

Imprinted Curvature Hypothesis: Constraint-Compatible Primordial Long Modes in General Relativity and Their Cosmological Tests

Watchara Phumdokmai

Independent Researcher, Bangkok, Thailand

ORCID: 0009-0008-5997-0534

Email: wp_biomed@hotmail.com

Abstract

The Imprinted Curvature Hypothesis (ICH) is formulated as a restricted cosmological initial-condition hypothesis within unmodified General Relativity (GR). It introduces no new gravitational dynamics, matter component, or homogeneous expansion law, but specifies a weak coherent ultra-long-wavelength contribution to otherwise standard adiabatic primordial curvature initial data. We construct an explicit family of ADM-constraint-satisfying initial data, give a standard adiabatic matter completion, identify the gauge and equivalence-principle boundaries of the construction, and propagate the same global parameter set through standard Einstein–Boltzmann transfer functions to cosmic microwave background (CMB), large-scale-structure, lensing, void, rare-object, and cosmic-web observables. A deterministic standing-wave benchmark is propagated with CAMB and independently checked with CLASS. The frozen Planck PR3 temperature-only scan over $2 \leq \ell \leq 29$ gives $\Delta\chi^2 = -17.745$ but a matched 2000-sky look-elsewhere probability $p = 0.454$, hence no evidence against isotropic Λ CDM. A subsequent noise-aware temperature-plus-polarization scan gives $\Delta\chi^2 = -24.990$ and $p = 0.02390$ in the unmasked full-map diagnostic, but this feature fails Planck half-mission replication ($p = 0.3565$) and becomes statistically ordinary under corrected polarization masks ($p = 0.103$ – 0.180) and masked split cross-validation ($p = 0.511$ – 0.614). A separate Planck PR3 fixed-parameter TT+TE+EE+lensing consistency audit and CAMB high- ℓ transfer check show that the tested ultra-long modes are confined to the lowest multipoles on the stored transfer grid; this is a CMB-survival statement, not a directional

likelihood detection. The present evidence therefore supports mathematical admissibility and a falsifiable cross-survey program, but not a CMB preference for ICH over Λ CDM.

Keywords: cosmology, general relativity, cosmological perturbation theory, initial conditions, cosmic microwave background, large-scale structure, weak lensing, cosmic web

1. Introduction

The standard cosmological model succeeds across a wide range of scales and epochs, so a proposal intended to operate alongside it should begin from that success rather than from a presumption that the dynamical laws must be replaced. The Imprinted Curvature Hypothesis (ICH) is deliberately narrow. It does not modify the Einstein field equations, replace the Friedmann–Lemaître–Robertson–Walker (FLRW) background, introduce an additional stress–energy component, or postulate a new propagating sector. Its question is instead one of admissible initial geometry: can the realized cosmological initial data contain a weak, ultra-long-wavelength, coherent curvature component whose subsequent evolution is entirely standard, but whose residual influence is statistically detectable in selected large-scale observables?

Long-wavelength curvature perturbations are standard objects in cosmological perturbation theory. Their gauge properties, conservation laws, separate-universe interpretation, and observable consequences are well established [4–6, 8, 9]. ICH does not claim otherwise. The proposed distinction lies in how a restricted subset of that familiar initial-data space is organized and tested. The hypothesis is that the realized low- k geometry contains a subdominant coherent component described by a small global parameter set and that the same parameters must survive confrontation with all relevant surveys. If direction, scale, phase, morphology, or interpretation can be re-tuned independently from anomaly to anomaly, the framework loses predictive content and collapses into post-hoc selection.

Earlier formulations of ICH emphasized conceptual boundaries, qualitative predictions, and the relation between initial intrinsic curvature, gauge-robust perturbation variables, and observables. The central constructive weakness was that the initial-curvature statement was not accompanied by a full constraint-satisfying initial-data family, while the observable mapping remained schematic. The present manuscript consolidates the mathematical

construction and all completed CMB validation stages into one self-contained paper. In particular, it integrates the frozen temperature-only Planck diagnostic, the corrected temperature-plus-polarization extension, half-mission and mask robustness tests, and a multi-spectrum CMB fixed-parameter/acoustic-scale consistency audit. Negative and superseded diagnostics are retained where they materially define the scientific boundary of the result.

The model comparison can be summarized as

$$H_0 : \quad \zeta(\mathbf{k}) = \zeta_G(\mathbf{k}), \quad (1)$$

$$H_1 : \quad \zeta(\mathbf{k}) = \zeta_G(\mathbf{k}) + \zeta_I(\mathbf{k}; \Theta_I), \quad (2)$$

where ζ_G is the ordinary stochastic adiabatic curvature field and ζ_I is a restricted coherent contribution to the same initial data. ICH changes the prior/realization class, not the local dynamical equations.

The intended hierarchy of claims is essential. First, the model must be posed consistently inside ordinary GR. Second, one restricted low- k realization must propagate into multiple observables without changing its global parameters. Third, that realization must be compared statistically with the standard null model after the full search volume and relevant systematics are included. Mathematical admissibility is not evidence that the Universe selected these initial conditions, and a suggestive object or individual residual is not a detection.

2. Position relative to standard long-mode physics

2.1. Standard long-mode restrictions

On an FLRW background, scalar perturbations admit gauge-invariant descriptions. On sufficiently large scales and for adiabatic perturbations, the curvature perturbation ζ is conserved to leading order in gradients [4–6]. Separate-universe and response approaches express the effect of sufficiently long modes on shorter-scale observables as a local background modulation plus finite-gradient density and tidal responses [8–10].

Three restrictions follow. A constant adiabatic scalar mode is not a locally measurable curvature anomaly. A pure gradient is locally removable in the appropriate long-wavelength limit and cannot by itself be advertised as a physical preferred direction. Measurable effects of an ultra-long mode must therefore enter through finite curvature/tidal derivatives, complete light-cone

projection terms, boundary conditions, or survey-scale correlations. ICH adopts these restrictions rather than attempting to evade them.

Likewise, the ADM initial-value formulation of GR already permits inhomogeneous initial spatial geometry subject to the Hamiltonian and momentum constraints [2, 3]. The existence of non-FLRW initial data is not a novelty claim. The scientific question is whether the restricted ICH family is mathematically admissible, observationally coherent, and distinguishable from ordinary cosmic variance.

2.2. The specific ICH restriction

A minimal deterministic benchmark is

$$\zeta_{\text{I}}(\boldsymbol{x}) = A_{\text{I}} \cos(k_{\text{I}} \hat{\boldsymbol{n}}_{\text{I}} \cdot \boldsymbol{x} + \phi_{\text{I}}), \quad (3)$$

with

$$\Theta_{\text{I}} = \{A_{\text{I}}, k_{\text{I}}, \hat{\boldsymbol{n}}_{\text{I}}, \phi_{\text{I}}\}, \quad |A_{\text{I}}| \ll 1. \quad (4)$$

The standing wave is a benchmark rather than a claim that nature must realize exactly one Fourier mode. A narrow band-limited low- k template can be used instead. The scientific discipline is that the parameterization and priors are specified before the decisive prospective test and the same Θ_{I} propagates across all channels.

The central degeneracy is therefore

$$\begin{array}{c} \text{restricted ICH long mode} \\ \stackrel{?}{=} \\ \text{allowed ordinary adiabatic long-mode realization.} \end{array} \quad (5)$$

The present paper does not resolve this by definition. It specifies the forward and statistical machinery required to resolve it with data.

2.3. ICH is not modified gravity

The gravitational action, Einstein tensor, gravitational coupling, local field content, and matter conservation law are unchanged. ICH does not require an anisotropic background metric or a new homogeneous expansion history. Its burden is therefore to show that a proposed signal can arise from allowed initial data and standard transfer physics while remaining subdominant to the successful baseline cosmology.

3. Model definition and domain of validity

Let Σ_* denote an early spacelike hypersurface fixed by a physical slicing condition. The original geometric definition writes its intrinsic scalar curvature as

$${}^{(3)}R(t_*, \mathbf{x}) = \overline{{}^{(3)}R}(t_*) + \epsilon\Phi(\mathbf{x}), \quad (6)$$

with the weak-curvature condition

$$\left| \frac{a_*^2}{6} \epsilon\Phi(\mathbf{x}) \right| \ll 1. \quad (7)$$

The function Φ is not an independent dynamical field and has no separate equation of motion.

The minimal model imposes the following restrictions: (i) Einstein gravity and standard stress–energy conservation are unchanged; (ii) the ICH contribution resides in admissible initial data rather than a new source term; (iii) the amplitude is weak and perturbative; (iv) spectral support is restricted to ultra-large scales; (v) exact FLRW/ Λ CDM is recovered continuously as $A_I \rightarrow 0$; (vi) the minimal scalar realization is adiabatic; (vii) physical claims are made only in a fixed physical slicing, gauge-invariant variables, or complete observables; (viii) the imprint does not create a compact localized mass component; (ix) the background distance–redshift relation is not replaced by an ICH-specific expansion law; and (x) one global Θ_I is shared across observational channels.

A flexible band-limited realization may be written

$$\zeta_I(\mathbf{k}; \Theta_I) = A_I W_I(k; k_I, \sigma_k) Q(\hat{\mathbf{k}}; \hat{\mathbf{n}}_I) e^{i\phi_k}, \quad (8)$$

with

$$W_I(k) = 0 \quad \text{for} \quad k > k_{\text{cut}}, \quad (9)$$

and the reality condition

$$\zeta_I(-\mathbf{k}; \Theta_I) = \zeta_I^*(\mathbf{k}; \Theta_I). \quad (10)$$

The angular template, bandwidth, and phase prescription are part of the model prior and must be frozen before the test to which they are applied.

4. Constructive initial data in unmodified General Relativity

We use metric signature $(-, +, +, +)$ and define

$$K_{ij} \equiv -\frac{1}{2}\mathcal{L}_{\mathbf{n}}\gamma_{ij}, \quad (11)$$

where n^μ is the future-directed unit normal. The spacetime dynamics remain

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu}, \quad \nabla_\mu T^{\mu\nu} = 0. \quad (12)$$

For induced metric γ_{ij} , extrinsic curvature K_{ij} , energy density $E = T_{\mu\nu}n^\mu n^\nu$, and momentum density $p_i = -\gamma_i^\mu T_{\mu\nu}n^\nu$, the ADM constraints are [2, 3]

$${}^{(3)}R + K^2 - K_{ij}K^{ij} = 16\pi G E + 2\Lambda, \quad (13)$$

$$D_j K^j_i - D_i K = 8\pi G p_i. \quad (14)$$

4.1. Explicit scalar initial-data family

On an early total-comoving/flow-orthogonal slice with a spatially flat reference background choose

$$\gamma_{ij}(t_*, \mathbf{x}) = a_*^2 e^{2\zeta_I(\mathbf{x})} \delta_{ij}, \quad (15)$$

and constant-mean-curvature extrinsic curvature

$$K_{ij}(t_*, \mathbf{x}) = -H_* \gamma_{ij}(t_*, \mathbf{x}). \quad (16)$$

The exact intrinsic scalar curvature is

$${}^{(3)}R = \frac{e^{-2\zeta_I}}{a_*^2} \left[-4\nabla^2 \zeta_I - 2(\nabla \zeta_I)^2 \right], \quad (17)$$

so at linear order

$${}^{(3)}R = -\frac{4}{a_*^2} \nabla^2 \zeta_I + \mathcal{O}(\zeta_I^2). \quad (18)$$

Because $K = -3H_*$ and $K_{ij}K^{ij} = 3H_*^2$, the Hamiltonian constraint is solved exactly on Σ_* by

$$E(\mathbf{x}) = \frac{3H_*^2}{8\pi G} - \frac{\Lambda}{8\pi G} + \frac{{}^{(3)}R(\mathbf{x})}{16\pi G}. \quad (19)$$

Since $K^j_i = -H_* \delta_i^j$ with spatially constant H_* , the momentum constraint is solved by

$$p_i = 0. \quad (20)$$

Thus $(\gamma_{ij}, K_{ij}, E, p_i)$ is an explicit constraint-satisfying family. This establishes admissibility, not uniqueness.

4.2. Adiabatic matter completion

Constraint satisfaction fixes total energy and momentum densities but not the microphysical decomposition of $T_{\mu\nu}$. To connect the ADM slice to standard Einstein–Boltzmann evolution, choose the ordinary cosmological species with no new ICH species, impose vanishing total momentum on the total-comoving slice, and initialize scalar perturbations adiabatically:

$$S_{AB} \equiv 3(\zeta_A - \zeta_B) = 0, \quad (21)$$

or equivalently on super-horizon scales

$$\frac{\delta\rho_A}{\rho_A + p_A} = \frac{\delta\rho_B}{\rho_B + p_B} \quad (22)$$

for all ordinary species, with $\sum_A \delta\rho_A = \delta E$ from Eq. (19). This gives a standard-matter realization suitable for an ordinary Einstein–Boltzmann solver [7].

4.3. Single-mode gradient suppression

For the standing wave,

$${}^{(3)}R_I \simeq \frac{4A_I k_I^2}{a_*^2} \cos(k_I \hat{\mathbf{n}}_I \cdot \mathbf{x} + \phi_I), \quad (23)$$

and, defining $\bar{E}_{\text{eff},*} = 3H_*^2/(8\pi G)$,

$$\frac{\delta E}{\bar{E}_{\text{eff},*}} \simeq \frac{2A_I}{3} \left(\frac{k_I}{a_* H_*} \right)^2 \cos(k_I \hat{\mathbf{n}}_I \cdot \mathbf{x} + \phi_I). \quad (24)$$

The k^2 suppression is essential: an exactly constant adiabatic mode has no local scalar-curvature signal in this construction.

4.4. No new local degree of freedom and constraint propagation

The construction changes the allowed initial data, not the local dynamical phase space of GR. The imprint variable is not supplied with a conjugate momentum or an independent evolution equation. Once gauge freedom and the standard constraints are accounted for, the propagating gravitational content is the ordinary metric content of GR. Promoting ζ_I or Φ to an independently propagating field would define a different theory.

For completeness, the standard $3 + 1$ evolution begins schematically with

$$(\partial_t - \mathcal{L}_\beta)\gamma_{ij} = -2NK_{ij}, \quad (25)$$

$$\begin{aligned} (\partial_t - \mathcal{L}_\beta)K_{ij} = & -D_i D_j N + N \left({}^{(3)}R_{ij} + KK_{ij} - 2K_{ik}K^k_j \right) \\ & + \text{standard matter terms.} \end{aligned} \quad (26)$$

where N and β^i are the lapse and shift. No ICH-specific evolution term is added. The contracted Bianchi identity and ordinary matter conservation,

$$\nabla_\mu G^{\mu\nu} = 0, \quad \nabla_\mu T^{\mu\nu} = 0, \quad (27)$$

provide the usual constraint-propagation structure. An initially constraint-satisfying realization is therefore evolved by the same Einsteinian system as any other admissible cosmological data.

5. Gauge, equivalence principle, persistence, and Weyl boundary

The defining hypersurface is not an arbitrary coordinate slice. For the scalar adiabatic benchmark it is taken to be the total-comoving/flow-orthogonal slicing, equivalently the uniform-density slicing at leading order on super-horizon scales. A sufficiently long mode may be expanded locally as

$$\zeta_L(\mathbf{x}) = \zeta_0 + b_i x^i + \frac{1}{2} C_{ij} x^i x^j + \dots \quad (28)$$

For an adiabatic long mode, constant and pure-gradient pieces are locally removable in the appropriate limit. Representative finite-gradient content is

$${}^{(3)}R_L \sim -4a^{-2}\nabla^2\zeta_L, \quad (29)$$

and the trace-free tidal tensor

$$\mathcal{K}_{ij} \equiv \left(\frac{\partial_i \partial_j}{\nabla^2} - \frac{\delta_{ij}}{3} \right) \delta_L. \quad (30)$$

A claimed ICH signal must therefore be formulated in gauge-invariant variables or a complete observable including the relevant projection terms [8].

For the curvature perturbation,

$$\dot{\zeta} = -\frac{H}{\rho + p} \delta p_{\text{nad}} + \mathcal{O}\left(\frac{k^2}{a^2 H^2}\right), \quad (31)$$

so under adiabatic super-horizon conditions $\dot{\zeta} \simeq 0$ [5, 6]. “Frozen” therefore means persistence of comoving large-scale curvature information, not a static metric, matter distribution, Weyl tensor, or numerical $^{(3)}R$. For the scalar benchmark, $^{(3)}R \propto a^{-2}$ while the comoving pattern in $\zeta_{\mathbf{l}}$ remains conserved in the stated regime.

A scalar intrinsic three-curvature modulation alone does not imply a nonzero spacetime Weyl tensor. Weyl curvature depends on the full spacetime initial data, including trace-free spatial-Ricci and extrinsic-curvature structure. The conformally flat scalar benchmark is therefore not proof of primordial Weyl curvature. Any Weyl/tidal extension must be specified separately and consistently with the full constraints.

6. Unified forward model

6.1. Standard transfer to matter and metric

In conformal Newtonian gauge,

$$ds^2 = a^2(\eta) \left[-(1 + 2\Psi)d\eta^2 + (1 - 2\Phi)d\mathbf{x}^2 \right]. \quad (32)$$

No ICH-specific transfer law is introduced. Standard Einstein–Boltzmann evolution gives

$$\delta_m(\mathbf{k}, z) = T_\delta(k, z)\zeta(\mathbf{k}), \quad (33)$$

$$\Phi(\mathbf{k}, z) = T_\Phi(k, z)\zeta(\mathbf{k}), \quad (34)$$

$$\Psi(\mathbf{k}, z) = T_\Psi(k, z)\zeta(\mathbf{k}), \quad (35)$$

with lensing potential

$$\Phi_W \equiv \frac{\Phi + \Psi}{2} = T_W(k, z)\zeta(\mathbf{k}). \quad (36)$$

On sufficiently large linear scales a useful benchmark is

$$\delta_m(\mathbf{k}, z) = \mathcal{M}(k, z)\zeta(\mathbf{k}), \quad \mathcal{M}(k, z) = \frac{2k^2 T(k) D(z)}{5\Omega_m H_0^2}, \quad (37)$$

subject to normalization conventions.

6.2. Common long-mode response

For an environmental observable O embedded in a long-wavelength density/tidal field,

$$\frac{O - \bar{O}}{\bar{O}} = R_1^O \delta_L + R_K^O \hat{q}_i \hat{q}_j \mathcal{K}_{ij} + \frac{1}{2} R_2^O \delta_L^2 + R_{K^2}^O \mathcal{K}_{ij} \mathcal{K}^{ij} + \dots \quad (38)$$

The coefficients are ordinary perturbative, separate-universe, simulation-calibrated, or hydrodynamic response coefficients, not new ICH constants [9, 10].

Where the full initial data contain a finite tidal component, the relevant quantity is not a new force but the ordinary trace-free tidal field, equivalently the electric-Weyl/tidal response after the complete spacetime solution is specified. For a host with directional statistic S_{sat} , a minimal environmental model is

$$\langle S_{\text{sat}} | \mathcal{K}_{ij} \rangle = \bar{S}_{\text{sat}} + R_{\text{sat}}^K \hat{s}_i \hat{s}_j \mathcal{K}_{ij} + \dots \quad (39)$$

ICH therefore permits a weak statistical preference for coherent satellite organization, anisotropic accretion, or filament-aligned substructure in selected environments; it does not predict a universal satellite plane and does not replace local halo or baryonic dynamics [30, 31].

7. CMB template and analytic benchmark

For $X \in \{T, E\}$, the deterministic CMB harmonic template is

$$a_{\ell m}^{X, \text{I}} = 4\pi i^\ell \int \frac{d^3 k}{(2\pi)^3} \zeta_{\text{I}}(\mathbf{k}; \Theta_{\text{I}}) \Delta_\ell^X(k) Y_{\ell m}^*(\hat{\mathbf{k}}). \quad (40)$$

For the real standing wave in Eq. (3),

$$a_{\ell m}^{X, \text{I}} = 2\pi A_{\text{I}} i^\ell \Delta_\ell^X(k_{\text{I}}) \left[e^{i\phi_{\text{I}}} + (-1)^\ell e^{-i\phi_{\text{I}}} \right] Y_{\ell m}^*(\hat{\mathbf{n}}_{\text{I}}). \quad (41)$$

In the Sachs–Wolfe limit,

$$\Delta_\ell^T(k) \simeq \frac{1}{5} j_\ell(k\chi_*), \quad (42)$$

which gives

$$C_\ell^{\text{I, SW}} = \frac{2\pi A_{\text{I}}^2}{25} j_\ell^2(k_{\text{I}}\chi_*) \left[1 + (-1)^\ell \cos(2\phi_{\text{I}}) \right]. \quad (43)$$

The parity/phase factor is a property of this benchmark, not a universal ICH prediction.

For a deterministic map-level template $t(\Theta_I)$ and standard covariance C_0 ,

$$-2 \ln \mathcal{L}_{\text{CMB}} = [d - t(\Theta_I)]^\dagger C_0^{-1} [d - t(\Theta_I)] + \ln \det C_0 + \text{const.} \quad (44)$$

If an ICH prescription instead modifies statistical covariance, one must use $C(\Theta_I)$ consistently; deterministic-mean and covariance models should not be mixed without an explicit hierarchy.

For an approximately scale-invariant standard curvature spectrum with amplitude A_s ,

$$\frac{\ell(\ell+1)C_\ell^{\text{G,SW}}}{2\pi} \simeq \frac{A_s}{25}. \quad (45)$$

A useful worst-phase fractional benchmark is

$$f_\ell^{\text{SW}} \equiv \frac{D_\ell^{\text{I}}}{D_\ell^{\text{G}}} \leq \frac{2\ell(\ell+1)A_{\text{I}}^2 j_\ell^2(x)}{A_s}, \quad x \equiv k_{\text{I}}\chi_*. \quad (46)$$

Using $A_s = 2.1 \times 10^{-9}$, $\chi_* = 13.9$ Gpc, $h = 0.674$, $\Omega_m = 0.315$, and $A_{\text{I}} = 5 \times 10^{-6}$ gives Table 1. The matter amplitudes use Eq. (37) with $T(k) \simeq 1$ at the very low k shown.

Table 1: Illustrative analytic sanity benchmark for one deterministic standing-wave amplitude $A_{\text{I}} = 5 \times 10^{-6}$. This is neither a Planck likelihood fit nor an N -body result.

$x = k\chi_*$	k [h Mpc $^{-1}$]	$ \delta_m^{\text{I}}(0) $	$\max f_\ell^{\text{SW}}$ (ℓ at max)
1	1.07×10^{-4}	6.50×10^{-7}	5.50×10^{-4} (2)
3	3.20×10^{-4}	5.85×10^{-6}	1.27×10^{-2} (2)
5	5.34×10^{-4}	1.63×10^{-5}	1.67×10^{-2} (4)
10	1.07×10^{-3}	6.50×10^{-5}	2.70×10^{-2} (8)

For this illustrative amplitude and $x \leq 10$, the worst-phase Sachs–Wolfe contribution remains below about 2.7% of the idealized scale-invariant plateau for $2 \leq \ell \leq 20$. This demonstrates numerical smallness only; it does not establish allowance by the full Planck likelihood.

8. Numerical CMB validation

The CMB program was deliberately staged. The frozen V4 temperature-only calculation established a reproducible baseline. A later temperature-plus-polarization extension introduced empirical noise and internal replication tests.

A subsequent multi-spectrum CMB consistency audit tested acoustic-scale compatibility without converting the deterministic directional template into an isotropic primordial-power modification. The stages are reported separately because they answer different questions and have different statistical strength.

8.1. Einstein–Boltzmann implementation boundary and numerical setup

For the deterministic directional template in Eq. (40), simply changing the isotropic primordial power spectrum $P_\zeta(k)$ would test a different model. The implementation instead applies standard Einstein–Boltzmann transfer functions to the deterministic mode, leaving Einstein gravity and the matter sector unchanged.

CAMB 2.0.4 was run at the Planck 2018 best-fit base- Λ CDM parameters

$$\begin{aligned} H_0 &= 67.32117 \text{ km s}^{-1} \text{ Mpc}^{-1}, & \omega_b &= 0.0223828, & \omega_c &= 0.1201075, \\ \tau &= 0.05430842, & A_s &= 2.100549 \times 10^{-9}, & n_s &= 0.9660499. \end{aligned} \quad (47)$$

with one massive-neutrino mass parameter $m_\nu = 0.06$ eV in the fixed-parameter TT+TE+EE+lensing consistency audit. The reference reproducibility environment used Python 3.12.14, NumPy 2.5.2, CAMB 2.0.4, CLASS/classy 3.3.4, healpy 1.20.0, Astropy 8.0.1, SciPy 1.18.1, pandas 3.0.5, and Matplotlib 3.11.1.

For temperature, the returned transfer functions were checked by reconstructing the ordinary spectrum through

$$C_\ell^{TT} = 4\pi \int d \ln k P_\zeta(k) \left| \Delta_\ell^T(k) \right|^2. \quad (48)$$

Over $2 \leq \ell \leq 29$, the reconstructed/native TT ratio has mean 1.000008, minimum 0.999910, and maximum 1.000094. For the later polarization extension, the CAMB E-mode transfer is converted to the physical EE/TE convention using the spin-2 factor

$$\mathcal{S}_\ell = \sqrt{(\ell+2)(\ell+1)\ell(\ell-1)}. \quad (49)$$

After this correction, the mean reconstruction ratios are 1.000008 (TT), 1.000027 (EE), and 1.000011 (TE), with mean absolute fractional errors 5.87×10^{-5} , 9.82×10^{-5} , and 1.02×10^{-4} , respectively.

CLASS 3.3.4 provides an independent low- ℓ baseline check. Over $2 \leq \ell \leq 29$, the raw CLASS/CAMB TT ratio has mean 1.000740, minimum 1.000324,

and maximum 1.002070. CLASS gives $\chi(z = 1089.8) = 13869.658$ Mpc, differing fractionally by 7.8×10^{-6} from the $\chi_* = 13869.55$ Mpc used in the scan. These residual differences are negligible for the present low- ℓ diagnostic.

8.2. Frozen Planck PR3 temperature-only baseline

The deterministic temperature template was evaluated with the exact CAMB $\Delta_\ell^T(k)$ for $2 \leq \ell \leq 29$, scanning

$$0.5 \leq x \equiv k_I \chi_* \leq 15, \quad \Delta x = 0.5, \quad (50)$$

together with phase and a HEALPix $N_{\text{side}} = 16$ direction grid. The primary diagnostic used the public Planck PR3 inpainted SMICA temperature map and a full-sky Gaussian harmonic covariance constructed from the Planck best-fit spectrum.

The raw maximum-likelihood improvement relative to $A_I = 0$ is

$$x = 14.5, \quad A_I = 2.006 \times 10^{-5}, \quad \Delta\chi_{\text{raw}}^2 = -17.7453, \quad (51)$$

with one equivalent Galactic-axis representation $(l, b) = (182.8125^\circ, 41.8103^\circ)$ and phase $\phi_I = 2.5782$ rad. Because scale, phase, and direction were scanned, the raw improvement cannot be interpreted with a fixed-template χ^2 distribution.

A matched look-elsewhere calibration was therefore performed using 2000 Gaussian Λ CDM realizations drawn from the same null covariance and scanned over the same search grid. The global probability of obtaining a best improvement at least as large as the SMICA value is

$$p_{\text{global}} = 0.45427 \pm 0.01113, \quad (52)$$

with null median best-scan $\Delta\chi^2 = -17.396$. The null hypothesis is therefore not rejected.

The raw feature is robust to Planck temperature component separation in the limited technical sense shown in Table 2. This argues against creation by a single component-separation pipeline, but the look-elsewhere result shows that such reproducibility does not make the feature statistically anomalous relative to ordinary low- ℓ cosmic variance.

The TT-only result is therefore classified as compatible but non-confirmatory.

Table 2: Frozen TT-only Planck PR3 component-map diagnostic. These are full-sky Gaussian-harmonic fits, not official Planck likelihood values.

Map	x_{best}	$A_{\text{I,best}}$	$\Delta\chi^2_{\text{raw}}$	(l, b) [deg]
SMICA	14.5	2.006×10^{-5}	-17.75	(182.8, 41.8)
NILC	14.5	1.994×10^{-5}	-17.52	(182.8, 41.8)
SEVEM	15.0	1.971×10^{-5}	-18.18	(182.8, 41.8)
Commander	14.5	2.021×10^{-5}	-18.01	(182.8, 41.8)

8.3. Corrected noise-aware T+E full-map diagnostic

The first T+E implementation exposed an E-mode normalization error. The result from that implementation is not used. After applying Eq. (49), the analysis incorporated an empirical half-mission half-difference (HMHD) noise estimate. The median estimated noise-to-CMB ratio is approximately 2.30×10^{-4} in TT but 5.34 in EE, making polarization-noise treatment essential.

With 10,000 matched Gaussian CMB-plus-noise simulations, the corrected full-map T+E result is

$$\Delta\chi^2 = -24.9902, \quad x = 11.0, \quad A_{\text{I}} = 2.700 \times 10^{-5}, \quad (l, b) = (315.0^\circ, 38.68^\circ), \quad (53)$$

with

$$p_{\text{full}} = 0.02390 \pm 0.00153. \quad (54)$$

This is an exploratory approximate matched-null result, not an official Planck lowE significance. The apparently low full-map p -value is therefore tested adversarially rather than interpreted as evidence.

8.4. Half-mission internal replication

Planck half-mission maps observe the same sky while providing partly independent instrumental noise and systematics. HM1 selects $x = 11$ and approximately $(315^\circ, 38.7^\circ)$, whereas HM2 selects $x = 15$ and approximately $(182.8^\circ, 41.8^\circ)$. Locking or refitting the template across halves produces ordinary matched-null statistics. The 10,000-simulation symmetric cross-refit probability is

$$p_{\text{cross,refit}} = 0.3565, \quad (55)$$

while a symmetric fixed-template calculation gives approximately $p = 0.404$. Thus the particular full-map directional feature does not internally replicate

and cannot presently be distinguished from chance, noise, or systematics by this test.

8.5. Multipole localization

A multipole jackknife shows that the fitted improvement is highly concentrated at the lowest multipoles. For the fixed full-map reference template, $\ell = 2\text{--}9$ contributes approximately $\Delta\chi^2 = -20.26$, $\ell = 10\text{--}19$ contributes -4.73 , and $\ell = 20\text{--}29$ contributes essentially zero. The $\ell = 2\text{--}14$ block alone recovers $x = 11$, the same $(315^\circ, 38.68^\circ)$ direction, and $\Delta\chi^2 \simeq -24.783$. The disjoint fixed contributions sum to the full -24.9902 result.

This localization is qualitatively consistent with the prior ICH expectation that a relevant primordial imprint should be an ultra-large-scale/low- ℓ phenomenon. However, the exact boundary $\ell = 14$ was identified from these Planck data and is not described as an exact preregistered cutoff. Free scans in isolated bands are not independent detections because they reuse the same sky and the reference was selected on the full data.

8.6. T+E component-separation robustness

A separate T+E component robustness run uses a common CAMB CMB covariance plus the SMICA HMHD empirical noise covariance for comparison only. The best-fit structures are summarized in Table 3. SMICA and NILC are close in scale and direction; Commander remains broadly nearby but shifts in latitude and scale; SEVEM selects the approximately opposite TT-like direction and is an outlier. These maps share the same underlying Planck sky and are not independent replications; no component-specific significance is assigned.

Table 3: Corrected T+E component-separation robustness diagnostic. The common weighting is used only for cross-component comparison.

Map	x_{best}	$A_{\text{I,best}}$	$\Delta\chi^2$	(l, b) [deg]
SMICA	10.5	2.460×10^{-5}	-20.12	(315.0, 32.8)
NILC	10.5	2.410×10^{-5}	-19.49	(315.0, 38.7)
SEVEM	15.0	2.191×10^{-5}	-25.17	(182.8, 41.8)
Commander	9.5	2.668×10^{-5}	-21.19	(320.6, 17.0)

8.7. Corrected polarization-mask robustness

An audit of the preliminary masked implementation found unequal pre-mask spectral support: observed maps had been low-passed through $\ell = 47$ while null simulations contained support only through $\ell = 14$. Because masking couples multipoles, those preliminary values are superseded. The corrected calibration gives both observed maps and matched-null simulations common pre-mask support through $\ell = 29$, while the analyzed statistic remains restricted to the data-localized band $\ell = 2\text{--}14$.

The SMICA polarization confidence mask was downgraded by pixel averaging to $N_{\text{side}} = 16$ and thresholded at 0.80, 0.90, and 0.95. For each threshold, 2000 matched null simulations were transformed through the same map-mask-harmonic procedure before truncation to $\ell = 2\text{--}14$. Results are shown in Table 4.

Table 4: Corrected T+E polarization-mask robustness. f_{sky} is the low-resolution retained sky fraction.

Mask threshold	f_{sky}	x	$\Delta\chi^2$	Direction (l, b) [deg]	Global p
0.80	0.8975	10.5	-18.914	(315.0, 38.68)	0.10295 ± 0.00679
0.90	0.8551	10.5	-16.798	(315.0, 38.68)	0.15242 ± 0.00804
0.95	0.8154	15.0	-15.693	(182.81, 41.81)	0.17991 ± 0.00859

The unmasked $p \simeq 0.024$ feature is therefore not robust under corrected masking.

8.8. Corrected masked half-mission cross-validation

The masked split analysis uses the same corrected pre-mask support through $\ell = 29$, with a shared Gaussian fiducial CMB and independent empirical half-mission noise for the two halves. Each half is transformed map $\rightarrow$ mask $\rightarrow a_{\ell m}$ and then truncated to $\ell = 2\text{--}14$ before the complete train/test scan. HM1 repeatedly selects approximately $(315.0^\circ, 38.68^\circ)$ with $x = 10.5$, whereas HM2 selects approximately $(182.81^\circ, 41.81^\circ)$ with $x = 15$.

These results provide no robust internal replication of a cosmological preferred direction.

8.9. Planck PR3 fixed-parameter $TT+TE+EE+lensing$ consistency and acoustic-scale audit

A separate audit combines the frozen V4 cosmology with the standard Planck PR3 power-spectrum likelihood baseline and CAMB propagation of

Table 5: Corrected masked half-mission cross-validation. All values are ordinary under the matched null.

Mask threshold	Cross-refit global p	Cross-fixed global p
0.80	0.51074 ± 0.01117	0.54173 ± 0.01114
0.90	0.61319 ± 0.01089	0.64468 ± 0.01070
0.95	0.61419 ± 0.01088	0.65567 ± 0.01062

the deterministic standing-wave template. The purpose is to test acoustic-scale consistency without conflating a deterministic directional mean with an isotropic covariance modification.

Using Cobaya 3.6.2 with CAMB 2.0.4, the fixed cosmology in Eq. (47) plus $m_\nu = 0.06$ eV and $A_{\text{Planck}} = 1$ was evaluated with the likelihood components `planck_2018_lowl.TT`, `planck_2018_lowl.EE`, `planck_2018_highl_plik`, `.TTTEEE_lite_native`, and `planck_2018_lensing.native`. The returned likelihood-specific χ^2 components are listed in Table 6. These are not reduced χ^2 values and are not an ICH-versus- Λ CDMsignificance; they establish the precision baseline at the frozen cosmological point.

Table 6: Planck PR3 likelihood components at the frozen baseline cosmology used for the fixed-parameter consistency audit.

Likelihood component	χ^2 contribution
low- ℓ TT	23.229745
low- ℓ EE	396.053460
high- ℓ TTTEEE Plik-lite	585.325211
lensing	8.912272
Total CMB	1013.520687

For each benchmark in Table 7, a phase-independent deterministic envelope was evaluated on the stored CAMB scalar transfer grid. TT and EE were normalized to their corresponding unlensed scalar baseline spectra; TE was normalized to $\sqrt{TT EE}$ to avoid ordinary TE zero crossings. For these ultra-low- k modes, CAMB 2.0.4 stores the sampled transfer for all $\ell \geq 30$ as exactly zero. This is reported only as *below CAMB stored transfer-grid numerical support*; it is not an analytic claim that the physical line-of-sight transfer function is identically zero.

Table 7: Stored CAMB transfer support for the three benchmark standing-wave realizations.

Benchmark	x	A_I	Last ℓ_T	Last ℓ_E
Post-V4 T+E full-map best scan	11.0	2.7×10^{-5}	17	17
V4 illustrative amplitude	10.0	5.0×10^{-6}	17	17
Frozen V4 TT best scan	14.5	2.006×10^{-5}	27	27

An independent spherical-Bessel calculation gives a geometric last-scattering projection sanity check. It is not a formal bound on the complete line-of-sight source integral, but it confirms severe high- ℓ geometric suppression. Table 8 shows representative values.

Table 8: Independent last-scattering projection sanity check. Ratios are relative to the low- ℓ peak of $|j_\ell(x)|^2$ for each benchmark.

Benchmark	x	$ j_{30} $	j_{30}^2/j_{pk}^2	$ j_{40} $	j_{40}^2/j_{pk}^2
T+E best scan	11.0	3.69×10^{-12}	1.03×10^{-21}	3.36×10^{-20}	8.48×10^{-38}
Illustrative	10.0	2.51×10^{-13}	4.00×10^{-24}	8.44×10^{-22}	4.51×10^{-41}
TT best scan	14.5	7.01×10^{-9}	5.70×10^{-15}	1.22×10^{-15}	1.72×10^{-28}

The combined multi-spectrum CMB result is therefore a survival and localization statement: the tested ultra-long modes are supported by the numerical transfer machinery only at the lowest multipoles at the sampled grid points, while the low- ℓ directional candidate itself does not survive the strongest internal robustness tests. No official Planck directional likelihood preference for ICH is claimed.

8.10. CMB interpretation and methodological boundary

The CMB analysis separates prediction consistency from statistical evidence. The fitted T+E structure is concentrated at very low ℓ , as qualitatively expected by ICH, and one candidate direction recurs in several treatments. However, the low p -value of the unmasked full-map scan is not stable under corrected masking and does not replicate across half-mission splits. The corrected mask calibration therefore strengthens the methodological conclusion rather than the signal claim.

The present suite is not the official Planck lowE pixel likelihood and does not use the full Planck end-to-end systematics simulation machinery. The empirical HMHD noise model is approximate, component maps are correlated,

and the exact $\ell = 2\text{--}14$ localization was identified after inspecting the same Planck sky. Cosmic variance is intrinsically severe at these multipoles. A genuinely new or held-out data set is required for independent confirmation. Within the ICH program, the principal discriminating test remains large-scale structure and cross-survey coherence rather than retrospective optimization of the same CMB realization.

9. Large-scale structure, lensing, voids, and rare objects

9.1. Galaxy clustering and redshift space

At leading linear order,

$$\delta_g^s(\mathbf{k}, z) = [b_1(z) + f(z)\mu^2] \delta_m(\mathbf{k}, z) + \epsilon_{\text{stoch}} + \text{relativistic/selection terms.} \quad (56)$$

A survey-facing prediction may be written schematically as

$$\mathbf{d}_{\text{LSS}}^{\text{pred}} = \mathbf{d}_0 + \mathbf{W}_{\text{survey}} \mathbf{M}_{\text{RSD}} \mathbf{T}_\delta \zeta_{\text{I}}(\Theta_{\text{I}}), \quad (57)$$

where $\mathbf{W}_{\text{survey}}$ includes mask, radial selection, coordinate convention, target selection, estimator response, and other window effects. An unwindowed sky template is not directly comparable to a survey statistic.

9.2. Weak lensing and voids

For sources at comoving distance χ_s , define

$$W_L(\chi, \chi_s) = \frac{2}{c^2} \frac{\chi(\chi_s - \chi)}{\chi_s}. \quad (58)$$

In the Born/flat-sky benchmark,

$$\kappa_{\text{I}}(\hat{\mathbf{n}}, \chi_s) = \int_0^{\chi_s} d\chi W_L(\chi, \chi_s) \nabla_\perp^2 \Phi_{W,\text{I}}(\chi \hat{\mathbf{n}}, z), \quad (59)$$

with a corresponding curved-sky generalization [16]. For the standing wave,

$$\kappa_{\text{I}} \propto -A_{\text{I}} k_{\text{I}}^2 (1 - \mu^2) \int d\chi W_L T_W(k_{\text{I}}, z) \cos(k_{\text{I}} \mu \chi + \phi_{\text{I}}). \quad (60)$$

A normalized tomographic template correlation is

$$r_{ij}^{\text{I}} = \frac{\langle \kappa_{\text{I},i} \kappa_{\text{I},j} \rangle}{\sqrt{\langle \kappa_{\text{I},i}^2 \rangle \langle \kappa_{\text{I},j}^2 \rangle}}. \quad (61)$$

For a stacked void observable $V(R, z)$,

$$V(R, z|\delta_L) \simeq V_0(R, z) [1 + R_V(R, z)\delta_L + \dots]. \quad (62)$$

Because the template is explicitly low- k , it does not generate a compact high-frequency lensing clump by itself. Sharp localized features attributed directly to the imprint would violate the model’s own support assumptions. The natural lensing target is a low-amplitude, spatially extended residual or a coherent cross-survey/tomographic mode, not a new population of compact peaks. Candidates must survive shear-calibration, mask, photometric-redshift, point-spread-function, and reconstruction null tests [24–26].

For void applications, the same low- k restriction implies smooth stacked residuals rather than sharp profile features [27–29]. A True Void is faint because it is genuinely low in inferred mass as well as luminosity, whereas an optically faint region with high lensing/clustering mass and node-like topology is not a void and belongs to the Dark-Node test described below. This mass–luminosity–topology distinction prevents optical emptiness from being confused with gravitational underdensity.

9.3. Rare objects and early growth

For halo abundance,

$$n_h(M, z|\delta_L) = \bar{n}_h(M, z) \left[1 + b_1^h \delta_L + \frac{1}{2} b_2^h \delta_L^2 + \dots \right]. \quad (63)$$

Large rare-halo bias allows a small long-mode modulation to produce a larger fractional effect in the extreme tail without requiring a large global change to structure growth [10]. This is the mathematical route to the ICH rare-early-object prediction class; it does not replace ordinary halo, gas, accretion, or selection physics. The null must include ordinary rare-tail statistics, halo bias, baryonic uncertainties, lensing magnification, stellar-population systematics, and selection effects [32–34].

9.4. Distance-related inference without a new expansion law

ICH does not replace the homogeneous Friedmann relation. At background order,

$$H_{\text{ICH}}(a) = H_{\text{base}}(a), \quad D_{\text{ICH}}^{(0)}(z) = D_{\text{base}}^{(0)}(z). \quad (64)$$

Any effect in distance-sensitive inference must enter through standard perturbative line-of-sight terms, environmental covariance, peculiar velocities,

lensing magnification, or survey selection rather than through a new zeroth-order expansion law:

$$D_{\text{obs}}(z, \hat{\mathbf{n}}) = D_{\text{base}}^{(0)}(z) + \delta D_{\text{std}}(z, \hat{\mathbf{n}}; \zeta_{\text{G}}) + \delta D_{\text{I}}(z, \hat{\mathbf{n}}; \Theta_{\text{I}}), \quad (65)$$

with $|\delta D_{\text{I}}| \ll D_{\text{base}}^{(0)}$. A robust requirement for a genuinely modified background expansion would therefore not support ICH as defined here; it would point to a different cosmological model.

10. Cosmic-web node environments and the population layer

10.1. Independent node definition

The node-environment extension is an observational hierarchy, not a new source term. To avoid circularity, the cosmic-web environment must be reconstructed independently of the external catalog used to test quasar, H I, black-hole, or gas occupation. A smoothed tidal/Hessian tensor can be defined as

$$T_{ij}(\mathbf{x}; R_s) = \partial_i \partial_j \nabla^{-2} \delta_R(\mathbf{x}), \quad (66)$$

with node classification based on eigenvalues/topology or a persistent-topology method such as DisPerSE [17]. Threshold, smoothing scale, and mass/connectivity ranking must be fixed without knowledge of which nodes host the external target objects.

A bounded schematic response for node membership may be written

$$\text{logit } P(N = 1 | \Theta_{\text{I}}) = \alpha_N + R_N \delta_L^{\text{I}} + R_N^K \mathcal{K}_{ij}^{\text{I}} \mathcal{K}_1^{ij} + \dots. \quad (67)$$

10.2. P11: Radiant Phase

The Radiant Phase is the luminous active face of the node program: an extreme quasar or active ultra-massive-black-hole candidate associated with a massive node environment. For an indicator Y_Q ,

$$\log \frac{P(Y_Q = 1)}{1 - P(Y_Q = 1)} = \alpha_Q + \beta_Q \cdot \mathbf{C}_Q + \beta_N N + \beta_{\text{I},Q} q_{\text{I}}(\mathbf{x}; \Theta_{\text{I}}). \quad (68)$$

The ICH-specific test is a residual $\beta_{\text{I},Q} > 0$ after predefined controls. The node field must be defined independently before testing quasar occupation.

10.3. P12: Veiled Phase and Dark Nodes

The Veiled Phase is an astrophysical/evolutionary state: a faded, dormant, obscured, or low-accretion remnant candidate. A Dark Node is an environment/class: a massive but underluminous cosmic-web node with high gravitational/topological rank relative to its luminosity. The two are not synonyms, and a Veiled object need not live in a Dark Node. Let Y_V represent the operational Veiled class and DN a Dark-Node indicator. Then

$$\log \frac{P(Y_V = 1)}{1 - P(Y_V = 1)} = \alpha_V + \beta_V \cdot \mathbf{C}_V + \beta_D DN + \beta_{I,V} q_I. \quad (69)$$

GR supplies the geometry/environment layer. Fading, dormancy, obscuration, accretion state, and H I content require ordinary astrophysical duty-cycle and selection models.

The Radiant–Veiled relation is evolutionary motivation, not a new field equation. If the luminous duty cycle is short compared with the time spent in low-accretion states, a conditional expectation is

$$N_{\text{Veiled}} \gg N_{\text{Radiant}}, \quad (70)$$

subject to redshift, selection depth, obscuration, gas supply, and duty-cycle completeness. No analogous abundance inequality is asserted for Dark Nodes themselves.

Operationally, a minimal Dark-Node candidate requires both (i) node-like topology or high filament connectivity and (ii) high lensing/clustering/richness/topology-inferred mass relative to electromagnetic luminosity. H I richness, low surface brightness, weak AGN emission, or optical faintness are supporting tracers only; none is sufficient without the mass–topology criterion.

10.4. P13: Baryonic Trace

For a baryonic observable B ,

$$B_{\text{obs}}(\mathbf{x}) = B_0(\mathbf{x}; \boldsymbol{\eta}_B) + R_B \delta_L^I + R_B^K \mathcal{K}_{ij}^I \mathcal{K}_I^{ij} + \epsilon_B. \quad (71)$$

The nuisance vector $\boldsymbol{\eta}_B$ contains ordinary gas physics, feedback, merger history, calibration, and selection. P13 is a residual topology-linked test after these controls, not a claim that GR alone predicts gas thermodynamics.

Relevant tracers may include X-ray emission, thermal/kinetic Sunyaev–Zel’dovich observables, UV absorption, FRB dispersion measures, Ly α nebulae, molecular or ionized gas, H I structures, radio jets/lobes, fossil cavities, metallicity, low-surface-brightness material, star-forming filaments, and future warm-hot intergalactic-medium reconstructions [35–37]. These are ordinary baryons. The geometric hypothesis enters only through the residual association with independently defined node environments after halo mass, gas physics, merger history, AGN feedback, redshift, and selection are controlled.

10.5. Combined node-environment logic and controls

P11–P13 are not three independent anomaly baskets. The strongest node-environment outcome is a population-level pattern in which (i) massive luminous nodes show a residual enhancement of extreme active systems, (ii) massive nodes contain a residual faded/dormant/obscured counterpart population, (iii) a high-contrast subset is massive but underluminous in the Dark-Node sense, and (iv) node-connected regions show feeding, feedback, depletion, heating, expulsion, inflow, or redistribution traces beyond ordinary expectations. The four signatures need not occur in the same individual node; the coherence is statistical across the independently reconstructed node population.

Illustrative cluster-control examples and toy counting thresholds are collected in Appendix D. They are not used to define the node classes or to set thresholds for the prospective test.

11. Pre-registered predictions P1–P13

The Euclid-testable prediction record was staged before the decisive Euclid DR1 test: P1–P10 were prepared on 17 January 2026, P11 was added on 1 February 2026, and P12–P13 were added on 30 April 2026 [1]. The exact archival title, version, and filename are given in that reference. The chronology is preserved rather than retroactively presenting all thirteen as simultaneous predictions. Table 9 gives the operational mapping used in the present formulation.

Table 9: ICH P1–P13 operational mapping. Outcome classification is given separately in Appendix C.

P	Prediction	Primary quantitative channel
P1	Non-evolving residual patterns	Tomographic residual-template coherence after standard evolution/projection
P2	Subdominant amplitude	Posterior on A_I and residual amplitudes
P3	Coherent weak-lensing residual modes	E-mode shear, low- ℓ modes, lensing–LSS cross-correlation
P4	Cross-survey persistence	Common- Θ_I joint likelihood/posterior overlap
P5	No distance–redshift modification	Standard background distances with perturbative residuals only
P6	Smooth void lensing	Void response and cross-survey stacks
P7	Cluster lensing: no new localized clumps	Low- k support plus lensing operator
P8	No high-frequency/small-scale imprint	$W_I(k > k_{\text{cut}}) = 0$
P9	Rare localized early-growth outliers	Biased rare-halo response, Eq. (63)
P10	Faint anisotropy within precision constraints	Oriented CMB/LSS template under isotropy bounds
P11	Radiant Phase	Independent node reconstruction plus extreme-quasar occupation, Eq. (68)
P12	Veiled Phase and Dark Nodes	Veiled occupation/Dark-Node contrast, Eq. (69)
P13	Baryonic Trace	Residual baryonic response, Eq. (71)

The strongest support would not be the accumulation of individually suggestive examples. It would be a common low- k realization compatible with precision CMB constraints that then predicts independent LSS, lensing, void, and node-environment features without refitting direction, phase, or scale to each data set.

12. Joint inference, model comparison, and falsification

Let

$$\mathcal{D} = \{D_{\text{CMB}}, D_{\text{lens}}, D_{\text{LSS}}, D_{\text{void}}, D_{\text{node}}, D_{\text{pop}}, D_{\text{baryon}}\}. \quad (72)$$

A master forward distribution can be written

$$P(\mathcal{D}|\Theta_{\text{I}}, \Theta_{\text{cos}}, \{\eta_s\}) = \mathcal{L}_{\text{joint}}[\mathcal{D} | \boldsymbol{\mu}(\Theta_{\text{I}}, \Theta_{\text{cos}}, \boldsymbol{\eta}), \boldsymbol{C}(\Theta_{\text{I}}, \Theta_{\text{cos}}, \boldsymbol{\eta})]. \quad (73)$$

Channel means contain the relevant transfer/projection/window operators acting on $\zeta_{\text{G}} + \zeta_{\text{I}}(\Theta_{\text{I}})$. Instrumental and astrophysical nuisances may be channel-specific, but Θ_{I} is global. Factorization into independent likelihoods is justified only when cross-channel covariances are negligible or explicitly conditioned away.

A Bayesian comparison uses

$$B_{10} = \frac{Z_1}{Z_0}, \quad Z_i = \int d\theta_i \mathcal{L}(\mathcal{D}|\theta_i, H_i) \pi(\theta_i|H_i), \quad (74)$$

while a frequentist implementation may use a likelihood-ratio statistic calibrated on complete H_0 mocks that repeat the scan over direction, phase, scale, and bandwidth. A useful cross-survey compatibility condition is

$$\bigcap_s \mathcal{C}_s(\Theta_{\text{I}}) \neq \emptyset. \quad (75)$$

No common region means the unified realization fails even if isolated channels look individually suggestive.

The standard limit is

$$A_{\text{I}} \rightarrow 0 \implies \zeta_{\text{I}} \rightarrow 0, \quad {}^{(3)}R_{\text{I}} \rightarrow 0, \quad a_{\ell m}^{\text{I}} \rightarrow 0, \quad \delta_m^{\text{I}} \rightarrow 0, \quad H_1 \rightarrow H_0. \quad (76)$$

A claimed ICH realization fails, or an individual prediction route fails, if the signal disappears after gauge-removable pieces and complete projection terms are handled; if it requires new gravity, extra stress–energy, or localized imprint mass; if high- k power is required; if CMB and late-time data admit no common Θ_{I} ; if survey-window or selection effects remove the signal; if P11–P13 vanish after independent node reconstruction and matched controls; or if the restricted template gains no model-comparison advantage over ordinary ΛCDM cosmic variance after prior-volume and look-elsewhere penalties.

The boundaries are part of the model definition. Weak amplitude requires the imprint contribution to remain subdominant to the baseline signal. Ultra-large-scale support requires scientifically relevant template power to remain below the adopted k_{cut} and not migrate into halo-scale structure. Smoothness forbids interpreting the imprint itself as a compact lensing or mass clump.

Frozen-imprint persistence permits redshift evolution of tracers, growth, and projection kernels but not arbitrary intrinsic sign or morphology changes of the underlying coherent pattern. Background preservation forbids rescuing the fit with a new homogeneous expansion law. Environmental selectivity requires topology-conditioned residuals rather than a uniform sky-wide enhancement.

A useful redshift-coherence diagnostic is to write a normalized structural template $S_X(z)$ for observable channel X and test whether

$$\text{Corr}[S_X(z_i), S_X(z_j)] \quad (77)$$

remains consistent with the evolution predicted by the same Θ_1 after standard transfer and selection effects. P1 does not require numerical constancy of an observable; it requires persistence of the underlying large-scale organization once ordinary redshift-dependent response is included.

13. Selected observational context

This section does not redefine the model and is deliberately limited to observations that are compatible with, useful for calibrating, or directly constraining the preregistered tests. None of the examples below is counted as a detection of ICH.

13.1. *Planck: precision boundary, not supporting evidence*

The integrated analysis in Sec. 8 sharpens the earlier Planck interpretation. The frozen TT-only preference is statistically ordinary after the look-elsewhere correction. The corrected T+E full-map candidate has a lower nominal matched-null probability, but the feature is not robust to corrected polarization masking and fails half-mission replication. The multi-spectrum CMB audit further shows that the tested ultra-long modes are strongly localized to low multipoles in CAMB's stored transfer representation. The resulting statement is therefore compatibility plus a precision boundary, not CMB evidence for ICH.

13.2. *Euclid high-redshift quasars: compatible, non-confirmatory*

Euclid reported 31 new quasars at $6.6 < z < 7.8$ from approximately 3000 deg^2 of early Wide Survey coverage, including 12 at $z \geq 7$ [18]. Their existence is compatible with the rare-object/environmental direction of the ICH program and enlarges the target sample relevant to P11. It is not a P11

test by itself. The preregistered test requires an independently reconstructed Euclid node field followed by controlled extreme-quasar occupation or cross-correlation. Using quasar locations to define the nodes would reverse the conditioning and create selection circularity.

13.3. *JWST/ASPIRE: qualitatively consistent, non-discriminating*

ASPIRE studied 25 luminous quasars at $z > 6.5$, identifying 487 [O III] emitters at $5.3 \lesssim z \lesssim 7.0$, of which 122 lie within $|\Delta v_{\text{los}}| < 1000 \text{ km s}^{-1}$ of the quasars; the reported enhancement is about 9.4 relative to the average density at other redshifts, with seven quasars tracing significant galaxy overdensities [19]. This is qualitatively consistent with the environmental direction of the Radiant-Phase prediction but is not discriminating evidence for ICH because the observations are quasar-centered and standard massive-halo/environment physics is already part of the null.

13.4. *FAST/FASHI dark-galaxy candidates: compatible tracer context*

The FASHI/FAST analysis of Monaci et al. compiled 70 dark-galaxy candidates within 50 Mpc under its catalog definition [20]. Such H I-rich or optically faint objects are compatible as external tracers for P12/P13, but the catalog is not a Dark-Node catalog in the ICH sense. A Dark Node requires independently inferred high gravitational/topological node rank relative to luminosity.

13.5. *Cluster systems as controls*

Complex merging clusters can show substantial separation between gas, galaxies, lensing mass, and projected structure. Such systems are useful controls for the operational Dark-Node criteria, but they are not evidence that ICH creates an additional compact mass component. In particular, the Bullet Cluster is treated as a control on mass-gas offsets in ordinary gravitational structure, consistent with P7’s requirement that the imprint itself not behave as a localized clump [21].

13.6. *Gravitational-wave memory: analogy only*

Gravitational-wave memory illustrates that GR admits persistent changes in geometric observables after transient dynamics [22, 23]. The analogy is conceptually useful but physically distinct:

$$\Delta h_{ij}^{\text{memory}} \neq 0 \not\Rightarrow \zeta_{\text{I}} \neq 0, \quad \zeta_{\text{I}} \neq 0 \not\Rightarrow \Delta h_{ij}^{\text{memory}} \neq 0. \quad (78)$$

No gravitational-wave statistic enters the ICH likelihood, and memory is not presented as empirical support.

14. Discussion and present limitations

14.1. What is established mathematically

The scalar backbone has four connected pieces. The ICH contribution is identified with a restricted part of standard curvature initial data rather than an extra field. Equations (15)–(20) give an explicit family satisfying the ADM constraints. The adiabatic matter completion connects that family to ordinary Einstein–Boltzmann evolution. The forward equations then provide a continuous chain from the initial template to observable likelihoods. This removes the objection that ICH is only a verbal curvature analogy. It does not establish that nature selects the restricted initial-condition family.

14.2. What is an ansatz rather than a GR theorem

The restricted low- k prior itself is the hypothesis. GR permits many initial data; it does not derive the preferred direction, phase, wavelength, bandwidth, or amplitude. Likewise, the node-environment occupation equations are phenomenological statistical models. GR supplies geometry and ordinary evolution, while quasar duty cycles, obscuration, gas supply, feedback, star formation, and survey selection must be modeled astrophysically.

14.3. The ordinary-long-mode degeneracy

A sufficiently general ordinary stochastic realization can contain unusual low- k features. ICH acquires independent content only if its restricted template compresses multiple observations more effectively than the standard ensemble after the full penalty for choosing that family. A strong outcome would leave a nonzero CMB-allowed region of Θ_I that, without re-optimization, predicts a weak directional/phase-coherent signal in independent LSS/lensing statistics and population residuals in independently defined node environments. Conversely, if Planck-allowed parameters predict nothing measurable at late times, if late-time residuals demand a different direction or scale, or if the evidence disappears when the template scan is repeated in null mocks, the unified realization fails.

14.4. What the completed CMB sequence establishes

The numerical CMB work now goes beyond the original V4 diagnostic. Exact CAMB temperature and polarization transfers are normalized and independently checked against native spectra, CLASS provides a low- ℓ baseline cross-check, the frozen TT-only scan is calibrated with matched look-elsewhere simulations, the T+E extension includes empirical polarization noise, half-mission and component-map robustness have been tested, the mask calibration has been corrected for mode-coupling support, masked half-mission cross-validation is complete, and a Planck fixed-parameter TT+TE+EE+lensing consistency baseline plus high- ℓ transfer audit has been performed.

The resulting scientific message is more restrictive than the unmasked T+E p -value alone: Planck contains an exploratory recurrent low- ℓ directional structure whose scale localization is qualitatively consistent with the ICH low- ℓ expectation, but the significance is not robust to corrected masks and the feature fails internal half-mission replication. It is therefore a prediction-consistent candidate feature, not CMB evidence for ICH.

14.5. Remaining precision-level extensions

The CMB validation sequence reported in this manuscript is complete at the level claimed here: it includes the diagnostic low- ℓ directional analysis, matched look-elsewhere calibration, polarization extension, robustness and internal replication tests, and a fixed-parameter TT+TE+EE+lensing consistency audit with a high- ℓ transfer check. What has not been attempted is an official Planck-likelihood precision parameter inference. In particular, the present work does not implement the official Planck lowE pixel likelihood with the full Planck end-to-end systematics simulation machinery, nor a full cosmological/nuisance posterior for the deterministic directional template. The empirical HMHD noise model remains approximate and does not fully represent spatially anisotropic Planck noise/systematics. These are precision-level extensions beyond the completed diagnostic validation, not missing steps in the numerical tests reported above.

The principal prospective tests are instead in independent late-time data. Quantitative LSS and weak-lensing confrontation requires actual survey masks, selection functions, estimators, photometric-redshift uncertainties, shape calibration, and realistic null mocks. The node program requires a frozen node classifier and matched controls; P13 additionally requires gas/feedback calibration or nuisance marginalization. These steps do not require modified

gravity, although sufficiently realistic nonlinear and hydrodynamic mock suites may eventually become HPC-scale tasks.

For the node program, progress means catalog-level definitions fixed without inspecting the external target population: node smoothing/topology threshold, node mass or connectivity rank, quasar/extreme-object criteria, Veiled-state criteria, the Dark-Node mass-to-light statistic, and baryonic-trace selections. Famous objects may illustrate a definition but must not set the threshold used to test it.

14.6. Interpretation if the restricted model survives

If the model survives both precision-CMB constraints and independent late-time tests, the appropriate conclusion remains narrow: an observationally useful part of the realized cosmological initial data may be compressible as a weak coherent low- k geometric component while the dynamical laws remain those of ordinary GR. Such a result would not by itself identify a microphysical origin, replace standard matter physics, or prove that every cosmological tension has a common cause. Its significance would be the cross-survey predictive value of a restricted initial-condition realization. Conversely, internal mathematical consistency without predictive advantage over the ordinary stochastic ensemble is not sufficient evidence for a distinct ICH component.

15. Conclusion

The Imprinted Curvature Hypothesis can be posed as a mathematically explicit, constraint-compatible, gauge-disciplined initial-condition model within unmodified General Relativity. The minimal construction introduces a weak restricted low- k curvature component ζ_I on a physically specified early hypersurface, embeds it in explicit ADM-constraint-satisfying initial data, completes the matter sector with standard adiabatic perturbations, and evolves the result using ordinary Einstein–Boltzmann transfer physics. Constant and pure-gradient adiabatic pieces are excluded from local physical claims; observable content begins with finite curvature/tidal structure and complete projection effects.

The same global parameters are required to propagate into CMB, LSS, RSD, weak lensing, voids, rare-object abundance, and independently reconstructed cosmic-web environments. The Radiant Phase, Veiled Phase, Dark Node, and Baryonic Trace are operational population classes or states, not

new physical sectors. P1–P13 remain linked by the weak, smooth, large-scale, background-preserving character of the hypothesis and by their staged preregistration chronology.

The integrated Planck analysis imposes a particularly important boundary. The frozen TT-only scan is ordinary under the matched Λ CDMnull. The corrected unmasked T+E scan finds a recurrent low- ℓ candidate with nominal $p \simeq 0.024$, but this significance is not stable under corrected polarization masks and the candidate fails half-mission replication. The multi-spectrum CMB audit shows that the tested ultra-long modes are strongly localized to the lowest multipoles in the stored CAMB transfer representation. Taken together, these results establish neither a preferred cosmological direction nor evidence favoring ICH over Λ CDM.

The decisive empirical question is therefore prospective and cross-survey: does a restricted coherent low- k realization survive precision CMB constraints and then predict independent late-time residuals with enough shared direction, scale, phase, and morphology to outperform an ordinary Gaussian adiabatic realization after all survey and model-complexity penalties? If not, the unified ICH realization fails. If a nonzero common region survives and predicts the preregistered signatures without per-survey retuning, the framework gains empirical support while remaining inside standard GR.

Appendix A. Bridge to the original $\epsilon\Phi$ notation

Combining Eqs. (6) and (18) on a flat reference background gives

$$\epsilon\Phi(\mathbf{x}) = -\frac{4}{a_*^2}\nabla^2\zeta_{\text{I}}(\mathbf{x}). \quad (\text{A.1})$$

For a single Fourier mode, $\epsilon\Phi \propto 4k_{\text{I}}^2 A_{\text{I}}/a_*^2$. The two parameterizations encode the same scalar initial-curvature content at linear order, but ζ_{I} is more convenient for standard cosmological transfer functions while $\epsilon\Phi$ retains the original geometric definition.

Appendix B. Constraint algebra, curvature decomposition, and standard-limit checks

The four-dimensional Riemann tensor admits the standard decomposition

$$R_{abcd} = C_{abcd} + \frac{1}{2}(g_{ac}R_{bd} + g_{bd}R_{ac} - g_{bc}R_{ad} - g_{ad}R_{bc}) - \frac{R}{6}(g_{ac}g_{bd} - g_{ad}g_{bc}). \quad (\text{B.1})$$

Exact FLRW has $C_{abcd} = 0$. This identity does not imply that a scalar modulation of ${}^{(3)}R$ automatically produces $C_{abcd} \neq 0$; that conclusion requires the full spacetime initial data.

For the constructive data of Eqs. (15)–(19),

$$K^2 - K_{ij}K^{ij} = 9H_*^2 - 3H_*^2 = 6H_*^2. \quad (\text{B.2})$$

Substitution into Eq. (13) gives

$${}^{(3)}R + 6H_*^2 = 16\pi GE + 2\Lambda, \quad (\text{B.3})$$

and hence Eq. (19). For $K^j_i = -H_*\delta_i^j$ and spatially constant H_* , $D_j K^j_i - D_i K = 0$, hence $p_i = 0$. The explicit family therefore satisfies both constraints on Σ_* . The limit $A_I \rightarrow 0$ gives $\zeta_I \rightarrow 0$, removes the ICH curvature perturbation, and continuously recovers the standard background initial data.

Appendix C. Prediction outcome matrix

“Supportive” in Table C.10 means compatible with the stated ICH channel after the baseline/null analysis; it does not mean unique evidence for ICH. “Tension” denotes an ambiguous or analysis-sensitive result. “Failure” means that the named prediction route is not supported under the model definition.

Table C.10: Operational outcome matrix for P1–P13.

P	Supportive/ compatible outcome	Tension/ambiguous outcome	Failure of route
P1	Same underlying large-scale morphology remains coherent across redshift after transfer/projection.	Coherence depends strongly on binning or reconstruction.	Intrinsic sign/orientation/morphology changes are required.
P2	Weak residual amplitude; no dominant correction to the baseline.	Posterior approaches the perturbative boundary.	A large/dominant new component is required.

P	Supportive/ compatible outcome	Tension/ambiguous outcome	Failure of route
P3	Low- ℓ lensing residual survives masks, calibration, and independent pipelines.	Present only in selected reconstructions.	Vanishes under routine lensing/systematics null tests.
P4	Compatible shared- Θ_I structure occurs in independent surveys.	Partial overlap only, sensitive to windows.	Survey-preferred templates are mutually inconsistent.
P5	Background distances remain standard; only weak perturbative residuals appear.	One inference pathway shows unresolved environment/calibration dependence.	Viability requires a new homogeneous expansion law.
P6	Smooth coherent void-lensing/environment modulation.	Strong void-finder dependence.	Sharp/localized or intrinsically rapidly evolving void feature is attributed to the imprint.
P7	Cluster-related ICH contribution remains smooth and subdominant.	Weak calibration trend is degenerate with baryons/projection.	The imprint must supply a compact additional mass peak.
P8	Candidate support is confined to low k /low ℓ .	Weak small-scale leakage remains analysis-dependent.	Claimed detection is dominated by small scales/high k .
P9	Extreme early objects show a residual environment-conditioned excess.	Rare-tail excess lacks clear environmental dependence.	A uniform global growth enhancement is required.
P10	Faint preferred-direction signal remains within CMB/LSS isotropy bounds.	Significance is mask/look-elsewhere sensitive.	Strong anisotropic expansion or precision-incompatible anisotropy is required.

P	Supportive/ compatible outcome	Tension/ambiguous outcome	Failure of route
P11	Extreme-quasar tail is enhanced in independently defined massive nodes after controls.	Trend is catalog/ node-definition sensitive.	Residual node enhancement vanishes after controls.
P12	Residual Veiled population and/or genuine high-mass underluminous Dark-Node environments survive controls.	Classification depends strongly on obscuration, mass proxy, or threshold.	Veiled residual vanishes and Dark-Node candidates are low mass or non-node-like.
P13	Node-connected baryonic traces retain a topology-linked residual after gas/ feedback controls.	Association depends strongly on tracer or gas model.	Standard halo/gas/ feedback/selection physics absorbs the residual.

Appendix D. Node-control examples and illustrative counting thresholds

The examples in this appendix are controls and intuition aids rather than evidence for ICH or definitions of the prospective node test. Complex cluster systems are useful because gas, galaxy light, lensing mass, and projected topology can separate for ordinary reasons. Abell 520 illustrates how a proposed high mass-to-light feature can depend strongly on independent lensing reconstruction, while Abell 2744 illustrates a highly connected massive node without thereby becoming an underluminous Dark Node [38–40]. The Bullet Cluster remains a control on merger-induced mass–gas displacement rather than an ICH explanation target [21]. Phoenix-like cool-core systems illustrate the richness of ordinary baryonic feeding/feedback physics that must be included in the P13 null rather than counted as ICH evidence [41–43].

For intuition only, suppose a predefined node sample obeys $K \sim \text{Binomial}(n, p_0)$ under the null. Table D.11 gives exact illustrative one-sided thresholds. Real analysis must replace this toy model

with survey mocks, correlated environments, selection functions, and look-elsewhere treatment.

Table D.11: Illustrative one-sided binomial count thresholds for a predefined node sample; not a forecast for the actual Euclid node catalog.

n	p_0	$K_{3\sigma}$	$K_{5\sigma}$
31	0.10	10	15
31	0.20	14	20
31	0.25	16	22
25	0.10	9	13
25	0.20	13	17
25	0.25	14	19

Appendix E. Equation taxonomy

The equations in this paper fall into four classes: standard GR/cosmological identities and dynamics; a constructive choice of admissible initial data; the ICH ansatz/prior defining the restricted low- k family; and observable mappings, response models, or statistical tests. The distinction is deliberate. No equation adds a new gravitational degree of freedom. The node-environment logistic relations and baryonic-response equations are phenomenological population models, not deductions of gas or black-hole astrophysics from the Einstein equations.

Appendix F. Validation audit trail and reproducibility boundary

The scientific conclusions use only the corrected diagnostics reported in Sec. 8. Two earlier calculations are retained here solely for auditability.

First, an early noise-free T+E calculation returned an apparently extreme $p < 0.0005$. That result is invalid because observed Planck polarization contains substantial noise/systematics while the null simulations omitted polarization noise. It is not used as evidence.

Second, preliminary mask robustness and masked split calculations used unequal pre-mask spectral support: observed maps were low-passed through $\ell = 47$ while null simulations contained support only through $\ell = 14$. Because masking couples multipoles, those preliminary results are superseded by the

corrected common-support-through- $\ell = 29$ calibrations. The superseded full-map mask probabilities were approximately 0.121, 0.140, and 0.151 for thresholds 0.80, 0.90, and 0.95; the corresponding preliminary masked cross-refit probabilities were approximately 0.522, 0.588, and 0.550. None of these preliminary values is used in the current inference.

The frozen TT reproducibility reference uses 2000 null simulations, seed 562026, HEALPix $N_{\text{side}} = 16$, and $x = 0.5, 1.0, \dots, 15.0$. The corrected noise-aware T+E run uses 10,000 matched simulations. The corrected mask and masked-split calibrations use 2000 matched simulations per mask threshold. Machine-readable null distributions, scan summaries, software pins, and checksums are preserved in the archived validation packages. The manuscript contains all numerical quantities required to understand the scientific claims; the archives are reproducibility records rather than hidden scientific dependencies.

Declarations

Funding. No external funding was received for this work.

Competing interests. The author declares no competing interests.

Data availability. The CMB validation uses publicly released Planck PR3 best-fit cosmological products, public PR3 SMICA/NILC/SEVEM/Commander products, and Planck half-mission maps. No new observational data were created. The remaining observational programs discussed in this work require the corresponding public or collaboration survey products.

Code availability. The reproducibility archives preserve the Python workflows used for the CAMB transfer reconstruction, CLASS-CAMB cross-check, frozen Planck low- ℓ template scan, matched Gaussian-null Monte Carlo, corrected noise-aware T+E scan, multipole jackknife, polarization-mask robustness, masked half-mission cross-validation, and the fixed-parameter TT+TE+EE+lensing/high- ℓ audit. The frozen TT reference package records the reference Monte Carlo seed, pinned software environment, SHA256 manifest, Docker configuration, and automated workflow. No custom code is required for the formal ADM derivations.

Author contributions. Single-author work. The author developed the conceptual framework, physical interpretation, prediction structure, and scientific conclusions, and takes responsibility for the final manuscript.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process. During the preparation and revision of this work, the author used OpenAI ChatGPT, including GPT-4o, GPT-5, GPT-5.2, and GPT-5.6 Sol, for language editing, structural organization, mathematical transcription, equation formatting, literature-search assistance, and consistency checking. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the accuracy, integrity, originality, and final content of the manuscript.

Author note. Earlier ICH materials and archival versions are available via Zenodo, DOI: 10.5281/zenodo.17672445. Project website: <https://ich-cosmology.org/>. The preregistered P1–P13 chronology predates the decisive Euclid DR1 test and is documented by the Version 1.2 pre-registration record cited in Sec. 11 [1].

References

- [1] W. Phumdokmai, *Pre-registered Euclid-Testable Predictions for the Imprinted Curvature Hypothesis (ICH)*, Version 1.2, prepared 17 January 2026; P11 added 1 February 2026; P12–P13 added 30 April 2026, archival pre-registration document, file ICH_Euclid_Predictions_PreRegistration_v1.2_2026-04-30.pdf (2026).
- [2] R. Arnowitt, S. Deser, C. W. Misner, The dynamics of general relativity, in: *Gravitation: An Introduction to Current Research*, Wiley, 1962; reprint arXiv:gr-qc/0405109.
- [3] E.ourgoulhon, 3+1 Formalism and Bases of Numerical Relativity, arXiv:gr-qc/0703035 (2007).
- [4] K. A. Malik, D. Wands, Cosmological perturbations, Phys. Rep. 475 (2009) 1–51, arXiv:0809.4944.
- [5] D. Wands, K. A. Malik, D. H. Lyth, A. R. Liddle, A new approach to the evolution of cosmological perturbations on large scales, Phys. Rev. D 62 (2000) 043527, arXiv:astro-ph/0003278.
- [6] S. Weinberg, Adiabatic modes in cosmology, Phys. Rev. D 67 (2003) 123504, arXiv:astro-ph/0302326.

- [7] C.-P. Ma, E. Bertschinger, Cosmological perturbation theory in the synchronous and conformal Newtonian gauges, *Astrophys. J.* 455 (1995) 7–25, arXiv:astro-ph/9506072.
- [8] L. Dai, E. Pajer, F. Schmidt, Conformal Fermi Coordinates, *JCAP* 11 (2015) 043, arXiv:1502.02011.
- [9] A. Barreira, F. Schmidt, Responses in large-scale structure, *JCAP* 06 (2017) 053, arXiv:1703.09212.
- [10] V. Desjacques, D. Jeong, F. Schmidt, Large-scale galaxy bias, *Phys. Rep.* 733 (2018) 1–193, arXiv:1611.09787.
- [11] J. Lesgourgues, The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview, arXiv:1104.2932 (2011).
- [12] D. Blas, J. Lesgourgues, T. Tram, The Cosmic Linear Anisotropy Solving System (CLASS) II: Approximation schemes, *JCAP* 07 (2011) 034, arXiv:1104.2933.
- [13] A. Lewis, A. Challinor, collaborators, CAMB: Code for Anisotropies in the Microwave Background, public code and documentation; see also arXiv:astro-ph/9911177.
- [14] Planck Collaboration, Planck 2018 results. V. CMB power spectra and likelihoods, *Astron. Astrophys.* 641 (2020) A5, arXiv:1907.12875.
- [15] Planck Collaboration, Planck 2018 results. VII. Isotropy and statistics of the CMB, *Astron. Astrophys.* 641 (2020) A7, arXiv:1906.02552.
- [16] A. Lewis, A. Challinor, Weak gravitational lensing of the CMB, *Phys. Rep.* 429 (2006) 1–65, arXiv:astro-ph/0601594.
- [17] T. Sousbie, The persistent cosmic web and its filamentary structure I: Theory and implementation, *Mon. Not. R. Astron. Soc.* 414 (2011) 350–383, arXiv:1009.4015.
- [18] D. Yang et al. (Euclid Collaboration), Euclid: Discovery of 31 new quasars at $6.6 < z < 7.8$, *Astron. Astrophys.* 711 (2026) A104, arXiv:2607.03432, doi:10.1051/0004-6361/202658883.

- [19] F. Wang et al., ASPIRE: The Environments and Dark Matter Halos of Luminous Quasars in the Epoch of Reionization, *Astrophys. J.* 1006 (2026) 39, arXiv:2602.04979, doi:10.3847/1538-4357/ae7bfa.
- [20] M. Monaci et al., FAST and dark: a catalogue of dark galaxy candidates within 50 Mpc, *Mon. Not. R. Astron. Soc.* 548 (2026) stag732, arXiv:2604.14699, doi:10.1093/mnras/stag732.
- [21] D. Clowe et al., A direct empirical proof of the existence of dark matter, *Astrophys. J. Lett.* 648 (2006) L109–L113, arXiv:astro-ph/0608407.
- [22] D. Christodoulou, Nonlinear nature of gravitation and gravitational-wave experiments, *Phys. Rev. Lett.* 67 (1991) 1486–1489, doi:10.1103/PhysRevLett.67.1486.
- [23] M. Favata, The gravitational-wave memory effect, *Class. Quantum Grav.* 27 (2010) 084036, arXiv:1003.3486.
- [24] A. Amon et al., Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration, *Phys. Rev. D* 105 (2022) 023514, arXiv:2105.13543.
- [25] G. Congedo et al. (Euclid Collaboration), Euclid preparation. LIII. LensMC, weak lensing cosmic shear measurement with forward modelling and Markov Chain Monte Carlo sampling, *Astron. Astrophys.* 691 (2024) A319, arXiv:2405.00669.
- [26] C. Heymans et al., KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints, *Astron. Astrophys.* 646 (2021) A140, arXiv:2007.15632.
- [27] Y.-C. Cai, N. Padilla, B. Li, S. Cole, Testing Gravity using Cosmic Voids, *Mon. Not. R. Astron. Soc.* 451 (2015) 1036–1055, arXiv:1410.1510.
- [28] A. Pisani et al., Cosmic voids: a novel probe to shed light on our Universe, arXiv:1903.05161 (2019).
- [29] C. Sánchez et al., Cosmic Voids and Void Lensing in the Dark Energy Survey Science Verification Data, *Mon. Not. R. Astron. Soc.* 465 (2017) 746–759, arXiv:1605.03982.

- [30] N. I. Libeskind et al., Planes of satellite galaxies and the cosmic web, *Mon. Not. R. Astron. Soc.* 452 (2015) 1052–1059, arXiv:1503.05915.
- [31] M. S. Pawlowski, Phase-space correlations among systems of satellite galaxies, *Galaxies* 9 (2021) 66, arXiv:2109.02654.
- [32] M. Boylan-Kolchin, Stress testing Λ CDM with high-redshift galaxy candidates, *Nat. Astron.* 7 (2023) 731–735, arXiv:2208.01611.
- [33] I. Labbé et al., A population of red candidate massive galaxies ~ 600 Myr after the Big Bang, *Nature* 616 (2023) 266–269, arXiv:2207.12446.
- [34] A. C. Carnall et al., The JWST EXCELS survey: Too much, too young, too fast? Ultra-massive quiescent galaxies at $3 < z < 5$, *Mon. Not. R. Astron. Soc.* 534 (2024) 325–347, arXiv:2405.02242.
- [35] F. Nicastro et al., Observations of the missing baryons in the warm-hot intergalactic medium, *Nature* 558 (2018) 406–409, arXiv:1806.08395.
- [36] J. P. Macquart et al., A census of baryons in the Universe from localized fast radio bursts, *Nature* 581 (2020) 391–395.
- [37] C. R. H. Walker et al., The dispersion measure contributions of the cosmic web, *Astron. Astrophys.* 683 (2024) A71, doi:10.1051/0004-6361/202347139.
- [38] A. Mahdavi et al., A Dark Core in Abell 520, *Astrophys. J.* 668 (2007) 806–814, arXiv:0706.3048.
- [39] D. Clowe et al., On Dark Peaks and Missing Mass: A Weak-Lensing Mass Reconstruction of the Merging Cluster System A520, *Astrophys. J.* 758 (2012) 128, arXiv:1209.2143.
- [40] S. Gallo et al., Tracing gaseous filaments connected to galaxy clusters: The case study of Abell 2744, *Astron. Astrophys.* 692 (2024) A200, arXiv:2407.10518.
- [41] M. McDonald et al., Deep Chandra, HST-COS, and Megacam observations of the Phoenix Cluster: extreme star formation and AGN feedback on hundred kiloparsec scales, *Astrophys. J.* 811 (2015) 111, arXiv:1508.05941.

- [42] C. Pinto et al., AGN feedback in the Phoenix cluster, *Mon. Not. R. Astron. Soc.* 480 (2018) 4113–4123, arXiv:1808.02872.
- [43] M. Reece et al., Directly imaging the cooling flow in the Phoenix Cluster, *Nature* 638 (2025) 360–364, arXiv:2502.08619.