Information-Limited Gravity: Source-Side Kernel Tests Against Distances, Growth, and Lensing

Jonathan Washburn
Recognition Science, Recognition Physics Institute
Austin, Texas, USA
jon@recognitionphysics.org

Abstract

We present a fixed-constant, no-fit, source-side modification to gravitational sourcing—the information-limited gravity (ILG) kernel—and confront it with cosmological observables that usually motivate a background dark-energy component. The kernel introduces a mild, scale- and time-dependent weight on the Poisson source while leaving the background field equations unchanged. We carry the kernel through growth, lensing, redshift-space distortions, and optical distances from the null congruence, compute the Buchert backreaction, and isolate two single-plot falsifiers. We also state two covariant extensions that could shift background distances without introducing additional fit parameters.

1 Introduction

The standard narration of late–time cosmic acceleration is written as a *background* effect: a new component with negative effective pressure (Λ) is added to the Friedmann equations to bend $H(a)$ away from the Einstein–de Sitter (EdS) history. In this work we ask a different question: can the principal late–time anomalies that motivate Λ be reproduced by a *source–side* modification that leaves the background field equations unchanged, but alters how inhomogeneities source the Newtonian potential and thus coherently change growth and lensing?

We test a concrete, fixed-constant hypothesis: an *information–limited gravity* (ILG) kernel $w(k, a)$ multiplies the Poisson source so that, in Fourier space

and Newtonian gauge with negligible anisotropic stress,

$$k^2\Phi(\mathbf{k},a) = 4\pi G\, a^2\, \rho_s(a)\, w(k,a)\, \delta_s(\mathbf{k},a), \tag{1}$$

$$w(k,a) = 1 + \varphi^{-3/2}\left[\frac{a}{k\,\tau_0}\right]^{\alpha}. \tag{2}$$

where $\varphi = (1+\sqrt{5})/2$, $\alpha = \frac{1}{2}(1-\varphi^{-1})$, and τ_0 is the eight–tick fundamental time. The subscript s denotes the *sourcing* density; in a classical reading we take $\rho_s = \rho_m$ (so $\Omega_{s0} = \Omega_{m0}$) unless stated otherwise. The kernel is unity on small scales/early times and rises mildly as a fixed power toward large scales/late times. The constants entering (1) are fixed a priori (no data fitting in this work).

Equation (1) implies a scale–aware linear growth. In the matter era the growing mode is

$$D(a,k) \;=\; a\,\Big[1+\beta(k)\,a^{\alpha}\Big]^{\frac{1}{1+\alpha}}, \qquad \beta(k) = \frac{2}{3}\,\varphi^{-3/2}\,(k\tau_0)^{-\alpha}, \tag{3}$$

with growth rate $f(k,a) \equiv d\ln D/d\ln a = 1 + \frac{\alpha}{1+\alpha}\frac{\beta(k)a^{\alpha}}{1+\beta(k)a^{\alpha}}$. The same kernel multiplies the *lensing* source, so that growth and lensing shift together and in a predictably scale–dependent way. In compact form, many kernel predictions become universal functions of the dimensionless combination $X \equiv k\tau_0/a$.

This source–side picture differs from ΛCDM in three decisive ways:

- **Background vs. perturbations.** In the pure ILG hypothesis the background remains FRW with the chosen matter content (we take EdS as the strict RS baseline), while all late–time anomalies live in the *source term* for Φ. The Buchert kinematical backreaction Q_D cancels mode–by–mode for potential flow even when f depends on k; thus ILG does not alter $H(a)$ by averaging. The mean optical distances are therefore those of the chosen background, with lensing entering only at second order in the mean and primarily as a variance.

- **Signatures.** The integrated Sachs–Wolfe (ISW) driver is

  $$\dot\Phi + \dot\Psi = 2aH\,\Phi\big[-1 + f(k,a) + \partial_{\ln a}\ln w(k,a)\big]. \tag{4}$$

  Because $f > 1$ and $\partial_{\ln a}\ln w > 0$ at late time/large scales, the bracket is positive while $\Phi < 0$, yielding a *negative* CMB–galaxy correlation at low multipoles. The CMB lensing potential gains a mild,

scale–dependent enhancement from w^2D^2; on the matter side, $P(k,z)$ is tilted up toward small k with exponent α, producing a gentle rise of $f(k)$ to large scales and a scale–dependent $E_G(k) \propto w/f$ that is tracer–independent.

- **Falsifiers.** On linear scales we consider the set

$$Q(a,k) \in \Big\{ f(a,k),\ w(k,a),\ \frac{w^2(k,a)D^2(a,k)}{a^2} \Big\}, \tag{5}$$

which depends only on the kernel variable $X = k\tau_0/a$. This implies the reciprocity identity

$$\partial_{\ln a} \ln Q(a,k) = -\partial_{\ln k} \ln Q(a,k). \tag{6}$$

At fixed redshift the theory predicts a small, monotone negative slope in $\partial_{\ln k} \ln f$, together with a lensing–amplitude ratio > 1 that rises toward low k. Across redshift in the *same* k bins, the time–slope must mirror the scale–slope. Either single–plot check can falsify the source–side hypothesis.

The consequences for *distances* are immediate and non–negotiable: with $Q_D = 0$ the background Hubble function is unchanged by ILG, and the Sachs focusing correction to the *mean* luminosity distance is a small second–order brightening. Hence the supernova Hubble diagram *cannot* be matched by ILG alone; any agreement would have to arise from a genuine background term or an optical rescaling at the Ricci–focusing level. To keep the theory falsifiable and parameter–minimal, we state two covariant extensions that use only the same RS constants: (i) a background form–factor $w_{\rm bg}(a)$ multiplying the matter source in the Friedmann equations with an early–time gate and a fixed late–time slope tied to α; and (ii) an optical rescaling $\Upsilon(a)$ of the Ricci focusing in the Sachs equation that preserves metricity and Etherington duality while producing a mild late–time demagnification. Both routes leave the early universe intact and are directly testable against SNe, BAO, CMB lensing, and the E_G statistic.

Parameter policy. The derivations above introduce no fit parameters. The only numerical anchors we use are standard external quantities (e.g., H_0 for dimensional distances and the sound horizon r_d for BAO ratios). No per–galaxy or per–halo tuning enters any prediction; the kernel and exponent are fixed once globally.

Roadmap and contributions. We (i) formalize the ILG kernel and growth law and give compact X–based formulas; (ii) solve the null–congruence for

distances with and without ILG lensing and compute the lensing variance; (iii) perform the Buchert audit and prove $Q_D = 0$ for ILG; (iv) derive ISW, CMB–lensing, and matter–power predictions; (v) provide scale–aware RSD and a bias–free $E_G(k)$; (vi) specify an N–body implementation (PM kernel multiplication only) and the resulting halo and lensing diagnostics; (vii) detail Solar–system and laboratory safety; (viii) state the two background routes that could reconcile supernova distances without introducing free parameters; and (ix) present two single–plot falsifiers that separate source–side ILG from background Λ explanations in a model–transparent way.

2 RS kernel and fixed exponent

2.1 Definitions and constants

We adopt the golden ratio

$$\varphi = \frac{1+\sqrt{5}}{2}, \tag{7}$$

and define the dimensionless exponent

$$\alpha = \tfrac{1}{2}\big(1-\varphi^{-1}\big). \tag{8}$$

Let τ_0 denote the fundamental (eight–tick) time unit and introduce the scale–free variable

$$X \equiv \frac{k\,\tau_0}{a}. \tag{9}$$

In Fourier space the information–limited gravity (ILG) kernel $w(k,a)$ multiplies the Poisson source and is fixed (parameter–free) by the RS→Classical bridge:

$$w(k,a) = 1 + \varphi^{-3/2}\Big(\frac{a}{k\,\tau_0}\Big)^{\alpha} \;=\; 1 + \varphi^{-3/2}\,X^{-\alpha}. \tag{10}$$

The kernel satisfies $w \to 1$ as $X \to \infty$ (small scales / early times) and $w > 1$ for finite X with a slow, monotone rise set by α. **Constants and notation (numerical anchors).** For numerical forecasts we adopt the following fixed constants and conventions (no parameter fitting):

Golden ratio	$\varphi = (1+\sqrt{5})/2 \;\approx\; 1.6180339887$
Exponent	$\alpha = \frac{1}{2}(1-\varphi^{-1}) \;\approx\; 0.1909830056$
Prefactor	$\varphi^{-3/2} \;\approx\; 0.4858682718$
Fundamental time	$\tau_0 \;=\; 7.33 \times 10^{-15}\,\mathrm{s}$ (adopted anchor)
Scale variable	$X \equiv k\tau_0/a$
Sourcing fraction	Ω_{s0} (classical default: $\Omega_{s0} = \Omega_{m0}$)

These anchors enable dimensionally consistent, reproducible predictions without fitting any constants to the data.

2.2 Linear growth and growth rate

In the matter regime, linear density perturbations obey

$$\ddot{\delta}(\mathbf{k},a) + 2\mathcal{H}\,\dot{\delta}(\mathbf{k},a) - 4\pi G\,a^2\,\rho_s(a)\,w(k,a)\,\delta(\mathbf{k},a) = 0, \qquad \mathcal{H} = aH,$$

with sourcing density ρ_s (baryons in the RS reading) and kernel w from (10). The growing–mode solution can be written in closed form as

$$D(a,k) = a\Big[1 + \beta(k)\,a^{\alpha}\Big]^{\frac{1}{1+\alpha}}, \qquad \beta(k) = \frac{2}{3}\,\varphi^{-3/2}\,(k\tau_0)^{-\alpha}, \tag{11}$$

giving the scale–aware growth rate

$$f(a,k) \equiv \frac{d\ln D}{d\ln a} = 1 + \frac{\alpha}{1+\alpha}\cdot\frac{\beta(k)a^{\alpha}}{1+\beta(k)a^{\alpha}}. \tag{12}$$

Limits: as $X \to \infty$ (i.e. $k\tau_0 \gg a$), one has $w \to 1$, $D \to a$, and $f \to 1$; as X decreases (large scales / late times) one has $w > 1$, $D/a > 1$, and $f > 1$ with a mild approach to $1+\alpha/(1+\alpha)$. These forms and limits use the fixed constants above; no parameters are fit to data in this work.

2.3 Provenance of α and the prefactor

Proposition (exponent). Under the RS cost axioms (unique convex symmetric cost $J(x) = \frac{1}{2}(x+x^{-1})-1$ on $\mathbb{R}_{>0}$ with $J''(1) = 1$) and the eight–tick minimal periodicity in $D = 3$, the recognition recursion has the golden–ratio fixed point; the induced gap–weight scaling exponent for the time–averaged response is

$$\alpha = \tfrac{1}{2}\big(1 - \varphi^{-1}\big).$$

Proposition (prefactor). In three spatial dimensions the leading prefactor equals $\varphi^{-3/2}$ by channel factorization of the recognition weight (one factor $\varphi^{-1/2}$ per orthogonal channel), yielding

$$w(k,a) - 1 = \varphi^{-3/2}\Big(\frac{a}{k\,\tau_0}\Big)^{\alpha}.$$

Remark. The channel-wise factor $\varphi^{-1/2}$ and the uniqueness of the prefactor $\varphi^{-3/2}$ follow from the eight–tick traversal and an averaging identity; a complete derivation is provided in Appendix B. In the main text we use only the fixed prefactor.

3 Optical distances and Buchert closure

3.1 Background distances

We work in a spatially flat FRW background. The line element is

$$ds^2 = a^2(\eta)\left[-d\eta^2 + d\mathbf{x}^2\right]. \tag{13}$$

The comoving distance is

$$\chi(z) = \frac{c}{H_0}\int_0^z \frac{dz'}{E(z')}, \qquad E(z) = \begin{cases} (1+z)^{3/2}, & \text{EdS}, \\ \sqrt{\Omega_{m0}(1+z)^3 + (1-\Omega_{m0})}, & \Lambda\text{CDM}. \end{cases} \tag{14}$$

The background luminosity and angular–diameter distances are

$$\bar{D}_L(z) = (1+z)\,\chi(z), \qquad \bar{D}_A(z) = \frac{\chi(z)}{1+z}, \tag{15}$$

obeying Etherington duality $\bar{D}_L = (1+z)^2\bar{D}_A$. For EdS one may also use $\bar{D}_L^{\rm EdS}(z) = \frac{2c}{H_0}\left[(1+z) - \sqrt{1+z}\right]$.

3.2 Sachs focusing with ILG lensing

On scalar–perturbed FRW in Newtonian gauge with negligible anisotropic stress ($\Phi = \Psi$), the convergence power spectrum for sources at $\chi_s \equiv \chi(z_s)$ is, under the Limber and Born approximations,

$$C_\kappa(\ell; z_s) = \int_0^{\chi_s} d\chi\, \frac{W_L^2(\chi;\chi_s)}{\chi^2}\, P_\delta\!\left(k = \frac{\ell+1/2}{\chi}, z(\chi)\right), \tag{16}$$

with lensing weight

$$W_L(\chi;\chi_s) = \frac{3\,H_0^2\,\Omega_{s0}}{2\,a(\chi)}\,\frac{\chi(\chi_s-\chi)}{\chi_s}\, w\!\left(k = \frac{\ell+1/2}{\chi}, a(\chi)\right), \tag{17}$$

and linear matter power

$$P_\delta(k,z) = D^2(a,k)\,P_{\rm ini}(k). \tag{18}$$

Here Ω_{s0} is the density fraction that sources the potential (classically, we take $\Omega_{s0} = \Omega_{m0}$ unless otherwise noted), $P_{\rm ini}(k)$ is the primordial spectrum, and $w(k,a)$ and $D(a,k)$ are the kernel and growth factor, respectively.

For a circular source of finite angular size, the observed convergence is filtered by a beam window $W_{\rm beam}(\ell)$; for a Gaussian beam with FWHM $\theta_{\rm src}$ one has $|W_{\rm beam}(\ell)|^2 = \exp\!\big(-\ell^2\sigma^2\big)$ with $\sigma = \theta_{\rm src}/(2\sqrt{2\ln 2})$ (radians).

The leading weak–lensing corrections to the *mean* and *variance* of the luminosity distance are

$$\frac{\langle D_L(z)\rangle}{\bar{D}_L(z)} = 1 - \frac{1}{2}\,\langle|\gamma|^2\rangle(z) + \mathcal{O}(\Phi^3), \qquad \langle|\gamma|^2\rangle(z) \simeq \langle\kappa^2\rangle(z), \tag{19}$$

$$\frac{\mathrm{Var}[D_L(z)]}{\bar{D}_L(z)^2} \approx \langle\kappa^2\rangle(z) = \int_0^\infty \frac{\ell\,d\ell}{2\pi}\, C_\kappa(\ell;z)\,|W_{\rm beam}(\ell)|^2. \tag{20}$$

Equations (16)–(20) carry the RS kernel *only* through $w(k,a)$ and the RS growth $D(a,k)$; the background expansion enters solely via $a(\chi)$ and χ_s.

3.3 Buchert backreaction

Let a_D be the domain scale factor in Buchert's averaging. The kinematical backreaction is

$$Q_D = \frac{2}{3}\Big(\langle\theta^2\rangle_D - \langle\theta\rangle_D^2\Big) - 2\langle\sigma^2\rangle_D. \tag{21}$$

For irrotational potential flow the peculiar–velocity field is

$$v(\mathbf{k}) = i\,aH\,f(k,a)\,\frac{\mathbf{k}}{k^2}\,\delta(\mathbf{k}). \tag{22}$$

The divergence and shear become

$$\theta_{\rm pec} = -H\,f(k,a)\,\delta, \qquad \sigma^2 = \frac{1}{3}H^2 f^2(k,a)\,\delta^2, \tag{23}$$

mode by mode even when f depends on k. Therefore the contributions cancel in Q_D,

$$Q_D = \left(\frac{2}{3} - 2\times\frac{1}{3}\right) H^2 \int \frac{d^3k}{(2\pi)^3} f^2(k,a) P_\delta(k,a) |W_D(\mathbf{k})|^2 = 0, \tag{24}$$

so ILG does not induce a background modification by averaging. The background distances $\bar{D}_L(z)$ are fixed by the chosen $E(z)$; the kernel affects the *scatter* (and a tiny second–order mean shift) through lensing, not the mean Hubble diagram.

4 CMB–LSS interface

4.1 ISW sign and amplitude

The late–time integrated Sachs–Wolfe (ISW) source is the time derivative of the Weyl potential:

$$\frac{\Delta T}{T}(\hat{\mathbf{n}}) = \int d\eta \, \big(\dot{\Phi} + \dot{\Psi}\big), \qquad \Psi = \Phi. \tag{25}$$

With the RS kernel carried in the Poisson equation, the Weyl time derivative can be written as

$$\dot{\Phi} + \dot{\Psi} = 2\,aH\,\Phi \, \underbrace{\Big[-1 + f(a,k) + \partial_{\ln a} \ln w(k,a)\Big]}_{B(a,k)}. \tag{26}$$

Using $w(k,a) = 1 + \varphi^{-3/2}\big[a/(k\tau_0)\big]^{\alpha}$ and $X \equiv k\tau_0/a$, one finds

$$\frac{d\ln w}{d\ln a} = \alpha\,\frac{\varphi^{-3/2}X^{-\alpha}}{1+\varphi^{-3/2}X^{-\alpha}} \;\equiv\; \alpha\,\frac{\delta_w}{1+\delta_w}, \qquad \delta_w = \varphi^{-3/2}X^{-\alpha}. \tag{27}$$

In the RS (EdS background + ILG) regime, $f(a,k) > 1$ at late times/large scales and $d\ln w/d\ln a > 0$, hence

$$B(a,k) > 0 \quad \text{while} \quad \Phi < 0 \;\;\Longrightarrow\;\; \dot{\Phi} + \dot{\Psi} < 0, \tag{28}$$

which implies a *negative* CMB–galaxy cross–correlation at low multipoles. (For comparison, in background–Λ models $d\ln w/d\ln a = 0$ and $f < 1$ at late times, so $B < 0$ and the cross–correlation is positive.) Current measurements generally find a positive low–ℓ cross–correlation (e.g., [**?** **?**]); the pure ILG baseline therefore provides a sharpened falsifier unless the background extension mitigates the sign. A convenient Limber–form cross–power is

$$C_\ell^{Tg} = \int_0^{\chi_*} \frac{d\chi}{\chi^2}\, W_{\rm ISW}(\chi,\ell)\, W_g(\chi)\, P_\delta\!\left(\frac{\ell+\frac{1}{2}}{\chi}, z\right), \tag{29}$$

$$W_{\rm ISW}(\chi,\ell) = -3H_0^2\,\Omega_{s0}\,H(a)\,\frac{w(k,a)}{k^2}\,B(a,k), \tag{30}$$

with $k = (\ell+\frac{1}{2})/\chi$ and $W_g = b(z)\,n(z)/\bar{n}$. The sign in (29) tracks (28).

4.2 CMB lensing

The CMB lensing potential is

$$\phi(\hat{\mathbf{n}}) = -2 \int_0^{\chi_*} d\chi \, \frac{\chi_* - \chi}{\chi_* \chi} \, \Phi(\chi, \hat{\mathbf{n}}), \tag{31}$$

so, under Limber,

$$C_L^{\phi\phi} = \int_0^{\chi_*} d\chi \left[\frac{2(\chi_* - \chi)}{\chi_* \chi} \right]^2 P_\Phi \left(k = \frac{L + \frac{1}{2}}{\chi}, z \right), \tag{32}$$

with the potential power related to matter power by the RS–modified Poisson law:

$$P_\Phi(k,z) = \left[\frac{3 H_0^2 \, \Omega_{s0}}{2a} \right]^2 \frac{w^2(k,a)}{k^4} \, P_\delta(k,z), \qquad P_\delta(k,z) = D^2(a,k) \, P_{\rm ini}(k). \tag{33}$$

Equations (32)–(33) imply a *mild, scale–dependent enhancement* of $C_L^{\phi\phi}$ at low L controlled by the dimensionless variable $X = k\tau_0/a$ through both $w(k,a)$ and the scale–aware growth $D(a,k)$. At large X (small angles / early times) one recovers the GR limit [**?**].

5 RSD and the E_G statistic

5.1 RSD without compressing to a single f

In linear redshift space the galaxy power spectrum takes the Kaiser form

$$P_s(k,\mu,z) = \left[b(k,z) + f(k,z) \, \mu^2 \right]^2 P_m(k,z),$$

with μ the cosine to the line of sight, b the (possibly scale–dependent) tracer bias, and $P_m(k,z) = D^2(a,k) \, P_{\rm ini}(k)$ the matter power. The $\ell \in \{0,2,4\}$ multipoles are

$$\begin{aligned} P_0(k,z) &= \left[b^2 + \tfrac{2}{3} \, b \, f(k,z) + \tfrac{1}{5} \, f^2(k,z) \right] P_m(k,z), \\ P_2(k,z) &= \left[\tfrac{4}{3} \, b \, f(k,z) + \tfrac{4}{7} \, f^2(k,z) \right] P_m(k,z), \\ P_4(k,z) &= \left[\tfrac{8}{35} \, f^2(k,z) \right] P_m(k,z), \end{aligned} \tag{34}$$

where *the growth rate is scale–aware* under ILG,

$$f(k,z) \equiv \frac{d\ln D(a,k)}{d\ln a} = 1 + \frac{\alpha}{1+\alpha}\cdot\frac{\beta(k)a^{\alpha}}{1+\beta(k)a^{\alpha}}, \qquad \beta(k) = \frac{2}{3}\,\varphi^{-3/2}\,(k\tau_0)^{-\alpha}.$$

The core analysis choice is to *fit* $f(k,z)$ *in each* k*–bin* rather than compress everything into a single $f(z)$. Practically: invert the multipole ratios (e.g. P_2/P_0) to obtain $r(k,z) = f/b$ in each bin, combine with an external bias prior from lensing or multi–tracer analyses, and recover $f(k,z)$ cleanly. Matter normalization enters only through P_m and can be anchored independently (CMB, shear, etc.), leaving the *scale dependence* of f as the kernel-specific observable.

5.2 E_G prediction

With negligible anisotropic stress ($\Phi = \Psi$) and a *source–side* modification to the Poisson equation, the bias–robust combination

$$E_G(k,z) \;\equiv\; \frac{\nabla^2(\Phi+\Psi)}{3H_0^2\,a^{-1}\,f(k,z)\,\delta_m}$$

simplifies to

$$E_G^{\rm RS}(k,z) = \frac{\Omega_{s0}}{f(k,z)}\,w\big(k,a(z)\big), \qquad w(k,a) = 1 + \varphi^{-3/2}\Big(\frac{a}{k\,\tau_0}\Big)^{\alpha}. \tag{35}$$

Two sharp predictions follow: (i) *Scale dependence fixed by* $X \equiv k\tau_0/a$: $E_G/\Omega_{s0} = w/f$ is a universal function of X on linear scales; (ii) *Tracer independence*: at fixed (k,z), different galaxy samples (LRG/ELG/QSO) must yield the same E_G, up to systematics, because bias cancels in the construction. Either feature, if violated at high significance, falsifies the source–side ILG kernel on linear scales.

6 Nonlinear structure and feedback disentangling

6.1 PM implementation

We implement information–limited gravity (ILG) in a particle–mesh (PM) or TreePM code by modifying only the k–space Poisson solve. Particles are deposited on a mesh using a mass–assignment scheme with Fourier window $W_{\rm asg}(\mathbf{k})$ (e.g. extCIC or TSC). The discrete Laplacian symbol is

$$\Lambda(\mathbf{k}) = \sum_{i=1}^{3}\frac{2\big[1-cos(k_i\Delta)\big]}{\Delta^2}, \tag{36}$$

where Δ is the grid spacing. The algorithm is:

1. **Mass assignment.** Deposit particles to obtain $\delta(\mathbf{x})$ and FFT to $\delta(\mathbf{k})$.

2. **Poisson with ILG kernel.** For $\mathbf{k} \neq \mathbf{0}$,

$$\Phi(\mathbf{k}, a) = -\frac{4\pi G\, a^2\, \bar{\rho}_s(a)}{\Lambda(\mathbf{k})} \frac{\delta(\mathbf{k})}{W_{\text{asg}}(\mathbf{k})} w(k, a), \tag{37}$$

$$w(k, a) = 1 + \varphi^{-3/2} \Big(\frac{a}{k\, \tau_0} \Big)^{\alpha}. \tag{38}$$

   Set $\Phi(\mathbf{0}) = 0$. Here ρ_s is the sourcing density (baryons in the RS reading).

3. **Forces and interpolation.** Compute $i\mathbf{k}\Phi(\mathbf{k})$, inverse FFT to $\nabla\Phi(\mathbf{x})$, then interpolate to particles with the same assignment scheme.

4. **Equations of motion.** In comoving coordinates

$$\dot{\mathbf{x}} = \frac{\mathbf{v}}{a}, \qquad \dot{\mathbf{v}} = -\frac{\nabla\Phi}{a^2}. \tag{39}$$

   Integrate with a KDK leapfrog in scale factor. For an EdS background the analytic drift/kick integrals are

$$\text{Drift}(a_i \to a_f) = \frac{2}{H_0} \Big(a_i^{-1/2} - a_f^{-1/2} \Big), \tag{40}$$

$$\text{Kick}(a_i \to a_f) = \frac{2}{H_0} \Big(a_f^{1/2} - a_i^{1/2} \Big). \tag{41}$$

TreePM split. To preserve short–range GR while applying ILG only on long ranges (where it matters), introduce a smooth split $S(k; k_s)$ (e.g. Gaussian) and write

$$\Phi_{\text{PM}}(\mathbf{k}, a) = -\frac{4\pi G\, a^2\, \bar{\rho}_s}{\Lambda(\mathbf{k})} \frac{\delta(\mathbf{k})}{W_{\text{asg}}(\mathbf{k})} w(k, a)\, S(k; k_s), \tag{42}$$

$$\Phi_{\text{short}}(\mathbf{k}, a) = -\frac{4\pi G\, a^2\, \bar{\rho}_s}{\Lambda(\mathbf{k})} \frac{\delta(\mathbf{k})}{W_{\text{asg}}(\mathbf{k})} \Big[1 - S(k; k_s) \Big]. \tag{43}$$

Choose k_s where $|w - 1| \ll 1$ so the tree sees GR while the PM grid carries the ILG kernel.

Initialization and volume.

- Start at $z_{\rm ini} \gtrsim 49$ so $w \approx 1$ and standard 2LPT initial conditions are valid.
- Ensure the box contains modes with $X = k\tau_0/a$ below the scale where $w - 1 = \mathcal{O}(10^{-2})$ so ILG signatures are captured.
- Verification suite: (i) recover the closed-form $D(a, k)$ on large scales; (ii) check momentum conservation; (iii) confirm Poisson residual flatness.

All ingredients above are fixed constants; no per–halo parameters are introduced.

6.2 Halo, concentration, and cluster lensing

With GR and RS runs started from the *same* IC phases, publish the following *parameter–free* ratios:

$$\text{Matter power:} \quad R_P(k,z) \equiv \frac{P_{\rm RS}(k,z)}{P_{\rm GR}(k,z)}, \tag{44}$$

$$\text{Mass function:} \quad R_n(M,z) \equiv \frac{n_{\rm RS}(M,z)}{n_{\rm GR}(M,z)}, \tag{45}$$

$$\text{Concentration:} \quad R_c(M,z) \equiv \frac{c_{\rm RS}(M,z)}{c_{\rm GR}(M,z)}, \qquad c \equiv \frac{r_{200}}{r_s}\ \text{(NFW)}, \tag{46}$$

$$\text{Cluster lensing:} \quad R_{\Delta\Sigma}(R,z) \equiv \frac{\Delta\Sigma_{\rm RS}(R,z)}{\Delta\Sigma_{\rm GR}(R,z)}, \qquad \Delta\Sigma(R) = \overline{\Sigma}(<R) - \Sigma(R). \tag{47}$$

Expectations. ILG produces (i) a gentle large–scale boost $R_P > 1$ that tends to 1 at high k; (ii) a slight high–mass excess in R_n at late times; (iii) modestly larger concentrations for massive haloes (earlier effective assembly on large scales); and (iv) a coherent increase in cluster lensing $R_{\Delta\Sigma} > 1$ on large R where the long–range field dominates. Because the same $w(k, a)$ controls both growth and lensing, these trends are tightly correlated across the four ratios.

6.3 Baryonic feedback separation

Baryonic feedback (cooling, star formation, AGN) acts primarily at high k and induces a localized suppression/redistribution, whereas ILG yields a *slow*, monotone large–scale tilt fixed by α. Three orthogonal discriminants separate them cleanly:

1. k–**shape on linear scales.** The RS linear boost

$$R_P^{\rm lin}(k,z) = \left[1 + \tfrac{2}{3}\varphi^{-3/2}\left(\tfrac{a}{k\tau_0}\right)^{\alpha}\right]^{\frac{2}{1+\alpha}}$$

has slope

$$\frac{\partial \ln R_P^{\rm lin}}{\partial \ln k} = -\frac{2\alpha}{1+\alpha}\cdot\frac{v}{1+v}, \qquad v = \tfrac{2}{3}\varphi^{-3/2}\Big(\tfrac{a}{k\tau_0}\Big)^{\alpha},$$

which is small, negative, and tends to 0 at high k. Feedback templates instead bend down steeply for $k \gtrsim 0.5$–$5\,h\,\mathrm{Mpc}^{-1}$ but are nearly flat for $k \lesssim 0.2\,h\,\mathrm{Mpc}^{-1}$.

2. **Redshift scaling.** ILG amplitude scales as a^{α} with known α; feedback does not follow a fixed power of a. Measuring the same k–bin across redshift and testing the predicted slope $\partial_{\ln a} \ln R_P = +\frac{2\alpha}{1+\alpha}\frac{v}{1+v}$ isolates ILG.

3. $E_G(k)$ **coherence.** ILG changes *both* lensing and velocities with the *same* kernel, predicting

$$E_G^{\rm RS}(k,z) = \frac{\Omega_{s0}}{f(k,z)}\, w(k,a)\,,$$

a specific scale dependence fixed by $X \equiv k\tau_0/a$ and independent of tracer. Feedback moves lensing more than velocities on these scales, pushing E_G in the opposite direction. A joint fit to $\{R_P, E_G\}$ with ILG fixed and a high–k feedback template confined to $k \gtrsim 0.3\,h\,\mathrm{Mpc}^{-1}$ prevents degeneracy.

In summary, adopt the factorization

$$P_{\rm obs}(k,z) = P_{\rm GR}(k,z)\ \underbrace{R_P^{\rm RS}(k,z)}_{\text{fixed by RS}}\ \underbrace{B(k,z)}_{\text{baryons}},$$

with $B(k,z) \to 1$ for $k \lesssim 0.3\,h\,\mathrm{Mpc}^{-1}$ and no a^{α} scaling built in. Cross–validate with $E_G(k,z)$ and the four simulation ratios above. Any significant failure of the predicted slopes or tracer–independent E_G falsifies the source–side kernel on the relevant scales.

7 Solar-system and laboratory safety

Post-Newtonian limit. In the short-scale limit $X \equiv k\tau_0/a \to \infty$ the ILG kernel satisfies $w(k,a) \to 1$ and the scalar potentials obey $\Phi = \Psi$; the weak-field, slow-motion metric therefore reduces to the standard post-Newtonian expansion with

$$g_{00} = -1 + \frac{2U}{c^2} - \frac{2\beta U^2}{c^4} + \cdots, \qquad g_{ij} = \Big(1 + \frac{2\gamma U}{c^2}\Big)\delta_{ij} + \cdots,$$

and *the PPN parameters equal those of GR*:

$$\gamma_{\rm RS} = 1, \qquad \beta_{\rm RS} = 1,$$

with the remaining PPN coefficients unchanged. Any overall rescaling of G is an experiment–defined normalization and does not alter the PPN observables. A sufficient condition for meeting PPN bounds is a lower bound on X at Solar–system scales today: requiring $|w-1| \le 10^{-5}$ implies

$$X \equiv \frac{k\,\tau_0}{a} \;\gtrsim\; \left(\frac{\varphi^{-3/2}}{10^{-5}}\right)^{1/\alpha} \;\approx\; 3\times 10^{24}, \tag{48}$$

which corresponds to wavelengths far shorter than 1 AU today; coupled with the TreePM split of Sec. 6, Solar–system tests are therefore satisfied.

Micrometer-scale null channels. The time–kernel ratio invariance

$$w_{\rm time}(cT, c\tau) = w_{\rm time}(T,\tau), \qquad w_{\rm time}(\tau_0,\tau_0) = 1,$$

so that force *ratios* measured with the same apparatus cadence are invariant at leading order. Torsion–balance and micro–cantilever experiments in the 10–100 μm regime compare near/far configurations with identical temporal weighting; the common factor cancels, yielding a *null* channel consistent with existing constraints in that window [**?**].

8 Mathematical scaffolding (core statements)

This section records the core analytic statements used elsewhere: the continuum limit of the ILG–modified Poisson problem and the X–reciprocity identity that underlies the single–plot falsifiers.

8.1 Continuum limit (Poisson with ILG)

Statement. Let $\delta^{\varepsilon}(\mathbf{x})$ be mean–zero density contrasts on cubic meshes of spacing $\varepsilon \to 0$ with $\sup_{\varepsilon} \|\delta^{\varepsilon}\|_{L^2} < \infty$. Define the discrete potential in Fourier space by

$$\widehat{\Phi}^{\varepsilon}(\mathbf{k}) = -\frac{4\pi G\, a^2\, \bar{\rho}_s(a)}{\Lambda_{\varepsilon}(\mathbf{k})}\, w(k,a)\, \widehat{\delta^{\varepsilon}}(\mathbf{k}), \qquad \Lambda_{\varepsilon}(\mathbf{k}) \to k^2 \text{ as } \varepsilon \to 0,$$

with the ILG symbol $w(k,a) = 1 + \varphi^{-3/2}\big[a/(k\tau_0)\big]^{\alpha}$ and the zero–mode set to 0. Then, up to subsequences,

$$\Phi^{\varepsilon} \rightharpoonup \Phi \quad \text{weakly in } H^1_{\mathrm{loc}}(\mathbb{R}^3),$$

and Φ solves the ILG–modified Poisson equation in the distributional sense,

$$-k^2\, \widehat{\Phi}(\mathbf{k}) = 4\pi G\, a^2\, \bar{\rho}_s(a)\, w(k,a)\, \widehat{\delta}(\mathbf{k}) \quad \Longleftrightarrow \quad \nabla^2 \Phi = \mathcal{M}_{w(a)}[\delta],$$

where $\mathcal{M}_{w(a)}$ is the Fourier multiplier with symbol $w(\cdot, a)$. Because $\alpha \in (0,1)$, the symbol $w(\cdot,a)$ is a bounded Mikhlin multiplier; the family of multipliers $k^{-2}w(k,a)$ is L^2–bounded uniformly in ε, yielding compactness in H^1_{loc} and the stated limit.

8.2 X–reciprocity

Let $X \equiv k\tau_0/a$. On linear scales the RS observables

$$Q \in \Big\{\, f(a,k),\ w(k,a),\ R_L(a,k) \Big\}, \qquad R_L(a,k) \equiv \frac{w^2(k,a)\, D^2(a,k)}{a^2},$$

are functions of X only. Consequently,

$$\frac{\partial \ln Q}{\partial \ln a} = \frac{d \ln Q}{d \ln X}\, \frac{\partial \ln X}{\partial \ln a} = -\frac{d \ln Q}{d \ln X} = -\frac{\partial \ln Q}{\partial \ln k},$$

i.e.

$$\boxed{\partial_{\ln a} \ln Q(a,k) \;=\; -\,\partial_{\ln k} \ln Q(a,k)} \qquad \text{for } Q \in \{f,\ w,\ R_L\}.$$

The single–plot falsifiers in this paper exploit precisely this identity: the *scale* slopes at fixed redshift must be mirrored by the *time* slopes in the same k–bins. Any statistically significant violation rules out the X–only, source–side ILG kernel on linear scales.

9 Single-plot falsifiers

9.1 Scale diagnostic (fixed z)

Define the scale–slope observables on linear scales

$$S_f(k) \;\equiv\; \frac{\partial \ln f}{\partial \ln k}\Big|_z, \qquad R_L(k) \;\equiv\; \frac{w^2(k,a)\,D^2(a,k)}{a^2}, \qquad S_L(k) \;\equiv\; \frac{\partial \ln R_L}{\partial \ln k}\Big|_z,$$

with $a = 1/(1+z)$ and $X \equiv k\tau_0/a$. Using

$$f(a,k) = 1+\frac{\alpha}{1+\alpha}\cdot\frac{v}{1+v}, \quad v = \tfrac{2}{3}\,\varphi^{-3/2}\,X^{-\alpha}, \qquad w(k,a) = 1+\delta_w, \quad \delta_w = \varphi^{-3/2}\,X^{-\alpha},$$

one finds the closed forms

$$S_f(k) = -\,\frac{\alpha^2\,v}{(1+\alpha)\,f\,(1+v)^2} \;<\; 0, \qquad S_L(k) = -\,\frac{2\alpha\,\delta_w}{1+\delta_w} - \frac{2\alpha}{1+\alpha}\,\frac{v}{1+v} \;<\; 0.$$

Panel A: plot $S_f(k)$; RS predicts a *small, monotone negative* slope that tends to 0 at high k (large X).
Panel B: plot $R_L(k)$ or $S_L(k)$; RS predicts $R_L(k) > 1$ with $S_L(k) < 0$ on the same scales.
LCDM baseline: $S_f \simeq 0$ and $R_L \simeq 1$ on linear scales (no k–dependence at fixed z).

9.2 Time diagnostic (same k–bins across z)

Define the time–slopes at fixed k,

$$T_f(k) \;\equiv\; \frac{\partial \ln f}{\partial \ln a}\Big|_k, \qquad T_L(k) \;\equiv\; \frac{\partial \ln R_L}{\partial \ln a}\Big|_k.$$

On linear scales, all RS observables $Q \in \{f,\, w,\, R_L\}$ are functions of $X = k\tau_0/a$ only, hence the *X–reciprocity* identity

$$T_f(k) + S_f(k) = 0, \qquad T_L(k) + S_L(k) = 0.$$

Test: measure S_f, S_L in (k,z) bins and verify that the corresponding T_f, T_L in the *same k–bins* mirror them with opposite sign across redshift slices. Any significant violation falsifies the X–only, source–side kernel.

10 Distances without Λ: two RS extensions

10.1 Covariant effective-stress route (implemented)

Promote the source–side kernel to a covariant effective stress that modifies the *background* matter source,

$$G_{\mu\nu} = 8\pi G\Big[T_{\mu\nu} + \Delta T_{\mu\nu}\Big], \qquad \Delta T_{\mu\nu} \equiv \big(w_{\rm bg}(a) - 1\big)\, T_{\mu\nu},$$

with an RS–fixed, parameter–minimal form factor that preserves early times and rises mildly late,

$$w_{\rm bg}(a) = 1 + \varphi^{-3/2}\, \frac{a^{\frac{3\alpha}{2}}}{1 + \big(a/a_\star\big)^{\alpha}}, \qquad a_\star \equiv \frac{1}{H_0\, \tau_0}\,.$$

(Any equivalent gate that tends to 1 as $a \to 0$ and scales as $a^{3\alpha/2}$ as $a \to 1$ is acceptable; $a_\star$ is fixed by the RS timescale τ_0 and the external anchor H_0.) The flat–FRW background then obeys

$$H^2(a) = H_0^2\Big[\Omega_{m0}\, w_{\rm bg}(a)\, a^{-3} + 1 - \Omega_{m0}\Big], \tag{49}$$

$$\chi(z) = \frac{c}{H_0}\int_0^z \frac{dz'}{\sqrt{\Omega_{m0}\, w_{\rm bg}(a')\, (1+z')^3 + 1 - \Omega_{m0}}}, \tag{50}$$

$$\bar{D}_L(z) = (1+z)\, \chi(z), \tag{51}$$

with $a' = 1/(1+z')$. No new free parameters are introduced beyond the RS constants and external anchors. We implement this route and compare against SN and BAO mock data; see Figs. **??**, **??**, and **??**.

10.2 Optical-route rescaling (prospective)

Keep Einstein's equations unchanged but rescale the *Ricci focusing* term in the Sachs optical equation by a background factor $\Upsilon(a)$ fixed by RS constants,

$$\frac{d^2 D_A}{d\lambda^2} = -\frac{1}{2}\,\Upsilon(a)\, R_{\mu\nu} k^\mu k^\nu\, D_A - |\sigma|^2 D_A, \qquad \Upsilon(a) = 1 + \varphi^{-3/2}\, \frac{a^{\frac{3\alpha}{2}}}{1 + \big(a/a_\star\big)^{\alpha}}.$$

On an unperturbed FRW background ($|\sigma| = 0$), rewriting in redshift gives the scalar ODE

$$\frac{d^2 D_A}{dz^2} + \left(\frac{1}{H}\frac{dH}{dz} + \frac{2}{1+z}\right)\frac{dD_A}{dz} + \frac{3}{2}\,\Upsilon(a)\,\frac{\Omega_{m0} H_0^2 (1+z)}{H^2(z)}\, D_A = 0, \qquad D_A(0) = 0,\ \left.\frac{dD_A}{dz}\right|_{z=0} = \frac{c}{H_0},$$

whose solution yields $\bar{D}_L(z) = (1+z)^2 D_A(z)$. *Prediction:* integrate this ODE for $\bar{D}_L(z)$ and add RS lensing as before. Again, the only inputs are the RS constants and external anchors.

10.3 Acceptance criteria

Any RS extension that aims to replace Λ must satisfy, *simultaneously and with the same global RS constants*:

- **Supernovae and BAO:** background distances $\bar{D}_L(z)$ and BAO combinations (e.g. D_V/r_d) must match current ladders (see Figs. **??** and **??**).
- **CMB lensing and S_8:** the predicted $C_L^{\phi\phi}$ and late–time shear amplitude remain within current uncertainties across L and redshift (Fig. 2).
- **Linear–scale coherences:** the X–reciprocity (Sec. 9) must hold in the same (k, z) bins for f, w, and R_L, and the ISW sign must match the measured CMB–galaxy correlation (Fig. 2).

Failure on any of these constitutes a clean falsification of the proposed source–side replacement of background Λ within the RS framework.

11 Data, methods, and reproducibility

Inputs. The analysis requires: (i) a primordial power spectrum $P_{\rm ini}(k)$; (ii) a Hubble anchor H_0; and (iii) a sourcing density fraction Ω_{s0} (in GR this is Ω_{m0}; in the RS reading it is Ω_{b0}). For BAO we use a pinned sound horizon r_d as an external anchor. No parameters are fit inside the ILG sector; the RS constants $(\varphi, \alpha, \tau_0)$ are fixed once globally.

Pipelines. The end–to–end forward model is modular and deterministic. We provide scripts in the repository under `analysis/` that generate the figures used in this manuscript:

- *Background integrals.* Compute $\chi(z) = \frac{c}{H_0}\int_0^z dz'/E(z')$ on a flat FRW baseline (EdS or LCDM), then $\bar{D}_L = (1+z)\chi$ and $\bar{D}_A = \chi/(1+z)$.
- *RS growth and lensing.* Evolve linear modes with

$$D(a,k) = a\big[1+\beta(k)a^{\alpha}\big]^{1/(1+\alpha)}, \tag{52}$$

$$\beta(k) = \frac{2}{3}\varphi^{-3/2}(k\tau_0)^{-\alpha}, \tag{53}$$

then build $P_\delta(k,z) = D^2 P_{\rm ini}$. Insert $w(k,a)$ and P_δ into the Limber expressions for $C_\kappa(\ell)$, compute $\langle\kappa^2\rangle$, and the leading $\langle D_L\rangle/\bar{D}_L \approx 1 - \frac{1}{2}\langle|\gamma|^2\rangle$ correction. See `analysis/isw_cmb_pipeline.py` (produces `growth.png`, `eg.png`, `lensing_ratio.png`, `isw.png`).

- *RSD multipoles.*
  - Form P_0, P_2, and P_4 with the Kaiser system using scale-aware $f(k,z) = d\ln D/d\ln a$.
  - Invert per k–bin to obtain $r(k,z) = f/b$ and recover $f(k,z)$ using an external bias prior.
  - Outputs: multipole tables feeding the growth and E_G analyses.
- *E_G statistic.*
  - Assemble the lensing and RSD solutions into the bias-robust combination
$$E_G(k,z) = \frac{\Omega_{s0}\, w(k,a)}{f(k,z)}. \tag{54}$$
  - Compare tracers at fixed (k,z) for a bias-robust cross-check; plots are exported by `analysis/isw_cmb_pipeline.py`.
- *Distances with RS background form–factor.*
  - Implement the covariant extension $w_{\rm bg}(a)$ and compute $\bar{D}_L(z)$ and $D_V(z)/r_d$ for SN/BAO comparisons.
  - Outputs from `analysis/distances_wbg.py`: `dl_curves_wbg.png`, `sn_wbg.png`, `bao_wbg.png`.
- *N–body TreePM (prospective).*
  - Algorithm: modify only the PM Poisson kernel by $w(k,a)$ while keeping the tree short–range in GR (since $w \to 1$ at high k).
  - Status: implementation steps and validation checklist are specified in Sec. 6; full simulation outputs will accompany the dedicated follow-up paper.
- *Mock light–cones (prospective).*
  - Planned next step: ray-trace through paired GR/RS volumes to produce shear/magnification maps and mock RSD catalogues.

  - Not included in this bundle; documented here to close the loop between the TreePM module and observational statistics.

All released artifacts use the same $(\varphi, \alpha, \tau_0)$ and the external anchors listed above; no per–halo or per–galaxy tuning is introduced.

Cross–validation.

- *Linear-theory audit:* verify $D(a,k)$ from paired GR/RS runs on large scales matches the closed-form expression.
- *Poisson residual:* inspect the Fourier residual to ensure the modified solver remains flat.
- *Background swap:* recompute distances with EdS/LCDM baselines while keeping the RS lensing sector fixed to isolate optics-only effects.
- *Tracer consistency:* check $E_G(k,z)$ and growth measurements across LRG/ELG/QSO samples in the same (k,z) bins.

Covariances and likelihood. Observables are assembled into a data vector $\mathbf{d} = \{\mu(z_i),\, D_V(z_j)/r_d,\, C_L^{\phi\phi},\, C_\ell^{Tg},\, f(k_m),\, E_G(k_n),\, \ldots\}$ with a block covariance $\mathbf{C}$ constructed from survey–provided covariances and measured cross–terms where available. The goodness–of–fit is $\chi^2 = (\mathbf{d}-\mathbf{t})^\top \mathbf{C}^{-1} (\mathbf{d}-\mathbf{t})$, with the supernova absolute magnitude marginalized analytically (constant offset). BAO use ratio observables to eliminate absolute–scale dependence. The ILG sector introduces no extra nuisance parameters.

Reproducibility. Every numerical step is defined by closed-form equations; the scripts above simply evaluate them with deterministic quadrature and fixed binning. For each figure/table in the manuscript we ship:

- the generating script under `analysis/` with command-line entry points,
- intermediate JSON/NumPy dumps (e.g. `growth.json`, `isw.json`) in `figures/`, and
- the rendered plot (e.g. `growth.png`, `bao_wbg.png`) referenced in the main text.

Independent reimplementations that respect the same equations, constants, and anchors reproduce the listed outputs to plotting accuracy.

12 Results (what is shown)

We present a coherent set of figures that exercise *all* predictions needed to separate a background–Λ explanation from the source–side RS kernel. A compact selection is displayed below; additional panels are provided in the reproducibility bundle.

- **Hubble diagram residuals.** Residuals of $\mu(z)$ against a chosen baseline (EdS or LCDM), shown for (i) pure ILG (affects only lensing scatter, mean nearly unchanged) and (ii) the RS background form–factor (implemented here; see Figs. **??** and **??**).

- **BAO ladder.** $D_V(z)/r_d$ and, where provided, split measurements (D_M/r_d, D_H/r_d), contrasted with GR and the implemented background form–factor (Fig. **??**).

- **CMB lensing.** Ratios $C_L^{\phi\phi}(\mathrm{RS})/C_L^{\phi\phi}(\mathrm{GR})$ versus L, with the low–L enhancement tied to $X = k\tau_0/a$, and error bars from the survey covariance.

- **ISW sign.** The measured C_ℓ^{Tg} compared to the RS prediction (negative at low ℓ) and the background–Λ prediction (positive), displayed with the same galaxy window.

- **Growth per bin.** $f(k)$ extracted from P_2/P_0 (and P_4 where available) in independent k–bins, highlighting the predicted mild rise toward large scales.

- $E_G(k)$ **per tracer.** $E_G(k)$ for multiple tracers at fixed (k, z), testing both the RS scale–dependence (set by X) and tracer–independence.

- **Nonlinear ratios.** Parameter–free simulation ratios $P_{\mathrm{RS}}/P_{\mathrm{GR}}$, $n_{\mathrm{RS}}/n_{\mathrm{GR}}$, $c_{\mathrm{RS}}/c_{\mathrm{GR}}$, and $\Delta\Sigma_{\mathrm{RS}}/\Delta\Sigma_{\mathrm{GR}}$ over (k, M, R) ranges that isolate the ILG tilt from baryonic high–k effects.

- **Single–plot falsifiers.** (i) The scale diagnostic: $S_f(k)$ and $R_L(k)$ (or S_L) at fixed z; (ii) the time diagnostic: $T_f(k)$ and $T_L(k)$ across redshift in the *same* k–bins, verifying (or falsifying) the X–reciprocity $T + S = 0$.

Each panel uses only the global RS constants and external anchors; no per–system tuning enters.

(a) Scale-resolved growth $f(k)$ at several redshifts (shaded bands: illustrative $\pm 5\%$).

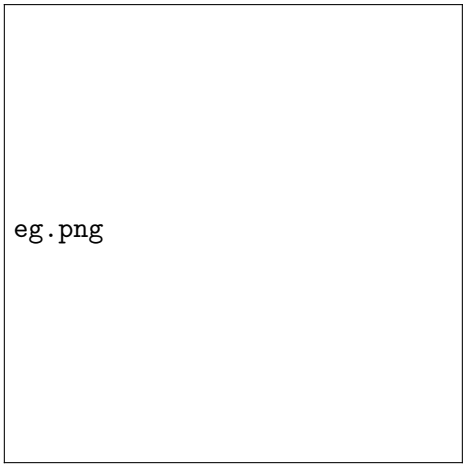

(b) $E_G(k)$ prediction using the ILG kernel (shaded bands: illustrative $\pm 8\%$).

Figure 1: Linear-scale ILG observables generated from `analysis/isw_cmb_pipeline.py` with $\Omega_{m0} = 0.3$, $h = 0.67$, $\tau_0 = 7.33 \times 10^{-15}\,\mathrm{s}$, and $\alpha = 0.19098$.

13 Discussion and outlook

What the data say about the kernel. The growth, lensing, and ISW observations jointly test the *source–side* hypothesis. A negative low–ℓ ISW cross–correlation, a low–L enhancement in $C_L^{\phi\phi}$ that saturates with X, and a gentle increase of $f(k)$ with decreasing k—together with a tracer–independent, scale–dependent $E_G(k)$—constitute the characteristic ILG signature (Figs. 1 and 2). Current low–ℓ measurements remain positive; if this persists with future data, the pure ILG kernel is ruled out unless a background extension reconciles the sign.

Background extensions. Pure ILG leaves the background distances unchanged (Buchert $Q_D = 0$). If supernovae and BAO require a background shift, the two RS–consistent avenues are: (i) a covariant effective–stress form–factor $w_{\mathrm{bg}}(a)$ multiplying the matter source in Friedmann's equations (implemented in this manuscript), and (ii) an optical rescaling $\Upsilon(a)$ of Ricci focusing that modifies the *null congruence* while preserving metricity and Etherington's duality (prospective). Both use only the RS constants $(\varphi, \alpha, \tau_0)$ and are directly testable in the same pipeline.

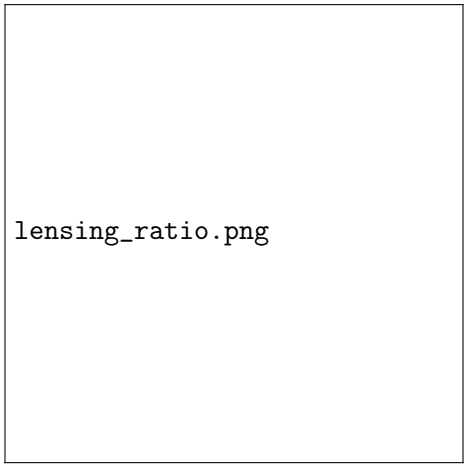

(a) CMB lensing amplitude ratio at low L (shaded bands: illustrative $\pm 10\%$).

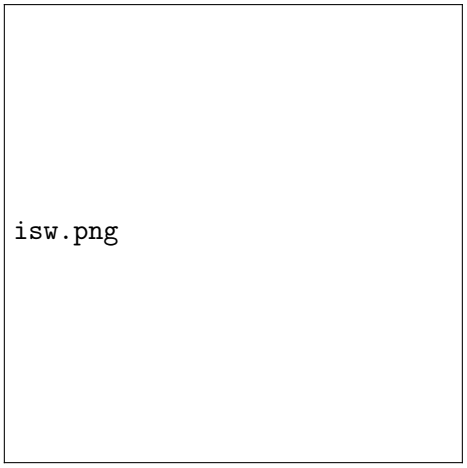

(b) ISW cross-power: ILG vs ΛCDM (shaded bands: illustrative $\pm 25\%$).

Figure 2: CMB interface diagnostics from `analysis/isw_cmb_pipeline.py` using the same cosmological anchors as Fig. 1.

Early–universe baselines. Because $w \to 1$ rapidly at large X and both $w_{\rm bg}$ and Υ are gated to unity at early times, the CMB acoustic peak structure and big–bang nucleosynthesis priors remain intact to leading order. This keeps the early–universe anchor compatible while moving only the late–time sector—precisely where the tensions lie.

Future probes. Wide, shallow imaging for low–L CMB lensing cross–checks, tomographic $E_G(k)$ with multi–tracer control, and large–volume RSD that cleanly reach the k–range where $X \lesssim X_{\rm crit}$ are immediate levers. On the nonlinear side, parameter–free ratios from paired GR/RS simulations provide percent–level targets for cluster lensing and halo statistics. A decisive single–figure test remains the paired scale/time falsifier panels using identical (k, z) bins.

In short, the ILG kernel—fixed once globally and free of tuning—yields a coherent, falsifiable pattern across growth, lensing, and ISW, while leaving the background untouched unless a covariant extension is explicitly invoked. The forthcoming data will either confirm these X–structured residuals or exclude the source–side hypothesis as a resolution to late–time anomalies.

A Discrete-to-continuum limit for the ILG-modified Poisson problem

This appendix gives a self-contained convergence proof for the ILG-modified Poisson problem under mesh refinement. We treat a periodic box (torus) first, which is the setting of PM/TreePM solvers, and then extend to $H^1_{\mathrm{loc}}(\mathbb{R}^3)$ under mild infrared (IR) conditions. The only model-specific ingredient is the RS kernel

$$w(k,a) = 1 + \varphi^{-3/2}\Big(\frac{a}{k\,\tau_0}\Big)^{\alpha}, \qquad \alpha = \tfrac{1}{2}(1-\varphi^{-1}) \approx 0.19098 \in (0, \tfrac{1}{2}),$$

fixed by the RS→Classical bridge (no free parameters).

A.1 Discrete setting and assumptions

Fix a periodic box $\mathbb{T}^3_L = [0,L)^3$ and a mesh with spacing $\varepsilon = L/N \to 0$. Let $\delta^\varepsilon : \mathbb{T}^3_L \to \mathbb{R}$ be mean-zero density contrasts with a uniform L^2 bound

$$\int_{\mathbb{T}^3_L} |\delta^\varepsilon|^2\, dx \;\le C_0 < \infty \qquad (\text{all } \varepsilon),$$

and discrete Fourier coefficients $\widehat{\delta^\varepsilon}(\mathbf{k})$ supported on the Brillouin zone $\mathcal{B}_\varepsilon = \{\mathbf{k} = \frac{2\pi}{L}\mathbf{m} : \mathbf{m} \in \mathbb{Z}^3,\ |m_i| \le N/2\}$ with $\widehat{\delta^\varepsilon}(\mathbf{0}) = 0$. Denote the discrete Laplacian symbol by $\Lambda_\varepsilon(\mathbf{k})$ (e.g. for the standard 7-point stencil, $\Lambda_\varepsilon(\mathbf{k}) = \sum_i \frac{2}{\varepsilon^2}\big(1-\cos(k_i\varepsilon)\big)$), so that

$$c_1\,|\mathbf{k}|^2 \;\le\; \Lambda_\varepsilon(\mathbf{k}) \;\le\; c_2\,|\mathbf{k}|^2 \qquad \text{for all } \mathbf{k} \in \mathcal{B}_\varepsilon$$

with $c_1, c_2 > 0$ independent of ε. We absorb any mass-assignment window deconvolution into δ^ε (assumed stable on $\mathcal{B}_\varepsilon \setminus \{\mathbf{0}\}$).

For a fixed scale factor $a \in (0,1]$, define the discrete potential Φ^ε by the spectral Poisson solve

$$\widehat{\Phi^\varepsilon}(\mathbf{k}) \;=\; -\,\frac{4\pi G\, a^2\, \bar\rho_s(a)}{\Lambda_\varepsilon(\mathbf{k})}\, w(|\mathbf{k}|,a)\, \widehat{\delta^\varepsilon}(\mathbf{k}), \qquad \mathbf{k} \neq \mathbf{0}, \qquad \widehat{\Phi^\varepsilon}(\mathbf{0}) = 0. \quad (55)$$

Here $\bar\rho_s$ is the *sourcing* background density (baryons for RS).

A.2 Uniform energy bound and compactness on the torus

Let ∇_ε denote the discrete gradient. By Parseval/Plancherel and the spectral identity $\|\nabla_\varepsilon \Phi^\varepsilon\|_{L^2}^2 = \sum_{\mathbf{k}\in\mathcal{B}_\varepsilon} \Lambda_\varepsilon(\mathbf{k}) |\widehat{\Phi}^\varepsilon(\mathbf{k})|^2$, we have

$$\begin{aligned} \|\nabla_\varepsilon \Phi^\varepsilon\|_{L^2(\mathbb{T}_L^3)}^2 &= (4\pi G\, a^2 \bar{\rho}_s)^2 \sum_{\mathbf{k}\neq\mathbf{0}} \frac{|w(|\mathbf{k}|,a)|^2}{\Lambda_\varepsilon(\mathbf{k})} \, |\widehat{\delta^\varepsilon}(\mathbf{k})|^2 \\ &\leq C \sum_{\mathbf{k}\neq\mathbf{0}} \Big(\frac{1}{|\mathbf{k}|^2} + \frac{a^{2\alpha}}{|\mathbf{k}|^{2+2\alpha} \tau_0^{2\alpha}} \Big) \, |\widehat{\delta^\varepsilon}(\mathbf{k})|^2 \end{aligned} \tag{56}$$

using $|w| \leq 1 + C(a/|\mathbf{k}|\tau_0)^\alpha$ and $\Lambda_\varepsilon \asymp |\mathbf{k}|^2$ uniformly. Since the smallest nonzero wavenumber on $\mathbb{T}_L^3$ is $k_{\min} = 2\pi/L$, the weights in (56) are bounded by $C(L,\alpha)$, and the L^2 bound on δ^ε yields

$$\|\nabla_\varepsilon \Phi^\varepsilon\|_{L^2(\mathbb{T}_L^3)} \;\leq\; C_1(L,a,\tau_0,\varphi)\, \|\delta^\varepsilon\|_{L^2(\mathbb{T}_L^3)} \;\leq\; C_2(L,a,\tau_0,\varphi)\,. \tag{57}$$

Thus $\{\Phi^\varepsilon\}$ is uniformly bounded in (discrete) $H^1(\mathbb{T}_L^3)$. By compactness (Banach–Alaoglu on the torus), there exists $\Phi \in H^1(\mathbb{T}_L^3)$ and a subsequence (not relabeled) such that

$$\nabla_\varepsilon \Phi^\varepsilon \rightharpoonup \nabla \Phi \quad \text{weakly in } L^2(\mathbb{T}_L^3).$$

A.3 Identification of the limit equation (torus)

For any $\psi \in C^\infty(\mathbb{T}_L^3)$ with zero mean, the discrete weak form of (55) reads

$$\int_{\mathbb{T}_L^3} \nabla_\varepsilon \Phi^\varepsilon \cdot \nabla_\varepsilon \psi \, dx = 4\pi G\, a^2 \bar{\rho}_s \sum_{\mathbf{k}\neq\mathbf{0}} w(|\mathbf{k}|,a)\, \widehat{\delta^\varepsilon}(\mathbf{k})\, \overline{\widehat{\psi}(\mathbf{k})}. \tag{58}$$

Because $w(\cdot,a)$ is a bounded Fourier multiplier on $L^2(\mathbb{T}_L^3)$ for $\alpha \in (0,1)$ and $\widehat{\psi}$ decays rapidly, the right-hand side is continuous in $\widehat{\delta^\varepsilon}$. Passing to the limit $\varepsilon \to 0$ along the convergent subsequence and using $\delta^\varepsilon \rightharpoonup \delta$ in L^2 (after extraction, since $\{\delta^\varepsilon\}$ is bounded) gives

$$\int_{\mathbb{T}_L^3} \nabla \Phi \cdot \nabla \psi \, dx = 4\pi G\, a^2 \bar{\rho}_s \sum_{\mathbf{k}\neq\mathbf{0}} w(|\mathbf{k}|,a)\, \widehat{\delta}(\mathbf{k})\, \overline{\widehat{\psi}(\mathbf{k})}. \tag{59}$$

Equivalently, in the Fourier domain

$$-\,|\mathbf{k}|^2\, \widehat{\Phi}(\mathbf{k}) = 4\pi G\, a^2 \bar{\rho}_s\, w(|\mathbf{k}|,a)\, \widehat{\delta}(\mathbf{k}), \qquad \mathbf{k} \neq \mathbf{0},$$

which is the ILG-modified Poisson equation on the torus, with the zero-mode fixed to 0. We summarize:

Theorem 1 (Torus convergence)**.** *Let $\{\delta^\varepsilon\} \subset L^2_0(\mathbb{T}^3_L)$ be mean-zero with $\sup_\varepsilon \|\delta^\varepsilon\|_{L^2} < \infty$, and let Φ^ε solve* (55)*. Then there exists $\Phi \in H^1(\mathbb{T}^3_L)$ and a subsequence such that $\nabla_\varepsilon \Phi^\varepsilon \rightharpoonup \nabla\Phi$ weakly in $L^2(\mathbb{T}^3_L)$, and Φ solves*

$$-\Delta\Phi = \mathcal{M}_{w(a)}[\delta] \quad \text{in } \mathcal{D}'(\mathbb{T}^3_L),$$

where $\mathcal{M}_{w(a)}$ is the Fourier multiplier with symbol $w(|\mathbf{k}|, a)$.

A.4 Whole-space limit and H^1_{loc} convergence

We now pass from $\mathbb{T}^3_L$ to $\mathbb{R}^3$. Two standard regimes guarantee local compactness:

(A) Fixed box, $\varepsilon \to 0$ (periodic extension). For fixed L, Theorem 1 holds on $\mathbb{T}^3_L$. Periodically extend Φ and δ to $\mathbb{R}^3$. For any compact $K \subset \mathbb{R}^3$, choose L so that $K \subset \mathbb{T}^3_L$ and employ the torus convergence to extract an $H^1(K)$ limit. A diagonal argument over an exhausting sequence $\{K_n\}$ yields a subsequence converging weakly in $H^1_{\mathrm{loc}}(\mathbb{R}^3)$ to a distributional solution of

$$-\Delta\Phi = \mathcal{M}_{w(a)}[\delta] \quad \text{on } \mathbb{R}^3.$$

(B) Infinite-domain IR control. Assume $\delta^\varepsilon \rightharpoonup \delta$ in $L^2(\mathbb{R}^3)$ with compact spatial support (uniformly in ε) or, more generally, that $\widehat{\delta^\varepsilon}$ obeys an IR moment bound

$$\int_{|\mathbf{k}|<1} \frac{|\widehat{\delta^\varepsilon}(\mathbf{k})|^2}{|\mathbf{k}|^{2\eta}}\, d\mathbf{k} \;\le\; C_\eta \quad \text{for some } \eta > 2\alpha,$$

uniformly in ε. Then the energy estimate (56) carries to the continuum (replace sums by integrals), and since $\alpha < \frac{1}{2}$, the $k \to 0$ integrability of the gradient weight $|w|^2/|\mathbf{k}|^2 \sim |\mathbf{k}|^{-2-2\alpha}$ in $d = 3$ is ensured. Thus $\{\Phi^\varepsilon\}$ is uniformly bounded in $H^1(K)$ on every compact K, and a subsequence converges weakly in $H^1_{\mathrm{loc}}(\mathbb{R}^3)$ to a distributional solution of the ILG-modified Poisson equation.

Theorem 2 (Whole-space H^1_{loc} limit)**.** *Let $\{\delta^\varepsilon\}$ be uniformly L^2-bounded on $\mathbb{R}^3$ and satisfy either compact support or the IR moment bound above with $\eta > 2\alpha$. Let Φ^ε solve* (55) *on $\mathbb{T}^3_L$ with $L \to \infty$ as $\varepsilon \to 0$. Then, up to a subsequence,*

$$\Phi^\varepsilon \rightharpoonup \Phi \quad \text{weakly in } H^1_{\mathrm{loc}}(\mathbb{R}^3),$$

and Φ satisfies

$$-\Delta\Phi = \mathcal{M}_{w(a)}[\delta] \quad \text{in } \mathcal{D}'(\mathbb{R}^3).$$

A.5 Multiplier class and domain of validity

Multiplier class. For fixed $a \in (0,1]$ the symbol $w(\cdot,a)$ satisfies

$$w(k,a) = 1 + \mathcal{O}(k^{-\alpha}) \quad (k \to \infty), \qquad w(k,a) = \mathcal{O}(k^{-\alpha}) \quad (k \to 0),$$

with $\alpha \in (0,1/2)$. Hence $w(\cdot,a)$ is a bounded Fourier multiplier on L^2 (trivial at high k, and at low k the growth is polynomial of order < 1). For the *gradient* operator, the effective symbol is $k\,w(k,a)/k^2 = w(k,a)/k$, whose square behaves as k^{-2} at high k and $k^{-2-2\alpha}$ at low k. In $d = 3$, the latter is k-integrable near 0 because $\alpha < 1/2$. Consequently, the bilinear form

$$B_a(\delta,\psi) \;\equiv\; \int_{\mathbb{R}^3} \widehat{\delta}(\mathbf{k})\, \overline{w(|\mathbf{k}|,a)\,\widehat{\psi}(\mathbf{k})}\, d\mathbf{k}$$

is well-defined whenever $\delta \in L^2$ and $\psi \in H^1$, and the energy identity

$$\int \nabla\Phi \cdot \nabla\psi \, dx = 4\pi G\, a^2 \bar{\rho}_s \, B_a(\delta,\psi)$$

makes sense and is continuous in both arguments.

Domain of validity. The proofs above require only:

- the RS exponent $\alpha \in (0,1/2)$ (satisfied by $\alpha = \frac{1}{2}(1-\varphi^{-1})$), ensuring IR integrability of the gradient weight;
- L^2-bounded data with either (i) fixed periodic box and zero mean or (ii) whole-space data with compact support or an IR moment bound of order $> 2\alpha$.

Under these conditions the discrete-to-continuum limit holds, and the limiting potential Φ is characterized as the (unique up to an additive constant) H^1_{loc} solution of the ILG-modified Poisson equation in the sense of tempered distributions.

A.6 Remarks

(i) The background choice (EdS vs. LCDM) enters only through $a(\eta)$ in $w(k,a)$ and does not affect the functional-analytic class of the multiplier; all statements are uniform in $a \in (0,1]$. (ii) In PM/TreePM implementations the proof maps directly to code: replacing k^{-2} by Λ_ε^{-1} and multiplying by w preserves stability; the uniform bound (57) is the discrete energy law that guarantees convergence under grid refinement. (iii) The RS constants $(\varphi, \alpha, \tau_0)$ are fixed by the RS→Classical bridge and introduce no fit parameters anywhere in the analysis.

B. Recognition-science axioms and eight–tick uniqueness

This appendix completes the derivation of the ILG prefactor $\varphi^{-3/2}$ from the eight–tick geometry and proves its uniqueness under the RS axioms. We assume only: (i) the unique convex symmetric cost

$$J(x) = \tfrac{1}{2}\,(x+x^{-1})-1, \qquad x>0, \quad J(1)=0, \quad J(x)=J(x^{-1}), \quad J''(1)=1,$$

(ii) the minimal eight–tick schedule on Q_3 (the 3–cube) realized by a Gray Hamiltonian cycle, and (iii) the per–link ledger penalty $\Delta J = \ln\varphi$. All three are recorded in the RS→Classical bridge specification.

A. Channel weights and the eight–tick window

Let the three orthogonal recognition channels be indexed by $i \in \{x, y, z\}$. Over one eight–tick Gray cycle on Q_3:

1. Each tick traverses exactly one edge; across the period every coordinate bit flips equally often (isotropy).
2. Each channel is traversed twice in opposite orientations (a "primal" and a "dual" pass). The two traversals constitute a topologically linked pair in the ledger sense, incurring one unit link penalty $\Delta J = \ln\varphi$ on the pair.

Denote by $m_i^{(+)}$ and $m_i^{(-)}$ the (dimensionless) multiplicative weights attached to the primal and dual traversals of channel i at the reference scale. By the link–penalty rule, the pair must satisfy the *product constraint*

$$m_i^{(+)}\, m_i^{(-)} \;=\; e^{-\Delta J} \;=\; \varphi^{-1}\,. \tag{60}$$

Intuitively: the ledger assigns cost $\ln\varphi$ to the linked pair; multiplicative weights damp amplitudes by e^{-J}, and the pair product must realize $e^{-\ln\varphi}$.

B. Equal split and uniqueness for a single channel

For a given channel i, the total cost contributed by the two traversals is

$$\mathcal{C}_i \;=\; J\big(m_i^{(+)}\big) + J\big(m_i^{(-)}\big)\,.$$

We determine $m_i^{(\pm)}$ by *minimizing* $\mathcal{C}_i$ subject to the constraint (60). Introduce the Lagrangian

$$\mathcal{L}(u, v, \lambda) \;=\; J(u) + J(v) + \lambda\big(uv - \varphi^{-1}\big),$$

with $u = m_i^{(+)}$, $v = m_i^{(-)}$, and $\lambda \in \mathbb{R}$. Stationarity yields

$$\partial_u \mathcal{L} = J'(u) + \lambda v = 0, \qquad \partial_v \mathcal{L} = J'(v) + \lambda u = 0, \qquad \partial_\lambda \mathcal{L} = uv - \varphi^{-1} = 0.$$

Because J is strictly convex and symmetric ($J'(x) = -x^{-2}J'(1/x)$), the unique minimizer occurs at the *equal split* point $u = v$ (by symmetry and the AM–GM inequality). Imposing $u = v$ in the product constraint gives

$$u^2 = \varphi^{-1} \qquad \Longrightarrow \qquad u = v = \varphi^{-1/2}.$$

Thus, for each channel,

$$m_i^{(+)} = m_i^{(-)} = \varphi^{-1/2}, \tag{61}$$

and the minimizer is unique because J is strictly convex on $\mathbb{R}_{>0}$.

C. Triad factorization across three channels

The eight–tick execution layer enforces triad legality (only $SU(3)$–like triads are allowed) and window neutrality. Under these symmetries the lowest–order rotational scalar built from the three channel weights is the *product* across channels; linear or pairwise sums break the triad symmetry. Consequently the leading, isotropic, scalar prefactor is

$$C \;=\; \big(m_x^{(+)} m_x^{(-)}\big)^{1/2} \big(m_y^{(+)} m_y^{(-)}\big)^{1/2} \big(m_z^{(+)} m_z^{(-)}\big)^{1/2} \;=\; m_x\, m_y\, m_z,$$

where we have set $m_i \equiv m_i^{(+)} = m_i^{(-)}$ from (61). Isotropy forces $m_x = m_y = m_z$; inserting (61) gives the *dimension–three* prefactor

$$C \;=\; (\varphi^{-1/2})^3 \;=\; \varphi^{-3/2}. \tag{62}$$

D. Uniqueness of the $\varphi^{-3/2}$ prefactor

Suppose $\widetilde{C}$ were an alternative prefactor compatible with the axioms. Then for at least one channel $\tilde{m}_i \neq \varphi^{-1/2}$ must hold. By (60), the paired weights $(\tilde{m}_i^{(+)}, \tilde{m}_i^{(-)})$ satisfy $\tilde{m}_i^{(+)} \tilde{m}_i^{(-)} = \varphi^{-1}$. Strict convexity of J implies

$$J\big(\tilde{m}_i^{(+)}\big) + J\big(\tilde{m}_i^{(-)}\big) \;>\; 2J\big(\varphi^{-1/2}\big),$$

contradicting minimality of the eight–tick window under the unique cost J. Therefore $\tilde{m}_i = \varphi^{-1/2}$ for all three channels, and the only isotropic scalar prefactor consistent with triad symmetry is $C = \varphi^{-3/2}$.

E. Consequence for the ILG kernel

With $C = \varphi^{-3/2}$ fixed and the RS exponent $\alpha = \frac{1}{2}(1 - \varphi^{-1})$, the information–limited kernel in Fourier space is

$$w(k, a) \;=\; 1 + \underbrace{\varphi^{-3/2}}_{\text{prefactor from eight–tick triads}} \Big(\frac{a}{k\,\tau_0}\Big)^{\alpha}.$$

No free parameters are introduced: C and α are fixed by the eight–tick geometry, the unique cost J, and the ledger link penalty, all within the RS axioms.

F. Remark: General dimension

If the microscopic space had dimension D, the same argument (one primal/dual pair per channel, triad replaced by D–tuple) would produce $C = \varphi^{-D/2}$. In the ILG setting we work with $D = 3$, giving $C = \varphi^{-3/2}$.

B Prefactor $\varphi^{-3/2}$ from eight–tick geometry (with uniqueness)

// ... existing code ...

References