Title

Covariance matrices for variance-suppressed simulations

Journal

Monthly Notices of the Royal Astronomical Society, 518(3)

ISSN

0035-8711

Authors

Zhang, Tony

Chuang, Chia-Hsun

Wechsler, Risa H

et al.

Publication Date

2022-12-02

DOI

10.1093/mnras/stac3261

MNRAS 000, 1–9 (2021) Preprint 21 March 2023 Compiled using MNRAS LATEX style file v3.0

Covariance matrices for variance-suppressed simulations

Tony Zhang,[1]★ Chia-Hsun Chuang,[2,3]† Risa H. Wechsler,[1,2,4] Shadab Alam,[5] Joseph DeRose,[6] Yu Feng,[7] Francisco-Shu Kitaura,[8,9] Marcos Pellejero-Ibanez,[10] Sergio Rodríguez-Torres,[11] Chun-Hao To,[12] Gustavo Yepes,[11,13] Cheng Zhao,[14]

[1]*Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA*
[2]*Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, CA 94305, USA*
[3]*Department of Physics and Astronomy, University of Utah, Salt Lake City, UT 84112, USA*
[4]*SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, CA 94025, USA*
[5]*Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK*
[6]*Physics Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*
[7]*Berkeley Center for Cosmological Physics, Department of Physics, University of California Berkeley, Berkeley, CA 94720, USA*
[8]*Instituto de Astrofísica de Canarias (IAC), C/Vía Láctea, s/n, E-38200 La Laguna, Tenerife, Spain*
[9]*Departamento Astrofísica, Universidad de La Laguna (ULL), E-38206 La Laguna, Tenerife, Spain*
[10] *Donostia International Physics Center (DIPC), Paseo Manuel de Lardizabal, 4, E-20018 Donostia-San Sebastián, Spain*
[11]*Departamento de Física Teórica, Módulo 8, Facultad de Ciencias, Universidad Autónoma de Madrid, E-28049 Madrid, Spain*
[12]*Center for Cosmology and AstroParticle Physics, Ohio State University, Columbus, OH 43210, USA*
[13]*CIAFF, Facultad de Ciencias, Universidad Autónoma de Madrid, E-28049 Madrid, Spain*
[14]*Institute of Physics, Laboratory of Astrophysics, Ecole Polytechnique Fédérale de Lausanne (EPFL), Observatoire de Sauverny, CH-1290 Versoix, Switzerland*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Cosmological N-body simulations provide numerical predictions of the structure of the Universe against which to compare data from ongoing and future surveys, but the growing volume of the Universe mapped by surveys requires correspondingly lower statistical uncertainties in simulations, usually achieved by increasing simulation sizes at the expense of computational power. It was recently proposed to reduce simulation variance without incurring additional computational costs by adopting fixed-amplitude initial conditions. This method has been demonstrated not to introduce bias in various statistics, including the two-point statistics of galaxy samples typically used for extracting cosmological parameters from galaxy redshift survey data, but requires us to revisit current methods for estimating covariance matrices of clustering statistics for simulations. In this work, we find that it is not trivial to construct covariance matrices analytically for fixed-amplitude simulations, but we demonstrate that EZMOCK (Effective Zel'dovich approximation mock catalogue), the most efficient method for constructing mock catalogues with accurate two- and three-point statistics, provides reasonable covariance matrix estimates for such simulations. We further examine how the variance suppression obtained by amplitude-fixing depends on three-point clustering, small-scale clustering, and galaxy bias, and propose intuitive explanations for the effects we observe based on the EZMOCK bias model.

Key words: methods: numerical – galaxies: haloes – large-scale structure of Universe – cosmology: theory – software: simulations

1 INTRODUCTION

The study of the large-scale structure of the Universe has become a precision science in recent years, thanks to such surveys as the photometric DES[1] (Dark Energy Survey, Dark Energy Survey Collaboration 2005) and the spectroscopic SDSS[2] (Sloan Digital Sky Survey, Blanton et al. 2017). The total volume of the Universe mapped with galaxy surveys will continue to increase as the next generation of large ground- and space-based experiments comes online, including DESI[3] (Dark Energy Spectroscopic Instrument;

★ E-mail: txz@stanford.edu
† E-mail: chuangch@stanford.edu

[1] http://www.darkenergysurvey.org

[2] http://www.sdss.org/sdss-surveys

[3] http://desi.lbl.gov/

Schlegel et al. 2011; Levi et al. 2013), 4MOST[4] (4-metre Multi-Object Spectroscopic Telescope; de Jong et al. 2012), HETDEX[5] (Hobby–Eberly Telescope Dark Energy Experiment; Hill et al. 2008), J-PAS[6] (Javalambre Physics of the Accelerating Universe Astrophysical Survey; Benitez et al. 2014), PFS[7](Subaru Prime Focus Spectrograph; Takada et al. 2014), LSST[8] (Legacy Survey of Space and Time; LSST Science Collaboration, 2009), Euclid[9] (Laureijs et al. 2011), and the Roman Space Telescope[10] (Spergel et al. 2013).

Theoretical modelling of galaxy clustering, including galaxy bias and peculiar motions, is crucial for testing cosmological models against observations. On large scales, these theoretical models are often based on perturbation theory with approximations, and their validation requires fully non-linear numerical solutions, generally in the form of N-body simulations tracking the full growth of structure over the Universe's history. These simulations must be sufficiently large compared to the volumes sampled in surveys while maintaining sufficient mass resolution to resolve the dark matter haloes that host the galaxies typically detected in surveys. As survey volumes continue to grow, so too must the size of these N-body simulations.

These dual requirements of simulation volume and mass resolution are difficult to meet with current computational power. Indeed, a single simulation with the required halo mass resolution ($\sim 1 \times 10^{11}\, M_\odot\, h^{-1}$, Gonzalez-Perez et al. 2017) covering the entire survey volume of DESI [$\sim 70\,(\mathrm{Gpc}\, h^{-1})^3$] would demand an enormous number of particles ($\gtrsim 16\,000^3$ in a box of side length $4\,\mathrm{Gpc}\, h^{-1}$). However, the largest N-body simulations to date, e.g. MillenniumXXL (Angulo et al. 2012), MICE (Fosalba et al. 2015), MultiDark (Klypin et al. 2016), Dark Sky (Skillman et al. 2014), OuterRim (Habib et al. 2016), FLAGSHIP (Potter et al. 2017), UNIT (Chuang et al. 2019), Uchuu (Ishiyama et al. 2021), and AbacusSummit (Maksimova et al. 2021), remain well below the particle numbers we require. This problem is compounded by the need to reduce the uncertainty in these simulations to levels well below the statistical uncertainty of observations in order to maximize the information extracted from survey data: a simulation with size merely *equal* to the survey volume is insufficient.

Angulo & Pontzen (2016) proposed to suppress the variance in these simulations by removing amplitude fluctuations in the various k-modes in the initial conditions of a simulation. Such fixed-amplitude initial conditions have been employed in hydrodynamical simulations (Villaescusa-Navarro et al. 2018) and have been tested on Lyman-α forest statistics (Anderson et al. 2019). Chuang et al. (2019) also tested the method on the clustering measurements of galaxy redshift surveys. All found that the variance-suppression method provides more precise predictions without introducing bias, enabling simulations with large effective volume at much lower computational cost. However, the reduced variance in these simulations still must be quantified for comparison with the variance in survey data (whose estimation is an important problem in its own right), and we must verify that the methods typically employed to estimate covariance matrices for the clustering statistics of galaxy catalogues can provide reliable estimates when those catalogues are derived from variance-suppressed simulations.

[4] http://www.4most.eu/
[5] http://hetdex.org
[6] http://j-pas.org
[7] https://pfs.ipmu.jp
[8] http://www.lsst.org/lsst/
[9] http://www.euclid-ec.org
[10] https://roman.gsfc.nasa.gov

Traditionally, covariance matrix estimation for both observation and simulation data has been done with mock galaxy catalogues. Various methods exist for generating such mock catalogues, but we broadly categorize them into two classes based on how they construct their halo catalogues. The first class defines haloes by applying a halo finder on simulated dark matter particles, and includes methods such as PTHALOS (Manera et al. 2012, 2014), PINOCCHIO (PINpointing Orbit-Crossing Collapsed Hierarchical Objects, Monaco et al. 2002, 2013), PEAK PATCH (Bond & Myers 1996), FASTPM (Feng et al. 2016), COLA (COmoving Lagrangian Acceleration simulation; Tassev et al. 2013; Izard et al. 2016), and GLAM (Klypin & Prada 2018). These methods tend to be memory-intensive, as they require a large number of particles in order to resolve haloes; indeed, FASTPM, COLA, and GLAM are essentially efficient N-body simulations. As a consequence, however, mock catalogues produced with such codes using fixed-amplitude initial conditions naturally yield good covariance matrix estimates for fixed-amplitude simulations.

The second class of methods populates haloes based on bias models using coarse-resolution dark matter density fields. In these models, halo creation is stochastic, but is calibrated to reproduce the clustering statistics of a reference catalogue (obtained, for example, from survey data or from a higher resolution simulation). This second class includes such methods as LOG-NORMAL (Coles & Jones 1991), PATCHY (PerturbAtion Theory Catalog generator of Halo and galaxY distributions; Kitaura et al. 2013, 2015), HALOGEN (Avila et al. 2015), QPM (quick particle mesh; White et al. 2013), BAM (Bias Assignment Method to generate mock catalogues; Balaguera-Antolínez et al. 2019), and EZMOCK (Effective Zel'dovich approximation mock catalogue; Chuang et al. 2015a).

The computational cost of this second class of methods is much lower, but the validity of their use in covariance matrix estimation is less obvious. Some past work exists validating these methods for covariance matrix estimation and comparing them with each other and with those of the first class (see e.g. Chuang et al. 2015b; Blot et al. 2019; Colavincenzo et al. 2019; Lippich et al. 2019), but more study is desired to fully validate their use in upcoming surveys, especially for covariance matrix estimation of observational data.

Still, these methods have gained traction for their key advantage of efficiency. EZMOCK in particular delivers reasonable accuracy at very low computational cost, and indeed, was employed extensively in the cosmological analysis of the final eBOSS galaxy sample (see e.g. Avila et al. 2020; Bautista et al. 2020; Gil-Marin et al. 2020; Hou et al. 2020; Kong et al. 2020; Mohammad et al. 2020; Neveux et al. 2020; Raichoor et al. 2020; Ross et al. 2020; Tamone et al. 2020; de Mattia et al. 2020; eBOSS Collaboration, 2021; Zhao et al. 2021a,b. It remains unclear, however, whether EZMOCK, or any method in this second class, can be extended to estimate covariance matrices of variance-suppressed simulations.

In this study, we show that EZMOCK can provide reasonable covariance matrix estimates for such simulations. The rest of this paper is organized as follows. We outline the simulations used in this study in Section 2, then present our findings in Section 3. In particular, we show that fixed-amplitude EZMOCK can reproduce the covariance matrices of a reference fixed-amplitude simulation in Section 3.1, then proceed in Section 3.2 to examine in turn the effects of three-point clustering, small-scale clustering, and galaxy bias on the level of variance suppression obtained by fixing amplitudes, showing in the process how EZMOCK can provide intuitive explanations for our observations. We summarize and conclude in Section 4.

2 SIMULATIONS

2.1 FASTPM catalogues

In order to study how EZMOCK reproduces the covariance matrix of a reference simulation, we require a large number of reference simulations whose covariance matrices we will estimate. To this end, we use a set of simulations produced with the FASTPM code (Feng et al. 2016). The FASTPM code relies on accelerated particle–mesh solvers, which have recently been shown to produce accurate halo populations (compared to full N-body calculations) when enhanced with various techniques (cf. the COLA code; Tassev et al. 2013). FASTPM employs a pencil domain-decomposition Poisson solver and a Fourier-space four-point differential kernel to compute the force. Given that FASTPM simulations are essentially N-body simulations, we use them as our reference in this study.

Our simulations, 200 in all, are a subset of the FASTPM simulations created by the UNIT project to test the behaviour of the variance-suppression method (Chuang et al. 2019). All of these simulations are publicly available.[11] Of the 200 simulations we use here, half were produced with fixed-amplitude initial conditions and half without. All are at $z = 1$. Each simulation was generated with 1024^3 particles in a box $1\,\mathrm{Gpc}\,h^{-1}$ on each side and simulated for 100 time-steps. A Friends-of-Friends halo finder was used to identify haloes; the minimum halo mass is $1.68 \times 10^{12}\,M_\odot\,h^{-1}$.

Throughout this work we use the same cosmological parameters as these FASTPM simulations: $\Omega_\mathrm{m} = 0.3089$, $h \equiv H_0/(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}) = 0.6774$, $n_s = 0.9667$, and $\sigma_8 = 0.8147$ (based on Planck Collaboration, 2016, table 4).

2.2 EZMOCK

EZMOCK (Chuang et al. 2015a) is a method for generating mock catalogues based on application of a bias model to a coarse-resolution dark matter field. The dark matter field is constructed from the Zel'dovich approximation (ZA) density field. EZMOCK absorbs non-linear effects and halo bias (i.e. linear, non-linear, deterministic, and stochastic bias) into an effective model with only a few free parameters, which can be efficiently calibrated with N-body simulations. We use the slightly modified version described in Baumgarten & Chuang (2018), consisting of the following three steps:

(i) **Generate the dark matter field.** In the Lagrangian formulation of cosmological fluid dynamics, we describe the motion of a particle originally at $\boldsymbol{q}$ to a position $\boldsymbol{x}$ at cosmic time t by a Lagrangian displacement field $\boldsymbol{\Psi}$:

$$\boldsymbol{x}(\boldsymbol{q}, t) = \boldsymbol{q} + \boldsymbol{\Psi}(\boldsymbol{q}, t). \tag{1}$$

The first-order Lagrangian perturbation theory solution to the equations of motion is given by the ZA

$$\boldsymbol{\Psi}(\boldsymbol{q}) = \int \frac{d^3\boldsymbol{k}}{(2\pi)^3} \mathrm{e}^{i\boldsymbol{k}\cdot\boldsymbol{q}} \frac{i\boldsymbol{k}}{k^2} \hat{\delta}(\boldsymbol{k}) \tag{2}$$

where $\hat{\delta}(\boldsymbol{k})$ is the fractional density perturbation in Fourier space (for a review, see e.g. Bernardeau et al. 2002).

We construct this displacement field in the ZA and extrapolate to the redshift of the reference halo or galaxy sample (in our case, $z = 1$). To generate our dark matter density field, we initialize a uniform square lattice of dark matter particles and evolve the particles according to $\boldsymbol{\Psi}$. We can apply fixed-amplitude initial conditions in this step by selecting the $\hat{\delta}(\boldsymbol{k})$ to have random phases but predetermined fixed amplitudes rather than drawing them from Gaussian distributions.

(ii) **Determine final object densities using the dark matter density and the halo probability distribution function (PDF).** We model the PDF of the final mock catalogue, the number of grid cells containing n objects as a function of n, as

$$P(n) \propto A^n, \tag{3}$$

normalized to give the desired final object number density (a free parameter). Greater dark matter density should yield greater object density. Thus, we might naively sort the grid cells $\boldsymbol{r}$ by increasing dark matter density $\rho_0(\boldsymbol{r})$ and assign each cell the number of objects it should contain: the first $P(0)$ will contain no objects, the next $P(1)$ will contain one, etc.

In practice, we use modified "densities" $\rho_s(\boldsymbol{r})$ for this rank-ordering procedure to allow more tuning of our clustering. Beginning with dark matter densities $\rho_0(\boldsymbol{r})$ in each cell obtained using the cloud-in-cells (CIC) particle assignment scheme (see e.g. Hockney & Eastwood 1981), cells with densities below some density cut ρ_c are assigned $\rho_s = 0$. Adjusting ρ_c allows us to modify the bispectrum of the catalogues.

On the other hand, if $\rho_0 \geq \rho_c$, we compute

$$\rho_s(\boldsymbol{r}) = (1 - \mathrm{e}^{-\rho_0(\boldsymbol{r})/\rho_a}) \begin{cases} 1 + G(\boldsymbol{r}) & G(\boldsymbol{r}) \geq 0 \\ \mathrm{e}^{G(\boldsymbol{r})} & G(\boldsymbol{r}) < 0 \end{cases} \tag{4}$$

where, for each cell, $G(\boldsymbol{r})$ is independently drawn from a Gaussian distribution of some specified width λ [typically left fixed, having observed in the course of previous work (Baumgarten & Chuang 2018) that its effect on clustering statistics was highly degenerate with that of the other parameters]. This procedure effectively introduces some scatter into the PDF mapping procedure described above. The non-linear function $1 - \mathrm{e}^{-\rho_0/\rho_a}$ introduces saturation behaviour at a characteristic density ρ_a (another free parameter), converging to 1 when $\rho_0 \gg \rho_a$.

(iii) **Assign objects.** The mock catalogue is populated by selecting at random a subset of the original dark matter particles such that each cell contains the number of objects assigned to it in the previous step.

We generate three sets of EZMOCK catalogues, each consisting of 2×1000 catalogues of size $(1\,\mathrm{Gpc}\,h^{-1})^3$ and number density $\overline{n} = 2.343 \times 10^6\,/(\mathrm{Gpc}/h)^3$ (chosen to match the FASTPM catalogues) generated on a 256^3 grid. Within each set, half the catalogues are produced using normal initial conditions and half using fixed-amplitude initial conditions.

The first set of EZMOCKS is a fiducial set with parameters chosen to reproduce the clustering statistics of the FASTPM catalogues ($A = 0.37$ and $\rho_a, \rho_c = 2.25, 0$ in units of objects per grid cell). The second set has matching two-point clustering but differs in its three-point clustering (i.e. bispectrum) ($A = 0.37$ and $\rho_a, \rho_c = 0.75, 1$). The third set maintains the linear-scale amplitude of the fiducial set but differs in small-scale clustering ($A = 0.3$, $\rho_a, \rho_c = 3.1, 0$). We fixed $\lambda = 10$.

3 RESULTS

3.1 Covariance matrix estimation

It is well known that the fractional uncertainty in a band-averaged power spectrum value $P(k)$ for a Gaussian random field under

[11] http://www.unitsims.org

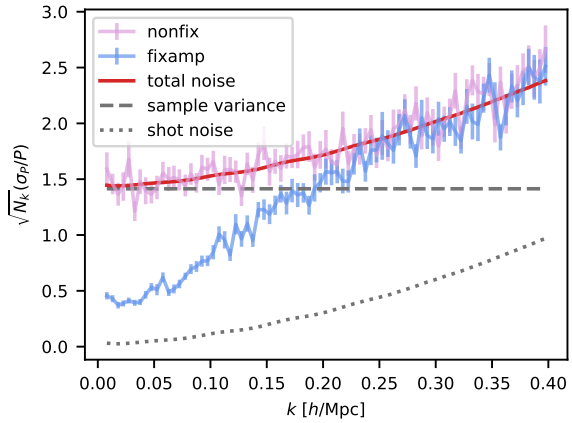

Figure 1. Comparison of fractional uncertainty in the power spectrum $P(k)$ for the two sets of FASTPM catalogues, one produced with (blue; "fixamp") and one produced without (pink; "nonfix") fixed-amplitude initial conditions, against the theoretical fractional uncertainty (red). We scale by $\sqrt{N_k}$, where $N_k \sim k^2$ is the number of modes in each k-bin. This theoretical uncertainty is the sum of a Gaussian component (grey dashed line) and a Poisson shot-noise component (grey dotted line). The theoretical uncertainty agrees with that of the non-fixed-amplitude FASTPM simulations very well, but that of the fixed-amplitude FASTPM simulations can neither be described by the theoretical uncertainty nor its Gaussian or Poissonian components at small k. The computation of error bars is described in Appendix A.

Poisson sampling can be decomposed into a sum of a Gaussian sample variance and a shot-noise contribution as

$$\frac{\sigma_{P(k)}}{P(k)} = \sqrt{\frac{2}{N_k}}\left(1 + \frac{1}{\bar{n}P(k)}\right), \tag{5}$$

where $N_k \sim k^2$ is the number of modes in the k-bin and $\bar{n}$ is the mean number density (see Feldman et al. 1994). Figure 1 confirms that this model accurately accounts for the power spectrum uncertainty in the non-fixed-amplitude FASTPM catalogues.

The fixed-amplitude FASTPM catalogues show suppressed variance at linear scales (small k). We might expect the residual uncertainty here to be dominated by a shot-noise contribution, as the initial dark matter power spectrum has no Gaussian sample variance (by construction), but Figure 1 shows that this is not the case. The excess variance here is also too great to be easily explained as the result of weak non-linear effects affecting the evolution of the dark matter field; Angulo & Pontzen (2016) found much lower dark matter power spectrum variance in their fixed-amplitude simulations at these k (see the residuals for individual fixed-amplitude simulations at $k < 0.1\,h\,\mathrm{Mpc}^{-1}$ in their fig. 2). Since halo presence is intimately connected to small-scale dark matter clustering, the residual small-k variance in the halo power spectrum in the fixed-amplitude FASTPM simulations may very well arise from large-k dark matter power spectrum variance. It is evidently non-trivial to build a theoretical model for the covariance matrix of two-point statistics of fixed-amplitude simulations, even at linear scales.

Can a method like EZMOCK, which produces mock catalogues based on a bias model, reproduce the covariance matrix of these fixed-amplitude catalogues? Yes. In Figure 2, we compare the power spectrum and two-point correlation function for a fiducial set of 2×100 EZMOCKS and the FASTPM catalogues. The parameters are the same between the fixed-amplitude ("fixamp") and non-fixed-amplitude ("nonfix") EZMOCKS. Thus, the second row shows that applying the fixed-amplitude condition does not bias our clustering statistics, consistent with previous studies (see e.g. Chuang et al. 2019). In the bottom panels, the variance suppression observed in the power spectrum of the FASTPM boxes is reproduced by the EZMOCKS. The variance suppression observed in the FASTPM boxes' correlation function measurements is also reproduced reasonably at large scales, though EZMOCK underestimates the suppression at smaller scales. This underestimation may be a limitation of the bias model underlying EZMOCK being "too stochastic" at small scales, as further tuning of fit parameters does not tend to improve this underestimation.

We also compute normalized covariance matrices (i.e. correlation matrices) of the power spectrum and correlation function for the full set of fiducial EZMOCKS (both fixed-amplitude and non-fixed-amplitude), as well as for the reference FASTPM catalogues. These correlation matrices are shown in Figure 3, where we observe that for both the power spectrum and the correlation function, the EZMOCK correlation matrix agrees with that of our FASTPM boxes.

The stochastic bias we observe in the fixed-amplitude EZMOCK catalogues is introduced through the scattering procedure used in the PDF mapping, as described in step (ii) of Section 2.2. The parameters involved are calibrated with the clustering measurements of the reference catalogue (in this case, FASTPM). We do *not* tune additional parameters to calibrate the covariance matrix separately. That the covariance matrices obtained from the EZMOCKS agree with those of the FASTPM simulations shows that the EZMOCK bias model reasonably reproduces the stochastic bias of the FASTPM haloes.

In Figure 3, we further note that for both the EZMOCKS and the FASTPM catalogues, the power spectrum correlation matrices are similar between the fixed-amplitude and non-fixed-amplitude simulations, as underscored by the cuts plotting the first off-diagonal correlations: fixed-amplitude initial conditions do not bias the mode coupling in Fourier space.

In contrast, the correlation matrices for the correlation functions are dramatically different between fixed-amplitude and non-fixed-amplitude catalogues. This difference is not surprising, however. Indeed, while one can compute the covariance matrix of the correlation function from that of the power spectrum by expressing the correlation function as a Fourier transform of the power spectrum, the correlation matrix of the correlation function cannot be determined from only the *correlation* matrix of the power spectrum: the magnitude of the variances in the power spectrum, suppressed in the fixed-amplitude mocks, must be considered when deriving the correlation matrix of the correlation function.

3.2 Influences on variance suppression

We proceed to investigate the impact of three-point clustering, small-scale clustering, and overall galaxy bias on the degree of variance suppression observed in our mock catalogues upon application of fixed-amplitude initial conditions, and provide intuitive explanations for our observations based on the EZMOCK bias model.

3.2.1 *Effects of higher order clustering*

To test the impact of higher order clustering on the variance suppression, we compare the fiducial set of EZMOCKS with the second set of 2000 EZMOCKS described in Section 2.2. The two sets have similar two-point statistics (within 2% for $k < 0.35\,\mathrm{Mpc}\,h^{-1}$), but their differing density cuts ρ_c produce different bispectra, as shown in the top panel of Figure 4. From Figure 5, the variance suppression from fixing amplitudes appears slightly worse with greater ρ_c (lower bispectrum).

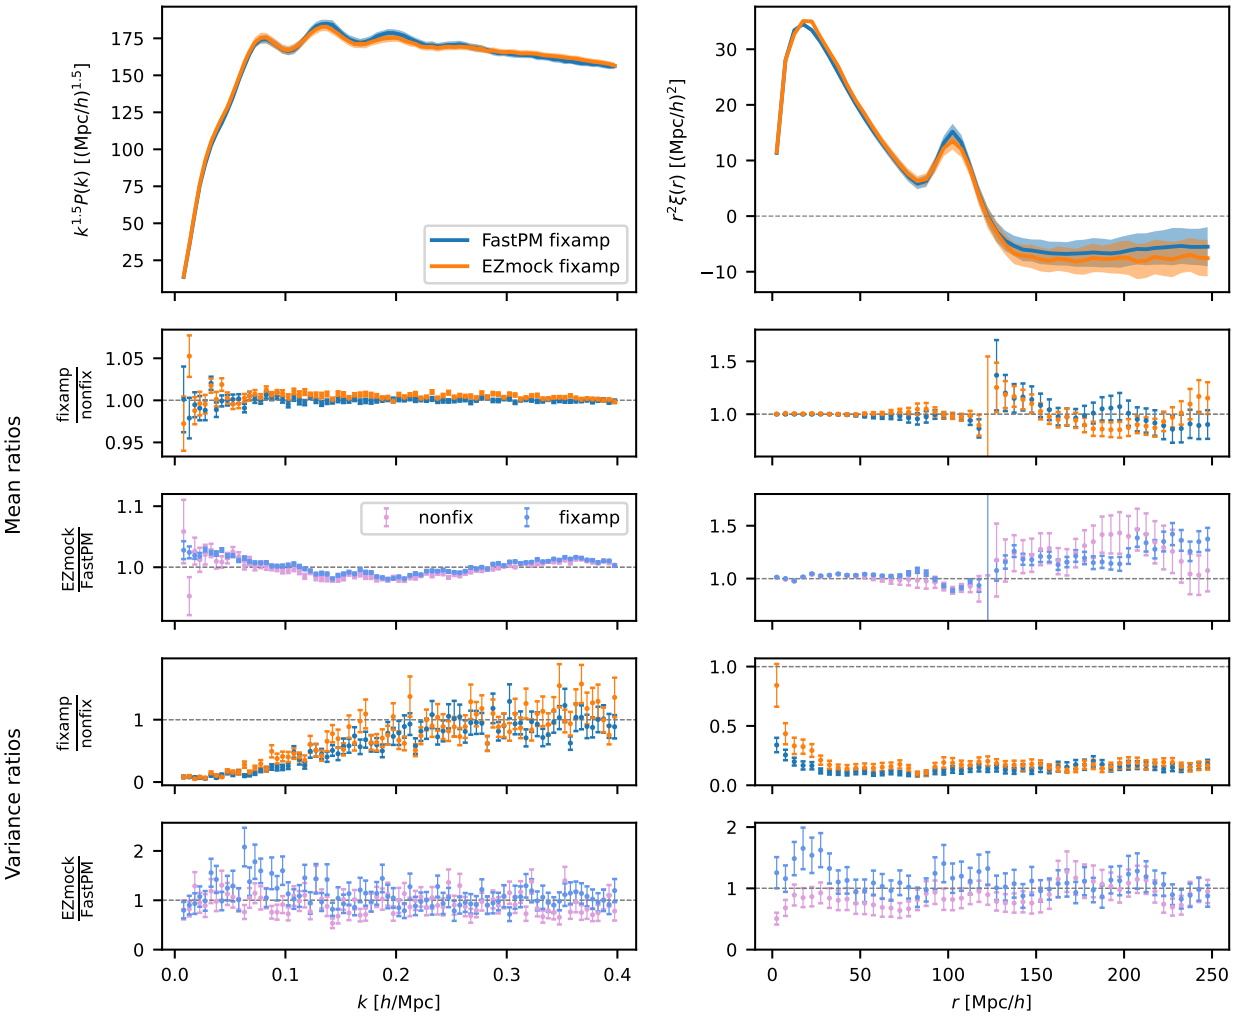

Figure 2. Comparison of FASTPM catalogues and 2×100 fitted EZMOCKS. Top panels show the mean power spectrum $P(k)$ and correlation function $\xi(r)$ of the fixamp FASTPM and EZMOCK catalogues (shaded bands show $\pm 1\sigma$). The next row, which shows the ratio of the $P(k)$ and $\xi(r)$ means for both FASTPM and EZMOCK for the fixamp catalogues and the corresponding means for the nonfix catalogues, indicates that the fixed-amplitude condition does not bias the mean values. The third row compares the means in a different way, plotting the ratio between the EZMOCKS and the FASTPM catalogues for each amplitude condition (i.e. nonfix EZMOCK vs. nonfix FASTPM and fixamp EZMOCK vs. fixamp FASTPM). We see that EZMOCK reproduces the clustering measurements of the FASTPM catalogues. The final two rows are analogous to the second and third, but show ratios of variances. It is evident that the EZMOCKS reproduce the variance suppression seen in the FASTPM simulations, especially at large scales. Standard errors are computed per Appendix A.

One might expect this behaviour from the perspective of the EZMOCK bias model. EZMOCK introduces scatter during the PDF mapping procedure (see step (ii) of Section 2.2) to reduce the clustering amplitude (galaxy linear bias): the larger the scatter, the lower the bias. However, when we apply a higher density cut, we increase the linear bias of the resulting EZMOCK, so to achieve the same two-point clustering as before, we must then reduce the linear bias by introducing greater scatter during PDF mapping, resulting in larger variance and weaker variance suppression. Indeed, as previously mentioned, the scatter procedure is the primary source of stochastic bias in EZMOCK; it introduces variance even when we apply fixed-amplitude initial conditions.

Incidentally, the middle panel of Figure 4, indicates that the fixed-amplitude condition yields no significant improvement in bispectrum uncertainty, as observed by Angulo & Pontzen (2016). The lower panel of Figure 4 indicates that the variances in the bispectrum are not sensitive to the mean value of the bispectrum.

3.2.2 *Effects of small-scale clustering*

We next explore the effect of adjusting small-scale (large k) clustering. We produced the third set of EZMOCKS described in Section 2.2 with a stronger PDF slope ($A = 0.3$ versus the original $A = 0.37$) and thus weaker small-scale clustering compared to the fiducial set. In order to isolate the effect of small-scale clustering, we adjusted the scatter parameter ρ_a to maintain the linear amplitude of the fiducial EZMOCKS. In Figure 6, we find that these new EZMOCKS, with weaker small-scale clustering, show stronger variance suppression.

The EZMOCK bias model again gives a natural explanation. The PDF mapping procedure inherently boosts the amplitude of density fluctuations in all k-modes (not necessarily by the same amount).

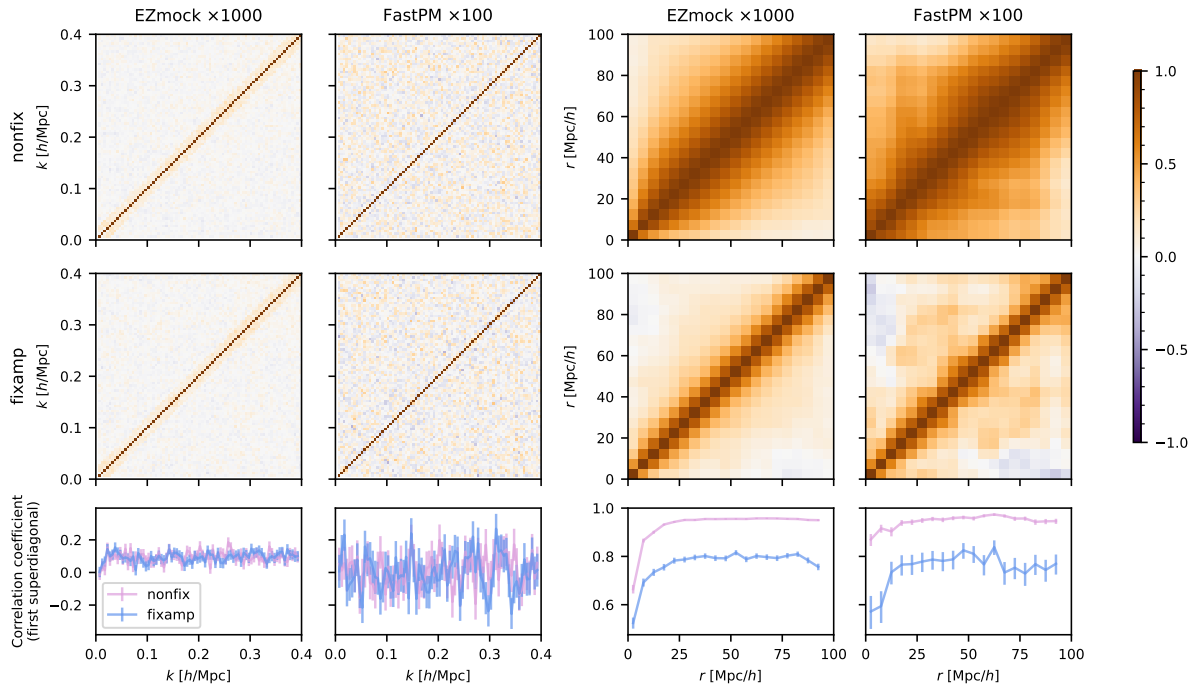

Figure 3. Comparison of correlation matrices for the two-point clustering statistics of four sets of simulations: the "nonfix" and the "fixamp" boxes for both the fiducial EZMOCKS and the reference FASTPM simulations. The EZMOCK (FASTPM) sets consist of 1000 (resp. 100) boxes each. The top two rows show correlation matrices for the nonfix and fixamp boxes. The correlation structure in the power spectrum is similar between the nonfix and fixamp boxes for both EZMOCK and FASTPM, but the fixamp correlation function values are less correlated than their nonfix counterparts, as confirmed by the lower panels, which show the first off-diagonal terms of the correlation matrices above [i.e. correlations between $P(k)$ and $P(k+\Delta k)$ on the left and likewise with the correlation function on the right]. We also find that the correlation structure between the corresponding sets of EZMOCK and FASTPM boxes is quite similar, up to the statistical uncertainty of the FASTPM boxes. Standard errors are computed per Appendix A.

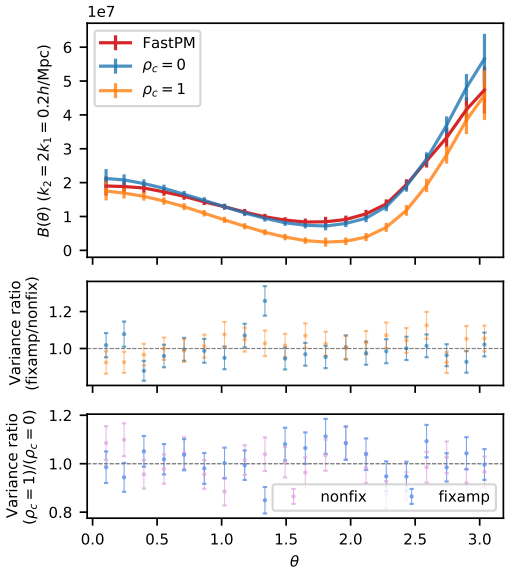

Figure 4. Top: Mean ($\pm 1\sigma$) of the bispectrum $B(\theta)$ ($k_2 = 2k_1 = 0.2\,h\,\mathrm{Mpc}^{-1}$) for fixamp EZMOCK catalogues with the same two-point clustering but differing density cuts ρ_c, as described in the text. Middle: Ratios of $B(\theta)$ variances between fixamp catalogues and their nonfix counterparts at both values of ρ_c, demonstrating that fixing amplitudes yields no improvement in bispectrum uncertainty. Bottom: Bispectrum variance ratios between $\rho_c = 1$ and $\rho_c = 0$ EZMOCKS for both the nonfix (pink) and fixamp (blue) initial conditions show that the variances in $B(\theta)$ are not sensitive to the mean values of $B(\theta)$. Standard errors computed per Appendix A.

The scatter procedure then reduces bias at the linear scale (small k), but does little to the small-scale clustering, which remains governed by the PDF we choose (i.e. A). So for our new set of EZMOCKS, with its lower small-scale clustering, the PDF mapping procedure does not increase the linear bias as much as it does in the fiducial set, thus requiring less scatter to reach the targeted linear bias. The result is the stronger variance suppression we observe.

3.2.3 *Effects of galaxy bias*

Finally, we investigate how galaxy bias affects variance suppression in fixed-amplitude catalogues. For each of the 200 FASTPM catalogues, we produce three subcatalogues with only half the number density of the full catalogues. The "heavy" and "light" subcatalogues are simply the top and bottom halves of the catalogues split by mass, while the "random" subcatalogue is a random 50% downsampling. We show the power spectra of these subcatalogues in Figure 7, where we find that the variance improvement in the "heavy" subcatalogues is noticeably stronger than in the "light" and "random" subcatalogues. Since the subcatalogues have equal number density, the shot-noise contributions to the variance are equal, and we conclude that larger bias leads to stronger variance suppression.

One might expect the opposite trend: stronger variance suppression when there is *less* bias. Indeed, we know that the underlying dark matter field (which has unity bias, lower than any of these galaxy samples) should show the greatest variance improvement when fixed-amplitude initial conditions are used: there should be no dark matter power spectrum variance in fixed-amplitude simulations, by definition.

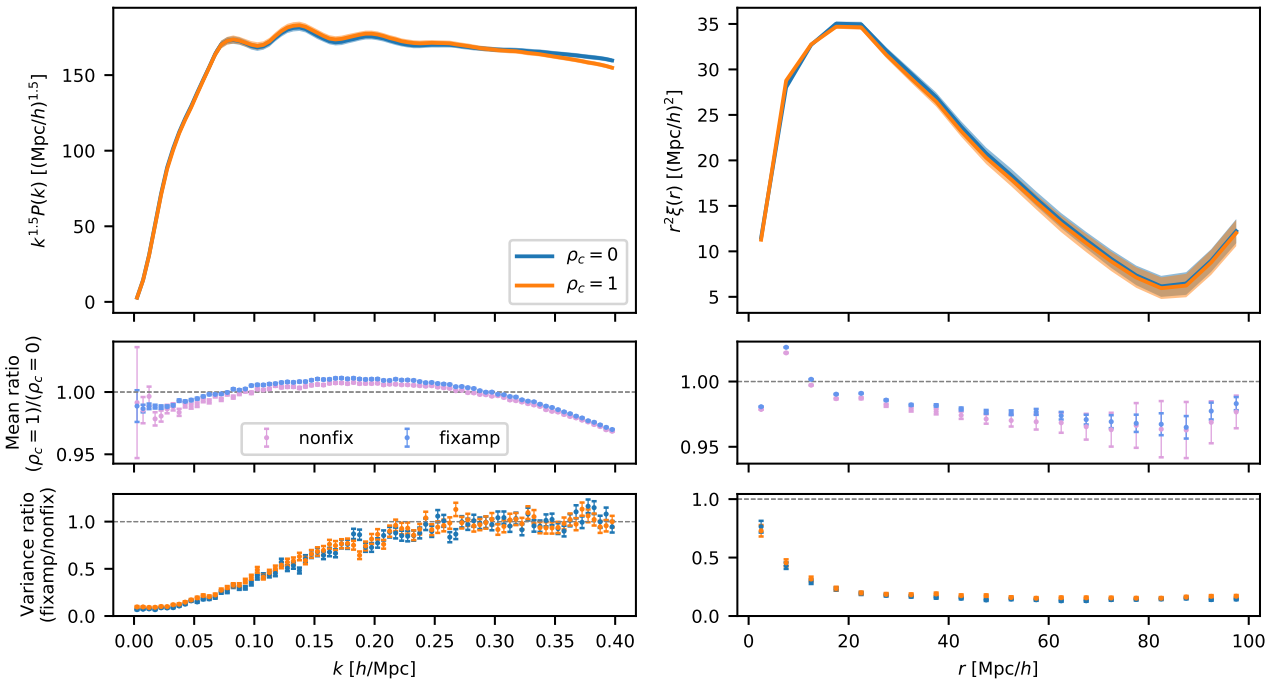

Figure 5. Comparison of two-point clustering statistics for two sets of 2000 EZMOCKS, both generated to match the FASTPM catalogues' two-point statistics, but with differing density cuts ρ_c giving different three-point clustering, as described in the text. Top: Mean power spectrum and correlation function for the fixed-amplitude mocks in each set. Middle: Ratio of the mean values of the statistics (for each amplitude condition) between the $\rho_c = 1$ and $\rho_c = 0$ mocks showing agreement to within a few percent in both Fourier and configuration space. Bottom: Variance suppression obtained by the fixed-amplitude condition for both sets. We find slightly stronger suppression in the set with $\rho_c = 0$. Standard errors are computed per Appendix A.

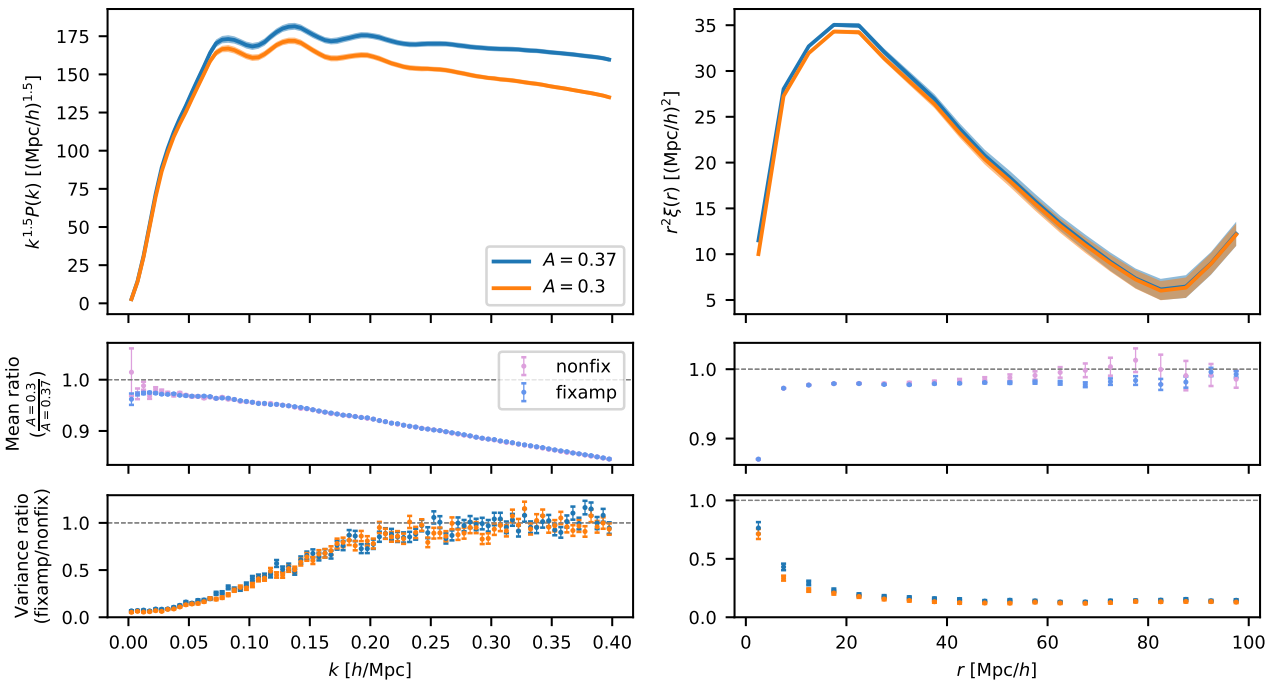

Figure 6. Two-point statistics for two sets of EZMOCKS generated with different values of A in Equation 3 to yield differing small-scale clustering but tuned to have matching large-scale clustering. Top: Mean statistics for the fixed-amplitude catalogues in each set (1000 each). Middle: Ratios between the mean values (for both initial conditions), showing comparable large-scale clustering but suppressed small-scale clustering for the $A = 0.3$ mocks. Bottom: Ratios of variances in the statistics between the fixed-amplitude catalogues and their non-fixed-amplitude counterparts, showing better variance suppression when there is less small-scale clustering. Standard errors are computed per Appendix A.

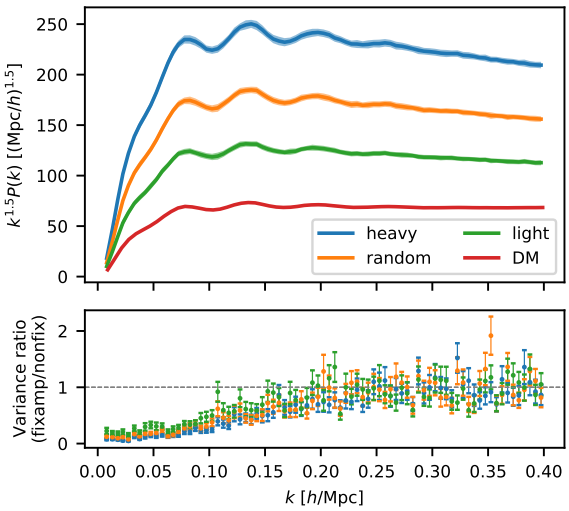

Figure 7. Top: Power spectra from subcatalogues of the FASTPM catalogues. The subcatalogues are created in three ways: by selecting the top and bottom halves by mass (yielding "heavy" and "light" catalogues) and by randomly downsampling by 50% ("random"). We plot mean power spectra ($\pm 1\sigma$) for the subcatalogues derived from the fixed-amplitude FASTPM catalogues, as well as the dark matter power spectrum ("DM") used in creating the FASTPM catalogues. Bottom: Ratios between the fixed-amplitude and non-fixed-amplitude variances, revealing stronger variance suppression in samples with larger bias. Standard errors are computed per Appendix A.

Interestingly, the EZMOCK bias model allows us to understand this apparent contradiction. Indeed, EZMOCK introduces scatter to *reduce* clustering amplitudes from some higher bias that the PDF mapping procedure would produce in the absence of scatter. To achieve a lower bias in the resulting EZMOCK catalogue, more scatter must be applied (and thus more stochastic bias introduced). So increased variances in clustering measurements are associated with *reduced* clustering amplitude.

4 CONCLUSIONS

In this study, we demonstrated that EZMOCK with fixed-amplitude initial conditions can be used to estimate covariance matrices for fixed-amplitude N-body simulations. We further investigated the behaviour of the variance suppression introduced by fixed-amplitude initial conditions. Our main findings can be summarized as follows:

(i) No simple analytical form exists for covariance matrices of fixed-amplitude catalogues.

(ii) Fixed-amplitude initial conditions do not bias clustering measurements in EZMOCK.

(iii) After calibrating EZMOCK with clustering measurements from a reference fixed-amplitude catalogue, the resulting covariance matrix is a good estimate for that of the reference without further calibration.

(iv) In fixed-amplitude catalogues, variance suppression is stronger with stronger three-point clustering, weaker when there is stronger small-scale clustering, and stronger when there is greater large-scale bias.

(v) The relative strength of the variance suppression in these cases can be understood with the EZMOCK bias model: the variance suppression due to fixed-amplitude initial conditions is reduced when EZMOCK requires greater scatter to reproduce the clustering statistics of a reference simulation.

As survey volumes continue to grow, variance-suppressed cosmological N-body simulations will become increasingly important in order to deliver large effective simulation volumes at reduced computational cost. Our work validates EZMOCK as an effective and efficient method for estimating covariance matrices for these variance-suppressed simulations, paving the way for the use of the variance-suppression technique in future simulations, and by extension, for much more computationally cost-effective cosmological analysis.

ACKNOWLEDGEMENTS

This work was supported in part by U.S. Department of Energy contracts to SLAC (DE-AC02-76SF00515). This research has made use of NASA's Astrophysics Data System and the arXiv preprint server.

Some of the computing for this project was performed on the Sherlock cluster at Stanford. We would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

Anderson L., Pontzen A., Font-Ribera A., Villaescusa-Navarro F., Rogers K. K., Genel S., 2019, ApJ, 871, 144
Angulo R. E., Pontzen A., 2016, MNRAS, 462, L1
Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, MNRAS, 426, 2046
Avila S., Murray S. G., Knebe A., Power C., Robotham A. S. G., Garcia-Bellido J., 2015, MNRAS, 450, 1856
Avila S., et al., 2020, MNRAS, 499, 5486
Balaguera-Antolínez A., Kitaura F.-S., Pellejero-Ibáñez M., Zhao C., Abel T., 2019, MNRAS, 483, L58
Baumgarten F., Chuang C.-H., 2018, MNRAS, 480, 2535
Bautista J. E., et al., 2020, MNRAS, 500, 736
Benitez N., et al., 2014, preprint (arXiv:1403.5237)
Bernardeau F., Colombi S., Gaztañaga E., Scoccimarro R., 2002, Phys. Rep., 367, 1
Blanton M. R., et al., 2017, AJ, 154, 28
Blot L., et al., 2019, MNRAS, 485, 2806
Bond J. R., Myers S. T., 1996, ApJS, 103, 1
Chuang C.-H., Kitaura F.-S., Prada F., Zhao C., Yepes G., 2015a, MNRAS, 446, 2621
Chuang C.-H., et al., 2015b, MNRAS, 452, 686
Chuang C.-H., et al., 2019, MNRAS, 487, 48
Colavincenzo M., et al., 2019, MNRAS, 482, 4883
Coles P., Jones B., 1991, MNRAS, 248, 1
Dark Energy Survey Collaboration 2005, arXiv e-prints, pp astro–ph/0510346
de Jong R. S., et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, Proc. SPIE Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV. p. 84460T
de Mattia A., et al., 2020, MNRAS, 501, 5616
eBOSS Collaboration, 2021, Phys. Rev. D, 103, 083533

Feldman H. A., Kaiser N., Peacock J. A., 1994, ApJ, 426, 23
Feng Y., Chu M.-Y., Seljak U., McDonald P., 2016, MNRAS, 463, 2273
Fosalba P., Crocce M., Gaztañaga E., Castander F. J., 2015, MNRAS, 448, 2987
Gil-Marin H., et al., 2020, MNRAS, 498, 2492
Gonzalez-Perez V., et al., 2017, MNRAS, 474, 4024
Habib S., et al., 2016, New Astron., 42, 49
Hill G. J., et al., 2008, in Kodama T., Yamada T., Aoki K., eds, ASP Conf. Ser. Vol. 399, Panoramic Views of Galaxy Formation and Evolution. p. 115
Hockney R. W., Eastwood J. W., 1981, Computer Simulation Using Particles. McGraw-Hill, New York
Hou J., et al., 2020, MNRAS, 500, 1201
Ishiyama T., et al., 2021, MNRAS, 506, 4210
Izard A., Crocce M., Fosalba P., 2016, MNRAS, 459, 2327
Kenney J. F., Keeping E. S., 1951, Mathematics of Statistics, Part Two, 2 edn. Van Nostrand, New York
Kitaura F.-S., Yepes G., Prada F., 2013, MNRAS, 439, L21
Kitaura F.-S., Gil-Marín H., Scóccola C. G., Chuang C.-H., Müller V., Yepes G., Prada F., 2015, MNRAS, 450, 1836
Klypin A., Prada F., 2018, MNRAS, 478, 4602
Klypin A., Yepes G., Gottlober S., Prada F., Hess S., 2016, MNRAS, 457, 4340
Kong H., et al., 2020, MNRAS, 499, 3943
LSST Science Collaboration, 2009, preprint (arXiv:0912.0201)
Laureijs R., et al., 2011, preprint (arXiv:1110.3193)
Levi M., et al., 2013, preprint (arXiv:1308.0847)
Lippich M., et al., 2019, MNRAS, 482, 1786
Maksimova N. A., Garrison L. H., Eisenstein D. J., Hadzhiyska B., Bose S., Satterthwaite T. P., 2021, MNRAS, 508, 4017
Manera M., et al., 2012, MNRAS, 428, 1036
Manera M., et al., 2014, MNRAS, 447, 437
Mohammad F. G., et al., 2020, MNRAS, 498, 128
Monaco P., Theuns T., Taffoni G., 2002, MNRAS, 331, 587
Monaco P., Sefusatti E., Borgani S., Crocce M., Fosalba P., Sheth R. K., Theuns T., 2013, MNRAS, 433, 2389
Neveux R., et al., 2020, MNRAS, 499, 210
Planck Collaboration, 2016, A&A, 594, A13
Potter D., Stadel J., Teyssier R., 2017, Comput. Astrophys. Cosmol., 4, 2
Raichoor A., et al., 2020, MNRAS, 500, 3254
Ross A. J., et al., 2020, MNRAS, 498, 2354
Schlegel D., et al., 2011, preprint (arXiv:1106.1706)
Skillman S. W., Warren M. S., Turk M. J., Wechsler R. H., Holz D. E., Sutter P. M., 2014, preprint (arXiv:1407.2600)
Spergel D., et al., 2013, preprint (arXiv:1305.5422)
Takada M., et al., 2014, PASJ, 66, R1
Tamone A., et al., 2020, MNRAS, 499, 5527
Tassev S., Zaldarriaga M., Eisenstein D., 2013, J. Cosmology Astropart. Phys., 1306, 036
Villaescusa-Navarro F., et al., 2018, ApJ, 867, 137
White M., Tinker J. L., McBride C. K., 2013, MNRAS, 437, 2594
Zhao C., et al., 2021a, MNRAS, 503, 1149
Zhao G.-B., et al., 2021b, MNRAS, 504, 33

APPENDIX A: UNCERTAINTY ESTIMATION

In this appendix, we briefly describe our computation of standard errors for variances and correlation coefficients.

Indeed, as variance suppression is central to our investigation, we must take care to compute standard errors on variances correctly. For data $x_1, \ldots, x_N$, we estimate the variance as

$$k_2 = \frac{1}{N-1}\sum_{i=1}^{N}(x_i - \bar{x})^2, \tag{A1}$$

where $\bar{x}$ is the sample mean $\frac{1}{N}\sum_i x_i$. This is the well-known minimum-variance symmetric unbiased estimator of the variance. Per Kenney & Keeping (1951, p. 189), an unbiased estimate of the variance of k_2 is given by

$$\frac{2k_2^2 + \frac{N-1}{N}k_4}{N+1}, \tag{A2}$$

where k_4 is the fourth k-statistic

$$k_4 = \frac{N^2[(N+1)m_4 - 3(N-1)m_2^2]}{(N-1)(N-2)(N-3)}, \tag{A3}$$

expressed in terms of the second and fourth sample moments about the mean $m_{2,4}$ ($m_n = \frac{1}{N}\sum_i (x_i - \bar{x})^n$). Taking the square root of the variance estimate A2 yields a standard error for our variances.

We estimate uncertainties on correlation coefficients r via the Fisher transformation $z = \tanh^{-1} r$, which produces a new random variable z with an approximately normal distribution with standard deviation $\delta z \approx 1/\sqrt{N-3}$ (Kenney & Keeping 1951, p. 222). To visualize any skewness in the distribution of r, we plot asymmetric error bars with limits $\tanh(z \pm \delta z)$.

This paper has been typeset from a TEX/LATEX file prepared by the author.