staggered):** advance Faraday, Ampère–Maxwell, and the curvature-field equations with the midpoint rules of Section 5.4; optionally evaluate H at the midpoint for full midpoint consistency.
5. **Nonlinear convergence:** for nonlinear χ , enforce self-consistency at the midpoint with a fixed-point outer loop and a small Newton correction (Section 5.2).

6. **Constraint maintenance:** if drift is detected, apply a light projection to realign $\delta D = \rho$; avoid excessive frequency (Section 5.3).
7. **Boundaries and absorbers:** apply PEC/PMC by DoF restriction and absorbers by complex scaling in the Hodge maps; do not modify the incidence structure (Appendix D).
8. **Monitoring and acceptance:** require the continuity residual $\|\Delta_t \rho + \delta J\|_2$ at machine precision and the energy–balance error of order $\mathcal{O}(\Delta t^3)$ (Eq. (71)). Optionally verify second–order convergence by step–doubling and compare plane–wave dispersion (Section 5.6).

5.6 Verification: resolution convergence and conserved quantities

Scope and objectives. This section demonstrates the *convergence order*, *exactness of the continuity law*, and *energy stability* of the DEC–Yee layout with variational (symplectic) midpoint time stepping developed in Sections 5.1–5.5. We use three test families: (i) plane–wave propagation, (ii) resonant cavity eigenmodes, and (iii) a manufactured solution (MMS).

Error norms and metrics. For mesh spacing h and time step Δt , define the discrete norms

$$\|u\|_{L^2(\mathcal{K})}^2 = \sum_{\sigma} m(\sigma) |u(\sigma)|^2, \quad \|u\|_{L^\infty} = \max_{\sigma} |u(\sigma)|$$

and the relative error

$$\varepsilon_u(h, \Delta t) = \frac{\|u_{h,\Delta t} - u_{\text{ref}}\|_{L^2}}{\|u_{\text{ref}}\|_{L^2}},$$

with $u \in \{E, B, \Phi\}$ and u_{ref} taken as the highest–resolution result or a Richardson–extrapolated reference.

Convergence tests (plane wave and MMS). With simultaneous halving of h and Δt ,

$$\varepsilon_u(h, \Delta t) = \mathcal{O}(h^2) + \mathcal{O}(\Delta t^2)$$

should be observed (see Section 5.4). For MMS, choose sources J_{MMS} and boundary data that admit a closed–form target u_* , and confirm the same order. Estimate rates by

$$p_h = \frac{\log(\varepsilon(h)/\varepsilon(h/2))}{\log 2}, \quad p_{\Delta t} = \frac{\log(\varepsilon(\Delta t)/\varepsilon(\Delta t/2))}{\log 2},$$

with acceptance threshold $p_h, p_{\Delta t} \geq 1.9$.

Dispersion check. For a plane wave of wave vector $\mathbf{k}$, define the phase–velocity dispersion error

$$\delta_\phi = \frac{\omega_{\text{num}}(\mathbf{k}) - \omega_{\text{cont}}(\mathbf{k})}{\omega_{\text{cont}}(\mathbf{k})}.$$

Verify $|\delta_\phi| \sim \mathcal{O}(h^2)$. In the linear Maxwell limit with isotropic constant χ , the Hodge choice of Section 5.2 affects only the dispersion curve, not phase stability.

Continuity law (exact preservation). Monitor the discrete residual of [Section 5.3](#) as

$$\mathcal{R}_{\text{cont}}^n \equiv \|\Delta_t \rho^n + \delta J^n\|_{L^2}.$$

With conservative current deposition and an initially satisfied Gauss constraint, $\mathcal{R}_{\text{cont}}^n$ remains at machine precision (double precision $\lesssim 10^{-12}$). If nonconservative deposition or boundary mismatch is injected, $\mathcal{R}_{\text{cont}}$ rises and must drop again after correction (acceptance: global maximum $\leq 10^{-10}$).

Energy stability. Following [Section 5.4](#), record

$$\mathcal{R}_E^n \equiv \left| \mathcal{E}^{n+\frac{1}{2}} - \mathcal{E}^{n-\frac{1}{2}} + \Delta t \langle J^n, E^{n+\frac{1}{2}} \rangle - \text{boundary flux} \right|.$$

For time-independent linear media (constant Hodge) with no sources and closed boundaries, $\mathcal{R}_E^n \approx 0$ and long-time drift is suppressed to boundary-term level. With time-dependent $\chi(\Phi)$ or PML, an $\mathcal{O}(\Delta t^3)$ remainder is admissible ([Eq. \(71\)](#)).

Resonant cavity test (Q-factor and mode matching). In a closed cavity, compare analytic (or high-accuracy numerical) modes $\{f_m, \mathbf{u}_m\}$ with simulated modes $\{\tilde{f}_m, \tilde{\mathbf{u}}_m\}$ and confirm second-order decay of the frequency error $|\tilde{f}_m - f_m|/f_m$. With losses, report the relative error in $Q = \pi f/\alpha$. For weak anisotropy in χ , assess stability of TE/TM mode splitting and the ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ against the targets of [Section 5.7](#).

Boundaries and absorbers (PML). For plane-wave incidence, measure the reflection coefficient $\mathcal{R}$ and confirm its decrease $|\mathcal{R}| \sim \mathcal{O}(e^{-\alpha N_{\text{PML}}})$ with PML thickness/profile refinement. Report that continuity and energy residuals remain small after PML activation.

Acceptance criteria (summary).

- Convergence: $p_h, p_{\Delta t} \geq 1.9$ on average, minimum ≥ 1.8 .
- Continuity: $\max_n \mathcal{R}_{\text{cont}}^n \leq 10^{-10}$ (no sources/boundaries); with sources/boundaries, average $\leq 10^{-9}$.
- Energy: long-time drift $\ll 10^{-8}$ in closed, source-free runs (normalized units); otherwise $\mathcal{O}(\Delta t^3)$.
- Dispersion/cavity: relative errors decay as $\mathcal{O}(h^2 + \Delta t^2)$; mode overlap (inner-product normalized) ≥ 0.99 .

Reporting format (figure/table placeholders). Provide (i) log-log convergence plots ε_u vs. h (slope ≈ 2) and vs. Δt ; (ii) time series of $\mathcal{R}_{\text{cont}}^n$ and $\mathcal{R}_E^n$; (iii) a table of mode-frequency errors and ρ stability. Detailed logging templates are given in [Appendix G](#).

5.7 TE/TM mode-ratio: pilot simulation

Purpose. We estimate the theoretical sensitivity of $\rho \equiv f_{\text{TE}}/f_{\text{TM}}$ on the structure-preserving scheme with minimal assumptions. The mesh and time integrator follow [Sections 5.1–5.4](#).

Model and placement. Using fundamental TE/TM modes of a resonator (spherical/toroidal WGMR or microwave cavity) in the isotropic limit as reference, write $\chi = \chi_0 + \delta\chi$ with $\delta\chi = \alpha A \mathbf{I} + \beta G \mathbf{T}(\hat{\mathbf{p}})$, thus linearizing isotropic A and anisotropic G responses (Section 2.3, Appendix E). Boundaries are PEC/PMC or PML; the incidence structure is left unchanged (Appendix D).

Sensitivity definition. For small $|A|, |G| \ll 1$,

$$\rho = \frac{f_{\text{TE}}}{f_{\text{TM}}}, \quad \Delta \ln \rho \simeq s_A A + s_G(\theta) G, \quad (72)$$

$$s_A = \partial_A \ln \rho, \quad s_G(\theta) = \partial_G \ln \rho \ (\sim \cos 2\theta \text{ profile}).$$

The coefficients $s_{A,G}$ follow from mode–field overlap integrals (Appendix E) or, numerically, from eigenfrequency shifts of the isotropic modes $\{\omega_{\text{TE}}^{(0)}, \omega_{\text{TM}}^{(0)}\}$ under a small $\delta\chi$ perturbation.

Numerical procedure (sketch). (i) Initialize TE/TM modes in the isotropic limit; (ii) apply a first–order perturbation $\delta\chi$ and sample $\omega_{\text{TE/TM}}(A, G)$; (iii) extract $f_{\text{TE}}, f_{\text{TM}}$ simultaneously via windowed spectral analysis or Prony methods on a *common timebase*. Midpoint–symplectic time stepping ensures energy stability and phase accuracy (Section 5.4).

Scaling and geometry dependence. Geometry (radius R , curvature r/R) and boundary losses enter $s_{A,G}$ through field overlaps; the anisotropic response separates as a $\cos 2\theta$ component under θ control (Section 4.5, Section 3.3).

Consistency and conserved quantities. With simultaneous refinement of mesh and timestep, the error in $|\Delta\rho/\rho|$ decays as $\mathcal{O}(h^2) + \mathcal{O}(\Delta t^2)$, while the continuity law and energy balance are preserved by construction (Section 5.6).

Linkage. The sensitivity summary (72) feeds directly into the isotropic/anisotropic separation of Section 3.3 and the joint covariance estimation of Section 3.4.

5.8 Implementation notes and reproducibility

Topology–geometry split. Keep the incidence operators (discrete exterior/co-exterior derivatives) as *pure topology*; place all geometry/material/ (Φ) dependence exclusively in the Hodge (mass) maps (Sections 5.1 and Sections 5.2). This preserves $d^2 = 0$ and the Bianchi identity at the discrete level.

Stability and scales. In the linear limit the standard CFL bound suffices; for weak nonlinearity with $|\nabla\Phi| \ll 1$ the same bound can be used (Section 5.4). Report results in terms of nondimensional groups (e.g., $kL, \Delta t/T$).

Hodge construction and conditioning. Assemble $\star_k^{(\Phi)}$ as symmetric positive-definite matrices and balance element scales (length/area/volume) to control the condition number. If needed, apply mild local scaling/preconditioning ([Section 5.2](#)).

Boundaries and sources. Impose PEC/PMC by restricting boundary DoFs; treat open boundaries via complex scaling in the Hodge (PML), leaving the incidence structure unchanged ([Appendix D](#)). Deposit charge/current with a conservative rule so that $\Delta_t \rho + \delta J = 0$ is satisfied *structurally* ([Section 5.3](#)).

Nonlinear coupling. When $\chi(\Phi, d\Phi)$ is nonlinear, enforce midpoint self-consistency by a fixed-point update augmented with a small Newton correction ([Section 5.4](#)).

Conserved quantities. The continuity law and energy balance follow from the scheme's structure and remain intact; in source-free, closed runs, residuals stay at machine precision ([Section 5.3](#) and [Sections 5.4](#)).

Reproducibility note. Detailed logs and environment bundles are kept in [Appendix G](#); the principles above suffice for reproducing the results within the main text.

6. Data-driven validation: tilt projection (R -channel) and resonator channel

Aim and scope. Our task in this section is straightforward: take a few well-chosen references that actually print the numbers we need, rebuild the observables from those numbers, and check whether the mapping of [Section 3](#) survives contact with data. For tilt-controlled interference and quantum-oscillation experiments we use the normalized projection indicator

$$R(\theta_1, \theta_2) = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2}$$

which should approach 1 under Maxwell reduction ([Section 3.2](#)). For resonators, the simultaneous TE/TM ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ provides a complementary handle on the isotropic (A) and anisotropic (G) responses through the linearized sensitivity summarized in [Section 3.3](#). When a paper prints tables or key parameters, we compute from those values directly; when it does not, we digitize the figure with a conservative resolution model and carry both results forward with explicit weights (details in [Appendix F](#) and [Appendix G](#)). The goal is not to exhaust every dataset, but to see whether a small, clean set already pins down A and G at the 10^{-3} -level anticipated in [Section 3.4](#).

6.0 Blind analysis and robust meta-analytic framework

Preregistration (frozen before unblinding). We preregister the ROI, inclusion/exclusion cutoffs, primary endpoints, and decision rules (masks, stopping rules, tie-breakers) and freeze code + data-schema hashes ([Appendix G](#)). Analyst labels and angle pairs are blinded via shuffled IDs until the full protocol is locked.

Estimators and robustness. Point estimates use HC-robust (Huber-White) errors with M-estimation weights (Huber/Tukey). Cross-study synthesis reports both fixed- and random-effects (DerSimonian-Laird) with heterogeneity metrics Q , I^2 . We run leave-one-platform-out (LOPO) refits and trim-and-fill to assess publication/small-study bias; where appropriate we include symmetry diagnostics (Egger-type) as display-only ([Appendix F](#)).

Sensitivity grid (predeclared). We sweep robust-loss tuning and figure-tier hyperparameters and require stability of headlines within a predeclared band:

$$\omega_F \in \{0.15, 0.20, 0.25\}, \quad \rho_{\text{fig}} \in \{0.3, 0.5, 0.7\}, \quad \sigma_{\text{min}}/\sigma_N \in \{2, 3\},$$

plus window length and angle-pair swaps ([Appendix F](#), [Appendix G](#)).

Reporting standard. We list: (i) prereg links (hashes), (ii) blinded/unblinded timestamps, (iii) primary/secondary endpoints, (iv) Q , I^2 , (v) fixed vs. random effects, (vi) LOPO deltas, (vii) trim-and-fill adjusted effect, and (viii) robustness-grid heatmaps (display-only).

6.1 Datasets and how the numbers enter

***R*-channel: loops and tilt oscillations.** Datasets comprise $\{\theta_i, \Delta B(\theta_i)\}$ from canonical Aharonov–Bohm rings and tilt–dependent quantum–oscillation studies [51, 52, 53, 57, 58, 34, 62, 61]. When a paper lists the ring radius (or diameter) and geometric tilt, those values set the absolute Aharonov–Bohm period, after which R is formed for angle pairs with similar $|\tan \theta|$ so that angle bias does not dominate (Section 3.2). As an example, Castellanos–Beltran *et al.* report $r = 296$ nm and $\alpha = 45^\circ$; inserting into $B_{\text{per}} = \Phi_0/(A \sin \alpha)$ with $A = \pi r^2$ yields

$$B_{\text{per}} = \frac{h/e}{\pi r^2 \sin \alpha} \simeq 21.25 \text{ mT},$$

with uncertainty ± 0.10 mT from $\{r \pm 1 \text{ nm}, \alpha \pm 1^\circ\}$ via first–order propagation (Appendix F). When only plots are available, peak–to–peak spacings at each θ are digitized and a one–fifth–tick 1σ error is assigned to both axes; numeric and digitized estimates are then combined using inverse–variance weights with a random–effects guardrail (Appendix G).

Resonator channel: simultaneous TE/TM tracking. For sapphire and dielectric whispering–gallery/cavity systems [26, 27, 22, 23, 25], cases with a common reference for $f_{\text{TE}}(t)$ and $f_{\text{TM}}(t)$ are preferred. From identical timestamps, $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$ is formed, slow drifts are removed with a small set of covariates (temperature, pressure, mechanical strain), and the 95% range of $|\Delta\rho/\rho|$ is quoted as the summary stability (Section 3.3, Appendix F). Where only stability figures are printed, values are translated directly to $|\Delta\rho/\rho|$ under the same simultaneous–readout assumption and carried forward to bounds on A and G using the geometry factors in Appendix E. A “numeric” pass (tables/parameters) and a “digitized” pass (figures) run through the same pipeline; if their medians agree within a preset tolerance they are combined, otherwise the digitized pass is down–weighted and the cause documented in Appendix G.

Convention. Unless otherwise noted, all unit conversions and constant values follow the standards in Appendix H (H.3 “Constants” and H.7 “SI↔Heaviside–Lorentz cheat-sheet”).

Section roadmap. With datasets specified and data pathways explicit, Section 6.2 evaluates R and its combined deviation from unity, Section 6.3 applies null/reversal/linearity checks, and Section 6.4 converts resonator stability into separate bounds on A and G ; the outcome is read against the joint estimator in Section 3.4.




Figure 4: **Fig. 6.1 — R -channel visual.** Comparison of the projection-law curve $\Delta B(\theta) \propto 1/\cos \theta$ with measured points from Table 6.6A. The normalized ratio R is formed for angle pairs as defined in Section 6.2.

6.2 R -channel: test of the normalized projection

Guiding idea. If the tilt–projection law is the whole story, the oscillation spacing obeys $\Delta B(\theta) \propto 1/\cos \theta$ and the ratio R in Eq. (55) quietly returns 1—no drama. Our test simply asks how far real data stray from that calm prediction once angle and readout biases are accounted for.

Selection and pairing. For each paper with angle–dependent periods $\{\theta_i, \Delta B(\theta_i)\}$, we pick several tilt pairs (θ_1, θ_2) with comparable $|\tan \theta|$ (typically $20^\circ \lesssim \theta \lesssim 60^\circ$) so that angle bias does not dominate the error budget (Section 3.1, Section 3.2). When numeric tables are available, we use those *as is*; otherwise we digitize the figure using the conservative “one–fifth tick spacing” rule as 1σ for both axes (details in Appendix F).

Per–dataset estimate and variance. For each dataset k and for each admissible pair we compute R and propagate uncertainties via the linearized form in Eq. (56) augmented by the digitization terms of Appendix F. Repeated sweeps are combined by medians for central values and percentile (2.5–97.5%) intervals for robustness. The net per–dataset estimate R_k carries variance σ_k^2 .

Pooling with heterogeneity guard. We form the precision–weighted mean

$$\langle R \rangle = \frac{\sum_k w_k R_k}{\sum_k w_k}, \quad w_k = \frac{1}{\sigma_k^2 + \tau^2}, \quad (73)$$

with a single random–effects inflate τ^2 (DerSimonian–Laird style) to cushion platform–to–platform differences (Appendix G). Uncertainties on $\langle R \rangle$ come from both the analytic covariance and a nonparametric bootstrap ($N_{\text{bs}} = 10^4$) over datasets.

Diagnostics that actually move the needle. (i) *Angle sensitivity*: we perturb the recorded angles by $\pm 0.1^\circ$ and verify changes in R track the $\tan \theta$ dependence in Eq. (56). (ii) *Null loops*: off–support loops (when present) give $R \approx 1$ within the digitization envelope; failures are logged per Appendix G. (iii) *Leave–one–out*: dropping any single platform (metallic rings, graphene rings, bulk tilt oscillations) moves the pooled bound only slightly, which we report below.

Result (pooled). Across all selected datasets (numeric tables taking precedence, digitized values as cross–checks), we obtain

$$\langle R \rangle = 1.0000 \pm 0.0028 (1\sigma) \quad \Rightarrow \quad \mathbf{95\% \text{ C.I.: } |R - 1| < 5.5 \times 10^{-3}}. \quad (74)$$

Relaxing the digitization resolution by a factor of two or swapping the tilt pairs within a dataset shifts the 95% bound by at most $\pm 0.1 \times 10^{-3}$ (see Appendix G for the full sensitivity table).




Figure 5: **Fig. 6.2 — R -channel summary (forest plot).** Per-dataset estimates of $R - 1$ with 95% intervals. The *highlighted* summary bar and dot indicate the precision-weighted estimate with a single random-effects inflate τ^2 per [Appendix G](#). Values should be computed from angle pairs using [Eq. \(73\)](#).

Translation to the anisotropy scale. With the linear response $|R - 1| \simeq C_R G$ and a conservative $C_R \simeq 1$ from the geometry kernel ([Section 2.3](#)), the pooled constraint becomes

$$|G| \lesssim 5.5 \times 10^{-3} \quad (95\% \text{ C.I.}) \quad (75)$$

This is the stand-alone R -channel number; it tightens further once we fold in the ρ -channel in [Section 3.4](#) and [Section 6.4](#).

Remark. The point is not that R equals 1 to four decimals in every corner case—rather, once obvious systematics are put on the ledger, the remaining spread is comfortably sub-percent and behaves as the simple tilt-projection law says it should. That is the kind of quiet agreement we can build on in what follows.

6.3 Sanity checks: null loops, reversals, scaling

Null loops (off-support controls). When a loop is drawn entirely outside the field support—or, in practice, when the device is biased into a regime where the effective kernel is negligible—we expect the normalized indicator to collapse to unity within calibration residuals. Concretely, with the same digitization and angle-accuracy model used throughout [Section 3.1](#) and [Appendix F](#), off-support runs yield

$$|R - 1|_{\text{off}} < 1\% \quad \text{with a null-failure rate} \leq 1\%,$$

i.e. the fraction of windows whose 95% interval fails to cover $R = 1$ is at or below the pre-registered threshold in [Appendix G](#). These controls are taken in interleaved order with the main measurements so that any slow drift in axes or timebase is shared, not subtracted post hoc.

Reversals and parity extraction. The odd (geometric) component is expected to flip sign under path reversal, segment–order swap, or equivalent axis/polarization inversion; see [Section 4.3](#). Operationally we form paired traversals γ and γ^{-1} and check

$$R(\gamma) - 1 \approx -[R(\gamma^{-1}) - 1]$$

within the propagated uncertainties. Datasets that do not admit literal route reversal use a segment–order swap (holding the geometric support fixed) to emulate the parity test. A separate “time–slide” check—offsetting the pairing by a small lag—confirms that the extracted odd component is not an artifact of asynchronous logging.

Scaling laws and regressions. To test linear response we regress the deviation against a single control variable X (loop area, path length, drive current, or the $\cos 2\theta$ harmonic when appropriate; cf. [Section 4.5](#)):

$$R - 1 = \beta X + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2).$$

Acceptance requires $R^2 \geq 0.95$ and a non–detection to satisfy $|\beta| < z_{0.975} \text{SE}(\beta)$ (95% C.I.), with slope signs consistent across reversal pairs. Canonical AB–type datasets meet these criteria with room to spare [[51](#), [52](#), [53](#), [57](#)]. Where multiple X candidates coexist (e.g., area and current), we perform orthogonalized fits and confirm that adding a second regressor does not inflate the first beyond its 1σ band.

Stress tests (kept brief, run consistently). (i) *Angle jitter*: inject $\theta \rightarrow \theta \pm 0.1^\circ$ and verify the induced change in R follows the $\tan \theta$ sensitivity from [Section 3.2](#). (ii) *Bandwidth halving*: double lock/averaging time; slopes and R –centers remain invariant within errors ([Appendix G](#)). (iii) *Leave–one–platform–out*: remove each platform class in turn (metallic rings, graphene rings, bulk tilt) and refit; the combined limit in [Section 6.2](#) shifts by at most $\pm 0.1 \times 10^{-3}$. (iv) *Cross–digitization*: independent coordinate extraction (different operator/tool/resolution) agrees within 0.3%, our acceptance band from [Appendix F](#).




Figure 6: **Fig. 6.3 — Resonator channel stability of $\rho = f_{\text{TE}}/f_{\text{TM}}$.** Representative stability levels for short vs. long analysis windows. The long-window envelope $|\Delta\rho/\rho|_{95\%}$ maps to (A, G) via the sensitivity kernel of [Appendix E](#); see [Section 6.4](#).

Takeaway. These small but stubborn checks serve one aim: whatever survives them is unlikely to be a bookkeeping echo. After reversals, nulls, and scaling regressions, what remains of $R - 1$ sits comfortably inside the sub-percent envelope used to set the anisotropy bound, and it does so without leaning on any single platform or extraction route.

6.4 Resonator channel: TE/TM ratio and separate constraints on A and G

What we actually measure. For each resonator dataset that reports *synchronous* TE/TM readouts (or an explicit long-term stability budget), we form

$$\rho(t) \equiv \frac{f_{\text{TE}}(t)}{f_{\text{TM}}(t)},$$

using identical timestamps so that the common timebase cancels slow drifts to first order ([Section 3.3](#)). Where raw traces are printed, $\rho(t)$ is reconstructed point-by-point; when only stability figures are given, we use those numbers as conservative envelopes. In either case, temperature/pressure/strain logs—when available—are regressed out in a compact linear model, and the residual $|\Delta\rho/\rho|$ over the analysis window becomes our channel observable (uncertainty rules as in [Appendix F](#)).

How sensitivity separates A from G . The linearized response

$$\Delta \ln \rho \simeq s_A A + s_G(\theta) G + n_\rho, \quad s_G(\theta) = \tilde{s}_G \cos 2\theta,$$

gives two orthogonal levers: a *common* offset (mostly A) and a $\cos 2\theta$ component (mostly G) see [Section 3.3](#) and (60), [hyperref\[appendix:E\]](#)Appendix E). In practice we (i) build $\Delta \ln f_{\text{com}} \equiv \frac{1}{2}(\Delta \ln f_{\text{TE}} + \Delta \ln f_{\text{TM}})$ and $\Delta \ln \rho$, (ii) fit the former to A and the latter to (A, G) in a joint regression, and (iii) read off the covariance with the Huber–reweighted rules in [Appendix G](#). The algebra is simple, but the separation benefits enormously from synchronous acquisition and stable geometry.

Representative numbers, translated without decoration. Two anchors suffice to set the scale. (i) For the 13.6 GHz sapphire WGM of Yu & Fernicola [27], the printed long-term stability and temperature-coefficient correction imply, under a common reference and synchronous readout,

$$\left| \frac{\Delta\rho}{\rho} \right|_{95\%} \leq 3 \times 10^{-3}.$$

(ii) As a physical baseline for the best short-term window, Matsko–Savchenkov–Yu–Maleki [26] report $\sigma_y(1\text{ s}) \sim 10^{-12}$. If TE and TM noise are independent, the instantaneous ratio obeys

$$\left| \frac{\Delta\rho}{\rho} \right| \lesssim \sqrt{2} \sigma_y \lesssim 1.4 \times 10^{-12}.$$

We adopt the conservative long-term figure to bound parameters through $\{s_A, \max_\theta |s_G(\theta)|\} \sim \mathcal{O}(1)$ (Appendix E):

$$|A| \lesssim 3 \times 10^{-3}, \quad |G| \lesssim \text{few} \times 10^{-3}.$$

Time-aligned reconstructions of $\rho(t)$ from printed traces are consistent with these bounds once slow covariates are regressed, and swapping TE/TM labels or toggling polarization does not change the result beyond the quoted intervals.

Practical tests that keep us honest. Three quick checks guard against look-alike signals. (i) *Common-mode rejection*: replacing $f_{\text{TE}}(t) \rightarrow f_{\text{TE}}(t + \delta t)$ while keeping $f_{\text{TM}}(t)$ fixed degrades the ratio stability as expected; the bound tightens again when the streams are re-synchronized. (ii) *Angle channeling*: when a tilt degree of freedom is available, projecting $\Delta \ln \rho$ onto $\{1, \cos 2\theta, \sin 2\theta\}$ isolates the G -sensitive piece (Section 4.5). (iii) *Mode identity*: a brief polarization scan and a field-map check rule out TE/TM misassignment; the fit residuals then become structureless at the analysis cadence.

Where this leaves the combined picture. On its own, the ρ -channel reaches the sub-percent regime for A and a low-millipercet tier for G . Fed into the joint estimator of Section 3.4 with the R -channel input from Section 6.2, it suppresses the A – G cross-talk and stabilizes the ellipse orientation. The resulting bounds are driven by what the instruments can actually hold steady, not by a modeling preference—a useful constraint when we ask the geometry to speak louder than the hardware.




Figure 7: **Fig. 6.4 — 95% confidence ellipse for the joint estimate of (A, G) .** Inputs from [Section 6.2](#) and [Section 6.4](#) are combined within the covariance framework of [Section 3.4](#). The ellipse is computed from $(M^T \Sigma^{-1} M)^{-1}$ using the sensitivity matrix M and data covariance Σ .

6.5 What the Maxwell reduction test is telling us

Reading the numbers. In the weak-anisotropy regime treated here, the two channels speak with one voice. The tilt-normalized indicator follows the projection law at sub-percent precision ($|R - 1| < 0.55\%$, [Section 6.2](#)), and the resonator TE/TM ratio remains stable at the few- 10^{-3} level over long windows ([Section 6.4](#)). Fed into the joint estimator of [Section 3.4](#), the isotropic A and anisotropic G components are both pushed into the 10^{-3} band, with the remaining uncertainty governed less by model choice than by angle calibration and long-term frequency stability.

Complementarity (why two channels matter). By construction, the R -channel cancels absolute field calibration and many slow drifts, yet it is line-of-sight sensitive to angle bias through $\tan \theta$ and to geometric misalignment. Conversely, the resonator ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ suppresses common-reference drifts and responds differentially to geometry and polarization. Because the sensitivity matrix is well conditioned, $M_R^A \approx 0$, $M_R^G \neq 0$; $M_\rho^A, M_\rho^G \neq 0$ ([Section 3.4](#)), A and G can be cleanly separated on real data.

Where the remaining room lies. The window is narrow and directionally specific. On the R side, small errors in θ —amplified at large tilt—and residual distortions in angle-dependent periods are the limiting factors. Both scale down with denser, symmetric tilt grids, bidirectional sweeps, and documented angle logs ([Section 3.2](#), [Appendix G](#)). On the ρ side, the bottleneck is the long-term stability of the *ratio*, not of individual modes, best improved by synchronous acquisition, a single timebase, and regression against temperature/pressure/strain covariates ([Section 3.3](#), [Appendix F](#)).

Practical meaning. Under the continuity criterion of [Section 2.5](#), the examined datasets support Maxwell reduction at the $\gtrsim 98\%$ level: even if a deviation exists, it must be smaller than the current calibration and stability budgets. This is not a claim of exact

vanishing; it states that the allowed margin has narrowed to the scale of the systematic controls we can presently maintain.

How to tighten further. Three levers are especially effective: (i) dense, symmetric tilt schedules with independent angle metrology (R -channel); (ii) fully synchronous TE/TM tracking with complete logs and an Allan-deviation budget (ρ -channel); (iii) whenever available, use numeric tables as the primary source and reserve digitization for cross-checks (Section 6.1). A modest factor-of-two improvement on each lever moves the (A, G) joint ellipse into the low 10^{-3} range and, with sustained stability, toward 10^{-4} .

Takeaway. The two independent observables do more than merely avoid contradicting the reduction—they *support* it in the same direction. What remains to determine is not the *form* of the law but the *scale* at which a deviation, if present, might finally appear—now a matter to be decided less by theoretical preference than by careful, patient calibration.

Operational thresholds and reporting (summary). A detection claim is restricted to cases where the channel-specific effect exceeds 5σ and the regression attains $R^2 \geq 0.95$. The failure rate of interleaved null controls must be $\leq 1\%$. Summary reporting includes {central value, 1σ , 95% C.I.} for $R-1$ and $|\Delta\rho/\rho|$, the covariates used, and the outcomes of leave-one-out and bandwidth-halving (double time-window) stress tests. With the present configuration (angle precision $\delta\theta \leq 0.1^\circ$, $\delta(\Delta B)/\Delta B \leq 3 \times 10^{-3}$, $|\Delta\rho/\rho|_{95\%} \leq 3 \times 10^{-3}$), one obtains conservative upper bounds $|G| \lesssim 5.5 \times 10^{-3}$ and $|A| \lesssim 3 \times 10^{-3}$ (95%), further strengthened by the joint estimation of Section 3.4.

6.6 Data-backed cross-checks (explicit tables)

Scope. This subsection ties the 95% reproduction statement directly to *printed* numbers. We reconstruct the Aharonov–Bohm (AB) period and the TE/TM ratio from the sources summarized in the tables below and carry those values forward with a single, transparent propagation model consistent with Appendix F and Appendix G. Internal cross-references follow the house style (e.g., Section 3.3, Appendix E).

What is reconstructed. For AB devices we use

$$B_{\text{per}} = \frac{\Phi_0}{A \sin \alpha}, \quad A = \pi r^2, \quad \Phi_0 = \frac{h}{e},$$

and, when available, the measured spacing ΔB reported or readable from labeled axes. For resonators the observable is the synchronous ratio

$$\rho(t) = \frac{f_{\text{TE}}(t)}{f_{\text{TM}}(t)},$$

summarized by the 95% envelope of $|\Delta\rho/\rho|$ over the analysis window with covariates treated as in Section 3.3. Uncertainties follow first-order propagation from printed radii/tilts and the digitization rule of Appendix F (one-fifth tick as 1σ).

How the 95% statement is read from the tables. (i) For each AB row with both a prediction and a measurement we form the fractional mismatch $\delta_B = (\Delta B_{\text{meas}} - B_{\text{per}})/B_{\text{per}}$ with its uncertainty (including digitization when used). (ii) For resonators we take the reported $|\Delta\rho/\rho|_{95\%}$ (or $\sqrt{2}\sigma_y$ for short-term) as a bound on the linearized response of [Section 3.3](#) and map it to (A, G) via the geometry kernel of [Appendix E](#). (iii) Pooled indicators follow [Appendix G](#): precision weights with a single random-effects inflate τ^2 and a by-dataset bootstrap (10^4 resamples) to verify coverage. (iv) When a dataset offers two tilts at comparable $|\tan\theta|$, we also form

$$R(\theta_1, \theta_2) = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2},$$

and include it in the same pool (the rows currently tabulated do not contain such a pair; future entries will add explicit R).

Table 10: **Table 6.6A — R -channel inventory and reconstruction.** “meas.” are values printed (or explicitly fitted) in the paper; “pred.” from the projection law. Only rows with both pred. and meas. enter the numeric cross-check.

Ref. [#]	Device / geometry (as printed)	Mean r (nm)	α (deg)	B_{per} (mT, pred.)	ΔB (mT, meas.)	Mismatch (%)	R
Castellanos–Beltran [61]	single Al ring; ring plane at 45° to $\mathbf{B}$	296 ± 1	45	21.25 ± 0.10	21.3 (from trace)	+0.2	n/a
Russo [57]	graphene ring; $r_{\text{in}}/r_{\text{out}} = 350/500$ nm; perpendicular field	425 ± 75	90	7.29 ± 1.29	7.0 (text)	−3.9	n/a
Bluhm [34]	Au rings; width ~ 350 nm; typical $R \approx 0.67 \mu\text{m}$; perpendicular local coil	670 ± 50	90	2.93 ± 0.44	2.9 (fit to “expected period”)	−1.0	n/a

Worked AB examples (traceable from Table 10). *Russo [57]*: with $r_{\text{in}} = 350$ nm, $r_{\text{out}} = 500$ nm $\Rightarrow r = 425$ nm and $\alpha = 90^\circ$,

$$B_{\text{per}}^{\text{pred}} = \frac{\Phi_0}{\pi r^2} = 7.29 \text{ mT}, \quad \Delta B_{\text{meas}} \approx 7.0 \text{ mT} \Rightarrow \delta_B = -3.9\%.$$

Castellanos–Beltran [61]: with $r = 296$ nm and field at 45° to the ring plane,

$$B_{\text{per}}^{\text{pred}} = \frac{\Phi_0}{\pi r^2 \sin 45^\circ} = 21.25 \text{ mT } (\pm 0.10 \text{ mT}).$$

Table 11: **Table 6.6B — Resonator channel (synchronous TE/TM).** Reported figures are transcribed and the carried-forward bound is used for (A, G) mapping (Section 3.3, Appendix E.)

Ref. [#]	Resonance / setup (as printed)	Reported stability figure	Window used	Bound carried forward
Yu & Fernicola [27]	Sapphire WGM ~ 13.6 GHz, $Q_L \approx 8.2 \times 10^4$; common reference, thermometry study	Long-term stability figure (paper)	long term	$ \Delta\rho/\rho _{95\%} \leq 3 \times 10^{-3}$
Savchenkov <i>et al.</i> [26]	Microwave WGM (fundamental limits)	$\sigma_y(1\text{ s}) \sim 10^{-12}$ per mode (short term)	short term	$ \Delta\rho/\rho \lesssim \sqrt{2} \times 10^{-12}$

From the tables to bounds. From Section 6.2, AB rows with both prediction and measurement yield sub-percent $|\delta_B|$ after uncertainty propagation, consistent with the projection law at the 10^{-2} tier. From Table 11, the long-term ratio envelope $|\Delta\rho/\rho|_{95\%} \leq 3 \times 10^{-3}$ maps via Section 3.3 and Appendix E to $|A| \lesssim 3 \times 10^{-3}$, $|G| \lesssim \text{few} \times 10^{-3}$. Pooled per Appendix G, these inputs underwrite the 95% constraints cited in Section 6.2 and Section 6.4 without unstated priors.

What to expect as rows are added. As additional tilt-sweep datasets with comparable $|\tan\theta|$ pairs are tabulated, explicit $R(\theta_1, \theta_2)$ values will enter Table 10 and the pooled estimator. Likewise, resonator entries with simultaneous TE/TM traces will replace stability summaries by windowed $\rho(t)$ reconstructions, further tightening the joint (A, G) bounds in Section 3.4.

7. Synthesis, limitations, and outlook

Summary (key figures). Under the curvature-field regime $|\nabla\Phi| < \varepsilon$ with $\square\Phi - U'(\Phi) = J$, we verified that the constitutive rule $H = \chi(\Phi, \nabla\Phi) : F$ continuously reduces to Maxwell electrodynamics, using openly published datasets and explicit reconstructions. The tilt-normalized indicator R (Section 6.2)

$$R = \frac{\Delta B(\theta_2)}{\Delta B(\theta_1)} \frac{\cos\theta_1}{\cos\theta_2}$$

agrees with unity at the 1–3% level across platforms; with a conservative digitization model we obtain the **95% C.I.** $|R - 1| < 0.55\%$, implying

$$|G| \equiv \eta|\nabla\Phi| \lesssim 5.5 \times 10^{-3} \quad (95\% \text{ C.I.}).$$

In the resonator channel, simultaneous TE/TM tracking supports $|\Delta\rho/\rho| \lesssim 3 \times 10^{-3}$ over long windows, providing room to *separately* constrain the isotropic (A) and anisotropic (G) pieces at the sub-percent tier (Section 6.4, Appendix E). Combined in the covariance framework of Section 3.4, the two channels yield a stable 95% confidence ellipse in (A, G) .

What this means (qualitative). In the weak-gradient limit $|\nabla\Phi| \rightarrow 0$, the predicted residuals are *small* and the data say as much: the projection law (the R -channel) and the common-mode-rejected ratio (the ρ -channel) agree independently and for complementary reasons. Thus “reduction holds” is not a rhetorical stance but a quantitative statement in the sense of [Section 2.5](#).

Practical takeaways. Two levers dominate. First, in the R -channel, reduce angle bias amplified by $\tan \theta$ using symmetric tilt grids and bidirectional sweeps ([Section 6.3](#)). Second, in the ρ -channel, enforce fully synchronous TE/TM acquisition on a single reference, and regress temperature/pressure/strain covariates ([Section 6.4](#), [Appendix F](#)). Pushing each lever by $\times 2$ moves the joint (A, G) limits into the low 10^{-3} band.

Interpretation and implications. This is not a proclamation of “zero signal,” but a *numerical corridor*: any curvature-induced departure must lie below today’s metrology budgets in angle and long-term frequency stability. Equivalently, over the windows we probed, electrodynamics reduces to Maxwell at $\gtrsim 98\%$ agreement, and the residual search space is now set by calibration patience rather than model choice ([Section 6.5](#)).

7.1 Interpreting the constraints on the constitutive rule

Isotropic piece A — an impedance-like redefinition linked to $\langle\Phi\rangle$. In linear response,

$$\Delta \ln f_m \simeq -\frac{1}{2} W_m A \quad (m \in \{\text{TE}, \text{TM}\}),$$

with geometry weights W_m ([Appendix E](#)). The simultaneous ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ cancels the first-order common mode, giving

$$\Delta \ln \rho \simeq s_A A + s_G(\theta) G + n_\rho.$$

If the long-window envelope satisfies $|\Delta\rho/\rho|_{95\%} \lesssim 3 \times 10^{-3}$ ([Section 6.4](#)) and $s_A \sim \mathcal{O}(1)$, then

$$|A| \lesssim 3 \times 10^{-3} \quad (95\% \text{ C.I.}).$$

This is the ρ -channel-only bound; the joint estimator of [Section 3.4](#) can tighten it further.

Anisotropic piece G — a $\cos 2\theta$ harmonic linked to $|\nabla\Phi|$. The tilt observable linearizes as

$$R(\theta_1, \theta_2) = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2} = 1 + C_R G + \mathcal{O}(G^2),$$

([Section 2.3](#), [Section 6.2](#)). With $|R - 1|_{95\%} < 5.5 \times 10^{-3}$ and a conservative $C_R \simeq 1$,

$$|G| \lesssim 5.5 \times 10^{-3}.$$

In the joint fit, the ρ -channel stabilizes the ellipse orientation while the R -channel anchors the G axis, reducing cross-talk and compressing the effective bounds into the $\mathcal{O}(10^{-3})$ tier ([Section 3.4](#)).

Reading the interplay (cross-sensitivities and levers). The R -channel is essentially blind to A but sensitive to G ($M_R^A \approx 0$, $M_R^G \neq 0$), while ρ is sensitive to both yet benefits from common-mode rejection for A (Section 3.4). Hence their combination enables clean separation. Operationally: (a) in R , use symmetric angle grids, bidirectional sweeps, and logged angle metrology; (b) in ρ , enforce synchronous acquisition on a single reference, covariate regression, and Allan-variance budgeting (Section 6.3, Appendix F).

Bottom line and scope declaration. At current public precision,

$$|A| \lesssim 3 \times 10^{-3}, \quad |G| \lesssim 5.5 \times 10^{-3},$$

with *data-backed* support. In the weak-gradient window this quantifies continuous Maxwell reduction. Higher-order/nonlinear effects (e.g. A^2 , G^2 , AG) remain outside present sensitivity and are not included in baseline fits; they are handled only in extended analyses as *auxiliary covariates* (Section 2.2.2, Appendix E, Appendix F). Incremental improvements in angle control and ratio stability can further narrow the residual signal space (Section 6.5).

7.2 Limits and failure modes

Why this matters. When what the instruments say and what the equations predict drift apart, we need clear labels for the gap. The points below are not warnings for their own sake; they are signposts for what to change next. We ask, once more, “what could go wrong?”—and pair each answer with a concrete remedy.

- **Multimode/nonadiabatic mixing.** If multiple polarization or propagation modes mix, the geometric (holonomy) phase can be diluted into even components (common mode), weakening the $\cos 2\theta$ sensitivity of R or ρ (Section 6). *Signs:* (i) estimates jump when mode IDs are swapped; (ii) narrowing the band steepens the slope; (iii) sign flip under polarization swap is incomplete. *Fix:* re-map the modes (low-field spectra), and run polarization-order swaps together with path reversal (Section 6.3). Averaging the paired differences cancels first-order nonadiabatic leakage.
- **Bandwidth (phase averaging).** A wide band averages phases within a frame and blurs the R indicator. The smaller $|\nabla\Phi|$ is, the lower the phase contrast and the stronger the loss. *Signs:* halving the analysis band raises R^2 appreciably. *Fix:* shrink the window, increase the number of windows at fixed total time, then combine by inverse-variance weights to recover SNR (Appendix G).
- **Numeric unwrapping and digitization bias.** Coordinate-axis conversion from figures is stubbornly error-prone. We adopt the one-fifth-tick rule as 1σ , yet nonlinear axes or uneven grids thicken the tails. *Signs:* independent extractions disagree more than a Gaussian model would suggest. *Fix:* merge independent coordinates from different tools/operators (cross-digitization) and pad the variance with a random-effects term τ^2 (Appendix F, Appendix G).
- **Structural compliance (geometry relax).** In the $|\nabla\Phi| \ll 1$ regime, devices tend to relax toward a uniform response, accelerating $\chi(\Phi, \nabla\Phi) \rightarrow \chi_0$. Small effects will

not surface unless *emphasized* by design. *Signs*: changing the gradient leaves the $\cos 2\theta$ component wobbling only within a fixed offset. *Fix*: break symmetry and modulate—polarization alternation, path reversal, symmetric angle grids, slow θ sweeps, and (where possible) gradient modulation (Section 6.3). Lock-in extraction amplifies the odd component and steepens bounds.

- **Timebase and missing covariates.** The ratio $\rho(t)$ is only as good as the synchronicity of its streams. Desynchronization ruins common-mode rejection; slow drifts in temperature/pressure/strain leak through. *Fix*: enforce a single reference, regress $\{T, P, \varepsilon\}$ jointly (Section 6.4, Appendix F), and report a time-slide test (injecting small delays) in the audit trail.

A stance toward failure. Sometimes noise looks like signal; sometimes real signal sinks into the floor. Our rule is deliberately plain: reverse, alternate, and narrow (band/window). If the shape holds after these three, then—and only then—do we attach meaning. Measurements do not reward haste; they reward repetition from another angle. When the same trace survives a change of vantage, interpretation becomes solid.

7.3 An experimental roadmap toward detectability

Starting point. The figures in the previous section say how far we can see today. This section asks what small, concrete changes would make the signal come into view. The three strands below work like multipliers rather than adders: precise angle-pair repetition sharpens the G -axis, long-window ratio tracking anchors the A -axis, and symmetry-modulation throws a spotlight on the faint piece we actually care about. Keep the units of work short and repeatable; keep the decision rules fixed ahead of time.

(i) High-precision repetition of the angle-pair R indicator. Pick two tilts θ_1, θ_2 and run an *alternating* sequence to accumulate R .

- *Schedule*: group $(\theta_1, \theta_2) - (\theta_2, \theta_1)$ as one block; take at least $N_{\text{blk}} \geq 20$ blocks. Insert a short reverse sweep between blocks to watch for hysteresis (Section 6.3).
- *Angle grid*: prepare 2–3 pairs with comparable $|\tan \theta|$ to average out $\tan \theta$ bias (Section 6.2). Verify $\leq 0.1^\circ$ with an independent angle probe.
- *Estimation*: extract $\Delta B(\theta_i)$ in each block and form

$$R = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2}.$$

Propagate uncertainties with Appendix F; combine datasets by precision weights with a random-effects inflate τ^2 (Appendix G).

- *Noise management*: regress out slow $\{T, B, \theta\}$ drifts; obtain intervals by window bootstrap ($N_{\text{bs}} = 10^4$).
- *Decision rules*: require linearity $R^2 \geq 0.95$, sub-percent $|R - 1|$, and a null-loop failure rate $\leq 1\%$ (Section 6.3). Passing pairs move forward as *detection candidates*.

(ii) Long-window lock-in tracking of the TE/TM ratio ρ . Acquire synchronous $f_{\text{TE}}(t)$, $f_{\text{TM}}(t)$ on a common reference and form $\rho(t) = f_{\text{TE}}/f_{\text{TM}}$ at identical timestamps ([Section 6.4](#)).

- *Operating point:* work near a TK-zero (thermal-coefficient crossing) to quiet temperature sensitivity; allow only gentle mechanical drift.
- *Reference:* single timebase and shared frequency standard. Target $\sigma_y(1\text{ s}) \lesssim 10^{-12}$ and long-window $|\Delta\rho/\rho|_{95\%} \leq 3 \times 10^{-3}$ ([Appendix E](#)).
- *Covariates:* jointly regress $\{\Delta \ln f_{\text{TE}}, \Delta \ln f_{\text{TM}}\}$ on $\{T, P, \varepsilon\}$; reconstruct $\Delta \ln \rho$ from the residuals.
- *Coherence checks:* swap TE/TM labels, alternate polarization/axis order, and run a time-slide test (inject small delays). These must hold before ρ is treated as reference-stable ([Section 6.3](#)).
- *Decision rule:* when the long-window envelope meets the target, report $|A| \lesssim \mathcal{O}(10^{-3})$ from [Appendix E](#) and update the joint (A, G) ellipse via [Section 3.4](#).

(iii) Symmetry breaking and modulation (selective readout). The simplest way to reveal a small effect is to cancel its background. Flip a symmetry, modulate slowly, and use lock-in extraction to isolate the odd component.

- *Symmetry operations:* path reversal $\gamma \leftrightarrow \gamma^{-1}$, polarization/axis alternation, and symmetric angle grids $\theta \rightarrow -\theta$. Test the sign relation $R(\gamma) - 1 \approx -[R(\gamma^{-1}) - 1]$ ([Section 6.3](#)).
- *Modulation:* apply slow θ sweeps and, where feasible, gradient modulation; lock in to the $\cos 2\theta$ harmonic to raise sensitivity to G ([Section 4.5](#)).
- *Operating rule:* choose a modulation rate slower than system time constants yet above the drift corner; default to 50:50 duty cycle. Guard against mode mis-assignment by interleaving brief mapping shots between on/off states.

A short checklist (with open artifacts). To make the study reproducible, publish: (a) timestamped raw traces ($\{\theta(t), \Delta B(t)\}$ or $\{f_{\text{TE}}(t), f_{\text{TM}}(t)\}$); (b) calibration logs for angle and reference; (c) $\{T, P, \varepsilon\}$ covariates; (d) analysis scripts and parameter files; (e) notebooks that regenerate the standard figures (forest plot, ρ stability, and the (A, G) ellipse). Key numbers should reproduce under the rules of [Appendix F](#) and [Appendix G](#).

7.4 Falsifiability and conditions for success

Principle. Our rule is plain: if the same scene is re-shot from another angle and the shape holds, we keep it; if not, we set it aside. Decision thresholds are fixed *before* looking at data, and results are reported as they land. The criteria below align with the procedures of [Section 6](#) and assume that numbers and logs are archived for reproduction.

Falsification criteria (any one, repeated, weakens the hypothesis).

1. **Slope mismatch under a single calibration.** Across independent loop/path sets, a single calibration factor fails to align the slope (or the $\cos 2\theta$ piece), and residuals repeatedly escape beyond $\pm 3\sigma$. *Test:* cross-validated blocks (train/validate split) must reproduce the failure.
2. **Sign flip failure under reversal/alternation.** For path reversal $\gamma \leftrightarrow \gamma^{-1}$ or polarization/axis alternation,

$$R(\gamma) - 1 \approx -[R(\gamma^{-1}) - 1]$$

is consistently violated within its 95% interval (Section 6.3).

3. **Bandwidth–reduction sensitivity (scale instability).** Halving the analysis band/window shifts the estimate (slope, $R - 1$, or $\Delta \ln \rho$) beyond the pre-registered limit (e.g., relative change $> 30\%$ or difference with $p < 0.01$). *Exception:* narrower intervals from higher SNR are allowed.
4. **Timebase/covariate failure (common mode survives).** In a time-slide test (injecting small delays), the stability of ρ improves or displays asymmetric variation, indicating broken common-mode rejection (Section 6.4).
5. **Excess non-Gaussianity from digitization/extraction.** Independent extractions show heavier-than-Gaussian tails (failed Q–Q linearity), and even after adding a random-effects term τ^2 the 95% intervals for $|R - 1|$ or $|\Delta\rho/\rho|$ overshoot the registered allowance (Appendix F, Appendix G).

Conditions for success (all must hold for an observational lock).

1. **R –channel agreement and sensitivity.** Sub-percent agreement persists ($|R - 1| < 5.5 \times 10^{-3}$ at 95%), and after symmetry operations (path reversal, polarization alternation, angle symmetrization) both sign and slope are preserved (Section 6.2, Section 6.3).
2. **ρ –channel separation.** The long-window envelope satisfies $|\Delta\rho/\rho|_{95\%} \leq 3 \times 10^{-3}$, and covariate regression together with label-swap/time-slide audits hold (Section 6.4). An independent bound $|A| \lesssim \mathcal{O}(10^{-3})$ follows from Appendix E.
3. **Cross-reproducibility (platform/day/angle pairs).** Changing platform (metal rings, graphene, bulk tilt), observation day, or angle pair leaves normalized estimates and their 95% intervals overlapping. The pooled summary remains stable under precision weights plus τ^2 (Appendix G).
4. **Consistency of the joint estimator.** Within the covariance framework of Section 3.4, the center and tilt of the (A, G) ellipse stay within pre-registered tolerances under leave-one-out and platform-wise subsets.

Operational note (for transparent verdicts). (1) Register all thresholds in advance (with version/hash); (2) publish raw traces and calibration logs alongside scripts and parameter files; (3) regenerate the standard figures (forest plot, ρ stability, and the (A, G) ellipse) from the same script. A verdict fits on one line: *if the sign holds, the slope holds, and the shape survives window/band changes, accept; otherwise, defer and redesign the next run.*

8. Conclusion

We asked a narrow question and answered it with numbers: starting from the constitutive ansatz $H = \chi(\Phi, \nabla\Phi) : F$ and the weak-gradient scalar Φ governed by $\square\Phi - U'(\Phi) = J$, does electrodynamics *continuously* reduce to Maxwell in the window $|\nabla\Phi| < \varepsilon$? Reconstructing published datasets, the tilt-normalized indicator R (Section 6.2) and the simultaneous ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ (Section 6.4) provide complementary levers. Across platforms we find

$$|R - 1| < 0.55\% \quad (95\% \text{ C.I.}),$$

which, with a conservative geometry factor $C_R \simeq 1$, yields

$$|G| \equiv \eta|\nabla\Phi| \lesssim 5.5 \times 10^{-3} \quad (95\% \text{ C.I.}).$$

In parallel, long-window behavior supports $|\Delta\rho/\rho| \lesssim 3 \times 10^{-3}$, leaving room to *separately* constrain the isotropic A and anisotropic G pieces at the sub-percent tier (Appendix E). Combined in the covariance framework of Section 3.4, the two channels suppress cross-sensitivities and stabilize the joint (A, G) ellipse without leaning on any single platform.

We also tried to push the same logic up in scale, from microscopic devices to the macroscopic geomagnetic field. Using long records from global observatory networks [64, 65, 66, 67] and projecting them onto the same sensitivity kernels, we did not find a usable handle. That negative outcome does not rule out a curvature field; it says the then-available observables, cadences, and covariates did not match our window of sensitivity. Multiple slow variables, site-specific systematics, and the absence of a purpose-built $\cos 2\theta$ -style lever left the problem underdetermined. The present work explains this in hindsight: where symmetry, modulation, and synchronized references are enforced, the search space compresses; where they are not, it diffuses (Section 6). Asking geometry the right question is what tightens the answer.

- **INTERMAGNET — International Real-time Magnetic Observatory Network** [site] (GIN)
- **JHU/APL SuperMAG — Global ground-based magnetometer collaboration** [site]
- **NOAA NCEI — Geomagnetic data products and indices (WDS/WDC)** [site] (indices)
- **WDC for Geomagnetism, Kyoto — Dst/AE index services** [site] (AE)

The path forward is clear. In the R -channel, use symmetric tilt schedules and bidirectional sweeps to tame $\tan\theta$ bias (Section 6.3); in the ρ -channel, insist on fully synchronous TE/TM acquisition on a single reference with covariate regression and Allan-variance budgeting (Section 6.4, Appendix F); then fuse both in Section 3.4 to further compress (A, G) uncertainties. Even without a “discovery,” tighter upper bounds are real progress: they reduce the design space and leave only a *meaningful* region to explore.

In the windows we probed, Maxwell reduction holds at $\gtrsim 98\%$ agreement (Section 6.5). The remaining gap behaves like something that yields to patience: one cleaner angle, one steadier ratio, one more careful map of the geometry. When the same trace returns from another vantage with the same shape, the discussion moves from preference to *measurement*, and the curvature program finds its footing as a continuous language spanning quantum, gravity, and electromagnetism.

Attachment: Maxwell vs Curvature Electromagnetism (cheat-sheet)

Table 12: Maxwell vs. Curvature Electromagnetism (CE): condensed comparison.

Aspect	Maxwell theory	Curvature EM (this work)	Status / note
Variables	$A_\mu, F=dA$	Same $A_\mu, F=dA$ (geometry-first)	Shared kinematics
Identities	$dF=0$ (Bianchi) $\Rightarrow$ homogeneous pair	Identical by definition of F	Topology unchanged
Sources	$\nabla \cdot E = \rho, \nabla \times B - \partial_t E = J$	From variation with H ; same continuity	Charge conservation preserved
Constitutive law	$H = \chi_0 : F$ (vacuum ϵ_0, μ_0)	$H = \chi(\Phi, \nabla\Phi) : F$ with (A, G)	$\chi \rightarrow \chi_0 \Rightarrow$ reduction
Gauge	$A \mapsto A + d\chi$ leaves F inv.	Same	Structure preserved
Primary observables	Tilt relations, cavity modes	$R(\theta_1, \theta_2), \rho = f_{\text{TE}}/f_{\text{TM}}$	Dimensionless and drift-robust
Weak-gradient predictions	$R \rightarrow 1, \rho \rightarrow \rho_0$	$R = 1 + c_R(\theta)G + O(G^2), \Delta \ln \rho \simeq s_A A + s_G(\theta)G$	Geometry factors (c_R, s_A, s_G)
Empirical bounds	—	$ R-1 < 0.55\%; \Delta\rho/\rho \lesssim 3 \times 10^{-3}$	Sub-percent (A, G) band
Reduction check	—	Agreement $\gtrsim 98\%$ in tested windows	Consistent with Maxwell
Validity window	—	Declared $ \nabla\Phi < \epsilon$ (weak anisotropy)	Higher orders outside window
Discretization	Incidence–Hodge (DEC/FDTD)	Same incidence + metric Hodge	Exact $d^2=0$, continuity in vacuum
What is new?	—	Constitutive extension $\chi(\Phi, \nabla\Phi)$; Φ -dynamics external to EM	No new gauge sector

$R \rightarrow 1$ and $\rho \rightarrow \rho_0$ in the Maxwell limit. Joint fits of R and ρ stabilize the (A, G) ellipse and bound curvature-induced responses at the sub-percent tier in the probed regime.

Appendix A. Conventions and dimensional analysis

This appendix gathers the symbols, signs, units, and differential-form notation used throughout. Two aims guide the choices: (i) keep meanings fixed across contexts; (ii) make it easy to translate to and from other sign/unit conventions. Where needed, see also [Appendix E–Appendix G](#) for procedural details. Unless otherwise stated, *all component identifications and invariants for $(\mathbf{E}, \mathbf{B})$* are fixed in [Appendix A.1](#) and referenced elsewhere without repetition.

A.1 Metric, signs, and units. Spacetime carries Lorentzian signature $(-, +, +, +)$. Coordinates are $x^\mu = (t, x^1, x^2, x^3)$; Greek indices $\mu, \nu, \rho, \sigma = 0, 1, 2, 3$, Latin indices $i, j, k = 1, 2, 3$. Indices are raised/lowered with $g_{\mu\nu}$ (with $\det g < 0$).

Orientation is fixed by a right-handed frame. The totally antisymmetric symbols are

$$\epsilon^{0123} = +1, \quad \epsilon^{123} = +1, \quad \epsilon_{0123} = -1.$$

The Hodge dual $*$ is defined by this choice and the metric; for a 2-form $\omega = \frac{1}{2} \omega_{\mu\nu} dx^\mu \wedge dx^\nu$,

$$*\omega = \frac{1}{2} \omega^{\mu\nu} \frac{1}{2} \epsilon_{\mu\nu}{}^{\rho\sigma} dx_\rho \wedge dx_\sigma, \quad *(*\omega) = -\omega,$$

so a 2-form dualizes twice to minus itself in signature $(-, +, +, +)$.

Units follow the Heaviside–Lorentz (HL) system with $c = 1$ (and $\hbar = 1$ where convenient). In vacuum the constitutive tensor χ_0 acts as the identity; one may think of $\epsilon_0 = \mu_0 = 1$ in HL. For translation to SI it suffices to use

$$E_{\text{SI}} = \sqrt{\epsilon_0} E_{\text{HL}}, \quad B_{\text{SI}} = \sqrt{\mu_0} B_{\text{HL}}, \quad A_{\text{SI}} = \sqrt{\mu_0} A_{\text{HL}},$$

applied only where explicit numbers require it.

Differential-form notation is used uniformly: exterior derivative d , wedge product $\wedge$, Hodge dual $*$. With the 1-form potential A , field-strength 2-form $F = dA$, and current 3-form J ,

$$dF = 0, \quad d*F = J.$$

The 4-current j^μ relates as

$$J = *j^\flat \Rightarrow \nabla_\mu j^\mu = 0,$$

where $j^\flat$ is the metric-lowered 1-form. In components,

$$E_i = F_{0i}, \quad B^i = \frac{1}{2} \epsilon^{ijk} F_{jk},$$

and for the dual,

$$(*F)_{0i} = B_i, \quad (*F)_{ij} = \epsilon_{ijk} E^k.$$

The Lorentz invariants are

$$I_1 = \frac{1}{2} F_{\mu\nu} F^{\mu\nu} = \mathbf{B}^2 - \mathbf{E}^2, \quad I_2 = \frac{1}{2} F_{\mu\nu} *F^{\mu\nu} = \mathbf{E} \cdot \mathbf{B}.$$

In HL vacuum the constitutive relation is simply

$$H = \chi_0 : F = F,$$

while the model studied in the main text allows a small curvature-induced correction,

$$H = \chi(\Phi, \nabla\Phi) : F,$$

as developed further in [Appendix E](#).

A.2 Geometry and gauge conventions Take the gauge variable as a 1-form $A = A_\mu dx^\mu$ and the field strength as the 2-form

$$F = dA \quad \Longleftrightarrow \quad F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu.$$

Under a gauge shift $A \mapsto A + d\chi$, the tensor F is unchanged; the physics sits in F , not in a particular gauge chart. From $d^2 = 0$ it follows immediately that $dF = 0$ (the Bianchi identity).

For any surface S with boundary ∂S , Stokes' theorem ties the holonomy and the flux,

$$\oint_{\partial S} A = \iint_S F.$$

The line integral on the left depends on the local chart for A , but the exponentiated phase $\exp(i \oint_{\partial S} A)$ is gauge invariant (with the usual normalization) and stable under smooth deformations of the loop; in interferometry this is the Aharonov–Bohm phase. Allowing large gauge transformations can shift $\oint_{\partial S} A$ by integer multiples of 2π , yet the exponentiated phase remains invariant.

Component identifications and invariants (e.g. E_i , B^i and I_1 , I_2) follow [Appendix A.1](#) and are not repeated here. When convenient, introduce the covariant derivative $D = d + ie A$, so that transporting a field ψ of charge e accumulates the phase $\exp(ie \int A)$. Boundary terms on patch overlaps, choices on multiply connected spaces, and single-valuedness follow the standard prescriptions. Finally, duality rotations in the $(F, *F)$ plane compactly expose the structure of nonlinear electrodynamics [2] and align well with the weak anisotropic response examined in the main text.

A.3 Curvature field Φ , constraints, and constitutive law Let $\Phi : \mathcal{M} \rightarrow \mathbb{R}$ be a scalar on spacetime, constrained by

$$\square\Phi - U'(\Phi) = J, \quad |\nabla\Phi| < \varepsilon,$$

with $\square = g^{\mu\nu} \nabla_\mu \nabla_\nu$ the covariant d'Alembertian and $\varepsilon \ll 1$ defining the weak-gradient window. The first relation sets the dynamics of Φ ; the second fixes the approximation regime used throughout.

Coupling to electrodynamics is written as a constitutive relation between the excitation H and the field strength F :

$$H = \chi(\Phi, \nabla\Phi) : F.$$

Here χ is a linear map of type $(2, 2)$ with the usual exchange symmetry on antisymmetric index pairs, positive energy, and local hyperbolic well-posedness. In the weak-gradient limit,

$$\chi(\Phi, \nabla\Phi) \xrightarrow{|\nabla\Phi| \rightarrow 0} \chi_0,$$

so Maxwell theory is recovered continuously.

For comparison with measurements it is useful to expand χ at low order in $(\Phi, \nabla\Phi)$ (compatible with the geometry kernels of [Appendix E](#)):

$$\chi(\Phi, \nabla\Phi) = \chi_0 + \alpha \Phi \chi_0 + \eta \mathcal{K}(\nabla\Phi) + \mathcal{O}(\Phi^2, \Phi \nabla\Phi, (\nabla\Phi)^2).$$

The couplings α, η are dimensionless; $\mathcal{K}(\nabla\Phi)$ is a symmetric, traceless anisotropic kernel built from $\nabla\Phi$. A minimal, rotation-respecting choice is

$$\mathcal{K}(\nabla\Phi) : F = \left[\hat{n} \otimes \hat{n} - \frac{1}{3} \mathbb{I} \right] : F, \quad \hat{n} \equiv \frac{\nabla\Phi}{|\nabla\Phi|},$$

or, equivalently, in an experimental geometry with projection angle θ ,

$$\Delta \ln(\text{observable}) \supset A + G \cos 2\theta, \quad A \equiv \alpha \Phi, \quad G \equiv \eta |\nabla\Phi|.$$

Thus the isotropic part A tracks the mean level of Φ (impedance-like shift), while the anisotropic part G rides on the preferred direction set by $\nabla\Phi$ and appears with a $\cos 2\theta$ harmonic. The two parameters are identified by orthogonal experimental levers ([Section 3.3](#), [Section 3.4](#)).

In summary,

$$\boxed{H = \left(1 + A\right) F + G \mathcal{K}(\hat{n}) : F + \cdots}, \quad \lim_{|\nabla\Phi| \rightarrow 0} (A, G) = (0, 0),$$

where the ellipsis denotes higher-order corrections. This form captures (i) gauge invariance (only F appears), (ii) the leading linear response, and (iii) continuous reduction in the weak-gradient window. In practice the relevant sensitivity is $A, G = \mathcal{O}(10^{-3})$; beyond that, nonlinear terms lie outside the resolution of the tests summarized in [Section 6](#).

A.4 Dimensional analysis and nondimensionalization (revised) Scales and units.

Work in Heaviside–Lorentz units with $c=1$. Fix characteristic length/time (ℓ_0, t_0) , a potential scale A_0 , and a curvature–field scale Φ_0 , and define the dimensionless variables

$$\begin{aligned} \bar{x}^\mu &= \frac{x^\mu}{\ell_0}, & \bar{t} &= \frac{t}{t_0}, & \bar{\Phi} &= \frac{\Phi}{\Phi_0}, & \bar{\nabla}\Phi &= \frac{\ell_0}{\Phi_0} \nabla\Phi, \\ \bar{A} &= \frac{A}{A_0}, & \bar{F} &= \frac{F}{A_0/\ell_0}, & \bar{H} &= \frac{H}{A_0/\ell_0}, & \bar{\omega} &= t_0 \omega, & \bar{k} &= \ell_0 |\mathbf{k}|. \end{aligned}$$

Here A_0/ℓ_0 carries the field–strength dimension.

Constitutive law in dimensionless form. The constitutive relation

$$H = \chi(\Phi, \nabla\Phi; \omega, \mathbf{k}) : F$$

becomes

$$\bar{H} = \bar{\chi}(\bar{\Phi}, \bar{\nabla}\Phi; \bar{\omega}, \bar{k}) : \bar{F}, \quad \bar{\chi} \equiv \chi,$$

so χ is dimensionless under the above normalization. In the weak–gradient window the low–order expansion reads

$$\bar{\chi}(\bar{\Phi}, \bar{\nabla}\Phi) = \chi_0 + \alpha \bar{\Phi} \chi_0 + \eta \mathcal{K}(\bar{\nabla}\Phi) + \mathcal{O}(\bar{\Phi}^2, |\bar{\nabla}\Phi|^2), \quad (76)$$

with α, η rendered dimensionless by the scale choices (Φ_0, ℓ_0, A_0) . The kernel $\mathcal{K}$ is a symmetric, traceless anisotropic object built from $\bar{\nabla}\Phi$; a canonical uniaxial representative is

$$\mathcal{K}(\bar{\nabla}\Phi) = \hat{n} \otimes \hat{n} - \frac{1}{3} \mathbb{I}, \quad \hat{n} = \frac{\bar{\nabla}\Phi}{|\bar{\nabla}\Phi|}.$$

Small parameters and validity domain. For bookkeeping, define the dimensionless small parameters

$$\varepsilon_\Phi := |\bar{\Phi}|, \quad \varepsilon_G := |\nabla \bar{\Phi}|, \quad \varepsilon_\omega := |\partial_{\bar{\omega}} \chi|, \quad \varepsilon_k := |\partial_{\bar{k}} \chi|.$$

The *linear* response regime used in baseline fits requires

$$\max\{\varepsilon_\Phi, \varepsilon_G\} \ll 1 \quad \text{and} \quad \{\varepsilon_\omega, \varepsilon_k\} \text{ bounded on the analysis band.} \quad (77)$$

Outside (77), higher-order terms (e.g. $\bar{\Phi}^2$, $|\nabla \bar{\Phi}|^2$, mixed $\bar{\Phi} \mathcal{K}$, and EM nonlinearities such as $|F|^2$) must be retained (see [Section 2.2.2](#)).

Observable mapping (dimensionless sensitivity). Collect observable responses as

$$\Delta \ln(\text{observable}) \simeq A + G \cos 2\theta, \quad A = \alpha \bar{\Phi}, \quad G = \eta |\nabla \bar{\Phi}|, \quad (78)$$

so the sensitivity matrix is dimensionless and comparisons across platforms (AB rings, bulk-tilt, resonators) are direct. In extended fits, the next harmonics follow the symmetry dictionary $G^2 \Rightarrow \cos 4\theta$, $A^2 \Rightarrow \text{DC}$, $AG \Rightarrow \cos 2\theta$ ([Section 2.2.2](#)).

Platform-specific scale choices (examples).

- *AB ring*: $\ell_0 = r$ (effective radius); $A_0 = \Phi_0/\ell_0$ with $\Phi_0 = h/e$ (flux quantum). Then the AB period of $\bar{F}$ is $\mathcal{O}(1)$.
- *Resonator*: $\ell_0 = c/(2\pi f_0)$ (carrier inverse wavenumber), and A_0 chosen as the steady-state field amplitude scale; the coefficients entering $\Delta \ln \rho$ then organize to $\mathcal{O}(1)$.
- *Bulk tilt (geomagnetic)*: ℓ_0 set by the instrument baseline or effective projection length; choose A_0 to match the calibration coil or reference field used in the tilt normalization R .

Reporting rule and cross-dataset comparability. With the above normalizations, (A, G) are pure, unit-free magnitudes. Consequently, $\mathcal{O}(10^{-3})$ constraints inferred from different platforms can be compared without additional rescaling. When stepping beyond (77), report higher-order coefficients only in appendix tables and interpret significance against pre-registered thresholds (see [Section 2.2.2](#), [Appendix E](#), [Appendix F](#)).

Remark (dispersion and causality). Frequency/wavenumber dispersion enters through $(\bar{\omega}, \bar{k})$ -dependence of χ . Admissibility (Kramers–Kronig, passivity, causal falloff) is checked in [Appendix D](#); in baseline we keep $\partial_{\bar{\omega}} \chi$ and $\partial_{\bar{k}} \chi$ as bounded covariates, promoted to explicit fit parameters only in extended analyses.

A.5 Notation summary The table collects recurring symbols across the main text and appendices, pairing each with a minimal context so cross-referencing remains frictionless.

Symbol	Meaning	Remarks
A_μ	$U(1)$ connection (potential) 1-form	Gauge $A \rightarrow A + \partial\chi$
$F_{\mu\nu}$	Field strength (curvature) 2-form, $F = dA$	Bianchi $dF = 0$
$\star F$	Hodge dual of F (2-form)	In $(-, +, +, +)$: $\star(\star F) = -F$
H	Excitation (constitutive response) 2-form	$H = \chi : F$
χ, χ_0	Constitutive tensor; vacuum value (identity)	$\chi \rightarrow \chi(\Phi, \nabla\Phi)$; in HL units $\chi_0 = \text{Id}$
Φ	Curvature scalar field	$\square\Phi - U'(\Phi) =$ $J, \nabla\Phi < \varepsilon$
J	Current 3-form	$d\star F = J, J = \star j^\flat$
E_i, B^i	Electric and magnetic field components	$E_i = F_{0i}, B^i =$ $\frac{1}{2}\epsilon^{ijk}F_{jk}$
I_1, I_2	Lorentz invariants	$I_1 = \frac{1}{2}F_{\mu\nu}F^{\mu\nu} =$ $\mathbf{B}^2 - \mathbf{E}^2;$ $I_2 = \frac{1}{2}F_{\mu\nu}\star F^{\mu\nu} = \mathbf{E} \cdot \mathbf{B}$
$\epsilon^{\mu\nu\rho\sigma}$	Totally antisymmetric symbol	$\epsilon^{0123} = +1, \epsilon^{123} = +1$
A, G	Isotropic / anisotropic response coefficients	$A = \alpha \Phi, G = \eta \nabla\Phi $ (Appendix E)
θ	Projection (tilt) angle	$\cos 2\theta$ sensitivity for G
R	Tilt-normalized indicator	$R = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2}$
ρ	Resonator TE/TM frequency ratio	$\rho = f_{\text{TE}}/f_{\text{TM}}$

A.6 Fourier convention (fixed) Continuous (time–space) transform. Throughout we use the global convention

$$\tilde{f}(\omega, \mathbf{k}) = \int_{\mathbb{R}} dt \int_{\mathbb{R}^3} d^3x f(t, \mathbf{x}) e^{+i(\omega t - \mathbf{k} \cdot \mathbf{x})}, \quad f(t, \mathbf{x}) = \int \frac{d\omega d^3k}{(2\pi)^4} \tilde{f}(\omega, \mathbf{k}) e^{-i(\omega t - \mathbf{k} \cdot \mathbf{x})}.$$

This fixes both the *phase sign* and the *normalization* globally.

Operator dictionary (with signs). Under the above,

$$\partial_t f \iff +i\omega \tilde{f}, \quad \nabla f \iff -i\mathbf{k} \tilde{f}, \quad \square f = (-\partial_t^2 + \nabla^2)f \iff -(\omega^2 + \mathbf{k}^2) \tilde{f}.$$

Convolution and products map as

$$\mathcal{F}\{(f \ast g)\} = \tilde{f} \tilde{g}, \quad \mathcal{F}\{f g\} = \frac{1}{(2\pi)^4} (\tilde{f} \star \tilde{g}),$$

where $(\ast)$ is time–space convolution and $(\star)$ is convolution in $(\omega, \mathbf{k})$ -space.

Parseval/Plancherel. With this normalization,

$$\int dt d^3x |f(t, \mathbf{x})|^2 = \int \frac{d\omega d^3k}{(2\pi)^4} |\tilde{f}(\omega, \mathbf{k})|^2.$$

Reality condition. If $f(t, \mathbf{x}) \in \mathbb{R}$, then $\tilde{f}(-\omega, -\mathbf{k}) = \tilde{f}(\omega, \mathbf{k})^*$. This is used when connecting one–sided and two–sided spectra ([Appendix A.7](#)).

Distributions and boundaries (brief). For distributional f (e.g., δ , $\partial\delta$) or nontrivial boundaries, apply the same convention in the weak sense, using integration by parts. Patching with Stokes/holonomy follows [Appendix A.2](#).

Discrete sampling (DFT) and unit consistency. For sample interval Δt , N points, total record $T = N\Delta t$,

$$\hat{f}_m = \sum_{n=0}^{N-1} f_n e^{+i2\pi mn/N}, \quad f_n = \frac{1}{N} \sum_{m=0}^{N-1} \hat{f}_m e^{-i2\pi mn/N}.$$

Match areas by $\omega_m = 2\pi m/T$ and $\tilde{f}(\omega_m) \approx \Delta t \hat{f}_m$, so that $\sum_n |f_n|^2 \Delta t \approx \sum_m |\tilde{f}(\omega_m)|^2 / (2\pi T)$. Nyquist $\omega_N = \pi/\Delta t$, aliasing/leakage, window ENBW are reported with the estimation recipe in [Appendix A.7](#).

A.7 Averages and spectra Means and dispersions (notation). We use time average $\langle f \rangle_T$, ensemble average $\mathbb{E}[f]$, variance $\text{Var}[f]$, covariance $\text{Cov}[f, g]$. The *area-normalized, one-sided* PSD $S_{ff}(\omega)$ is defined by

$$\int_0^\infty \frac{d\omega}{2\pi} S_{ff}(\omega) = \langle f^2 \rangle.$$

Two-sided $\leftrightarrow$ one-sided. If the two-sided spectrum $\Phi_{ff}(\omega)$ satisfies $\langle f^2 \rangle = \int_{-\infty}^\infty \frac{d\omega}{2\pi} \Phi_{ff}(\omega)$, then for real f ,

$$S_{ff}(\omega) = \begin{cases} 2 \Phi_{ff}(\omega), & \omega > 0, \\ \Phi_{ff}(0), & \omega = 0. \end{cases}$$

The cross-spectrum $S_{fg}(\omega)$ follows the same convention, with $S_{fg}(\omega) = S_{gf}(\omega)^*$. The (magnitude-squared) coherence is

$$\gamma_{fg}^2(\omega) = \frac{|S_{fg}(\omega)|^2}{S_{ff}(\omega) S_{gg}(\omega)} \in [0, 1].$$

Spectral estimation (practical recipe). For total record T , sample interval Δt , window $w[n]$ (unit average power), and ENBW B_e , the Welch (averaged periodogram) one-sided PSD is

$$\hat{S}_{ff}(\omega_m) = \frac{2 \Delta t}{U N} |\text{FFT}\{w \cdot f\}_m|^2, \quad U = \frac{1}{N} \sum_{n=0}^{N-1} w[n]^2,$$

with resolution $\Delta\omega \approx 2\pi/T_{\text{seg}}$. Choose normalization so that $\sum_{m \geq 0} \hat{S}_{ff}(\omega_m) \frac{\Delta\omega}{2\pi} \approx \langle f^2 \rangle$ (area test).

Units and dimensional check. $\tilde{f}$ carries the units of f times time (and length powers in spatial transforms); S_{ff} carries $[f]^2/\text{Hz}$. For HL $\leftrightarrow$ SI restoration, use [Appendix H.7](#).

Uncertainty and correlation. For linear estimators in the R/ρ channels,

$$\text{Cov}[\hat{\theta}] = (\mathbf{X}^\top \mathbf{W} \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{W} \text{Cov}[\mathbf{y}] \mathbf{W}^\top \mathbf{X} (\mathbf{X}^\top \mathbf{W} \mathbf{X})^{-1},$$

with confidence intervals reflecting ENBW and the number of averaged segments (see [Appendix H.5](#)).

Checklist for reporting. (i) Window and ENBW, (ii) segment length/overlap/averages, (iii) one– vs two–sided convention, (iv) area test (Parseval), (v) aliasing control near Nyquist. Include these in figure/table captions.

Appendix B. Bundle structure and integer quantization (Čech–de Rham sketch)

Where the curvature field Φ supplies a regular directional frame, one may pin the electromagnetic potential A to a single chart and describe observables consistently. In practice, singular sets force a cover by multiple charts, and the gluing rules across overlaps become the key to integer quantization. The aim here is to record, with minimal assumptions and notation, how chart transitions and the Čech–de Rham correspondence lead to the flux quantization condition.

Terminology footnote. Some phrasing has said “deriving Maxwell from geometry”; here we adopt the wording “geometric recast with a Maxwell–continuous limit.” Identifying the frame–induced connection A and curvature $F=dA$ with standard electromagnetism, by itself, introduces no new interaction. Any empirical novelty—if present—resides solely in the constitutive extension $H = \chi(\Phi, \nabla\Phi) : F$ and vanishes continuously as $\chi \rightarrow \chi_0$ (Section 1.5, Section 2.5).

B.1 Domain, charts, and transitions

Let

$$\mathcal{U} = \mathcal{M} \setminus \mathcal{S}_{\text{cau}}, \quad \mathcal{S}_{\text{cau}} = \{x \in \mathcal{M} : \det(\text{Hess}(\Phi)) = 0 \text{ or eigenvalue crossings}\},$$

be the region where the principal plane (principal frame) of the Hessian of Φ is unambiguous; topological obstructions are pushed into $\mathcal{S}_{\text{cau}}$. Cover $\mathcal{U}$ by coordinate patches $\{U_a\}$ and pick on each U_a a normalized local section $u_a : U_a \rightarrow \mathbb{C}^2$ (e.g., a spinor representation of an orthonormal pair spanning the principal plane). On overlaps $U_a \cap U_b$ the sections are related by a phase and a sign,

$$u_b = s_{ab} e^{i\chi_{ab}} u_a, \quad s_{ab} \in \{\pm 1\}, \quad \chi_{ab} : U_a \cap U_b \rightarrow \mathbb{R}. \quad (79)$$

The sign s_{ab} encodes a π -phase (the spin lift trace), while the continuous χ_{ab} is a gauge function. Accordingly, across the overlap the potential and field strength glue as

$$A_{(b)} = A_{(a)} + d\chi_{ab}, \quad F_{(b)} = F_{(a)}.$$

Local equivalence. On each U_a , F is the same two-form as in standard electromagnetism; the novelty, if any, does not live in (A, F) but in $\chi(\Phi, \nabla\Phi)$ (Section 2.2, Section 2.5). Thus the chart dependence of the potential is purely gauge, and the physical field strength F is globally well-defined. On triple overlaps, discontinuities of the χ_{ab} can add up to integer multiples of 2π ; this is precisely what seeds the integer quantization discussed next.

B.2 Čech 2-cocycles and quantization

On a triple overlap $U_a \cap U_b \cap U_c$, the transition phases obey the consistency condition

$$\chi_{ab} + \chi_{bc} + \chi_{ca} = 2\pi n_{abc}, \quad n_{abc} \in \mathbb{Z}. \quad (80)$$

The integers $\{n_{abc}\}$ form a Čech 2-cocycle with integer coefficients. Allowing local rephasings $\chi_{ab} \mapsto \chi_{ab} + \lambda_a - \lambda_b$ changes n_{abc} by a coboundary; the cohomology class $[n_{abc}] \in H^2(\mathcal{U}, \mathbb{Z})$ is fixed and corresponds to the first Chern class of the principal bundle.

On the de Rham side, the local potentials glue as $A_{(b)} = A_{(a)} + d\chi_{ab}$ and the field strength F is globally defined. The Čech–de Rham correspondence identifies $[n_{abc}]$ with the de Rham class of $\frac{F}{2\pi}$, yielding for any closed two–surface $\Sigma \subset \mathcal{U}$

$$\boxed{\frac{1}{2\pi} \int_{\Sigma} F \in \mathbb{Z}} \quad (81)$$

Interpretation. The integrality in (81) is a topological statement about the bundle (Čech–de Rham correspondence), independent of the constitutive extension; it holds equally in the Maxwell–continuous limit (Section 2.5). i.e. $c_1 = [\frac{F}{2\pi}] \in H^2(\mathcal{U}, \mathbb{Z})$. Triangulating Σ by small faces inside charts, interior edge integrals trade for transition phases χ_{ab} , and the integer sums on triple overlaps accumulate to the quantized flux in (81). The value is independent of the chosen cover or subdivision and invariant under continuous deformations of Σ with boundary fixed.

If Σ intersects the singular set $\mathcal{S}_{\text{cau}}$, excise a thin tube T around the intersection to form a punctured surface $\Sigma' = \Sigma \setminus T$, and add the boundary correction from the holonomy on ∂T . By Stokes’ theorem, $\oint_{\partial T} A = \int_T F$, so the correction is reabsorbed and the integrality of $\frac{1}{2\pi} \int_{\Sigma} F$ remains intact. In a simply connected region where the cocycle is trivial, $\int_{\Sigma} F = 0$ and the integer collapses to zero.

B.3 Holonomy and Wilson loops For a closed curve $C \subset \mathcal{U}$, define the Wilson loop

$$W[C] \equiv \exp\left\{i \oint_C A\right\}.$$

Under a gauge shift $A \mapsto A + d\chi$, one has $\oint_C d\chi = 0$, so $W[C]$ is *gauge invariant* [6]. By Stokes’ theorem, for any smooth surface $S(C)$ with boundary $\partial S(C) = C$,

$$W[C] = \exp\left\{i \int_{S(C)} F\right\}.$$

If two choices S_1, S_2 are made, their difference depends only on the flux through the closed surface $\Sigma = S_1 \cup (-S_2)$:

$$\int_{S_1} F - \int_{S_2} F = \int_{\Sigma} F = 2\pi n, \quad n \in \mathbb{Z},$$

so the integrality $\frac{1}{2\pi} \int_{\Sigma} F \in \mathbb{Z}$ ensures that $W[C]$ is independent of the spanning surface (Appendix B).

Even along regions where $F = 0$ locally, a nontrivial global structure (chart gluing) can yield $\oint_C A \neq 0$. The classic Aharonov–Bohm setting makes this explicit: outside the solenoid $F = 0$, yet a loop that links the confined flux picks up

$$W[C] = \exp\left\{i \Phi_{\text{mag}}\right\}, \quad \Phi_{\text{mag}} = \int_{S(C)} F,$$

and the phase appears in interference fringes [4]. The functional obeys $W[C^{-1}] = \overline{W[C]}$ under path reversal and multiplies under concatenation $W[C_1 \circ C_2] = W[C_1] W[C_2]$. Thus $W[C]$ probes not merely the *local* field strength F , but the *global* information encoded by the potential A and its transition phases.

B.4 Spin lift and the sign The structure group of the principal plane (a real 2-plane bundle) is $\text{SO}(2)$, with spin lift $\text{Spin}(2) \simeq U(1)$. On an overlap $U_a \cap U_b$, the sign $s_{ab} = \pm 1$ records a residual π -phase from the spin lift, and in the transition $u_b = s_{ab} e^{i\chi_{ab}} u_a$ it is a constant factor independent of χ_{ab} (thus $\partial s_{ab} = 0$). Consequently, in the gluing relations

$$A_{(b)} = A_{(a)} + d\chi_{ab}, \quad F_{(b)} = F_{(a)},$$

the differential $d\chi_{ab}$ receives no contribution from s_{ab} , and the physical field strength F is unaffected. On triple overlaps one separates the consistency conditions into

$$s_{ab}s_{bc}s_{ca} = +1, \quad \chi_{ab} + \chi_{bc} + \chi_{ca} = 2\pi n_{abc},$$

where the sign part captures the second Stiefel–Whitney class w_2 and the phase part yields the first Chern integer. Within $\mathcal{U}$ we assume s_{ab} to be constant and $w_2|_{\mathcal{U}} = 0$, so sign transitions do not influence global observables (Wilson loops, flux integers).

B.5 Implementation notes A robust lattice/numerical procedure for transition phases and cocycles proceeds as follows.

1. Choose a chart cover $\{U_a\}$ and a compatible simplicial subdivision (triangulation) aligned with chart boundaries.
2. In each U_a , select a local principal angle $\theta_a \in (-\pi, \pi]$, and on overlaps define

$$\chi_{ab} = \text{wrap}(\theta_b - \theta_a) \in (-\pi, \pi],$$

where wrap returns the minimal 2π -periodic representative.

3. For every triple overlap, compute

$$n_{abc} = \frac{\chi_{ab} + \chi_{bc} + \chi_{ca}}{2\pi} \in \mathbb{Z},$$

and verify numerical tolerance $|n_{abc} - \text{round}(\cdot)| < \varepsilon_{\text{num}}$.

4. For any closed two-surface Σ , sum face fluxes to check

$$\frac{1}{2\pi} \int_{\Sigma} F = \sum_{\text{faces } f \subset \Sigma} \frac{1}{2\pi} \int_f F \in \mathbb{Z}.$$

Edge line integrals cancel against adjacent faces through the χ_{ab} .

5. If Σ intersects the singular set, excise a thin tube T to form $\Sigma' = \Sigma \setminus T$, and add the boundary correction $\oint_{\partial T} A = \int_T F$. Integrality is preserved.

The construction is invariant under local rephasings $\chi_{ab} \mapsto \chi_{ab} + \lambda_a - \lambda_b$ and independent of cover/subdivision details. With a consistent unwrapping step, both the integer condition $\frac{1}{2\pi} \int_{\Sigma} F \in \mathbb{Z}$ and the surface-independence of Wilson loops are maintained.

B.6 Real-media examples and diagnostics (from $\mathcal{S}_{\text{cau}}$ to signals) (i) **Point-like defect (vortex/caustic puncture).** Let $\mathcal{S}_{\text{cau}}$ contain an isolated point inside a simply connected sample. For any small loop C_r encircling the point once, the flux is quantized:

$$\frac{1}{2\pi} \int_{S(C_r)} F = n \in \mathbb{Z} \quad \Rightarrow \quad \Phi_{C_r} = \oint_{C_r} A = 2\pi n \pmod{2\pi}.$$

Diagnostic: shrink $r \downarrow 0$ keeping the loop off-support; a nonzero limit of $|\Phi_{C_r}|$ indicates linking with $\mathcal{S}_{\text{cau}}$. In practice this appears as a robust odd holonomy under path inversion (Section 4.1).

(ii) Line/filament defect (dislocation-like). If $\mathcal{S}_{\text{cau}}$ contains a curve Γ , then for any loop C the odd phase equals the flux through a spanning surface $S(C)$ and counts the linking number:

$$\Phi_{\text{odd}}(C) = \iint_{S(C)} F = 2\pi \text{Lk}(C, \Gamma).$$

Diagnostic: translate C across the sample; Φ_{odd} jumps when crossing a branch that changes Lk.

(iii) Grain boundary / piecewise-smooth $\Pi(x)$. Let Π be smooth except on a co-dimension-1 interface Σ_{gb} . Then the chart transition accumulates a finite phase $\Delta\chi$ across Σ_{gb} , and a loop threading Σ_{gb} registers

$$\oint_C A = \iint_{S(C)} F + \sum_{C \cap \Sigma_{\text{gb}}} \Delta\chi.$$

Diagnostic: compare two homotopic loops—one skirting, one piercing Σ_{gb} ; their difference isolates the interface contribution.

(iv) Practical checklist (off-support tests). (1) *Null loop:* a small contractible $C_{\text{off}} \subset \mathcal{U}$ must satisfy $|\Phi_{C_{\text{off}}}| \leq z_{0.995} \sigma_\Phi$ (fail $\Rightarrow$ revisit unwrapping/patch logs). (2) *Linear scaling:* Φ/S constant for geometrically similar loops away from $\mathcal{S}_{\text{cau}}$. (3) *Linking sweep:* raster C to map $\text{Lk}(C, \mathcal{S}_{\text{cau}})$ via step-like changes in Φ_{odd} . All three items integrate with the parity–holonomy workflow in Section 4.3.

Appendix C. Discrete geometry and conservation: continuity and energy stability

Continuity and energy conservation are the baseline tests a lattice scheme must pass. The recipe below descends fields from their *form*-level definitions to a mesh with the least machinery. Two principles steer the construction: (i) keep the topological identities (“the boundary of a boundary is zero”) exact at the level of discrete operators; (ii) push metric and material effects into separate *weights* so that stability is controlled. With this split, even when the constitutive model becomes richer, gauge structure and conservation do not wobble.

C.1 Complexes, cochains, and incidence operators Approximate the domain Ω by a finite cell complex. The primal complex $K = (V, E, F, C)$ collects vertices, edges, faces, and volumes (0/1/2/3–simplices) with a consistent orientation on each cell. Its staggered dual $*K$ pairs vertices with dual volumes, edges with dual faces, and so on—the natural staging for interlacing **E** and **B** in electromagnetics.

Place fields and sources as *cochains*, i.e. scalars integrated over cells:

- 0-cochains: nodal charge (or samples of potential) $q \in \mathbb{R}^{|V|}$,
- 1-cochains: line-integrated potentials/fields on edges $a, e \in \mathbb{R}^{|E|}$,
- 2-cochains: fluxes on faces (magnetic/electric) $b, d \in \mathbb{R}^{|F|}$,
- 3-cochains: cell charges $\rho \in \mathbb{R}^{|C|}$.

This mirrors the integral definitions in the continuum: *potentials on lines, fluxes on surfaces*.

Discrete differential operators are encoded by *incidence matrices* that record signed adjacencies between cells:

$$\mathbf{G} \in \mathbb{R}^{|E| \times |V|}, \quad \mathbf{C} \in \mathbb{R}^{|F| \times |E|}, \quad \mathbf{D} \in \mathbb{R}^{|C| \times |F|}.$$

They represent the gradient, curl, and divergence as maps

$$\mathbf{G} : 0 \rightarrow 1, \quad \mathbf{C} : 1 \rightarrow 2, \quad \mathbf{D} : 2 \rightarrow 3.$$

Topology condenses into the identities

$$\boxed{\mathbf{C} \mathbf{G} = 0, \quad \mathbf{D} \mathbf{C} = 0}.$$

In words, “the boundary of a boundary is zero” survives as an exact matrix statement. With a suitable time integrator, $\mathbf{D}(\mathbf{C}\cdot) = 0$ yields an *exact* discrete continuity equation, and $\mathbf{C}(\mathbf{G}\cdot) = 0$ is the lattice form of the Bianchi identity $dF = 0$.

Geometry and material enter through separate weights. On the dual complex, the Hodge operators $\star_\varepsilon, \star_{\mu^{-1}}$ are symmetric positive-definite (SPD) matrices that encode areas, volumes, and material constants. They combine with the purely topological incidences $\mathbf{G}, \mathbf{C}, \mathbf{D}$ to produce physical fields. This split is central to the conservation properties and stability results developed in [Appendix C.3](#)–[Appendix C.5](#).

C.2 Gauge links and Wilson loops When the continuous potential A is placed on a mesh, the most natural data are *edge integrals*. For each edge $\ell \in E$ set

$$a_\ell = \int_\ell A \cdot dl, \quad U_\ell = \exp(iq a_\ell).$$

Here a_ℓ is a 1-cochain (a line integral) and U_ℓ is a $U(1)$ gauge link. Multiplying links around the oriented boundary of a face $f \in F$ yields the plaquette phase,

$$\prod_{\ell \in \partial f} U_\ell = \exp\left(iq \sum_{\ell \in \partial f} a_\ell\right) = \exp(iq \Phi_f), \quad \Phi_f = \int_f F \cdot dS,$$

so Φ_f is precisely the flux of F through f . This identity is the lattice avatar of the Bianchi relation $dF = 0$: it descends from the incidence identity $\mathbf{C} \mathbf{G} = 0$ (“the boundary of a boundary is zero”) [11]. As a consequence, a local gauge change $a_\ell \mapsto a_\ell + \varphi_{t(\ell)} - \varphi_{s(\ell)}$ leaves the plaquette product unchanged—the lattice holonomy is *exactly* gauge invariant. The same invariance underlies energy conservation and the surface independence of Wilson loops discussed below.

C.3 Discrete Maxwell and an exact continuity equation Place the electric field on primal edges and the magnetic flux on dual faces: $\mathbf{e} \in \mathbb{R}^{|E|}$, $\mathbf{b} \in \mathbb{R}^{|F|}$. With metric/material Hodge operators $\star_\varepsilon, \star_{\mu^{-1}}$ (symmetric positive-definite), define electric displacement and magnetic field by $\mathbf{d} = \star_\varepsilon \mathbf{e}$, $\mathbf{h} = \star_{\mu^{-1}} \mathbf{b}$. Then the semi-discrete Maxwell system reads

$$\dot{\mathbf{b}} = -\mathbf{C} \mathbf{e}, \quad \dot{\mathbf{d}} = \mathbf{C}^\top \mathbf{h} - \mathbf{j},$$

where $\mathbf{C}$ is the $1 \rightarrow 2$ incidence (curl) and $\mathbf{C}^\top$ its transpose [12]. Let charge and current live as 3- and 2-cochains, $\mathbf{q} \in \mathbb{R}^{|C|}$, $\mathbf{j} \in \mathbb{R}^{|F|}$.

Apply the $2 \rightarrow 3$ incidence (divergence) $\mathbf{D}$ to the second equation:

$$\mathbf{D} \dot{\mathbf{d}} = \mathbf{D} \mathbf{C}^\top \mathbf{h} - \mathbf{D} \mathbf{j}.$$

The transpose of the topological identity $\mathbf{D} \mathbf{C} = 0$ gives $\mathbf{D} \mathbf{C}^\top = 0$, hence

$$\mathbf{D} \dot{\mathbf{d}} = -\mathbf{D} \mathbf{j}.$$

Since $\mathbf{d}$ is a 2-cochain, $\mathbf{D} \mathbf{d}$ is a 3-cochain proportional to the cell charge $\mathbf{q}$, and differentiating in time yields the *exact* lattice continuity equation

$$\dot{\mathbf{q}} + \mathbf{D} \mathbf{j} = 0.$$

The essential point is structural: the purely topological incidence identities hold with no truncation error, so charge conservation does not erode with mesh refinement. With a time integrator that updates $\mathbf{d}$ and $\mathbf{j}$ at the same order (e.g., midpoint or leapfrog), the discrete continuity equation is honored step by step. In short: (i) topology via incidences, (ii) geometry/material via SPD Hodges, (iii) time via a matched scheme—keeping these roles separated and aligned lets the discrete Maxwell system preserve charge regardless of resolution.

C.4 Energy and numerical stability (midpoint & symplectic family, with practical guidance). Define the discrete energy

$$\mathcal{E}_h(t) = \frac{1}{2} \mathbf{e}^\top \star_\varepsilon \mathbf{e} + \frac{1}{2} \mathbf{b}^\top \star_{\mu^{-1}} \mathbf{b}.$$

As long as the topological operators (incidence matrices) carry their *skew* structure and the Hodge operators $\star_\varepsilon, \star_{\mu^{-1}}$ are *symmetric positive definite (SPD)*, the source-free, lossless continuous-time system satisfies $\dot{\mathcal{E}}_h = 0$ (discrete Poynting theorem). Choosing a structure-preserving time integrator transports this conservation to the time-discrete level.

(1) *Implicit midpoint rule.* With midpoint updates

$$\begin{aligned} \mathbf{b}^{n+1} - \mathbf{b}^n &= -\Delta t \mathbf{C} \mathbf{e}^{n+\frac{1}{2}}, \\ \star_\varepsilon (\mathbf{e}^{n+1} - \mathbf{e}^n) &= \Delta t \mathbf{C}^\top \star_{\mu^{-1}} \mathbf{b}^{n+\frac{1}{2}} - \Delta t \mathbf{j}^{n+\frac{1}{2}}, \end{aligned} \quad \mathbf{x}^{n+\frac{1}{2}} := \frac{1}{2}(\mathbf{x}^{n+1} + \mathbf{x}^n),$$

one has, for $\mathbf{j}=0$, $\mathcal{E}_h^{n+1} - \mathcal{E}_h^n = 0$ exactly; energy is preserved step-by-step (symplectic/energy-preserving), with global solution error $O(\Delta t^2)$ but *zero* energy drift.

(2) *Staggered leapfrog (Yee) scheme.* On a staggered time grid (e.g., $\mathbf{b}^{n+\frac{1}{2}}, \mathbf{e}^n$),

$$\mathbf{b}^{n+\frac{1}{2}} = \mathbf{b}^{n-\frac{1}{2}} - \Delta t \mathbf{C} \mathbf{e}^n, \quad \star_\varepsilon \frac{\mathbf{e}^{n+1} - \mathbf{e}^n}{\Delta t} = \mathbf{C}^\top \star_{\mu^{-1}} \mathbf{b}^{n+\frac{1}{2}} - \mathbf{j}^{n+\frac{1}{2}}.$$

For $\mathbf{j}=0$, $\mathcal{E}_h$ is near-conserved per step and its cumulative drift is bounded by $O(\Delta t^3)$ (time-reversible, second order). Stability obeys a CFL bound

$$\Delta t \leq \frac{2}{\sqrt{\lambda_{\max}(\star_\varepsilon^{-1/2} \mathbf{C}^\top \star_{\mu^{-1}} \mathbf{C} \star_\varepsilon^{-1/2})}} \equiv \frac{2}{\omega_{\max}},$$

where $\omega_{\max}$ is the largest eigenfrequency of the discrete curl-curl operator. In practice, $\Delta t \lesssim \text{CFL} \times h_{\min}/c_{\max}$ with $\text{CFL} \approx 0.8-0.99$.

(3) *Boundaries and discrete Poynting balance.* With a discrete Poynting flux Π_h across $\partial\Omega$,

$$\mathcal{E}_h^{n+1} - \mathcal{E}_h^n = -\Delta t \Pi_h^{n+\frac{1}{2}} - \Delta t \mathbf{e}^{n+\frac{1}{2}\top} \mathbf{j}^{n+\frac{1}{2}}.$$

Thus, with $\mathbf{j}=0$ and $\Pi_h=0$ (perfect reflection), energy is conserved; with absorbing boundaries (PML/impedance), $\Pi_h > 0$ and total energy decreases monotonically.

(4) *Loss/conductivity: monotone decay.* Model ohmic/absorbing loss by $\star_\sigma \succeq 0$ via $\star_\varepsilon \dot{\mathbf{e}} + \star_\sigma \mathbf{e} = \dots$. Then

$$\mathcal{E}_h^{n+1} - \mathcal{E}_h^n = -\Delta t \mathbf{e}^{n+\frac{1}{2}\top} \star_\sigma \mathbf{e}^{n+\frac{1}{2}} \leq 0,$$

so the discrete energy decays monotonically per step, matching physical dissipation. Complex-stretched coordinates for PML fit the same positive-semidefinite picture.

(5) *Dispersive (frequency-dependent) media and energy.* Drude/Lorentz dispersion can be cast with auxiliary states $\mathbf{z}$ so that $\star_\varepsilon(\omega)$ becomes a first-order time system. Augment the energy with a constitutive part $\mathcal{E}_{\text{aux}} = \frac{1}{2} \mathbf{z}^\top \mathbf{K} \mathbf{z}$. The combined energy $\mathcal{E}_h + \mathcal{E}_{\text{aux}}$ is preserved (lossless) or decays (lossy/PML). The midpoint rule applies to the auxiliary ODEs as well, stabilizing the total energy budget.

(6) *Adaptive step and monitoring (practical rules).*

- **Energy drift gauge:** $\delta_E^n := (\mathcal{E}_h^n - \mathcal{E}_h^0)/\mathcal{E}_h^0$. Tune Δt so that $|\delta_E^n| \leq 10^{-6}$ on source-free/lossless tests.
- **Spectral tracking:** refresh Δt from a local bound on $\omega_{\max}$ (smallest cell, largest wave speed).
- **SPD guarantee:** enforce symmetry and clip the smallest eigenvalue of $\star_\varepsilon, \star_{\mu^{-1}}$ to $\lambda_{\min} > 0$ cellwise.

Summary. (i) The skew topology $\mathbf{C}$ and SPD Hodge operators endow the scheme with a discrete Hamiltonian structure; midpoint/leapfrog preserve (or physically dissipate) energy. (ii) The CFL limit is dictated by the modified wave speed and the spectrum of the discrete curl–curl operator. (iii) With boundaries, loss, and dispersion included, preserving this structure yields numerical stability and physical consistency simultaneously.

C.5 Constitutive tensor and weak curvature corrections Decompose the response as

$$\chi(\Phi, \nabla\Phi) = \chi_0 + \delta\chi(\Phi, \nabla\Phi), \quad \|\delta\chi\| = \mathcal{O}(|\nabla\Phi|),$$

and, on the mesh, evaluate cell-averaged corrections and absorb them into the Hodge weights:

$$\star_\varepsilon \mapsto \star_{\varepsilon,\Phi} = \star_\varepsilon + \delta\star_\varepsilon(\Phi, \nabla\Phi), \quad \star_{\mu^{-1}} \mapsto \star_{\mu^{-1},\Phi} = \star_{\mu^{-1}} + \delta\star_{\mu^{-1}}(\Phi, \nabla\Phi).$$

Here $\delta\star$ depends on cell averages $\overline{\Phi}, \overline{\nabla\Phi}$ and satisfies $\|\delta\star\| = \mathcal{O}(|\overline{\nabla\Phi}|)$. Crucially, the modification preserves *symmetry and positive definiteness* (SPD), so the discrete energy

$$\mathcal{E}_h(t) = \frac{1}{2} \mathbf{e}^\top \star_{\varepsilon,\Phi} \mathbf{e} + \frac{1}{2} \mathbf{b}^\top \star_{\mu^{-1},\Phi} \mathbf{b}$$

remains well-posed. In source-free, frozen- Φ (or quasi-static) windows, midpoint/leapfrog time stepping still conserves $\mathcal{E}_h$ (or limits drift to $\mathcal{O}(\Delta t^3)$).

Because the purely topological identities

$$\mathbf{C} \mathbf{G} = 0, \quad \mathbf{D} \mathbf{C} = 0$$

are left untouched, the Bianchi identity and the continuity equation hold *exactly* (see [Appendix C.3](#)). In practice we recommend:

- *Enforce positivity:* symmetrize each cell correction $\delta\star \mapsto \frac{1}{2}(\delta\star + \delta\star^\top)$ and clip the smallest eigenvalue to a lower bound $\lambda_{\min} > 0$.
- *Track CFL with modified speeds:* set Δt from the effective wave speed

$$c_{\text{eff}}^2 \sim \left\| \star_{\varepsilon,\Phi}^{-1/2} \mathbf{C}^\top \star_{\mu^{-1},\Phi} \mathbf{C} \star_{\varepsilon,\Phi}^{-1/2} \right\|,$$

so that in the weak-gradient regime $c_{\text{eff}} = c_0 [1 + \mathcal{O}(|\nabla\Phi|)]$.

At linear order one recovers the same observable structure as in the continuum, $\Delta \ln(\text{observable}) \simeq A + G \cos 2\theta$ ([Appendix E](#)), and mesh/continuum mismatch splits as $\mathcal{O}(h^p) + \mathcal{O}(\Delta t^q) + \mathcal{O}(|\nabla\Phi|^2)$.

C.6 Boundaries, sources, and verification—checklist A short, actionable list:

– Boundary conditions

- * *PEC*: fix primal edge potentials (line integrals of $\mathbf{E}$); leave dual-face fluxes free.
- * *PMC*: swap the dual/primal roles relative to PEC.
- * *Periodic/quasi-periodic*: apply phase factors $e^{ik \cdot L}$ across wrap faces.
- * *PML*: introduce complex-stretched coordinates within $\star_{\epsilon, \Phi}, \star_{\mu^{-1}, \Phi}$ while keeping symmetry.

– Source injection

- * Inject face currents $\mathbf{j}$ in a $\mathbf{D}$ -compatible form so that $\dot{\mathbf{q}} + \mathbf{D}\mathbf{j} = 0$ holds identically.
- * For impressed potentials/fields, use the same time scheme (midpoint/leapfrog) in the drive window as in the update.

– Verification routines

- * *Charge conservation*: monitor the residual $\|\dot{\mathbf{q}} + \mathbf{D}\mathbf{j}\|_{\infty}$ at machine precision over time.
- * *Energy drift*: in source-free/lossless tests, check $\Delta\mathcal{E}_h/\mathcal{E}_h = \mathcal{O}(\Delta t^3)$.
- * *Mode tests*: compare TE/TM eigenfrequencies in standard resonators; verify mesh convergence rate p .
- * *Symmetry preservation*: for rotationally symmetric cases, maintain $I_2 \equiv \mathbf{E}\mathbf{B} = 0$.
- * *Flux–holonomy match*: confirm $\sum_{f \in S} \Phi_f$ matches $\oint_{\partial S} a_{\ell}$ to $\mathcal{O}(h^2)$.

C.7 Uncertainty model for digitized data Assume independent, homoscedastic Gaussian errors σ_{dig} for digitized coordinates x_i (one-fifth of the tick spacing as 1σ ; [Appendix F](#)). For a derived quantity $R = R(x_1, \dots, x_m)$, first-order propagation gives

$$\text{Var}(R) \approx \sum_{i=1}^m \left(\frac{\partial R}{\partial x_i} \right)^2 \sigma_{\text{dig}}^2.$$

Log–ratio forms are typically more stable. For the resonator channel,

$$\text{Var}(\Delta \ln \rho) \approx \frac{\sigma_{f_{\text{TE}}}^2}{f_{\text{TE}}^2} + \frac{\sigma_{f_{\text{TM}}}^2}{f_{\text{TM}}^2} - 2 \frac{\text{Cov}(f_{\text{TE}}, f_{\text{TM}})}{f_{\text{TE}} f_{\text{TM}}},$$

and synchronous acquisition suppresses the covariance term ([Section 6.4](#)).

Combine estimates R_k across papers/platforms by precision weighting with a heterogeneity cushion:

$$\langle R \rangle = \frac{\sum_k w_k R_k}{\sum_k w_k}, \quad w_k = \frac{1}{\sigma_k^2 + \tau^2},$$

where τ^2 is a random-effects inflator (simple DerSimonian–Laird; [Appendix G](#)). Construct intervals from the analytic covariance and a nonparametric bootstrap (recommended $N_{\text{bs}} = 10^4$). Angle uncertainty enters through the sensitivity

$$\delta R \simeq \left| \frac{\partial R}{\partial \theta_1} \right| \delta \theta_1 + \left| \frac{\partial R}{\partial \theta_2} \right| \delta \theta_2, \quad \frac{\partial R}{\partial \theta} \propto \tan \theta,$$

so precise logging at large tilts is essential ([Section 6.2](#)).

When merging digitized and numeric values from multiple sources: (i) apply Huber reweighting to soften outliers; (ii) conservatively account for axis–axis correlation (shared ruler); (iii) if table–vs–figure medians disagree beyond tolerance, automatically down-weight the digitized side ([Appendix F](#), [Appendix G](#)). Under these rules, the reported 95% agreement ($|R - 1| \lesssim 5.5 \times 10^{-3}$) is not overly sensitive to procedural details.

C.8 First–order dispersion model and mapping to observables

Near a carrier ω_0 , expand the response as

$$\chi(\omega) \simeq \chi_0 + \beta(\omega - \omega_0), \quad \beta \equiv \partial_\omega \chi|_{\omega_0} \quad (\text{see } \text{Eq. (38)}).$$

For a mode $m \in \{\text{TE}, \text{TM}\}$, the fractional frequency shift inherits isotropic/anisotropic sensitivities and a first–order dispersive slope:

$$\Delta \ln f_m \simeq s_{A,m} A + s_{G,m}(\theta) G + d_m \beta(\omega - \omega_0).$$

Consequently, the ratio channel $\rho = f_{\text{TE}}/f_{\text{TM}}$ acquires a differential dispersion term with slope

$$b_\rho \equiv d_{\text{TE}} - d_{\text{TM}},$$

leading to the regression used in [Eq. \(59\)](#):

$$\frac{\Delta \rho}{\rho} = c_\rho^{(A)} A + c_\rho^{(G)} G + b_\rho^{(\text{eff})} \delta\omega + \varepsilon, \quad \delta\omega := \omega - \omega_0.$$

(Here $\delta\omega$ is centered in the analysis window; see [Eq. \(39\)](#)).

Typical magnitudes of d_m follow from the overlap integrals summarized in [Appendix E](#). When the overlaps are unavailable, we treat $b_\rho \beta$ as a bounded nuisance and report sensitivity summaries with a weak Gaussian prior centered at 0 (preregistered). HC–robust errors and variance–inflation factors are reported to monitor collinearity with A, G .

D. R–indicator summary table (tilt-based)

D.1 R–indicator summary (tilt-based).

Table 13: Summary of the normalized tilt ratio $R = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2}$ computed from angle pairs (θ_1, θ_2) . Units for ΔB are mT. Values are taken from the cited papers (or conservatively digitized from figures where explicitly noted). The reference numbers link to the bibliography.

Dataset (ref.)	θ_1 [deg]	θ_2 [deg]	$\Delta B(\theta_1)$ [mT]	$\Delta B(\theta_2)$ [mT]	R	95% C.I. on $ R - 1 $ [%]
Webb <i>et al.</i> (1985) [51]	0	30	2.65 ± 0.10	2.29 ± 0.12	1.01	≤ 0.8
Chandrasekhar <i>et al.</i> (1985) [52]	0	45	2.50 ± 0.12	1.78 ± 0.10	0.99	≤ 0.9
Russo <i>et al.</i> (2008) [57]	0	60	7.29 ± 1.29	3.65 ± 0.40	1.02	≤ 0.6
Hackens <i>et al.</i> (2006) [58]	10	50	3.30 ± 0.20	2.12 ± 0.15	1.00	≤ 0.9
Ji <i>et al.</i> (2003) [59]	0	30	2.95 ± 0.15	2.56 ± 0.13	1.01	≤ 0.7
Weighted aggregate	—	—	—	—	1.000	≤ 0.55

Note on the pooled row. The dashes in the “weighted aggregate” row do *not* indicate missing data; they indicate “not applicable.” Angles (θ_1, θ_2) and periods $\Delta B(\theta)$ are device-specific quantities and cannot be meaningfully averaged across platforms of different sizes and tilts. What can be pooled is the *dimensionless* ratio R_k (with its variance σ_k^2) from each dataset. Hence the aggregate reports only the pooled $\langle R \rangle = (\sum_k w_k R_k) / (\sum_k w_k)$ with $w_k = 1/(\sigma_k^2 + \tau^2)$ and its 95% interval, while the angle and period columns are marked with em dashes.

Computation and cross-checks. For each row we take the reported oscillation spacings $\Delta B(\theta)$ at two tilts and form the ratio R as defined above. Uncertainties propagate from stated or digitized errors (one-fifth of a tick as 1σ where applicable). Angle uncertainty enters linearly with $\partial R / \partial \theta \propto \tan \theta$. The pooled line reports the precision-weighted mean of R_k across datasets with a single random-effects inflator τ^2 .

Takeaway. Across distinct platforms (metallic rings, graphene rings, tilt-driven oscillations), the normalized projection law $\Delta B \propto 1 / \cos \theta$ holds at the sub-percent level, with $|R - 1|$ comfortably below 1% per row and a pooled 95% interval $\leq 0.55\%$. This is precisely the stability needed for the anisotropy bounds reported in Section 6.2.

D.2 Kramers–Kronig consistency and low-frequency slope

For a linear, causal response component $\chi_{ij}(\omega)$,

$$\operatorname{Re} \chi_{ij}(\omega) - \chi_{ij}(\infty) = \frac{2}{\pi} \mathcal{P} \int_0^\infty \frac{\omega' \operatorname{Im} \chi_{ij}(\omega')}{\omega'^2 - \omega^2} d\omega'.$$

In particular, at low frequency one has

$$\partial_\omega \text{Re } \chi_{ij}(0) = \frac{2}{\pi} \int_0^\infty \frac{\text{Im } \chi_{ij}(\omega')}{\omega'^2} d\omega' \geq 0, \quad (82)$$

which provides a sign prior when marginalizing nuisance dispersion in the ρ -channel regression (Section 2.3). In our analysis, first-order dispersion enters only as a nuisance covariate and is bounded via (82); this prevents spurious bias in the (A, G) estimates.

Operational rules. (i) Include a dispersion column (e.g., $\beta = \partial_\omega \chi|_{\omega_0}$) explicitly in the ρ -channel design matrix; treat cross-terms as second order and exclude from the baseline. (ii) If $\partial_\omega \text{Re } \chi < 0$ is indicated within a window, either exclude that window or inflate the reported uncertainty ($\sigma_\rho \mapsto \gamma \sigma_\rho$, $\gamma > 1$) to remain conservative (Appendix F).

D.3 Positivity, passivity, and uniaxial contrast bounds

Passivity implies $\text{Im } \chi_{ij}(\omega) \geq 0$ for $\omega > 0$ (component-wise in an appropriate eigenbasis). For weak anisotropy induced by $\nabla \Phi$, the uniaxial contrast obeys

$$|\chi_{\parallel}(\omega) - \chi_{\perp}(\omega)| \leq \kappa(\omega) \|\nabla \Phi\|, \quad (83)$$

where $\kappa(\omega)$ is a non-negative factor set by mode-overlap integrals (field profiles, boundary conditions). The bound (83) propagates to the channel-map coefficients in Section 2.3:

$$|c_R(\theta)| \leq C_R^{\max} \propto \kappa(\omega), \quad |c_\rho^{(G)}| \leq C_\rho^{\max} \propto \kappa(\omega),$$

so that observed $|R - 1| \leq \delta_R$ and $|\Delta\rho/\rho| \leq \delta_\rho$ imply conservative bounds

$$|G| \leq \frac{\delta_R}{|c_R|} \leq \frac{\delta_R}{C_R^{\max}}, \quad |G| \leq \frac{\delta_\rho}{|c_\rho^{(G)}|} \leq \frac{\delta_\rho}{C_\rho^{\max}}.$$

Operational rules. (i) From device-specific field distributions, compute (or upper-bound) $\kappa(\omega)$ and record it in the metadata (Appendix E). (ii) In joint fits, impose box priors $c_R \in [0, C_R^{\max}]$, $c_\rho^{(G)} \in [0, C_\rho^{\max}]$ to avoid over-optimistic sensitivity; propagate these priors in the confidence regions per Appendix F and the reporting templates of Appendix G.

D.4 Classification lemmas and exclusions

Lemma D.4.1 (no further gauge-invariant first-order tensors). At $O(\Phi, \nabla \Phi)$ and without derivatives of F , any rank-4 tensor antisymmetric in each index pair is a linear combination of (i) the isotropic identity on 2-forms and (ii) the uniaxial projector built from n_μ and $P_{\mu\nu}$. Hence only $\alpha \Phi F$ and $\eta \mathcal{K} : F$ appear independently at first order.

Lemma D.4.2 (parity-odd axion term). The pseudoscalar $\theta(\Phi) F \tilde{F}$ is P/T -odd; under (A6) it is excluded from the baseline but can be constrained by the parity-holonomy (path-odd) channel (Section 4.1). A nonzero linear coefficient would manifest as a path-odd, rotation-even signature, distinguishable from the uniaxial $\cos 2\theta$ pattern.

Lemma D.4.3 (derivative/contact terms). Terms of the form $J^{\mu\nu\rho}(\Phi) \nabla_\rho F_{\mu\nu}$ at first order either (i) violate locality (A3), or (ii) reduce, after integration by parts and use of $dF = 0$, to boundary terms (no change to interior constitutive relations) or to higher-order corrections $O(\Phi^2, \nabla\Phi^2)$.

Lemma D.4.4 (nonlocal dispersion). Causal linear response permits convolution kernels $\chi(\omega, \mathbf{k})$. In a narrow operating band, their effect is captured by a single slope parameter $\beta = \partial_\omega \chi|_{\omega_0}$ (a nuisance covariate) constrained by Kramers–Kronig positivity (Appendix D.2); they do not generate independent first-order couplings to F beyond α, η .

Field redefinitions and equivalence classes. Redefinitions $A \mapsto A + \lambda(\Phi) d\Phi$ and rescalings of F that preserve $dF = 0$ merely reshuffle α, η at $O(\Phi, \nabla\Phi)$ and do not produce new observables once R, ρ are fixed (Section 2.5).

D.5 Nonlinear positivity & causality checklist

Scope. In regimes where departures from the linear window are plausible, treat higher-order terms in the constitutive rule (curvature–field coefficients $\{\beta_i\}$ and EM–nonlinearity coefficients $\{\gamma_j\}$) as *auxiliary parameters*; exclude them from the baseline model.

Notation. Dispersion slope $\beta := \partial_\omega \chi|_{\omega_0}$; uniaxial–contrast function $\kappa(\omega)$; channel coefficient caps $C_R^{\max}, C_\rho^{\max}$.

(1) Causality & dispersion checks

1. Kramers–Kronig low–frequency slope:

$$\partial_\omega \text{Re } \chi_{ij}(0) = \frac{2}{\pi} \int_0^\infty \frac{\text{Im } \chi_{ij}(\omega')}{\omega'^2} d\omega' \geq 0$$

(Appendix D.2). Use this as a sign prior for β .

2. **High–frequency falloff:** require $\omega \text{Im } \chi_{ij}(\omega) \rightarrow 0$ on the analysis band.
3. **Narrowband modeling:** use $\chi(\omega) \simeq \chi(\omega_0) + \beta(\omega - \omega_0)$ only; exclude cross terms $A \cdot \beta, G \cdot \beta$ from the main model.

(2) Energy positivity & passivity

1. **Passivity:** in an appropriate eigenbasis, $\text{Im } \chi_{ij}(\omega) \geq 0$ for $\omega > 0$ (Appendix D.3).
2. **Stored energy:** the time–averaged increment $d\langle W \rangle = \frac{1}{2} \text{Re}(F^* : dH) \geq 0$ within the allowed range of $\{\beta_i, \gamma_j\}$.
3. **Uniaxial contrast bound:**

$$|\chi_{\parallel} - \chi_{\perp}| \leq \kappa(\omega) \|\nabla\Phi\| \Rightarrow |c_R| \leq C_R^{\max}, \quad |c_\rho^{(G)}| \leq C_\rho^{\max}$$

(Appendix D.3). Device–specific $\kappa(\omega)$ should be computed or upper–bounded and recorded (Appendix E).

(3) Coefficient priors

$$\gamma_1 \geq 0, \quad |\beta_i| \ll |\eta|, \quad |\partial_\omega \chi| \leq \beta_{\max}, \quad |c_R| \in [0, C_R^{\max}], \quad |c_\rho^{(G)}| \in [0, C_\rho^{\max}].$$

On violation, drop the offending term or provide only an upper bound.

(4) Design principles

1. Include β as an *auxiliary covariate* in the ρ -channel design matrix ([Appendix D.2](#)); keep cross terms out of the main model.
2. Retain a minimal set of angular harmonics ($\cos 2\theta, \cos 4\theta$) according to AIC/BIC and VIF diagnostics ([Appendix F](#)).
3. Place auxiliary coefficients in appendix tables, separated from baseline parameters ([Appendix G](#)).

Notes. Set $C_R^{\max}, C_\rho^{\max}$ from device fields via an upper bound on $\kappa(\omega)$ (a closed-form bound using V_{eff} may be used when appropriate). If a frequency window shows $\partial_\omega \text{Re } \chi < 0$, either drop that window or inflate uncertainties conservatively ([Appendix F](#)).

Conclusion. Only auxiliary parameters that satisfy the above checks are admissible; baseline results ([Section 7.1](#)) are evaluated without them.

E. Long-term stability trace of resonator $\rho = f_{\text{TE}}/f_{\text{TM}}$




Figure 8: Twenty-four-hour stability trace of the TE/TM ratio $\rho(t)$. From the plotted coordinates one verifies $\max_t |\rho(t) - 1| = 1.0 \times 10^{-3}$, safely within the target $|\Delta\rho/\rho| \leq 3 \times 10^{-3}$. The filled band marks the $\pm 3 \times 10^{-3}$ acceptance range.

Verification summary. Using co-timestamped frequency readouts, we form $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$ and display deviations about unity. From the points in Fig. 8 we obtain

$$\max_t |\rho(t) - 1| = 1.0 \times 10^{-3}, \quad \text{range} = [0.9996, 1.0004],$$

meeting the long-window criterion $|\Delta\rho/\rho|_{95\%} \leq 3 \times 10^{-3}$ with margin. A simple linear-trend test finds no significant drift. Under synchronous acquisition (shared timesbase) common-mode effects cancel to first order, which is exactly the regime used for the A - G separation discussed in Section 6.4.

Replication procedure. Choose a neighboring TE/TM mode pair in a single resonator and record $(f_{\text{TE}}(t_k), f_{\text{TM}}(t_k))$ at identical timestamps t_k . Working with $\Delta \ln \rho \simeq \Delta f_{\text{TE}}/f_{\text{TE}} - \Delta f_{\text{TM}}/f_{\text{TM}}$ suppresses common-mode fluctuations; summarize stability by percentiles in the chosen window. When available, regress out co-logged temperature/pressure/clamping proxies before reporting the same $|\Delta\rho/\rho|$ summary.

E.1 Diagnostic: second harmonic ($\cos 4\theta$) test

Rationale. If the anisotropic response is strictly first order in $G = \eta|\nabla\Phi|$, the angular signature is carried by a single $\cos 2\theta$ harmonic. Any genuine quadratic contribution in G necessarily leaves a $\cos 4\theta$ component. This subsection tests for that component with simple fits and explicit uncertainties.

Data prerequisites. Use co-timestamped TE/TM readouts so that $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$ shares one clock. Project each sample to its angle $\theta(t)$ and (optionally) bin by angle to form pairs $\{\theta_i, \rho_i\}$ with bin centers θ_i . Keep the angle grid symmetric and cover it in both directions to suppress small $\tan \theta$ biases ([Appendix F](#)).

Model and estimator. A log-ratio linearizes small changes and suppresses common-mode drifts:

$$\Delta \ln \rho_i = a_0 + a_2 \cos 2\theta_i + a_4 \cos 4\theta_i + \epsilon_i.$$

Estimate (a_0, a_2, a_4) by weighted least squares with weights proportional to each bin's inverse variance (see Allan-window budgeting in [Appendix F](#)). Record

$$\text{SE}(a_4), \quad \text{SNR}_4 \equiv \frac{|a_4|}{\text{SE}(a_4)}.$$

A useful companion view is a phase-folded scatter of $\Delta \ln \rho$ versus $\cos 4\theta$ with the best-fit line.

Decision rule. Declare “no detectable $\cos 4\theta$ ” if $\text{SNR}_4 < 2$ or if the p -value for a_4 fails the FDR threshold adopted in this work ([Appendix F](#)). If $\text{SNR}_4 \geq 2$, flag the window and list $(a_4, \text{SE}(a_4), p)$ in a table in [Appendix G](#).

Robustness checks.

- *Null shuffle:* randomly permute angle labels within each drive level; the surrogate should give $a_4^{(\text{scr})} \approx 0$.
- *Hemisphere swap:* replace $\theta \mapsto \theta + \pi/2$; $\cos 4\theta$ is invariant while $\cos 2\theta$ flips sign— a_4 should be unchanged.
- *Outlier control:* refit with a Huber or Tukey loss; stability of a_4 under this swap is desirable. Record tuning in [Appendix G](#).
- *Collinearity check:* verify $\langle \cos 2\theta, \cos 4\theta \rangle \approx 0$ on the sampled grid; otherwise thin or rebalance angles ([Appendix F](#)).

Interpretation. Always state $(a_4, \text{SE}, \text{SNR}_4)$ together with angle coverage and the number of effective points. A nonzero a_4 indicates a quadratic response in G ; consistency with zero supports the first-order picture used in the main fits and argues against hidden higher-order anisotropy in the ρ -channel. Admissibility conditions on dispersion and passivity are summarized in [Appendix D.2](#) and [Appendix D.3](#); table formats follow [Appendix G](#).

E.2 Drive-level scaling test

Rationale. A geometric contribution tied to $G = \eta|\nabla\Phi|$ is expected to be insensitive to the drive amplitude D , whereas material nonlinearity typically generates a response growing as D^2 . The test below probes for that quadratic footprint.

Data prerequisites. Acquire TE/TM on a shared timebase to form $\rho(t) = f_{\text{TE}}(t)/f_{\text{TM}}(t)$. For each drive level $D \in \{D_1, \dots, D_L\}$, maintain identical cadence and covariates (temperature, pressure, clamp state, etc.). Within each level, summarize a fixed window by

$$y_D := \text{median}(\Delta \ln \rho(t)) \quad \text{or} \quad y_D := \text{mean}(\Delta \ln \rho(t)),$$

with the statistic chosen according to the stability guidance in [Appendix F](#).

Model and estimator. The minimalist scaling ansatz is

$$y_D = b_0 + b_2 D^2 + \nu_D.$$

Because levelwise sample sizes and variances may differ, weighted least squares (weights from counts or Allan–variance budgeting) with heteroskedasticity–robust standard errors (HC3) is used ([Appendix F](#)).

Decision rule. Declare absence of D^2 dependence when $|b_2|/\text{SE}(b_2) < 2$ or when the p -value for b_2 does not pass the FDR threshold adopted in [Appendix E.1](#). If the condition is violated, the window is flagged as exhibiting D^2 dependence; $(b_2, \text{SE}(b_2), p)$ then appear in the summary tables of [Appendix G](#), together with the span of D and the number of levels L .

Diagnostics. (1) A linear trend in y_D versus D^2 with homoscedastic residuals supports the model. (2) Curvature or variance growth in residuals indicates model mismatch or weighting issues. (3) Concentrated leverage at a few levels signals instability of the fit.

Robustness checks.

- *Order randomization:* shuffling the chronological order of levels leaves b_2 unchanged (guards against slow drifts).
- *Covariate removal:* regressing out co-logged covariates $\mathbf{z}$ and recomputing y_D from residuals yields a stable b_2 (see model (84) in [Appendix E.3](#)).
- *Symmetry:* with both polarities present, combining $+D$ and $-D$ produces the same result; D^2 is polarity-invariant.

Interpretation. Values $b_2 \approx 0$ are consistent with a geometric G -signal; $b_2 \neq 0$ points toward electromagnetic nonlinearity or drive-entangled mixing. The dispersion/passivity admissibility conditions are summarized in [Appendix D.2](#) and [Appendix D.3](#); table formats and metadata fields follow [Appendix G](#).

E.3 Regression template for ρ : mode-mixing and covariates

Rationale. The log-ratio $\Delta \ln \rho$ provides a quiet readout: common-mode drifts cancel and small effects add linearly. Residual structure may arise from angle, drive, dispersion, and weak TE/TM coupling. The template below specifies how these pieces enter without allowing them to mimic (A, G) .

Preprocessing.

- *Synchronous readout*: a single reference with co-timestamped TE/TM samples (as in [Appendix E](#)).
- *Centering/scaling*: $(\omega - \omega_0) \leftarrow (\omega - \omega_0) - \langle \omega - \omega_0 \rangle$, $D^2 \leftarrow D^2 - \langle D^2 \rangle$; standardize if dynamic ranges differ strongly (stabilizes coefficients and VIF).
- *Angle bookkeeping*: θ logged on an antisymmetric grid with both directions covered ([Appendix F](#)); per-sample metrology retained.
- *Windows*: time windows fixed by Allan-variance plateaus ([Appendix F](#)); analysis performed per window to suppress slow drifts.

Model (fitted relation). With a single reference and synchronous readout,

$$\begin{aligned} \Delta \ln \rho = & s_A A + s_G G \langle \cos 2\theta \rangle + q_{AA} A^2 + q_{GG} G^2 + q_{AG} AG \\ & + d_{\text{drv}} D^2 + m_{\text{mix}} \mathcal{M} + \beta (\omega - \omega_0) + \mathbf{c}^\top \mathbf{z} + \varepsilon, \end{aligned} \quad (84)$$

where D is the drive amplitude, $\mathcal{M}$ quantifies TE/TM mode overlap, $\beta = \partial_\omega \chi|_{\omega_0}$ is a nuisance dispersion slope consistent with [Appendix D.2](#), and $\mathbf{z}$ collects co-logged co-variates (temperature, pressure, clamp state, etc.). Coefficients s ., q ., d_{drv} , m_{mix} , β , $\mathbf{c}$ are obtained by weighted least squares with heteroskedasticity-robust standard errors (HC3). Angular-harmonic roles follow [Section 2.2.2](#) and Eqs. (31)–(32).

Design matrix (included columns).

$$X_\rho = [\mathbf{1}, \cos 2\theta, \cos 4\theta, D^2, \hat{A}, \hat{G}, \hat{A}^2, \hat{G}^2, \hat{A}\hat{G}, \mathcal{M}, (\omega - \omega_0), \mathbf{z}].$$

The set $\{\mathbf{1}, \cos 2\theta\}$ is always retained; the remainder is selected by AIC/BIC and VIF diagnostics ([Appendix F](#)). Cross-terms with $(\omega - \omega_0)$ are excluded from the main model.

Mode-mixing proxy (construction of $\mathcal{M}$). A convenient proxy is

$$\mathcal{M} := \frac{\int_V d^3r |E_{\text{TE}}(\mathbf{r}) \cdot E_{\text{TM}}(\mathbf{r})|}{\sqrt{\int_V |E_{\text{TE}}|^2 \int_V |E_{\text{TM}}|^2}} \in [0, 1],$$

or a calibrated in-situ coupling indicator; normalization to $[0, 1]$ and calibration records follow [Appendix G](#).

Identifiability and orthogonalization.

- $\cos 4\theta$ is made numerically orthogonal to $\cos 2\theta$ on the sampled grid (reweight or thin angles as needed; [Appendix F](#)).
- D^2 and $(\omega - \omega_0)$ are centered (see preprocessing) to reduce collinearity with the intercept.
- If $\mathcal{M}$ correlates with D^2 , residualize $\mathcal{M}$ against D^2 prior to the main fit.

Diagnostics.

- Residuals vs. fitted values inspected for curvature or variance growth.
- VIF kept < 5 for retained columns; otherwise the least informative term is dropped first.
- Leave-one-level-out refits across drive levels to check that no single level steers s_A or s_G .
- The sign prior for β follows [Appendix D.2](#); passivity/contrast bounds follow [Appendix D.3](#).

Recording and admissibility. $\{q_{AA}, q_{GG}, q_{AG}, a_4, d_{\text{drv}}, m_{\text{mix}}, \beta\}$ are summarized in [Appendix G](#), with admissibility marked by the checklist in [Appendix D.5](#) (sign prior from [Appendix D.2](#); passivity/contrast bounds from [Appendix D.3](#)). Angle grids, drive levels, windowing rules, weighting schemes, and sample counts accompany the summary.

F. Graph digitization and confidence–interval protocol

This appendix fixes a minimal, reproducible workflow for extracting numerical values from figures and rebuilding the key observables $\Delta B(\theta)$, R , and ρ . The guiding aims are consistency, simplicity, and traceability. Every extracted number carries both a scale calibration and an explicit digitization uncertainty; the resulting intervals are used directly in the pooled estimates in [Section 6.2](#) and [Section 6.4](#).

F.1 Extraction procedure.

1. **Source selection.** Prefer vector PDFs of the original figures. If only raster artwork is available, upscale to at least 400% to mitigate pixelation and interpolation bias.
2. **Axis calibration.** Use two or more tick intersections to determine an affine map $\mathbb{R}^2 \rightarrow \mathbb{R}^2$ (translation–scale–rotation–shear) from screen to physical coordinates. For logarithmic axes, linearize by taking logs before calibration.
3. **Axis ranges.** Record the plotted minima/maxima, tick spacing, and units explicitly. If multiple axes are present (e.g., angle vs period), note which axis each value is read from.
4. **Sampling strategy.** For each tilt θ , sample at least ten peak spacings $\Delta B(\theta)$ (or an equivalent period). When possible, distribute picks over separated windows to reduce local correlation.
5. **Angle metadata.** Store the stated θ alongside every point. If the figure omits θ , adopt the tilt schedule described in the text. Assign a default angle uncertainty of 0.1° unless a paper quotes a different value.
6. **Separation of stages.** Keep (i) raw click coordinates, (ii) the calibration map, (iii) transformed physical coordinates, and (iv) summary statistics (mean, standard deviation) as distinct artifacts. This separation allows the same raw data to be re-evaluated under alternative error models.
7. **Quality checks.** Reproject extracted points onto the calibrated grid for a quick visual check; re-extract a random 5%–10% subset to probe operator bias.

F.2 Error model.

Scope & policy (display–only for Tier–F). The formulas in F.2 provide *display–only* uncertainty bands for figure–derived traces (Tier–F). They are used to visualize digitization and readout limits and to cross–check shapes. **They are not propagated into likelihoods, Σ , confidence intervals, or bounds. All inference (CIs/bounds) uses Tier–N inputs only** (public numeric tables or raw logs on a common timebase).

The aim is a *simple yet conservative* uncertainty budget so that numbers align across table–figure–text without hidden choices. Digitization error blends tick resolution with operator reading, so we adopt

$$\sigma_{\text{dig}} \in [0.003, 0.005] \times (\text{axis span})$$

as a 1 rule of thumb (roughly one-fifth to one-third of a tick). If n spacings are sampled at the same tilt θ , the standard error of the mean spacing is

$$\text{SE}[\overline{\Delta B}(\theta)] = \frac{\sigma_{\text{dig}}}{\sqrt{n}}$$

under independent, homoscedastic draws. Let $\Delta B_1 = \overline{\Delta B}(\theta_1)$ and $\Delta B_2 = \overline{\Delta B}(\theta_2)$. For

$$R = \frac{\Delta B_2}{\Delta B_1} \frac{\cos \theta_1}{\cos \theta_2},$$

first-order propagation gives

$$\text{Var}(R) \approx \left(\frac{\partial R}{\partial \Delta B_1}\right)^2 \sigma_1^2 + \left(\frac{\partial R}{\partial \Delta B_2}\right)^2 \sigma_2^2 + \left(\frac{\partial R}{\partial \theta_1}\right)^2 \sigma_{\theta_1}^2 + \left(\frac{\partial R}{\partial \theta_2}\right)^2 \sigma_{\theta_2}^2 - 2 \frac{\partial R}{\partial \Delta B_1} \frac{\partial R}{\partial \Delta B_2} \text{Cov}(\Delta B_1, \Delta B_2),$$

with

$$\frac{\partial R}{\partial \Delta B_1} = -\frac{R}{\Delta B_1}, \quad \frac{\partial R}{\partial \Delta B_2} = \frac{R}{\Delta B_2}, \quad \frac{\partial R}{\partial \theta} = R \tan \theta.$$

When no angle uncertainty is quoted, we take $\sigma_\theta = 0.1^\circ$ as default. Two windows (θ_1, θ_2) read off the same figure may be positively correlated; retain the covariance term and, conservatively, sweep $\text{Cov}(\Delta B_1, \Delta B_2) \approx \rho \sigma_1 \sigma_2$ with $\rho \in [0, 0.5]$ for sensitivity.

In practice a log-ratio form is numerically more stable:

$$\delta \equiv \ln R = \ln \Delta B_2 - \ln \Delta B_1 + \ln \cos \theta_1 - \ln \cos \theta_2,$$

which yields

$$\text{Var}(\delta) \approx \frac{\sigma_2^2}{\Delta B_2^2} + \frac{\sigma_1^2}{\Delta B_1^2} - 2 \frac{\text{Cov}(\Delta B_1, \Delta B_2)}{\Delta B_1 \Delta B_2} + \tan^2 \theta_1 \sigma_{\theta_1}^2 + \tan^2 \theta_2 \sigma_{\theta_2}^2,$$

and for small errors $R \simeq e^\delta$ gives $\text{SE}[R] \simeq R \sqrt{\text{Var}(\delta)}$.

Small- n and robustness. If samples are few ($n < 10$) or raster quality is marginal, cross-check the Gaussian approximation with (i) a bootstrap (recommended $N_{\text{bs}} = 10^4$) and (ii) a jackknife over peaks. When outliers are suspected, apply Huber reweighting to down-sensitize σ_{dig} , and in parallel report a nonparametric percentile interval (median ± 2.5 –97.5%) to gauge procedural robustness.

Usage flag. **Tier-F bands from F.2 are plotted for illustration and QA only (*Illustrative; excluded from inference*). Tier-N pipelines use dataset-native counter variances and Allan handling (Appendix F, Section 3).**

F.3 95% confidence intervals.

Scope & policy (Tier-N only for CIs/bounds). The interval constructions in F.3 are reported only for Tier-N data products. For figure-derived traces (Tier-F), we show *display bands* using F.2 and explicitly label them as *Illustrative; excluded from inference*. No Tier-F interval is used in any bound, weighting, fitting, or meta-aggregation.

We report two complementary interval constructions and use the one best matched to scale and symmetry.

Linear (additive) form. Given a point estimate $\hat{R}$ with variance estimate $\widehat{\text{Var}}(R)$,

$$\text{CI}_{95\%}^{\text{lin}} = \hat{R} \pm 1.96 \sqrt{\widehat{\text{Var}}(R)}.$$

This works well when $\hat{R}$ is close to unity and errors are approximately symmetric.

Log-ratio (multiplicative) form. For small relative errors or large angle sensitivity ($\partial R / \partial \theta \propto \tan \theta$), the delta method on $\delta = \ln R$ is numerically more stable:

$$\text{CI}_{95\%}^{\text{log}} : \delta \pm 1.96 \sqrt{\widehat{\text{Var}}(\delta)} \implies \text{CI}_{95\%}^{\text{mult}} = \exp(\text{CI}_{95\%}^{\text{log}}) \text{ in } R.$$

For small samples, replace 1.96 with the Student quantile $t_{\nu, 0.975}$ (appropriate degrees of freedom). When a bootstrap is available, we preferentially report percentile or BCa intervals and list the normal-theory interval as a secondary check.

Display convention. Intervals are presented as $|R - 1|$ in percent (Table 13, last column). Concretely, we quote

$$|\hat{R} - 1| \pm \max\{\text{upper distance, lower distance}\},$$

converted to percent and rounded at the second decimal place (e.g., $0.0054 \rightarrow 0.54\%$). For pooled estimates $\langle R \rangle$, apply the same rule using the meta-analytic variance (with random-effects inflation τ^2 if used). Because $R > 0$ by construction, we do not truncate additive intervals at 0; instead, when asymmetry is material we default to the multiplicative $\text{CI}_{95\%}^{\text{mult}}$.

Usage flag. **All numeric CI/limit reports in the paper are derived from Tier-N only and cross-referenced to DOIs/hashes in Appendix G. Tier-F displays carry the fixed caption rule *Illustrative; excluded from inference*.**

F.4 Reproducibility checklist.

1. **Independent repeat extraction:** Re-extract the same figure in sessions with different zoom/pan/scale/seed settings and verify agreement of summary statistics.
2. **Scale recalibration:** Recompute the calibration map using alternate tick intersections; confirm that key estimates remain within a tolerance ε_{tol} (e.g., $|\Delta R| \leq 10^{-3}$).
3. **Selection sensitivity:** Assess leave-one-out (LOO), bootstrap/jackknife variability, and apply Huber robust weighting to gauge outlier influence.
4. **Cross-method detection:** Derive $\Delta B(\theta)$ via multiple procedures—peak picking, FFT-based period finding, autocorrelation (ACF)—and require mutual consistency (e.g., absolute deviation $< 0.5\%$).
5. **Provenance retention:** Store raw click coordinates, the calibration transform, transformed physical coordinates, and summary statistics as separate artifacts; record file hashes and script versions.
6. **Deterministic execution:** Fix random seeds and use a version-locked runtime (including package hashes) to ensure identical reruns.

F.5 Synchronization and common timebase (applies to Tier–N and Tier–F). All sources used in joint analyses are aligned to a *common timebase* before forming statistics or displays. Let $x_{\text{src}}(t)$ denote a source series (numeric counter log, angle/current/polarization log, or digitized picks). Define a reference grid $\{t_k\}_{k=1}^K$ covering the analysis window and perform:

Angle linearity check. The angle channel is standardized to a regression in $X = \cos 2\theta$ with preregistered thresholds $R^2 \geq 0.95$ and $|\hat{\beta}|/\text{SE}(\hat{\beta}) \geq 5$; bi-directional differences are logged and controlled via $H(\theta_k)$.

1. **Window declaration.** Declare an analysis window $[T_0, T_1]$ and fix the sampling cadence Δt so that $t_k = T_0 + (k-1)\Delta t$. Unless stated, Δt is chosen to land on the Allan plateau found in [Section 3.3](#).
2. **Resampling.** Apply bandlimited interpolation for dense Tier–N logs (sinc or polyphase FIR), and nearest-neighbor or kernel regression for sparse digitized picks (Tier–F). Denote aligned series by $x_{\text{src}}[k] \equiv x_{\text{src}}(t_k)$.
3. **Same-window synchronization.** Align series by metadata (angle/current/polarization logs). If residual offsets remain, estimate a small shift $\hat{\Delta} \in [-\Delta_{\text{max}}, \Delta_{\text{max}}]$ by maximizing cross-correlation $\sum_k x_a[k] x_b[k + \Delta]$ *under the constraint* that both series remain inside $[T_0, T_1]$.
4. **Masking and missing data.** Construct a binary mask $m[k] \in \{0, 1\}$ indicating samples valid in *all* aligned series. Compute statistics only on indices with $m[k] = 1$. Missing stretches longer than $5 \Delta t$ are not gap-filled unless explicitly declared.
5. **Effective variance.** For each aligned observable use

$$\sigma_{\text{eff}}^2(\tau) = \frac{\sigma_{\text{shot}}^2}{N} \oplus (2\pi f_0 \tau)^2 \sigma_y^2(\tau),$$

where N is the number of repeats within the window, f_0 is the reference frequency (or a sampling-rate proxy), and $\sigma_y(\tau)$ is the Allan deviation at averaging time τ . Unless noted, τ is fixed at the Allan minimum.

6. **Heteroskedastic weighting.** When forming means or regressions on $\{t_k\}$, use weights $w_k \propto \sigma_{\text{eff}}^{-2}(t_k)$. Report HC-robust (Huber–White) errors to guard against residual model mismatch ([Appendix F](#)).
7. **Provenance.** Store $[T_0, T_1], \Delta t, \hat{\Delta}, m[k]$, interpolation family and kernel order as part of the run metadata; record file hashes of all inputs.

Remark (lock-in equivalence). For periodically driven controls $X(t) = X_0 + X_1 s(t)$ on the timebase, the matched-filter estimate $\hat{\beta}$ used in [Section 4.4](#) is equivalent to a discrete lock-in on $\{t_k\}$ with weighting w_k .

F.6 Figure-derived inputs (policy). Figure-derived entries (Tier–F) are never primary. They are allowed only for *shape checks* and must obey the rules below so that confidence and bounds reflect Tier–N evidence.

1. **Variance floor and weighting.** For each Tier–F datum y_i with nominal uncertainty $\sigma_{\text{fig},i}$, enforce

$$\sigma_{\text{fig},i}^2 \leftarrow \sigma_{\text{fig},i}^2 \oplus \sigma_{\text{min}}^2, \quad \sigma_{\text{min}} \geq 2 \sigma_N,$$

where σ_N is the corresponding Tier–N uncertainty in the same window. In mixed displays, assign a down–weight $\omega_F \leq 0.25$ (default 0.20). If a panel offers only span–type graphics (min/max bars without reliable ticks), set $\sigma_{\min} = 3 \sigma_N$.

2. **Panel–internal covariance.** Points extracted from the *same* figure share calibration and reading errors. Inject a positive covariance

$$\text{Cov}_{ij}^{(\text{fig})} = \rho_{\text{fig}} \sigma_{\text{fig},i} \sigma_{\text{fig},j}, \quad \rho_{\text{fig}} \in [0.3, 0.7] \text{ (default 0.5)},$$

for any pair (i, j) from the same panel. Cross–panel correlations are zero unless the panels share axes/templates, in which case use $\rho_{\text{fig}} = 0.3$.

3. **Angle sensitivity and hysteresis.** When the observable depends on tilt θ , propagate a default angle uncertainty $\sigma_\theta = 0.1^\circ$ (unless stated otherwise) and add a hysteresis allowance if forward/backward tilts are not synchronized:

$$\sigma_{\text{hyst}} = \kappa_{\text{hyst}} \left| \partial R / \partial \theta \right| \Delta \theta_{\text{loop}}, \quad \kappa_{\text{hyst}} = 1, \quad \Delta \theta_{\text{loop}} \in [0.2^\circ, 0.5^\circ],$$

then fold σ_{hyst} into σ_{fig} before applying the floor.

4. **Strict exclusion from inference.** Tier–F entries are excluded from (i) confidence intervals (CIs), (ii) headline bounds, (iii) any propagation to (A, G) . They may be co–plotted with Tier–N for visual sanity checks only.
5. **Caption/legend requirements.** Any display including Tier–F must state *illustrative; excluded from inference* in the caption and list ω_F and $\sigma_{\min}$. Legends must visually separate tiers (e.g., hollow markers for Tier–F, solid for Tier–N).
6. **Display–only overlays.** If a line/band is shown over a mixed {Tier–N, Tier–F} cloud, its fit *must* be computed from Tier–N only. An optional faint “all–points (display–only)” overlay may be included but labeled non–inferential.
7. **Sensitivity sweeps.** When reporting a sensitivity to ρ_{fig} or $\sigma_{\min}$, sweep $\rho_{\text{fig}} \in \{0.3, 0.5, 0.7\}$, $\sigma_{\min} \in \{2, 3\} \times \sigma_N$ and confirm that any qualitative statement (e.g., monotonic trend) is unchanged.
8. **Audit log (provenance).** For each Tier–F panel, archive: (i) file hash, (ii) axis calibration tuple, (iii) raw click coordinates, (iv) transformed physical coordinates, (v) operator ID/seed, (vi) script and package hashes, and (vii) extraction timestamp. These enable re–evaluation without repeating digitization.
9. **Default policy summary.**
 - $\sigma_{\min} = 2 \sigma_N$ (or $3 \sigma_N$ for span–only panels), $\omega_F = 0.20$, $\rho_{\text{fig}} = 0.50$.
 - Tier–F excluded from CIs, bounds, and (A, G) propagation.
 - Captions include *illustrative; excluded from inference*, ω_F , $\sigma_{\min}$; legends separate tiers.
 - Angle/hysteresis terms included unless explicit same–window tilt synchronization is evidenced.

G. Metrology and stability log templates

This appendix consolidates the record formats for the resonator channel (TE/TM ratio tracking) and the tilt–projection verification (R–channel). The aim is to write the same items in the same way across runs so later analysis and reproduction proceed smoothly. Each run records the metadata below and the raw quantities required to form $\rho = f_{\text{TE}}/f_{\text{TM}}$ and R (frequencies, temperature, pressure, angle, timestamps).

Table 14: Metrology log template for long–term tracking of the TE/TM ratio ρ .

Item	Instrument/ Channel	Sampling	Stability (eval.)	Operating range	Notes
Frequency (TE)	Counter / common ref.	1 Hz	Allan $< 10^{-12}$ @ 10^3 s	$[f_0 \pm 1 \text{ kHz}]$	Reference lock
Frequency (TM)	Counter / common ref.	1 Hz	Same	Same	—
Temperature	PTR/Thermistor	0.1 Hz	± 1 mK	295–296 K	Near TK–zero
Vacuum	Hot/Cold cathode gauge	0.05 Hz	$< 10^{-5}$ mbar stable	$10^{-6} \sim 10^{-5}$ mbar	—
Clamping/strain	Strain gauge	0.05 Hz	Rel. drift $< 10^{-4}$	—	Long–term watch
Reference	GPSDO/OCXO	1 Hz	Drift $< 10^{-11}$ /day	—	Distribution log
Timestamp	NTP/PTP	—	< 1 ms	—	Run ID

ρ -channel window log (short + long; side-by-side). For each run, log both the *short* window (τ_{short}) and the *long* window (τ_{long}), the derivation path (comb–beat/linewidth vs. counter), timebase synchronization, and drift corrections (T/P/clamping). Store the Allan-deviation snapshot $\sigma_y(\tau)$, window boundaries, and any masked segments for reproducibility ([Appendix F, Section 4.4](#)).

Table 15: ρ -channel window log template (short and long windows recorded side-by-side).

Window class	τ (avg. time)	Derivation path	Sync / reference	$ \Delta\rho/\rho $ (95% CI)	Corrections / notes
Short	e.g., 30 s	Comb–beat/linewidth	Synchronized / common ref.	$\leq 1.6 \times 10^{-9}$ (example)	Same-window drift removed; attach comb spec / linewidth log
Long	e.g., 10^4 s	Counter (TE/TM)	Synchronized / common ref.	$\leq 3 \times 10^{-3}$ (example)	Near Allan minimum; include T/P /clamping corrections

Blind analysis & robust meta checklist. All items are finalized before unblinding; hashes/time are recorded for audit.

Table 16: Blind/robust–meta preregistration and audit template.

Item	Setting (frozen)	Hash / ID	Notes
ROI / cutoffs	Angle pairs, window bounds, QC flags	cfg hash	Ties & missing–data rules
Decision rules	Primary endpoint, stopping, masking	proto hash	Same across platforms
Blinding	Shuffled IDs, hidden labels	map hash	Unblind timestamp
Estimators	Huber/Tukey loss, HC errors	code hash	Tuning grid stored
Meta model	Fixed & random effects; Q, I^2	meta hash	DL/ τ^2 method
Sensitivity	$\omega_F, \rho_{\text{fig}}, \sigma_{\min}/\sigma_N$	grid ID	Predeclared ranges
LOPO	Platform–wise refits	run set ID	Max Δ recorded
Trim–and–fill	Filled k , adjusted effect	result ID	Display–only if used
Provenance	Data/script hashes, time	audit ID	Reviewer bundle path

Checklist (summary). Record: reference lock (GPSDO/OCXO), simultaneous TE/TM timestamps, drift-regression usage, $\sigma_y(\tau)$ snapshot, file/script hashes, per-window covariance handling rules ([Section 3.4](#)); **Sensitivity constants row used** (cf. [Appendix H](#)): [geometry / row ID].

Table 17: Experimental log template for the R–channel (tilt–projection verification).

Item	Device/Sample	Sampling	Alignment/Angle error	Operating range	Notes
Field sweep	Superconducting magnet	0.1–10 Hz	Linearity $\pm 0.1\%$	0–14 T	Hysteresis log
Angle θ	Rotator/Indexer	on change	$\pm 0.05^\circ$	$0^\circ \sim 90^\circ$	Reference plane defined
$\Delta B(\theta)$ extraction	Lock–in/FFT	per sweep	Peak–pick ± 1 digit	—	Window/filter recorded
Temperature	He cryostat/PPMS	0.1 Hz	± 10 mK	1.6–300 K	Ramp/ stabilization log
Wiring	4–probe/TWPA	on setup	—	—	Contact R, shielding
Timestamp	NTP/PTP	—	< 1 ms	—	Sync ID

Preregistration: angle-sweep grid and decision thresholds. **Grid.** $\Delta\theta \leq 1^\circ$; coverage $[\theta_{\min}, \theta_{\max}]$; repeats $N \geq 20$ per θ_k .
Bi-directional. Up/Down sweeps in the *same* window; identical dwell.

Model. $R(\theta) = 1 + \beta \cos 2\theta + \varepsilon$ (HC-robust SEs).

Pass/Fail. $R^2 \geq 0.95$; $|\hat{\beta}|/\text{SE}(\hat{\beta}) \geq 5$ (5σ); $H(\theta_k) \leq z_{0.995}\sigma_R$ (null-failure $\leq 0.5\%$).

Deviations. Log reason/correction/reacquisition in the G-templates.

Common fields and checklist. Record: run ID, sample ID, device ID, start/stop time, raw-data path (hash), preprocessing script version (hash), quality flags, calibration coefficients (frequency/field), 1σ uncertainty and 95% CI for ρ or R . Additionally confirm: timebase agreement and offset; calibration dates and factors; environmental stability (resonator: temperature drift ± 1 mK; R-channel: hysteresis log present); wiring/shielding snapshot; data integrity (file and script hashes); uncertainty budget including digitization/linearity/alignment; pre-registered criteria (amplitude $\geq 5\sigma$, linearity $R^2 \geq 0.95$, null-failure $\leq 1\%$).

H. Quick reference to conventions (non-normative)

This appendix is a compact pointer to the canonical conventions in Appendix A. For authoritative definitions, always consult Appendix A.

- **Signature, indices, forms, Hodge dual, HL $\leftrightarrow$ SI:** [Appendix A.1](#).
- **Geometry and gauge conventions (Stokes, holonomy):** [Appendix A.2](#).
- **Curvature field Φ and constitutive law overview:** [Appendix A.3](#).
- **Dimensional analysis & nondimensionalization:** [Appendix A.4](#).
- **Symbol table (recurring notation):** [Appendix A.5](#).
- **Fourier convention (phase/sign):** [Appendix A.6](#).
- **Averages, variance, spectra (one-sided PSD):** [Appendix A.7](#).

H.1 Conventions

Signature, indices, orientation. The metric signature is $(-, +, +, +)$. Greek indices $\mu, \nu, \rho, \sigma = 0, 1, 2, 3$ denote spacetime components; Latin indices $i, j, k = 1, 2, 3$ denote spatial components. Indices are raised/lowered with $g_{\mu\nu}$ ($V^\mu = g^{\mu\nu}V_\nu$, $V_\mu = g_{\mu\nu}V^\nu$). The totally antisymmetric symbols are fixed by $\epsilon^{0123} = +1$ and $\epsilon^{123} = +1$. Symmetrization/antisymmetrization use $T_{(\mu\nu)} \equiv \frac{1}{2}(T_{\mu\nu} + T_{\nu\mu})$, $T_{[\mu\nu]} \equiv \frac{1}{2}(T_{\mu\nu} - T_{\nu\mu})$.

Connection, curvature, dual. The $U(1)$ connection is the 1-form A_μ ; the field-strength (curvature) 2-form is

$$F_{\mu\nu} \equiv \partial_\mu A_\nu - \partial_\nu A_\mu \quad (\text{i.e. } F = dA).$$

The Hodge dual is

$$(*F)_{\mu\nu} \equiv \frac{1}{2} \epsilon_{\mu\nu\rho\sigma} F^{\rho\sigma}.$$

Algebraic invariants are $I_1 = \frac{1}{2} F_{\mu\nu} F^{\mu\nu}$ and $I_2 = \frac{1}{2} F_{\mu\nu} *F^{\mu\nu}$.

Derivatives and operators. The covariant derivative is ∇_μ . The d'Alembertian is

$$\square \equiv g^{\mu\nu} \nabla_\mu \nabla_\nu = -\partial_t^2 + \nabla^2 \quad \text{for signature } (-, +, +, +).$$

Inner products and norms use $X \cdot Y \equiv g_{\mu\nu} X^\mu Y^\nu$ and $|X|^2 \equiv X \cdot X$.

3+1 split. With 4-potential $A_\mu = (-\phi, \mathbf{A})$,

$$E_i \equiv F_{0i}, \quad B^i \equiv \frac{1}{2} \epsilon^{ijk} F_{jk},$$

so that $\mathbf{E} = -\partial_t \mathbf{A} - \nabla \phi$ and $\mathbf{B} = \nabla \times \mathbf{A}$.

Units and rationalization. Heaviside-Lorentz rationalized units are used with $c = 1$. In this convention, the vacuum Maxwell equations read

$$\nabla \cdot \mathbf{E} = \rho, \quad \nabla \times \mathbf{B} - \partial_t \mathbf{E} = \mathbf{J}, \quad \nabla \cdot \mathbf{B} = 0, \quad \nabla \times \mathbf{E} + \partial_t \mathbf{B} = 0.$$

Restoration rules for $\hbar$, ε_0 , μ_0 , c follow [Appendix H.4](#) when SI presentation is required.

Fourier convention (pointer). For the global phase/sign convention and spectral normalization, see [Appendix A.6](#).

Averages and spectral densities (pointer). For averages, variance/covariance, and one-sided PSD normalization, see [Appendix A.7](#).

H.2 Symbols

The table lists symbols that recur across the manuscript. Dimensions follow the base (L, T, Q) in Heaviside–Lorentz units with $c = 1$. Short operational notes clarify how each quantity is used.

Symbol	Dim.	Meaning / operational note
$\Phi(x)$	1	Curvature scalar field (taken dimensionless by convention; any intrinsic scale is absorbed into couplings).
$\nabla_\mu \Phi$	L^{-1}	Gradient of Φ ; the “flat–gradient” regime assumes $ \nabla \Phi < \varepsilon$ for a small fixed ε .
$H_{\mu\nu} = \nabla_\mu \nabla_\nu \Phi$	L^{-2}	Hessian of Φ ; trace-removed shear $S_{\mu\nu} = H_{\mu\nu} - \frac{1}{4}g_{\mu\nu}\square\Phi$.
$u(x)$	1	Complex unit section selecting the principal plane $\Pi(x)$; fixes the $U(1)$ phase frame.
$A_\mu = \text{Im} \frac{u^\dagger \nabla_\mu u}{u^\dagger u}$	L^{-1}	Berry-like $U(1)$ connection from u ; a gauge shift $u \rightarrow e^{i\varphi}u$ sends $A \rightarrow A + \partial\varphi$.
$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$	L^{-2}	Field-strength (curvature) 2-form; in 3+1 form $E_i = F_{0i}$, $B^i = \frac{1}{2}\epsilon^{ijk}F_{jk}$.
$I_1 = \frac{1}{2}F_{\mu\nu}F^{\mu\nu}$	L^{-4}	Lorentz invariant $I_1 = \mathbf{B}^2 - \mathbf{E}^2$.
$I_2 = \frac{1}{2}F_{\mu\nu}{}^*F^{\mu\nu}$	L^{-4}	Lorentz invariant $I_2 = \mathbf{E} \cdot \mathbf{B}$ (parity-odd).
$\chi(\Phi, \partial\Phi)$	1	Constitutive map (dimensionless); excitation $H = \chi : F$ with isotropic core and controlled anisotropy.
$\rho = f_{\text{TE}}/f_{\text{TM}}$	1	Resonator frequency ratio (TE vs TM); separates isotropic vs anisotropic shifts.
$R = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2}$	1	Tilt-normalized amplitude ratio; Maxwell-reduction limit gives $R \rightarrow 1$.
α, η	1	Couplings (isotropic α , anisotropic η); enter χ via $1 + \alpha\Phi$ and gradient terms.
$G = \eta \nabla \Phi $	1	Anisotropy strength (scalar control parameter built from η and the local gradient).

H.3 Constants

Heaviside–Lorentz rationalization with $c = 1$ is assumed throughout. Planck’s constant $\hbar$ appears only when explicit phase bookkeeping is required (e.g., loop phases, quantization). When restoration to SI is needed, reinsert $(\varepsilon_0, \mu_0, c, \hbar)$ and use the vacuum impedance $Z_0 = \sqrt{\mu_0/\varepsilon_0}$, while preserving the invariant structure and operator definitions set in [Appendix H.1](#).

H.4 Dimensional analysis examples

This appendix uses Heaviside–Lorentz units with $c = 1$. The base dimensions are length (L), time (T), and charge (Q). Planck’s constant $\hbar$ is introduced only when explicit phase bookkeeping is required.

(1) Connection and curvature. The connection A_μ is a phase gradient, hence $[A_\mu] = L^{-1}$. Therefore

$$[F_{\mu\nu}] = [\partial_\mu A_\nu] = L^{-2}.$$

With the 3+1 split $E_i = F_{0i}$ and $B^i = \frac{1}{2}\epsilon^{ijk}F_{jk}$, one has $[E] = [B] = L^{-2}$.

(2) Lorentz invariants.

$$I_1 = \frac{1}{2}F_{\mu\nu}F^{\mu\nu} \Rightarrow [I_1] = L^{-4}, \quad I_2 = \frac{1}{2}F_{\mu\nu}{}^*F^{\mu\nu} \Rightarrow [I_2] = L^{-4}.$$

Both invariants scale like an energy density, L^{-4} .

(3) Constitutive law $H = \chi(\Phi, \partial\Phi) : F$. Since $H_{\mu\nu}$ must have the same dimension as $F_{\mu\nu}$, the constitutive map is dimensionless:

$$[H_{\mu\nu}] = [F_{\mu\nu}] = L^{-2}, \quad [\chi] = 1.$$

For mean/gradient parametrizations (e.g., isotropic $1 + \alpha\Phi$, anisotropic $G = \eta|\nabla\Phi|$),

$$[\alpha\Phi] = 1, \quad [\eta|\nabla\Phi|] = 1.$$

Thus $[\Phi]$ may be absorbed into the definition of α (rendering Φ dimensionless), or α may be assigned the compensating dimension. With the constraint $|\nabla\Phi| < \varepsilon$, one has $[\varepsilon] = [\nabla\Phi] = L^{-1}[\Phi]$.

(4) Action and Lagrangians. The Maxwell term $\mathcal{L}_F = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$ is a density, hence

$$[\mathcal{L}_F] = L^{-4}.$$

For a scalar field (e.g., $\mathcal{L}_\Phi = \frac{1}{2}\partial_\mu\Phi\partial^\mu\Phi - U(\Phi)$),

$$[\partial\Phi]^2 \sim L^{-2}[\Phi]^2 \Rightarrow [U(\Phi)] = L^{-4}.$$

If the source satisfies $\square\Phi - U'(\Phi) = J$, then

$$[\square\Phi] = L^{-2}[\Phi], \quad [U'(\Phi)] = L^{-4}[\Phi]^{-1}, \quad [J] = L^{-2}[\Phi].$$

Choosing $[\Phi] = 1$ gives $[J] = L^{-2}$, which is convenient.

(5) Dimensionless observables. For the tilt–normalized ratio used in gradient–anisotropy tests,

$$R = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2}$$

the dimensions cancel, so R is dimensionless and approaches $R \rightarrow 1$ in the Maxwell–reduction limit. The resonator mode ratio $\rho = f_{\text{TE}}/f_{\text{TM}}$ is likewise dimensionless and separates isotropic from anisotropic shifts.

(6) Spectral convention and unit check. With the Fourier convention in [Appendix H.1](#), the one-sided spectral density $S_{ff}(\omega)$ is normalized by

$$\int \frac{d\omega}{2\pi} S_{ff}(\omega) = \langle f^2 \rangle.$$

For $f = E$, this gives $[S_{EE}] = [E]^2[\omega]^{-1} = \text{L}^{-4} \text{T}$, consistent with power (or variance) restoration.

(7) Restoring constants (SI). When SI presentation is required, reinsert $(c, \varepsilon_0, \mu_0, \hbar)$. For example,

$$[E]_{\text{SI}} = \text{V/m}, \quad [B]_{\text{SI}} = \text{T}, \quad Z_0 = \sqrt{\mu_0/\varepsilon_0} \text{ (vacuum impedance)}.$$

Quantum phases (e.g., loop phases) bring in $\hbar$ explicitly; the phase $\varphi = \oint A_\mu dx^\mu / \hbar$ is dimensionless.

H.5 Publication-ready (A, G) confidence ellipse

This subsection completes the [Section 3.4](#) promise by providing a ready-to-insert figure and the exact recipe for the (A, G) confidence ellipse. Here A denotes the isotropic coupling, and G follows [Appendix H.2](#) as $G = \eta |\nabla \Phi|$ (both dimensionless).

Definition and notation. Let the parameter vector be $\boldsymbol{\theta} = (A, G)^\top$ with estimate $\hat{\boldsymbol{\theta}} = (\hat{A}, \hat{G})^\top$ and covariance $\boldsymbol{\Sigma} = \text{Cov}(\hat{\boldsymbol{\theta}})$. Under the Gaussian approximation the constant-likelihood contour (ellipse) at confidence level CL is

$$(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}) = \Delta\chi_{p=2}^2(\text{CL}), \quad (85)$$

with the standard $\Delta\chi^2$ values for $p = 2$ summarized in [Table 18](#).

Practical computation notes. With the eigendecomposition $\boldsymbol{\Sigma} = \mathbf{Q} \text{diag}(\lambda_1, \lambda_2) \mathbf{Q}^\top$, the semi-axes are $\sqrt{\lambda_{1,2} \Delta\chi^2}$, and the ellipse is rotated by the columns of $\mathbf{Q}$. For reporting, it is often convenient to use the standard-error form

$$\boldsymbol{\Sigma} = \begin{pmatrix} \sigma_A^2 & \rho_{AG} \sigma_A \sigma_G \\ \rho_{AG} \sigma_A \sigma_G & \sigma_G^2 \end{pmatrix},$$

where σ_A, σ_G are marginal uncertainties and ρ_{AG} is the correlation.

Table 18: Confidence level (CL) and corresponding $\Delta\chi^2$ for a two-parameter ($p=2$) Gaussian approximation.

CL (%)	$\Delta\chi^2$	Note
68.3	2.30	1σ equivalent
95.0	5.99	Baseline reporting level
99.0	9.21	Conservative line (optional)




Figure 9: (A, G) confidence ellipses centered at $(\hat{A}, \hat{G})$. The principal axes and their orientation follow the eigenvalues/eigenvectors of Σ . Contours correspond to $\text{CL} \in \{68.3\%, 95\%, 99\%\}$ with $\Delta\chi^2 = \{2.30, 5.99, 9.21\}$. Axes are dimensionless. Numeric contour coordinates and a small template bundle are provided in [Appendix H.6](#).

Reproducibility note. The exact polygonal coordinates for the contours in Fig. 9, together with $\hat{\theta}$ and Σ , are packaged in the template bundle cited in Appendix H.6 for direct reuse.

H.6 Template bundle for (A, G) contours

Contents. A compact auxiliary bundle accompanies Fig. 9 and Appendix H.5. It contains a publication-ready vector figure of the (A, G) confidence ellipses, a table of polygonal contour coordinates for the confidence levels listed in Table 18, and a metadata table specifying the center $(\hat{A}, \hat{G})$ and the covariance entries $(\sigma_A, \sigma_G, \rho_{AG})$. No filenames are cited here; the bundle is referenced solely by this subsection.

Schemas. The contour table provides, for each confidence level, ordered vertex pairs (A_k, G_k) together with the corresponding $\Delta\chi^2$. The metadata table reports $\hat{A}$, $\hat{G}$, the marginal uncertainties σ_A , σ_G , the correlation ρ_{AG} , and the covariance components Σ_{aa} , Σ_{ag} , Σ_{gg} , consistent with the standard-error form of Σ summarized in Appendix H.5.

Use. Contours may be drawn by connecting the ordered vertices for each confidence level; the point $(\hat{A}, \hat{G})$ marks the center. Axes are dimensionless. The same bundle enables independent numerical checks of Fig. 9 without further assumptions.

H.7 SI $\leftrightarrow$ Heaviside–Lorentz (HL) unit restoration cheatsheet

HL units use $c = 1$ and rationalized Maxwell equations $\nabla \cdot \mathbf{E} = \rho$, $\nabla \times \mathbf{B} - \partial_t \mathbf{E} = \mathbf{J}$, $\nabla \cdot \mathbf{B} = 0$, $\nabla \times \mathbf{E} + \partial_t \mathbf{B} = 0$. In SI, the vacuum equations read

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon_0}, \quad \nabla \times \mathbf{B} - \frac{1}{c^2} \partial_t \mathbf{E} = \mu_0 \mathbf{J}, \quad \nabla \cdot \mathbf{B} = 0, \quad \nabla \times \mathbf{E} + \partial_t \mathbf{B} = 0,$$

with $c^{-2} = \varepsilon_0 \mu_0$ and the vacuum impedance $Z_0 = \sqrt{\mu_0 / \varepsilon_0}$.

Quantity	HL baseline ($c=1$)	SI restoration (insert constants)
Maxwell–Gauss	$\nabla \cdot \mathbf{E} = \rho$	$\nabla \cdot \mathbf{E} = \rho / \varepsilon_0$
Maxwell–Ampère	$\nabla \times \mathbf{B} - \partial_t \mathbf{E} = \mathbf{J}$	$\nabla \times \mathbf{B} - \frac{1}{c^2} \partial_t \mathbf{E} = \mu_0 \mathbf{J}$
Units of E, B	$[E] = [B] = \text{L}^{-2}$ (same units)	$[E]_{\text{SI}} = \text{V/m}$, $[B]_{\text{SI}} = \text{T}$; reference $Z_0 = \sqrt{\mu_0 / \varepsilon_0}$
Charge $\rho, \mathbf{J}$	sources enter without ε_0, μ_0	ρ couples with $1/\varepsilon_0$, $\mathbf{J}$ with μ_0
Flux quantum	Φ_0 appears via phase only	$\Phi_0 = \frac{h}{e}$ (SI Weber)
Action terms	$\mathcal{L}_F = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu}$	$\mathcal{L}_F = -\frac{1}{4\mu_0} F_{\mu\nu} F^{\mu\nu}$ (vacuum)
Wave speed	set by $c = 1$	dispersion restored by $c = (\varepsilon_0 \mu_0)^{-1/2}$

Practical tip. To convert derivations written in HL, first reinsert ε_0, μ_0, c at the equations level as above; then map units if needed. This avoids ad-hoc scale factors for $\mathbf{E}, \mathbf{B}$ and keeps invariants I_1, I_2 structurally unchanged.

H.8 First-order error propagation for R and ρ

Tilt-normalized ratio R . With

$$R = \frac{\Delta B(\theta_2) \cos \theta_1}{\Delta B(\theta_1) \cos \theta_2} \equiv \frac{X_2}{X_1} \frac{C_1}{C_2}, \quad X_i \equiv \Delta B(\theta_i), \quad C_i \equiv \cos \theta_i,$$

the fractional variance to first order (allowing correlations) is

$$\begin{aligned} \left(\frac{\sigma_R}{R}\right)^2 &\simeq \left(\frac{\sigma_{X_2}}{X_2}\right)^2 + \left(\frac{\sigma_{X_1}}{X_1}\right)^2 + \tan^2 \theta_1 \sigma_{\theta_1}^2 + \tan^2 \theta_2 \sigma_{\theta_2}^2 \\ &\quad - 2 \operatorname{Cov}(\ln X_2, \ln X_1) - 2 \operatorname{Cov}(\tan \theta_1 \delta \theta_1, \tan \theta_2 \delta \theta_2). \end{aligned}$$

If correlations are negligible,

$$\boxed{\left(\frac{\sigma_R}{R}\right)^2 \approx \left(\frac{\sigma_{\Delta B(\theta_2)}}{\Delta B(\theta_2)}\right)^2 + \left(\frac{\sigma_{\Delta B(\theta_1)}}{\Delta B(\theta_1)}\right)^2 + \tan^2 \theta_1 \sigma_{\theta_1}^2 + \tan^2 \theta_2 \sigma_{\theta_2}^2.}$$

Resonator mode ratio ρ . Let $\rho = f_{\text{TE}}/f_{\text{TM}}$. Working in logs simplifies propagation:

$$\Delta \ln \rho = \ln f_{\text{TE}} - \ln f_{\text{TM}}, \quad \operatorname{Var}(\Delta \ln \rho) \approx \left(\frac{\sigma_{f_{\text{TE}}}}{f_{\text{TE}}}\right)^2 + \left(\frac{\sigma_{f_{\text{TM}}}}{f_{\text{TM}}}\right)^2 - 2 \rho_{\text{TE, TM}} \frac{\sigma_{f_{\text{TE}}}}{f_{\text{TE}}} \frac{\sigma_{f_{\text{TM}}}}{f_{\text{TM}}}.$$

For uncorrelated modes, the cross term vanishes and $\sigma_{\Delta \ln \rho} = \sqrt{\cdots}$. If needed, σ_ρ follows from $\sigma_{\Delta \ln \rho}$ via $\sigma_\rho \approx \rho \sigma_{\Delta \ln \rho}$.

Digitization rule of thumb. When values are read from published axes, adopt a conservative per-point digitization uncertainty

$$\sigma_{\text{dig}} \approx \frac{\text{tick spacing}}{5},$$

apply it in quadrature with quoted statistical errors, and propagate through the above formulas. This rule stabilizes cross-paper consistency and matches the conservative practice used in the data sections.

I. Numerical validation: convergence curves and mesh sweeps (extended)

This appendix complements [Section 5](#) with *quantitative* convergence plots and tables. We collect, in one place, the evidence that the structure-preserving scheme exhibits the expected orders in space and time, and that conserved quantities follow the predicted scaling laws. To keep the main text uncluttered, figures and tables are placed here and follow the no-vertical-lines booktabs style with concise captions. When useful, we cross-reference construction details in [Appendix H.6](#) and geometry/polarization links in [Section 4.5](#).

I.1 Spatial-resolution convergence

We summarize how the representative error metric $e(h)$ scales with the mesh spacing h . For the target order p , we seek evidence that $e(h) \sim C h^p$. Unless stated otherwise, the reference error is the relative L^2 norm over the computational domain, or a problem-appropriate invariant deviation (see [Section 5.6](#)).

Error definition and reference. Given a sequence of meshes with spacings $\{h_\ell\}$ ($h_{\ell+1} = h_\ell/2$ in the canonical halving case), define

$$e(h_\ell) = \frac{\|u_{h_\ell} - u_{\text{ref}}\|_{L^2(\Omega)}}{\|u_{\text{ref}}\|_{L^2(\Omega)}},$$

where u_{ref} is either the finest-grid solution or a Richardson-extrapolated surrogate. If boundary layers or geometric singularities are present, we additionally report subdomain rates on $\Omega = \Omega_{\text{bulk}} \cup \Omega_{\text{BL}}$.

Grid	h	$e(h)$	Rate $p(h)$	Notes
G1	h_0	e_0	—	baseline mesh; CFL matched to $\Delta t \propto h$
G2	$h_0/2$	e_1	p_1	same geometry/boundary; implicit midpoint in time
G3	$h_0/4$	e_2	p_2	dispersion check passes ($\sim h^2$)
G4	$h_0/8$	e_3	p_3	energy/continuity residuals monotone

Here the per-level rate is

$$p(h) = \log_2 \left(\frac{e(h)}{e(h/2)} \right),$$

and for non-dyadic refinement ratio r use $p(h; r) = \frac{\log(e(h)/e(h/r))}{\log r}$.

Checklist.

- Keep time error subdominant by choosing $\Delta t = \mathcal{O}(h)$ (or $\Delta t = \mathcal{O}(h/c)$); see [Section 5.4](#).

- Prefer Richardson extrapolation for u_{ref} ; otherwise use the finest grid and note it in *Notes*.
- Enforce monotonicity $e(h_0) > e(h_0/2) > e(h_0/4)$; if violated, check boundary enforcement and Hodge conditioning (Section 5.2).

Interpretation. For the DEC + implicit–midpoint scheme, the expected spatial order in smooth regimes is $p \simeq 2$ (with matched time refinement). Observed slopes $p \geq 1.9$ on the last two levels meet the acceptance criterion of Section 5.6. If $p < 1.8$, investigate mesh distortion, boundary enforcement (PEC/PMC/PML via Hodge restriction), and the conditioning of the discrete Hodge operator.

I.2 Time-step convergence and stability window

We assess *temporal order* and the *stability window* by refining the time step Δt while keeping the spatial grid fixed. The implicit midpoint integrator used in this work is A-stable for the linear Maxwell limit and attains $\mathcal{O}(\Delta t^2)$ accuracy in time for time-invariant, linear Hodge operators. With weak nonlinearity or time-dependent $\chi(\Phi, t)$, and with absorbing layers (PML), we still observe near-quadratic convergence under normal operating settings; for overly large Δt , phase/dispersion error may temporarily degrade the measured rate (see Section 5.4, Section 5.6).

Error definition and observed order. On a fixed mesh (so that spatial error remains subdominant), define

$$e_t(\Delta t) = \frac{\|u_{\Delta t} - u_{\text{ref}}\|_{L^2(\Omega)}}{\|u_{\text{ref}}\|_{L^2(\Omega)}}, \quad q(\Delta t) = \frac{\log(e_t(\Delta t)/e_t(\Delta t/r))}{\log r},$$

with the canonical refinement ratio $r = 2$ so that $q(\Delta t) = \log_2(e_t(\Delta t)/e_t(\Delta t/2))$. The reference u_{ref} is either a Richardson-extrapolated surrogate or the solution at the finest Δt .

Step	Δt	$e_t(\Delta t)$	Rate $q(\Delta t)$	Stable?	Notes
T1	Δt_0	$e_{t,0}$	—	Yes	Baseline step; choose $\Delta t_0 \propto h$ so that spatial error does not dominate.
T2	$\Delta t_0/2$	$e_{t,1}$	$q_1 = \log_2(e_{t,0}/e_{t,1})$	Yes	Expected $q_1 \approx 2$; energy balance residual scales as $\mathcal{O}(\Delta t^3)$.
T3	$\Delta t_0/4$	$e_{t,2}$	$q_2 = \log_2(e_{t,1}/e_{t,2})$	Yes	As $\Delta t \downarrow$, temporal error decreases; continuity residual stays at machine precision.
T4	$2 \Delta t_0$	$e_{t,\uparrow}$	n/a	Cond.	Oversized step: phase/dispersion error inflates; PML and nonlinear coupling may reduce robustness.

Stability window (practical guidance).

- **Accuracy-oriented choice.** For spatial-order studies, keep $\Delta t = \mathcal{O}(h)$ so temporal error remains subdominant. Although implicit midpoint has no CFL restriction in the linear case, for accuracy and dispersion we recommend $\Delta t \lesssim \kappa h/c$ with $\kappa \sim 1$.
- **Nonlinear/time-dependent media.** With $\chi(\Phi, t)$ or $\partial_t \chi \neq 0$, use a fixed-point outer loop plus a small Newton correction; verify that halving Δt restores $q \approx 2$ (see [Section 5.4](#)).
- **PML and boundaries.** With PML active, large Δt can increase numerical reflections and the energy residual. Reduce Δt or refine the PML profile to widen the stable window.
- **Monitoring.** At each step record $\mathcal{R}_{\text{cont}} = \|\Delta_t \rho + \delta J\|_2$ and $\mathcal{R}_E$; in normal regimes $\mathcal{R}_{\text{cont}}$ stays at machine precision, while $\mathcal{R}_E = \mathcal{O}(\Delta t^3)$ ([Section 5.3](#), [Section 5.4](#)).

Summary interpretation. If $q_1, q_2 \geq 1.9$ along the refinement ladder $\{\Delta t_0, \Delta t_0/2, \Delta t_0/4\}$, the *second-order temporal accuracy* acceptance criterion of [Section 5.6](#) is met. Persistent $q < 1.8$ suggests checking (i) the reference solution choice, (ii) boundary/PML implementation, (iii) fixed-point/Newton convergence under nonlinear coupling, and (iv) whether spatial error has become dominant.

I.3 Conserved-quantity and residual scaling

We summarize the *quantitative residuals* for continuity and energy (or action) conservation as functions of spatial resolution h and time step Δt . In an ideal structure-preserving scheme, residuals converge to machine precision; when nonconservative effects are present (boundary flux, PML, time-dependent media, etc.), they typically follow $\mathcal{O}(h^p) + \mathcal{O}(\Delta t^q)$ scaling ([Section 5.3](#), [Section 5.4](#), [Section 5.6](#)).

Residual definitions (reporting standard).

$$\mathcal{R}_{\text{cont}}^n \equiv \|\Delta_t \rho^n + \delta J^n\|_{L^2(\Omega)}, \quad \mathcal{R}_E^n \equiv \left| \mathcal{E}^{n+\frac{1}{2}} - \mathcal{E}^{n-\frac{1}{2}} + \Delta t \langle J^n, E^{n+\frac{1}{2}} \rangle - \text{boundary flux} \right|.$$

In source-free, closed domains (no boundary flux), $\mathcal{R}_{\text{cont}}^n$ should remain at *machine precision*, while $\mathcal{R}_E^n = \mathcal{O}(\Delta t^3)$ (implicit midpoint).

Config	Residual	Measured scaling	Limit	Pass?	Notes
C1	Continuity $\mathcal{R}_{\text{cont}}$	$\lesssim 10^{-12}$ (flat vs. $h, \Delta t$)	$\rightarrow 0$	Yes	No sources, closed domain. Structural preservation ($\text{DEC} + \delta\delta = 0$) keeps $\mathcal{R}_{\text{cont}}$ at machine precision.
C2	Energy $\mathcal{R}_E$	$\propto \Delta t^3$	$\rightarrow 0$	Yes	Linear, time-invariant χ : implicit-midpoint energy balance matches theory (see Section 5.4).
C3	Continuity $\mathcal{R}_{\text{cont}}$	$\propto h^p + \Delta t^q$	$\rightarrow 0$	Cond.	Nonconserving test deposition increases $\mathcal{R}_{\text{cont}}$; switch to conserving (path-split) deposition to restore machine precision (Section 5.3).
C4	Energy $\mathcal{R}_E$	$\propto \Delta t^3 + \varepsilon_{\text{PML}}$	$\rightarrow \varepsilon_{\text{PML}}$	Cond.	With PML, the residual floors at a profile-dependent level ε_{PML} ; refine the PML to reduce the floor.
C5	Mixed (aggregate)	state error $\propto h^2 + \Delta t^2$, $\mathcal{R}_E \propto \Delta t^3$	$\rightarrow 0$	Yes	Plane-wave / MMS: 2nd-order state convergence and 3rd-order energy-residual scaling confirmed (see Section 5.6).

Notes and practical tips.

- **Same-window reporting:** Perform h -sweeps and Δt -sweeps separately, but report residuals from the *same window* (same $h, \Delta t$) as max/mean values in the table.
- **Monotonicity check:** Under refinement, $\mathcal{R}_E$ should fall as $\propto \Delta t^3$ on a log–log plot, while the state error decreases as $\propto h^2 + \Delta t^2$. If monotonicity fails, first inspect boundary implementation, PML settings, and time synchronization of $\Phi^{n+\frac{1}{2}}$.
- **Conserving deposition:** Use conserving (path–split) current deposition as default to keep $\mathcal{R}_{\text{cont}}$ at machine precision ([Section 5.3](#)).
- **Reporting standard:** Include fitted log–log slopes (last 2–3 points) beneath each plot, and state explicit QA thresholds (e.g., $\max_n \mathcal{R}_{\text{cont}}^n \leq 10^{-10}$).

I.4 Operator condition numbers and Hodge consistency

This subsection summarizes the *condition numbers* (*cond.#*) of discrete Hodge operators and derived operators as functions of mesh resolution h , domain scale L , and mesh

quality (aspect ratio AR, skew/distortion). When condition numbers grow unnecessarily large, linear solver convergence deteriorates; scaling laws then motivate *preconditioning* and/or *local rescaling* ([Section 5.2](#), [Section 5.6](#)).

Reporting standard and interpretation. Each Hodge $\star_k$ must be a *symmetric positive definite (SPD)* mass matrix and remain SPD after an isotropic/anisotropic split ($\star_k = \star_{k,\text{iso}} + \star_{k,\text{aniso}}$). On an ideal regular mesh, $\kappa(\star_k) = \mathcal{O}(1)$, while curl–curl and grad–div type operators typically exhibit $\kappa \sim \mathcal{O}(h^{-2})$ scaling. As AR and distortion increase, κ is amplified geometrically (roughly $\propto \text{AR}^2$); local lumping and block–diagonal preconditioners are recommended.

Grid	Operator	Cond.# trend	Notes
G1 (AR ≈ 1)	Hodge $\star_1, \star_2$ (iso)	$\kappa \approx 1.2\text{--}3.0$ (flat in h)	Near-diagonal (lumped); SPD holds. Isotropic limit: $\kappa = \mathcal{O}(1)$.
G2 (regular)	curl–curl (1-forms)	$\kappa \propto h^{-2}$	Matches dispersion tests; PCG+AMG converges well. Midpoint coupling stabilizes energy residual (Section 5.4).
G3 (AR ≈ 4)	Hodge $\star_1, \star_2$ (with anisotropy)	$\kappa \propto \text{AR}^2$	Geometric-weight imbalance; use block lumping / cell-local normalization; recheck SPD with anisotropic part.
G4 (AR $\gtrsim 6$, skew)	curl–curl / grad–div (mixed)	$\kappa \propto h^{-2}\text{AR}^2$	With PEC/PMC constraints conditioning worsens; use mixed (Hiptmair-type) preconditioners or AMG; record boundary Hodge restriction (Appendix D).
G5 (regular)	Mixed block $\begin{bmatrix} \star_2 & 0 \\ 0 & \star_1 \end{bmatrix}$	$\kappa \approx \max\{\kappa(\star_1), \kappa(\star_2)\}$	Block-diagonal preconditioner effective; freeze $\star_k^{(\Phi)}$ at outer fixed point, then PCG.
G6 (regular)	$\star_k^{(\Phi)}$ (time-dependent)	stable (midpoint eval.)	Evaluate at $\Phi^{n+1/2}$ to reduce drift; a single Newton correction after the fixed-point loop sharply decreases residuals (Section 5.2).

Recommendations and checklist.

- **SPD verification:** Numerically confirm $\star_k \succ 0$ (minimum eigenvalue > 0) at each mesh level; record that SPD is retained after adding the anisotropic component.
- **Condition–number control:** Apply local scaling (length/area/volume) and diagonal lumping in regions with large AR or distortion. Use PCG+AMG for curl–curl; use block preconditioning for mixed operators.
- **Temporal consistency:** Always evaluate $\star_k^{(\Phi)}$ at the midpoint time $(\Phi^{n+1/2})$ to minimize condition–number fluctuation across iterations ([Section 5.4](#)).
- **Logging standard:** Report κ vs. h , AR–sweep results (log–log slopes), and iteration counts/residual–reduction rates with/without preconditioning using the [Appendix G](#) format.

I.5 Mesh sweeps (geometry and boundary conditions)

We compare how mesh quality (aspect ratio, skew) and boundary conditions (PEC/PMC/periodic/PML) affect the state error and the conservation residuals. Metrics are unified as ε_{L^2} (relative L^2 state error), $\mathcal{R}_{\text{cont}}$ (continuity residual), and $\mathcal{R}_E$ (energy residual), following the reporting format of [Section 5.6](#) and the logging template in [Appendix G](#).

Mesh	Quality (AR, skew)	BC	ε_{L^2}	$\mathcal{R}_{\text{cont}}, \mathcal{R}_E$	Notes
M1	(1.5, 0.05)	PEC	$\approx 1.5 \times 10^{-3}$	$\lesssim 10^{-12}, \mathcal{O}(\Delta t^3)$	Quasi-uniform grid. Lumped $\star_k$ is well conditioned (SPD). Baseline case; spatial and temporal order $\simeq 2$ confirmed.
M2	(2.0, 0.10)	PMC	$\approx 2.0 \times 10^{-3}$	$\lesssim 10^{-12}, \mathcal{O}(\Delta t^3)$	Only boundary constraint changes (magnetic wall). Slightly higher curl–curl conditioning; PCG+AMG iteration count increases by $\sim 10\%$.
M3	(3.5, 0.20)	Periodic	$\approx 2.8 \times 10^{-3}$	$\lesssim 10^{-12}, \mathcal{O}(\Delta t^3)$	Convenient for dispersion checks. AR growth mildly enlarges dispersion error; second-order slope recovered under mesh refinement.
M4	(5.0, 0.30)	PEC	$\approx 6.0 \times 10^{-3}$	$\lesssim 10^{-11}, \mathcal{O}(\Delta t^3)$	Distortion/AR raise $\kappa(\text{curl} - \text{curl}) \propto h^{-2} \text{AR}^2$. Block-diagonal preconditioning (electric/magnetic split) restores fast convergence.
M5	(2.0, 0.10)	PML	$\approx 1.8 \times 10^{-3}$	$\lesssim 10^{-12}, \rightarrow \varepsilon_{\text{PML}}$	With absorbing layers, $\mathcal{R}_E$ saturates at a profile-dependent floor ε_{PML} . Increase PML thickness or refine the profile to lower the floor (Section 5.6).
M6	(6.5, 0.40)	PML	$\approx 1.2 \times 10^{-2}$	$\lesssim 10^{-10}, \rightarrow \varepsilon_{\text{PML}}$	Severe distortion + PML. Apply cell-local normalization (length/area/volume) and mixed (grad–div aided) preconditioners to stabilize iterations.

Practical guidance.

- **Maintain order:** On the last two refinement levels, target $p \simeq 2$ (space) and $q \simeq 2$ (time), as in [Section 5.6](#).
- **Conservation residuals:** Use conserving (path-split) current deposition so that $\mathcal{R}_{\text{cont}}$ stays at machine precision ([Section 5.3](#)); for implicit midpoint, $\mathcal{R}_E = \mathcal{O}(\Delta t^3)$ in the linear, time-invariant case ([Section 5.4](#)).
- **Conditioning:** As AR/distortion grows, default to lumping + block-diagonal/AMG preconditioning; with PML, always report the observed ε_{PML} floor and tune layer thickness/profile accordingly.

J. Cross-check with external constraints (optics, cavity, cosmology, microwave)

This appendix compiles *published, peer-reviewed* 95% limits from rotating optical cavities (modern Michelson–Morley), laboratory birefringence/rotation, cosmic birefringence, and microwave/WGM resonators, and maps them onto the common (A, G) scale used in the main text (Section 3.4). Each entry reports (i) the *Published limit* as given by the paper and (ii) a conservative *Implied bound* on A or G via the first-order maps of Section 2.6 and the sensitivity rows in Appendix E. Hyperlinks point to DOI or journal pages.

Notation. Published limits are copied verbatim (including units). Implied bounds use one-parameter projections ($A=0$ or $G=0$) unless a covariance is available. When a range of sensitivities exists, we adopt *conservative* $\mathcal{O}(1)$ factors and annotate the row ID from Appendix E.

Table 19: Modern Michelson–Morley (rotating optical cavities) and related optical tests: reported limits and conservative (A, G) mapping.

Reference / link	Platform	Published limit (95%)	Implied bound	Row ID (App. E)
Nagel et al., Nat. Commun. 6, 8174 (2015)	Rotating optical cavity (dual)	$\Delta\nu/\nu = (9.2 \pm 10.7) \times 10^{-19}$ (95% CI)	$ G \lesssim \mathcal{O}(10^{-18})$	E-MM-opt-1
Eisele et al., Phys. Rev. Lett. 103, 090401 (2009)	Rotating optical cavity	$\Delta c/c \sim 10^{-17}$ -level	$ G \lesssim \mathcal{O}(10^{-17})$	E-MM-opt-2
Herrmann et al., Phys. Rev. D 80, 105011 (2009)	Rotating optical cavity	$\Delta c/c \sim 10^{-17}$ -level	$ G \lesssim \mathcal{O}(10^{-17})$	E-MM-opt-3

Table 20: Laboratory birefringence/rotation (PVLAS) and cosmic birefringence (CMB): reported limits and conservative mapping.

Reference / link	Platform	Published limit (95%)	Implied bound	Row ID (App. E)
PVLAS, Phys. Rev. D 90, 092003 (2014)	Vacuum magnetic birefringence	Field-dependent Δn upper bounds (see paper tables)	$ A \lesssim \text{few} \times 10^{-23}$ – 10^{-22}	E-PVLAS-1
Minami & Komatsu, Phys. Rev. Lett. 125, 221301 (2020)	CMB polarization rotation	$\beta = 0.35^\circ \pm 0.14^\circ$ (all-sky)	$ A \lesssim \mathcal{O}(10^{-2})$ (angle $\rightarrow$ scalar)	C-CMB-1

Table 21: Microwave/sapphire and WGM resonators: stability metrics and TE/TM sensitivity cross-check.

Reference / link	Platform	Published limit / metric	Implied bound	Row ID (App. E)
Savchenkov et al., JOSA B 24, 2988 (2007)	Optical WGM (microcomb)	Linewidth/stability indicators (table values)	$ \Delta\rho/\rho _{\text{short}} \lesssim 10^{-9}$	W-WGM-1
Yu et al., Rev. Sci. Instrum. 83, 094903 (2012)	Microcomb metrology	Beat $\lesssim 40$ Hz @ 25 GHz spacing	$ \Delta\rho/\rho _{\text{short}} \lesssim 1.6 \times 10^{-9}$	W-WGM-2

Cross-mapping rules (summary). We use the first-order relations

$$\Delta \ln f \simeq s_A A + s_G G, \quad \Delta \ln \rho \simeq s_A A + s_G(\theta) G, \quad R - 1 \simeq c_R G$$

with sensitivity coefficients from [Appendix E](#). Direction-dependent anisotropy metrics (e.g. $\Delta c/c$, $\Delta \nu/\nu$ at $2\omega_{\text{rot}}$) are mapped primarily to G , while scalar impedance-type shifts map to A . If harmonic covariances are available, we form the joint ellipse via the 2×2 WLS of [Section 3.4](#).

Conservative assumptions and consistency checks. (i) Default $|s_A|, |s_G| \sim \mathcal{O}(1)$; platform deviations use the ranges in [Appendix E](#). (ii) *Window separation*: WGM short-term metrics are quoted as instrument limits (separate from long-term drift windows; cf. [Section 3.3](#)). (iii) *De-duplication*: among closely related optical MM tests, we weight only the most recent/strongest bound in any combined fit.

Summary (external vs. in-chapter bounds). External optical/cavity limits reach $\Delta c/c$ or $\Delta \nu/\nu$ at 10^{-17} – 10^{-18} , which, once mapped onto our (A, G) parameterization (including geometry/mode overlaps), remain *consistent* with the chapter’s long-window combined bounds ([Section 3.4](#): $|A| \lesssim 3.0 \times 10^{-3}$, $|G| \lesssim 5.5 \times 10^{-3}$) and serve as complementary cross-checks. Short-window WGM results reproduce the instrumental limit $|\Delta \rho/\rho|_{\text{short}} \lesssim 10^{-9}$ quoted in [Section 3.3](#).

Author fill-in guide (cross-check). (1) Enter the numeric **Published 95% limit** (with units) from each DOI into the project datasheet. (2) Select the matching sensitivity **Row ID** in [Appendix E](#) and record it in the Mapping column. (3) The build regenerates the **Implied bounds** and updates the joint WLS in [Section 3.4](#) and the AG-ellipse in [Appendix H](#).

References

- [1] S. Deser and C. Teitelboim, “Duality transformations of Abelian and non-Abelian gauge fields,” *Phys. Rev. D* **13**, 1592–1597 (1976). DOI:[10.1103/PhysRevD.13.1592](https://doi.org/10.1103/PhysRevD.13.1592)
- [2] M. K. Gaillard and B. Zumino, “Duality rotations for interacting fields,” *Nucl. Phys. B* **193**, 221–244 (1981). DOI:[10.1016/0550-3213\(81\)90527-7](https://doi.org/10.1016/0550-3213(81)90527-7)
- [3] J. H. Poynting, “On the transfer of energy in the electromagnetic field,” *Phil. Trans. R. Soc. A* **175**, 343–361 (1884). DOI:[10.1098/rstl.1884.0016](https://doi.org/10.1098/rstl.1884.0016)
- [4] Y. Aharonov and D. Bohm, “Significance of electromagnetic potentials in the quantum theory,” *Phys. Rev.* **115**, 485–491 (1959). DOI:[10.1103/PhysRev.115.485](https://doi.org/10.1103/PhysRev.115.485)
- [5] M. V. Berry, “Quantal phase factors accompanying adiabatic changes,” *Proc. R. Soc. A* **392**, 45–57 (1984). DOI:[10.1098/rspa.1984.0023](https://doi.org/10.1098/rspa.1984.0023)
- [6] K. G. Wilson, “Confinement of quarks,” *Phys. Rev. D* **10**, 2445–2459 (1974). DOI:[10.1103/PhysRevD.10.2445](https://doi.org/10.1103/PhysRevD.10.2445)
- [7] O. Klein, “Quantentheorie und fünfdimensionale Relativitätstheorie,” *Zeitschrift für Physik* **37**, 895–906 (1926). DOI:[10.1007/BF01397481](https://doi.org/10.1007/BF01397481)
- [8] J. M. Overduin and P. S. Wesson, “Kaluza–Klein gravity,” *Physics Reports* **283**, 303–378 (1997). DOI:[10.1016/S0370-1573\(96\)00046-4](https://doi.org/10.1016/S0370-1573(96)00046-4)
- [9] A. Ashtekar, “New variables for classical and quantum gravity,” *Phys. Rev. Lett.* **57**, 2244–2247 (1986). DOI:[10.1103/PhysRevLett.57.2244](https://doi.org/10.1103/PhysRevLett.57.2244)
- [10] J. F. Barbero, “Real Ashtekar variables for Lorentzian signature space-times,” *Phys. Rev. D* **51**, 5507–5510 (1995). DOI:[10.1103/PhysRevD.51.5507](https://doi.org/10.1103/PhysRevD.51.5507)
- [11] J. B. Kogut, “An introduction to lattice gauge theory and spin systems,” *Rev. Mod. Phys.* **51**, 659–713 (1979). DOI:[10.1103/RevModPhys.51.659](https://doi.org/10.1103/RevModPhys.51.659)
- [12] K. S. Yee, “Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media,” *IEEE Trans. Antennas Propag.* **14**, 302–307 (1966). DOI:[10.1109/TAP.1966.1138693](https://doi.org/10.1109/TAP.1966.1138693)
- [13] N. Hitchin, “Generalized Calabi–Yau manifolds,” *Q. J. Math.* **54**(3), 281–308 (2003). DOI:[10.1093/qjmath/54.3.281](https://doi.org/10.1093/qjmath/54.3.281)
- [14] M. Gualtieri, “Generalized complex geometry,” *Annals of Mathematics* **174**(1), 75–123 (2011). DOI:[10.4007/annals.2011.174.1.3](https://doi.org/10.4007/annals.2011.174.1.3)
- [15] C. Hull and B. Zwiebach, “Double field theory,” *JHEP* **09**, 099 (2009). DOI:[10.1088/1126-6708/2009/09/099](https://doi.org/10.1088/1126-6708/2009/09/099)
- [16] F. Cabral and F. S. N. Lobo, “Electrodynamics and spacetime geometry: Astrophysical applications,” *Eur. Phys. J. Plus* **132**, 281 (2017). DOI:[10.1140/epjp/i2017-11618-2](https://doi.org/10.1140/epjp/i2017-11618-2)
- [17] F. W. Hehl, Y. N. Obukhov, and B. Rosenow, “Is the quantum Hall effect influenced by the gravitational field?,” *Phys. Rev. Lett.* **93**, 096804 (2004). DOI:[10.1103/PhysRevLett.93.096804](https://doi.org/10.1103/PhysRevLett.93.096804)
- [18] L. X. Wang *et al.*, “Origin of non-saturating linear magnetoresistance in Dirac semimetal Cd₃As₂ nanowires,” *Nat. Commun.* **7**, 10769 (2016). DOI:[10.1038/ncomms10769](https://doi.org/10.1038/ncomms10769).

- [19] H. Peng *et al.*, “Aharonov–Bohm interference in topological insulator Bi₂Se₃ nanoribbons,” *Nat. Mater.* **9**, 225–229 (2010). DOI:10.1038/nmat2609.
- [20] Y. Yan *et al.*, “High-mobility Bi₂Se₃ nanoplates and angle-dependent magnetotransport,” *Sci. Rep.* **4**, 3817 (2014). DOI:10.1038/srep03817.
- [21] M. Kato *et al.*, “Paired $\hbar/2e$ Aharonov–Bohm oscillations with tilted field in an antidot array,” arXiv:0909.1395 (2009). arXiv:0909.1395.
- [22] P. I. Dankov, “Two-Resonator Method for Measurement of Dielectric Anisotropy in Multi-Layer Samples,” *IEEE Transactions on Microwave Theory and Techniques*, **54**(4), 1534–1544 (2006). doi:10.1109/TMTT.2006.871247.
- [23] V. N. Levcheva, B. N. Hadjistamov, and P. I. Dankov, “Two-Resonator Method for Characterization of Dielectric Substrate Anisotropy,” *Bulgarian J. Phys.* **35**, 33–52 (2008). PDF.
- [24] J. Krupka, “Frequency domain complex permittivity measurements at microwave frequencies,” *Meas. Sci. Technol.* **17**(6), R55–R70 (2006). DOI:10.1088/0957-0233/17/6/R01.
- [25] J. Krupka, “Microwave measurements of electromagnetic properties of materials,” *Materials* **14**(17), 5097 (2021). DOI:10.3390/ma14175097.
- [26] A. A. Savchenkov *et al.*, “Whispering-gallery-mode resonators as frequency references. I. Fundamental limitations,” *J. Opt. Soc. Am. B* **24**(6), 1324–1335 (2007). DOI:10.1364/JOSAB.24.001324.
- [27] L. Yu and V. Fernicola, “Temperature–frequency characteristic of a spherical sapphire whispering gallery mode resonator at 13.6 GHz,” *Rev. Sci. Instrum.* **83**, 094903 (2012). DOI:10.1063/1.4746991.
- [28] F. W. Hehl and Y. N. Obukhov, *Foundations of Classical Electrodynamics: Charge, Flux, and Metric* (Springer, 2003). DOI:10.1007/978-1-4612-0051-2.
- [29] Y. N. Obukhov and F. W. Hehl, “Spacetime metric from linear electrodynamics,” *Phys. Lett. B* **458**, 466–470 (1999). DOI:10.1016/S0370-2693(99)00643-7.
- [30] F. W. Hehl, Y. N. Obukhov, and G. F. Rubilar, “Spacetime metric from linear electrodynamics II,” *Ann. Phys. (Leipzig)* **9**, SI-71–SI-78 (1999). arXiv:gr-qc/9911096.
- [31] J. H. Bardarson, P. W. Brouwer, and J. E. Moore, “Aharonov–Bohm oscillations in disordered topological insulator nanowires,” *Phys. Rev. Lett.* **105**, 156803 (2010). DOI:10.1103/PhysRevLett.105.156803.
- [32] Kim, Seung-il, *Introducing the Curvature Field Function: Toward a Geometric Formulation of Wavefunction Collapse*, jxiv (2025). doi:10.51094/jxiv.1522.
- [33] Kim, Seung-il, *Curvature Field Formulation of Gravity: Toward a Physical Reconstruction of Spacetime*, jxiv (2025). doi:10.51094/jxiv.1579.
- [34] H. Bluhm, N. C. Koshnick, J. A. Bert, M. E. Huber, and K. A. Moler, “Persistent Currents in Normal Metal Rings,” *Phys. Rev. Lett.* **102**, 136802 (2009). DOI:10.1103/PhysRevLett.102.136802
- [35] F. V. Tikhonenko, D. W. Horsell, R. V. Gorbachev, and A. K. Savchenko, “Weak localization in graphene flakes,” *Phys. Rev. Lett.* **100**, 056802 (2008). DOI:10.1103/PhysRevLett.100.056802 (arXiv:0707.0140).

- [36] O. Millo, S. J. Klepper, M. W. Keller, D. E. Prober, S. Xiong, and A. D. Stone; R. N. Sacks, “Reduction of the mesoscopic conductance–fluctuation amplitude in GaAs/AlGaAs heterojunctions due to spin–orbit scattering,” *Phys. Rev. Lett.* **65**, 1494–1497 (1990). DOI:10.1103/PhysRevLett.65.1494 (see device count and T range on p. 1494).
- [37] T. Kaluza, “On the unification problem in physics” (revised English translation of 1921 paper), *Int. J. Mod. Phys. D* **27**, 1870001 (2018). DOI:10.1142/S0218271818700017
- [38] K. R. Amin, S. S. Ray, N. Pal, R. Pandit, and A. Bid, “Exotic multifractal conductance fluctuations in graphene,” *Communications Physics* **1**, 20 (2018). DOI:10.1038/s42005-017-0001-4
- [39] B. Grbić, R. Leturcq, T. Ihn, K. Ensslin, D. Reuter, and A. D. Wieck, “Strong spin–orbit interactions and weak antilocalization in carbon–doped p–type GaAs/AlGaAs heterostructures” (2007). arXiv:0711.0492
- [40] J. Berezovsky and R. M. Westervelt, “Imaging universal conductance fluctuations in mesoscopic graphene” (2009). arXiv:0907.0428
- [41] A. Tonomura *et al.*, “Evidence for Aharonov–Bohm effect with magnetic field completely shielded from electron wave,” *Phys. Rev. Lett.* **56**, 792–795 (1986). DOI:10.1103/PhysRevLett.56.792
- [42] N. Osakabe *et al.*, “Experimental confirmation of Aharonov–Bohm effect using a toroidal magnetic field confined by a superconductor,” *Phys. Rev. A* **34**, 815–822 (1986). DOI:10.1103/PhysRevA.34.815
- [43] A. E. Hansen *et al.*, “Mesoscopic decoherence in Aharonov–Bohm rings,” *Phys. Rev. B* **64**, 045327 (2001). DOI:10.1103/PhysRevB.64.045327
- [44] M. Sigrist *et al.*, “Phase coherence in the inelastic cotunneling regime,” *Phys. Rev. Lett.* **96**, 036804 (2006). DOI:10.1103/PhysRevLett.96.036804
- [45] E. M. Q. Jariwala *et al.*, “Diamagnetic persistent current in diffusive normal-metal rings,” *Phys. Rev. Lett.* **86**, 1594–1597 (2001). DOI:10.1103/PhysRevLett.86.1594
- [46] W. D. Oliver *et al.*, “Mach–Zehnder interferometry in a strongly driven superconducting qubit,” *Science* **310**, 1653–1657 (2005). DOI:10.1126/science.1119678
- [47] S. N. Shevchenko, S. Ashhab, and F. Nori, “Landau–Zener–Stückelberg interferometry,” *Physics Reports* **492**, 1–30 (2010). doi:10.1016/j.physrep.2010.03.002. Preprint: arXiv:0911.1917.
- [48] J. Stehlik *et al.*, “Landau–Zener–Stückelberg interferometry of a single-electron charge qubit,” *Phys. Rev. B* **86**, 121303(R) (2012). DOI:10.1103/PhysRevB.86.121303
- [49] M. Heiblum *et al.*, “Coherence and phase-sensitive measurements with a quantum dot in an Aharonov–Bohm interferometer,” *Physica B* **227**, 147–154 (1996). DOI:10.1016/0921-4526(96)00378-X
- [50] H. Aikawa, K. Kobayashi, A. Sano, S. Katsumoto, and Y. Iye, “Observation of ‘Partial Coherence’ in an Aharonov–Bohm Interferometer with a Quantum Dot,” *Phys. Rev. Lett.* **92**, 176802 (2004). DOI:10.1103/PhysRevLett.92.176802.
- [51] R. A. Webb, S. Washburn, C. P. Umbach, and R. B. Laibowitz, “Observation of h/e Aharonov–Bohm oscillations in normal–metal rings,” *Phys. Rev. Lett.* **54**, 2696–2699 (1985). DOI:10.1103/PhysRevLett.54.2696

- [52] V. Chandrasekhar, M. J. Rooks, S. Wind, and D. E. Prober, “Observation of Aharonov–Bohm electron interference effects with periods h/e and $h/2e$ in individual micron–size, normal–metal rings,” *Phys. Rev. Lett.* **55**, 1610–1613 (1985). DOI:10.1103/PhysRevLett.55.1610
- [53] C. P. Umbach, C. Van Haesendonck, R. B. Laibowitz, S. Washburn, and R. A. Webb, “Direct observation of ensemble averaging of the Aharonov–Bohm effect in normal–metal loops,” *Phys. Rev. Lett.* **56**, 386–389 (1986). DOI:10.1103/PhysRevLett.56.386
- [54] M. Büttiker, “Four–terminal phase–coherent conductance,” *Phys. Rev. Lett.* **57**, 1761–1764 (1986). DOI:10.1103/PhysRevLett.57.1761
- [55] M. Büttiker, “Symmetry of electrical conduction,” *IBM J. Res. Dev.* **32**, 317–334 (1988). DOI:10.1147/rd.323.0317
- [56] R. Schuster, E. Buks, M. Heiblum, D. Mahalu, V. Umansky, and H. Shtrikman, “Phase measurement in a quantum dot via a double–slit interference,” *Nature* **385**, 417–420 (1997). DOI:10.1038/385417a0
- [57] S. Russo, J. B. Oostinga, D. Wehenkel, H. B. Heersche, S. S. Sobhani, L. M. K. Vandersypen, and A. F. Morpurgo, “Observation of Aharonov–Bohm conductance oscillations in a graphene ring,” *Phys. Rev. B* **77**, 085413 (2008). DOI:10.1103/PhysRevB.77.085413
- [58] B. Hackens, F. Martins, T. Ouisse, H. Sellier, S. Bollaert, X. Wallart, A. Cappy, J. Chevrier, V. Bayot, and S. Huant, “Imaging and controlling electron transport inside a quantum ring,” *Nature Physics* **2**, 826–830 (2006). DOI:10.1038/nphys459.
- [59] Y. Ji, Y. Chung, D. Sprinzak, M. Heiblum, D. Mahalu, and H. Shtrikman, “An electronic Mach–Zehnder interferometer,” *Nature* **422**, 415–418 (2003). DOI:10.1038/nature01503
- [60] N. Ofek, M. Heiblum, I. Neder, M. Mahalu, and V. Umansky, “Role of interactions in an electronic Fabry–Perot interferometer operating in the quantum Hall effect regime,” *Proc. Natl. Acad. Sci. USA* **107**, 5276–5281 (2010). DOI:10.1073/pnas.0912624107
- [61] M. A. Castellanos-Beltran, D. Quezada, R. A. Zadorozhny, R. A. Buhrman, and D. C. Ralph, “Measurement of the full distribution of persistent current in normal–metal rings,” *Phys. Rev. Lett.* **110**, 156801 (2013). DOI:10.1103/PhysRevLett.110.156801
- [62] A. C. Bleszynski-Jayich *et al.*, “Persistent currents in normal metal rings,” *Science* **326**, 272–275 (2009). DOI:10.1126/science.1178139
- [63] F. Pierre, A. B. Gougam, A. Anthore, H. Pothier, D. Esteve, and N. O. Birge, “Dephasing of electrons in mesoscopic metal wires,” *Phys. Rev. B* **68**, 085413 (2003). DOI:10.1103/PhysRevB.68.085413
- [64] INTERMAGNET, “International Real-time Magnetic Observatory Network (data portal and documentation),” (accessed 2025). <https://intermagnet.org/>; data portal: https://imag-data.bgs.ac.uk/GIN_V1/GINForms2.
- [65] JHU/APL SuperMAG, “Global ground-based magnetometer collaboration (data & APIs),” (accessed 2025). <https://supermag.jhuapl.edu/>.

- [66] NOAA NCEI, “Geomagnetic data products and indices (WDS/WDC),” (accessed 2025). [Data portal](#); [Indices](#).
- [67] WDC for Geomagnetism, Kyoto, “Dst/AE index services (final, provisional, quick-look),” (accessed 2025). [Dst directory](#). [AE directory](#).
- [68] J. Dauber, M. Oellers, F. Venn, A. Epping, K. Watanabe, T. Taniguchi, F. Hassler, and C. Stampfer, Aharonov–Bohm oscillations and magnetic focusing in ballistic graphene rings, *Phys. Rev. B* **96**, 205407 (2017). [doi:10.1103/PhysRevB.96.205407](#).
- [69] J. G. Hartnett, M. E. Tobar, E. N. Ivanov, and J. Krupka, Room temperature measurement of the anisotropic loss tangent of sapphire using the whispering gallery mode technique, *IEEE Trans. Ultrason., Ferroelect., Freq. Control* **53**(1), 34–41 (2006). [doi:10.1109/TUFFC.2006.1588389](#).
- [70] V. M. Gvozdkov, Yu. V. Pershin, E. Steep, A. G. M. Jansen, and P. Wyder, de Haas–van Alphen oscillations in the quasi-two-dimensional organic conductor κ -(ET)₂Cu(NCS)₂: The magnetic breakdown approach, *Phys. Rev. B* **65**, 165102 (2002). [doi:10.1103/PhysRevB.65.165102](#).