

PAPER • OPEN ACCESS

Constraining primordial non-Gaussianity from DESI DR1 quasars and *Planck* PR4 CMB lensing

To cite this article: S. Chiarenza *et al* JCAP05(2026)098

View the [article online](#) for updates and enhancements.

You may also like

- [Constraining primordial non-Gaussianity from DESI quasar targets and Planck CMB lensing](#)
Alex Krolewski, Will J. Percival, Simone Ferraro et al.
- [Constraining primordial non-Gaussianity with DESI 2024 LRG and QSO samples](#)
E. Chaussidon, C. Yèche, A. de Mattia et al.
- [Optimal constraints on Primordial non-Gaussianity with the eBOSS DR16 quasars in Fourier space](#)
Marina S. Cagliari, Emanuele Castorina, Marco Bonici et al.

Constraining primordial non-Gaussianity from DESI DR1 quasars and *Planck* PR4 CMB lensing

S. Chiarenza et al.

Full author list at the end of the paper

E-mail: schiaren@uwaterloo.ca

ABSTRACT: We present the first measurement of local-type primordial non-Gaussianity from the cross-correlation between 1.2 million spectroscopically confirmed quasars from the first data release (DR1) of the Dark Energy Spectroscopic Instrument (DESI) and the *Planck* PR4 CMB lensing reconstructions. The analysis is performed in three tomographic redshift bins covering $0.8 < z < 3.5$, covering a sky fraction of $\sim 20\%$. We adopt a catalog-based pseudo- C_ℓ estimator and apply linear imaging weights validated on noiseless mocks. Compared to previous analyses using photometric quasar samples, our results benefit from the high purity of the DESI spectroscopic sample, the reduced noise of PR4 lensing, and the absence of excess large-scale power in the spectroscopic quasar auto-correlation. Fitting simultaneously for the non-Gaussianity parameter f_{NL} and the linear bias amplitude in each redshift bin, we obtain $f_{\text{NL}} = 2^{+28}_{-34}$ for a response parameter $p = 1.6$, and $f_{\text{NL}} = 6^{+20}_{-24}$ for $p = 1.0$. These results improve the constraints on f_{NL} by $\sim 35\%$ compared to the previous analysis based on the Legacy Imaging Survey DR9. Additionally, we derive an optimal weighting scheme to maximize the constraining power. In this case, and assuming $p = 1.6$, we obtain $f_{\text{NL}} = 19^{+25}_{-31}$. Our results demonstrate the statistical power of DESI quasars for probing inflationary physics, and highlight the promise of future DESI data releases.

KEYWORDS: cosmological parameters from LSS, inflation, physics of the early universe, cosmological parameters from CMBR

ARXIV EPRINT: [2512.17865](https://arxiv.org/abs/2512.17865)



Contents

1	Introduction	1
2	Theory	3
3	Data	7
3.1	DESI DR1 spectroscopically confirmed quasars	8
3.2	<i>Planck</i> PR4 CMB lensing map	9
4	Measuring angular power spectra	9
4.1	Monte Carlo normalization correction for the <i>Planck</i> lensing maps	11
5	Imaging systematic weights validation	12
5.1	The DESI systematic weights	12
5.2	Correcting for excess mode removal	12
6	Optimal redshift weighting for f_{NL}	17
7	Results and discussion	19
7.1	The optimal weights results	23
7.2	Fitting a single bias amplitude for all three bins	25
7.3	Impact of the bias relation	25
7.4	The impact of the primordial power spectrum	26
8	Conclusions	28
A	Derivation of the optimal f_{NL} weights for $C_\ell^{\kappa g}$	30
B	Restricting to quasars with $z < 3.1$	31
	Author List	39

1 Introduction

Inflation is the leading framework for explaining the emergence of cosmic structure. A key prediction of this theory is that the initial curvature perturbations are almost perfectly Gaussian [1–3]. This prediction is supported by precise measurements of the Cosmic Microwave Background (CMB) temperature and polarization [4]. Still, many inflation models allow small departures from Gaussianity, constraining these deviations, known as primordial non-Gaussianity (PNG), is a powerful direct test of early-Universe physics. For example, in single-field slow-roll (SFSR) inflation the expected level of local PNG is small, $f_{\text{NL}}^{\text{loc}} \sim \mathcal{O}(10^{-2})$ [3, 5]. Multi-field scenarios, on the other hand, can naturally reach $f_{\text{NL}}^{\text{loc}} \sim \mathcal{O}(1)$ [6]. Achieving order-unity precision on $f_{\text{NL}}^{\text{loc}}$ is thus an important goal in ruling out significant parts of the parameter space of inflation.

The best current limit comes from the *Planck* bispectrum, $f_{\text{NL}} = -0.9 \pm 5.1$ [7]. This is already close to the cosmic-variance limit for the CMB, and future CMB experiments will be able to improve it by only a factor of two [8]. Stronger constraints must therefore come from large-scale structure (LSS).

Local PNG introduces a scale-dependent bias in the clustering of galaxies [9], which manifests as an enhancement of the quasars (QSOs) power spectrum, $P(k)$, on the very large scales. The effect scales as $\propto k^{-2}$ and grows with redshift, making quasars particularly powerful tracers: their high-redshift distribution provides access to the large cosmic volumes and low- k modes where the signal is most pronounced. In fact, the best constraint to date from LSS comes from the 3D power spectrum analysis done on the combination of DESI DR1 QSO and LRG samples, yielding $f_{\text{NL}} = -3.6^{+9.0}_{-9.1}$ [10]. However, measuring galaxy clustering on very large scales is inherently challenging. The signal fluctuations are small in amplitude, and systematic effects can introduce significant spurious power, particularly on the largest scales [11–18]. As a result, all large-scale structure constraints on f_{NL} to date have been limited by systematic uncertainties rather than fundamental statistical noise [19].

Cross-correlations between galaxy surveys and CMB lensing provide a robust alternative to galaxy auto-correlations for probing primordial non-Gaussianity. The main advantage comes from the fact that the dominant sources of noise and systematics in galaxy surveys are generally uncorrelated with those in CMB experiments [20–23]. This property makes cross-correlations especially powerful for controlling large-scale systematics and isolating the true cosmological signal [23, 24]. In fact, while systematics can bias the galaxy auto-correlation signal, in the cross-correlation they primarily act to increase the statistical variance without introducing a systematic shift. We note that, in principle, correlated systematics could exist: galactic emission can bias lensing reconstruction and is correlated with extinction, but no significant effects have been detected so far. Moreover, the CMB lensing reconstruction can suffer from a potential contamination from extragalactic foregrounds, the main one being the Cosmic Infrared Background (CIB) and thermal Sunyaev-Zel’dovich (tSZ) effect, particularly when the reconstruction is derived from temperature data [25]. However, such effects are expected to be small for our measurement, and no significant bias has been reported in similar cross-correlation analyses [26]. As a supporting test, we performed the measurement using the *Planck* polarization-only lensing reconstruction (available within the PR4 data release), which is strongly unaffected by CIB and tSZ contamination. While the polarization-only data yield much larger statistical uncertainties, the recovered signal shows no significant deviation from our baseline result. Another benefit of CMB lensing cross-correlations is their sensitivity to high redshifts, where the f_{NL} signature is strongest. The lensing kernel for CMB lensing peaks at $z \approx 2$, overlapping well with the redshift distribution of quasars and other high-redshift tracers. There have been several recent applications of this idea [27–29]. In particular the constraint from the Quiaia photometric quasar sample cross-correlated with *Planck* PR4 lensing [30] yielded $\sigma_{f_{\text{NL}}} \approx 25$ [28]. The reported constraint using DESI LRGs is of $\sigma_{f_{\text{NL}}} \approx 40$ [27]. Finally, the analysis of the photometric DESI Legacy Survey quasar sample cross-correlated with *Planck* PR4 lensing presented in [29], which forms the basis of the approach used in this work, obtained $\sigma_{f_{\text{NL}}} \approx 45$. Looking ahead, forecasts suggest that ongoing spectroscopic surveys like DESI, and future Stage 4 surveys like SPHEREx and LSST, combined with CMB lensing, are advancing toward the era of $\sigma_{f_{\text{NL}}} < 1$ [31–38].

Driven by this motivation, we present the first measurement of local-type non-Gaussianity from the cross-correlation between DESI DR1 quasars and *Planck* PR4 CMB lensing maps [4]. Our analysis focuses exclusively on the cross-correlation to minimize contamination from auto-correlation systematics, providing a clean and complementary test of the standard LSS analysis which employs the 3D power spectrum of high redshift tracers. The analysis builds upon that by [29]: DESI DR1 spectroscopically confirmed quasars are employed instead of the photometric targets in the DESI Legacy Survey [39], so the sample is affected by fewer imaging systematics. *Planck* PR4 lensing maps are employed, which have 10–20% less noise than the *Planck* 2018 maps [40] due to use of additional data and more optimal filtering and analysis methods [4, 30]. Our constraining power now reaches $\sigma_{f_{\text{NL}}} \sim 20$, showing that the purer quasar sample and the new pixel-free pipeline improve the constraining power by a factor $\gtrsim 2$ over [29], who find $\sigma_{f_{\text{NL}}} \sim 45$ despite the $\sim 2\times$ larger sky fraction.

In section 2, we describe the theory necessary to compute angular correlation functions at low ℓ . In section 3, we describe the quasar and CMB lensing data, and in section 4, we describe our angular power spectrum pipeline. In section 5, we validate our pipeline on mocks, to verify that we do not over-correct when mitigating the effect of imaging systematics. In section 6 the implementation of optimal redshift weights to maximize the f_{NL} signal in the cross-correlation is discussed. Finally, in section 7, we present the results, and in section 8 we compare them to previous f_{NL} constraints. Throughout this paper, we fix the cosmological parameters to the *Planck* 2018 flat Λ CDM model [41] with $h = 0.6766$, $A_s = 2.105 \times 10^{-9}$, $n_s = 0.9665$, $\Omega_m = 0.3096$, $\Omega_b = 0.049$, one neutrino with mass 0.06 eV, and $\sigma_8 = 0.8102$, but we also test the impact of freeing the primordial power spectrum (i.e., A_s and n_s) on the constraints we obtain on f_{NL} .

2 Theory

Local PNG is characterized by the dimensionless parameter f_{NL} :

$$\Phi_{\text{NG}}(\mathbf{x}) = \phi_{\text{G}}(\mathbf{x}) + f_{\text{NL}} \left(\phi_{\text{G}}^2(\mathbf{x}) - \langle \phi_{\text{G}}^2 \rangle \right), \quad (2.1)$$

where ϕ_{G} is a Gaussian random field representing the primordial gravitational potential [42]. This local transformation introduces mode-coupling between long and short wavelengths, which manifests as a scale-dependent bias in large-scale structure tracers [9, 11, 43]. The physical origin of PNG-induced scale-dependent bias lies in the modulation of halo formation by long-wavelength potential fluctuations. Taking the Laplacian of eq. (2.1), we see that the presence of PNG increases the density in the peaks of the density field:

$$\delta_{\text{NG}} = \delta_{\text{G}}(1 + 2f_{\text{NL}}\phi_{\text{G}}) \quad (2.2)$$

This shifts the effective collapse threshold:

$$\delta_c \rightarrow \delta_c(1 - 2f_{\text{NL}}\phi_{\text{G}}) \quad (2.3)$$

And, following [9], one can show that the resulting scale-dependent correction to the linear bias $b_1(z)$ is:

$$b(k, z) = b_1(z) + \frac{b_{\Phi}(z)}{T_{\Phi \rightarrow \delta}(k, z)} f_{\text{NL}}, \quad (2.4)$$

where $b_1(z)$ is the linear bias of the tracer, b_Φ is the PNG bias, giving the response to the presence of local PNG of the tracer, and the transfer function $T_{\Phi \rightarrow \delta}(k, z)$ is the transfer function between the primordial gravitational field ϕ_G and the matter density perturbation, computed as:

$$T_{\Phi \rightarrow \delta}(k, z) = \sqrt{\frac{P_{\text{lin}}(k, z)}{P_\Phi(k)}} \quad \text{with} \quad P_\Phi(k) = \frac{9}{25} \frac{2\pi^2}{k^3} A_s \left(\frac{k}{k_{\text{pivot}}} \right)^{n_s-1}. \quad (2.5)$$

Hence, through the Poisson equation, $T_{\Phi \rightarrow \delta}(k, z)$ has the well-known scale dependence:

$$T_{\Phi \rightarrow \delta}(k, z) \propto k^2 T_{\Phi \rightarrow \Phi}(k, z) \quad (2.6)$$

where $T_{\Phi \rightarrow \Phi}(k, z)$ is the usual total matter transfer function. Since f_{NL} is always paired with b_Φ , a measurement of scale-dependent bias measures the product $b_\Phi f_{\text{NL}}$. Nonetheless, this degeneracy does not affect the significance of a potential non-zero detection, which would remain a robust indication of the presence of PNG [44, 45]. The PNG bias b_Φ quantifies the logarithmic response of the galaxy number density to a change in amplitude of the matter clustering. It is defined as:

$$b_\Phi = \frac{\partial \log \bar{n}}{\partial \log \sigma_8}. \quad (2.7)$$

However, the theoretical modeling of b_Φ is a widely discussed topic [44–47] that goes beyond the scope of this present work. Therefore, we will follow the standard prescription, also known as universality relation [11]:

$$b_\Phi(z) = 2\delta_c(b_1(z) - p) \quad (2.8)$$

where $\delta_c \approx 1.686$ is the critical spherical collapse density threshold and $b_1(z)$ is the redshift-dependent linear bias. The most typical choices for the response parameter p is 1 for a mass-selected sample, while for a sample dominated by recent mergers, such as quasars, $p = 1.6$ is a more appropriate choice [11, 48]. We will report constraints for both prescriptions.

In this work, we probe $\Delta b(k)$ using the matter-galaxy cross-power spectrum, $P_{gm}(k)$. Specifically, since CMB lensing is a 2-dimensional projected field, our observable is the angular cross-power spectrum:

$$C_\ell^{\kappa\delta} = \frac{2}{\pi} \int dz_1 W_\delta(z_1) \int dz_2 \frac{W_\kappa(z_2)}{\chi^2(z_2)} \int dk P_{mm}(k, z_1, z_2) j_\ell(k\chi(z_1)) j_\ell(k\chi(z_2)) \quad (2.9)$$

where $W_\delta(z)$ is the number counts kernel:

$$W_\delta(z) = \frac{H(z)}{c} b_1(z) n(z), \quad (2.10)$$

$n(z)$ being the quasar target redshift distribution, and $W_\kappa(z)$ is the CMB lensing kernel

$$W_\kappa(z) = \frac{3}{2} \frac{H_0^2}{c^2} \Omega_m \chi(z) (1+z) \left(1 - \frac{\chi(z)}{\chi(z_\star)} \right) \quad (2.11)$$

with $\chi(z_\star)$ the comoving distance to the surface of last scattering, at $z_\star = 1090$. When including the scale-dependent bias induced by f_{NL} , we replace the linear galaxy bias $b(z)$ with a

scale- and redshift-dependent term $b(k, z)$, as defined in eq. (2.4). This assumes a linear galaxy bias, valid over the scales used here ($4 < \ell < 300$, i.e. $k \lesssim 0.1 \, h \, \text{Mpc}^{-1}$ for $0.8 < z < 3.5$).

In accordance with other PNG analysis within DESI, we fix the redshift evolution of the bias to follow the functional form of [10], which captures the expected increase of quasar bias with redshift. In practice, this redshift dependence is integrated over in the model, and therefore modifies the effective redshift at which the measurement is interpreted. We introduce a free normalization parameter b_0^i for each tomographic bin, which preserves the redshift evolution but allows the overall amplitude to vary, resulting in discontinuities between bins. In section 7.3 we demonstrate that f_{NL} depends only on the effective bias amplitude of each bin, as the constraints remain consistent across all alternative bias evolution models considered.

Eq. (2.9) is a numerically challenging three-dimensional integral over oscillatory spherical Bessel functions. The Limber approximation, typically employed to simplify the problem, approximates the spherical Bessel functions as Dirac delta functions in their first peak [49]:

$$j_\ell(k\chi) \rightarrow \sqrt{\frac{\pi}{2\ell+1}} \delta_D(\ell + 1/2 - k\chi). \quad (2.12)$$

The three-dimensional integral is reduced to a one-dimensional integral in k or χ , but this approximation is not valid on the very large angular scales, which happen to contain the most f_{NL} information. Therefore, we evaluate our theory model using **Blast**¹ [50, 51], an algorithm for calculating angular power spectra without employing the Limber approximation or assuming a scale independent growth rate, based on the use of Chebyshev polynomials. The code assumes that the power spectrum can be factorized as

$$P(k) = P_{\text{lin}}(k) + (P_{\text{nl}} - P_{\text{lin}})(k) \quad (2.13)$$

The non-linear part is negligible until $\ell > 200$, where the Limber approximation works very well. Therefore, for multipoles $\ell < 200$, we compute the signal using the full non-Limber expression evaluated using $P_{\text{lin}}(k)$, and then add the non-linear correction $(P_{\text{nl}} - P_{\text{lin}})$ evaluated with Limber approximation. For $\ell > 200$, **Blast** uses the Limber approximation with $P_{\text{nl}}(k)$.

The modeling also includes the contributions from lensing magnification and redshift-space distortions. Lensing magnification accounts for the changes on the background density caused by foreground structures. This alters the observed number counts and is effectively correlated with CMB lensing, leading to an extra contribution to the cross-correlation:

$$C_\ell^{\kappa\mu} = \frac{2}{\pi} \int dz_1 \frac{W_\kappa(z_1)}{\chi^2(z_1)} \int dz_2 \frac{W_\mu(z_2)}{\chi^2(z_2)} \int dk P_{mm}(k, z_1, z_2) j_\ell(k\chi(z_1)) j_\ell(k\chi(z_2)) \quad (2.14)$$

where the magnification bias kernel $W_\mu(z)$ is

$$W_\mu(z) = \frac{3}{2} \frac{H_0^2}{c^2} \Omega_m \chi(z) (1+z) \int_z^{+\infty} dz' n(z') (5s(z') - 2) \left(1 - \frac{\chi(z)}{\chi(z')} \right) \quad (2.15)$$

The inner integral runs over the distribution of galaxies, and $s \equiv \frac{d \log_{10} n}{dm}$ [52] is the response of the number density n to achromatic changes in the brightness dm . In section 3, we describe how s is estimated for our sample.

¹<https://github.com/sofiachiarenza/Blast.jl>.

Redshift space distortions also contribute to the observed number counts by causing galaxies to be observed in a different redshift shell due to their peculiar velocities:

$$C_\ell^{\kappa\text{RSD}} = \frac{2}{\pi} \int dz_1 W_\kappa(z_1) \int dz_2 W_{\text{RSD}}(z_2) \int dk P_{mm}(k, z_1, z_2) j_\ell(k\chi(z_1)) j_\ell''(k\chi(z_2)) \quad (2.16)$$

The RSD kernel can be expressed as:

$$W_{\text{RSD}}(z) = \frac{H(z)}{c} f(z) n(z) \quad (2.17)$$

where $f(z)$ is the logarithmic derivative of the growth rate with respect to scale factor, $f \equiv \frac{d \ln D}{d \ln a}$, and the spherical Bessel function is replaced with its second derivative. The only sensitivity to f_{NL} is through $C_\ell^{\kappa\delta}$, as neither $C_\ell^{\kappa\mu}$ nor $C_\ell^{\kappa\text{RSD}}$ depend on the galaxy bias. As discussed, f_{NL} leaves an imprint in the clustering signal through the scale dependent bias, defined in eq. (2.4). In the presence of PNG, the number counts kernel, defined in eq. (2.10), becomes:

$$W_{\delta+\text{PNG}}(z) = \frac{H(z)}{c} (b_1(z) + \Delta b(k, z)) n(z). \quad (2.18)$$

This can be factorized:

$$W_{\delta+\text{PNG}}(z) = W_\delta(z) + W_{\text{PNG}}(z), \quad (2.19)$$

with

$$W_{\text{PNG}}(z) = \frac{H(z)}{c} f_{\text{NL}} b_\Phi(z) n(z). \quad (2.20)$$

Keeping in mind that **Blast** performs the inner k -integral first, and that this step is affected by the presence of the transfer function in the scale-dependent bias, it can be worked out that the PNG contribution to the cross-correlation is:

$$C_\ell^{\kappa\text{PNG}} = \frac{2}{\pi} \int dz_1 \frac{W_\kappa(z_1)}{\chi^2(z_1)} \int dz_2 W_{\text{PNG}}(z_2) \int dk P_\Phi(k) T_{\Phi \rightarrow \delta}(k, z_1) j_\ell(k\chi_1) j_\ell(k\chi_2), \quad (2.21)$$

where the matter power spectrum has been written as $P_{mm}(k, z_1, z_2) = P_\Phi(k) T_{\Phi \rightarrow \delta}(k, z_1) T_{\Phi \rightarrow \delta}(k, z_2)$. The reason why eq. (2.21) contains only one transfer function is that the scale-dependent bias scales as $\Delta b(k, z) \propto 1/T_{\Phi \rightarrow \delta}(k, z)$, which cancels the second transfer-function factor from P_{mm} . Effectively, we can write our final model as:

$$C_\ell^{\kappa g} = C_\ell^{\kappa\delta} + C_\ell^{\kappa\mu} + C_\ell^{\kappa\text{RSD}} + C_\ell^{\kappa\text{PNG}} \quad (2.22)$$

with $C_\ell^{\kappa\text{PNG}} \propto f_{\text{NL}}$. For a more technical description of the algorithm, see [51].

Figure 1 shows the fractional contributions of magnification, RSD and PNG to the total $C_\ell^{\kappa g}$ signal with $f_{\text{NL}} = 50$ in the first tomographic bin ($0.8 < z < 2.1$). The fiducial values of $b_0^i = 1$ and $s = 0.099$ are assumed. The magnification term is a considerable fraction ($\sim 15\%$) of the clustering term, and rises in a scale-dependent way that is approximately degenerate with f_{NL} at $\ell > 30$, emphasizing the importance of low multipoles. Despite this degeneracy, the magnification bias slope is measured sufficiently accurately that it does not worsen the f_{NL} constraints. The RSD term is very subdominant, only rising to 1% of the fiducial model at $\ell < 6$. Finally, the f_{NL} contribution dominates for $\ell < 10$ and constitutes a non negligible part of the signal up to $\ell \approx 100$. Although $f_{\text{NL}} = 50$ is a very high and unrealistic value for this parameter, it shows how a non-Limber evaluation of the theory model is crucial to probe the PNG signal.




Figure 1. *Left:* Total quasar-CMB lensing cross-correlation in the first tomographic bin ($0.8 < z < 2.1$). The fiducial model with $f_{\text{NL}} = 0$ is presented in blue, and with $f_{\text{NL}} = 50$ in magenta. *Right:* Contributions of each term to eq. (2.22) as a fraction of the fiducial model with $f_{\text{NL}} = 50$. Negative terms are shown as dashed lines.




Figure 2. Data used for the cross-correlation (same as in ref. [53]). Panel (a) shows the *Planck* PR4 lensing convergence map, κ , together with its smoothed mask, obtained by applying a Gaussian filter with 1° FWHM. Panel (b) displays the DESI DR1 quasar number counts for the full sample spanning the redshift interval $0.8 \leq z \leq 3.5$, while the corresponding completeness mask is shown in panel (c). For visualization, we adopt a `HEALPIX` resolution of $N_{\text{side}} = 128$, whereas all computations are performed at $N_{\text{side}} = 2048$. All maps are displayed in a Mollweide projection and are presented in the Galactic coordinate system.

3 Data

We cross-correlate quasars from the DESI survey with lensing mass maps obtained by the *Planck* satellite, namely the PR4 convergence maps [4]. The data are visualized in figure 2. In section 3.1 we summarize the quasar sample and in section 3.2 we briefly present the employed CMB lensing data.




Figure 3. Normalized redshift distribution $n(z)$ of the DESI DR1 spectroscopic quasar sample. The different colors identify the three non-overlapping redshift bins.

Label	z -range	N_{qso}	Shot Noise	$\bar{z}$	z_{eff}	s_{μ}	$f_{\text{sky}}^{\text{PR4} \times \text{DR1}}$
g_1	$0.8 \leq z < 2.1$	856 831	2.6×10^{-6}	1.49	1.44	0.099	19.8
g_2	$2.1 \leq z < 2.5$	194 754	11.2×10^{-6}	2.28	2.27	0.185	18.9
g_3	$2.5 \leq z \leq 3.5$	171 806	12.7×10^{-6}	2.85	2.75	0.244	18.6

Table 1. Overview of DESI DR1 quasar samples used for the cross-correlation with *Planck* CMB lensing. We list the three redshift bins $\{g_1, g_2, g_3\}$, their redshift ranges, quasar counts, shot-noise levels, mean redshift $\bar{z}$, effective redshift z_{eff} , magnification bias s_{μ} , and the sky-fraction overlap with the PR4 CMB-lensing convergence maps, $f_{\text{sky}}^{\text{PR4} \times \text{DR1}}$ (in percent).

3.1 DESI DR1 spectroscopically confirmed quasars

Our analysis employs a spectroscopic catalogue of 1 223 391 quasars in the redshift interval $0.8 \leq z \leq 3.5$ taken from the DESI DR1 quasar sample [54] released with the survey’s first public data release [55]. DESI is a highly multiplexed fiber-fed spectrograph located on the 4m Mayall telescope at the Kitt Peak National Observatory [56, 57]. It has the capability to record up to 5 000 spectra in a single exposure [58–62]. Throughout a eight-year campaign, DESI aims to map up to $17\,000 \text{ deg}^2$, which will provide precise measurements of baryon acoustic oscillations (BAO) and redshift-space distortions (RSD) from both galaxies and quasars across $0 \lesssim z \lesssim 3.5$ [58, 63]. Following successful validation of the survey [64] and an early data release [65], the data from the first year has already provided competitive BAO constraints [66–68], accurate RSD measurements [69], and the currently strongest large-scale structure limit on primordial non-Gaussianity, derived from the quasar power spectrum [10].

Table 1 summarizes the DR1 quasar subsamples, while figure 3 illustrates their redshift distributions alongside the tomographic binning scheme used in this analysis. This scheme is consistent with the accompanying study of DESI DR1 quasars provided in [53]. The sample is divided into three redshift bins to retain radial information. The first bin corresponds to the DESI fiducial range, $0.8 < z < 2.1$, while the remaining quasars, extending up to $z = 3.5$,

are split into two approximately equally populated bins, $2.1 < z < 2.5$ and $2.5 < z < 3.5$. We utilize the clustering catalogues described in [70, 71] together with redshifts produced by the DESI spectroscopic pipeline [72–75].

Table 1 also reports the measured values of magnification bias for the three tomographic bins. We measure the magnification bias slope, $s \equiv d \log_{10} N / dm$, by perturbing the Legacy Imaging Survey DR9 photometry [39] uniformly by ± 0.05 mag, re-running quasar target selection, and measuring the resulting change in number density. The values obtained from brightening and dimming are consistent, so we adopt their mean as our fiducial slope. Since quasars are dominated by point sources and selected using total magnitudes rather than fiber magnitudes, it is an excellent approximation to treat lensing magnification as a simple brightening or dimming of their flux. This procedure measures the slope using the observed (post-magnification) magnitude distribution, but the induced bias is negligible, as magnification only scatters a small fraction of quasars across the flux limit. We find that the slope varies with redshift, and therefore use the value appropriate for each bin: $s = 0.099$ for $0.8 < z < 2.1$, $s = 0.185$ for $2.1 < z < 2.5$, and $s = 0.244$ for $2.5 < z < 3.5$.

3.2 *Planck* PR4 CMB lensing map

Gravitational lensing by large-scale structure deflects CMB photons, imprinting subtle distortions in the temperature and polarization fields that can be exploited to reconstruct the lensing potential [see, e.g., 76, for a review]. For this work we employ the *Planck* PR4 CMB lensing convergence maps, κ , which are publicly available² [4]. Starting from the global minimum-variance (GMV) spherical-harmonic coefficients $\kappa_{\ell m}$, we resample the data onto a HEALPix grid with NSIDE=2048 [77].

The PR4 release benefits from the NPIPE processing pipeline [4], which provides improved, uniformly processed CMB maps that serve as inputs to the lensing reconstruction. The reconstruction itself was upgraded in PR4: the collaboration adopted a GMV estimator that incorporates temperature-polarization correlations [78] and implemented an anisotropic filtering and local-noise weighting scheme [30]. Combined with the $\sim 8\%$ additional CMB observations obtained during satellite repointing, these improvements yield an overall $\sim 20\%$ increase in signal-to-noise. As a result, the PR4 lensing map remains signal-dominated up to multipoles $\ell \simeq 70$, compared with $\ell \simeq 40$ for PR3. The effective reconstruction noise power spectrum, $N_{\ell}^{\kappa\kappa}$, supplied by the *Planck* collaboration, is incorporated in the analytical covariance matrix, described in section 4. The lensing mask retains a sky fraction of $f_{\text{sky}} = 67.1\%$.³ Its overlap with the tomographic DESI DR1 quasar catalogue is summarized in table 1.

4 Measuring angular power spectra

We measure clustering directly from the un-pixelised DESI quasar catalogue using the catalog-based *pseudo- C_{ℓ}* estimator [79, 80]. Our analysis utilizes the NaMaster implementation⁴ [80],

²<https://pla.esac.esa.int> and https://github.com/carronj/planck_PR4_lensing.

³The sky fraction is defined as $f_{\text{sky}} = \sum_{i=1}^{N_{\text{pix}}} x_i / N_{\text{pix}}^{\text{tot}}$, where $N_{\text{pix}}^{\text{tot}} = 12 N_{\text{side}}^2 \approx 5 \times 10^7$ and x_i is the value of the mask in pixel i .

⁴<https://github.com/LSSTDESC/NaMaster>.

which builds upon the algorithm described in [81] and leverages the DUCC library⁵ for maximal efficiency [e.g., 82]. In the standard pixelized approach, galaxies and randoms are first binned onto a high-resolution HEALPix grid, and the overdensity field is formed as $1 + \delta_{\text{gal}}$ by dividing the two maps. While conceptually straightforward, this division can lead to numerical instabilities in pixels where the survey completeness is very low. In early tests with this pixel-based procedure, we observed a significant excess of power on large scales in the measured angular power spectra, which could not be accounted for by known imaging systematic effects. By carefully tracking the origin of this excess, we found it was caused by instabilities in the overdensity maps calculation in low-completeness pixels. Switching to the catalog-based *pseudo- C_ℓ* estimator completely resolved this issue, yielding robust measurements across all scales. We analyze angular scales in the range $4 \leq \ell \leq 300$. The binning scheme uses $\Delta\ell = 10$ for $\ell \leq 100$, where most of the signal resides, and $\Delta\ell = 50$ for $\ell \geq 100$, yielding 14 bins in total. Finally, we stress that the NaMaster estimator automatically subtracts the Poisson shot noise from the measured auto-power spectra. We have verified that the shot-noise levels estimated internally by the code agree with the theoretical expectations obtained from the number density in each redshift bin (table 1).

Furthermore, the pixelized approach is subject to aliasing and pixel-window function effects. Mitigating these problems usually requires ultra-fine grids (with an arbitrary choice of NSIDE) and very large random catalogues, driving up both memory and computational costs. The catalog-based pipeline instead adopts the “FKP” approach from three-dimensional analyses by working with the linear difference $n_g - \alpha n_r$ instead of the ratio of data over randoms catalogs. Because this operation is linear, we can simply subtract the direct spherical-harmonic transforms of data and scaled randoms, avoiding numerical instabilities and pixel-related problems. In the limit that the objects are point-like, the spherical harmonic transform of these fields is simply a sum of spherical harmonics evaluated at the positions of the objects, $Y_{\ell m}(\theta_i, \phi_i)$, making the evaluation of the estimator straightforward. For a detailed description of the algorithms, the reader is referred to [79, 81].

Pixel-free estimators offer a promising approach for future analyses involving cross-correlations with discrete tracers, however the corresponding analytic covariance calculations for catalog-based fields are still under development [83]. Therefore, following the approach used, for example, in [28, 84], we adopt the standard analytic Gaussian, window-convolved covariance approximation implemented in the `gaussian_covariance` method of the NaMaster package to estimate the errors in our analysis. Since this covariance routine requires a pixelized map of the data field, we construct such a map from the data $a_{\ell m}$ ’s obtained with the catalog-based estimator, generating HEALPix maps at NSIDE = 2048. Moreover, this approach also requires some theory spectra, which, as discussed in section 2, we generate using Blast [50, 51]. The code gives predictions for $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_j}$ (where we add the shot noise contribution as in table 1 to $C_\ell^{g_i g_i}$) in the fiducial cosmology. For the lensing power spectrum $C_\ell^{\kappa\kappa}$ we use the publicly available PR4 lensing spectra with the provided lensing reconstruction noise $N_{\kappa\kappa}$. As in [53], we use an iterative approach: we fit the cosmological and nuisance parameters with the covariance matrix evaluated with the theory curves just described, and use the resulting best-fit parameters to recalculate the fiducial spectra and tune the covariance matrices.

⁵<https://gitlab.mpcdf.mpg.de/mtr/ducc>.




Figure 4. Monte Carlo normalization correction applied to the cross-spectra $C_\ell^{\kappa g_i}$ which amounts to an approximately constant amplitude change of $\approx 8\%$.

4.1 Monte Carlo normalization correction for the *Planck* lensing maps

A well-known effect in CMB lensing reconstruction is a misnormalization of the reconstructed field caused by the masks and anisotropic filtering applied in order to perform the reconstruction. The mode coupling introduced in this procedure is not modeled into the cross-spectra obtained with the **NaMaster** algorithm. As a result, the reconstructed convergence $\hat{\kappa}$ does not have the same normalization as the true convergence κ , and a correction is required [85–87].

In practice, this correction is determined through simulations: the basic idea is to compare the cross-correlation between the input convergence (appropriately masked) and the reconstruction with the known input power spectrum. This yields a Monte Carlo normalization factor that is dependent on the mask and the footprint:

$$A_\ell^{\text{MC}} = \frac{C_\ell^{\kappa_{\text{in}}, \kappa - \text{mask}} \kappa_{\text{in}, g - \text{mask}}}{C_\ell^{\hat{\kappa} \kappa_{\text{in}, g - \text{mask}}}}, \quad (4.1)$$

where $\hat{\kappa}$ is the masked CMB lensing reconstruction, $\kappa_{\text{in}, \kappa - \text{mask}}$ is the input lensing convergence masked with the lensing mask, and $\kappa_{\text{in}, g - \text{mask}}$ is the input convergence masked using the galaxy mask. An unbiased estimate of the CMB lensing cross-spectrum is then obtained as

$$C_\ell^{\hat{\kappa} \text{MC} g} = C_\ell^{\hat{\kappa} g} A_\ell^{\text{MC}}. \quad (4.2)$$

Following [53], we implement this correction in a simulation-based, mode-by-mode fashion. Specifically, we generate a suite of CMB lensing reconstructions using the appropriate mask:

$$T_L \equiv \frac{\sum_\ell W_{L\ell} \sum_i^{N_{\text{sim}}} \sum_{m=-\ell}^\ell \{M^\kappa \kappa^i\}_{\ell m} \{M^g \kappa^i\}_{\ell m}^*}{\sum_\ell W_{L\ell} \sum_i^{N_{\text{sim}}} \sum_{m=-\ell}^\ell \{M^\kappa \hat{\kappa}^i\}_{\ell m} \{M^g \kappa^i\}_{\ell m}^*}, \quad (4.3)$$

where W_L indicates the bandpower window functions, M denotes the mask, κ the input (and $\hat{\kappa}$ the reconstructed) convergence, i indexes the simulations, and $XY_{\ell m} \equiv \int d^2\hat{\mathbf{n}} Y_{\ell m}^*(\hat{\mathbf{n}}) X(\hat{\mathbf{n}}) Y(\hat{\mathbf{n}})$ with $Y_{\ell m}$ the spherical harmonics. For the *Planck* PR4 maps over the DESI DR1 footprint, this correction is at the level of $\approx 8\%$ on the scales of interest (see figure 4).

5 Imaging systematic weights validation

Spatial variations in imaging quality and foregrounds imprint spurious modulations in the angular number density of DESI targets. If left uncorrected, these modulations contaminate the largest-scale Fourier modes that are most sensitive to primordial non-Gaussianity.

5.1 The DESI systematic weights

DESI adopts the template-fitting approach introduced in recent large-scale structure surveys. The technique derives a per-object weight, w_{sys} , by quantifying how the target surface density varies with each imaging attribute and then correcting for those dependencies. More specifically, DESI attaches three multiplicative weights to every object:

$$w_{\text{tot}} = w_{\text{comp}} w_{\text{sys}} w_{\text{zfail}}, \quad (5.1)$$

ensuring that the ratio of weighted data to weighted randoms remains flat across the footprint once all known observational selection effects are removed. Completeness weights w_{comp} correct for the probability that a target is allocated a fiber, and redshift failure weights, w_{zfail} , mitigate spatial fluctuations in redshift success produced by variations in exposure time, focal-plane position, and hardware status. The imaging systematics weights w_{sys} , the focus of this section, remove spurious target-density fluctuations that track imaging conditions such as depth, seeing, dust, and stellar density. Full implementation details for the first two weights are given in [71].

Below, we discuss only the imaging systematic weights, since this is the element we test extensively for its impact on the quasar angular power spectrum. The other two weights, completeness and redshift-failure, are not a concern for our analysis. Systematic weights have many possible choices, including different templates and implementations, and can have a large impact on the results. By contrast, it is well understood how to correctly account for completeness. Redshift-failure weights have a much smaller impact than the systematic errors and therefore do not significantly affect the analysis [71, 88]. The DESI catalog in the latest released version (v1.5), comes with three possible sets of systematic weights: the neural net method `sysnet`, denoted `WEIGHT_SN` in the catalogs; the random forest method `regressis`, denoted `WEIGHT_RF`; the linear method applied to eBOSS [88] LSS catalogs, identified as `WEIGHT_IMLIN`. In the present work, we only consider two different sets of linear weights as it is known that both machine learning based weights, namely the `sysnet` neural network [89] and the `regressis`⁶ random forest weights [90], over-correct for excess large scale power introduced by systematics, especially in the case of the quasar sample [10, 71, 90]. The second set of linear weights is also computed with `regressis` at `HEALPix` map level using a linear regression algorithm. Note that, compared to the catalog linear weights, the linear regression here does not fit the data to binned statistics but rather the fluctuation at `HEALPix` map level. The first are identified as `WEIGHT_IMLIN`, the latter as `WEIGHT_Linear`.

5.2 Correcting for excess mode removal

As described above, we test two sets of linear systematic weights that are derived from the same imaging templates but differ in their implementation. Accurately assessing the performance

⁶<https://github.com/echaussidon/regressis>.

of such weights is crucial, as they can over-correct and remove large-scale clustering signal, thereby weakening the f_{NL} constraints. The correction for this mode removal by reweighting modes is related to the angular integral constraint (AIC) correction [70]. This effect has been shown to be particularly severe for neural network-based weighting schemes [10], which we therefore do not consider in this analysis. Nevertheless, it remains important to quantify any residual bias for the two linear weighting schemes tested here. To quantify the impact of our systematic-mitigation weighting scheme, we use custom-generated mocks with no imaging systematic contamination, so that tracer densities are, on average, uncorrelated with the imaging features. Any deviation in the angular power spectra after applying the systematic weights, will be due to the mitigation method itself and must be corrected on the measurement on the data.

5.2.1 Creating the mocks

We have created 500 uncontaminated, catalog-based lognormal mocks that accurately reproduce the DESI DR1 footprint and clustering properties. This number of realizations is sufficient to ensure that the uncertainty associated with the mock-based correction remains well below the statistical errors across the relevant scales. The first step is to generate correlated Gaussian maps for the three redshift bins used in the analysis ($0.8 < z < 2.1$, $2.1 < z < 2.5$, $2.5 < z < 3.5$) and for the CMB lensing map. This is done using the **healpy** package [91], specifically the **synfast** function, which synthesizes maps at a chosen resolution from input angular power spectra C_ℓ . The theoretical C_ℓ are computed with the **GLASS**⁷ package [92], which allows the generation of lognormal fields with prescribed two-point statistics. Because the lognormal transformation, defined as:

$$f(X; \lambda) = \lambda (e^X - 1), \quad (5.2)$$

is nonlinear and modifies the underlying statistics of the field, **GLASS** determines, through an iterative algorithm, the Gaussian input power spectrum G_ℓ that will yield the target C_ℓ after applying the transformation. In our case, we set the shift-parameter λ to 1 and transform the three Gaussian fields representing the tomographic bins. The target angular spectra C_ℓ are computed with **Blast** [50, 51] and the bias is adjusted to match the data, ensuring that the final mocks reproduce the same large-scale clustering statistics as the data. We convert the pixelized lognormal maps into mock catalogs by Poisson-sampling the number of galaxies in each pixel according to its overdensity. Each mock galaxy is then assigned sky coordinates (RA, DEC) by randomly perturbing the position around the pixel center using an effective pixel radius $\theta_{\text{pix}} = \sqrt{A_{\text{pix}}}/2$. During this process, the maps are normalized so that the resulting mock catalogs contain approximately the same number of objects per redshift bin as the data. As this process is very approximate on small scales, we generate the maps at a very high resolution (**NSIDE** = 8192) to ensure that, on the relevant scales ($\ell < 300$), the resulting angular clustering would not be affected by this sampling procedure.

The final step to obtain realistic mocks is to apply angular cuts and downsampling to reproduce the survey footprint and completeness. The DESI DR1 footprint exhibits substantial small-scale structure, as the survey completeness remains low at this stage. To

⁷<https://github.com/glass-dev/glass>.




Figure 5. *Left:* A zoom into the survey footprint. In darker blue, the DESI DR1 quasars, while in light blue are the galaxies making up one of the 500 mocks. The agreement between the two footprint is very good, confirming the success of our procedure to replicate the DESI survey geometry. *Right:* Average angular power spectrum $C_\ell^{\kappa g}$ over 500 mocks with the corresponding 1σ error bars. As desired, the mocks have a power spectrum that matches the input theory curve.

reproduce these variations as accurately as possible, we perform four successive downsampling steps, designed to account for the main observational effects impacting the real survey, while avoiding the computational cost of running each mock through the `fiberassign` software [93] and detailed veto masks. In particular, we select areas of the sky covered during DR1, remove areas around bright stars and regions affected by bad imaging systematic conditions. Finally, we account for the survey completeness by constructing completeness maps directly from existing survey random catalogs. Specifically, we use the official DESI DR1 quasar random files, which encode the spatial completeness pattern of the survey, together with the Legacy Survey DR9 randoms [39], which are uniform over the sky. Both sets of randoms are pixelized at $\text{NSIDE} = 2048$, and the ratio of the two maps provides an estimate of the angular completeness. Each mock catalog is then downsampled according to this completeness map, ensuring that the mocks reproduce the same spatial variations in target completeness as the data. This procedure is correct up to the resolution of the maps, which is enough for our analysis, restricted to scales $\ell < 300$. The results of these footprint and clustering validations are shown in figure 5. The left panel demonstrates that the mocks accurately reproduce the DESI DR1 survey geometry, as evidenced by their excellent agreement with the official quasar randoms. This confirms that the completeness correction procedure described above effectively transfers the angular selection function to the mock catalogs. The right panel shows that the average cross-correlation $C_\ell^{\kappa g}$ measured from 500 mocks is consistent with the input theoretical prediction, within the expected statistical uncertainties. Together, these tests verify that our catalog-based lognormal mocks correctly reproduce both the survey footprint and the target clustering statistics.




Figure 6. Left: Average angular auto power spectrum C_ℓ^{gg} of 500 mocks with different linear weighting schemes, in the first redshift bin considered in the analysis ($0.8 < z < 2.1$). Right: Ratio of the weighted to unweighted power spectra. The ratio deviates from 1 on large scales, showing that the weights are removing signal from the analysis. The shaded grey area highlights the 1% region around unity.

5.2.2 Test on systematics-free mocks

To evaluate the significance of mode removal using either the linear systematic correction weight or those creating using the linear regressor within the **regressis** package, we measure the angular power spectra of the mocks using each set of weights and compare them to the unweighted case (where all weights are set to unity). As the mocks do not have systematic contamination, a suppression of large-scale power in the weighted cases would indicate that the weights are spuriously removing cosmological signal instead of systematics, requiring a corresponding correction. We examine how the systematic weights impact both the galaxy auto-spectrum, C_ℓ^{gg} , which is most sensitive to observational systematics and contributes to the covariance, and the cross-spectrum with CMB lensing, $C_\ell^{\kappa g}$, which is used in our main analysis. Figure 6 shows the results for the galaxy autocorrelation C_ℓ^{gg} , indicating that both sets of weights are removing power on very large scales ($\ell < 20$). We use the ratio, displayed on the right side of the plot, as a multiplicative correction to the power spectra measured on the data. The impact is more pronounced for the “catalog” linear weights, and the underlying cause is currently under investigation, as it must reside in the implementation of the weights. Figure 7 shows the corrected autocorrelation measurement: we note that, even after accounting for angular mode removal, the auto power spectrum does not show signs of excess large scale power, a huge improvement with respect to the first analysis using the DESI DR9 Legacy Survey [29]. This can be attributed to the fact that the DESI DR1 spectroscopic catalog has a much higher purity compared to the photometric targets in the Legacy Survey and to the employment of the new catalog based estimator, which is more stable and robust to numerical instabilities and pixel related issues, as discussed in section 4. In this analysis we will adopt the more conservative and robust approach of using only the cross-correlation to constrain f_{NL} . Nonetheless, this result is encouraging as there is no need to model excess large scale power in the covariance, as in the previous analysis. Moreover, it




Figure 7. Auto-correlation angular power spectra measured from the catalog with the two possible sets of weights and corrected for mode removal. The error bars are computed from an analytical gaussian covariance computed using the **NaMaster** package, as discussed in section 4.




Figure 8. Left: Average angular cross-power spectrum $C_\ell^{\kappa g}$ of 500 mocks with different linear weighting schemes (pixel-based, labeled as “regressis,” or object-based, labeled as “catalog”), in the first redshift bin considered in the analysis ($0.8 < z < 2.1$). Right: Ratio of the weighted to unweighted power spectra. In this case, no significant deviation from 1 is observed. The shaded grey area highlights the 1% region around unity.

indicates that we are going in the direction of being able to safely use the auto-correlation information for the f_{NL} analysis in future DESI data releases.

Figure 8 shows no sign of over-subtraction of angular modes in the CMB lensing cross-correlation, which we employ in this work. Hence, we decided to not correct the measurement, an example of which is shown in figure 9.




Figure 9. Cross-correlation angular power spectra measured from the catalog with the two possible sets of weights. In this case, no correction is applied as there is no sign of angular mode removal introduced in the cross correlation. The error bars come from the analytic estimation of the covariance matrix performed using the `NaMaster` package, described in section 4.

6 Optimal redshift weighting for f_{NL}

In the present work, we tested the use of optimal redshift weights for the angular two-point statistics, with the aim of improving constraining power of the analysis. The idea of optimal redshift weighting was first developed in [94, 95] and is now routinely applied in three-dimensional power spectrum analyses. For reference, the optimal f_{NL} weight for the galaxy monopole is given by [94, 96, 97]:

$$w_{\hat{P}_0}(z) = w_{\text{FKP}}^2 b_{\Phi}(z) D(z) \left[b(z) + \frac{f(z)}{3} \right], \quad (6.1)$$

where w_{FKP} is the standard FKP weight [98]. Since this is a weight for the power spectrum, the corresponding per-galaxy weight is its square root.

In practice, this corresponds to cross-correlating two weighted fields, respectively proportional to $w_{\text{FKP}}(b(z) - p)$ and $w_{\text{FKP}} \left[b(z) + \frac{f(z)}{3} \right]$, as in [97]. This weighting scheme enhances the constraining power on f_{NL} by explicitly accounting for the redshift dependence of the PNG signal. We extend this framework to the angular cross-correlation between CMB lensing and galaxy overdensity, $C_{\ell}^{\kappa g}$, and derive an optimal estimator for this observable. The details of the derivation are given in appendix A. We find that the optimal redshift weight that maximizes the f_{NL} sensitivity of $C_{\ell}^{\kappa g}$ is:

$$w_{\text{opt}}(z) = (b(z) - p) w_{\text{FKP}}, \quad (6.2)$$

normalized such that $\int w_{\text{opt}}(z) dz = 1$. As in the power spectrum case, the resulting weights effectively up-weight higher-redshift galaxies, where the f_{NL} signature is strongest. As one




Figure 10. Redshift distribution of the DESI DR1 quasars. The plot shows the $n(z)$ of the three tomographic bins in which the data is divided for this analysis, in order of increasing redshift. The solid lines represent the fiducial $n(z)$, and the dashed lines show the redshift distribution after weighting the galaxies with the weights in eq. (6.2). The displayed redshift distributions are normalized in each bin.

might expect, we only get one power of the FKP weight here, as opposed to the two powers in the power spectrum case.

The implementation of these weights is straightforward. We evaluate the optimal weights for each galaxy in the catalog using the expressions derived above. The CMB lensing kernel W_κ is defined in eq. (2.11), while b_Φ is given in eq. (2.8), with the linear bias $b(z)$ following the model of [10]. The growth factor $D(z)$ is computed using the Boltzmann solver CAMB [99], and the FKP weights [98] are already provided as an attribute of the galaxy catalog. In the default computation of the redshift distribution $n(z)$, each galaxy is weighted according to eq. (5.1) before binning. When including the optimal weighting scheme, the total weight assigned to each galaxy becomes $w = w_{\text{tot}} w_{\text{opt}} = w_{\text{comp}} w_{\text{sys}} w_{\text{zfail}} w_{\text{opt}}$, effectively modifying the inferred $n(z)$. As shown in figure 10, the optimal weights up-weight the high-redshift part of the sample, where the sensitivity to f_{NL} is greatest. This occurs because galaxies at higher redshift are more affected by a potential PNG signal, which enhances their large scale clustering. On the other hand, at the lowest redshifts the weights become negative. In the first redshift bin (in particular in the range $0.8 < z < 0.9$), some galaxies receive negative weights. This is expected and does not pose any issue for the analysis: the negative values simply indicate that galaxies in this range would respond to a stronger PNG signal with a reduction in clustering amplitude. Measuring f_{NL} from such an absence of clustering is intrinsically more difficult, and the optimal weighting appropriately accounts for this effect.

To apply the optimal weights consistently throughout the analysis, we must consider all components of the pipeline that are affected. A change in the galaxy redshift distribution modifies the theoretical predictions, since $n(z)$ enters the number counts kernel $W_\delta(\chi)$ defined in eq. (2.10). Consequently, the theory curves used in the analytic covariance computation (see section 4) must also be updated. The analysis performed with optimal weights therefore employs a covariance matrix recomputed following the same procedure described in section 4, but using the weighted $n(z)$. Finally, the weighting scheme must also be applied to the data when measuring the angular power spectra. Figure 11 illustrates the effect of the optimal weights on the measured $C_\ell^{\kappa g}$ and on the corresponding theoretical predictions. The comparison between the default configuration (defined in section 5) and the weighted case




Figure 11. $C_\ell^{\kappa g}$ in the first tomographic bin. In navy, the measurement and theory curve in the default configuration (i.e., using the `regressis` linear weights, see section 5). In light green, the case where the optimal weights are applied. This plot showcases the impact of the optimal weights on the pipeline.

highlights how both the measurement and the theoretical model respond to the modified redshift distribution.

7 Results and discussion

In this section we present tomographic constraints on primordial non-Gaussianity obtained from the cross-correlation between the *Planck* PR4 CMB lensing convergence maps and spectroscopically confirmed quasars from DESI DR1. Our analysis covers the redshift range $0.8 \leq z \leq 3.5$, divided into three tomographic bins, and includes a total of 1,223,391 quasars across 7,200 deg^2 of sky. We use the angular power spectrum estimator described in section 4 to measure the cross-spectra $C_\ell^{\kappa g}$. Our parameter inference is performed with `Turing.jl`⁸ [100], which provides an interface to define probabilistic models in terms of explicit priors and our differentiable `Blast.jl` likelihood. Posterior exploration is carried out using the No-U-Turn Sampler (NUTS) algorithm [101], a self-tuning Hamiltonian Monte Carlo sampler that efficiently leverages gradient information and has been widely adopted for cosmological inference. In this baseline configuration, we do not employ the optimal weighting scheme, which we instead treat as an analysis variation in section 7.1. The main results is presented in table 2, visualized in figure 12 and discussed in the following text. Throughout this section we also present results for various analysis variations and tests that we performed, like the inclusion of optimal f_{NL} weights and different redshift cuts and bins. Everything is summarized in figure 17, showcasing the robustness of the analysis. As discussed in section 1, uncertainties in the theoretical prediction of the non-Gaussian bias coefficient b_Φ create a

⁸<https://github.com/TuringLang/Turing.jl>.




Figure 12. As in figure 9, the scatter points represent the measurements on the data performed with two possible weighting schemes: the default choice **regressis** in green, and the alternative linear weighting scheme implemented in the DESI catalog in purple. The dashed lines represent the corresponding best fit curves when assuming $p = 1.6$. This baseline result does not utilize the optimal weighting scheme.

degeneracy between f_{NL} and b_Φ . In general, one could constrain the product $f_{\text{NL}} b_\Phi$ directly, without assuming a specific relation between b_Φ and b_1 [see e.g. 44], since a detection of the product would remain physically meaningful. Our goal here, however, is not to explore this broader parameterization, but to assess how the inferred f_{NL} shifts under reasonable choices for the $b_\Phi(b_1)$ relation. We therefore consider two benchmark cases commonly used in the literature [see e.g. 10, 28, 29]: $p = 1.6$, corresponding to a recent-merger scenario, and $p = 1.0$, corresponding to the universality assumption. We impose flat priors on all parameters, with $f_{\text{NL}} \sim \mathcal{U}(-200, 200)$ and $b_i \sim \mathcal{U}(0, 2)$.

To assess the robustness of our results, we performed a series of analysis variations and validation tests, described below. Figure 12 presents the measured angular power spectra together with the best-fit model for the two sets of systematic weights considered. In the third tomographic bin, the lowest multipole ($\ell = 6.5$) shows a deviation of $\approx 2\sigma$ from the best-fit prediction. This feature is consistent with statistical noise and does not dominate the total χ^2 . However, we know that the third bin is the one where the quasar sample has lower purity, so it is possible that some remaining systematic in the CMB lensing maps could correlate with uncorrected systematics in the sample resulting in the large scale extra power. For this reason, we run the analysis in the default configuration removing that point, which is located at scales very relevant to constrain f_{NL} . The result in this case is $f_{\text{NL}} = -22^{+27}_{-34}$ as opposed to $f_{\text{NL}} = 2^{+28}_{-34}$ when all the points are considered. The error bars stay consistent, but there is a $\approx 0.8\sigma$ shift in the best fit value. We decided to keep the analysis as it is and not exclude that point from the pipeline as the evidence for doing that is not compelling enough to avoid falling into confirmation bias.

The marginalized constraints corresponding to each configuration are summarized in table 2. The best-fit values are consistent within 1σ across all weighting schemes, and the reduced χ^2 values confirm the overall quality of the fits. If we use $p = 1$ instead of $p = 1.6$, the constraints become tighter: this is to be expected as, in that case, the quasars are more sensitive to variation in f_{NL} . The same results can be visualized in figure 13, which shows the contours for the 4 free parameters. The constraints on the bias scaling amplitudes b_0^i are not

p	regressis linear weights	catalog linear weights
1.6	$f_{\text{NL}} = 2^{+28}_{-34}$ ($\chi^2 = 56$)	$f_{\text{NL}} = -14^{+27}_{-33}$ ($\chi^2 = 53$)
1.0	$f_{\text{NL}} = 6^{+20}_{-24}$ ($\chi^2 = 54$)	$f_{\text{NL}} = -4^{+19}_{-22}$ ($\chi^2 = 53$)

Table 2. Constraints on f_{NL} for different choices of p and weighting schemes (calculated without the additional optimal f_{NL} weights described in section 6. The results from the fiducial scenario are reported in the top left corner. The effective number of degrees of freedom in the analysis is 38, giving reduced χ^2 values of 1.47 and 1.39 respectively.




Figure 13. Corner plots showing the constraints for the 4 free parameters in the analysis, namely f_{NL} and a bias amplitude for each of the 3 tomographic bins. *Left*: constraints in the fiducial configuration, using the **regressis** linear weights. *Right*: results of the analysis performed with a different choice of systematic weights, i.e. the ones implemented in the official DESI catalog.

affected by variations in the analysis: table 3 shows the marginalized posteriors for the bias amplitudes for the two possible choices of systematic weights. Those results have been obtained with $p = 1.6$. Interestingly, the bias amplitude that we fit for the third bin comes out to be $\approx 20\text{--}30\%$ lower than expected, with a best fit value of $b_0^3 = 0.68$, very robust to variations in the analysis. This discrepancy has emerged in other DESI analyses [53], however the reason is still unknown and is currently being investigated. Nonetheless, since this fact can raise doubts on the validity of the assumed bias relation [10], we test the impact of that choice in section 7.3.

To test the impact of the tomographic redshift bins on the constraints, we also performed the analysis using a single redshift bin $0.8 < z < 3.5$. The measurement on the data in this case is displayed in figure 14, while the marginalized constraints for f_{NL} are reported in table 4. The constraining power in this case deteriorates, indicating how the redshift binning or the inclusion of optimal weights, as discussed in section 7.1, is important in extracting f_{NL} information.

	regressis linear weights	catalog linear weights
b_0^1	1.13 ± 0.05	1.13 ± 0.05
b_0^2	1.14 ± 0.13	1.17 ± 0.12
b_0^3	0.68 ± 0.12	0.68 ± 0.12

Table 3. Bias amplitude constraints for the two possible choices of linear weights tested throughout this paper. The default choice are the **regressis** linear weights, but the weight choice does not have an impact on the bias amplitudes as shown in this table.




Figure 14. $C_\ell^{\kappa g}$ when the quasars are not divided into tomographic bins.

p	regressis linear weights	catalog linear weights
1.6	$f_{\text{NL}} = 17_{-40}^{+30} (\chi^2 = 11)$	$f_{\text{NL}} = 4_{-41}^{+33} (\chi^2 = 11)$
1.0	$f_{\text{NL}} = 12_{-28}^{+24} (\chi^2 = 11)$	$f_{\text{NL}} = 3_{-27}^{+23} (\chi^2 = 11)$

Table 4. Constraints on f_{NL} for the analysis without tomographic bins. Again, different choices of p and weighting schemes have been tested and the results from the fiducial scenario are reported in the top left corner. The effective number of degrees of freedom in the analysis is 12, giving reduced χ^2 values of 0.91. This corresponds to a p -value of 0.5, indicating that the model provides a good fit to the data.

Optimal weights	regressis linear weights	catalog linear weights
No	$f_{\text{NL}} = 2_{-34}^{+28} (\chi^2 = 56)$	$f_{\text{NL}} = -14_{-33}^{+27} (\chi^2 = 53)$
Yes	$f_{\text{NL}} = 19_{-31}^{+25} (\chi^2 = 62)$	$f_{\text{NL}} = 9_{-31}^{+25} (\chi^2 = 59)$

Table 5. Comparison of the results with and without applying the optimal weighting scheme. The fiducial value of $p = 1.6$ is assumed.

Optimal weights	regressis linear weights	catalog linear weights
No	$f_{\text{NL}} = 17_{-40}^{+30} (\chi^2 = 14)$	$f_{\text{NL}} = 4_{-41}^{+33} (\chi^2 = 14)$
Yes	$f_{\text{NL}} = 10_{-35}^{+28} (\chi^2 = 14)$	$f_{\text{NL}} = -4_{-36}^{+28} (\chi^2 = 14)$

Table 6. Comparison of the results with and without applying the optimal weighting scheme for the analysis with a single tomographic bin. The fiducial value of $p = 1.6$ is assumed.

7.1 The optimal weights results

In existing literature [10, 94, 97, 102] about optimal weighting schemes to maximize the f_{NL} signal, it was found that the improvement of constraining power coming from this technique is around 8–10% in the case of the 3D power spectrum. To test this in our own pipeline, we compare our baseline results to the case where the optimal weighting scheme is also applied. As the marginalized constraints in table 5 show, we find a similar gain (≈ 5 –10%), however we also observe relevant shifts in the best fit values. We also test the impact of optimal weights in the single bin analysis, whose results are displayed in table 6: while the error bars only marginally improve, again the best fit values for f_{NL} suffer from quite significant shifts (from 0 to 20, $\approx 0.7\sigma$, and from -16 to 10, $\approx 0.8\sigma$). Despite the substantial shifts in the posterior means, the mean obtained with optimal weighting falls within the 1σ range of our default analysis. Nonetheless, these shifts suggest that the data is quite noisy (as illustrated in figure 12) and characterized by large error bars. This noise could potentially obscure the advantages of employing optimal weights. Furthermore, given that the f_{NL} signal is dependent on redshift and that we are dividing our dataset into three tomographic bins, it is possible we are already extracting the majority of the available information.

To investigate these hypotheses regarding the performance of the optimal weights, we conducted the analysis using a synthetic noiseless data vector created with $f_{\text{NL}} = 0$ and $b_0^i = 1$. This was done in two different setups: the default case with 3 tomographic redshift bins, and a simplified case without tomography, which utilized a single redshift bin ranging from $z = 0.8$ to $z = 3.5$. We tested the impact of the optimal weights in these two scenarios, obtaining the constraints reported in table 7 and visualized in figure 15. The results for the bias amplitudes are not reported because, as expected, the optimal weights for f_{NL} do not affect the bias constraints, which are completely consistent to that reported in table 3 for the analysis without optimal weights.

The results of this test are very interesting: in fact, in the case of the tomographic analysis, when using noiseless data the optimal weights are providing a ≈ 5 –10% improvement in the constraining power, similar to what we see in the analysis with real data. On the

Optimal weights	3 redshift bins (default)	Single bin ($0.8 < z < 3.5$)
No	$f_{\text{NL}} = 0^{+28}_{-35}$	$f_{\text{NL}} = 0^{+40}_{-54}$
Yes	$f_{\text{NL}} = 0^{+25}_{-32}$	$f_{\text{NL}} = 0^{+35}_{-42}$

Table 7. Marginalized constraints for f_{NL} when assessing the impact of the tomographic bins and noisy data on the optimal weights effectiveness.




Figure 15. Results of the test using a noiseless data vector to assess the performance of the optimal weights. Left: Analysis in the default configuration, where the dataset is split into three redshift bins. In this case, the optimal weights (pink contours) do not improve the f_{NL} constraint relative to the baseline (blue). Right: Results of the analysis performed without tomography, grouping the full quasar sample into a single bin. Here, the error bars on f_{NL} are $\approx 10\%$ tighter when applying the optimal weights (pink) compared to the unweighted case (blue). In both panels, **regressis** linear weights with $p = 1.6$ were used.

other hand, when we employ only one redshift bin, then using optimal weights tightens the constraints by $\approx 15\text{--}20\%$ in this noiseless data vector test. This suggests that, by binning our data, we are already extracting most of the redshift information that helps in better constraining f_{NL} . Overall, the optimal weighting scheme is straightforward to implement and computationally inexpensive, making it a promising extension of the analysis. However, since the tomographic approach already achieves comparable performance and the current measurements are still limited by systematic uncertainties, the practical gain from applying optimal weights is modest. For this reason, we adopt the results without optimal weighting as our fiducial choice. Looking ahead, the method is expected to become increasingly valuable for future DESI data releases and other Stage-IV surveys, where a $\sim 10\text{--}15\%$ improvement in constraining power could play a significant role in reaching the target precision of $\sigma_{f_{\text{NL}}} < 1$.

Bias	regressis linear weights	catalog linear weights
3 biases	$f_{\text{NL}} = 2_{-34}^{+28} (\chi^2 = 56)$	$f_{\text{NL}} = -14_{-33}^{+27} (\chi^2 = 53)$
1 bias	$f_{\text{NL}} = -3_{-28}^{+25} (\chi^2 = 70)$	$f_{\text{NL}} = -20 \pm 25 (\chi^2 = 66)$

Table 8. Results for the test of fitting a single bias amplitude b_0 for all three bins, compared to the baseline configuration where we allow for each redshift bin to have a free amplitude b_0^i . When fitting a single bias amplitude the reduced χ^2 rises from 1.47 to 1.70, indicative of a worst fit. The results reported here were obtained using $p = 1.6$ and without employing the optimal weights.

7.2 Fitting a single bias amplitude for all three bins

As an alternative to fitting a bias amplitude per tomographic bin, we tried to fit the same amplitude b_0 to all three bins. The resulting marginalized f_{NL} constraints are presented in the second row of table 8. These constraints are compared with the fiducial results. The resulting error bars on f_{NL} are tighter because the analysis has more degrees of freedom (40 instead of 38); however, as indicated by the increase in the reduced χ^2 from 1.4 to 1.7, indicating that this fit is not as good as the fiducial one. This issue arises because the third tomographic bin has been observed to prefer a much lower amplitude ($b_0^3 = 0.68$ compared to $b_0^{1,2} = 1.1$), making it difficult to accurately represent this with a single amplitude fit. In fact the best fit values for b_0 in this scenario are $b_0 = 1.07 \pm 0.05$ and $b_0 = 1.08 \pm 0.05$ for the **regressis** and catalog weights respectively. The decision of having a free amplitude per tomographic bin as our fiducial configuration was driven by this goodness-of-fit test, and for consistency with other similar analyses [28, 53]. Nevertheless, the notably low amplitude in the third bin demonstrates that our baseline bias model does not fully capture the clustering properties of the quasars, thereby motivating the additional tests presented in section 7.3.

7.3 Impact of the bias relation

The baseline analysis reports some inconsistencies emerging with the third tomographic bin: first of all, the lowest multipole $\ell = 6.5$ shows a $> 2\sigma$ deviation from the best fit model which might be a sign of some remaining uncorrected systematic. Moreover, the marginalized posterior for the bias amplitude b_0^3 is not compatible with 1 at $\approx 3\sigma$ level, indicating that the bias we assumed is incorrect for high redshift quasars. Despite this discrepancy emerging in other DESI analysis, it is important to test its impact on the f_{NL} marginalized posterior, the focus of this work. As a first test, we tried to completely remove the third bin from the analysis, performing it with only the first two tomographic bins ($0.8 < z < 2.1$, and $2.1 < z < 2.5$). The resulting marginalized constraint is $f_{\text{NL}} = -15_{-35}^{+29}$. Noticeably, the best fit value shifts by $\approx 0.5\sigma$ compared to the baseline result ($f_{\text{NL}} = 2_{-34}^{+28}$), but the constraining is almost unchanged, suggesting that the third bin could be removed from the analysis without any substantial loss of information. As this was discovered later in the analysis, we do not change our baseline to this case to avoid confirmation bias and to stay consistent with the analysis presented in [53].

The marginalized constraints for the bias amplitudes b_0^i , reported in table 3, are indicative of the fact that the bias evolution we choose does not represent well the bias of our QSO sample. In order to better describe our sample, we tested two alternative bias relations:

Parameter	Constant $b(z)$	Linear Interpolation $b(z)$
f_{NL}	-3^{+28}_{-35}	2^{+28}_{-33}
b_0^1	1.01 ± 0.05	1.02 ± 0.05
b_0^2	1.12 ± 0.12	1.16 ± 0.13
b_0^3	1.07 ± 0.17	1.16 ± 0.18

Table 9. Constraints on the parameters obtained from the analyses performed with different prescriptions for the bias model $b(z)$. In both cases, the reduced χ^2 of the fit is 1.5.

we find the best fit effective bias of each tomographic bin by evaluating our bias model at the effective redshift of each bin (those values are reported in table 1), with b_0^i being the marginalized posteriors (reported in table 3). The resulting values are: $b(z_{\text{eff}}^1) = 2.46$, $b(z_{\text{eff}}^2) = 3.77$, $b(z_{\text{eff}}^3) = 2.79$. Starting from these values, we construct a constant bias model: in each bin, the bias is constant and set to the values just reported. Alternatively, we define a redshift dependent bias model $b(z)$ by linearly interpolating between the three values. In both cases, we still allow for a free bias amplitude in each tomographic bin. Table 9 reports the marginalized posteriors for the 4 free parameters of the analysis. The f_{NL} constraints are completely consistent with the baseline choice for the bias, showing deviations in the best fit values $< 0.1\sigma$ and the same constraining power, showing how f_{NL} constraints are mostly insensitive to the choice of bias evolution. As expected for how the alternative bias models are set up, the bias amplitudes b_0^i in those cases are compatible with 1.

7.4 The impact of the primordial power spectrum

The parameter f_{NL} primarily affects the largest cosmic scales, where the observed power spectrum is most sensitive to primordial fluctuations. Consequently, it can be partially degenerate with the fundamental parameters of the primordial power spectrum, namely the amplitude A_s and the spectral index n_s . Allowing these parameters to vary freely improves the propagation of uncertainties and provides a more robust characterization of the constraints on f_{NL} . In this extended analysis, we introduced A_s and n_s as additional free parameters within our differentiable likelihood framework implemented in **Blast**. This flexibility made the joint sampling of all parameters computationally efficient and straightforward. While flat priors were assumed for the standard cosmological and nuisance parameters, we applied Gaussian priors on A_s and n_s following the *Planck* 2018 results [41]:

$$A_s = (2.101 \pm 0.034) \times 10^{-9}, \quad n_s = 0.9649 \pm 0.0042.$$

The analysis employed the default configuration: **regressis** linear weights with $p = 1.6$, one bias amplitude per tomographic bin, and no application of the optimal weighting scheme. The resulting posterior distributions are shown in figure 16, and the corresponding parameter constraints are reported in table 10. As expected, enlarging the parameter space to include A_s and n_s slightly degrades the precision on f_{NL} , increasing its uncertainty by approximately 10%. The best-fit value, however, remains consistent with that obtained in the fiducial analysis (table 2). The same trend is observed for the bias parameters: their uncertainties broaden, but




Figure 16. Contours for the default analysis, but performed with two extra free parameters: A_s and n_s .

Parameter	68% limits
f_{NL}	1^{+33}_{-36}
A_s	$(2.098 \pm 0.032) \cdot 10^{-9}$
n_s	0.9648 ± 0.0040
b_0^1	1.101 ± 0.074
b_0^2	1.12 ± 0.13
b_0^3	0.66 ± 0.14

Table 10. Constraints on the parameters obtained in our fiducial analysis, but with two extra free parameters: A_s and n_s .

the central values agree well with those reported in table 3. This confirms that the inclusion of primordial parameters does not introduce significant tension, and that the treatment of degeneracies between f_{NL} and the primordial power spectrum is under good control.

8 Conclusions

In this work, we have established constraints on primordial non-Gaussianity by cross-correlating the *Planck* PR4 CMB lensing maps with quasars from DESI DR1 that have been spectroscopically confirmed. Using scales $4 < \ell < 300$, we measure the parameter f_{NL} , which encodes information about local-type primordial non-Gaussianity. We focus on the cross-correlation signal, as it is more robust to systematic errors, whose presence only increases the noise of the measurement. While the quasar autocorrelation does not exhibit significant signs of contamination (figure 7), such as the large-scale excess power previously observed in the DESI Legacy Survey photometric quasar sample [29], for this work we restrict the analysis to $C_{\ell}^{\kappa g}$, with plans to incorporate C_{ℓ}^{gg} in future data releases. To ensure that the systematic weights applied to the quasars do not remove true clustering signal, we tested them extensively on noiseless mocks (section 5).

In our fiducial setup, we use the **regressis** linear weights (section 5), adopt $p = 1.6$, and fit simultaneously for f_{NL} and a free bias amplitude b_0^i for each of the three tomographic redshift bins. This analysis yields $f_{\text{NL}} = 2_{-34}^{+28}$, while for $p = 1.0$, we find $f_{\text{NL}} = 6_{-24}^{+20}$. In comparison to the analysis of photometric quasars from the DESI Legacy Survey [29, 39], our constraints are $\sim 35\%$ tighter. The improvement can be attributed to three key factors: the use of a spectroscopic quasar sample, which is purer and less affected by systematics; the employment of *Planck* PR4 lensing maps that exhibit $\approx 20\%$ lower noise compared to the 2018 release; and the application of a new catalog-based estimator (refer to section 4), which effectively avoids numerical instabilities and pixelization complications [79, 81]. Furthermore, while the quasar auto-correlation C_{ℓ}^{gg} in [29] was significantly contaminated by systematics, we detect no excess large-scale power in the DESI DR1 sample. As a result, we do not need any additional noise modeling in the covariance matrix.

A similar analysis to the one presented here was previously performed using the *Quaia* quasar sample [28] and combining $C_{\ell}^{\kappa g}$ and C_{ℓ}^{gg} to obtain $f_{\text{NL}} = -28_{-24}^{+26}$ for $p = 1.6$ and $f_{\text{NL}} = -20_{-18}^{+19}$ for $p = 1.0$, and $f_{\text{NL}} = -13.8_{-25}^{+26.7}$ when using the cross-correlation alone (and $p = 1$). *Quaia* is a photometric catalog of 1.3 million quasars covering nearly the full sky. Notably, our DESI DR1 spectroscopic sample, which only covers $\sim 17\%$ of the sky and relies exclusively on $C_{\ell}^{\kappa g}$, achieves comparable constraining power. This highlights the impressive constraining power of DESI, which will only grow with upcoming data releases. At the same time, our constraints are weaker than the tightest bounds from LSS, which come from the 3D power spectrum of DESI DR1 QSOs and LRGs, which yield $f_{\text{NL}} = -3 \pm 9$ [10].

In section 6, we derived the optimal weights for $C_{\ell}^{\kappa g}$ to maximize the f_{NL} signal. This optimal weighting scheme improves the constraints by 5–10%, similar but somewhat smaller than the gain typically achieved in three-dimensional power spectrum analyses. Furthermore, the best fits shift by ~ 0.5 – 1σ . Tests with a synthetic data vector and alternative tomographic binning (section 7.1) suggest that most of the additional information from the redshift evolution of the PNG signal is already captured by splitting the sample into tomographic




Figure 17. Constraints on f_{NL} obtained in our fiducial analysis (shaded horizontal band), and adopting alternative analysis strategies that test the robustness of our results. The numerical values shown here are listed in table 2–8 and described in section 7.

bins, which is the baseline configuration of this analysis. Finally, we test the impact of freeing the primordial power spectrum parameters A_s and n_s , finding that our f_{NL} constraints remain completely consistent.

We presented a cosmological analysis of the large-scale clustering of DESI DR1 spectroscopically confirmed quasars, aimed at constraining the f_{NL} parameter, while a full cosmological interpretation of the sample is provided in the companion paper [53]. Our constraints on PNG, which proved to be very stable against many analysis variations (see figure 17), are statistically weaker than those from three-dimensional power-spectrum analyses, as the measurement remains noise-dominated on the large scales that carry most of the information. However, this situation will improve significantly with future DESI data releases, which will offer greater statistical power, cover a larger sky area, and access additional scales where the f_{NL} signal is strongest. The inclusion of the quasar auto-correlation will further enhance constraining power, provided that imaging systematics are well understood and under control.

Acknowledgments

The authors are grateful to Emanuele Castorina for his insightful feedback and for providing the derivation that helped refine the optimal weighting scheme presented in this work. WP acknowledges support from the Natural Sciences and Engineering Research Council of Canada (NSERC), [funding reference number RGPIN-2025-03931] and from the Canadian Space Agency. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Colleges and Universities. This research was enabled in part by support provided by Compute Ontario (computeontario.ca) and the Digital Research Alliance of Canada (alliancecan.ca).

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the

U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF’s National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Humanities, Science and Technology of Mexico (CONAHCYT); the Ministry of Science, Innovation and Universities of Spain (MICIU/AEI/10.13039/501100011033), and by the DESI Member Institutions: <https://www.desi.lbl.gov/collaborating-institutions>.

The DESI Legacy Imaging Surveys consist of three individual and complementary projects: the Dark Energy Camera Legacy Survey (DECaLS), the Beijing-Arizona Sky Survey (BASS), and the Mayall z-band Legacy Survey (MzLS). DECaLS, BASS and MzLS together include data obtained, respectively, at the Blanco telescope, Cerro Tololo Inter-American Observatory, NSF’s NOIRLab; the Bok telescope, Steward Observatory, University of Arizona; and the Mayall telescope, Kitt Peak National Observatory, NOIRLab. NOIRLab is operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation. Pipeline processing and analyses of the data were supported by NOIRLab and the Lawrence Berkeley National Laboratory. Legacy Surveys also uses data products from the Near-Earth Object Wide-field Infrared Survey Explorer (NEOWISE), a project of the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration. Legacy Surveys was supported by: the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy; the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility; the U.S. National Science Foundation, Division of Astronomical Sciences; the National Astronomical Observatories of China, the Chinese Academy of Sciences and the Chinese National Natural Science Foundation. LBNL is managed by the Regents of the University of California under contract to the U.S. Department of Energy. The complete acknowledgments can be found at <https://www.legacysurvey.org/>. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies. The authors are honored to be permitted to conduct scientific research on I’oligam Du’ag (Kitt Peak), a mountain with particular significance to the Tohono O’odham Nation.

A Derivation of the optimal f_{NL} weights for $C_\ell^{\kappa g}$

In this appendix, we derive the optimal weighting scheme for the cross-correlation between a galaxy field and a matter-like tracer, such as the CMB lensing convergence κ , by generalizing the original derivation in [94].

We begin by considering the real-space galaxy overdensity field, $\delta_g(\mathbf{x}_1)$, and the real-space matter overdensity field, $\delta_m(\mathbf{x}_2)$. Assuming that the covariance matrices are diagonal in pixel space, we can write their inverses as $C_g^{-1} = 1/\sigma_g^2$ and $C_m^{-1} = 1/\sigma_m^2$ for the galaxy and matter fields, respectively (see the discussion around eq. 2.8 in [94]). Then, the optimal

quadratic estimator q for f_{NL} is:

$$q = \frac{1}{2} \frac{\delta_g(\mathbf{x}_1)}{\sigma_g^2} \left(\frac{\partial}{\partial f_{\text{NL}}} \langle \delta_g(\mathbf{x}_1) \delta_m(\mathbf{x}_2) \rangle \Big|_{f_{\text{NL}}=0} \right) \frac{\delta_m(\mathbf{x}_2)}{\sigma_m^2}. \quad (\text{A.1})$$

The derivative of the cross-correlation signal with respect to f_{NL} is given by:

$$\left(\frac{\partial \langle \delta_g(\mathbf{x}_1) \delta_m(\mathbf{x}_2) \rangle}{\partial f_{\text{NL}}} \Big|_{f_{\text{NL}}=0} \right) = \int \frac{d^3k}{(2\pi)^3} \frac{D(z_1)D(z_2)}{D(z_1)} (b(z_1) - p) \tilde{\alpha}(k) P_m(k, z_0) e^{i\mathbf{k} \cdot (\mathbf{x}_2 - \mathbf{x}_1)}, \quad (\text{A.2})$$

where $P_m(k, z_0)$ is the matter power spectrum at a reference redshift z_0 , and the transfer function $\tilde{\alpha}(k)$ is defined as:

$$\tilde{\alpha}(k) \equiv \frac{3\Omega_m H_0^2 \delta_c}{c^2 k^2 T(k)}. \quad (\text{A.3})$$

We can sum over all positions in the sky to obtain:

$$q = \frac{1}{2} \int \frac{d^3k}{(2\pi)^3} \tilde{\alpha}(k) P_m(k, z_0) \int d^3x_1 e^{-i\mathbf{k} \cdot \mathbf{x}_1} \frac{\delta_g(\mathbf{x}_1)}{\sigma_g^2} (b(z_1) - p) \int d^3x_2 e^{i\mathbf{k} \cdot \mathbf{x}_2} \frac{\delta_m(\mathbf{x}_2)}{\sigma_m^2} D(z_2). \quad (\text{A.4})$$

At this stage, we could insert the kernels W_g and W_κ . We tested the impact of adding those kernels to the weights and they did not impact the constraining power, so we decided to not include them in the final weighting scheme. Finally, one can apply the plane-wave expansion, expand the Legendre polynomials in spherical harmonics and integrate over $d\Omega_k/4\pi$ to obtain:

$$q = \frac{4\pi}{2} \int \frac{dk k^2}{2\pi^2} \tilde{\alpha}(k) P_m(k, z_0) \sum_{\ell, m} \left[\int d\Omega_1 \frac{\delta_g(\mathbf{x}_1)}{\sigma_g^2} Y_{\ell m}(\hat{\mathbf{x}}_1) \right] \times \left[\int d\Omega_2 \frac{\delta_m(\mathbf{x}_2)}{\sigma_m^2} Y_{\ell m}^*(\hat{\mathbf{x}}_2) \right]. \quad (\text{A.5})$$

This expression demonstrates that the optimal measurement of f_{NL} is obtained by applying a redshift-dependent weight of $w_g(z) \propto (b(z) - p)$ to the galaxy field and a weight of $w_m(z) \propto D(z)$ to the matter field, results that are consistent with [94, 95].

In our specific analysis, the matter field is provided by the CMB lensing convergence κ , which is a projected field integrated over a broad range of redshifts. Because we lack 3D information for the CMB lensing field κ , we cannot apply the optimal $D(z)$ weight to it. In conclusion, the resulting weights used in this work for $C_\ell^{\kappa g}$ are:

$$w_{\text{opt}}(z) = (b(z) - p) w_{\text{FKP}}, \quad (\text{A.6})$$

where w_{FKP} is the standard inverse-noise weighting, which is obtained from the inverse galaxy covariance C_g^{-1} [94].

B Restricting to quasars with $z < 3.1$

The best f_{NL} measurement from LSS available to date was performed in [10]. In that analysis, only the DESI quasars in the range $0.8 < z < 3.1$ were considered. This excludes $\approx 32,500$ quasars from the sample. That choice was motivated by the fact that the covariance matrix




Figure 18. Measured angular cross-power spectrum C_ℓ^{kg} in the third tomographic bin, changing the maximum redshift. *Left:* measurement performed with the default weighting scheme, the **regressis** linear weights. *Right:* Same measurement, but the systematic weights are the linear weights implemented in the catalog. See section 5 for a more detailed explanation of the weights.

p	regressis linear weights	catalog linear weights
1.6	$f_{\text{NL}} = -10^{+29}_{-35} \ (\chi^2 = 48)$	$f_{\text{NL}} = -23^{+28}_{-34} \ (\chi^2 = 48)$
1	$f_{\text{NL}} = -1^{+20}_{-24} \ (\chi^2 = 48)$	$f_{\text{NL}} = -9^{+20}_{-23} \ (\chi^2 = 48)$

Table 11. Results for the analysis performed with a value of $z_{\text{max}} = 3.1$ instead of $z_{\text{max}} = 3.5$.

is estimated from simulations, the **EZMocks**, which have a simulation box that is sufficient to emulate the quasars up to $z_{\text{max}} = 3.1$ without repeating the box. We test the impact of this choice to our analysis: we keep the default configuration of 3 tomographic bins, fit a bias amplitude for each bin and we do not apply optimal weights. This analysis change affects the measurement, theory curve and therefore the covariance matrix, which is re-computed and tuned following the procedure outlined in section 4. As in the previous cases, results are presented for variation of systematic weights choice and values of the parameter p in table 11. Removing the higher redshift part of the sample does not affect the error bars, but has an effect on the best fit values, which move around of $\approx 1/3 \sigma$ in the fiducial scenario. It is interesting to look at the effect on the bias amplitudes: the first two bins do not show a relevant shift, with the best fit values being completely compatible with the ones showed in table 3. However, the best fit amplitude in the third bin is $b_3 = 0.62 \pm 0.13$, a value shift of $\approx 0.5 \sigma$. Finally, the χ^2 in this case is lower than in the default scenario (48 vs 56). The difference is driven by the first multipole ($\ell = 6.5$) in the third tomographic bin, as shown in figure 18. Excluding the higher redshift part of the quasar sample causes the measurement to shift by $\approx 1 \sigma$, with a significant impact of the χ^2 value. The reasons for this shift are unknown and currently under investigation. From the figure, it is also possible to see why a lower bias is preferred, in fact, the measurement in the last multipole ($\ell = 275$) is lower for both weighting schemes, explaining the preference for a lower bias value.

Data Availability Statement. This article has associated data in a data repository. The data used in this work are publicly available as part of DESI Data Release 1 (see <https://data.desi.lbl.gov/doc/releases/>). The data points corresponding to the figures, as well as the chains required to reproduce the corner plots, are available on Zenodo at <https://zenodo.org/records/19471607>. All codes and packages employed in this analysis are publicly accessible and are referenced throughout the paper.

References

- [1] J.M. Bardeen, P.J. Steinhardt and M.S. Turner, *Spontaneous Creation of Almost Scale-Free Density Perturbations in an Inflationary Universe*, *Phys. Rev. D* **28** (1983) 679 [INSPIRE].
- [2] T. Falk, R. Rangarajan and M. Srednicki, *Dependence of density perturbations on the coupling constant in a simple model of inflation*, *Phys. Rev. D* **46** (1992) 4232 [astro-ph/9208002] [INSPIRE].
- [3] V. Acquaviva, N. Bartolo, S. Matarrese and A. Riotto, *Second order cosmological perturbations from inflation*, *Nucl. Phys. B* **667** (2003) 119 [astro-ph/0209156] [INSPIRE].
- [4] PLANCK collaboration, *Planck intermediate results. LVII. Joint Planck LFI and HFI data processing*, *Astron. Astrophys.* **643** (2020) A42 [arXiv:2007.04997] [INSPIRE].
- [5] J.M. Maldacena, *Non-Gaussian features of primordial fluctuations in single field inflationary models*, *JHEP* **05** (2003) 013 [astro-ph/0210603] [INSPIRE].
- [6] N. Bartolo, E. Komatsu, S. Matarrese and A. Riotto, *Non-Gaussianity from inflation: Theory and observations*, *Phys. Rept.* **402** (2004) 103 [astro-ph/0406398] [INSPIRE].
- [7] PLANCK collaboration, *Planck 2018 results. IX. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **641** (2020) A9 [arXiv:1905.05697] [INSPIRE].
- [8] CMB-S4 collaboration, *CMB-S4 Science Book, First Edition*, arXiv:1610.02743 [DOI:10.2172/1352047] [INSPIRE].
- [9] N. Dalal, O. Doré, D. Huterer and A. Shirokov, *The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects*, *Phys. Rev. D* **77** (2008) 123514 [arXiv:0710.4560] [INSPIRE].
- [10] E. Chaussidon et al., *Constraining primordial non-Gaussianity with DESI 2024 LRG and QSO samples*, *JCAP* **06** (2025) 029 [arXiv:2411.17623] [INSPIRE].
- [11] A. Slosar et al., *Constraints on local primordial non-Gaussianity from large scale structure*, *JCAP* **08** (2008) 031 [arXiv:0805.3580] [INSPIRE].
- [12] J.-Q. Xia et al., *Constraints on Primordial Non-Gaussianity from Large Scale Structure Probes*, *JCAP* **08** (2011) 033 [arXiv:1104.5015] [INSPIRE].
- [13] N. Nikoloudakis, T. Shanks and U. Sawangwit, *Clustering analysis of high-redshift Luminous Red Galaxies in Stripe 82*, *Mon. Not. Roy. Astron. Soc.* **429** (2013) 2032 [arXiv:1204.3609] [INSPIRE].
- [14] A.R. Pullen and C.M. Hirata, *Systematic effects in large-scale angular power spectra of photometric quasars and implications for constraining primordial nongaussianity*, *Publ. Astron. Soc. Pac.* **125** (2013) 705 [arXiv:1212.4500] [INSPIRE].
- [15] B. Leistedt and H.V. Peiris, *Exploiting the full potential of photometric quasar surveys: Optimal power spectra through blind mitigation of systematics*, *Mon. Not. Roy. Astron. Soc.* **444** (2014) 2 [arXiv:1404.6530] [INSPIRE].

- [16] T. Giannantonio et al., *Improved Primordial Non-Gaussianity Constraints from Measurements of Galaxy Clustering and the Integrated Sachs-Wolfe Effect*, *Phys. Rev. D* **89** (2014) 023511 [[arXiv:1303.1349](#)] [[INSPIRE](#)].
- [17] B. Leistedt, H.V. Peiris and N. Roth, *Constraints on Primordial Non-Gaussianity from 800 000 Photometric Quasars*, *Phys. Rev. Lett.* **113** (2014) 221301 [[arXiv:1405.4315](#)] [[INSPIRE](#)].
- [18] S. Ho et al., *Sloan Digital Sky Survey III photometric quasar clustering: probing the initial conditions of the Universe*, *JCAP* **05** (2015) 040 [[arXiv:1311.2597](#)] [[INSPIRE](#)].
- [19] M. Rezaie et al., *Local primordial non-Gaussianity from the large-scale clustering of photometric DESI luminous red galaxies*, *Mon. Not. Roy. Astron. Soc.* **532** (2024) 1902 [[arXiv:2307.01753](#)] [[INSPIRE](#)].
- [20] K.M. Smith, O. Zahn and O. Doré, *Detection of Gravitational Lensing in the Cosmic Microwave Background*, *Phys. Rev. D* **76** (2007) 043510 [[arXiv:0705.3980](#)] [[INSPIRE](#)].
- [21] C.M. Hirata et al., *Correlation of CMB with large-scale structure: II. Weak lensing*, *Phys. Rev. D* **78** (2008) 043520 [[arXiv:0801.0644](#)] [[INSPIRE](#)].
- [22] T.-C. Chang, U.-L. Pen, K. Bandura and J.B. Peterson, *Hydrogen 21-cm Intensity Mapping at redshift 0.8*, *Nature* **466** (2010) 463 [[arXiv:1007.3709](#)] [[INSPIRE](#)].
- [23] J. Rhodes et al., *Exploiting Cross Correlations and Joint Analyses*, *Astropart. Phys.* **63** (2015) 42 [[arXiv:1309.5388](#)] [[INSPIRE](#)].
- [24] T. Giannantonio and W.J. Percival, *Using correlations between CMB lensing and large-scale structure to measure primordial non-Gaussianity*, *Mon. Not. Roy. Astron. Soc.* **441** (2014) L16 [[arXiv:1312.5154](#)] [[INSPIRE](#)].
- [25] A. Baleato Lizancos et al., *A halo model of extragalactic contamination to CMB lensing, delensing, and cross-correlations*, *JCAP* **11** (2025) 031 [[arXiv:2507.03859](#)] [[INSPIRE](#)].
- [26] G. Piccirilli et al., *Growth history and quasar bias evolution at $z < 3$ from Quaia*, *JCAP* **06** (2024) 012 [[arXiv:2402.05761](#)] [[INSPIRE](#)].
- [27] J.R. Bermejo-Climent et al., *Constraints on primordial non-Gaussianity from the cross-correlation of DESI luminous red galaxies and Planck CMB lensing*, *Astron. Astrophys.* **698** (2025) A177 [[arXiv:2412.10279](#)] [[INSPIRE](#)].
- [28] G. Fabbian, D. Alonso, K. Storey-Fisher and T. Cornish, *Constraints on primordial non-Gaussianity from Quaia*, *JCAP* **02** (2026) 056 [[arXiv:2504.20992](#)] [[INSPIRE](#)].
- [29] DESI collaboration, *Constraining primordial non-Gaussianity from DESI quasar targets and Planck CMB lensing*, *JCAP* **03** (2024) 021 [[arXiv:2305.07650](#)] [[INSPIRE](#)].
- [30] J. Carron, M. Mirmelstein and A. Lewis, *CMB lensing from Planck PR4 maps*, *JCAP* **09** (2022) 039 [[arXiv:2206.07773](#)] [[INSPIRE](#)].
- [31] U. Seljak, *Extracting primordial non-gaussianity without cosmic variance*, *Phys. Rev. Lett.* **102** (2009) 021302 [[arXiv:0807.1770](#)] [[INSPIRE](#)].
- [32] SPHEREx collaboration, *Cosmology with the SPHEREx All-Sky Spectral Survey*, [[arXiv:1412.4872](#)] [[INSPIRE](#)].
- [33] D. Yamauchi, K. Takahashi and M. Oguri, *Constraining primordial non-Gaussianity via a multitracers technique with surveys by Euclid and the Square Kilometre Array*, *Phys. Rev. D* **90** (2014) 083520 [[arXiv:1407.5453](#)] [[INSPIRE](#)].

- [34] D. Karagiannis et al., *Constraining primordial non-Gaussianity with bispectrum and power spectrum from upcoming optical and radio surveys*, *Mon. Not. Roy. Astron. Soc.* **478** (2018) 1341 [[arXiv:1801.09280](#)] [[INSPIRE](#)].
- [35] M. Schmittfull and U. Seljak, *Parameter constraints from cross-correlation of CMB lensing with galaxy clustering*, *Phys. Rev. D* **97** (2018) 123540 [[arXiv:1710.09465](#)] [[INSPIRE](#)].
- [36] S. Ferraro et al., *Inflation and Dark Energy from Spectroscopy at $z > 2$* , *Bull. Am. Astron. Soc.* **51** (2019) 72 [[arXiv:1903.09208](#)] [[INSPIRE](#)].
- [37] D. Gualdi, H. Gil-Marín and L. Verde, *Joint analysis of anisotropic power spectrum, bispectrum and trispectrum: application to N-body simulations*, *JCAP* **07** (2021) 008 [[arXiv:2104.03976](#)] [[INSPIRE](#)].
- [38] DESI collaboration, *A Spectroscopic Road Map for Cosmic Frontier: DESI, DESI-II, Stage-5*, [arXiv:2209.03585](#) [[DOI:10.2172/1887254](#)] [[INSPIRE](#)].
- [39] DESI collaboration, *Overview of the DESI Legacy Imaging Surveys*, *Astron. J.* **157** (2019) 168 [[arXiv:1804.08657](#)] [[INSPIRE](#)].
- [40] PLANCK collaboration, *Planck 2018 results. VIII. Gravitational lensing*, *Astron. Astrophys.* **641** (2020) A8 [[arXiv:1807.06210](#)] [[INSPIRE](#)].
- [41] PLANCK collaboration, *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, *Astron. Astrophys.* **641** (2020) A1 [[arXiv:1807.06205](#)] [[INSPIRE](#)].
- [42] E. Komatsu and D.N. Spergel, *Acoustic signatures in the primary microwave background bispectrum*, *Phys. Rev. D* **63** (2001) 063002 [[astro-ph/0005036](#)] [[INSPIRE](#)].
- [43] V. Desjacques and U. Seljak, *Primordial non-Gaussianity from the large scale structure*, *Class. Quant. Grav.* **27** (2010) 124011 [[arXiv:1003.5020](#)] [[INSPIRE](#)].
- [44] A. Barreira, *Can we actually constrain f_{NL} using the scale-dependent bias effect? An illustration of the impact of galaxy bias uncertainties using the BOSS DR12 galaxy power spectrum*, *JCAP* **11** (2022) 013 [[arXiv:2205.05673](#)] [[INSPIRE](#)].
- [45] A. Barreira, *On the impact of galaxy bias uncertainties on primordial non-Gaussianity constraints*, *JCAP* **12** (2020) 031 [[arXiv:2009.06622](#)] [[INSPIRE](#)].
- [46] M. Biagetti et al., *Verifying the consistency relation for the scale-dependent bias from local primordial non-Gaussianity*, *Mon. Not. Roy. Astron. Soc.* **468** (2017) 3277 [[arXiv:1611.04901](#)] [[INSPIRE](#)].
- [47] J.M. Sullivan, T. Prijon and U. Seljak, *Learning to concentrate: multi-tracer forecasts on local primordial non-Gaussianity with machine-learned bias*, *JCAP* **08** (2023) 004 [[arXiv:2303.08901](#)] [[INSPIRE](#)].
- [48] P. Breiding et al., *Powerful Radio-Loud Quasars are Triggered by Galaxy Mergers in the Cosmic Bright Ages*, [arXiv:2305.11804](#) [[DOI:10.3847/1538-4357/ad19db](#)].
- [49] D.N. Limber, *The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field*, *Astrophys. J.* **117** (1953) 134.
- [50] S. Chiarenza, M. Bonici, W.J. Percival and M. White, *BLAST: Beyond Limber Angular power Spectra Toolkit. A fast and efficient algorithm for 3×2 pt analysis*, [arXiv:2410.03632](#) [[DOI:10.33232/001c.127038](#)] [[INSPIRE](#)].
- [51] S. Chiarenza et al., *Evolving Blast.jl into a fast and differentiable toolkit for full-physics 6×2 pt cosmology*, in preparation (2026).

- [52] SDSS collaboration, *Detection of cosmic magnification with the Sloan Digital Sky Survey*, *Astrophys. J.* **633** (2005) 589 [[astro-ph/0504510](#)] [[INSPIRE](#)].
- [53] R. de Belsunce et al., *Cosmology from Planck CMB lensing and DESI DR1 quasar tomography*, *JCAP* **10** (2025) 077 [[arXiv:2506.22416](#)] [[INSPIRE](#)].
- [54] E. Chaussidon et al., *Target Selection and Validation of DESI Quasars*, *Astrophys. J.* **944** (2023) 107 [[arXiv:2208.08511](#)] [[INSPIRE](#)].
- [55] DESI collaboration, *Data Release 1 of the Dark Energy Spectroscopic Instrument*, *Astron. J.* **171** (2026) 285 [[arXiv:2503.14745](#)] [[INSPIRE](#)].
- [56] DESI collaboration, *The DESI Experiment Part II: Instrument Design*, [arXiv:1611.00037](#) [[INSPIRE](#)].
- [57] DESI collaboration, *Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **164** (2022) 207 [[arXiv:2205.10939](#)] [[INSPIRE](#)].
- [58] DESI collaboration, *The DESI Experiment Part I: Science, Targeting, and Survey Design*, [arXiv:1611.00036](#) [[INSPIRE](#)].
- [59] DESI collaboration, *The Robotic Multiobject Focal Plane System of the Dark Energy Spectroscopic Instrument (DESI)*, *Astron. J.* **165** (2023) 9 [[arXiv:2205.09014](#)] [[INSPIRE](#)].
- [60] DESI collaboration, *The Optical Corrector for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **168** (2024) 95 [[arXiv:2306.06310](#)] [[INSPIRE](#)].
- [61] C. Poppett et al., *Overview of the Fiber System for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **168** (2024) 245.
- [62] DESI collaboration, *Survey Operations for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **166** (2023) 259 [[arXiv:2306.06309](#)] [[INSPIRE](#)].
- [63] DESI collaboration, *The DESI Experiment, a whitepaper for Snowmass 2013*, [arXiv:1308.0847](#) [[INSPIRE](#)].
- [64] DESI collaboration, *Validation of the Scientific Program for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **167** (2024) 62 [[arXiv:2306.06307](#)] [[INSPIRE](#)].
- [65] DESI collaboration, *The Early Data Release of the Dark Energy Spectroscopic Instrument*, *Astron. J.* **168** (2024) 58 [[arXiv:2306.06308](#)] [[INSPIRE](#)].
- [66] DESI collaboration, *DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations*, *JCAP* **02** (2025) 021 [[arXiv:2404.03002](#)] [[INSPIRE](#)].
- [67] DESI collaboration, *DESI 2024 IV: Baryon Acoustic Oscillations from the Lyman alpha forest*, *JCAP* **01** (2025) 124 [[arXiv:2404.03001](#)] [[INSPIRE](#)].
- [68] DESI collaboration, *DESI 2024 III: baryon acoustic oscillations from galaxies and quasars*, *JCAP* **04** (2025) 012 [[arXiv:2404.03000](#)] [[INSPIRE](#)].
- [69] DESI collaboration, *DESI 2024 VII: cosmological constraints from the full-shape modeling of clustering measurements*, *JCAP* **07** (2025) 028 [[arXiv:2411.12022](#)] [[INSPIRE](#)].
- [70] DESI collaboration, *DESI 2024 II: sample definitions, characteristics, and two-point clustering statistics*, *JCAP* **07** (2025) 017 [[arXiv:2411.12020](#)] [[INSPIRE](#)].
- [71] A.J. Ross et al., *The construction of large-scale structure catalogs for the Dark Energy Spectroscopic Instrument*, *JCAP* **01** (2025) 125 [[arXiv:2405.16593](#)] [[INSPIRE](#)].
- [72] DESI collaboration, *The Spectroscopic Data Processing Pipeline for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **165** (2023) 144 [[arXiv:2209.14482](#)] [[INSPIRE](#)].

- [73] S. Bailey et al., *Redrock: Spectroscopic Classification and Redshift Fitting for the Dark Energy Spectroscopic Instrument*, in preparation (2026).
- [74] DESI collaboration, *Performance of the Quasar Spectral Templates for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **166** (2023) 66 [[arXiv:2305.10426](#)] [[INSPIRE](#)].
- [75] A. Anand et al., *Archetype-based Redshift Estimation for the Dark Energy Spectroscopic Instrument Survey*, *Astron. J.* **168** (2024) 124 [[arXiv:2405.19288](#)] [[INSPIRE](#)].
- [76] A. Lewis and A. Challinor, *Weak gravitational lensing of the CMB*, *Phys. Rept.* **429** (2006) 1 [[astro-ph/0601594](#)] [[INSPIRE](#)].
- [77] K.M. Górski et al., *HEALPix: A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](#)] [[INSPIRE](#)].
- [78] A.S. Maniyar et al., *Quadratic estimators for CMB weak lensing*, *Phys. Rev. D* **103** (2021) 083524 [[arXiv:2101.12193](#)] [[INSPIRE](#)].
- [79] K. Wolz, D. Alonso and A. Nicola, *Catalog-based pseudo- C_ℓ 's*, *JCAP* **01** (2025) 028 [[arXiv:2407.21013](#)] [[INSPIRE](#)].
- [80] LSST DARK ENERGY SCIENCE collaboration, *A unified pseudo- C_ℓ framework*, *Mon. Not. Roy. Astron. Soc.* **484** (2019) 4127 [[arXiv:1809.09603](#)] [[INSPIRE](#)].
- [81] A. Baleato Lizancos and M. White, *Harmonic analysis of discrete tracers of large-scale structure*, *JCAP* **05** (2024) 010 [[arXiv:2312.12285](#)] [[INSPIRE](#)].
- [82] M. Reinecke, S. Belkner and J. Carron, *Improved cosmic microwave background (de-)lensing using general spherical harmonic transforms*, *Astron. Astrophys.* **678** (2023) A165 [[arXiv:2304.10431](#)] [[INSPIRE](#)].
- [83] N. Tessore and A. Hall, *Shot noise in clustering power spectra*, [arXiv:2507.03749](#) [[DOI:10.33232/001c.145919](#)] [[INSPIRE](#)].
- [84] M. Maus et al., *A joint analysis of 3D clustering and galaxy $\times$ CMB-lensing cross-correlations with DESI DR1 galaxies*, *JCAP* **11** (2025) 077 [[arXiv:2505.20656](#)] [[INSPIRE](#)].
- [85] ACT collaboration, *The Atacama Cosmology Telescope: Cosmology from Cross-correlations of unWISE Galaxies and ACT DR6 CMB Lensing*, *Astrophys. J.* **966** (2024) 157 [[arXiv:2309.05659](#)] [[INSPIRE](#)].
- [86] A. Benoit-Lévy et al., *Full-sky CMB lensing reconstruction in presence of sky-cuts*, *Astron. Astrophys.* **555** (2013) A37 [[arXiv:1301.4145](#)] [[INSPIRE](#)].
- [87] J. Carron, *Real-world CMB lensing quadratic estimator power spectrum response*, *JCAP* **02** (2023) 057 [[arXiv:2210.05449](#)] [[INSPIRE](#)].
- [88] eBOSS collaboration, *The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Large-scale structure catalogues for cosmological analysis*, *Mon. Not. Roy. Astron. Soc.* **498** (2020) 2354 [[arXiv:2007.09000](#)] [[INSPIRE](#)].
- [89] M. Rezaie, H.-J. Seo, A.J. Ross and R.C. Bunescu, *Improving Galaxy Clustering Measurements with Deep Learning: analysis of the DECaLS DR7 data*, *Mon. Not. Roy. Astron. Soc.* **495** (2020) 1613 [[arXiv:1907.11355](#)] [[INSPIRE](#)].
- [90] E. Chaussidon et al., *Angular clustering properties of the DESI QSO target selection using DR9 Legacy Imaging Surveys*, *Mon. Not. Roy. Astron. Soc.* **509** (2021) 3904 [[arXiv:2108.03640](#)] [[INSPIRE](#)].
- [91] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in Python*, *J. Open Source Softw.* **4** (2019) 1298 [[INSPIRE](#)].

- [92] N. Tessore et al., *GLASS: Generator for Large Scale Structure*, [arXiv:2302.01942](#) [[DOI:10.21105/astro.2302.01942](#)] [[INSPIRE](#)].
- [93] D. Bianchi et al., *Characterization of DESI fiber assignment incompleteness effect on 2-point clustering and mitigation methods for DR1 analysis*, *JCAP* **04** (2025) 074 [[arXiv:2411.12025](#)] [[INSPIRE](#)].
- [94] eBOSS collaboration, *Redshift-weighted constraints on primordial non-Gaussianity from the clustering of the eBOSS DR14 quasars in Fourier space*, *JCAP* **09** (2019) 010 [[arXiv:1904.08859](#)] [[INSPIRE](#)].
- [95] E.-M. Mueller, W.J. Percival and R. Ruggeri, *Optimizing primordial non-Gaussianity measurements from galaxy surveys*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 4160 [[arXiv:1702.05088](#)] [[INSPIRE](#)].
- [96] E.-M. Mueller et al., *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: constraining modified gravity*, *Mon. Not. Roy. Astron. Soc.* **475** (2018) 2122 [[arXiv:1612.00812](#)] [[INSPIRE](#)].
- [97] M.S. Cagliari, E. Castorina, M. Bonici and D. Bianchi, *Optimal constraints on Primordial non-Gaussianity with the eBOSS DR16 quasars in Fourier space*, *JCAP* **08** (2024) 036 [[arXiv:2309.15814](#)] [[INSPIRE](#)].
- [98] H.A. Feldman, N. Kaiser and J.A. Peacock, *Power spectrum analysis of three-dimensional redshift surveys*, *Astrophys. J.* **426** (1994) 23 [[astro-ph/9304022](#)] [[INSPIRE](#)].
- [99] A. Lewis, A. Challinor and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, *Astrophys. J.* **538** (2000) 473 [[astro-ph/9911177](#)] [[INSPIRE](#)].
- [100] T.E. Fjelde et al., *Turing.jl: A General-Purpose Probabilistic Programming Language*, *ACM Trans. Probab. Mach. Learn.* **1** (2025) 1.
- [101] M.D. Hoffman and A. Gelman, *The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo*, *J. Mach. Learn. Res.* **15** (2014) 1593 [[arXiv:1111.4246](#)] [[INSPIRE](#)].
- [102] eBOSS collaboration, *Primordial non-Gaussianity from the completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey II: measurements in Fourier space with optimal weights*, *Mon. Not. Roy. Astron. Soc.* **514** (2022) 3396 [[arXiv:2106.13725](#)] [[INSPIRE](#)].
- [103] <https://data.desi.lbl.gov/doc/releases/>.
- [104] <https://zenodo.org/records/19471607>.

Author List

S. Chiarenza^{1,2,3,*}, A. Krolewski^{1,2,3}, M. Bonici^{1,2,3}, E. Chaussidon⁴, R. de Belsunce⁴,
W.J. Percival^{1,2,3}, J. Aguilar⁴, S. Ahlen⁶, A. Baleato Lizancos^{4,5}, D. Bianchi^{7,8}, D. Brooks⁹,
T. Claybaugh⁴, A. Cuceu⁴, K.S. Dawson¹⁰, A. de la Macorra¹¹, P. Doel⁹, S. Ferraro^{4,5,12},
A. Font-Ribera¹³, J.E. Forero-Romero^{14,15}, E. Gaztañaga^{16,17,18}, S. Gontcho A Gontcho^{4,19},
G. Gutierrez²⁰, H.K. Herrera-Alcantar^{21,23}, K. Honscheid^{23,24}, D. Huterer²⁵, M. Ishak²⁶,
R. Joyce²⁷, D. Kirkby²⁸, A. Kremin⁴, O. Lahav⁹, C. Lamman²⁴, M. Landriau⁴,
L. Le Guillou²⁹, M.E. Levi⁴, M. Manera^{30,13}, P. Martini^{23,24,31}, A. Meisner²⁷,
R. Miquel^{13,32}, S. Nadathur¹⁷, J.A. Newman³³, G. Niz^{34,35}, N. Palanque-Delabrouille^{22,4},
C. Poppett^{4,36,5,12}, F. Prada³⁷, I. Pérez-Ràfols³⁸, G. Rossi³⁹, E. Sanchez⁴⁰, D. Schlegel⁴,
M. Schubnell²⁵, H. Seo⁴¹, J. Silber⁴, D. Sprayberry²⁷, G. Tarlé²⁵, B.A. Weaver²⁷, C. Yèche²²,
R. Zhou⁴, H. Zou⁴²

¹ Department of Physics and Astronomy, University of Waterloo,
200 University Ave W, Waterloo, ON N2L 3G1, Canada

² Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada

³ Waterloo Centre for Astrophysics, University of Waterloo,
200 University Ave W, Waterloo, ON N2L 3G1, Canada

⁴ Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, U.S.A.

⁵ Department of Physics, University of California, Berkeley,
366 LeConte Hall MC 7300, Berkeley, CA 94720-7300, U.S.A.

⁶ Department of Physics, Boston University, 590 Commonwealth Avenue, Boston, MA 02215, U.S.A.

⁷ Dipartimento di Fisica “Aldo Pontremoli”, Università degli Studi di Milano,
Via Celoria 16, I-20133 Milano, Italy

⁸ INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy

⁹ Department of Physics & Astronomy, University College London,
Gower Street, London, WC1E 6BT, U.K.

¹⁰ Department of Physics and Astronomy, The University of Utah,
115 South 1400 East, Salt Lake City, UT 84112, U.S.A.

¹¹ Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica,
Ciudad Universitaria, Cd. de México C. P. 04510, México

¹² University of California, Berkeley, 110 Sproul Hall #5800 Berkeley, CA 94720, U.S.A.

¹³ Institut de Física d’Altes Energies (IFAE), The Barcelona Institute of Science and Technology,
Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain

¹⁴ Departamento de Física, Universidad de los Andes,
Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia

¹⁵ Observatorio Astronómico, Universidad de los Andes,
Cra. 1 No. 18A-10, Edificio H, CP 111711 Bogotá, Colombia

¹⁶ Institut d’Estudis Espacials de Catalunya (IEEC),
c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain

¹⁷ Institute of Cosmology and Gravitation, University of Portsmouth,
Dennis Sciamia Building, Portsmouth, PO1 3FX, U.K.

¹⁸ Institute of Space Sciences, ICE-CSIC, Campus UAB,
Carrer de Can Magrans s/n, 08913 Bellaterra, Barcelona, Spain

¹⁹ University of Virginia, Department of Astronomy, Charlottesville, VA 22904, U.S.A.

²⁰ Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, U.S.A.

²¹ Institut d’Astrophysique de Paris. 98 bis boulevard Arago. 75014 Paris, France

²² IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

²³ Center for Cosmology and AstroParticle Physics, The Ohio State University,
191 West Woodruff Avenue, Columbus, OH 43210, U.S.A.

- ²⁴ *Department of Physics, The Ohio State University,
191 West Woodruff Avenue, Columbus, OH 43210, U.S.A.*
- ²⁵ *Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, U.S.A.*
- ²⁶ *Department of Physics, The University of Texas at Dallas,
800 W. Campbell Rd., Richardson, TX 75080, U.S.A.*
- ²⁷ *NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, U.S.A.*
- ²⁸ *Department of Physics and Astronomy, University of California, Irvine, 92697, U.S.A.*
- ²⁹ *Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE),
FR-75005 Paris, France*
- ³⁰ *Departament de Física, Serra Hünter, Universitat Autònoma de Barcelona,
08193 Bellaterra (Barcelona), Spain*
- ³¹ *Department of Astronomy, The Ohio State University,
4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, U.S.A.*
- ³² *Institució Catalana de Recerca i Estudis Avançats,
Passeig de Lluís Companys, 23, 08010 Barcelona, Spain*
- ³³ *Department of Physics & Astronomy and Pittsburgh Particle Physics, Astrophysics, and Cosmology
Center (PITT PACC), University of Pittsburgh, 3941 O'Hara Street, Pittsburgh, PA 15260, U.S.A.*
- ³⁴ *Departamento de Física, DCI-Campus León, Universidad de Guanajuato,
Loma del Bosque 103, León, Guanajuato C. P. 37150, México*
- ³⁵ *Instituto Avanzado de Cosmología A. C.,
San Marcos 11 - Atenas 202. Magdalena Contreras. Ciudad de México C. P. 10720, México*
- ³⁶ *Space Sciences Laboratory, University of California, Berkeley, 7 Gauss Way, Berkeley, CA 94720, U.S.A.*
- ³⁷ *Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain*
- ³⁸ *Departament de Física, EEBE, Universitat Politècnica de Catalunya,
c/Eduard Maristany 10, 08930 Barcelona, Spain*
- ³⁹ *Department of Physics and Astronomy, Sejong University,
209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea*
- ⁴⁰ *CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain*
- ⁴¹ *Department of Physics & Astronomy, Ohio University, 139 University Terrace, Athens, OH 45701, U.S.A.*
- ⁴² *National Astronomical Observatories, Chinese Academy of Sciences,
A20 Datun Road, Chaoyang District, Beijing, 100101, P.R. China*

* *Corresponding author*