

Linear systematics mitigation in galaxy clustering in the Dark Energy Survey Year 1 Data

Erika L. Wagoner¹,[★] Eduardo Rozo,¹[★] Xiao Fang²,³ Martín Crocce³,³ Jack Elvin-Poole^{4,5} and Noah Weaverdyck^{6,7} (DES Collaboration)

¹Department of Physics, University of Arizona, 1118 E. Fourth Street, Tucson, AZ 85721, USA

²Department of Astronomy and Steward Observatory, University of Arizona, 933 N Cherry Ave, Tucson, AZ 85719, USA

³Institut de Ciències de l'Espai, IEEC-CSIC, Campus UAB, Carrer de Can Magrans, s/n, E-08193 Bellaterra, Barcelona, Spain

⁴Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210, USA

⁵Department of Physics, The Ohio State University, Columbus, OH 43210, USA

⁶Department of Physics, University of Michigan, 450 Church St, Ann Arbor, MI 48109-1040, USA

⁷Leinweber Center for Theoretical Physics, University of Michigan, 450 Church St, Ann Arbor, MI 48109-1040, USA

Accepted 2021 March 4. Received 2021 March 3; in original form 2020 September 28

ABSTRACT

We implement a linear model for mitigating the effect of observing conditions and other sources of contamination in galaxy clustering analyses. Our treatment improves upon the fiducial systematics treatment of the Dark Energy Survey (DES) Year 1 (Y1) cosmology analysis in four crucial ways. Specifically, our treatment (1) does not require decisions as to which observable systematics are significant and which are not, allowing for the possibility of multiple maps adding coherently to give rise to significant bias even if no single map leads to a significant bias by itself, (2) characterizes both the statistical and systematic uncertainty in our mitigation procedure, allowing us to propagate said uncertainties into the reported cosmological constraints, (3) explicitly exploits the full spatial structure of the galaxy density field to differentiate between cosmology-sourced and systematics-sourced fluctuations within the galaxy density field, and (4) is fully automated, and can therefore be trivially applied to any data set. The updated correlation function for the DES Y1 redMaGiC catalogue minimally impacts the cosmological posteriors from that analysis. Encouragingly, our analysis does improve the goodness-of-fit statistic of the DES Y1 3×2 pt data set ($\Delta\chi^2 = -6.5$ with no additional parameters). This improvement is due in nearly equal parts to both the change in the correlation function and the added statistical and systematic uncertainties associated with our method. We expect the difference in mitigation techniques to become more important in future work as the size of cosmological data sets grows.

Key words: methods: data analysis – methods: statistical – galaxies: photometry – dark energy – cosmology: observations.

1 INTRODUCTION

Galaxies are an important tool for studying the distribution of matter in the Universe and testing cosmological models. The correlation between the shapes of galaxies can be used to measure the weak gravitational lensing field, which is a direct measure of the total mass along lines of sight (see e.g. Mellier 1999; Refregier 2003). Because galaxies are a biased tracer of the mass, their clustering also carries within it cosmological information (e.g. Nishimichi et al. 2020). In addition, the tangential distortion of background galaxies around the position of foreground galaxies – usually referred to as galaxy–galaxy lensing – can be used to study the correlation between the foreground galaxies and the matter around them (e.g. Sheldon et al. 2004; Mandelbaum et al. 2006; Yoo & Seljak 2012). The combination of all three measurements, sometimes referred to as ‘ 3×2 pt’ for the use of three 2-point functions, can break the degeneracy between the galaxy bias and the clustering amplitude of matter (see e.g. Krause &

Eifler 2017, and references therein). These 3×2 pt analyses are the focus of ongoing surveys such as the Kilo Degree Survey (KiDS; Hildebrandt et al. 2017), the Dark Energy Survey (DES; The Dark Energy Survey Collaboration 2005, 2016), and the Hyper Suprime-Camera Survey (HSC; Aihara et al. 2018). Many other analyses have also applied a similar combined-probes approach to a variety of data sets (see e.g. Mandelbaum et al. 2013; More et al. 2015; Abbott et al. 2018; van Uitert et al. 2018; Joudaki et al. 2018). These types of combined analyses can also help to mitigate systematics that impact only one of the three 2-point measurements.

Wide field stage II dark energy experiments [such as the Sloan Digital Sky Survey; York et al. 2000, the WiggleZ DES (Drinkwater et al. 2010)] and the Canada–France–Hawaii Telescope Legacy Survey [Cuillandre et al. 2012 and stage III dark energy experiments (e.g. KiDS, DES, HSC, and eBOSS)] have provided imaging and spectra for hundreds of millions of galaxies, and stage IV experiments such as the Dark Energy Spectroscopic Instrument (Levi et al. 2013), the Rubin Observatory’s Legacy Survey of Space and Time (Ivezic et al. 2019), the *Nancy Grace Roman Space Telescope* (Spergel et al. 2015), and *Euclid* (Laureijs et al. 2011) are expected to increase that number substantially. As the number of observed galaxies in-

* E-mail: wagoner47@email.arizona.edu (ELW); erozo@email.arizona.edu (ER)

creases, the statistical uncertainty on measurements made with them decreases. Consequently, our understanding and treatment of the systematic effects that impact galaxy clustering measurements must be improved if the uncertainties on the inferred cosmological parameters from such galaxy surveys are to remain statistics dominated.

There are a large number of potential contaminants that can result in this type of coherent fluctuations, e.g. star–galaxy separation, stellar occultation, extinction, and variations in observing conditions such as airmass or sky brightness. Differentiating between the true cosmologically sourced fluctuations and those caused by such survey properties has been the subject of many studies over the years (see e.g. Suchyta et al. 2016; Elvin-Poole et al. 2018; Nicola et al. 2020; Rezaie et al. 2020; Weaverdyck & Huterer 2020, and references therein). Rezaie et al. (2020) identify three broad categories of mitigation techniques: (a) Monte Carlo simulation of fake objects, (b) mode projection, and (c) regression.

The first of these methods, involving injecting artificial sources into real images, is extremely promising. It results in forward modelling the survey selection mask imposed by real imaging properties. Examples of this method include Bergé et al. (2013) and Suchyta et al. (2016). However, this technique is computationally expensive, and therefore less utilized than the other methods.

Techniques utilizing mode projection typically involve down-weighting the spatial modes that are strongly correlated with survey properties by assigning a large variance to them. This technique has been explained and utilized in, e.g. Rybicki & Press (1992) and Nicola et al. (2020). The variance of the estimated clustering increases as more survey properties are considered unless a threshold is used to limit the number of survey property maps. However, using such a threshold has been shown to introduce a bias in the resulting two-point function (Elsner, Leistedt & Peiris 2016).

Regression-based techniques attempt to model the impact of the survey properties on the galaxy density, fitting the parameters of the model by cross-correlating the galaxies and systematic fluctuations or by using a least-squares estimate. For instance, Ross et al. (2011) and Ho et al. (2012) fit for the impact of observing conditions in the correlation function and power spectrum, respectively. The disadvantage with this method is that any spurious correlation between the 2-point function of the galaxies and the survey properties will result in a correction, even if the fluctuations are not spatially related. This makes it easy to overcorrect for systematic fluctuations that may bias the resulting correlation function estimate, although Weaverdyck & Huterer (2020) show how the pseudo-Cl implementation of mode projection can be interpreted as an ordinary least-squares regression approach that accounts for this overcorrection.

As part of the analysis of the DES Y1 ‘Gold’ data release, Elvin-Poole et al. (2018, hereafter Paper I) also fit for the impact of survey properties, but using one of the alternative suggestions from Ross et al. (2011) of applying the corrections one at a time in order to account for potential correlations between different sources of systematic fluctuations. Briefly, the method of Paper I is as follows: the average number of galaxies per pixel N_{gal} is measured for all pixels with a survey property value s within a bin $s \in [s_{\text{min}}, s_{\text{max}}]$ for one of the survey property maps, relative to the average number of galaxies per pixel in all pixels ($\langle N_{\text{gal}} \rangle$). A model is fit across all bins of the survey property values, and the $\Delta\chi^2$ for this model compared to a null test where $N_{\text{gal}}/\langle N_{\text{gal}} \rangle = 1$ is calculated. The significance of the survey property map is defined by comparing this $\Delta\chi^2$ to the 68th percentile of the equivalent quantity measured in 1000 contaminated Gaussian mock catalogues. This procedure is repeated for each survey property map and the maps are ranked by significance. A

correction is applied for the most significant map to the measurements of $N_{\text{gal}}/\langle N_{\text{gal}} \rangle$, and the significance of each map is re-calculated. To avoid overcorrection, this iterative process continues until none of the survey property maps have a significance above some target threshold. However, it is not necessarily the case that the effects of the various survey properties can be separated in this manner. For instance, this method precludes the possibility that significant systematic fluctuations can arise from the coherent contribution of multiple sources of systematics despite each individual survey property map being negligible by itself. Also, the analysis in Paper I included the spatial structure of the galaxy distribution only through the covariance in the galaxy densities binned by survey property. The analysis method introduced in this paper explicitly incorporates the density and spatial separations of neighboring pixels for determining the coefficients of the fluctuations sourced by survey properties: it is a much finer-grained look at that spatial structure.

Several other recent studies have attempted to use the regression-based technique directly with the galaxy density field while incorporating the spatial structure of the galaxy density field (see e.g. Prakash et al. 2016; Delubac et al. 2017). However, as discussed in Rezaie et al. (2020), these models are also often vulnerable to overcorrection. The regression method used by Rezaie et al. (2020) differs from previous regression-based techniques in that it does not assume a functional form for the impact of the survey properties on the observed galaxy density. Instead, Rezaie et al. (2020) rely on a neural network approach and feature selection to achieve accurate systematic corrections without overcorrection. However, this method fails to propagate the statistical and systematic uncertainties due to the correction into the error budget of the galaxy clustering signal.

In this paper, we implement an improved version of the linear model described in Prakash et al. (2016). Relative to that work, we reduce the number of free parameters by one by enforcing the condition that in the absence of systematic fluctuations, the observed galaxy density field will be equal to the true galaxy density field with a mean of zero (i.e. we do not include the constant term in equations 13 and 14 of that paper as a free parameter in our model). Our analysis explicitly incorporates the spatial clustering signal of the galaxy density field in an iterative approach, and mock catalogues are used to calibrate and correct for the residual bias due to overcorrection. The combination of using a Markov chain Monte Carlo (MCMC) to fit our model and utilizing mock catalogues to correct for the bias allows us to estimate both the statistical and systematic uncertainty of our systematics-corrected galaxy correlation function. Our procedure therefore correctly inflates the error budget associated with the measurement of the galaxy correlation function, enabling us to trivially propagate these uncertainties into cosmological constraints downstream. We apply our model to the DES Y1 Gold redMaGiC catalogue, and compare our results to those from Paper I and Abbott et al. (2018).

The paper is organized as follows: in Section 2, we describe the redMaGiC catalogue and the survey properties we use. We describe our method in Section 3. The generation of our mock catalogues and the results of the validation in the mocks are discussed in Section 4. We determine the impact of our systematics correction on the uncertainty in the correlation function in Section 5. Our results are presented in Section 6, and we summarize our findings in Section 7.

2 DATA

We will estimate and correct for systematic-sourced fluctuations in the density of the DES Year 1 redMaGiC galaxy sample (Paper I).

Table 1. The redshift binning with information about the number of galaxies and number density both from the DES Y1 analysis and the current analysis. The second and fourth columns are the total number of galaxies in each of the redshift bins, while the third and fifth give the galaxy density per square arcminute. Note that there is a change in the mask in going from the Y1 counts and number density of columns two and three to our own in columns four and five, which reduces the area by ~ 3.5 per cent.

z range	Y1 N_{gal}	Y1 n_{gal} (arcmin $^{-2}$)	N_{gal}	n_{gal} (arcmin $^{-2}$)
$0.15 < z < 0.3$	63 719	0.0134	61621	0.0134
$0.3 < z < 0.45$	163 446	0.0344	157800	0.0344
$0.45 < z < 0.6$	240 727	0.0506	231649	0.0505
$0.6 < z < 0.75$	143 524	0.0302	138450	0.0302
$0.75 < z < 0.9$	42 275	0.0089	40812	0.0089

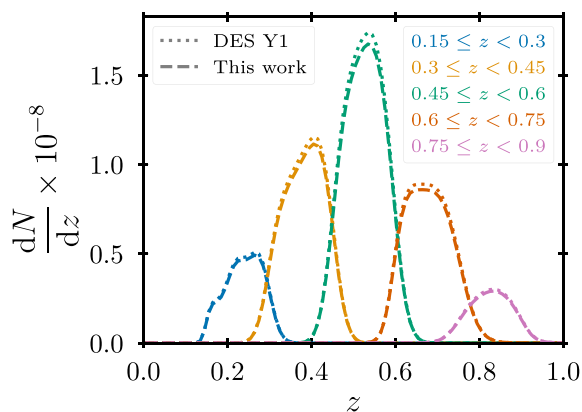


Figure 1. The redshift distribution for each redshift bin, found by stacking Gaussian distributions with mean and standard deviation equal to the redMaGiC redshift and error. The dashed lines are the distributions with our new mask, and the dotted lines of the corresponding colors are the corresponding distributions from Paper I in the same bins. The curves are not normalized, so differences in height are from the number of galaxies in the bin.

We use the same redshift binning as the Y1 analysis, shown here in Table 1, along with the number count and galaxy density in each bin. As described in Section 3, our analysis leads us to remove survey regions with large systematic-sourced fluctuations. This cut removes ~ 3.5 per cent of the fiducial Y1 redMaGiC footprint, for a final area of ≈ 1274 deg 2 . The counts and galaxy density after our systematic cut is shown in the fourth and fifth columns of Table 1. Fig. 1 compares the redshift distributions in each bin before and after the systematics cuts. The dotted lines of various colors are the distributions for the full redMaGiC sample, while the dashed lines of the same colour are the distributions in the same bin after cutting based on systematics. The distributions are not normalized, so differences in height are caused by the difference in the number of galaxies before and after the cut.

We estimate the galaxy correlation function using the Landy & Szalay (1993) estimator:

$$\hat{w}(\theta) = \frac{DD - 2DR + RR}{RR}, \quad (1)$$

where DD , DR , and RR are the number of pairs of galaxies with angular separation θ given a galaxy sample D and a random catalogue

R . We measure the number of pairs using `TreeCorr`¹ (Jarvis, Bernstein & Jain 2004), and our random catalogue is the same one used by Paper I, except for the fact that we remove the random points in the survey regions excluded by our analysis.

We consider a total of 18 potential sources of systematics for the observed galaxy correlation function. Each of these is represented as a map that is pixelated on the sky using the `HEALPix`² (Górski et al. 2005) pixelization scheme. The majority of the maps we consider are imaging properties from the DES Y1 ‘GOLD’ catalogue release (Drlica-Wagner et al. 2018). In each of the four bands ($griz$), we have maps of

- (i) total exposure time;
- (ii) mean PSF full width at half-maximum;
- (iii) mean sky brightness, due to e.g. the moon; and
- (iv) mean airmass.

For all mean quantities, the value on a pixel in a given band is computed as the weighted mean over all exposures in that band that contribute to the pixel. The exposure time is instead the sum of the exposure times for each exposure contributing to the pixel. Unlike Paper I, we do not include any depth maps, as these depend on the other imaging properties in a complicated way, and therefore are not linearly independent from the other imaging properties – including the depth maps would be double counting the other imaging properties and would likely increase any overcorrection biases that might exist. We therefore have 16 imaging property maps. We also consider contamination due to foreground stars, for which we use the stellar density map described in section 5 of Paper I. Galactic extinction is included using the dust opacity map from the Planck Collaboration (Planck Collaboration XI 2014). Both stellar density and extinction were considered in Paper I, but were not found to have significant correlation with the galaxy density and thus were ultimately excluded from that correction. We include both here because we do not want to preclude the possibility that they could still add coherently with other potential sources of contamination and thus impact the observed galaxy density. We note that the Planck dust map is known to have some positive correlation with the galaxy density in some redshift bins, as do most other Galactic dust maps (see e.g. Chiang & Ménard 2019). Collectively, we refer to the set of 18 imaging property, stellar density, and Galactic extinction maps as ‘survey property maps’ or SPs. Where necessary, we use the routines of `healpy` (Zonca et al. 2019) for manipulating both survey property and density maps.

3 METHOD

We determine the impact of observing conditions in the clustering of galaxies by relying on the spatial structure of the survey properties. Specifically, we estimate the extent to which the galaxy density maps are contaminated by systematic fluctuations by measuring the extent to which the galaxy density map traces the various survey property maps.

We begin by constructing a low-resolution ($N_{\text{side}} = 128$) map of the galaxy density field. This choice limits the number of empty pixels to a small percentage (≤ 10 per cent) of the total pixels. Working at this resolution, the average number of galaxies per pixel is ≥ 10 at all redshifts.

¹<http://ascl.net/1508.007>

²<https://healpix.sourceforge.net>

We degrade the resolution of our survey property maps to match the resolution of our galaxy density map, properly accounting for the masked portions of every pixel. Specifically, the degraded survey property map \mathcal{S}' is related to the original survey property map \mathcal{S} via

$$\mathcal{S}'^j = \frac{\sum_{i \in j} \mathcal{S}^i f^i}{\sum_{i \in j} f^i}, \quad (2)$$

where $f^i \in [0, 1]$ is the fraction of pixel i (at the original map resolution) that is detected in the footprint. The sums are over all high resolution pixels i that fall within low resolution pixel j . We also degrade the pixel fraction map f^i , such that the fraction f^{ij} of low-resolution pixel j in the footprint is related to the high resolution fraction by

$$f^{ij} = \frac{1}{\tilde{N}} \sum_{i \in j} f^i, \quad (3)$$

where \tilde{N} is the number of high-resolution pixels within a low-resolution pixel.

The degraded survey property maps are transformed into standardized fluctuation maps as follows. Let \mathcal{S}'_β be the value of survey property map β on low resolution pixel j . We define the mean $\bar{\mathcal{S}}_\beta$ and fluctuation scale $\hat{\sigma}_\beta$ of \mathcal{S}'_β via

$$\bar{\mathcal{S}}_\beta \equiv \sum_{j=1}^{N_{\text{pix}}} f^{ij} \mathcal{S}'_\beta / \sum_{j=1}^{N_{\text{pix}}} f^{ij} \quad (4)$$

and

$$\hat{\sigma}_\beta \equiv 1.4826 \text{MAD}(\mathcal{S}'_\beta). \quad (5)$$

The median absolute deviation in equation (5) is

$$\text{MAD}(\mathcal{S}'_\beta) \equiv \sum_{j=1}^{N_{\text{pix}}} \left| \mathcal{S}'_\beta - \text{med}(\mathcal{S}'_\beta) \right| / N_{\text{mask}},$$

where N_{mask} is the number of pixels not removed by the mask. The ‘fluctuation scale’ $\hat{\sigma}_\beta$ defined above is an estimator of the standard deviation for Gaussian fluctuations, but its value is more robust to outliers than estimates based on the sample variance. The standardized fluctuation map for survey property β is defined as

$$S_\beta^j \equiv \frac{\mathcal{S}'_\beta - \bar{\mathcal{S}}_\beta}{\hat{\sigma}_\beta}. \quad (6)$$

Rather than working with the fluctuation maps themselves, we construct an orthogonal map eigenbasis as follows. We assume the survey properties on each pixel are an independent random realization from an N_{maps} -dimensional distribution. We find the covariance matrix \mathbf{C} of the standardized maps at the fit resolution, where

$$\mathbf{C}_{\alpha\beta} = \langle (S_\alpha - \langle S_\alpha \rangle) (S_\beta - \langle S_\beta \rangle) \rangle,$$

and $\langle \cdot \rangle$ is the spatial average over all observed pixels. We define the rotation matrix \mathbf{R} from the eigenvectors of \mathbf{C} such that

$$\mathbf{C} = \mathbf{R} \mathbf{D} \mathbf{R}^\top,$$

where \mathbf{D} is a diagonal matrix with the eigenvalues of \mathbf{C} along the diagonal. The rotated and standardized survey property value for map α on pixel j is

$$s_\alpha^j \equiv \mathbf{R}_{\alpha\beta}^\top S_\beta^j. \quad (7)$$

Each s_α^j is therefore a linear combination of the fluctuations in the original SP maps $\{\mathcal{S}'_\beta\}$ on a given pixel. For the rest of the paper,

unless otherwise noted, the term ‘SP’ refers to the eigenmap s_α^j of equation (7) rather than the original survey property map \mathcal{S}'_β .

Since fluctuations in the density field cannot be sensitive to a constant non-zero SP value – any non-zero constant would simply shift the mean value of the galaxy density field – the observed galaxy density must only depend upon the fluctuations of the SPs. Thus, we write $\delta_{\text{obs}}^j \equiv \delta_{\text{obs}}(\{s_\alpha^j\})$, where $\{s_\alpha^j\}$ is a vector containing the value of pixel j across all SP maps α . Expanding around $\{s_\alpha^j\} = \vec{0}$ to first order, we have

$$\delta_{\text{obs}}^j(\{s_\alpha^j\}) \approx \delta_{\text{true}}^j + \sum_{\alpha} a_{\alpha} s_{\alpha}^j, \quad (8)$$

where the coefficient a_{α} is the derivative of δ_{obs} with respect to s_{α} at $\{s_\alpha^j\} = \vec{0}$. Note that any impact on the monopole of the galaxy density field by the survey properties gets absorbed into the mean observed galaxy density, and therefore has no impact on the galaxy fluctuations. Since our expansion is at first order, we can ignore the monopole as any impact with couplings to the linear perturbations would be second order. In the expansion, we have used the fact that $\delta_{\text{obs}}^j(\{s_\alpha^j\} = \vec{0}) = \delta_{\text{true}}^j$, where δ_{true}^j is the true galaxy overdensity on pixel j . We have also assumed that the impact of SP on the galaxy density field is local: the SP in pixel j only impact the galaxy density at pixel j .

Our task is to find the set of coefficients $\{a_{\alpha}\}$ in equation (8). We do this by fitting the likelihood $P(\vec{\delta}_{\text{obs}} | \vec{\delta}_{\text{sys}})$ of the observed overdensity map given the systematics map $\vec{\delta}_{\text{sys}} \equiv \sum_{\alpha} a_{\alpha} \vec{s}_{\alpha}$, where the vector symbol denotes the full map. As discussed below, our procedure allows for covariance between pixels, so that this likelihood distribution does not in general reduce to a product over all pixels. We assume a Gaussian likelihood for $\vec{\delta}_{\text{obs}}$. This explains why it is important for the mean number of galaxies in the galaxy density map to be large. We test our sensitivity to using a Gaussian distribution in Section 4.2. The ensemble average over realizations of the observed density field at fixed systematics is simply

$$\langle \vec{\delta}_{\text{obs}} \rangle = \vec{\delta}_{\text{sys}}. \quad (9)$$

We can thus write our Gaussian likelihood for $\vec{\delta}_{\text{obs}}$ as

$$\ln P(\vec{\delta}_{\text{obs}} | \vec{\delta}_{\text{sys}}) = -\frac{1}{2} \log |\boldsymbol{\Sigma}^{\text{obs}}| - \frac{1}{2} (\vec{\delta}_{\text{obs}} - \vec{\delta}_{\text{sys}})^\top (\boldsymbol{\Sigma}^{\text{obs}})^{-1} (\vec{\delta}_{\text{obs}} - \vec{\delta}_{\text{sys}}), \quad (10)$$

where we have dropped all constant terms, and again

$$\vec{\delta}_{\text{sys}} = \sum_{\alpha} a_{\alpha} s_{\alpha}^j. \quad (11)$$

The model parameters characterizing $\vec{\delta}_{\text{sys}}$ are the coefficients a_{α} for each survey property, which we aim to recover from the data. With this notation, both $\vec{\delta}_{\text{sys}}$ and $\vec{\delta}_{\text{obs}}$ are vectors of length N_{pix} and $\boldsymbol{\Sigma}^{\text{obs}}$ is an $N_{\text{pix}} \times N_{\text{pix}}$ matrix, where N_{pix} is the number of pixels within the footprint (i.e. the number of observed pixels).

The covariance matrix for our likelihood can be written as the sum of two terms:

$$\boldsymbol{\Sigma}^{\text{obs}} = \boldsymbol{\Sigma}^{PN} + \boldsymbol{\Sigma}^{SV}. \quad (12)$$

The first term contains the Poisson noise in the density field, and takes the form

$$\boldsymbol{\Sigma}^{PN}_{jk} = \sigma_g^2 \delta_{jk},$$

where σ_g is a constant for which we can fit and δ_{jk} is the Kronecker delta. It will become clear shortly why we allow σ_g to be an unknown

constant, rather than fixing it to the Poisson expectation. The second term in equation (12) accounts for the sample variance.

We fit for our SP coefficients in two iterations. During the first iteration, we assume there is no sample variance, so that Σ^{obs} is diagonal. In this case, we can analytically solve for the variance σ_g^2 and coefficients $\{a_\alpha\}$ that minimize the likelihood in equation (10) by solving the simultaneous set of equations obtained when setting all of the partial derivatives with respect to the survey parameter coefficients and σ_g^2 to zero. We are also able to find the 19×19 D parameter covariance matrix analytically as the inverse of the Hessian matrix evaluated at the minimum – we use this parameter covariance matrix (excluding the row and column corresponding to σ_g^2) in the second iteration to select random starting locations within the 18D parameter space.

Once we complete our first iteration, we use our results to estimate $\hat{\delta}_{\text{true}}$. We then define Σ^{SV} via

$$\Sigma^{\text{SV}}_{jk} = (1 - \delta_{jk}) \hat{w}_{\text{true}}(\theta_{jk}),$$

where \hat{w}_{true} is the correlation function of our estimated true overdensity field $\hat{\delta}_{\text{true}}$ and θ_{jk} is the angular separation between pixels j and k . We artificially set the diagonal elements of Σ^{SV} to zero because we cannot differentiate between the sample variance and Poisson noise within a single pixel. This also explains why we treated σ_g as an unknown constant: σ_g is really the sum of the Poisson and zero-offset sample variance terms. We therefore continue to use the σ_g obtained from the minimization in the first iteration as the only term on the diagonal of Σ^{obs} in the second iteration.

We use the resulting ‘Poisson’ and sample variance noise estimates to refit for the coefficients of each of the SP parameters. In the second iteration, we use an MCMC algorithm (specifically `emcee`; Foreman-Mackey et al. 2013) to sample our parameter space and estimate the posterior distribution while holding both Σ^{PN} and Σ^{SV} fixed. Our best-fitting coefficients after the second iteration are the mean parameter values from the chain.³ To check for convergence, we look at the shift in the coefficients between the first and second halves of each chain relative to the error from the chain. We find a median shift (over all 18 parameters) of 0.19, 0.29, 0.18, 0.26, and 0.14 for redshift bins 1 through 5, respectively, and the worst convergence in any single parameter for each redshift bin is 0.60, 0.72, 0.55, 0.53, and 0.34. We have verified that using the coefficients from the second iteration to update Σ^{SV} and performing a second MCMC (i.e. getting a third iteration of the coefficients) does not have a significant impact on our results.

Once we have our coefficients, we correct for the effect of systematic fluctuations on the correlation function. We do so by defining weights for each galaxy based on the systematics map value on the pixel containing the galaxy. For calculating galaxy weights, we use the systematics map at a resolution of $N_{\text{side}} = 4096$. While we must fit at low resolution to ensure that our likelihood is roughly Gaussian, the fundamental assumption of our method is that survey properties only produce local modulations of the galaxy density field. Since our model is linear, all the local modulations add together when smoothing to go to lower resolution, so the relation between the survey properties and the galaxy density must be the

same at low and high resolution. We standardize and rotate the high-resolution maps as we did with the low-resolution maps, but we use the mean, fluctuation scale, and rotation matrix determined from the low-resolution maps for the purposes of defining the high-resolution eigen-maps. This is critical, as the definition of the maps must match that employed in our fits. The weight for a galaxy on high-resolution pixel i is

$$w^i = \frac{1}{1 + \sum_{\alpha} a_{\alpha} s_{\alpha}^i}. \quad (13)$$

We refer to the correlation function measured using these weights as w_{corr0} . As previously mentioned, when calculating the systematics-corrected correlation function, we also exclude any galaxies on pixels with $\delta_{\text{sys}}^i > 0.2$. This should restrict us to only areas of the sky where our first-order approximation is valid. The resulting footprint is ~ 3.5 per cent smaller than the original Y1 footprint, and a total of 23 359 galaxies are removed across all redshift bins. We expect complications due to the interpolation to higher resolution to be small as we find that only ~ 12.8 per cent (~ 0.8 per cent) of galaxies have a weight that differs from unity by more than 10 per cent (20 per cent) before applying the cut based on δ_{sys}^i .

The above procedure tends to overcorrect the data for the impact of SPs. We calibrate the amount of overcorrection in the correlation function from our method using mock galaxy catalogues, and use these to de-bias our procedure, which will result in an updated systematics-corrected correlation function estimate w_{corr1} . The details of this de-biasing are presented in the next section. We describe how we incorporate statistical and systematic uncertainties due to our correction in the error budget of the observed correlation function in Section 5.

4 METHODOLOGY VALIDATION WITH MOCK CATALOGUES

There are three potential sources of systematic bias in our analysis. These are, in no particular order, (i) the first-order approximation from equation (8) is not accurate; (ii) the Gaussian likelihood is not correct; and (iii) the estimates of the SP coefficients are noisy and too much correlation is removed from the data, an effect usually referred to as overcorrection. As mentioned in Section 2, we restrict our final data set to pixels where the linear prediction of the SP-sourced galaxy density fluctuations are ≤ 0.2 . This serves to minimize potential biases from non-linear responses in the systematics correction. We test the robustness of our methodology to non-Gaussian fields and noise by testing it on lognormal mock galaxy catalogues. We further use these catalogues to calibrate the bias in our method due to overcorrection.

4.1 Mock catalogue generation

To create our lognormal mock catalogues, we use the fiducial cosmological parameters from Paper I: $\Omega_m = 0.295$, $A_s = 2.260574 \times 10^{-9}$, $\Omega_b = 0.0468$, $h = 0.6881$, and $n_s = 0.9676$. We run `CAMB` (Lewis, Challinor & Lasenby 2000; Howlett et al. 2012) and `Halofit_Takahashi` (Smith et al. 2003; Takahashi et al. 2012) using `CosmoSIS` (Zuntz et al. 2015) to compute the angular galaxy clustering power spectrum. We then use this power spectrum to generate a lognormal random field for the true galaxy overdensity, δ_{true} , in each of our five redshift bins via the code `psydocl`.⁴ This

³We run our chain with 36 walkers for 1000 steps each. We do not use a burn-in when fitting to the real data as we generate the initial positions by drawing from a multivariate Gaussian with a mean and covariance matrix given by the coefficients and parameter covariance from the first iteration. We use a burn-in of 300 steps per walker when fitting to mock catalogues.

⁴<https://bitbucket.org/niallm1/psydocl/src/master/>

galaxy density field is generated at high resolution ($N_{\text{side}} = 4096$). When appropriate (i.e. depending on the test being pursued, see below), we add systematic fluctuations to the galaxy density field using our linear model. We then calculate the expected number of galaxies in each pixel, taking into account the masked fraction in each pixel. Finally, we randomly place N galaxies within each pixel, where N is a Poisson realization of the expected number of galaxies.

We generate 100 independent realizations of δ_{true} for each redshift bin. Each realization is then used to create two mock catalogues, one with no SP contamination and another with SP applied using the best fit coefficients from our analysis of the DES Y1 data set. We refer to these as uncontaminated and contaminated mocks, respectively. Note that while both the uncontaminated and contaminated mocks share the same underlying overdensity fields, they have different Poisson realizations.

We use our methodology from Section 3 to estimate the impact of SPs in our mock galaxy catalogues, and compare the resulting corrected correlation function to the underlying true mock galaxy correlation function. To increase computational efficiency, we restrict our mock catalogues to the final mask employed in our analysis of the DES Y1 galaxies. That is, we do not re-apply the $\delta_{\text{sys}}^i \leq 0.2$ cut in every mock. Doing so would have forced us to recompute random pairs for every mock due to slight differences in the final footprint. Because systematic fluctuations are linear in the mock catalogue by construction, this additional restriction has no bearing on the conclusions drawn from our simulations. Unfortunately, this also means our mock catalogues do not allow us to test how sensitive our method is to non-linear contamination.

We test whether our contaminated mock galaxy catalogues have comparable levels of SP contamination to the data as follows. For the data and both sets of mock galaxy catalogues, we compute the raw observed correlation function, and the corrected correlation function w_{corr0} as described in Section 3. We then calculate the difference between these two correlation functions in all three cases.

The blue solid line in Fig. 2 shows the biased systematic correction of the DES Y1 redMaGiC data computed using the first iteration of our method, while the orange-dashed line is the mean correction from the 100 contaminated mock galaxy catalogues. The green-dashed line is the mean of the uncontaminated galaxy catalogues. The width of the bands shows the sample standard deviation for each of the two sets of mocks. It is immediately apparent that the amplitude of the systematic correction in our uncontaminated mocks is significantly smaller than that of the data in redshift bins 3, 4, and 5. That is to say, we have robustly detected the presence of systematic fluctuations in the DES Y1 data set. More generally, the correction derived from our contaminated mocks is comparable to that in the data, particularly for the redshift bins that exhibit strong systematic fluctuations. Thus, Fig. 2 provides evidence that the contaminated mock galaxy catalogues used in our analysis are a reasonable match to the data.

4.2 Methodology validation: recovery of the SP coefficients

We fit for the SP coefficients in both sets of 100 mocks for each redshift bin, for a total of 1000 independent mock catalogues to be analysed. Because we know the SP coefficients used to generate the mocks, we can test whether we correctly recover the input coefficients with our analysis. To do so, we calculate the χ^2 of the mean coefficients estimated from our posterior and the input for each mock. That is, for each mock catalogue ν we compute

$$\chi_\nu^2 = (\{\hat{a}_\alpha\}_\nu - \{a_\alpha\}_{0,\nu})^\top \hat{\mathbf{C}}_\nu^{-1} (\{\hat{a}_\alpha\}_\nu - \{a_\alpha\}_{0,\nu}), \quad (14)$$

where $\{a_\alpha\}_{0,\nu}$ is the input vector of 18 coefficients used in generating mock catalogue ν , $\{\hat{a}_\alpha\}_\nu$ is the mean vector of the posterior from our analysis for mock ν with length 18, and $\hat{\mathbf{C}}_\nu$ is the parameter covariance matrix estimated from the MCMC chain for mock ν with dimensions 18×18 . We show the distribution of the χ_ν^2 statistics for all 1000 mocks as the blue histogram in Fig. 3. For reference, the green line is the expected χ^2 distribution for 18 degrees of freedom, 18 being the number of SPs. It is clear that the distribution of χ^2 values is biased relative to our expectation.

Hartlap, Simon & Schneider (2007) pointed out that noise in the covariance matrix biases χ^2 statistics. In our case, the noise in the covariance matrix is only partly due to a finite number of realizations in the MCMC: noise in the data will also generate noise in the empirically estimated covariance matrix, which will in turn bias the recovered χ^2 . In the absence of a first principles prescription for the expected bias in our analysis, we adopt an ad hoc correction by demanding the average χ^2 over all our simulations be equal to the number of degrees of freedom in the problem (18). That is, we de-bias every χ^2 value by dividing it by the factor $\lambda \equiv 22.33/18 = 1.24$. The resulting distribution is shown as the orange histogram in Fig. 3, which is now an excellent match to expectations.

As discussed in Hartlap et al. (2007), the bias due to noise in the covariance matrix estimate propagates into the parameter posteriors. Consequently, we increase the statistical uncertainty in our recovered corrections for the correlation function by a factor of $\sqrt{1.24}$. The fact that our recovered distribution of χ^2 values matches expectation implies that we are successfully recovering the input systematic coefficients within our re-scaled noise estimate.

4.3 Overcorrection calibration

The orange-dashed line and shaded band in Fig. 4 show the mean and 1σ region for the difference between the observed and true correlation functions of our 100 independent systematics-contaminated mock catalogues, in units of the statistical uncertainty of the DES Y1 analysis. The 1σ region is computed as the error on the mean. The blue solid line and shaded band are the same as the orange, but for the systematics-corrected correlation function with no bias correction (i.e. w_{corr0}). While there is a significant improvement when going from no correction to our systematics correction, it is also clear that our method somewhat overcorrects the data.

We seek to calibrate the amount of overcorrection for our method based on the results from Fig. 4. However, *note that the level of overcorrection is itself sensitive to the input amount of contamination*. This is apparent in Fig. 5, which shows the mean and error on the mean of the overcorrection for both uncontaminated (orange) and contaminated (blue) mock galaxy catalogues.

We use the results in Fig. 5 to reduce the impact of overcorrection, and to characterize the remaining systematic uncertainty associated with this effect. Because we see that the level of overcorrection is sensitive to the amount of contamination and we do not know the actual contamination level in the data, we must account for this sensitivity when we de-bias. The contaminated and uncontaminated mocks represent the two extreme possibilities for the data, so we de-bias our correlation functions using the mean of the overcorrection measured in the contaminated and uncontaminated mocks. That is, we define

$$\Delta w(\theta) \equiv \frac{1}{2} [(w_{\text{corr0}}^{\text{cont}}(\theta) - w_{\text{true}}(\theta)) + (w_{\text{corr0}}^{\text{uncont}}(\theta) - w_{\text{true}}(\theta))], \quad (15)$$

where $w_{\text{corr0}}^{\text{cont}}(\theta)$ is the systematics-corrected correlation function at θ for the contaminated mock galaxy catalogues prior to de-biasing, and

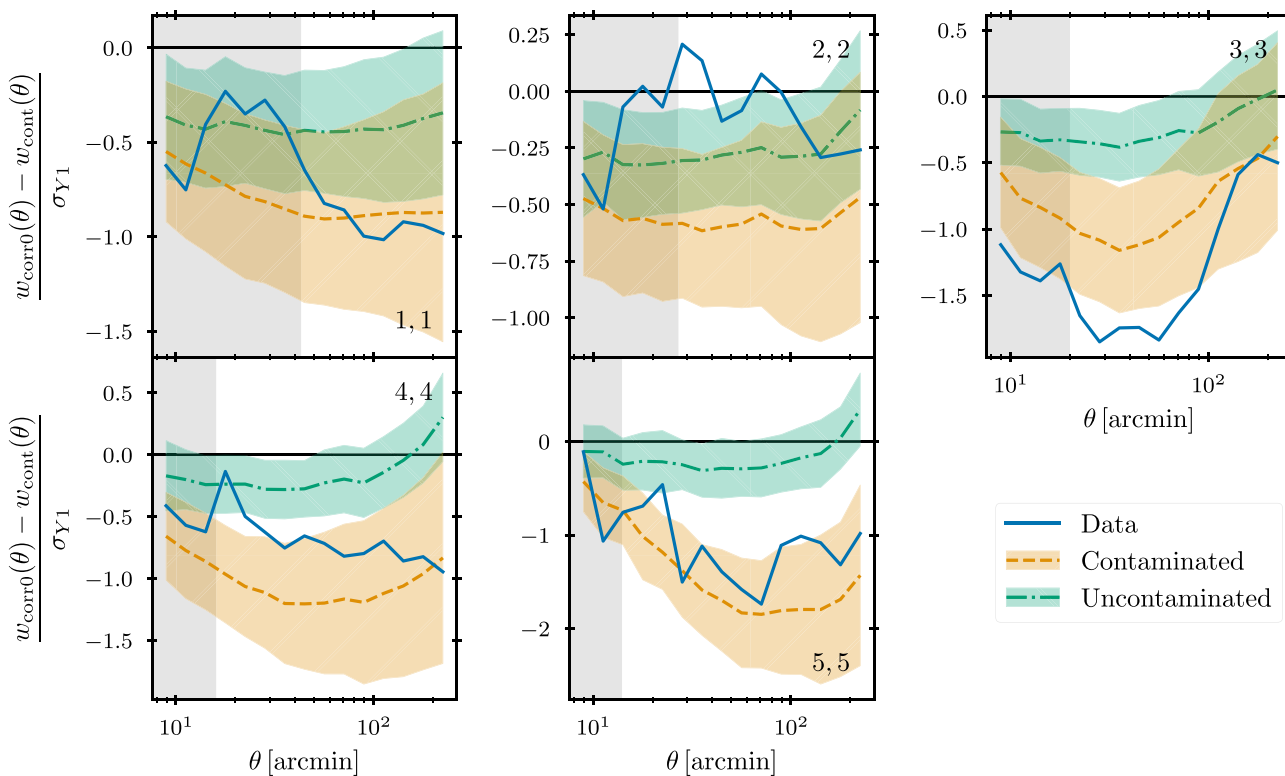


Figure 2. Comparison of the bias between the systematics-corrected ($w_{\text{corr}0}$) and uncorrected (w_{cont}) correlation functions for the DES Y1 data and the uncontaminated and contaminated mocks, relative to the DES Y1 errors (see text for details). The blue solid line is the result for the data. The mean and sample standard deviation for the contaminated mocks is shown as the orange-dashed line and the orange-shaded region, while the green dash-dotted line and green-shaded region show the same for the uncontaminated mocks. These error regions do not include the correction factor discussed in Section 4.2. By eye, we see 1σ agreement between the contaminated mocks and the data for three of the five redshift bins, and 2σ agreement in bins 2 and 3. The grey-shaded region is once again the small scale cut used by Paper I.




Figure 3. The distribution of χ^2 , as defined in equation (14), for all contaminated and uncontaminated mocks in all redshift bins. The blue histogram is the original distribution. The orange histogram is the result of re-scaling every χ^2 by $18/\langle\chi^2\rangle$. The green line is the expected χ^2 distribution with 18 degrees of freedom, for reference. Note that both histograms are normalized.

$w_{\text{corr}0}^{\text{uncont}}$ is the equivalent quantity computed for the uncontaminated mock galaxy catalogues. The average $\langle \cdot \rangle$ above is over the simulated data sets. The difference between the two terms in equation (15) is indicative of the systematic uncertainty of this bias correction, as explained in Section 5. Given Δw , we define an updated systematics-corrected correlation function $w_{\text{corr}1}$ via

$$w_{\text{corr}1}(\theta) \equiv w_{\text{corr}0}(\theta) - \Delta w(\theta). \quad (16)$$

The green dash-dotted line and shaded band in Fig. 4 show the mean and the 1σ region for the difference between our updated systematics-corrected correlation function estimates $w_{\text{corr}1}$ and the true correlation function, as estimated using 100 contaminated mock catalogues. Recall that the y-axis is scaled in units of the purely statistical uncertainty of the DES Y1 analysis. It is clear from the figure that while a residual bias remains, the amplitude and uncertainty is much smaller than the statistical uncertainties for the DES Y1 data set. Moreover, the true underlying correlation function is within the expected errors in the measurement.

5 THE IMPACT OF SYSTEMATICS REMOVAL ON THE NOISE

The covariance matrix used in Paper I when fitting the galaxy clustering signal was solely based upon theoretical considerations, as described in Krause et al. (2017). In particular, it accounted only for

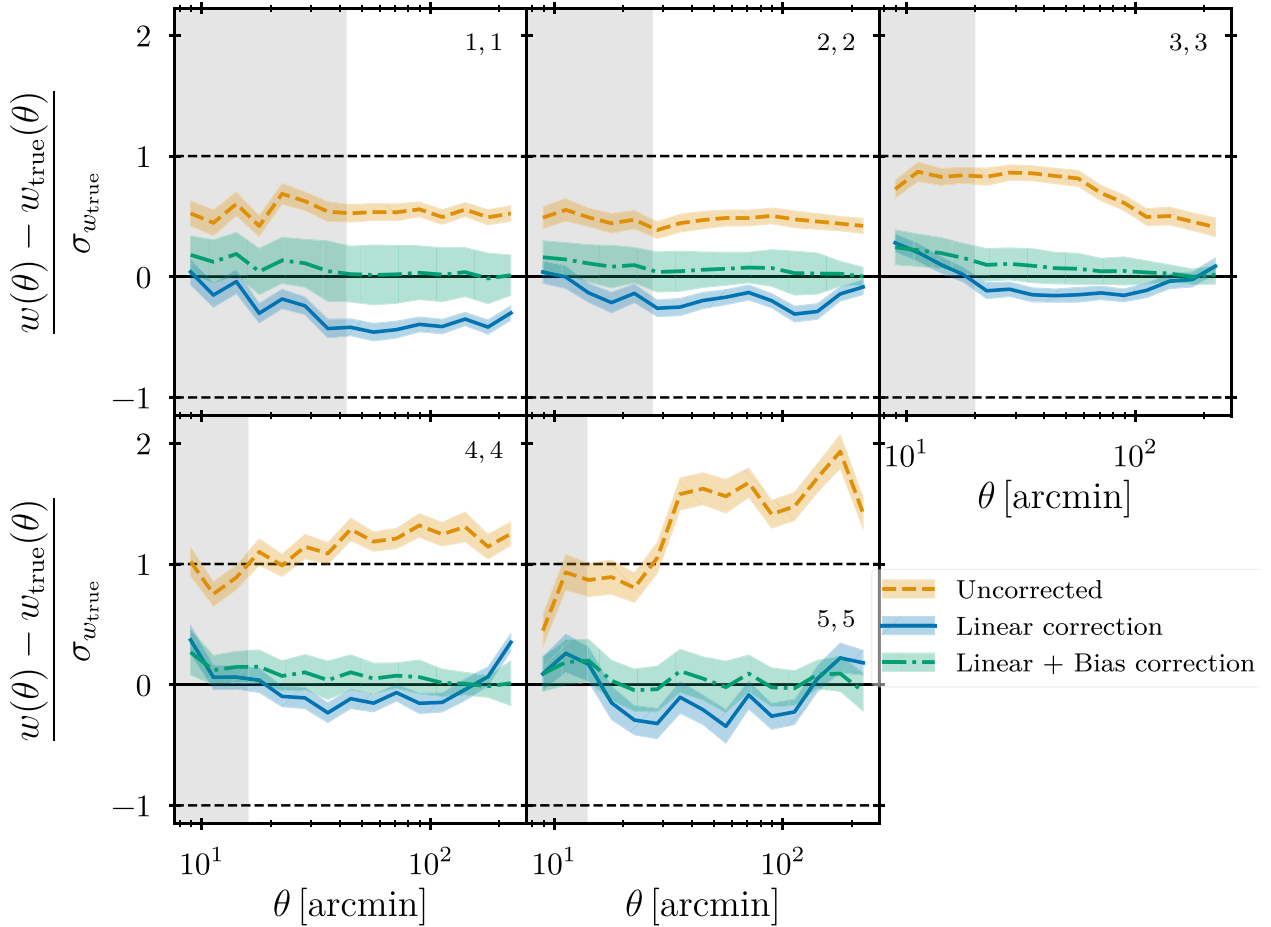


Figure 4. The difference between the various correlation functions for the contaminated mocks and the true correlation function. The orange dashed line shows the offset for the correlation function without any corrections. The blue solid line shows the offset when the systematics weights are applied, but no bias correction is used. The green dash–dotted line is the final offset, with both the systematics weights and the bias correction. Each line is the mean for the 100 mocks, and the shaded regions are the error on the mean. Note that the offset is also divided by the sample standard deviation of the true correlation function. We only scales with $\theta > 8$ arcmin for clarity. The grey-shaded region shows the small scale cut used by Paper I, so any scales within that region will not impact the cosmology results.

Poisson noise and sample variance in the galaxy density field, where the latter includes both Gaussian and connected terms, as well as the supersample covariance contribution. In practice, removing the imprint of systematic fluctuations on the galaxy density field carries with it additional uncertainty that needs to be propagated into the covariance matrix used to analyse the data. We now characterize this additional noise contribution.

We start with the statistical uncertainty in our method, the uncertainty due to the noise in our estimates of the linear coefficients of the SP maps. Because we use an MCMC to fit for the coefficients describing the impact of SPs, we can readily sample the posterior distribution to obtain realizations of the coefficients. For each such set of coefficients, we calculate the systematics-corrected correlation function $w_{\text{corr}0}$, resulting in many realizations of systematics-corrected correlation functions. We calculate the covariance matrix from these realizations, and re-scale it by the factor of $\langle \chi^2 \rangle / 18 = 1.24$ from the discussion in Section 4.2. This defines the statistical covariance matrix \mathbf{C}^{stat} , which characterizes statistical uncertainties in the systematics correction. We note that our estimation does not allow for statistical covariance in the SP corrections across the redshift bins. Since the noise in the coefficients of the SP maps depends on the galaxy

density field, which will be correlated across bins, this is not true in detail. However, we expect the relative quality of the redMaGiC photometric redshifts implies that any such correlation is small, particularly when propagated on to the SP coefficients.

In the process of de-biasing our systematics-corrected correlation functions, i.e. going from $w_{\text{corr}0}$ to $w_{\text{corr}1}$, our corrections may remove clustering modes from the correlation function. This will in turn remove some of the sample variance from $w_{\text{corr}1}$, particularly at large scales. To test for this possibility, we compared the variance in $w_{\text{corr}1}$ as measured in our simulations to the quantity

$$\hat{\text{Var}}(w_{\text{corr}1}) \equiv \text{Var}(w_{\text{true}}) + \text{diag } \mathbf{C}^{\text{stat}},$$

where \mathbf{C}^{stat} is the statistical covariance matrix including the re-scaling from the previous paragraph. The blue-shaded region in Fig. 6 shows the ratio of $\text{Var}(w_{\text{true}}) + \text{diag } \mathbf{C}^{\text{stat}}$ and $\hat{\text{Var}}(w_{\text{corr}1})$ averaged over all redshift bins. The width of the band represents the 1σ region around the mean. We see that at small scales the variance in our simulations is consistent with our error estimate. By contrast, we overestimate the error at large scales, likely due to the removal of clustering modes in our systematic correction algorithm. We have

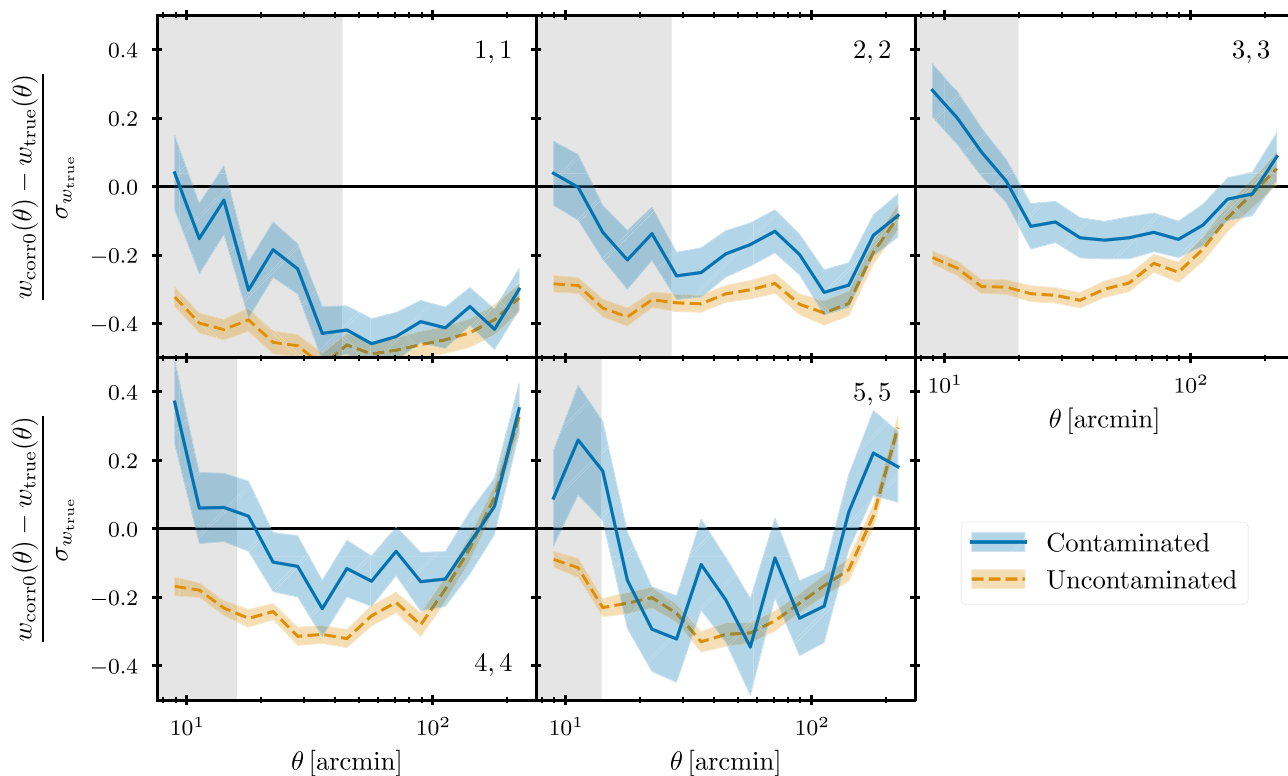


Figure 5. Bias in the systematics-corrected correlation function, relative to the sample standard deviation of the true correlation function. The orange-dashed line shows the mean bias for the 100 uncontaminated mocks, and the orange-shaded region is the error on the mean. Similarly, the blue solid line and shaded region are the mean and the error on the mean for the 100 contaminated mocks. Note that there is a non-trivial bias even for the contaminated mocks indicating that we are overcorrecting for SPs. We only show scales with $\theta > 8$ arcmin for clarity. The grey-shaded regions are once again the small scale cuts from Paper I.

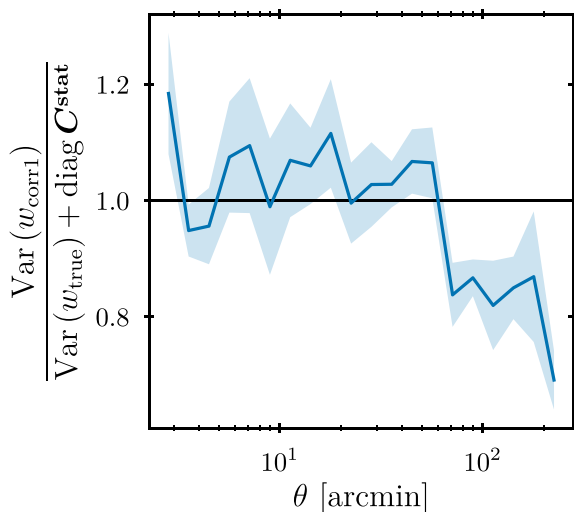


Figure 6. A comparison of $\text{Var}(w_{\text{corr1}})$ in our simulations and our revised error estimate $\text{Var}(w_{\text{true}}) + \text{diag } \mathbf{C}^{\text{stat}}$. Our revised error estimate correctly describes the variance of the simulations at small scales, but slightly overestimates the noise at large scales. This likely reflects reduced sample variance due to removal of clustering modes at large scales.

found that accounting for this reduced variance tends to make the covariance matrix of the 3×2 pt data vector non-invertible due to the cross-terms between probes. Because the systematic bias in the variance estimate is small, and because our revised errors have little impact on cosmological posteriors (see below), we will leave the problem of how to adequately model the resulting decrease in variance to future work.

The systematic uncertainty associated with our de-biasing procedure of Section 4.3 is calculated as the sum in quadrature of two distinct terms. The first term sets the systematic uncertainty to half the amplitude of the applied correction, i.e. large corrections will result in large uncertainties. The second term accounts for the difference in the amount of overcorrection inferred from the contaminated and uncontaminated mocks. If the inferred overcorrections are vastly different, the resulting mean correction should be assigned a large uncertainty. This uncertainty is set to half the difference between the overcorrection inferred from the contaminated and uncontaminated mocks. The corresponding covariance matrix characterizing these systematic uncertainties takes the form

$$\mathbf{C}^{\text{sys}}_{ab} \equiv \frac{1}{4} [\Delta w(\theta_a) \Delta w(\theta_b) + \delta w(\theta_a) \delta w(\theta_b)], \quad (17)$$

where a and b are the index angular bins, and where we have defined

$$\delta w(\theta) \equiv \frac{1}{2} [\langle w_{\text{corr0}}^{\text{cont}}(\theta) - w_{\text{true}}(\theta) \rangle - \langle w_{\text{corr0}}^{\text{uncont}}(\theta) - w_{\text{true}}(\theta) \rangle].$$

As in equation (15), the average $\langle \cdot \rangle$ above is over all simulated data sets.

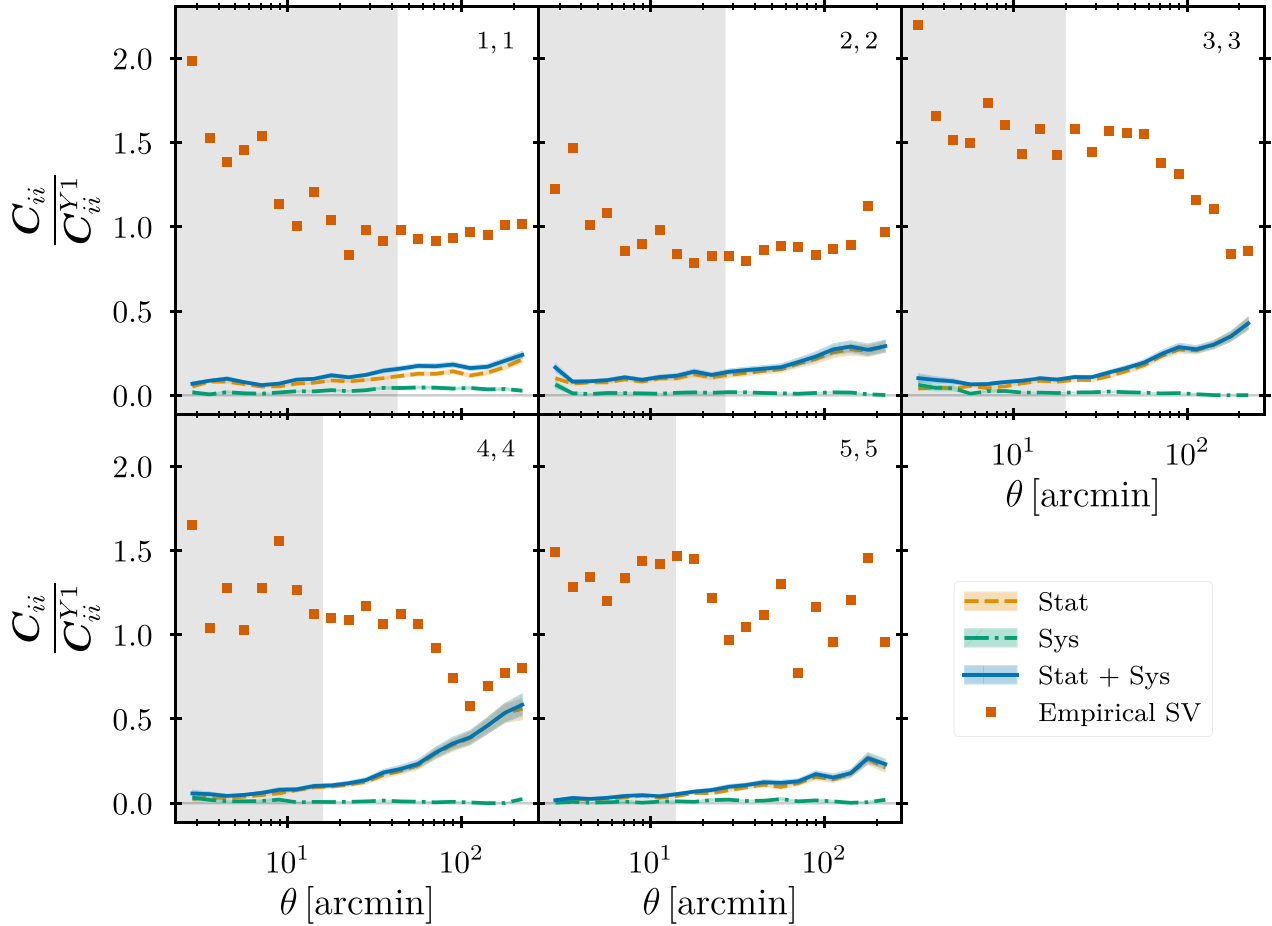


Figure 7. A comparison of the diagonal elements from various components of the covariance matrix relative to the diagonal elements of the theoretical covariance matrix utilized in Paper I. In all cases, the denominator of the quantity on the y-axis is the diagonal elements of C^{Y1} . The numerator for the blue solid line corresponds to our updated error estimates, including both statistical (orange dashed) and systematic (green dash-dotted) errors. The red squares show the corresponding empirical estimates of the variance in our simulations. The grey-shaded region is the small-scale cut from Paper I, and will not impact the cosmology results.

The final covariance matrix estimate for the data is $C^{Y1} + C^{\text{stat}} + C^{\text{sys}}$, where C^{Y1} is the theoretical covariance matrix used in Paper I. The green dash-dotted line and band in Fig. 7 show the mean and the uncertainty of the ratio between the diagonal elements of C^{sys} , as defined in equation (17), to the diagonal elements of C^{Y1} . The orange-dashed line and the band is the same ratio but for C^{stat} . We have checked that increasing the number of realizations used to estimate C^{stat} does not significantly change our measured covariance. The combination of the systematic and statistical covariance relative to the Y1 covariance is shown as the blue solid line and band. The grey-shaded region in each panel shows the region excluded by the small-scale cuts for the cosmology analysis in Paper I, for which our changes will not impact the inferred cosmological parameters. While uncertainties in our de-biasing procedure for overcorrection are negligible, we see that the statistical uncertainties in our systematics mitigation algorithm start to become comparable to statistical uncertainties in the correlation function at large scales.

6 RESULTS

As a brief summary of Sections 3, 4, and 5, we assume fluctuations in SPs introduce artificial galaxy fluctuations through a local linear

response. We calibrate these response coefficients using the observed galaxy density maps and SP maps, and use them to remove the impact of systematic fluctuations in the galaxy density field. Using mock galaxy catalogues, we demonstrate that our method results in some small amount of overcorrection, which we calibrate. We further characterize the additional statistical and systematic uncertainty introduced by our systematics-mitigation algorithm. We now apply our full systematics-correction algorithm to the DES Y1 data set.

In Fig. 8, we show the angular correlation function in each of the five redshift bins using our systematics weights and bias correction as the blue circles, with errors from the combined $C^{Y1} + C^{\text{stat}} + C^{\text{sys}}$ covariance matrix. For comparison, we also show the correlation function without correction and the systematics-corrected correlation function from Paper I. We note that in arriving at our updated correlation function, there is a small change in the mask to mitigate the impact of non-linear systematic fluctuations, so that the areas over which the correlation functions are computed are not precisely the same. The bottom panel in each figure shows the difference of each of the correlation function relative to the systematics-corrected estimate of Paper I. We see that the two different methods for estimating systematic corrections are in excellent agreement relative to the statistical uncertainty of the DES Y1 data set. Nevertheless, some

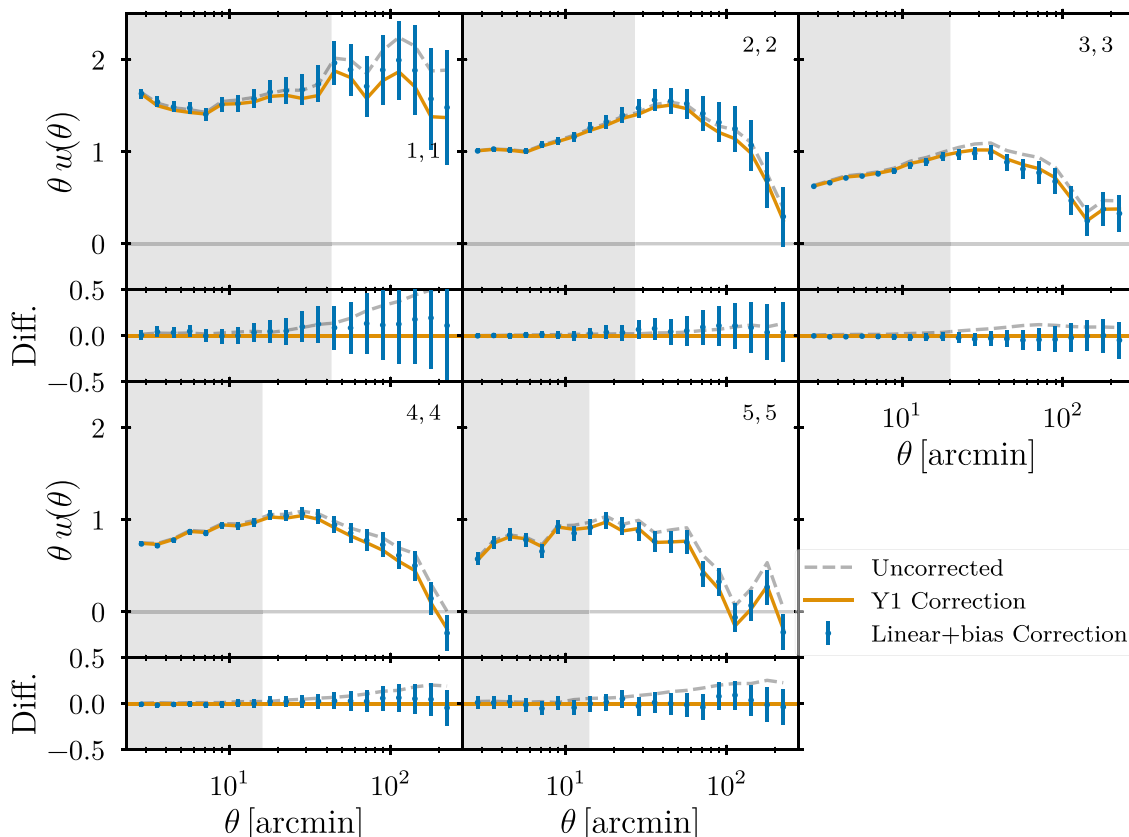


Figure 8. The correlation function in each redshift bin for the DES Y1 redMaGiC galaxies. The grey-dashed line is the correlation function without correcting for SPs. The orange solid line is the systematics-corrected correlation from Paper I. The blue points are the de-biased correlation function using our linear model weights, and the error bars are obtained from the full ($C^{Y1} + C^{\text{stat}} + C^{\text{sys}}$) covariance matrix. Note that while the grey and orange lines are computed with the DES Y1 mask, the blue points use our restricted mask with $\delta_{\text{sys}} \leq 0.2$, resulting in ~ 3.5 per cent less area.

small differences are clearly present. It is interesting to note that in the second redshift bin, our correction results in slightly *more* correlation than the uncorrected correlation function, rather than *less*. This boost is due to the overcorrection de-biasing procedure calibrated in the mocks.

To quantify the difference in the correlation functions from the two different weighting methods, Table 2 shows the χ^2 statistic for the DES Y1 correlation function and our correlation function, namely

$$\chi^2 = (w_{Y1}(\theta) - w_{\text{corr}}(\theta))^T \mathbf{C}^{-1} (w_{Y1}(\theta) - w_{\text{corr}}(\theta)),$$

where the choice of covariance matrix \mathbf{C} used requires some discussion (see below). In calculating χ^2 , we exclude any angular bins that are removed with the small-scale cut (the grey regions in Fig. 8). The number of remaining angular bins after the small-scale cut is shown in the last column of the table. The difference between the correlation functions should not be subject to Poisson noise or sample variance, as these are the same for both correlation functions. Therefore, in the second column of Table 2, we show the χ^2 when we use $\mathbf{C} = \mathbf{C}^{\text{stat}} + \mathbf{C}^{\text{sys}}$, with the probability to exceed the given χ^2 shown in parentheses. While in principle this comparison should also be subject to the uncertainty due to the method of Paper I, that paper demonstrated that the uncertainties in their systematics correction did not impact the cosmological priors and therefore those uncertainties were not characterized. Consequently, our comparison does not account for the uncertainty in the Y1 systematics correction. It is clear that our weights method results

Table 2. The χ^2 for the systematics-corrected correlation function from Paper I and this work in each redshift bin. The last column is the number of angular bins used to calculate the χ^2 , which are the bins outside the small scale cut represented by the grey-shaded regions in Fig. 8. The second column is the χ^2 when including only the uncertainty from the systematics correction, while the third column is the χ^2 relative to the full covariance matrix. The fourth and fifth columns are the same as the second and third, but the $\delta_{\text{sys}} > 0.2$ mask is applied to the galaxies with the DES Y1 weights. The numbers in parentheses in the second and fourth columns show the probability to exceed the χ^2 given the number of angular bins in the last column (the probability to exceed is ~ 1.0 for all bins in both the third and fifth columns). Our updated correlation function is only consistent with that from Paper I in the first redshift bin, but it is consistent with the correlation function with our mask and the Y1 weights in bins 1, 3, and 5.

z bin	$\chi^2_{\text{stat+sys}}$	χ^2_{tot}	$\chi^2_{\text{stat+sys}}$	χ^2_{tot}	# Bins
1	18.02(0.02)	0.463	9.563(0.30)	0.286	8
2	120.0(0.00)	2.04	57.03(0.00)	1.37	10
3	97.46(0.00)	0.900	26.71(0.01)	0.502	11
4	45.65(0.00)	0.922	67.91(0.00)	0.611	12
5	344.6(0.00)	1.89	22.82(0.04)	0.354	13

in a correlation function that is formally inconsistent with that of the Y1 analysis assuming zero uncertainty from the Y1 weights method. However, the size of the cosmology contours is sensitive to the full covariance matrix $\mathbf{C}^{Y1} + \mathbf{C}^{\text{stat}} + \mathbf{C}^{\text{sys}}$. The third column in Table 2

shows the χ^2 when we use the full covariance matrix for \mathbf{C} . Notice that in this case, the $\chi^2/\text{dof} \leq 0.1$ for most redshift bins. This result explicitly demonstrates that the difference in the correlation function produced by the two methods is small relative to the statistical uncertainty.

While our updated correlation function is inconsistent with the correlation function presented in Paper I when excluding the statistical uncertainty, the difference between them is actually sourced by two effects: the difference in the weights produced by the two corrections, and the difference in the footprint. In particular, removal of the pixels with $\delta_{\text{sys}}^i > 0.2$ improves the agreement between the two. To determine whether the corrections are consistent after accounting for the different masks, we recompute the galaxy correlation function using the fiducial Y1 weights over our new mask. The χ^2 statistics for this comparison are shown in the fourth and fifth columns of Table 2. As before, the $\chi_{\text{tot}}^2/\text{dof} \lesssim 0.1$ in all bins when using the full covariance matrix. However, the correlation functions with the different weights in this case are in much better agreement. We take the two systematic corrections to be consistent with one another if the probability to exceed the observed χ^2 between them is at least 0.01. Based on this definition, the corrections for redshift bins 1, 3, and 5 are consistent with each other, as opposed to only the first redshift bin when the masks were different. The second and fourth bins are inconsistent in both cases.

We use our new de-biased systematics-corrected correlation function and full $\mathbf{C}^{\text{Y1}} + \mathbf{C}^{\text{stat}} + \mathbf{C}^{\text{sys}}$ covariance matrix in combination with the cosmic shear and galaxy–galaxy lensing data vectors and covariance matrices from the DES Y1 cosmology analysis (Abbott et al. 2018) to re-run the DES $3 \times 2\text{pt}$ cosmology analysis. The resulting cosmology contours for Ω_m , A_s , and S_8 are shown in blue in Fig. 9. For our analysis, we use an updated version of `CosmoLike` (Krause & Eifler 2017; Fang, Eifler & Krause 2020) and use `emcee` (Foreman-Mackey et al. 2013) as our sampler. The result with this pipeline and our updated data vector and covariance matrix are shown in red in Fig. 9.

As we use a different pipeline and sampler than the fiducial Y1 analysis of Abbott et al. (2018), it is unclear how much of the difference between the red and blue contours in Fig. 9 is because of our changes to the data vector and covariance matrix and how much is a reflection of the differences in the modelling pipeline. We therefore show as the black-dashed lines in Fig. 9 the results of using the updated `CosmoLike` pipeline when run on the fiducial Y1 data vector and covariance matrix. The differences between the red and black contours are due to the difference in the estimated correlation function and its corresponding covariance matrix. It is clear that our weighting method does not have a significant impact on the cosmological inference relative to the Y1 analysis. This is expected given that both the difference in the correlation functions with the two different weighting methods and the uncertainty in our systematic correction are small relative to the statistical uncertainty of the measurement.

The black and red text above the histogram of S_8 in Fig. 9 show the minimum χ^2 values for the fiducial Y1 data vector and our updated data vector, respectively, for each data vector compared to the model with 444 degrees of freedom (see Abbott et al. 2018). The minimum χ^2 in each case is $-2\log L_{\text{max}}$ at the maximum-likelihood point in the MCMC chain. It is encouraging to see that even though our method does not significantly change the cosmological inference, it does result in a significant improvement in the goodness of fit ($\Delta\chi^2 = -6.5$ with no additional parameters). This improvement

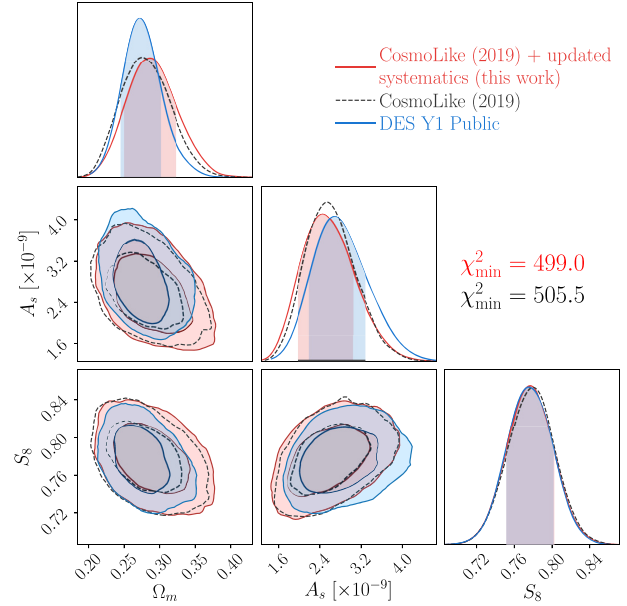


Figure 9. A comparison of the cosmology contours for the $3 \times 2\text{pt}$ analysis, with each 2D contour showing the 68 per cent and 95 per cent confidence levels, and the shaded regions in the 1D plots signifying the 68 per cent confidence level. The blue contours are the public DES Y1 results as in Abbott et al. (2018). The red contours are the results with our new correlation function and updated covariance matrix. Note that the blue and red contours use a different version of `CosmoLike` and different samplers. The black-dashed lines also show the contours using the DES Y1 data vector, but using the same version of `CosmoLike` and same sampler as was used to generate the red contours. The minimum χ^2 for the DES Y1 data vector and our updated data vector are shown as the black and red text, respectively, with 444 degrees of freedom.

in the χ^2 is due to both the increased error from our systematics correction and the shifts in the data vector that occur when replacing the Y1 weighting method with ours. To show that this is the case, we consider the calculation of the best fit χ^2 with our updated data vector and covariance matrix, which we now write as

$$\chi_{\text{new}}^2 = \left(\vec{d}_{\text{Y1}} + \vec{\Delta} - \vec{m}_{\text{Y1}} \right)^\top (\mathbf{C}^{\text{Y1}} + \delta\mathbf{C})^{-1} \left(\vec{d}_{\text{Y1}} + \vec{\Delta} - \vec{m}_{\text{Y1}} \right),$$

where \vec{d}_{Y1} is the original data vector from the Y1 analysis, \vec{m}_{Y1} is the best-fitting model vector from the original Y1 analysis, $\delta\mathbf{C} \equiv \mathbf{C}^{\text{stat}} + \mathbf{C}^{\text{sys}}$ is the change in the covariance matrix, and $\vec{\Delta}$ is the change to the difference between the data vector and best fit model vector introduced by our weights method. Note that this means that $\vec{\Delta}$ is sensitive to both the change in the data vector as well as changes to the best-fitting parameters. We can expand this equation around $\delta\mathbf{C} = 0$, dropping terms that are beyond first-order in $\delta\mathbf{C}$ as well as terms involving $\Delta^\top \delta\mathbf{C}$. Doing so, we find

$$\begin{aligned} \Delta\chi^2 \approx & \vec{\Delta}^\top (\mathbf{C}^{\text{Y1}})^{-1} \left[2 \left(\vec{d}_{\text{Y1}} - \vec{m}_{\text{Y1}} \right) + \vec{\Delta} \right] \\ & - \left[(\mathbf{C}^{\text{Y1}})^{-1} \left(\vec{d}_{\text{Y1}} - \vec{m}_{\text{Y1}} \right) \right]^\top \delta\mathbf{C} \left[(\mathbf{C}^{\text{Y1}})^{-1} \left(\vec{d}_{\text{Y1}} - \vec{m}_{\text{Y1}} \right) \right]. \end{aligned}$$

The first term in this expression gives the $\Delta\chi^2$ resulting from changing the data vector and the difference in the resulting best-fitting model vector. The second term is the $\Delta\chi^2$ caused by the change to the covariance matrix from our systematics correction. We find $\Delta\chi^2 \approx -3.6$ for the first term and $\Delta\chi^2 \approx -3.2$ for the second.

From this, we conclude that both the shift in the data vector (and resulting shift in the best fit) and the increased uncertainty due to our systematics correction contribute to the improvement in the fit.

7 CONCLUSIONS

We have presented a method for using a linear model to mitigate the effect of systematic fluctuations in galaxy clustering analyses due to observing conditions. Our method uses a Gaussian likelihood to fit the linear model to the observed galaxy overdensity and SP maps on each pixel. Our analysis explicitly incorporates the fact that neighboring pixels in the sky are correlated using an iterative approach: our first iteration uses a diagonal covariance matrix, while the second builds a non-diagonal covariance matrix from the systematics-corrected correlation function estimated from the first iteration. We further use mock catalogues to calibrate the remaining overcorrection bias, which we then remove from the data correlation function.

We apply our methodology to the DES Y1 redMaGiC data set. Our method has four important advantages relative to that adopted in the DES 3×2 pt analysis presented in Abbott et al. (2018), namely

(i) Our method does not require that decisions be made with regards to which survey properties matter and which do not. This also allows for the possibility of multiple survey properties ‘conspiring’ to create an observationally significant signal without any single systematic reaching that threshold.

(ii) Our method properly increases the error budget of the galaxy correlation function estimate by accounting for the statistical and systematic uncertainty associated with systematic mitigation. We have found doing so non-trivially impacts the goodness-of-fit statistic of the best-fitting cosmological model.

(iii) Our method explicitly incorporates clustering information from neighboring pixels in our calibration of the impact of survey properties on the galaxy density field.

(iv) Our method is fully automated: it can be run from start to finish with minimal supervision, enabling for quick turn around for future data sets, with no extra tuning.

While our updated systematics-corrected correlation function in the DES Y1 data set is formally inconsistent with that of Paper I, the two are in good agreement relative to the level of statistical uncertainty in DES Y1, and the corrections are consistent in redshift bins 1, 3, and 5 when applied over the same footprint. Because the statistical uncertainty in the measurement is larger than the uncertainty in our correction, we observe no significant impact on the cosmological inference using a data vector with our systematics weights relative to the Y1 3×2 pt cosmology analysis. Encouragingly, however, we do see an improvement in the goodness of fit, which is caused by both the change to the data vector with our new weights and the increased error from our systematics correction. We also expect the difference in the data vector and the uncertainty in the correction to become more important in the near future as the large number of galaxies observed by upcoming surveys decreases the statistical uncertainty in galaxy clustering measurements.

ACKNOWLEDGEMENTS

ELW and ER were supported by the DOE grant DE-SC0015975. XF is supported by NASA ROSES ATP 16-ATP16-0084 grant. Some calculations in this paper use High Performance Computing (HPC) resources supported by the University of Arizona TRIF, UITS, and RDI and maintained by the UA Research Technologies department.

This paper has gone through internal review by the DES collaboration.

Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft and the Collaborating Institutions in the DES.

The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l’Espai (IEEC/CSIC), the Institut de Física d’Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, NFS’s NOIRLab, the University of Nottingham, The Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, Texas A&M University, and the OzDES Membership Consortium.

Based in part on observations at Cerro Tololo Inter-American Observatory at NSF’s NOIRLab (NOIRLab Prop. ID 2012B-0001; PI: J. Frieman), which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

The DES data management system is supported by the National Science Foundation under Grants AST-1138766 and AST-1536171. The DES participants from Spanish institutions are partially supported by MICINN under grants ESP2017-89838, PGC2018-094773, PGC2018-102021, SEV-2016-0588, SEV-2016-0597, and MDM-2015-0509, some of which include ERDF funds from the European Union. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. Research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Program (FP7/2007-2013) including ERC grant agreements 240672, 291329, and 306478. We acknowledge support from the Brazilian Instituto Nacional de Ciência e Tecnologia (INCT) e-Universe (CNPq grant 465376/2014-2).

This manuscript has been authored by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Some of the results in this paper have been derived using the `healpy` and `HEALPIX` packages. This research has used `ASTROPY`,⁵

⁵<http://www.astropy.org>

a community-developed core PYTHON package for Astronomy (Astropy Collaboration 2013, 2018).

DATA AVAILABILITY

The DES Y1 redMaGiC catalogue is available for download at <https://des.ncsa.illinois.edu/releases/y1a1/key-catalogues/key-redmagic>. The DES observing condition maps, excluding dust and stellar density, are available for download at <https://des.ncsa.illinois.edu/releases/y1a1/gold/systematics>.

The dust map is available from *Planck* at https://irsa.ipac.caltech.edu/data/Planck/release_1/all-sky-maps/previews/HFI_CompMap_ThermalDustModel_2048_R1.20/index.html.

The stellar density catalogue is not publicly available but can be constructed from the DES Y1GOLD badmask (<https://des.ncsa.illinois.edu/releases/y1a1/gold/footprint>) and the first public data release (<https://des.ncsa.illinois.edu/releases/dr1>). It may also be made available upon request with permission of the DES Collaboration.

REFERENCES

- Abbott T. M. C. et al., 2018, *Phys. Rev. D*, 98, 043526
 Aihara H. et al., 2018, *PASJ*, 70, S4
 Astropy Collaboration, 2013, *A&A*, 558, A33
 Astropy Collaboration, 2018, *AJ*, 156, 123
 Bergé J., Gamper L., Réfrégier A. r., Amara A., 2013, *Astron. Comput.*, 1, 23
 Chiang Y.-K., Ménard B., 2019, *ApJ*, 870, 120
 Cuillandre J.-C. J. et al., 2012, in Peck A. B., Seaman R. L., Comeron F., eds, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 8448, Observatory Operations: Strategies, Processes, and Systems IV, p. 84480M
 Dark Energy Survey Collaboration, 2016, *MNRAS*, 460, 1270
 Delubac T. et al., 2017, *MNRAS*, 465, 1831
 Drinkwater M. J. et al., 2010, *MNRAS*, 401, 1429
 Drlica-Wagner A. et al., 2018, *ApJS*, 235, 33
 Elsner F., Leistedt B., Peiris H. V., 2016, *MNRAS*, 456, 2095
 Elvin-Poole J. et al., 2018, *Phys. Rev. D*, 98, 042006
 Fang X., Eifler T., Krause E., 2020, *MNRAS*, 497, 2699
 Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, 125, 306
 Górski K. M., Hivon E., Banday A. J., Wand elt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
 Hildebrandt H. et al., 2017, *MNRAS*, 465, 1454
 Ho S. et al., 2012, *ApJ*, 761, 14
 Howlett C., Lewis A., Hall A., Challinor A., 2012, *J. Cosmol. Astropart. Phys.*, 2012, 027
 Ivezić Ž. et al., 2019, *ApJ*, 873, 111
 Jarvis M., Bernstein G., Jain B., 2004, *MNRAS*, 352, 338
 Joudaki S. et al., 2018, *MNRAS*, 474, 4894
 Krause E., Eifler T., 2017, *MNRAS*, 470, 2100
 Krause E. et al., 2017, preprint ([arXiv:1706.09359](https://arxiv.org/abs/1706.09359))
 Landy S. D., Szalay A. S., 1993, *ApJ*, 412, 64
 Laureijs R. et al., 2011, preprint ([arXiv:1110.3193](https://arxiv.org/abs/1110.3193))
 Levi M. et al., 2013, preprint ([arXiv:1308.0847](https://arxiv.org/abs/1308.0847))
 Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
 Mandelbaum R., Seljak U., Kauffmann G., Hirata C. M., Brinkmann J., 2006, *MNRAS*, 368, 715
 Mandelbaum R., Slosar A., Baldauf T., Seljak U., Hirata C. M., Nakajima R., Reyes R., Smith R. E., 2013, *MNRAS*, 432, 1544
 Mellier Y., 1999, *ARA&A*, 37, 127
 More S., Miyatake H., Mandelbaum R., Takada M., Spergel D. N., Brownstein J. R., Schneider D. P., 2015, *ApJ*, 806, 2
 Nicola A. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 044
 Nishimichi T., D’Amico G., Ivanov M. M., Senatore L., Simonović M., Takada M., Zaldarriaga M., Zhang P., 2020, *Phys. Rev. D*, 102, 123541
 Planck Collaboration XI, 2014, *A&A*, 571, A11
 Prakash A. et al., 2016, *ApJS*, 224, 34
 Refregier A., 2003, *ARA&A*, 41, 645
 Rezaie M., Seo H.-J., Ross A. J., Bunescu R. C., 2020, *MNRAS*, 495, 1613
 Ross A. J. et al., 2011, *MNRAS*, 417, 1350
 Rybicki G. B., Press W. H., 1992, *ApJ*, 398, 169
 Sheldon E. S. et al., 2004, *AJ*, 127, 2544
 Smith R. E. et al., 2003, *MNRAS*, 341, 1311
 Spergel D. et al., 2015, preprint ([arXiv:1503.03757](https://arxiv.org/abs/1503.03757))
 Suchyta E. et al., 2016, *MNRAS*, 457, 786
 Takahashi R., Sato M., Nishimichi T., Taruya A., Oguri M., 2012, *ApJ*, 761, 152
 The Dark Energy Survey Collaboration, 2005, preprint ([astro-ph/0510346](https://arxiv.org/abs/astro-ph/0510346))
 van Uitert E. et al., 2018, *MNRAS*, 476, 4662
 Weaverdyck N., Huterer D., 2021, *MNRAS*
 Yoo J., Seljak U., 2012, *Phys. Rev. D*, 86, 083504
 York D. G. et al., 2000, *AJ*, 120, 1579
 Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, *J. Open Source Softw.*, 4, 1298
 Zuntz J. et al., 2015, *Astron. Comput.*, 12, 45

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