

## Improving constraints on baryon acoustic oscillations with field-level inference

Ivana Babić,<sup>\*</sup> Fabian Schmidt<sup>†</sup>, and Beatriz Tucci<sup>‡</sup>

*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85748 Garching, Germany*



(Received 4 June 2025; accepted 4 March 2026; published 25 March 2026)

We present results of field-level inference of the baryon acoustic oscillation (BAO) scale  $r_s$  on rest-frame dark matter halo catalogs. Our field-level constraint on  $r_s$  is obtained by explicitly sampling the initial conditions along with the bias and noise parameters via the LEFTfield EFT-based forward model. Comparing with a standard reconstruction pipeline applied to the same data and over the same scales, the field-level constraint on the BAO scale improves by a factor of  $\sim 1.2$ – $1.4$  over standard BAO reconstruction. We point to a surprisingly simple source of additional information.

DOI: [10.1103/rkzj-q18k](https://doi.org/10.1103/rkzj-q18k)

**Introduction.** The baryon acoustic oscillation (BAO) feature is currently one of the prime science targets of galaxy redshift surveys. This feature is so important because it allows us to measure the distance-redshift relation  $d_A(z)$  as well as the Hubble rate  $H(z)$  as a function of redshift, thus probing the evolution of dark energy, as well as the mass of neutrinos. The unique large-scale oscillatory pattern of the BAO feature (of order 100 Mpc in comoving units) allows for these robust measurements by avoiding significant degeneracies with other cosmological parameters, as well as the bias parameters describing the relation between galaxies and matter on large scales. Conversely, the broadband shape of galaxy clustering as a function of scale probes the evolution of the growth of structure in the Universe, as well as the above-mentioned bias parameters. Importantly, the BAO scale measurement is strictly relative to  $r_s$ , the sound horizon at the baryon drag epoch, which is well constrained assuming standard early-Universe evolution, but could be modified in the presence of early dark energy or additional relativistic degrees of freedom before recombination. So, keeping the late-time expansion history constrained, one can turn the BAO scale measurement around into a constraint on  $r_s$  and thus early-Universe physics. The potential of BAO measurements is highlighted by the recent DESI DR2 measurements, demonstrating distance constraints relative to  $r_s$  at the level of 0.5%–1% [1].

While the BAO feature, thanks to its shape and large-scale nature, is largely preserved in the late-time galaxy distribution, it is both damped and slightly shifted by the nonlinear evolution of large-scale structure. In particular, the large-scale displacements of matter and galaxies under gravity essentially smooth out the feature. The damping effect leads to a loss in the statistical precision of the BAO measurement from galaxy clustering, because it reduces the contrast of the feature. Moreover, the small systematic shift induced in the inferred scale is no longer negligible given the substantial volume covered by current galaxy surveys.

In order to deal with these issues, the current state of the art is to perform a so-called *BAO reconstruction*, where the large-scale displacements responsible for the BAO damping and shift are estimated from the galaxy distribution itself [2–4]. Then, the observed galaxy positions are moved by the opposite of the estimated displacement, thus attempting to reverse the damping effect on the BAO feature. Finally, the BAO scale is estimated from the power spectrum of the galaxy density field measured in terms of the displaced positions. The detailed algorithm is described in Sec. B in Supplemental Material (SM) [5]. While elegant and easy to implement, the standard BAO reconstruction approach suffers from several drawbacks. First, the displacement estimate is necessarily based on assumptions of a fiducial cosmology and a value for the linear galaxy bias parameter. Second, standard BAO reconstruction does not remove the effects of nonlinear galaxy formation (nonlinear bias), which, as we will see, also play a role in the precision of the BAO scale measurement. Finally, the displacement is estimated from the galaxy density field only at linear order (Zel’dovich approximation); improvements on this approximation have been suggested in the literature [6–17], but so far have not been applied to large redshift surveys.

Here, we present an alternative approach to BAO inference that addresses all of these shortcomings: field-level Bayesian

<sup>\*</sup>Contact author: [I.Babic@physik.uni-muenchen.de](mailto:I.Babic@physik.uni-muenchen.de)

<sup>†</sup>Contact author: [fabians@mpa-garching.mpg.de](mailto:fabians@mpa-garching.mpg.de)

<sup>‡</sup>Contact author: [tucci@mpa-garching.mpg.de](mailto:tucci@mpa-garching.mpg.de)

*Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI. Open access publication funded by the Max Planck Society.*

inference (FLI). This approach jointly estimates the BAO scale, large-scale displacements, and bias parameters in a consistent Bayesian framework, by explicitly sampling the initial conditions, i.e., the linear matter density field  $\delta^{(1)}(\mathbf{k})$ , where the BAO feature is imprinted. More precisely, the BAO scale enters the inference via the structure of the prior on the initial conditions, which is determined by the linear power spectrum  $P_L(k|r_s)$  which depends on  $r_s$  [Eqs. (3) and (4)]. The accuracy of the BAO inference is then limited only by the accuracy of the *forward model* for the observed galaxy density field,  $\delta_g[\delta^{(1)}]$  (the results presented here will be in fact on simulated dark matter halo catalogs). The latter is in fact likely to be the key challenge for field-level inference [18–20]: since we cannot simulate realistic galaxy catalogs in a cosmological volume, and certainly not at a computational cost that would allow for field-level inference, model misspecification is inherent in any forward model, and the robustness of the results against such misspecification must be validated. Unlike the standard reconstruction approach, field-level inference uses the entire available dataset, and a field-level inference of the BAO scale that is robust to model misspecification is a nontrivial demand.

The forward model employed here, LEFTfield [21], is based on a Lagrangian formulation of the effective field theory (EFT) of large-scale structure [22–24]. This model incorporates the dependence of the galaxy density on all local observables of a comoving observer, while enforcing Einstein’s equivalence principle, which ensures that matter and galaxies comove on large scales [25]. Indeed, this fact allows for the inference of large-scale displacements, and thus is the basis also for the standard BAO reconstruction. Unlike the model employed in standard BAO reconstruction, our forward model includes all contributions up to third order in perturbation theory, and is able to predict the halo density field accurately up to wave numbers of  $k \sim 0.2 \, h \text{ Mpc}^{-1}$  [26–29].

We should highlight previous studies that performed field-level inference of distances from large-scale structure. Reference [30] presented an inference of cosmology via Alcock-Paczynski distortions on mock data generated from the forward model, demonstrating substantial improvement over the power spectrum without reconstruction. Recently, Ref. [31] demonstrated field-level BAO-scale inference using LEFTfield on mock catalogs generated by the EFT model at a higher scale than used in the inference, showing robustness to model mismatch.

Here, we demonstrate for the first time that the model can successfully perform field-level inference of the BAO scale on fully nonlinear tracers: dark matter halos in N-body simulations. Given that our forward model is agnostic to the details of the tracer considered—and based only on the equivalence principle—we expect that our results generalize to actual galaxies as well (see [32,33] for field-level results on simulated galaxies) [5].

*Data.* We consider two halo samples. The SNG sample consists of main halos in the  $\log_{10} M_{200m} = 12.5\text{--}14.0 \, h^{-1} M_\odot$  mass range, identified with the ROCKSTAR halo finder [34] at redshift  $z = 0.50$  in an N-body, gravity-only simulation. This sample has a mean comoving number density of  $\bar{n} = 1.3 \times 10^{-3} \, (h^{-1} \text{ Mpc})^{-3}$ . The simulation assumes a flat Lambda cold dark matter ( $\Lambda$ CDM) cosmology, encompasses a comoving volume  $L^3 = (2000 \, h^{-1} \text{ Mpc})^3$  and contains  $N_{\text{particle}} = 1536^3$  particles of mass  $M_{\text{particle}} = 1.8 \times 10^{11} \, h^{-1} M_\odot$  [27]. The second, Uchuu halo sample, consists of main halos in the  $\log_{10} M_{200m} = 12.0\text{--}13.5 \, h^{-1} M_\odot$  mass range, identified with ROCKSTAR at redshift  $z = 1.03$  ( $\bar{n} = 3.6 \times 10^{-3} \, (h^{-1} \text{ Mpc})^{-3}$ ) in the Uchuu simulation [35], which assumes a different flat  $\Lambda$ CDM cosmology. This simulation spans the same volume  $L^3 = (2000 \, h^{-1} \text{ Mpc})^3$  while offering a much higher mass resolution,  $N_{\text{particle}} = 12800^3$  particles of mass  $M_{\text{particle}} = 3.27 \times 10^8 \, h^{-1} M_\odot$ . Notice that the initial conditions for the Uchuu simulations were not available to the authors. Throughout, the model is evaluated at the fixed redshift of the respective sample, and we drop time and redshift arguments for clarity in the following.

*Field-level EFT of biased tracers.* The EFT of large-scale structure provides a perturbative framework within which the galaxy or halo density field  $\delta_g$  can be systematically expanded, order by order in perturbations, as

$$\delta_g[\delta^{(1)}](\mathbf{k}) = \sum_{\mathcal{O}} b_{\mathcal{O}} O[\delta^{(1)}](\mathbf{k}) + \epsilon(\mathbf{k}). \quad (1)$$

At each given order, there is a finite number of galaxy bias operators  $O[\delta^{(1)}]$ , each associated with a coefficient  $b_{\mathcal{O}}$ . This includes  $b_\delta \delta[\delta^{(1)}](\mathbf{k})$ , where  $\delta[\delta^{(1)}]$  is the *forward-evolved* matter density, which is the well-known linear relation that is employed in the standard BAO reconstruction. Notably, Eq. (1) includes several additional nonlinear bias terms, namely the complete set of six bias operators at second and third order; see Sec. A1 (SM) [5] for the detailed description. Note that the EFT counterterms for matter itself are a subset of the bias terms and hence already included in Eq. (1). The bias operators are jointly computed with the nonlinear evolution, for which we adopt second-order Lagrangian perturbation theory (2LPT).

The second term on the rhs of Eq. (1) is the stochastic contribution encapsulating the random nature of small-scale fluctuations and the discrete nature of galaxies. These are described by the noise field  $\epsilon$ , which to leading order is Gaussian with rms

$$\sigma_\epsilon(k) = \sigma_{\epsilon,0} [1 + \sigma_{\epsilon,k^2} k^2], \quad (2)$$

where the subleading contribution  $\propto k^2$  captures the finite extent of galaxy-forming regions.

*Field-level forward model.* We forward model the bias fields  $O$  in Eq. (1) starting from Gaussian initial conditions parametrized via a unit Gaussian random field  $\hat{s} \sim \mathcal{N}(0, 1)$ , discretized on a grid of size  $N_{\text{grid}}$ . The linear density field is related to  $\hat{s}$  by

$$\delta^{(1)}(\mathbf{k}, z|r_s) = f(k, r_s) \left[ \frac{N_{\text{grid}}^3}{L^3} P_{\text{L}}^{\text{fid}}(k, z) \right]^{1/2} \hat{s}(\mathbf{k}), \quad (3)$$

where  $L$  is the simulation box side length,  $N_{\text{grid}}$  is the size of the initial conditions grid, and  $P_{\text{L}}^{\text{fid}}(k, z)$  is the linear power spectrum in the fiducial (simulation) cosmology. We implement a change in the BAO scale  $r_s$  via the function  $f(k, r_s)$  defined by

$$f(k, r_s) = \frac{T_{\text{BAO}}^2(k|r_s)}{T_{\text{BAO}}^2(k|r_{s,\text{fid}})},$$

$$T_{\text{BAO}}^2(k|r_s) = 1 + A \sin(kr_s + \phi) \exp(-k/k_D), \quad (4)$$

where  $r_{s,\text{fid}}$  is the comoving sound horizon in the fiducial cosmology (see Sec. A2 (SM) [5] for how we determine the value of  $r_{s,\text{fid}}$  and of the transfer function parameters  $A$ ,  $\phi$  and  $k_D$ ). Thus, unlike what is done in BAO scale inferences on actual data, where the angular diameter distance to a given redshift is varied, we instead vary the BAO scale in the linear density field. We adopt this approach, since it enables a change in the BAO scale while retaining the periodic boundary conditions within the simulation volume, which is a significant technical simplification. We stress that the exact same approach for varying the BAO scale is used in both the field-level and standard reconstruction analyses; thus the comparison is consistent, and we expect it to be an accurate indicator for what would be obtained when varying the distance instead.

Our forward model employs the Lagrangian, EFT-based forward model of cosmological density fields, LEFTfield [21]. Following the EFT principle, LEFTfield evolves all cosmological (plus auxiliary) fields up to a finite EFT cutoff  $\Lambda$  [24,26,36]. Specifically, we implement a cubic sharp- $k$  filter on the initial conditions. LEFTfield computes both  $O = O(\hat{s})$  and  $\partial O / \partial \hat{s}$ , the latter of which is essential for the Hamiltonian Monte Carlo technique employed. We refer to [20,21,26,29] for details on the implementation and validation. The forward model adopted here corresponds to the third-order Lagrangian bias forward model presented as alternative in Supplemental Material of [37], where comparable posteriors on the parameter  $\sigma_8$  were reported from both Eulerian and Lagrangian bias expansions. The Lagrangian bias expansion avoids an additional cut on the evolved density field, yielding improved numerical convergence properties [29], and thus is preferred.

*Field-level BAO inference.* In the field-level Bayesian inference pipeline, we evaluate and sample from an explicit field-level likelihood  $\mathcal{L}_{\text{FLI}}$ . Following [38], our analyses

assume Gaussianity of galaxy stochasticity and analytically marginalize over  $\epsilon$ . This leads to a Gaussian log-likelihood for an observed galaxy field  $\delta_g^{\text{obs}}$  of the following form [36,39]:

$$\mathcal{L}_{\text{FLI}}(\delta_g^{\text{obs}}|\hat{s}, r_s, \{b_O\}, \{\sigma_{\epsilon,i}\}) = -\frac{1}{2} \sum_{\mathbf{k} \neq 0}^{|k| < k_{\text{max}}} \left[ \ln 2\pi\sigma_{\epsilon}^2(k) + \frac{1}{\sigma_{\epsilon}^2(k)} \left| \delta_g^{\text{obs}}(\mathbf{k}) - \sum_O b_O O[\hat{s}, r_s](\mathbf{k}) \right|^2 \right]. \quad (5)$$

The  $\sum_{\mathbf{k} \neq 0}^{|k| < k_{\text{max}}}$  amounts to a spherical sharp- $k$  filter which only includes Fourier modes of the data up to  $k_{\text{max}}$ . We also exclude the  $\mathbf{k} = \mathbf{0}$  mode, as it is degenerate with the effective mean density of the sample. For the analyses presented here, we mostly adopt  $k_{\text{max}} = \Lambda$  (we will also discuss results with  $\Lambda > k_{\text{max}}$ ), but note that the initial conditions filter  $\Lambda$  is cubic and thus includes modes up to  $\sqrt{3}\Lambda$ . We expand  $\sum_O b_O O$  to third order in the galaxy bias operators  $O$ , and further analytically marginalize over all bias coefficients except  $b_\delta$  (this is possible thanks to the simple dependence of Eq. (5) on the  $b_O$  [40]) assuming weakly informative Gaussian priors (see Sec. C (SM) [5]).

The final explicitly sampled FLI parameter space consists of  $\{r_s, \{\hat{s}\}, b_\delta, \sigma_{\epsilon,0}, \sigma_{\epsilon,k^2}\}$ , where  $\hat{s}$  is a three-dimensional grid of size  $N_{\text{grid}}^3$ . To explore this high-dimensional posterior, following [20,41], we employ two Monte Carlo Markov chain sampling methods: Hamiltonian Monte Carlo [42] for  $\{\hat{s}\}$ —leveraging the differentiability of the LEFTfield forward models—and slice sampling [43] for  $\{r_s, b_\delta, \sigma_{\epsilon,0}, \sigma_{\epsilon,k^2}\}$ . The posterior validation, effective sample size, and correlation lengths are presented in Sec. D (SM) [5].

*BAO inference from the postreconstruction power spectrum.* For comparison and benchmarking, we aim to replicate a standard BAO reconstruction pipeline, taking care not to include modes beyond the cutoff scale  $k_{\text{max}}$  considered in the FLI analysis. For this purpose, we assign the tracer catalog to a density grid using the same grid resolution and assignment scheme used in the FLI analysis. We then filter the assigned tracer density with a Gaussian filter on the scale  $R = \Lambda^{-1}$  before estimating the displacement. After reassigning the tracer positions to their estimated Lagrangian positions, we evaluate the postreconstruction power spectrum  $P_{\text{p-rec}}(k)$  up to the same scale  $k_{\text{max}} = \Lambda$  used in the FLI. We fit a model including a broadband part parametrized with five free parameters, and obtain the best-fit and error bar on  $r_s$  via Monte Carlo sampling using a Gaussian power spectrum likelihood. The details are described in Sec. B (SM) [5].

We note that we fix the linear bias  $b_1$  in the displacement estimation, and any uncertainty on  $b_1$  is not propagated into the error on  $r_s$ . Further, the power spectrum likelihood

assumes a Gaussian covariance, and as such ignores any additional correlations induced by mode coupling (this approximation has been validated for DESI in Ref. [44]). In contrast, both of these sources of uncertainty are incorporated in the posterior on  $r_s$  obtained in the FLI analysis. In [31] we have validated both analyses on mock data, and showed that standard reconstruction and FLI agree where they should, namely in the case of a linearly biased tracer of the 1LPT (Zel'dovich) density field with fixed linear bias.

**Results.** Our main results are shown in Fig. 1, where we show, as a function of the maximum wave number included in the analysis, the inferred posterior mean and RMS for  $r_s$  relative to the fiducial value  $r_{s,\text{fid}}$ . Unbiased inference thus corresponds to a value consistent with 1. While the standard reconstruction approach yields unbiased BAO constraints, FLI yields slightly biased results for the two higher cutoffs in the SNG sample. This bias can be traced back to our choice of  $\Lambda = k_{\text{max}}$ . To ensure that the EFT bias terms can completely absorb the  $\Lambda$  dependence of the model, one should ideally choose  $\Lambda \gg k_{\text{max}}$  [45,46], which however is numerically costly. We have subsequently performed inferences choosing  $\Lambda = (3/2)k_{\text{max}}$  for the SNG sample. The results (Fig. 1) indicate that this choice removes the systematic shift in  $r_s$  [47]. FLI results for the higher-redshift Uchuu sample show a smaller systematic bias, which we expect will likewise be improved when increasing  $\Lambda/k_{\text{max}}$ .

In Fig. 1 we also indicate the ratio of the 68% CL error on  $r_s$  between the postreconstruction power spectrum analysis and the FLI result. Despite taking fully into account the reconstruction uncertainties and marginalizing over all EFT bias terms, FLI generally improves over the BAO scale uncertainty in the standard reconstruction analysis, for both halo samples, by 10%–40%. The only exception is the intermediate cutoff for the SNG sample, where FLI is also most significantly biased. This data point should be revisited with the  $\Lambda > k_{\text{max}}$  inference in the future. Notice that the absolute error bars depend sensitively on  $k_{\text{max}}$ , and current analyses typically use larger  $k_{\text{max}}$  values than presented here. However, we expect our conclusions to continue to hold, or become even stronger, at higher  $k_{\text{max}}$  (see [37]), provided we remain within the range of validity of the EFT approach.




FIG. 1. Inferred BAO scale using FLI (blue circles) and postreconstruction power spectrum (orange boxes) as a function of  $k_{\text{max}}$ . The upper panel shows results for the SNG halo sample, while the lower shows the Uchuu halo sample. We also indicate the ratio in BAO scale uncertainties from both methods; FLI increases the BAO scale precision by 10%–40%, with the exception of the intermediate cutoff in the SNG case. The blue stars show FLI results using the same scales of the data, but with increased  $\Lambda$  (see text for discussion).

*Where does the additional information come from?* The information gain demonstrated in Fig. 1 can have several causes. We investigate this analytically by considering the limit of infinitely informative (zero-noise) data following [46,48]. Section E (SM) [5] derives the Fisher information on the BAO scale in both FLI and reconstruction approaches in this simplified setting. The result is

$$F_{r_s r_s}^{\text{FLI}} = - \left\langle \frac{\partial^2}{\partial r_s^2} \ln \mathcal{P}_{\text{FLI}}[\{b_O\}, r_s | \delta_g] \Big|_{r_{s,\text{fid}}} \right\rangle = \frac{1}{2} \sum_{\mathbf{k}}^{\Lambda} \frac{1}{[P_L(k|r_{s,\text{fid}})]^2} \left( \frac{\partial P_L(k|r_{s,\text{fid}})}{\partial r_s} \right)^2, \quad (6)$$

$$F_{r_s r_s}^{\text{rec-P(k)}} = - \left\langle \frac{\partial^2}{\partial r_s^2} \ln \mathcal{P}_{\text{rec-P(k)}}[r_s | \delta_g] \Big|_{r_{s,\text{fid}}} \right\rangle = \frac{1}{2} \sum_{\mathbf{k}}^{\Lambda} \frac{1}{[P_{\text{p-rec}}(k|r_{s,\text{fid}})]^2} \left( \frac{\partial P_{\text{p-rec}}(k|r_s)}{\partial r_s} \right)^2. \quad (7)$$

Comparing Eqs. (7) and (6) we can identify the following factors that contribute to increased information on the BAO scale at the field level:

- (1) *Improved reconstruction quality:* in FLI, the inferred initial density field  $\delta_{g,\text{det}}^{-1}[\delta_g, \{b_O\}]$  is a representative sample from the Gaussian distribution characterized

by the linear power spectrum  $P_L(k|r_s)$ , provided that the forward model correctly describes the data on the range of scales considered. The inferred initial density field thus carries the full BAO information encoded in  $(\partial P_L/\partial r_s)^2$ . On the other hand, in the standard reconstruction-based analysis, this is replaced by  $(\partial P_{\text{p-rec}}/\partial r_s)^2$ , which could be suppressed relative to  $\partial P_L/\partial r_s$ , for example because the 1LPT-based reconstruction does not restore the BAO wiggles as well as the 2LPT forward model employed in FLI. To test this, we have also performed an inference with the 1LPT forward model for matter, keeping all aspects of the forward model including the bias expansion the same. This case yields very similar constraints on  $r_s$  as the fiducial inference using the 2LPT forward model, which indicates that reconstruction quality is not the dominant effect on the scales considered here.

- (2) *Reduced variance*: the variance of each mode in the FLI is controlled by  $P_L(k|r_s)$ . This is illustrated by the maximum *a posteriori* (MAP) sample of the initial conditions in the FLI analysis shown in Fig. 2, which differs from the denominator  $P_{L,\text{no-wiggle}}$  only via the BAO oscillations. On the other hand, the variance in the standard reconstruction analysis is determined by  $P_{\text{p-rec}}(k|r_s)$ . Since the standard reconstruction approach only reverts large-scale displacements,  $P_{\text{p-rec}}$  is not close to  $P_L$ , but has a substantially increased broadband part due to contributions from nonlinear evolution and nonlinear bias (notice that the linear bias drops out of the zero-noise Fisher information in both cases). This broadband contribution is clearly visible in Fig. 2 (shown there for the SNG sample; the corresponding figure for the Uchuu sample is shown in Sec. D (SM) [5]). In fact, this increased variance, estimated from Fig. 2, can by itself explain at least a significant part of the observed information gain. The non-Gaussian contributions to the covariance of  $P_{\text{p-rec}}$  neglected in Eq. (7) generally work in the same direction.

The simplified Fisher analysis above did not consider that the BAO feature in the postreconstruction power spectrum is extracted via template fitting. The additional free parameters in the template can in principle further reduce the information over that in Eq. (7). Figure 1 in SM [5] indicates that this should not be a dramatic degradation, but could be noticeable at the  $\sim 10\%$  level. One might wonder then if it is possible to improve the standard BAO reconstruction to also remove the broadband contribution, i.e. to make  $P_{\text{p-rec}}$  as close as possible to  $P_L$ . However, this is not easy to do, since the broadband shape arises from nonlinear mode coupling across a broad range of scales due to both nonlinear evolution and nonlinear bias. It would thus not only require a full inverse model of galaxy clustering, but also information on the nonlinear bias parameters. FLI on the other hand jointly infers




FIG. 2. Ratio of the pre- and postreconstruction halo power spectrum of the SNG sample to the smooth, “no-wiggle” part of the linear power spectrum multiplied by the best-fit  $b_1^2$ . The BAO feature is enhanced postreconstruction, but there is a significant broadband enhancement over the linear power spectrum. On the other hand, FLI reconstructs the linear density field and thus “measures” the BAO scale without a broadband enhancement. This is illustrated by the power spectrum of the maximum *a posteriori* initial conditions sample from one of the SNG FLI chains (green dotted).

the higher-order bias parameters (essentially using higher-order  $n$ -point functions of the data [46]) with the BAO scale, so that the power spectrum of the inferred initial density field is indeed very close to  $P_L(k|r_s)$  (Fig. 2). This conclusion is also supported by the test on linearly biased 1LPT mocks presented in [31].

*Summary and discussion.* Field-level inference of the BAO scale is the unique method that allows for a consistent inference of cosmological distances and parameters while reaping the benefits of BAO reconstruction. As such, it is worth pursuing even if it yields comparable error bars to standard reconstruction approaches, which keep cosmological parameters in the reconstruction step fixed. Here, we have presented the first field-level BAO inference on fully nonlinear tracers. We have found that, compared to a standard BAO reconstruction pipeline using the same scales of the data, an improvement on the BAO uncertainty by up to 40% is achievable. The more accurate forward model employed in FLI enables several sources of information gain. The likely dominating factor for our results here is the incorporation of nonlinear bias (coupled with nonlinear evolution), which is responsible for a significant contribution to the postreconstruction broadband power spectrum and hence error bar.

In this context, it is worth discussing the assumption of a fixed broadband linear power spectrum shape in the analysis here, which, it should be stressed, is assumed consistently in both FLI and reconstruction analyses. In an application to real data, one would exploit the fact that the FLI analysis implements a consistent Bayesian inference

framework, and would thus jointly vary all cosmological parameters that describe the linear power spectrum as well as the distance-redshift relation in the analysis. Such an inference would consistently combine BAO, broadband, and growth-of-structure information.

One important constraint of the EFT-based FLI analysis is that it is restricted to perturbative scales. On the other hand, standard BAO analyses often fit the BAO feature to smaller, nonlinear scales, which are no longer under perturbative control. Still, we stress that we have not yet fully exploited the range of scales accessible to the perturbative treatment in the results presented here, especially at redshifts  $z \gtrsim 1$ . Moreover, when pushing to smaller yet perturbative scales, the improved BAO sharpening due to the higher-order gravity forward model could lead to additional improvements. It would further be interesting to validate our analysis on simulated galaxy, rather than halo catalogs. Even more importantly, the inclusion of redshift-space distortions following [49] is not only essential for the

application to real data, but could lead to further gains of field-level BAO constraints over standard reconstruction. We will pursue this in upcoming work.

Our postprocessing pipeline is empowered by ROOT, pyROOT, GetDist, NUMPY, EMCEE, and MATPLOTLIB.

*Acknowledgments.* We thank Andrija Kostić, Minh Nguyen, Martin Reinecke, Ariel Sánchez, Julia Stadler, and Uroš Seljak for helpful discussions. F. S. acknowledges discussions and collaborations with Aquila consortium [50] members on previous works that eventually led to these results. Our analyses are performed on the freya and orion clusters, maintained by the Max Planck Computing & Data Facility [51], as well as the ada cluster hosted and maintained by MPA.

*Data availability.* No data were created or analyzed in this study.

- 
- [1] M. Abdul-Karim *et al.* (DESI Collaboration), [arXiv:2503.14738](#).
  - [2] D. J. Eisenstein, H. Seo, E. Sirko, and D. N. Spergel, *Astrophys. J.* **664**, 675 (2007).
  - [3] D. J. Eisenstein, H. Seo, and M. White, *Astrophys. J.* **664**, 660 (2007).
  - [4] N. Padmanabhan, M. White, and J. D. Cohn, *Phys. Rev. D* **79**, 063523 (2009).
  - [5] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/rkzj-q18k> for details on the forward model, priors, posterior validation, and the Fisher calculation.
  - [6] Y. Shi, M. Cautun, and B. Li, *Phys. Rev. D* **97**, 023505 (2018).
  - [7] Y. Noh, M. White, and N. Padmanabhan, *Phys. Rev. D* **80**, 123501 (2009).
  - [8] S. Tashev and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **10** (2012) 006.
  - [9] A. Burden, W. J. Percival, and C. Howlett, *Mon. Not. R. Astron. Soc.* **453**, 456 (2015).
  - [10] M. Schmittfull, Y. Feng, F. Beutler, B. Sherwin, and M. Y. Chu, *Phys. Rev. D* **92**, 123522 (2015).
  - [11] X. Wang, H.-R. Yu, H.-M. Zhu, Y. Yu, Q. Pan, and U.-L. Pen, *Astrophys. J.* **841**, L29 (2017).
  - [12] M. Schmittfull, T. Baldauf, and M. Zaldarriaga, *Phys. Rev. D* **96**, 023505 (2017).
  - [13] C. Modi, Y. Feng, and U. Seljak, *J. Cosmol. Astropart. Phys.* **10** (2018) 028.
  - [14] M. McQuinn, *J. Cosmol. Astropart. Phys.* **06** (2021) 024.
  - [15] C. J. Shallue and D. J. Eisenstein, *Mon. Not. R. Astron. Soc.* **520**, 6256 (2023).
  - [16] X. Chen, F. Zhu, S. Gaines, and N. Padmanabhan, *Mon. Not. R. Astron. Soc.* **523**, 6272 (2023).
  - [17] L. Parker, A. E. Bayer, and U. Seljak, *J. Cosmol. Astropart. Phys.* **09** (2025) 039.
  - [18] N.-M. Nguyen, F. Schmidt, G. Lavaux, and J. Jasche, *J. Cosmol. Astropart. Phys.* **03** (2021) 058.
  - [19] P. Villanueva-Domingo and F. Villaescusa-Navarro, *Astrophys. J.* **937**, 115 (2022).
  - [20] A. Kostić, N.-M. Nguyen, F. Schmidt, and M. Reinecke, *J. Cosmol. Astropart. Phys.* **07** (2023) 063.
  - [21] F. Schmidt, *J. Cosmol. Astropart. Phys.* **04** (2021) 033.
  - [22] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, *J. Cosmol. Astropart. Phys.* **07** (2012) 051.
  - [23] J. J. M. Carrasco, M. P. Hertzberg, and L. Senatore, *J. High Energy Phys.* **09** (2012) 082.
  - [24] S. M. Carroll, S. Leichenauer, and J. Pollack, *Phys. Rev. D* **90**, 023518 (2014).
  - [25] V. Desjacques, D. Jeong, and F. Schmidt, *Phys. Rep.* **733**, 1 (2018).
  - [26] F. Schmidt, G. Cabass, J. Jasche, and G. Lavaux, *J. Cosmol. Astropart. Phys.* **11** (2020) 008.
  - [27] F. Schmidt, *J. Cosmol. Astropart. Phys.* **04** (2021) 032.
  - [28] I. Babić, F. Schmidt, and B. Tucci, *J. Cosmol. Astropart. Phys.* **08** (2022) 007.
  - [29] J. Stadler, F. Schmidt, and M. Reinecke, *J. Cosmol. Astropart. Phys.* **04** (2025) 089.
  - [30] D. K. Ramanah, G. Lavaux, J. Jasche, and B. D. Wandelt, *Astron. Astrophys.* **621**, A69 (2019).
  - [31] I. Babić, F. Schmidt, and B. Tucci, *J. Cosmol. Astropart. Phys.* **11** (2025) 066.
  - [32] A. Barreira, T. Lazeyras, and F. Schmidt, *J. Cosmol. Astropart. Phys.* **08** (2021) 029.
  - [33] E. Krause *et al.* (Beyond-2pt Collaboration), *Astrophys. J.* **990**, 99 (2025).

- [34] P. S. Behroozi, R. H. Wechsler, and H.-Y. Wu, *Astrophys. J.* **762**, 109 (2013).
- [35] T. Ishiyama *et al.*, *Mon. Not. R. Astron. Soc.* **506**, 4210 (2021).
- [36] F. Schmidt, F. Elsner, J. Jasche, N. M. Nguyen, and G. Lavaux, *J. Cosmol. Astropart. Phys.* **01** (2019) 042.
- [37] N.-M. Nguyen, F. Schmidt, B. Tucci, M. Reinecke, and A. Kostić, *Phys. Rev. Lett.* **133**, 221006 (2024).
- [38] G. Cabass and F. Schmidt, *J. Cosmol. Astropart. Phys.* **07** (2020) 051.
- [39] G. Cabass and F. Schmidt, *J. Cosmol. Astropart. Phys.* **04** (2020) 042.
- [40] F. Elsner, F. Schmidt, J. Jasche, G. Lavaux, and N.-M. Nguyen, *J. Cosmol. Astropart. Phys.* **01** (2020) 029.
- [41] J. Jasche and G. Lavaux, *Astron. Astrophys.* **625**, A64 (2019).
- [42] R. Neal, in *Handbook of Markov Chain Monte Carlo* (Chapman and Hall, New York, 2011), pp. 113–162.
- [43] R. M. Neal, [arXiv:physics/0009028](https://arxiv.org/abs/physics/0009028).
- [44] D. Forero-Sánchez *et al.*, *J. Cosmol. Astropart. Phys.* **04** (2025) 055.
- [45] H. Rubira and F. Schmidt, *J. Cosmol. Astropart. Phys.* **01** (2024) 031.
- [46] F. Schmidt, *J. Cosmol. Astropart. Phys.* **09** (2025) 056.
- [47] Due to the limited sample size, we do not show the inferred error bar for these inferences, but we expect them to be very similar to the  $\Lambda = k_{\max}$  cases.
- [48] G. Cabass, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **109**, 043526 (2024).
- [49] J. Stadler, F. Schmidt, M. Reinecke, and M. Esposito, *J. Cosmol. Astropart. Phys.* **11** (2025) 055.
- [50] <https://www.aquila-consortium.org/>.
- [51] <https://www.mpcdf.mpg.de/>.