

***Euclid*: Field-level inference of primordial non-Gaussianity and cosmic initial conditions★**

A. Andrews^{1,2,★}, J. Jasche^{3,4}, G. Lavaux⁵, F. Leclercq⁵, F. Finelli^{1,2}, Y. Akrami^{6,7}, M. Ballardini^{8,1,9},
 D. Karagiannis^{10,11}, J. Valiviita^{12,13}, N. Bartolo^{14,15,16}, G. Cañas-Herrera^{17,18}, S. Casas^{19,20},
 B. R. Granett²¹, F. Pace^{22,23,24}, D. Paoletti^{1,2}, N. Porqueres²⁵, Z. Sakr^{26,27,28}, D. Sapone²⁹,
 N. Aghanim³⁰, A. Amara³¹, S. Andreon²¹, C. Baccigalupi^{32,33,34,35}, M. Baldi^{36,1,37}, S. Bardelli¹,
 D. Bonino²⁴, E. Branchini^{38,39,21}, M. Brescia^{40,41,42}, J. Brinchmann^{43,44}, S. Camera^{22,23,24},
 V. Capobianco²⁴, C. Carbone⁴⁵, J. Carretero^{46,47}, M. Castellano⁴⁸, G. Castignani¹, S. Cavuoti^{41,42},
 A. Cimatti⁴⁹, C. Colodro-Conde⁵⁰, G. Congedo⁵¹, C. J. Conselice⁵², L. Conversi^{53,54}, Y. Copin⁵⁵,
 F. Courbin^{56,57,58}, H. M. Courtois⁵⁹, A. Da Silva^{60,61}, H. Degaudenzi⁶², G. De Lucia³³,
 A. M. Di Giorgio⁶³, J. Dinis^{60,61}, F. Dubath⁶², C. A. J. Duncan⁵², X. Dupac⁵⁴, S. Dusini¹⁵, M. Farina⁶³,
 S. Farrens⁶⁴, F. Faustini^{65,48}, S. Ferriol⁵⁵, M. Frailis³³, E. Franceschi¹, S. Galeotta³³, B. Gillis⁵¹,
 C. Giocoli^{1,37}, P. Gómez-Alvarez^{66,54}, A. Grazian¹⁶, F. Grupp^{67,68}, S. V. H. Haugan⁶⁹, W. Holmes⁷⁰,
 F. Hormuth⁷¹, A. Hornstrup^{72,73}, P. Hudelot⁵, S. Ilić^{74,27}, K. Jahnke⁷⁵, M. Jhabvala⁷⁶, B. Joachimi⁷⁷,
 E. Keihänen⁷⁸, S. Kermiche⁷⁹, A. Kiessling⁷⁰, B. Kubik⁵⁵, M. Kunz⁸⁰, H. Kurki-Suonio^{12,13}, S. Ligi²⁴,
 P. B. Lilje⁶⁹, V. Lindholm^{12,13}, I. Lloro⁸¹, E. Maiorano¹, O. Mansutti³³, O. Marggraf⁸², K. Markovic⁷⁰,
 M. Martinelli^{48,83}, N. Martinet⁸⁴, F. Marulli^{85,1,37}, R. Massey⁸⁶, E. Medinaceli¹, S. Mei⁸⁷, Y. Mellier^{3,5},
 M. Meneghetti^{1,37}, E. Merlin⁴⁸, G. Meylan⁵⁶, M. Moresco^{85,1}, L. Moscardini^{85,1,37}, C. Neissner^{88,47},
 S.-M. Niemi¹⁷, J. W. Nightingale⁸⁹, C. Padilla⁸⁸, S. Paltani⁶², F. Pasian³³, K. Pedersen⁹⁰, V. Pettorino¹⁷,
 S. Pires⁶⁴, G. Polenta⁶⁵, M. Poncet⁹¹, L. A. Popa⁹², L. Pozzetti¹, F. Raison⁶⁷, R. Rebolo^{50,93,94},
 A. Renzi^{14,15}, J. Rhodes⁷⁰, G. Riccio⁴¹, E. Romelli³³, M. Roncarelli¹, R. Saglia^{68,67}, A. G. Sánchez⁶⁷,
 B. Sartoris^{68,33}, M. Schirmer⁷⁵, P. Schneider⁸², T. Schrabbach⁹⁵, A. Secroun⁷⁹, E. Sefusatti^{33,32,34},
 S. Serrano^{96,97,98}, C. Sirignano^{14,15}, G. Sirri³⁷, L. Stanco¹⁵, J. Steinwagner⁶⁷, P. Tallada-Crespí^{46,47},
 A. N. Taylor⁵¹, I. Tereno^{60,99}, R. Toledo-Moreo¹⁰⁰, F. Torradeflot^{47,46}, I. Tutusaus²⁷, L. Valenziano^{1,2},
 T. Vassallo^{68,33}, G. Verdoes Kleijn¹⁰¹, A. Veropalumbo^{21,39,102}, Y. Wang¹⁰³, J. Weller^{68,67}, G. Zamorani¹,
 E. Zucca¹, C. Burigana^{104,2}, V. Scottez^{3,105}, A. Spurio Mancini^{106,107}, and M. Viel^{32,33,35,34,108}

(Affiliations can be found after the references)

Received 17 January 2025 / Accepted 2 February 2026

ABSTRACT

A primary target of the *Euclid* space mission is to constrain early-universe physics by searching for deviations from a primordial Gaussian random field. A significant detection of primordial non-Gaussianity would rule out the simplest models of cosmic inflation and transform our understanding of the origin of the Universe. This paper forecasts how well field-level inference of galaxy redshift surveys can constrain the amplitude of local primordial non-Gaussianity, $f_{\text{NL}}^{\text{local}}$, within a Bayesian hierarchical framework, in the upcoming *Euclid* data. We designed and simulated mock datasets and performed Markov chain Monte Carlo analyses using a full-field forward modelling approach. By including the formation history of the cosmic matter field in the analysis, the method takes into account all available probes of primordial non-Gaussianity, and goes beyond statistical summary estimators of $f_{\text{NL}}^{\text{local}}$. Probes include, for example, two-point and higher-order statistics, peculiar velocity fields, and scale-dependent galaxy biases. Furthermore, the method simultaneously handles systematic survey effects, such as selection effects, survey geometries, and galaxy biases. The forecast shows that, using simulated *Euclid* data, the method can achieve a precision of $\sigma(f_{\text{NL}}^{\text{local}}) = 2.6$ (68.3% confidence level), assuming a grid resolution of $\Delta L = 31.25 h^{-1} \text{ Mpc}$ and a cut-off scale of $k_{\text{NF}} = 0.1 h \text{ Mpc}^{-1}$. We also provide data products, including realistic simulations with non-zero values of $f_{\text{NL}}^{\text{local}}$ and maps of adiabatic curvature fluctuations. The results underscore the feasibility and advantages of field-level inference to constrain $f_{\text{NL}}^{\text{local}}$ in galaxy redshift surveys. Our approach consistently captures all the information available in the large-scale structure to constrain $f_{\text{NL}}^{\text{local}}$, and resolves the degeneracy between early-universe physics and late-time gravitational effects, while mitigating the impact of systematic and observational effects.

Key words. methods: data analysis – methods: statistical – cosmological parameters – early Universe – large-scale structure of Universe – inflation

★ This paper is published on behalf of the Euclid Consortium.

★★ Corresponding author: adam.andrews@inaf.it

1. Introduction

One of the major tasks of modern cosmology is to understand the origin of the cosmic structure and the nature of the physical processes that governed the beginning of our Universe (Bartolo et al. 2004; Biagetti 2019; Achúcarro et al. 2022). The current canonical mechanism, cosmic inflation, generates quantum fluctuations from one or more quantum fields. These fields drove an epoch of quasi-exponential cosmic expansion at the beginning of the Universe (Starobinsky 1980; Guth 1981). Standard inflationary theory predicts primordial fluctuations as adiabatic, (almost) Gaussian, and (nearly) scale invariant. Examples include single-field slow-roll models with quantum vacuum fluctuations as initial conditions, which induce tiny departures from Gaussianity (Salopek & Bond 1990; Gangui et al. 1994; Acquaviva et al. 2003; Maldacena 2003; Creminelli & Zaldarriaga 2004; Byrnes et al. 2010; Creminelli et al. 2011; Baldauf et al. 2011b). To test these predictions, ongoing and upcoming cosmological surveys aim to further constrain early-universe physics by searching for deviations from a primordial Gaussian random field, in particular of the local type (LSST Science Collaborations 2009; Doré et al. 2014; Amendola et al. 2018; Euclid Collaboration: Blanchard et al. 2020). A significant detection of such a signal would radically transform our understanding of the early Universe, as it would hint toward more complex inflationary models, involving multiple fields (see e.g. Chen 2010; Komatsu 2010; Alvarez et al. 2014; Finelli et al. 2018; Celoria & Matarrese 2018; Meerburg et al. 2019).

The deviations from a primordial Gaussian random field are described by primordial non-Gaussianity (PNG). Potential sources of PNG include nonlinearity of gravity during inflation, inflation self-interactions, and additional, yet unknown, light or heavy quantum fields, with various models predicting higher levels of PNG with respect to the standard single-field slow-roll models (Falk et al. 1993; Gangui et al. 1994; Maldacena 2003; Bartolo et al. 2004; Chen 2010; Byrnes & Choi 2010; Arkani-Hamed & Maldacena 2015; Meerburg et al. 2019; Chen et al. 2022; Green et al. 2024). To the lowest order, local PNG is parameterised by the nonlinearity parameter $f_{\text{NL}}^{\text{local}}$ (see Eq. (3) for the definition). The perturbation of PNG induces a global rescaling of the primordial gravitational potential, leading to a multitude of effects and probes that can be used to measure $f_{\text{NL}}^{\text{local}}$ (Scoccimarro 2000; Komatsu & Spergel 2001; Verde et al. 2001; Scoccimarro et al. 2004; Komatsu et al. 2009; Byrnes et al. 2009; Chen 2010; Biagetti 2019). A subset of these probes has been used in observations of the cosmic microwave background (CMB; Planck Collaboration XXIV 2014a; Planck Collaboration XXII 2014b; Planck Collaboration XX 2016b; Planck Collaboration XVII 2016a; Planck Collaboration VI 2020c; Planck Collaboration X 2020a; Planck Collaboration IX 2020b), and the cosmic large-scale structures (LSS; Castorina et al. 2019; Mueller et al. 2021; D’Amico et al. 2025; Cabass et al. 2022) to constrain PNG.

To date, the tightest constraint on $f_{\text{NL}}^{\text{local}}$ (hereafter written as f_{NL}) has been obtained by CMB observations of the *Planck* satellite, which finds $f_{\text{NL}} = -0.9 \pm 5.1$ at a 68.3% confidence interval (CI, Planck Collaboration IX 2020b). Although the CMB is a powerful cosmological probe, its information content on PNG has been depleted, because large-scale temperature observations are expected to have reached the cosmic-variance limit¹.




Fig. 1. Inference power of BORG on $f_{\text{NL}}^{\text{local}}$ in mock *Euclid* data, compared to the inferred value of *Planck* (Planck Collaboration IX 2020b), and current state-of-the-art constraints from the large-scale structure (Cagliari et al. 2024; Chaussidon et al. 2025b). The BORG results are based on the main results in this paper and indicate that field-level inference will be able to provide state-of-the-art constraints on $f_{\text{NL}}^{\text{local}}$ with *Euclid* data.

In contrast, next-generation three-dimensional galaxy surveys can provide additional information by covering large cosmic volumes probing the largest scales of the cosmic matter distribution (Biagetti 2019; McQuinn 2021; Achúcarro et al. 2022).

Among the tightest constraints on local PNG from the LSS are the results of the Dark Energy Spectroscopic Instrument (DESI), the first data release, which finds $f_{\text{NL}} = -3.6^{+9.0}_{-9.1}$ (with $p = 1.6$ for the quasar sample Chaussidon et al. 2025b). For quasars only, a well-studied dataset is the SDSS-IV/eBOSS catalogue (DR 16, quasar sample Mueller et al. 2021; Cagliari et al. 2024). For this dataset, Mueller et al. (2021) find $|f_{\text{NL}}| < 21$ at 68.3% CI ($p = 1.6$), and Cagliari et al. (2024) find $-4 < f_{\text{NL}} < +27$ at 68.3% CI ($p = 1.0$). Another example is Castorina et al. (2019), who measured an earlier data release (DR14, quasar sample) to constrain f_{NL} to $-26 < f_{\text{NL}} < +14$ at 68.3% CI ($p = 1.0$, where p is the assembly bias parameter which determines the amplitude of the scale-dependent bias effect; see Sect. 2.3 for a detailed discussion).

As a side note, the two-dimensional CMB information can be cross-correlated with the three-dimensional cosmic LSS (Giannantonio & Percival 2014; Euclid Collaboration: Ilić et al. 2022; McCarthy et al. 2023; Krolewski et al. 2024). In this way, PNG can be measured without cosmic variance (Seljak 2009; Schmittfull & Seljak 2018; Ballardini et al. 2019; Barreira & Krause 2023; Karagiannis et al. 2024; Barberi Squarotti et al. 2024; Bermejo-Climent et al. 2025), with forecasts that demonstrate an improvement large enough to reduce the error to the above-mentioned science goal (Münchmeyer et al. 2019).

So far, current LSS analyses have been based on statistical estimators that are sensitive to two- and three-point correlation functions. These estimators are thus sensitive to the large-scale effect on the power spectrum and the small-scale effect on the bispectrum. The scale-dependent bias effect depends on the initial bispectrum that, in the local model, has a large signal in squeezed configurations that correlates large scales with small scales responsible for halo formation (Dalal et al. 2008;

expected to reach the science target of $\sigma(f_{\text{NL}}) = 1$ (Meerburg et al. 2019).

¹ Some improvements may arise from small-scale polarisation measurements (Baumann et al. 2009; Abazajian et al. 2016; Duivenvoorden et al. 2020; Kalaja et al. 2021), but these are not




Fig. 2. Flow chart illustrating the forward model implemented in BORG (Jasche & Wandelt 2013; Andrews et al. 2023). The forward model connects a set of initial conditions to a model prediction. This output can then be compared to the data at the field level through a likelihood evaluation. The parameter under each box represents the output of the box and what is provided to the next step of the forward model. The parameters above some boxes represent the additional input of each computation. We especially highlight the inclusion of the $f_{\text{NL}}^{\text{local}}$ parameter in the primordial gravitational potential and in the bias model.

LoVerde et al. 2008; Matarrese & Verde 2008; Carbone et al. 2008; Verde & Matarrese 2009; Scoccimarro et al. 2012). As a result, the scale-dependent bias effect is the most informative and crucial probe for measuring f_{NL} with galaxy surveys (Uhlemann et al. 2018; Mueller et al. 2019; Karagiannis et al. 2021; Biagetti et al. 2022; Giri et al. 2025; Riquelme et al. 2023; Brown et al. 2025; Peron et al. 2024; Yip et al. 2024; Jung et al. 2024). However, it is heavily affected by large-scale contamination and systematic survey effects (Leistedt et al. 2014; Rezaie et al. 2021; Mueller et al. 2021; Rezaie et al. 2024; Chaussidon et al. 2025a), which require the application of data cleaning techniques to provide unbiased measurements of f_{NL} . On the other hand, the bispectrum of galaxies is sensitive to small-scale effects from the imprint of PNG onto the matter field (Baldauf et al. 2011a; Karagiannis et al. 2018; Friedrich et al. 2020; Goldstein et al. 2022, 2024; Chen et al. 2025). However, these primordial perturbations are intertwined with the effects from late-time structure formation, meaning that the signal of interest is non-trivial to disentangle in the data. Interestingly, the bispectrum can decipher the different shapes of PNG, for example, the local, equilateral, and orthogonal modes. We refer to the literature for more information on these modes (Babich et al. 2004; Scoccimarro et al. 2012; Regan et al. 2012; Planck Collaboration XXIV 2014a; Schmidt et al. 2015; Planck Collaboration XVII 2016a; Karagiannis et al. 2018; Planck Collaboration IX 2020b).

Ongoing and upcoming galaxy redshift surveys such as *Euclid* (Laureijs et al. 2011; Amendola et al. 2018; Euclid Collaboration: Blanchard et al. 2020; Euclid Collaboration: Scaramella et al. 2022; Euclid Collaboration: Ballardini et al. 2024; Euclid Collaboration: Mellier et al. 2025; Euclid Collaboration: Finelli et al. 2026) are designed with the aim, among others, of probing early-universe physics by providing an unprecedented amount of data. However, data from these missions are expected to be affected by the systematic and observational effects of the survey (Graham et al. 2018). Consequently, strategies for mitigating and modelling these effects, such as instrumentation noise and astrophysical contamination, are crucial for handling the largest scales. Neglecting to address these factors can introduce biases in the results, particularly the constraints of f_{NL} (Huterer et al. 2013; Leistedt et al. 2014; Ho et al. 2015; Jasche & Lavaux 2017, 2019).

To solve the open issues mentioned above and go beyond statistical summary estimators, we applied a Bayesian field-level inference approach to constrain PNG. This method uses a forward modelling approach to connect the primordial matter fluctuations of the early Universe with the three-dimensional galaxy distribution at the field level (Jasche & Wandelt 2013). In this way, we leverage the complete formation history of the Universe in a holistic manner to constrain f_{NL} in the observed data. With this novel approach, we provide a complementary and independent method for optimally measuring deviations from Gaussian initial conditions. In fact, this method allows us, among other things, to simultaneously

– Naturally incorporate all probes of PNG in the LSS²;
 – Disentangle the late-time effects of nonlinearities from early-universe signals of PNG;
 – Marginalise unknown nuisance parameters and large-scale foreground contamination³;
 – Take into account survey systematic effects, for example, survey geometry and instrumentation noise.

In Andrews et al. (2023), we developed a field-level inference method to measure f_{NL} in galaxy surveys. We forecasted that for a simplified Stage IV survey, the method can achieve up to $\sigma(f_{\text{NL}}) = 5.7$, with $\sigma(f_{\text{NL}})$ denoting the 68.3% CI.

In this paper, we improve on our previous work in several aspects. We apply a more advanced forward model, which goes beyond the forward model used in Andrews et al. (2023). This includes a more realistic structure formation model and additional higher-order bias terms, to account for finer small-scale physics in galaxy formation. In addition, more realistic survey specifications of *Euclid* are incorporated, leading to more realistic survey aspects than in the previous work. This includes, for example, a more accurate survey mask, a radial selection function, and galaxy number counts. The combination of these modifications leads to generally improved forecasts of f_{NL} inference of spectroscopic data from the *Euclid* space telescope. Our preliminary results indicate that our method can constrain f_{NL} to the order of $\sigma(f_{\text{NL}}) = 2.6$ (68.3% CI) for a realistic *Euclid* survey, as can be seen in Fig. 1.

The paper is structured as follows. In Sect. 2, we provide an overview of the method and the algorithmic design choices. Emphasis has been placed on the galaxy bias model employed, the scale-dependent bias model, and on how the adiabatic curvature fluctuation maps were generated. In Sect. 3, we provide descriptions of the data generation setup and the datasets generated. Specifications for *Euclid* mock data and objectives of the *Euclid* mission with respect to PNG are included here.

² Examples include scale-dependent bias, higher-order statistics in the density field, mass distribution of tracers, and velocity field imprints.

³ See, for example, Porqueres et al. (2019b) and Lavaux et al. (2019).

We finalise the section with the list of tests included in this project. We present the results in Sect. 4, which cover convergence results, results on f_{NL} , and test-specific results. Finally, we summarise in Sect. 5, and discuss future work in Sect. 6.

2. Method

The primary goal of this paper is to forecast how well field-level inference can jointly constrain f_{NL} and cosmic initial conditions in spectroscopic *Euclid* mock data. In this context, field-level inference uses the entirety of the 3D cosmic matter field and its formation history to optimally extract the available information from the data to constrain cosmology, such as primordial fluctuations or cosmological parameters (Jasche & Wandelt 2013; Leclercq & Heavens 2021; Baumann & Green 2022; Nguyen et al. 2021, 2024; Beyond-2pt Collaboration 2025). This is achieved by constructing a data model that forward models an arbitrary set of initial conditions, so that the corresponding predicted galaxy field can be directly compared with the data at the field level (Jasche et al. 2010; Jasche & Wandelt 2013; Seljak et al. 2017; Schmidt et al. 2019; Kostić et al. 2022; Porqueres et al. 2022; Andrews et al. 2023; Porqueres et al. 2023; Bayer et al. 2023; Stopyra et al. 2024). To test the performance of field-level inference, we set up a series of mock datasets, all emulating the survey features of the *Euclid* mission, including selection effects, window function, and galaxy number counts. Then, we analysed the mock datasets by running Markov chain Monte Carlo (MCMC) analyses in a Bayesian hierarchical framework. The MCMC runs produced plausible values of the primordial matter fluctuations, f_{NL} , and marginalised nuisance parameters, all conditioned on the observed data. These outputs constitute the main results to make the said forecasts on f_{NL} and the cosmic initial conditions.

In this section, we provide a more detailed description of the field-level inference algorithm used (Sect. 2.1). Next, we describe the applied PNG model (Sect. 2.2), and the galaxy bias model used (Sect. 2.3). In the following, we provide a brief description of how we ran the MCMC analyses. Finally, we outline the generation of inferred 3D maps of adiabatic curvature fluctuations (Sect. 2.5), which is a new data product presented for the first time, to our knowledge, in this paper.

2.1. Overview of BORG

The Bayesian Origin Reconstruction from Galaxies (BORG) algorithm is a Bayesian hierarchical inference framework and is designed for the analysis of cosmic structure in cosmological surveys through forward modelling of three-dimensional galaxy fields (Jasche & Wandelt 2013; Jasche et al. 2015; Lavaux & Jasche 2016; Jasche & Lavaux 2019; Lavaux et al. 2019). The forward model in BORG aims to recreate the physical process in which the galaxies formed and observed, as closely as possible to the underlying physical process. In other words, the data model connects the three-dimensional primordial matter field to the observed distribution of galaxies, effectively reformulating the inverse problem of inferring the initial conditions into a statistical forward problem. In practice, we still have to simplify the model that provides the local galaxy abundances through statistical mapping between the matter field and the galaxy distribution. Thus, the objective of BORG is, given the assumed forward model, to explore the joint posterior distribution of initial conditions (denoted as ϵ), cosmological parameters, and nuisance parameters, as constrained by the observed data. The problem can be formulated in the form of a joint pos-




Fig. 3. Flow chart depicting the sampling scheme in BORG. The sampling scheme can be divided into four major parts: Sampling the 3D initial conditions, sampling the galaxy bias parameters (one set for each tracer catalogue), sampling the f_{NL} parameter, and finally saving the data products and restarting the cycle. Each sub-box depicts the conditional posterior from which the sample is drawn, and the sampling technique that is used.

terior distribution $\mathcal{P}_{\text{post}}(\epsilon, f_{\text{NL}}, \{b_i^g\} | N_g^O)$:

$$\mathcal{P}_{\text{post}}(\epsilon, f_{\text{NL}}, \{b_i^g\} | N_g^O) \propto \mathcal{P}_f(f_{\text{NL}}) \mathcal{P}_\epsilon(\epsilon) \mathcal{P}_b(\{b_i^g\}) \mathcal{P}_{\text{like}}(N_g^O | \epsilon, f_{\text{NL}}, \{b_i^g\}), \quad (1)$$

where N_g^O are the observed data in the form of galaxy counts, $\{b_i^g\}$ are the bias parameters, and $\mathcal{P}_f, \mathcal{P}_\epsilon, \mathcal{P}_b$ are the prior distributions. The prior in the white-noise field $\mathcal{P}_\epsilon(\epsilon)$ is a Gaussian with zero mean and unit standard variance. The likelihood distribution $\mathcal{P}_{\text{like}}(N_g^O | \epsilon, f_{\text{NL}}, \{b_i^g\})$ is defined by the data model.

The data model forward evolves a set of initial conditions to the corresponding predicted galaxy field, in the form of galaxy number counts. The forward model starts by simulating the gravitational progression of the matter field over time through a structure formation model. This process yields a predicted realisation of the late-time dark matter field, which is populated using a galaxy bias model (Sect. 2.3, Andrews et al. 2023). The window function and survey selection effects are taken into account by evaluating the survey footprint and radial selection functions at each voxel. We note that while the sky map in this work consists of a binary selection function in the form of the survey footprint, BORG is able to account for more complex selection functions, such as those in the form of relative probabilities in each sky pixel (Lavaux et al. 2019). The resulting predicted galaxy field is compared to the observed galaxy field through a likelihood distribution. A schematic overview of the data model used in the BORG algorithm in this paper is given in Fig. 2. One feature to be specifically pointed out in the data model is that the primordial perturbation with f_{NL} enters the data model at a different step from the evaluation of the structure formation model. This allows the algorithm to break the degeneracy between gravitational nonlinearities and primordial signals (Baumann & Green 2022; Andrews et al. 2023). To emphasise, a major advantage

of relying on forward model analysis to infer PNG is that more information is available, beyond what is available for standard methods (Leclercq & Heavens 2021; Nguyen et al. 2024).

To sample the joint posterior spanned by the data model, the algorithm relies on MCMC analysis. The complete multivariate distribution (Eq. (1)) is handled using a Gibbs sampling approach. The sampling scheme consists of a mixture of Hamiltonian Monte Carlo (HMC, Duane et al. 1987; Betancourt 2017) and slice sampling techniques (Hastings 1970; Neal 2003, 2011). To reiterate, the joint posterior distribution includes the three-dimensional initial density fields, cosmological parameters (i.e. f_{NL}), and the galaxy bias parameters for each galaxy catalogue (see Sect. 3.2). Exploring this joint posterior distribution allows BORG to effectively leverage all available information in the data for optimal parameter constraints, while also marginalising over nuisance parameters, given the data model and its resolution. Thus, through an iterative MCMC analysis, BORG performs a statistically rigorous analysis, allowing us to quantify the significance of inferred quantities (Jasche & Wandelt 2013). The sampling scheme is given in Fig. 3.

For more details on the forward model and likelihood, we refer to the previous publication (Andrews et al. 2023). For a description of the structure formation model, second-order Lagrangian perturbation theory (2LPT), we refer to similar work (Jasche & Wandelt 2013; Jasche et al. 2015; Lavaux & Jasche 2016; Lavaux et al. 2019; Tsaprazi et al. 2022).

2.2. Model of local primordial non-Gaussianity f_{NL}

Before applying the primordial perturbation with f_{NL} , we first convolved the Gaussian white noise field ϵ with the primordial transfer function $T_G(k)$. The field ϵ is the set of initial conditions for the forward model shown in Fig. 2. The transfer function $T_G(k)$ scaled the white noise field so that it has the properties of the primordial power spectrum. In this work, we first generated the adiabatic curvature perturbation $\mathcal{R}$, with covariance set by the amplitude A_s , and, following the discrete primordial power spectrum, the resulting Gaussian potential field ϕ_G then has the following covariance:

$$\langle \hat{\mathcal{R}}_{\mathbf{a}} \hat{\mathcal{R}}_{\mathbf{b}}^* \rangle = V \delta_{\mathbf{a},\mathbf{b}}^K A_s \frac{2\pi^2}{k_a^3} \left(\frac{k_a}{k_{\text{pivot}}} \right)^{n_s-1}, \quad (2)$$

with $V = L^3$ the volume of the data cube, $\mathbf{a}$ and $\mathbf{b}$ mesh indices, and $\delta_{\mathbf{a},\mathbf{b}}^K$ the Kronecker delta. The amplitude A_s set the variance of the primordial curvature perturbation $\mathcal{R}$, which we then converted to the Bardeen potential during matter domination via $\phi_g = \frac{3}{5} \mathcal{R}$. This Gaussian potential field ϕ_g , evaluated at $a \rightarrow 0$, provided the physical initial conditions for the forward model and contained full 3D information on the primordial matter fluctuations. The primordial gravitational potential in real space is then calculated with an inverse Fourier transform, with which we computed the perturbed primordial gravitational potential Φ_{NG} . Generally, PNG is the deviation of the primordial gravitational potential from Gaussian statistics. To the lowest order, these deviations are parameterised by the quantity f_{NL} . Although the deviation can have different shapes (Karagiannis et al. 2018), in this paper we focus on forecasting the constraining power of the local form (or squeezed shape). We parameterised f_{NL} through the matter-era Bardeen potential (Hodges et al. 1990; Kofman 1991; Salopek & Bond 1990; Gangui et al. 1994; Verde et al.




Fig. 4. *Euclid* sky map. This sky map illustrates the observed (yellow) and masked (red) regions for the *Euclid* survey of this project. The survey mask is a result of the observation strategy of the *Euclid* mission (Euclid Collaboration: Scaramella et al. 2022). We point out that each tracer catalogue uses one single survey strategy but extends outwards at different redshifts (as illustrated in Fig. 5).

2000; Wang & Kamionkowski 2000; Komatsu & Spergel 2001):

$$\Phi_{\text{NG}} = \phi_g + f_{\text{NL}} \left(\phi_g^2 - \langle \phi_g^2 \rangle \right), \quad (3)$$

where ϕ_g is a field described by Gaussian statistics, evaluated at the Lagrangian position. We considered f_{NL} to be a constant parameter, independent of scale and any other parameter, including time. It is also the lowest order of PNG. We leave the incorporation of the higher-order expansion of PNG for future implementations (Jeong & Komatsu 2009; Roth & Porciani 2012; Leistedt et al. 2014).

2.3. Galaxy bias model

In galaxy survey analysis, galaxy formation is typically described as a functional relationship of the dark matter field and bias parameters (Assassi et al. 2015; Desjacques et al. 2018a). Specifically, we treated galaxies as ‘biased’ tracers of the dark matter field, meaning that they shared similar clustering statistics and properties. However, the true relationship is unknown and constitutes one of the most important unresolved problems in LSS cosmology (for an exhaustive review, see Desjacques et al. 2018a). Thus, while noting that the forecasts are highly dependent on the estimated relationship, we assumed a next-to-leading order bias model with scale-dependent bias components (Assassi et al. 2015; Barreira 2020). The motivation is that since we are still at relatively large scales ($>60 h^{-1}$ Mpc), a relationship between the dark matter field and the galaxy field can be described as a linear function. By further including additional terms beyond the linear bias, we allowed the model to account for nonlinear features as well.

The complete bias model for a given tracer population g adopted in this paper is defined as

$$\rho_g(z, \mathbf{q}) = \langle N_g^0 \rangle \left[1 + b_1 \delta_m(z) + \frac{b_2}{2} \delta_m^2(z) + b_K K^2(z) + b_\phi f_{\text{NL}} \phi_g(\mathbf{q}) + b_{\phi,\delta} f_{\text{NL}} \delta_m(z) \phi_g(\mathbf{q}) \right], \quad (4)$$

where $\langle N_g^0 \rangle$ is the mean number of observed galaxies, b_1 is the linear bias, b_2 is the second-order bias, b_K is the tidal field bias, and $K^2(z) = \text{tr}[K_{ij}^2(z)]$, where $K_{ij}(z) \equiv (\partial_i \partial_j / \nabla^2 - \delta_{ij}^K / 3) \delta_m(z)$ is the long-wavelength tidal field (Lazeyras et al. 2021;




Fig. 5. *Euclid* radial selection function. This plot displays the normalised radial selection, $dN(z)/(d\Omega dz)$, for the four galaxy catalogues in this project. Notice how the tracer catalogues do not overlap but rather cover separate regions in the mock universe.

Barreira et al. 2021)⁴. Scale-dependent bias terms b_ϕ and $b_{\phi,\delta}$ are defined and further discussed in Sect. 2.3. We note that $\phi(\mathbf{q})$ is evaluated in Lagrangian space, while $\delta_m(z)$ is evaluated in a space that includes redshift-space distortions, which we call redshift space. In this project, we analyse the data at constant redshifts and thus do not account for effects that arise due to observing the galaxies on their light cones. We leave the inclusion of light-cone effects when inferring PNG to a future publication.

We assumed a fixed noise level for the Gaussian distribution of galaxies, set to $\sigma_g^2 = \langle N_g^0 \rangle$ (Andrews et al. 2023). In an upcoming publication, we will assess both the noise-level assumption and the likelihood model for describing the distribution of galaxies at the voxel level. Additionally, we will explore how the constraints on f_{NL} depend on these choices. The ground truth values of the galaxy parameters (used to generate the mock data) can be found in Table 1.

We mention that there is also the possibility of using Effective Field Theory (EFT) to model the galaxy bias formalism and likelihood in a field-level inference approach (Schmidt et al. 2019, 2020; Schmidt 2021; Babić et al. 2022; Tucci & Schmidt 2024; Stadler et al. 2023; Babić et al. 2025; Nguyen et al. 2024; Stadler et al. 2025a,b; Kostić et al. 2023). For a review of the galaxy bias problem, the interested reader is referred to the literature (Desjacques et al. 2018a).

Scale-dependent bias model

In the model of local PNG considered, the primordial perturbation gives rise to a scale-dependent imprint on the biased tracer populations. This is due to the coupling of short- and long-wavelength modes in a non-zero f_{NL} universe (Dalal et al. 2008; Slosar et al. 2008; Matarrese & Verde 2008; Carbone et al. 2008; Verde & Matarrese 2009). Thus, the PNG adds a scale-dependent contribution to the galaxy bias relation between galaxies and the underlying primordial gravitational field, which scales as $\propto k^{-2}$ (Dalal et al. 2008; Slosar et al. 2008; Matarrese & Verde 2008; Carbone et al. 2008; Verde & Matarrese 2009), which is the most prominent on the largest scales. In this paper, we adopted the universal mass function (Barreira 2022c; Lucie-Smith et al. 2023; Fondi et al. 2024; Gutiérrez Adame et al. 2024) for a non-zero f_{NL} universe.

⁴ K_{ij} is computed in Fourier space, where the differential operators correspond to multiplications by components of the wavevector k .

This allowed us to model the bias parameter of the primordial gravitational potential b_ϕ as a function of the linear bias b_1

$$b_\phi = 2\delta_c(b_1 - p), \quad (5)$$

with $\delta_c = 1.686$ being the spherical critical overdensity in an Einstein–de Sitter universe (Percival 2005), and p a tracer-dependent parameter. For this forecast, we fixed p to 0.55 or 1 for each tracer population (Barreira et al. 2020; Cabass et al. 2022). Additionally, we incorporated the bias for the cross-field term $b_{\phi,\delta}$, adopting the parameterisation of Barreira (2022b) and Cabass et al. (2022)⁵,

$$b_{\phi,\delta} = b_\phi - b_1 + 1 + \delta_c \left[b_2 - \frac{8}{21}(b_1 - 1) \right]. \quad (6)$$

The problem of accurately modelling b_ϕ and $b_{\phi,\delta}$ remains an unresolved challenge within the cosmological community (Biagetti 2019; Barreira 2022a; Achúcarro et al. 2022; Barreira & Krause 2023; Sullivan et al. 2023; Gutiérrez Adame et al. 2024; Fondi et al. 2024; Ding et al. 2024; Sullivan & Chen 2025; Kvasiuk et al. 2025). Improving the precision of the models for these bias parameters is crucial for robust and unbiased inference of f_{NL} (Moradinezhad Dizgah et al. 2021; Barreira 2022c; Lazeyras et al. 2023). Further model-driven investigations of the treatment of b_ϕ and $b_{\phi,\delta}$ will result in more robust and comprehensive models that better capture the underlying physical processes. In this paper, we adopt the universal mass approximation as a practical choice (Barreira 2020, 2022c). However, we emphasise that this assumption is not fundamental to the method itself, but as advances are achieved in the modelling of b_ϕ and $b_{\phi,\delta}$, the forward model will be revised accordingly. For a more in-depth discussion of the problem, see Moradinezhad Dizgah et al. (2021) and Barreira (2022c).

2.4. Running the MCMC analysis

The BORG algorithm performs a large-scale MCMC to explore the joint posterior distribution of Eq. (1), given the mock datasets. We briefly touch on the details of the MCMC analysis performed in this paper. We follow the prescription as in Ramanah et al. (2019) and Andrews et al. (2023), which provide more details.

To ensure that the sampler can sample from the target posterior distribution, we initialised the initial conditions ϵ at a randomly chosen point, set at one-tenth of the overall amplitude. The bias parameters are initialised at nine-tenths of their ground truth values, and f_{NL} is shifted by +5. The prior distribution on f_{NL} is a Gaussian distribution with $\mu_{f_{NL}} = 0$, $\sigma(f_{NL}) = 100$. This design is motivated by the choice to have a broad and non-informative prior.

We started the burn-in phase of the MCMC runs by exclusively sampling the initial conditions ϵ , with f_{NL} and the bias parameters kept constant to their starting values. This first step continued until the amplitude of the initial conditions fluctuated around the prior expectation. This was monitored by the power spectra estimated from Markov samples. For an example of such a plot, see the results of similar work (e.g. Porqueres et al. 2019a; Ramanah et al. 2019; Porqueres et al. 2022; Andrews et al. 2023). Next, we continued the run with

⁵ We note that alternative parameterisations of $b_{\phi,\delta}$ exist in the literature, for example Moradinezhad Dizgah et al. (2021) and D’Amico et al. (2025).

sampling bias parameters, where each catalogue was assigned its own set, which were allowed to converge and stabilise. Finally, we included the sampling of f_{NL} . To ensure the inclusion of only post-burn-in samples, additional 5000 samples were generated before including MCMC samples in the analysis. At this point, BORG explored the parameter space of plausible large-scale structure realisations, spanned by ϵ , f_{NL} , and the bias parameters.

The MCMC chains were run until convergence, as determined by the Gelman–Rubin statistic $\hat{R}$ (Gelman & Rubin 1992). The Gelman–Rubin statistic is calculated by dividing the sequential samples of the chains into M distinct sets, each with an equal number of samples, where M typically ranges from 4 to 8. The sets were separated by N discarded samples, where N denotes the number of samples required to decorrelate f_{NL} , to ensure statistical independence. More details on the autocorrelation length of f_{NL} for each run can be found in Appendix B. When the threshold $|\hat{R}| \leq 1.05$ was reached, the chain was considered to have converged.

It should be noted that a proper Gelman–Rubin test assumes independent MCMC chains. In our case, a single MCMC chain has been split into several chains for the evaluation of the Gelman–Rubin test, due to limitations in computational resources. By doing this, we acknowledge the risks associated with this, for example, the sampler getting stuck in a local minimum or biasing our results. However, from investigating the convergence in the other diagnostic results (e.g. correlation lengths, corner plots, and estimates on uncertainty of uncertainty), we deem these risks to be negligible.

2.5. Generating 3D maps of adiabatic curvature fluctuations

Our field-level inference method, in addition to providing measurements of f_{NL} , inferred the primordial gravitational potential. Given a posterior sample of Φ_{NG} , we can generate a 3D map of the adiabatic curvature perturbation, or $\mathcal{R}$ map (Planck Collaboration XVII 2016a). From Eqs. (2) and (3), we computed $\Phi_{\text{NG}}(\mathbf{q})$ and then related it to the corresponding curvature field via the matter-era relation⁶.

$$\mathcal{R}(\mathbf{q}) = \frac{5}{3} \Phi_{\text{NG}}(\mathbf{q}), \quad (7)$$

where $\mathbf{q}$ is the vector in Lagrangian space (Komatsu & Spergel 2001; Okamoto & Hu 2002; Lyth & Wands 2002; Sasaki et al. 2006). Since BORG sampled the posterior of $\Phi_{\text{NG}}(\mathbf{q})$, we also estimated the corresponding uncertainties across the chain. The resulting $\mathcal{R}$ maps are presented in Sect. 4.2, together with a detailed description of the dataset. For a 2D reconstruction of the adiabatic curvature field, see Sect. 6.3 in Planck Collaboration XVII (2016a).

3. Data and data generation

3.1. The Euclid Wide Survey and local primordial non-Gaussianity

The Euclid Wide Survey has among its goals to probe the expansion history and evolution of our Universe (Scaramella et al. 2014; Euclid Collaboration: Scaramella et al. 2022). To perform this task, *Euclid* will observe a region of 15 000 square degrees,

⁶ On super-Hubble scales, the general relation between the curvature perturbation and the Bardeen potential is $\mathcal{R}_k = -[(5 + 3w)/(3 + 3w)] \Phi_k$, with w the equation-of-state parameter of the dominant component. During matter domination ($w = 0$) this reduces to the standard convention $\mathcal{R} = \frac{5}{3} \Phi$, which we adopt throughout this work.

Table 1. Detailed specifications of the spectroscopic *Euclid* mock data used in this paper.

Cat.	b_1	b_2	b_K	b_ϕ^*	$b_{\phi,\delta}^*$	$z_{\text{min}}/z_{\text{max}}$	n_{gal}
1	1.30	−0.74	−0.17	2.53	0.79	0.9/1.1	6.86
2	1.38	−0.70	−0.22	2.80	0.99	1.1/1.3	5.58
3	1.46	−0.66	−0.26	3.13	1.23	1.3/1.5	4.21
4	1.54	−0.60	−0.31	3.33	1.44	1.5/1.8	2.61

Notes. These specifications form the basis for the four tracer catalogues generated for all of the runs in this paper. The galaxy bias parameters used in this galaxy bias model are the linear bias (b_1), second-order bias (b_2), and tidal field bias (b_K). Scale-dependent bias parameters b_ϕ and $b_{\phi,\delta}$ have values derived based on Eqs. (5) and (6), with $p = 0.55$. The values of the galaxy bias parameters are based on Table 1 in Yankelevich & Porciani (2019), and the number densities are based on Table 2 in Euclid Collaboration: Blanchard et al. (2020). The galaxy density, n_{gal} , is given in units of $10^{-4} (h^3 \text{Mpc}^{-3})$. * Derived values.

over a redshift range of $0.9 < z < 1.8$. Over the next six years, it will observe up to 30 million spectroscopic redshift galaxies with high precision ($\sigma_z \approx 0.001$), which can be used for galaxy clustering studies.

Among several cosmological measurements, inferring PNG is one of the primary objectives of the *Euclid* mission. More specifically, we focus solely on the local shape of PNG (Eq. (3)). The earliest goals were set at $\sigma(f_{\text{NL}}) \approx 2$ (68.3% CI). This forecast is based on statistical information from two-point correlation functions measured by spectroscopic redshift galaxies and for a ground truth value of $f_{\text{NL}}^{\text{gt}} = 0$ (Laureijs et al. 2011). However, more recent Fisher forecasts estimate that $\sigma(f_{\text{NL}})$ around 4 to 5 (68.3% CI) is achievable, given the accuracy of the spectroscopic redshift measurements and updated galaxy counts (Amendola et al. 2018). These results include marginalisation over the galaxy bias parameters, nuisance parameters, and other cosmological parameters (Giannantonio et al. 2012; Amendola et al. 2018). We acknowledge that the constraining power of the forecasts is highly dependent on the measured linear galaxy bias b_1 and the relationship $b_\phi(b_1)$, due to the scale-dependent bias effect (as discussed in Sect. 2.3). That being said, the constraints on $f_{\text{NL}} \times b_\phi$ will be largely insensitive to this uncertainty (Barreira 2022a). Unless unspecified, we use the value of $p = 0.55$ (Euclid Collaboration: Finelli et al. 2026).

3.2. Euclid specifications

We base the mock datasets on the forecast specifications of *Euclid*, meaning that we made use of the survey features of the *Euclid* mission to generate the mock data. Examples include sky completeness coverage, radial selection effects, bias parameter values, and galaxy counts.

In Fig. 4, we illustrate the sky completeness map used. We emphasise that the same completeness map is used for all four galaxy catalogues, and has a total sky coverage of roughly 15 000 square degrees. Also, while the sky completeness map used in this forecast is the survey geometry footprint of the *Euclid* survey (Euclid Collaboration: Scaramella et al. 2022), BORG has the ability to use more complex sky completeness masks in its data model (Jasche & Lavaux 2019; Lavaux et al. 2019; Andrews et al. 2023).

The radial selection functions are plotted in Fig. 5. The observed number counts are split into four different galaxy populations, each population corresponding to a catalogue. This

Table 2. Overview of the runs included in this project.

Run #	Resolution	k_{NF}	$f_{\text{NL}}^{\text{gt}}$	Note
1	250	0.0125	0	Low resolution
2	125	0.025	0	Medium resolution
3	62.5	0.05	0	High resolution
4	31.25	0.1	0	Highest resolution
5	62.5	0.05	5	Different $f_{\text{NL}}^{\text{gt}}$
6	125	0.025	0	$p = 1$
7	125	0.025	0	Sample $b_\phi, b_{\phi, \delta}$

Notes. The runs are designed to fulfil the tests outlined in Sect. 3.4. Run #4 is the main run of the paper (bolded in the table), which uses the most realistic settings for a future *Euclid* study. The computed resolutions and values of k_{NF} are given in $h^{-1} \text{Mpc}$ and $h \text{Mpc}^{-1}$, respectively. The ground truth value used to generate the mock data is denoted as $f_{\text{NL}}^{\text{gt}}$. For the resulting inferred f_{NL} values with uncertainties, see Table 3.

design choice is based on the redshift binning of the *Euclid* forecasts (Euclid Collaboration: Blanchard et al. 2020). Here, it can be seen that each subsequent catalogue reaches further into the observable universe, covering a nonoverlapping redshift range $0.9 < z < 1.8$. The number counts decrease towards higher redshifts.

The values of the galaxy bias parameters and number densities are listed in Table 1. These bias values were chosen as in Table 1 of Yankelevich & Porciani (2019), which constitutes the current estimate of the bias values as a function of redshift. The number densities were based on Table 3 of Euclid Collaboration: Blanchard et al. (2020). The values for the scale-dependent bias parameters were derived with Eqs. (5) and (6).

3.3. Mock data generation

To generate mock data, we follow similar procedures as described in previous works (Jasche & Kitaura 2010; Jasche & Wandelt 2013; Ramanah et al. 2019; Andrews et al. 2023). In general, these mock datasets were generated by running the forward model on a set of randomly drawn initial conditions. Thus, the same forward model is used for both mock data generation and the inference process, maintaining no model mismatch. We provide the details below, in a step-by-step description.

1. The evaluation of the physics forward model was prepared for a cubic Cartesian box of side length $L = 8000 h^{-1} \text{Mpc}$ and $N_{\text{grid}} = 32, 64, 128, \text{ or } 256$, yielding grid resolutions in the range of $\Delta L = 250 h^{-1} \text{Mpc}$, $\Delta L = 125 h^{-1} \text{Mpc}$, $\Delta L = 62.5 h^{-1} \text{Mpc}$, and $\Delta L = 31.25 h^{-1} \text{Mpc}$, respectively.
2. A random three-dimensional field ϵ , with zero mean and unit standard deviation, was generated. Given this white-noise field, a primordial density field was computed by applying the primordial power spectrum (Eq. (2)), perturbing it with the f_{NL} parameter (Eq. (3)), and then apply the cosmological transfer function as provided by CLASS (Lesgourgues & Tram 2014). This produced the linear matter field δ_L , which was the starting point for the gravitational structure formation model.
3. To reduce the sample variance of the particle distribution, we oversampled the initial density by a factor of 2 per dimension, resulting in a total number of $(2N)^3$ simulation particles. Particles were evolved to the present epoch using




Fig. 6. Field-level results for inferring f_{NL} for a high-resolution run (Run #4). The figure illustrates that the method can find a unimodal marginalised distribution of f_{NL} that best explains the data, with the ground truth value $f_{\text{NL}}^{\text{gt}}$ (solid line) within the 68.3% CI of the estimated value. The value in the box is the maximum of the distribution in the range of 68.3%. We point out that these results are based on the universal mass function, which means that the scale-dependent bias parameters are fixed to the expressions in Eqs. (5) and (6).

2LPT. Redshift-space distortions were then applied by performing an additional displacement of the particles along the line of sight, proportional to their 2LPT peculiar velocities, thereby mapping the particle distribution from real space to redshift space. The particles were subsequently assigned to a three-dimensional Cartesian grid via the Cloud-In-Cell (CIC) kernel to obtain the present-day three-dimensional matter density field δ_m in redshift space.

4. To emulate a biased galaxy distribution, we applied a next-to-leading order, scale-dependent galaxy bias (as described in Sect. 2.3) to the forward-modelled density field. The output is the galaxy field, wherein the galaxy counts in each voxel are characterised by a Gaussian distribution. Detailed specifications are organised in Table 1, including the parameter choices for the galaxy bias model.
5. Finally, the radial selection functions and the survey geometry were applied to the simulated galaxy field to emulate the observational effects of the survey.

We used the set of best-fit cosmological parameters ($\Omega_m = 0.3153$, $\Omega_\Lambda = 0.6847$, $\Omega_b = 0.0493$, $h = 0.6736$, $A_s = 2.1 \times 10^{-9}$, $n_s = 0.9649$) from *Planck* (Planck Collaboration VI 2020c) to calculate the cosmological power spectrum and transfer functions. A summary of the specifications, including detailed parameter choices for the galaxy bias model, is provided in Table 1. We provide a rendering of one of the mock datasets in Appendix C. We note that by relying on 2LPT to model structure formation we forego including small-scale physics that capture information at higher-order correlation functions beyond the bispectrum. To accurately and completely capture these higher-order effects at small-scales, in future publications, we will rely on N -body solvers, such as tCOLA (Tassev et al. 2013), particle mesh (Jasche & Lavaux 2019), or field-level emulators (Doeser et al. 2024). Since the data analysis in this paper is self-consistent with the mock data generation – using the same forward model for both inference and mock data creation – it captures all the information contained in the mock datasets.

Finally, we note that our inference formally includes modes close to the Nyquist frequency of the initial conditions grid. These scales are not fully converged, since their true nonlinear evolution would in reality be influenced by unresolved fluctuations at smaller scales. While this does not introduce a bias in our analysis, as the treatment is self-consistent, the unconverged modes remain well below the noise level and are therefore not expected to affect the inference. However, their inclusion may

still lead to an optimistic estimate of the constraining power on f_{NL} . A more conservative treatment would involve downweighting these modes through, for example, a scale-dependent noise prescription, which we leave for future work.

3.4. Overview of runs

In this section, we provide a complete list of the runs that we performed for this paper. Although the main aim of this paper is to forecast f_{NL} measurements with mock *Euclid* galaxy surveys, we are also interested in investigating the performance of the method in various configurations. To achieve this, we vary the specifications in the data and the analysis (e.g. resolution, marginalisation, and parameters values). The main questions for each test, together with the design of the runs, are outlined in the following list. For a concise overview of the runs in this project, we refer to Table 2.

1. Resolution study

The imprint of PNG affects the full cosmic matter field, both at large and small scales. To test the method's constraining power as a function of scale, we set up three different runs. The first at coarser grid resolution ($\Delta L = 250 h^{-1}$ Mpc, Run #1), one at medium grid resolution ($\Delta L = 125 h^{-1}$ Mpc, Run #2), one at finer grid resolution ($\Delta L = 62.5 h^{-1}$ Mpc, Run #3), and one at higher grid resolution ($\Delta L = 31.25 h^{-1}$ Mpc, Run #4). By increasing the voxel resolution, we allow the algorithm to have more degrees of freedom in describing the 3D cosmic matter field. Therefore, we expect that the method can use more small-scale information in the LSS to constrain f_{NL} . We also perform these runs as a benchmark as we adjust other parameters, for example, changing the sampling scheme or parameter values.

2. $f_{\text{NL}} = 0$, $f_{\text{NL}} = 5$

We aim to investigate whether the algorithm's constraining power depends on the fiducial amplitude of PNG. To test this, we generated two identical mock datasets that differ only by their values of f_{NL} : one with $f_{\text{NL}} = 0$ (Run #3) and the other with $f_{\text{NL}} = 5$ (Run #5), both with finer grid resolutions. In this way, we test the algorithm roughly in the upper 68.3% CI of the *Planck* 2018 measurement and the other in null detection (Planck Collaboration IX 2020b). This also allows us to check if the algorithm can accurately retrieve a non-zero ground truth value of $f_{\text{NL}}^{\text{gt}}$.

3. $p = 1$

We aim to investigate how much the algorithm's constraining power depends on the amplitude of the scale-dependent bias effect. To test this, we generate and analyse mock data that have p (Eq. (5)) set to 1 (Run #6) instead of 0.55 (which is the default for the other runs). Since this value of p is larger, there is a weaker scale-dependent bias amplitude in this analysis, meaning that the expectation is that the inferred uncertainty in f_{NL} will be larger.

4. Sampling b_ϕ , $b_{\phi,\delta}$

As described in Sect. 2.3, PNG gives rise to a scale-dependent bias effect that is modelled by b_ϕ and $b_{\phi,\delta}$. In the other runs, we fix b_ϕ and $b_{\phi,\delta}$ to the expressions of the universal mass function (Eqs. (5) and (6)). In this test, we want to test whether the algorithm is able to jointly sample f_{NL} , b_ϕ and $b_{\phi,\delta}$, in the presence of priors. Therefore, in Run #7, we include the sampling of b_ϕ and $b_{\phi,\delta}$, in addition to the initial conditions ϵ , f_{NL} , and the other bias parameters. In the mock data, the ground truth values of b_ϕ and $b_{\phi,\delta}$ are set to the universal expressions of the mass function. The priors for b_ϕ and $b_{\phi,\delta}$ are Gaussians centred on universal mass function expressions (as a

Table 3. Inferred f_{NL} for each run.

Run	$f_{\text{NL}}^{\text{gt}}$	$\langle f_{\text{NL}} \rangle$	$\sigma(f_{\text{NL}})$	$\Delta[\sigma(f_{\text{NL}})]$	k_{NF}	Resol.
1	0	3.2	10.8	0.21	0.0125	250
2	0	-5.4	6.3	0.45	0.025	125
3	0	1.2	4.3	0.19	0.05	62.5
4	0	0.3	2.6	0.23	0.1	31.25
5	5	9.3	4.4	0.21	0.5	62.5
6	0	13.1	10.4	0.47	0.025	125
7	0	-2.5	8.2	0.63	0.025	125

Notes. The table illustrates how different exchanges in the physics model and setup affect the constraint power in the data to infer f_{NL} . Examples include an increase in resolution or a change of the structure formation model. We note that all inferred values of f_{NL} are within the 68.3% CI of the ground truth $f_{\text{NL}}^{\text{gt}}$ values. This suggests that the method can reliably infer f_{NL} from the data. We have also included the uncertainty of the uncertainty estimates, denoted as $\Delta[\sigma(f_{\text{NL}})]$, using the batch means method (Fishman & Yarberry 1997). As a reminder, for Runs #1–#6 we assume the universal mass function, while for Run #7 we sample the scale-dependent bias parameters (with a prior centred as in Eqs. (5) and (6)). The computed resolutions and values of k_{NF} are given in h^{-1} Mpc and $h \text{ Mpc}^{-1}$, respectively.

function of b_1 and b_2), with standard deviations set to 40% of their values: $\mathcal{P}_\phi(b_\phi) = \mathcal{G}(b_\phi^{\text{UMA}}, 0.4b_\phi^{\text{UMA}})$, and $\mathcal{P}_{\phi\delta}(b_{\phi,\delta}) = \mathcal{G}(b_{\phi,\delta}^{\text{UMA}}, 0.4b_{\phi,\delta}^{\text{UMA}})$, where b_ϕ^{UMA} and $b_{\phi,\delta}^{\text{UMA}}$ are the expressions in Eqs. (5) and (6). We note that b_1 and b_2 used to evaluate b_ϕ^{UMA} and $b_{\phi,\delta}^{\text{UMA}}$ correspond to the values in the current state of the MCMC chain, rather than the ground truth values of b_ϕ^{UMA} and $b_{\phi,\delta}^{\text{UMA}}$. We are mainly interested in evaluating the performance of the method, in terms of $\sigma(f_{\text{NL}})$, and investigating possible correlations between bias parameters b_ϕ and $b_{\phi,\delta}$, and f_{NL} . If successful, this run further highlights the flexibility of the field-level inference approach in adjusting the galaxy bias model, and to sample bias parameters arising due to primordial effects.

4. Results

The main goal of this work is to infer the marginal posterior distribution of f_{NL} for each generated mock dataset. In this way, we provide both forecasts of how well BORG can constrain f_{NL} with the Euclid Wide Survey, and test the inference power under a variety of configurations. The results have been summarised in Table 3. For each run, we obtain an ensemble of samples, each containing plausible values of f_{NL} , the initial conditions, and the nuisance parameters, given the mock data. From these ensembles, we calculate the ensemble mean and uncertainty of f_{NL} and include these in the table. We also compute the uncertainty of the uncertainty, which quantifies the margin of error. For completeness, we also include the ground truth $f_{\text{NL}}^{\text{gt}}$ values, the resolution, and the corresponding k_{max} , which we define as the Nyquist frequency, k_{NF}^7 . We highlight that each inferred ensemble mean $\langle f_{\text{NL}} \rangle$ is within the 68.3% CI of the ground truth $f_{\text{NL}}^{\text{gt}}$ value.

The main run of this project, Run #4, constitutes the most realistic *Euclid* mock data. With it, we infer f_{NL} at a voxel

⁷ Here we take k_{max} to be the Nyquist frequency associated with the grid resolution. Modes near this scale are affected by discretisation and do not carry the full physical information, but this definition provides a convenient and consistent cut-off for comparison purposes.




Fig. 7. Field-level results for inferring f_{NL} for all runs. The first four runs, Runs #1–#4, are included in panel a. The last three runs, Runs #5–#7, are included in panel b. The marginal distribution of Run #5 has been shifted to be relative to the ground truth $f_{\text{NL}}^{\text{gt}} = 0$.

resolution of $31.25 h^{-1} \text{Mpc}$ ($k_{\text{NF}} = 0.1 h \text{Mpc}^{-1}$). We illustrate the inferred posterior distribution in Fig. 6, with which we find $f_{\text{NL}} = 0.3 \pm 2.6$. This marginal posterior distribution contains the entirety of the information available in the data, given the physics model and resolution.

Furthermore, for a visual comparison between the runs, we include Fig. 7, with all the inferred marginal posterior distributions of f_{NL} . To maintain legibility, we have divided the seven posteriors into two subplots.

4.1. Further tests

Below, we present the results for the tests outlined in Sect. 3.4. In Sect. 4.2, we describe and present the results of generating maps of adiabatic curvature fluctuations.

4.1.1. Resolution study

We discuss the results of the resolution investigation, in which we test the constraining power of BORG as a function of the available small scales. This is tested by comparing Runs #1–4. As can be seen, by increasing the resolution from 32^3 to 64^3 , we improve the results by roughly 42%. Furthermore, by increasing the resolution from 64^3 to 128^3 , we improve the results by roughly 32%. Finally, by increasing the resolution from 128^3 to 256^3 , we improve the constraining power by approximately 40%. We note that we are still at the mildly linear regime ($k_{\text{NF}} = 0.1 h \text{Mpc}^{-1}$), but we use a forward model that can account for nonlinearities in the data. This means that we can, in principle, further increase the resolution while still being able to model the emerging nonlinear small-scale physics. However, due to the computational cost, we leave this investigation to a future project.

4.1.2. $f_{\text{NL}} = 0$ and $f_{\text{NL}} = 5$

We outline the results for the performance of the algorithm for a non-zero ground truth value of $f_{\text{NL}}^{\text{gt}}$. We generate two differ-

ent mock datasets with the same white noise and galaxy bias parameters but with different f_{NL} values. These two are Run #3 ($f_{\text{NL}}^{\text{gt}} = 0$) and Run #5 ($f_{\text{NL}}^{\text{gt}} = 5$). The results show that the constraints for the two runs are similar, which indicates that the algorithm's performance is largely independent of the underlying f_{NL} value. Moreover, the uncertainty of the uncertainty std [$\sigma(f_{\text{NL}})$] is similar for the two runs, providing additional confirmation of this. The results indicate that our algorithm does not have a strong bias or performance issue related to ground truth $f_{\text{NL}}^{\text{gt}}$, which means that we expect the same sensitivity for a null signal or a primordial signal. We leave the exploration of sensitivity to large ground truth values of $f_{\text{NL}}^{\text{gt}}$ to a future project, which previous work has shown can influence the estimated uncertainty (Creminelli et al. 2007; Liguori et al. 2007).

4.1.3. Amplitude of scale-dependent bias effect

We discuss the results of weakening the effect of the scale-dependent bias in the model, by changing the parameter p . We perform two runs at the same resolution, but with differing values of p from 0.55 (Run #2) to 1 (Run #6). The results show that by increasing p we reduce the constraining power in f_{NL} . In fact, the change in p from 0.55 to 1 results in a decrease in the constraints of roughly 39%. This outcome highlights the dependence of the algorithm forecast on changes in scale-dependent bias parameters.

4.1.4. Sampling of b_ϕ and $b_{\phi,\delta}$

With this investigation, we want to test the flexibility of BORG to include the joint sampling of the scale-dependent bias parameters, including priors. This is relevant since it is still not fully understood how to model the scale-dependent bias effect (Barreira & Krause 2023). We present the results of Run #7, in which we include the sampling of b_ϕ and $b_{\phi,\delta}$, in addition to the initial conditions, f_{NL} , and the other parameters of the galaxy




Fig. 8. Field-level results for inferring f_{NL} , when simultaneously also sampling b_ϕ and $b_{\phi,\delta}$. Corner plot for f_{NL} and bias parameters, for Run #7, catalogue 4. Although the priors (described in Sect. 3.4) keep the inferred value of f_{NL} centering around the expected value of 0, the possible degeneracies with b_ϕ and $b_{\phi,\delta}$ are still explored. Thus, while the results indicate that our field-level inference method can jointly sample f_{NL} together with b_ϕ and $b_{\phi,\delta}$ (in the presence of priors), more work is needed to stabilise the region of explored f_{NL} values. The priors on b_ϕ and $b_{\phi,\delta}$ is centred around their Gaussians centred on universal mass function expressions, with standard deviations at 40% of that value: $\mathcal{P}_\phi(b_\phi) = \mathcal{G}(b_\phi^{\text{UMF}}, 0.4b_\phi^{\text{UMF}})$, and $\mathcal{P}_{\phi\delta}(b_{\phi,\delta}) = \mathcal{G}(b_{\phi,\delta}^{\text{UMF}}, 0.4b_{\phi,\delta}^{\text{UMF}})$, where b_ϕ^{UMF} and $b_{\phi,\delta}^{\text{UMF}}$ are the expressions in Eqs. (5) and (6). The corner plots for the other catalogues and for Run #4 can be found in Appendix B.

bias. The results show that BORG is still able to provide unbiased constraints on f_{NL} , but that the constraints degrade by approximately 20%. In Fig. 8, we provide a visualisation of the corner plot of the sampled f_{NL} and galaxy bias parameters of the fourth catalogue in the Markov chain. The corner plots for the other catalogues and the full correlation matrix are presented and discussed in Appendix B. Although the results show promise, further developments and choice of priors (for example the ones presented in Fondi et al. 2024; Gutiérrez Adame et al. 2024) will be investigated in a future publication.

4.2. Maps of adiabatic fluctuations

In addition to measuring f_{NL} , our field-level inference method infers the initial 3D conditions of the data. With these sets of plausible initial conditions, we generate maps of the 3D adiabatic curvature fluctuations of the post-inflationary universe. To generate maps of adiabatic curvature fluctuations, we apply Eq. (7) to a subset of the inferred samples of Run #4. The subset consists of every tenth sample from the Markov chain, starting from the first sample after the burn-in phase has concluded.




Fig. 9. Illustrations of the inferred adiabatic curvature fluctuations. For each saved sample, we have a set of initial conditions ϵ that produce a plausible set of model predictions, constrained by the data. Each set of initial conditions corresponds to a field of adiabatic curvature fluctuations (Eqs. (2) and (7)), which are the input to the structure formation model. By computing these fluctuations for a subset of the chain, the method can provide an expected estimate of the fluctuation of the adiabatic curvature along with uncertainty. We highlight the edge of the survey window with the dotted lines, meaning that voxels outside of the inner regions in the final observed field contain no observations. In the top-left plot, we have included the ground truth field of adiabatic curvature fluctuations used to generate the mock data, averaged over the x direction. In the top-right plot, we have the expectation value of the adiabatic curvature fluctuations averaged over the x direction. In the bottom-left plot, we have the corresponding uncertainty averaged over the x direction. Lastly, in the bottom-right plot, we include the absolute difference between the mean inferred field and the ground truth field.

By analysing the posterior ensemble, we can compute the average $\mathcal{R}$ field and the corresponding variance. We mention that we also include the inferred values of f_{NL} in the evaluation of the perturbed primordial gravitational potential Φ_{NG} . These esti-

mated statistical 3D fields constitute a novel data product that the method enables.

In Fig. 9, we render the 2D projections (in the x direction) of the resulting 3D maps. We also include the ground truth $\mathcal{R}$

field of the mock data and the absolute residual between the inferred and ground truth field. The upper-left panel illustrates the ground truth map of $\mathcal{R}$ provided by the mock data. The upper-right panel illustrates the average inferred $\mathcal{R}$ values. The bottom-left panel illustrates the uncertainty in the inferred adiabatic curvature fluctuations, whereas the bottom-right panel depicts the residual between the ground truth and the inferred. To highlight the effect of the selection function, the edges of the survey are marked with dotted lines, with voxels outside the edge completely masked in Eulerian space. The voxels inside the edges are regions of the data that are unmasked or only partially masked. However, we emphasise that BORG is capable of extrapolating information into unobserved voxels using the physics-informed data model (Jasche & Wandelt 2013; Leclercq et al. 2015; Jasche et al. 2015; Leclercq et al. 2017). In summary, these maps offer a comprehensive representation of plausible adiabatic curvature fluctuations and capture the statistical properties and (at least the three-point) correlation functions of these fluctuations.

4.2.1. Additional diagnostics of the runs

Before concluding, we briefly mention the additional diagnostics and results of the runs provided in the appendix. In Appendix B we present correlation matrices, correlation lengths, and corner plots of f_{NL} and the bias parameters for Run #4 and #7. In Appendix C, we highlight the generated mock data and illustrate the inferred final density fields and compare them with their ground truth representations. We present further tests of the adiabatic curvature fluctuations, including plots, in Appendix A.

5. Discussion

The detection of PNG would have profound implications for our understanding of the early Universe and the inflationary paradigm. Local PNG can serve as a powerful probe to test the single-field inflation hypothesis (Falk et al. 1993; Gangui et al. 1994; Komatsu & Spergel 2001; Maldacena 2003; Bartolo et al. 2004; Chen 2010; Biagetti 2019; Meerburg et al. 2019; Green et al. 2024). In this context, the local f_{NL} parameter provides a convenient parameterisation to quantify the lowest order of PNG. Projects such as *Euclid* aim at constraining f_{NL} with high-precision redshift surveys (Laureijs et al. 2011; Amendola et al. 2018; Euclid Collaboration: Finelli et al. 2026). To forecast the detectable level of PNG in *Euclid* data, we apply a field-level inference method on *Euclid*-like spectroscopic mock data to constrain f_{NL} . This approach allows us to naturally and jointly use all of the information available in the data to provide measurements of PNG and the primordial matter fluctuations.

To generate and infer f_{NL} from the mock data, we use BORG, which is a Bayesian hierarchical field-level inference algorithm designed to analyse galaxy redshift surveys. The data model for this project, illustrated in Fig. 2, forward models the primordial matter fluctuations to a predicted observation of the 3D galaxy field, so that the initial conditions can be directly inferred from the data. The data model captures all the effects of the perturbation of the primordial gravitational potential with f_{NL} and the scale-dependent bias effect on the final observable. We point out that we still assume the universal mass approximation (Moradinezhad Dizgah & Keating 2019; Barreira 2022b) and lowest-order f_{NL} contributions. The data model will be updated as more progress is made in modelling the scale-dependent bias effect. Our approach can handle a variety of observational and systematic effects in the forward model, providing robust

inferences of PNG and the early Universe (Jasche & Lavaux 2017; Porqueres et al. 2019b; Lavaux et al. 2019). Moreover, the explicit physics-informed data model allows the method to maintain interpretability (Jasche & Wandelt 2013; Porqueres et al. 2019b) and perform posterior predictive tests of the results (Jasche & Lavaux 2019; Lavaux et al. 2019).

In this study, we use *Euclid* forecast specifications to generate mock data, such as the number of galaxies, number of catalogues, observed volume, and redshift binning, among others. To generate the mock data, we employed a forward model on a white-noise field, including effects such as structure formation effects, redshift selection functions, and sky masks, allowing us to simulate realistic observations. Throughout the study, we generate and analyse mock data under a variety of conditions, for example, tests with different ground truth $f_{\text{NL}}^{\text{gt}}$ values, and with different resolutions. The investigations paint a comprehensive picture of the impact of various factors on the inference process.

This study demonstrates the successful sampling of the parameter f_{NL} using mock *Euclid* data. All the runs carried out for this analysis result in inferred values of f_{NL} within the 68.3% CI of the ground truth $f_{\text{NL}}^{\text{gt}}$ value. These constraints demonstrate the reliability and robustness of the inference methodology. In particular, at a resolution of $31.25 h^{-1} \text{ Mpc}$, we achieve a constraining power of $\sigma(f_{\text{NL}}) = 2.6$ (68.3% CI). The resolution study carried out in this paper demonstrates the ability of our method to use information on small scales to constrain f_{NL} . In addition, we present detailed maps of primordial adiabatic curvatures, or $\mathcal{R}$ maps, and their corresponding uncertainties, providing a comprehensive rendering of the inferred initial conditions. In general, these results highlight the potential of using field-level inference to constrain primordial physics with *Euclid* data.

We forecast how well BORG can constrain f_{NL} in mock *Euclid* spectroscopic datasets. We highlight our major findings below and thereafter provide a summary.

- The field-level inference framework successfully handles *Euclid*-like specifications, accommodating the large survey volume, the number of galaxies, and various physical effects and noise properties. The data model used includes the generation of primordial gravitational potential, perturbation with local PNG, running a structure formation computation, applying a bias model, and evaluating the likelihood of galaxy formation.
- Our primary finding is that, at a resolution of $31.25 h^{-1} \text{ Mpc}$ ($k_{\text{NF}} = 0.1 h \text{ Mpc}^{-1}$), our method achieves an uncertainty of $\sigma(f_{\text{NL}}) = 2.6$ for $f_{\text{NL}} = 0$ at 68.3% CI, assuming the universal mass function and $p = 0.55$.
- Our method can infer the parameter f_{NL} with high precision at multiple resolutions, different fiducial f_{NL} values, and sampling configurations, demonstrating the consistency and versatility of the methodology.
- As we increase the resolution, our algorithm provides tighter constraints on f_{NL} , confirming the results in Andrews et al. (2023, see Figure 7). Thus, we expect additional improvements by further increasing the resolution, especially since our data model is capable of handling nonlinearities beyond the considered k_{NF} .
- When the amplitude of the scale-dependent bias effect is weakened, from $p = 0.55$ to $p = 1$, the constraints on the inferred f_{NL} decrease by approximately 40% (68.3% CI).
- We perform a run by also sampling the scale-dependent bias parameters b_ϕ and $b_{\phi, \delta}$, with priors centred on Eqs. (5) and (6). The run provides unbiased constraints on f_{NL} , but with an increase in $\sigma(f_{\text{NL}})$ by approximately 20% (68.3% CI).

- At a resolution of $31.25 h^{-1}$ Mpc, we generate maps of adiabatic curvature fluctuations from inferred initial conditions, offering valuable data products for conducting further investigations into the early Universe.
- We showcase additional convergence tests and data products, for example corner plots, correlation lengths, and inferred cosmic fields, in Appendices B–C.

In conclusion, our method provides a complementary and independent approach to statistical summary estimators for constraining primordial physics in galaxy redshift surveys. By using the full formation history of the Universe in the data model, field-level inference opens up the possibility to perform optimal measurements of PNG up to the given resolution and data model. With these measurements, the scientific community will be able to significantly reduce the parameter space of plausible models for the inflationary universe. In this way, our method will contribute to the scientific success of the *Euclid* mission by enabling it to excel in one of its primary research objectives.

6. Future work

In this section, we provide a discussion on the further testing and development of the method for inferring primordial physics in galaxy redshift survey data, with a focus on the following aspects: (1) validation against the *Euclid* flagship simulation (Potter et al. 2017; Euclid Collaboration: Castander et al. 2025), (2) handling observational effects such as relativistic effects and foreground systematic effects, and (3) inclusion of additional inflationary model parameters and (4) cosmological parameters.

6.1. Method validation with the Euclid flagship simulation

Before analysing the upcoming observed data, further tests of the adopted data model against more complex data are required, for example against the *Euclid* flagship simulation. The *Euclid* flagship simulation is a gravity-only dark matter simulation consisting of $12\,600^3$ dark matter particles within the size of $3780 h^{-1}$ Mpc. With the method, one can use the simulated halo catalogue as a substitute for galaxies, and thus perform further tests on the data model and the performance of the algorithm in constraining f_{NL} . Examples include further testing of the structure formation model and bias model in BORG, and how well they can capture the results of the full simulation. Furthermore, analysing N -body simulations with non-zero f_{NL} values is another way to validate the implemented PNG model (Jung et al. 2022; Coulton et al. 2023; Jung et al. 2023a,b; Fondi et al. 2024; Hadzhiyska et al. 2024). In addition, another insightful investigation is to test how changes in the data model, such as the choice of structure formation model, impact the algorithm's inference power on f_{NL} . Successfully further testing the algorithm with N -body simulations is a crucial step before the subsequent application on the upcoming real data.

6.2. Handling more complex observational effects

Real observational data are subject to various survey systematic and observational effects. Future work will therefore have to incorporate and mitigate these effects in the analysis framework. One example is the consideration of general relativistic effects, which imprint an effect similar to PNG on large scales. Another is the modelling of light-cone effects, which account for the evolution of the observable universe over cosmic time (Bruni et al. 2012; Jeong et al. 2012; Schmidt et al. 2013; Bertacca et al. 2014a,b; Jeong & Schmidt 2015; Yoo 2014; Yoo & Gong

2016; Koyama et al. 2018; Desjacques et al. 2018b; Umeh et al. 2019; Lavaux et al. 2019; Wang et al. 2020; Maartens et al. 2021; Martinez-Carrillo et al. 2021; Castorina & Di Dio 2022; Enríquez et al. 2022; Shiveshwarkar et al. 2023; Rossiter et al. 2025; Addis et al. 2025). By properly handling these effects, the algorithm aims to break the degeneracy between these effects and PNG, in such a way that one can ensure that the method accurately captures the relevant information from observed data. Furthermore, foreground systematic effects, such as contamination from Galactic emissions or instrumental artefacts, can significantly bias the measured cosmological parameters if they are not taken into account (Leistedt et al. 2014; Rezaie et al. 2021; Mueller et al. 2021; Rezaie et al. 2024). Developing robust methods to identify and marginalise these systematic effects is a key question in the *Euclid* mission. In addition, a careful treatment of velocity bias is required when modelling tracer dynamics in redshift space. While negligible on the largest scales relevant for PNG, velocity bias and related higher-order RSD effects become important at intermediate scales ($k_{\text{max}} > 0.05 h \text{ Mpc}^{-1}$), particularly if one aims to perform the inference at higher resolutions, such as the resolutions of run #4. Extending or adjusting the bias model to properly capture these effects will be crucial to ensure unbiased cosmological constraints (Stadler et al. 2023, 2025a,b). Lastly, tests incorporating photometric redshift datasets will also be conducted (Jasche & Wandelt 2012; Tsaprazi et al. 2023).

6.3. Inclusion of additional inflationary model parameters

Inflationary models offer a rich framework for understanding the early Universe and its subsequent evolution. To further constrain the space of plausible inflationary models, the data model can be extended to include additional inflationary model parameters beyond the parameter f_{NL} . For example, the data model could include other parameters such as the local trispectrum non-Gaussianity parameter g_{NL} (Okamoto & Hu 2002; Sasaki et al. 2006; Jeong & Komatsu 2009; Leistedt et al. 2014; Roth & Porciani 2012; Shiveshwarkar et al. 2024; Pardede et al. 2023), the running-of-the-scalar index α_s (Fedeli et al. 2010; Planck Collaboration X 2020a; Germán 2021), and other shapes of PNG, for example, the equilateral non-Gaussianity parameter $f_{\text{NL}}^{\text{equi}}$ (Babich et al. 2004; Scoccimarro et al. 2012; Regan et al. 2012; Planck Collaboration XXIV 2014a; Schmidt et al. 2015; Planck Collaboration XVII 2016a; Karagiannis et al. 2018; Planck Collaboration IX 2020b; Karagiannis et al. 2020; Baumann & Green 2022). Such implementations would be subject to validation tests and mock data studies before making predictions on how well field-level inference methods can constrain such parameters in the large-scale structure. In short, by incorporating additional primordial parameters, the data model would be able to perform a more thorough exploration of the inflationary paradigm in the data, and thus allowing the method to further distinguish different inflationary models.

6.4. Joint sampling of cosmological parameters

Another test of interest is to explore the parameter space of the Λ CDM model together with non-Gaussian cosmic initial conditions. By doing so, one would be able to assess the efficacy of field-level inference in the context of jointly sampling other cosmological parameters together with f_{NL} . Examples of such parameters include Ω_{m} , w_0 (Ramanah et al. 2019), and σ_8 (Schmidt 2021; Porqueres et al. 2021, 2022, 2023; Nguyen et al. 2024; Beyond-2pt Collaboration 2025). However, extending the

analysis to include additional degrees of freedom exposes the algorithm to potential parameter degeneracies, which could degrade the constraints in f_{NL} . Based on previous work, we anticipate a marginal decline in performance, projected to be within the limit of 10% (Jung et al. 2023a,b). Confirmation of these expectations will be made through future mock data tests, within the context of *Euclid* simulations and other datasets.

Acknowledgements. We thank Fabian Schmidt for valuable feedback on the manuscript. JJ acknowledges support by the Swedish Research Council (VR) under the project 2020-05143 – ‘Deciphering the Dynamics of Cosmic Structure’ and from the contract ASI/ INAF for the Euclid mission n.2018-23-HH.0. GL acknowledge financial support from the Centre National d’Etudes Spatiales (project GCEUCLID), and the Simons Foundation collaboration programme ‘Learning the Universe’. FF, MB, and DP acknowledge partial financial support from the contract ASI/ INAF for the Euclid mission n.2018-23-HH.0 and from the INFN InDark initiative. The computations and data handling were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) and the Swedish National Infrastructure for Computing (SNIC) at Tetralith partially funded by the Swedish Research Council through grant agreements no. 2022-06725 and no. 2018-05973. This research utilised the HPC facility supported by the Technical Division of the Department of Physics, Stockholm University. MB acknowledges financial support from the INFN InDark initiative and from the COSMOS network through the ASI (Italian Space Agency) Grants 2016-24-H.0 and 2016-24-H.1-2018, as well as 2020-9-HH.0 (participation in LiteBIRD phase A). JV was supported by Ruth och Nils-Erik Stenbäcks Stiftelse and Research Council of Finland grant 347088. FP acknowledges partial support from the INFN grant InDark and the Departments of Excellence grant L.232/2016 of the Italian Ministry of University and Research (MUR) and the FCT project with ref. number PTDC/FIS-AST/0054/2021. We acknowledge the use of the following packages: NumPy (Harris et al. 2020), Matplotlib (Hunter 2007), GetDist (Lewis 2019), and HEALPix (Gorski et al. 2005). This work is done within the Aquila Consortium (<https://www.aquila-consortium.org/>). The Euclid Consortium acknowledges the European Space Agency and a number of agencies and institutes that have supported the development of *Euclid*, in particular the Agenzia Spaziale Italiana, the Austrian Forschungsförderungsgesellschaft funded through BMIMI, the Belgian Science Policy, the Canadian Euclid Consortium, the Deutsches Zentrum für Luft- und Raumfahrt, the DTU Space and the Niels Bohr Institute in Denmark, the French Centre National d’Etudes Spatiales, the Fundação para a Ciência e a Tecnologia, the Hungarian Academy of Sciences, the Ministerio de Ciencia, Innovación y Universidades, the National Aeronautics and Space Administration, the National Astronomical Observatory of Japan, the Nederlandse Onderzoeksschool Voor Astronomie, the Norwegian Space Agency, the Research Council of Finland, the Romanian Space Agency, the Swiss Space Office (SSO) at the State Secretariat for Education, Research, and Innovation (SERI), and the United Kingdom Space Agency. A complete and detailed list is available on the *Euclid* web site (www.euclid-ec.org/consortium/community/).

References

- Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, arXiv e-prints [arXiv:1610.02743]
- Achúcarro, A., Biagetti, M., Braglia, M., et al. 2022, arXiv e-prints [arXiv:2203.08128]
- Acquaviva, V., Bartolo, N., Matarrese, S., et al. 2003, *Nucl. Phys. B*, **667**, 119
- Addis, C., Guandalini, C., & Clarkson, C. 2025, *JCAP*, **04**, 080
- Alvarez, M., Baldauf, T., Bond, J. R., et al. 2014, arXiv e-prints [arXiv:1412.4671]
- Amendola, L., Appleby, S., Avgoustidis, A., et al. 2018, *Liv. Rev. Relat.*, **21**, 2
- Andrews, A., Jasche, J., Lavaux, G., et al. 2023, *MNRAS*, **520**, 5746
- Arkani-Hamed, N., & Maldacena, J. 2015, arXiv e-prints [arXiv:1503.08043]
- Assassi, V., Baumann, D., & Schmidt, F. 2015, *JCAP*, **12**, 043
- Babić, I., Schmidt, F., & Tucci, B. 2022, *JCAP*, **08**, 007
- Babić, I., Schmidt, F., & Tucci, B. 2025, *JCAP*, **11**, 066
- Babich, D., Creminelli, P., & Zaldarriaga, M. 2004, *JCAP*, **08**, 009
- Baldauf, T., Seljak, U., & Senatore, L. 2011a, *JCAP*, **04**, 006
- Baldauf, T., Seljak, U., Senatore, L., et al. 2011b, *JCAP*, **10**, 031
- Ballardini, M., Matthewson, W. L., & Maartens, R. 2019, *MNRAS*, **489**, 1950
- Barberi Squarotti, M., Camera, S., & Maartens, R. 2024, *JCAP*, **02**, 043
- Barreira, A. 2020, *JCAP*, **12**, 031
- Barreira, A. 2022a, *JCAP*, **11**, 013
- Barreira, A. 2022b, *JCAP*, **01**, 033
- Barreira, A. 2022c, *JCAP*, **04**, 057
- Barreira, A., & Krause, E. 2023, *JCAP*, **10**, 044
- Barreira, A., Cabass, G., Schmidt, F., et al. 2020, *JCAP*, **12**, 013
- Barreira, A., Lazeyras, T., & Schmidt, F. 2021, *JCAP*, **08**, 029
- Bartolo, N., Komatsu, E., Matarrese, S., et al. 2004, *Phys. Rep.*, **402**, 103
- Baumann, D., & Green, D. 2022, *JCAP*, **08**, 061
- Baumann, D., Jackson, M. G., Adshead, P., et al. 2009, *AIP Conf. Proc.*, **1141**, 10
- Bayer, A. E., Seljak, U., & Modi, C. 2023, arXiv e-prints [arXiv:2307.09504]
- Bermejo-Climent, J. R., Demina, R., Krolewski, A., et al. 2025, *A&A*, **698**, A177
- Bertacca, D., Maartens, R., & Clarkson, C. 2014a, *JCAP*, **09**, 037
- Bertacca, D., Maartens, R., & Clarkson, C. 2014b, *JCAP*, **11**, 013
- Betancourt, M. 2017, arXiv e-prints [arXiv:1701.02434]
- Beyond-2pt Collaboration (Krause, E., et al.) 2025, *ApJ*, **990**, 99
- Biagetti, M. 2019, *Galaxies*, **7**, 71
- Biagetti, M., Calles, J., Castiblanco, L., et al. 2022, *JCAP*, **10**, 002
- Brown, Z., Demina, R., Adame, A. G., et al. 2025, *MNRAS*, **543**, 2078
- Bruni, M., Crittenden, R., Koyama, K., et al. 2012, *Phys. Rev. D*, **85**, 041301
- Byrnes, C. T., & Choi, K.-Y. 2010, *Adv. Astron.*, **2010**, 724525
- Byrnes, C. T., Choi, K.-Y., & Hall, L. M. H. 2009, *JCAP*, **02**, 017
- Byrnes, C. T., Gerstenlauer, M., Nurmi, S., et al. 2010, *JCAP*, **10**, 004
- Cabass, G., Ivanov, M. M., Philcox, O. H. E., et al. 2022, *Phys. Rev. Lett.*, **129**, 021301
- Cagliari, M. S., Castorina, E., Bonici, M., et al. 2024, *JCAP*, **08**, 036
- Carbone, C., Verde, L., & Matarrese, S. 2008, *ApJ*, **684**, L1
- Castorina, E., & Di Dio, E. 2022, *JCAP*, **01**, 061
- Castorina, E., Hand, N., Seljak, U., et al. 2019, *JCAP*, **09**, 010
- Celoria, M., & Matarrese, S. 2018, arXiv e-prints [arXiv:1812.08197]
- Chaussidon, E., de Mattia, A., Yèche, C., et al. 2025a, *JCAP*, **01**, 135
- Chaussidon, E., Yèche, C., de Mattia, A., et al. 2025b, *JCAP*, **06**, 029
- Chen, X. 2010, *Adv. Astron.*, **2010**, 638979
- Chen, X., Ebadi, R., & Kumar, S. 2022, *JCAP*, **08**, 083
- Chen, X., Padmanabhan, N., & Eisenstein, D. J. 2025, *JCAP*, **08**, 055
- Coulton, W. R., Villaescusa-Navarro, F., Jamieson, D., et al. 2023, *ApJ*, **943**, 64
- Creminelli, P., & Zaldarriaga, M. 2004, *JCAP*, **10**, 006
- Creminelli, P., Senatore, L., & Zaldarriaga, M. 2007, *JCAP*, **03**, 019
- Creminelli, P., D’Amico, G., Musso, M., et al. 2011, *JCAP*, **11**, 038
- Dalal, N., Doré, O., Huterer, D., et al. 2008, *Phys. Rev. D*, **77**, 123514
- D’Amico, G., Lewandowski, M., Senatore, L., et al. 2025, *Phys. Rev. D*, **111**, 063514
- Desjacques, V., Jeong, D., & Schmidt, F. 2018a, *Phys. Rep.*, **733**, 1
- Desjacques, V., Jeong, D., & Schmidt, F. 2018b, *JCAP*, **12**, 035
- Ding, S., Lavaux, G., & Jasche, J. 2024, *A&A*, **690**, A236
- Doerer, L., Jamieson, D., Stopyra, S., et al. 2024, *MNRAS*, **535**, 1258
- Doré, O., Bock, J., Ashby, M., et al. 2014, arXiv e-prints [arXiv:1412.4872]
- Duane, S., Kennedy, A. D., Pendleton, B. J., et al. 1987, *Phys. Lett. B*, **195**, 216
- Duivenvoorden, A. J., Meerburg, P. D., & Freese, K. 2020, *Phys. Rev. D*, **102**, 023521
- Enríquez, M., Hidalgo, J. C., & Valenzuela, O. 2022, *JCAP*, **03**, 048
- Euclid Collaboration (Ballardini, M., et al.) 2024, *A&A*, **683**, A220
- Euclid Collaboration (Blanchard, A., et al.) 2020, *A&A*, **642**, A191
- Euclid Collaboration (Castander, F. J., et al.) 2025, *A&A*, **697**, A5
- Euclid Collaboration (Finelli, F., et al.) 2026, *A&A*, in press, <https://doi.org/10.1051/0004-6361/202556600>
- Euclid Collaboration (Ilić, S., et al.) 2022, *A&A*, **657**, A91
- Euclid Collaboration (Mellier, Y., et al.) 2025, *A&A*, **697**, A1
- Euclid Collaboration (Scaramella, R., et al.) 2022, *A&A*, **662**, A112
- Falk, T., Rangarajan, R., & Srednicki, M. 1993, *ApJ*, **403**, L61
- Fedeli, C., Finelli, F., & Moscardini, L. 2010, *MNRAS*, **407**, 1842
- Finelli, F., Bucher, M., Achúcarro, A., et al. 2018, *JCAP*, **04**, 016
- Fishman, G., & Yarberry, S. 1997, *INFORMS J. Comput.*, **9**, 3
- Fondi, E., Verde, L., Villaescusa-Navarro, F., et al. 2024, *JCAP*, **02**, 048
- Friedrich, O., Uhlemann, C., Villaescusa-Navarro, F., et al. 2020, *MNRAS*, **498**, 464
- Gangui, A., Lucchin, F., Matarrese, S., et al. 1994, *ApJ*, **430**, 447
- Gelman, A., & Rubin, D. B. 1992, *Stat. Sci.*, **7**, 457
- Germán, G. 2021, *Int. J. Mod. Phys. D*, **30**, 2150038
- Giannantonio, T., & Percival, W. J. 2014, *MNRAS*, **441**, L16
- Giannantonio, T., Porciani, C., Carron, J., et al. 2012, *MNRAS*, **422**, 2854
- Giri, U., Münchmeyer, M., & Smith, K. M. 2025, *Phys. Rev. D*, **112**, 123544
- Goldstein, S., Esposito, A., Philcox, O. H. E., et al. 2022, *Phys. Rev. D*, **106**, 123525
- Goldstein, S., Philcox, O. H. E., Hill, J. C., et al. 2024, *Phys. Rev. D*, **109**, 043515
- Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Graham, M. L., Connolly, A. J., Ivezić, Ž., et al. 2018, *AJ*, **155**, 1
- Green, D., Guo, Y., Han, J., et al. 2024, *JCAP*, **05**, 090
- Guth, A. H. 1981, *Phys. Rev. D*, **23**, 347
- Gutiérrez Adame, A., Avila, S., Gonzalez-Perez, V., et al. 2024, *A&A*, **689**, A69

- Hadzhiyska, B., Garrison, L. H., Eisenstein, D. J., et al. 2024, *Phys. Rev. D*, **109**, 103530
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, **585**, 357
- Hastings, W. K. 1970, *Biometrika*, **57**, 97
- Ho, S., Agarwal, N., Myers, A. D., et al. 2015, *JCAP*, **05**, 040
- Hodges, H. M., Blumenthal, G. R., Kofman, L. A., et al. 1990, *Nucl. Phys. B*, **335**, 197
- Hunter, J. D. 2007, *Comput. Sci. Eng.*, **9**, 90
- Huterer, D., Cunha, C. E., & Fang, W. 2013, *MNRAS*, **432**, 2945
- Jasche, J., & Kitaura, F. S. 2010, *MNRAS*, **407**, 29
- Jasche, J., & Lavaux, G. 2017, *A&A*, **606**, A37
- Jasche, J., & Lavaux, G. 2019, *A&A*, **625**, A64
- Jasche, J., & Wandelt, B. D. 2012, *MNRAS*, **425**, 1042
- Jasche, J., & Wandelt, B. D. 2013, *MNRAS*, **432**, 894
- Jasche, J., Kitaura, F. S., Wandelt, B. D., et al. 2010, *MNRAS*, **406**, 60
- Jasche, J., Leclercq, F., & Wandelt, B. D. 2015, *JCAP*, **01**, 036
- Jeong, D., & Komatsu, E. 2009, *ApJ*, **703**, 1230
- Jeong, D., & Schmidt, F. 2015, *Class. Quant. Grav.*, **32**, 044001
- Jeong, D., Schmidt, F., & Hirata, C. M. 2012, *Phys. Rev. D*, **85**, 023504
- Jung, G., Karagiannis, D., Liguori, M., et al. 2022, *ApJ*, **940**, 71
- Jung, G., Karagiannis, D., Liguori, M., et al. 2023a, *ApJ*, **948**, 135
- Jung, G., Raverri, A., Baldi, M., et al. 2023b, *ApJ*, **957**, 50
- Jung, G., Raverri, A., Liguori, M., et al. 2024, *ApJ*, **976**, 109
- Kalaja, A., Meerburg, P. D., Pimentel, G. L., et al. 2021, *JCAP*, **04**, 050
- Karagiannis, D., Lazanu, A., Liguori, M., et al. 2018, *MNRAS*, **478**, 1341
- Karagiannis, D., Slosar, A., & Liguori, M. 2020, *JCAP*, **11**, 052
- Karagiannis, D., Fonseca, J., Maartens, R., et al. 2021, *Phys. Dark Universe*, **32**, 100821
- Karagiannis, D., Maartens, R., Fonseca, J., et al. 2024, *JCAP*, **03**, 034
- Kofman, L. 1991, *Phys. Scr. Vol. T*, **36**, 108
- Komatsu, E. 2010, *Class. Quant. Grav.*, **27**, 124010
- Komatsu, E., & Spergel, D. N. 2001, *Phys. Rev. D*, **63**, 063002
- Komatsu, E., Afshordi, N., Bartolo, N., et al. 2009, arXiv e-prints [arXiv:0902.4759]
- Kostić, A., Jasche, J., Ramanah, D. K., et al. 2022, *A&A*, **657**, L17
- Kostić, A., Nguyen, N.-M., Schmidt, F., et al. 2023, *JCAP*, **07**, 063
- Koyama, K., Umeh, O., Maartens, R., et al. 2018, *JCAP*, **07**, 050
- Krolewski, A., Percival, W. J., Ferraro, S., et al. 2024, *JCAP*, **03**, 021
- Kvasiuk, Y., Münchmeyer, M., & Smith, K. 2025, *Phys. Rev. D*, **112**, 023540
- Laureijs, R., Amiaux, J., Arduini, S., et al. 2011, arXiv e-prints [arXiv:1110.3193]
- Lavaux, G., & Jasche, J. 2016, *MNRAS*, **455**, 3169
- Lavaux, G., Jasche, J., & Leclercq, F. 2019, arXiv e-prints [arXiv:1909.06396]
- Lazeyras, T., Barreira, A., & Schmidt, F. 2021, *JCAP*, **10**, 063
- Lazeyras, T., Barreira, A., Schmidt, F., et al. 2023, *JCAP*, **01**, 023
- Leclercq, F., & Heavens, A. 2021, *MNRAS*, **506**, L85
- Leclercq, F., Jasche, J., Lavaux, G., & Wandelt, B. 2015, arXiv e-prints [arXiv:1512.02242]
- Leclercq, F., Jasche, J., Lavaux, G., et al. 2017, *JCAP*, **06**, 049
- Leistedt, B., Peiris, H. V., & Roth, N. 2014, *Phys. Rev. Lett.*, **113**, 221301
- Lesgourgues, J., & Tram, T. 2014, *JCAP*, **09**, 032
- Lewis, A. 2019, arXiv e-prints [arXiv:1910.13970]
- Liguori, M., Yadav, A., Hansen, F. K., et al. 2007, *Phys. Rev. D*, **76**, 105016
- Lo Verde, M., Miller, A., Shandera, S., et al. 2008, *JCAP*, **04**, 014
- LSST Science Collaborations (Abell, P. A., et al.) 2009, arXiv e-prints [arXiv:0912.0201]
- Lucie-Smith, L., Barreira, A., & Schmidt, F. 2023, *MNRAS*, **524**, 1746
- Lyth, D. H., & Wands, D. 2002, *Phys. Lett. B*, **524**, 5
- Maartens, R., Jolicœur, S., Umeh, O., et al. 2021, *JCAP*, **04**, 013
- Maldacena, J. 2003, *J. High Energy Phys.*, **5**, 013
- Martinez-Carrillo, R., Hidalgo, J. C., Malik, K. A., et al. 2021, *JCAP*, **12**, 025
- Matarrese, S., & Verde, L. 2008, *ApJ*, **677**, L77
- McCarthy, F., Madhavacheril, M. S., & Maniyar, A. S. 2023, *Phys. Rev. D*, **108**, 083522
- McQuinn, M. 2021, *JCAP*, **06**, 024
- Meerburg, P. D., Green, D., Flauger, R., et al. 2019, *BAAS*, **51**, 107
- Moradinezhad Dizgah, A., & Keating, G. K. 2019, *ApJ*, **872**, 126
- Moradinezhad Dizgah, A., Biagetti, M., Sefusatti, E., et al. 2021, *JCAP*, **05**, 015
- Mueller, E.-M., Percival, W. J., & Ruggeri, R. 2019, *MNRAS*, **485**, 4160
- Mueller, E. M., Rezaie, M., Percival, W. J., et al. 2021, arXiv e-prints [arXiv:2106.13725]
- Münchmeyer, M., Madhavacheril, M. S., Ferraro, S., et al. 2019, *Phys. Rev. D*, **100**, 083508
- Neal, R. 2011, *Handbook of Markov Chain Monte Carlo*, 1st edn. (Chapman and Hall/CRC), 113
- Neal, R. M. 2003, *Ann. Stat.*, **31**, 705
- Nguyen, N.-M., Schmidt, F., Lavaux, G., et al. 2021, *JCAP*, **03**, 058
- Nguyen, N.-M., Schmidt, F., Tucci, B., et al. 2024, *Phys. Rev. Lett.*, **133**, 221006
- Okamoto, T., & Hu, W. 2002, *Phys. Rev. D*, **66**, 063008
- Pardede, K., Di Dio, E., & Castorina, E. 2023, *JCAP*, **09**, 030
- Percival, W. J. 2005, *A&A*, **443**, 819
- Peron, M., Jung, G., Liguori, M., et al. 2024, *JCAP*, **07**, 021
- Planck Collaboration XXIV. 2014a, *A&A*, **571**, A24
- Planck Collaboration XXII. 2014b, *A&A*, **571**, A22
- Planck Collaboration XVII. 2016a, *A&A*, **594**, A17
- Planck Collaboration XX. 2016b, *A&A*, **594**, A20
- Planck Collaboration X. 2020a, *A&A*, **641**, A10
- Planck Collaboration IX. 2020b, *A&A*, **641**, A9
- Planck Collaboration VI. 2020c, *A&A*, **641**, A6
- Porqueres, N., Jasche, J., Lavaux, G., et al. 2019a, *A&A*, **630**, A151
- Porqueres, N., Kodi Ramanah, D., Jasche, J., et al. 2019b, *A&A*, **624**, A115
- Porqueres, N., Heavens, A., Mortlock, D., et al. 2021, *MNRAS*, **502**, 3035
- Porqueres, N., Heavens, A., Mortlock, D., et al. 2022, *MNRAS*, **509**, 3194
- Porqueres, N., Heavens, A., Mortlock, D., et al. 2023, arXiv e-prints [arXiv:2304.04785]
- Potter, D., Stadel, J., & Teyssier, R. 2017, *Comput. Astrophys. Cosmol.*, **4**, 2
- Ramanah, D. K., Lavaux, G., Jasche, J., et al. 2019, *A&A*, **621**, A69
- Regan, D. M., Schmittfull, M. M., Shellard, E. P. S., et al. 2012, *Phys. Rev. D*, **86**, 123524
- Rezaie, M., Ross, A. J., Seo, H.-J., et al. 2021, *MNRAS*, **506**, 3439
- Rezaie, M., Ross, A. J., Seo, H.-J., et al. 2024, *MNRAS*, **532**, 1902
- Riquelme, W., Avila, S., García-Bellido, J., et al. 2023, *MNRAS*, **523**, 603
- Rossiter, S. J., Camera, S., Clarkson, C., et al. 2025, *JCAP*, **07**, 055
- Roth, N., & Porciani, C. 2012, *MNRAS*, **425**, L81
- Salopek, D. S., & Bond, J. R. 1990, *Phys. Rev. D*, **42**, 3936
- Sasaki, M., Välväita, J., & Wands, D. 2006, *Phys. Rev. D*, **74**, 103003
- Scaramella, R., Mellier, Y., Amiaux, J., et al. 2014, *Proc. Int. Astron. Union*, **10**, 375
- Schmidt, F. 2021, *JCAP*, **04**, 032
- Schmidt, F., Jeong, D., & Desjacques, V. 2013, *Phys. Rev. D*, **88**, 023515
- Schmidt, F., Chisari, N. E., & Dvorkin, C. 2015, *JCAP*, **10**, 032
- Schmidt, F., Elsner, F., Jasche, J., et al. 2019, *JCAP*, **01**, 042
- Schmidt, F., Cabass, G., Jasche, J., et al. 2020, *JCAP*, **11**, 008
- Schmittfull, M., & Seljak, U. 2018, *Phys. Rev. D*, **97**, 123540
- Scoccimarro, R. 2000, *ApJ*, **542**, 1
- Scoccimarro, R., Sefusatti, E., & Zaldarriaga, M. 2004, *Phys. Rev. D*, **69**, 103513
- Scoccimarro, R., Hui, L., Manera, M., et al. 2012, *Phys. Rev. D*, **85**, 083002
- Seljak, U. 2009, *Phys. Rev. Lett.*, **102**, 021302
- Seljak, U., Aslanyan, G., Feng, Y., et al. 2017, *JCAP*, **12**, 009
- Shiveshwarkar, C., Brinckmann, T., Loverde, M., et al. 2023, *Phys. Rev. D*, **108**, 103538
- Shiveshwarkar, C., Brinckmann, T., & Loverde, M. 2024, *JCAP*, **05**, 094
- Slosar, A., Hirata, C., Seljak, U., et al. 2008, *JCAP*, **08**, 031
- Stadler, J., Schmidt, F., & Reinecke, M. 2023, *JCAP*, **10**, 069
- Stadler, J., Schmidt, F., & Reinecke, M. 2025a, *JCAP*, **04**, 089
- Stadler, J., Schmidt, F., Reinecke, M., et al. 2025b, *JCAP*, **11**, 055
- Starobinsky, A. A. 1980, *Phys. Lett. B*, **91**, 99
- Stoppra, S., Peiris, H. V., Pontzen, A., et al. 2024, *MNRAS*, **527**, 1244
- Sullivan, J. M., & Chen, S.-F. 2025, *JCAP*, **03**, 016
- Sullivan, J. M., Prijon, T., & Seljak, U. 2023, *JCAP*, **08**, 004
- Tassev, S., Zaldarriaga, M., & Eisenstein, D. J. 2013, *JCAP*, **06**, 036
- Tsaprazi, E., Nguyen, N.-M., Jasche, J., et al. 2022, *JCAP*, **08**, 003
- Tsaprazi, E., Jasche, J., Lavaux, G., & Leclercq, F. 2023, arXiv e-prints [arXiv:2301.03581]
- Tucci, B., & Schmidt, F. 2024, *JCAP*, **05**, 063
- Uhlemann, C., Pajer, E., Pichon, C., et al. 2018, *MNRAS*, **474**, 2853
- Umeh, O., Koyama, K., Maartens, R., et al. 2019, *JCAP*, **05**, 020
- Verde, L., & Matarrese, S. 2009, *ApJ*, **706**, L91
- Verde, L., Wang, L., Heavens, A. F., et al. 2000, *MNRAS*, **313**, 141
- Verde, L., Jimenez, R., Kamionkowski, M., et al. 2001, *MNRAS*, **325**, 412
- Wang, L., & Kamionkowski, M. 2000, *Phys. Rev. D*, **61**, 063504
- Wang, M. S., Beutler, F., & Bacon, D. 2020, *MNRAS*, **499**, 2598
- Yankelevich, V., & Porciani, C. 2019, *MNRAS*, **483**, 2078
- Yip, J. H. T., Biagetti, M., Cole, A., et al. 2024, *JCAP*, **09**, 034
- Yoo, J. 2014, *Phys. Rev. D*, **90**, 123507
- Yoo, J., & Gong, J.-O. 2016, *Phys. Lett. B*, **754**, 94

¹ INAF-Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Via Piero Gobetti 93/3, 40129 Bologna, Italy

² INFN-Bologna, Via Irnerio 46, 40126 Bologna, Italy

³ Institut d'Astrophysique de Paris, 98bis Boulevard Arago, 75014 Paris, France

⁴ Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, Stockholm SE-106 91, Sweden

- ⁵ Institut d'Astrophysique de Paris, UMR 7095, CNRS, and Sorbonne Université, 98 bis boulevard Arago, 75014 Paris, France
- ⁶ Instituto de Física Teórica UAM-CSIC, Campus de Cantoblanco, 28049 Madrid, Spain
- ⁷ CERCA/ISO, Department of Physics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, USA
- ⁸ Dipartimento di Fisica e Scienze della Terra, Università degli Studi di Ferrara, Via Giuseppe Saragat 1, 44122 Ferrara, Italy
- ⁹ Istituto Nazionale di Fisica Nucleare, Sezione di Ferrara, Via Giuseppe Saragat 1, 44122 Ferrara, Italy
- ¹⁰ School of Physics and Astronomy, Queen Mary University of London, Mile End Road, London E1 4NS, UK
- ¹¹ Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town 7535, South Africa
- ¹² Department of Physics, P.O. Box 64, 00014 University of Helsinki, Finland
- ¹³ Helsinki Institute of Physics, Gustaf Hållströmin katu 2, University of Helsinki, Helsinki, Finland
- ¹⁴ Dipartimento di Fisica e Astronomia “G. Galilei”, Università di Padova, Via Marzolo 8, 35131 Padova, Italy
- ¹⁵ INFN-Padova, Via Marzolo 8, 35131 Padova, Italy
- ¹⁶ INAF-Osservatorio Astronomico di Padova, Via dell'Osservatorio 5, 35122 Padova, Italy
- ¹⁷ European Space Agency/ESTEC, Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands
- ¹⁸ Institute Lorentz, Leiden University, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands
- ¹⁹ Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, 52056 Aachen, Germany
- ²⁰ Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth PO1 3FX, UK
- ²¹ INAF-Osservatorio Astronomico di Brera, Via Brera 28, 20122 Milano, Italy
- ²² Dipartimento di Fisica, Università degli Studi di Torino, Via P. Giuria 1, 10125 Torino, Italy
- ²³ INFN-Sezione di Torino, Via P. Giuria 1, 10125 Torino, Italy
- ²⁴ INAF-Osservatorio Astrofisico di Torino, Via Osservatorio 20, 10025 Pino Torinese (TO), Italy
- ²⁵ CEA Saclay, DFR/IRFU, Service d'Astrophysique, Bat. 709, 91191 Gif-sur-Yvette, France
- ²⁶ Institut für Theoretische Physik, University of Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany
- ²⁷ Institut de Recherche en Astrophysique et Planétologie (IRAP), Université de Toulouse, CNRS, UPS, CNES, 14 Av. Edouard Belin, 31400 Toulouse, France
- ²⁸ Université St Joseph; Faculty of Sciences, Beirut, Lebanon
- ²⁹ Departamento de Física, FCFM, Universidad de Chile, Blanco Encalada 2008, Santiago, Chile
- ³⁰ Université Paris-Saclay, CNRS, Institut d'astrophysique spatiale, 91405 Orsay, France
- ³¹ School of Mathematics and Physics, University of Surrey, Guildford, Surrey GU2 7XH, UK
- ³² IFPU, Institute for Fundamental Physics of the Universe, via Beirut 2, 34151 Trieste, Italy
- ³³ INAF-Osservatorio Astronomico di Trieste, Via G. B. Tiepolo 11, 34143 Trieste, Italy
- ³⁴ INFN, Sezione di Trieste, Via Valerio 2, 34127 Trieste TS, Italy
- ³⁵ SISSA, International School for Advanced Studies, Via Bonomea 265, 34136 Trieste TS, Italy
- ³⁶ Dipartimento di Fisica e Astronomia, Università di Bologna, Via Gobetti 93/2, 40129 Bologna, Italy
- ³⁷ INFN-Sezione di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
- ³⁸ Dipartimento di Fisica, Università di Genova, Via Dodecaneso 33, 16146 Genova, Italy
- ³⁹ INFN-Sezione di Genova, Via Dodecaneso 33, 16146 Genova, Italy
- ⁴⁰ Department of Physics “E. Pancini”, University Federico II, Via Cinthia 6, 80126 Napoli, Italy
- ⁴¹ INAF-Osservatorio Astronomico di Capodimonte, Via Moiariello 16, 80131 Napoli, Italy
- ⁴² INFN section of Naples, Via Cinthia 6, 80126 Napoli, Italy
- ⁴³ Instituto de Astrofísica e Ciências do Espaço, Universidade do Porto, CAUP, Rua das Estrelas, PT4150-762 Porto, Portugal
- ⁴⁴ Faculdade de Ciências da Universidade do Porto, Rua do Campo de Alegre, 4150-007 Porto, Portugal
- ⁴⁵ INAF-IASF Milano, Via Alfonso Corti 12, 20133 Milano, Italy
- ⁴⁶ Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Avenida Complutense 40, 28040 Madrid, Spain
- ⁴⁷ Port d'Informació Científica, Campus UAB, C. Albareda s/n, 08193 Bellaterra (Barcelona), Spain
- ⁴⁸ INAF-Osservatorio Astronomico di Roma, Via Frascati 33, 00078 Monteporzio Catone, Italy
- ⁴⁹ Dipartimento di Fisica e Astronomia “Augusto Righi” – Alma Mater Studiorum Università di Bologna, Viale Berti Pichat 6/2, 40127 Bologna, Italy
- ⁵⁰ Instituto de Astrofísica de Canarias, Calle Vía Láctea s/n, 38204 San Cristóbal de La Laguna, Tenerife, Spain
- ⁵¹ Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK
- ⁵² Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK
- ⁵³ European Space Agency/ESRIN, Largo Galileo Galilei 1, 00044 Frascati, Roma, Italy
- ⁵⁴ ESAC/ESA, Camino Bajo del Castillo, s/n., Urb. Villafranca del Castillo, 28692 Villanueva de la Cañada, Madrid, Spain
- ⁵⁵ Université Claude Bernard Lyon 1, CNRS/IN2P3, IP2I Lyon, UMR 5822, Villeurbanne F-69100, France
- ⁵⁶ Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, 1290 Versoix, Switzerland
- ⁵⁷ Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona (IEEC-UB), Martí i Franquès 1, 08028 Barcelona, Spain
- ⁵⁸ Institució Catalana de Recerca i Estudis Avançats (ICREA), Passeig de Lluís Companys 23, 08010 Barcelona, Spain
- ⁵⁹ UCB Lyon 1, CNRS/IN2P3, IUF, IP2I Lyon, 4 rue Enrico Fermi, 69622 Villeurbanne, France
- ⁶⁰ Departamento de Física, Faculdade de Ciências, Universidade de Lisboa, Edifício C8, Campo Grande, PT1749-016 Lisboa, Portugal
- ⁶¹ Instituto de Astrofísica e Ciências do Espaço, Faculdade de Ciências, Universidade de Lisboa, Campo Grande, 1749-016 Lisboa, Portugal
- ⁶² Department of Astronomy, University of Geneva, ch. d'Ecogia 16, 1290 Versoix, Switzerland
- ⁶³ INAF-Istituto di Astrofisica e Planetologia Spaziali, via del Fosso del Cavaliere, 100, 00100 Roma, Italy
- ⁶⁴ Université Paris-Saclay, Université Paris Cité, CEA, CNRS, AIM, 91191 Gif-sur-Yvette, France
- ⁶⁵ Space Science Data Center, Italian Space Agency, via del Politecnico snc, 00133 Roma, Italy
- ⁶⁶ FRACTAL S.L.N.E., calle Tulipán 2, Portal 13 1A, 28231 Las Rozas de Madrid, Spain
- ⁶⁷ Max Planck Institute for Extraterrestrial Physics, Giessenbachstr. 1, 85748 Garching, Germany
- ⁶⁸ Universitäts-Sternwarte München, Fakultät für Physik, Ludwig-Maximilians-Universität München, Scheinerstrasse 1, 81679 München, Germany
- ⁶⁹ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, 0315 Oslo, Norway
- ⁷⁰ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
- ⁷¹ Felix Hormuth Engineering, Goethestr. 17, 69181 Leimen, Germany
- ⁷² Technical University of Denmark, Elektrovej 327, 2800 Kgs. Lyngby, Denmark
- ⁷³ Cosmic Dawn Center (DAWN), Denmark
- ⁷⁴ Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France

- ⁷⁵ Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany
- ⁷⁶ NASA Goddard Space Flight Center, Greenbelt, MD 20771, USA
- ⁷⁷ Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, UK
- ⁷⁸ Department of Physics and Helsinki Institute of Physics, Gustaf Hållströmin katu 2, 00014 University of Helsinki, Finland
- ⁷⁹ Aix-Marseille Université, CNRS/IN2P3, CPPM, Marseille, France
- ⁸⁰ Université de Genève, Département de Physique Théorique and Centre for Astroparticle Physics, 24 quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland
- ⁸¹ NOVA optical infrared instrumentation group at ASTRON, Oude Hoogeteensedijk 4, 7991PD Dwingeloo, The Netherlands
- ⁸² Universität Bonn, Argelander-Institut für Astronomie, Auf dem Hügel 71, 53121 Bonn, Germany
- ⁸³ INFN-Sezione di Roma, Piazzale Aldo Moro, 2 – c/o Dipartimento di Fisica, Edificio G. Marconi, 00185 Roma, Italy
- ⁸⁴ Aix-Marseille Université, CNRS, CNES, LAM, Marseille, France
- ⁸⁵ Dipartimento di Fisica e Astronomia “Augusto Righi” – Alma Mater Studiorum Università di Bologna, via Piero Gobetti 93/2, 40129 Bologna, Italy
- ⁸⁶ Department of Physics, Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK
- ⁸⁷ Université Paris Cité, CNRS, Astroparticule et Cosmologie, 75013 Paris, France
- ⁸⁸ Institut de Física d’Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona), Spain
- ⁸⁹ School of Mathematics, Statistics and Physics, Newcastle University, Herschel Building, Newcastle-upon-Tyne NE1 7RU, UK
- ⁹⁰ DARK, Niels Bohr Institute, University of Copenhagen, Jagtvej 155, 2200 Copenhagen, Denmark
- ⁹¹ Centre National d’Études Spatiales – Centre spatial de Toulouse, 18 avenue Edouard Belin, 31401 Toulouse Cedex 9, France
- ⁹² Institute of Space Science, Str. Atomistilor, nr. 409 Măgurele, Ilfov 077125, Romania
- ⁹³ Departamento de Astrofísica, Universidad de La Laguna, 38206 La Laguna, Tenerife, Spain
- ⁹⁴ Consejo Superior de Investigaciones Científicas, Calle Serrano 117, 28006 Madrid, Spain
- ⁹⁵ Universität Innsbruck, Institut für Astro- und Teilchenphysik, Technikerstr. 25/8, 6020 Innsbruck, Austria
- ⁹⁶ Institut d’Estudis Espacials de Catalunya (IEEC), Edifici RDIT, Campus UPC, 08860 Castelldefels, Barcelona, Spain
- ⁹⁷ Atlantis, University Science Park, Sede Bld 48940, Leioa-Bilbao, Spain
- ⁹⁸ Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain
- ⁹⁹ Instituto de Astrofísica e Ciências do Espaço, Faculdade de Ciências, Universidade de Lisboa, Tapada da Ajuda, 1349-018 Lisboa, Portugal
- ¹⁰⁰ Universidad Politécnica de Cartagena, Departamento de Electrónica y Tecnología de Computadoras, Plaza del Hospital 1, 30202 Cartagena, Spain
- ¹⁰¹ Kapteyn Astronomical Institute, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands
- ¹⁰² Dipartimento di Fisica, Università degli studi di Genova, and INFN-Sezione di Genova, via Dodecaneso 33, 16146 Genova, Italy
- ¹⁰³ Infrared Processing and Analysis Center, California Institute of Technology, Pasadena, CA 91125, USA
- ¹⁰⁴ INAF, Istituto di Radioastronomia, Via Piero Gobetti 101, 40129 Bologna, Italy
- ¹⁰⁵ ICL, Junia, Université Catholique de Lille, LITL, 59000 Lille, France
- ¹⁰⁶ Department of Physics, Royal Holloway, University of London, TW20 0EX, UK
- ¹⁰⁷ Mullard Space Science Laboratory, University College London, Holmbury St Mary, Dorking, Surrey RH5 6NT, UK
- ¹⁰⁸ ICSC – Centro Nazionale di Ricerca in High Performance Computing, Big Data e Quantum Computing, Via Magnanelli 2, Bologna, Italy

Appendix A: Additional tests of the adiabatic curvature fluctuations

To further test our generated data products, we compute how well they relate to the ground truth adiabatic curvature field. We compare the power spectra of the ensemble-predicted adiabatic curvature fluctuations with the ground truth. The results can be seen in Fig. A.1. Furthermore, in each voxel, we have stored three values: i) the ground truth value (from the mock data itself), ii) the mean inferred estimate, and iii) the uncertainty of the estimate, which is given in terms of $\sigma(\mathcal{R})$. Thus, in the range of voxels, we evaluate whether the ground truth value is within the CIs 68.3%, 95.4%, or 99.7%. Our results show that in the 256^3 voxels, 69.7 % are in the $1\sigma(\mathcal{R})$ range, 95.6 % are in the $2\sigma(\mathcal{R})$ range and 99.8 % are in the $3\sigma(\mathcal{R})$ range. The test also shows that roughly 49 % of the inferred voxels fall above the ground truth value, while 51 % of the inferred voxels fall below the ground truth value. Thus, our inferred $\mathcal{R}$ maps are representative of the ground truth $\mathcal{R}$ map, to the expected confidence.

Figs. A.2– A.5: We present the ensemble statistics of the adiabatic curvature fluctuation maps, in the Mollweide projection, for a distance of $r = 2250 h^{-1}$ Mpc. The fields shown are the averages of the ground truth, the mean inferred, the standard deviation of the inferred, and the residuals between the ground truth and inferred in a single direction. The fields are multiplied by the selection value in each direction, effectively setting the masked directions to zero.




Fig. A.1. Ensemble power spectra statistics of the inferred adiabatic curvature fluctuations, relative to the ground truth. The grey region is the 68% scatter around the mean power spectrum of the ensemble.




Fig. A.2. Mollweide projection of the ground truth adiabatic curvature fluctuation map. The projection is computed for a distance of $r = 2250 h^{-1}$ Mpc, for an observer placed in the centre of the cube, and multiplied by the window selection function.




Fig. A.3. Similar to Fig. A.2, but for the mean inferred adiabatic curvature fluctuation map.




Fig. A.4. Similar to Fig. A.2, but for the uncertainty of the inferred adiabatic curvature fluctuation map.




Fig. A.5. Similar to Fig. A.2, but for the residual adiabatic curvature fluctuation map, defined as $\mathcal{R}^{\text{ground truth}} - \mathcal{R}^{\text{mean inferred}}$.

Appendix B: Tests of the algorithm

This section presents additional plots to offer further insight into the algorithm's performance. The plots encompass autocorrela-




Fig. B.1. Correlation length of all chains, illustrating the rate at which samples in the various chains achieve independence. The typical correlation length for the chains is around 10 000 samples when bias parameters are also sampled.




Fig. B.2. Correlation matrix of Run #4 (the primary run). The correlation matrix illustrates the pairwise relationships among variables, with colour-coding indicating the strength and direction of correlations, aiding in the identification of patterns and dependencies within the dataset. The results show little to no correlation, except for a mild anti-correlation between f_{NL} and the linear bias values.

tion lengths, correlation matrices, and corner plots that illustrate the relationships between f_{NL} and the galaxy bias parameters.

Figure B.1 shows the autocorrelation lengths for f_{NL} from all chains. The autocorrelation length is a metric used in MCMC algorithms to assess how quickly the samples in the chain become independent. Longer lengths indicate slower convergence, necessitating more iterations for reliable parameter estimates. On the contrary, shorter lengths denote faster convergence with reduced reliance on past samples. This metric is vital for evaluating MCMC efficiency, which impacts parameter estimation speed and computational requirements. In particular, the correlation length of f_{NL} in the chains is approximately 10 000 samples for each chain, excluding Run #5.

The correlation matrix summarises the relationships between variables in a dataset. It shows the magnitudes and direction




Fig. B.3. Correlation matrix of Run #7 (which includes the sampling of the scale-dependent bias parameters, b_ϕ and $b_{\phi,\delta}$). Colour coding indicates the strength and direction of the correlations, illustrating the little to no correlation between f_{NL} and the non-linear bias parameters, including the scale-dependent bias parameters.

of linear associations between pairs of variables. High positive values indicate strong positive correlations, while high negative values imply strong negative correlations. A correlation close to zero suggests a weak or no linear relationship. This matrix can also be used to identify potential degeneracies. The Pearson correlation coefficient r_{ij} is defined as:

$$r_{ij} = \frac{\text{cov}(X_i, X_j)}{\sigma_i \sigma_j}. \quad (\text{B.1})$$

Here, X_i and X_j are the bias parameters or f_{NL} , $\text{cov}(X_i, X_j)$ is the covariance, and σ_i and σ_j are the standard deviations. As can be seen in Figs. B.2 and B.3, the galaxy bias parameters exhibit little or no correlations with f_{NL} .

On the same note, the corner plot depicts the interaction between the f_{NL} and galaxy bias parameters and offers a visual representation of their joint distribution and correlations. This plot, visualised in Figs. B.4–B.7, B.8–B.10, and 8 displays the marginal distributions of each parameter on the diagonal and their joint distributions on the off-diagonal. It illustrates how changes in one parameter are related to changes in the other, providing information on potential relationships and regions of interest between f_{NL} and galaxy bias.

Figure B.11 shows trace plots of the inferred values of f_{NL} for the seven different runs, together with the corresponding effective sample sizes (ESS). Each panel displays the sampled values over the chains, along with the posterior mean and the 1σ interval. All chains fluctuate around a well-defined mean within the prior range, and no signs of long-term drifts or divergences are visible, indicating stable convergence. The ESS values, reported in the final panel, confirm that despite differences in chain length and correlation structure, the effective number of independent samples remains sufficiently large for reliable statistical inference. In particular, the main run, Run #4, shows good mixing properties. Overall, the combined diagnostics suggest that the inference of f_{NL} is robust across different simulation setups, and that the recovered values are statistically stable.




Fig. B.4. Corner plot for f_{NL} and bias parameters, for Run #4, catalogue 1. The corner plot displays the joint distributions and marginal distributions of the variables in the multidimensional dataset covered by f_{NL} and the bias parameters. Each subplot captures the relationships between pairs of variables, offering an overview of the dataset structure and dependencies. For the main run, there are few to no degeneracies in the bias parameters.

Appendix C: Inferred density field and data projections

This section showcases plots of the inferred density fields and a comprehensive comparison by including the ground truth mock data fields.

Fig. C.1: We present the averages of the mock data field in three different directions. Notice how the data are cut due to the combination of the radial selection function and the completeness mask.

Fig. C.2: We present the averages of the ground truth, the mean inferred, the variance of the inferred, and the residuals of the present-day dark matter field, δ_{m} , in a single direction.

To emphasise, the core outcome of our method is the inferred f_{NL} distribution; these inferred cosmic fields, with uncertainty estimates, stem from the byproduct of the field-level inference approach.




Fig. B.5. Corner plot for f_{NL} and bias parameters, for Run #4, catalogue 2. Similar to Fig. B.4, there are little to no degeneracies between f_{NL} and the bias parameters.




Fig. B.6. Corner plot for f_{NL} and bias parameters, for Run #4, catalogue 3. Similar to Fig. B.4, there are little to no degeneracies between f_{NL} and the bias parameters.




Fig. B.7. Corner plot for f_{NL} and bias parameters, for Run #4, catalogue 4. Similar to Fig. B.4, there are few to no degeneracies between f_{NL} and the bias parameters.




Fig. B.8. Corner plot for f_{NL} and bias parameters, for Run #7, catalogue 1. Similar to Fig. B.4, there are little to no degeneracies between f_{NL} and the bias parameters.




Fig. B.9. Corner plot for f_{NL} and bias parameters, for Run #7, catalogue 2. Similar to Fig. B.4, there are little to no degeneracies between f_{NL} and the bias parameters.




Fig. B.10. Corner plot for f_{NL} and bias parameters, for Run #7, catalogue 3. Similar to Fig. B.4, there are few to no degeneracies between f_{NL} and the bias parameters.




Fig. B.11. Trace plots of the inferred values of f_{NL} for the seven runs. Each panel shows the thinned chains, displaying only every 1,000th sample, with the raw chain in red, the posterior mean in black (solid line), and the 1σ interval in gray (dashed lines). The last panel summarises the effective sample sizes (ESS) for all runs, which are computed for the entire, unthinned chain. All chains exhibit stable convergence around well-defined mean values, with sufficiently large ESS to ensure reliable inference.




Fig. C.1. Averaged projections of the mock data fields. The colour bar displays the number of galaxies in each pixel, which contains the sum of all galaxies in the summed-over axis. The image is intended to demonstrate the effects of the window function on the observed data and how the method can account for it. The galaxy field projected here is used for Run #4. Each pixel covers a width of $8000 h^{-1} \text{ Mpc} / 256 = 31.25 h^{-1} \text{ Mpc}$.




Fig. C.2. Averaged projections of the ground truth and statistical summaries of the inferred fields. The edge of the survey is highlighted with dotted lines, which means that voxels outside the edge are not observed. In the left panel, the ground truth density field is plotted. In the middle panel, the mean of the ensemble of the inferred fields is plotted. In the right panel, the standard deviation of the ensemble of the inferred fields is plotted. The image illustrates the method's capability to recover the ground truth density field within the regions of observed data. We note that the voxels within the window selection function have less uncertainty and larger inferred means. The inferred fields are the product of Run #4. Each pixel covers a width and height of $8000 h^{-1} \text{ Mpc} / 256 = 31.25 h^{-1} \text{ Mpc}$.