

The impact of our peculiar motion on primordial non-Gaussianity measurements using the LIGER4GAL framework

Bartolomeo Bottazzi Baldi*

*Dipartimento di Fisica e Astronomia Galileo Galilei, Università di Padova, 35131, Padova, Italy and
INFN Sezione di Padova, I-35131, Padova, Italy*

Mohamed Yousry Elkhatab

*Dipartimento di Fisica – Sezione di Astronomia, Università di Trieste, Via Tiepolo 11, 34131 Trieste, Italy
INAF, Osservatorio Astronomico di Trieste, Via Tiepolo 11, I-34131 Trieste, Italy
INFN, Sezione di Trieste, Via Valerio 2, 34127 Trieste TS, Italy and
IFPU, Institute for Fundamental Physics of the Universe, via Beirut 2, 34151 Trieste, Italy*

Daniele Bertacca

*Dipartimento di Fisica e Astronomia Galileo Galilei, Università di Padova, 35131, Padova, Italy
INFN Sezione di Padova, I-35131, Padova, Italy and
INAF, Osservatorio Astronomico di Padova, Italy*

Cristiano Porciani

Argelander-Institut für Astronomie, Auf dem Hügel 71, D-53121 Bonn, Germany

Current and forthcoming galaxy surveys will map the observable Universe with unprecedented depth, sky coverage, and precision. These maps are affected by relativistic redshift-space distortions (RSDs), which become increasingly relevant on ultra-large scales. Accurate modelling of these relativistic RSDs is essential to avoid systematic biases in key cosmological measurements, such as primordial non-Gaussianity (PNG). To address this, we introduce an updated implementation of the LIGER method, LIGER4GAL, which incorporates all linear-order relativistic RSDs directly at the tracer level of high-resolution N-body simulations. We demonstrate that LIGER4GAL improves upon previous iterations of the LIGER method by reproducing the expected non-linear clustering while maintaining accuracy for relativistic RSDs on large scales. We use the updated code to generate a DESI-like sample of luminous red galaxies from the Huge MultiDark Planck simulation. By measuring the power spectrum multipoles of this sample with and without the imprint of relativistic RSDs, we assess the impact of relativistic effects on measurements of the local PNG signal (f_{nl}). We find that the omission of the “finger-of-the-observer” (sourced by the peculiar velocity of the observer) effect in the power spectrum modelling can bias measurements of f_{nl} by more than $1 (0.25) \sigma_{f_{\text{nl}}}$ in 40% (80%) of the possible realizations of the universe if scales down to $k_{\text{min}} = 0.0015 h/\text{Mpc}$ are included.

CONTENTS

I. Introduction	2	B. Impact of Relativistic RSDs on cosmological parameters	11
II. The LIGER method	3	1. Local PNG	11
A. Relativistic redshift space distortions	3	2. Relativistic effects impact on full-shape analysis	14
B. LIGER	4	V. Conclusions	15
1. LIGER4DM	4	Acknowledgments	17
2. LIGER4GAL	4	References	17
III. High-Res Mocks	5	A. Lightcone construction	19
A. HMDPL simulation	6	1. Halo trajectories	19
B. Galaxy assignment	6	2. Gravitational potential computation	19
C. Low resolution light cones	8	B. Survey functions numerical estimation	20
IV. Results	9	1. H-FS Catalogue	20
A. Validation	9	2. LRG-FS catalogue	21
1. Angular clustering	9	C. Estimators	21
2. The FOTO signal	9	1. Angular power spectrum	21
		2. 3D power spectrum	21

* bartolomeo.bottazzibaldi@phd.unipd.it

I. INTRODUCTION

Forthcoming and ongoing surveys such as DESI [4], Euclid [76], and SPHEREx [41] are designed to map the large-scale structure of the Universe in unprecedented volume, depth, and precision. These missions aim to advance our understanding of fundamental physics by probing the nature of Dark Energy (DE) and Dark Matter (DM), constraining the initial conditions of the Universe, and testing the validity of general relativity on cosmological scales.

Typically in galaxy redshift surveys, we assume a homogeneous and isotropic universe, modelled by a Friedmann-Lemaître-Robertson-Walker (FLRW) metric, to convert from redshifts to distances. However, the inhomogeneities of the observed Universe arising from structure formation introduce additional contributions to the observed redshift beyond the purely cosmological component, which are collectively referred to as redshift-space distortions (RSDs). RSDs are typically divided into two categories. The first comprises local effects, such as the leading-order contribution from the peculiar velocities of source and observer galaxies, first identified by Kaiser [61], as well as distortions arising from the local gravitational potential at both the emitter and the observer. The second category concerns line-of-sight integrated effects, such as weak lensing [72, 73, 91], which is the dominant integral contribution. Furthermore, there are other integrated terms such as the integrated Sachs-Wolfe effect [90] and the Shapiro time delay [94]. We collectively refer to all contributions beyond the source peculiar velocity term (which is routinely included in cosmological analyses) as relativistic RSDs. The impact of relativistic RSDs on the galaxy overdensity has been investigated both at linear order [e.g. 17, 22, 33, 60, 104] and at second order [e.g. 14, 16, 40, 100, 105]. In particular, several studies have quantified their effect on observables such as the multipoles of the power spectrum [31, 45, 82] and the two-point correlation function [16, 24, 46, 67, 86, 98].

The main aim of this study is the quantitative analysis of systematic biases in key cosmological measurements that arise from neglecting the impact of relativistic RSDs. Firstly, we focus on measurements of primordial non-Gaussianity (PNG) derived from large-scale-structure (LSS) analyses. PNG serves as a pivotal probe of the physics of the early Universe, offering a direct window into the inflationary mechanisms that generate the primordial fluctuations, which in turn seed structure formation [6, 11]. In the local form (typically described by the non-linearity parameter f_{nl}), PNG induces a scale-dependent correction to the bias of tracers of the underlying DM field, which becomes most prominent on large scales [36, 38, 53, 71, 97]. However, relativistic effects impact the observed power-spectrum signal in this same scale regime, where some of the contributions follow a similar scale dependence as the PNG signature [e.g. 26, 29, 51, 87]. Ignoring these effects can thus lead to systematic biases in f_{nl} measurements [e.g. 8, 27, 55].

Secondly, we focus on constraints of the growth rate of structure obtained through standard RSD analyses [83]. These studies extract cosmological information by analysing the n -point statistics of tracers using perturbative models that in-

corporate structure formation, biasing, and velocity-induced RSDs up to linear and quasi-linear scales. One widely adopted approach is the Effective Field Theory of LSS [EFT, for a review, see 59], which has been applied to datasets such as the BOSS survey [42] and, more recently, in the DESI collaboration [74]. However, in its standard formulation, EFT-based RSD modelling accounts only for the velocity contribution, neglecting relativistic RSDs. In this work, we examine the validity of that assumption and quantify the potential impact of missing contributions, focusing on constraints of the power spectrum amplitude (A_s), the DM density parameter (ω_c) and the dimensionless Hubble parameters (h) within a vanilla Λ CDM model.

To measure the impact of relativistic RSDs on these analyses, we generate mock galaxy catalogues that include relativistic RSDs up to linear order using the LIGER method and use these catalogues for cosmological analysis. To that end, we implement a new version of the method, hereafter LIGER4GAL, designed to generate lightcones with improved accuracy at non-linear scales of the matter overdensity while maintaining the linear implementation of relativistic RSDs.

Previous iterations of the LIGER method [23, 45, 46] relied on large, low-resolution simulations and employed a simple linear biasing scheme to construct galaxy mocks tailored to specific surveys. While this approach was effective at reproducing the clustering statistics on large scales, it failed in describing the clustering on smaller scales due to the resolution of the underlying simulations, and to the biasing scheme implemented. In contrast, LIGER4GAL is applied directly to halo catalogues, enabling the use of high-resolution simulations and significantly extending the validity of the resulting mock catalogues into the non-linear regime, which allows us to perform realistic cosmological analysis of the dataset. To populate haloes with galaxies, we also provide a toolkit (which is compatible with the `halotools` framework [57]) that ensures consistency with the relativistic treatment of LIGER4GAL, by imprinting the relativistic corrections on each galaxy's position and magnification, accounting also for their peculiar velocities with respect to the host halo.

In this work, we apply the new code to one realization from the MultiDark simulation suite [62, 85, 88], to generate galaxy catalogues that follow the expected distribution of DESI luminous red galaxies (LRGs) in both full-sky and partial-coverage configurations. Specifically, we make use of the Huge MultiDark Planck (HMDPL) simulation, which is a DM-only simulation with a box size of 4 Gpc/ h and a mass resolution of $m = 7.9 \times 10^{10} M_\odot/h$. These catalogues allow us to assess the impact of not accounting for relativistic RSDs on inferred cosmological parameters in two contexts: (i) full-shape fits of the power spectrum multipoles using the EFT model, and (ii) PNG analyses of the power spectrum multipoles.

This paper is structured as follows. In Sect. II, we first review the theoretical framework for relativistic RSDs and their implementation in LIGER, then introduce LIGER4GAL, highlighting its improvements over the earlier low-resolution version. In Sect. III, we describe the LIGER4GAL approach, and its application to the HMDPL simulation, producing DESI-like LRG catalogues. Sect. IV presents the valida-

tion of LIGER4GAL against theoretical predictions and the previous version of the LIGER method (hereafter LIGER4DM), followed by an assessment of the impact of relativistic effects—particularly the FOTO signal—on the inference of various cosmological parameters. We summarize and conclude in Sect. V.

Throughout this paper, we adopt Einstein’s summation convention and define the space-time metric tensor to have the signature $(-, +, +, +)$. Greek indices refer to space-time components (i.e. run from 0 to 3), while Latin indices label spatial components (i.e. run from 1 to 3). Furthermore, the Dirac delta and the Kronecker delta functions are denoted by the symbols δ^D and δ^K , respectively. Our Fourier-transform convention is $\tilde{f}(\mathbf{k}) = \int f(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}} d^3x$, and $f(\mathbf{x}) = 1/(2\pi)^3 \int \tilde{f}(\mathbf{k}) e^{+i\mathbf{k}\cdot\mathbf{x}} d^3k$. Finally, the symbol c denotes the speed of light in vacuum.

II. THE LIGER METHOD

A. Relativistic redshift space distortions

Assuming a homogeneous and isotropic Friedmann–Lemaître–Robertson–Walker (FLRW) universe, we can use the redshift measured in galaxy surveys, z_{obs} , as a proxy for the comoving distance. This relation can be written as

$$x = \int_0^{z_{\text{obs}}} \frac{c}{H(z)} dz, \quad (1)$$

where $H(z)$ denotes the Hubble parameter. Combined with the observed angular positions on the sky, $\mathbf{n} \equiv \mathbf{x}/x$, we are able to construct maps of the Universe. Such maps, which are model dependent due to Eq. (1), are then used for various cosmological statistical studies.

However, Eq. (1) neglects the impact of inhomogeneities encountered by light along its trajectory from source to observer. These inhomogeneities alter the apparent position of the source, as first demonstrated in the seminal work of Kaiser [61]. Using linear perturbation theory, Kaiser has shown that the leading-order correction on large scales – arising from the peculiar velocities of the source galaxies – induces a line-of-sight-dependent distortion to the observed galaxy overdensity δ_g , with respect to the (unobservable) real-space galaxy overdensity $\delta_{g,r}$. Nevertheless, the fact that we observe galaxies on our past light cone implies that the photons are not influenced only by local peculiar velocities but also by the entire distribution of matter they encounter. Relativistic linear treatments have shown that additional effects—such as the integrated Sachs–Wolfe effect, weak gravitational lensing, and Shapiro time delay—also affect δ_g . These corrections typically become significant at large scales [e.g., see 22, 33, 60, 104]. The numerical modelling of these effects, collectively called relativistic RSDs is the object of this study. We summarize the linear theoretical treatment of these effects in this section. We limit the implementation of relativistic corrections at linear order, as we are mainly interested in their impact on the

two-point statistics of galaxy clustering, where linear theory suffices.

To explicitly quantify the impact of relativistic RSDs, a specific gauge choice is required. In this work, we adopt the scalar-restricted Poisson gauge, i.e.

$$ds^2 = a^2(\eta) \left[-(1 + 2\Psi) c^2 d\eta^2 + (1 - 2\Phi) \delta_{ij}^K dx^i dx^j \right], \quad (2)$$

where η is the conformal time, a is the scale factor, the two potentials Φ and Ψ are the (dimensionless) Bardeen potentials [9]. Using the perturbed geodesic equation, it is possible to define a coordinate map that relates the observed four-vector, x_s^μ , to the real-space four-vector, x_r^μ . The coordinate map $\Delta x^\mu = x_r^\mu - x_s^\mu$, is given by [22, 33, 60, 104]

$$\begin{aligned} \Delta x^0 &= \frac{c}{\mathcal{H}} \delta \ln a, \\ \Delta x^i &= -x \left[n_s^i (\Phi_o + \Psi_o) + \frac{v_o^i}{c} - \frac{n_s^i (n_s^j v_j)_o}{c} \right] - \frac{c n_s^i}{\mathcal{H}} \delta \ln a \\ &\quad + \int_0^x \frac{(x - x')}{c} \left[n_s^i \partial_0 (\Phi + \Psi) - \delta_j^i \partial^j (\Phi + \Psi) \right] dx' \\ &\quad + 2n_s^i \int_0^x (\Phi + \Psi) dx', \end{aligned} \quad (3)$$

where $\partial_0 = \partial/\partial\eta$, $\partial_i = \partial/\partial x^i$, and

$$\delta \ln a = \left[\frac{(v_e^i - v_o^i)}{c} \cdot n_{s,i} - (\Phi_e - \Phi_o) - \int_0^x \frac{\partial^0 (\Phi + \Psi)}{c} dx' \right]. \quad (4)$$

Here v^i denotes the peculiar velocity vector, $\mathcal{H} = \partial_0 \ln a$ is the conformal Hubble parameter and the subscripts “o” and “e” refer to quantities evaluated at the observer’s and emitter’s positions, respectively.¹

Cosmological perturbations also alter the solid angle under which galaxies are seen by distant observers, thus enhancing, or decreasing their apparent flux [e.g. 25]. In terms of the luminosity distance, d_L , the magnification of a galaxy is defined as

$$\mathcal{M} = \left(\frac{d_L}{\bar{d}_L} \right)^{-2}, \quad (5)$$

where $\bar{d}_L$ denotes the luminosity distance in the background model universe evaluated at z_{obs} . At linear order [e.g. 16, 33],

$$\begin{aligned} \mathcal{M}(x) &= 1 + 2\Phi_e - 2 \left(1 - \frac{c}{\mathcal{H}x} \right) \delta \ln a - 2 \frac{v_o^i}{c} \cdot n_{s,i} \\ &\quad + 2\kappa - \frac{2}{x} \int_0^x (\Phi + \Psi) d\tilde{x}. \end{aligned} \quad (6)$$

where κ is the weak-lensing convergence

$$\kappa(x) = \frac{1}{2} \int_0^x (x - \tilde{x}) \frac{\tilde{x}}{x} \nabla_{\perp}^2 (\Phi + \Psi) d\tilde{x}. \quad (7)$$

¹ Equation 4 ignores corrections at the observer’s location of the scale factor a , as they only impact the mean number density and have no effect on the observed overdensity [15].

Equipped with the distortions in the observed position and flux, we are able to derive how the observed galaxy overdensity of a catalogue of a particular tracer is affected. To that end, let us define the average comoving density assuming flux-limited tracer

$$\bar{n}_g(z) = \bar{n}(L_{\text{lim}}(z), z), \quad (8)$$

where $L_{\text{lim}}(z)$ is the luminosity limit associated with a flux cut f_{cut} at a given redshift z .

Following the definitions in [16, 45], we can then quantify the rate of change with redshift and the luminosity limit of the survey via the evolution

$$\mathcal{E}(z) = - \left. \frac{\partial \ln \bar{n}(L_{\text{min}}, z)}{\partial \ln(z+1)} \right|_{L_{\text{min}}=L_{\text{min}}(z)} \quad (9)$$

and magnification

$$Q(z) = - \left. \frac{\partial \ln \bar{n}(L_{\text{min}}, z)}{\partial \ln L_{\text{min}}} \right|_{L_{\text{min}}=L_{\text{min}}(z)} \quad (10)$$

biases, respectively. Putting all the distortions together, the impact on the observed overdensity can be written as [22, 33, 60, 104]

$$\begin{aligned} \delta_g(\mathbf{x}) = & \delta_{g,r} - \frac{1}{\mathcal{H}} \frac{\partial(\mathbf{v}_e \cdot \mathbf{n})}{\partial x} - 2(1-Q)\kappa \\ & + \left[2 - \mathcal{E} + \frac{\mathcal{H}'}{\mathcal{H}^2} + \frac{2(1-Q)c}{x\mathcal{H}} \right] \frac{\mathbf{v}_o}{c} \cdot \mathbf{n} \\ & + \left[\mathcal{E} - 2Q - \frac{\mathcal{H}'}{\mathcal{H}^2} - \frac{2(1-Q)c}{x\mathcal{H}} \right] \\ & \times \left[\frac{\mathbf{v}_e}{c} \cdot \mathbf{n} - (\Phi_e - \Phi_o) - \int_0^x \frac{(\Phi' + \Psi')}{c} d\tilde{x} \right] \\ & - 2(1-Q)\Phi_e + \Psi_e + \frac{\Phi'_e}{\mathcal{H}} + (3-\mathcal{E}) \frac{\mathcal{H}\phi_v}{c^2} \\ & + \frac{2(1-Q)}{x} \int_0^x (\Phi + \Psi) d\tilde{x}, \end{aligned} \quad (11)$$

where we denote the linear velocity potential by ϕ_v . Equation (11) analytically describes the impact of relativistic RSDs on the observed overdensity. The second term denotes the standard Kaiser correction and represents the leading-order contribution. The next-to-leading order is the weak lensing term [e.g. 72, 91], which is particularly significant in photometric surveys [68]. Of particular relevance to this study is the subsequent term, sourced by peculiar velocity of the observer. This contribution has received increasing attention due to its potential to be isolated and used as a probe of the observer's motion [8, 15, 46], thereby offering insight into the apparent tension between the kinematic dipole inferred from early-Universe CMB measurements [5], and the one measured from late-Universe quasar and radio catalogues [e.g., 21, 93]. The remaining terms are proportional to the gravitational potential and are sub-leading in the standard clustering statistics [i.e. the power spectrum and two-point-correlation function 47]. However, several studies have explored the possibility of isolating these contributions, including the Doppler velocity term [23, 79].

B. LIGER

The LIGER [Light cones with General Relativity, 23, 44, 45] method is a numerical technique for generating galaxy catalogues of the Universe on the past light cone of an observer with a Λ CDM cosmology. The method uses the coordinate transformation between real and redshift space (see Eqs. 3) and the magnification (see Eq. 6) to calculate the observed positions and fluxes of tracers in Newtonian N-body simulations. In this section, we present an updated version of the LIGER method for generating light cones from high-resolution simulations after briefly reviewing the low-resolution implementation of LIGER employed in previous studies (hereafter LIGER4DM).

1. LIGER4DM

In this section, we summarise the field-based methodology introduced in Elkhatab et al. [45, 46]. This approach was designed to generate galaxy mocks that span very large volumes while keeping computational costs low, yet still preserving the accuracy of large-scale clustering statistics. LIGER4DM applies the coordinate map to the DM particles of a given simulation. For each particle, it computes its perturbed position, magnification as well as the change in redshift (see Eq. 4). The accompanying `buildcone` toolkit then uses a cloud-in-cell algorithm to construct the matter density contrasts $\delta_{\text{DM},r}(\mathbf{x}, z)$ and $\delta_{\text{DM}}(\mathbf{x}, z)$ out of the real and redshift-space positions, respectively. At linear order, we can also construct field-level functions for $\mathcal{M}$ and $\delta \ln a$ mass-weighted averages [see appendix A2 in 45]. Equipped with these four functions, we create a galaxy density field using ²

$$n_g(\mathbf{x}) \approx \bar{n}_g(x) [(b_1(x) - 1) \delta_{\text{DM},r}(\mathbf{x}) + \delta_{\text{DM}}(\mathbf{x}) + \mathcal{E}(x) \delta \ln a(\mathbf{x}) + Q(x) (\mathcal{M}(\mathbf{x}) - 1)], \quad (12)$$

which can be directly derived from Eq. (11) by replacing the real-space galaxy overdensity by the product of the $\delta_{\text{DM},r}(\mathbf{x})$ and the linear bias $b_1(x)$. Equation (12) permits the use of the same DM light cone to produce multiple galaxy maps tailor-made for distinct surveys by changing the functions $\bar{n}_g, b, Q$ and $\mathcal{E}$, which we denote as “survey functions”. A flowchart of this procedure is shown on the left side of Fig. 1.

2. LIGER4GAL

The approach described in Sect. II B 1 performs well for studies at large-scale structures. However, by design it is less accurate in describing the clustering signal on quasi-linear scales. This limitation arises from the use of the linear relation outlined in Eq. (12), which neglects the non-linear contributions that arise in the matter-tracer bias relation.

² With a slight abuse of notation, we equate the composite functions $b_1(z(x)), Q(z(x))$ & $\mathcal{E}(z(x))$ and the functions $b_1(x), Q(x)$ & $\mathcal{E}(x)$.

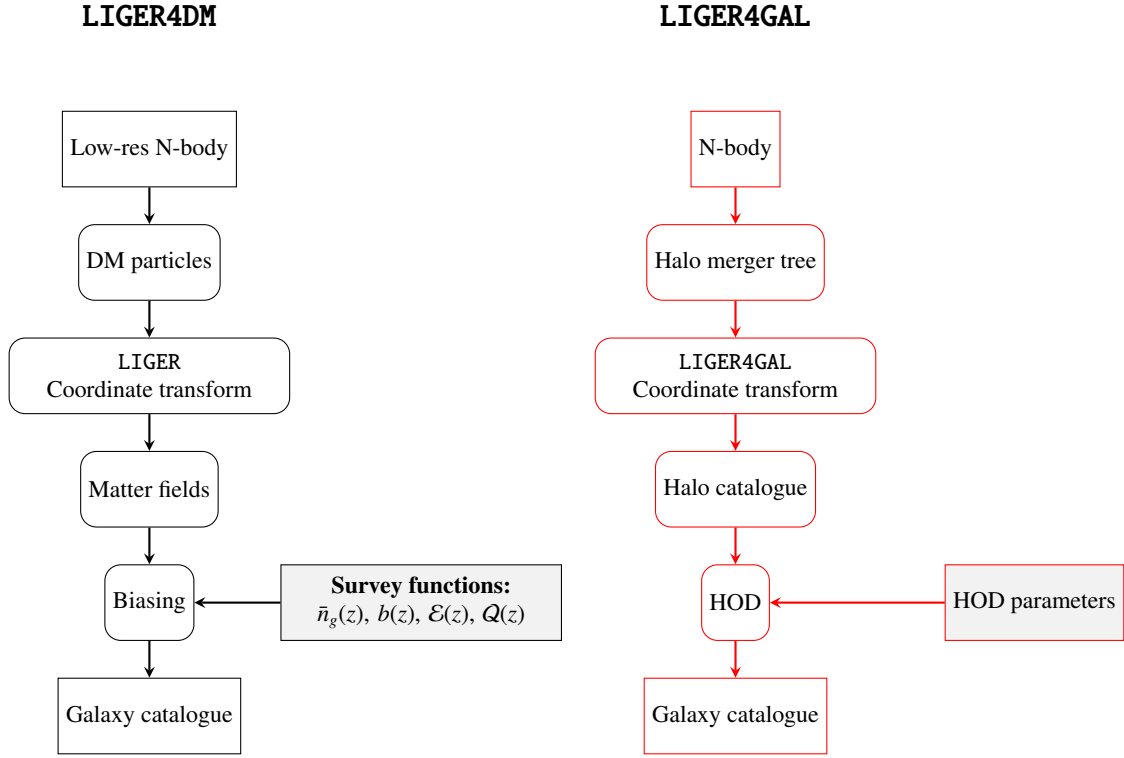


FIG. 1: *Left*: Field approach schematic of LIGER4DM. *Right*: LIGER4GAL approach implemented in this work.

To circumvent this issue, in this work we apply the LIGER method directly to biased tracers of the DM field, without relying on Eq. (12), which integrates easily with our numerical framework and facilitates the implementation of relativistic effects in standard analyses, where significantly more information about tracers is retained compared to DM. We thus present a new version of the LIGER method (LIGER4GAL) that directly post-processes large, high-resolution Newtonian simulations, avoiding the need to rely on the biasing procedure. By shifting directly the tracer under study (instead of the simulation particles), we can generate a catalogue that is consistent with both the large-scale signature arising from relativistic RSDs, and with the clustering signal of the input simulation up to non-linear scales.

The general workflow of LIGER4GAL is as follows. Starting from a halo catalogue, we reconstruct the corresponding halo merger tree (see appendix A 1) and the gravitational potential field using the available DM snapshots (see appendix A 2). Subsequently, LIGER4GAL employs the merger tree and the gravitational potential to produce a halo catalogue in redshift space by evaluating Eqs. (3), (4), and (6), in accordance with the LIGER method, i.e., by interpolating the redshift-space position of each particle to the intersection between the halo worldline and the observer’s past light cone. Lastly, for a given HOD model, we populate each halo with galaxies, correcting for their peculiar velocities by applying the shifts described in Eqs. (15) and (17) to their positions and magnifications. This procedure, detailed in Sect. III B, yields the final galaxy catalogue. The right side of Fig. 1 shows the flowchart

of this procedure.

Similar to previous applications of the LIGER method, we construct series of mock catalogues that include subsets of the full set of relativistic RSDs, which allows us to assess the importance of individual corrections. The catalogues are labelled as follows.

The label $\mathcal{R}$ denotes real-space catalogue, where no RSDs are included, $\mathcal{V}$ implies that only the peculiar-velocity terms of the tracers are applied. $\mathcal{G}$ indicates that all effects are applied, however, the observer is assumed to be at rest with respect to the cosmic microwave background rest frame, which we will label as Cosmic Rest Frame (CRF), i.e. $\mathbf{v}_o = 0$. Finally $\mathcal{O}$: all effects are applied including the observer’s velocity terms, which in this work we assume to have magnitude and direction as measured by the Planck satellite [5].

III. HIGH-RES MOCKS

In this section, we discuss our application of LIGER4GAL on the HMDPL simulation [62, 85, 88], which we use to generate a halo catalogue in redshift space with the full imprint of linear relativistic RSDs. We then describe how we populate this halo catalogue with galaxies, generating a Luminous Red Galaxy mock catalogue that is consistent with our relativistic treatment.

A. HMDPL simulation

Simulation boxes large enough to encompass the full footprint of Stage-IV galaxy surveys often lack the resolution necessary to support the construction of semi-analytic galaxy catalogues. One possible solution to circumvent that problem is to replicate smaller boxes to cover larger volumes. However, replications can introduce artificial periodicity, leading to spurious features in the clustering signal - particularly at the large scales relevant for probing relativistic effects [78]. In this work, we instead follow an approach based on the construction of DM halo lightcone catalogues, which are then populated with galaxies using the Halo Occupation Distribution framework [HOD, 13]. Specifically, we make use of the HMDPL simulation, part of the MultiDark simulation suite [85, 88]. HMDPL is a DM-only simulation with a box size of 4 Gpc/h, comprising 4096^3 particles with a mass resolution of $m = 7.9 \times 10^{10} M_\odot/h$. The simulation adopts a Planck 2014 cosmology [3] with $(\Omega_m, \Omega_b, \Omega_\Lambda, \sigma_8, n_s, h) = (0.307, 0.048, 0.693, 0.829, 0.96, 0.678)$.

From the HMDPL simulation, we import the two available DM particle snapshots at the redshifts $z = 0.49$ and $z = 0.00$ as well as the ROCKSTAR halo catalogues [12] from the CosmoSim [88] database. The halo catalogues are available in a series of snapshots, roughly uniformly distributed in the scale factor a in the range $0.1 \leq a \leq 1$ (corresponding to $8.58 \geq z \geq 0$).

We construct halo trajectories from the merger tree by tracking them across snapshots accounting for mergers and interruptions. Once each trajectory is identified, we determine its intersection with the observer's light cone and use a cubic spline to interpolate the computed shifts and magnifications at the intersection point. We use the virial mass definition of Bryan & Norman [28] as a proxy for the halo mass. We consider all halos with masses that exceed the threshold mass $M_{\text{cut}} = 10^{12} M_\odot/h$ at any redshift³ (see Appendix A 1 for further details). We denote the resulting fullsky halo lightcone by H-FS.

The other ingredient necessary for LIGER4GAL is the gravitational potential field. Since the HMDPL simulation provides only two DM snapshots, we compute the gravitational potential at these snapshots using spectral methods and employ linear theory to evolve it temporally to any redshift required by the code (see appendix A 2 for more details and validation of this approach). After constructing the trajectories of all halos and calculating the gravitational potential field, we employ LIGER4GAL to generate four halo catalogues that incorporate relativistic RSDs at varying levels (see Sect. II B 2). In addition to the galaxy catalogues derived from these halo catalogues (see Sect. III B), we also construct a mass-selected halo catalogue (see Appendix A 1 for details), used in the validation exercises (see Sect. IV A).

B. Galaxy assignment

In this section, we describe how these halos are then populated with galaxies. As relativistic effects become relevant on scales larger than the size of DM haloes, a Newtonian treatment of this process is sufficient. The key requirement is to maintain consistency with the relativistic RSDs treatment of LIGER4GAL, which primarily requires accounting for the additional distortions arising from the fact that populated galaxies do not necessarily share the exact velocities or line-of-sight positions of their parent halos.

We populate the halo catalogues with luminous red galaxies (LRGs), as their observed distribution in the DESI spectroscopic survey mostly lies within the redshift range covered by our halo catalogues [1]. For the LRG sample we adopt the HOD model used by the DESI collaboration [106]. This HOD model [originally introduced in 108] relates the mean occupation of central galaxies to the mass of a given halo, M , via

$$\bar{n}_{\text{cen}}^{\text{LRG}}(M) = f_{\text{ic}} \frac{1}{2} \text{erfc} \left[\frac{\log_{10}(M_{\text{cut}}/M)}{\sqrt{2}\sigma_{\log M}} \right], \quad (13)$$

while the satellite occupation is given by

$$\bar{n}_{\text{sat}}^{\text{LRG}}(M) = \left[\frac{M - M_0}{M_1} \right]^\alpha \bar{n}_{\text{cen}}^{\text{LRG}}(M). \quad (14)$$

Here, the HOD model is characterized by six parameters, $f_{\text{ic}}, M_{\text{cut}}, \sigma_{\log M}, \alpha, M_1, M_0$, whose values are taken from the vanilla model fit listed in Table 3 of Yuan et al. [107]. These values are extracted from the DESI 1% data at the redshift bins $z \in [0.4, 0.6]$ and $z \in [0.6, 0.8]$. We, however, interpolate their values for each redshift when populating the halo lightcone⁴, following the approach described in Hadzhiyska et al. [56]. We stress that for the procedure of populating haloes with galaxies we rely on the cosmological redshifts of the haloes, not the observed ones, as in this context z is interpreted as a time coordinate of the Universe evolution, and not as an observable.

After assigning a central galaxy and a number of satellite galaxies to each halo, we generate their phase-space (precisely, coordinate and velocity) positions as follows. Each central galaxy is assigned the same phase-space position as its host halo. For the satellites, we assume an isotropic spatial distribution within the halo, following a Navarro-Frenk-White [NFW, 81] radial profile without bias. The mass-concentration relation is adopted from [58], via the COLOSSUS code [39]. For simplicity, we further assume that the satellites are in virial equilibrium within the host halo potential and follow isotropic orbits. Under these assumptions,

³ This requires some care, as a halo near the threshold may initially have a lower mass and later grow to exceed M_{cut} . To account for such cases, we first apply a lower threshold of $5 \times 10^{11} M_\odot/h$, identify the trajectories of all haloes, and then impose the final mass cut of $10^{12} M_\odot/h$.

⁴ We apply the parameter f_{ic} , which accounts for potential incompleteness in the dataset by effectively reducing the average occupation number of each halo, slightly differently in our pipeline. Rather than modifying the average occupation directly, we implement incompleteness by randomly subsampling the galaxy sample after generation, retaining only a fraction f_{ic} of the galaxies

the satellite galaxies' radial velocity distribution can be characterized by its second moment, which is obtained by solving the Jeans equation [see Eq. 24 in 80]. Allowing us to assign positions $\mathbf{X}_g$ and velocities $\mathbf{V}_g$ to each of the satellite galaxies relative to their respective host haloes. Although the choice of satellite occupation can significantly affect the clustering signal on scales where the one-halo term is important [impacting both the overall clustering and nonlinear RSDs, see e.g., 35], it has negligible effect on the large-scale regime where the two-halo term dominates, which is the primary focus of this work. We note, however, that the consistency of the LIGER4GAL pipeline is independent of this stage in the catalogue production. In principle, more sophisticated methods for assigning phase-space positions can be easily incorporated and implemented. In fact, the numerical implementation of the HOD on the LIGER4GAL catalogues relies heavily on the HaloTools library [57], which allows for the straightforward implementation of complex occupation models.

Equipped with $\mathbf{X}_g$ and $\mathbf{V}_g$, we correct the observed positions of satellite galaxies using

$$\mathbf{x}_{g,s} = \mathbf{x}_{h,s} + \mathbf{X}_g - \frac{(\mathbf{n}_{h,s} \cdot \mathbf{v}_h)}{\mathcal{H}} \mathbf{n}_{h,s} + \frac{[\mathbf{n}_{g,s} \cdot (\mathbf{V}_g + \mathbf{v}_h)]}{\mathcal{H}} \mathbf{n}_{g,s}. \quad (15)$$

where the subscripts g & h denote the galaxies and halos, respectively. In particular, since this correction is purely radial, the observed galaxy angular positions in the sky $\mathbf{n}_{g,s}$ can be obtained before the shift is applied, i.e.,

$$\mathbf{n}_{g,s} = \frac{\mathbf{x}_{h,s} + \mathbf{X}_g}{|\mathbf{x}_{h,s} + \mathbf{X}_g|}. \quad (16)$$

Similarly, for the satellite galaxy magnification, we implement the correction

$$\mathcal{M}_{g,s} = \mathcal{M}_{h,s} + 2 \left(1 - \frac{c}{\mathcal{H}|\mathbf{x}_{h,s}|} \right) \frac{\mathbf{n}_{h,s} \cdot \mathbf{v}_h}{c} - 2 \left(1 - \frac{c}{\mathcal{H}|\mathbf{x}_{g,s}|} \right) \frac{\mathbf{n}_{g,s} \cdot (\mathbf{V}_g + \mathbf{v}_h)}{c}. \quad (17)$$

In deriving Eqs. (15) and (17), we assume that the potential-related terms are the same for all members of a given halo. We then account for the velocity terms by removing only the halo peculiar velocity contributions from Eqs. (3) and (17), to then re-add them with the velocity and line of sight of the member galaxy. We show a schematic of this procedure in Figure (2). Since in this work we do not probe scales smaller than $\approx 40 \text{ Mpc}/h$, and the maximum virial radius of haloes appearing in the simulation is $\approx 2.5 \text{ Mpc}/h$, this is an appropriate approximation.

We incorporate magnification bias through a weighting scheme applied to all galaxies [56].⁵ Specifically, for a given




FIG. 2: Schematic summarising how the haloes (Section II B 2) and then galaxies (Section III B) are shifted to build the light cones. We first apply the shift $\Delta \mathbf{x}_h$ (shown with a green arrow, see Eq. 3) to the halo positions, including both local and integrated terms. Then, we perform the peculiar velocity correction $\Delta \mathbf{x}_{g,v}$ (red arrows) described by Eq.(15) to each galaxy. In an analogue way, we produce a global magnification term $\mathcal{M}_{h,s}$ for each halo (see Eq. 6), that is applied to each galaxy inhabiting it after correcting for the individual galaxy peculiar velocities with Eq.(17).

population with magnification bias $\mathcal{Q}$, each galaxy is assigned a weight

$$w_g = [1 + (\mathcal{M}_{g,obs} - 1) \mathcal{Q}], \quad (18)$$

when computing the N -point statistics of the sample. We note that this weighting approach differs from the LIGER4DM approach used in Eq. (12). The latter relates the magnification field [computed from DM particles using a mass-assignment scheme, see appendix A2 in 45] to the observed overdensity field, whereas in our case the weights are applied directly to individual galaxies.

As for the magnification bias function $\mathcal{Q}$, we adopt the values for the LRG sample reported in Zhou et al. [109]⁶ for the two redshift bins, following the same procedure as for the HOD parameters.

At the end of this procedure, we obtain four full-sky LRG-like galaxy catalogues in the redshift bin $z \in [0.4, 0.8]$ which include varying levels of relativistic RSDs while preserving the nonlinear clustering accuracy of the HMDPL simulation. We use the LRG galaxy catalogues to produce DESI-like ones by applying the angular mask shown in Fig. 3 as well as accounting for the incompleteness parameter f_{ic} which is set to one in the full-sky case. For clarity, we adopt the following nomenclature to differentiate the catalogues: TYPE-MASK-RSDs. The first entry specifies the type of

⁵ Another possible approach is to explicitly sample galaxy luminosities [e.g., through subhalo abundance matching techniques, 64], compute the magnified flux, and then apply a flux cut. For a simple flux-limited selection, both approaches yield the same effect on clustering statistics. The more detailed approach becomes relevant only for complex selection functions.

⁶ Zhou et al. [109] define the magnification bias as $s = d \log_{10} N_g / dm$, where m is the apparent magnitude. That parametrization can be related to our definition via $\mathcal{Q} = 5/2 s$.




FIG. 3: Mollweide projection in ecliptic coordinates of the DESI Year-5 like sample we generated, used to construct the LRG-DESI and BCLRG-DESI catalogues. The blue cross marks the direction of the observer’s velocity adopted in this work [5], while the opposite direction is indicated by the red square. The galactic plane is shown as a dashed black line.

tracer (H for the halo catalogue, and LRG for galaxies), the second denotes the geometrical footprint (FS for full-sky and DESI for the DESI footprint), and the third indicates the type of RSDs applied, with $\text{RSDs} \in \{O, V, G, R\}$ (see Sect. II B 2).

To ensure that our catalogues accurately depict a DESI-like LRG sample, we compare various properties of the LRG-FS to the LRG sample published in the DESI 1% data [see Table 3 in 107]. The comparison is presented in Table I. Both the linear bias parameter b_1 (see appendix B 2 for how the parameter is estimated) and the satellite fraction f_{sat} are compatible between the two DESI 1% data and the LRG-DESI catalogues. For the average host mass $\log_{10} \bar{M}_h$, our values are slightly larger than those of the DESI 1% data, in particular in the later bin. This discrepancy is most likely sourced by the mass resolution of the HMDPL simulations that lead to an underestimation in the Halo Mass function (HMF) in the low-mass end (see appendix A 1). As for the average number density, after accounting for the sample incompleteness, we find that $\bar{n} \approx 5.4 \times 10^{-4} h^3 \text{Mpc}^{-3}$ for the LRG-DESI catalogues, which is consistent with typical expectations for this redshift range [see, e.g., 109]. Alongside the LIGER4GAL version of the code, we make available the toolkit that implements the galaxy generation pipeline: starting from a LIGER halo catalogue as input, it produces tracer populations, allowing the selection of different HOD models.

C. Low resolution light cones

To compare the LIGER4GAL version of the code with the LIGER4DM version, we generate a set of light cones following the procedure described in Sect. II B 1. For the required DM-only simulations, we utilise the MONOFONIC code [77] to produce the simulation snapshots using third-order Lagrangian perturbation theory (3LPT). We run 64 simulations, adopting the same cosmological parameters as the HMDPL run (see Sect. III A), each using 256^3 particles of




FIG. 4: Halo and LRG galaxies survey functions. With a grey dotted line we show the radial density and the linear and evolution biases of the full sky halo sample H, while with a black continuous line and a red-dashed line we show the same quantities along with the magnification bias, for the LRG-like samples in the full-sky (FS) and DESI setting respectively. The radial densities of the two LRG samples differ due to the application of the incompleteness factor of Eq. (13) only to the latter. As mentioned in Sect. IV A, the halo magnification bias is null, thus not shown.

mass $m_{\text{dm}} = 3.25 \times 10^{14} M_{\odot}/h$ within a periodic box of side length $4000 \text{Mpc}/h$. For each simulation, snapshots are stored at the same redshifts as in the HMDPL run, to ensure consistency with the HMDPL catalogues.

We then use LIGER4DM to generate a mock catalogue of the halo sample (hereafter BCH-FS)⁷, and of the LRG galaxy sample. We estimate the survey functions of the two tracers from their number counts power spectrum and radial selection function, following the approach described in Appendix B. In Fig. 4 we show the survey functions for the BCH-FS halo sample (grey dotted line), the LRG-FS (black continuous line), and LRG-DESI (red dashed line) catalogues. The halo catalogues do not constitute a flux-selected sample, so the selection function $\bar{n}(z)$ does not depend on any luminosity cut L_{cut} . As a consequence, the magnification bias is zero for the halo catalogues.

⁷ We add the label BC (for buildcone) to differentiate between the low-resolution and high-resolution catalogues.

Redshift bin	Model	$\bar{n}_g [h^3 \text{Mpc}^{-3}]$	b_1	f_{sat}	$\log_{10}(\bar{M}_h/(M_\odot/h))$
$z \in [0.4, 0.6]$	DESI 1%	—	$1.94^{+0.04}_{-0.04}$	$0.089^{+0.013}_{-0.010}$	$13.42^{+0.02}_{-0.02}$
	LRG-FS	5.5×10^{-4}	1.90	0.100	13.46
$z \in [0.6, 0.8]$	DESI 1%	—	$2.11^{+0.03}_{-0.04}$	$0.104^{+0.013}_{-0.010}$	$13.26^{+0.02}_{-0.02}$
	LRG-FS	6.4×10^{-4}	2.01	0.122	13.34

TABLE I: The DESI 1% LRG sample parameters [as reported in Table 5 of 107] are compared to the values derived from our mock catalogues (LRG-FS). Specifically, we report the linear bias, b_1 , satellite fraction f_{sat} , and the average halo mass per galaxy $\log_{10} \bar{M}_h$ (in units of $M_\odot/h$). Additionally, for the mock catalogue, the average galaxy number density $\bar{n}_g$ in each redshift bin is provided.

IV. RESULTS

In this section, we outline the main results of this work. Firstly, we quantify the improvements of our LIGER4GAL implementation through a comparison with the LIGER4DM method and with theoretical predictions. Secondly, by performing a standard f_{nl} inference from the power spectrum of our galaxy catalogues, we quantify the bias in the measurements of local primordial non-Gaussianity that arises from neglecting relativistic RSDs and - in particular - our peculiar velocity in the theoretical modelling of the galaxy power spectrum. Lastly, we perform a full-shape analysis of the power spectrum to assess the extent of the bias sourced by relativistic RSDs in the inference of cosmological parameters in a Λ CDM universe.

A. Validation

1. Angular clustering

Following the approach adopted in Elkhatab et al. [45], we use the angular power spectrum, C_ℓ , to validate the clustering properties of our mocks against the theoretical predictions, which we calculate using the CAMB [70] code. We measure the C_ℓ for both the H-FS and BCH-FS halo samples in the redshift bin $z \in [0.5, 0.6]$, using the pseudo- C_ℓ (PCL) estimator [84, see also C 1 for a brief summary]. We consider a full-sky survey, generating HEALPIX maps with 12×1024^2 pixels, which allows us to access the C_ℓ far beyond the scales of interest. For the CAMB prediction, we include linear biasing as well as all linear relativistic RSDs (the latter only for comparison with $\mathcal{O}$ mocks) through the survey functions of BCH-FS shown in Fig. 4.

We show the C_ℓ comparison for the $\mathcal{R}$ (left panel) and $\mathcal{O}$ (right panel) mocks in Fig. 5. The top panels show the C_ℓ predicted by the CAMB linear model (grey dotted line) and its non-linear [75] counterpart. Theoretical lines are compared with the spectra estimated from the H-FS (black dashed line) and BCH-FS catalogues. For the latter, we display the average spectrum with its R.M.S. scatter (red dashed line and pink shaded region). The bottom panels show the residuals with respect to the non-linear theoretical prediction. In both cases,

the large-scale signal is consistent between the two LIGER implementations, showing residuals comparable to the scatter of the BCH-FS catalogues. At smaller scales, the LIGER4DM approach overestimates the clustering signal due to the limited resolution of the linear biasing procedure on which it relies, whereas the LIGER4GAL implementation continues to provide an accurate description. As we reach $\ell = 300$, the non-linear model also ceases to reproduce the halo power spectrum, as expected.

2. The FOTO signal

The FOTO signal is an oscillatory imprint in the power spectrum sourced by the v_o term in Eq. (11). Its contribution to the monopole of the power spectrum for a full-sky survey is given by [45, 46]

$$P_{0,\text{FOTO}}(k) = \frac{16\pi^2}{3} \frac{v_o^2}{H_0^2} \frac{I_1^2(k)}{4\pi \int_{r_1}^{r_2} r^2 \bar{n}_g^2 dr}, \quad (19)$$

where

$$I_1(k) = \int_{r_1}^{r_2} \frac{r \bar{n}_g \alpha_0}{a H/H_0} j_1(kr) dr, \quad (20)$$

with the integration limits corresponding to the boundaries of the redshift shell, and

$$\alpha_0(z) = 2(1 - \mathcal{Q}) + \left[3 - \mathcal{E} - \frac{d \ln H}{d \ln(1+z)} \right] \frac{r H}{c(1+z)}. \quad (21)$$

We present the FOTO signal for our galaxy catalogues in Fig. 6, obtained by taking the difference between the $P_0(k)$ measurements from the $\mathcal{O}$ and $\mathcal{G}$ catalogues, i.e. $\Delta P_0 \approx P_{0,\text{FOTO}}$. The spectra are computed by embedding the catalogues in a box with side length of $L_{\text{FFT}} = 16000 \text{ Mpc}/h$, ensuring that the large-scale behaviour is adequately captured. We estimate the spectra using the `pypower` code [2] that implements the standard Feldman-Peacock-Kaiser [hereafter, FKP, 50, 103, see appendix C for more details]. We show the signal of the LRG-FS catalogues as red dots, with the pink-shaded region is bordered by the 16th and 84th percentiles (hereafter, σ_{68} region) of the FOTO signal computed from the BCLRG-FS catalogues. For the masked-sky case,




FIG. 5: *Upper* : We show the halo angular power spectra of the H-FS catalogues (black dashed lines) and of the low-resolution BCH-FS catalogues, where mean of the mocks is shown with a red dashed line and the R.M.S. with a pink-shaded region). Spectra are shown for the $z \in [0.5, 0.6]$ bin, with $\mathcal{R}$ and $\mathcal{O}$ mocks in the left and right panels. CAMB linear (grey dotted lines) and halofit (grey dash-dotted) predictions are overplotted. *Lower* : Residuals relative to the CAMB non-linear prediction.

the spectrum of the LRG-DESI catalogue is shown as black crosses, while the corresponding σ_{68} region derived from the BCLRG-DESI mocks is represented by the grey-shaded region. We also select two realizations from the BCLRG-FS mocks that exhibit a “typical” and a strong FOTO signal are illustrated with dotted and solid lines, respectively. For comparison, we also include the theoretical prediction (see Eq. 19) for the full-sky case. It is clear from the figure that the theoretical prediction is consistent with the numerical signal from both catalogues.

The signal has interesting features in the masked sky case. Firstly, the survey mask modifies the oscillatory pattern by shifting both the amplitudes and the peak positions as it mixes between the k -modes. Secondly, a higher fraction of the realizations produces a negative FOTO signal w.r.t. the full-sky case, because the survey mask also induces mixing between the multipoles of the power spectrum [e.g., 102], and the higher multipoles of the FOTO signal, which are not positive definite [46]. Additionally, the incompleteness factor f_{ic} applied only to masked mocks introduces a different redshift dependence, further modifying the signal.

The comparison clearly indicates a stronger FOTO effect in full-sky mocks, where the mean signal is larger and the uncertainties are smaller, due to the larger volume coverage. However, the signal of individual realizations in the masked-case varies strongly. Consequently, for the forthcoming analysis, we adopt the LRG-FS signal as a “typical” scenario, since it is an accurate description of the average signal. For the masked-sky case, we consider three spectra: the first is computed from the LRG-DESI- $\mathcal{O}$ catalogue, which exhibits a very weak FOTO signal. In addition, we construct two distinct




FIG. 6: We plot the FOTO signal for the power spectrum monopole of different LRG samples. The full-sky LRG-FS measurement is shown as red dots, while the σ_{68} region derived from the low-resolution BCLRG-FS catalogues is shown as the pink-shaded area. The corresponding masked-sky case LRG-DESI and BCLRG-DESI are represented respectively by black crosses and grey-shaded regions. Additionally, a typical and a strong masked-sky FOTO signals are shown as dotted and solid lines, respectively. Finally, the theoretical prediction for the full-sky case is plotted as a dash-dotted line.

power spectrum measurements by imprinting the FOTO signature from a “typical” BCLRG-DESI realization (i.e. within

TABLE II: Different labels for the mocks used in this work. While $\mathcal{R}$, $\mathcal{V}$, $\mathcal{G}$ and $\mathcal{O}$ refer to the different levels of GR corrections that we implement in each sample, $\mathcal{O}_t$ and $\mathcal{O}_s$ indicate two specific realizations of the full treatment $\mathcal{O}$ that we have chosen to analyse more in depth.

Mock	Description
$\mathcal{R}$	Real space
$\mathcal{V}$	Peculiar velocities only
$\mathcal{G}$	Full GR in the CRF
$\mathcal{O}$	Full GR with v_0 from Aghanim et al. [5]
$\mathcal{O}_t$	Selected typical realization of $\mathcal{O}$
$\mathcal{O}_s$	Selected strong realization of $\mathcal{O}$

the 1σ range of the possible realizations in the scales of interest) and a strong BCLRG-DESI realization onto the $\mathcal{G}$ catalogue of the LRG-DESI mocks. The power spectrum multipoles after this imprint are then

$$P_{\ell}^{\mathcal{O}_i}(k) = P_{\ell}^{\mathcal{G}}(k) + \Delta P_{\ell}^i(k), \quad (22)$$

where $P_{\ell}^{\mathcal{G}}$ are the power spectrum multipoles of the LRG-DESI- $\mathcal{G}$ catalogue, and ΔP_{ℓ}^i is the FOTO signals obtained from one of the two BCLRG-DESI realizations considered. We denote these two signals as LRG-DESI- $\mathcal{O}_t$ and LRG-DESI- $\mathcal{O}_s$, respectively. In Table II, we show the different types of mocks that are used in the analysis of the LRG-DESI sample as well as the relativistic RSDs included in these catalogues (see also Sect. II B 2).

B. Impact of Relativistic RSDs on cosmological parameters

We now turn to the central aim of this work: evaluating the impact of relativistic RSDs on cosmological parameter measurements using the galaxy catalogues we generated. As these effects affect the large-scale galaxy clustering signal, neglecting them in the theoretical modelling of the observables can bias measurements of cosmological parameters and lead to false detection of certain signatures. In particular, we focus on the FOTO signal and the bias it induces in estimates of Local PNG, as well as its influence on parameters inferred from the full-shape fit of the power spectrum, when not properly accounted for in the theoretical modelling of the galaxy power spectrum.

1. Local PNG

In this section, we measure the bias in measurements of local PNG from the observed power spectrum due to relativistic RSDs. We characterize the deviation from Gaussianity of a random field ϕ_{NG} by the following parametrization [see e.g. 63]:

$$\phi_{\text{NG}} = \phi + f_{\text{NL}}(\phi^2 - \langle \phi \rangle^2), \quad (23)$$

where ϕ is a Gaussian random field and f_{NL} denotes the amplitude of the deviation from a Gaussian distribution. The local PNG contribution at the initial conditions modifies the observed power spectrum through its impact on the linear bias of tracers of DM [see e.g. 71, 97]. Our approach to measuring the induced bias in f_{NL} is as follows. We extract f_{NL} from the LRG-FS and LRG-DESI catalogues w/o relativistic RSDs (from the $\mathcal{O}$ and $\mathcal{V}$ catalogues), using standard modelling that accounts only for velocity distortions.

Since the base HMDPL simulations (see Sec. III A) are generated with Gaussian initial conditions, i.e., $f_{\text{NL}} = 0$, any significant detection of f_{NL} in the $\mathcal{O}$ catalogues w.r.t. the $\mathcal{V}$ catalogues is attributed to systematic biases due to neglecting relativistic RSDs in the theoretical model. Comparing to the $\mathcal{V}$ catalogues, rather than directly to the fiducial value, $f_{\text{NL}} = 0$, allows us to disentangle relativistic RSDs from other possible large-scale systematics, such as wide angle effects or the radial integral constraint [see e.g. 2, 37, for methods to account for this effect]. We thus define the quantity

$$\Delta f_{\text{NL}} = f_{\text{NL}}^{\mathcal{O}} - f_{\text{NL}}^{\mathcal{V}}, \quad (24)$$

where $f_{\text{NL}}^{\mathcal{O}/\mathcal{V}}$ are measured from the $\mathcal{O}$ and $\mathcal{V}$ catalogues respectively.

We adopt the following parametrization for the redshift-space power spectrum previously used in measurements of f_{NL} [32, 34]:

$$P^{\text{Th}}(k, \mu) = \frac{\{b_1(z_{\text{eff}}) + [b_{\Phi}(z_{\text{eff}})/T_{\Phi \rightarrow \delta}(k, z_{\text{eff}})] f_{\text{NL}} + f(z_{\text{eff}})\mu^2\}^2}{\left[1 + \frac{1}{2}(k\mu\Sigma_s)^2\right]^2} \times P_{\text{lin}}(k, z_{\text{eff}}) + s_{\text{n},0} P_0^{\text{SN}}, \quad (25)$$

where $P_{\text{lin}}(k, z_{\text{eff}})$ is the linear matter power spectrum evaluated at the effective redshift z_{eff} , $T_{\Phi \rightarrow \delta}$ is the transfer function between the primordial gravitational field Φ and $\delta_{\text{DM},r}$, and b_{Φ} is the non-Gaussian contribution to the bias of the tracer. The velocity RSD contribution is included in the model via the Kaiser correction, $f\mu^2$, for RSDs [61], where f is the growth rate of structures, and μ is the cosine of the angle between the $\mathbf{k}$ -mode considered and a constant line-of-sight. The parameter $s_{\text{n},0}$ captures deviations from Poissonian shot-noise and P_0^{SN} is the shotnoise of the measurement (see appendix C). Finally, a finger-of-god (FoG) damping factor is included to partially account for non-linear RSDs modulated by the parameter Σ_s .

We parametrize b_{Φ} as

$$b_{\Phi} = 2\delta_c [b_1(z) - p], \quad (26)$$

where $\delta_c = 1.686$ is the critical overdensity in the spherical collapse model [see e.g. 95], and p encodes the merger history. The appropriate value of p remains a subject of debate [see e.g. 10]. As this work adopts an HOD that depends solely on halo-mass, we assume $p = 1$, following the approach used for the LRG sample in the DESI DR1 analysis by Chaussidon et al. [34]. Regardless of this choice, we will quantify the biases in f_{NL} in units of its uncertainty $\sigma_{f_{\text{NL}}}$. Since both these

quantities scale with $1/b_\Phi$, we therefore expect our results to be independent from this assumption.

We estimate z_{eff} using

$$z_{\text{eff}} = \frac{\sum_i w_i^2 \bar{n}_{g,i} z_i}{\sum_i w_i^2 \bar{n}_{g,i}}, \quad (27)$$

where z_i is the redshift of the i -th galaxy, $\bar{n}_{g,i}$ is the sample selection function, and w_i are the FKP weights (see Eq. C6), evaluated at the galaxy positions. We note that, in analyses of f_{nl} , the standard FKP weights are often replaced with optimal weights to enhance the constraining power on f_{nl} [see 32]. However, we adopt the FKP weights in this work, as the redshift-dependent weighting scheme has been shown to have a negligible impact on the f_{nl} constraints for the LRG sample [see Figs. 8 and 9a in 34].

We account for the different geometries of the LRG-FS and LRG-DESI through the mixing matrix formalism. That approach uses the random catalogue to calculate a mixing matrix, $\mathcal{W}_{\ell,\ell'}$, which relates the observed power spectrum multipoles to the theoretical model via [18, 19, 47, 102]

$$P_{\text{obs},\ell}(k) = \sum_{\ell'=0,2,4} \int_0^\infty k'^2 \mathcal{W}_{\ell,\ell'}(k,k') P_{\ell'}^{\text{Th}}(k') dk' - \frac{P_\ell^W(k)}{P_0^W(0)} \sum_{\ell'=0,2,4} \int_0^\infty k'^2 \mathcal{W}_{\ell,\ell'}(0,k') P_{\ell'}^{\text{Th}}(k') dk'. \quad (28)$$

Where the first term accounts for the mixing between the different modes due to the survey geometry and the second term corrects for the global integral constraint. We compute the mixing matrices from the random catalogues using the CONVO code⁸ (Salvalaggio et al. in prep.). Finally, the theoretical model multipoles are obtained by the Legendre expansion of Eq. (25)

$$P_\ell^{\text{Th}} = \frac{2\ell+1}{2} \int_{-1}^1 P^{\text{Th}}(k,\mu) \mathcal{L}_\ell(\mu) d\mu. \quad (29)$$

We implement the model of Eq. (28) and fit it to our measurements of the power spectrum monopole and quadrupole, varying four parameters $\theta = [b_1, f_{\text{nl}}, s_{\text{n},0}, \Sigma_p]$. For this analysis, we focus on the full redshift bin of our mock catalogue, $z \in [0.4, 0.8]$, as it offers the strongest constraining power. The power spectra are estimated using an FFT box of size 8000 Mpc/h, with a k -binning of approximately $\Delta k \approx 10^{-3} h/\text{Mpc}$. Following [34], we restrict the analysis to scales up to $k_{\text{max}} = 8 \times 10^{-2} h/\text{Mpc}$, in order to mitigate non-linear effects that may not be adequately captured by the simple FOG model adopted in this work. For the parameter inference, we employ the importance nested sampling algorithm for Bayesian posterior reconstruction, implemented in the Nautilus package [65]. The resulting parameter chains are then analysed and visualized using the Python package

TABLE III: Priors used in this work for the parameter inference of the primordial non-Gaussianity. The symbol $\mathcal{U}$ represents a uniform distribution.

Parameter	Prior
b_1	$\mathcal{U}(0.25, 6.0)$
f_{nl}	$\mathcal{U}(-500, +500)$
$s_{\text{n},0}$	$\mathcal{U}(-1, 2)$
$\Sigma_p [\text{Mpc } h^{-1}]$	$\mathcal{U}(0, 20.0)$

getdist [69]. We assume a Gaussian likelihood and compute the covariance using the THECOV code [7, see appendix C for more details]. Finally, we adopt wide uniform priors for the model parameters, shown in Table III.

For the data vectors, we repeat the fits under multiple scenarios. First, we fit the spectra of the $\mathcal{V}$ and $\mathcal{O}$ realizations of LRG-FS catalogues. Second, we consider the partial-sky LRG-DESI catalogues, analysing the $\mathcal{V}$, $\mathcal{O}$, and $\mathcal{O}_s$ realizations. All data vectors employed in the fits are shown in Figure 7, with the LRG-FS spectra presented in the left panel and the LRG-DESI spectra presented in the right panel. The error bars are derived from the analytical covariance, under the assumption that the covariance is unaffected by the inclusion of relativistic effects. However, the error bars differ between the LRG-FS and LRG-DESI spectra due to the different volumes of the two catalogues. In the LRG-FS case (left panel), we notice that the difference between the $\mathcal{O}$ and $\mathcal{V}$ measurements becomes increasingly evident as we decrease k , while for $k > 10^{-2} h/\text{Mpc}$ they overlap. In the LRG-DESI on the other hand, the $\mathcal{O}$ and $\mathcal{V}$ overlap for most of the k -range, while the $\mathcal{O}_s$ has a higher signal at small k . We note an apparent discrepancy between the LRG-DESI measurements and the best-fit model around $k = 10^{-2} h/\text{Mpc}$. This feature appears consistently in all mock measurements for the masked-sky case but is absent in the full-sky (LRG-FS) case. However, both the masked- and full-sky analyses yield comparable goodness-of-fit values (χ^2), indicating that this deviation is consistent with statistical fluctuations. Furthermore, because the feature persists independently of the level of relativistic effects included, it does not bias our constraints on f_{nl} .

Figure 8 presents the estimated posterior distributions in the $f_{\text{nl}} - b_1$ plane. In the LRG-FS case (upper-right panel), the power spectrum multipoles constrain f_{nl} with a half-width of the 68% credibility interval of $\sigma_{f_{\text{nl}}} \approx 9$. Including relativistic effects in the mocks (comparing the solid blue and solid green curves) induces a bias of $\Delta f_{\text{nl}} \approx 18 = 2\sigma_{f_{\text{nl}}}$ in the marginalized posterior means. For the partial-sky LRG-DESI case (lower-left panel), the f_{nl} parameter is constrained with $\sigma_{f_{\text{nl}}} \approx 16$. The weaker constraint results from the reduced sky coverage. We test two realizations of the FOTO signal: the $\mathcal{O}$ case (green dashed line), corresponding to a weaker FOTO signal (see blue crosses in Fig. 7), and the LRG-DESI case (red dashed line), corresponding to a stronger FOTO signal (see blue stars in Fig. 7). While the former leads to negligible bias in the f_{nl} measurement (where the bias is negative, consistent with the fact that for that realization $\Delta P_0 < 0$ at large scales, see Fig. 6), the latter induces

⁸ Available at <https://gitlab.com/jacoposalvalaggio/convo>.




FIG. 7: *Upper* : The power spectrum monopole (upper panels) and quadrupole (middle panels) measurements are shown for the full-sky LRG-FS catalogues (left panels) and the masked-sky LRG-DESI catalogues (right panels). The spectra of both the $\mathcal{O}$ (crosses) and $\mathcal{V}$ (barred points) are shown as well as the “strong” LRG-DESI- $\mathcal{O}_s$ (stars) realization in the partial sky case on the right. The error bars are identical for all realizations with the same geometry and are displayed only for the $\mathcal{O}$ mocks. Solid, dashed, and dotted lines indicate the best-fit models for the $\mathcal{O}$, $\mathcal{V}$, and $\mathcal{O}_s$ (right panel only) mocks, respectively. *Lower* : Residuals of the measurements and their corresponding best-fit models in units of standard deviation.

a shift of $\Delta f_{\text{nl}} \approx 24 = 1.5 \sigma_{f_{\text{nl}}}$. We compare the measurements of each catalogue with their respective best-fit model in Fig. 7. The lower panel shows the residuals between the data and the model in units of the standard deviation. In all cases, the data remain within 3σ over most of the k -range, except at the largest scales of the LRG-FS- $\mathcal{O}$ and LRG-DESI- $\mathcal{O}_s$ measurements, where the residuals are driven by the high FOTO signal.

The amplitude of the FOTO signal is both oscillatory and rapidly decreasing with k . To assess the scale at which the bias induced by this signal is removed, we repeat the parameter inference while varying only the minimum wavenumber k , included in the analysis. We compute the quantity $\Delta f_{\text{nl}}/\sigma_{f_{\text{nl}}}$ for the different pairs shown in Fig. 8 as a function of k_{min} . The resulting biases are displayed in Fig. 9 for LRG-FS- $\mathcal{O}$ (red points), LRG-DESI- $\mathcal{O}$ (grey crosses), and LRG-DESI- $\mathcal{O}_s$ (black stars). We also present the relative constraining power of the measurements, quantified as $\sigma_{f_{\text{nl}}}/\sigma_0$, shown with lines following the same colour scheme, where $\sigma_0 = \sigma_{f_{\text{nl}}}(k_{\text{min}} = 10^{-3})$ in the lower panel. For the LRG-FS case, the bias falls below $0.5 \sigma_{f_{\text{nl}}}$ for $k_{\text{min}} > 2 \times 10^{-3} h/\text{Mpc}$, while it remains negligible at all scales for the LRG-DESI- $\mathcal{O}$ realization. In contrast, for the partial-sky LRG-DESI- $\mathcal{O}_s$ realization, the bias persists up to $k_{\text{min}} \approx 4 \times 10^{-3} h/\text{Mpc}$. We emphasize that this behaviour depends on the survey selection (see Eq. 21) and the redshift bin considered; thus, for different selection, a dedicated study is required to establish the appropriate k_{min} . However, increasing the k_{min} to avoid this bias entirely nearly doubles the size of the constraints on the f_{nl} as

shown in the lower panel of Fig. 9.

As illustrated in Figs. 6 and 9, different realizations of the FOTO signal can produce significantly different biases in the f_{nl} measurements. To account for this, we repeat the analysis across all 64 FOTO realizations, each time imprinting the signal onto the $\mathcal{G}$ -sample power spectrum multipoles of the LRG-DESI measurements. In Fig. 10, we present the fractional percentage of realizations, $R(x, k_{\text{min}})$, that result in a bias in f_{nl} exceeding $x, \sigma_{f_{\text{nl}}}$ as a function of k_{min} . For $x = 1$ (blue dotted line), we find that approximately 20% of the realizations produce a bias larger than the measurement error when including scales below $0.003 h/\text{Mpc}$. As progressively larger modes are excluded, $R(x, k_{\text{min}})$ converges to zero. Similar behaviour is observed for the $x = 0.5$ (red dashed line) and $x = 0.25$ (black solid line) cases, with higher values of $R(x, k_{\text{min}})$ that decay less rapidly as x decreases. In particular, for all values of x , the convergence of $R(x, k_{\text{min}})$ towards zero is not monotonic. This is consistent with Fig. 9, where the oscillatory nature of the FOTO signal implies that including larger scales can, in some cases, slightly reduce the bias in f_{nl} .

Several remarks help contextualize the interpretation of the preceding results. Firstly, we ignore the impact of various angular systematics [e.g., Milky Way extinction and variations of the photometric zero-point 48, 49, 66, 89], which strongly affect power-spectrum measurements on large scales, as our goal is to demonstrate the significance of the FOTO signal even in the best-case scenario. Secondly, Chaussidon et al. [34] found that the geometrical systematic mitigation techniques of the DESI LRG sample remained valid




FIG. 8: Constraints on f_{nl} and b_1 obtained from the full-sky LRG-FS (top right, solid lines) and masked-sky LRG-DESI (bottom left, dashed lines) catalogues, with all relativistic effects (blue) and velocity-only (green) mocks. The fiducial $f_{\text{nl}} = 0$ value is marked by the grey dotted line. Moreover, for the partial sky case, we also show the constraints obtained from the “strong” LRG-DESI- $\mathcal{O}_s$ realization in red.

down to $k_{\text{min}} = 10^{-3} h/\text{Mpc}$, and the imaging systematics mitigation down to $k_{\text{min}} = 3 \times 10^{-3} h/\text{Mpc}$. However, in their measurements they conservatively limited the analysis to $k_{\text{min}} = 6 \times 10^{-3} h/\text{Mpc}$. Since for this scale cut we find that there is a probability $\sim 20\%$ to present biases larger than $0.2 \sigma_{f_{\text{nl}}}$, and taking into account that the DR1 sky coverage they probed is smaller than the Y5-like mask used in this work, we do not expect the FOTO to have substantially biased their LRG measurements. Thirdly, in this analysis, we choose to vary the type of mocks while keeping the model restricted to the peculiar velocity contributions, rather than using only the $\mathcal{O}$ mocks and explicitly adding the FOTO effect in the model. This approach is motivated by two considerations: (i) it allows us to explore different realizations of the FOTO signal, rather than only its expectation value, and (ii) accounting analytically for the survey geometry is complex [see 46], since the FOTO effect generates both odd and even multipoles that extend beyond $\ell = 4$. Nevertheless, for a Gaussian likelihood with the same covariance, the average shift in f_{nl} computed by varying the mocks is equivalent to the shift that would be induced by including the FOTO effect directly in the model. Finally, we note that the FOTO signal of our sample has an average signal-to-noise ratio (S/N) of around 1 for the partial-sky case. A higher S/N [as shown for the Euclid Spectroscopic Sample, where $\text{S/N} \approx 4, 45$] would lead to stronger biases in the f_{nl} constraints.




FIG. 9: *Upper:* Bias in the f_{nl} measurement due to the FOTO signal plotted as a function of the minimum k employed in the analysis, for the full-sky sample LRG-FS- $\mathcal{O}$ (red dots), and the masked-sky samples LRG-DESI- $\mathcal{O}$ (black crosses) and LRG-DESI- $\mathcal{O}_s$ (black stars). The bias is shown in units of the standard deviation of the measurement, which differ between the DESI-like and full-sky samples. *Lower:* Change in constraining power as a function of k_{min} presented as $\sigma_{f_{\text{nl}}}/\sigma_0$, where $\sigma_0 = \sigma_{f_{\text{nl}}}(k_{\text{min}} = 10^{-3})$. The colour scheme matches the upper panel.




FIG. 10: Fraction of masked-sky BCLRG-DESI mock realizations yielding a bias $|\Delta f_{\text{nl}}/\sigma_{f_{\text{nl}}}| > x$ for different thresholds x , as a function of the k_{min} adopted in the analysis. The grey vertical line shows the k_{min} used for the LRG f_{nl} constraints of Chaussidon et al. [34].

2. Relativistic effects impact on full-shape analysis

In this subsection we examine the impact of relativistic RSDs on the cosmological parameters inferred from the full-shape analysis of the power spectrum. To that end, we infer the cosmological parameters from the LRG-DESI catalogue

TABLE IV: The prior distributions used for parameter inference in the EFT model. Cosmological parameters are separated by a horizontal line from the bias and nuisance parameters. The symbols $\mathcal{N}$ and $\mathcal{U}$ represent Gaussian and uniform distributions, respectively.

Parameter	Prior
h	$\mathcal{U}(0.5, 1.0)$
ω_c	$\mathcal{U}(0.085, 0.155)$
A_s	$\mathcal{U}(1.4, 2.6)$
n_s	$\mathcal{N}(0.96, 0.0041)$
ω_b	$\mathcal{N}(0.02218, 0.00055)$
b_1	$\mathcal{U}(0.25, 4.0)$
b_2	$\mathcal{U}(-5, 5)$
γ_2	$\mathcal{U}(-5, 5)$
γ_{21}	Coevolution
c_0 [Mpc ²]	$\mathcal{N}(0, 50)$
c_2 [Mpc ²]	$\mathcal{N}(0, 50)$
c_4 [Mpc ²]	$\mathcal{N}(0, 50)$
c_{nl0} [Mpc ⁴]	$\mathcal{N}(0, 50)$
N_0^p	$\mathcal{N}(0, 3)$
N_{20}^p [Mpc ²]	$\mathcal{N}(0, 50)$

both with and without relativistic RSDs, directly measuring their impact parameter inference. For the theoretical modeling of the power spectrum, we adopt the effective field theory (EFT) of the large-scale structure, which describes the observables of the galaxy clustering through perturbative expansions of the nonlinear density and velocity fields [for a recent review, see 59]. The EFT model divides the theoretical power spectrum into three main contributions. The first contribution stems from the dynamical, bias, and RSD terms computed at the leading and next-to-leading order in standard perturbation theory (SPT). That contribution is parameterized by the cosmological parameters $\{h, \omega_c, \omega_b, A_s, n_s\}$ and the bias parameters $\{b_1, b_2, \gamma_2, \gamma_{21}\}$. The second arises from the stochastic components of the density and velocity fields, represented by the stochastic shot-noise parameters: $\{N_0^p, N_{20}^p, N_{22}^p\}$. Whereas the previous contributions are already present in SPT, EFT introduces a set of parameters that encode small-scale corrections not captured by SPT. These corrections are implemented via “counterterms”, $c_0, c_2, c_4, c_{\text{nl0}}$.

We compute the model predictions via the Gaussian-process emulator comet [43], with priors summarized in Table IV. All parameters are varied except γ_{21} and N_{22}^p , with γ_{21} set by the coevolution relation

$$\gamma_{21} = \frac{2}{21}(b_1 - 1) + \frac{6}{7}\gamma_2, \quad (30)$$

and N_{22}^p fixed to zero. We perform parameter inference performed using the importance nested sampling algorithm implemented in Nautilus, and the chains are analyzed and visualized with getdist.

We perform parameter inference using the monopole and

quadrupole of the LRG-DESI power spectrum over the range $k \in (2 \times 10^{-3}, 0.15) h/\text{Mpc}$, with bin width $dk = 3 \times 10^{-3} h/\text{Mpc}$. We employ the $\mathcal{V}$, and $\mathcal{O}$ catalogues to measure the impact of all relativistic RSDs, along with two variations of the $\mathcal{O}$ catalogue in which we imprint a typical ($\mathcal{O}_t$) and a strong ($\mathcal{O}_s$) FOTO signal (dotted and solid black lines in Fig. 11). We show the inferred cosmological parameters in Fig. 11. We can deduce from the figure that adding relativistic RSDs with a weak FOTO signal has little impact on the parameter constraints, as we can see from the overlapping contours of the $\mathcal{V}$ and $\mathcal{O}$ mocks. The typical and strong FOTO signal, on the other hand, induce a systematic shift in all three parameters shown. That shift increases in amplitude as the FOTO signal strength increases. We note that the posterior means are consistent with the maximum-likelihood estimate derived from the MCMC chain at the few-percent level, suggesting that prior-volume effects are negligible [e.g. 30, 96]. To identify the k -range at which this bias is significant, we use the Figure-of-Bias (FoB).

We define the FoB as

$$\text{FoB} = \left[(\theta_a - \theta_{\mathcal{V}}) S_a^{-1} (\theta_a - \theta_{\mathcal{V}})^T \right]^{1/2}, \quad (31)$$

where, θ_a are vectors containing the posterior means of the three cosmological parameters, the subscript denotes the type of catalogue $a \in \{\mathcal{O}, \mathcal{O}_t, \mathcal{O}_s\}$, and S_a is the parameter covariance matrix, computed from each MCMC chain using the getdist package. Assuming that the covariance does not change between the different chains and that the posterior is a Gaussian distribution, the FoB quantifies the covariance-weighted difference between the posterior means of the mocks with relativistic RSDs and their velocity-only counterpart.

The covariance matrix converts the parameter differences into levels of significance, such that a value of $\text{FoB} = 1.88 (2.83)$ can be interpreted as corresponding to the 68% (95%) credibility levels. We recompute the chains varying the minimum scales considered and plot the FoB in Fig. 12. As seen from the figure, the FOTO signal leads to a 2σ bias at the largest scales considered for the LRG-DESI- $\mathcal{O}_s$ mocks. However, that bias is removed by excluding the largest scales from parameter inference, showing that as we consider the typical range of scales for a full-shape analysis, the impact of the FOTO is negligible.

V. CONCLUSIONS

Stage-IV galaxy surveys will measure redshifts and angular positions for tens of millions of galaxies over a large fraction of the sky, spanning redshift ranges that allow us to probe clustering on previously inaccessible scales. In constructing galaxy maps from these surveys, it is standard to assume that galaxies populate an unperturbed background universe when converting redshifts into comoving distances. However, galaxy positions and fluxes are influenced by the inhomogeneities encountered by the light along its path from source to observer, leading to various RSDs. While the leading RSD effect, arising from peculiar velocity differences between




FIG. 11: Corner plot showcasing shifts in cosmological constraints due to different relativistic effects from the LRG-DESI catalogues.




FIG. 12: FoB values obtained from different catalogues w.r.t. the $\mathcal{V}$ mocks are plotted against the minimum scale included in the parameter inference.

sources and the observer Kaiser [61], is routinely accounted for, additional relativistic contributions—including Doppler terms, gravitational lensing, Sachs–Wolfe effects, and time delays—have often been neglected. These corrections become increasingly relevant on scales approaching the Hubble radius. In this work, we present a new implementation of the LIGER method, LIGER4GAL, that incorporates relativistic effects directly at the tracer level in mock halo/galaxy catalogues derived from high-resolution Newtonian N-body simulations. Our method makes it possible to post-process high-resolution simulations by imprinting relativistic RSDs by using the halo merger tree to produce a halo catalogue with the effects included.

We also developed a toolkit that populates the generated halo catalogue with galaxies or other tracers, consistently with the relativistic RSDs imprinted by LIGER. For a chosen halo population model (such as HOD) the toolkit generates galaxy positions that incorporate the full relativistic treatment of super-halo scales, while accounting for peculiar velocity

distortions of individual galaxies, on sub-halo scales. These contributions affect not only their line-of-sight positions but also the Doppler-induced contribution to magnification bias, as described in equations 15 and 17.

We designed this approach to align with the formalism of the HaloTools library, ensuring straightforward adaptability to a variety of halo population models⁹.

We applied our method to the HMDPL simulation, constructing a full-sky halo catalogue that reaches $z = 0.8$. We compared the LIGER4GAL implementation to the previous field-based one (LIGER4DM), and showed that we can extend the accuracy of the produced catalogues to capture the non-linearities of the halo clustering signal, where LIGER4DM fails. Having access to the non-linear scales allowed us to perform realistic measurements of clustering observables, such as the galaxy power spectrum, which we in turn used to perform standard cosmological inference analyses, to assess the impact of large-scale relativistic effects in the inference process.

Then, relying on the HOD formalism, we generated a DESI-like LRG sample that presents the imprint of relativistic RSDs at large scales. We utilised a full-sky and a partial-sky sample of that catalogue incorporated with different levels of relativistic RSDs to assess the impact of relativistic contributions which are generally not accounted for in the power spectrum models used in the standard full-shape analysis, as well as PNG inference via galaxy clustering.

In particular, we assessed the impact of the observer’s peculiar motion (the FOTO effect) on the inference of the f_{nl} parameter. By focusing directly on the bias Δf_{nl} caused by the FOTO, we were able to disentangle our results from other possible effects that could impact the analysis, such as the radial integral constraint [37], or exploring different values of the *true* underlying f_{nl} . We found that this contribution leads to a bias in the measurement of f_{nl} both for the full- and partial-sky surveys, which varies with the largest scale included in the analysis. In particular, we found that for a DESI-like (partial sky) analysis, the FOTO signal biases f_{nl} by more than $1\sigma_{f_{\text{nl}}}$ in approximately 40% of the realizations of the universe when including scales down to $k_{\text{min}} = 1.5 \times 10^{-3} h/\text{Mpc}$; this probability drops to 10% if scales below $k_{\text{min}} < 5 \times 10^{-3} h/\text{Mpc}$ are excluded. Full-shape analyses of the power spectrum monopole and quadrupole using the EFT model indicate that A_s , ω_c , and h are biased when including scales below $k_{\text{min}} < 4 \times 10^{-3} h/\text{Mpc}$, but this bias disappears if these large scales are excluded. Due to the size of available simulations, our analysis was limited to an LRG-like sample up to $z = 0.8$. Extending to higher-redshift populations requires larger simulations capable of producing wide light-cone catalogues. We note that the FOTO signal can vary substantially with survey selection, and consequently, the biases discussed here may be amplified for different populations or redshift ranges. Nonetheless, the toolkit introduced here can be applied directly to generate light-cones with relativistic RSDs, provided sufficiently large simulations with halo merger trees are available.

⁹ Available at: <https://gitlab.com/liger-cosmo/liger-hod>

Since upcoming measurements of f_{nl} will probe larger regions of the Universe and include increasingly larger scales (e.g. DESI Y5, Euclid DR3), the FOTO contribution will become progressively more significant. It is therefore important to incorporate this signal into both the power-spectrum model and the covariance matrices, along with other wide-angle systematics such as Milky Way extinction, and large-scale flux variations. Incorporating these effects into mock catalogues is relatively straightforward, allowing for a robust estimation of their impact on parameter constraints. Applying the same treatment in analytic models is more challenging, as it requires careful accounting of the mixing between scales and multipoles—including odd multipoles—induced by the survey geometry. Nevertheless, analytic approaches to the FOTO signal could offer the additional advantage of simultaneously constraining the observer’s velocity and PNG, or of marginalizing over the velocity parameters. We leave the development of such strategies to future work.

ACKNOWLEDGMENTS

We thank Benedict Bahr-Kalus, Pierluigi Monaco, Emiliano Sefusatti and Licia Verde for the useful suggestions and

discussion during the first distribution of this manuscript. DB acknowledges support from the COSMOS network (www.cosmosnet.it) through ASI (Italian Space Agency) Grants 2016-24-H.0, 2016-24-H.1-2018 and 2020-9-HH.0. We acknowledge the use of computational resources from the University of Padova Strategic Research Infrastructure Grant 2017: “CAPRI: Calcolo ad Alte Prestazioni per la Ricerca e l’Innovazione”. This paper is supported by the PRIN 2022 PNRR project “Space-based cosmology with Euclid: the role of High-Performance Computing” (code no. P202259YAF), funded by European Union – Next Generation EU. We acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) and the Partnership for Advanced Supercomputing in Europe (PRACE, www.prace-ri.eu) for funding the MultiDark simulation project by providing computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre (LRZ, www.lrz.de). The Bolshoi simulations have been performed within the Bolshoi project of the University of California High-Performance AstroComputing Center (UC-HiPACC) and were run at the NASA Ames Research Center.

-
- [1] Adame, A. G. et al. 2025, JCAP, 07, 017
 - [2] Adame, A. G. et al. 2025, JCAP, 09, 008
 - [3] Ade, P. A. R. et al. 2014, Astron. Astrophys., 571, A16
 - [4] Aghamousa, A. et al. 2016 [1611.00036]
 - [5] Aghanim, N. et al. 2020, Astron. Astrophys., 641, A1
 - [6] Alvarez, M. et al. 2014 [1412.4671]
 - [7] Alves, O. & DESI Collaboration. 2024, (in prep.)
 - [8] Bahr-Kalus, B., Bertacca, D., Verde, L., & Heavens, A. 2021, JCAP, 11, 027
 - [9] Bardeen, J. M. 1980, Phys. Rev. D, 22, 1882
 - [10] Barreira, A., Cabass, G., Schmidt, F., Pillepich, A., & Nelson, D. 2020, Journal of Cosmology and Astroparticle Physics, 2020, 013
 - [11] Bartolo, N., Komatsu, E., Matarrese, S., & Riotto, A. 2004, Physics Reports, 402, 103–266
 - [12] Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013, Astrophys. J., 762, 109
 - [13] Berlind, A. A. & Weinberg, D. H. 2002, Astrophys. J., 575, 587
 - [14] Bertacca, D. 2015, Classical and Quantum Gravity, 32, 195011
 - [15] Bertacca, D. 2020, International Journal of Modern Physics D, 29, 2050085
 - [16] Bertacca, D., Maartens, R., & Clarkson, C. 2014, Journal of Cosmology and Astroparticle Physics, 2014, 037–037
 - [17] Bertacca, D., Maartens, R., Raccañelli, A., & Clarkson, C. 2012, JCAP, 2012, 025
 - [18] Beutler, F. & McDonald, P. 2021, Journal of Cosmology and Astroparticle Physics, 2021, 031
 - [19] Beutler, F. et al. 2014, Mon. Not. Roy. Astron. Soc., 443, 1065
 - [20] Bianchi, D., Gil-Marín, H., Ruggeri, R., & Percival, W. J. 2015, Mon. Not. Roy. Astron. Soc., 453, L11
 - [21] Böhme, L., Schwarz, D. J., Tiwari, P., et al. 2025 [2509.16732]
 - [22] Bonvin, C. & Durrer, R. 2011, Phys. Rev. D, 84, 063505
 - [23] Borzyszkowski, M., Bertacca, D., & Porciani, C. 2017, Mon. Not. Roy. Astron. Soc., 471, 3899
 - [24] Breton, M.-A., Rasera, Y., Taruya, A., Lacombe, O., & Saga, S. 2019, Monthly Notices of the Royal Astronomical Society, 483, 2671, arXiv:1803.04294 [astro-ph]
 - [25] Broadhurst, T. J., Taylor, A. N., & Peacock, J. A. 1995, The Astrophysical Journal, 438, 49
 - [26] Bruni, M., Crittenden, R., Koyama, K., et al. 2012, Physical Review D, 85, 041301, arXiv:1106.3999 [astro-ph]
 - [27] Bruni, M., Crittenden, R., Koyama, K., et al. 2012, Phys. Rev. D, 85, 041301
 - [28] Bryan, G. L. & Norman, M. L. 1998, Astrophys. J., 495, 80
 - [29] Camera, S., Maartens, R., & Santos, M. G. 2015, Monthly Notices of the Royal Astronomical Society: Letters, 451, L80, arXiv:1412.4781 [astro-ph]
 - [30] Carrilho, P., Moretti, C., & Pourtsidou, A. 2023, Journal of Cosmology and Astroparticle Physics, 2023, 028
 - [31] Castorina, E. & Dio, E. d. 2022, Journal of Cosmology and Astroparticle Physics, 2022, 061, arXiv:2106.08857 [astro-ph]
 - [32] Castorina, E., Hand, N., Seljak, U., et al. 2019, Journal of Cosmology and Astroparticle Physics, 2019, 010
 - [33] Challinor, A. & Lewis, A. 2011, Phys. Rev. D, 84, 043516
 - [34] Chaussidon, E., Yèche, C., de Mattia, A., et al. 2024 [2411.17623]
 - [35] Cooray, A. & Sheth, R. 2002, Physics Reports, 372, 1–129
 - [36] Dalal, N., Doré, O., Huterer, D., & Shirokov, A. 2008, Physical Review D, 77
 - [37] de Mattia, A. & Ruhlmann-Kleider, V. 2019, JCAP, 08, 036
 - [38] Desjacques, V., Crocce, M., Scoccimarro, R., & Sheth, R. K. 2010, Physical Review D, 82
 - [39] Diemer, B. 2018, Astrophys. J. Suppl., 239, 35

- [40] Dio, E. D., Durrer, R., Marozzi, G., & Montanari, F. 2014, *Journal of Cosmology and Astroparticle Physics*, 2014, 017–017
- [41] Doré, O., Bock, J., Ashby, M., et al. 2015 [1412.4872]
- [42] d’Amico, G., Gleyzes, J., Kokron, N., et al. 2020, *Journal of Cosmology and Astroparticle Physics*, 2020, 005–005
- [43] Eggemeier, A., Camacho-Quevedo, B., Pezzotta, A., et al. 2022, *Monthly Notices of the Royal Astronomical Society*, 519, 2962
- [44] Elkhatab, M. Y., Bertacca, D., Porciani, C., et al. 2025, *Astronomy & Astrophysics*, 697, A85
- [45] Elkhatab, M. Y., Porciani, C., & Bertacca, D. 2021, *Mon. Not. Roy. Astron. Soc.*, 509, 1626
- [46] Elkhatab, M. Y., Porciani, C., & Bertacca, D. 2025, *JCAP*, 03, 044
- [47] Elkhatab, M. Y. et al. 2025, *Astron. Astrophys.*, 697, A85
- [48] Elsner, F., Leistedt, B., & Peiris, H. V. 2016, *Mon. Not. Roy. Astron. Soc.*, 456, 2095
- [49] Euclid Collaboration Monaco, P., Elkhatab, M. Y., Granett, B. R., et al. 2025, Euclid preparation. Controlling angular systematics in the Euclid spectroscopic galaxy sample
- [50] Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, *Astrophys. J.*, 426, 23
- [51] Foglieni, M., Pantiri, M., Di Dio, E., & Castorina, E. 2023, *Phys. Rev. Lett.*, 131, 111201
- [52] Frigo, M. & Johnson, S. 2005, *Proceedings of the IEEE*, 93, 216
- [53] Giannantonio, T. & Porciani, C. 2010, *Phys. Rev. D*, 81, 063530
- [54] Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *Astrophys. J.*, 622, 759
- [55] Guedezounme, S. L., Jolicoeur, S., & Maartens, R. 2025, *JCAP*, 07, 063
- [56] Hadzhiyska, B. et al. 2023, *MNRAS*, 525, 4367
- [57] Hearin, A. P. et al. 2017, *Astron. J.*, 154, 190
- [58] Ishiyama, T. et al. 2021, *Mon. Not. Roy. Astron. Soc.*, 506, 4210
- [59] Ivanov, M. M. 2022 [2212.08488]
- [60] Jeong, D., Schmidt, F., & Hirata, C. M. 2012, *Phys. Rev. D*, 85, 023504
- [61] Kaiser, N. 1987, *MNRAS*, 227, 1
- [62] Klypin, A., Yepes, G., Gottlöber, S., Prada, F., & Heß, S. 2016, *MNRAS*, 457, 4340
- [63] Komatsu, E. & Spergel, D. N. 2001, *Physical Review D*, 63
- [64] Kravtsov, A. V., Berlind, A. A., Wechsler, R. H., et al. 2004, *Astrophys. J.*, 609, 35
- [65] Lange, J. U. 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 3181
- [66] Leistedt, B. et al. 2016, *Astrophys. J. Suppl.*, 226, 24
- [67] Lepori, F., Schulz, S., Tutusaus, I., et al. 2025, *Astronomy & Astrophysics*, 694, A321
- [68] Lepori, F., Tutusaus, I., Viglione, C., et al. 2022, *A&A*, 662, A93
- [69] Lewis, A. 2025, *JCAP*, 08, 025
- [70] Lewis, A. & Challinor, A. 2011, *CAMB: Code for Anisotropies in the Microwave Background*, *Astrophysics Source Code Library*, record ascl:1102.026
- [71] Matarrese, S. & Verde, L. 2008, *The Astrophysical Journal*, 677, L77–L80
- [72] Matsubara, T. 2000, *The Astrophysical Journal*, 537, L77–L80
- [73] Matsubara, T. & Szalay, A. S. 2001, *The Astrophysical Journal*, 556, L67, arXiv:astro-ph/0105493
- [74] Maus, M., Lai, Y., Noriega, H. E., et al. 2024 [2404.07272]
- [75] Mead, A. J., Brieden, S., Tröster, T., & Heymans, C. 2021, *Monthly Notices of the Royal Astronomical Society*, 502, 1401–1422
- [76] Mellier, Y., Abdurro’uf, Acevedo Barroso, J. A., et al. 2025, *Astronomy & Astrophysics*, 697, A1
- [77] Michaux, M., Hahn, O., Rampf, C., & Angulo, R. E. 2020, *Mon. Not. Roy. Astron. Soc.*, 500, 663
- [78] Monaco, P., Paribelli, G., Elkhatab, M. Y., et al. 2025, arXiv e-prints, arXiv:2507.12116
- [79] Montano, F. & Camera, S. 2024, *Physics of the Dark Universe*, 46, 101570
- [80] More, S., Bosch, F. C. v. d., & Cacciato, M. 2009, *Mon. Not. Roy. Astron. Soc.*, 392, 917
- [81] Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, *Astrophys. J.*, 490, 493
- [82] Noorikuhani, M. & Scoccimarro, R. 2023, *Physical Review D*, 107, 083528, arXiv:2207.12383 [astro-ph]
- [83] Peacock, J. A., Cole, S., Norberg, P., et al. 2001, *Nature*, 410, 169–173
- [84] Peebles, P. J. E. 1973, *Astrophys. J.*, 185, 413
- [85] Prada, F., Klypin, A. A., Cuesta, A. J., Betancort-Rijo, J. E., & Primack, J. 2012, *Mon. Not. Roy. Astron. Soc.*, 423, 3018
- [86] Raccanelli, A., Bertacca, D., Dore, O., & Maartens, R. 2014, *Journal of Cosmology and Astroparticle Physics*, 2014, 022, arXiv:1306.6646 [astro-ph]
- [87] Raccanelli, A., Montanari, F., Bertacca, D., Doré, O., & Durrer, R. 2016, *Journal of Cosmology and Astroparticle Physics*, 2016, 009–009
- [88] Riebe, K. et al. 2013, *Astron. Nachr.*, 334, 691
- [89] Ross, A. J., Percival, W. J., Sánchez, A. G., et al. 2012, *Monthly Notices of the Royal Astronomical Society*, 424, 564–590
- [90] Sachs, R. K. & Wolfe, A. M. 1967, *Astrophys. J.*, 147, 73
- [91] Sasaki, M. 1989, *Mon. Not. Roy. Astron. Soc.*, 240, 415
- [92] Scoccimarro, R. 2015, *Physical Review D*, 92
- [93] Secrest, N. J., Hausegger, S. v., Rameez, M., et al. 2021, *The Astrophysical Journal Letters*, 908, L51
- [94] Shapiro, I. I. 1964, *Phys. Rev. Lett.*, 13, 789
- [95] Sheth, R. K. & Tormen, G. 1999, *Monthly Notices of the Royal Astronomical Society*, 308, 119–126
- [96] Simon, T., Zhang, P., Poulin, V., & Smith, T. L. 2023, *Physical Review D*, 107
- [97] Slosar, A., Hirata, C., Seljak, U., Ho, S., & Padmanabhan, N. 2008, *Journal of Cosmology and Astroparticle Physics*, 2008, 031
- [98] Tansella, V., Jelic-Cizmek, G., Bonvin, C., & Durrer, R. 2018, *Journal of Cosmology and Astroparticle Physics*, 2018, 032
- [99] Tinker, J. L., Robertson, B. E., Kravtsov, A. V., et al. 2010, *Astrophys. J.*, 724, 878
- [100] Umeh, O., Jolicoeur, S., Maartens, R., & Clarkson, C. 2017, *Journal of Cosmology and Astroparticle Physics*, 2017, 034–034
- [101] Wadekar, D. & Scoccimarro, R. 2020, *Phys. Rev. D*, 102, 123517
- [102] Wilson, M. J., Peacock, J. A., Taylor, A. N., & de la Torre, S. 2017, *MNRAS*, 464, 3121
- [103] Yamamoto, K., Nishioka, H., & Suto, Y. 1999, *The Astrophysical Journal*, 527, 488, arXiv:astro-ph/9908006
- [104] Yoo, J., Fitzpatrick, A. L., & Zaldarriaga, M. 2009, *Phys. Rev. D*, 80, 083514
- [105] Yoo, J. & Zaldarriaga, M. 2014, *Physical Review D*, 90
- [106] Yuan, S. et al. 2024, *Mon. Not. Roy. Astron. Soc.*, 530, 947
- [107] Yuan, S. et al. 2024, *Mon. Not. Roy. Astron. Soc.*, 530, 947
- [108] Zheng, Z., Berlind, A. A., Weinberg, D. H., et al. 2005, *Astrophys. J.*, 633, 791

Appendix A: Lightcone construction

In this appendix, we outline the procedure for constructing the halo lightcone and computing the gravitational potential at the relevant redshifts for the LIGER4GAL method.

1. Halo trajectories

To construct the observed past light cone from the simulation, we implement a cubic spline interpolation procedure for halo positions as they intersect with the observer’s lightcone. We first select an observer position at the $z = 0$ snapshot and assign it a peculiar velocity; if none is specified, the code adopts the velocity of the nearest halo to ensure consistency with the local environment. For each halo, we reconstruct its full trajectory across all simulation timesteps, starting from the earliest snapshots until a merger occurs or the trajectory is interrupted. Subsequently, we identify the instant at which the halo intersects the observer’s past light cone and apply the interpolation scheme of [45], while additionally tracking the halo mass evolution between snapshots.

Regarding the definition of halo trajectories, several caveats are considered to ensure a consistent catalogue. When a merger between two halos occurs at snapshot i (where $i = 0$ corresponds to $z = 0$ and $z_i > z_{i-1}$) we treat the Most Massive Progenitor (MMP) and the secondary merging halos differently. For the MMP, we continue to follow the evolution of its descendant in snapshot $i - 1$. For non-MMP halos, we add a final step to their trajectory in snapshot $i - 1$, keeping their mass fixed but assigning them the position and velocity of the descendant. Although the halo merger occurred in between the two snapshot, it is not possible to determine if it occurred before or after the lightcone intersection. In this work, we assume the former case, so that the two haloes intersect the lightcone as a single halo.

Another situation of interest occurs when halos without progenitors are first identified in a given snapshot i and do not intersect the light cone in subsequent snapshots. In such cases, no information is available on whether they may have crossed the light cone between snapshots i and $i + 1$ ($z_i < z_{i+1}$). This can happen either because the halo actually formed in snapshot i , or because the halo finder failed to correctly link it to its progenitor in snapshot $i + 1$. To address this issue, we initially excluded these halos; however, this approach systematically underestimated the radial density of halos with $M_{\text{vir}} \geq 5 \times 10^{12} M_{\odot}/h$ compared to snapshot-based estimates. The actual importance of this effect depends on the merger-tree definition of the simulation under consideration and may be less impactful in other simulations. Consequently, we opt to always account for the possibility of a light-cone intersection between i and $i + 1$. We do so by extrapolating the halo position at the snapshot $i + 1$ by assuming a motion with constant velocity $\mathbf{v}_i$, taken to be the halo’s velocity at the snapshot i . Under this prescription, the extrapolated position assigned




FIG. 13: In the upper panel, we compare the halo radial number density of the H-FS within the redshift bin $z \in [0.2, 0.7]$ (red dashed curve) to the values computed at each simulation snapshot (dotted markers). The lower panel shows the relative difference between these quantities.

to these halos is $\mathbf{x}_{i+1} = \mathbf{x}_i + \mathbf{v}_i \cdot (t_{i+1} - t_i)$, where t_i is the proper time of the snapshot. To validate the latter approach, we compare the radial densities of the H-FS with the radial densities obtained from the halo snapshots. The results are shown in Fig. 13. Each red dot corresponds to the radial density of a snapshot halo catalogue at the corresponding radial distance, and they agree (within a few percent) with the dashed curve, which represents the radial density of the H-FS catalogue.

Finally, we validate our halo mass interpolation procedure in Fig. 14, where we compare the halo mass function (HMF) obtained from a thin shell of the H-FS (black solid line) with the HMF estimated from a snapshot of the simulation at the mean redshift of the shell (gray dashed line). For comparison, we also plot the HMF model of Tinker et al. [99] (hereafter Tinker10) evaluated at the snapshot redshift. As evident in the figure, the LRG-FS1 HMF agrees with the snapshot HMF across the entire mass range. Both HMFs deviate from Tinker10 at the low-mass end due to the mass resolution limit of the HMDPL simulation.

2. Gravitational potential computation

The underlying principle of the LIGER method is that, in a Λ CDM cosmology at linear order, the Bardeen potentials (see Eq. 3) coincide, i.e. $\Phi = \Psi$, and can be computed by solving the Poisson equation, with the matter overdensity from the input simulations serving as the source term [see 23, 46]. As the HMDPL simulation provides only two publicly available particle snapshots, we used linear theory to model the time evolution of the gravitational potential, allowing us to derive the potential at the redshifts of the available halo snapshots. The procedure is as follows. First, we computed the potential at one of the matter snapshots using spectral methods imple-




FIG. 14: In the upper panel, we compare the HMF of H-FS lightcone (black solid) computed within a narrow redshift bin centred at $z = 0.46$ with thickness $\Delta z = 0.06$ to the snapshot HMF at $z = 0.46$ (grey dashed). We also plot the corresponding Tinker10 model prediction (red dashed line). In the bottom, we show ratio of the between the lightcone and snapshot HMFs.

mented with the FFTW library [52]. The potential can then be obtained at any scale factor a in a given grid cell located at position $\mathbf{x}_i$, we use the linear relation

$$\Phi(\mathbf{x}_i, a) = \frac{a_0 D^{(+)}(a)}{a D^{(+)}(a_0)} \Phi(\mathbf{x}_i a_0), \quad (\text{A1})$$

where $D^{(+)}$ is the linear growth factor of density perturbations and a_0 denotes the reference scale factor of one of the available snapshots.

We test the validity of the linear approximation using the simulation data by computing the gravitational potential from the two available particle snapshots on a grid with $N_{\text{grid}} = 256^3$, and comparing the average of their ratio with the theoretical prediction from Eq. (A1). The mean ratio, estimated over all grid cells where neither of the two potentials vanishes, is

$$\hat{r} \equiv \left\langle \frac{\Phi(z_1)}{\Phi(z_2)} \right\rangle \Big|_{\Phi(z_1), \Phi(z_2) \neq 0} = 1.15496, \quad (\text{A2})$$

while the theoretical prediction reads

$$r = \frac{D^{(+)}(a_1) a_2}{D^{(+)}(a_2) a_1} = 1.15477, \quad (\text{A3})$$

that implies a sub-percent fractional error of $(\hat{r} - r)/r = 0.016\%$.

We also test whether the linear solution accurately recovers the potential power spectrum, $P_\Phi(k)$, at different redshifts. In Fig. 15, we compare the power spectrum of the gravitational potential computed directly from the snapshot at redshift z_2 (black solid line) with the linear extrapolation (red dashed)




FIG. 15: Power spectrum comparison of the gravitational potential Φ obtained respectively from the particle distribution at the HMDPL snapshot $z = 0.49$ (black-continuous line) and from the linear theory extrapolation of equation A1, starting from the $z = 0.0$ snapshot (red-dashed line). In grey we show the Nyquist frequency of the measurement.

line) obtained using Eq. A1. The two estimates show excellent agreement across all scales up to the Nyquist frequency of the grid, $k_{\text{Ny}} \approx 0.20 h \text{ Mpc}^{-1}$. We therefore conclude that the linear approximation is sufficient for our purposes. Consequently, we extrapolate the potential to all redshifts by utilizing the potential at $z = 0$ and applying Eq. A1. Finally, we note that this assumption is not strict: the potential extrapolation is not required for the implementation of our approach and can be relaxed when full particle snapshots (or sufficiently large subsamples) are available at each redshift.

Appendix B: Survey functions numerical estimation

In this appendix, we briefly describe the procedure implemented to estimate the survey functions of the H-FS, LRG-FS, and LRG-DESI. These survey functions [precisely $\bar{n}(z)$, $b_1(z)$, $\mathcal{E}(z)$ and $Q(z)$]. In particular, for the H-FS catalogue, we compute the survey functions from the simulation snapshots, while for the LRG-FS and LRG-DESI sample we work directly with the lightcone data.

1. H-FS Catalogue

As the halos are available in snapshot format, we can compute the survey functions directly at each redshift from the snapshots. We compute the $\bar{n}(z_i)$ (where i runs over the different snapshots) simply by counting the haloes belonging to the sample and dividing by the simulation comoving volume. For the linear bias, we compute from the power spectra ratio. For a given snapshot at a redshift z_i , and a sample of mass-selected halos, we use the linearly evolved matter power

spectrum $P_m(k_j, z_i)$ (see appendix A 2) and estimate the halo power spectrum $P_m(k_j, z_i)$, where k_j are the discrete wavenumbers at which the power spectrum is sampled. Then, we estimate the linear halo bias using an average of the power spectra ratio

$$b_1(z_i) = \frac{1}{N_k} \sum_{k_j}^{k_j < k_0} \sqrt{\frac{P_h(k_j, z_i)}{P_m(k_j, z_i)}} \quad (\text{B1})$$

where k_0 is an upper limit for k that we choose so that the linear bias relation holds and $N_k = \sum_{k_j}^{k_j < k_0}$.

For the magnification and evolution bias, we use [see appendix B in 45]

$$-\frac{d \ln \bar{n}(z)}{d \ln(1+z)} = 2Q(z) \left[1 + \frac{(1+z)}{H(z)x(z)} \right] + \mathcal{E}(z). \quad (\text{B2})$$

As for the halo sample $Q(z) = 0$ by construction, we directly note that

$$\mathcal{E}(z) = -\frac{d \ln \bar{n}(z)}{d \ln(1+z)}. \quad (\text{B3})$$

2. LRG-FS catalogue

We use an analogous approach to construct the survey functions for the LRG-FS and LRG-DESI catalogues, albeit applied directly on the lightcone, since the galaxies are painted directly onto the lightcone halo catalogue. To compute $\bar{n}$, we count the galaxies in redshift bins of size $\delta z = 0.04$, and divide by the comoving volume associated with each corresponding spherical shell.

For the linear bias, we use the angular power spectra ratio. Specifically, we compute the ratio between the angular power spectra of the LRG-FS- $\mathcal{R}$ catalogues in tomographic redshift bins, $C_\ell(z_i)$ (where z_i is the median redshift of the bin), and a theoretical linear angular power spectrum computed using CAMB. The linear bias is then obtained by averaging over a range of multipoles:

$$b_1(z_i) = \frac{1}{N_\ell} \sum_{\ell}^{\ell < \ell_0} \sqrt{\frac{C_\ell^{\mathcal{R}}}{C_\ell^{\text{Linear}}}}, \quad (\text{B4})$$

where ℓ_0 is the maximum multipole chosen such that the linear biasing relation holds, and $N_\ell = \sum_{\ell}^{\ell < \ell_0}$.

For the magnification bias, we use the estimates from Zhou et al. [109], where the complex selection function of the LRGs is already taken into account. We interpolate their measurements to each galaxy's redshift.

There is, however, a slight subtlety in the computation of the evolution bias. Since the LRG galaxies are selected using flux cuts across several bands, we account for magnification effects through a weighting scheme (see Eq. 18). This implies that the logarithmic derivative in Eq. B2 must be applied to $\bar{n}_g = \bar{w}_g(z) \bar{n}_g(z)$, where $\bar{n}_g(z)$ denotes the average density of the catalogues without the magnification weights, and $\bar{w}_g(z)$ is

the average value of the weight $w_g(\mathbf{x})$ at a given redshift. Consequently, we estimate $\mathcal{E}(z)$ from the logarithmic total derivative of the galaxy number density *before* applying the magnification weights.

Appendix C: Estimators

In this appendix, we review the estimators utilised the two-point statistics studied in this work.

1. Angular power spectrum

For the estimation the angular power spectrum, we make use of the HEALPix algorithm [Hierarchical Equal Area iso-Latitude Pixelisation 54] to partition the sky into $N_{\text{pix}} = 12 \times N_{\text{side}}^2$ pixels. For each tracer sample i , we compute the projected counts within a redshift bin, $N_g^i(\Omega)$, and then estimate the projected density contrast as

$$\Sigma^i(\Omega) = \frac{N_g^i(\Omega)}{\bar{N}_g^i} - 1, \quad (\text{C1})$$

where $\bar{N}_g^i$ is the average of the projected counts over the survey footprint. Then, exploiting the fast spherical harmonic decomposition of HEALPix, we compute

$$a_{\ell,m}^i = \int \Sigma^i(\Omega) Y_{\ell m}^*(\Omega) d^2\Omega, \quad (\text{C2})$$

which we use to estimate the auto and cross angular power spectra with the pseudo- C_ℓ estimator [84]

$$C_\ell^{ij} = \frac{1}{w_p^2(2\ell+1)f_{\text{sky}}} \sum_{m=-\ell}^{\ell} a_{\ell m}^i a_{\ell m}^{j*} - \frac{\delta_K^{i,j}}{\bar{N}_g^i}, \quad (\text{C3})$$

where f_{sky} is the sky fraction covered by the survey footprint, and w_p is a pixelisation correction factor.

2. 3D power spectrum

For the (3D) power spectrum multipoles, we first construct the FKP field [50]:

$$F(\mathbf{x}) = \frac{w(\mathbf{x})}{\sqrt{A}} \left[n_g(\mathbf{x}) - \alpha n_r(\mathbf{x}) \right], \quad (\text{C4})$$

where $A = \int w^2(\mathbf{x}) \bar{n}^2 d^3x$ is a normalization factor, and the fields n_g & n_r denote the number density of the data and random catalogues, respectively. Moreover,

$$\alpha = \frac{\int w(\mathbf{x}) n_g(\mathbf{x}) d^3x}{\left[\int w(\mathbf{x}) n_r(\mathbf{x}) d^3x \right]} \quad (\text{C5})$$

rescales the random density to match the data catalogue number densities. Finally, the FKP weights w are given by

$$w(\mathbf{x}) = \mathcal{I}(\mathbf{x}) [1 + \bar{n}(\mathbf{x}) \mathcal{P}_0]^{-1}, \quad (\text{C6})$$

where $\bar{n}(x)$ is the mean density of tracers in our catalogue at a distance x , $\mathcal{I}$ is the indicator function, equal to 1 inside the survey volume and to 0 elsewhere, and the parameter $\mathcal{P}_0$ represents an approximate value of the power spectrum at the scales under study. In this work, we estimate $\bar{n}(x)$ from our random catalogue, evaluating it within radial bins of thickness $\delta r = 15 \text{ Mpc}/h$ and interpolating it with a cubic spline. For each of our samples, we generate random catalogues 50× as dense as our data ($\alpha \approx 1/50$), sampling their distribution using the shuffling method, which assigns randomly selected redshifts from the original dataset and uniformly distributed angular positions within the survey mask. For the typical value of the power spectrum, we adopt $\mathcal{P}_0 = 10\,000 \text{ Mpc}^3 h^{-3}$.

To compute the power spectrum, we use the `pypower` code, employed in recent DESI data analyses. This implementation follows the FFT-based Yamamoto–Bianchi formalism [20, 92, 103]:

$$P_\ell(k) = (2\ell + 1) \int F_\ell(\mathbf{k}) F_0(-\mathbf{k}) \frac{d^2\Omega_k}{4\pi} - P_\ell^{\text{SN}}(k), \quad (\text{C7})$$

where

$$F_\ell(\mathbf{k}) = \int F(\mathbf{x}) e^{-i\mathbf{k}\cdot\mathbf{x}} \mathcal{L}_\ell(\hat{\mathbf{k}} \cdot \hat{\mathbf{x}}) d^3x, \quad (\text{C8})$$

and

$$P_\ell^{\text{SN}}(k) = \frac{(1 + \alpha)}{A} \int w^2(\mathbf{x}) n_g(\mathbf{x}) \mathcal{L}_\ell(\hat{\mathbf{k}} \cdot \hat{\mathbf{x}}) d^3x. \quad (\text{C9})$$

Lastly, for all likelihood evaluations in the various parameter inference tests, we estimate the covariance matrices of the power spectra following the analytical approach of Wadekar & Scoccimarro [101]. In particular, using the `THECOV` code [7], we compute the covariance matrices for our power spectrum measurements, including the effects of the survey geometry. For simplicity, we only consider the Gaussian component of the covariance. We compute the covariance matrices for both the LRG-FS and LRG-DESI catalogues. We note that these covariance matrices do not account for the imprint of relativistic RSDs present in our dataset as the aim of our analysis is to investigate the biases that arise when neglecting such effects.