Journal of Cosmology
and Astroparticle
Physics

PAPER • **OPEN ACCESS**

Assembly bias and local Primordial non-Gaussianity from DESI DR1 quasars

To cite this article: E. Fondi *et al* JCAP06(2026)062

Journal of Cosmology and Astroparticle Physics
An IOP and SISSA journal

Received: *February 19, 2026*
Accepted: *April 30, 2026*
Published: *June 22, 2026*

Assembly bias and local Primordial non-Gaussianity from DESI DR1 quasars

E. Fondi, L. Verde et al.

Full author list at the end of the paper

E-mail: emanuelefondi@icc.ub.edu, liciaverde@icc.ub.edu

Abstract: The analysis of the large-scale clustering of quasars (QSO) observed by the Dark Energy Spectroscopic Instrument (DESI) represents a promising avenue for constraining local Primordial non-Gaussianity (PNG), parameterized by $f_{\rm NL}$. The signal to be constrained is the scale-dependent bias induced in the 2-point clustering of the considered tracer sample. The resulting constraints on $f_{\rm NL}$, however, are fully degenerate with the local PNG bias parameter b_ϕ, dependent on the assembly bias parameter p. Using IllustrisTNG hydrodynamical simulations, we select a QSO sample reflecting the selection criteria and properties of DESI QSOs, and provide a robust prior for p, and thus for b_ϕ, building on the findings of Fondi et al. 2024. We find a distribution with mean $\bar{p} \simeq 1.4$ with weak redshift dependence, stable to selection noise and consistent with the expected recent merger history typical of quasar-hosting halos. By comparing with the CAMELS simulations we demonstrate that this prior is robust to astrophysical assumptions and cosmic variance. Finally, applying this prior to the DESI DR1 dataset, we derive updated constraints on local PNG, obtaining $f_{\rm NL} = -3.3 \pm 9.2$.

Keywords: cosmological parameters from LSS, cosmological simulations, non-gaussianity, redshift surveys

ArXiv ePrint: 2602.12357

https://doi.org/10.1088/1475-7516/2026/06/062

JCAP06(2026)062

Contents

1 Introduction

The inflationary paradigm represents the standard framework for explaining the early evolution of the Universe and the origin of cosmic structures. However, the nature of inflation is still an open question and different scenarios can be constrained by probing the statistical properties of primordial fluctuations. While single-field slow-roll inflation predicts nearly Gaussian primordial fluctuations [1], more complex scenarios can generate measurable amounts of primordial non-Gaussianity (PNG). Among various types of PNG, the *local-type* [2, 3] emerges naturally in multifield inflationary models and is usually parameterized as:

$$\Phi(\mathbf{x}) = \phi(\mathbf{x}) + f_{\rm NL}\left(\phi^2(\mathbf{x}) - \langle\phi^2(\mathbf{x})\rangle\right), \tag{1.1}$$

where Φ is the primordial gravitational potential, ϕ is a Gaussian random field, and $f_{\rm NL}$ quantifies the amplitude of non-Gaussianity [4, 5]. Cosmic Microwave Background (CMB) anisotropies measurements from Planck [6] currently provide the tightest constraints on local PNG: $f_{\rm NL} = -0.9 \pm 5.1$. Nevertheless, the information accessible through the two-dimensional CMB map, especially on large angular scales, is approaching the cosmic variance limit. On the other hand, the statistical power of the three-dimensional distribution of tracers probed by galaxy surveys is increasing, having the potential to drastically improve constraints on PNG. Crucially, local PNG induces correlations between small-scale density fluctuations and long-wavelength potential modes, resulting in a distinctive scale-dependent bias in the clustering of dark matter halos and galaxies on large scales [7–11]. This feature represents a clear signature that galaxy and quasar surveys can exploit to constrain multifield inflation [12–17].

In practice, the large-scale power spectrum is sensitive to the product $f_{\rm NL}\, b_\phi$, where b_ϕ measures the response of tracer number density fluctuations to large-scale primordial potential fluctuations. Therefore, extracting robust constraints on $f_{\rm NL}$ critically depends on the knowledge of b_ϕ. Its conventional prediction, widely adopted in the literature, can be expressed as [7–9]:

$$b_\phi = 2\delta_c(b_1 - p), \tag{1.2}$$

where b_1 is the linear bias and $\delta_c \approx 1.686$ is the spherical collapse threshold. Typically, p is fixed to $p = 1$, in which case eq. (1.2) is known as *universality relation*. This relation is valid only for a *fair* sample of halos of mass M, the probability distribution of which is given by a universal mass function [18–20]. However, departures from universality are expected due to *assembly bias*, whenever the selection of tracers depends on quantities beyond dark matter halo mass [21–27]. For example, for recently merged halos, early work found a value of $p = 1.6$ [9].

This result was obtained as a particular case of the extended Press-Schechter (ePS) formalism. This framework, in which the conditional mass function is used to incorporate merger history information beyond mass alone [18, 28, 29], was then generalized to model p as a function of arbitrary halo formation redshift and validated against N-body simulations [30, 31].

Quasi-stellar objects (QSO, or quasars), which are assumed to form in recently merged halos, constitute a biased selection relative to a mass-selected sample. Consequently, the choice $p = 1.6$ was adopted to analyze the clustering of quasars from the final data release of the extended Baryon Oscillation Spectroscopic Survey (eBOSS), yielding $-23 < f_{\rm NL} < 21$ at 68% confidence level [32]. More recently, the first-year data release of the Dark Energy Spectroscopic Instrument (DESI) [33–48] has provided new measurements using its quasar sample, improving both the statistical power and control over systematics [49].

DESI analysis adopted two values for p: a universality choice, $p = 1$, and a recent-merger choice, $p = 1.6$, implying two different values of b_ϕ at fixed b_1. Since the scale-dependent bias signature constrains only the product $f_{\rm NL}\, b_\phi$, $f_{\rm NL}$ is completely degenerate with b_ϕ; consequently, DESI reports two separate constraints on $f_{\rm NL}$, one for each value of p [49].

Motivated by these considerations, in this paper we use the state-of-the-art IllustrisTNG [50, 51] hydrodynamical simulations and the CAMELS simulations [52, 53] to construct a QSO sample reflecting the selection properties of the DESI survey. Leveraging the ePS framework and the methodology of ref. [31], we calibrate a robust and physically motivated prior $\Pi(p)$ for the parameter p, thus breaking the degeneracy between $f_{\rm NL}$ and b_ϕ. Using $\Pi(p)$, we then derive a single updated constraint on local PNG from the DESI DR1 QSO dataset.

The paper is structured as follows: section 2 presents the theoretical background, introducing the ePS prescription that will be used to model p from the formation redshift. In section 3, we build a DESI-like QSO sample in IllustrisTNG and study its merger history, providing a distribution for p and testing its robustness to selection effects. Section 4 validates the prior with CAMELS, quantifying sensitivity to cosmic variance and quasar physics. In section 5, we apply the simulation-based prior to DESI DR1 and report the updated constraint on $f_{\rm NL}$. Finally, conclusions are outlined in section 6. Appendix A documents the redshift dependence of the prior.

JCAP06(2026)062

2 Theoretical background

The large-scale modulation induced by local PNG can be described analytically in the peak-background split picture [9]. In this approach, the long mode acts a background rescaling of the local amplitude of small-scale clustering (quantified by σ_8). As a result, this influences halo abundance, and the way halo number density responds to a change in σ_8 is quantified by $b_\phi \equiv 2\frac{\partial \ln n}{\partial \ln \sigma_8}$.

Therefore, b_ϕ can be predicted if we know how n (the halo mass function) changes under a rescaling of power spectrum amplitude. For a universal mass function, one finds the relation in eq. (1.2) with $p = 1$. If a sample of halos is preferentially selected by a secondary property, the universality assumption breaks down. Different formation histories can alter their abundance at fixed halo mass. In what follows, we parameterize this through the formation redshift z_f, and we should then consider a joint mass function $n = n(M, z_f) = n(M)P_{z_f}(z_f|M)$, where P_{z_f} is the conditional mass function. As a result, the b_ϕ of halos of mass M, observed at z_o, that had accreted a fraction f of their final mass by formation time z_f is given by

$$b_\phi(M, z_o, z_f, f) = 2\frac{\partial \ln n(M, z_o)}{\partial \ln \sigma_8} + 2\frac{\partial \ln P_{z_f}(fM, z_f|M, z_o)}{\partial \ln \sigma_8} = 2\delta_c(b_1 - p). \tag{2.1}$$

Effectively, the value of b_ϕ may differ from the universal value for a subpopulation of halos, because their abundance $n(M, z_f)$ does not scale with σ_8 in the same way as the global $n(M)$ does. For example, halos that assembled recently through major mergers exhibit a lower b_ϕ [7, 9], corresponding to $p > 1$. Within the extended Press-Schechter (ePS) paradigm [28–30, 54] it is possible to provide a prediction for b_ϕ, by writing down an analytical expression for P_{z_f}. This can be derived by describing halo mass accretion as a diffusion problem between redshift-dependent collapse thresholds [28, 29]. The solution can be written in a simplified form [30] by introducing the variable

$$\omega_f \equiv \frac{\delta_c \left(\frac{1}{D(z_f)} - \frac{1}{D(z_o)}\right)}{\sqrt{\sigma^2(fM_o) - \sigma^2(M_o)}}, \tag{2.2}$$

that we use to parameterize assembly bias and is proportional to $z_f - z_o$ at fixed M_o [31]. This variable depends on halo properties such as mass and formation redshift, as well as on cosmology-dependent quantities $D(z)$ and $\sigma(M)$. In terms of this variable, for $f = 0.5$ [28–30] the ePS prediction for the conditional mass function reads

$$P_{\omega_f} = 2\omega_f \mathrm{erfc}\left(\frac{\omega_f}{\sqrt{2}}\right) \tag{2.3}$$

where $\mathrm{erfc}(\omega_f) = 1 - \mathrm{erf}(\omega_f)$ is the complementary error function. Early tests of the analytical expression of P_{ω_f} in eq. (2.3) against N-body simulations were performed in refs. [30, 54]. These works found good agreement between eq. (2.3) and simulated halos with mass $M_o \gtrsim 2 \times 10^{13}\, h^{-1} M_\odot$. More recent results [31] based on IllustrisTNG showed that at higher redshift, the range of validity extends towards lower masses: at $z = 1$, halos with $M_o \gtrsim 5 \times 10^{12}\, h^{-1} M_\odot$ are accurately described by eq. (2.3). At $z = 0$, this analytical prediction breaks down, due to P_{ω_f} taking a dependence on M_o and z_o. In this case, a fitting function with a similar form, presented in [31], accurately reproduces the results. Although

JCAP06(2026)062

ω_f depends on cosmology, eq. (2.3) is effectively universal. The expression was originally derived for an EdS cosmology, yet it accurately describes results from ΛCDM simulations with different cosmologies [30, 31], at least for the range of masses and redshift we are considering. The cosmology dependence of p is fully captured in the variable ω_f, which is proportional to $z_f - z_o$. Because this framework provides a physically motivated and quantitatively accurate description of the non-Gaussian halo assembly bias [30, 31], we adopt it here.

The step that requires care is the accurate identification of the key characteristics of the host halos corresponding to the tracers' sample selected observationally. This will make it possible to construct an accurate and physically motivated prior $\Pi(p)$. In what follows, we focus on DESI QSO sample, which is characterized by a HOD-inferred mean halo mass of $M^{\rm QSO} = 4.6 \times 10^{12}\, h^{-1} M_\odot$ at redshift $z = 1.5$ [55], falling in the regime of validity of the analytic solution.

From eq. (2.1), using eq. (2.3) and the change of variable (2.2), we finally obtain

$$p = 1 - \frac{1}{\delta_c}\left[\frac{\omega_f\, e^{-\omega_f^2/2}}{\sqrt{2\pi}\,\mathrm{erfc}\left(\frac{\omega_f}{\sqrt{2}}\right)} - 1\right]. \tag{2.4}$$

As expected, for $\omega_f = 0$, corresponding to observing halos exactly at their formation $z_f = z_o$, we obtain $p = 1.6$. This represents a hard upper limit for the distributions of p we will show in the next sections.

In the absence of a secondary dependence beyond the halo mass, the derivative of P_{z_f} with respect to σ_8 in eq. (2.1) vanishes, so the second term in eq. (2.4) drops out and we recover universality, $p = 1$.

3 Quasars and the DESI-like sample in IllustrisTNG

Quasars are highly accreting, luminous active galactic nuclei (AGN), powered by accretion onto supermassive black holes. These energetic phenomena allow quasars to serve as tracers of large-scale structure, and their luminosity is directly related to the rate of matter accretion onto the black hole (BH). The accretion efficiency defines two different types of BH feedback: *radio mode* and *quasar mode*, which can be characterized through the Eddington ratio [56–58], defined as

$$\chi = \frac{\dot{M}_{\rm BH}}{\dot{M}_{\rm Edd}}, \tag{3.1}$$

where $\dot{M}_{\rm BH}$ is the BH accretion rate and $\dot{M}_{\rm Edd}$ is the Eddington accretion rate, given by:

$$\dot{M}_{\rm Edd} = \frac{4\pi G\, m_p}{\sigma_{\rm T}\, c}\frac{M_{\rm BH}}{\epsilon_r}, \tag{3.2}$$

with $M_{\rm BH}$ being the BH mass, ϵ_r the radiative efficiency (equal to 0.2 in IllustrisTNG) and the rest are physical constants: G denotes Newton constant, m_p the proton mass, c the speed of light and σ_T the Thomson scattering cross-section. Ref. [56] implemented a threshold of $\chi = 0.01$ to define the quasar mode, with higher χ values corresponding to more luminous quasars.

In this section, we make use of the TNG300-1 hydrodynamical simulation from the IllustrisTNG suite [51], which includes prescriptions for galaxy formation and black hole feedback, to study the properties of the dark matter halos hosting a sample of quasars mimicking the DESI QSO sample. In IllustrisTNG, the quasar mode is defined by the threshold $\chi = \min\left[\chi_0\left(10^{-8}M_\odot^{-1}\,M_{\rm BH}\right)^{\beta}, 0.1\right]$ with $\chi_0 = 0.002$ and $\beta = 2$ [59]. In other words, objects with $\chi > 0.1$ in IllustrisTNG are always in quasar mode.

In order to select a DESI-like QSO sample, as the mean redshift of DESI QSO sample is $z_o = 1.5$, we start by considering the $z_o = 1.5$ halo catalog of TNG300-1. We first apply a halo mass cut $\log_{10}(M/h^{-1}M_\odot) > 12.2$ to roughly match the inferred host mass of DESI QSO sample [55]. Then, to select quasar-hosting halos from this sample, we adopt a conservative approach and select the BHs with the highest χ (above $\chi = 0.1$, such that they are in quasar mode). This follows from the assumption that the most efficient, and thus luminous quasars, are likely observed by surveys like DESI. In order to match DESI QSO number density at $z_o = 1.5$, we select a total of $N_q = 258$ quasars from the TNG300-1 subhalo catalog. Then, we load the merger trees of the quasar-hosting halos and we measure their z_f by tracing their mass accretion backwards, up to the redshift at which their mass was $0.5M_o$. The variable ω_f, as defined in eq. (2.2), is then computed for each halo. Finally, the parameter p, related to assembly bias, is determined using the ePS prediction in eq. (2.4).

For comparison, we also consider a sample of 258 randomly selected halos with a similar mass distribution of the QSO hosts. This control sample is selected by minimizing the discrepancy between the two halo mass distributions, ensuring any differences in p are due to formation history, not mass.

The top panel of figure 1 shows the halo mass distributions of the two samples. By construction they are very similar: we obtain a QSO sample with mean halo mass $\log_{10}\bar{M} = 12.65$ and a control sample with $\log_{10}\bar{M} = 12.69$.

On the other hand, their p distributions are very different, as shown in the bottom panel of figure 1. In particular, when the halos are selected randomly (i.e., the sample is purely mass selected and samples all merger histories), the mean p value is very close to 1 ($\bar{p} = 0.99$), matching the prediction of the universality relation. In contrast, the p distribution of the QSO sample is skewed towards large values (with $\bar{p} = 1.41$), reflecting the fact that quasars predominantly form in halos with recent merger events. In the ePS model, p is bounded from above by $p \leq 1.6$, so the QSO sample sits close to this upper limit and the distribution develops a long tail towards lower p values. Theoretical models [60] have shown that mergers drive gas inflows that fuel rapid BH accretion, producing luminous quasars. Our findings are consistent with this picture, and this demonstrates that p is effectively determined by the merger history rather than by halo mass.

3.1 Robustness to selection effects

The width of the p distribution in figure 1 arises from the intrinsic scatter in the formation histories of the quasar-hosting halos within the QSO sample. The sample is constructed in a somewhat idealized way by imposing a sharp threshold for quasar selection, including the top N_q objects ranked by χ, up to some minimum $\chi_{\rm min}$.

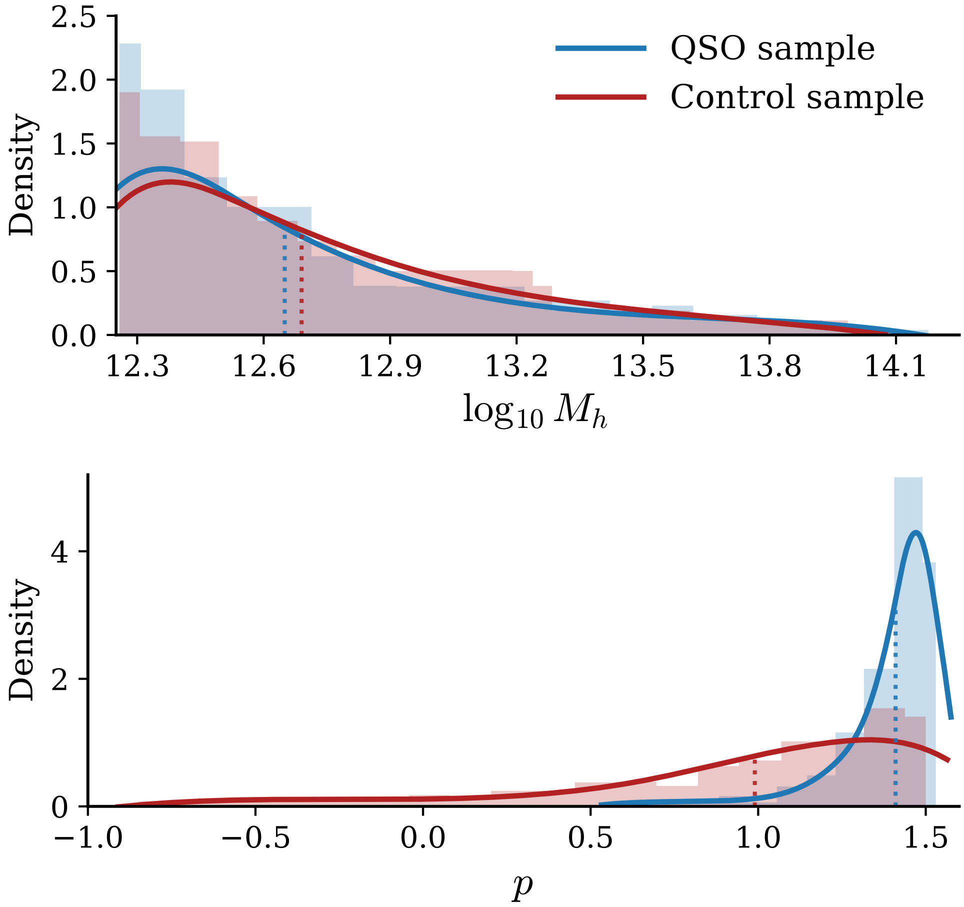

Figure 1. Distributions of host halo mass (top panel) and the p parameter (bottom panel) for the QSO (in blue) and control (in red) samples. Solid curves show kernel density estimates (KDE) of the normalized histograms and vertical dotted lines indicate the mean. While the two samples have (by construction) very similar halo mass distributions, the quasar hosts are skewed toward larger values of p, with $\bar{p} = 1.41$, while $\bar{p} = 0.99$ for the control sample.

In reality, the DESI target selection and observational cuts may exclude some quasars or may introduce luminosity-dependent biases, where some quasars appear brighter than others, leading to imperfect selection. We test the impact of selection uncertainty on the p distribution by introducing randomness in the sample selection.

Specifically, instead of enforcing a strict threshold, we introduce a probabilistic selection method. We rank the quasars by χ and define χ^i the Eddington ratio of the i-th quasar in the ranking. For this quasar, we impose a selection probability

$$P(\chi^i) = 1 - G(\chi^i), \tag{3.3}$$

where G is a Gaussian distribution centered at $\chi_{\min}$ with width σ. This mimics a scenario where some quasars below the threshold might scatter into the sample and vice versa. We vary σ to achieve different fractions r of objects replaced by this randomness (with larger σ meaning more deviation from the sharp cut).

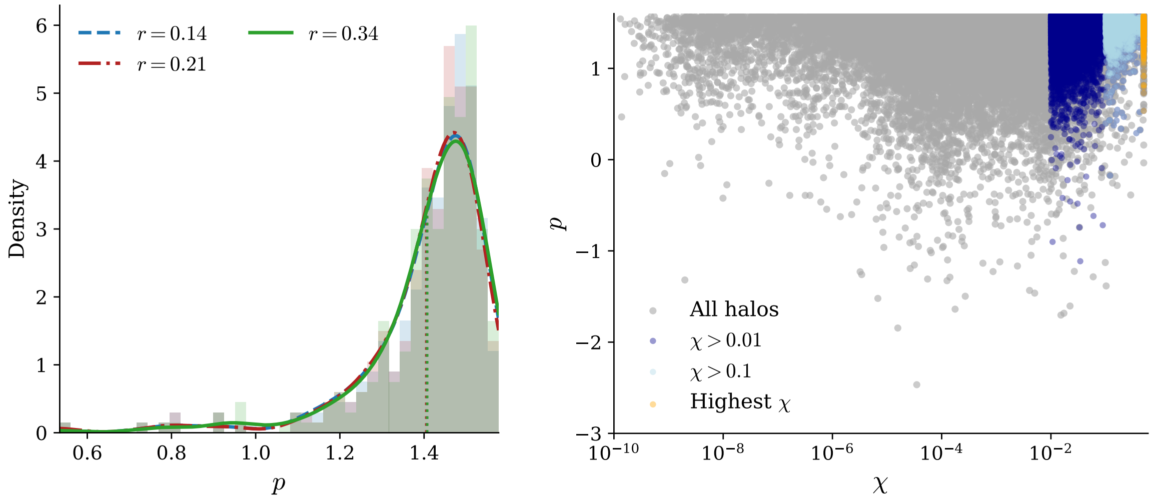

Figure 2. Left: Effect of selection noise on the p distribution of the QSO sample. Faint histograms show the empirical densities and solid curves the KDE; dashed vertical lines mark the mean of each distribution. Different curves correspond to different values of r, the fraction of quasars replaced due to selection effects. The mean p value remains stable despite introducing noise into the selection process. Right: scatter of p versus Eddington ratio χ. Grey points show all halos, while colored points highlight the Eddington-selected subsamples ($\chi > 0.01$, $\chi > 0.1$ — threshold used in the literature to define the quasar mode — and the N_q highest χ).

The left panel of figure 2 shows the resulting distributions of p for different values of r. Remarkably, we find that even when 34% of the quasars are changed with respect to the original sample, the distribution remains stable, with the same mean $\bar{p} = 1.41$.

This result suggests that moderate stochasticity or biases in quasar target selection do not significantly affect the inferred p for the sample. In fact, the host halo merger history (and therefore the p parameter) are similar for all halos hosting bright quasars, with a weak dependence on the specific values of the QSO Eddington accretion rate, as we show explicitly next.

In the right panel of figure 2 we report p as a function of χ for the full sample of halos at $z_o = 1.5$, highlighting the selections for the thresholds $\chi = 0.01$ and 0.1 discussed above. As the quasar regime is approached ($\chi \gtrsim 10^{-2}$), p increases with χ and reaches the recent-merger expectation for the highest-χ objects. In particular, it is remarkable how the broad $\chi > 0.1$ selection has the same distribution as the DESI-like QSO sample, although being $\sim$7 times denser. This explains why the fluctuations observed in the left panel are small: objects in quasar mode have similar formation histories, leading to a stable value of p, independent of small variations in selection criteria.

4 Validation with CAMELS simulations: quantifying the dependence on assumptions and subgrid physics

To assess the robustness of our results to cosmic variance and uncertainties in astrophysical modeling — particularly those affecting black hole growth and quasar activity — we complement our analysis with the CAMELS suite of simulations [52, 53]. CAMELS (Cosmology and

Astrophysics with MachinE Learning Simulations) is a set of cosmological hydrodynamical simulations designed to explore the impact of both cosmological and astrophysical parameters on structure formation, with a strong focus on subgrid physics.

In this work, we use the suite of CAMELS simulations employing IllustrisTNG galaxy formation model. A broader comparison with other CAMELS hydrodynamical models, such as CAMELS-SIMBA and CAMELS-ASTRID, would certainly be valuable. However, at the time this work is carried out, these sets are only available in 25 Mpc/h boxes, which are too small for the application considered here. We therefore leave a broader multi-model comparison to future work. From the CAMELS-TNG suite, we specifically consider two sets of simulations:

- The **CV** set, which consists of 26 simulations run with identical cosmological and astrophysical parameters but different initial conditions. This allows us to quantify the impact of cosmic variance, treating the 26 realizations as independent sub-volumes of the Universe.

- The **1P** set, which consists of simulations where astrophysical and cosmological parameters are varied one at a time, while the remaining ones are held fixed. Each parameter is sampled at 4 values different from the fiducial IllustrisTNG counterpart. We focus on variations of 9 parameters controlling AGN feedback, BH growth and QSO activity, corresponding to a total of 36 realizations. The parameters, together with their range of variation and description, are listed in table 1. This set allows us to explore how sensitive our QSO p distribution is to the assumed parameterization in IllustrisTNG, effectively marginalizing over the implementation of AGN and QSO physics.

Although the full CAMELS-1P suite also includes realizations where cosmological parameters are varied, we do not use them here. Our goal is to quantify the impact of subgrid astrophysical modelling on the p prior at fixed cosmology, consistently with the DESI DR1 analysis we update, which assumes a fiducial cosmology for the linear power spectrum. The robustness to assumptions about the underlying cosmology has been demonstrated in e.g., [31].

Each CAMELS volume (with size $L = 50\,h^{-1}$Mpc) [53] is $\sim$ 68 times smaller than TNG300-1. As a consequence, to match DESI number density at $z \sim 1.5$, we select the top 4 quasars per box after rank-ordering them by their χ. This ensures that we always select the most highly accreting quasars in a way that is robust to changes in AGN and QSO subgrid parameters, while matching the DESI QSO abundance. By combining the different realizations, we then obtain a total of $N_{\rm CV} = 104$ and $N_{\rm 1P} = 144$ objects for the CV and 1P sets, respectively.

The stacked distribution of p values across CAMELS-CV and CAMELS-1P sets is shown in figure 3, where we also report the IllustrisTNG distribution, for comparison. The CAMELS-CV distribution closely follows the TNG-300 result, showing a small impact of cosmic variance.

The 1P distribution is instead broadened by the presence of objects with slightly different merger histories than in the fiducial case. This is a result of the different AGN feedback parameterizations, as well as the QSO activities, which affect their selection. However, this effect is relatively small and the overall shape of the distribution is preserved, with a small shift of the mean from $\bar{p} = 1.41$ to $\bar{p} = 1.38$.

Parameter	Fiducial	Range	Description
$A_{\rm AGN1}$	1	[0.25, 4]	Low-accretion feedback energy
$A_{\rm AGN2}$	20	[10, 40]	Feedback burstiness, f_{re} in [59]
$M_{\rm BH,0}$	8	[2.5, 25]	BH seed mass (units of $10^5 M_\odot$)
$A_{\dot{M}_{\rm BH}}$	1	[0.25, 4.0]	$\dot{M}_{\rm BH}$ multiplier
$A_{\dot{M}_{\rm Edd}}$	1	[0.1, 10.0]	$\dot{M}_{\rm Edd}$ multiplier
$\epsilon_{f,\rm high}$	0.1	[0.025, 0.4]	High-accretion feedback energy, eq. 7 in [59]
ϵ_r	0.2	[0.05, 0.8]	Radiative efficiency
χ_0	0.002	[0.000063, 0.063]	χ threshold amplitude, eq. 5 in [59]
β	2	[0.0, 4.0]	χ threshold index, eq. 5 in [59]

Table 1. Astrophysical parameters varied in the CAMELS-1P set. Each parameter is varied individually while others are held fixed. In the second column, we report the fiducial IllustrisTNG value, while the range of variation of the parameters in CAMELS is reported in the third one. A brief description is displayed in the last column.

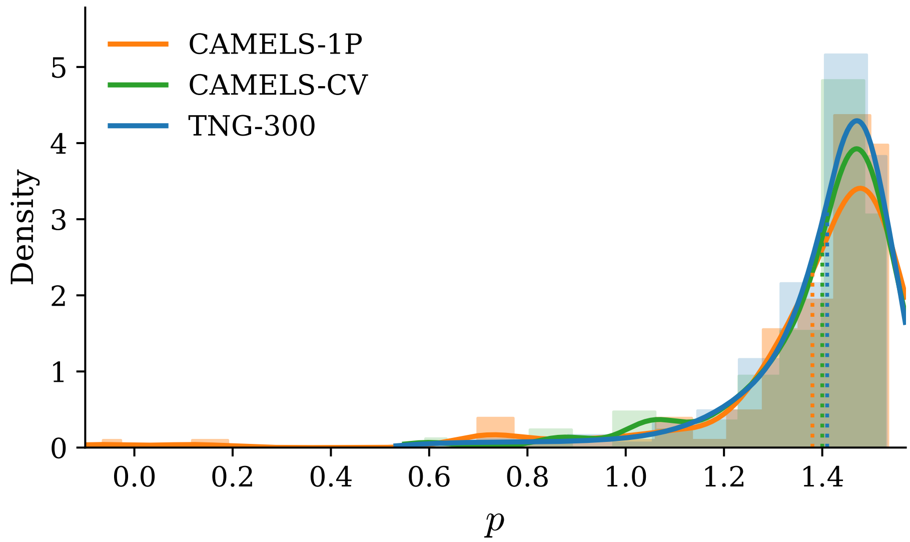

Figure 3. Empirical distributions of the parameter p for DESI-like QSO selections in simulations. Filled histograms show the stacked samples; smooth curves are KDE estimates. Blue dashed: TNG-300; green: CAMELS-CV (cosmic-variance set); orange: CAMELS-1P (astrophysics-variation set). The vertical dotted lines mark the mean of the distributions.

The results confirm that the enhancement in p due to recent merger history is not a byproduct of a specific choice of subgrid physics. Even after marginalizing over astrophysical parameters, the p distribution remains peaked at $p > 1$, in line with expectations for quasar samples dominated by recently merged halos. This reinforces the robustness of our prior on p and supports its application in cosmological analyses based on DESI QSO clustering.

5 Updated DESI DR1 constraints

In this section we use our simulation-based prior $\Pi(p)$ — the probability distribution of p for the relevant tracers — to update DESI DR1 $f_{\rm NL}$ constraints from the QSO sample. Constraints on $f_{\rm NL}$ are obtained from power spectrum measurements using the model

$$P(k,\mu) = \frac{\left[b_1(z_{\rm eff}) + \frac{b_\phi(z_{\rm eff})}{T_{\Phi\to\delta}(k,z_{\rm eff})}\, f_{\rm NL} + f_g(z_{\rm eff})\,\mu^2\right]^2}{\left[1+\frac{1}{2}\,(k\mu\Sigma_s)^2\right]^2} P_{\rm lin}(k,z_{\rm eff}) + s_{n,0}\,, \tag{5.1}$$

which includes the Kaiser effect and a damping factor, determined by Σ_s, to model redshift space distortions. Here, f_g is the linear growth rate and $s_{n,0}$ is a nuisance parameter left free in the fit to account for any residual shot-noise contribution after subtracting the Poisson shot noise from the measured power spectrum. The redshift-dependent quantities are evaluated at the effective redshift $z_{\rm eff}$, which depends on the *weighting scheme* adopted for the power spectrum estimator. We refer the reader to ref. [49] for more details. For the scope of this work, it is sufficient to mention that the DR1 analysis [49] adopted an optimal quadratic estimator (OQE) weighting scheme, designed to optimize the constraining power on $f_{\rm NL}$. In this case, quasars are weighted by

$$\tilde{w}(z) = b_1(z) - p_w\,, \tag{5.2}$$

where the linear bias $b_1(z)$ is given by a fitting function calibrated in ref. [49] and p_w denotes the assumed value of the assembly-bias parameter p when constructing the OQE weights. DR1 analysis used two choices $p_w = 1$ and $p_w = 1.6$, which yield slightly different effective redshifts for the QSO sample, $z_{\rm eff} = 1.93$ and $z_{\rm eff} = 2.08$, respectively. Correspondingly, we repeat the procedure outlined in the previous sections to build the prior on p from the snapshots closest to these redshifts, both for IllustrisTNG and CAMELS. In appendix A, we report the redshift dependence of the p distribution in IllustrisTNG, which is found to be very weak.

As a first test, we start from the DR1 MCMC chains obtained from sampling the joint posterior of $[b_1, b_\phi f_{\rm NL}]$. We propagate our prior on p by marginalization: for each chain sample $(b_1, b_\phi f_{\rm NL})$ and for each p_j sampled from the empirical prior $\Pi(p_j)$ we compute $b_\phi(p_j) = 2\delta_c(b_1 - p_j)$ and obtain $f_{{\rm NL},j} = (b_\phi f_{\rm NL})/b_\phi(p_j)$. Stacking all j yields the marginalized posterior for $f_{\rm NL}$. The left panel of figure 4 shows the prior $\Pi(p)$ at the two effective redshifts. Each Π is built by combining the DESI-like QSO samples from IllustrisTNG and CAMELS at that $z_{\rm eff}$ (i.e., we simply pool the two p catalogs). For the marginalization we use directly the N raw p values from this combined sample, so the result is independent of histogram binning or KDE bandwidth; this is equivalent to drawing N samples from an analytic prior but avoids any kernel/parametric choices. This combination naturally captures the spread induced by astrophysical uncertainty in CAMELS.

On the right panel, we display the $f_{\rm NL}$ posterior from the $[b_1, b_\phi f_{\rm NL}]$ chains with $p_w = 1.6$ OQE weights ($z_{\rm eff} \simeq 2.1$), for three different assumptions. Specifically, we verify that marginalizing over the full prior or simply fixing p to the prior mean yields negligible differences, the resulting constraint being $f_{\rm NL} = -3.3 \pm 9.2$, with shifts below 3%.As a robustness test, we deliberately apply the $z_{\rm eff} \simeq 1.9$ prior to the same chains. This is not meant as a self-consistent inference setup, but rather as a way to quantify how much the

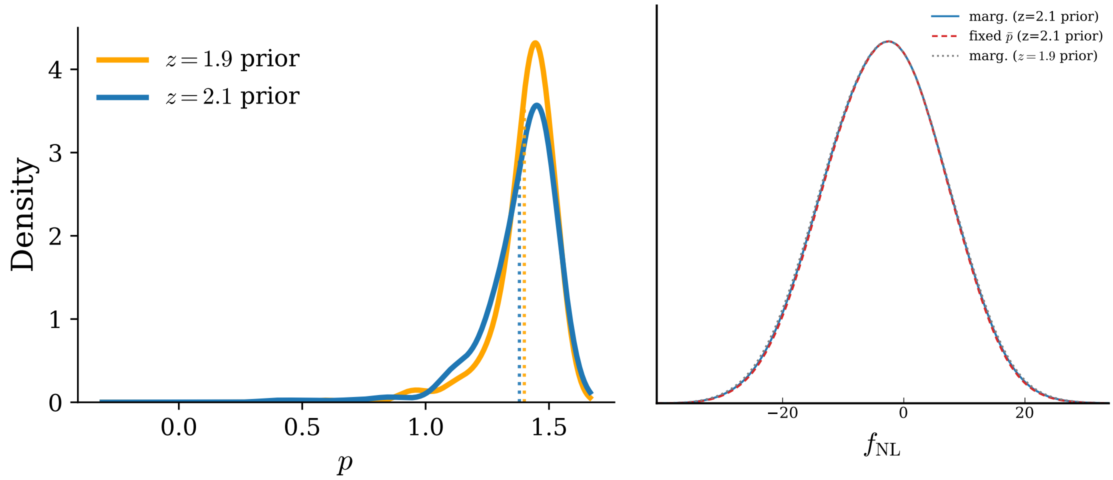

Figure 4. Left: Prior distributions of the assembly-bias parameter p built from the DESI-like QSO selections in IllustrisTNG+CAMELS (combined) at the two OQE effective redshifts. The curves show KDE-smoothed distributions for $z_{\rm eff} \simeq 1.9$ (orange) and $z_{\rm eff} \simeq 2.1$ (blue); vertical dotted lines mark the corresponding means. Right: One-dimensional posterior for $f_{\rm NL}$ from DESI DR1 QSOs $[b_1, b_\phi f_{\rm NL}]$ chains with $p_w = 1.6$ OQE weights. We compare marginalizing over the corresponding $z_{\rm eff} \simeq 2.1$ prior on p (solid blue), fixing p to the mean of that prior (dashed red), and, as a robustness test only, marginalizing over the neighboring $z_{\rm eff} \simeq 1.9$ prior (dotted gray) to quantify the impact of the weak redshift dependence of $\Pi(p)$.

mild redshift dependence of the prior propagates to the final $f_{\rm NL}$ constraint. The statistically indistinguishable differences are a consequence of the current level of precision, related to other sources of uncertainties. Despite the skewness of $\Pi(p)$, replacing the full prior by a fixed value $p = \langle p \rangle$ leads to indistinguishable constraints on $f_{\rm NL}$, and choosing the median instead of the mean would also have negligible impact on DR1 constraints. Finally, figure 5 shows the contours associated with the DESI DR1 $[b_1, b_\phi f_{\rm NL}]$ chains, post-processed by conditioning on the prior mean value of the assembly-bias parameter, $p = \langle p \rangle$.

Figure 5 highlights the key parameter degeneracies and the impact of the OQE weighting choice. As expected, the $(b_\phi f_{\rm NL}, f_{\rm NL})$ panel exhibits a tight correlation since $f_{\rm NL}$ and b_ϕ are largely degenerate before imposing a prior on p. The b_1 posteriors are shifted between weightings: b_1 is systematically larger for $p_w = 1.6$ than for $p_w = 1$, consistent with the higher effective redshift of the $p_w = 1.6$ sample and the known increase of bias with redshift. The 1D $f_{\rm NL}$ posteriors show a dependence on the weighting ($p_w = 1$ vs. 1.6), reflecting the modest change in $z_{\rm eff}$.

In principle the slight inconsistency between the choices of p_w and p can bias the recovered $f_{\rm NL}$. This is not an issue at the current level of precision, as the bias is at the $1/10\ \sigma$ level. Re-running the full power-spectrum pipeline with $p_w = \langle p \rangle$ would require recomputing the weighted power spectrum, window matrix and mock catalogues, which is computationally expensive. Instead, we use the existing DR1 measurements and interpret the shift between the two posteriors with $p_w = 1$ and $p_w = 1.6$ as an empirical upper bound on the possible bias.

We have verified that replacing the exact prior mean $\langle p \rangle$ by the rounded value $p = 1.4$ leads to indistinguishable constraints on $f_{\rm NL}$; this is the value quoted in table 2 for simplicity, where

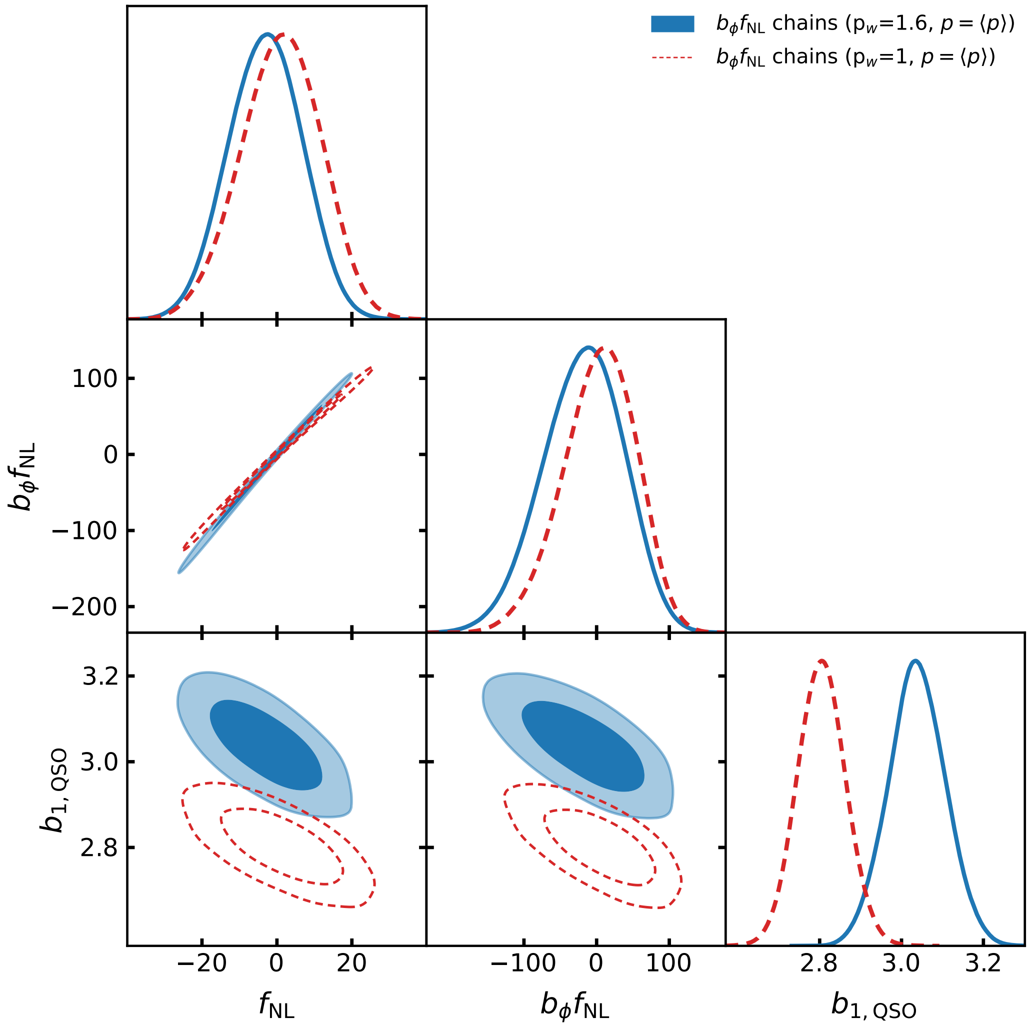

Figure 5. Corner plot comparing the different weighting choices for DESI DR1 QSOs $f_{\rm NL}$ inference. Solid contours correspond to analyses run with $p_w = 1.6$; while dashed contours use $p_w = 1$. In both cases, the plotted constraints are obtained by post-processing the $[b_\phi f_{\rm NL}, b_1]$ chains and conditioning on $p = \langle p \rangle$.

we compare our updated constraints with the original DESI DR1 results for both choices of OQE weights. For future DESI data releases, our results motivate adopting $p_w = 1.4$ directly in the measurement and mock pipelines, thereby ensuring a fully self-consistent treatment of p.

6 Conclusions

Large-scale structure constraints on local PNG rely on the scale-dependent bias induced in two-point clustering, which is sensitive to the product $f_{\rm NL}\, b_\phi$. Crucially b_ϕ depends on the sample linear bias b_1 which can be estimated from the data themselves, and the parameter p which is usually not known and notoriously hard to measure. The standard universality relation, eq. (1.2) with $p = 1$, is only valid for fair, halo mass-selected samples, while realistic

p_w	p	$f_{\rm NL}$
1.6	1.4	-3.3 ± 9.2
1.6	1.6	-2^{+11}_{-10} [49]
1	1.4	2.4 ± 11.7
1	1	$3.5^{+10.7}_{-7.4}$ [49]

Table 2. $f_{\rm NL}$ constraints obtained in this work and in the DESI DR1 analysis of ref. [49]. For each value of the OQE weighting parameter p_w, we report our result using the prior-informed value $p = 1.4$, corresponding to the rounded mean of the inferred p distribution, and the corresponding DR1 result where p was set equal to p_w.

tracer selections are generally affected by assembly bias, having $p \neq 1$. In DESI DR1 this uncertainty was handled by fixing p to either 1 or 1.6, yielding two separate $f_{\rm NL}$ constraints.

In this work, we have built a physically motivated, simulation-based prior for the assembly bias parameter p of DESI quasars and used it to obtain a single, robust $f_{\rm NL}$ constraint from DESI DR1 QSO sample, $f_{\rm NL} = -3.3 \pm 9.2$.

From a DESI-like selection in IllustrisTNG TNG300-1, quasar hosts are preferentially recent mergers, yielding a p distribution skewed to values > 1 with mean $\bar{p} \simeq 1.4$ at $z \simeq 1.5$. This result is robust to selection noise (we find no shift in $\bar{p}$ even when $\sim 34\%$ of objects are replaced) and is consistent with previous assumptions on the merger history of quasar-hosting halos. We have validated our prior with CAMELS simulations, which allow us to assess the dependence on the specific implementation of subgrid physics in IllustrisTNG. The results show a negligible impact from cosmic variance and only a mild broadening when varying AGN/QSO subgrid parameters, with the mean shifting by 1.4% to $\bar{p} \simeq 1.38$. The redshift evolution of the prior is weak enough over the range relevant for DR1 to make it negligible at the current level of precision afforded by the data. Together, these tests support the use of an empirical prior $\Pi(p)$ for DESI QSOs, built by combining IllustrisTNG and CAMELS distributions, which we report in section 5.

We propagate this prior through $b_\phi = 2\delta_c\,(b_1 - p)$ in the DESI DR1 analysis of [49]. Because the OQE weights used in [49] depend on the assumed value of the p parameter, p_w defined in eq. (5.2), we rebuild $\Pi(p)$ at the two effective redshifts corresponding to the DR1 choices ($p_w = 1$ and $p_w = 1.6$). Post-processing the DR1 $[b_1, b_\phi f_{\rm NL}]$ chains by marginalizing over $\Pi(p)$ yields $f_{\rm NL}$ posteriors that are indistinguishable from those obtained by re-running the DESI DR1 sampler, fixing p to the prior mean and sampling $[b_1, f_{\rm NL}]$ directly.

In table 2 we compare our updated constraints with the original DESI DR1 results [49]. For our fiducial choice of OQE weights, $p_w = 1.6$, adopting the prior-informed value $p = 1.4$ yields slightly tighter constraints, due to a larger value of b_ϕ, and vice-versa for $p_w = 1$. This comparison shows that updating the value of p using our simulation-based prior has only a modest impact at the current DR1 precision, while providing a more physically motivated description of the QSO sample that will become increasingly relevant for future DESI data releases.

Acknowledgments

EF thanks Shy Genel and Francisco Villaescusa-Navarro for providing access to the CAMELS simulations. EF acknowledges the support from "la Caixa" Foundation (ID 100010434, code LCF/BQ/DI21/ 11860061). Funding for this work was partially provided by project PID2022-141125NB-I00MCIN/ AEI/10.13039/501100011033 and "Center of Excellence Maria de Maeztu" award to the ICCUB CEX2024-001451-M funded by MICIU/AEI/10.13039/ 501100011033.

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 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 and Innovation of Spain (MICINN), 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 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.

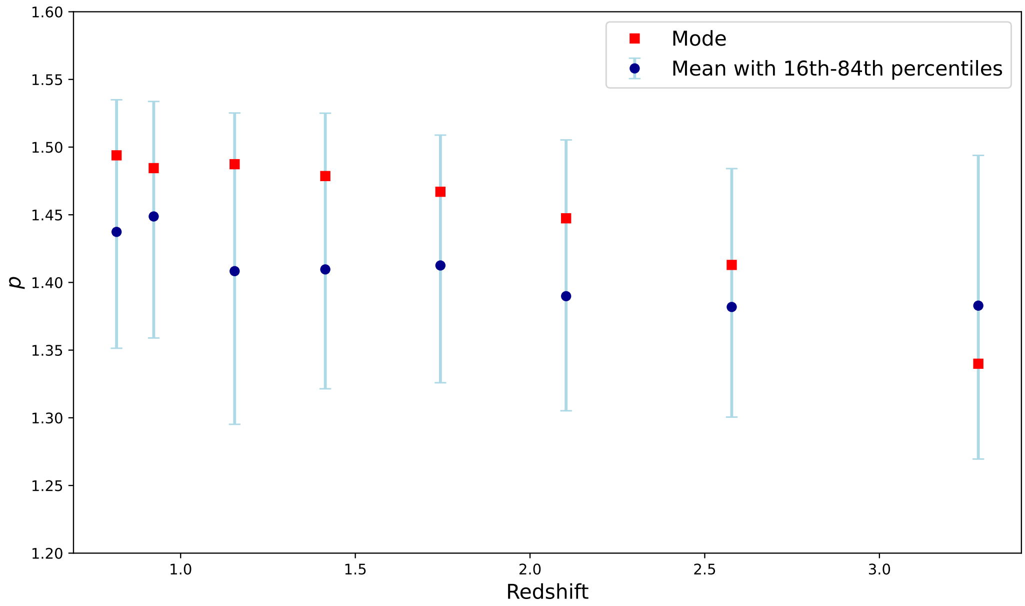

Figure 6. Evolution of p as a function of redshift. At each redshift, the red point represents the mode of the distribution, while the light blue bar denotes the 16th-84th percentiles range around the mean value, shown in dark blue.

A Redshift dependence

To assess whether the parameter p exhibits any redshift dependence, we repeat the same procedure described in section 3 at different redshifts. Specifically, we select quasar-hosting halos at various redshifts and determine their p values using eq. (2.4).

In order to match the DESI number density $n^{\rm QSO}(z)$ at each redshift, we use the tabulated values from DESI Year 1 (Y1) data and select $N_q = n^{\rm QSO}(z) \cdot V$ quasars, where V is the volume of the simulation box.

As previously observed in figure 1, the distribution of p is non-Gaussian. To summarize the variation of p across redshifts, in figure 6 we report the mean and the 16th-84th percentile range of the distributions, as well as the mode.

From figure 6, we observe a weak dependence of p on redshift, with the mean of the distribution remaining stable. The trend observed in figure 6 appears coherent with expectations, since the merger history of quasars is more intense at redshifts below $z \sim 2$, leading to an expected increase in p towards 1.6 as z decreases. Conversely, at high redshifts, where quasars typically exhibit weaker merger activity, one would naively expect p to decrease. The observed redshift dependence is very weak, especially compared to the width of the distribution, and in many practical applications it can be neglected.

Code Availability Statement. 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 are available on Zenodo at 10.5281/zenodo.18469267 in compliance with the DESI data management plan.

JCAP06(2026)062

References

[1] J.M. Maldacena, *Non-Gaussian features of primordial fluctuations in single field inflationary models*, *JHEP* **05** (2003) 013 [astro-ph/0210603] [INSPIRE].

[2] A. Gangui, F. Lucchin, S. Matarrese and S. Mollerach, *The three point correlation function of the cosmic microwave background in inflationary models*, *Astrophys. J.* **430** (1994) 447 [astro-ph/9312033] [INSPIRE].

[3] F. Bernardeau and J.-P. Uzan, *NonGaussianity in multifield inflation*, *Phys. Rev. D* **66** (2002) 103506 [hep-ph/0207295] [INSPIRE].

[4] L. Verde, L.-M. Wang, A. Heavens and M. Kamionkowski, *Large scale structure, the cosmic microwave background, and primordial non-gaussianity*, *Mon. Not. Roy. Astron. Soc.* **313** (2000) L141 [astro-ph/9906301] [INSPIRE].

[5] E. Komatsu and D.N. Spergel, *Acoustic signatures in the primary microwave background bispectrum*, *Phys. Rev. D* **63** (2001) 063002 [astro-ph/0005036] [INSPIRE].

[6] Planck collaboration, *Planck 2018 results. IX. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **641** (2020) A9 [arXiv:1905.05697] [INSPIRE].

[7] 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].

[8] S. Matarrese and L. Verde, *The effect of primordial non-Gaussianity on halo bias*, *Astrophys. J. Lett.* **677** (2008) L77 [arXiv:0801.4826] [INSPIRE].

[9] A. Slosar et al., *Constraints on local primordial non-Gaussianity from large scale structure*, *JCAP* **08** (2008) 031 [arXiv:0805.3580] [INSPIRE].

[10] N. Afshordi and A.J. Tolley, *Primordial non-gaussianity, statistics of collapsed objects, and the Integrated Sachs-Wolfe effect*, *Phys. Rev. D* **78** (2008) 123507 [arXiv:0806.1046] [INSPIRE].

[11] P. McDonald, *Primordial non-Gaussianity: large-scale structure signature in the perturbative bias model*, *Phys. Rev. D* **78** (2008) 123519 [arXiv:0806.1061] [INSPIRE].

[12] T. Giannantonio et al., *Constraining primordial non-Gaussianity with future galaxy surveys*, *Mon. Not. Roy. Astron. Soc.* **422** (2012) 2854 [arXiv:1109.0958] [INSPIRE].

[13] M. Alvarez et al., *Testing Inflation with Large Scale Structure: Connecting Hopes with Reality*, arXiv:1412.4671 [INSPIRE].

[14] R. de Putter and O. Doré, *Designing an Inflation Galaxy Survey: how to measure $\sigma(f_{\rm NL}) \sim 1$ using scale-dependent galaxy bias*, *Phys. Rev. D* **95** (2017) 123513 [arXiv:1412.3854] [INSPIRE].

[15] 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].

[16] S. Ferraro, N. Sailer, A. Slosar and M. White, *Snowmass2021 Cosmic Frontier White Paper: Cosmology and Fundamental Physics from the three-dimensional Large Scale Structure*, arXiv:2203.07506 [INSPIRE].

[17] A. Achúcarro et al., *Inflation: Theory and Observations*, arXiv:2203.08128 [INSPIRE].

[18] J.R. Bond, S. Cole, G. Efstathiou and N. Kaiser, *Excursion set mass functions for hierarchical Gaussian fluctuations*, *Astrophys. J.* **379** (1991) 440 [INSPIRE].

[19] W.H. Press and P. Schechter, *Formation of galaxies and clusters of galaxies by selfsimilar gravitational condensation*, *Astrophys. J.* **187** (1974) 425 [INSPIRE].

[20] R.K. Sheth and G. Tormen, *Large scale bias and the peak background split*, *Mon. Not. Roy. Astron. Soc.* **308** (1999) 119 [astro-ph/9901122] [INSPIRE].

[21] A. Barreira et al., *Galaxy bias and primordial non-Gaussianity: insights from galaxy formation simulations with IllustrisTNG*, *JCAP* **12** (2020) 013 [arXiv:2006.09368] [INSPIRE].

[22] 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].

[23] T. Lazeyras, A. Barreira, F. Schmidt and V. Desjacques, *Assembly bias in the local PNG halo bias and its implication for f_{NL} constraints*, *JCAP* **01** (2023) 023 [arXiv:2209.07251] [INSPIRE].

[24] 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].

[25] A. Barreira and E. Krause, *Towards optimal and robust f_nl constraints with multi-tracer analyses*, *JCAP* **10** (2023) 044 [arXiv:2302.09066] [INSPIRE].

[26] A.G. Adame et al., *PNG-UNITsims: Halo clustering response to primordial non-Gaussianities as a function of mass*, *Astron. Astrophys.* **689** (2024) A69 [arXiv:2312.12405] [INSPIRE].

[27] L.A. Perez, S. Genel, E. Krause and R.S. Somerville, *The Impact of Galaxy Formation on Galaxy Biasing, and Implications for Primordial non-Gaussianity Constraints*, arXiv:2602.04987 [INSPIRE].

[28] C. Lacey and S. Cole, *Merger rates in hierarchical models of galaxy formation*, *Mon. Not. Roy. Astron. Soc.* **262** (1993) 627 [INSPIRE].

[29] C.G. Lacey and S. Cole, *Merger rates in hierarchical models of galaxy formation. 2. Comparison with N body simulations*, *Mon. Not. Roy. Astron. Soc.* **271** (1994) 676 [astro-ph/9402069] [INSPIRE].

[30] B.A. Reid et al., *Non-Gaussian halo assembly bias*, *JCAP* **07** (2010) 013 [arXiv:1004.1637] [INSPIRE].

[31] E. Fondi et al., *Taming assembly bias for primordial non-Gaussianity*, *JCAP* **02** (2024) 048 [arXiv:2311.10088] [INSPIRE].

[32] 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].

[33] DESI collaboration, *The DESI Experiment, a whitepaper for Snowmass 2013*, arXiv:1308.0847 [INSPIRE].

[34] DESI collaboration, *The DESI Experiment Part I: Science,Targeting, and Survey Design*, arXiv:1611.00036 [INSPIRE].

[35] DESI collaboration, *The DESI Experiment Part II: Instrument Design*, arXiv:1611.00037 [INSPIRE].

[36] DESI collaboration, *Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument*, *Astron. J.* **164** (2022) 207 [arXiv:2205.10939] [INSPIRE].

[37] DESI collaboration, *The Robotic Multiobject Focal Plane System of the Dark Energy Spectroscopic Instrument (DESI)*, Astron. J. **165** (2023) 9 [arXiv:2205.09014] [INSPIRE].

[38] DESI collaboration, *The Optical Corrector for the Dark Energy Spectroscopic Instrument*, Astron. J. **168** (2024) 95 [arXiv:2306.06310] [INSPIRE].

[39] DESI collaboration, *The Spectroscopic Data Processing Pipeline for the Dark Energy Spectroscopic Instrument*, Astron. J. **165** (2023) 144 [arXiv:2209.14482] [INSPIRE].

[40] DESI collaboration, *Survey Operations for the Dark Energy Spectroscopic Instrument*, Astron. J. **166** (2023) 259 [arXiv:2306.06309] [INSPIRE].

[41] DESI collaboration, *Target Selection and Validation of DESI Luminous Red Galaxies*, Astron. J. **165** (2023) 58 [arXiv:2208.08515] [INSPIRE].

[42] DESI collaboration, *Validation of the Scientific Program for the Dark Energy Spectroscopic Instrument*, Astron. J. **167** (2024) 62 [arXiv:2306.06307] [INSPIRE].

[43] DESI collaboration, *The Early Data Release of the Dark Energy Spectroscopic Instrument*, Astron. J. **168** (2024) 58 [arXiv:2306.06308] [INSPIRE].

[44] DESI collaboration, *DESI 2024 III: baryon acoustic oscillations from galaxies and quasars*, JCAP **04** (2025) 012 [arXiv:2404.03000] [INSPIRE].

[45] DESI collaboration, *DESI 2024 IV: Baryon Acoustic Oscillations from the Lyman alpha forest*, JCAP **01** (2025) 124 [arXiv:2404.03001] [INSPIRE].

[46] DESI collaboration, *DESI 2024 II: sample definitions, characteristics, and two-point clustering statistics*, JCAP **07** (2025) 017 [arXiv:2411.12020] [INSPIRE].

[47] DESI collaboration, *DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations*, JCAP **02** (2025) 021 [arXiv:2404.03002] [INSPIRE].

[48] C. Poppett et al., *Overview of the Fiber System for the Dark Energy Spectroscopic Instrument*, Astron. J. **168** (2024) 245.

[49] E. Chaussidon et al., *Constraining primordial non-Gaussianity with DESI 2024 LRG and QSO samples*, JCAP **06** (2025) 029 [arXiv:2411.17623] [INSPIRE].

[50] D. Nelson et al., *First results from the IllustrisTNG simulations: the galaxy colour bimodality*, Mon. Not. Roy. Astron. Soc. **475** (2018) 624 [arXiv:1707.03395] [INSPIRE].

[51] D. Nelson et al., *The IllustrisTNG simulations: public data release*, Comput. Astrophys. Cosmol. **6** (2019) 2 [arXiv:1812.05609] [INSPIRE].

[52] CAMELS collaboration, *The CAMELS project: Cosmology and Astrophysics with MachinE Learning Simulations*, Astrophys. J. **915** (2021) 71 [arXiv:2010.00619] [INSPIRE].

[53] S. Genel et al., *Learning the Universe with the 2nd Generation of CAMELS: Varying 35 parameters of the IllustrisTNG model in* $(50\,Mpc/h)^3$ *boxes*, arXiv:2606.10038 [INSPIRE].

[54] F.C. Bosch, *The universal mass accretion history of cdm haloes*, Mon. Not. Roy. Astron. Soc. **331** (2002) 98 [astro-ph/0105158] [INSPIRE].

[55] S. Yuan et al., *The DESI one-per cent survey: exploring the halo occupation distribution of luminous red galaxies and quasi-stellar objects with AbacusSummit*, Mon. Not. Roy. Astron. Soc. **530** (2024) 947 [arXiv:2306.06314] [INSPIRE].

[56] D. Sijacki, V. Springel, T. Di Matteo and L. Hernquist, *A unified model for AGN feedback in cosmological simulations of structure formation*, Mon. Not. Roy. Astron. Soc. **380** (2007) 877 [arXiv:0705.2238] [INSPIRE].

JCAP06(2026)062

[57] D. Sijacki et al., *The Illustris simulation: the evolving population of black holes across cosmic time*, *Mon. Not. Roy. Astron. Soc.* **452** (2015) 575 [arXiv:1408.6842] [INSPIRE].

[58] A.K. Bhowmick et al., *Probing the $z \gtrsim 6$ quasars in a universe with IllustrisTNG physics: impact of gas-based black hole seeding models*, *Mon. Not. Roy. Astron. Soc.* **516** (2022) 138 [arXiv:2205.05717] [INSPIRE].

[59] R. Weinberger et al., *Simulating galaxy formation with black hole driven thermal and kinetic feedback*, *Mon. Not. Roy. Astron. Soc.* **465** (2017) 3291 [arXiv:1607.03486] [INSPIRE].

[60] P.F. Hopkins, L. Hernquist, T.J. Cox and D. Kereš, *A Cosmological Framework for the Co-Evolution of Quasars, Supermassive Black Holes, and Elliptical Galaxies: I. Galaxy Mergers & Quasar Activity*, *Astrophys. J. Suppl.* **175** (2008) 356 [arXiv:0706.1243] [INSPIRE].

Author List

E. Fondi[1,*], L. Verde[1,2,*], E. Chaussidon[3], J. Aguilar[3], S. Ahlen[4], S. BenZvi[5], D. Bianchi[6,7], D. Brooks[8], T. Claybaugh[3], A. Cuceu[3], A. de la Macorra[9], P. Doel[8], S. Ferraro[3,10], J.E. Forero-Romero[11,12], E. Gaztañaga[13,14,15], S. Gontcho A Gontcho[16], G. Gutierrez[17], H.K. Herrera-Alcantar[18,19], D. Huterer[20,21], M. Ishak[22], R. Joyce[23], A. Kremin[3], O. Lahav[8], C. Lamman[24], M. Landriau[3], L. Le Guillou[25], M. Manera[26,27], P. Martini[28,29,24], A. Meisner[23], R. Miquel[2,27], S. Nadathur[14], N. Palanque-Delabrouille[19,3], W.J. Percival[30,31,32], F. Prada[33], I. Pérez-Ràfols[34], G. Rossi[35], L. Samushia[36,37,38], E. Sanchez[39], D. Schlegel[3], D. Sprayberry[23], G. Tarlé[21], B.A. Weaver[23], H. Zou[40]

[1] *Institut de Ciències del Cosmos, Universitat de Barcelona (ICCUB), c. Martí i Franquès, 1, 08028 Barcelona, Spain*

[2] *Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, 08010 Barcelona, Spain*

[3] *Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, U.S.A.*

[4] *Department of Physics, Boston University, 590 Commonwealth Avenue, Boston, MA 02215, U.S.A.*

[5] *Department of Physics & Astronomy, University of Rochester, 206 Bausch and Lomb Hall, P.O. Box 270171, Rochester, NY 14627-0171, U.S.A.*

[6] *Dipartimento di Fisica "Aldo Pontremoli", Università degli Studi di Milano, Via Celoria 16, I-20133 Milano, Italy*

[7] *INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy*

[8] *Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, U.K.*

[9] *Instituto de Física, Universidad Nacional Autónoma de México, Circuito de la Investigación Científica, Ciudad Universitaria, Cd. de México C. P. 04510, México*

[10] *University of California, Berkeley, 110 Sproul Hall #5800 Berkeley, CA 94720, U.S.A.*

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

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

[13] *Institut d'Estudis Espacials de Catalunya (IEEC), c/ Esteve Terradas 1, Edifici RDIT, Campus PMT-UPC, 08860 Castelldefels, Spain*

[14] *Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth, PO1 3FX, U.K.*

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

[16] *University of Virginia, Department of Astronomy, Charlottesville, VA 22904, U.S.A.*

[17] *Fermi National Accelerator Laboratory, PO Box 500, Batavia, IL 60510, U.S.A.*

[18] *Institut d'Astrophysique de Paris. 98 bis boulevard Arago. 75014 Paris, France*

[19] *IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France*

[20] *Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109, U.S.A.*

[21] *University of Michigan, 500 S. State Street, Ann Arbor, MI 48109, U.S.A.*

[22] *Department of Physics, The University of Texas at Dallas, 800 W. Campbell Rd., Richardson, TX 75080, U.S.A.*

[23] *NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, U.S.A.*

[24] *The Ohio State University, Columbus, 43210 OH, U.S.A.*

[25] *Sorbonne Université, CNRS/IN2P3, Laboratoire de Physique Nucléaire et de Hautes Energies (LPNHE), FR-75005 Paris, France*

[26] *Departament de Física, Serra Húnter, Universitat Autònoma de Barcelona, 08193 Bellaterra (Barcelona), Spain*

JCAP06(2026)062

[27] *Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Edifici Cn, Campus UAB, 08193, Bellaterra (Barcelona), Spain*
[28] *Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, U.S.A.*
[29] *Department of Astronomy, The Ohio State University, 4055 McPherson Laboratory, 140 W 18th Avenue, Columbus, OH 43210, U.S.A.*
[30] *Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*
[31] *Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada*
[32] *Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*
[33] *Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain*
[34] *Departament de Física, EEBE, Universitat Politècnica de Catalunya, c/Eduard Maristany 10, 08930 Barcelona, Spain*
[35] *Department of Physics and Astronomy, Sejong University, 209 Neungdong-ro, Gwangjin-gu, Seoul 05006, Republic of Korea*
[36] *Abastumani Astrophysical Observatory, Tbilisi, GE-0179, Georgia*
[37] *Department of Physics, Kansas State University, 116 Cardwell Hall, Manhattan, KS 66506, U.S.A.*
[38] *Faculty of Natural Sciences and Medicine, Ilia State University, 0194 Tbilisi, Georgia*
[39] *CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain*
[40] *National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Road, Chaoyang District, Beijing, 100101, P.R. China*

[*] *Corresponding author*