

Baryon acoustic oscillation theory and modelling systematics for the DESI 2024 results

S.-F. Chen^{1,★†}, C. Howlett^{2,★}, M. White^{3,4}, P. McDonald⁴, A. J. Ross⁵, H.-J. Seo⁶, N. Padmanabhan⁷, J. Aguilar⁴, S. Ahlen⁸, S. Alam⁹, O. Alves¹⁰, U. Andrade^{10,11}, R. Blum¹², D. Brooks¹³, X. Chen⁷, S. Cole¹⁴, K. Dawson¹⁵, A. de la Macorra¹⁶, A. Dey¹², Z. Ding¹⁷, P. Doel¹³, S. Ferraro^{4,18}, A. Font-Ribera^{13,19}, D. Forero-Sánchez²⁰, J. E. Forero-Romero^{21,22}, C. Garcia-Quintero^{23,24}, E. Gaztañaga^{25,26,27}, S. G. A Gontcho⁴, M. M. S Hanif¹⁰, K. Honscheid^{5,28,29}, T. Kisner⁴, A. Kremin⁴, A. Lambert⁴, M. Landriau⁴, M. E. Levi⁴, M. Manera^{19,30}, A. Meisner¹², J. Mena-Fernández³¹, R. Miquel^{19,32}, A. Munoz-Gutierrez¹⁶, E. Paillas^{33,34}, N. Palanque-Delabrouille^{4,35}, W. J. Percival^{33,34,36}, A. Pérez-Fernández^{16,37}, F. Prada³⁸, M. Rashkovetskyi²³, M. Rezaie³⁹, A. Rosado-Marin⁶, G. Rossi⁴⁰, R. Ruggeri^{2,41}, E. Sanchez⁴², D. Schlegel⁴, J. Silber⁴, G. Tarlé¹⁰, M. Vargas-Magaña¹⁶, B. A. Weaver¹², J. Yu²⁰, S. Yuan⁴³, R. Zhou⁴ and Z. Zhou⁴⁴

Affiliations are listed at the end of the paper

Accepted 2024 September 3. Received 2024 September 3; in original form 2024 February 22

ABSTRACT

This paper provides a comprehensive overview of how fitting of baryon acoustic oscillations (BAO) is carried out within the upcoming Dark Energy Spectroscopic Instrument’s (DESI) 2024 results using its DR1 data set, and the associated systematic error budget from theory and modelling of the BAO. We derive new results showing how non-linearities in the clustering of galaxies can cause potential biases in measurements of the isotropic (α_{iso}) and anisotropic (α_{ap}) BAO distance scales, and how these can be effectively removed with an appropriate choice of reconstruction algorithm. We then demonstrate how theory leads to a clear choice for how to model the BAO and develop, implement, and validate a new model for the remaining smooth-broad-band (i.e. without BAO) component of the galaxy clustering. Finally, we explore the impact of all remaining modelling choices on the BAO constraints from DESI using a suite of high-precision simulations, arriving at a set of best practices for DESI BAO fits, and an associated theory and modelling systematic error. Overall, our results demonstrate the remarkable robustness of the BAO to all our modelling choices and motivate a combined theory and modelling systematic error contribution to the post-reconstruction DESI BAO measurements of no more than 0.1 per cent (0.2 per cent) for its isotropic (anisotropic) distance measurements. We expect the theory and best practices laid out here to be applicable to other BAO experiments in the era of DESI and beyond.

Key words: cosmological parameters – distance scale – large-scale structure of Universe – cosmology: theory.

1 INTRODUCTION

Constraining the expansion history of the Universe is important both for improving our empirical cosmographical knowledge and for the constraints it places on the constituents of the Universe and our theories of gravity. Within the standard Friedmann–Lemaître–Robertson–Walker (Friedmann 1922; Lemaître 1931; Robertson 1935; Walker 1937) model, the scale factor $a(t)$ – is the only degree of freedom at the background level, determining distance relations over cosmic history. Constraining its behaviour has been a major goal of cosmology for close to a century. One of the most robust and accurate ways of determining $a(t)$ is through measurements of

baryon acoustic oscillations (BAO) in the distribution of galaxies and gas. Oscillations in the baryon-photon fluid prior to decoupling leave an imprint in both the cosmic microwave background (CMB) and the matter density with a characteristic length-scale, related to the distance r_d a sound wave can travel prior to the baryon-drag epoch, just after recombination. Measurements of this scale at low redshift in large-scale structure provide a ‘standard ruler’ for constraining the expansion history of the Universe, both by measuring the angular diameter distance [$D_A(z)$] and the *Hubble* parameter ($H(z)$) to a specific redshift. This ruler can be made ‘absolute’ by connecting the BAO scale back to early universe physics, for example, through constraints from the CMB (Peebles & Yu 1970; Sunyaev & Zeldovich 1970; Dodelson & Schmidt 2020; Huterer 2023) or big-bang nucleosynthesis (BBN; Cooke, Pettini & Steidel 2018).

The BAO method has become a workhorse of modern cosmology and provides some of our tightest current constraints on the

* E-mail: sfschen@ias.edu (SC); c.howlett@uq.edu.au (CH)

† Both authors contributed equally to this work.

cosmological expansion history (Cole et al. 2005; Eisenstein et al. 2005; Beutler et al. 2011; Blake et al. 2011; Alam et al. 2021). Relying as it does on a feature at large spatial separations that is little affected by non-linear evolution or the complex astrophysics of galaxy formation, it is robust and theoretically well understood. For this reason, it is generally regarded as a low-systematics method for constraining cosmology. Indeed, significant theoretical effort over the past two decades has led to a fairly mature understanding of the effect of the small shifts and modulations to the BAO signal incurred by large-scale structure physics which we recap and build on in this work (Section 2).

The purpose of this paper is to use our theoretical understanding to generate a theory and modelling systematics budget, and suggest fitting best practices, for BAO measurements in the Dark Energy Spectroscopic Instrument (DESI). DESI is a Stage-IV spectroscopic instrument (DESI Collaboration 2016, 2022) that will deliver galaxy BAO measurements with unprecedented precision up to $z \approx 2$ from galaxy clustering alone, with even greater reach when Lyman- α data are included at higher redshift. At the end of its five-year run DESI is expected to yield measurements on the (isotropic) BAO scale at a cumulative 0.2 per cent precision (DESI Collaboration 2016, 2024d).¹ Already, galaxies from the early data release in the luminous red galaxy (LRG) sample and bright galaxy sample (BGS), representing a small fraction of the total DESI data, yield BAO detections at the few percent level (Moon et al. 2023; DESI Collaboration 2024e).

This paper is part of a set focused on the analysis of clustering in DESI Data Release 1 (DR1; DESI Collaboration in preparation-a, b), and in support of the key paper presenting the main galaxy BAO measurements for that data set (DESI Collaboration 2024a). The DR1 data further contain BAO information in the Lyman- α forest (DESI Collaboration 2024b), and the cosmological implications of the joint galaxy and Lyman- α BAO measurements are discussed in DESI Collaboration (2024c). Information beyond the BAO – particularly the full-shape information in the galaxy two-point function – are studied in DESI Collaboration (in preparation-c, d, e). While the aggregate isotropic BAO scale precision from analysis of DR1 is closer to 0.5 per cent, our conclusions and recommendations are somewhat general and aimed towards the full DESI sample. Our understanding of galaxy clustering in the standard model of cosmology (lambda-cold dark matter, Λ CDM) gives us a baseline of known theoretical systematics that need to be considered in order to make robust measurements of the BAO scale, even before unknown modifications due to nonstandard physics are considered.

The structure of the paper is as follows: in Section 2, we begin with a broad overview of the BAO signal as observed in large-scale structure simulations. In Sections 3 and 4, we review the theory of the BAO and BAO reconstruction, computing error estimates due to theory systematics, filling in extant gaps in the literature where necessary, and motivating our baseline form for fitting the BAO signal. In Section 5, we introduce a cubic-spline based method for marginalizing over the non-BAO broad band of galaxy clustering and harmonize broad-band models in Fourier and configuration space. Section 6 introduces the N -body simulations and inference pipeline BARRY implementing the suggested analysis choices in this paper, and

the systematic effects of various numerical choices are examined in Section 7. We conclude in Section 8.

2 THE OBSERVED BARYON ACOUSTIC OSCILLATION SIGNAL: AN OVERVIEW

BAOs manifest as an oscillatory, or wiggly, imprint in the linear matter power spectrum $P_{\text{lin}}(k)$. The physics is very well understood and can be computed to high accuracy by numerical evolution of the coupled Einstein, fluid and Boltzmann equations as is done, for example, by CAMB (Lewis, Challinor & Lasenby 2000; Howlett et al. 2012) and CLASS (Blas, Lesgourgues & Tram 2011). To recap the basic physics, prior to recombination, when the densities are high, baryon-photon scattering is rapid compared with the traveltime across a wavelength. This allows the behaviour of the baryon and photons to be described as a single fluid, whose continuity and Euler equations reduce to a driven harmonic oscillator with slowly varying natural frequency. During this ‘tight-coupling’ phase, the baryon perturbations hence also undergo harmonic motion with very slowly decaying amplitude (Peebles & Yu 1970; Sunyaev & Zeldovich 1970). These ‘BAOs’ become ‘frozen-in’ at decoupling when the photons release their hold on the baryons, setting the oscillation frequency r_d , and the matter perturbations begin to grow as $\propto a$. The combination of oscillatory densities and velocities which project onto the growing mode gives the final post-recombination spectrum.² This then grows in a scale-independent fashion until the present day, up to small corrections due to massive neutrinos which start-off relativistic but transition to non-relativistic at some point post-recombination. The BAOs themselves are superposed upon the smooth power spectrum, whose shape is primarily determined by the epoch of matter-radiation equality. The amplitude of the oscillations depends upon the driving force (gravitational potentials) and the baryon-to-photon ratio. Larger oscillations come from higher baryon density ω_b and/or smaller total-matter density ω_m . Additionally, light species (e.g. neutrinos) which propagate ahead of the sound waves also exert a gravitational influence on the BAO which can shift their phase (Bashinsky & Seljak 2004), and any isocurvature modes in the initial conditions can also alter the shape of the observed signal (Zunckel et al. 2011); the impact of some these effects on DESI BAO measurements is discussed in Pérez-Fernández et al. (2024), but we note that they are also more directly constrained by the CMB itself.

Given a linear matter power spectrum P_{lin} , the redshift-space galaxy power spectrum in linear theory and in the plane-parallel or distant-observer approximation, is given by (Kaiser 1987)

$$P_{gg}(k, \mu) = (b_1 + f\mu^2)^2 P_{\text{lin}}(k), \quad (1)$$

where $f = d \ln D / d \ln a \approx \Omega_m^{0.55}$ is the growth rate and b_1 is the linear (scale-independent, deterministic) bias. Since the linear spectrum is otherwise smooth, and neither the non-linearities of structure/galaxy formation nor survey systematics are expected to produce features with the same frequency, the aim of BAO measurements in galaxy surveys is to independently extract the oscillatory part of this linear signal from the overall clustering of galaxies and use it as a standard ruler. The BAO feature is a localized peak in the correlation function at a scale determined by the sound horizon at the baryon-drag epoch (r_d), or equivalently as a series of oscillations in the power

¹Where the cumulative or ‘aggregate’ BAO precision is the statistical uncertainty we would expect in the BAO scale if all DESI galaxies were put into the same redshift bin and used to produce a single BAO measurement. It is equivalent to an inverse sum of the BAO precision across all our independent measurement bins.

²At low k , this is a mixture of density and velocity modes but at higher k it is primarily the velocity mode which is why the higher harmonics of the BAO are out of phase with the peaks in the CMB, the latter of which arise primarily from photon density perturbations.

spectrum with Fourier frequency $2\pi/r_d$. At a given redshift, z , the line-of-sight (LOS) and transverse determinations of this ruler fix $H(z)r_d$ and $r_d/D_A(z)$, respectively. Modern spectroscopic surveys typically convert the measured redshifts and sky positions of galaxies into comoving coordinates using a fiducial cosmology (usually flat Λ CDM with a fiducial $\Omega_{m,\text{fid}} \approx 0.3$), in which the multipoles of the power spectrum $P(k, \mu) = \sum_\ell P_\ell(k)\mathcal{L}_\ell(\mu)$ and correlation functions $\xi(r, \mu)$ can be computed in physical units. These two-point functions are then fitted against a BAO ‘template’ that is scaled in the LOS and transverse directions (Padmanabhan & White 2008) to extract the so-called dilation parameters

$$\alpha_{\parallel} = \frac{H^{\text{fid}}(z)r_s^{\text{tem}}}{H(z)r_s}, \quad \alpha_{\perp} = \frac{D_A(z)r_s^{\text{tem}}}{D_A^{\text{fid}}(z)r_s}, \quad (2)$$

where ‘fid’ and ‘tem’ refer to quantities in the fiducial and template cosmologies, respectively. It is also common to see the BAO dilation parameters recast into a purely isotropic term $\alpha_{\text{iso}} \equiv \alpha_{\parallel}^{1/3}\alpha_{\perp}^{2/3}$ and an anisotropic term $\alpha_{\text{ap}} \equiv (1 + \epsilon)^3 \equiv \alpha_{\parallel}/\alpha_{\perp}$. We will use these all interchangeable in this work.

Let us describe the observation and extraction of the BAO in galaxy surveys outlined above in more detail, beginning with the dynamical model. In order to do so, it is helpful to split the power spectrum into a BAO template consisting of the wiggles alone P_w and a ‘no-wiggle’ portion P_{nw} containing the broad-band power which we shall discuss further below (Section 7.4)

$$P_{\text{lin}}(k) = P_{nw}(k) + f_b P_w(k). \quad (3)$$

In the definition of the wiggle/no-wiggle split above we have included in the wiggly component a pre-factor of the baryon fraction $f_b \approx 15$ per cent. This small parameter controls the size of the BAO signal relative to the overall clustering signal, and we will use it as a convenient power-counting parameter to denote the suppression of contributions involving higher powers of the wiggle component. However, since f_b is degenerate with the amplitude of P_w we will generally include it in P_w for convenience in the rest of the paper, except when it is useful for book-keeping like here. The measurement of the linear BAO signal can then be thought of as the measurement of the linear dependence of the observed galaxy power spectrum, P_{gg} , on f_b .³ Schematically, for the observed two-point function we can write

$$P_{gg}(k, \mu) = (P_{gg})_{f_b=0} + f_b \left(\frac{dP_{gg}}{df_b} \right)_{f_b=0} + \mathcal{O}(f_b^2). \quad (4)$$

Based on this structure models of the BAO typically take the form

$$P_{gg}(k, \mu) = \mathcal{B}(k, \mu)P_{nw}(k) + \mathcal{C}(k, \mu)P_w(k) + \mathcal{D}(k, \mu), \quad (5)$$

where the term proportional to $\mathcal{C}(k, \mu)$ is the leading BAO signal $dP_{gg}/df_b = \mathcal{C}(k, \mu)P_w$ we want to measure, with $\mathcal{C}(k, \mu) = (b + f\mu^2)^2$, in linear theory. The rough non-oscillatory broad-band shape is contained in the $\mathcal{B}(k, \mu)$ term, given by the Kaiser formula in linear theory. The final term $\mathcal{D}(k, \mu)$ represents any remaining contributions

³There is some ambiguity in this ‘split’ as, in addition to generating the acoustic oscillations, the fact that the baryon perturbations oscillate rather than grow between matter-radiation equality and photon-baryon decoupling means the potentials evolve as they would in a universe with only dark matter $\Omega_m = \Omega_c$ (i.e. the baryons don’t contribute to the gravitational collapse). This alters the shape of P_{lin} just to the right of the matter-radiation equality peak leading to a faster decrease than would otherwise occur. Ultimately, this means that $P_{nw}(k)$ in a Universe with baryons is not the same as the power spectrum of pure CDM.

to the observed power spectrum, which are largely due to non-linear coupling of the smooth linear spectrum and observational systematics. It may also absorb higher order (e.g. $(f_b)^2$) terms in the BAO, though these will tend to be both rather suppressed and oscillatory. The systematics we will discuss in this paper largely have to do with the goodness of this approximation; specifically we will consider the extent to which non-linear oscillatory terms modify the $\mathcal{C}(k, \mu)$ signal we wish to measure, and how well the remaining terms can be absorbed by $\mathcal{D}(k, \mu)$ under the assumption of smoothness.

Given one adopts a fiducial cosmology for converting galaxy redshifts to distances, and a template linear spectrum with corresponding r_d^{tem} and P_w^{tem} , the assumption is then that any of the scale information carried in the BAO wiggles can be accounted for with a simple scaling $P_w(k) \sim P_w^{\text{tem}}(r_d k/r_d^{\text{tem}})$ in equation (5). Allowing for $H(z)$ and $D_A(z)$ to vary independently away from the fiducial cosmology, such that

$$\mathbf{k}_{\text{obs}} = \left(\frac{D_A(z)}{D_A^{\text{fid}}(z)} \mathbf{k}_{\perp}^{\text{true}}, \frac{H^{\text{fid}}(z)}{H(z)} k_{\parallel}^{\text{true}} \right) \quad (6)$$

leads to the conclusion that the free parameters required to ‘match’ the template to the observed BAO are precisely the dilation parameters defined in equation (2), with smooth changes to the broad band and the shape and normalization of the wiggles captured by marginalizing over $\mathcal{B}, \mathcal{C}, \mathcal{D}$. This borrows philosophically from the Alcock–Paczynski effect (Alcock & Paczynski 1979), which describes the dilation of scales and angles when analysing galaxy positions assuming a fiducial cosmology that differs from the true underlying cosmology of the Universe.

It is worth noting that not all of the cosmological dependence of the BAO signal (e.g. its amplitude, damping or phase shifts) is captured by this rescaling – indeed, some of these, such as the phase shift, cannot even be captured by other parameters like galaxy bias or broad-band modelling in standard BAO fits. In general, the physical effects that change P_{lin} also change the CMB anisotropies, often more significantly (Eisenstein & White 2004), and tend to be disfavoured by existing measurements (e.g. Planck Collaboration VI 2020). The systematic errors incurred by these additional cosmology-dependent effects are characterized in an accompanying DESI paper (Pérez-Fernández et al. 2024), so we will only note here that the extraction of $\alpha_{\parallel, \perp}$ in BAO fits has been shown to be rather robust to these (Vargas-Magaña et al. 2014; Bernal et al. 2020; Carter et al. 2020), while referring interested readers to that upcoming companion paper for further details in the context of DESI Y1. In addition, the definition of P_w is not unique and several different template extraction methods exist in the literature. Properly constructed, these P_w should only differ by smooth functions of k that are degenerate with the marginalized broad-band terms, and we will explore the extent to which this is true in Section 7.4.

Our interest in the rest of this paper will be in quantifying the effects of dynamical non-linearities and modelling choices on measuring $\alpha_{\parallel, \perp}$ so as to form a theoretical error budget for DESI BAO analyses. In doing so we will work primarily in the plane-parallel approximation and validate our models against the DESI Y1 ‘cubic mocks’ based on single-redshift snapshots of N -body simulations in periodic, cubic boxes. This isolates the ‘theoretical systematics’ from those depending upon observing geometry or observational non-idealities. However, before moving into these investigations it is important to describe a few additional details necessary to make contact with observations.

First, galaxy surveys measure clustering not in translationally invariant cubic boxes but within defined angular and redshift windows. In the power spectrum this effect, along with wide-angle contribu-

tions to common power spectrum estimators, can be accounted for by convolving theory predictions with suitably computed window matrices (Feldman, Kaiser & Peacock 1994; Blake, Carter & Koda 2018; Beutler & McDonald 2021). Since these effects should ideally be exactly computable given the galaxy window function we will not consider their systematics here but relegate them to the DESI Y1 catalogue papers. Secondly, the measured two-point function is not sampled from a single point but rather a weighted sum across the redshift range $P = \sum_i w_i P(z_i)$; to leading order this is taken into account by taking the BAO measurement to be at the effective redshift, but in cosmological analyses it is also possible to take into account higher order terms, for example,

$$\hat{\alpha}_{\parallel,\perp} = \alpha_{\parallel,\perp}(z_{\text{eff}}) + \frac{1}{2} \left(\frac{d^2 \alpha_{\parallel,\perp}}{dz^2} \right)_{z_{\text{eff}}} \sigma_z^2 + \dots, \quad (7)$$

where σ_z^2 is the mean-square width of the redshift bin. We discuss this further, and define the effective redshift, in Appendix A. Finally, in our modelling of the BAO signal we have ignored various large-scale effects including wide-angle effects, unequal time effects, and relativistic effects. These typically scale as $(H/k)^2$ in the power spectrum and so are extremely suppressed on the range of scales we fit, since our fiducial setup has $k_{\text{min}} = 0.02 h \text{ Mpc}^{-1}$. Many of these effects can produce out-of-phase contributions but their effect on the inferred BAO scale is rather limited (typically less than 0.05 per cent; Appendix B). As such, we will not further consider them in the rest of this work.

3 NON-LINEAR BARYON ACOUSTIC OSCILLATIONS IN GALAXY CLUSTERING

Non-linear structure formation and galaxy bias significantly modify the observed shape and phase of the linear BAO signal, especially pre-reconstruction (Bharadwaj 1996; Meiksin, White & Peacock 1999; Eisenstein, Seo & White 2007a; Crocce & Scoccimarro 2008; Matsubara 2008; Seo et al. 2008; Smith, Scoccimarro & Sheth 2008; Padmanabhan & White 2009; Sherwin & Zaldarriaga 2012; Carlson, Reid & White 2013). Broadly speaking, we can split the effects of these non-linearities into two categories: (1) the damping, or modulation, of the amplitude of the linear wiggles and (2) the shift in the observed BAO scale through out-of-phase contributions to galaxy clustering.

The most important contributions to both of these effects come from the displacements Ψ of galaxies on large scales, which are well described by perturbation theory (PT). Within Lagrangian perturbation theory (LPT), which models structure formation via the trajectories of galaxies, $\mathbf{x}(\mathbf{q}, t) = \mathbf{q} + \Psi(\mathbf{q}, t)$ where \mathbf{q} is the initial Lagrangian position of the galaxy, the galaxy power spectrum in real space can be written as (Matsubara 2008; Carlson et al. 2013)

$$P(k) = \int d^3 \mathbf{q} e^{i\mathbf{k}\cdot\mathbf{q} - \frac{1}{2} k_i k_j A_{ij}(\mathbf{q})} \left(1 + 2b_1^L i k_i \langle \delta(\mathbf{q}) \Delta_i \rangle + (b_1^L)^2 \xi_{\text{lin}}(\mathbf{q}) + 2b_1^L b_2^L i k_i \langle \delta(\mathbf{q}) \Delta_i \rangle \xi_{\text{lin}}(\mathbf{q}) + \dots \right), \quad (8)$$

where $\Delta = \Psi(\mathbf{q}) - \Psi(\mathbf{0})$ is the pairwise displacement between two galaxies separated by \mathbf{q} , $A_{ij} = \langle \Delta_i \Delta_j \rangle$ is its variance, and ξ_{lin} is the linear matter correlation function. The numbers b_n^L are the free parameters in the Lagrangian bias expansion of galaxies, related to the canonical Eulerian ones through simple linear transformations, for example, $b_1 = 1 + b_1^L$ (Desjacques, Jeong & Schmidt 2018). Throughout this work, we will operate under the assumption that galaxies are biased tracers of CDM and baryons but not (massive) neutrinos (Castorina et al. 2015). This has been shown to be an

excellent approximation in simulations – close to the 0.1 per cent level for the galaxy power spectrum in redshift space on the scales we are interested in – and as such we will operate within this regime throughout the text (Bayer, Banerjee & Seljak 2022). We note that the discussion in this section is not meant to constitute an exhaustive treatment of galaxy clustering in PT, and nor is equation (8) meant to be the complete expression for the one-loop galaxy power spectrum; the intent is rather to pick out the pieces of physics most relevant to the BAO signal in order to estimate their systematic effect on BAO measurements.⁴

3.1 Real space

Let us begin with modifications to the linear BAO signal in real space due to only one power of the baryon fraction, that is, which involves contributions of order $\mathcal{O}(P_w \times P_{nw}^n)$. Physically, we can think of these as distortions to the BAO feature due to gravitational non-linearities of the smooth component.

We can see the origin of both the ‘in-phase’ BAO damping and ‘out-of-phase BAO’ shift directly in equation (8). In order to do so it is again useful to split the linear correlation function ξ_{lin} into its contributions from the wiggle and no-wiggle components $\xi_{w(nw)}$ – the former has a sharp peak at r_d . Focusing on the b_1^L piece we can see that the wiggle component can therefore be approximated as (Vlah et al. 2016, see also e.g. Baldauf et al. 2015; Senatore & Zaldarriaga 2015; Blas et al. 2016 for discussions outside of LPT)

$$(b_1^L)^2 \int d^3 \mathbf{q} e^{i\mathbf{k}\cdot\mathbf{q} - \frac{1}{2} k_i k_j A_{ij}(\mathbf{q})} \xi_w(\mathbf{q}) \approx (b_1^L)^2 e^{-\frac{1}{2} k^2 \Sigma_{\text{NL}}^2} \int d^3 \mathbf{q} e^{i\mathbf{k}\cdot\mathbf{q}} \xi_w(\mathbf{q}) = b_1^2 e^{-\frac{1}{2} k^2 \Sigma_{\text{NL}}^2} P_w(k)$$

where Σ_{NL}^2 is the angular average of $A_{ij}(\mathbf{q})$ for $q = r_d$. At first order in LPT, that is, in the Zeldovich approximation, this is given by⁵

$$\Sigma_{\text{NL}}^2 = \frac{2}{3} \int \frac{dk}{2\pi^2} P(k) [1 - j_0(kr_d)]. \quad (9)$$

This gives us the damping of the BAO by the bulk displacements of galaxies. Similar reasoning shows that any ‘wiggle’ component of the power spectrum is similarly damped (Vasudevan et al. 2019; Beutler et al. 2019a; Chen et al. 2020). An important feature of equation (9) is that the damping comes from modes with wavelengths smaller than the BAO, that is, with $k \gtrsim 1/r_d$; physically, this reflects that galaxies separated by r_d are moved coherently by modes larger than the BAO scale, but incoherently by those smaller. In Λ CDM universes similar to our own the square displacement Σ_{NL}^2 is dominated by linear displacements close to the $1/r_d$ cut-off, whose effects are sufficiently large that they have to be kept in the exponential, or resummed, instead of being perturbatively expanded. However, we

⁴See Appendix C and onwards for a more complete accounting for galaxies in redshift space and post-reconstruction.

⁵The saddle-point approximation here involves an angular integral dependent on the relative orientations of \hat{q} and \hat{k} . The approximation we have made here is to only resum the isotropic component, but other resummation choices also exist in the literature, for example, Baldauf et al. (2015) and Blas et al. (2016) have an integrand proportional to $1 - j_0 + 2j_2$ instead, corresponding to resumming the $\hat{q} = \hat{k}$ component. The differences in the BAO amplitude in one-loop PT using these two resummation choices are well below the percent level (see e.g. figs 3 and 5 of Chen, Vlah & White 2020) while in the BAO fitting only case their difference should be well within the priors of Σ_{NL}^2 set in the analysis (see Section 7.2). For a more detailed exploration of the angular integral in LPT see Chen, Vlah & White (2024).

stress that while BAO damping in Λ CDM universes similar to our own is dominated by contributions from relatively long-wavelength modes as described above, small-scale, non-linear, displacements can also contribute. We discuss the importance of these effects when fitting BAO, especially considering redshift-space distortions and reconstruction, further below.

The BAO shift effect can be similarly derived by approximating the non-linear effect of no-wiggle contributions by their values near $q = r_d$ when they multiply well-localized wiggly components such as ξ_w . Let us consider the $b_1 b_2$ term in equation (8). In this case, writing to linear order

$$\langle \delta(\mathbf{q}) \Delta_i \rangle = -\frac{1}{3} q_i \sigma^2(q), \quad (10)$$

where $\sigma^2(q)$ is the mean-square linear density fluctuation in a sphere of radius q ,⁶ we have

$$\begin{aligned} & 2i b_1^L b_2^L i k_i \int d^3 \mathbf{q} e^{i \mathbf{k} \cdot \mathbf{q}} \langle \delta(\mathbf{q}) \Delta_i \rangle \xi_w(\mathbf{q}) \\ & \approx -\frac{2}{3} b_1^L b_2^L \sigma_s^2 k_i \nabla_{k_i} \int d^3 \mathbf{q} e^{i \mathbf{k} \cdot \mathbf{q}} \xi_w(\mathbf{q}) = -\frac{2}{3} \sigma_s^2 b_1^L b_2^L \frac{dP_w}{d \ln k}. \end{aligned} \quad (11)$$

In the above, we have defined $\sigma_s^2 = \sigma^2(r_d)$ to be the mean-square density at the peak of the wiggly correlation function. This directly translates into a physical shift in α_{iso} and, sign-included, has the interpretation that galaxies with positive b_2 are more likely to form in overdense regions, and galaxies sitting in an overdensity will tend to flow towards each other, shrinking the average observed BAO (Crocco & Scoccimarro 2008; Padmanabhan & White 2009; Sherwin & Zaldarriaga 2012). Importantly, this effect is caused by overdensities due to modes with wavelengths longer than r_d . While we have focused on the phase shift due to b_2 , all other quadratic contributions to clustering will produce similar shifts; working in Eulerian perturbation theory (EPT), Sherwin & Zaldarriaga (2012) showed that the real-space shift for a biased tracer with quadratic density bias is⁷

$$\begin{aligned} \Delta \alpha_{\text{iso, NL}} &= \frac{47}{105} \left(1 + \frac{70 b_2}{47 b_1} \right) \sigma_s^2 \\ &\approx 0.3 \text{ per cent} \left(1 + \frac{70 b_2}{47 b_1} \right) \left(\frac{D(z)}{D(0)} \right)^2, \end{aligned} \quad (12)$$

⁶This correlator, $\langle \delta \Delta_i \rangle$, describes the mean inflow of matter towards an overdensity and is more commonly written in the LPT literature as (Carlson et al. 2013)

$$U_i(\mathbf{q}) = -\hat{q}_i \int \frac{dk k}{2\pi^2} P_{\text{lin}}(k) j_1(kq).$$

To relate the two expressions, we note that

$$j_1(kq) = \frac{1}{3} (kq) W_q(k),$$

where $W_q = 3(\sin(kq)/(kq) - \cos(kq))/(kq)^2$ is the Fourier transform of the spherical top hat window. This is a consequence of Birkhoff's theorem, that is, that the inflow velocity depends only on the enclosed mass at radius q .

⁷We note that Sherwin & Zaldarriaga (2012) presented two different expressions for the BAO phase shift, the one we have quoted here and a larger one in the correlation function. However, the two expressions are in fact equivalent under Fourier transform and, since the covariance is (more) diagonal in Fourier space and therefore corrections therein more directly translated into systematic shifts in parameters, we use the former expression throughout this text. In addition, we note that the precise weighting of the k modes in σ_s^2 stems from our LPT calculation.

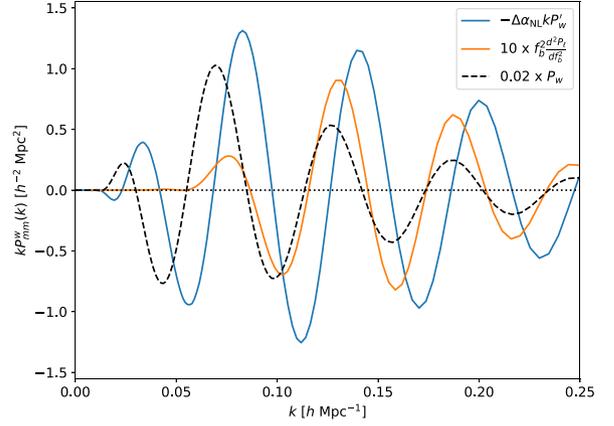


Figure 1. A summary of oscillatory, non-linear, effects that lead to shifts in the measured BAO position in the real-space one-loop matter power spectrum. Exact expressions are given in equations (11)–(13). For reference, the dashed line is the linear theory wiggle power spectrum P_w scaled down to 2 per cent of its nominal amplitude, and the f_b^2 contribution amplified by a factor of 10, for ease of comparison.

where $D(z)$ is the linear growth factor and we have evaluated the shift in the fiducial DESI cosmology.

The BAO feature can also couple with itself to produce additional out-of-phase contributions to the signal. These are contributions of the order $\mathcal{O}(P_w^n)$ and hence f_b^n . For example, at quadratic order we can write the contribution in Fourier space

$$\frac{d}{d f_b^2} P(k) = \int \frac{d^3 \mathbf{p}}{(2\pi)^3} K_2(\mathbf{p}, \mathbf{k} - \mathbf{p}) P_w(p) P_w(|\mathbf{k} - \mathbf{p}|), \quad (13)$$

where K_2 is the quadratic PT kernel for galaxy clustering. Since each power of P_w is sinusoidal with frequency r_d , their convolution in Fourier space can produce a frequency dependence distinct from that in linear theory. In the case of non-linear matter clustering or linear bias we have that K_2 is the quadratic matter kernel $F_2(\mathbf{p}_1, \mathbf{p}_2)$ which peaks when $\mathbf{p}_1 = \mathbf{p}_2 = \mathbf{k}/2$, producing an out-of-phase oscillatory signal $\sin(r_d k/2)^2 \sim \cos(r_d k)$ (Padmanabhan & White 2009). Since this signal only has support within one BAO wiggle of $k/2$, it scales roughly as $f_b^2 (\pi/r_d)^2 k P_w(k/2)^2$, compared to the non-linear shift derived from equations (11) and (12) which scales as $\propto \sigma_s^2 \times k P_w'$ – it is thus subdominant since $k/2$ will typically appear after the peak of the power spectrum and because it contains an extra power of the baryon fraction f_b .

A comparison of both of the BAO shift contributions we have discussed, in the case of matter clustering, is shown in Fig. 1. Note that unlike in the case of a pure sine wave the actual BAO wiggles in the linear power spectrum lead to a $d^2 P/d f_b^2$ that also contains a significant in-phase component, further suppressing the contribution of this effect to the total BAO shift.

In addition to the above effects, which describe the clustering of galaxies assuming dark matter and baryons can be treated as a single matter fluid on cosmological scales, the very same physics that generates the BAO produces residual differences in dark matter and baryon clustering that remain even after decoupling. The dominant effects can be written as (Schmidt 2016; Chen, Castorina & White 2019a)

$$\delta_g \ni b_{\delta bc} \delta_{bc} + b_{\theta bc} \theta_{bc} + b_{v,2} v_{bc}^2, \quad (14)$$

where we have defined the relative linear perturbations as $X_{bc} = X_b - X_c$, with δ and θ the linear densities and velocity divergences of the baryons and dark matter, respectively. The linear relative

densities and biases are constant and decay with redshift as $a^{-1}(z)$, respectively. The third term is the so-called relative velocity effect which was first described in Tseliakhovich & Hirata (2010) as the effect of supersonic baryon-dark matter velocities on gravitational collapse. Unlike the first two terms, it appears only at one-loop order in PT, but can be enhanced as the relative velocity squared is a dimensionful parameter – studies of the relative velocity effect have often used the value $b_{v,2} \sim 0.01\sigma_{bc}^{-2}(z)$ where σ_{bc}^2 is the mean-square fluctuation of relative velocity perturbations, which correspond to roughly 30 km s^{-1} flows at recombination, based on early investigations into this effect.

Compared to the usual galaxy biases, b_1 , b_2 etc., the sizes of these relative bias parameters are somewhat less well understood, though some recent progress has been made. Barreira et al. (2020) measure the relative density bias $b_{\delta_{bc}}$ using the separate-universe technique in hydrodynamical simulations and find distinct trends for dark matter haloes (with halo mass) and galaxies (with stellar mass). While the trend in haloes can be well modelled by the change in amplitude of local fluctuations due to baryon density, as can halo-mass selected galaxies, the trend in galaxies selected by stellar mass is sensitive to other baryonic physics and exhibits an opposite (positive) trend with mass, particularly at the redshifts $z < 2$ that DESI will observe. None the less for all of these competing effects, the inferred bias values for both relevant haloes and galaxies are well described by $|b_{\delta_{bc}}| \lesssim 1$ for the galaxies relevant to DESI. The linear relative velocity divergence bias is less well understood but physical estimates through excursion sets and other methods (Blazek, McEwen & Hirata 2016; Schmidt 2016) suggest $|b_{\theta_{bc}}| \lesssim 7(1+z)^{-1}$, where the redshift dependence cancels the decay of peculiar velocities. The size of the relative velocity effect at DESI redshifts is fairly poorly understood, with estimates as high as $b_{v,2} \sim 0.01\sigma_{bc}^{-2}$ or as low as $10^{-5}\sigma_{bc}^{-2}$ for DESI-like galaxies (Yoo, Dalal & Seljak 2011; Blazek et al. 2016; Schmidt 2016). A direct estimate of the effect of streaming velocities on baryonic infall into haloes gives an order $(10^5 M_\odot / M)^{2/3}$ effect (Blazek et al. 2016), which yields $b_{v,2}$ much closer to the lower estimate for DESI galaxies, while reaching the higher value often used in the literature requires that galaxies today retain significant memory of star formation at early times (Yoo & Seljak 2013).

Fig. 2 shows the leading contributions to the galaxy power spectrum of each of these terms in real space at $z = 0$ (Yoo et al. 2011; Blazek et al. 2016; Schmidt 2016; Chen et al. 2019a), compared to linear theory wiggles. When the relative velocity effect ($b_{v,2}$; blue, orange, and green curves) is the size of the most generous estimates quoted above it is the leading oscillatory contribution, with anticorrelated wiggles up to one-quarter the size of the matter oscillations. The leading relative density contribution ($b_{\delta_{bc}}$; red) is also quite substantial, and its importance grows as a function of redshift since δ_{bc} is stationary, while δ_m decays as $D(z)$ with redshift. The relative velocity divergence ($b_{\theta_{bc}}$; purple) on the other hand is suppressed by up to two orders of magnitude compared to these other contributions and thus its contribution to any BAO shift should be negligible. It is worth noting that all of the shown relative bias contributions contain significant or even dominant in-, or exactly out-of-phase contributions with oscillations with the opposite sign to the BAO. These oscillations mainly modify the amplitude of the BAO rather than shift it; if these are present at a significant level in DESI galaxies they would lead to modifications of their clustering (e.g. inconsistencies between the BAO and broad-band amplitudes) which could be constrained in full-shape fits (Blazek et al. 2016; Beutler, Seljak & Vlah 2017b; Slepian et al. 2018). These constraints could then be used to reduce the theoretical error on the BAO.

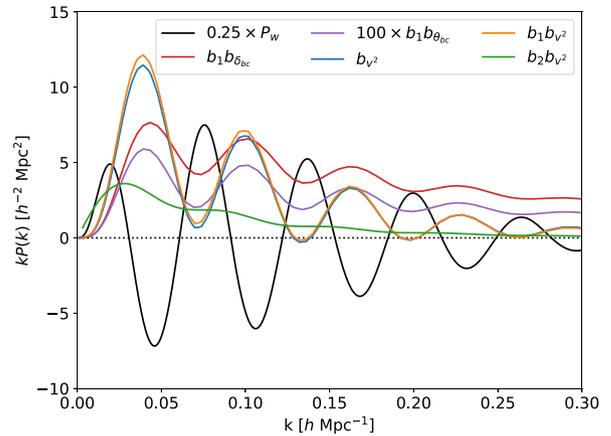


Figure 2. Contributions to the real-space galaxy power spectrum at $z = 0$ from relative baryon-dark matter perturbations, compared to the matter P_w (black). All three leading effects show prominent oscillatory features independent from the P_w , with the largest contribution from the relative density perturbation for values $b_{\delta_{bc}} \sim 1$ (red). The contributions due to the relative velocity divergence ($b_{\theta_{bc}}$) (purple) are quite small and are scaled here by a factor of 100 for visibility. The relative velocity effect terms ($\propto b_{v,2}$) assume an effect of order $0.01\sigma_{bc}^{-2}$, a very pessimistic choice at the upper end of the range of estimates. Contributions due to the cross-correlation of the $b_{v,2}$ term with matter, linear, and quadratic bias are shown in blue, orange, and green, respectively, with the former two showing particularly prominent oscillations at this value of $b_{v,2}$.

Overall, the unknown amplitude of these effects means they must be accounted for, conservatively, in our theoretical systematic error budget. However, this requires further consideration of these effects in redshift space, which will be carried out next.

3.2 Redshift space

We now extend the discussion in the previous section to include redshift-space distortions. In LPT, redshift-space distortions can be modelled by including an extra displacement along the LOS \hat{n} proportional to the peculiar velocity, that is, $\Psi_s = \Psi + \hat{n} (\hat{n} \cdot \Psi) / (aH)$ (Matsubara 2008; Carlson et al. 2013). Within linear theory this is equivalent to taking $\Psi_{s,i} = R_{ij} \Psi_j$, where $R_{ij} = \delta_{ij} + f \hat{n}_i \hat{n}_j$ and $f(z)$ is the linear growth rate.

The effect of redshift-space distortions on the BAO damping is straightforward. Since the exponentiated displacements all get multiplied by R_{ij} , the damping factor acquires an anisotropy given by (Eisenstein et al. 2007a; Matsubara 2008; Ivanov & Sibiryakov 2018)

$$-\frac{1}{2} k_i k_j R_{ia} R_{jb} \Sigma_{\text{NL}}^2 \delta_{ab} = -\frac{1}{2} K^2 \Sigma_{\text{NL}}^2, \quad K_i = R_{ij} k_j, \quad (15)$$

such that $K^2 = (1 + f(2 + f)\mu^2)k^2$. As in real space, small-scale non-linear effects can also affect the damping of the BAO peak. In practice, the most significant contribution will be from the so-called Fingers of God (FoG) – random motions, and therefore redshift-space displacements – due to virial motions of galaxies within haloes (Jackson 1972). To leading order, these contribute an additional $-\frac{1}{2}k^2\mu^2\Sigma_{\text{FoG}}^2$ to the exponential damping. Since virial velocities come from a range of halo masses with different dispersions (Diaferio & Geller 1996; Sheth 1996; White 2001) their distribution is not Gaussian (Peebles 1976; Davis & Peebles 1983). Their overall effect may be better described by other probability distributions (Seljak & McDonald 2011; Chen et al. 2021), such as the, for example,

Lorentzian form sometimes used in the literature (Peacock 1992; Park et al. 1994; Peacock & Dodds 1994; Ballinger, Peacock & Heavens 1996). We will quantify the impact of these choices in Section 7.2.

The effect on the non-linear BAO shift is somewhat more complicated. Following the same reasoning as in real space, we can easily come to the conclusion that the BAO shift should become anisotropic and enhanced along the LOS. This is because redshift-space distortions from peculiar velocities streaming towards overdensities will cause galaxies separated by an overdensity along the LOS to appear closer to each other than those separated perpendicular to such an overdensity. An exact calculation of all the anisotropic non-linearities proportional to the large-scale overdensity σ_s^2 yields (Appendix C)

$$P_{\text{shift}} = -(1/315)\sigma_s^2 k P'_w(k)(b_1 + f\mu^2)(141b_1 + 210b_2 + 112b_s + f(93 + 330b_1 + 210b_2 + 112b_s - 42(1 + b_1)f)\mu^2 + 3(80 + 77b_1 - 28f)f^2\mu^4 + 189f^3\mu^6), \quad (16)$$

where we have also included the possible effect of a tidal bias b_s . Evidently, the shift can no longer be mapped into a single $\Delta\alpha_{\text{iso}}$ as it could in real space, though it is worth noting that the total shift for modes perpendicular to the LOS ($\mu = 0$) remains the same, while the shift along the LOS becomes significantly enhanced.

In order to translate equation (16) into systematic shifts in α_{\parallel} , α_{\perp} we conduct a simple Fisher analysis assuming a cosmic-variance limited survey and Gaussian covariances with a fiducial model given by the Kaiser form plus a polynomial broad-band model as described in Section 5. Defining the systematic error vector $\epsilon = (P_{\text{shift}}^0, P_{\text{shift}}^2)$ to be the multipoles of equation (16) we have that the shift in α 's is given by

$$\Delta\alpha_{\parallel,\perp}^{\text{NL}} = F_{\alpha_{\parallel,\perp},b}^{-1} t_i^b C_{ij}^{-1} \epsilon_j, \quad (17)$$

where C_{ij}^{-1} is the inverse covariance, t_i^b are the linear templates for the free parameters p_b , and F_{ab} is the Fisher matrix. For simplicity we assume the co-evolution relations $b_2 = b_2^L + \frac{8}{21}b_1^L$ and $b_s = -\frac{2}{3}b_1^L$, where the Lagrangian b_2^L is fixed to be a function of b_1^L according to the Sheth–Tormen mass function (Sheth, Mo & Tormen 2001).

Fig. 3 shows the expected shifts in the BAO scale in the above setup as a function of b_1 and redshift. The solid lines show the expected BAO shift in the absence of any non-linear biases: they roughly reproduce the real-space prediction (black dotted) in the case of α_{\perp} but are up to 3–4 times larger for α_{\parallel} . This makes sense since the perpendicular part of equation (16) is identical to the real-space prediction.

The dashed lines in Fig. 3 show the result of adding in nonzero quadratic density and shear biases using peak-background split predictions. These show significant deviations away from the linear bias-only predictions; roughly, for $b_1^L < 1$ the quadratic bias becomes negative, reducing the effect, and vice versa for more biased tracers, which cluster more strongly around overdensities. The bias relations used here are known to qualitatively reproduce the behaviour of dark matter haloes in N -body simulations (Abidi & Baldauf 2018; Lazeyras & Schmidt 2018); however, they are not guaranteed to hold for the galaxy samples observed in DESI, and as such the scatter between the different lines in the plot should be thought of as an indicator of the possible systematic scatter due to the unknown size of non-linear biasing. Conversely, a full-shape fit of galaxy clustering aimed not just at extracting the BAO signal and where all PT parameters are varied automatically incorporates this uncertainty into its parameter constraints. At least in the case of pre-reconstruction, where PT models are quite mature, an alternative procedure to fitting the BAO may be to employ such a model, as one would do in a ‘full’ fit of the galaxy power spectrum, while

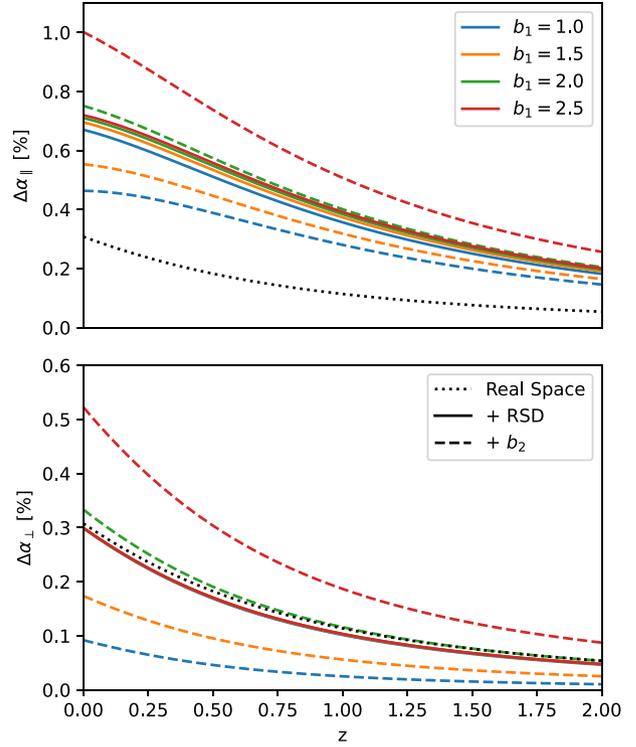


Figure 3. Fisher forecasts for the BAO shift in redshift space due to non-linear clustering as a function of redshift z and linear bias b_1 , assuming a signal-dominated survey. Solid lines show results for linear bias only, while dashed lines show the shift when a quadratic bias parameter, predicted from the peak-background split, is included, showing that the scatter in potential shifts due to the unknown non-linear bias parameters is substantial. The dotted black line shows the shift in real space, which is similar in size to $\Delta\alpha_{\perp}$ in most cases but substantially smaller than $\Delta\alpha_{\parallel}$.

simultaneously marginalizing over broad-band terms (Section 2) as is typically done in BAO fits. Beyond standard BAO measurements, we note that these shifts have a larger effect on $\alpha_{\parallel,\perp}$ than the expected shift due to light species such as neutrinos, which has emerged as a topic of recent interest (e.g. Baumann et al. 2019), and therefore must be carefully marginalized over when looking for ‘exotic’ effects in BAO physics; on the other hand, as we will see, they are very well mitigated by reconstruction.

Converting the error estimates above into α_{iso} and α_{ap} , the former unsurprisingly is subject to non-linear shifts spanning the geometric average $\Delta\alpha_{\parallel}/3 + 2\Delta\alpha_{\perp}/3$, such that the range of scatter due to the unknown size of non-linear biases is directly translated over from Fig. 3. In the case that the non-linear biases are given by the peak background split and Lagrangian coevolution we have, roughly, $\Delta\alpha_{\text{iso}} \approx 0.5\sigma_s^2(1 + (b_1 - 1))D(z)^2$, or roughly 0.3 per cent for a bias one tracer at redshift zero, while neglecting any non-linear Eulerian bias gives $\approx 0.9\sigma_s^2D(z)^2$. The scatter between these two again is indicative of the uncertainty induced in pre-reconstruction measurements by non-linearities. On the other hand, it is interesting to note that this scatter essentially cancels in α_{ap} such that we have $\Delta\alpha_{\text{ap}} \approx 1.2\sigma_s^2f(z)D(z)^2$, corresponding to approximately a 0.4 per cent shift at $z = 0$, to within 0.1 per cent accuracy over the entire range of biases and redshifts shown, with only minor variations increasing with non-linear bias. This suggests that a direct correction based on this estimate can be performed for α_{ap} when measuring it from pre-reconstruction BAO. We caution that the fitting forms above are approximate and intended for error estimation.

We can also apply the same Fisher formalism to extract the BAO shifts due to the quadratic coupling of P_w with itself and the relative velocity effect. Redshift-space distortions similarly enhance the size of the former along the LOS; however, since this effect ultimately has a rather distinguishable scale dependence compared to the BAO wiggles and their derivatives the contribution to the measured BAO scale is extremely suppressed – for a bias $b = 2$ tracer at $z = 0.8$ the effect is of order $f_b^2 \times 0.1$ per cent, which is entirely negligible. This result justifies neglecting any order f_b^2 effects in BAO fits even though they are nominally only suppressed by one power of the baryon fraction $f_b \approx 15$ per cent.

Finally, we compute the redshift-space power spectrum contributions of the relative velocity effect using the Lagrangian formalism of Chen et al. (2019a), converting their real-space expressions into redshift space as in Vlah & White (2019) and Chen et al. (2021, see also Schmidt 2016). By operating within LPT we avoid having to perform wiggle/no-wiggle splits on the relative velocity spectra to model their non-linear damping; the standard algorithms to perform this split are optimized for the matter power spectrum and cannot typically be used ‘out of the box’ in this case where the wiggles are also more pronounced. Keeping corrections only to first order in the relative bias parameters we can express the resulting BAO shifts from equation (17) as $\Delta\alpha_{\parallel,\perp} = b_i(d\alpha_{\parallel,\perp}/db_i)$ for $b_i = \{b_{\delta_{bc}}, b_{\theta_{bc}}, b_{v^2}\}$. In addition to the shift in best-fitting parameters, the presence of these relative bias terms can also change the strength of BAO constraints by modifying the BAO signal amplitude with their in-phase components. However, this effect will be avoided so long as the covariance matrices and model can match the amplitude of the observed BAO.

Fig. 4 shows the resulting redshift-space BAO shift per unit relative bias. The top panel shows the effect of the relative overdensity bias δ_{bc} . The magnitude of the shift diminishes towards lower redshift and higher biases and is opposite in sign for $\alpha_{\text{iso,ap}}$. This is as expected, since the relative component stays constant with redshift unlike the growing adiabatic component, so perpendicular to the LOS the relative size of the effect scales as $b_{\delta_{bc}}/(b_1 D(z))$, while along the LOS, the shift is somewhat reduced by the presence of redshift-space distortions sourced by matter, such that the shift in α_{ap} is negative. We note that while Barreira et al. (2020) find that $b_{\delta_{bc}}$ is typically of order one for both galaxies and haloes and the expected shifts are therefore of the same order of magnitude as the derivatives shown in the top panel, they also find that $b_{\delta_{bc}}$ tends to increase with halo mass. Given that the exact value of this bias is unknown as a conservative estimate we should therefore expect a $\sim 0.05 D^{-1}(z)$ per cent shift in the BAO from this effect.

The bottom panel similarly shows the expected shift per $0.01b_{v^2}$ as a function of bias and redshift. Since these terms exist at one-loop order it is necessary to make some assumptions about non-linear biasing; we have therefore adopted the peak-background split values of b_2 like in Fig. 3, though our results in this case are not very sensitive to this choice as the relative velocity effect is mostly sourced by long-wavelength modes. Here again the signal decreases as a function of linear bias, but unlike the relative density bias, it stays roughly constant with redshift since the one-loop terms contributing to it scale similarly to the linear adiabatic modes. Given the lack of a clear measurement in either data or simulations of the value of b_{v^2} , or a clear theoretical consensus in the literature about its expected size, the derivatives plotted can be thought of as a worst-case scenario for this bias. If we take a log-uniform prior on b_{v^2} spanning the range 10^{-5} to $10^{-2}\sigma_{bc}^{-2}$ the mean shift is roughly $0.02b_1^{-1}$ per cent on α_{iso} with negligible redshift dependence and a significantly smaller anisotropic shift. However, since the straightforward estimate based on halo mass yields of a few times 10^{-5} , absent discovering an

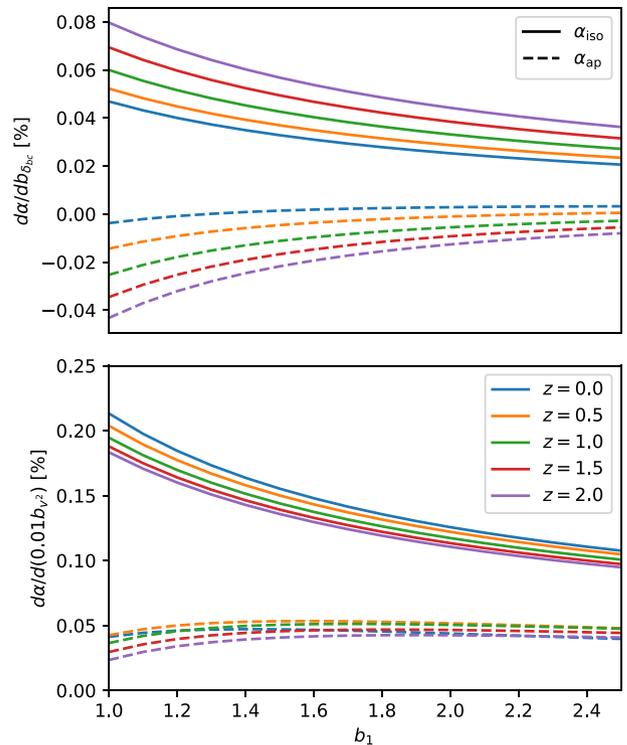


Figure 4. BAO shifts from relative velocity effect as a function of linear bias b_1 and redshift z , computed for (top) the relative density bias $b_{\delta_{bc}} = 1$ and (bottom) $b_{v^2} = 0.01\sigma_{bc}^{-2}$. The normalization for b_{v^2} here is on the upper range of estimates in the literature, with estimates in the literature spanning three orders of magnitude smaller – we adopt $10^{-3}\sigma_{bc}^{-2}$ as the fiducial mean in this work. We have omitted the shift due to the relative velocity divergence since its effect is expected to be much smaller than the two shown.

astrophysical mechanism that preserves the relative velocity signal in early star formation in DESI target galaxies it seems safe to assume this is also a conservative estimate. Further, while the shift due to a b_{v^2} in the upper range of reported values in the literature would be comparable to the statistical precision of the full Y5 DESI data, we expect that full-shape analyses in the meantime will be able to place stronger bounds on its value, since data from the BBaryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013) survey already come quite close to constraining these values (Yoo & Seljak 2013; Beutler et al. 2017b).

4 RECONSTRUCTION

4.1 The standard reconstruction algorithm

In Section 3, we considered the effects of unadulterated non-linearities in structure formation on the BAO signal in galaxy clustering. However, as described in the Introduction, the bread-and-butter of BAO measurements in spectroscopic surveys like DESI concerns the BAO signal in galaxy clustering post-reconstruction. The reconstruction algorithm presented in Eisenstein et al. (2007b) amounts to a non-linear transformation of the observed galaxy density field to enhance and unbiased the BAO signal – the purpose of this subsection is to review these effects in some detail using PT (see Noh, White & Padmanabhan 2009; Padmanabhan, White & Cohn 2009; Schmittfull et al. 2015; White 2015; Hikage, Koyama & Heavens 2017; Ding et al. 2018; Chen et al. 2019b; Hikage, Koyama & Takahashi 2020, for earlier discussions).

The first step of the reconstruction algorithm is to generate the reconstructed displacement $\Psi_r(\mathbf{s})$ at a given redshift-space position \mathbf{s} . Schematically, this is done by dividing the galaxy density field by $b + f\mu^2$ after smoothing the observed galaxy density field by a Gaussian filter $\mathcal{S}(k) = e^{-k^2 R^2/2}$ with R sufficiently large that the smoothed field is well described by linear theory. In this case, the linear Zeldovich displacement can be reconstructed as

$$\Psi_r(\mathbf{k}) = \frac{i\mathbf{k}}{k^2} \left(\frac{\mathcal{S}(k)\delta_g(\mathbf{k})}{b + f\mu^2} \right) \approx \Psi_{\text{Zel}}(\mathbf{k}). \quad (18)$$

In practice, survey geometry and other effects mean that the simple procedure above cannot be followed literally (i.e. the un-windowed $\delta_g(\mathbf{k})$ cannot be directly observed and divided by the Kaiser factor). In the below, however, we will mostly use the approximation in equation (18) except when describing potential biases due to imperfect reconstruction, leaving the details of the reconstruction algorithm itself to our companion ‘optimal reconstruction’ papers (Chen in preparation; Paillas et al. 2024).

Given Ψ_r , the reconstruction algorithm proceeds to shift both the observed galaxies and randoms by the negative reconstructed displacement in order to cancel the non-linear damping due to long-wavelength displacements, and the reconstructed density field is defined to be the difference between the displaced galaxies and shifted randoms, $\delta_r = \delta_d - \delta_s$. Subtracting off the overdensity of the shifted randoms restores the linear clustering power removed by displacing the galaxies. Using number conservation, the overdensities of the galaxies and randoms post-reconstruction can be expressed as (Schmittfull et al. 2015)

$$1 + \delta_{d,s}(\mathbf{x}_r) = \int d^3\mathbf{x} \delta_D(\mathbf{x}_r - \mathbf{x} - \Psi_r(\mathbf{x})) (1 + \delta_{d,s}(\mathbf{x})). \quad (19)$$

Here, \mathbf{x}_r is the position of galaxies and randoms post-reconstruction, and the densities $\delta_{d,s}(\mathbf{x})$ refer to their pre-reconstruction densities, that is, $\delta_d(\mathbf{x}) = \delta_g(\mathbf{x})$ and $\delta_s(\mathbf{x}) = 0$. Equivalently we have in Fourier space

$$\delta_{d,s}(\mathbf{k}) = \int d^3\mathbf{x} e^{-i\mathbf{k}\cdot\mathbf{x}} (e^{-i\mathbf{k}\cdot\Psi_r(\mathbf{x})} (1 + \delta_{d,s}(\mathbf{x})) - 1). \quad (20)$$

In order to make a connection with LPT we can further substitute $\mathbf{x} = \mathbf{q} + \Psi(\mathbf{q})$, such that the integral over \mathbf{x} is replaced with \mathbf{q} and $\mathbf{x}_r = \mathbf{q} + \Psi(\mathbf{q}) + \Psi_r(\mathbf{q} + \Psi)$ (White 2015; Chen et al. 2019a).

In order to include redshift-space distortions the reconstructed displacement is multiplied by the matrix R_{ij} (Section 3.2). Where this boost factor is applied varies across the literature; in this paper, and generally in DESI Y1, we will focus on two conventions:

(i) *RecIso*: galaxies shifted by $-\mathbf{R}\Psi_r$, randoms shifted by $-\Psi_r$

$$\delta_d = (b_1 + f\mu^2 - (1 + f\mu^2)S)\delta, \quad \delta_s = -S\delta,$$

$$\delta_r = (b_1 + (1 - S)f\mu^2)\delta,$$

(ii) *RecSym*: galaxies shifted by $-\mathbf{R}\Psi_r$, randoms shifted by $-\mathbf{R}\Psi_r$

$$\delta_d = (b_1 + f\mu^2 - (1 + f\mu^2)S)\delta, \quad \delta_s = -(1 + f\mu^2)S\delta,$$

$$\delta_r = (b_1 + f\mu^2)\delta,$$

where we have also written down the linear theory predictions for each scheme. The former convention, RecIso, was chosen by the SDSS (Sloan Digital Sky Survey), BOSS, and eBOSS collaborations (Padmanabhan et al. 2012; Alam et al. 2017, 2021); since linear redshift-space distortions are sourced by the same linear velocities as in $\mathbf{R}\Psi_r$, removing them from the displaced galaxies without doing the same to the randoms leads to the removal of linear redshift-space distortions on large scales and ‘isotropizes’ the clustering.

The second convention, RecSym, was introduced in White (2015) and treats randoms and galaxies ‘symmetrically’ such that the linear theory broad band is preserved. It is important to note that (i) the anisotropic clustering still contains BAO information, so subtracting it may result in weaker BAO constraints for a fixed shot noise and (ii) the cancellation of RSD depends on the extent to which the true and fiducial cosmologies match, such that even given the true Zeldovich displacement the anisotropic term would be more like $(f - S f_{\text{fid}})\mu^2$, leaving a residual that could be confused for anisotropic BAO from the Alcock–Paczynski effect. RecSym is on the other hand robust against any such mistakes to leading order since both particles and randoms are moved identically.

4.2 Non-linear modelling of the reconstructed field

The reconstructed field δ_r carries a somewhat different set of non-linearities than the original galaxy density δ_g , as expected since the reconstruction algorithm was designed specifically to reduce the non-linear damping of the BAO signal. The non-linear modelling of the reconstructed BAO is rather non-trivial and, while much progress has been made in recent years including efforts to model the post-reconstruction matter field (Hikage et al. 2017, 2020) and galaxies in the Zeldovich approximation (White 2015; Chen et al. 2019b), to the best of the authors’ knowledge no complete model taking into account both non-linear dynamics and bias as well as dependences on the reconstruction algorithm itself (e.g. due to the fiducial cosmology assumed) has yet been presented. None the less, it is possible to understand the effect of non-linearities post-reconstruction using the same methods as the pre-reconstruction case. In the below, we will outline general arguments. More detailed calculations, particularly in the more tractable case of matter, will be presented in Appendix D.

4.2.1 Real space

Let us first examine the damping of the BAO peak post-reconstruction. For the displaced galaxies, we can write the following version of equation (20):

$$\delta_d(\mathbf{k}) = \int d^3\mathbf{q} e^{i\mathbf{k}\cdot(\mathbf{q} + \Psi(\mathbf{q}) + \Psi_r(\mathbf{q} + \Psi))} (1 + \delta_g(\mathbf{q})), \quad (21)$$

where we have dropped delta functions in \mathbf{k} for brevity. Compared to the pre-reconstruction case, we can see that the featureless displacement that smooths the power spectrum is now given by $\Psi(\mathbf{q}) + \Psi_r(\mathbf{q}) \approx (1 - S) * \Psi$ to leading order, that is, the damping parameter is now

$$\Sigma_{\text{rec}}^2 = \frac{2}{3} \int \frac{dk}{2\pi^2} (1 - S)^2 P(k)(1 - j_0(kr_d)). \quad (22)$$

The case of the shifted field is similar: at first sight, it looks like the shifted field starts out undisplaced so that the only BAO damping comes from the shift we apply to the randoms; however, since the shift field is computed from the Eulerian density field it already comes pre-displaced, such that its ultimate damping scale is the same as in the displaced field (Appendix D2). Another way to see this is that the full reconstructed density field is given via equation (20) as (see also Shirasaki et al. 2021; Sugiyama et al. 2021)

$$\delta_{\text{rec}}(\mathbf{k}) = \int d^3\mathbf{x} e^{-i\mathbf{k}\cdot\mathbf{x} - i\mathbf{k}\cdot\Psi_r(\mathbf{x})} \delta_g(\mathbf{x}) \quad (23)$$

and, since $\delta_g(\mathbf{x})$ can be expressed as a sum of bias operators at the Lagrangian position \mathbf{q} (Mirbabayi, Schmidt & Zaldarriaga 2015) it follows that the total shift is given in equation (22). This result that the BAO post-reconstruction is still damped by a simple single-exponential form, which we derive for the first time, is in contrast

to previous derivations in the literature using PT and techniques from effective field theory (EFT) (Ding et al. 2018; Chen et al. 2019b).

We thus have that in real space the damping form is equivalent to the pre-reconstruction case except with a modified damping parameter, that is,

$$P(k) = b^2 \left(P_{nw}(k) + P_w(k) e^{-\frac{1}{2} k^2 \Sigma_{\text{rec}}^2} \right). \quad (24)$$

In addition, since the shift field is reconstructed from noisy galaxy densities it will add an additional uncorrelated damping proportional to the shot noise P_{shot}/\bar{n} (White 2010)

$$\Sigma_{\text{shot}}^2 = \frac{2}{3} \frac{P_{\text{shot}}}{b^2 \bar{n}} \int \frac{dk}{2\pi^2} \mathcal{S}(k)^2 (1 - j_0(kr_d))^2 \approx \frac{1}{6\pi^{3/2}} \frac{P_{\text{shot}}}{b^2 \bar{n} R}, \quad (25)$$

where we have assumed that $R \ll r_d$. Here, P_{shot} describes the deviation of the galaxy stochasticity away from a Poisson distribution, where we would have $P_{\text{shot}} = 1$. For galaxies with inverse comoving density $3 \times 10^3 h^{-3} \text{Mpc}^3$ this is equal to $\Sigma_{\text{shot}}^2 = 6b^{-2} R_{15}^{-1} P_{\text{shot}} h^{-2} \text{Mpc}^2$ where R_{15} is the reconstruction smoothing in units of $15 h^{-1} \text{Mpc}$. Again, it is important to note that this will not in general be the only additional contribution to the post-reconstruction damping, but its presence is a good gauge of the size of non-linear corrections to the Zeldovich damping. Some illustrative values of the expected damping as a function and smoothing scale are shown in Fig. 5 – while in general the damping can be somewhat reduced by decreasing R , other non-linearities eventually make this choice suboptimal. In the case of bias values and redshifts representative of the DESI quasar (QSO) sample, we can see that the extra damping incurred by the shot noise brings the total post-reconstruction damping above the pre-reconstruction one for essentially all smoothing scales. However, for a suitably large R the smoothing scales are rather comparable, while the non-linear BAO shift is still cancelled post-reconstruction such that performing reconstruction may still be preferred.

The effect of reconstruction on the non-linear BAO shift follows a similar logic. Reconstruction explicitly removes, up to the smoothing scale R , the Zeldovich displacement terms that produce the shift; since the shift depends on modes $k \lesssim 1/r_d \ll 1/R$ the cancellation is essentially perfect if the linear displacement is properly estimated (see Appendix D1 for an explicit calculation). However, mis-estimating the linear bias b_1 can lead to residual displacements post-reconstruction, leading to errors roughly of order $(\Delta b_1/b_1) \Delta \alpha_{\text{iso, NL}}$, where Δb_1 is the difference between the true linear bias and that assumed during the application of the reconstruction algorithm. It is important to note that the relevant quantity here is the error on the linear bias and not its product with σ_8 . The latter is typically extremely well constrained by the data while, absent external priors, the constraints on the former are essentially dictated by the errors on σ_8 itself. Conceptually, this is because we can only cancel the shift as well as we can know the amplitude of matter fluctuations within spheres of the BAO radius r_d .

Reconstruction also has a mixed effect on the relative biases. The contributions due to $b_{\delta_{bc}, \theta_{bc}}$ are essentially linear BAO's that are out of phase with the matter one, and reconstruction thus simply sharpens them in the same way, taking the damping from $\Sigma_{\text{NL}}^2 \rightarrow \Sigma_{\text{rec}}^2$. Since both signals are identically strengthened, reconstruction is not expected to reduce the theoretical error due to this effect. On the other hand, the $b_{v,2}$ contributions are due to its coupling with quadratic non-linearities in the galaxy density, which is in fact dominated by the large-scale displacements cancelled by reconstruction as v_{bc}^2 is dominated by long-wavelength modes (Blazek et al. 2016). While we thus expect reconstruction to reduce the contributions of the

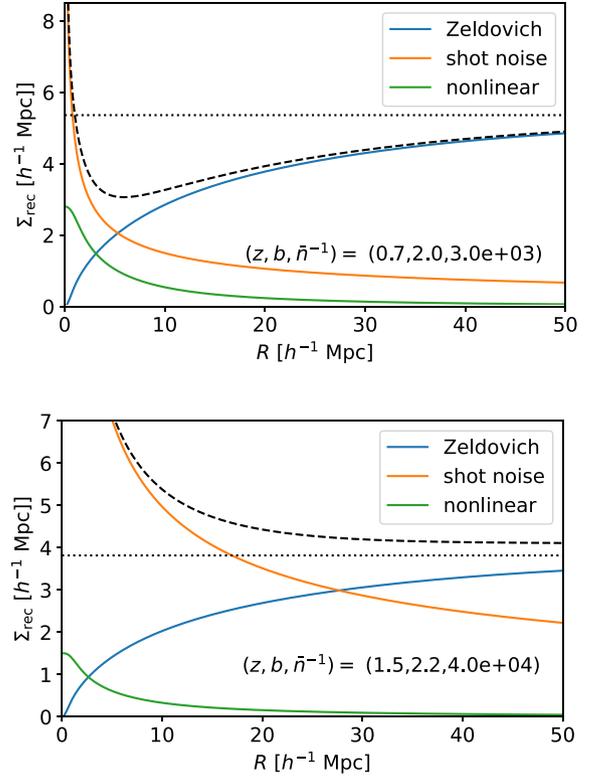


Figure 5. Contributions to the post-reconstruction BAO damping due to Zeldovich displacements (blue), shot noise (orange), and non-linearities (green) as a function of smoothing scale R , with the total shown in the black dashed line. The black dotted line shows the full Zeldovich displacement at the BAO scale, that is, the pre-reconstruction damping scale. The green curve is estimated by computing the mean-square reconstructed displacement (Ψ_r^2) using the one-loop power spectrum for the value of bias shown, with non-linear Eulerian biases assumed to be zero. The top and bottom panels show representative dampings for the DESI LRG and QSO samples, respectively, with the shot noise value reported in [$h^{-3} \text{Mpc}^3$].

relative velocity effect to the BAO – and indeed calculations within the Zeldovich approximation show that its oscillatory components are reduced by roughly a factor of two post-reconstruction – a complete calculation of the post-reconstruction relative velocity effect is outside the scope of this work. A conservative estimate is therefore to use the pre-reconstruction shifts as shown in Fig. 4 in our error budget.

4.2.2 Redshift space: fitting form and algorithm choice

In the case of RecSys, the arguments in real space carry over to redshift space essentially identically; since both the reconstructed displacement and the Zeldovich displacements responsible for the non-linear BAO shift and BAO damping are multiplied by the same matrix \mathbf{R} , reconstruction simply cancels the shift equally in redshift space and we get $k^2 \Sigma_{\text{rec}}^2 \rightarrow K^2(\mu) \Sigma_{\text{rec}}^2$ as in the pre-reconstruction case. This motivates the fitting form

$$\mathcal{C}(k, \mu) = (b + f\mu^2)^2 \exp \left[-\frac{1}{2} k^2 (\mu^2 \Sigma_{\parallel}^2 + (1 - \mu^2) \Sigma_{\perp}^2) \right] \quad (26)$$

both pre- and post-reconstruction, with the understanding that the phase shift on top of this will be larger pre-reconstruction, and that $\Sigma_{\parallel, \perp}^2$ are roughly given by the Zeldovich prediction with uncertainty due to non-linearities and imperfect knowledge of the

true cosmology. In particular, as with the linear bias, the extent to which these cancellations happen depends on the fidelity with which we can constrain the linear growth rate f , such that we should expect a remaining theoretical error of size $(\Delta f/f)\Delta\alpha_{\perp,\parallel}$. Both of these errors are related to our certainty on the underlying matter clustering amplitude σ_8 , which we can conservatively place a 5 per cent uncertainty on given current discrepancies between the primary CMB and late-universe probes, though we note that differences up to 10 per cent with the primary CMB exist, for example, from galaxy redshift-space clustering (Abdalla et al. 2022). We remind the reader that, while Zeldovich displacements form the dominant contribution to BAO damping, small-scale non-linearities can also contribute and will tend to be equivalent to equation (26) only at leading order. We investigate the effect of including such effects (e.g. FoG) into BAO fits in Section 7.2, finding that they have negligible impact even at DESI precision. More generally, these effects beyond leading order can be mitigated by limiting the range of scales fit, especially in the power spectrum.

In the case of RecIso however some care must be taken, because only the real-space portion of the Zeldovich displacement is cancelled for the shifted randoms. In this case, it is reasonable to assume that much (though not all) of the contributions due to the anisotropic BAO shift will remain, though the $\mu = 0$ contribution should still be cancelled. We show this explicitly for matter in Appendix D3. Similarly, while the displaced galaxies will still have their non-linear damping reduced to $K^2\Sigma_{\text{rec}}^2$ the shifted fields and their cross-spectra with galaxies will now be damped according to

$$\begin{aligned}\Sigma_{ss}^2(\mu) &= \Sigma_{\text{rec}}^2 + f^2\mu^2\Sigma_{\text{NL}}^2 + 2f\mu^2\Sigma_x^2, \\ \Sigma_{ds}^2(\mu) &= (1 + f\mu^2)\Sigma_{\text{rec}}^2 + f(1 + f)\mu^2\Sigma_x^2 + f^2\mu^2\Sigma_S^2,\end{aligned}\quad (27)$$

where we have defined $\Sigma_x^2 = \frac{2}{3} \int \frac{dk}{2\pi^2} (1 - \mathcal{S}) P_{\text{lin}}(k) (1 - j_0(r_d k))$ as the cross-correlation between the pre- and post-reconstruction Zeldovich displacements and $\Sigma_S^2 = \frac{1}{3} \int \frac{dk}{2\pi^2} \mathcal{S}^2 P_{\text{lin}}(k)$ is the mean-square Zeldovich displacement after smoothing by \mathcal{S} . The distinct pairwise displacements above between the shifted galaxies and randoms in RecIso mean that the reconstructed power spectrum in this case does not have BAO which are damped by a single exponential, but rather piecewise according to equation (27), which can be also seen from the fact that equation (23) does not hold if the galaxies and randoms are not moved by the same amount. The above expressions reduce to the pre-reconstruction displacement when $\mathcal{S} \rightarrow 0$, and the post-reconstruction displacement in real space if $f \rightarrow 0$, but in general the BAO signal in RecIso will be more damped than it is in RecSym, and the non-linear shift larger, because of the incomplete cancellation of displacements along the LOS. We thus see that RecIso suffers from (i) a partially uncanceled non-linear BAO shift, (ii) more damping, (iii) less signal, and (iv) greater sensitivity to systematics in reconstruction than RecSym.

5 BROAD-BAND MODELS IN FOURIER AND CONFIGURATION SPACE

Thus far, we have developed the theory for the effects of non-linear structure formation on the oscillatory BAO signal. The main contributions come from large-scale displacements which smoothly alter the envelope of the damped BAO wiggles and shift the observed BAO. These can be effectively removed through reconstruction, leading to a robust BAO signal that can be only weakly biased through second-order bias effects and (for pre-reconstruction or under RecSym convention) damped by a single exponential damping term. However, non-linearities – which are modified, but not

removed, by reconstruction – not to mention systematics in the data and differences in the fiducial and true cosmologies, also contribute to the broad-band shape of the observed galaxy two-point function. Since the express purpose of BAO fits is to isolate the geometric information in the oscillatory part of the power spectrum only, the ability to marginalize over any broad-band differences is an integral part of any BAO modelling.

An essential feature of such a broad-band model is that it must fit any feature in the power spectrum with frequency in Fourier space lower than that of the BAO, thereby isolating the BAO information from the rest of the clustering signal. This is equivalent in configuration space to saying that it must remove any possible features localized at radii smaller than the BAO scale, as we will discuss in detail in Section 5.2. It is important to emphasize that this is done regardless of the source (fiducial cosmology, observational systematic or non-linearity) of the broad-band contribution and, conversely, if a systematic contributes to clustering in a way that is not sufficiently smooth it would need to be specifically modelled in order to mitigate its effect on the BAO.

The broad-band terms in our fiducial BAO model (equation 5) are contained in $\mathcal{B}(k, \mu)P_{nw} + \mathcal{D}(k, \mu)$. In principle \mathcal{D} can be any sufficiently flexible fitting form non-degenerate with the BAO wiggles, though in practice it is useful to isolate a ‘baseline’ contribution expected from linear theory alone described by $\mathcal{B}(k, \mu)$. A further role of $\mathcal{D}(k, \mu)$ is to account for any residual broad-band term in $P_w(k)$ – ideally this would represent only the BAO wiggles, but in practice this separation cannot be performed on the linear power spectrum exactly (leading to, i.e. differences in the BAO templates using different methods, see Section 7.4 and Fig. 17). In this case, a good choice of $\mathcal{D}(k, \mu)$ ensures consistent results regardless of the method used to generate $P_w(k)$. A particularly simple, well-motivated choice for $\mathcal{B}(k, \mu)$ is (Kaiser 1987; Peacock & Dodds 1994; Cole, Fisher & Weinberg 1995)

$$\mathcal{B}(k, \mu) = (b + f\mu^2)^2 \left(1 + \frac{1}{2}k^2\mu^2\Sigma_S^2 \right)^{-2}, \quad (28)$$

where in addition to the linear theory prediction we have included a rough ‘FoG’ parametrization through a single free velocity dispersion parameter Σ_S in order to capture the rapid damping of the power spectrum along the LOS due to small-scale virial motions.

Past BAO analyses (e.g. Anderson et al. 2014; Beutler et al. 2017a; Bautista et al. 2021) have tended to rely on parametrizations of $\mathcal{D}(k, \mu)$ with angular multipoles expanded in polynomials of the wavenumber k in powers between some chosen $n_{\text{min,max}}$. The logic is that polynomials of sufficiently low order cannot mimic the oscillatory BAO signal. In configuration space, past analyses have also employed polynomials in r of sufficiently low degree with the expectation that the BAO peak in the correlation function cannot be reproduced by polynomials (Ross et al. 2017; Bautista et al. 2021). While historically demonstrated to be robust, this approach suffers from a few limitations: (1) there is no clear criteria for what orders of polynomial terms to include – which is a function of both the scale cuts and noise of the analysis (the polynomial cannot have more critical points than the number of BAO wiggles detected) – and even if there were, there is no clear relation between the order of polynomial and smoothness of the function; and (2) while low-order polynomials cannot mimic wiggles in Fourier space, it is easy for them to reproduce a peak in configuration space. Due to these limitations, the stability of any particular prescription for the broad band given a particular set of data was only empirical. In fact, the earliest broad-band models used to fit BAO in the SDSS were based on roughly corresponding powers of r^n and k^{-3-n} to fit the correlation

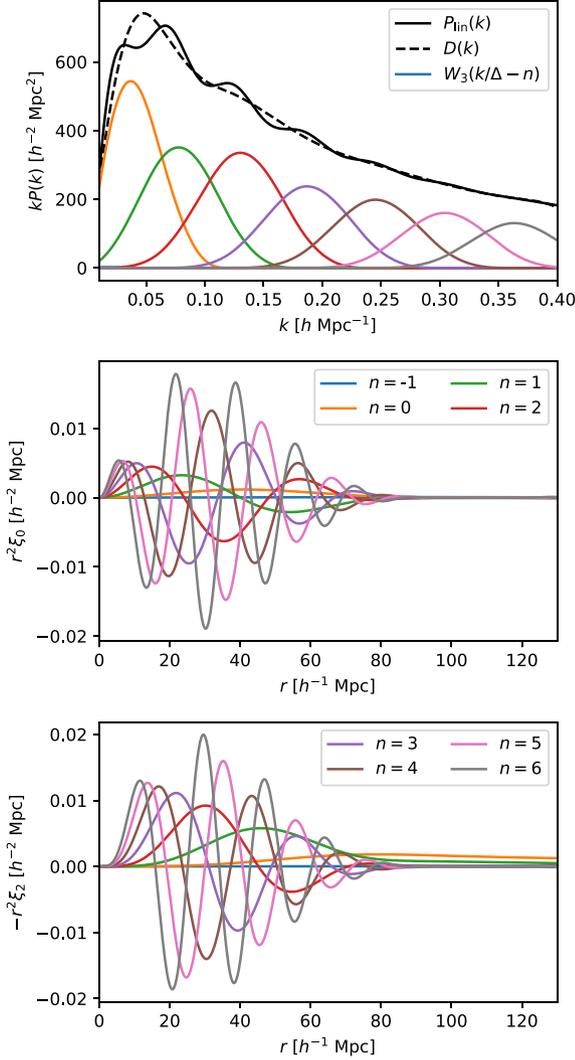


Figure 6. Smooth spline basis for broad-band fitting. Top: the basis in Fourier space, multiplied by a_n coefficients fit to the linear power spectrum (black solid). The resulting broad-band fit, which broadly matches the power spectrum, but does not contain BAO wiggles, is shown in the black-dashed line. Middle: the same spline functions without the coefficients a_n Hankel transformed to give the correlation function monopole. The first two basis functions have smooth large-scale contributions, while the rest are oscillatory and confined to $r < 2\pi/\Delta$. The small wiggles at the largest r are due to ringing from the discontinuous higher derivatives of the spline basis. Bottom: same as the middle plot, but for the quadrupole. The first three basis functions in this case have smooth large-scale contributions.

function and power spectrum, respectively (Anderson et al. 2012; Xu et al. 2012) but as the data improved the number of polynomial terms in the power spectrum increased while those in the correlation function stayed roughly constant. In this work, we instead develop and validate a new unified treatment of the broad band $\mathcal{D}(k, \mu)$ for the power spectrum (and correlation function) that avoids these limitations. While the arguments in this section are quite general, we note that their assumptions are implicitly and explicitly tested in many accompanying DESI 2024 BAO papers – we refer readers interested in the effects of fiducial cosmology and observational systematics on the broad-band model developed to Pérez-Fernández et al. (2024) and DESI Collaboration (2024a).

5.1 Fourier space

Our starting point will be to develop a broad-band model in Fourier space $\mathcal{D}_\ell(k)$, where we allow for distinct terms for each clustering multipole ℓ . The requirements on \mathcal{D}_ℓ are that they should not be degenerate with the BAO signal itself; this is really a requirement on the data that any theoretical or observational systematics not have features that can ‘look’ like the BAO. In practice, this requires that \mathcal{D}_ℓ be smoother than the BAO scale, or rather that its values be uncorrelated beyond some scale $\Delta > \pi/r_d$, which is half the wavelength of the BAO by the Nyquist sampling theorem.

A convenient way to do this is to parametrize the broad band using a cubic spline basis with bases separated by a suitably chosen Δ , that is,

$$\mathcal{D}_\ell(k) = \sum_{n=-1}^{\infty} a_{\ell,n} W_3\left(\frac{k}{\Delta} - n\right). \quad (29)$$

Here W_3 is the ‘piecewise cubic spline’ (fourth-order) extension of the counts-in-cell interpolation kernel (Hockney & Eastwood 1988; Jeong 2010)

$$W_3(x) = \frac{1}{6} \begin{cases} 4 - 6|x|^2 + 3|x|^3 & |x| \leq 1, \\ 8 - 12|x| + 6|x|^2 - |x|^3 & 1 < |x| \leq 2, \\ 0 & |x| > 2, \end{cases} \quad (30)$$

which has the convenient property that $\mathcal{D}_\ell(k > 0) = 1$ when $a_{\ell,n} = 1$. The first few kernels in this basis fit to the linear power spectrum, where we have chosen to use $\Delta = 0.06 \text{ h Mpc}^{-1} > \pi/r_d$, are shown in Fig. 6. The model is able to match the broad band of the power spectrum without reproducing the BAO signal itself.

A nice feature of this parametrization of the broad band is that the free coefficients $a_{\ell,n}$ in this model are by construction never degenerate with the BAO signal, and thus cannot lead to numerical instability when fitting the BAO scale. As an extreme example, even if $n\Delta > k_{\text{max}}$ such that $a_{\ell,n}$ is unconstrained, it will not interfere with the BAO signal within the range of fit where BAO are detected, but rather simply sample its prior. As such there is no fine-tuning required for the number of free parameters unlike in the polynomial case.

Finally, we note that in the Gaussian power spectrum likelihood typically used to constrain the BAO, any linear parameter in the power spectrum model, for example, the $a_{\ell,n}$ ’s, can be analytically marginalized instead of being sampled directly. This is equivalent to adding an additional theoretical contribution to the covariance matrix if one does not care about the actual value of the marginalized-over parameters (Taylor & Kitching 2010, see also Section 6.2). Since we allow $a_{\ell,n}$ to take on arbitrary values, our prescription is equivalent to allowing infinite error at each wavenumber $P_\ell(k)$, with the proviso that the error across k points are correlated when they are within Δ of each other, and the error on different multipoles is uncorrelated. A related theoretical-error-based approach to fitting the BAO was proposed in Philcox et al. (2020), who impose an exponentially decaying correlation with Δ of the error but more narrowly estimate its size according to one-loop PT. As we will now see, while the spline broad band induces a slightly counterintuitive theoretical covariance that is technically infinite at each k point, by performing a simple coordinate (Fourier) transform the theoretical error can instead be strongly localized below some scale in configuration space.

5.2 Configuration space

The spline basis introduced in Fourier space is composed of smooth and well-localized functions which can be readily Hankel

transformed into correlation functions. Fig. 6 shows their $j_{0,2}$ transforms: these basis functions turn into a basis of oscillatory functions in configuration space with zeros around $r = \pi/\Delta$, after which they rapidly decay. These functions have arbitrarily high frequencies below this radius as a function of n . This is to be expected since the functions in Fourier space have widths given by Δ and are centred at $n\Delta$. These functions are given by

$$\Delta^3 B_{\ell,n}(r\Delta) = i^2 \int \frac{dk k^2}{2\pi^2} W_3\left(\frac{k}{\Delta} - n\right) j_\ell(kr) \quad (31)$$

and we give their exact form in Appendix E. Fitting these functions with free amplitudes in configuration space, that is,

$$\tilde{D}_\ell^{\text{spline}}(r) = \Delta^3 \sum_{n=-1}^{\infty} a_{\ell,n} B_{\ell,n}(r\Delta), \quad (32)$$

nulls any signal in the correlation function, BAO or otherwise, up to $r_{\min} \approx \pi/\Delta$, since the functions sample all available Fourier frequencies up to that scale. While there are, in principle, an infinite number of broad-band parameters in this expansion, their effects are localized to small r ; we therefore have that the spline broad-band model in Fourier space is equivalent to setting an r_{\min} when performing the fit in configuration space rather than including any smooth broad-band (e.g. polynomial) model. The two exceptions to this rule are the two basis functions $n = [0, 1]$ in the quadrupole, which have some smooth dependence out to large r . Note that the $B_{2,-1}$ term also has nontrivial support at large k but it is degenerate with the k_{\min} contributions discussed below due to being confined to low k only. We test the equivalence of these Fourier and configuration space configurations in Section 7 (see in particular the discussion around Fig. 9). In particular, we note that while there is some ‘spillover’ due to the finite localization of these functions to small r , the equivalence at the level of posteriors is quite insensitive to these numerical choices.

Perhaps surprisingly, however, another (implicit) component of the Fourier-space broad-band model does make it necessary to introduce a smooth polynomial broad-band model in configuration space. In order to protect against unknown large-scale data systematics we typically only include scales above a certain k_{\min} in the power spectrum when fitting the BAO. The Fourier transforms of these uncontrolled large-scale multipoles can be written as an expansion in even powers of the radius r (odd powers are cancelled when computing even multipoles), such that

$$\tilde{D}_\ell(r > r_{\min}) = \tilde{a}_0 + \sum_{n=1}^{\infty} \tilde{a}_n \left(\frac{rk_{\min}}{2\pi}\right)^{2n}. \quad (33)$$

For typical k_{\min} , for example, $k_{\min} = 0.02 h \text{ Mpc}^{-1}$, then $2\pi/k_{\min} \approx 300 h^{-1} \text{ Mpc}$ and the higher order terms in this series get progressively smaller at the BAO scale. We can hence truncate this to first order, $n = 1$.

Overall, both the considerations above lead to a model for the broad band of the monopole and quadrupole of the correlation function given by

$$\tilde{D}_0(r) = \tilde{a}_{0,0} + \tilde{a}_{0,1} \left(\frac{rk_{\min}}{2\pi}\right)^2, \quad (34)$$

$$\begin{aligned} \tilde{D}_2(r) = \tilde{a}_{2,0} + \tilde{a}_{2,1} \left(\frac{rk_{\min}}{2\pi}\right)^2 \\ + \Delta^3 (a_{2,0} B_{2,0}(r\Delta) + a_{2,1} B_{2,1}(r\Delta)) \end{aligned} \quad (35)$$

with $B_{2,0}(r\Delta)$ and $B_{2,1}(r\Delta)$ given by equation (31).

6 NUMERICAL VALIDATION WITH CONTROL VARIATE SIMULATIONS

Having provided a comprehensive theoretical motivation for how the BAO can be modelled and fit given a set of power spectrum or correlation function measurements, we now turn to validating this model using a set of precise simulated measurements with a known cosmological model. In addition to what has been presented so far, there are a number of choices of how exactly one builds the full model that we will also test, many of which are required due to subtle differences between the model we have presented here and what has been done in previous works. The full extent of the things we aim to test/understand are:

- (i) How do the BAO constraints compare between the previous polynomial-based broad-band model and the new spline-based model?
- (ii) What is the impact of different prior choices for the BAO/FoG damping parameters Σ_{\perp} , Σ_{\parallel} , and Σ_s ?
- (iii) What is the impact of including the FoG term in the wiggle component of the model as well as the smooth component?
- (iv) What are the optimum scale cuts to use for fitting the power spectrum and correlation function?
- (v) How does the choice of method used to perform the wiggle/no-wiggle split impact the BAO constraints?
- (vi) Does dilating the full power spectrum model or just the wiggle component make a difference?
- (vii) Does our updated model of the correlation function—consistently treating it as the Hankel transform of the dilated power spectrum model with the same bias and RSD parameters; and accounting for the measurement binning—change the BAO constraints compared to previous methods?
- (viii) Does the choice of Markov chain Monte Carlo (MCMC) sampler used to perform the fit impact the constraints?

Combining all of these effects together, we will arrive at a single consensus value for the modelling systematic error within the DESI BAO analyses, to be combined with the theoretical systematic error from earlier.

6.1 Simulations

For all the numerical tests in this work, aimed towards motivating a modelling systematic error budget for DESI, we use a suite of 25 high-precision measurements of the power spectrum and correlation function. These are obtained from control variate (CV) realizations of a set of cubic simulations populated with a set of galaxies mimicking the clustering of the DESI LRG sample. Here, we give a brief summary of the relevant details of these simulations, and refer the reader to our companion papers, Garcia-Quintero et al. (2024) and Mena-Fernández et al. (2024) for a more complete description.

Our base cubic simulations, the AbacusSummit⁸ suite (Maksimova et al. 2021) are a set of over 140 simulations across 97 different cosmological models, each with a box size $(2 h^{-1} \text{ Gpc})^3$ containing 6192^3 particles; for an effective particle mass of $2 \times 10^9 h^{-1} M_{\odot}$. In this work, we restrict to the ‘cosm000’ (Planck Collaboration VI 2020), Λ CDM-based cosmological model,⁹ and a single simulation snapshot with redshift $z = 0.8$, which corresponds closely to the peak of the DESI LRG sample number density. The simulations are

⁸<https://abacussummit.readthedocs.io/en/latest/>

⁹Specifically, with $\omega_b = 0.02237$, $\omega_c = 0.12$, $h = 0.6736$, $A_s = 2.083 \times 10^{-9}$, $n_s = 0.9649$, and a single massive neutrino with mass $m_\nu = 0.06 \text{ eV}$.

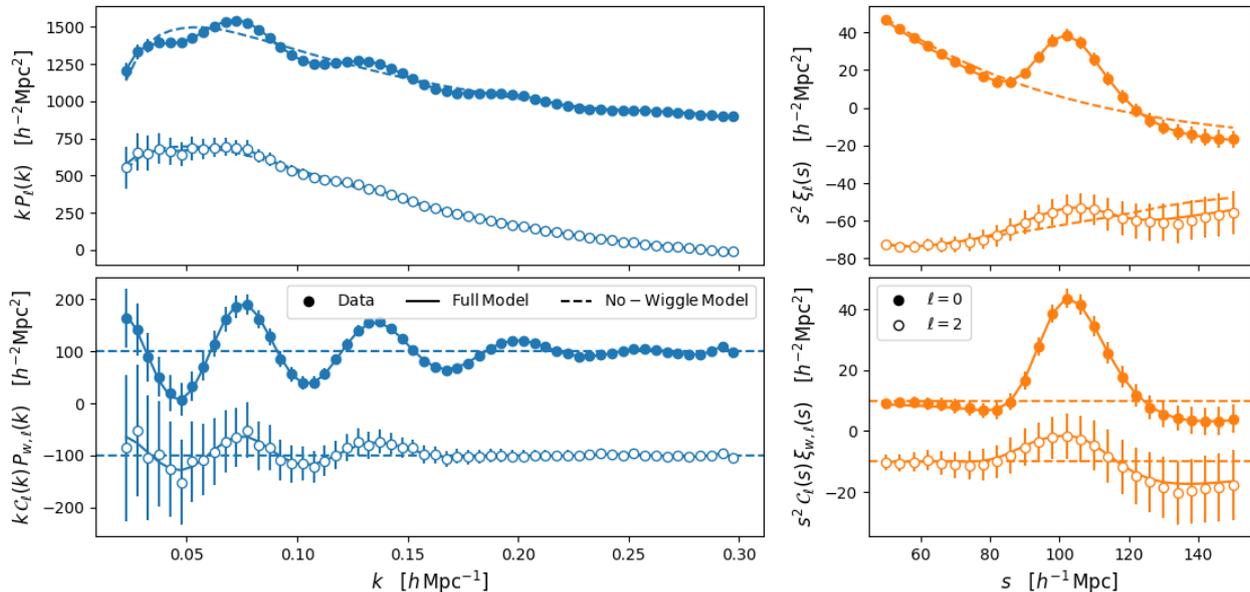


Figure 7. The mean clustering signal, taken from 25 CV simulations of DESI’s LRG sample, used for all the modelling tests in this work. In all panels closed/open points show the mean measured monopole and quadrupole clustering respectively, and error bars are those for a single realization (i.e. not the error on the mean). Left panels are the power spectrum and right are the correlation function. For both, the top panels show the full signal, including the best-fitting full- and no-wiggle models (solid/dashed lines, respectively) using our fiducial fitting methodology (see Section 7). The bottom panels show only the remaining wiggle component of the multipole data/models after subtracting the no-wiggle model (i.e. showing only the second term on the right-hand side of equation (5), multiplied by k , and arbitrarily offset by ± 100 (± 10) for the power spectrum (correlation function) for visualization purposes.

populated with the ‘vanilla’ Halo Occupation Distribution model described in Mena-Fernández et al. (2024), leading to a sample with constant number density $\sim 1 \times 10^{-3} h^3 \text{Mpc}^{-3}$. Finally, we apply a reconstruction algorithm to these data, following RecSym convention, and other algorithmic choices presented in Chen (in preparation) and Paillas et al. (2024).

From these, a set of more precise clustering measurements is obtained using the CV technique (Chartier et al. 2021; Ding et al. 2022; Kokron et al. 2022; DeRose et al. 2023). The CV technique works by subtracting off the intrinsic sample variance in simulated cosmological observables shared with highly correlated but more cheaply computable surrogate observables – in the case of reconstruction, Hadzhiyska et al. (2023) showed that a reasonable such surrogate is the linear two-point function computed from the same initial conditions.¹⁰ We refer the interested reader to that work for further details, particularly as pertains to the implementation of CV in the context of DESI and AbacusSummit.

The uncertainties on our clustering measurements are obtained using a set of 1000 approximate EZmock simulations (Chuang et al. 2015) with the same volume, particle number density, and HOD as the AbacusSummit simulations. Reconstruction is applied to these approximate simulations in the same way as for the AbacusSummit realizations. Validation of these covariance matrices and comparisons with other approaches for estimating them can be found in Alves & DESI Collaboration (in preparation), Forero Sánchez & DESI Collaboration (in preparation), and Rashkovetskyi et al. (2024). Because we use a covariance matrix estimated from a finite number of simulations, we also debias the inverse covariance matrix and account

for noise in its estimation following Hartlap, Simon & Schneider (2007) and Percival et al. (2014, 2022).

Fig. 7 shows the mean power spectrum and correlation function of the 25 realizations along with error bars for a single realization and our best-fitting fiducial BAO model (see Section 7). This figure shows the extremely precise set of measurements available for us to test our modelling assumptions – and when fitting the mean of 25 realizations, we further reduce the covariance matrix by a factor 25 compared to that shown in the figure to represent the error on the mean. When fitting a single of the 25 individual LRG realizations, the resulting precision on the BAO constraints is ~ 0.3 per cent on α_{iso} and ~ 1 per cent on α_{ap} . Fitting the mock mean, with reduced covariance, gives a precision of ~ 0.1 per cent on α_{iso} and ~ 0.5 per cent on α_{ap} . For comparison, the expected aggregate precision on the BAO constraints for $0.0 < z < 1.1$ for DESI Y1 and the full DESI sample are 0.51 per cent (1.78 per cent) and 0.24 per cent (0.68 per cent) for α_{iso} (α_{ap}), respectively (DESI Collaboration 2024d). Hence, a single one of our simulations is sufficient to investigate systematics at the level of $\sim 2\times$ the precision of DESI’s Y1 measurements, while averaging over all 25 realizations is sufficient to determine systematics for the full DESI sample.

6.2 Numerical implementation in BARRY

To test the various theoretical systematics and modelling choices identified above, we use the modular BAO-fitting code BARRY (Hinton, Howlett & Davis 2020). This code is designed to be fully modular, easily enabling comparative fits to the same data when different modelling choices are made. This is done by specifying a series of unique model and data set pairs in one script, which are then submitted automatically as unique jobs to a supercomputing cluster. Hinton et al. (2020) used this architecture to study the impact of different perturbation-theory inspired models on the isotropic BAO

¹⁰That this is a good choice can be further understood from Section 4 that shows that the reconstructed density, particularly in RecSym, has its displacement terms reduced to $(1 - \mathcal{S}) * \Psi^{(1)}$. This displacement term is what leads to most of the decorrelation with the initial conditions.

signal (Seo et al. 2016; Beutler et al. 2017a; Ding et al. 2018; Noda, Peloso & Pietroni 2020; Chen et al. 2021), finding excellent agreement between the different models. However, as demonstrated in Section 4, we find that a more nuanced and complete treatment of theory leads to a single clear choice of reconstruction convention (RecSym), and a single BAO damping term. We hence fix our model to the one presented in Section 5. All of the remaining modelling choices outlined at the start of this section can be investigated by simple changes to the input parameters within BARRY.

To facilitate this, significant improvements have been made to the code relative to that presented in Hinton et al. (2020) which we will now describe. For all other base aspects of the code, we refer the reader to this reference. In contrast to Hinton et al. (2020), the latest version of BARRY implements anisotropic BAO fitting as default, reverting back to a standard isotropic method only when specified for particular model-data set pairs.

It is important to note that an ‘isotropic’ BAO fit, both within BARRY and in general, is not the same as running an ‘anisotropic’ fit to only the monopole with fixed $\epsilon = 0$. The reason for this is due to the impact of the survey window function. Both the window function convolution and modelling of wide-angle effects are implemented via matrix multiplications as described in Beutler & McDonald (2021), which also allows the fitting of the odd multipoles of the power spectrum if they are provided. Even when fitting only the monopole, the presence of a survey window function means that power leaks from higher order multipoles into the monopole and vice versa, requiring the full set of multipoles to be modelled and the window function for all even and odd multipoles to be used. When an isotropic fit is performed, the higher order multipoles are assumed to be zero and the part of the window function responsible for describing the leakage of power between these and the monopole is ignored. This is a fundamental difference in the modelling that is retained in BARRY for historical reasons and comparison to results from previous surveys, but in this work, we only consider the correct ‘anisotropic’ approach of modelling all higher order multipoles regardless of the exact parameters, scales or data multipoles being fit.

A second feature that is important to highlight is the ability to analytically marginalize over the broad-band terms during the BAO fitting. As these enter at linear order, such marginalization can be done fully (by modifying the likelihood used during MCMC fitting or optimization) or partially (by using a standard Gaussian likelihood and finding the best-fitting broad-band terms at each step of the fitting). Appendix F discusses this further and demonstrates the consistency between these approaches. All results presented hereon use full analytic marginalization, which we argue is the formally correct choice, although partial analytic marginalization is retained as an option in the code for comparison to previous works which used this approach (i.e. Ross et al. 2017; Bautista et al. 2021).

The full code used to perform all the tests in this work is available at <https://github.com/Samreay/Barry>.

7 RESULTS

This section presents the results of our various systematic modelling tests. Unless otherwise stated, we fit the post-reconstruction data with Gaussian priors on the BAO and FoG damping parameters, $\Sigma_{\perp} = \mathcal{N}(2, 1)$, $\Sigma_{\parallel} = \mathcal{N}(5, 2)$, and $\Sigma_s = \mathcal{N}(2, 2)$ on scales $0.02 h \text{ Mpc}^{-1} \leq k \leq 0.30 h \text{ Mpc}^{-1}$ and $50 h^{-1} \text{ Mpc} < s < 150 h^{-1} \text{ Mpc}$ for the power spectrum and correlation function respectively. Both the linear galaxy bias and RSD parameters are allowed to vary within wide flat prior ranges. We dilate only the

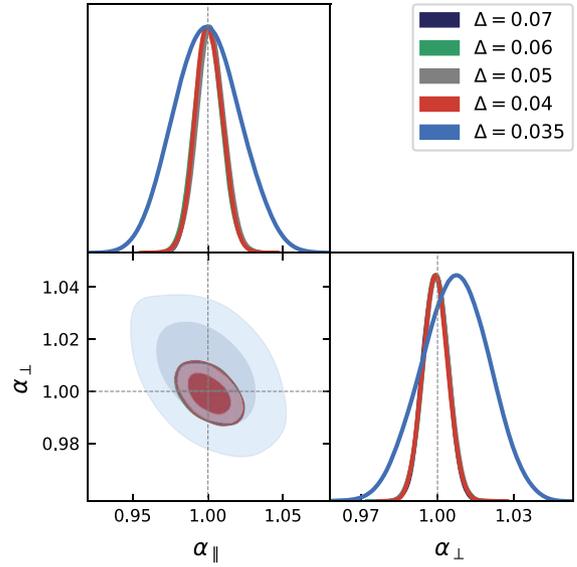


Figure 8. Fits to the mean power spectrum with varying levels of Δ . The constraints on the BAO parameters $\alpha_{\parallel, \perp}$ are remarkably stable to Δ up to $k \sim 0.04 h \text{ Mpc}^{-1}$; at around $\Delta \sim 0.035 h \text{ Mpc}^{-1}$, just shy of 10 per cent above the Nyquist frequency, the constraints rapidly degrade.

wiggle component of the model, and the FoG damping applies only to the smooth component. The wiggle/no-wiggle split is performed using the method of Wallisch (2018) and all our fits use the Nautilus sampler (Lange 2023).

7.1 Validating the spline-based broad-band model

Our first set of tests validates the new spline-based broad-band method. First, we test the dependence of our fit results to the smoothing scale Δ . Intuitively, splines with too small a Δ will be able to reproduce the BAO signal, while splines with too large a Δ will lack sufficient freedom to fit the broad band. This is especially clear in configuration space where the former corresponds with r_{\min} overlapping the BAO peak, while the latter corresponds with r_{\min} extending too far out from the BAO peak. Fortunately, no fine tuning in Δ is required as fits using the spline method are very stable over a range of Δ until the BAO scale is reached.

Fig. 8 shows fits to the mean LRG power spectrum with varying values of Δ . The constraints on $\alpha_{\parallel, \perp}$ are extremely stable to a broad range of Δ , and only begin to rapidly degrade around $\Delta \sim 0.035 h \text{ Mpc}^{-1}$ or roughly 10 per cent above the Nyquist frequency of the BAO wiggles. The central values of the fits shift by less than 0.1 per cent in all cases and the best-fitting χ^2 changes only by around the amount implied by the increase in free spline parameters with support within the fitting range ($k_{\max} = 0.3 h \text{ Mpc}^{-1}$). Since there is no apparent bias due to using a slightly wider spacing Δ but a clear degradation of the overall signal well before the Nyquist frequency is reached we pick roughly twice this frequency, $\Delta = 0.06 h \text{ Mpc}^{-1}$, as our fiducial choice. At $k_{\max} = 0.3 h \text{ Mpc}^{-1}$ this corresponds to seven free parameters per multipole.

Next, we confirm that the Fourier-space broad-band model is equivalent to setting $r_{\min} = \pi/\Delta$ with a minimal set of spline coefficients and large-scale polynomial terms. To do so, we perform two fits: one where the spline broad band including 20 bases with $\Delta = 0.06 h \text{ Mpc}^{-1}$ is Hankel transformed and the correlation function fit from $0 h^{-1} \text{ Mpc} < r < 130 h^{-1} \text{ Mpc}$ (summing from $n = -1$ to 19 in equation 32), and another where only the first two

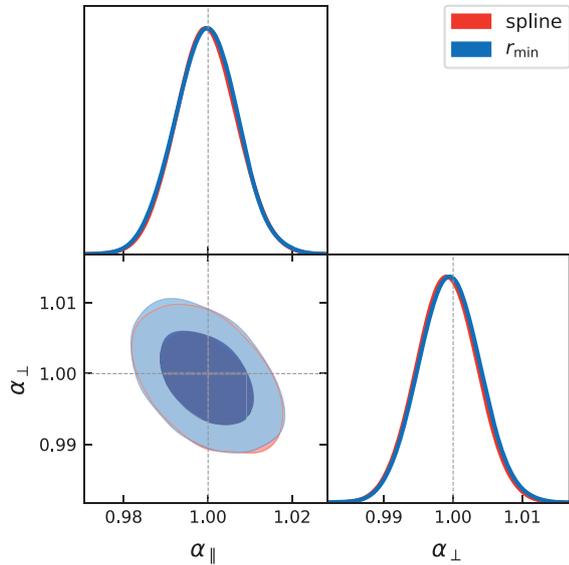


Figure 9. Fits to the mean correlation function using a wide scale range $0 h^{-1} \text{ Mpc} < r < 130 h^{-1} \text{ Mpc}$ and the full set of Hankel-transformed spline coefficients compared to when an r_{min} cut of $80 h^{-1} \text{ Mpc}$ is applied and only the $n = [0, 1]$ spline bases in the quadrupole are used. The excellent consistency between the two demonstrates that our default method of fitting the BAO removing small scales and unnecessary spline components is sufficient.

spline bases in the quadrupole with large-scale support in ξ_2 are kept (using only the $n = 0$ and 1 terms in equation 32), fitting to $80 h^{-1} \text{ Mpc} < r < 130 h^{-1} \text{ Mpc}$, where $80 h^{-1} \text{ Mpc}$ is chosen as the r_{min} where the spline basis has essentially no support. The two constraints are shown in Fig. 9. The BAO constraints are essentially identical, up to percent level differences in constraining power, despite the radial fitting range differing by close to a factor of two. Both the means and variances of the posteriors on the linear theory amplitudes b and f also agree at the roughly 5 per cent level, which is notable given the very different radial ranges over which the data are fit. This suggests that these linear-theory amplitude parameters are mostly fitting the amplitude of the BAO only, while sensitivity to the amplitude of the smooth component is removed either by r_{min} or by fitting the spline basis.

Finally, we would like to check that the fits are robust against various levels of noise. In order to perform this check we fit the data with the covariance rescaled by factors of $A = 1/4, 1/2, 1, 2, \text{ and } 4$. This corresponds approximately to rescaling the volume by the inverse factor – at $A = 1/4$ and 4 these rescalings correspond to more volume than all the DESI Y1 samples combined and less volume than the CMASS sample of the BOSS survey (Reid et al. 2016). The results in Fourier and configuration space are shown in Fig. 10. The fits are remarkably stable across these rescalings, with the error bars on $\alpha_{\parallel, \perp}$ scaling as \sqrt{A} to a very good approximation, demonstrating that the spline broad band does not sap BAO information out of the signal at any signal to noise. Indeed, we have tested that this remains true even for a $16\times$ rescaling, beyond which the BAO scale can essentially only be constrained to order unity.

7.2 Impact of BAO/FoG damping parameter choices

The BAO damping parameters Σ_{\parallel} and Σ_{\perp} play a key role in describing the shape of the measured BAO feature, which in turn impacts how well its peak position/wavelength can be determined.

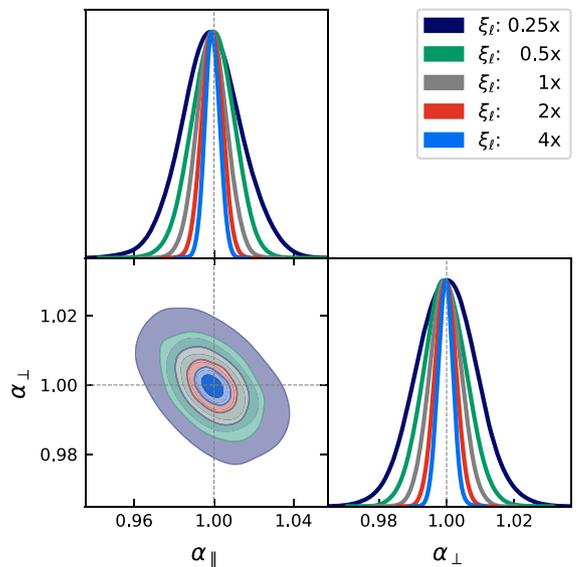
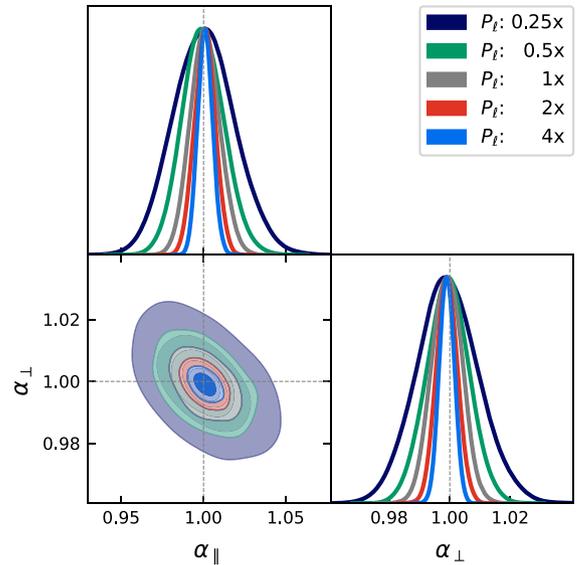


Figure 10. Fits to the mean power spectrum (top) and correlation function (bottom) when the covariance matrix is re-scaled between $0.25\times$ and $4\times$. The excellent consistency of the constraints demonstrates that our default spline-basis never removes BAO information regardless of the SNR level.

Furthermore, Σ_s is designed to account for non-linear RSD in the broad-band shape of the clustering but may interact with the BAO signal depending on how it is implemented. Previous studies (Alam et al. 2017; Gil-Marín et al. 2020; Bautista et al. 2021) have typically fixed the values of these parameters based on fits to a set of mocks, additionally demonstrating that fits to the data remain robust if different fixed values are used, or the parameters are freed up entirely. However, the treatment of Σ_s has not always been consistent, and it is not clear to what extent the priors on the three Σ parameters impact BAO constraints at DESI precision.

To test, this we first investigate how fixing these to different values changes the BAO constraints. Fig. 11 shows the bias $\Delta\alpha_{\text{iso}} = \alpha_{\text{iso}} - 1$ and similar for α_{ap} when we fit the mock mean clustering. Our results show that fixing the choice of damping parameters is somewhat risky. While the constraints on the BAO dilation parameters are remarkably

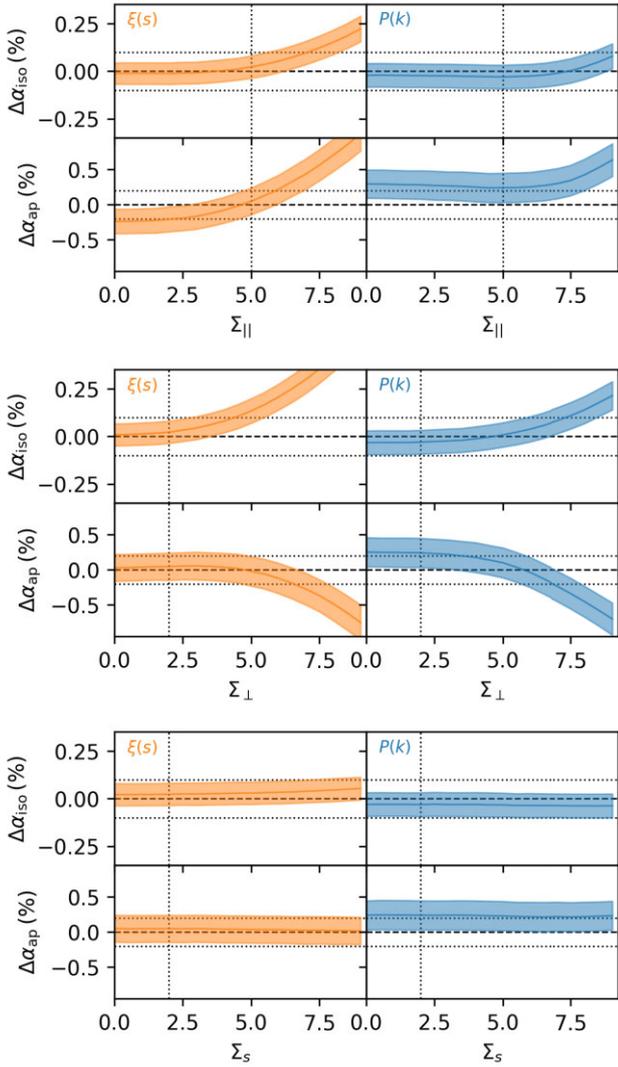


Figure 11. Percentage bias in the BAO dilation parameters as a function of the fixed damping scale used in the model when fitting the mean of the 25 LRG simulations. Each band shows the mean and standard deviation in the BAO parameters (relative to the expected value of 1) from fitting our mock mean measurements. Left panels show fits to the correlation function, while right panels are fits to the power spectrum. In all cases, the vertical lines show our fiducial choice of damping scales, while the horizontal lines are placed at characteristic values of 0.1 per cent for α_{iso} and 0.2 per cent for α_{ap} to guide the eye. These results suggest that the BAO model can accommodate small inaccuracies in the fiducial choice when fixing the damping scales, but that larger-than-ideal biases occur if the fiducial value is more than $\sim 2 h^{-1}$ Mpc from the correct choice. This motivates allowing for some freedom in the damping scale, that is, a Gaussian prior centred on the fiducial value with width $1 - 2 h^{-1}$ Mpc.

robust to moderate changes in the choices of damping scales, both the power spectrum and correlation function show systematic variation of increasing bias in the α 's when the BAO damping Σ 's are increased significantly – biases of order 0.1 per cent and 0.2 per cent for α_{iso} and α_{ap} respectively occur when the Σ values are changed by more than $\sim 2 h^{-1}$ Mpc from their fiducial values. Similar trends have been identified in previous works, see, for example, table B2 in Ross et al. (2017). We see no notable impact from varying Σ_s on either the correlation function or power spectrum, as expected given our default model decouples it from the wiggly power spectrum.

The bias is more pronounced for the combination of α_{iso} with Σ_{\perp} than with Σ_{\parallel} , while α_{ap} shows a complex behaviour from variations in either damping parameter. This can be understood from theoretical considerations – Taylor expanding the full BAO damping term in equation (26) and integrating with respect to the $\ell = 0$ and 2 Legendre polynomials allows us to write effective BAO damping scales for the monopole and quadrupole as

$$\Sigma_{\ell=0}^2 = \frac{1}{1 + 2\beta + \frac{3}{5}\beta^2} \left[\left(1 + \frac{6}{5}\beta + \frac{15}{35}\beta^2 \right) \Sigma_{\parallel}^2 + \left(2 + \frac{4}{5}\beta + \frac{6}{35}\beta^2 \right) \Sigma_{\perp}^2 \right], \quad (36)$$

$$\Sigma_{\ell=2}^2 = \frac{1}{2\beta + \frac{6}{7}\beta^2} \left[\left(1 + \frac{12}{7}\beta - \frac{5}{7}\beta^2 \right) \Sigma_{\parallel}^2 - \left(1 - \frac{2}{7}\beta - \frac{1}{7}\beta^2 \right) \Sigma_{\perp}^2 \right], \quad (37)$$

where $\beta = f/b$ is the linear RSD parameter. Information on α_{iso} arises primarily from the monopole and so bias in this parameter actually occurs when variations in Σ_{\parallel} or Σ_{\perp} cause a change in $\Sigma_{\ell=0}$. For the LRG mocks used here, where $\beta \approx 0.35$, the same constant change applied to Σ_{\perp}^2 has a $\sim 1.5\times$ greater impact on the damping scale in the monopole, compared to the same change applied to Σ_{\parallel}^2 . This is consistent with the larger bias seen in α_{iso} when Σ_{\perp} is varied compared to Σ_{\parallel} . Conversely, the damping scale in the quadrupole is more affected by changes in Σ_{\parallel} , but can also become too small if Σ_{\perp} is increased enough at fixed Σ_{\parallel} . As the quadrupole is the primary source of information on α_{ap} , considering variations in $\Sigma_{\ell=2}^2$ can explain the more nuanced behaviour of the α_{ap} as a function of both $\Sigma_{\parallel,\perp}$. We note that the dependence on β means that the picture will be slightly quantitatively different for other tracers, though qualitatively similar.

Overall, this suggests that some freedom should be given to the BAO damping parameters. We recommend using a Gaussian prior, centred on theory calculations such as in Fig. 5 or from fitting simulations, but with width $1 - 2 h^{-1}$ Mpc to allow for the presence of non-linearities not captured by theory or deviations between the mocks and data which may otherwise cause small 0.1 per cent – 0.2 per cent biases. When fitting the damping scale to simulations we advocate for fitting the BAO in high-precision simulated mocks (e.g. control-variate DESI mocks) using equation (26) with the cosmology set to ‘truth’ (i.e. $\alpha_{\perp,\parallel} = 1$) – previous works often instead fit to the ‘propagator’, that is, the correlation between the initial conditions and (reconstructed) galaxy density, but we note that this quantity is not strictly the same as the BAO damping since it is not cut-off at the BAO scale as in equation (9).

An alternative is to allow the BAO damping parameters to freely vary between wide flat priors such that the values are purely informed by the data. However, in other tests as part of this work, on more weakly constraining simulations than the CVs used here, we found that this generally reduces constraining power and has the potential to bias the BAO constraints. In the literature, nearly all previous BAO studies have elected to fix the damping parameters; of the studies we referred to only Gil-Marín et al. (2020) tested varying these, finding consistent BAO constraints, but we caution that the combined BOSS+eBOSS LRG data set used there is particularly well constraining. We also note that future studies could also investigate whether it is more useful to place priors on $\Sigma_{\ell=0}$ and $\Sigma_{\ell=2}$ rather than Σ_{\perp} and Σ_{\parallel} as our results suggest that it is changing along these degeneracy directions that ultimately cause biases in the BAO scale.

Having determined a reasonable choice of prior for the damping parameters, we then test for any impact of including the FoG damping term in the wiggly component of the model as well as the smooth

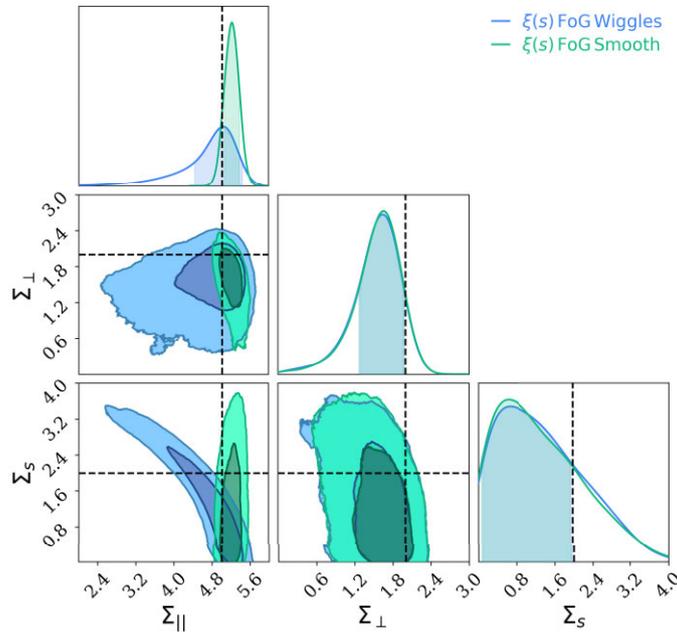


Figure 12. Constraints on the BAO and FoG damping parameters when fitting the mock mean correlation function with fixed $\alpha_{\text{iso}} = \alpha_{\text{ap}} = 1$ and including the FoG damping in only the smooth model component ($\xi(s)$ FoG smooth), or allowing it to multiply both the smooth and wiggle components ($\xi(s)$ FoG wiggles). The dashed lines represent the central values we use for the priors on these parameters in our other fits. Including the FoG damping in only the smooth model component is preferable as it removes a clear degeneracy between the BAO and FoG damping scales.

component, that is, including the Lorentzian term from equation (28) in equation (26). As a first demonstration of why we advocate for decoupling the FoG damping from the wiggle component of the model, we show in Fig. 12 that doing so effectively removes the degeneracy between Σ_{\parallel} and Σ_s , when fitting the mean of our mocks. This makes it far easier to enforce or determine suitable priors on these parameters in the manner described above and improves the convergence of the fitting. This breaking of the degeneracy occurs even when the α 's are allowed to vary in concert with the damping parameters.

Secondly, we look at the impact of this on the BAO constraints themselves by fitting each of our 25 individual mock realizations and then looking at the distribution of mock-to-mock differences in the α 's for the two different modelling configurations. The results are shown in Fig. 13, where we find that regardless of the choice of clustering statistic, or whether a polynomial-based or our new spline-based broad-band model is used, the largest average difference is ≤ 0.02 per cent and 0.11 per cent for α_{iso} and α_{ap} respectively. Albeit small, this latter difference is detected with 2σ significance and so should be accounted for in our final systematic error budget. That said, the largest difference seen in any individual realization is ~ 0.25 per cent (0.75σ) for α_{iso} and 0.75 per cent (0.75σ) for α_{ap} , indicating that while on average the results are robust to the choice of where the FoG damping is included, this choice could lead to statistical fluctuations on individual datasets dependent on the noise in the data.

7.3 Impact of different scale cuts

We next turn to verify the appropriate range of scales to include in the BAO fit. For both the power spectrum and correlation function, we fit the mock mean correlation with a variety of choices for the minimum and maximum scale cuts and measure the bias in the dilation parameters, and their uncertainties relative to our default.

These are shown in Figs 14 and 15 for the power spectrum and correlation function, respectively.

For the power spectrum, we find that the difference from unity exceeds 0.1 per cent and 0.2 per cent for α_{iso} and α_{ap} , respectively, for a couple of narrow ranges of k_{max} . The k_{max} at which this occurs is not the same for the two dilation parameters, and, more importantly, including even more non-linear scales does not make the bias worse. There is no significant trend of increasing bias when increasing/decreasing either k_{min} or k_{max} . As such, we attribute these regions with slight deviations from unity to noise in the simulations, rather than a bias in the model – looking at the measurements in Fig. 7, one can see a slight excess of power at around $k = 0.29h \text{ Mpc}^{-1}$ which inconveniently coincides with our fiducial choice of scales.

Turning to the uncertainty on the BAO parameters using the power spectrum, we do see a clear weakening of the constraints by up to 50 per cent when restricting the fitting to larger scales $k_{\text{max}} = 0.2h \text{ Mpc}^{-1}$ and some small 10 per cent improvement when including smaller scales $k_{\text{max}} = 0.4h \text{ Mpc}^{-1}$. There is also perhaps a very weak trend as k_{min} is increased and starts to encroach on the first BAO wiggle. It is worth reminding that as this is from a fit to the mean mock measurements, the uncertainties on the clustering measurements are significantly smaller than we expect for DESI DR1 and comparable to the aggregated BAO across $0.0 < z < 1.1$ for the full DESI survey. This exacerbates the trends in the uncertainties, and so we conclude that our fiducial choice of fitting scales $0.02h \text{ Mpc}^{-1} < k < 0.30h \text{ Mpc}^{-1}$ is reasonable – it is not worth extending the fit to smaller scales for DESI DR1 or its upcoming DESI 2024 results. But, as we see no evidence of additional bias, this may be worth testing for the full DESI data.

For the correlation function, we find that the results are remarkably stable to the choice of scale cuts, with a small increase in bias but a decrease in error as s_{min} is reduced. The bias reaches 0.1 per cent in α_{iso} for $s_{\text{min}} < 28h^{-1} \text{ Mpc}$ across the range of s_{max} we test, and ~ 0.15 per cent in α_{ap} when a maximum fitting scale

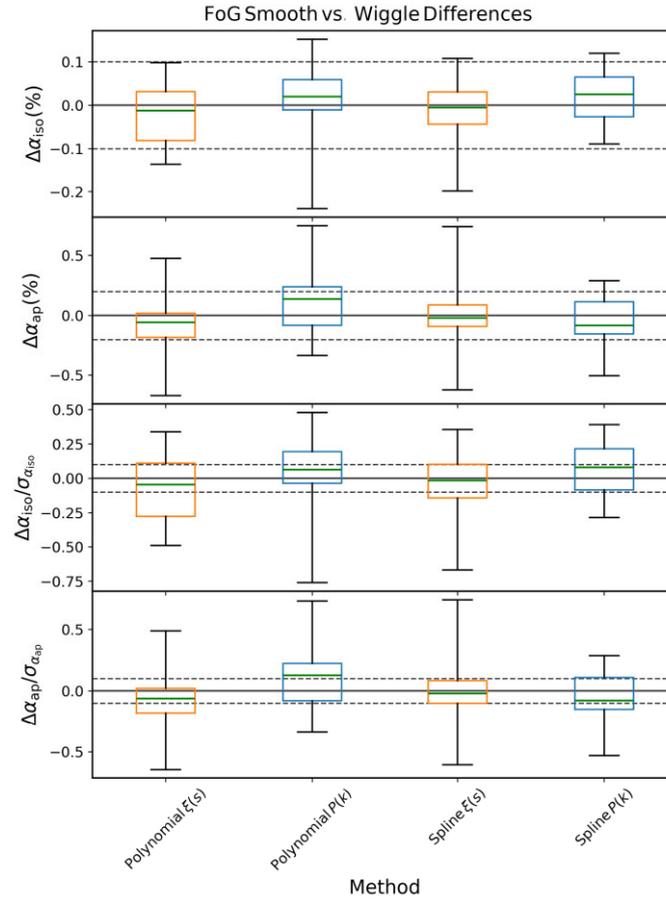


Figure 13. Difference in the BAO constraints between our 25 mock realizations fit with and without the FoG damping applied to the wiggle component in addition to the smooth component (without is our default). Each box-and-whisker plot shows the median difference over the individual mocks (green line) the interquartile range (blue/orange boxes for power spectrum/correlation function, respectively) and the largest individual differences (black error bars). These plots are shown for both the polynomial and spline-based broad-band methods. The top two panels show the percentage differences between the same realization, while the bottom two panels show the difference relative to the average statistical uncertainty on the BAO parameter for a single realization. The horizontal lines are to guide the eye and are placed at 0.1 for $\Delta\alpha_{\text{iso}}$ and the error-normalized differences, and 0.2 per cent for $\Delta\alpha_{\text{ap}}$.

$s_{\text{max}} < 130 h^{-1} \text{ Mpc}$ is also applied. When the upper fitting range is reduced to this, there is enough freedom in the quadrupole model (particularly in its amplitude/BAO damping) that one starts to affect the BAO information if additional scales are not included on either side of the BAO feature with which to constrain the broad-band model. Beyond this, any choice of fitting scale seems as good as any other, with < 10 per cent relative change in the uncertainties across all choices, so we opt to fix to our fiducial choice of $50 h^{-1} \text{ Mpc} < s < 150 h^{-1} \text{ Mpc}$.

7.4 Impact of de-wiggling methodology

When performing a BAO fit, a choice must be made of how to isolate the wiggle and no-wiggle components in the linear power spectrum. A number of different algorithms/approaches for this have been developed; using fitting formulae (Eisenstein & Hu 1998), filters designed for signal processing (such as the Savitsky–Golay filter, B-splines, or a Gaussian filter as used in Hinton et al. 2017), or methods that numerically search for changes in the gradient in the power spectrum or its Fourier transform (Wallisch 2018; Brieden, Gil-Marín & Verde 2022). As shown in Fig. 16, the wiggle component of the model arising from these different choices of algorithm can differ

appreciably and, without a clear theoretical reason to prefer one over the other, the impact of this choice must be explored.

To do this, we follow the same procedure used previously and fit our 25 individual LRG mock realizations using each of the different template extraction methods. We then take the mock-to-mock difference between the constraints relative to our fiducial Wallisch (2018) method. The results are shown in Fig. 17. Despite the clear differences in the templates, the systematic difference in the BAO constraints is < 0.02 per cent (< 0.09 per cent) for α_{iso} (α_{ap}). However, as with the case where we investigate the impact of including the FoG damping in both the smooth and wiggle model components, individual realizations can show larger statistical variations when the template is varied – up to 0.2 per cent (0.75σ) for α_{iso} , and 0.5 per cent (0.5σ) for α_{ap} . The correlation function appears slightly more sensitive to these variations than the power spectrum, consistent with the larger variations between templates over our fitting range seen in the correlation function of Fig. 16.

7.5 Impact of dilating the smooth model component or not

BAO fitting constraints arise from dilating the scales in the two-point clustering signal by the α parameters and then comparing the dilated model (equation 5) to the data. In our fiducial setup, we dilate only

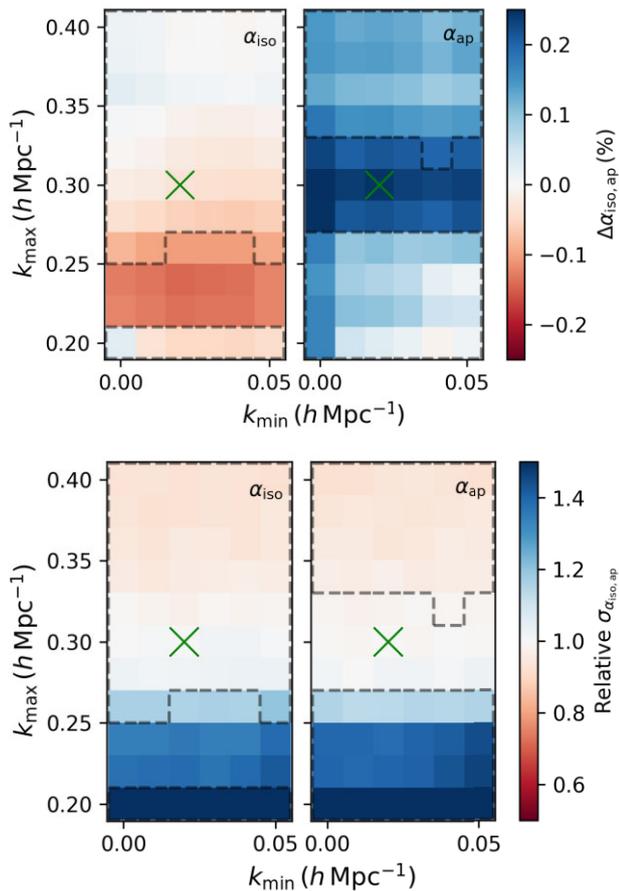


Figure 14. The impact of varying minimum and maximum scale cuts (k_{\min} and k_{\max}) on BAO dilation parameters extracted from the mock mean LRG power spectrum. The top panel is coloured according to the percentage deviation of the fitted α 's from unity. The bottom panel is the uncertainty of the α 's relative to our fiducial fitting range (indicated by the green cross). In both cases, the left (right) panel is for α_{iso} (α_{ap}). The regions enclosed by the dashed line in both of the left panels are where the deviation in α_{iso} from unity is < 0.1 per cent. For the right panels, it denotes a deviation in α_{ap} of < 0.2 per cent.

the wiggle term $C(k, \mu)P_w(k)$, evaluating the remainder of the model at the undilated scales/angles k and μ . This ensures that constraints on the α 's arise only from the wiggle component. However, one could also apply the dilation to the smooth component of the model. Previous results in the literature have often taken this latter approach (e.g. Ross et al. 2017; Beutler et al. 2017a; Gil-Marín et al. 2020; Bautista et al. 2021).

Following the same approach as above, fitting the 25 individual LRG mocks with both approaches and looking at the distribution of differences in the BAO constraints, we find that the choice of whether or not to dilate the smooth model component in addition to the wiggles results in a systematic error of < 0.02 per cent (< 0.09 per cent) for α_{iso} (α_{ap}). The largest individual differences are ~ 0.15 per cent (0.5σ) for α_{iso} , and 0.75 per cent (0.75σ) for α_{ap} .

7.6 Impact of different methods for transforming from the power spectrum to the correlation function

In this work, we provide a fully consistent/equivalent model for the correlation function and power spectrum, treating the former purely

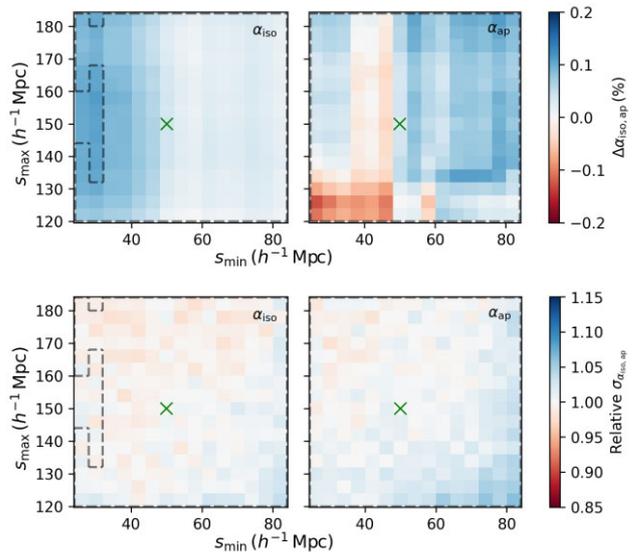


Figure 15. As for Fig. 14, but showing BAO constraints from the correlation function when varying the minimum and maximum fitting scales s_{\min} and s_{\max} .

as the Hankel transform of the latter, and applying all our modelling choices and BAO dilation to the power spectrum. The new spline broad-band method also ensures a consistent modelling between the two. However, some previous analyses (Anderson et al. 2014; Ross et al. 2017) instead transformed the undilated power spectra, without galaxy bias, and applied both the BAO dilation and galaxy bias terms in configuration space. They also used two different galaxy bias terms for the two correlation function multipoles rather than including a physical angular-dependent RSD with free parameter f or β . Equations (13)–(17) and accompanying text in Ross et al. (2017) describe this process.

We estimate the impact of the differences between these two methods by implementing both and fitting the individual LRG realizations using the spline-based broad band. We find a negligible systematic difference in α_{iso} , but a detection of a small difference of 0.12 ± 0.03 per cent in α_{ap} . The fact that the two approaches differ mainly in their modelling of the quadrupole and the difference occurs in α_{ap} indicates the presence of a real difference, not just noise. Although we believe that our new approach is preferable, we hence accommodate for this difference in our total systematic error on α_{ap} . On individual realizations, we find that the largest difference between the two methods is 0.04 per cent (0.15σ) for α_{iso} , and 0.4 per cent (0.4σ) for α_{ap} . Despite the small detection of systematic differences when averaging over all 25 realizations, this is consistent with the statistical differences on individual realizations seen from other modelling choices.

7.7 Impact of different parameter sampling methods

The last choice we investigate is less about the evaluation of the modelling for the BAO, but rather how the constraints on the α 's are obtained numerically. BARRY exposes a number of well-known sampling algorithms with which one can fit the data and recover posterior samples (and possibly evidences) on the dilation parameters – these include the publicly available packages EMCEE, DYNESTY, ZEUS, and NAUTILUS (Foreman-Mackey et al. 2013; Speagle 2020;

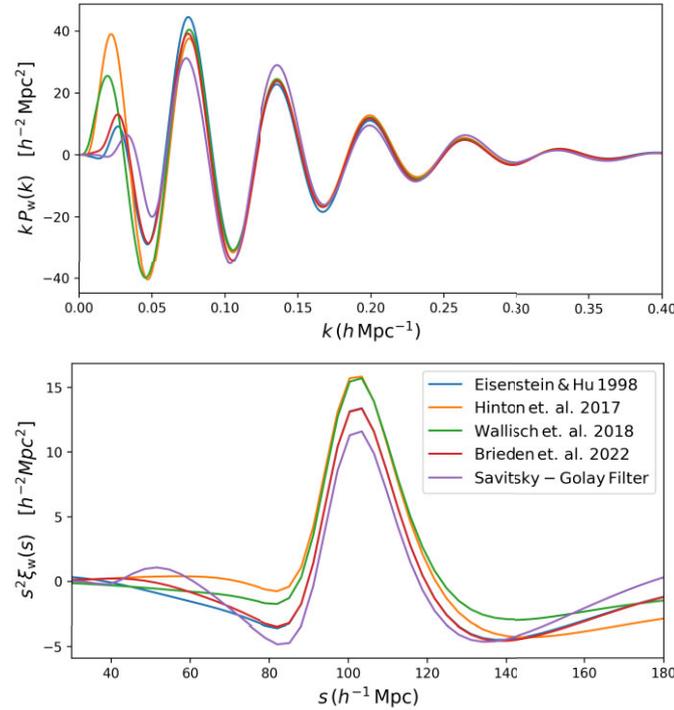


Figure 16. A comparison of the BAO template extracted using a variety of different literature approaches applied to the same linear power spectrum.

Karamanis, Beutler & Peacock 2021; Lange 2023), each with automated convergence checking.

Fig. 18 shows the marginalized posterior means on the α 's for our individual realizations for each of the samplers we test, relative to those obtained using NAUTILUS. We also investigate the width of the 68 per cent and 95 per cent confidence intervals on the α 's. Given these samplers are all well tested and use reasonable automated convergence checks we do not expect and do not see any systematic differences. What is interesting to quantify however is that there is statistical differences between the posterior means for individual realizations depending on the choice of sampler – up to 0.1 per cent and 0.5 per cent on α_{iso} and α_{ap} respectively. In all cases, we see much better convergence in the 68 per cent and 95 per cent confidence intervals than the posterior mean, consistent with statistical sampling noise around the peak of the posterior.

We measured the speed to convergence of these various samplers across the 25 realizations and found that NAUTILUS performed significantly faster than other methods. It also has the benefit of returning the Bayesian evidence as well as posterior samples. For these reasons, we adopted NAUTILUS as our default method, including in this paper.

7.8 Summary of methodology systematics and differences between our new method and the (e)BOSS methodology

Overall, our tests in this Section have shown that BAO measurements are remarkably robust to model choices not specified by theory. We find that, if our suggested best practices (summarized further below) are followed, we expect a systematic error of no more than 0.05 per cent in α_{iso} and 0.12 per cent in α_{ap} to arise from any modelling choices. For clarity compared to past studies and as a guide for the future, we now reiterate our best practices, highlighting differences with previous literature:

(i) For BAO reconstruction, previous works (Anderson et al. 2012, 2014; Xu et al. 2012; Alam et al. 2017, 2021) largely followed RecIso convention. In contrast, we have demonstrated a clear theoretical preference for the RecSym reconstruction convention [see Section 4.1 for a description of these two conventions, and equations (24) and (26) for the corresponding RecSym propagators]. Our companion papers Chen (in preparation) and Paillas et al. (2024) further detail how this can be applied optimally in the context of DESI.

(ii) We use the spline-based broad-band model validated in Section 7.1. Previous analyses used polynomials (of slightly varying degrees and functional forms) to fit the broad band. We find the spline method has better theoretical motivation, and to be less arbitrary in its choice of free parameters. For our default fitting of the monopole and quadrupole, between the ranges of $k = 0.02 - 0.30 h \text{ Mpc}^{-1}$ and $s = 50 - 150 h^{-1} \text{ Mpc}$ and with a spline width $\Delta = 0.06 h \text{ Mpc}^{-1}$, this gives 14 free broad-band parameters for the power spectrum and 6 for the correlation function. We again clarify that the broad band between the two has the same theoretical basis – the difference in number of parameters stems only from the behaviour of the Hankel transform and the different fitting ranges used for the two sets of clustering measurements.

(iii) We recommend choosing the BAO damping parameters based on theoretical calculations and allowing these to vary within reasonably tight priors ($1 - 2 h^{-1} \text{ Mpc}$). For DESI, we generally use Gaussian priors of width $1 h^{-1} \text{ Mpc}$ for Σ_{\perp} and $2 h^{-1} \text{ Mpc}$ for Σ_{\parallel} and Σ_s , corresponding to approximately 20 per cent variations from their expected values. Central values for these priors are chosen based on fits to simulated data and theoretical considerations. More info on the priors adopted for our fiducial data analysis can be found in DESI Collaboration (2024a). We found that completely fixing the BAO damping has the potential to lead to biased α 's if the wrong values are chosen, while allowing them to vary freely can

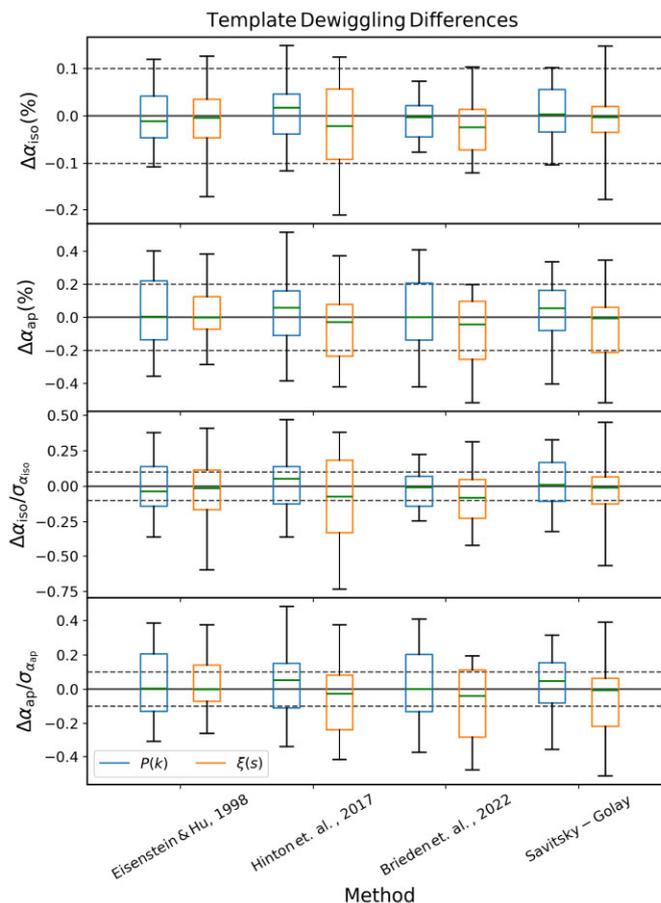


Figure 17. As for Fig. 13, but showing the differences in the BAO constraints between our 25 mock realizations using different BAO template extraction methods. In all cases, the template cosmology and linear power spectrum are the same and the differences are taken relative to our fiducial (Wallisch 2018) method.

also cause bias, and generally weak constraints, in the case of noisy data. Accounting for shot noise and non-linearities is able to give robust analytic predictions. Previous BOSS/eBOSS analyses fixed these based on fits to simulations (although found consistent results compared to when these were allowed to freely vary).

(iv) We apply the FoG damping to only the smooth model components, allowing for no interaction with the BAO wiggles. This makes interpreting the impact of the BAO damping parameters simpler. Previous models applied the FoG term equally to both the wiggle and no-wiggle components.

(v) The DESI BAO template is constructed following the method of Wallisch (2018). Previous BOSS and eBOSS analyses have used varying methods for extracting the template (i.e. Eisenstein & Hu 1998; Kirkby et al. 2013). Our tests do not point to a single correct choice, but we find a maximum systematic error associated with this choice of only 0.02 per cent and 0.09 per cent for α_{iso} and α_{ap} , respectively when following our fiducial fitting methodology.

(vi) We dilate only the BAO wiggles by the α 's. Previous analyses generally dilated both the smooth and wiggle components of the model. We believe the new method to be preferable as it guarantees only information from the BAO is captured. Given the large amount of freedom in the broad band used in previous analyses, we do not expect any information was included from the full shape of the clustering with the previous method, but our model choice ensures this remains true even if the broad-band model is less flexible.

(vii) We give both the galaxy bias b and RSD term β flat priors, finding that difficulties determining the appropriate central value of other prior choices can lead to poor fits on (a small number) of individual realizations. Previous analyses typically used more informative (but reasonably wide) Gaussian priors on one or both of these.

(viii) We provide a fully consistent/equivalent model for the correlation function and power spectrum, treating the former purely as the Hankel transform of the latter, and applying all our modelling choices and BAO dilation to the power spectrum. BOSS analyses instead transformed the undilated power spectra, without galaxy bias, and applied both the BAO dilation and galaxy bias terms in configuration space. They also used two different galaxy bias terms for the correlation function multipoles rather than the RSD parameter β . We find the new, consistent, modelling method to be preferable, and find 0.12 per cent of systematic difference in α_{ap} between the two methods.

8 CONCLUSIONS

BAOs are one of the premier probes of the cosmic expansion history, especially at high redshifts, providing percent level and below measurements of cosmological distances up to $z = 2$. Since BAO are an oscillatory signal in Fourier space, with a distinct and fast oscillation frequency set by the sound horizon r_d at the baryon drag epoch, its imprint on large-scale structure cannot easily be reproduced by

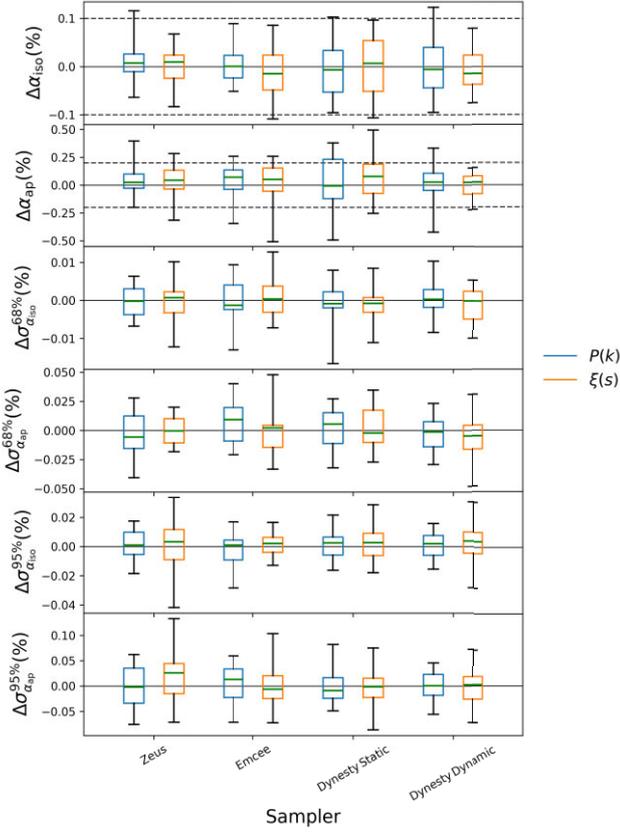


Figure 18. Difference in the BAO constraints between our 25 mock realizations fit with different MCMC/Nested samplers, relative to those obtained using NAUTILUS. Each box-and-whisker plot shows the median difference over the individual mocks (green line) the interquartile range (blue/orange boxes for power spectrum/correlation function, respectively) and the largest individual differences (black error bars). The top two panels show the percentage differences between the same realization, the middle two panels show the differences between the width of the 68 per cent confidence intervals (i.e. in the 1σ bounds), and the bottom two panels show the same for the 95 per cent confidence intervals. The horizontal lines are to guide the eye and are placed at 0.1 for $\Delta\alpha_{\text{iso}}$ and 0.2 per cent for $\Delta\alpha_{\text{ap}}$.

either survey systematics or non-linearities, making the measured BAO scale remarkably robust. Indeed, in the galaxy power spectrum, the latter typically induces shifts of no more than a percent, shifts which are further reduced by reconstruction. On the other hand, modern spectroscopic surveys, both planned and currently online, promise to probe galaxy clustering with larger volumes and greater number densities than ever before, leading to BAO measurements with statistical power verging upon the bias naively induced by these systematic effects. Indeed, the DESI, which will soon be releasing its DR1 is projected to yield a cumulative 0.2 per cent measurement of the BAO scale when its five-year run is completed.

Given the above, our goal in this paper has been to perform a thorough accounting of the theory and modelling systematics incurred in measuring the non-linear BAO in galaxy clustering. Our approach has been twofold: first, we reviewed the theoretical literature modelling the BAO signal in galaxies, filling in holes, and updating results where necessary, in order to both estimate the shift in the BAO scale due to non-linearities and suggest an appropriate fitting form for the observed signal. Secondly, we test these methods on a set of high-precision mocks whose sample variance has been reduced using Zeldovich CVs, paying particular

attention to systematic effects of numerical choices not uniquely specified by the theory.

In the first part of the paper, we (re)examine the PT of the BAO, particularly as pertains to the non-linear shift of the non-linear BAO peak and its damping by long-wavelength displacements. Two of the largest contributions to the former are (i) a non-linear ‘contraction’ of the BAO due to density modes larger than r_d , whose effect in redshift space we derive and (ii) imprints of relative baryon-dark matter perturbations on galaxy clustering. The former effect is substantially larger in redshift space, where spectroscopic surveys naturally operate, than previous calculations in real-space suggest, but is significantly diminished after reconstruction, up to an error induced by a mismatch of the fiducial and true matter clustering amplitude $\Delta\sigma_8^2/\sigma_8^2$. The impact of these two non-linear effects can be computed as a function of bias and redshift and added to the theoretical error budget of BAO measurements. Post reconstruction, assuming that the matter clustering on large scales $\propto \sigma_8$ is known to 5 per cent, effect (i) leads to shifts in the BAO scale of no more than 0.1 per cent in the regimes relevant for DESI, while effect (ii) leads to biases of less than 0.05 per cent.

While it has long been known that reconstruction reduces the non-linear damping of the BAO in galaxy clustering due to relatively long-wavelength modes, there has been significant confusion in the literature about the appropriate way to model their effect analytically, with many works suggesting a complicated scenario where the displaced galaxies and randoms have different damping scales (e.g. Ding et al. 2018; Chen et al. 2019b). Here we show that in the RecSym scheme, where both galaxies and randoms are moved by the same redshift-space displacement, this is in fact not the case; the long modes in fact produce a simple and unified exponential damping across both the shifted randoms and galaxies. On the other hand, we show that the Reclso scheme adopted by previous galaxy surveys like BOSS and eBOSS results in a complicated damping form where the effect of modes longer than the BAO are not controlled, leading to non-negligible shifts in the measured BAO in the process. For this reason, the RecSym scheme is the preferred one, adopted for DESI Y1 and which we recommend for future BAO analyses, though we note that in practice the differences are relatively small, especially at the statistical level of DESI year one and existing surveys, and we refer the reader to the accompanying references, Chen (in preparation) and Paillas et al. (2024) for further validation on simulations. In addition, in the process of preparing this paper, we learned that Sugiyama (2024) had reached a similar conclusion on the cancellation of long-wavelength modes in reconstruction schemes where galaxies and randoms are moved by the same reconstructed displacement. Where our paper focused on deriving the infrared (IR) resummation of the reconstructed BAO and the cancellation of the mode-coupling induced phase shift, Sugiyama (2024) instead performed a more detailed and general look at the cancellation of long-wavelength displacements in the power spectrum as a whole. We refer readers more interested in the latter aspect to that work, which we have coordinated to appear at a similar time with this one and which should be rather complementary.

In the second part of the paper, we investigate the impact of a number of other modelling choices, using numerical fits to a set of highly precise simulations. We arrive at a simple set of best practices for how to model the BAO, enumerated in Section 7.8. We demonstrate that if these are followed we expect a conservative systematic modelling error budget of 0.05 per cent and 0.12 per cent on α_{iso} and α_{ap} , respectively. Any larger deviations from the expected α values can be attributed to measurement noise. As will be

Table 1. Different contributions to the DESI BAO theory and modelling systematic error budget for α_{iso} and α_{ap} considered in this work. The section column denotes where the discussion of this contribution primarily occurs. In the theory cases, where the exact contributions depend on the nature of the tracers (galaxy bias, redshift, etc.), we have estimated a conservative value considering all DESI tracers/redshift bins.

Name/description	Section	$\sigma_{\alpha_{\text{iso}}}$ (per cent)	$\sigma_{\alpha_{\text{ap}}}$ (per cent)
Non-linear mode coupling	Sections 3.1 and 4.2	< 0.1	< 0.1
Relative velocity effects	Sections 3.1 and 4.2	< 0.05	< 0.05
Broad-band modelling	Section 7.2	< 0.02	0.11
BAO wiggle extraction	Section 7.4	< 0.02	< 0.09
Dilating smooth versus wiggle	Section 7.5	< 0.02	< 0.09
Modelling $\xi(s)$ from $P(k)$	Section 7.6	< 0.01	0.12
Combined	Section 8	0.1	0.2

demonstrated in upcoming publications, analysis of the DESI Y1 data follows these best-practices.

Table 1 summarizes our systematic error contributions from both theory and modelling. Combining our two prongs of investigation we find that the combined systematics in BAO fitting can be controlled, conservatively, to within 0.1 per cent for α_{iso} and 0.2 per cent for α_{ap} . We arrive at these numbers by summing in quadrature all the detected systematics in Table 1 and rounding up to account for the unresolved systematics for which we have only upper limits. These are the values that will be adopted for modelling and theory systematics in the DESI Y1 analysis and are substantially less than the aggregate $0.0 < z < 1.1$ precision on the BAO expected from this data. They are also more than a factor of 2 less than the expected aggregate BAO precision for the full DESI sample, indicating that even with this exceptional data set, we expect to remain in a regime where the uncertainties are dominated by statistics rather than systematics.

Looking to the future, it is clear that in order to further validate modelling systematics for the BAO beyond that which has been shown here, a method to test our models beyond running large numbers of N -body simulations will have to be found. First, the statistical errors even from fitting sample-variance reduced DESI mocks in this work prohibit us from clearly identifying (i.e. at $\geq 3\sigma$) systematic errors at levels below ~ 0.1 per cent. Producing more simulations with the required precision and fidelity to overcome this will be rather costly. Secondly, we are in a realm of accuracy where the potential errors induced by galaxy clustering physics that cannot be easily introduced into N -body simulations are now important and, in addition, numerical simulations struggle due to approximations in the initial condition generation, finite computational resources, and numerical approximations made during the evolution and post-processing of the simulations themselves (Angulo & Hahn 2022; Ding et al. 2022; Grove et al. 2022). While sufficient for DESI, addressing this is of the utmost importance if we aim to determine the limiting systematic precision with post-DESI BAO measurements.

ACKNOWLEDGEMENTS

We thank Naonori Sugiyama for useful discussions, particularly on the perturbative treatment of infrared effects in reconstruction. We thank Alejandro Aviles and Florian Beutler for graciously agreeing to serve on the internal DESI review committee and providing helpful feedback and additionally thank Will Percival and Sesh Nadathur for useful comments.

SC acknowledges the support of the National Science Foundation at the Institute for Advanced Study. CH acknowledges support from the Australian Government through the Australian Research Council's Laureate Fellowship (project FL180100168) and Dis-

covery Project (project DP20220101395) funding schemes. H-JS acknowledges support from the U.S. Department of Energy, Office of Science, Office of High Energy Physics under grant nos DE-SC0019091 and DE-SC0023241. H-JS also acknowledges support from Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under contract no. DE-AC02-05CH1123 during the sabbatical visit. The author list was produced using the publicly available `mkauthlist`.¹¹

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under contract no. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under contract no. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technology Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Science and Technology of Mexico (CONACYT); the Ministry of Science and Innovation of Spain (MICINN), and by the DESI Member Institutions: <https://www.desi.lbl.gov/collaborating-institutions>. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U.S. National Science Foundation, the U.S. Department of Energy, or any of the listed funding agencies.

The authors are honoured to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

DATA AVAILABILITY

Data from the plots in this paper are available on Zenodo (<https://zenodo.org/records/10685759>) as part of DESI's Data Management Plan.

REFERENCES

- Abdalla E. et al., 2022, *J. High Energy Astrophys.*, 34, 49
- Abidi M. M., Baldauf T., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 029
- Alam S. et al., 2017, *MNRAS*, 470, 2617
- Alam S. et al., 2021, *Phys. Rev. D*, 103, 083533
- Alcock C., Paczynski B., 1979, *Nature*, 281, 358

¹¹<https://github.com/DarkEnergySurvey/mkauthlist>

- Anderson L. et al., 2012, *MNRAS*, 427, 3435
- Anderson L. et al., 2014, *MNRAS*, 441, 24
- Angulo R. E., Hahn O., 2022, *Living Rev. Comput. Astrophys.*, 8, 1
- Baldauf T., Mirbabayi M., Simonović M., Zaldarriaga M., 2015, *Phys. Rev. D*, 92, 043514
- Ballinger W. E., Peacock J. A., Heavens A. F., 1996, *MNRAS*, 282, 877
- Barreira A., Cabass G., Nelson D., Schmidt F., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 005
- Bashinsky S., Seljak U., 2004, *Phys. Rev. D*, 69, 083002
- Baumann D., Beutler F., Flauger R., Green D., Slosar A., Vargas-Magaña M., Wallisch B., Yèche C., 2019, *Nat. Phys.*, 15, 465
- Bautista J. E. et al., 2021, *MNRAS*, 500, 736
- Bayer A. E., Banerjee A., Seljak U., 2022, *Phys. Rev. D*, 105, 123510
- Bernal J. L., Smith T. L., Boddy K. K., Kamionkowski M., 2020, *Phys. Rev. D*, 102, 123515
- Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, *Phys. Rep.*, 367, 1
- Beutler F., McDonald P., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 031
- Beutler F. et al., 2011, *MNRAS*, 416, 3017
- Beutler F. et al., 2017a, *MNRAS*, 464, 3409
- Beutler F., Seljak U., Vlah Z., 2017b, *MNRAS*, 470, 2723
- Beutler F., Biagetti M., Green D., Slosar A., Wallisch B., 2019a, *Phys. Rev. Res.*, 1, 033209
- Beutler F., Castorina E., Zhang P., 2019b, *J. Cosmol. Astropart. Phys.*, 2019, 040
- Bharadwaj S., 1996, *ApJ*, 472, 1
- Blake C. et al., 2011, *MNRAS*, 415, 2892
- Blake C., Carter P., Koda J., 2018, *MNRAS*, 479, 5168
- Blas D., Lesgourgues J., Tram T., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 034
- Blas D., Garny M., Ivanov M. M., Sibiryakov S., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 028
- Blazek J. A., McEwen J. E., Hirata C. M., 2016, *Phys. Rev. Lett.*, 116, 121303
- Bridle S. L., Crittenden R., Melchiorri A., Hobson M. P., Kneissl R., Lasenby A. N., 2002, *MNRAS*, 335, 1193
- Brieden S., Gil-Marín H., Verde L., 2022, *J. Cosmol. Astropart. Phys.*, 2022, 024
- Carlson J., Reid B., White M., 2013, *MNRAS*, 429, 1674
- Carter P., Beutler F., Percival W. J., DeRose J., Wechsler R. H., Zhao C., 2020, *MNRAS*, 494, 2076
- Castorina E., Di Dio E., 2022, *J. Cosmol. Astropart. Phys.*, 2022, 061
- Castorina E., White M., 2018, *MNRAS*, 476, 4403
- Castorina E., Carbone C., Bel J., Sefusatti E., Dolag K., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 043
- Chartier N., Wandelt B., Akrami Y., Villaescusa-Navarro F., 2021, *MNRAS*, 503, 1897
- Chen S.-F., Castorina E., White M., 2019a, *J. Cosmol. Astropart. Phys.*, 2019, 006
- Chen S.-F., Vlah Z., White M., 2019b, *J. Cosmol. Astropart. Phys.*, 2019, 017
- Chen S.-F., Vlah Z., White M., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 035
- Chen S.-F., Vlah Z., Castorina E., White M., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 100
- Chen S.-F., Vlah Z., White M., 2024, preprint (arXiv:2406.00103)
- Chuang C.-H., Kitaura F.-S., Prada F., Zhao C., Yepes G., 2015, *MNRAS*, 446, 2621
- Cole S., Fisher K. B., Weinberg D. H., 1995, *MNRAS*, 275, 515
- Cole S. et al., 2005, *MNRAS*, 362, 505
- Cooke R. J., Pettini M., Steidel C. C., 2018, *ApJ*, 855, 102
- Crocce M., Scoccimarro R., 2008, *Phys. Rev. D*, 77, 023533
- d'Amico G., Gleyzes J., Kokron N., Markovic K., Senatore L., Zhang P., Beutler F., Gil-Marín H., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 005
- DESI Collaboration, 2016, preprint (arXiv:1611.00036)
- DESI Collaboration, 2022, *AJ*, 164, 207
- DESI Collaboration, 2024a, preprint (arXiv:2404.03000)
- DESI Collaboration, 2024b, preprint (arXiv:2404.03001)
- DESI Collaboration, 2024c, preprint (arXiv:2404.03002)
- DESI Collaboration, 2024d, *AJ*, 167, 62
- DESI Collaboration, 2024e, *AJ*, 168, 58
- Davis M., Peebles P. J. E., 1983, *ApJ*, 267, 465
- Dawson K. S. et al., 2013, *AJ*, 145, 10
- DeRose J., Chen S.-F., Kokron N., White M., 2023, *J. Cosmol. Astropart. Phys.*, 2023, 008
- Desjacques V., Jeong D., Schmidt F., 2018, *Phys. Rep.*, 733, 1
- Diaferio A., Geller M. J., 1996, *ApJ*, 467, 19
- Ding Z., Seo H.-J., Vlah Z., Feng Y., Schmittfull M., Beutler F., 2018, *MNRAS*, 479, 1021
- Ding Z. et al., 2022, *MNRAS*, 514, 3308
- Dodelson S., Schmidt F., 2020, *Modern Cosmology*. Academic Press, London
- Eisenstein D. J., Hu W., 1998, *ApJ*, 496, 605
- Eisenstein D., White M., 2004, *Phys. Rev. D*, 70, 103523
- Eisenstein D. J. et al., 2005, *ApJ*, 633, 560
- Eisenstein D. J., Seo H.-J., White M., 2007a, *ApJ*, 664, 660
- Eisenstein D. J., Seo H.-J., Sirko E., Spergel D. N., 2007b, *ApJ*, 664, 675
- Feldman H. A., Kaiser N., Peacock J. A., 1994, *ApJ*, 426, 23
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
- Friedmann A., 1922, *Z. Phys.*, 10, 377
- García-Quintero C. et al., 2024, preprint (arXiv:2404.03009)
- Gil-Marín H. et al., 2020, *MNRAS*, 498, 2492
- Grimm N., Scaccabarozzi F., Yoo J., Biern S. G., Gong J.-O., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 064
- Grove C. et al., 2022, *MNRAS*, 515, 1854
- Hadzhiyska B. et al., 2023, *Open J. Astrophys.*, 6, 38
- Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
- Hikage C., Koyama K., Heavens A., 2017, *Phys. Rev. D*, 96, 043513
- Hikage C., Koyama K., Takahashi R., 2020, *Phys. Rev. D*, 101, 043510
- Hinton S. R. et al., 2017, *MNRAS*, 464, 4807
- Hinton S. R., Howlett C., Davis T. M., 2020, *MNRAS*, 493, 4078
- Hockney R. W., Eastwood J. W., 1988, *Computer Simulation using Particles*. Hilger, Bristol
- Howlett C., Lewis A., Hall A., Challinor A., 2012, *J. Cosmol. Astropart. Phys.*, 2012, 027
- Huterer D., 2023, *A Course in Cosmology*. Cambridge Univ. Press, Cambridge, UK
- Ivanov M. M., Sibiryakov S., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 053
- Jackson J. C., 1972, *MNRAS*, 156, 1P
- Jeong D., 2010, PhD thesis, University of Texas, Austin
- Kaiser N., 1987, *MNRAS*, 227, 1
- Karamanis M., Beutler F., Peacock J. A., 2021, *MNRAS*, 508, 3589
- Kirkby D. et al., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 024
- Kokron N., Chen S.-F., White M., DeRose J., Maus M., 2022, *J. Cosmol. Astropart. Phys.*, 2022, 059
- Lange J. U., 2023, *MNRAS*, 525, 3181
- Lazeyras T., Schmidt F., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 008
- Lemaître G., 1931, *MNRAS*, 91, 483
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Maksimova N. A., Garrison L. H., Eisenstein D. J., Hadzhiyska B., Bose S., Satterthwaite T. P., 2021, *MNRAS*, 508, 4017
- Matsubara T., 2008, *Phys. Rev. D*, 78, 083519
- Meiksin A., White M., Peacock J. A., 1999, *MNRAS*, 304, 851
- Mena-Fernández J. et al., 2024, preprint (arXiv:2404.03008)
- Mirbabayi M., Schmidt F., Zaldarriaga M., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 030
- Moon J. et al., 2023, *MNRAS*, 525, 5406
- Noda E., Peloso M., Pietroni M., 2020, *Phys. Dark Univ.*, 29, 100579
- Noh Y., White M., Padmanabhan N., 2009, *Phys. Rev. D*, 80, 123501
- Padmanabhan N., White M., 2008, *Phys. Rev. D*, 77, 123540
- Padmanabhan N., White M., 2009, *Phys. Rev. D*, 80, 063508
- Padmanabhan N., White M., Cohn J. D., 2009, *Phys. Rev. D*, 79, 063523
- Padmanabhan N., Xu X., Eisenstein D. J., Scalzo R., Cuesta A. J., Mehta K. T., Kazin E., 2012, *MNRAS*, 427, 2132
- Paillass E. et al., 2024, preprint (arXiv:2404.03005)
- Park C., Vogeley M. S., Geller M. J., Huchra J. P., 1994, *ApJ*, 431, 569
- Peacock J. A., 1992, *MNRAS*, 258, 581
- Peacock J. A., Dodds S. J., 1994, *MNRAS*, 267, 1020
- Peebles P. J. E., 1976, *Ap&SS*, 45, 3

Peebles P. J. E., Yu J. T., 1970, *ApJ*, 162, 815
 Percival W. J. et al., 2014, *MNRAS*, 439, 2531
 Percival W. J., Friedrich O., Sellentin E., Heavens A., 2022, *MNRAS*, 510, 3207
 Pérez-Fernández A. et al., 2024, preprint (arXiv:2406.06085)
 Philcox O. H. E., Ivanov M. M., Simonović M., Zaldarriaga M., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 032
 Planck Collaboration VI, 2020, *A&A*, 641, A6
 Raccanelli A., Vlah Z., 2023, *Phys. Rev. D*, 108, 043537
 Rashkovetskyi M. et al., 2024, preprint (arXiv:2404.03007)
 Reid B. et al., 2016, *MNRAS*, 455, 1553
 Reimberg P., Bernardeau F., Pitrou C., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 048
 Robertson H. P., 1935, *ApJ*, 82, 284
 Ross A. J. et al., 2017, *MNRAS*, 464, 1168
 Schmidt F., 2016, *Phys. Rev. D*, 94, 063508
 Schmittfull M., Feng Y., Beutler F., Sherwin B., Chu M. Y., 2015, *Phys. Rev. D*, 92, 123522
 Seljak U., McDonald P., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 039
 Senatore L., Trevisan G., 2018, *J. Cosmol. Astropart. Phys.*, 2018, 019
 Senatore L., Zaldarriaga M., 2015, *J. Cosmol. Astropart. Phys.*, 2015, 013
 Seo H.-J., Siegel E. R., Eisenstein D. J., White M., 2008, *ApJ*, 686, 13
 Seo H.-J., Beutler F., Ross A. J., Saito S., 2016, *MNRAS*, 460, 2453
 Sherwin B. D., Zaldarriaga M., 2012, *Phys. Rev. D*, 85, 103523
 Sheth R. K., 1996, *MNRAS*, 279, 1310
 Sheth R. K., Mo H. J., Tormen G., 2001, *MNRAS*, 323, 1
 Shirasaki M., Sugiyama N. S., Takahashi R., Kitaura F.-S., 2021, *Phys. Rev. D*, 103, 023506
 Slepian Z. et al., 2018, *MNRAS*, 474, 2109
 Smith R. E., Scoccimarro R., Sheth R. K., 2008, *Phys. Rev. D*, 77, 043525
 Speagle J. S., 2020, *MNRAS*, 493, 3132
 Sugiyama N., 2024, preprint (arXiv:2402.06142)
 Sugiyama N. S., Saito S., Beutler F., Seo H.-J., 2021, *MNRAS*, 501, 2862
 Sunyaev R. A., Zeldovich Y. B., 1970, *Ap&SS*, 7, 3
 Taylor A. N., Kitching T. D., 2010, *MNRAS*, 408, 865
 Tsaliakhovich D., Hirata C., 2010, *Phys. Rev. D*, 82, 083520
 Vargas-Magaña M. et al., 2014, *MNRAS*, 445, 2
 Vasudevan A., Ivanov M. M., Sibiriyakov S., Lesgourgues J., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 037
 Vlah Z., White M., 2019, *J. Cosmol. Astropart. Phys.*, 2019, 007
 Vlah Z., Seljak U., Yat Chu M., Feng Y., 2016, *J. Cosmol. Astropart. Phys.*, 2016, 057
 Walker A. G., 1937, *Proc. London Math. Soc.*, s2–42, 90
 Wallisch B., 2018, PhD thesis, University of Cambridge, UK
 White M., 2001, *MNRAS*, 321, 1
 White M., 2010, preprint (arXiv:1004.0250)
 White M., 2015, *MNRAS*, 450, 3822
 Xu X., Padmanabhan N., Eisenstein D. J., Mehta K. T., Cuesta A. J., 2012, *MNRAS*, 427, 2146
 Yoo J., Seljak U., 2013, *Phys. Rev. D*, 88, 103520
 Yoo J., Dalal N., Seljak U., 2011, *J. Cosmol. Astropart. Phys.*, 2011, 018
 Zunckel C., Okouma P., Muya Kasanda S., Moodley K., Bassett B. A., 2011, *Phys. Lett. B*, 696, 433

APPENDIX A: REDSHIFT EVOLUTION

The measured power spectrum (or correlation function) is a weighted average over the entire redshift bin, given by

$$\hat{P}(\mathbf{k}) = \sum_i w_i P(\mathbf{k}(z_i), z_i). \quad (\text{A1})$$

The corresponding wavenumber \mathbf{k} varies with redshift since, if the fiducial and true cosmologies do not exactly match, the Alcock-Paczynski (AP) parameters translating between distances therein will evolve with time (see equation 6). Within each redshift bin, the measured power spectrum is primarily sensitive to clustering at the effective redshift $z_{\text{eff}} = \sum_i w_i z_i$ at which linear trends in the time

evolution cancel, that is,

$$\hat{P}(\mathbf{k}) = P(\mathbf{k}(z_{\text{eff}})) + \frac{1}{2}((z - z_{\text{eff}})^2) \left(\frac{d^2 P}{dz^2} \right)_{z=z_{\text{eff}}} + \dots, \quad (\text{A2})$$

where angular brackets denote the weighted mean so that the leading correction is proportional to the variance of the redshift distribution σ_z^2 . This motivates the interpretation of the measured AP parameters from BAO as distances at z_{eff} with a theoretical error set by the width of the redshift bin.

Let us try to gain some heuristic understanding of the kind of errors induced in the AP parameters by the finiteness of σ_z . All else being equal we can write $P = P(\alpha_{\parallel}, \alpha_{\perp})$ such that for example

$$\frac{dP}{dz} = \frac{\partial P}{\partial \alpha_{\parallel}} \frac{d\alpha_{\parallel}}{dz} + \frac{\partial P}{\partial \alpha_{\perp}} \frac{d\alpha_{\perp}}{dz}.$$

The derivatives $\partial P / \partial \alpha_{\parallel, \perp}$ are precisely the templates for the AP parameters and, if the incorrect effective redshift were used and the linear redshift evolution not entirely cancelled, the measured distances would be off by $\Delta \alpha_{\parallel, \perp} = (d\alpha_{\parallel, \perp} / dz) \Delta z$. This would clearly be undesirable since $\Delta z \sim 0.1$ and, within even the baseline Λ CDM model we might expect $d\alpha_{\parallel, \perp} / dz \sim \Delta \Omega_m / \Omega_m \sim 0.1$, leading to percent-level systematic errors.

We can repeat the same analysis for the case where z_{eff} were correctly determined, such that the leading error is given by the second derivative. Let us further assume that, while the first derivatives lead to phase shifts in the BAO wiggles, the second derivatives lead to only in-phase corrections that do not bias distance measurements.¹² Within this approximation we have that

$$\frac{d^2 P}{dz^2} \sim \frac{\partial P}{\partial \alpha_{\parallel}} \frac{d^2 \alpha_{\parallel}}{dz^2} + \frac{\partial P}{\partial \alpha_{\perp}} \frac{d^2 \alpha_{\perp}}{dz^2}$$

which we can again directly translate into an error on the AP parameters as

$$\Delta \alpha_{\parallel, \perp} = \frac{1}{2} \left(\frac{d^2 \alpha_{\parallel, \perp}}{dz^2} \right)_{z=z_{\text{eff}}} \sigma_z^2. \quad (\text{A3})$$

For $\sigma_z \sim 0.1$ and for cosmologies wherein distances are within 10 per cent of the fiducial one this translates to errors of less than 0.1 per cent, but even in the case of larger redshift bins or models with more freedom, this shift can be exactly computed, since σ_z depends only on the redshift distribution and the second derivative can be computed given an cosmology, and included when performing cosmological inference.

APPENDIX B: WIDE-ANGLE EFFECTS

In the main body of the text, we have operated almost exclusively within the plane-parallel approximation where the LOS can be treated as a constant vector \hat{n} . This approximation is an excellent one for most of the scales on which the BAO is measured, with corrections scaling as negative powers $x = k\chi$ where χ is the comoving distance (Reimberg, Bernardeau & Pitrou 2016; Castorina & White 2018). These corrections to the plane parallel limit are due to (i) the physical LOS along which galaxies are redshifted being not parallel but radial, (ii) the radial evolution of the galaxy selection function $d(\chi^2 \bar{n}) / d \ln \chi$, and (iii) anisotropies incurred by the numerical implementation of the power spectrum estimator. The leading effect due to (iii) is a practical rather than physical one and

¹²Roughly speaking we can imagine $P_w \sim \cos(\alpha k r_d)$ such that the second derivative is again a cosine, now multiplied by k^2 .

can be absorbed into the window function (Castorina & White 2018; Beutler, Castorina & Zhang 2019b) – since we adopt this formalism in DESI we will focus on the first two effects (DESI Collaboration in preparation-f), whose leading order contributions are proportional to x^{-2} .

Since these wide-angle effects are on large scales we can evaluate them assuming the underlying dynamics are linear, in which case the first effect (i) gives (Castorina & White 2018; Beutler et al. 2019b)

$$\xi_0(s, \chi) \supset - \left(\frac{4f^2}{45} \xi_0^0 + \frac{f(9b_1 + f)}{45} \xi_0^2 \right) \left(\frac{s}{\chi} \right)^2$$

$$\xi_2(s, \chi) \supset \left(\frac{4f^2}{45} \xi_0^0 + \frac{f(189b_1 + 53f)}{441} \xi_0^2 - \frac{4f^2}{245} \xi_0^4 \right) \left(\frac{s}{\chi} \right)^2$$

while the selection function effect gives in addition

$$\xi_0(s, \chi) \supset \frac{4f^2}{3\chi^2} \xi_{-2}^0 + \frac{2f(b_1 - f)}{3\chi} \xi_{-1}^1 \left(\frac{s}{\chi} \right)^2$$

$$\xi_2(s, \chi) \supset - \frac{8f^2}{3\chi^2} \xi_{-2}^2 - \frac{8f(5b_1 + f)}{15\chi} \xi_{-1}^1 \left(\frac{s}{\chi} \right)$$

$$+ \frac{4f^2}{5\chi} \xi_{-1}^3 \left(\frac{s}{\chi} \right) \quad (\text{B1})$$

in the limit of constant number density \bar{n} .¹³ In the above, we have defined the generalized correlation functions

$$\xi_n^\ell = \int \frac{dk}{2\pi^2} k^n P(k) j_\ell(k s) \quad (\text{B2})$$

The contributions to the power spectrum multipoles are related to the correlation function expressions above via the usual Hankel transforms. For a realistic survey, these effects have to be averaged over LOS vectors $\vec{\chi}$, though for the exploratory purposes of this section, we will assume a fixed $\chi = \chi(z)$.

In order to estimate the bias these additional contributions cause in standard BAO fits, we perform a simple Fisher forecast as we did for non-linear effects in Section 3. These effects are largest at low redshift, where $k\chi$ is smallest, and when the galaxy bias is low. For this reason, let us consider the case of a sample-variance limited survey with $b_1 = 1$ at $z = 0.1$, that is, the lowest redshift boundary for the BGS sample. Even in this case, the shifts to both $\alpha_{\parallel, \perp}$ are uniformly less than 0.05 per cent using our fiducial setup with $k_{\min} = 0.02 h \text{ Mpc}^{-1}$. The biggest shifts are due to the selection function contribution since the odd Bessel functions are out of phase with the even ones. None the less, while the wide-angle effects can reach sizes of up to several percent of the total signal at k_{\min} their sizes are not amplified by the large BAO scale as is the non-linear phase shift derived in Section 3, and furthermore since they roughly decay with wavenumber as k^{-2} their overall shape is very different from a shift in the BAO scale $\propto k P'_w(k)$.

Let us comment briefly on some other large-scale effects on the observed galaxy two-point function not covered in the main text, such as unequal time and relativistic effects (see e.g. Grimm et al. 2020; Castorina & Di Dio 2022; Raccanelli & Vlah 2023 for recent examples). In general, these effects are suppressed on subhorizon scales like the wide-angle effects above and, moreover, do not have the exact shape of shifts in the BAO scale. For example, Raccanelli & Vlah (2023) recently showed that unequal time effects can contribute

¹³We take this limit in order to provide a rough estimate the size of the effect on BAO measurements, but note that for a spectroscopic survey like DESI the redshift dependence of \bar{n} can be measured to very high precision so the actual contributions could be calculated very precisely if desired.

a term proportional to $k P'(k)$; however, this term is suppressed on large scales and is proportional to an additional factor of $(H(z)/k)^2$ thus it contributes negligibly to the measured BAO scale. As we do not expect any of the large-scale effects on the galaxy power spectrum to significantly bias DESI BAO measurements, we leave their detailed analysis to future work.

APPENDIX C: THE NON-LINEAR BAO SHIFT IN REDSHIFT SPACE

In order to compute the non-linear BAO shift in redshift space in EPT, we need to compute the galaxy density to second-order

$$\delta_{g,s}^{(2)} = \int_{\mathbf{p}} Z_2(\mathbf{p}, \mathbf{k} - \mathbf{p}) \delta_{\text{lin}}(\mathbf{p}) \delta_{\text{lin}}(\mathbf{k} - \mathbf{p}),$$

where the second-order redshift-space kernel is given in Fourier space as (see e.g. Bernardeau et al. 2002)

$$Z_2(\mathbf{p}_1, \mathbf{p}_2) = \frac{1}{2} b_2 + b_s \left(\frac{(\mathbf{p}_1 \cdot \mathbf{p}_2)^2}{p_1^2 p_2^2} - \frac{1}{3} \right)$$

$$+ b_1 F_2(\mathbf{p}_1, \mathbf{p}_2) + f \mu^2 G_2(\mathbf{p}_1, \mathbf{p}_2)$$

$$+ \frac{f k \mu}{2} \left(\frac{\hat{n} \cdot \mathbf{p}_1}{p_1^2} (b_1 + f \mu_2^2) + \frac{\hat{n} \cdot \mathbf{p}_2}{p_2^2} (b_1 + f \mu_1^2) \right) \quad (\text{C1})$$

and we have defined $\mu_n = \hat{n} \cdot \hat{p}_n$ and $\hat{k} = \mathbf{p}_1 + \mathbf{p}_2$, with $\mu = \hat{n} \cdot \hat{k}$. We have also used the standard definitions of the matter density and velocity divergence kernels for F_2 and G_2 and will return to the cubic kernels Z_3 when discussing BAO damping below.

The largest out-of-phase contributions to the BAO in the galaxy power spectrum are sourced by the square of the quadratic contribution above (Crocce & Scoccimarro 2008; Padmanabhan & White 2009; Sherwin & Zaldarriaga 2012)

$$P_{g,s}^{(22)}(\mathbf{k}) = 2 \int_{\mathbf{p}} Z_2(\mathbf{p}, \mathbf{k} - \mathbf{p})^2 P_{\text{lin}}(\mathbf{p}) P_{\text{lin}}(\mathbf{k} - \mathbf{p}).$$

Specifically, this term contains contributions from modes longer than the BAO proportional to a shift, that is, logarithmic derivative, of the BAO template, for example,

$$\int_{|\mathbf{p}| \lesssim 1/r_d} \left(\frac{\mathbf{k} \cdot \mathbf{p}}{p^2} \right) P_{\text{lin}}(\mathbf{p}) \left(-\frac{\mathbf{k} \cdot \mathbf{p}}{k} \right) P'_w(k), \quad (\text{C2})$$

where the first factor comes from the displacement $\Psi \cdot \nabla \delta$ and the second one comes from expanding $P(|\mathbf{k} - \mathbf{q}|)$ for small \mathbf{q} . The leading terms will be those that scale as p^0 in the integrand like the one above. The bound on \mathbf{p} comes the fact that $P(|\mathbf{k} - \mathbf{q}|)$ begins to oscillate for larger \mathbf{q} . Note that there are in fact two long-wavelength regions in the 22 integral since both \mathbf{p} and $\mathbf{k} - \mathbf{p}$ can be small. Additionally, the inclusion of redshift-space distortions can introduce powers of the LOS angle μ into the above integral since displacements are amplified along \hat{n} .

In order to compute all such terms to one-loop order in the redshift-space galaxy power spectrum we use MATHEMATICA, setting up our calculation in the spherical coordinate system

$$\hat{k} = \{0, 0, 1\}, \quad \hat{n} = \{\sqrt{1 - \mu^2}, 0, \mu\}$$

$$\hat{p} = \{\sqrt{1 - \nu^2} \cos \phi, \sqrt{1 - \nu^2} \sin \phi, \nu\}, \quad (\text{C3})$$

such that the ϕ and ν dependence can be performed analytically in the small \mathbf{p} limit. Specifically, we Taylor-series expand Z_2^2 and $P_{\text{lin}}(\mathbf{k} - \mathbf{p})$ in $p = |\mathbf{p}|$, perform the angular integrals in ν, ϕ , and compute the derivative of this integral with respect to $P'_w(k)$, leading to equation (16).

APPENDIX D: MODELLING RECONSTRUCTION IN EULERIAN PERTURBATION THEORY

In Appendix C, we outlined the calculation of the non-linear BAO shift in the galaxy power spectrum. This shift is, however, known to be significantly reduced by reconstruction, which we will show analytically using PT in this appendix.

D1 Post-reconstruction BAO shift in real space

We begin by considering the case of matter in real space. The smoothed matter density yields displacements $\Psi_r^{(n)}(\mathbf{x}) = \nabla^{-1}(\mathcal{S} * \delta_m^{(n)}(\mathbf{x}))$ which can be used to compute post-reconstruction densities as

$$(2\pi)^3 \delta_D(\mathbf{k}) + \delta_a(\mathbf{k}) = \int d^3x e^{-i\mathbf{k}\cdot\mathbf{x} - i\mathbf{k}\cdot\Psi_r(\mathbf{x})} (1 + \delta_a(\mathbf{x})) \\ = \sum_{n=0}^{\infty} \frac{(-i)^n}{n!} \text{FT} [(\mathbf{k} \cdot \Psi_r)^n (1 + \delta_a)](\mathbf{k}). \quad (\text{D1})$$

Here, $\delta_a(\mathbf{x})$ is the pre-reconstruction density, for example, $\delta_d(\mathbf{x}) = \delta_m(\mathbf{x})$ and $\delta_s(\mathbf{x}) = 0$. We then have at second order (Hikage et al. 2017)

$$K_d^{(2)}(\mathbf{p}_1, \mathbf{p}_2) = (1 - \mathcal{S}_k) F_2(\mathbf{p}_1, \mathbf{p}_2) - \frac{1}{2} \left(\frac{\mathbf{k} \cdot \mathbf{p}_1}{p_1^2} \mathcal{S}_{p_1} + \frac{\mathbf{k} \cdot \mathbf{p}_2}{p_2^2} \mathcal{S}_{p_2} \right) \\ + \frac{1}{2} \left(\frac{\mathbf{k} \cdot \mathbf{p}_1}{p_1^2} \right) \left(\frac{\mathbf{k} \cdot \mathbf{p}_2}{p_2^2} \right) \mathcal{S}_{p_1} \mathcal{S}_{p_2} \\ K_s^{(2)}(\mathbf{p}_1, \mathbf{p}_2) = -\mathcal{S}_k F_2(\mathbf{p}_1, \mathbf{p}_2) + \frac{1}{2} \left(\frac{\mathbf{k} \cdot \mathbf{p}_1}{p_1^2} \right) \left(\frac{\mathbf{k} \cdot \mathbf{p}_2}{p_2^2} \right) \mathcal{S}_{p_1} \mathcal{S}_{p_2}, \quad (\text{D2})$$

where we have used the shorthand $\mathcal{S}_p = \mathcal{S}(p)$.

We can now proceed to compute the non-linear BAO shift as in the pre-reconstruction case with these kernels by considering the products $K_a^{(2)}(\mathbf{p}, \mathbf{k} - \mathbf{p}) K_b^{(2)}(\mathbf{p}, \mathbf{k} - \mathbf{p})$. If the smoothing scale of \mathcal{S} is significantly smaller than the BAO scale r_d then for the long-wavelength modes \mathbf{p} contributing to the shift we can use $\mathcal{S}_p \approx 1$ and $\mathcal{S}_{\mathbf{k}-\mathbf{p}} \approx \mathcal{S}_k$. For example, performing this procedure for the autospectrum of the shifted randoms yields

$$P_{ss} \supset k P'_w(k) \int_p \left(-\frac{1}{3} \mathcal{S}_p^2 \mathcal{S}_{k-p}^2 + \frac{82}{105} \mathcal{S}_p \mathcal{S}_{k-p} \mathcal{S}_k - \frac{47}{105} \mathcal{S}_k^2 \right) P_{\text{lin}}(\mathbf{p}) \\ \approx k P'_w(k) \int_p \left(-\frac{1}{3} \mathcal{S}_k^2 + \frac{82}{105} \mathcal{S}_k^2 - \frac{47}{105} \mathcal{S}_k^2 \right) P_{\text{lin}}(\mathbf{p}) = 0, \quad (\text{D3})$$

that is, a vanishing contribution to the BAO phase shift. Computing this contribution for the dd and ds spectra similarly yields that they vanish.

We can also repeat the above calculations for galaxies rather than matter. In this case, we have instead $\Psi_r^{(n)}(\mathbf{x}) = \nabla^{-1}(\mathcal{S} * \delta_g^{(n)}(\mathbf{x})/b_1)$ and $\delta_d(\mathbf{x}) = \delta_g(\mathbf{x})$, which we can nonetheless plug back into equation (D1) to derive $K_{d,s}^{(2)}$ for post-reconstruction galaxies. In this case, an equivalent calculation to the above yields again that the non-linear shift is exactly cancelled by the contributions introduced by reconstruction. This is true regardless of the details of the non-linear biasing since the non-linear shift is a product of the bias and the quadratic shift term $\Psi \cdot \nabla \delta$; the latter is always removed by reconstruction given that the correct b_1 is used. Indeed, the above calculation can be redone when

$$\Psi_r = \nabla^{-1} \left(\frac{\mathcal{S} * \delta_g(\mathbf{x})}{b_1^{\text{fid}}} \right) \quad (\text{D4})$$

allowing for $b_1^{\text{fid}} \neq b_1$. Doing so yields a shift

$$\Delta\alpha \propto \left(\frac{b_1^{\text{fid}}}{b_1} - 1 \right) \sigma_s^2 \quad (\text{D5})$$

where the proportionality constant is of the same order as the usual non-linear shift pre-reconstruction, that is, the residual shift in the BAO post-reconstruction is proportional to how well the ‘true’ matter overdensity can be estimated to produce the displacements. It is worth noting that the post-reconstruction shift left over due to incorrect assumptions will generally be in addition scale dependent and may not be entirely degenerate with BAO measurements, so that our estimates here should be taken to be conservative.

D2 Post-reconstruction BAO damping in real space

We can also extend the PT analysis above to look at the damping of the BAO after reconstruction. More specifically, we can use one-loop PT to examine what long-wavelength contributions get enhanced by the BAO such that they need to be resummed.

Before we do so it is useful to recall what happens pre-reconstruction. In this case we have for example in the matter power spectrum that (Bernardeau et al. 2002)

$$P_{mm}^{\text{loop}}(\mathbf{k}) = \int_p 2F_2(\mathbf{p}, \mathbf{k} - \mathbf{p})^2 P_{\text{lin}}(p) P_{\text{lin}}(|\mathbf{k} - \mathbf{p}|) \\ + 6F_3(\mathbf{k}, \mathbf{p}, -\mathbf{p}) P_{\text{lin}}(p) P_{\text{lin}}(k).$$

Both F_2^2 and F_3 have poles proportional to $1/p_n^2$ in the long-wavelength limit where one of their arguments $p_n \rightarrow 0$, such that their individual contributions will be quite sensitive to these IR contributions. However, if P_{lin} is a smooth function we can approximate $P_{\text{lin}}(|\mathbf{k} - \mathbf{p}|) \approx P_{\text{lin}}(k)$ to leading order, and with some algebra it is possible to show that the poles at $\mathbf{p}, \mathbf{k} - \mathbf{p} \rightarrow 0$ in the F_2^2 piece cancel the $\mathbf{p} \rightarrow 0$ pole in the F_3 piece. On the other hand, if the power spectrum contains a non-smooth component such as the BAO ($\sim \sin(r_d k)$) then the derivative of P_{lin} with respect to k is enhanced by the oscillation frequency r_d . In this case, the infrared contribution can be written as (Baldauf et al. 2015; Senatore & Trevisan 2018)

$$P_{mm}^{\text{loop}} \supset - \int_p^{p \ll k} \frac{(\mathbf{k} \cdot \mathbf{p})^2}{p^4} P_{\text{lin}}(p) (P_w(k) - P_w(|\mathbf{k} - \mathbf{p}|)) \\ = - \int d^3x e^{-i\mathbf{k}\cdot\mathbf{x}} \xi_w(\mathbf{x}) \int_{p, p \ll k} \frac{(\mathbf{k} \cdot \mathbf{p})^2}{p^4} P_{\text{lin}}(p) (1 - e^{i\mathbf{p}\cdot\mathbf{x}}) \\ \approx -\frac{1}{2} k_i k_j \int d^3x e^{-i\mathbf{k}\cdot\mathbf{x}} \xi_w(\mathbf{x}) [A_{ij}(\mathbf{x})]_{x=r_d} \\ \approx -\frac{1}{2} k^2 \Sigma_{\text{NL}}^2 P_w(k), \quad (\text{D6})$$

where ξ_w is the Fourier transform of the P_w and we have used that ξ_w is localized at r_d to perform a saddle point approximation between the second and third lines. This is none other than the one-loop contribution to the BAO damping $e^{-\frac{1}{2} k^2 \Sigma_{\text{NL}}^2 P_w}$ and uses the same saddle-point approximation in LPT described in Section 3, with the result that the BAO is damped by (numerically) large contributions from long-wavelength modes.

We can perform the same calculation post-reconstruction. For brevity, we will only consider the case of the shifted randoms in matter reconstruction in order to highlight why our result in Section 4 differs from the existing literature – the calculation for shifted galaxies, etc., is entirely analogous. In order to do so we need

to compute the third-order contribution in equation (D1)

$$K_s^{(3)}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3) = -S_k F_3(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3) + \frac{1}{3} \left(\frac{(\mathbf{k} \cdot \mathbf{p}_1)(\mathbf{k} \cdot \mathbf{p}_{23})}{p_1^2 p_{23}^2} S_1 S_{23} + \text{et. cycl.} \right) - \frac{1}{6} \left(\frac{(\mathbf{k} \cdot \mathbf{p}_1)(\mathbf{k} \cdot \mathbf{p}_2)(\mathbf{k} \cdot \mathbf{p}_3)}{p_1^2 p_2^2 p_3^2} \right) S_1 S_2 S_3,$$

where we have used numerical indices to indicate arguments by \mathbf{p}_n and used the shorthand $\mathbf{p}_{nm} = \mathbf{p}_n + \mathbf{p}_m$. Plugging this into the one-loop power spectrum then yields

$$P_{ss}^{\text{loop}} = \int_p 2K_s^{(2)}(\mathbf{p}, \mathbf{k} - \mathbf{p})^2 P_{\text{lin}}(p) P_{\text{lin}}(|\mathbf{k} - \mathbf{p}|) + 6K_s^{(3)}(\mathbf{k}, \mathbf{p}, -\mathbf{p}) P_{\text{lin}}(p) P_{\text{lin}}(k) \supset -S_k^2 \int_p^{p \ll k} \frac{(\mathbf{k} \cdot \mathbf{p})^2}{p^4} (1 - S_p)^2 P_{\text{lin}}(p) \times (P_w(k) - P_w(|\mathbf{k} - \mathbf{p}|)). \quad (\text{D7})$$

This is identical to the integral except that equation (1) there is a factor of S_k^2 in front, since P_{ss} is given by $S^2 P_{\text{lin}}$ at linear order and more (2) the integrand is suppressed by $(1 - S_p)^2$ internally. This is because, as described in Section 4, the shifted field is in fact displaced by bulk motions given to leading order by $(1 - S) * \Psi$ even though the randoms themselves are only moved by $-S * \Psi$. This is also reflected by the fact that the non-linear BAO shift derived above cancels independently in P_{ss} by virtue of the $\Psi \cdot \nabla \delta$ contribution cancelling there as well. Similar calculations show that the same result occurs in both $P_{dd,ds}$ as well, such that all of the pieces of the reconstructed power spectrum have the same BAO damping. The case for galaxies is also entirely similar.

D3 Non-linear BAO in redshift space

Finally, let us extend our treatment of the non-linear BAO post-reconstruction at one loop in EPT to include redshift-space distortions. The order n contribution to the reconstructed displacement is given by

$$\Psi_{r,s}^{(n)}(\mathbf{k}) = - \left(\frac{i\mathbf{k}}{k^2} \right) \left(\frac{\delta_{g,s}^{(n)}(\mathbf{k})}{b_1 + f\mu^2} \right). \quad (\text{D8})$$

In RecSym both the galaxies and randoms are shifted by this displacement multiplied by the redshift-space transformation matrix $R_{ij} = \delta_{ij} + f\hat{n}_i \hat{n}_j$. Let us again consider the case of matter such that $b_1 = 1$ for simplicity. In this case, the quadratic kernel for the shifted randoms is now

$$Z_s^{(2)}(\mathbf{p}_1, \mathbf{p}_2) = (-ik_i) R_{ij} \left(- \left(\frac{ik_j}{k^2} \right) \left(\frac{Z^{(2)}(\mathbf{p}_1, \mathbf{p}_2)}{1 + f\mu^2} \right) \right) S_k - \frac{1}{2} \left(k_i R_{ij} \left(\frac{-ip_{1,j}}{p_1^2} \frac{Z_1(\mathbf{p}_1)}{1 + f\mu_1^2} \right) \right) \times \left(k_i R_{ij} \left(\frac{-ip_{2,j}}{p_2^2} \frac{Z_1(\mathbf{p}_2)}{1 + f\mu_2^2} \right) \right) S_1 S_2, \quad (\text{D9})$$

where as usual $\mathbf{k} = \mathbf{p}_1 + \mathbf{p}_2$ and $\mu_n = \hat{n} \cdot \hat{p}_n$. We have included the dependence on Z_1 explicitly, but note that $Z_1(\mathbf{p}_n)/(1 + f\mu_n^2) = 1$. Repeating the real-space calculation using the coordinate system of equation (C3), and explicitly integrating over the angular dependence of the long mode \mathbf{p} , it is straightforward, if laborious, to show that the non-linear shift proportional to $\sigma_s^2 k P'_w(k)$ vanishes. In addition, we can also see that the arguments about the damping scale being set by $(1 - S) * \Psi$ also hold in redshift space since the relevant (resummed)

contribution comes from powers of the redshift-space displacement sourced by $S_{p_n} \delta^{(1)}(\mathbf{p}_n)/(b_1 + f\mu_n^2)$. These smooth displacements cancel the exponentiated Zeldovich displacements on large scales at all orders, such that their total contribution can be resummed as claimed in the main text. The case of displaced galaxies is entirely analogous.

Finally, let us repeat the calculation for RecIso. In RecIso, the particles are shifted by the displacement without multiplying by R_{ij} , and we can reproduce this in our calculations by taking $R_{ij} \rightarrow \delta_{ij}$ in the above. Performing the same calculation yields a BAO shift in RecIso of

$$\left(\frac{7f^3(4 - 9\mu^2)\mu^4 + f^2(14 - 73\mu^2)\mu^2 - 12f\mu^2}{105(1 + f\mu^2)} \right) \sigma_s^2 W^2(k) k P'_w(k)$$

representing a nonzero shift unlike RecSym.

APPENDIX E: HANKEL TRANSFORM OF BROAD-BAND SPLINE MODEL

The two relevant Hankel-transformed spline coefficients for the correlation function were evaluated to

$$B_{2,0}(x) = \frac{2}{x^6} \left(x^3 \text{Si}(x) - 2x^3 \text{Si}(2x) + x^2 \cos(x) - x^2 \cos(2x) - x \sin(x) - 16 \cos(x) + 4 \cos(2x) + x \sin(x) \cos(x) + 12 \right) B_{2,1}(x) = -\frac{1}{2x^3} \left(6x^3 \text{Si}(x) - 32x^3 \text{Si}(2x) + 27x^3 \text{Si}(3x) + 8x^2 + 6x^2 \cos(x) - 16x^2 \cos(2x) + 9x^2 \cos(3x) - 6x \sin(x) + 8x \sin(2x) - 3x \sin(3x) - 96 \cos(x) + 64 \cos(2x) - 16 \cos(3x) + 48 \right) \quad (\text{E1})$$

where $\text{Si}(x)$ is the sine integral $\int_0^x dt \sin(t)/t$.

APPENDIX F: ANALYTIC MARGINALIZATION OF THE BROAD BAND

Our full model for the BAO contains a large number of nuisance parameters designed to recover unbiased constraints on the two parameters we actually care about (α_{iso} and α_{ap}) in the presence of other physics that affects the shape of the galaxy clustering, survey selection effects and measurement noise.

In our base model for the galaxy clustering however, most of these nuisance parameters enter at linear order, and so can be marginalized over analytically at the level of the likelihood without needing to be sampled, expensively and numerically, as part of an MCMC or Nested sampling algorithm (Bridle et al. 2002; Taylor & Kitching 2010; d'Amico et al. 2020). This is useful as we are typically not interested in the constraints on these other parameters.

To demonstrate how this works, we rewrite the multipole version of our clustering model from equation (5) into two parts, where the first combines the wiggle and no-wiggle model components, and the latter is the broad band, making clear its dependence on the linear nuisance parameters $a_{\ell,n}$

$$P_{gg,\ell}(k) = \tilde{P}_{gg,\ell}(k) + \sum_n a_{\ell,n} \frac{\partial \mathcal{D}_\ell(k)}{\partial a_{\ell,n}}. \quad (\text{F1})$$

When performing a BAO fit, this model is then convolved with the survey window function (matrix $\mathbf{W}_{\ell\ell'}$) to give a new vector

$\mathbf{P}_\ell^W = \mathbf{W}_{\ell\ell'} \mathbf{P}_{gg,\ell'} = \tilde{\mathbf{P}}_\ell^W + \sum_n a_{\ell,n} \mathbf{D}_{\ell,n}^W$ where we have still kept the windowed model in two parts, those that do and do not depend on the linear order nuisance parameters, respectively. The matrix $\mathbf{D}_{\ell,n}^W$ then gives a set of purely linear templates in the power spectrum model. These templates can in general depend on the non-linear degrees of freedom (e.g. bias and dilation parameters) though they do not in our particular case.

Concatenating the different multipoles, we then take the common approach of using a Gaussian likelihood to describe our data ($\hat{\mathbf{P}}$) given the model with parameters θ ,

$$\mathcal{L}(\hat{\mathbf{P}}|\theta) \propto \text{Exp}\left[-\frac{1}{2}[\hat{\mathbf{P}} - \mathbf{P}^W]\mathbf{C}^{-1}[\hat{\mathbf{P}} - \mathbf{P}^W]^T\right]. \quad (\text{F2})$$

The act of marginalizing over the nuisance parameters, before or after sampling them, is akin to performing the integral

$$\mathcal{L}(\hat{\mathbf{P}}|\Omega) \propto \int d\mathbf{a} \mathcal{L}(\hat{\mathbf{P}}|\theta) P(\mathbf{a}), \quad (\text{F3})$$

where $P(\mathbf{a})$ are the prior probabilities for the various $a_{\ell,n}$ and Ω denotes all the other free parameters of the model. When using the Gaussian likelihood in equation (F2) and a flat (unbounded) prior for $P(\mathbf{a})$, this integral has an exact analytic solution,

$$\mathcal{L}(\hat{\mathbf{P}}|\Omega) \propto \frac{1}{|\mathbf{F}|^{1/2}} \text{Exp}\left[-\frac{1}{2}[\hat{\mathbf{P}} - \tilde{\mathbf{P}}^W]\tilde{\mathbf{C}}^{-1}[\hat{\mathbf{P}} - \tilde{\mathbf{P}}^W]^T\right], \quad (\text{F4})$$

where

$$\mathbf{F} = \mathbf{D}^W \mathbf{C}^{-1} \mathbf{D}^{W,T}, \quad (\text{F5})$$

is a square matrix with length equal to the number of nuisance parameters, and

$$\tilde{\mathbf{C}}^{-1} = \mathbf{C}^{-1} - \mathbf{C}^{-1} \mathbf{D}^{W,T} \mathbf{F}^{-1} \mathbf{D}^W \mathbf{C}^{-1}. \quad (\text{F6})$$

This has the interpretation that the full covariance is given by the data covariance plus a theory covariance, such that the total precision matrix is reduced. Similar analytic solutions can be obtained for the cases where $P(\mathbf{a})$ takes a Gaussian, in which cases one simply replaces

$$\mathbf{F} \rightarrow \mathbf{F} + \Sigma, \quad (\text{F7})$$

where Σ is the covariance matrix for the Gaussian prior on \mathbf{a} .

Given the above expressions, one can also identify the best-fitting values of the linear order nuisance parameters given the remainder of the model $\tilde{\mathbf{P}}^W$,¹¹

$$\mathbf{a}_{\text{best-fit}} = \mathbf{F}^{-1} \mathbf{D}^W \tilde{\mathbf{C}}^{-1} [\hat{\mathbf{P}} - \tilde{\mathbf{P}}^W]^T. \quad (\text{F8})$$

The improvement in χ^2 obtained by using these best-fitting parameters is equal to the reduction in the precision matrix above.

If the linear templates in \mathbf{D}^W are independent of the non-linear degrees of freedom, as is the case for the BAO broad band, an additional simplification occurs. In this case \mathbf{F} and $\tilde{\mathbf{C}}^{-1}$ can be pre-computed and so evaluation of the likelihood takes no additional time compared to the unmarginalized case. Combining with the above discussion this allows one to also take the ‘partially’ marginalized approach – evaluating the best-fitting broad-band parameters at each sampling iteration, multiplying by \mathbf{D}^W and adding them to the rest of the model to evaluate the standard χ^2 . This is the approach taken in some past literature (Ross et al. 2017; Bautista et al. 2021) and gives the same marginalized posterior on the parameters of interest in this special case. However, as can be seen from the description above, it requires more numerical operations per likelihood evaluation and so we argue the direct marginalization is preferable.

In either case, the benefits to sampling and run-time are substantial when including analytic marginalization. For our default model, when fitting the LRG mocks the number of free parameters in the sampling is reduced from 21 to 7. As such, we adopt full analytic marginalization over the broad-band parameters as our default.

Analytic marginalization can be carried out similarly for the correlation function, and further inspection shows that the galaxy bias (squared) could also be analytically marginalized over if desired. However, we do not include this latter option in our default methodology as tests on noisier data suggest that allowing the squared galaxy bias to go negative (as must be allowed in the case of a flat, unbounded prior) can give incorrect fits to the BAO in the correlation function, where the ‘dip’ in the correlation function on scales smaller than the BAO is mistaken for (a flipped version of) the BAO itself. The method of analytic marginalization does not lend itself to mixed priors, and so an alternative for further investigation would be to investigate Gaussian priors on the galaxy bias and broad-band terms.

¹Institute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA

²School of Mathematics and Physics, The University of Queensland, Brisbane, QLD 4072, Australia

³Department of Physics, University of California, Berkeley, 366 LeConte Hall MC 7300, Berkeley, CA 94720-7300, USA

⁴Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

⁵Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

⁶Department of Physics and Astronomy, Ohio University, Athens, OH 45701, USA

⁷Physics Department, Yale University, PO Box 208120, New Haven, CT 06511, USA

⁸Physics Department, Boston University, 590 Commonwealth Avenue, Boston, MA 02215, USA

⁹Department of Theoretical Physics, Tata Institute of Fundamental Research, Homi Bhabha Road, Mumbai 400005, India

¹⁰Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

¹¹Leinweber Center for Theoretical Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109-1040, USA

¹²NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA

¹³Department of Physics and Astronomy, University College London, Gower Street, London, WC1E 6BT, UK

¹⁴Institute for Computational Cosmology, Department of Physics, Durham University, South Road, Durham DH1 3LE, UK

¹⁵Department of Physics and Astronomy, The University of Utah, 115 South 1400 East, Salt Lake City, UT 84112, USA

¹⁶Instituto de Física, Universidad Nacional Autónoma de México, Cd. de México C.P. 04510, México

¹⁷Department of Astronomy, School of Physics and Astronomy, Shanghai Jiao Tong University, Shanghai 200240, China

¹⁸Berkeley Center for Cosmological Physics, Department of Physics, University of California, Berkeley, 110 Sproul Hall 5800 Berkeley, CA 94720, USA

¹⁹Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, E-08193 Bellaterra Barcelona, Spain

²⁰Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

²¹Departamento de Física, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio Ip, CP 111711, Bogotá, Colombia

²²Observatorio Astronómico, Universidad de los Andes, Cra. 1 No. 18A-10, Edificio H, CP 111711, Bogotá, Colombia

²³Center for Astrophysics | Harvard and Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA

²⁴Department of Physics, The University of Texas at Dallas, Richardson, TX 75080, USA

²⁵Institut d'Estudis Espacials de Catalunya (IEEC), E-08034 Barcelona, Spain

²⁶*Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciana Building, Portsmouth, PO1 3FX, UK*

²⁷*Institute of Space Sciences, ICE-CSIC, Campus UAB, Carrer de Can Magrans s/n, E-08913 Bellaterra, Barcelona, Spain*

²⁸*Department of Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA*

²⁹*The Ohio State University, Columbus, OH 43210, USA*

³⁰*Departament de Física, Serra Húnter, Universitat Autònoma de Barcelona, E-08193 Bellaterra (Barcelona), Spain*

³¹*Laboratoire de Physique Subatomique et de Cosmologie, 53 Avenue des Martyrs, F-38000 Grenoble, France*

³²*Institució Catalana de Recerca i Estudis Avançats, Passeig de Lluís Companys, 23, E-08010 Barcelona, Spain*

³³*Department of Physics and Astronomy, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*

³⁴*Department of Physics and Astronomy, Waterloo Centre for Astrophysics, University of Waterloo, 200 University Ave W, Waterloo, ON N2L 3G1, Canada*

³⁵*IRFU, CEA, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France*

³⁶*Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, ON N2L 2Y5, Canada*

³⁷*Max Planck Institute for Extraterrestrial Physics, Gießenbachstraße 1, D-85748 Garching, Germany*

³⁸*Instituto de Astrofísica de Andalucía (CSIC), Glorieta de la Astronomía, s/n, E-18008 Granada, Spain*

³⁹*Department of Physics, Kansas State University, 116 Cardwell Hall, Manhattan, KS 66506, USA*

⁴⁰*Department of Physics and Astronomy, Sejong University, Seoul 143-747, Korea*

⁴¹*Centre for Astrophysics and Supercomputing, Swinburne University of Technology, P.O. Box 218, Hawthorn, VIC 3122, Australia*

⁴²*CIEMAT, Avenida Complutense 40, E-28040 Madrid, Spain*

⁴³*SLAC National Accelerator Laboratory, Menlo Park, CA 94305, USA*

⁴⁴*National Astronomical Observatories, Chinese Academy of Sciences, A20 Datun Rd., Chaoyang District, Beijing, 100012, PR China*

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.