

Towards Precision Cosmology on the Largest Observable Scales

by

Jessica L. Muir

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in the University of Michigan
2018

Doctoral Committee:

Professor Dragan Huterer, Chair
Professor August E. Evrard
Professor Jeffrey J. McMahon
Professor Christopher J. Miller
Professor Aaron T. Pierce

Jessica L. Muir

jlmuir@umich.edu

ORCID iD: 0000-0002-7579-770X

©Jessica L. Muir 2018

Acknowledgements

Numerous individuals have contributed to my growth as a scientist during the five years that I have spent as a PhD student at the University of Michigan. They have helped make this thesis possible. What follows is an incomplete and somewhat self-indulgent accounting of my appreciation.

First of all, I would like to thank my advisor Dragan Huterer, for his scientific guidance, for cultivating a collaborative group environment, and for his active support for my continuing career as a cosmologist. Many thanks as well to the rest of my committee members for agreeing to serve on my committee and for sharing their time and expertise.

I would also like to thank my collaborators and group-mates who contributed to the projects described in this thesis, including Noah Weaverdyck, Saroj Adhikari, Franz Elsner, and Gary Bernstein. Also, many thanks to Alessandro Manzotti for being willing to discuss his ISW analysis with me, and to Diego Mollinari for answering my many questions about the Planck large-angle CMB parity asymmetry analysis. Many thanks also to Joe Zuntz and Eric Baxter for getting me up to speed with CosmoSIS and the DES parameter estimation pipeline.

In order to complete the analyses presented in this thesis, I made use of a number of software tools, including an Anaconda [1] installation of Python, as well as the packages SciPy [2], NumPy [3], Matplotlib [4], and Healpy, the Python implementation of Healpix [5]. I made use of Jupyter [6] notebooks for parts of the analyses. Results associated with alignments in the large-angle CMB were derived using the multipole vector code available at <http://www.phys.cwru.edu/projects/mpvectors/>, which was developed with support from the US Department of Energy. Two-point function calculations and cosmological parameter estimations for the DES analyses were done using CosmoSIS [7].

Of course, my life in graduate school has been shaped by many things outside of just research. I am also grateful for the work done by the administrative, IT, maintenance, and janitorial staff associated with the Physics department, which provided me with a well-supported and comfortable environment in which I could focus on science. Thanks the

members of the Society for Women in Physics, especially those who have served with me on SWIP's executive board for the past three years, for being such a supportive and thoughtful community within the department. More generally, I have greatly appreciated having the chance to work with the numerous individuals (in SWIP or otherwise) who have strived to build a sense of community within our department and to ensure that it reflects inclusive values. I have been consistently impressed by their hard work and dedication, and feel lucky to be able to call them colleagues. Thanks also to the Graduate Employees Organization for advocating for graduate student interests on a University-wide scale. Off campus, of course, I am incredibly grateful for the number of friends I made while in Ann Arbor whose company has made my time here richer and more fulfilling.

I also have much appreciation for the many individuals who have influenced my path to this point. Thanks to my master's advisor at the University of Cambridge, Anne-Christine Davis, for being willing to take me on as a one-year MPhil student when I otherwise struggled to find a research group. Though doing a masters in science communication (my other option for that year) certainly would have been fun, doing research with Anne helped me define my research interests, and gave me tools to be more successful as a PhD student. Thanks to my undergraduate research mentors at Michigan State University, Stuart Tessmer, Joey Huston, and Mark Voit, for their guidance and patience as I stumbled my way towards learning what physics was as a field. Many additional thanks to Joey, Debbie Barratt and others who pulled the strings necessary for me to be able stay on at MSU for the year following graduation when I was being treated for Hodgkin's lymphoma. Being able to have a somewhat normal routine and to continue my physics education during that time meant a lot to me, personally and professionally. Going even further back, thanks to whoever it was at the University of Michigan who listed Stephen Hawking's "A Brief History of Time" and Brian Green's "An Elegant Universe" on a brochure given to visiting high school seniors in 2006. Though I ultimately decided to go to MSU for college, I read those books because I saw them listed on that pamphlet, and doing so was one of the main reasons I ended up studying physics in the first place.

Last but certainly not least, the support and encouragement from my family contributed immeasurably to getting me to this point. Thanks to my parents, Don and Shelly, for encouraging my curiosity from an early age, for their unfaltering support of my education, and for driving in from Romeo to hear my outreach talks whenever they could. Thanks also to my sisters, Sarah, and Carolyn, for their continual support and encouragement.

Preface

The analyses described in this thesis reflect research projects I worked on during my time as a PhD student for which I took a leading role.

Most of the discussion in Chapter 3 about ISW reconstruction was previously published in Ref. [8]. The work described in that chapter was continued in a follow-up study, published as Ref. [9]. I have not included that second paper in this thesis because, though I was actively involved in both the analysis and writing, the project was led by Noah Weaverdyck.

Chapter 4 has been adapted from a draft of a paper on CMB anomaly covariances that was coauthored with Saroj Adhikari and Dragan Huterer. That paper was submitted to a journal as Ref. [10].

The work described in Chapter 5 reflects validation tests of a blinding scheme that is currently being used for the DES Year 3 analysis. The study grew out of preliminary work done by Franz Elsner, though I performed all of the analysis presented in this thesis. Some form of the results in that chapter will eventually be published as a DES methods paper, and the discussion at the beginning of the chapter has been adapted from an early draft coauthored by Gary Bernstein and Dragan Huterer.

Table of Contents

Acknowledgements	ii
Preface	iv
List of Tables	vii
List of Figures	viii
List of Appendices	xvii
Abstract	xviii
Chapter 1 Introduction	1
1.1 Context: Cosmology as a data-driven science	1
1.2 Problems addressed in this work	5
1.3 Plan of thesis	7
Chapter 2 Overview of ΛCDM	8
2.1 Basic definitions	9
2.2 Expansion history	10
2.3 Structure growth	14
2.4 LSS observables	18
2.5 Summary	24
Chapter 3 ISW signal reconstruction	26
3.1 Methods	28
3.2 The effect of survey properties	36
3.3 Cross-check with Manzotti and Dodelson (2014)	44

3.4	The effect of mismodeling redshift distributions	46
3.5	The effect of photometric calibration errors	52
3.6	Implications for cosmic alignments	57
3.7	Conclusion	60
Chapter 4 The covariance of large-angle CMB anomalies		63
4.1	Methods	66
4.2	Features Studied	71
4.3	Results: Anomaly covariances	84
4.4	Conclusions	97
Chapter 5 Multi-probe blinding for DES		99
5.1	Method	102
5.2	The DES 3×2 pt analysis	110
5.3	Validation of two-point function blinding for DES: Methods	116
5.4	Validation of two-point function blinding for DES: Results	126
5.5	Conclusions	137
Chapter 6 Closing remarks		141
Appendices		144
Bibliography		164

List of Tables

2.1	Cosmological parameters and other quantities defined in this chapter.	25
4.1	Summary of large angle CMB features studied in this work.	72
5.1	Fiducial parameter values, their prior ranges if they are varied in this work, and a note on when they are varied for our blinding study (which will be described in Section 5.3). In the last column, “shift” means we vary $\Delta\Theta$ in that direction when creating blinding factors, “obs.” means we vary the parameter when we draw Θ_{obs} for “true” cosmologies, and “fit” means we vary that parameter when performing parameter estimation fits. If no range is given for a parameter, it is held fixed. All fiducial values and priors for varied parameters are chosen to match the settings for the DES Y1 3×2 pt analysis.	113
5.2	List of realizations which exceed $\Delta\chi^2 > 30$ for each set of Maxlike fits to blinded data. The actual $\Delta\chi^2$ values for each realization is shown in square brackets.	131
5.3	Cosmological parameters used to generate data vectors and blinding factors or selected realizations. Recall that Θ_{obs} is the set of “true” cosmological parameters used to produce the unblinded data vector, and $\Delta\Theta$ is the parameter shift associated with the blinding transformation. Roughly speaking, we expect the best-fit parameters for unblinded data to be $\Theta_{\text{unbl}} \approx \Theta_{\text{obs}}$ and for blinded data to be $\Theta_{\text{bl}} \approx \Theta_{\text{obs}} + \Delta\Theta$	134
5.4	Comparison between Maxlike and Multinest fits showing χ^2 for simulated data vector realizations with the worst $\Delta\chi^2$ for additive blinding. The differences in the $\Delta\chi^2$ are relative to the unblinded fit in the same row.	136

List of Figures

2.1	Illustration of shear coordinates. Here the blue ellipses represent a pair of galaxies separated by angle θ , and γ_1 and γ_2 are shear coordinates defined γ_1 is parallel to their separation vector. The coordinates γ_t and γ_\times are defined to measure shear components that are parallel to and perpendicular, respectively, to the line connecting the two galaxies.	23
3.1	Flowchart of ISW reconstruction pipeline.	29
3.2	Scatter plot comparing the true (simulated) ISW signal, on the horizontal axis, to the reconstructed ISW signal, on the vertical axis, for a single realization assuming each of five different depths of the survey. Each data point corresponds to one pixel on an NSIDE=32 map. If there was a perfect reconstruction, all points would fall on the dotted line.	36
3.3	Histograms of the correlation between true and reconstructed ISW maps ρ (left panel), and the typical size of residuals relative to that of the true ISW map fluctuations s (right panel). These plots show the results of 10,000 simulations for surveys of various depths, with their dn/dz distributions shown in arbitrary units as an inset in the left plot. The solid and dashed vertical lines show the theoretical expectation value and measured average, respectively, for the statistic in question.	37
3.4	Un-normalized redshift distributions for the six redshift bins studied, with photometric-redshift uncertainty $\sigma(z) = 0.05(1 + z)$. Because these distributions are not yet normalized (they neglect the denominator of Eq. (3.24)), the area under the curves gives an idea of the relative number of galaxies in each bin. The dotted line shows the ISW kernel in arbitrary units.	39
3.5	Histogram of ρ values measured from 10,000 map realizations for selected binning strategies. The inset shows the un-normalized dn/dz distributions for the sets of redshift bins considered.	40

3.6	Theoretical expectation value for ρ computed for different redshift binning strategies and levels of photometric-redshift uncertainties. The colored bars and corresponding numbers on the left side of the plot are schematic labels for how the galaxies are divided into redshift bins. Different colored points show the effect of different photo- z uncertainties. The “X” points with blue horizontal error bars show the mean and standard deviation of ρ extracted from the histograms in Fig. 3.5.	41
3.7	How filtering out angular scales with $\ell > \ell_{\min}$ affects reconstruction of ISW map. The data points show the mean and standard deviation of ρ , the correlation coefficient between true and reconstructed ISW maps, observed in 10,000 realizations, while the line shows the value of $\langle \rho \rangle$ computed analytically. 43	43
3.8	How changing \bar{n} affects reconstruction of the ISW map. The data points show the mean and standard deviation of ρ , the correlation coefficient between true and reconstructed ISW maps, observed in 10,000 realizations. The line shows the value of $\langle \rho \rangle$ computed analytically.	44
3.9	Left panel: Redshift distributions of surveys, chosen to match the LSS surveys studied in Ref. [11]. The DES 2 bin case is produced by merging the two highest redshift bins for the DES 3 bin case, and so except for the region around $z \sim 1$, the blue lines are directly behind the red lines. Right panel: Histogram of ρ found for 10,000 simulations of surveys with redshift distributions shown in the left panel. Values of $\bar{\rho}$ from Ref. [11] are shown by the arrows along the top of the plot. The observed discrepancies are likely due to different amounts of simulated galaxy shot noise.	45
3.10	Impact of mismodeling survey depth on the expected correlation between the true and reconstructed ISW maps $\langle \rho \rangle$ (left panel) and the ratio of the average size of residuals to that of ISW map features $\langle s \rangle$ (right panel). The true value of the parameter z_0 , which controls the depth of the survey, is fixed at 0.7, while the values used for reconstruction are shown on the x -axis. The blue circular points show results from our standard reconstruction pipeline, while the gray diamond-shaped points (directly behind the blue points in the ρ plot) show results when we skip the \bar{b} -fitting step. The y -axis is linear within one tick mark of zero; otherwise, it has logarithmic scaling.	47

3.11	Impact of mismodeling redshift-dependent bias on the expected correlation between the true and reconstructed ISW maps $\langle \rho \rangle$ (left panel), and the typical size of residuals relative to that of ISW map features $\langle s \rangle$ (right panel). The bias is modeled as $b(z) = 1 + b_2(1 + z)^2$ with the true value fixed at $b_2 = 0.5$ and the values used in the ISW estimator shown on the x -axis. Both axes have logarithmic scaling except in regions within one tick mark of zero, where they are linear. The blue circular points show results from our standard reconstruction pipeline, while the gray diamond-shaped points (directly behind the blue points in the ρ plot) show results when we skip the \bar{b} -fitting step.	48
3.12	Impact of mismodeling the fraction x of galaxies subject to catastrophic photo- z errors on the expected correlation between the true and reconstructed ISW maps $\langle \rho \rangle$ (left panel), and the typical size of residuals relative to that of ISW map features $\langle s \rangle$ (right panel). Both axes have logarithmic scaling except in regions within one tick mark of zero, where they are linear. The blue and brown circular points show results from our standard reconstruction pipeline when the true value of x is 0.1 and 0.01 respectively. The gray diamond-shaped points (directly behind the other points in the ρ plot and the blue points in the s plot) show results when we skip the $bar{b}$ -fitting step.	50
3.13	The change in the galaxy angular power spectrum $C_\ell \equiv C_\ell^{\text{gal-gal}}$ in response to (left to right) changes in survey depth, characterized via z_0 in Eq. (3.10); the redshift dependence of bias, modeled by varying the b_2 parameter in $b(z) = 1 + b_2(1 + z)^2$; and the fraction x of galaxies subject to catastrophic photo- z errors. These plots show the ratio of galaxy autopower relative to that of a reference survey.	52
3.14	The effect of photometric calibration errors on reconstruction quality. We show results for the correlation coefficient between true and reconstructed ISW maps (top panel) and for the typical size of map residuals relative to the variance of the true ISW map (bottom panel). The lines show the expectation from theory, considering only additive contributions from calibration errors, while the data points show the mean and standard deviation from 10,000 simulated map realizations. The shaded regions show the current and projected levels of control over residual calibration errors discussed in Section 3.5.1.	56

- 3.15 Exploration of whether raising ℓ_{\min} can mitigate the impact of photometric calibration errors on ISW signal reconstruction. The top panel shows the mean and standard deviation of ρ , the correlation between the true and reconstructed ISW maps, measured from 10,000 simulations. The bottom panel shows the fractional change in ρ relative to the case with no calibration errors. Points for different values of ℓ_{\min} are staggered so that the errors bars are legible; each cluster of three points shares the same value of $\text{var}[c]$ 57
- 3.16 The relationship between the true (x -axis) and reconstructed (y -axis) angular momentum dispersion $(\Delta L)_{2+3}^2$, defined in the text, for the combined quadrupole and octopole in 10,000 randomly generated ISW maps. Results are shown for two alternate survey depths: our fiducial LSS survey with $z_0 = 0.7$ (red points) and $z_0 = 0.3$ (black points), which have correlation coefficients 0.58 and 0.11 respectively. The gray region denotes $(\Delta L)_{2+3}^2$ as high or higher than measured in WMAP and Planck CMB maps, while the diagonal line is where the true and reconstructed values match. See the text for details. . . . 59
- 4.1 Comparison of the different angular power spectrum measurements described in Section 4.1.3 (left) and their corresponding angular correlation functions (right). The collection of grey lines behind them are from the full sky pseudo- C_ℓ measurements of the first 100 synfast simulations. The black dotted line shows theoretical expectation, and the gray dotted lines show the 68% confidence level cosmic-variance errors. 69
- 4.2 Plot of the lower-tail probability p for the parity asymmetry as a function of the largest multipole ℓ_{\max} considered, for various SMICA power spectrum measurement and simulation ensemble combinations. Solid lines show the probabilities for the SMICA map assessed relative to the synfast simulations, and dashed lines show them relative to the FFP simulations. For the QML and full-sky pseudo- C_ℓ SMICA measurements, the simulations are measured using full-sky pseudo- C_ℓ 's. For the cut-sky pseudo- C_ℓ SMICA measurements, cut-sky pseudo- C_ℓ simulation measurements are used. The vertical line denotes the ℓ_{\max} value at which where Planck XVI (I&S) found the most anomalous parity statistic; see text for details. 78

4.3	Summary plot of one-dimensional probability distributions of features associated with large-angle CMB anomalies. Grey histograms show results for simulations, measured on either cut-sky (using the UT78 mask) or full-sky measurements as indicated by grey text in each panel. Vertical lines are measurements of the real sky, with their single-tail probabilities in the corresponding color. Arrows by the p -values indicate our fiducial measurement choices.	83
4.4	Relationships between large-angle CMB features (see Table 4.1 descriptions of these quantities). Gray contours show the 1, 2, and 3σ confidence regions based on measurements from our ensemble of 100,000 synfast simulations, using the same measurement choices that produced the gray histograms in Fig. 4.3. The 1D histograms on the diagonal are the same as in the left column of Fig. 4.3. The marked data points for SMICA measurements and theoretical expectations are equivalent to the vertical lines in Fig. 4.3. Note that the statistics based on the phases of the $a_{\ell m}$, S_{QO} and A_{LV} , do not have the corresponding theoretical expectations because they cannot be computed analytically from C_ℓ	85
4.5	Left: The feature covariance matrix S_{syn} measured from the synfast simulation ensemble. Right: The difference ΔS (given in Eq. (4.24)), between the covariances measured from the FFP and synfast simulations in units of its sampling error $\sigma_{\text{syn}}^{(1000)}$ estimated from sets 1,000 synfast realizations (defined in Eq. (4.25)).	88
4.6	Left: Absolute difference between the covariance matrix measured from 1000 FFP simulations and from 100,000 synfast simulations. Right: Sample variance error bars for covariance matrix entries measured from 100 subsamples of the 100,000 synfast realizations. The right panel of Fig. 4.5 shows the ratio between the two panels of this figure.	88
4.7	Left: Eigenvectors of the anomaly covariance for our fiducial set of 10^5 synfast simulations. Each column is one PC; the one with index 1 points in the direction of the data's maximum variance. The rows indicate their components in the direction of the initial large-angle CMB quantities. Right: Fractional residual variance for these PCs and the eigenvalues associated with each PC as a fraction of the sum of all eigenvalues	91

4.8	PCA results for the FFP simulations and comparison to synfast results. Top left: Eigenvectors of the anomaly covariance for the set of 1000 FFP simulations. These are the principal components. Top right: Differences between the PC's obtained using the FFP simulations and the fiducial Gaussian simulations. Bottom: Absolute and fractional differences between the Fractional Residual Variance (FRV) when studying the anomaly covariance obtained from the FFP rather than synfast simulations.	94
4.9	Summary of feature statistics projected into our PCA basis. The row labels indicate the basis vectors of the PC basis, which are equivalent to the unit-eigenvectors of the synfast (left) and FFP (right) simulations' feature covariance matrix. The gray histograms show the distribution of the components of simulation realizations in the direction of each PC. The red lines show the same projection of our fiducial SMICA map measurements (using the measurement methods whose p -values in Fig. 4.3 are denoted by an arrow). The red numbers in the top right corner of each panel show the percentage of simulations that are more extreme than the corresponding SMICA measurement.	96
5.1	Top: Map from Ref. [201] of the region on the sky covered by the analysis of the DES Year 1 data, shown relative to the outline of the full five-year footprint in black. Bottom: DES Year 1 constraints from Ref. [12], showing constraints on $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$ and Ω_m , in comparison to and combined with Planck CMB constraints.	101
5.2	Cartoon of model and data spaces that we consider when thinking about how to blind an analysis. \mathcal{M} is the space of all viable model parameter sets Θ , which projects onto the data vector space $\mathcal{D}_{\text{Pri}} \subset \mathcal{D}$, where \mathcal{D} is the space of all possible data vectors. \mathcal{M}_{Pre} describe the region in parameter space where experiments have some preconceived expectation for Θ to be, which project onto the \mathcal{D}_{Pre} of data vectors. An effective blinding transformation must have the possibility of moving the observed data vector \mathbf{d} in or out of \mathcal{D}_{Pre} without moving it out of \mathcal{D}_{Pri}	104

5.3	Plot from Ref. [12] of the redshift distributions for lens and source galaxies in the DES Y1 3×2 pt analysis. The vertical colored bands show the nominal redshift range of each bin, while the lines show the estimated redshift distribution. The black lines show the unbinned total distribution. We adopt these same dn/dz functions for our Y3 blinding tests.	110
5.4	Probability distribution for χ^2 values expected for realizations of our data vector, which have 448 degrees of freedom (457 data points – 9 varied parameters). The vertical black line and shaded regions show the mean, and 1, 2, and 3σ regions for the χ^2 distribution, computed for the limit where the χ^2 distribution with many degrees of freedom $k \gg 1$, approaches a Gaussian with mean $\mu = k$ and a variance $\sigma^2 = 2k$	116
5.5	Locations in parameter space of the realizations used for the blinding scheme validation tests. The red circles show Θ_{shift} , the blue circles show Θ_{obs} , the dashed green lines show the location of the reference cosmology, and the dotted gray lines show the bounds on the allowed cosmology shifts used to generate blinding factors. The solid grey contours show the expected constraining power of the DES Year 3 analysis.	120
5.6	Example of effect of blinding on a subset of the simulated DES 3×2 pt data vector for realization 1, which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$. The rest of the data vector is shown in Fig. 5.7. This plot shows the impact of additive (green) and multiplicative (blue) blinding on $\xi_{\pm}(\theta)$ for a simulated data vector (black). The dashed black line shows the simulated data vector without noise added. The orange band around the black line shows the statistical errors projected for DES Year 3 (the square root of the covariance’s diagonal entries), and the vertical gray bands show the scales that are excluded by scale cuts. Each subplot shows the cross correlations between a different pair of redshift bins labeled by numbers in white boxes. The bin label numbers go up in order of increasing redshift.	122
5.7	Example of effect of blinding on a subset of the simulated DES 3×2 pt data vector for realization 1, which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$. This plot displays the $w(\theta)$ and $\gamma_t(\theta)$ components of the same data vector shown in Fig. 5.6.	123
5.8	Differences between the blinded and unblinded data vector components shown in Fig. 5.6, for realization 1 which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$	124

5.9	Differences between the blinded and unblinded data vector components shown in Fig. 5.7, for realization 1 which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$	125
5.10	Upper panel: Histograms of χ^2 values obtained from Maxlike fits to simulated noisy data vectors. The black histogram shows the distribution for the unblinded data, while the blue and green histograms are for additive and multiplicative blinding, respectively. Lower panel: histogram of $\Delta\chi^2$ values, using the same colors as the upper panel. The vertical dashed gray line at $\Delta\chi^2 = 30$ marks the change that will move the realization by 1σ in the probability distribution for χ^2 expected for the 3×2 pt data vector.	128
5.11	Upper panel: Histograms of χ^2 values obtained from Maxlike fits to simulated noisy data vectors. The black histogram shows the distribution for the unblinded data, while the blue and green histograms are for additive and multiplicative blinding, respectively. Lower panel: histogram of $\Delta\chi^2$ values, using the same colors as the upper panel. The vertical dashed gray line at $\Delta\chi^2 = 30$ marks the change that will move the realization by 1σ in the probability distribution for χ^2 expected for the 3×2 pt data vector. The wide axis range is chosen to show the outlier for multiplicative blinding at $\Delta\chi^2 = 369$. Fig. 5.12 reproduces this information with a decreased axis range.	129
5.12	Same data as the bottom panel of Fig. 5.11, but not displaying the high $\Delta\chi^2$ outlier for multiplicative blinding, in order to show more detail along the horizontal axis.	130
5.13	Scatter plots of $\Delta\chi^2$ versus the parameter shifts $\Delta\Theta$ used to produce blinding factors. The top row shows results for noisy data, while the bottom row shows the noiseless case.	132
5.14	Relationships between shift in best fit cosmological parameters due to additive (left plot) and multiplicative (right plot) blinding for noiseless data and the parameter shifts used to generate blinding factors. The colors shows the shift in χ^2 between blinded and unblinded data, with the high and low-ends of the color scale truncated to make the plots easier to read. Solid gray lines show where $\Delta\Theta^{\text{out}} = \Delta\Theta$, and the grey numbers show the correlation coefficients R for the data points in each subplot.	134
5.15	Parameter constraints for realization 1. Dashed black lines show the cosmology at which the unblinded data was simulated. The blinding factors used shifts $\Delta\sigma_8 = 0.074$ and $\Delta w = +0.43$	138

D.1	Parameter constraints for realization 54. Dashed black lines show the cosmology at which the unblinded data was simulated. The blinding factors used shifts $\Delta\sigma_8 = -0.009$ and $\Delta w = +0.49$. The dashed line for b_1 cannot be seen because it is at $b_1 = 1.24$, significantly below where the countour is.	158
D.2	Parameter constraints for realization 54. Dashed black lines show the cosmology at which the unblinded data was simulated. The blinding factors used shifts $\Delta\sigma_8 = +0.033$ and $\Delta w = -0.36$. The sampler is clearly failing for the multiplicatively blinded data vector.	159
D.3	Plot of $\xi_{\pm}(\theta)$ for realization 54. The dotted and solid black lines show the noiseless and noisy unblinded data, respectively. The solid green line shows the additively blinded noisy data vector, and the solid blue line shows the same for multiplicative blinding. The dotted green and blue lines shows the model prediction computed at the Maxlike best fit parameters for their corresponding blinded data vector. Vertical gray bands show scale cuts, and the shaded region around the unblinded data shows the DES Y3 expected statistical errors. Each subplot shows the cross correlations between a different pair of redshift bins labeled by numbers in white boxes. The bin label numbers go up in order of increasing redshift.	160
D.4	Plot of $w(\theta)$ and $\gamma_t(\theta)$ for realization 54, using the same conventions as Fig. 5.6.	161
D.5	Plot of $\xi_{\pm}(\theta)$ for realization 69, using the same conventions as Fig. 5.6. . . .	162
D.6	Plot of $w(\theta)$ and $\gamma_t(\theta)$ for realization 69, using the same conventions as Fig. 5.6.	163

List of Appendices

Appendix A Derivation of LSS C_ℓ's	144
A.1 Projected density	144
A.2 Conversion to Fourier space	146
A.3 Angular cross correlations	147
A.4 The Limber approximation	148
A.5 Window function for the ISW signal	149
Appendix B Calibration error formalism	151
Appendix C Large-noise limit of ISW reconstruction s statistic	154
Appendix D Additional plots for blinding tests	157

Abstract

One of the primary goals of cosmology is to extract information about fundamental physics from observations of the universe on large scales. With the advent of increasingly large cosmological datasets (and correspondingly small statistical uncertainties), future progress in the field will fundamentally be determined by our ability to account for systematic errors, understand model predictions, and protect analyses from sources of bias. This thesis presents three projects in this theme. We first examine the interpretation of large-angle features of the Cosmic Microwave Background (CMB) which are statistically unlikely in Λ CDM. Particularly, we study some of these features might be due to the imprint of large scale structure (LSS) through the integrated Sachs-Wolfe (ISW) effect, testing the reliability of a method for ISW signal reconstruction based on LSS data. Using simulated ISW and LSS maps, we show that direction-dependent calibration errors are by far the most limiting survey systematic and that for current levels of calibration control, reconstructed ISW maps from existing data are almost entirely noise. Thus, current data cannot be reliably used to separate primordial and late-time contributions to CMB features. We additionally use ensembles of simulated CMB maps to perform a comprehensive study of the covariance between eight features associated with commonly studied large-angle anomalies. The latter part of this thesis shifts its focus to LSS data, introducing a new method for blinding the Dark Energy Survey's (DES) combined analysis of correlations between galaxy positions and weak lensing shear. This technique, which works by modifying LSS tracers' two-point correlation functions, will be used to prevent experimenter bias from influencing DES' precision measurements of dark energy. We demonstrate using simulated DES Year 3 data that this method successfully shifts the analysis pipeline's output cosmological parameters while preserving the ability to check for systematic errors. Given this, the technique will be used to blind the DES Year 3 multi-probe cosmology analyses, and as implemented, will be the most sophisticated blinding strategy for a cosmology analysis to date. Together, these three projects represent varied ways characterizing and developing tools for ensuring future tests of Λ CDM are both precise and accurate.

Chapter 1

Introduction

1.1 Context: Cosmology as a data-driven science

The field of cosmology, broadly, encompasses the effort to use our knowledge of physics to describe the large-scale properties of the Universe, and in turn to use observations of those large-scale properties to learn about fundamental physics. Progress in the field has been the result of a repeated process through which, as the precision of measurements has improved, they have revealed surprising results which have required new, often exotic components, in our cosmological model. In turn, at each iteration, theoretical advances have provided the tools to quantitatively explain — or at least parameterize — the new physics and to guide the development of more precise observational probes.

This interplay between theory and observation is exemplified in the two developments which arguably mark the beginning of modern cosmology: the description of gravity in Einstein’s theory of general relativity (GR), first published in 1915 [13], and Edwin Hubble’s discovery, published in 1929, that the universe is expanding [14]. In GR, the dynamics of gravity are governed by Einstein’s equation,

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu}. \tag{1.1}$$

In this equation, c is the speed of light and $G_{\mu\nu}$ is the Einstein tensor. It is constructed out of second derivatives of the metric $g_{\mu\nu}$, and as such describes the behavior of space-time, while the stress-energy tensor $T_{\mu\nu}$ describes the mass and energy contents of that space-time. Edwin Hubble’s discovery was based on his observations of Cepheid variable stars in other galaxies. By using the Leavitt Law [15] to convert measurements of the star’s period of variation to its absolute luminosity, Hubble was able to use the comparison of the stars’

absolute and apparent magnitudes to measure distances d to their host galaxies. He was also able to draw on previous measurements by Vesto Slipher [16–19] of the galaxies’ velocities relative to Earth based on of Doppler redshifting their spectra. (See e.g. Refs. [20, 21] for a historical overview of these measurements.) A galaxy’s velocity determined in this way is given by $v = cz$, where $z = (\lambda_{\text{obs}} - \lambda_{\text{em}})/\lambda_{\text{em}}$ is the fractional shift between the wavelengths at which light from the galaxy is observed (obs) and emitted (em), for $v \ll c$. Hubble found that there was linear relationship between those distances and recession velocities, now known as Hubble’s Law,

$$v = cz = H_0 d. \tag{1.2}$$

Hubble’s Law has a natural explanation within GR: space itself is expanding. In the years since Einstein’s initial publication of Eq. (1.1), Alexander Friedmann [22, 23], Georges Lemaître [24, 25], Howard P. Robertson [26–28], and Arthur Geoffrey Walker [29] independently derived what is now known as the FLRW metric, a solution to Einstein’s equations which satisfies the assumptions of homogeneity (spatial uniformity) and isotropy (the lack of a preferred direction). Given those symmetry requirements (which, to good approximation are obeyed by the properties of the Universe on scales greater than about 10 Mpc), the background evolution of the Universe can be described entirely in terms of the dynamics a time-dependent scale factor $a(t)$. This scale factor relates physical distance d of a galaxy to its distance r in terms of a set of comoving coordinates via $d = a(t)r$. Neglecting small perturbations due to peculiar velocities, r will be constant, so the galaxy’s velocity relative to us will be $v = \dot{a}r$, where the overdot represents a time derivative. Given this, it is straightforward to show that the Hubble constant in Eq. (1.2) can be written as

$$H_0 = \left. \frac{\dot{a}}{a} \right|_{t_0}, \tag{1.3}$$

where t_0 is the present time.

As GR provided a quantitative description for physics behind Hubble’s Law (which will be described in more detail in Chapter 2), Hubble’s observations paved the way for further theoretical advances by providing empirical information about what the Universe is actually like. During the early days of GR, the prevailing wisdom was that Universe should be static at large scales, so Einstein added a term to his equation to ensure that it would admit static cosmological solutions to take the form [30]

$$G_{\mu\nu} + g_{\mu\nu}\Lambda = \frac{8\pi G}{c^4} T_{\mu\nu}, \tag{1.4}$$

where the cosmological constant Λ can be associated with a constant energy density $\rho = \Lambda c^2 / (8\pi G)$. If tuned to a specific value, Λ could balance the gravitational effects which would otherwise cause the Universe to expand or contract, making $\dot{a} = 0$. Thus, when Hubble’s observations revealed something fundamentally new about the Universe — that it was, in fact, expanding — they showed that this balancing act was unnecessary and that the real universe was described by a dynamic $a(t)$.

Cosmology has advanced significantly in the near-century since, with data-driven discoveries continuing to provide indications of what kind of (sometimes exotic or new) physics needs to be included in our cosmological model. Two more recent examples relevant to this thesis include the introduction of inflation and dark energy.

Inflation describes a period of rapid expansion during the Universe’s first fraction of a second, most simply described in a class of “slow-roll” models. In these models, a scalar field called the inflaton is able to generate exponential expansion with $\ln a(t) \propto t$ if its potential energy is greater than its kinetic energy. First postulated to explain the striking uniformity in temperature of the Cosmic Microwave Background (CMB) radiation, as well as the Universe’s flat (Euclidean) geometry [31], inflation additionally predicts that quantum fluctuations in the inflaton will seed small density perturbations which later evolve under the influence of gravity to generate temperature fluctuations in the CMB and later large scale structure (LSS) [32]. Thus, the premise of inflation was bolstered when the spectra of CMB temperature fluctuations were found to be consistent with the predictions of slow-roll inflation.

Dark energy is the name given to the component of the Universe’s mass-energy budget that is causing it to expand at an accelerating rate. This acceleration was discovered by two teams [33, 34] in the late 1990’s, when they used observations of distant Type Ia supernovae to update Hubble’s measurements of the distance-redshift relation. Because the absolute luminosity of Type Ia supernovae can be calibrated based on their light curves, they served as “standard candle” distance indicators analogous to Hubble’s Cepheids. The fact that supernovae are much brighter than Cepheids means they are visible at much greater distances, and so they allowed the teams probe the departures from the linear distance-redshift relation in Eq. (1.2) caused by time evolution of the expansion rate $H(t)$. These measurements revealed that $\dot{H} > 0$, which is contrary to what one would expect in a universe containing only gravitating matter and radiation. This acceleration thus requires new physics, which could take the form a new contribution to the energy density of the universe, which is what we refer to when we say dark energy, or perhaps a modification to GR at cosmological scales. The simplest viable model which can explain the observed expansion history is to reintroduce

the same cosmological constant that appears in Eq. (1.4).

1.1.1 Λ CDM as the standard cosmological model

The revelations described above and others, based on information from a variety of observational probes, including galaxy surveys, measurements of the CMB, and supernova surveys have led to the development of a standard cosmological model known as Λ CDM. Here we will use Λ CDM as a shorthand for a paradigm which has three main components:

- It describes the present-day composition of the Universe as roughly 5% baryonic matter, 25% cold (non-relativistic) dark matter which interacts gravitationally but not (or at least very weakly) through other forces, and 70% dark energy described by a cosmological constant. (The name Λ CDM references this composition.)
- It includes general relativity as a theory of gravity.
- It posits that initial random, nearly Gaussian density fluctuations are seeded by a period of slow-roll inflation.

1.1.2 Precision cosmology: Testing Λ CDM

Though Λ CDM is a powerful tool for modeling the properties of the Universe on large scales, there remain a number of unanswered questions about its components. These include: What is the particle identity of dark matter? What is the physics behind inflation? Is dark energy a cosmological constant, some kind of dynamic field, or a sign that GR needs to be extended? In order to make progress towards answering these, the cosmology community has invested in a number of increasingly precise cosmological measurements. The motivation for this is that by making careful comparisons between observables the predictions of Λ CDM, we can search for indications as to where the simplest version of our cosmological standard breaks down, which could provide clues for how to build a more fundamental description of, for example, the physics of inflation or dark energy.

Thus far, Λ CDM has stood up to rigorous tests from a wide range of observables. The current generation of cosmological experiments include CMB measurements from Planck¹ [35], ACT² [36–38], and SPT³ [39, 40], the JLA⁴ [41] supernovae survey, the spectroscopic galaxy

¹www.esa.int/Our_Activities/Space_Science/Planck

²act.princeton.edu

³pole.uchicago.edu

⁴supernovae.in2p3.fr/sdss_snls_jla/ReadMe.html

survey BOSS⁵ [42,43], and photometric weak lensing surveys KiDS⁶ [44,45] and DES⁷ [12,46]. These surveys have been able to place constraints with precision on the order of $\sim 1\text{--}10\%$ on the values of cosmological parameters, and have already enabled powerful tests of Λ CDM through tests of the consistency of different experiments’ constraints. (See e.g. the comparison of Planck and DES results in Ref. [12].) The statistical power of cosmological data will only improve as the next generation of surveys, such as Simons Observatory⁸, CMB-S4⁹ [47], DESI¹⁰ [48], LSST¹¹ [49], Euclid¹² [50], SKA¹³ [51], and WFIRST¹⁴ [52], come online.

However, as statistical errors on cosmological measurements shrink, systematic errors and modeling uncertainties will become more significant by comparison. For example, models of dark energy which introduce modifications to GR will predict changes to the clustering properties of matter which are most pronounced at small ($\lesssim 10$ Mpc) scales. Detecting these effects is complicated by the fact that predicting the evolution of density fluctuations which cannot be modeled as linear perturbations is difficult, even in Λ CDM. Failing to account for that theoretical uncertainty when checking small-scale LSS clustering information against model predictions could potentially lead to a false detections of deviation from GR. Thus, broadly, in order for us to be able to perform meaningful tests of Λ CDM— and therefore, to probe fundamental physics — at a level commensurate with the growing statistical power of cosmological experiments, we must carefully assess and account for how analysis choices, systematics, and modeling errors affect our analyses. Characterizing and mitigating systematics for cosmological measurements in service of the quest to learn about fundamental physics is the unifying theme of this thesis.

1.2 Problems addressed in this work

The projects described in this thesis are primarily concerned with the analysis of two observables: temperature fluctuations of the Cosmic Microwave Background (CMB) and Large Scale Structure (LSS) mapped using photometric galaxy surveys.

For the most part, observations of the statistical properties of temperature fluctuations

⁵www.sdss3.org/surveys/boss.php

⁶kids.strw.leidenuniv.nl/

⁷www.darkenergysurvey.org

⁸simonsobservatory.org

⁹cmb-s4.org

¹⁰desi.lbl.gov/

¹¹www.lsst.org

¹²www.euclid-ec.org

¹³www.skatelescope.org

¹⁴wfirst.gsfc.nasa.gov

of the CMB have been incredibly consistent with the predictions of Λ CDM, but there are a handful of features at the largest observable scales which have attracted attention as being statistically unlikely in that standard model. Some of these so-called large-angle anomalies could be hinting at new, perhaps isotropy-breaking, physics associated with inflation, but they could also be the result of systematics associated with astrophysical foregrounds or data processing. They could also simply be statistical flukes which have received attention due to *a posteriori* choices of how to measure the CMB sky. Disentangling these possible causes has been a topic of discussion since these large-angle anomalies were first noted in WMAP data.

One line of inquiry in the study of large-angle CMB anomalies has been associated with efforts to separate truly primordial contributions to CMB temperature anisotropies from those imprinted by late-time effects. In particular, the Integrated Sachs-Wolfe (ISW) effect, a process by which the energy of CMB photons is modulated as they pass through evolving LSS potential wells, contributes to temperature fluctuations on large angular scales. In order to investigate whether certain large-angle anomalies — for example, the alignment between the temperature quadrupole and octopole — can be attributed to the ISW effect, several groups have used estimators based on LSS data to try to separate late and early time components of the large-angle CMB anisotropies. If removing ISW signal contributions from the CMB temperature map reduces the significance of an anomaly, the origin of that anomalous feature is unlikely to be associated with primordial, inflationary physics. We ask the question: How reliable is the ISW reconstruction method used to perform that separation? In Chapter 3 present an investigation into the reliability of the ISW signal estimator involved. Using simulated CMB and LSS maps, we study how the accuracy of the reconstructed ISW signal changes as a function of LSS survey properties, and comment on how our findings should inform discussion of large angle anomalies.

Chapter 4 is also concerned with large-angle CMB anomalies. We present a comprehensive study of the expected independence, in the context of Λ CDM, of a set of commonly studied large-angle CMB features associated with anomalies. We do so by numerically measuring their covariance from ensembles of simulated maps. By analyzing the structure of covariances extracted from simulations, we investigate the impact of the data processing pipelines on the relationship between large-angle anomalies and use a principal component analysis as an alternate look at the ways in which the real CMB temperature map is unusual. By performing detailed examination of the predictions of Λ CDM in the space of these *a posteriori*-selected large-angle CMB features, the goal of this project is to to inform future studies which use large-angle anomalies as motivation for developing inflationary models.

The rest of the thesis will focus entirely on LSS survey analysis in the context of the Dark Energy Survey (DES). As the constraining power of LSS surveys, including DES, has improved, it has become increasingly important to protect against experimenter bias by blinding their analyses. DES is an ongoing photometric galaxy survey which, when complete, will analyze galaxy positions and shapes over a region of the sky with an area of 5000 deg^2 . By performing a combined analysis of several observables — galaxy clustering, weak lensing, supernovae, and galaxy clusters — DES will be able to place tight constraints on cosmology, and in particular will be able to perform a powerful test on whether or not dark energy is well described by a cosmological constant. However, the same fact that drives DES’s constraining power, that it uses the combined analysis of multiple probes, makes it difficult to blind the survey’s analysis pipeline in a way that preserves our ability to use the consistency of different probes as a validation test. As an answer to this challenge, we have developed a novel scheme for blinding combined-probe cosmological analyses based on the manipulation of two-point correlation functions for galaxy clustering and shear. This blinding strategy has already been implemented in the analysis pipeline for the DES Year-3 analysis. In Chapter 5 we will introduce the blinding method, and will present the results of tests demonstrating its performance.

1.3 Plan of thesis

The plan of this document is as follows. Chapter 2 will be devoted to providing general background information about the modeling ingredients of Λ CDM and the observables used to test it. In subsequent chapters we will present the background, methods, and results of the projects covered in this thesis. Chapter 3 will discuss ISW signal reconstruction. Then, in Chapter 4 we will present the large-angle anomaly covariance study. Chapter 5 will be entirely concerned with LSS measurements: it will describe the plan for and validation of the DES multi-probe blinding strategy. Finally, in Chapter 6 we will conclude, summarizing the original results presented in this thesis and the prospects for their outlook in future research.

Chapter 2

Overview of Λ CDM

Our description of the Universe within Λ CDM will rely on two key assumptions: that, to good approximation, the properties of the Universe are homogenous and isotropic, and that structures can be well described as perturbations about a uniform background. These properties of homogeneity (lack of preferred location) and isotropy (lack of preferred direction) have been found to be consistent with observations on scales $\gtrsim 10\text{Mpc}$, though they can and continue to be tested experimentally. Similarly, predictions made using perturbation theory are in good agreement with observations of structure evolution.

Keeping these properties in mind, we will begin in Section 2.1 by defining some terms and parameters used to describe the expanding Universe and the matter and energy contained within it. Then in Section 2.2 we will give an overview of how the “zeroth order” background dynamics of the Universe’s background expansion history can be modeled using GR. Section 2.3 will then focus on the description of perturbations to that background, summarizing the findings of linear perturbation theory and how that connects to the observables discussed in subsequent chapters. In Section 2.5 we will summarize by presenting a table of parameter definitions.

As the discussions in each of the following chapters will be relatively self-contained, the goal of this overview will not be to give a comprehensive introduction to the physics of Λ CDM. Rather, our focus will be on defining quantities, summarizing important results and sketching how they are obtained, in order to build a common vocabulary for the later chapters. This discussion draws on the information in a number of textbooks (Refs. [53–56]), as well as sets of lecture notes (Refs. [57–59]).

2.1 Basic definitions

The expansion of the Universe is described in terms of a scale factor $a(t)$, which by convention is equal to one at the present time t_0 . It relates physical distances x (defined so $dx = c dt$, where c is the speed of light) to comoving distances r via $dx = a(t) dr$. As was mentioned in the description of Hubble's Law above, cosmological expansion causes light emitted by distant objects to be redshifted. The redshift is quantified as the fractional difference between the light's wavelength when it is emitted, λ_{em} , and its observed wavelength λ_{obs} :

$$z \equiv \frac{\lambda_{\text{obs}} - \lambda_{\text{em}}}{\lambda_{\text{em}}}. \quad (2.1)$$

The cosmological redshift z can be equivalently interpreted as Doppler reddening due to the time derivative of the source's physical distance, or as the photons' wavelength being stretched to as space expands. It can be related to the scale factor via $a = (1 + z)^{-1}$.

The Hubble parameter $H(t) \equiv \dot{a}/a$ quantifies the logarithmic time derivative of the scale factor, while the Hubble constant H_0 refers to its present-day value, which has been measured to be roughly $H_0 \approx 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Because most cosmological information in measurements of expansion history and structure growth depend on the time dependence of $H(t)$ rather than its normalization, it is common to put the value of H_0 into distance units in terms of the parameter $h = H_0 / [100 \text{ km s}^{-1} \text{ Mpc}^{-1}]$.

Because light travels at a finite velocity, at distances where cosmological redshift is much larger than Doppler contributions from peculiar velocities,

1. a source's redshift z ,
2. the time t_e that the observed light was emitted,
3. the scale factor at that time $a(t_e)$,
4. and the source's comoving distance r

can all be used interchangeably as a radial distance or time coordinates. They can be related to one another via

$$r = \int_0^r dr' = \int_{t_e}^{t_0} \frac{c dt}{a(t)} = \int_0^z \frac{c dz'}{H(z')}. \quad (2.2)$$

Thus, $r = 0$, $z = 0$ and $a = 1$ correspond to nearby sources and the present time t_0 , while large r , large z and small a correspond to distant objects and early times.

The composition of the Universe is described in terms of the densities of different components, with ρ_r , ρ_c , ρ_b , ρ_ν representing the density of radiation, cold dark matter, baryons,

and neutrinos, respectively, and with $\rho_m \equiv \rho_c + \rho_b + \rho_\nu$ representing the total matter density. (These quantities can be converted to energy densities simply by multiply by a factor of c^2 .) The pressure associated with each component is defined as p with the same set of subscripts. The density and pressure associated with dark energy can be defined as ρ_{DE} and p_{DE} . We will use ρ and p with no subscript to refer to the total energy density and pressure.

The various energy densities are commonly written in units of a critical density such that for mass-energy component X ,

$$\Omega_X = \frac{\rho_X}{\rho_{\text{crit}}} \quad \text{for} \quad \rho_{\text{crit}} = \frac{3H^2}{8\pi G}. \quad (2.3)$$

The critical density is defined so that in a Universe with flat geometry (which we will assume throughout this work),

$$1 = \Omega = \sum_X \Omega_X. \quad (2.4)$$

A subscript 0 (as in ρ_{X0} or Ω_{X0}) will be used to denote densities at the present time.

2.2 Expansion history

The background dynamics of the universe are governed by GR, which describes the relationship between the behavior of spacetime and the matter and energy occupying through Einstein's Equation, given above in in Eq. (1.1). Under the assumptions of spatial homogeneity and isotropy, the stress energy tensor components take the form

$$T_{\mu\nu} = \rho u_\mu u_\nu + p(g_{\mu\nu} + c^{-2}u_\mu u_\nu) \quad (2.5)$$

which describes a perfect fluid with energy density ρ , pressure p , and four-velocity u_μ . where u_μ is the four-velocity of the fluid. That fluid describes the behavior of the average *total* energy density and pressure of all components — the sum of contributions from dark matter, baryons, etc. The most general metric $g_{\mu\nu}$ that obeys homogeneity and isotropy is the FLRW metric, which, using the convention where the Minkowski metric is $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$, has the interval

$$ds^2 = g_{\mu\nu}^{\text{FLRW}} dx^\mu dx^\nu = -c^2 dt^2 + a^2(t) \left[\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right], \quad (2.6)$$

where r is the comoving distance in the radial direction, θ and ϕ are angular coordinates, and k a constant describing a uniform curvature. For a flat universe, $k = 0$ and the FLRW

metric becomes

$$g_{\mu\nu}^{\text{FLRW}} = \begin{pmatrix} -c^2 & 0 & 0 & 0 \\ 0 & a^2(t) & 0 & 0 \\ 0 & 0 & a^2(t) & 0 \\ 0 & 0 & 0 & a^2(t) \end{pmatrix}. \quad (2.7)$$

For the remainder of this thesis we will assume flat geometry.

Inserting $g_{\mu\nu}^{\text{FLRW}}$ and $T_{\mu\nu}$ from Eq. (2.5) into Einstein's equation (Eq. (1.4)) causes it to reduce to the two Friedmann equations,

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho + \frac{\Lambda c^2}{3} \quad (2.8)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3}\left(\rho + \frac{3p}{c^2}\right) + \frac{\Lambda c^2}{3}. \quad (2.9)$$

As equations of motion for the scale factor $a(t)$ the Friedmann equations relate the expansion history of the Universe to its mass-energy contents.

We can additionally describe the relationship between the energy density of component X and the scale factor using the continuity equation,

$$\frac{d\rho}{dt} + 3H\left(\rho + \frac{p}{c^2}\right) = 0. \quad (2.10)$$

This can be derived through manipulation of the Friedmann equations, or through energy conservation in general relativity by setting the covariant derivative of $T_{\mu\nu}$ equal to zero. The continuity equation will hold for both the total density of the Universe, as well as for each individual component in the limit where interactions between radiation, matter, and dark energy can be neglected.

We can analytically find the relationship between the physical density ρ_X of mass-energy component X and the scale factor if we define the equation of state parameter $w_X = p_x/(\rho_x c^2)$. Using this definition, we can solve the continuity equation to show

$$\rho_X(a) = \rho_X(1) a^{-3(1+w_X)}. \quad (2.11)$$

Because on the large scales matter can be treated as pressureless, $w_m = 0$. We can see that this gives us the expected relation between matter density and physical volume, $\rho_m \propto a^{-3}$. (That is to say, the number of particles per comoving volume is conserved.) For radiation $w_r = \frac{1}{3}$ gives $\rho_r \propto a^{-4}$. One can show that this is in line predictions based on thermodynamics for the relationship between volume and the energy density of radiation, and can additionally

associate it with a time evolution of the radiation's average blackbody temperature,

$$\bar{T}(a) = a^{-1}\bar{T}_0. \quad (2.12)$$

2.2.1 Dark Energy

By examining Eqs. (2.8) and (2.9), one can show that the cosmological constant Λ can be associated with a density and pressure,

$$\rho_\Lambda \equiv \Lambda/(8\pi G) \quad \text{and} \quad p_\Lambda \equiv -\Lambda c^2/(8\pi G). \quad (2.13)$$

From this, we can infer that the fluid description of the cosmological constant as the equation of state parameter $w_\Lambda = -1$, and so, by Eq. (2.11), ρ_Λ is independent of a . Thus, the physical manifestation of the cosmological constant is a time-independent, spatially-uniform energy density.

It is common to define the parameter w , without a subscript, to be the dark energy equation of state,

$$w \equiv w_{DE} = \frac{p_{DE}}{\rho_{DE}c^2}. \quad (2.14)$$

By leaving w as a free parameter, we can parameterize the extent to which dark energy is well described by a cosmological constant by checking whether observational constraints are consistent with $w = -1$. This slight generalization of Λ CDM is commonly referred to as w CDM.

We noted in Chapter 1 that dark energy is the name given to whatever mass-energy component of the Universe is causing its expansion rate to accelerate, so more generally, we can use the Friedmann equations to place requirements on w . In a universe where dark energy dominates ($\rho \approx \rho_{DE}$), according to Eq. (2.9) the requirement for acceleration $\ddot{a} > 0$ becomes

$$w < -\frac{1}{3}. \quad (2.15)$$

2.2.2 Inflation

As was mentioned in Chapter 1, inflation refers to a period of accelerating expansion that occurred during the Universe's first fraction of a second. Models of inflation were initially developed to help solve the horizon problem. Stated roughly, the horizon problem is the question of why the Universe on large scales is observed to be so uniform, with fluctuations in CMB temperature smaller than one part in 10^5 , when, given our knowledge of the expansion

history of the Universe accounted for in the Friedmann equations above, patches of the CMB farther than about two degrees from one another on the sky have never been in causal contact.

We can demonstrate why an early period of accelerating expansion can solve the horizon problem if we examine the expression for the particle horizon $r_p(t)$, the greatest comoving distance a photon could have traveled by time t : [58, 59]

$$r_p = \int_{t_i=0}^t \frac{c dt'}{a'} = c \int_{a_i}^{a(t)} \frac{d \ln a}{aH}. \quad (2.16)$$

Referring to the first Friedmann equation (Eq. (2.8)), given the mass-energy components of the Universe that we measure today,

$$H(a) = H_0 \sqrt{\Omega_{r0} a^{-4} + \Omega_{m0} a^{-3} + \Omega_{\Lambda 0}}. \quad (2.17)$$

Let us examine the behavior of the integrand in Eq. (2.16) at early times. In the limit of small a , Eq. (2.17) becomes $H(a) \approx H_0 \sqrt{\Omega_{r0} a^{-4}} \propto a^{-2}$. This means that the integrand $(aH)^{-1} \propto a$, which goes to zero at early times. Because of this, the integral in Eq. (2.16) will be dominated by late-time contributions, and $r_p \approx (aH)^{-1}$. (This can be shown explicitly for a Universe containing a single fluid component whose density evolution is governed by Eq. (2.11).)

The incredible uniformity of the Universe on large scales could be explained if all points we observe have been in causal contact with one another at some point in the past. This will be the case if we introduce new physics that causes the integral of Eq. (2.16) to receive large contributions from early times. In other words, if $(aH)^{-1}$, were large at early times and shrank, before more familiar physics expansion caused it to grow again according to Eq. (2.17), the horizon problem will be solved. It can be shown with a few lines of calculations that requiring a period where $\frac{d}{dt}[(aH)^{-1}] < 0$ is equivalent to having a period of accelerated expansion, with $\ddot{a} > 0$.

For a model of inflation to solve the horizon problem it must both predict an early period of early accelerated expansion, and predict that the period lasts long enough to put the entire observable Universe in causal contact. Slow-roll inflation is the class of simple models which achieve this through the dynamics of a scalar inflaton field ϕ . It can be shown that

for a scalar with potential $V(\phi)$, the field's energy density and pressure are

$$\rho_\phi c^2 = \frac{1}{2\hbar c} \dot{\phi}^2 + V(\phi) \quad (2.18)$$

$$p_\phi = \frac{1}{2\hbar c} \dot{\phi}^2 - V(\phi), \quad (2.19)$$

where \hbar is the reduced Planck's constant and ϕ is defined to have units of mass. We can use these expressions to compute the inflaton's equation of state, and referencing Eq. (2.15), can show that $\ddot{a} > 0$ if

$$\frac{1}{\hbar c} \dot{\phi}^2 < V(\phi). \quad (2.20)$$

Similarly, we can obtain the equation of motion for ϕ by inserting Eqs. (2.18) and (2.19) into the Friedmann equations,

$$0 = \ddot{\phi} + 3H\dot{\phi} + \hbar c^3 V', \quad (2.21)$$

where the prime on V indicates a partial derivative with respect to ϕ . By manipulating those equations we can show that inflation will occur, and last sufficiently long, if $\dot{\phi}^2$ and $\ddot{\phi}$ are small relative to other scales in the problem. These requirements can be translated into conditions on the inflaton potential $V(\phi)$ through the slow roll parameters

$$\epsilon \equiv \frac{M_{pl}^2}{2} \left(\frac{V'}{V} \right)^2 \quad \text{and} \quad \eta \equiv M_{pl}^2 \frac{|V''|}{V}, \quad (2.22)$$

where $M_{pl} = \sqrt{\hbar c / (8\pi G)}$ is the reduced Planck energy. Inflation will occur if $\epsilon \ll 1$ and $|\eta| \ll 1$. Slow roll inflation refers to models satisfying these requirements.

One can additionally show, given an inflationary model, that quantum fluctuation in the inflaton field during inflation will seed small fluctuations in curvature, which in turn will seed adiabatic density fluctuations. Adiabatic density fluctuations (with $\delta\rho_r \propto \delta\rho_b \propto \delta\rho_m$) are thus a generic prediction of inflationary models. The form of $V(\phi)$ will determine the statistical properties of those density fluctuations. In the slow roll limit they will be generically close to Gaussian and will predict a fluctuation power spectrum of a particular form, which will be defined quantitatively in Eq. (2.29) below.

2.3 Structure growth

A sketch of the history of structure in the Universe proceeds as follows. During inflation, quantum fluctuations in the inflaton field seeded small fluctuations in density. When density

perturbations on a given length scale enter the horizon, they begin to evolve under the influence of gravity. Generally gravity causes structure in dark matter to continually grow — that is, the spatial distribution of ρ_c became more clumpy — with its growth rate determined by the matter density and the Universe’s expansion rate. In the early Universe, baryonic matter was ionized and coupled to photons, so the gravitational pull into overdense regions was counteracted by the baryon-photon plasma’s radiation pressure, causing the plasma to undergo oscillations known as Baryon Acoustic Oscillations (BAO). When the Universe first cooled enough to allow stable atoms to form, at around 300,000 years, the Universe became neutral in a process known as recombination. The subsequent decoupling of the photons and baryons released the radiation that we observe today as the CMB. Density fluctuations in the primordial baryon-photon fluid were imprinted in the variations of the blackbody temperature of CMB photons. After recombination, baryons fell into dark matter potential wells, and as density contrasts continued to grow over time, they eventually formed galaxies in high density regions. Thus, given a model which describes

1. the initial fluctuation spectrum,
2. expansion history,
3. gravitational effects,
4. and the interactions between radiation, baryons, and dark matter,

we can predict the statistical properties of both CMB temperature fluctuations and late-time density inhomogeneities.

2.3.1 Power spectra

We generally describe the properties and evolution of structure in our Universe in terms of two point correlation functions. These correlation functions are the most relevant quantities for comparing the predications of models to observations because in Λ CDM we expect fluctuations in CMB temperature and matter density (at least at large scales) to be Gaussian or nearly Gaussian. For a Gaussian field, its two point functions contain all of the information its statistical properties of the field.

Structure in the map of CMB blackbody temperature $T(\hat{\mathbf{n}})$, where $\hat{\mathbf{n}}$ indicates the direction on the sky, can be quantified in terms of its angular correlation function

$$C(\theta) = \langle T(\hat{\mathbf{n}}_1)T(\hat{\mathbf{n}}_2) \rangle \quad \text{where } \hat{\mathbf{n}}_1 \cdot \hat{\mathbf{n}}_2 = \cos \theta. \quad (2.23)$$

Equivalently, one can decompose the CMB temperature map in spherical harmonics,

$$T(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \quad (2.24)$$

and define the angular power spectrum, which for isotropic fluctuations takes the form

$$\langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell\ell'} \delta_{mm'} C_{\ell}. \quad (2.25)$$

The Fourier space power spectrum (C_{ℓ}) is more directly related to the underlying physics than the real space correlations ($C(\theta)$) because the solutions of the Boltzmann equations that must be solved to compute it are independent for each multipole. Given this, the power spectrum C_{ℓ} is most commonly the focus of efforts to compare the observed CMB to the predictions of Λ CDM or alternative models.

In describing the distribution of matter at late times, we instead use the spectrum $P(k, z)$ of matter density fluctuations. It is defined in terms of perturbations δ to the average matter density $\bar{\rho}_m$, referenced by the field

$$\delta(\mathbf{r}, t) = \frac{\rho_m(\mathbf{r}, t) - \bar{\rho}_m(t)}{\bar{\rho}_m(t)}. \quad (2.26)$$

The Fourier transform of the density perturbations is,

$$\delta(\mathbf{k}, t) = \int \frac{d^3\mathbf{r}}{(2\pi)^{3/2}} \delta(\mathbf{r}, t) e^{i\mathbf{k}\cdot\mathbf{r}}, \quad (2.27)$$

and, assuming the statistical properties of δ are homogeneous and isotropic, the matter power spectrum is defined

$$\langle \delta(\mathbf{k}, t) \delta^*(\mathbf{k}', t) \rangle = (2\pi)^3 \delta^{(3)}(\mathbf{k} - \mathbf{k}') P(k, t). \quad (2.28)$$

Here, $\delta^{(3)}$ is the Dirac delta function and $k \equiv |\mathbf{k}|$.

Slow roll inflation predicts an initial power spectrum of the form

$$P(k) = A_s \left(\frac{k}{k_*} \right)^{n_s}, \quad (2.29)$$

where the A_s is the amplitude of fluctuations at pivot scale k_* and the spectral index $n_s = 1 - 6\epsilon + 2\eta$ is slightly smaller than 1. (The slow-roll parameters ϵ and η are defined

above in Eq. (2.22).) The form of the matter power spectrum will be significantly altered, particularly on small scales (high k) due to the physics which govern structure growth. Some of the characteristics of these alterations can be derived analytically, but in general the full calculation is fairly involved and requires numerically solving coupled Boltzmann equations.

In practice, predictions for C_ℓ and $P(k, z)$ for Λ CDM or alternative models can be computed using a Boltzmann code such as CAMB.¹ [60, 61].

2.3.2 Linear evolution of LSS

Given the matter power spectrum at some time t , say, obtained using a model of slow roll inflation and CAMB, we can use linear perturbation theory to describe its evolution analytically. This works on scales where density fluctuations are small compared to the average density, or $\delta \ll 1$. One can combine the Poisson equation, the Euler equation describing the conservation of momentum for fluids, and the continuity equation for matter density to obtain the linear growth equation,

$$\ddot{\delta} + 2H\dot{\delta} - 4\pi\bar{\rho}_m(t)\delta = 0. \quad (2.30)$$

With some manipulation, by drawing on our knowledge from Section 2.2 to determine the time evolution of H , we can show that

$$\delta(t) \propto \begin{cases} \ln t & \text{when } \rho_r \gg \rho_m, \rho_\Lambda \\ a(t) & \text{when } \rho_m \gg \rho_r, \rho_\Lambda \\ \text{const.} & \text{when } \rho_\Lambda \gg \rho_m, \rho_r \end{cases}. \quad (2.31)$$

Given this scale-independent behavior on scales where the dynamics of density fluctuations can be described using linear perturbation theory, it is common to define the linear growth function $D(a)$ and growth rate $f(a)$ as

$$D(t) \equiv \frac{\delta(\mathbf{r}, t)}{\delta(\mathbf{r}, t_0)} \quad (2.32)$$

$$f(t) \equiv \frac{d \ln D}{d \ln a}. \quad (2.33)$$

Solving Eq. (2.30) will determine these functions.

In practice, to compute the power spectrum, we would use CAMB or similar software to

¹<http://camb.info>

determine the form of $P(k, z_0)$ at some redshift z_0 , then evolve it linearly via

$$P(k, z) = \frac{D^2(z)}{D^2(z_0)} P(k, z_0). \quad (2.34)$$

Here and later when z appears as an argument (as in e.g. $P(k, z)$ or $\delta(\mathbf{r}, z)$) it is meant to indicate that the function is evaluated at the time $t(z)$. Once one moves beyond linear perturbation theory, the calculations for the evolution of structure becomes significantly more involved. We can estimate the effect of non-linear growth on $P(k, z)$ using semi-analytic models which are fit to N-body simulations, or emulation via simulations. A commonly used example of this is the halofit model described in Ref. [62]. Additional corrections to the power spectrum due to baryonic physics [63] or modifications to GR [64] may also be applied using semi-analytic models which have been similarly calibrated on simulations.

In order to determine on which length scales linear perturbation theory is expected to be accurate, one can refer to the quantity $\sigma^2(R)$, the variance of the mass inside a sphere of comoving radius R . It can be written in terms of the matter power spectrum via

$$\sigma^2(R, z) = \frac{1}{2\pi^2} \int_{-\infty}^{\infty} k^2 dk P(k, z) \left(\frac{3j_1(kR)}{kR} \right)^2 \quad (2.35)$$

where $j_\ell(x)$ is a spherical Bessel function and the quantity inside the large parentheses is the Fourier-space window function for a spherical top-hat filter. In other words, $\sigma(R)$ provides a measure of how clumpy matter in the Universe is on length scale R . If $\sigma(R) \ll 1$, the density distribution is relatively smooth and linear perturbation theory should do a good job at describing structure growth. Otherwise, nonlinear corrections may be needed. The quantity

$$\sigma_8 \equiv \sigma(R = 8h^{-1}\text{Mpc}, z = 0) \quad (2.36)$$

has been measured to be $\sim \mathcal{O}(1)$ and is commonly used as a parameter associated with amplitude of the density perturbation spectrum, analogous to A_s (see Eq. (2.29)).

2.4 LSS observables

Calculations within Λ CDM or its extensions produce predictions for $P(k, z)$, the power spectrum for matter fluctuations. However, when we study LSS observationally, we must rely on measurements of tracers of the total matter distribution, rather than observing the total matter distribution, which is dominated by fluctuations in dark matter, directly. The

LSS tracers that this thesis will focus on are galaxy positions and weak gravitational lensing.

2.4.1 Galaxy clustering

Because galaxies form in regions of high density, they trace the peaks of the total density distribution. We describe the relationship between the fluctuations in galaxy number density ρ_g (with perturbation field $\delta_g(\mathbf{r}, t)$ and its Fourier transform $\delta_g(\mathbf{k}, t)$ defined in the same way as for δ) and those in total matter density in terms of a bias $b(z, k)$, defined so

$$\delta_g(\mathbf{k}, z) = b(z, k) \delta(\mathbf{k}, z). \quad (2.37)$$

This means that the relationship between the galaxy power spectrum and the matter power spectrum is

$$P_g(k, z) = b^2(z, k) P(k, z). \quad (2.38)$$

On linear scales, and for Gaussian random primordial fluctuations, bias is scale independent, so it can be written solely as a function of z .

When we measure the distribution of galaxies using a LSS survey, it is relatively easy to determine their position on the sky to high accuracy. Finding the radial component of their three dimensional positions is more challenging. This can be determined by measuring a galaxy's redshift z from its spectrum, and simply taking care to make the corrections to $P_g(k, z)$ to account for peculiar velocities associated with structure formation. This is the strategy taken by spectroscopic surveys like BOSS. However, collecting enough photons to obtain an accurate spectrum for a given galaxy takes a significant amount of time. In order to be able to quickly survey galaxies over a very large volume, photometric surveys (like DES) use an alternative strategy to obtain three-dimensional information about galaxy populations.

In photometric galaxy surveys, galaxies are imaged using a small number of color filters. Then, an algorithm, based on the analysis of a training set of galaxies for which there are both spectroscopic and photometric data, is used to estimate the redshift of each galaxy based on its relative flux in the different filters. When determined using this method, the redshift of any individual galaxy will be poorly determined, but the distribution along the line of sight of galaxies in a tomographic redshift bin can be estimated with reasonable accuracy. Therefore, in order to compare the predictions of Λ CDM to observations, we need to compute the angular power spectra of projected number density for galaxies in different tomographic bins.

If dn^i/dz is the redshift distribution of the direction-averaged angular density of galaxies associated with tomographic bin i , normalized so that

$$1 = \int dz \frac{dn^i}{dz}, \quad (2.39)$$

then the fluctuations in projected number density $\delta_n^i(\hat{\mathbf{n}})$ can be written (assuming linear bias) as

$$\delta_g^i(\hat{\mathbf{n}}) = \int dz \frac{dn^i}{dz} b(z) \delta(\hat{\mathbf{n}} r(z), z). \quad (2.40)$$

The angular power for the correlation between δ_g^i and δ_g^j — that is, the fluctuations of projected galaxy number density in redshift bins i and j — using the Limber approximation (see Section A.4 of Appendix A) by

$$C_{gg}^{ij}(\ell) = \int dz \frac{H(z)}{c r^2(z)} W_g^i(z, k) W_g^j(z, k) P(k, z) \Big|_{k=(\ell+\frac{1}{2})/r(z)}. \quad (2.41)$$

Here

$$W_g^i(z, k) \equiv \frac{dn^i}{dz} b^i(z, k) \quad (2.42)$$

is a weight function which translates between the fluctuation in the density of tracers and the fluctuations of total matter density. These expressions are derived in Appendix A in both its exact and Limber approximation forms. In practice, galaxy correlations are typically measured from data in terms of their real space correlations,

$$w^{ij}(\theta) = \sum_{\ell} \frac{2\ell+1}{4\pi} P_{\ell}(\cos \theta) C_g^{ij}(\ell), \quad (2.43)$$

where $P_{\ell}(x)$ is a Legendre polynomial of order ℓ . The two-point function $w^{ij}(\theta)$ measures the excess probability compared to random of finding a pair of galaxies separated by angle θ .

2.4.2 Weak lensing

Weak gravitational lensing measurements extract information about LSS by measuring the correlations in galaxy shapes due to the deflection of light from distant galaxies by lensing due to intervening LSS. As photons from distant galaxies travel through the Universe, their path will be deflected by the potential fluctuations associated with LSS along the line of sight. This causes shears, or distortions, to their apparent shapes. Because we do not know

source galaxies' intrinsic shapes, we cannot determine the shear of a given galaxy explicitly, but the fact that source galaxies that are closer together will be sheared by similar lower redshift structures allows us to use statistical correlations between galaxy shapes in order to measure the influence of the total projected mass density.

We can quantify the amount of distortion due to gravitational lensing in terms of a lensing convergence $\kappa(\hat{\mathbf{n}})$. The convergence is defined in terms of the Laplacian with respect to angle on the sky of a lensing potential ψ ,

$$\kappa = \frac{1}{2} \nabla_{\theta}^2 \psi, \quad (2.44)$$

where ψ is defined so that its gradient $\nabla_{\theta} \psi = \theta_s - \theta_o$ is a two-dimensional vector describing the angular deflection between a source's observed position on the sky θ_o and its true position θ_s . The lensing potential for a source plane at comoving distance r_s is [57, 65]

$$\psi(\hat{\mathbf{n}}, r_s) = \frac{2}{c^2} \int_0^{r_s} dr \frac{r_s - r}{r_s r} \Phi(\hat{\mathbf{n}} r, t(r)), \quad (2.45)$$

where Φ is the gravitational potential, and so the associated convergence [66]

$$\kappa(\hat{\mathbf{n}}, r_s) = \frac{1}{c^2} \int_0^{r_s} dr \frac{r(r_s - r)}{r_s} \nabla_{\mathbf{r}}^2 \Phi(\hat{\mathbf{n}} r, t(r)) \quad (2.46)$$

$$= \frac{3H_0^2 \Omega_{m0}}{2c^2} \int_0^{r_s} dr \frac{r(r_s - r)}{r_s} \frac{\delta(\mathbf{r}, t)}{a(r)} \quad (2.47)$$

is proportional to the projected mass density. In the limit where the gravitational lensing mass is confined to a thin plane at distance r_l , the convergence κ is physically defined as a dimensionless surface mass density, $\kappa = \Sigma / \Sigma_{\text{cr}}$, where $\Sigma_{\text{cr}} = c^2(4\pi G)^{-1} r_s / (r_l(r_s - r_l))$ is a critical density characterizing the lens-source system.

Note that above we have used the Poisson equation in comoving coordinates,

$$\nabla^2 \Phi(\mathbf{r}, t) = \frac{3H_0^2 \Omega_{m0}}{2a(t)} \delta(\mathbf{r}, t), \quad (2.48)$$

and to go from Eq. (2.44) to Eq. (2.46), we made the substitution,

$$\nabla_{\theta}^2 \Phi = r^2 \nabla_{\mathbf{r}_{\perp}}^2 \Phi \approx r^2 \nabla_{\mathbf{r}}^2 \Phi. \quad (2.49)$$

This substitution appears to neglect the contribution to $\nabla_{\mathbf{r}}^2$ from the line-of-sight derivatives, but the line-of-sight integral ensures that this can be done without loss of accuracy: it can

be shown that the component of the integrand $\propto \partial^2 \Phi / \partial r^2$ integrates to zero.

The effective convergence for a population of sources i (e.g. associated with a specific tomographic redshift bin) with normalized redshift distribution dn^i/dz defined as in Eq. (2.39), the effective lensing convergence will be

$$\kappa^i(\hat{\mathbf{n}}) = \frac{3H_0^2 \Omega_{m0}}{2c^2} \int_0^{r_h} dr_s \frac{H(z_s)}{c} \frac{dn^i}{dz_s} \int_0^{r_s} dr \frac{r(r_s - r)}{r_s} \frac{\delta(\mathbf{r}, t)}{a(r)} \quad (2.50)$$

$$= \frac{3H_0^2 \Omega_{m0}}{2c} \int_0^\infty dz_s \int_0^\infty dz \frac{r(z)(r(z_s) - r(z))}{r(z_s)} \frac{dn^i}{dz_s} \frac{\delta(\hat{\mathbf{n}} r(z), z)}{H(z) a(z)}. \quad (2.51)$$

Having converted this expression to an integral over z allows us, by comparison to Eqs. (2.40) and (2.42) to define a weight function for shear convergence,

$$W_\kappa^i(z) = \left(\frac{3H_0^2 \Omega_{m0}}{2c} \right) \left(\frac{r(z)}{a(z) H(z)} \right) \int_z^\infty dz' \frac{dn^i}{dz'} \frac{r(z') - r(z)}{r(z')}, \quad (2.52)$$

such that

$$\kappa^i(\hat{\mathbf{n}}) = \int dz W_\kappa^i(z) \delta(\hat{\mathbf{n}} r(z), z). \quad (2.53)$$

This can be used to compute the angular power spectrum of correlations, using the Limber approximation, between the shear convergence of sources in tomographic bins i and j ,

$$C_{\kappa\kappa}^{ij}(\ell) = \int dz \frac{H(z)}{c r^2(z)} W_\kappa^i(z) W_\kappa^j(z) P(k, z) \Big|_{k=(\ell+\frac{1}{2})/r(z)}, \quad (2.54)$$

as well as between the projected galaxy density for bin i and the convergence for bin j ,

$$C_{g\kappa}^{ij}(\ell) = \int dz \frac{H(z)}{c r^2(z)} W_g^i(z) W_\kappa^j(z) P(k, z) \Big|_{k=(\ell+\frac{1}{2})/r(z)}. \quad (2.55)$$

The notation used here can be related to that in Ref. [67] simply by converting the z integrals to comoving distance via $dz = H(z)dr/c$ and relating their weight functions to ours via $q_X^i(k, r) = W_X^i(k, r) H(z)/c$.

When studying the effects of lensing for a galaxy survey analysis, rather than measuring the convergence directly, we estimate the shear distortions to galaxy shapes. Shear is a two-component quantity $\gamma = (\gamma_1, \gamma_2)$, defined in terms of derivatives of the lensing potential




Figure 2.1: Illustration of shear coordinates. Here the blue ellipses represent a pair of galaxies separated by angle θ , and γ_1 and γ_2 are shear coordinates defined γ_1 is parallel to their separation vector. The coordinates γ_t and γ_x are defined to measure shear components that are parallel to and perpendicular, respectively, to the line connecting the two galaxies.

with respect to angular coordinates θ_1 and θ_2 on the sky as

$$\gamma_1 = \frac{1}{2} \left(\frac{\partial^2 \psi}{\partial \theta_1^2} - \frac{\partial^2 \psi}{\partial \theta_2^2} \right), \quad \text{and} \quad \gamma_2 = \frac{\partial^2 \psi}{\partial \theta_1 \partial \theta_2}. \quad (2.56)$$

Given a catalog of galaxy images, we can estimate the shear two-point functions by averaging the components of galaxies' measured ellipticities that are perpendicular or parallel to the line connecting galaxy pairs on the sky. The relevant shear quantities for these measurements is therefore the tangential shear γ_t and the cross shear γ_x . If the θ_1 - θ_2 basis is set up so that the two galaxies in a pair are along the θ_1 axis, these components are defined so $\gamma_t = -\gamma_1$ and $\gamma_x = -\gamma_2$. This is illustrated in Fig. 2.1. Shear-shear correlation studies most commonly make use of the two-point functions, [68, 69]

$$\xi_{\pm} = \langle \gamma_t \gamma_t \rangle \pm \langle \gamma_x \gamma_x \rangle. \quad (2.57)$$

It can be shown that in the weak lensing regime, where $\kappa \ll 1$ and $\gamma \equiv \sqrt{\gamma_1^2 + \gamma_2^2} \ll 1$, shear correlations can be simply related to convergence power spectra [66]. In the flat-sky approximation, where sums over spherical harmonics are converted to two-dimensional Fourier modes, the predicted angular correlations between the shears of galaxies in tomographic bins

i and j are given by [67]

$$\xi_+^{ij}(\theta) = \int \frac{d\ell \ell}{2\pi} J_0(\ell\theta) C_{\kappa\kappa}^{ij}(\ell), \quad (2.58)$$

$$\xi_-^{ij}(\theta) = \int \frac{d\ell \ell}{2\pi} J_4(\ell\theta) C_{\kappa\kappa}^{ij}(\ell), \quad (2.59)$$

where $J_m(x)$ is a Bessel function of the first kind of order m . Similarly, the cross correlation between galaxy positions in bin i and tangential shears of source galaxies in bin j is

$$\gamma_t^{ij}(\theta) = \int \frac{d\ell \ell}{2\pi} J_2(\ell\theta) C_{g\kappa}^{ij}(\ell). \quad (2.60)$$

It can be shown that, for the precision of current weak lensing surveys, the differences introduced to these quantities by the flat sky approximations are negligible compared to statistical uncertainties [65, 67].

2.5 Summary

The parameters of the standard cosmological model introduced in this chapter are summarized in Table 2.1.

Parameter	Description
z	Redshift
$H(t)$	Hubble rate, $d \ln a / dt$
H_0	Hubble parameter, $d \ln a / dt _{t_0}$
h	$H_0 / [100 \text{ km s}^{-1} \text{ Mpc}^{-1}]$
Ω_m	Total matter density in units of ρ_{crit}
Ω_b	Baryon density in units of ρ_{crit}
Ω_c	Cold dark matter density in units of ρ_{crit}
Ω_ν	Neutrino density in units of ρ_{crit}
Λ	Cosmological constant
Ω_{DE}	Dark energy density in units of ρ_{crit}
w	Dark energy equation of state, equal to -1 for Λ
n_s	Spectral index of primordial power spectrum
A_s	Amplitude of primordial power spectrum
σ_8^2	Variance of matter density in spheres of radius $8h^{-1} \text{ Mpc}$
$w^{ij}(\theta)$	Galaxy-galaxy correlation between z -bins i and j
$\xi_{\pm}^{ij}(\theta)$	Shear-shear angular correlation between z -bins i and j
$\gamma_t^{ij}(\theta)$	Galaxy-shear angular correlation between z -bins i and j

Table 2.1: Cosmological parameters and other quantifies defined in this chapter.

Chapter 3

ISW signal reconstruction

As cosmic microwave background (CMB) photons travel from the last scattering surface to our detectors, they can experience a frequency shift beyond that which is guaranteed by the expansion of the universe. This additional effect is a result of the fact that gravitational potential fluctuations associated with large-scale structure (LSS) evolve with time when the Universe is not fully matter dominated. Consequently, the CMB photons are subject to a direction-dependent temperature modulation which is proportional to twice the rate of change in the potential integrated along the line of sight. This modulation is known as the integrated Sachs-Wolfe (ISW) effect [70]. Its magnitude in direction $\hat{\mathbf{n}}$ on the sky was worked out in the classic Sachs-Wolfe paper [71] to be

$$\left. \frac{\Delta T}{\bar{T}} \right|_{ISW}(\hat{\mathbf{n}}) = \frac{2}{c^2} \int_{t_*}^{t_0} dt \frac{\partial \Phi(\mathbf{r}, t)}{\partial t}, \quad (3.1)$$

where t_0 is the present time. t_* is that of recombination, c is the speed of light, \mathbf{r} is the position in comoving coordinates, and Φ is the gravitational potential.

The ISW effect introduces a weak signal at very large scales (low multipoles) in the CMB angular power spectrum. It carries important information about dark energy [72, 73], particularly its clustering properties that are often parametrized by the dark energy speed of sound. It also potentially offers useful information about the nature of dark energy, as modified gravity theories have unique ISW signatures [74]. However, the fact that the largest CMB multipoles are subject to cosmic variance severely limits how much information can be gleaned from the ISW given the CMB temperature measurements alone.

We are able to observe the ISW effect because the dependence of the ISW signal on the time derivative of the potential results in a large-angle cross-correlation between LSS tracers and CMB temperature. This was first pointed out by Crittenden & Turok [75],

who further suggested cross-correlation between CMB temperature anisotropy $(\delta T/T)_{\text{ISW}}(\hat{\mathbf{n}})$ and galaxy positions, $(\delta N/N)(\hat{\mathbf{n}}')$, as a statistic through which to detect the ISW effect. This cross-correlation signal was detected shortly thereafter [76] and was later confirmed by many teams who found cumulative evidence of about 4σ using a number of different LSS tracers [77–91]. Comprehensive surveys of recent results can be found in Refs. [89, 91, 92]. While the detection of the ISW effect itself provides independent evidence for dark energy at high statistical significance, prospects for using it to constrain the cosmological parameters are somewhat limited [93].

The ISW *map*, $(\delta T/T)_{\text{ISW}}(\hat{\mathbf{n}})$, is also of interest in its own right. By assuming a cosmological model, one can construct an estimator using theoretical cross-correlations in combination with LSS data. Because the ISW signal represents a late-universe contribution to the CMB anisotropy, measuring and subtracting it from observed temperature fluctuations would allow us to isolate the (dominant) early-universe contributions to the CMB. If this procedure could be done reliably, it would have immediate implications for our understanding of the cosmological model.

For example, the ISW signal has been identified as a potential contributor to large-angle CMB features which have been reported to be in tension with the predictions of Λ CDM [94]. A reconstructed ISW map would clarify whether some component of the CMB anomalies (discussed further in Section 3.6) become stronger or weaker when evaluated on the early-universe-only contribution to the CMB. A few studies [95, 96] have already explored this. To study the impact of ISW contributions on CMB anomalies, Ref. [96] uses WMAP data with 2MASS and NVSS, while Ref. [95] uses 2MASS alone.

The late-time ISW also provides a contaminant to the measurement of primordial non-Gaussianity from CMB maps. Because both the ISW effect and gravitational lensing trace LSS, they couple large- and small-scale modes of the CMB, resulting in a nonprimordial contribution to the bispectrum. Recent analyses [97] have corrected for this by including a theoretical template for the ISW-lensing bispectrum in analyses of primordial non-Gaussianity. Reconstructing and subtracting the ISW contribution from the CMB temperature maps could provide an alternative method for removing ISW-lensing bias when studying primordial non-Gaussianity [98]. More generally, understanding how reliably the ISW map can be reconstructed from large-scale structure information impacts our understanding of how the late universe affects our view of the primordial CMB sky.

Before reconstruction can be done reliably, however, we must understand how systematics associated with the input data impact the ISW estimator’s accuracy. Previous works have explored this to some extent, looking at how reconstruction quality is affected by the inclusion

of different input data sets [11, 91, 99], masks [91, 99] and, to a limited degree, the influence of uncertainties in cosmological and bias models [99]. Additionally, Ref. [100] studied how systematics like redshift uncertainties and photometric calibration change the signal to noise of the ISW effect’s *detection*. That being said, there remain a number of systematics inherent to galaxy survey data which have not yet been subject to detailed analysis in the context of ISW map reconstruction. We aim to address this.

In this analysis, we use simulated ISW and LSS maps to identify which survey properties are important for ISW reconstruction and to quantify their effects on the reconstructed maps. We begin by studying how survey depth, redshift binning strategy, and the minimum measured multipole ℓ_{\min} influence reconstruction quality in the absence of systematics. Using these results as a baseline, we then explore two broad classes of systematics: ways one can mismodel the redshift distribution of LSS sources, and direction-dependent photometric calibration errors that can result from, for example, contamination by stars. In doing this, we demonstrate that while the impact of redshift-distribution modeling errors is mostly negligible, even small levels of photometric calibration errors can have devastating effects on ISW signal reconstruction accuracy. We also briefly discuss the implications of our results for analysis of whether the ISW signal contributes to the observed alignments between large-angle multipoles of the CMB temperature map.

The chapter is organized as follows. In Section 3.1 we discuss our general procedure for the ISW map reconstruction and assessment of the accuracy in this procedure. In Section 3.2, we describe the properties of the surveys that we will consider, and draw on the results to interpret a cross-check, presented in Section 3.3, of the results for ISW reconstruction-accuracy results in Ref. [11]. Then, in Sections 3.4 and 3.5 we discuss the effect of various systematic errors on the ISW map reconstruction. We conclude in Section 3.7.

3.1 Methods

We perform a number of studies examining how survey properties and systematics affect the accuracy of reconstructed ISW maps. These studies all follow this general pipeline:

- Select a fiducial cosmological model and specifications of the LSS survey.
- Compute the “true” angular cross-power C_{ℓ}^{XY} for ISW and LSS maps, assuming the fiducial cosmology and survey specifications.
- Use the true C_{ℓ}^{XY} to generate correlated Gaussian realizations of the true ISW signal and corresponding LSS maps.

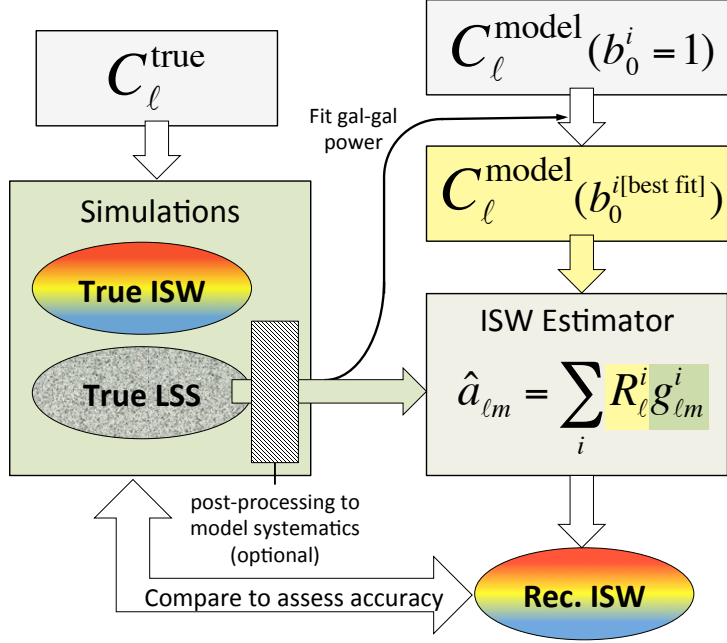


Figure 3.1: Flowchart of ISW reconstruction pipeline.

- If applicable, postprocess the galaxy maps to model direction-dependent systematic effects.
- Construct an estimator for the ISW signal using the simulated galaxy maps and a set of “model” C_ℓ^{XY} which may or may not match those used to generate the simulations.
- Compare the reconstructed ISW signal to the true ISW map and evaluate the accuracy of the reconstruction.

This section will introduce some of the theoretical tools needed for this analysis.

3.1.1 Theoretical cross-correlations

The angular cross-power between ISW and galaxy maps serves as input for both the simulation and reconstruction processes used in the following sections. Given maps X and Y , the expression for the angular cross-power between them is

$$C_\ell^{XY} = \frac{2}{\pi} \int dk k^2 P(k) I_\ell^X(k) I_\ell^Y(k) \quad (3.2)$$

where $P(k)$ is the matter power spectrum at $z = 0$, and the transfer function $I_\ell^X(k)$ is written

$$I_\ell^X(k) \equiv \int_0^\infty dz D(z) W^X(z, k) j_\ell(kr). \quad (3.3)$$

Here, $r \equiv r(z)$ represents comoving radius; $j_\ell(x)$ is a spherical Bessel function; and $D(z)$, which is normalized to one at $z = 0$, describes the linear growth of matter fluctuations. The function $W^X(k, z)$ is tracer-specific window function that encapsulates the relationship between the tracer X and underlying dark matter fluctuations δ . Note that this expression looks different from that given in Section 2.4.1 because it is the exact expression for C_ℓ^{XY} , while Eq. (2.41) (which will be reiterated below) uses the Limber approximation. Both of these expressions are derived in Appendix A.

The tracers relevant to our studies are the ISW signal and galaxy number density. The ISW window function is

$$W^{\text{ISW}}(z, k) = [\Theta(z_{\text{max}} - z)] \left[\frac{3H_0^2 \Omega_m}{c^2 k^2} \right] (1 - f(z)), \quad (3.4)$$

where Θ is the Heaviside step function. In this expression, derived in Section A.5 of Appendix A, the term in square brackets comes from when the Poisson equation is used to relate potential fluctuations to dark matter density, Ω_m is the matter density in units of the critical density, and H_0 is the present-day Hubble parameter. The appearance of the growth rate $f(z) \equiv d \ln D / d \ln a$ comes from the time derivative in Eq. (3.1). To compute the full ISW contribution, one would integrate to the redshift of recombination, $z_{\text{max}} = z_*$. In this work, though, we are interested only in the late ISW effect, so we can set $z_{\text{max}} = 15$ without a loss in accuracy.

Each survey (and each redshift bin within a given survey) will have its own window function. For a map of galaxy number density fluctuations, it is the same as the function $W_g(k, z)$ used in Section 2.4.1,

$$W^{\text{gal}}(z, k) = b(z) \frac{dn}{dz}. \quad (3.5)$$

Reiterating the discussion from Chapter 2, in this expression, $b(z)$ represents linear bias, which we assume is scale independent. The function dn/dz describes the redshift distribution of the observed sources, encapsulating information about how their physical density varies with redshift as well as survey volume and selection effects. It is normalized so it integrates

to one. Galaxy shot noise is included by adding a contribution to its autopower spectrum,

$$C_\ell^{\text{gal-gal}} \rightarrow C_\ell^{\text{gal-gal}} + \bar{n}^{-1} \quad (3.6)$$

where \bar{n} is the average number density of sources per steradian. In summary, to simulate a given galaxy survey, we need $b(z)$, describing how clustered its sources are relative to dark matter; dn/dz , describing how the observed sources are distributed along the line of sight; and \bar{n} , the average number density of sources per steradian.

For $\ell > 20$, we use the Limber approximation to compute C_ℓ^{XY} . This dramatically reduces the computation time and gives results that are accurate to within about 1% [101]. In this approximation, the cross-correlations become

$$C_\ell^{XY} = \int dz \frac{H(z) D^2(z)}{c r^2(z)} [P(k) W^X(k, z) W^Y(k, z)]_{k=k_\ell}, \quad (3.7)$$

where $k_\ell = (\ell + \frac{1}{2})/r(z)$ and $H(z)$ is the Hubble parameter.

We developed an independent code to calculate the cross-power spectra C_ℓ^{XY} and have extensively tested its accuracy for various survey redshift ranges against the publicly available CLASS code [102], which is accurate to about 0.1%.

3.1.2 Simulating LSS maps

As we care only about large-angle ($\ell \lesssim 100$) features, we model the ISW signal and galaxy number density fluctuations as correlated Gaussian fields. To simulate them, we compute the relevant angular auto- and cross-power C_ℓ 's and then use the `synalm` function from Healpy [5] to generate appropriately correlated sets of spherical harmonic coefficients $g_{\ell m}$. These components are defined via the spherical harmonic expansion of the number density of sources in the i th LSS map,

$$\left[\frac{\delta N}{N} \right]^i(\hat{\mathbf{n}}) = \sum_{\ell m} g_{\ell m}^i Y_{\ell m}(\hat{\mathbf{n}}). \quad (3.8)$$

For each study using simulated maps, we generate 10,000 map realizations. We use Healpix with NSIDE=32 and compute C_ℓ up to $\ell_{\text{max}} = 95$, guided by the relation¹ $\ell_{\text{max}} = 3(\text{NSIDE}) - 1$. Unless we state otherwise, our ISW reconstructions include multipole information down to $\ell_{\text{min}} = 2$.

¹See, for example the Healpix user guide at <https://healpix.jpl.nasa.gov/html/facilitiesnode7.htm>.

All of our analyses are for full-sky data and our fiducial cosmological model is Λ CDM, with parameter values from best-fit Planck 2015,

$$\{\Omega_c h^2, \Omega_b h^2, \Omega_\nu h^2, h, n_s\} = \{0.1188, 0.0223, 0, 0.6774, 0.9667\}. \quad (3.9)$$

3.1.2.1 Fiducial survey

We model our fiducial galaxy survey on what is expected for Euclid [103]. With its large sky coverage and deep redshift distribution, the Euclid survey has been identified as a promising tool for ISW detection [100, 104] and it is reasonable to assume that these properties will also make it a good data set to use for ISW reconstruction. We therefore adopt the redshift distribution used in Ref. [105],

$$\frac{dn}{dz} = \frac{3}{2z_0^2} z^2 \exp[-(z/z_0)^{-1.5}] \quad (3.10)$$

which has a maximum at $z_{\text{peak}} \simeq 1.21z_0$. We adopt $z_0 = 0.7$ and $\bar{n} = 1 \times 10^9$ sources per steradian. For binning studies (see Section 3.2.2) we assume a photo- z redshift uncertainty of $\sigma(z) = 0.05(1+z)$. Our fiducial bias is $b(z) = 1$. We explicitly state below whenever these fiducial values are varied for our tests.

3.1.3 ISW estimation

In contrast to ISW detection studies mentioned in the introduction to this chapter, in which the correlation between observed galaxy and CMB maps is measured as a signal, for ISW signal reconstruction, we use the theoretical cross correlations for a fiducial cosmological model in order to estimate the ISW contributions to CMB temperature fluctuations given observed LSS density fluctuations. We use the optimal estimator derived in Ref. [11] to reconstruct the ISW signal from LSS maps. Because we are interested in quantifying the impact of galaxy survey systematics, in this work we focus on the case where only galaxy maps are used as input. We thus neglect the part of the estimator that includes CMB temperature information and write

$$\hat{a}_{\ell m}^{\text{ISW}} = \sum_i^n R_\ell^i g_{\ell m}^i. \quad (3.11)$$

Here $\hat{a}_{\ell m}$ is the optimal estimator for the ISW map component, $g_{\ell m}^i$ is the observed spherical component of LSS tracer i , and n is the number of LSS tracers considered. The operator

$$R_\ell^i \equiv -N_\ell [D_\ell^{-1}]_{\text{ISW}-i} \quad (3.12)$$

is the reconstruction filter applied to the i th LSS map. It is constructed from the covariance matrix D_ℓ between ISW and LSS tracers,

$$D_\ell = \begin{pmatrix} C_\ell^{\text{ISW,ISW}} & C_\ell^{\text{LSS}_1,\text{ISW}} & \dots & C_\ell^{\text{LSS}_n,\text{ISW}} \\ C_\ell^{\text{LSS}_1,\text{ISW}} & C_\ell^{\text{LSS}_1,\text{LSS}_1} & \dots & C_\ell^{\text{LSS}_1,\text{LSS}_n} \\ \vdots & \vdots & \ddots & \vdots \\ C_\ell^{\text{LSS}_n,\text{ISW}} & C_\ell^{\text{LSS}_1,\text{LSS}_n} & \dots & C_\ell^{\text{LSS}_n,\text{LSS}_n} \end{pmatrix}. \quad (3.13)$$

The term $N_\ell^{-1} \equiv (D_\ell^{-1})_{11}$ estimates the reconstruction variance.

Note that for reconstruction using a single LSS map this reduces to a Wiener filter, a signal-processing method commonly used to separate signals and noise with known auto- and cross-correlations:

$$\hat{a}_{\ell m}^{\text{ISW}} \xrightarrow{\text{single LSS}} \frac{C_\ell^{\text{ISW-gal}}}{C_\ell^{\text{gal-gal}}} g_{\ell m}. \quad (3.14)$$

In the subsequent discussion, we will refer to the correlations appearing in D_ℓ (and thus the reconstruction filters R_ℓ^i) as C_ℓ^{model} . This is to distinguish them from the correlations used to generate the simulations, which we will call C_ℓ^{true} . We adopt this convention because if we were reconstructing the ISW signal based on real data, C_ℓ^{true} would be the correlations determined by the true underlying physics of the universe, while C_ℓ^{model} would be computed theoretically based on our best knowledge of cosmological parameters and the properties of the input LSS tracers.

Setting $C_\ell^{\text{model}} = C_\ell^{\text{true}}$ represents a best-case scenario where we have perfect knowledge of the physics going into the calculations outlined in Section 3.1.1. Incorrect modeling will break that equality, causing the estimator in Eq. (3.11) to become suboptimal. Our analysis of LSS systematics in Sections 3.4 and 3.5 will fundamentally be an examination of how different manifestations of the of $C_\ell^{\text{model}} \neq C_\ell^{\text{true}}$ mismatch impact reconstruction.

3.1.4 Fitting for effective galaxy bias

Our pipeline actually contains an additional step, which as we will see in later sections, helps protect against some systematics: before constructing the ISW estimator, we fit the galaxy maps for a constant bias.

When performing this procedure, the first step of our reconstruction process is to measure the galaxy autopower spectrum from the observed galaxy map, $C_\ell^{\text{gal}(\text{obs})}$. This will be subject to cosmic variance scatter about $C_\ell^{\text{gal}(\text{true})}$ and so will be realization dependent. We then perform a linear fit for a constant \bar{b} satisfying

$$C_\ell^{\text{gal}(\text{obs})} = \bar{b}^2 C_\ell^{\text{gal}(\text{model})}. \quad (3.15)$$

We then scale the model power spectra:

$$\begin{aligned} C_\ell^{\text{gal}} &\rightarrow \bar{b}^2 C_\ell^{\text{gal}}, \\ C_\ell^{\text{gal-ISW}} &\rightarrow \bar{b} C_\ell^{\text{gal-ISW}}, \\ C_\ell^{\text{gal}_i\text{-gal}_j} &\rightarrow \bar{b}^i \bar{b}^j C_\ell^{\text{gal}_i\text{-gal}_j}. \end{aligned} \quad (3.16)$$

If there are no systematics affecting our measurements, $C_\ell^{\text{gal}(\text{true})} = C_\ell^{\text{gal}(\text{model})}$, so \bar{b} will be close to 1. When a galaxy bias is modeled as a constant, b_0 , for each galaxy map, this scaling will exactly correct for any mismatch between the value used in the simulations and that in the model used to construct the ISW estimator:

$$\bar{b} = b_0^{\text{true}} / b_0^{\text{model}}. \quad (3.17)$$

Outside the case of constant bias, there is not a direct correspondence between \bar{b} and the parameters of the bias model. (It corresponds to the ratio between weighted averages of $b(z)^{\text{true}}$ and $b(z)^{\text{model}}$.) However, the procedure for fitting for and scaling by \bar{b} is well defined and makes our estimator robust against systematics which shift C_ℓ 's by a multiplicative constant, including mismodeled $b(z)$ and dn/dz . We will demonstrate this in Section 3.4.

3.1.5 Evaluating reconstruction accuracy

We will use two statistics to quantify the accuracy of reconstructed ISW maps. Primarily, we will use the correlation coefficient between the true ISW signal $T^{\text{ISW}}(\hat{\mathbf{n}})$ and the reconstructed ISW map $T^{\text{rec}}(\hat{\mathbf{n}})$. For a given realization we compute this as

$$\rho = \frac{\langle T^{\text{ISW}} T^{\text{rec}} \rangle_{\text{pix}}}{\sigma_{\text{ISW}} \sigma_{\text{rec}}}, \quad (3.18)$$

where $\langle \rangle_{\text{pix}}$ indicates an average over pixels, and σ_X is the variance of map X .

We can approximate the theoretical expectation value for ρ using the cross-power between

maps,

$$\langle \rho \rangle = \frac{\sum_{\ell i} (2\ell + 1) R_{\ell}^i C_{\ell}^{\text{ISW}-i}}{\langle \sigma_{\text{rec}} \rangle \langle \sigma_{\text{ISW}} \rangle}, \quad (3.19)$$

where the indices i and j label LSS maps and

$$\langle \sigma_{\text{ISW}} \rangle = \sqrt{\sum_{\ell} (2\ell + 1) C_{\ell}^{\text{ISW}}} \quad (3.20)$$

$$\langle \sigma_{\text{rec}} \rangle = \sqrt{\sum_{\ell ij} (2\ell + 1) R_{\ell}^i R_{\ell}^j C_{\ell}^{ij}} \quad (3.21)$$

are the standard deviations of the temperature maps. In deriving this expression, we assumed $\langle \sigma^{-1} \rangle = \langle \sigma \rangle^{-1}$ and that the various factors in this expression are uncorrelated. We will see later that this is a reasonably accurate approximation to make, as it gives values which are in good agreement with simulation results.

One can see by examining Eqs. (3.18) and (3.19) that ρ is sensitive to the reconstruction of phases but insensitive to changes in the overall amplitude of the reconstructed ISW map. Because of this, though $\rho \rightarrow 1$ is generally indicative of a more accurate reconstruction, this quantity does not capture all important information about reconstruction quality. We therefore also consider a complementary statistic which is sensitive to amplitude, defined

$$s = \frac{\langle (T^{\text{ISW}} - T^{\text{rec}})^2 \rangle_{\text{pix}}^{1/2}}{\sigma_{\text{ISW}}}. \quad (3.22)$$

The quantity s measures how the average size of errors in the reconstructed signal compares to that of fluctuations in the true ISW map. As with ρ , we can compute its expectation value,

$$\langle s \rangle = \frac{\sqrt{\langle \sigma_{\text{rec}} \rangle^2 + \langle \sigma_{\text{ISW}} \rangle^2 - 2 \sum_{\ell i} (2\ell + 1) R_{\ell}^i C_{\ell}^{\text{ISW}-i}}}{\langle \sigma_{\text{ISW}} \rangle}. \quad (3.23)$$

Because the bias-fitting procedure discussed in Section 3.1.4 corrects for amplitude differences, for most of the scenarios we study, ρ and s effectively contain the same information. For this reason, we will primarily use ρ as our quality statistic and will only show results for s when it contributes new insight.

Throughout this chapter we will use angled brackets to indicate the theoretical expectation values for these statistics, and an overbar to indicate averages computed from simulations.




Figure 3.2: Scatter plot comparing the true (simulated) ISW signal, on the horizontal axis, to the reconstructed ISW signal, on the vertical axis, for a single realization assuming each of five different depths of the survey. Each data point corresponds to one pixel on an NSIDE=32 map. If there was a perfect reconstruction, all points would fall on the dotted line.

3.2 The effect of survey properties

Before studying the effects of systematics, it is instructive to explore how LSS survey properties impact ISW signal reconstruction in the ideal, $C_\ell^{\text{model}} = C_\ell^{\text{true}}$, scenario. This has already been done to some extent in Refs. [11], [91], and [99].

Our studies in this section will serve two primary purposes. First, they will provide a straightforward demonstration of our pipeline and the reconstruction quality statistics introduced in Section 3.1.5. More importantly, they will serve as a baseline for our analysis of systematics in Sections 3.4 and 3.5: Our goal is *not* to find optimized survey properties for ISW signal reconstruction, though our results might serve as a rough guide for doing so. Rather, we want to study how shifting, for example, survey depth or redshift binning strategy affects ISW reconstruction in the best-case scenario (no with systematic errors) so that we can better understand the impact of what happens when those errors are introduced.




Figure 3.3: Histograms of the correlation between true and reconstructed ISW maps ρ (left panel), and the typical size of residuals relative to that of the true ISW map fluctuations s (right panel). These plots show the results of 10,000 simulations for surveys of various depths, with their dn/dz distributions shown in arbitrary units as an inset in the left plot. The solid and dashed vertical lines show the theoretical expectation value and measured average, respectively, for the statistic in question.

3.2.1 Varying survey depth

The first property we examine is survey depth. We model this by changing the value of z_0 in our fiducial dn/dz [Eq. (3.10)] while holding all other survey properties fixed. We look at values $\Delta z = \pm 0.1$ on either side of our fiducial $z_0 = 0.7$, plus a redshift distribution comparable to DES [106] with $z_0 = 0.5$ and the even-shallower $z_0 = 0.3$.

Figure 3.2 shows a pixel-by-pixel comparison between the reconstructed and true ISW signal for a single representative realization. We can see that the deeper surveys have data-points more tightly clustered around the $T_{\text{rec}}^{\text{ISW}} = T_{\text{true}}^{\text{ISW}}$ diagonal and correspondingly higher values of ρ .

We find that this pattern holds, if noisily, in the full ensemble of simulated maps. Figure 3.3 shows histograms of ρ for the same surveys, with their dn/dz distributions shown in an inset. In it, the sample average $\bar{\rho}$ of Eq. (3.18) and theoretical expectation value $\langle \rho \rangle$ are plotted as dashed and solid vertical lines, respectively. We find that though $\langle \rho \rangle$ tends to be lower than $\bar{\rho}$, the difference between them is much smaller than the scatter in the data, and that the ordering of $\langle \rho \rangle$ values for the different surveys is consistent with the results from simulations. We take this to mean that the more computationally efficient $\langle \rho \rangle$ is a slightly biased but reasonably reliable indicator of the ISW reconstruction quality.

Looking at the data, we also note that the scatter in the individual ρ distributions is

large compared to the difference between their mean values. This tells us that, while $\langle \rho \rangle$ (or $\bar{\rho}$) values succeed in predicting how ISW reconstruction quality from different surveys will compare on average, they are a relatively poor predictor of how surveys will compare for any individual realization.

For illustrative purposes, in Fig. 3.3, we also show a histogram for the values of statistic s which, recall, is mainly sensitive to the amplitude accuracy in the map reconstruction – measured from the same simulations. We see that (as expected) surveys with larger $\bar{\rho}$ have smaller \bar{s} and that the surveys with $\bar{\rho} \sim 0.9$ correspond to $\bar{s} \sim 0.4$. This tells us that even in the best maps that we study here, errors in the reconstructed ISW temperature are a little over one-third of the amplitude of true ISW signal fluctuations.

We keep the mean source number density \bar{n} fixed for this analysis, so that any differences we observe in reconstruction quality are due only to how the redshift distributions are sampled, not to the fact that a deeper survey will observe a larger number of sources. We argue that this is well motivated because the only way \bar{n} enters our calculations is via shot noise, and we have set it to a large enough value so that its contributions are negligible on large, ISW-relevant scales.

3.2.2 Redshift binning strategy

Here we study how different strategies for binning galaxy data affect the reconstruction. For each bin with $z_i \leq z < z_{i+1}$, we model the redshift distribution by weighting the survey’s overall distribution dn^{tot}/dz with a window function $F_i(z)$ and scale the total number density accordingly:

$$\frac{dn^i}{dz} = \frac{\frac{dn^{\text{tot}}}{dz} F_i(z)}{\int_0^\infty \frac{dn^{\text{tot}}}{dz} F_i(z) dz}, \quad (3.24)$$

$$\bar{n}^i = \bar{n}^{\text{tot}} \times \left[\int_0^\infty \frac{dn^{\text{tot}}}{dz} F_i(z) dz \right]. \quad (3.25)$$

We can then compute C_ℓ^{XY} using the expressions in Section 3.1.1, treating each redshift bin as an individual map (X or Y).

Photometric redshift uncertainties will cause sharp divisions in observed redshift to be smoothed when translated to spectroscopic redshift. As in Ref. [11] we therefore model the

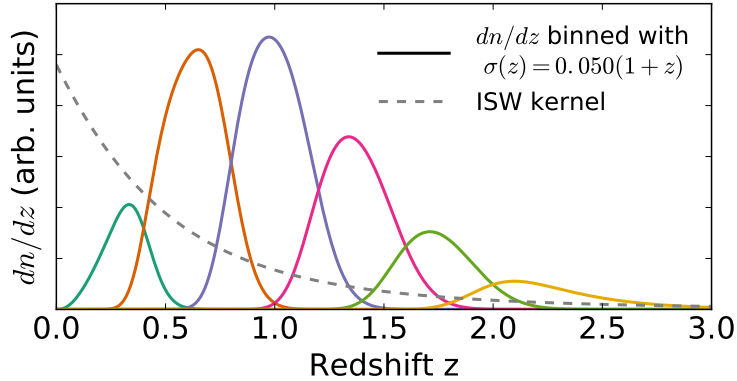


Figure 3.4: Un-normalized redshift distributions for the six redshift bins studied, with photometric-redshift uncertainty $\sigma(z) = 0.05(1+z)$. Because these distributions are not yet normalized (they neglect the denominator of Eq. (3.24)), the area under the curves gives an idea of the relative number of galaxies in each bin. The dotted line shows the ISW kernel in arbitrary units.

effect of photometric uncertainties $\sigma(z)$ via

$$F_i(z) = \frac{1}{2} \left[\operatorname{erfc} \left(\frac{z_i - z}{\sigma(z)\sqrt{2}} \right) - \operatorname{erfc} \left(\frac{z_{i+1} - z}{\sigma(z)\sqrt{2}} \right) \right], \quad (3.26)$$

which effectively acts as a smoothed top-hat window in z . We use the standard form for photometric-redshift uncertainty

$$\sigma(z) = \sigma_{z0} \times (1+z). \quad (3.27)$$

For reference, Euclid forecasts consider $\sigma_{z0} = 0.05$ a requirement and give $\sigma_{z0} = 0.03$ as a reach goal [103, 107].

In order to understand how binning affects ISW reconstruction, we split our fiducial redshift distribution into the six bins shown in Fig. 3.4 and compute all possible auto- and cross-correlations between them. We then use the relations from Ref. [108] to compute C_ℓ^{XY} for cases where two or more adjacent bins are merged.

To check that our understanding of reconstruction statistics holds for surveys with multiple redshift bins, we simulated 10,000 map realizations for three configurations: the one-bin fiducial case, the six-bin case, and a three-bin case with edges at $z \in [0, 0.8, 1.6, 3.5]$. For all of these, we used $\sigma_{z0} = 0.05$. The results, shown in Fig. 3.5, reveal that though binning slightly improves the reconstruction quality, it does not dramatically change the shape of

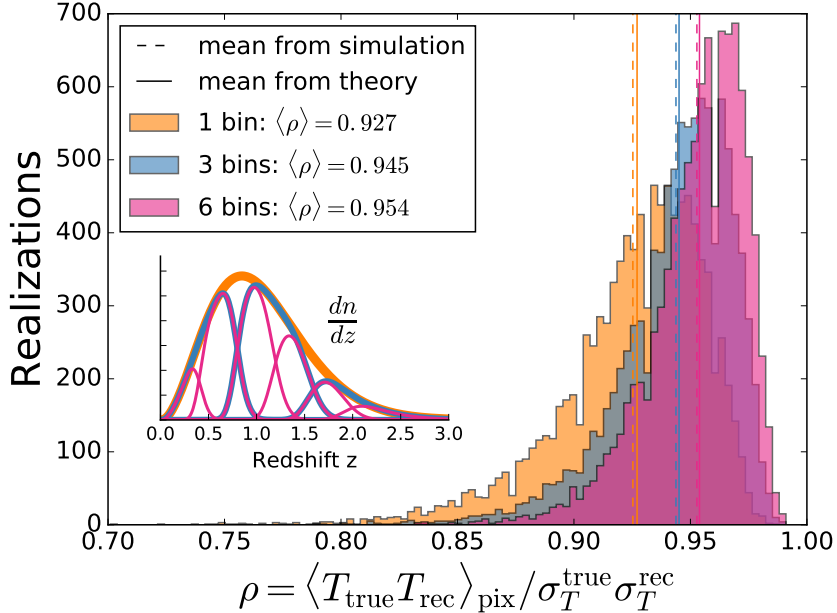


Figure 3.5: Histogram of ρ values measured from 10,000 map realizations for selected binning strategies. The inset shows the un-normalized dn/dz distributions for the sets of redshift bins considered.

the ρ distribution, nor the relationship between $\langle \rho \rangle$ and $\bar{\rho}$.

We see that splitting data into redshift bins improves our ISW reconstruction, if only slightly: the correlation between the reconstructed and true map shifts by $\Delta\rho \lesssim 0.03$. This change is smaller than the observed scatter in ρ and is comparable to that produced in the previous section by shifting the survey depth by $\Delta z = \pm 0.1$ about $z_0 = 0.7$. This improvement could be due to gains in three-dimensional information, or to the fact that we are now using multiple LSS maps with uncorrelated noise.

Reassured that $\langle \rho \rangle$ is still a reliable statistic, we compute it for all 32 possible combinations of the six bins from Fig. 3.4. The results are shown in Fig. 3.6. In this figure, the bars labeling the y -axis schematically illustrate the binning configurations, with different colors corresponding to different numbers of bins. The data points show $\langle \rho \rangle$ for various values of σ_{z_0} , while the X-shaped points with error bars show the mean and standard deviations extracted from the histograms in Fig. 3.4.

We note a couple of patterns in the results. First, for a fixed number of bins, the reconstruction tends to be better if we place finer divisions at high redshift. Also, having a smaller photometric-redshift uncertainty actually slightly degrades the reconstruction rather than improving it. This implies that combining maps with redshift distributions which




Figure 3.6: Theoretical expectation value for ρ computed for different redshift binning strategies and levels of photometric-redshift uncertainties. The colored bars and corresponding numbers on the left side of the plot are schematic labels for how the galaxies are divided into redshift bins. Different colored points show the effect of different photo- z uncertainties. The “X” points with blue horizontal error bars show the mean and standard deviation of ρ extracted from the histograms in Fig. 3.5.

overlap more tend to lead to better reconstructions. This could be due a multitracer effect, in that overlap between bins means that we are sampling the same potential fluctuations with multiple source populations. However, it is also possible this is due to how our model of $\sigma(z)$ affects the shapes of the redshift distributions, which itself only an approximation for a survey’s true photo- z uncertainties. Given the small size of these effects, one should be cautious about assigning them much physical significance.

Last, we observe a shift $\Delta\rho$ due to changes in binning that is smaller than what is found in the work by Ref. [11] by about a factor of 3. Because their simulated DES-like survey is shallower than our fiducial survey and the relationship between $\Delta\rho$ and $\bar{\rho}$ is nonlinear (e.g., a shift from 0.98 to 0.99 is more significant than one from 0.28 to 0.29), this does not necessarily mean that our results are incompatible. As a cross-check, we performed additional simulations similar to those analyzed in Ref. [11]. Our results, discussed in Section 3.3, support this.

3.2.3 Varying ℓ_{\min} of reconstruction

For most of the studies presented in this chapter, we reconstruct and assess the accuracy of ISW maps using all multipoles with $2 \leq \ell \leq 95$. This range is chosen because $\ell = 2$ is the lowest multipole typically considered for CMB analysis and $\ell = 95$ is the maximum multipole which corresponds to angular scales larger than the size of individual pixels in NSIDE=32 Healpix maps. In this section, we study the effect of changing ℓ_{\min} .

When we perform ISW map reconstruction, we enforce ℓ -range requirements in three ways. First, when we construct the ISW estimator shown in Eq. (3.11), we set all R_ℓ^i not satisfying $\ell_{\min} \leq \ell \leq \ell_{\max}$ to be zero, so the reconstructed map contains no information from multipoles outside that range. Second, when analyzing simulations, we remove the same ℓ values from maps before computing ρ . Likewise, when we analytically compute $\langle\rho\rangle$ as shown in Eq. (3.19), we restrict the sum over multipole to $\ell_{\min} \leq \ell \leq \ell_{\max}$. In other words, when we show $\rho_{\ell \geq \ell_{\min}}$, we are showing the result for an ISW map reconstructed for a limited range of ℓ values, evaluated by considering only those multipoles.

The results of this analysis are shown in Fig. 3.7. Here we show the correlation coefficient between true and reconstructed maps $\rho_{[\ell \geq \ell_{\min}]}$ as a function of the minimum multipole used in the reconstruction. The solid line is the theoretical expectation value, while the data points with error bars show results from simulations. We find that ρ increases with the minimum multipole out to $\ell_{\min} \sim 5$, after which it begins to very gradually decrease with ℓ_{\min} . Increasing ℓ_{\min} also decreases the scatter in ρ measured across realizations.




Figure 3.7: How filtering out angular scales with $\ell > \ell_{\min}$ affects reconstruction of ISW map. The data points show the mean and standard deviation of ρ , the correlation coefficient between true and reconstructed ISW maps, observed in 10,000 realizations, while the line shows the value of $\langle \rho \rangle$ computed analytically.

We interpret these trends to be the result of a competition between cosmic variance and the fact that most ISW information (power and cross-power) is at small multipoles. That is, removing the lowest few multipoles (out to $\ell \simeq 4$) from the analysis largely removes noise due to cosmic variance, while removing further multipoles largely removes ISW information. This can be seen by noting that the slope of the line Fig. 3.7 becomes negative for $\ell_{\max} > 5$. This has implications for efforts to reconstruct ISW maps from data; if we only care about small-angle features, it can be worth ignoring a few low- ℓ modes in order to get a more accurate reconstruction. Conversely, if we want to study how the ISW signal contributes to the CMB quadrupole and octopole, we must recognize that reconstruction quality will be necessarily less predictable.

Because cosmic variance of the ISW C_ℓ has a nontrivial relationship with the value and scatter of ρ , one cannot make a direct connection between ℓ_{\min} and how f_{sky} affects reconstruction, as is done in the ISW signal-to-noise detection studies (e.g., Ref. [104]). To understand how sky coverage affects reconstruction, one should perform simulations using the mask appropriate for a given survey. We refer the reader to Ref. [99] for an analysis of how ISW signal reconstruction is affected by survey masks.

We also looked at the impact of varying ℓ_{\max} but found that the correlation coefficient ρ is insensitive to it, and therefore do not show it.




Figure 3.8: How changing \bar{n} affects reconstruction of the ISW map. The data points show the mean and standard deviation of ρ , the correlation coefficient between true and reconstructed ISW maps, observed in 10,000 realizations. The line shows the value of $\langle \rho \rangle$ computed analytically.

3.2.4 Varying \bar{n}

Additionally, we studied how the level of galaxy shot noise affects reconstruction. For this test, we varied the number density of sources, \bar{n} , for our fiducial survey and introduced it to both C_ℓ^{true} and C_ℓ^{model} according to Eq. (3.6). Our results are shown in Fig. 3.8.

We find that as long as $\bar{n} \gtrsim 1 \text{ arcmin}^{-2} \approx 10^7 \text{ sr}^{-1}$, shot noise will have a negligible impact on reconstruction. Note that this requirement is easily satisfied by essentially all photometric surveys (e.g., for DES or Euclid, $n \simeq (10 - 30) \text{ arcmin}^{-2}$). However, the quality of the reconstruction degrades rapidly for lower values of number density; once $\bar{n} \lesssim 10^{-3} \text{ arcmin}^{-2} \approx 10^4 \text{ sr}^{-1}$, the reconstruction contains effectively no information about the true ISW map. Therefore, ISW reconstruction from spectroscopic galaxy surveys, as well as galaxy cluster samples, may be subject to degradations due to high shot noise.

3.3 Cross-check with Manzotti and Dodelson (2014)

Here we perform a crosscheck of our reconstruction procedure against Manzotti and Dodelson (2014) [11] (MD). In their paper, MD perform simulations for an NVSS-like survey and a DES-like survey in two- and three-binned configurations. We attempt to simulate ISW reconstruction for similar surveys.




Figure 3.9: Left panel: Redshift distributions of surveys, chosen to match the LSS surveys studied in Ref. [11]. The DES 2 bin case is produced by merging the two highest redshift bins for the DES 3 bin case, and so except for the region around $z \sim 1$, the blue lines are directly behind the red lines. Right panel: Histogram of ρ found for 10,000 simulations of surveys with redshift distributions shown in the left panel. Values of $\bar{\rho}$ from Ref. [11] are shown by the arrows along the top of the plot. The observed discrepancies are likely due to different amounts of simulated galaxy shot noise.

For the NVSS-like survey, we use the analytic dn/dz distribution given by MD, integrating between $0.01 \leq z \leq 6$ when computing its C_ℓ . The redshift distributions used for these simulations are shown in the left panel of Fig. 3.9. For the DES-like survey, we adjusted the parameters in our fiducial dn/dz model by eye so that the three-binned case is similar to that shown in MD’s relevant figure. For the three-binned case, we place bin edges at $z \in [0.1, 0.5, 1.0, 1.6]$. Because MD do not describe how the two-binned case is divided, we somewhat arbitrarily place the bin edges at $z \in [0.1, 0.5, 1.6]$. Like MD, we include multipoles $3 \leq \ell \leq 80$ in our analysis. We leave \bar{n} at our fiducial value of 10^9 for all of these surveys. This value was selected based on an assumption that shot noise contributions would be negligible, but we note below that this is likely not the case.

The right panel of Fig. 3.9 shows a histogram of the ρ values for 10,000 map realizations in our study, with the values from MD shown with arrows. We find that our $\bar{\rho}$ values are systematically higher than, but not wildly incompatible with those in MD. It is hard to specifically identify a cause for this without more information, but the discrepancy is most likely due to differences in the amount of Poisson noise we add to our galaxy maps. We note, for example, that we can get our $\langle \rho \rangle$ for the NVSS-like survey to roughly match the MD value if we reduce our simulation’s \bar{n} to $\sim 5 \times 10^5$. If we set \bar{n} to the value reported for

NVSS by MD, $\bar{n} = 5 \times 10^4 \text{ sr}^{-1} \approx 16 \text{ deg}^{-2}$, we get a lower value of $\langle \rho \rangle = 0.22$.

The shift between the two- and three-bin DES surveys in our simulations is larger than the $\Delta\rho \sim 0.03$ seen in the binning study of Section 3.2.2. This supports our hypothesis that $\bar{\rho}$ shifts more easily at lower ρ values. The fact that our observed shift is still only about half the size of that by MD is probably also due to the fact that we are finding larger $\bar{\rho}$ values than they do. In other words, since we measure higher $\bar{\rho}$ (probably because we include less shot noise than MD), a comparable degradation in reconstruction quality will produce a smaller shift in $\bar{\rho}$.

3.4 The effect of mismodeling redshift distributions

Large-scale structure surveys are subject to a variety of systematic errors that limit the extent to which LSS tracers can be used to probe dark matter, dark energy, and primordial physics. These systematics can be astrophysical, instrumental, or theoretical in origin. Concretely, in this work, they include anything that makes $C_\ell^{\text{model}} \neq C_\ell^{\text{true}}$, which will cause the estimator given in Eq. (3.11) to become suboptimal — both in the sense that it will no longer be a maximum-likelihood estimator for the true ISW signal, that the errors may potentially biasing reconstructed ISW map. Our goal is to study these LSS systematics generally, without requiring specific information about a LSS survey (e.g., wavelengths at which it observes the sky). We do this by considering two broad classes of LSS systematics:

1. Mismodeling of the distribution of LSS sources along the line of sight.
2. Direction-dependent calibration errors.

Our studies will give us some insight into which, and how much, systematics need to be controlled if one wishes to use LSS data to reconstruct a map of the ISW signal. This section will focus on the first class of systematic, while Section 3.4.3 will focus on the second.

In the context of ISW map reconstruction, it would be reasonable to guess that accurate knowledge of galaxy redshifts is important for our ability to correctly associate the observed number density fluctuations on the sky with the three-dimensional gravitational potential fluctuations which source the ISW signal in a redshift-dependent way. Uncertainties about redshift distributions are a pervasive class of systematics affecting LSS surveys, which have already been studied by numerous authors (e.g., Refs. [109, 110]) in the context of cosmological parameter measurements from photometric surveys. Here we study how redshift modeling errors affect the ISW reconstruction accuracy.




Figure 3.10: Impact of mismodeling survey depth on the expected correlation between the true and reconstructed ISW maps $\langle \rho \rangle$ (left panel) and the ratio of the average size of residuals to that of ISW map features $\langle s \rangle$ (right panel). The true value of the parameter z_0 , which controls the depth of the survey, is fixed at 0.7, while the values used for reconstruction are shown on the x -axis. The blue circular points show results from our standard reconstruction pipeline, while the gray diamond-shaped points (directly behind the blue points in the ρ plot) show results when we skip the \bar{b} -fitting step. The y -axis is linear within one tick mark of zero; otherwise, it has logarithmic scaling.

For the purposes of this discussion, we define redshift uncertainties broadly as anything that makes the galaxy window function (Eq. (3.5)) used in our ISW estimator different from that which describes the true line-of-sight distribution of objects we observe on the sky. We study three specific cases of this: the mismodeling of a survey’s median redshift, redshift-dependent bias, and the fraction of catastrophic photometric-redshift errors. In each case, we identify a parameter which controls the survey characteristic in question. Then, choosing a true (simulation) value for that parameter, we perform reconstructions using several mismodeled values as input to the ISW estimator. This allows us to and look at how the theoretical expectation values of our quality statistics respond relative the best, correctly modeled case.

Let us place these shifts in context by referring to previous sections. In an ideal scenario with no systematic errors, changing the survey depth parameter (see Section 3.2.1) from the fiducial $z_0 = 0.7$ to 0.6 (0.8) causes $\langle \rho \rangle$ to change by -3% (+1.5%) and $\langle s \rangle$ by +20% (-10%). Also, splitting our fiducial survey into 6 redshift bins (in Section 3.2.2) improves $\langle \rho \rangle$ by 3% relative to the one-bin case.

3.4.1 Median redshift

We begin by studying how reconstruction accuracy responds when we construct the ISW estimator using the wrong median LSS source redshift. Though the parameter z_0 in the dn/dz distribution given in Eq. (3.10) is lower than z_{median} , raising or lowering it will have




Figure 3.11: Impact of mismodeling redshift-dependent bias on the expected correlation between the true and reconstructed ISW maps $\langle \rho \rangle$ (left panel), and the typical size of residuals relative to that of ISW map features $\langle s \rangle$ (right panel). The bias is modeled as $b(z) = 1 + b_2(1 + z)^2$ with the true value fixed at $b_2 = 0.5$ and the values used in the ISW estimator shown on the x -axis. Both axes have logarithmic scaling except in regions within one tick mark of zero, where they are linear. The blue circular points show results from our standard reconstruction pipeline, while the gray diamond-shaped points (directly behind the blue points in the ρ plot) show results when we skip the \bar{b} -fitting step.

a similar effect as shifting the median of the distribution. We thus use z_0 as a proxy for median redshift. We compute C_ℓ^{true} with z_0 fixed at its fiducial value of 0.7, and vary the z_0 values used to compute C_ℓ^{model} .

Figure 3.10 shows the fractional change in our reconstruction statistics when the value of z_0 used for reconstruction is shifted from its true value by $\pm 1\%$, $\pm 10\%$, $\pm 20\%$, $\pm 30\%$, and $\pm 50\%$. We see that even for large shifts in z_0 (with correspondingly dramatic mismatches between the true and model dn/dz) the fractional change in ρ is less than $\mathcal{O}(10^{-3})$. The effect on s is also small; for all but the most extreme points, the fractional change in the size of residuals $\langle s \rangle$ is less than 10%.

To understand this lack of sensitivity of z_0 , it is instructive to note that varying z_0 changes C_ℓ by a nearly scale-independent amplitude, as can be seen in the power spectrum visualizations in Fig. 3.13. As we observed in Section 3.1.5, ρ , the correlation coefficient between true and reconstructed ISW maps, is insensitive to overall shifts in the map amplitude. The fact that it does not respond strongly to these changes in z_0 is thus not surprising. The statistic $\langle s \rangle$, which measures the size of residuals, is sensitive to changes in amplitude, however. The fact that it also displays small fractional changes illustrates the importance of the bias-fitting procedure described in Section 3.1.4. Because the effects of mismodeling z_0 are degenerate with shifts in constant bias, fitting for \bar{b} protects our reconstruction against this kind of systematic.

For comparison, we compute $\langle \rho \rangle$ and $\langle s \rangle$ while neglecting the bias-fitting step and show

the results as gray points in Fig. 3.10. We see no change in the ρ plot (the gray points are directly behind the blue ones), reflecting the fact that ρ is insensitive to constant multipliers. In the s plot, we see that the bias-fitting procedure suppresses the size of the reconstruction errors by about an order of magnitude.

To summarize, we find that the quality of the ISW reconstruction is much less dependent on our knowledge of the survey’s median redshift than naively expected. The median redshift mostly changes the normalization of the C_ℓ , but so does the galaxy bias (which, recall, is to a good approximation scale independent at the large scales we are studying). By fitting for the bias parameter in the angular power spectrum—something that is typically done in LSS surveys regardless of their application—one effectively also fits for z_0 . As a result, the combination of the galaxy bias and survey depth that enters the amplitude of the C_ℓ is fit to the correct value.

3.4.2 Redshift-dependent bias

Here, we study what happens if the redshift dependence of the galaxy bias is modeled incorrectly. Using the functional forms for the bias of high-redshift tracers given in Ref. [91] for guidance, we parametrize the redshift dependence of the bias via

$$b(z) = b_0(1 + b_2(1 + z)^2). \quad (3.28)$$

For this study, we set $b_0 = 1$ and vary b_2 , noting that Ref [91] uses $b_2 \sim 0.5$ for sources in NVSS and WISE-AGN.

In the expression for C_ℓ , $b(z)$ appears inside the same integrand as dn/dz , so changes to $b(z)$ have an effect similar to altering the LSS source redshift distribution. The results here, shown in Fig. 3.11, are thus similar to what was seen in the previous section. Increasing b_2 mostly just increases the overall amplitude of the galaxy C_ℓ ’s, so the reconstruction is not very sensitive to b_2 once we fit for \bar{b} . For example, if the true value of b_2 is 0.5 and we reconstruct the ISW signal assuming no redshift dependence ($b_2 = 0$), the fractional change in $\langle\rho\rangle$ is $\mathcal{O}(10^{-4})$ and the fractional change in $\langle s\rangle$ is $\mathcal{O}(10^{-2})$. The reason the \bar{b} -fitting step has a larger effect here than in the z_0 study above is probably because the normalization requirements of dn/dz somewhat limit the size of C_ℓ amplitude shifts, whereas $b(z)$ has no such normalization scaling.

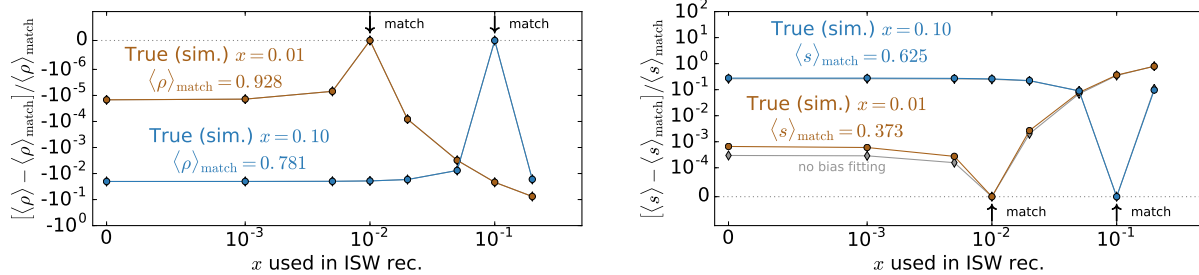


Figure 3.12: Impact of mismodeling the fraction x of galaxies subject to catastrophic photo- z errors on the expected correlation between the true and reconstructed ISW maps $\langle \rho \rangle$ (left panel), and the typical size of residuals relative to that of ISW map features $\langle s \rangle$ (right panel). Both axes have logarithmic scaling except in regions within one tick mark of zero, where they are linear. The blue and brown circular points show results from our standard reconstruction pipeline when the true value of x is 0.1 and 0.01 respectively. The gray diamond-shaped points (directly behind the other points in the ρ plot and the blue points in the s plot) show results when we skip the *barb*-fitting step.

3.4.3 Catastrophic photo- z error rate

Galaxies in photometric-redshift surveys are also subject to so-called *catastrophic* photometric-redshift errors—cases where the true redshift is misestimated by a significant amount [110, 111]. This is a distinct effect from the photo- z uncertainty modeled in the binning tests in Section 3.2.2, which causes a redshift bin selected using sharp cuts in photo- z to occupy a smoothed distribution in the spectroscopic redshift. Rather, for galaxies suffering catastrophic photo- z errors, the photometric-redshift finding algorithms have failed, and the spectroscopic redshift corresponding to a given photo- z is effectively randomized. The reasons for this are not fully understood, but, like the conventional photo- z error case, the rate and outcome of catastrophic errors depend strongly on the number of photometric filters and their relation to the spectral features that carry principal information about the redshift.

In the absence of detailed, survey-specific information about the photometric pipeline, we model catastrophic redshift errors by randomly assigning the true redshift of a fraction x of the galaxies in our sample (e.g., $x = 0.01$ means that one in a hundred galaxies has a catastrophic photo- z error). We implement this by modifying the redshift distribution of each bin i to

$$\frac{d\tilde{n}^i}{dz} = (1 - x) \frac{dn^i}{dz} + x \bar{n}^i [\Theta(z - z_{\min}) - \Theta(z_{\max} - z)], \quad (3.29)$$

where x is the fraction of galaxies suffering catastrophic errors, dn^i/dz is the redshift distribution of bin i without catastrophic errors, and Θ is the Heaviside step function. The added

term on the right models the fact that, of the \bar{n}^i galaxies assigned to that photometric-redshift bin, $x\bar{n}^i$ of them have spectroscopic redshifts which are randomized across the full range of the survey. For our analysis, we choose the range of these randomized redshifts to be $z \in [z_{\min}, z_{\max}] = [0.01, 2.5]$. In practice, we significantly smooth the edges of the step function to avoid numerical artifacts in our C_ℓ calculations.

For this study, we use two different true (simulation) catastrophic photo- z fractions: $x = 0.01$ and 0.1 ; these values roughly bracket the currently achieved levels of catastrophic outliers in current surveys (e.g., CFHTLenS [112]). Figure 3.12 shows the fractional change in $\langle\rho\rangle$ and $\langle s\rangle$ when the ISW estimator is constructed assuming various values of x , with true $x = 0.01$ and $x = 0.1$ shown in blue and brown lines, respectively.

Our results show us two things. First, though mismodeling x results in more significant changes than what was seen for the survey depth and redshift-dependent bias, the shifts are still relatively small; in the worst-case scenarios, $\langle\rho\rangle$ shifts by less than 10% and $\langle s\rangle$ shifts by about 20%. Second, the constant-bias-fitting step of our pipeline does not provide protection against mismodeled catastrophic photo- z error rates. This is because the dn/dz modification in Eq. (3.29) alters C_ℓ in a scale-dependent way, as can be seen in Fig. 3.13.

To check whether catastrophic photo- z errors are more damaging when LSS data are binned in redshift, we ran a similar analysis for a case where the fiducial dn/dz was split into three redshift bins. We observed fractional changes in the quality statistics similar to those seen for the one-bin case, so we conclude that our results are roughly independent of the binning strategy.

In summary, we find that properly modeling a survey’s catastrophic photo- z error fraction is more important for preserving ISW reconstruction quality than either its depth or redshift-dependent bias but that, overall, reconstruction is relatively robust against these kinds of errors.

3.4.4 Interpretation: Impact of mismodeling dn/dz on C_ℓ

Figure 3.13 shows how galaxy-galaxy and galaxy-ISW power spectra respond to changes in the parameters discussed in Section 3.4. We study the effect of survey depth by shifting the parameter z_0 in Eq. (3.10), redshift dependence of bias by changing b_2 in Eq. (3.28), and the fraction of galaxies x subject to catastrophic photometric-redshift errors via Eq. (3.29).

We see that changing z_0 and b_2 shifts C_ℓ by a mostly scale-independent factor. As noted in Section 3.4, this is why systematics related to mismodeling depth and bias redshift dependence have only a small effect on ISW reconstruction quality. It is also why fitting for




Figure 3.13: The change in the galaxy angular power spectrum $C_\ell \equiv C_\ell^{\text{gal-gal}}$ in response to (left to right) changes in survey depth, characterized via z_0 in Eq. (3.10); the redshift dependence of bias, modeled by varying the b_2 parameter in $b(z) = 1 + b_2(1+z)^2$; and the fraction x of galaxies subject to catastrophic photo- z errors. These plots show the ratio of galaxy autopower relative to that of a reference survey.

scale-independent bias \bar{b} via

$$C_\ell^{\text{gal(obs)}} = \bar{b}^2 C_\ell^{\text{gal(model)}}, \quad (3.30)$$

as can be seen in Figs. 3.10 and 3.11, provides some protection against these systematics.

In contrast, changing the catastrophic photo- z fraction x by more than about 0.01 significantly changes the low- ℓ shape of C_ℓ . This explains why mismodeling x has a relatively larger (though still small) impact on ISW reconstruction quality and why constant bias fitting does not mitigate this effect as much.

3.5 The effect of photometric calibration errors

Photometric calibration errors are a very general class of systematics that cause the magnitude limit of a survey to vary across the sky. This introduces direction-dependent number density variations which do not correspond to fluctuations in physical matter density, thus biasing the observed galaxy power spectrum. Examples of photometric calibration errors include atmospheric blurring, unaccounted-for Galactic dust, and imperfect star-galaxy separation, among other things. A number of recent LSS observations have found a significant excess of power at large scales [113–118], suggesting the presence of this kind of error.

We adopt a parametrization of calibration errors from Ref. [119], who presented a systematic study of the effects of calibration errors and requirements on their control for cosmological parameter estimates. See also Refs. [120–122] for other approaches. We model photometric calibration errors in terms of a calibration error field $c(\hat{\mathbf{n}})$ which modifies the

observed number density N^{obs} via

$$N^{\text{obs}}(\hat{\mathbf{n}}) = (1 + c(\hat{\mathbf{n}})) N(\hat{\mathbf{n}}). \quad (3.31)$$

This kind of direction-dependent “screen” is straightforward to implement on the level of maps but complicates the process of computing the theoretical expectation value for our statistics, $\langle \rho \rangle$ and $\langle s \rangle$. Because multiplicative effects introduce mixing between spherical components of the galaxy maps, there is a nontrivial relationship between the power spectra for the true galaxy distribution, the observed galaxy distribution, and the calibration error field $c(\hat{\mathbf{n}})$. (See, for example, Refs. [119, 122].) To make calculations tractable, we use the fact that calibration error effects will be dominated by additive contributions at large angular scales and estimate

$$[C_\ell^{XY}]^{\text{obs}} \approx \frac{C_\ell^{XY} + C_\ell^{\text{cal}XY} - \delta_{\ell 0} c_{00}^X c_{00}^Y}{(1 + c_{00}^X/\sqrt{4\pi})(1 + c_{00}^Y/\sqrt{4\pi})}. \quad (3.32)$$

Here, $C_\ell^{\text{cal}XY}$ is the cross-power between calibration error fields affecting maps X and Y . The $c_{00}^X \equiv (C_{\ell=0}^{\text{cal}X})^{1/2}$ terms are their monopoles, which contribute by shifting \bar{n}^X . We derive this expression in Appendix B.

Note that this modification is only applied to C_ℓ^{true} . We wish to study the impact of calibration errors which not accounted for in our model, so we will always (when analyzing simulations or calculating quality statistic expectation values) compute C_ℓ^{model} without including calibration error effects.

For this analysis, we adopt a functional form for the calibration error field power spectrum,

$$C_\ell^{\text{cal}} = \begin{cases} \alpha^{\text{cal}} \exp[-(\ell/10)^2] & \text{if } \ell \leq 30 \\ 0 & \text{otherwise} \end{cases} \quad (3.33)$$

where α^{cal} is a normalization constant set to fix the variance of $c(\hat{\mathbf{n}})$ to a desired value. The variance is given by

$$\text{var}[c] \equiv \langle c^2(\hat{\mathbf{n}}) \rangle = (4\pi)^{-1} \sum_\ell (2\ell + 1) C_\ell^{\text{cal}}. \quad (3.34)$$

The form of Eq. (3.33) is inspired by power spectrum estimates for maps of dust extinction corrections and magnitude limit variations in existing surveys. (See Figs. 5 and 6 in Ref. [119]) Using this power spectrum, we generate independent Gaussian realizations of $c(\hat{\mathbf{n}})$ which are then combined with our simulated galaxy maps according to Eq. (3.31). (Though the phases of the calibration errors due to real dust maps will not be random, simulating

errors as Gaussian should still give us a reasonably good indication of the impact of calibration errors with a given variance on ISW reconstruction.) These postprocessed maps are used as input for ISW reconstruction.

3.5.1 Context: Current and future levels of calibration error

To put our results in context, it is useful to identify what values of variance in the calibration field $\text{var}[c]$ are expected from current and future surveys. Here we emphasize that we are talking about *residual* calibration errors—that is, calibration errors which are not properly corrected for and thus can cause biases in cosmological inferences.

Above, we defined these errors in terms of variations in the number of observed galaxies. To relate this to variations in a survey’s limiting magnitude, we must multiply the magnitude variations by a factor of $\ln(10)s(z)$, where $s(z) \equiv d \log_{10} N/dm|_{m_{\text{lim}}}$ is the survey-dependent faint-end slope of the luminosity function; see Eq. (30) in Ref. [119]. We adopt $s(z) \simeq 0.3$ estimated from the simulations of Ref. [123], assuming a median galaxy redshift $z \sim 0.75$. This means that the conversion factor is $\ln(10)s(z) \sim 1$, and variance in calibration is roughly equal to that in the limiting magnitude, $c(\hat{\mathbf{n}}) \equiv (\delta N/N)(\hat{\mathbf{n}}) \simeq (\delta m)_{\text{lim}}$.

With these assumptions, the smallest currently achievable variance of the calibration error $c(\hat{\mathbf{n}})$ is of order $\text{var}[c] \sim 10^{-3}$ (e.g., Fig. 14 in Ref. [120]). For example, residual limiting magnitude variations in the SDSS DR8 survey are at the level of 0.03 mag [124], again implying that $\text{var}[c] \simeq 10^{-3}$. Note that, while the impressive SDSS “uber-calibration” to 1% [125] would imply an order of magnitude smaller variance, this might be difficult to achieve in practice because there are sources of calibration error that come from the analysis of the survey and are not addressed in the original survey calibration. We show the current levels of residual calibration errors value as a blue vertical band in Fig. 3.14, spanning a range between the optimistic level associated with the SDSS uber calibration to the more conservative $\text{var}[c] = 10^{-3}$.

In the same figure, we also show the *future* control of calibration errors required to ensure that they do not contribute appreciably to cosmological parameter errors—e.g., those in dark energy and primordial non-Gaussianity. This range, forecasted assuming final DES data and adopted from Ref. [119], is shown as a green band spanning $\text{var}[c] \sim 10^{-6}$ – 10^{-5} . The lower bound is set by the requirement that the bias to cosmological parameter estimates be smaller than their projected errors, while 10^{-5} is chosen as an intermediate value between that and $\text{var}[c] = 10^{-4}$, which introduces unacceptable levels of bias. (See Fig. 4 of Ref. [119].) These should be viewed as only rough projections, as the precise requirements depend on the faint-

end slope $s(z)$ of the source luminosity function, the cosmological parameters in question, and the shape of the calibration field’s power spectrum C_ℓ^{cal} .

3.5.2 Results for ISW reconstruction

We find that even small levels of calibration error can have a significant impact on ISW reconstruction quality — moreso than any of the other systematics studied in this chapter. Figure 3.14 shows how the correlation between true and reconstructed maps, ρ , and the reconstructed map residuals, s , respond to different levels of calibration error.

Reconstruction quality starts to degrade when $\text{var}[c] \sim 10^{-6}$, which roughly corresponds to the same 0.1% magnitude calibration required to achieve cosmic-variance-limited ISW detection [100]. At this level, we see ρ begin to move away from its best-case (no calibration error) value and the s plot shows that residuals are comparable in amplitude to fluctuations due the true ISW signal.

Once the calibration error power starts to dominate over the galaxy autopower, occurring around $\text{var}[c] \sim 10^{-4}$, the reconstruction contains little information about the true ISW signal. Here, the scatter in ρ overlaps with zero and we see that the reconstructed map residuals approach a constant value. We present an explanation in Appendix C for why we expect this to occur.

Comparing these numbers to the shaded bands, we see that, with current levels of calibration error control, we have little hope of accurately reconstructing the ISW signal with galaxy survey data alone. Encouragingly, though, the levels of control required to obtain unbiased cosmological parameter estimates from next-generation surveys [119] are precisely the levels needed for accurate ISW reconstruction.

We note that the additive-error-only theory calculations show good agreement with our results from simulations, and so can be useful as a computationally efficient indicator of when calibration errors become important. In light of this, we also computed $\langle \rho \rangle$ and $\langle s \rangle$ using a power law spectrum, $C_\ell^{\text{cal}} \propto \ell^{-2}$, in order to check how sensitive our results are to the shape of the calibration error field’s power spectrum. This more sharply peaked spectrum caused reconstruction quality to start degrading at a slightly smaller $\text{var}[c]$ compared to the Gaussian model, but otherwise showed similar results. This can likely be explained by the fact that the power law C_ℓ^{cal} reaches higher values at low ℓ for a given field variance, which means it can start dominating over true galaxy power at those multipoles earlier.

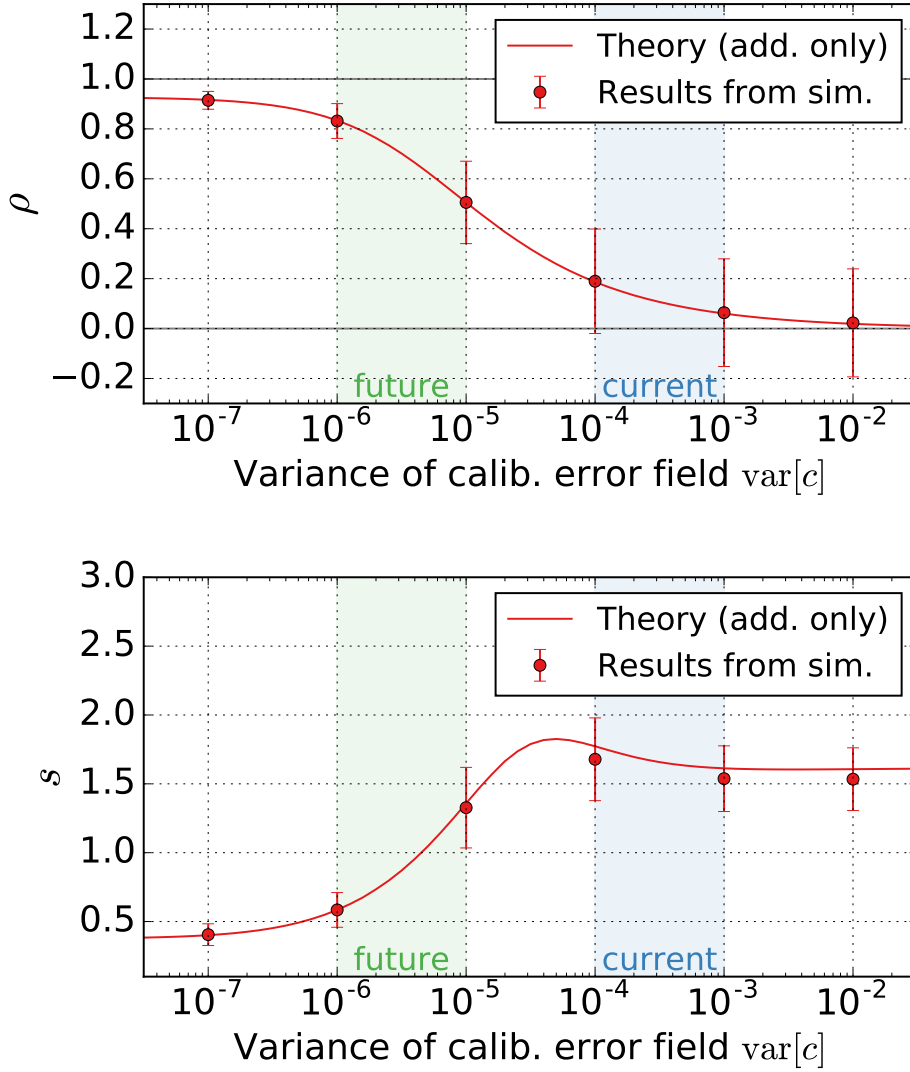


Figure 3.14: The effect of photometric calibration errors on reconstruction quality. We show results for the correlation coefficient between true and reconstructed ISW maps (top panel) and for the typical size of map residuals relative to the variance of the true ISW map (bottom panel). The lines show the expectation from theory, considering only additive contributions from calibration errors, while the data points show the mean and standard deviation from 10,000 simulated map realizations. The shaded regions show the current and projected levels of control over residual calibration errors discussed in Section 3.5.1.




Figure 3.15: Exploration of whether raising ℓ_{\min} can mitigate the impact of photometric calibration errors on ISW signal reconstruction. The top panel shows the mean and standard deviation of ρ , the correlation between the true and reconstructed ISW maps, measured from 10,000 simulations. The bottom panel shows the fractional change in ρ relative to the case with no calibration errors. Points for different values of ℓ_{\min} are staggered so that the errors bars are legible; each cluster of three points shares the same value of $\text{var}[c]$.

3.5.2.1 Mitigation by raising ℓ_{\min}

Because calibration error fields tend to have the most power on large scales, we looked at whether raising ℓ_{\min} can mitigate their impact. Our results, shown in Fig. 3.15, show that raising ℓ_{\min} from 2 to 3 or 5 causes the error bars denoting the scatter in ρ to cross zero at a higher value of $\text{var}[c]$. However, this effect is small, and we conclude that raising ℓ_{\min} provides only limited protection against calibration errors.

3.6 Implications for cosmic alignments

Over the past 15 years, as the full-sky CMB maps provided by the WMAP and Planck experiments became available, increasing evidence has been found for anomalies at large angular scales. In particular, angular correlations at scales above 60 deg on the sky seem to be

missing, while the quadrupole and octopole moment of the CMB anisotropy are aligned both mutually and with the geometry and the direction of motion of Solar System. The origin for the anomalies is not well understood at this time; they could be caused by astrophysical systematic errors or foregrounds or cosmological causes (e.g. departures from simple inflationary scenarios), or they could be a statistical fluctuation, albeit a very unlikely one. The anomalies have most recently been reviewed in Ref. [94].

Some authors [95, 96] have commented on the fact that current efforts to “peel off” the ISW contribution from the CMB maps indicate that the significance of some CMB anomalies is “significantly reduced” once the ISW contribution is subtracted. If true, this statement implies that the observed anomalies are either due to features in the ISW map or caused by an accidental alignment of the early- and late-time CMB anisotropy [126]. In any case, statements on how the primordial and late CMB combine to produce the anomalies clearly depend on the fidelity of the reconstructed ISW contribution to the CMB, which is the subject of our work.

Our goal here is not to carry out a full investigation of the ISW map reconstruction’s effect on the anomalies’ significance. Instead, we would like to simply build intuition on how much imperfect reconstruction affects inferences about the anomalies.

To that end, we pose the following question: if we *assume* for the moment that an ISW map reconstructed using available LSS data happens to show a significant quadrupole-octopole alignment, what is the likelihood that the true ISW map is actually aligned? Note that we in no way imply that the ISW-only alignment scenario is a favored model for the observed CMB anomalies. We simply want to study how robust certain properties of the ISW map, particularly the phase structure of the anisotropies in the map, are to the reconstruction process.

To study the alignments, we adopt the (normalized) angular momentum dispersion maximized over directions on the sky, defined as [127, 128]

$$(\Delta L)_{2+3,\text{true}}^2 \equiv \max_{\hat{\mathbf{n}}} \left(\frac{\sum_{m=-\ell}^{\ell} m^2 |a_{\ell m}(\hat{\mathbf{n}})|^2}{\ell^2 \sum_{m=-\ell}^{\ell} |a_{\ell m}(\hat{\mathbf{n}})|^2} \right) \quad (3.35)$$

where $a_{\ell m}(\hat{\mathbf{n}})$ are expansion coefficients of the map in a coordinate system where the z -axis is in the $\hat{\mathbf{n}}$ direction. Hence, the maximization is performed over all directions $\hat{\mathbf{n}}$; note that only the numerator of the expression in angular parentheses depends on the direction, and see Section 5.6 of Ref. [128] for the algorithm to efficiently compute the maximization. Intuitively, high values of the angular momentum indicate significant planarity of the $\ell = 2$ and $\ell = 3$ modes as well as their mutual alignment.

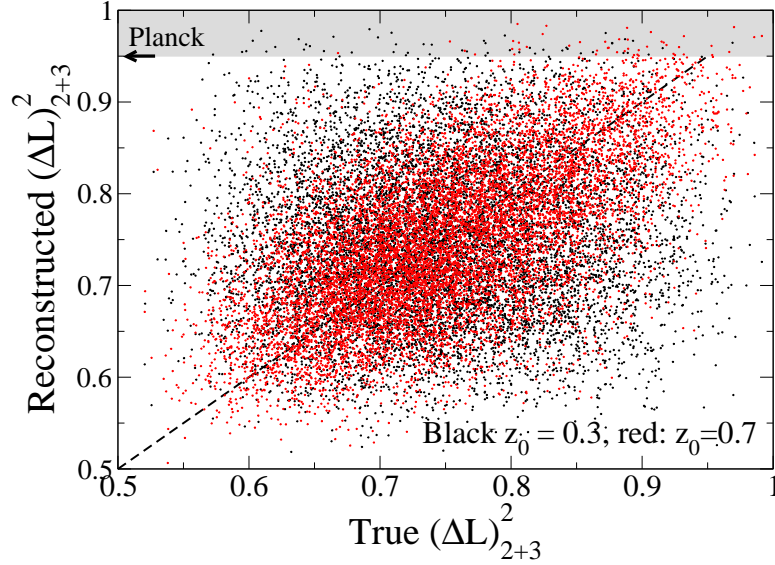


Figure 3.16: The relationship between the true (x -axis) and reconstructed (y -axis) angular momentum dispersion $(\Delta L)_{2+3}^2$, defined in the text, for the combined quadrupole and octopole in 10,000 randomly generated ISW maps. Results are shown for two alternate survey depths: our fiducial LSS survey with $z_0 = 0.7$ (red points) and $z_0 = 0.3$ (black points), which have correlation coefficients 0.58 and 0.11 respectively. The gray region denotes $(\Delta L)_{2+3}^2$ as high or higher than measured in WMAP and Planck CMB maps, while the diagonal line is where the true and reconstructed values match. See the text for details.

We set up the following pipeline:

- Start with 10,000 random realizations of the true ISW map and the corresponding LSS maps (so that each LSS map contains gravitational potential field that produces the corresponding ISW map).
- For each true ISW map, measure the angular momentum dispersion $(\Delta L)_{2+3,\text{true}}^2$ defined in Eq. (3.35).
- Reconstruct each map assuming a fiducial LSS survey and repeat the calculation to get a set of $(\Delta L)_{2+3,\text{rec}}^2$.
- Make a scatter plot of $(\Delta L)_{2+3,\text{rec}}^2$ vs $(\Delta L)_{2+3,\text{true}}^2$, which will show how much and in which direction reconstruction biases the alignment information.

The results are summarized in Fig. 3.16. There we show how the inferred angular momentum dispersion of the combined quadrupole and octopole is affected by reconstruction

for 10,000 randomly generated ISW maps. The x -axis shows the value for the true ISW map, while the y -axis shows values reconstructed from our fiducial LSS survey at two alternate depths, $z_0 = 0.7$ (red points) and 0.3 (black points). We find that the true and reconstructed angular momentum dispersions are not very correlated, having a correlation coefficient of only 0.58 for $z_0 = 0.7$ and 0.11 for $z_0 = 0.3$.

We also denote the value for the angular momentum dispersion of the WMAP/Planck full map, which includes both primordial and late-time ISW contributions, at $(\Delta L)_{2+3}^2 \simeq 0.95$. (The precise value varies slightly depending on the map. [94, 128]) Of the $z_0 = 0.7$ (0.3) reconstructed maps which have $(\Delta L)_{2+3}^2$ as high as or higher than the WMAP and Planck CMB maps (points falling in the shaded gray region), only $7/70 = 10\%$ ($1/70 = 2\%$) have corresponding true maps which satisfy the same high angular momentum dispersion criterion.

Investigating the implications of the ISW reconstruction on the inferences about the alignments of primordial-only and ISW-only maps in depth is beyond the scope of this analysis. Nevertheless, our simple test indicates that at least the quadrupole-octopole alignment *in the ISW-only maps* is not very robust under ISW reconstruction using realistic LSS maps, even without taking into account calibration and other systematic errors. This is because the filter in Eq. (3.11), which is a maximum likelihood estimator for the ISW signal map given input LSS maps, has significant variance in its reconstruction of individual modes.

3.7 Conclusion

In this work we use simulated ISW and LSS maps to study the accuracy of ISW signal reconstructions performed using LSS data as input. In particular, we study how systematics associated with galaxy surveys affect the ISW map reconstruction. We measure reconstruction accuracy using two quality statistics: ρ , the correlation coefficient between the true and reconstructed ISW maps, and s , the rms error in the reconstructed map relative to the rms of true ISW map features.

In the absence of systematics, we find that increasing survey depth improves these statistics (brings ρ closer to 1 and lowers s), though the shifts in their average values are small compared to their scatter. Similarly, splitting the survey data into redshift bins leads to moderate improvement. The reconstruction quality improvement due to increasing survey depth by $\Delta z = 0.1$ is comparable to that gained by splitting into three redshift bins: both lead to improvement $\Delta\bar{\rho} \sim 0.02$, or $\Delta\bar{\rho}/\bar{\rho} \sim 2\%$. We also find that reconstruction can be slightly improved if we are willing to neglect the reconstruction of very low- ℓ multipoles; increasing our fiducial $\ell_{\min} = 2$ to 5 results in $\Delta\bar{\rho} \sim 0.01$ and a reduction in the scatter of ρ

by about a factor of 2. Last, we find that galaxy shot noise has a negligible impact as long as $\bar{n} \gtrsim 1 \text{ arcmin}^{-2} \approx 10^7 \text{ sr}^{-1}$. These results provided a baseline comparison for our studies of systematics.

The first class of systematics we study are those associated with mismodeling the line-of-sight distribution of LSS sources. By examining what happens to reconstruction quality when different galaxy window functions are used for the ISW-estimator input C_ℓ^{model} than for the simulation-generating C_ℓ^{true} , we find that ISW signal reconstruction is robust against these kinds of errors. We study the mismodeling of survey depth and redshift-dependent bias and find that fractional shifts in $\langle \rho \rangle$ are less than $\mathcal{O}(10^{-4})$ for all but the most extreme cases. Inaccurately estimating the fraction of catastrophic photo- z errors results in a larger shift, which depends on the true fraction, but at worst this degrades $\langle \rho \rangle$ by about a percent. Reconstruction quality is likely to be similarly insensitive to other direction-independent modeling uncertainties; for example, the choice of cosmological parameter values and maybe models of modified gravity.

The fact that we fit data for a constant galaxy bias is the key to this robustness. This is because the modeling errors discussed above change the galaxy spectrum by a mostly scale-independent amplitude which is degenerate with a shift in constant bias \bar{b} . Thus, the more a given systematic changes the shape (rather than amplitude) of galaxy C_ℓ , the more of an impact it will have on ISW signal reconstruction.

We find that photometric calibration errors are by far the most important systematic to control if one wants to construct a map of the ISW signal from LSS data. For the reconstructed ISW map to contain accurate information about the true ISW signal, calibration-based variations in number density must be controlled so that the calibration error field c , defined via $N^{\text{obs}}(\hat{\mathbf{n}}) = (1 + c(\hat{\mathbf{n}}))N(\hat{\mathbf{n}})$, has a variance less than 10^{-4} . Even at that level, which is optimistic for current surveys, the reconstruction quality is significantly degraded compared to the case with no systematics. For the model we studied, in order to keep that degradation smaller than $\mathcal{O}(10\%)$, calibration errors must be controlled so that $\text{var}[c] \lesssim 10^{-6}$. This is a similar level to what is required to avoid biasing cosmological parameter estimates made with future survey data. Prospects for mitigation of these effects by neglecting low ℓ multipoles are limited.

Our initial motivation for studying the accuracy of reconstructed ISW signal maps was to assess whether they can be used to determine whether certain large-angle CMB anomalies are sourced by features in the ISW signal. Given this, we additionally considered, even in the case of a good overall signal reconstruction, how well the standard ISW estimator (Eq. (3.11)) is able to reconstruct the ISW contribution to the alignment between the CMB quadrupole

and octopole. We do this by comparing the level of alignment, parametrized in terms of the angular momentum dispersion of the $\ell = 2, 3$ modes of true and reconstructed ISW maps. We find that, even in the absence of systematics, the angular momentum dispersion, and thus amount of quadrupole-octopole alignment, was only weakly correlated between these maps. For example, the values of true and reconstructed angular momentum dispersion had a correlation coefficient of only 0.58 for our fiducial survey. Therefore, precise alignments of structures in the ISW map cannot reliably be recovered, given the standard signal estimator with only LSS data as input. This means that comparing the significance of alignments in the original CMB temperature map to that of a map with the reconstructed ISW signal subtracted is not a reliable way to determine whether anomalous large angle alignments are due to features imprinted by the ISW effect.

These results have implications for current and future attempts to reconstruct the ISW signal. Most significantly, they tell us that understanding the level and properties of residual calibration errors in LSS maps is vital to assessing the accuracy of reconstructions made using those maps as input. Given the current levels of calibration error control, at face value our results would seem to imply that reconstruction using existing data is hopeless. Thus, a productive avenue for future work would be to modify the ISW reconstruction pipeline to make it more robust against calibration errors, by including them in the ISW estimator’s noise modeling or by some other method. Since the presence of uncorrected calibration errors will cause one to underestimate galaxy-galaxy noise, it would also be worth turning a critical eye toward how calibration uncertainties affect the evaluation of ISW detections’ signal to noise.

We note that using multiple cross-correlated LSS data sets—which map the same potential fluctuations but are presumably subject to different systematics—will mitigate the impact of calibration errors, as will combining LSS maps with CMB temperature and polarization data. The results of the binning test in Section 3.2.2 provide provisional evidence for this, though for that study it is not possible to disentangle the effects of noise mitigation from those of adding tomographic information. An interesting extension to this work would thus be to explore in more detail whether and to what extent using multiple LSS maps protects ISW reconstruction against calibration errors. Studying the combination of multiple surveys introduces a number of new questions: one might study, for example, how the strength of correlation between galaxy maps influences the improvement in reconstruction due to their combination, or what happens when calibration errors for multiple maps are correlated. These questions were explored further in a follow-up study, published as Ref. [9].

Chapter 4

The covariance of large-angle CMB anomalies

In this chapter we will examine in more detail the properties of the large-angle CMB anomalies which served as a motivation for the ISW reconstruction study in the previous chapter. To put this discussion into context, it is important to note that the spectacular maps of the cosmic microwave background (CMB) anisotropy that have been made over the past few decades have revolutionized our understanding of the universe, and rejuvenated efforts to test fundamental processes in the early and late universe. As we described in Chapters 2 and 3, the CMB maps are overall in a very good agreement with the six-parameter spatially flat Λ CDM model specified by the energy densities of dark matter and baryons, the amplitude and spectral index of primordial scalar fluctuations, the reionization optical depth, and the expansion rate (Hubble constant) [129]. Shortly after the WMAP experiment's data were released, however, several surprising coincidences were noticed on large angular scales. In particular, the WMAP maps of temperature anisotropies exhibit low variance, a lack of correlation on the largest angular scales, alignment between various low multipole moments [127], alignment between those low multipole moments and the motion and geometry of the Solar System [130], a hemispherical power asymmetry [131], a preference for odd parity modes [132], and an unexpectedly large cold spot in the Southern hemisphere [133]. Planck data [134] largely confirmed the presence of these features. For a review of the CMB anomalies, see [94].

While these large-angle CMB anomalies have remained an active area of study over the years, it is difficult to draw firm conclusions from the study of features of the CMB at very large angles, mainly due to the significant cosmic variance at those scales. Moreover, the *a posteriori* nature of their observation, as well as the generally good fit of data to

the standard cosmological model, means that large-angle anomalies do not in themselves provide compelling evidence for beyond- Λ CDM physics [135]. Rather, in a time where nearly all cosmological observations have been in remarkable agreement with the predictions of Λ CDM, the statistically unlikely large-angle features have attracted attention because of the tantalizing possibility that one or some of them might have cosmological origins [136–139]. If that were the case, due to e.g. an isotropy-breaking mechanism in the early universe, the feature in question could provide insight into the physics of inflation.

It is important, however, to consider other explanations for anomalous large-scale CMB features: they could be artifacts of instrumental or astrophysical systematics, or they could simply be unlikely fluctuations in the standard isotropic model. Much of the study of large-angle anomalies has thus been focused on disentangling these three logical possibilities: whether large-angle CMB anomalies are cosmological, or are due to systematic or statistical flukes. Better understanding of the anomalies in the future will be driven by observations of new quantities on very large spatial scales, such as CMB polarization [140–144] and lensing [145], and large-scale structure [146]. Whether or not new insights about the early universe become readily apparent, studying large-angle anomalies has and will continue to provide us with an opportunity to build a deeper understanding of our measurements of the large-angle CMB.

One largely unanswered question is how the observed anomalies are related to one another. If we observe one unlikely feature, does that make us less surprised to find another? Roughly speaking, (positively) correlated anomalies imply a smaller overall joint significance than if they are uncorrelated. Full understanding of the anomalies thus enables an accurate accounting of the likelihood for the joint observation of unexpected features. While such a knowledge of the joint likelihood would not remove the nature of the *a posteriori* choice for the statistics to investigate and observe, it would effectively prevent double-counting (or, if the anomalies are anti-correlated, under-counting) of the anomaly statistics. Furthermore, the anomaly covariance enables quantifying of the anomaly ‘atoms’, i.e. a set of mutually independent features out of which all anomalous observations can be derived [94].

Previous work on the covariance of CMB temperature anisotropy anomalies has mostly been limited to studying pairs of anomalies. For example in Ref. [147], the authors show that missing power at large scales quantified by $S_{1/2}$ and the quadrupole-octopole alignment are not correlated in Λ CDM (such a conclusion was also reached, albeit for the full-sky-only analysis, by Ref. [148]). The lack of correlation between hemispherical power asymmetry and the quadrupole-octopole alignment in Λ CDM are demonstrated in Ref. [149]. In Ref. [150], the authors indicate a possible connection between the lack of power and the odd-multipole

preference anomaly. In particular, they claim — based on an analytical argument — that the odd-multipole preference can be a phenomenological cause of the lack of large-angle correlation. Additionally, in Ref. [151] the authors find a correlation between the low large-angle power and the low value of the CMB quadrupole. Ref. [152] explores the relationship between low power at large angles and the amplitude of the CMB quadrupole and octopole, while Ref. [153] studies its relationship with the quadrupole and octopole phases. Here we aim to take a more global view by studying the relationship between all of these features simultaneously.

In this work, we will use ensembles of simulated CMB temperature maps to empirically characterize, in the context of Λ CDM, the covariance between a collection of features associated with commonly-studied large-angle anomalies. Our analysis proceeds in three general steps. First, we will measure the quantities associated with those features and confirm that the comparison between our measurements of the real CMB sky and simulations reproduce previous findings about the real sky’s anomalousness. Next, we will study the distribution of the simulation ensembles in the space defined by the “anomaly feature” quantities to find their covariances. In doing so, we will investigate the impact of foregrounds and survey properties by comparing the results obtained from simple Gaussian simulations of the CMB temperature map to the more realistic Planck Full Focal Plane simulations [154]. Finally, we will use the measured feature covariances perform a principal component analysis in order to further characterize the ways in which large-angle CMB map properties are expected to vary in Λ CDM, and in which the observed CMB sky is unusual. We emphasize that the goal of this analysis is to gain a deeper understanding, rather than to do any explicit model-comparison.

The rest of the paper is organized as follows. In Section 4.1 we introduce our methods in detail, including the description of maps and masks adopted, of our simulation ensembles and of our power spectrum measurements. In Section 4.2, we outline the statistical description of the eight large-scale features that we study in this work while reporting the statistical significance of the features with respect to our simulation ensembles. The main results of our work — the measurement of covariances between the anomalies — is presented in Section 4.3, along with discussion of a PCA analysis. We summarize and conclude in Section 4.4.

4.1 Methods

We begin by introducing basic terminology and notation that describes CMB anisotropy. Temperature fluctuations can be expanded in a harmonic series as

$$T(\hat{\mathbf{n}}) = \sum_{\ell} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}), \quad (4.1)$$

where $\hat{\mathbf{n}}$ is the direction on the sky and the complex coefficients $a_{\ell m}$ contain all information about the temperature field. For statistically isotropic fluctuations, the expectation of the two-point correlation function in $a_{\ell m}$ drastically simplifies and only depends on ℓ

$$\langle a_{\ell m} a_{\ell' m'}^* \rangle = \delta_{\ell \ell'} \delta_{m m'} C_{\ell}. \quad (4.2)$$

If the fluctuations are statistically isotropic and Gaussian, the angular power spectrum, C_{ℓ} , contains all statistical information about the temperature field.

It is useful to additionally define the real-space angular correlation function for CMB temperature fluctuations as

$$C(\theta) = \langle T(\hat{\mathbf{n}}_1) T(\hat{\mathbf{n}}_2) \rangle \quad (4.3)$$

$$= \frac{1}{4\pi} \sum_{\ell} (2\ell + 1) C_{\ell} P_{\ell}(\cos \theta) \quad (4.4)$$

where $\hat{\mathbf{n}}_1 \cdot \hat{\mathbf{n}}_2 = \cos \theta$ and $P_{\ell}(x)$ is a Legendre polynomial.

4.1.1 Maps and masks

In the course of this analysis we will use several data products from the 2015 data release provided on the Planck Legacy Archive¹. For transparency and reproducibility, when relevant we identify the names of specific files used in footnotes.

Though the primary product of this project will be a study of the covariance between anomalies as measured from simulation ensembles, we also use Planck data to quantify the values of the anomaly statistics described above for the real, observed CMB sky. Our purpose in doing this will be twofold. First, it will allow us to compare our assessment of how anomalous features of the observed CMB sky are against the probabilities reported in the literature. Additionally, employing the same code to measure statistics from the real

¹pla.esac.esa.int

Planck data as from our simulations will allow us to place our real CMB sky for reference in the multi-dimensional feature space examined in Section 4.2 and 4.3.

For map-based statistics we use the SMICA [155] map² from the 2015 Planck data release. Though the Commander map, which separates CMB signal from foreground microwave emission using a pixel-based parametric method (as opposed to SMICA, which is based on a linear combination of Planck’s various frequency channels with multipole-dependent weights), should more properly be used for the analysis of very large scale features, past studies [134, 156, 157] have found that the significance of the various anomalies does not depend strongly on which component separation method is used. Therefore, using the SMICA map should be sufficient for our purposes.

Because we care only about large-angle features, we will work with maps that are at a resolution of $N_{\text{side}} = 64$ which are smoothed with a Gaussian beam of 160 arcmin. We therefore downgrade the resolution of the SMICA CMB temperature map, which is provided at $N_{\text{side}} = 1024$, following the prescription described in [156]. We do this by first extracting its spherical components $a_{\ell m}$ using the HEALPix³ [5] function `map2alm`. Then, again using HEALPix, we get the harmonic space representation of the Gaussian beam b_ℓ and pixel window functions p_ℓ corresponding to the full width half maximum (FWHM) and pixel resolution, respectively, of both the input and output maps. By combining these together, we obtain the downgraded harmonic coefficients,

$$a_{\ell m}^{\text{out}} = \frac{b_\ell^{(\text{out})} p_\ell^{(\text{out})}}{b_\ell^{(\text{in})} p_\ell^{(\text{in})}} a_{\ell m}^{\text{in}}. \quad (4.5)$$

Using the HEALPix function `alm2map` we then convert back to pixel space to obtain the downgraded map. We refer to Table 1 in Ref. [156] for the appropriate beam FWHM values: 160 arcmin for $N_{\text{side}} = 64$, and 10 arcmin for $N_{\text{side}} = 1024$. (For other parts of this study we will also use the conversions: 5 arcmin for $N_{\text{side}} = 2048$ and 640 arcmin for $N_{\text{side}} = 16$.)

When we use a mask, we adopt the UT78 common mask⁴, which is identified in Ref. [134] as the one that should be used for the analysis of Planck temperature maps. UT78 is the union of the masks for Planck’s four methods for separating CMB signal from foreground emission (SMICA, NILC, SEVEM, and Commander). This mask is provided as an $N_{\text{side}} = 2048$ map of zeros and ones, where zeros represent masked pixels and ones signify unmasked pixels. To downgrade the mask to $N_{\text{side}} = 64$, we follow the same procedure described in

²COM_CMB_IQU-smica_1024_R2.02_full.fits

³<http://healpix.sourceforge.net>

⁴COM_Mask_CMB-IQU-common-field-MaskInt.2048_R2.01.fits, field 0.

Eq. (4.5), then threshold the resulting map so that all pixels with a value ≤ 0.9 are marked as masked. This reduces the mask from its original $f_{\text{sky}} = 0.78$ to 0.67. (When we use $N_{\text{side}} = 16$ maps for one of the anomalies studied below, it reduces further to $f_{\text{sky}} = 0.58$.)

4.1.2 Simulated ensembles

Our primary simulation ensemble will be a set of 100,000 noiseless Gaussian CMB temperature maps generated using the `synfast` function in `healpy`. Gaussian temperature map realizations are drawn using the power spectrum for Planck best-fit theory prediction for the power spectrum.⁵ The maps are produced at $N_{\text{side}} = 64$ with FWHM=160 arcmin Gaussian smoothing, and with the `pixwin` argument set to `True`. These settings were chosen to make the simulated maps have properties consistent with the downgraded SMICA temperature maps described above. These straightforward-to-implement simulations, which we refer to as the “synfast simulations,” will allow us to obtain the statistics of fluctuations associated with the CMB signal only. They do not, however, include a number of foreground- and survey-related effects that are present in the Planck SMICA map.

In order to explore whether those effects influence the relationship between anomalies, we repeat our analysis on the publicly available ensemble of Planck Full Focal Plane (FFP) simulations [154] which have been processed using the SMICA component separation pipeline. Specifically, we use the FFP8.1 CMB sky and noise maps, which we add together before downgrading to $N_{\text{side}} = 64$. The FFP simulations include the physical effects of astrophysical foregrounds, gravitational lensing, Doppler modulation, and frequency-dependent Rayleigh scattering effects. They also model the Planck mission’s scanning strategy, detector response, beam shape, and data reduction pipeline. Additionally, a small, frequency-dependent intensity quadrupole has been added to the FFP simulations to account for an uncorrected residual in the data from the dipole-induced Doppler quadrupole identified in Ref. [158]. Note that because we use the FFP8.1 rather than FFP8 simulations, we do not need to rescale the CMB components of the simulations by the factor of 1.0134 that was applied in Ref. [134].

4.1.3 Power spectrum measurements

There are several methods that one can use to measure the angular power spectrum of a map of CMB temperature fluctuations. Using different methods generally will cause variations in

⁵Provided on the Planck Legacy Archive in the file `COM.PowerSpect_CMB-base-plikHM-TT-lowTEB-minimum...-theory_R2.02.txt`

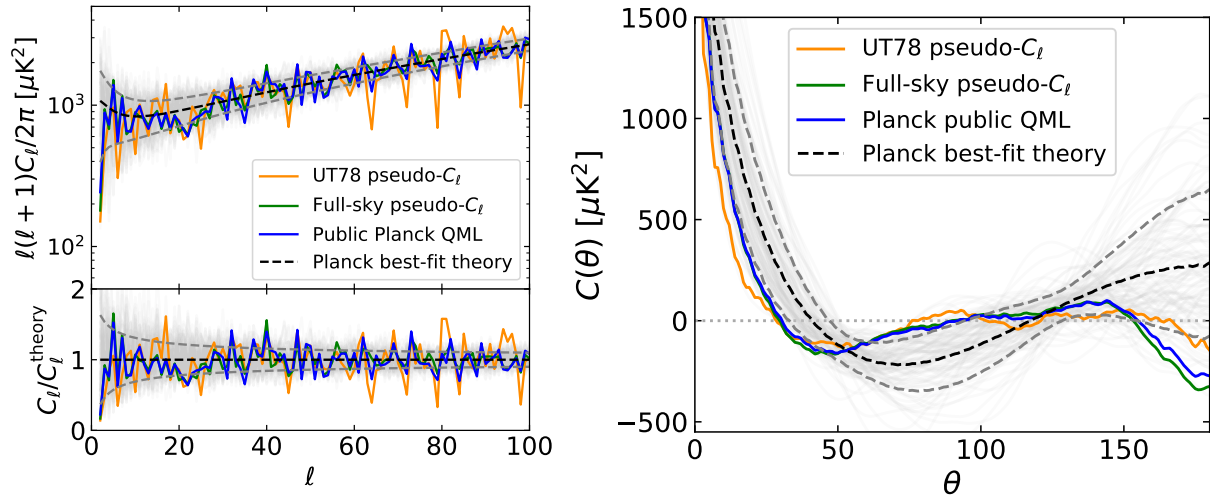


Figure 4.1: Comparison of the different angular power spectrum measurements described in Section 4.1.3 (left) and their corresponding angular correlation functions (right). The collection of grey lines behind them are from the full sky pseudo- C_ℓ measurements of the first 100 synfast simulations. The black dotted line shows theoretical expectation, and the light grey dotted lines show the 68% confidence level cosmic-variance errors.

the estimate for C_ℓ , and thus the choice of how to measure C_ℓ can impact anomaly statistics.

Our analysis mainly relies on pseudo- C_ℓ estimates for the power spectrum. For full-sky measurements these will give unbiased estimates of the true power spectrum simply by averaging the observed spherical harmonics,

$$\hat{C}_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{m=+\ell} |a_{\ell m}|^2. \quad (4.6)$$

If a mask is used to remove contaminated parts of the map, additional care must be taken, as is described in e.g. Ref. [157]. In practice we measure cut-sky pseudo- C_ℓ 's using the `polspice` algorithm⁶ [159], which removes the monopole and dipole of the masked map, measures the angular correlation function $C(\theta)$ of the unmasked part of the sky, and then integrates to obtain C_ℓ ,

$$C_\ell = 2\pi \int_{-1}^1 C(\theta) P_\ell(\cos \theta) d \cos \theta. \quad (4.7)$$

⁶The `polspice` software can be found at <http://www2.iap.fr/users/hivon/software/PolSpice/>. we run it using the settings `subav=YES`, `subdipole=YES`, `apodizesigma=NO`, and `pixelfile=NO`.

In contrast to these cut-sky pseudo- C_ℓ 's, which only estimate the statistical properties of the unmasked parts of the sky, quadratic maximum likelihood (QML) methods can be used to estimate the statistical properties of the entire sky. It will make the most sense for us to study the statistical properties of certain large-angle features in terms of the observed SMICA map's QML power spectrum. For this, we do not implement our own QML power spectrum estimator, but instead use the public QML spectrum provided by the Planck team [160]. The Planck QML power spectrum was obtained using the `BoLPol` software [161] applied to the Commander component-separated map (and mask) at $N_{\text{side}} = 16$ for multipoles $\ell = 2 - 29$, and the `Plik` likelihood applied to measured pseudo- C_ℓ 's for $\ell \geq 30$. The low- ℓ power spectrum estimation uses the Commander mask, which has $f_{\text{sky}} = 0.94$ and therefore leaves available much more of the sky available for analysis than the UT78 common mask. The high- ℓ power spectrum likelihood uses galactic masks, described in Appendix A of Ref. [160], which leave less available sky than those used with component-separated CMB maps.

To summarize, the three C_ℓ measurement strategies we will examine are:

- **Full-sky pseudo- C_ℓ 's:** We computed them using `polspice` based on a map with $N_{\text{side}} = 64$, with the monopole and dipole subtracted.
- **UT78 pseudo- C_ℓ 's:** We computed them using `polspice` with the same settings as the full sky case, using the $N_{\text{side}} = 64$ version of the Planck UT78 common mask.
- **Planck public QML C_ℓ 's:** These estimates are provided for low⁷ and high⁸- ℓ on the Planck Legacy archive.

In Fig. 4.1 we compare the angular power spectra derived from the SMICA CMB temperature map using these three methods, as well as the corresponding angular correlation functions derived using Eq. (4.3). We also show theoretical predictions using the Planck best-fit model, along with the 68% confidence region for cosmic variance. We can see that the full-sky pseudo- C_ℓ 's (green lines) and the Planck public QML C_ℓ 's (blue lines) are similar to one another, while the cut-sky pseudo- C_ℓ 's have a larger variance. The cut-sky pseudo- C_ℓ 's higher variance does not have a significant physical meaning: it is due to the fact that the cut-sky measurements are based on fewer observed modes and so have higher cosmic variance, possibly in combination some impact of masking effects on the translation of real-space measurements to fourier space.

Our $N_{\text{side}} = 64$ resolution implies that we can study multipoles up to $\ell_{\text{max}} = 3N_{\text{side}} - 1 = 191$ (though for practical purposes, pixelization effects become apparent for pseudo- C_ℓ

⁷COM_PowerSpect_CMB-TT-loL-full_R2.02.txt

⁸COM_PowerSpect_CMB-TT-hiL-full_R2.02.txt

measurements at $\ell \sim 150$ for full sky measurements and $\ell \sim 100$ for cut-sky). Since we will be focusing on scales $\ell < 100$, this choice of N_{side} is sufficiently high.

4.2 Features Studied

Here we study eight characteristics of the large-angle CMB temperature maps. These features, which are summarized in Table 4.1, include some of the most prominently discussed CMB anomalies. This set of features is not intended to be comprehensive⁹ Rather, our intention is to focus on a representative sample that will allow us to develop an understanding of the large-angle CMB’s statistical properties in Λ CDM.

Broadly, we classify features based on whether they depend entirely on information in the isotropic two-point statistics, or whether they require map or $a_{\ell m}$ -based information. We adopt this classification to aid in our interpretation of their covariances: because the two-point function anomalies are all functions of the same angular power spectrum, their respective definitions directly imply some *a priori* expectations for their covariances. The same is not necessarily true for the isotropy-breaking anomalies due to the stochastic nature of the $a_{\ell m}$ upon which these statistics are based.

In this section we introduce the features that we study. For each one, we will define the quantity that we use to measure it, briefly introduce relevant findings from previous studies, and discuss how those findings compare to our measurements. We will mainly perform these comparisons against Ref. [134], the Planck 2015 paper on the isotropy and statistics of the CMB, which we will henceforth refer to as Planck XVI (I&S). Unless otherwise noted, their anomaly measurements were done using a QML C_ℓ estimator on UT78 cut-sky maps evaluated at $N_{\text{side}} = 64$. Their measurements of the real sky were done on the SMICA temperature map, and they evaluated statistics based on the FFP8 simulations. Note that because we use a different ensemble of simulations, as well as a different power spectrum measurement technique, we expect our findings for anomaly statistics to be similar to the Planck XVI (I&S) results, but not necessarily to *exactly* match them.

Following Planck XVI (I&S), we will quantify how unusual (or not) the SMICA temperature map appears compared to simulations using p -value, defined to be equal to the fraction of simulations in a given ensemble that return more extreme values than the real sky. As part of each feature description, we will note whether and how measurement choices (be-

⁹For example, we do not include the statistics from Ref. [162] which quantify alignments between different combinations of even and odd parity multipoles. We also do not include the “cold spot” [163] which, being a localized feature at smaller scales, does not naturally belong to the set of large-angle features studied here.

Depends on	Quantity	Description	Multipoles	Section
2pt funcs. only	$S_{1/2}$	Amount of angular power at $\theta > 60^\circ$	2 – 100	4.2.1.1
	C_2	Quadrupole amplitude	2	4.2.1.2
	C_3	Octopole amplitude	3	4.2.1.3
	σ_{16}^2	Variance of δT at $N_{\text{side}} = 16$	2 – 47	4.2.1.4
	R_{27}	Ratio of power between even and odd ℓ	2 – 27	4.2.1.5
	$C(\pi)$	Angular correlation at $\theta = 180^\circ$	2 – 191	4.2.1.6
Phases of $a_{\ell m}$	S_{QO}	Quadrupole-octopole alignment	2, 3	4.2.2.1
	A_{LV}	Hemispherical power asymmetry	2 – 191	4.2.2.2

Table 4.1: Summary of large angle CMB features studied in this work.

tween Planck QML vs. pseudo- C_ℓ , cut-sky vs. full-sky) affect those statistics, and will take care to identify which of those choices are used in our anomaly covariance studies presented in the next Section (4.3). These single-feature results are summarized in Fig. 4.3, which is described in more detail in Section 4.2.3.

4.2.1 Features depending on two-point functions only

We first study the six features that are fundamentally a function of the angular clustering power.

4.2.1.1 $S_{1/2}$: Large-angle power

First, we measure power in large angular scales of the temperature map using the $S_{1/2}$ statistic, defined as the integral of the square of the angular correlation function $C(\theta)$ over angles between 60° and 180° [164]

$$S_{1/2} = \int_{-1}^{1/2} [C(\theta)]^2 d(\cos \theta). \quad (4.8)$$

It measures the deviation of $C(\theta)$ from zero at angles greater than 60° . The inclusion of this statistic is motivated by the lack of power at large angular scales $\theta \gtrsim 60^\circ$ first observed by COBE [165], and later confirmed by WMAP [164] and Planck [134, 157]. Though there are several ways of quantifying this lack of large-angle correlation, we adopt $S_{1/2}$ because it is the most commonly used.

In practice, to measure $S_{1/2}$ for a temperature map, we first measure the angular power

spectrum and then calculate it in harmonic space via [152]

$$S_{1/2} = \frac{1}{(4\pi)^2} \sum_{\ell, \ell'} (2\ell + 1)(2\ell' + 1) C_\ell I_{\ell, \ell'} \left(\frac{1}{2}\right) C_{\ell'}. \quad (4.9)$$

Here the matrix $I_{\ell, \ell'}$ is defined as

$$I_{\ell, \ell'}(x) = \int_{-1}^x P_\ell(x') P_{\ell'}(x') dx', \quad (4.10)$$

but is in practice computed using the recursion relation in Appendix A of Ref. [152]. We sum over values $\ell = 2 - 100$.

When analyzing the SMICA map at $N_{\text{side}} = 64$ with the UT78 mask Planck XVI (I&S) reports a low value for $S_{1/2}$ with a lower tail probability¹⁰ of $p = 0.4\%$. That is to say, they find that only 0.4% of simulations have a lower value of $S_{1/2}$ than the SMICA map. Our cut-sky $S_{1/2}$ measurements give similar probabilities: $p = 0.7\%$ for the fiducial synfast simulations and $p = 0.5\%$ for the FFP simulations.¹¹

The lower-tail probability of the observed sky's $S_{1/2}$ value depends dramatically on the method used to measure the angular power spectrum, increasing to 8% for full sky pseudo- C_ℓ 's and to 6% for Planck public QML C_ℓ 's (which, recall, effectively reconstruct the full-sky anisotropy field). This is consistent with results from previous studies [152, 157, 166] which have shown that the relatively small amount of (non-zero) correlations on the full sky are dominated by contributions from pixels close to the galactic mask.

4.2.1.2 C_2 : Quadrupole amplitude

We additionally study C_2 , the quadrupole of temperature fluctuations, which was first found to be low in COBE [167] data, and later in WMAP [127, 168] and Planck [160]. Analyses have shown that the lowness of the quadrupole is not particularly significant [169–171], so its value or lower-tail probabilities are not generally reported explicitly in the literature. Given this, we do not directly compare our measurement of C_2 to previous results, but do include it as one of our statistics, in order to study its covariance with the low angular power at large angles and other features.

¹⁰Value from Table 13 of Ref. [134]

¹¹These p -values are weakly sensitive to whether the C_ℓ 's are corrected for resolution according to Eq. (4.11): with that correction, the p -values for $S_{1/2}$ go down to 0.6% for the synfast simulations and to 0.4% for the FFP simulations. We opt not to make that correction when computing $S_{1/2}$ and $C(\pi)$ because doing so introduces significant noise contributions at high multipoles and makes the sums involved overly sensitive to our choice of ℓ_{max} .

Our one-dimensional study of C_2 's statistics reflect the findings in the literature. Our fiducial choice for the quadrupole is to adopt the Planck QML C_2 to represent the observed value, while for the simulations we calculate the full-sky C_ℓ 's and adopt the corresponding quadrupole. In order to make the simulation measurements more directly comparable to the QML power spectrum, we apply a correction for the $N_{\text{side}} = 64$ maps' beam and pixel window functions via

$$C_\ell = (b_\ell^{(64)} p_\ell^{(64)})^{-2} C_\ell^{\text{polspice}}. \quad (4.11)$$

Here, $b_\ell^{(64)}$ and $p_\ell^{(64)}$ are the harmonic components of the beam and pixel window functions for the $N_{\text{side}} = 64$ input map.

We find that C_2 has a lower-tail probability of 5% using the synfast simulations, and 6% compared to the FFP simulations. Pseudo- C_ℓ measurements of the SMICA map give slightly lower probabilities, with $p \sim 2\%$ for either full- or cut-sky measurements.

4.2.1.3 C_3 : Octopole amplitude

Though the value of the observed CMB temperature C_3 amplitude is not anomalous (e.g. [171]) — to the contrary, we find the Planck QML C_3 to be close to the median the values from our simulated ensembles — we also include it in our study. We do so because its behavior in relation to other features has the potential to be interesting. For example, Ref. [153] points out that contributions from the quadrupole and octopole seem to be canceling the power from the rest of the sky, and that a measure of large-angle power becomes less anomalous when their contributions to the correlation function are removed. Additionally, Ref. [152] finds that relationship between several of the lowest multipoles, certainly more than the just the quadrupole, is responsible for the low observed $S_{1/2}$. Given this, we include C_3 in our analysis because the relationship between C_3 , C_2 and $S_{1/2}$ may reveal some interesting structure.

We perform our fiducial measurement of C_3 in the same way as for C_2 : we use the Planck QML C_ℓ 's for the observed temperature map, and the beam-and-pixelization-corrected (according to Eq. (4.11)) full-sky C_ℓ measurements from simulations. Compared to the synfast and FFP simulations, the p -values for both the QML and full-sky SMICA measurements are 47 – 49%, while the cut-sky octopole is lower, with $p \sim 15\%$.

4.2.1.4 σ_{16}^2 : Variance at $N_{\text{side}} = 16$

We study another indicator of large-angle power via σ_{16}^2 , the variance of unmasked pixels of a low resolution, $N_{\text{side}} = 16$ temperature map. The variance of CMB temperature maps,

especially at low spatial resolution, has been observed to be anomalously low in analyses of both WMAP [172, 173] and Planck [134] data. Planck measured the variance of unmasked pixel values with various N_{side} , finding the lowest investigated value, $N_{\text{side}} = 16$, with a p -value¹² of 0.5%, to give the most anomalously low variance. They found that the variance tends to become lower as the mask is extended to cover more of the sky, and that the statistical significance of its lowness persists when different foreground subtraction methods are applied.

To measure σ_{16}^2 for a given CMB temperature map, we first downgrade the map from $N_{\text{side}} = 64$ to $N_{\text{side}} = 16$. We also downgrade the UT78 mask, but go directly from the original $N_{\text{side}} = 2048$ resolution to $N_{\text{side}} = 16$ in order to make the resulting sky fraction consistent with that used in the Planck study. We then simply compute the variance of all unmasked pixels.

Though we measure σ_{16}^2 through a pixel-based method, given an angular power spectrum C_ℓ we can predict its expectation value for full-sky measurements via

$$\langle \sigma_{16}^2 \rangle(C_\ell) = \frac{1}{4\pi} \sum_{\ell} (2\ell + 1) C_\ell (b_\ell^{(16)} p_\ell^{(16)})^2 \quad (4.12)$$

where $b_\ell^{(16)}$ and $p_\ell^{(16)}$ are the beam and pixel window functions corresponding to $N_{\text{side}} = 16$. This expression will allow us to compare our map-based measurements of σ_{16}^2 to the predictions from the Planck best fit theory C_ℓ 's as well as the Planck public QML C_ℓ 's.

It is worth noting that our method of measuring σ_{16}^2 is different from that used to quantify map variance in the WMAP and Planck analyses. Those analyses use an estimator [172] to isolate the cosmological contribution to the variance of a normalized version of the temperature map, in which each pixel value has been divided by its expected dispersion from both cosmological temperature fluctuations and noise. Because of this, our reported numbers for σ_{16}^2 will be much larger than the normalized variances reported in Planck XVI (I&S), but the statistical distribution of variances should be similar, to the extent that noise contributions to variance can be approximated as direction-independent.¹³

For measurements of σ_{16}^2 we would like to exclude pixels that may contain residual foregrounds; so we focus on its cut-sky value for both from the SMICA map and simulations. We find the SMICA σ_{16}^2 to be low compared to simulations, with single-tail probability of $p = 0.8\%$ and 0.5% for the **synfast** and FFP simulations, respectively. Thus, our cut-sky

¹²This value is taken from Table 12 of Ref. [134].

¹³Though in principle the noise dispersion can vary with position on the sky due to beam effects and weights used to construct component separated maps, those effects are expected to be small [173].

FFP p -value exactly matches that in Planck XVI (I&S). The σ_{16}^2 expectation value from the Planck public QML C_ℓ 's and our our full-sky pseudo- C_ℓ measurements are very similar, with $p = 20\%$ when compared to either simulation ensemble.

4.2.1.5 R_{27} : Parity asymmetry at low ℓ

We use the statistic R_{27} to quantify large-angle parity asymmetry of the CMB temperature map. It has been noted that, at low ℓ , the CMB maps have more power in odd multipoles than even. This was observed in the WMAP 3, 5, and 7 year data [174–176] as well as in Planck [134]. We quantify this asymmetry using the same estimator as Planck XVI (I&S),

$$R_{\ell_{\max}} = \frac{D_+(\ell_{\max})}{D_-(\ell_{\max})} \quad (4.13)$$

where

$$D_{+,-} = \frac{1}{\ell_{\text{tot}}^{+,-}} \sum_{\ell=2, \ell_{\max}}^{+,-} \frac{\ell(\ell+1)}{2\pi} C_\ell, \quad (4.14)$$

where the plus and minus indicate sums over even (parity-symmetric) and odd (parity-antisymmetric) multipoles, respectively. The $R_{\ell_{\max}}$ statistic is therefore a ratio of the parity-even over parity-odd multipole band-powers. The factor of $\ell(\ell+1)/(2\pi)$ is used because the theoretical prediction for $\ell(\ell+1)(2\pi)^{-1}C_\ell$ is approximately scale-independent out to multipoles of $\ell \lesssim 50$, and thus predicts $R_{\ell_{\max}} \sim 1$ over that range.

Because $R_{\ell_{\max}}$ is directly based on the power spectrum, we will focus on its measurements from the Planck QML power spectrum, and then compare them to full-sky pseudo- C_ℓ measurements in simulations. As in the case of C_2 and C_3 , we correct for the impact of the simulations' resolution on the power spectrum using Eq. (4.11).

For our covariance studies, we will focus on the behavior for $\ell_{\max} = 27$, as that multipole range gives the most anomalously low value of $R_{\ell_{\max}}$ in the Planck XVI (I&S) analysis¹⁴, with a single-tail probability $p = 0.2\%$ for the SMICA map. We find the SMICA map's R_{27} to be notably less anomalous: measurements of the Planck QML power spectrum give single tail probabilities of $p = 3\%$ and 2% when compared to the `synfast` and FFP simulations, respectively.

Given this p -value discrepancy, we investigated how $R_{\ell_{\max}}$ depends on ℓ_{\max} and power spectrum measurement technique. Results of this investigation, shown in Fig. 4.2, reveal

¹⁴Value taken from text associated with Fig. 20 of Ref [134]. Though that text actually reports $\ell = 28$ to give the lowest $R_{\ell_{\max}}$ p -value, this is due to a typographical error, and we confirmed with that section's author that the minimum p -value is actually at $\ell = 27$. [177]

that the significance of parity asymmetry heavily depends on the choice of mask and power spectrum measurement method. The differences between our results and those of Planck XVI (I&S) can therefore be explained by the fact that their $R_{\ell_{\max}}$ measurements are based on QML C_ℓ measurements of the UT78 cut-sky map, which has a smaller sky fraction than the maps used to produce the Planck public QML power spectrum [177].¹⁵ We obtain a p -value closest to that reported for R_{27} in Planck XVI (I&S) using full-sky pseudo- C_ℓ measurement of the SMICA map compared to FFP simulations (0.6%). However, in order to be consistent with our treatment of the other purely power-spectrum based features, we will use the public Planck QML C_ℓ 's compared to full-sky pseudo- C_ℓ simulation measurements as our fiducial choices for measuring R_{27} .

4.2.1.6 $C(\pi)$: Two-point correlation at $\theta = 180^\circ$

We next consider the angular correlation function of CMB temperature evaluated at 180° , which we will refer to as $C(\pi)$. We include it in the hope that it will help clarify the relationship between other features. Our motivation comes from the fact that $C(\theta)$, which is otherwise fairly flat at large angles, drops to negative values at $\theta \simeq \pi$. This dip has been observed in both WMAP and Planck data, and can be seen in the colored curves on the right-hand side of Fig. 4.1. By its definition we expect the value $C(\pi)$ to be related to the missing large-angle correlations statistic $S_{1/2}$, as well as to the R_{27} measurement of parity asymmetry. This can be seen by comparing the definition for $R_{\ell_{\max}}$ in Eqs. 4.13 to

$$C(\pi) = \sum_{\ell=2}^{\infty} (-1)^\ell \frac{2\ell+1}{4\pi} C_\ell. \quad (4.15)$$

We will of course investigate these correlations quantitatively further below.

We measure $C(\pi)$ by using cut-sky pseudo- C_ℓ measurements of the SMICA map and simulations to obtain measured power spectra (as we do for $S_{1/2}$), transforming to real space angular correlations via Eq. (4.15). We compute the sum over multipoles $\ell = 2 - 100$. For the cut-sky measurements, the SMICA map's $C(\pi)$ is lower than for the majority of simulations, but is not particularly anomalous, with $p \sim 11\%$ for both the **synfast** and FFP simulation ensembles. The p -value goes down to 5–6% if we instead compare SMICA map measurements using the public Planck QML C_ℓ 's to the same set of simulation measurements, and 3–4% for full-sky pseudo- C_ℓ 's. As this measure is not commonly studied in the literature,

¹⁵We confirmed that we were able to replicate the Planck results when using their UT78 QML C_ℓ values, which are an intermediate data product of their analysis, but restrict this study to only publicly available data.




Figure 4.2: Plot of the lower-tail probability p for the parity asymmetry as a function of the largest multipole ℓ_{\max} considered, for various SMICA power spectrum measurement and simulation ensemble combinations. Solid lines show the probabilities for the SMICA map assessed relative to the synfast simulations, and dashed lines show them relative to the FFP simulations. For the QML and full-sky pseudo- C_ℓ SMICA measurements, the simulations are measured using full-sky pseudo- C_ℓ 's. For the cut-sky pseudo- C_ℓ SMICA measurements, cut-sky pseudo- C_ℓ simulation measurements are used. The vertical line denotes the ℓ_{\max} value at which where Planck XVI (I&S) found the most anomalous parity statistic; see text for details.

we do not compare this to any reported values.

4.2.2 Features depending on $a_{\ell m}$ phases

We also consider two features that depend on the phases of CMB temperature harmonic coefficients $a_{\ell m}$ — i.e. which cannot be measured solely from the two-point correlations the maps. These features are associated with reported anomalies that could indicate a possible departure from the assumption of statistical isotropy of CMB temperature fluctuations: the quadrupole-octopole alignment and hemispherical power asymmetry.

4.2.2.1 S_{QO} : Quadrupole-octopole alignment

The CMB temperature’s quadrupole and octopole were first observed to be planar and aligned in Ref. [127]. There are a number of possible ways to denote the directionality, and thus alignment, of multipoles. To do so, we will follow the approach presented in Ref. [178] and make use of Maxwell multipole vectors. Multipole vectors are a representation of a function on a sphere; while they are at some level equivalent to spherical harmonics in that role, their relationship to the $a_{\ell m}$ is highly nonlinear in a way that makes them are particularly well suited to studying the directionality of patterns on the sky.

For each multipole ℓ , there are ℓ corresponding multipole vectors $\mathbf{v}^{(\ell,i)}$, where $i \in \{1, 2, \dots, \ell\}$. Roughly speaking, the more the power of temperature fluctuations associated with a given multipole is concentrated in a plane, the more its associated multipole vectors will be confined to a plane, and the more the oriented-area vectors defined by their cross products,

$$\mathbf{w}^{(\ell,i,j)} \equiv \pm(\mathbf{v}^{(\ell,i)} \times \mathbf{v}^{(\ell,j)}), \quad (4.16)$$

will line up in a direction normal to that plane. Moreover, planarity (as opposed to simply orientation along a direction) of the temperature multipole will cause the multipole vectors $\mathbf{v}^{(\ell,i)}$ and $\mathbf{v}^{(\ell,j)}$ to be at large angles relative to each other, enhancing the magnitude of $\mathbf{w}^{(\ell,i,j)}$. Thus, we can use the extent to which the object oriented vectors for two multipoles point in similar directions to measure how much the power from those ℓ -modes are aligned.

The statistic S_{QO} takes advantage of this property to quantify the quadrupole-octopole alignment. It is the normalized sum of the dot products of the quadrupole oriented-area vector $\mathbf{w}^{(2,1,2)}$ with the three octopole oriented-area vectors $\mathbf{w}^{(3,i,j)}$: [128, 179]

$$S_{\text{QO}} = \frac{1}{3} \sum_{\{i,j\}} |\mathbf{w}^{(2,1,2)} \cdot \mathbf{w}^{(3,i,j)}| \quad (4.17)$$

where $\{i, j\}$ can be $\{1, 2\}$, $\{2, 3\}$, or $\{3, 1\}$. Given this, larger values of S_{QO} indicate more alignment and planarity in the $\ell = 2$ and $\ell = 3$ modes of the temperature maps.

Because multipole vectors are defined in terms of $a_{\ell m}$, this measurement can only be done on full-sky maps. To measure S_{QO} for a temperature map, we first use the HEALPix function `map2alm` to measure its $a_{\ell m}$, and then use the procedure¹⁶ described in Appendix A of Ref. [178] to extract the multipole vectors for $\ell = 2 - 3$. We then combine them via Eqs. (4.16) and (4.17) to get S_{QO} .

We find that the S_{QO} value measured from the SMICA map is larger than that from most simulations, with a p -value (here, upper-tail probability) of 0.4% when compared with either our synfast or FFP simulation ensembles. This is consistent with the results in Ref. [126]. They found the SMICA map from the Planck 2013 data release [180] had a larger S_{QO} , with a p -value¹⁷ of 0.54% compared to an ensemble of 10^6 simulations analogous to our synfast simulations, which use constrained realizations to in-paint masked regions. The fact that our p -value is so similar to theirs indicates that simply measuring S_{QO} on full-sky maps (as we do) rather than doing in-painting does not significantly affect the large-scale alignment behavior.

4.2.2.2 A_{LV} : Hemispherical power asymmetry via local-variance dipole

Finally, we include a measure of the level of asymmetry in temperature power between two hemispheres of the sky. This is studied because one hemisphere of the observed CMB sky has been noted to have more power than the other [131, 181], which can be modeled by a dipole modulation of temperature fluctuations at large angular scales [182, 183]. Following Refs. [184, 185], we quantify hemispherical power asymmetry using a local variance map, which measures the size of temperature fluctuations within disks of radius θ centered on each of its pixels. By measuring the dipole of a local-variance map, we can quantify the direction and magnitude of any hemispherical power asymmetry in a computationally inexpensive way. Additionally, we can probe the scale dependence of the effect by varying the angular size θ of the disks used to create the local variance map.

More formally, if $\hat{\mathbf{n}}_i$ is the location of the i th pixel of the input temperature map $T(\hat{\mathbf{n}})$ from which the monopole and dipole of unmasked pixels have been removed, we can write

¹⁶Calculations were performed using code `mpd_decomp.py` provided at <http://www.phys.cwru.edu/projects/mpvectors/>.

¹⁷Value from Table 7 of Ref. [126]

the local-variance map $\sigma_\theta^2(\hat{\mathbf{n}})$ as

$$\sigma_\theta^2(\hat{\mathbf{n}}) = \frac{1}{N[\mathcal{D}_\theta(\hat{\mathbf{n}})]} \sum_{i \in \mathcal{D}_\theta(\hat{\mathbf{n}})} [T(\hat{\mathbf{n}}_i) - \bar{T}_\theta(\hat{\mathbf{n}})]^2, \quad (4.18)$$

where $\mathcal{D}_\theta(\hat{\mathbf{n}})$ is the set of unmasked pixels within angle θ of direction $\hat{\mathbf{n}}$, $N[\mathcal{D}_\theta(\hat{\mathbf{n}})]$ is the number of pixels in that set, and $\bar{T}_\theta(\hat{\mathbf{n}})$ is their average temperature.

In practice, we measure the dipole of a dimensionless, weighted version of the local-variance map,

$$\tilde{\sigma}_\theta^2(\hat{\mathbf{n}}) = \frac{w(\hat{\mathbf{n}})}{\bar{w}} \times \frac{\sigma_\theta^2(\hat{\mathbf{n}}) - \mu_\theta(\hat{\mathbf{n}})}{\mu_\theta(\hat{\mathbf{n}})}. \quad (4.19)$$

In this expression, $w(\hat{\mathbf{n}})$ is a dimensionless weight map (defined below), \bar{w} is its average over unmasked pixels, and $\mu_\theta(\hat{\mathbf{n}})$ is the mean computed by averaging the local variance maps of an ensemble of simulated CMB temperature maps.

We choose the weight function $w(\hat{\mathbf{n}})$ to be

$$w(\hat{\mathbf{n}}) = \frac{1}{\text{var}[\sigma_\theta^2(\hat{\mathbf{n}})]} \times \left\{ \frac{1}{N_{\text{pix}}} \sum_{i=1}^{N_{\text{pix}}} \text{var}[\sigma_\theta^2(\hat{\mathbf{n}}_i)] \right\} \quad (4.20)$$

where we define, if α labels the simulation realization and N_{sim} is the number of simulations,

$$\text{var}[\sigma_\theta^2(\hat{\mathbf{n}})] = \frac{1}{N_{\text{sim}}} \sum_{\alpha=1}^{N_{\text{sim}}} \left[\frac{\sigma_\theta^{2(\alpha)}(\hat{\mathbf{n}}) - \mu_\theta(\hat{\mathbf{n}})}{\mu_\theta(\hat{\mathbf{n}})} \right]^2. \quad (4.21)$$

This inverse variance weighting suppresses the impact of noisy regions of the input temperature map, as long as those noise contributions are modeled in the simulations. In the limit that noise properties are direction-independent, the weight factor $w(\hat{\mathbf{n}})/\bar{w}$ will approach 1.

In our work, we measure A_{LV} from the UT78 cut-sky for both the SMICA and simulation maps. We fix the disk radius to be $\theta = 8^\circ$, which is the scale previously found to produce the most anomalous local variance dipole, and compute the local variance maps at a resolution of $N_{\text{side}} = 16$. The amplitude of the dipole of a normalized local-variance map, A_{LV} , is then obtained by using the Healpix function `remove_dipole`. Following Planck XVI (I&S), we include only disks for which at least 90% of the input pixels are unmasked.

Our measurements return a local-variance dipole amplitude with $A_{\text{LV}} = 0.22$ when $\sigma_\theta^2(\hat{\mathbf{n}})$ is normalized using the synfast simulations, and $A_{\text{LV}} = 0.21$ using the FFP simulations. These values give an upper tail probability of $p = 1\%$ when compared to either set of simulations. Both of these values are notably larger than the Planck XVI (I&S) findings of

$A_{LV} \sim 0.044$ for the SMICA map, with a p -value of 0.1%.¹⁸ In investigating this discrepancy, we found that the value of A_{LV} is highly sensitive to the resolution of the input temperature maps, with lower resolution input maps tending to give larger dipole amplitudes. Because we compute local variances using our fiducial set of $N_{\text{side}} = 64$ maps, while Planck XVI (I&S) uses $N_{\text{side}} = 2048$ input maps, we believe this resolution dependence explains the difference.

We also note that Doppler dipole modulation included in the FFP simulations [186] (but not the synfast simulations) will generate a small power asymmetry which contributes to the local-variance power asymmetry. However, that contribution is expected to be negligible for the scales $\ell \lesssim 191$ that we are investigating [187]. The fact that the p -values from comparisons to the synfast and FFP simulations are nearly identical is in line with that expectation.

4.2.3 Summary: Individual anomaly measurements

To summarize, we have defined quantities associated with eight properties of the CMB temperature map which are either found to be statistically unlikely in the observed sky or which are expected shed shed light on the relationship between statistically unlikely features. For each feature, we have described our technique for measuring it from real and simulated CMB maps. By analyzing the resulting data associated with each feature individually, we have verified (where applicable) that our measurement of the SMICA map’s p -values (single-tail probabilities) relative to the simulation ensembles are consistent with previous findings.

Fig. 4.3 shows a summary plot of these anomaly statistics measured from our ensembles of 100,000 synfast (left column), and 1,000 FFP (right column) simulations. In each panel, the grey histogram shows the distribution of simulation measurements, which are either made based on full-sky maps, or the cut UT78 sky, as indicated by the gray text in the lower right corner. These data are what will be used in subsequent sections to study the relationship between features in the context of isotropic Λ CDM. The vertical lines show feature measurements done using the Planck public QML C_ℓ ’s (blue), pseudo- C_ℓ ’s extracted from the full-sky SMICA map (green), UT78 cut-sky pseudo- C_ℓ ’s for the SMICA map (orange), and the theoretical expectation based on Planck’s best-fit parameter values (dashed black). Note that the two statistics that depend on the phases (S_{QO} and A_{LV}) do not have a corresponding measurement from either the published QML or the best-fit theory C_ℓ , as they can not be related to the angular power spectrum. The QML and theory values for σ_{16}^2 are computed using Eq. (4.12). The p -values for these measurements are shown in the same

¹⁸Values taken from Fig. 27 (A_{LV} value) and Table 20 (p -value) of Ref. [134].

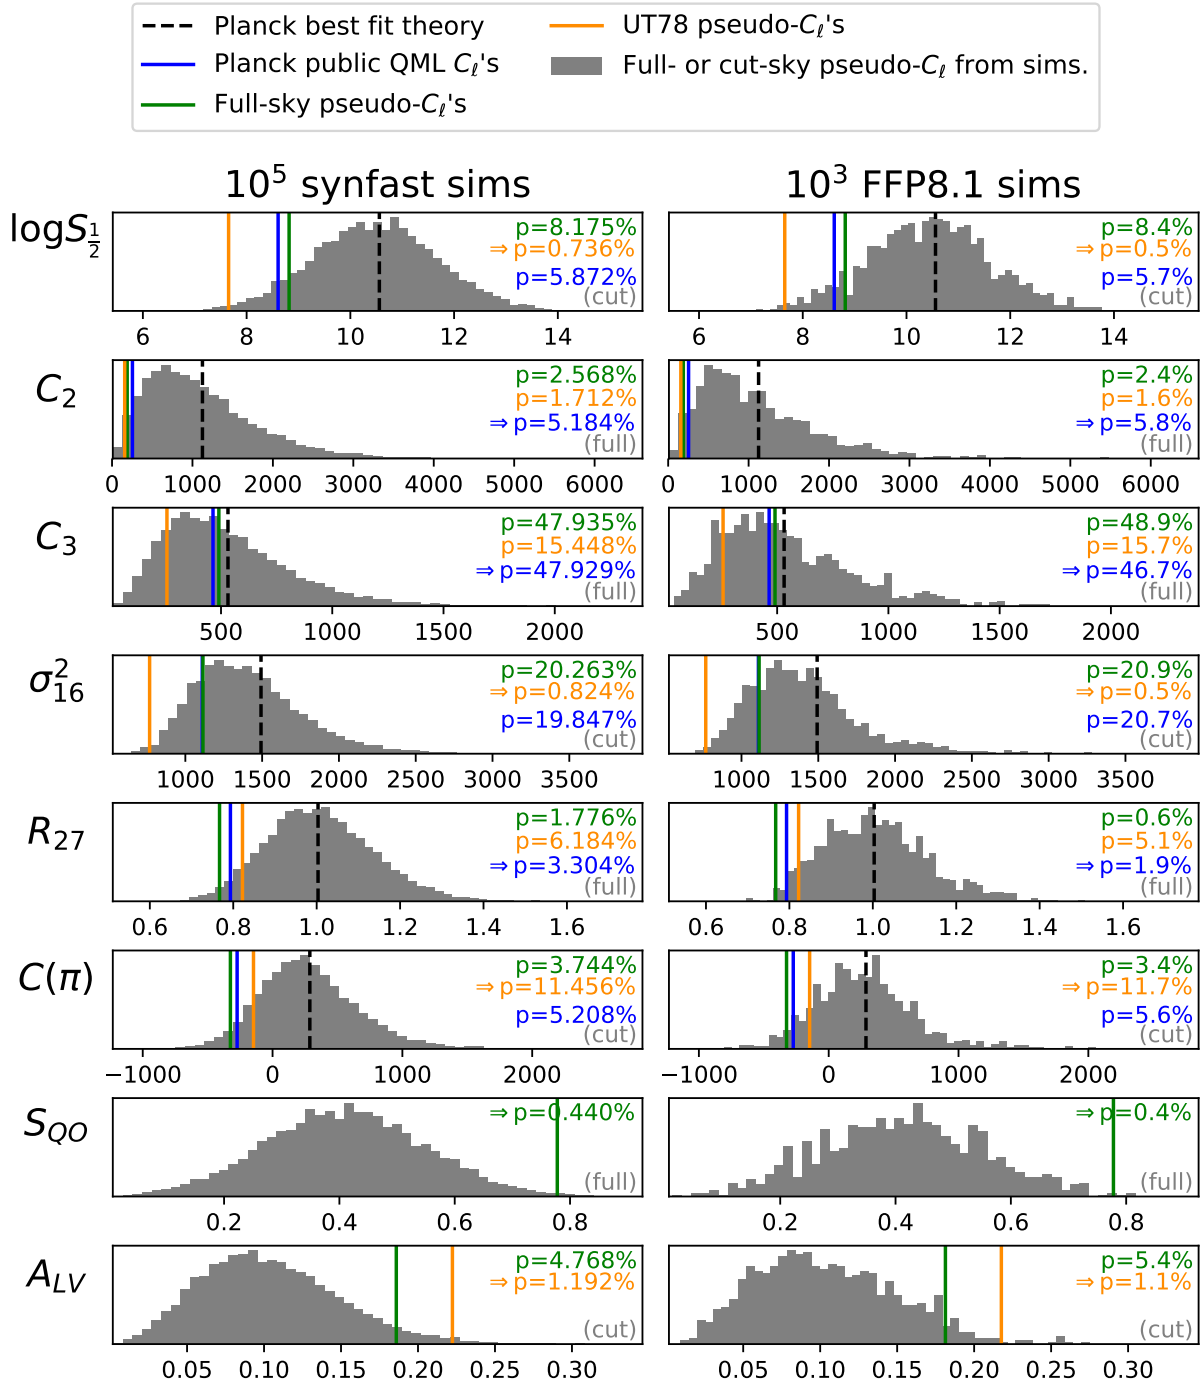


Figure 4.3: Summary plot of one-dimensional probability distributions of features associated with large-angle CMB anomalies. Grey histograms show results for simulations, measured on either cut-sky (using the UT78 mask) or full-sky measurements as indicated by grey text in each panel. Vertical lines are measurements of the real sky, with their single-tail probabilities in the corresponding color. Arrows by the p -values indicate our fiducial measurement choices.

color on the right-hand side of each panel, with an arrow indicating which measurement we think is most relevant for the feature in question. These choices and the findings for each quantity are discussed in detail above.

To put the p -values in context, the 1, 2, and 3σ error bars for a normal distribution correspond to single tail probabilities of $p = 16\%$, 2.2% , and 0.014% , respectively. Thus, of the features we study, the SMICA values for $S_{1/2}$, C_2 , σ_{16}^2 , S_{QO} , and A_{LV} are $2 - 3\sigma$ unlikely compared to our simulations, while R_{27} and $C(\pi)$ are between 1 and 2σ , and C_3 is not in tension. For the most part our measured p -values are consistent with previous findings. Where they are not, we can explain the discrepancies in terms of the power spectrum measurement technique (for R_{27}) or the initial map resolution (for A_{LV}).

We can also use the results in Fig. 4.3 to make some general observations about the impact of different measurement choices. Outside of small differences which are within the reasonable range of sampling error, the one-dimensional p -values comparing measurements of the real sky measurements relative to synfast simulations are in good agreement with those comparing the real sky to FFP simulations. We additionally note that in general the Planck public QML C_ℓ 's give results that are very similar to full-sky pseudo- C_ℓ measurements. The fact that the cut-sky pseudo- C_ℓ SMICA measurements have comparably lower $S_{1/2}$, C_2 , C_3 , and σ_{16}^2 are consistent with previous studies which have found that the observed lack of power in the CMB sky is more severe in regions further from the galactic mask.

4.3 Results: Anomaly covariances

We are now ready to tackle the main goal of this chapter and study the relationships between the large-angle CMB temperature map properties which have been examined individually above. The features are: the integrated power of temperature fluctuations at angles $\theta > 60^\circ$ ($S_{1/2}$), the quadrupole amplitude (C_2), the octopole amplitude (C_3), the variance of the temperature map evaluated at resolution $N_{\text{side}} = 16$ (σ_{16}^2), the parity statistic R with maximum multipole of $\ell = 27$ (R_{27}), the angular power spectrum at 180° ($C(\pi)$), the quadrupole-octopole alignment (S_{QO}), and the amplitude of the hemispherical power asymmetry (A_{LV}). Using our measurements of these quantities for our synfast and FFP simulations, we will determine their covariances in order to build an understanding of how they are related under the assumption of isotropic Λ CDM. We will do so in three stages, first describing the relationship between pairs of features measured from the synfast simulations, then comparing the covariance matrices for the synfast and FFP simulations, and finally, further exploring the covariance structure by performing a principal component analysis.

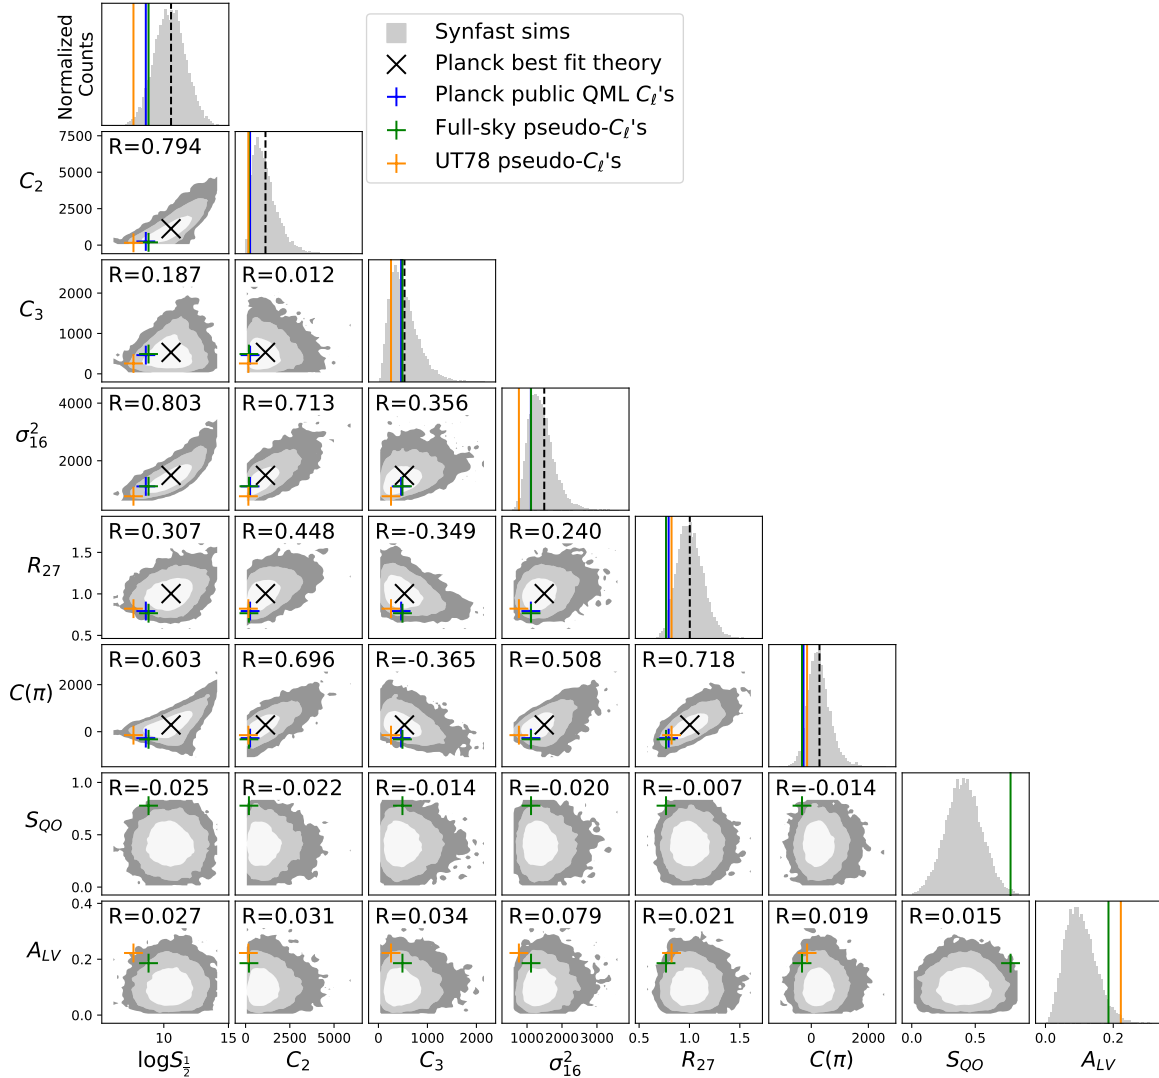


Figure 4.4: Relationships between large-angle CMB features (see Table 4.1 descriptions of these quantities). Gray contours show the 1, 2, and 3 σ confidence regions based on measurements from our ensemble of 100,000 synfast simulations, using the same measurement choices that produced the gray histograms in Fig. 4.3. The 1D histograms on the diagonal are the same as in the left column of Fig. 4.3. The marked data points for SMICA measurements and theoretical expectations are equivalent to the vertical lines in Fig. 4.3. Note that the statistics based on the phases of the $a_{\ell m}$, S_{Q0} and A_{LV} , do not have the corresponding theoretical expectations because they cannot be computed analytically from C_ℓ .

We begin by inspecting how our fiducial set of 100,000 synfast simulations are distributed in the eight-dimensional space defined by the parameters $S_{1/2}$, C_2 , C_3 , σ_{16}^2 , R_{27} , $C(\pi)$, S_{QO} and A_{LV} . Fig. 4.4 shows the relationships between pairs of those quantities; the methods we use to measure them are described in Section 4.2. The diagonal panels display the same one-dimensional statistics information as the left column of Fig. 4.3, with the gray histograms showing the distribution from simulations and the vertical lines showing the measurements of the SMICA map and theoretical expectations based on the Planck’s best-fit cosmological parameters.

In the off-diagonal panels, the grey contours indicate the 1, 2, and 3σ confidence regions based on simulation data. The contour locations were determined for each panel as follows. First, simulation data was used to make a two-dimensional histogram with 50 bins along each axis. Next, the histogram was smoothed using a Gaussian filter with a width corresponding to one bin. The smoothed histogram was then thresholded at constant-count (constant-probability) surfaces so that 68% of the input realizations fall inside the 1σ contours, etc. The number at the top displays the correlation coefficient R of the two quantities shown in the panel, computed based on the simulation samples. Measurements of the observed SMICA map and theory predictions are shown using colored crosses.

Examining Fig. 4.4 we can make a few general observations. First, there are notable covariances and structure to the relationships between most of the features that depend only on two-point functions, but not between C_2 and C_3 , nor between the $pa_{\ell m}$ -phase-dependent quantities (S_{QO} and A_{LV}) and any of the other quantities. This is expected given the isotropic Λ CDM model used to generate the simulations. We additionally note that the distribution of simulation points in this eight-dimensional space is decidedly non-Gaussian; this is due to the asymmetric limits on the quantities measured, as well as the (in some cases) non-linear dependence of quantities on C_ℓ .

More specifically, the covariances of the two-point-function-based quantities can be understood by how they depend on the power spectrum components C_ℓ . In isotropic Λ CDM, we expect the power at different multipoles ℓ to be independent, and correspondingly we see little covariance between C_2 and C_3 . The positive correlation between σ_{16}^2 and either of these amplitudes is straightforward, given Eq. (4.12): all else being equal, adding power to low ℓ is expected to increase the variance at large scales. Similarly, increasing C_2 adds to even- ℓ power and increasing C_3 adds to odd- ℓ power, so we expect and see that the correlation of the parity measure R_{27} to be positively and negatively correlated with C_2 and C_3 , respectively.

Looking at $C(\pi)$ allows us to clarify the relationship between the parity properties and $S_{1/2}$. We note that, given the parity properties of spherical harmonics and referencing

Eq. (4.15), the contributions from even- ℓ modes to $C(\pi)$ will be $\propto (2\ell + 1)C_\ell$, while odd- ℓ contributions are $\propto -(2\ell + 1)C_\ell$. Thus $C(\pi)$ is effectively another way of characterizing the parity properties of the large-angle CMB. Accordingly we see that $C(\pi)$ has a strong positive correlation with R_{27} and with C_2 and a somewhat weaker negative correlation with C_3 . Since $C(\pi)$ is the measurement of the angular correlation at $\theta = 180^\circ$, small $S_{1/2}$ values require that $|C(\pi)|$ be close to zero. The sharp triangular structure of the contours in the $C(\pi)$ - $S_{1/2}$ plane reflect this behavior. (The fact that the triangle is asymmetric about $C(\pi) = 0$ can be understood in terms of the fact that the simulations are based on the Λ CDM Planck best-fit power spectrum, which has $C(\pi) > 0$, because the largest temperature anisotropy probed is the quadrupole $\ell = 2$ mode which generates a dominant positive term in Eq. (4.15).) Accordingly, the $S_{1/2}$ - R_{27} contours show an echo of that triangular structure, such that for a fixed R_{27} value, lowering $S_{1/2}$ will make it more anomalous. This provides an intuitive way to understand the result derived analytically in Ref. [150].

We note that the panel showing the cross correlation between C_2 and $\log S_{1/2}$ is comparable to that studied in Ref. [151]. We find a looser relation between the two quantities than is found in that work. The reason for this difference is that we mix masking choices for our fiducial synfast map statistics, using full-sky measurements for the C_2 and cut-sky measurements for $S_{1/2}$. We verify that if we measure both quantities on either full-sky or cut-sky maps, the shape of our contours closely resemble the distribution in Ref. [151], which uses UT78 cut-sky measurements of both features.

The results from FFP simulations look visually similar to those in Fig. 4.4. Given this, we do not show the scatter plot for that ensemble. Instead, below we examine the significance of the difference between the feature covariances based on the synfast and FFP sets of simulations.

4.3.1 Covariance structure comparison

Here we measure the covariance between large-angle features measured from our simulated CMB temperature maps, and will compare the covariances extracted from the synfast and FFP simulations. By making this comparison, we can gauge whether the survey properties modeled in the FFP but not in the synfast simulations affect the relationship between the features studied. This in turn could potentially provide insight into whether those survey properties influence observed anomalies in the SMICA map (though the fact that no one has yet found a convincing systematics-based explanation for any of them makes this unlikely).

For a given ensemble of n simulations, we represent each realization as a d -dimensional




Figure 4.5: Left: The feature covariance matrix S_{syn} measured from the synfast simulation ensemble. Right: The difference ΔS (given in Eq. (4.24)), between the covariances measured from the FFP and synfast simulations in units of its sampling error $\sigma_{\text{syn}}^{(1000)}$ estimated from sets 1,000 synfast realizations (defined in Eq. (4.25)).




Figure 4.6: Left: Absolute difference between the covariance matrix measured from 1000 FFP simulations and from 100,000 synfast simulations. Right: Sample variance error bars for covariance matrix entries measured from 100 subsamples of the 100,000 synfast realizations. The right panel of Fig. 4.5 shows the ratio between the two panels of this figure.

vector \mathbf{x} , where $d = 8$ is the number of large-angle quantities measured ($S_{1/2}$, C_2 , C_3 , etc.). Before measuring the covariance matrix, we center and normalize the data so the j th vector component of realization i becomes

$$\tilde{x}_i^{(j)} = (x_i^{(j)} - \bar{x}^{(j)})/\sigma^{(j)}, \quad (4.22)$$

where $\bar{x}^{(j)}$ and $\sigma^{(j)}$ are the mean and standard deviation of the j th quantity being measured, respectively. This ensures that the covariance structure is not dominated by the fact the characteristic size of some of the quantities we study are simply much larger than others (e.g. σ_{16}^2 compared to A_{LV}). It will also mean that the covariance we measure will be equivalent to the correlation coefficients appearing in the subplots of Fig. 4.5.

Once the data are preprocessed, we define a $d \times n$ -dimensional matrix $X = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, where each column corresponds to one realization. This allows us to concisely write the measured $d \times d$ -dimensional feature covariance matrix as

$$S = \frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{x}}_i \otimes \tilde{\mathbf{x}}_i = \frac{1}{n} X X^T. \quad (4.23)$$

We show the covariance matrix for our ensemble of 100,000 synfast simulations S_{syn} on the left side of Fig. 4.5.

We would like to study how the covariance matrix derived from the ensemble of 1,000 FFP simulations S_{FFP} differs from that measured from our 100,000 synfast simulations S_{syn} . To make that comparison meaningful, we must ensure that differences we see are not an artifact of the smaller number of FFP simulations. On the right side of Fig. 4.5 we therefore show the relative difference for covariance matrix entries S_{ij} ,

$$(\Delta S)_{ij} = \frac{[(S_{\text{FFP}})_{ij} - (S_{\text{syn}})_{ij}]}{[\sigma_{\text{syn}}^{(1000)}]_{ij}}, \quad (4.24)$$

where the denominator $[\sigma_{\text{syn}}^{(1000)}]_{ij}$ is sampling error for when S_{ij} is measured from a set of 1,000 synfast simulations.

We estimate $\sigma_{\text{syn}}^{(1000)}$ based on $N = 100$ subdivisions of the 100,000 synfast simulations. This allows us to measure the covariance matrix $S_{\text{syn}}^{(1000, \alpha)}$ for each subsample $\alpha \in \{1, \dots, N\}$. For each entry ij of the matrix, we can then compute the mean over subsamples $[\bar{S}_{\text{syn}}^{(1000)}]_{ij}$,

as well as the sample variance,

$$[\sigma_{\text{syn}}^{(1000)}]_{ij}^2 = \frac{1}{N-1} \sum_{\alpha=1}^N ([S_{\text{syn}}^{(1000,\alpha)}]_{ij} - (\bar{S}_{\text{syn}}^{(1000)})_{ij})^2. \quad (4.25)$$

Thus, assuming the errors on the covariance matrix entries are Gaussian, the values plotted in the right panel of Fig. 4.5 show the difference between the FFP and synfast covariances in units of their 1,000-synfast-realization-based standard deviation. Plots of the absolute difference $S_{\text{FFP}} - S_{\text{syn}}$ and the sampling error $\sigma_{\text{syn}}^{(1000)}$ are shown for completeness in Fig. 4.6.

We see that there are several moderately significant differences between the feature covariances derived using the FFP and synfast simulations. The most prominent of these are between the quadrupole-octopole alignment S_{QO} and $\log S_{1/2}$, C_2 , C_3 , and σ_{16}^2 , with differences ranging from $(2-3.6)\sigma$. There is also a $2.3\text{-}\sigma$ difference in the $S_{1/2}-C_2$ entry, and several other less significant differences in the range $(1-2)\text{-}\sigma$.

Noting that the largest ΔS entries involve the quadrupole, we hypothesized that these differences might be driven by the kinematic quadrupole, which is partially simulated in the FFP maps but not in the synfast simulations. To test this idea, we created an alternative version synfast ensemble where the Doppler quadrupole (DQ) correction a_{2m}^{DQ} given in Table 3 of Ref. [126] is added to each map¹⁹. When we use this synfast+DQ ensemble to reproduce Fig. 4.5 there are no significant changes in the feature covariance matrix or its differences from the FFP feature covariance. If the $S_{\text{FFP}} - S_{\text{syn}}$ differences were mainly driven by the kinematic quadrupole present in the FFP maps, we would expect that adding a DQ correction to the synfast simulations would significantly change the structure of ΔS . Because it does not, we conclude that modeled foregrounds or survey properties other than the kinematic quadrupole are likely to be driving the differences between S_{FFP} and S_{syn} .

4.3.2 Principal component analysis

We next use a principal component analysis (PCA) to investigate whether large-angle CMB anomalies can be reduced to a few fundamental “building blocks” — features, or combinations thereof, which explain the ways that the observed CMB sky is unusual compared to our ensembles of simulations. This search is motivated by the desire to better understand the relationships between the features expected in isotropic ΛCDM , with the hope that this might provide insight into the physical origin of anomalous features in the observed CMB sky.

¹⁹This DQ correction is slightly different than that included in the FFP simulations, which model only the residual frequency-dependent portion of the kinematic quadrupole which is not removed during the Planck map processing.




Figure 4.7: Left: Eigenvectors of the anomaly covariance for our fiducial set of 10^5 synfast simulations. Each column is one PC; the one with index 1 points in the direction of the data's maximum variance. The rows indicate their components in the direction of the initial large-angle CMB quantities. Right: Fractional residual variance for these PCs and the eigenvalues associated with each PC as a fraction of the sum of all eigenvalues

Ref. [94] conjectured that there are three such building blocks in the CMB maps observed by WMAP and Planck: missing large-angle power, alignments between the low multipoles, and dipolar modulation of the CMB (which is a model roughly equivalent to the hemispherical asymmetry studied here). We now have an opportunity to quantitatively test this conjecture by using our measurements of simulation ensembles. By finding the simulation data's principal components (PCs) in the eight-dimensional feature space we consider, we can determine which linear combinations of features explain most of the covariance structure discussed in Section 4.3.1. It is our hope that studying the position of the SMICA in this PC basis will allow us to learn something about the ways in which the observed CMB temperature map is unusual.

PCA is a dimensionality reduction technique which works by identifying the directions in a d -dimensional parameter space along which a set of data points have the maximum variance. The PCs defined sequentially: the first PC corresponds to the direction in which our simulation realizations have the most variance; the second PC corresponds to the direction of maximum variance after the components of the data in the direction of PC 1 are projected out; and so on. In practice, the PCs are obtained by finding the eigenvectors of the data's

covariance matrix. Therefore, in our analysis, we determine the eigenvectors of the covariance matrices S derived above in Eq. (4.23) to get to obtain PCs which are unit-length linear combinations of the quantities $S_{1/2}$, C_2 , C_3 , σ_{16}^2 , R_{27} , $C(\pi)$, S_{QO} , and A_{LV} . The first principal component is the eigenvector with the largest eigenvalue, the second PC has the second largest eigenvalue, and so on.

PCA works as a dimensionality reduction technique because we can capture much of the information about the input data's variance by projecting it onto the first $d' \leq d$ PC directions. Heuristically, the fraction of the information that is retained in this projection is equal to ratio between the sum of the first d' covariance matrix eigenvalues to the sum of all d eigenvalues. To quantify the relative importance of the various PCs, we adopt the complement of this quantity, the fractional residual variance (FRV), the fraction of the variance that is not captured by the first d' PCs. It is given by the expression

$$\text{FRV} \equiv 1 - \frac{\sum_{i=1}^{d'} \lambda_i}{\sum_{i=1}^d \lambda_i}, \quad (4.26)$$

where the eigenvalues have been ordered so that $\lambda_i \geq \lambda_{i+1}$.

Fig 4.7 shows the properties of the PCs derived from our ensemble of 100,000 synfast simulations. The left panel show PCs (the eigenvectors of S_{syn}), with each column corresponding to one PC, and the rows corresponding to their components associated with each of the original eight quantities. The right panel shows the fractional residual variance as a function of the number of PCs retained in the analysis as well as the individual contribution of each PC to its sum.

Looking at the FRV plot, we find that 42% of the simulations' variance is in the direction of the first PC, which quantifies the missing large-angle correlations and has comparable coefficients of the same sign in C_2 , $S_{1/2}$, σ_{16}^2 and $C(\pi)$. Another 20% is the PC 2 direction, which largely lies in the direction of C_3 , along with less dominant contributions by features correlated with the octopole. The next two PCs quantify the sum and difference of the quadrupole-octopole alignment S_{QO} and the hemispherical asymmetry statistic A_{LV} , capturing 13% and 12% of the data's variance, respectively. These first four PCs together explain about 90% of variation in the space of the (eight) features.

Studying the eigenvectors themselves, we find that interpretation of the first four PCs is fairly straightforward because they can be associated with input quantities which are largely independent of one another.

The first PC corresponds to missing large-angle correlations. It is dominated by C_2 , $S_{1/2}$, σ_{16}^2 and $C(\pi)$ (and, to an extent, R_{27}). These particular features are positively correlated,

with correlation coefficients R varying between about 0.5 and 0.8 (see Fig. 4.4), so it is mathematically expected that they would form a principal component whose eigenvector components have same signs and comparable amplitudes, as we observe in PC 1. One can intuitively understand the relationship between these features by noting how their quantities will change if we change the quadrupole amplitude, all else being equal. Raising the quadrupole will increase large-angle power, and will increase the power observed in even-parity modes, and so it make sense that $S_{1/2}$, σ_{16}^2 , $C(\pi)$, and R_{27} will all increase. Thus, one interpretation of the first PC is that it picks out direction in our feature space that is roughly aligned with variations in the quadrupole amplitude.

Next, PC 2 is dominated by the octopole. Given the correlations of C_3 and the other statistics in Fig. 4.4, it is unsurprising that PC 2 receives moderate contributions from $S_{1/2}$ and σ_{16}^2 with the same sign as C_3 , and R_{27} and $C(\pi)$, with the opposite sign. As with PC 1, we can also understand this in terms of how other quantities will respond if we raise or lower C_3 without changing power at other multipoles. More octopole power will generally add to large-angle power, increasing $S_{1/2}$ and σ_{16}^2 , but specifically through odd parity contributions, and so it will lower R_{27} and $C(\pi)$.

The fact that first and second PCs can be associated with the CMB temperature quadrupole and octopole amplitudes, respectively, is unsurprising. We know the lowest multipoles dominate the properties of the CMB temperature map at large angles. We also know that C_2 and C_3 are independent in isotropic Λ CDM, so they correspond to orthogonal directions in our feature space.

The third and fourth PCs are associated with the $a_{\ell m}$ -phase dependent quantities, S_{QO} and A_{LV} . We note that their associated covariance eigenvalues λ_3 and λ_4 are nearly equal, so the ordering of PC 3 and PC 4 is somewhat arbitrary. This reflects the fact that the correlation between S_{QO} and A_{LV} is very small, and means that using PC 3 and PC 4 together is basically equivalent to just defining two unit vectors in the A_{LV} and S_{QO} directions.

The structure of the fifth through eighth PC resists simple interpretation, as they are determined by the relationships between the non-independent quantities, after the variation of the data in the direction of the first four PCS (roughly C_2 , C_3 , S_{QO} and A_{LV}) are projected out. One could infer, for example, that because PC 5 has C_2 and C_3 components with different signs, that it might capture some information about whether the power from the quadrupole and octopole cancel one another, but this is far from clear. PCs 6-8 all have small eigenvalues with $\lambda_6 \sim \lambda_7 \sim \lambda_8$, so their order is somewhat arbitrary and basically just divide up whatever degrees of freedom are left after the first five PCs are removed.

We also performed a PCA on the FFP simulation data. Fig. 4.8 shows the FFP PC's,


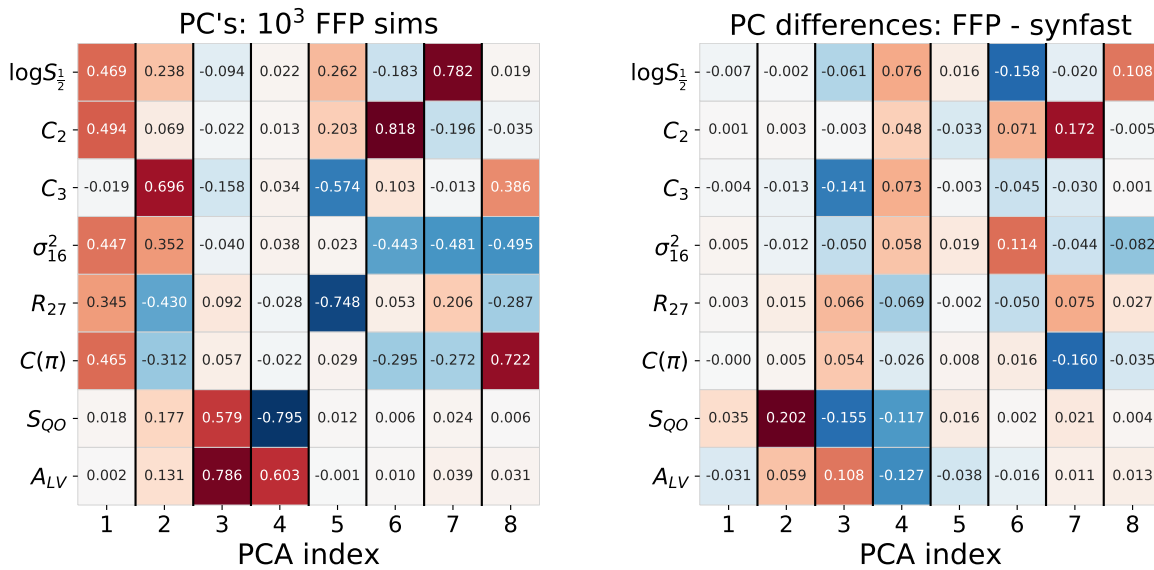


Figure 4.8: PCA results for the FFP simulations and comparison to synfast results. Top left: Eigenvectors of the anomaly covariance for the set of 1000 FFP simulations. These are the principal components. Top right: Differences between the PC's obtained using the FFP simulations and the fiducial Gaussian simulations. Bottom: Absolute and fractional differences between the Fractional Residual Variance (FRV) when studying the anomaly covariance obtained from the FFP rather than synfast simulations.

the difference between the FFP and synfast PC’s, and the shows the differences between their FRV functions. Generally, we find that the FFP results are very similar to those for the synfast simulations. The small differences between them largely reflect their covariance matrices’ differences discussed in Section 4.3.1, but are otherwise difficult to interpret.

We next calculate the probability of the features projected to the PC basis; this can inform whether linear combinations of the features that tend to “come together” in simulated skies are particularly anomalous (or not) when observed on our CMB sky. We proceed as follows: for each simulation, our measurement of the quantities corresponds to a vector in our eight-dimensional feature space. By taking the dot product of that vector with each of the PCs, we find the components of that simulation’s vector in the new PC basis. The grey histograms in Fig. 4.9 show the resulting distributions of the statistics projected to the PC basis, as calculated by synfast (left column) and FFP (right column) simulations. The red vertical lines correspond to the statistics calculated on our CMB sky, again projected to the PC basis.

Fig. 4.9 shows some instructive trends. Since the first principal component (PC 1) is a linear combination of the features that encode the missing angular correlations (low C_2 , $S_{1/2}$, σ_{16}^2 and $C(\pi)$), it makes sense that its probability is low. However the fact that this probability is *lower* than that for any of the individual features (p -value = 0.064%) further indicates that, *given* the lowness of one of its constituent statistics (for example, the quadrupole), the other aforementioned features that make up PC 1 are still lower than expected in Λ CDM. Next, the probability of PC 2 is not anomalous ($p = 47\%$), which is unsurprising given that it largely reflects the rather average C_3 . The PC 3 probability, however, is surprisingly high ($p = 0.032\%$), which is the smallest p-value among all PCs. The extremely high value of the statistics projected to PC 3 comes from the fact that this principal component is largely a sum of S_{QO} and A_{LV} , which are both high on our sky *but uncorrelated* ($R = 0.015$) in Λ CDM. Hence, PC 3 is the sum of two high-valued statistics, and is itself very high. In contrast, PC 4 is largely a difference between the same two high-valued statistics (S_{QO} and A_{LV}), and is itself therefore average ($p = 39\%$). Finally, the higher PCs also do not shed significant further light on the structure of the statistics of the features we study.

In concluding this section, we note that PCA as a method is only able to capture linear structures in the data. Because the relationship between many of the quantities we measure are nonlinear by definition, the PCAs will therefore capture only partial information about the structure in the simulations’ distribution.

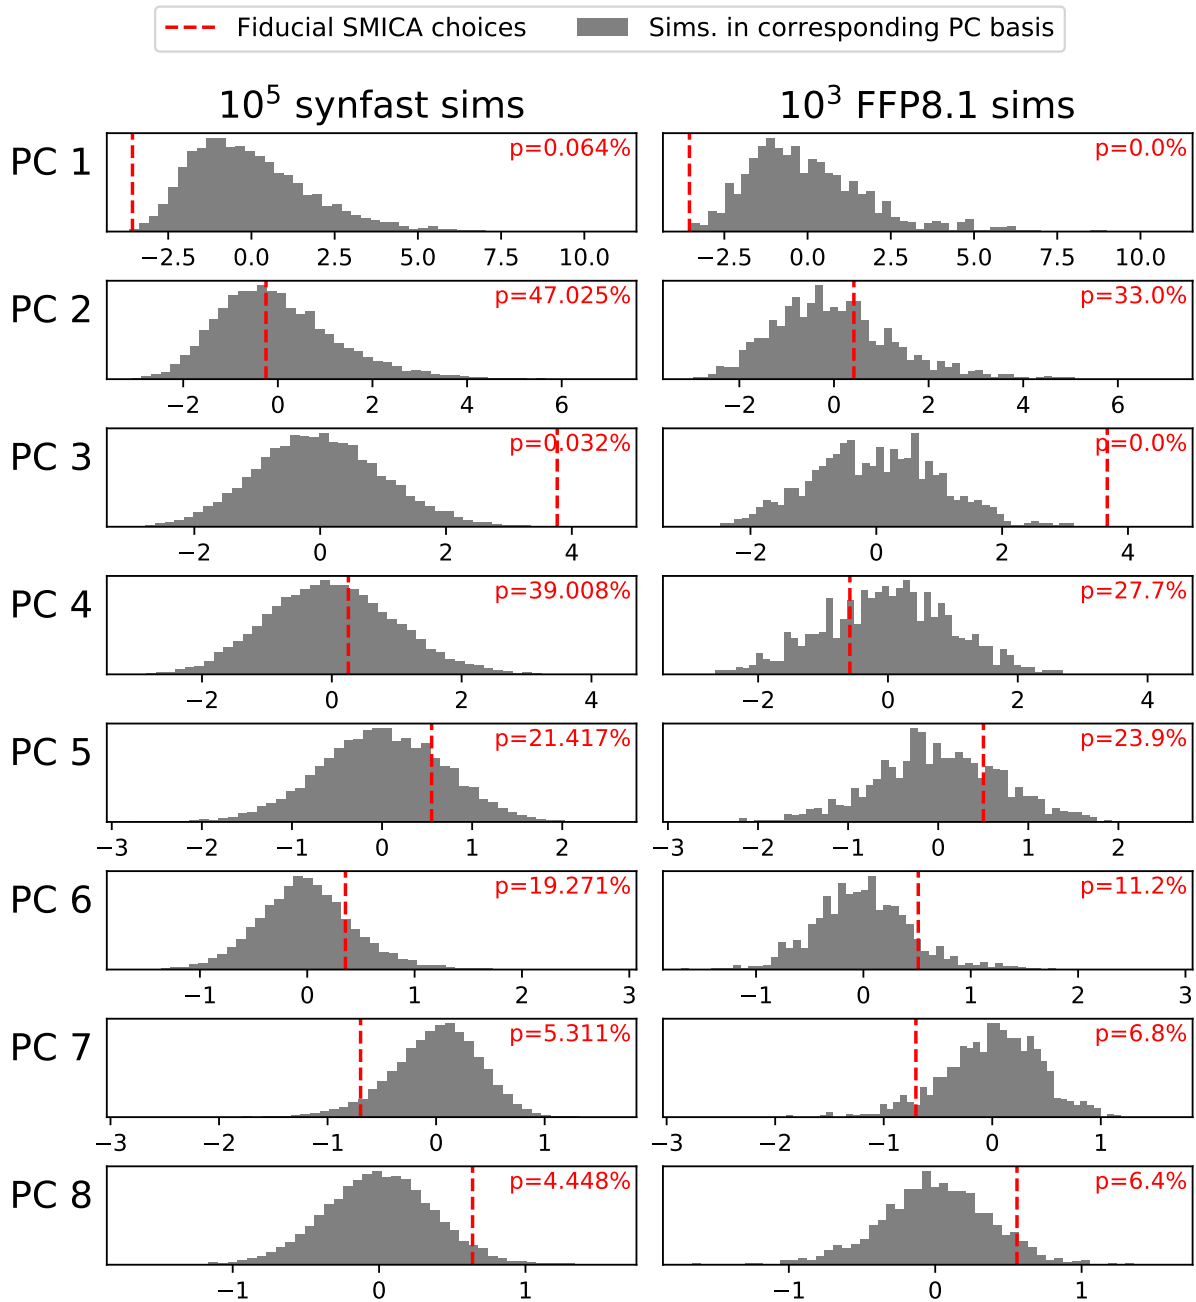


Figure 4.9: Summary of feature statistics projected into our PCA basis. The row labels indicate the basis vectors of the PC basis, which are equivalent to the unit-eigenvectors of the synfast (left) and FFP (right) simulations' feature covariance matrix. The gray histograms show the distribution of the components of simulation realizations in the direction of each PC. The red lines show the same projection of our fiducial SMICA map measurements (using the measurement methods whose p -values in Fig. 4.3 are denoted by an arrow). The red numbers in the top right corner of each panel show the percentage of simulations that are more extreme than the corresponding SMICA measurement.

4.4 Conclusions

In this chapter we have studied the relationships between a set of large-angle CMB features expected in the Λ CDM model, with the goal of better understanding the interdependence of large-angle temperature anomalies observed in WMAP and Planck data. In particular, we have studied the eight features measured through the quantities defined in Section 4.2: the integrated power of temperature fluctuations at angles $\theta > 60^\circ$ ($S_{1/2}$), the quadrupole amplitude (C_2), the octopole amplitude (C_3), the variance of the temperature map evaluated at resolution $N_{\text{side}} = 16$ (σ_{16}^2), the parity statistic R with maximum multipole of $\ell = 27$ (R_{27}), the angular power spectrum at 180° ($C(\pi)$), the quadrupole-octopole alignment (S_{QO}), and the amplitude of the hemispherical asymmetry (A_{LV}). The first six of these features depend on the angular power spectrum and quantify various aspects of angular clustering at large scales, while the last two depend on phases of the $a_{\ell m}$ and quantify large-angle alignments and parity asymmetry observed in CMB maps.

Our analysis was based on on measurements of two ensembles of Λ CDM simulations: 100,000 noiseless Gaussian CMB temperature maps generated using the `synfast` function in `healpy`, and 1000 full focal plane (FFP8.1) simulations provided by the Planck team that contain astrophysical foregrounds and other physical artifacts expected in the observed sky. We began by using these ensembles to find the probability of the observed values of each feature in Λ CDM, which allowed us to study the impact of analysis choices on the feature statistics and to make sure we could recover results from previous work. We found generally excellent agreement between the statistics of the features based upon our two sets of synthetic maps, and summarized the results in Fig. 4.3.

Then, selecting a fiducial set of analysis choices, in Section 4.3 we used those same simulation measurements to fulfill the principal goal of this project by calculating the correlation between the eight features studied. Fig. 4.4 shows, for the first time, a complete covariance of features associated with the most commonly discussed large-angle CMB anomalies. Our results confirm and quantify various aspects of the features that were previously either guessed or calculated in isolation. For example, the quadrupole C_2 , the missing large-angle correlations statistic $S_{1/2}$, and the variance σ_{16}^2 are all positively correlated and largely uncorrelated to the phase-dependent features. The phase-dependent features — the quadrupole-octopole alignment S_{QO} and the hemispherical asymmetry statistic A_{LV} — are uncorrelated both each other and with all other features studied. Less trivially, we find that the covariance between S_{QO} and several other features, is significantly higher when measured from FFP simulations than from `synfast` (though still low compared to the covariances between other features), and

that introducing a kinetic quadrupole correction to the synfast simulations has little impact on that difference.

We then diagonalized the measured covariance matrix, obtaining the principal components of features' expected distribution in Λ CDM. This allowed us to quantify whether most of the information about how the simulations vary in our eight-dimensional feature space is retained in some smaller number of PCs. We find that 42% of the simulations' variance is in the direction of the first PC, which quantifies the missing large-angle correlations and has comparable coefficients of the same sign in C_2 , $S_{1/2}$, σ_{16}^2 and $C(\pi)$. Another 20% is the PC 2 direction, which largely lies in the direction of C_3 , along with less dominant contributions by features correlated with the octopole. The next two PCs quantify the sum and difference of the quadrupole-octopole alignment S_{QO} and the hemispherical asymmetry statistic A_{LV} , capturing 13% and 12% of the data's variance, respectively. These first four PCs together explain about 90% of variation in the space of the (eight) features.

It is important to remind ourselves that apart from the few (generally $2 - 3\sigma$) anomalies discussed here and elsewhere, the Λ CDM model describes most of the current cosmological observations with immense success. Given the significant cosmic variance inherent in the largest angular scales of the CMB, as well as the absence of concrete models that are clearly preferred over the Λ CDM model, we should be wary of putting too much weight on these anomalies as motivations for new physics. However, given the success of Λ CDM, any observational clues as to how to build a more fundamental description of the physics of inflation and dark energy will (initially at least) take the form of small deviations from the predictions of the model [188]. Given this, we should certainly take a careful look at reported tensions and anomalies, making sure we understand how assumptions related to modeling and analysis affect their significance.

It is in this spirit that this work contributes to the discussion of large-angle CMB anomalies: by understanding in detail how observed features are related in Λ CDM, we can better assess the independent ways in which our observed CMB sky is unusual, and thus whether they might provide clues about beyond- Λ CDM physics. An interesting potential avenue for future work is to study how the covariance between the anomalies changes when assuming underlying models that are extensions of, or alternatives to, Λ CDM.

Chapter 5

Multi-probe blinding for DES

We now turn our attention from CMB observables to measurements of the late Universe, focusing in particular on the blind analysis of LSS survey data. The practice of blinding against human bias in data analysis is standard in many areas of science. The idea is to prevent the scientists from biasing their analysis toward results that are theoretically expected or, more generally, deemed by them to be likely or correct. In experimental particle physics strategies for blinding are manifold and have been honed since their earliest application decades ago [189]. Blinding strategies in particle physics include hiding the signal region, offsetting parameters in the analysis by a hidden constant, and adding or removing events from the analysis (for a review, see Ref. [190]).

Blinding started to be applied to astrophysics and cosmology only relatively recently. The first application to cosmology was described in Ref. [191], which reports on an analysis of magnitude-redshift data of type Ia supernovae. In that study, the full analysis was performed with unknown offsets added to the key cosmological parameters, Ω_M and Ω_Λ , until unblinding revealed final parameter values. Many type Ia supernova analyses have adopted some variation of this blinding approach since (e.g. Refs. [41, 192–194]). More recently, blinding has been regularly applied to analyses involving strong gravitational lensing [195, 196], as well as cosmological inferences from weak gravitational lensing observations [197–200].

Our goal here is to develop a blinding technique that can be used for cosmological inferences from the Dark Energy Survey (DES)¹ [46]. DES is an ongoing photometric galaxy survey based at the Blanco 4-meter telescope at the Cerro Tololo Inter-American Observatory in Chile. Designed to obtain precise constraints on cosmological parameters, particularly those associated with the physics of dark energy, when complete, DES will produce a catalog of galaxy shapes and locations over 5000 deg^2 out to $z \sim 1.3$ (roughly 8 billion light years).

¹<http://www.darkenergysurvey.org>

Cosmology results for the analysis of the first year of DES data were made public in August of 2017 [12]. This Year 1 (Y1) analysis was based on measurements in the 1350 deg² square degree area shown in green in the upper panel of Figure 5.1. In it, cosmological parameters were constrained through the combined analysis of three two-point functions: the angular correlation of galaxy positions via $w(\theta)$, the angular correlation of weak lensing shears via $\xi_{\pm}(\theta)$, and the cross correlation between galaxies and shears via $\gamma_t(\theta)$. Constraint contours from this so-called 3×2 pt analysis are shown for the parameter plane where DES has the most constraining power in the bottom panel of Figure 5.1. In that figure, $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$ is a reparameterization of σ_8 chosen to remove degeneracies with Ω_m . Though the centers of the Planck and DES contours are slightly offset, the constraints from the two datasets are statistically consistent. The fact that early-Universe CMB constraints, extrapolated forward using Λ CDM to be compared to DES’ late-Universe LSS measurements is a mark of the success of Λ CDM. This agreement will be further tested in the future.

The next major DES cosmology analysis will be of the survey’s first three years of data, which, covering the full 5000 deg² footprint, will have more than three times the area of the Y1 analysis. Its constraints will therefore have significantly smaller statistical error contributions. Watching how the consistency of Planck and DES constraints change in the DES Year 3 (Y3) analysis will be one of the most significant near-future tests of Λ CDM. In order to take full advantage of the survey’s statistical power, however, we must ensure that our analysis of DES data is robust against systematic errors, including experimenter bias. Given the scrutiny that will be applied to how DES Y3 constraints compare to previous cosmology results (e.g. from Planck), as well as the fact that Λ CDM selects a special value of $w = -1$, blinding is particularly important for precision tests of dark energy.

Determining how to blind the DES cosmology analysis is far from trivial, however. Direct application of techniques from experimental particle physics is not feasible due to numerous important differences between cosmological observations and particle-physics experiments. First, most major tests of cosmology do not use statistics in the form of events counts in a feature space. Second, there is no clear division of the data space into a “signal” region that can be hidden vs. a “control” region that can be used for all validation tests. Thirdly, the DES cosmological inferences are now produced by combination of multiple “probes,” i.e. summary statistics of diverse forms of measurement of different classes of objects. Much of DES’ constraining power comes from its ability to use the combined analysis of multiple observable probes to break degeneracies between cosmological nuisance parameters: for example, analyzing the galaxy-galaxy correlations $w(\theta)$ and the galaxy-shear correlations $\gamma_t(\theta)$ together breaks the degeneracy between galaxy bias $b(z)$ and σ_8 [202]. This


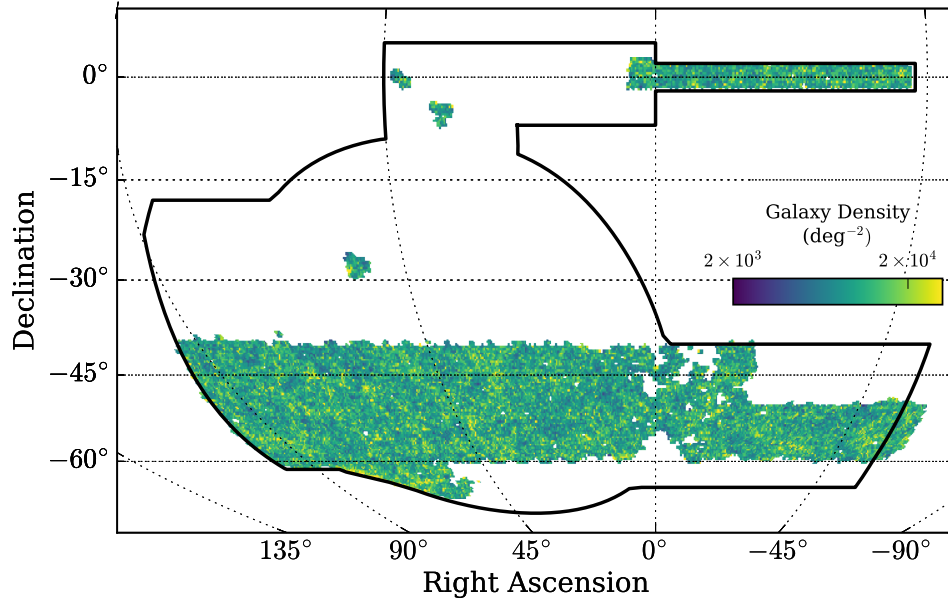


Figure 5.1: Top: Map from Ref. [201] of the region on the sky covered by the analysis of the DES Year 1 data, shown relative to the outline of the full five-year footprint in black. Bottom: DES Year 1 constraints from Ref. [12], showing constraints on $S_8 \equiv \sigma_8 \sqrt{\Omega_m/0.3}$ and Ω_m , in comparison to and combined with Planck CMB constraints.

same fact makes blinding the DES multi-probe analysis challenging because any performing any simple blinding operation to data associated with one observable will cause the total multi-probe data vector to become inconsistent with any viable cosmology.

A new approach is therefore necessary for effective blinding of multi-probe inferences. We describe here an approach based on the manipulation of the entire multi-probe data vector, which for the DES 3×2 pt analysis will be the collection of two-point functions $\xi_{\pm}(\theta)$, $\gamma_t(\theta)$, and $w(\theta)$ measured for a variety of redshift and angular bins. In this strategy, a blinding procedure would be applied to the data vector when it is generated from DES catalogs. That blinded data vector would then be used in all validation tests or cosmological inferences, and only when all tests are passed would the analysis pipeline be rerun with an unblinded version of the data vector.

The outline of the chapter is as follows. First, in Section 5.1 we describe our proposed blinding methodology and the motivation behind it. Then, to provide context, Section 5.2 will introduce the DES 3×2 pt analysis pipeline, describing the procedure for computing model predictions for its data vector. In Section 5.3 we describe the plan for using simulated data to test the performance of this blinding method for the DES Y3 3×2 pt analysis, and in Section 5.4 we present the results. We conclude in Sec. 5.5.

5.1 Method

Broadly speaking, the goal of blinding is to change or hide the output of an analysis in a way that still allows experimenters to effectively perform validation checks on the analysis pipeline. In this section we will discuss the requirements a blinding scheme must fulfill in order to be effective, and then, given that context, will introduce our proposed blinding transformation. Though we are specifically developing and testing the performance of a blinding scheme for the DES 3×2 pt analysis, the ideas we present could in principle be applied to any analysis, so we will frame most of the discussion in terms of a generic data vector. Specifically, Section 5.1.1 will develop a vocabulary for describing the shifts in parameter space needed to overcome experimenter bias, Section 5.1.2 will discuss the interaction between blinding and pipeline validation tests, Section 5.1.3 will introduce our proposed blinding transformation, and Section 5.1.4 will define a metric for assessing its performance.

5.1.1 Prior and prejudice

Let us assume that the experiment produces a vector \mathbf{d} of observed quantities, and we wish to constrain the parameters Θ of a model $\hat{\mathbf{d}}(\Theta)$ for these data. The parameters can include astrophysical and instrumental nuisances as well as the cosmological parameters of interest. There will always be some prior probability, $\text{Pri}(\Theta)$, that expresses the physical bounds of our model (e.g. $\Omega_m > 0$) and results of trusted previous experimentation. In a Bayesian view, the purpose of the experiment is to produce a likelihood function $\mathcal{L}(\mathbf{d}|\Theta)$ that is combined with the prior to produce a posterior measure of belief across the model space, $P(\Theta|d) \propto \mathcal{L}(\mathbf{d}|\Theta) \text{Pri}(\Theta)$. One easily-visualized variant of the prior is to have it be uniform over some model space \mathcal{M} of Θ and zero elsewhere, i.e \mathcal{M} encompasses all parameter vectors considered feasible.

The experimenters may additionally harbor prejudices about the “correct” values of the parameters, for instance that they should agree with some theoretical prejudice such as a flat Universe, or that they should agree with some previous experiment that one is trying to confirm. We can express these prejudices with another (albeit, harder to quantify) probability function $\text{Pre}(\Theta)$. It could for example be a uniform distribution over some region $\mathcal{M}_{\text{Pre}} \subset \mathcal{M}$. Note that in this framing, one must make a decision regarding previous experiments’ results: either we accept them as true and place them in Pri ; or we are using their comparison to our results as a test of our model, in which case we must be wary of confirmation bias and should place them in Pre .

The danger of experimenter bias arises when choices about the analysis process are made, consciously or otherwise, on the basis of whether the experiment’s results conform to the prejudices, i.e. whether $\Theta \in \mathcal{M}_{\text{Pre}}$. To confound the experimenter bias, a blinding procedure will apply a transformation $\mathbf{d} \rightarrow \tilde{\mathbf{d}} = B(\mathbf{d})$ to the data before the experimenters perform analyses. The first critical property of B is therefore that it must destroy the experimenter’s ability to know whether the data are consistent or inconsistent with their prejudices. For example, if we take the maximum-posterior parameter values for blinded and unblinded data

$$\Theta_{\text{unbl}} = \underset{\Theta}{\text{argmax}} \{P(\Theta|\mathbf{d})\} \tag{5.1}$$

$$\Theta_{\text{bl}} = \underset{\Theta}{\text{argmax}} \{P(\Theta|B(\mathbf{d}))\}, \tag{5.2}$$

then there must be a non negligible chance that either

$$\text{Pre}(\Theta_{\text{unbl}}) \gg \text{Pre}(\Theta_{\text{bl}}) \quad \text{or} \quad \text{Pre}(\Theta_{\text{unbl}}) \ll \text{Pre}(\Theta_{\text{bl}}). \tag{5.3}$$




Figure 5.2: Cartoon of model and data spaces that we consider when thinking about how to blind an analysis. \mathcal{M} is the space of all viable model parameter sets Θ , which projects onto the data vector space $\mathcal{D}_{\text{Pri}} \subset \mathcal{D}$, where \mathcal{D} is the space of all possible data vectors. \mathcal{M}_{Pre} describe the region in parameter space where experiments have some preconceived expectation for Θ to be, which project onto the \mathcal{D}_{Pre} of data vectors. An effective blinding transformation must have the possibility of moving the observed data vector \mathbf{d} in or out of \mathcal{D}_{Pre} without moving it out of \mathcal{D}_{Pri} .

A simple graphical illustration is given in Figure 5.2: if we define \mathcal{D}_{Pre} as the region of data space \mathcal{D} produced by parameter values within the prejudice region \mathcal{M}_{Pre} , then the blinding transformation must be able to move data into and out of this region. The experimenters should believe that this is possible, but not know for certain whether it has happened.

5.1.2 Preserving the ability to check for errors

In addition to obscuring the true parameter output of an analysis, an effective blinding scheme must still allow experimenters to, before unblinding, examine the data \mathbf{d} to uncover errors in their analysis procedures. A validation test is one whose failure indicates that data could not have been produced by any allowed parameters $\Theta \in \mathcal{M}$. For a non-trivial validation test to exist, there must be some redundancy in the data, namely the region \mathcal{D}_{Pri} comprising the image of \mathcal{M} under the model mapping must be a proper subset of the full space of possible data, as is illustrated in Fig. 5.2. The most general possible validation test is, in fact, to ask whether $\mathbf{d} \in \mathcal{D}_{\text{Pri}}$ (for the case when $\text{Pri}(\Theta)$ is uniform over \mathcal{M}). More generically, we can imprecisely state this in terms of a requirement on the posterior,

$$\max_{\Theta \in \mathcal{M}} \mathcal{L}(\mathbf{d}|\Theta) \text{Pri}(\Theta) \sim \left\langle \max_{\Theta \in \mathcal{M}} \mathcal{L}(\mathbf{d}'|\Theta) \text{Pri}(\Theta) \right\rangle, \quad (5.4)$$

where the average on the right is taken over realizations of data \mathbf{d}' drawn from the distributions predicted by $\Theta \in \mathcal{M}$. In other words, if our model (for both physics and systematics) is able to accurately describe \mathbf{d} , the goodness of fit for the prediction of *some* set of parameters should be comparable to that for realizations of simulated data generated using the same model.

A blinding transformation B should not alter the conclusions of validation tests. There are a number of ways of stating this requirement. Generally, $B(d)$ should map the allowed region \mathcal{D}_{Pri} onto itself, and likewise for its complement, the disallowed region $\tilde{\mathcal{D}}_{\text{Pri}}$. Additionally, the maximum-posterior values from Equations 5.2 should obey

$$\frac{\mathcal{L}(\mathbf{d}_{\text{bl}}|\Theta_{\text{bl}})}{\mathcal{L}(\mathbf{d}_{\text{unbl}}|\Theta_{\text{unbl}})} \approx 1. \quad (5.5)$$

A transformation satisfying these requirements will ensure that blinding will not alter experimenters' judgement about whether there are flaws in the data.

Sometimes the validation tests are more practically expressed as some projection of the data onto a “null test” $T(\mathbf{d})$ such that

$$T(\mathbf{d}) = 0 \quad \forall \Theta \in \mathcal{M}. \quad (5.6)$$

Many kinds of validation tests fall into this paradigm. For example, if T projects onto the B mode (divergence-free) of weak lensing, it should be zero within errors. We can measure the difference between data vectors split by some property uncorrelated with extragalactic signals, such as seeing. Another very generic test is to run the parameter inference on two subsets of the data vector and check that the results are consistent with common Θ . Allowances must of course be made for the expected noise in the null test output at fixed Θ . Generally speaking, a useful blinding transformation must yield $T(\mathbf{d}_{\text{blind}}) = 0$ within errors if and only if $T(\mathbf{d}) = 0$ within errors.

When defining validation criteria it is important to carefully specify the range of models where those criteria are considered viable. For example, suppose we compare the high- and low-redshift halves of a supernova Hubble diagram. If both halves are fit with a Λ CDM model and the data truly are from a Λ CDM universe, then analyzing the two halves separately should produce consistent cosmology results, making this comparison is a useful validation test. If, however, the universe is not described by Λ CDM, then the high/low z split can yield inconsistent results in the absence of processing errors. (Similarly, we can note that the original discovery of dark energy was effectively a demonstration that fitting supernovae

data with $\Lambda = 0$ produced this kind of mismatch.) This example underlines an important point: defining validation tests requires one to make implicit modeling choices, and defining a blinding procedure which preserves the result of those tests can only produce shifts in model space which respects those choices. When constructing an analysis pipeline, it is therefore important to carefully consider what measurements will be considered signals for testing Λ CDM and which can be used as checks on the performance of the analysis pipeline. In other words, defining validation criteria, and a blinding scheme that preserves their results, requires one to specify the space of models that are considered viable.

With a transformation B in hand that satisfies both Eq. (5.3) and Eq. (5.5), experimenters can work with \mathbf{d}_{bl} until all validation tests are passed, remaining ignorant of the nature of the transformation. Once the analysis pipeline has been validated, the decision can be made to unblind by re-running the analysis using the true, unblinded data \mathbf{d} for parameter inference.

5.1.3 Proposed data vector transformation

The first choice to make in picking the blinding transformation B is what data will be transformed. In DES, the data start as pixel values; then are converted to cataloged fluxes, shapes, etc.; then to summary statistics such as the tomographic weak lensing correlation functions $\xi_{\pm}(\theta)$, or the counts $N(\lambda, z)$ of galaxy clusters; and finally to the parameter estimates $\hat{\Theta}$ themselves. The simplest case is simply for B to operate on the $\hat{\Theta}$, as has been done in previous experiments, such as the DES Year 1 analyses, where an unknown shift was applied to cosmological parameter values in any human-readable results. The risk of accidental revelation of the true parameter estimates is high, however, if the blinding code is mistakenly omitted. The temptation for experimenters to peek at the true results is also high when the “curtain” is so thin. Furthermore, in this scenario, the blinding is defeated if anyone plots a theoretical model for the prejudicial model atop the summary statistics.

As a rule of thumb, the earlier in the analysis process the blinding is applied, the stronger the protection against confirmation bias. This is because blinding data earlier in the pipeline means it requires more steps to produce unblinded results in a form that an experimenter could recognize as conforming to their biases or not. Blinding by alteration of the pixel data is probably impossible, apart from substituting an entire set of simulated data for the real one. Blinding at the catalog level is possible in some cases, e.g. when only one or two summary statistics are going to be derived from the catalog. This has already been done for the DES shear-only analyses for both the Year 1 and Year 3 analyses, by scaling each galaxy ellipticity by some unknown multiplicative factor. For a multi-probe experiment,

however, such as DES' combined probe analysis with hundreds of summary statistics, we have discovered no catalog transformation that preserve the validity of the data. We therefore opt to transform the summary statistics.

Our blinding transformation is as follows: let $\hat{d}_i(\Theta)$ be the theoretically computed (noiseless) value of element i of the data vector for parameters Θ . We choose a known reference model Θ_{ref} and a blind shift $\Delta\Theta$ in the cosmological parameters to yield

$$\Theta_{\text{shift}} = \Theta_{\text{ref}} + \Delta\Theta. \quad (5.7)$$

The blinding operation is a simple modification of each element d_i of the data vector, which can either be done additively via

$$B(d_i) = d_i + f_i^{(\text{add})}, \quad (5.8)$$

$$f_i^{(\text{add})} = \hat{d}_i(\Theta_{\text{shift}}) - \hat{d}_i(\Theta_{\text{ref}}), \quad (5.9)$$

or multiplicatively via

$$B(d_i) = f_i^{(\text{mult})} d_i, \quad (5.10)$$

$$f_i^{(\text{mult})} = \frac{\hat{d}_i(\Theta_{\text{shift}})}{\hat{d}_i(\Theta_{\text{ref}})}. \quad (5.11)$$

If the expected noise level on \mathbf{d} does not vary much across the parameter shift $\Delta\Theta$, then it is true by construction that the B will map data generated at Θ_{obs} into viable data for $\Theta_{\text{bl}} = \Theta_{\text{obs}} + \Delta\Theta$ if the truth (Θ_{obs}) is sufficiently close to the reference cosmology (Θ_{ref}).

5.1.4 Considerations for evaluating performance

Showing that these kinds of blinding transformations work for $\Theta_{\text{obs}} \neq \Theta_{\text{ref}}$, and for realistic noise levels requires validation, however. Even for noiseless data, it need not be true that B maps data for an arbitrary $\Theta_{\text{obs}} \in \mathcal{M}$ into viable data for some Θ_{bl} . The imperfection in B can be quantified using Equation 5.5:

$$\Delta\chi^2 \equiv \min_{\Theta \in \mathcal{M}} \{-\log \mathcal{L}(B(\mathbf{d})|\Theta)\} - \min_{\Theta \in \mathcal{M}} \{-\log \mathcal{L}(\mathbf{d}|\Theta)\}. \quad (5.12)$$

If indeed the blinding transformation can have $\Delta\Theta$ larger than the prejudicial region \mathcal{M}_{Pre} which conserving validation results by $\Delta\chi^2 \lesssim 1$, then it satisfies all of our blinding needs.

We can examine this metric explicitly for the case of a Gaussian likelihood which has

$$\chi^2(d, \Theta) = -\log \mathcal{L}(\mathbf{d}|\Theta) = \frac{1}{2} \sum_{ij} \left(d_i - \hat{d}_i(\Theta) \right) [\text{Cov}^{-1}]_{ij} \left(d_j - \hat{d}_j(\Theta) \right) - \log \mathcal{L}_0, \quad (5.13)$$

where Cov is the data covariance, which may depend on Θ , $\mathcal{L}_0 \equiv [(2\pi)^N |\text{Cov}|]^{-1/2}$ is the likelihood's normalization, and the indices i and j run over all N data vector elements. It will be convenient for this discussion to decompose the data into signal and noise components, $\mathbf{d} = \mathbf{s} + \mathbf{n}$, where the signal $\mathbf{s} \equiv \hat{\mathbf{d}}(\Theta_{\text{obs}})$ and the noise may have some dependence on model parameters. Inserting these quantities into Eq. (5.12), and assuming that to good approximation \mathcal{L}_0 is independent of Θ , we can show that

$$\Delta\chi^2 = \Delta\chi_{\text{signal}}^2 + \Delta\chi_{\text{noise}}^2 \quad (5.14)$$

where

$$\Delta\chi_{\text{signal}}^2 = \frac{1}{2} \sum_{ij} \left[(s'_i - \hat{d}_i) [\text{Cov}^{-1}]_{ij} (s'_j - \hat{d}_j) \Big|_{\Theta_{\text{bl}}} - (s_i - \hat{d}_i) [\text{Cov}^{-1}]_{ij} (s_j - \hat{d}_j) \Big|_{\Theta_{\text{unbl}}} \right] \quad (5.15)$$

$$\Delta\chi_{\text{noise}}^2 = \frac{1}{2} \sum_{ij} \left[(n'_i [\text{Cov}^{-1}]_{ij} n'_j - n_i [\text{Cov}^{-1}]_{ij} n_j) \right]. \quad (5.16)$$

Here we have defined

$$B(d_i) = s'_i + n'_i \quad (5.17)$$

such that for additive blinding according to Eq. (5.8),

$$s'_i = s_i + f_i^{(\text{add})} \quad \text{and} \quad n'_i = n_i, \quad (5.18)$$

and for multiplicative blinding according to Eq. (5.10),

$$s'_i = f_i^{(\text{mult})} s_i \quad \text{and} \quad n'_i = f_i^{(\text{mult})} n_i. \quad (5.19)$$

In the case of noiseless data, we have $d_j = s_j = \hat{d}_j(\Theta_{\text{obs}})$, and we will have the best fit parameters $\Theta_{\text{unbl}} = \Theta_{\text{obs}}$. Given this, we can see that $\chi^2 = 0$ when evaluated at the maximum likelihood parameter values Θ_{unbl} . Even in the case of data with noise, where it is possible that $\Theta_{\text{unbl}} \neq \Theta_{\text{obs}}$, our expectation value over many realizations is that contributions to $\Delta\chi_{\text{signal}}^2$ from the unblinded data will be close to zero. Thus, $\Delta\chi_{\text{signal}}^2$, or equivalently $\Delta\chi^2$ in the noiseless case, will quantify the size of residuals of the best fit model to the blinded

signal.

In the presence of noise, we must address a difference in the behavior of additive and multiplicative blinding. We can see that additive blinding according to Eq. (5.8) will not affect the noise contribution, so $n'_i = n_i$ and $\Delta\chi_{\text{noise}}^2 = 0$. In contrast, multiplicative blinding will scale the noise to $n_i \rightarrow f_i^{(\text{mult})}n_i$. If we do not otherwise modify the likelihood calculation, this will degrade the quality of fit to the blinded data.

This motivates us to add an additional step to multiplicative blinding: in addition to scaling the data vector according to Eq. (5.10), we scale the data covariance via

$$B(\text{Cov}_{ij}) = f_i^{(\text{mult})}\text{Cov}_{ij}f_j^{(\text{mult})}. \quad (5.20)$$

For covariance entries dominated by noise such that $\text{Cov}_{ij} \approx \langle n_i n_j \rangle$, this scaling will capture the effect of blinding and ensure that $\Delta\chi_{\text{noise}}^2 \approx 0$. If there are significant signal contributions to Cov_{ij} , the interpretation of this scaling is less clear and its correction for blinding's impact on the noise contributions will be more imperfect.

The advancing precision in cosmological experiments works in favor of both the additive and multiplicative blinding schemes. The values of $\Delta\Theta$ necessary to confound confirmation bias, as well as the span $\Theta_{\text{obs}} - \Theta_{\text{ref}}$ of the viable model space, will shrink linearly with the typical parameter uncertainty σ_{Θ} in our experiments. Additionally, as the statistical power of surveys increase, the noise contributions \mathbf{n} will become less significant.

To summarize, the blinding procedure is:

1. Choose a reference cosmology (and nuisance parameters) Θ_{ref} in the middle of the range of models considered feasible truths.
2. Select a (blind) shift $\Delta\Theta$ from a distribution broader than the preconceptions causing the confirmation bias. For example if there is a theoretical prejudice for $w = -1$, then Δw should be capable of shifts 4–5 \times the experiment's uncertainty in w .
3. For each summary statistic d_i being used for cosmological inference, calculate the blinding factor f_i from either Eq. (5.9) or Eq. (5.11).
4. Hide the real data d_i and give experimenters the blinded values (either $d_i + f_i^{(\text{add})}$ or $f_i^{(\text{mult})}d_i$) with which to conduct all validation tests. If performing multiplicative blinding, also scale the covariance matrix provided to experimenters according to Eq. (5.20).
5. After passing validation tests, unblind by using the original data \mathbf{d} for a final inference of Θ .




Figure 5.3: Plot from Ref. [12] of the redshift distributions for lens and source galaxies in the DES Y1 3×2 pt analysis. The vertical colored bands show the nominal redshift range of each bin, while the lines show the estimated redshift distribution. The black lines show the unbinned total distribution. We adopt these same dn/dz functions for our Y3 blinding tests.

5.2 The DES 3×2 pt analysis

Using the discussion above as a guide, we will test this blinding strategy explicitly for simulated DES Y3 3×2 pt data in Section 5.4. As background for those tests, this section will give an overview of the pieces of DES analysis relevant to simulating 3×2 pt data vectors and performing parameter estimation. As the analysis choices for Y3 are not yet fixed, we will approximate the Y3 data vector and covariance by using the same analysis choices and survey as the ongoing Y1 analysis, but with the sky coverage increased from 1350 deg^2 to the full 5000 deg^2 .

The 3×2 pt data vector is based on observations of two populations of galaxies. Positions are measured for a set of so-called lens galaxies which have been selected to have small photo- z errors and which have been carefully checked for residual systematics. In the Y1 analysis, this population consisted of 650,000 bright red sequence galaxies which are selected as part of the redMaGiC catalog [203]. Their redshift distribution is shown in the upper panel of Fig. 5.3. Cosmic shears are measured from a larger population of source galaxies. For the Y1 analysis the source galaxies included 26 million objects selected from the Y1

Gold catalog [201] and their shapes are measured using one of two algorithms known as Metacalibration and Im3shape [204]. The lens galaxies are divided into five redshift bins, and source galaxies are divided into four redshift bins. Their dn/dz distributions are shown in the bottom panel of Fig. 5.3.

The full 3×2 pt data vector consists of three types of two-point functions measured from the lens and source catalogs. The galaxy-galaxy correlations are measured as autocorrelations within each lens bin, producing a set of functions $w^i(\theta)$ for $i = 1-5$. Shear-shear correlations are measured for all auto and cross correlations of the source bins, producing functions $\xi_+^{ij}(\theta)$ and $\xi_-^{ij}(\theta)$ for $i = 1-4$ and $1 \leq j \leq i$. (Since these are correlations between the same tracer, $\xi_{\pm}^{ij} = \xi_{\pm}^{ji}$, we require $j \leq i$ to avoid including redundant information.) The galaxy-shear cross correlations are measured between all combinations of the five lens bins and four source bins, producing $\gamma_t^{ij}(\theta)$ for $i = 1-5$ and $j = 1-4$. (Note that these are cross correlations between *different* tracers, so $\gamma_t^{ij} \neq \gamma_t^{ji}$.) All of these two-point functions are measured for twenty logarithmically spaced angular bins between 2.5 and 250 arcmin. We adopt the same angular bins, and also use the same scale cuts as the DES Y1 analysis [67] in order to remove scales which are impacted by modeling uncertainties associated with nonlinear physics. These scale cuts can be seen as vertical gray bands in Figs. 5.6 and 5.7. The resulting data vector used for parameter estimation has 457 data points.

The DES 3×2 pt likelihood is modeled as a multivariate Gaussian. Its covariance has significant off-diagonal contributions, since the same source and lens populations are used to calculate multiple two-point function data vector entries. Computing the 3×2 pt covariance is a rather involved process, and so we do not compute it ourselves for this analysis. Instead, we adapt the covariance matrix that was previously analytically computed for the Y1 analysis using Cosmolike [205] via the method described in Ref. [67]. To approximate the Y3 covariance, we simply scale the Y1 covariance by a factor of $0.27 = 1350/5000$ to account for Y3's increased survey area. This survey-area scaling correctly modifies the Gaussian parts of the covariance, but it does not properly scale the non-Gaussian contributions [206]. Thus, this is only a rough approximation for the Y3 covariance. It should be sufficient for our purposes, however. Though in principle the 3×2 pt covariance depends on the model parameters, it has been shown [207] that the covariance's cosmology dependence can be neglected without significantly affecting parameter constraints. Because of this, in the DES Y1 3×2 pt analysis and in this work, we treat the data covariance matrix as constant (modulo any scaling that is applied as part of the blinding process) and do not vary it when performing parameter searches.

In addition to cosmological parameters, the model required to make predictions for the $3 \times$

2pt data vector includes a set of nuisance parameters used to account for various systematic uncertainties. Again, we adopt the same parameter choices as the DES Y1 3×2 pt analysis, as it is described in Ref. [67]. We use a constant galaxy bias b_i for each lens redshift bin, where $i = 1-5$ labels the bin. Additionally, in order to model the effect of uncertainties in photo- z estimation, a parameter Δz_i is introduced for all five lens bins and all four source bins, which produces an redshift offset for their dn/dz distributions. Another set of nuisance parameters quantify shear calibrations, defined so the measured shear for a galaxy is $\gamma_{\text{meas}} = (1+m)\gamma_{\text{true}}$. Following the Y1 analysis, we assign one shear calibration nuisance parameter per source galaxy bin, which modify the two-point functions via

$$\xi_{\pm}^{ij}(\theta) \rightarrow (1+m_i)(1+m_j)\xi_{\pm}^{ij}(\theta), \quad \text{and} \quad \gamma_t^{ij}(\theta) \rightarrow (1+m_j)\gamma_t^{ij}(\theta). \quad (5.21)$$

The last set of nuisance parameters we include are A_{IA} , α_{IA} , and $z_0^{(IA)}$, which are input for a model of how intrinsic (as opposed to lensing-induced) alignments between galaxies affect their observed two-point functions. They modify the shear convergence weight function (Eq. (2.52)) according to a linear alignment model via

$$W_{\kappa}^i(z) \rightarrow W_{\kappa}^i(z) - \left[A_{IA} \left(\frac{1+z}{1+z_0} \right)^{\alpha_{IA}} \frac{C_1 \rho_{m0}}{D(z)} \right] \frac{dn^i}{dz}, \quad (5.22)$$

where $C_1 = 0.0134/\rho_{\text{crit}}$ is a normalization constant calibrated based on previous observations [208].

The fiducial values for all of these nuisance parameters, as well as cosmological parameters, are shown in Table 5.2. In the fiducial DES Y1 3×2 pt analysis, the number of neutrinos $N_{\text{massive } \nu}$ and $N_{\text{massless } \nu}$ (chosen to sum to the standard model effective number of neutrinos N_{eff}), the optical depth of the CMB τ , and $z_0^{(IA)}$ were fixed, while the rest of the parameters shown in the table were varied. Thus, the Y1 3×2 pt parameter estimation was done for a total of 26 free parameters. The Y1 parameter estimation used Gaussian priors on the various photo- z shifts Δz_i and shear calibrations m_i , and flat priors on the rest of the parameters.

For our validation studies, we will use a slightly modified version of the DES Y1 analysis pipeline, run using software CosmoSIS² [7], in order to compute model predictions for the 3×2 pt data vector. (The modules, or calculation steps, which comprise that pipeline are enumerated below in Section 5.2.1.) In order to simulate more realistic data, we generate a noise realization with its entries drawn from the multivariate Gaussian distribution defined

²bitbucket.org/joezuntz/cosmosis/wiki/Home

	Parameter	Fid. value	Prior range	Varied during
blinded parameters {	σ_8	0.837	[0.1, 0.9]	shift, obs, fit
	w	-1.0	[-2.0, -0.0]	shift, obs., fit
lens galaxy bias {	Ω_m	0.295	[0.1, 2.0]	obs., fit
	h	0.6882	[0.2, 1.0]	obs., fit
	b_1	1.45	[0.8, 2.5]	obs., fit
	b_2	1.55	[0.8, 2.5]	obs., fit
	b_3	1.65	[0.8, 2.5]	obs., fit
	b_4	1.8	[0.8, 2.5]	obs., fit
	b_5	2.0	[0.8, 2.5]	obs., fit
	Ω_b	0.0468	-	-
n_s	0.9676	-	-	
	$\Omega_\nu h^2$	6.166×10^{-4}	-	-
	$N_{\text{massive } \nu}$	3	-	-
	$N_{\text{massless } \nu}$	0.046	-	-
shear	τ	0.08	-	-
calibrations {	m_1, m_2, m_3, m_4	0.012	-	-
intrinsic alignment model {	A_{IA}	0.0	-	-
	α_{IA}	0.0	-	-
	$z_0^{(IA)}$	0.62	-	-
photo- z bias for source galaxies {	$\Delta z_1^{\text{source}}$	-0.002	-	-
	$\Delta z_2^{\text{source}}$	-0.015	-	-
	$\Delta z_3^{\text{source}}$	0.007	-	-
	$\Delta z_4^{\text{source}}$	0.018	-	-
photo- z bias for lens galaxies {	Δz_1^{lens}	0.002	-	-
	Δz_2^{lens}	0.001	-	-
	Δz_3^{lens}	0.003	-	-
	Δz_4^{lens}	0.0	-	-
	Δz_5^{lens}	0.0	-	-

Table 5.1: Fiducial parameter values, their prior ranges if they are varied in this work, and a note on when they are varied for our blinding study (which will be described in Section 5.3). In the last column, “shift” means we vary $\Delta\Theta$ in that direction when creating blinding factors, “obs.” means we vary the parameter when we draw Θ_{obs} for “true” cosmologies, and “fit” means we vary that parameter when performing parameter estimation fits. If no range is given for a parameter, it is held fixed. All fiducial values and priors for varied parameters are chosen to match the settings for the DES Y1 3×2 pt analysis.

by the data covariance matrix, and add that to the model data vector. The dotted and solid black lines in Figs. 5.6 and 5.7 show an example of one such simulated data vector.

5.2.1 DES 3×2 pt CosmoSIS pipeline

The software CosmoSIS provides a modular framework for cosmological parameter estimation. Calculations within it are built around the idea of a pipeline, the set of all calculations needed to go from a set of model parameters Θ to a prediction for the two-point function data vector $\hat{\mathbf{d}}$, and if desired, an evaluation of the likelihood function for a measured realization of the data \mathbf{d} . If simply run once, the pipeline can be used simply to compute $\hat{\mathbf{d}}$, but CosmoSIS’s main functionality is to run the pipeline many times through a sampler, some algorithm which performs parameter estimation by computing the posterior at many points in parameter space.

The calculations that make up a CosmoSIS pipeline are divided into discrete steps, where each is performed by a separate software module. We will compute the data vectors using the nearly same settings as in Refs. [12, 67]. Here we present a brief description of the modules we use to compute the 3×2 pt data vector:

1. **consistency** – Checks that input cosmological parameters are consistently defined, and that enough parameters are defined to perform calculations.
2. **camb**: Runs CAMB [60, 61] to compute linear matter power spectrum $P(k, z)$ for a grid of redshifts.
3. **sigma8_rescale**: Scales amplitude of $P(k, z)$ to be consistent with input σ_8 value. The inclusion of this module is the only way in which our pipeline differs from the Y1 3×2 pt analysis, which samples over A_s rather than σ_8 .
4. **halofit**: Use halofit [62] to correct $P(k, z)$ for the effects of nonlinear structure growth.
5. **growth**: Computes the linear growth function $D(z)$ and growth rate $f(z)$ for a grid of z values.
6. **extrapolate**: Extrapolates $P(k, z)$ linearly in log-space up to a specified k_{\max} to allow integration bounds to be extended.
7. **fits_nz**: Reads in dn/dz distributions for all redshift bins from a specified data file.
8. **lens_photoz_bias**: Uses the nuisance parameters Δz_i to shift the redshift bins i for lens galaxy populations to model the effects of photo- z redshift estimation errors.

9. `source_photoz_bias`: The same as above, but for the source galaxy bins.
10. `unbiased_galaxies`: Creates arrays for the galaxy power spectrum by copying the matter power spectrum.
11. `bias_neutrinos`: Computes a small, scale dependent correction to galaxy bias due to the effect of massive neutrinos.
12. `multiply_pk`: Incorporates the neutrino bias corrections into the galaxy power spectrum.
13. `IA`: Uses a linear alignment model to create arrays for the effective power spectrum corrections that will be used to compute the modifications to $C_{\kappa\kappa}(\ell)$ and $C_{g\kappa}(\ell)$ due to intrinsic alignments.
14. `ia_z_field`: Scale the intrinsic alignment power spectrum contributions according to a power law in $(1+z)$ with slope α_{IA} .
15. `pk_to_cl`: Projects three dimensional power spectrum $P(k, z)$ to various angular power spectra $C_{xx}(\ell)$ using the Limber approximation.
16. `bin_bias`: Applies a constant bias b_i for $i=1-5$ for each of the lens galaxy bins.
17. `add_intrinsic`: Adds corrections from intrinsic alignments to $C_{\kappa\kappa}(\ell)$ and $C_{g\kappa}(\ell)$.
18. `shear_m_bias`: Models shear measurement bias using the shear calibration nuisance parameters m_i for $i=1-4$ for each of the source redshift bins.
19. `2pt_gal`: Computes $w(\theta)$ from $C_{gg}(\ell)$ using the function `tpstat_via_hankel` from the `nicaea` software³ [209].
20. `2pt_gal_shear`: Computes $\gamma_t(\theta)$ from $C_{g\kappa}(\ell)$.
21. `2pt_shear`: Computes ξ_{\pm} from $C_{\kappa\kappa}(\ell)$.
22. `2pt_like`: Given computed theory data vector, use measured data vector, the data covariance, and priors to evaluate the posterior. (This is only done if the pipeline is being used for parameter estimation.)

This pipeline takes approximately 30 seconds to run. The computing time is dominated by the call to CAMB.

³www.cosmostat.org/software/nicaea

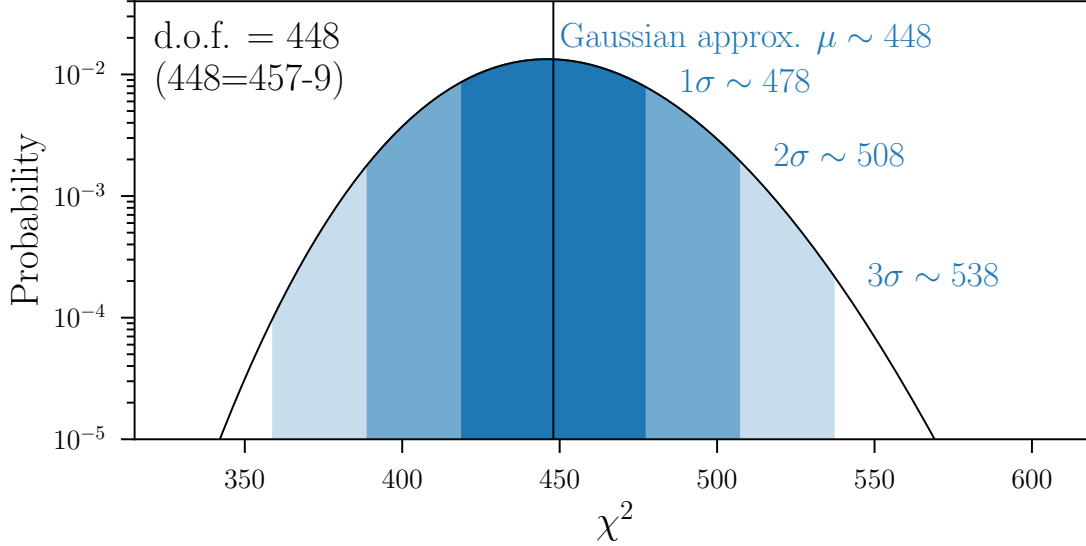


Figure 5.4: Probability distribution for χ^2 values expected for realizations of our data vector, which have 448 degrees of freedom (457 data points – 9 varied parameters). The vertical black line and shaded regions show the mean, and 1, 2, and 3σ regions for the χ^2 distribution, computed for the limit where the χ^2 distribution with many degrees of freedom $k \gg 1$, approaches a Gaussian with mean $\mu = k$ and a variance $\sigma^2 = 2k$.

5.3 Validation of two-point function blinding for DES: Methods

Our primary goal here is to determine whether we trust this method, for some choice of settings, enough to use it to blind the data vectors for the DES Y3 3×2 pt analysis. We will use $\Delta\chi^2$ as the main indicator of that. One of the key validation tests for the 3×2 pt pipeline will be to check the value of χ^2 before unblinding. Given the probability distribution of χ^2 values expected for the experiment — roughly, the number of data points minus the number of free parameters (ignoring the effect of priors on the nuisance parameters) — we can assess the probability p to exceed the measured χ^2 value, assuming our model accurately accounts for the physics and systematics of our experiment. If p is larger than some predetermined threshold, say 0.01, the pipeline passes the validation test. If not, more testing of the internal consistency of the data and analysis will be required before unblinding. Thus, our metric for assessing the performance of our blinding method will be the size of the shifts in p that it generates.

Fig. 5.4 shows the probability distribution for χ^2 values expected for the studies presented in this thesis: we have 457 data vector entries and fit with 9 varied parameters, so there

are $457 - 9 = 448$ degrees of freedom. Since a χ^2 distribution with k degrees of freedom for $k \gg 1$ can be approximated as a Gaussian with mean $\mu = k$ and variance $\sigma^2 = 2k$, we can easily compute that we expect our χ^2 distribution to be centered on $\chi^2 = 448$ and to have a standard deviation of $\sigma = 30$. Thus, we will use the question of whether two-point function blinding, for some choice of settings, can sure shifts in χ^2 that are smaller than 1σ , so with $\Delta\chi^2 < 30$.

In order to test this, we generate an ensemble of 100 simulated DES-Y3 like data vectors and blind them according to the method presented in Section 5.1.3. By performing parameter estimation on the simulated data vectors, we can examine the relationship between the best fit parameters to the blinded data Θ_{bl} ; the best fit parameter for the original, unblinded data Θ_{unbl} ; the shift $\Delta\Theta$ used to generate the blinding factors; the distance in parameter space between the true parameters Θ_{obs} used to generate the observed data; and the goodness-of-fit metric $\Delta\chi^2$ given in Eq. (5.12).

In doing these tests, we will compare the performance of additive and multiplicative blinding, and will do so for both noiseless data and for the case where a noise realization has been added to the simulated data vector. This section will describe the tools and considerations used to perform these studies, and then Section 5.4 will present the results.

5.3.1 Cosmology selection

Let us recall from Section 5.1.3 that our blinding procedure works by modifying the data vector \mathbf{d} using a blinding factor f which is built out of the theoretical predictions for the data vector $\hat{\mathbf{d}}(\Theta)$ evaluated at a reference set of parameters Θ_{ref} and a shifted set $\Theta_{\text{shift}} = \Theta_{\text{ref}} + \Delta\Theta$. For additive blinding the blinding factor consists of the difference $f^{(\text{add})} = \hat{\mathbf{d}}(\Theta_{\text{shift}}) - \hat{\mathbf{d}}(\Theta_{\text{ref}})$, while for multiplicative blinding it is the ratio $f^{(\text{mult})} = \hat{\mathbf{d}}(\Theta_{\text{shift}})/\hat{\mathbf{d}}(\Theta_{\text{ref}})$. If we would like test the effect of this transformation for simulated data, for each realization there are three sets of cosmological parameters we must select. We need to choose a set of parameters, which we will call Θ_{obs} , which are used to generate the simulated, “observed” unblinded data. We additionally need to select a reference cosmology Θ_{ref} , and by selecting the blinding shift $\Delta\Theta$, the shifted parameter set Θ_{shift} .

We would like to study the performance of our two-point function blinding method for a variety of choices for the shift $\Delta\Theta$, as well as for a range of offsets between the Θ_{ref} and the “true” cosmologies. Given this, our approach will be to fix Θ_{ref} at fiducial parameter values, and to draw Θ_{shift} and Θ_{obs} for each realization from distributions in parameter space.

We will focus on blinding the two cosmological parameters which are most at risk for

experimenters bias: the amplitude of matter clustering σ_8 , which will be checked against Planck constraints, and the dark energy equation of state parameter w , which has the special value of $w = -1$ in Λ CDM. Therefore, the shift $\Delta\Theta$ will be zero for all parameters other than σ_8 and w . We will refer to this set of two parameters as the *blinded parameters*.

For Θ_{obs} , and for our parameter estimation fits, we will limit the parameters that vary to σ_8 , w , Ω_m , h , and the five galaxy bias parameters b_i for $i = 1-5$. This choice is motivated by both by the fact that we will be trying to find the best fit for a large number of simulated data vectors, and must reduce the dimensionality of our parameter space in order to make fits computationally feasible with the Maxlike sampler. (We will perform 600 parameter fits in total: for each of the 100 realizations, we have both noisy and noiseless versions of data vectors, and for each simulated data vector we will have an unblinded, additively blinded, and multiplicatively blinded version.)

The probability distributions used to select Θ_{shift} and Θ_{obs} are defined in terms of the expected constraining power of the DES Y3 analysis, which we characterize through a Fisher analysis. Using the CosmoSIS Fisher⁴ sampler, we compute the Fisher information matrix with entries

$$F_{ij} = \sum_{ab} \frac{\partial d_a}{\partial \Theta_i} [\text{Cov}^{-1}]_{ab} \frac{\partial d_b}{\partial \Theta_j}, \quad (5.23)$$

where the data vector and covariance for the Y3 3×2 pt analysis are computed according to the description in Section 5.2. The inverse of the Fisher matrix $F^{-1} \equiv C_\Theta$ is a Gaussian approximation of the parameter covariance matrix. When we compute the Fisher matrix for Y3, we vary σ_8 , w , Ω_m , h , b_1 , b_2 , b_3 , b_4 , and b_5 and fix all other parameters. This choice to fix so many parameters means this procedure will somewhat underestimate the Y3 statistical errors.

We draw realizations of Θ_{obs} from a multivariate distribution in parameter space, centered on Θ_{ref} and with the Fisher covariance C_Θ . This distribution is truncated by the prior range: if any draw returns Θ_{obs} outside the prior ranges used in the Y1 3×2 pt analysis, we discard that parameter set and draw another. The fiducial parameter values and their priors are shown in Table 5.2.

We also select Θ_{shift} from a truncated multivariate Gaussian distribution centered on Θ_{ref} . For the distribution's covariance, we multiply the Fisher covariance by a scaling factor $\alpha > 1$ to get $C'_\Theta = \alpha C_\Theta$, and then generate shifts using just the rows and columns of C'_Θ corresponding to our blinded parameters, σ_8 and w . This effectively marginalizes over all other parameters. We truncate the Θ_{shift} distribution according to a more conservative

⁴bitbucket.org/joezuntz/cosmosis/wiki/samplers/fisher

prior region than Θ_{obs} , chosen to try and ensure that the blinded best fit parameters Θ_{bl} will remain inside the priors region for our parameter estimation searches. We choose the bounds for allowed Θ_{shift} somewhat arbitrarily to require that $w \in [-1.5, 0.5]$ and $\sigma_8 \in [0.714, 0.954]$ (this is the fiducial $\sigma_8 \pm$ three times the estimated Y3 marginalized error).

In order to allow the blinding procedure to produce shifts in parameter space that are large compared to expected constraints, we select α by requiring that Θ_{ref} will be $\geq N\sigma$ away from Θ_{shift} , as defined by the Gaussian distribution with covariance C_{Θ} , about half of the time. Constant $\chi^2 = \chi_0^2$ contours for a Gaussian distribution with covariance C'_{Θ} will be equivalent to $\chi^2 = \alpha\chi_0^2$ contours for a distribution with covariance C_{Θ} . Thus, using the χ^2 probability distribution for two free parameters (w and σ_8), we set

$$\alpha = \frac{\chi_2^2[4\sigma]}{\chi_2^2[50\%]} = \frac{19.33}{1.39} = 13.91. \quad (5.24)$$

Here we use the notation $\chi_2^2[4\sigma]$ to represent the value for which, for two-degrees of freedom corresponding to a two-dimensional parameter plane, the χ^2 probability to exceed is equivalent to that of a point 4σ away from the mean of a Gaussian distribution. Similarly, we define $\chi_2^2[50\%]$ to be the χ^2 value which for two free parameters has the probability to exceed of 0.5. This covariance scaling is equivalent to scaling the axes of confidence ellipses by a factor of $\sqrt{\alpha} = 3.72$. This scaling combined with the truncation described above produces what is effectively a flat distribution within Θ_{shift} 's allowed range in the $w - \sigma_8$ plane.

For each of 100 realizations, we draw parameter sets Θ_{obs} and Θ_{shift} according to these probability distributions. The resulting distribution of points in parameter space is shown for a selection of the varied parameters in Fig. 5.5. In that Figure, dashed green lines show the location of Θ_{ref} , red points show the location of our Θ_{shift} draws, and blue points show the Θ_{obs} draws. The gray ellipses show the 1σ marginalized Y3 3×2 pt constraints according to our Fisher analysis, and the dotted gray lines show the bounds used to truncate our Θ_{shift} distribution.

5.3.2 Data vector generation

For each realization, we use the pipeline described in Section 5.2 to compute data vector predictions $\hat{\mathbf{d}}(\Theta_{\text{obs}})$, $\hat{\mathbf{d}}(\Theta_{\text{ref}})$, and $\hat{\mathbf{d}}(\Theta_{\text{shift}})$. We combine $\hat{\mathbf{d}}(\Theta_{\text{ref}})$ and $\hat{\mathbf{d}}(\Theta_{\text{shift}})$ to get blinding factors $f^{(\text{add})}$ and $f^{(\text{mult})}$ according to Eqs. (5.9) and (5.11). To generate a noisy ‘‘observed’’ data vector \mathbf{d} , we use the Y3 3×2 pt data covariance to generate a realization of Gaussian noise and add it to $\hat{\mathbf{d}}(\Theta_{\text{obs}})$. The prediction $\hat{\mathbf{d}}(\Theta_{\text{obs}})$ becomes the noiseless ‘‘observed’’ data

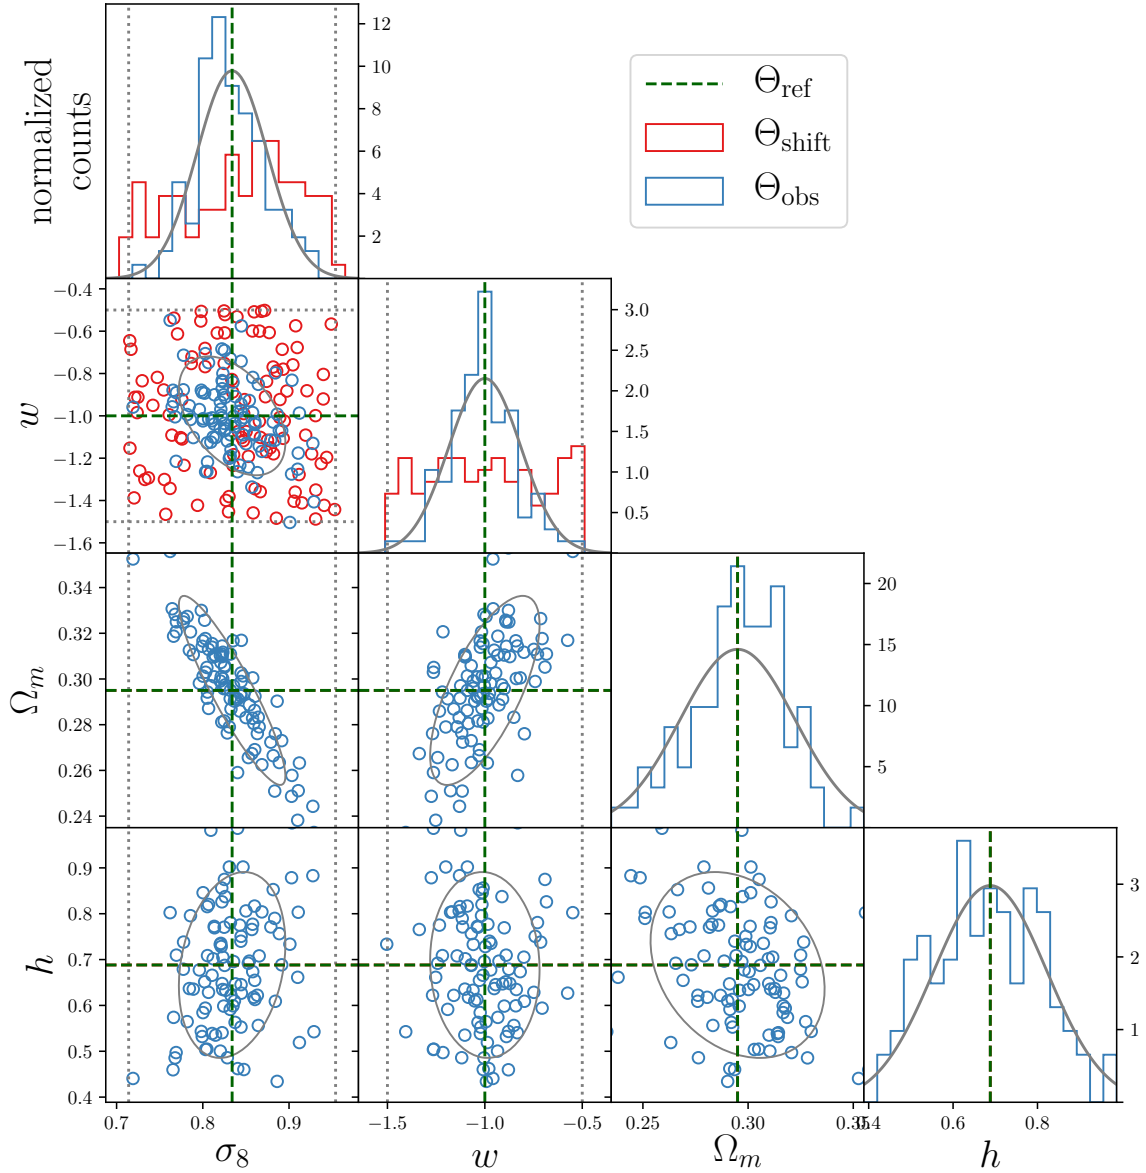


Figure 5.5: Locations in parameter space of the realizations used for the blinding scheme validation tests. The red circles show Θ_{shift} , the blue circles show Θ_{obs} , the dashed green lines show the location of the reference cosmology, and the dotted gray lines show the bounds on the allowed cosmology shifts used to generate blinding factors. The solid grey contours show the expected constraining power of the DES Year 3 analysis.

vector.

As an example, we show data vectors for our first simulated realization in Figs. 5.6 and 5.7. In those plots, the dotted black lines show $\hat{\mathbf{d}}(\Theta_{\text{obs}})$, and the solid black lines show the noisy data vector \mathbf{d} . The blinded versions of the noisy data vector are shown with colored lines, with solid green showing the result of additive blinding and dashed blue showing the result of multiplicative blinding. The blinding shift $\Delta\Theta$ used in this realization is $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$. Figs. 5.8 and 5.9 show the difference between the blinded and unblinded data vectors in order to more clearly display the effects of blinding factors in more detail. For additive blinding, the difference will exactly be the blinding factor $f^{(\text{add})}$. We can see that multiplicative blinding results in similar changes to the data vector, except for where the data vector is close to zero.

5.3.3 Parameter estimation with the Maxlike sampler

We would like to perform parameter estimation on our simulated data vectors in order to assess the χ^2 goodness of fit at the location in parameter space with the maximum posterior. Because we care only about the maximum likelihood point (with flat priors, this is equivalent to the maximum posterior point) and not about obtaining accurate constraint contours, and because we need to be able to perform this search for a large number (600) of data vectors, we will use a numerical optimization routine, rather than a Markov chain Monte Carlo (MCMC) or similar chain-based method. Specifically, we use CosmoSIS’ MaxLike sampler⁵, which is a wrapper for the `scipy.optimize.minimize` function⁶. This sampler uses the Nelder-Mead Simplex algorithm [210] to find the point in parameter space which maximizes the likelihood. The parameter bounds that we use for this search are listed in Table 5.2.

We should treat the outputs of the Maxlike sampler with caution, as it is known to fail in high-dimensional spaces, and to be sensitive to the starting point chosen to begin its search in parameter space. This issue is the main reason why we perform fits in a 9-dimensional subset of the full 26-dimensional DES parameter space. We will try to help the sampler by using the fact that we know the “true” or expected parameter values to provide the sampler with good guesses as the starting points for our searches. When fitting unblinded data, we start the sampler at Θ_{obs} , the parameter set used to generate the data vector. When fitting data which has been blinded using the parameter shift $\Delta\Theta = \Theta_{\text{shift}} - \Theta_{\text{ref}}$, we start the sampler at $\Theta_{\text{obs}} + \Delta\Theta$. Using these settings, a typical run of the Maxlike sampler on noisy data takes about half an hour on a single core.

⁵bitbucket.org/joezuntz/cosmosis/wiki/samplers/maxlike

⁶docs.scipy.org/doc/scipy-0.13.0/reference/generated/scipy.optimize.minimize.html

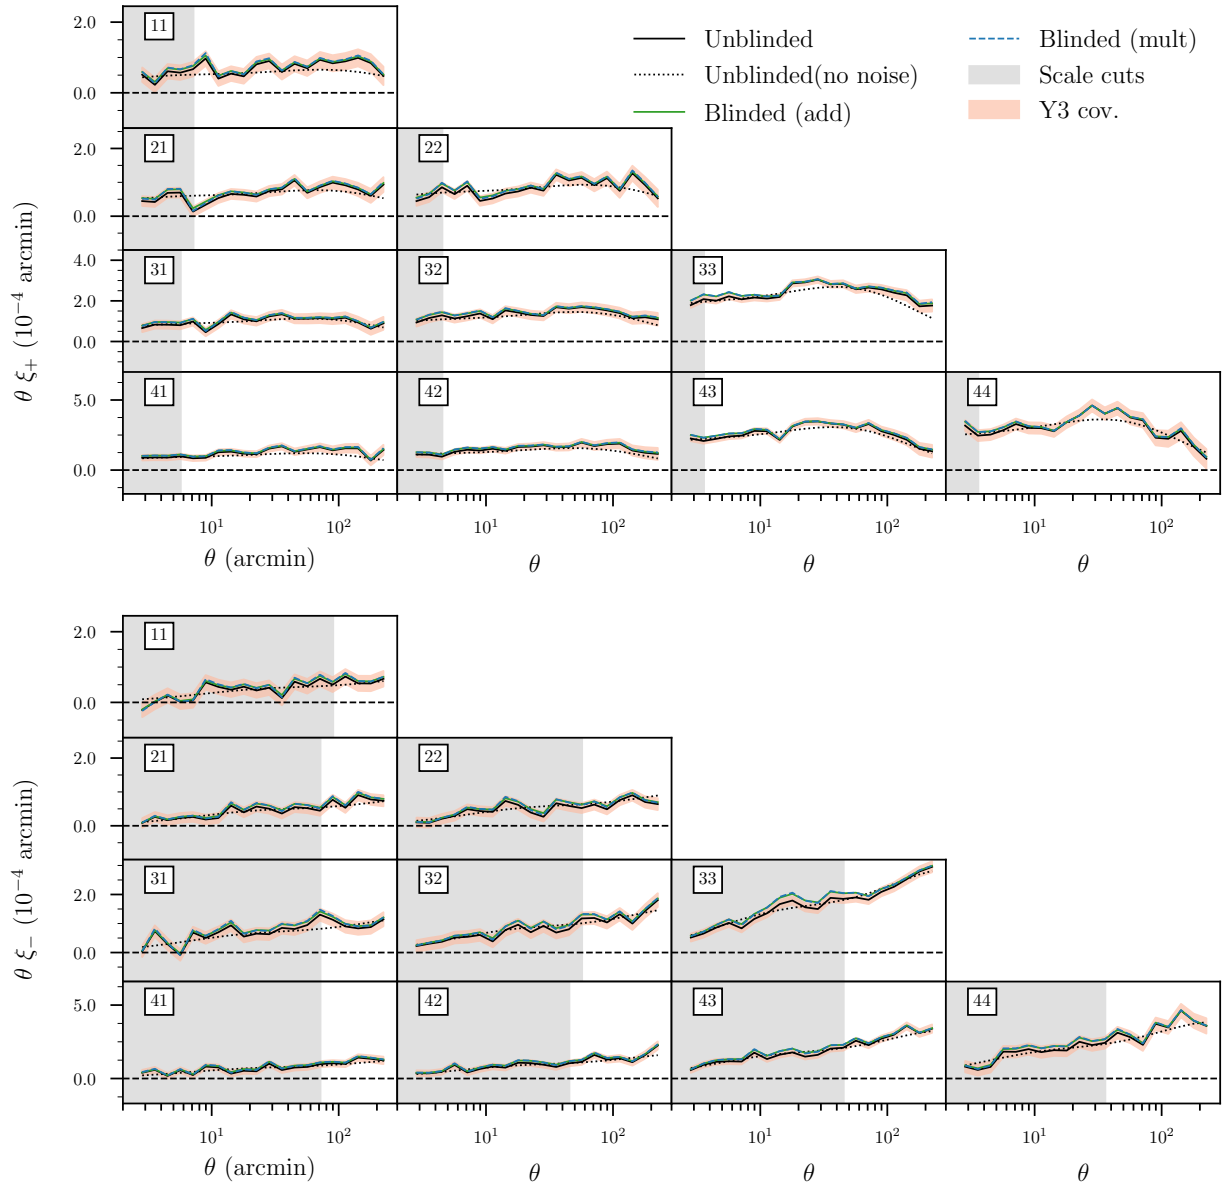


Figure 5.6: Example of effect of blinding on a subset of the simulated DES 3×2 pt data vector for realization 1, which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$. The rest of the data vector is shown in Fig. 5.7. This plot shows the impact of additive (green) and multiplicative (blue) blinding on $\xi_{\pm}(\theta)$ for a simulated data vector (black). The dashed black line shows the simulated data vector without noise added. The orange band around the black line shows the statistical errors projected for DES Year 3 (the square root of the covariance’s diagonal entries), and the vertical gray bands show the scales that are excluded by scale cuts. Each subplot shows the cross correlations between a different pair of redshift bins labeled by numbers in white boxes. The bin label numbers go up in order of increasing redshift.

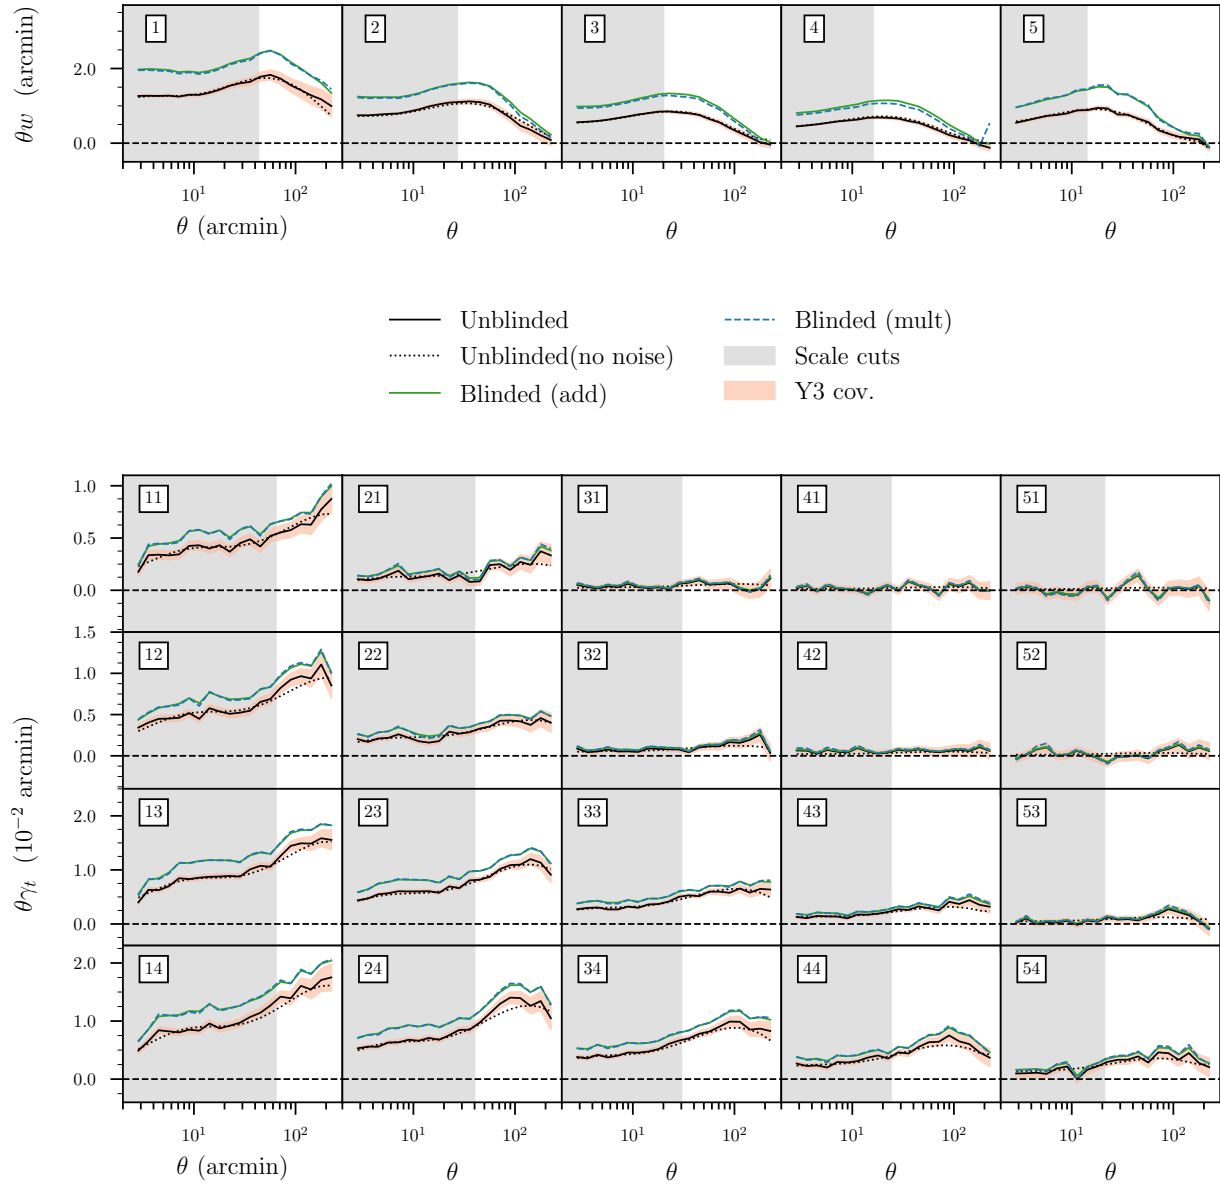


Figure 5.7: Example of effect of blinding on a subset of the simulated DES 3×2 pt data vector for realization 1, which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$. This plot displays the $w(\theta)$ and $\gamma_t(\theta)$ components of the same data vector shown in Fig. 5.6.

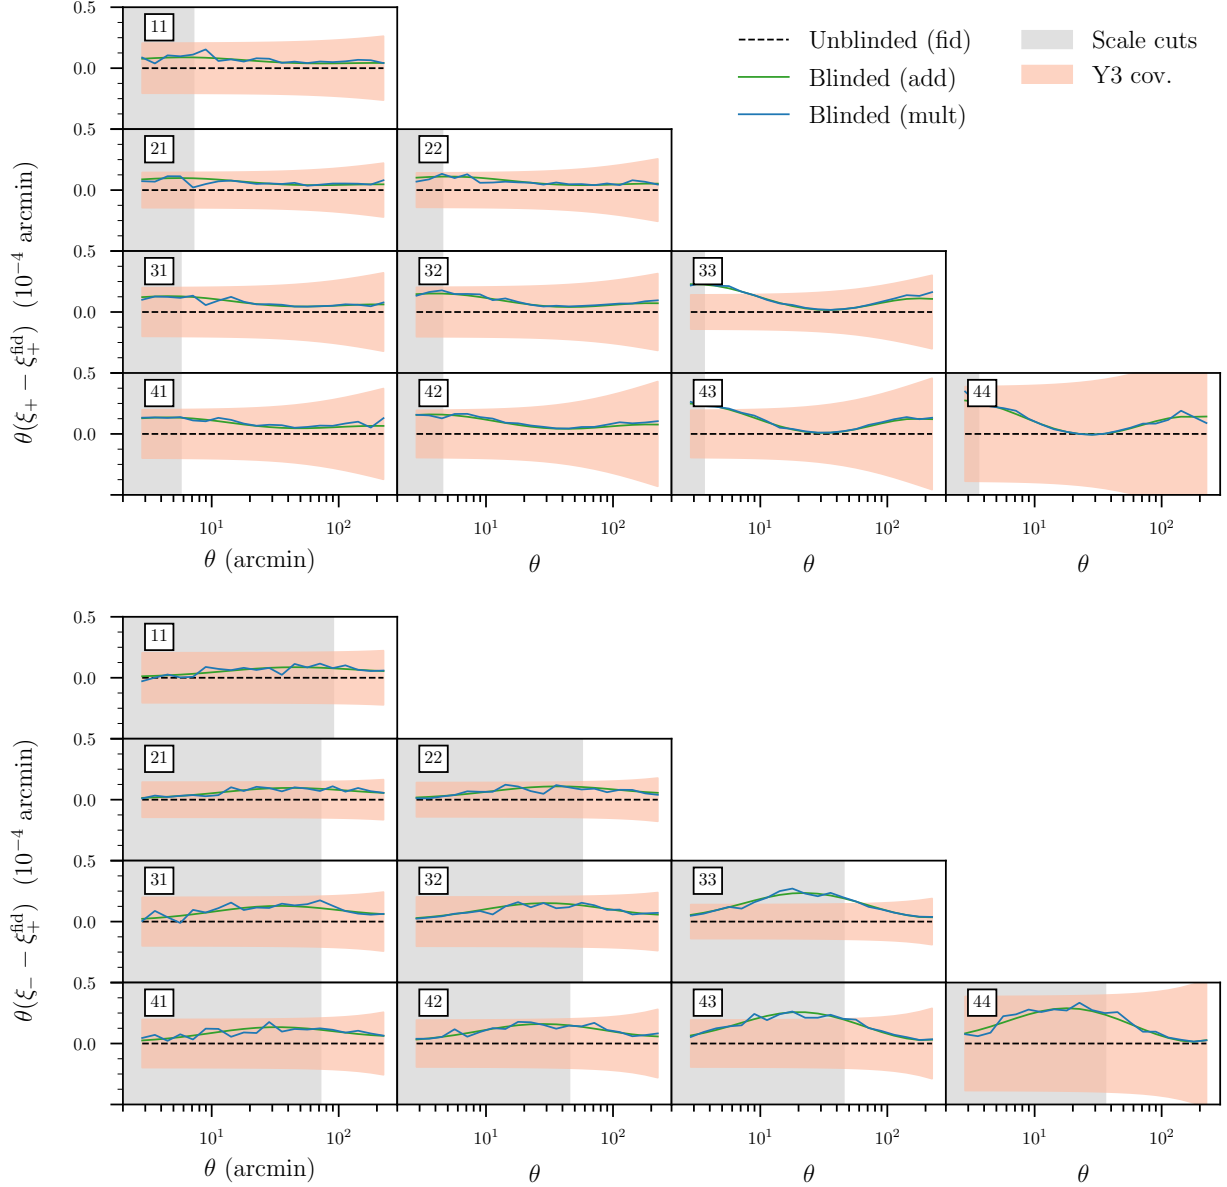


Figure 5.8: Differences between the blinded and unblinded data vector components shown in Fig. 5.6, for realization 1 which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$.

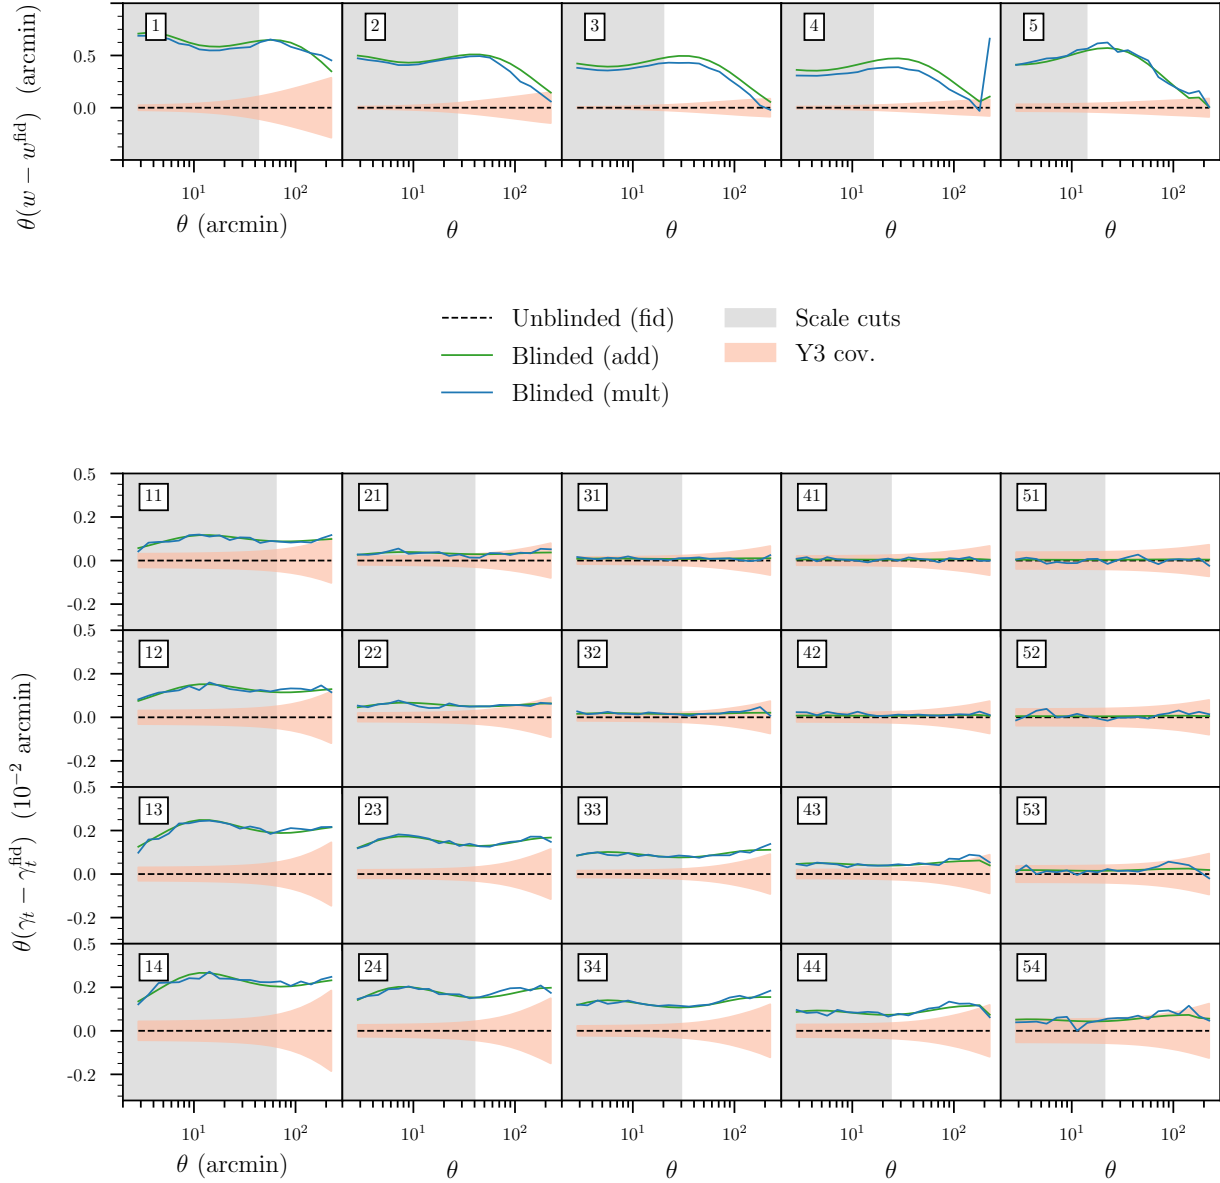


Figure 5.9: Differences between the blinded and unblinded data vector components shown in Fig. 5.7, for realization 1 which had a blinding shift of $\delta\sigma_8 = +0.074$ and $\delta_w = +0.43$.

5.3.4 Parameter estimation with the Multinest sampler

Though it is not feasible to do for all 600 data vector fits, for a few selected realizations we will also perform parameter estimation using the Multinest sampler⁷ [211–213]. Multinest is an MCMC-like nested sampling routine which uses a set of “live points” that step strategically through parameter space to produce a chain of posterior evaluations. That chain can be used to estimate the shape of constant-posterior surfaces and evaluate Bayesian evidences. Its accuracy can be specified by two input parameters: the tolerance on the computed Bayesian evidence, which is used as a stopping criterion for the chain, and the efficiency, the ratio between the number of points accepted and used for posterior estimation to those sampled. We use the Multinest implementation in CosmoSIS⁸, which is the primary sampler used to obtain parameter constraints for the DES Y1 cosmology analyses.

Since we are mainly concerned with finding χ^2 at the maximum likelihood set of parameters, we run Multinest with low-resolution settings, using 250 live points, an efficiency of 0.8, and a tolerance of 0.1. These settings mean that chains will end after about 140,000 posterior evaluations, the mean and parameter covariances from these runs will be accurate to a few percent [214], and constraint contours plotted based on the chains will have rough edges. These runs, using the same varied parameters and priors as our Maxlike runs take about five hours when run in parallel on 144 cores.

5.4 Validation of two-point function blinding for DES: Results

Our discussion of the results for blinding scheme tests will have three parts. First, in Section 5.4.1 we will examine the distribution of χ^2 and $\Delta\chi^2$ values for the ensemble of fits we have done using the Maxlike sampler. The rest of the studies will be in service of developing a deeper understanding of those distributions. In Section 5.4.2 we will study how $\Delta\chi^2$ depends on the shift $\Delta\Theta$ used to generate blinding factors, and in Section 5.4.3 we will study how $\Delta\Theta$ correlates with the resulting shifts in best fit parameters. Finally, in Section 5.4.4 we will compare the Maxlike parameter estimation results to those from Multinest for a few selected realizations in order to assess whether some high $\Delta\chi^2$ values may be due to the failure of the Maxlike sampler.

⁷ccpforge.cse.rl.ac.uk/gf/project/multinest/

⁸bitbucket.org/joezuntz/cosmosis/wiki/samplers/multinest

5.4.1 Distribution of χ^2 for fits

The upper panel of Fig. 5.10 shows histograms of the χ^2 values measured from our ensemble of 100 noisy unblinded data vectors as a black histogram. Behind it, the χ^2 histograms from fits to the blinded version of those data vectors, with the additive blinding case shown in green and multiplicative blinding (which as a fiducial choice include the covariance scaling of Eq. (5.20)) shown in blue. The lower panel shows histograms in corresponding colors for $\Delta\chi^2$ measured for each realizations.

The χ^2 results look initially promising. The additive and multiplicative blinding histograms are fairly similar to that for unblinded data, if perhaps slightly more skewed towards high χ^2 values. The additive and multiplicative blinding histograms are also fairly similar to one another, though we do note that one realization has an anomalously high χ^2 value for the multiplicative fit. Studying the histograms in bottom panel of Fig. 5.10, we can compare the $\Delta\chi^2$ distributions to our desired cap of $\Delta\chi^2 = 30$, which identifies shifts of 1σ in the χ^2 probability distribution for our data vector and is marked with a gray dashed line. We find that neither additive nor multiplicative blinding quite satisfies the requirement that $\Delta\chi^2 < 30$ for all realizations. However, the vast majority of realizations (99 for additive blinding, and 95 for multiplicative blinding) do satisfy that requirement. All additive blinding realizations represent χ^2 shifts of less than 2σ ($\Delta\chi^2 < 60$), as do all but two multiplicative blinding realizations.

Fig. 5.11 shows the same information for noiseless data vectors. Recall from our discussion in Section 5.1.4 that the $\Delta\chi^2$ for noiseless data quantifies the residuals between the best fit model and the *signal* part of the blinded data vector. It is therefore these $\Delta\chi^2$ distributions for noiseless data that quantify how well two-point function which are scaled according to our blinding method mimic valid data vectors within our cosmological mode. Again here we see that most but not all realizations have $\Delta\chi^2 < 30$: now 98 out of 100 realizations for additive blinding and 95 realizations for multiplicative blinding. We note that in general the $\Delta\chi^2$ histograms for additive blinding and the fiducial multiplicative blinding procedure look similar to the case with noisy data vectors. Because there is no noise to be scaled, for these fits the multiplicative blinding returns fairly similar results whether or not the data covariance is scaled, but with less severe outliers.

We note that the fact that the black histogram includes non-zero values is a reminder that we cannot entirely trust the results from the Maxlike sampler, since the noiseless unblinded data vectors with no noise exactly match theory predictions and so should all have $\chi^2 = 0$. However, distributions of χ^2 obtained with Maxlike can still be informative even if the sampler is failing to find the true maximum likelihood (minimum χ^2) point in parameter

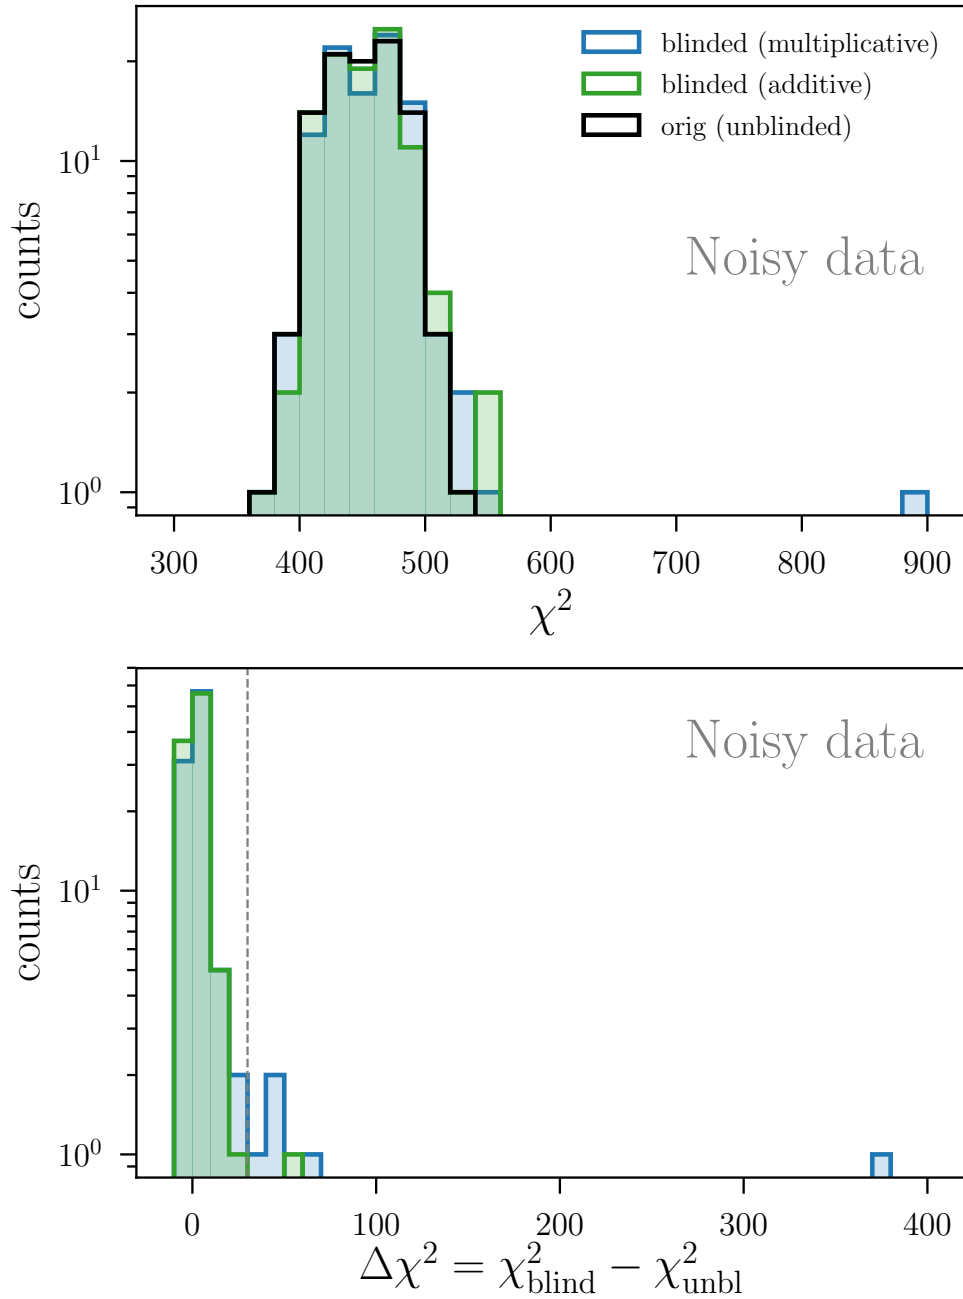


Figure 5.10: Upper panel: Histograms of χ^2 values obtained from Maxlike fits to simulated noisy data vectors. The black histogram shows the distribution for the unblinded data, while the blue and green histograms are for additive and multiplicative blinding, respectively. Lower panel: histogram of $\Delta\chi^2$ values, using the same colors as the upper panel. The vertical dashed gray line at $\Delta\chi^2 = 30$ marks the change that will move the realization by 1σ in the probability distribution for χ^2 expected for the $3 \times 2\text{pt}$ data vector.

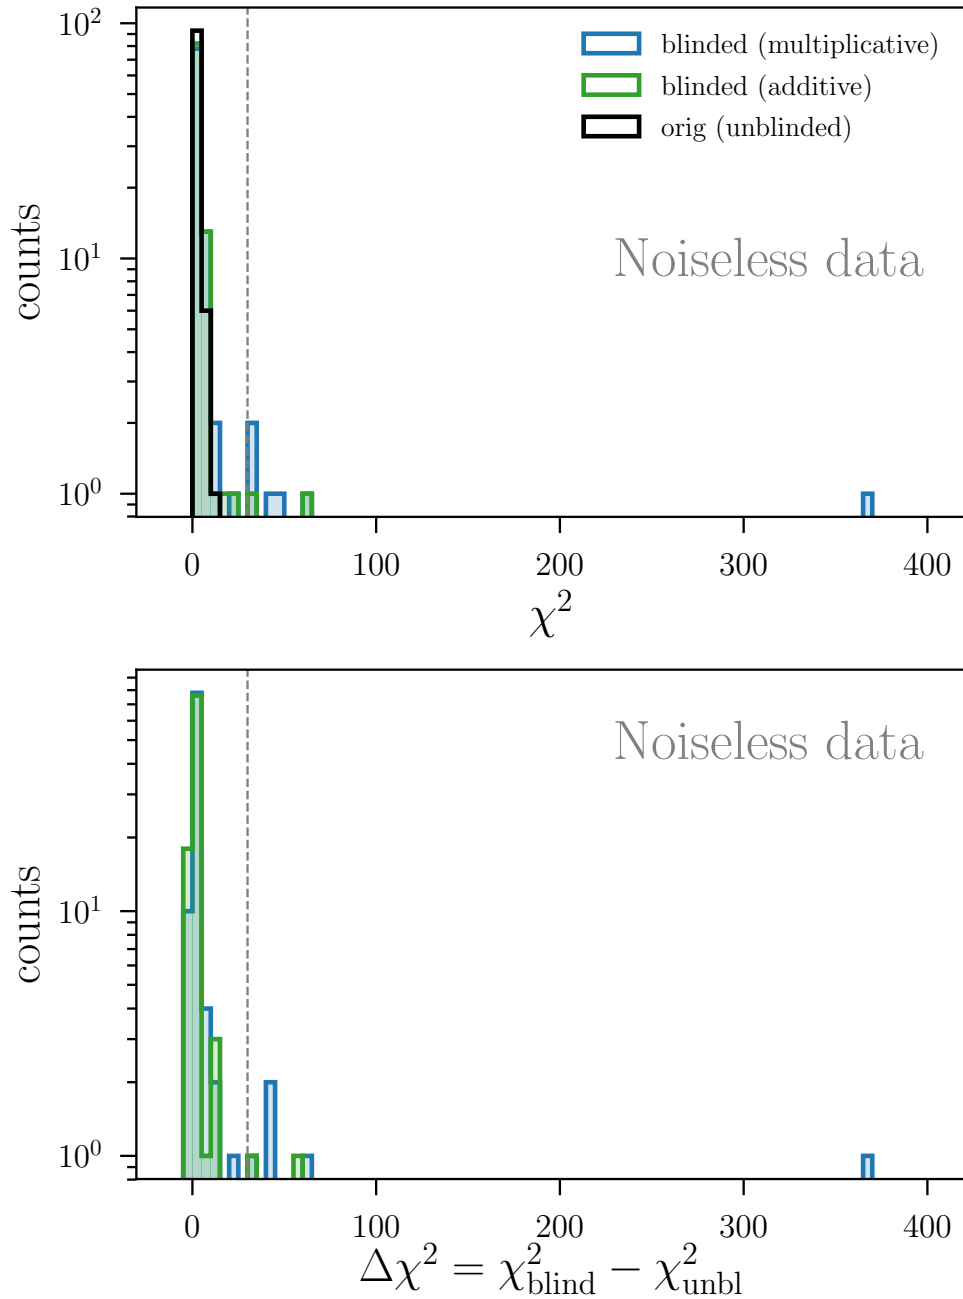


Figure 5.11: Upper panel: Histograms of χ^2 values obtained from Maxlike fits to simulated noisy data vectors. The black histogram shows the distribution for the unblinded data, while the blue and green histograms are for additive and multiplicative blinding, respectively. Lower panel: histogram of $\Delta\chi^2$ values, using the same colors as the upper panel. The vertical dashed gray line at $\Delta\chi^2 = 30$ marks the change that will move the realization by 1σ in the probability distribution for χ^2 expected for the $3 \times 2\text{pt}$ data vector. The wide axis range is chosen to show the outlier for multiplicative blinding at $\Delta\chi^2 = 369$. Fig. 5.12 reproduces this information with a decreased axis range.

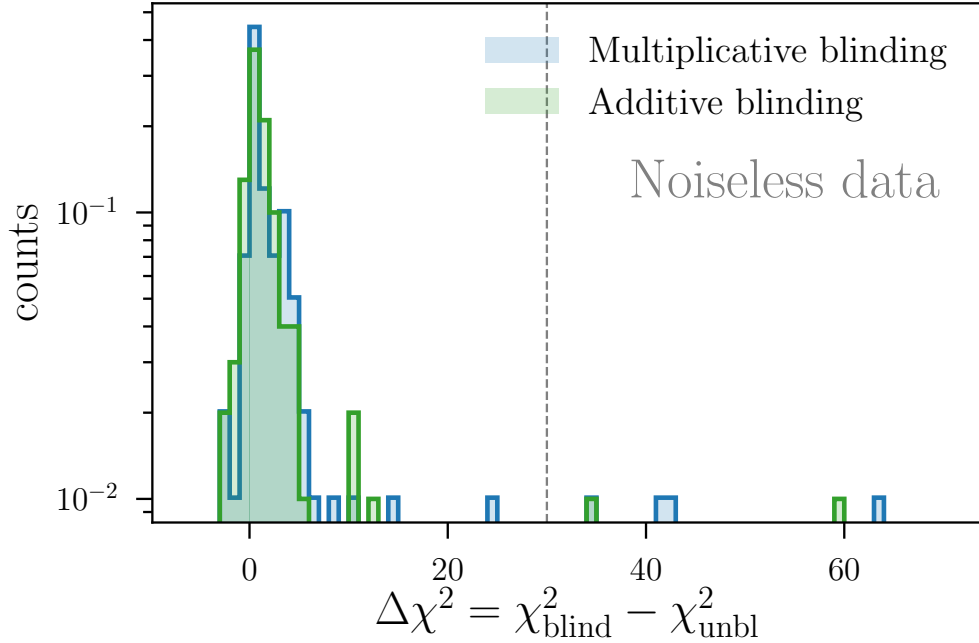


Figure 5.12: Same data as the bottom panel of Fig. 5.11, but not displaying the high $\Delta\chi^2$ outlier for multiplicative blinding, in order to show more detail along the horizontal axis.

space, since χ^2 can only ever be over-estimated relative to that of the global best-fit point.

Fig. 5.11 shows version of the same noiseless-data $\Delta\chi^2$ histogram for additive blinding alone, allowing us to reduce the range of the horizontal axis to examine its distribution in more detail. We see that there are only two realizations with $\Delta\chi^2 > 30$, one only slightly above 30 and the other at about $\Delta\chi^2 \approx 60$. The rest of the realizations have $\Delta\chi^2 \lesssim 15$. This is promising: if the couple of realizations with high $\Delta\chi^2$ can be explained, it appears additive blinding results in data vectors that look to good approximation like a valid data vector for some point in w CDM parameter space.

To begin investigating what is driving the observed high $\Delta\chi^2$ values, in Table 5.2 we note the realization numbers for each type of fit (with noise or noiseless, additive blinding or multiplicative blinding) which have $\Delta\chi^2 > 30$. We see the same realization numbers appearing for several data vector types. For example, the same four realizations have the highest $\Delta\chi^2$ values for multiplicative blinding, with or without noise. Additionally, realization 54 has the highest $\Delta\chi^2$ for both noise-added and noiseless additive blinding, as well as the third highest $\Delta\chi^2$ for both multiplicative blinding cases. We also note that the actual $\Delta\chi^2$ value for a given realization and blinding method is not significantly affected by whether or not noise is included in the simulation.

Data vector type		# above	Realization numbers with $\Delta\chi^2 > 30$ [$\Delta\chi^2$]
Noise	add bl.	1	54 [55.3]
	mult bl.	5	69 [374.0], 19 [60.7], 54 [46.5], 80 [44.2], 13 [32.6]
No Noise	add bl.	2	54 [59.3], 77 [34.8]
	mult bl.	5	69 [368.5], 19 [63.8], 54 [42.9], 80 [41.7], 77 [34.3]

Table 5.2: List of realizations which exceed $\Delta\chi^2 > 30$ for each set of Maxlike fits to blinded data. The actual $\Delta\chi^2$ values for each realization is shown in square brackets.

5.4.2 Relationship between $\Delta\chi^2$ and $\Delta\Theta$

The goal of these validation tests is to ascertain whether it will be appropriate to apply this blinding technique to the real DES Y3 data vector. There will be only one realization of the true data vector, and ideally we would only apply one blinding factor to it (to avoid stochastically sampling the allowed distribution of blinding shifts). If blinding causes the χ^2 of a fit to the real data to become large, it will impact the results of systematics tests of the analysis pipeline. This means that it is not sufficient to simply note that most of the simulated realizations we study in the blinding validation tests presented here have $\Delta\chi^2$ smaller than our threshold of 30. We also need to understand what is driving the high $\Delta\chi^2$ values of the outliers in our set of 100 simulated realizations, so that we can prevent the real data from experiencing a large blinding-induced χ^2 shift. Thus, this and the two following subsections will work towards understand what drives the high $\Delta\chi^2$ values for fits to certain realizations.

Here we test the hypothesis that these poor fits might be coming from certain input blinding parameter shifts $\Delta\Theta$. We do so in Fig. 5.13 by plotting $\Delta\chi^2$ against the input shifts in $\Delta\sigma_8$ and Δw . In those plots, we use green points to denote additive blinding and blue points to represent multiplicative blinding. For reference, we mark $\Delta\chi^2 = 30$ with a horizontal dashed gray line.

Studying Fig. 5.13, we see that there is no coherent pattern for how $\Delta\chi^2$ relates to blinding parameter shifts. Though we note that there seem to be more high values at large blinding shifts in w (in either direction), this is not consistently the case: there are some $\Delta\chi^2 > 30$ for smaller Δw , and there are also plenty of realizations with high Δw which have small $\Delta\chi^2$. We also checked, but do not show here, similar plots of $\Delta\chi^2$ vs. $\Theta_{\text{obs}} + \Delta\Theta$, positing that high $\Delta\chi^2$ values may be associated with certain combinations of the draw for Θ_{obs} and for Θ_{shift} . Again, no identifiable pattern emerged. Therefore, the the high $\Delta\chi^2$ values of the outlier points seen in the histograms of Section 5.4.1 are not caused (solely, at least) by blinding shifts to or from certain parts of parameter space.

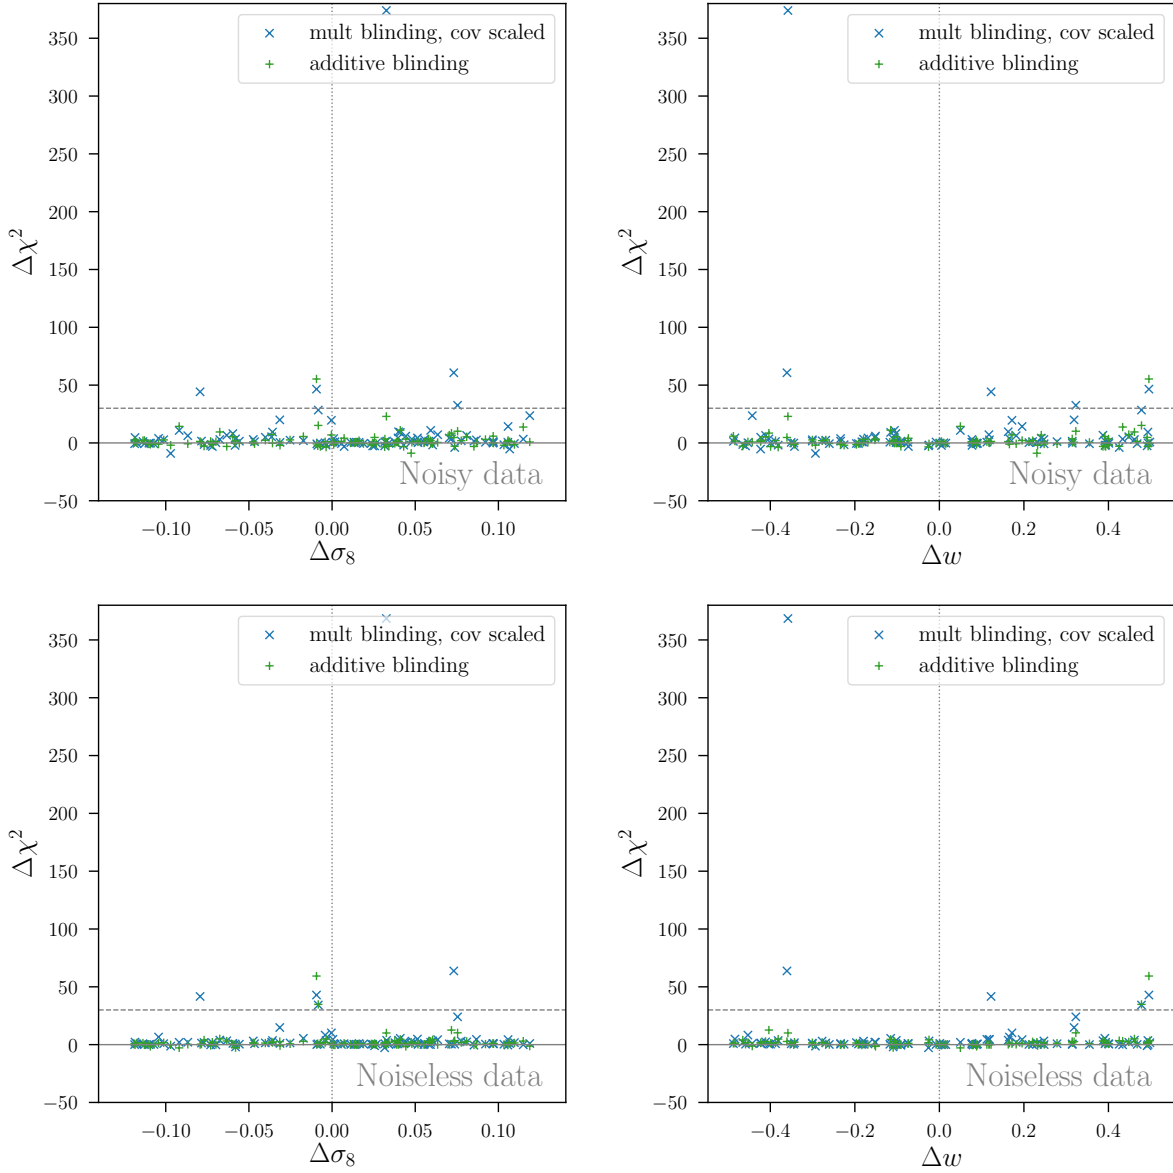


Figure 5.13: Scatter plots of $\Delta\chi^2$ versus the parameter shifts $\Delta\Theta$ used to produce blinding factors. The top row shows results for noisy data, while the bottom row shows the noiseless case.

5.4.3 Relationship between input and output parameter shifts

Next we look at how the shift in best-fit parameters due to blinding is related to the parameter shifts used to produce the blinding factors, as well as to the $\Delta\chi^2$ value of each realization. Fig. 5.14 shows this relationship through a series of scatter plots which have $\Theta_{\text{bl}} - \Theta_{\text{unbl}}$ on their vertical axes and the blinding shifts $\Delta\Theta$ along the horizontal axes. In showing parameters that have been fit, we show the best-fit shifts for the two parameters we are trying to blind, w and σ_8 , as well as Ω_m in order to examine how unblinded parameters are affected by the blinding procedure. In the plots, the colors of the points indicate their $\Delta\chi^2$ values and the gray numbers show the correlation coefficient between the two quantities shown in each subplot. Since the resulting scatter looks similar for noisy and noiseless data, we show results for the noiseless case only. These plots allow us to draw two main conclusions.

First, we note that the input blinding parameter shift is a good predictor of the resulting output shift in best fit parameters. We see that, as expected, the resulting shift in the best fit value for w and σ_8 is highly correlated with the shifts used to produce blinding factors. The w shifts have a correlation coefficient of $R = 0.97$ for additive blinding and 0.96 for multiplicative blinding. The σ_8 shifts have correlation coefficients of 0.98 for additive blinding and 0.97 for multiplicative blinding. Thus, the correlations are higher for additive blinding, but only slightly. There is not significant correlation between blinding shifts and the best fit values for other parameters. This is a desired outcome, as it indicates that a blinding transformation with parameter shift $\Delta\Theta$ corresponds reasonably well to how the data vector changes if the true cosmological parameters used are shifted by the same vector in parameter space.

We present these correlation results with a note of caution, however. Recall that we used $\Theta_{\text{obs}} + \Delta\Theta$ as our starting guess for the Maxlike sampler. It is possible that if the Maxlike sampler is failing for some fits, that choice could bias the correlations in w and σ_8 in Fig. 5.14 to be closer to $R = 1$ than is actually the case.

The second main result is that here is no obvious pattern for where the high $\Delta\chi^2$ values occur in these two-dimensional projections of our parameter space, other than the tendency noted in the previous section for high χ^2 be at high $|\Delta w|$. We observe similar correlations, and a similar lack of pattern for the location of high $\Delta\chi^2$ values when we study plots of Θ_{bl} versus $\Theta_{\text{obs}} + \Delta\Theta$ (which are not shown).

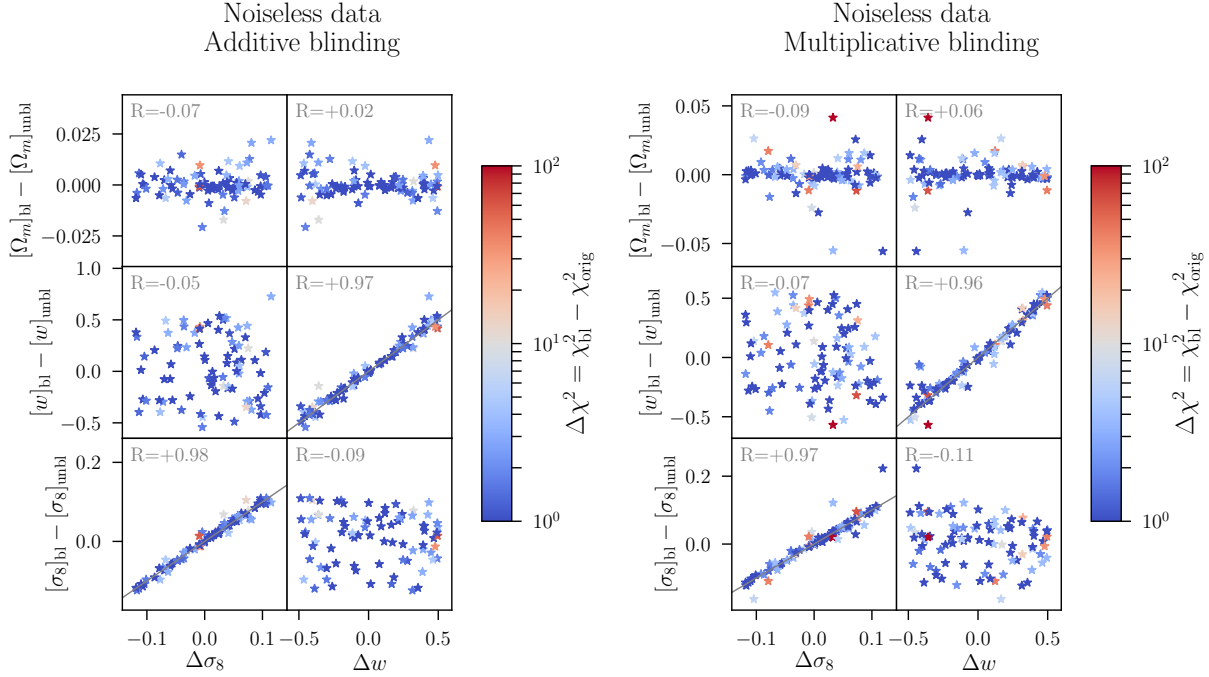


Figure 5.14: Relationships between shift in best fit cosmological parameters due to additive (left plot) and multiplicative (right plot) blinding for noiseless data and the parameter shifts used to generate blinding factors. The colors shows the shift in χ^2 between blinded and unblinded data, with the high and low-ends of the color scale truncated to make the plots easier to read. Solid gray lines show where $\Delta\Theta^{\text{out}} = \Delta\Theta$, and the grey numbers show the correlation coefficients R for the data points in each subplot.

Rlzn.	param. set	w	σ_8	Ω_m	h	b_1	b_2	b_3	b_4	b_5
	Θ_{ref}	-1.00	0.834	0.295	0.688	1.45	1.55	1.65	1.80	2.00
1	Θ_{obs}	-1.34	0.857	0.267	0.766	1.44	1.66	1.76	1.95	2.20
	$\Delta\Theta$	+0.43	+0.074							
	$\Theta_{\text{obs}} + \Delta\Theta$	-0.911	0.931							
54	Θ_{obs}	-0.99	0.912	0.263	0.519	1.24	1.41	1.50	1.63	1.74
	$\Delta\Theta$	+0.49	-0.009							
	$\Theta_{\text{obs}} + \Delta\Theta$	-0.49	0.903							
69	Θ_{obs}	-1.09	0.832	0.295	0.362	1.51	1.56	1.68	1.80	2.06
	$\Delta\Theta$	-0.36	+0.033							
	$\Theta_{\text{obs}} + \Delta\Theta$	-1.44	0.865							

Table 5.3: Cosmological parameters used to generate data vectors and blinding factors or selected realizations. Recall that Θ_{obs} is the set of “true” cosmological parameters used to produce the unblinded data vector, and $\Delta\Theta$ is the parameter shift associated with the blinding transformation. Roughly speaking, we expect the best-fit parameters for unblinded data to be $\Theta_{\text{unbl}} \approx \Theta_{\text{obs}}$ and for blinded data to be $\Theta_{\text{bl}} \approx \Theta_{\text{obs}} + \Delta\Theta$.

5.4.4 Multinest results for selected realizations

Now that we have verified that there is no causal relation between which realizations have high $\Delta\chi^2$ and where their parameter sets Θ_{obs} or Θ_{shift} occur in parameter space, we now explore the possibility that high $\Delta\chi^2$ values may be due to failures of our maximum likelihood finder, the Maxlike sampler. If this is the case, it means that the $\Delta\chi^2$ of the observed outliers may be over-estimated, and therefore that the performance of our blinding method is actually better than indicated by the histograms in Section 5.4.1.

We will test the performance of the Maxlike sampler by using the Multinest sampler to perform more robust (and more computationally expensive) fits to data vectors for a few selected realizations. Specifically, we will run Multinest on the six relevant data vectors (noisy and noiseless, for unblinded, additively blinded, and multiplicatively blinded) for realizations 1, 54, and 69. We select realizations 54 and 69 because they have the highest $\Delta\chi^2$ values for additive and multiplicative blinding, respectively, as we previously noted in Table 5.2. Realization 1 is included as a “control” comparison, as an example of a parameter draw which results in small $\Delta\chi^2$. (Its associated data vectors and blinding factors are visualized in Figs. 5.6, 5.7, 5.8, and 5.9.) For reference, we show the input cosmological parameter associated with these three realizations in Table 5.3.

In Table 5.4 we collect information that will allow us to compare χ^2 information for these three realizations. For all of the data vectors fit with Multinest, we show the measured χ^2 value for the unblinded, additively blinded, and multiplicatively blinded data vectors, as well as the $\Delta\chi^2$ values where relevant, comparing the values obtained from Maxlike and Multinest. As would be expected for a more robust sampler, the χ^2 values returned by Multinest are nearly always lower than those from Maxlike. Similarly, using Multinest instead of Maxlike causes nearly all of the $\Delta\chi^2$ values to decrease, sometimes fairly dramatically. Based on these results, it is reasonable to assume that most, if not all of our realizations with $\Delta\chi^2 > 30$ are suffering from failures of the Maxlike sampler.

This means that if we were able to locate the true maximum posterior point for all realizations, the results for the histograms we showed in Section 5.4.1 would improve. Notably, when we use Multinest on realization 54, which had the highest Maxlike $\Delta\chi^2$ value for additive blinding, its $\Delta\chi^2$ drops to below our threshold of 30 for both noisy and noiseless data vectors, and for both additive and multiplicative blinding. This suggests that all additively blinded data vectors are likely below that threshold. For the multiplicatively blinded data vectors, this is not necessarily the case. Note that the multiplicatively blinded data vector for realization 69 was the data vector that with Maxlike returned the anomalously high value of $\Delta\chi^2 = 368$, which appears as an outlier in Figs. 5.10 and 5.11. With Multinest that value

	Rlzn.	Sampler	unblinded	add. bl.		mult. bl.	
			χ^2	χ^2	$\Delta\chi^2$	χ^2	$\Delta\chi^2$
Noise	1	maxlike	455.5	452.6	-2.8	451.5	-3.9
		multinest	453.4	450.9	-2.5	451.9	-1.5
	54	maxlike	487.6	542.9	55.3	534.1	46.5
		multinest	479.7	499.5	19.7	504.2	24.4
	69	maxlike	518.2	541.1	23.0	892.2	374.0
		multinest	517.8	527.0	9.3	608.3	90.5
No Noise	1	maxlike	1.7	5.1	3.4	2.1	0.4
		multinest	1.3	1.9	0.6	1.9	0.5
	54	maxlike	4.3	63.6	59.3	47.2	42.9
		multinest	1.5	11.8	10.3	26.3	24.8
	69	maxlike	0.8	10.9	10.1	369.3	368.5
		multinest	0.82	9.0	8.1	80.9	80.0

Table 5.4: Comparison between Maxlike and Multinest fits showing χ^2 for simulated data vector realizations with the worst $\Delta\chi^2$ for additive blinding. The differences in the $\Delta\chi^2$ are relative to the unblinded fit in the same row.

lowers significantly, but remains high at $\Delta\chi^2 = 80$. We will see below that Multinest’s constraint contours for this data vector are hitting the prior bounds for w , something about the application of this particular multiplicative blinding shift too this particular data vectors is causing problems. Given this, we conclude that for the DES Y3 analysis, additive blinding is less likely to cause a spurious large χ^2 value, and is therefore a better choice for application of our blinding technique to real data.

Having run Multinest on these data vectors affords us with the possibility of seeing how blinding affects the constraint contours for a given data vector. Fig. 5.15 shows these constraints in a subset of our parameter space. The red contours show constraints for the noiseless, unblinded data vector, and the dashed black lines show the location of the parameters Θ_{obs} used to generate it. The green and blue contours show constraints after that noiseless data vector has been blinded additively and multiplicatively, respectively. Referencing the input parameter shifts listed in Table 5.3, we see that the contours are centered at slightly lower σ_8 and higher w compared to what we would expect if $\Theta_{\text{bl}} = \Theta_{\text{obs}} + \Delta\Theta$ (which would give $w = -0.91$ and $\sigma_8 = 0.931$). When we extract the maximum posterior points from the Multinest chains, we get $w = -0.70$ and $\sigma_8 = 0.899$ for the additively blinded data vector and $w = -0.78$ and $\sigma_8 = 0.926$ for the multiplicatively blinded data vector. This is consistent with the fact that the correlations we observed in Fig. 5.14 have some scatter. These offsets are not a problem for the blinding methods, since we do

not care what output cosmology the blinded data vector is centered at (as long as it is not too close to the prior boundaries), just that it looks like it could be a valid data realization for *some* set of parameters. In Appendix D we show similar contours for realizations 54 and 69, as well as plots of their data vectors.

5.5 Conclusions

In this chapter we have introduced and tested a new method for blinding a DES multi-probe cosmology analysis. Our method works by modifying the two-point correlation functions which make up the DES data vector used for parameter estimation. We began by discussing blinding in a general way, noting that it must fulfill two main requirements in order to be effective. First, blinding must destroy the experimenters' ability to know whether or not the output of their analysis is consistent with their expectations. Second, blinding must preserve the experimenters' ability to perform validation checks on their analysis pipeline. We quantify a general assessment of how well this second requirement is fulfilled using the statistic $\Delta\chi^2$, the difference between the χ^2 of the blinded data evaluated at its best-fit point in parameter space and that of the unblinded data at its best-fit point. We then describe the specifics of our proposed blinding scheme, where theoretical predictions for the data vector evaluated at two sets of parameters Θ_{ref} and $\Theta_{\text{shift}} = \Theta_{\text{ref}} + \Delta\Theta$ are used to modify the observed data. For additive blinding, we compute a blinding factor by taking the difference between the shifted and reference cosmology predictions, and add that to the observed data vector. For multiplicative blinding, we multiply each data vector entry by the ratio between the shifted and reference cosmology predictions, and additionally scale the data covariance to account for how this multiplication scales noise contributions. After introducing the DES $3 \times 2\text{pt}$ data vector, which consists of galaxy-galaxy, shear-shear, and galaxy-shear two-point correlation functions for a number of redshift and angular bins, we use simulated data vectors to test performance of both additive and multiplicative blinding on the DES $3 \times 2\text{pt}$ analysis pipeline.

We generate noiseless and noise-added versions of an ensemble of 100 simulated $3 \times 2\text{pt}$ data vectors, and used the Maxlike sampler to find the best fit parameter values for each of them. By doing the same for additively and multiplicatively blinded version of those data vectors, we were able to assess the distribution of $\Delta\chi^2$ values. We found that the vast majority of realizations had $\Delta\chi^2 < 30$, the threshold marking the shift in χ^2 that will produce a 1σ equivalent change in probability in the χ^2 -distribution expected for the DES data vector. We performed additional tests to explore whether the couple of realizations

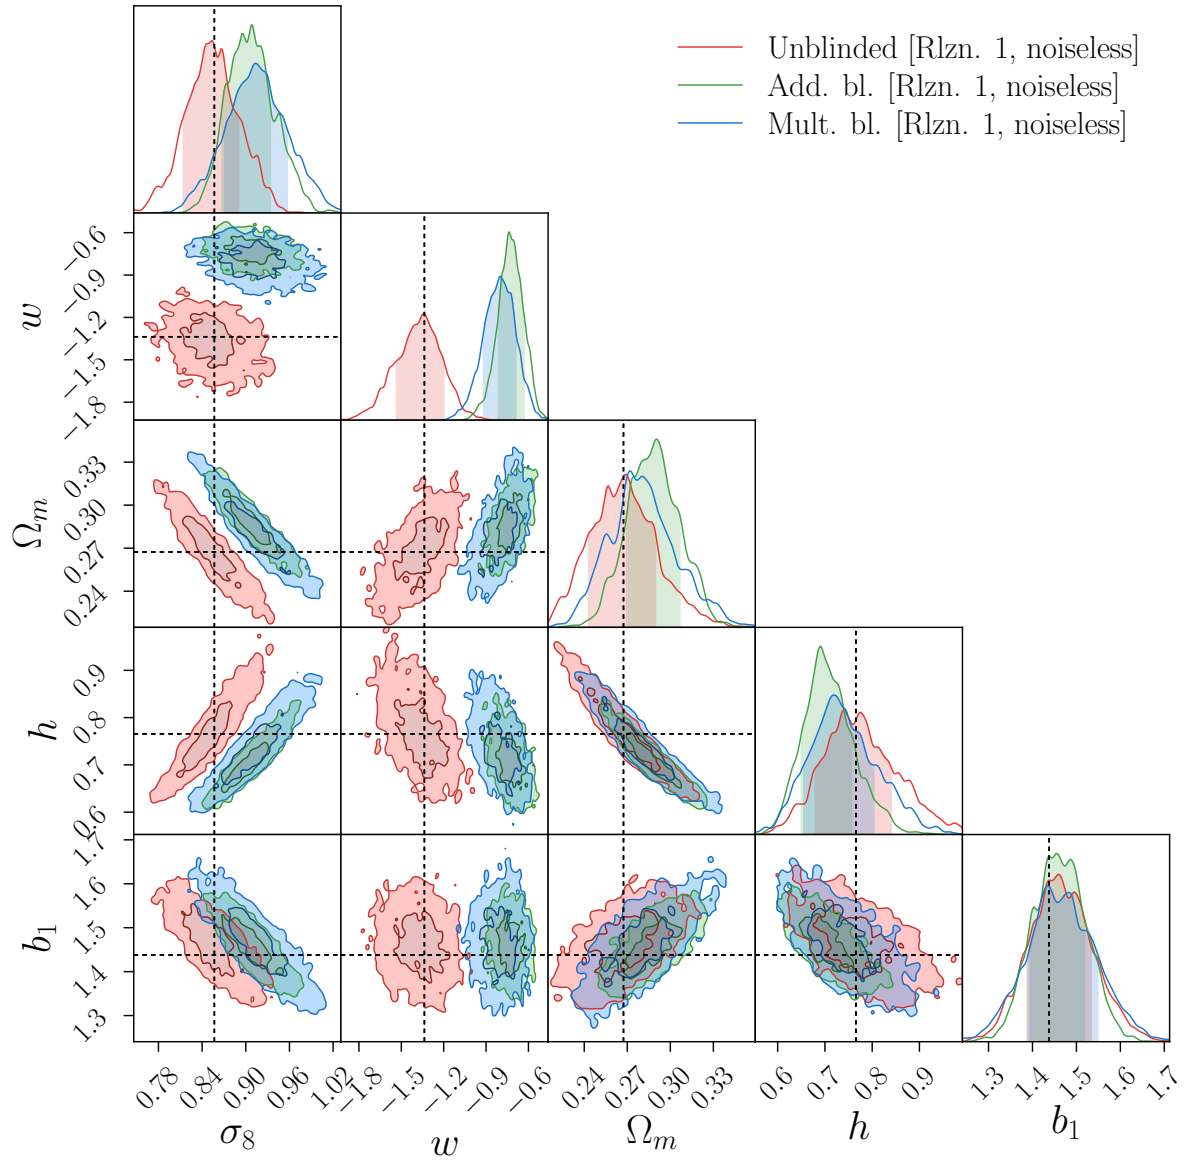


Figure 5.15: Parameter constraints for realization 1. Dashed black lines show the cosmology at which the unblinded data was simulated. The blinding factors used shifts $\Delta\sigma_8 = 0.074$ and $\Delta w = +0.43$.

with $\Delta\chi^2 > 30$ could be explained by where either their blinding shift $\Delta\Theta$ or true cosmology Θ_{obs} is in parameter space, and found no clear relation. Follow-up tests for a few selected realizations using a more robust but more computationally expensive parameter estimation method suggests that most of the high χ^2 values we observe are the result of our maximum likelihood finder failing to find the true best-fit point in parameter space. Based on these tests, we are reasonably confident that all additive blinding realizations actually have $\Delta\chi^2 < 30$. We additionally confirm that our blinding method can effectively shift the constraint contours by significant distances in parameter space, enough to blind the DES comparison to Planck, or the measurements of the dark energy equation of state parameter w 's comparison to the Λ CDM value of -1.

Based on the tests described in this chapter, additive blinding of σ_8 and w using this two-point-function-based method was deemed reliable enough to be used in the DES Y3 3×2 pt analysis. It is currently implemented as part of the data vector measurement pipeline, and has already been applied to preliminary measurements of a small subset of the Y3 shear measurements. This means that the data in the files of two-point function summary information which are input for Y3 parameter estimation will have been blinded in this way.

Even as this blinding method is being implemented in the DES Y3 3×2 pt analysis, there is a need for further study. In future tests, we will perform further checks on the specific version of this method that is implemented in the DES pipeline, specifically checking on whether our fiducial choices for nuisance parameters can affect its performance. In all of the tests presented above, the non-varied nuisance parameters are equal for both the ‘‘observed’’ data and for blinding factor generation. When this procedure is applied to real data, there is no guarantee that this will be the case, so it will be important to understand whether mismatches in systematics-related parameters affect the performance of the blinding method.

Testing the performance of this method on other kinds of summary statistics will also be a subject of future investigation. One of the benefits of how our method is implemented is that its blinding factors are entirely determined by the reference cosmology Θ_{ref} and the shift $\Delta\Theta$. Beyond that, it simply makes use of existing analysis infrastructure: if you are able to predict a data vector given a model (which you need to be able to do to perform parameter estimation), you can generate a blinding factor that should produce a shift in best fit parameters $\sim \Theta$. That is why the method is able to blind the different components of the 3×2 pt analysis in a consistent way and why, in principle it can be extended to the combined analysis of other probes. Particularly salient probes within other DES analyses are the so-called 5×2 pt data vector, which adds CMB lensing information to the 3×2 pt probes, and galaxy cluster mass functions.

The work presented in this chapter thus represents the development and initial validation of a robust and flexible multi-probe blinding technique which will serve as an important tool for combined analyses within DES and beyond. As cosmology as a field leverages the combined analysis of increasingly large datasets to tighten constraints on cosmological parameters, the need for blinding to protect against experimenter bias will only become more critical. Thus, the blinding strategy developed here has the potential to make significant contributions to future precision tests of Λ CDM.

Chapter 6

Closing remarks

Because Λ CDM as our standard cosmological model has proved to be in precise agreement with the vast majority of observations to date, even small or moderate tensions with its predictions naturally attract attention. Given this, for any observed anomalies we must make sure we understand our tools for assessing the potential sources of tension, as well as any uncertainties in our model predictions for associated observables. Additionally, when preparing for future analyses, we must take care to protect results from all sources of bias which could cause us to either make spurious detections or miss subtle hints of new physics. It is in this spirit that this dissertation presents three projects which contribute to the effort to characterize uncertainties and mitigate biases in precision tests of Λ CDM.

Chapter 3 assessed the accuracy with which we can use large scale structure survey information to disentangle late- and early-time contributions to CMB temperature anisotropies by reconstructing a map of the integrated Sachs-Wolfe signal. As this separation would allow us to distinguish which properties of the large-angle CMB are truly primordial, it could therefore inform the discussion of whether certain large-angle CMB anomalies are clues about the physics of inflation. Unfortunately, our analysis of simulated maps, modeling a survey with properties similar to the near-future Euclid experiment, showed that the accuracy of the reconstructed integrated Sachs-Wolfe signal is severely limited by our ability to separate physical density fluctuations of galaxy density on large scales from calibration errors. Our results suggest that for current levels of calibration control, reconstructed integrated Sachs-Wolfe maps are mostly noise, and so they can provide little or no information about whether specific large-angle CMB features are sourced by the integrated Sachs-Wolfe signal. Promisingly, we note that the level of calibration needed for future surveys like WFIRST to perform unbiased measurements of dark energy properties, will also allow for accurate integrated Sachs-Wolfe signal reconstruction. Thus, improving the separation of primordial and

late-universe contributions to CMB temperature anisotropies provides additional motivation for improving the calibration of future large scale structure surveys.

Chapter 4 delved further into the discussion of large-angle CMB temperature anomalies. Using measurements of a large number of Gaussian simulations, in the context of Λ CDM, we studied the covariance between eight large-angle features which are associated with commonly studied anomalies. By comparing measurements from our fiducial simulation ensemble to those based on the Planck Full Focal Plane simulations, we found that the relationship between the quadrupole-octopole alignment and other, power-spectrum based, quantities is significantly affected by the astrophysical or survey systematics that are modeled in the Full Focal Plane simulations. Additionally, we used a principal component analysis to determine the directions in our space of measured features which maximize the variance of simulated Λ CDM maps. Projecting the real CMB sky in this new principal-component basis allowed us to confirm previous suggestions that the cumulative effect of the lack of large-angle CMB power is more anomalous than any one of the individual quantities used to measure it. Though the results of this analysis are fairly unsurprising, its big-picture approach serves to tie together a number of previous results, and by doing so will provide context for future studies of large angle anomalies — for example, using large-angle measurements of CMB polarization.

Finally, in Chapter 5 we turn our focus to the analysis of galaxy survey data, specifically in the context of DES. The increasing size of large scale structure surveys has improved their ability to place precise constraints on the parameters describing dark energy. Much of their constraining power comes from our ability to simultaneously analyze multiple observable probes, which provide complementary information and therefore break degeneracies between cosmological parameters and nuisance parameters associated with systematic uncertainties. As with any precision experiment, it is important to ensure that the results of the DES 3×2 pt combined analysis of galaxy clustering and weak lensing are not subject to confirmation bias. Motivated by this, we introduced a new technique for consistently and flexibly blinding data from multiple observable probes, and assessed its performance for simulated DES Year 3 data. A version of this technique is currently being used for the DES Year 3 analysis, and it will likely be used in other future DES multi-probe tests of Λ CDM. Because DES will serve as a test-bed for the developing analysis techniques for next-generation surveys like LSST and Euclid, in addition to placing interesting cosmological constraints in its own right, this blinding strategy has the potential to serve as a useful tool for many future precision measurements of dark energy.

The aim of all these analyses has been to ensure that analysis pipelines we use to assess the

agreement of Λ CDM with observations of the Universe on large scales are robust, so that the accuracy of the insights they provide about fundamental physics is commensurate with their statistical precision. By the standard of the revolutionary discoveries that took place during the last decades of the twentieth century, which helped enshrine Λ CDM as our standard cosmological model, the recent advances in cosmology have been relatively incremental. Continued progress will depend on our ability to subject Λ CDM to unprecedentedly rigorous tests, and to do this we must take full advantage of the projected increase in the constraining power of near future and next generation cosmological experiments. These tests have the potential to provide new insights into the physics of inflation, distinguish between dynamic dark energy and a cosmological constant, test whether general relativity accurately describes gravity on all scales, and probe the physical properties of dark matter. Thus, in coming decades, it will be exciting to see what careful and deliberate analysis of new cosmological data can reveal about the Universe.

Appendix A

Derivation of LSS C_ℓ 's

Here demonstrate a derivation of the relationship between the three-dimensional matter power spectrum $P(k, z)$ and the angular power spectra for LSS tracers, such as the projected galaxy density fluctuations.

A.1 Projected density

Let $\rho_X(\hat{\mathbf{n}}, z)$ be the comoving volume density of the density of some population X of tracers. (Note that this same derivation will hold for computing the power spectrum of mass density.) Defining $\bar{\rho}(z)$ be the average comoving volume density at the time associated with redshift z , we write the density fluctuations δ as

$$\delta(\mathbf{r}, z) = \frac{\rho(\mathbf{r}, z)}{\bar{\rho}(z)} - 1, \quad (\text{A.1})$$

where \mathbf{r} is a position in comoving coordinates, and the z arguments indicates the function should be evaluated at time $t(z)$, such that

$$\int_t^{t_0} \frac{c dt'}{a(t')} = \int_0^z \frac{c dz'}{H(z)}. \quad (\text{A.2})$$

Next, let $n(\hat{\mathbf{n}})$ be the projected density of tracers in direction $\hat{\mathbf{n}}$ on the sky, and let \bar{n} be the average number density per steradian. Fluctuations in projected density are then

$$\delta_n(\hat{\mathbf{n}}) = \frac{n(\hat{\mathbf{n}})}{\bar{n}} - 1. \quad (\text{A.3})$$

We would like to write $\delta_n(\hat{\mathbf{n}})$ in terms of $\delta(\mathbf{r}, z)$. For conciseness, from here forward we will

write the arguments of three-dimensional functions so that, for example

$$\delta(\hat{\mathbf{n}}, z) \equiv \delta(\hat{\mathbf{n}} r(z), z) = \delta(\mathbf{r}, z). \quad (\text{A.4})$$

We can begin by relating $n(\hat{\mathbf{n}})$ to $\rho(\hat{\mathbf{n}}, z)$ via

$$n(\hat{\mathbf{n}}) = \int dz \frac{c r_A^2(z)}{H(z)} \rho(\hat{\mathbf{n}}, z) = \int dz \frac{c r^2(z)}{H(z)} \bar{\rho}(z) (1 + \delta(\hat{\mathbf{n}}, z)). \quad (\text{A.5})$$

Here $r_A(z)$ is the comoving angular diameter distance, which in a flat Universe is equal to the comoving radial coordinate $r(z)$. The term $c r^2/H$ appears in order to convert the units of the volume element: a factor of c/H comes from the change of the radial integral from dr to dz , and a factor of r^2 is associated with the conversion from tracers per comoving area to tracers per steradian. From here forward we will suppress the integral bounds for conciseness.

We relate \bar{n} to $n(\hat{\mathbf{n}})$ by averaging over all directions on the sky. This gives,

$$\bar{n} = \frac{1}{4\pi} \int d\Omega_{\hat{\mathbf{n}}} n(\hat{\mathbf{n}}) = \frac{1}{4\pi} \int dz \frac{c r^2(z)}{H(z)} \frac{1}{4\pi} \int d\Omega \rho(\hat{\mathbf{n}}, z) \equiv \int dz \frac{d\bar{n}}{dz} \quad (\text{A.6})$$

In the last part of this expression we have defined the quantity $d\bar{n}/dz$ to be the direction-averaged projected number of tracers in a redshift shell $[z, z + dz]$. We can relate this to the average volume density $\bar{\rho}$, by writing down expressions for the total number of tracers in the shell:

$$4\pi \frac{d\bar{n}}{dz} dz = 4\pi \bar{\rho}(z) \frac{c r^2}{H} dz \quad \Rightarrow \quad \frac{d\bar{n}}{dz} = \bar{\rho}(z) \frac{c r^2}{H}. \quad (\text{A.7})$$

Revisiting our expression for $\delta_n(\hat{\mathbf{n}})$, we can now write,

$$\delta_n(\hat{\mathbf{n}}) = \frac{n(\hat{\mathbf{n}}) - \bar{n}}{\bar{n}} \quad (\text{A.8})$$

$$= \left[\int dz \frac{c r^2(z)}{H(z)} \bar{\rho}(z) \right]^{-1} \int dz \frac{c r^2(z)}{H(z)} \bar{\rho}(z) \delta(\hat{\mathbf{n}}, z) \quad (\text{A.9})$$

$$= \left[\int dz \frac{d\bar{n}}{dz} \right]^{-1} \int dz \frac{d\bar{n}}{dz} \delta(\hat{\mathbf{n}}, z). \quad (\text{A.10})$$

To make the notation more concise, we represent normalized version of the redshift distribution dn/dz via

$$\frac{dn}{dz} \equiv \frac{d\bar{n}}{dz} \left[\int dz \frac{d\bar{n}}{dz} \right]^{-1}. \quad (\text{A.11})$$

The expression for fluctuations in projected density then becomes

$$\delta_n(\hat{\mathbf{n}}) = \int dz \frac{dn}{dz} \delta(\hat{\mathbf{n}}, z). \quad (\text{A.12})$$

A.2 Conversion to Fourier space

Now we will convert to Fourier space. Because we are interested in the power spectrum of distributions projected onto the sky, we will work in spherical coordinates. We define $g_{\ell m}^X$ to be the spherical components of the map $\delta n^X(\hat{\mathbf{n}})$, where X labels the map, so

$$\delta n^X(\hat{\mathbf{n}}) = \sum_{\ell m} g_{\ell m}^X Y_{\ell m}(\hat{\mathbf{n}}), \quad (\text{A.13})$$

for spherical harmonic functions $Y_{\ell m}$.

In spherical coordinates, the Fourier transform of the density perturbations $\delta^X(\mathbf{x}, t)$ can be rewritten using a Fourier Bessel decomposition,

$$\delta^X(\mathbf{r}, t) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} \delta^X(\mathbf{k}, t) e^{-i\mathbf{k}\cdot\mathbf{r}} = \sqrt{\frac{2}{\pi}} \sum_{\ell m} \int k dk \delta_{\ell m}^X(k, t) j_{\ell}(kr) Y_{\ell m}(\hat{\mathbf{n}}). \quad (\text{A.14})$$

where $j_{\ell}(x)$ is a Spherical Bessel function.

Referencing Eq. (A.12), we can then relate the functions $\delta_{\ell m}^X(k)$ to the projected Fourier components $g_{\ell m}^X$ via

$$g_{\ell m}^X = \sqrt{\frac{2}{\pi}} \int dz \frac{dn}{dz} \int k dk j_{\ell}(kr) \delta_{\ell m}^X(k, t(z)). \quad (\text{A.15})$$

Additionally, we can use the Rayleigh expansion,

$$e^{-i\mathbf{k}\cdot\mathbf{r}} = e^{-i(kr)\hat{\mathbf{k}}\cdot\hat{\mathbf{n}}} = 4\pi \sum_{\ell m} (-i)^{\ell} j_{\ell}(kr) Y_{\ell m}^*(\hat{\mathbf{k}}) Y_{\ell m}(\hat{\mathbf{n}}), \quad (\text{A.16})$$

to show (after a few lines of calculation) that

$$\delta_{\ell m}^X(k, t) = \sqrt{\frac{\pi}{2}} \frac{(-i)^{\ell}}{2\pi^2} k \int d\Omega_{\hat{\mathbf{k}}} \delta^X(\mathbf{k}, t) Y_{\ell m}^*(\hat{\mathbf{k}}). \quad (\text{A.17})$$

Putting Eqs. (A.15) and (A.17) together, we can write

$$g_{\ell m}^X = \frac{(-i)^\ell}{2\pi^2} \int dz \int k^2 dk j_\ell(kr(z)) \int d\Omega_{\hat{\mathbf{k}}} Y_{\ell m}^*(\hat{\mathbf{k}}) \frac{dn}{dz} \delta^X(\mathbf{k}, t(z)). \quad (\text{A.18})$$

A.3 Angular cross correlations

We are now ready to compute the angular power spectra. Assuming isotropy, the angular power of correlations between maps X and Y is given via

$$C_\ell^{XY} \delta_{\ell'\ell} \delta_{m'm} = \langle [g_{\ell'm'}^X]^* g_{\ell m}^Y \rangle. \quad (\text{A.19})$$

Thus, to get an expression for C_ℓ^{XY} we can use Eq. (A.18) to write an expression for $\langle [g_{\ell'm'}^X]^* g_{\ell m}^Y \rangle$. This expression will be messy and will involve six integrals (over the z , k , and $\hat{\mathbf{k}}$ variables associated with each of the two maps X and Y):

$$\langle [g_{\ell'm'}^X]^* g_{\ell m}^Y \rangle = \frac{1}{4\pi^4} \int dz \int dz' \int k^2 j_\ell(kr(z)) dk \int (k')^2 j_\ell(k'r(z')) dk' \quad (\text{A.20})$$

$$\times \int d\Omega_{\hat{\mathbf{k}}} \int d\Omega_{\hat{\mathbf{k}'}} Y_{\ell m}(\hat{\mathbf{k}}) Y_{\ell m}^*(\hat{\mathbf{k}'}) \frac{dn^X}{dz} \frac{dn^Y}{dz'} \langle [\delta^X(\mathbf{k}, t(z))]^* \delta^Y(\mathbf{k}', t(z')) \rangle \quad (\text{A.21})$$

We can simplify this expression if we can assume that the density fluctuations of tracer X can be related to the underlying matter fluctuations δ via some function, where we use z as a time coordinate:

$$\delta^X(\mathbf{k}, z) = B(k, z) \delta(\mathbf{k}, z). \quad (\text{A.22})$$

For a galaxy population, this function $B(k, z)$ will simply be the galaxy bias, which takes the form $b(z)$ on linear scales. We will leave the expression in this general form however, as this same formalism can be used to compute the power spectrum for other tracers, such as the ISW signal. Additionally, assuming we are working in the regime where linear growth holds, per Eq. (2.32) we can write the density fluctuation at redshift at time t in terms of its form at the present time t_0 as

$$\delta(\mathbf{k}, t) = D(t) \delta(\mathbf{k}, t_0). \quad (\text{A.23})$$

Thus, we can see the expectation value in Eq. (A.20) becomes

$$\langle [\delta^X(\mathbf{k}, t(z))]^* \delta^Y(\mathbf{k}', t(z')) \rangle = D(z)D(z')B(k, z)B(k', z') \langle [\delta(\mathbf{k}, t_0)]^* \delta(\mathbf{k}', t_0) \rangle \quad (\text{A.24})$$

$$= (2\pi)^3 \delta^{(3)}(\mathbf{k} - \mathbf{k}') D(z)D(z') B^X(k, z) B^Y(k', z') P(k, t_0), \quad (\text{A.25})$$

where we have inserted the matter power spectrum defined in Eq. (2.28).

Inserting Eq. (A.25) into Eq. (A.20), and using the facts that

$$\int d^3\mathbf{k} = \int k^2 dk \int d\Omega_{\hat{k}}, \quad \text{and} \quad (\text{A.26})$$

$$\int d\Omega_{\hat{k}} |Y_{\ell m}(\hat{k})|^2 = 1, \quad (\text{A.27})$$

we can derive

$$C_\ell^{XY} = \frac{2}{\pi} \int dz \int dz' \int k^2 dk j_\ell(kr(z)) j_\ell(kr(z')) \quad (\text{A.28})$$

$$\times \frac{dn^X}{dz} \frac{dn^Y}{dz'} D(z)D(z') B^X(k, z) B^Y(k, z') P(k, z=0). \quad (\text{A.29})$$

To get this in the form in which it appears in Chapter 1, we define the window function of tracer X to be

$$W^X(k, z) \equiv \frac{dn^X}{dz} B(k, z), \quad (\text{A.30})$$

and additionally define an associated transfer function

$$I_\ell^X(k) \equiv \int dz D(z) W^X(k, z) j_\ell(kr(z)). \quad (\text{A.31})$$

Inserting these functions, we obtain

$$C_\ell^{XY} = \frac{2}{\pi} \int dk k^2 I_\ell^X(k) I_\ell^Y(k) P(k, 0). \quad (\text{A.32})$$

A.4 The Limber approximation

For high ℓ , we can compute C_ℓ^{XY} with reasonable accuracy using the Limber approximation, which uses the relation [65, 101, 215]

$$j_\ell(x) \approx \sqrt{\frac{\pi}{2(2\ell+1)}} [\delta^{(1)}(\ell + \frac{1}{2} - x)], \quad (\text{A.33})$$

where $\delta^{(1)}$ is the one-dimensional Dirac delta function. This gives the angular cross power between bins X and Y as

$$C_\ell^{XY} = \int dz \frac{H(z)}{cr^2(z)} W^X(k, z) W^Y(k, z) P(k, z) \Big|_{k=(\ell+\frac{1}{2})/r(z)}. \quad (\text{A.34})$$

The factor of H/c comes from the factor of $|\partial r/\partial z|^{-1}$ that appears when one integrates the delta function approximation for $j_\ell(kr(z'))$ over z' . Because the Limber approximation involves fewer integrals, it is much faster computationally than Eq. (A.32). It has an accuracy of about 1% for $\ell \gtrsim 20$, depending on the details of the redshift bin width and mean z [101].

Within the Limber approximation, one can include any nonlinear corrections to clustering simply by substituting the nonlinear power spectrum $P_{NL}(k, z)$, for example one obtained using `halofit` [?, 62] for the linear $P(k, z)$ in Eq. (A.34). How to include nonlinear effects is less straightforward for the exact expression for C_ℓ^{XY} in Eq. (A.32), but one could approach it by promoting the linear growth factor $D(z)$ to become a scale dependent transfer function $D_{NL}(k, z)$ which models the effects of nonlinear growth.

A.5 Window function for the ISW signal

Though the ISW is not a projected density, the angular correlation for its signal can be written in the same form as the C_ℓ 's described here. We just need to do some additional work to find its window function, $W^{\text{ISW}}(k, z)$.

As was noted in Chapter 3, the temperature perturbations associated with the ISW signal can be related to the integral along the line of sight of the time derivative of the gravitational potential $\Phi(\mathbf{r}, t)$. Letting t_* be the time of last scattering, when CMB photons were released, their temperature modulation due to the ISW effect will be

$$\frac{\Delta T}{\bar{T}} \Big|_{\text{ISW}}(\hat{\mathbf{n}}) = \frac{2}{c^2} \int_{t_*}^{t_0} dt \frac{\partial \Phi(\mathbf{r}, t)}{\partial t}. \quad (\text{A.35})$$

We will begin by using the Poisson equation in comoving coordinates,

$$\nabla^2 \Phi(\mathbf{r}, t) = \frac{3}{2} \frac{H_0^2 \Omega_{m0}}{a} \delta(\mathbf{r}, t), \quad (\text{A.36})$$

to relate the potential to density fluctuations. Writing $\Phi(\mathbf{r}, t)$ in terms of its Fourier transform,

$$\Phi(\mathbf{r}, t) = \int d^3 \mathbf{k} \tilde{\Phi}(\mathbf{k}, t) e^{-i\mathbf{k} \cdot \mathbf{r}}, \quad (\text{A.37})$$

and inserting that expression into Eq. (A.36), we can show that

$$-k^2 \tilde{\Phi}(\mathbf{k}, t) = \frac{3}{2} \frac{H_0^2 \Omega_{m0}}{a} \delta(\mathbf{k}, t). \quad (\text{A.38})$$

Plugging this into the expression for the ISW signal given in Eq. (A.35), using the linear growth relation $\delta(\mathbf{k}, t) = D(t)\delta(\mathbf{k}, t_0)$ and noting that

$$\frac{\partial}{\partial t} \left(\frac{D(t)}{a} \right) = \dot{a} \frac{\partial}{\partial a} \left(\frac{D(a)}{a} \right) = H(f-1) \left(\frac{D(a)}{a} \right), \quad (\text{A.39})$$

we obtain the relation

$$\left. \frac{\Delta T}{\bar{T}} \right|_{ISW}(\hat{\mathbf{n}}) = -\frac{3H_0^2 \Omega_{m0}}{c^2} \int_{t_*}^{t_0} dt \int \frac{d^3 \mathbf{k}}{(2\pi)^3} e^{-i\mathbf{k}\cdot\mathbf{r}} \frac{\delta(\mathbf{k}, t) H(f-1)}{a}. \quad (\text{A.40})$$

We would like to make this expression look like Eq. (A.18). We will do so by converting the integral to be over redshift ($dt = -adz/H$), reversing the integration bounds, and using the Rayleigh expansion to turn the Fourier exponent into a Bessel function and spherical harmonics. Defining $a_{\ell m}^{\text{ISW}}$ to the spherical components of the ISW signal, we obtain

$$a_{\ell m} = \frac{(-i)^\ell}{2\pi^2} \int_0^{z_*} dz \int k^2 dk j_\ell(kr(z)) \int d\Omega_{\hat{k}} Y_{\ell m}^*(\hat{k}) \left[\frac{3H_0^2 \Omega_{m0} (1-f(z))}{c^2 k^2} \delta(\mathbf{k}, z) \right]. \quad (\text{A.41})$$

Comparing this to Eq. (A.18), we see we can define the kernel for the ISW signal in terms of an effective density distribution

$$\frac{dn^{\text{ISW}}}{dz} \delta^{\text{ISW}}(\mathbf{k}, z) = \Theta(z_* - z) \left[\frac{3H_0^2 \Omega_{m0} (1-f(z))}{c^2 k^2} \right] \delta(\mathbf{k}, z), \quad (\text{A.42})$$

where $\Theta(x)$ is the Heaviside step function.

Thus, we see that the window function for the ISW signal is

$$W^{\text{ISW}}(k, z) = \Theta(z_* - z) \frac{3H_0^2 \Omega_{m0} (1-f(z))}{c^2 k^2} \quad (\text{A.43})$$

Appendix B

Calibration error formalism

In Section 3.5 we study the impact of photometric calibration errors on ISW signal reconstruction. We model them using a direction-dependent calibration error field $c(\hat{\mathbf{n}})$ via

$$N^{\text{obs}}(\hat{\mathbf{n}}) = (1 + c(\hat{\mathbf{n}}))N(\hat{\mathbf{n}}), \quad (\text{B.1})$$

where $\hat{\mathbf{n}}$ is the direction on the sky, N^{obs} is the observed number of galaxies, and N is the true number of galaxies. Here, we present the calculations necessary to describe how this modifies the galaxy C_ℓ and which we used above to predict how calibration errors will impact our reconstructions quality statistics. Our notation follows that by Ref. [119].

We will define fluctuations in the true and observed number density as δ and δ^{obs} , respectively, and write them in terms of spherical components,

$$\delta(\hat{\mathbf{n}}) = \frac{N(\hat{\mathbf{n}})}{\bar{n}} - 1 \equiv \sum_{\ell m} g_{\ell m} Y_{\ell m}(\hat{\mathbf{n}}) \quad (\text{B.2})$$

$$\delta^{\text{obs}}(\hat{\mathbf{n}}) = \frac{N^{\text{obs}}(\hat{\mathbf{n}})}{\bar{n}^{\text{obs}}} - 1 \equiv \sum_{\ell m} t_{\ell m} Y_{\ell m}. \quad (\text{B.3})$$

Additionally, we will define a parameter ϵ to relate the true and observed average number densities,

$$\bar{n}^{\text{obs}} = \bar{n}(1 + \epsilon), \quad (\text{B.4})$$

and use $c_{\ell m}$ to denote the spherical components of the calibration error field $c(\hat{\mathbf{n}})$. Each galaxy map can have its own calibration error field, and so we will use superscripts (e.g., $g_{\ell m}^i$, $c_{\ell m}^i$, and $t_{\ell m}^i$) to denote components associated with LSS map i .

Our goal is to find a relation between the observed galaxy power T_ℓ^{ij} , the true power C_ℓ^{ij} , and the properties of the calibration error field $C_\ell^{\text{cal},ij}$. To do this, we start by relating the

spherical components of the fields. We note that observed number density fluctuations are

$$\delta^{\text{obs}}(\hat{\mathbf{n}}) = \frac{\delta + c + \delta c - \epsilon}{(1 + \epsilon)}, \quad (\text{B.5})$$

where we suppress the $\hat{\mathbf{n}}$ arguments to simplify notation. After some algebra, we can write

$$t_{\ell m}^i = (1 + \epsilon^i)^{-1} \left[-\sqrt{4\pi} \delta_{\ell 0} \epsilon^i + g_{\ell m}^i + c_{\ell m}^i \right. \quad (\text{B.6})$$

$$\left. + \sum_{\ell_1 \ell_2 m_1 m_2} c_{\ell_2 m_2}^i g_{\ell_1 m_1}^i R_{m m_1 m_2}^{\ell \ell_1 \ell_2} \right]. \quad (\text{B.7})$$

In this expression, $\delta_{\ell 0}$ is a Kronecker delta, and the multiplicative term

$$R_{m m_1 m_2}^{\ell \ell_1 \ell_2} \equiv \int d\Omega Y_{\ell m}^*(\hat{\mathbf{n}}) Y_{\ell_2 m_2}(\hat{\mathbf{n}}) Y_{\ell_1 m_1}(\hat{\mathbf{n}}) \quad (\text{B.8})$$

is related to Wigner-3j symbols.

We define the cross-power between two observed maps via

$$T_{\ell}^{ij} \equiv \sum_m \frac{\langle t_{\ell m}^i t_{\ell m}^{j*} \rangle}{2\ell + 1} \quad (\text{B.9})$$

and that of the calibration error fields as

$$C_{\ell}^{\text{cal},ij} \equiv \sum_m \frac{c_{\ell m}^i c_{\ell m}^{j*}}{2\ell + 1}. \quad (\text{B.10})$$

Note that these definitions do not preclude the possibility that the $c_{\ell m}$ could introduce correlations between different (ℓ, m) modes. The fact that we only show correlations between modes with matching ℓ and m reflects the (potentially biased) measurement that would be made even if one assumes that they do not.

The expression for T_{ℓ}^{ij} in terms of $g_{\ell m}$, $c_{\ell m}$ is fairly involved, though it can be simplified to some extent using Wigner-3j symbol identities. For the purposes of this study, we approximate it by only including additive components—that is, neglecting all terms containing $R_{m m_1 m_2}^{\ell \ell_1 \ell_2}$. Doing this, and using the fact that

$$\langle \epsilon^i \rangle = \frac{c_{00}^i}{\sqrt{4\pi}} = \sqrt{\frac{C_{\ell=0}^{\text{cal},i}}{4\pi}}, \quad (\text{B.11})$$

we write

$$T_\ell^{ij} = \frac{C_\ell^{gij} + C_\ell^{cij} - \delta_{\ell 0} c_{00}^i c_{00}^j}{(1 + c_{00}^i/\sqrt{4\pi})(1 + c_{00}^j/\sqrt{4\pi})}. \quad (\text{B.12})$$

This is the expression given in Eq. (3.32) and is what is used to compute expectations values of ISW reconstruction quality statistics in Section 3.5.

Appendix C

Large-noise limit of ISW reconstruction s statistic

In Section 3.5.2, and particularly in Fig. 3.14, we saw that as the amplitude of calibration error fluctuations gets large the ratio between the rms of reconstructed map residuals and the rms of the true ISW map, s , approaches a constant value. Here we outline why this occurs.

Recall from Eq. (3.23) that our theoretical estimator $\langle s \rangle$ is written

$$\langle s \rangle = \frac{\sqrt{\langle \sigma_{\text{rec}} \rangle^2 + \langle \sigma_{\text{ISW}} \rangle^2 - 2 \sum_{\ell i} (2\ell + 1) R_{\ell}^i \tilde{C}_{\ell}^{\text{ISW}-i}}}{\langle \sigma_{\text{ISW}} \rangle}, \quad (\text{C.1})$$

where

$$\begin{aligned} \langle \sigma_{\text{ISW}} \rangle &= \sqrt{\sum_{\ell} (2\ell + 1) \tilde{C}_{\ell}^{\text{ISW}}}, \text{ and} \\ \langle \sigma_{\text{rec}} \rangle &= \sqrt{\sum_{\ell ij} (2\ell + 1) R_{\ell}^i R_{\ell}^j \tilde{C}_{\ell}^{ij}}. \end{aligned} \quad (\text{C.2})$$

In the case with a single LSS map, which we focus on here for simplicity, the reconstruction filter is

$$R_{\ell}^i = \frac{C_{\ell}^{\text{gal-ISW}}}{C_{\ell}^{\text{gal}}}. \quad (\text{C.3})$$

For clarity, and in contrast with the notation in the main text, here we use tildes (as in \tilde{C}_{ℓ}) to denote the C_{ℓ}^{true} which are associated with observed or simulated maps. The C_{ℓ} with no tilde will be the C_{ℓ}^{model} used to construct the ISW estimator.

Let us examine how the various terms scale as we increase the amplitude of calibration errors. As the level of calibration errors—or any form of noise—gets large,

$$\tilde{C}_\ell^{\text{gal}} \xrightarrow{\text{large } A} C_\ell^{\text{noise}} \propto A \quad (\text{C.4})$$

where C_ℓ^{noise} is the noise power spectrum and A is a measure of its amplitude. The observed ISW power $\tilde{C}_\ell^{\text{ISW}}$ and ISW-galaxy cross-power $\tilde{C}_\ell^{\text{gal-ISW}}$ will not depend on A .

For the calibration error studies in Section 3.5, we focused on the case of residual calibration errors, which are not accounted for in the ISW estimator. In this scenario, any excess in observed power will be interpreted as a bias and fit for via

$$\bar{b}^2 C_\ell^{\text{gal}} = \tilde{C}_\ell^{\text{gal}}, \quad (\text{C.5})$$

according to the procedure described in Section 3.1.4. Because C_ℓ^{gal} is independent of A , the resulting best fit value will be $\bar{b}^{\text{fit}} \propto \sqrt{A}$. The model $C_\ell(\bar{b}^{\text{fit}})$ scales accordingly,

$$C_\ell^{\text{gal}}(\bar{b}^{\text{fit}}) \propto A, \quad (\text{C.6})$$

$$C_\ell^{\text{gal-ISW}}(\bar{b}^{\text{fit}}) \propto \sqrt{A}, \quad (\text{C.7})$$

$$R_\ell \propto \frac{1}{\sqrt{A}}. \quad (\text{C.8})$$

Examining the terms in Eq. (C.1), we see that $\langle \sigma_{\text{rec}} \rangle$ and $\langle \sigma_{\text{ISW}} \rangle$ will approach constants as A grows, while the cross-term will go to zero like $A^{-1/2}$. Thus, in the case of unmodeled noise contributions to the galaxy maps, in the limit of large noise,

$$\langle s \rangle \xrightarrow{\text{large } A} \frac{\sqrt{\langle \sigma_{\text{rec}} \rangle^2 + \langle \sigma_{\text{ISW}} \rangle^2}}{\langle \sigma_{\text{ISW}} \rangle}. \quad (\text{C.9})$$

This is a constant greater than 1, in agreement with our results in the right panel of Fig. 3.14.

In contrast, if the C_ℓ used in the ISW estimator correctly model the level of galaxy noise—as occurs in the shot noise tests in Section 3.2.4—the best fit bias parameter \bar{b}^{fit} will remain close to 1. In that case, the fact that noise is properly accounted for means that

$$C_\ell^{\text{gal}} = \tilde{C}_{\text{ell}}^{\text{gal}} \propto A \quad (\text{C.10})$$

while all other C_ℓ and \tilde{C}_ℓ are independent of A . In this case, as the noise power dominates

over that of galaxies, the estimator operator goes to zero according to

$$R_\ell \propto \frac{1}{A}. \quad (\text{C.11})$$

This means that for large levels of properly modeled noise, the reconstructed map amplitude goes to zero. This causes $\langle \sigma_{\text{rec}} \rangle$ and the cross-term in $\langle s \rangle$ to go to zero and so the reconstruction residuals are just a measure of the true ISW map:

$$\langle s \rangle \xrightarrow{\text{large } A} \frac{\sqrt{\langle \sigma_{\text{ISW}} \rangle^2}}{\langle \sigma_{\text{ISW}} \rangle} = 1. \quad (\text{C.12})$$

Appendix D

Additional plots for blinding tests

Here we present additional plots associated with the Multinest tests discussed in Section 5.4.4. Figs. D.1 and D.2 show the parameter estimation contours produced from running Multinest on noiseless data vectors for realization 54 and 69, respectively. Studying them, we note first that, even for the unblinded data for both of these realizations, the parameter constraints are biased with respect to the nominal “true” value of b_1 . Because the b_1 marginalized constraints for all realizations plotted (these two and realization 1 shown in Chapter 5 seems to be centered on its fiducial value, this may be a sign of a bug in the data vector generation code: it is possible that b_1 is not actually being varied. Even if that is the case, however, it will not affect our $\Delta\chi^2$ considerations.

In contrast to the results for realization 1 (shown in Fig. 5.15), the contours for additive and multiplicative blinding are significantly offset from one another for realization 54, particularly in the w and b_1 directions. The additive blinding case for realization 69 also shows reasonable looking contours, though they are also not centered on the region we expect based on the input cosmology Θ_{obs} and shift $\Delta\Theta$. The fit to the multiplicatively blinded data vector for realization 69 is clearly failing: the constraint contours are hitting the lower prior bounds for w . This suggests that some characteristic of these data vectors makes fitting them difficult, though it is unclear what that is.

To gain insight into properties of these “badly behaved” realizations, we additionally plot the data vectors for realization 54 in Figs. D.3 and D.4, and for realization 69 in Figs. D.5 and D.6. On these plots, we show the predictions for the best fit model Θ_{bl} (found using Maxlike and blinded noisy data) with dotted lines. We also display the Maxlike $\Delta\chi^2$ values in gray text on the $\xi_{\pm}(\theta)$ plots. The fact that the Maxlike fits are failing can be most obviously seen when looking at the galaxy-galaxy correlations $w(\theta)$.




Figure D.1: Parameter constraints for realization 54. Dashed black lines show the cosmology at which the unblinded data was simulated. The blinding factors used shifts $\Delta\sigma_8 = -0.009$ and $\Delta w = +0.49$. The dashed line for b_1 cannot be seen because it is at $b_1 = 1.24$, significantly below where the contour is.




Figure D.2: Parameter constraints for realization 54. Dashed black lines show the cosmology at which the unblinded data was simulated. The blinding factors used shifts $\Delta\sigma_8 = +0.033$ and $\Delta w = -0.36$. The sampler is clearly failing for the multiplicatively blinded data vector.

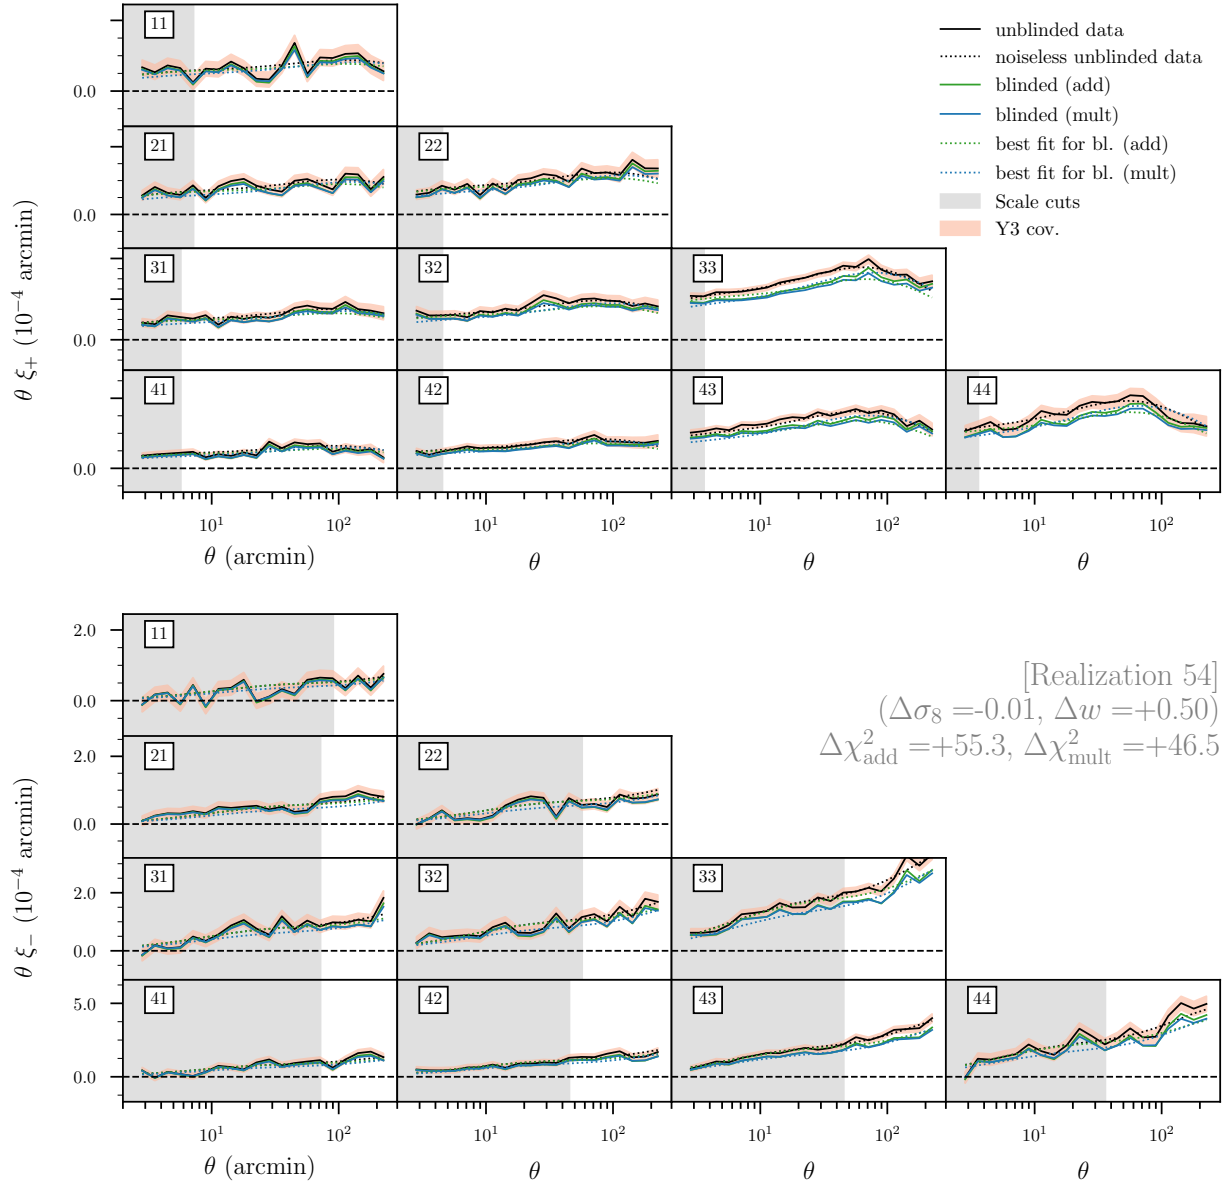


Figure D.3: Plot of $\xi_{\pm}(\theta)$ for realization 54. The dotted and solid black lines show the noiseless and noisy unblinded data, respectively. The solid green line shows the additively blinded noisy data vector, and the solid blue line shows the same for multiplicative blinding. The dotted green and blue lines shows the model prediction computed at the Maxlike best fit parameters for their corresponding blinded data vector. Vertical gray bands show scale cuts, and the shaded region around the unblinded data shows the DES Y3 expected statistical errors. Each subplot shows the cross correlations between a different pair of redshift bins labeled by numbers in white boxes. The bin label numbers go up in order of increasing redshift.

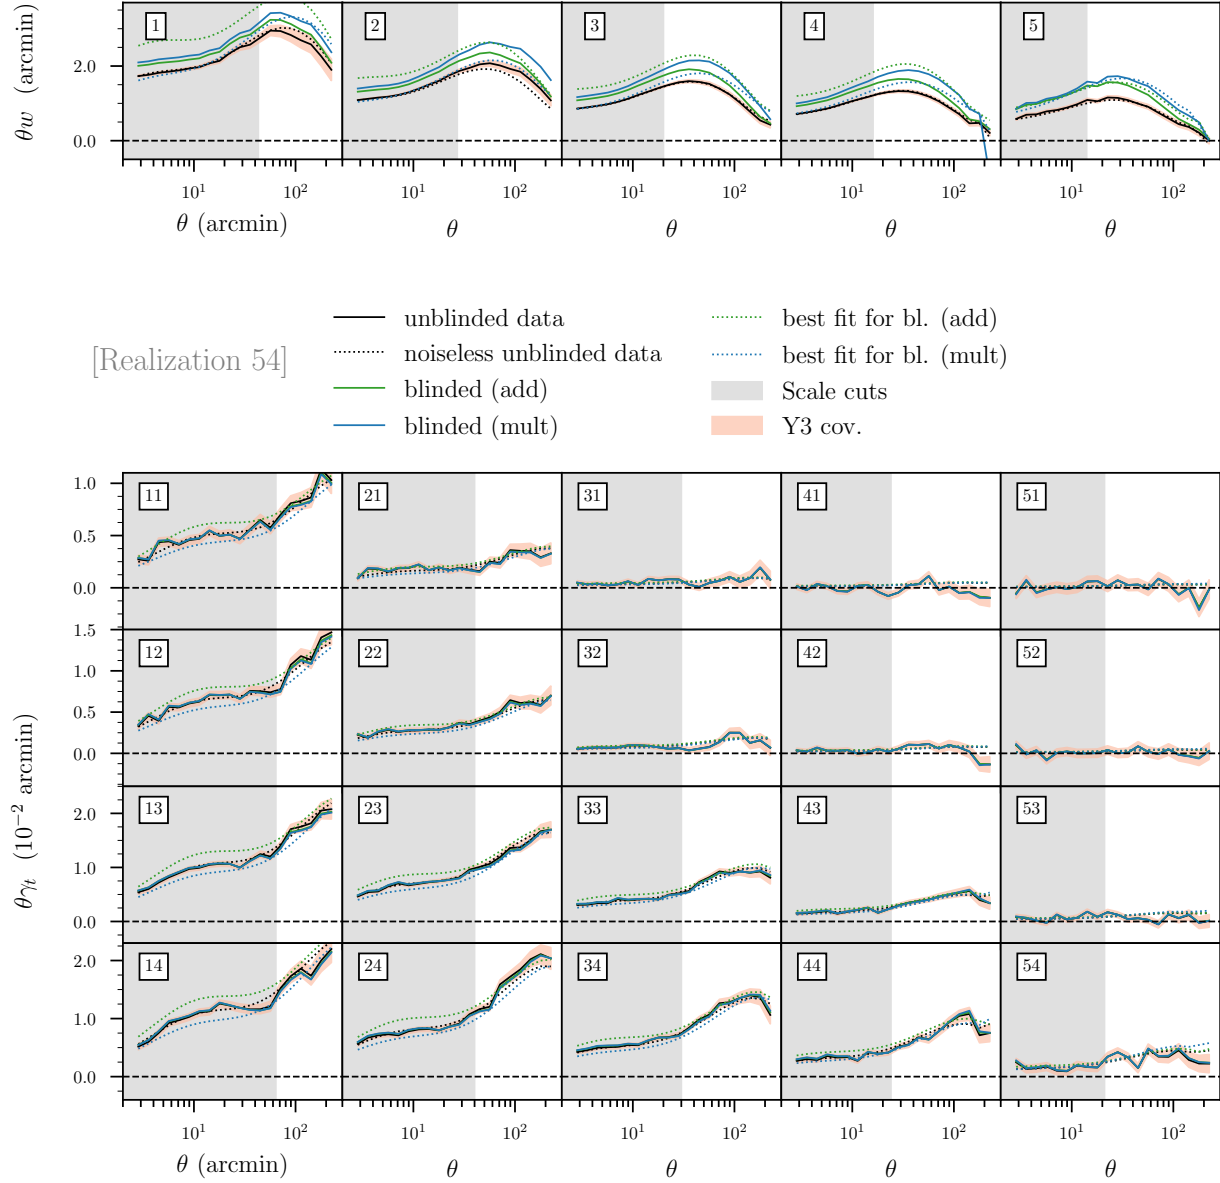


Figure D.4: Plot of $w(\theta)$ and $\gamma_t(\theta)$ for realization 54, using the same conventions as Fig. 5.6.

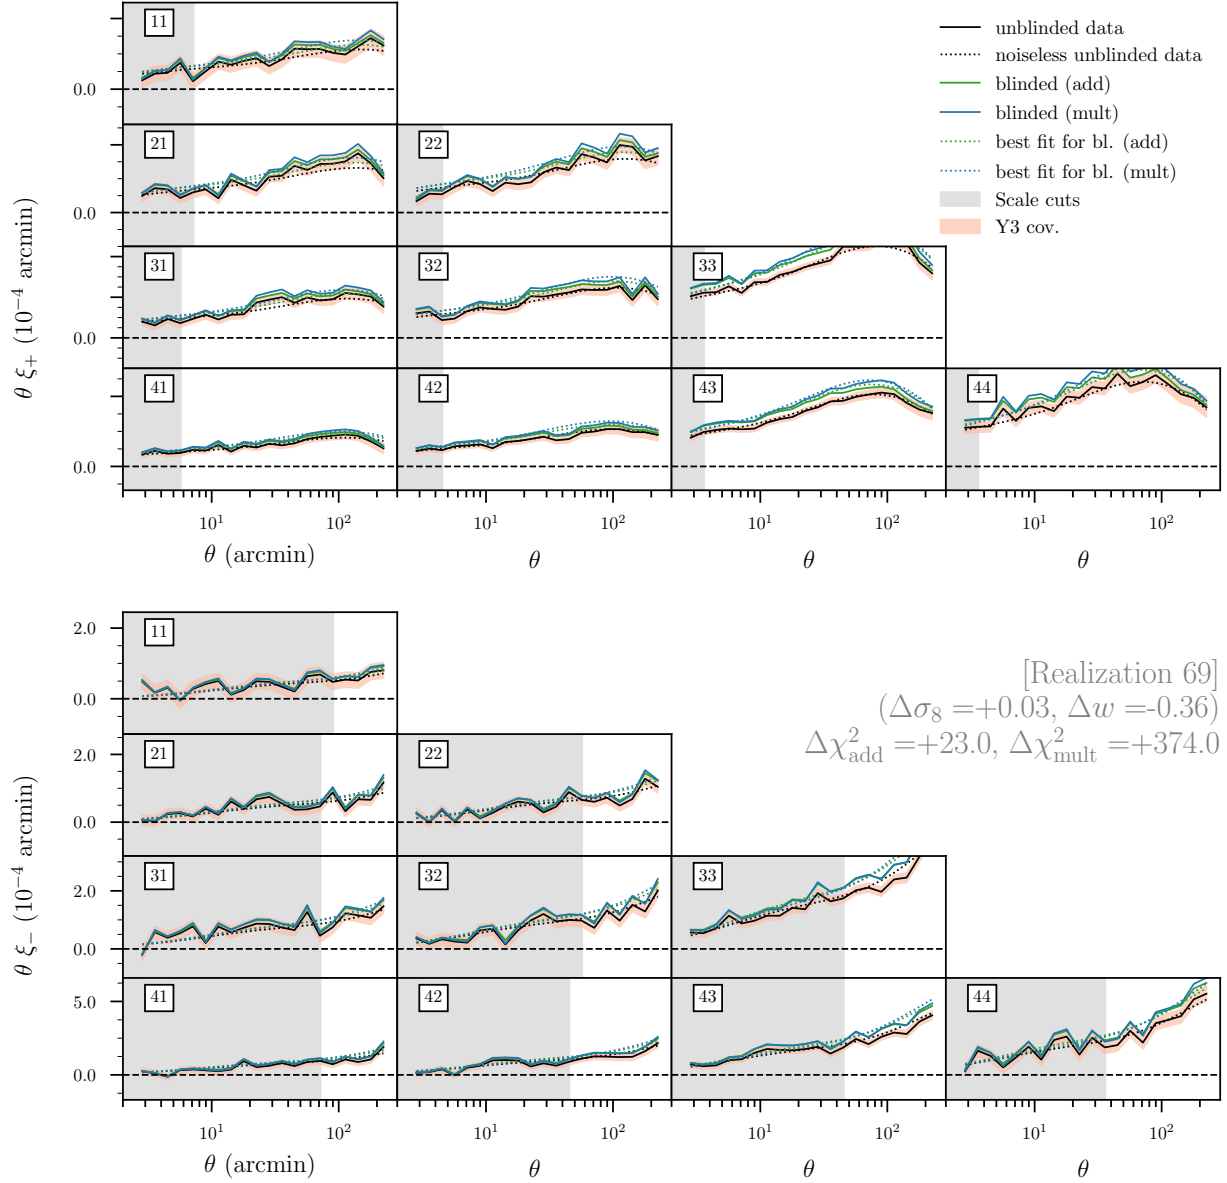


Figure D.5: Plot of $\xi_{\pm}(\theta)$ for realization 69, using the same conventions as Fig. 5.6.

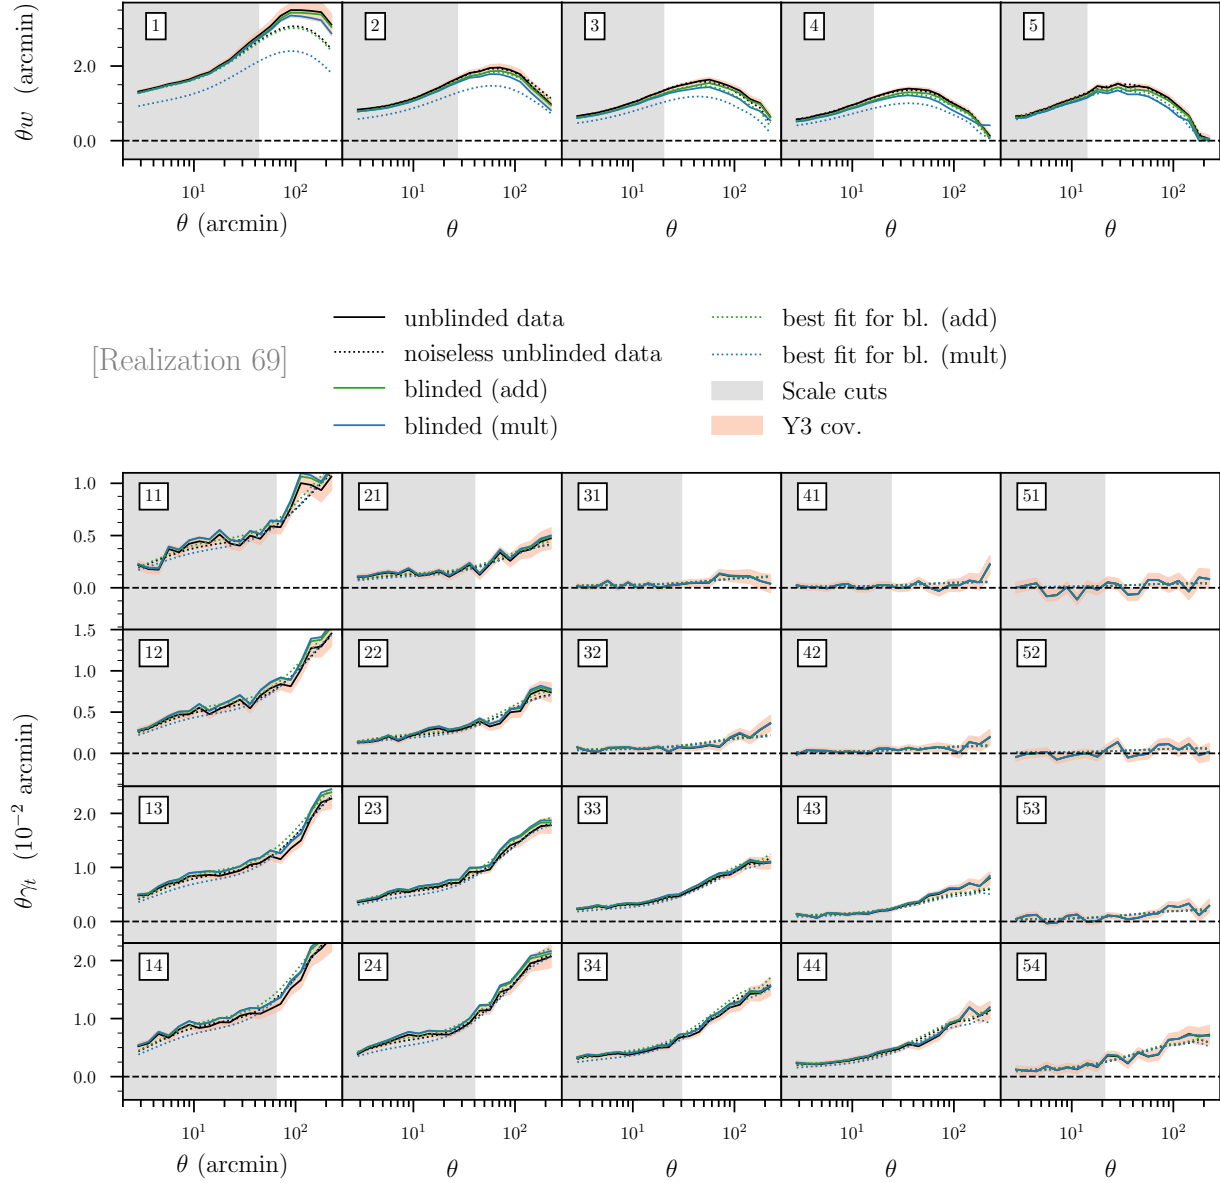


Figure D.6: Plot of $w(\theta)$ and $\gamma_t(\theta)$ for realization 69, using the same conventions as Fig. 5.6.

Bibliography

- [1] Anaconda, Inc., *Anaconda software distribution.*, 2018. [Computer Software. Version 4.3.22].
- [2] E. Jones, T. Oliphant, P. Peterson, et al., *SciPy: Open source scientific tools for Python*, 2001–.
- [3] T. E. Oliphant, *A guide to NumPy*. Trelgol Publishing, USA, 2006.
- [4] J. D. Hunter, *Matplotlib: A 2D Graphics Environment*, Computing in Science & Engineering **9** (2007) 90–95.
- [5] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman, *HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, Astrophys. J. **622** (2005) 759–771, [[astro-ph/0409513](#)].
- [6] F. Pérez and Granger, *IPython: A system for interactive scientific computing*, Computing in Science & Engineering **9** (2007) 21–29.
- [7] J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson, S. Bridle, S. Sehrish, and J. Kowalkowski, *CosmoSIS: Modular Cosmological Parameter Estimation*, Astron. Comput. **12** (2015) 45–59, [[arXiv:1409.3409](#)].
- [8] J. Muir and D. Huterer, *Reconstructing the integrated Sachs-Wolfe map with galaxy surveys*, Phys. Rev. **D94** (2016), no. 4 043503, [[arXiv:1603.06586](#)].
- [9] N. Weaverdyck, J. Muir, and D. Huterer, *Integrated Sachs-Wolfe map reconstruction in the presence of systematic errors*, Phys. Rev. **D97** (2018), no. 4 043515, [[arXiv:1709.08661](#)].
- [10] J. Muir, S. Adhikari, and D. Huterer, *The Covariance of CMB anomalies*, [arXiv:1806.02354](#).
- [11] A. Manzotti and S. Dodelson, *Mapping the Integrated Sachs-Wolfe Effect*, Phys.Rev. **D90** (2014), no. 12 123009, [[arXiv:1407.5623](#)].
- [12] **DES** Collaboration, T. M. C. Abbott et al., *Dark Energy Survey Year 1 Results: Cosmological Constraints from Galaxy Clustering and Weak Lensing*, [arXiv:1708.01530](#).

- [13] A. Einstein, *The Field Equations of Gravitation*, Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.) **1915** (1915) 844–847.
- [14] E. Hubble, *A relation between distance and radial velocity among extra-galactic nebulae*, Proc. Nat. Acad. Sci. **15** (1929) 168–173.
- [15] H. S. Leavitt and E. C. Pickering, *Periods of 25 Variable Stars in the Small Magellanic Cloud*, Harvard Obs. Circ. **173** (1912) 1–3.
- [16] V. M. Slipher, *The Radial Velocity of the Andromeda Nebula*, Lowell Obs. Bulletin **2** (1913) 56.
- [17] V. M. Slipher, *Spectroscopic Observations of Nebulae*, Popular Astronomy **23** (1915) 21.
- [18] V. M. Slipher, *Nebulae*, Proc. Am. Phil. Soc. **56** (1917) 403–409.
- [19] V. M. Slipher, *Two Nebulae with Unparalleled Velocities*, Popular Astronomy **29** (1921) 128.
- [20] R. W. Smith, *The Origins of the Velocity-Distance Relation*, Journal for the History of Astronomy **10** (1979) 133–165.
- [21] L. A. Thompson, *Vesto Slipher and the First Galaxy Redshifts*, [arXiv:1108.4864](https://arxiv.org/abs/1108.4864).
- [22] A. Friedman, *On the Curvature of space*, Z. Phys. **10** (1922) 377–386. [Gen. Rel. Grav.31,1991(1999)].
- [23] A. Friedmann, *On the possibility of a world with constant negative curvature of space*, Z. Phys. **21** (1924) 326–332. [Gen. Rel. Grav.31,2001(1999)].
- [24] G. Lemaitre, *A Homogeneous Universe of Constant Mass and Growing Radius Accounting for the Radial Velocity of Extragalactic Nebulae*, Annales Soc. Sci. Bruxelles A **47** (1927) 49–59. [Gen. Rel. Grav.45,no.8,1635(2013)].
- [25] G. Lemaitre, *The Expanding Universe*, Mon. Not. Roy. Astron. Soc. **91** (1931) 490–501.
- [26] H. P. Robertson, *Kinematics and World-Structure*, Astrophys. J. **82** (1935) 284–301.
- [27] H. P. Robertson, *Kinematics and World-Structure. 2*, Astrophys. J. **83** (1935) 187–201.
- [28] H. P. Robertson, *Kinematics and World-Structure. 3*, Astrophys. J. **83** (1936) 257–271.
- [29] A. G. Walker, *On Milne's Theory of WorldStructure*, Proceedings of the London Mathematical Society **s2-42** (1937), no. 1 90–127.

- [30] A. Einstein, *Cosmological Considerations in the General Theory of Relativity*, Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.) **1917** (1917) 142–152.
- [31] A. H. Guth, *The Inflationary Universe: A Possible Solution to the Horizon and Flatness Problems*, Phys. Rev. **D23** (1981) 347–356.
- [32] A. H. Guth and S. Y. Pi, *Fluctuations in the New Inflationary Universe*, Phys. Rev. Lett. **49** (1982) 1110–1113.
- [33] **Supernova Cosmology Project** Collaboration, S. Perlmutter et al., *Measurements of Omega and Lambda from 42 high redshift supernovae*, Astrophys. J. **517** (1999) 565–586, [[astro-ph/9812133](#)].
- [34] **Supernova Search Team** Collaboration, A. G. Riess et al., *Observational evidence from supernovae for an accelerating universe and a cosmological constant*, Astron. J. **116** (1998) 1009–1038, [[astro-ph/9805201](#)].
- [35] **Planck** Collaboration, R. Adam et al., *Planck 2015 results. I. Overview of products and scientific results*, Astron. Astrophys. **594** (2016) A1, [[arXiv:1502.01582](#)].
- [36] **ACTPol** Collaboration, S. Naess et al., *The Atacama Cosmology Telescope: CMB Polarization at $200 < \ell < 9000$* , JCAP **1410** (2014), no. 10 007, [[arXiv:1405.5524](#)].
- [37] R. J. Thornton et al., *The Atacama Cosmology Telescope: The polarization-sensitive ACTPol instrument*, Astrophys. J. Suppl. **227** (2016), no. 2 21, [[arXiv:1605.06569](#)].
- [38] **ACTPol** Collaboration, T. Louis et al., *The Atacama Cosmology Telescope: Two-Season ACTPol Spectra and Parameters*, JCAP **1706** (2017), no. 06 031, [[arXiv:1610.02360](#)].
- [39] **SPT** Collaboration, T. de Haan et al., *Cosmological Constraints from Galaxy Clusters in the 2500 square-degree SPT-SZ Survey*, Astrophys. J. **832** (2016), no. 1 95, [[arXiv:1603.06522](#)].
- [40] **SPT-3G** Collaboration, B. A. Benson et al., *SPT-3G: A Next-Generation Cosmic Microwave Background Polarization Experiment on the South Pole Telescope*, Proc. SPIE Int. Soc. Opt. Eng. **9153** (2014) 91531P, [[arXiv:1407.2973](#)].
- [41] **SDSS** Collaboration, M. Betoule et al., *Improved cosmological constraints from a joint analysis of the SDSS-II and SNLS supernova samples*, Astron. Astrophys. **568** (2014) A22, [[arXiv:1401.4064](#)].
- [42] **BOSS** Collaboration, K. S. Dawson et al., *The Baryon Oscillation Spectroscopic Survey of SDSS-III*, Astron. J. **145** (2013) 10, [[arXiv:1208.0022](#)].
- [43] **BOSS** Collaboration, S. Alam et al., *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample*, Mon. Not. Roy. Astron. Soc. **470** (2017), no. 3 2617–2652, [[arXiv:1607.03155](#)].

- [44] H. Hildebrandt et al., *KiDS-450: Cosmological parameter constraints from tomographic weak gravitational lensing*, Mon. Not. Roy. Astron. Soc. **465** (2017) 1454, [[arXiv:1606.05338](#)].
- [45] S. Joudaki et al., *KiDS-450: Testing extensions to the standard cosmological model*, Mon. Not. Roy. Astron. Soc. **471** (2017) 1259, [[arXiv:1610.04606](#)].
- [46] **DES** Collaboration, T. Abbott et al., *The Dark Energy Survey*, [astro-ph/0510346](#).
- [47] **CMB-S4** Collaboration, K. N. Abazajian et al., *CMB-S4 Science Book, First Edition*, [arXiv:1610.02743](#).
- [48] **DESI** Collaboration, *DESI Final Design Report. 1. Science, Targeting, and Survey Design*, .
- [49] **LSST** Collaboration, P. Marshall et al., *Science-Driven Optimization of the LSST Observing Strategy*, [arXiv:1708.04058](#).
- [50] L. Amendola et al., *Cosmology and fundamental physics with the Euclid satellite*, Living Rev. Rel. **21** (2018), no. 1 2, [[arXiv:1606.00180](#)].
- [51] C. L. Carilli and S. Rawlings, *Science with the Square Kilometer Array: Motivation, key science projects, standards and assumptions*, New Astron. Rev. **48** (2004) 979, [[astro-ph/0409274](#)].
- [52] **WFIRST** Collaboration, O. Dor et al., *WFIRST Science Investigation Team "Cosmology with the High Latitude Survey" Annual Report 2017*, [arXiv:1804.03628](#).
- [53] B. Ryden, Introduction to cosmology. Cambridge University Press, 2016.
- [54] S. Dodelson, Modern Cosmology. Academic Press, Amsterdam, 2003.
- [55] M. S. Longair, Galaxy Formation. Astronomy and Astrophysics Library. Springer, Heidelberg, Germany, 2008.
- [56] R. M. Wald, General Relativity. Chicago Univ. Pr., Chicago, USA, 1984.
- [57] D. Huterer, *Physics 526: Cosmology II (Lecture notes)*, 2016.
- [58] D. Baumann, *Cosmology (Lecture notes)*, 2016.
- [59] D. Baumann, *Inflation (TASI 09 Lecture Notes)*, [arXiv:0907.5424](#).
- [60] A. Lewis, A. Challinor, and A. Lasenby, *Efficient computation of CMB anisotropies in closed FRW models*, Astrophys. J. **538** (2000) 473–476, [[astro-ph/9911177](#)].
- [61] C. Howlett, A. Lewis, A. Hall, and A. Challinor, *CMB power spectrum parameter degeneracies in the era of precision cosmology*, JCAP **1204** (2012) 027, [[arXiv:1201.3654](#)].

- [62] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri, *Revising the Halofit Model for the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **761** (2012) 152, [[arXiv:1208.2701](#)].
- [63] A. Mead, J. Peacock, C. Heymans, S. Joudaki, and A. Heavens, *An accurate halo model for fitting non-linear cosmological power spectra and baryonic feedback models*, *Mon. Not. Roy. Astron. Soc.* **454** (2015), no. 2 1958–1975, [[arXiv:1505.07833](#)].
- [64] A. Mead, C. Heymans, L. Lombriser, J. Peacock, O. Steele, and H. Winther, *Accurate halo-model matter power spectra with dark energy, massive neutrinos and modified gravitational forces*, *Mon. Not. Roy. Astron. Soc.* **459** (2016), no. 2 1468–1488, [[arXiv:1602.02154](#)].
- [65] P. Lemos, A. Challinor, and G. Efstathiou, *The effect of Limber and flat-sky approximations on galaxy weak lensing*, *JCAP* **1705** (2017), no. 05 014, [[arXiv:1704.01054](#)].
- [66] M. Bartelmann and P. Schneider, *Weak gravitational lensing*, *Phys. Rept.* **340** (2001) 291–472, [[astro-ph/9912508](#)].
- [67] **DES** Collaboration, E. Krause et al., *Dark Energy Survey Year 1 Results: Multi-Probe Methodology and Simulated Likelihood Analyses*, Submitted to: *Phys. Rev. D* (2017) [[arXiv:1706.09359](#)].
- [68] **DES** Collaboration, M. A. Troxel et al., *Dark Energy Survey Year 1 Results: Cosmological Constraints from Cosmic Shear*, [[arXiv:1708.01538](#)].
- [69] P. Schneider, L. Van Waerbeke, and Y. Mellier, *B-modes in cosmic shear from source redshift clustering*, *Astron. Astrophys.* **389** (2002) 729–741, [[astro-ph/0112441](#)].
- [70] W. Hu and N. Sugiyama, *The Small scale integrated Sachs-Wolfe effect*, *Phys. Rev.* **D50** (1994) 627–631, [[astro-ph/9310046](#)].
- [71] R. Sachs and A. Wolfe, *Perturbations of a cosmological model and angular variations of the microwave background*, *Astrophys.J.* **147** (1967) 73–90.
- [72] R. Bean and O. Dore, *Probing dark energy perturbations: The Dark energy equation of state and speed of sound as measured by WMAP*, *Phys. Rev.* **D69** (2004) 083503, [[astro-ph/0307100](#)].
- [73] J. Weller and A. M. Lewis, *Large scale cosmic microwave background anisotropies and dark energy*, *Mon. Not. Roy. Astron. Soc.* **346** (2003) 987–993, [[astro-ph/0307104](#)].
- [74] Y.-S. Song, I. Sawicki, and W. Hu, *Large-Scale Tests of the DGP Model*, *Phys. Rev.* **D75** (2007) 064003, [[astro-ph/0606286](#)].
- [75] R. G. Crittenden and N. Turok, *Looking for Lambda with the Rees-Sciama effect*, *Phys.Rev.Lett.* **76** (1996) 575, [[astro-ph/9510072](#)].

- [76] S. Boughn and R. Crittenden, *A Correlation of the cosmic microwave sky with large scale structure*, Nature **427** (2004) 45–47, [[astro-ph/0305001](#)].
- [77] P. Fosalba, E. Gaztanaga, and F. Castander, *Detection of the ISW and SZ effects from the CMB-galaxy correlation*, Astrophys. J. **597** (2003) L89–92, [[astro-ph/0307249](#)].
- [78] **WMAP** Collaboration, M. R.olta et al., *First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Dark energy induced correlation with radio sources*, Astrophys. J. **608** (2004) 10–15, [[astro-ph/0305097](#)].
- [79] P.-S. Corasaniti, T. Giannantonio, and A. Melchiorri, *Constraining dark energy with cross-correlated CMB and large scale structure data*, Phys. Rev. **D71** (2005) 123521, [[astro-ph/0504115](#)].
- [80] N. Padmanabhan, C. M. Hirata, U. Seljak, D. Schlegel, J. Brinkmann, and D. P. Schneider, *Correlating the CMB with luminous red galaxies: The Integrated Sachs-Wolfe effect*, Phys. Rev. **D72** (2005) 043525, [[astro-ph/0410360](#)].
- [81] P. Vielva, E. Martinez-Gonzalez, and M. Tucci, *WMAP and NVSS cross-correlation in wavelet space: ISW detection and dark energy constraints*, Mon. Not. Roy. Astron. Soc. **365** (2006) 891, [[astro-ph/0408252](#)].
- [82] J. D. McEwen, P. Vielva, M. P. Hobson, E. Martinez-Gonzalez, and A. N. Lasenby, *Detection of the integrated Sachs-Wolfe effect and corresponding dark energy constraints made with directional spherical wavelets*, Mon. Not. Roy. Astron. Soc. **376** (2007) 1211–1226, [[astro-ph/0602398](#)].
- [83] T. Giannantonio, R. G. Crittenden, R. C. Nichol, R. Scranton, G. T. Richards, A. D. Myers, R. J. Brunner, A. G. Gray, A. J. Connolly, and D. P. Schneider, *A high redshift detection of the integrated Sachs-Wolfe effect*, Phys. Rev. **D74** (2006) 063520, [[astro-ph/0607572](#)].
- [84] A. Cabre, P. Fosalba, E. Gaztanaga, and M. Manera, *Error analysis in cross-correlation of sky maps: Application to the ISW detection*, Mon. Not. Roy. Astron. Soc. **381** (2007) 1347, [[astro-ph/0701393](#)].
- [85] A. Rassat, K. Land, O. Lahav, and F. B. Abdalla, *Cross-correlation of 2MASS and WMAP3: Implications for the Integrated Sachs-Wolfe effect*, Mon. Not. Roy. Astron. Soc. **377** (2007) 1085–1094, [[astro-ph/0610911](#)].
- [86] T. Giannantonio, R. Scranton, R. G. Crittenden, R. C. Nichol, S. P. Boughn, et al., *Combined analysis of the integrated Sachs-Wolfe effect and cosmological implications*, Phys. Rev. **D77** (2008) 123520, [[arXiv:0801.4380](#)].
- [87] S. Ho, C. Hirata, N. Padmanabhan, U. Seljak, and N. Bahcall, *Correlation of CMB with large-scale structure: I. ISW Tomography and Cosmological Implications*, Phys. Rev. **D78** (2008) 043519, [[arXiv:0801.0642](#)].

- [88] J.-Q. Xia, M. Viel, C. Baccigalupi, and S. Matarrese, *The High Redshift Integrated Sachs-Wolfe Effect*, JCAP **0909** (2009) 003, [[arXiv:0907.4753](#)].
- [89] T. Giannantonio, R. Crittenden, R. Nichol, and A. J. Ross, *The significance of the integrated Sachs-Wolfe effect revisited*, Mon.Not.Roy.Astron.Soc. **426** (2012) 2581–2599, [[arXiv:1209.2125](#)].
- [90] **Planck Collaboration** Collaboration, P. Ade et al., *Planck 2013 results. XIX. The integrated Sachs-Wolfe effect*, Astron.Astrophys. **571** (2014) A19, [[arXiv:1303.5079](#)].
- [91] **Planck Collaboration** Collaboration, P. Ade et al., *Planck 2015 results. XXI. The integrated Sachs-Wolfe effect*, [arXiv:1502.01595](#).
- [92] F.-X. Dupe, A. Rassat, J.-L. Starck, and M. Fadili, *Measuring the Integrated Sachs-Wolfe Effect*, Astron.Astrophys. **534** (2011) 51, [[arXiv:1010.2192](#)].
- [93] W. Hu and R. Scranton, *Measuring Dark Energy Clustering with CMB-Galaxy Correlations*, Phys. Rev. **D70** (2004) 123002, [[astro-ph/0408456](#)].
- [94] D. J. Schwarz, C. J. Copi, D. Huterer, and G. D. Starkman, *CMB Anomalies after Planck*, Class. Quant. Grav. **33** (2016), no. 18 184001, [[arXiv:1510.07929](#)].
- [95] C. Francis and J. Peacock, *An estimate of the local ISW signal, and its impact on CMB anomalies*, Mon.Not.Roy.Astron.Soc. **406** (2010) 14, [[arXiv:0909.2495](#)].
- [96] A. Rassat, J. L. Starck, and F. X. Dupe, *Removal of two large scale Cosmic Microwave Background anomalies after subtraction of the Integrated Sachs Wolfe effect*, Astron.Astrophys. **557** (2013) A32, [[arXiv:1303.4727](#)].
- [97] **Planck** Collaboration, P. A. R. Ade et al., *Planck 2015 results. XVII. Constraints on primordial non-Gaussianity*, [arXiv:1502.01592](#).
- [98] J. Kim, A. Rotti, and E. Komatsu, *Removing the ISW-lensing bias from the local-form primordial non-Gaussianity estimation*, JCAP **1304** (2013) 021, [[arXiv:1302.5799](#)].
- [99] L. Bonavera, R. B. Barreiro, A. Marcos-Caballero, and P. Vielva, *On the recovery of ISW fluctuations using large-scale structure tracers and CMB temperature and polarization anisotropies*, Mon. Not. Roy. Astron. Soc. **459** (2016), no. 1 657–672, [[arXiv:1602.05893](#)].
- [100] N. Afshordi, *Integrated Sachs-Wolfe effect in cross - correlation: The Observer’s manual*, Phys. Rev. **D70** (2004) 083536, [[astro-ph/0401166](#)].
- [101] M. LoVerde and N. Afshordi, *Extended Limber Approximation*, Phys. Rev. **D78** (2008) 123506, [[arXiv:0809.5112](#)].

- [102] J. Lesgourgues, *The Cosmic Linear Anisotropy Solving System (CLASS) I: Overview*, [arXiv:1104.2932](#).
- [103] **EUCLID** Collaboration, R. Laureijs et al., *Euclid Definition Study Report*, [arXiv:1110.3193](#).
- [104] M. Douspis, P. G. Castro, C. Caprini, and N. Aghanim, *Optimising large galaxy surveys for ISW detection*, *Astron. Astrophys.* **485** (2008) 395, [[arXiv:0802.0983](#)].
- [105] N. Martinet, J. G. Bartlett, A. Kiessling, and B. Sartoris, *Constraining cosmology with shear peak statistics: tomographic analysis*, *Astron. Astrophys.* **581** (2015) A101, [[arXiv:1506.02192](#)].
- [106] **DES** Collaboration, M. R. Becker et al., *Cosmic Shear Measurements with DES Science Verification Data*, [arXiv:1507.05598](#).
- [107] A. Cimatti, R. Laureijs, B. Leibundgut, S. Lilly, R. Nichol, A. Refregier, P. Rosati, M. Steinmetz, N. Thatte, and E. Valentijn, *Euclid Assessment Study Report for the ESA Cosmic Visions*, [arXiv:0912.0914](#).
- [108] W. Hu, *Power spectrum tomography with weak lensing*, *Astrophys. J.* **522** (1999) L21–L24, [[astro-ph/9904153](#)].
- [109] Z.-M. Ma, W. Hu, and D. Huterer, *Effect of photometric redshift uncertainties on weak lensing tomography*, *Astrophys. J.* **636** (2005) 21–29, [[astro-ph/0506614](#)].
- [110] G. Bernstein and D. Huterer, *Catastrophic photometric redshift errors: weak lensing survey requirements*, *Mon. Not. Roy. Astron. Soc.* **401** (2010) 1399, [[arXiv:0902.2782](#)].
- [111] A. P. Hearin, A. R. Zentner, Z. Ma, and D. Huterer, *A General Study of the Influence of Catastrophic Photometric Redshift Errors on Cosmology with Cosmic Shear Tomography*, *Astrophys. J.* **720** (2010) 1351–1369, [[arXiv:1002.3383](#)].
- [112] H. Hildebrandt et al., *CFHTLenS: Improving the quality of photometric redshifts with precision photometry*, *Mon. Not. Roy. Astron. Soc.* **421** (2012) 2355, [[arXiv:1111.4434](#)].
- [113] A. R. Pullen and C. M. Hirata, *Systematic effects in large-scale angular power spectra of photometric quasars and implications for constraining primordial nongaussianity*, *Publications of the Astronomical Society of the Pacific*, Volume 125, Issue 928, pp. **705-718** (2013) [[arXiv:1212.4500](#)].
- [114] S. Ho, A. Cuesta, H.-J. Seo, R. de Putter, A. J. Ross, et al., *Clustering of Sloan Digital Sky Survey III Photometric Luminous Galaxies: The Measurement, Systematics and Cosmological Implications*, *Astrophys. J.* **761** (2012) 14, [[arXiv:1201.2137](#)].

- [115] S. Ho, N. Agarwal, A. D. Myers, R. Lyons, A. Disbrow, et al., *Sloan Digital Sky Survey III Photometric Quasar Clustering: Probing the Initial Conditions of the Universe using the Largest Volume*, 1311.2597 (2013) [[arXiv:1311.2597](#)].
- [116] N. Agarwal, S. Ho, and S. Shandera, *Constraining the initial conditions of the Universe using large scale structure*, JCAP **1402** (2014) 038, [[arXiv:1311.2606](#)].
- [117] T. Giannantonio, A. J. Ross, W. J. Percival, R. Crittenden, D. Bacher, et al., *Improved Primordial Non-Gaussianity Constraints from Measurements of Galaxy Clustering and the Integrated Sachs-Wolfe Effect*, Phys.Rev. **D89** (2014) 023511, [[arXiv:1303.1349](#)].
- [118] N. Agarwal, S. Ho, A. D. Myers, H.-J. Seo, A. J. Ross, et al., *Characterizing unknown systematics in large scale structure surveys*, JCAP **1404** (2014) 007, [[arXiv:1309.2954](#)].
- [119] D. Huterer, C. E. Cunha, and W. Fang, *Calibration errors unleashed: effects on cosmological parameters and requirements for large-scale structure surveys*, Mon.Not.Roy.Astron.Soc. **432** (2013) 2945, [[arXiv:1211.1015](#)].
- [120] B. Leistedt, H. V. Peiris, D. J. Mortlock, A. Benoit-Lévy, and A. Pontzen, *Estimating the large-scale angular power spectrum in the presence of systematics: a case study of Sloan Digital Sky Survey quasars*, Mon.Not.Roy.Astron.Soc. **435** (2013) 1857–1873, [[arXiv:1306.0005](#)].
- [121] B. Leistedt and H. V. Peiris, *Exploiting the full potential of photometric quasar surveys: Optimal power spectra through blind mitigation of systematics*, Mon.Not.Roy.Astron.Soc. **444** (2014) 2, [[arXiv:1404.6530](#)].
- [122] D. L. Shafer and D. Huterer, *Multiplicative errors in the galaxy power spectrum: self-calibration of unknown photometric systematics for precision cosmology*, Mon. Not. Roy. Astron. Soc. **447** (2015) 2961, [[arXiv:1410.0035](#)].
- [123] S. Jouvel et al., *Designing Future Dark Energy Space Mission: I. Building Realistic Galaxy Spectro-Photometric Catalogs and their first applications*, Astron. Astrophys. **504** (2009) 359, [[arXiv:0902.0625](#)].
- [124] E. S. Rykoff, E. Rozo, and R. Keisler, *Assessing Galaxy Limiting Magnitudes in Large Optical Surveys*, Submitted to: Astrophysical Journal (2015) [[arXiv:1509.00870](#)].
- [125] N. Padmanabhan et al., *An Improved Photometric Calibration of the Sloan Digital Sky Survey Imaging Data*, Astrophys. J. **674** (2008) 1217–1233, [[astro-ph/0703454](#)].
- [126] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, *Large-scale alignments from WMAP and Planck*, Mon. Not. Roy. Astron. Soc. **449** (2015), no. 4 3458–3470, [[arXiv:1311.4562](#)].

- [127] A. de Oliveira-Costa, M. Tegmark, M. Zaldarriaga, and A. Hamilton, *The Significance of the largest scale CMB fluctuations in WMAP*, Phys. Rev. **D69** (2004) 063516, [[astro-ph/0307282](#)].
- [128] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, *On the large-angle anomalies of the microwave sky*, Mon. Not. Roy. Astron. Soc. **367** (2006) 79–102, [[astro-ph/0508047](#)].
- [129] **Planck** Collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, Astron. Astrophys. **594** (2016) A13, [[arXiv:1502.01589](#)].
- [130] C. Copi, D. Huterer, D. Schwarz, and G. Starkman, *The Uncorrelated Universe: Statistical Anisotropy and the Vanishing Angular Correlation Function in WMAP Years 1-3*, Phys. Rev. **D75** (2007) 023507, [[astro-ph/0605135](#)].
- [131] H. K. Eriksen, F. K. Hansen, A. J. Banday, K. M. Gorski, and P. B. Lilje, *Asymmetries in the Cosmic Microwave Background anisotropy field*, Astrophys. J. **605** (2004) 14–20, [[astro-ph/0307507](#)]. [Erratum: Astrophys. J. 609,1198(2004)].
- [132] K. Land and J. Magueijo, *Is the Universe odd?*, Phys. Rev. **D72** (2005) 101302, [[astro-ph/0507289](#)].
- [133] P. Vielva, E. Martinez-Gonzalez, R. B. Barreiro, J. L. Sanz, and L. Cayon, *Detection of non-Gaussianity in the WMAP 1 - year data using spherical wavelets*, Astrophys. J. **609** (2004) 22–34, [[astro-ph/0310273](#)].
- [134] **Planck** Collaboration, P. A. R. Ade et al., *Planck 2015 results. XVI. Isotropy and statistics of the CMB*, Astron. Astrophys. **594** (2016) A16, [[arXiv:1506.07135](#)].
- [135] C. L. Bennett et al., *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Are There Cosmic Microwave Background Anomalies?*, Astrophys. J. Suppl. **192** (2011) 17, [[arXiv:1001.4758](#)].
- [136] C. R. Contaldi, M. Peloso, L. Kofman, and A. D. Linde, *Suppressing the lower multipoles in the CMB anisotropies*, JCAP **0307** (2003) 002, [[astro-ph/0303636](#)].
- [137] C. Gordon, W. Hu, D. Huterer, and T. M. Crawford, *Spontaneous isotropy breaking: a mechanism for cmb multipole alignments*, Phys. Rev. **D72** (2005) 103002, [[astro-ph/0509301](#)].
- [138] E. F. Bunn and A. Bourdon, *Contamination cannot explain the lack of large-scale power in the cosmic microwave background radiation*, Phys. Rev. **D78** (2008) 123509, [[arXiv:0808.0341](#)].
- [139] S. Adhikari, S. Shandera, and A. L. Erickcek, *Large-scale anomalies in the cosmic microwave background as signatures of non-gaussianity*, Physical Review D **93** (jan, 2016) 023524, [[arXiv:1508.06489](#)].

- [140] C. Dvorkin, H. V. Peiris, and W. Hu, *Testable polarization predictions for models of CMB isotropy anomalies*, Phys. Rev. **D77** (2008) 063008, [[arXiv:0711.2321](#)].
- [141] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, *Large-Angle CMB Suppression and Polarization Predictions*, Mon. Not. Roy. Astron. Soc. **434** (2013) 3590–3596, [[arXiv:1303.4786](#)].
- [142] A. Yoho, S. Aiola, C. J. Copi, A. Kosowsky, and G. D. Starkman, *Microwave Background Polarization as a Probe of Large-Angle Correlations*, Phys. Rev. **D91** (2015), no. 12 123504, [[arXiv:1503.05928](#)].
- [143] M. O’Dwyer, C. J. Copi, L. Knox, and G. D. Starkman, *CMB- S_4 and the Hemispherical Variance Anomaly*, Mon. Not. Roy. Astron. Soc. **470** (2017), no. 1 372–378, [[arXiv:1608.02234](#)].
- [144] D. Contreras, J. P. Zibin, D. Scott, A. J. Banday, and K. M. Grski, *Testing physical models for dipolar asymmetry with CMB polarization*, [arXiv:1704.03143](#).
- [145] A. Yoho, C. J. Copi, G. D. Starkman, and A. Kosowsky, *Probing Large-Angle Correlations with the Microwave Background Temperature and Lensing Cross Correlation*, Mon. Not. Roy. Astron. Soc. **442** (2014), no. 3 2392–2397, [[arXiv:1310.7603](#)].
- [146] C. Gibelyou, D. Huterer, and W. Fang, *Detectability of large-scale power suppression in the galaxy distribution*, Phys. Rev. **D82** (2010) 123009, [[arXiv:1007.0757](#)].
- [147] D. Sarkar, D. Huterer, C. J. Copi, G. D. Starkman, and D. J. Schwarz, *Missing Power vs low- l Alignments in the Cosmic Microwave Background: No Correlation in the Standard Cosmological Model*, Astropart. Phys. **34** (2011) 591–594, [[arXiv:1004.3784](#)].
- [148] A. Rakic and D. J. Schwarz, *Correlating anomalies of the microwave sky: The Good, the Evil and the Axis*, Phys. Rev. **D75** (2007) 103002, [[astro-ph/0703266](#)].
- [149] L. Polastri, A. Gruppuso, and P. Natoli, *CMB low multipole alignments in the Λ CDM and Dipolar models*, JCAP **1504** (2015) 018, [[arXiv:1503.01611](#)].
- [150] J. Kim and P. Naselsky, *Lack of angular correlation and odd-parity preference in CMB data*, Astrophys. J. **739** (2011) 79, [[arXiv:1011.0377](#)].
- [151] R. Knight and L. Knox, *The low level of correlation observed in the CMB sky at large angular scales and the low quadrupole variance*, [arXiv:1705.01178](#).
- [152] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, *No large-angle correlations on the non-Galactic microwave sky*, Mon. Not. Roy. Astron. Soc. **399** (2009) 295–303, [[arXiv:0808.3767](#)].

- [153] A. Hajian, *Analysis of the apparent lack of power in the cosmic microwave background anisotropy at large angular scales*, [astro-ph/0702723](#).
- [154] **Planck** Collaboration, P. A. R. Ade et al., *Planck 2015 results. XII. Full Focal Plane simulations*, *Astron. Astrophys.* **594** (2016) A12, [[arXiv:1509.06348](#)].
- [155] **Planck** Collaboration, R. Adam et al., *Planck 2015 results. IX. Diffuse component separation: CMB maps*, *Astron. Astrophys.* **594** (2016) A9, [[arXiv:1502.05956](#)].
- [156] **Planck** Collaboration, P. A. R. Ade et al., *Planck 2013 results. XXIII. Isotropy and statistics of the CMB*, *Astron. Astrophys.* **571** (2014) A23, [[arXiv:1303.5083](#)].
- [157] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, *Lack of large-angle TT correlations persists in WMAP and Planck*, *Mon. Not. Roy. Astron. Soc.* **451** (2015), no. 3 2978–2985, [[arXiv:1310.3831](#)].
- [158] M. Kamionkowski and L. Knox, *Aspects of the cosmic microwave background dipole*, *Phys. Rev.* **D67** (2003) 063001, [[astro-ph/0210165](#)].
- [159] G. Chon, A. Challinor, S. Prunet, E. Hivon, and I. Szapudi, *Fast estimation of polarization power spectra using correlation functions*, *Mon. Not. Roy. Astron. Soc.* **350** (2004) 914, [[astro-ph/0303414](#)].
- [160] **Planck** Collaboration, N. Aghanim et al., *Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters*, *Astron. Astrophys.* **594** (2016) A11, [[arXiv:1507.02704](#)].
- [161] A. Gruppuso, A. De Rosa, P. Cabella, F. Paci, F. Finelli, P. Natoli, G. de Gasperis, and N. Mandolesi, *New estimates of the CMB angular power spectra from the WMAP 5 yrs low resolution data*, *Mon. Not. Roy. Astron. Soc.* **400** (2009) 463–469, [[arXiv:0904.0789](#)].
- [162] P. K. Aluri, J. P. Ralston, and A. Weltman, *Alignments of parity even/odd-only multipoles in CMB*, [arXiv:1703.07070](#).
- [163] P. Vielva, *A Comprehensive overview of the Cold Spot*, *Adv. Astron.* **2010** (2010) 592094, [[arXiv:1008.3051](#)].
- [164] **WMAP** Collaboration, D. N. Spergel et al., *First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Determination of cosmological parameters*, *Astrophys. J. Suppl.* **148** (2003) 175–194, [[astro-ph/0302209](#)].
- [165] G. Hinshaw, A. J. Banday, C. L. Bennett, K. M. Gorski, A. Kogut, C. H. Lineweaver, G. F. Smoot, and E. L. Wright, *2-point correlations in the COBE DMR 4-year anisotropy maps*, *Astrophys. J.* **464** (1996) L25–L28, [[astro-ph/9601061](#)].

- [166] A. Gruppuso, *Two-point correlation function of Wilkinson Microwave Anisotropy Probe 9-yr data*, Mon. Not. Roy. Astron. Soc. **437** (2014), no. 3 2076–2082, [[arXiv:1310.2822](#)].
- [167] C. L. Bennett, A. Banday, K. M. Gorski, G. Hinshaw, P. Jackson, P. Keegstra, A. Kogut, G. F. Smoot, D. T. Wilkinson, and E. L. Wright, *Four year COBE DMR cosmic microwave background observations: Maps and basic results*, Astrophys. J. **464** (1996) L1–L4, [[astro-ph/9601067](#)].
- [168] **WMAP** Collaboration, C. L. Bennett et al., *First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: Preliminary maps and basic results*, Astrophys. J. Suppl. **148** (2003) 1–27, [[astro-ph/0302207](#)].
- [169] I. J. O’Dwyer, H. K. Eriksen, B. D. Wandelt, J. B. Jewell, D. L. Larson, K. M. Gorski, A. J. Banday, S. Levin, and P. B. Lilje, *Bayesian power spectrum analysis of the first-year WMAP data*, Astrophys. J. **617** (2004) L99–102, [[astro-ph/0407027](#)].
- [170] M. Chu, H. K. Eriksen, L. Knox, K. M. Gorski, J. B. Jewell, D. L. Larson, I. J. O’Dwyer, and B. D. Wandelt, *Cosmological parameter constraints as derived from the Wilkinson Microwave Anisotropy Probe data via Gibbs sampling and the Blackwell-Rao estimator*, Phys. Rev. **D71** (2005) 103002, [[astro-ph/0411737](#)].
- [171] A. Slosar and U. Seljak, *Assessing the effects of foregrounds and sky removal in WMAP*, Phys. Rev. **D70** (2004) 083002, [[astro-ph/0404567](#)].
- [172] M. Cruz, P. Vielva, E. Martinez-Gonzalez, and R. B. Barreiro, *Anomalous variance in the WMAP data and Galactic Foreground residuals*, Mon. Not. Roy. Astron. Soc. **412** (2011) 2383, [[arXiv:1005.1264](#)].
- [173] C. Monteserin, R. B. B. Barreiro, P. Vielva, E. Martinez-Gonzalez, M. P. Hobson, and A. N. Lasenby, *A low CMB variance in the WMAP data*, Mon. Not. Roy. Astron. Soc. **387** (2008) 209–219, [[arXiv:0706.4289](#)].
- [174] J. Kim and P. Naselsky, *Anomalous parity asymmetry of WMAP power spectrum data at low multipoles: is it cosmological or systematics?*, Phys. Rev. **D82** (2010) 063002, [[arXiv:1002.0148](#)].
- [175] J. Kim and P. Naselsky, *Anomalous parity asymmetry of the Wilkinson Microwave Anisotropy Probe power spectrum data at low multipoles*, Astrophys. J. **714** (2010) L265–L267, [[arXiv:1001.4613](#)].
- [176] A. Gruppuso, F. Finelli, P. Natoli, F. Paci, P. Cabella, A. De Rosa, and N. Mandolesi, *New constraints on parity symmetry from a re-analysis of the wmap-7 low-resolution power spectra*, Monthly Notices of the Royal Astronomical Society **411** (2011), no. 3 1445–1452.
- [177] D. Molinari. personal communication, Jan, 2018.

- [178] C. J. Copi, D. Huterer, and G. D. Starkman, *Multipole vectors - A New representation of the CMB sky and evidence for statistical anisotropy or non-Gaussianity at $2 \leq l \leq 8$* , Phys. Rev. **D70** (2004) 043515, [[astro-ph/0310511](#)].
- [179] D. J. Schwarz, G. D. Starkman, D. Huterer, and C. J. Copi, *Is the low- l microwave background cosmic?*, Phys. Rev. Lett. **93** (2004) 221301, [[astro-ph/0403353](#)].
- [180] **Planck** Collaboration, P. A. R. Ade et al., *Planck 2013 results. I. Overview of products and scientific results*, Astron. Astrophys. **571** (2014) A1, [[arXiv:1303.5062](#)].
- [181] F. K. Hansen, P. Cabella, D. Marinucci, and N. Vittorio, *Asymmetries in the local curvature of the WMAP data*, Astrophys. J. **607** (2004) L67–L70, [[astro-ph/0402396](#)].
- [182] H. K. Eriksen, A. J. Banday, K. M. Gorski, F. K. Hansen, and P. B. Lilje, *Hemispherical power asymmetry in the three-year Wilkinson Microwave Anisotropy Probe sky maps*, Astrophys. J. **660** (2007) L81–L84, [[astro-ph/0701089](#)].
- [183] J. Hoftuft, H. K. Eriksen, A. J. Banday, K. M. Gorski, F. K. Hansen, and P. B. Lilje, *Increasing evidence for hemispherical power asymmetry in the five-year WMAP data*, Astrophys. J. **699** (2009) 985–989, [[arXiv:0903.1229](#)].
- [184] Y. Akrami, Y. Fantaye, A. Shafieloo, H. K. Eriksen, F. K. Hansen, A. J. Banday, and K. M. Gorski, *Power asymmetry in WMAP and Planck temperature sky maps as measured by a local variance estimator*, Astrophys. J. **784** (2014) L42, [[arXiv:1402.0870](#)].
- [185] S. Adhikari, *Local variance asymmetries in Planck temperature anisotropy maps*, Mon. Not. Roy. Astron. Soc. **446** (2015), no. 4 4232–4238, [[arXiv:1408.5396](#)].
- [186] **Planck** Collaboration, N. Aghanim et al., *Planck 2013 results. XXVII. Doppler boosting of the CMB: Eppur si muove*, Astron. Astrophys. **571** (2014) A27, [[arXiv:1303.5087](#)].
- [187] M. Quartin and A. Notari, *On the significance of power asymmetries in Planck CMB data at all scales*, JCAP **1501** (2015), no. 01 008, [[arXiv:1408.5792](#)].
- [188] D. Scott, *The Standard Model of Cosmology: A Skeptic's Guide*, 2018. [[arXiv:1804.01318](#)].
- [189] K. Arisaka et al., *Improved upper limit on the branching ratio $B(K0L \rightarrow \mu^+ e^-)$* , Phys. Rev. Lett. **70** (1993) 1049–1052.
- [190] J. R. Klein and A. Roodman, *Blind analysis in nuclear and particle physics*, Ann. Rev. Nucl. Part. Sci. **55** (2005) 141–163.

- [191] **Supernova Cosmology Project** Collaboration, A. J. Conley et al., *Measurement of $\Omega(m)$, $\Omega(\lambda)$ from a blind analysis of Type Ia supernovae with CMAGIC: Using color information to verify the acceleration of the Universe*, *Astrophys. J.* **644** (2006) 1–20, [[astro-ph/0602411](#)].
- [192] **Supernova Cosmology Project** Collaboration, M. Kowalski et al., *Improved Cosmological Constraints from New, Old and Combined Supernova Datasets*, *Astrophys. J.* **686** (2008) 749–778, [[arXiv:0804.4142](#)].
- [193] N. Suzuki et al., *The Hubble Space Telescope Cluster Supernova Survey: V. Improving the Dark Energy Constraints Above $z \gtrsim 1$ and Building an Early-Type-Hosted Supernova Sample*, *Astrophys. J.* **746** (2012) 85, [[arXiv:1105.3470](#)].
- [194] **Supernova Cosmology Project** Collaboration, D. Rubin et al., *Unity: Confronting Supernova Cosmologies Statistical and Systematic Uncertainties in a Unified Bayesian Framework*, *Astrophys. J.* **813** (2015), no. 2 137, [[arXiv:1507.01602](#)].
- [195] S. H. Suyu et al., *Two accurate time-delay distances from strong lensing: Implications for cosmology*, *Astrophys. J.* **766** (2013) 70, [[arXiv:1208.6010](#)].
- [196] S. H. Suyu et al., *H0LiCOW I. H0 Lenses in COSMOSGRAB's Wellspring: program overview*, *Mon. Not. Roy. Astron. Soc.* **468** (2017), no. 3 2590–2604, [[arXiv:1607.00017](#)].
- [197] C. Heymans et al., *CFHTLenS: The Canada-France-Hawaii Telescope Lensing Survey*, *Mon. Not. Roy. Astron. Soc.* **427** (2012) 146, [[arXiv:1210.0032](#)].
- [198] A. von der Linden et al., *Weighing the Giants I. Weak-lensing masses for 51 massive galaxy clusters: project overview, data analysis methods and cluster images*, *Mon. Not. Roy. Astron. Soc.* **439** (2014), no. 1 2–27, [[arXiv:1208.0597](#)].
- [199] K. Kuijken et al., *Gravitational Lensing Analysis of the Kilo Degree Survey*, *Mon. Not. Roy. Astron. Soc.* **454** (2015), no. 4 3500–3532, [[arXiv:1507.00738](#)].
- [200] C. Blake et al., *RCSLenS: Testing gravitational physics through the cross-correlation of weak lensing and large-scale structure*, *Mon. Not. Roy. Astron. Soc.* **456** (2016), no. 3 2806–2828, [[arXiv:1507.03086](#)].
- [201] **DES** Collaboration, A. Drlica-Wagner et al., *Dark Energy Survey Year 1 Results: The Photometric Data Set for Cosmology*, *Astrophys. J. Suppl.* **235** (2018), no. 2 33, [[arXiv:1708.01531](#)].
- [202] **DES** Collaboration, J. Prat et al., *Dark Energy Survey Year 1 Results: Galaxy-Galaxy Lensing*, [[arXiv:1708.01537](#)].
- [203] **DES** Collaboration, J. Elvin-Poole et al., *Dark Energy Survey Year 1 Results: Galaxy clustering for combined probes*, [[arXiv:1708.01536](#)].

- [204] **DES** Collaboration, J. Zuntz et al., *Dark Energy Survey Year 1 Results: Weak Lensing Shape Catalogues*, Submitted to: Mon. Not. Roy. Astron. Soc. (2017) [[arXiv:1708.01533](#)].
- [205] E. Krause and T. Eifler, *cosmolike cosmological likelihood analyses for photometric galaxy surveys*, Mon. Not. Roy. Astron. Soc. **470** (2017), no. 2 2100–2112, [[arXiv:1601.05779](#)].
- [206] B. Joachimi, P. Schneider, and T. Eifler, *Analysis of two-point statistics of cosmic shear. 3. Covariances of shear measures made easy*, Astron. Astrophys. (2007) [[arXiv:0708.0387](#)]. [Astron. Astrophys.477,43(2008)].
- [207] T. Eifler, P. Schneider, and J. Hartlap, *Dependence of cosmic shear covariances on cosmology - Impact on parameter estimation*, Astron. Astrophys. **502** (2009) 721–731, [[arXiv:0810.4254](#)].
- [208] S. Bridle and L. King, *Dark energy constraints from cosmic shear power spectra: impact of intrinsic alignments on photometric redshift requirements*, New J. Phys. **9** (2007) 444, [[arXiv:0705.0166](#)].
- [209] M. Kilbinger et al., *Dark energy constraints and correlations with systematics from CFHTLS weak lensing, SNLS supernovae Ia and WMAP5*, Astron. Astrophys. **497** (2009) 677, [[arXiv:0810.5129](#)].
- [210] J. A. Nelder and R. Mead, *A Simplex Method for Function Minimization*, Comput. J. **7** (1965) 308–313.
- [211] F. Feroz and M. P. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, Mon. Not. Roy. Astron. Soc. **384** (2008) 449, [[arXiv:0704.3704](#)].
- [212] F. Feroz, M. P. Hobson, and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, Mon. Not. Roy. Astron. Soc. **398** (2009) 1601–1614, [[arXiv:0809.3437](#)].
- [213] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, [arXiv:1306.2144](#).
- [214] J. Zuntz. Comments in CosmoSIS initialization file template, 2017.
- [215] N. Afshordi, Y.-S. Loh, and M. A. Strauss, *Cross - correlation of the Cosmic Microwave Background with the 2MASS galaxy survey: Signatures of dark energy, hot gas, and point sources*, Phys. Rev. **D69** (2004) 083524, [[astro-ph/0308260](#)].