UC Berkeley UC Berkeley Previously Published Works

Title

Precision redshift-space galaxy power spectra using Zel'dovich control variates

Permalink

https://escholarship.org/uc/item/2fm1x4p7

Journal

Journal of Cosmology and Astroparticle Physics, 2023(02)

ISSN

1475-7516

Authors

DeRose, Joseph Chen, Shi-Fan Kokron, Nickolas <u>et al.</u>

Publication Date

2023-02-01

DOI

10.1088/1475-7516/2023/02/008

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <u>https://creativecommons.org/licenses/by/4.0/</u>

Peer reviewed

Precision Redshift-Space Galaxy Power Spectra using Zel'dovich Control Variates

Joseph DeRose^{*a*} Shi-Fan Chen^{*b*} Nickolas Kokron^{*c,d*} Martin White^{*a,e,f*}

^aPhysics Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

 $^c\mathrm{Department}$ of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA

^dKavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA

^eBerkeley Center for Cosmological Physics, Department of Physics, UC Berkeley, CA 94720, USA

^fDepartment of Physics, University of California, Berkeley, CA 94720

E-mail: jderose@lbl.gov, sfschen@ias.edu, kokron@stanford.edu, mwhite@berkeley.edu

Abstract. Numerical simulations in cosmology require trade-offs between volume, resolution and run-time that limit the volume of the Universe that can be simulated, leading to sample variance in predictions of ensemble-average quantities such as the power spectrum or correlation function(s). Sample variance is particularly acute at large scales, which is also where analytic techniques can be highly reliable. This provides an opportunity to combine analytic and numerical techniques in a principled way to improve the dynamic range and reliability of predictions for clustering statistics. In this paper we extend the technique of Zel'dovich control variates, previously demonstrated for 2-point functions in real space, to reduce the sample variance in measurements of 2-point statistics of biased tracers in redshift space. We demonstrate that with this technique, we can reduce the sample variance of these statistics down to their shot-noise limit out to $k \sim 0.2 h \text{Mpc}^{-1}$. This allows a better matching with perturbative models and improved predictions for the clustering of e.g. quasars, galaxies and neutral Hydrogen measured in spectroscopic redshift surveys at very modest computational expense. We discuss the implementation of ZCV, give some examples and provide forecasts for the efficacy of the method under various conditions.

^bInstitute for Advanced Study, 1 Einstein Drive, Princeton, NJ 08540, USA

Contents

1	Introduction	1
2	Control variates: a worked example	3
3	The Zel'dovich approximation in redshift space	4
4	Simulations and galaxy models	8
5	 Results and discussion 5.1 Biased tracers 5.2 Satellite galaxies 5.3 Comparison to a linear theory control variate 5.4 Variance reduction forecasts 	9 9 11 13 14
6	Conclusions	15
A	Exact redshift-space power spectrum in the Zel'dovich approximation	18
в	Accounting for incomplete μ sampling	19
\mathbf{C}	Estimating β and ρ_{xc}	20

1 Introduction

One of the main tools in the employ of the modern cosmologist is the numerical simulation of structure formation. Such simulations are used in a broad range of cosmological studies, notably as tests of analytical models of large-scale structure and as forward models for cosmological observations. In both cases, one wishes to reduce statistical and systematic errors on measurements made from these simulations.

In the case of virtually all numerical simulations, run-time scales proportionally to the number of resolution elements employed, raised to some power (e.g. [1-3]). Thus, for a fixed number of resolution elements, statistical and systematic errors become anti-correlated, since statistical errors scale with volume, and systematic errors scale roughly with the size of the resolution elements employed. Because of this scaling, methods for reducing simulation statistical errors without increasing simulation volume are highly desirable, as then one can decrease simulation statistical error without increasing run-time or systematic error. Furthermore, one often wishes to run simulations at a number of different cosmologies in order to perform inference tasks. In this case, reducing the statistical error of individual simulations allows for a broader range of cosmologies to be simulated at fixed computational cost, thus making the models built from said simulations more accurate.

In the context of large-scale structure studies, a number of methods for reducing the statistical error of N-body simulations have been developed. Perhaps most common is the 'fixed amplitude' method [4, 5], where instead of initializing a simulation with a linear density field whose amplitudes are drawn from a Rayleigh distribution, one chooses the amplitude of each Fourier mode to be the expectation value determined from the linear power spectrum. In

doing so, one can reduce the variance of statistics measured from fixed amplitude simulations, with the largest reduction in variance coming in the linear regime and degrading significantly in non-linear or stochastic regimes. In addition to this technique, one often runs a second, 'paired', simulation with the same fixed amplitude initial conditions, but with modes 180° out of phase with the first simulation [5, 6]. Doing so cancels out the next-leading-order contribution to the variance of two-point matter statistics. There is now a large literature investigating the performance of these techniques in the context of various cosmological models [7], and for various statistics beyond the matter power spectrum [6, 8–13]. These techniques clearly produce non-Gaussian initial conditions, but biases incurred to two- and three-point functions from this violation have been shown to be small for many cases of interest [8, 9, 13]. Furthermore, these techniques render the simulations unusable for the estimation of covariance matrices.

More recently, the method of control variates has been introduced in the context of cosmology. Control variates have been studied for many years in the statistics literature as a method for reducing the variance of Monte Carlo estimators in general [14]. When estimating the mean of a random variable of interest via Monte Carlo, if there exists a correlated random variable, i.e. the control variate, whose mean is precisely known, then one can construct an estimator for the mean of the variable of interest that has drastically improved convergence properties. This method was introduced in the cosmology literature under the name 'Convergence Acceleration by Regression and Pooling' (CARPool) in order to reduce the variance of summary statistics, such as N-point statistics [15, 16], and covariance matrices [17, 18] measured from cosmological simulations. This method does not require the simulator to alter any properties of the initial conditions, e.g. by introducing non-Gaussianity, and with a well chosen control variate one can also reduce the variance of measured statistics to a significantly larger degree than what is seen with paired and fixed simulations, particularly beyond the linear regime.

In [15–18], approximate or particle-mesh N-body simulations were used as control variates for higher resolution N-body simulations. In this case, the dominant cost of the control variate method was estimating the mean of the control variate, which was still performed via Monte Carlo. In the case of [16], who used FastPM simulations [19] for their control variate, estimation of the mean required 500 FastPM simulations, using approximately 24 million NERSC CPU-hours. This choice of control variate, while having the beneficial properties mentioned above, is more expensive than the method of pairing and fixing.

This additional expense can be circumvented if an alternative control variate is chosen, such that it has an analytically calculable mean. In [20], we investigated the use of just such an alternative: the Zel'dovich approximation (ZA) [21, 22]. This choice, which we call Zel'dovich control variates (ZCV), comes with the benefit of a mean prediction that is known analytically to very high precision. Furthermore, when realized with the same initial conditions, ZA matter density fields are significantly correlated with their full *N*-body counterparts out to $k \sim 1 h \text{Mpc}^{-1}$. In [20], we showed that the ZA was an effective control variate for realspace matter, biased tracer auto- and cross-power spectra, and hybrid effective field theory (HEFT) basis spectra [23–25]. Depending on the statistic in question, ZCV produced variance reduction factors of 10² to 10⁶ at $k < 0.2 h \text{Mpc}^{-1}$.

In this work, we further extend the ZCV method to two-point statistics in redshift space. In section 2, we give an overview of the method of control variates, and work through a simplified example demonstrating how to apply ZCV to redshift-space matter power spectra. In section 3 we discuss our methodology for producing analytic and Monte Carlo ZA predictions for biased tracers in redshift space, and show that they agree with each other. Section 4 describes the simulations that we use in this work, and section 5 summarizes our main results. We conclude in section 6, with technical details related to our analytic predictions and estimators relegated to appendices.

2 Control variates: a worked example

In this section we provide a brief overview of the steps required to reduce the variance of the redshift-space matter power spectrum, $\hat{P}_{mm,\ell}(k)$, using the ZA as a control variate (throughout this work, we will make use of \hat{P} to denote measured power spectra). In the following sections, we will describe each step in greater detail and investigate corrections beyond these steps, and their impact on the resulting variance reduction.

Control variates are a general technique that can be applied to reduce the variance of a random variable, X, when a correlated random variable, C, with known mean, μ_c , is available. In this case, a new variable, Y, can be defined as

$$Y = X - \beta (C - \mu_c).$$

This is an unbiased estimator of X regardless of what β is set to, as $\langle C - \mu_c \rangle = 0$ (some caveats are discussed in appendix C). If we wish to minimize the variance of Y, it can be shown that the optimal choice for β is

$$\beta^{\star} = \frac{\operatorname{Cov}[X, C]}{\operatorname{Var}[C]}.$$
(2.1)

In this case, the variance of Y is given by

$$\operatorname{Var}[Y] = \operatorname{Var}[X](1 - \frac{\operatorname{Cov}[X, C]}{\operatorname{Var}[C]\operatorname{Var}[X]} + \beta^{2}\operatorname{Var}[\mu_{c}])$$
$$= \operatorname{Var}[X](1 - \rho_{xc}^{2} + \beta^{2}\operatorname{Var}[\mu_{c}]).$$

where ρ_{xc} is the cross-correlation coefficient between X and C, and $\operatorname{Var}[\mu_c]$ is the uncertainty on the mean of C. The first line holds so long as μ_c and β are not correlated with each other, nor with X. Typically, what has been done in the control variate literature in the context of cosmology is to estimate μ_c via Monte Carlo realizations of C, in which case the dominant cost of the control variate technique is this process. On the other hand, if μ_c is known analytically, then this additional expense can be circumvented, and $\operatorname{Var}[Y] = \operatorname{Var}[X](1 - \rho_{xc}^2)$.

In cosmology, X is often a measurement, such as the redshift-space matter power spectrum, $\hat{P}_{\text{mm},\ell}(k)$, estimated from an expensive simulation. In this case, we would like to choose C such that it is highly correlated with the redshift-space matter field, while still having an analytically known mean. In [20], we showed that the ZA satisfies both of these requirements for real space fields, and in this work, we demonstrate that this also holds in redshift space, for matter, halos and simulated galaxies.

The algorithm for using the ZA to reduce the variance of $\hat{P}_{\text{mm},\ell}(k)$ measured from an N-body simulation is as follows:

1. Generate an initial linear density field $\delta(\mathbf{k})$ using the same random seed as used to run the *N*-body simulation, and use this to compute Zel'dovich displacements, $\Psi(\mathbf{q})$.

- 2. Compute the redshift-space Zel'dovich density field $\delta_s^{\text{ZA}}(\mathbf{k}, z)$ at the desired redshift z using the linear growth factor, D(z), and growth rate, f(z).
- 3. Measure the ZA and N-body matter auto-power spectra, and ZA–N-body cross-power spectra, $\hat{P}_{\text{mm},\ell}^{zz}(k), \hat{P}_{\text{mm},\ell}^{nn}(k)$, and $\hat{P}_{\text{mm},\ell}^{zn}(k)$ and estimate β using a disconnected approximation for $\text{Cov}[\hat{P}_{\text{mm},\ell}^{zz}(k), \hat{P}_{\text{mm},\ell}^{nn}(k)]$.
- 4. Analytically compute the mean Zel'dovich redshift-space auto-power spectrum, $P_{\text{mm},\ell}^{zz}(k)$, and construct the reduced variance matter power spectrum via $\hat{P}_{\text{mm},\ell}^{*,nn}(k) = \hat{P}_{\text{mm},\ell}^{nn}(k) \beta \left(\hat{P}_{\text{mm},\ell}^{zz}(k) P_{\text{mm},\ell}^{zz}(k)\right)$.

The results of this process are summarized in figure 1. The blue curves in the top panels display the monopole and quadrupole moments of the redshift-space matter power spectrum measured from our N-body simulation, with their fractional errors, estimated using a disconnected covariance approximation, shown in blue in the bottom panels. In solid black, we also display the ZA matter power spectra produced with the same initial conditions. On large scales, these curves are almost indistinguishable, indicating the large amount of correlation that exists between the ZA and N-body matter fields in redshift space. The black dashed line displays the analytic prediction for $P_{\text{mm},\ell}^{zz}(k)$, whose computation we will briefly overview in §3. The orange curve in the bottom panel displays the fractional error on $\hat{P}_{\text{mm},\ell}^{*,nn}(k)$, which is reduced with respect to the blue curve by a factor of $1 - \rho_{xc}^2$. ρ_{xc} is estimated from a single simulation as described in section 5.1 and appendix C. We see that for the redshift-space matter power spectrum we are able to reduce the errors on our measurements from a $(1050 h^{-1} \text{ Mpc})^3$ simulation to yield fractional errors, represented by the shaded gray regions in the bottom panel [26–28].

For the purposes of demonstration, we also display the effect of applying the ZCV method to correlation functions in figure 2. Here, we apply Hankel transforms to the redshift-space matter power spectra in figure 1, extrapolating at $k < 0.01 h \text{Mpc}^{-1}$ using our analytic ZA model, and at k > 4 using a power law extrapolation. The comparison between the correlation functions with and without ZCV applied is quite stark, with the large scales matching ZA nearly perfectly in the ZCV case, while significant noise is apparent in its absence. Correlation functions can be estimated more accurately by using ZCV to reduce the variance of the un-binned three-dimensional redshift-space power spectrum, Fourier transforming to threedimensional correlation functions, and measuring multipoles. This would avoid any reliance on interpolation and extrapolation. Doing so is readily possible with the methods presented in this work.

In the following sections, we will elaborate on this example, demonstrating the efficacy of ZCV for redshift-space galaxy and halo power spectra. In order to do this, we will introduce bias operators in our Zel'dovich field in section 5.1. We have also glossed over a number of technical details that are important for obtaining an unbiased control variate estimator, such as window function deconvolution and the estimation of β that are discussed in detail in Appendix B and C, respectively.

3 The Zel'dovich approximation in redshift space

Within a cold dark matter universe, structure formation is entirely determined by the displacements $\Psi(q, a)$, or trajectories, of phase-space elements over cosmic time. The displacements



Figure 1. (Top) Comparison of redshift-space matter power spectrum multipoles measured from our fiducial $(1050 h^{-1} \text{ Mpc})^3 N$ -body simulation (blue), a numerical realization of the ZA using the same initial linear density field as the N-body simulation (black), and the N-body spectrum with reduced variance using the black curve as a control variate (orange). For reference, we also display the mean prediction for the Zel'dovich matter power spectrum (black dashed). On large scales, the variance reduced N-body spectra agree nearly perfectly with the noiseless analytic ZA prediction. (Bottom) Fractional errors, computed via a disconnected covariance approximation (blue), and the variance reduced version of these errors (orange). The variance reduction is given by $1 - \rho_{xc}^2$. The shaded gray region below 0.5% fractional error is approximately where we expect the contribution from systematic errors in simulations to become appreciable.



Figure 2. Matter correlation function multipoles with (orange) and without (blue) applying the ZCV method, compared to our analytic ZA theory (dashed). We have Hankel transformed the measured power spectra in order to produce correlation functions as described in the text. A significant reduction in noise is apparent in the correlation functions that have ZCV applied.

translate between the initial positions q of these elements and their positions $x = q + \Psi(q, a)$

at scale factor a. The matter density δ_m is obtained by counting the particles that end up at position \boldsymbol{x} at a, i.e.

$$1 + \delta_m(\boldsymbol{x}, a) = \int d^3 \boldsymbol{q} \ F(\boldsymbol{q}) \ \delta_D(\boldsymbol{x} - \boldsymbol{q} - \Psi(\boldsymbol{q}, a)), \tag{3.1}$$

where the weight given to each dark matter particle is F(q) = 1, so that the non-linear clustering of matter can be entirely written in terms of the statistics of the displacements. For biased tracers like galaxies we need to additionally weight each q by functions of the initial conditions using the so-called bias expansion (see e.g. [29, 30]) such that

$$F(\boldsymbol{q}) = 1 + b_1 \delta_0 + b_2 \left(\delta_0(\boldsymbol{q})^2 - \langle \delta_0^2 \rangle \right) + b_s \left(s_0(\boldsymbol{q})^2 - \langle s_0^2 \rangle \right) + \dots$$
(3.2)

where $s^2 = s_{ij}s_{ij}$ is the local shear field. From eq. 3.1 we can then see that the final biased tracer field, $\delta_t(\boldsymbol{x}, a)$, can be expressed in terms of an expansion in the advected initial fields,

$$1 + \delta_t(\boldsymbol{x}, a) = \sum_{\mathcal{O}_i \in m, \delta_0, \, \delta_0^2, \, s_0^2, \, \dots} b_{\mathcal{O}_i} \mathcal{O}_i(\boldsymbol{x}, a) \,, \tag{3.3}$$

where the advected operators O_i are given in both configuration and Fourier space by

$$\mathcal{O}_{i}(\boldsymbol{x}, a) = \int d^{3}\boldsymbol{q} \ \mathcal{O}_{i}(\boldsymbol{q}) \ \delta_{D}(\boldsymbol{x} - \boldsymbol{q} - \Psi(\boldsymbol{q}, a))$$
$$\mathcal{O}_{i}(\mathbf{k}, a) = \int d^{3}\boldsymbol{q} \ e^{-i\mathbf{k}\cdot(\boldsymbol{q} + \Psi(\boldsymbol{q}))} \ \mathcal{O}_{i}(\boldsymbol{q})$$
(3.4)

where we will denote the advected initial fields as $\mathcal{O}_i(\boldsymbol{x}, a)$. Furthermore, for notational convenience, we have set $\mathcal{O}_m(\boldsymbol{q}) = 1$, $b_{\mathcal{O}_m} = 1$, and thus $\mathcal{O}_m(\boldsymbol{x}, a) = \delta_m(\boldsymbol{x}, a)$.

There exist many ways to compute the displacements, Ψ . Perhaps the most straightforward is to directly simulate the dynamics of structure formation using N-body simulations. In this case, the Lagrangian position q is the grid position set in the initial conditions, and Ψ is simply given by the trajectory of the N-body particle over the simulated time. Within Λ CDM universes close to our own, however, the largest contribution to Ψ comes from large-scale bulk flows that are extremely well captured by first-order Lagrangian perturbation theory (LPT), also known as the Zel'dovich approximation (ZA) [21]. The advection of dark matter particles away from their initial positions q by the Zel'dovich displacement is the leading contribution to the decorrelation of large scale structure with the initial conditions [31], and the cancellation of this effect by using Zel'dovich power spectra as our control variate is the main advantage of ZCV over linear initial condition control variates, as we show in §5.3.

In the ZA, the displacement is set by the initial gravitational potential gradient experienced by the particle, given by Poisson's equation to be $k^2 \Phi \propto \delta_m$, such that

$$\Psi^{(1)}(\mathbf{k}, a) = \frac{i\mathbf{k}}{k^2} D(a) \ \delta_0(\mathbf{k}) \tag{3.5}$$

where D(a) is the linear-theory growth factor. Likewise, peculiar velocities are linearly proportional to $\mathbf{v}^{(1)} = f\mathcal{H}\Psi^{(1)}$, where $f = d\ln D/d\ln a$ is the linear growth rate and \mathcal{H} is the conformal Hubble parameter. The peculiar velocity is important because, in most situations, the distance to a galaxy is inferred from its cosmological redshift, which receives a contribution from the peculiar line-of-sight velocity $\mathbf{u} = \mathbf{v}_{\parallel}/\mathcal{H}$. Within the ZA this is equivalent to multiplying the displacement by a constant matrix [32]

$$\Psi_{s,i}^{(1)} = R_{ij}\Psi_j^{(1)}, \quad R_{ij} = \delta_{ij} + f\hat{n}_i\hat{n}_j.$$
(3.6)

where \hat{n} is the line-of-sight unit vector.

In this paper we will be mostly concerned with the redshift-space 2-point function in configuration space and Fourier space, i.e. the correlation function and power spectrum. From Equation 3.1 one can show [33] that the biased tracer power spectrum is given by

$$P_s^{tt}(\mathbf{k}) = \sum_{\mathcal{O}_i, \mathcal{O}_j} b_{\mathcal{O}_i} b_{\mathcal{O}_j} P_{ij, s}(\mathbf{k})$$
(3.7)

where we have defined the basis spectra

$$P_{ij,s}(\mathbf{k})(2\pi)^{3}\delta_{D}(\mathbf{k}+\mathbf{k}') = \left\langle \mathcal{O}_{i,s}(\mathbf{k})\mathcal{O}_{j,s}(\mathbf{k}')\right\rangle.$$
(3.8)

The operators $\mathcal{O}_{i,s}(\mathbf{k})$ are the redshift-space counterparts of $\mathcal{O}_i(\mathbf{k})$, defined such that Ψ is replaced by Ψ_s in their construction. We make predictions for the ensemble mean of these statistics with the ZeNBu code¹.

For a given realization of the initial conditions on a grid we can numerically construct these Zel'dovich operators and their spectra by realizing the linear density field δ_0 on a cubic grid, and constructing the relevant initial condition fields, δ_0^2 , s_0^2 , from δ_0 . We can then numerically compute the displacements using Eqs. 3.5 and 3.6, which we can use to perform the advection integral in Eq. 3.3, using $\Psi(\mathbf{q})$ and $\Psi_s(\mathbf{q})$ for real and redshift-space quantities, respectively. Predicting biased tracer power spectra is then as simple as measuring the cross-power spectra between the advected operators and performing the sum in Eq. 3.7. We make use of CLASS [34, 35] to compute linear power spectra, growth and growth rates, and Monofonic [36] for our grid-based predictions of δ_0 and Ψ . The code to compute the rest of the quantities required for this work is publicly available².

We can also compute these spectra analytically without sample variance. Within the ZA, the expectation value of the exponentiated displacements in eq. 3.4 can be solved exactly since the displacements are purely Gaussian; for the bias expansion in Equation 3.2 this was done in real space to all orders in ref. [20]. In order to perform the same calculation in redshift space, it is necessary to transform the displacements via the matrix R_{ij} , or alternatively by transforming the Fourier space wavenumber by the same amount, since $\mathbf{k}^T (\mathbf{R}\Delta) = (\mathbf{k}\mathbf{R})^T \Delta$. Numerical methods to compute the quantities in the latter choice of coordinates were developed in [37–40]; roughly, the azimuthally symmetric real space integrals are replaced with ones accounting for the line-of-sight dependence in redshift space. We review these developments in the specific context of the ZA in Appendix A.

The validity of the Zel'dovich control variate method depends on our ability to very accurately predict the ensemble average of the grid-based realizations, averaging over realizations of δ_0 . In [20], we showed that we were able to do this in real space after smoothing δ_0 using a Gaussian kernel with a width equal to the mesh scale, $N\pi/L$ where N is the mesh size and L is the side length of the simulation cube. We find that we are able to predict our redshiftspace grid-based results with our analytic code to comparable accuracy after performing this same smoothing procedure, using two times the smoothing scale, i.e. $k_{\text{smooth}} = N\pi/(2L)$, as shown in figure 3. We were not able to achieve a comparable level of accuracy for the hexadecapole, and so do not treat it in this work. Care must be taken to correct for artifacts in our grid-based measurements imparted due to discrete **k** sampling in our simulations. This is discussed in Appendix B.

¹https://github.com/sfschen/ZeNBu

²https://github.com/kokron/anzu



Figure 3. The difference between our grid-based ZA realizations and their analytic predictions for $P_{ij,\ell}(k)$ up to linear order in δ_0 . These measurements are averaged over 60 random seeds, and coarsely binned in k to suppress variance. We have offset the measurements for the different basis spectra in k to aid in presentation. The lines displayed are $0.0025P_{ij,\ell}(k)$, i.e. they represent a 0.25% error in each spectrum. Away from zero-crossings, the difference between our grid-based and analytic ensemble mean predictions is consistent with or less than this, and thus any uncertainty due to our ability to predict the mean of $P_{ij,\ell}(k)$ is negligible. We find consistent behavior for all other basis spectra, but we do not display them in order to facilitate presentation. On large scales, a proper treatment of the mixing of multipoles due to discreteness effects in our simulation is of great importance for obtaining this level of agreement, see app. B.

4 Simulations and galaxy models

In this work, we employ the Aemulus ν simulations [41], a suite of CDM and massive neutrino simulations run in a similar configuration to the Aemulus α simulations [42], but with an expanded cosmological parameter space, and a greater number of realizations. Each simulation has a volume of $(1.05h^{-1}\text{Gpc})^3$, evolving 1400³ CDM particles, and the same number of neutrino particles. The simulations are initialized at z = 12 with 3rd-order Lagrangian perturbation theory, using a version of the Monofonic code [36] modified to account for the presence of neutrinos on the evolution of δ_{cb} [43]. Further details and convergence tests of these simulations will be presented in [41]. In this work, we make use of the α 000 simulation with a cosmology of $n_s = 0.97$, $H_0 = 67.0$, w = -1.0, $\omega_b = 0.0223$, $\omega_c = 0.12$, $\sum M_{\nu} = 0.07 \text{eV}$, $\sigma_8 = 0.81$. For simplicity, we will conflate matter and δ_{cb} power spectra in this work, as the difference between the two is unimportant for all of the findings presented here. Halo finding was performed using the ROCKSTAR spherical overdensity halo finder [44] and M_{200b} strict spherical overdensity (SO) masses.

When populating galaxies in our simulation, we make use of the halo occupation distribution (HOD) formalism [45, 46]. In particular, we present tests using galaxies populated with HODs consistent with the DESI luminous red galaxy (LRG) sample, making use of the HOD functional form presented in [47]. We take as parameters of this model the best-fit values for the second redshift bin presented in [48]. For the LRG sample, we make use of the z = 0.56 output from our simulation, yielding a number density of $\bar{n} = 3 \times 10^{-4} h^3 \text{Mpc}^{-3}$.

5 Results and discussion

Here we demonstrate the effectiveness of ZCV for halo and galaxy samples, and discuss some of the complications that arise in these cases. We begin by investigating the impact of biasing, followed by satellite galaxies and their accompanying addition of stochastic small scale velocities. We will see that even in the presence of these complications, we can reduce the sample variance of redshift-space power spectra down to the shot-noise limit until $k \sim 0.2 h \text{Mpc}^{-1}$. We then compare the ZCV method to a simpler method using linear theory as a control variate. Finally, we take advantage of the demonstrated performance of ZCV to forecast the amount of variance reduction we expect to achieve for a number of key galaxy samples that will be used for cosmological studies in the coming years. For most of the samples we consider, we find variance reduction factors of $\sim 10 - 100 \times$.

5.1 Biased tracers

In section 2, we introduced the ZCV method in the context of the redshift-space matter power spectrum. While this application may be useful in some contexts, it represents an idealized case. Instead, one is often interested in measuring power spectra of biased tracers such as halos or galaxies. In this section, we will investigate the extent to which ZA predictions are correlated with these statistics. Importantly, we wish to disentangle the effects that biasing, shot noise, and stochastic small scale velocities, i.e. the finger of god (FoG), have on the efficacy of the ZCV method. In order to do so, we will first consider ZCV in the context of power spectra of halos in bins of mass. By varying halo mass, we aim to isolate the impact that halo bias, both linear and non-linear, as well as shot noise, has on the performance of the ZCV algorithm.

First, we consider the ability of the Zel'dovich matter power spectrum to act as a control variate for the power spectra of halo samples split into three logarithmically spaced bins in mass between $10^{13} h^{-1} M_{\odot}$ and $10^{14.5} h^{-1} M_{\odot}$. We have chosen relatively high mass halos to emphasize the role of shot noise, which leads to decorrelation and hence reduces the effectiveness of the control variates technique. In order to quantify the effectiveness of the ZA as a control variate, we measure the cross correlation coefficient, ρ_{xc} , between the redshift-space power spectra of these halo samples, and the redshift-space Zel'dovich matter power spectrum. We compute this quantity as

$$\rho_{xc} = \frac{\operatorname{Cov}[P_{\ell}^{tt}(k), P_{\ell}^{zz}(k)]}{\sqrt{\operatorname{Var}[\hat{P}_{\ell}^{tt}(k)]\operatorname{Var}[\hat{P}_{\ell}^{zz}(k)]}}$$
(5.1)

where $\hat{P}_{\ell}^{tt}(k)$ are the multipoles of the redshift-space power spectrum of the (biased) tracer under consideration, and $\hat{P}_{\ell}^{zz}(k)$ are the same for the ZA field that we employ as a control variate. The fractional variance reduction afforded by a particular control variate is equal to $1-\rho_{xc}^2$, making this the main quantity of interest in our study. We approximate the covariances and variances in eq. 5.1, keeping only their disconnected contributions, as described in app. C. In this case, the Zel'dovich field is simply the unweighted matter field, but we will make use of the same method for measuring ρ_{xc} when we include bias operators in the ZA. We expect the disconnected approximation to break down at high k, but for $k < 0.2 h \text{Mpc}^{-1}$ it holds to very high accuracy [49]. As we will see, above $k \sim 0.2 h \text{Mpc}^{-1}$ the ZA fields decorrelate from the fields of interest, and so break-downs in this approximation are not important for the conclusions of this work.



Figure 4. Cross-correlation coefficient between the redshift-space ZA matter power spectrum and various tracer power spectra, including the non-linear matter field, and halos in three bins of mass as denoted in the legend. The different panels show the cross correlation between different multipoles of the redshift-space power spectrum. The dashed lines are the cross-correlation coefficient that we would expect if the tracer field were perfectly correlated with the ZA field apart from a constant shot-noise contribution. The different multipoles exhibit very similar behavior, with more biased, higher shot-noise samples decorrelating more. As discussed later, including higher-order bias operators allows us to saturate the shot-noise limit for all samples to $k \sim 0.2 h \text{Mpc}^{-1}$.

Figure 4 displays ρ_{xc} for each of the three halo mass bins that we consider. For reference, we also compare ρ_{xc} between the nonlinear redshift-space matter power spectrum, $\hat{P}_{\text{mm},\ell}(k)$, and the Zeldovich redshift-space matter power spectrum, $\hat{P}_{\text{mm},\ell}^{zz}(k)$, in black. The main aspect of note is that ρ_{xc} decreases monotonically with halo mass, and this decrease becomes even more pronounced with respect to the $\hat{P}_{\text{mm},\ell}$ case at high k. The dashed lines in this figure show the expectation for ρ_{xc} in the limiting case where the halo power spectra correlate perfectly with their ZA counterparts except for the shot-noise contribution to the former (we describe how we estimate the shot noise momentarily). These dashed lines are the best one can do without incorporating a control variate with the same shot-noise properties as the tracer sample in question. Finding such a control variate is highly nontrivial, as is analytically or cheaply computing its mean, as the shot-noise level depends on a number of things other than the number density of the sample in question, such as halo exclusion, and non-linear biasing effects [50–52].

Now we proceed to investigate the extent to which including bias operators in our Zel'dovich field increases ρ_{xc} for the halo power spectra discussed above. In order to include additional bias operators in the ZA, we must fit for the bias coefficients of each tracer sample. We do this at the field level, minimizing the residuals between the Zel'dovich and halo fields in real space by minimizing

$$S = \langle \epsilon(\boldsymbol{x})^2 \rangle_{|\mathbf{k}| < k_{\max}}$$

$$\approx \int_{|\mathbf{k}| < k_{\max}} \frac{d^3 \mathbf{k}}{(2\pi)^3} \left| \delta^t(\mathbf{k}) - \delta^z_m(\mathbf{k}) - \sum_{\mathcal{O}_i} b_{\mathcal{O}_i} \mathcal{O}_i(\mathbf{k}) \right|^2,$$

following the procedure outlined in section 2.3 of [52], but replacing the HEFT operators with their ZA counterparts. We perform fits including the $\delta_0(\mathbf{k})$, $\delta_0^2(\mathbf{k})$ and $s_0^2(\mathbf{k})$ fields. This procedure also allows us to precisely determine the shot-noise contribution to the tracer power spectra, because at low-k, shot noise is equivalent to the error power spectrum:

$$P^{\epsilon\epsilon}(k) = \langle \epsilon(\mathbf{k})\epsilon(-\mathbf{k}) \rangle$$

= $\langle [\delta^t(\mathbf{k}) - \delta^z(\mathbf{k})][\delta^t(-\mathbf{k}) - \delta^z(-\mathbf{k})] \rangle$,

where we have applied the shorthand $\delta^{z}(\mathbf{k}) = \delta_{m}^{z}(\mathbf{k}) + \sum_{\mathcal{O}_{i}} b_{\mathcal{O}_{i}} \mathcal{O}_{i}(\mathbf{k}).$

We perform fits to $k_{\text{max}} = 0.15 \, h \text{Mpc}^{-1}$, and average $\hat{P}^{\epsilon\epsilon}(k)$ over this range in order to estimate the shot-noise values. We have confirmed that our shot-noise estimates do not change when including the effects of nonlinear displacements, e.g. by replacing the Zel'dovich fields used above with their HEFT counterparts. We make use of this estimate of shot noise, rather than the Poissonian expectation, because it incorporates the effect of higher-order biasing and halo exclusion [52]. Poissonian shot-noise estimates would mostly overestimate the impact of shot noise on the performance of ZCV for our halo and galaxy samples, as halo exclusion effects tend to reduce the level of shot noise compared to the Poisson expectation for the halo mass range used here, leading to lower levels of uncorrelated noise than the naive estimate.

With an estimate of the shot noise of the tracer in question, we can then construct our estimate of the shot-noise limit of ρ_{xc} as

$$\rho_{xc}^{\text{SN. limit}} = \frac{\text{Var}_{\text{noSN}}[\hat{P}_{\ell}^{tt}(k)]}{\sqrt{\text{Var}[\hat{P}_{\ell}^{tt}(k)]\text{Var}_{\text{noSN}}[\hat{P}_{\ell}^{tt}(k)]}}$$
(5.2)

where $\operatorname{Var}_{\operatorname{noSN}}[\hat{P}_{\ell}^{tt}(k)]$ is estimated by subtracting $\langle \hat{P}^{\epsilon\epsilon}(\mathbf{k}) \rangle_{|\mathbf{k}| < k_{\max}}$ from the monopole in the disconnected approximation for the variance.

Figure 5 demonstrates the extent to which adding bias operators to our ZA control variate can optimize the sample variance reduction properties of the ZCV method. The top and bottom panels demonstrate two different regimes. The top shows the behavior of ZCV for a relatively low-bias halo sample, with bias values of $b_{\delta_0} = 1.06$, $b_{\delta_0^2} = -0.14$ and $b_{s_0^2} = 0.26$. We see that including additional bias operators makes a relatively minor difference. Using the unbiased (i.e. matter) ZA redshift-space power spectrum nearly saturates the shot-noise limit, and operators beyond linear bias yield no additional improvement.

The bottom panel displays a much higher bias halo sample, with $b_{\delta_0} = 3.72$, $b_{\delta_0^2} = 9.09$ and $b_{s_0^2} = -3.0$. Here we see large improvements going from the unbiased ZA, to linear bias and quadratic bias, with the latter saturating the shot-noise limit to $k \sim 0.2 h \text{Mpc}^{-1}$. Somewhat unsurprisingly, we never observe improvements when adding in tidal bias. This is simply explained by the fact that the basis spectra, $P_{ij}(k)$, that contain the tidal field operator in eq. 3.7 are highly subdominant to the linear and quadratic bias terms on scales where the ZA is still correlated with our halo fields. As such, it seems unlikely that tidal bias operators will ever be required for an optimal ZCV implementation.

5.2 Satellite galaxies

Having seen that ZCV is able to saturate the shot-noise limit for halos, we now turn to investigating whether this still holds for galaxy samples populated via HODs. In particular, one might expect the addition of satellite galaxies to boost stochastic contributions from nonperturbative small scale velocities, i.e. the FoG, thus decorrelating our ZA control variate from the galaxy field more than what we found for halos. We consider three different samples for this test. First, we examine the LRG-like HOD described in Sec. 4, taken from [48]. We



Figure 5. The effect of including additional bias terms in the ZA on the cross-correlation coefficient between multipoles of the redshift-space Zel'dovich and halo power spectra. We show the two extreme halo mass bins: $13 < \lg M \le 13.5$ (top) and $14 < \lg M \le 14.5$ (bottom). For reference, we display the expectation for the correlation coefficient if all of the signal in the field were perfectly correlated (and only shot noise contributed to decorrelation) in solid black. We saturate the shot-noise limit in both bins, although including higher order bias operators is much more important in doing so for the higher mass sample.

also consider a galaxy sample with two times the satellite fraction of the LRG-like sample, i.e. $f_{\rm sat} = 0.3$ compared to the original $f_{\rm sat} = 0.15$, which we have achieved by setting $\lg M_1 = 13.6$ and leaving the remaining HOD parameters fixed. Finally, we use a sample containing only the central galaxies from the LRG-like HOD, thus minimizing the FoG effect but otherwise keeping the mean halo mass nearly the same.

Figure 6 investigates the extent to which these samples correlate with a Zel'dovich control variate that includes linear and quadratic bias operators. The top panel shows ρ_{xc} for these samples, compared to dashed lines representing each respective samples' shot-noise limit. Here, we see that the samples including satellite galaxies decorrelate from our control variate more rapidly than the central-galaxy-only sample, suggesting that non-linearities in the form of non-linear bias and small scale velocity contributions do lead to additional decorrelation, as expected. Somewhat surprisingly, the amount of additional decorrelation is not particularly sensitive to what multipole we are looking at, although some slight additional decorrelation is apparent for $\ell = 2$ with respect to $\ell = 0$. The fact that the behavior is so similar between the multipoles can largely be attributed to the fact that $P_0(k)$ is the dominant contribution to our disconnected covariances. We have also examined the convergence of our covariance matrices with respect to ℓ_{max} in App. C, demonstrating that our findings are not highly sensitive to our choice of $\ell_{\text{max}} = 4$.

The bottom panels then show the ratio of ρ_{xc} to the shot-noise limit for each sample, where the dashed lines in the bottom left panel show the equivalent quantity in real space. At $k < 0.2 h \text{Mpc}^{-1}$, the solid and dashed lines are nearly identical, although at higher k the realspace power spectra remain more correlated with our control variate than their redshift-space counterparts. Thus, it is apparent that non-linear velocities do degrade the performance of the ZCV method, but only slightly.



Figure 6. (*Top*) Comparison of ρ_{xc} for an LRG-like galaxy sample, the central galaxies from that sample, and an HOD with two times the number of satellite galaxies as the fiducial HOD. All samples use a control variate that includes linear and quadratic bias. The shot-noise limit for each sample is shown by the dashed lines. To $k \sim 0.2 h \text{Mpc}^{-1}$, we find that we are able to reduce the variance of all the samples to this limit. At higher k, the extent to which we are able to reduce the variance is sensitive to decorrelation between the Zel'dovich and galaxy fields beyond the effect of shot noise. (*Bottom*) The ratio of ρ_{xc} to the shot-noise expectation for each sample. Dashed lines represent this value for the real-space power spectra. The real space measurements exhibit less decorrelation beyond $k \sim 0.2 h \text{Mpc}^{-1}$, indicating that additional decorrelation is present in redshift space due to, e.g., the finger of god effect.

5.3 Comparison to a linear theory control variate

A common strategy for mitigating sample variance in simulations that has been used for many years is to divide the measurement of interest by the power spectrum of the linear density field used to initialize the simulation, i.e.

$$Y = \frac{X}{C}\mu_c$$

where X is the simulation measurement in question, C is the linear matter power spectrum measured from the initial conditions of the simulation, and μ_c is the noiseless linear power spectrum. This estimator, known as the "ratio control variate" estimator [14], is biased because $\langle X/C \rangle \neq \langle X \rangle / \langle C \rangle$.

A similar, but unbiased estimator using the same measured quantities is to use the measured linear power spectrum as a replacement for the ZA in the control variate estimator that we have been using for the rest of this work [6, 14]. It is interesting to consider how well this linear theory control variate performs compared to ZCV.

In fig. 7, we make just such a comparison using our fiducial LRG HOD. We include the Kaiser factor [53], $(1 + b_{\delta_0} + f \mu^2)^2$, in our linear theory control variate in order to optimize its performance. For ZCV, we use a combination of linear and quadratic bias. We see that for scales with $k < 0.05 h \text{Mpc}^{-1}$, linear theory performs as well as ZCV, but for k larger than this the performance degrades significantly. Thus, unless one only cares about reducing variance on these very large scales, then ZCV performs significantly better and should be preferred given the small difference in computational cost between the two.



Figure 7. (*Top*) Here we compare the performance of the ZCV method to an alternative linear control variate using the LRG HOD as a test case. We see that the linear control variate performs well out to $k \sim 0.05 h \text{Mpc}^{-1}$, but rapidly degrades in performance on smaller scales with respect to ZCV. Furthermore, ρ_{xc} is noisier for the linear control variate, exacerbating the bias discussed in app. C. (*Bottom*) A direct comparison of ZCV and linear control variates to the shot-noise limit. Again we see the linear control variate performs significantly worse.

5.4 Variance reduction forecasts

In the previous sections we demonstrated that we can reduce the variance of biased tracer power spectra in real and redshift space to their shot-noise limit to $k \sim 0.2 h \text{Mpc}^{-1}$. Using this fact, we now proceed to forecast the variance reduction that we can expect for a number of galaxy samples of interest for ongoing and upcoming galaxy surveys such as DESI [54], and planned future spectroscopic surveys [55, 56]. We focus on forecasting two quantities. The first is $1-\rho_{xc}^2$, which is the effective decrease in variance that is delivered by the ZCV method. We make our forecasts for real- and redshift-space galaxy auto-power spectra as well as for galaxy-matter cross-power spectra. With these, and matter power spectra shown in Fig. 1, we have all the statistics necessary for two-point function analyses of galaxy surveys and their cross-correlations with lensing. We also show $\sigma[P(k)]/P(k)$ assuming a $(2h^{-1}\text{Gpc})^3$ volume used in this work in order to demonstrate that for the vast majority of galaxy samples, one never needs to simulate a larger volume than this in order to make statistical simulation error subdominant to the systematic error floor associated with current N-body codes.

In order to perform these forecasts, we assume that we can achieve the shot-noise limit for ρ_{xc} , and thus can predict it given a model for the power spectrum in question, its covariance, and shot-noise level. For simplicity, we assume that only linear bias contributes to the signal, since for $k < 0.2 h \text{Mpc}^{-1}$ higher order bias operators contribute at the tens of percent level for reasonable bias values. Furthermore, the covariances are approximated using their disconnected parts, and the shot noise is assumed to be given by the Poisson expectation. We forecast $1 - \rho_{xc}^2$ and $\sigma[P(k)]/P(k)$ for five samples: DESI Bright Galaxy Survey (BGS), LRGs, emission line galaxies (ELGs) and quasars (QSOs), as well as for Lyman-alpha emitters (LAEs) which may form the workhorse sample for future spectroscopic surveys [57]. For BGS, we assume $b_{\delta_0} = 0$, i.e. an Eulerian bias of 1, a number density of $\bar{n} = 0.01 h^3 \text{Mpc}^{-3}$ and an effective redshift of $z_{\text{eff}} = 0.3$, comparable to the $M_r < -19.5$ sample from SDSS [58], but at higher redshift. For DESI LRGs [59], we use $b_{\delta_0} = 1.2$ and $z_{\text{eff}} = 0.56$, consistent with the HOD model used in this paper, and $\bar{n} = 3 \times 10^{-4} h^3 \text{Mpc}^{-3}$. For ELGs, we take $b_{\delta} = 0.7$, $z_{\text{eff}} = 1.25$, and $\bar{n} = 3.8 \times 10^{-4} h^3 \text{Mpc}^{-3}$ based on [54, 60]. For QSOs, we assume $b_{\delta} = 2$, $z_{\text{eff}} = 2$, and $\bar{n} = 2.6 \times 10^{-5} h^3 \text{Mpc}^{-3}$ based on [61]. Finally, for LAEs we set $b_{\delta} = 1.5$, $z_{\text{eff}} = 2.5$, and $\bar{n} = 4 \times 10^{-4} h^3 \text{Mpc}^{-3}$ [62].

Figure 8 shows the results of these forecasts. The top panel shows $1/(1 - \rho_{xc}^2)$, which is the effective increase in volume as a function of scale that the ZCV method provides. We see marked improvements in effective volume for all the samples considered, other than the DESI QSOs which are almost entirely shot-noise dominated at all relevant scales. Unsurprisingly, the sample that performs best is DESI BGS, with its extremely high number density. $P_{gm}(k)$ also exhibits a greater variance reduction for all samples, as the shot-noise contribution to its covariance is reduced with respect to galaxy density auto-correlations. Above $k \sim 0.2 h \text{Mpc}^{-1}$, the approximation that ZCV achieves the shot-noise limit on ρ_{xc} may break depending on the satellite fraction and non-linear biasing behavior of the galaxy samples under consideration, and so our forecasts become less trustworthy above these scales.

The solid lines in the bottom panel show the fractional errors that we expect for each sample when applying ZCV, to be compared to the raw fractional errors shown by the dashed lines. The grayed out region roughly represents the level at which we expect systematic errors in our simulations to dominate over statistical errors [26–28]. We have assumed a simulation volume of $(2h^{-1}\text{Gpc})^3$ and a binning of $\Delta k = 0.01$ to compute the errors. We see that with this volume we achieve the systematic error floor by $k \sim 0.07 h \text{Mpc}^{-1}$ for all the samples considered other than QSOs, thus suggesting that for the majority of galaxy samples that will be observed in the coming years, it will not be important to simulate more than a volume of $(2h^{-1}\text{Gpc})^3$ in order to remove statistical error from simulated two-point galaxy measurements on scales with $k > 0.1 h \text{Mpc}^{-1}$. At larger scales, perturbation theory is known to work exquisitely well [40, 63].

6 Conclusions

In this work, we have extended the Zel'dovich control variate (ZCV) technique to redshift space. In [20], we demonstrated the utility of the Zel'dovich approximation (ZA) as a control variate for real-space measurements. The great benefit of ZCV in real space comes from our ability to inexpensively produce realizations of the ZA, and to analytically predict the mean of those realizations, obviating the two main costs usually associated with the control variate technique. In this work, we showed that we can do so equally well for multipoles of the redshift-space power spectrum of biased tracers, achieving the theoretical limit for variance reduction imposed by the shot noise of those tracers.

In section 2, we gave an overview of the ZCV technique in its simplest form, providing a recipe for applying it to the redshift-space matter power spectrum, which we then expanded to include a treatment for biased tracers in the following sections. In section 3, we described how we make redshift-space power spectrum predictions in the ZA, both on a grid for a particular realization of a linear density field, and for noiseless analytical calculations of the ensemble mean. Section 4 introduced the simulations and HOD that we used to illustrate the ZCV technique. In section 5, we proceeded to demonstrate the extent to which ZCV can reduce the variance of biased tracers by measuring the cross-correlation coefficient between our ZA control variates and power spectra of halos and galaxies measured from an N-body



Figure 8. (*Top*) The factor by which the variance of $\hat{P}_{gg,\ell}(k)$ (left), $\hat{P}_{gg}(k)$ (middle), and $\hat{P}_{gm}(k)$ (right) would be reduced using ZCV for a variety of galaxy samples, as listed in the legend. We have assumed that we can reduce the variance with ZCV up to the shot-noise limit. This assumption should hold to $k \sim 0.2 h \text{Mpc}^{-1}$, as seen in fig. 6. The number densities, redshifts and linear bias values assumed for these samples are described in the text. (*Bottom*) The fractional error on each statistic for the same galaxy samples as in the top, with (solid) and without (dashed) using ZCV, assuming a volume of $(2h^{-1}\text{Gpc})^3$. The gray shaded region roughly represents the level at which we expect systematics in simulations to become the dominant error. Thus, for most applications a volume of $(2h^{-1}\text{Gpc})^3$ is sufficient to remove statistical error from simulated measurements of two-point clustering for galaxy surveys.

simulation. This cross-correlation coefficient, ρ_{xc} , is the key statistic that determines the performance of the ZCV method, as the amount of variance reduction achieved by ZCV is given by $1 - \rho_{xc}^2$.

In section 5.1 we showed that even with the simplest unbiased ZA control variate we are able to achieve a significant cross-correlation with redshift-space halo power spectra. Adding linear and quadratic bias operators to our ZA control variate significantly improves on this performance for more biased samples, while including tidal bias offers little to no additional gain. Including satellite galaxies somewhat degrades the cross-correlation at $k > 0.2 h \text{Mpc}^{-1}$, depending on the satellite fraction of the galaxy sample in question. Nevertheless, when using a Zel'dovich control variate that incorporates linear and quadratic bias, we are able to achieve the theoretical maximum cross-correlation in the presence of shot noise out to $k \sim 0.2 h \text{Mpc}^{-1}$ for all the tracers considered in this work. Here we have focused on z = 0.55 as a fiducial redshift for our tests, but we expect that this performance will improve at higher redshift. At lower redshifts it may degrade, but the volumes probed at redshifts z < 0.5 are small enough to be simulated at relatively low cost. In section 5.3, we compared ZCV to a linear theory control variate, showing that ZCV performs significantly better for $k > 0.05 h \text{Mpc}^{-1}$

In section 5.4 we proceeded to forecast the factor by which the variance of clustering measurements made using a number of galaxy samples of relevance to upcoming cosmological measurements can be reduced using ZCV. Figure 8 shows this variance reduction or equivalently, the effective increase in simulation volume. It also shows the expected fractional error on measurements made from a hypothetical $(2h^{-1}\text{Gpc})^3$ simulation. When using ZCV, it is

clear that running simulations with volumes larger than $(2h^{-1}\text{Gpc})^3$ is unnecessary for most applications involving two-point functions of biased tracers, and their cross-correlation with matter in either real or redshift space. This has broad implications for a variety of simulation use cases. For example, ZCV will drastically reduce the cost of producing mock galaxy catalogs that have the sufficient statistical precision to test theoretical models at the accuracy required by upcoming surveys. Furthermore, ZCV will allow those wishing to run suites of simulations for the purposes of emulation to run more simulations at higher resolution than would otherwise be possible, thus making the task of interpolating between simulated cosmologies significantly simpler.

There are still a number of applications that we imagine to be similarly suitable to the ZCV method. Nearest to the current application, one can imagine reducing the variance of reconstructed power spectra or correlation functions [64, 65]. It is possible that an optimal variance reduction of such measurements may even be achieved with just linear theory, given that much of the displacement responsible for the decorrelation with the linear initial conditions is removed by the reconstruction algorithm.

Another potential application is to use ZCV to reduce the variance of measurements made in lightcone simulations [66, 67] in order to produce variance reduced measurements with realistic observational systematics. This will likely require the construction of a pipeline to produce ZA lightcones to pair with their N-body counterparts, but otherwise the same ZCV procedure presented here should be appropriate.

Given the limitations on ZCV imposed by shot noise that we have observed here, it is clear that applications to samples with low shot noise are ideal. Measurements of the power spectra of neutral hydrogen, Lyman-alpha flux, or other low shot-noise fields may be particularly suitable for the ZCV method. To what extent these measurements remain cross-correlated with ZA control variates must still be checked.

Finally, higher order correlation functions are a particularly appealing target, both because they have become a common statistic measured in surveys and because of their utility in predicting covariance matrices beyond the disconnected approximation. Going beyond the ZA may be necessary for such applications, and analytic methods for predicting these higher-order functions are numerically complicated to evaluate in LPT. Nevertheless, mean predictions for these statistics in LPT can still be computed via brute force Monte-Carlo, potentially providing a path forward. We leave these additional avenues to future work.

Acknowledgments

The authors thank Pat McDonald for useful conversations during the preparation of this work. J.D. is supported by the Lawrence Berkeley National Laboratory Chamberlain Fellowship. S.C. is supported by the Bezos Membership at the Institute for Advanced Study. N.K. is supported by the Gerald J. Lieberman Fellowship. M.W. is supported by the DOE. This research has made use of NASA's Astrophysics Data System and the arXiv preprint server. This research is supported by the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy 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. Calculations and figures in this work have been made using the SciPy Stack [68–70]. Power spectrum measurements were made with pypower.

A Exact redshift-space power spectrum in the Zel'dovich approximation

In ref. [20], the present authors derived exact expressions for the real-space power spectrum within the ZA for tracers up to quadratic order in the bias. Using the cumulant theorem for Gaussian variables, the power spectrum can be put into the form

$$P(\mathbf{k}) = \sum_{\mathcal{O}_{\alpha}, \mathcal{O}_{\beta}} b_{\mathcal{O}_{\alpha}} b_{\mathcal{O}_{\beta}} \int d^{3}\boldsymbol{q} \ e^{i\mathbf{k}\cdot\boldsymbol{q} - \frac{1}{2}k_{i}k_{j}A_{ij}(\boldsymbol{q})} \ F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}(\mathbf{k}, \boldsymbol{q})$$
(A.1)

where the sum is over the operators m, δ_0 , δ_0^2 and s_0^2 , and the correlator in the exponent $A_{ij} = \langle \Delta_i \Delta_j \rangle$ is the second moment of the pairwise displacement $\Delta = \Psi(\mathbf{q}) - \Psi(\mathbf{0})$. For example, the (δ_0^2, s_0^2) contribution is given by

$$F_{\delta_0^2 s_0^2}(\mathbf{k}, \mathbf{q}) = 2E_{ab}E_{ab} - 4k_i k_j U_i E_{ab}B_{jab} + k_i k_j k_k k_l U_i U_j B_{kab}B_{lab}$$
(A.2)

where we are using i, j, k, l, ... to denote indices on displacements Ψ_i and a, b, c, d, ... to denote indices on shears s_{ab} in correlators like U_i and B_{iab} . These correlators are functions of the Lagrangian pairwise separation \boldsymbol{q} by construction, e.g. $U_i = \langle \Psi_i(\boldsymbol{q})\delta_0(\boldsymbol{0})\rangle$ is the correlation function between the Lagrangian displacement and the initial density. A comprehensive list of $F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}$ are given in Equation A.7 of ref. [20].

The functions $F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}$, thus defined, can be decomposed into scalar products of \hat{k} and \hat{q} ; that is, defining $\mu_{\boldsymbol{q}} = \hat{k} \cdot \hat{q}$, we can write

$$F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}} = F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}^{(0)}(k,q) + F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}^{(1)}(k,q)\mu_{\boldsymbol{q}} + F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}^{(2)}(k,q)\mu_{\boldsymbol{q}}^{2} + \dots$$
(A.3)

Importantly, with the $\mu_{\boldsymbol{q}}$ dependence factored out, the coefficients $F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}^{(n)}$ are scalar functions of the vector magnitudes k, q only. Similarly, we can write by symmetry $A_{ij}(\boldsymbol{q}) = X(q)\delta_{ij} + Y(q)\hat{q}_i\hat{q}_j$. The angular dependence of Equation A.1 can then be explicitly separated out as

$$P(\mathbf{k}) = \sum_{\mathcal{O}_{\alpha}, \mathcal{O}_{\beta}} b_{\mathcal{O}_{\alpha}} b_{\mathcal{O}_{\beta}} \sum_{n} \int dq \ q^{2} \ K_{n}(k, q) \ e^{-\frac{1}{2}k^{2}X(q)} \ F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}^{(n)}(k, q).$$
(A.4)

The angular kernel is given by

$$K_{n}(q) = \int d\phi \ d\mu_{q} \ e^{ikq\mu_{q} - \frac{1}{2}k^{2}Y(q)\mu_{q}^{2}} \ \mu_{q}^{n}$$
(A.5)

independent of $\mathcal{O}_{\alpha}, \mathcal{O}_{\beta}$ and can be expressed as a series of spherical bessel functions [71, 72], and the spherical coordinates are set up such that \hat{k} is at the zenith.³.

As described in Section 3, the only difference incurred by moving to redshift space is that all the displacements must now be multiplied by the matrix \mathbf{R} or, equivalently, since all displacements are dotted with wavevectors, all wavevectors other than the one in the Fourier transform $(\mathbf{k} \cdot \mathbf{q})$ must be transformed into $K_i = R_{ij}k_j$, such that now [38, 39]

$$P_{s}(\mathbf{k}) = \sum_{\mathcal{O}_{\alpha}, \mathcal{O}_{\beta}} b_{\mathcal{O}_{\alpha}} b_{\mathcal{O}_{\beta}} \int d^{3}\boldsymbol{q} \ e^{i\mathbf{k}\cdot\boldsymbol{q} - \frac{1}{2}K_{i}K_{j}A_{ij}(\boldsymbol{q})} F_{\mathcal{O}_{\alpha}\mathcal{O}_{\beta}}(\mathbf{K}, \boldsymbol{q}).$$
(A.6)

³See Equation F.1 of ref. [73] for an exact expression.

In this case it is sensible to choose a new system of spherical coordinates with \hat{K} at the zenith; then, the angular dependence can be captured simply by replacing the real-space kernel in Equation A.4 with

$$K_{n,s}(q) = \int d\phi \ d\mu_{\boldsymbol{q}} \ e^{i\mathbf{k}\cdot\boldsymbol{q} - \frac{1}{2}K^2Y(q)\mu_{\boldsymbol{q}}^2} \ \mu_{\boldsymbol{q}}^n.$$
(A.7)

The evaluation of this kernel was performed explicitly in the above references, and the reader is directed to them for further details.

B Accounting for incomplete μ sampling

Measurements of power spectrum multipoles in a periodic box with finite volume incur errors due to mixing between multipoles. This mixing is caused by the sparse sampling in μ imposed by the discrete nature of the wave modes sampled by the mesh used to measure the simulated power spectra. These artifacts are most significant for large scales and multipoles with $\ell > 0$, where the Legendre polynomials become highly oscillatory and so fine μ sampling becomes more important. This is analogous to the mixing between multipoles (and wavenumbers) caused by incomplete sky coverage in observational data [74], often referred to as the survey window function effect.

To see how to correct for this effect, we can write our power spectra as

$$\hat{P}_{\ell}(k_i) = N_i^{-1} \sum_{\mathbf{k}} B_i(k) \mathcal{L}_{\ell}(\mu) |\delta(\mathbf{k})|^2$$
(B.1)

$$= N_i^{-1} \sum_{\mathbf{k},\ell'} B_i(k) \mathcal{L}_\ell(\mu) \mathcal{L}_{\ell'}(\mu) P_{\ell'}(k)$$
(B.2)

where $k = |\mathbf{k}|$, and $\mu = \hat{\mathbf{n}} \cdot \hat{\mathbf{k}}$, with $\hat{\mathbf{n}}$ the line-of-sight unit vector. $B_i(k)$ is a top-hat around k_i , and $N_i = \sum_{\mathbf{k}} B_i(k)$. The response of this estimator to a change in the uncoupled, noiseless theory, $P_{\ell'}(k')$, is:

$$W_{\ell,\ell'}^{i} = \frac{\partial \langle P_{\ell}(k_i) \rangle}{\partial P_{\ell'}(k')} = N_i^{-1} \sum_{\mathbf{k}} B_i(k) \mathcal{L}_{\ell}(\mu) \mathcal{L}_{\ell'}(\mu) \delta^D(k-k') \,. \tag{B.3}$$

With this window matrix, we can correctly account for the simulation window function when making theory predictions. Alternatively, by assuming that $P_{\ell'}(k')$ is constant over the width of the bin, k_i , then we have

$$\bar{W}^{i}_{\ell,\ell'} = \frac{\partial \langle P_{\ell}(k_i) \rangle}{\partial P_{\ell'}(k_i)} = N_i^{-1} \sum_{\mathbf{k}} B_i(k) \mathcal{L}_{\ell}(\mu) \mathcal{L}_{\ell'}(\mu) , \qquad (B.4)$$

where $P_{\ell'}(k_i) = N_i^{-1} \sum_{k'} B_i(k') P_{\ell'}(k')$. Then we have

$$P_{\ell}^{*}(k_{i}) = (\bar{W}^{-1})_{\ell,\ell'}^{i} \hat{P}_{\ell'}(k_{i}), \qquad (B.5)$$

where repeated indices are summed over. $\hat{P}_{\ell}^*(k_i)$ is now an unbiased estimator of $\hat{P}_{\ell}(k_i)$ to the extent that our piecewise constant approximation holds and we have included a sufficient number of multipoles in eq. B.5.

Whether one should deconvolve this effect from the simulation measurements, or convolve the theory with the window matrix depends on the application under consideration. If the simulation measurements are binned finely in k, then deconvolution is appropriate, because the assumption that the theory is piecewise constant likely holds to a good approximation. When one wishes to bin more coarsely in k, then convolving a finely sampled theory prediction with the window is more appropriate. In this work we deconvolve the window from all the multipole measurements made, in order to ensure that the ensemble mean of our Zel'dovich realizations is equal to the noiseless, uncoupled mean predicted by our analytic ZA model.

C Estimating β and ρ_{xc}

We have seen that accurate estimates of β and ρ_{xc} are of significant importance, β for obtaining an optimal control variate estimator, and ρ_{xc} for accurately predicting the variance reduction factor provided by said control variate. We can see from eqs. 2.1 and 5.1 that β and ρ_{xc} both rely on accurate estimates of the variance of the control variate, Var[C], and the covariance between the control variate and the statistic whose variance we wish to reduce, Cov[X, C]. ρ_{xc} additionally depends on the variance of the latter, Var[X].

In this work, we have employed the disconnected approximation to estimate covariances, using power spectra measured from our simulations where required. For example, when measuring β for real space power spectra, our estimator reads:

$$\beta = \frac{P^{tz}(k)^2}{P^{zz}(k)^2} \tag{C.1}$$

where $\hat{P}^{tz}(k)$ is the measured cross-power spectrum between the tracer in question and our ZA control variate, and $\hat{P}^{zz}(k)$ is the auto-power spectrum of the ZA control variate. Because the tracer field and Zel'dovich field come from the same initial noise realization, the variance of the two spectra largely cancels out, yielding a low-noise estimate of β .

For multipoles of the redshift-space cross-power spectrum between two fields a and b, this disconnected covariance reads

$$\operatorname{Cov}\left[P_{\ell}^{aa}(k), P_{\ell'}^{bb}(k)\right] = \frac{(2\ell+1)(2\ell'+1)}{V_{\text{obs}}} \frac{1}{4\pi k^2 \Delta k} \int d\mu \ \mathcal{L}_{\ell}(\mu) \mathcal{L}_{\ell'}(\mu) \ P^{ab}(k,\mu)^2 \qquad (C.2)$$

Avoiding the full generality of Wigner-3*j* symbols, we can approximate $P^{ab}(k,\mu) = P_0\mathcal{L}_0 + P_2\mathcal{L}_2 + P_4\mathcal{L}_4$ in order to compute the necessary covariances. Thus, our estimates of β and ρ_{xc} depend on the maximum ℓ that we include in our multipole expansion of $P^{ab}(k,\mu)$. As a cross-check that we are not particularly sensitive to our chosen $\ell_{\max} = 4$, we have re-estimated β and ρ_{xc} setting $\ell_{\max} = 2$ for the "2x sats." HOD depicted in fig. 6. This should be a worst-case scenario for the tracers considered in this work, as it exhibits the most significant hexadecapole moment. As seen in the bottom two panels of fig. 9, for $k < 0.4 h \text{Mpc}^{-1}$ the impact is consistent with zero. For $k > 0.4 h \text{Mpc}^{-1}$ we observe $\sim 10 - 20\%$ effects, but these are mitigated by the fact that we apply a damping factor to β that we discuss presently.

The disconnected approximation that we assume for covariances breaks down when nonlinear mode coupling becomes significant [75]. This leads to a significant overestimation of β , compared to the value that we would measure if we had access to the true covariance of the measurements in question. In order to mitigate this and avoid adding extra variance to our simulations at high-k where our ZA surrogate is no longer strongly correlated, we damp β to zero by applying a tanh function:

$$F(k;k_0,\Delta_k) = \frac{1}{2} \left[1 - \tanh\left(\frac{k-k_0}{\Delta_k}\right) \right].$$
(C.3)

using $k_0 = 0.618 h \text{Mpc}^{-1}$ and $\Delta_k = 0.167 h \text{Mpc}^{-1}$, which were determined by [20] to bring β into agreement with measurement made using an ensemble of simulations. In the future we will redetermine these constants for specific use cases as needed. The upshot of this damping is that we set $\beta = 0$ for nearly the same k range as where we observe significant variation when varying ℓ_{max} .

There is also a subtle bias that enters into the control variate estimator when reducing the variance of measurements from the same simulation that is used to estimate β . Using real-space auto-power spectra as an example, this can be seen by the fact that the ensemble average

$$\langle y \rangle = \langle \hat{P}^{tt}(k) - \hat{\beta}(\hat{P}^{zz}(k) - P^{zz}(k)) \rangle \tag{C.4}$$

$$= \left\langle \hat{P}^{tt}(k) - \frac{\hat{P}^{tz}(k)^{2}}{\hat{P}^{zz}(k)^{2}} (\hat{P}^{zz}(k) - P^{zz}(k)) \right\rangle$$
(C.5)

$$= \left\langle \hat{P}^{tt}(k) - \left(\frac{\hat{P}^{tz}(k)^2}{\hat{P}^{zz}(k)} - \frac{\hat{P}^{tz}(k)^2}{\hat{P}^{zz}(k)^2} P^{zz}(k)\right) \right\rangle \neq \left\langle \hat{P}^{tt}(k) \right\rangle$$
(C.6)

where hatted variables are measurements from individual realizations, and the ensemble average is taken over many realizations. In order to mitigate this bias, we smooth our estimates of β using a third-order Savitsky-Golay filter with a window length of 21. For each k bin, this filter fits a third order polynomial to the 20 adjacent k bins. Thus it is possible to calculate that the effective number of points that each smoothed point receives contributions from is 9.3. In the regime where ZCV yields significant gains, these points are largely uncorrelated. Thus, this process serves to suppress any small correlation between β and $\hat{P}^{zz}(k)$ that might prevent $\hat{\beta}(\hat{P}^{zz}(k) - P^{zz}(k)))$ from averaging to zero.⁴ Furthermore, since $\hat{P}^{tz,tt}$ are highly

$$\langle y_i \rangle = \sigma - c.$$
 (C.7)

Now, suppose we instead estimate the correlation by taking the mean across bins, i.e. $\beta = \frac{1}{N} \sum_{i} \beta_{i}$. In this case we can see that

$$y'_i = x_i - \beta \mu_i, \quad \langle y_i \rangle = \sigma - \frac{c}{N},$$
 (C.8)

i.e. the bias is suppressed by the number of samples. Moreover, we can see that in the case that we are interested in the *overall* amplitude of the x_i (e.g. σ_8) that

$$Y = \frac{1}{N} \sum_{i} y_{i}, \quad \langle Y \rangle = \sigma - \frac{c}{N}$$
(C.9)

similarly has a suppressed bias. It is worth noting however that the additional correlation between β and μ_i will increase the noise of y'_i , i.e.

$$\langle y'_i y'_i \rangle_c \ni \langle \beta \mu_i \rangle^2 = \frac{c^2}{N^2}$$
 (C.10)

⁴Let us consider the effect of smoothing within a simple toy model. Consider a set of N random variables x_i with associated control variates μ_i with zero mean, as well as estimated β_i , such that we can form the variance-reduced combination $y_i = x_i - \beta_i \mu_i$. Here we imagine x_i is the power spectrum in bin *i*, so for simplicity let us assume $\langle x_i \rangle = \sigma$ and that different bins are uncorrelated; in particular, let us assume that β_i and μ_i can be correlated, as discussed above, such that $\langle \beta_i \mu_j \rangle = c \, \delta_{ij}$. In this case we see that y_i is biased:



Figure 9. (Top) A depiction of our procedure for estimating β from our simulations, in the case of the "2x sats." HOD shown in fig. 6. The raw measurements assuming a disconnected covariance approximation are shown in blue. We then apply a damping function (eq. C.3) to these measurements in order to take β to 0 where we expect this disconnected approximation to break down. We additionally smooth β with a Savitsky-Golay filter in order to debias our control variate estimator. (*Middle*) The fractional error that is made in β by not including the hexadecapole in our covariance estimates. This error is negligible in the regime that we do not damp β to zero. Note, that we include the hexadecapole for the rest of this work, and that we expect inclusion of even higher multipoles in our covariance to have even less of an impact. (*Bottom*) Same as the middle panel but for ρ_{xc} . Again, the error incurred is negligible.

correlated, their ratio β should be significantly less correlated with \hat{P}^{zz} than each individually, suppressing the bias due to correlations of the control variate with β . This is why the ρ_{xc} 's shown in this paper are far less noisy than the spectra themselves, with noise primarily due to stochasticity on large scales. The process of smoothing and damping β described here is depicted in the top panel of fig. 9.

Alternatively, if one wishes to avoid these complications and set β to some anzatz, e.g. $\beta = 1$, as it must be on large scales when including linear and quadratic bias, then the variance of y is

$$\operatorname{Var}[Y|\beta] = \operatorname{Var}[X] + \beta^2 \operatorname{Var}[C] - 2\beta \rho_{xc} \sigma_x \sigma_c = \sigma_c^2 (\beta - \beta^*)^2 + \operatorname{Var}[Y|\beta^*]$$

and is thus degraded from the optimal choice of $\beta = \beta^*$ by

$$\frac{\operatorname{Var}[Y|\beta]}{\operatorname{Var}[Y|\beta^*]} = \frac{1 + \beta^2 \frac{\operatorname{Var}[C]}{\operatorname{Var}[X]} - 2\beta \rho_{xc} \frac{\sigma_c}{\sigma_x}}{1 - \rho_{xc}^2} = 1 + \frac{\sigma_c^2}{\sigma_x^2} \frac{(\beta - \beta^*)^2}{1 - \rho^2} \,.$$

If we take the LRG-like HOD as an example, we find $\operatorname{Var}[Y|\beta = 1]/\operatorname{Var}[Y|\beta^*] \simeq 1.05$ at $k \simeq 0.2 \, h \mathrm{Mpc}^{-1}$. Thus, one can eliminate the need to estimate β , apply smoothing, etc., by setting $\beta = 1$ without incurring a large penalty in variance, although there is no guarantee of this fact. This simplifies the pipeline and saves approximately 25% of the computing cost required by ZCV as there is no longer a need to compute cross correlations between the Zel'dovich and tracer fields.

however this too is suppressed by N and, in our case, further by the smallness of the correlation between β and μ as argued in the text.

References

- V. Springel, R. Pakmor, O. Zier and M. Reinecke, Simulating cosmic structure formation with the GADGET-4 code, Mon. Not. R. Astron. Soc. 506 (2021) 2871 [2010.03567].
- [2] L.H. Garrison, D.J. Eisenstein, D. Ferrer, N.A. Maksimova and P.A. Pinto, *The ABACUS cosmological N-body code*, Mon. Not. R. Astron. Soc. 508 (2021) 575 [2110.11392].
- [3] V. Springel, E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh, Mon. Not. R. Astron. Soc. 401 (2010) 791 [0901.4107].
- [4] J.A. Peacock and S.J. Dodds, Non-linear evolution of cosmological power spectra, Mon. Not. R. Astron. Soc. 280 (1996) L19 [astro-ph/9603031].
- [5] R.E. Angulo and A. Pontzen, Cosmological N-body simulations with suppressed variance, Mon. Not. R. Astron. Soc. 462 (2016) L1 [1603.05253].
- [6] A. Pontzen, A. Slosar, N. Roth and H.V. Peiris, Inverted initial conditions: Exploring the growth of cosmic structure and voids, Phys. Rev. D 93 (2016) 103519 [1511.04090].
- [7] S. Avila and A. Gutierrez Adame, Validating galaxy clustering models with Fixed & Paired and Matched-ICs simulations: application to Primordial Non-Gaussianities, arXiv e-prints (2022) arXiv:2204.11103 [2204.11103].
- [8] F. Villaescusa-Navarro, S. Naess, S. Genel, A. Pontzen, B. Wandelt, L. Anderson et al., Statistical Properties of Paired Fixed Fields, Astrophys. J. 867 (2018) 137 [1806.01871].
- [9] C.-H. Chuang, G. Yepes, F.-S. Kitaura, M. Pellejero-Ibanez, S. Rodríguez-Torres, Y. Feng et al., UNIT project: Universe N-body simulations for the Investigation of Theoretical models from galaxy surveys, Mon. Not. R. Astron. Soc. 487 (2019) 48 [1811.02111].
- [10] L. Anderson, A. Pontzen, A. Font-Ribera, F. Villaescusa-Navarro, K.K. Rogers and S. Genel, Cosmological Hydrodynamic Simulations with Suppressed Variance in the Lyα Forest Power Spectrum, Astrophys. J. 871 (2019) 144.
- [11] A. Klypin, F. Prada and J. Byun, Suppressing cosmic variance with paired-and-fixed cosmological simulations: average properties and covariances of dark matter clustering statistics, Mon. Not. R. Astron. Soc. 496 (2020) 3862 [1903.08518].
- [12] Y. Zhang, A.R. Pullen, S. Alam, S. Singh, E. Burtin, C.-H. Chuang et al., Testing general relativity on cosmological scales at redshift z ~ 1.5 with quasar and CMB lensing, Mon. Not. R. Astron. Soc. 501 (2021) 1013 [2007.12607].
- [13] F. Maion, R.E. Angulo and M. Zennaro, Statistics of biased tracers in variance-suppressed simulations, Journal of Cosmology and Astro-Particle Physics 2022 (2022) 036 [2204.03868].
- [14] A.B. Owen, Monte Carlo theory, methods and examples (2013).
- [15] N. Chartier, B. Wandelt, Y. Akrami and F. Villaescusa-Navarro, CARPool: fast, accurate computation of large-scale structure statistics by pairing costly and cheap cosmological simulations, Mon. Not. R. Astron. Soc. 503 (2021) 1897 [2009.08970].
- [16] Z. Ding, C.-H. Chuang, Y. Yu, L.H. Garrison, A.E. Bayer, Y. Feng et al., The DESI N-body Simulation Project - II. Suppressing sample variance with fast simulations, Mon. Not. R. Astron. Soc. 514 (2022) 3308 [2202.06074].
- [17] N. Chartier and B.D. Wandelt, CARPool covariance: fast, unbiased covariance estimation for large-scale structure observables, Mon. Not. R. Astron. Soc. 509 (2022) 2220 [2106.11718].
- [18] N. Chartier and B.D. Wandelt, Bayesian control variates for optimal covariance estimation with pairs of simulations and surrogates, Mon. Not. R. Astron. Soc. 515 (2022) 1296 [2204.03070].
- [19] Y. Feng, M.-Y. Chu, U. Seljak and P. McDonald, FASTPM: a new scheme for fast simulations of dark matter and haloes, MNRAS 463 (2016) 2273 [1603.00476].

- [20] N. Kokron, S.-F. Chen, M. White, J. DeRose and M. Maus, Accurate predictions from small boxes: variance suppression via the Zel'dovich approximation, Journal of Cosmology and Astro-Particle Physics 2022 (2022) 059 [2205.15327].
- [21] Y.B. Zel'Dovich, Gravitational instability: an approximate theory for large density perturbations., Astronomy & Astrophysics 500 (1970) 13.
- [22] M. White, The Zel'dovich approximation, Mon. Not. R. Astron. Soc. 439 (2014) 3630 [1401.5466].
- [23] C. Modi, S.-F. Chen and M. White, Simulations and symmetries, Mon. Not. R. Astron. Soc. 492 (2020) 5754 [1910.07097].
- [24] N. Kokron, J. DeRose, S.-F. Chen, M. White and R.H. Wechsler, The cosmology dependence of galaxy clustering and lensing from a hybrid N-body-perturbation theory model, Mon. Not. R. Astron. Soc. 505 (2021) 1422 [2101.11014].
- [25] M. Zennaro, R.E. Angulo, M. Pellejero-Ibáñez, J. Stücker, S. Contreras and G. Aricò, The BACCO simulation project: biased tracers in real space, arXiv e-prints (2021) arXiv:2101.12187 [2101.12187].
- [26] A. Schneider, R. Teyssier, D. Potter, J. Stadel, J. Onions, D.S. Reed et al., Matter power spectrum and the challenge of percent accuracy, Journal of Cosmology and Astro-Particle Physics 2016 (2016) 047 [1503.05920].
- [27] L.H. Garrison, M. Joyce and D.J. Eisenstein, Good and proper: self-similarity of N-body simulations with proper force softening, Mon. Not. R. Astron. Soc. 504 (2021) 3550 [2102.08972].
- [28] C. Grove, C.-H. Chuang, N.C. Devi, L. Garrison, B. L'Huillier, Y. Feng et al., The DESI N-body simulation project - I. Testing the robustness of simulations for the DESI dark time survey, Mon. Not. R. Astron. Soc. 515 (2022) 1854 [2112.09138].
- [29] T. Matsubara, Nonlinear perturbation theory with halo bias and redshift-space distortions via the Lagrangian picture, Phys. Rev. D 78 (2008) 083519 [0807.1733].
- [30] V. Desjacques, D. Jeong and F. Schmidt, *Large-scale galaxy bias*, Phys. Rep. **733** (2018) 1 [1611.09787].
- [31] N.E. Chisari and A. Pontzen, Unequal time correlators and the Zel'dovich approximation, Phys. Rev. D 100 (2019) 023543 [1905.02078].
- [32] T. Matsubara, Resumming cosmological perturbations via the Lagrangian picture: One-loop results in real space and in redshift space, Phys. Rev. D 77 (2008) 063530 [0711.2521].
- [33] J. Carlson, B. Reid and M. White, Convolution Lagrangian perturbation theory for biased tracers, Mon. Not. R. Astron. Soc. 429 (2013) 1674 [1209.0780].
- [34] J. Lesgourgues, The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview, arXiv e-prints (2011) arXiv:1104.2932 [1104.2932].
- [35] D. Blas, J. Lesgourgues and T. Tram, The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes, Journal of Cosmology and Astro-Particle Physics 2011 (2011) 034 [1104.2933].
- [36] M. Michaux, O. Hahn, C. Rampf and R.E. Angulo, Accurate initial conditions for cosmological N-body simulations: minimizing truncation and discreteness errors, Mon. Not. R. Astron. Soc. 500 (2021) 663 [2008.09588].
- [37] A.N. Taylor and A.J.S. Hamilton, Non-linear cosmological power spectra in real and redshift space, Mon. Not. R. Astron. Soc. 282 (1996) 767 [astro-ph/9604020].
- [38] Z. Vlah and M. White, Exploring redshift-space distortions in large-scale structure, Journal of Cosmology and Astro-Particle Physics 2019 (2019) 007 [1812.02775].

- [39] S.-F. Chen, Z. Vlah and M. White, The reconstructed power spectrum in the Zeldovich approximation, Journal of Cosmology and Astro-Particle Physics 2019 (2019) 017 [1907.00043].
- [40] S.-F. Chen, Z. Vlah, E. Castorina and M. White, Redshift-space distortions in Lagrangian perturbation theory, Journal of Cosmology and Astro-Particle Physics 2021 (2021) 100 [2012.04636].
- [41] J. DeRose and the Aemulus Project, The Aemulus ν Simulations, in prep. (2022).
- [42] J. DeRose, R.H. Wechsler, J.L. Tinker, M.R. Becker, Y.-Y. Mao, T. McClintock et al., The AEMULUS Project. I. Numerical Simulations for Precision Cosmology, Astrophys. J. 875 (2019) 69 [1804.05865].
- [43] W. Elbers, C.S. Frenk, A. Jenkins, B. Li and S. Pascoli, Higher order initial conditions with massive neutrinos, Mon. Not. R. Astron. Soc. 516 (2022) 3821 [2202.00670].
- [44] P.S. Behroozi, R.H. Wechsler and C. Conroy, The Average Star Formation Histories of Galaxies in Dark Matter Halos from z = 0-8, Astrophys. J. 770 (2013) 57 [1207.6105].
- [45] U. Seljak, Analytic model for galaxy and dark matter clustering, Mon. Not. R. Astron. Soc. 318 (2000) 203 [astro-ph/0001493].
- [46] A.A. Berlind and D.H. Weinberg, The Halo Occupation Distribution: Toward an Empirical Determination of the Relation between Galaxies and Mass, Astrophys. J. 575 (2002) 587 [astro-ph/0109001].
- [47] Z. Zheng, A.L. Coil and I. Zehavi, Galaxy Evolution from Halo Occupation Distribution Modeling of DEEP2 and SDSS Galaxy Clustering, Astrophys. J. 667 (2007) 760 [astro-ph/0703457].
- [48] R. Zhou, J.A. Newman, Y.-Y. Mao, A. Meisner, J. Moustakas, A.D. Myers et al., The clustering of DESI-like luminous red galaxies using photometric redshifts, Mon. Not. R. Astron. Soc. 501 (2021) 3309 [2001.06018].
- [49] D. Wadekar and R. Scoccimarro, Galaxy power spectrum multipoles covariance in perturbation theory, Phys. Rev. D 102 (2020) 123517 [1910.02914].
- [50] T. Baldauf, U. Seljak, R.E. Smith, N. Hamaus and V. Desjacques, Halo stochasticity from exclusion and nonlinear clustering, Phys. Rev. D 88 (2013) 083507 [1305.2917].
- [51] N. Hamaus, U. Seljak, V. Desjacques, R.E. Smith and T. Baldauf, *Minimizing the stochasticity of halos in large-scale structure surveys*, *Phys. Rev. D* 82 (2010) 043515 [1004.5377].
- [52] N. Kokron, J. DeRose, S.-F. Chen, M. White and R.H. Wechsler, Priors on red galaxy stochasticity from hybrid effective field theory, Mon. Not. R. Astron. Soc. 514 (2022) 2198 [2112.00012].
- [53] N. Kaiser, Clustering in real space and in redshift space, Mon. Not. R. Astron. Soc. 227 (1987)
 1.
- [54] DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L.E. Allen et al., The DESI Experiment Part I: Science, Targeting, and Survey Design, arXiv e-prints (2016) arXiv:1611.00036 [1611.00036].
- [55] D.J. Schlegel, S. Ferraro, G. Aldering, C. Baltay, S. BenZvi, R. Besuner et al., A Spectroscopic Road Map for Cosmic Frontier: DESI, DESI-II, Stage-5, arXiv e-prints (2022) arXiv:2209.03585 [2209.03585].
- [56] D.J. Schlegel, J.A. Kollmeier, G. Aldering, S. Bailey, C. Baltay, C. Bebek et al., The MegaMapper: A Stage-5 Spectroscopic Instrument Concept for the Study of Inflation and Dark Energy, arXiv e-prints (2022) arXiv:2209.04322 [2209.04322].

- [57] D.J. Schlegel, J.A. Kollmeier, G. Aldering, S. Bailey, C. Baltay, C. Bebek et al., The MegaMapper: A Stage-5 Spectroscopic Instrument Concept for the Study of Inflation and Dark Energy, arXiv e-prints (2022) arXiv:2209.04322 [2209.04322].
- [58] I. Zehavi, Z. Zheng, D.H. Weinberg, M.R. Blanton, N.A. Bahcall, A.A. Berlind et al., Galaxy Clustering in the Completed SDSS Redshift Survey: The Dependence on Color and Luminosity, Astrophys. J. 736 (2011) 59 [1005.2413].
- [59] R. Zhou, B. Dey, J.A. Newman, D.J. Eisenstein, K. Dawson, S. Bailey et al., Target Selection and Validation of DESI Luminous Red Galaxies, arXiv e-prints (2022) arXiv:2208.08515 [2208.08515].
- [60] A. Raichoor, J. Moustakas, J.A. Newman, T. Karim, S. Ahlen, S. Alam et al., Target Selection and Validation of DESI Emission Line Galaxies, arXiv e-prints (2022) arXiv:2208.08513 [2208.08513].
- [61] E. Chaussidon, C. Yèche, N. Palanque-Delabrouille, D.M. Alexander, J. Yang, S. Ahlen et al., Target Selection and Validation of DESI Quasars, arXiv e-prints (2022) arXiv:2208.08511 [2208.08511].
- [62] A.A. Khostovan, D. Sobral, B. Mobasher, J. Matthee, R.K. Cochrane, N. Chartab et al., The clustering of typical Ly α emitters from z ~ 2.5-6: host halo masses depend on Ly α and UV luminosities, Mon. Not. R. Astron. Soc. 489 (2019) 555 [1811.00556].
- [63] T. Nishimichi, G. D'Amico, M.M. Ivanov, L. Senatore, M. Simonović, M. Takada et al., Blinded challenge for precision cosmology with large-scale structure: Results from effective field theory for the redshift-space galaxy power spectrum, Phys. Rev. D 102 (2020) 123541 [2003.08277].
- [64] D.J. Eisenstein, H.-J. Seo, E. Sirko and D.N. Spergel, Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak, Astrophys. J. 664 (2007) 675 [astro-ph/0604362].
- [65] N. Padmanabhan, M. White and J.D. Cohn, Reconstructing baryon oscillations: A Lagrangian theory perspective, Phys. Rev. D 79 (2009) 063523 [0812.2905].
- [66] P. Fosalba, M. Crocce, E. Gaztañaga and F.J. Castander, The MICE grand challenge lightcone simulation - I. Dark matter clustering, Mon. Not. R. Astron. Soc. 448 (2015) 2987 [1312.1707].
- [67] J. DeRose, R.H. Wechsler, M.R. Becker, M.T. Busha, E.S. Rykoff, N. MacCrann et al., The Buzzard Flock: Dark Energy Survey Synthetic Sky Catalogs, arXiv e-prints (2019) arXiv:1901.02401 [1901.02401].
- [68] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau et al., Array programming with NumPy, Nature 585 (2020) 357 [2006.10256].
- [69] P. Virtanen, R. Gommers, T.E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau et al., SciPy 1.0: fundamental algorithms for scientific computing in Python, Nature Methods 17 (2020) 261 [1907.10121].
- [70] J.D. Hunter, Matplotlib: A 2D Graphics Environment, Computing in Science and Engineering 9 (2007) 90.
- [71] Z. Vlah, U. Seljak and T. Baldauf, Lagrangian perturbation theory at one loop order: Successes, failures, and improvements, Phys. Rev. D 91 (2015) 023508 [1410.1617].
- [72] Z. Vlah, E. Castorina and M. White, The Gaussian streaming model and convolution Lagrangian effective field theory, Journal of Cosmology and Astro-Particle Physics 2016 (2016) 007 [1609.02908].
- [73] S.-F. Chen, Z. Vlah and M. White, Consistent modeling of velocity statistics and redshift-space distortions in one-loop perturbation theory, Journal of Cosmology and Astro-Particle Physics 2020 (2020) 062 [2005.00523].

- [74] F. Beutler and P. McDonald, Unified galaxy power spectrum measurements from 6dFGS, BOSS, and eBOSS, Journal of Cosmology and Astro-Particle Physics **2021** (2021) 031 [2106.06324].
- [75] A.J.S. Hamilton, C.D. Rimes and R. Scoccimarro, On measuring the covariance matrix of the non-linear power spectrum from simulations, Mon. Not. R. Astron. Soc. 371 (2006) 1188 [astro-ph/0511416].