

The Participatory Decoder: Architecture for the Participatory Horizon Programme — Part III

Gregory O’Grady*

April 27, 2026

Abstract

Persistent large-angle anomalies in observational cosmology, including suppressed CMB correlations beyond 60° , excess source-count dipoles, and quadrupole–octopole alignment with the kinematic dipole and planarity (‘Axis of Evil’), share an observer-dependent character that primordial mechanisms should not produce. The participatory-horizon programme proposes that these anomalies arise because the observer’s causal diamond acts as an information aperture whose geometry constrains which correlations are realised as classical records, while the standard Λ CDM cosmology is retained in full. The companion PMH framework [1] derives an angular recording filter $h^2(\ell)$ and a radial kernel $k_L(x)$ from the conformal geometry of the causal diamond, and the companion Lindblad paper [2] constructs the local irreversible recording dynamics with canonical collapse operator $\hat{R}_a = \delta \hat{K}_a / \Sigma_W$. Both papers identified a variance channel, m -dependent modulation from anisotropic position-space recording, but deferred its future quantitative development.

The central question addressed is whether position-space recording on the causal diamond can produce the observed quadrupole–octopole alignment and low- ℓ planarity, and if so, what the framework predicts for the pattern of m -dependent variance modulation in the low- ℓ CMB sky.

The framework yields exact, parameter-free variance coefficients, $\kappa_{20} = 2/7$, $\kappa_{21} = 1/7$, $\kappa_{22} = -2/7$, locked to the kinematic dipole direction. The decoder concentrates Fisher information on zonal ($m = 0$) modes, suppressing their posterior variance while leaving sectoral ($|m| = \ell$) posterior variance relatively enhanced, thereby producing the geometry of the planar morphology observed in the CMB quadrupole. A parity selection rule, following from the quadratic structure of the Lindblad hazard, forces the canonical decoder to be antipodally symmetric: it cannot produce a diagonal even/odd C_ℓ parity asymmetry. A common-axis theorem shows that independent bias of the quadrupole and octopole axes toward \hat{v} enhances their mutual alignment without requiring direct $\ell = 2 \leftrightarrow 3$ covariance coupling.

The amplitude falls short of the observational target by 15–17 orders of magnitude, and this gap is shown to be structural within the canonical recording architecture. The geometric predictions, i.e., the $\kappa_{\ell m}$ variance pattern, the parity selection rule, and the common-axis theorem, hold regardless of the amplitude and constitute parameter-free predictions of the framework.

1 Introduction

Large-angle anomalies in the cosmic microwave background (CMB) and in extragalactic source counts have resisted conventional explanations for over two decades. CMB two-point correlations are suppressed beyond 60° [3–5], radio and mid-infrared source-count dipoles exceed

*Professor, University of Auckland, New Zealand. Parts of the computational workflow and manuscript preparation used AI-assisted tools under human oversight.

the kinematic expectation by factors of two to four [6–8], and the quadrupole and octopole principal axes cluster near each other and near the kinematic dipole direction [9–12]. This latter anomaly, the quadrupole–octopole alignment, sometimes called the ‘Axis of Evil’, is the primary target of this paper. This alignment has been quantified through multiple statistics including the angular-momentum dispersion, area vectors, and multipole-vector mutual angles, with reported significances typically in the range 10^{-2} to 10^{-3} [4, 11], though interpretation and statistical relevance remains debated. Recent work examining the joint distribution of several large-angle anomaly statistics also argues that their collective occurrence may be substantially less probable than any individual p -value suggests [12].

A peculiar feature shared by a number of these anomalies is their association with the local observer’s horizon and motion, with the preferred directions clustering near the kinematic dipole, a pattern that primordial mechanisms would not naturally produce [3, 4, 9]. The participatory-horizon programme proposes that these anomalies arise not from new cosmological physics but from the observer’s finite causal domain acting as an information aperture [1, 13, 14]. The standard Λ CDM cosmology is retained in full, but the programme adds a recording layer defined by the geometry of the observer’s causal diamond, whose axisymmetry about the observer’s velocity axis generates the anomaly pattern in the realised sky.

Two preceding papers developed the core recording architecture for the programme. The Participatory Modular Hamiltonian (PMH) companion [1] pairs the Casini–Huerta–Myers modular Hamiltonian [15] with a local recording map on the observer’s causal diamond, deriving an angular recording filter $h^2(\ell) = 1 - \exp[-\ell(\ell + 1)/\ell_c^2]$ and a radial kernel $k_L(x) = x^L(1 - x^2)^4$ from the conformal geometry of the diamond and the thermal structure of the modular Hamiltonian on hyperbolic space H^3 . The angular filter resolves the low-power anomaly, shifting the $S_{1/2}$ statistic from the 8.3rd to the 49.7th percentile of Λ CDM realisations [14]; the radial kernel addresses the source-count dipole tension [16, 17]. The Lindblad companion [2] constructed the local irreversible recording dynamics on regulated modular cells, proved that the canonical collapse operator $\hat{R}_a = \delta\hat{K}_a/\Sigma_W$ is unique at leading order within its stated class, and showed that the PMH recording kernel is the square root of the Lindblad excess rate, grounding the pointwise signal-to-noise prescription as a consequence of the dynamics rather than an independent hypothesis.

Both papers identified the requirement for a third structural layer. When the Lindblad collapse triggers at a specific angular cell, it conditions all multipole components $a_{\ell m}$ contributing to the modular-energy density at that cell. If the recording rate is isotropic, this conditioning averages out and all m -components within a sector are treated equally. But the kinematic dipole makes the recording rate anisotropic — cells near the velocity axis \hat{v} trigger preferentially, conditioning $m = 0$ modes while leaving $|m| \geq 1$ modes comparatively unconstrained [2]. This m -dependent variance modulation is a direct consequence of the position-space structure of the Lindblad dissipator combined with the anisotropy of the cosmological signal; it requires no new postulates beyond irreversible first-hit recording (P_{rec}). Quantitative development was explicitly deferred to future work, and is now resumed in the present paper.

The importance of the decoder layer extends beyond the quadrupole–octopole alignment. Neither the sector gate $h^2(\ell)$ nor the cell-level trigger rate Γ_a alone determines the m -resolved coefficients, axis statistics, or parity structure that enter observational anomaly diagnostics. The decoder closes this gap by specifying the retained classical record, the local measurement equation, and the Fisher information kernel that connects them. The primary observational target is the quadrupole–octopole alignment, but the structural results, i.e., the parity selection rule and the m -sum obstruction in particular, also determine what the canonical architecture can and cannot explain, shaping the path forward for anomalies not yet within the programme’s scope, notably the CMB parity anomaly.

This paper delivers two main outcomes: exact, parameter-free geometric predictions for

the variance structure of the low- ℓ sky, and a quantification of the amplitude gap problem that has accompanied the programme since its inception [16, 18]. The construction proceeds in three stages. First, the measurement equation relating the pattern of fired modular cells to the low- ℓ sky is derived from the Lindblad collapse operator, confirming that the angular structure $Y_{\ell m}(\hat{n}_a)$ is preserved at leading order. Second, the parity-resolved Fisher information kernel decomposes through Gaunt integrals into an even channel with exact variance coefficients $\kappa_{20} = 2/7$, $\kappa_{21} = 1/7$, $\kappa_{22} = -2/7$, locked to the kinematic dipole direction; a parity selection rule forces the fired-cell density to be antipodally symmetric, and a common-axis theorem shows that independent bias of the quadrupole and octopole axes toward \hat{v} enhances their mutual alignment without requiring direct $\ell = 2 \leftrightarrow 3$ covariance coupling. Third, the canonical amplitude is shown to fall approximately fifteen to seventeen orders of magnitude below the observational target, and this gap is structural within the canonical recording architecture; the geometric predictions hold regardless.

This paper is the third of a companion triplet constructing the canonical recording architecture of the participatory horizon programme. The PMH companion [1] determines what is recordable through the angular filter and radial kernel; the Lindblad companion [2] determines how recording happens through the canonical collapse operator; the present paper determines what the observer sees through the decoder’s variance geometry, parity rule, and common-axis mechanism. Figure 1 provides a programme-level overview.

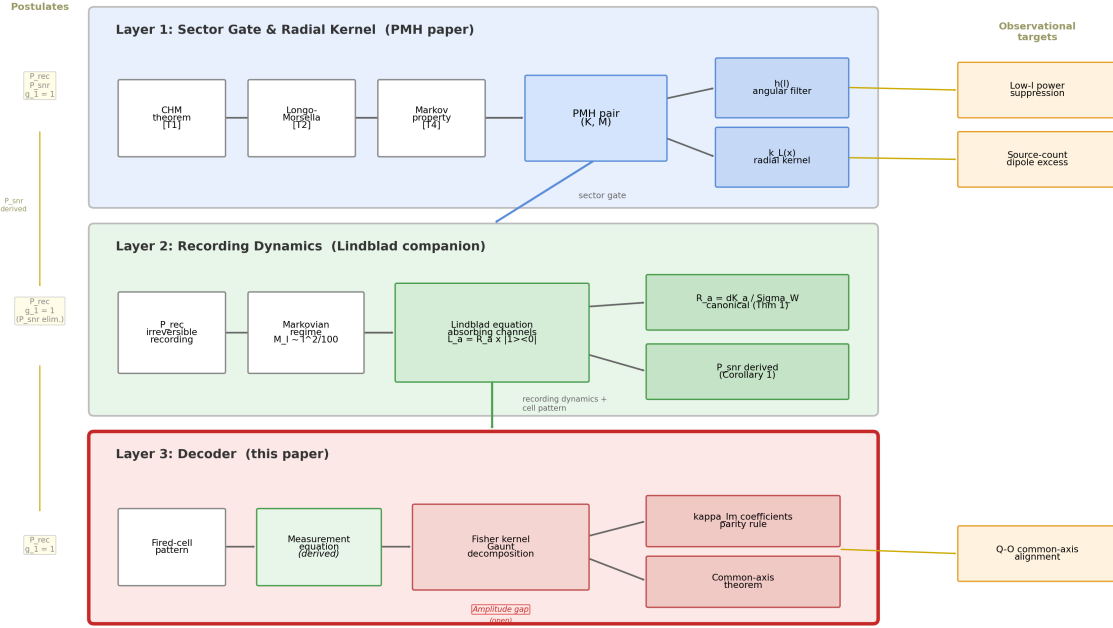


Figure 1: **Programme architecture of the participatory horizon engine.** Three layers answer successive questions: Layer 1 (the PMH companion [1]) determines *what is recordable* through the angular filter $h^2(\ell)$ and the radial kernel $k_L(x)$; Layer 2 (the Lindblad companion [2]) determines *how recording happens* through the canonical collapse operator and absorbing channels; Layer 3 (this paper) determines *what the observer sees* through the parity-resolved Fisher kernel and common-axis mechanism. Solid arrows denote theorem-level or derived connections; dashed arrows denote downstream observational targets or postulate flow. The postulate count decreases from top to bottom: P_{snr} is eliminated at Layer 2 (Corollary 1 of the Lindblad companion). The measurement equation at Layer 3 is derived in this paper (Section 3.3); the amplitude of the variance channel remains the principal open quantity. The present paper develops Layer 3 (red border).

2 Inputs and Architecture

The decoder builds on two preceding layers. Figure 2 provides an overview of the decoder’s logic, from the polar-first recording mechanism through the geometric predictions to the observable sky; Figure 1 places the decoder within the programme’s three-layer architecture.

The angular recording filter

$$h^2(\ell) = 1 - \exp\left[-\frac{\ell(\ell+1)}{\ell_c^2}\right] \quad (1)$$

determines whether each angular sector enters the classical record [1]. This filter is m -independent: it treats all m -components within a given sector equally. The radial recording kernel $k_L(x) = x^L(1-x^2)^4$ determines how recording fidelity varies with conformal depth.

The Lindblad master equation provides the cell-level recording dynamics [2]. Each regulated modular cell a carries an absorbing collapse operator $\hat{L}_a = \hat{R}_a \otimes |1\rangle\langle 0|$, where the system-side operator is the canonical modular-energy fluctuation

$$\hat{R}_a = \frac{\delta \hat{K}_a}{\Sigma W}, \quad (2)$$

unique at leading order within its stated class. The cell recording rate is $\Gamma_a = \gamma_0 \langle \hat{R}_a^2 \rangle_\rho$, and the informative excess from a coherent signal of angular pattern L is

$$\Delta\Gamma(\hat{n}, x) \propto \left[\frac{S_L(x)}{\sigma_{\mathcal{K}}(x)} \right]^2 Y_{L0}(\hat{n})^2, \quad (3)$$

which is quadratic in the signal amplitude. This quadratic structure has immediate consequences for the parity of the decoder (Section 4).

Neither the sector gate $h^2(\ell)$ nor the trigger rate Γ_a determines the m -resolved coefficients or axis statistics that enter anomaly diagnostics. The decoder closes this gap through three sub-layers: the retained classical record \mathcal{R} (what information survives from each fired cell), the forward model (the measurement equation relating the record to the low- ℓ sky), and the Fisher information kernel connecting them. The conformal-field setting of the PMH and Lindblad companions is inherited throughout; the geometric results of this paper depend only on the angular structure $Y_{\ell m}(\hat{n}_a)$ and are robust to the details of the conformal identification.

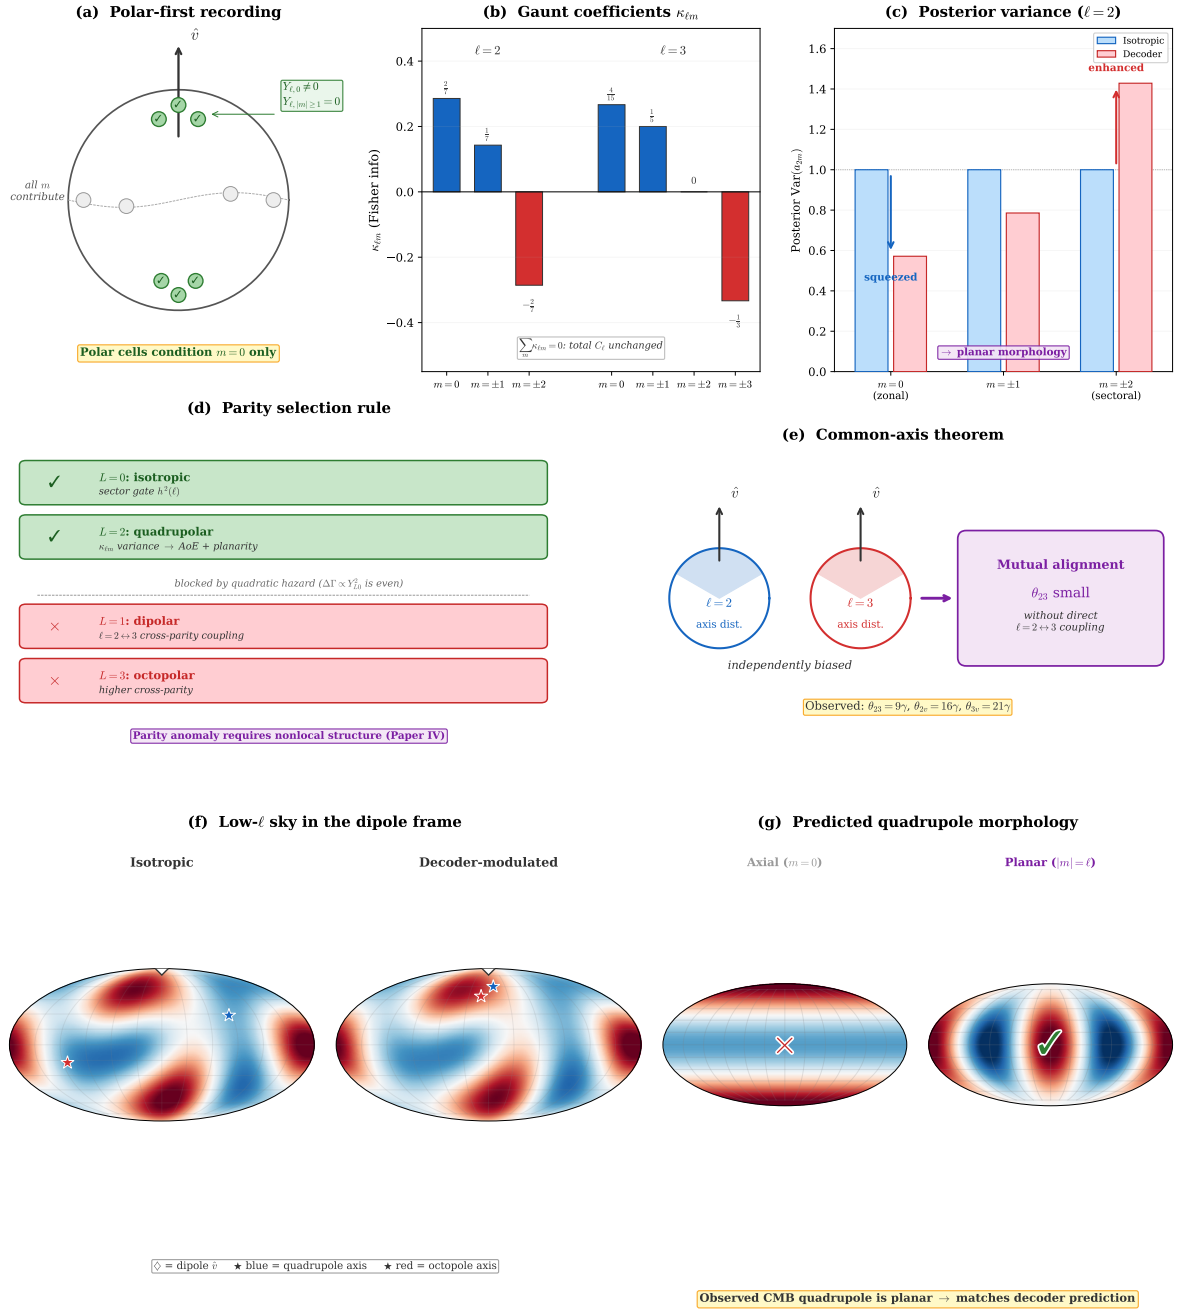


Figure 2: Conceptual overview of the Participatory Decoder. (a) Polar-first recording: cells near \hat{v} fire preferentially; at the poles only $m = 0$ modes are conditioned. (b) Exact Gaunt coefficients $\kappa_{\ell m}$: positive values indicate enhanced Fisher information, summing to zero at each ℓ . (c) Posterior variance at $\ell = 2$: the Fisher–variance inversion squeezes zonal ($m = 0$) posterior variance and enhances sectoral ($|m| = \ell$), producing planar morphology. (d) Parity selection rule: even channels open, odd channels blocked by the quadratic hazard. (e) Common-axis theorem: independent bias toward \hat{v} produces mutual alignment. (f) $\ell = 2, 3$ sky: isotropic (left) vs decoder-modulated (right); diamonds mark \hat{v} , stars mark quadrupole (blue) and octopole (red) axes. (g) Predicted morphology: the decoder favours planar over axial configurations, matching the observed CMB quadrupole.

3 The Canonical Decoder

This section defines the canonical retained record, constructs its forward model, and derives the Fisher information kernel from first principles.

3.1 The retained record

The Lindblad first-hit dynamics produces a natural record, effectively comprising a competing race among active cells, in which one cell fires at each stage and is then removed. The canonical retained record consists of:

$$\mathcal{R} = \{(k_1, \Delta t_1), \dots, (k_N, \Delta t_N)\}, \quad (4)$$

where $k_j \in A_j$ is the winner-cell identity at stage j (drawn from the active set A_j) and Δt_j is the inter-hit waiting time. No signed or analog cell-level readout is included in the canonical record; such extensions are discussed in Appendix Appendix A.

This is the natural record for a Poisson competing-risks process: each active cell has hazard rate $\lambda_a(\theta)$, the cell with the highest cumulative hazard triggers first, and the elapsed time is exponentially distributed with rate $\Lambda_j = \sum_{a \in A_j} \lambda_a$.

3.2 The forward model

The local hazard at cell a depends on the low- ℓ sky $\theta \equiv \{a_{\ell m}\}$ through the modular-energy density at the cell's angular position \hat{n}_a and conformal depth x_a :

$$\eta_a(\theta) = \sum_{\ell m} a_{\ell m} Y_{\ell m}(\hat{n}_a) g_\ell(x_a), \quad (5)$$

where $g_\ell(x) = s_\ell(x)/\sigma_\mathcal{K}(x)$ is the dimensionless coupling constant (the ratio of the sector- ℓ signal profile to the local rms noise in flat coordinates).

The total hit probability for the canonical threshold model is

$$w_H(\eta) = p_+(\eta) + p_-(\eta) = P(|z| > \tau \mid \eta), \quad (6)$$

where $z = \eta + \xi$ with $\xi \sim N(0, \sigma^2)$ and τ is the recording threshold. Near the isotropic operating point $\eta = 0$:

$$w_H(\eta) = 2q + \frac{u\phi}{\sigma^2} \eta^2 + O(\eta^4), \quad (7)$$

with $u = \tau/\sigma$, $q = \Phi(-u)$ the tail probability, and $\phi = \phi(u)$ the standard normal density evaluated at u . The absence of a linear term confirms that w_H is even in η .

In the exact derived treatment below, the proxy radial coupling $g_\ell(x)$ is replaced by the ball-averaged response $g_\ell^{\text{dec}}(\rho)$ derived from the Lindblad cell form factor; the proxy should be understood as its leading small-cell approximation.

3.3 The derived measurement equation

The proxy response basis $\psi_a \sim Y_{\ell m}(\hat{n}_a) g_\ell(x_a)$ of the preceding subsection left the radial coupling g_ℓ as an unspecified profile. We now derive g_ℓ from the Lindblad collapse operator and the spectral decomposition of the modular field on H^3 , following the regulated-cell construction of the Lindblad companion [2].

Spectral decomposition of the cell observable. The vacuum-subtracted modular-energy density on H^3 decomposes in the Plancherel basis as

$$\delta\hat{\mathfrak{E}}_H(\rho, \Omega) = \sum_{\ell m} \int_0^\infty ds \hat{q}_{\ell sm} u_{\ell s}(\rho) Y_{\ell m}(\Omega), \quad (8)$$

where $s \in [0, \infty)$ labels the continuous radial spectrum of the scalar Laplacian on H^3 , $u_{\ell s}(\rho)$ are the regular radial mode functions, and $\hat{q}_{\ell sm}$ are spectral coefficients. The smeared cell observable at regulated geodesic-ball cell $W_a = B_\varepsilon(X_a) \subset H^3$ of proper radius ε is then

$$\delta\hat{K}_a = \int_{W_a} dV_{H^3} \delta\hat{\mathfrak{E}}_H = \sum_{\ell m} \int_0^\infty ds B_{\ell sm}(a) \hat{q}_{\ell sm}, \quad (9)$$

with the cell form factor $B_{\ell sm}(a) = \int_{W_a} dV_{H^3} u_{\ell s}(\rho) Y_{\ell m}(\Omega)$.

Exact factorisation for a geodesic-ball cell. Define the ball-average operator $(A_\varepsilon f)(X) = V_\varepsilon^{-1} \int_{B_\varepsilon(X)} f(Y) dV_{H^3}(Y)$, where $V_\varepsilon = \pi(\sinh 2\varepsilon - 2\varepsilon)$ is the volume of a geodesic ball of radius ε on H^3 . Because A_ε is defined by averaging over a geodesic ball, it commutes with all isometries of H^3 and therefore with the Laplacian Δ_{H^3} . It follows that A_ε acts diagonally on each Laplacian eigenspace: for any mode $\Psi_{\ell sm} = u_{\ell s} Y_{\ell m}$,

$$A_\varepsilon \Psi_{\ell sm} = m_\varepsilon(s) \Psi_{\ell sm}, \quad (10)$$

with a scalar multiplier depending only on s , not on ℓ or m . The form factor therefore factorises exactly:

$$\boxed{B_{\ell sm}(a) = V_\varepsilon m_\varepsilon(s) u_{\ell s}(\rho_a) Y_{\ell m}(\hat{n}_a)} \quad (11)$$

All angular structure resides in $Y_{\ell m}(\hat{n}_a)$; all radial spectral filtering resides in the product $m_\varepsilon(s) u_{\ell s}(\rho_a)$.

The spectral multiplier. Using the zonal spherical function $\phi_s(r) = \sin(sr)/(s \sinh r)$ on H^3 , the multiplier evaluates to the closed form

$$m_\varepsilon(s) = \frac{4\pi}{V_\varepsilon} \frac{\sin(s\varepsilon) \cosh \varepsilon - s \cos(s\varepsilon) \sinh \varepsilon}{s(1+s^2)}. \quad (12)$$

Its small-cell expansion is $m_\varepsilon(s) = 1 - (1+s^2)\varepsilon^2/10 + O(\varepsilon^4)$, consistent with the identity $A_\varepsilon f = f + (\varepsilon^2/10) \Delta_{H^3} f + O(\varepsilon^4)$ applied to a Laplacian eigenfunction. Crucially,

$$m_\varepsilon(0) = \frac{4\pi}{V_\varepsilon} (\varepsilon \cosh \varepsilon - \sinh \varepsilon) > 0. \quad (13)$$

The spectral multiplier does not vanish at threshold: the canonical geodesic-ball readout couples to the $s = 0$ continuum.

The derived decoder response. Introduce the source-side transfer $\langle \hat{q}_{\ell sm} \rangle_\theta = T_{\ell s} \theta_{\ell m}$, where $\theta_{\ell m}$ are the latent low- ℓ sky variables and $T_{\ell s}$ encodes the spectral projection of the coherent cosmological signal. The angular factorisation of the cell response is derived from the Lindblad operator and the H^3 spectral decomposition. The radial response becomes computable once the source transfer $T_{\ell s}$ is specified; here it is identified with the PMH signal profile, importing the source-side closure $g_1 = 1$ from the preceding papers [1, 2]. Identifying

$T_{\ell s}$ with the spectral coefficients of the PMH signal profile $S_\ell(\rho) = \int ds T_{\ell s} u_{\ell s}(\rho)$ gives the decoder response as a ball-average of the PMH signal:

$$g_\ell^{\text{dec}}(\rho) = \frac{V_\varepsilon}{\Sigma_W} (A_\varepsilon S_\ell)(\rho) \quad (14)$$

so that the mean cell response is

$$\eta_a(\theta) = \sum_{\ell m} \theta_{\ell m} Y_{\ell m}(\hat{n}_a) g_\ell^{\text{dec}}(\rho_a). \quad (15)$$

This confirms the angular structure of the proxy forward model (5): the response at cell a is a sum over $\theta_{\ell m} Y_{\ell m}(\hat{n}_a)$, with a radial profile g_ℓ^{dec} that is no longer a free function but a derived spectral overlap.

Relation to the recording kernel. The PMH signal profile $S_\ell(x) = x^\ell(1-x^2)$ and the recording kernel $k_L(x) = x^L(1-x^2)^4$ of the PMH companion [1] differ by the noise normalisation: $k_L(x) = S_L(x)/\sigma_{\mathcal{K}}(x)$, where $\sigma_{\mathcal{K}}(x) \propto (1-x^2)^{-3}$ is the local rms modular-energy fluctuation in flat coordinates. The decoder response (14) ball-averages the unnormalised signal S_ℓ ; the factor $1/\Sigma_W$ in the definition supplies the cell-level noise normalisation. In the small-cell limit, and after the flat-coordinate pullback relating Σ_W to $\sigma_{\mathcal{K}}(x)$, g_ℓ^{dec} is proportional to the signal-to-noise profile $k_\ell(x) = S_\ell(x)/\sigma_{\mathcal{K}}(x)$. The amplitude estimate of Section Appendix E uses $k_1 \equiv k_L(x)|_{L=1, x=1/3} \simeq 0.208$, which is the peak of the normalised recording kernel.

Small-cell expansion of the decoder response. For the PMH signal $S_\ell(x) = x^\ell(1-x^2)$ (Section 2.6.1 of the PMH companion [1]), the sector- ℓ hyperbolic Laplacian evaluates to $\Delta_{H^3, \ell} S_\ell = \frac{1}{2} x^\ell(1-x^2)[(\ell+1)x^2 - (\ell+3)]$. The small-cell expansion is therefore

$$g_\ell^{\text{dec}}(x) = \frac{V_\varepsilon}{\Sigma_W} x^\ell(1-x^2) \left[1 + \frac{\varepsilon^2}{20} ((\ell+1)x^2 - (\ell+3)) + O(\varepsilon^4) \right]. \quad (16)$$

At the $L = 1$ recording-kernel peak $x = 1/3$, the correction is $-0.189\varepsilon^2$: the ball average slightly smooths the PMH signal but does not qualitatively alter its shape.

Consequences for the decoder geometry. The factorisation (11) on H^3 is exact for geodesic-ball cells. Corrections enter only when this is pulled back to the finite angular extent of cells in the observer frame, giving $O((\varepsilon/\rho_a)^2)$ departures from the idealised point-centre response. Because the angular structure $Y_{\ell m}(\hat{n}_a)$ is preserved at leading order, the Gaunt decomposition of Section 5.1, the $\kappa_{\ell m}$ coefficients of Section 5.2, the parity selection rule (Proposition 1), and the common-axis theorem (Theorem 1) are all preserved under the replacement of the proxy by the derived form factor. Corrections enter at $O((\varepsilon/\rho_a)^2)$ through the finite angular extent of the cell, where ε is the cell proper radius.

Note on subsequent sections. All Fisher-kernel and posterior-variance results in Sections 3.4–3.7 use the exact factorised form factor (11), with the proxy forward model of Section 3.2 understood as its leading small-cell approximation.

The structural conclusion for the amplitude gap. The nonvanishing of $m_\varepsilon(0)$ (13) means that the canonical local geodesic-ball readout does not annihilate the low- s thermal continuum on H^3 . The decoder noise floor $\sigma_{\text{dec}, \ell}^2$ therefore receives finite contributions from all spectral modes, not only from a narrow band. This strongly indicates that the amplitude gap identified in Section Appendix E is a structural feature of the canonical recording architecture, not an artefact of the proxy forward model. The gap is addressed further in Section Appendix E with the corrected accounting.

3.4 Exact stage likelihood

At stage j , the surviving active cells A_j compete with hazards $\lambda_a(\theta) = \lambda_{*,a} w_H(\eta_a(\theta))$. The joint probability of winner identity k_j and waiting time Δt_j is

$$p(k_j, \Delta t_j | \theta) = \lambda_{k_j}(\theta) \exp\left[-\Delta t_j \Lambda_j(\theta)\right], \quad \Lambda_j(\theta) = \sum_{a \in A_j} \lambda_a(\theta). \quad (17)$$

The full-trajectory log-likelihood is therefore

$$\log p(\mathcal{R} | \theta) = \sum_{j=1}^N \left[\log \lambda_{k_j}(\theta) - \Delta t_j \Lambda_j(\theta) \right] + \text{const.} \quad (18)$$

3.5 Score vector

Define the local score vector at cell a :

$$g_{a,i} = \partial_i \log \lambda_a = \frac{w'_H(\eta_a)}{w_H(\eta_a)} \psi_{a,i}, \quad \psi_{a,i} = \partial_i \eta_a = Y_{\ell m}(\hat{n}_a) g_\ell(x_a), \quad (19)$$

where the index i runs over the low- ℓ sky components $\{a_{\ell m}\}$. The stage- j score is

$$S_i^{(j)} = g_{k_j,i} - \Delta t_j \sum_{a \in A_j} \lambda_a g_{a,i}. \quad (20)$$

The first term is the winner's local score; the second comes from the exponential waiting-time contribution and involves the hazard-weighted average over all active cells.

3.6 Fisher information

The canonical record retains only whether a cell fired, not which threshold branch (+ or -) was crossed. The relevant local Fisher density is therefore that of the unsigned hit probability $w_H(\eta)$:

$$\chi_H(\eta) = \frac{[w'_H(\eta)]^2}{w_H(\eta)}. \quad (21)$$

Since w_H is even in η (7), $w'_H(0) = 0$ and the Fisher density vanishes at the isotropic operating point. Sensitivity to the low- ℓ sky arises only at the kinematic-dipole operating point $\eta_v(\hat{n}_a) \propto v k_1 Y_{10}(\hat{n}_a)$, where $w'_H(\eta_v) \neq 0$. Expanding $\eta_a = \eta_v(a) + \delta\eta_a$ with $\delta\eta_a = \sum_{\ell \geq 2, m} \theta_{\ell m} \psi_{\ell m}(a)$, the stage- j Fisher information for the low- ℓ parameters is

$$F_{ij}^{(j)} = \frac{1}{W_j} \sum_{a \in A_j} \lambda_{*,a} \chi_H(\eta_v(a)) \psi_{a,i} \psi_{a,j}, \quad (22)$$

where $W_j = \sum_{a \in A_j} \lambda_{*,a} w_H(\eta_v(a))$.

A crucial property of χ_H is that it is even in η_v :

$$\chi_H(-\eta_v) = \chi_H(\eta_v), \quad (23)$$

which follows directly from $w_H(-\eta) = w_H(\eta)$. Since $\eta_v \propto Y_{10}(\hat{n})$ and $Y_{10}(-\hat{n}) = -Y_{10}(\hat{n})$, the product $\chi_H(\eta_v(\hat{n}))$ is antipodally symmetric, reinforcing the parity selection rule of Section 4.

Remark 1. *If the observer additionally retains which threshold branch was crossed, the Fisher density decomposes into the sign-resolved form $\chi(\eta) = [p'_+(\eta)]^2/p_+(\eta) + [p'_-(\eta)]^2/p_-(\eta) \geq \chi_H(\eta)$, which is strictly stronger. This extension is analysed in Appendix Appendix A; the canonical results of this paper use only χ_H .*

The competition subtraction. If inter-hit times are discarded and only winner identities are retained, the Fisher changes. For winner identities only, the score at stage j is $g_{k_j,i} - \bar{g}_{j,i}$ where $\bar{g}_{j,i} = \sum_a \pi_{j,a} g_{a,i}$ is the winner-density average. This produces a *centred* Fisher:

$$F_{ij}^{\text{id}} = \sum_{j=1}^N \text{Cov}_{\pi_j}(g_i, g_j), \quad (24)$$

which subtracts the mean outer product $\bar{g}_{j,i}\bar{g}_{j,j}$ from the uncentred kernel. This subtraction is the exact mathematical imprint of competition: the winner's information is defined relative to the loser pool.

For the canonical record with timing retained:

$$F_{ij}^{\text{time}} = \sum_{j=1}^N \mathbb{E}_{\pi_j}[g_i g_j] = F_{ij}^{\text{id}} + \sum_{j=1}^N \bar{g}_{j,i} \bar{g}_{j,j}. \quad (25)$$

The inter-hit times restore the mean outer product, giving a strictly stronger decoder. However, in the small-anisotropy regime, $\bar{g}_{j,i} \sim O(\rho_2)$ for i corresponding to $\ell = 2$ modes, so the competition subtraction affects the even channel only at $O(\rho_2^2)$ — a negligible correction when $\rho_2 \ll 1$.

3.7 Dense-active-set approximation

In the limit of many active cells per stage (each stage sampling from the same effective angular density), the sum over stages gives

$$F_{ij} \approx N_{\text{hit}} K_{ij}, \quad K_{ij} = \int d\Omega \nu(\hat{n}) \psi_i(\hat{n}) \psi_j(\hat{n}), \quad \nu(\hat{n}) = \frac{\lambda_*(\hat{n}) \chi_H(\eta_v(\hat{n}))}{W}, \quad (26)$$

where $\nu(\hat{n}) = \lambda_*(\hat{n}) \chi_H(\eta_v(\hat{n}))/W$ is the effective Fisher measure and N_{hit} is the effective number of informative hits. This is the kernel that enters the parity-resolved decoder theorem (Section 5).

4 The Parity Selection Rule

The Lindblad hazard (7) is quadratic in the collapse operator \hat{R}_a . Combined with the evenness of the Fisher kernel (23), this forces the canonical decoder to be structurally *even*.

4.1 The quadratic hazard and antipodal symmetry

Proposition 1 (Parity selection rule). *For the dipole-driven operating point ($L = 1$), the informative excess recording rate (3) is antipodally symmetric. The normalised fired-cell density $\pi(\hat{n})$ therefore admits only even Legendre moments:*

$$\pi(\hat{n}) = \frac{1}{4\pi} \sum_{L \geq 0} (2L + 1) \rho_L P_L(\hat{n} \cdot \hat{v}), \quad \rho_1 = \rho_3 = \rho_5 = \dots = 0. \quad (27)$$

Proof. The excess rate (3) depends on $Y_{10}(\hat{n})^2$. Under antipodal reflection $\hat{n} \rightarrow -\hat{n}$, the dipole transforms as $Y_{10}(\hat{n}) \rightarrow -Y_{10}(\hat{n})$; hence $Y_{10}(\hat{n})^2 \rightarrow Y_{10}(\hat{n})^2$. The isotropic baseline γ_0 is trivially invariant. Therefore $\Gamma(\hat{n}) = \Gamma(-\hat{n})$, which implies $\pi(-\hat{n}) = \pi(\hat{n})$.

The Legendre expansion of an antipodally symmetric function on S^2 contains only even- L terms, since $P_L(-\mu) = (-1)^L P_L(\mu)$ and the expansion coefficients of an even function vanish for odd L . \square

Remark 2 (Generality and scope). *The parity selection rule applies to the coherent-displacement contribution to the informative excess rate, for which $\Delta\Gamma_a \propto \langle \hat{R}_a \rangle^2$. It is a leading-order result of the canonical coherent kinematic-dipole selector, not a theorem for arbitrary perturbations of the modular state. A more general state perturbation could alter the covariance part of $\langle \hat{R}_a^2 \rangle$, potentially introducing additional angular structures at higher order.*

4.2 The even Fisher kernel

The unsigned Fisher density $\chi_H(\eta)$ defined in (21) is even:

Proposition 2. $\chi_H(-\eta) = \chi_H(\eta)$ for all η .

Proof. Since $w_H(-\eta) = w_H(\eta)$, differentiation gives $w'_H(-\eta) = -w'_H(\eta)$. Therefore $\chi_H(-\eta) = [w'_H(-\eta)]^2/w_H(-\eta) = [-w'_H(\eta)]^2/w_H(\eta) = \chi_H(\eta)$. \square

4.3 Physical consequence

The combined result of Propositions 1 and 2 is that the full decoder kernel K_{ij} (26) has only even- L content in the Gaunt decomposition, since both the fired-cell density $\pi(\hat{n})$ and the Fisher density $\chi_H(\eta_v(\hat{n}))$ are antipodally symmetric. This means:

Corollary 1. *For the canonical hazard-only record (winner identities and inter-hit times), the decoder kernel couples ℓ to ℓ' only when $\ell + \ell'$ is even. In particular, the odd cross-parity coupling $\ell = 2 \leftrightarrow 3$ is identically zero at leading order.*

This is a strong structural constraint. The canonical decoder generates within-sector m -dependent variance modulation (the $\kappa_{\ell m}$ pattern, producing planarity and independent axis bias toward \hat{v}) but not direct cross-sector $\ell = 2 \leftrightarrow 3$ covariance coupling. The common-axis theorem (Section 6) shows that independent axis bias suffices for mutual alignment without such coupling. Cross-parity effects, if present, would require a richer retained record (Appendix Appendix A).

4.4 Scope of the parity rule

The parity selection rule applies to the leading-order hazard driven by the kinematic dipole. Higher-order corrections (e.g., from the kinematic quadrupole or from non-Markovian effects) can introduce small even and odd corrections, but these are suppressed by additional powers of v . The rule is a leading-order structural result, not an exact symmetry.

5 The Even Decoder Channel

The even channel is the leading-order decoder effect for the canonical record. This section derives the Gaunt decomposition, computes the exact $\kappa_{\ell m}$ coefficients, and establishes the universal tensor anisotropy law.

5.1 Gaunt decomposition

Take the low- ℓ response basis $\psi_a \sim Y_{\ell m}(\hat{n}_a)$ and expand the effective Fisher measure $\nu(\hat{n}) = \lambda_* \chi_H(\eta_v(\hat{n}))/W$ in Legendre polynomials. Since $\nu(\hat{n})$ is antipodally symmetric at leading order (Propositions 1 and 2), only even L moments appear, and the kernel decomposes as

$$K_{\ell m, \ell' m'} = \delta_{mm'} \sum_{L \geq 0} \rho_L^{\text{eff}} \mathcal{G}_{\ell m, \ell' m'}^{(L)}, \quad (28)$$

where ρ_L^{eff} is the effective L -th Legendre moment of the combined density $\pi \chi$, and the Gaunt coefficient is

$$\mathcal{G}_{\ell m, \ell' m}^{(L)} = \int d\Omega P_L(\hat{n} \cdot \hat{v}) Y_{\ell m}(\hat{n}) Y_{\ell' m}^*(\hat{n}). \quad (29)$$

This integral is nonzero only when the standard selection rules are satisfied:

1. same m only: $m = m'$;
2. triangle rule: $|\ell - \ell'| \leq L \leq \ell + \ell'$;
3. parity: $\ell + \ell' + L$ must be even.

For the even channel ($L = 0, 2, 4, \dots$), the parity condition requires $\ell + \ell'$ to be even, so the even channel couples ℓ to ℓ (diagonal) and ℓ to $\ell \pm 2$ (cross-sector, same parity). The odd channel ($L = 1, 3, \dots$) would couple ℓ to $\ell \pm 1$ (cross-parity), but this is absent from the canonical decoder by Corollary 1.

5.2 Explicit Gaunt calculation for $L = 2$

The leading anisotropic contribution comes from $L = 2$. Using the decomposition $P_2(\hat{n} \cdot \hat{v}) = (4\pi/5) \sum_M Y_{2M}(\hat{n}) Y_{2M}^*(\hat{v})$ and the three- $Y_{\ell m}$ integral

$$\int d\Omega Y_{\ell_1 m_1} Y_{\ell_2 m_2} Y_{\ell_3 m_3} = \sqrt{\frac{(2\ell_1 + 1)(2\ell_2 + 1)(2\ell_3 + 1)}{4\pi}} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ m_1 & m_2 & m_3 \end{pmatrix}, \quad (30)$$

the diagonal Gaunt coefficient for sector ℓ , azimuthal order m , at $L = 2$ evaluates to

$$\mathcal{G}_{\ell m, \ell m}^{(2)} = \frac{1}{4\pi} \kappa_{\ell m}, \quad (31)$$

where

$$\kappa_{\ell m} = (-1)^m (2\ell + 1) \begin{pmatrix} \ell & 2 & \ell \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} \ell & 2 & \ell \\ -m & 0 & m \end{pmatrix}. \quad (32)$$

Evaluating the $3j$ symbols explicitly:

Quadrupole ($\ell = 2$). The $3j$ symbol $\begin{pmatrix} 2 & 2 & 2 \\ 0 & 0 & 0 \end{pmatrix} = -\sqrt{2/35}$. The m -dependent structure comes from $\begin{pmatrix} 2 & 2 & 2 \\ -m & 0 & m \end{pmatrix}$, giving:

$$\kappa_{20} = \frac{2}{7}, \quad \kappa_{21} = \frac{1}{7}, \quad \kappa_{22} = -\frac{2}{7}. \quad (33)$$

Note that $\sum_m \kappa_{2m} = 2/7 + 2 \times 1/7 + 2 \times (-2/7) = 0$: the even channel redistributes variance across m without changing the total sector power.

Octopole ($\ell = 3$). Similarly:

$$\kappa_{30} = \frac{4}{15}, \quad \kappa_{31} = \frac{1}{5}, \quad \kappa_{32} = 0, \quad \kappa_{33} = -\frac{1}{3}. \quad (34)$$

Again $\sum_m \kappa_{3m} = 0$.

Axis convention. Throughout this paper, the ‘preferred axis’ of a multipole refers to the principal eigenvector of the angular-momentum dispersion tensor, which is the normal to the plane of maximum power concentration. A multipole ‘biased toward \hat{v} ’ therefore has its power concentrated in a great-circle belt perpendicular to \hat{v} , not in lobes along \hat{v} .

Physical interpretation. The signs tell a clear story: $\kappa_{\ell,0} > 0$ and $\kappa_{\ell,|\ell|} < 0$ for both $\ell = 2$ and $\ell = 3$. Modes aligned with the preferred axis \hat{v} ($m = 0$) receive more Fisher information from the decoder than modes perpendicular to it ($|m| = \ell$). Through the posterior variance equation (35), this translates into *suppressed* zonal posterior variance and *enhanced* sectoral posterior variance: the decoder squeezes the $m = 0$ component while leaving high- $|m|$ modes relatively unconstrained. The resulting morphology is planar, with power concentrated in a great-circle belt perpendicular to \hat{v} rather than along the axis — consistent with the observed planarity of the CMB quadrupole.

5.3 Posterior variance modulation

Let $\Sigma_\ell^{(0)} = C_\ell/(1 + \lambda_\ell)$ be the isotropic posterior variance, where $\lambda_\ell = N_{\text{hit}} C_\ell/(4\pi \sigma_{\text{dec}}^2)$ is the isotropic information load. To leading order in ρ_2 , the m -dependent posterior variance is

$$\frac{\delta\Sigma_{\ell m}}{\Sigma_\ell^{(0)}} \approx -5\rho_2 \kappa_{\ell m} \frac{\lambda_\ell}{1 + \lambda_\ell}. \quad (35)$$

This is the central quantitative result of the even channel. The modulation is:

- proportional to ρ_2 (the $L = 2$ Legendre moment of the effective Fisher measure ν ; see Section 5.1 and Appendix ??);
- proportional to $\kappa_{\ell m}$ (the parameter-free Gaunt coefficient);
- modulated by the factor $\lambda_\ell/(1 + \lambda_\ell)$, which saturates at unity for $\lambda_\ell \gg 1$.

5.4 Universal tensor anisotropy

The angular-momentum dispersion tensor $T_{ij}^{(\ell)} = \sum_m a_{\ell m}^2$ (angular-momentum weights) $_{ij}$ has its principal eigenvector defining the sector's preferred axis. Under the $Q_{\ell m}$ variance perturbation (35), the mean tensor anisotropy is

$$A_\ell \equiv \frac{\langle T_{zz}^{(\ell)} \rangle - \langle T_{xx}^{(\ell)} \rangle}{\text{Tr}\langle T^{(\ell)} \rangle} = \frac{\rho_2}{2} \frac{\lambda_\ell}{1 + \lambda_\ell}. \quad (36)$$

Accordingly, the ℓ -dependent Gaunt factors in the decoder kernel and the ℓ -dependent prefactors in the tensor anisotropy cancel exactly, producing a formula that is the same for all ℓ . The only sector-dependence enters through the information load λ_ℓ , which varies between sectors because C_ℓ is ℓ -dependent.

5.5 Axis susceptibilities

The axis-bias coefficient β_ℓ relates the tensor anisotropy to the actual axis distribution:

$$\beta_\ell \approx s_\ell A_\ell, \quad (37)$$

where s_ℓ is the *axis susceptibility*: the ratio of the axis-distribution P_2 moment to the mean tensor anisotropy, computed from Monte Carlo realisations of the anisotropic Gaussian distribution. We find:

$$s_2 \approx 1.4, \quad s_3 \approx 1.5. \quad (38)$$

Both are order unity: the nonlinear axis-finding map neither amplifies nor suppresses the variance-channel seed.

6 The Common-Axis Theorem

A key simplification emerges when the Axis of Evil alignment problem is reframed at the axis level rather than the coefficient level.

6.1 Statement and proof

Theorem 1 (Common-axis alignment). *Let the quadrupole and octopole axis distributions be independently axisymmetric about the same preferred direction \hat{v} , with Legendre expansions*

$$f_\ell(\hat{n}) = \frac{1}{4\pi} \sum_{L \geq 0} (2L+1) a_L^{(\ell)} P_L(\hat{n} \cdot \hat{v}), \quad a_0^{(\ell)} = 1. \quad (39)$$

Then the probability density of their mutual cosine $\gamma = \hat{n}_2 \cdot \hat{n}_3$ is

$$p(\gamma) = \frac{1}{2} \sum_{L \geq 0} (2L+1) a_L^{(2)} a_L^{(3)} P_L(\gamma). \quad (40)$$

Proof. The joint density of the two axes is $f_2(\hat{n}_2) f_3(\hat{n}_3)$ (independence). The mutual cosine is the inner product $\gamma = \hat{n}_2 \cdot \hat{n}_3$. Write $\hat{n}_2 \cdot \hat{n}_3 = P_1(\cos \Theta_{23})$ and apply the addition theorem for spherical harmonics:

$$P_L(\hat{n}_2 \cdot \hat{n}_3) = \frac{4\pi}{2L+1} \sum_{M=-L}^L Y_{LM}(\hat{n}_2) Y_{LM}^*(\hat{n}_3). \quad (41)$$

The marginal distribution of γ is obtained by integrating $f_2 f_3$ against $\delta(\hat{n}_2 \cdot \hat{n}_3 - \gamma)$ over both sphere variables. By the Funk–Hecke theorem, the integral of an axisymmetric function against a Legendre polynomial yields the corresponding Fourier–Legendre coefficient. Evaluating term by term and using the orthogonality of the Y_{LM} produces (40). \square

Remark 3. *For headless axes (where \hat{n} and $-\hat{n}$ are identified), the distribution is folded to $|\gamma|$ and the even Legendre moments are unchanged. The leading $L = 2$ alignment enhancement used below is therefore identical for directed and headless axes.*

6.2 Lowest-order consequence

For headless axis statistics, the leading anisotropy is $L = 2$. Writing $f_\ell(\hat{n}) \approx (4\pi)^{-1}[1 + 5\beta_\ell P_2(\hat{n} \cdot \hat{v})]$:

$$p(\gamma) \approx \frac{1}{2} [1 + 5\beta_2\beta_3 P_2(\gamma)]. \quad (42)$$

Corollary 2. *Mutual alignment enhancement is quadratic in the individual axis biases. Direct $\ell = 2 \leftrightarrow 3$ covariance coupling is sufficient but not necessary for correlated quadrupole–octopole alignment: independent common-axis bias suffices.*

6.3 Target amplitude

For a headless axis distribution on S^2 , the isotropic cone probability is $P_{\text{iso}}(\theta_{23} < 30^\circ) \approx 0.134$. From (42), the enhancement is

$$\Delta P \approx \frac{5}{2} \beta_2 \beta_3 (u_0 - u_0^3), \quad u_0 = \cos 30^\circ \approx 0.866, \quad (43)$$

giving $\Delta P \approx 0.54 \beta_2 \beta_3$. A 10% absolute enhancement requires $\beta_2 \beta_3 \sim 0.19$, hence each $\beta \sim 0.4$ – 0.5 . From (36) and (37) with $s_\ell \sim 1.5$, this translates to an observable product target of $\mathcal{P} \sim 0.5$, requiring the effective information load $\lambda_{\text{eff}} \gtrsim 1$ — a target against which the internal amplitude must be measured (Appendix ??).

7 Observational Comparison

The decoder framework makes parameter-free predictions about the *shape* of the low- ℓ variance anisotropy, independent of the overall amplitude. This section compares those predictions with the observed CMB.

7.1 The $Q_{\ell m}$ pattern in the dipole-aligned frame

The coefficients $\kappa_{\ell m}$ predict a specific hierarchy of m -dependent variance aligned with the kinematic dipole direction $\hat{v} = (l, b) = (264^\circ, 48^\circ)$ in Galactic coordinates. In the \hat{v} -aligned frame:

- For $\ell = 2$: the ratio $\kappa_{20}/\kappa_{22} = -1$ predicts that $\text{Var}(a_{2,0})$ and $\text{Var}(a_{2,\pm 2})$ are modulated in *opposite directions* by the same amount. The $m = 0$ mode receives more Fisher information and its posterior variance is suppressed; the $|m| = 2$ modes receive less Fisher information and their posterior variance is enhanced.
- For $\ell = 3$: $\kappa_{30}/\kappa_{33} = -4/5$ gives a similar but not identical ratio. The intermediate mode $m = 2$ is unaffected ($\kappa_{32} = 0$), providing a null test.

The observed CMB quadrupole is anomalously planar, with power concentrated in a great-circle belt rather than along an axis [10, 11]. The decoder’s posterior variance equation (35) predicts exactly this morphology: $\kappa_{20} > 0$ suppresses zonal posterior variance while $\kappa_{22} < 0$ enhances sectoral posterior variance, favouring planar configurations. A quantitative comparison requires fixing the amplitude ρ_2 .

7.2 Axis alignment

The observed quadrupole and octopole axes lie within approximately 10° – 25° of the kinematic dipole direction, depending on the statistic used [11, 12]. The common-axis theorem (Theorem 1) provides a natural framework for this alignment: if both axes are independently biased toward \hat{v} , their mutual alignment follows without direct coefficient-space coupling.

The key qualitative prediction is that the preferred direction should be \hat{v} — the kinematic dipole — not an arbitrary sky direction. This is a consequence of the Lindblad excess rate (3) being driven by the kinematic dipole signal, which is the only cosmological anisotropy that enters the modular-energy density at leading order.

7.3 Hierarchy prediction

The parity selection rule predicts a specific ordering of anomaly types:

1. *Power suppression* (from $h^2(\ell)$): the strongest effect, already confirmed at $2\Delta \ln \mathcal{L} = 7.50$.
2. *Within-sector variance geometry* (even channel): the $Q_{\ell m}$ pattern, at $O(\rho_2)$.
3. *Cross- ℓ alignment coupling* (odd channel, if present): requires a signed readout; the systematic seed is at most $O(v^3)$ because the leading Gaunt integral vanishes by orthogonality (Appendix Appendix A).

Within the canonical unsigned hazard architecture, this hierarchy is fixed by the quadratic structure of the Lindblad hazard. It is broadly consistent with the observational ordering, in which low- ℓ power suppression is the most significant anomaly and direct cross-parity statistics are less significant [4].

8 The Amplitude Gap

The geometric results of Sections 5–6 hold for any value of the decoder anisotropy ρ_2 . The observational relevance of the variance channel depends on whether ρ_2 is large enough for the

effects to be detectable.

For the kinematic dipole selector, the excess recording rate scales as $\Delta\Gamma/\Gamma_0 \propto v^2 k_1^2$, where $v \simeq 1.234 \times 10^{-3}$ is the observer’s peculiar velocity and $k_1 \simeq 0.208$ is the recording kernel at peak depth. The P_2 projection of the squared dipole pattern $Y_{10}^2 = (1/4\pi)[1 + 2P_2(\cos\theta)]$ gives a normalised $L = 2$ Legendre moment

$$\rho_2 = \frac{2(vk_1)^2}{5} \approx 2.6 \times 10^{-8}. \quad (44)$$

The full decomposition of this amplitude into geometric anisotropy and effective information load is given in Appendix ??.

The observable variance modulation is proportional to the product of the geometric anisotropy and the effective information load (Appendix ??). The geometric anisotropy $\rho_F = 2/5$ is order unity and parameter-free; the effective information load is suppressed by $(vk_1)^2 \sim 10^{-8}$ and by the sector-projected decoder noise floor, giving an observable product $\mathcal{P} \sim 10^{-17}$ to 10^{-15} — a gap of 15–17 orders of magnitude from the target $\mathcal{P} \sim 0.5$. This gap is structural: the canonical geodesic-ball readout has $m_\varepsilon(0) > 0$ (Proposition ??), so no choice of cell radius or depth can dramatically reduce the decoder noise floor. Nine routes to internal amplification were investigated without success; three closure scenarios are catalogued in Appendix ??.

The $\kappa_{\ell m}$ coefficients are exact Gaunt integrals, the parity selection rule follows from the quadratic Lindblad hazard, and the common-axis theorem is a statement of pure probability. These constitute parameter-free structural predictions that can be tested against the observed CMB independently of any amplitude assumption. The amplitude gap does not invalidate the framework; it constrains where new physics must enter if the AoE directional anomalies are to be derived rather than accommodated.

9 Derivation Ledger

A derivation ledger cataloguing closed results, structural predictions, and open problems is provided in Appendix ??.

10 Discussion

Wheeler famously described his vision of the universe as a kind of ‘self-excited circuit’, illustrating this with his iconic U-diagram, in a manner “meant to inspire thought” [13, 19, 20]. The idea has often been regarded as philosophically profound yet scientifically elusive, and the ambition of this programme has been to reduce this vision to practice, by specific application to anomalies, paradoxes and puzzles in the foundations of physics [13]. The current companion triplet, of which this paper is the third part, now offers the first explicit mathematical form. The PMH companion [1] derives what the observer’s causal diamond can record (the angular filter $h^2(\ell)$ and the radial kernel $k_L(x)$); the Lindblad companion [2] derives how recording happens (canonical local collapse with a unique leading-order operator $\hat{R}_a = \delta\hat{K}_a/\Sigma_W$); while the present paper derives what the observer *decodes* (the $\kappa_{\ell m}$ variance geometry, the parity selection rule, and the common-axis alignment mechanism). Together these drive a self-excited circuit: the Λ CDM cosmology defines the observer’s causal diamond, whose geometry determines the recording architecture, which in turn determines which correlations of the primordial sky are realised as classical records. The loop from cosmos to observer to registered phenomenon is completed and continuously repeated, without modifying the primordial spectrum or the background cosmology. The feedback that breaks the isotropy of the recorded sky is the observer’s own motion.

The framework now addresses four features of the anomaly landscape that primordial mechanisms should not naturally produce. The angular filter resolves the low- ℓ power suppression, shifting $S_{1/2}$ from the 8.3rd to the 49.7th percentile of Λ CDM realisations [3, 5, 14]. The radial kernel predicts a universal dipole excess along the observer’s velocity axis, consistent with the cross-survey analysis of CatWISE2020 and radio source-count data [6, 7, 16, 17]. The decoder derived here produces the geometric skeleton of the quadrupole–octopole alignment and the observed planar morphology of the CMB quadrupole: exact m -dependent variance coefficients locked to the kinematic dipole direction, a parity selection rule from the quadratic Lindblad hazard, and a common-axis theorem showing that independent bias of the quadrupole and octopole axes toward \hat{v} generates their mutual alignment without direct $\ell = 2 \leftrightarrow 3$ covariance coupling [9–12].

Several alternative mechanisms for the quadrupole–octopole alignment have been proposed, including spontaneous isotropy-breaking models [21], Bianchi VII_h cosmologies [22], and primordial preferred-direction modifications to the inflationary power spectrum [23]. Each produces some form of m -dependent variance modulation, but the preferred direction and its amplitude were free parameters. The distinguishing feature of the PMH approach is that the preferred direction, the functional form of the variance anisotropy, and the anomaly hierarchy are all determined by the causal diamond geometry.

The structural predictions for the Axis of Evil may be summarised as follows. The preferred direction is derived to be \hat{v} , the kinematic dipole. The mutual alignment of the quadrupole and octopole axes follows from the common-axis theorem (Theorem 1) without cross-parity covariance coupling. The m -dependent variance hierarchy ($\kappa_{\ell,0} > 0$, $\kappa_{\ell,|\ell|} < 0$) is exact and parameter-free, with $\kappa_{32} = 0$ as a null test. The $\kappa_{\ell m}$ pattern concentrates Fisher information on zonal ($m = 0$) modes, suppressing their posterior variance while leaving sectoral ($|m| = \ell$) posterior variance relatively enhanced; the predicted morphology is therefore planar (power concentrated in a great-circle belt perpendicular to \hat{v}), consistent with the observed planarity of the CMB quadrupole [11, 12]. The parity selection rule also fixes a strict anomaly hierarchy: power suppression (from $h^2(\ell)$) is the strongest effect, within-sector variance geometry (the $\kappa_{\ell m}$ pattern) is next, at $O(\rho_2)$, and cross-parity coupling is suppressed to at most $O(v^3)$. This ordering may match the observational pattern, in which the low- ℓ power deficit is the most significant anomaly and direct cross-parity statistics are weaker [4]. These geometric predictions are exact and parameter-free; the amplitude gap (Appendix ??) remains the programme’s principal open problem.

The even/odd channel decomposition is relevant to another current open problem in this programme — the CMB parity anomaly [4, 24]. The even decoder channel treats all ℓ symmetrically through the universal anisotropy law (36), and the parity selection rule (Proposition 1) proves that the canonical unsigned kinematic-dipole hazard cannot produce an even/odd C_ℓ asymmetry at leading order. Resolving the parity anomaly therefore requires an extension beyond the canonical local architecture, and is the next logical target for an extension of the programme.

Two foundational principles deserve special emphasis. First, while the current proposal gives the observer a special status in symmetry breaking, the framework explicitly preserves the Copernican principle. The preferred direction is not a cosmological axis but a kinematic one, arising from the observer’s peculiar motion relative to the CMB rest frame. Every observer at every location would derive the same architecture with \hat{v} replaced by their own velocity vector; an observer in the Andromeda galaxy would therefore find an identical recording geometry aligned with her own peculiar motion, and a correspondingly distinct but congruent set of low- ℓ anomalies [18]. The universe in this model has no centre, no preferred point, and no preferred ‘observer’ [13]. Second, no retrocausality is involved. The Lindblad master equation is a standard Markovian quantum channel operating in the forward time direction; the modular Hamiltonian provides a thermal state on the causal diamond, not a signalling

channel to the past. The decoherence structure of the absorbing-flag construction (Section 3.1 of the Lindblad companion [2]) ensures that once a cell fires, the bit is irreversibly classical — exactly the kind of ‘irreversible amplification’ that the Bohr-Wheeler postulate requires for a phenomenon to become ‘real’ [13, 19, 25]. As summarised by Wheeler, “The observer does not influence the past. Instead, by his choice of question, he decides about what feature of the object he shall have the right to make a clear statement” [25].

While this programme now has its first complete mathematical architecture, several limitations should be noted. The amplitude gap between the internal variance channel and the observational target remains the programme’s hardest open problem (Appendix ??). The clean mathematical footing applies most directly to massless conformally coupled fields; extension to massive species and the gravitational sector is not established [1]. On the observational side, full-sky polarisation from LiteBIRD [26] would test whether the angular filter applies universally across CMB fields, while spectroscopic tomography such as offered by SPHEREx [27] and Euclid [28] could test the radial kernel’s depth-dependent prediction [8, 17].

In summary, this paper completes the canonical recording architecture that the participatory-horizon programme set out to build. The three layers of recordability, dynamics, and decoding now form a closed mathematical engine grounded in the conformal geometry of the observer’s causal diamond. The engine’s geometric predictions are exact and parameter-free, while its amplitude remains an open frontier. What began as an exploration of Wheeler’s philosophical intuition about a self-excited circuit is therefore now emerging as a falsifiable framework: the $\kappa_{\ell m}$ pattern, the parity selection rule, and the common-axis theorem offer concrete predictions that the next generation of CMB and large-scale-structure surveys could probe.

References

- [1] Gregory O’Grady. The participatory modular Hamiltonian: Architecture for the participatory horizon programme — part I. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.19713513>, 2026.
- [2] Gregory O’Grady. The participatory Lindblad equation: Architecture for the participatory horizon programme — part II. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.19722842>, 2026.
- [3] Craig J. Copi, Dragan Huterer, Dominik J. Schwarz, and Glenn D. Starkman. Large-angle anomalies in the CMB. *Advances in Astronomy*, 2010:847541, 2010.
- [4] Dominik J. Schwarz, Craig J. Copi, Dragan Huterer, and Glenn D. Starkman. CMB anomalies after Planck. *Classical and Quantum Gravity*, 33:184001, 2016.
- [5] M. Billi, R. B. Barreiro, and E. Martínez-González. The anomaly of the CMB power with the latest Planck data. *Journal of Cosmology and Astroparticle Physics*, 2024(07):080, 2024.
- [6] Nathan J. Secrest, Sebastian von Hausegger, Mohamed Rameez, Roya Mohayaee, and Subir Sarkar. A challenge to the standard cosmological model. *The Astrophysical Journal Letters*, 937(2):L31, 2022. doi: 10.3847/2041-8213/ac88c0.
- [7] L. Böhme, D. J. Schwarz, P. Tiwari, M. Pashapour-Ahmadabadi, B. Bahr-Kalus, M. Bilicki, C. L. Hale, C. S. Heneka, and T. M. Siewert. Overdispersed radio source counts and excess radio dipole detection. *Physical Review Letters*, 135(20):201001, 2025.
- [8] Nathan J. Secrest, Sebastian von Hausegger, Mohamed Rameez, Roya Mohayaee, and Subir Sarkar. Colloquium: The cosmic dipole anomaly. *Reviews of Modern Physics*, 97(4):041001, 2025. doi: 10.1103/RevModPhys.97.041001.

- [9] Kate Land and João Magueijo. Examination of evidence for a preferred axis in the cosmic radiation anisotropy. *Physical Review Letters*, 95:071301, 2005.
- [10] A. de Oliveira-Costa, M. Tegmark, M. Zaldarriaga, and A. Hamilton. The significance of the largest scale CMB fluctuations in WMAP. *Physical Review D*, 69:063516, 2004.
- [11] Craig J. Copi, Dragan Huterer, Dominik J. Schwarz, and Glenn D. Starkman. Large-scale alignments from WMAP and Planck. *Monthly Notices of the Royal Astronomical Society*, 449(4):3458–3470, 2015. doi: 10.1093/mnras/stv501.
- [12] Joann Jones, Craig J. Copi, Glenn D. Starkman, and Yashar Akrami. Strong evidence against a statistically isotropic universe. *arXiv preprint*, 2023. v3, revised March 2026; submitted to Phys. Rev. D.
- [13] Gregory O’Grady. Extending Wheeler’s participatory universe: A conceptual framework for a ‘Measureverse’. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.17956983>, 2025.
- [14] Gregory O’Grady. The participatory horizon as an angular recording filter: A derived CMB suppression scale. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.19647157>, 2026.
- [15] Horacio Casini, Marina Huerta, and Robert C. Myers. Towards a derivation of holographic entanglement entropy. *Journal of High Energy Physics*, 2011(05):036, 2011. doi: 10.1007/JHEP05(2011)036.
- [16] Gregory O’Grady. Participatory horizons and the cosmic dipole anomaly. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.19528259>, 2026.
- [17] Gregory O’Grady. Participatory horizons and the cosmic dipole anomaly II: A cross-survey analysis. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.19658674>, 2026.
- [18] Gregory O’Grady. Kinematic participatory horizons and the axis of evil. Preprint, Zenodo. <https://doi.org/10.5281/zenodo.17922164>, 2026.
- [19] John A. Wheeler. The ‘past’ and the ‘delayed-choice’ double-slit experiment. In A. R. Marlow, editor, *Mathematical Foundations of Quantum Theory*, pages 9–48. Academic Press, 1978.
- [20] John A. Wheeler and Wojciech H. Zurek. *Quantum Theory and Measurement*. Princeton University Press, 1984.
- [21] Christopher Gordon, Wayne Hu, Dragan Huterer, and Tom Crawford. Spontaneous isotropy breaking: A mechanism for CMB multipole alignments. *Physical Review D*, 72:103002, 2005. doi: 10.1103/PhysRevD.72.103002.
- [22] T. R. Jaffe, A. J. Banday, H. K. Eriksen, K. M. Górski, and F. K. Hansen. Bianchi type VII_h models and the WMAP 3-year data. *Astronomy & Astrophysics*, 460(2):393–396, 2006.
- [23] Lotty Ackerman, Sean M. Carroll, and Mark B. Wise. Imprints of a primordial preferred direction on the microwave background. *Physical Review D*, 75:083502, 2007.
- [24] Jaiseung Kim and Pavel Naselsky. Anomalous parity asymmetry of WMAP 7-year power spectrum data at low multipoles: Is it cosmological or systematics? *Physical Review D*, 82(6):063002, 2010. doi: 10.1103/PhysRevD.82.063002.

- [25] William A. Miller and John Archibald Wheeler. Delayed-choice experiments and Bohr's elementary quantum phenomenon. In S. Kamefuchi et al., editors, *Foundations of Quantum Mechanics in the Light of New Technology*, pages 140–152. Kokubunji, Tokyo, 1984.
- [26] LiteBIRD Collaboration, E. Allys, K. Arnold, J. Aumont, et al. Probing cosmic inflation with the LiteBIRD cosmic microwave background polarization survey. *Progress of Theoretical and Experimental Physics*, 2023(4):042F01, 2023. doi: 10.1093/ptep/ptac150.
- [27] Olivier Doré, J. Bock, M. Ashby, P. Capak, A. Cooray, R. de Putter, T. Eifler, N. Flagey, Y. Gong, S. Habib, and K. Heitmann. Cosmology with the SPHEREx all-sky spectral survey, 2014. arXiv:1412.4872 [astro-ph.CO].
- [28] Euclid Collaboration, Y. Mellier, et al. Euclid. I. overview of the Euclid mission. *Astronomy & Astrophysics*, 697:A1, 2025. doi: 10.1051/0004-6361/202450810.

Appendix A The Signed-Readout Extension

When cell a fires, the post-jump system state is

$$\rho' = \frac{\hat{R}_a \rho \hat{R}_a^\dagger}{\text{Tr}(\hat{R}_a \rho \hat{R}_a^\dagger)}. \quad (45)$$

The operator $\hat{R}_a = \delta \hat{K}_a / \Sigma_W$ is linear in the modular-energy fluctuation and carries its sign. If the observer retains this signed information as $S_a = \text{sgn}(z_a) \cdot H_a$ (where $H_a = \mathbf{1}\{|z_a| > \tau\}$ is the binary hit), the hybrid record (H_a, S_a) has a clean parity split:

$$\mathbb{E}[H_a | \eta] = 2q + \frac{u\phi}{\sigma^2} \eta^2 + O(\eta^4) \quad (\text{even in } \eta), \quad (46)$$

$$\mathbb{E}[S_a | \eta] = \frac{2\phi}{\sigma} \eta + O(\eta^3) \quad (\text{odd in } \eta). \quad (47)$$

The sign-resolved Fisher density $\chi(\eta) = [p'_+(\eta)]^2/p_+(\eta) + [p'_-(\eta)]^2/p_-(\eta)$ is also even in η , since $p_-(\eta) = p_+(-\eta)$ implies that swapping the sign of η merely relabels the two branches. The odd channel therefore appears in the *realised posterior mean*, not in the expected covariance:

$$\delta\hat{\theta} \approx \Sigma_0 \frac{M(u)}{\sigma} \sum_{j=1}^N s_j \psi_{k_j}, \quad (48)$$

where $M(u) = \phi(u)/\Phi(-u)$ is the Mills ratio.

The leading systematic source term for odd sectors involves $\int \pi_H Y_{\ell m} Y_{10} d\Omega$. For $\ell = 3, m = 0$: $\int Y_{30} Y_{10} d\Omega = 0$ by orthogonality. The nonzero contribution enters only through the $\rho_2 P_2$ correction to π_H , making the systematic octopole seed $O(v^3)$.

The cross-coupling Gaunt coefficients for the $\ell = 2 \leftrightarrow 3$ block at $L = 1$ are:

$$c_0 = \frac{3}{\sqrt{35}} \approx 0.507, \quad c_1 = \sqrt{\frac{8}{35}} \approx 0.478, \quad c_2 = \frac{1}{\sqrt{7}} \approx 0.378. \quad (49)$$

Appendix B The Bingham Trajectory Statistic

For the headless hazard class $\lambda(\hat{u} | \hat{n}) = \lambda_0 \exp[\alpha P_2(\hat{u} \cdot \hat{n})]$, the full trajectory likelihood reduces to

$$p(\hat{n} | \mathcal{T}) \propto \exp\left[\frac{3\alpha}{2} \hat{n}^T M_{\mathcal{T}} \hat{n}\right], \quad (50)$$

where

$$M_{\mathcal{T}} = \sum_{j=1}^N \left(\hat{u}_j \hat{u}_j^T - \frac{1}{3} I \right) \quad (51)$$

is the traceless scatter tensor. This is a Bingham posterior on the axis. The mean concentration is $\kappa_{\text{traj}} = N\alpha^2/5$ for small α . With $\alpha \sim v^2 k_1^2 \sim 10^{-8}$, reaching $\kappa \sim 0.5$ requires $N \sim 2.5 \times 10^{16}$ effective hits — not achievable for the low- ℓ sector-projected trajectory.

Appendix C Threshold Tradeoff Bound

Threshold-tuning trades hit count against per-hit anisotropy. The effective even-channel budget satisfies:

$$p_{\text{hit}} \rho_2 \leq \frac{2u\phi(u)}{3\sigma^2} \varepsilon^2 \leq 0.161 \frac{\varepsilon^2}{\sigma^2}, \quad (52)$$

with maximum at $u \approx 1$. No threshold strategy can turn a small ε into a large even channel.

Appendix D Analog Readout Bound

For the continuous analog readout $z = x$ conditioned on $|x| > \tau$, the odd mean response is $a_1(u) = 2[q(u) + u\phi(u)]$, which is maximised at $u = 0$ with $a_1(0) = 1$. The even information budget is bounded by $0.178 \varepsilon^2/\sigma^2$. Thresholding only discards information; the full Gaussian readout is optimal.

Appendix E The Amplitude Gap

The geometric results of Sections 5–6 hold for any value of the effective Fisher anisotropy ρ_2 . In the canonical unsigned model, the Fisher density $\chi_H(0) = 0$: the entire Fisher kernel is proportional to $(vk_1)^2 Y_{10}^2$, giving a geometric anisotropy $\rho_F = 2/5$ (order unity and parameter-free) while the absolute information load is suppressed by $(vk_1)^2 \sim 10^{-8}$ and by the sector-projected decoder noise. The resulting observable product is $\mathcal{P} \sim 10^{-17}$ to 10^{-15} , a gap of 15–17 orders from the target $\mathcal{P} \sim 0.5$. This gap is structural within the canonical recording architecture (Appendix ??).

The $\kappa_{\ell m}$ coefficients are exact Gaunt integrals, the parity selection rule follows from the quadratic Lindblad hazard, and the common-axis theorem is a statement of pure probability. These constitute parameter-free structural predictions that can be tested against the observed CMB independently of any amplitude assumption. The amplitude gap does not invalidate the framework; it constrains where new physics must enter if the AoE directional anomalies are to be derived rather than accommodated.

Scenario A: Source-level closure. The absolute Fisher scale is enhanced to produce $\lambda_{\ell\text{eff}} \sim O(1)$ by additional spurion structure beyond the kinematic dipole — for example, sector-interference coupling between the peculiar bulk flow and the local gravitational potential gradient, projected onto the recording kernel. This requires new source-level physics within Λ CDM but does not modify the recording dynamics. The nine mechanisms above bound the kinematic-dipole contribution at $O(v^2)$; closing the information-load gap from source-level physics alone would require an enhancement mechanism stronger than any yet identified.

Scenario B: Detector-level closure. The information load λ_ℓ is enhanced to $O(1)$ by a reformulation of the recording dynamics that replaces the canonical collapse operator with a richer measurement — for example, a compensated or derivative-type readout for which the form factor $B_{\ell sm}(a) \sim s^p$ ($p \geq 1$) at small s , annihilating the infrared thermal pole. This would require revision of the canonical collapse class (Theorem 1 of the Lindblad companion [2]), abandoning the uniqueness proof that is one of the programme’s strongest results.

Scenario C: The framework provides geometry, not amplitude. The internal engine provides the steering — the parameter-free $\kappa_{\ell m}$ variance pattern, the parity selection rule, the common-axis alignment with \hat{v} — but the amplitude of the observed directional anomalies receives contributions from sources outside the canonical architecture: mask-reconstruction asymmetry, foreground residuals, or statistical variance of the Λ CDM realisation. In this scenario the PMH decoder acts as a geometric lens that shapes whatever amplitude is present into the observed pattern, while the amplitude itself is not a prediction of the framework. The programme’s core successes — the angular filter $h^2(\ell)$ and the source-count dipole kernel $k_L(x)$, where the amplitudes are well-matched — remain the strongest predictions. Scenario C is the most conservative of the three and is adopted as the working position of this paper.

Derivation ledger

Closed results. C1. The measurement equation is derived: the exact form factor factorises as $B_{\ell sm}(a) = V_\varepsilon m_\varepsilon(s) u_{\ell s}(\rho_a) Y_{\ell m}(\hat{n}_a)$ (Section 3.3). C2. The decoder noise floor is closed: $m_\varepsilon(0) > 0$ (Proposition ??); the canonical local readout couples to the IR continuum. C3. The amplitude is quantified: $\mathcal{P} \sim 10^{-17}$ to 10^{-15} , a gap of 15–17 orders from the target ~ 0.5 (Section Appendix E).

Structural predictions. S1. The $Q_{\ell m}$ variance pattern aligns with \hat{v} . S2. The hierarchy even-before-odd follows from the parity selection rule. S3. Mutual alignment arises from independent common-axis bias (Theorem 1). S4. $\kappa_{32} = 0$ is an exact null prediction. S5. Any observed amplitude exceeding $O(v^2)$ requires either an external contribution or a richer measurement equation.

Open problems. O1. Whether the signed readout constitutes a physical retained record (Appendix Appendix A). O2. Source spurion catalogue: what coherent perturbations beyond the kinematic dipole contribute to the excess recording rate. O3. Whether the AoE directional anomalies are derivable from the framework or require external physics.