

Eigen-decomposition of covariance matrices: An application to the BAO linear point

Jaemyoung (Jason) Lee^{*}

*Department of Physics and Astronomy, University of Pennsylvania,
Philadelphia, Pennsylvania 19104, USA*

Farnik Nikakhtar

Department of Physics, Yale University, New Haven, Connecticut 06511, USA

Aseem Paranjape

Inter-University Centre for Astronomy and Astrophysics, Ganeshkhind, Pune 411007, India

Ravi K. Sheth

Center for Particle Cosmology, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA



(Received 10 July 2024; accepted 10 October 2024; published 12 November 2024)

The baryon acoustic oscillations (BAO) feature in the two-point correlation function (TPCF) of discrete tracers such as galaxies is an accurate standard ruler. The covariance matrix of the TPCF plays an important role in determining how the precision of this ruler depends on the number density and clustering strength of the tracers, as well as the survey volume. An eigen-decomposition of this matrix provides an objective way to separate the contributions of cosmic variance from those of shot-noise to the statistical uncertainties. For the signal-to-noise levels that are expected in ongoing and next-generation surveys, the cosmic variance eigen-modes dominate. These modes are smooth functions of scale, meaning that they are insensitive to the modest changes in binning that are allowed if one wishes to resolve the BAO feature in the TPCF; they provide a good description of the correlated residuals which result from fitting smooth functional forms to the measured TPCF; they motivate a simple but accurate approximation for the uncertainty on the linear point (LP) estimate of the BAO distance scale. This approximation allows one to quantify the precision of the BAO distance scale estimate without having to generate a large ensemble of mock catalogs and explains why: the uncertainty on the LP does not depend on the functional form fitted to the TPCF or the binning used; the LP is more constraining than the peak or dip scales in the TPCF; the evolved TPCF is less constraining than the initial one, so that reconstruction schemes can yield significant gains in precision.

DOI: [10.1103/PhysRevD.110.103515](https://doi.org/10.1103/PhysRevD.110.103515)

I. INTRODUCTION

Some of the tightest constraints on the cosmological distance-redshift relation come from the baryon acoustic oscillations (BAO) feature in the pair correlation function [1–5]. This has led to interest in the precision with which the two-point correlation function (TPCF) can be measured, and how this precision translates into uncertainties on the distance scale estimate. As a result, there is significant interest in understanding the covariance between pair counts on different scales.

On the ~ 140 Mpc scales of most interest to BAO cosmology, the Gauss-Poisson approximation to the covariance is rather accurate [6,7]. In this approximation, three different terms contribute: one is a purely cosmological

term, the other is a pure shot-noise term, and the third is a combination of the two. The first part of our paper is devoted to a study of the relative importance of these terms. We address this by rotating the covariance matrix into diagonal form, checking how many eigenvectors contribute significantly to the total covariance, and then looking at those eigenvectors. This provides a simple way of seeing which term dominates and when, as well as for understanding the shapes of the eigenvectors. We use this insight to explore how binning of the pair counts (width of the rectangular bins, or different bin shapes) affects the structure of the covariance matrix. Our methodology is similar in spirit to the Karhunen-Loève decomposition [8], in which a stochastic process is represented by orthonormal basis functions and uncorrelated random coefficients. The Karhunen-Loève decomposition was applied to cosmic microwave background (CMB) maps [9–11] as well as

^{*}Contact author: astjason@sas.upenn.edu

galaxy redshift surveys [12] in the 1990s for optimal data compression around the time of the first large-scale cosmological surveys. Reference [13] provides a generalization of this method in anticipation of subsequent generation surveys where the amount of data collected is too large for an uncompressed maximum likelihood analysis. They found that for both CMB maps and redshift-space distortions, a compression by a factor of ~ 10 is achievable by keeping the first $\sim 10\%$ of eigenmodes from the covariance matrix.

The second half of this paper applies these insights to a particular estimator of the BAO scale: the linear point (LP). The LP feature in the correlation function of dark matter or galaxies can be used as a standard cosmological ruler [2,14–21]. The LP lies midway between the peak and dip values in $\xi(r)$, the two-point correlation function:

$$r_{\text{LP}} \equiv \frac{r_{\text{peak}} + r_{\text{dip}}}{2}. \quad (1)$$

Evidently, the precision with which r_{LP} can be estimated from data depends on the covariance between the r_{peak} and r_{dip} estimates. In turn, this depends on the covariance matrix of the measurements, which depends on the widths of bins in which pairs were counted (or, more generally, the bin shapes themselves). This has led to significant computational efforts simply to determine the optimal bin width [22,23].

Moreover, r_{LP} is typically estimated by fitting a pre-determined functional form to the measured ξ (e.g. polynomials, Chebyshev polynomials, generalized half-integer Laguerre functions). The associated error bars would then appear to be closely tied to this functional form (e.g. Eqs. (2.6) and (9) in [16] and [18]). However, in practice, provided that the fits are good, neither the r_{LP} estimates nor their error bars depend strongly on which functional form is fit. Our analysis of the covariance matrix allows us to provide a rather general estimate of the expected precision that is not tied to a particular functional form. It also allows us to address a closely related question. In principle, the inflection point r_{infl} , the scale on which $d^2\xi/dr^2 = 0$, could also be used as a standard rod [14]. Previous work has suggested that it is less robust than the LP [23,24]; our analysis provides some insight into why this is so.

This paper is organized as follows. Section II describes how the eigenvalues and eigenvectors of the TPCF covariance matrix change as the shot-noise increases, and then uses this to provide a simple estimate of the error on the LP. Section III shows how our results depend on the binning. Section IV summarizes our conclusions.

Although we focus on the linear point scale, the scale r_0 where the pair correlation function crosses zero, $\xi(r_0) = 0$, has also been proposed for use as a standard ruler [25]. The Appendix shows the results of applying our analysis to r_0 .

II. METHODS AND RESULTS

Because the neighboring bins of the TPCF amplitudes are correlated, the covariance matrix of the bin counts is not diagonal. Here below, we describe in detail how we use the structure of the covariance matrix to estimate realistic error bars on the BAO distance scale.

Where necessary to illustrate our results, we will use a comoving volume of $5 \times (1.024)^3 h^{-3} \text{Gpc}^3$ in a flat ΛCDM model with $(\Omega_{m,0}, \Omega_{b,0}) = (0.281, 0.046)$, and $(h, n_s, \sigma_8) = (0.697, 0.971, 0.842)$ as in [19]. The associated linear theory values of r_{LP} and r_{infl} for the dark matter are 97.154 and 97.635 $h^{-1} \text{Mpc}$, respectively. For easy comparison with Refs. [18,19], we focus on biased tracers at $z = 0.5$ [also denoted as $z = 0.5057$ following [19]]. Although we explore other combinations of number density and clustering strength, our fiducial choice has $\bar{n} = 3.2 \times 10^{-4} / (h^{-1} \text{Mpc})^3$ and linear bias factor $b \approx 1.97$, which is similar to the Baryon Oscillation Spectroscopic Survey (BOSS) and the Dark Energy Spectroscopic Instrument (DESI) survey [26,27]. The combination $\bar{n}P(k_{\text{max}})$, where k_{max} is the scale on which $P(k) = b^2 P_{\text{Lin}}(k)$ is maximum, is sometimes used as a crude measure of whether the BAO clustering signal is dominated by shot-noise. Our fiducial choice has $\bar{n}P(k_{\text{max}}) \sim 5$; shot-noise dominates for values smaller than unity. While we provide our formalism in terms of the real-space correlation function, we use the redshift-space monopole in our figures to show that our methodology is valid even under redshift-space distortions. This makes $b_{\text{eff}} = \sqrt{(b^2 + 2bf/3 + f^2/5)}$ with $f = d \ln D / d \ln a$ being the linear theory growth rate [28]. With $b \sim 1.97$ at $z \sim 0.5$, $b_{\text{eff}} \sim 2.23$.

A. Gauss-Poisson approximation to covariance matrix

We begin with the two-point correlation function, which is related to the power spectrum $P(k)$ by

$$\xi(s) \equiv \int \frac{dk}{k} \frac{k^3 P(k)}{2\pi^2} j_0(ks). \quad (2)$$

A crude model for nonlinear evolution sets $P(k) = P_{\text{Lin}}(k)e^{-k^2 \Sigma^2}$ [e.g. [29]], so the nonlinear correlation function is a “smeared” version of the linear one. At $z = 0.5$ in our fiducial cosmology, $\Sigma \approx 4.7 h^{-1} \text{Mpc}$ for the real-space TPCF; it is slightly larger, $\Sigma_{\text{eff}} \approx 6.3 h^{-1} \text{Mpc}$, for the biased redshift-space monopole [24,30].

If the correlation function is estimated by counting pairs separated by $s \pm \Delta s/2$ in a discrete set of N_{tot} particles distributed in a volume V_s , then the TPCF covariance matrix described by the “Gauss-Poisson” approximation is given by [6,7,16,31]

$$C_{ij} = \frac{2}{V_s} \int_0^\infty \frac{dk k^2}{2\pi^2} \bar{J}_0(ks_i) \bar{J}_0(ks_j) \left[P(k) + \frac{1}{\bar{n}} \right]^2, \quad (3)$$

where $\bar{n} \equiv N_{\text{tot}}/V_s$ is the survey number density and \bar{j}_0 is the bin-averaged spherical Bessel function,

$$\bar{j}_0(ks_i) \equiv \frac{4\pi}{V_{s_i}} \int_{s_i-\Delta s/2}^{s_i+\Delta s/2} s^2 j_0(ks) ds, \quad (4)$$

with the volume $V_{s_i} = 4\pi/3(s_{i,\text{max}}^3 - s_{i,\text{min}}^3)$, s_i being the midpoint of bin i , and Δs being the bin size. The shot-noise only term proportional to $1/\bar{n}^2$ only contributes when $i = j$, i.e., to the error bar in a single bin. The other two terms describe the covariance between bins, and come from “cosmic variance.” This covariance must be accounted for when estimating the uncertainty on the BAO scale.

B. Eigen-decomposition of the covariance matrix

To illustrate our results, we now use Eq. (3) to generate C_{ij} , for 30 nonoverlapping bins of $\Delta r = 2 h^{-1}$ Mpc, running from $[60 - 120] h^{-1}$ Mpc, with the fiducial values of background cosmology, redshift, survey volume, biased tracer number density, and clustering strength mentioned at the start of this section.

Next, we diagonalize C_{ij} . The eigenvectors, which we denote $\Lambda_i(s)$, provide a set of orthogonal shape functions, whose relative importance is set by the eigenvalues λ_i . Before we consider the interplay of P and \bar{n} in determining these eigen-modes, note that the survey volume V_s only appears as an overall scaling. It scales the eigenvalues up and down but keeps their ratios fixed, and does not affect the eigen-shapes. That said, V_s is important because it does *not* enter in the definition of ξ itself. Therefore, larger V_s means that the eigen-modes will have smaller amplitudes compared to ξ . This will be important below.

Before we look at the shapes, Fig. 1 shows the fractional contribution of the eigenvalues λ_i (ordered from largest to smallest) to the total variance. The various curves show different choices for the relative contributions of “signal” $P(k)$ and “shot-noise” $1/\bar{n}$.

The lowermost curve is for the pure shot-noise limit (we have set $P(k) = 0$). This case is analytic: C_{ii} is diagonal, with entries $(2/\bar{n}V)/(\bar{n}V_i)$, where V_i is the volume of the i th bin. For bins of width $\Delta r = r_{i+1} - r_i$ that are equally spaced with spacing $\epsilon\Delta r$ (typically $\epsilon = 1$ but we will see later why the more general case is interesting) where r_i is the lower bound of the i th bin, $V_i = (4\pi/3)(\Delta r)^3 \times [(j\epsilon - (\epsilon - 1)/2)^3 - (j\epsilon - 1 - (\epsilon - 1)/2)^3]$ where $j = i + 60/(\epsilon\Delta r)$ and the bin centers are given by $\frac{(r_i+r_{i+1})}{2} = 60 + (i - 1)\epsilon\Delta r + \epsilon\Delta r/2$ since our bins start at $60 h^{-1}$ Mpc. Note that all eigenvectors matter. These eigenvectors are delta functions, one for each bin, centered on the middle of the bin.

In contrast, the uppermost curve shows the case in which $1/\bar{n} = 0$: here C_{ij} is completely determined by the P^2 term in Eq. (3). Notice that now the variance is dominated by just a few eigenvalues/eigenvectors. The intermediate curves




FIG. 1. Fractional contribution of the sorted eigenvalues λ_j to the total variance, $\sum_i \lambda_j / \sum_{i=1}^{30} \lambda_j$, for a range of choices of the ratio of “signal” ($P(k)$ amplitude) to “noise” ($1/\bar{n}$). Fiducial and idealized no shot-noise cases are very similar; more than 95% of the total variance comes from the first 3 modes. More modes matter when the noise dominates. The smooth dashed curve shows our analytic expression for the pure noise case: $(2/\bar{n}V)/(\bar{n}V_i)$. Dotted curve is for $b \rightarrow b - 1$ but fiducial shot-noise.

show different choices for the $1/\bar{n}$. The “fiducial” choice ($\bar{n}P(k_{\text{max}}) \approx 5$) is very similar to the no-noise case. However, as the noise increases, more modes begin to matter.¹

Figure 2 shows the corresponding eigenvectors. Notice that the n th eigenvector has $n - 1$ zero-crossings, at least for the first few n . It is striking that the mode with 4 zero-crossings divides the $60\text{--}120 h^{-1}$ Mpc range up into patches that are approximately the size of the BAO feature itself. Presumably this is because the same $P(k)$ appears in both ξ and C_{ij} . Dashed curves are for the case of no shot-noise, and solid curves have the fiducial shot-noise. Recall that these are the cases that are dominated by the first few eigenvalues, and the associated eigenvectors are very similar and very smooth. The smoothness is consistent with the expectation that terms contributing to cosmic variance should be smooth functions of scale. However, the similarity is particularly interesting here: it suggests that the eigenvectors for the no-noise case remain interesting even in the presence of fiducial noise.

The other higher-order eigenvectors, which contribute little to the total variance, are more strongly modified by the presence of shot-noise. To gain some insight, recall that the pure shot-noise eigenvectors would be a set of delta functions, each centered on a bin. But, when P is significant, these delta functions are now approximately rotated into the

¹The cross term in C_{ij} is also analytic: it is a smoothed version of ξ , with smoothing depending on scales i and j , but the expression is lengthy so we have not written it explicitly here.




FIG. 2. The first 9 eigenvectors $\sqrt{\lambda_i}\Lambda_i(s)$ at $z = 0.5057$, for default shot-noise $\bar{n} = 3.2 \times 10^{-4}/(h^{-1} \text{ Mpc})^3$ (solid) and no shot-noise (dashed). The first 3 modes (left hand panel), which contribute most to the total variance, are much less sensitive to shot-noise. Higher order modes (right hand panel, note the factor of 4 difference in the y-axis scale) are essentially zero when there is no shot-noise, and resemble delta-functions which peak at different scales, followed by oscillations, when shot-noise is present.

basis in which the P^2 term is diagonal. This mixes the delta functions, and is why these higher order modes display oscillations. We will exploit this relatively clean separation into cosmic variance vs shot-noise dominated eigen-modes in the next section.

C. Eigen-decomposition of correlation function realizations

We can write one realization Ξ of the real-space TPCF as

$$\Xi(s) = \xi(s) + \sum_i g_i \Lambda_i(s), \quad (5)$$

where $\xi(s)$ is given by Eq. (2), the g_i 's are independent Gaussian random variates with variance λ_i and mean zero, while $\Lambda_i(s)$ are the eigenvectors of C_{ij} . Hence, the terms other than $\xi(s)$ represent the (correlated) scatter around the mean.

The symbols in Fig. 3 show how the shape of Ξ changes as more modes are added to ξ , for one realization where we have assumed the fiducial noise and bias. The changes are relatively mild because V_s is sufficiently large that the λ_i are small. To highlight the differences as more modes are added, the symbols in Fig. 4 show the total residual from the mean, $\Xi(s) - \xi(s) = \sum_i g_i \Lambda_i(s)$, for this same realization. The other curves show the contribution from modes 1 to 4 ($\sum_{i=1}^4 g_i \Lambda_i(s)$), and from 5 onward ($\sum_{i=5}^{30} g_i \Lambda_i(s)$). Clearly, the first 4 modes capture most of the residual, including the small change in shape, while the sum of modes 5 and onward is mostly uncorrelated noise (small amplitude oscillations around zero). This just illustrates what Fig. 1 showed: the higher order modes are not particularly important.

D. Truncation of modes and basis functions for fitting the TPCF

Because we see that the sum of the first 4 modes captures most of the residual, while the remaining modes are mostly




FIG. 3. $\xi(s)$, $\Xi(s)$, and Eq. (5) with the first 2 and 4 modes. For a survey volume smaller than our fiducial V_s of $\sim 5(h^{-1} \text{ Gpc})^3$, the deviations of the other 3 curves from the blue $\xi(s)$ curve would be larger.




FIG. 4. Total residual $\Xi(s) - \xi(s)$ for one TPCF realization (symbols), and the contributions to it from modes 1 to 4 (dot-dashed), and from 5 onward (dashed).

“noise,” it is interesting to consider restricting the sum in Eq. (5) to include only the first ~ 4 modes. Evidently, this removes the shot-noise dominated fluctuations from the realization of Ξ , leaving a smoother curve. In essence, this is the smooth curve one is after when “fitting” the correlation. This is a nontrivial statement, since ξ + the first few eigenvectors, while smoother than the full Ξ , will not generally have the same shape as ξ . Nevertheless, since typical datasets were designed (i.e. V_s is large enough) so that cosmic variance does not dominate, the amplitude of these “cosmic variance” modes is small compared to the amplitude of ξ itself, so the correction to the shape is small (cf. Fig. 3).

In the same vein, suppose one is interested in derivatives of the correlation function. Although

$$\Xi' = \xi' + \sum_i g_i \Lambda'_i, \quad (6)$$

if we include all 30 terms in the sum, then this will be like differentiating a single measurement of the correlation function. However, it is well known that one should *not* differentiate a noisy measurement; rather, one should first fit a smooth functional form to the measurement and then differentiate the fit. In the present context, our model for this procedure is to assert that one is *not* interested in $d\Xi/ds$ when the sum includes all 30 terms; rather, one should only include the first 4 (really, the ones which account for, say, 90% of the variance).

The virtue of this point of view is that this shape is clearly determined *solely* by the shape of $P(k)$; in particular, it makes no reference to the set of basis functions which one wishes to fit to Ξ (simple polynomials? Laguerre functions? etc.). This is attractive, since a reasonable concern is whether the set of basis functions which worked for one underlying $P(k)$ will also work for another (polynomials for one, Laguerres for another?). Here, the point is that one should think of the eigenvectors Λ_i as being the most appropriate set of basis functions, since these are clearly determined by the shape of $P(k)$.

E. Model-independent error estimates on the BAO distance scales

We will now use this insight to discuss how one might quantify uncertainties on estimates of the BAO scales r_{LP} , r_{Dip} and r_{Peak} . We also consider r_{infl} , the scale on which the second derivative vanishes, as an alternative to r_{LP} .

We start with Eq. (6), but restrict the sum to the first few terms (the ones which account for, say, 90% of the variance). Next, note that the scale s_{max} where $\Xi' = 0$ is not necessarily the same as r_{max} , where $\xi' = 0$. Assuming $\Delta_{\text{max}} \equiv s_{\text{max}} - r_{\text{max}} \ll 1$, we have when $\Xi' = 0$,

$$\begin{aligned} 0 &= \xi'(s_{\text{max}}) + \sum_i g_i \Lambda'_i(s_{\text{max}}) \\ &\approx \xi'(r_{\text{max}}) + \Delta_{\text{max}} \xi''(r_{\text{max}}) + \sum_i g_i \Lambda'_i(s_{\text{max}}) \\ &\approx \Delta_{\text{max}} \xi''(r_{\text{max}}) + \sum_i g_i [\Lambda'_i(r_{\text{max}}) + \Delta_{\text{max}} \Lambda''_i(r_{\text{max}})], \end{aligned} \quad (7)$$

where we have used that $\xi' = 0$ at r_{max} . Isolating Δ_{max} yields

$$\frac{\Delta_{\text{max}}}{r_{\text{max}}} = - \frac{\sum_i g_i (\Lambda'_i(r_{\text{max}})/\xi'')}{1 + \sum_i g_i (\Lambda''_i(r_{\text{max}})/\xi'')}. \quad (8)$$

In practice, on the peak and dip scales, $(\Lambda'_i/\xi'') \leq 0.1$ or so. Since $\langle g_i^2 \rangle = \lambda_i$ is also small, we can neglect the term in the denominator and approximate

$$\left\langle \frac{\Delta_{\text{max}}^2}{r_{\text{max}}^2} \right\rangle \approx \sum_i \lambda_i \left(\frac{\Lambda'_i(r_{\text{max}})}{r_{\text{max}} \xi''(r_{\text{max}})} \right)^2. \quad (9)$$

Equation (9) shows that the root-mean-square (rms) of Δ_{max} increases as $|\xi''|$ decreases. Since the nonlinear TPCF is more smeared (i.e. less curved at the peak and dip scales) than the linear theory TPCF [recall discussion of Eq. (2)], we expect the uncertainty in the peak scale to be larger in the evolved field (at lower redshifts). We return to this point in the next subsection.

The same logic can be applied to the dip scale, so

$$\left\langle \frac{\Delta_{\text{min}}^2}{r_{\text{min}}^2} \right\rangle \approx \sum_i \lambda_i \left(\frac{\Lambda'_i(r_{\text{min}})}{r_{\text{min}} \xi''(r_{\text{min}})} \right)^2. \quad (10)$$

Hence, just as for Δ_{min} , the rms of Δ_{min} in the evolved TPCF is larger than the linear theory value.

Finally, for the error on the LP scale,

$$\Delta_{\text{LP}} \equiv \frac{\Delta_{\text{max}} + \Delta_{\text{min}}}{2}, \quad (11)$$

we have

$$\langle \Delta_{\text{LP}}^2 \rangle \approx \frac{\langle \Delta_{\text{max}}^2 \rangle}{4} + \frac{\langle \Delta_{\text{min}}^2 \rangle}{4} + \sum_i \frac{\lambda_i \Lambda'_i(r_{\text{max}}) \Lambda'_i(r_{\text{min}})}{2 \xi''(r_{\text{max}}) \xi''(r_{\text{min}})}. \quad (12)$$

Because $\xi''(r_{\text{max}})$ and $\xi''(r_{\text{min}})$ have opposite signs (by definition), the variance of Δ_{LP} is smaller than either Δ_{Dip} or Δ_{Peak} . This demonstrates why the LP is a more precise probe than either r_{Dip} and r_{Peak} .

Similarly for the inflection point, if r_{infl} is the scale where $\xi'' = 0$ and s_{infl} is the scale where $\Xi'' = 0$ (when the sum which defines Ξ is truncated to only include the eigenvalues whose eigenvectors contribute $\sim 90\%$ of the variance),

we would set

$$0 = \xi''(s_{\text{infl}}) + \sum_i g_i \Lambda_i''(s_{\text{infl}}) \approx \Delta_{\text{infl}} \xi'''(r_{\text{infl}}) + \sum_i g_i [\Lambda_i''(r_{\text{infl}}) + \Delta_{\text{infl}} \Lambda_i'''(r_{\text{infl}})], \quad (13)$$

where $\Delta_{\text{infl}} = s_{\text{infl}} - r_{\text{infl}}$. Thus,

$$\frac{\Delta_{\text{infl}}}{r_{\text{infl}}} = -\frac{\sum_i g_i (\Lambda_i''/r_{\text{infl}})/\xi'''}{1 + \sum_i g_i (\Lambda_i'''/\xi''')} \approx -\sum_i g_i \frac{\Lambda_i''}{r_{\text{infl}} \xi'''}, \quad (14)$$

so

$$\left\langle \frac{\Delta_{\text{infl}}^2}{r_{\text{infl}}^2} \right\rangle \approx \sum_i \lambda_i \left(\frac{\Lambda_i''(r_{\text{infl}})}{r_{\text{infl}} \xi'''(r_{\text{infl}})} \right)^2. \quad (15)$$

F. Comparison with standard method

To see how well this works, we first estimate the four scales (peak, dip, LP and inflection point) in the standard way [e.g. [14–16, 18, 19, 30]]: we made 100 mock realizations of the measurement [Eq. (5)], fitted a 7th order polynomial to each, and estimated the various scales by differentiating the fit. The rms scatter of each scale satisfies $\sigma_{\text{Dip}} > \sigma_{\text{Peak}} > \sigma_{\text{infl}} > \sigma_{\text{LP}}$: the LP is the most precise, followed by the inflection point, etc. This is consistent with previous work [e.g. Tables 4–7 of [19]]. The actual values are shown as the horizontal dashed lines in Fig. 5. Previous work has shown that it does not matter if one fits 7th order generalized Laguerre functions instead [18].

Now we turn to our eigen-mode-based estimates. Equations (10), (9), (14), and (12) show that these depend on the number of modes that are included in the relevant




FIG. 5. Fractional uncertainties on the dip, peak, inflection point and LP scales [square root of Eqs. (10), (9), (14), (12)], as a function of the number of eigen-modes that are included (symbols). Horizontal dashed lines show the corresponding uncertainties estimated in the standard way (see main text).

sums. The symbols in Fig. 5 show how these estimates increase as more modes are added. We previously argued that one should not include the higher order modes (because one should not differentiate noisy data); these are the ones for which the error estimate starts to diverge. The plateau at intermediate values indicates that the error estimate is not very sensitive to exactly how many modes are included, provided we have enough modes, and are not including the ones which are dominated by shot-noise. [This plateau is not an artifact produced by approximation Eq. (9); it is present even if we use Eq. (8).] In effect, this plateau provides an objective measure of how many modes should be included to accurately model the scatter between realizations, analogous to how Fig. 1 provides an objective way of deciding which modes are most important for a single realization. Indeed, previous work has noted that the inflection point is slightly less robust than the LP: here, this is indicated by the fact that it has a shorter plateau.

It is reassuring that not only do the plateau values reproduce the qualitative trends shown by the standard method, they are within 80% of the fitted rms values in all cases. Some of this discrepancy arises from the assumption that the scale which the standard method identifies as being the peak, say t_{max} (t for template), may differ slightly from s_{max} (our eigen-mode based estimate of the peak scale). As a result, the rms of $t_{\text{max}} - r_{\text{max}}$ differs from the rms of $s_{\text{max}} - r_{\text{max}}$; evidently, this difference is small.

Recall that the standard method results do not depend on the functional form that was fit to the binned TPCF. In effect, our analysis shows why: the low order eigenvectors Λ_i represent the covariance around any good fit to the measurements which is not “fitting the noise.”

We noted that, because the BAO feature in ξ becomes more smeared at late times, ξ'' decreases, so we expect the uncertainty on the LP scale to increase [cf. Eqs. (9)–(12)]. Figure 6 tests this: it shows the same comparison as in Fig. 5, but now with $P(k) \propto P_{\text{Lin}}(k)$ (i.e. $\Sigma = 0$, no smearing) for determining both $\xi_0(s)$ and C_{ij} when using Eq. (5) to produce 100 realizations of $\Xi(s)$. Setting $\Sigma \rightarrow 0$ changes the BAO feature in ξ_0 dramatically, and C_{ij} less so: e.g. the total variance is about 15% larger, and eigenvalues 3–7 contribute considerably more to the total variance compared to when $\Sigma \neq 0$. The change in ξ_0 means that a 7th order polynomial is no longer a good fit, so, for the “standard” analysis, we used a 9th order polynomial. Comparison of the dashed lines here with those in the previous figure shows that the rms is decreased by a factor of about 2 to 3. This is especially true for the peak which is most affected by the smearing.² Notice that this decrease is

²It may seem surprising that, in contrast to the evolved field, the peak and LP scales in the linear field are measured with similar precision. This is mainly a signal-to-noise issue. Recall that the original reason for working with the LP was not its precision, but its robustness to evolution/smearing [14].




FIG. 6. Same as Fig. 5, but with $\Sigma = 0$ (no smearing), for both $\xi_0(s)$ and the covariance matrix. Comparison with Fig. 5 (note the difference in y-axis range) shows the potential gains in precision which come from working with the reconstructed BAO signal.

reproduced by our eigen-mode estimates. The onset of the plateau is delayed from about mode 4 to about mode 7 or so, since now it is the first 7 modes which contribute to most of the variance. This decrease demonstrates the potential gains which come from using the reconstructed TPCF rather than the smeared one to measure the BAO scales: for the LP, this results in a precision of better than 0.4% as opposed to 0.8%. Our analysis has allowed us to estimate this improvement without having to run simulations.

Some reconstruction methods move the observed biased tracers back to (an estimate of) their initial positions [31–33]. The TPCF is then measured using these reconstructed positions. In this case, the number density is unchanged but $b \rightarrow b - 1$: typically, the BAO feature is sharper, but the amplitude is smaller [31]. Our methodology allows an estimate of the precision of the LP distance scale in this reconstructed signal as follows. Since $\bar{n}(b - 1)^2 P_{\text{Lin}}$ is smaller than $\bar{n}b^2 P_{\text{Lin}}$, the dotted curve in Fig. 1 suggests that more modes will be needed before we converge to a plateau as in Fig. 5. Although we do not show it here, we have checked that this is indeed the case. In addition, the precision of the reconstructed feature is slightly *less* precise, less constraining, than the original measurement. In other words, reconstruction yields no significant gain in precision (of course, it does reduce the bias in the mean value, bringing the LP closer to its linear theory value). To realize the potential for increased precision shown in Fig. 6, one must combine reconstructed fields, as discussed in [33].

We conclude that our methodology is an efficient way of determining accurate uncertainties on the linear point estimate of the BAO distance scale. In particular, since our analysis suggests that, for reasonable/realistic values of the shot-noise, the relevant eigenvalues and eigenvectors are entirely determined by the ratio of $P^2(k)$ and the survey




FIG. 7. Dependence of eigenvalues on bin size and spacing. Wider bins have smaller total variance, so we have normalized each eigenvalue to show its fractional contribution to the total. The first 4 normalized eigenvalues are very similar between all 3 curves, while the first 4 values for non-overlapping 6 h^{-1} Mpc bins are essentially the same as the overlapping 6 h^{-1} Mpc bins.

volume V_s , and they scale as $(b\sigma_8)^2/\sqrt{V_s}$. Our curves assumed $(b, \sigma_8, V_s) = (1.97, 0.84, 5 h^{-3} \text{Gpc}^3)$, so for other values, the fractional error is given by scaling the numbers in Fig. 5 by $(b_{\text{eff}}/2.23)^2(\sigma_8/0.84)^2\sqrt{5/V_s}$. This is not quite right, because it ignores the fact that the smearing also depends on σ_8 , but it is a useful first guess. Additionally, we note that the number of eigen-modes that captures a sufficient amount of information depends on $bP(k)$, V_s , as well as \bar{n} . There are two features in our analysis that show this number: (1) the change of the slope in Fig. 7 where we plot the eigenvalues by mode and (2) the plateaus in Figs. 5, 6, and 9 that show σ_{BAO} (which is what we really care about). Although we do not provide an explicit expression here, both features suggest that the first 5 eigen-modes (which contributes to $\sim 95\%$ of the total covariance) are sufficient for our choice of (b, V_s, \bar{n}) , which are reasonable values for current and next-generation surveys. While our fiducial choice of binning is 30 non-overlapping bins of 2 h^{-1} Mpc, our results are rather insensitive to binning, with 5 eigen-modes being sufficient for 15 bins of 4 h^{-1} Mpc, as well as 10 bins of 6 h^{-1} Mpc, the reason for which we will see in Sec. III.

III. DEPENDENCE OF EIGENVECTORS ON BINNING

The previous section considered the structure of the covariance matrix of a binned estimate of the TPCF. In that analysis, the original bin size (and shape) was fixed. How is the analysis modified if we change the binning?

In what follows, we first show that, in general, the eigenvectors of the binned covariance matrix are not simply binned versions of the original eigenvectors. Nevertheless, the first few eigenvectors are unchanged by the variations in

binning, permitted by the requirement that one be able to detect the BAO feature in the first place.

A. Analytic analysis

Let \mathbf{x} denote our list of bins, \mathbf{y} the list of measured bin amplitudes, $\mathbf{C} \equiv \langle \mathbf{y}\mathbf{y}^T \rangle$ the covariance matrix of the measurements, and λ_i and $\mathbf{\Lambda}_i$ its eigenvalues and eigenvectors. With some abuse of notation, let $\boldsymbol{\lambda}$ denote the diagonal matrix with the eigenvalues along the diagonal. If \mathbf{V} is a square matrix whose columns are the eigenvectors $\mathbf{\Lambda}$, then

$$\mathbf{C} = \mathbf{V}\boldsymbol{\lambda}\mathbf{V}^{-1} = \mathbf{V}\boldsymbol{\lambda}\mathbf{V}^T, \quad (16)$$

where the final equality follows because \mathbf{C} is real and symmetric.

Suppose we bin so that $\mathbf{y}_B \equiv \mathbf{B}\mathbf{y}$, with \mathbf{B} being a square matrix and each side having the same dimension as \mathbf{y} . For example,

$$\mathbf{B} = \frac{1}{3} \begin{pmatrix} 1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 1 \end{pmatrix} \quad (17)$$

would correspond to averaging the y values in the bins on either side of each x -bin. Note that, when this is done, one typically works with a sparser set of \mathbf{x} values, so as to not “double-count.” The analysis below is more transparent when \mathbf{x}_B and \mathbf{x} , and hence \mathbf{y}_B and \mathbf{y} , have the same length.

If \mathbf{C}_B denotes the covariance matrix of the binned measurements, then

$$\begin{aligned} \mathbf{C}_B &\equiv \langle \mathbf{B}\mathbf{y}(\mathbf{B}\mathbf{y})^T \rangle = \mathbf{C}\mathbf{B}\mathbf{B}^T \\ &= \mathbf{B}(\mathbf{V}\boldsymbol{\lambda}\mathbf{V}^{-1})\mathbf{B}^T = (\mathbf{B}\mathbf{V})\boldsymbol{\lambda}(\mathbf{V}^T\mathbf{B}^T) \\ &= (\mathbf{B}\mathbf{V})\boldsymbol{\lambda}(\mathbf{B}\mathbf{V})^T. \end{aligned} \quad (18)$$

If $(\mathbf{B}\mathbf{V})^T = (\mathbf{B}\mathbf{V})^{-1}$, the expression above would be the eigenvalue decomposition of \mathbf{C}_B , making it appear that the eigenvalues of \mathbf{C}_B are the same as those of \mathbf{C} , and the eigenvectors are simply those of \mathbf{C} , binned using \mathbf{B} . At face value, this is sensible: if the eigenvectors were smooth on scales smaller than the ‘bin width’ then they will be unchanged by—essentially invariant to—the binning.

However, notice that

$$\mathbf{B}\mathbf{V}(\mathbf{B}\mathbf{V})^T = \mathbf{B}(\mathbf{V}\mathbf{V}^T)\mathbf{B}^T = \mathbf{B}\mathbf{B}^T \quad (19)$$

is not diagonal (i.e., although \mathbf{V} is an orthonormal basis, $\mathbf{V}^T = \mathbf{V}^{-1}$, the same is not true for $\mathbf{B}\mathbf{V}$). This means that we should *not* think of $\mathbf{B}\mathbf{V}$ as being the eigenvectors. If we

use \mathbf{V}_B to denote the eigenvectors of \mathbf{C}_B ,

$$\mathbf{C}_B = \mathbf{V}_B\boldsymbol{\lambda}_B\mathbf{V}_B^{-1}, \quad (20)$$

then it is natural to ask: How different are the vectors which make up \mathbf{V}_B from those of $\mathbf{B}\mathbf{V}$?

B. Numerical analysis

Heuristically, we expect that if the binning remains smaller than the typical size of features in the eigenvectors, then they will be unchanged by binning. This should be particularly true for the primary “cosmic variance” dominated eigenvectors; the shot-noise dominated eigenvectors oscillate more, but we argued that they are not interesting anyway. Therefore, we expect the estimates of the BAO distance scale and their uncertainties should not depend on how the TPCF was binned, provided this binning is not wider than the BAO feature itself. (If the bins are too wide, they will not provide a good description of the BAO feature anyway.)

Figures 7 and 8 show the result of two explicit tests. The first increases the bin size, but keeps the bin centers and hence the number of bins the same. (As a result, neighboring bins are more correlated, but this just means that \mathbf{C}_B , which has the same dimension as in the previous section, is less diagonal than the original \mathbf{C} .) This corresponds to the $\epsilon \neq 1$ case mentioned previously (Sec. II B). Figure 7 shows the eigenvalues when we increase Δr by factors of 2 and 3 ($\epsilon = 1/2$ and $\epsilon = 1/3$), respectively. The first few eigenvalues, which dominate the total variance, are indistinguishable from the original ones, but the more shot-noise dominated modes are affected. In particular, for wider bins the shot-noise is smaller, so these shot-noise dominated modes contribute less to the total variance. To remove the fact that the total variance is reduced, we normalized each set of eigenvalues by their total: this shows explicitly that, as the bin size increases, the shot-noise dominated modes contribute a smaller fraction of the total variance.

The larger symbols show the eigenvalues when $\Delta r \rightarrow 3\Delta r$ but the bins do not overlap ($\epsilon = 1$). In this case, there are $3\times$ fewer eigenvalues, so the total variance is obviously different. Nevertheless, the fractional contribution of the first 10 modes to the variance is similar to that for the overlapping bins (of the same width). Clearly, for the cosmic-variance dominated modes which contribute most to the total variance, the binning does not matter.

Figure 8 shows that this is also true for the eigenvectors. The big symbols show the case in which $\Delta r \rightarrow 3\Delta r$ but now the bins do not overlap (so $\epsilon = 1$ even for this larger bin size; for clarity, we do not show the intermediate case where $\Delta r \rightarrow 2\Delta r$). Again, the eigenvectors which dominate the total variance (the first ~ 5) are indistinguishable from the original ones. This is slightly nontrivial since now \mathbf{C}_B is 10×10 rather than 30×30 , but the leading eigenvectors are unchanged. Hence, the uncertainties on distance




FIG. 8. First 9 eigenvectors $\sqrt{\lambda_i}\Lambda_i(s)$ at $z = 0.5057$, for bin sizes $2 h^{-1}$ Mpc and $6 h^{-1}$ Mpc with the same bin centers. Larger symbols show the eigenvectors for the $6 h^{-1}$ Mpc bins that do not overlap (so there are $3\times$ fewer bin centers, different from the overlapping bins).

scale estimates provided in the previous section will be unchanged: they do *not* depend on the binning, at least for fiducial values of the shot-noise.

More generally, if the convolution kernel which defines the binning does not erase features in the original (cosmic variance dominated) eigenvectors, then these eigenvectors will not depend on the exact bin shape. For example, this will certainly be true if the off-diagonal entries in Eq. (17) are less than unity. Similarly, counting pairs in, e.g., Gaussian-like bins rather than in rectangles will not change our conclusions.

IV. DISCUSSION AND CONCLUSIONS

We presented an eigen-decomposition of the Gauss-Poisson approximation to the covariance matrix of the two-point correlation function [Eq. (3)] and assessed the importance of the power spectrum-dominated modes that trace cosmic variance as opposed to the modes which are dominated by shot-noise. For a fiducial cosmology and noise-levels that are consistent with current and next-generation surveys, the cosmic variance eigen-modes account for most of the total variance of the TPCF (Fig. 1). They are also smoother than the shot-noise dominated modes (Fig. 2), so they are insensitive to the modest changes in binning that are allowed if one wishes to resolve the BAO feature in the TPCF (Figs. 7 and 8).

We argued that, as a result, the cosmic variance eigen-modes alone should provide a good description of the correlated residuals which result from fitting smooth functional forms to the measured TPCF. We provided a simple [Eq. (12)] but accurate (Fig. 5) approximation for the uncertainty on the linear point estimate of the BAO distance scale which explains why the uncertainty is greater in the evolved field than in linear theory (Fig. 6), allows one to quantify the gains from working with the reconstructed signal, and does not depend on the functional form fitted to the TPCF or the binning used. It also provides insight into why the LP is more robust than the inflection point, and why both are more precise distance indicators than the peak

or dip scales. Perhaps most importantly, our approximation allows one to quantify the precision of the BAO distance scale estimate without having to generate a large ensemble of mock catalogs. Therefore, it should be useful for estimating the gains in precision which come from making measurements in the reconstructed field (which are often quoted), after marginalizing over the unknown cosmological model (a step which is often ignored).

Our approach also provides a realistic estimate of the precision of the zero-crossing scale (Fig. 9), showing that the uncertainty on it depends on how steeply ξ crosses zero. Although it is considerably less precise than either r_{LP} or r_{infl} for realistic galaxy surveys, our analysis vastly simplifies the process of quantifying the synergy between joint analyses of r_{LP} and r_0 in the same dataset.

In practice, our analysis exploits the appearance of extended plateaus in plots of how the precision in the distance scale indicator varies as more (sorted) eigenmodes are included (cf. Figs. 5, 6 and 9), since it is these plateau heights that correspond to the uncertainty estimates that are returned from fitting basis functions to the measurements. (If a plateau is not seen, then this is an indication that the contributions from the cosmic variance and shot-noise terms to the full uncertainty are not so easily separated, and this makes it harder to pinpoint the correspondence between our analysis and the standard one.) These plateaus suggest that it should be possible to write down a prescription for determining the optimal number of eigen-modes which should be used in cosmological analyses.

While we do not perform cosmological analysis to infer cosmological parameters or evaluate cosmological models in this work, our approach can be used to rapidly measure the BAO scale and its uncertainties accurately given a dataset. This will enhance the utility of the LP that previous works have already shown: (1) the distance scale obtained from the LP is consistent with the standard template-fitting method, while being more model agnostic [2] and (2) the resulting cosmological parameters inferred from the LP are consistent with the standard approach with 20 to 30%

larger uncertainties under a flat Λ CDM cosmology [21]. Furthermore, the rapid estimation of the BAO scale and its uncertainties means that performing multiple BAO analyses assuming a variety of different cosmological models will be possible at reasonable computational cost, which will be especially important in the light of the potential time-evolving dark energy results from [5,34]. In future work, we hope to combine our eigen-mode analysis with the Bayesian framework of [20] utilizing the Laguerre reconstruction methodology of [18,24] for a completely model-agnostic BAO analysis.

ACKNOWLEDGMENTS

J.L. was supported by DOE Grant No. DE-FOA-0002424 and NSF Grant No. AST-2108094. F.N. gratefully acknowledges support from the Yale Center for Astronomy and Astrophysics Prize Postdoctoral Fellowship. A.P. and R.K.S. are grateful to the ICTP, Trieste for its hospitality in summer 2024, and R.K.S. is grateful to the EAFIR, Kigali for its hospitality when this work was completed. The research of A.P. is supported by the Associates Scheme of ICTP, Trieste.

APPENDIX: APPLICATION TO THE ZERO-CROSSING SCALE

With some care, the zero-crossing scale, r_0 , can also be used as a standard ruler [25,35]. In practice, estimating r_0 is similar to the linear point scale: one fits the pair counts with suitably chosen basis functions before estimating the zero-crossing from the fit. Therefore, the analysis in the main text can be modified to provide estimates of the uncertainty on r_0 as follows.

If s_0 is the scale on which $\Xi(s) = 0$, and r_0 is where $\xi(r) = 0$, we have

$$0 = \xi(r_0) + \Delta_0 d\xi/dr + \sum_i g_i [\Lambda_i(r_0) + \Delta_0 d\Lambda/dr], \quad (\text{A1})$$

where $\Delta_0 = s_0 - r_0$. Since $\xi(r_0) = 0$, dividing throughout by $\xi'(r_0)$ and then assuming $\Lambda'(r_0)/\xi'(r_0) \ll 1$, yields

$$\frac{\Delta_0}{r_0} \approx -\sum_i g_i \frac{\Lambda_i(r_0)}{r_0 \xi'(r_0)}, \quad (\text{A2})$$

so

$$\left\langle \frac{\Delta_0^2}{r_0^2} \right\rangle \approx \sum_i \lambda_i \left[\frac{\Lambda_i(r_0)}{r_0 \xi'(r_0)} \right]^2. \quad (\text{A3})$$




FIG. 9. Same as Fig. 5, but now showing the fractional uncertainties on the zero-crossing scale (gray) as well.

Notice that the uncertainty is smaller if ξ' is bigger. This is sensible: the zero-crossing of a curve is easier to detect if the curve is steep. In practice, this means that more biased tracers are to be preferred for this measurement.

Figure 9 shows how well this compares to the standard approach (of fitting a polynomial to the pair counts over the range $100 - 160 h^{-1} \text{Mpc}$ in bins of width $\Delta r = 2 h^{-1} \text{Mpc}$) as the number of modes that is included in the sum increases. The figure is in the same format as Fig. 5 in the main text, to highlight the similarity of the analysis. The height of the plateau region defined by the gray symbols is our eigen-mode estimate of the uncertainty on r_0 . It tracks that from the standard (polynomial fitting) method quite well (grey dashed), showing that it provides reliable estimates for analyses of the zero-crossing scale.

For r_0 , there are two striking differences compared to our analyses of the scales (r_{LP} etc.) that we highlighted in the main text. First, the plateau region—which the main text argued is a proxy for the uncertainty on the distance scale—is reached after only the first couple of eigen-modes. For r_{LP} etc., more eigen-modes are needed before convergence to a plateau is seen, and then the plateau lasts for a smaller range of eigen-modes before the variance increases again (mainly due to shot-noise dominated modes). Second, the uncertainty on r_0 is much larger than on r_{LP} etc. This quantifies results which are obvious by eye in Ref. [35]. In this respect, r_{LP} is the better distance scale indicator than r_0 . Nevertheless, there are synergies associated with having two distance scale estimates from the same measurement which we will explore elsewhere.

- [1] D. J. Eisenstein, I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol, R. Scranton, H.-J. Seo, M. Tegmark, Z. Zheng *et al.*, *Astrophys. J.* **633**, 560 (2005).
- [2] S. Anselmi, G. D. Starkman, P.-S. Corasaniti, R. K. Sheth, and I. Zehavi, *Phys. Rev. Lett.* **121**, 021302 (2018).
- [3] S. Alam, M. Aubert, S. Avila, C. Balland, J. E. Bautista, M. A. Bershad, D. Bizyaev, M. R. Blanton, A. S. Bolton, J. Bovy *et al.*, *Phys. Rev. D* **103**, 083533 (2021).
- [4] T. Abbott, M. Adamow, M. Agüena, S. Allam, O. Alves, A. Amon, F. Andrade-Oliveira, J. Asorey, S. Avila, D. Bacon *et al.*, *Phys. Rev. D* **110**, 063515 (2024).
- [5] A. Adame, J. Aguilar, S. Ahlen, S. Alam, D. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.*, [arXiv:2404.03002](https://arxiv.org/abs/2404.03002).
- [6] R. E. Smith, R. Scoccimarro, and R. K. Sheth, *Phys. Rev. D* **77**, 043525 (2008).
- [7] J. N. Grieb, A. G. Sánchez, S. Salazar-Albornoz, and C. Dalla Vecchia, *Mon. Not. R. Astron. Soc.* **457**, 1577 (2016).
- [8] K. Karhunen, *Ann. Acad. Sci. Fenn. Ser. AI, Mat.* **37**, 1 (1947).
- [9] J. R. Bond, *Phys. Rev. Lett.* **74**, 4369 (1995).
- [10] E. F. Bunn and N. Sugiyama, *Astrophys. J.* **446**, 49 (1995).
- [11] E. F. Bunn, [arXiv:astro-ph/9607088](https://arxiv.org/abs/astro-ph/9607088).
- [12] M. S. Vogeley and A. S. Szalay, *Astrophys. J.* **465**, 34 (1996).
- [13] M. Tegmark, A. N. Taylor, and A. F. Heavens, *Astrophys. J.* **480**, 22 (1997).
- [14] S. Anselmi, G. D. Starkman, and R. K. Sheth, *Mon. Not. R. Astron. Soc.* **455**, 2474 (2016).
- [15] S. Anselmi, P.-S. Corasaniti, A. G. Sanchez, G. D. Starkman, R. K. Sheth, and I. Zehavi, *Phys. Rev. D* **99**, 123515 (2019).
- [16] G. Parimbelli, S. Anselmi, M. Viel, C. Carbone, F. Villaescusa-Navarro, P. Corasaniti, Y. Rasera, R. Sheth, G. Starkman, and I. Zehavi, *J. Cosmol. Astropart. Phys.* (2021) 009.
- [17] M. O'Dwyer, S. Anselmi, G. D. Starkman, P.-S. Corasaniti, R. K. Sheth, and I. Zehavi, *Phys. Rev. D* **101**, 083517 (2020).
- [18] F. Nikakhtar, R. K. Sheth, and I. Zehavi, *Phys. Rev. D* **104**, 043530 (2021).
- [19] J. J. Lee, B. Fiorini, F. Nikakhtar, and R. K. Sheth, [arXiv:2406.09379](https://arxiv.org/abs/2406.09379).
- [20] A. Paranjape and R. K. Sheth, *Mon. Not. R. Astron. Soc.* **517**, 4696 (2022).
- [21] M. He, C. Zhao, and H. Shan, *Mon. Not. R. Astron. Soc.* **525**, 1746 (2023).
- [22] M. Vargas-Magaña, S. Ho, X. Xu, A. G. Sánchez, R. O'Connell, D. J. Eisenstein, A. J. Cuesta, W. J. Percival, A. J. Ross, E. Aubourg *et al.*, *Mon. Not. R. Astron. Soc.* **445**, 2 (2014).
- [23] S. Anselmi, P.-S. Corasaniti, G. D. Starkman, R. K. Sheth, and I. Zehavi, *Phys. Rev. D* **98**, 023527 (2018).
- [24] F. Nikakhtar, R. K. Sheth, and I. Zehavi, *Phys. Rev. D* **104**, 063504 (2021).
- [25] F. Prada, A. Klypin, G. Yepes, S. E. Nuza, and S. Gottloeber, [arXiv:1111.2889](https://arxiv.org/abs/1111.2889).
- [26] S. Alam, F. D. Albareti, C. A. Prieto, F. Anders, S. F. Anderson, T. Anderton, B. H. Andrews, E. Armengaud, É. Aubourg, S. Bailey *et al.*, *Astrophys. J. Suppl. Ser.* **219**, 12 (2015).
- [27] A. Adame, J. Aguilar, S. Ahlen, S. Alam, D. Alexander, M. Alvarez, O. Alves, A. Anand, U. Andrade, E. Armengaud *et al.*, [arXiv:2404.03000](https://arxiv.org/abs/2404.03000).
- [28] N. Kaiser, *Mon. Not. R. Astron. Soc.* **227**, 1 (1987).
- [29] M. Crocce and R. Scoccimarro, *Phys. Rev. D* **77**, 023533 (2008).
- [30] A. Paranjape and R. K. Sheth, *Mon. Not. R. Astron. Soc.* **526**, 700 (2023).
- [31] F. Nikakhtar, R. K. Sheth, B. Lévy, and R. Mohayaee, *Phys. Rev. Lett.* **129**, 251101 (2022).
- [32] F. Nikakhtar, N. Padmanabhan, B. Lévy, R. K. Sheth, and R. Mohayaee, *Phys. Rev. D* **108**, 083534 (2023).
- [33] F. Nikakhtar, R. K. Sheth, N. Padmanabhan, B. Lévy, and R. Mohayaee, *Phys. Rev. D* **109**, 123512 (2024).
- [34] T. Abbott, M. Acevedo, M. Agüena, A. Alarcon, S. Allam, O. Alves, A. Amon, F. Andrade-Oliveira, J. Annis, P. Armstrong *et al.*, *Astrophys. J. Lett.* **973**, L14 (2024).
- [35] S. Gaines, F. Nikakhtar, N. Padmanabhan, and R. K. Sheth, [arXiv:2408.00072](https://arxiv.org/abs/2408.00072).