

Dark Energy Survey Year 6 results: Redshift calibration of the MagLim++ lens sample

G. Giannini^{1,2,*}, A. Alarcon,¹ W. d'Assignies,³ G. M. Bernstein,⁴ M. A. Troxel,⁵ C. Chang,^{6,2} B. Yin,⁵ A. Amon,⁷ J. Myles,⁷ N. Weaverdyck,^{8,9} A. Porredon,^{10,11} D. Anbajagane,² S. Avila,¹¹ K. Bechtol,¹² M. R. Becker,¹³ J. Blazek,¹⁴ M. Crocce,^{15,1} D. Gruen,¹⁶ M. Rodriguez-Monroy,^{17,18} C. Sánchez,^{3,19} I. Sevilla-Noarbe,¹¹ M. Aguena,^{20,21} S. Allam,²² O. Alves,²³ F. Andrade-Oliveira,²⁴ D. Bacon,²⁵ S. Bocquet,¹⁶ D. Brooks,²⁶ R. Camilleri,²⁷ A. Carnero Rosell,^{28,21,29} J. Carretero,³ R. Cawthon,³⁰ L. N. da Costa,²¹ M. E. da Silva Pereira,³¹ T. M. Davis,²⁷ J. De Vicente,¹¹ D. L. DePoy,³² S. Desai,³³ H. T. Diehl,²² S. Dodelson,^{22,6,2} P. Doel,²⁶ C. Doux,^{4,34} A. Drlica-Wagner,^{2,22,6} J. Elvin-Poole,³⁵ S. Everett,³⁶ A. E. Evrard,^{37,23} B. Flaugher,²² J. Frieman,^{6,22,2} J. García-Bellido,³⁸ M. Gatti,^{2,1} E. Gaztanaga,^{25,1,15} P. Giles,³⁹ R. A. Gruendl,^{40,41} G. Gutierrez,²² K. Herner,²² S. R. Hinton,²⁷ D. L. Hollowood,⁴² K. Honscheid,^{43,44} D. Huterer,²³ D. J. James,⁴⁵ K. Kuehn,^{46,47} O. Lahav,²⁶ S. Lee,^{48,49} H. Lin,²² J. L. Marshall,³² J. Mena-Fernández,^{34,50} F. Menanteau,^{41,40} R. Miquel,^{3,51} J. Muir,^{52,53} R. L. C. Ogando,⁵⁴ D. Petravick,⁴¹ A. A. Plazas Malagón,^{55,56} J. Prat,^{57,58} M. Raveri,⁵⁹ E. S. Rykoff,^{56,55} S. Samuroff,^{3,14} E. Sanchez,¹¹ D. Sanchez Cid,^{24,11} T. Shin,⁶⁰ M. Smith,⁶¹ E. Suchyta,⁶² M. E. C. Swanson,⁴¹ G. Tarle,²³ D. Thomas,²⁵ C. To,⁶ D. L. Tucker,²² V. Vikram,⁶³ and M. Yamamoto^{5,7}

(DES Collaboration)

¹*Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, 08193 Barcelona, Spain*

²*Kavli Institute for Cosmological Physics, University of Chicago, Chicago, Illinois 60637, USA*

³*Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, 08193 Bellaterra (Barcelona) Spain*

⁴*Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*

⁵*Department of Physics, Duke University, Durham, North Carolina 27708, USA*

⁶*Department of Astronomy and Astrophysics, University of Chicago, Chicago, Illinois 60637, USA*

⁷*Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, New Jersey 08544, USA*

⁸*Berkeley Center for Cosmological Physics, Department of Physics, University of California, Berkeley, California 94720, USA*

⁹*Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA*

¹⁰*Ruhr University Bochum, Faculty of Physics and Astronomy, Astronomical Institute, German Centre for Cosmological Lensing, 44780 Bochum, Germany*

¹¹*Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain*

¹²*Physics Department, 2320 Chamberlin Hall, University of Wisconsin-Madison, 1150 University Avenue Madison, Wisconsin 53706-1390, USA*

¹³*Argonne National Laboratory, 9700 South Cass Avenue, Lemont, Illinois 60439, USA*

¹⁴*Department of Physics, Northeastern University, Boston, Massachusetts 02115, USA*

¹⁵*Institut d'Estudis Espacials de Catalunya (IEEC), 08034 Barcelona, Spain*

¹⁶*University Observatory, LMU Faculty of Physics, Scheinerstraße 1, 81679 Munich, Germany*

¹⁷*Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, 28049 Madrid, Spain*

¹⁸*Laboratoire de physique des 2 infinis Irène Joliot-Curie, CNRS Université Paris-Saclay, Bât. 100, F-91405 Orsay Cedex, France*

¹⁹*Departament de Física, Universitat Autònoma de Barcelona (UAB), 08193 Bellaterra, Barcelona, Spain*

²⁰*INAF-Osservatorio Astronomico di Trieste, via G. B. Tiepolo 11, I-34143 Trieste, Italy*

²¹*Laboratório Interinstitucional de e-Astronomia—LIneA, Avenida Pastor Martin Luther King Jr, 126 Del Castilho, Nova América Offices, Torre 3000/sala 817 CEP: 20765-000, Brazil*

²²*Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, Illinois 60510, USA*

²³*Department of Physics, University of Michigan, Ann Arbor, Michigan 48109, USA*

²⁴*Physik-Institut, University of Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland*

²⁵*Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth, PO1 3FX, United Kingdom*

²⁶*Department of Physics and Astronomy, University College London, Gower Street, London, WC1E 6BT, United Kingdom*

²⁷*School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia*

²⁸*Universidad de La Laguna, Departamento de Astrofísica, E-38206 La Laguna, Tenerife, Spain*

²⁹*Instituto de Astrofísica de Canarias, E-38205 La Laguna, Tenerife, Spain*

³⁰*Oxford College of Emory University, Oxford, Georgia 30054, USA*

³¹*Hamburger Sternwarte, Universität Hamburg, Gojenbergsweg 112, 21029 Hamburg, Germany*

- ³²*George P. and Cynthia Woods Mitchell Institute for Fundamental Physics and Astronomy, and Department of Physics and Astronomy, Texas A&M University, College Station, Texas 77843, USA*
- ³³*Department of Physics, IIT Hyderabad, Kandi, Telangana 502285, India*
- ³⁴*Université Grenoble Alpes, CNRS, LPSC-IN2P3, 38000 Grenoble, France*
- ³⁵*Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, Ontario N2L 3G1, Canada*
- ³⁶*California Institute of Technology, 1200 East California Boulevard, MC 249-17, Pasadena, California 91125, USA*
- ³⁷*Department of Astronomy, University of Michigan, Ann Arbor, Michigan 48109, USA*
- ³⁸*Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, 28049 Madrid, Spain*
- ³⁹*Department of Physics and Astronomy, Pevensey Building, University of Sussex, Brighton, BN1 9QH, United Kingdom*
- ⁴⁰*Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana, Illinois 61801, USA*
- ⁴¹*Center for Astrophysical Surveys, National Center for Supercomputing Applications, 1205 West Clark Street, Urbana, Illinois 61801, USA*
- ⁴²*Santa Cruz Institute for Particle Physics, Santa Cruz, California 95064, USA*
- ⁴³*Department of Physics, The Ohio State University, Columbus, Ohio 43210, USA*
- ⁴⁴*Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, Ohio 43210, USA*
- ⁴⁵*Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, Massachusetts 02138, USA*
- ⁴⁶*Australian Astronomical Optics, Macquarie University, North Ryde, NSW 2113, Australia*
- ⁴⁷*Lowell Observatory, 1400 Mars Hill Road, Flagstaff, Arizona 86001, USA*
- ⁴⁸*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr., Pasadena, California 91109, USA*
- ⁴⁹*Department of Physics and Astronomy, Ohio University, Clipping Labs, Athens, Ohio 45701*
- ⁵⁰*Aix Marseille University, CNRS/IN2P3, CPPM, Marseille, France*
- ⁵¹*Institució Catalana de Recerca i Estudis Avançats, E-08010 Barcelona, Spain*
- ⁵²*Perimeter Institute for Theoretical Physics, 31 Caroline St. North, Waterloo, Ontario N2L 2Y5, Canada*
- ⁵³*Department of Physics, University of Cincinnati, Cincinnati, Ohio 45221, USA*
- ⁵⁴*Observatório Nacional, Rua Gal. José Cristino 77, Rio de Janeiro, Rio de Janeiro—20921-400, Brazil*
- ⁵⁵*Kavli Institute for Particle Astrophysics and Cosmology, P.O. Box 2450, Stanford University, Stanford, California 94305, USA*
- ⁵⁶*SLAC National Accelerator Laboratory, Menlo Park, California 94025, USA*
- ⁵⁷*Nordita, KTH Royal Institute of Technology and Stockholm University, Hannes Alfvéns väg 12, SE-10691 Stockholm, Sweden*
- ⁵⁸*University of Copenhagen, Dark Cosmology Centre, Juliane Maries Vej 30, 2100 Copenhagen O, Denmark*
- ⁵⁹*Department of Physics, University of Genova and INFN, Via Dodecaneso 33, 16146, Genova, Italy*
- ⁶⁰*Department of Physics and Astronomy, Stony Brook University, Stony Brook, New York 11794, USA*
- ⁶¹*Physics Department, Lancaster University, Lancaster, LA1 4YB, United Kingdom*
- ⁶²*Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA*
- ⁶³*Central University of Kerala, Periyar, Kerala 671320, India*



(Received 25 March 2026; accepted 26 May 2026; published 7 July 2026)

In this work, we derive and calibrate the redshift distribution of the MagLim++ lens galaxy sample used in the Dark Energy Survey Year 6 (DES Y6) 3×2 pt cosmology analysis. The 3×2 pt analysis combines galaxy clustering from the lens galaxy sample and weak gravitational lensing. The redshift distributions are inferred using the SOMPZ method—a self-organizing map framework that combines deep-field multiband photometry, wide-field data, and a synthetic source injection (Balrog) catalog. Key improvements over the DES Year 3 (Y3) calibration include a noise-weighted SOM metric, an expanded Balrog catalog, and an improved scheme for propagating systematic uncertainties, which allows us to generate $O(10^8)$ redshift realizations that collectively span the dominant sources of uncertainty. These realizations are then combined with independent clustering-redshift measurements via importance sampling. The resulting calibration achieves typical uncertainties on the mean redshift of 1%–2%, corresponding to a 20%–30% average reduction relative to DES Y3. We compress the $n(z)$ uncertainties into a small number of orthogonal modes for use in cosmological inference. Marginalizing over these modes leads to only a minor degradation in cosmological constraints. This analysis establishes the MagLim++ sample as a robust lens sample for precision cosmology with DES Y6 and provides a scalable framework for future surveys.

DOI: [10.1103/mkwr-bfls](https://doi.org/10.1103/mkwr-bfls)

*Contact author: giulia.giannini15@gmail.com

I. INTRODUCTION

The success of modern cosmology builds upon a combination of probes that measure the geometry and structure growth of the Universe [1]. For structure growth, the most constraining and robust structure measurements in modern galaxy surveys rely on the two-point auto- and cross-correlation between galaxy density and galaxy weak lensing. This combination, colloquially referred to as the “ 3×2 pt” probes, consistently combines cosmic shear (the autocorrelations of galaxy weak lensing), galaxy-galaxy lensing (the cross-correlation between galaxy density and galaxy weak lensing), and galaxy clustering (the autocorrelation of galaxy density). The three probes have different degeneracy directions in the model and share a number of nuisance parameters; this unique feature allows it to self-consistently calibrate the nuisance parameters and result in tighter and more robust constraints. Many modern galaxy surveys have demonstrated the effectiveness of the 3×2 pt probes [2–4].

One of the key ingredients of the 3×2 pt probes is the distance, or redshift, to the galaxies used for the measurements. In a 3×2 pt analysis we define two galaxy samples according to their role: a *lens* sample, used as a tracer of the galaxy density field, and a *source* sample, used as a tracer of the weak lensing shear field. While these samples are often constructed to be disjoint, in principle any galaxy sample can serve as either, and some analyses use the same galaxies in both roles [5]. Due to the different nature of the samples, the treatment in how we derive the distances to these samples often differs. This paper focuses on deriving and characterizing the redshift distribution for the lens galaxy sample, the MagLim++ sample, which is defined in [6] and will be used in the dark energy survey (DES) Year 6 (Y6) 3×2 pt cosmology analysis [7].

The lens galaxy sample, used in both the galaxy clustering and the galaxy-galaxy lensing measurements, is designed to be a relatively bright galaxy sample that traces the dark matter density in a way that is easy to model. Since lens galaxy clustering probes the matter density field through a redshift kernel set by the width of the lens sample’s $n(z)$ (as opposed to weak lensing, which is sensitive to a much broader kernel), a desired property for the sample is to have narrow redshift distributions [8] so as to maximize the information in the radial direction. In addition, as shown in [8,9], for current datasets the precision with which we can calibrate the redshift distribution of the lens sample can have a greater impact on cosmological constraints than simply increasing its number density. In other words, improving $n(z)$ calibration often yields larger gains than adding more galaxies, since inaccurate redshift estimates can bias the results even in a high-density sample.

As the lens samples are typically bright galaxies, historically not as much attention has been placed on understanding their redshift distributions compared to the source sample. In the DES Y3 analysis, a large amount of

work went into developing a new methodology for redshift estimation of the source sample, combining photometric information [self-organizing map photometric redshift (SOMPZ) [10]], clustering information (wz, [11]), and geometric information (shear ratio, [12]). The SOMPZ framework is able to characterize the redshift distribution and uncertainty of the source sample in a fully Bayesian way, utilizing three sets of photometry from the DES wide field, the DES deep field and the spectroscopic sample (see Sec. II); the wz pipeline uses independent clustering information to further constrain the redshift distribution; finally the shear ratio technique provides an additional independent check on the high-redshift galaxies using weak lensing itself.

Unlike the source sample, most previous work using a lens sample for either galaxy-galaxy lensing or galaxy-clustering opts to use spectroscopic samples if there is good overlap with lensing data [13–15], since precise redshifts are especially beneficial. Although DES overlaps with several spectroscopic surveys, the coverage at this point is not sufficient to define a spectroscopic lens sample large enough for competitive cosmological analyses. We therefore construct photometric lens samples from bright galaxies with reliable photometric redshifts. In the fiducial DES Y3 analysis, we looked at a redshift-dependent magnitude-limited sample MagLim [16]. The redshift distribution of that sample was first determined using directional neighborhood fitting (DNF) [17] and then calibrated using wz [18].

In a follow-up paper using a similar SOMPZ+WZ technology developed for the source sample [19], we found that adopting the more principled redshift calibration method resulted in $\sim 0.4\sigma$ shift of the cosmological constraints combining galaxy clustering and galaxy-galaxy lensing in the $S_8 - \Omega_m^{-1}$ plane, highlighting the importance of Bayesian redshift calibration.

The goal of this paper is to present the redshift distribution and its uncertainty for the DES Y6 MagLim++ lens sample: this will be the lens sample used in the legacy 3×2 pt analysis from DES. We describe the methodology as well as data updates to [19], and assess their impact on cosmological constraints. The main updates include: (1) improvements in data products throughout the pipeline, including a cleaner lens sample (MagLim++), a larger synthetic source injection catalog (Balrog, [20]) that covers the entire footprint, and an enhanced redshift sample; (2) the Y6 imaging is about twice as deep as Y3, which reduces magnitude errors, although for the relatively bright MagLim++ lenses ($i < 22.2$) this effect is modest, it still provides more precise color measurements; (3) an updated

¹ $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$, where σ_8 is the normalization of the present-time, linear matter-power spectrum smoothed on $8h^{-1}$ Mpc scales, and Ω_m is the ratio of the present-time matter energy density to the critical energy density.

TABLE I. Reference table for the series of four papers that describe the full redshift treatment for the DES Y6 3×2 pt analysis.

Shorthand	Reference	Description
Y6ShearPZ	[65]	Source photo- z calibration
Y6PosPZ	<i>This work</i>	Lens photo- z calibration
Y6WZ	[66]	Clustering redshift
Y6Mode	[67]	Redshift marginalization method

SOMPZ algorithm that incorporates the photometric uncertainty when assigning galaxies into the self-organizing map (SOM) [21,22]; (4) a more general prescription in marginalizing the uncertainties in the redshift sample; and (5) a more general parametrization for the uncertainties in the redshift distribution.

This paper is one of the series of four papers that will describe the full redshift treatment for the DES Y6 3×2 pt analysis, summarised in Table I: (1) [Y6ShearPZ] describes the derivation of the SOMPZ redshift estimates for the source galaxy sample; (2) this paper describes the derivation of the SOMPZ redshift estimates for the lens galaxy sample; (3) [Y6WZ] describes how the clustering information (WZ) is incorporated into the redshift estimates of (1)

and (2); and (4) [Y6Modes] describes a novel approach used in the Y6 analysis to sample over the redshift uncertainty during the cosmological inference. Altogether, the four papers give the complete description of the Y6 redshift estimates. Figure 1 illustrates the overall redshift workflow for the DES Y6 3×2 pt analysis. In this work, MagLim++ serves as the target sample in the flowchart.

The paper is structured as follows. In Sec. II we describe the data products used in this work, including both DES data as well as external redshift catalogs. In Sec. III we describe the methodology we have adopted to estimate and calibrate the redshift distribution of the MagLim++ sample based on the SOMPZ method. Section IV describes how the different contributions to the uncertainties on the redshift distributions are estimated. In Sec. V we describe how the SOMPZ results are combined with the clustering redshift information presented in [Y6WZ], and how we marginalize over the uncertainties following [Y6Modes]. We discuss some of the major improvements of this work compared to DES Y3 and the impact on cosmological constraints in Sec. VI. We conclude in Sec. VII.

II. DATA

This section outlines the datasets used in our analysis, including the photometric galaxy sample, redshift




FIG. 1. Flowchart summarizing the DES Y6 3×2 pt redshift calibration pipeline. The target sample of this study is MagLim++ (Sec. II), with redshift information provided by the redshift sample (Sec. II B) and deep fields combined with Balrog simulations (Sec. II). Photometric data are then mapped using self-organizing maps (SOMs) for deep and wide fields (Sec. III), transferring the redshift information to the target galaxies. Sampling over the associated uncertainties (Sec. IV) produces realizations $n_i^{\text{SOMPZ}}(z)$. Clustering information (from [Y6WZ]) provides additional constraints via importance sampling (Sec. V), yielding $n_i^{\text{SOMPZ+WZ}}(z)$. These realizations are then projected onto modes (Sec. V C; see also [Y6Mode]), enabling accurate marginalization over the redshift distribution in cosmological inference. This paper details the SOMPZ derivation, its combination with WZ, the mode-projection framework, and robustness tests through cosmological inference, while the WZ measurement itself is presented in [Y6WZ].

calibration inputs, and external datasets employed for validation and cross-correlation.

A. DES data products

The DES is a photometric galaxy survey that took data with the Dark Energy Camera (DECam, [23]) on the 4-m Blanco Telescope at the Cerro Tololo Inter-American Observatory in Chile during the period of August 2013 to January 2019. This paper uses the full six years of data. We describe the main ingredients used in this paper below.

- (i) DES Y6 `Gold` catalog: Described in [24], this catalog is the foundation for most of the Y6 data products. It includes the final coadded *grizY* photometry, object classifications, masks, and quality flags. Of the different flux estimates provided in this catalog, we use the `Fitvd` fits to each object. The full catalog contains 448 million galaxies and 120 million stars.
- (ii) DES Y3* deep field catalog: We use an updated version (hence the *) of the catalog originally presented in [25]. In particular, we start from the same set of detected objects as [25] and then remeasure the eight-band photometry for each object using an updated fitting algorithm. We note that the update is only in the fitting algorithm and *not* in the images. The motivation for developing this updated pipeline is to enable its application to both the deep- and wide-field `Gold` fluxes. We apply a number of quality cuts to this deep-field sample, following [25], to remove objects in regions of poor image quality or regions without eight-band coverage. Finally, following previous redshift calibration work in DES, we discard any deep-field galaxies that are never detected in our wide field; this is determined using the synthetic source catalog described below. Figure 2 depicts the four deep fields and their overlap with redshift samples (see next Sec. II B).
- (iii) DES Y6 synthetic source catalog (Balrog): In this work, we use the Balrog catalog derived in [20] to connect the deep-field photometry to the wide-field photometry, which we refer to as the “transfer function” in this work. Galaxies from the deep-field catalog are injected into the real charge-coupled device images, with each galaxy being injected multiple times, and the modified images are processed through the entire DES pipeline. As a result, for each deep-field object we can compute the distribution of its measured, wide-field properties, i.e., the transfer function. Compared to the dataset in DES Y3, the Y6 dataset spans the entire survey footprint (described in [26]) and is further optimized for improving the precision on estimates of the transfer function. The underlying injection and processing pipeline remain largely unchanged from Y3, aside from minor updates such as the use of PSFs in the full

field of view models [27] and an improved weighting scheme, which are described in Section 3.4 of [20].

- (iv) The MagLim++ lens catalog: The MagLim++ sample in DES Y6 is described in [6], and consists of the same basic magnitude limited selection from DES Y3 [16], with additional refinements to increase sample purity and mitigate systematic contamination. We only summarize the main characteristics of the sample here, and refer the reader to [6] for more details.

As in DES Y3, the primary selection is performed on galaxies in the `Gold` catalog using the following *i*-band limiting magnitude threshold that depends on the predicted mean photometric redshift of each galaxy, estimated from the DNF algorithm [17],

$$i < 4 \times z_{\text{MEAN}} + 18, \quad i > 17.5. \quad (1)$$

This cut is designed to create a magnitude-limited sample with a smooth redshift-dependent selection function, balancing between density and photometric redshift accuracy. The selection has been optimized in [9], with the goal of improving the cosmological constraints, as quantified by the figure of merit on Ω_m , σ_8 , and the dark energy equation of state, w . We then split the sample into six tomographic bins from $z_{\text{MEAN}} = 0.2$ to $z_{\text{MEAN}} = 1.05$, with bin edges $[0.20, 0.40, 0.55, 0.70, 0.85, 0.95, 1.05]$. For DES Y6, two additional selection criteria were introduced to further refine the sample. First, the redshift-bin optimized star-galaxy classifier of [6] is applied to remove residual stellar contamination. This leverages near-infrared (NIR) data from `unWISE` [28] to identify the optimal color cuts in $(r - z, z - W1)$ to separate stars and galaxies in each redshift bin, complementing the morphological `EXTMASH` star-galaxy classifier in the `Gold` sample, and removing $\sim 0.6\%$ – 3.3% of objects from different bins.

Second, [6] apply a novel nonparametric selection based on SOM. Briefly, the SOM serves as an unsupervised clustering and dimensionality reduction algorithm that maps the four-dimensional color space onto a 2D grid, see Sec. III for details. It is trained on the four-band *griz* photometry of the sample to remove objects that occupy compact regions of color space with large photo-*z* dispersion. Such dispersion comes from low precision in photo-*z* estimation in this region as a result of degeneracies in the color-redshift relation or limitations in the training data. In either case, redshifts of galaxies in such regions are more likely to be mismodeled and thus function as contaminants to the sample.

Three complementary measures are used to identify compact regions in the SOM with high levels of photo-*z* dispersion at either the galaxy or cell level,

and objects in these cells are removed from the selection (amounting to 2.4% of the total). Cross-matching these removed objects with the VISTA survey reveals that they are primarily quasistellar objects, (QSOs) [29].

These two improvements in MagLim++ result in a cleaner lens sample in three ways: (1) reducing stellar contamination using a combination of morphological and redshift-bin optimized NIR color-based criteria, (2) removing contamination from redshift-interloper regions through precise color cuts of regions with unreliable photometric redshifts, and (3) sharpening the clustering signal by removing objects more likely to populate the tails of the $n(z)$.

These refinements lead to a more homogeneous sample with reduced systematic uncertainties, enhancing its suitability for cosmological analyses.

B. The redshift catalogs

Our analysis relies on the use of galaxy samples with known redshift and deep-field photometry. To this end, we use catalogs of both high-resolution spectroscopic and multiband photometric redshifts, and combine them into a redshift calibration sample, prioritizing information coming from spectroscopic surveys. In addition, we develop a model to propagate the uncertainty in our redshift calibration arising from residual biases and uncertainty in the original redshift catalogs (quantified in Sec. IV C).

Figure 2 shows all the DES deep fields and the distribution of different photometric and spectroscopic datasets that provide redshift information for the MagLim++ sample. Primarily, we use redshift information coming from the COSMOS and X3 DES deep fields. The spectroscopic catalog we use contains both public and private spectra from multiple surveys compiled by [30], and references therein. A majority of these spectra come from: zCOSMOS [31], C3R2 [32–34], VVDS [35], VIPERS [36], and DESI EDR [37]. In addition to spectroscopic surveys, we use three multiband photo- z catalogs: (1) the COSMOS2020 photometric redshift catalog [38], which includes 30 broad, intermediate, and narrow bands covering the UV, optical, and IR regions of the electromagnetic spectrum in the COSMOS field [39]; (2) the PAUS + COSMOS 66-band photometric redshift catalog [40], also in the COSMOS field, from the combination of PAU Survey data in 40 narrow-band filters [41,42] and 26 COSMOS2015 bands excluding the midinfrared [43]; and finally (3) the PAUSW1 [44] catalog, which is part of the ~ 50 deg² complete PAUS survey that overlaps with the CFHTLS-Wide W1 field and the DES X3 deep field.

Figure 3 shows the i -band magnitude distribution of the galaxies in the redshift sample, split by tomographic bin (panels) and color coded by the dataset providing the redshift. Each galaxy is weighted by its Balrog-based detection and selection probability into the corresponding




FIG. 2. Four DES deep fields used for our redshift analysis, which includes overlapping deep DES *ugriz* bands and VIDEO or UltraVISTA JHK bands, as compiled from [25]. Black points indicate DES deep-field galaxies with no redshift information, while points with colors show galaxies with spectroscopy or from COSMOS2020, PAUS + COSMOS or PAUSW1, as indicated in the legend.

MagLim++ tomographic bin. This weighting ensures that the redshift sample reflects the same selection as the lensing sample, thus enabling unbiased calibration. The figure also highlights the diversity of redshift sources




FIG. 3. Distribution of redshift samples as a function of the deep field DES i -band magnitude. Each galaxy in these stacked histograms is weighted by the Balrog probability of detection and selection into each Maglim sample (different rows). For details on the definition of the “bulge plus disk, fixed ratio” (BDF) galaxy profile see Ref. [25].

used, including both spectroscopic surveys (e.g., C3R2, VIPERS, ZCOSMOS, DESI) and high-quality photometric catalogs (e.g., PAUS, COSMOS2020). Notably, the transition in redshift source composition with increasing bin number and magnitude indicates the varying depth and availability of redshift information across the sample.

An important consideration when building a redshift calibration sample is its completeness and redshift quality. An incomplete redshift sample will result in a biased $n(z)$ inference due to selection effects (e.g. Refs. [45–47]) that can be particularly challenging to forward model. Only galaxies with reliable redshifts—whether spectroscopic or from COSMOS2020 photo- z catalogs—are included. We emphasize that we restrict our calibration to the regions where these samples are complete across the full magnitude range of interest, minimizing the risk of selection biases that would otherwise impact the inferred $n(z)$. For this reason, to construct the redshift sample we only use areas in the deep fields with complete redshift coverage for all galaxies that could be observed as MagLim++ galaxies. This results in excluding fields such as C3 and E2, which contain very little redshift information, as visible in Fig. 2. In contrast, both the COSMOS and X3 fields have a large spatial coverage. In particular, we use COSMOS-field galaxies to inform all tomographic bins, while we only use X3-field galaxies to inform the first two tomographic bins. This is motivated by the quality of the redshift information available in each field; while COSMOS has high quality photometric redshifts for all galaxies down to $i \sim 25$, in the X3 field this is only true down to $i \sim 20.5$. For reference, the maximum (observed) magnitude for each tomographic bin is $i_{\text{max}} = [19.6, 20.2, 20.8, 21.4, 21.8, 22.2]$ [Eq. (1)], which motivates only using the X3 field for the first two tomographic bins.

Figure 4 shows the redshift performance of the different photo- z catalogs used in the COSMOS and X3 fields as a function of deep DES BDF i -band magnitude. To evaluate this performance, we compute $(z_{\text{phot}} - z_{\text{spec}})/(1 + z_{\text{spec}})$ for each object and derive the median, the scatter (denoted as σ_{68}), and the outlier rate of the resulting distributions. The top panel shows the median redshift bias, the middle panel displays the scatter σ_{68} , and the bottom panel reports the outlier rate—defined as the fraction of galaxies for which $|z_{\text{phot}} - z_{\text{spec}}| > 0.1(1 + z_{\text{spec}})$. When we have redshift information from more than one source, we prioritize the most reliable one, as detailed in Fig. 4. In the COSMOS field, we first use spectroscopic redshifts, then PAUS + COSMOS, and finally COSMOS2020. Within the COSMOS2020 release there are two types of photometric catalogs (Classic and Farmer) each providing two redshift estimates, from Lephare and EAZY.

Based on our assessment of the redshift quality in Fig. 4, we use COSMOS2020 Classic-EAZY for galaxies with $i < 22.5$, and Farmer-LePhare for galaxies with $i > 22.5$. The Classic-EAZY combination shows the lowest median bias and smallest scatter at bright magnitudes, while Farmer-




FIG. 4. Redshift quality of different redshift catalogs used for the redshift calibration of MagLim++ samples, relative to spec- z measurements for the same objects, using $\Delta_z/(1 + z_{\text{spec}})$. We define outliers as having $|\Delta_z|/(1 + z_{\text{spec}}) > 0.1$. The redshift catalogs shown are PAUS + COSMOS; PAUS + W1; and then Classic-EAZY and Farmer-Lephare from the COSMOS2020 release (see text). Relative to each other, Classic-EAZY performs better at $i < 22.5$, while Farmer-Lephare performs better at $i > 22.5$, in terms of median bias and photo- z error (top and middle panels). The PAUS-W1 is used for $i < 20.5$, where it shows excellent photo- z precision and low bias.

LePhare performs better at the faint end, at least up to $i < 24$, which is the relevant deep magnitude range for MagLim++. In the X3 field, we prioritize spectroscopic redshifts, followed by PAUS-W1. PAUS-W1 galaxies are only used for $i < 20.5$, where their photometric redshifts show low bias and excellent precision. As a result, the X3 field contributes to the first two MagLim++ tomographic bins alongside COSMOS. For the remaining four MagLim++ bins, we rely solely on the COSMOS field, as X3 lacks reliable redshift coverage at fainter magnitudes (see Fig. 3).

III. SOM-BASED REDSHIFT INFERENCE METHODOLOGY

Our calibration process to determine the redshift distributions of the MagLim++ galaxies in DES Y6 incorporates both SOM-based photometric redshift estimates and clustering redshift constraints. This section provides a comprehensive account of the SOMPZ method, refining the

approach explored by [48], and previously employed in Y3 for the source galaxies [49,50], and later applied to the lens galaxies [19]. The SOMPZ effort for the MagLim++ sample is fully documented in this work.

A. Self-organizing maps for photometric redshift calibration

The SOM is an unsupervised machine learning algorithm that projects high-dimensional data onto a two-dimensional grid while maintaining the relative proximity of similar data points. Its key strength is in capturing the structure of complex datasets by organizing inputs with similar properties into neighboring cells.

In this work, the input to the SOM consists of galaxy flux measurements across multiple photometric bands. Each cell on the map represents a set of galaxies with similar observed fluxes. Owing to redshift degeneracies, a given cell does not correspond to a single redshift value but instead maps to a distribution of redshifts. Beyond clustering galaxies with similar fluxes, the SOM provides a topological representation of the data, preserving relationships between neighboring cells. This structure allows us to visualize transitions between different galaxy populations and identify distinct features such as Lyman breaks or contamination from stars. These properties help refine data selection and assess potential systematics.

The SOM process consists of two stages:

- (1) Training: The SOM is initialized with a grid of cells, each with a representative set of flux values. During training, galaxies are sequentially mapped to the closest-matching cell based on a predefined metric. The winning cell and its neighboring cells are then updated to better match the input flux values, with updates decreasing over time to stabilize the mapping. This results in a structured representation where neighboring cells contain similar galaxy populations.
- (2) Assignment: Once trained, the SOM can be used to classify new galaxies by assigning them to the closest-matching cell, effectively categorizing galaxies into different phenotypes based on their photometric properties.

To determine which SOM cell most closely matches a galaxy, DES Y6 adopts the revised distance metric of [21], applied to data in [22]. Unlike the Y3 formulation, which compared raw color-magnitude vectors at equal weight, the Y6 metric works directly in flux space and explicitly accounts for photometric uncertainties: bands with low signal-to-noise ratio (S/N) are down weighted, so noisy measurements cannot dominate the match. This improvement is crucial for faint sources and still beneficial, though subtler, for the comparatively bright MagLim++ sample. This construction boosts the influence of high-quality bands, suppresses noisy measurements, and remains insensitive to absolute flux differences.

Besides the new metric, the Y6 SOM employs non-periodic map boundaries, better reflecting real flux behavior, and adds the g band to the riz set used in Y3, further sharpening the color space mapping [22]. Consequently, the DES Y6 SOM yields a smoother and more physically meaningful organization of galaxies in color space, improving the stability of photometric-redshift calibration.

1. Redshift calibration via deep and wide SOMs

The SOMPZ methodology was first developed by DES in [49]. This technique utilizes SOMs to classify galaxies into distinct phenotypes based on their observed properties (e.g., flux, colors). The redshift calibration of each phenotype is achieved by incorporating galaxies with high-quality redshift measurements into the same SOM. This approach employs two separate SOMs: one representing the color space of the sample needing calibration (*wide SOM*) and another constructed from the DES deep fields (*deep SOM*). A direct calibration of MagLim++ using a single SOM would be inadequate, as broadband DECam photometry lacks the precision needed to resolve certain redshift-type degeneracies. The overlap of the DES deep fields with spectroscopic and narrow-band photometric surveys provides additional wavelength coverage, including near-ultraviolet and near-infrared bands (u, J, H, K_s), which helps mitigate these degeneracies. More details are provided in Sec. II.

To bridge the deep SOM and the wide SOM, an additional intermediary is needed to transfer redshift information across different observational conditions. This transfer function is constructed using Balrog simulations described in Sec. II and implemented through a Bayesian framework, ensuring consistency between the deep and wide observations.

The two-tiered SOM redshift calibration method was applied to observational data for the first time using the Y3 dataset [19,50]. The calibration pipeline follows these steps:

- (1) Train a deep SOM using deep-field photometry ($ugrizJHK_s$ bands) and a wide SOM with wide-field photometry ($griz$ bands).
- (2) Assign galaxies from the deep sample to the deep SOM and those from the wide sample to the wide SOM, yielding probability distributions $p(c)$ and $p(\hat{c})$ of a galaxy being assigned to a specific deep cell c or wide cell $\hat{c}$, respectively.
- (3) Cross-match galaxies with known redshifts from spectroscopic and high-quality photometric samples (e.g., COSMOS, PAUS) to deep SOM cells, deriving the redshift distribution of each deep cell $p(z|c)$.
- (4) Assign Balrog galaxies to the wide SOM, connecting the deep and wide photometric spaces through $p(c|\hat{c})$.

Combining this information, the final redshift distribution for each tomographic bin is computed as,




FIG. 5. Visualization of the six deep SOMs, one for each tomographic bin (ordered from left to right and top to bottom, i.e., 1–2 on the first row, 3–4 on the second, 5–6 on the third), each composed of 144 cells (12×12). Left: average redshift for each deep SOM cell c . Right: standard deviation on the redshift distribution for each deep SOM cell c . The black cells in the deep SOM are due to the lack of spectroscopic information in those regions of the color space, i.e., there are no galaxies in the redshift samples that were assigned to those cells.

$$p(z) = \sum_{\hat{c}} \sum_c p(z|c) p(c) p(c|\hat{c}) p(\hat{c}), \quad (2)$$

where we can rewrite this equation in terms of the data products used,

$$p(z) = \sum_{\hat{c}} \sum_c \underbrace{p(z|c)}_{\text{Redshift}} \underbrace{p(c)}_{\text{Deep}} \underbrace{\frac{p(c, \hat{c})}{p(c)p(\hat{c})}}_{\text{Balrog}} \underbrace{p(\hat{c})}_{\text{Wide}}. \quad (3)$$

This formulation effectively reweights the deep-field redshift sample to match the selection function of the wide-field survey, ensuring a representative redshift distribution for weak lensing analyses.

The choice of SOM resolution follows the approach adopted in DES Y3: since the MagLim++ sample did not change significantly in color or depth, we apply the same configuration as in Y3—using a 12×12 grid for each deep SOM and a 32×32 grid for each wide SOM. Because the SOM is trained and populated using the full deep photometric sample, some cells—especially those corresponding to rare or extreme color-magnitude combinations—may contain only galaxies without redshift measurements, even though the redshift sample itself is complete within its

footprint. These empty cells are therefore a consequence of incomplete color-space overlap between the photometric and redshift samples, not of missing redshift measurements for galaxies in the latter.

This setup balances the need for resolution with the requirement that the fraction of galaxies mapping to deep cells lacking redshift information remains below 1%, a condition we verified also holds in Y6.

Figures 5 and 6 showcase the six deep SOMs and six wide SOMs, respectively. The left panel of Fig. 5 displays the mean redshift ($\langle z|c \rangle$) associated with each deep SOM cell for each tomographic bin. As expected, the redshift structure across the SOM reflects the underlying color-magnitude distribution of galaxies: contiguous regions in SOM space typically correspond to galaxies with similar photometric properties and, consequently, similar redshifts. Each tomographic bin highlights distinct regions of the SOM, illustrating how the binning aligns with the SOM’s learned structure. The progression of redshift patterns across bins confirms the consistency of the binning strategy and demonstrates the SOM’s effectiveness in organizing photometric observables in redshift space. The right panel of Fig. 5 instead presents the corresponding redshift scatter, $\sigma(z|c)$, for each deep SOM cell. Cells with low $\sigma(z|c)$




FIG. 6. Visualization of the six wide SOMs, one for each tomographic bin, each composed of 1024 cells (32×32). Left: average redshift for each wide SOM cell $\hat{c}$. Right: standard deviation on the redshift distribution for each wide SOM cell $\hat{c}$.

indicate regions of high redshift precision, where the photometric observables tightly constrain the redshift distribution. Conversely, cells with high redshift scatter often trace color-magnitude degeneracies or lower signal-to-noise regimes, where photometric redshifts become less reliable. These maps are useful for identifying regions of high redshift uncertainty that may be down weighted or excluded in the final analysis.

To better understand the origin and impact of these high-scatter cells, we examine their occupancy by both the deep Balrog photometric sample and the spectroscopic redshift sample in Fig. 15. This analysis shows that high- $\sigma(z|\hat{c})$ cells are not generally associated with low occupancy or poor spectroscopic sampling. Instead, they reflect intrinsic limitations in the photometric data: in these regions, galaxies with similar photometric properties have broad or multimodal redshift distributions that cannot be cleanly separated, even with the deep multiband photometry. We further quantify the effect of these cells by showing, in Fig. 16, the number of galaxies per $\sigma(z|\hat{c})$ bin for both the deep and wide samples. These histograms confirm that high-scatter cells are extremely rare and contribute negligibly to the overall redshift calibration. For more details, see Appendix A.

Figure 6 shows the corresponding information for the wide SOMs. The left panel displays the average redshift per each wide SOM cell, $\langle z | \hat{c} \rangle$, while the right panel shows the redshift scatter $\sigma(z | \hat{c})$. Although the wide data have increased photometric noise, clear redshift gradients are still visible across SOM space, indicating that the wide photometry retains significant discriminatory power. Compared to the deep SOMs, the wide SOMs show smoother and more coherent features, with well-defined regions of increasing redshift. The redshift scatter $\sigma(z | \hat{c})$ in the right panel shows that the highest uncertainties tend to lie at the edges and corners of the SOMs—regions that typically correspond to low-density or degenerate areas of the color-magnitude space.

We further assess the potential impact of these high- $\sigma(z|\hat{c})$ regions in Fig. 16, which shows the number of wide-field galaxies per value of $\sigma(z|\hat{c})$. This distribution confirms that only a small fraction of the wide sample lies in SOM cells with large redshift scatter, reinforcing the conclusion that the overall redshift uncertainty remains well controlled even in the wide data.

In this work, the training phase was performed using a subset of the deep sample, which differed from the fiducial one by 1%. Catastrophic outliers were removed by

eliminating objects with nonphysical colors (i.e. colors required to be > -1). This step was necessary to prevent these extreme cases from disproportionately influencing the training process. The entire deep sample was then assigned to the deep SOM, ensuring that redshift characterization remained stable. Although this process may introduce double peaks in some SOM cells, the impact is expected to be smaller than the risk of completely misclassified cells.

IV. CHARACTERIZATION OF SOURCES OF UNCERTAINTY

In this section we describe the process we employ to characterize the uncertainties in the SOMPZ redshift distribution. We consider the following four sources of uncertainties:

- (i) *Sample variance (SV)* of the deep and redshift samples, due to the limited area of the deep fields ($\sim 9 \text{ deg}^2$). This is modeled via the *three-step Dirichlet* (3sDir) analytical model described in Sec. IV A.
- (ii) *Shot noise (SN)*, arising from the finite number of galaxies in the deep and redshift samples. This is also captured within the 3sDir framework.
- (iii) *Photometric calibration uncertainty (ZPU)*, introduced by errors in the deep-field photometric zero points. This is modeled by perturbing the deep photometry and rerunning the SOMPZ pipeline, as described in Sec. IV B.
- (iv) *Redshift sample uncertainty (RU)*, which accounts for residual biases in the redshift sample itself. This is modeled via coherent shifts in magnitude-redshift bins based on spectroscopic comparisons, as described in Sec. IV C.

Compared to the DES Y3 framework in [19], two sources of uncertainty are no longer included in this work. The uncertainty introduced by discretizing the continuous color space was gauged to be negligible in Y3. Similarly, the uncertainty in Balrog, which was also estimated as negligible in Y3, is even less significant given that the Y6 Balrog sample is approximately five times larger than in Y3 (see Sec. II). Potential systematic effects—e.g., from mismatches between injected and real sources—are expected to be subdominant, as the injection and measurement pipeline is largely unchanged from Y3 aside from minor improvements in PSF modeling and weighting.

Following its application to Y3 data, the 3sDir method models the sample variance and shot noise uncertainties on the $n(z)$, by explicitly sampling fluctuations in both the redshift calibration sample and the deep field, based on finite galaxy counts in each. This results in a set of $n(z)$ realizations that reflect expected statistical fluctuations, without additional systematic shifts in redshift or color space. ZPU and RU are then incorporated on top of 3sDir, but differently than the latter, their handling has been changed in Y6: ZPU is modeled by perturbing deep-field

photometry and recomputing SOM assignments, affecting how galaxies are mapped to redshift. RU is modeled by applying empirical redshift shifts of the narrow-band photometric redshift sample derived from spectroscopic comparisons, coherently across magnitude-redshift bins.

These perturbations can alter the effective shape of $n(z)$ in a way that either broadens or narrows it, depending on the direction and magnitude of the induced shifts. In some bins, they may introduce coherent changes that reduce the high-redshift tails or shift galaxies toward narrower redshift distributions. Adding an extra category of uncertainties (e.g., redshift, zero point) generally increases the uncertainty on the mean and width of the redshift distribution. The magnitude of this effect can vary by uncertainty type and may differ between the two metrics: in some cases, only the mean or only the width shows a significant increase (e.g., $>10\%$). This is because uncertainty propagation is a nonlinear process, encoded at the SOM $n(z)$ level rather than directly at the level of the mean and width. As a result, it impacts the full shape of the redshift distribution, with the mean and width being affected only indirectly as derived quantities. This highlights the limitation of using individual summary statistics—such as the mean or width—to represent the total uncertainty in a redshift distribution, as these metrics can fail to capture coherent variations in shape introduced by different uncertainty sources. In contrast, ensembles of $n(z)$ realizations fully represent these variations and their impact on cosmological inference, as illustrated in Fig. 17 and Table V, where RU and ZPU shift the dominant modes and their priors.

The final ensemble of $n(z)$ realizations from SOMPZ is shown in Fig. 7, which collectively spans the region where the true redshift distribution is likely to land. We generate these by generating 100 realizations for each of the ZPU and RU sampled values that parametrize the corresponding uncertainty, and then generating 1,000,000 additional realizations for each of the 100 SOMPZ $n(z)$ estimates using the 3sDir model. This yields a total of 100,000,000 redshift distributions realizations per tomographic bin, capturing the full propagated uncertainties for the MagLim++ lens sample.

Figure 8 presents the variance in the mean redshift $\langle z \rangle$ (left panel) and width parametrized as σ_{68}^2 (right panel) across six tomographic bins, decomposed by the source of the uncertainty. This comparison allows us to assess the relative contribution of SV, SN, RU, and ZPU to the total redshift uncertainty. The contribution of each uncertainty source to the variance of the mean and width is computed by differencing the variance of ensembles of

² σ_{68} is defined as half the width of the interval that encloses the central 68.25% of the redshift probability distribution. It is computed as $(z_{84.075} - z_{15.825})/2$, where z_p is the redshift at the p th percentile of the cumulative distribution. This provides a robust measure of the redshift scatter, analogous to the standard deviation for non-Gaussian or asymmetric distributions.




FIG. 7. Ensemble of 100,000,000 $n(z)$ distributions for the MagLim++ lens sample, for each of the six tomographic bins. The spread across realizations reflects the four sources of uncertainty included in cosmological analyses: sample variance, shot noise, zero point, and redshift uncertainty. The legend reports also the mean redshift and its uncertainty.




FIG. 8. Relative contribution of each source of uncertainty for the mean (left) and the width (right) for each of the six tomographic bins. We show the individual contributions from sample variance (SV, blue), shot noise (SN, light blue), redshift sample uncertainty (RU, purple), and photometric zero point uncertainty (ZPU, green). No single source dominates all bins or metrics, underscoring the need to propagate all uncertainty components when modeling $n(z)$. In bin 6, where the redshift distribution is broader and more sensitive to magnitude-dependent shifts, ZPU drives a particularly large uncertainty on the width. The magnitude of the effect varies by source and bin, and in some cases only the width shows a significant increase (e.g., $>10\%$). This reflects the nonlinear nature of uncertainty propagation at the SOM $n(z)$ level, which impacts the full distribution shape, with the mean and width being affected only indirectly as derived quantities. To illustrate the relative magnitude of each source of uncertainty in each bin, and its variation with redshift, we rescale the total variance in this figure to match the combined uncertainty (see Table III).

$n(z)$ realizations with and without that source included (e.g., $\sigma_{SV}^2 = \text{Var}[SV + SN] - \text{Var}[SN]$), with any residual between the total variance and the sum of individual terms assigned to a cross-term accounting for nonlinear interactions. In the following subsections, we describe the methodology used to model each component and interpret its behavior across bins in the context of the plots and numerical values.

In what follows, we describe in detail how each of the four uncertainty components—SV, SN, ZPU, and RU—is modeled and propagated.

A. Sample variance and shot noise

Most redshift distributions in photometric surveys are inherently uncertain due to two dominant sources of error: sample variance and shot noise. Sample variance arises from the fact that redshift calibration samples are typically drawn from small patches of the sky, meaning they do not perfectly represent the global distribution of galaxies due to fluctuations from large-scale structure. Shot noise, on the other hand, is a consequence of the finite number of galaxies in a given calibration sample, leading to Poisson fluctuations in redshift distributions. To properly propagate these uncertainties into cosmological analyses, we employ the *three-step Dirichlet resampling (3sDir)* method developed in [21], which generates realizations of the redshift distribution that include both sources of uncertainty. This method was applied to the DES Y3 dataset for the calibration of the source galaxy redshifts [50] and of the lens galaxy redshifts [19].

The 3sDir approach treats redshift calibration as a hierarchical process, in which the overall redshift distribution of a galaxy sample is built by combining the redshift distributions of smaller, distinct subpopulations. These subpopulations are typically defined based on observational or photometric properties (such as color or magnitude), and each is associated with its own underlying redshift distribution. The final, observed redshift distribution is then obtained by appropriately weighting and summing these individual components, reflecting the relative abundance of each subpopulation in the full sample. Instead of assuming a fixed redshift distribution for a given calibration sample, we perturb it by drawing three sequential Dirichlet samples, each capturing a different contribution to the total uncertainty. The Dirichlet distribution is chosen because it is a natural way to model probability distributions while ensuring that all values remain positive and normalized to one. Each sampled realization of the redshift distribution is thus self-consistent, properly incorporating both shot noise and sample variance in a statistically robust manner.

The three steps of 3sDir correspond to different levels of uncertainty propagation, where the output of each step serves as the input for the next. These also require to define a *phenotype* as a single SOM cell and a *superphenotype* as a broader region of the SOM constructed by merging

adjacent cells with similar color-redshift properties. The superphenotypes are defined by grouping deep SOM cells that are close in redshift, such that the superphenotypes become nearly disjoint in redshift space (details in [21]). The three steps are:

- (i) Superphenotype sampling ($p(T)$): First, we sample the fraction of galaxies in each superphenotype T . This step introduces large-scale sample variance by allowing fluctuations in the relative abundance of different galaxy populations across the sky. A Dirichlet sample is drawn for $p(T)$, giving a realization of how many galaxies belong to each superphenotype.
- (ii) Redshift given superphenotype ($p(z|T)$): Next, for each superphenotype, we sample the redshift distribution within that category. This introduces sample variance within each galaxy type, capturing the fact that different realizations of the universe will have fluctuations in redshift distributions due to large-scale structure. The effect of large-scale structure is implicitly encoded in the redshift distributions derived from a specific calibration field (e.g., COSMOS), which reflect the density fluctuations present in that region. Drawing a second Dirichlet sample for $p(z|T)$ introduces variation consistent with the expected cosmic variance from limited-area observations, while ensuring the sampled distributions remain properly normalized.
- (iii) Phenotype given redshift ($p(t|z, T)$): Finally, we sample the phenotype t distribution within each redshift bin. Here, the term “redshift bin” refers to fine-grained bins ($\Delta z = 0.05$) used internally in the construction of the redshift distribution. The redshift separation of $\Delta z = 0.05$ is chosen to minimize the correlation between bins, which allows us to sample each redshift bin independently in 3sDir (see e.g. Fig. 14 in [21]). This step captures shot noise, as it represents the Poisson fluctuations in the number of galaxies assigned to each fine-scale classification. A third Dirichlet sample is drawn for $p(t|z, T)$, distributing galaxies within each redshift bin among the phenotypes while preserving normalization.

Putting all components together, the 3sDir method generates a full redshift-phenotype joint distribution modeled as,

$$f(z, t) = \sum_T p(T) p(z|T) p(t|z, T), \quad (4)$$

where each component is sampled via a Dirichlet distribution as described above. This hierarchical factorization enables a structured propagation of uncertainty from large-scale sample variance [through $p(T)$ and $p(z|T)$] and shot noise [through $p(t|z, T)$].

Since we are interested in the marginalized redshift distribution, we sum over all phenotypes,

$$f(z) = \sum_t f(z, t) = \sum_T p(T) p(z|T). \quad (5)$$

This marginalization ensures that the final $f(z)$ properly reflects both sample variance and shot noise while remaining agnostic to the internal phenotypic structure.

Since each realization of the redshift distribution is generated sequentially in this manner, every sampled realization includes all sources of uncertainty simultaneously; it is not the case that some samples contain only sample variance and others only shot noise. Instead, across an ensemble of realizations, we recover the full range of fluctuations in the redshift distribution. The expectation value of these realizations converges to the redshift distribution measured in the redshift sample, meaning that 3sDir introduces a statistically consistent uncertainty around the input distribution without biasing it. This ensures that the propagated uncertainties reflect both sample variance and shot noise, while preserving the original mean z .

By applying 3sDir to redshift calibration samples, we obtain a statistically robust set of resampled redshift distributions that can be directly incorporated into cosmological analyses. The hierarchical nature of 3sDir makes it particularly well-suited for methods that rely on multitiered calibration steps, such as our SOMPZ method.

The 3sDir method, as applied here, assumes that the redshift sample used for calibration originates entirely from the COSMOS field. In practice, high-quality redshift information is also available in other deep fields, particularly in X3, which contributes spectroscopic and narrow-band redshifts at the bright end. However, these external redshifts often come from heterogeneous surveys with complex, color- and magnitude-dependent selection functions that are difficult to model accurately. To avoid introducing potential biases in the calibration due to these poorly characterized selection effects, we conservatively choose to exclude non-COSMOS redshifts from the uncertainty model. While this increases the sample variance—especially for the brighter MagLim++ lens galaxies, for which the COSMOS coverage is less dominant (62%, 65%, 65%, 75%, 85%, 87%, from lowest to highest bin)—it ensures that the estimated uncertainty in $p(z|c)$ remains robust. The choice therefore leads to a more conservative (i.e., broader) uncertainty estimate by avoiding the risk of underestimating redshift uncertainties due to unmodeled systematics in non-COSMOS fields.

We expect the impact of this approximation on the uncertainty estimate to remain small, since the dominant source of variance captured by 3sDir arises from the limited area of COSMOS itself. Nonetheless, future extensions of 3sDir could incorporate multiple deep fields with redshift coverage, allowing a more accurate characterization of the sampling variance and potential cross-field systematics.

As shown in Fig. 8, SV (dark blue) dominates the uncertainty on the mean redshift and width of $n(z)$ at low

redshift, especially in Bin 1. In contrast, SN (light blue) becomes increasingly important with redshift and overtakes SV in the higher bins—most notably in Bin 6. This is likely due to a combination of broader redshift PDFs and lower effective number densities in high-redshift regions of the SOM, especially after selection and matching cuts. The relative balance of SV and SN also differs slightly between the mean and width of $n(z)$, with SN contributing more prominently to the width in the high-redshift bins. These trends support the use of 3sDir as a practical and flexible method for propagating statistical uncertainties from the calibration sample, since it naturally captures the interplay between sample variance and shot noise across tomographic bins.

B. Deep fields photometric calibration

Accurate photometric redshift estimation relies on precise flux measurements across multiple bands, making it sensitive to small systematic errors in photometric calibration—commonly referred to as ZPU. In DES Y6, these uncertainties arise from imperfect calibration across the deep fields (COSMOS, X3, C3, E2), leading to systematic offsets in the measured fluxes. Such offsets can distort the flux-based mapping between deep and wide field galaxies in the SOMPZ pipeline, introducing biases in the inferred redshift distributions and, consequently, in cosmological constraints. To account for this, we simulate zero point shifts in each deep field and propagate their impact through the entire photo- z estimation pipeline. The resulting variations in the lens galaxy redshift distributions are then marginalized over in the cosmological analysis.

In practice, we fix the zero point of the COSMOS field as a reference and apply independent shifts to the zero points of the X3, C3, and E2 fields. For each realization, we draw independent zero point shifts in all bands ($ugrizYJHK_s$), perturb the corresponding galaxy fluxes and flux errors, and reassign the galaxies to the deep SOM based on their perturbed fluxes. The SOM-to-redshift mapping is then recomputed to produce a new $n(z)$. This reassignment captures how calibration errors can distort the mapping from flux space to SOM cell and, in turn, to redshift.

A simplifying assumption in this procedure is that the dominant impact of zero point uncertainties arises from the u band. This assumption is motivated by the fact that the u band has the largest calibration uncertainties and is not included in either the wide field observations or the Balrog simulations used to derive the deep-to-wide transfer function. The zero-point uncertainties (in magnitude space) are: 0.055 for the u band; 0.005 for the g , r , i , and z bands; and 0.008 for the J , H , and K_s bands (see Ref. [25] for details). As a result, we do not recompute Balrog for each perturbed realization, assuming that ZPU-induced changes to the transfer function are negligible.

To efficiently sample the multidimensional space of possible zero point shifts, we use Latin hypercube sampling

(LHS) in quantile space. LHS is a stratified, space-filling sampling technique that divides the range of each parameter into equal-probability intervals and draws one sample per interval, ensuring uniform coverage along each dimension. Unlike simple random sampling, this method avoids clustering and enables efficient exploration of high-dimensional parameter spaces with relatively few realizations. It balances efficiency and diversity in the combinations of shifts applied across different bands and fields. We generate 100 such realizations from the LHS sampling, each corresponding to a unique combination of zero point shifts across the deep fields and bands, perturbing the redshift distributions $p(z|c)$ and transfer functions $p(c, \hat{c})$. For each of these 100 realizations, we generate 1,000,000 3sDir realizations, obtaining the 100 million.

The contribution of photometric zero point uncertainties (green) shows a striking redshift dependence, as seen in Fig. 8. Bin 6 exhibits the largest ZPU-driven variance in both $\langle z \rangle$ and σ_{68} —by a wide margin—while Bin 2 shows a moderate contribution, particularly to the width. This pattern suggests that at high redshift, small shifts in photometric calibration more readily cause galaxies to cross SOM cell boundaries with steep redshift gradients, magnifying the impact on the inferred redshift distribution. The increased sensitivity in Bin 6 may also reflect the complex flux-redshift degeneracies and broader intrinsic scatter in the SOM mapping at faint magnitudes. By contrast, the larger uncertainties seen in Bin 2 are primarily driven by photo- z degeneracies around $z \sim 0.4$, coinciding with the $g-r$ filter transition, where small color changes can shift galaxies between competing redshift solutions.

Interestingly, this trend is not monotonic. In Bin 5—despite its proximity in redshift to Bin 6—the ZPU contribution is anomalously low. In fact, we observe that the standard deviation of the redshift realizations in the 3sDir + ZPU case is slightly smaller than in the 3sDir-only baseline. This counterintuitive result likely reflects non-linear interactions in the SOM assignment, where flux perturbations lead to a narrower redshift distribution by chance. Since uncertainty components do not combine linearly under this mapping, and the apparent reduction is not physically meaningful, we conservatively set the ZPU contribution to zero in Bin 5 when reporting decomposed uncertainties.

C. Photometric redshift biases

Achieving a representative and complete redshift sample for calibration is essential for accurate estimation of the redshift distribution $n(z)$. A sample that does not fully populate the color-magnitude space of the target population can lead to biased estimates due to selection effects [45]. In DES Y6, we mitigate this by leveraging high-quality photometric redshifts from narrow-band imaging surveys, as described in Sec. II B. While these narrow-band photo- z estimates are typically less precise than spectroscopic

redshifts on a per-galaxy basis, their uncertainties remain small compared to the much broader redshift resolution of the wide-field *griz* photometry. Crucially, these samples offer high completeness and wide coverage in the relevant color space.

Nevertheless, photometric redshifts are inherently subject to systematic uncertainties, arising from measurement errors, mismatches in template libraries, and limitations in the photo- z algorithms themselves [47,51,52]. If not accounted for, such biases can distort the inferred $n(z)$ and propagate into the cosmological inference. We therefore implement a redshift uncertainty calibration procedure designed to model systematic redshift shifts in a way that reflects these uncertainties while preserving the statistical properties of the underlying redshift distributions.

In DES Y3, the lens redshift distributions were derived using a combination of three heterogeneous calibration catalogs (spectroscopic compilations, PAU + COSMOS, and COSMOS30), with three realizations constructed by ranking the catalogs differently and assigning the redshift from the highest-ranked source available for each galaxy [19,50]. The final ensemble of realizations was obtained by averaging over these three samples with equal prior weight, which provided coverage across multiple calibration sets but did not explicitly propagate systematic uncertainties associated with incomplete or heterogeneous information.

To address this limitation, we adopt a different strategy for propagating redshift uncertainties. Rather than perturbing the redshifts of individual galaxies independently—which would mainly introduce random fluctuations that tend to average out across the large number of objects and therefore underestimate the impact of systematic effects—we instead apply coherent shifts defined at the level of magnitude-redshift bins. This method captures correlated redshift errors among galaxies with similar properties and more faithfully reflects the kind of biases expected from photo- z systematics. The procedure begins by matching galaxies from each narrow-band catalog (Cosmos Classic, Cosmos Farmer, PAUS Cosmos, PAUS W1, see Sec. II B for more details) to available spectroscopic redshifts. The matched galaxies are then binned in both magnitude and redshift. For each bin i , we compute the median redshift bias and the redshift scatter using the distribution of $dz \equiv z_{\text{photo}z} - z_{\text{spec}z}$ values in that bin,

$$b(i) = \text{median}(dz), \quad (6)$$

$$s(i) = \langle \sigma(dz^{(k)}) \rangle, \quad (7)$$

where $\sigma(dz^{(k)})$ is the sample standard deviation of the k th bootstrap resampling of the dz values in bin i , and the average $\langle \cdot \rangle$ is taken over 1000 bootstrap realizations. This approach provides a robust estimate of the scatter of the distribution, mitigating the effects of sampling noise and

accounting for uncertainty due to limited statistics in each bin.

The bias and scatter values are used to define the range of systematic shifts applied to galaxies in each bin. Specifically, we draw shifts Δz from a Gaussian prior characterized by,

$$\Delta z \sim \mathcal{N}(\mu = b(i), \sigma = s(i)),$$

where μ is the median redshift bias and σ is the scatter estimated as described above. All galaxies in a given bin receive the same shift.

We apply the same Latin hypercube sampling strategy described in Sec. IV B to sample coherent shifts across all magnitude-redshift bins. This approach ensures broad and efficient coverage of the uncertainty space while avoiding redundant or biased combinations of perturbations. Each realization of the redshift uncertainty is constructed as follows:

- (1) A systematic shift Δz is drawn for each magnitude-redshift bin using LHS.
- (2) All galaxies in a given magnitude-redshift bin are perturbed by the corresponding Δz .
- (3) The resulting redshift distribution $n(z)$ is computed from the modified sample.

This process is repeated N times (in our case, $N = 100$), generating an ensemble of perturbed redshift distributions. The final $n(z)$ uncertainty is characterized by the spread across these realizations.

Figure 9 summarizes the properties of the photometric redshift samples used in this calibration. It shows the number of galaxies per magnitude-redshift bin (left column), the median photometric redshift bias (center column), and the scatter in photometric redshift estimates (right column). Each row corresponds to one of the photometric redshift samples. The median bias and scatter values shown in the figure are used to define the means and widths of the Gaussian prior from which the systematic shifts Δz are drawn via LHS.

The choice of $\sigma = s(i)$, i.e., using the photo- z scatter in each bin to set the width of the Gaussian prior, is intentionally conservative. This is especially important for faint galaxies, where spectroscopic redshift coverage is sparse and the true systematic uncertainty may be underconstrained. For brighter galaxies with more robust spectroscopic calibration, alternative or tighter priors could be used, but in this work we apply a uniform treatment across all magnitude bins for consistency.

By applying systematic shifts at the joint magnitude-redshift bin level, this method avoids artificially underestimating uncertainties that can result from random, uncorrelated perturbations canceling out. This approach differs with the methodology adopted in DES Y3, where redshift sample uncertainty was modeled by constructing three separate redshift realizations based on different catalogs

(spectroscopic, PAUS + COSMOS, and COSMOS30), and propagating them through the full pipeline with equal weights. In that framework, the systematic uncertainty was captured by comparing results across different samples, each assumed to be independently valid. While this method accounted for catalog-level systematics, it treated galaxy redshifts as fixed once assigned from a particular catalog and did not explicitly model correlated redshift errors as a function of galaxy properties. In contrast, the DES Y6 approach implements a continuous and flexible model of redshift uncertainty, applying coherent perturbations across magnitude-redshift bins and sampling the space of possible biases via Latin hypercube sampling. This forward-modeling-compatible framework allows for a more granular treatment of systematic shifts and can be easily extended or refined. In future work, this could include assigning more informative priors to specific bins based on spectroscopic coverage or combining coherent and stochastic perturbations to capture both structured and residual uncertainties.

RU (purple in Fig. 8) shows a nontrivial pattern across tomographic bins, with distinct differences in its impact on the mean and width of $n(z)$. For the mean redshift, RU contributes most significantly in Bins 2 and 6, while in Bin 1 the estimated change in variance relative to the 3sDir baseline is slightly negative. The prominent RU contribution in Bin 2 is again consistent with the difficulties of resolving the $z \sim 0.4$ degeneracies associated with the $g-r$ filter transition, which amplify the sensitivity of this bin to limited redshift-sample statistics.

For the width σ_{68} , the strongest contributions appear in Bins 3 and 6, where magnitude-redshift degeneracies and limited spectroscopic coverage make the inferred distributions more sensitive to redshift calibration shifts. In contrast, RU has minimal impact on the width in the other bins. The small negative values observed in Bin 1 do not correspond to a true variance—since variance is by definition non-negative—but arise because we compute differential contributions relative to the 3sDir baseline. Owing to the nonlinear nature of the SOM mapping, coherent perturbations can sometimes redistribute galaxies in a way that reduces the apparent spread of the redshift distribution. These negative values have no impact on the propagated uncertainties though, which are always derived from the full set of realizations. We only set the Bin 1 contribution to zero when displaying the relative breakdown in Fig. 8, to avoid misinterpretation.

V. FINAL REDSHIFT DISTRIBUTIONS AND UNCERTAINTIES FOR MagLim++

The output of the SOMPZ framework described in the previous section is 100,000,000 realizations of the redshift distributions for each of the six tomographic bins, consistent with our estimated uncertainties from the photometric information. These are shown in Fig. 7. To derive the




FIG. 9. Distribution and photometric redshift performance across magnitude and redshift bins, for the four photometric redshift samples Cosmos Classic, Cosmos Farmer, PAUS Cosmos, and PAUS W1. Left: number of galaxies (with spectra) in each magnitude-redshift bin. Center: median photometric redshift bias ($z_{\text{photoz}} - z_{\text{specz}}$). Right: scatter of photometric redshift estimates. On the left column, we indicate with a white color cells that have less than 10 galaxies. For those cells, we extrapolate the values of median bias and photo- z scatter from the value of the closest (magnitude, redshift) cell with more than 10 galaxies, which explains the patterns shown in the and rightmost column. For example, the relatively large and constant median biases seen at high redshift in most panels, come from the 1–2 mag- z cells. Upon deeper inspection, these large median biases come from those cells being dominated by large outlier fractions, coming from confusing the Balmer and Lyman break.

final redshift distributions and uncertainties used for the cosmological inference, we first incorporate additional information from spatial clustering (Sec. VA), and then derive a representation of the final realizations that can be efficiently sampled in an Monte-Carlo Markov Chain framework (Sec. VC). The resulting calibration achieves typical uncertainties on the mean redshift at the 1%–2% level, corresponding to an average improvement of $\sim 20\%$ – 30% relative to DES Y3. The uncertainties on the width are likewise reduced in most bins, with improvements ranging from $\sim 20\%$ to $\sim 45\%$. This improvement

does not arise from a single component, but rather from the combination of several methodological and data-related advances in the Y6 analysis.

At the SOM-PZ level, improvements are primarily driven by reductions in methodological uncertainty terms relative to Y3. The substantially larger Balrog sample reduces shot noise in the estimation of the deep-to-wide transfer function, particularly in sparsely populated regions of the deep SOM. In addition, the spectroscopic portion of the redshift calibration sample was revised to better match the requirements of this analysis, restricting the calibration to a subset

of spectroscopic surveys and redshift-quality selections deemed most robust for this application. The Y6 analysis also adopts a more flexible treatment of redshift-sample uncertainty than Y3, modeling coherent redshift systematics as a function of galaxy properties through perturbations applied in magnitude-redshift bins. Additionally, the updated construction of the MagLim++ sample yields a cleaner and more homogeneous galaxy selection, improving the mapping between the deep and wide samples used in SOMPZ. Comparing the SOMPZ-only calibration between Y3 and Y6, we find an average reduction of $\sim 22\%$ in the uncertainty on the mean redshift, while the uncertainty on the distribution width is reduced by $\sim 8\%$ on average. These improvements are primarily driven by reductions in the methodological uncertainty contributions, while the sample variance component remains largely unchanged.

The reduction is not fully uniform across tomographic bins. The highest-redshift bin already had the largest uncertainties in Y3 and remains the least constrained bin in Y6. This likely reflects the intrinsic difficulty of calibrating this redshift range, together with the more complete uncertainty propagation adopted in Y6. In particular, the Y3 calibration for this bin relied on a combination of heterogeneous catalogs with equal prior weight (see Sec. IV C), which may have led to a less conservative estimate of the uncertainty. We therefore interpret the slightly larger Bin 6 error bars in Y6 as a more faithful accounting of the limitations in this regime, rather than a loss of constraining power due to data quality.

Additional information is then incorporated from clustering-redshift (wz) measurements, which further improves the uncertainties of both the mean redshift and the overall shape of the distributions. The impact of wz on the mean-redshift uncertainties is comparable to that observed in DES Y3, providing an additional reduction of $\sim 20\%$ on average across tomographic bins. In Y6, however, the wz constraint on the distribution widths is stronger, with typical improvements of $\sim 40\%$. This sensitivity to the width is likely aided by the extended redshift coverage of the reference sample, including QSOs, which help constrain the high-redshift tail of the distributions. The wz information is incorporated consistently within the importance-sampling framework mentioned in the following sections.

Table II provides an overview of the different elements entering the MagLim++ redshift calibration pipeline, indicating for each step whether it is introduced here or in an accompanying DES Y6 analysis paper.

A. Incorporating clustering information

As outlined in the introduction, the final redshift calibration methodology consists of the statistical combination of photometric and clustering information. In fact, we leverage cross-correlations between lens galaxies and spectroscopic reference samples to refine the redshift

TABLE II. Summary of the main components of this work and their relation to companion DES papers within the Y6 3×2 pt analysis.

Component	Defined/performed in
MagLim++ lens sample definition	[6]
Redshift sample definition	This work
Deep samples	[25]
Deep and wide SOMs	This work
SOMPZ transfer function	This work, from the Balrog catalog [20]
MagLim++ SOMPZ calibration	This work
SOMPZ uncertainty estimation	This work
wz measurement	[Y6WZ]
Combination of wz and SOMPZ $n(z)$	[Y6WZ]
Mode projection methodology	[Y6Modes]
Mode projection of the combined $n(z)$	This work
Tests on simulated data vectors	This work

distributions from SOMPZ further. The clustering redshift (wz) technique relies on the spatial association of photometric sources with spectroscopic samples. For example, we measure the small-scale angular cross-correlations between our target sample and data from the baryonic oscillation spectroscopic survey (BOSS) [53,54] and the extended-baryon oscillation spectroscopic survey (eBOSS) [55–57].

The full implementation of this method for both lens and source samples—including the modeling framework, systematic treatment, and application to DES Y6—was developed and carried out in [Y6WZ], on which this work builds. We summarize the key components here for completeness.

We model the observed angular correlation function w_{ur} between photometric unknown samples (indexed by u) and spectroscopic reference samples (indexed by r) as,

$$\hat{w}_{ur} = f(b_u, b_r, n_r(z), n_u(z), s_u), \quad (8)$$

where b_x represents galaxy bias, $n_x(z)$ denotes redshift distributions, and s_u encapsulates systematic effects such as magnification. The likelihood function for clustering redshifts is then given by,

$$\ln \mathcal{L}(w_{ur} | b_u, b_r, n_r, n_u, s_u) \propto (w_{ur} - \hat{w}_{ur})^T C_w^{-1} (w_{ur} - \hat{w}_{ur}), \quad (9)$$

where C_w is the covariance of the cross-correlation measurements. We obtain posterior constraints on $n_u(z)$ by reweighting the SOMPZ realizations with the clustering-redshift likelihood, which already accounts for nuisance parameters such as galaxy biases and magnification.

To statistically combine the photometric modeling from SOMPZ with the external constraints from wz, we adopt an importance sampling scheme that reweights the SOMPZ

realizations based on their consistency with the clustering-redshift measurements. This approach treats the ensemble of SOMPZ redshift realizations as a prior over plausible $n(z)$ shapes, and uses the WZ-derived likelihood to assign weights to each realization.

While clustering-based redshift estimates have been used in the past to construct $n(z)$ directly, this typically requires assumptions about galaxy bias evolution and can suffer from limited redshift resolution. In contrast, importance sampling allows us to incorporate WZ information without reconstructing the $n(z)$ from scratch or introducing additional parametric modeling. Thus, importance sampling offers a principled and computationally efficient way to combine these two sources of information. The resulting reweighted ensemble preserves the complex, non-Gaussian uncertainty structure of SOMPZ while integrating empirical constraints from large-scale structure. This yields a statistically consistent posterior distribution that reflects both photometric and clustering information. The final $n(z)$ distributions used in cosmological inference are drawn from this reweighted set.

B. Redshift distribution preprocessing

To avoid unnecessarily increasing the runtime of the cosmological inference, we “pile up” the high-redshift tail by assigning all probability at $z > 3$ to a single bin at $z = 3$. This removes a negligible and noisy component of the distribution that would otherwise slow down the Monte Carlo chains without contributing useful information. This truncation removes the numerically unstable, low-weight tail of $n(z)$ and ensures all realizations are defined on a common redshift grid, simplifying the subsequent cosmological inference without significantly affecting the clustering signal, since the clipped high-redshift tail contributes negligibly to the galaxy-galaxy correlations.

Second, the redshift distributions are then rebinned onto a basis of triangular kernels, which we called *sawtooth*. This transformation replaces the original binned $n(z)$ with a piecewise-linear approximation formed by overlapping triangular functions. The resulting representation is continuous and differentiable, which facilitates accurate computation of the derivatives of the cosmological data vector with respect to redshift shifts. While this rebinning introduces small artifacts in the form of low-amplitude wiggles, particularly near sharp features in the $n(z)$, these effects are negligible and do not impact the resulting cosmological inference. The triangular basis also ensures normalization and consistency across all realizations used in the compression.

C. Efficient sampling for the final redshift distributions

The estimation of photometric redshift distributions for galaxy samples introduces a high-dimensional set of nuisance parameters, as each redshift realization corresponds to

a different possible shape of the redshift distribution $n(z)$. Directly sampling over these high-dimensional realizations in a cosmological analysis is computationally infeasible, both because of the sheer dimensionality and the difficulty of defining a smooth, physically meaningful prior over such a space.

To address this, we employ a mode projection technique that reduces the effective dimensionality of the problem by identifying a basis of variations in $n(z)$ that significantly affect cosmological observables—such as the two-point correlation functions—while projecting out directions that leave the likelihood nearly unchanged. This approach ensures that only those variations in the redshift distribution that matter for cosmology are retained, greatly improving sampling efficiency without sacrificing accuracy.

In DES Y3, a similar dimensionality reduction was attempted using the Hyperrank method [58], which reparametrizes the space of redshift realizations using a Gaussian process latent space and ranks directions by their cosmological importance. While effective, Hyperrank requires learning a mapping from redshift realizations to a latent space and back, which can introduce approximations and technical complexity. In contrast, the mode projection method used here directly operates in the observable space and selects modes based on their impact on the likelihood, allowing for a more transparent and targeted marginalization over redshift uncertainty.

This approach follows the method outlined and validated in [Y6Modes], where the high-dimensional set of nuisance parameters is compressed into a much lower-dimensional set of modes. Specifically, the redshift distribution is expressed as a weighted sum of basis functions,

$$n(z) = \sum_k n_k b_k(z), \quad (10)$$

where the coefficients n_k are treated as nuisance parameters. Given a set of sampled realizations of $n(z)$, we determine a transformation that maps these into a low-dimensional space spanned by a set of modes u , such that only the variations that significantly affect cosmological observables are retained. The compression follows a linear transformation,

$$u = E(n - \bar{n}), \quad (11)$$

where $\bar{n}$ is the mean redshift distribution across realizations, and E is a projection matrix that isolates the directions in redshift space to which the observables (our two point correlation functions) are most sensitive. The transformation matrix is determined from the Fisher information matrix of the observables (i.e., the data vector), such that the retained modes are those that are best constrained by the data given its covariance. These modes correspond to directions in observable space with the highest signal-to-noise ratio,




FIG. 10. Cumulative discarded χ^2 as a function of the number of retained modes. The sharp decline indicates that most of the cosmologically relevant variance is captured in just a few modes, justifying our truncation choice.

which may or may not coincide with the directions that maximize constraints on cosmological parameters. This distinction arises because the procedure optimizes information content in the data space rather than directly in parameter space, although in practice many of the best-constrained data modes also carry significant cosmological information. The original redshift distribution can be reconstructed from the compressed representation as,

$$n(z) \approx \bar{n}(z) + \sum_{i=1}^M u_i U_i(z), \quad (12)$$

where $U_i(z)$ are the dominant modes of variation. The number of retained modes, M , is chosen to ensure that the discarded directions contribute negligibly to the likelihood function.

The role of the mode projection technique is illustrated in Figs. 10 and 11, which show the structure and impact of the dominant modes used in the redshift uncertainty model. Additional examples, including the visual representation of the modes themselves, are provided in Appendix B.

Figure 10 illustrates the number of modes retained in our analysis by showing the cumulative discarded χ^2 as a function of the number of modes included. This metric quantifies the statistical significance of any biases in the model for the observables that can be caused by the fluctuations in $n(z)$ that are in the discarded modes. We retain only the number of modes necessary to keep the cumulative discarded χ^2 across tomographic bins below a defined threshold of $\chi^2 < 0.15$. A threshold $\ll 1$ is chosen to ensure that any biases or variability in the inferred parameters are well below the 1σ variation attributable to

measurement errors. The rapid drop in χ^2 with increasing mode count demonstrates that only a small number of modes are needed to preserve nearly all sensitivity to redshift uncertainties, justifying the truncation strategy and confirming the efficiency of the compression. For the SOMPZ-only uncertainty, this criterion leads to four retained modes, while for the SOMPZ + WZ case, only three modes are needed.

Figure 11 compares the original sampled $n(z)$ distributions to those reconstructed using only the retained modes for both SOMPZ-only and SOMPZ + WZ uncertainties. The close agreement demonstrates that the dominant statistical properties of the redshift distributions are preserved despite a strong dimensionality reduction: for example, keeping $M = 3$ modes per tomographic bin reduces the number of free parameters per realization from $6 \times 300 = 1800$ (one per histogram bin) to $6 \times 3 = 18$ mode amplitudes, a factor of ~ 100 reduction. The mean and modes basis (each of length 300) are fixed and stored once; only the mode coefficients vary between realizations. These reconstructions visually compare to the uncompressed $n(z)$ shown in Fig. 7. The bottom panels show the residuals $(n_{\text{recon}} - n_{\text{orig}})/[1 + \langle n_{\text{orig}} \rangle]$. A value of 0.1 in this metric corresponds to differences well below the statistical uncertainty in the cosmological analysis, and is therefore tolerable. A small wiggle is visible in the residuals: as previously mentioned, this feature is introduced during the conversion to the sawtooth binning and does not originate from the mode compression itself. Its amplitude is negligible relative to the overall uncertainty and has no discernible impact on cosmological constraints.

This method provides a way to account for redshift uncertainty while significantly reducing computational complexity. Instead of defining a prior directly over the high-dimensional space of redshift distributions $n(z)$, we parametrize the uncertainty in the cosmological inference using the u coefficients. A prior over u can be efficiently learned using density estimation techniques (e.g., Gaussian mixtures, kernel density estimation) applied to a suite of $n(z)$ realizations. This low-dimensional representation preserves the dominant modes of variation in $n(z)$, allowing the inference pipeline to marginalize over redshift uncertainty with significantly reduced computational cost, while retaining all relevant cosmological information.

For source redshifts, the formalism naturally accommodates correlations between tomographic bins. For lens redshifts, where the MagLim++ binning was explicitly designed to produce independent redshift distributions, we confirm that the mode amplitudes remain uncorrelated in the cosmological posteriors, justifying the per-bin treatment of uncertainties. This ensures an accurate and efficient propagation of redshift calibration errors without introducing unnecessary nuisance correlations. Additional validation of this uncorrelated treatment is provided in Appendix B.




FIG. 11. Comparison of the redshift distributions for the six tomographic bins of MagLim++, reconstructed from the modes. The blue lines represent the average SOMPZ $n(z)$ reconstructed from the modes, with the $\pm 1\sigma$ and $\pm 2\sigma$ bands showcasing the realizations. The orange lines and bands represent the SOMPZ + WZ $n(z)$ reconstructed from modes. Solid and dashed styles are used solely to differentiate tomographic bins in regions where the curves overlap. These are the distributions marginalized over in the cosmological inference. The lower panels show the scatter of the relative residuals across realizations, defined as $(n_{\text{recon}} - n_{\text{orig}})/(1 + \langle n_{\text{orig}} \rangle)$, which suppresses artificial amplification in low-density regions. The reconstructed distributions successfully preserve the key statistical features of the original samples while significantly reducing dimensionality. The small wiggles visible in the residuals are not due to the mode compression but are introduced during the rebinning of $n(z)$ using the triangle kernel representation.

VI. IMPACT ON COSMOLOGY

In this section, we quantify how different assumptions about the redshift distributions and their associated uncertainties affect the inference of cosmological parameters. To this end, we use noiseless synthetic data vectors constructed of the three two-point functions—galaxy clustering $w(\theta)$, galaxy-galaxy lensing $\gamma_t(\theta)$, and cosmic shear $\xi_{\pm}(\theta)$ —generated using real data redshift distributions. By repeating the inference with different $n(z)$ estimates or

different marginalization schemes for redshift uncertainties, we can isolate their impact on the cosmological constraints.

In particular, we evaluate the impact of the redshift uncertainty parametrization, contrasting the fiducial Y6 mode-based approach against the simpler shift-and-stretch model adopted in the Y3 analysis (Sec. VI A). We also compare the use of lens redshift distributions estimated from SOMPZ alone to those obtained from the fiducial SOMPZ + WZ combination (Sec. VI B).

TABLE III. Summary of values for systematic uncertainties and center values for mean ($\langle z \rangle$, top panel) and width (σ_{68} , bottom panel) for the $n(z)$ distributions. The various components are computed as described in Sec. IV and as they are not completely independent they do not sum up to the total value. The values related to SOMPZ and SOMPZ + WZ refer to Fig. 7. The final $n(z)$ which has been used in the cosmological analysis is the bottom line.

	Mean $\langle z \rangle$					
	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5	Bin 6
Uncertainty	$z \in [0.2, 0.4]$	$z \in [0.4, 0.55]$	$z \in [0.55, 0.7]$	$z \in [0.7, 0.85]$	$z \in [0.85, 0.95]$	$z \in [0.95, 1.05]$
SOMPZ	0.3079 ± 0.0078	0.4343 ± 0.0106	0.6270 ± 0.0093	0.7771 ± 0.0076	0.8992 ± 0.0081	0.9970 ± 0.0152
SOMPZ + WZ	0.3059 ± 0.0063	0.4352 ± 0.0077	0.6235 ± 0.0073	0.7777 ± 0.0065	0.9027 ± 0.0062	1.0112 ± 0.0131
	Width σ_{68}					
	$z \in [0.2, 0.4]$	$z \in [0.4, 0.55]$	$z \in [0.55, 0.7]$	$z \in [0.7, 0.85]$	$z \in [0.85, 0.95]$	$z \in [0.95, 1.05]$
SOMPZ	0.0769 ± 0.0049	0.0947 ± 0.0059	0.0637 ± 0.0049	0.0756 ± 0.0039	0.0653 ± 0.0050	0.0970 ± 0.0099
SOMPZ + WZ	0.0793 ± 0.0027	0.0916 ± 0.0040	0.0605 ± 0.0021	0.0692 ± 0.0021	0.0604 ± 0.0028	0.0889 ± 0.0088

Each synthetic data vector is computed at fixed cosmological parameters and includes all relevant two-point functions evaluated using the fiducial angular scale cuts defined in [59]. Cosmological inference is then performed using the CosmoSIS pipeline [60], assuming a flat Λ CDM (cold dark matter) cosmology with five free parameters: $\{\Omega_m, \sigma_8, h, n_s, \Omega_b\}$, and fixing the sum of neutrino masses $\sum m_\nu$ to the minimum allowed value of 0.06 eV. We adopt a linear galaxy bias model for the lens sample and include the full set of nuisance parameters used in the DES Y6 analysis:

- (i) one linear galaxy bias parameter per lens bin;
- (ii) intrinsic alignment parameters from the tidal alignment and tidal torquing (TATT) [61] model: $\{a_1, a_2, \eta_1, \eta_2, b_{TA}\}$;
- (iii) one shear calibration parameter m_i per source bin;
- (iv) one magnification slope parameter α_i per lens bin;
- (v) source sample redshift uncertainty parameters: seven mode coefficients (for all bins); and
- (vi) lens redshift uncertainty parameters: either three orthogonal redshift mode coefficients per tomographic bin (fiducial case), or one shift and one stretch parameter per lens tomographic bin, depending on the test.

We adopt flat priors for all cosmological parameters and informative priors for nuisance parameters, listed in Table IV. In particular, the priors on magnification and intrinsic alignments follow those derived in the DES Y6 modeling framework [59]. The priors on the calibration parameters of the source galaxy samples, the multiplicative shear bias m and the source redshift distribution modes, are treated as correlated. We assume the redshift distributions and corresponding priors as estimated from data: for the MagLim++ lens sample, we use the results from this work; for the METADTECT source sample [62], we use the calibration from [Y6ShearPZ]. In both cases, the redshift distributions are combined with WZ constraints following the procedure described in [Y6WZ]. The covariance matrix is computed using CosmoCov [63] at the same cosmology as the data vector, and includes contributions from shape noise

and galaxy shot noise. Parameter inference is performed using the nested sampler NAUTILUS [64], and convergence is assessed using internal stopping criteria based on the stability of the likelihood bounds and the effective sample size, as implemented by default in CosmoSIS.

A. Characterization of the redshift uncertainties

In the DES Y3 analysis, the uncertainty in the MagLim redshift distributions was modeled using a two-parameter approach. Specifically, for each tomographic bin, we used a fiducial redshift distribution $\bar{n}(z)$, derived by averaging over the ensemble of redshift distribution realizations. The uncertainty was parametrized in terms of a shift in the mean redshift, Δz , and a stretch in the width of the distribution, Δw . The priors on these parameters were derived by fitting the fiducial $n(z)$ to the $n(z)$ obtained from a clustering redshift measurement. These parameters were sampled from Gaussian distributions,

$$\Delta_z^i \sim \mathcal{N}(0, \Delta_z), \quad \sigma_z^i \sim \mathcal{N}(0, \sigma_z), \quad (13)$$

where Δ_z and σ_z were estimated from the scatter among the realizations. The fiducial $\bar{n}(z)$ was then modified at each step of the inference via,

$$n_{\text{shifted}}^i(z) = \bar{n}(z - \Delta_z^i),$$

$$n_{\text{shifted and stretched}}^i(z) = \frac{1}{\sigma_z^i} n_{\text{shifted}}^i(z) \left(\frac{z - \langle z \rangle}{\sigma_z^i} + \langle z \rangle \right). \quad (14)$$

While this method captured leading-order uncertainties in the redshift distributions, it imposed restrictive assumptions on their form, allowing only global shifts and rescalings. This can overlook more structured or localized variations in $n(z)$ that arise in realistic scenarios.

In Y6, we move beyond the shift/stretch model and adopt the more flexible mode-based marginalization framework described in Sec. VC and detailed in [Y6Modes]. This approach decomposes redshift uncertainties into a set of

TABLE IV. Parameters and their priors used in the fiducial MagLim Λ CDM and w CDM analyses. The parameter w is fixed to -1 in Λ CDM. Square brackets denote a flat prior, while parentheses denote a Gaussian prior of the form $\mathcal{N}(\mu, \sigma)$. The priors on the calibration parameters of the source galaxy samples—namely the multiplicative shear bias m and the source redshift distribution modes—are correlated.

Parameter	Fiducial	Prior
Cosmology		
Ω_m	0.31	[0.1, 0.6]
$A_s 10^9$	1.831	[0.5, 5.0]
n_s	0.965	[0.93, 1.00]
w	-1.0	Fixed
Ω_b	0.051	[0.03, 0.07]
h_0	0.69	[0.58, 0.80]
$m_\nu [\text{eV}]$	0.77	[0.06, 0.6]
Linear galaxy bias		
b_i	1.54	[0.8, 3.0]
b_i	1.81	[0.8, 3.0]
b_i	1.85	[0.8, 3.0]
b_i	1.76	[0.8, 3.0]
b_i	1.93	[0.8, 3.0]
b_i	1.90	[0.8, 3.0]
Lens magnification		
C_1	0.43	$\mathcal{N}(1.58, 0.51)$
C_2	0.30	$\mathcal{N}(0.30, 0.48)$
C_3	1.75	$\mathcal{N}(1.75, 0.39)$
C_4	1.94	$\mathcal{N}(1.94, 0.35)$
C_5	1.56	$\mathcal{N}(1.56, 0.71)$
C_6	2.96	$\mathcal{N}(2.96, 0.95)$
Lens photo z		
$u_{i,m}^l (i \in [1, 6], m \in [1, 3])$	0	$\mathcal{N}(0, 1) \in [-3, 3]$
Source photo z		
$u_m^l (m \in [1, 7])$	0	$\mathcal{N}(0, 1) \in [-3, 3]$
Intrinsic alignment		
z_0	0.3	Fixed
A_1	0.13	$[-1, 3]$
A_2	-0.2	$[-3, 3]$
η_1	2.0	$\mathcal{N}(0.0, 3.0) \in [-5, 5]$
η_2	2.0	$\mathcal{N}(0.0, 3.0) \in [-5, 5]$
b_{TA}	1.0	fixEd
Shear calibration		
m^1	0.0	$[-0.1, 0.1]$
m^2	0.0	$[-0.1, 0.1]$
m^3	0.0	$[-0.1, 0.1]$
m^4	0.0	$[-0.1, 0.1]$

orthogonal modes U_i derived from an ensemble of redshift realizations, which represent principal directions of variation in the redshift distributions. For each tomographic bin of the SOMPZ + WZ $n(z)$, we marginalize over three corresponding mode coefficients $u_{i,j}$ (4 modes for SOMPZ

only) that parametrize the amplitude of variation along each mode. These mode coefficients serve as redshift nuisance parameters in the cosmological inference. This enables us to capture a much broader class of redshift uncertainties without imposing strong priors on their functional form. The priors and posteriors of the leading mode coefficients $u_{i,0}$ are shown in Appendix B. As a minor numerical artifact of the linear mode combination, $n(z)$ can become slightly negative in the very low-density tails. We verified that clipping these excursions to zero has a negligible impact on cosmological constraints; see Appendix C.

In Fig. 12, we compare the cosmological constraints obtained using these two approaches. Specifically, we marginalize over the source redshift uncertainties using the SOMPZ + WZ modes, and compare the resulting S_8 – Ω_m posteriors when applying either the shift/stretch method or the mode-based marginalization for the lens sample. For the 2×2 pt analysis, the joint analysis of galaxy clustering and galaxy-galaxy lensing, the contours are nearly identical, whereas for 3×2 pt the shift/stretch method yields slightly tighter constraints—likely due to a mild underestimation of uncertainties by the parametric model. The marginalized posterior values are listed in Table V. The close agreement between the two approaches offers a useful consistency check, but the mode-based framework is ultimately preferred: it captures a broader range of redshift variations, ensures consistency across samples and probes, and provides a more robust basis for future analyses.

B. Impact on cosmology from the MagLim++ redshift uncertainties

The accuracy of photometric redshift estimates directly impacts cosmological parameter constraints in the 2×2 pt analysis. In this section, we quantify how uncertainties in the MagLim++ redshift distribution propagate into cosmological inference. As described in the previous section, we perform this test using a simulated data vector, enabling us to isolate the impact of redshift uncertainties in a controlled setting.

Figure 13 shows constraints in the (Ω_m, S_8) plane under three different treatments of the lens redshift distribution: (i) assuming a fixed $n(z)$; (ii) marginalizing over SOMPZ-derived modes; and (iii) marginalizing over SOMPZ + WZ modes. In all cases, the source redshift distribution is treated using SOMPZ + WZ modes.

The three chains are centered around a common region in parameter space, but the posterior widths depend on the treatment of redshift uncertainties. As expected, the fixed- $n(z)$ case yields the tightest constraints, since it assumes perfect knowledge of the lens redshift distribution. Marginalizing over SOMPZ modes introduces additional uncertainty and broadens the contours. Including WZ information alongside SOMPZ typically helps recover some constraining power. This pattern is clearly observed




FIG. 12. Comparison of marginalization methods for photometric redshift uncertainties for 2×2 pt (left) and 3×2 pt (right). The figure contrasts the shift and stretch approach (green) with the mode-based marginalization method used in the fiducial analysis (blue). While the resulting S_8 – Ω_m constraints for 2×2 pt are very similar, for 3×2 pt the shift and stretch method yields slightly tighter contours, suggesting that it may underestimate the true redshift uncertainties, as by construction the mode-based framework captures a broader range of redshift distribution variations in a data-driven and flexible manner, providing a more accurate and robust marginalization approach for current and future analyses.

in the 2×2 pt analysis. For the 3×2 pt case (right panel), the constraints from SOMPZ and SOMPZ + WZ are nearly identical. We note that, given the already good calibration of the lens $n(z)$, the cosmological constraints are not very responsive to variations in the lens $n(z)$, rendering the additional WZ information redundant relative to the SOMPZ-only case. The corresponding marginalized constraint values for each case are listed in Table V.

We also note that the posteriors are not centered exactly on the true cosmology of the simulation. This is expected, and likely reflects projection effects in the inference process. In a high-dimensional parameter space, even when the true model lies within the allowed region, marginalized posteriors in low-dimensional projections (such as the (Ω_m, S_8) plane) may appear offset from the truth.

These results illustrate both the importance of properly treating redshift uncertainties and the utility of combining complementary redshift calibration methods. Including WZ information not only improves the fidelity of the redshift model, but also helps recover constraining power that would otherwise be lost through marginalization.

As an additional test, we verify the robustness of the compressed likelihood pipeline to inconsistencies between the redshift distributions used to generate the theory data vector and the redshift modes used for compression. We perform a set of internal validation tests.

We consider three configurations:

- (1) The baseline configuration, where both the theory data vector and the compression modes are computed using the SOMPZ + WZ redshift distributions.

TABLE V. Marginalized constraints on the cosmological parameters σ_8 , S_8 , and Ω_m for the simulated analysis described in Secs. VI B and VI A. For each parameter we report the mean of the posterior and the 68 percent confidence level.

Analysis	Method	Marginalization	σ_8	S_8	Ω_m
3×2 pt	SOMPZ + WZ	Fixed	0.732 ± 0.024	0.763 ± 0.009	0.326 ± 0.017
3×2 pt	SOMPZ + WZ	Shift and stretch	0.726 ± 0.026	0.762 ± 0.010	0.331 ± 0.019
3×2 pt	SOMPZ	Modes	0.726 ± 0.031	0.762 ± 0.011	0.332 ± 0.023
3×2 pt	SOMPZ + WZ	Modes	0.726 ± 0.030	0.762 ± 0.011	0.331 ± 0.022
2×2 pt	SOMPZ + WZ	Fixed	0.719 ± 0.031	0.758 ± 0.015	0.334 ± 0.020
2×2 pt	SOMPZ + WZ	Shift and stretch	0.711 ± 0.032	0.753 ± 0.017	0.338 ± 0.022
2×2 pt	SOMPZ	Modes	0.714 ± 0.038	0.756 ± 0.018	0.338 ± 0.026
2×2 pt	SOMPZ + WZ	Modes	0.713 ± 0.033	0.756 ± 0.017	0.339 ± 0.022




FIG. 13. Impact of photo- z uncertainty modeling on cosmological parameter constraints (Ω_m , σ_8 , and S_8) in 2×2 pt analysis. The different contours show the effect of fixing the $n(z)$ for both lenses and sources, fixing $n(z)$ only for lenses while marginalizing over SOMPZ + WZ modes for sources, using SOMPZ-only modes for lenses, and marginalizing over SOMPZ + WZ modes for both lenses and sources. Increasing marginalization freedom shifts and broadens the posteriors, highlighting the impact of redshift uncertainties on parameter constraints.

- (2) A test configuration where the theory data vector is computed using SOMPZ redshift distributions, while the compression modes are derived from SOMPZ + WZ.
 - (i) A second test configuration where the theory data vector uses SOMPZ + WZ, while the compression modes are derived from SOMPZ.

Figure 14 compares the Ω_m – S_8 constraints obtained with three different choices for the lens redshift distribution used in the data vector and in the theory predictions. The fiducial case, in which both the compressed data vector and the model employ the SOMPZ + WZ redshifts, delivers the tightest contours. If the data vector is built with SOMPZ redshifts while the theory still uses SOMPZ + WZ $n(z)$, the posteriors broaden modestly, with a distance of 0.14σ between the two contours, indicating that a less informative compression basis degrades the cosmological precision. A similar, though slightly smaller, widening appears when the roles are reversed—i.e., the data vector retains SOMPZ + WZ while the theory adopts SOMPZ $n(z)$ —with the corresponding distance reduced to 0.008σ .

We have also verified that the nuisance parameters most directly coupled to the lens $n(z)$ —the galaxy bias and magnification slopes—remain stable across these tests. The corresponding posteriors are shown in Appendix D.

VII. DISCUSSION

Photometric redshift (photo- z) estimation for lens galaxies in the DES 3×2 pt analysis has undergone

substantial refinement over the course of the survey. In the SV and Y1 analyses, the lens galaxy sample was based on redMaGiC [68], a selection of red-sequence galaxies with tightly controlled photo- z estimates. The redMaGiC algorithm provides high-precision photometric redshifts, calibrated using spectroscopic samples and characterized by a typical scatter of $\sigma_z/(1+z) \sim 0.015$, making it well suited for early cosmological analyses.

In DES Y3, a new sample of lens galaxies known as MagLim was introduced [9], defined by magnitude and optimized for cosmological signal to noise in 2×2 pt and 3×2 pt analyses. For this sample, the fiducial photo- z estimation was performed using the DNF algorithm [17], with redshift distributions constructed for each tomographic bin. To account for biases in the DNF estimates, the redshift distributions were empirically corrected using clustering redshift (wz) measurements, by applying a shift and stretch in each bin to align the DNF $n(z)$ with the wz estimates [18]. These corrected distributions, along with Gaussian priors on the shift and stretch parameters, were propagated through the cosmological inference pipeline. A similar approach was taken for the DES Y6 BAO analysis [69], where the DNF $n(z)$ served as the baseline shape, and was shifted and stretched to match wz (at $z_{\text{ph}} < 1$) or VIPERS (at higher z) [70].

In Y3, sources and lenses photo- z calibrations approaches were different. In fact, the SOMPZ + WZ methodology had been developed and validated specifically for the source galaxies only [49,50]. Motivated by its success in the source sample, this approach was applied for




FIG. 14. Impact of mismatches between the redshift distributions used in the theory data vector and the redshift modes used for data compression on cosmological constraints. We show the Ω_m - S_8 posterior contours for three runs: using SOMPZ redshift distributions with SOMPZ + WZ compression modes (red), using SOMPZ + WZ redshift distributions with SOMPZ modes (green), and using SOMPZ + WZ redshift distributions with matched SOMPZ + WZ modes (blue), which represents our fiducial analysis. We find that the cosmological constraints are insensitive to moderate mismatches between the redshift model used to generate the theory vector and the one used to construct the compression basis. In particular, the combination of SOMPZ redshifts with SOMPZ + WZ modes yields results consistent with the fiducial run, confirming the robustness of the method.

the first time to a lens galaxy sample in a dedicated reanalysis of the MagLim redshift distributions [19]. We found that using SOMPZ + WZ instead of the fiducial DNF + WZ-corrected redshifts resulted in a shift of approximately 0.4σ in the Ω_m - S_8 parameter plane, indicating the importance of more robust redshift calibration. Given that the SOMPZ + WZ method provides a more comprehensive and flexible estimation of $n(z)$ with a clearer uncertainty model, it was adopted as the fiducial redshift calibration approach for the DES Y6 cosmological analysis.

In this paper, we calibrate the redshift distributions and associated uncertainties for the Y6 lens galaxy sample, which enter the modeling of galaxy clustering and galaxy-galaxy lensing observables and represent a key input to the DES Y6 3×2 pt cosmological analysis ([7]); we also evaluate their impact on parameter constraints using simulated likelihood analyses. In Y6, there have been even further refinements to both the upstream data products, and the methodology, which together contribute to better redshift calibration. In terms of improved data samples,

- (i) the Y6 imaging is about twice as deep as in Y3, resulting in lower magnitude errors. For the bright MagLim lenses ($i < 22.2$) this has only a minor effect on the photo- z calibration, but it provides more precise color measurements and contributes to the overall robustness of the redshift inference;
- (ii) there have been many advances in the curation of the MagLim++ sample described in [6], including less contamination from stars and QSOs, better control of photo- z outliers, and better masking [26];
- (iii) the number of objects in the Balrog simulations generated from synthetic source injection is almost five times larger than Year 3 and now covers the entire footprint. This translates into a more precise and less biased transfer function between the deep and the wide fields; and
- (iv) significant improvements to the redshift calibration sample were achieved by prioritizing the most accurate redshift per object and restricting to deep-field regions with complete coverage, minimizing selection biases.

In addition, this work incorporates significant advances in the methodology over the previous version from Y3:

- (i) an improved SOM algorithm is used for redshift calibration [22], with better handling of photometric uncertainties, nonperiodic boundaries, and g -band inclusion—leading to a more realistic and robust color space representation;
- (ii) we introduced a more comprehensive framework to estimate redshift sample uncertainties, capturing residual biases and catalog incompleteness; and
- (iii) a new approach based on [Y6Modes] is taken to parametrize the uncertainty in the redshift distributions that allows for more flexible and general variations in the distributions, rather than relying on simple shift-and-stretch parametrizations.

Figure 11 displays the mode-reconstructed redshift distributions, $n(z)$, that serve as the effective inputs to the cosmological parameter inference. Compared with the redshift distributions of the MagLim sample used in [19], the relative contribution of each source of uncertainty differs significantly. In Y3, sample variance and shot noise were typically the dominant sources of uncertainty on the mean redshift across all bins, with subdominant contributions from redshift sample incompleteness and photometric zero point variations. In Y6, the breakdown is more nuanced: no single source dominates across all bins or metrics, with redshift sample and zero point uncertainties becoming more relevant in Y6, in particular the latter is the dominant uncertainty in bin 6. This more complex structure reflects the improved modeling of systematic effects in Y6, as well as the use of a deeper and more representative redshift sample, resulting in a more complete and accurate propagation of redshift uncertainties across the tomographic bins.

We investigate the impact of the specific parametrization used to model redshift uncertainties by comparing the traditional shift-and-stretch approach to our new framework, which projects uncertainties onto modes that most affect cosmological inference. We find that the resulting cosmological constraints are similar in both the 2×2 pt and 3×2 pt analyses, though the shift-and-stretch method yields slightly tighter constraints in 3×2 pt. This suggests that the shift-and-stretch parametrization may underestimate the true redshift uncertainties. In contrast, the mode-based approach offers a more flexible and principled treatment that is better suited for capturing complex, data-driven variations in the redshift distributions and for extending to joint analyses.

We include these redshift distributions in a DES Y6-like simulated likelihood analysis. We analyze either the 2×2 pt (galaxy clustering + galaxy-galaxy lensing) or 3×2 pt (2×2 pt + cosmic shear) data vectors under Λ CDM cosmology. Fixing everything else to the fiducial analysis choice and only varying the treatment of the lens redshift, we find that overall, the lens redshift uncertainties enlarge the overall constraints on S_8 by 20% for both 2×2 pt and 3×2 pt when including SOMPZ information; including WZ yields tighter constraints on S_8 of 5% for 2×2 pt, while it remains stable for 3×2 pt. The marginalized posterior values are summarized in Table V.

This work is part of a coordinated effort to calibrate redshifts across all galaxy samples used in the DES Y6 3×2 pt cosmological analysis. In addition to the lens redshift calibration presented here, companion papers detail the calibration of source galaxy redshifts using the SOMPZ + WZ methodology [Y6ShearPZ], the derivation of clustering redshift measurements for both source and lens samples [Y6WZ], and the development of the mode-based uncertainty parametrization [Y6Modes]. Together, these studies provide a unified and self-consistent framework for treating redshift uncertainties across the full 3×2 pt data vector.

In summary, the DES Y6 redshift calibration for the MagLim++ lens sample marks a significant advance in the precision and robustness of photometric redshift treatment for cosmological analyses. Through improved data products, a carefully curated redshift sample, and a principled framework that propagates uncertainties directly into cosmological inference, we ensure that redshift errors are accurately modeled and consistently integrated into the analysis. This approach strengthens the reliability of the DES Y6 3×2 pt results and establishes a scalable foundation for redshift calibration in next-generation multiprobe surveys such as the Vera C. Rubin Observatory's Legacy Survey of Space and Time and the Euclid mission.

ACKNOWLEDGMENTS

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 Dark Energy Survey. 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, NSF 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 NSF Cerro Tololo Inter-American Observatory at NSF 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 No. AST-1138766 and No. AST-1536171. The DES participants from Spanish institutions are partially supported by MICINN under Grants No. PID2021-123012, No. PID2021-128989, No. PID2022-141079, No. PID2023-153229NA-I00, No. SEV-2016-0588, No. CEX2020-001058-M, and No. CEX2020-001007-S, some of which include ERDF funds from the European Union. I. F. A. E. is partially funded by the CERCA program of the Generalitat de Catalunya. We acknowledge support from the Brazilian Instituto Nacional de Ciência e Tecnologia (INCT) do e-Universo (CNPq Grant No. 465376/2014-2). Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the

National Aeronautics and Space Administration (80NM0018D0004). This document was prepared by the DES Collaboration using the resources of the Fermi National Accelerator Laboratory (Fermilab), a U.S. Department of Energy, Office of Science, Office of High Energy Physics HEP User Facility. Fermilab is managed by Fermi Forward Discovery Group, LLC, acting under Contract No. 89243024CSC000002. A. Alarcon acknowledges support from the Spanish Ministerio de Ciencia, Innovación y Universidades (MICIU), Project No. PID2023-153229NA-I00. The project that gave rise to these results received the support of a fellowship from “la Caixa” Foundation (ID 100010434). The fellowship code is LCF/BQ/PI23/11970028. Funding for the Sloan Digital Sky Survey IV has been provided by the Alfred P. Sloan Foundation, the U.S. Department of Energy Office of Science, and the Participating Institutions. S.D.S.S. acknowledges support and resources from the Center for High-Performance Computing at the University of Utah. SDSS [71] is managed by the Astrophysical Research Consortium for the Participating Institutions of the SDSS Collaboration including the Brazilian Participation Group, the Carnegie Institution for Science, Carnegie Mellon University, Center for Astrophysics | Harvard & Smithsonian (CfA), the Chilean Participation Group, the French Participation Group, Instituto de Astrofísica de Canarias, The Johns Hopkins University, Kavli Institute for the Physics and Mathematics of the Universe (IPMU)/University of Tokyo, the Korean Participation Group, Lawrence Berkeley National Laboratory, Leibniz Institut für Astrophysik Potsdam (AIP), Max-Planck-Institut für Astronomie (MPIA Heidelberg), Max-Planck-Institut für Astrophysik (MPA Garching), Max-Planck-Institut für Extraterrestrische Physik (MPE), National Astronomical Observatories of China, New Mexico State University, New York University, University of Notre Dame, Observatorio Nacional/MCTI, The Ohio State University, Pennsylvania State University, Shanghai Astronomical Observatory, United Kingdom Participation Group, Universidad Nacional Autónoma de México, University of Arizona, University of Colorado Boulder, University of Oxford, University of Portsmouth, University of Utah, University of Virginia, University of Washington, University of Wisconsin, Vanderbilt University, and Yale University. W. A. acknowledges support from the MICINN Projects No. PID2019-111317 GB-C32, No. PID2022-141079NB-C32 as well as predoctoral program AGAUR-FI ajuts (2024 FI-I 00692) Joan Oró.

All authors contributed to this paper and/or carried out infrastructure work that made this analysis possible. G. G. performed almost all analysis and manuscript preparation. A. A. contributed to analysis and manuscript preparation, and prepared the redshift sample. C. C. contributed with manuscript preparation. G. B. and M. A. T. developed the

modes pipeline. W. A. ran the clustering- z measurement. G. B. also performed the importance sampling. B. Y. and A. A. contributed to analysis interpretation as members of the DES Y6 redshift WG. N. W., J. M., and A. P. contributed to manuscript preparation as collaboration internal reviewers. The remaining authors have made contributions to this paper that include, but are not limited to, the construction of DECam and other aspects of collecting the data; data processing and calibration; catalog creation; developing broadly used methods, codes, and simulations; running the pipelines and validation tests; and promoting the science analysis.

DATA AVAILABILITY

The data that support the findings of this article are openly available [72].

APPENDIX A: SOM

Figure 15 shows the number of galaxies (left) and the number of secure spectroscopic redshifts (right) assigned to each deep SOM cell. These maps illustrate how the deep SOM is populated by the full photometric sample and how well it is sampled by the redshift calibration set. While the spectroscopic coverage is somewhat heterogeneous across the SOM, we find no systematic correlation between $\sigma(z|c)$ and either the photometric or spectroscopic cell occupancy, consistent with the interpretation of redshift scatter provided in Sec. III A. We note that the occupancy of the deep SOM cells is highly uneven, with a small number of cells containing a large fraction of the photometric sample. While this might appear concerning, these high-occupancy cells are typically located in regions of color-magnitude space with well-constrained redshifts and abundant spectroscopic coverage. To further assess the impact of high- $\sigma(z|c)$ cells, Fig. 16 shows the number of galaxies per value of $\sigma(z|c)$, for both the deep and wide SOMs. These histograms allow us to quantify the fraction of galaxies that fall into poorly calibrated SOM regions, defined here as cells with large redshift scatter. As shown in Fig. 16, the vast majority of galaxies in these cells have low $\sigma(z|c)$ values, confirming that their contribution to the redshift distribution is well calibrated. Moreover, the SOMPZ framework naturally accounts for this imbalance by weighting each cell’s contribution according to its population, ensuring that sparsely populated or poorly calibrated regions have minimal impact on the final $n(z)$ estimate.

In the deep sample, we find that the vast majority of galaxies occupy SOM cells with $\sigma(z|c) < 0.1$, and fewer than 1% of galaxies lie in cells with $\sigma(z|c) > 0.2$. This demonstrates that, despite the presence of a small number of cells with high redshift uncertainty, these regions contribute negligibly to the overall redshift distribution. The result reflects the effectiveness of the deep SOM in




FIG. 15. SOM cell occupancy and spectroscopic coverage for the deep sample. Left panel: number of galaxies in each deep SOM cell. Right panel: number of galaxies with secure spectroscopic redshifts in each deep SOM cell. These maps are used to assess the sampling of the SOM in both photometric and spectroscopic space. As discussed in the main text, high redshift scatter is not systematically correlated with low occupancy or poor spec- z coverage.

organizing galaxies according to photometric properties that are predictive of redshift.

In the wide sample, the distribution of $\sigma(z|c)$ is narrower, with a sharper dropoff beyond ~ 0.15 . This may seem counterintuitive given the increased photometric noise in the wide field; however, it is a natural outcome of the fact that redshift estimates for wide galaxies are inherited from their deep counterparts. In particular, the wide sample is dominated by cells where the deep-redshift calibration is precise, and relatively few wide galaxies fall into high-scatter regions of the deep SOM.

Overall, this analysis confirms that both the deep and wide samples are predominantly associated with SOM regions where redshift uncertainties are well controlled. High- $\sigma(z|c)$ cells are rare and contain only a small fraction of galaxies, indicating that the global $n(z)$ calibration is robust against localized failures in redshift estimation.

APPENDIX B: MODES

As described in Sec. V C, we model redshift uncertainties using a basis of orthogonal perturbation modes applied to the fiducial redshift distributions $n(z)$. Each coefficient $u_{i,j}$ modulates the amplitude of the j th mode in tomographic bin i and is treated as a free parameter in the 3×2 pt likelihood.

Figure 17 presents the leading modes $U_i(z)$, which define the dominant directions of variation in $n(z)$ retained by the compression. These modes correspond to cosmologically relevant deformations of the redshift distribution. Since they are derived from the Fisher information matrix of the observables, they are ordered by their impact on the likelihood: the first mode produces the largest change in observables, the second the next largest, and so on. In practice, the leading modes often resemble intuitive variations such as shifts in the mean redshift, changes in the width, or distortions in the tails, although this interpretation depends on the specific sensitivity of the observables used. Their smooth structure and orthogonality under the Fisher metric make them well suited for parametrizing redshift uncertainty in a basis aligned with cosmological sensitivity.

Figure 18 shows the prior and posterior distributions of $u_{i,0}$, meaning the leading mode, for the six tomographic bins. The priors, centered at zero, are derived from the variance of redshift distributions estimated through the SOMPZ + WZ calibration and encode the level of uncertainty in each mode. The posteriors are obtained by sampling over the redshift, cosmological, and nuisance parameters in the likelihood analysis.

We find that the leading mode of bin 2 is well constrained by the simulated data, reflecting sensitivity to




FIG. 16. Number of galaxies per value of redshift scatter $\sigma(z|c)$ for each tomographic bin. The left panel shows the deep SOM, and the right panel shows the wide SOM. Each point corresponds to a SOM cell; the x axis indicates the scatter $\sigma(z|c)$ associated with that cell, while the y axis shows the number of galaxies assigned to it. The deep Balrog sample displays a broader range of scatter values, but high-scatter cells contain very few galaxies. In the wide sample (right panel), the distribution is narrower and peaks at lower values of $\sigma(z|c)$, reflecting the fact that most wide-field galaxies are assigned to well-calibrated cells in the deep SOM. This confirms that redshift uncertainties are well controlled for the majority of galaxies in both samples.

large-scale shifts in the redshift distribution. We are not showing higher-order modes, since they remain completely unconstrained and are dominated by the prior in all bins. This supports our decision to truncate the mode expansion to a small number of dominant modes per bin.

The independent treatment of redshift uncertainties for the MagLim lens bins is well justified by both the binning strategy and empirical tests. The MagLim sample was designed to optimize 2×2 pt constraints, with each tomographic bin carefully selected to maximize cosmological information while maintaining well-separated redshift distributions. Since the binning is explicitly based on galaxy properties such as magnitude and color, any systematic redshift uncertainties are expected to primarily affect individual bins rather than introduce strong correlations across them. This is in contrast to the source redshift uncertainties, where modes are applied globally to capture correlated shifts across tomographic bins.

A key test to validate this assumption is to examine the behavior of the mode parameters u_i which parametrize the

redshift uncertainty within each bin. If the uncertainties were significantly correlated between bins, we would expect to see strong posterior correlations between these parameters. However, as shown in Fig. 18, the posterior distributions of the u_i parameters remain uncorrelated across different tomographic bins, confirming that the redshift uncertainties in each bin behave independently. This result is consistent with the calibration approach used in MagLim++, where redshift uncertainties are sampled independently across three deep fields using a Latin hypercube and then propagated to the final redshift distributions. While this process does not strictly enforce full independence, it ensures that any correlations arising from the calibration procedure remain small. The lack of significant correlations in the posteriors further supports the use of independent nuisance parameters for the MagLim++ lens bins, as incorporating artificial correlations would only add unnecessary complexity without improving the accuracy of the uncertainty marginalization.




FIG. 17. Modes $U_i(z)$ retained in the compression. These modes define the primary directions of variation in $n(z)$, capturing shifts in peak position, width, and other structural changes. Higher-order modes, which contribute minimally to the observable impact, are discarded. Left: SOMPZ only, right: SOMPZ + WZ.

APPENDIX C: IMPACT OF CLIPPING REDSHIFT DISTRIBUTIONS TO ENFORCE POSITIVITY

As discussed in Sec. V C, our lens redshift distributions are modeled as perturbations around a fiducial $n(z)$ using a set of orthogonal modes. Because the reconstructed distributions result from a linear combination of these modes, they can occasionally take on slightly negative values in low-density regions, particularly at the tails of the distribution.

To test whether these small negative excursions affect cosmological inference, we run a version of the 2×2 pt likelihood in which the reconstructed $n(z)$ is clipped to zero wherever it is negative. We then compare the resulting constraints to those obtained using the original, unclipped distributions.

Figure 19 shows that the contours are virtually unchanged between the clipped and unclipped cases, confirming that these small artifacts do not impact the final results.

APPENDIX D: IMPACT ON GALAXY BIAS AND LENS MAGNIFICATION PARAMETERS

In Sec. V B we examined the robustness of cosmological constraints to mismatches between the redshift distributions used to construct the compressed data vector and those used to define the compression basis. Since both galaxy bias and magnification slope parameters are coupled to the lens $n(z)$, we extend that analysis here by inspecting the corresponding posteriors.




FIG. 18. Prior (grey) and posterior (blue) distributions for the amplitudes of the main redshift modes, $u_{i,0}$, for all six tomographic bins. Each $u_{i,0}$ parameter controls the amplitude of the dominant mode of redshift uncertainty in bin i , as derived from the SOMPZ + WZ calibration. The priors reflect this propagated uncertainty, while the posteriors are constrained by the 3×2 pt data. The plot shows that the $u_{1,0}$ parameter (corresponding to bin 2) is noticeably more constrained by the 3×2 pt data compared to the other bins, likely because this bin overlaps more strongly with the regions where the lensing and clustering signals have the highest sensitivity to redshift calibration.




FIG. 19. Impact of clipping the reconstructed lens redshift distributions at $n(z) > 0$ when using a mode-based representation. We compare cosmological constraints from 2×2 pt measurements using lens $n(z)$ reconstructed from modes with and without enforcing positivity. The blue contours correspond to the default case where no clipping is applied, and the orange contours correspond to the case where the reconstructed $n(z)$ is clipped at zero. The constraints are very similar, indicating that the small negative values introduced by the mode-based reconstruction do not impact the final results.




FIG. 20. Posterior constraints on the linear galaxy bias parameters for the three data-vector/compression-basis combinations tested in Sec. VIB. Results are consistent across cases, indicating that moderate mismatches in the lens $n(z)$ treatment do not significantly affect the inference of galaxy bias.

Figure 20 shows the marginalized constraints on the six linear galaxy bias parameters for the three configurations considered in Sec. VIB: DV(SOMPZ)-Basis(SOMPZ + WZ), DV(SOMPZ + WZ)-Basis(SOMPZ), and the fiducial DV(SOMPZ + WZ)-Basis(SOMPZ + WZ). We find that the inferred biases are stable across the three cases, with shifts always well within the prior widths, in line with the values reported in Sec. VIB.

For the magnification parameters, the posteriors in the three configurations exhibit only minimal shifts—considerably smaller than those seen for the galaxy bias parameters. Since the results are nearly identical across cases, we do not include a separate figure.

- [1] D. H. Weinberg, M. J. Mortonson, D. J. Eisenstein, C. Hirata, A. G. Riess, and E. Rozo, *Phys. Rep.* **530**, 87 (2013).
 [2] C. Heymans *et al.*, *Astron. Astrophys.* **646**, A140 (2021).

- [3] DES Collaboration, *Phys. Rev. D* **105**, 023520 (2022).
 [4] S. Sugiyama *et al.*, *Phys. Rev. D* **108**, 123521 (2023).

- [5] A. N. Salcedo, T. Eifler, and P. Behroozi, *Astrophys. J.* **983**, 183 (2025).
- [6] N. Weaverdyck, M. Rodríguez-Monroy *et al.*, [arXiv:2601.14484](#).
- [7] DES Collaboration, [arXiv:2601.14559](#).
- [8] D. Tanoglidis, C. Chang, and J. Frieman, *Mon. Not. R. Astron. Soc.* **491**, 3535 (2020).
- [9] A. Porredon (DES Collaboration), *Phys. Rev. D* **103**, 043503 (2021).
- [10] J. Myles (DES Collaboration), *Mon. Not. R. Astron. Soc.* **519**, 1792 (2023).
- [11] M. Gatti and G. Giannini (DES Collaboration), *Mon. Not. R. Astron. Soc.* **510**, 1223 (2022).
- [12] C. Sánchez and J. Prat (DES Collaboration), *Phys. Rev. D* **105**, 083529 (2022).
- [13] R. Mandelbaum, A. Slosar, T. Baldauf, U. Seljak, C. M. Hirata, R. Nakajima, R. Reyes, and R. E. Smith, *Mon. Not. R. Astron. Soc.* **432**, 1544 (2013).
- [14] A. Leauthaud *et al.*, *Mon. Not. R. Astron. Soc.* **467**, 3024 (2017).
- [15] S. Alam *et al.*, *Mon. Not. R. Astron. Soc.* **470**, 2617 (2017).
- [16] A. Porredon (DES Collaboration), *Phys. Rev. D* **106**, 103530 (2022).
- [17] J. De Vicente, E. Sánchez, and I. Sevilla-Noarbe, *Mon. Not. R. Astron. Soc.* **459**, 3078 (2016).
- [18] R. Cawthon (DES Collaboration), *Mon. Not. R. Astron. Soc.* **513**, 5517 (2022).
- [19] G. Giannini (DES Collaboration), *Mon. Not. R. Astron. Soc.* **527**, 2010 (2024).
- [20] D. Anbajagane *et al.*, *Open J. Astrophys.* **8**, 65 (2025).
- [21] C. Sánchez, M. Raveri, A. Alarcon, and G. M. Bernstein, *Mon. Not. R. Astron. Soc.* **498**, 2984 (2020).
- [22] A. Campos *et al.*, [arXiv:2408.00922](#).
- [23] B. Flaugher (DES Collaboration), *Astron. J.* **150**, 150 (2015).
- [24] K. Bechtol *et al.* (DES Collaboration), *Astrophys. J. Suppl. Ser.* **282**, 62 (2026).
- [25] W. G. Hartley (DES Collaboration), *Mon. Not. R. Astron. Soc.* **509**, 3547 (2022).
- [26] M. Rodríguez-Monroy, N. Weaverdyck *et al.*, [arXiv:2509.07943](#).
- [27] T. Schutt *et al.* (DES Collaboration), *Open J. Astrophys.* **8**, 26 (2025).
- [28] E. F. Schlafly, A. M. Meisner, and G. M. Green, *Astrophys. J. Suppl. Ser.* **240**, 30 (2019).
- [29] H. J. McCracken *et al.*, *Astron. Astrophys.* **544**, A156 (2012).
- [30] J. Gschwend *et al.*, *Astron. Comput.* **25**, 58 (2018).
- [31] S. J. Lilly *et al.*, *Astrophys. J. Suppl. Ser.* **184**, 218 (2009).
- [32] D. C. Masters, D. K. Stern, J. G. Cohen, P. L. Capak, J. D. Rhodes, F. J. Castander, and S. Paltani, *Astrophys. J.* **841**, 111 (2017).
- [33] D. C. Masters, D. K. Stern, J. G. Cohen, P. L. Capak, S. A. Stanford, N. Hernitschek, A. Galametz, I. Davidzon, J. D. Rhodes, D. Sanders, B. Mobasher, F. Castander, K. Pruett, and S. Fotopoulou, *Astrophys. J.* **877**, 81 (2019).
- [34] S. A. Stanford (Euclid Collaboration), *Astrophys. J. Suppl. Ser.* **256**, 9 (2021).
- [35] O. Le Fèvre *et al.*, *Astron. Astrophys.* **559**, A14 (2013).
- [36] M. Scodeggio *et al.*, *Astron. Astrophys.* **609**, A84 (2018).
- [37] DESI Collaboration, *Astron. J.* **168**, 58 (2024).
- [38] J. R. Weaver *et al.*, *Astrophys. J. Suppl. Ser.* **258**, 11 (2022).
- [39] N. Scoville *et al.*, *Astrophys. J. Suppl. Ser.* **172**, 1 (2007).
- [40] A. Alarcon *et al.*, *Mon. Not. R. Astron. Soc.* **501**, 6103 (2021).
- [41] C. Padilla *et al.*, *Astron. J.* **157**, 246 (2019).
- [42] M. Eriksen *et al.*, *Mon. Not. R. Astron. Soc.* **484**, 4200 (2019).
- [43] C. Laigle *et al.*, *Astrophys. J. Suppl. Ser.* **224**, 24 (2016).
- [44] D. Navarro-Gironés *et al.*, *Mon. Not. R. Astron. Soc.* **534**, 1504 (2024).
- [45] W. G. Hartley (DES Collaboration), *Mon. Not. R. Astron. Soc.* **496**, 4769 (2020).
- [46] D. Gruen and F. Brimiouille, *Mon. Not. R. Astron. Soc.* **468**, 769 (2017).
- [47] J. A. Newman and D. Gruen, *Annu. Rev. Astron. Astrophys.* **60**, 363 (2022).
- [48] M. Carrasco Kind and R. J. Brunner, *Mon. Not. R. Astron. Soc.* **438**, 3409 (2014).
- [49] R. Buchs (DES Collaboration), *Mon. Not. R. Astron. Soc.* **489**, 820 (2019).
- [50] J. Myles, A. Alarcon *et al.*, [arXiv:2012.08566](#).
- [51] H. Hildebrandt *et al.*, *Astron. Astrophys.* **523**, A31 (2010).
- [52] R. Beck, C. A. Lin, E. E. O. Ishida, F. Gieseke, R. S. de Souza, M. V. Costa-Duarte, M. W. Hattab, and A. Krone-Martins, *Mon. Not. R. Astron. Soc.* **468**, 4323 (2017).
- [53] S. A. Smee *et al.*, *Astron. J.* **146**, 32 (2013).
- [54] K. S. Dawson *et al.*, *Astron. J.* **145**, 10 (2013).
- [55] K. S. Dawson *et al.*, *Astron. J.* **151**, 44 (2016).
- [56] R. Ahumada *et al.*, *Astrophys. J. Suppl. Ser.* **249**, 3 (2020).
- [57] S. Alam *et al.*, *Mon. Not. R. Astron. Soc.* **504**, 4667 (2021).
- [58] J. P. Cordero (DES Collaboration), *Mon. Not. R. Astron. Soc.* **511**, 2170 (2022).
- [59] D. Sanchez-Cid *et al.*, [arXiv:2601.14859](#).
- [60] J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson, S. Bridle, S. Sehrish, and J. Kowalkowski, *Astron. Comput.* **12**, 45 (2015).
- [61] J. A. Blazek, N. MacCrann, M. A. Troxel, and X. Fang, *Phys. Rev. D* **100**, 103506 (2019).
- [62] M. Yamamoto, M. R. Becker *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **543**, 4156 (2025).
- [63] X. Fang, E. Krause, and T. Eifler, *Mon. Not. R. Astron. Soc.* **497**, 2699 (2020).
- [64] J. U. Lange, *Mon. Not. R. Astron. Soc.* **525**, 3181 (2023).
- [65] B. Yin *et al.* (DES Collaboration), [arXiv:2510.23566](#).
- [66] W. d'Assignies *et al.* (DES Collaboration), [arXiv:2510.23565](#).
- [67] G. Bernstein *et al.* (DES Collaboration), *Astrophys. J.* **1002**, 95 (2026).
- [68] E. Rozo *et al.*, *Mon. Not. R. Astron. Soc.* **461**, 1431 (2016).
- [69] DES Collaboration, *Phys. Rev. D* **110**, 063515 (2024).
- [70] J. Mena-Fernández (DES Collaboration), *Phys. Rev. D* **110**, 063514 (2024).
- [71] The SDSS website is <https://www.sdss.org>.
- [72] DES Collaboration, DES Y6 data products, <https://des.ncsa.illinois.edu/releases> (2026).