



The CatWISE2020 Quasar Dipole: A Reassessment of the Cosmic Dipole Anomaly

Masroor Bashir^{1,2} , Pravabati Chingangbam^{1,2} , and Stephen Appleby³ ¹ Indian Institute of Astrophysics, Koramangala II Block, Bangalore 560 034, India; masroor.bashir@iiap.res.in² Department of Physics, Pondicherry University, R.V. Nagar, Kalapet, 605 014 Puducherry, India³ Asia Pacific Center for Theoretical Physics, Pohang 37673, Republic of Korea

Received 2025 November 2; revised 2026 April 24; accepted 2026 April 24; published 2026 May 22

Abstract

The Ellis–Baldwin test probes the cosmological principle by comparing the kinematic cosmic microwave background dipole with the Doppler-driven dipole in the number counts of extragalactic radio sources. Recent analysis of the CatWISE2020 quasar catalog has reported a number-count dipole amplitude exceeding the kinematic expectation at 4.9σ significance. We present a comprehensive reassessment of this test using the same dataset, incorporating major sources of uncertainty in the statistical inference. We employ a simulation framework based on the FLASK package, using lognormal realizations of the large-scale structure, the quasar clustering bias, the survey’s radial selection function, and its exact sky coverage. Our simulations account for the kinematic dipole, the intrinsic clustering dipole, the shot noise, and the survey geometry effects. The analysis yields a revised significance of 3.63σ in the absence of a clustering dipole and 3.44σ with a randomly oriented clustering dipole. When the clustering dipole is aligned with the kinematic dipole, the significance decreases further to 3.27σ . Although the anomaly is reduced, it cannot be explained solely by the clustering dipole or mode coupling from the survey mask. We further assess the dipole measurement robustness by fitting models with successively higher-order multipoles up to $\ell = 4$. Partial sky coverage induces mode coupling, shifting the dipole estimate to higher values when the octopole is included and inflating its variance as additional modes are incorporated, reflected in the increasing condition number of the estimator. This behavior highlights a bias–variance tradeoff inherent in multipole fitting on partial-sky data.

Unified Astronomy Thesaurus concepts: [Cosmological principle \(2363\)](#); [Large-scale structure of the universe \(902\)](#); [Cosmic microwave background radiation \(322\)](#); [Cosmic isotropy \(320\)](#)

1. Introduction

The cosmological principle, which proposes that the Universe is statistically homogeneous and isotropic on sufficiently large scales, forms the bedrock of modern cosmology. This principle is directly tested through observations of the cosmic microwave background (CMB) and large-scale structure (see, e.g., P. Kumar Aluri et al. 2023). The dipole anisotropy observed in the CMB temperature has been conventionally interpreted as a kinematic effect due to our peculiar motion relative to the CMB rest frame. According to the most recent measurement, the velocity inferred from CMB dipole measurements is $369.82 \pm 0.11 \text{ km s}^{-1}$, in the direction $(l, b) = (264^\circ.021, 48^\circ.253)$ (Planck Collaboration et al. 2020a).

An important consequence of this kinematic interpretation is that it must be observable not only in the CMB but also as a corresponding Doppler-induced dipole anisotropy in flux-limited all-sky surveys of sources. The Ellis–Baldwin test can provide a crucial verification of this interpretation by searching for a Doppler-induced dipole in the number counts of distant radio sources (G. F. R. Ellis & J. E. Baldwin 1984). Given a population of sources with a homogeneous and isotropic sky distribution, an observer’s peculiar velocity induces a dipolar modulation in the flux and the number of sources per unit solid angle, from which a kinetic dipole amplitude can be inferred and compared with the value obtained from the CMB.

With the availability of a large number of sources at radio and infrared frequencies, it became feasible to carry out the Ellis–Baldwin test. The 1.4 GHz NRAO VLA Sky Survey (NVSS) data (J. J. Condon et al. 1998) allowed astronomers to measure the dipole, providing an independent check of the kinematic interpretation (C. Blake & J. Wall 2002; A. K. Singal 2011; C. Gibelyou & D. Huterer 2012). Subsequent analysis by M. Rubart & D. J. Schwarz (2013), C. A. P. Bengaly et al. (2018), A. K. Singal (2019), T. M. Siewert et al. (2021), A. K. Singal (2023), J. D. Wagenfeld et al. (2023), M. Land-Strykowski et al. (2025), and L. Böhme et al. (2025) consistently reported a dipole amplitude exceeding the CMB prediction, suggesting potential tensions with the standard kinematic interpretation.

This mystery deepened with more studies. N. J. Secrest et al. (2021) measured a dipole in the CatWISE2020 quasar catalog that was more than twice the CMB expectation at the 4.9σ significance level and aligned at $\sim 27^\circ$ with respect to the CMB dipole direction. L. Dam et al. (2023) confirmed this anomalously large signal using Bayesian inference. Further analysis of the same dataset was carried out by P. Tiwari et al. (2023), testing the data against FLASK simulations (H. S. Xavier et al. 2016), by incorporating quasar clustering properties. They found a clustering dipole amplitude of $D_{\text{clus}} = 0.81 \times 10^{-3}$ ($C_1 = 0.91 \times 10^{-6}$). While substantial, this clustering component cannot fully account for the observed amplitude excess. Notably, for angular scales corresponding to multipoles $\ell \geq 10$ (approximately $>18^\circ$), the active galactic nucleus (AGN)/quasar clustering agrees with Λ CDM expectations. A newer catalog, presented in N. J. Secrest et al. (2022), extended the Wide-field Infrared Survey Explorer (WISE) selection to

include ~ 1.6 million sources. Their analysis yielded a slightly lower but still anomalously high dipole significance of 4.4σ . A recent reassessment of the newer catalog by S. von Hausegger et al. (2026) reported a higher clustering dipole amplitude, $D_{\text{clus}} = 1.25 \times 10^{-3}$ ($C_1 = 2.21 \times 10^{-6}$), suggesting that the clustering contribution in the extended catalog may be larger than previously inferred, though still insufficient to fully account for the observed dipole.

More recent analysis by A. Abghari et al. (2024) specifically demonstrates that the large sky cut ($\sim 52.6\%$) in the CatWISE2020 catalog induces strong mode coupling between low-order multipoles. Consequently, if low-order multipoles (in their example, $\ell = 3$) have amplitudes comparable to the dipole, this coupling can bias the inferred dipole amplitude.

This coupling can result in dipole amplitude errors large enough that consistency with the CMB dipole cannot be ruled out, highlighting the critical challenge of disentangling a true cosmological signal from mask-induced systematics. The conclusions of A. Abghari et al. (2024) were challenged in N. Secrest et al. (2025), with questions raised about the magnitude of the higher- ℓ -mode amplitudes adopted.

In this work, we perform a reassessment of the reported dipole anomaly by implementing a comprehensive treatment of uncertainties potentially underestimated in the previous analyses discussed above. We choose to use the older CatWISE2020 quasar catalog presented in N. J. Secrest et al. (2021), containing approximately 1.3 million sources, for two main reasons. First, this enables a direct and consistent comparison with earlier studies—most notably, N. J. Secrest et al. (2021), which reports the highest dipole significance for a single quasar catalog ($\sim 4.9\sigma$), as well as with subsequent works that relied on the same dataset. Second, our primary objective is to study and quantify the impact of clustering-related systematics on the dipole amplitude and significance when applied to quasar-number-count surveys on a partial sky. This is a broader methodological question, and our conclusions extend beyond a specific catalog. Our analysis is facilitated by the fact that the clustering properties of this particular dataset have already been rigorously characterized by P. Tiwari et al. (2023), including its bias model, radial selection function, and intrinsic clustering dipole, all of which directly inform our simulations.

We systematically quantify the impact of key systematic effects on the significance level of the CatWISE2020 quasar dipole, including: shot-noise contamination, higher-order clustering contributions, the amplitude and orientation of the intrinsic clustering dipole, mask-induced mode coupling, and the propagation of cosmic variance uncertainties in the clustering component. By incorporating these effects through simulations, we derive an updated statistical significance that reflects a more complete error budget, providing an updated evaluation of the dipole’s consistency with the kinematic interpretation.

This paper is structured as follows. Section 2 discusses the details of the CatWISE2020 quasar catalog and the sky-mask geometry used for the analysis. Section 3 presents the formalism for the number-count dipole and outlines the least-squares estimator employed to measure the dipole amplitude. In Section 4, we analyze some sources of uncertainties and bias that affect the dipole measurement. This includes a quantification of the bias introduced by mode mixing due to incomplete sky coverage and an assessment of

the cosmic variance inherent to the clustering signal. Section 5 describes our procedure for generating mock quasar catalogs, which are essential for statistical validation. Section 6 presents our main results, where we quantify the statistical significance of the observed dipole under a series of distinct astrophysical scenarios. Finally, Section 7 provides a discussion of our findings and the implications of our results.

2. CatWISE2020 Quasar Catalog

The quasar sample used in this study is the catalog constructed by N. J. Secrest et al. (2021) from the CatWISE2020 catalog (P. R. M. Eisenhardt et al. 2020; F. Marocco et al. 2021). CatWISE2020 is derived from data from the WISE all-sky satellite survey (E. L. Wright et al. 2010), which imaged the sky in four infrared bands: W_1 ($3.4 \mu\text{m}$), W_2 ($4.6 \mu\text{m}$), W_3 ($12 \mu\text{m}$), and W_4 ($22 \mu\text{m}$). The CatWISE2020 catalog prioritizes the W_1 and W_2 bands, due to their superior sensitivity and depth, being 95% complete for sources brighter than $W_1 \leq 17.4$ mag and $W_2 \leq 17.2$ mag.

The N. J. Secrest et al. (2021) selection applies the mid-infrared color cut $W_1 - W_2 \geq 0.8$ to identify AGNs. This criterion identifies 61.9 ± 5.4 AGNs per square degree to a depth of $W_2 \approx 15.0$ mag, with a reliability of $\leq 95\%$ (D. Stern et al. 2012). Deeper samples can achieve densities of $130 \pm 4 \text{ deg}^{-2}$ for $W_2 \leq 17.11$ with a decreased reliability of $\leq 90\%$ (R. J. Assef et al. 2013). To ensure high-quality photometry for cosmological analysis, the sample also implements a strict magnitude cut of $9.0 < W_1 < 16.4$ to remove bright, saturated sources and faint objects affected by uneven coverage.

Ecliptic Bias Correction. The ecliptic bias is a systematic bias in the source count, resulting in a lower density of sources per pixel near the ecliptic poles and a lower density near the ecliptic equator. We correct for this artifact following N. J. Secrest et al. (2021), using their linear fit as a correction and weighting function in our calculations.

Sky Mask. For our analysis, we employ the CatWISE2020 mask, depicted in the top panel of Figure 1. This mask is constructed by first applying a galactic cut ($|b| < 30^\circ$) that removes 50% of the sky to avoid dust and stellar contamination from the Milky Way. Furthermore, 291 specific extragalactic point sources and extended objects are masked to eliminate contamination from local bright sources and nearby galaxies, removing an additional $\approx 2.6\%$ of the sky. The resulting total effective sky fraction is $f_{\text{sky}} \approx 0.474$ or 47.4% sky coverage.

We compute the areal density of this sample by creating a map using the HEALPIX framework (K. M. Górski et al. 2005) with a resolution parameter $N_{\text{side}} = 64$. The resulting masked density map is shown in the bottom panel of Figure 1. The spectral properties of the sample, characterized by a mean spectral index of $\langle \alpha \rangle = 1.26$ and a mean logarithmic flux slope of $\langle x \rangle = 1.7$, are consistent with a predominantly AGN-dominated population.

3. The Number-count Dipole

Let $N(>S_\nu, \hat{n})$ denote the sky-projected number-count density of a given distribution of observed cosmological “sources” (e.g., the CMB and number-count densities of quasars, galaxies, and X-ray sources, etc.), in the sky direction \hat{n} . S_ν denotes the minimum flux above which a source is detected.




Figure 1. Top: sky mask for the CatWISE2020 quasar sample that excludes the Galactic plane ($|b| < 30^\circ$), regions around bright stars, and other areas with poor photometry, removing 52.6% of the sky. Bottom: sky distribution of the CatWISE2020 quasar catalog used for dipole anisotropy studies.

$N(\hat{n}, >S_\nu)$ can be expanded as

$$N(>S_\nu, \hat{n}) = \bar{N}(1 + \mathbf{D} \cdot \hat{n}) + \mathcal{O}(\ell > 1), \quad (1)$$

where $\bar{N} = N_{\text{tot}}/4\pi f_{\text{sky}}$ is the monopole term given by the average number density, with $N_{\text{tot}}(>S_\nu)$ being the total number of observed sources within the survey area, and $\mathcal{O}(\ell > 1)$ denoting terms of higher multipoles higher than $\ell = 1$. For the CMB and binned flux-limited number counts averaged over a wide redshift range, the higher-multipole terms are usually small—the former because fluctuations in the early Universe are perturbatively small, and the latter because the act of smoothing large-scale structure along the line of sight over gigaparsec scales washes out small-scale clustering information. Hence, for these particular datasets, the monopole and dipole are the dominant modes.

The dipole vector, \mathbf{D} , represents the directional anisotropy in the angular distribution of cosmological sources. It is a vector sum of three distinct contributing vector quantities:

$$\mathbf{D} = \mathbf{D}_{\text{kin}} + \mathbf{D}_{\text{clus}} + \mathbf{D}_{\text{SN}}. \quad (2)$$

The physical origin of each component is distinct. The first one, \mathbf{D}_{kin} , is the *kinematic dipole* that results from the observer’s peculiar motion relative to an assumed cosmological rest frame. For a population of sources with a homogeneous and isotropic sky distribution and a flux density spectral index α (where $S_\nu \propto \nu^{-\alpha}$), a peculiar observer velocity \mathbf{v} induces a dipolar modulation in the number of sources per unit solid angle, N , and in their flux densities, S_ν . N is commonly parameterized as a power law near the flux

limit S_ν of a survey, as $N(>S_\nu) \propto S_\nu^{-x}$. Then, the resulting amplitude of the dipole in N caused by the Doppler boosting and aberration is given by G. F. R. Ellis & J. E. Baldwin (1984):

$$D_{\text{kin}} = [2 + x(1 + \alpha)]\beta, \quad (3)$$

up to order β in the limit $\beta \ll 1$, where $\beta = v/c$.

\mathbf{D}_{clus} is the *clustering dipole*, which originates from large-scale matter inhomogeneities whose gravitational potentials modulate tracer distributions, generating an asymmetry that is intrinsic to the large-scale matter distribution. Last, \mathbf{D}_{SN} , is the *shot-noise dipole*, which is a statistical artifact stemming from finite-source Poisson sampling, scaling as $\sim 1/\sqrt{N}$ in magnitude.

Expanding $N(\hat{n}, S_\nu)$ in spherical harmonics as $N(\hat{n}, S_\nu) = \sum_{\ell m} a_{\ell m} Y_{\ell m}$, the power in each multipole mode is $C_\ell = \sum_m |a_{\ell m}|^2 / (2\ell + 1)$. For full-sky data with $f_{\text{sky}} = 1$, we have $|\mathbf{D}|^2 \propto C_1$.

For each dipole component i (kin, clus, or SN), we have $|\mathbf{D}_i|^2 \propto C_{1,i}$, where $C_{1,i}$ is the power of the $\ell = 1$ component of the spherical harmonic expansion. The total dipole amplitude as derived by C. Gibelyou & D. Huterer (2012) is written as

$$|\mathbf{D}|^2 = \frac{9C_1}{4\pi}, \quad (4)$$

where C_1 has contributions from all three components. For observed data, survey masks are employed to exclude regions contaminated by Galactic sources and other unwanted objects with $f_{\text{sky}} < 1$, breaking the sky’s spherical symmetry. The asymmetry due to the survey mask leads to the leakage of power from higher-order multipoles ($\ell \geq 2$) into the estimated dipole through mode mixing. Consequently, each dipole component becomes dependent on multipoles higher than $\ell = 1$, as $\mathbf{D}_i \equiv \mathbf{D}_i(C_{1,i}, C_{2,i}, \dots)$. The dependence of each \mathbf{D}_i on the multipole hierarchy is then determined by the mask geometry and the underlying power spectrum.

For a generic random field, $f(\hat{n})$, we define its overdensity field as

$$\Delta f = \frac{(f(\hat{n}) - \bar{f})}{\bar{f}}, \quad (5)$$

where \bar{f} denotes its mean value. In the sections that follow, we will use harmonic coefficients of both f and Δf , depending on the context.

3.1. Estimator for the Number-count Dipole

For the sky pixelized into N_{pix} number of pixels, we use the following model for the number-count density in each pixel indexed by $p = 0, 1, \dots, N_{\text{pix}} - 1$ (corresponding to the sky direction \hat{n}_p), given by N. J. Secrest et al. (2021):

$$N(\hat{n}_p) = N_0 + \mathbf{D}_1 \cdot \hat{n}_p + \epsilon_p. \quad (6)$$

N_0 is the parameter for the monopole \bar{N} , and $\mathbf{D}_1 = (D_{1,x}, D_{1,y}, D_{1,z}) = \mathbf{D}N_0$ is the dipole vector in Cartesian coordinates. The last term ϵ_p encapsulates contributions to the number-count density $N(\hat{n}_p)$ that are not described by the simple monopole (N_0) and dipole ($\mathbf{D}_1 \cdot \hat{n}_p$) components. These include the higher-order multipole contributions mentioned in Equation (1) and other sources of error. Masked pixels are excluded from the fitting procedure and assigned NaN values

to prevent them from influencing Equation (6) for the unmasked sky.

We bin the source catalog to construct the random field $N(\hat{n}_p)$. From this, we extract the monopole and the dipole by performing a linear least-squares fit between $N(\hat{n}_p)$ and the right-hand side of Equation (6), keeping up to the dipole term. The best-fit values for the monopole and dipole components, contained in the vector $\boldsymbol{\theta} = (N_0, D_{1,x}, D_{1,y}, D_{1,z})$, are found by minimizing the function

$$\chi^2(\boldsymbol{\theta}) = \sum_p w_p [N(\hat{n}_p) - (N_0 + \mathbf{D}_1 \cdot \hat{n}_p)]^2, \quad (7)$$

where the sum is over all unmasked pixels, and w_p are optional pixel weights (in this case, $w_p = 1$). We solve this system using the `fit_dipole` routine in the HEALPY package (A. Zonca et al. 2019). The best-fit Cartesian dipole vector is then normalized by the monopole to yield the dipole vector:

$$\mathbf{D} = \frac{1}{N_0} \begin{pmatrix} D_{1,x} \\ D_{1,y} \\ D_{1,z} \end{pmatrix}. \quad (8)$$

Equivalently, the dipole amplitude $D = |\mathbf{D}|$ can be described using the angular power spectrum C_ℓ as (A. Abghari et al. 2024)

$$D = 3 \sqrt{\frac{C_1}{C_0}}. \quad (9)$$

This connects the parameter-based fit (N_0, \mathbf{D}) to the harmonic-based power spectrum (C_0, C_1).

The fitting procedure outlined above, specifically minimizing Equation (7), ignores the ϵ_p term in Equation (6). Different physical and numerical effects will contribute to ϵ_p , which in turn add uncertainties into the measured values of N_0 and \mathbf{D}_1 . While the simultaneous fit helps to isolate these components, it does not fully eliminate bias from higher-order moments. These uncertainties will be discussed in the next section.

Throughout this work, we make a distinction between a ‘‘model,’’ referring to the multipoles being fit to the quasar data, and an ‘‘estimator,’’ meaning the least-squares fitting procedure as described in this section.

4. Uncertainties and Bias in the Estimation of the Number-count Dipole

In this section, we examine how masking and cosmic variance broaden the uncertainties and potentially introduce bias into the measured dipole. Using the pseudo- C_ℓ mode-coupling matrix formalism for masked sky, developed in E. Hivon et al. (2002), we quantify the level of leakage of power across different modes for the CatWISE2020 mask. Although a similar analysis was presented in A. Abghari et al. (2024), we include it here for the completeness of our analysis. Our goal is to clarify the underlying mechanism in a manner directly connected to the systematics investigated in this work. The insights gained here will be incorporated into the construction of realistic mock quasar-number-count simulations in Section 5.

4.1. Mask-induced Effects—Mode Mixing

While masking removes major systematic contaminants, it simultaneously distorts the underlying signal, by breaking the

sky’s isotropy and coupling power between different spherical harmonic modes. Consequently, the presence of higher-order multipoles, such as those from large-scale structure clustering and residual systematic gradients, can project into the dipole mode during the fitting procedure, thereby biasing the estimated dipole vector.

The observed harmonic coefficients $\tilde{a}_{\ell m}$ on the cut sky are a convolution of the true full-sky harmonic coefficients $a_{\ell' m'}$ (of $N(\hat{n}, S_\nu)$). This relationship is encoded in the coupling kernel $K_{\ell m}^{\ell' m'}$ (E. Hivon et al. 2002):

$$\tilde{a}_{\ell m} = \sum_{\ell' m'} K_{\ell m}^{\ell' m'} a_{\ell' m'}. \quad (10)$$

The effect of the mask on the angular power spectrum is characterized by the mode-coupling matrix $M_{\ell \ell'}$. The observed power spectrum $\langle \tilde{C}_\ell \rangle$ is a weighted sum of the true spectrum $\langle C_{\ell'} \rangle$:

$$\langle \tilde{C}_\ell \rangle = \sum_{\ell'} M_{\ell \ell'} \langle C_{\ell'} \rangle. \quad (11)$$

The mode-coupling matrix is determined by the power spectrum of the mask itself, as

$$M_{\ell \ell'} = \frac{2\ell' + 1}{4\pi} \sum_L (2L + 1) \mathcal{W}_L \begin{pmatrix} \ell & \ell' & L \\ 0 & 0 & 0 \end{pmatrix}^2, \quad (12)$$

where \mathcal{W}_L is the angular power spectrum of the mask, given by

$$\mathcal{W}_L = \frac{1}{2L + 1} \sum_M |w_{LM}|^2, \quad (13)$$

with w_{LM} being the harmonic coefficients of the spherical harmonic expansion of the mask.

A major focus for this study is the matrix elements $M_{1\ell'}$, which quantify the leakage of power from higher-order multipoles $\ell' > 1$ into the measured dipole ($\ell = 1$). Figure 2 demonstrates the nature of the leakage.

These plots are similar in terms of information content to Figure 6 of A. Abghari et al. (2024). The top panel shows $M_{\ell \ell'}$ as a 2D heat map. The prominent off-diagonal elements demonstrate significant power leakage from higher multipoles (ℓ') into lower multipoles (ℓ). The middle panel provides a 3D representation of the same matrix, visually emphasizing the strength and pattern of this leakage. Both plots highlight how power from a wide range of scales couples with the dipole ($\ell = 1$), posing a significant challenge for an unbiased dipole measurement. The behavior of $M_{1\ell'}$, plotted in the bottom panel, reveals two fundamental characteristics of this leakage for the CatWISE2020 mask:

1. *Decrease with ℓ' .* The amplitude of $M_{1\ell'}$ diminishes with increasing ℓ' , as the coupling to small-scale modes is low.
2. *Parity Effect.* The leakage is significantly stronger for odd- ℓ' multipoles compared to even ones. (See the inset in the bottom panel of Figure 2.)

Notably, when compared to an idealized symmetrical 50% galactic cut, the more complex CatWISE2020 mask with its additional small-scale masking produces a markedly higher level of leakage from even multipoles. Although this effect remains at the $\sim 1\%$ level for the quadrupole ($\ell' = 2$) and hexadecapole ($\ell' = 4$) values considered here, the result

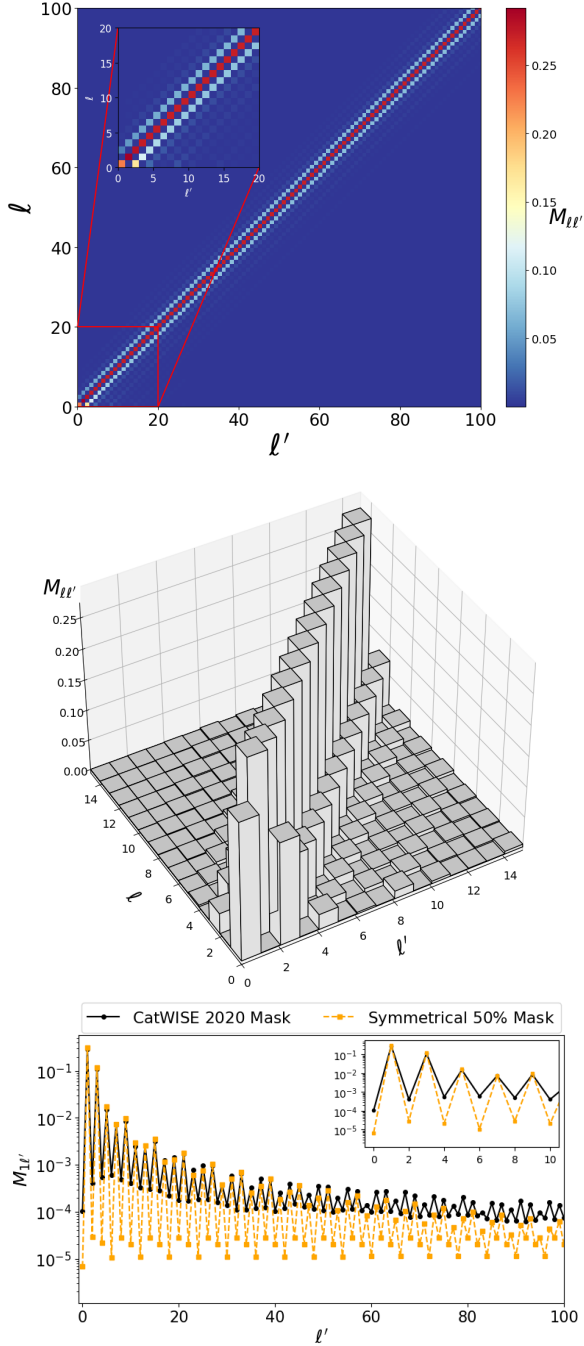


Figure 2. Top and middle panels: the mode-coupling matrix $M_{ll'}$ for the CatWISE2020 survey mask, shown as a 2D heat map (top) and as a 3D representation (middle). The prominent off-diagonal elements demonstrate significant power leakage from higher multipoles (l') into lower multipoles (l). Bottom panel: $M_{1l'}$ plotted versus l' , illustrating the leakage of power from the multipole l' into the measured dipole ($l = 1$). The plot clearly shows the decay of the coupling with increasing l' and the enhanced leakage from odd multipoles compared to even ones (see the inset). For comparison, the coupling for a simple symmetrical galactic cut is also shown.

highlights that even small changes to the mask geometry can affect the statistical significance of the recovered dipole, making an accurate characterization of the mask critically important.

To quantitatively understand the contamination of the dipole measurement by higher-order multipoles, we conducted a simulation analysis involving 1000 realizations of a sky map

containing a dipole signal ≈ 0.007 in the CMB dipole direction.⁴ For each realization, we systematically injected a single higher-multipole power—specifically, the quadrupole (C_2), octopole (C_3), hexadecapole (C_4), or dotricontapole (C_5)—with progressively increasing amplitude. The initial dipole amplitude D_i (before injection) and the final amplitude D_f (after injection) were measured after applying three distinct masks to the simulations: a 10° symmetrical galactic cut, a 30° symmetrical galactic cut, and the asymmetric CatWISE2020 mask. For each mask, we computed the fractional deviation $(D_f - D_i)/D_i$ across all realizations, and analyzed its mean and 1σ dispersion as a function of the injected power. This quantifies the bias induced by the mask in the recovered dipole amplitude.

The results for $(D_f - D_i)/D_i$, summarized in Figure 3, exhibit striking parity-dependent behavior. For even multipoles (C_2 and C_4 ; panels (a) and (c)), the mean fractional deviation remains consistent with zero, indicating no systematic bias. However, the scatter is more pronounced for the CatWISE2020 mask, revealing noise enhancement due to its asymmetry. In contrast, injections of odd multipoles (C_3 and C_5 ; panels (b) and (d)) produce a distinct positive linear trend in $\langle (D_f - D_i)/D_i \rangle$, confirming a leakage of power that systematically inflates the dipole estimate and skews the underlying dipole probability density function (PDF) toward larger values. Notably, the realistic CatWISE2020 mask amplifies both effects, yielding greater variance for even multipoles and bias for odd ones. This underscores the critical role of mask geometry in modulating spurious mode coupling.

4.2. Uncertainties Due to Cosmic Variance

The angular kinematic dipole power, $C_{1,\text{kin}}$, is a deterministic quantity for a particular observer, arising from the observer’s peculiar motion relative to the cosmic rest frame. In contrast, the clustering dipole, $C_{1,\text{clus}}$, is stochastic, originating from the particular distribution of matter in the large-scale structure within our observable volume. As a result, its measured amplitude is inherently uncertain due to the scatter expected from observing one finite realization of a random field, or cosmic variance.

The cosmic variance uncertainty for the clustering dipole is given by the formula

$$\Delta C_{1,\text{clus}} = \sqrt{\frac{2}{3f_{\text{sky}}}} C_{1,\text{clus}}. \quad (14)$$

For $f_{\text{sky}} = 0.474$, the relative uncertainty at 1σ is approximately 118.6% of $C_{1,\text{clus}}$. Consequently, the 1σ and 2σ intervals are:

$$1\sigma : C_{1,\text{clus}} \pm 1.186 C_{1,\text{clus}}, \quad (15)$$

$$2\sigma : C_{1,\text{clus}} \pm 2.372 C_{1,\text{clus}}. \quad (16)$$

This large uncertainty propagates to the total observed dipole, significantly contributing to the error budget and complicating the extraction of the deterministic kinematic component. The exceptionally high cosmic variance for limited sky coverage underscores the challenge in achieving precise measurements

⁴ We repeated our analysis with the dipole randomly directed on the sky, which did not alter our results.

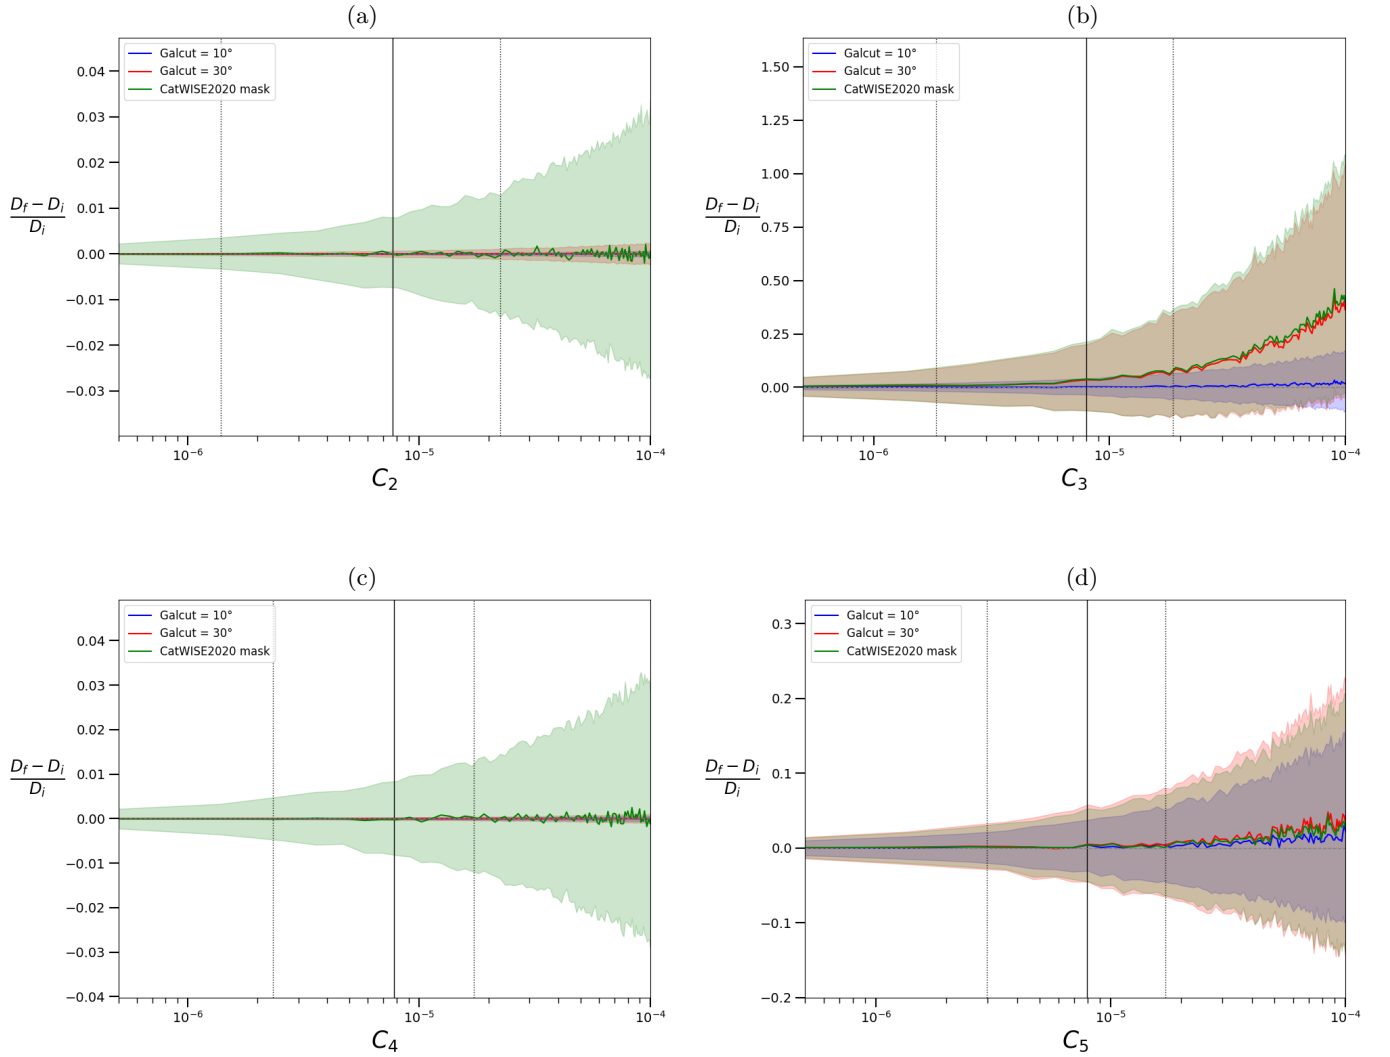


Figure 3. The fractional deviation of the measured dipole amplitude, $(D_f - D_i)/D_i$, as a function of the injected multipole power. The solid lines represent the mean across 1000 realizations, while the shaded regions indicate the $\pm 1\sigma$ uncertainty intervals. Results are shown for the CatWISE2020 mask (orange) and symmetrical 10° (green) and 30° (blue) galactic cuts. Panels (a) and (c) show the results for even multipoles (C_2 and C_4), where no systematic trend is observed. Panels (b) and (d) show the results for odd multipoles (C_3 and C_5), demonstrating a clear positive correlation between the injected power and the fractional dipole deviation. For comparison, the solid and dotted black vertical lines represent the median C_l and their corresponding 95% confidence limits, from FLASK simulations. Note that the black vertical lines shows the *full-sky* power spectrum (C_l) computed directly from the overdensity simulation maps via *healpy*.

of dipole anisotropies and highlights the need for careful modeling that takes these systematic effects into account.

5. Mock Simulations

The statistical significance of the dipole measured from an observed source catalog can be obtained by comparing with the dipole PDF obtained from mock simulations. For accurate comparison, all physical effects that contribute to the dipole must be modeled and included in the simulations. An important shortcoming of the simulations used in N. J. Secrest et al. (2021) is that they included the kinematic dipole and shot-noise component, but not the clustering dipole (deeming it negligibly small) or correlated higher-order modes from large-scale structures. By representing fluctuations only with a white-noise field of equal power across all multipoles, these simulations did not account for the bias that occurs from clustering power, especially from the odd multipoles, that leaks into the measured dipole due to incomplete sky coverage.

It was shown in P. Tiwari et al. (2023) that the angular power spectrum of CatWISE2020 for intermediate and high ℓ agrees well with the Λ CDM expectation. This underscores the need to include these clustering properties in mock simulations.

Here, we extend beyond previous works with improved modeling of the dipole contributions. We generate mock simulations of radio source catalogs incorporating the following: (a) the kinematic dipole; (b) shot-noise dipole systematics; (c) intrinsic higher- ℓ large-scale structure and shot-noise contributions; and (d) the clustering dipole, in this order.

(a) *Incorporating the Kinematic Dipole.* The kinematic dipole component, arising from our motion relative to the CMB rest frame, is simulated directly in pixel space for computational efficiency and physical clarity, following N. J. Secrest et al. (2022). For each pixel i on the sphere, we compute the Doppler factor

$$\delta_i = \frac{(1 + \beta \cos \theta_i)}{\sqrt{1 - \beta^2}}, \quad (17)$$

where θ_i is the angular separation between the pixel center and the direction of the CMB dipole. β is set to 1.23×10^{-3} , corresponding to the measured CMB dipole velocity of $v = 369.82 \text{ km s}^{-1}$ in the direction $(l, b) \approx (264, 48)$ (Planck Collaboration et al. 2020a). The expected number of sources in each pixel is then given by

$$m_i = \delta_i^{2+x(1+\alpha)} N_0, \quad (18)$$

which incorporates the combined effects of relativistic aberration and Doppler boosting, modifying the observed source distribution due to our peculiar motion relative to the CMB rest frame (G. F. R. Ellis & J. E. Baldwin 1984). Here, N_0 represents the sky-averaged monopole density, estimated from the total number count of sources in the survey. Note that the expression on the right-hand side leads to the expression for D_{kin} given in Equation (3), up to order β . The values of the parameters $x = 1.7$ and $\alpha = 1.26$ are taken as the mean values derived from the CatWISE2020 survey catalog, ensuring our simulations reflect the empirical characteristics of the observed quasar population. This process produces a map dominated by the monopole and dipole, with comparatively low power in higher multipoles ($\ell > 1$).

(b) *Incorporating Shot-noise Dipole Contributions.* To generate source counts, we populate each pixel by drawing a random number from a Poisson distribution with mean m_i , following N. J. Secrest et al. (2022). This process introduces the appropriate shot noise corresponding to the discrete nature of the source detection in flux-limited surveys. Shot noise contributes power across all multipoles, adding a scale-independent variance proportional to $1/\bar{N}$. This effect perturbs the initial monopole and dipole signal, introducing additional uncertainties into the measured dipole amplitude and direction, thereby increasing the error budget of the observed dipole.

From the resulting maps, we extract the exact harmonic coefficients for the monopole (a_{00}) and the three components of the dipole ($a_{1,-1}$, $a_{1,0}$, $a_{1,+1}$), using a standard spherical harmonic transform. Here, we isolate only the monopole ($\ell = 0$) and dipole ($\ell = 1$) modes, by setting all coefficients for $\ell \geq 2$ to zero, producing a map that contains purely the monopole, kinematic dipole signal, and its associated shot-noise dipole. This is done so as to not double count the $\ell \geq 2$ contributions of shot noise that are added in step (c).

(c) *Incorporating Higher-multipole Contributions.* We simulate higher-order multipoles ($\ell \geq 2$) using the FLASK package (H. S. Xavier et al. 2016), which generates lognormal random fields tomographically across 35 redshift bins from $z = 0$ to $z = 3.5$ ($\Delta z = 0.1$). The angular auto- ($z_i = z_j$) and cross- ($z_i \neq z_j$) matter power spectra $C_\ell(z_i, z_j)$, which quantify matter clustering across and within redshift bins, are computed using the Boltzmann code CLASS for a fiducial Λ CDM cosmology (Planck Collaboration et al. 2020b). To connect the underlying matter distribution to the observable quasar clustering, we employ a redshift-dependent best-fit bias model $b(z) = 1.54 + 0.53z + 0.50z^2$, from P. Tiwari et al. (2023), and the redshift distribution of the mean number of quasar sources, denoted by $n(z)$, is modeled as $n(z) \propto z^{a_1} \exp[-(z/a_2)^{a_3}]$, where the parameters $a_1 = 0.86$, $a_2 = 1.48$, and $a_3 = 2.36$, which define the shape, peak, and tail of the distribution, are adopted from P. Tiwari et al. (2023).

Let C_ℓ^{SN} denote the angular power spectrum of the shot noise, and let $C_\ell^{\Delta N}$ denote the angular power spectrum of the




Figure 4. Shot-noise-subtracted angular power spectrum ($C_\ell^{\Delta N} - C_\ell^{\text{SN}}$) of the CatWISE2020 quasar distribution overdensity field (ΔN) compared with masked FLASK simulations. The red points show the measured power spectrum from the real data, computed using NaMaster and averaged over five multipole bins. The solid black line represents the mean from 1000 FLASK mock realizations, while the shaded blue region indicates the $\pm 2\sigma$ uncertainty derived from the simulations. The red dotted horizontal line represents the constant contribution from shot noise (C_ℓ^{SN}) over all modes. The faint colored symbols display individual power spectra from the first five FLASK realizations as generated in step (c) of Section 4. The CatWISE2020 measurements lie within the 2σ confidence interval for multipoles $\ell > 5$, indicating consistency with the Λ CDM-based FLASK simulations on intermediate and small angular scales. This figure contains the same information as Figure 6 of P. Tiwari et al. (2023). Note that the overdensity field includes a normalization of $1/\bar{N}^2$, producing the ~ 4 orders of magnitude lower amplitude in this figure relative to Figure 3. The shot-noise subtraction and masking further reduce the amplitude. Thus, the spectra satisfy $C_\ell^{\Delta N} = C_\ell/\bar{N}^2$, where C_ℓ is the angular power spectrum of the CatWISE2020 number density field, even though both reflect the same underlying clustering.

overdensity field ΔN of the CatWISE2020 data (as defined in Equation (5)). Figure 4 shows the shot-noise-subtracted power spectrum $C_\ell^{\Delta N} - C_\ell^{\text{SN}}$. The measured power spectrum agrees well with FLASK simulations within the 2σ confidence interval for $\ell \geq 5$. This agreement validates that the amplitudes of the higher multipoles considered in the simulations here are reasonable and representative of the CatWISE2020 quasar data. This comparison reveals power in higher multipoles that was previously unaccounted for in dipole significance calculations, particularly the substantial clustering contribution at scales $3 \leq \ell < \ell_{\text{max}}$ that can leak into the dipole measurement through mask effects. In general, we use $\ell_{\text{max}} = 3N_{\text{side}} - 1$ for our calculations. For the results that we will present in Section 6, we test the consistency of our analysis against the variation of ℓ_{max} , by systematically decreasing it. We find that $\ell_{\text{max}} \simeq 64$ is sufficient to capture higher-multipole effects, and increasing it beyond this value does not significantly alter our results.

FLASK first produces Gaussian random fields for each redshift shell, then applies a lognormal transformation to capture non-Gaussianity. Correlations between shells are preserved via the input power spectra. The line-of-sight integration across these shells yields a full-sky map, which is converted into a number-count map by Poisson sampling. The number of quasars $N_q(i, j)$ in each pixel i of the redshift shell j is given by

$$N_q(i, j) = \bar{N}_q(j)[1 + \delta_q(i, j)], \quad (19)$$

where $\bar{N}_q(j)$ is the expected number of quasars per pixel in the redshift shell j in a homogeneous Universe, and $\delta_q(i, j)$ is the quasar density field generated according to input power spectra. This introduces shot-noise and discrete sampling

effects. After generating the FLASK map, we compute its spherical harmonic coefficients and explicitly set the monopole ($\ell = 0$) and dipole ($\ell = 1$) modes to zero, retaining only coefficients for $\ell \geq 2$. This yields a map containing pure clustering (plus its shot noise) starting at $\ell = 2$.

(d) *Incorporating the Clustering Dipole Contribution.* The clustering dipole amplitude used in our simulations is $D_{\text{clus}} = 0.81 \times 10^{-3}$ ($C_1 = 0.91 \times 10^{-6}$), adopted from P. Tiwari et al. (2023), who derived this value using their best-fit quasar redshift distribution $n(z)$ and bias evolution $b(z)$ mentioned in step (c) above. We generate this clustering dipole as an independent map containing only $\ell = 1$ modes with the specified amplitude, subject to cosmic variance, as characterized in Section 4.2.

Note that the value of D_{clus} used here is higher than that determined by N. J. Secrest et al. (2021; $D_{\text{clus}} = 0.24 \times 10^{-3}$). Their value is obtained with the assumption of a constant bias ($b(z) = 1$).

The final synthetic quasar number density map is a superposition of these carefully separated independent components. This construction ensures a clean separation of components with no double counting: the kinematic and shot-noise dipole originates exclusively from steps (a) + (b), the clustering dipole is added independently, and the cosmological clustering and associated shot noise ($\ell \geq 2$) comes solely from FLASK, with its monopole and dipole modes explicitly set to zero:

$$\begin{aligned}
 N(\hat{n}) = & \underbrace{a_{00} Y_{00}(\hat{n})}_{\text{Monopole } (N_0)} + \underbrace{\sum_{m=-1}^1 a_{1m}^{\text{kin}} Y_{1m}(\hat{n})}_{\text{CMB Kinematic Dipole } (D_{\text{kin}})} \\
 & + \underbrace{\sum_{m=-1}^1 a_{1m}^{\text{SN}} Y_{1m}(\hat{n})}_{\text{SN Dipole } (D_{\text{SN}})} + \underbrace{\sum_{m=-1}^1 a_{1m}^{\text{clus}} Y_{1m}(\hat{n})}_{\text{Clustering Dipole } (D_{\text{clus}})} \\
 & + \underbrace{\sum_{\ell=2}^{\ell_{\text{max}}} \sum_{m=-\ell}^{\ell} a_{\ell m}^{\text{SN+clus}} Y_{\ell m}(\hat{n})}_{\ell \geq 2 \text{ Clustering+SN}}. \quad (20)
 \end{aligned}$$

The different superscripts on $a_{\ell m}$ s indicate the different contributions. To summarize: step (a) gives the first two terms; step (b) gives the third term; step (c) gives the fifth (last) term; and finally step (d) gives the fourth term. This equation captures the deterministic, signal-dominated kinematic dipole alongside the stochastic, lognormal fluctuations from large-scale structure and shot-noise systematics. The fluctuations of the simulated $N(\hat{n})$ then provide a way to determine the statistical significance of the observed CatWISE2020 dipole amplitude. In Figure 6 of Appendix A, we present one realization of each component, along with the composite $N(\hat{n})$.

In the next section, the statistical significance of the observed CatWISE2020 dipole amplitude will be tested against different sets of simulations, where we keep only specific contributions to Equation (20). These are:

- S0: contains the monopole, CMB kinematic dipole, and a shot-noise component including all $1 \leq \ell \leq \ell_{\text{max}}$ modes. These mocks reproduce the assumptions of N. J. Secrest et al. (2021).
- S1: includes all contributions to Equation (20), except the clustering dipole (the fourth term on the right-hand side). In particular, it adds higher-order ($\ell \geq 2$)

shot-noise and clustering moments generated with FLASK, but omits the clustering dipole.

- S2: expands on S1 by adding a clustering dipole with amplitude taken from P. Tiwari et al. (2023). Its direction is chosen randomly on the sky, reflecting the assumption that large-scale structure is uncorrelated with our motion.
- S3: identical to S2, except that the clustering and CMB kinematic dipoles are aligned, representing the possibility of a common physical origin related to our local motion. The amplitude of each clustering dipole realization is varied in accordance with the cosmic variance. The alignment is also motivated by findings such as those of O. T. Oayda et al. (2024), who observed that clustering dipole directions in data from the NVSS and the Rapid ASKAP Continuum Survey appear to be aligned with the CMB dipole direction.
- S4: similar to S2, but with the clustering power ($C_{1,\text{clus}}$) increased to its upper 2σ limit ($D_{\text{clus}}^{2\sigma} = 1.48 \times 10^{-3}$); its direction remains random. This scenario is hypothetical and tests the implications of a higher value of D_{clus} with random orientation.
- S5: in this scenario, the clustering dipole direction is aligned with the CMB kinematic dipole, while the amplitude is constant and the same as in S4. This scenario is again hypothetical and tests the implications of a higher value of D_{clus} with orientation aligned with the CMB.

These simulation sets are designed to isolate the impacts of different physical inputs on the inferred statistical significance of the observed dipole amplitude. Specifically, S0 is used to verify consistency with N. J. Secrest et al. (2021) under the same assumptions. S1–S3 test the influence of the expected clustering dipole, which is respectively absent, randomly oriented, or aligned with the kinematic dipole in these three cases. Finally, S4 and S5 explore the sensitivity of the measured dipole significance to a clustering dipole with larger amplitude that is randomly oriented or aligned with the kinematic dipole.

6. Results—Statistical Significance

The observed CatWISE2020 dipole amplitude, which we denote by D^{obs} , is obtained to be 15.54×10^{-3} , in agreement with the value reported by N. J. Secrest et al. (2021). This value was obtained using the standard monopole plus dipole ($M + D$) model defined in Section 3.1. To quantify the statistical significance of this observed value, we performed a hypothesis test against each suite of mock simulations prepared as described in Section 5. We employed the CatWISE2020 mask on the simulations to ensure consistent sky coverage with the observed dataset. Then we calculated the statistical significance for each set by counting how many simulations have dipole amplitude values that match or exceed D^{obs} . The results are described below.

S0 Simulations. This set has the same inputs as the simulations used by N. J. Secrest et al. (2021) and serves to reproduce their results, following their analysis pipeline. The top left panel of Figure 5 shows the PDF of the dipole amplitudes measured from the simulations. The CatWISE2020 value is marked in red. We find that 9 out of 10^7 simulations have dipole values that match or exceed D^{obs} , giving a p -value $\approx 9 \times 10^{-7}$, or a significance of 4.82σ . This confirms their

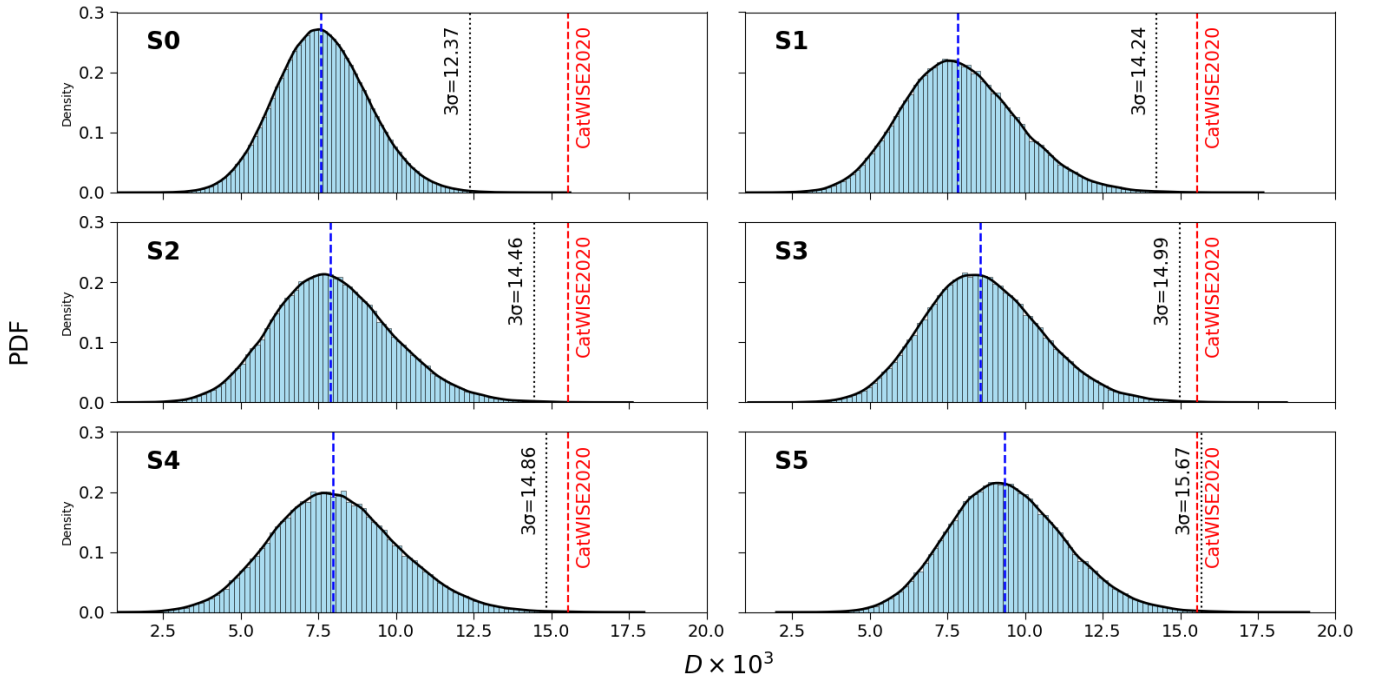


Figure 5. PDFs of the dipole amplitude ($\times 1000$) from six simulation scenarios. Across all subplots, the CatWISE2020 catalog dipole amplitude ($D^{\text{obs}} \times 10^3 = 15.54$) is indicated by the vertical red dashed line. The blue dashed and black dotted vertical lines represent the median and 3σ confidence level, respectively, specific to each simulation set.

reported high level of significance and verifies that our analysis is consistent with their study, before we introduce other physical effects into the simulations.

From these simulations, we obtain the ratio of the median values of C_1 and C_3 to be ~ 18.01 . The C_3 power is contributed by shot noise alone, and it is roughly a factor of 18 smaller than C_1 .

S1 Simulations. The top right panel of Figure 5 shows the PDF of the dipole amplitudes measured from the simulations. We observe that the PDF is broadened, accompanied by a slight shift of the median toward higher values relative to the S0 case. This behavior can be attributed to the inclusion of stochastic clustering modes at multipoles $\ell \geq 2$, with the dominant contribution arising from $\ell = 3$, as discussed in Section 4.

Of 10^5 simulations, 13 exhibit dipole amplitudes equal to or exceeding the observed value D^{obs} . This corresponds to a p -value of 1.30×10^{-4} , equivalent to a statistical significance of 3.63σ .

The ratio of the median values of C_1 and C_3 is obtained to be approximately 8.71, which is roughly double the value in the S0 case. The increase of the C_3 power results from contributions from *both* shot noise and clustering. Nevertheless, C_3 remains approximately an order of magnitude smaller than C_1 .

S2 Simulations. The left panel of the middle row of Figure 5 shows the PDF of the dipole amplitudes measured from the simulations. We observe a mild shift of the median relative to the S0 case, similar to S1. There is also a small broadening of the PDF due to the introduction of the stochastic clustering dipole relative to the PDFs of the S0 and S1 cases.

We obtain 27 out of 10^5 simulations matching or exceeding D^{obs} . This yields a p -value of 2.70×10^{-4} , corresponding to a significance of 3.44σ .

S3 Simulations. The right panel of the middle row of Figure 5 shows the PDF of the dipole amplitudes measured from these simulations. We observe a shift of the median, relative to the previous cases, which can be attributed to the alignment between the clustering and kinetic dipole directions.

The number of simulations matching or exceeding D^{obs} increases to 58 out of 10^5 , corresponding to a p -value of 5.80×10^{-4} and reducing the significance to 3.27σ .

S4 Simulations. The bottom left panel of Figure 5 shows the PDF of the dipole amplitudes measured from these simulations. There is a very mild broadening (not easily discernible) of the PDF relative to S2. This can be attributed to the enhanced clustering dipole amplitude compared to the value used in S2.

In this case, 53 out of 10^5 simulations have dipole amplitudes that match or exceed D^{obs} . This yields $p = 5.30 \times 10^{-4}$, corresponding to a significance of 3.27σ .

S5 Simulations. The bottom right panel of Figure 5 shows the PDF of the dipole amplitudes measured from the S5 simulations. We observe a shift of the median relative to all previous cases, which can be attributed to the alignment between the clustering and kinematic dipole directions, as well as the enhanced amplitude of the clustering dipole. Of 10^5 realizations, 175 yield dipole amplitudes equal to or exceeding the observed value D^{obs} ($p = 1.75 \times 10^{-3}$), corresponding to a significance of 2.93σ . This indicates that the observed dipole remains a statistically significant anomaly even when all model parameters are pushed to their most extreme, yet still plausible, values.

In the five scenarios S1–S5 analyzed here, the decrease in statistical significance relative to S0 arises from both the broadening of the PDFs and the shifting of their medians toward higher values. This behavior is naturally expected when the clustering dipole and higher-order multipoles are included, since their stochastic nature contributes positively to

Table 1
Results of Statistical Significance Tests from 10^5 Simulations for S1–S3 and 10^7 Realizations for the S0 Case

Simulation	Physical Inputs	Clustering Dipole Direction	n	p -value	Significance (σ)
S0	Kinematic Dipole + Shot Noise (N. J. Secrest et al. 2021)	...	9	9.00×10^{-7}	4.82σ
S1	Baseline + Higher-order Modes, No Clustering Dipole	...	13	1.30×10^{-4}	3.63σ
S2	Baseline + Higher-order Modes + Clustering Dipole	Random	27	2.70×10^{-4}	3.44σ
S3	Baseline + Higher-order Modes + Clustering Dipole	Aligned with CMB Dipole	58	5.80×10^{-4}	3.27σ

Note. We define “Baseline” here as the sum of the monopole (N_0), the kinematic dipole (D_{kin}), and the shot-noise dipole (D_{SN}) and higher-order contributions: the first three terms on the right-hand side of Equation (20). n is the number of realizations that presented a dipole larger than the CatWISE2020 value. The S0 case replicates the original N. J. Secrest et al. (2021) analysis, confirming their high significance.

the total variance of the measured dipole. The extents of the broadening and median shifting depend on the relative amplitudes of the dipole and higher multipoles. If the clustering dipole and higher multipoles should be weaker than the values adopted in our analysis (P. Tiwari et al. 2023), the inferred statistical significance of D^{obs} would be correspondingly higher.

The statistical significance values for the main physically motivated scenarios (S1, S2, and S3) are summarized in Table 1 and constitute the main results of this paper. The value for the case of S0 is included for comparison. In conclusion, while the inclusion of a clustering dipole, particularly one that is aligned with the CMB dipole and affected by cosmic variance, increases the probability of observing a large dipole amplitude, the CatWISE2020 result remains a statistically significant outlier at approximately the 3.3σ level or higher across all tested scenarios.

To assess the robustness of the measurement of the observed dipole value, we further employed a suite of alternative models extending beyond the one defined in Section 3.1, sequentially including higher-order multipoles in the least-squares fitting. The full analysis, including discussions of their numerical stability, and the resulting dipole values and their significance, is presented in Appendix B. In summary, while alternative models yield broadly consistent dipole directions, the inclusion of higher multipoles, particularly the octopole, increases the fitted dipole amplitude and markedly increases the variance of the measurement, as indicated by the condition number of the normal-equation matrix.

7. Conclusion and Discussion

This work presents a reassessment of the cosmic dipole anomaly using the CatWISE2020 quasar catalog. We have demonstrated the importance of accounting for several sources of uncertainties, including a rigorous treatment of the mode coupling introduced by the sky mask, which systematically biases the dipole estimate. Further, we have developed mock simulations that account for kinematic, clustering, and shot-noise effects.

Our primary finding is a revised estimate of the statistical significance of the CatWISE2020 dipole, summarized in Table 1. The reduction from 4.9σ to 3.63σ in the case of no clustering dipole and to 3.27σ in the case of CMB-aligned clustering are notable drops. However, it is clear that the observed anomalously large dipole amplitude cannot be explained purely by the physical effects (clustering) and numerical artifacts (mode coupling) studied in this work, and it remains a challenge to the standard cosmological model. This conclusion is bolstered by similarly anomalous findings in

other independent datasets (J. Colin et al. 2017; N. J. Secrest et al. 2022).

Our analysis improves upon previous studies, most notably, that of N. J. Secrest et al. (2021), by incorporating a more complete treatment of systematic effects. While the earlier works correctly accounted for the shot-noise dipole, our study includes two missing components: the intrinsic clustering dipole and the influence of higher-order clustering modes. These inclusions are important. The sky mask induces mode coupling, but without the significant power from intrinsic clustering in the simulation, the previous significance estimates were large. Furthermore, we have systematically studied the bias introduced by the mask itself on the estimated dipole value and have employed the FLASK simulation package to generate mock catalogs that replicate the survey’s statistical properties, leading to the improvement of the significance assessment. We also emphasize that should the amplitude of the clustering effects be lower than the values measured by P. Tiwari et al. (2023; which are used here), the inferred statistical significance of D^{obs} would increase, approaching the level reported by N. J. Secrest et al. (2021). Note that for a direct comparison with the results of N. J. Secrest et al. (2022), we will need to repeat our analysis using their updated and expanded quasar catalog. We plan to undertake this as part of a follow-up study.

Finally, we systematically quantify how the statistical inference depends on the model itself, showing that including higher multipoles in the fitting increases the condition number and can obscure the true level of tension. Crucially, we find that the choice of which low-order multipoles are simultaneously fitted constitutes a significant systematic. High condition numbers that indicate mask-induced mode coupling can artificially suppress the measured significance by inflating the variance of the dipole PDF distribution, as described in Appendix C.

S. von Hausegger et al. (2026) recently carried out a Bayesian-inference-based parametric fit and determined the dipole, quadrupole, and octopole amplitudes using the extended catalog consisting of 1.6 million sources from CatWISE2020. The inferred dipole amplitude value varies slightly with the input model (only-dipole, dipole+quadrupole, or dipole+quadrupole+octopole). They asked whether the anomalously large value of D^{obs} can be due to ignoring the octopole term, and find that including an octopole term slightly increases the observed dipole rather than decreasing it. They also state that the measured octopole value is consistent with noise, based on a 3σ detection criterion. Our work differs from theirs, in that our focus is on the determination of the statistical significance of D^{obs} (measured using Equation (7)) based on a

frequentist hypothesis test when the clustering dipole and higher-order clustering modes are included in mock simulations. Our results are broadly consistent with their analysis: the primary effect of including the octopole is to broaden the posteriors on the dipole, reducing its overall significance compared to the case in which only a dipole is fit to the data.

A shortcoming of the present analysis is that contamination from stellar and other unwanted sources in the quasar catalog is not accounted for. Such contamination can potentially alter the assessed significance level. The color cut of $W1 - W2 \geq 0.8$ produces a highly reliable ($\sim 95\%$) quasar catalog (D. Stern et al. 2012), but the impact of the remaining contaminants on the amplitude of the CatWISE2020 dipole has yet to be explored. In future work, it will be important to incorporate such contamination, either by injecting into the simulations or by utilizing more refined observational datasets with robust unwanted source removal. Furthermore, a future study involving a morphological-level comparison between the CatWISE2020 survey and our mock simulation can provide additional information. The large-scale behavior of the one-point function of the data has been studied in A. Antony et al. (2026), but more sophisticated techniques can be utilized to leverage more information. Such validation is essential for developing simulations that accurately replicate the higher-order clustering statistics of the observed quasar population. Moreover, the least-squares estimator used here is susceptible to bias from mask-induced mode coupling and confusion with the intrinsic clustering dipole. Hence, there is a need for improved methodologies to disentangle these different effects (T. Nadolny et al. 2021).

Finally, estimations of the matter dipole have been extended to sources such as galaxies and quasars identified from optical observations. A tomographic dipole estimator applied to Sloan Digital Sky Survey spectroscopic data was found to broadly agree with the CMB dipole (P. da Silveira Ferreira & V. Marra 2024; P. Tiwari et al. 2024). Similarly, a Bayesian analysis of the Quiaia quasar sample (K. Storey-Fisher et al. 2024), after applying conservative sky cuts, found the dipole to be consistent with the CMB (V. Mittal et al. 2024), though this consistency may be driven by the larger noise in the reduced sample. A. K. Singal (2024, 2025) also analyzed the Quiaia data using a different selection function and concluded that the peculiar velocity of the solar system is 4–5 times higher than the value inferred from the CMB dipole, with the direction pointing toward the Galactic center within 1σ . These results collectively highlight that conclusions drawn from matter dipole analyses remain sensitive to the choice of dataset and the methodology adopted, and the associated issues continue to be actively examined by the community (V. Mittal et al. 2025, O. T. Oayda

et al. 2025). It is expected that forthcoming data releases with high-source-count densities, such as from the Rubin telescope (P. A. Abell et al. 2009; R. Braun et al. 2019), will enable a decisive measurement of the dipole, its redshift dependence, and the existence of the anomaly (DESI Collaboration et al. 2016; L. Amendola et al. 2018; B. P. Crill et al. 2020). Understanding the possible origin of such a dipole, and performing further tests of the statistical isotropy of large-scale structure, remain outstanding challenges in cosmology.

Acknowledgments

M.B. gratefully acknowledges Iqra Hamid Bhat for insightful discussions; Nidharssan S. and Vikranth P. for valuable suggestions; and Sriram K. and Arav B. J. for discussions on source systematics and mode mixing. The authors thank N. Secrest, S. von Hausegger, and S. Sarkar for critical comments on this work. Computations were performed on the Nova cluster at the Indian Institute of Astrophysics, Bengaluru. This work made use of HEALPIX⁵ (K. M. Górski et al. 2005), Healpy (A. Zonca et al. 2019), Matplotlib (J. D. Hunter 2007), NaMaster (D. Alonso et al. 2019), and CLASS (J. Lesgourgues 2011). The authors thank N. Secrest et al. for making their data publicly available. P.C. and S.A. gratefully acknowledge the hospitality of the Korea Institute for Advanced Study, Seoul, where part of this work was carried out.

Appendix A Visualization of Mock Simulations

For visual clarity, Figure 6 illustrates the individual components that contribute to each simulation of $N(\hat{n})$ (Equation (20)). The left column shows the monopole component, which is uniform on the sky. The middle column displays the three dipole components: the CMB kinetic dipole (top), a shot-noise dipole realization (middle), and a clustering dipole realization aligned with the CMB dipole direction (bottom). The right column shows the higher-multipole contributions from shot noise and clustering. At the bottom, the composite number-count density map, obtained by summing all components, is shown.

The complete maps $N(\hat{n})$ include fluctuations arising from the shot-noise and clustering dipoles, as well as higher-order multipoles. For full-sky data, only the kinetic, shot-noise, and clustering dipole contributions would be relevant; however, masking induces mode coupling and makes the amplitudes of all fluctuations relevant for assessing the significance of the observed dipole in CatWISE2020.

⁵ <http://healpix.sourceforge.net/>

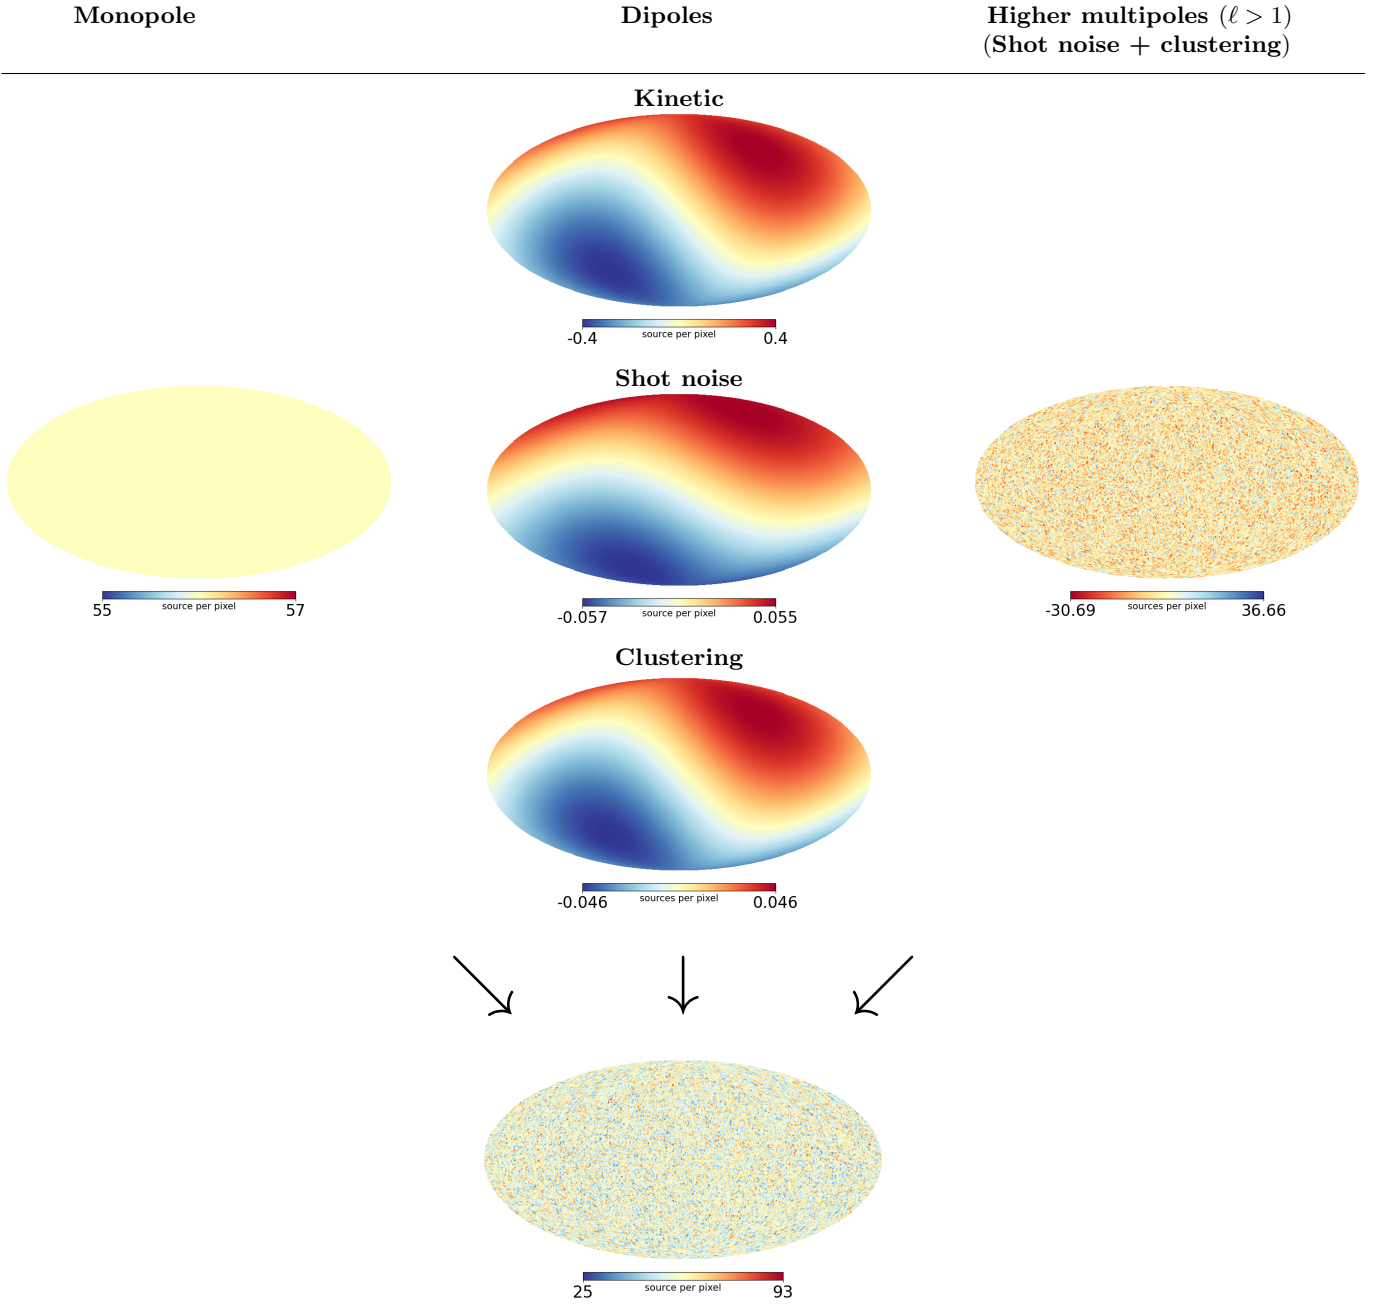


Figure 6. Left column: monopole component of the number-count density map. Middle column: dipole contributions—kinematic (top), shot noise (middle), and clustering aligned with the CMB dipole direction (bottom). Right column: combined higher-multipole contributions from shot noise and clustering. Bottom: final simulated number-count density map, $N(\hat{n})$, as defined in Equation (20), obtained by summing all the above components.

Appendix B

Estimation of Quasar Dipole when Including Higher Multipoles

To estimate the direction and amplitude of the apparent dipole from the CatWISE2020 quasar sample, we performed a series of sequential, low-multipole least-squares fits directly on the number density maps.

The spherical harmonic templates $Y_{\ell m}$, masked to match the survey geometry, were used to form a linear model for the observed map:

$$\mathbf{d}_{\text{model}} = \sum_{\ell \in S} a_{\ell m} \mathbf{Y}_{\ell m} = \mathbf{X} \mathbf{a},$$

where S is the set of multipoles included in the fit, \mathbf{a} is the vector of corresponding spherical harmonic coefficients $a_{\ell m}$, and \mathbf{X} is the design matrix whose columns are the mode vectors.

The coefficients \mathbf{a} were obtained by minimizing the residual between the data vector \mathbf{d} (the pixelated number density map) and the model:

$$\min_{\mathbf{a}} \|\mathbf{d} - \mathbf{X} \mathbf{a}\|^2.$$

This leads to the normal equations

$$\mathbf{X}^T \mathbf{X} \mathbf{a} = \mathbf{X}^T \mathbf{d},$$

Table 2
Dipole Estimates Obtained from Sequential Low-multipole Fits of the CatWISE2020 Quasar Map

Model	D/M	l (deg)	b (deg)	Offset from CMB (deg)	Condition Number (κ)
M + D	0.01554	238.22	28.80	27.78	5.05×10^0
M + D + Q	0.01545	237.86	29.18	27.68	3.48×10^1
M + D + Q + O	0.01971	237.22	24.04	32.15	5.66×10^1
M + D + Q + O + $\ell = 4$	0.01970	236.57	24.38	32.20	2.28×10^2

Note. The dipole amplitude is quoted relative to the monopole, D/M . The offset denotes the great-circle distance from the CMB dipole direction (l, b) = (264.0, 48.3). The condition number refers to the normal-equation matrix $X^T X$.

with the solution

$$\mathbf{a} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{d}.$$

To monitor the stability of the inferred dipole against potential contamination from higher multipoles, particularly important given the partial sky coverage and the resulting mode coupling, we performed a sequence of fits with an increasing number of low multipoles. Specifically, we fit models containing: the monopole and dipole (M+D); the monopole, dipole, and quadrupole (M+D+Q); the monopole, dipole, quadrupole, and octopole (M+D+Q+O); and the monopole through $\ell = 4$ (M+D+Q+O+ $\ell = 4$). For each fit, we computed the condition number⁶ (D. A. Belsley et al. 1980) of the Gram matrix that is inverted to infer \mathbf{a} : $\kappa(\mathbf{X}^T \mathbf{X})$. The condition number increases as more modes are included in the fitting procedure, and we discuss the impact of this phenomenon on our results in the following section. We do not impose any specific threshold on the condition number as a criterion for rejecting a model. All values of κ reported here are well below the regime of numerical ill-conditioning. Here, the condition number acts as a diagnostic: a larger κ indicates stronger statistical correlations between the fitted multipole templates induced by the mask, which in turn inflates the variance of the recovered dipole amplitude and reduces its measured significance.

The results of this sequential fitting procedure are summarized in Table 2.

The dipole estimate remains stable when only the monopole and dipole are included. However, the inclusion of the octopole ($\ell = 3$) shifts the amplitude by approximately 27% and the direction by roughly 5° , coinciding with an order of magnitude increase in the condition number. This indicates that the dipole estimate becomes sensitive to mask-induced mode coupling when higher multipoles are simultaneously fit, highlighting the importance of the multipole selection when fitting to the data. Consequently, the (M+D) fit, which has a low condition number, is adopted as our fiducial model for the CatWISE2020 dipole.

Appendix C Significance Estimates from Different Models

The statistical significance of the dipole amplitude anomaly reported for the CatWISE2020 quasar catalog is highly

sensitive to the specific mathematical model used to extract the dipole from the sky map. While our primary analysis (Section 6) employed a simple monopole and dipole (M+D) fit, we systematically tested the impact of including higher multipoles in the fitting procedure.

This sensitivity is particularly pronounced for the M+D+Q+O model. This choice, characterized by a high condition number (Appendix B), arises from strong mask-induced coupling between the dipole and higher-order multipoles, particularly $\ell = 5, 7$, as shown in case of $\ell = 5$ in Figure 7, and beyond. Consequently, stochastic variations in these higher- ℓ modes within the simulated maps produce amplified variance in the estimated dipole amplitude. This inflated variance in the null distribution dilutes the statistical extremity of the observed CatWISE2020 dipole, thereby suppressing its measured significance.

We recomputed the p -values for the S1, S2, and S3 simulation suites using two alternative models: M+D+Q and M+D+Q+O. The results are presented in Table 3.

The results reveal a critical systematic: the choice of model included in the fit dramatically alters the inferred significance. The simple M+D and M+D+Q models yield consistent results, with significances between $\sim 3.2\sigma$ and $\sim 3.6\sigma$ across the physical scenarios. In contrast, the M+D+Q+O model reduces the significance to approximately 2.3σ – 2.4σ , a difference of approximately 1σ .

This reduction stems from a fundamental bias–variance tradeoff inherent to fitting multipoles to masked sky data. Including higher multipoles like the octopole in the fit reduces the bias from that specific multipole, but simultaneously increases the variance of the recovered dipole due to enhanced mask-induced effects. In particular, when the M+D+Q+O model is applied to simulations containing realistic higher-order clustering modes (S1–S3), it yields a broader distribution of recovered dipole amplitudes. Consequently, the observed CatWISE2020 amplitude, while large, becomes a less extreme outlier, given the inflated variance, leading to a substantially higher p -value. Future work should therefore focus on systematic stress tests of dipole estimators under controlled conditions, including varying mask geometries, noise levels, and injected higher-order multipoles, in order to identify choices that provide the most reliable balance between bias and variance on a partial sky.

⁶ The condition number C of a nonsingular matrix A is defined as $C = s_1/s_2$, where s_1 and s_2 are the largest and smallest singular values; large C indicates numerical ill-conditioning.




Figure 7. Recovered dipole amplitudes from simulations compared to the expected CMB dipole amplitude. In both panels, the same sky mask is applied. Left panel: the map includes only a test-input $C_3 \approx 0.1 C_{1,kin}$ (octopole) component added to the kinematic dipole. Right panel: the map includes only a test-input $C_5 \approx 0.1 C_{1,kin}$ component added to the kinematic dipole. Each panel shows box plots comparing the M + D model (fitting only the monopole and dipole) with the M + D + Q + O model (fitting up to the octopole). This comparison demonstrates how including higher multipoles in the fit affects the recovery of the dipole amplitude.

Table 3
Impact of Model Choice on the Significance of the CatWISE2020 Dipole Excess

Simulation Suite	Model	n	p -value	Significance (σ)
S1 (No Clust. Dipole)	M+D ($\kappa = 5.01$)	13	1.30×10^{-4}	3.63
	M+D+Q ($\kappa = 34.8$)	20	2.00×10^{-4}	3.53
	M+D+Q+O ($\kappa = 56.6$)	797	7.97×10^{-3}	2.41
S2 (Random Clust. Dipole)	M+D ($\kappa = 5.01$)	27	2.70×10^{-4}	3.44
	M+D+Q ($\kappa = 34.8$)	32	3.20×10^{-4}	3.41
	M+D+Q+O ($\kappa = 56.6$)	817	8.17×10^{-3}	2.40
S3 (Aligned Clust. Dipole)	M+D ($\kappa = 5.01$)	58	5.80×10^{-4}	3.27
	M+D+Q ($\kappa = 34.8$)	71	7.10×10^{-4}	3.19
	M+D+Q+O ($\kappa = 56.6$)	1125	1.13×10^{-2}	2.28

Note. The reported p -values and significance (σ) are shown for three simulation suites (S1–S3), using the M+D (primary), M+D+Q, and M+D+Q+O models. n denotes the number of simulations (out of 100,000) with a dipole amplitude \geq the observed value.

ORCID iDs

Masroor Bashir <https://orcid.org/0009-0002-2935-0496>

Pravabati Chingangbam <https://orcid.org/0000-0002-7385-8273>

Stephen Appleby <https://orcid.org/0000-0001-8227-9516>

References

- Abell, P. A., Allison, J., Anderson, S. F., et al. 2009, LSST Science Book, Version 2.0 (Fermi National Accelerator Laboratory)
- Abghari, A., Bunn, E. F., Hergt, L. T., et al. 2024, *JCAP*, 2024, 067
- Alonso, D., Sanchez, J., & Slosar, A. 2019, *MNRAS*, 484, 4127
- Amendola, L., Appleby, S., Avgoustidis, A., et al. 2018, *LRR*, 21, 2
- Antony, A., Appleby, S., Matthewson, W. L., & Shafieloo, A. 2026, *JCAP*, 2026, 001
- Assef, R. J., Stern, D., Kochanek, C. S., et al. 2013, *ApJ*, 772, 26
- Belsley, D. A., Kuh, E., & Welsch, R. E. 1980, Wiley Series in Probability and Statistics (Wiley)
- Bengaly, C. A. P., Maartens, R., & Santos, M. G. 2018, *JCAP*, 2018, 031
- Blake, C., & Wall, J. 2002, *Natur*, 416, 150
- Böhme, L., Schwarz, D. J., Tiwari, P., et al. 2025, *PhRvL*, 135, 201001
- Braun, R., Bonaldi, A., Bourke, T., Keane, E., & Wagg, J. 2019, arXiv:1912.12699
- Colin, J., Mohayaee, R., Rameez, M., & Sarkar, S. 2017, *MNRAS*, 471, 1045
- Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, *AJ*, 115, 1693
- Crill, B. P., Werner, M., Akesson, R., et al. 2020, *SPIE*, 11443, 1144301
- da Silveira Ferreira, P., & Marra, V. 2024, *JCAP*, 2024, 077
- Dam, L., Lewis, G. F., & Brewer, B. J. 2023, *MNRAS*, 525, 231
- DES Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036
- Eisenhardt, P. R. M., Marocco, F., Fowler, J. W., et al. 2020, *ApJS*, 247, 69
- Ellis, G. F. R., & Baldwin, J. E. 1984, *MNRAS*, 206, 377
- Gibelyou, C., & Huterer, D. 2012, *MNRAS*, 427, 1994
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
- Hivon, E., Górski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, 567, 2
- Hunter, J. D. 2007, *CSE*, 9, 90
- Kumar Aluri, P., Cea, P., Chingangbam, P., et al. 2023, *CQGra*, 40, 094001
- Land-Strykowski, M., Lewis, G. F., & Murphy, T. 2025, *MNRAS*, 543, 3229
- Lesgourgues, J. 2011, arXiv:1104.2932
- Marocco, F., Eisenhardt, P. R. M., Fowler, J. W., et al. 2021, *ApJS*, 253, 8
- Mittal, V., Oayda, O. T., & Lewis, G. F. 2024, *MNRAS*, 527, 8497
- Mittal, V., Oayda, O. T., & Lewis, G. F. 2025, *OJAp*, 8, 143
- Nadolny, T., Durrer, R., Kunz, M., & Padmanabhan, H. 2021, *JCAP*, 2021, 009
- Oayda, O. T., Mittal, V., & Lewis, G. F. 2025, *MNRAS*, 537, 1
- Oayda, O. T., Mittal, V., Lewis, G. F., & Murphy, T. 2024, *MNRAS*, 531, 4545
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, *A&A*, 641, A1
- Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020b, *A&A*, 641, A6
- Rubart, M., & Schwarz, D. J. 2013, *A&A*, 555, A117
- Secrest, N., von Hausegger, S., Rameez, M., Mohayaee, R., & Sarkar, S. 2025, *RvMP*, 97, 041001
- Secrest, N. J., Hausegger, S., Rameez, M., et al. 2021, *ApJL*, 908, L51
- Secrest, N. J., von Hausegger, S., Rameez, M., Mohayaee, R., & Sarkar, S. 2022, *ApJL*, 937, L31
- Siewert, T. M., Schmidt-Rubart, M., & Schwarz, D. J. 2021, *A&A*, 653, A9

- Singal, A. K. 2011, [ApJL](#), 742, L23
- Singal, A. K. 2019, [PhRvD](#), 100, 063501
- Singal, A. K. 2023, [MNRAS](#), 524, 3636
- Singal, A. K. 2024, [MNRAS](#), 532, L1
- Singal, A. K. 2025, [NatSR](#), 15, 31805
- Stern, D., Assef, R. J., Benford, D. J., et al. 2012, [ApJ](#), 753, 30
- Storey-Fisher, K., Hogg, D. W., Rix, H.-W., et al. 2024, [ApJ](#), 964, 69
- Tiwari, P., Schwarz, D. J., Zhao, G.-B., et al. 2024, [ApJ](#), 975, 279
- Tiwari, P., Zhao, G.-B., & Nusser, A. 2023, [ApJ](#), 943, 116
- von Hausegger, S., Secret, N., Desmond, H., et al. 2026, [MNRAS](#), 546, [stag201](#)
- Wagenveld, J. D., Klöckner, H.-R., & Schwarz, D. J. 2023, [A&A](#), 675, A72
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, [AJ](#), 140, 1868
- Xavier, H. S., Abdalla, F. B., & Joachimi, B. 2016, [MNRAS](#), 459, 3693
- Zonca, A., Singer, L., Lenz, D., et al. 2019, [JOSS](#), 4, 1298