UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title

Towards accurate cosmology from CMB secondaries and large-scale structure surveys

Permalink

https://escholarship.org/uc/item/5zm4970p

ISBN

9798293892945

Author

Sailer, Noah

Publication Date

2025-08-01

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Towards accurate cosmology from CMB secondaries and large-scale structure surveys

by

Noah Sailer

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Martin White, Chair
Senior Scientist Simone Ferraro
Professor Adrian Lee
Professor Daniel Weisz

Summer 2025

Towards accurate cosmology from CMB secondaries and large-scale structure surveys

Copyright 2025
by
Noah Sailer

Abstract

Towards accurate cosmology from CMB secondaries and large-scale structure surveys

by

Noah Sailer

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Martin White, Chair

The past decade has seen the rise of high-precision cosmological observations accompanied by emerging tensions among various datasets when interpreted within the standard cosmological model (ΛCDM). We investigate how joint analyses of secondary anisotropies in the cosmic microwave background (CMB) and large-scale structure (LSS) surveys can be leveraged to improve the robustness of cosmological inference.

In Part I, we develop novel methods for mitigating foreground contamination in CMB lensing reconstruction. Extragalactic foregrounds such as the thermal Sunyaev-Zel'dovich (tSZ) effect and the cosmic infrared background (CIB) can bias lensing estimators, particularly on small scales. We construct bias-hardened quadratic estimators, both in temperature and polarization, that null specific foreground profiles with minimal degradation in signal-to-noise. We further generalize temperature-based shear estimators to polarization, finding that the shear-only estimator is a special case of a more general class of large-lens-suppressed estimators. In tandem, we explore optimal multifrequency methods, quantifying the tradeoffs between noise and bias in the presence of extragalactic foregrounds.

Part II applies these tools to current datasets, addressing a subset of the aforementioned tensions. We perform a cross-correlation analysis of the latest CMB lensing maps from Planck and ACT with the DESI legacy imaging data, yielding tomographic constraints on linear structure growth. We find overall consistency with ΛCDM, both in the overall amplitude and time-evolution of structure growth. We then present methods to mitigate lensing biases in the mean patchy screening signal around DESI galaxies and show that, when these biases are taken into account, patchy screening is consistent with the baryon distribution inferred from the kinetic SZ effect. Finally, we quantify the impact of the optical depth τ on parameter inference from the CMB, arguing that a higher value of τ would alleviate the preference for beyond-ΛCDM physics by current CMB and BAO data.

Part III projects forward to the next generation of large-scale structure surveys. We demon-

strate that combining high-redshift galaxy surveys with CMB data from Simons Observatory or a stage-4 CMB survey will enable percent-level constraints on the neutrino mass sum, curvature, light relics, the expansion and growth histories, (early) dark energy, modified gravity and primordial non-Gaussianity. A new method is proposed for using 21 cm clustering to measure τ independently of large-scale CMB measurements. It is shown that future 21 cm interferometers have the potential to measure τ to a higher precision than a cosmic-variance limited CMB survey while at the same time mitigating biases arising from the miss-modeling of the small-scale astrophysics relevant for reionization.

Altogether, this work advances the extraction of cosmological information from secondary CMB effects and LSS, addresses key theoretical and observational systematics, and lays a foundation for robust analyses from upcoming datasets.

To my high school teacher, Chris McGrath, who taught me to “embrace the struggle.”

Contents

Contents ii

List of Figures vi

List of Tables xxv

1 A review and an overview 1
- 1.1 Our current best cosmological model . . . 2
- 1.2 Secondary anisotropies in the CMB . . . 7
- 1.3 Observed tracers of the large-scale structure . . . 11
- 1.4 Cosmological tensions . . . 13
- 1.5 Scope and outline of this thesis . . . 15

I Foreground mitigation in CMB lensing measurements 17

2 Geometric foreground mitigation I: tSZ profile-hardening for temperature based estimators 18
- 2.1 Introduction . . . 19
- 2.2 Foreground-hardened quadratic estimators . . . 20
- 2.3 Source QE and source-hardened lensing QE . . . 24
- 2.4 Biases to CMB lensing & foreground hardening . . . 26
- 2.5 Discussion and conclusions . . . 34

3 Geometric foreground mitigation II: bias-hardened and shear-only estimators in polarization 36
- 3.1 Introduction . . . 37
- 3.2 Linear responses and quadratic estimators . . . 38
- 3.3 Bias-hardened estimators . . . 41
- 3.4 Large-lens suppressed estimators . . . 47
- 3.5 Noise properties . . . 50
- 3.6 Bias reduction . . . 53
- 3.7 Discussion and conclusions . . . 57

4 Optimal multifrequency weighting **58**
4.1 Introduction . . . 59
4.2 CMB lensing quadratic estimators . . . 59
4.3 Internal linear combination (ILC) methods . . . 60
4.4 Reconstruction noise . . . 62
4.5 Bias to lensing reconstruction from simulations . . . 65
4.6 Minimizing lensing bias and noise . . . 67
4.7 Conclusions . . . 68

II Tenuous tensions **70**

5 Tomographic structure growth I: cross-correlation of DESI Luminous Red Galaxies with CMB lensing **71**
5.1 Introduction . . . 72
5.2 Data . . . 75
5.3 Power spectra, window functions and covariances . . . 80
5.4 Modeling . . . 82
5.5 Likelihood and pipeline checks . . . 89
5.6 LRG systematics tests . . . 96
5.7 Results . . . 106
5.8 Comparison with previous weak lensing analyses . . . 117
5.9 Discussion and conclusions . . . 120

6 Tomographic structure growth II: cross-correlation of the DESI Bright Galaxy Survey with CMB lensing and galaxy magnification **125**
6.1 Introduction . . . 126
6.2 Data . . . 127
6.3 Likelihoods . . . 129
6.4 Results . . . 133
6.5 Discussion and conclusions . . . 134

7 Baryon distribution I: lens-hardened patchy screening estimators **137**
7.1 Introduction . . . 137
7.2 Patchy screening estimators . . . 138
7.3 Lens-hardened stacked estimators . . . 140
7.4 Comparison with simulations . . . 145
7.5 Comments on ACT × unWISE . . . 146
7.6 Discussion and conclusions . . . 148

8 Baryon distribution II: a self-consistent comparison of the kSZ effect and patchy screening **150**

8.1 Introduction . . . 151
8.2 Data . . . 153
8.3 Methods . . . 154
8.4 Results . . . 158
8.5 Discussion and conclusions . . . 164

9 Dispuτable: the high cost of a low optical depth 166
9.1 Introduction . . . 167
9.2 Data . . . 168
9.3 The optical depth's impact on cosmological 'tensions' . . . 169
9.4 Inferring the optical depth without reference to large scale CMB . . . 171
9.5 Discussion and conclusions . . . 172

III Prospects for future surveys 175

10 Fundamental physics from high redshift large-scale structure surveys 176
10.1 Introduction . . . 176
10.2 Surveys and probes . . . 177
10.3 Forecasting formalism . . . 186
10.4 Results . . . 198
10.5 Observational challenges . . . 222
10.6 Conclusions . . . 225

11 Measuring the optical depth from the epoch of reionization 21cm signal 228
11.1 Introduction . . . 229
11.2 Extracting the optical depth from 21cm clustering . . . 232
11.3 The 21cm clustering signal in Lagrangian perturbation theory . . . 235
11.4 Forecasts . . . 237
11.5 Discussion and conclusions . . . 245

A Geometric foreground mitigation I 249
A.1 Foreground power spectra . . . 249
A.2 Hybrid estimator . . . 249
A.3 Derivation of Eq. (2.20) . . . 250
A.4 tSZ-like profile . . . 252
A.5 Noise price of bias-hardening increases with profile size . . . 252
A.6 Noise and cross-correlation checks . . . 255

B Geometric foreground mitigation II 257
B.1 Expansions in the large-lens limit . . . 257
B.2 Explicit expressions for the global minimum variance weights . . . 259
B.3 Cleaning combined estimators . . . 260

B.4 Bias hardening as constrained minimization . . . 262
B.5 Derivation of polarization-based LLS estimators . . . 264
B.6 Worked example: explicitly hardening the TE estimator . . . 265
B.7 Expected biases to Simons Observatory . . . 266

C Optimal multifrequency weighting **268**
C.1 Cancellations . . . 268
C.2 Proof that ILC minimizes the lensing noise for all QEs . . . 269
C.3 Attempts at numerically minimizing the loss functions: approximations . . . 270
C.4 Dependence of the lensing biases on the masking . . . 272
C.5 Dependence of the lensing biases on $\ell_{\mathrm{max},T}$ . . . 272

D Tomographic structure growth I **274**
D.1 Data and code availability . . . 274
D.2 Signal to noise . . . 274
D.3 Parameter-based comparison of analysis choices . . . 275
D.4 Analytic marginalization and Bayesian PTEs . . . 277
D.5 Supplemental constraints . . . 279

E Baryon distribution I **281**
E.1 Response functions of more general stacked estimators . . . 281
E.2 Fourier transform of the signed and thresholded estimator . . . 284
E.3 Mean field bias . . . 285

F Baryon distribution II **287**
F.1 Simulation test of patchy screening and the effect of CMB lensing . . . 287
F.2 Comparison of the signal-to-noise . . . 288

G Cosmology at high redshift **290**
G.1 Primordial figure of merit (N_{modes}) . . . 290
G.2 Covariance approximations . . . 292
G.3 BAO forecasting . . . 293
G.4 Lensing auto- vs. cross-correlation . . . 295
G.5 Empirical checks . . . 297
G.6 Python implementation . . . 298

Bibliography **300**

List of Figures

1.1 *Top left:* Mollweide projection of intensity fluctuations in the cosmic microwave background (CMB) as measured by the Planck satellite [473]. *Bottom left:* Mollweide projection of the number of DESI luminous red galaxies per square degree [640]. Both the CMB intensity and observed galaxy counts are nearly isotropic on large scales, in agreement with the Copernican principle. *Right:* The mean intensity of the CMB as a function of wavelength as measured by the COBE FIRAS instrument [234]. The discovery that the CMB follows a near-perfect blackbody distribution was crucial in establishing the hot big bang model, in which the early Universe was in thermal equilibrium. Figure from `https://www.astro.ucla.edu/~wright/CMB.html`. 3

1.2 *Left:* The angular power spectrum ($D_\ell \equiv \ell(\ell+1)C_\ell/2\pi$) of CMB temperature fluctuations measured by the Planck satellite [592]. *Right:* The power spectrum monopole and quadrupole of the DESI luminous red galaxy sample [183]. Acoustic oscillations in the primordial photon-baryon fluid are imprinted as wiggles in both the CMB and galaxy power spectra whose wavelength is given by the sound horizon at the time of recombination. . . 6

1.3 A summary of CMB secondary anisotropies along with their order-of-magnitude estimates. Figure from ref. [89]. 8

1.4 Illustration of the spacetime region influencing galaxy formation (shaded tube). Time flows upward while the horizontal axis corresponds to spatial hypersurfaces. The line connecting Lagrangian coordinate $\boldsymbol{q} = \boldsymbol{x}_{\rm fl}(0)$ to the observed Eulerian position $\boldsymbol{x} = \boldsymbol{x}_{\rm fl}(\tau)$ represents the worldline of the galaxy. The small width $R_* \ll (aH)^{-1}$ of the tube reflects the assumption that galaxy formation is only influenced by physics within a limited spatial neighborhood. Figure from [189]. 13

1.5 Comparison of advected operators appearing in the bias expansion with the fully nonlinear 21 cm density field from simulations. Figure from [58]. . . . 14

2.1 Bias-hardening against point source foregrounds. The solid curve is the determinant of the matrix in Eq. (2.10), which is the ratio of the bias-hardened estimator's noise to the standard QE's noise. It is close to unity on most scales, implying that the noise penalty for bias hardening is minimal. The dashed curve is the ratio of the SNR of the point source estimator to the relative bias to C_L^κ from the source trispectrum. This ratio is roughly 10 at all scales. Thus in the regime where the bias to the standard QE is significant, the source estimator reconstructs the sources with a high SNR, and the source-hardened estimator can remove the bias due to sources with little noise cost. The spike is due to a zero crossing in the response $\mathcal{R}_L$ at $L \sim 2500$. 23

2.2 Noises of the Point Source Hardened (PSH), Profile Hardened (PH), Point source and Profile Hardened (PPH), Shear, and standard quadratic (QE) estimators for $\ell_{\mathrm{max,T}} = 3500$. In black is the lensing signal C_L^κ. Hardening against a single source (PSH,PH) marginally increases the noise over QE, whereas hardening against multiple sources (PPH) results in a larger noise penalty. We also note that the noise cost of bias-hardening increases with the size of the source profile u_ℓ, which is discussed further in Appendix A.5. 27

2.3 Relative bias on the cross with LSST at $\ell_{\mathrm{max,T}} = 3500$. The gray boxes are the statistical uncertainties when reconstructing $C_L^{\kappa\delta_g}$ with the standard QE. We find that hardening against point sources (PSH) nulls or dramatically reduces the bias from radio point sources and CIB, whereas hardening against a tSZ-like profile (PH) nulls or significantly reduces the bias from tSZ and kSZ. Hardening against both (PPH) effectively nulls the bias from all foregrounds, resulting in a bias that is roughly an order of magnitude smaller than the bias to the standard QE. 30

2.4 Bias to the lensing auto-correlation at $\ell_{\mathrm{max,T}} = 3500$. The gray boxes are the statistical uncertainties when reconstructing C_L^κ with the standard QE. Hardening against point sources (PSH) significantly reduces the primary and trispectrum bias from radio point sources and CIB, whereas hardening against a tSZ-like profile significantly reduces these biases from tSZ. Simultaneously hardening against both (PPH) effectively nulls the primary bias. Bias-hardening has no significant effect on the secondary bias. 31

2.5 Relative bias on the lensing amplitude and the cross-correlation with LSST as a function of $\ell_{\mathrm{max,\,T}}$. In gray is the 1σ error for each estimator $1/\mathrm{SNR}(A_{\mathrm{lens}})$. CIB and tSZ are the main sources of bias to the standard QE. The Point Source Hardened (PSH) estimator drastically reduces the bias from CIB, whereas the Profile Hardened (PH) estimator reduces the bias from tSZ. The Point source and Profile Hardened (PPH) estimator reduces the bias from both tSZ and CIB to a level comparable to Shear while at a lower noise cost. 32

2.6 Total signal-to-noise of the lensing amplitude (top) and the cross correlation with LSST (bottom). The black markers denote the highest $\ell_{\mathrm{max,\,T}}$ where the bias is less than 1σ. We find that the Profile Hardened (PH) estimator reconstructs the lensing amplitude with the highest signal to noise: a $\sim 50\%$ improvement over the standard QE. We find that the Point Source Hardened (PSH) and Point source and Profile Hardened (PPH) reconstruct the cross with similar SNRs: a $\sim 50\%$ improvement over the standard QE. The Hybrid estimator is a combination of Shear and QE, and is described in greater detail in Appendix A.2. We note that the color-scale saturates at 2σ, and that the bias to the standard QE can be $> 10\sigma$ at $\ell_{\mathrm{max,T}} \gtrsim 3000$. 33

3.1 *Top:* The angular averaged temperature-only lensing weights ($\times 2\pi N^{\kappa}_{\mathbf{0}}\ell$) in the large-lens limit ($L \to 0$) when hardening against a point source (blue), hardening against a Gaussian profile with a $2'$ FHWM (purple) and when choosing the weights to minimize the variance (red). Note that since the point source hardened estimator has zero response to point-sources (when $w_\ell = 1$) the area under the blue curve vanishes. *Bottom:* The response to a Gaussian foreground w_ℓ with a full width half max θ_{FHHM}, normalized to the response of the MV estimator to point sources. In both plots we take $\ell_{\mathrm{max},T} = 3500$. See section 3.5 for details regarding the instrumental assumptions. 47

3.2 The top panel shows the reconstruction noise (per Fourier mode) when reconstructing on a single pair of maps. The colors denote the pair XY. Solid, dashed and dashed-dotted curves represent minimum variance (MV), point source hardened (PSH) or large-lens suppressed (LLS) estimators respectively. $C^{\kappa\kappa}_L$ is shown in black. The bottom two panels show the percent noise cost for the LLS and PSH estimators relative to the appropriate MV estimator. 51

3.3 Lensing reconstruction noise for the optimal linear combination of single-pair LLS (orange) and point source hardened (PSH, red) estimators. The noise for the global MV (GMV) estimator is shown in blue. For reference we show the MV EB noise (purple) and $C^{\kappa\kappa}$ (black). 52

3.4 *Top two panels:* Solid lines show the 1σ error of the lensing amplitude ($A_L \equiv C_L^{\kappa\kappa,\mathrm{measured}}/C_L^{\kappa\kappa,\mathrm{true}}$) for each single-pair minimum variance (MV) estimator. Colors denote pairs of temperature and polarization. Dashed lines show the simulated (absolute) biases to the MV estimators due to radio point sources, with the shaded regions denoting the 1σ scatter of the mean. *Bottom:* Error (solid) and bias (dashed) to the lensing amplitude for the optimal linear combination of MV single-pair estimators. The dotted curve shows the bias for an optimal combination of point source hardened estimators. See sections 3.5 and 3.6 for details regarding the experimental configuration and simulations respectively. 54

3.5 Simulated relative biases to the CMB lensing power spectrum from point sources (assuming a 5 mJy flux cut and 3% polarization fraction) for single-pair lensing estimators with $\ell_{\mathrm{max},T} = 3500$ and $\ell_{\mathrm{max},P} = 5000$, and CMB-S4 like sensitivity. Columns correspond to different pairs of temperature and polarization, while the rows indicate primary, secondary or trispectrum biases. Biases for the Minimum Variance (MV), Large-Lens Suppressed (LLS) and Point Source Hardened (PSH) estimators are shown in blue, orange and red respectively. Note that for the TE LLS estimator we only show biases for $L < 120$ since this estimator's noise dramatically increases at smaller scales (see Fig. 3.2), and that the TE, TB primary biases vanish analytically. The assumed L-binning is depicted in Fig. 3.4, and the shaded regions denote the 1σ scatter of the mean biases across 40 cutouts. We stress that the TT estimator here only shows the contribution from point sources: the overall bias is significantly larger when contributions from the CIB, tSZ and kSZ are included [519]. 55

4.1 The top panel shows the temperature power spectrum for different multifrequency combinations, compared to the noiseless lensed CMB (black curve). Compared to the minimum-variance combination (ILC, red), the deprojection of tSZ (orange) comes at a larger cost in noise than that of CIB (green). Jointly deprojecting tSZ and CIB (blue) causes a more than ten-fold noise increase at $\ell \sim 3000$, where most of the lensing information resides. We note that the significant increase in power at large scales from joint or tSZ deprojection is primarily due to a large increase in atmospheric noise. The middle and bottom panels show the tSZ and CIB power spectra, respectively, compared to the single-frequency case (150 GHz). For tSZ (middle), the standard ILC and the CIB deprojection cause a significant increase. For CIB (bottom), the ILC very effectively reduces the power; however, the tSZ-deprojection enhances it by more than an order of magnitude. 63

4.2 As illustrated in the top panel, the standard ILC is the linear combination of multifrequency maps which minimizes the lensing noise, for the standard QE and all the other estimators considered in this chapter. As expected from Fig. 4.1, deprojecting CIB (green) comes at almost no cost in lensing noise, whereas deprojecting tSZ (yellow) multiplies the noise by ~ 2. Despite causing a more than ten-fold increase in temperature noise, joint deprojection only causes a four-fold increase in lensing noise. This can be understood by counting the number of signal-dominated temperature modes. The bottom panel shows the fractional bias to the CMB lensing auto-spectrum for the various frequency combinations. The gray boxes denote the bandpower errors for the standard ILC. For the ILC, the bias is largely dominated by the enhanced tSZ signal. While the joint deprojection successfully nulls all biases, simply deprojecting tSZ or CIB does not. 64

4.3 *Left:* The total |bias| and noise on the lensing amplitude $A_{\rm lens}$ for the various ILC combinations. The biases for the standard QE are shown in blue, Shear in red, the Point Source Hardened estimator in green, and the Profile Hardened estimator in purple. *Right:* The same as the left panel, but for the amplitude $A_{\rm cross}$ of the cross-correlation with the mock LSST sample. Circles denote the minimum variance (MV) ILC, squares (triangles) denote CIB (tSZ) deprojection, while diamonds denote the simultaneous deprojection of both the CIB and tSZ. The lines connecting MV ILC to either CIB, tSZ or joint deprojection follow Eq. (4.19), with step sizes $\Delta t = 0.2$. Throughout we sum over $20 < L < 1000$ when calculating the bias and noise for either $A_{\rm lens}$ or $A_{\rm cross}$, and take $f_{\rm sky} = 0.4$. 65

5.1 Contour plots of the peak normalized SNR(k, z) for the galaxy-CMB lensing cross-correlations ($C_\ell^{\kappa g}$, left) and the galaxy auto-correlations (C_ℓ^{gg}, right) when taking $\ell_{\rm max} = 600$. SNR(k, z) is defined as the contribution to the signal-to-noise ratio per $dz\, d\ln k$, such that the total (signal-to-noise ratio)^2 is SNR$^2 = \int dz\, d\ln k$ SNR$^2(k, z)$. See Appendix D.2 for a more detailed discussion. The colors of the contours correspond to the redshift bin (see Fig. 5.2 for the redshift distributions and Fig. 5.3 for the measurements themselves) while the shading indicates the value of each contour (from lighter to darker these values are 0.15, 0.3, 0.45, 0.6 0.75, and 0.9). The brown solid line corresponds to $k = 600/\chi(z)$, while $k_{\rm nl} = [\int dk P_{\rm lin}(k)/6\pi^2]^{-1/2}$ (grey dot dashed line) and $k_{\rm HEFT} \simeq 0.6\, h\,{\rm Mpc}^{-1}$ (grey dashed line) is the maximum scale for which we expect HEFT to be valid to the percent level [327] for LRG-like galaxies. 74

5.2 *Left*: The imaging footprints of our LRG sample (blue, orange, red) and the ACT lensing footprint (purple). The SFD $E(B-V)$ map [529] is shown in greyscale. *Right*: In the top panel we show the (spectroscopically-calibrated) normalized redshift distributions of the four LRG samples (labeled by color) averaged over the full photometric footprint from [643]. The bottom three panels show the variation of the redshift distribution across imaging regions: $\Delta\phi(z) = \phi^X(z) - \phi^{\rm full}(z)$ where $X =$ North, DECaLS, or DES. 77

5.3 In the top row we plot the LRG cross-correlation ($C_\ell^{\kappa g}$) with *Planck* PR3 (blue), PR4 (orange) and ACT DR6 (red), while in the third row we plot the LRG auto-spectra (C_ℓ^{gg}). With the exception of the PR3 cross-correlation (computed following §5.3) these spectra and their covariances are taken from the companion paper [320] (and cross-validated in Appendix D.3). The hybrid covariance computed in the companion paper [320] finds that neighboring bandpowers are correlated by at most 4%. The black solid lines are the best-fit HEFT prediction when jointly fitting PR4 and DR6 using all four redshift bins. We show the residuals with the best-fit prediction in the second and fourth rows. For each measurement we quote its individual χ^2 value (indicated by color for the cross-correlation measurements). The grey shaded regions correspond to the scales omitted from our fiducial HEFT fits to the galaxy auto- and cross-correlation with *Planck*, while the shaded red regions are additionally excluded in our fits to the ACT cross-correlation (see Table 5.2). There are 9 data points for each *Planck* cross-correlation, 8 for each ACT cross-correlation and 7 for each galaxy auto-correlation. We estimate the goodness of fit in §5.7. 81

5.4 Estimation of "beyond Limber" effects within linear theory for both the galaxy auto- and cross-correlation (left and right). The middle panels show the ratio of the Limber approximation to the full integral, while the bottom panels show the ratio when evaluating the power spectra at a fixed redshift $z_{\rm eff}$, defined in Eq. (5.11). The shaded bands correspond to 1 ± 0.01. . . . 88

5.5 We estimate the size of "volume effects" for our fiducial analysis choices by fitting to a noiseless model prediction assuming a PR4 (blue), DR6 (red) or joint PR4+DR6 (purple) covariance. We consider fitting to individual redshift bins (z_i on the x-axis) or jointly fitting to all redshift bins (All). We show the marginal mean (circles), $\pm 1\sigma$ errors (caps) and best-fit value ($\times$'s) for $S_8^\times \equiv \sigma_8(\Omega_m/0.3)^{0.4}$ (top), S_8 (middle) and σ_8 (bottom). When quoting σ_8 constraints we include a mock BAO prior (see §5.5). The horizontal dashed lines indicate the cosmological parameters ($\Omega_m = 0.286$ and $\sigma_8 = 0.82$) used in the model prediction. 94

5.6 Cosmological parameter recovery when fitting to mock galaxy and CMB lensing measurements from the Buzzard simulations. In the left plot we show $S_8^\times \equiv \sigma_8(\Omega_m/0.3)^{0.4}$ (top), S_8 (middle) and σ_8 (bottom) constraints for HEFT (circles) and linear theory (squares). In purple we adopt a PR4+DR6 covariance, while in black we use a covariance representative of the Buzzard simulation. We include a mock BAO prior when quoting σ_8 constraints (see §5.5). In the right panel we show the model-independent measurements of α_8 (see Eq. 5.8). The location of the best fit is indicated with a $\times$. Dashed lines indicate the values ($\Omega_m = 0.286$, $\sigma_8 = 0.82$, $\alpha_8 = 1$) of the Buzzard cosmology. 96

5.7 The relative variation in LRG power spectra (first and third rows) and their cross-correlation with PR4 (second and fourth rows) across different imaging footprints. We divide by the (binned) C_ℓ's used for covariance estimation (see §5.3). Columns indicate the redshift bin. In blue we show the difference between DECaLS and DES, while in orange we show the difference between North and DES. The gray regions are excluded from the fitting range. The PTEs for the cross-correlation measurements are sensible: 0.58, 0.44, 0.62, 0.55 for DECaLS−DES and 0.73, 0.11, 0.25, 0.84 for North−DES. The PTEs for the auto-correlation measurements are: 5.1×10^{-11}, 1.1×10^{-3}, 0.33, 5.1×10^{-8} for DECaLS−DES and 1.1×10^{-3}, 0.075, 0.85 and 9.7×10^{-8} for North−DES. The small PTEs indicate significant variations in the LRG sample's statistical properties on different imaging footprints. The solid (blue and orange) lines show the expected variation if the galaxies were drawn from the same 3D statistical distribution (taken to be the fiducial best fit to PR4+DR6) but projected to 2D using the different redshift distributions shown in Fig. 5.2. 98

5.8 The relative difference in LRG clustering between the NGC and SGC regions of the DECaLS footprint. We divide by the (binned) C_ℓ's used for covariance estimation (see §5.3). The gray regions indicate our fitting range. The PTEs are 0.62, 0.63, 0.32 and 0.43 respectively. 100

5.9 The relative difference in our fiducial LRG power spectra (first row) and their cross-correlation with PR4 (second row) with those measured with an additional DEC $> -15°$ mask applied. We divide by the (binned) C_ℓ's used for covariance estimation (see §5.3). The gray regions indicate our fitting range. The PTEs for the cross-correlations are 0.36, 0.13, 0.59 and 0.53 respectively. With the exception of z_3, the PTEs for the auto-correlations are clearly inconsistent with fluctuations: 2.7×10^{-5}, 9.2×10^{-5}, 0.24 and 2.5×10^{-13} respectively. 101

5.10 The relative change in the LRG auto spectra when adopting a stricter $E(B-V)$ cut (from 0.15 to 0.05; blue), adopting a stricter stellar density cut (2500 to 1500 deg^{-2}; orange), removing the SFD $E(B-V)$ map [529] as a template when constructing systematics weights (green), and neglecting the systematics weights altogether (red). Different panels indicate the redshift bin. The gray vertical regions indicate our fitting range, while the gray horizontal band indicates the interval $[-0.01, 0.01]$. 101

5.11 The correlation ($\equiv C_\ell^{ab}/\sqrt{C_\ell^{aa}C_\ell^{bb}}$) between the fiducial systematics weight maps (panels) and galaxy density contrast maps (colored lines and filled circles) after the systematics weights have been applied. We additionally show the correlations between the weight maps and *Planck* PR4 convergence map (blue). The correlations between the weight maps and galaxy maps without any systematics weights applied are shown by the pale lines with $\times$ markers. The gray vertical region indicates our fitting range for the galaxy auto-spectra. The horizontal gray band corresponds to the interval $[-0.02, 0.02]$, which is of the same order as the observed point-to-point scatter in the measured correlations. 102
5.12 The measured galaxy cross-spectra (circles with errorbars). Black lines indicate the predicted (Eq. 5.14) cross-spectra from the best-fit to our fiducial PR4+DR6 analysis (§5.7), neglecting correlated shot noise. We also show the mean (blue line) and $\pm1\sigma$ interval (shaded blue) associated with the PPD of that same fit (again neglecting shot noise). Note that in some panels (e.g. $C_\ell^{g_3 g_4}$) the mean and best-fit curves overlap. By contrast, the cross-correlations between neighboring bins are all low compared to the prediction. This effect is not yet fully understood, but we subtract a $1/\ell$ contribution from the best fit prediction to obtain the dashed black curves and then check the impact of marginalizing over a bias of that form in §5.7. In red we show the magnification contribution to the best fit, i.e. the last two terms in Eq. (5.14). 103
5.13 Connecting the dots between the White et al. [625] S_8 constraints (far left) and those obtained from our fiducial PR3 reanalysis (far right). Highlighted in orange are the four significant changes in the analysis that are predominantly responsible for the shift in S_8. The shaded bands correspond to the $\pm1\sigma$ region of our fiducial PR3 constraints. Best-fit values are indicated by $\times$'s. See the text in §5.7 for a detailed discussion of the analysis changes. . 107
5.14 Fiducial cosmological constraints (without BAO) when analyzing each redshift bin separately (colors), both for *Planck* PR4 (left) and ACT DR6 (right). The dashed lines correspond to $S_8 = 0.75$ and 0.85. 109
5.15 *Left:* Cosmological constraints when adopting our fiducial analysis choices. We jointly analyze the four redshift bins and their cross-correlations with ACT DR6 (blue), *Planck* PR4 (red) and their combination (black). *Right:* Same as before, but now with BAO data included [84, 504, 28]. 110
5.16 Marginal S_8 (top) and σ_8 (bottom) constraints when fitting to each redshift bin individually, or jointly fitting to all redshift bins. We show the constraints from PR4 in blue, DR6 in red and their combination in purple. The location of the best-fit value is indicated with $\times$'s. The dashed black curves correspond to the $\pm1\sigma$ interval associated with our joint PR4 + DR6 fit (i.e. Eqs. 5.22 and 5.23). 111

5.17 A comparison of our fiducial HEFT $\sigma_8(z)$ constraints (black) to those obtained from a "model independent" (blue, §5.4) or linear theory (red, §5.4) approach. We note that there is no linear theory point for z_1 as we do not include this bin in our fiducial linear theory analysis. In all scenarios we jointly fit to the cross-correlation with *Planck* PR4 and ACT DR6 in each individual redshift bin, and use the fiducial scale cuts and priors listed in Tables 5.2 and 5.3 respectively. We include a BAO prior for both the HEFT and linear theory constraints, while for the model-independent constraints the cosmology is fixed to the values listed in Table 5.3. We offset the linear theory and model independent constraints by $\Delta z = \pm 0.015$ respectively from the effective redshifts for clarity. The black dashed line is the predicted ΛCDM evolution assuming a *Planck* 2018 cosmology [475]. In gray we plot $0.4 + c\,dN/dz$ for each galaxy sample, where c is an arbitrary constant, to illustrate the redshift origin of the signal. Following §5.7 we estimate 9, 13 and 28% correlations between the $z_1 - z_2$, $z_2 - z_3$ and $z_3 - z_4$ amplitude measurements respectively. 113

5.18 Variations (or lack thereof) in our baseline PR4+DR6 S_8 constraints (without BAO) for both HEFT (blue) and linear theory (orange) with different data subsets, priors and systematic mitigation. Best fit values are indicated with ×'s. 115

5.19 Variations in S_8 constraints (without BAO) when additionally masking the LRG sample with the North, DECaLS, DES, DEC $> -15°$, or ACT DR6 lensing mask. Best-fit values are indicated with ×'s. The horizontal lines and shaded bands correspond to mean and $\pm 1\sigma$ intervals of the fiducial constraints for each scenario (PR4 or DR6; HEFT or linear theory). We suspect that the discrepancies between the mean and best fit S_8 values when analyzing the individual imaging footprints with HEFT are the result of noisy minima and volume effects. 116

5.20 Linear theory DR6 S_8 constraints (without BAO) for alternative lensing reconstruction scenarios. "T and P only" refer to temperature and polarization only lensing reconstruction respectively; "CIB deproj." is short for a CIB deprojected internal linear combination of the 90 and 150 GHz channels; while "40% mask" is short for a more restrictive Galactic mask that is applied to the primary CMB maps before performing the lensing reconstruction. Best-fit values are indicated by ×'s. The shaded region corresponds to the $\pm 1\sigma$ interval associated with our fiducial linear theory DR6 analysis. We find statistically consistent results across different frequencies, temperature and polarization, different galactic masks, and combinations of the above. . . . 117

5.21 A comparison of our results with recent tomographic structure growth measurements obtained from cross-correlations of CMB lensing and galaxy positions. In blue we show the results obtained from the blue and green unWISE samples (correlated with PR4 and DR6) [223], while in orange we show the results obtained from the lowest two redshift bins of the *Quaia* quasar sample (correlated with PR4) [469]. Following §5.7 we estimate 9, 13 and 28% correlations between the $z_1 - z_2$, $z_2 - z_3$ and $z_3 - z_4$ amplitude measurements respectively. 118
5.22 A compilation of S_8 and σ_8 constraints from cross-correlations of galaxy positions and CMB lensing (purple and blue); cross-correlations of CMB lensing with galaxy lensing and their positions (orange); galaxy lensing (green); the CMB lensing power spectrum (red) and the primary CMB (black). We differentiate "early" and "late" time measurements with the black dashed line. Constraints derived from this chapter are shown in purple, with the purple band representing the $\pm 1\sigma$ contour of our fiducial PR4+DR6 analysis. See the text in §5.8 for a detailed discussion. 119

6.1 Relative difference in the growth (solid) and expansion (dashed) history between best-fit DESI ΛCDM (blue) or $w_0 w_a$CDM (green) [184] and best-fit *Planck*+ACT ΛCDM [475, 362] cosmologies. Curves are computed with `CLASS` [97]. 126
6.2 We show CMB lensing cross-correlations ($C_\ell^{\kappa g}$) in the top two panels, lens-source cross-correlations ($C_\ell^{\mu g}$) in the third and fourth panels, and galaxy auto-correlations (C_ℓ^{gg}) in the bottom panel. Colors indicate the galaxy sample. Black shaded regions are excluded by our scale cuts for all three samples; red regions are additionally excluded for both BGS bins; while the blue regions are additionally excluded for the lowest redshift BGS bin. Solid (dot dashed) lines indicate best-fit predictions from our BGS–CMB lensing cross-correlation (galaxy magnification) analysis. Note that the BGS best-fit C_ℓ^{gg} curves from both analyses lie on top of one another. We have shifted the data (and best-fit curves) by $\Delta\ell = \pm 5$ and rescaled the C_ℓ^{gg} errors by a factor of 5 for clarity. 128
6.3 Structure growth from the joint analysis of BAO and CMB lensing cross-correlations with DESI BGS (dashed purple) and LRG [320, 511] (solid purple) imaging data, *unWISE* galaxy catalogs [223] (blue) and *Quaia* quasars [469] (orange). The shaded purple band represents the $\pm 1\sigma$ credible interval for our combined BGS & LRG fit (bottom row of Table 6.3). We plot our magnification constraint in pink and the *Planck*+ACT ΛCDM evolution of $\sigma_8(z)$ in black [592, 503, 362]. The shaded regions on the bottom of the plot are proportional to the redshift distribution of each galaxy sample. 135

7.1 The three coefficients appearing in Eq. (7.9) for the experimental setup discussed in §7.4, with the color labeling the coefficient. Solid (dashed) lines correspond to positive (negative) values. 142
7.2 The isotropic response to lensing (Maya blue) and to patchy screening (Amaranth pink) for the experimental setup described in §7.4. In the top and bottom panels we fix $r = 2$ and 8 arcmin respectively, while solid and dashed lines correspond to T_l and T_s being evaluated at offset or identical positions. Note that the response to lensing is multiplied by 10. In light gray we plot the CMB lensing cross-correlation measured from the AbacusSummit simulations for a DESI LRG-like set of mock galaxies [270]. The dotted line shows the product of the high-pass filter, the beam B_L and the Bessel function of the first kind $J_0(Lr)$. 144
7.3 The predicted lensing bias from stacked (catawba) and mean (black lines) hardening for a DESI LRG-like set of mock galaxies built from the AbacusSummit simulations. We find excellent agreement between both hardening approaches and the signal obtained when stacking on lensed CMB maps with no patchy screening signal (light catawba). Dashed and solid curves correspond to T_l and T_s being evaluated at identical or offset positions respectively. The profiles have been binned into 9 annuli whose 10 radial edges are linearly spaced between $r = 0$ and 10 arcmin. 146
7.4 Mean lensing bias to the real space stacked estimator for the combined *unWISE* blue and green sample (solid black line). In blue we show the unfiltered measurement from [154]. 147

8.1 Estimates of the optical depth of DESI LRG groups using the two estimators defined in Eq. 8.12 and Eq. 8.13 and applying them to the ACT DR5 combined f090 and f150 maps. We show that the two estimates of $\tau_{\rm hi}$ are in very good agreement if we take into consideration the contribution from CMB lensing to the 'patchy screening' estimator (i.e., Eq. 8.13), which makes up about two thirds of the total signal. The lensing contribution is assessed using the N-body simulation, AbacusSummit (see Section 8.4, for more details). A noteworthy feature is that the high-pass filtering of the CMB map makes the profiles with high, low and no baryonic feedback look very similar to each other (see also Ref. [154]), rendering the comparison of profile shapes at high significance somewhat challenging given the signal-to-noise of the data. 159

8.2 Optical depth signal in the 'patchy screening' measurement from DESI LRG × ACT DR5, obtained by stacking via the estimator in Eq. 8.13. To make the comparison with simulations and the kSZ, we add the lensing contribution to the kSZ and simulations at the mean halo mass of the galaxy sample, $\log M \approx 13.3$, using the AbacusSummit weak lensing maps. The measured τ profile has a lower amplitude compared with TNG300-1, implying that the true baryonic feedback is stronger than the one predicted in TNG300-1. The profiles are more akin to the Illustris τ profile, implying that the amount of feedback Illustris predicts, which is large compared with the majority of hydro simulations, reflects more accurately the extent to which gas is expelled in the Universe through processes such as AGN feedback. 162

8.3 CAP (*top*) and high-pass-filtered (*bottom*) profiles measured in IllustrisTNG (TNG300-1) around high-mass galaxies from 2D maps of the optical depth, τ (dotted), and the kSZ effect (solid) at $z = 0.5$. We note that the profiles are not convolved with the ACT beam and that the maps contain no primary CMB or lensing contributions. The units of the signal in the top panel are steradians, as the CAP filter measures the difference between the integrated signal in a disk and in a ring around it (see Eq. 8.7). The quantity on the y-axis can thus be thought of as a cumulative baryon profile, whereas $\tau^{\rm hi}$ in the bottom panel is unitless, as it corresponds to the mean high-pass filtered optical depth. The three colors correspond to stacking on the default LRG sample ('All galaxies'), stacking on only the centrals belonging to that sample, or around 80% of the galaxies ('Centrals only'), and stacking on all galaxies, but at the location of their host halo centers to avoid miscentering ('All centrals'). Most noteworthy is the fact that the kSZ and τ lie on top of each other, suggesting that they measure the same underlying quantity of optical depth. 163

9.1 *Left:* Constraints on τ and $\Omega_{\rm m}$ from the joint analysis of high-ℓ CMB and CMB lensing. The dashed yellow line $\left(\Omega_{\rm m}(\tau) = 0.349 - 0.505\,\tau\right)$ approximates the $\tau - \Omega_{\rm m}$ degeneracy direction. *Right:* ΛCDM constraints from BAO (blue) or CMB high-ℓ primary and lensing with a τ prior (pink/black). 169

9.2 Constraints on $M_{\nu,\rm eff}$ from the joint analysis of BAO and CMB data when including a τ prior ($\tau = 0.06$ in pink and $\tau = 0.09$ in black) or low-ℓ data from *Planck* (orange). 171

9.3 Constraints on the $w_0 w_a$ model from the combination of high-ℓ primary CMB data, CMB lensing, DESI BAO and a prior on the optical depth. Locations of posterior maxima are indicated by triangles. 172

9.4 When discarding the low-ℓ data from *Planck* we infer $\tau = 0.090 \pm 0.012$ within ΛCDM from the combination of high-ℓ CMB, CMB lensing and BAO, or $\tau = 0.095 \pm 0.014$ for ΛCDM$+\sum m_\nu$ with a physical prior of $\sum m_\nu \geq 0.06$ eV. 173

10.1 $b^2 n$ for each of the LSS surveys considered in our forecasts. This combination of linear bias and number density is proportional to the signal-to-noise ratio in linear theory, at fixed $\boldsymbol{k}$ (§10.3). For both PUMA and HIRAX, the "number density" is taken to be the inverse of the thermal + shot noise. We note that while the fiducial PUMA design would enable observations down to $z = 0.2$ (dotted lines), we neglect this region in our forecasts, as it will be well measured by upcoming surveys such as DESI or Euclid which we already include. 178

10.2 $N_{\rm modes}$ as a function of $z_{\rm max}$ for PUMA (-5K and -32K) and MegaMapper ($2 < z < z_{\rm max}$). For PUMA, we consider both optimistic and pessimistic foreground models (§10.3), which are the boundaries of the shaded regions. In green are the number of modes for Simons Obervatory and CMB-S4 like experiments with 7 μK-arcmin and 1.5 μK-arcmin noise levels respectively, both with a 1.4′ beam and $f_{\rm sky} = 0.4$. 179

10.3 Measurements of the linear matter power spectrum at $z = 0$. For both MegaMapper and PUMA-32K we show constraints for 15 linearly space k-bins between $0.1\,h\,{\rm Mpc}^{-1} \lesssim k \lesssim 1\,h\,{\rm Mpc}^{-1}$. This figure is adapted from refs. [473, 126] using https://github.com/marius311/mpk_compilation/. . . 180

10.4 *Left:* In blue we plot several values of $\bar{n}/n_{\rm LBG}$, where $n_{\rm LBG}$ is the idealized LBG number density with $m < 24.5$. For comparison, we plot the number density of the fiducial MegaMapper survey in orange. *Right:* Comoving volume per dz per sky area ($\chi^2 d\chi/dz$) as a function of redshift. The shaded blue, purple, and (hatched) grey regions represent the volume covered by Euclid, MegaMapper and PUMA respectively, excluding sky coverage. . . 183

10.5 *Left*: $k_{\rm max}(z)$ and $\ell_{\rm max}(z) = k_{\rm max}(z)\chi(z)$, assuming that $k_{\rm max}$ is some constant fraction of $k_{\rm nl}(z) = \Sigma^{-1}(z)$. *Right*: Domain of integration in the $k_\parallel - k_\perp$ plane at $z \simeq 3.9$. The grey region is excluded by the $k < k_{\rm nl}$ constraint. For 21-cm surveys the blue region is excluded by the (pessimistic) foreground wedge. The red region is excluded by FoG-like effects, chosen so that $N_0 \sigma_v^2 (k\mu)^2$ is less than 20% of the fiducial power. In this figure we plot the red region for a MegaMapper-like survey (it would shift to higher $k_\parallel$ for a 21-cm survey). 190

10.6 *Left:* Lensing noise N_ℓ^κ for Planck, Simons Observatory and CMB-S4. The convergence power spectrum is shown in black. *Middle:* The galaxy auto-correlation $C_\ell^{g_i g_i}$ for MegaMapper, where g_i is the LBG sample in the $i^{\rm th}$ redshift bin: $z \in [2.0, 2.8], [2.8, 3.5], [3.5, 4.3]$, or $[4.3, 5.0]$. *Right:* The convergence-galaxy cross-spectra in each redshift bin. 194

10.7 Some exemplary logarithmic derivatives at $z = 2.5$, $\mu = 0$. The solid curves are calculated using Lagrangian perturbation theory, the dashed curves are calculated using linear theory. These derivatives do not account for the A-P (§10.3) effect. 195

10.8 $\Sigma(z)H(z)/[c(1+z)]$ as a function of redshift. If the redshift uncertainties $\sigma_z/(1+z)$ are near or above this curve, then they will suppress modes that could carry valuable cosmological information. The left y-axis shows the requirement on $\sigma_z/(1+z)$ in percent, while the right y-axis converts that to a velocity scale, in $100\,\mathrm{km\,s^{-1}}$. 198
10.9 For each parameter X on the perimeter, the corresponding radius is the ratio of the uncertainty σ_X from Planck+SO+DESI (PSD) alone to the uncertainty from Planck+SO+DESI+high redshift ($z > 2$) data. The constraints on ΛCDM are calculated assuming our base model (see §10.3) while each extension to our base model (M_ν, $N_{\rm eff}$, Ω_k) is treated as a single parameter extension. These constraints include both full shape and lensing data (see §10.3 for our conventions when combining surveys). For PSD+MegaMapper we include MegaMapper×SO lensing, while for all other configurations DESI×SO lensing is included. In green (orange) we show constraints for optimistic (pessimistic) foreground assumptions (§10.3). 199
10.10 Error on the A-P parameters from the reconstructed power spectrum, which can be interpreted as relative errors on $D_A(z)/r_d$ and $r_d H(z)$ respectively. The boundaries of the shaded regions denote optimistic/pessimistic foreground assumptions for the 21-cm surveys. In the top panels we show the errorbars for the optimistic case. 201
10.11 *Left:* Errors on $\sigma_8(z)$ when derived from ΛCDM, assuming a $\sigma(\omega_b) = 0.0005$ prior from BBN. *Right:* Errors on $\sigma_8(z)$ from full shape data when calculated using the fixed shape procedure, without any external priors. The errorbars with smaller caps include cross-correlations with SO lensing, but do not include any information from the convergence auto-spectrum. All 21-cm measurements assume optimistic foregrounds. The gray bands extend from 1 ± 0.01. 205
10.12 Constraints on M_ν as a function of $k_{\rm max}/k_{\rm nl}$ from the primary CMB and full shape data. The curves that are labeled "(with lensing)" include information from DESI×S4 lensing ($\kappa\kappa$, κg, gg). In both plots we set $k_{\rm max} = k_{\rm nl}$ for DESI, and only vary $k_{\rm max}/k_{\rm nl}$ for the high redshift survey of interest. *Left:* Constraints for a MegaMapper-like survey that observes a redshift-independent fraction $\bar{n}/n_{\rm LBG}$ of LBGs from $2 < z < 5$ with $f_{\rm sky} = 0.34$. *Right*: Constraints for PUMA. The edges of the shaded regions represent optimistic and pessimistic foreground assumptions. 207
10.13 $N_{\rm eff}$ constraints from Planck+SO (primary CMB) +DESI+PUMA-NK (full shape) for variable dish number N and $z_{\rm max}$. No lensing data is included in this figure. In the right plot we fix $k_{\rm max}/k_{\rm nl} = 1$, while in the left plot we fix $N = 32$. For reference, the constraint from Planck+SO (primary CMB) +DESI (full shape) is 0.044. 210

10.14 Constraints on local primordial non-Gaussianity as a function of $k_{\rm min}$ and $k_{\rm max}$, using the full shape power spectrum without any CMB priors. For SPHEREx we set $\sigma_z/(1+z) = 0.05$, while for all other surveys we set $\sigma_z = 0$. The vertical lines are the lowest possible $k_{\rm min}$ for each survey: $2\pi/V^{1/3}$, where V is the volume of the survey (including sky coverage). 212
10.15 Constraints on $A_{\rm lin}$ for various fiducial values of $\omega_{\rm lin}$, assuming $\phi_{\rm lin} = \pi/2$. For all surveys we include a Planck+SO prior on ΛCDM. The slight "bump" in the constraints at $\omega_{\rm lin} \sim 100\ h^{-1}$Mpc is due to the degeneracy with the BAO frequency. 213
10.16 The absolute error on the dark energy fraction ($\Omega_{\rm DE}$) achievable by current and future LSS surveys. 214
10.17 Change in the linear matter power spectrum (at $z = 2$) when including Early Dark Energy (EDE) for various values of $f_{\rm EDE}$ and $1.5 < \log_{10}(z_c) < 6.5$. $\theta_i = 2.83$, and all standard cosmological parameters are held at their fiducial ΛCDM values. For $z_c \lesssim 30$, the ratio $P^{\rm lin}_{\rm EDE}(k)/P^{\rm lin}_{\Lambda{\rm CDM}}(k)$ is roughly given by $1 - f_{\rm EDE}$. 216
10.18 Constraints on the maximum amplitude of early dark energy ($f_{\rm EDE}$) as a function of the time at which EDE peaks z_c, assuming $\theta_i = 2.83$. We include a Planck+SO prior on ΛCDM for all experiments. In the left panel we show constraints from full shape (FS) measurements only, while in the right panel we include a prior on ΛCDM and the nuisance parameters in Table 10.2 from SO lensing and cross-correlations with the respective galaxy surveys, following the convention in §10.3. 217
10.19 Error on the gravitational slip parameter γ from LSS×CMB lensing cross-correlations, LSS full shape data, and the primary CMB. For comparison, we show the constraint on γ from Planck [473] at $z \approx 1100$ (recombination) to the right of the plot. 218
10.20 Constraints on ΛCDM and extensions from four different combinations of experiments, using primary CMB and clustering data only. We include a $\sigma(\tau) = 0.002$ prior for the green bars. *Left*: The shading/width denotes which parameters are allowed to vary when fitting for ΛCDM (e.g. for Base ΛCDM+M_ν we keep $N_{\rm eff}$ and Ω_k fixed). *Right*: The shading denotes the number N of Parameter Extensions (NPE) that are simultaneously allowed to vary in our fits, starting with the relevant parameter on the x-axis and continuing along the cycle $(M_\nu, N_{\rm eff}, \Omega_k)$. E.g. for Ω_k, 2PE is equivalent to Ω_k, M_ν. 220

10.21 Constraints on ΛCDM and several one-parameter extensions when using the full model (FM), relative to the linear model (LM, red) or extended linear model (ELM, blue). In this plot we only include the primary CMB and full shape clustering data. Note that the largest radius in the two leftmost panels is 1.8, but 2.4 in the right panel. For the linear model, the values for n_s extend well beyond these limits: from left to right the ratios are 3.1, 6.8, and 11.6. 221

10.22 *Left*: Constraints on ΛCDM and several one-parameter extensions from MegaMapper full shape data vs redshift uncertainty, with $k_{\rm max}/k_{\rm nl} = 1$. The constraints are measured relative to $\sigma_z = 0$. *Right*: Change in constraints as a function of $k_{\rm max}/k_{\rm nl}$, relative to $k_{\rm max}/k_{\rm nl} = 1$. In both plots we include a prior on $\{\Lambda{\rm CDM}, N_{\rm eff}, \Omega_k, M_\nu\}$ from Planck + SO. The constraint on $f_{\rm EDE}$ is with fixed $\log_{10}(z_c) = 3.61$, while the constraint on $A_{\rm lin}$ is with fixed $\omega_{\rm lin} = 147\ h^{-1}\,{\rm Mpc}$. 223

10.23 *Left:* Relative improvement to Planck + SO + MegaMapper-like parameters as (MegaMapper) area is traded for depth, for fixed observing time. We assume that $m_{\rm UV}(t) = 24.5 + 2.5\log_{10}(\sqrt{t})$ and $f_{\rm sky}(t) = 0.34/t$, and scan through $0.4 < t < 8$. *Right:* Degradation to Planck + SO + MegaMapper-like constraints with a shallower u-band depth. Here the MegaMapper-like survey has a default 24.5 limiting magnitude, and detects 50% of the LBGs for a given $m_{\rm UV}$. 224

10.24 Constraints on the amplitude of local primordial non-Gaussianity with a cut on $k_\parallel$ modes. In the left panel we set $k_{\rm max} = 0.02\ h\,{\rm Mpc}^{-1}$, while in the right panel we set $k_{\rm max} = 0.05\ h\,{\rm Mpc}^{-1}$. In both panels we take $k_{\rm min} = 0.001\ h\,{\rm Mpc}^{-1}$ 224

11.1 The dimensionless 21cm power spectrum ($\Delta^2 = k^3P/2\pi^2$) as a function of redshift for our fiducial model (solid lines, see §11.4) and the thermal noise per mode (steeply rising, dashed lines) for the case of 1000 14m dishes with a 10 year integration time (first column of Table 11.1). The step in thermal noise (vertical lines) at low-k is due to a absence of baselines capable of measuring modes with $k_\perp\chi/2\pi < D/\lambda_{21}^{\rm obs}$, where D is the diameter of the dishes. The panels show $\mu = 0.1$ (left), $\mu = 0.5$ (middle) and $\mu = 0.9$ (right) with colors indicating redshift. Note that the spectrum is quite anisotropic (the amplitudes vary significantly with μ, particularly at lower redshifts), with the shape encoding valuable information about the cosmology and reionization history. 230

11.2 *Left:* Error on $\ln \overline{T}_b$ (normalized to $\mu_{\rm min} = 0$) after marginalizing over the linear bias, as a function of $\mu_{\rm min}$, following the simplified scenario outlined in §11.1. The colors denote the assumed fiducial value of the linear bias. *Right:* Foreground wedge for our fiducial scenario ($N_w = 3$; Eq. 11.18) as a function of redshift, assuming 14 m diameter dishes (see §11.4 for a more detailed discussion of the foreground wedge). 231

11.3 *Left*: In the top panel we show our fiducial reionization history (blue) and the best fit Weibull function (cyan). The bottom panel shows the residual $|\delta x| < 0.02$. This residual results in a systematic error of $|\delta\tau| = 0.0005$. The edges of shaded bands represent the δz by which one can shift the reionization history and change the optical depth by $|\delta\tau| = 0.001$. In green we show the (relative) cumulative contribution to the optical depth for redshifts $< z$. *Right*: The fiducial linear Eulerian bias b, the effective bubble size $R_{\rm eff}$ and a representative evolution of $k_{\rm max}$ (see Eq. (11.17)), where $k_\star = 1\,h\,{\rm Mpc}^{-1}$. The grey solid, dashed and dotted lines correspond to $X_\alpha = 3$, 4, 5 respectively. Such limits ensure that for $k < k_{\rm max}$ the corrections are perturbative and the higher-order counterterms are negligible. 240

11.4 Errors on the optical depth (times a thousand) for a noiseless survey covering $7 < z < 15$ over half the sky. The survey has been split into 16 redshift bins with $\Delta z = 0.5$. The fiducial scale cuts are $k_{\rm max} = 0.5\,h\,{\rm Mpc}^{-1}$, $k_\parallel^{\rm min} = 0.05\,h\,{\rm Mpc}^{-1}$ and $\mu_{\rm min} = 0.7$, and we include a prior on ΛCDM from Planck primary CMB measurements. 241

A.1 Foreground power spectra at 148 GHz after being masked with a matched-filter with a flux cut of 5 mJy and a mask patch radius of 3'. We assume a 6 μK-arcmin noise, and a 1.4' beam. 250

A.2 tSZ-like profile $u_{\boldsymbol{\ell}}$. For $\ell \lesssim 60$ we approximate the profile as a flat curve. . 253

A.3 **Shaded (blue to yellow) region:** Noise when bias hardening against a single Gaussian profile $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = e^{\sigma^2 L^2/2} e^{-\sigma^2 \ell^2/2} e^{-\sigma^2 |\boldsymbol{L}-\boldsymbol{\ell}|^2/2}$ for $\sigma \in [0', 2.4']$. σ is monotonically increasing from blue to yellow. Note that $\sigma = 0'$ corresponds to hardening against point sources (PSH). **Red:** Noise of the standard QE. 254

A.4 Verification of the noises of each estimator. The auto-correlations (blue) of the estimators $\hat{\kappa}, \hat{s}$, and $\hat{\kappa}^{\rm PSH}$ when run on a Gaussian random field with power spectrum $C_L^{\rm tot}$. In black are the theory curves. 255

A.5 Cross-correlating the true convergence κ with $\hat{\kappa}$, $\hat{s}$, and $\hat{\kappa}^{\rm BH}$ when run on a lensed CMB map. We see that both the standard QE $\hat{\kappa}$ and the point source hardened estimator $\hat{\kappa}^{\rm PSH}$ recover the lensing signal C_L^κ when correlated with the true convergence. The cross correlation of the point source estimator $\hat{s}$ with the true lensing convergence gives $N_L^s C_L^\kappa \mathcal{R}_L$, which is verified in the middle plot. The blue points are positive, whereas the red points are negative. The response has a zero crossing at $L \sim 2500$. 255

A.6 Cross-correlating the true source map s with $\hat{\kappa}$, $\hat{s}$, and $\hat{\kappa}^{\rm PSH}$ when run on the true point source map s. The cross-correlation of $\hat{\kappa}$ with s gives $N_L^\kappa C_L^s \mathcal{R}_L$, which is verified in the plot on the left. $\hat{s}$ recovers the source signal C_L^s when cross correlated with the true source map. Since the point-source hardened estimator is designed to have zero response to point sources, the cross-correlation of $\hat{\kappa}^{\rm PSH}$ and s should be zero, which is consistent with the result found in the rightmost plot. 256

B.1 Lensing reconstruction noise for the optimal linear combination of single-pair LLS estimators (solid orange) and for the Global LLS estimator (GLLS, dashed orange). For reference we show the reconstruction noise for the MV EB estimator (purple) and the convergence power spectrum (black). . . . 263

B.2 The same as the top two panels of Fig. 3.4, but with Simons Observatory-like sensitivity instead of CMB-S4. 267

C.1 The same as Fig. 4.3, but with $|\text{bias}(C_L^\kappa)|$ replacing $\text{bias}(C_L^\kappa)$ in Eq. (4.17), and similarly for the cross-correlation with the mock LSST sample. 268

C.2 The same as Fig. 4.3, using the more agressive masking technique described in App. C.4. When calculating σ, we neglect any differences in sky coverage due to more of the sky being masked (8% of the sky as opposed to 3% with the fiducial mask used in the main text), as this is a negligible effect. 271

C.3 Noise vs bias for several $\ell_{\text{max},T}$'s. For $\ell_{\text{max},T} = 3000$, 3500, and 4000 we plot the line going from the standard ILC to joint deprojection, with the linestyle denoting the $\ell_{\text{max},T}$. For $\ell_{\text{max},T} = 1500$, 2000, and 2500 we plot the noise and bias at the standard ILC for each estimator (denoted by the circles). Note that for $\ell_{\text{max},T} = 1500$, the Shear estimator's noise is larger than 10% (3%) for the auto- (cross-) correlation, extending beyond the domain of the plot. 273

D.1 Monte-Carlo normalization corrections for the baseline PR3 (left), PR4 (middle) and DR6 (right) CMB lensing reconstructions for different LRG masks. For PR3 and PR4 we show the normalization corrections for the full LRG footprint (blue) and the intersection of the LRG footprint with the Northern (orange), DECaLS (green) and DES (red) imaging regions. Note in particular the large $\mathcal{O}(10\%)$ variations for PR4, which arise from the optimal anisotropic filtering adopted by [119]. 276

D.2 Linear theory constraints from PR4+DR6 using our fiducial measurements, window functions, and hybrid covariance vs those obtained following §5.3 and Appendix D.3. 276

D.3 A comparison of analytic (blue) and brute-force marginalization (orange) over shot noise and counterterms. The true values are indicated by black dashed lines. We fit to the z_3 model prediction (see §5.5) using our fiducial HEFT model and a PR4-like covariance. In this figure we have adopted wider priors for b_s and α_x than used in our fiducial analysis (see text of §D.4). 278
D.4 Parameter constraints from the baseline PR4+DR6 fit. 280

E.1 The pole structure of $\mathcal{W}[k, T]$ for $k\,T_c = 0.5$. 285

F.1 Simulated high-pass-filtered profiles of the 'patchy screening' effect as implemented in Eq. 8.13. Maps are generated by applying the 'transfer function method' to the AbacusSummit huge box density maps between $z = 0.45$ and 1 to match as closely as possible the LRG sample redshift support (see the text in App. F.1 for more details). Stacks are performed on all halos in that redshift range above $10^{13.07}$ such that the mean halo mass matches that of the LRG sample, $10^{13.3}$ in units of $M_\odot/h$. 287
F.2 Comparison of the signal-to-noise per radial bin for our measurements using ACT and DESI data of the 'patchy screening' signal, resulting from τ and lensing, and the kSZ effect, resulting from τ after reconstructing the line-of-sight velocity field. 289

G.1 Corrections to the number densities and volume from the effects considered by ref. [611] for six of the surveys considered in this chapter, when $P_{\rm fid} = 10000\ h^{-3}{\rm Mpc}^3$. In blue is the "true" number density $\bar{n}(z)$ of each survey. In orange we plot I_{44}/I_{34} at the center of each redshift bin. In red is $\sqrt{I_{44}/I_{24}}$. In green we plot the ratio $V I_{44}/I_{22}^2$. 293
G.2 In the top two pannels we compare our distance constraints (solid), calculated using the method of §10.3, to those found using the standard procedure of ref. [547] (dashed). The ratio of the two methods is shown in the bottom two panels. 294
G.3 Constraints from cross-correlations of SO lensing with HIRAX-ideal (greay), PUMA-ideal (light blue), a version of PUMA-ideal that extends down to $z = 0.2$ (dark blue), and PUMA-ideal with non-linear nuisance terms held fixed (green). In red we plot the constraints from the SO lensing convergence autospectrum. To illustrate how the degeneracies from lensing and full shape measurements differ, we plot the constraints from PUMA-32K full shape data in yellow. All contours share the same Planck + SO prior on ΛCDM+M_ν. . 296
G.4 Planck + SO + MegaMapper (full shape only) constraints as the N_2 cutoff is varied. This is with $\log_{10}(z_c) = 3.61$ and $\omega_{\rm lin} = 147\ h^{-1}\,{\rm Mpc}$. 297

List of Tables

2.1 Comparison of the total signal-to-noise ratio on the amplitude of the CMB lensing auto-spectrum (first column) and cross-spectrum with LSST-like galaxies (second column). In all cases, the maximum temperature multipole $\ell_{\rm max,\,T}$ is selected to guarantee a foreground bias smaller than 1σ. . . 34

2.2 Relative bias (in both % and σ units) and SNR for each estimator at $\ell_{\rm max,T} = 3000$. The PH and PPH estimators reconstruct the auto- and cross-correlations respectively with the highest SNR while achieving a sub-percent level bias. 34

3.1 The linear response $g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ for all unique combinations of XY (rows) and Z (columns). Note that all entries in the table should be multiplied by a geometric factor $u_{\boldsymbol{\ell}}u_{\boldsymbol{L}-\boldsymbol{\ell}}/u_{\boldsymbol{L}}$. For brevity we have defined $\mathcal{A} \equiv \langle s_i^3\rangle/\langle s_i^2\rangle$, $\mathcal{B} \equiv \langle p_i^2 s_i^3\rangle/\langle p_i^2 s_i^2\rangle$, and $\mathcal{C} \equiv \langle p_i^2 s_i^3\rangle/2\langle s_i^2\rangle$. 45

3.2 In the first column we show scaling of the field level biases for each single-pair MV estimator relative to the MV TT field-level bias. Here $p_{\rm source}$ refers to the polarization fraction of the foreground, while $p_{\rm CMB} \sim \sqrt{C_\ell^{EE}/C_\ell^{TT}}$, where ℓ should be evaluated where the lensing weights receive the majority of their support ($\ell \sim 3000$). In the second column we show the scaling for the field-level bias suppression when the MV estimators are substituted for LLS estimators. Note that there is no LLS TB estimator to leading order in L/ℓ, while the LLS EB estimator is equivalent to the MV EB estimator. . 50

5.1 A summary of the LRG sample properties. The effective redshift $z_{\rm eff}$ is calculated with Eq. (5.11). We list the best-fit Eulerian linear bias $b_1^E = 1+b_1^L$ (where b_1^L is the Lagrangian linear bias) values from our fiducial PR4+DR6 analysis. The mean redshift $\bar{z}$, standard deviation Δz, number count slope s_μ, surface density $\bar{n}_\theta$ and shot noise $\mathrm{SN}^{\rm 2D}$ are taken from [643]. 78

5.2 The fiducial scale cuts used for each of our modeling choices. We drop the lowest redshift bin when using linear theory, and note that $\ell_{\rm min}$ and $\ell_{\rm max}$ refer to bandpower bin edges as opposed to their "centers" (as was done in [625]). 90

5.3 Sampled parameters and their associated priors for each of our modeling choices. Parameters with argument z_i (for $i = 1, 2, 3, 4$) are redshift dependent, e.g. $b_s^L(z_1)$ and $b_s^L(z_2)$ are treated as separate parameters that are individually sampled. With the exception of shot noise and number count slope, the priors for all nuisance parameters are identical in each redshift bin. We center the prior means for $s_\mu(z_i)$ and $\mathrm{SN}^{\mathrm{2D}}(z_i)$ around those listed in Table 5.1. $\mathcal{U}(X, Y)$ denotes a uniform prior between X and Y, $\mathcal{N}(\mu, \sigma)$ is a Gaussian prior with mean μ and standard deviation σ, and $\mathcal{N}_r(\mu, \sigma) \equiv \mathcal{N}(\mu, \sigma\mu)$. Priors are in **bold font** when the posterior of a nuisance parameter is (approximately) prior-dominated. 91

6.1 A summary of the BGS physical properties (effective redshift, magnification bias, Poisson shot noise), scale cuts and priors on nuisance parameters for our tomographic CMB lensing analysis. $\mathcal{N}(\mu, \sigma)$ and $\mathcal{U}(\min, \max)$ indicate Gaussian and uniform distributions respectively, while $\mathcal{N}_r(\mu, \sigma/\mu) \equiv \mathcal{N}(\mu, \sigma)$. The linear Eulerian bias is given by $b_1^E = 1 + b_1^L$. 129

6.2 Tomographic CMB lensing constraints (formatted as: mean $\pm 1\sigma$ [best fit]) from our fits to noiseless synthetic data (top three rows) and our fits to the `Buzzard` mock measurements (bottom three rows). In all cases we use the same covariance as in our fiducial analysis. The scatter in best-fits around the fiducial values for the noiseless model prediction are representative of the minimizer's precision. 131

6.3 Constraints (formatted as: mean $\pm 1\sigma$ [best fit]) from cross-correlating *Planck*+ACT CMB lensing with several combinations of the DESI BGS and LRG photo-z bins. For all combinations we include DESI Y1 BAO data. 133

9.1 Summary of our findings. Here we define ✗ as "inconsistent" at $> 2\sigma$, and ✔ as "consistent" within 2σ. 173

10.1 *Left:* A list of the spectroscopic galaxy surveys considered in our forecasts, along with their redshift ranges, samples and sky coverage. *Right:* A list of 21-cm experiments, their redshift range and sky coverage. 185

10.2 A list of the free parameters in our base model, which are marginalized over in all forecasts unless explicitly stated otherwise. In the first set of rows we list the ΛCDM parameters and their fiducial values. In the second set of rows we list the nuisance terms that we use to model redshift-space clustering. The last row shows the counterterm introduced when including lensing cross-correlations. Here δ_g is the non-linear density contrast of the tracer with power spectrum P_{gg}, δ_{cb} is the non-linear density contrast of CDM+baryons, P_{cb}^{Zel} is the power spectrum of the CDM+baryons within the Zeldovich approximation, $s_{ij} \equiv (\partial_i\partial_j/\partial^2 - \delta_{ij}/3)\delta_{cb}$ is the shear field, $s^2 \equiv s_{ij}s_{ij}$, and $k_* = 1\,h\,\mathrm{Mpc}^{-1}$. 193

10.3 1σ uncertainties on selected ΛCDM parameters and extensions from a combination of primary CMB, large-scale clustering ($P_{gg}(\boldsymbol{k})$), and CMB×LSS lensing data. For all configurations that include 21-cm data, we give constraints for optimistic/pessimistic foreground assumptions. The Planck constraints in the first row are the diagonal components of the measured TT,TE,EE+lowE covariance [475]. These would be the true constraints if the Planck likelihood were a perfect Gaussian, and agree with the constraints quoted in ref. [475] to within 5% (with the exception of M_ν, where non-Gaussianities in the likelihood increase the true σ_{M_ν} to 130 meV). Everywhere else in this chapter, Planck refers to the CMB prior described in §10.2. In rows $2-8$ (LSS alone) we include a conservative $\sigma(\omega_b) = 0.0005$ prior from BBN [475]. PSD is short for Planck + SO + DESI, while LSD is short for LiteBIRD + S4 + DESI. "Everything bagel" is short for LSD + PUMA-32K (with optimistic foreground assumptions) + MegaMapper + MegaMapper×S4 lensing, and neglects any correlations between 21-cm and LBG clustering. 203

10.4 Percent increase to PSD + MegaMapper constraints on ΛCDM (columns 2-7) and several one-parameter extensions (columns 8-11) when the base model is extended by a single parameter. The constraint on $f_{\rm EDE}$ is for fixed $\log_{10}(z_c) = 3.61$ and $\theta_i = 2.83$. 219

11.1 Forecasts for the uncertainty on the optical depth and the sum of the neutrino masses for different experimental configurations. All forecasts on the optical depth include a prior from Planck and assume $f_{\rm sky} = 0.5$. The neutrino mass constraints all include Planck, SO and the emission line galaxy sample of DESI (PSD) following ref. [512]. The baseline constraints for these are shown in the last two columns. Columns 2-9 show the impact of varying each parameter one at a time from our fiducial to a more pessimistic scenario. As a more futuristic example we include an SKA-like forecast, in which the collecting area is increased to $1\,{\rm km}^2$, which increases the sensitivity significantly over our fiducial scenario. 244

G.1 Default and new fiducial values for the nuisance parameters of the MiniMapper survey. 298

G.2 Constraints on selected ΛCDM parameters and extensions for Planck + MiniMapper. In each row we change the relevant nuisance parameter's fiducial value to the new value listed in Table G.1. The fiducial values of the remaining nuisance parameters are left unchanged. In the last row we show the constraints when all of the nuisance terms fiducial values' are set to the default values in Table G.1. 298

Acknowledgments

First and foremost I would like to thank my advisors, Martin White and Simone Ferraro, for their guidance, generosity, and impressive patience over the past five years. Their combined wisdom, and at times polar-opposite approaches toward science, has not only redefined my understanding of cosmology, but how to think critically, ask better questions and occasionally trust my physical intuition.

Second, I would like to thank those who inspired me to pursue a career in cosmology in the first place. I took my first cosmology course with Rachel Bean (a whirlwind introduction) followed by Nick Battaglia's inaugural grad-level course. I left both of those classes fascinated but deeply confused, and I'm immensely grateful to Eve Vavagiakis and Michael Niemack for enabling me to pursue that curiosity by taking me on as a clueless undergrad researcher. Due in large part to their incredible mentorship, I entered grad school with the ambition of becoming an instrumentalist. Adrian Lee graciously mentored me as I spent my first year (theoretically) tinkering with tiny lenses. It was by sheer chance that during this time I met Simone and began working on lensing estimators as a "side" project, which (after many prior failures) eventually turned into my first successful research project. The following summer Simone (virtually) introduced me to Martin and, contrary to my initial ambitions, I spent the remainder of my PhD morphing into something between an analyst, a theorist, and perhaps most importantly, a skeptic.

My first few years at Berkeley were greatly enriched by the mentorship of several (former) Berkeley postdocs. I'd especially like to thank Emmanuel Schaan for enthusiastically entertaining my mind-numbing algebra and Emanuele Castorina for his patience and uplifting humor when debugging my early forecasts. I thank both Simone and Emmanuel for introducing me to Blake Sherwin and Omar Darwish and for the engaging collaborations that followed. Thank you Marius and Vanessa for making me feel welcome at my first few BCCP journal clubs, and for your fruitful insights over pizza.

I would also like to express my gratitude to the postdocs and senior scientists at Berkeley Lab. In particular I thank Boryana Hadzhiyska, Joe DeRose and Rongpu Zhou for your wisdom and guidance throughout several collaborations. Thank you Satya Gontcho A Gontcho, for your inspiring takes on both research and life; Antón Baleato Lizancos, for your joyful and inclusive approach to science; David Valcin, for introducing me to both globular clusters and the best spots for sweets in Oakland; Roger de Belsunce, for our wing foiling adventures; and Richard Feder, for our ambitiously small lensing goals. I am incredibly privileged to have been surrounded by pioneers in our field throughout my time in Berkeley, and in particular I would like to thank David Schlegel, Julien Guy, Uroš Seljak, Liang Dai, Claire Poppett, and Patrick McDonald for their time and instruction.

I've had the pleasure of researching alongside an outstanding group of fellow students. I'd like to express my utmost thanks to Stephen Chen for his non-perturbative knowledge of LPT and guidance over the years. I'd like to thank the rest of the (slightly) more senior students − Jamie Sullivan, Biwei Dai, Byeonghee Yu and Adrian Bayer − for their comradery and for fostering a welcoming atmosphere in the Campbell corridor. Thank you Haruki for

scrutinizing my calculations and correcting them on multiple occasions, and to Mark, Alexa and Henry for your engaging questions and companionship.

I was fortunate enough to visit Perimeter Institute several times during the last year of my PhD. I thank Kendrick Smith and Selim Hotinli for their invitations, hospitality and enthusiasm towards hardcore Wigner algebra. I also thank Alex Krowleski, Zach Weiner and Neal Dalal for both technical discussions and half-baked ideas that I hope will someday evolve into fully-fledged projects. My visits to PI have reinvigorated my fascination with cosmology on multiple occasions, and for that I am incredibly grateful.

My graduate experience wouldn't have been the same had I not met Gerrit Farren, Beatriz Tucci and Sam Goldstein at the ICTP summer school, who I am now honored to call my colleagues and friends. I look forward to our future collaborations and misadventures, whether they be at the blackboard or the bars.

Lastly, I thank my friends − Madge and Tylor − for our weekly Survivor nights, spontaneous overseas adventures, and for your unwavering kindness and support through everything. You've been my constants these past six years, and I'll forever be grateful for the intersection of our worldlines.

Chapter 1

A review and an overview

History of paradigm shifts
1965 Discovery of CMB (age 10)
1982 Inflation (quantum origin of fluctuations) (age 27)
1992 Discovery of CMB anisotropies (age 37)
1998 Discovery of dark energy (age 43)
Nothing in the last 25 years!

-*George Efstathiou (age 67)*, 241st AAS meeting

Our current best cosmological model (ΛCDM) has withstood decades of observational tests, becoming the standard framework to describe the Universe's large-scale evolution. After a rapid succession of revolutionary discoveries from the mid 60's to late 90's it might seem that the current cosmological paradigm has settled into a steady state. However, subtle tensions in recent datasets have perhaps hinted at departures from the standard ΛCDM model. It is of upmost importance to scrutinize the robustness of current tensions to determine what, if anything, needs to be amended to our current understanding of the large-scale Universe.

Tensions aside, ΛCDM still leaves fundamental questions unanswered. It successfully accounts for phenomena ranging from big bang nucleosynthesis to the anisotropies of the cosmic microwave background and large-scale structure, but the true nature of both dark energy and dark matter remains unknown. Beyond the model's titular components lie further mysteries − such as the particle(s) responsible for inflation and the complex astrophysics of reionization − which continue to motivate intense study.

Ongoing and future observational efforts − including DESI [186], the Simons Observatory [22], Euclid [334], SPHEREx [197], Rubin Observatory's Legacy Survey of Space and Time (LSST) [368] and LiteBIRD [281] − are poised to provide the data needed to rigorously test these hints of beyond-ΛCDM physics. In what follows we summarize the current status of ΛCDM and its associated tensions, presenting the background necessary to understand the techniques proposed in the following chapters to robustly interpret current and future cosmological datasets.

1.1 Our current best cosmological model

Below we present a lightning review of ΛCDM. For a more comprehensive and pedagogical treatment of the standard cosmological model, its theoretical foundations, detailed observational tests, and proposed theoretical solutions to current tensions we defer the reader to refs. [194, 424, 618, 340, 192, 10].

Thermal and expansion history

Our understanding of cosmology is grounded in several key observations: our Universe is expanding [302], nearly homogeneous and isotropic on sufficiently large scales (i.e. the Copernican principle), and was once in thermal equilibrium (see Fig. 1.1). As the Universe expands its temperature cools according to $T \propto a^{-1}$ where the scale factor $a(t)$ denotes the relative size of the Universe as a function of time (see Eq. (1.1) below). As the Universe expands and cools, it passes through a series of important thermal transitions. Weak interactions freeze out around $T \sim 1\,\mathrm{MeV}$, setting the neutron-to-proton ratio. Shortly thereafter at around $T \sim 0.1 - 1$ MeV nuclear reactions commence, producing the light elements in a process known as big bang nucleosynthesis (BBN). The resulting primordial abundances of light nuclei such as ^{2}H, ^{3}He, ^{4}He and ^{7}Li are sensitive to the baryon density and expansion rate at that time, and their agreement with observations provides strong evidence for the hot big bang model [612]. At redshift $z \equiv a^{-1} - 1 \sim 1100$ the temperature drops below $0.3\,\mathrm{eV}$, allowing electrons and protons to combine into neutral hydrogen in a process known as recombination. After recombination photons propagate *nearly* freely (but not entirely, see §1.2) and form the cosmic microwave background (CMB), which today is observed as a near perfect blackbody with temperature $T_\gamma = 2.725\,\mathrm{K}$.

The homogeneous Universe is governed by the Friedmann-Lemaître-Robertson-Walker (FLRW) metric,

$$ds^2 = -dt^2 + a^2(t)\left[\frac{dr^2}{1-kr^2} + r^2 d\Omega^2\right], \tag{1.1}$$

where k determines the mean spatial curvature and $d\Omega^2$ denotes the metric of a 2-sphere. The expansion of the Universe is captured by the Friedmann equation, which relates the Hubble parameter $H = \dot{a}/a$, where $\dot{a} \equiv da/dt$, to the Universe's energy density:

$$H^2 = \frac{\rho_{\rm tot}}{3M_{\rm pl}^2} \quad \text{with} \quad \rho_{\rm tot} = \underbrace{\rho_r}_{\text{radiation}} + \underbrace{\rho_m}_{\text{matter}} + \underbrace{\rho_\Lambda}_{\text{dark energy}} + \underbrace{\rho_k}_{\text{curvature}}, \tag{1.2}$$

where we're introduced the Planck mass $M_{\rm pl} = (8\pi G)^{-1/2}$. The composition of the Universe includes several components distinguished by the scaling of their energy density. Radiation, sourced primarily by CMB photons and relativistic neutrinos, dominates at early times and scales as $\rho_r \propto a^{-4}$ due to both number density dilution and redshifting of energy. In addition, the early Universe may have contained additional light relics that were relativistic during the radiation-dominated epoch and have since decoupled. Their presence is parametrized

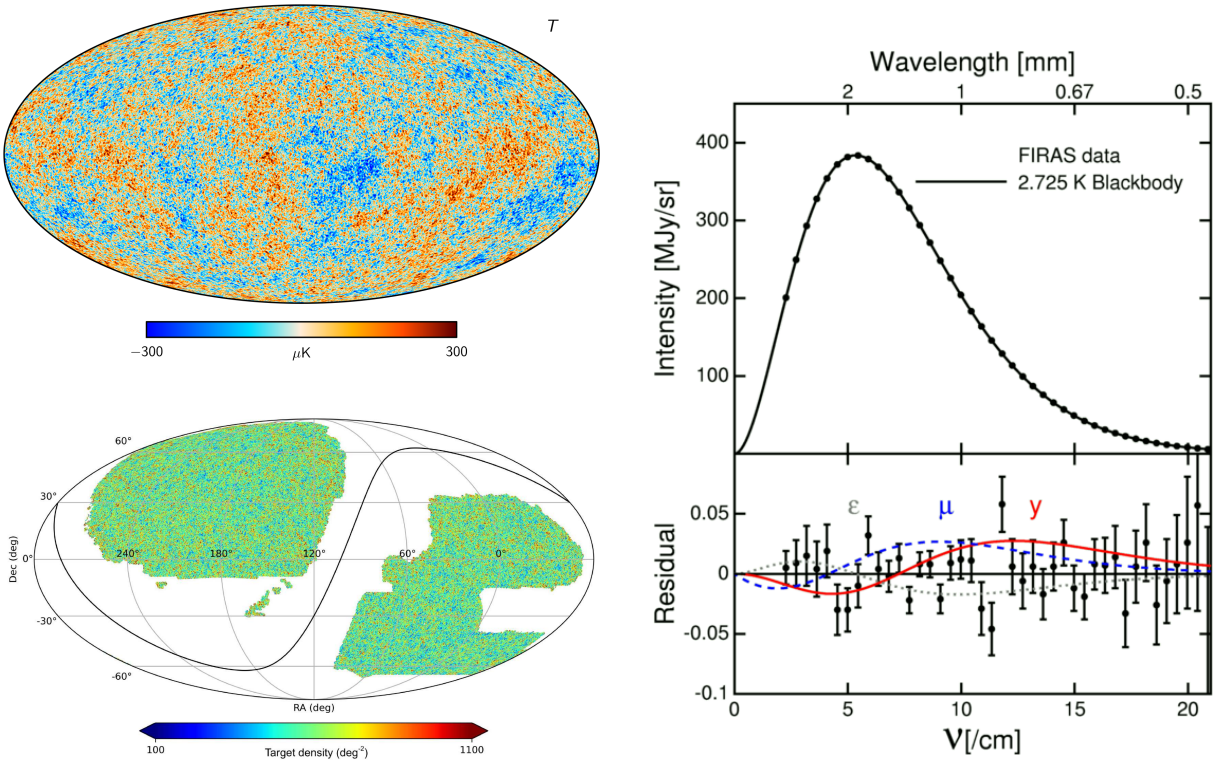


Figure 1.1: *Top left:* Mollweide projection of intensity fluctuations in the cosmic microwave background (CMB) as measured by the Planck satellite [473]. *Bottom left:* Mollweide projection of the number of DESI luminous red galaxies per square degree [640]. Both the CMB intensity and observed galaxy counts are nearly isotropic on large scales, in agreement with the Copernican principle. *Right:* The mean intensity of the CMB as a function of wavelength as measured by the COBE FIRAS instrument [234]. The discovery that the CMB follows a near-perfect blackbody distribution was crucial in establishing the hot big bang model, in which the early Universe was in thermal equilibrium. Figure from https://www.astro.ucla.edu/~wright/CMB.html.

as the effective number of neutrino species, $N_{\rm eff}$, which modifies the total radiation energy density as

$$\rho_r = \rho_\gamma \left[1 + \frac{7}{8}\left(\frac{4}{11}\right)^{4/3} N_{\rm eff}\right] \quad \text{with} \quad \rho_\gamma = \frac{\pi^2}{15} T_\gamma^4 \simeq \frac{0.260}{a^4}\,\frac{\rm eV}{\rm cm^3}, \tag{1.3}$$

where ρ_γ is the energy density of the CMB and $N_{\rm eff} \simeq 3.044$ in the Standard Model of particle physics. Deviations from this value would be a signature of new physics and are the target of ongoing and future CMB and large-scale structure surveys. Meanwhile non-relativistic species such as baryons and cold dark matter (CDM) scale as $\rho_m \propto a^{-3}$, reflecting simple dilution of number density in an expanding volume. Matter dominates the Universe's energy density between redshift $z \simeq 3400$ and $z \simeq 0.3$ after which a component with a negative effective pressure, dark energy, comes to dominate. In the simplest case dark energy is modeled as a cosmological constant Λ with a constant energy density ρ_Λ. Observations of distant Type Ia supernovae in the late 1990s provided the first direct evidence for dark energy [501, 464], and the interpretation in terms of a cosmological constant is consistent with a broad range of other probes including the CMB and large-scale structure. Finally, spatial curvature contributes effectively as an energy component with $\rho_k \propto a^{-2}$, but measurements from the CMB and large-scale structure constrain its contribution to be very close to zero, consistent with a flat Universe. The standard ΛCDM model implicitly takes $\rho_k = 0$.

Fluctuations

The discovery of a near isotropic CMB led to the *horizon problem:* how is it that the CMB temperature is isotropic to one part in 10^5 everywhere on the sky when, within the hot big bang model, the subtended angle[1] of the causal horizon at the time of recombination is only ~ 1 degree? Inflation solves this problem by exponentially expanding a tiny, causally connected patch of spacetime so that it becomes larger than our observable Universe.

Inflation models the very early Universe as small fluctuations about a de Sitter background, which is a special case of the FLRW metric (Eq. 1.1) with $k = 0$ and $a(t) \propto e^{Ht}$. In the simplest models inflation is driven by a single scalar field for which, in an appropriate gauge, the metric during inflation takes the form

$$ds^2 = \underbrace{\frac{1}{H^2(\tau)\,\tau^2}}_{=a^2(\tau)} \Big(d\tau^2 - \big(1 - 2\,\mathcal{R}(\tau, \boldsymbol{x})\big) d\boldsymbol{x}^2\Big) \quad \text{for} \quad \begin{array}{l} \tau \in (-\infty, 0) \\ \boldsymbol{x} \in \mathbb{R}^3, \end{array} \tag{1.4}$$

where $\mathcal{R} \ll 1$ corresponds to the spatial curvature of $\tau =$ const. hypersurfaces, τ (defined through $a\,d\tau = dt$) is the conformal time and H is presumed to evolve slowly for an extended period of time. The latter two conditions are satisfied if the slow roll parameters $\epsilon \equiv -\dot{H}/H^2$ and $\eta \equiv 2M_{\rm pl}^2\,\ddot{H}/H$ are small. On scales much larger than the horizon ($k|\tau| \ll 1$) quantum

[1] as observed by earthlings

fluctuations in the curvature perturbations are "frozen out," forcing them into a squeezed state with nearly all of the uncertainty placed in the curvature amplitude, whereas the curvature "momentum" ($\mathcal{R}' \equiv d\mathcal{R}/d\tau$) and its commutator with the amplitude decay like $\langle \mathcal{R}'\mathcal{R}' \rangle \propto [\mathcal{R}, \mathcal{R}'] \propto (k\tau)^2$. This results in a quantum-to-classical transition, so that the primordial curvature perturbations can be interpreted as stochastic random variables with zero momentum and an amplitude drawn from a Gaussian distribution with variance

$$\langle \mathcal{R}_{\boldsymbol{k}} \mathcal{R}_{\boldsymbol{k}'} \rangle = (2\pi)^3 \delta^D(\boldsymbol{k} + \boldsymbol{k}') \underbrace{\frac{2\pi^2}{k^3} A_s \left(\frac{k}{k_\star}\right)^{n_s - 1}}_{\equiv P_{\mathcal{R}}(k)} \quad \text{where} \quad \begin{array}{l} A_s = H^2/8\pi^2\epsilon \\ n_s = 1 + 2\eta - 4\epsilon \\ k_\star = 0.05 \text{ Mpc}^{-1}. \end{array} \tag{1.5}$$

In addition to scalars, inflation generically predicts a stochastic background of gravitational waves (tensor modes) whose power spectrum amplitude $A_t = r\, A_s$, where $r = 16\,\epsilon$ is the tensor-to-scalar ratio. A detection of primordial gravitational waves would reveal the energy scale of inflation, and is a primary science driver of current and future CMB experiments [22, 1, 281].

After inflation ends a reheating phase transfers the energy stored in the field(s) responsible for inflation into a thermal bath of radiation and matter. While the precise details of reheating remain elusive current observations are consistent the process being adiabatic, such that the initial conditions for each species is uniquely set by the curvature perturbation. After re-entering the horizon during radiation domination, perturbations in the radiation and matter evolve under gravity and microphysical interactions according to the Einstein-Boltzmann equations, imprinting acoustic oscillations in the tightly coupled photon-baryon fluid. These baryonic acoustic oscillations (BAO) are pressure supported up to the time of recombination, producing a preferential scale (the sound horizon at recombination) that is now observed as wiggles in the power spectra of CMB fluctuations and galaxy number counts (see Fig. 1.2).

After recombination the evolution of small perturbations to the matter density are governed entirely by gravity. The *linear* evolution of the matter density contrast $\delta_m \equiv \delta\rho_m/\bar{\rho}_m$ satisfies the second-order differential equation

$$\delta_m'' + aH\delta_m' - \frac{a^2\bar{\rho}_m}{M_{\rm pl}^2}\delta_m = 0. \tag{1.6}$$

The two independent solutions to this equation define the growing (D_+) and decaying (D_-) mode functions, commonly expressed in terms of the scale factor a as

$$D_+(a) = H(a)\int_0^a \frac{da'}{a'^3 H(a')}, \quad D_-(a) = H(a). \tag{1.7}$$

Matter perturbations grow efficiently during matter domination ($D_+(a) \propto a$). In particular, for sufficiently small scale ($k \gtrsim 0.02\ h\,\text{Mpc}^{-1}$) modes which re-entered the horizon earlier

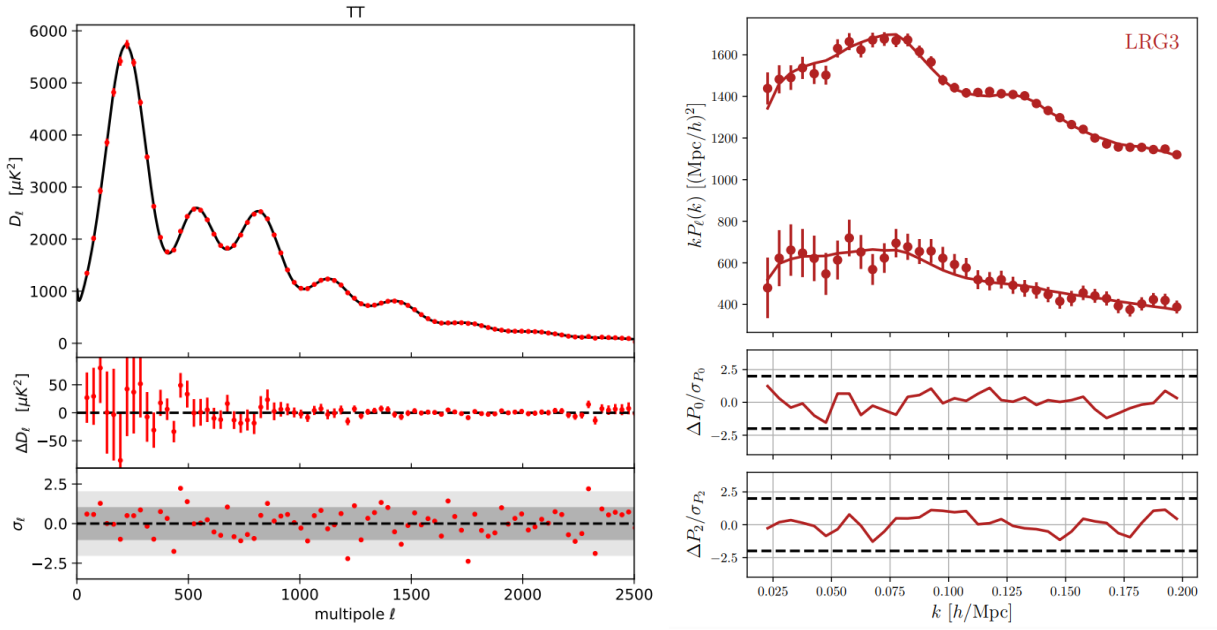


Figure 1.2: *Left:* The angular power spectrum ($D_\ell \equiv \ell(\ell+1)C_\ell/2\pi$) of CMB temperature fluctuations measured by the Planck satellite [592]. *Right:* The power spectrum monopole and quadrupole of the DESI luminous red galaxy sample [183]. Acoustic oscillations in the primordial photon-baryon fluid are imprinted as wiggles in both the CMB and galaxy power spectra whose wavelength is given by the sound horizon at the time of recombination.

and hence had a longer time to grow, this growth becomes large enough to merit a beyond-linear treatment. For modes deep inside the horizon ($k \gg aH$) the time-time component of Einstein's equations reduces the Poisson equation

$$\nabla^2\Phi = \frac{a^2\bar{\rho}_m}{M_{\rm pl}^2}\delta_m. \tag{1.8}$$

The remaining equations governing the nonlinear evolution of δ_m can be derived from the conservation of the energy-momentum tensor $\nabla_\mu T^{\mu\nu} = 0$. In particular, the $\nu = 0$ component enforces mass conservation (the continuity equation)

$$\delta_m' + \nabla\cdot[(1+\delta_m)\boldsymbol{v}] = 0, \tag{1.9}$$

while the spatial components $\nu = i$ enforce momentum conservation (the Euler equation)

$$\boldsymbol{v}' + \mathcal{H}\boldsymbol{v} + (\boldsymbol{v}\cdot\nabla)\boldsymbol{v} = -\nabla\Phi, \tag{1.10}$$

where $\boldsymbol{v}$ is the peculiar velocity field. Together Eqs. (1.8)–(1.10) form a closed system used to capture the transition from linear growth to nonlinear gravitational collapse. In practice

one solves these equations either perturbatively[2] within the mildly-nonlinear regime or using N-body simulations, yielding the complex non-linear structures observed today in tracers of the large-scale structure.

To connect the non-linear evolution of matter to biased tracers in redshift-space (§1.3) it is useful to recast Eqs. 1.8–1.10 in terms of the *Lagrangian* framework. In this framework each fluid element is labeled by its initial comoving coordinate $\boldsymbol{q}$, while the dynamics are specified by a displacement field $\boldsymbol{\Psi}$, such that the comoving position of the fluid at comoving time τ is given by

$$\boldsymbol{x}(\boldsymbol{q},\tau) = \boldsymbol{q} + \boldsymbol{\Psi}(\boldsymbol{q},\tau). \tag{1.11}$$

The displacement field itself evolves according to

$$\boldsymbol{\Psi}'' + \mathcal{H}\boldsymbol{\Psi}' = -\nabla_x \Phi, \tag{1.12}$$

where $\mathcal{H} = a'/a$ is the conformal Hubble parameter and Φ is the gravitational potential, which is sourced by the matter distribution through the Poisson equation (Eq. 1.8). Finally, the nonlinear matter density contrast is determined from the displacement field by mass conservation:

$$1 + \delta_m(\boldsymbol{x},t) = \int d^3\boldsymbol{q}\, \delta^D\big(\boldsymbol{x} - \boldsymbol{q} - \boldsymbol{\Psi}(\boldsymbol{q},t)\big). \tag{1.13}$$

The nonlinearity in the equation of motion for $\boldsymbol{\Psi}$ is encoded in the fact that the gravitational potential Φ must be evaluated at the advected position $\boldsymbol{x} = \boldsymbol{q} + \boldsymbol{\Psi}$, and the non-linear relationship between $\Phi \sim k^{-2}\delta_m$ and $\boldsymbol{\Psi}$.

1.2 Secondary anisotropies in the CMB

Secondaries are not Secondary,
Secondaries are Primary
a tale of Old Light,
only some of which is Ancient

-*Dick Bond*, New Physics from Old Light workshop

After recombination CMB photons traverse the entire observable Universe *nearly* unhindered, however, interactions between the photons and the intervening matter generate additional temperature and polarization fluctuations known as *secondary anisotropies*. Secondaries encode valuable information about the late-time Universe, complementing the primordial information imprinted at recombination.

Secondary anisotropies can broadly be classified into two categories: those arising from gravitational effects vs those arising from scattering events with free electrons produced during the reionization epoch. A summary of these effects along with order of magnitude estimates is given in Fig. 1.3. In the following discussion we give a lightning review of weak gravitational lensing, patchy screening and the kinetic Sunyaev-Zel'dovich effects − all of

[2]Perturbative approaches were rigorously explored in a previous Berkeley thesis [141].

which feature prominently in the subsequent chapters of this thesis. We conclude this section with a very brief review of the remaining secondaries featured in Fig. 1.3.

Useful estimates		
$\theta_{\rm lens} \sim 1', \ \int \dot{\Phi}\, d\eta \sim 10^{-4}, \ \frac{k_B T_e}{m_e c^2} \sim \left(\frac{v_{\rm th}}{c}\right)^2 \approx 0.01, \ \frac{v_{\rm bulk,rot}}{c} \sim 10^{-3}, \ \frac{\delta T_0}{T_0} \sim a_2 \sim 10^{-5}$		
Effect	$\delta T^{\rm effect}/T_0 \propto$	**Information**
Gravitational secondaries		
Lensing	$\theta_{\rm lens}$	Mapping the total matter, cosmological constraints
ISW, Rees-Sciama	$\int \dot{\Phi}\, d\eta$	Dark energy, curvature
Moving Lens	$v_{\rm bulk}\perp$	Transverse velocity field
Scattering secondaries		
tSZ, relativistic tSZ	$\tau(v_{\rm th}/c)^{2,4}$	Thermal pressure, temperature
kSZ, rotational kSZ	$\tau(v_{\rm bulk}\parallel/c)$	Density, LOS velocity field
Polarized SZ	$\tau(v_{\rm bulk}\perp/c)^2, \tau a_2$	Density, remote CMB quadrupole
Patchy screening	$\tau(\delta T_0/T_0)$	Density

Figure 1.3: A summary of CMB secondary anisotropies along with their order-of-magnitude estimates. Figure from ref. [89].

Gravitational lensing

Gravitational lensing arises from the deflection of CMB photons by the large-scale distribution of matter between the surface of last scattering and Earth. Because the deflection field traces the integrated gravitational potential along the line of sight, it is sensitive to the amplitude and growth of matter fluctuations across cosmic time. This makes lensing a powerful probe of the matter power spectrum and a key observable for constraining the sum of neutrino masses, as massive neutrinos suppress small-scale structure in a scale-dependent way. Lensing is sensitive to the Weyl (rather than Newtonian) potential, which when analyzed alongside non-relativistic tracers such as galaxies enables constraints on modified gravity theories that alter the relationship between the two potentials. Collectively, these features make CMB lensing a uniquely valuable tool for probing otherwise invisible components of the Universe.

In the weak lensing regime the observed CMB temperature and polarization fields are simply remapped by a deflection angle $\vec{\alpha}(\hat{n})$

$$X_{\rm lensed}(\hat{n}) = X_{\rm unlensed}(\hat{n} + \vec{\alpha}(\hat{n})), \tag{1.14}$$

where X denotes either a temperature or polarization component (e.g. Q or U). Typical deflection angles are $\mathcal{O}(1\,\text{arcmin})$ and so it is reasonable to treat α as a small perturbation. To linear order one can write the deflection field as a pure gradient

$$\vec{\alpha}(\hat{n}) = \vec{\nabla}\phi(\hat{n}), \tag{1.15}$$

where the lensing potential ϕ is a line-of-sight projection of the gravitational potential Φ weighted by a lensing kernel, which for a flat Universe takes the form

$$\phi(\hat{n}) = -2\int_0^{\chi_*} d\chi\, \frac{\chi_* - \chi}{\chi_*\chi}\Phi(\chi\hat{n}), \tag{1.16}$$

where χ is the comoving distance and χ_* indicates the distance to the surface of last scattering.

Gravitational lensing is imprinted in the CMB in several notable ways. First, lensing effectively blurs out the BAO feature imprinted at the time of recombination, smoothing the acoustic peaks in the temperature and polarization power spectra. Second, lensing partially converts the parity-even (E-mode) component of the polarization field produced by Thompson scattering at recombination into a parity-odd (B-mode) component, which must be taken into account when searching for primordial gravitational waves. Finally, while the unlensed CMB and lensing potential themselves are (nearly) Gaussian fields, the lensed CMB (Eq. 1.14) is a non-linear function of the lensing potential and is thus no longer Gaussian. In particular, the lensed CMB has a non-zero bispectrum with any LSS tracer g, which in the squeezed limit takes the form

$$\langle T_{\text{lensed}} T_{\text{lensed}}\, g\rangle \sim C^{TT}\, C^{\phi g}. \tag{1.17}$$

Given prior knowledge of the primordial temperature power spectrum C^{TT}, one can therefore reconstruct the lensing potential's correlations with LSS through quadratic combinations of the observed CMB [300]. Lensing, however, is not the only component of the microwave sky to produce a non-zero TTg bispectrum. Extragalactic foregrounds, which are highly non-linear (and hence non-Gaussian) and correlated with large-scale structure tracers, also contribute to $\langle TTg\rangle$ and must be taken into account to recover unbiased estimates of the lensing potential.

Patchy screening

During and after the epoch of reionization CMB photons undergo Thompson scattering with free electrons in the (optically thin) gas surrounding galaxies and clusters. These electrons scatter CMB photons in or out of the line of sight direction of a given observer. The fraction of CMB photons scattered out of the line-of-sight direction $\hat{n}$ is given by $1 - \text{Exp}[-\tau(\hat{n})]$, where we've introduced the optical depth

$$\tau(\hat{n}) = \sigma_T \int d\chi\, a(\chi)\, n_e(\chi), \tag{1.18}$$

where σ_T is the Thompson scattering cross-section, $a(\chi)$ is the scale factor as a function of comoving distance and $n_e(\chi)$ is the free-electron number density. Due to conservation of photons an equal number of photons scatter into our line of sight as those which scatter out of it, replacing a small fraction of the temperature fluctuation in direction $\hat{n}$ with the sky-averaged CMB signal, damping the primordial temperature fluctuations according to

$$T_{\rm observed}(\hat{n}) = e^{-\tau(\hat{n})} T_{\rm primordial}(\hat{n}). \tag{1.19}$$

Eq. (1.19) is known as *screening* of the CMB. Of particular relevance for constraining the distribution of baryons (and their associated feedback processes) is the *patchy screening* effect, which is simply the contribution to Eq. (1.19) arising from anisotropies in the optical depth field. Just as for CMB lensing, patchy screening modulates the primordial CMB fluctuations and can thus be reconstructed using a squeezed bispectrum analogous to Eq. (1.17). In fact, optimal estimators of patchy screening are biased by lensing and vice-versa. The methods developed in Part I of this thesis to mitigate foregrounds in lensing measurements are adapted to "lens-harden" the patchy screening measurements in Part II, enabling a robust inference of the distribution of baryons around DESI galaxies.

The kinetic Sunyaev-Zel'dovich effect

The kinetic Sunyaev-Zel'dovich (kSZ) effect arises from the Doppler shift of CMB photons scattered by free electrons with bulk peculiar velocities $\vec{v}_e$ along the line of sight. Unlike thermal effects, the kSZ preserves the CMB's blackbody spectrum and appears as a frequency-independent temperature fluctuation:

$$T_{\rm kSZ}(\hat{n}) = -\sigma_T \int d\chi\, a(\chi)\, n_e(\chi)\, \vec{v}_e \cdot \hat{n}, \tag{1.20}$$

where σ_T is the Thomson scattering cross-section, $a(\chi)$ is the scale factor, $n_e(\chi)$ is the free-electron number density and $\vec{v}_e$ is the peculiar velocity field.

There are two main contributions to the kSZ signal: the homogeneous kSZ, sourced by free electrons in the fully ionized late-time Universe, and the patchy kSZ, arising during the epoch of reionization when ionized bubbles first formed around luminous sources. Measurements of the patchy contribution constrain the timing and duration of reionization, while cross-correlating the the later-time kSZ with galaxy surveys allow for reconstruction of the velocity field and inference of the baryon distribution.

The kSZ effect uniquely probes ionized gas surrounding galaxies and clusters, particularly in low-density environments that are otherwise difficult to observe. In Part II of this thesis, we compare kSZ-based inferences of baryon distributions to those derived from patchy screening, finding consistency between the two approaches.

Additional secondaries

Beyond the effects discussed above, several other secondary anisotropies probe the late-time structure through gravitational and scattering mechanisms. The **integrated Sachs–Wolfe**

(ISW) effect results from the decay of gravitational potentials when the universe deviates from matter domination, which occurs both during both radiation domination (early ISW) and dark energy domination (late ISW). The latter imprints large-scale CMB temperature anisotropies correlated with the distribution of large-scale structure, providing a complementary probe of cosmic acceleration. The **Rees–Sciama effect** is the non-linear counterpart to ISW, and arises when potentials grow nonlinearly during gravitational collapse, contributing small anisotropies on arcminute scales. The **moving lens effect**, sourced by the transverse motion of massive halos with respect to the CMB rest frame, induces a characteristic dipolar temperature distortion.

Several scattering secondaries also appear in the form of spectral and polarization distortions. The **thermal Sunyaev–Zel'dovich (tSZ) effect** results from inverse Compton scattering off hot electrons in galaxy clusters and is a sensitive tracer of the thermal pressure of baryons. **Relativistic corrections (rSZ)** to the tSZ effect become important in systems with highly energetic electrons. The **rotational SZ (rSZ) effect** arises from coherent rotational motion in halos, imprinting a small dipolar anisotropy orthogonal to the line-of-sight velocity component. Though suppressed relative to the standard kSZ signal, it provides a potential probe of a halo's angular momentum. Finally, the **polarized SZ (pSZ) effect**, generated by Thomson scattering of the local CMB quadrupole by free electrons, induces linear polarization patterns that encode information about the large-scale CMB quadrupole at remote locations.

1.3 Observed tracers of the large-scale structure

Extracting cosmological information from large-scale structure surveys requires modeling how direct observables (e.g. galaxies or 21 cm emission) trace the underlying matter distribution. This task is highly non-trivial due to the vast range of scales and highly non-linear physics relevant for e.g. the formation of galaxies. Perhaps the most ambitious approach for doing so is to run high-resolution hydrodynamic simulations. However, even the most advanced simulations to-date are only able to resolve masses down to $\sim 10^7\, M_\odot$ and are thus forced to invoke add-hoc "sub-grid physics" to match observational data. Computational cost aside, the heuristic models invoked by pure simulation-based approaches will forever leave open the question: "for a given survey, how flexible of a subgrid prescription is 'flexible enough' to accurately estimate cosmological parameters?"

An alternative and more principled approach to model observed tracers of the matter, known as the *bias expansion*, was originally introduced by McDonald and Roy [395]. A key insight to the bias expansion is that while galaxy formation occurs over long timescales ($t \sim H^{-1}$) the spatial region responsible for developing a galaxy's host dark matter halo is small compared to the Hubble scale. This idea is illustrated in Fig. 1.4, where the shaded tube with comoving (spatial) size $R_* \ll (aH)^{-1}$ corresponds to the region of spacetime capable of influencing the evolution of a proto-galaxy along its trajectory from initial position $\boldsymbol{x}_{\text{fl}}(0)$ to its observed position $\boldsymbol{x}_{\text{fl}}(\tau)$. During matter-domination the only relevant (observable)

dynamical variables are the tidal tensor $\nabla_i\nabla_j\Phi$ and velocity gradients $\nabla_i v_j$. Thus we can think of the galaxy density at position $\boldsymbol{x}_{\rm fl}(\tau)$ as a functional $F[\nabla_i\nabla_j\Phi,\, \nabla_i v_j]$, where $\nabla_i\nabla_j\Phi$ and $\nabla_i v_j$ are evaluated anywhere within a spatial distance $\sim R_*$ from the worldline of the galaxy. Provided that we work in the limit where $\delta_m \propto \nabla^2\Phi$ and kR_* are small we can expand the functional as a double series in perturbations and derivatives. In the context of Lagrangian perturbation theory this is accomplished by weighting the delta function appearing in Eq. (1.13) by the functional F, which to second order in perturbations and first order in derivatives yields

$$\begin{aligned} 1+\delta_g(\boldsymbol{x},t) &= \int d^3\boldsymbol{q}\, F(\boldsymbol{q})\, \delta^D\big(\boldsymbol{x}-\boldsymbol{q}-\boldsymbol{\Psi}(\boldsymbol{q},t)\big) \\ F(\boldsymbol{q}) &= 1 + b_1\delta_m(\boldsymbol{q}) + b_2\delta_m^2(\boldsymbol{q}) + b_s s^2(\boldsymbol{q}) + b_{\nabla^2}\nabla^2\delta_m(\boldsymbol{q}) + \epsilon(\boldsymbol{q}), \end{aligned} \tag{1.21}$$

where we've introduced the shear field $s_{ij} \equiv (\nabla_i\nabla_j\nabla^{-2} - \delta^K_{ij}/3)\delta_m$ and a contribution ϵ to account for small-scale stochasticity (i.e. shot noise). The time-dependent coefficients appearing in the bias expansion (Eq. 1.21) encode the large-scale impact of the small-scale physics responsible for galaxy formation. These bias coefficients must be scalar quantities to ensure that all observers agree on their values regardless of their position or orientation (i.e. the Copernican principle). This requirement forces the operators (i.e. products of $\nabla_i\nabla_j\Phi$, $\nabla_i v_j$ and their derivatives) allowed in the bias expansion to also be scalars, simplifying the expansion significantly at any given order. Marginalizing over these coefficients immunizes cosmological analyses to systematic uncertainties associated with the small-scale physics relevant for galaxy formation.

To capture the evolution of δ_g, one can use Lagrangian perturbation theory (LPT) as a systematic framework to evolve the displacement field $\boldsymbol{\Psi}$ into the mildly nonlinear regime. However, for low-bias tracers at late times, the scale associated with nonlinear gravitational evolution (i.e. the rms displacement of a dark matter particle) can be significantly larger than the length scale R_* associated with the bias expansion. In this case, a hybrid effective field theory (HEFT) approach [412] − which replaces the perturbative LPT displacements in Eq. (1.21) with fully nonlinear displacements $\boldsymbol{\Psi}$ from N-body simulations − can extract more cosmological information than LPT alone, thanks to its ability to accurately model a broader range of scales. Fig. 1.5 illustrates this strategy by showing various operators (pink/yellow) in the bias expansion after being advected using simulation-based displacements. An appropriate linear combination of the displaced operators can be used to accurately recover the non-linear tracer of interest, which in Fig. 1.5 is the high-redshift 21 cm density field (black/yellow).

An additional strength of the Lagrangian formalism is its natural treatment of *redshift-space distortions* (RSD), which arise due to peculiar velocities along the line of sight. These distortions can be consistently incorporated at the level of the displacement field via

$$\boldsymbol{\Psi}_s = \boldsymbol{\Psi} + \frac{\hat{\boldsymbol{n}}(\boldsymbol{v}\cdot\hat{\boldsymbol{n}})}{a\,H}, \tag{1.22}$$

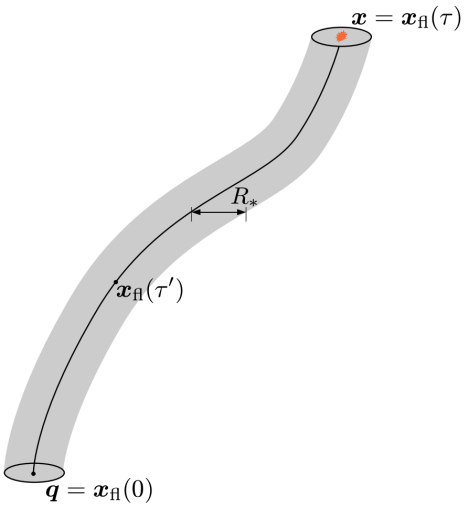


Figure 1.4: Illustration of the spacetime region influencing galaxy formation (shaded tube). Time flows upward while the horizontal axis corresponds to spatial hypersurfaces. The line connecting Lagrangian coordinate $\boldsymbol{q} = \boldsymbol{x}_{\rm fl}(0)$ to the observed Eulerian position $\boldsymbol{x} = \boldsymbol{x}_{\rm fl}(\tau)$ represents the worldline of the galaxy. The small width $R_* \ll (aH)^{-1}$ of the tube reflects the assumption that galaxy formation is only influenced by physics within a limited spatial neighborhood. Figure from [189].

where $\hat{\boldsymbol{n}}$ is the line-of-sight direction. This mapping enables a unified treatment of both redshift-space and projected clustering, which is essential for precision cosmology from spectroscopic surveys and their cross-correlation with CMB secondaries.

1.4 Cosmological tensions

Despite the remarkable success of the standard ΛCDM model in explaining a broad range of cosmological observations, several persistent tensions between different datasets have emerged. These discrepancies have sparked considerable interest as they could hint at physics beyond the ΛCDM model or perhaps unresolved systematic errors.

One of the most prominent tensions is the *Hubble constant* (H_0) discrepancy. Measurements of H_0 derived from the cosmic microwave background (CMB) anisotropies, particularly from the *Planck* satellite under the assumption of ΛCDM [475], yield a value of about $67.4 \pm 0.5\,\mathrm{km\ s^{-1}Mpc^{-1}}$. In contrast the SH0ES collaboration [500] finds a significantly higher value near $73\,\mathrm{km\ s^{-1}Mpc^{-1}}$ from local measurements of Cepheid-calibrated supernovae. This ~ 4–5σ tension has persisted and even strengthened with newer data, however, a detailed discussion of the Hubble tension and proposed resolutions is beyond the scope of this thesis and will not be discussed further here.

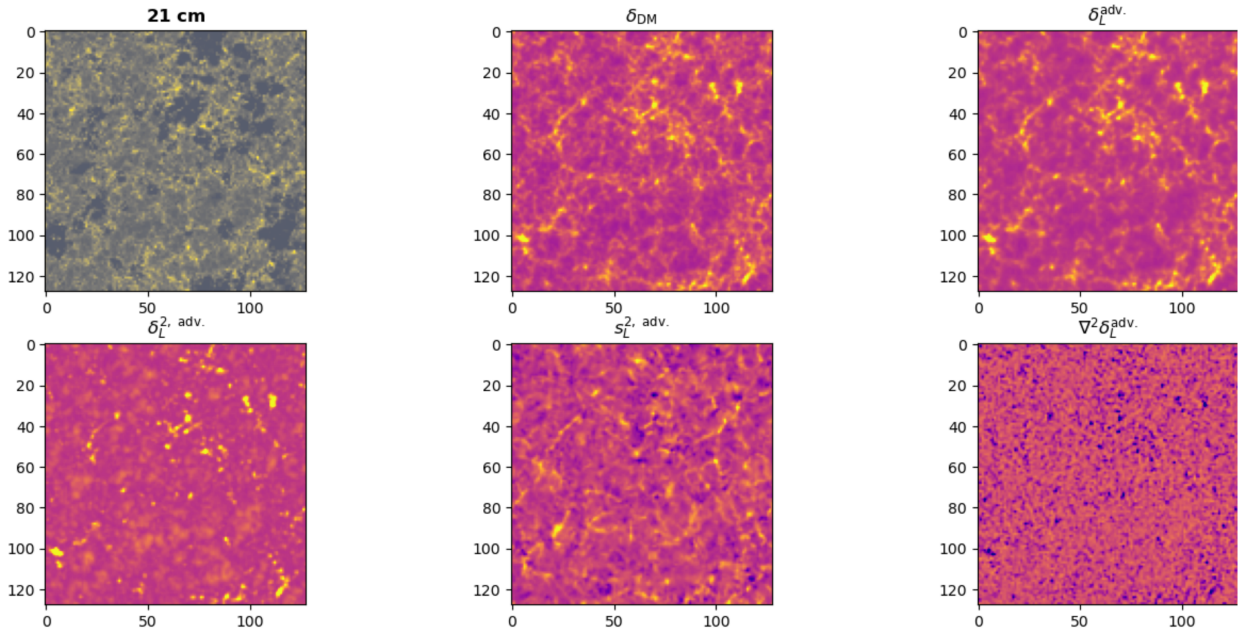


Figure 1.5: Comparison of advected operators appearing in the bias expansion with the fully nonlinear 21 cm density field from simulations. Figure from [58].

Another key tension involves the amplitude of matter fluctuations, often quantified by the parameter $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$, where σ_8 is the RMS of matter fluctuations on $8\,h^{-1}$ Mpc scales and Ω_m is the matter density parameter. CMB-derived constraints from *Planck* favor a higher S_8 value compared to previous measurements from weak gravitational lensing surveys such as the Dark Energy Survey (DES) [6] and the Kilo-Degree Survey (KiDS) [286], although more recent analyses that take baryonic feedback into account find better consistency with Planck [567]. In Part II we address this tension with CMB lensing cross-correlations, which have the advantage of being largely immune to the small-scale systematic uncertainties present in galaxy-lensing analyses.

Secondary anisotropies in the CMB provide complementary approaches to measure the baryon distribution around galaxies and clusters. Measurements of the kinetic Sunyaev-Zel'dovich (kSZ) effect stacked around DESI galaxies suggest larger baryonic feedback and hence a smoother baryon profiles than was previously believed to be possible [265]. In contrast, previous analyses of patchy screening point to more pronounced spatial fluctuations in the baryon density field [536, 154], which at face value are in $\sim 7\sigma$ tension with the profiles inferred from the kSZ effect. Reconciling this apparent discrepancy is the subject of chapters 7 and 8.

Recent baryon acoustic oscillation (BAO) data from the Dark Energy Spectroscopic Instrument (DESI), when analyzed alongside current CMB data, show a slight preference for evolving dark energy [182] over a cosmological constant. The discovery of evolving dark

energy would revolutionize our basic understanding of fundamental physics. However, as we discuss in detail in chapter 9, the significance of this preference crucially hinges on the accuracy of the optical depth to reionization inferred from large-scale CMB polarization measurements, whose precision is currently systematics-limited.

While these tensions have generated widespread interest in the community, the analyses in this thesis suggest that several of them may be less robust than commonly portrayed. The S_8 discrepancy, for example, is shown to be absent in cross-correlations of CMB lensing with large-scale structure tracers. Likewise, the tension between kSZ and patchy screening measurements is shown to be sourced from unaccounted for lensing biases, whereas the preference for evolving dark energy is argued to instead indicate systematics in large-scale polarization measurements. Through prioritizing cosmological data that are both observationally clean and well understood from a theoretical perspective, the results presented here argue for a more cautious interpretation these tensions.

1.5 Scope and outline of this thesis

This thesis is organized into three parts, each addressing a key aspect of extracting accurate cosmological information from secondary anisotropies in the cosmic microwave background (CMB) and their cross-correlations with large-scale structure (LSS) surveys.

Part I (§2–§4) develops and analyzes techniques to mitigate astrophysical foreground contamination in CMB lensing reconstruction. Foregrounds such as the thermal Sunyaev-Zel'dovich (tSZ) effect and cosmic infrared background (CIB) bias lensing estimators and can significantly affect both auto- and cross-correlation measurements. Chapter 2 introduces geometrically motivated bias-hardened quadratic estimators that reduce foreground bias while maintaining low reconstruction noise. Chapter 3 extends these methods to polarization-based estimators, including large-lens-suppressed and shear-only variants. Chapter 4 investigates optimal multifrequency weighting schemes to minimize lensing bias and noise using internal linear combination (ILC) techniques, including the deprojection of specific foreground components.

Part II (§5–§9) applies the tools developed in Part I to perform cosmological inference using cross-correlations between CMB secondaries and galaxy surveys. Chapter 5 presents tomographic measurements of the growth of structure using cross-correlations between CMB lensing maps and DESI luminous red galaxies (LRGs). Chapter 6 complements this analysis with cross-correlations of CMB lensing and galaxy magnification using the DESI bright galaxy survey (BGS). Chapters 7 and 8 explore the distribution of baryons through the patchy screening and the kinetic Sunyaev-Zel'dovich (kSZ) effects, presenting the first self-consistent comparisons between the two. Chapter 9 investigates the role of the optical depth to reionization in current tensions between CMB and baryonic acoustic oscillation (BAO) measurements.

Part III (§10–§11) looks toward the next generation of high-redshift cosmological surveys. Chapter 10 forecasts the potential of surveys such as MegaMapper and PUMA to

constrain ΛCDM and several extensions. These forecasts incorporate both galaxy clustering and cross-correlations with future CMB lensing datasets. Chapter 11 introduces a novel method to measure the optical depth using the clustering of the 21 cm signal from the epoch of reionization, providing a new path to break degeneracies in CMB-based cosmological inference.

Taken together, the work presented in this thesis demonstrates the utility of secondary CMB anisotropies and their cross-correlations with LSS in advancing precision cosmology while maintaining accurate inference. It offers methodological innovations that improve robustness to systematic contamination, presents measurements and subsequent analyses that address current cosmological tensions, and provides forecasts to motivate and guide future observational efforts.

Part I

Foreground mitigation in CMB lensing measurements

Chapter 2

Geometric foreground mitigation I: tSZ profile-hardening for temperature based estimators

This chapter was originally published as

> Noah Sailer, Emmanuel Schaan, and Simone Ferraro. "Lower bias, lower noise CMB lensing with foreground-hardened estimators". In: *Phys. Rev. D* 102.6 (2020), p. 063517. DOI: 10.1103/PhysRevD.102.063517. arXiv: 2007.04325 [astro-ph.CO]

Extragalactic foregrounds in temperature maps of the Cosmic Microwave Background (CMB) severely limit the ability of standard estimators to reconstruct the weak lensing potential. These foregrounds are not fully removable by multi-frequency cleaning or masking and can lead to large biases if not properly accounted for. For foregrounds made of a number of unclustered point sources, an estimator for the source amplitude can be derived and deprojected, removing any bias to the lensing reconstruction. We show with simulations that all of the extragalactic foregrounds in temperature can be approximated by a collection of sources with identical profiles, and that a simple bias hardening technique is effective at reducing any bias to lensing, at a minimal noise cost. We compare the performance and bias to other methods such as "shear-only" reconstruction, and discuss how to jointly deproject any arbitrary number of foregrounds, each with an arbitrary profile. In particular, for a Simons Observatory-like experiment foreground-hardened estimators allow us to extend the maximum multipole used in the reconstruction, increasing the overall statistical power by $\sim 50\%$ over the standard quadratic estimator, both in auto and cross-correlation. We conclude that source hardening outperforms the standard lensing quadratic estimator both in auto and cross-correlation, and in terms of lensing signal-to-noise and foreground bias.

2.1 Introduction

CMB lensing [342, 279] is rapidly becoming one of the most powerful tools in a cosmologist's toolkit. With the ability to measure matter fluctuations to high redshift, no photometric redshift uncertainties and its sensitivity to crucial cosmological parameters, sub-percent precision measurements have the potential to greatly increase our understanding of the Universe.

Current and upcoming wide-field CMB experiments such as AdvACT [283], SPT-3G [80] and Simons Observatory [22] will heavily rely on reconstruction from CMB temperature fluctuations (as opposed to polarization). For future polarization-dominated experiments like CMB-S4 [3] and CMB-HD [540, 542], temperature lensing will still contribute a non-negligible part of the total signal-to-noise.

Temperature maps contain significant contamination from extragalactic foregrounds such as the thermal and kinematic Sunyaev-Zel'dovich effects (tSZ and kSZ), the Cosmic Infrared Background (CIB), and both radio and infrared point sources. Lensing reconstruction on contaminated maps creates significant biases in both the auto- and cross-correlation of the reconstructed field, which in turn can bias our inference of cosmology [597, 446, 374, 227, 524, 71]. These biases can be up to several tens of percent for upcoming surveys, thus severely limiting our ability to use CMB lensing for precision cosmology.

Several techniques have been proposed and implemented to mitigate the impact of extragalactic foregrounds. Some are based on symmetries, such as the recently proposed shear-only reconstruction [524] (and related generalizations such as hybrid and general multipole estimators), while others make use of multi-frequency maps to reduce the impact of particular foregrounds [374, 165].

Here we revisit and extend a bias hardening technique first proposed in [446, 434] and applied to real data in [21], and compare it to shear-only reconstruction and the hybrid estimator of [524]. We show that if we approximate extragalactic foregrounds as a collection of sources with identical profiles, an estimator for their amplitude can be obtained and hardened against in an optimal way to mitigate their impact on the reconstructed lensing convergence. With the use of realistic and correlated simulations of all of the foregrounds, we show that this approximation is excellent and leads to a very large suppression in bias, allowing for the use of considerably smaller scales in the reconstruction.

The remainder of the chapter is organized as follows: In Section 2.2 we review the general formalism for bias hardening against a single foreground, then generalize the procedure to an arbitrary number of foregrounds, each with arbitrary profiles. In Section 2.3 we specialize to the case where foregrounds are point sources, and in Section 2.4 we numerically explore the biases and the performance of mitigation through bias hardening or shear-only estimators. We conclude with Section 2.5.

2.2 Foreground-hardened quadratic estimators

In this section, we first review the derivation of the standard lensing quadratic estimator (QE), and rephrase it in terms of the linear response of the map covariance to the lensing convergence. We build on this formalism to construct foreground estimators, and finally null the linear response of the lensing QE to these foregrounds, as in [446, 434].

Lensing quadratic estimator, bispectrum & linear response

The observed CMB temperature T is the sum of the *lensed* primary CMB $T^{\rm CMB}$ and the foregrounds s, which will be approximated by a collection of sources in this chapter: $T = T^{\rm CMB} + s$. The lensed CMB temperature $T^{\rm CMB}$ field is statistically invariant under translations. As a result, the off-diagonal covariance $\langle T^{\rm CMB}_{\boldsymbol{\ell}} T^{\rm CMB}_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle$ with $\boldsymbol{L} \neq 0$ is zero. However, fixing the lensing convergence Fourier mode $\kappa_{\boldsymbol{L}}$ breaks this statistical homogeneity, and produces non-zero off-diagonal covariances:

$$\begin{aligned} \langle T^{\rm CMB}_{\boldsymbol{\ell}} T^{\rm CMB}_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle_{\text{at fixed } \kappa_{\boldsymbol{L}}} &= f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\kappa_{\boldsymbol{L}} + \mathcal{O}\left(\kappa^2_{\boldsymbol{L}}\right) \\ \text{with} \quad f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &\equiv \frac{2\boldsymbol{L}}{L^2}\cdot\left[\boldsymbol{\ell} C^0_\ell + (\boldsymbol{L}-\boldsymbol{\ell}) C^0_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right], \end{aligned} \tag{2.1}$$

where C^0_ℓ is the *unlensed* power spectrum. From this, we obtain an unbiased lensing estimator (to first order in κ) from the ratio $T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}/f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$. The standard lensing QE [300] is simply the inverse-variance weighted average of all these "building block" estimators, summing over $\boldsymbol{\ell}$ at fixed $\boldsymbol{L}$. This leads to the usual result:

$$\begin{aligned} \hat{\kappa}_{\boldsymbol{L}} &= N^{\kappa}_{\boldsymbol{L}} \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}} \\ (N^{\kappa}_{\boldsymbol{L}})^{-1} &= \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \\ F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2 C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}. \end{aligned} \tag{2.2}$$

In order to generalize this construction to foreground quadratic estimators, we consider Eq. (2.1) and note that $f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ can be seen as the linear response of the two point correlator $\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle_{\text{at fixed } \kappa_{\boldsymbol{L}}}$ to the lensing convergence mode $\kappa_{\boldsymbol{L}}$:

$$f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{\delta\langle T^{\rm CMB}_{\boldsymbol{\ell}} T^{\rm CMB}_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle_{\text{at fixed } \kappa_{\boldsymbol{L}}}}{\delta\kappa_{\boldsymbol{L}}}(\kappa_{\boldsymbol{L}} = 0). \tag{2.3}$$

Furthermore, by multiplying both sides of Eq. (2.1) with $\kappa_{-\boldsymbol{L}}$, we see that this linear response can be expressed as the ratio of the following bispectrum and power spectrum:

$$f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{\langle T^{\rm CMB}_{\boldsymbol{\ell}} T^{\rm CMB}_{\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{-\boldsymbol{L}}\rangle}{\langle \kappa_{\boldsymbol{L}} \kappa_{-\boldsymbol{L}}\rangle}. \tag{2.4}$$

This final expression is our starting point to derive foreground quadratic estimators.

Foreground quadratic estimators

Following the argument above, anytime a field s (e.g. a foreground, the mask, the beam) has a non-zero bispectrum $\langle s_{\boldsymbol{\ell}} s_{\boldsymbol{L}-\boldsymbol{\ell}} s_{-\boldsymbol{L}}\rangle$, we can define the following linear response[1]:

$$f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{\langle s_{\boldsymbol{L}} s_{\boldsymbol{L}-\boldsymbol{\ell}} s_{-\boldsymbol{L}}\rangle}{\langle s_{\boldsymbol{L}} s_{-\boldsymbol{L}}\rangle} \tag{2.5}$$

such that

$$\langle s_{\boldsymbol{\ell}} s_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle_{\text{at fixed } s_{\boldsymbol{L}}} = f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s_{\boldsymbol{L}} + \mathcal{O}\left((s_{\boldsymbol{L}})^2\right). \tag{2.6}$$

We can then define the "building block" foreground estimator, unbiased to first order in the foreground s, via $T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}/f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$. Finally, we inverse-variance weight these estimators for all $\boldsymbol{\ell}$, at fixed $\boldsymbol{L}$, to obtain the standard minimum variance quadratic estimator for s. In practice, evaluating this estimator requires knowledge of the foreground non-Gaussianity, through its bispectrum. We note however that no assumption about the foreground trispectrum is used.

Bias-hardened lensing quadratic estimators

In this section we review the derivation of the bias-hardened (BH) lensing estimator of [446, 434] for a single foreground, relate the noise of the BH estimator to the noise of the standard minimum variance QE, and generalize the procedure for an arbitrary number of foregrounds.

Hardening against a single foreground

In the presence of a single foreground s, the off-diagonal covariances of the observed temperature map take the form

$$\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle = f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{\boldsymbol{L}} + f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s_{\boldsymbol{L}} \tag{2.7}$$

to lowest order in κ and s. In this equation the average is taken at fixed $\kappa_{\boldsymbol{L}}$ and $s_{\boldsymbol{L}}$, which will be implied from here on out for notational economy. Since both $\kappa_{\boldsymbol{L}}$ and $s_{\boldsymbol{L}}$ produce off-diagonal covariances, the standard quadratic estimators for κ and s acquire a bias. These biases can be written in the form [446]:

$$\begin{pmatrix}\langle\hat{\kappa}_{\boldsymbol{L}}\rangle\\ \langle\hat{s}_{\boldsymbol{L}}\rangle\end{pmatrix} = \begin{pmatrix}1 & N^{\kappa}_{\boldsymbol{L}}\mathcal{R}_{\boldsymbol{L}}\\ N^s_{\boldsymbol{L}}\mathcal{R}_{\boldsymbol{L}} & 1\end{pmatrix}\begin{pmatrix}\kappa_{\boldsymbol{L}}\\ s_{\boldsymbol{L}}\end{pmatrix} \tag{2.8}$$

where the response $\mathcal{R}_{\boldsymbol{L}}$ is defined to be

$$\mathcal{R}_{\boldsymbol{L}} = \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} \frac{f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\text{tot}}_{\ell} C^{\text{tot}}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}. \tag{2.9}$$

[1] We note that this expression only accounts for the non-Gaussianity of the sources. In principle, the sources are themselves lensed which results in an additional smaller correlation and κ dependence of the mask [410]. For simplicity, we ignore this effect in this chapter.

We see that the bias to $\hat{\kappa}$ is proportional to s, and vice-versa. This allows us to create new bias-hardened estimators, which are the unique linear combinations of $\hat{\kappa}$ and $\hat{s}$ that null these biases

$$\begin{pmatrix} \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \\ \hat{s}^{\rm BH}_{\boldsymbol{L}} \end{pmatrix} = \begin{pmatrix} 1 & N^{\kappa}_{\boldsymbol{L}}\mathcal{R}_{\boldsymbol{L}} \\ N^{s}_{\boldsymbol{L}}\mathcal{R}_{\boldsymbol{L}} & 1 \end{pmatrix}^{-1} \begin{pmatrix} \hat{\kappa}_{\boldsymbol{L}} \\ \hat{s}_{\boldsymbol{L}} \end{pmatrix}. \tag{2.10}$$

By taking an average of Eq. (2.10), and inserting Eq. (2.8) on the right hand side, we see that bias-hardened estimators indeed satisfy $\langle \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \rangle = \kappa_{\boldsymbol{L}}$ and $\langle \hat{s}^{\rm BH}_{\boldsymbol{L}} \rangle = s_{\boldsymbol{L}}$.

Noise of the hardened lensing QE

By construction, the weights of the standard QE are chosen to minimize its variance. Since the BH estimator's weights differ from the standard QE's, bias-hardening comes at a price in noise, which we quantify in this section. We start with Eq. (2.10), the definition of the BH estimator, from which we compute the variance

$$\left\langle \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \hat{\kappa}^{\rm BH}_{-\boldsymbol{L}} \right\rangle = \Big[\langle \hat{\kappa}_{\boldsymbol{L}} \hat{\kappa}_{-\boldsymbol{L}} \rangle + \left(N^{\kappa}_{\boldsymbol{L}}\mathcal{R}_{\boldsymbol{L}}\right)^2 \langle \hat{s}_{\boldsymbol{L}} \hat{s}_{-\boldsymbol{L}} \rangle - 2N^{\kappa}_{\boldsymbol{L}}\mathcal{R}_{\boldsymbol{L}} \mathrm{Re}\, \langle \hat{\kappa}_{\boldsymbol{L}} \hat{s}_{-\boldsymbol{L}} \rangle \Big] \Big/ \left(1 - N^{\kappa}_{\boldsymbol{L}} N^{s}_{\boldsymbol{L}} \mathcal{R}^2_{\boldsymbol{L}}\right)^{-2}. \tag{2.11}$$

In this equation we've used the identity $N^{\kappa}_{\boldsymbol{L}} = N^{\kappa}_{-\boldsymbol{L}}$, and we've assumed that $f^{s}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ is even in both of its arguments, which forces both the noise of the source estimator and the response to be even functions of $\boldsymbol{L}$. This assumption is valid when modeling the field s as a collection of individual sources with profiles $u_{\boldsymbol{\ell}}$ (since $u_{\boldsymbol{\ell}} = u_{-\boldsymbol{\ell}}$ for any physical emission profile), as in section 2.3.

The cross-correlation of $\hat{\kappa}$ and $\hat{s}$ is proportional to the response function, and is explicitly given by

$$\left\langle \hat{\kappa}_{\boldsymbol{L}} \hat{s}_{\boldsymbol{L}'} \right\rangle = (2\pi)^2 \delta^{D}_{\boldsymbol{L}+\boldsymbol{L}'} N^{\kappa}_{\boldsymbol{L}} N^{s}_{\boldsymbol{L}} \mathcal{R}_{\boldsymbol{L}}. \tag{2.12}$$

Combining Eq. (2.11) and (2.12) results in the simple expression

$$N^{\kappa^{\rm BH}}_{\boldsymbol{L}} = N^{\kappa}_{\boldsymbol{L}} \left(1 - N^{\kappa}_{\boldsymbol{L}} N^{s}_{\boldsymbol{L}} \mathcal{R}^2_{\boldsymbol{L}}\right)^{-1}. \tag{2.13}$$

We see that the noise cost for bias-hardening is completely determined by the determinant of the matrix in Eq. (2.10). The value of this determinant is plotted in Fig. 2.1 for the case of point source foregrounds ($f^{s}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = 1$). Nearly identical calculations yield the noise of the bias-hardened source estimator and the noise cross-correlation of the two bias-hardened estimators. The results are

$$\begin{aligned} N^{s^{\rm BH}}_{\boldsymbol{L}} &= N^{s}_{\boldsymbol{L}} \left(1 - N^{\kappa}_{\boldsymbol{L}} N^{s}_{\boldsymbol{L}} \mathcal{R}^2_{\boldsymbol{L}}\right)^{-1} \\ \langle \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \hat{s}^{\rm BH}_{\boldsymbol{L}} \rangle &= -\langle \hat{\kappa}_{\boldsymbol{L}} \hat{s}_{\boldsymbol{L}} \rangle \left(1 - N^{\kappa}_{\boldsymbol{L}} N^{s}_{\boldsymbol{L}} \mathcal{R}^2_{\boldsymbol{L}}\right)^{-1}. \end{aligned} \tag{2.14}$$

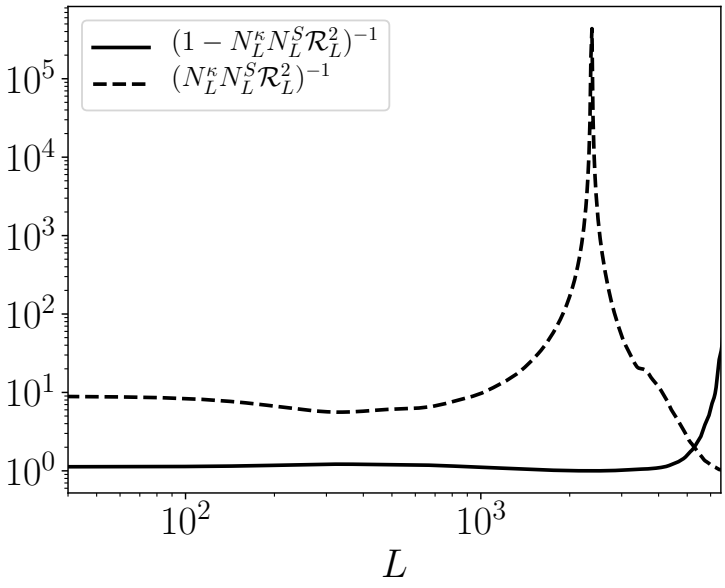


Figure 2.1: Bias-hardening against point source foregrounds. The solid curve is the determinant of the matrix in Eq. (2.10), which is the ratio of the bias-hardened estimator's noise to the standard QE's noise. It is close to unity on most scales, implying that the noise penalty for bias hardening is minimal. The dashed curve is the ratio of the SNR of the point source estimator to the relative bias to C_L^κ from the source trispectrum. This ratio is roughly 10 at all scales. Thus in the regime where the bias to the standard QE is significant, the source estimator reconstructs the sources with a high SNR, and the source-hardened estimator can remove the bias due to sources with little noise cost. The spike is due to a zero crossing in the response $\mathcal{R}_L$ at $L \sim 2500$.

Simultaneously hardening against n foregrounds

It is straightforward to generalize the bias-hardening procedure for n foreground fields $s_{1,\boldsymbol{L}}, \cdots, s_{n,\boldsymbol{L}}$. In this case the observed map has off-diagonal covariances

$$\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle = f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{\boldsymbol{L}} + f^{s_1}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s_{1,\boldsymbol{L}} + \cdots + f^{s_n}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s_{n,\boldsymbol{L}} \tag{2.15}$$

to lowest order in κ and the foreground fields. Just as before, this average is taken at fixed $\kappa_{\boldsymbol{L}}, s_{1,\boldsymbol{L}}, \cdots, s_{n,\boldsymbol{L}}$. First, we write down the standard minimum variance quadratic estimators

for each of the fields:

$$\begin{aligned}\hat{X}_{\boldsymbol{L}} &= N^X_{\boldsymbol{L}} \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}} \\ \left(N^X_{\boldsymbol{L}}\right)^{-1} &= \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \\ F^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{f^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}},\end{aligned} \tag{2.16}$$

where $X = \kappa, s_1, \cdots, s_n$. We then write down the $(n+1)\times(n+1)$ matrix that quantifies the biases to all the estimators, analogous to Eq. (2.8). Finally, we invert this matrix to obtain the bias-hardened estimators

$$\begin{pmatrix} \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \\ \hat{s}^{\rm BH}_{1,\boldsymbol{L}} \\ \vdots \\ \hat{s}^{\rm BH}_{n,\boldsymbol{L}} \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & N^\kappa_{\boldsymbol{L}}\mathcal{R}^{\kappa,s_1}_{\boldsymbol{L}} & \cdots & N^\kappa_{\boldsymbol{L}}\mathcal{R}^{\kappa,s_n}_{\boldsymbol{L}} \\ N^{s_1}_{\boldsymbol{L}}\mathcal{R}^{s_1,\kappa}_{\boldsymbol{L}} & 1 & \cdots & N^{s_1}_{\boldsymbol{L}}\mathcal{R}^{s_1,s_n}_{\boldsymbol{L}} \\ \vdots & \vdots & \vdots & \vdots \\ N^{s_n}_{\boldsymbol{L}}\mathcal{R}^{s_n,\kappa}_{\boldsymbol{L}} & N^{s_n}_{\boldsymbol{L}}\mathcal{R}^{s_n,s_1}_{\boldsymbol{L}} & \cdots & 1 \end{pmatrix}^{-1}}_{\equiv \boldsymbol{M}^{-1}_{\boldsymbol{L}}} \begin{pmatrix} \hat{\kappa}_{\boldsymbol{L}} \\ \hat{s}_{1,\boldsymbol{L}} \\ \vdots \\ \hat{s}_{n,\boldsymbol{L}} \end{pmatrix} \tag{2.17}$$

where the generalized response is defined to be

$$\mathcal{R}^{X,Y}_{\boldsymbol{L}} = \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} \frac{f^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^Y_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}. \tag{2.18}$$

Just as in the previous section, the response $\mathcal{R}^{X,Y}_{\boldsymbol{L}}$ is proportional to the cross-correlation of the two fields X and Y. When $X = Y$, the response is simply the inverse of the noise, which is manifested in Eq. (2.17) by the 1's on the diagonal. Writing the matrix $\boldsymbol{M}_{\boldsymbol{L}}$ as

$$\boldsymbol{M}_{\boldsymbol{L}} = \begin{pmatrix} 1 & \boldsymbol{w}^T_{\boldsymbol{L}} \\ \boldsymbol{v}_{\boldsymbol{L}} & \boldsymbol{A}_{\boldsymbol{L}} \end{pmatrix}, \tag{2.19}$$

where $\boldsymbol{A}_{\boldsymbol{L}}$ is a $n \times n$ matrix and $\boldsymbol{w}_{\boldsymbol{L}}, \boldsymbol{v}_{\boldsymbol{L}}$ are n-dimensional column vectors, we find that the noise of the BH lensing estimator is simply

$$N^{\kappa^{\rm BH}}_{\boldsymbol{L}} = \frac{\det(\boldsymbol{A}_{\boldsymbol{L}})}{\det(\boldsymbol{M}_{\boldsymbol{L}})} N^\kappa_{\boldsymbol{L}}. \tag{2.20}$$

A derivation of this equation is given in Appendix A.3.

2.3 Source QE and source-hardened lensing QE

Halo model for sources

In the previous section we reviewed how to bias-harden against any number of foregrounds. The only input to this procedure is the linear response $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ of each foreground, which can

be calculated if the foreground's bispectrum and power spectrum are known. To estimate these quantities we adopt a halo model, writing each foreground $s(\boldsymbol{x})$ as a sum of individual sources with profiles $u(\boldsymbol{x})$:

$$s(\boldsymbol{x}) = \sum_i s_i u(\boldsymbol{x} - \boldsymbol{x}_i), \tag{2.21}$$

where s_i and $\boldsymbol{x}_i$ are the flux and position of the i'th source, and the profile $u(\boldsymbol{x})$ is assumed to be source independent. These sources are assumed to be a Poisson sampling of the matter density field. Because they trace matter on large scales, the source field s is correlated with the true lensing convergence field. As we show below, this correlation causes a bias in the CMB lensing cross- and auto-correlations. With this model, the linear response takes the form

$$f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{\langle s_{\boldsymbol{\ell}} s_{\boldsymbol{L}-\boldsymbol{\ell}} s_{-\boldsymbol{L}}\rangle}{\langle s_{\boldsymbol{L}} s_{-\boldsymbol{L}}\rangle} = \frac{\langle s_i^3\rangle}{\langle s_i^2\rangle}\frac{u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}}{u_{\boldsymbol{L}}}, \tag{2.22}$$

where in the second equality we assumed that the sources are sparse, such that the bispectrum and power spectrum are both shot noise dominated, and the source clustering is negligible. In this case the linear response simplifies to a separable function of $\boldsymbol{\ell}$ and $\boldsymbol{L}-\boldsymbol{\ell}$, and is proportional to the ratio of the third to the second moment of the individual source fluxes.

We can choose the normalization of the source estimator $\hat{s}$ to remove the constant factor $\langle s_i^3\rangle/\langle s_i^2\rangle$. Doing so makes $\hat{s}$ an estimator of $s' \equiv s\langle s_i^2\rangle/\langle s_i^3\rangle$, which has linear response $f^{s'}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}/u_{\boldsymbol{L}}$. This simplifies the expression of the source estimator to

$$\begin{aligned} \hat{s}_{\boldsymbol{L}} &= N^s_{\boldsymbol{L}} \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} \frac{T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \frac{u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}}{u_{\boldsymbol{L}}} \\ (N^s_{\boldsymbol{L}})^{-1} &= \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} \frac{1}{2C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \left[\frac{u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}}{u_{\boldsymbol{L}}}\right]^2, \end{aligned} \tag{2.23}$$

which can be quickly evaluated using FFT. We emphasize that the weights and normalization of the BH estimator are left unchanged by this renormalization of $\hat{s}$. Therefore the bias-hardening procedure requires no knowledge of the source amplitude, and the normalization of the profile u is completely arbitrary. For this reason we will drop the primes from here-on-out, defining $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}/u_{\boldsymbol{L}}$.

Noise cost and importance of the non-Gaussianity

Consider a single source component s. If s were Gaussian, it would be indistinguishable from map noise, and the source estimator $\hat{s}$ would not be able to pick it out. Here, we show how the non-Gaussianity of the source map controls both the signal-to-noise (SNR) of the source estimator, and the bias from sources to CMB lensing. For simplicity let s be a Poisson distributed point source field ($f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = u_{\boldsymbol{\ell}} = 1$). In this case the mean number of sources $\bar{n}$ determines the non-Gaussianity; as $\bar{n}$ goes to infinity, the statistics converges to that of a Gaussian.

The point source estimator is an unbiased (to lowest order) estimator of s, so the SNR per Fourier mode is simply :

$$\mathrm{SNR}^s_L = \langle \hat{s}_{\boldsymbol{L}} \hat{s}_{-\boldsymbol{L}} \rangle / N^s_L \tag{2.24}$$
$$= \mathcal{T}^s / N^s_L. \tag{2.25}$$

Thus the SNR for the source estimator scales with the scale-independent source trispectrum $\mathcal{T}^s$: it is larger for low source densities which corresponds to a higher non-Gaussianity. As the number of sources increases, the field becomes increasingly Gaussian, and can no longer be distinguished from Gaussian noise in the map. The expression above is valid for mildly non-Gaussian source field, as is the case in reality. For the very highly non-Gaussian case, which is not relevant for the CMB, the source non-Gaussianity contributes also to the noise, leading to a saturation of the SNR to a finite value as $\bar{n}$ goes to zero.

Even if the source field and lensing convergence were independent, point sources would cause a bias to the standard QE due to its trispectrum:

$$\frac{\text{bias to } C^\kappa_L \text{ from } \mathcal{T}^s}{N^\kappa_L} = \frac{\mathcal{T}^s}{N^\kappa_L} \left(N^\kappa_L \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^\kappa_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \right)^2 = \mathcal{T}^s N^\kappa_L \mathcal{R}^2_L. \tag{2.26}$$

Here again, the source non-Gaussianity controls the bias to lensing via the $1/\bar{n}$ scaling. From equations (2.25) and (2.26), we see that the ratio of the SNR of the source estimator to the relative bias to lensing from the source trispectrum is $(N^\kappa_L N^S_L \mathcal{R}^2_L)^{-1}$. This ratio cancels out the non-Gaussianity of the sources and is around 10 for most scales of interest, as shown in Fig 2.1. This means that subtracting the lensing bias due to the source trispectrum will always cause only a small loss in lensing SNR. Indeed, as shown in Fig 2.1 (solid line), the lensing noise for the standard QE and the bias hardened estimators are very similar, meaning that subtracting the source bias causes a negligible increase in the lensing noise. This is consistent with the findings of [221].

Finally, the non-Gaussianity of the point sources also contributes an extra noise term in the lensing reconstruction, which we haven't included. However, this noise is only important when the bias to CMB lensing is also important. Thus, by making sure that the non-Gaussian point sources (and other foregrounds) do not significantly bias the lensing signal, we are making sure that their noise contribution is also small.

2.4 Biases to CMB lensing & foreground hardening

In this section we quantify the biases for three different bias-hardening schemes, and compare these results to the standard QE and Shear estimators. We consider two source components: point-sources $f^{\mathrm{PS}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = 1$ and sources with tSZ-like profiles $f^{\mathrm{tSZ}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}} / u_{\boldsymbol{L}}$. The profile $u_{\boldsymbol{\ell}}$ is calculated by taking the square root of the tSZ power spectrum, which is measured from the simulations. Our three different bias-hardening schemes harden against one or

both of these source components. The Point Source Hardened (PSH) estimator hardenens against point sources; the Profile Hardened (PH) estimator hardens against a tSZ-like cluster profile, as described in App. A.4; and the Point source and Profile Hardened (PPH) estimator simultaneously hardens against both. We show the noise power spectra for these various estimators in Fig. 2.2, and performed various checks of our code in App. A.6.

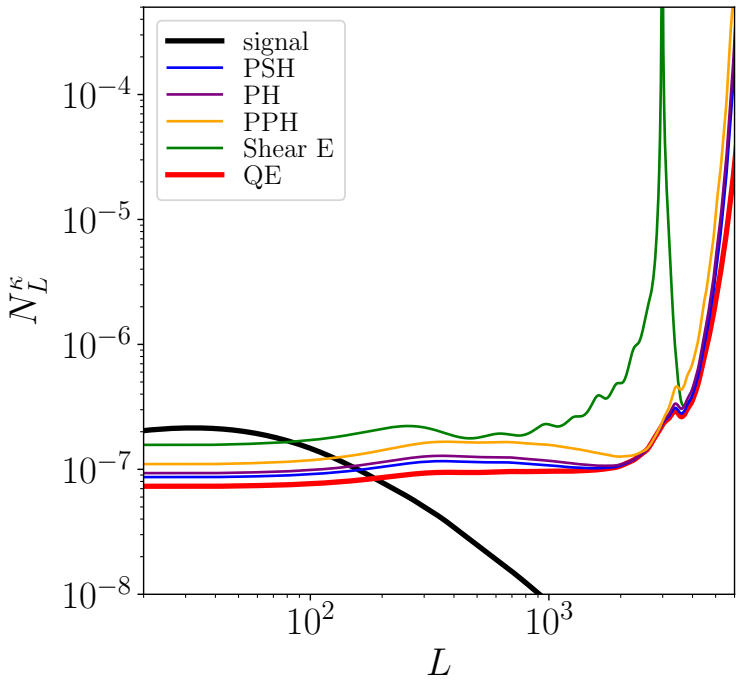


Figure 2.2: Noises of the Point Source Hardened (PSH), Profile Hardened (PH), Point source and Profile Hardened (PPH), Shear, and standard quadratic (QE) estimators for $\ell_{\rm max,T} = 3500$. In black is the lensing signal C_L^κ. Hardening against a single source (PSH,PH) marginally increases the noise over QE, whereas hardening against multiple sources (PPH) results in a larger noise penalty. We also note that the noise cost of bias-hardening increases with the size of the source profile $u_{\boldsymbol{\ell}}$, which is discussed further in Appendix A.5.

Bias to CMB lensing from non-Gaussian foregrounds

The presence of a non-Gaussian foreground component s in the temperature map produces three bias terms to CMB lensing, referred to as primary, secondary and trispectrum terms [524, 446, 227, 597]. Indeed, if $T = T^{\rm CMB} + s$, then the quadratic estimator $\mathcal{Q}$ applied to T can be expanded bilinearly as:

$$\mathcal{Q}[T,T] = \mathcal{Q}[T^{\rm CMB},T^{\rm CMB}] + 2\mathcal{Q}[T^{\rm CMB},s] + \mathcal{Q}[s,s], \tag{2.27}$$

assuming the quadratic estimator is symmetric (or has been symmetrized in its arguments).

In cross-correlation with a mass tracer g, this simply leads to the bispectrum bias:

$$\langle g\mathcal{Q}\rangle = \underbrace{\langle g\kappa_{\rm CMB}\rangle}_{\text{lensing signal}} + \underbrace{\langle g\mathcal{Q}[s,s]\rangle}_{\text{bispectrum bias}}. \tag{2.28}$$

For the auto-correlation, more terms arise:

$$\begin{aligned}\langle \mathcal{Q}\mathcal{Q}\rangle = &\underbrace{\langle \kappa_{\rm CMB}\kappa_{\rm CMB}\rangle}_{\text{lensing signal}} + \underbrace{2\langle \mathcal{Q}[T^{\rm CMB},T^{\rm CMB}]\mathcal{Q}[s,s]\rangle}_{\text{Primary bispectrum bias}} \\ &+ \underbrace{4\langle \mathcal{Q}[T^{\rm CMB},s]\mathcal{Q}[T^{\rm CMB},s]\rangle}_{\text{Secondary bispectrum bias}} + \underbrace{\langle \mathcal{Q}[s,s]\mathcal{Q}[s,s]\rangle}_{\text{Trispectrum bias}}.\end{aligned} \tag{2.29}$$

The primary and secondary bias terms are due to the non-Gaussianity of the foreground and its correlation with the true CMB lensing signal. They are integrals of the $\kappa_{\rm CMB}ss$ bispectrum. In particular, predicting them requires knowing the redshift distribution of the foreground. On the other hand, the trispectrum bias is present whether or not the sources are correlated with CMB lensing, and only depend on the statistics (the trispectrum) of the projected 2d foreground map.

In what follows, we compute these biases for the various quadratic lensing estimators of interest. We follow the approach of [524], using the simulations from [541].

Effectiveness of foreground hardening: expectations

Here we explain which of these bias terms are nulled with bias-hardening. The only knowledge of foregrounds used in the bias hardening process is the shape of the foreground bispectrum and power spectrum, through f^s. In particular, the procedure does not assume anything about the $ss\kappa_{\rm CMB}$ bispectrum, which causes the primary and secondary biases, nor about the $ssss$ trispectrum, which causes the trispectrum bias. Why would bias hardening help reducing foreground biases at all then? The answer is somewhat subtle.

The bias hardening procedure nulls the linear response of the lensing estimator to the foreground s, i.e. the bispectrum $\langle \mathcal{Q}^{\rm BH}[s,s]\ s\rangle$. This is very similar but not identical to the primary bispectrum bias $\langle \mathcal{Q}^{\rm BH}[s,s]\ \kappa_{\rm CMB}\rangle$. The two bispectra are all the more similar as the foreground map and the true lensing map are highly correlated. Thus in principle, the primary bispectrum bias is only partially reduced by the bias-hardening. In practice though, as we show below with simulated foregrounds, this reduction is very large.

Although the secondary contraction is also a bispectrum of the form $ss\kappa_{\rm CMB}$, it comes from a different contraction, $\langle \mathcal{Q}[T,s]\mathcal{Q}[T,s]\rangle$, which the bias hardened estimator was not designed to null. For this reason, we do not expect bias hardening to reduce the secondary contraction. Indeed, we show below that the various lensing estimators have a similar level of secondary bias.

The case of the trispectrum bias is interesting. In general, if $\mathcal{T}_{\boldsymbol{\ell}_1,\boldsymbol{\ell}_2,\boldsymbol{\ell}_3,\boldsymbol{\ell}_4}$ is the source trispectrum, the trispectrum bias to lensing is given by

$$\langle \mathcal{Q}[s,s]\mathcal{Q}[s,s]\rangle_c = N_{\boldsymbol{L}}^2 \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \int \frac{d^2\boldsymbol{\ell}'}{(2\pi)^2} F_{\boldsymbol{\ell}',-\boldsymbol{L}-\boldsymbol{\ell}'} \mathcal{T}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}',-\boldsymbol{L}-\boldsymbol{\ell}'}. \tag{2.30}$$

If the lensing weights of the bias-hardened estimator are denoted as $F^{\rm BH}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$, then the zero linear response to the foreground can be translated as:

$$\int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^{\kappa^{\rm BH}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{s}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = 0. \tag{2.31}$$

This does not null the trispectrum bias in general. However, it does in the case of Poisson sources with identical profiles. Indeed, in this case, the trispectrum

$$\mathcal{T}_{\boldsymbol{\ell}_1,\boldsymbol{\ell}_2,\boldsymbol{\ell}_3,\boldsymbol{\ell}_4} = \bar{n}\langle s^4\rangle u_{\boldsymbol{\ell}_1} u_{\boldsymbol{\ell}_2} u_{\boldsymbol{\ell}_3} u_{\boldsymbol{\ell}_4}$$

is separable, such that the trispectrum bias becomes:

$$\bar{n}\langle s^4\rangle N_{\boldsymbol{L}}^2 \left[\int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^{\rm BH}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}\right]^2. \tag{2.32}$$

Since the response $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}/u_{\boldsymbol{L}} \propto u_{\boldsymbol{\ell}} u_{\boldsymbol{L}-\boldsymbol{\ell}}$, we see that Eq. (2.31) implies that the trispectrum bias is exactly nulled. However, it is no longer generally true if the sources are clustered, or if there is a distribution of profile sizes or shapes. Because our profile hardening assumes a single fiducial source profile, the trispectrum bias is also not exactly nulled when the foreground sources have a range of profiles, as is the case for the tSZ.

In summary, the PH estimator should significantly reduce the tSZ trispectrum bias, while the PSH estimator should exactly null the point source trispectrum. Both should largely reduce the primary bias, with no *a priori* expectation for the secondary bias.

For bias hardening against more complex foregrounds, such as clustered sources, the effectiveness of this simple bias hardening is not guaranteed. However, if the statistical properties of the source distribution is known, it is nonetheless possible to extend our formalism to harden against non-linear clustering. This has been explored in the context of line intensity mapping in [237], and we believe that a similar treatment could be helpful for mitigating biases to CMB lensing. In what follows, we quantify this effect precisely, by running the hardened estimators on foreground simulations with a realistic clustering.

Effectiveness of foreground hardening: simulations

To quantity the biases of each estimator for each of the extragalactic foregrounds, we use simulated maps [541] of the CIB, tSZ, kSZ, radio point sources, and $\kappa_{\rm CMB}$ at 148 GHz. We rescale each of the foregrounds (0.38 for CIB, 0.7 for tSZ, 0.82 for kSZ, 1.1 for radio PS) to match [200], then subtract the mean from each foreground, and finally mask each foreground

using a matched filter for point sources with a flux cut of 5 mJy and mask patch radius 3'. The resulting power spectra are plotted in Appendix A.1.

For CMB measurements, we consider an experiment similar to Simons Observatory [22], with a white noise level of 6 μK-arcmin and a Gaussian beam with full-width at half maximum of 1.4′.

As in [524], we re-weight the halos from [541] to match the redshift distribution of the LSST gold sample ($dn/dz \propto (z/0.24)^2 e^{-z/0.24}/0.48$) [368] and obtain a projected galaxy number density count δ_g. Altogether, a set of these maps (foregrounds, $\kappa_{\rm CMB}$, and δ_g) allow us to calculate the biases to the cross-correlation of CMB lensing with galaxy number density counts $C_L^{\kappa\delta_g}$, as well as the primary, secondary, and trispectrum biases to the lensing auto-correlation $C_L^{\kappa\kappa}$. We run the estimators on 81 flat square cutouts from these maps, obtaining the biases plotted in Fig. 2.3 and 2.4.

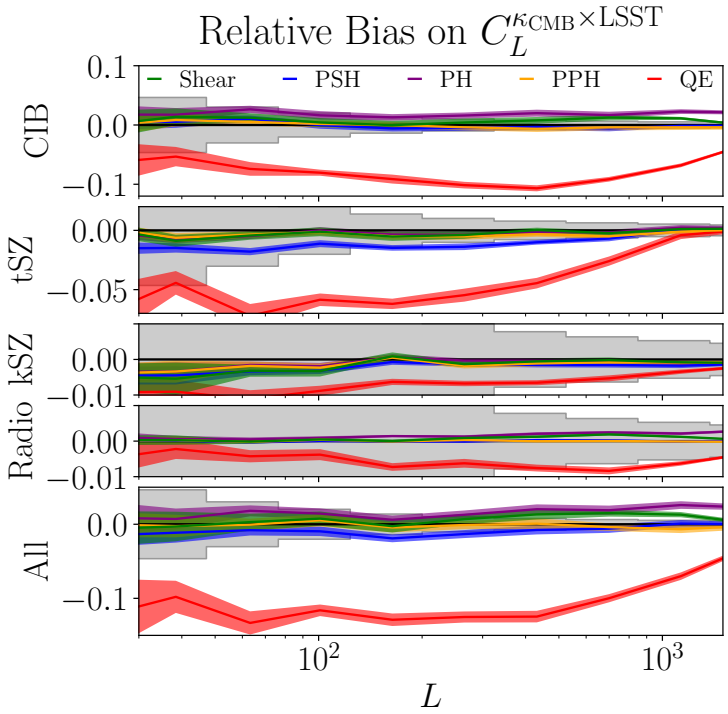


Figure 2.3: Relative bias on the cross with LSST at $\ell_{\rm max,T} = 3500$. The gray boxes are the statistical uncertainties when reconstructing $C_L^{\kappa\delta_g}$ with the standard QE. We find that hardening against point sources (PSH) nulls or dramatically reduces the bias from radio point sources and CIB, whereas hardening against a tSZ-like profile (PH) nulls or significantly reduces the bias from tSZ and kSZ. Hardening against both (PPH) effectively nulls the bias from all foregrounds, resulting in a bias that is roughly an order of magnitude smaller than the bias to the standard QE.

For the cross-correlation of CMB lensing with galaxy counts (Fig. 2.3), we find that the standard minimum variance quadratic estimator acquires a $\sim 10\%$ bias from the CIB and tSZ for reasonable values of $\ell_{\mathrm{max,T}}$. Hardening against a single source component with a tSZ-like profile (PH) effectively nulls the bias from both tSZ and kSZ. Hardening against just point sources (PSH) is complementary, effectively nulling the radio point source and CIB biases. When simultaneously hardening against both tSZ-like profiles and point sources (PPH), we're able to null the biases from all extragalactic foregrounds, creating a highly robust estimator of the cross-correlation.

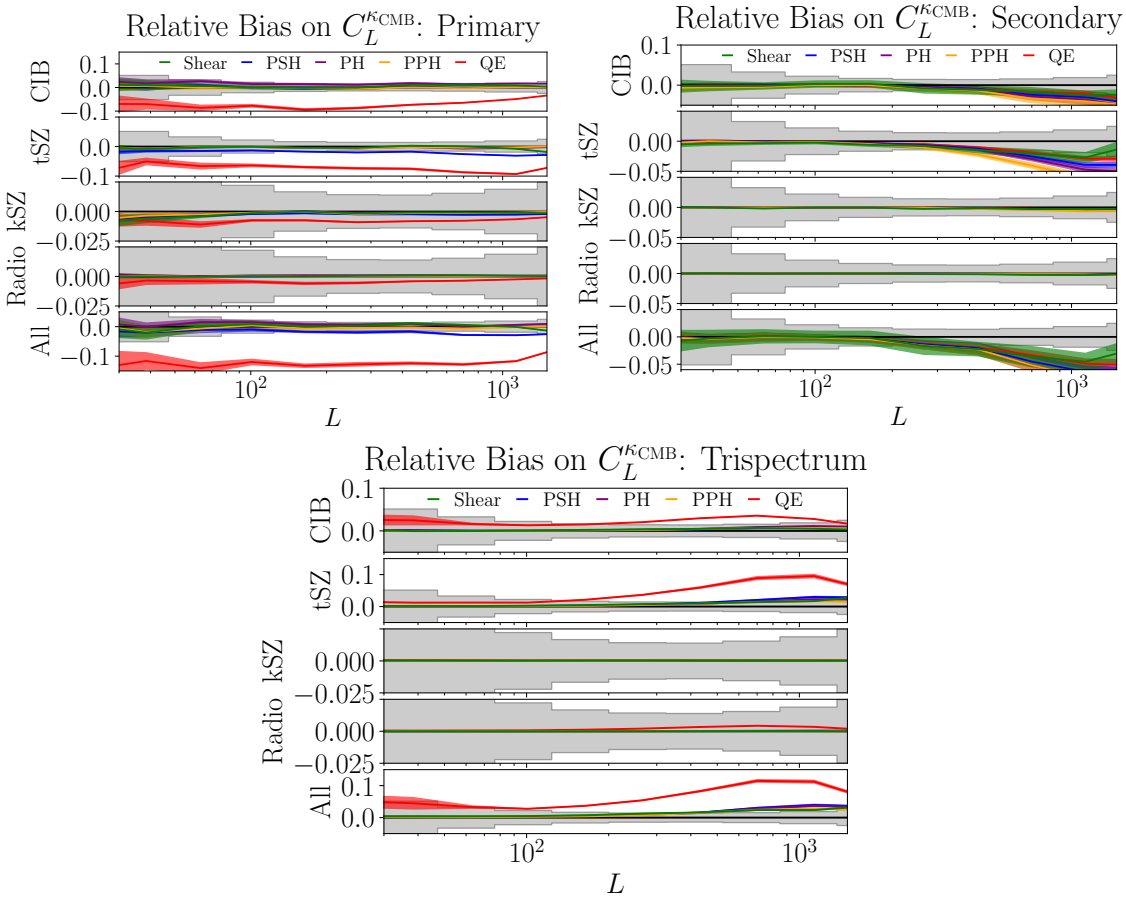


Figure 2.4: Bias to the lensing auto-correlation at $\ell_{\mathrm{max,T}} = 3500$. The gray boxes are the statistical uncertainties when reconstructing C_L^{κ} with the standard QE. Hardening against point sources (PSH) significantly reduces the primary and trispectrum bias from radio point sources and CIB, whereas hardening against a tSZ-like profile significantly reduces these biases from tSZ. Simultaneously hardening against both (PPH) effectively nulls the primary bias. Bias-hardening has no significant effect on the secondary bias.

For the CMB lensing auto-correlation we again find that the standard QE acquires a $\sim$

10% bias, primarily due to the CIB and tSZ. This is consistent with the results found in [446]. In Fig. 2.4 we see that the PH estimator nulls the primary bias from tSZ, and significantly reduces the remaining primary and trispectrum biases. Likewise the PSH estimator nulls the primary bias from point sources, while also significantly reducing remaining primary and trispectrum biases. The PPH estimator essentially nulls all primary biases, and reduces the trispectrum bias to a level comparable to Shear. As discussed in the previous section, there is no reason to expect a significant reduction to the secondary biases from bias-hardening. In practice we find that the secondary bias to bias-hardened estimators is slightly worse than (but comparable to) the bias to the standard QE. We note that the secondary bias is negative, whereas the bias from the trispectrum is positive. This results in a cancellation in the overall bias to the lensing power, allowing higher values of $\ell_{\mathrm{max,T}}$ than one might expect from looking at the secondary and trispectrum biases alone.

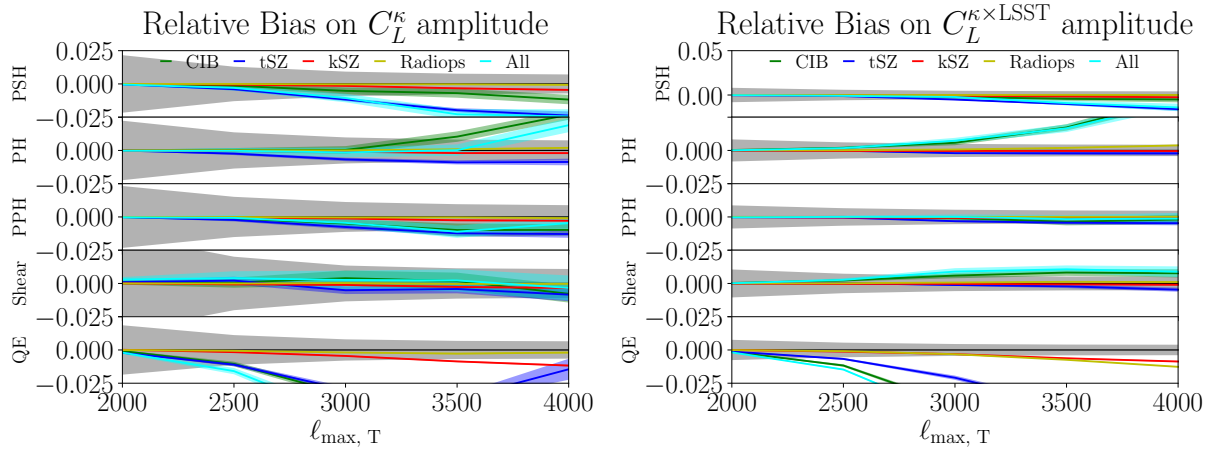


Figure 2.5: Relative bias on the lensing amplitude and the cross-correlation with LSST as a function of $\ell_{\mathrm{max,\,T}}$. In gray is the 1σ error for each estimator $1/\mathrm{SNR}(A_{\mathrm{lens}})$. CIB and tSZ are the main sources of bias to the standard QE. The Point Source Hardened (PSH) estimator drastically reduces the bias from CIB, whereas the Profile Hardened (PH) estimator reduces the bias from tSZ. The Point source and Profile Hardened (PPH) estimator reduces the bias from both tSZ and CIB to a level comparable to Shear while at a lower noise cost.

As a figure of merit for the performance of each estimator, we use the signal-to-noise ratio (SNR) of the lensing amplitude A_{lens}, defined for each L as $C_L^{\kappa,m}/C_L^\kappa$, where $C_L^{\kappa,m}$ and C_L^κ are the measured and true lensing power spectra respectively. By inverse variance weighting the measurement of A_{lens} for each L, we obtain an estimator for the lensing amplitude:

$$\hat{A}_{\mathrm{lens}} = \int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{C_L^{\kappa,m}}{C_L^\kappa} \frac{(C_L^\kappa)^2}{\sigma_L^2} \Bigg/ \int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2} \qquad (2.33)$$
$$\text{with } \sigma_L^2 = 2(C_L^\kappa + N_L)^2.$$

Since $A_{\rm lens}$ has a fiducial value of 1, the SNR and bias of $A_{\rm lens}$ are (assuming $f_{\rm sky} = 1$)

$$\begin{aligned} {\rm SNR}^2(A_{\rm lens}) &= 4\pi \int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2} \\ {\rm bias}(A_{\rm lens}) &= \int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{{\rm bias}(C_L^\kappa)}{C_L^\kappa} \frac{(C_L^\kappa)^2}{\sigma_L^2} \Big/ \int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2}. \end{aligned} \tag{2.34}$$

Similar definitions for the amplitude of the CMB-LSST cross-correlation can be made by replacing the lensing power C_L^κ with $C_L^{\kappa\delta_g}$ and by setting $\sigma_L^2 = (C_L^\kappa + N_L^\kappa)(C_L^{\delta_g} + N_L^{\delta_g}) + (C_L^{\kappa\delta_g})^2$ in the equations above. We show the best SNR achievable in auto and cross-correlation for the various estimators, while maintaining the foreground bias below 1σ, in Tab. 2.1.

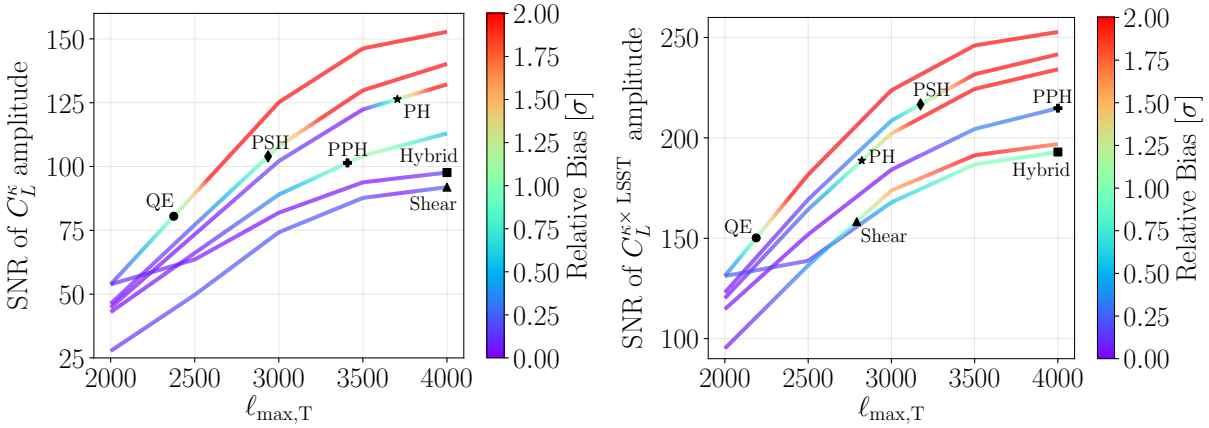


Figure 2.6: Total signal-to-noise of the lensing amplitude (top) and the cross correlation with LSST (bottom). The black markers denote the highest $\ell_{\rm max,\,T}$ where the bias is less than 1σ. We find that the Profile Hardened (PH) estimator reconstructs the lensing amplitude with the highest signal to noise: a $\sim 50\%$ improvement over the standard QE. We find that the Point Source Hardened (PSH) and Point source and Profile Hardened (PPH) reconstruct the cross with similar SNRs: a $\sim 50\%$ improvement over the standard QE. The Hybrid estimator is a combination of Shear and QE, and is described in greater detail in Appendix A.2. We note that the color-scale saturates at 2σ, and that the bias to the standard QE can be $> 10\sigma$ at $\ell_{\rm max,T} \gtrsim 3000$.

For the auto-correlation, we find that the Profile Hardened (PH) estimator achieves the highest SNR before bias($A_{\rm lens}$) becomes larger than the noise $\sigma = 1/{\rm SNR}(A_{\rm lens})$, as shown in Fig. 2.6. This is a $\sim 50\%$ improvement over the standard QE. The PH estimator's 1σ bias threshold occurs at $\ell_{\rm max,T} \sim 3700$. From Fig. 2.5, we see that this high value of $\ell_{\rm max,T}$ is due to a cancellation in the tSZ and CIB biases. A more conservative stopping point would be when any one of the foregrounds crosses the 1σ threshold, which would be around $\ell_{\rm max,T} \sim 3200$ for the PH estimator. We note that even at this lower $\ell_{\rm max,T}$, the PH estimator

still has the highest SNR and is still a $\sim 30\%$ improvement over QE, which crosses the 1σ bias threshold at $\ell_{\rm max,T} \sim 2400$.

For the cross-correlation, the PSH and PPH achieve similar SNRs at the 1σ bias threshold, improving upon QE by $\sim 40\%$. We note that the PPH has a stable bias across the entire range of realistic $\ell_{\rm max,T}$'s, making it a conservative yet competitive option for reconstructing the cross-correlation.

We conclude by noting that for an experiment with $\ell_{\rm max,T} \sim 3000$, the PH and PPH estimators reconstruct the auto- and cross-correlations respectively with the highest SNR while achieving a sub-percent level bias, as shown in Tab. 2.2.

	SNR in auto ($\ell_{\rm max,T}$)	SNR in cross ($\ell_{\rm max,T}$)
PSH	104 (2936)	217 (3175)
PH	126 (3705)	189 (2822)
PPH	101 (3408)	215 (4000)
Shear	92 (4000)	158 (2792)
QE	81 (2375)	150 (2188)

Table 2.1: Comparison of the total signal-to-noise ratio on the amplitude of the CMB lensing auto-spectrum (first column) and cross-spectrum with LSST-like galaxies (second column). In all cases, the maximum temperature multipole $\ell_{\rm max,\ T}$ is selected to guarantee a foreground bias smaller than 1σ.

	Rel. bias to auto (cross) [%]	Rel. bias to auto (cross) [σ]	SNR on C_L^κ ($C_L^{\kappa\times\rm LSST}$)
PSH	-1.1 (-0.2)	-1.14 (-0.46)	108 (209)
PH	-0.1 (0.7)	-0.15 (1.47)	102 (202)
PPH	-0.4 (0.1)	-0.39 (0.01)	89 (184)
Shear	0.3 (0.9)	0.19 (1.55)	74 (173)
QE	-4.9 (-5.0)	-6.15 (-11.16)	125 (223)

Table 2.2: Relative bias (in both % and σ units) and SNR for each estimator at $\ell_{\rm max,T} = 3000$. The PH and PPH estimators reconstruct the auto- and cross-correlations respectively with the highest SNR while achieving a sub-percent level bias.

2.5 Discussion and conclusions

In this study we revisited foreground-hardened lensing quadratic estimators, described how to harden against arbitrary profiles, and extended the formalism to simultaneously harden against an arbitrary number of extragalactic foregrounds. We showed that the point source

hardened estimator (PSH) reduces not only the bias from radio point sources, but also from the cosmic infrared background, the thermal and the kinematic Sunyaev-Zel'dovich effects.

Motivated by the extended nature of tSZ clusters, we modify the point source hardening technique to deproject the lensing bias from uncorrelated clusters (PH estimator), or simultaneously point sources and extended clusters (PPH estimator). We propose a simple and sufficient method to approximate the input cluster profile from the data itself, by considering the observed tSZ power spectrum from the same data.

The signal-to-noise cost from deprojecting the point source bias is small ($\sim 15\%$ for Point Source Hardening), and is more than compensated by enabling the use of higher multipoles: overall, the SNR on the lensing power spectrum is increased by $\sim 56\%$ when using Profile Hardening (PH), while for cross-correlations, both PSH and PPH provide a $\sim 45\%$ increase in SNR compared to QE. Part of the improvement of PH and PPH estimators comes from the increase in the maximum multipole that can be included in the lensing reconstruction.

The PSH, PH and PPH estimators improve the lensing SNR over the shear and hybrid shear estimators of [524], although at the cost of a slightly larger bias in cross-correlation. They also provide consistency checks for the standard quadratic estimator, having a smaller foreground bias for a given $\ell_{\rm max}$ in temperature. Overall, these estimators outperform the standard lensing quadratic estimator both in auto and cross-correlation, both in terms of the lensing SNR and in terms of the foreground bias. We recommend their use in the CMB lensing analyses of upcoming temperature-dominated CMB data such as the Simons Observatory. Future surveys that are dominated by polarization reconstruction such as CMB-S4 [3], will still benefit from improved robustness, since temperature reconstruction will contribute a non-trivial fraction to the SNR. We note that a realistic analysis will require modeling of the higher order biases $N^{(i)}$ [221] that appear in lensing reconstruction and will likely be non-negligible for the next generation of surveys.

Finally, foreground hardening can and should be implemented in conjunction with other methods of foreground mitigation, such as masking [597] and multi-frequency cleaning in one or both legs of the estimator. We shall report on the optimal combination of these methods in chapter 4.

Acknowledgments

We thank Anthony Challinor, Colin Hill, Antony Lewis, Mathew Madhavacheril, Blake Sherwin, Omar Darwish, Alex van Engelen, Neelima Sehgal, David Spergel, Martin White and the participants of the Sussex CMB lensing workshop for all their feedback. We give a special thanks to Zoom for facilitating this collaboration throughout the COVID-19 pandemic. E.S. is supported by the Chamberlain fellowship at Lawrence Berkeley National Laboratory. S.F. is supported by the Physics Division of Lawrence Berkeley National Laboratory. This chapter used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Chapter 3

Geometric foreground mitigation II: bias-hardened and shear-only estimators in polarization

This chapter was originally published as

> Noah Sailer, Simone Ferraro, and Emmanuel Schaan. "Foreground-immune CMB lensing reconstruction with polarization". In: *Phys. Rev. D* 107.2 (2023), p. 023504. DOI: 10.1103/PhysRevD.107.023504. arXiv: 2211.03786 [astro-ph.CO]

Extragalactic foregrounds are known to generate significant biases in temperature-based CMB lensing reconstruction. Several techniques, which include "source hardening" and "shear-only estimators" have been proposed to mitigate contamination and have been shown to be very effective at reducing foreground-induced biases. Here we extend both techniques to polarization, which will be an essential component of CMB lensing reconstruction for future experiments, and investigate the "large-lens" limit analytically to gain insight on the origin and scaling of foreground biases, as well as the sensitivity to their profiles. Using simulations of polarized point sources, we estimate the expected bias to both Simons Observatory and CMB-S4 like (polarization-based) lensing reconstruction, finding that biases to the former are minuscule while those to the latter are potentially non-negligible at small scales ($L \sim 1000 - 2000$). In particular, we show that for a CMB-S4 like experiment, an optimal linear combination of point-source hardened estimators can reduce the (point-source induced) bias to the CMB lensing power spectrum by up to two orders of magnitude, at a $\sim 4\%$ noise cost relative to the global minimum variance estimator.

3.1 Introduction

Point sources and other foregrounds induce biases [597, 227] to standard quadratic estimators of the weak lensing potential due to the fact that foregrounds are non-Gaussian and correlated with the lensing potential. The impact of foregrounds on the CMB temperature lensing reconstruction has been explored extensively in the literature and a number of mitigation techniques have been proposed. These range from decreasing the multipole range used in the reconstruction to more sophisticated "geometric methods" that employ the different symmetries of foregrounds and lensing.

Hardening against point sources in temperature-based reconstruction has been proposed in [434, 446, 436, 519] and has been shown to reduce the foreground biases by an order of magnitude with very modest noise cost of order 10%. A generalization to hardening against a fixed profile (profile hardening) was proposed in [519] and has been shown to be even more effective when the mean profile of tSZ halos is hardened against. An even more general "foreground-immune" estimator is the shear-only estimator [524], which suppresses the response to arbitrary azimuthally symmetric profiles. Further generalizations to "multipole" estimators have also been proposed in [524].

Moreover, frequency-based methods to reduce biases have been proposed in [374, 520], and the optimal combination of multi-frequency and geometric methods have been explored in [166]. Extragalactic foregrounds can also affect delensing of the $B-$modes produced by primordial gravitational waves, and the techniques mentioned above appear to be very effective at reducing these biases [52].

In this chapter we extend two of the most promising geometric techniques used in temperature reconstruction (i.e. bias-hardening and shear-only reconstruction) to polarization. We note that these can be used in addition to multi-frequency foreground reduction to further mitigate foreground biases.

CMB polarization dominates the statistical power of lensing reconstruction for map noise $\Delta_T \lesssim 5\,\mu$K-arcmin, and therefore will be the principal channel in the future. For current experiments, polarization already dominates reconstruction for SPT-3G [80] and will play an important role in Simons Observatory [22] and CMB-S4 [2]. It is therefore important to characterize and mitigate possible biases induced by foregrounds in polarization.

First of all, we note that extragalactic foregrounds in polarization are thought to be simpler than in temperature. The largest sources of extragalactic foregrounds, such as emission from the Cosmic Infrared Background (CIB) or the thermal Sunyaev-Zel'dovich effect (tSZ), are mostly unpolarized (although a small polarized signal is present at higher order [224, 190]), and therefore their effect on lensing reconstruction from polarization is likely negligible. However, bright infrared and radio point sources are known to be polarized at the $\sim$ few % level [351, 168, 263], inducing potential biases. Therefore we'll concentrate on mitigating the effect of polarized point sources, while noting that the techniques we present here have also been shown to be very effective against CIB and tSZ in temperature, and therefore are likely to also work more generally in polarization.

In this chapter we focus on extragalactic foregrounds. The methods developed to mitigate

these biases may also prove useful for galactic foregrounds, whose biases are potentially important and their study will be the subject of future work.

The remainder of the chapter is organized as follows: in section 3.2, we outline the theory of quadratic estimators, including multi-frequency and minimum-variance combinations. In section 3.3, we extend the bias-hardening techniques to polarization, while in section 3.4 we explore generalizations to the "shear-only" estimators, highlighting important differences in polarization. In section 3.5 we study the noise properties of these new estimators and in section 3.6 we investigate the bias reduction obtained. We conclude in section 3.7.

3.2 Linear responses and quadratic estimators

Here we review the basic theory of quadratic estimators in temperature and polarization. We separately consider the cases of reconstructing the convergence from a single pair of maps and the case of summing over all pairs, and derive the minimum variance weights for each case. For simplicity we work under the flat sky approximation throughout.

Lensed fields and linear responses

The lensed CMB temperature $\tilde{T}_{\boldsymbol{\ell}}$ is related to the unlensed field $T_{\boldsymbol{\ell}}$ *to linear order* in the lensing convergence κ as follows [342]:

$$\tilde{T}_{\boldsymbol{\ell}} = T_{\boldsymbol{\ell}} - 2\int_{\boldsymbol{\ell}'} \frac{\boldsymbol{\ell}'\cdot(\boldsymbol{\ell}-\boldsymbol{\ell}')}{(\boldsymbol{\ell}-\boldsymbol{\ell}')^2}\kappa_{\boldsymbol{\ell}-\boldsymbol{\ell}'}T_{\boldsymbol{\ell}'}, \tag{3.1}$$

where $\int_{\boldsymbol{\ell}_1\cdots\boldsymbol{\ell}_n} \equiv \int d^2\boldsymbol{\ell}_1\cdots d^2\boldsymbol{\ell}_n/(2\pi)^{2n}$. Throughout we use tildes to distinguish lensed fields from their unlensed counterparts. The analogous expression for the lensed polarization fields is:

$$\tilde{E}_{\boldsymbol{\ell}} \pm i\tilde{B}_{\boldsymbol{\ell}} = E_{\boldsymbol{\ell}} - 2\int_{\boldsymbol{\ell}'} \frac{\boldsymbol{\ell}'\cdot(\boldsymbol{\ell}-\boldsymbol{\ell}')}{(\boldsymbol{\ell}-\boldsymbol{\ell}')^2}e^{\pm 2i(\phi_{\boldsymbol{\ell}'}-\phi_{\boldsymbol{\ell}})}\kappa_{\boldsymbol{\ell}-\boldsymbol{\ell}'}E_{\boldsymbol{\ell}'}, \tag{3.2}$$

where the angle $\phi_{\boldsymbol{\ell}}$ is defined through $\boldsymbol{\ell} = \ell\cos(\phi_{\boldsymbol{\ell}})\hat{\boldsymbol{x}}_1 + \ell\sin(\phi_{\boldsymbol{\ell}})\hat{\boldsymbol{x}}_2$, $\{\hat{\boldsymbol{x}}_1, \hat{\boldsymbol{x}}_2\}$ are orthonormal basis vectors, and we've assumed that there are no primordial $B-$modes.

The lensed CMB temperature and polarization fields are statistically isotropic, and as a result $\langle \tilde{X}_{\boldsymbol{\ell}}\tilde{Y}_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle = 0$ for $L > 0$, where $\langle\cdots\rangle$ denotes an average over realizations of the primary CMB *and* the lensing field, and $X, Y \in \{T, E, B\}$. However, fixing the Fourier mode $\kappa_{\boldsymbol{L}}$ breaks statistical isotropy, resulting in off diagonal covariances of the form:

$$\langle \tilde{X}_{\boldsymbol{\ell}}\tilde{Y}_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle' = f^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\,\kappa_{\boldsymbol{L}} + \mathcal{O}(\kappa^2) \quad \text{for} \quad L > 0, \tag{3.3}$$

where $f^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ is the linear response of lensed $\tilde{X}, \tilde{Y}$ fields to the lensing convergence, and $\langle\cdots\rangle'$ denotes an average over realizations of the primary CMB and all Fourier modes of the lensing

field *except*[1] for $\boldsymbol{L}$. Note that the responses trivially satisfy $f^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}'} = f^{YX}_{\boldsymbol{\ell}',\boldsymbol{\ell}}$.

From Eqs. (3.1) and (3.2) one obtains[2]:

$$\begin{aligned}
f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{2\boldsymbol{L}}{L^2}\cdot\left[\boldsymbol{\ell}C^{TT}_{\ell} + (\boldsymbol{L}-\boldsymbol{\ell})C^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right] \\
f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{2\boldsymbol{L}}{L^2}\cdot\left[\boldsymbol{\ell}C^{TE}_{\ell}\cos(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}) + (\boldsymbol{L}-\boldsymbol{\ell})C^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right] \\
f^{TB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{2\boldsymbol{L}}{L^2}\cdot\boldsymbol{\ell}C^{TE}_{\ell}\sin(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}) \\
f^{EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{2\boldsymbol{L}}{L^2}\cdot\left[\boldsymbol{\ell}C^{EE}_{\ell} + (\boldsymbol{L}-\boldsymbol{\ell})C^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right]\cos(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}) \\
f^{EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{2\boldsymbol{L}}{L^2}\cdot\boldsymbol{\ell}C^{EE}_{\ell}\sin(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}) \\
f^{BB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= 0,
\end{aligned} \tag{3.4}$$

where $C^{TT}_{\ell}, C^{TE}_{\ell}$, and C^{EE}_{ℓ} are the unlensed power spectra, and $\theta_{\boldsymbol{\ell},\boldsymbol{\ell}'} \equiv \phi_{\boldsymbol{\ell}} - \phi_{\boldsymbol{\ell}'}$ is the angle between $\boldsymbol{\ell}$ and $\boldsymbol{\ell}'$.

Minimum variance estimators

Eq. (3.3) is the starting point for deriving a quadratic estimator (QE) for the lensing convergence. By dividing both sides of Eq. (3.3) by the linear response, one can build a simple unbiased QE $\hat{\kappa}^{XY}_{\boldsymbol{\ell},\boldsymbol{L}}$ for each $\boldsymbol{\ell}$: $\hat{\kappa}^{XY}_{\boldsymbol{\ell},\boldsymbol{L}} = \tilde{X}_{\boldsymbol{\ell}}\tilde{Y}_{\boldsymbol{L}-\boldsymbol{\ell}}/f^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$, such that $\langle\hat{\kappa}^{XY}_{\boldsymbol{\ell},\boldsymbol{L}}\rangle' = \kappa_{\boldsymbol{L}}$ to linear order. The most general QE is obtained by taking an arbitrary weighted average of the $\hat{\kappa}^{XY}_{\boldsymbol{\ell},\boldsymbol{L}}$'s, provided that the weights sum to unity (to keep the estimator unbiased). Below we consider separately the specialized case where only a single pair of maps are used for the lensing reconstruction, and the more general scenario of summing over all available pairs. While the former is slightly suboptimal in terms of noise, it is still practically useful to have a collection of individual QEs for performing cross-checks and diagnosing systematics.

[1] We could have also defined $\langle\cdots\rangle'$ as an average over the primary CMB with the entire lensing field fixed (not just the $\boldsymbol{L}$'th Fourier mode) and Eq. (3.3) would still be valid. In 3.3 our definition of only fixing the $\boldsymbol{L}$'th Fourier mode will be relevant for constructing the foreground QE, and we choose to adopt a common notation throughout for convenience.

[2] Note that $\phi_{-\boldsymbol{\ell}} \cong \phi_{\boldsymbol{\ell}} + n\pi$ for any odd integer n. Since all angles in Eq. (3.4) are multiplied by 2, we are always free to change the sign of $\boldsymbol{\ell}$ in $\phi_{\boldsymbol{\ell}}$, which we have done to make the responses look as symmetric as possible.

Single pair of maps

Given a single pair of maps $\tilde{X}, \tilde{Y}$ one can construct a generic unbiased QE for the lensing convergence:

$$\begin{aligned} \hat{\kappa}^{XY}_{\boldsymbol{L}} &= N^{XY}_{\boldsymbol{L}} \int_{\boldsymbol{\ell}} F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tilde{X}_{\boldsymbol{\ell}} \tilde{Y}_{\boldsymbol{L}-\boldsymbol{\ell}} \\ \text{where} \quad (N^{XY}_{\boldsymbol{L}})^{-1} &= \int_{\boldsymbol{\ell}} F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}, \end{aligned} \tag{3.5}$$

where $F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ are some weights, which are in principle arbitrary. The Gaussian component of the variance of $\hat{\kappa}^{XY}$, defined such that $\langle \hat{\kappa}^{XY}_{\boldsymbol{L}} \hat{\kappa}^{XY}_{\boldsymbol{L}'} \rangle_c = (2\pi)^2 \delta^D_{\boldsymbol{L}+\boldsymbol{L}'} \mathcal{N}^{XY}_{\boldsymbol{L}}$, is given by

$$\mathcal{N}^{XY}_{\boldsymbol{L}} = (N^{XY}_{\boldsymbol{L}})^2 \int_{\boldsymbol{\ell}} F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \left[F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tilde{C}^{XX}_{\ell} \tilde{C}^{YY}_{|\boldsymbol{L}-\boldsymbol{\ell}|} + F^{XY}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}} \tilde{C}^{XY}_{\ell} \tilde{C}^{XY}_{|\boldsymbol{L}-\boldsymbol{\ell}|} \right], \tag{3.6}$$

where $\tilde{C}^{XY}_{\ell}$ is the *total observed* power spectrum, i.e. the power spectra of the *lensed* fields, including instrument noise and foregrounds. In the equation above we have made the simplifying assumption that the weights are real, so that $F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = F^{XY}_{-\boldsymbol{\ell},\boldsymbol{\ell}-\boldsymbol{L}}$. From here on out we will always assume that the weights and linear responses are real.

A common and optimistic choice for the weights is to minimize $\mathcal{N}^{XY}_{\boldsymbol{L}}$. This can be accomplished using e.g. functional derivatives, giving [300]:

$$F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{\tilde{C}^{YY}_{\ell} \tilde{C}^{XX}_{|\boldsymbol{L}-\boldsymbol{\ell}|} f^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{XY}_{\ell} \tilde{C}^{XY}_{|\boldsymbol{L}-\boldsymbol{\ell}|} f^{XY}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}}}{\tilde{C}^{XX}_{\ell} \tilde{C}^{YY}_{|\boldsymbol{L}-\boldsymbol{\ell}|} \tilde{C}^{YY}_{\ell} \tilde{C}^{XX}_{|\boldsymbol{L}-\boldsymbol{\ell}|} - \left(\tilde{C}^{XY}_{\ell} \tilde{C}^{XY}_{|\boldsymbol{L}-\boldsymbol{\ell}|} \right)^2}. \tag{3.7}$$

Note that in the limit where $\tilde{X}$ and $\tilde{Y}$ are uncorrelated (or when $\tilde{X} = \tilde{Y}$) the minimum variance weights reduce to $F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \propto f^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} / \tilde{C}^{XX}_{\ell} \tilde{C}^{YY}_{|\boldsymbol{L}-\boldsymbol{\ell}|}$ up to some generic prefactor (any function of $\boldsymbol{L}$) which the QE (Eq. (3.5)) is insensitive to. Neglecting the cross-correlation between e.g. T and E is often a very good approximation, and doing so makes Eq. (3.5) amenable to FFTs. Of course, this approximation to the weights still leaves the estimator unbiased, although slightly suboptimal. In what follows we always make the approximation of neglecting the cross-correlation in the minimum variance weights for a single pair of maps.

Summing over all pairs

Given a collection of lensed maps $\{\tilde{M}^i\}$ one can build the general QE:

$$\begin{aligned} \hat{\kappa}_{\boldsymbol{L}} &= N^{\kappa}_{\boldsymbol{L}} \sum_{ij} \int_{\boldsymbol{\ell}} F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tilde{M}^i_{\boldsymbol{\ell}} \tilde{M}^j_{\boldsymbol{L}-\boldsymbol{\ell}} \\ \text{where} \quad (N^{\kappa}_{\boldsymbol{L}})^{-1} &= \sum_{ij} \int_{\boldsymbol{\ell}} F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}, \end{aligned} \tag{3.8}$$

where the linear response $f^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ to the lensing convergence is defined in an analogous way as Eq. (3.3). Note that in Eq. (3.8) we are summing over all pairs (e.g. both ij and ji) and that the weights satisfy $F^{ij}_{\boldsymbol{\ell},\boldsymbol{\ell}'} = F^{ji}_{\boldsymbol{\ell}',\boldsymbol{\ell}}$. The index i is quite general. For example, i could simply index T, E and B; or $\tilde{M}^i$ could be a temperature map with i indexing the frequency; or even more generally i could index a pair of indices (X, ν) where $X \in \{T, E, B\}$ and ν is a frequency.

The variance of Eq. (3.8) is:

$$\mathcal{N}^\kappa_{\boldsymbol{L}} = 2(N^\kappa_{\boldsymbol{L}})^2 \sum_{ijmn} \int_{\boldsymbol{\ell}} F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} F^{mn}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tilde{C}^{im}_{\ell} \tilde{C}^{jn}_{|\boldsymbol{L}-\boldsymbol{\ell}|}, \tag{3.9}$$

where the covariance matrix $\tilde{C}^{ij}_\ell$ is defined as $\langle \tilde{M}^i_{\boldsymbol{\ell}} \tilde{M}^j_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle = (2\pi)^2 \delta^D_{\boldsymbol{L}} \tilde{C}^{ij}_\ell$.

The weights which minimize $\mathcal{N}^\kappa_{\boldsymbol{L}}$ are given by:

$$F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{1}{2} \sum_{mn} \left(\tilde{\boldsymbol{C}}^{-1}_\ell \right)^{im} \left(\tilde{\boldsymbol{C}}^{-1}_{|\boldsymbol{L}-\boldsymbol{\ell}|} \right)^{jn} f^{mn}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}, \tag{3.10}$$

where we have chosen the normalization of the weights so that $\mathcal{N}^\kappa_{\boldsymbol{L}} = N^\kappa_{\boldsymbol{L}}$. Note that unlike for the single pair case, here the true minimum variance QE is always FFT-able. We provide explicit expressions for these weights in Appendix B.2.

Note that in the case where $\tilde{M}^i$ is a temperature map with i indexing a frequency, Eq. (3.10) reduces to the familiar form: $w^i_\ell w^j_{|\boldsymbol{L}-\boldsymbol{\ell}|} f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} / 2\tilde{C}^{\rm ILC}_\ell \tilde{C}^{\rm ILC}_{|\boldsymbol{L}-\boldsymbol{\ell}|}$, where $\tilde{C}^{\rm ILC}_\ell$ is the power spectrum of the minimum variance internal linear combination (MVILC) of temperature maps, and w^i_ℓ are the MVILC weights. In the case where i indexes T, E, and B (at a single frequency) Eq. (3.10) recovers the global minimum variance estimator (GMV), originally derived in [289]. As was pointed out in [382], the GMV has a slightly lower noise than the minimum variance linear combination of the MV single-pair estimators $\hat{\kappa}^{XY}_{\boldsymbol{L}}$.

3.3 Bias-hardened estimators

The most general QE takes the form of Eq. (3.8), which is unbiased when the maps $\tilde{M}^i$ are composed of a lensed CMB signal and Gaussian noise, so that $\langle \tilde{M}^i_{\boldsymbol{\ell}} \tilde{M}^j_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle' = f^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{\boldsymbol{L}}$. If instead the maps are contaminated by a foreground s^i, the covariance $\langle \tilde{M}^i_{\boldsymbol{\ell}} \tilde{M}^j_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle'$ picks up additional off-diagonal contributions, which to lowest order in κ and s^i take the form[3] [446, 519]:

$$\langle \tilde{M}^i_{\boldsymbol{\ell}} \tilde{M}^j_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle' = f^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{\boldsymbol{L}} + \sum_k g^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s^k_{\boldsymbol{L}}, \tag{3.11}$$

where as before f^{ij} is the linear response to lensing in the absence of foregrounds, while $g^{ij,k}$ is the linear response to the foreground in the absence of lensing[4]: $\langle s^i_{\boldsymbol{\ell}} s^j_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle' = \sum_k g^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s^k_{\boldsymbol{L}}$.

[3] Note that in the average $\langle \cdots \rangle'$ we are fixing the $\boldsymbol{L}$'th Fourier mode of both the lensing convergence and the foreground.

[4] Note that one could equivalently define $g^{ij,k}$ as the linear response of $\langle \tilde{M}^i \tilde{M}^j \rangle'$ to s^k in the absence of lensing.

The additional term in Eq. (3.11) induces an additive bias to Eq. (3.8) which is proportional to the foreground:

$$\langle \hat{\kappa}_{\boldsymbol{L}} \rangle' = \kappa_{\boldsymbol{L}} + \sum_k \underbrace{\left(N^{\kappa}_{\boldsymbol{L}} \sum_{ij} \int_{\boldsymbol{\ell}} F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} g^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \right)}_{\equiv \mathcal{R}^{\kappa,k}_{\boldsymbol{L}}} s^k_{\boldsymbol{L}}, \tag{3.12}$$

where we have defined the response $\mathcal{R}^{\kappa,i}_{\boldsymbol{L}}$ of the lensing estimator to the foreground in the i'th map.

The intuition behind bias-hardening goes as follows: just as we can construct a convergence QE $\hat{\kappa}$ we can also construct a foreground QE $\hat{s}^i$ which can in turn be used to subtract off the bias to Eq. (3.12). The only ingredients that we need are the foreground linear response and a set of weights $G^{ij,k}$:

$$\begin{aligned} \hat{s}^i_{\boldsymbol{L}} &= N^i_{\boldsymbol{L}} \sum_{jk} \int_{\boldsymbol{\ell}} G^{jk,i}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tilde{M}^j_{\boldsymbol{\ell}} \tilde{M}^k_{\boldsymbol{L}-\boldsymbol{\ell}} \\ \text{where} \quad (N^i_{\boldsymbol{L}})^{-1} &= \sum_{jk} \int_{\boldsymbol{\ell}} G^{jk,i}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} g^{jk,i}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}. \end{aligned} \tag{3.13}$$

When evaluated on the contaminated maps, $\hat{s}^i$ picks up field level biases proportional to the foreground in the other maps ($m \neq i$), and to κ:

$$\begin{aligned} \langle \hat{s}^i_{\boldsymbol{L}} \rangle' &= \sum_m \left(N^i_{\boldsymbol{L}} \sum_{jk} \int_{\boldsymbol{\ell}} G^{jk,i}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} g^{jk,m}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \right) s^m_{\boldsymbol{L}} + \left(N^i_{\boldsymbol{L}} \sum_{jk} \int_{\boldsymbol{\ell}} G^{jk,i}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{jk}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \right) \kappa_{\boldsymbol{L}} \\ &\equiv \sum_m \mathcal{R}^{i,m}_{\boldsymbol{L}} s^m_{\boldsymbol{L}} + \mathcal{R}^{i,\kappa}_{\boldsymbol{L}} \kappa_{\boldsymbol{L}}, \end{aligned} \tag{3.14}$$

where $\mathcal{R}^{i,i}_{\boldsymbol{L}} = 1$. By recasting Eqs. (3.12) and (3.14) in a matrix form one can construct a suitable linear combination of $\hat{\kappa}$ and $\hat{s}^i$ to null the linear order field level biases:

$$\begin{pmatrix} \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \\ \hat{s}^{{\rm BH},1}_{\boldsymbol{L}} \\ \vdots \\ \hat{s}^{{\rm BH},n}_{\boldsymbol{L}} \end{pmatrix} = \begin{pmatrix} 1 & \mathcal{R}^{\kappa,1}_{\boldsymbol{L}} & \cdots & \mathcal{R}^{\kappa,n}_{\boldsymbol{L}} \\ \mathcal{R}^{1,\kappa}_{\boldsymbol{L}} & 1 & \cdots & \mathcal{R}^{1,n}_{\boldsymbol{L}} \\ \vdots & \vdots & \ddots & \vdots \\ \mathcal{R}^{n,\kappa}_{\boldsymbol{L}} & \mathcal{R}^{n,1}_{\boldsymbol{L}} & \cdots & 1 \end{pmatrix}^{-1} \begin{pmatrix} \hat{\kappa}_{\boldsymbol{L}} \\ \hat{s}^1_{\boldsymbol{L}} \\ \vdots \\ \hat{s}^n_{\boldsymbol{L}} \end{pmatrix}. \tag{3.15}$$

By taking an average of Eq. (3.15) and plugging in Eqs. (3.12) and (3.14), we see that the bias-hardened estimators satisfy: $\langle \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \rangle' = \kappa_{\boldsymbol{L}}$ and $\langle \hat{s}^{{\rm BH},i}_{\boldsymbol{L}} \rangle' = s^i_{\boldsymbol{L}}$.

In practice we choose the foreground weights G and lensing weights F to minimize the variance of $\hat{s}^i$ and $\hat{\kappa}$ respectively. In Appendix B.4 we show that this choice of weights is equivalent to the minimum variance lensing estimator constrained to have zero response to foregrounds (i.e. the optimal bias-hardened estimator).

For the case of using a single pair of maps for reconstruction we therefore take $G^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = g^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}/\tilde{C}^{ii}_{\ell}\tilde{C}^{jj}_{|\boldsymbol{L}-\boldsymbol{\ell}|}$ and $F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = f^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}/\tilde{C}^{ii}_{\ell}\tilde{C}^{jj}_{|\boldsymbol{L}-\boldsymbol{\ell}|}$, where we have neglected the cross-correlation $\tilde{C}^{ij}$ to make the estimators FFT-able. As an illustrative example we explicitly bias-harden the TE estimator in Appendix B.6 using the foreground linear responses derived in section 3.3.

Foreground model

The only input required to construct the foreground QE is the linear response. Below we outline our model for the foregrounds. We use this model to compute the linear response in the following subsection.

We assume that the foreground s is a sum of $N_{\rm sources}$ sources with identical profiles, and further assume that polarization fractions and polarization angles do not vary significantly over a single source, so that the intensity and polarization of the foreground take the form:

$$\begin{aligned} s^I(\boldsymbol{x}) &= \sum_i s^I_i u(\boldsymbol{x}-\boldsymbol{x}_i) \\ s^Q(\boldsymbol{x}) \pm i s^U(\boldsymbol{x}) &= \sum_i s^I_i p_i e^{\pm 2i\psi_i} u(\boldsymbol{x}-\boldsymbol{x}_i) \end{aligned} \tag{3.16}$$

where u is the profile, and $s^I_i, p_i, \psi_i, \boldsymbol{x}_i$ is the flux, polarization fraction, polarization angle, and position of the i'th source respectively. For simplicity we will assume that the sources emit no circular polarization[5]. From the definition of the E and B fields [636, 342]

$$E_{\boldsymbol{\ell}} \pm iB_{\boldsymbol{\ell}} = -\big[Q_{\boldsymbol{\ell}} \pm iU_{\boldsymbol{\ell}}\big] e^{\mp 2i\phi_{\boldsymbol{\ell}}} \tag{3.17}$$

the T, E, and B fields for the sources take the form:

$$\begin{aligned} s^T_{\boldsymbol{\ell}} &= \sum_i s_i e^{-i\boldsymbol{\ell}\cdot\boldsymbol{x}_i} u_{\boldsymbol{\ell}} \\ s^E_{\boldsymbol{\ell}} \pm i s^B_{\boldsymbol{\ell}} &= -\sum_i s_i p_i e^{\pm 2i(\psi_i - \phi_{\boldsymbol{\ell}})} e^{-i\boldsymbol{\ell}\cdot\boldsymbol{x}_i} u_{\boldsymbol{\ell}}, \end{aligned} \tag{3.18}$$

where s_i is the flux in the appropriate units.

To compute the responses we must make some assumptions about the joint PDF of the sources $P(\{s_i, p_i, \psi_i, \boldsymbol{x}_i\})$. As we will see in the next subsection, this is easily accomplished with the following simple set of assumptions:

- Each source is statistically independent.

[5] While the primary CMB is expected to have a negligible level of circular polarization in absence of new physics [418], physical processes such as synchrotron emission (which is responsible for radio point sources) can have a circularly-polarized component, which may be potentially detectable by future experiments at low frequencies ($\nu \lesssim 10$ GHz) [321]. It appears that circular polarization should be negligible at CMB frequencies, and it will not be explored further in this chapter.

- Positions and polarization angles are statistically independent, both of each other, and independent of polarization fractions and fluxes.
- Positions and polarization angles are uniformly distributed. This key assumption neglects the clustering of sources, and is appropriate for the polarized radio and infrared sources we consider here (on the small scales relevant for lensing reconstruction).

These assumptions force the joint PDF to factorize as[6]:

$$P\left(\{s_i, p_i, \psi_i, \boldsymbol{x}_i\}\right) = P\left(N_{\rm sources}\right) \times \prod_{i=1}^{N_{\rm sources}} \mathcal{U}\left(\boldsymbol{x}_i\right) \mathcal{U}\left(\psi_i\right) \mathrm{LF}\left(s_i\right) P\left(p_i | s_i\right). \tag{3.19}$$

Here, "LF" is the source luminosity function, normalized to integrate to unity, and $\mathcal{U}$ denotes a uniform distribution (in coordinate or angle).

Note that by forcing each source to be independent and uniformly distributed on the sky, any n-point function of the foreground will only involve one independent sum over sources, since

$$\langle \mathrm{Exp}[-i(\boldsymbol{\ell}_1 \cdot \boldsymbol{x}_{i_1} + \cdots + \boldsymbol{\ell}_n \cdot \boldsymbol{x}_{i_n})] \rangle$$

vanishes unless $i_1 = \cdots = i_n$ and $\sum_i \boldsymbol{\ell}_i = 0$, assuming that no non-trivial subset of the $\boldsymbol{\ell}_i$'s sum to zero, which will always be the case for the responses below. That is, our assumptions create the usual shot noise n-point functions.

Foreground linear response

Our goal is to compute the linear response $g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ of the covariance $\langle s^X_{\boldsymbol{\ell}} s^Y_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle'$ to the field $s^Z_{\boldsymbol{L}}$ using the simple foreground model outlined in the previous subsection.

We first note that the source cross-power spectra of T, E and B vanish, because the polarization angles are uniform and independent:

$$\langle s^X_{\boldsymbol{L}} s^Y_{-\boldsymbol{L}} \rangle \propto \delta^K_{XY}. \tag{3.20}$$

A more rigorous argument for this orthogonality goes as follows: when $X = T$ and $Y = E$ or B, the LHS of Eq. (3.20) vanishes since $\langle e^{\pm 2i\psi_i} \rangle = 0$. When $X = E$ and $Y = B$, we average over $\langle \cos(2\psi_i - 2\phi_{\boldsymbol{L}}) \sin(2\psi_i - 2\phi_{-\boldsymbol{L}}) \rangle \propto \sin(2\phi_{\boldsymbol{L}} - 2\phi_{-\boldsymbol{L}}) = 0$. Eq. (3.20) follows from these two observations.

With Eq. (3.20) at our disposal, it is straightforward to compute the responses. We start with the definition of the response g: $\langle s^X_{\boldsymbol{\ell}} s^Y_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle' = \sum_Z g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\, s^Z_{\boldsymbol{L}}$, we multiply both sides by $s^{Z'}_{-\boldsymbol{L}}$, and take an average:

$$g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{\langle s^X_{\boldsymbol{\ell}} s^Y_{\boldsymbol{L}-\boldsymbol{\ell}} s^Z_{-\boldsymbol{L}} \rangle}{\langle s^Z_{\boldsymbol{L}} s^Z_{-\boldsymbol{L}} \rangle}. \tag{3.21}$$

[6] In practice, we do not need to assume a specific distribution for the number of sources $N_{\rm sources}$ to compute the response. However, as soon as the sources are independent, the number of sources in any subpatch of the map will be Poisson distributed.

Using this expression and the foreground PDF (Eq. (3.19)), we can explicitly compute all unique 18 responses, which are given in Table 3.1. Similar responses to s^Q and s^U were recently derived in Appendix E of [18].

	$Z = T$	E	B
$XY = TT$	$\mathcal{A}$	0	0
TE	0	$\mathcal{B}\cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})$	$\mathcal{B}\sin(2\theta_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{L}})$
TB	0	$\mathcal{B}\sin(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})$	$\mathcal{B}\cos(2\theta_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{L}})$
EE	$\mathcal{C}\cos(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}})$	0	0
EB	$\mathcal{C}\sin(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}})$	0	0
BB	$\mathcal{C}\cos(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}})$	0	0

Table 3.1: The linear response $g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ for all unique combinations of XY (rows) and Z (columns). Note that all entries in the table should be multiplied by a geometric factor $u_{\boldsymbol{\ell}}u_{\boldsymbol{L}-\boldsymbol{\ell}}/u_{\boldsymbol{L}}$. For brevity we have defined $\mathcal{A} \equiv \langle s_i^3\rangle/\langle s_i^2\rangle$, $\mathcal{B} \equiv \langle p_i^2 s_i^3\rangle/\langle p_i^2 s_i^2\rangle$, and $\mathcal{C} \equiv \langle p_i^2 s_i^3\rangle/2\langle s_i^2\rangle$.

As an illustrative example we will calculate the response of TE to B for the specialized case of point sources ($u_{\boldsymbol{\ell}} = 1$). We first compute the bispectrum $\langle s^T_{\boldsymbol{\ell}} s^E_{\boldsymbol{L}-\boldsymbol{\ell}} s^B_{-\boldsymbol{L}}\rangle$, given by:

$$\begin{aligned}&\sum_i \langle s_i^3 p_i^2\rangle \int \frac{d\psi_i}{2\pi} \cos(2(\psi_i - \phi_{\boldsymbol{L}-\boldsymbol{\ell}}))\sin(2(\psi_i - \phi_{-\boldsymbol{L}}))\\ &= \frac{\langle N_{\rm sources}\rangle}{2}\langle s_i^3 p_i^2\rangle \underbrace{\sin(2(\phi_{\boldsymbol{L}-\boldsymbol{\ell}} - \phi_{-\boldsymbol{L}}))}_{=\sin(2\theta_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{L}})}.\end{aligned} \tag{3.22}$$

A similar calculation gives the power spectrum $\langle s^B_{\boldsymbol{L}} s^B_{-\boldsymbol{L}}\rangle$:

$$\begin{aligned}&\sum_i \langle s_i^2 p_i^2\rangle \int \frac{d\psi_i}{2\pi} \sin(2(\psi_i - \phi_{\boldsymbol{L}}))\sin(2(\psi_i - \phi_{-\boldsymbol{L}}))\\ &= \frac{\langle N_{\rm sources}\rangle}{2}\langle s_i^2 p_i^2\rangle \underbrace{\cos(2(\phi_{\boldsymbol{L}} - \phi_{-\boldsymbol{L}}))}_{=1}.\end{aligned} \tag{3.23}$$

The response $g^{TE,B}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ is given by the ratio of Eq. (3.22) to Eq. (3.23), corresponding to the value reported in Table 3.1. Note that the factor $\langle N_{\rm sources}\rangle$ cancels, which is why we do not explicitly need to assume the distribution of the number of sources.

Picking the wrong profile

Here we explore the expected bias to a bias-hardened estimator (hardened against a profile $u_{\boldsymbol{\ell}}$) from a foreground with profile $w_{\boldsymbol{\ell}} \neq u_{\boldsymbol{\ell}}$. For brevity we will consider the case of temperature-only reconstruction in the large-lens limit ($L \to 0$).

For the case of temperature-only reconstruction the optimal bias-hardened lensing weights take the form:

$$F^{\rm BH}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = F^{\rm MV}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \frac{\int_{\boldsymbol{\ell}'} F^{\rm MV}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'} u_{\ell'} u_{|\boldsymbol{L}-\boldsymbol{\ell}'|}}{\int_{\boldsymbol{\ell}'} G^{\rm MV}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'} u_{\ell'} u_{|\boldsymbol{L}-\boldsymbol{\ell}'|}} G^{\rm MV}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}, \tag{3.24}$$

as can be shown with Eq. (3.15). In Eq. (3.24) we have assumed that the input profile u_ℓ is isotropic, and $F^{\rm MV}$ ($G^{\rm MV}$) are the minimum variance temperature-only weights for the lensing convergence (foreground):

$$\begin{aligned} F^{\rm MV}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} / 2\tilde{C}^{TT}_\ell C^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|} \\ G^{\rm MV}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \mathcal{A} u_\ell u_{|\boldsymbol{L}-\boldsymbol{\ell}|} / u_L 2\tilde{C}^{TT}_\ell C^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}. \end{aligned} \tag{3.25}$$

As an aside, note that the bias-hardened weights (Eq. (3.24)) are insensitive to an overall normalization of $G^{\rm MV}$. Hence bias-hardening the TT estimator requires no knowledge of the source amplitude $\mathcal{A}$, or the normalization of the profile. In Appendix B.3 we show that this is always the case when bias-hardening any single-pair lensing estimator, but is no longer true when bias-hardening the global minimum variance estimator.

Suppose that the true foreground has an isotropic profile $w_\ell \neq u_\ell$. In the $L \to 0$ limit the response of a generic lensing estimator $\hat{\kappa}$ (with temperature weights F^κ and normalization N^κ) to the true foreground is:

$$\begin{aligned} \mathcal{R}^{\kappa,w}_{\boldsymbol{0}} &= N^\kappa_{\boldsymbol{0}} \int_{\boldsymbol{\ell}} F^\kappa_{\boldsymbol{\ell},-\boldsymbol{\ell}}\, g^{TT,T}_{\boldsymbol{\ell},-\boldsymbol{\ell}} \\ &= \mathcal{A} \int \frac{d\ell}{(2\pi)^2} w^2_\ell \left[N^\kappa_{\boldsymbol{0}} \ell \int d\theta_{\boldsymbol{L},\boldsymbol{\ell}} F^\kappa_{\boldsymbol{\ell},-\boldsymbol{\ell}} \right], \end{aligned} \tag{3.26}$$

such that the field level bias to $\langle \hat{\kappa}_{\boldsymbol{0}} \rangle'$ is $\mathcal{R}^{\kappa,w}_{\boldsymbol{0}} s^T_{\boldsymbol{0}}$. Here we have assumed that the profile is normalized so that $\lim_{\ell \to 0} u_\ell = 1$.

In the bottom panel of Fig. 3.1 we show the values of $|\mathcal{R}^{\kappa,w}_{\boldsymbol{0}}|$ when the true foreground profile w_ℓ is a Gaussian with a full width half maximum (FWHM) given by $\theta_{\rm FWHM}$ (in arcmin), and consider three cases for F^κ: minimum variance weights (red), hardening against a Gaussian profile (purple), and hardening against a point source (blue). The angular averaged values of these weights are plotted in the top panel of Fig. 3.1.

We find that point source hardening reduces the response to foregrounds smaller than ~ 2 arcmin by an order of magnitude relative to the MV weights. By instead hardening against a Gaussian with a non-zero width, one can further reduce (or precisely null) the response to foreevolvegrounds whose angular size is within $\sim 40\%$ of the size of the profile. In particular the tSZ-profile hardened estimator [519] (cyan), which is well approximated by the Gaussian-hardened estimator with a 2 arcmin FWHM, reduces the response to point sources by roughly a factor of 5 relative to the MV estimator.

Finally, if the foreground is larger than a few arcmin, the response is always small regardless of the assumed profile or technique, since the primary CMB dominates dominates on these scales (i.e. the lensing weights have no support for $\ell < 1500$).

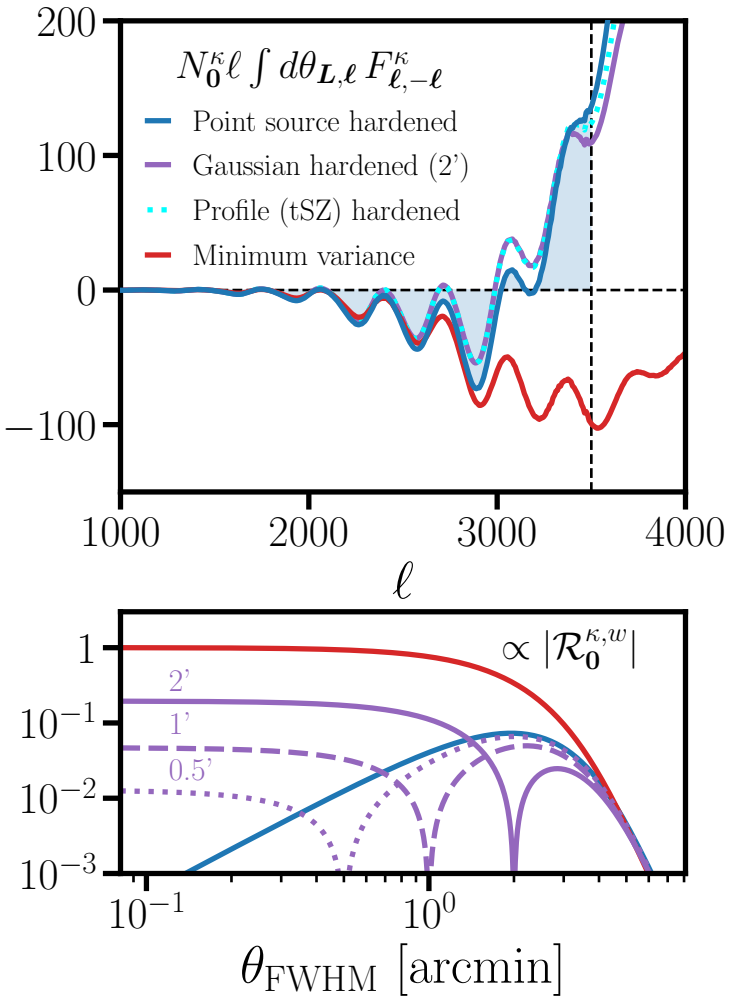


Figure 3.1: *Top:* The angular averaged temperature-only lensing weights ($\times 2\pi N_{\mathbf{0}}^{\kappa}\ell$) in the large-lens limit ($L \to 0$) when hardening against a point source (blue), hardening against a Gaussian profile with a $2'$ FHWM (purple) and when choosing the weights to minimize the variance (red). Note that since the point source hardened estimator has zero response to point-sources (when $w_\ell = 1$) the area under the blue curve vanishes. *Bottom:* The response to a Gaussian foreground w_ℓ with a full width half max $\theta_{\rm FHHM}$, normalized to the response of the MV estimator to point sources. In both plots we take $\ell_{\mathrm{max},T} = 3500$. See section 3.5 for details regarding the instrumental assumptions.

3.4 Large-lens suppressed estimators

In this section we study the "large-lens" or "squeezed" limit ($L \ll \ell$) analytically. A similar study for the temperature-only case has given rise to "shear-only" estimators [524], which

are highly effective at reducing foreground biases in this limit. Here we extend this study to polarization, highlighting important differences compared to temperature.

As in section 3.3, define $\mathcal{R}_{\boldsymbol{L}}^{XY,Z}$ as the response of a single-pair lensing estimator to a foreground in the Z'th map:

$$(N_{\boldsymbol{L}}^{XY})^{-1}\,\mathcal{R}_{\boldsymbol{L}}^{XY,Z} = \int \frac{\ell d\ell}{(2\pi)^2} \int d\theta_{\boldsymbol{L},\boldsymbol{\ell}}\, F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{XY}\, g_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{XY,Z}, \tag{3.27}$$

so that the field level bias to $\langle \hat{\kappa}_{\boldsymbol{L}}^{XY} \rangle'$ due to $s_{\boldsymbol{L}}^{Z}$ is $\mathcal{R}_{\boldsymbol{L}}^{XY,Z} s_{\boldsymbol{L}}^{Z}$. The inverse of the normalization can similarly be expressed as:

$$(N_{\boldsymbol{L}}^{XY})^{-1} = \int \frac{\ell d\ell}{(2\pi)^2} \int d\theta_{\boldsymbol{L},\boldsymbol{\ell}}\, F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{XY} f_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{XY}. \tag{3.28}$$

The large-lens suppressed (LLS) estimators are derived by choosing the weights to suppress the field level bias by $\mathcal{O}((L/\ell)^2)$ relative to the lensing signal, where $\ell \sim 3000$ is evaluated in the domain where the lensing estimator receives the majority of its signal (see the top panel of Fig. 3.1). This can be accomplished by forcing Eq. (3.27) to be suppressed by $\mathcal{O}((L/\ell)^2)$ while simultaneously retaining a $\mathcal{O}(1)$ scaling in the inverse normalization.

We first consider the temperature-only case. By Taylor expanding the angular-dependent piece of foreground linear response $g_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{TT,T} = \mathcal{A} u_\ell u_{\boldsymbol{L}-\boldsymbol{\ell}}/u_L$ in powers of L/ℓ (see Appendix B.1) we find:

$$g_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{TT,T} = \mathcal{A}\frac{u_\ell}{u_L}\left\{ u_\ell - L u_\ell' \cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{1}{4}\left(\frac{L}{\ell}\right)^2 \left[D_\ell^+ + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) D_\ell^- \right] u_\ell + \mathcal{O}\left((L/\ell)^3\right) \right\}, \tag{3.29}$$

with corrections going as $\mathcal{O}\left((L/\ell)^3\right)$, and we have defined the operators $D_\ell^\pm \equiv \ell^2 d^2/d\ell^2 \pm d/d\ln\ell$. Likewise, the linear response of the CMB to lensing also factorizes into a nice multipole form [486, 524]:

$$f_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{TT} = C_\ell^{TT}\left[\frac{d\ln \ell^2 C_\ell^{TT}}{d\ln\ell} + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) \frac{d\ln C_\ell^{TT}}{d\ln\ell} \right] \tag{3.30}$$

with corrections going as $\mathcal{O}(L/\ell)$.

In a similar manner one can expand the lensing weights into multipoles:

$$F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{TT} = \sum_{m=0}^{\infty} \left(A_{L,\ell}^{(m)} \cos(m\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + B_{L,\ell}^{(m)} \sin(m\theta_{\boldsymbol{L},\boldsymbol{\ell}}) \right), \tag{3.31}$$

and in turn express the coefficients $A^{(m)}$ and $B^{(m)}$ as a Taylor expansion in L/ℓ. From Eq. (3.29), we see that setting $A^{(0)} = \mathcal{O}((L/\ell)^2)$, $A^{(1)} = \mathcal{O}(L/\ell)$ forces the response to foregrounds (3.27) to be suppressed by $\mathcal{O}((L/\ell)^2)$. From Eq. (3.30), we see that retaining a $\mathcal{O}(1)$ scaling in Eq. (3.28) requires choosing $A^{(2)} = \mathcal{O}(1)$. The remaining coefficients are

chosen to minimize the noise of the LLS estimator in the large-lens limit. The large-lens noise (Eq. (3.6) as $L \to 0$) of the temperature-only LLS estimator takes the form:

$$\mathcal{N}_0^{TT} = \frac{\int \frac{\ell d\ell}{2\pi} (C_\ell^{TT})^2 \left[\sum_{m=2}^{\infty} \left(A_{0,\ell}^{(m)}\right)^2 + \sum_{m=1}^{\infty} \left(B_{0,\ell}^{(m)}\right)^2\right]}{\left[\int \frac{\ell d\ell}{4\pi} A_{0,\ell}^{(2)} \frac{dC_\ell^{TT}}{d\ln\ell}\right]^2}. \tag{3.32}$$

Note that the $A^{(2)}$-term is the only term contributing to the normalization (since we set $A_{0,\ell}^{(0)} = 0$). Thus a non-zero value for the remaining coefficients ($A^{(m)}$ for $m > 2$ and all $B^{(m)}$) only increases the noise of the LLS estimator. To minimize the large-lens noise, we therefore choose the remaining coefficients to be $\mathcal{O}(L/\ell)$, and solve for the value of $A_{0,\ell}^{(2)}$ which minimizes Eq. (3.32). Doing so forces the LLS estimator to take the form:

$$F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{\mathrm{LLS},TT} = \frac{1}{2(\tilde{C}_\ell^{TT})^2} \frac{dC_\ell^{TT}}{d\ln\ell} \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \mathcal{O}(L/\ell), \tag{3.33}$$

recovering the shear-only estimator [524]. As was shown in [491], while the shear-only estimator suppresses the bias by $\mathcal{O}((L/\ell)^2)$, the bias *is not precisely nulled* for a source with a generic profile, even if the profile is azimuthally symmetric. We further note that Eq. (3.29) also contains derivatives of the (Fourier-transformed) profile, which additionally suppress biases for small sources. In the special case of point sources, for which these derivatives vanish, the shear-only estimator is non-perturbatively unbiased.

We explicitly generalize this procedure for all pairs of temperature and polarization maps in Appendix B.5. Our results are summarized by:

$$\begin{aligned}
F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{\mathrm{LLS},TE} &= \frac{C_\ell^{TE}}{\tilde{C}_\ell^{TT}\tilde{C}_\ell^{EE}} \frac{d\ln\ell^2|C_\ell^{TE}|}{d\ln\ell} + \mathcal{O}(L/\ell) \\
F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{\mathrm{LLS},TB} &= \mathcal{O}(L/\ell) \\
F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{\mathrm{LLS},EE} &= \frac{1}{2(\tilde{C}_\ell^{EE})^2} \frac{dC_\ell^{EE}}{d\ln\ell} \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \mathcal{O}(L/\ell) \\
F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{\mathrm{LLS},EB} &= F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{\mathrm{MV},EB} = \frac{f_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{EB}}{\tilde{C}_\ell^{EE}\tilde{C}_{|\boldsymbol{L}-\boldsymbol{\ell}|}^{BB}}.
\end{aligned} \tag{3.34}$$

One can continue this procedure to obtain LLS estimators to any order in L/ℓ. Including higher order corrections will in principle make the estimators less noisy, however, by construction the approximations in Eqs. (3.33) and (3.34) are only slightly suboptimal in the regime where the large-lens expansion is valid.

As indicated by Eq. (3.34), we find that there is no leading order LLS TB estimator. As shown in Appendix B.1 and B.5, this is because $f^{TB} \propto g^{TB,E} \propto \sin(2\theta)$ to lowest order in the large-lens limit, i.e. foregrounds and lensing are indistinguishable to lower order. On the contrary we find that the MV EB estimator already satisfies the LLS requirements.

	$\mathcal{O}$(MV biases)	$\mathcal{O}$(LLS suppression)
TT	1	$L^2 u''_\ell/u_\ell$, $L^2 u'_\ell/\ell u_\ell$
TE	$p_{\rm source}/p_{\rm CMB}$	$L^2 u''_\ell/u_\ell$, $L^2 u'_\ell/\ell u_\ell$
TB	$p_{\rm source}/p_{\rm CMB}$	–
EE	$(p_{\rm source}/p_{\rm CMB})^2$	$(L/\ell)^2$
EB	$(L/\ell)^2 \times (p_{\rm source}/p_{\rm CMB})^2$	–

Table 3.2: In the first column we show scaling of the field level biases for each single-pair MV estimator relative to the MV TT field-level bias. Here $p_{\rm source}$ refers to the polarization fraction of the foreground, while $p_{\rm CMB} \sim \sqrt{C_\ell^{EE}/C_\ell^{TT}}$, where ℓ should be evaluated where the lensing weights receive the majority of their support ($\ell \sim 3000$). In the second column we show the scaling for the field-level bias suppression when the MV estimators are substituted for LLS estimators. Note that there is no LLS TB estimator to leading order in L/ℓ, while the LLS EB estimator is equivalent to the MV EB estimator.

We show the scaling of the field-level biases for all single-pair estimators in Table 3.2. In the first column we show the scaling of the minimum variance (MV) field-level bias (relative to the MV TT bias) with the polarization fractions of the foreground and CMB. In the second column we show the scaling of the field-level bias suppression for each LLS estimator. Note that while all LLS estimators reduce the bias by $\mathcal{O}((L/\ell)^2)$, the TT and TE estimators additionally reduce the bias by factors involving derivatives of the Fourier transform of the profile (e.g. $\ell u'_\ell \equiv \ell\, du_\ell/d\ell$). For the special case of point sources ($u_\ell = 1$) these derivatives vanish, and the (TT, TE) LLS estimators are unbiased to all orders in L/ℓ.

3.5 Noise properties

Here we compute the variance of the estimators discussed in sections 3.2, 3.3 and 3.4. We consider a full-sky CMB experiment at a single frequency (150 GHz) with stage-4 like sensitivity ($\Delta_T = 1$ μK-arcmin, $\Delta_P = \sqrt{2}\Delta_T$), and assume a Gaussian beam with a 1.4 arcmin FWHM. Throughout we take $\ell_{{\rm max},T} = 3500$ and $\ell_{{\rm max},P} = 5000$.

In our total temperature power spectrum $\tilde{C}_\ell^{TT}$ we include the lensed CMB, instrument noise, and contributions from extragalactic foregrounds (tSZ, kSZ, CIB and radio point sources) using `LensQuEst`[7], which adopts templates for the extragalactic foreground power spectra from [200]. With the exception of tSZ $\times$ CIB, we neglect the cross-correlation of the components listed above. For polarization-only spectra and for cross-correlations of temperature and polarization, we include the lensed CMB and instrument noise (where appropriate) in our total power spectra, but neglect extragalactic foregrounds. We use `symlens`[8] to implement our estimators and to compute noise curves.

[7] https://github.com/EmmanuelSchaan/LensQuEst
[8] https://github.com/simonsobs/symlens

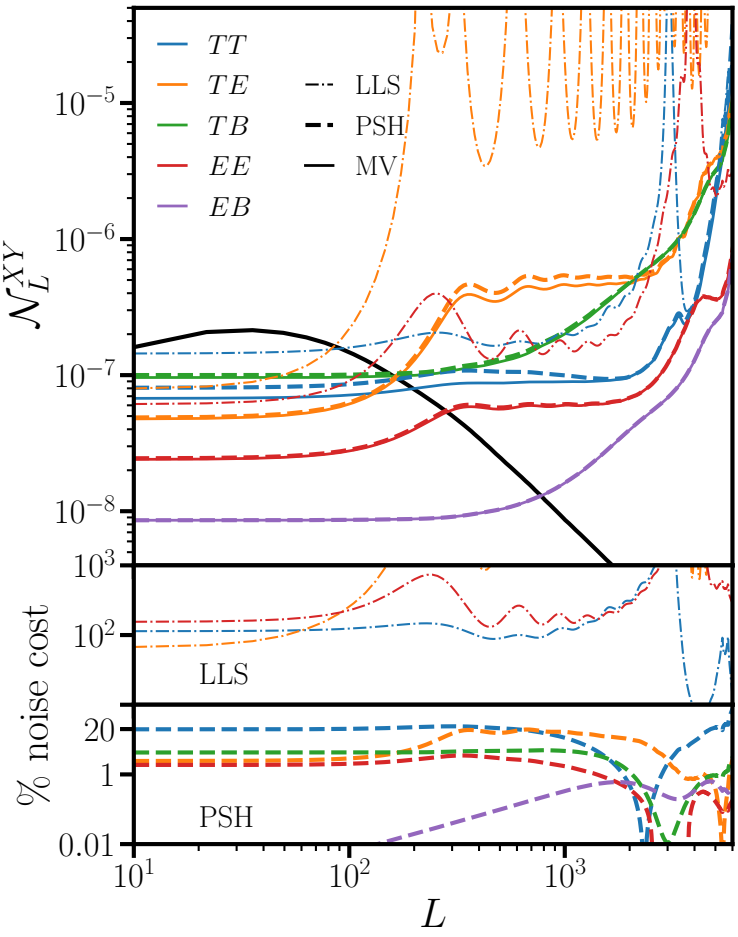


Figure 3.2: The top panel shows the reconstruction noise (per Fourier mode) when reconstructing on a single pair of maps. The colors denote the pair XY. Solid, dashed and dashed-dotted curves represent minimum variance (MV), point source hardened (PSH) or large-lens suppressed (LLS) estimators respectively. $C_L^{\kappa\kappa}$ is shown in black. The bottom two panels show the percent noise cost for the LLS and PSH estimators relative to the appropriate MV estimator.

We first consider the variance of single-pair estimators, given by Eq. (3.6) and plotted in Fig. 3.2. We see that point source hardening comes at a $1-20\%$ noise cost (relative to the MV estimators) for most pairs of maps, and that bias-hardening polarization-based estimators generically comes at a lower noise cost than the temperature-only case. In particular, the noise cost for hardening EB against point sources is less than 0.1% for all signal dominated scales ($L < 1000$). By contrast, the single-pair LLS estimators come at more modest factor of ~ 2 noise cost for low L. As shown in Appendix B.5, all three non-trivial LLS estimators (TT, TE, EE) are only sensitive to one of the two terms (monopole and quadrupole) in

the lensing linear response in the large lens limit. Since these two terms are comparable in magnitude, the LLS estimators are only sensitive to roughly half of the lensing signal, providing some intuition for the approximate factor of 2 cost in noise[9].

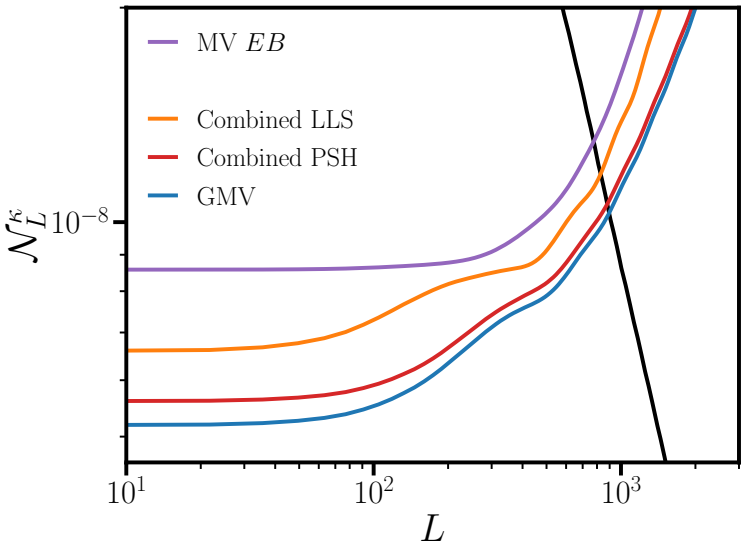


Figure 3.3: Lensing reconstruction noise for the optimal linear combination of single-pair LLS (orange) and point source hardened (PSH, red) estimators. The noise for the global MV (GMV) estimator is shown in blue. For reference we show the MV EB noise (purple) and $C^{\kappa\kappa}$ (black).

We next consider the reconstruction noise for combined estimators, shown in Fig. 3.3. In blue we show the reconstruction noise for the Global Minimum Variance estimator (i.e. Eq. (3.9)), whose weights are explicitly given by Eq. (B.7). In orange and red we show the reconstruction noises for optimal linear combinations (see Appendix B.3) of the single-pair LLS and point source hardened (PSH) estimators respectively. In particular, we find that near the transition of signal- to noise-domination ($L \sim 1000$) the combined PSH estimator's noise is only 4% larger than the GMV's. The combined LLS estimator comes at more modest $\sim 20\%$ noise penalty.

In Appendix B.3 we explore the feasibility of constructing global forms of the LLS and PSH estimators, however, we find that an optimal linear combination of single-pair estimators is both more practical and nearly optimal in both cases.

[9]That is, from Eq. (3.32) we see that the large-lens noise very roughly scales as (# of multipoles)/(# of multipoles)^2, assuming that each multipole is of a similar size (and sign).

3.6 Bias reduction

Biases to the CMB lensing power spectrum can be decomposed into primary, secondary and trispectrum components [227, 446]. For simplicity consider the case of a single-pair estimator $\hat{\kappa}^{XY}_{\boldsymbol{L}}$ which is used to measure $\hat{C}^{\kappa\kappa}_L \propto \langle \hat{\kappa}^{XY}_{\boldsymbol{L}} \hat{\kappa}^{XY}_{-\boldsymbol{L}} \rangle$ from a set of maps containing the lensed CMB $(\tilde{X}, \tilde{Y})$ and the foregrounds of interest (s^X, s^Y). Schematically, let $\hat{\kappa}[X, Y]$ denote reconstruction with the estimator $\hat{\kappa}^{XY}_{\boldsymbol{L}}$ on the maps $X_{\boldsymbol{\ell}}$, $Y_{\boldsymbol{L}-\boldsymbol{\ell}}$. Using this notation the primary, secondary and trispectrum biases take the form:

$$\begin{aligned}
\text{Primary} &= 2\big\langle \kappa\, \hat{\kappa}[s^X, s^Y] \big\rangle \\
\text{Secondary} &= 2\big\langle \hat{\kappa}[\tilde{X}_0, s^Y] \hat{\kappa}[\tilde{X}_1, s^Y] \big\rangle + 2\big\langle \hat{\kappa}[s^X, \tilde{Y}_0] \hat{\kappa}[s^X, \tilde{Y}_1] \big\rangle \\
&\quad + 2\big\langle \hat{\kappa}[s^X, \tilde{Y}_0] \hat{\kappa}[\tilde{X}_1, s^Y] \big\rangle \\
&\quad + 2\big\langle \hat{\kappa}[s^X, \tilde{Y}_1] \hat{\kappa}[\tilde{X}_0, s^Y] \big\rangle \\
\text{Trispectrum} &= \big\langle \hat{\kappa}[s^X, s^Y] \hat{\kappa}[s^X, s^Y] \big\rangle_c,
\end{aligned} \tag{3.35}$$

where we have expanded $\tilde{X} = \tilde{X}_0 + \tilde{X}_1 + \cdots$ in powers of κ to remove noise bias in our calculation of the secondary [524]. The Gaussian component of the foreground four-point function is subtracted analytically to obtain the trispectrum, as indicated by the c subscript.

From field-level to power spectrum bias suppression

Both bias-hardened and LLS estimators aim to suppress the field-level response to foregrounds, which schematically take the form $\int Fg$, where F is the lensing weights and g is the foreground linear response. Here we relate these field-level responses to the primary and trispectrum biases to the CMB lensing power spectrum.

The primary bias ultimately arises from a non-zero $\langle s_{\boldsymbol{\ell}} s_{\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{-\boldsymbol{L}} \rangle$ bispectrum. Since it is trivially the case that $\langle s_{\boldsymbol{\ell}} s_{\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{-\boldsymbol{L}} \rangle = \langle \langle s_{\boldsymbol{\ell}} s_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle' \kappa_{-\boldsymbol{L}} \rangle$, one can expand this bispectrum in terms of the linear foreground response and the foreground-lensing cross-power spectrum:

$$\langle s^X_{\boldsymbol{\ell}} s^Y_{\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{-\boldsymbol{L}} \rangle \propto \sum_Z g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\, \langle s^Z_{\boldsymbol{L}} \kappa_{-\boldsymbol{L}} \rangle. \tag{3.36}$$

As a consequence the primary bias is proportional to the field-level response. Thus choosing the weights F to null the field-level response also nulls the primary bias.

Likewise, for a foreground obeying the assumptions discussed in Section 3.3, we also find that the trispectrum can be expressed as a sum over power spectra, this time weighted by two factors of the foreground linear response:

$$\langle s^X_{\boldsymbol{\ell}} s^Y_{\boldsymbol{L}-\boldsymbol{\ell}} s^M_{\boldsymbol{\ell}'} s^N_{-\boldsymbol{L}-\boldsymbol{\ell}'} \rangle \propto \sum_Z g^{XY,Z}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\, g^{MN,Z}_{\boldsymbol{\ell}',-\boldsymbol{L}-\boldsymbol{\ell}'}\, \langle s^Z_{\boldsymbol{L}} s^Z_{-\boldsymbol{L}} \rangle, \tag{3.37}$$

when at least one of the maps is a temperature map. Thus, for these cases, the trispectrum bias is proportional to the field-level response squared. Polarization-only trispectra also

contain a term proportional to either $\cos(2\Phi)$ or $\sin(2\Phi)$, where $\Phi \equiv \phi_{\boldsymbol{\ell}} + \phi_{\boldsymbol{L}-\boldsymbol{\ell}} - \phi_{\boldsymbol{\ell}'} - \phi_{-\boldsymbol{L}-\boldsymbol{\ell}'}$, in addition to those in Eq. (3.37). In practice, we find that these extra terms are highly subdominant to those in Eq. (3.37), as shown in Fig. 3.5.

One can also obtain expressions analogous to Eq. (3.36) for the secondary bias, which is sourced by the same $\langle ss\kappa \rangle$ bispectrum. However, the angular structure of the secondary bias differs significantly from that of the primary, and as a result there is no simple relation between the field-level response and the secondary bias. Thus we expect nulling the field-level bias to null both the primary and trispectrum, but not the secondary. We outline a new approach to null the secondary in Appendix B.4, but leave further development to future work.

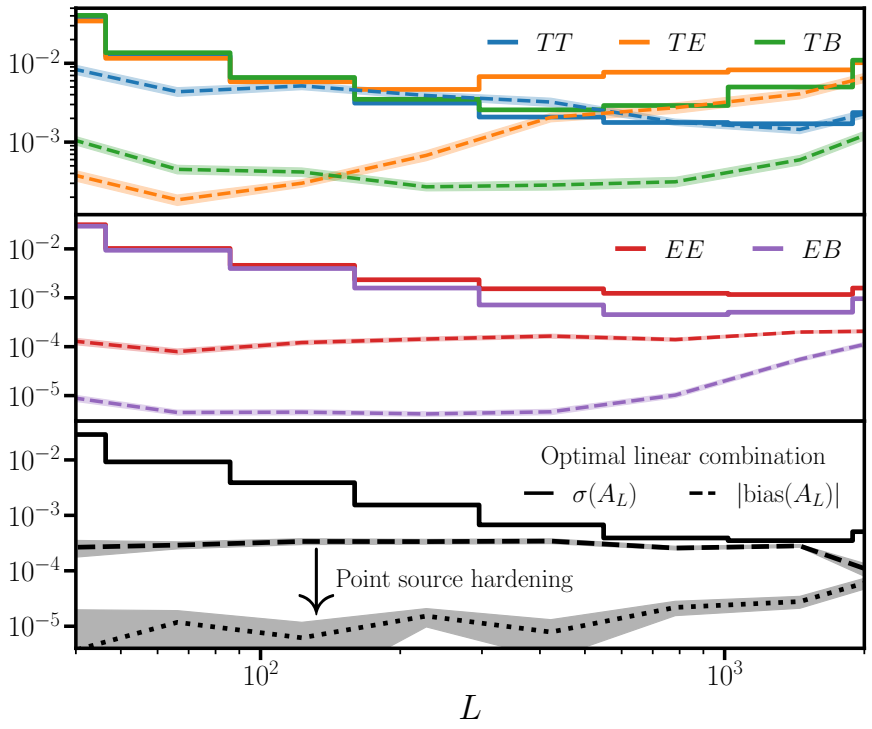


Figure 3.4: *Top two panels:* Solid lines show the 1σ error of the lensing amplitude ($A_L \equiv C_L^{\kappa\kappa,\text{measured}}/C_L^{\kappa\kappa,\text{true}}$) for each single-pair minimum variance (MV) estimator. Colors denote pairs of temperature and polarization. Dashed lines show the simulated (absolute) biases to the MV estimators due to radio point sources, with the shaded regions denoting the 1σ scatter of the mean. *Bottom:* Error (solid) and bias (dashed) to the lensing amplitude for the optimal linear combination of MV single-pair estimators. The dotted curve shows the bias for an optimal combination of point source hardened estimators. See sections 3.5 and 3.6 for details regarding the experimental configuration and simulations respectively.

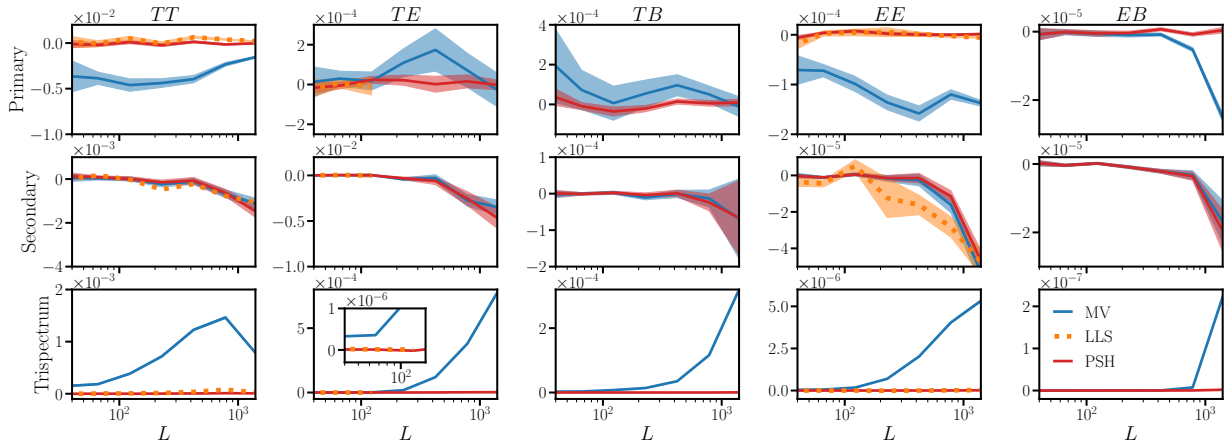


Figure 3.5: Simulated relative biases to the CMB lensing power spectrum from point sources (assuming a 5 mJy flux cut and 3% polarization fraction) for single-pair lensing estimators with $\ell_{\mathrm{max},T} = 3500$ and $\ell_{\mathrm{max},P} = 5000$, and CMB-S4 like sensitivity. Columns correspond to different pairs of temperature and polarization, while the rows indicate primary, secondary or trispectrum biases. Biases for the Minimum Variance (MV), Large-Lens Suppressed (LLS) and Point Source Hardened (PSH) estimators are shown in blue, orange and red respectively. Note that for the TE LLS estimator we only show biases for $L < 120$ since this estimator's noise dramatically increases at smaller scales (see Fig. 3.2), and that the TE, TB primary biases vanish analytically. The assumed L-binning is depicted in Fig. 3.4, and the shaded regions denote the 1σ scatter of the mean biases across 40 cutouts. We stress that the TT estimator here only shows the contribution from point sources: the overall bias is significantly larger when contributions from the CIB, tSZ and kSZ are included [519].

Power spectrum biases: simulations

To numerically estimate the biases in Eq. (3.35) we use the WebSky simulations of both the lensing convergence [566] and radio galaxies [351] (in temperature). These simulations accurately reproduce number counts measured from current CMB and radio telescopes, and have a realistic correlation between the galaxies and lensing convergence. Following [524, 519] we mask sources brighter than 5 mJy by fulling excluding the pixels associated with those sources.

We generate mock polarization data by assigning random polarization angles (sampled from a uniform distribution) to each source. For simplicity we assume that all sources have the same 3% polarization fraction. The measured biases presented below tend to scale simply with moments of the polarization fraction distribution, e.g. the EE trispectrum bias $\propto \langle p^4 \rangle$. Thus one can scale our results to any desired polarization fraction distribution.

We divide a subsample of the WebSky simulation into 40 square 15 degree cutouts using

`pixell`[10]. For each cutout we generate mock Gaussian CMB maps with `LensQuEst` and lens these maps using the corresponding WebSky convergence cutout. We run the lensing reconstruction for each single-pair estimator under the flat sky approximation using `symlens`, assuming the S4-like configuration described in section 3.5.

First, in top two panels of Fig. 3.4 we show the sensitivity[11] to the lensing amplitude of each single-pair minimum variance (MV) estimator, along with the total bias (primary + secondary + trispectrum) from polarized point sources. We note that these biases are significant ($\sim 0.5-1\sigma$) for both the TT and TE estimators, while for the remaining single-pair estimators they can approach $\sim 0.1\sigma$ at small ($L > 500$) scales. The bias to the optimal linear combination of single-pair MV estimators (bottom panel) is $\sim 1\sigma$ for $L > 500$, and non-negligible on larger scales. These biases are suppressed by up to two orders of magnitude when one instead takes an optimal linear combination of point source hardened estimators, ensuring bias $< \sigma/50$ for $L < 500$ and $< \sigma/20$ for $500 < L < 1000$. We show an analogous figure for a Simons Observatory-like survey in Appendix B.7, and again find that the biases to TT and TE are significant ($\sim 1\sigma$), while for the remaining estimators the biases are negligible ($< \sigma/100$).

In Fig. 3.5 we show the individual (relative) primary, secondary and trispectrum biases for the Minimum Variance (MV; blue), Point Source Hardened (PSH; orange), and Large-Lens Suppressed estimators (LLS; red). We find that both point source hardening and large-lens suppression effectively null the primary and trispectrum biases for all single-pair estimators, but have little impact on the secondary bias, as expected. Since the total TT, TB, EE, and EB biases receive sizeable contributions from the primary and trispectrum terms, one significantly reduces the overall bias for these estimators. However, the TE bias is dominated by the secondary $\langle s^T s^T \kappa \rangle$ contribution, and thus the TE bias is only marginally reduced using these techniques. If necessary one can null this secondary contribution using the technique outlined in Appendix B.4. As indicated in the bottom panel of Fig. 3.4, the TE secondary is a subdominant contribution to the bias from the optimal combination of PSH estimators, and we therefore expect that an optimal linear combination of PSH estimators is sufficient for a robust and precise measurement of $C^{\kappa\kappa}$ with future surveys.

We note that there is another population of "Infrared" (IR) point sources which are dominant at higher frequency and have not been included in the simulations described here, since their polarization fraction is expected to be smaller than for radio sources [168, 263]. The same techniques discussed here should also mitigate the impact of polarized IR sources with similar effectiveness.

[10] https://github.com/simonsobs/pixell

[11] The sensitivities quoted here are for the assumed $L-$binning in Fig. 3.4, and would be more significant if broader bins were used.

3.7 Discussion and conclusions

In this chapter we generalized previous work to polarization, showing that geometric methods are effective at reducing the bias to CMB lensing in both temperature and polarization, even on single-frequency maps. We used state-of-the art simulations of radio point sources to estimate the size of these biases and the effectiveness of our mitigation strategies. We have found that the point-source bias in polarization is relatively small for upcoming experiments such as Simons Observatory (Appendix B.7), but it can become potentially significant on small scales for future low-noise experiments (Fig. 3.4). We have shown that point source hardened estimators, which come at a negligible noise cost, can very effectively control these biases for all experiments in the foreseeable future. We find that LLS estimators are equally as effective at reducing these biases, however, generically come at a significantly higher noise cost.

We note that patchy reionization is also expected to give a contribution to the temperature fluctuations (through the kSZ effect) and to polarization, through patchy screening and scattering [202]. While we expect the biases induced by reionization to be subdominant (this has been shown explicitly for temperature in [114]), we note that that the majority of the reionization bubbles will be sub-beam in size, and therefore we expect PSH to decrease any possible bias. A full study of this is left to future work.

In this study we have not explored the impact of *galactic* foregrounds and potential mitigation strategies. In [73], the authors show that foreground-induced biases can be significant in polarization, and consider multi-frequency cleaning as a partial mitigation strategy. While we note that the simple geometry of point sources allowed us to analytically write estimators with zero response to them, some of the same techniques presented here may prove useful when dealing with galactic foregrounds. The lensing-induced distortion provides a consistency between the different angular multipoles discussed in Section 3.4. Due to their complex geometry, galactic foregrounds are unlikely to match such a consistency and therefore they may be distinguishable using geometric methods such as the ones discussed in this chapter. Finally, it's worth noting that the same dust seen in emission in the CMB maps also causes extinction at optical wavelengths, potentially modulating the observed galaxy field and creating biases to cross-correlations. This further motivates this line of research, which is left to a future investigation.

Acknowledgments

We thank Julien Carron, Antón Baleato Lizancos, Zack Li, Marius Millea, Toshiya Namikawa and Giuseppe Puglisi for useful discussions. NS and SF are supported by the Physics Division of Lawrence Berkeley National Laboratory and by the U.S. Department of Energy (DOE), Office of Science, under contract No. DE-AC02-05CH11231. ES was supported by the Chamberlain fellowship at Lawrence Berkeley National Laboratory, and is now supported by SLAC National Accelerator Laboratory.

Chapter 4

Optimal multifrequency weighting

This chapter was originally published as

> Noah Sailer, Emmanuel Schaan, Simone Ferraro, Omar Darwish, and Blake Sherwin. "Optimal multifrequency weighting for CMB lensing". In: *Phys. Rev. D* 104.12 (2021), p. 123514. DOI: 10.1103/PhysRevD.104.123514. arXiv: 2108.01663 [astro-ph.CO]

Extragalactic foregrounds in Cosmic Microwave Background (CMB) temperature maps lead to significant biases in CMB lensing reconstruction if not properly accounted for. Combinations of multifrequency data have been used to minimize the overall map variance (internal linear combination, or ILC), or specifically null a given foreground, but these are not tailored to CMB lensing. In this chapter, we derive an optimal multifrequency combination to jointly minimize CMB lensing noise and bias. We focus on the standard lensing quadratic estimator, as well as the "shear-only" and source-hardened estimators, whose responses to foregrounds differ. We show that an optimal multifrequency combination is a compromise between the ILC and joint deprojection, which nulls the thermal Sunyaev-Zel'dovich (tSZ) and Cosmic Infrared Background (CIB) contributions. In particular, for a Simons Observatory-like experiment with $\ell_{\mathrm{max},T} = 3000$, we find that profile hardening alone (with the standard ILC) reduces the bias to the lensing power amplitude by 40%, at a 20% cost in noise, while the bias to the cross-correlation with a LSST-like sample is reduced by nearly an order of magnitude at a 10% noise cost, relative to the standard quadratic estimator. With a small amount of joint deprojection the bias to the profile hardened estimator can be further reduced to less than half the statistical uncertainty on the respective amplitudes, at a 20% and 5% noise cost for the auto- and cross-correlation respectively, relative to the profile hardened estimator with the standard ILC weights. Finally, we explore possible improvements with more aggressive masking and varying $\ell_{\mathrm{max},T}$.

4.1 Introduction

CMB lensing [94, 342, 279] allows us to reconstruct the projected mass distribution all the way to the surface of last scattering, enabling cosmological inference (e.g. neutrino masses, dark matter and dark energy) using a very well-understood source redshift and redshift weighting. Extragalactic foregrounds, which are correlated with the lensing potential, are likely to be one of the main limiting factors [597, 446, 374, 227, 524, 71, 519] in temperature-based CMB lensing reconstruction, producing significant biases to the standard minimum variance quadratic estimator (QE, [300]), which in turn bias both the reconstructed lensing power spectrum and the cross-correlations with tracers of the large-scale structure. In this chapter we consider two techniques to mitigate these biases: geometric-based methods that distinguish lensing from foregrounds through their different spatial symmetries (shear-only reconstruction [524] and bias-hardening [519, 446]), and multifrequency based methods that exploit the different frequency scaling of the CMB and extragalactic foregrounds [175, 374]. In particular, we explore the *lensing-optimized* linear combinations of frequency maps to minimize a combination of bias and noise, for the standard QE, as well as for shear and bias-hardened estimators. Another promising method to control foreground biases to CMB lensing is the gradient-cleaned [374] or symmetrized estimator [165, 457]. We explore these methods in a companion paper [166]. The remainder of the chapter is organized as follows: in §4.2 we introduce the standard, shear-only, and bias-hardened quadratic estimators; in §4.3 we review multifrequency methods for noise and bias mitigation; in §4.4 (4.5) we outline the calculation of the noise (bias) of a general quadratic estimator; in §4.6 we combine geometric- and frequency-based techniques to minimize a combination of lensing noise and bias; we conclude in §4.7.

4.2 CMB lensing quadratic estimators

In the presence of a foreground $s_{\boldsymbol{\ell}}$, the observed temperature covariance receives off-diagonal contributions from lensing and the foreground[1] (see Eqs. 3-7 in [519]):

$$\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}} \rangle = f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \kappa_{\boldsymbol{L}} + f^{s}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} s_{\boldsymbol{L}} \tag{4.1}$$

to lowest order in κ and s. The standard quadratic estimator (QE, [300]) neglects the foreground term and seeks the minimum variance solution for $\kappa_{\boldsymbol{L}}$ in Eq. (4.1). By neglecting the foreground contribution, it can lead to a biased answer [524, 519].

On small scales where the foreground contributions are large, they are well approximated by a collection of unclustered halos or galaxies with azimuthally-symmetric profiles. In this regime, the shear-only κ estimator, which is insensitive to such isotropic terms has been shown to significantly reduce all foreground-induced biases [524].

[1] $T_{\boldsymbol{\ell}} = T^{\rm CMB}_{\boldsymbol{\ell}} + s_{\boldsymbol{\ell}}$, where $T^{\rm CMB}_{\boldsymbol{\ell}}$ is the *lensed* primary CMB, and the average is taken over realizations of the unlensed CMB at fixed $\kappa_{\boldsymbol{L}}$ and $s_{\boldsymbol{L}}$.

Bias-hardening [446, 434, 519] seeks to measure the foreground contribution directly and "subtract off" the bias to the lensing convergence $\kappa_{\boldsymbol{L}}$. Since both $\kappa_{\boldsymbol{L}}$ and $s_{\boldsymbol{L}}$ produce off-diagonal covariances, the standard quadratic estimators for κ and s are biased [446, 519]:

$$\begin{pmatrix} \langle \hat{\kappa}_{\boldsymbol{L}} \rangle \\ \langle \hat{s}_{\boldsymbol{L}} \rangle \end{pmatrix} = \begin{pmatrix} 1 & N^{\kappa}_{\boldsymbol{L}} \mathcal{R}_{\boldsymbol{L}} \\ N^{s}_{\boldsymbol{L}} \mathcal{R}_{\boldsymbol{L}} & 1 \end{pmatrix} \begin{pmatrix} \kappa_{\boldsymbol{L}} \\ s_{\boldsymbol{L}} \end{pmatrix} \tag{4.2}$$

where the response $\mathcal{R}_{\boldsymbol{L}}$ is defined in [519]. A simple matrix inversion leads to unbiased estimates of the lensing and foreground:

$$\begin{pmatrix} \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \\ \hat{s}^{\rm BH}_{\boldsymbol{L}} \end{pmatrix} = \begin{pmatrix} 1 & N^{\kappa}_{\boldsymbol{L}} \mathcal{R}_{\boldsymbol{L}} \\ N^{s}_{\boldsymbol{L}} \mathcal{R}_{\boldsymbol{L}} & 1 \end{pmatrix}^{-1} \begin{pmatrix} \hat{\kappa}_{\boldsymbol{L}} \\ \hat{s}_{\boldsymbol{L}} \end{pmatrix}. \tag{4.3}$$

We assume that the foreground s can be approximated by a collection of sources with identical profiles $u_{\boldsymbol{\ell}}$ [519]. Explicitly, we assume $s_{\boldsymbol{\ell}} = \sum_i s_i e^{i\boldsymbol{\ell}\cdot\boldsymbol{x}_i} u_{\boldsymbol{\ell}}$, where s_i ($\boldsymbol{x}_i$) is the amplitude (position) of the i'th source. In this chapter we consider either hardening against point sources [446, 434, 519] (point source hardening, or PSH), where the profile is taken to be a delta function ($u_{\boldsymbol{\ell}} = 1$), or against a tSZ-like profile [519] (profile hardening, or PH), where the profile is taken to be the square root of the tSZ power spectrum at 150 GHz.

4.3 Internal linear combination (ILC) methods

Let $(\boldsymbol{T}_{\boldsymbol{\ell}})_i$ denote the observed map in the i-th frequency channel. We assume that $\boldsymbol{T}_{\boldsymbol{\ell}}$ can be written in the form:

$$\boldsymbol{T}_{\boldsymbol{\ell}} = \mathbf{1} T^{\rm CMB}_{\ell} + \boldsymbol{n}_{\boldsymbol{\ell}} + \sum_s \boldsymbol{A}_s s_{\boldsymbol{\ell}}, \tag{4.4}$$

where $\mathbf{1} = (1, 1, \cdots, 1)^T$ and $T^{\rm CMB}_{\boldsymbol{\ell}}$ is the *lensed* CMB. The "noise" $\boldsymbol{n}_{\boldsymbol{\ell}}$ receives contributions from the detectors' thermal fluctuations, the atmosphere, and galactic dust. The extragalactic foregrounds $s_{\boldsymbol{\ell}}$ include tSZ, kSZ, CIB, and radio point sources. In Eq. (4.4) we assume that each map has been calibrated to have unit response to the lensed CMB. The lensed CMB in each frequency channel is thus determined by a single template $T^{\rm CMB}_{\boldsymbol{\ell}}$. Similarly, we assume that each foreground s follows a single spatial template $s_{\boldsymbol{\ell}}$, up to a frequency dependence $\boldsymbol{A}_s$. Our normalized linear estimator for the lensed CMB takes the form $\hat{T}_{\boldsymbol{\ell}} = \boldsymbol{w}^T_{\ell} \boldsymbol{T}_{\boldsymbol{\ell}}$ for some weights $\boldsymbol{w}_{\ell}$, which we are free to choose subject to the constraint $\boldsymbol{w}^T_{\ell} \mathbf{1} = 1$. The noise of the estimator $\hat{T}_{\ell}$ is

$$C^{\rm tot}_{\ell} = \boldsymbol{w}^T_{\ell} \boldsymbol{C}_{\ell} \boldsymbol{w}_{\ell}, \tag{4.5}$$

where $(2\pi)^2 \delta^D_{\mathbf{0}} \boldsymbol{C}_{\ell} = \langle \boldsymbol{T}_{\boldsymbol{\ell}} \boldsymbol{T}^T_{-\boldsymbol{\ell}} \rangle$ is the covariance matrix of the maps $\boldsymbol{T}_{\boldsymbol{\ell}}$.

In this chapter we consider a Simons Observatory-like experiment, which will observe the CMB in six frequency channels: 27, 39, 93, 145, 225, and 280 GHz. We follow [22] in modeling the detector noise and atmospheric contributions to $\boldsymbol{C}_{\boldsymbol{\ell}}$, assuming the "goal" noise levels. The theory power spectra and SEDs for the galactic dust and extragalactic foregrounds are taken from [200].

Before deriving the optimal linear combination for CMB lensing, we start by reviewing the case of the standard ILC and deprojection. These two limiting cases, commonly used in CMB data analysis, correspond to two extremes. The standard ILC focuses only on the total map noise, without trying to reduce any specific foreground bias, whereas deprojection completely nulls a given foreground (in so far as its frequency dependence is known exactly), regardless of the noise penalty. As we show below, our CMB lensing-optimized ILC is a compromise between these two extremes.

Standard ILC

The standard harmonic ILC is defined to minimize the total map variance from noise and foregrounds. Minimizing $C_\ell^{\rm tot}$ subject to the constraint $\boldsymbol{w}_\ell^T\mathbf{1}=1$ can be achieved by introducing a Lagrange multiplier λ_ℓ and minimizing the Lagrangian:

$$\mathcal{L}\left[w_\ell,\lambda_\ell\right]\equiv C_\ell^{\rm tot}+\lambda_\ell\left(1-\boldsymbol{w}_\ell^T\mathbf{1}\right). \tag{4.6}$$

Setting $\partial\mathcal{L}/\partial\boldsymbol{w}_\ell=0$ forces $\boldsymbol{w}_\ell\propto\boldsymbol{C}_\ell^{-1}\mathbf{1}$. The normalization is fixed by the constraint $\boldsymbol{w}_\ell^T\mathbf{1}=1$, from which we recover the standard result [216, 581]:

$$\boldsymbol{w}_\ell^{\rm ILC}=\frac{\boldsymbol{C}_\ell^{-1}\mathbf{1}}{\mathbf{1}^T\boldsymbol{C}_\ell^{-1}\mathbf{1}}. \tag{4.7}$$

While the ILC minimizes the total power spectrum of the combined map, there is no guarantee that the power spectrum of individual foreround components will be reduced, so long that the foregrounds are a subdominant source of power. Even if the foreground power spectra are reduced, the ILC it is not designed to reduce the bispectrum or trispectrum of the foregrounds, which in turn source the biases to lensing. Thus it is possible for a foreground to have a subdominant power spectrum, but a dominant bispectrum or trispectrum, in which case the ILC does nothing to suppress the resulting bias.

Deprojection

In certain applications, one may need to null a specific foreground. For instance, tSZ measurements (y-maps and stacks) are contaminated by thermal dust emission, if unaccounted for. In this context, one may use constrained ILC maps with CIB deprojection [497, 376]. As shown in Fig. 4.1, one can deproject one or several foregrounds to remove the corresponding biases, at the cost of increasing the temperature map noise. Deprojection amounts to choosing the weights $\boldsymbol{w}_\ell$ to have zero response to foreground s: $\boldsymbol{w}_\ell^T\boldsymbol{A}_s=0$. Minimizing the map noise subject to this additional constraint leads to minimizing

$$\mathcal{L}\left[w_\ell,\lambda_\ell,\sigma_\ell\right]\equiv C_\ell^{\rm tot}+\lambda_\ell\left(1-\boldsymbol{w}_\ell^T\mathbf{1}\right)-\sigma_\ell\boldsymbol{w}_\ell^T\boldsymbol{A}_s. \tag{4.8}$$

Setting $\partial\mathcal{L}/\partial\boldsymbol{w}_\ell=0$ forces $\boldsymbol{w}_\ell\propto\boldsymbol{C}_\ell^{-1}\left(\mathbf{1}+\sigma_\ell\boldsymbol{A}_s/\lambda_\ell\right)$. Enforcing the constraint $\boldsymbol{w}_{\boldsymbol{\ell}}^T\boldsymbol{A}_s=0$ fixes the coefficient:

$$\sigma_\ell/\lambda_\ell=-\frac{\mathbf{1}^T\boldsymbol{C}_\ell^{-1}\boldsymbol{A}_s}{\boldsymbol{A}_s^T\boldsymbol{C}_\ell^{-1}\boldsymbol{A}_s}. \tag{4.9}$$

Finally, $\boldsymbol{w}_\ell^T \mathbf{1} = 1$ fixes the normalization. This technique can be trivially extended to simultaneously deproject more than one foreground ($\boldsymbol{w}_\ell^T \boldsymbol{A}_{s_i} = 0$ for $i = 1, \cdots, n$) by introducing additional Lagrange multipliers.

4.4 Reconstruction noise

A general quadratic estimator $\hat{\kappa}_{\boldsymbol{L}}$ of the lensing convergence $\kappa_{\boldsymbol{L}}$ takes the form

$$\hat{\kappa}_{\boldsymbol{L}}[T, T'] = N_{\boldsymbol{L}} \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} T_{\boldsymbol{\ell}} T'_{\boldsymbol{L}-\boldsymbol{\ell}} \tag{4.10}$$

where $N_{\boldsymbol{L}}$ is some normalization, chosen so that the estimator has unit response to the signal $\langle \hat{\kappa}_{\boldsymbol{L}} \rangle = \kappa_{\boldsymbol{L}} +$ bias, and $F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ are the weights, which are in principle arbitrary. In this chapter, both maps T and T' have unit response to the CMB, such that the normalization is

$$N_{\boldsymbol{L}}^{-1} = \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}. \tag{4.11}$$

This estimator applied to the map $\hat{T}_{\boldsymbol{\ell}}$ has noise

$$\begin{aligned} \mathcal{N}_L[\boldsymbol{w}] &= \langle \hat{\kappa}_{\boldsymbol{L}}[\hat{T}, \hat{T}] \hat{\kappa}_{-\boldsymbol{L}}[\hat{T}, \hat{T}] \rangle \big/ (2\pi)^2 \delta^D_{\mathbf{0}} \\ &= 2N_{\boldsymbol{L}}^2 \int \frac{d^2\boldsymbol{\ell}}{(2\pi)^2} F^2_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|} \end{aligned} \tag{4.12}$$

for $L > 0$. In going from the first to second line of Eq. (4.12), we have only kept the disconnected contributions to the auto-correlation, as is customary in lensing noise $N^{(0)}$ calculations. We also assumed that F is symmetric and even in both of its arguments.

Note that both the normalization $N_{\boldsymbol{L}}$ and the weights $F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ could in principle depend on the weights $\boldsymbol{w}_{\boldsymbol{\ell}}$. In this analysis, the spatial weights $F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ depend on the map noise $C^{\rm tot}_\ell$, which in turn depend on the frequency weights $\boldsymbol{w}_\ell$. In App. C.2, we show that the ILC is the frequency combination which minimizes the lensing noise, for all the estimators we consider here. The top panel of Fig. 4.2 shows the lensing noise for the standard QE for various frequency combinations, with $\ell_{\mathrm{max},T} = 3000$.

Lensing amplitude noise

The lensing amplitude $A_{\rm lens}$ is defined to be the ratio of the measured convergence auto-spectrum to the truth, with fiducial value 1. An estimator for the lensing amplitude can be easily constructed for each mode L: $\hat{A}_{{\rm lens},L} = \hat{C}^\kappa_L / C^\kappa_L$, where $\hat{C}^\kappa_L$ denotes the measured convergence auto-spectrum. Given a range of modes, the minimum variance estimator takes the form

$$\hat{A}_{\rm lens} = \left(\int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{\hat{C}^\kappa_L}{C^\kappa_L} \frac{(C^\kappa_L)^2}{\sigma^2_L} \right) \Big/ \left(\int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{(C^\kappa_L)^2}{\sigma^2_L} \right), \tag{4.13}$$

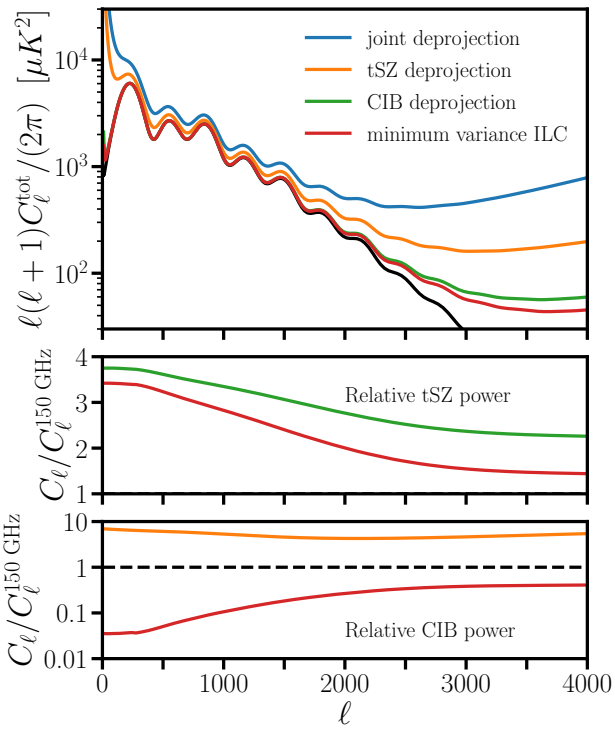


Figure 4.1: The top panel shows the temperature power spectrum for different multifrequency combinations, compared to the noiseless lensed CMB (black curve). Compared to the minimum-variance combination (ILC, red), the deprojection of tSZ (orange) comes at a larger cost in noise than that of CIB (green). Jointly deprojecting tSZ and CIB (blue) causes a more than ten-fold noise increase at $\ell \sim 3000$, where most of the lensing information resides. We note that the significant increase in power at large scales from joint or tSZ deprojection is primarily due to a large increase in atmospheric noise. The middle and bottom panels show the tSZ and CIB power spectra, respectively, compared to the single-frequency case (150 GHz). For tSZ (middle), the standard ILC and the CIB deprojection cause a significant increase. For CIB (bottom), the ILC very effectively reduces the power; however, the tSZ-deprojection enhances it by more than an order of magnitude.

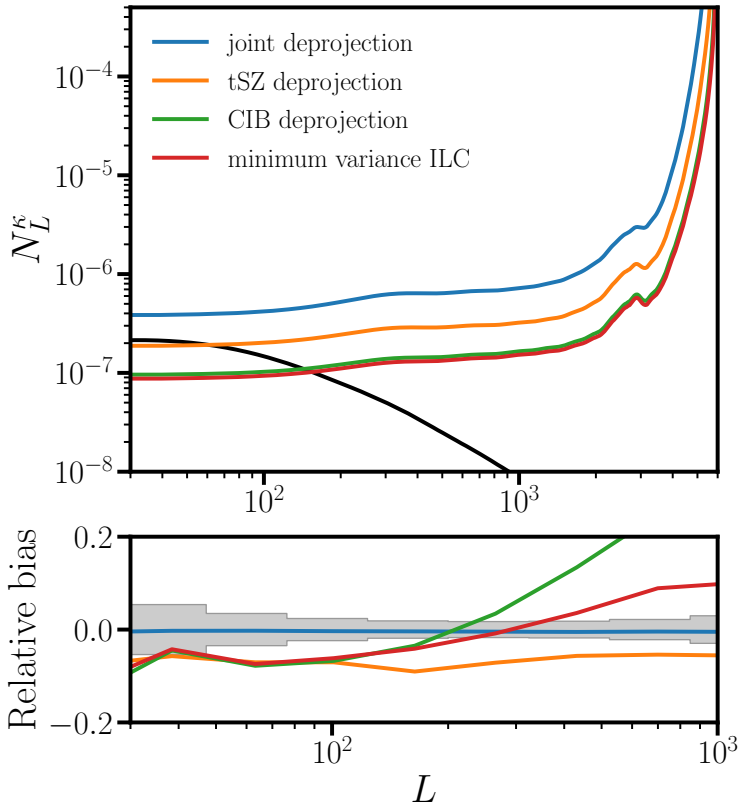


Figure 4.2: As illustrated in the top panel, the standard ILC is the linear combination of multifrequency maps which minimizes the lensing noise, for the standard QE and all the other estimators considered in this chapter. As expected from Fig. 4.1, deprojecting CIB (green) comes at almost no cost in lensing noise, whereas deprojecting tSZ (yellow) multiplies the noise by ~ 2. Despite causing a more than ten-fold increase in temperature noise, joint deprojection only causes a four-fold increase in lensing noise. This can be understood by counting the number of signal-dominated temperature modes. The bottom panel shows the fractional bias to the CMB lensing auto-spectrum for the various frequency combinations. The gray boxes denote the bandpower errors for the standard ILC. For the ILC, the bias is largely dominated by the enhanced tSZ signal. While the joint deprojection successfully nulls all biases, simply deprojecting tSZ or CIB does not.

where $\sigma_L^2 = 2(C_L^\kappa + N_L^\kappa)^2$. This is simply an inverse variance weighting of the estimators for each Fourier mode.

Here we simply aim to minimize the lensing noise:

$$\sigma^2_{A_{\rm lens}} = \frac{1}{4\pi f_{\rm sky}} \left(\int \frac{d^2\boldsymbol{L}}{(2\pi)^2} \frac{(C_L^\kappa)^2}{\sigma_L^2} \right)^{-1} \tag{4.14}$$

Minimizing $\sigma_{A_{\rm lens}}$ is equivalent to maximizing the integral in Eq. (4.14), which has a strictly

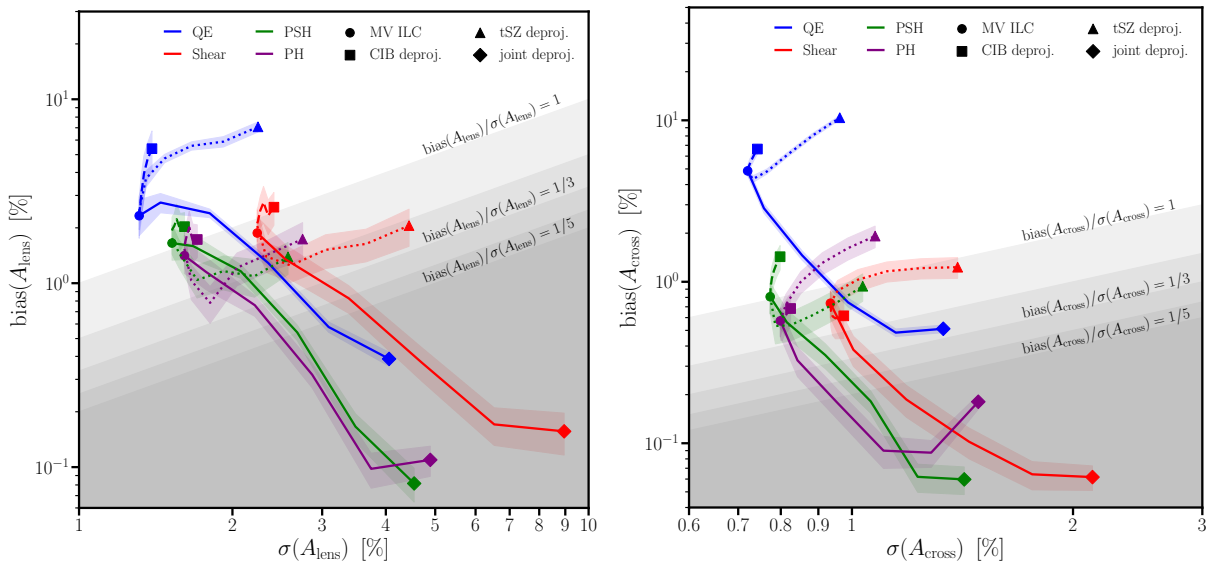


Figure 4.3: *Left:* The total |bias| and noise on the lensing amplitude $A_{\rm lens}$ for the various ILC combinations. The biases for the standard QE are shown in blue, Shear in red, the Point Source Hardened estimator in green, and the Profile Hardened estimator in purple. *Right:* The same as the left panel, but for the amplitude $A_{\rm cross}$ of the cross-correlation with the mock LSST sample. Circles denote the minimum variance (MV) ILC, squares (triangles) denote CIB (tSZ) deprojection, while diamonds denote the simultaneous deprojection of both the CIB and tSZ. The lines connecting MV ILC to either CIB, tSZ or joint deprojection follow Eq. (4.19), with step sizes $\Delta t = 0.2$. Throughout we sum over $20 < L < 1000$ when calculating the bias and noise for either $A_{\rm lens}$ or $A_{\rm cross}$, and take $f_{\rm sky} = 0.4$.

positive integrand. Thus maximizing the integral is equivalent to maximizing the integrand for each $\boldsymbol{L}$, which in turn is equivalent to minimizing σ_L^2, or N_L^κ. Therefore the weights that minimize the noise of the lensing amplitude $\sigma_{A_{\rm lens}}$ are equivalent to the weights that minimize the lensing reconstruction noise N_L^κ, which is given by the standard ILC.

4.5 Bias to lensing reconstruction from simulations

Method

To estimate the non-Gaussian lensing biases from extragalactic foregrounds, we follow the procedure in [519, 524]. We ignore the lensed-foreground bias, due to the fact that the foregrounds are emitted at cosmological distances and are themselves lensed [410]. The non-Gaussian foreground bias to the lensing auto-spectrum can be decomposed into a primary,

secondary and trispectrum terms [524, 446, 227, 597]:

$$\begin{aligned}\langle \mathcal{Q}\mathcal{Q}\rangle = &\underbrace{\langle \kappa_{\rm CMB}\kappa_{\rm CMB}\rangle}_{\text{lensing signal}} + \underbrace{2\langle \mathcal{Q}[T^{\rm CMB},T^{\rm CMB}]\mathcal{Q}[s,s]\rangle}_{\text{Primary bispectrum bias}} \\ &+ \underbrace{4\langle \mathcal{Q}[T^{\rm CMB},s]\mathcal{Q}[T^{\rm CMB},s]\rangle}_{\text{Secondary bispectrum bias}} + \underbrace{\langle \mathcal{Q}[s,s]\mathcal{Q}[s,s]\rangle}_{\text{Trispectrum bias}},\end{aligned} \tag{4.15}$$

where $\mathcal{Q}[T,T']$ represents any (symmetric) quadratic estimator evaluated on the maps T and T'. For the cross-correlation of CMB lensing with an external tracer g of the mass (e.g., galaxies or galaxy lensing), only the primary bispectrum bias appears:

$$\langle g\mathcal{Q}\rangle = \underbrace{\langle g\kappa_{\rm CMB}\rangle}_{\text{lensing signal}} + \underbrace{\langle g\mathcal{Q}[s,s]\rangle}_{\text{bispectrum bias}} . \tag{4.16}$$

In both cases, the bispectrum biases (primary and secondary) appear because the foregrounds are simultaneously non-Gaussian and correlated with the true CMB lensing signal. The trispectrum bias in auto-correlation is present for any non-Gaussian foreground, regardless of any potential correlation with CMB lensing.

In what follows, we estimate each of these bias terms separately for the total contribution from extragalactic foregrounds (tSZ + CIB + kSZ + radio PS). We apply the method presented in [524] identically. This uses the non-Gaussian foreground simulations of [541] at 150 GHz, which are then rescaled to any other frequency, given the expected SED of each component from [200], as implemented in `MicroCoSM`[2]. We choose the tracer g to be a LSST-like sample constructed from reweighting the simulated halos, as in [524, 519], to obtain the correct redshift distribution and linear bias.

In all cases, the foreground maps from [541] are masked before the lensing reconstruction. The mask is obtained by detecting the individual point sources with flux density higher than 5 mJy in the 150 GHz map using a matched filter, and masking the foreground map with a 3 arcmin radius disks around each source. We then add the unmasked lensed CMB map to the foreground map. This simulates the effect of inpainting, or source template subtraction, as expected for SO. We also explore a more aggressive masking strategy in App. C.4, which could reduce the bias to the profile hardened estimator by as much as 30%, provided that the effects of the mean field produced by this more aggressive mask are kept under control. Masking the point sources may also mask the high κ peaks, which would bias the signal low [220, 358], but inpainting or template subtraction should circumvent this problem.

Another important method for avoiding foreground contamination is to discard the highest multipoles in the input temperature map. Unless otherwise indicated, we assume that only modes with $\ell \leq \ell_{\rm max,}T = 3000$ are included. In App. C.5, we explore the sensitivity of our results to this choice.

[2] https://github.com/EmmanuelSchaan/MicroCoSM

Lensing amplitude bias

The bias to the lensing amplitude $A_{\rm lens}$ is given by

$$\text{bias}\,(A_{\rm lens}) = \frac{\int d^2\boldsymbol{L}\, C_L^\kappa\, \text{bias}\,(C_L^\kappa)\,/\sigma_L^2}{\int d^2\boldsymbol{L}\,(C_L^\kappa)^2/\sigma_L^2}, \tag{4.17}$$

where the bias to the lensing auto-correlation is the sum of the primary, secondary and trispectrum terms, as defined in Eq. (4.15). Analogous expressions for the noise and bias of the amplitude $A_{\rm cross}$ of the cross-correlation of CMB lensing with an external tracer g may be obtained by substituting C_L^κ with $C_L^{\kappa g}$ and setting $\sigma_L^2 = (C_L^\kappa + N_\ell^\kappa)(C_L^g + N_L^g) + (C_L^{\kappa g})^2$ in Eqs. (4.14) and (4.17).

We note that cancellations could potentially occur in Eq. (4.17) between the biases at different lensing multipoles. In App. C.1 we explore how our optimal ILC weights change if we replace $\text{bias}(C_L^\kappa)$ with $|\text{bias}(C_L^\kappa)|$ in the expression above, which removes these cancellations.

4.6 Minimizing lensing bias and noise

In this section we piece together the preceding results and show that an optimal multifrequency linear combination for minimizing a combination of noise and bias (for both $A_{\rm lens}$ and $A_{\rm cross}$) lies between the standard ILC and joint deprojection.

Method

Ideally, one would solve for the scale and frequency-dependent weights $\boldsymbol{w}_\ell$ which minimize some function $\mathcal{L}$ of the variance and bias of $A_{\rm lens}$ (or $A_{\rm cross}$). A reasonable choice for $\mathcal{L}$ is:

$$\mathcal{L}[\boldsymbol{w}_\ell] = \sigma^2(A_{\rm lens}) + f_b^2 \text{bias}^2(A_{\rm lens}), \tag{4.18}$$

where f_b is a tuning parameter for favoring either noise ($f_b = 0$) or bias ($f_b \to \infty$). In practice, however, solving for $\boldsymbol{w}_\ell$ via a brute force minimization of $\mathcal{L}$ is computationally intractable. The problem can be greatly simplified with a few well-motivated approximations; however, these turn out not to be accurate enough for our purposes (see App. C.3). We circumvent this issue with the following ansantz:

$$\boldsymbol{w}_\ell(t) = t\boldsymbol{X}_\ell + (1-t)\boldsymbol{w}_\ell^{\rm ILC}, \tag{4.19}$$

which smoothly connects the standard ILC to some form of deprojection ($\boldsymbol{X}_\ell = \boldsymbol{w}_\ell^{\text{CIB deproj.}}$, $\boldsymbol{w}_\ell^{\text{tSZ deproj.}}$, or $\boldsymbol{w}_\ell^{\text{joint deproj.}}$), and reduces the dimensionality of the parameter space from $5 \times$ (# of ℓ-bins) to 1. In what follows we compute the bias and noise for each estimator along the line segment $t \in [0, 1]$, in steps of $\Delta t = 0.2$. The optimal value of t could in principle be found through minimizing Eq. (4.18) for fixed f_b. However, the value of f_b is arbitrarily chosen to reduce the bias to some desirable level, which can just as well be

accomplished by choosing t to be the smallest value such that bias$(A_{\rm lens})/\sigma(A_{\rm lens})$ is below some desirable cutoff. This is similar in spirit to the approach taken in deriving the partially constrained ILC [13]. However, our cutoff is implemented directly at the bias level, while the cutoff in [13] is implemented at the foreground power level.

Results

Our results are summarized in Fig. 4.3, where we plot the bias vs noise for each estimator as we vary the weights $\boldsymbol{w}_\ell$ between the standard ILC and some form of deprojection $\boldsymbol{X}_\ell$, following the prescription of Eq. (4.19). We first note that deprojecting a single foreground (tSZ or CIB) often boosts the overall bias. As shown in bottom two panels Fig. 4.1, deprojecting tSZ significantly boosts the CIB power spectrum (and vise versa), increasing the overall bias for both the auto and cross-correlation in many cases.

We find that hardening against a tSZ-like profile outperforms the remaining estimators for any bias/noise cutoff[3], both in auto- and cross-correlation. Profile hardening at ILC ($t = 0$, purple circles in Fig. 4.3) reduces the bias to the lensing amplitude by 40%, at a 20% cost in noise, while the bias to the cross-correlation is reduced by nearly an order of magnitude at a 10% noise cost, relative to the standard quadratic estimator at ILC (blue circles in Fig. 4.3).

For all estimators, comfortably reducing the bias below $\sigma(A_{\rm lens})/2$ (or $\sigma(A_{\rm cross})/2$) requires partially deprojecting both the tSZ and CIB. We find that the noise cost for partial joint deprojection is much more significant for the auto-correlation than for the cross-correlation. In particular, for the profile hardened estimator, reducing the bias below $\sigma/2$ via partial joint deprojection corresponds to $t \sim 0.2$, resulting in a 20% and 5% noise penalty for the auto- and cross-correlation respectively, relative to profile hardening at ILC[4]. Further improvements can be made by pushing to a higher $\ell_{{\rm max},T}$ and moving closer to deprojection (higher t), or through a more aggressive masking scheme, as discussed in Apps. C.5 and C.4 respectively.

4.7 Conclusions

CMB lensing is quickly becoming one of the most powerful cosmological tools. Percent level measurements from the next generation of experiments require excellent control of extragalactic foregrounds, especially from temperature anisotropies. Traditional analyses have used multifrequency foreground reduction methods when possible. Here we combined geometric methods with multifrequency information to optimize the lensing reconstruction. In

[3]This statement is no longer true near joint deprojection ($t \simeq 1$). In this regime, the PSH estimator has a lower bias (and noise) than the PH estimator, since in this case the foreground biases are primarily sourced by radio point sources (as opposed to tSZ or CIB), which the former estimator is designed to null. In practice, however, one does not need to push to $t \simeq 1$ to reduce the bias to an acceptable level.

[4]Relative to the standard QE at ILC, the noise penalty for reducing the bias below $\sigma/2$ using profile hardening and partial joint deprojection is 50% and 15% for the auto- and cross-correlation respectively.

doing so, we maximize the amount of information that can be extracted and the robustness of the measurement, while minimizing the potential biases. We also note that any multi-frequency linear combination alone will not reduce the bias from kSZ, since it preserves the black-body nature of the CMB. While kSZ bias can be significant [227], it is quite effectively reduced by the shear-only or bias hardened estimators [524, 519]. Moreover, direct measurement of kSZ by combining low-redshift galaxy catalogs together with high-resolution CMB maps [525] can potentially provide templates for subtraction at the map level. We find that the standard ILC can lead to significant biases, which can be partially mitigated with geometric methods such as profile hardening, which alone can reduce the bias by 40% (in auto-correlation) or almost an order of magnitude (in cross-correlation) with minimal $(10-20\%)$ increase in noise. The bias can be further reduced to any desirable bias/noise cutoff via partial joint deprojection of the tSZ and CIB, with the auto-correlation suffering a higher relative noise penalty (factor of ~ 4) than the cross-correlation. We find that deprojection of tSZ or CIB alone do not reduce the biases, due to an enhancement in the component that is not deprojected. Finally, we have explored possible improvements with more aggressive masking (App. C.4) and varying $\ell_{\mathrm{max},T}$ (App. C.5).

While here we have focused on temperature reconstruction, similar bias hardening techniques should be applicable to extragalactic foregrounds in polarization as well, where the biases are expected to be less severe. We leave a full exploration of this to future work.

Acknowledgments

We thank Anthony Challinor, Jo Dunkley, Giulio Fabbian, Colin Hill, Akito Kusaka, Antony Lewis, Mathew Madhavacheril, Neelima Sehgal, Alexander Van Engelen and Martin White for useful discussions in the preparation of this chapter. N.S. is supported by the NSF. E.S. is supported by the Chamberlain fellowship at Lawrence Berkeley National Laboratory. S.F. is supported by the Physics Division of Lawrence Berkeley National Laboratory. This chapter used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Part II

Tenuous tensions

Chapter 5

Tomographic structure growth I: cross-correlation of DESI Luminous Red Galaxies with CMB lensing

This chapter was originally published as

> Noah Sailer et al. "Cosmological constraints from the cross-correlation of DESI Luminous Red Galaxies with CMB lensing from Planck PR4 and ACT DR6". In: *JCAP* 06 (2025), p. 008. DOI: 10.1088/1475-7516/2025/06/008. arXiv: 2407.04607 [astro-ph.CO]

We infer the growth of large scale structure over the redshift range $0.4 \lesssim z \lesssim 1$ from the cross-correlation of spectroscopically calibrated Luminous Red Galaxies (LRGs) selected from the Dark Energy Spectroscopic Instrument (DESI) legacy imaging survey with CMB lensing maps reconstructed from the latest *Planck* and ACT data. We adopt a hybrid effective field theory (HEFT) model that robustly regulates the cosmological information obtainable from smaller scales, such that our cosmological constraints are reliably derived from the (predominantly) linear regime. We perform an extensive set of bandpower- and parameter-level systematics checks to ensure the robustness of our results and to characterize the uniformity of the LRG sample. We demonstrate that our results are stable to a wide range of modeling assumptions, finding excellent agreement with a linear theory analysis performed on a restricted range of scales. From a tomographic analysis of the four LRG photometric redshift bins we find that the rate of structure growth is consistent with ΛCDM with an overall amplitude that is $\simeq 5-7\%$ lower than predicted by primary CMB measurements with modest ($\sim 2\sigma$) statistical significance. From the combined analysis of all four bins and their cross-correlations with *Planck* we obtain $S_8 = 0.765 \pm 0.023$, which is less discrepant with primary CMB measurements than previous DESI LRG cross *Planck* CMB lensing results. From the cross-correlation with ACT we obtain $S_8 = 0.790^{+0.024}_{-0.027}$, while when jointly analyzing *Planck* and ACT we find $S_8 = 0.775^{+0.019}_{-0.022}$ from our data alone and $\sigma_8 = 0.772^{+0.020}_{-0.023}$

with the addition of BAO data. These constraints are consistent with the latest *Planck* primary CMB analyses at the $\simeq 1.6-2.2\sigma$ level, and are in excellent agreement with galaxy lensing surveys.

5.1 Introduction

The evolution of large scale structure (LSS) fluctuations offers a unique window into fundamental physics, the formation of galaxies and their associated clusters [460, 194]. ΛCDM accurately predicts the evolution of matter perturbations from their primordial seeds to the present day on sufficiently large scales. Consequentially, large scale late time structure growth measurements can be used as a powerful consistency check of ΛCDM conditioned on primary CMB observations. More generally, structure growth is sensitive to extensions of the standard cosmological model including but not limited to dark matter interactions (e.g. [110, 111, 44, 76, 43, 333]), deviations from general relativity (e.g. [621]) and modifications to the expansion history tracing back to deep within the radiation dominated epoch (e.g. [512, 251]).

There is a wide range of observational handles on the amplitude of low redshift density fluctuations, conventionally parameterized by σ_8 (or $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$) which by definition sets the amplitude of *linear* matter density fluctuations at the present day. Analyses of cluster counts [20, 5, 521, 247, 100], peculiar velocity surveys [297, 522, 565], Sunyaev-Zel'dovich (SZ) effects [573, 293], redshift space distortions (RSD) [309, 162, 152, 594, 135, 308, 639, 326, 205, 294] and gravitational lensing measurements [285, 286, 363, 6, 39, 538, 421, 411, 571, 164, 349, 7, 25, 119, 493] cover a broad range of scales with different systematic uncertainties and assumptions required to infer σ_8 (or S_8) from each observable. Several recent analyses of low redshift tracers report S_8 constraints that are slightly lower than predicted from ΛCDM conditioned on primary CMB data from *Planck*, albeit at modest statistical significance. Examples include an analysis of cluster counts from DES Y1 observations [5]; the peculiar velocity field derived from the Democratic Samples of Supernovae [565]; reanalyses of BOSS full shape and post reconstruction data [309, 162, 152, 135, 308] using Effective Field Theory (EFT) based models; galaxy shear and its cross-correlations with galaxy positions from KiDS [286, 363], DES Y3 [6, 39, 538], and HSC [411, 571, 164, 349] data; correlations between *Planck* CMB lensing and DESI Legacy Survey galaxies [323, 278, 625], *Planck* lensing and unWISE galaxies [330], *Planck*+SPT lensing and DES galaxies (including shear) [128, 8], *Planck*+ACT lensing and KiDS shear [502], ACT lensing and DES Y3 galaxies [385]; and combinations of the above [244, 140, 132], all of which find S_8 constraints that are $\sim 5-10\%$ lower than *Planck* at modest $\simeq 2-3\sigma$ significance (with the exception of DES Y1 cluster counts, which claim a $\sim 5\sigma$ tension).

The collection of these low S_8 constraints (among others, see e.g. [10] for a more comprehensive review) has become known as the S_8 "tension." However, not all low redshift probes prefer low S_8 values. For example, an analysis of tSZ clusters identified from *Planck* and ACT data [293], SPT clusters with DES and HST weak lensing [100], ACT CMB lensing

correlated with BOSS galaxies [165], *Planck* CMB lensing correlated with DES Y1 galaxy positions and shear [630], and *Planck* CMB lensing correlated with Quaia quasars [469] are all consistent with *Planck* to well within 1σ, while the latest eROSITA cluster analysis [247] finds $S_8 = 0.86 \pm 0.01$. Moreover, recent reanalyses of datasets claiming low S_8 values have reported less discrepant results than found previously. Examples include a reanalysis of DES Y1 clusters [521] that properly forward models cluster selection effects, a joint analysis of DES and KiDS data [7] (which finds that their S_8 constraints vary at the $\simeq 1\sigma$ level depending on the linear alignment model assumed), and a reanalysis of CMB lensing (*Planck* PR4 + ACT DR6) correlated with unWISE galaxies [223]. In addition, it has been noted [639] that recent EFT-based analyses of BOSS data, which favor lower S_8 values than more traditional RSD analyses [205, 294] or halo model based methods [326], are susceptible to "volume effects" that complicate the interpretation of marginal posteriors (exaggerating the "standard" parameter based tension metric).

Notably, analyses of the CMB lensing power spectrum $C_\ell^{\kappa\kappa}$ from the latest *Planck* [119] and ACT [377] data are also in excellent agreement with the *Planck* primary CMB. CMB lensing directly probes the (Weyl) potential on predominantly linear scales and at late times ($z \simeq 0.5 - 5$) with a well-characterized source distribution, making it straightforward to model. Additionally, the CMB lensing convergence is measured by utilizing very well understood statistical properties of the primary CMB, such that its calibration is known analytically as a function of cosmological parameters. Together these properties make CMB lensing arguably the most pristine probe of low redshift matter fluctuations. However, given that CMB lensing is a projected probe over a wide range of redshifts, it is difficult to make precise statements about specific late-time epochs (e.g. $0.4 < z < 1$) from CMB lensing alone. This in contrast to the tomographic CMB lensing analyses mentioned above (with mixed S_8 results) and explored in this chapter, which isolate the CMB lensing contribution from a desired redshift range via cross-correlations with a second LSS tracer.

In particular, a previous cross-correlation analysis of DESI Luminous Red Galaxies (LRGs) with *Planck* PR3 CMB lensing reported $S_8 = 0.725 \pm 0.030$ [625], roughly 3σ lower than favored by the primary (*Planck* 2018) CMB. Motivated by the availability of lower noise CMB lensing maps from *Planck* [119] and ACT [493], several improvements to the LRG sample [643], a more accurate treatment of mode couplings arising from e.g. masks in CMB lensing estimators, and the development of a Hybrid EFT emulator [181] we perform an improved analysis of the cross-correlation of DESI LRGs with CMB lensing. The LRG sample has subpercent stellar contamination, is highly robust to the systematic weight treatment, and has a spectroscopically calibrated redshift distribution, while our fiducial HEFT model robustly marginalizes over the large-scale impact of small-scale astrophysical uncertainties, making our analysis tailored to mitigate systematic uncertainties. In Fig. 5.1 we show the contribution to the signal-to-noise ratio (SNR) for each of our measurements (see §5.3 and the companion paper [320]) per unit redshift and wavenumber for our fiducial analysis choices. Our data are primarily sensitive to redshifts $0.4 \lesssim z \lesssim 1$ and for all measurements 70% of the SNR comes from $k < 0.3\,h\,\mathrm{Mpc}^{-1}$. Given that we marginalize over higher-order astrophysical uncertainties, raw SNR (particularly at higher ℓ) does not neces-

sarily translate to a higher precision S_8 measurement, making our constraint predominately sensitive to linear scales ($k < 0.1\,h\,\mathrm{Mpc}^{-1}$).

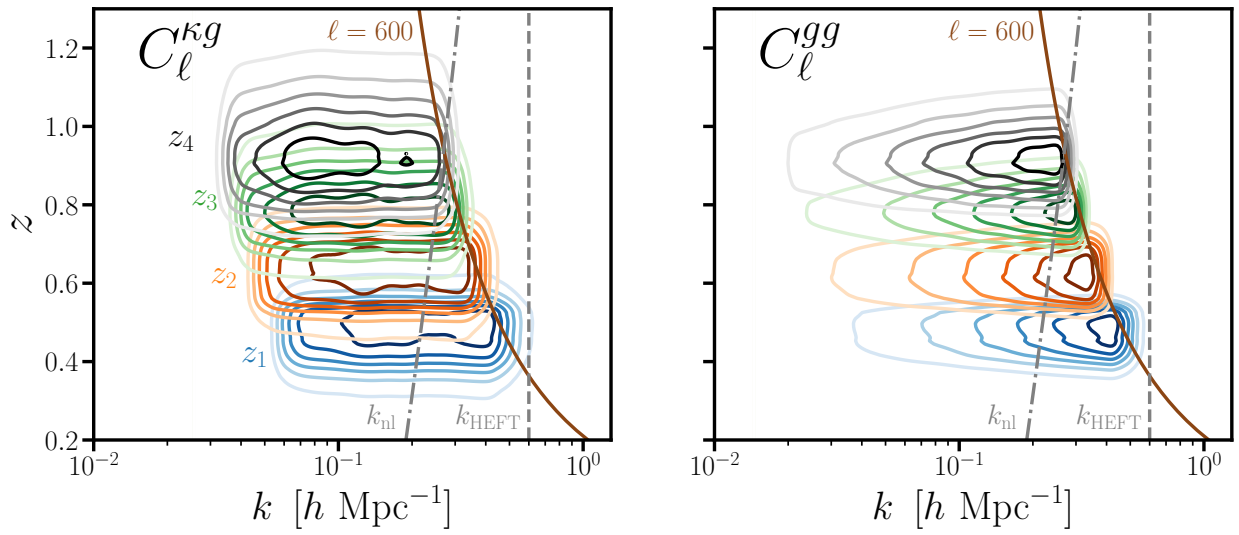


Figure 5.1: Contour plots of the peak normalized SNR(k, z) for the galaxy-CMB lensing cross-correlations ($C_\ell^{\kappa g}$, left) and the galaxy auto-correlations (C_ℓ^{gg}, right) when taking $\ell_{\rm max} = 600$. SNR(k, z) is defined as the contribution to the signal-to-noise ratio per $dz\, d\ln k$, such that the total (signal-to-noise ratio)^2 is $\mathrm{SNR}^2 = \int dz\, d\ln k\ \mathrm{SNR}^2(k, z)$. See Appendix D.2 for a more detailed discussion. The colors of the contours correspond to the redshift bin (see Fig. 5.2 for the redshift distributions and Fig. 5.3 for the measurements themselves) while the shading indicates the value of each contour (from lighter to darker these values are 0.15, 0.3, 0.45, 0.6 0.75, and 0.9). The brown solid line corresponds to $k = 600/\chi(z)$, while $k_{\rm nl} = [\int dk P_{\rm lin}(k)/6\pi^2]^{-1/2}$ (grey dot dashed line) and $k_{\rm HEFT} \simeq 0.6\ h\,\mathrm{Mpc}^{-1}$ (grey dashed line) is the maximum scale for which we expect HEFT to be valid to the percent level [327] for LRG-like galaxies.

In a companion paper [320] we present the fiducial power spectra measurements, the hybrid covariance used in our fiducial analysis, consistency tests exploring the robustness of the cross-correlation of DESI galaxies with the ACT DR6 lensing map and constraints on our best-constrained parameter combination $S_8^\times \equiv \sigma_8(\Omega_m/0.3)^{0.4}$ using the likelihood presented here. In this chapter we cross-validate the power spectra measurements with an independent pipeline, investigate the robustness of the DESI LRG sample and its cross-correlation with *Planck* PR4, present our theory model, likelihood implementation and validation, and our primary cosmological constraints on S_8 and σ_8, including the addition of BAO data.

The remainder of this chapter is organized as follows. In §5.2 we summarize the data used in our analysis. In §5.3 we outline the methodology for ancillary power spectra measurements and covariance estimation used in our *Planck* PR3 reanalysis and systematics tests. We discuss the modeling (including alternatives to HEFT) in §5.4 and present our likelihood

and associated tests in §5.5. In §5.6 we perform additional systematics tests for the galaxy auto-spectra and the cross-correlation with *Planck* PR4. Our main cosmological results are given in §5.7. We discuss our results in the context of previous constraints in §5.8 and conclude with §5.9.

5.2 Data

Our analysis utilizes a photometric sample of Luminous Red Galaxies (LRGs) from the DESI Legacy Imaging Survey DR9 [640, 191, 643] and CMB lensing convergence maps reconstructed from *Planck* [25, 119] and Atacama Cosmology Telescope (ACT) [493, 372, 377] data. DESI is a highly multiplexed spectroscopic survey that is capable of measuring 5000 objects at once [187, 552, 409] and is currently operating on the Mayall 4-meter telescope at Kitt Peak National Observatory [185]. DESI is currently conducting a five-year survey and will obtain spectra for approximately 40 million galaxies and quasars [186, 264, 527], enabling constraints on the nature of dark energy through its impact on the universe's expansion history [341].

Some key properties of the LRG sample are summarized in Table 5.1. The LRG footprint, which is illustrated in the left panel of Fig. 5.2, covers[1] 18200 deg^2 with a surface density of 500 deg^{-2} with approximately 16600 and 7900 deg^2 of overlap with *Planck* PR4 and ACT DR6 respectively. In §5.2 we describe the LRG sample, including the photometric selection, systematic weights, redshift distribution and magnification bias estimation. We briefly discuss the *Planck* and ACT CMB lensing maps in §5.2 and §5.2 respectively, and refer the reader to the companion paper [320] and refs. [25, 119, 493] for a more detailed discussion.

Luminous Red Galaxies

The LRG samples are a subset of the Legacy Survey (LS) DR9 [191], which is currently being used for targeting by the DESI spectroscopic redshift survey. The data are comprised of optical (g, r, z) imaging from the Beijing–Arizona Sky Survey (BASS [645]), Mayall z-band Legacy Survey (MzLS [191]), Dark Energy Camera Legacy Survey (DECaLS [191]) and the Dark Energy Survey (DES [4]), as well as four mid-infrared ($W1 - W4$) bands from the Wide-field Infrared Survey Explorer (WISE [629]).

Photometric selection: We use the "Main LRG sample" whose photometric selection is described in detail in [640]. Briefly, the selection is composed of three color cuts on extinction-corrected g, r, z and $W1$ magnitudes to mitigate stellar contamination, remove galaxies below $z \lesssim 0.4$ and produce a roughly constant number density out to $z \sim 0.8$, in

[1] These numbers differ from Table 1 of [643], which take into account the fractional coverage in each pixel, whereas we treat each pixel as observed (1) or unobserved (0). The sky coverage quoted in the text are the relevant numbers for the clustering analysis performed here. The fractional coverage has been accounted for in the systematics weights and is reflected e.g. in the (increased) shot noise of the galaxies.

addition to a cut in z-band fiber magnitude to produce a tail in the redshift distribution extending just beyond $z = 1$. Due to differences in the photometry of the different imaging surveys, the selection in the "Northern" (BASS+MzLS) and "Southern" (DECaLS+DES) regions differ slightly in their implementation (see Eqs. 1 & 2 of [640]) in an effort to homogenize the LRG sample across the full footprint.

Photometric redshifts: The LRG sample has been subdivided into four photometric redshift bins [643] that we label z_1 through z_4. We note that the photometric redshifts presented in [643] differ slightly from those released in [641] and used in previous analyses (e.g. [625]). A detailed description of the photo-z algorithm and its latest improvements is presented in Appendix B of [643]. Most notably, the spectroscopic training data have been updated to include redshifts from DESI's Survey Validation [188] and Early Data Release [16]. We also note that the definition of the photo-z bins in the Northern and Southern regions differ slightly in an effort to homogenize the sample (see Table 2 of [643]).

Image quality cuts: We use the LRG masks publicly available at this URL [643], which have been constructed to mitigate contamination from unwanted imaging artifacts, stars, and other undesired astrophysical contaminants (e.g. large galaxies, star clusters, and planetary nebulae). These include the BRIGHT, GALAXY and CLUSTER bit masks used in the LS DR9, in addition to a "veto" mask constructed from unWISE [404], WISE [629], Gaia/Tycho-2 [240, 241], and visually-inspected LS DR9 data to remove problematic regions not captured by the LS DR9 bit masks. For a more detailed description of the veto mask, see Section 2.4 and Appendix D of [640]. With this set of photometric selections and masks, only 0.3% of the LRGs have been classified as stars by DESI's current spectroscopic data. In addition the following image quality cuts have been imposed: pixels with $E(B-V) \geq 0.15$ (determined with the extinction map from ref. [529], hereafter SFD) are masked; pixels where the stellar density is larger than 2500 deg^{-2} (determined by the Gaia [429] stellar density map) are masked; and only pixels with at least two exposures in the g, r, and z bands are included.

Imaging weights: Ref. [643] removes observational modulations to the LRG density maps following a standard procedure that we now describe. Weights are constructed under the assumption that variations in the local mean LRG density due to instrumental and galactic effects can be modeled as a random downsampling of the "true" LRG population. Under this assumption these density trends can be removed via division of the galaxy density map by a suitable weight map reflecting this downsampling factor. In practice, these weights are taken to be a linear combination of a set of templates whose coefficients are chosen via linear regression to remove trends in LRG density with the templates. The fiducial weights provided in [643] use seven templates: depth and seeing in the three optical bands (g, r, z) and $E(B-V)$ [529]. Ref. [643] also provides a set of weights that do not use an $E(B-V)$ template, which we use in §5.6 to assess the importance of Galactic extinction and CIB contamination in the SFD map as a potential systematic. On the scales relevant for our analysis the systematic weights impact the galaxy auto-spectra by less than a percent (Fig. 5.10).

Map making: We use the LRG density contrast maps publicly available at this URL.

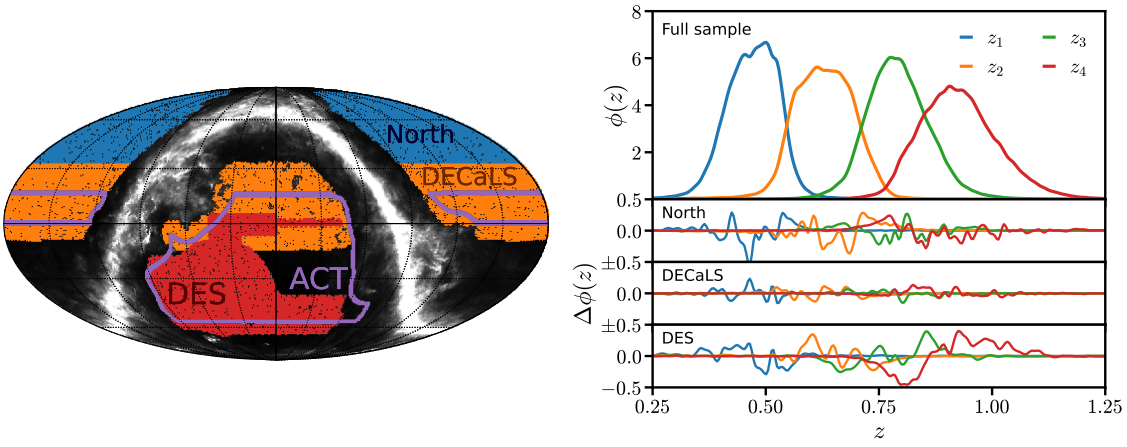


Figure 5.2: *Left*: The imaging footprints of our LRG sample (blue, orange, red) and the ACT lensing footprint (purple). The SFD $E(B-V)$ map [529] is shown in greyscale. *Right*: In the top panel we show the (spectroscopically-calibrated) normalized redshift distributions of the four LRG samples (labeled by color) averaged over the full photometric footprint from [643]. The bottom three panels show the variation of the redshift distribution across imaging regions: $\Delta\phi(z) = \phi^X(z) - \phi^{\rm full}(z)$ where X = North, DECaLS, or DES.

Here we briefly summarize the map making performed by ref. [643]. The LRGs were binned into HEALPix [252] pixels with `nSide = 2048`. Likewise a pixelized map of weighted randoms was constructed using the previously described imaging weights. The masks and image quality cuts (discussed above) were then applied to both the LRG and weighted random maps. In an effort to mitigate the presence of many small holes in the LRG mask, pixels whose (weighted) random density was less than 1/5 of the mean (weighted) random density over the footprint satisfying the image quality cuts were additionally masked. We follow [625, 643] and choose to not apodize the LRG masks. The LRG density contrast is defined as the LRGs divided by the weighted randoms, mean subtracted and normalized to the mean.

Redshift distribution: The redshift distributions of the four photometric LRG samples have been directly calibrated in [643] using 2.3 million LRG redshifts from DESI's Survey Validation [188] and first year (Y1) of observations. These redshift distributions account for the aforementioned imaging weights, masks and image quality cuts, as well as the spectroscopic systematics weights for DESI Y1 observations. In addition, each spectroscopic galaxy is weighted by its inverse success rate (as defined in section 4.4 of [640]) to mitigate the impact of redshift failures[2]. The normalized redshift distributions of the four LRG samples, averaged over the full imaging footprint (determined by area weighting the redshift distributions calibrated on the North, DECaLS and DES footprints), is shown in the top right panel

[2]The spectroscopic sample used to calibrate these redshift distributions has a redshift failure rate of 1.3% and catastrophic failure rate (defined as deviating by more than 1000 km/s from the true redshift) of 0.2%.

Sample	$z_{\rm eff}$	b_1^E	$\bar{z}$	Δz	s_μ	$\bar{n}_\theta$ [deg^{-2}]	$10^6\,{\rm SN}^{\rm 2D}$
z_1	0.470	1.8	0.470	0.063	0.972	81.9	4.07
z_2	0.625	2.0	0.628	0.074	1.044	148.1	2.25
z_3	0.785	2.2	0.791	0.078	0.974	162.4	2.05
z_4	0.914	2.4	0.924	0.096	0.988	148.3	2.25

Table 5.1: A summary of the LRG sample properties. The effective redshift $z_{\rm eff}$ is calculated with Eq. (5.11). We list the best-fit Eulerian linear bias $b_1^E = 1 + b_1^L$ (where b_1^L is the Lagrangian linear bias) values from our fiducial PR4+DR6 analysis. The mean redshift $\bar{z}$, standard deviation Δz, number count slope s_μ, surface density $\bar{n}_\theta$ and shot noise ${\rm SN}^{\rm 2D}$ are taken from [643].

of Fig. 5.2. In the bottom three panels we show the variations of the redshift distribution in the three imaging regions. We propagate these variations to ΔC_ℓ^{gg} in §5.6, and examine their impact on cosmological constraints in §5.7.

Number count slope: We make use of the number count slopes s_μ measured in [643] for the "combined" sample, which are listed in Table 5.1. The logarithmic change to the total number of LRGs for a small change in magnitude due to magnification is defined as $s_\mu = d\log_{10} N/dm$. These are measured via finite difference by shifting the (g, r, z, $W1$) LRG magnitudes by $\delta m = \pm 0.01$, computing the appropriate z-band fiber magnitude shift $\delta m_{\rm fiber}$ for each galaxy (see Appendix C of [643] for details regarding the $\delta m_{\rm fiber}$ calculation), reapplying the photometric selection, recomputing the photometric redshifts with the shifted magnitudes and galaxy sizes, and finally rebinning the LRGs into photometric redshift bins.

Changes to the LRG sample since White et al. [625]: These are listed in Appendix A of [643] and briefly summarized here. The LRG photo-z's have been improved using more training data, the redshift distributions are estimated with a larger number of spectra, the zero point offset for DEC $< -29.25\circ$ in the r and z bands has been corrected using the offset map from [642], and the imaging systematics no longer include a template for WISE $W1$ depth (due to a lack of an observed density trend with $W1$ depth and a discrepancy in the $W1$ depths in the data and randoms). With these changes, the LRG auto-correlation differs by at most 2% from the measurements used in [625], which is primarily driven by the improved photo-z's and hence narrower redshift distributions.

Planck CMB lensing

We consider both *Planck* 2018 [25] (PR3) and PR4 [119] CMB lensing measurements in this work. For both releases we low pass filter the CMB lensing convergence multipoles[3] ($\hat{\kappa}_{\ell m}$)

[3] As in [625] we multiply the reconstructed CMB lensing multipoles $\hat{\kappa}_{\ell m}$ by ${\rm Exp}[-(\ell/2500)^6]$. This filter impacts the cross-correlation measurements by less than 0.02% for $\ell < 600$, thus we neglect this filtering in our forward model. We do however include the filter in the fiducial CMB lensing noise curves used for our covariance.

and rotate them from Galactic to equatorial coordinates using `healpy`'s `rotate_alm` routine before constructing `nSide = 2048` HEALPix convergence maps via `alm2map`. Following [625] we apodize the (binary) CMB lensing masks with a 30 arcmin "C2" filter using the `NaMaster` [35] routine `mask_apodization` before rotating the masks from galactic to equatorial coordinates with the `healpy` routine `rotate_map_alms`.

Public Release 3 (PR3): We reanalyze the cross-correlation with PR3 to make direct contact with a previous analysis [625] (which found $S_8 = 0.725 \pm 0.030$) using a similar sample of LRGs. In §5.7 we quantify the impact of the improved LRG sample, new binning scheme, and updated modeling choices on these results. The *Planck* PR3 lensing measurement [25] reconstructs the lensing convergence from a linear combination of minimum variance quadratic estimators [445] whose inputs are Wiener-filtered and inverse-variance-weighted `SMICA` [27] temperature and polarization multipoles with $100 \leq \ell \leq 2048$. Specifically, we use the minimum variance lensing reconstruction[4] along with the provided CMB lensing mask and effective reconstruction noise curve $N_\ell^{\kappa\kappa}$.

Public Release 4 (PR4): The *Planck* PR4 lensing map [119] is reconstructed using a more optimal Global Minimum Variance estimator [382] from the latest NPIPE processing pipeline, which uses $\sim 8\%$ additional measurement time and a more optimal (anisotropic) filtering scheme (using the same scales $100 \leq \ell \leq 2048$ as PR3). The resulting CMB lensing map is signal dominated (per mode) out to[5] $\ell \sim 70$, compared with $\ell \sim 40$ with the previous PR3 map, resulting in an overall 20% improvement to the signal-to-noise ratio and a 42σ detection significance of the CMB lensing auto-correlation. Specifically, we use the reconstructed convergence multipoles, mask and reconstruction noise curve provided on NERSC[6].

ACT DR6 CMB lensing

Our headline analysis uses the latest CMB lensing measurement[7] (which detected the CMB lensing auto-correlation with a 43σ significance) from the Atacama Cosmology Telescope (ACT) DR6 [493, 372, 377], which we low pass filter[8] before constructing a κ map using `healpy`'s `alm2map` routine with `nSide = 2048`. The data used to construct this map consists of nighttime observations made through 2021 at 98 and 150 GHz. The CMB lensing convergence is reconstructed from a (nearly optimal) linear combination of temperature- and polarization-based quadratic estimators, which in turn have been separated into a linear

[4] COM_Lensing_4096_R3.00 from the *Planck* legacy archive.

[5] We use context rather than different symbols (e.g. with L instead of ℓ, as in the companion paper [320]) to differentiate CMB lensing and primary CMB multipoles, since we use L in §5.3 to denote bandpowers. From §5.3 onward ℓ always refers to a CMB lensing multipole.

[6] See github.com/carronj/planck_PR4_lensing for more information. Specifically, we use PR4_klm_dat_p.fits, mask.fits.gz and PR4_nlkk_p.dat.

[7] We use the baseline reconstructed multipoles ($\hat{\kappa}_{\ell m}$, kappa_alm_data_act_dr6_lensing_v1_baseline.fits) and corresponding mask (mask_act_dr6_lensing_v1_healpix_nside_4096_baseline.fits) that are available at this URL. We downgrade the mask from an nSide of 4096 to 2048 using healpy's ud_grade routine.

[8] We set $\hat{\kappa}_{\ell m} = 0$ for $\ell > 3000$, and leave modes with $\ell \leq 3000$ untouched.

combination of estimators measured with disjoint data splits [375] to mitigate noise biases. All estimators use the scales $600 < \ell < 3000$ to mitigate contamination from Galactic emission and extragalactic foregrounds. To further mitigate extragalactic contamination in the temperature-based reconstruction, a profile-hardened estimator [434, 446, 519] has been adopted, which has been shown to efficiently suppress tSZ- and CIB-induced biases without requiring a finely-tuned model for the mean tSZ intensity profile [517].

5.3 Power spectra, window functions and covariances

The bandpowers and accompanying covariance used in our fiducial *Planck* PR4 and ACT DR6 analyses are discussed in detail in the companion paper [320] and illustrated in Fig. 5.3. In this chapter we independently measure these power spectra and analytically estimate their covariance (following a slightly different approach from the companion paper [320]) as a cross-check. We compare cosmological constraints from the two approaches in Appendix D.3 and find negligible differences between them (in particular, the S_8 constraints are identical to three decimal places). In sections 5.5, 5.6 and 5.7 we perform systematics tests, estimate volume effects, fit to mocks, reanalyze the *Planck* PR3 cross-correlation and perform a suite of parameter-based consistency checks − all of which require additional measurements and covariance estimation that were done using the methods summarized below.

Masking and Fourier transforms are both linear operations. It follows that the measured power spectrum[9] of a masked LSS tracer ($\tilde{C}_\ell$) is linearly related to the underlying power spectrum (C_ℓ). Typically one bins the measured spectrum into bandpowers (which we label $\tilde{C}_L$, where capital L stands for bandpower), such that $\langle \tilde{C}_L \rangle = \sum_\ell W_{L\ell} C_\ell$ where the "window function" $W_{L\ell}$ is a functional of the mask(s), definition of the bandpowers and any other (optional) linear operations that the user applies. We measure "mask-deconvolved" bandpowers[10] following the MASTER algorithm [290] as implemented in `NaMaster` [35]. Specifically we use the `compute_full_master` routine with uniform weights and 33 non-overlapping bin edges that are linearly spaced in $\sqrt{\ell}$ across the fiducial ($20 \leq \ell < 600$) analysis range:

$$[10, 20, 44, 79, 124, 178, 243, 317, 401, 495, 600, 713, 837, 971, 1132, 1305, 1491, 1689,$$
$$1899, 2122, 2357, 2605, 2865, 3137, 3422, 3719, 4028, 4350, 4684, 5030, 5389, 5760, 6143].$$

We follow the default `NaMaster` convention that each bandpower is inclusive (exclusive) for the minimum (maximum) bin edge, i.e. for two bin edges $L^{\rm edge}_{\rm min}$ and $L^{\rm edge}_{\rm max}$ the bandpower is given the average over all C_ℓ's satisfying $L^{\rm edge}_{\rm min} \leq \ell < L^{\rm edge}_{\rm max}$, such that our bin edges define 32 *disjoint* bandpowers. Despite the fact that we only use modes with $20 \leq \ell < 600$ in our analysis we compute bandpowers down to $L_{\rm edge} = 10$ and up to $L_{\rm edge} =$ `3*nSide-1`

[9] We use tildes to differentiate a measured from an underlying power spectrum.

[10] Defined such that $W_{L\ell}$ is a tophat in ℓ for each L when the underlying power spectrum is piece-wise constant in each bandpower.

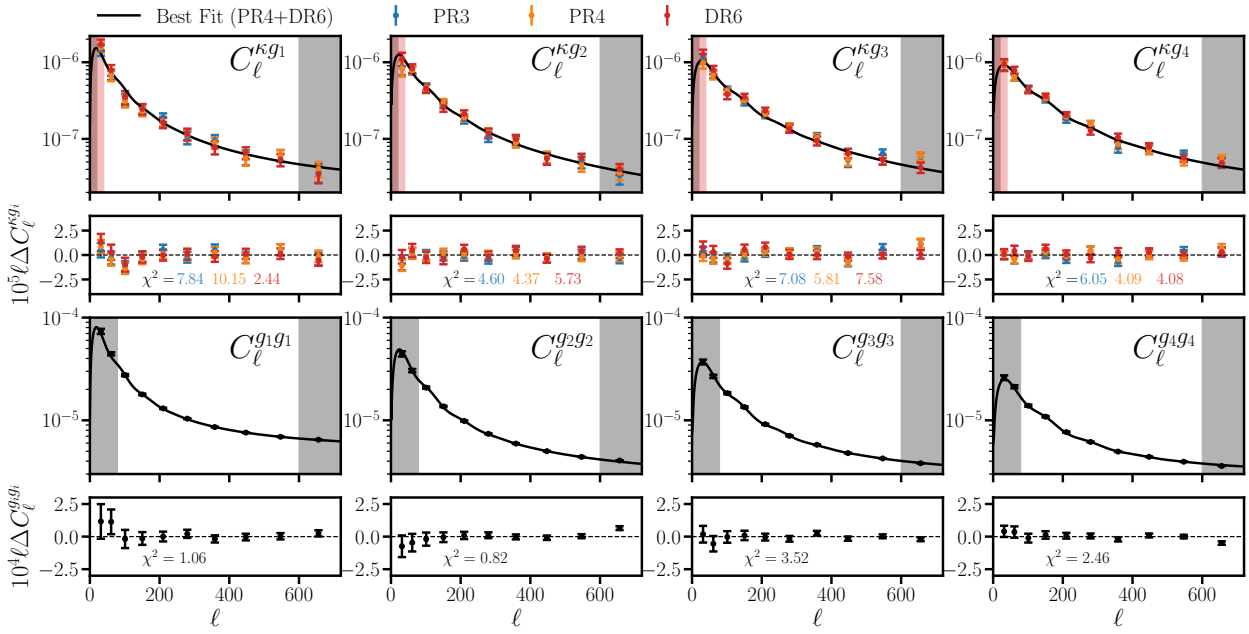


Figure 5.3: In the top row we plot the LRG cross-correlation ($C_\ell^{\kappa g}$) with *Planck* PR3 (blue), PR4 (orange) and ACT DR6 (red), while in the third row we plot the LRG auto-spectra (C_ℓ^{gg}). With the exception of the PR3 cross-correlation (computed following §5.3) these spectra and their covariances are taken from the companion paper [320] (and cross-validated in Appendix D.3). The hybrid covariance computed in the companion paper [320] finds that neighboring bandpowers are correlated by at most 4%. The black solid lines are the best-fit HEFT prediction when jointly fitting PR4 and DR6 using all four redshift bins. We show the residuals with the best-fit prediction in the second and fourth rows. For each measurement we quote its individual χ^2 value (indicated by color for the cross-correlation measurements). The grey shaded regions correspond to the scales omitted from our fiducial HEFT fits to the galaxy auto- and cross-correlation with *Planck*, while the shaded red regions are additionally excluded in our fits to the ACT cross-correlation (see Table 5.2). There are 9 data points for each *Planck* cross-correlation, 8 for each ACT cross-correlation and 7 for each galaxy auto-correlation. We estimate the goodness of fit in §5.7.

to mitigate numerical artifacts in the `NaMaster` implementation. Window functions are obtained from the `get_bandpower_windows` routine, which are computed out to $\ell_{\rm max} = 6143$. In §5.7, when "convolving" our theory prediction with the window function we truncate this sum at $\ell_{\rm max} = 1000$. For the relevant bandpowers in our analysis, the window functions are $< 10^{-6}$ for $\ell > 1000$.

After measuring a cross-correlation with a CMB lensing map we apply a "normalization" correction. The implementation of this correction differs slightly from that used in the companion paper [320], which we discuss in Appendix D.3.

We estimate the bandpower covariance with `NaMaster`'s `gaussian_covariance` routine, which in our case requires fiducial spectra for the CMB lensing power spectrum $C_\ell^{\kappa\kappa}$, CMB lensing reconstruction noise $N_\ell^{\kappa\kappa}$, cross-correlation between CMB lensing estimators (relevant only when combining PR4 and DR6), the galaxy power spectra $C_\ell^{g_i g_j}$ (for $i, j = 1, 2, 3, 4$) and their cross-correlation with CMB lensing $C_\ell^{\kappa g_i}$. Fiducial spectra are provided up to $\ell_{\rm max} = 3000$. We take $C_\ell^{\kappa\kappa}$ to be the fiducial CMB lensing power spectrum used in the ACT DR6 lensing simulations [493], and $N_\ell^{\kappa\kappa}$ to be the effective (sky-averaged) lensing reconstruction noise for a given CMB experiment (see sections 5.2 and 5.2). The cross-correlation between the *Planck* PR4 and ACT DR6 lensing reconstructions is calculated in the companion paper [320] (see also [493, 223]) using the FFP10 simulations [24]. Fiducial curves for $C_\ell^{g_i g_i}$ and $C_\ell^{\kappa g_i}$ are obtained by fitting to each measurement *individually* as we now describe[11]. For each measured $\tilde{C}_L$ (either $\tilde{C}_L^{g_i g_i}$ or $\tilde{C}_L^{\kappa g_i}$ for a single i) and HEFT prediction (see §5.4) C_ℓ we assign a loss function: $\sum_L (\tilde{C}_L - \sum_\ell W_{L\ell} C_\ell)^2 / \tilde{C}_L^2$, where the sum over ℓ runs from 0 to 3000 and the sum over bandpowers runs up to $L_{\rm edge,\,max} = 2865$. In the HEFT prediction we hold the cosmology fixed to that used in the ACT DR6 lensing simulations[12] and use the measured redshift distributions of the LRGs (Fig. 5.2). A fiducial theory prediction is obtained by varying the nuisance terms to minimize the loss function[13]. Fiducial curves for the galaxy cross-spectra ($C_\ell^{g_i g_j}$ for $i \neq j$) are obtained via Eq. (5.14) by linearly interpolating the best-fit nuisance terms from the galaxy auto fits with redshift.

We have checked that replacing the loss function with a Gaussian likelihood (using an analytic covariance with fiducial curves estimated via the preceding paragraph) and repeating the procedure above results in subpercent changes to the best-fit $C_\ell^{g_i g_j}$'s and $C_\ell^{\kappa g_i}$'s.

5.4 Modeling

Aemulus ν

The models discussed in the following subsections are built on the `Aemulus` ν simulations [181], which are a set of 150 N-body simulations spanning a broad range of cosmologies (in particular, $1.1 < 10^9 A_s < 3.1$ and $0.08 < \omega_c < 0.16$) including both massive neutrinos and variations to the dark energy equation of state. The simulations treat the CDM+baryon (cb) and neutrino fields as two sets of 1400^3 collisionless particles that are evolved in a $1.05h^{-1}$Gpc box, yielding a cb mass resolution of $3.51 \times 10^{10} (\Omega_{cb}/0.3) h^{-1} M_\odot$. The initial displacements ($z_{\rm ini} = 12$) of the cb field are calculated using third-order Lagrangian Perturbation Theory (as implemented in `monofonic` [406, 275]) from an initial Gaussian density field whose power

[11] We use the cross-correlation with PR4 when estimating covariances. We note that the exact procedure for obtaining fiducial theory spectra is not important provided that they are in good agreement with the data (once binned with the respective window functions). For example, [625] simply fit a polynomial to the data.

[12] $A_s = 2.15086 \times 10^{-9}$, $\omega_c = 0.1203058$, $n_s = 0.9625356$, $h = 0.6702393$, $\omega_b = 0.02219218$, $M_\nu = 60$ meV

[13] As a technical detail, we actually minimize the sum of the loss function and the χ^2 associated with the nuisance parameter priors (listed in Table 5.3). Doing so stabilizes the minimization.

spectrum is taken to be $P_{cb}(k, z_{\rm ini}) = D_{cb}(k, z_{\rm ini})^2 P_{cb,\rm lin}(k, z = 0)$, where $P_{cb,\rm lin}(k, z = 0)$ is computed with `CLASS` [97] and the scale-dependent growth factor is computed using first-order Newtonian fluid approximation as implemented in `zwindstroom` [212], while the initial conditions for the neutrinos are set with `fastDF` [213]. The cb and neutrino particles are evolved from $z_{\rm ini} = 12$ to $z = 0$ using a modified version of `Gadget-3` [226] that enables modifications and includes relativistic corrections to the background evolution $H(z)$.

Hybrid Effective Field Theory

We adopt a Hybrid Effective Field Theory (HEFT) model [413, 327, 268, 637, 461, 440] in our fiducial analysis. Within this framework the "initial" (Lagrangian) galaxy density is modeled as a biased tracer of the Lagrangian CDM+baryon density. Proto-galaxies at Lagrangian position $\boldsymbol{q}$ are advected to their later (real-space) positions $\boldsymbol{x}(t) = \boldsymbol{q} + \boldsymbol{\Psi}_{cb}(\boldsymbol{q}, t)$ following the (non-linear) displacement vector $\boldsymbol{\Psi}_{cb}(\boldsymbol{q}, t)$ of the CDM+baryon field [601, 122, 121, 600, 364, 426, 231, 140] that is measured from simulations. From number conservation, the (real-space) late-time galaxy density contrast is given by

$$1 + \delta_g(\boldsymbol{x}) = \int d^3\boldsymbol{q}\, F(\boldsymbol{q})\, \delta^D(\boldsymbol{x} - \boldsymbol{q} - \boldsymbol{\Psi}_{cb}) \tag{5.1}$$

where the weight function $F(\boldsymbol{q})$ characterizes the initial galaxy density fluctuations and δ^D is the Dirac delta function. Here we take $F(\boldsymbol{q})$ to be a perturbative expansion in scalar combinations of the tidal ($\partial_i\partial_j\Phi$) and velocity ($\partial_i v_j$) tensors to next-to-leading order in derivatives and perturbations [395]

$$F(\boldsymbol{q}) = 1 + b_1^L \delta_{cb}(\boldsymbol{q}) + \frac{b_2^L}{2}\left(\delta_{cb}^2(\boldsymbol{q}) - \langle\delta_{cb}^2\rangle\right) + b_s^L\left(s_{cb}^2(\boldsymbol{q}) - \langle s_{cb}^2\rangle\right) + \frac{b_{\nabla^2}^L}{4}\left(\nabla^2\delta_{cb}(\boldsymbol{q}) - \langle\nabla^2\delta_{cb}\rangle\right) + \mathcal{E}(\boldsymbol{q}) \tag{5.2}$$

where $s_{cb}^2 = \sum_{ij} s_{cb}^{ij} s_{cb}^{ij}$ is the square of the shear field $s_{cb}^{ij} = (\partial_i\partial_j/\partial^2 - \delta_{ij}^K/3)\delta_{cb}$ and $\mathcal{E}(\boldsymbol{q})$ is a small-scale stochastic component. Within this approximation the late-time galaxy density contrast takes the form

$$\delta_g(\boldsymbol{k}) = \left(1 - \frac{b_{\nabla^2}^L k^2}{4}\right)\delta_{cb}(\boldsymbol{k}) + \sum_X b_X^L F_X(\boldsymbol{k}) + \mathcal{E}(\boldsymbol{k}) \tag{5.3}$$

where $X \in \{1, 2, s\}$, $F_X(\boldsymbol{k})$ corresponds to the respective term (post-advection and Fourier transformed) in the weight function, and we have approximated the post-advection $\nabla^2\delta_{cb}$ term as $-k^2\delta_{cb}(\boldsymbol{k})$. The galaxy power spectrum is then

$$P_{gg}(k) = \left(1 - \frac{\alpha_a k^2}{2}\right) P_{cb}(k) + 2\sum_X b_X^L P_{cb,X}(k) + \sum_{XY} b_X^L b_Y^L P_{XY}(k) + \mathrm{SN}^{\rm 3D} \tag{5.4}$$

where P_{cb} is the power spectrum of the CDM+baryon field, $P_{cb,X}$ is the cross-correlation of the CDM+baryon field with F_X, P_{XY} is the cross-spectrum of F_X and F_Y, $\mathrm{SN}^{\rm 3D}$ is a

shot noise contribution, we introduced the "counterterm" α_a for the auto-correlation, and we have neglected $\mathcal{O}(k^2\delta_{cb}^3)$ contributions. Likewise the cross-correlation of the galaxies with the matter is approximately

$$P_{gm}(k) = \left(1 - \frac{\alpha_x k^2}{2}\right) P_{cb,m}(k) + \sum_X b_X^L P_{m,X}(k) \tag{5.5}$$

where $P_{cb,m}$ is the cross spectrum of the CDM+baryon field with the total matter field, $P_{m,X}$ is the cross correlation of the total matter field with F_X, and we have defined the counterterm α_x for the cross-correlation.

The power spectra ($P_{cb}, P_{cb,X}$, P_{XY}, $P_{cb,m}$, and $P_{m,X}$) appearing in Eqs. (5.4) and (5.5) are calculated in the `Aemulus` ν simulations described in §5.4. The sample variance of these measurements on perturbative scales has been suppressed using the Zel'dovich density field (implemented using the `ZeNBu` and `velocileptors` [136] codes) as a control variate[14]. We use the `Aemulus` ν emulator in our calculations, which has been shown to be accurate to within 0.25% for the scales ($k \lesssim 0.6\,h\,\mathrm{Mpc}^{-1}$) and redshifts ($z \lesssim 1$) relevant for our analysis [181].

While we have introduced two counterterms (α_x and α_a) in Eqs. (5.4) and (5.5), these parameters are not independent in the limit that the bias expansion (Eq. 5.2) accurately describes the *field-level* galaxy distribution and the simulated $\delta_m(\boldsymbol{x})$ accurately describes the true matter distribution[15]. The former approximation is valid up to baryonic feedback and the finite mass resolution of the simulations. Both of these effects are small on the scales of interest and have the qualitative impact of damping small-scale power that we approximate by introducing an effective "derivative bias" in the matter distribution: $\delta_m(\boldsymbol{k}) \to \left(1 - \tilde{\epsilon}k^2/4\right)\delta_m(\boldsymbol{k})$. Baryonic feedback scenarios capable of resolving the S_8-tension in galaxy-lensing measurements (Fig 6 of ref. [40], C-OWLS AGN with $\log_{10}(\Delta T_{\rm heat}/\mathrm{K}) = 8.7$ at $z = 0$) require a $\sim 10\%$ suppression to the matter power spectrum at $k \simeq 0.4\,h\,\mathrm{Mpc}^{-1}$ corresponding to $\tilde{\epsilon} \simeq 1.3\,h^{-2}\mathrm{Mpc}^2$, while the effective grid scale in the `Aemulus` ν simulations is $\approx 1\,h^{-1}\mathrm{Mpc}$ corresponding to $\tilde{\epsilon} \simeq 1\ h^{-2}\mathrm{Mpc}^2$. With the inclusion of the effective derivative bias in the matter distribution the $\mathcal{O}(k^2)$ components of the galaxy-auto and cross-correlation with

[14]Specifically, [181] estimates a power spectrum (e.g. $P_{X,Y}$) using $\hat{P}^{\rm CV}_{X,Y} \equiv \hat{P}^{\rm Aemulus}_{X,Y} - \beta(\hat{P}^{\rm LPT}_{X,Y} - P^{\rm LPT}_{X,Y})$, where $\hat{P}^{\rm Aemulus}_{X,Y}$ is the measured power spectrum from the Aemulus ν simulation, $\hat{P}^{\rm LPT}_{X,Y}$ is the measured power spectrum of the (Zel'dovich) LPT evolution of the same initial conditions used in the simulation and $P^{\rm LPT}_{X,Y}$ is the analytic prediction for the ensemble average (which is known to machine precision). The variance of $\hat{P}^{\rm CV}_{X,Y}$ is (optimally) suppressed by a factor of $1-\rho^2_{\rm Aemulus,LPT}$ relative to the variance of $\hat{P}^{\rm Aemulus}_{X,Y}$ if one chooses $\beta = \mathrm{Cov}[\hat{P}^{\rm Aemulus}_{X,Y}, \hat{P}^{\rm LPT}_{X,Y}]/\mathrm{Var}[\hat{P}^{\rm LPT}_{X,Y}]$, where $\rho_{\rm Aemulus,LPT}$ is the correlation coefficient of the Aemulus and LPT measurements.

[15]This is analogous to multi-tracer studies where the coefficients of counterterm-like contributions to the cross-correlation of two biased tracers can be expressed in terms of counterterm coefficients appearing in the two individual auto-spectra (shown explicitly in Appendix C of [204]).

matter are

$$\begin{aligned} P_{gg}(k) &\supset 2 \times \left\langle \left[-\frac{b_{\nabla^2}^L k^2}{4} \delta_{cb}(\boldsymbol{k}) \right] \times \left[\delta_{cb}(\boldsymbol{k}) + b_1^L F_1(\boldsymbol{k}) \right]^* \right\rangle' \\ &= -\frac{b_1^E b_{\nabla^2}^L k^2}{2} P_{cb}(k) \\ P_{gm}(k) &\supset \left\langle \left[\left(1 - \frac{b_{\nabla^2}^L k^2}{4} \right) \delta_{cb}(\boldsymbol{k}) + b_1^L F_1(\boldsymbol{k}) \right] \left(1 - \frac{\tilde{\epsilon} k^2}{4} \right) \delta_m^*(\boldsymbol{k}) \right\rangle' \\ &\supset -\frac{(b_1^E \tilde{\epsilon} + b_{\nabla^2}^L) k^2}{4} P_{cb,m}(k) \end{aligned} \tag{5.6}$$

where we neglect $\mathcal{O}(k^2\delta^3)$ contributions (in particular, $F_1(\boldsymbol{k}) = \delta_{cb}(\boldsymbol{k})$ under this approximation), we introduced the Eulerian bias $b_1^E \equiv 1 + b_1^L$, the symbol $\supset$ denotes a subset of the terms appearing in the power spectra predictions (in this case, the k^2 terms), and primed brackets correspond to an ensemble average with the momentum-conserving delta function removed, e.g. $\langle \delta(\boldsymbol{k}_1)\delta(\boldsymbol{k}_2)\rangle \equiv (2\pi)^3 \delta^D(\boldsymbol{k}_1 + \boldsymbol{k}_2)\langle \delta(\boldsymbol{k}_1)\delta(\boldsymbol{k}_2)\rangle'$. Comparing Eq. (5.6) to Eqs. (5.4) and (5.5) we see that $\alpha_a = b_1^E b_{\nabla^2}^L$ and $2\alpha_x = b_1^E \tilde{\epsilon} + b_{\nabla^2}^L$, which in turn implies

$$\alpha_x = \frac{\alpha_a}{2 b_1^E} + \epsilon, \tag{5.7}$$

where $\epsilon \equiv b_1^E \tilde{\epsilon}/2$ ($b_1^E/2 \simeq 1$ for the LRGs considered here) has a "typical" value of $\simeq 1\,h^{-2}\mathrm{Mpc}^2$, which we marginalize over with an informative Gaussian prior centered at zero with width $\simeq 2\,h^{-2}\mathrm{Mpc}^2$ in our fiducial analysis (see §5.5). We note that this prior is sufficiently broad to permit an enhancement of small-scale power (within one prior σ) rather than a suppression, even when taking into account the suppression already present in the model prediction arising from the effective grid scale in the `Aemulus` ν simulations.

Linear theory

Non-linear corrections to the matter power spectrum are $\leq 1.5\%$ for $k \leq 0.1\,h\,\mathrm{Mpc}^{-1}$ and $z > 0.3$. Provided that the length scale associated with the bias expansion (e.g. the Lagrangian radius of LRG host halos) is smaller than that associated with non-linear (gravitational) dynamics[16], one can accurately model the galaxy distribution using linear theory on these scales, which we consider as an alternative to our fiducial HEFT model in §5.7.

Rather than using pure linear theory for the galaxy cross- and auto-correlations, we choose to use the HEFT prediction with $b_2^L = b_s^L = 0$, and adjust our scale cuts such that $\ell_{\max}(z) \simeq \chi(z) \times (0.1\,h\,\mathrm{Mpc}^{-1})$ (see §5.5 and Table 5.2). This is done purely for convenience (e.g. the emulator is faster than a `CLASS` evaluation), since swapping the HEFT predictions

[16] We note that this approximation is implicitly made when fitting to higher k with HEFT than with "pure" perturbation theory. A caveat to this claim is that some higher-order contributions, such as $\langle \delta^2 \delta^2 \rangle$, are non-zero in the $k \to 0$ limit. However, these contributions are scale-independent in this limit and thus degenerate with shot noise.

for linear theory on these scales would have a negligible impact on our constraints. To approximately account for the residual systematic error in P_{gg} and P_{gm} from neglecting higher-order contributions on these scales we choose to marginalize over α_a and α_x independently with an informative prior chosen such that a 1.5% correction is allowed at $k = 0.1\,h\,\mathrm{Mpc}^{-1}$ at 1σ (see §5.5 and Table 5.3). As for our fiducial model, we use the non-linear P_{mm} prediction when calculating magnification contributions (§5.4).

Model independent constraints

As yet another alternative to HEFT and linear theory, we consider a "model independent" (or "fixed shape") approach to constrain $\sigma_8(z)$. Under this approach we use the same scale cuts as for our linear theory analysis, fix the cosmology, and allow the linear bias for the galaxy auto ($b_1^{L,a}$) and cross ($b_1^{L,x}$) to differ by a factor

$$\alpha_8 \equiv \frac{1 + b_1^{L,x}}{1 + b_1^{L,a}}. \tag{5.8}$$

Within linear theory one constrains σ_8 schematically through the ratio $C^{\kappa g}/\sqrt{C^{gg}} \sim \alpha_8 \sigma_{8,\,\mathrm{fid}}$, where $\sigma_{8,\,\mathrm{fid}}$ is the value of σ_8 for the fiducial cosmology. In §5.7 we consider fitting to each redshift bin independently, and interpret our constraint on $\alpha_8\,\sigma_{8,\,\mathrm{fid}}(z_i)$ as a model independent measurement of $\sigma_8(z_i)$.

From 3D to 2D

We use the Kaiser-Limber approximation [352, 317] (hereafter Limber approximation) to model the angular galaxy auto- and cross-correlation:

$$\begin{aligned} C_\ell^{gg} &= \int_0^\infty \frac{d\chi}{\chi^2}\left[[W^g(\chi)]^2\, P_{gg}\,(k,z) + 2W^g(\chi)W^\mu(\chi)P_{gm}(k,z) + [W^\mu(\chi)]^2\, P_{mm}(k,z)\right] \\ C_\ell^{\kappa g} &= \int_0^\infty \frac{d\chi}{\chi^2} W^\kappa(\chi)\left[W^g(\chi)P_{gm}(k,z) + W^\mu(\chi)P_{mm}(k,z)\right] \end{aligned} \tag{5.9}$$

where $k = (\ell + 1/2)/\chi$ [365] and z is implicitly a function of χ. The projection kernels for the galaxies, magnification contribution, and CMB lensing are [603, 61, 417]

$$\begin{aligned} W^g(\chi) &= H(z)\phi(z) \\ W^\mu(\chi) &= (5s_\mu - 2)\frac{3}{2}\Omega_m H_0^2(1+z)\int_z^\infty dz' \frac{\chi(z)(\chi(z') - \chi(z))}{\chi(z')}\phi(z') \\ W^\kappa(\chi) &= \frac{3}{2}\Omega_m H_0^2(1+z)\frac{\chi(\chi_* - \chi)}{\chi_*} \end{aligned} \tag{5.10}$$

where $\phi(z) \propto dN/dz$ is the normalized ($\int dz\phi(z) = 1$) redshift distribution, χ_* (z_*) is the comoving distance (redshift) to the surface of last scattering, and s_μ is the number

count slope. Due to the narrow redshift distributions of the LRG bins, the magnification contributions to the galaxy auto-spectra amount to less than a percent of the total signal, while for the cross-correlation with CMB lensing the magnification contribution is at most $\simeq 10\%$ for the highest redshift bin (and smaller for the remaining bins). We compute background quantities (i.e Ω_m, χ_*, $H(z)$, and $\chi(z)$) relevant for Limber integration from our sampled cosmological parameters (§5.5 and Table 5.3) using `CLASS` [97].

In practice, we neglect the evolution of P_{gg} and P_{gm} in the Limber integrals, and evaluate them at the effective redshift

$$z_{\rm eff} = \frac{\int d\chi\, z(\chi)\, [W^g(z)]^2/\chi^2}{\int d\chi [W^g(z)]^2/\chi^2}. \tag{5.11}$$

This approximation, which we expect to be accurate to subpercent precision for our LRG sample (see Fig. 5.4) due to its narrow redshift distributions, has the advantage of making our analysis largely agnostic to the assumed redshift evolution of the galaxy nuisance parameters within each redshift bin. Instead, we treat e.g. $b_1^L(z_i)$ as a single number evaluated at the effective redshift for the i'th redshift bin.

Deviations from the Limber approximation are primarily relevant on large scales where linear theory is a good approximation. Within this approximation and neglecting magnification bias, the spherical harmonic coefficients $g_{\ell m}$ of the projected galaxy density contrast can be split in two pieces $g^{\rm real}_{\ell m} + g^{\rm RSD}_{\ell m}$ where [233, 451]

$$\begin{aligned} g^X_{\ell m} &\equiv \int \frac{d^3k}{(2\pi)^3} Y^*_{\ell m}(\hat{\boldsymbol{k}}) \int d\chi \mathcal{W}^X_\ell(k,\chi)\delta_g(\boldsymbol{k}) \\ \mathcal{W}^{\rm real}_\ell(k,\chi) &= 4\pi i^\ell W^g(\chi) j_\ell(k\chi) \\ \mathcal{W}^{\rm RSD}_\ell(k,\chi) &= 4\pi i^\ell \frac{1}{k}\frac{dW^g}{d\chi} j'_\ell(k\chi)\beta(z), \end{aligned} \tag{5.12}$$

$\delta_g(\boldsymbol{k})$ is the real-space 3D galaxy density contrast, $\beta(z) = f(z)/b_1^E(z)$ where $f(z)$ is the linear growth rate and $j'_\ell(x) = \partial_x j_\ell(x)$ is the derivative of the spherical Bessel function of the first kind. The cross-correlation $C_\ell^{XY} \equiv \langle g^X_{\ell m} g^{Y*}_{\ell m}\rangle$ is given by the "full integral"

$$C_\ell^{XY} = \int \frac{k^2 dk}{(2\pi)^3} \int d\chi d\chi' \mathcal{W}^X_\ell(k,\chi)\mathcal{W}^Y_\ell(k,\chi') P_{gg}(k;\chi,\chi') \tag{5.13}$$

where $\langle \delta_g(\boldsymbol{k},\chi)\delta_g(\boldsymbol{k}',\chi')\rangle \equiv (2\pi)^3\delta^D(\boldsymbol{k}+\boldsymbol{k}')P_{gg}(k;\chi,\chi')$.

In Fig. 5.4 we show the full integral for both the auto- and cross-correlations assuming linear theory: $P_{gg}(k;\chi,\chi') = D(z)b(z)D(z')b(z')P_0(k)$ where $D(z)$ is the linear growth factor normalized to $D(0)=1$ and $P_0(k)$ is the cb linear power spectrum at $z=0$ (we are ignoring shot noise). We take $b(z)$ to be a linear interpolation of the best-fit values listed in Table 1 of [625]. We also show the Limber approximations, including the effective redshift approximation. For $\ell \lesssim 75$, the largest beyond Limber effect is due to RSD, while above $\ell \gtrsim 75$ the largest effect is due to the effective redshift approximation, which is still $< 1\%$ and neglected in our analysis.

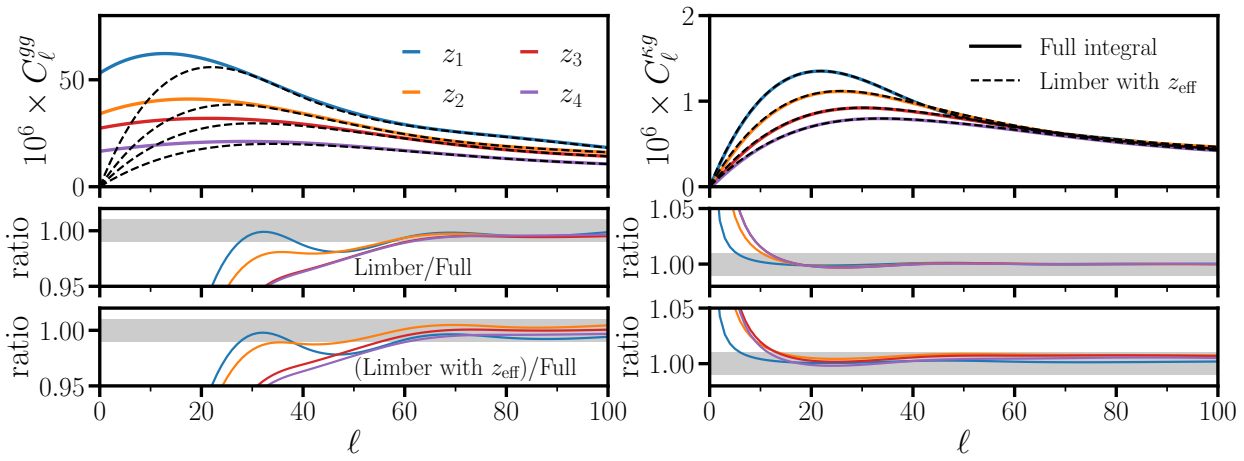


Figure 5.4: Estimation of "beyond Limber" effects within linear theory for both the galaxy auto- and cross-correlation (left and right). The middle panels show the ratio of the Limber approximation to the full integral, while the bottom panels show the ratio when evaluating the power spectra at a fixed redshift $z_{\rm eff}$, defined in Eq. (5.11). The shaded bands correspond to 1 ± 0.01.

Pixelization

The LRGs considered here have been placed on a HEALPix grid[17]. This pixelization has the approximate effect of convolving the galaxies with the pixel window function, but in addition gives rise to aliasing of power [167, 539, 53, 515]. The combination of these effects can be well approximated by taking $C_\ell^{\kappa g} \to w_\ell C_\ell^{\kappa g}$ and $C_\ell^{gg} \to w_\ell^2 (C_\ell^{gg} - \mathrm{SN}^{\rm 2D}) + \mathrm{SN}^{\rm 2D}$, where w_ℓ is the pixel window function (as determined by `healpy`'s `pixwin` for the appropriate `nSide`). Previous analyses (e.g. [625, 330]) have approximately corrected for the pixel window function at the data level, which requires assuming a fiducial value for the projected shot noise. This is somewhat undesirable (but still reasonable) given that one fits for the shot noise in practice. Here we opt to forward model the impact of pixelization using the aforementioned approximation. For `nSide = 2048` the window function differs from one by less than 0.4% for $\ell < 600$, and as a result this improved treatment of the pixel window function has a minuscule impact on our final constraints.

[17] See [53] for an alternative approach that bypasses placing galaxies on a discrete grid, and hence the pixel window function and aliasing.

5.5 Likelihood and pipeline checks

Parameter inference is performed with the `cobaya` [586, 587] sampling framework. We make our likelihood publicly available[18] (`MaPar`) and summarize its code structure in Appendix D.1. Chains are sampled using `cobaya`'s Markov Chain Monte Carlo Metropolis sampler [345, 343] and are considered converged when the Gelman-Rubin [245, 108] statistic satisfies $R - 1 \leq 0.01$. We obtain marginal distributions with `GetDist` [344] with the first 30% of the chains removed as burn-in. Best-fit points are obtained using `cobaya`'s default minimizer (see Appendix D.4 for a discussion of analytic minimization for linear model parameters).

We adopt a Gaussian likelihood throughout and analytically marginalize over all parameters that linearly appear in our theory prediction: $\alpha_a(z_i)$, $\alpha_x(z_i)$ and shot noise. In our fiducial analysis analytic marginalization reduces the number of directly sampled parameters from 30 (2 cosmological and 7×4 nuisance parameters) to 18 which dramatically improves the convergence time of our chains. We discuss our implementation of analytic marginalization in Appendix D.4. In the same Appendix we explicitly verify that our implementation of analytic marginalization is in agreement with the brute-force approach.

In §5.7 we place tight constraints on σ_8 after including a BAO prior, which efficiently breaks the $\Omega_m - \sigma_8$ degeneracy. To mimic this in both our "volume effects" tests (§5.5) and fits to mock data (§5.5) we include a mock BAO prior when quoting σ_8 constraints. We construct this prior using the covariances from [84, 504, 28] and adjust the central value of each distance measurement to the predicted ΛCDM value assuming a Buzzard cosmology (see `MaPar/mocks/mockBAO/`).

Scale cuts

Throughout §5.7 we use (a subset of) the bandpowers discussed in §5.3 spanning the range $20 \leq \ell < 600$. The fiducial scale cuts used for each our modeling choices (sections 5.4, 5.4 and 5.4) are summarized in Table 5.2. As discussed in §5.4, we adopt the Limber approximation and neglect the impact of redshift space distortions in our model, which thus acts as a systematic on large scales. For both the galaxy auto- and cross-correlation with *Planck*, we adopt an $\ell_{\rm min}$ (20 for κg and 79 for gg) where the "beyond Limber" corrections are less than a percent (see Fig. 5.4). On large scales the ACT DR6 lensing reconstruction is contaminated by a mean field contribution that greatly exceeds the CMB lensing signal for $\ell \lesssim 30$ (e.g. Fig. 9 of [493]), making it difficult accurately correct for in this regime. Following suit with the DR6 auto-correlation measurement [493] (which adopted $\ell_{\rm min} = 40$) we choose to adopt a larger $\ell_{\rm min} = 44$ for the ACT cross-correlation to mitigate spurious correlations with the mean field.

At high ℓ our scale cuts vary with the model being used. Previous work [327] has found that our fiducial HEFT model (§5.4) is accurate to subpercent precision for $k \leq$

[18] with the exception of the neural network weights for the HEFT emulator, which will be made public upon the publication of an upcoming DESI-DES Y3 cross-correlation analysis [129].

		Planck $C_\ell^{\kappa g}$				ACT $C_\ell^{\kappa g}$				C_ℓ^{gg}			
		z_1	z_2	z_3	z_4	z_1	z_2	z_3	z_4	z_1	z_2	z_3	z_4
HEFT (fiducial)	$\ell_{\rm min}$	20	20	20	20	44	44	44	44	79	79	79	79
	$\ell_{\rm max}$	600	600	600	600	600	600	600	600	600	600	600	600
Linear theory	$\ell_{\rm min}$	–	20	20	20	–	44	44	44	–	79	79	79
	$\ell_{\rm max}$	–	178	243	243	–	178	243	243	–	178	243	243
Model independent	$\ell_{\rm min}$	20	20	20	20	44	44	44	44	79	79	79	79
	$\ell_{\rm max}$	178	178	243	243	178	178	243	243	178	178	243	243

Table 5.2: The fiducial scale cuts used for each of our modeling choices. We drop the lowest redshift bin when using linear theory, and note that $\ell_{\rm min}$ and $\ell_{\rm max}$ refer to bandpower bin edges as opposed to their "centers" (as was done in [625]).

$0.6\,h\,{\rm Mpc}^{-1}$ when fitting to dark matter halos with the same characteristic masses ($10^{12.5} < \log_{10} M_h/M_\odot < 10^{13}$) and redshifts as expected from recent HOD fits [635] for the host halos of our LRG sample. This motivates $\ell_{\rm max} \simeq \chi(z_{\rm min}) \times (0.6\,h\,{\rm Mpc}^{-1})$ where $z_{\rm min}$ is the lower edge of a given redshift bin. In particular, for the first redshift bin $z_{\rm min} \sim 0.35$ and thus $\ell_{\rm max} \simeq 600$. In principle one could reliably extend to a higher $\ell_{\rm max}$ for the higher redshift samples, however in practice we expect limited gains from higher ℓ due to shot noise, lensing reconstruction noise, and degeneracies with galaxy bias. For simplicity we adopt a redshift-independent $\ell_{\rm max} = 600$ in our fiducial HEFT analysis for both the galaxy auto- and cross-correlation.

As discussed in §5.4 we tune our linear theory scale cuts such that $\ell_{\rm max} \simeq \chi(z_{{\rm eff},i}) \times (0.1\,h\,{\rm Mpc}^{-1})$, where the differences between the linear and non-linear matter power spectrum are $\leq 1.5\%$. For the lowest redshift bin applying this scale cut removes all but one bandpower for the galaxy auto-correlation ($79 < \ell < 124$), which we opt to discard when quoting linear theory constraints. For bins z_2, z_3 and z_4 we adopt $\ell_{\rm max} = 178$, 243, 243 corresponding to $k_{\rm max} \simeq 0.11$, 0.13, $0.11\,h\,{\rm Mpc}^{-1}$. We use the same $\ell_{\rm max}$'s used in our linear theory fits for our model independent constraints. Unlike for linear theory, we consider a model independent $\sigma_8(z_1)$ measurement in the lowest redshift bin, for which we take $\ell_{\rm max} = 178$. In §5.5 we verify that all of our models recover unbiased results on simulations when adopting these scale cuts.

Fiducial cosmology, sampled parameters and priors

Our data (Fig. 5.3) are particularly powerful at extracting the relative amplitude between $C_\ell^{\kappa g}$ and $\sqrt{C_\ell^{gg}}$ on large scales, roughly corresponding to S_8. Since our goal is to assess the consistency of a low redshift structure growth measurement with that predicted within ΛCDM conditioned on primary CMB data, we fix the remaining relevant ΛCDM parameters to their

Planck 2018 [475] mean values[19]. Specifically, we fix the baryon abundance $\Omega_b h^2 = 0.02236$, spectral index $n_s = 0.9649$, $\Omega_m h^3 = 0.09633$ as a proxy for the angular acoustic scale θ_* [463, 475] and additionally fix the sum of the neutrino masses to $\sum m_\nu = 0.06$ eV. We directly sample the (log) primordial power-spectrum amplitude $\ln(10^{10}A_s)$ and dark matter abundance $\Omega_c h^2$ with uniform $\mathcal{U}(2,4)$ and $\mathcal{U}(0.08,0.16)$ priors respectively. When quoting model-independent constraints these parameters are fixed to 3.045 and 0.1202 respectively [475] and we vary α_8 (Eq. 5.8) with a $\mathcal{U}(0.5,1.5)$ prior. We note that the $\ln(10^{10}A_s)$ prior extends beyond the range of the `Aemulus` ν simulations (§5.4), however, our data are sufficiently constraining such that for nearly all[20] of the results presented in §5.7 the 2σ credible intervals lie within the `Aemulus` ν range.

Parameter	Description	Prior or fixed value		
...	...	HEFT (fiducial)	Linear theory	Model independent
$\ln(10^{10}A_s)$	ln(primordial amplitude)	$\mathcal{U}(2,4)$	$\mathcal{U}(2,4)$	3.045
$\Omega_c h^2$	dark matter abundance	$\mathcal{U}(0.08,0.16)$	$\mathcal{U}(0.08,0.16)$	0.1202
α_8	Eq. (5.8)	–	–	$\mathcal{U}(0.5,1.5)$
$b_1^L(z_i)$	linear (Lagrangian) bias	$\mathcal{U}(0,3)$	$\mathcal{U}(0,3)$	$\mathcal{U}(0,3)$
$b_2^L(z_i)$	quadratic bias	$\mathcal{U}(-5,5)$	0	0
$b_s^L(z_i)$	shear bias	$\boldsymbol{\mathcal{N}(0,1)}$	0	0
$s_\mu(z_i)$	number count slope	$\boldsymbol{\mathcal{N}}$**(Tab. 5.1, 0.1)**	$\boldsymbol{\mathcal{N}}$**(Tab. 5.1, 0.1)**	$\boldsymbol{\mathcal{N}}$**(Tab. 5.1, 0.1)**
$\mathrm{SN}^{\mathrm{2D}}(z_i)$	projected shot noise	$\boldsymbol{\mathcal{N}_r}$**(Tab. 5.1, 0.3)**	$\boldsymbol{\mathcal{N}_r}$**(Tab. 5.1, 0.3)**	$\boldsymbol{\mathcal{N}_r}$**(Tab. 5.1, 0.3)**
$\alpha_a(z_i)$ $[h^{-2}\,\mathrm{Mpc}^2]$	counterterm (auto)	$\mathcal{N}(0,50)$	$\boldsymbol{\mathcal{N}(0,3)}$	$\boldsymbol{\mathcal{N}(0,3)}$
$\alpha_x(z_i)$ $[h^{-2}\,\mathrm{Mpc}^2]$	counterterm (cross)	Eq. (5.7)	$\boldsymbol{\mathcal{N}(0,3)}$	$\boldsymbol{\mathcal{N}(0,3)}$
$\epsilon(z_i)$ $[h^{-2}\,\mathrm{Mpc}^2]$	Eq. (5.7)	$\boldsymbol{\mathcal{N}(0,2)}$	–	–

Table 5.3: Sampled parameters and their associated priors for each of our modeling choices. Parameters with argument z_i (for $i = 1,2,3,4$) are redshift dependent, e.g. $b_s^L(z_1)$ and $b_s^L(z_2)$ are treated as separate parameters that are individually sampled. With the exception of shot noise and number count slope, the priors for all nuisance parameters are identical in each redshift bin. We center the prior means for $s_\mu(z_i)$ and $\mathrm{SN}^{\mathrm{2D}}(z_i)$ around those listed in Table 5.1. $\mathcal{U}(X,Y)$ denotes a uniform prior between X and Y, $\mathcal{N}(\mu,\sigma)$ is a Gaussian prior with mean μ and standard deviation σ, and $\mathcal{N}_r(\mu,\sigma) \equiv \mathcal{N}(\mu,\sigma\mu)$. Priors are in **bold font** when the posterior of a nuisance parameter is (approximately) prior-dominated.

Priors on galaxy-induced nuisance parameters vary with the model being used and are summarized in Table 5.3. In all scenarios we place uniform priors on the linear Lagrangian bias (in each redshift bin) between 0 and 3, place a Gaussian prior on shot noise centered around its Poisson value (Table 5.1) with a (relative) 30% width, and a Gaussian prior on

[19] We use the values from the TT,TE,EE+lowE column in Table 2 of [475], which contains no low redshift information modulo lensing and other secondary anisotropies.

[20] The one exception is when analyzing z_1 independently without a BAO prior. Restricting the $\ln(10^{10}A_s)$ prior to $\mathcal{U}(2.4,3.4)$ (corresponding to the Aemulus ν range) truncates the $\sigma_8-\Omega_m$ contour shown in Fig. 5.14, however, the mean S_8 is largely unaffected ($\Delta S_8 = 0.003$, corresponding to a $< \sigma/10$ shift).

the number count slopes centered around their measured values [643] (Table 5.1) with width 0.1, which is $\simeq 10\times$ larger than the Poisson errors estimated by ref. [643].

For our fiducial HEFT model we adopt a uniform prior on (b_2^L) between -5 and 5. This prior is largely "uninformative" in the sense that the data are readily capable of distinguishing e.g. $b_2^L = -5$ from $b_2^L = +5$, resulting in marginal posteriors that are significantly narrower than the prior (see Fig. D.4 in Appendix D.5). We find that b_s^L is highly-degenerate with b_2^L (see Fig. D.3 in Appendix D.4) and that the best-fit values of b_s^L to mock LRGs (§5.5) are typically of $\mathcal{O}(0.1)$. For these reasons we apply an informative $\mathcal{N}(0,1)$ Gaussian prior on each $b_s^L(z_i)$, and explore the impact of widening this prior in §5.7. We adopt a wide Gaussian prior with mean zero and width 50 $h^{-2}\,\mathrm{Mpc}^2$ on the (auto) counterterm $\alpha_a(z_i)$, while the (cross) counterterm $\alpha_x(z_i)$ is determined via Eq. (5.7) with a Gaussian prior on ϵ centered at 0 with width 2 $h^{-2}\,\mathrm{Mpc}^2$. The prior on the former is largely uninformative – $\alpha_a = 50\ h^{-2}\,\mathrm{Mpc}^2$ corresponds to a $\mathcal{O}(1)$ change in P_{gg} at $k = 0.2\,h\,\mathrm{Mpc}^{-1}$ (see Fig. D.3 in Appendix D.4) – while the latter is an informative prior on the $\alpha_a - \alpha_x$ relationship, where the value of 2 $h^{-2}\,\mathrm{Mpc}^2$ has been chosen to reflect uncertainties arising from `Aemulus` ν's finite grid resolution and potential baryonic feedback (see the discussion in §5.4).

For the linear theory and model independent constraints we set $b_2^L = b_s^L = 0$. We vary α_a and α_x independently with informative priors on both parameters to roughly account for the residual systematic error arising from non-linear evolution. Motivated by the discussion in §5.4 we adjust our prior such that a 1.5% variation in P_{gm} (or P_{gg}) at $k = 0.1\,h\,\mathrm{Mpc}^{-1}$ is allowed within one (prior) sigma, corresponding to a 3 $h^{-2}\,\mathrm{Mpc}^2$ width.

The priors discussed above are physically plausible (for the LRG sample [641]) while at the same time sufficiently restrictive to mitigate "volume effects" (see §5.5) arising from unconstrained degeneracy directions, which complicate the interpretation of marginal posteriors. In part due to the rise in EFT-like models being applied to cosmological datasets (for which there are dozens of poorly constrained nuisance parameters and hence sizable "volume effects"), the discussion of priors has recently become an active topic in the literature [639, 271, 196, 467, 105, 389]. Previous work has found that a Jeffreys prior [313], or "partial" Jeffreys prior [271, 196] have been effective in reducing volume effects. Others have suggested imposing a "perturbativity prior" [105] on higher-order corrections to penalize regions of parameter space for which higher-order corrections are non-perturbative. We choose not to explore these approaches in this chapter. However we comment on the expected size of volume effects in §5.5.

Code checks

As a check of our Limber code we replaced the 3D power spectrum predictions in Eq. (5.9) with a linear bias model where $P_{mm}(k,z)$ is calculated using `CLASS`' [97] default version of `HaloFit` [576]. When integrating over χ we include the redshift evolution of the power spectrum and set $b(z) = 2$, $s_\mu(z) = 1$ and dN/dz to the measured redshift distribution of the third (z_3) LRG photo-z bin. We found that our predicted C_ℓ^{gg} and $C_\ell^{\kappa g}$ agreed with `CCL`'s [146] `angular_cl` method to within $\simeq 0.3\%$, which is of the same order as difference between

`CLASS`' default version of `HaloFit` and `CCL`'s default non-linear matter power spectrum. We repeated this exercise with $s_\mu(z) = 0.4$ to explicitly test our calculation of the magnification contribution and found similar agreement.

As an additional check, we compared our (fiducial) C_ℓ's predicted with the `Aemulus` ν emulator to those predicted using `velocileptors` [136] and found excellent $\mathcal{O}(0.1\%)$ agreement (i.e. of order the emulator error) in the $\ell \to 0$ limit (more precisely, $\ell < 100$) for a broad range of nuisance term values, ensuring that our definitions of e.g. b_2^L are consistent with those used in previous works [625], and that our summation of power spectrum monomials (e.g. Eq. 5.4) has been performed self-consistently.

"Volume effect" estimation

The high accuracy and flexibility of our fiducial HEFT model comes at the cost of introducing 16 additional nuisance parameters in our fiducial setup: b_2^L, b_s^L, α_a, α_x in each of the four redshift bins. Given the large dimensionality of our posterior (30 parameters when analyzing all four bins) and the highly-compressed nature of our dataset (power spectra of projected LSS tracers), one may naturally worry that our analysis is susceptible to "volume effects": shifts in the marginalized posteriors of cosmological parameters away from their values at the maximum *a posteriori* (MAP) sourced by poorly-constrained degeneracy directions with other parameters that cover a large volume of parameter space [271, 389].

In particular, volume effects were present in the previous PR3 cross-correlation analysis [625] for which the S_8 mean (0.725) was roughly 1σ lower than the best fit value. As discussed in [625], the primary culprit for these volume effects was a degeneracy with the counterterm α_x, which was treated as a separate parameter from α_a and marginalized over with a broad and largely-uninformative prior. Here we adopt a prior relating α_a and α_x (Eq. 5.7). This prior, along with the extended dynamic range that HEFT affords over "pure" perturbation theory (increasing $\ell_{\rm max}$ to 600) efficiently mitigate the impact of volume effects on our analysis.

To estimate the size of residual volume effects we fit to a noiseless model prediction. We generate this prediction with our fiducial HEFT model (§5.4) using the measured redshift distributions and magnification biases of the LRGs. We adopt the Buzzard cosmology (§5.5) in our prediction and adjust the remaining nuisance parameters to roughly match the data[21]. Finally we "convolve" the model predictions (see §5.4) with the pixel window function and bin them into bandpowers using the same binning as for the data (§5.5). Fig. 5.5 summarizes our results when fitting to the model prediction using our fiducial analysis choices (HEFT with scales cuts and priors listed in Tables 5.2 and 5.3 respectively). We consider the cases of fitting to each redshift bin individually and jointly fitting to all bins (x-axis), and for each scenario consider a DR6, PR4 or joint PR4+DR6 covariance (blue, red and purple). In all scenarios, volume effects (difference between the mean and truth) are $< \sigma/3$, while for

[21]Specifically, we take $b_1^L = (0.872, 1.01, 1.19, 1.27)$, $10 \times b_2^L = (-2.37, -1.63, 0.553, 2.96)$, $10^3 \times b_s^L = (3.82, -2.03, -20.7, -22.7)$, $\alpha_a = (-2.79, -3.86, -5.71, -6.45)$, and $10^6 \times \mathrm{SN}^{\rm 2D} = (4.01, 2.24, 2.09, 2.32)$ in each redshift bin respectively, and fix $\alpha_x(z) = \alpha_a(z)/2(1 + b_1^L(z))$.

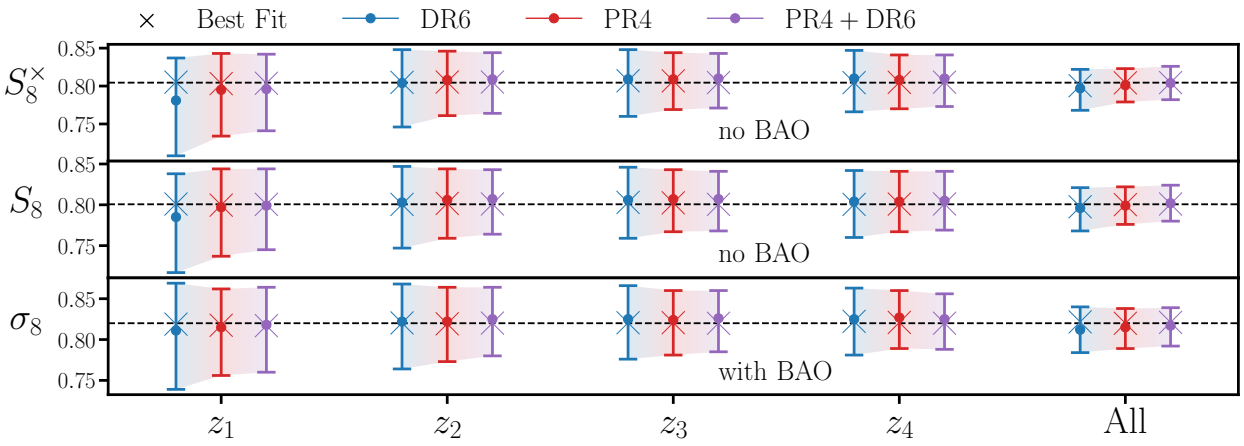


Figure 5.5: We estimate the size of "volume effects" for our fiducial analysis choices by fitting to a noiseless model prediction assuming a PR4 (blue), DR6 (red) or joint PR4+DR6 (purple) covariance. We consider fitting to individual redshift bins (z_i on the x-axis) or jointly fitting to all redshift bins (All). We show the marginal mean (circles), $\pm 1\sigma$ errors (caps) and best-fit value ($\times$'s) for $S_8^\times \equiv \sigma_8(\Omega_m/0.3)^{0.4}$ (top), S_8 (middle) and σ_8 (bottom). When quoting σ_8 constraints we include a mock BAO prior (see §5.5). The horizontal dashed lines indicate the cosmological parameters ($\Omega_m = 0.286$ and $\sigma_8 = 0.82$) used in the model prediction.

the "baseline" scenario of PR4+DR6 with all four redshift bins the volume effects are all $\leq \sigma/10$.

We additionally checked that our linear theory (§5.4) and our model-independent (§5.4) approaches yield negligible volume effects on cosmological parameters when fitting to a model prediction using a joint PR4+DR6 covariance. This is unsurprising given that these models have far fewer nuisance parameters than our fiducial HEFT model.

While the tests considered here provide an estimate for the expected "volume effect" size, they are not exhaustive in the sense that the size (and direction) of these shifts can in principle depend on the noise (or systematic) realization in the data (e.g. [145]). To quickly diagnose the presence of these shifts, we always report the location of the best-fit in addition to the mean when quoting marginal posteriors.

Fits to Buzzard mocks

We assess the accuracy of each of our modeling choices by fitting to mock data constructed from the `Buzzard` simulations [178, 180, 616]. We briefly summarize these simulations and the mock LRG sample below, and refer the reader to ref. [179] and the references above for a more detailed discussion.

The `Buzzard` DM particle lightcone is constructed from three separate N-body simulations spanning $0 < z < 0.32$, $0.32 < z < 0.84$ and $0.84 < z < 2.35$ respectively. Particles (1400^3 to 2048^3 particles depending on the redshift bin) are evolved with `Gadget-2` [564] in a box size ranging from $1.05 - 4\,h^{-1}\,\mathrm{Gpc}$, yielding an effective halo mass resolution of $0.5 \times 10^{13}\,h^{-1}M_\odot$ for $z \leq 0.32$ and $10^{13}\,h^{-1}M_\odot$ for $z \lesssim 2$. Galaxy positions, velocities, rest-frame magnitudes and SEDs are assigned with the `Addgals` algorithm [616]. The magnitudes and SEDs are then passed through DECam-like bandpasses that closely match those applied to the data, with slight modifications to better match the observed projected number density. Photometric redshifts are assigned to each galaxy by adding a Gaussian photometric error $\sigma_z(z)$ to the true redshift z, where $\sigma_z(z)$ has been calibrated with the photo-z errors measured by ref. [643]. Galaxies are then assigned to photometric redshift bins defined similarly to those presented in ref. [643]. The CMB lensing convergence map and magnification contributions are constructed within the Born approximation from the matter distribution in the simulation.

The galaxy auto- and cross-correlations with CMB lensing are binned into the same bandpowers discussed in §5.5 using an independent pipeline rather than the one detailed detailed in the companion paper [320] and §5.3. We adopt the same scale cuts as in Table 5.2 and adjust the central values of the priors on the number count slope and shot noise to the measured values in the simulations[22]. Otherwise the priors are identical to those listed in Table 5.3 for each modeling choice. We adjust the fiducial values of $n_s = 0.96$, $\Omega_b h^2 = 0.02254$, $\Omega_m h^3 = 0.098098$ and $\sum m_\nu = 0$ to those used in the Buzzard simulations. We use the measured redshift distributions of the simulated LRGs when fitting to the mock measurements.

The mock power spectra are measured from seven quarter-sky cutouts, corresponding to $\simeq 4\times$ the sky coverage of the cross-correlation with *Planck* PR4. The mock CMB lensing maps contain no lensing reconstruction noise. We analytically estimate the covariance of these mock measurement using the methods discussed in §5.3, where the power spectrum of the mock CMB lensing map is taken to be $C_\ell^{\kappa\kappa}$ without reconstruction noise and the "mask" is taken to be the full sky. We multiply the resulting covariance matrix by 4/7 to account for the Buzzard sky-coverage.

Our results are summarized in Fig. 5.6, where we show the posteriors of $S_8^\times \equiv \sigma_8(\Omega_m/0.3)^{0.4}$ (top), S_8 (middle) and σ_8 (bottom) for our fiducial HEFT (errorbars with circles) and linear theory (errorbars with squares) models, while for the model-independent constraints (right panel) we show α_8 constraints (Eq. 5.8). For both HEFT and linear theory we consider fitting to each redshift bin individually or jointly fitting to all bins (x-axis), while for the model independent approach we only consider the former (y-axis). Since we do not consider fitting to z_1 with linear theory in §5.7 we correspondingly choose not to plot z_1 linear theory constraints in Fig. 5.6. The true values of cosmological parameters are indicated by the thin black dashed lines.

In black we show the constraints when adopting a simulation-like covariance for which

[22] $s_\mu(z_i) = 1.04, 0.97, 0.81, 0.80$ and $10^6\,\mathrm{SN}^{2\mathrm{D}} = 3.7786, 1.7858, 1.9084, 2.5209$

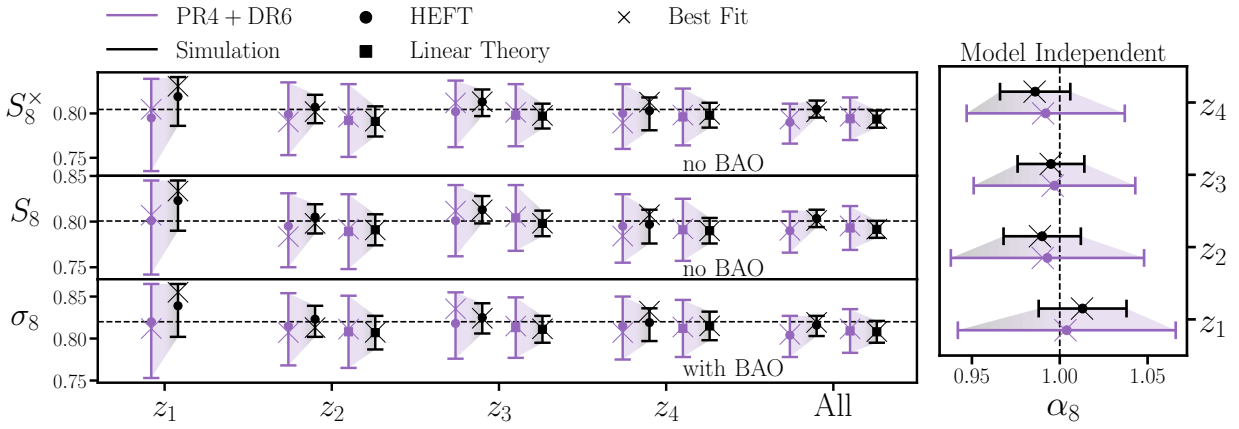


Figure 5.6: Cosmological parameter recovery when fitting to mock galaxy and CMB lensing measurements from the Buzzard simulations. In the left plot we show $S_8^\times \equiv \sigma_8(\Omega_m/0.3)^{0.4}$ (top), S_8 (middle) and σ_8 (bottom) constraints for HEFT (circles) and linear theory (squares). In purple we adopt a PR4+DR6 covariance, while in black we use a covariance representative of the Buzzard simulation. We include a mock BAO prior when quoting σ_8 constraints (see §5.5). In the right panel we show the model-independent measurements of α_8 (see Eq. 5.8). The location of the best fit is indicated with a $\times$. Dashed lines indicate the values ($\Omega_m = 0.286$, $\sigma_8 = 0.82$, $\alpha_8 = 1$) of the Buzzard cosmology.

$\sim 1\sigma$ level fluctuations in the posteriors are expected, which is consistent with the observed scatter. We show the constraints for a joint PR4+DR6 covariance (still using the same mock measurements) in purple, for which we naively expect $\sim 0.5\sigma$ fluctuations, in addition to (small, see §5.5) volume shifts. This is again consistent with the observed scatter in the mock fits. In particular, when jointly fitting to all four redshift bins using our fiducial HEFT model with a joint PR4+DR6 covariance our recovered cosmological parameters are within 0.7σ of their true values. We conclude that each of the considered models (with their fiducial scale cuts and priors) show no evidence of bias, and no signs of significant volume effects.

5.6 LRG systematics tests

In addition to the null tests presented in the companion paper [320], which primarily (but not exclusively) test for systematics in the ACT DR6 CMB lensing map, here we present an additional set of systematics tests for the LRG auto- and cross-correlation with the baseline *Planck* PR4 κ map.

A significant challenge for current and future large-scale structure surveys is to ensure the uniformity of a galaxy sample's physical properties on different regions of the sky. The selection criterion and systematics weights used for our LRG samples [641] were primarily tuned to correct for correlations in LRG density $\bar{n}(\hat{\boldsymbol{n}})$ with a set of systematic templates.

While enforcing a uniform $\bar{n}(\hat{\boldsymbol{n}})$ is a necessary condition, this does not necessarily ensure the uniformity of other physical properties, e.g. redshift distributions [54] or linear bias, which have the potential to bias cosmological inference if these effects are large and not properly taken into account. In sections 5.6, 5.6, 5.6, and 5.6 we quantify these effects by examining variations in the LRG auto- and cross-correlation with *Planck* PR4 on different footprints[23]. In §5.6 we quantify the impact of systematic weights on our measurements. Finally in §5.6 we examine the consistency of the measured galaxy cross-spectra $C_\ell^{g_i g_j}$ with those predicted from the galaxy auto- and cross-correlation with CMB lensing alone. We find significant evidence for variations in the LRGs' physical properties on different footprints (§5.6), which we expect to be our leading source of systematic error. In §5.7 we estimate the impact of these variations on cosmological parameters, finding at most $\simeq 0.2$ shifts to our fiducial S_8 constraints (some of which is statistical).

Imaging footprints

We first consider variations with the different imaging footprints (North, DECaLS and DES), which are shown in Fig. 5.2. The North mask is defined by the intersection of DEC $> 32.375\circ$ and the NGC. The (binary) DES mask is defined to be non-zero where there is a positive detection fraction in any of the DES DR2 *grizY* bands. The DECaLS mask is defined to be non-zero everywhere except North and DES, with an additional cut on declination DEC $> -15°$. Finally we multiply each imaging region mask by the fiducial LRG mask, producing binary masks for each region that are subsets of the full LRG footprint.

We (re)mask the LRG samples with each of the imaging region masks and measure their auto- and cross-correlation with PR4 using the methods discussed in §5.3. We approximate the covariance of these measurements using the analytic methods described in §5.3, from which we compute the variance of the difference: $\mathrm{Var}[X-Y] = \mathrm{Var}[X]+\mathrm{Var}[Y]-2\mathrm{Cov}[X,Y]$.

Our results are summarized in Fig. 5.7, where we show the relative variation in C_ℓ^{gg} (first and third row) and $C_\ell^{\kappa g}$ (second and fourth row). In blue we show the relative difference between DECaLS and DES, while in orange we plot the difference between North and DES. We find that C_ℓ^{gg} is consistent across the DECaLS and North regions over our fiducial analysis range. However, we find significant variations between the DES footprint and elsewhere for all redshift bins except z_3 (i.e. the blue and orange points differ from zero in the same way). For z_1 this difference is primarily relevant on small scales and is thus likely due to mismatched shot noise. For z_2 and z_4 this difference is also relevant on large scales ($\ell \simeq 150$). For comparison we show the expected variation from changes in the redshift distribution alone (solid lines), which are not large enough to explain the observed discrepancies (especially for z_4). We hypothesize that the LRGs in the DES region have a slightly larger linear bias. This could potentially be the byproduct of the DES region's deeper $z-$band depth, which z_4 is

[23]Some of the footprints considered here (e.g. the Northern imaging region) have either minimal or no overlap with the ACT lensing footprint, which is why we only choose to examine variations in the PR4 cross-correlation. The robustness of the DR6 cross-correlation to additional footprint variations is explored in the companion paper [320].

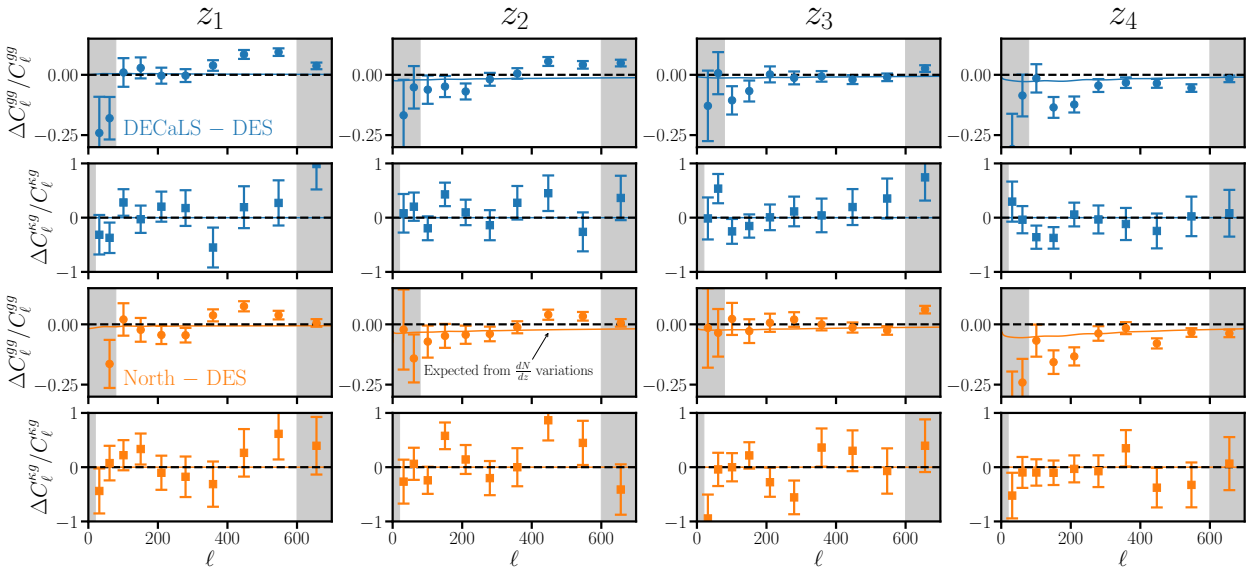


Figure 5.7: The relative variation in LRG power spectra (first and third rows) and their cross-correlation with PR4 (second and fourth rows) across different imaging footprints. We divide by the (binned) C_ℓ's used for covariance estimation (see §5.3). Columns indicate the redshift bin. In blue we show the difference between DECaLS and DES, while in orange we show the difference between North and DES. The gray regions are excluded from the fitting range. The PTEs for the cross-correlation measurements are sensible: 0.58, 0.44, 0.62, 0.55 for DECaLS−DES and 0.73, 0.11, 0.25, 0.84 for North−DES. The PTEs for the auto-correlation measurements are: 5.1×10^{-11}, 1.1×10^{-3}, 0.33, 5.1×10^{-8} for DECaLS−DES and 1.1×10^{-3}, 0.075, 0.85 and 9.7×10^{-8} for North−DES. The small PTEs indicate significant variations in the LRG sample's statistical properties on different imaging footprints. The solid (blue and orange) lines show the expected variation if the galaxies were drawn from the same 3D statistical distribution (taken to be the fiducial best fit to PR4+DR6) but projected to 2D using the different redshift distributions shown in Fig. 5.2.

most sensitive to. This is reflected in the systematics weight maps, where the DES footprint is clearly visible (see Fig. 8 of [643]). Lower photometric noise in the DES region would make it less likely for fainter objects to scatter into the sample, resulting in a net increase in the DES region's linear bias. Alternatively, it may be that the redshift distribution in the southern half of the DES footprint (DEC $< -15°$) differs from that on the footprint where we have direct spectroscopic calibration of the redshift distribution (see §5.6). We have repeated this exercise when additionally applying a stricter $E(B-V) \leq 0.05$ and stellar-density (< 1500 per square degree) cuts and found qualitatively similar results, suggesting that this behavior is not sourced by Galactic contamination.

We note, however, that the variations in the cross-correlation with *Planck* PR4 are con-

sistent with fluctuations (see the caption of Fig. 5.7 for the PTEs) but note that the SNR of the cross-correlation is low relative to the galaxy auto. Moreover, since our constraint on S_8 is not proportional to C_ℓ^{gg} but instead depends essentially on the ratio $C_\ell^{\kappa g}/\sqrt{C_\ell^{gg}}$, biases to S_8 (or σ_8) may be significantly smaller than those impacting C_ℓ^{gg}. For example, if a LSS tracer on two disjoint regions of the sky (with sky coverage $f_{\rm sky,\,1}$ and $f_{\rm sky,\,2}$) has linear bias b_1 and b_2 respectively on each region, then the galaxy-auto spectrum is sensitive to the effective bias $b^2_{\rm eff,\,a} = \sum_i f_{\rm sky,\,i} b_i^2 / \sum_i f_{\rm sky,\,i}$ while the cross-correlation with matter (CMB lensing) probes $b_{\rm eff,\,x} = \sum_i f_{\rm sky,\,i} b_i / \sum_i f_{\rm sky,\,i}$. In the limit that the bias variation is small, say $b_2 = b_1(1+\epsilon)$, the ratio $b_{\rm eff,\,x}/\sqrt{b^2_{\rm eff,\,a}} = 1 + \mathcal{O}(\epsilon^2)$. Thus the expected bias to S_8 is of order ϵ^2. In Fig. 5.7 we see $\sim 10\%$ changes to C^{gg} between North and DES for z_4, suggesting a $\sim 5\%$ change in the linear bias, for which the corresponding bias to S_8 is a quarter of a percent which is significantly smaller than our statistical errorbars (roughly 2.5%, see §5.7).

We explore the impact of these variations at the parameter-level in §5.7 where we treat the North, DECaLS and DES regions as individual samples with their own corresponding nuisance parameters, and using the redshift distributions calibrated for each individual footprint. This prescription mitigates bias to S_8 from variations in linear bias or redshift distributions across footprints[24]. We find consistent results when treating the entire LRG footprint as a single sample as we do when treating each imaging region as its own sample. For simplicity we adopt the former as the fiducial scenario in §5.7.

Northern and Southern galactic caps

Here we test for variations in the LRG clustering in the Northern and Southern galactic hemispheres. Since we have already established (§5.6) that there are significant deviations in C_ℓ^{gg} across different imaging regions (i.e. DES vs elsewhere), a strict comparison of the intersection of our full LRG mask with the NGC and SGC would yield qualitatively similar results to §5.6. To isolate the impact of NGC vs SGC, we instead choose to compare the LRG clustering on the northern and southern regions of the DECaLS footprint alone, which we show in Fig. 5.8. The PTEs (listed in the caption) of these differences are all reasonable, suggesting that variations across different galactic caps are consistent with fluctuations. The same is true for the cross-correlation with PR4 (and DR6, see the companion paper [320]), which we have opted not to plot in Fig. 5.8.

Above and below DEC $= -15°$

The LRG redshift distribution has been directly calibrated with DESI spectroscopy on a subset of the full (imaging) footprint that approximately corresponds to DEC $> -15°$ (see e.g. Fig. 10 of [643]). When treating the full LRG footprint as a single sample we implicitly assume that the redshift distribution below DEC $< -15°$ (or other physical properties of

[24] This isn't necessarily true for variations (in linear bias or redshift distribution) within each footprint [54]. We expect these to be negligibly small for our analysis.

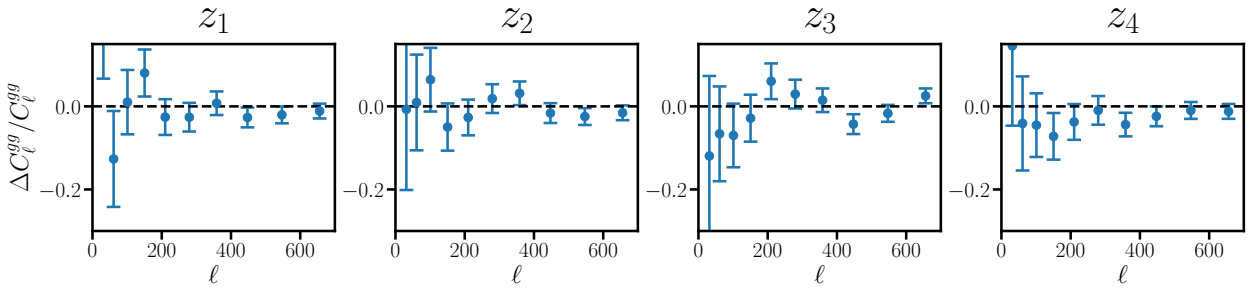


Figure 5.8: The relative difference in LRG clustering between the NGC and SGC regions of the DECaLS footprint. We divide by the (binned) C_ℓ's used for covariance estimation (see §5.3). The gray regions indicate our fitting range. The PTEs are 0.62, 0.63, 0.32 and 0.43 respectively.

the sample) is consistent with that measured on the calibrated subset. Here we empirically test for variations above and below DEC $= -15°$ by comparing our fiducial auto- and PR4 cross-correlation measurements with those obtained when additionally masking pixels with DEC $< -15°$. These results are summarized in Fig. 5.9. As in §5.6 we find that the variations in the PR4 cross-correlation are consistent with statistical fluctuations (see the caption of Fig. 5.9 for PTEs). The variations to the galaxy auto-correlation are statistically significant, however the (relative) magnitude of these variations are milder than found in §5.6 (at most 5%). We explore the impact of these variations at the parameter-level in §5.7 and find statistically consistent results with and without masking the DEC $< -15°$ region.

Stricter extinction and stellar density cuts

To check for potential signs of galactic contamination we consider making stricter cuts on extinction and stellar density. These results are shown in Fig. 5.10. In blue we show the relative variation in C_ℓ^{gg} when we additionally mask regions with $E(B-V) \geq 0.05$ using the $E(B-V)$ map from [529]. In orange we show the variations when masking regions where the stellar density exceeds 1500 deg^{-2} using the map from ref. [429]. In both scenarios we find that the galaxy power spectra vary by $\simeq 1\%$ on the scales of interest. We do not consider estimating the covariance of the measurements here as they are highly correlated, making our analytic estimation to the covariance matrix considerably less accurate. Regardless of the PTEs associated with these variations, their impact on our cosmological results would be negligible.

Impact of systematic weights

The systematics weights published by ref. [643] are designed to remove spatial trends in the mean galaxy density with a set of seven templates: the SFD $E(B-V)$ map [529] in

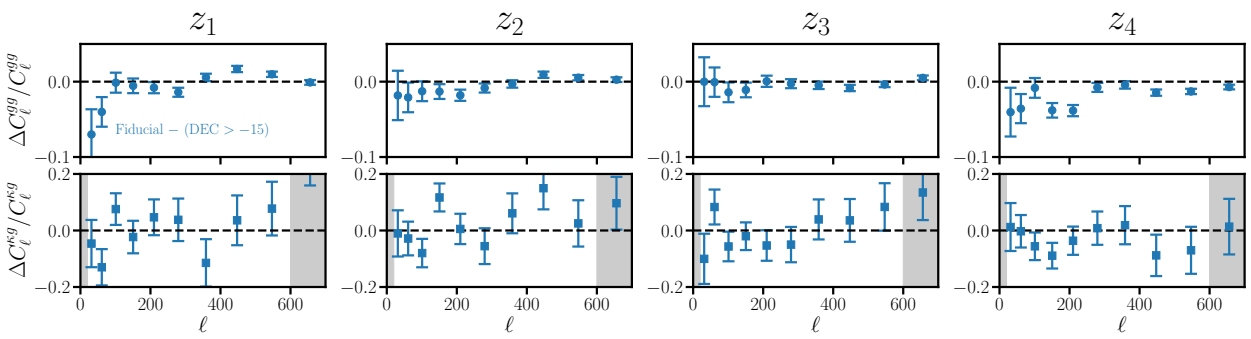


Figure 5.9: The relative difference in our fiducial LRG power spectra (first row) and their cross-correlation with PR4 (second row) with those measured with an additional DEC $> -15°$ mask applied. We divide by the (binned) C_ℓ's used for covariance estimation (see §5.3). The gray regions indicate our fitting range. The PTEs for the cross-correlations are 0.36, 0.13, 0.59 and 0.53 respectively. With the exception of z_3, the PTEs for the auto-correlations are clearly inconsistent with fluctuations: 2.7×10^{-5}, 9.2×10^{-5}, 0.24 and 2.5×10^{-13} respectively.

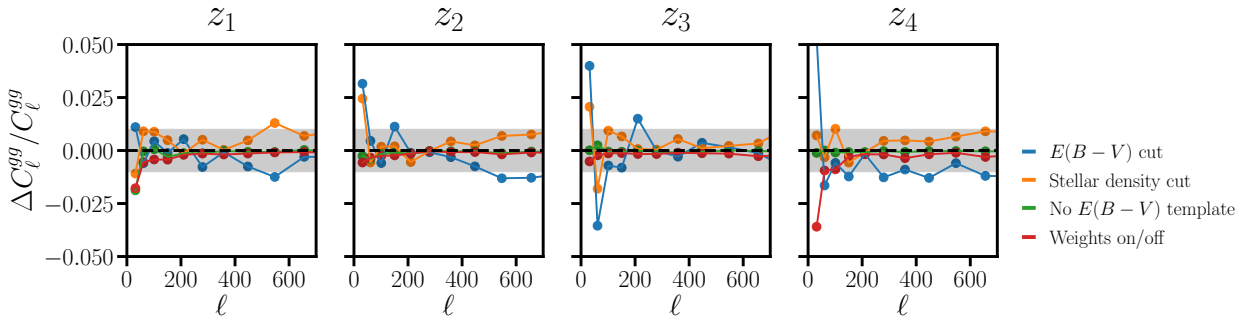


Figure 5.10: The relative change in the LRG auto spectra when adopting a stricter $E(B-V)$ cut (from 0.15 to 0.05; blue), adopting a stricter stellar density cut (2500 to 1500 deg^{-2}; orange), removing the SFD $E(B-V)$ map [529] as a template when constructing systematics weights (green), and neglecting the systematics weights altogether (red). Different panels indicate the redshift bin. The gray vertical regions indicate our fitting range, while the gray horizontal band indicates the interval $[-0.01, 0.01]$.

addition to seeing and depth in each of the three optical (g, r, z) bands. Here we quantify the impact of these weights on the galaxy auto-spectra, and in addition quantify the expected systematic error arising from residual observational trends.

Ref. [643] provides weights that only use seeing and depth as templates. We show the variations in the galaxy power spectra with and without including $E(B-V)$ as a systematics template in Fig. 5.10 (green points), while the red points in Fig. 5.10 show the impact of

neglecting the systematics weights completely. These variations are at most a percent on the scales of interest, thus even modest (e.g. $\sim 10\%$) errors in the systematics weight map would be expected to only yield $\sim 0.1\%$ changes to our fiducial power spectra. As the power spectra and masks of these sets of maps are nearly identical, it is unclear how one would estimate the covariance between them (even numerically). For this reason we do not attempt to assign errorbars to this plot. Regardless of the statistical significance of these deviations, they are small enough that they would have no significant impact on our results.

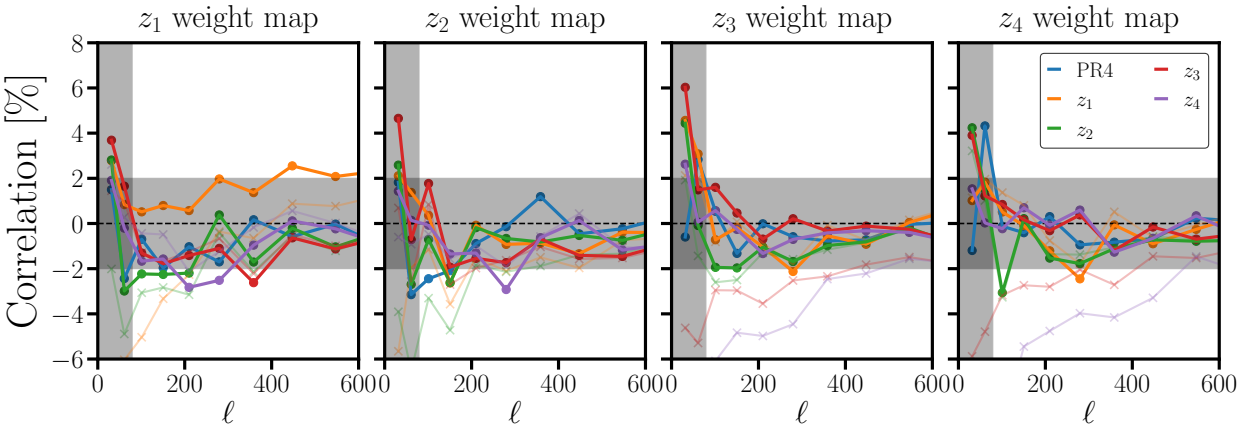


Figure 5.11: The correlation ($\equiv C_\ell^{ab}/\sqrt{C_\ell^{aa}C_\ell^{bb}}$) between the fiducial systematics weight maps (panels) and galaxy density contrast maps (colored lines and filled circles) after the systematics weights have been applied. We additionally show the correlations between the weight maps and *Planck* PR4 convergence map (blue). The correlations between the weight maps and galaxy maps without any systematics weights applied are shown by the pale lines with $\times$ markers. The gray vertical region indicates our fitting range for the galaxy auto-spectra. The horizontal gray band corresponds to the interval $[-0.02, 0.02]$, which is of the same order as the observed point-to-point scatter in the measured correlations.

In the ideal limit that systematics weight maps perfectly remove observational trends in the mean galaxy density, correlations between the weight maps and the (systematics-corrected) galaxy density contrast maps should vanish. We plot correlations between the systematic weight maps (panels) and the galaxy samples (colors) in Fig. 5.11. On the scales relevant for our analysis we find $\simeq 1-2\%$ residual correlations between the weight maps and galaxy samples. In Fig. 5.11 we also show the correlations with the galaxy samples before the systematics weight were applied (pale lines with $\times$'s). On large scales ($100 \lesssim \ell \lesssim 200$) we find that applying systematics weights reduces the correlation between the galaxy samples and their respective weights maps by $\sim 3-6\%$. As discussed previously (Fig. 5.10), this $\sim 3-6\%$ reduction in correlation results in subpercent changes to the galaxy auto-spectra. Extrapolating this trend, we expect the residual $\simeq 1-2\%$ correlations with the weight maps to bias our galaxy auto-spectra by less than a third of a percent on the scales of interest.

Galaxy cross-spectra

As a consistency check we compare the measured galaxy cross-spectra to predictions from our fiducial model, fitted from the galaxy auto-spectra and cross-correlations with CMB lensing. The galaxy cross-spectra could in principle be used to better constrain number count slopes and improve the degeneracy-breaking between cosmological and bias parameters. However, given the heightened sensitivity of these cross-spectra to potential mischaracterization of the tails in the redshift distributions, and that the magnitude of the systematic contamination is expected to be similar to that in the galaxy auto-spectra while the cosmological signal is $\simeq 10\times$ smaller, we choose not to fit to the galaxy cross-spectra in §5.7.

We show the measured galaxy cross-spectra in Fig. 5.12. The errorbars of these measurements are estimated following §5.3. We note that there is negligible redshift overlap between the lowest and highest redshift bins, however $C_\ell^{g_1 g_4}$ is non-zero (primarily) due to correlations of z_1 galaxies with the matter magnifying the z_4 sample. We detect a non-zero value of $C_\ell^{g_1 g_4}$, and hence a detection of magnification, with a statistical significance of 8σ on the fiducial scale range ($79 \leq \ell < 600$) adopted for our fits to the galaxy auto-spectra.

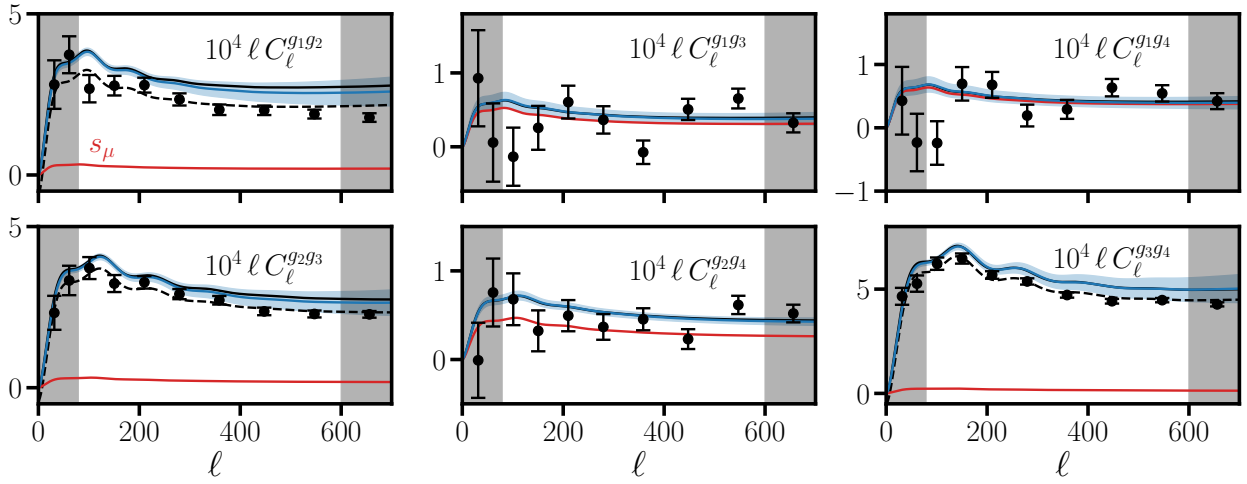


Figure 5.12: The measured galaxy cross-spectra (circles with errorbars). Black lines indicate the predicted (Eq. 5.14) cross-spectra from the best-fit to our fiducial PR4+DR6 analysis (§5.7), neglecting correlated shot noise. We also show the mean (blue line) and $\pm 1\sigma$ interval (shaded blue) associated with the PPD of that same fit (again neglecting shot noise). Note that in some panels (e.g. $C_\ell^{g_3 g_4}$) the mean and best-fit curves overlap. By contrast, the cross-correlations between neighboring bins are all low compared to the prediction. This effect is not yet fully understood, but we subtract a $1/\ell$ contribution from the best fit prediction to obtain the dashed black curves and then check the impact of marginalizing over a bias of that form in §5.7. In red we show the magnification contribution to the best fit, i.e. the last two terms in Eq. (5.14).

We predict the cross-spectra from fits to the galaxy auto- and cross-correlation with CMB lensing alone as follows. Assuming that the galaxies in the four photometric redshift bins can be thought of as subsamples of a single smoothly evolving galaxy sample, the cross-correlation between two photometric redshift bins is given by

$$C_\ell^{g_i g_j} = \int \frac{d\chi}{\chi^2} \left[W^{g_i} W^{g_j} P_{gg} + \left(W^{g_i} W^{\mu_j} + W^{g_j} W^{\mu_i} \right) P_{gm} + W^{\mu_i} W^{\mu_j} P_{mm} \right] \tag{5.14}$$

within the Limber approximation[25], where we have suppressed the explicit dependence of projection kernels on comoving distance and power spectra on wavenumber and redshift for simplicity. The projection kernels W^{g_i} and W^{μ_i} are given by Eq. (5.10) where ϕ is replaced with the redshift distribution of the appropriate sample. Given a set of nuisance parameters (for each effective redshift) we can construct a plausible evolution for $b_1^L(z)$, $b_2^L(z)$, $b_s^L(z)$, $\alpha_a(z)$, $s_\mu(z)$, and $\epsilon(z)$ by interpolating their values at each effective redshift[26]. Along with a set of cosmological parameters, this interpolation then determines the non-stochastic contribution to $C_\ell^{g_i g_j}$.

Using the best-fit parameters to our baseline PR4+DR6 analysis (§5.7) we predict $C_\ell^{g_i g_j}$ following the recipe above, which are shown as solid black lines in Fig. 5.12. We also show the mean (blue line) and $\pm 1\sigma$ (shaded blue) interval associated with the posterior predictive distribution (PPD) of the same baseline fit. This is obtained from 16384 draws of the posterior. We select 256 random samples from the chain, and generate 64 Monte-Carlo realizations of the linear parameters for each sample following the methods discussed in Appendix D.4. We find that the measured $C_\ell^{g_1 g_2}$ is $\sim 20\%$ smaller than predicted. We find smaller deficits for the remaining two cross-spectra of neighboring redshift bins: approximately 12% and 8% for $C_\ell^{g_2 g_3}$ and $C_\ell^{g_3 g_4}$ respectively. The remaining cross-spectra, whose signals all receive large (or are dominated by) magnification contributions, are all within reasonable agreement with the data[27]. As we discuss below and in §5.7, even if these deficits are the result of systematic contamination in the data we expect these contaminants to have a negligible impact on our cosmological constraints.

There are several important subtleties to bear in mind when comparing these predictions to the measured cross-spectra. First, in Eq. (5.14) we assume that the physical properties of the galaxy samples depend only on redshift, which may not necessarily be the case. For example, the evolution of the linear bias in the high-z tail of z_1 may disagree with the linear bias evolution in the low z tail of z_2, which would not be captured by Eq. (5.14). Additionally, we assume a simple linear interpolation of nuisance parameters, while the true redshift evolution may be more complex. Second, we should in principle allow for a free shot noise component for each cross-spectrum − especially for neighboring bins with overlapping redshift distributions − although Fig. 5.12 suggests that a negative shot noise

[25] Using Eq. (5.12) we estimate that beyond-Limber corrections to the galaxy cross-spectra are $\leq 2\%$ for $\ell > 79$.

[26] We use scipy's interp1d with kind='linear' and its default extrapolation for $z < 0.47$ and $z > 0.91$.

[27] There is a single point in $C^{g_1 g_3}$ at $\ell \simeq 350$ that fluctuated $\simeq 2.8\sigma$ lower than the prediction.

contribution would be required to alleviate the observed deficits. With these subtleties in mind, we caution the reader that the predictions made with Eq. (5.14) are somewhat crude and should be interpreted as such.

On the other hand, it may be the case that these deficits are the result of some form of systematic modulation in the data. First, these deficits could suggest miscalibrated magnification contributions, although we find this unlikely given that the magnification contribution for neighboring bins is small, and that the data for the remaining cross-spectra with significant magnification contributions are in reasonable agreement with our predictions. Second, the tails of the redshift distributions may be slightly miscalibrated. To get a sense for the error in the tails we swapped the fiducial redshift distributions to those calibrated on DES footprint[28] [643] and found $\simeq 2$, 15 and 8% variations in $C_\ell^{g_1g_2}$, $C_\ell^{g_2g_3}$ and $C_\ell^{g_3g_4}$ respectively. Third, the probability that a galaxy is assigned to e.g. bin 1 or bin 2 at a fixed redshift may be modulated by an observational effect (e.g. extinction) that could lead to negative correlations between neighboring redshift bins. Fourth, the deficits could indicate the presence of a systematic contaminant (not captured by the systematic weights) that is negatively correlated between neighboring redshift bins, or in the maximally pessimistic scenario, a contaminant that has added power to the galaxy auto-spectra, resulting in an overestimate of the cross-spectra.

Out of an abundance of caution, we entertain the maximally pessimistic scenario in more detail. We find that the residuals between the best-fit predictions and measured galaxy cross-spectra are well fit by a $1/\ell$ contribution (black dashed lines in Fig. 5.12) with the magnitude of the coefficients ranging from -4×10^{-5} to -6×10^{-5}. A $\Delta C_\ell^{g_ig_i} = 5 \times 10^{-5}/\ell$ absolute bias corresponds to a relative bias of at most 3% to any of the galaxy auto-spectra. In the simplistic picture where S_8 is derived from the ratio $C^{\kappa g}/\sqrt{C^{gg}}$, a relative 3% scale-invariant increase (neglecting shot noise) in C^{gg} would bias S_8 low by 1.5%, corresponding to a $\simeq 0.6\sigma$ shift for our fiducial analysis. However, this crude calculation overestimates the resulting bias to S_8 since the systematic contamination has a different scale-dependence than a rescaling of S_8, such that not all of the $\sim 3\%$ bias would be projected to cosmological constraints. We confirm this suspicion in §5.7 by adding $1/\ell$ terms to the galaxy power spectra and marginalizing over their coefficients with $\mathcal{N}(0, 10^{-4})$ priors. Doing so yields negligible (0.05σ) shifts in our fiducial linear theory S_8 constraints.

We conclude that the observed deficits are unlikely to source significant biases to our cosmological constraints. These tests highlight the powerful ability of spectroscopically-calibrated galaxy samples to self-consistently probe systematic contamination, and test for immunity to it at the cosmological parameter level (§5.7). This is in contrast to photometric samples, for which systematic contaminants in the galaxy cross-spectra are more difficult to identify in the presence of noisy redshift distributions.

[28] We repeated this using the redshift distributions calibrated on the North and DECaLS footprints but found smaller variations than for DES.

5.7 Results

Blinding

We blinded both the cross-correlation with ACT DR6 and all cosmological parameters derived from ACT data until all of the scale cuts (§5.5) and priors (§5.5) were finalized, null and systematics tests were satisfactorily passed (see the companion paper [320] as well as sections 5.6 and 5.7), and our likelihood had been tested with the `Buzzard` mocks (§5.5). We emphasize that none of our fiducial analysis choices have changed post unblinding. We defer a detailed discussion of the full blinding policy to the companion paper [320], and further note here that scale cuts and priors were chosen based on rough guidelines from previous works (e.g. [327]) and were finalized before running tests on the latest `Buzzard` mock measurements, rather than using the `Buzzard` mocks to calibrate these choices. This improves the mock tests as a cross-check for two reasons: (1) the fits to the latest mock measurements were done semi-blindly[29] and (2) calibrating e.g. priors with the *same set* of mocks used to validate the likelihood would in general reduce the scatter of the recovered cosmological parameters around their true values.

Planck PR3 reanalysis

White et al. [625] found $S_8 = 0.725 \pm 0.030$ [0.751] (including BAO) from the cross correlation of *Planck* PR3 CMB lensing with a previous version of the LRG sample considered here, where here and below the best-fit value is given in square brackets. With the latest LRG sample and the same CMB lensing map, our fiducial analysis choices yield

$$\textit{Planck}\ \text{PR3:}\ S_8 = 0.762 \pm 0.023\ [0.761] \tag{5.15}$$

from our data alone, and

$$\textit{Planck}\ \text{PR3:}\ \sigma_8 = 0.758^{+0.023}_{-0.026}\ [0.765] \tag{5.16}$$

once BAO data from 6dF [84], BOSS DR7 [504] and DR12 [28] are included. The decrease in statistical error (from 0.030 to 0.023) is the result of the informative prior on the $\alpha_a - \alpha_x$ relationship (Eq. 5.7). There are four significant changes (orange points in Fig. 5.13) in our analysis that are responsible for the $\simeq 1\sigma$ upward shift in the *mean* value of S_8:

- consistent use of neutrino masses (see below): $\Delta S_8 = +0.017$
- including a "normalization" correction to the PR3 cross-correlation: $\Delta S_8 = +0.015$
- adding an informative prior on the $\alpha_a - \alpha_x$ relationship (Eq. 5.7), which efficiently mitigates "volume effects" without shifting the best-fit significantly: $\Delta S_8 = +0.014$

[29] in the sense that our analysis choices were fixed before running on the finalized measurements and that none of these choices changed after running the mock tests

- updating the LRG sample: $\Delta S_8 = -0.012$

We provide a more thorough description of additional changes below. Their impacts on S_8 constraints are summarized in Fig. 5.13.

To smoothly connect our results with those found by ref. [625] we begin by reanalyzing the same data with nearly identical analysis choices. We use the power spectra, window functions, covariance and redshift distributions provided at this link. We adjust our scale cuts (for both the auto- and cross-correlation) to $\ell_{\rm min}(z_i) = 25$ and $\ell_{\rm max}(z_i) = 275, 325, 375, 425$ for each redshift bin respectively[30]. Following [625] we fix $n_s = 0.97$, $\omega_b = 0.022$, $\sum m_\nu = 0.06$ eV, $100\,\theta_\star = 1.04109$ and directly sample $\ln(10^{10} A_s)$ and $\Omega_c h^2$ with the priors listed in Table 5.3. We fix $b_s(z_i) = 0$, vary α_x directly (rather than ϵ) with a $\mathcal{N}(0, 50)$ prior and adjust the central value of our priors for shot noise and magnification bias to those listed in Table 1 of [625]. Priors on the remaining nuisance parameters are the same as those listed in Table 5.3. Following [625] we include supernovae [537] and BAO [84, 504, 28] data by default as an effective prior on Ω_m. With these analysis choices we find $S_8 = 0.742^{+0.026}_{-0.029}$ [0.755] (second point in Fig. 5.13). We have repeated this exercise with $\Omega_m h^3 = 0.09633$ fixed instead of $\theta_\star$, and using `velocileptors` instead of Aemulus ν for the power spectrum predictions and found negligible changes: $S_8 = 0.743 \pm 0.028$ [0.758] and $S_8 = 0.744 \pm 0.028$ [0.757] respectively (third and fourth points in Fig. 5.13). In what follows we always fix $\Omega_m h^3$ instead of $\theta_\star$, and continue to use the HEFT emulator rather than `velocileptors`.

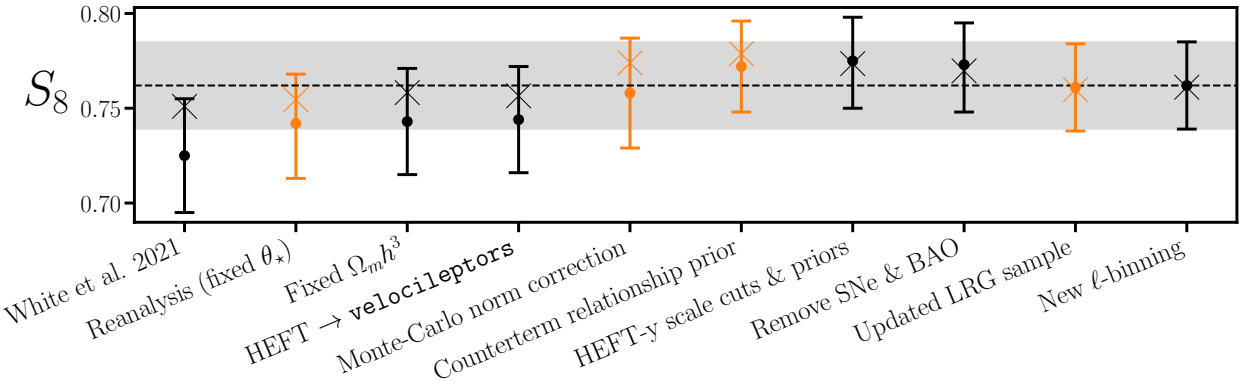


Figure 5.13: Connecting the dots between the White et al. [625] S_8 constraints (far left) and those obtained from our fiducial PR3 reanalysis (far right). Highlighted in orange are the four significant changes in the analysis that are predominantly responsible for the shift in S_8. The shaded bands correspond to the $\pm 1\sigma$ region of our fiducial PR3 constraints. Best-fit values are indicated by $\times$'s. See the text in §5.7 for a detailed discussion of the analysis changes.

We note that the S_8 mean of our reanalysis is $\Delta S_8 = +0.017$ larger ($\sim 0.5\sigma$) than that found by ref. [625]. This difference is attributed to an inconsistent use of neutrino

[30] We note that ref. [625] quotes $\ell_{\rm max}$'s in terms of ℓ bin centers, whereas here we refer to bin edges.

masses. The 3D power spectrum emulators used by [625] did not include massive neutrinos, while elsewhere in the likelihood (i.e. for computing background quantities and σ_8) the neutrino mass was set to 0.06 eV. The lack of massive neutrinos in the emulators results in "extra power" that is compensated for by lowering S_8. The size of the resulting bias can be estimated as $\Delta S_8 = 0.5\Delta P(k)/P(k) \simeq -4f_\nu$ [299, 337, 338], which for $\Omega_m = 0.3$, $h = 0.67$ and $\sum m_\nu = 0.06$ eV gives $\Delta S_8 = -0.019$, in good agreement with the observed offset.

The PR3 cross-correlation used in [625] did not include a Monte Carlo "normalization" correction [223]. We compute this correction following Appendix D.3. Applying this to the White et al. data increases S_8 by approximately 0.5σ to $S_8 = 0.758 \pm 0.029$ [0.774] (fifth point in Fig. 5.13). Next we consider imposing the same prior on the $\alpha_a - \alpha_x$ relationship (Eq. 5.7) as used in our fiducial analysis, which gives $S_8 = 0.772 \pm 0.024$ [0.779] (sixth point in Fig. 5.13). We attribute the $\Delta S_8 = +0.014$ shift in S_8 mean to the mitigation of "volume effects", and note that we observe a significantly smaller (0.005) shift in the best-fit value. Next we consider updating the scale cuts and priors on higher order biases to approximately match those used in our fiducial analysis. We adjust the values of fixed cosmological parameters to those listed in §5.5 and adopt the fiducial priors listed in Table 5.3, with the exception of shot noise and magnification whose central values are still centered on those listed in Table 1 of [625]. We set $\ell^{gg}_{\rm min}(z_i) = 75$, $\ell^{\kappa g}_{\rm min}(z_i) = 25$ and $\ell_{\rm max}(z_i) = 575$ for both the galaxy auto- and cross-correlation. Applying these changes gives $S_8 = 0.775^{+0.023}_{-0.025}$ [0.774] (seventh point in Fig. 5.13). This is in very good agreement with the previous result and suggests that variations in scale cuts and priors on higher order bias parameters have an insignificant impact on our S_8 constraints. We explore this in more detail in §5.7. We next consider removing the supernovae and BAO priors (eighth point in Fig. 5.13), which has a negligible impact on our results: $S_8 = 0.773^{+0.022}_{-0.025}$ [0.770].

We remeasured the cross-correlation of PR3 with the latest version of our LRG sample [643] using the same ℓ-binning scheme as in [625]. We use the redshift distributions provided by [643] and recompute the window functions and covariance matrix following §5.3. The new data prefer a slightly lower value of S_8 when adopting the same scale cuts as in the previous paragraph: $S_8 = 0.761 \pm 0.023$ [0.760]. Finally, we switch from linear $\Delta\ell = 50$ bins to the $\sqrt{\ell}$ bins discussed in §5.2 and adjust our scale cuts to match those in Table 5.2, such that the analysis choices now exactly match those used in our fiducial PR4 analysis below. Updating the ℓ binning and scale cuts has a negligible impact on our constraints: $S_8 = 0.762 \pm 0.023$ [0.761].

Planck PR4, ACT DR6 and their combination

In this section we present results using our fiducial HEFT model (§5.4), scale cuts (§5.5, Table 5.2) and priors (§5.5, Table 5.3). In Fig. 5.14 we show the cosmological constraints obtained when fitting to each redshift bin individually (without BAO), considering separately the cases of cross-correlating with *Planck* PR4 (left) and ACT DR6 (right). The $\Omega_m - \sigma_8$ contours from each redshift bin are all in agreement to well within 1σ, as apparent in Fig. 5.14.

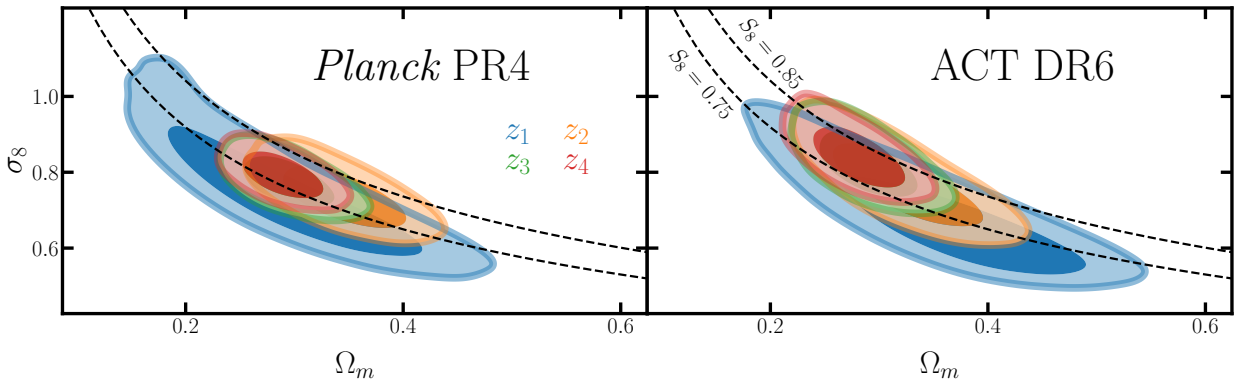


Figure 5.14: Fiducial cosmological constraints (without BAO) when analyzing each redshift bin separately (colors), both for *Planck* PR4 (left) and ACT DR6 (right). The dashed lines correspond to $S_8 = 0.75$ and 0.85.

When jointly analyzing all four galaxy auto-spectra and their cross-correlations with *Planck* PR4 we find

$$\textit{Planck}\ \text{PR4:}\ S_8 = 0.765 \pm 0.023\ [0.764], \tag{5.17}$$

from our measurements alone (i.e. no BAO), where the best-fit value is quoted in square brackets. We note that this is in very good agreement with the *Planck* PR3 result (Eq. 5.15), and that while the PR4 lensing reconstruction has improved significantly over PR3, we interestingly do not find a corresponding improvement in our S_8 constraint. In combination with BAO data from 6dF [84] SDSS DR7 [504] and DR12 [28], we are able to break the $\Omega_m - \sigma_8$ degeneracy, yielding

$$\textit{Planck}\ \text{PR4:}\ \sigma_8 = 0.762 \pm 0.023\ [0.765]. \tag{5.18}$$

We find a slight improvement in the PR4 σ_8 constraint over PR3 (cf. Eq. 5.16).

When jointly analyzing the four galaxy auto-spectra and their cross-correlations with ACT DR6 (see Fig. 5.3) we find

$$\text{ACT DR6:}\ S_8 = 0.790^{+0.024}_{-0.027}\ [0.789] \tag{5.19}$$

while with the addition of BAO data we find

$$\text{ACT DR6:}\ \sigma_8 = 0.787^{+0.025}_{-0.029}\ [0.794] \tag{5.20}$$

We note that the DR6 constraints are weaker than PR4 despite the lower reconstruction noise of the DR6 map due to the larger overlap between the PR4 footprint and LRGs.

We show the PR4 and DR6 $\sigma_8 - \Omega_m$ contours (without BAO) in the left panel of Fig. 5.15. The difference between the ACT DR6 and Planck PR4 S_8 constraint is $\Delta S_8 = 0.025$. As

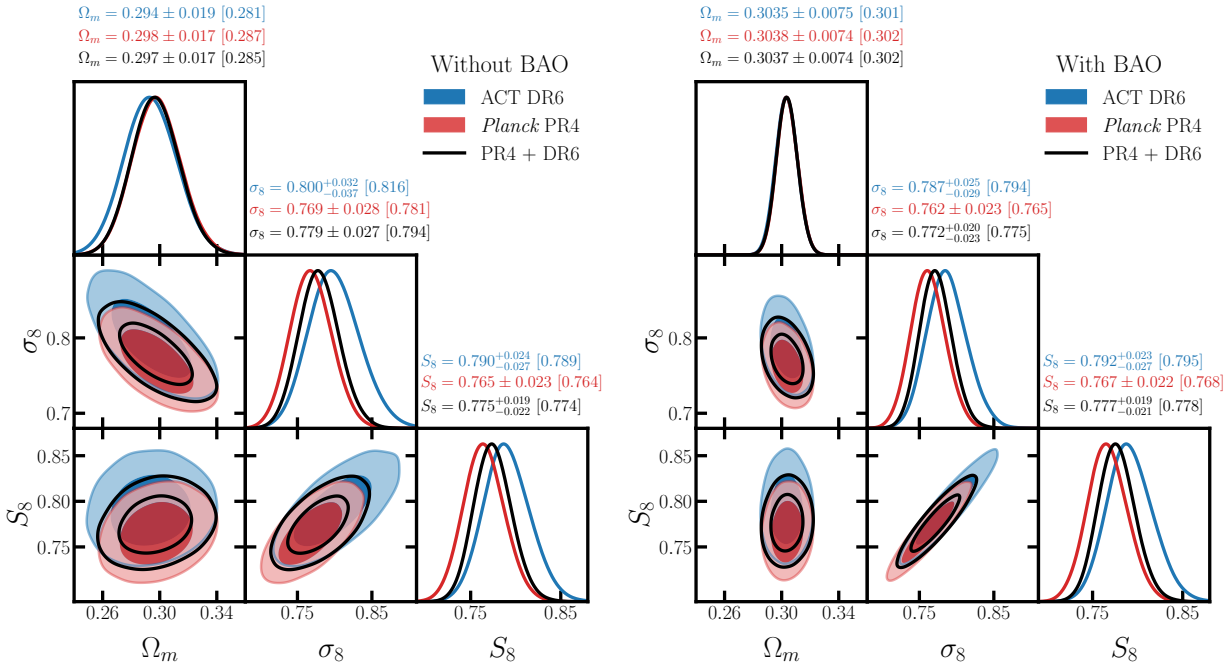


Figure 5.15: *Left:* Cosmological constraints when adopting our fiducial analysis choices. We jointly analyze the four redshift bins and their cross-correlations with ACT DR6 (blue), *Planck* PR4 (red) and their combination (black). *Right:* Same as before, but now with BAO data included [84, 504, 28].

shown in Fig. 6 of the companion paper [320], the PR4 and DR6 cross-correlation measurements are correlated by at most $\simeq 40\%$ on the scales ($100 \lesssim \ell \lesssim 300$) where most of the S_8 constraint is derived. Within linear theory, where the constraint on S_8 is schematically driven by the ratio $C_\ell^{\kappa g}/\sqrt{C_\ell^{gg}}$, we expect the correlation coefficient of the PR4 and DR6 S_8 measurements to be nearly identical to the correlations of the associated cross-spectra. Taking $r = 0.4$, $\sigma_{\rm PR4} = 0.023$ and $\sigma_{\rm DR6} = 0.0255$ we estimate the error on ΔS_8 to be $[\sigma^2_{\rm PR4} + \sigma^2_{\rm DR6} - 2r\sigma_{\rm PR4}\sigma_{\rm DR6}]^{1/2} \simeq 0.027$. This estimate suggests that the PR4 and DR6 S_8 constraints are consistent with one another to within 1σ, as visually suggested by Fig. 5.15.

Having established the consistency of the *Planck* PR4 and ACT DR6 measurements we consider jointly analyzing the two datasets. When analyzing each redshift bin independently we find

$$\begin{aligned} S_8 &= 0.727^{+0.037}_{-0.054}\ [0.726],\ 0.787^{+0.036}_{-0.043}\ [0.780],\ 0.783^{+0.033}_{-0.039}\ [0.775],\ 0.789 \pm 0.032\ [0.789] \\ \sigma_8 &= 0.717^{+0.039}_{-0.051}\ [0.717],\ 0.776^{+0.035}_{-0.047}\ [0.769],\ 0.781^{+0.032}_{-0.038}\ [0.779],\ 0.790 \pm 0.030\ [0.791] \end{aligned} \tag{5.21}$$

without and with the addition of BAO data for S_8 and σ_8 respectively. When jointly fitting the four galaxy auto-spectra, four cross-correlations with PR4, and four cross-correlations

with DR6 we find

$$\boxed{\text{Combined: } S_8 = 0.775^{+0.019}_{-0.022}\ [0.774]} \tag{5.22}$$

without any additional BAO information and

$$\boxed{\text{Combined: } \sigma_8 = 0.772^{+0.020}_{-0.023}\ [0.775]} \tag{5.23}$$

once BAO data are included. We find $\sim 10\%$ improvements to both the S_8 and σ_8 errorbars over *Planck* PR4 alone from the addition of ACT DR6.

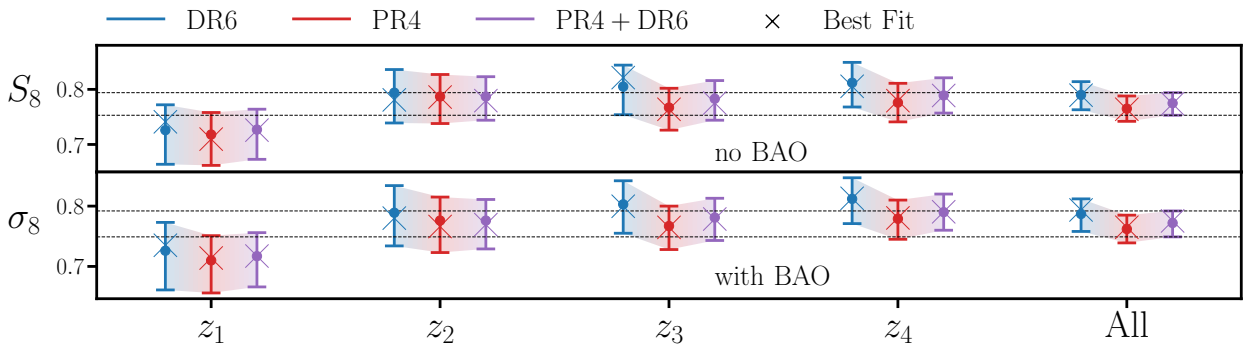


Figure 5.16: Marginal S_8 (top) and σ_8 (bottom) constraints when fitting to each redshift bin individually, or jointly fitting to all redshift bins. We show the constraints from PR4 in blue, DR6 in red and their combination in purple. The location of the best-fit value is indicated with $\times$'s. The dashed black curves correspond to the $\pm 1\sigma$ interval associated with our joint PR4 + DR6 fit (i.e. Eqs. 5.22 and 5.23).

We summarize our fiducial results in Fig. 5.16, where we show the marginal S_8 (without BAO) and σ_8 (with BAO) constraints for PR4, DR6 and their combination (in blue, red and purple respectively) for each redshift bin separately, and for the combination of all redshift bins (x-axis). We note that the DR6 values (blue) are consistently larger than the PR4 (red) results. This is qualitatively consistent with the results found in the DR6 auto-spectrum analysis [377, 493] and the cross-correlation with unWISE galaxies [223]. As in [625] we find that z_1 prefers a slightly lower S_8 (and σ_8) value than the higher redshift samples, however, we have yet to find evidence (see §5.7 for parameter-based tests) that this preference is driven by systematics.

Correlations between $S_8(z)$ measurements

To estimate correlations between the $S_8(z)$ constraints listed in Eq. (5.21) we reanalyzed our joint PR4+DR6 dataset (including all four redshift bins) with an expanded set of cosmological parameters: $\Omega_c h^2$ and four $A_s(z_i)$'s that set the primordial power spectrum amplitude in each redshift bin. We sample each $\ln(A_s(z_i))$ directly with a uniform $\mathcal{U}(2, 4)$ prior. We define

$S_8(z_i)$ as the ΛCDM-derived value of S_8 given $A_s(z_i)$ and Ω_c (in addition to the remaining fixed cosmological parameters), serving as a proxy for the value of S_8 derived from the i'th redshift bin alone. Following this procedure we obtain

$$S_8(z_i) = 0.728^{+0.038}_{-0.051},\ 0.785^{+0.034}_{-0.048},\ 0.784^{+0.031}_{-0.038},\ 0.793 \pm 0.031, \tag{5.24}$$

which is in excellent agreement with the S_8 constraints derived from each redshift bin (cf. Eq. 5.21). From the posterior we obtain a covariance matrix for the $S_8(z_i)$'s from which we estimate 9, 13 and 28% correlations between $S_8(z_1) - S_8(z_2)$, $S_8(z_2) - S_8(z_3)$ and $S_8(z_3) - S_8(z_4)$ respectively. The remaining correlations (e.g. $z_1 - z_3$) are all $\leq 3\%$. In the limit that S_8 is derived from the ratio $C_\ell^{\kappa g}/\sqrt{C_\ell^{gg}}$ on linear scales, one would expect the correlations between $S_8(z_i)$ and $S_8(z_j)$ to approximately be equal to the correlation between $C_\ell^{\kappa g_i}$ and $C_\ell^{\kappa g_j}$ at $\ell \sim 100$. As illustrated in Fig. 6 of the companion paper [320] this is indeed the case: there are $\sim$ 15, 15, and 25% correlations between $C_\ell^{\kappa g_1} - C_\ell^{\kappa g_2}$, $C_\ell^{\kappa g_2} - C_\ell^{\kappa g_3}$ and $C_\ell^{\kappa g_3} - C_\ell^{\kappa g_4}$ respectively.

To connect the $S_8(z_i)$ constraints to our combined S_8 measurement (Eq. 5.22) we define the following derived quantity, which is an optimal linear combination of the individual amplitudes: $\hat{S}_8 \equiv \sum_{ij} \mathtt{C}^{-1}_{ij} S_8(z_j) / \sum_{ij} \mathtt{C}^{-1}_{ij}$, where $\mathtt{C}_{ij}$ is the covariance matrix of the $S_8(z_i)$'s obtained from the posterior. We find $\hat{S}_8 = 0.778 \pm 0.021$ which is in excellent agreement with Eq. (5.22).

Goodness of fit and consistency with *Planck* ΛCDM

The best-fit HEFT prediction to the combined (PR4 + DR6) analysis (without BAO) is shown in Fig. 5.3. In each panel we quote the χ^2 value for each individual measurement. The best-fit $\chi^2 = 53.1$ for 96 data points and 30 free parameters, 16 of which are partially prior-dominated. Following Appendix D.4 we estimate PTE = 96%. We note that here and in the companion paper [320] we approximate the LRGs as Gaussian random fields when estimating covariance matrices. This approximation likely underestimates the covariance between high-ℓ bandpowers and hence overestimates the PTE. We leave an investigation of the trispectrum contribution to the covariance matrix to future work, and note that for our fiducial linear theory analysis (for which these contributions are considerably smaller and likely negligible) we estimate a lower PTE = 83%.

We assess the consistency of our results with a *Planck* ΛCDM cosmology following two approaches. First we consider the "standard" approach of citing tension metrics for a given cosmological parameter, defined [in units of σ] as the difference between the means divided by the square root of the quadrature sum of the respective errorbars (i.e. ignoring correlations in the datasets). Our combined S_8 constraint (Eq. 5.22) is 2.3σ lower than that preferred by *Planck* 2018 [475] primary CMB data (TT,TE,EE+lowE) and 1.8σ (2.2σ) lower than the latest PR4 analysis from Tristram et al. [592] (Rosenberg et al. + `SRoll2` low ℓ EE [503, 452]). These "tensions" are less significant when comparing σ_8 values. Our combined σ_8 constraint (Eq. 5.23) is 1.8σ lower than PR3 and 1.6σ (1.7σ) lower than the aforementioned PR4 analyses.

Our second approach is to fix the cosmology to *Planck* 2018 (see §5.5 and Table 5.3) and to compare the resulting best-fit and PTEs to that obtained when $\ln(10^{10}A_s)$ and $\Omega_c h^2$ are allowed to float freely. We find $\Delta\chi^2 = 5.0$ between the best-fit with fixed and free cosmology, while with fixed cosmology we estimate PTE = 93% following Appendix D.4. We conclude that the *Planck* 2018 cosmology yields a perfectly reasonable fit to our data.

Alternative models

Here we present cosmological constraints from the alternative models discussed in §5.4 and §5.4 using our fiducial scale cuts (Table 5.2) and priors (Table 5.3). These results are summarized in Fig. 5.17 and discussed in detail below.

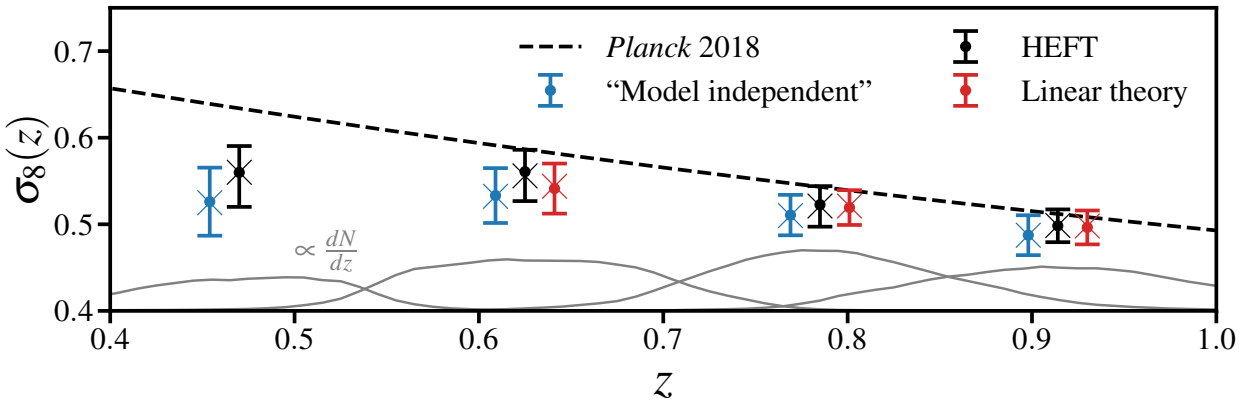


Figure 5.17: A comparison of our fiducial HEFT $\sigma_8(z)$ constraints (black) to those obtained from a "model independent" (blue, §5.4) or linear theory (red, §5.4) approach. We note that there is no linear theory point for z_1 as we do not include this bin in our fiducial linear theory analysis. In all scenarios we jointly fit to the cross-correlation with *Planck* PR4 and ACT DR6 in each individual redshift bin, and use the fiducial scale cuts and priors listed in Tables 5.2 and 5.3 respectively. We include a BAO prior for both the HEFT and linear theory constraints, while for the model-independent constraints the cosmology is fixed to the values listed in Table 5.3. We offset the linear theory and model independent constraints by $\Delta z = \pm 0.015$ respectively from the effective redshifts for clarity. The black dashed line is the predicted ΛCDM evolution assuming a *Planck* 2018 cosmology [475]. In gray we plot $0.4 + c\,dN/dz$ for each galaxy sample, where c is an arbitrary constant, to illustrate the redshift origin of the signal. Following §5.7 we estimate 9, 13 and 28% correlations between the $z_1 - z_2$, $z_2 - z_3$ and $z_3 - z_4$ amplitude measurements respectively.

We first consider "linear theory" constraints (§5.4). When fitting to each redshift bin independently without a BAO prior, we find

$$S_8 = 0.761 \pm 0.042\ [0.763],\ 0.796 \pm 0.034\ [0.796],\ 0.797 \pm 0.032\ [0.799] \tag{5.25}$$

for z_2, z_3, z_4 respectively when jointly fitting to the PR4 and DR6 cross-correlations, or $S_8 = 0.787 \pm 0.024$ [0.789] from the combination of bins $z_2 - z_4$. In combination with BAO data we find

$$\sigma_8 = 0.749 \pm 0.040\ [0.753],\ 0.776 \pm 0.030\ [0.779],\ 0.787 \pm 0.031\ [0.790] \tag{5.26}$$

from the individual redshift bins, and $\sigma_8 = 0.774 \pm 0.024$ [0.779] when jointly fitting to $z_2 - z_4$. We note that the linear theory analysis prefers slightly higher S_8 value than our fiducial HEFT analysis (cf. Eq. 5.22). This preference is driven by the exclusion of the lowest redshift bin (by default) in our linear theory fits, which prefers an amplitude that is $\sim 1\sigma$ lower than the remaining bins (see Fig. 5.16). When dropping the lowest redshift bin in our combined HEFT analysis, we find $S_8 = 0.785 \pm 0.023$ [0.781], which is in better agreement with our linear theory results. Next we consider the model independent approach discussed in §5.4. When fitting to the combination of PR4 and DR6, we find

$$\alpha_8(z_i) = 0.830{\pm}0.062\ [0.828],\ 0.909{\pm}0.054\ [0.907],\ 0.940{\pm}0.043\ [0.939],\ 0.952{\pm}0.045\ [0.952] \tag{5.27}$$

for the four redshift bins respectively. In particular, the constraint from the lowest redshift bin is 2.7σ lower than 1, while the remaining redshift bins are consistent with 1 at the $1-2\sigma$ level.

A comparison of the different modeling choices is given in Fig. 5.17, where we have rescaled the HEFT and linear theory σ_8 constraints (including BAO) from each redshift bin by $\sigma_8(z_{\mathrm{eff},\,i})/\sigma_8(0)$ and the model-independent constraints by $\sigma_8(z_{\mathrm{eff},\,i})$, where $\sigma_8(z)$ is computed assuming a *Planck* 2018 cosmology (see §5.5 and Table 5.3). We note that the HEFT and linear constraints agree to within $\simeq 0.5\sigma$, while the model-independent results are consistently lower than our fiducial HEFT constraints by $\simeq 1\sigma$.

Parameter based consistency tests

Here we present variations in our fiducial (§5.7) and linear theory (§5.7) S_8 constraints (without BAO) when adopting different analysis choices beyond the alternative models discussed in §5.7 and when analyzing different data subsets.

First we consider variations in scale cuts and priors in our fiducial PR4+DR6 analysis, which are summarized in the left half of Fig. 5.18. When decreasing ℓ_{max} for both the galaxy auto- cross-correlation with CMB lensing from 600 to 401, increasing ℓ_{min} from 20 (44) to 79 for the cross-correlation with *Planck* (ACT) lensing, increasing the ℓ_{min} for the galaxy auto from 79 to 124, or adopting a wider $\mathcal{U}(-5, 5)$ prior on b_s in each redshift bin, the mean S_8 value shifts by at most 0.4σ from our fiducial constraint (Eq. 5.22). In Fig. 5.18 we also consider a more aggressive analysis where we adopt a more restrictive $\ell_{\mathrm{max}} = 401$ but fix the counterterms to zero, reduce the width of the magnification bias prior from 0.1 to 0.05, and adopt highly informative $\mathcal{N}(0, 0.6)$ priors on b_2 and b_s. Doing so gives $S_8 = 0.788 \pm 0.018$ [0.788], which is $\simeq 0.7\sigma$ (in units of the aggressive errorbar) larger than our fiducial result. The aggressive constraints are more directly comparable to those found

with our fiducial priors but with a reduced $\ell_{\rm max} = 401$ (second point in Fig. 5.18), for which we find $S_8 = 0.783 \pm 0.022$ [0.778], suggesting that our results are fairly insensitive to the assumed priors on higher order nuisance parameters.

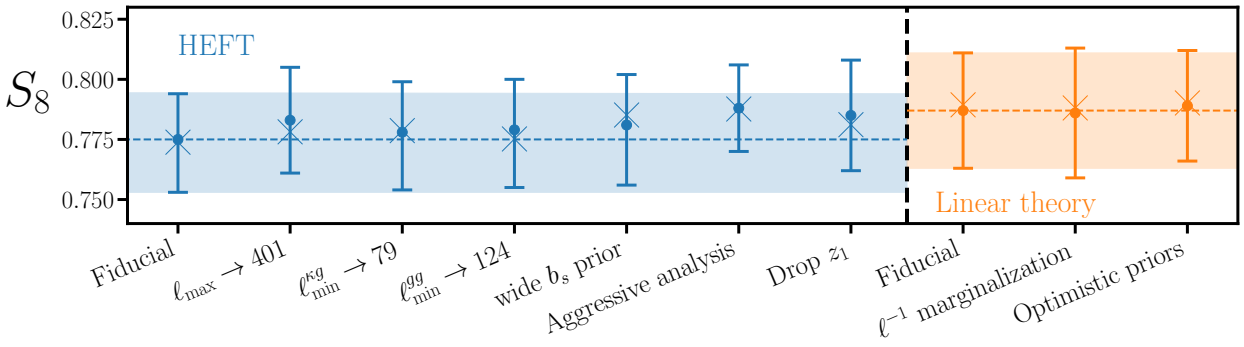


Figure 5.18: Variations (or lack thereof) in our baseline PR4+DR6 S_8 constraints (without BAO) for both HEFT (blue) and linear theory (orange) with different data subsets, priors and systematic mitigation. Best fit values are indicated with $\times$'s.

As discussed in §5.7 our fiducial PR4+DR6 linear theory result (leftmost orange point Fig. 5.18) is in good agreement with our fiducial HEFT constraint, but is more directly comparable with our HEFT analysis when dropping the first redshift bin (rightmost blue point in Fig. 5.18), for which we find even better agreement. Motivated by discrepancies between the inferred and measured galaxy cross-spectra (§5.6 and Fig. 5.12), we consider adding a ℓ^{-1} contribution (second orange point in Fig. 5.18) to C_ℓ^{gg} and marginalizing over its amplitude with a $\mathcal{N}(0, 10^{-4})$ prior. Doing so has a negligible impact on the mean and best fit S_8 values of our linear theory fits, which suggests that even if "systematic contamination" were to blame for the discrepancies observed in the galaxy cross spectra, these contaminants have a negligible impact on our cosmological results. Finally, we performed linear theory fits with more optimistic priors on shot noise (Gaussian width decreased from 30% to 10%), magnification bias (Gaussian width decreased from 0.1 to 0.05) and counterterms (Gaussian width decreased from 3 to 0.5) and found negligible changes.

Motivated by the statistically significant variations in the galaxy auto-spectra on different imaging footprints (§5.6) we fit to each of these measurements individually as a parameter based test. These result are summarized in Fig. 5.19. We first consider fitting the galaxy auto- and PR4 cross-correlation measured on the North, DECaLS and DES footprints using our fiducial HEFT model (blue points in Fig. 5.19) with the redshift distributions calibrated on each respective footprint [643] (see Fig. 5.2). For all three footprints we find results that are statistically consistent with our fiducial HEFT constraint. Next we fit to the data measured on the DEC $> -15°$ subset of the LRG footprint, finding $S_8 = 0.759 \pm 0.024$ [0.759] which is in very good agreement with our fiducial PR4 result (cf. Eq. 5.17), suggesting that potential variations in the redshift distributions north and south of DEC $= -15°$ have a small impact on our final results.

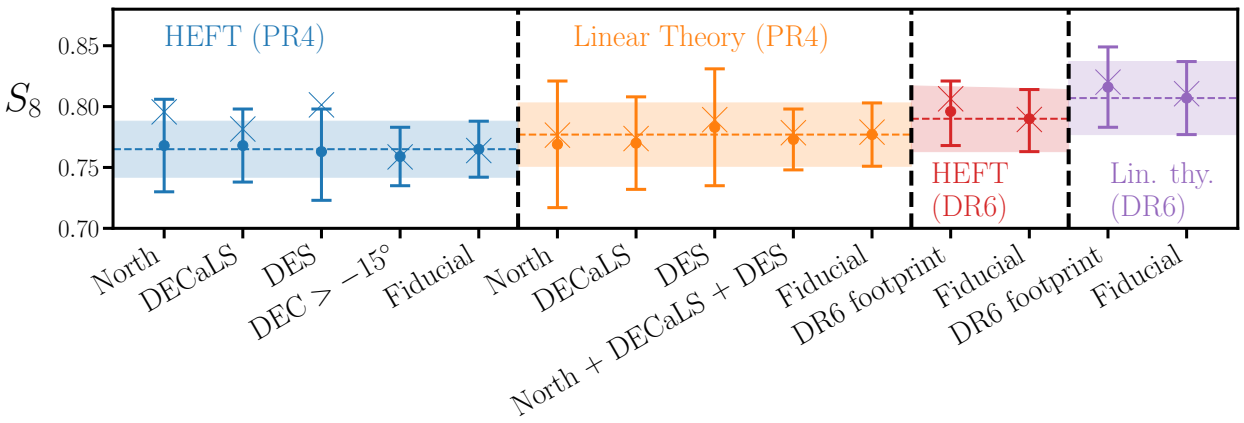


Figure 5.19: Variations in S_8 constraints (without BAO) when additionally masking the LRG sample with the North, DECaLS, DES, DEC $> -15°$, or ACT DR6 lensing mask. Best-fit values are indicated with $\times$'s. The horizontal lines and shaded bands correspond to mean and $\pm 1\sigma$ intervals of the fiducial constraints for each scenario (PR4 or DR6; HEFT or linear theory). We suspect that the discrepancies between the mean and best fit S_8 values when analyzing the individual imaging footprints with HEFT are the result of noisy minima and volume effects.

With our fiducial HEFT model we find that the individual footprints prefer a best-fit S_8 value that is $\simeq 0.5 - 1\sigma$ larger than the mean. When holding A_s fixed to the best-fit of the fiducial PR4 cross-correlation analysis, the best-fit χ^2 for each of the individual footprints increases by at most 1.5. Given the modest shifts in χ^2 we hypothesize that the difference in best-fit and mean S_8 values seen in Fig. 5.19 are the result of noisy minima. We have repeated the exercise of fitting to each imaging footprint using linear theory (orange points) and find results that are consistent with our fiducial linear theory PR4 constraint: $S_8 = 0.777 \pm 0.026$ [0.779]. Unlike for the HEFT model, here we do not observe significant differences between the best fit and mean S_8 values. We suspect that these minima are more stable than for the HEFT model, due to fewer nuisance parameters (30 vs $\simeq$ 14). We also performed a joint linear theory fit to the North, DECaLS and DES data (12 galaxy auto- and 12 cross-spectra) where we marginalize over nuisance terms for each sample individually. Doing so yields $S_8 = 0.773 \pm 0.025$ [0.779] which is in extremely good agreement with our fiducial result, suggesting that variations in the LRG linear bias across footprints has a negligible bias on our final results, which is in reasonable agreement with the back-of-the-envelope calculation made in §5.6. As a final footprint-variation test we reanalyze the cross-correlation with ACT DR6 when additionally masking the LRGs by the DR6 lensing mask, which ensures that the galaxy auto- and cross-correlation are probing the same effective linear bias (to leading order). When using a HEFT (red) or linear theory (purple) model our S_8 constraints shift by less than 0.3σ.

The companion paper [320] performs an extensive set of bandpower-level null tests for

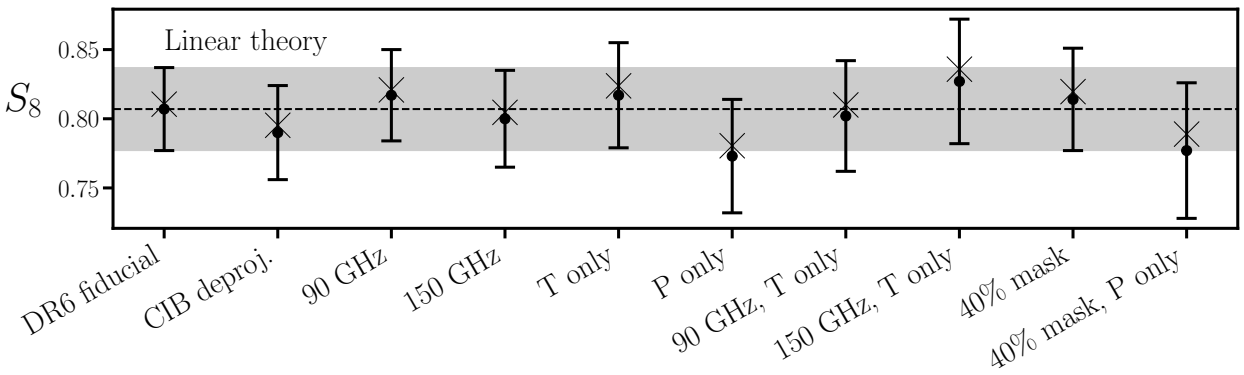


Figure 5.20: Linear theory DR6 S_8 constraints (without BAO) for alternative lensing reconstruction scenarios. "T and P only" refer to temperature and polarization only lensing reconstruction respectively; "CIB deproj." is short for a CIB deprojected internal linear combination of the 90 and 150 GHz channels; while "40% mask" is short for a more restrictive Galactic mask that is applied to the primary CMB maps before performing the lensing reconstruction. Best-fit values are indicated by $\times$'s. The shaded region corresponds to the $\pm 1\sigma$ interval associated with our fiducial linear theory DR6 analysis. We find statistically consistent results across different frequencies, temperature and polarization, different galactic masks, and combinations of the above.

the cross-correlation with DR6. Here we present parameter-based tests (using linear theory for computational efficiency) for a subset of these checks. To improve the value of these tests as a cross-check, we have independently measured (following §5.3) the cross-correlation with the publicly available alternative DR6 lensing maps. These results are summarized in Fig. 5.20. We find statistically consistent results across different frequencies, temperature and polarization, different galactic masks, and combinations of the above.

5.8 Comparison with previous weak lensing analyses

We compare our structure growth measurements with other recent tomographic constraints from correlations of CMB lensing and galaxy positions in Fig. 5.21. In purple we show the evolution of S_8 (and σ_8) inferred from our baseline PR4+DR6 analysis, where as in §5.7 we rescale the amplitude constraint from each redshift bin by the factor $\sigma_8(z)/\sigma_8(z=0)$ calculated with a *Planck* 2018 ΛCDM cosmology. We also show the mean (purple line) and $\pm 1\sigma$ (shaded purple region) interval for the ΛCDM evolution of $S_8(z) \equiv \sigma_8(z)\sqrt{\Omega_m/0.3}$ and $\sigma_8(z)$ as inferred from the PPD of our fiducial PR4+DR6 analysis (including all four redshift bins). In blue we show the $S_8(z)$ and $\sigma_8(z)$ constraints derived from the blue and green unWISE samples from a recent cross-correlation with the same CMB lensing maps used in this chapter (Farren et al. [223]), which have been rescaled to the mean redshifts

0.6 and 1.1 respectively. In orange we show the $\sigma_8(z)$ constraints obtained from the cross-correlation of the PR4 κ map with lowest two redshift bins of *Quaia* quasars (Piccirilli et al. [469]).

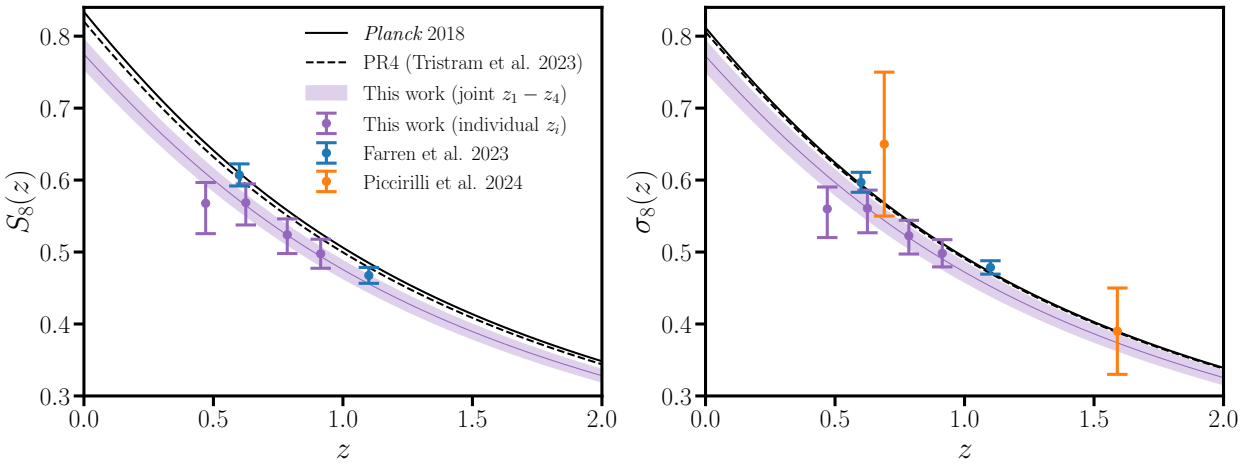


Figure 5.21: A comparison of our results with recent tomographic structure growth measurements obtained from cross-correlations of CMB lensing and galaxy positions. In blue we show the results obtained from the blue and green unWISE samples (correlated with PR4 and DR6) [223], while in orange we show the results obtained from the lowest two redshift bins of the *Quaia* quasar sample (correlated with PR4) [469]. Following §5.7 we estimate 9, 13 and 28% correlations between the $z_1 - z_2$, $z_2 - z_3$ and $z_3 - z_4$ amplitude measurements respectively.

We compare our combined constraints to other weak-lensing based measurements in Fig. 5.22. As in §5.7 we follow the standard parameter-based approach when quoting "tensions" with other datasets. We quote tensions for the S_8 parameter by default, however, for analyses listed in Fig. 5.22 that do not report S_8 measurements we instead quote the tension for σ_8. We caution the reader that this tension metric neglects correlations between datasets (typically resulting in an underestimation of the tension) and implicitly assumes that the posterior is Gaussian. Beyond this, the assumptions (including priors) made to model these datasets vary considerably with the observable and analyst.

We first compare our results to those also derived from cross-correlations of CMB lensing and galaxy positions (blue points in Fig. 5.22). With the exception of White et al. [625] (see §5.7), our analysis is most directly comparable to Farren et al. [223], which analyzes the same CMB lensing maps with a similar (in spirit) EFT-inspired model. From the combination of the blue and green unWISE samples ref. [223] finds $S_8 = 0.810 \pm 0.015$ (without BAO) which differs from our fiducial result by 1.4σ (cf. Eq. 5.22). We note that our errorbar is $\simeq 37\%$ larger than that obtained from the unWISE sample. This difference is partially driven by the larger number density and broader redshift coverage of the unWISE galaxies,

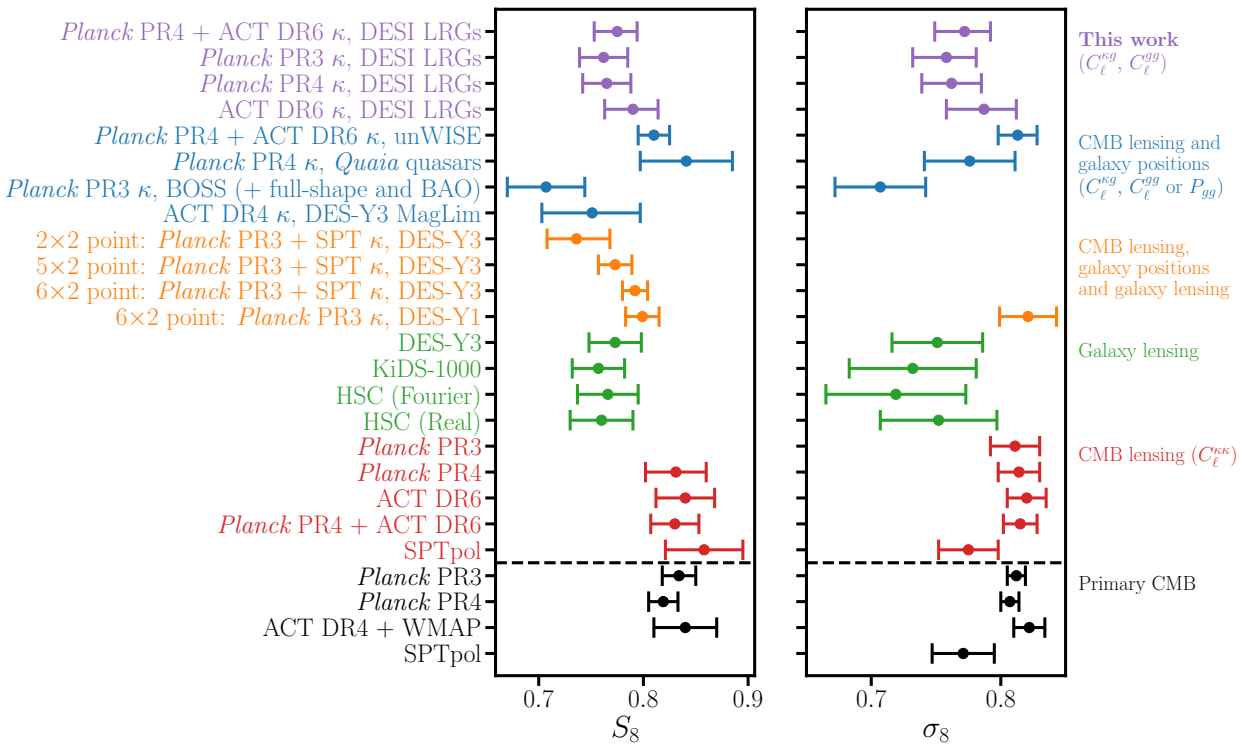


Figure 5.22: A compilation of S_8 and σ_8 constraints from cross-correlations of galaxy positions and CMB lensing (purple and blue); cross-correlations of CMB lensing with galaxy lensing and their positions (orange); galaxy lensing (green); the CMB lensing power spectrum (red) and the primary CMB (black). We differentiate "early" and "late" time measurements with the black dashed line. Constraints derived from this chapter are shown in purple, with the purple band representing the $\pm 1\sigma$ contour of our fiducial PR4+DR6 analysis. See the text in §5.8 for a detailed discussion.

however, we also emphasize that ref. [223] adopts highly informative priors on higher order corrections to the bias expansion (in particular, both counterterms were fixed to zero). In §5.7 we consider a more "aggressive analysis" whose scale cuts and priors were chosen to roughly match those used by ref. [223]. With these choices we find that our errorbar is $\sim 20\%$ larger than ref. [223]'s, suggesting that this higher precision is driven just as much by the priors assumed as it is by the larger raw signal-to-noise of the unWISE dataset. We find a similar $\sim 1.4\sigma$ difference with the cross-correlation of PR4 with *Quaia* quasars [469], which found $S_8 = 0.841 \pm 0.044$. We find a slightly larger $\sim 1.6\sigma$ difference with the S_8 measurement derived from cross-correlating PR3 with BOSS galaxies [140] (including full-shape and post-reconstruction measurements), however, this analysis did not include a "normalization" correction to the cross-correlation measurement and is susceptible to volume effects, which when taken into account would reduce the significance of this tension. Our

results are in very good ($\sim 0.5\sigma$) agreement with those obtained from cross-correlating DES-Y3 MagLim galaxies with the ACT DR4 lensing map [385].

Our results are in good agreement with those derived from cross-correlation with CMB lensing and galaxy lensing, with differences ranging from $\sim 0.1-1.1\sigma$. The orange points in Fig. 5.22 show the constraints obtained from the "2×2 point" analysis [128] of correlations between PR3+SPT CMB lensing with DES-Y3 galaxy positions and lensing ($\langle\delta_g\kappa\rangle + \langle\gamma\kappa\rangle$), constraints derived from the "5×2 point" analysis that adds correlations between galaxy lensing and their positions (including each auto-correlation), and the "6×2 point" analysis that in addition to the above adds the CMB lensing auto-correlation [8]. We also show constraints from an alternative 6×2 point analysis using PR3 and DES-Y1 data that marginalizes over baryonic feedback uncertainties using a principle component-based approach [630], which is in good agreement with the DES-Y3 result.

Our results are in even better agreement (within $\sim 0.6\sigma$) with those obtained from galaxy lensing alone. The green points in Fig. 5.22 show the constraints obtained from a reanalysis of DES-Y3 [6], KiDS-1000 [286] and HSC (both in Fourier [164] and real [349] space) data when adopting the priors on cosmological parameters used in the recent ACT DR6 auto-correlation analysis [377]. The constraints plotted in Fig. 5.22 are taken from Table 2 of ref. [377].

Our results are consistent to within 1.4, 1.6, 1.9, 1.8, and 2.0σ with those obtained from *Planck* PR3 [25], PR4 [119], ACT DR6 [377], PR4+DR6 and SPTpol [88, 453] analyses of the CMB lensing power spectrum respectively, indicated by the red points in Fig. 5.22. We emphasize that the CMB lensing power spectrum probes different redshifts (with the signal peaking around $z\sim2$) and scales than our tomographic analysis, leaving open the possibility that non-standard structure growth for $z<1$ could consistently reconcile these (mildly significant) differences. We defer the reader to §5.7 for a discussion of the consistency of our results with *Planck* primary CMB measurements, which are also shown in Fig. 5.22 (black points). We find a similar level of agreement ($\sim 1.8\sigma$) with primary CMB measurements from ACT DR4 + WMAP [26], and even better agreement with measurements from SPTpol [284].

5.9 Discussion and conclusions

We infer the amplitude of late time large scale structure from the cross-correlation of DESI legacy survey LRGs with the latest CMB lensing maps from *Planck* and ACT. The code base used to perform this analysis is publicly available (`MaPar`) and discussed in more detail in Appendix D.1. From the joint analysis of all four LRG redshift bins we obtain a 2.6% constraint of $S_8 = 0.775^{+0.019}_{-0.022}$ [0.774], while with the addition of BAO data we obtain a 2.7% constraint of $\sigma_8 = 0.772^{+0.020}_{-0.023}$ [0.775] (Eqs. 5.22 and 5.23). These results are less discrepant with primary CMB measurements than was found in a previous LRG cross-correlation analysis [625] (see §5.7 for a discussion). Our S_8 measurement is 5% (7%) lower than obtained from *Planck* PR4 (PR3) primary CMB data with a statistical significance

ranging from $1.8-2.3\sigma$ (§5.7). Our results are consistent with the decade-long claims of a low S_8 from weak galaxy shear measurements, however, unlike for shear our data's preference for a low S_8 cannot be explained by baryonic feedback [40] as our signal is derived primarily from the linear regime. When jointly analyzed with primary CMB measurements, our data would have a mild preference for beyond-ΛCDM physics (see e.g. [10]) that suppresses structure evolution on linear scales and at late times.

Through individually analyzing four photometric-redshift bins we constrain the evolution of structure growth across $0.4 \lesssim z \lesssim 1$. As indicated by Fig. 5.21 we find a rate of growth that is largely consistent with a ΛCDM prediction conditioned on primary CMB observations with an overall normalization that is $\simeq 5-7\%$ lower. As was found in White et al. [625], the constraints derived from the lowest redshift bin are most discrepant with a *Planck* ΛCDM prediction. The lowest redshift bin is also the least statistically constraining, and as such it is inconclusive if this deviation is the result of a statistical fluctuation or a preference for a slower rate of structure growth for $z \lesssim 0.5$. These scenarios could potentially be distinguished through the cross-correlation of CMB lensing with lower-redshift samples (albeit with modest precision due to the limited dynamic range of linear scales and limited overlap with the CMB lensing kernel at these redshifts), such as the DESI Bright Galaxy Survey [274], which we examine in the next chapter.

In our fiducial analysis we adopt a Hybrid EFT (HEFT) model (§5.4) that approximates the early-time galaxy distribution as a perturbative bias expansion that is advected to the present day using the non-linear CDM displacements measured from the `Aemulus` ν simulations [181]. In §5.5 we verify that our likelihood recovers unbiased cosmological results from a set of realistic mock measurements for both our fiducial HEFT model and alternatives presented in §5.7. We marginalize over both quadratic and lowest-order derivative corrections to the bias expansion with relatively wide priors. Doing so robustly regulates the cosmological information that can be obtained from smaller scales, such that our fiducial cosmological constraints are derived primarily from the linear regime ($k < 0.1\,h\,\mathrm{Mpc}^{-1}$). We explicitly demonstrate this point by performing a linear theory analysis (§5.7) whose scale cuts were chosen such that only modes with $k < 0.1\,h\,\mathrm{Mpc}^{-1}$ efficiently contribute to the signal, and find excellent agreement with our fiducial results with a comparable statistical precision.

Our data are largely immune to systematic uncertainties faced by photometric surveys. CMB lensing has a well-characterized source distribution and is measured using very well-understood statistical properties of primordial fluctuations, while the redshift distributions of the four LRG photometric redshift bins were spectroscopically calibrated by ref. [643]. To further ensure the robustness of our measurements we performed an extensive set of systematics, null and parameter-based consistency tests. In a companion paper [320] we present a suite of bandpower-level null tests for the ACT cross-correlation measurement and confirm with simulations that the expected bias from extragalactic foregrounds is negligible. In this chapter we thoroughly inspect the uniformity of the LRG sample's physical properties by measuring their power spectra on different footprints (see sections 5.6 to 5.6), confirm that our results are robust to stricter cuts on extinction and stellar contamination (§5.6) and quantify the (negligible) bias arising from residual correlations between the systematic

weights and (systematics-corrected) LRG maps (§5.6). We verify that our results are robust to variations in the scale cuts and priors assumed in §5.7. While the vast majority of these tests find no significant evidence for systematic contamination, there are two notable exceptions: (1) we find statistically significant variations in the high-z LRG auto-correlations on the DES footprint vs elsewhere (§5.6) and (2) the measured galaxy cross-spectra (§5.6) of neighboring redshift bins are slightly smaller than predicted from fits to the galaxy auto- and cross-correlation with CMB lensing. In passing, we detect a non-zero magnification signal by cross-correlating the lowest and highest LRG bins (which have negligible overlap in their redshift distributions) at 8σ significance. In §5.7 we perform a set of parameter-based consistency tests which suggest that neither of these potential sources of systematic contamination significantly bias our cosmological results. We encourage future work to further investigate these findings to improve the fidelity of the LRG sample.

In the near future DESI will obtain redshifts for the majority of LRGs used in this chapter, improving both the calibration of the redshift distribution and enabling more accurate treatments of systematic contaminants. These data will also be vital to improving the systematic characterization (e.g. through indirect redshift distribution calibration) of other photometric samples, which in particular may prove useful for diagnosing the (mild) discrepancy with the latest unWISE analysis [223]. In addition, these data will enable "full-shape" analyses of the 3D galaxy-power spectrum that can (somewhat) independently constrain $\sigma_8(z)$ from redshift-space distortions. Full-shape analyses will serve as a valuable consistency check of our results, and can be used to assess the dynamical consistency of relativistic and non-relativistic tracers within General Relativity when combined with CMB lensing cross-correlations. A joint analysis including LRG full-shape and BAO data will be presented in future work.

Future galaxy surveys such as Rubin Observatory's Legacy Survey of Space and Time (LSST) [310], Roman [563], Euclid [37] and other next-generation imaging and spectroscopic surveys will further refine these cross-correlation measurements due to their greater number densities and wider sky areas. At the same time, these cross-correlations will benefit significantly from lower noise CMB lensing maps produced from upcoming high-resolution telescopes such as Simons Observatory [22] and CMB-S4 [1]. The methods presented here can be readily extended to these datasets, enabling even more precise constraints on low redshift structure growth while maintaining an acceptable level of accuracy.

Acknowledgments

We would like to thank Alex Krowleski, Antón Baleato Lizancos, Niall MacCrann, Bruce Partridge, John Peacock, David Spergel, and Noah Weaverdyck for helpful discussions during the preparation of this manuscript. NS is supported by the Office of Science Graduate Student Research (SCGSR) program administered by the Oak Ridge Institute for Science and Education for the DOE under contract number DE-SC0014664. JK acknowledges support from NSF grants AST-2307727 and AST-2153201. SF is supported by Lawrence Berkeley

National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. MM acknowledges support from NSF grants AST-2307727 and AST-2153201 and NASA grant 21-ATP21-0145.

IAC acknowledges support from Fundación Mauricio y Carlota Botton and the Cambridge International Trust. EC acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 849169). JD acknowledges support from NSF award AST-2108126. CEV received the support of a fellowship from "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/EU22/11930099. GSF acknowledges support through the Isaac Newton Studentship and the Helen Stone Scholarship at the University of Cambridge. GSF furthermore acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant agreement No. 851274). KM acknowledges support from the National Research Foundation of South Africa. CS acknowledges support from the Agencia Nacional de Investigación y Desarrollo (ANID) through Basal project FB210003. This research has made use of NASA's Astrophysics Data System and the arXiv preprint server.

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of High-Energy Physics, under Contract No. DE-AC02-05CH11231, and by the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility under the same contract. Additional support for DESI was provided by the U.S. National Science Foundation (NSF), Division of Astronomical Sciences under Contract No. AST-0950945 to the NSF's National Optical-Infrared Astronomy Research Laboratory; the Science and Technologies Facilities Council of the United Kingdom; the Gordon and Betty Moore Foundation; the Heising-Simons Foundation; the French Alternative Energies and Atomic Energy Commission (CEA); the National Council of Science and Technology of Mexico (CONACYT); the Ministry of Science and Innovation of Spain (MICINN), and by the DESI Member Institutions: https://www.desi.lbl.gov/collaborating-institutions.

The DESI Legacy Imaging Surveys consist of three individual and complementary projects: the Dark Energy Camera Legacy Survey (DECaLS), the Beijing-Arizona Sky Survey (BASS), and the Mayall z-band Legacy Survey (MzLS). DECaLS, BASS and MzLS together include data obtained, respectively, at the Blanco telescope, Cerro Tololo Inter-American Observatory, NSF's NOIRLab; the Bok telescope, Steward Observatory, University of Arizona; and the Mayall telescope, Kitt Peak National Observatory, NOIRLab. NOIRLab is operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation. Pipeline processing and analyses of the data were supported by NOIRLab and the Lawrence Berkeley National Laboratory. Legacy Surveys also uses data products from the Near-Earth Object Wide-field Infrared Survey Explorer (NEOWISE), a project of the Jet Propulsion Laboratory/California Institute of Technology, funded by the National Aeronautics and Space Administration. Legacy Surveys was supported by: the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy; the National Energy Research Scientific Computing Center, a DOE Office of Science User

Facility; the U.S. National Science Foundation, Division of Astronomical Sciences; the National Astronomical Observatories of China, the Chinese Academy of Sciences and the Chinese National Natural Science Foundation. LBNL is managed by the Regents of the University of California under contract to the U.S. Department of Energy. The complete acknowledgments can be found at https://www.legacysurvey.org/.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the U. S. National Science Foundation, the U. S. Department of Energy, or any of the listed funding agencies.

The authors are honored to be permitted to conduct scientific research on Iolkam Du'ag (Kitt Peak), a mountain with particular significance to the Tohono O'odham Nation.

Support for ACT was through the U.S. National Science Foundation through awards AST-0408698, AST-0965625, and AST-1440226 for the ACT project, as well as awards PHY-0355328, PHY-0855887 and PHY-1214379. Funding was also provided by Princeton University, the University of Pennsylvania, and a Canada Foundation for Innovation (CFI) award to UBC. ACT operated in the Parque Astronómico Atacama in northern Chile under the auspices of the Agencia Nacional de Investigación y Desarrollo (ANID). The development of multichroic detectors and lenses was supported by NASA grants NNX13AE56G and NNX14AB58G. Detector research at NIST was supported by the NIST Innovations in Measurement Science program. Computing for ACT was performed using the Princeton Research Computing resources at Princeton University, the National Energy Research Scientific Computing Center (NERSC), and the Niagara supercomputer at the SciNet HPC Consortium. SciNet is funded by the CFI under the auspices of Compute Canada, the Government of Ontario, the Ontario Research Fund–Research Excellence, and the University of Toronto. We thank the Republic of Chile for hosting ACT in the northern Atacama, and the local indigenous Licanantay communities whom we follow in observing and learning from the night sky.

Chapter 6

Tomographic structure growth II: cross-correlation of the DESI Bright Galaxy Survey with CMB lensing and galaxy magnification

This chapter was originally published as

> Noah Sailer, Joseph DeRose, Simone Ferraro, Shi-Fan Chen, Rongpu Zhou, Martin White, Joshua Kim, and Mathew Madhavacheril. "Evolution of structure growth during dark energy domination: Insights from the cross-correlation of DESI galaxies with CMB lensing and galaxy magnification". In: *Phys. Rev. D* 111.10, 103540 (May 2025), p. 103540. DOI: 10.1103/27rg-tq8z. arXiv: 2503.24385 [astro-ph.CO]

We use a Hybrid Effective Field Theory (HEFT) model to constrain the evolution of low-redshift ($z \lesssim 0.4$) matter fluctuations by cross-correlating DESI Bright Galaxy Survey (BGS) legacy imaging with the latest CMB lensing maps from *Planck* and ACT. Our tomographic BGS analysis finds that the evolution and amplitude of matter fluctuations align with CMB-conditioned ΛCDM predictions. When including DESI Baryon Acoustic Oscillation (BAO) measurements we obtain $\sigma_8 = 0.876^{+0.051}_{-0.067}$ from BGS alone. Jointly analyzing BGS and Luminous Red Galaxy (LRG) cross-correlations with the same CMB lensing maps yields $\sigma_8 = 0.791 \pm 0.021$. As a complementary approach we isolate the galaxy magnification signal from the cross-correlation of non-overlapping BGS and LRG photometric redshift bins, ruling out the null-magnification hypothesis at 11σ. For the first time, we constrain quasi-linear structure growth from the (finite-difference calibrated) galaxy magnification signal and find $\sigma_8 = 0.720 \pm 0.047$ (stat.) ± 0.050 (sys.) assuming a linear bias model and including BAO data.

6.1 Introduction

The linear growth of matter perturbations is a precise prediction of the standard cosmological model. While variations in the dark energy (DE) equation of state are imprinted on the low-redshift growth history via Hubble drag, in practice this effect is small relative to $\delta H/H$ for the cosmologies favored by the DESI Y3 BAO data [184] (see Fig. 6.1). Thus the low-redshift evolution of $\sigma_8(z)/\sigma_8$ offers a complementary view of DE which is instead primarily sensitive to DE interactions [e.g. 554, 331, 46, 193, 369, 43, 483, 333] or modified gravity alternatives [311, 315, 314, 37, 557], while a comparison of the present-day σ_8 with the primordial amplitude A_s probes e.g. the neutrino mass sum [103, 299, 339, 255, 367] or deviations in the expansion history at higher redshifts [137, 512].

At low redshifts where DE is relevant the universe exhibits large nonlinearities, limiting the statistical precision achievable from perturbative scales. Higher precision can be attained by incorporating smaller-scale measurements such as galaxy shear [286, 363, 6, 39, 538, 411, 571, 164, 349] but this comes at the cost of modeling increasingly complex and non-linear processes. Accurately modeling these processes is essential for robust inference. Recognition of systematic effects – such as baryonic feedback [525, 91, 265] and intrinsic alignments [535, 99, 608, 133, 129] – has led to greater reported uncertainties from these probes, bringing them to a level comparable to those obtainable from perturbative scales.

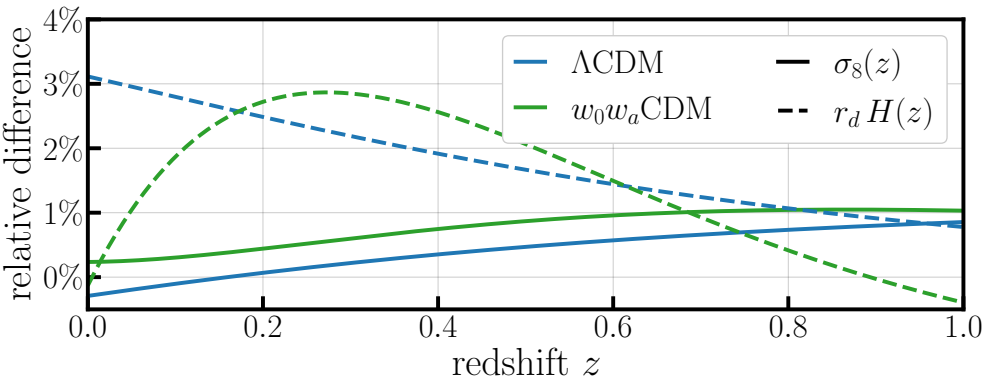


Figure 6.1: Relative difference in the growth (solid) and expansion (dashed) history between best-fit DESI ΛCDM (blue) or w_0w_aCDM (green) [184] and best-fit *Planck*+ACT ΛCDM [475, 362] cosmologies. Curves are computed with `CLASS` [97].

In this chapter we constrain $\sigma_8(z)$ using two approaches: (1) tomographic cross-correlations of DESI imaging data with CMB lensing and (2) the correlation of galaxy positions with galaxy magnification. While the former is more robust to non-linear systematic effects, the latter serves as a new probe of the low-redshift clustering on large scales that is complementary to CMB lensing and galaxy shear measurements[1]. We

[1]See [631] for an alternative approach correlating source fluxes (rather than positions) with the lens

describe the data and models used in each of these analyses in §6.2 and §6.3 respectively. Our results are presented in §6.4 and we conclude in §6.5.

6.2 Data

We use both the Bright Galaxy Survey (BGS) [274] and Luminous Red Galaxy (LRG) [640, 643] samples from the DESI Legacy Imaging Surveys DR9 [191] alongside the latest CMB lensing maps from *Planck* PR4 [119] and ACT DR6 [493, 372, 377]. The BGS and LRG imaging data cover 18200 deg^2 with mean angular number densities of 950 and 500 deg^{-2} respectively, while the CMB lensing data are signal dominated per mode out to $\ell \simeq 80$ (*Planck*) and $\ell \simeq 200$ (ACT). The imaging data have 16600 and 7900 deg^2 of overlap with the *Planck* and ACT lensing footprints, yielding 25 and 50σ detections of a non-zero CMB lensing cross-correlation for the BGS and LRG samples on the scales used in our analysis.

Ref. [643] divides the LRG sample into four photometric redshift bins whose redshift distributions are calibrated using 2.3 million DESI redshifts, while ref. [129] splits the BGS sample into two spectroscopically calibrated photo-z bins following an analogous approach. The redshift distributions of the six photo-z bins are plotted in purple in Fig. 6.3. Number count slopes ($s_\mu \equiv d\log_{10} N/dm$) for each bin were computed via finite differences following ref. [643] and are listed in Table 6.1. An extensive set of systematics checks were performed for the LRGs by refs. [640, 643, 320, 511] which found negligible systematic error at the level of cosmological parameters, and a subset of these tests for the BGS sample were performed by [129]. Refs. [620, 490] estimate that with (tSZ) profile-hardening [519, 517] residual extragalactic foreground contamination in the ACT DR6 κ map yield $\leq 0.02\sigma$ biases when cross-correlating with the BOSS LOWZ and Hang et al. [278] samples, which together cover the BGS redshift range. We deem those tests sufficient for our purposes, and further note that the BGS sample is notably brighter (hence smaller systematic modulation is expected) and is $\simeq 3\times$ less constraining than the LRGs.

We use `NaMaster` [35] to measure angular power spectra and compute their associated window functions. We adopt the same bandpowers and analytic approximation to the covariance as in ref. [511]. We apply a normalization correction to the raw CMB lensing cross-correlation measurements to approximately account for the nontrivial mode couplings introduced by masks and filters [223, 515]. In Fig. 6.2 we show the cross-correlation of the two BGS photo-z bins with CMB lensing (first two panels), the cross-correlation of BGS and the lowest redshift LRG bin (serving as lenses) with the two highest redshift LRG bins (serving as sources; third and fourth panels) and the BGS auto-correlations (bottom panel). On small scales we adopt an $\ell_{\rm max}$ such that the approximate $k_{\rm max} = 0.6\ h\,{\rm Mpc}^{-1}$ at the effective redshift of each photo-z bin. On these scales we expect our fiducial model (§6.3) to be accurate to subpercent precision [327]. On large scales we choose $\ell_{\rm min}$ such that both beyond-Limber effects (including redshift-space distortions) and the residual mean field contributions to the CMB lensing estimator are negligible [511].

sample on highly non-linear scales ($k \lesssim 50\ h\,{\rm Mpc}^{-1}$).

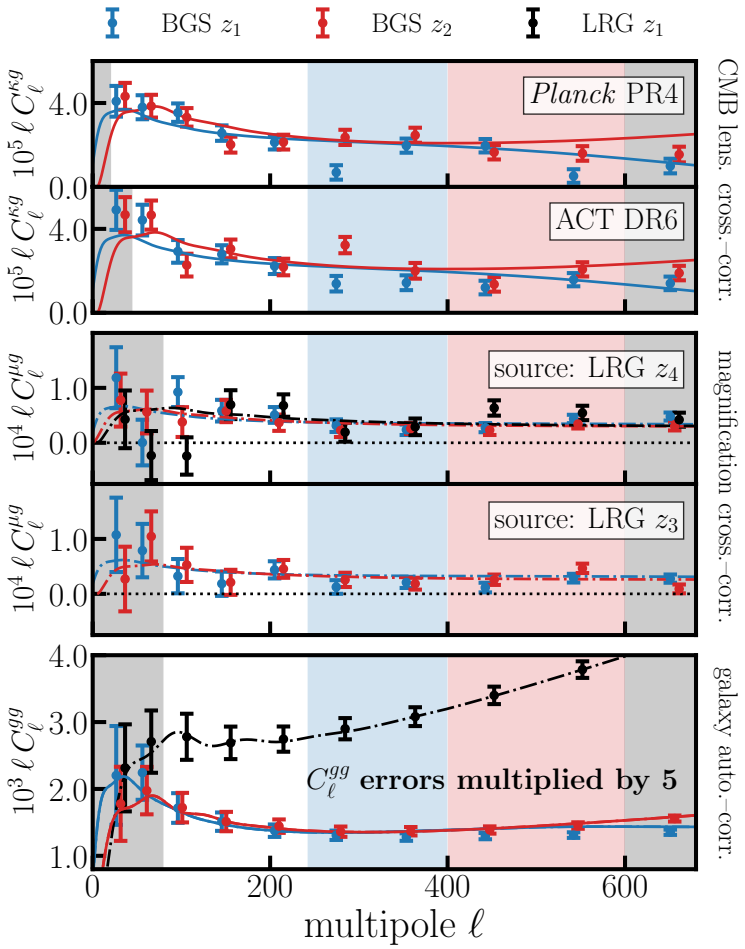


Figure 6.2: We show CMB lensing cross-correlations ($C_\ell^{\kappa g}$) in the top two panels, lens-source cross-correlations ($C_\ell^{\mu g}$) in the third and fourth panels, and galaxy auto-correlations (C_ℓ^{gg}) in the bottom panel. Colors indicate the galaxy sample. Black shaded regions are excluded by our scale cuts for all three samples; red regions are additionally excluded for both BGS bins; while the blue regions are additionally excluded for the lowest redshift BGS bin. Solid (dot dashed) lines indicate best-fit predictions from our BGS–CMB lensing cross-correlation (galaxy magnification) analysis. Note that the BGS best-fit C_ℓ^{gg} curves from both analyses lie on top of one another. We have shifted the data (and best-fit curves) by $\Delta\ell = \pm 5$ and rescaled the C_ℓ^{gg} errors by a factor of 5 for clarity.

Our measurements are primarily sensitive to $S_8 \equiv \sigma_8\sqrt{\Omega_m/0.3}$. To break the degeneracy between σ_8 and Ω_m and to mitigate volume effects we include the DESI Year 1 BAO data [15] which effectively acts as a prior on the present day matter density ($\Omega_m = 0.295 \pm 0.015$).

sample	physical properties			scale cuts				priors on nuisance parameters						
	$z_{\rm eff}$	$\bar{s}_\mu$	$10^7\,{\rm SN}^{\rm 2D}_{\rm pois}$	$\ell^{\rm auto}_{\rm min}$	$\ell^{\rm Planck}_{\rm min}$	$\ell^{\rm ACT}_{\rm min}$	$\ell_{\rm max}$	$b_1^L(z_{\rm eff})$	$b_2^L(z_{\rm eff})$	$b_s^L(z_{\rm eff})$	$\alpha_a(z_{\rm eff})$	$\alpha_x(z_{\rm eff})$	${\rm SN}^{\rm 2D}$	s_μ
BGS z_1	0.211	0.81	4.63	79	20	44	243	$\mathcal{U}(-1,2)$	$\mathcal{U}(-5,5)$	$\mathcal{N}(0,1)$	$\mathcal{N}(0,50)$	$\mathcal{N}\left(\frac{\alpha_a}{2b_1^E},2\right)$	$\mathcal{N}_r({\rm SN}^{\rm 2D}_{\rm pois},0.3)$	$\mathcal{N}(\bar{s}_\mu,0.1)$
BGS z_2	0.352	0.80	9.18	79	20	44	401	....			same as above		....	

Table 6.1: A summary of the BGS physical properties (effective redshift, magnification bias, Poisson shot noise), scale cuts and priors on nuisance parameters for our tomographic CMB lensing analysis. $\mathcal{N}(\mu,\sigma)$ and $\mathcal{U}(\min,\max)$ indicate Gaussian and uniform distributions respectively, while $\mathcal{N}_r(\mu,\sigma/\mu)\equiv\mathcal{N}(\mu,\sigma)$. The linear Eulerian bias is given by $b_1^E=1+b_1^L$.

6.3 Likelihoods

Our analysis of the CMB lensing cross-correlation with the BGS sample (§6.4) closely mirrors the approach taken by refs. [320, 511] for the DESI LRGs, for which the code (`MaPar`) is publicly available. The likelihood used for our magnification-derived growth constraints (§6.4) follows a nearly identical structure but with the lens-source cross-correlation replacing the CMB lensing cross-correlation. Below we briefly summarize the model and priors adopted for the BGS and LRG samples. We then show that our analysis choices yield acceptably small "volume effects" and faithfully recover the input cosmology to a set of BGS-like simulations.

Model, priors and sampling

We model the galaxy density contrast using a Hybrid Effective Field Theory (HEFT) [413, 327, 268, 637, 461, 440] implemented with the `Aemulus` ν simulations [181]. This implementation includes linear (b_1^L), quadratic (b_2^L and b_s^L) and leading-order derivative ($b_{\nabla^2}^L$) contributions to the galaxy bias expansion which enables a rigorous marginalization over the small-scale astrophysics relevant for galaxy formation. We use the Limber approximation [352, 317, 365] to compute the galaxy auto-correlation (C_ℓ^{gg}) and cross-correlation ($C_\ell^{\kappa g}$) with CMB lensing. We adopt an effective redshift ($z_{\rm eff}$) approximation when evaluating the three dimensional galaxy-galaxy (P_{gg}) and galaxy-matter (P_{gm}) power spectra within the Limber integrals. Due to the BGS sample's narrow dN/dz's (see Fig. 6.3) the effective redshift approximation has subpercent accuracy for physically reasonable galaxy bias evolutions. The effective redshift approximation dramatically simplifies the modeling, since bias parameters need only to be specified at each $z_{\rm eff}$ rather than their evolution within each photo-z bin. For additional details about the model and its implementation, including the treatment of pixelization and magnification, we refer the reader to §4 of ref. [511].

We model the cross-correlation of source and lens galaxies following §6.6 of ref. [511]. For simplicity, we have chosen only to include lens-source pairs with negligible redshift overlap in our magnification-derived growth constraints. In this case, the cross-correlation between

the source sample μ and the lens g is very well approximated by

$$C_\ell^{\mu g} \simeq \int \frac{d\chi}{\chi^2} W^\mu(\chi) W^g(\chi) P_{gm}\Big(k = \frac{\ell + 1/2}{\chi}, z_{\rm eff}\Big), \tag{6.1}$$

where $W^\mu \propto (5s_\mu - 2)$ is the magnification kernel of the source, $W^g \propto dN/dz$ is the projection kernel of the lens and $z_{\rm eff}$ is the effective redshift of the lens. Provided that the magnification bias s_μ of the source sample is known more precisely than the amplitude of the lens-source cross-correlation one can heuristically measure S_8 from the ratio $C_\ell^{\mu g}/\sqrt{C_\ell^{gg}}$. We use Eq. (6.1) as our lens-source model in §6.4 to derive growth constraints from the galaxy data alone.

We adopt a Gaussian likelihood and analytically marginalize over nuisance parameters which enter linearly into the model (see Appendix D of [511]) to decrease the convergence time of our chains. Posteriors are sampled with `cobaya`'s [586, 587] Markov Chain Monte Carlo Metropolis sampler [345, 343]. Plots of marginal distributions are obtained with the `GetDist` [344] package. We estimate best-fit values using `cobaya`'s default minimizer.

As in [320, 511] we fix $n_s = 0.9649$, $\Omega_b h^2 = 0.02236$, $\Omega_m h^3 = 0.09633$ and $N_{\rm eff} = 3.044$ to their *Planck* 2018 mean values[2] [475] and fix the neutrino mass sum to $M_\nu = 60$ meV. We sample $\ln(10^{10} A_s)$ and $\Omega_c h^2$ with uniform $\mathcal{U}(2, 4)$ and $\mathcal{U}(0.08, 0.16)$ priors respectively. We place uniform $\mathcal{U}(-1, 2)$ priors on the linear (Lagrangian) bias parameters (b_1^L) for both BGS photo-z bins. Priors on higher order bias parameters, counterterms, and shot noise are set in an identical fashion to those used in the fiducial analysis of refs. [320, 511] and are listed in Table 6.1. We marginalize over the magnification bias with the priors listed in Table 6.1 for CMB lensing-derived growth constraints (§6.4) while for magnification-derived constraints (§6.4) we hold magnification bias fixed to the central values measured by refs. [643, 129].

Likelihood tests

CMB lensing tomography: To estimate the expected size of volume (projection) effects in our CMB lensing analysis we analyze a noiseless data vector generated from our fiducial model with cosmological parameters fixed to those used for the `Buzzard` simulations (in particular, $S_8 = 0.801$ and $\sigma_8 = 0.820$). We assume a linear bias of $b_1^L(z_{\rm eff}) = 0.0, 0.2$ [129] for the two redshift bins, use the measured magnification bias and shot noise listed in Table 6.1 (under "physical properties") and set all higher order biases to 0 in our fiducial model prediction. Throughout we use the same covariance here as in our fiducial BGS analysis.

We list marginal S_8 posteriors obtained from the synthetic cross- and auto-correlation alone (i.e. without BAO information) in the first column of Table 6.2. We find that when fitting to both BGS bins the mean S_8 is $\simeq 0.8\sigma$ lower than the input value. The offsets are smaller when fitting to the individual samples: $\simeq 0.5\sigma$ and $\simeq 0.2\sigma$ for the lower and higher redshift bins respectively.

[2] Our results are largely insensitive to n_s: increasing n_s from 0.965 to 0.984 increases the inferred S_8 from our baseline analysis by 0.003 ($\simeq 0.14\sigma$).

One could in principle mitigate volume effects by adopting informative priors on galaxy nuisance parameters. As an extreme case we fixed magnification bias and all higher order contributions to their true values so that the only nuisance parameters varied are linear bias and shot noise, and still found the mean S_8 to be $\simeq 0.4\sigma$ lower than the input value when fitting to both BGS bins. An alternative approach to mitigate these shifts is to include BAO data when quoting BGS-only S_8 constraints. To demonstrate this we generated a noiseless DESI Y1 BAO data vector evaluated at the `Buzzard` cosmology. When including the mock BAO data (second and third columns of Table 6.2) we find that mean S_8 and σ_8 values are within 0.25σ of their input values when analyzing the BGS bins individually and together; thus we always include BAO in §6.4 for both our S_8 and σ_8 constraints.

	S_8 without BAO	S_8 with BAO	σ_8 with BAO
Fid. value	0.801	0.801	0.820
	noiseless model prediction		
BGS z_1	$0.782^{+0.050}_{-0.092}$ [0.806]	$0.800^{+0.067}_{-0.098}$ [0.798]	$0.819^{+0.069}_{-0.100}$ [0.817]
BGS z_2	$0.792^{+0.049}_{-0.065}$ [0.793]	$0.808^{+0.048}_{-0.071}$ [0.826]	$0.827^{+0.050}_{-0.073}$ [0.836]
BGS (z_1 & z_2)	$0.765^{+0.041}_{-0.051}$ [0.802]	$0.793^{+0.044}_{-0.054}$ [0.798]	$0.810^{+0.047}_{-0.057}$ [0.816]
	`Buzzard` simulations		
BGS z_1	$0.776^{+0.052}_{-0.089}$ [0.803]	$0.800^{+0.060}_{-0.092}$ [0.788]	$0.819^{+0.062}_{-0.095}$ [0.805]
BGS z_2	$0.781^{+0.042}_{-0.055}$ [0.806]	$0.766^{+0.043}_{-0.065}$ [0.787]	$0.783^{+0.046}_{-0.068}$ [0.804]
BGS	$0.740^{+0.039}_{-0.050}$ [0.790]	$0.762^{+0.040}_{-0.050}$ [0.801]	$0.778^{+0.042}_{-0.053}$ [0.819]

Table 6.2: Tomographic CMB lensing constraints (formatted as: mean $\pm 1\sigma$ [best fit]) from our fits to noiseless synthetic data (top three rows) and our fits to the `Buzzard` mock measurements (bottom three rows). In all cases we use the same covariance as in our fiducial analysis. The scatter in best-fits around the fiducial values for the noiseless model prediction are representative of the minimizer's precision.

To test our likelihood in a more realistic setting we next analyze a mock BGS sample built on the `Buzzard v2.0` suite of N-body simulations. We refer the reader to [178, 179] and the references therein [564, 177, 180, 616] for a detailed description of the sample's construction. The galaxy auto- and cross-correlation with CMB lensing are measured on seven quarter-sky simulations and binned into bandpowers in an identical fashion to the data. When fitting to these mock measurements we adopt the same scale cuts as in §6.2, use the mock window functions and redshift distributions, and adjust the fiducial values of fixed cosmological parameters to match those used in the `Buzzard` simulations, while the priors for sampled parameters remain identical to those listed in Table 6.1.

The Buzzard simulations cover $\simeq 4\times$ the area of the imaging data. Thus $\simeq 0.5\sigma$ fluctuations about the input cosmology are expected when fitting to the Buzzard measurements with the covariance used in our fiducial BGS analysis. To maintain this expectation when including BAO we add Gaussian noise to the mock BAO data described above whose covariance is $1/4\times$(the DESI Y1 BAO covariance). Results of the Buzzard tests are listed in the last three rows of Table 6.2. When fitting to the first (second) BGS bin only all of the posteriors are within 0.5σ (0.7σ) of the input cosmology. When jointly analyzing both bins we find 0.9σ (1.3σ) differences with (without) the inclusion of mock BAO data. From the fits to a noiseless model prediction we expect that $\simeq 0.25\sigma$ ($\simeq 0.8\sigma$) of these differences can be attributed to volume effects. We conclude that our mock tests are consistent within $\simeq 0.7\sigma$ (or smaller) fluctuations about the input cosmology when volume effects are taken into account, which are mild ($< 0.25\sigma$) with the inclusion of BAO data.

Galaxy magnification: We repeated the volume effect test for our galaxy magnification analysis (§6.4) on a synthetic data vector and found $\simeq 2\sigma$ shifts in S_8 and σ_8 for the priors used in the CMB lensing analysis (Table 6.1), despite including synthetic BAO data and fixing magnification bias[3]. Given the modest SNR = 11 of the magnification signal we chose to instead adopt a linear bias model. For the linear bias model volume effects are at most 0.3σ when BAO data are included.

We estimate the systematic error from beyond linear biasing as follows. We fix the cosmological parameters of our joint BGS & LRG likelihood (§6.4) to the Buzzard cosmology, shot noise to its Poisson estimate and number count slopes to their measured finite difference values. We then obtain best-fit values for the remaining nuisance parameters. Using these values we generate a synthetic data vector that includes realistic beyond linear biasing, which is then fit using our linear bias magnification likelihood including synthetic BAO. Doing so yields $\sigma_8 = 0.762 \pm 0.052$ [0.776]. Given that volume effects can account for $(\Delta\sigma_8)_{\rm vol.} = -0.015$ of this shift away from the true $\sigma_8 = 0.82$, we estimate a systematic shift $(\Delta\sigma_8)_{\rm sys.} \simeq -0.05$ from beyond linear biasing. This shift is comparable to the statistical precision and so we report both statistical and systematic errors in §6.4.

In our magnification model we neglect the redshift evolution of s_μ within the source photo-z bins, which ref. [643] found to be significant for the BOSS CMASS sample. If we instead assume a linear evolution of $s_\mu(z)$ whose slope is estimated via finite differences[4] our constraints shift by $\Delta\sigma_8 = -0.002$. This systematic is subdominant to beyond-linear biasing and we therefore neglect it.

[3] In practice the Poisson errors of the finite difference method are much smaller ($\sim 1\%$, see ref. [643]) than the precision of our magnification-derived growth constraint.

[4] Specifically we set $s_\mu(z) = \bar{s}_\mu + \bar{s}'_\mu(z - \bar{z})$ for both LRG source samples, where $\bar{z}$ is the mean redshift, $\bar{s}_\mu$ is is the mean slope measured by ref. [643] and $\bar{s}'_\mu \simeq -0.4$, 1.2 for LRG z_3 and z_4 respectively. We estimate $\bar{s}'_\mu$ by splitting each photo-z bin in half and measuring $\bar{s}_\mu$ for each half.

6.4 Results

In §6.4 and §6.4 we present growth constraints from CMB lensing tomography and galaxy-magnification respectively. To simplify the interpretation (i.e. mitigate volume effects, see §6.3) of our results we always include DESI Y1 BAO data. To mitigate confirmation bias in our tomographic CMB lensing analysis we added a 5% Gaussian scatter to the mean and best-fit cosmological parameter values when producing preliminary versions of Fig. 6.3 until all scale cuts and priors had been finalized.

CMB lensing tomography

We summarize our growth constraints from the cross-correlation of *Planck*+ACT CMB lensing with the DESI BGS and LRG samples in Table 6.3. Constraints from each individual BGS photo-z bin are listed in the first two rows. We jointly analyze both photo-z bins to obtain the third row. For comparison we quote the joint LRG constraints [320, 511] in the fourth row, which have been updated to include DESI Y1 BAO. In the final row we list the joint BGS & LRG constraints.

	S_8	σ_8
BGS z_1	$0.974^{+0.072}_{-0.110}$ [0.940]	$0.981^{+0.074}_{-0.110}$ [0.946]
BGS z_2	$0.785^{+0.038}_{-0.061}$ [0.768]	$0.792^{+0.040}_{-0.063}$ [0.776]
BGS	$0.870^{+0.050}_{-0.064}$ [0.892]	$0.876^{+0.051}_{-0.067}$ [0.899]
LRG [320, 511]	$0.775^{+0.019}_{-0.022}$ [0.779]	0.781 ± 0.022 [0.788]
BGS & LRG	0.788 ± 0.020 [0.788]	0.791 ± 0.021 [0.796]

Table 6.3: Constraints (formatted as: mean $\pm 1\sigma$ [best fit]) from cross-correlating *Planck*+ACT CMB lensing with several combinations of the DESI BGS and LRG photo-z bins. For all combinations we include DESI Y1 BAO data.

For BGS-only (third row, Table 6.3) we find a best-fit $\chi^2_{\kappa g,gg} = 30.2$ and $\chi^2_{\rm BAO} = 12.8$. The corresponding best-fit prediction for each BGS auto- and CMB lensing cross-correlation are plotted in Fig. 6.2 (solid lines). With the addition of LRG data (last row, Table 6.3) we find a best-fit $\chi^2_{\kappa g,gg} = 89.1$ and $\chi^2_{\rm BAO} = 12.8$. Our likelihood consists of 30 (126) data points and is fit with 16 (44) free parameters when analyzing the BGS (BGS & LRG) data, while the DESI Y1 BAO likelihood has 12 data points. We emphasize that some of our free parameters (e.g. s_μ, b_s, α_x) are prior dominated, and as such it is nontrivial to estimate the effective number of degrees of freedom. Instead we quote Bayesian PTEs [246] following Appendix B of ref. [511]. For BGS (BGS & LRG) we estimate PTE = 14% (77%).

Growth from galaxy magnification

For our magnification-derived constraints we jointly analyze three galaxy auto-spectra (fifth panel of Fig. 6.2) alongside five lens-source cross-correlation measurements (third and fourth panels). The lens-source pairs have been chosen to have no redshift overlap, thus their corresponding cross-correlations vanish in the absence of magnification. From our five lens-source pairs we obtain SNR = 11 for the scale cuts discussed in §6.2 and shown in Fig. 6.2. When combined with DESI Y1 BAO these measurements yield

$$S_8 = 0.715 \pm 0.046 \text{ (stat.)} \pm 0.050 \text{ (sys.)}$$
$$\sigma_8 = 0.720 \pm 0.047 \text{ (stat.)} \pm 0.050 \text{ (sys.)}$$

assuming a linear bias model and fixed number count slope s_μ. The systematic (sys.) error quoted above comes from beyond-linear biasing and is estimated following §6.3. Our best-fit predictions are shown in the bottom three panels of Fig. 6.2 (dot dashed). We find a best-fit $\chi^2_{\mu g,gg} = 33.8$ and $\chi^2_{\rm BAO} = 12.7$ for $38 + 12$ data points. We estimate the Bayesian PTE to be 40%.

We only analyze far-separated lens-source pairs so that the only "cosmological" correlation is expected to arise from magnification. However, imaging systematics can potentially correlate bins that are disjoint in redshift; this will need to be studied and characterized for future analyses. For the present work, the modest constraining power together with the tests performed by refs. [320, 511, 129], which found no evidence of residual imaging systematics, lends confidence in the robustness of these results.

6.5 Discussion and conclusions

We constrain the low redshift evolution of linear structure growth from the cross-correlation of DESI BGS imaging data with the latest CMB lensing maps from *Planck* and ACT. With the inclusion of BAO we find that σ_8 is consistent with CMB-conditioned ΛCDM predictions to within $\simeq 1.1\sigma$, while S_8 is slightly more discrepant ($\simeq 1.7\sigma$) due to the CMB preferring a higher Ω_m [475, 362, 278, 490, 184]. This is true both when analyzing the BGS data on their own or when jointly analyzed with the tomographic LRG measurements from refs. [320, 511]. Our tomographic analysis suggest that the rate of growth at redshifts $z \lesssim 1$ is also consistent with a CMB-conditioned ΛCDM cosmology. Recent CMB lensing cross-correlations with higher redshift tracers such as *unWISE* galaxies [223] or *Quaia* quasars [469] suggest that this trend continues to higher redshifts (see Fig. 6.3).

Since low-z linear growth is relatively insensitive to time-dependent DE (Fig. 6.1) our findings do not contradict the preference for evolving DE by the DESI Y3 BAO data [184]. Our tomographic measurements provide a complementary approach to BAO observations for constraining DE properties beyond its equation of state, such as DE interactions or modified gravity alternatives to DE.

We note that our results are in tension with those recently obtained from the joint analysis of the 3D clustering of BOSS galaxies and their cross-correlation with CMB lensing [132].

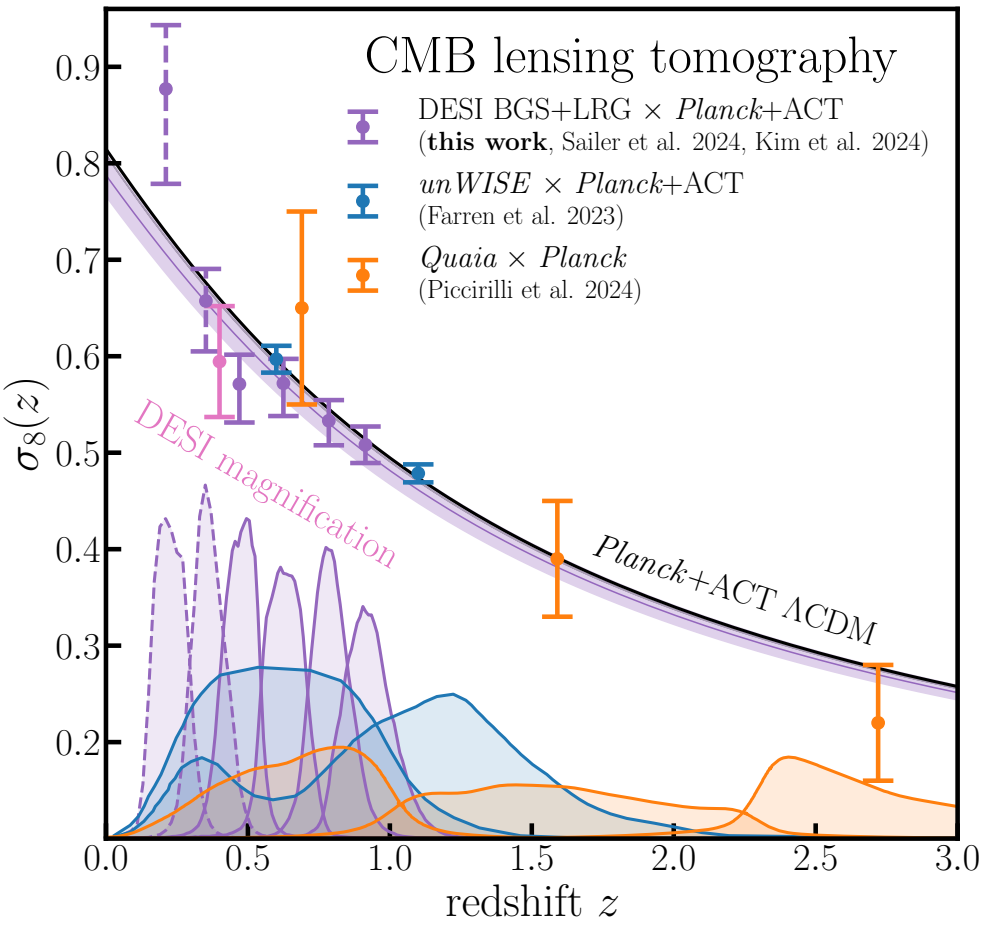


Figure 6.3: Structure growth from the joint analysis of BAO and CMB lensing cross-correlations with DESI BGS (dashed purple) and LRG [320, 511] (solid purple) imaging data, *unWISE* galaxy catalogs [223] (blue) and *Quaia* quasars [469] (orange). The shaded purple band represents the $\pm 1\sigma$ credible interval for our combined BGS & LRG fit (bottom row of Table 6.3). We plot our magnification constraint in pink and the *Planck*+ACT ΛCDM evolution of $\sigma_8(z)$ in black [592, 503, 362]. The shaded regions on the bottom of the plot are proportional to the redshift distribution of each galaxy sample.

We suspect that this tension is driven by unresolved systematics in the BOSS data and we leave an in-depth investigation to future work. In the near future, we will present a joint analysis of our measurements with the 3D clustering of DESI galaxies [390] which will be more directly comparable to ref. [132].

We analyze galaxy magnification as a new probe of structure growth. Given the modest precision of our magnification measurement, we adopt a linear bias model which yields

a $\simeq 9\%$ constraint[5] of σ_8 that is consistent with CMB-conditioned ΛCDM predictions to within $\simeq 1\sigma$. The precision obtainable from galaxy magnification is limited by (the cosmic variance associated with) the source sample auto-correlation, which is $\simeq 10\times$ larger than the lens-source cross-correlation. The SNR of the magnification signal may be improved by tuning the photometric selection to maximize the source number count slope while simultaneously decreasing the source linear bias. Finally, a joint analysis of magnification and CMB lensing cross-correlations could be used to calibrate magnification bias independently of finite difference estimates, which we leave to future work.

Acknowledgments

We thank William Coulton, Gerrit Farren, Fiona McCarthy, Blake Sherwin and David Spergel for useful discussions during the preparation of this manuscript. NS and SF are supported by Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. SC acknowledges the support of the National Science Foundation at the Institute for Advanced Study through NSF/PHY 2207583. JK and MM acknowledge support from NSF grants AST-2307727 and AST-2153201. MM additionally acknowledges NASA grant 21-ATP21-0145.

[5]We add the statistical and systematic errors in quadrature.

Chapter 7

Baryon distribution I: lens-hardened patchy screening estimators

This chapter was originally published as

> Noah Sailer, Boryana Hadzhiyska, and Simone Ferraro. "Bias hardened estimators of patchy screening profiles". In: *arXiv e-prints* (June 2025). arXiv: 2506.17217 [astro-ph.CO]

Detecting anisotropic screening of the cosmic microwave background (CMB) holds the promise of revealing the distribution of gas in the Universe, characterizing the complex processes of galaxy formation and feedback, and studying the epoch of reionization. Estimators for inhomogeneous screening, including some recently proposed small-scale (stacked) estimators, are quadratic or higher order in the CMB temperature or polarization fields and are therefore subject to contamination from CMB lensing. We review the origin of this lensing bias and show that, when stacking on *unWISE* galaxies, the expected lensing bias dominates the signal if left unmitigated. Hardening techniques that null the lensing bias have been proposed for standard quadratic estimators, whereas only approximate methods have been proposed for stacked estimators. We review these techniques and apply the former to stacked estimators, presenting several strategies (including the optimal strategy) to null lensing contamination when stacking on any large-scale structure (LSS) tracer.

7.1 Introduction

Studying the diffuse gas in the Universe is of paramount importance to characterize the complex processes behind galaxy formation and feedback. Moreover, the ionized gas amounts to almost 16% of the mass in the Universe, making understanding its location essential for the cosmological interpretation of weak lensing measurements from current and future surveys. High resolution and low noise observations of the cosmic microwave background (CMB) make it possible to use the CMB as a "backlight" to study the gas at low redshift. Several

effects imprint the properties of the gas on the small-scale CMB. In particular, two of the Sunyaev-Zel'dovich (SZ) effects [572] are proportional to the gas density: the kinematic SZ (kSZ) effect is the Doppler shift of a photon's energy due to a galaxy's peculiar velocity, and is proportional to both the gas density and the large scale peculiar velocity. The patchy screening (or "blurring" SZ) effect is the suppression of the CMB anisotropies by Thomson scattering with free electrons, and is proportional to both the gas density and the large-scale CMB temperature or polarization fluctuations. While kSZ has been used to measure the gas properties for quite some time [523, 525, 265], patchy screening has been less studied in the literature: [202] showed that quadratic estimators similar to those of CMB lensing can be used to reconstruct the optical depth τ field, while [568] showed that CMB lensing can be an important contribution and suggested a method of lens-hardening to mitigate that problem[1]. Later forecasts show that the τ field can be used to characterize the epoch of reionization [250] and the gas around late-time galaxies [507].

Following similar work for estimators of the kSZ effect [347, 523] and cluster lensing [292], stacking estimators for τ from patchy screening were proposed by refs. [536, 154]. These stacked estimators are quadratic (or higher order) in the underlying CMB temperature field, and therefore are sensitive to the effect of CMB lensing. In this chapter, we explain the origin of the lensing contamination to stacked estimators and discuss three possible mitigation strategies with varying levels of assumptions made about the lensing field.

The remainder of the chapter is organized as follows: in §7.2 we briefly review several patchy screening estimators. In §7.3 we present three approaches to harden these estimators against lensing, which we validate using simulations in §7.4. We comment on the importance of lensing biases for *unWISE* galaxies in §7.5 and conclude in §7.6. In the next chapter we present a self-consistent comparison of kSZ- and patchy screening-inferred electron profiles taking into account the lensing bias explored here.

Notation: Since we primarily work on small scales, we adopt the flat-sky approximation throughout. We let $\boldsymbol{r}$ and $\boldsymbol{\ell}$ correspond to angular positions and momenta respectively. We adopt the following convention for Fourier transforms

$$T_{\boldsymbol{\ell}} = \int d^2\boldsymbol{r}\, e^{-i\boldsymbol{\ell}\cdot\boldsymbol{r}}\, T(\boldsymbol{r}) \quad \text{and} \quad T(\boldsymbol{r}) = \int_{\boldsymbol{\ell}} e^{i\boldsymbol{\ell}\cdot\boldsymbol{r}}\, T_{\boldsymbol{\ell}}$$

where $\int_{\boldsymbol{\ell}} = \int d^2\boldsymbol{\ell}/(2\pi)^2$. Throughout we let $\langle\cdots\rangle$ denote a full ensemble average and $\langle\cdots\rangle'$ an average over primary CMB realizations in the presence of a fixed late-time large scale structure (LSS) realization.

7.2 Patchy screening estimators

Thomson scattering of CMB photons in and out of the line of sight damps primary CMB temperature fluctuations according to:

$$T(\boldsymbol{r}) = e^{-\tau(\boldsymbol{r})} T_0(\boldsymbol{r}) \simeq T_0(\boldsymbol{r}) - \tau(\boldsymbol{r}) T_0(\boldsymbol{r}), \tag{7.1}$$

[1] Ref. [568] also derives a lensing estimator that is unbiased to patchy screening.

where T is the observed CMB temperature fluctuation, T_0 is the primary CMB and $\tau \ll 1$ is the optical depth field. The modulation of the primary CMB by the optical depth produces (among other things) a non-zero $\langle TT\tau\rangle$ bispectrum. The τ-induced bispectrum can equivalently be phrased in terms of a *linear response* function: when averaging over primary CMB realizations in the presence of a *fixed* τ, the covariance $\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle'$ acquires "off diagonal" (or anisotropic) contributions proportional to the optical depth [202]:

$$\begin{aligned} &\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle' = (2\pi)^2 \delta^D_{\boldsymbol{L}} C_\ell^{TT} + f^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\, \tau_{\boldsymbol{L}} + \mathcal{O}(\tau^2) \\ &\text{with} \quad f^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = -\left[C_\ell^{TT} + C^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right], \end{aligned} \tag{7.2}$$

where f^τ is the *linear response* of the CMB to patchy screening. To relate f^τ to the bispectrum, multiply both sides of Eq. (7.2) by $\tau_{-\boldsymbol{L}}$ and perform a full ensemble average, yielding $f^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}} \tau_{-\boldsymbol{L}}\rangle / \langle \tau_{\boldsymbol{L}} \tau_{-\boldsymbol{L}}\rangle$.

Optimal quadratic estimator

To reconstruct patchy screening from the observed CMB maps ref. [202] proposed a quadratic estimator (QE) $\hat{\tau}$, which takes the general form

$$\hat{\tau}_{\boldsymbol{L}} = \int_{\boldsymbol{\ell}} F^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}, \tag{7.3}$$

for a set of weights $F^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$. To ensure that the estimator has unit response to patchy screening (i.e. $\langle\hat{\tau}\rangle' = \tau$) we enforce the normalization constraint $\int_{\boldsymbol{\ell}} F^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = 1$. In particular, the minimum variance weights satisfying the normalization constraint are given by [300]

$$F^{\tau,\mathrm{MV}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{f^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C_\ell^{\mathrm{tot}} C^{\mathrm{tot}}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \Bigg/ \int_{\boldsymbol{\ell}} \frac{\left(f^\tau_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\right)^2}{2C_\ell^{\mathrm{tot}} C^{\mathrm{tot}}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}, \tag{7.4}$$

where C_ℓ^{tot} is the total observed temperature power spectrum.

Stacked estimators

Since the ionized gas surrounding galaxies has a $\mathcal{O}(1)$ arcmin extent, the optical depth field is primarily composed of small scale modes. On the other hand, Silk damping largely erases these modes from the primary CMB. Motivated by this insight, ref. [536] proposed using a long-short split to estimate the optical depth field. To do so we construct low- and high-pass filtered CMB maps with isotropic filters W^l_ℓ and W^s_ℓ respectively, which we colloquially refer to as long and short filters. The long map $T^l_{\boldsymbol{\ell}} \equiv W^l_\ell T_{\boldsymbol{\ell}} \simeq T^0_{\boldsymbol{\ell}}$ is predominately sensitive to the primary CMB, while the short map $T^s_{\boldsymbol{\ell}} \equiv W^s_\ell T_{\boldsymbol{\ell}} \simeq -\int_{\boldsymbol{\ell}'} \tau_{\boldsymbol{\ell}'} T^0_{\boldsymbol{\ell}-\boldsymbol{\ell}'}$ is predominantly sensitive to the modulation of the primary CMB by patchy screening. Ref. [536] uses this to construct an optical depth map estimator from the ratio $-T_s(\boldsymbol{r})/T_l(\boldsymbol{r})$ which under idealistic conditions is significantly less noisy than the optimal QE on small scales.

Alternatively, we can use the long-short split to construct a real space QE: $-T_l(\boldsymbol{r})T_s(\boldsymbol{r})/\langle T_l^2\rangle$. While the real space QE is suboptimal (i.e. its weights differ from Eq. 7.4) it has the advantage of being computationally simpler than a full sky QE formalism, is readily amenable to public stacking codes[2] [525] and can be easily adapted to perform oriented stacking [following e.g. 63, 19, 361, 360]. When stacking the real space QE on a set of galaxy positions $\boldsymbol{r}_i$ we obtain a *real space stacked estimator* $\widehat{\mathcal{T}}$ for the optical depth profile, which can be (schematically) interpreted as the average optical depth at a displacement $\boldsymbol{r}$ away from a galaxy in the sample:

$$\boxed{\widehat{\mathcal{T}}(\boldsymbol{r}) = \frac{-1}{\langle T_l^2\rangle N_g}\sum_{i=1}^{N_g} T_l(\xi\boldsymbol{r}+\boldsymbol{r}_i)T_s(\boldsymbol{r}+\boldsymbol{r}_i),} \tag{7.5}$$

where we introduced a number ξ that specifies the offset between where T_l and T_s are evaluated. When comparing with simulations in §7.4 we consider only two cases: $\xi = 0$ or 1, corresponding to the long map being evaluated at the galaxy center or a distance r away from the galaxy center respectively. When r is smaller than the scale associated with the long-short split $T_l(\boldsymbol{r})$ is roughly a constant and the estimator's response to patchy screening only weakly varies with $\mathcal{O}(1)$ changes in ξ. We explicitly show this in §7.3.

7.3 Lens-hardened stacked estimators

Weak gravitational lensing of the CMB effectively shifts CMB pixels from position $\boldsymbol{r}$ to $\boldsymbol{r}+\boldsymbol{\nabla}\phi(\boldsymbol{r})$, where ϕ is the lensing potential [342]. To linear order this shift is equivalent to adding the term $\boldsymbol{\nabla}\phi\cdot\boldsymbol{\nabla}T_0$ to Eq. (7.1). Following the same logic as §7.2, the lensing term induces its own off-diagonal contributions to the covariance $\langle T_{\boldsymbol{\ell}}T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle' = f^{\tau}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\tau_{\boldsymbol{L}} + f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\kappa_{\boldsymbol{L}} + \cdots$, where $\kappa_{\boldsymbol{L}} = L^2\phi_{\boldsymbol{L}}/L^2$ is the CMB lensing convergence and

$$f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{2\boldsymbol{L}}{L^2}\cdot\left[\boldsymbol{\ell}C_\ell^{TT} + (\boldsymbol{L}-\boldsymbol{\ell})C^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right] \tag{7.6}$$

is the associated linear response to lensing [300]. Since the stacked estimator is quadratic (or higher order, see Appendix E.1) in the CMB temperature it is subject to biases from CMB lensing [568]. Here we explore three strategies to mitigate CMB lensing biases to stacked patchy screening estimators with varying levels of assumptions made about the lensing field.

Field level hardening

Rather than the real space stacked estimator (Eq. 7.5), one could in principle use the minimum variance QE (Eqs. 7.3 & 7.4) to estimate the optical depth profile by averaging

[2] e.g. https://github.com/EmmanuelSchaan/ThumbStack

$\hat{\tau}^{\rm MV}(\boldsymbol{r}+\boldsymbol{r}_i) = \int_{\boldsymbol{L}} e^{i\boldsymbol{L}\cdot(\boldsymbol{r}+\boldsymbol{r}_i)}\hat{\tau}^{\rm MV}_{\boldsymbol{L}}$ over galaxy positions. Following the discussion above, the minimum variance patchy screening estimator acquires a bias from lensing $\langle\hat{\tau}^{\rm MV}_{\boldsymbol{L}}\rangle' = \tau_{\boldsymbol{L}} + R^{\tau\kappa}_L\kappa_{\boldsymbol{L}}$, where we introduced the response function

$$R^{\tau\kappa}_L = \int_{\boldsymbol{\ell}} F^{\tau,{\rm MV}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}. \tag{7.7}$$

To remove this bias ref. [568] constructs a minimum variance CMB lensing estimator $\hat{\kappa}^{\rm MV}$, which can be obtained by relabeling $\tau \to \kappa$ in Eqs. (7.3) and (7.4). The minimum variance lensing estimator has a bias proportional to the optical depth: $\langle\hat{\kappa}^{\rm MV}_{\boldsymbol{L}}\rangle' = \kappa_{\boldsymbol{L}} + R^{\kappa\tau}_L\tau_{\boldsymbol{L}}$. By taking an appropriate linear combination of the minimum variance τ and κ estimators, ref. [568] showed that one can subtract off the lensing bias while retaining unit response to the optical depth [see also 434, 436, 446, 519, 435]. Alternatively one can obtain a bias hardened τ estimator by imposing $\int_{\boldsymbol{\ell}} F^{\tau}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = 0$ in addition to the normalization constraint and performing a constrained minimization [517]. Both approaches yield the same result for the optimal lens-hardened weights:

$$\boxed{F^{\tau,{\rm BH}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{F^{\tau,{\rm MV}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - R^{\tau\kappa}_L F^{\kappa,{\rm MV}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{1 - R^{\tau\kappa}_L R^{\kappa\tau}_L}.} \tag{7.8}$$

With these modified weights, the average of $\hat{\tau}^{\rm BH}(\boldsymbol{r}+\boldsymbol{r}_i)$ over galaxy positions is an unbiased estimate of the optical depth profile.

It is straightforward to generalize this procedure to simultaneously harden against both lensing and unclustered point sources [435]. A tedious but straightforward calculation shows that the resulting bias hardened QE $\hat{\tau}^{\rm BH}$ can be expressed as a series of filtering and Fourier transforms:

$$\hat{\tau}^{\rm BH}_{\boldsymbol{L}} = A^{\rm wf,\ ivar}_L \mathcal{F}_{\boldsymbol{L}}\Big[T^{\rm wf}\times T^{\rm ivar}\Big] + A^{\Delta{\rm wf,\ ivar}}_L \mathcal{F}_{\boldsymbol{L}}\Big[T^{\rm wf}\times\left(\nabla^2 T\right)^{\rm ivar} - \left(\nabla^2 T\right)^{\rm wf}\times T^{\rm ivar}\Big] + A^{\rm ivar,ivar}_L \mathcal{F}_{\boldsymbol{L}}\Big[\left(T^{\rm ivar}\right)^2\Big], \tag{7.9}$$

where $T^{\rm wf}_{\boldsymbol{\ell}} = C^{TT}_\ell\, T_{\boldsymbol{\ell}}/C^{\rm tot}_\ell$ and $T^{\rm ivar}_{\boldsymbol{\ell}} = T_{\boldsymbol{\ell}}/C^{\rm tot}_\ell$ are Wiener- and inverse variance-filtered maps respectively, $(\nabla^2 T)_{\boldsymbol{\ell}} = -\ell^2 T_{\boldsymbol{\ell}}$ is a gradient map and $\mathcal{F}_{\boldsymbol{L}}[\cdots]$ indicates a Fourier transform of a real-space map. The coefficients $A^{\rm wf,\ ivar}$, $A^{\Delta{\rm wf,\ ivar}}$ and $A^{\rm ivar,\ ivar}$ (see Fig. 7.1) are given by

$$\begin{aligned}
\frac{A^{\rm wf,ivar}_L}{N^{\tau,{\rm BH}}_L} &= \frac{r^{\tau s}_L r^{\kappa s}_L - (r^{\kappa s}_L)^2}{r^{\kappa\kappa}_L r^{ss}_L} - \frac{r^{\tau\kappa}_L}{r^{\kappa\kappa}_L} - 1,\\
\frac{A^{\Delta{\rm wf,ivar}}_L}{N^{\tau,{\rm BH}}_L} &= \frac{1}{L^2}\left[\frac{r^{\tau s}_L r^{\kappa s}_L}{r^{\kappa\kappa}_L r^{ss}_L} - \frac{r^{\tau\kappa}_L}{r^{\kappa\kappa}_L}\right],\\
\frac{A^{\rm ivar,ivar}_L}{N^{\tau,{\rm BH}}_L} &= \frac{r^{\tau\kappa}_L r^{\kappa s}_L}{r^{\kappa\kappa}_L r^{ss}_L} - \frac{r^{\tau s}_L}{r^{ss}_L},\\
\left(N^{\tau,{\rm BH}}_L\right)^{-1} &= r^{\tau\tau}_L\left[1 - \frac{(r^{\tau\kappa}_L)^2}{r^{\tau\tau}_L r^{\kappa\kappa}_L} - \frac{(r^{\tau s}_L)^2}{r^{\tau\tau}_L r^{ss}_L} - \frac{(r^{\kappa s}_L)^2}{r^{\kappa\kappa}_L r^{ss}_L} + 2\frac{r^{\tau\kappa}_L r^{\tau s}_L r^{\kappa s}_L}{r^{\tau\tau}_L r^{\kappa\kappa}_L r^{ss}_L}\right],\\
r^{XY}_L &= \int_{\boldsymbol{\ell}} \frac{f^X_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\, f^Y_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\rm tot}_\ell C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}.
\end{aligned} \tag{7.10}$$

Here the superscript "s" refers to point sources, for which the corresponding response function is scale invariant (i.e. $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = 1$). The quantity r_L^{XY} can be thought of as the unnormalized response of the MV X estimator to Y, while the properly normalized response is given by $R_L^{XY} = r_L^{XY}/r_L^{XX}$. Note that the simplified scenario of only hardening against lensing (Eq. 7.8) can be obtained by setting $r^{\tau s} = r^{\kappa s} = 0$ in the equation above.

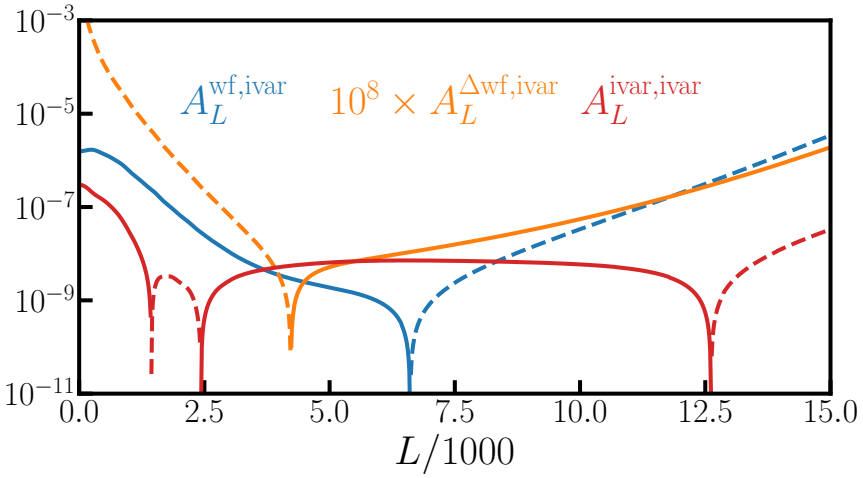


Figure 7.1: The three coefficients appearing in Eq. (7.9) for the experimental setup discussed in §7.4, with the color labeling the coefficient. Solid (dashed) lines correspond to positive (negative) values.

Stacked hardening

Field level hardening (§7.3) removes the leading order lensing bias via an internal quadratic estimate of the lensing convergence. On the small scales relevant for patchy screening measurements going beyond the QE will be more optimal and accurate for both estimating the patchy screening signal and removing any lensing contamination [see e.g. 298, 292, 408, 90, 510]. While in this chapter we restrict ourselves to "stacked estimates" of patchy screening (e.g. Eq. 7.5), in this section we will assume that we are given a general *unbiased estimate* $\hat{\kappa}(\boldsymbol{r})$ of the convergence field which could come from e.g. a QE, a maximum likelihood approach, or external LSS data as a lensing template. Below we present an analogous method to lens-harden the real space stacked estimator $\hat{\mathcal{T}}$ (Eq. 7.5) using this unbiased estimate of the convergence field, which we call *stacked hardening*.

To connect stacked hardening to the previous section we simply take $R_L^{\kappa\tau} \to 0$ in Eq. (7.8). That is, the lensing estimator is presumed to be unbiased to patchy screening. In this limit there is a simple interpretation of how one estimates the lensing bias to the stacked minimum variance QE: filter the convergence estimate by the response function $R_L^{\tau\kappa}$, transform back to real space and stack on the galaxy positions.

We generalize this procedure to the real space stacked estimator in Appendix E.1, where we compute the linear response $\mathcal{R}^{\mathcal{T}\kappa}_{\boldsymbol{L}}(\boldsymbol{r})$ of the stacked estimator to the lensing convergence. Following the same recipe as before we obtain an estimator for the lensing bias by filtering the convergence estimator and averaging over galaxy positions:

$$\boxed{\begin{aligned}&\text{lensing bias}\,(\boldsymbol{r}) = \frac{1}{N_g}\sum_{i=1}^{N_g}\hat{\kappa}^F(\boldsymbol{r}_i,\boldsymbol{r}),\\ &\hat{\kappa}^F(\boldsymbol{r}_i,\boldsymbol{r}) = \int_{\boldsymbol{L}} e^{i\boldsymbol{L}\cdot\boldsymbol{r}_i}\mathcal{R}^{\mathcal{T}\kappa}_{\boldsymbol{L}}(\boldsymbol{r})\hat{\kappa}_{\boldsymbol{L}},\\ &\mathcal{R}^{\mathcal{T}\kappa}_{\boldsymbol{L}}(\boldsymbol{r}) = \frac{-1}{\langle T_l^2\rangle}\int_{\boldsymbol{\ell}} e^{i\left(\boldsymbol{L}-(1-\xi)\boldsymbol{\ell}\right)\cdot\boldsymbol{r}} W^s_{|\boldsymbol{L}-\boldsymbol{\ell}|}W^l_\ell f^\kappa_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}.\end{aligned}} \tag{7.11}$$

Note that this procedure can be straightforwardly adapted to oriented stacking by taking $\hat{\kappa}^F(\boldsymbol{r}_i,\boldsymbol{r}) \to \hat{\kappa}^F(\boldsymbol{r}_i,\mathsf{R}_i\,\boldsymbol{r})$, where R_i is a rotation matrix for the i'th galaxy. Additionally, one can easily generalize this approach to more optimally weight the galaxy sample.

For the case of an isotropic profile or when stacking with random orientations the response function simplifies to a function of two scalars (L, r) rather than two vectors $(\boldsymbol{L},\boldsymbol{r})$. In Fig. 7.2 we plot the isotropic response to lensing (Maya blue) and the isotropic response to patchy screening (Amaranth pink) for the analysis setup described in §7.4. In the top and bottom panels we fix $r = 2$ and 8 arcmin respectively, while solid and dashed lines correspond to T_l and T_s being evaluated at offset or identical positions in Eq. (7.5).

We emphasize that when T_l and T_s are evaluated at the same position the response to lensing (dashed blue) is confined to a smaller range of scales than when T_l is evaluated at the galaxy center (solid blue), leaving the stacked estimator nearly unbiased on scales $L \gtrsim 3000$ for $r \simeq 2$ arcmin for the setup discussed in §7.4. We therefore recommend that when using a stacked estimator one evaluates T_l and T_s at the same position. The residual lensing bias can be subtracted following Eq. (7.11) or further suppressed (but not exactly nulled) by high-pass filtering the stacked map, which "manually" sets $\mathcal{R}^{\mathcal{T}\tau}_L = \mathcal{R}^{\mathcal{T}\kappa}_L = 0$ below some cutoff scale $L < L_c$.

As is evident from Fig. 7.2 the stacked response to patchy screening is very well approximated by the product of W^s_L and the beam B_L (see §7.4 for a discussion of the beam) up to a geometric J_0 factor, implying that the stacked estimator reconstructs a high pass filtered $\tau(\boldsymbol{r})$ convolved with the beam to a very good approximation [536].

Mean hardening

Field-level (§7.3) and stacked hardening (§7.3) require accurate small scale estimates of the lensing field. This can be seen from Fig. 7.2, where the response to lensing can be non-negligible for $L \lesssim 8000$ (4000) when T_l and T_s are evaluated at offset (identical) positions. Faithfully reconstructing lensing on these scales with QEs is challenging due to contamination from extragalactic foregrounds [597, 227] and higher-order effects such as post-Born corrections [74]. In particular the ACT DR6 lensing map [493] is only reliable to $L_{\rm max} \sim 3000$,

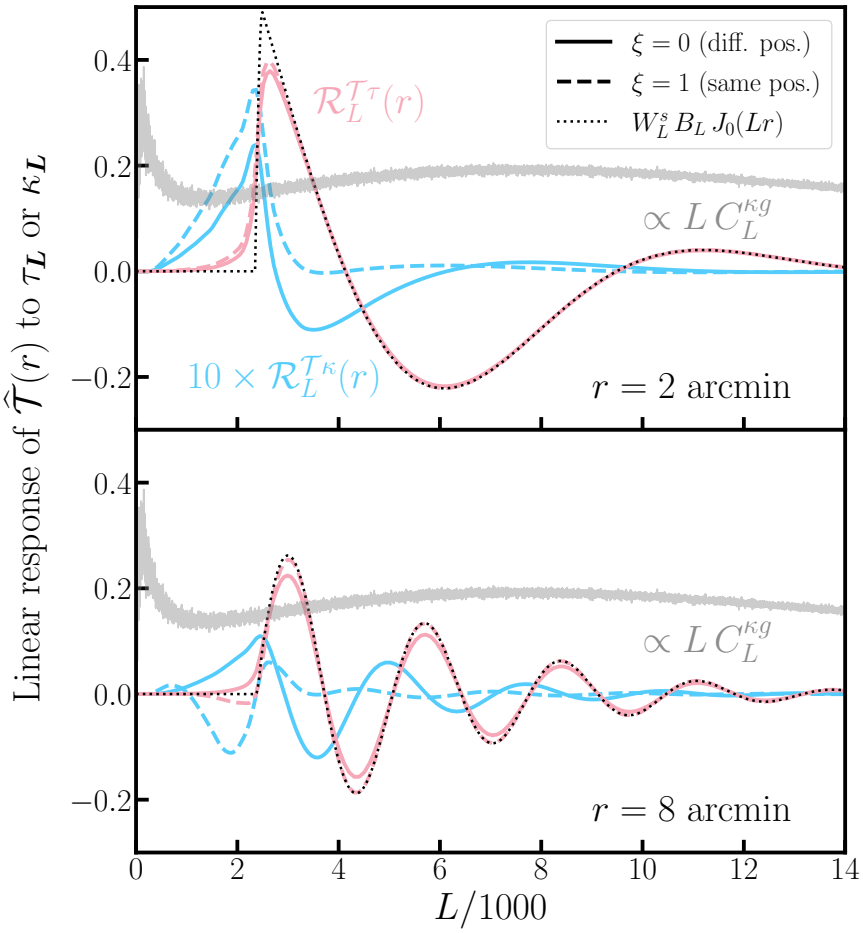


Figure 7.2: The isotropic response to lensing (Maya blue) and to patchy screening (Amaranth pink) for the experimental setup described in §7.4. In the top and bottom panels we fix $r = 2$ and 8 arcmin respectively, while solid and dashed lines correspond to T_l and T_s being evaluated at offset or identical positions. Note that the response to lensing is multiplied by 10. In light gray we plot the CMB lensing cross-correlation measured from the AbacusSummit simulations for a DESI LRG-like set of mock galaxies [270]. The dotted line shows the product of the high-pass filter, the beam B_L and the Bessel function of the first kind $J_0(Lr)$.

while in principle one could push to $L_{\rm max} = 6000$ for their baseline reconstruction scales ($\ell_{\rm max} = 3000$). Furthermore, on small scales the QE quickly becomes suboptimal. While more optimal small-scale estimators have been proposed [e.g. 298, 292, 408, 90, 510] these estimators are still subject to biases (primarily in temperature [492]) from extragalactic foregrounds.

In the absence of an accurate high resolution CMB lensing map one can instead estimate the lensing bias from an estimate of the large scale CMB lensing-galaxy cross-correlation,

which we call *mean hardening*. For simplicity we only consider the case of stacking with random orientations so that $\mathcal{R}^{\mathcal{T}^\kappa}_{\boldsymbol{L}}(\boldsymbol{r})$ simplifies to a function of the scalars (L, r). In Appendix E.1 we show that the average lensing bias to the real space stacked estimator is

$$\boxed{\langle \text{lensing bias}\,(r)\rangle = \int \frac{L\,dL}{2\pi}\mathcal{R}^{\mathcal{T}^\kappa}_{L}(r)C^{\kappa g}_{L},} \tag{7.12}$$

where $C^{\kappa g}_L$ is the CMB lensing-galaxy cross-correlation. Given an accurate measurement of $C^{\kappa g}_L$ on "large scales" (e.g. $L \lesssim 3000$) one can predict the small scale power from e.g. a power law extrapolation, halo model or hydrodynamic simulation. While this approach comes with the cost of introducing a parametric model in the highly nonlinear regime, it has the benefit of being less sensitive to foreground contamination and higher order effects in the lensing reconstruction.

7.4 Comparison with simulations

We test both stacked (§7.3) and mean (§7.3) hardening using a CMB lensing map and DESI LRG-like set of mock galaxies [for details see 270] built from the AbacusSummit simulations [269]. We adopt the following filters for the long and short maps respectively:

$$\begin{aligned} W^l_\ell &= \begin{cases} 1 & \ell < 2000 \\ \cos\left(\frac{\pi}{300}(\ell - 2000)\right) & \text{otherwise} \\ 0 & \ell > 2150 \end{cases} \\ W^s_\ell &= \begin{cases} 0 & \ell < 2350 \\ \sin\left(\frac{\pi}{300}(\ell - 2350)\right) & \text{otherwise} \\ 1 & \ell > 2500. \end{cases} \end{aligned} \tag{7.13}$$

We assume that the short map is convolved with a Gaussian beam $B_\ell = \text{Exp}[-(\theta_{\text{FWHM}}\ell)^2/16\ln 2]$ with a full width half max (FWHM) of 1.4 arcmin while the long map is assumed to be beam-deconvolved. To include the beam-convolution of the short map we substitute $W^s_\ell \to B_\ell W^s_\ell$ when computing the stacked response function (Eq. 7.11).

In the next chapter we compute the profile obtained when lensed CMB maps with no patchy screening signal are stacked on LRG-like galaxies. We plot the result in Fig. 7.3 (light catawba) for the case where T_l and T_s are evaluated at offset ($\xi = 0$, tracing the solid curve) or identical positions ($\xi = 1$, tracing dashed). We use this measurement as a baseline from which to compare the predicted lensing bias from the stacked (catawba) and mean (black lines) hardening approaches. We implement stacked hardening using the simulated convergence map while for mean hardening we use the cross-correlation measured from the

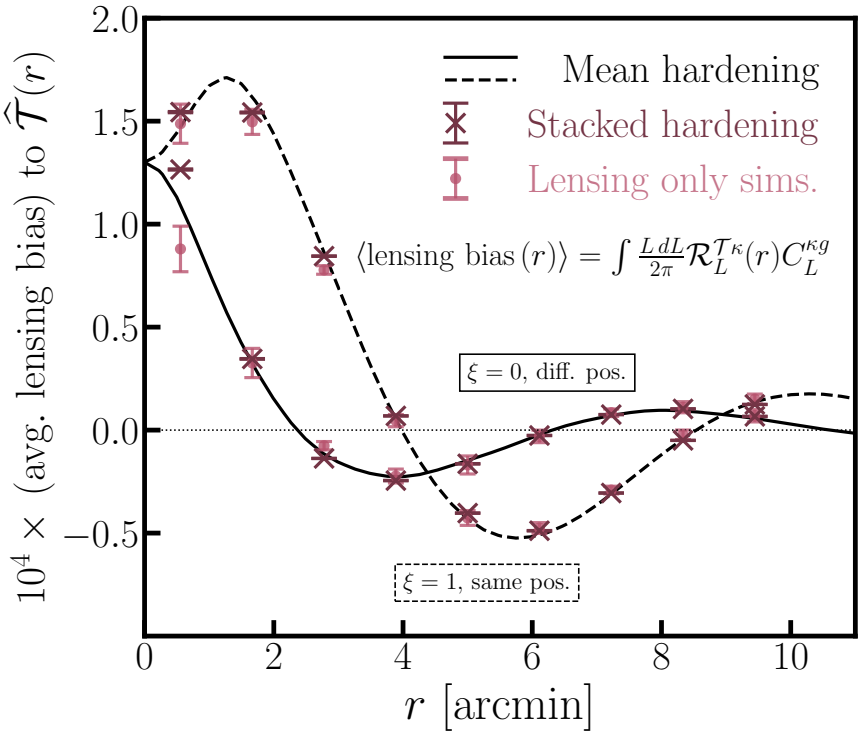


Figure 7.3: The predicted lensing bias from stacked (catawba) and mean (black lines) hardening for a DESI LRG-like set of mock galaxies built from the AbacusSummit simulations. We find excellent agreement between both hardening approaches and the signal obtained when stacking on lensed CMB maps with no patchy screening signal (light catawba). Dashed and solid curves correspond to T_l and T_s being evaluated at identical or offset positions respectively. The profiles have been binned into 9 annuli whose 10 radial edges are linearly spaced between $r = 0$ and 10 arcmin.

simulation (plotted in light gray in Fig. 7.2). With the exception of the first radial bin[3] we find excellent agreement between all three approaches.

7.5 Comments on ACT × unWISE

Refs. [536, 154] proposed a modification to the stacked estimator, called the "signed and thresholded" estimator, in which the long map T_l appearing the numerator of Eq. (7.5) is replaced by its sign and cutouts with $|T_l(\boldsymbol{r}_i)| < 40$ μK are discarded. This is equivalent to replacing T_l with the weight $\mathcal{W}[T_l] \equiv \text{Sign}[T_l]H(|T_l| - T_c)$, where $H(x)$ is the Heaviside

[3]This is due to two effects: the bandpower window function has the largest impact on the comparison of the smooth to binned curves for the first radial bin, and in practice [270] effectively applies a 0.5 arcmin smoothing to the galaxies when doing stacked hardening.

step function and $T_c = 40$ μK. To ensure that $\widehat{\mathcal{T}}$ is correctly normalized one must also substitute $\langle T_l^2 \rangle \to \langle \mathcal{W}[T_l] T_l \rangle$ in the denominator of Eq. (7.5). In Appendix E.1 we show that the linear response of this more general estimator to lensing (or any other extragalactic foreground) is identical to the $\mathcal{W}[T] = T$ case[4] provided that $\mathcal{W}[T]$ has a Fourier transform, which we compute for the signed and thresholded case in Appendix E.2. Thus, the leading order extragalactic foreground bias to the signed and thresholded estimator is identical to the standard stacked estimator (Eq. 7.5).

As an alternative to lens-hardening (§7.3) one can suppress (but not precisely null) the lensing bias by filtering the stacked estimator to downweight modes with large response to lensing [154, see also §7.3]. While this approach is formally suboptimal it has the advantage of being simpler to implement. We caution that since the clustering amplitude of a LSS tracer varies more slowly with the characteristic tracer mass than the mass itself, a lower-mass sample from e.g. LSST will presumably have a larger relative lensing bias than the more massive *unWISE* galaxies. Thus, the precise filtering required for this approach to remain unbiased depends non-trivially on the galaxy sample considered, whereas lens-hardening works "out of the box" for any LSS tracer.

In Fig. 7.4 we show the (unfiltered) profile obtained from stacking ACT DR6 temperature maps on *unWISE* blue and green galaxies [329, 526] using the long and short filters from §7.4 [154]. We note that the expected lensing bias (black) for the *unWISE* sample dominates the signal if left unmitigated, which we estimate as follows.

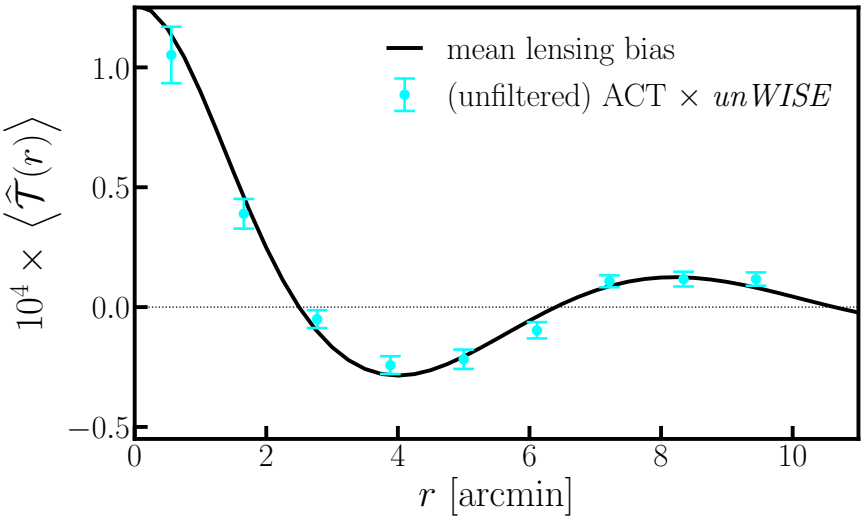


Figure 7.4: Mean lensing bias to the real space stacked estimator for the combined *unWISE* blue and green sample (solid black line). In blue we show the unfiltered measurement from [154].

[4]Heuristically, this can be "derived" by Taylor expanding [580] the weight function $\mathcal{W}[T]$.

We model the cross-correlation of the combined *unWISE* green and blue sample with CMB lensing using the Kaiser-Limber approximation [352, 317, 365], a linear bias model and the non-linear matter power spectrum from the Aemulus ν simulations [181]. We use the fiducial redshift distributions from [223] for the individual blue and green samples and add them to obtain the redshift distribution of the combined sample. We fix the cosmological parameters to their *Planck* 2018 [475] values and approximate the linear bias of the combined sample as the number-density weighted average of the fiducial blue and green linear biases from [330]. Likewise, we approximate the redshift evolution of the combined sample's magnification bias as the number-density weighted average of the (redshift independent) magnification biases measured in [329] for the individual blue and green samples. The numerical implementation of the $C_L^{\kappa g}$ calculation is identical to that used in [320, 511]. We assume that the short map is convolved with a 1.6 arcmin beam and evaluate T_l at the galaxy center, corresponding to $\xi = 0$ in Eq. (7.5).

We find excellent agreement between the unfiltered profile and expected lensing bias, and hence no detection of patchy screening, despite not fitting any free parameters in our estimation of $C_L^{\kappa g}$. We subtract off the expected lensing bias in each radial bin (including the bandpower window) to obtain a debiased profile. We model the electron profile as a Gaussian with peak height τ_0 and a FWHM of 2 arcmin that is then filtered by the response function ($\mathcal{R}^{\mathcal{T}\tau}$ with $\xi = 0$, see Fig. 7.2). When fit to the debiased profile we find $\tau_0 < 1.1 \times 10^{-4}$ at 68% confidence.

7.6 Discussion and conclusions

Both stacked (7.3) and mean hardening (7.3) are viable methods to mitigate the lensing induced bias with different assumptions made about the lensing field. The former is more data driven and straightforward to adapt to oriented stacking and or optimal galaxy weighting, while the latter is more parametric and less sensitive to extragalactic foreground contamination. In §7.4 we show that both approaches accurately reproduce the lensing bias found from simulated DESI LRG-like galaxies and lensed CMB maps.

Consistent with previous results [568], we find that without mitigation the lensing contamination dominates the measured profile, making bias hardening essential for a robust screening measurement. Applying mean hardening to previous measurements of ACT DR6 temperature maps stacked on *unWISE* galaxies yields an upper bound on the screening signal consistent with expectations from recent kSZ measurements [see the next chapter 270].

While in this chapter we have focused on the leading order lensing bias, additional biases can arise from higher-order lensing effects and other extragalactic foregrounds. We leave a thorough investigation of these contaminants to future work, noting that in Appendix §E.1 we show that the linear response of a generic class of stacked estimators (including the signed and thresholded estimator) to any extragalactic foreground is identical to that found for the "standard" stacked estimator (Eq. 7.5).

For realistic surveys with partial sky coverage the stacked estimator is also biased by a "mean field" which must be subtracted from Eq. (7.5) to obtain an unbiased profile. We compute the mean field bias in Appendix E.3 and explicitly show that it can be estimated by stacking on random positions that are uniformly distributed within the galaxy survey mask. To obtain a precise estimate of the mean field relative to the measured profile one should ensure that the number of random positions is significantly larger than the number of galaxies used when stacking. An application of optimal lens-hardening and mean field subtraction will be the subject of future work.

Acknowledgments

We thank William Coulton and Emmanuel Schaan for useful discussions during the preparation of this manuscript. N.S. and S.F. are supported by Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. B.H. thanks the Miller Institute for supporting her postdoctoral research.

Chapter 8

Baryon distribution II: a self-consistent comparison of the kSZ effect and patchy screening

This chapter was originally published as

> Boryana Hadzhiyska, Noah Sailer, and Simone Ferraro. "Mapping the gas density with the kinematic Sunyaev-Zel'dovich and patchy screening effects: a self-consistent comparison". In: *arXiv e-prints* (June 2025). arXiv: 2506.17379 [astro-ph.CO]

The secondary anisotropies of the cosmic microwave background (CMB) provide a wealth of astrophysical and cosmological information. Pairing measurements of the CMB temperature map obtained by DR5 of the Atacama Cosmology Telescope (ACT) with the large-scale structure imaging survey conducted by the Dark Energy Spectroscopic Instrument, DECaLS DR9, we investigate two effects that are sensitive to the gas density τ: kinematic Sunyaev-Zel'dovich (kSZ) and 'patchy screening' (also known as 'anisotropic screening'). In particular, we measure the stacked profiles around Luminous Red Galaxies (LRGs) at a mean redshift of $z \approx 0.7$. We detect the kSZ signal at 7.2σ, and we find a signal at $\sim 4.1\sigma$ for the patchy screening estimator, which is in excess relative to the kSZ signal. We attribute this excess to contamination from CMB lensing. We demonstrate the effect of lensing using N-body simulations, and we show that the screening signal is dominated by it. Accounting for lensing, our measurement places a 95% upper bound on the optical depth of the Extended DESI LRG sample of $\tau < 2.5\ 10^{-4}$ for a mean value of the sample of $\tau \approx 1.6\ 10^{-4}$. Furthermore, via hydro simulations, we show that the underlying optical depth signal measured by both effects (after removing the CMB lensing contribution) is in perfect agreement when adopting either a Compensated Aperture Photometry (CAP) filter or a high-pass filter. Consistent with previous measurements, we see evidence for excess baryonic feedback around DESI LRGs in the patchy screening measurement. Provided both

effects are measured with high signal-to-noise, one can measure the amplitude ratio between them, which is proportional to the root-mean-square velocity of the host halo sample, and place constraints on velocity-sensitive models such as modified gravity and phantom dark energy.

8.1 Introduction

A long-standing pursuit in astrophysics, mapping the distribution of baryons (gas) in the Universe holds the key to unraveling the formation and thermodynamic evolution of galaxy groups and clusters [448, 432, 397, 64, 422]. Moreover, their elusive connection to the underlying dark matter poses a significant challenge for cosmological analyses such as weak lensing studies, for which the effect of baryons remains one of the most major unknowns [40, 45]. Thus, inventorying the gas in low-density regions, known as the warm-hot intergalactic medium (WHIM), not only aids in resolving the so-called "missing baryon" problem [466, 238, 551, 113, 66], but also lets us disentangle astrophysical effects from manifestations of exotic cosmological models.

While X-ray studies have successfully characterized the properties of gas near the centers of massive clusters [e.g., 397, 508, 158, 173, 206, 638, 454], doing so for lower-mass systems and for gas located on the outskirts of groups and clusters has proved a more difficult task. Among the most powerful tools for studying the properties of gas is the measurement of the Sunyaev-Zel'dovich (SZ) effect, which comes in two flavors. The thermal SZ (tSZ) effect stems from the inverse Compton scattering of cosmic microwave background (CMB) photons by the hot thermal gas within groups and clusters. Its magnitude, directly proportional to the electron pressure integrated along the line-of-sight, offers insights into the thermodynamic conditions of these cosmic structures, especially close to their hot and dense centers. In contrast, the kinematic SZ (kSZ) effect results from the encounters between CMB photons and free electrons moving in bulk relative to the CMB rest-frame. Thus, the kSZ effect is well-suited for tracing the spatial distribution of baryons, even in the outskirts of galaxy groups and clusters (e.g., see [93, 423] for a review on the SZ effect).

Another potent probe for detecting the gas in groups and clusters is the "patchy screening" (also sometimes referred to as "anisotropic screening") effect, resulting from the Thomson scattering of CMB photons off free electrons on their path from the surface of last scattering to the observer. This leads to a characteristic damping of the primary CMB anisotropies that is spatially anisotropic [447, 605, 301, 195, 465, 261, 201, 202]. On small scales of order arcminutes, these anisotropies are proportional to the product between the optical depth integrated along the line-of-sight and the primary CMB temperature fluctuations.

$$\left(\frac{\Delta T}{T_{\rm CMB}}\right)_{\rm kSZ} \approx -\tau_{\rm gal}\left(\frac{v_r}{c}\right) \tag{8.1}$$

where $\tau_{\rm gal} = \int a n_e \sigma_T d\chi$ is the optical depth of the galaxy in question and v_r is the radial

velocity.

$$\left(\frac{\Delta T}{T_{\rm CMB}}\right)_{\rm kSZ,\ rms} \sim 0.1\mu{\rm K}\left(\frac{M_{\rm vir}}{10^{13}M_\odot/h}\right)\left(\frac{v_r^{\rm rms}}{300{\rm km/s}}\right) \tag{8.2}$$

$$\left(\frac{\Delta T}{T_{\rm CMB}}\right)_{\rm scr} \approx -\tau_{\rm gal}\left(\frac{\Delta T}{T_{\rm CMB}}\right)_{\rm primary} \tag{8.3}$$

Since for the primary CMB the rms fluctuation is $(\Delta T)_{\rm rms} \approx 110\mu$K, we obtain

$$\left(\frac{\Delta T}{T_{\rm CMB}}\right)_{\rm scr,\ rms} \sim 0.04\mu{\rm K}\left(\frac{M_{\rm vir}}{10^{13}M_\odot/h}\right) \tag{8.4}$$

Since velocity evolves slowly with redshift and is always close to $\sim$ 300 km/s at the redshifts of interest, we conclude that the patchy signal is always a factor of ≈ 24 smaller than kSZ for any galaxy mass.

Both the kSZ and the patchy screening effects are prominent during the epoch of reionization, when the cosmic ionization fraction experiences large spatial fluctuations, and in the late Universe, when structure collapses under the influences of gravity and leads to large peaks in the electron density [e.g. 506, 225, 435, 507]. The benefit of both probes is that they are directly proportional to the electron (gas) density along the line-of-sight and unlike X-rays and the tSZ have no dependence on other thermodynamic quantities, yielding a direct measure of the projected baryon profiles of galaxy groups and clusters. Because their weightings are different: the kSZ effect depends on the line-of-sight velocity, whereas patchy screening depends on the size of the primary CMB anisotropy, these probes are highly complementary and give us a way of validating conclusions about the gas distribution, calibrating the effect of systematics such as reconstruction noise and CIB/tSZ contamination, and estimating the characteristic line-of-sight velocities of galaxy groups and clusters. Deriving the velocity field is instrumental for many cosmological applications, as it can allow us to stress test modified gravity models and differentiate between different dark energy scenarios. Additionally, on very large scales, $k \lesssim 0.01$, primordial non-Gaussianity models of the local-type predict a scale-dependent signal, which can be measured with much higher accuracy if one has access to the joint tracer density and velocity fields [425].

Measuring the late-time kSZ and the patchy screening signals requires access to a high-resolution, large-area CMB experiment and a large-scale structure (LSS) survey. Ground based CMB experiments, such as the Atacama Cosmology Telescope (ACT), South Pole Telescope, the Simons Observatory, CMB-HD, and CMB-S4 [80, 283, 22, 582, 1], are thus suitable on small scales and can be complemented by low-resolution, large-area CMB surveys such as *Planck* and *WMAP* [471, 79]. To bring these two effects into light, we typically cross-correlate the CMB temperature map with a galaxy survey, such as the Baryon Oscillation Spectroscopic Survey (BOSS) [170, 29], the Dark Energy Survey (DES) [6], and the Dark Energy Spectroscopic Instrument [186].

Over the years, a number of different measurements of the kSZ effect have been performed using a variety of techniques such as the pairwise method [277, 172, 562, 570, 116, 348],

projected fields [287, 229, 332, 102], matched filter optimization [347], and the velocity reconstruction stacking method [523, 525, 578, 380, 91, 265]. Recently, Refs. [154, 536] performed a stacked measurement of the patchy screening effect using ACT and unWISE data [526, 329, 330] and put an upper bound on the optical depth.

In this chapter, we make the first self-consistent measurement of the kSZ and patchy screening effects on the same galaxy sample, provided by the DESI imaging survey, in an effort to test and validate both methods independently as well as place constraints on the velocity field along the line-of-sight to which the kSZ effect is sensitive. This chapter is organized as follows. In Section 8.3, we describe the CMB experiment ACT and the LSS survey DESI before proceeding to discuss the techniques used for measuring the kSZ and patchy screening effects on ACT and DESI data. In Section 8.4, we summarize our main results of comparing the baryon profiles derived using both probes and comment on their consistency and the benefits of having access to the velocity field. Finally, in Section 8.5 we present some concluding remarks and discuss the future of joint CMB and LSS analysis.

8.2 Data

In this section, we summarize the observational and simulation data sets used in this study.

Dark Energy Spectroscopic Instrument

The Dark Energy Spectroscopic Instrument is a robotic, fiber-fed, highly multiplexed spectroscopic telescope that operates on the Mayall 4-meter telescope at Kitt Peak National Observatory [185]. DESI can obtain simultaneous spectra of almost 5000 objects over a $\sim 3^\circ$ field [187, 552, 409] and is currently conducting a five-year dark energy survey of about a third of the sky [341]. The end product will comprise spectra for approximately 40 million galaxies and quasars [186].

Here, we use the Extended photometric sample of Luminous Red Galaxies (LRGs) [640, 643] from the DESI Legacy Imaging Survey, which was used to select DESI targets from three telescopes: Blanco for Dark Energy Camera Legacy Survey (DECaLS), Mayall for the Mayall z-band Legacy Survey (MzLS), and Bok for the Beijing–Arizona Sky Survey (BASS). We also employ the photometric redshifts presented in [643] for Data Release 9 (DR9).

Atacama Cosmology Telescope

This chapter employs the f090 and f150 single-drequency maps [155, 431] from Data Release 5 (DR5) of the Atacama Cosmology Telescope (ACT), a 6m telescope that was located in the Atacama Desert of Chile and measured the CMB from 2007 to 2022. The DR5 data comprise multifrequency observations from 2017 to 2022 covering roughly a third of the sky at three frequency bands: f090 (77–112 GHz), f150 (124–172 GHz), and f220 (182–277 GHz), and uses *Planck* data on large scales [585, 377]. The observational program of DR5 targeted

the 'wide' field. For this work, we use only the night-time portion of the data taken in the first five observing seasons (2017-2021). The ACT maps are produced in the plate-carrée (CAR) projection scheme. This analysis uses the first version of the ACT DR5 maps, dr5.01.

We apply a mask to the ACT maps that effectively discards all galaxies within 10 arcmin of a point source or a cluster [493]. In addition, we remove galaxies around which the measured filtered temperature decrements are more than a 5σ outlier in any of the bins. That way, we minimize the impact of especially bright objects and avoid introducing bias to the stacked profiles.

8.3 Methods

In this section, we introduce two methods for measuring the baryon profiles around DESI galaxies using ACT data: stacking of the kSZ effect and stacking of the patchy screening effect. We make use of the publicly available pipeline `ThumbStack` [1].

Filtering of the CMB map

In this chapter, we make use of a high-pass filtered map for the kSZ stacking measurement and both a high-pass as well as a low-pass filtered map for the patchy screening measurement, following v1 of Ref. [154].

To obtain the small-scale temperature map, we apply the following filter in harmonic space, $f_\ell^{\rm hi}$,

$$f_\ell^{\rm hi} = \begin{cases} 0 & \text{if } \ell < 2350 \\ \sin\left(\frac{(\ell-2350)\pi}{300}\right) & \text{otherwise} \\ 1 & \text{if } \ell > 2500. \end{cases} \tag{8.5}$$

For the large-scale temperature map used in the patchy screening measurement, we adopt the filter $f_\ell^{\rm lo}$,

$$f_\ell^{\rm lo} = \begin{cases} 1, & \text{if } \ell < 2000 \\ \cos\left(\frac{(\ell-2000)\pi}{300}\right) & \text{otherwise} \\ 0 & \text{if } \ell > 2150. \end{cases} \tag{8.6}$$

The cutoff scales are chosen to be non-overlapping so that the two maps are not strongly correlated with each other, and the smooth roll-off is applied so as to reduce the ringing from the harmonic space transforms. The location of the cut, $\ell \sim 2000$, ensures that the large-scale temperature sign is preserved and simplifies the form of the estimator by eliminating the need to include an inverse-variance filter.

[1] https://github.com/EmmanuelSchaan/ThumbStack

Stacking of the kSZ effect

To obtain the stacked profiles of LRGs through the kSZ effect, we need two ingredients: an estimate of the line-of-sight velocity, and a measurement of the CMB temperature decrement at the location of each galaxy.

We measure the CMB temperature decrements, $\mathcal{T}_i(\theta_d)$, around each galaxy i using a 'meanring' aperture photometry filter defined as:

$$\mathcal{T}(\theta_d) = \mathcal{N}^{-1}(\theta_d) \int_{\theta_j}^{\theta_{j+1}} \mathrm{d}^2\theta\, T^{\rm hi}(\theta), \tag{8.7}$$

where $T^{\rm hi}(\theta)$ is the high-pass filtered CMB temperature map, and the normalization is chosen as:

$$\mathcal{N}(\theta_d) = \int_{\theta_j}^{\theta_{j+1}} \mathrm{d}^2\theta. \tag{8.8}$$

In other words, we measure the mean temperature fluctuation in a ring with inner radius θ_j and an outer radius θ_{j+1} at each radial bin, d. Similarly to Ref. [154], we choose 9 radial bins (i.e., $j = 0 \cdots 9$), spanning between 0 and 10 arcmin. For the LRGs, 1 arcmin corresponds roughly to 0.5 Mpc/h ($\bar{z} \approx 0.7$). Since we obtain the measurement in several radial bins, we have sufficient information to reconstruct the projected gas profiles unlike the case of using a single bin for the aperture photometry measurement or a matched filter.

The other commonly used filter, which has been adopted for several kSZ analyses (see e.g., Refs. [525, 265]), is the compensated aperture photometry (CAP) filter. Similarly to the 'meanring' filter, we vary the aperture radius θ_d, measuring the temperature fluctuation $\mathcal{T}(\theta_d)$ as:

$$\mathcal{T}(\theta_d) = \int_0^{\theta_d} d^2\theta\, \delta T(\theta) - \int_{\theta_d}^{\sqrt{2}\theta_d} d^2\theta\, \delta T(\theta), \tag{8.9}$$

where δT are the temperature decrements measured around each galaxy on the unfiltered CMB temperature map.

The second necessary ingredient is an estimate (reconstruction) of the velocity field, which we infer by solving the linearized continuity equation in redshift space:

$$\nabla \cdot \boldsymbol{v} + \frac{f}{b} \nabla \cdot [(\boldsymbol{v} \cdot \hat{\boldsymbol{n}})\hat{\boldsymbol{n}}] = -aHf\frac{\delta_g}{b} \tag{8.10}$$

where δ_g is the observed galaxy number overdensity, $H(z)$ is the redshift-dependent Hubble parameter, f is the logarithmic growth rate, defined as $f \equiv d\ln(D)/d\ln(a)$ with $D(a)$ the growth factor and a the scale factor. Here, we assume that the galaxy overdensity δ_g is related to the matter overdensity in redshift space, δ, by a linear bias factor, b, such that $\delta_g = b\delta$.

To estimate the individual galaxy velocities, we adopt the standard reconstruction method adopted by many BAO analyses, which yields the first-order galaxy displacement

field and can be converted into an estimate for the velocity via:

$$\boldsymbol{v}^{\rm rec}(\boldsymbol{x}) = f(z)a(z)H(z)\psi(\boldsymbol{x}). \tag{8.11}$$

In principle, this gives us access to the 3D reconstructed velocity field, but the component that matters for kSZ analysis is only the line-of-sight one, which we denote by $v_\parallel^{\rm rec} \equiv v^{\rm rec}$. In particular, we apply the `MultiGrid` implementation to solve the linearized system of partial differential equations via the package 'pyrecon'[2] [622].

Once we have attained the reconstructed line-of-sight velocity and temperature decrements for each galaxy, we apply the velocity-weighted, uniform-mean estimator from [525]:

$$\hat{\tau}_{\rm kSZ}(\theta_d) = -\frac{1}{r}\frac{v_{\rm rms}^{\rm rec}}{c}\frac{\sum_i T_i^{\rm hi}(\theta_d)(v_{{\rm rec},i}/c)}{\sum_i (v_{{\rm rec},i}/c)^2} T_{\rm CMB}, \tag{8.12}$$

where the sum is over all galaxies, $T_{\rm CMB} = 2.7255$ is the mean CMB temperature, $v_{\rm rms}^{\rm rec}$ is the rms of the radial component of the reconstructed velocities, $v_{{\rm rec},i}$, c is the speed of light, and $r \equiv \langle v^{\rm rec} v^{\rm true}\rangle / (v_{\rm rms}^{\rm rec} v_{\rm rms}^{\rm true})$ is the cross-correlation coefficient between the reconstructed and true velocity, with the averaging being done over all galaxies. Typically [525, 265], this coefficient is evaluated from mock simulations, and the value Ref. [266] find for a photometric survey with DESI-like specifications and some cleaning of the redshift outliers is $r = 0.3 \pm 0.03$. If this normalization is not applied, the measured profiles and their covariance will be underestimated by $1/r$ and $1/r^2$, respectively, relative to the true optical depth.

Similarly to Ref. [265], we ensure that the number of galaxies in each velocity bin is symmetric around the mean by random downsampling, which avoids unwanted biases from massive clusters and guarantees that in the absence of a kSZ signal, our estimator yields zero mean signal. Since we sum over an equal number of galaxies with a positive and a negative sign per absolute value bin in the velocity, additive contaminants such as the cosmic infrared background (CIB) and the tSZ effect cancel, as has been demonstrated in Refs. [523, 525, 265].

Stacking of the patchy screening effect

In analogy with kSZ estimators, small-scale, real-space estimators have been proposed, with the potential of allowing patchy screening analyses to be performed with the same pipelines as kSZ. Here we use a similar estimator to v1 of [154, 536]:

$$\hat{\tau}_{\rm patchy}(\theta_d) = \frac{-\sum_i {\rm Sign}[T_i^{\rm lo}] T_i^{\rm hi}(\theta_d)}{N_g \langle |T^{\rm lo}| \rangle}, \tag{8.13}$$

where the sum and averaging is done over all N_g galaxies in the sample. We note that unlike Ref. [154] where $T_i^{\rm lo}$ is allowed to vary as a function of radius for each object, we take the sign of the low-pass filtered map at the pixel corresponding to the galaxy right ascension and

[2] https://github.com/cosmodesi/pyrecon

declination. We find that this choice makes a negligible difference and prevents accruing a bias in some fringe cases.

Similarly to the kSZ effect, this form of the estimator (that depends on the large-scale sign) is in principle immune to foreground emission such as tSZ and CIB. To further suppress the amount of residual foreground contamination to the final profiles, we randomly downsample the galaxy samples such that the number of positively and negatively weighted regions is equal. We also apply a cut on the amplitude of the low-pass-filtered temperature of $|T^{\rm lo}| > 40\,\mu$K. This is done since when that amplitude is small, the value of $T^{\rm lo}$ might be dominated by foregrounds (that can change the sign of $T^{\rm lo}$) rather than the primary CMB, which is what is needed for this measurement.

Covariance matrix

We estimate the covariance matrix for the stacked profiles by calculating the covariance between the radial bins per galaxy of $\tau_{\rm kSZ}$ and $\tau_{\rm patchy}$ and then simply dividing by $\sqrt{N_g}$. We note that because this calculation assumes that the galaxy profiles are independent of each other (and we know that many of the cutouts are overlapping), this covariance ends up underestimating the true covariance by around 10% as demonstrated in Ref. [154] when comparing with simulations. Since the main goal of this study is to compare the two measurements (kSZ and patchy screening) for the same sample of objects, we are less preoccupied with this effect as long as it affects both measurements equally.

Simulations

To put our findings into context, we compare the profiles we measure via ACT and DESI with the profiles measured in the hydrodynamical simulation IllustrisTNG. To this end, we construct 2D maps from the simulations of the kSZ or the optical depth, then convolve them with a 1.4 arcmin beam, and finally apply filtering as described in Section 8.3.

The IllustrisTNG galaxy formation model features primordial and metal line cooling, a subgrid model for star formation and the interstellar medium, mass return from stars via active galactic nuclei (AGN) and supernovae, an effective model for galactic winds, as well as a model for the formation, growth, and feedback from supermassive black holes [619, 470]. At a baryonic mass resolution of $1.1 \times 10^7 {\rm M}_\odot$, the highest-resolution IllustrisTNG simulation TNG300-1 reproduces well many properties of observed galaxies and galaxy clusters.

IllustrisTNG outputs a number of useful quantities: gas cell mass (m), electron abundance (x), and internal energy (ϵ). Assuming a primordial hydrogen mass fraction of $X_H = 0.76$, we compute the volume-weighted electron number density, $n_{\rm e}$, for each gas particle i as

$$V_i n_{{\rm e},i} = x_i m_i \frac{X_H}{m_p} \tag{8.14}$$

where m_p the proton mass and k is the Boltzmann constant. We then compute the 2D maps of the kSZ and the optical depth by binning the gas particles into a (10000, 10000) grid, so

that the optical depth in cell j is given by:

$$\tau_j = \sigma_T A_j^{-1} \sum_{i \in A_j} V_i n_{\mathrm{e},i}, \tag{8.15}$$

and the momentum of the electron density is:

$$b_j = \sigma_T A_j^{-1} \sum_{i \in A_j} V_i n_{\mathrm{e},i} v_i / c, \tag{8.16}$$

where A_j is the area of each grid cell (of size 0.03 Mpc), σ_T is the Thomson cross section, and c is the speed of light.

8.4 Results

In this section, we present our main findings of comparing the kSZ and patchy screening signals measured on both data and simulations.

Patchy screening and kSZ comparison

In Fig. 8.1, we compare the signal measured from both effects, 'patchy screening' and kSZ, on DESI LRG groups by applying the two estimators defined in Eq. 8.12 and Eq. 8.13 to the ACT DR5 combined f090 and f150 maps. We do so by filtering the f090 single-frequency map by the ratio of the two beams such that it becomes effectively rebeamed to the same beam size as the f150 map. The significance of detection using the two estimators is 7.2σ and 4.1σ, respectively, and the measured values of χ^2 are 52.5 and 29.9, with 9 degrees-of-freedom. We obtain the signal-to-noise ratio as follows:

$$\mathrm{SNR} = \sqrt{\chi^2_{\mathrm{null}} - \chi^2_{\text{best-fit}}}, \tag{8.17}$$

where the χ^2 statistic is defined as

$$\chi^2_{\mathrm{model}} = (d-m)^\top \mathrm{C}^{-1} (d-m). \tag{8.18}$$

In the case of χ^2_{null}, $m = 0$, whereas for the best-fit model, we adopt a rescaling of the simulation curves detailed in App. F.1 with a single free parameter for the amplitude, i.e. $m = A\ s$, where s is the simulation curve of galaxies with a mean mass of $10^{13.3}\ M_\odot/h$ [635]. We note that because of the high-pass filtering, most of the information regarding the amount of feedback, halo mass, satellite fraction, and beam is contained in the amplitude rather than the shape of the curve. This is different from the CAP filter, for which the shape and amplitude convey independent information, as will be discussed in the next section.

To convert the kSZ signal measured on the temperature maps into optical depth, we do as follows:

$$\tau^{\mathrm{hi}}_{\mathrm{kSZ}} = T^{\mathrm{hi}}_{\mathrm{kSZ}} / T_{\mathrm{CMB}}\ c/v_{\mathrm{rms}} \frac{1}{r}, \tag{8.19}$$

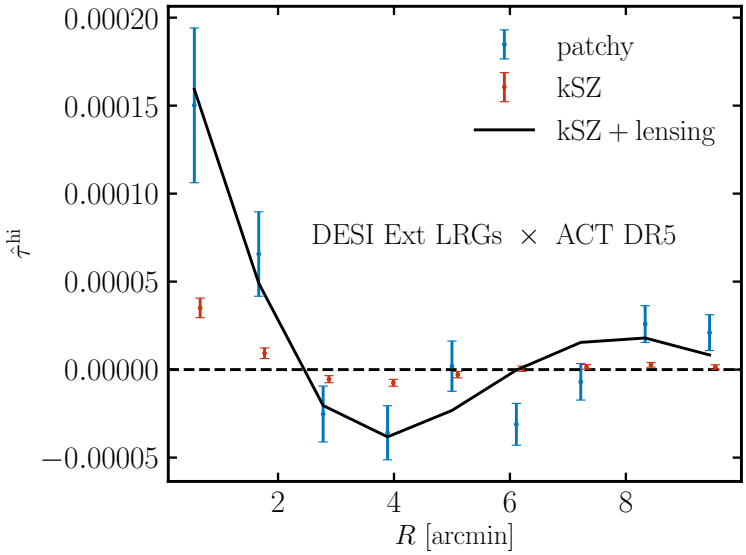


Figure 8.1: Estimates of the optical depth of DESI LRG groups using the two estimators defined in Eq. 8.12 and Eq. 8.13 and applying them to the ACT DR5 combined f090 and f150 maps. We show that the two estimates of $\tau_{\rm hi}$ are in very good agreement if we take into consideration the contribution from CMB lensing to the 'patchy screening' estimator (i.e., Eq. 8.13), which makes up about two thirds of the total signal. The lensing contribution is assessed using the N-body simulation, AbacusSummit (see Section 8.4, for more details). A noteworthy feature is that the high-pass filtering of the CMB map makes the profiles with high, low and no baryonic feedback look very similar to each other (see also Ref. [154]), rendering the comparison of profile shapes at high significance somewhat challenging given the signal-to-noise of the data.

where the rms of the sample is assumed to be 300 km/s, and the cross-correlation coefficient between the true and reconstructed velocity is $r = 0.25$, which we verify in simulations for a photometric LRG-like sample in DR9 with no photo-z cleaning performed on it as done in Ref. [265].

We see that the amplitude of the kSZ signal appears to make up one third of what the estimator in Eq. 8.13 yields, and we attribute the difference between the two to weak lensing of the CMB giving a large contribution to the patchy screening estimator as we discuss next. This significantly lowers the detection significance of 4.1σ and therefore, we quote an upper bound instead.

To put an upper bound on the optical depth detected through the 'patchy screening' effect, we take the best-fit value of the amplitude parameter A (see Eq. 8.18), which minimizes the difference between data and model. We note that as we are interested in the pure τ signal, here we work with the lensing-free curve, which can either be obtained by subtracting the lensing contribution from the 'patchy screening' curve or by directly using the kSZ profile.

The final result is virtually unchanged. The upper limit is then given by

$$A + N\sigma_A = A + \frac{N}{\sqrt{d^\top \mathrm{C}^{-1} d}}, \tag{8.20}$$

where σ_A is the variance of A and $N = 1.64$ for a 95% confidence level. We find that $A = 0.54$ and $\sigma_A = 0.61$. To convert this into a physically meaningful quantity, we note that the mean optical depth for the sample used in the theory prediction (mean halo mass of $10^{13.3}\ M_\odot/h$) is $\approx 1.6\ 10^{-4}$, including the one- and two-halo terms [272]. Thus, our measurement places a 95% upper bound on the optical depth of the Extended DESI LRG sample of $\tau < 2.5\ 10^{-4}$.

CMB lensing contribution to the stacked patchy screening estimator

CMB lensing couples large and small scales, generally biasing patchy screening estimators. This was first pointed out by [568], who noted that the lensing bias is generally larger than the signal itself, and proposed "lensing hardening" as a potential solution, which is noted to lead to a non-negligible noise cost. The lensing contribution to real-space estimators has not yet been quantified (to our knowledge), and as discussed below, it is expected to be very significant. We discussed the origin of the lensing bias and several mitigation strategies in the previous chapter [515].

The lensing contribution is estimated by adopting the publicly available convergence maps of AbacusSummit (namely, the all-sky 'huge' simulation `AbacusSummit_huge_c000_ph201`) and lensing a random Gaussian realization of the angular power spectrum, C_ℓ, from `CAMB` using `pixell` [273]. We then stack at the location of the AbacusSummit halos on the light cone [269] in the redshift range corresponding to the DESI LRGs: $z = 0.45$ to 1. Note that the generated lensing maps do not contain a τ signal and can thus be independently added to a pure τ curve, as done in Fig. 8.1. Finally, we select a low-mass threshold, so that the mean halo mass roughly matches that of the LRG host halos, $\log M \approx 13.3$ in units of $M_\odot/h$. We perform an additional test to check that we recover the correct relative amplitude between the lensing and the τ signals and reaffirm our finding (App. F.1).

Naively, one might expect that the matter profiles measured by lensing would be different from the gas profiles that make up a fraction of the total 'patchy screening' signal, especially on the outskirts of groups and clusters, where we expect the gas to be pushed out compared with the dark matter. However, we note that because of the high-pass filtering of the CMB map, information on these larger scales is largely erased, making the shapes of the profiles with high, low and no baryonic feedback look virtually indistinguishable from each other (see also Ref. [154]) given the signal-to-noise of the data.

Instead, the amplitude now contains information about the gas distribution (provided lensing is accounted for in the 'patchy screening' case) and thus feedback, while also being sensitive to the satellite fraction and the mean halo mass of the sample. This sensitivity comes from the fact that stacking at the location of satellite galaxies leads to more extended

profiles compared with the case of stacking at the center of the halos (see Fig. 8.3), which then curbs the power at high-ℓ. We note that the beam of a CMB experiment has the same effect, as it puffs out the mass distribution.

Similarly, at fixed halo mass (and we remind the reader that both the τ and lensing signals are proportional to halo mass), galaxy groups with little feedback would have a higher amplitude than galaxy groups with a large amount of feedback. As we increase the halo mass (keeping the baryonic properties and satellite fraction fixed), the amplitude increases proportionally to the halo mass until the size becomes comparable to the size of the high-pass filter threshold applied to the CMB map, at which point some of the amplitude will be lost, as it leaks to larger scales than allowed by the filtering.

However, if one can pinpoint the halo occupation distribution of the sample, then the gas feedback can be constrained. Additionally, provided that the effect of CMB lensing is well understood, the comparison between the kSZ $\tau^{\rm hi}$ amplitude and that of the bias-hardened 'patchy screening' $\tau^{\rm hi}$ could yield a powerful probe of the typical velocities of the galaxy sample host halos and can also be used to calibrate r independently of N-body simulations [499, 266], providing a long-sought solution to the 'optical-depth degeneracy' problem of disentangling astrophysical quantities (τ) from cosmological ones ($v_{\rm rms}$) [36].

Size of the baryonic feedback

First, we inspect the optical depth signal in the 'patchy screening' measurement from DESI LRG $\times$ ACT DR5, obtained by stacking via the estimator in Eq. 8.13. This is shown in Fig. 8.2. For all other samples (i.e., the kSZ data measurement and the τ simulation curves), we add the lensing contribution at the mean halo mass of the galaxy sample, $\log M \approx 13.3$, using the AbacusSummit weak lensing maps. The two hydro simulation curves are obtained by painting gas properties onto the AbacusSummit light cone maps via the method outlined in App. F.1. We see that the τ profile measured in the data has a much lower amplitude compared with TNG300-1, implying that the true baryonic feedback is stronger than the one predicted in TNG300-1. In fact, we see that the data appears to be much more similar in shape and amplitude to the Illustris τ profile, implying that the amount of feedback Illustris predicts, which is large compared with the majority of hydro simulations, reflects more accurately the extent to which gas is expelled in the Universe through processes such as AGN feedback. Reassuringly, we find great consistency with previous kSZ measurement analyses [38, 265, 91, 393, 267], providing further evidence for strong baryonic feedback in the LRG sample.

We note that while the transfer function method (App. F.1) does an extremely good job of mimicking the gas distribution in a hydro simulation at scales larger than $R \sim 1$ Mpc/h ($k \lesssim 1\ h$/Mpc in Fourier space) [359], on smaller scales, where the cross-correlation coefficient between the dark matter density and τ becomes significantly smaller than one, the painted-on gas density will deviate from the gas density in the hydro simulation. We note that this affects the first couple of radial bins of the 'painted' Illustris curve, which end up having higher amplitude compared with the 'hydro' Illustris curve (we have verified that to be

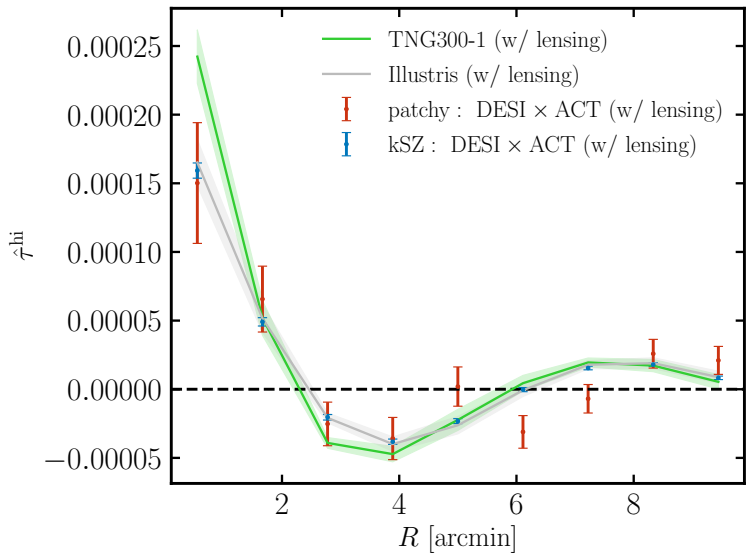


Figure 8.2: Optical depth signal in the 'patchy screening' measurement from DESI LRG × ACT DR5, obtained by stacking via the estimator in Eq. 8.13. To make the comparison with simulations and the kSZ, we add the lensing contribution to the kSZ and simulations at the mean halo mass of the galaxy sample, $\log M \approx 13.3$, using the AbacusSummit weak lensing maps. The measured τ profile has a lower amplitude compared with TNG300-1, implying that the true baryonic feedback is stronger than the one predicted in TNG300-1. The profiles are more akin to the Illustris τ profile, implying that the amount of feedback Illustris predicts, which is large compared with the majority of hydro simulations, reflects more accurately the extent to which gas is expelled in the Universe through processes such as AGN feedback.

the case at lower masses, for which we have a representative sample of halos in the hydro simulation).

Significance of the CAP filter

In the top panel of Fig. 8.3, we show the CAP profiles of high-mass galaxies (i.e., with $log M_* > 10.45$) measured in IllustrisTNG (TNG300-1) obtained by stacking on 2D maps of the optical depth and the kSZ effect at $z = 0.5$. We show the results for the default LRG sample ('All galaxies'), the centrals belonging to that sample ('Centrals only'), and the halo centers for all galaxies ('All centrals'). This quantity can be thought of as a cumulative baryon profile. We note that these maps do not have the primary or lensed CMB signal. Furthermore, on very small scales (∼1 arcmin), the transfer function method breaks down, as the cross-correlation coefficient between the dark matter field and the gas field in the Illustris simulation becomes small ($r \ll 1$).

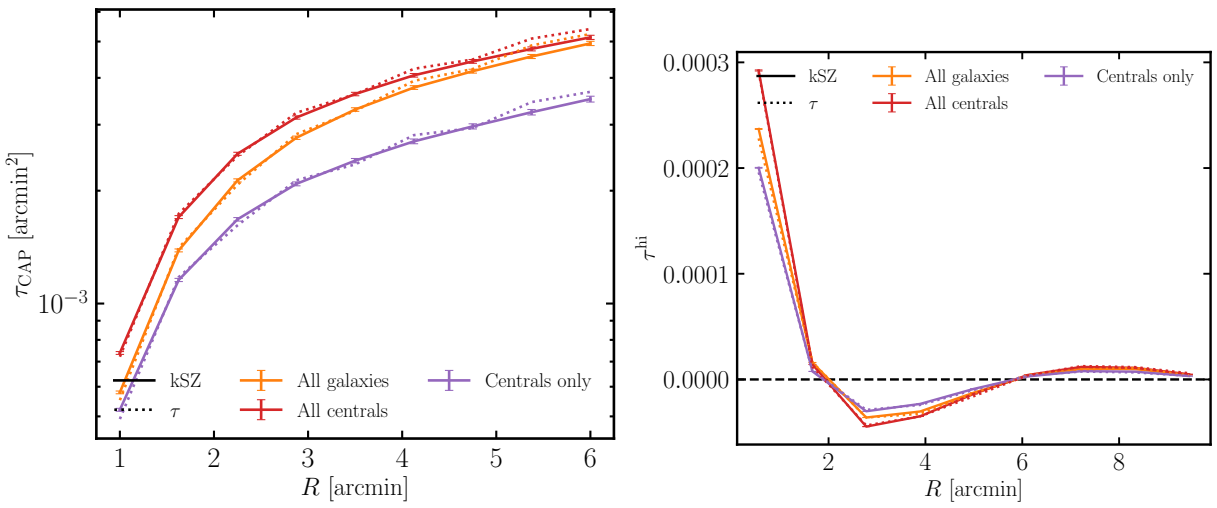


Figure 8.3: CAP (*top*) and high-pass-filtered (*bottom*) profiles measured in IllustrisTNG (TNG300-1) around high-mass galaxies from 2D maps of the optical depth, τ (dotted), and the kSZ effect (solid) at $z = 0.5$. We note that the profiles are not convolved with the ACT beam and that the maps contain no primary CMB or lensing contributions. The units of the signal in the top panel are steradians, as the CAP filter measures the difference between the integrated signal in a disk and in a ring around it (see Eq. 8.7). The quantity on the y-axis can thus be thought of as a cumulative baryon profile, whereas $\tau^{\rm hi}$ in the bottom panel is unitless, as it corresponds to the mean high-pass filtered optical depth. The three colors correspond to stacking on the default LRG sample ('All galaxies'), stacking on only the centrals belonging to that sample, or around 80% of the galaxies ('Centrals only'), and stacking on all galaxies, but at the location of their host halo centers to avoid miscentering ('All centrals'). Most noteworthy is the fact that the kSZ and τ lie on top of each other, suggesting that they measure the same underlying quantity of optical depth.

Unsurprisingly, at large aperture, the 'All centrals' and 'All galaxies' curves reach the same value of total gas mass. In further comparing the 'All galaxies' sample with the 'All centrals,' we see that the main difference is that 'All centrals' appears slightly flatter. This is due to the fact that the non-zero satellite fraction of the 'All galaxies' sample leads to miscentering, which artificially makes the stacked profiles more extended. The amplitude, as expected, is the same as the mean halo mass remains unchanged.

On the other hand, the 'Centrals only' sample has slightly lower amplitude as when stacking on centrals only, we have discarded satellite galaxies, which are more likely to reside within massive groups and clusters, and have thus lowered the mean halo mass of the sample. That is also the main reason for the difference in shape: namely, the mean virial radius of that sample is smaller, making the profiles flatter at fixed apparent size on the sky. The error bars here are calculated as the standard deviation of the measured curves from the

sample divided by the square root of objects in the sample. The CAP filter cannot be applied to the 'patchy screening' estimator due to leakage of the large-scale CMB fluctuations (see also App. F.2), so when studying that effect, we need to adopt high-pass filtered profiles.

In the bottom panel, analogously to the CAP case, we study the high-pass filtered profiles. Since in this case, we are taking the average of the signal in rings of increasing distance, the highest point of the profile sits in the center of the halo, as expected, where most of the gas resides. The dip below zero and then rise of the signal at large radii is due to the ringing effect at imposing the cutoff on the 2D map. As discussed previously, the amplitude of the high-pass filtered signal is lowered in the presence of satellite galaxies, as the density profile becomes more puffed out and the high-pass filter takes away some of that power on large scales.

Reassuringly, for both filters (CAP and high-pass) the kSZ and τ lie on top of each other, indicating that they measure the same underlying quantity of optical depth. The role of the CAP filter is important, as it removes the uncorrelated part of the signal, which dominates the optical depth map, leaving the contribution of the targeted galaxies only. Similarly, the high-pass filter gets rid of the largest-scale contributions both from the primary CMB and also from random structure along the line-of-sight. In the case of the kSZ, that is less of an issue as due to the velocity decorrelation along the line-of-sight, the signal picks up predominantly the contribution from the targeted object [272].

8.5 Discussion and conclusions

In this chapter, we pursue a comparison between the stacked measurements of the gas density using the kinematic Sunyaev-Zel'dovich effect and the 'patchy screening' effect. Both of these probes are sensitive to the gas distribution and can thus help us inventory the missing baryons in the Universe. In principle, the patchy screening effect is about 25 times smaller than the kSZ effect, but it can be applied to 2D surveys without any knowledge of the line-of-sight velocity field, which is necessary for the kSZ effect. As such, the combination of these two effects provides a powerful way of disentangling the astrophysical quantity, optical depth, from the cosmological one, peculiar velocity, breaking the so-called 'optical depth degeneracy' and yielding an independent measurement of the line-of-sight velocity.

Here, we work with the Extended LRG sample of the DESI imaging survey DR9, which has previously been used for measuring the gas density profiles using the kSZ effect. We then high-pass filter the CMB temperature map and apply the estimator proposed in v1 of Ref. [154] (Eq. 8.13) and the kSZ estimator (Eq. 8.12) to consistently compare the two. Our main findings are summarized as follows:

- In Fig. 8.1, we present a high-significance detection of both effects: namely, we measure the kSZ signal at 7.2σ and the 'patchy screening' signal at 4.1σ, the latter of which turns out to be contaminated by CMB lensing.

- We show that the 'patchy screening' estimator (Eq. 8.13) receives a large contribution from CMB lensing, which makes up about two thirds of the measured signal for the chosen sample. While this chapter uses simulations to estimate the effect of lensing, in the preceding chapter [515], the focus is on deriving an analytical expression to account for this contamination.

- As explored in App. F.2, after removing the lensing portion, the signal-to-noise for τ coming from 'patchy screening' of 4.1σ is lowered significantly, so we can quote an upper bound of $\tau < 2.5\ 10^{-4}$ for our sample, which we assume to have a mean optical depth of $\tau \approx 1.6\ 10^{-4}$. We note that the effect is still measurable with future surveys such as the Vera Rubin Observatory (LSST) [310].

- By performing stacked measurements on hydro simulations of pure kSZ and pure τ maps (i.e., without the inclusion of other CMB primary and secondary effects), we are reassured that the kSZ estimator (Eq. 8.12) is indeed measuring the correct underlying quantity of optical depth, τ.

In the future, we aim to explore ways of combining these two potent measurements to acquire a detailed picture of the baryon distribution in the Universe. The ultimate constraints on the thermodynamics of galaxy groups and clusters and by extension the effects on the small-scale matter power spectrum will come from obtaining a multi-scale multi-wavelength view of the Universe through additional powerful probes such as tSZ, X-rays and Fast Radio Bursts (FRBs).

Acknowledgments

We thank Matt Johnson for a useful discussion during the preparation of this manuscript. N.S. and S.F. are supported by Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Chapter 9

Dispuτable: the high cost of a low optical depth

This chapter was originally published as

> Noah Sailer, Gerrit S. Farren, Simone Ferraro, and Martin White. "Dispuτable: the high cost of a low optical depth". In: *arXiv e-prints* (Apr. 2025). arXiv: 2504.16932 [astro-ph.CO]

Recent Baryonic Acoustic Oscillation (BAO) measurements from the Dark Energy Spectroscopic Instrument (DESI) are mildly discrepant (2.2σ) with the Cosmic Microwave Background (CMB) when interpreted within ΛCDM. When analyzing these data with extended cosmologies this inconsistency manifests as a $\simeq 3\sigma$ preference for sub-minimal neutrino mass or evolving dark energy. It is known that the preference for sub-minimal neutrino mass from the suppression of structure growth could be alleviated by increasing the optical depth to reionization τ. We show that, because the CMB-inferred τ is negatively correlated with the matter fraction, a larger optical depth resolves a similar preference from geometric constraints. Optical depths large enough to resolve the neutrino mass tension ($\tau \sim 0.09$) reduce the preference for evolving dark energy from $\simeq 3\sigma$ to $\simeq 1.5\sigma$ and increase the CMB-inferred values of n_s and H_0 to 0.968 ± 0.004 and 67.94 ± 0.44 km/s/Mpc respectively. Conversely, within ΛCDM the combination of DESI BAO, high-ℓ CMB and CMB lensing yields $\tau = 0.090 \pm 0.012$. The required increase in τ is in $\simeq 3-5\sigma$ tension with *Planck* low-ℓ polarization data when taken at face value. While there is no evidence for systematics in the large-scale *Planck* data, τ remains the least well-constrained ΛCDM parameter and is far from its cosmic variance limit. The importance of τ for several cosmological measurements strengthens the case for future large-scale CMB experiments as well as direct probes of the epoch of reionization.

9.1 Introduction

Recent measurements of the Cosmic Microwave Background (CMB) and Baryon Acoustic Oscillations (BAO), have highlighted a number of discrepancies within the standard flat ΛCDM cosmological model. For example, DESI DR2 BAO, together with CMB measurements from *Planck* show a $\sim 3\sigma$ preference for a time-varying (dynamical) dark energy model [184]. These stem in part from parameter inconsistencies (such as in the matter density $\Omega_{\rm m}$ and the Hubble parameter H_0) in the comparison between the values inferred at early times by the CMB and at later times by BAO within ΛCDM.

Moreover, attempts at measuring the mass of neutrinos through their effect on cosmological observables [338, 367] have revealed surprising features: when imposing a positivity prior, the constraints are tighter than a Fisher information analysis would suggest. If a positivity prior is not imposed, and the "effective" neutrino mass is allowed to be negative, the minimum mass implied by oscillation experiments ($\sum m_\nu \geq 0.06$ eV [218]) is excluded at $\sim 3\sigma$ [211, 214, 255, 156, 370]. These discrepancies can be partially alleviated by e.g. dynamical dark energy [211] or new physics in the neutrino sector [156]. However, given the paradigm shift this would imply, it is important to explore other possibilities.

It has long been known that most measurements of neutrino mass from the suppression of structure growth are sensitive to the assumed value of the optical depth to reionization τ [32, 634, 248, 367, 156, 31]. Moreover, constraints on τ are also correlated with other cosmological parameters (e.g. $\Omega_{\rm m}$ and H_0) and therefore τ influences constraints derived from the expansion history, including both neutrino mass and the time-evolution of dark energy [248, 31].

Typically measured through large-scale polarization of the CMB at $\ell \lesssim 20$, the last "official" *Planck* analysis gives $\tau = 0.0544 \pm 0.0073$ [475] based on 2018 PR3 data. Subsequent analyses of the same data find a slightly higher value $\tau \approx 0.06$ [78, 452, 41]. Independent analyses of *Planck* PR4 data yield $\tau = 0.0533 \pm 0.0074$ [503] and $\tau = 0.058 \pm 0.006$ [592]. While great progress in the understanding and control of systematics has been made in the past few years, this remains an extremely challenging measurement. As noted in the *Planck* legacy paper [473], the polarization on large scales is more than two orders of magnitude smaller than the corresponding temperature anisotropies, making it difficult to control instrumental systematic effects and foregrounds while retaining enough sensitivity to the signal of interest (see Section 6.6 of [473] or further discussions in refs [452, 171, 478, 593, 503, 592]).

Moreover, the interpretation of the polarization signal depends to some extent on the reionization history: τ is almost independent of the detailed history for "fast" models of reionization [14], but a non-trivial fractional ionization at high redshift could potentially affect the CMB measurement of τ [282]. We don't attempt to study shifts in τ given different reionization histories, but we note that no preference for a high redshift ionized component is apparent in the latest *Planck* data, and that an ionized component at $z > 15$ appears to be disfavored by the large-scale CMB polarization [407, 475], although this statement is somewhat model dependent [306].

Direct measurements of the Lyman-α forest and other observables tightly constrain the low-z half of reionization [396, 259, 258, 561, 386, 626, 613] (and are consistent with a "late and fast" reionization); however, the higher-z tail has very few direct constraints. Observations of high redshift galaxies from the James Webb Space Telescope (JWST) are consistent with a non-zero ionization fraction at $z \gtrsim 10$ (but with significant uncertainty), potentially leading to a larger optical depth [427, 433, 577]. New physics, such as dark matter annihilation or decay [357], or accretion onto primordial black holes [77, 498], could also increase the ionization fraction at even higher redshift, potentially affecting CMB measurements of τ.

Other techniques, such as using the redshifted 21cm line from neutral hydrogen [355, 513] or the kinematic Sunyaev-Zel'dovich (kSZ) effect [65, 456, 312, 441, 559, 230, 36], might shed light on the matter in the future. The absence of a detection of a large kSZ power spectrum or trispectrum appears to favor a relatively "short" reionization, though this is model dependent and currently relatively weak [496, 494, 371].

To explore the sensitivity of the tensions discussed here to τ, we exclude low-ℓ polarization [474, 176, 592, 78] from our fits, and replace it with a Gaussian prior on τ. We study two scenarios: one with prior $\tau = 0.06 \pm 0.006$, which reproduces the *Planck* constraints very closely, as well as a hypothetical higher value $\tau = 0.09 \pm 0.006$. This latter prior is, at face value, inconsistent with the *Planck* low-ℓ measurement at 3.8σ significance.[1] Such a large shift is likely not reconciled by changes in the reionization history alone; however, given the considerations outlined above, we think that it is interesting to explore the consequences of a hypothetical higher value of τ.

Following previous work [475, 634, 248, 367, 156], we also consider the inverse problem: assuming a flat ΛCDM model and a physical prior on neutrino masses from oscillation experiments, we solve for τ (without the inclusion of the low-ℓ polarization), and show that the inferred τ is in tension with the *Planck* measurement, thus suggesting that within flat ΛCDM, some of the internal inconsistencies can be reinterpreted as a "τ tension."

Given that modeling errors could potentially be at play we choose to only include observables that can be modeled using linear theory (primary CMB, BAO) or with extremely mild non-linearities and no additional free parameters (CMB lensing). The inclusion of other observables such as galaxy clustering, galaxy weak lensing, cross-correlations, or supernovae is worthwhile but requires care and will be the subject of future work.

9.2 Data

We adopt the shorthand **"DESI"** for the DR2 BAO data [184] and **"CMB"** for the combination of *Planck* PR4 `CamSpec` high-ℓ [208, 503] primary temperature and polarization measurements with PR4 [119] and ACT DR6 [493, 372, 377] lensing data. We substitute

[1] This takes the uncertainty of the prior into account. Determination of τ from other experiments such as the WMAP 9 year measurement $\tau = 0.089 \pm 0.014$ [79] or CLASS $\tau = 0.053^{+0.018}_{-0.019}$ [350] (in cross-correlation with *Planck*) have considerably larger uncertainty and are consistent with both scenarios.

low-ℓ CMB data[2] with Gaussian priors on the optical depth τ with mean 0.06 or 0.09 and uncertainty $\sigma_\tau = 0.006$, corresponding to the approximate σ_τ from low-ℓ *Planck* data [475]. We use the shorthand **"$\boldsymbol{\tau = 0.06}$"** and **"$\boldsymbol{\tau = 0.09}$"** for these priors respectively. To directly compare with previous results in the literature we choose not to include the ACT DR6 primary CMB likelihood [362, 115, 430], while noting excellent agreement between ACT and *Planck* on the relevant parameters and the absence of a τ measurement from ACT.

9.3 The optical depth's impact on cosmological 'tensions'

When jointly fitting to CMB lensing and high-ℓ primary CMB the inferred optical depth and matter fraction are negatively correlated within ΛCDM, as illustrated in the left panel of Fig. 9.1. Thus increasing the optical depth brings the CMB-preferred Ω_m into better agreement with BAO measurements. We illustrate this in the right panel of Fig. 9.1, where we plot the DESI ΛCDM constraints (blue) alongside those derived from the CMB when additionally including a $\tau = 0.06$ (pink) or $\tau = 0.09$ (black) prior. Following section VI of ref. [184] we compute the parameter-level χ^2 using the best fit $(\Omega_\mathrm{m}, H_0 r_d)$ values from

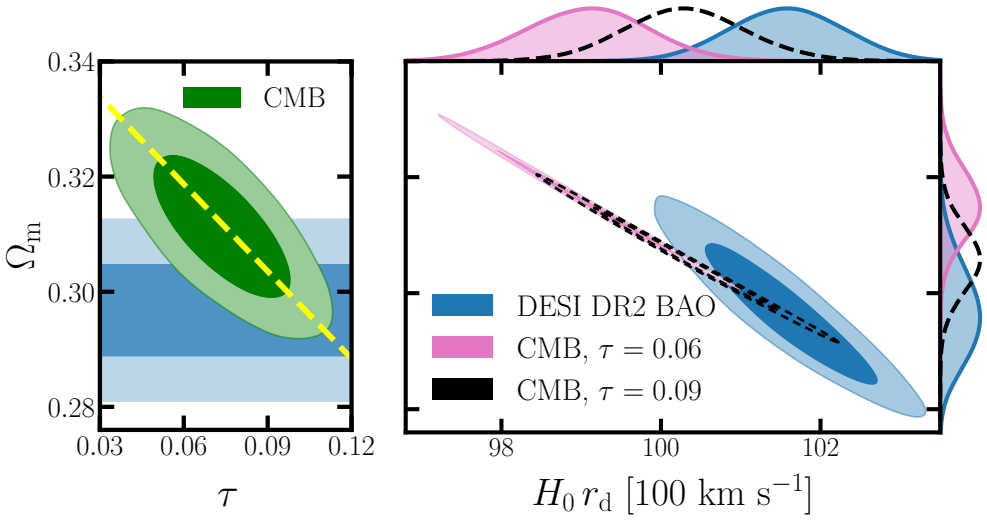


Figure 9.1: *Left:* Constraints on τ and Ω_m from the joint analysis of high-ℓ CMB and CMB lensing. The dashed yellow line $\left(\Omega_\mathrm{m}(\tau) = 0.349 - 0.505\,\tau\right)$ approximates the $\tau - \Omega_\mathrm{m}$ degeneracy direction. *Right:* ΛCDM constraints from BAO (blue) or CMB high-ℓ primary and lensing with a τ prior (pink/black).

DESI and the CMB, finding $\chi^2 = 6.55$ and 2.07 when including a $\tau = 0.06$ or 0.09 prior

[2]We additionally discard low-ℓ temperature for simplicity. We have explicitly checked that CMB-derived ΛCDM constraints shift by at most 0.2σ with the inclusion of low-ℓ temperature.

respectively. The corresponding ΛCDM "tension" between the BAO and CMB data reduces from 2.1σ to 0.9σ.

Neutrino mass

The combination of primary CMB, CMB lensing and BAO has shown a preference for neutrino masses in tension with lower bounds from oscillation experiments [214, 156, 255, 367]. To assess the consistency of the inferred neutrino mass sum with the lower bound $\sum m_\nu \geq 0.06$ eV we implement a phenomenological model in which the neutrino mass is replaced by an effective parameter $M_{\nu,\mathrm{eff}}$. While there isn't a unique way to extend the model to negative values of $M_{\nu,\mathrm{eff}}$ [214, 156, 255], here we use the following simple prescription: For any observable $X(\Sigma\, m_\nu)$, we take

$$X(M_{\nu,\mathrm{eff}}) = X(0) + \mathrm{sgn}(M_{\nu,\mathrm{eff}})\Big[X(|M_{\nu,\mathrm{eff}}|) - X(0)\Big]. \tag{9.1}$$

This implementation recovers the physical neutrino model for $M_{\nu,\mathrm{eff}} > 0$ and linearly extrapolates to negative "masses." This model was shown to agree well with a more sophisticated treatment based on replacing the neutrino energy density with an effective energy density, allowed to take on negative values, at all orders in perturbation theory [214]. We are able to reproduce the results based on the latter method in [211] with good accuracy. When including large-scale polarization data from *Planck* we find $M_{\nu,\mathrm{eff}} = -0.105^{+0.049}_{-0.062}$ (cf. $-0.101^{+0.047}_{-0.056}$ [211]).

When replacing the low-ℓ *Planck* data with Gaussian priors centered on $\tau = 0.06$ and $\tau = 0.09$ we find $M_{\nu,\mathrm{eff}} = -0.08^{+0.05}_{-0.06}$ and $M_{\nu,\mathrm{eff}} = 0.01 \pm 0.06$ respectively (see Fig. 9.2). While the former exhibits a moderate $\sim 2.4\sigma$ tension with the minimum mass allowed in the normal hierarchy, the latter is consistent with minimum mass, normal hierarchy neutrinos at $\sim 0.9\sigma$. Both values of τ provide similarly good fits to the data with a $\Delta\chi^2_{\mathrm{MAP}} = 1.8$ in favor of the smaller of the two.

Our results are consistent with refs. [156, 367], who note that for $\tau \sim 0.09$ the resulting increase in A_s shifts the neutrino mass inferred from the suppression of structure towards its minimal value. The negative correlation between τ and Ω_m (Fig. 9.1) supports this trend since the "geometric" preference for negative neutrino mass weakens as the CMB-inferred Ω_m decreases.

Note that other ways of alleviating the tension between Ω_m inferred from the CMB and BAO, like the decaying dark matter model presented in [370], do not alleviate the discrepancy in both methods of inferring the neutrino mass, since they increase the amplitude of matter density fluctuations inferred from CMB lensing and therefore potentially exacerbate any excess lensing compared to ΛCDM expectations.

Evolving dark energy

The DESI collaboration [184] finds a 3.1σ preference for the $w_0 w_a$ dynamical dark energy model [144, 353] over ΛCDM based on a joint analysis of BAO and CMB (primary and

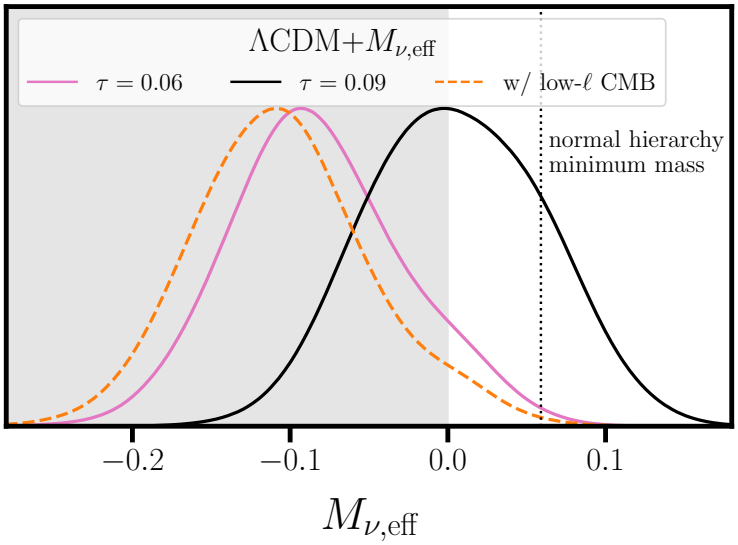


Figure 9.2: Constraints on $M_{\nu,\mathrm{eff}}$ from the joint analysis of BAO and CMB data when including a τ prior ($\tau = 0.06$ in pink and $\tau = 0.09$ in black) or low-ℓ data from *Planck* (orange).

lensing) measurements. When substituting a $\tau = 0.06$ prior in place of the *Planck* low-ℓ data we find $\Delta\chi^2_{\mathrm{MAP}} = -10.75$ between the best-fit w_0w_aCDM and ΛCDM cosmologies. This corresponds to a 2.8σ preference for the former, in reasonable agreement with the DESI DR2 results. As is shown in Fig. 9.3, the best fit (w_0, w_a) shifts from $(-0.5, -1.5)$ to $(-0.8, -0.6)$ when increasing the mean of the τ prior from 0.06 to 0.09. With $\tau = 0.09$ we find $\Delta\chi^2_{\mathrm{MAP}} = -3.82$, corresponding to a mild, 1.4σ, preference for evolving dark energy.

9.4 Inferring the optical depth without reference to large scale CMB

We have shown that several widely discussed preferences for beyond-ΛCDM physics arising from CMB and BAO data are substantially weakened when a larger optical depth is assumed. Conversely, we may ask the question what value of τ is preferred by the data when combined within a given model and whether this value is consistent with the one inferred from low-ℓ CMB polarization. For this purpose we combine the high-ℓ CMB, CMB lensing, and BAO data, but leave τ to vary with a uniform prior in the range 0.02 to 0.20. We consider two scenarios, one in which the neutrino mass sum is fixed to its minimal mass ($\sum m_\nu = 0.06$ eV) and one where it may to take on any value allowed by oscillation constraints ($\sum m_\nu \geq 0.06$

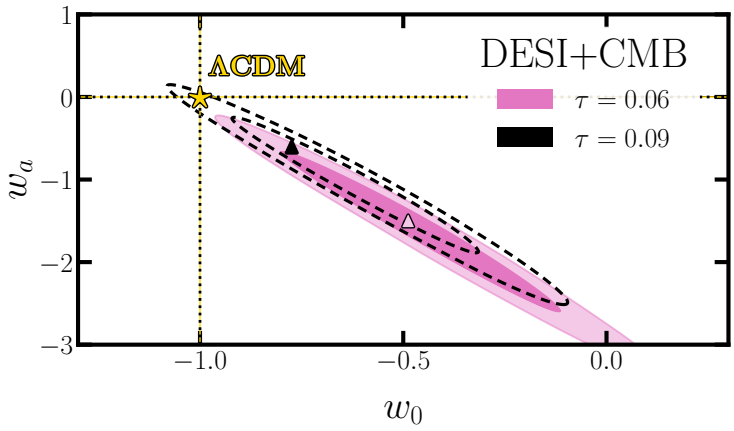


Figure 9.3: Constraints on the $w_0 w_a$ model from the combination of high-ℓ primary CMB data, CMB lensing, DESI BAO and a prior on the optical depth. Locations of posterior maxima are indicated by triangles.

eV). We find

$$\begin{aligned} \tau &= 0.090 \pm 0.012 \quad (\Lambda\text{CDM}) \\ \tau &= 0.095 \pm 0.014 \quad (\Lambda\text{CDM} + \Sigma m_\nu), \end{aligned} \tag{9.2}$$

as illustrated in Fig. 9.4. These values are in $\sim 3\sigma$ disagreement with the optical depth inferred from the full CMB data set (including large scale polarization) in combination with CMB lensing and BAO[3], and are consistent with those found by refs. [248, 367].

9.5 Discussion and conclusions

We have explored the cosmological consequences of a hypothetical higher value of τ and shown that some of the recently-discussed data inconsistencies and tensions with ΛCDM can be alleviated by a higher value of τ. We summarize our main findings in Table 9.1 and further note that with a $\tau = 0.09$ prior the CMB-inferred values of n_s and H_0 increase to 0.968 ± 0.004 and 67.94 ± 0.44 km/s/Mpc within ΛCDM [see also 248, 31]. We have not included data from supernovae, galaxy weak lensing, galaxy clustering or cross-correlations. We note that type Ia supernovae are themselves in mild discrepancy with ΛCDM [109, 509, 9]; however, their joint analysis with the CMB is also affected by the discussion in this chapter.

We caution that at present we have no evidence to doubt the CMB polarization measurements of τ. However, given the difficulties of the measurement and the paradigm shift in terms of fundamental physics that some of the proposed solutions would imply, we think

[3]This estimate takes into account that the data used in our analysis are a subset of this full CMB data set and is based on adding the *Planck* PR3 low-ℓ likelihoods to our fiducial data [254].

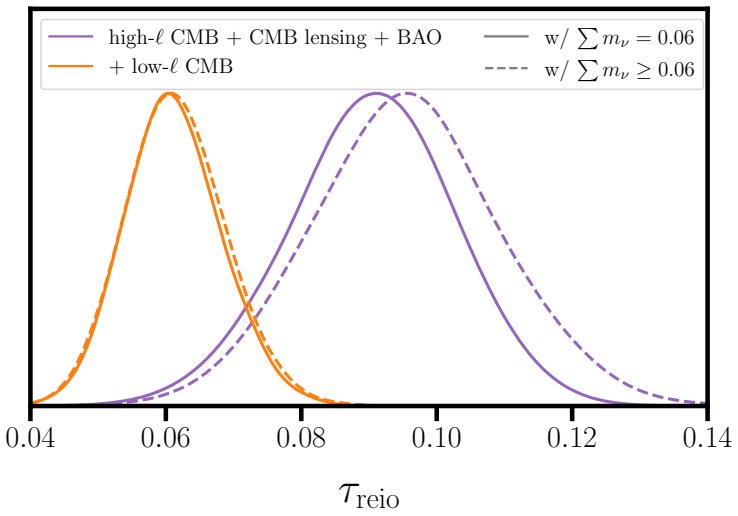


Figure 9.4: When discarding the low-ℓ data from *Planck* we infer $\tau = 0.090 \pm 0.012$ within ΛCDM from the combination of high-ℓ CMB, CMB lensing and BAO, or $\tau = 0.095 \pm 0.014$ for $\Lambda\text{CDM} + \sum m_\nu$ with a physical prior of $\sum m_\nu \geq 0.06$ eV.

Consistent with?	$\tau = 0.06$	$\tau = 0.09$
Planck low-ℓ polarization	✔	✗✗[4]
WMAP low-ℓ polarization	✔	✔
BAO within ΛCDM	✗	✔
$\sum m_\nu \geq 0.06$ eV	✗	✔
Cosmological constant Λ	✗	✔

Table 9.1: Summary of our findings. Here we define ✗as "inconsistent" at $> 2\sigma$, and ✔as "consistent" within 2σ.

that it is important to explore all possibilities. A combination of a statistical fluctuation and mild systematics in one or more of the datasets might also be able to reconcile the results. In fact, τ values as low as 0.08 can bring all of the discrepancies discussed to $< 2\sigma$.

Finally, we have highlighted the crucial role that τ plays not only for neutrino masses, but also for exploring new physics with cosmological observations. Given that current observational uncertainties remain well above the cosmic variance limit [614], we believe that there is a strong scientific motivation for future experiments designed to achieve higher precision, as well as for more direct measurements of the epoch of reionization by observing high redshift galaxies, the 21cm signal or the kSZ effect, among others.

Acknowledgments

We thank Liang Dai, Roger de Belsunce, George Efstathiou, Lloyd Knox, Arthur Kosowsky, Marilena Loverde, Gabriel Lynch, Matthew McQuinn, Seshadri Nadathur, Alexander Reeves, Uroš Seljak, Blake Sherwin and Zachary Weiner for useful discussions. NS additionally thanks Mr. May. NS, GSF, and SF are supported by Lawrence Berkeley National Laboratory and the Director, Office of Science, Office of High Energy Physics of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Software

Scripts used to organize this analysis are publicly available: `disputauble`[5]. This chapter makes use of the `matplotlib` [303], `numpy` [280] and `scipy` [604] packages. We use the Boltzman code `CAMB` [346, 296] to compute theory spectra, and use `GetDist` [344] and `Cobaya` [586] for inference.

[5] https://github.com/NoahSailer/disputauble

Part III

Prospects for future surveys

Chapter 10

Fundamental physics from high redshift large-scale structure surveys

This chapter was originally published as

> Noah Sailer, Emanuele Castorina, Simone Ferraro, and Martin White. "Cosmology at high redshift — a probe of fundamental physics". In: *JCAP* 12.12 (2021), p. 049. DOI: 10.1088/1475-7516/2021/12/049. arXiv: 2106.09713 [astro-ph.CO]

An observational program focused on the high redshift ($2 < z < 6$) Universe has the opportunity to dramatically improve over upcoming LSS and CMB surveys on measurements of both the standard cosmological model and its extensions. Using a Fisher matrix formalism that builds upon recent advances in Lagrangian perturbation theory, we forecast constraints for future spectroscopic and 21-cm surveys on the standard cosmological model, curvature, neutrino mass, relativistic species, primordial features, primordial non-Gaussianity, dynamical dark energy, and gravitational slip. We compare these constraints with those achievable by current or near-future surveys such as DESI and Euclid, all under the same forecasting formalism, and compare our formalism with traditional linear methods. Our Python code FishLSS − used to calculate the Fisher information of the full shape power spectrum, CMB lensing, the cross-correlation of CMB lensing with galaxies, and combinations thereof − is publicly available.

10.1 Introduction

Large Scale Structure (LSS), as probed by fluctuations in the Cosmic Microwave Background (CMB) radiation or large galaxy surveys, provides one of our best windows on fundamental physics. Topics include, but are not limited to, constraints on the expansion history and curvature [557, 575], primordial non-Gaussianity [403, 34], features in the power spectrum (primordial [555] or induced [288, 307, 163, 324, 137]) or running of the spectral index [228,

125], models of inflation [476], dark energy [472, 558], dark matter interactions [1], light relics and neutrino mass [203, 256] and modified gravity [311, 314, 37, 557]. Unfortunately, while many of these extensions of the standard model are well motivated, we do not have strong guidance from theory on which of these many probes is most likely to turn up evidence of new, beyond standard model, physics. However, well understood phenomenology allows us to forecast where our sensitivities to such physics will be highest and our inference cleanest. To work where the inference is cleanest and the noise lowest we should push to high redshift.

Continuous advances in detector technology and experimental techniques are pushing us into a new regime, enabling mapping of large-scale structure in the redshift window $2 < z < 6$ using both relativistic and non-relativistic tracers. This will allow us to probe the metric, particle content and *both* epochs of accelerated expansion (Inflation and Dark Energy domination) with high precision. Moving to higher redshift allows us to take advantage of four simultaneous trends. First, a wider lever arm in redshift leads to rotating degeneracy directions among cosmological parameters and thus tighter constraints. Second, the volume on the past lightcone increases dramatically, leading to much tighter constraints on sample-variance dominated modes. Third, the degree of non-linearity is smaller, so the field is better correlated with the early Universe. Fourth, high precision perturbative models built around principles familiar from high-energy particle physics become increasingly applicable. Indeed at high redshift, we increasingly probe long wavelength modes which are linear or quasi-linear and carry a great deal of "unprocessed" information from the early Universe.

In this chapter we forecast the constraints on cosmological parameters and extensions to the standard model that could be achieved by an observational program focused on the high redshift Universe. We compare the reach of different experiments using a Fisher matrix formalism that builds upon recent advances in cosmological perturbation theory. Our Python code `FishLSS`[1] used to create these forecasts is publicly available, and is described in greater detail in Appendix G.6.

Since at high redshift the non-linear scale shifts to higher k, perturbative models provide a robust method for forecasting that allows for consideration of scale dependent bias, stochastic small-scale power and the decorrelation of the initial and observed modes (see also ref. [124]). We begin in §10.2, where we list observational probes and discuss our assumptions about their number density and linear bias, and describe the experimental parameters of the (future) surveys for which we forecast. Details of the forecasting formalism we employ are provided in §10.3 and our results are presented in §10.4. We explore how these results depend on observational limitations in §10.5. We conclude in §10.6.

10.2 Surveys and probes

In this section we describe the various surveys and probes considered in our forecasts, as summarized in Table 10.1. These include spectroscopic galaxy surveys, which observe some

[1] https://github.com/NoahSailer/FishLSS

combination of emission line galaxies (ELGs), Hα galaxies, or Lyman break galaxies (LBGs), along with interferometers dedicated to 21-cm intensity mapping, and CMB observatories.

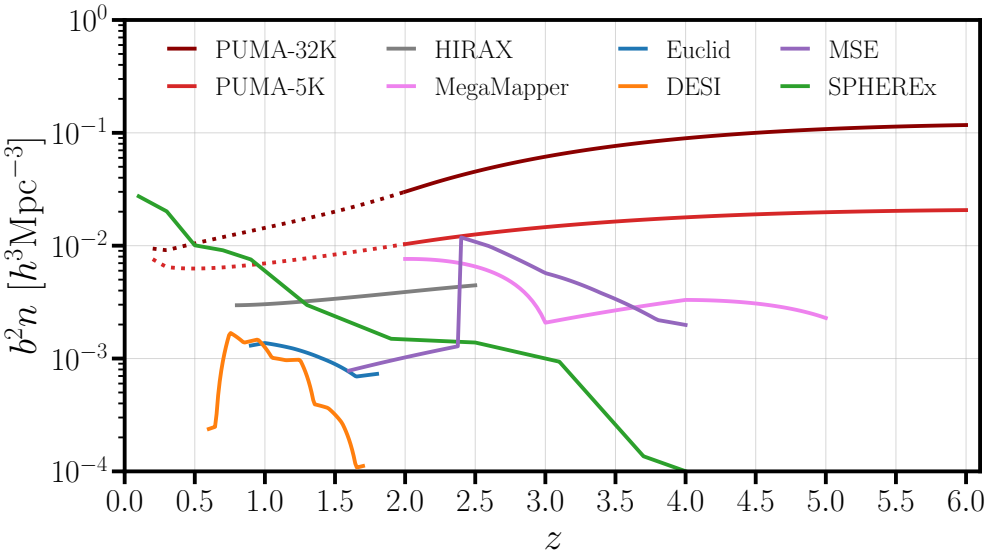


Figure 10.1: b^2n for each of the LSS surveys considered in our forecasts. This combination of linear bias and number density is proportional to the signal-to-noise ratio in linear theory, at fixed $\boldsymbol{k}$ (§10.3). For both PUMA and HIRAX, the "number density" is taken to be the inverse of the thermal + shot noise. We note that while the fiducial PUMA design would enable observations down to $z = 0.2$ (dotted lines), we neglect this region in our forecasts, as it will be well measured by upcoming surveys such as DESI or Euclid which we already include.

In the following subsections we discuss our assumptions about the linear bias and number density of each LSS probe. These values are summarized in Fig. 10.1, where we plot the effective number density b^2n of each survey. This figure already provides some baseline intuition for the results of the following forecasts. In particular, we see that MegaMapper has a comparable noise to the Euclid Hα sample, but covers a much broader redshift range, surveying a volume ~ 2.5 times as large, as illustrated in Fig. 10.4. Similarly, PUMA (-5K and -32K) has a lower noise than either MegaMapper, Euclid or DESI, and its survey volume is ~ 3.2 times that of Euclid[2]. In addition to lower shot noise and more volume, high redshift ($z > 2$) experiments have access to more "linear" modes, i.e. modes that can be modeled very accurately using perturbation theory. In Fig. 10.2 we plot the "effective number of linear (primordial) modes" ($N_{\rm modes}$) that MegaMapper, PUMA, and two upcoming CMB observatories will have access to. For a precise definition of $N_{\rm modes}$, as well as a discussion of its interpretation and caveats on the comparison between CMB and LSS experiments using

[2]This rough factor ignores differences in sky coverage.

this metric, see Appendix G.1. In particular, we see that MegaMapper will have access to roughly 25% more modes than SO, and that PUMA-32K will have access to $\gtrsim 4$ times more modes than CMB-S4, assuming optimistic foregrounds. Our definition of $N_{\rm modes}$ can be easily modified to forecast the error on a measurement of the linear matter power spectrum $P_m(k, z = 0)$, as shown in Fig. 10.3. In particular we see that MegaMapper and PUMA will be capable of measuring the power spectrum out to $k = 0.9\ h\,\mathrm{Mpc}^{-1}$ with an order of magnitude higher precision than Planck at $k \sim 0.2\ h\,\mathrm{Mpc}^{-1}$, despite having finer-grained k-bins. These estimates suggest that cosmological constraints from future, high-redshift LSS surveys will be competitive with CMB measurements. As we show in §10.4, combining the two dramatically improves constraints on ΛCDM parameters, $N_{\rm eff}$, the neutrino masses and other extensions of the standard cosmological model.

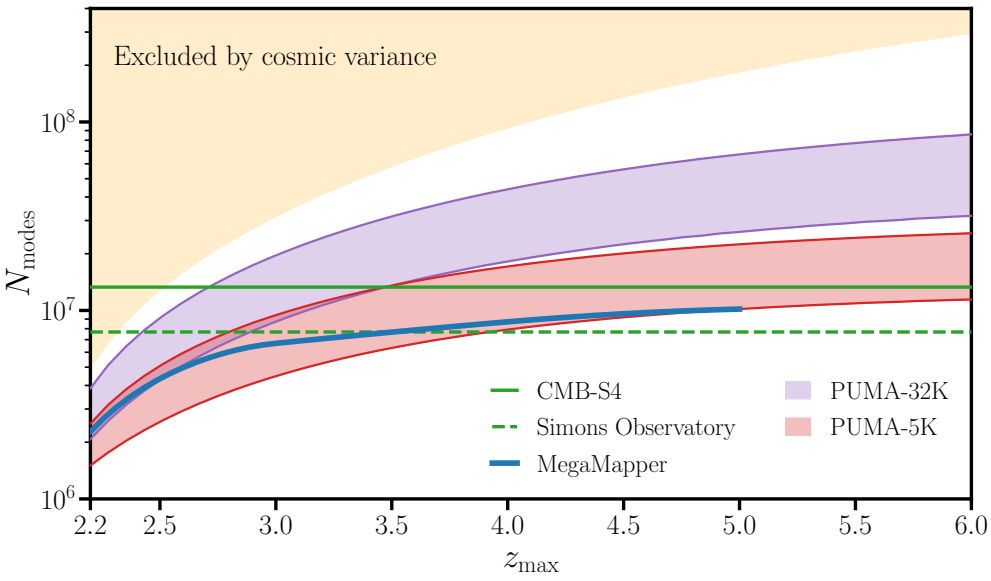


Figure 10.2: $N_{\rm modes}$ as a function of $z_{\rm max}$ for PUMA (-5K and -32K) and MegaMapper ($2 < z < z_{\rm max}$). For PUMA, we consider both optimistic and pessimistic foreground models (§10.3), which are the boundaries of the shaded regions. In green are the number of modes for Simons Obervatory and CMB-S4 like experiments with 7 μK-arcmin and 1.5 μK-arcmin noise levels respectively, both with a $1.4'$ beam and $f_{\rm sky} = 0.4$.

We note that the forecasts presented in this chapter do not include constraints on cosmological parameters coming from cosmic shear or galaxy-galaxy lensing. Both quantities will be measured to exquisite precision by upcoming surveys such as Euclid, the Rubin Observatory and the Roman Space Telescope. However, most of their cosmological constraints from lensing will come from $z < 2$, due to the difficulty in defining high number density samples with well-characterized photometric redshifts and shape measurements at high redshift. Moreover, due to the relatively large redshift uncertainties compared to a spectroscopic survey, detailed features in the power spectrum (such as the BAO) are suppressed, together with

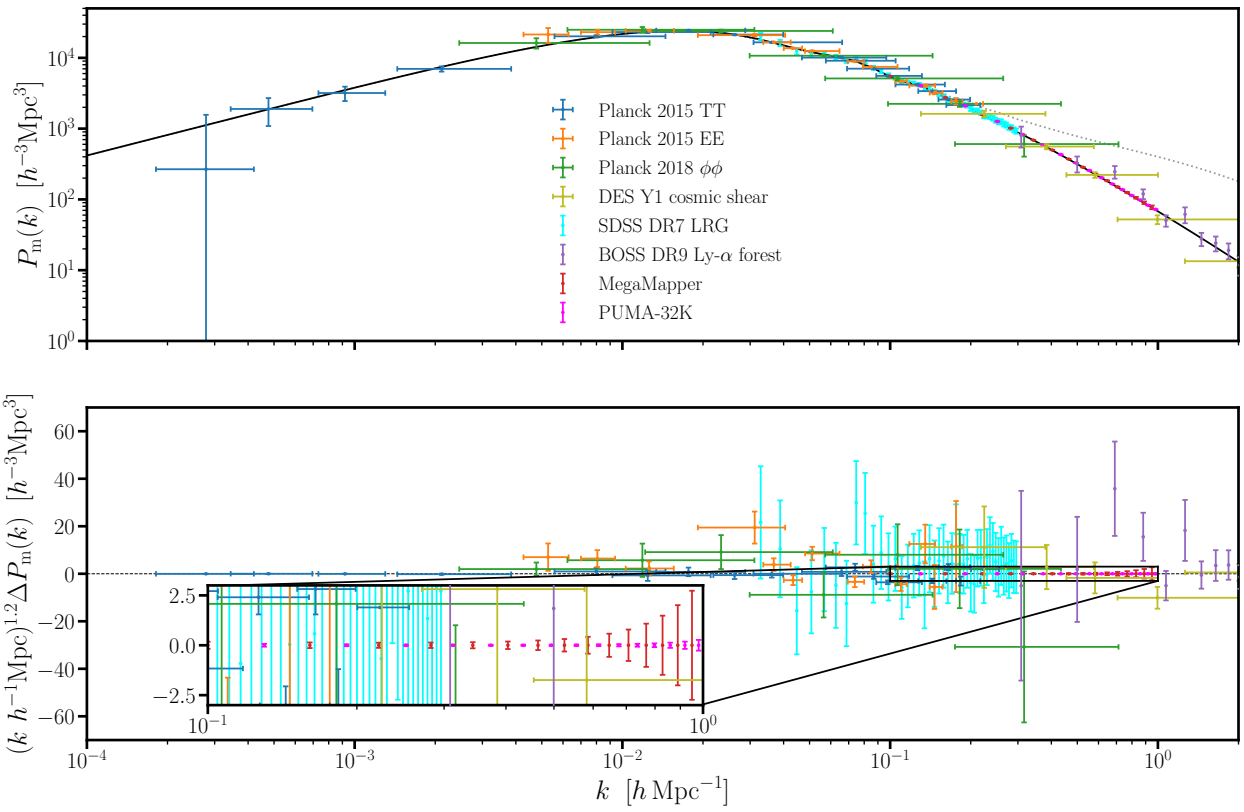


Figure 10.3: Measurements of the linear matter power spectrum at $z = 0$. For both MegaMapper and PUMA-32K we show constraints for 15 linearly space k-bins between $0.1\ h\,{\rm Mpc}^{-1} \lesssim k \lesssim 1\ h\,{\rm Mpc}^{-1}$. This figure is adapted from refs. [473, 126] using https://github.com/marius311/mpk_compilation/.

the majority of the modes in the radial direction. Therefore, the analysis is often performed in a small number (order 10) of 2-dimensional tomographic bins. For this reason, photometric surveys are very effective at measuring the amplitude of structure and/or lensing at low redshift, while spectroscopic surveys can measure the 3-dimensional power spectrum and hence obtain growth and expansion measurements directly. Thus the high-redshift spectroscopic surveys studied here are highly complementary to other upcoming and planned photometric experiments.

Near-term galaxy surveys

The next five years will see a massive influx of LSS data, with spectroscopic galaxy surveys including DESI, Euclid, and SPHEREx coming online. Together these surveys will map out the low redshift universe, taking spectra of hundreds of millions of galaxies.

The Dark Energy Spectroscopic Instrument (DESI) [186] is actively taking data on

the 4 m Mayall Telescope. Using its five thousand fibers and 10 spetrographs, DESI will take spectra of tens of millions of Emission Line Galaxies (ELGs), Luminous Red Galaxies (LRGs), and quasars over the next ~ 5 years. Since ELGs comprise the largest sample of galaxies in the DESI catalogue, providing a significant fraction of the constraining power for most cosmological parameters, we only include the ELG sample in our forecasts. ELGs are active star-forming galaxies which exhibit strong nebular emission lines that originate from gas in the ionized regions around massive stars. This nebular emission includes the [OII] doublet, a distinct spectral signature that can be used to accurately determine a galaxy's redshift. DESI's spectrographs cover a spectral range of 360 nm to 980 nm, enabling [OII] measurements out to $z \sim 1.7$. We convert the mean values of $dN_{\rm ELG}/dz$ (per square degree) in Table 2.3 of ref. [186] to comoving number densities using $n_{\rm ELG} = (180/\pi)^2 \chi^{-2} (d\chi/dz)^{-1} dN_{\rm ELG}/dz$. These values assume a [OII] flux cut of approximately 8×10^{-17} erg cm^{-2} s^{-1}, and account for both selection cuts and target success rates. We follow ref. [186] and approximate the ELG linear bias as $b_{\rm ELG} = 0.84/D(z)$, where $D(z)$ is the linear growth factor normalized to unity at $z = 0$.

The Euclid satellite [37] is a near-infrared space telescope that is currently under developement, with a nominal launch date in 2022. Over its six-year long mission Euclid will measure spectra of tens of millions of galaxies out to $z \sim 2$, mapping out roughly 15,000 square degrees. Euclid consists of a 1.2 m mirror coupled to three imaging and spectroscopic instruments (with wavelength ranges 500-800 nm, 920-1250 nm, and 1250-1850 nm), enabling observations of Hα-emitters from $0.9 < z < 1.8$. For our forecasts we use the reference case number densities $n_{\rm H\alpha}$ and linear biases $b_{\rm H\alpha}$ listed in Table 3 of ref. [95](see also [305]).

The Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx) is a space-based, low resolution spectral survey, with a current launch date in 2024 [197]. It will image $\sim$ 450 million galaxies in both the optical and near-infrared, primarily targeting low redshift ($z < 1$) objects and searching for evidence of primordial non-Gaussianity. We take the number densities and biases from ref. [197], which were split into five samples according to their maximum redshift uncertainty. While the redshift uncertainty can be relatively large for a subset of the combined sample, the key science goal for SPHEREx (i.e. primordial non-Gaussianity) is relatively insensitive to photo-z uncertainties. Moreover SPHEREx was designed to have a very low fraction of redshift outliers, and therefore is especially well-suited to this kind of measurement. Thus we simply add the number densities of the individual samples to get the number density of the combined sample. For the bias of the combined sample we take the appropriate weighted average: $\sum_i b_n n_i / \sum_i n_i$, where b_i and n_i are the bias and number density of the i'th sample.

Future galaxy surveys

While the construction and deployment of various low reshift ($z < 2$) facilities are currently underway, designs for new high-redshift surveys are actively being developed. In this chapter we consider two of these futuristic surveys in detail: MegaMapper and MSE. These surveys

would exploit dropout selection to measure galaxies out to $z \sim 5$, providing cosmological constraints that are competitive with upcoming CMB surveys (see Fig. 10.2).

MegaMapper [528] is a proposed highly-multiplexed spectroscopic instrument that is designed to survey galaxies at high redshift ($2 < z < 5$), covering 14,000 square degrees. Its 6.5 m telescope aperture, $\sim$ 20K fibers, and five year observation time would yield galaxy number densities $> 10^{-4} h^3\,\mathrm{Mpc}^{-3}$ across its redshift range. MegaMapper's primary target would be Lyman Break Galaxies (LBGs), which are abundant, well studied and well understood, representing massive, actively star-forming galaxies with a luminosity approximately proportional to their stellar mass. A sizeable fraction have bright emission lines which could facilitate obtaining a redshift [628].

For the fiducial MegaMapper experiment we use the LBG number densities listed in Table 2 of ref. [228], which are shown in Fig. 10.4. In §10.4 we also consider an idealized MegaMapper-like experiment that detects some redshift-independent fraction of LBGs $\bar{n}/n_{\rm LBG}$ with apparent magnitude $m < 24.5$. This is slightly unrealistic, but provides an easy way to see in what manner our forecasts are sensitive to noise levels. The "idealized" number densities are calculated using a Schechter function of absolute magnitude:

$$\Phi(M_{\rm UV}) = \left(\frac{\ln 10}{2.5}\right) \phi^\star 10^{-0.4(1+\alpha)(M_{\rm UV}-M^\star_{\rm UV})} \exp\left(-10^{-0.4(M_{\rm UV}-M^\star_{\rm UV})}\right). \tag{10.1}$$

We use the best fit values to the UV luminosity function for $\alpha, M^\star_{\rm UV}$, and ϕ^*, which are listed in Table 3 of ref. [628]. The number density is obtained via $n_{\rm LBG} = \int_{-\infty}^{M_c} \Phi(M) dM$, where the cutoff M_c corresponds to galaxies with apparent magnitude $m = 24.5$:

$$M_c \simeq 24.5 - 5\log_{10}\left(\frac{D_L(z)}{10\,\mathrm{pc}}\right) + 2.5\log_{10}(1+z). \tag{10.2}$$

For both the fiducial and idealized MegaMapper surveys, we adopt the bias model of ref. [628], in which the bias of galaxies with apparent magnitude m is approximated to linear order in m and to second order in the scale factor:

$$b(m,z) = A(m)(1+z) + B(m)(1+z)^2 \tag{10.3}$$

with $A(m) = -0.98(m-25) + 0.11$ and $B(m) = 0.12(m-25) + 0.17$. We assume the LBG bias is dominated by the faintest (and most numerous) galaxies, approximating $b_{\rm LBG}(z) = b(m = 24.5, z)$. These approximations to the number density and bias are consistent with the idealised version of MegaMapper considered in ref. [228]. As shown in Fig. 10.4, the fiducial MegaMapper number density is within the range $0.3 \lesssim \bar{n}/n_{\rm LBG} < 1$, depending upon the redshift. In §10.4 we show that taking $\bar{n}/n_{\rm LBG} = 0.5$ yields constraints similar to those from the fiducial MegaMapper survey.

The Maunakea Spectroscopic Explorer (MSE) [583], planned for first light in 2029, will cover over 10,000 square degrees. MSE will couple a 11.25 m mirror with a 1.5 square degree field of view to $\sim$ 4000 fibers, feeding to spectrographs that cover 360 to 1300 nm. This

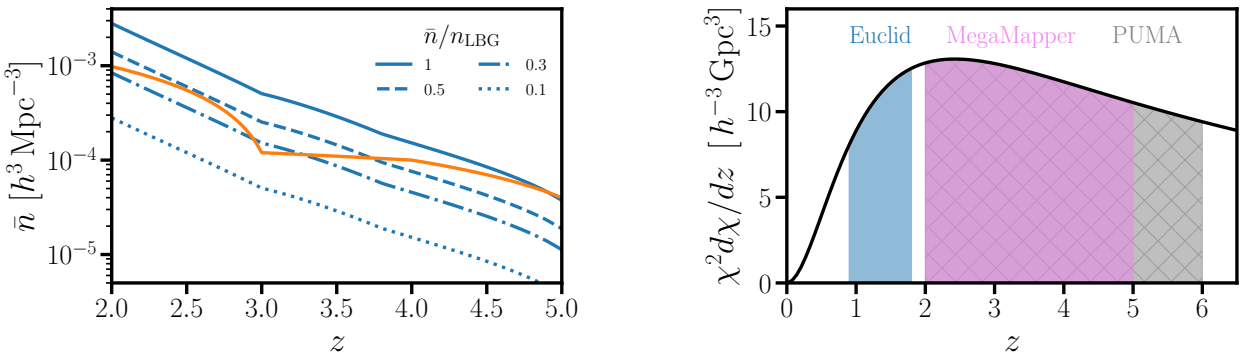


Figure 10.4: *Left:* In blue we plot several values of $\bar{n}/n_{\rm LBG}$, where $n_{\rm LBG}$ is the idealized LBG number density with $m < 24.5$. For comparison, we plot the number density of the fiducial MegaMapper survey in orange. *Right:* Comoving volume per dz per sky area ($\chi^2 d\chi/dz$) as a function of redshift. The shaded blue, purple, and (hatched) grey regions represent the volume covered by Euclid, MegaMapper and PUMA respectively, excluding sky coverage.

design enables the detection of ELGs out to $z \sim 2.4$, and LBGs between $2.4 < z < 4$. For the ELG sample we use the number densities listed in Table 1 of ref. [462], and take $b_{\rm ELG}(z) = D(0)/D(z)$, where $D(z)$ is the linear growth factor. To be consistent with our assumptions regarding LBGs, we use the same linear bias $b(m = 24.5, z)$ for both MegaMapper and the MSE LBG sample. For the LBG sample we estimate the number density via $\int_{-\infty}^{M_c} \Phi(M)E(M)dM$, where Φ and M_c are defined in Eqs. (10.1) and (10.2) respectively, and $E(M)$ is MSE's LBG target efficiency. We take the efficiency to be an average over the templates shown in Figure 2 of ref. [462], assuming that 40% of the LBGs have a negative equivalent width (EW), 30% have $0 <$ EW < 20 and 30% have EW > 20.

In addition to the aforementioned surveys, several high redshift spectroscopic surveys have been proposed or are set to launch within the coming decade, including GAUSS [96], SpecTel [215, 169] and the Roman telescope [563]. We provide a brief description of these surveys below, however, we do not attempt forecasts for these experiments as their survey designs are still in flux.

The Gravitation And the Universe from large Scale-Structures (GAUSS) [96] telescope is a proposed large-scale structure focused space-based mission. GAUSS would consist of a 4 m aperture telescope coupled to both imaging and spectroscopic instruments, covering both the optical and infrared (~ 500 to 5000 nm). The survey would measure spectra of tens of billions of galaxies out to $z \sim 5$, achieving number densities larger than Euclid's across its entire redshift range.

SpecTel [215, 169] is a proposed spectroscopic survey in the southern hemisphere, making it ideally situated for cross-correlations with SO or CMB-S4. SpecTel would couple an 11.4 m dish (with a 5 square degree FoV) to 60,000 fibers − enabling 120 million fiber exposures (each $\geq 4,000$ seconds long) over its survey period − which feed to spectrographs

covering 360 to 1330 nm. Its design would permit observations of LBGs and Lyman Alpha Emitters (LAEs) out to redshift ~ 5, with number densities 2 to 5 times those achievable by a MegaMapper-like survey, depending upon the redshift bin. In §10.4 we explore how MegaMapper's parameter constraints improve with increased LBG number density, which can be interpreted as a rough estimate for SpecTel constraints.

The Roman telescope [563] is a space-based NASA mission set to launch in 2025, carrying a spectrometer that covers the visible to near-infrared spectrum. Under the straw-man proposal for a $2,000\,\mathrm{deg}^2$ survey, constraints from Roman would not be competitive with those from DESI or the other proposed surveys we investigate in detail. Study of trade-offs which would improve Roman's performance in these forecasts is beyond the scope of this chapter and would require a better understanding of achievable depth and area for the galaxy redshift survey given its current constraints.

21-cm intensity mapping

Proof-of-concept interferometers dedicated to 21-cm mapping are currently being deployed [57, 143, 598], while designs for future larger-scale facilities are actively under development [438, 558, 11]. We consider two 21-cm surveys in detail: the Hydrogen Intensity and Real-time Analysis eXperiment (HIRAX) [438] and the Packed Ultra-Wideband Mapping array (PUMA) [558]. HIRAX is a 400–800 MHz radio interferometer that is currently under development for deployment in South Africa, while PUMA is a proposed ultra-wideband, low-resolution and transit interferometric radio telescope operating at 200–1100 MHz. PUMA's fiducial telescope design is a hexagonal array of 32K 6 m dishes, designed to survey half of the sky ($f_{\rm sky} = 0.5$) over five years. We also consider a smaller, proof-of-concept design with only 5K dishes (PUMA-5K). In our noise model (see below), we assume a hex-packed array with a 50% fill factor. This fill factor is accounted for by doubling the number of detectors and setting the observation time to $t_{\rm int} = 5/4$ years. All other couplings, efficiencies, temperatures etc are set to the values quoted in ref. [42]. We use the same noise model for HIRAX, but with a 1K, fully filled array, and $t_{\rm int} = 4$ years.

Below $z \simeq 6$ the majority of the neutral Hydrogen in the Universe resides in dense, self-shielded regions within galaxies [602]. While the details of how HI traces the dark matter are not well understood, current observations and numerical simulations suggest that we can approximate the HI mass within a given halo with $M_{\rm HI}(M_h; z) = A(z) M_h e^{-M_{\rm min}/M_h}$, where $M_{\rm min} = 5 \times 10^9 h^{-1} M_\odot$ [131]. Using this relation, along with a halo number density and bias function, we can then approximate the linear bias of HI by weighting the halo bias with the halo energy density:

$$b_{\rm HI}(z) = \frac{\int_0^\infty b(M_h; z) n(M_h; z) M_{\rm HI}(M_h; z) dM_h}{\int_0^\infty n(M_h; z) M_{\rm HI}(M_h; z) dM_h}. \tag{10.4}$$

We adopt the same conventions as ref. [131] and use the Sheth-Tormen [550] halo bias and mass function.

There are two sources of noise in HI observations: a shot noise due to the finite sampling of the continuous HI field, and the thermal noise from the instrument. For all HI surveys we adopt the thermal noise model described in the appendix of ref. [42] (see also [112, 546]), along with the shot noise from ref. [123, 443]. Further discussion of observational and data-analysis issues related to interferometric, 21-cm intensity mapping can be found in ref. [356].

Experiment	$z_{\rm min}$, $z_{\rm max}$	Sample	$f_{\rm sky}$
DESI [186]	0.6, 1.7	ELG	0.34
Euclid [37, 49, 95]	0.9, 1.8	Hα	0.36
SPHEREx [197, 198]	0.1,4.3	low-z galaxies	0.65
MegaMapper [528]	2, 5	LBG	0.34
MSE [583]	1.6, 4	ELG + LBG	0.24

Experiment	$z_{\rm min}$, $z_{\rm max}$	Sample	$f_{\rm sky}$
PUMA [558]	2, 6	21-cm	0.50
HIRAX [438]	0.8, 2.5	21-cm	0.36

Table 10.1: *Left:* A list of the spectroscopic galaxy surveys considered in our forecasts, along with their redshift ranges, samples and sky coverage. *Right:* A list of 21-cm experiments, their redshift range and sky coverage.

CMB experiments

In addition to LSS probes, we consider external priors from CMB surveys, with configurations similar to the Planck [473], Simons Observatory (SO; [242]) CMB-S4 [1] and LiteBIRD [569] surveys. To obtain the Planck Fisher matrix, we use the bands, beams and sensitivities from Table 4 of ref. [473]. We include both temperature and polarization anisotropies, and we adjust the minimum $\ell_{\rm min,\,P}$ in polarization to match the published value of uncertainty on the optical depth $\sigma(\tau) = 0.007$. We have checked that this method gives good agreement with the published uncertainties on the standard cosmological parameters, while noting that the main purpose of this chapter is not forecasting parameters for the next generation of CMB experiments, but rather to highlight the improvements from high-redshift LSS surveys. Because of this, we consider a simplified setup in which detector noise is treated as white noise uncorrelated between different frequency channels, and we do not include any effects from the atmosphere on large scales (beyond the scale cut $\ell_{\rm min} = 30$), and do not include any instrumental or foreground nuisance parameters. For SO and CMB-S4 we take $\ell_{\rm max,\,T} = 3000$ in temperature and $\ell_{\rm max,\,P} = 5000$ in polarization, and $\ell_{\rm min,\,T,\,P} = 30$ for both. Since SO and CMB-S4 are ground based and not expected to measure the very largest scales due to contamination from the atmosphere, we always add a low-ℓ Planck (or LiteBIRD) Fisher matrix to SO or CMB-S4. This is accomplished by splitting the SO/S4/Planck forecasts into high- and low-ℓ Fisher matrices (with cutoff $\ell = 30$), and adding the low-ℓ Planck Fisher matrix to the high-ℓ SO/S4 Fisher matrix. This simplistic addition is justified since the two Fisher matrices use a disjoint set of multipoles, and are therefore largely uncorrelated (we

expect any mode coupling from lensing or non-Gaussian foregrounds to be small). A fiducial model for the foregrounds is added to the total power spectrum in temperature, following the model of ref. [200] (with the modifications listed in ref. [242]). We neglect any foregrounds in polarization, noting that point source masking should be very effective at mitigating their effect on small scales. Unless stated otherwise, we use the 2-point function of the unlensed primary CMB temperature and polarization fluctuations for the derivatives with respect to the cosmological parameters, and do not include CMB lensing reconstruction. This allows us to simply add the Fisher matrices for CMB and LSS probes. Another feature of using the unlensed CMB power spectrum is that it should provide accurate forecasts from high-ℓ temperature. If we included lensing the damping tail would be proportional to the lensing amplitude, and therefore largely degenerate with the foreground amplitude and shape, as well as baryon effects in the lensing power spectrum which is not marginalized over here. This can lead to artificially tight constraints [392]. We also consider a futuristic LiteBIRD-like space based CMB experiment, capable of reaching the cosmic variance limit $\sigma(\tau) \approx 0.002$ from $\ell < 30$. For this we take a 1.4μK-arcmin experiment in polarization with a 33 arcmin Gaussian beam covering 65% of the sky. Although not considered in this chapter, we note that an instrument like PICO [276] would also be highly complementary to LSS probes.

10.3 Forecasting formalism

In this section we outline our forecasting procedure, based on the Fisher formalism. We examine three primary observables of interest: primary anisotropies in the CMB (as described in §10.2), the redshift space galaxy power spectrum (§10.3), and CMB lensing (§10.3). We take ΛCDM as our fiducial cosmology, with ΛCDM values given in Table 10.2 and $N_{\rm eff} = 3.046$, consistent with refs. [473, 475]. We assume a normal neutrino hierarchy with minimal mass: $m_{\nu_1} = 0$, $m_{\nu_2} = 0.01$ eV and $m_{\nu_3} = 0.05$ eV. Unless stated otherwise the neutrino mass $M_\nu = \sum_i m_{\nu_i}$ and $N_{\rm eff}$ are held fixed in all forecasts. A list of all of the free parameters in our base model is given in Table 10.2.

Full shape measurements

Both historically and currently, most forecasts use the linear power spectrum and scale-independent bias, or the non-linear matter power spectrum computed using fitting formulae, to model the power spectrum of a biased tracer [95, 235, 186, 525, 634, 532, 443, 107, 596, 33, 558]. In this chapter, we use more sophisticated models based on perturbation theory and a general bias expansion (see refs. [124, 147, 104, 142] for related work on forecasts and ref. [401] for inclusion of scale-dependent biases in the halo model). We consider a general quadratic bias model and non-linear evolution up to one-loop order in perturbation theory. On large scales and at high redshifts, perturbation theory provides an accurate model that allows us to take into account scale-dependent bias, broadening of the acoustic peaks and mode-

coupling due to non-linear evolution. Including these effects should improve the reliability of the forecasts. We further discuss the differences in linear vs. non-linear forecasting in §10.4.

To include the non-linear effects and more complex bias model we use the `velocileptors` `LPT_RSD` module [131, 137, 134], which takes the linear power spectrum and biases as inputs, and self-consistently calculates the non-linear redshift-space power spectrum $P_{gg}^{\text{1-loop}}(k,\mu)$ to one-loop order in perturbation theory. We use the linear CDM+baryon power spectrum $P_{cb}^{\text{lin}}(k,z)$ as our input to `velocileptors`, calculated using `CLASS` [97], along with a second order biasing scheme:

$$\delta_g(\boldsymbol{x},z) = b\,\delta_{cb}(\boldsymbol{x},z) + \frac{1}{2}b_2\,\delta_{cb}^2(\boldsymbol{x},z) + b_s\,s^2(\boldsymbol{x},z) \tag{10.5}$$

where δ_g is the non-linear tracer density contrast, δ_{cb} is the non-linear CDM+baryon density contrast, and $s_{ij} = (\partial_i\partial_j/\partial^2 - \delta_{ij}/3)\delta_{cb}$ is the shear field. To better connect our formalism to previous literature here, and elsewhere in the text, these biases are in the Eulerian frame[3]. However, the `LPT_RSD` module computes $P_{gg}^{\text{1-loop}}(\boldsymbol{k})$ using Lagrangian perturbation theory, which has been shown to perform well in these situations [134] and allows for a simple treatment of models with complex features in the power spectrum [137]. To compute the Lagrangian biases, which are inputs to `velocileptors`, we perform the following rotation [127, 51]:

$$b^{\text{L}} = b - 1 \qquad b_2^{\text{L}} = b_2 - \frac{8}{21}(b-1) \qquad b_s^{\text{L}} = b_s + \frac{2}{7}(b-1). \tag{10.6}$$

We use the linear biases given in §10.2. The fiducial values of b_2 and b_s are set by assuming $b_2^{\text{L}} = b_s^{\text{L}} = 0$. This choice is only a rough approximation whose validity is pending more data. In all forecasts we marginalize over these non-linear biases, making our results mostly insensitive to their fiducial values (see Appendix G.5 for a more detailed discussion).

To model the effects of small-scale physics, we include a handful of counterterms (α_n) and stochastic contributions (N_n) in our fiducial power spectra:

$$P_{gg}(\boldsymbol{k}) = P_{gg}^{\text{1-loop}}(\boldsymbol{k}) + \left(\alpha_0 + \alpha_2\mu^2 + \alpha_4\mu^4\right)(k/k_*)^2\,P_{cb}^{\text{Zel}}(\boldsymbol{k}) + N_0 + N_2(\mu k)^2 + N_4(\mu k)^4, \tag{10.7}$$

where N_2 and N_4 encode small-scale velocities (i.e. FoG effects and redshift errors, if appropriate), N_0 is the shot noise, $k_* = 1\,h\,\text{Mpc}^{-1}$, P_{cb}^{Zel} is the CDM+baryon power spectrum in the Zel'dovich approximation, and $\mu = \hat{\boldsymbol{k}}\cdot\hat{\boldsymbol{n}}$ is the cosine of the angle between the wavevector and the line of sight. Throughout we approximate $P_{gg}(k,\mu) \simeq \sum_{\ell=0}^{4}\mathcal{L}_\ell(\mu)P_{gg,\ell}(k)$ by its first three non-vanishing multiples, where $\mathcal{L}_\ell(\mu)$ is the Legendre polynomial of degree ℓ. For the surveys, redshifts and k-ranges considered in this chapter we find that this approximation is good to 2%.

We determine the fiducial value of α_0 by fitting the `velocileptors` output to `CLASS`' default version of `HaloFit` [576]. The fiducial values for the higher order stochastic terms (α_2, α_4) are set to zero purely for convenience. We take the Poisson value $\bar{n}_{\text{gal}}^{-1}$ for the

[3]We assume cubic bias operators are fixed and generated by time evolution in Lagrangian perturbation theory.

fiducial value of N_0 and set $N_2 = -N_0\sigma_v^2$, where σ_v is a typical velocity dispersion. For galaxy surveys we set the velocity dispersion to 100 km/s, converted to comoving length units: $\sigma_v = (1+z)(100\,\mathrm{km/s})/H(z)$. Finger-of-god effects in 21-cm maps are much smaller, driven by a small number of satellites (occupying ∼10% of the sample) with typical rms velocities ≈ 40 km/s [602]. Thus for 21-cm surveys, we choose a fiducial velocity dispersion of 10 km/s ($\sim \sqrt{0.1}\,40$ km/s). The fiducial value of N_4 is set to zero purely for convenience. In all forecasts we marginalize over these counterterms and stochastic contributions, making our results largely insensitive to their fiducial values, as discussed in Appendix G.5.

The modeling of the power spectrum for 21-cm surveys is slightly more complex than for galaxy surveys. Radio interferometers measure the 21-cm signal in intensity units, so that the power spectrum has units of (temperature)^2 × volume:

$$P_{\rm HI}^{\rm obs}(\boldsymbol{k}) = T_b^2\,[P_{\rm HI}(\boldsymbol{k}) + P_{\rm SN}] + P_{\rm N}, \tag{10.8}$$

where $P_{\rm HI}(\boldsymbol{k})$ is the non-linear HI power spectrum (taking the same form as Eq. 10.7, with $P_{\rm SN}$ replacing N_0), $P_{\rm SN}$ is the shot-noise coming from halos that contribute to the HI signal, P_N is the instrumental noise [42], and T_b is the mean 21-cm brightness temperature [232]:

$$T_b(z) = 188\,h(1+z)^2E^{-1}(z)\Omega_{\rm HI}(z)\,\mathrm{mK}\,, \tag{10.9}$$

where $E(z) \equiv H(z)/H_0$ and $\Omega_{\rm HI}(z) \equiv \rho_{\rm HI}(z)/\rho_c(z)$ is the cosmic density of neutral hydrogen, taken from the fitting formula of ref. [159]. In all forecasts that include 21-cm surveys we marginalize over the 21-cm brightness temperature by default as its true value is still largely uncertain.

Because our fiducial model includes massive neutrinos, we need to specify their perturbative treatment in our setup. It is by now well established that halos and galaxies can be well approximated as biased tracers of the dark matter and baryon fluids only [601, 122, 121, 600, 364, 426, 231]. For this reason our real-space power spectrum calculation goes through as for the no-neutrino case if we make the trivial replacement of the total matter field with P_{cb}, the dark matter plus baryon spectrum, as described above. We also assume the Green's function associated with higher order perturbative solutions is very well approximated by the Einstein-de-Sitter one, which implies the structure of the loops remains unchanged. This is an excellent approximation for the small fiducial value of neutrino masses considered in this chapter [48]. In principle one needs to consider the scale-dependence of the growth rate, which we shall neglect for the same reason mentioned above.

For our full shape forecasts we work at the power spectrum level. As in Eq. (10.7), let $P_{gg}(k,\mu,z)$ denote the observed nonlinear redshift-space power spectrum of the matter tracer. Let i index a (μ_i, k_i) pair, e.g. $P_{gg,i}(z) = P_{gg}(k_i,\mu_i,z)$. We divide each survey into linearly spaced redshift bins $\{B_m;\, m = 1,2,\cdots\}$, and calculate the corresponding Fisher

matrix in each bin as[4]:

$$F_{ab}(B_m) = \sum_{ij} \frac{\partial P_{gg,i}(z)}{\partial \theta_a} C_{ij}^{-1}(z) \frac{\partial P_{gg,j}(z)}{\partial \theta_b}\bigg|_{z=z_m}$$
$$\text{where} \quad C_{ij}(z) = \delta_{ij}^K \frac{4\pi^2}{k_i^3 V_m \, d\mu \, d\ln k} \left[P_{gg,i}(z)\right]^2 . \tag{10.10}$$

Here V_m is the comoving volume of bin B_m, which implicitly includes a factor of $f_{\rm sky}$, $z_m = (z_{\rm max} + z_{\rm min})/2$ is the central redshift in bin B_m, and θ_a represents either a cosmological or a nuisance parameter. Throughout we assume that $P_{gg}(k, \mu, z)$ is even in μ, and use Simpson's rule to evaluate Eq. (10.10) with 2,000 logarithmically spaced k-bins (ranging from 5×10^{-4} to 1 $h\,\mathrm{Mpc}^{-1}$) and 100 linearly spaced μ-bins (ranging from 0 to 1). We limit the domain of integration by multiplying the summand by an appropriate window function. To calculate the Fisher matrix of the full survey we make the standard approximation that correlations between non-overlapping redshift bins are negligible, so that the final Fisher matrix is obtained by summing over redshift bins:

$$F_{ab} = \sum_m F_{ab}(B_m). \tag{10.11}$$

In Eq. (10.10) we neglect any temporal or spatial variation in the number density within each bin. This effect can be included analytically [611], and can lead to significant error in the forecast for sufficiently large redshift bins, or for sufficiently low number densities. However, as we show in Appendix G.2, for the surveys and redshift bins considered in this chapter these effects are negligibly small. We thus ignore this correction, and take Eq. (10.10) at face value.

In all forecasts we take $k_{\rm min} = \max\left[0.003\,h\,\mathrm{Mpc}^{-1}, 2\pi/V_m^{1/3}\right]$ for each redshift bin. Since the clustering signal is low for $k < k_{\rm min}$, our results would not change significantly if a slightly different $k_{\rm min}$ were assumed, with the exception of local primordial non-Gaussianity, as discussed in Sec. 10.4. Unless explicitly stated otherwise we take $k_{\rm max}(z) = k_{\rm nl}(z)$, where $k_{\rm nl}(z) = \Sigma^{-1}(z)$ is the non-linear scale set by the RMS displacement in the Zel'dovich approximation:

$$\Sigma^2(z) = \int \frac{dk}{6\pi^2} P^{\rm lin}(k, z), \tag{10.12}$$

where $P^{\rm lin}(k, z)$ is the linear matter power spectrum. For 21-cm surveys, we also include a foreground wedge [42] that further constrains the limits of integration:

$$k_\parallel > \max\left[k_\perp \frac{\chi(z) H(z)}{c(1+z)} \sin(\theta_w(z)) \;,\; k_\parallel^{\rm min}\right] \quad \text{where} \quad \theta_w(z) = N_w \frac{1.22}{2\sqrt{0.7}} \frac{\lambda_{\rm obs}}{D_{\rm phys}} \tag{10.13}$$

where $\lambda_{\rm obs} = (1+z)\,21\,\mathrm{cm}$ is the observed HI wavelength, $D_{\rm phys}$ is the physical diameter of the dish, $k_\parallel = k\mu$ and $k_\perp = k\sqrt{1-\mu^2}$. We take $D_{\rm phys} = 6$ m for both HIRAX and

[4]The shot noise $N_0(z) = 1/n_{\rm gal}(z)$ is included in our definition of the observed power spectrum, so it does not appear explicitly in our covariance.

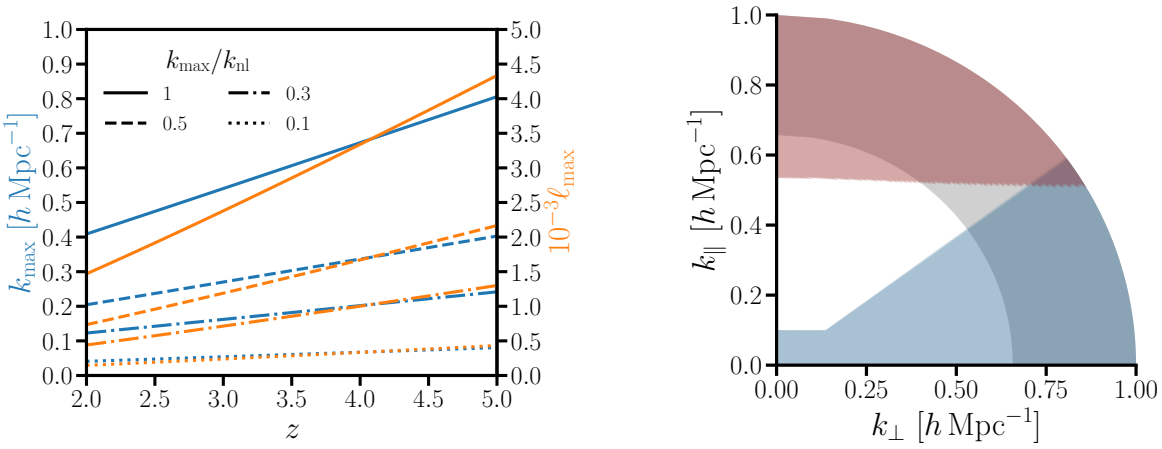


Figure 10.5: *Left*: $k_{\rm max}(z)$ and $\ell_{\rm max}(z) = k_{\rm max}(z)\chi(z)$, assuming that $k_{\rm max}$ is some constant fraction of $k_{\rm nl}(z) = \Sigma^{-1}(z)$. *Right*: Domain of integration in the $k_\parallel - k_\perp$ plane at $z \simeq 3.9$. The grey region is excluded by the $k < k_{\rm nl}$ constraint. For 21-cm surveys the blue region is excluded by the (pessimistic) foreground wedge. The red region is excluded by FoG-like effects, chosen so that $N_0\sigma_v^2(k\mu)^2$ is less than 20% of the fiducial power. In this figure we plot the red region for a MegaMapper-like survey (it would shift to higher $k_\parallel$ for a 21-cm survey).

PUMA. For each 21-cm survey we consider "optimistic" and "pessimistic" foreground cases, in which $N_w = 1$, $k_\parallel^{\rm min} = 0.01\,h\,{\rm Mpc}^{-1}$ and $N_w = 3$, $k_\parallel^{\rm min} = 0.1\,h\,{\rm Mpc}^{-1}$ respectively. This is consistent with the definitions in ref. [42]. Finally, we note that non-perturbative effects like Fingers-of-God (FoG) limit the information that we can extract from high values of μ. The importance of FoG is not known a priori, but roughly speaking they can be accounted for by a stochastic term $N_0\sigma_v^2(k\mu)^2$ (see Eq. 10.7) on sufficiently large scales. We therefore make the conservative choice of integrating only over modes where this term is less than 20% of the fiducial power $P_{gg}(k,\mu)$. We have checked that this choice in cutoff has an insignificant impact on our forecasts, as discussed in Appendix G.5 and shown in Fig. G.4. In Fig. 10.5 we show both the foreground wedge (for the pessimistic case) and the region excluded by FoG-like effects for our fiducial MegaMapper survey.

CMB lensing and cross-correlations

Weak lensing of the CMB is a tracer of the matter distribution integrated along the line of sight [342], and is one of the central science cases for upcoming CMB surveys [242, 569, 1, 81]. Throughout we work under both the Born and Limber approximations so that C_ℓ^{XY} may be expressed as a single integral:

$$C_\ell^{XY} = \int_0^{\chi_*} d\chi \frac{W^X(\chi)W^Y(\chi)}{\chi^2} P_{\delta_X\delta_Y}\left(k_\perp = \frac{\ell+1/2}{\chi}, k_\parallel = 0\right), \tag{10.14}$$

where X and Y denote either the CMB lensing convergence κ or the galaxy sample g, and $\delta_\kappa = \delta_m$. In Eq. (10.14) we have included the lowest order correction to the Limber approximation, replacing $\ell \to \ell + 1/2$, which increases the accuracy to $\mathcal{O}(\ell^{-2})$ [366]. In principle both approximations could be relaxed, but in practice we do not expect them to have a large impact on the parameters constraints [365, 328, 485, 384, 219, 221, 101]. The projection kernels take the form:

$$W^\kappa(\chi) = \frac{3}{2}\Omega_m H_0^2(1+z)\frac{\chi(\chi_* - \chi)}{\chi_*} \quad \text{and} \quad W^g(\chi) \propto H(z)\frac{dN}{dz}, \tag{10.15}$$

where W^g is normalized so that $\int d\chi W^g(\chi) = 1$, and χ_* is the comoving distance to the surface of last scattering. For simplicity we neglect any contribution from magnification bias, as this will be measured by future LSS surveys to high enough accuracy to be treated as an essentially fixed parameter [131].

The nonlinear real space power spectra (P_{gg} and P_{gm}) are calculated to one-loop in LPT using `velocileptors`' `cleft_fftw.py`, which includes a second order biasing scheme (b, b_2, b_s), two counterterms (α_0 and α_x) for the auto- and cross-correlations, and a shot noise contribution to P_{gg}. Note that lensing cross-correlations introduces only one new nuisance term (α_x), and that P_{gm} and P_{gg} manifestly have different dependencies on the bias parameters. Thus lensing cross-correlations play a role in breaking degeneracies, improving the constraints on parameters which are highly degenerate with bias (e.g. M_ν), as discussed in Appendix G.4 (see Fig. G.3). The fiducial values for the biases, counterterms and shot noise are identical to those for the full shape power spectrum, and are listed in Table 10.2.

For both the cross-correlation $C_\ell^{\kappa g}$ and the galaxy auto-correlation C_ℓ^{gg} we approximate Eq. (10.14) with 100 integration points that are linearly spaced in the redshift interval corresponding to the galaxy sample g. In principle the power spectra should be evaluated at each integration point, but this is computationally expensive to do in practice. To speed up our calculations we split the redshift interval into m linearly-spaced bins, each of which contain $100/m$ integration points. Within each bin the power spectrum is evaluated at a single effective redshift [416]:

$$z_{\rm eff}^{XY} = \frac{\int d\chi W^X(\chi) W^Y(\chi) z(\chi)/\chi^2}{\int d\chi W^X(\chi) W^Y(\chi)/\chi^2}, \tag{10.16}$$

which is chosen to cancel the linear order correction in the expansion of $P_{\delta_X\delta_Y}(k,z)$ about $z_{\rm eff}^{XY}$. We choose m to be the smallest integer that yields bins with widths $\Delta z \leq 0.2$, resulting in $\mathcal{O}(0.1\%)$ accuracy for $30 < \ell < 500$ and $1 < z < 6$. Achieving a similar accuracy for the convergence power spectrum via direct integration is computationally inefficient. We instead calculate $C_\ell^{\kappa\kappa}$ with `CLASS`, using its default version of `HaloFit` to model nonlinearities.

We consider both SO and S4-like experiments when providing lensing constraints. For SO we use the noise N_ℓ^κ from the SO noise calculator while for S4, we use the lensing noise curve described on the CMB-S4 Wiki. These curves are plotted in Fig. 10.6. Both of these noises include atmospheric contributions, and are calculated with an iterated minimum

variance (MV) quadratic estimator and MV Internal Linear Combination (ILC) of both CMB temperature ($30 < \ell_{\rm T} < 3000$) and polarization ($30 < \ell_{\rm P} < 5000$) data.

In our forecasts we work at the power spectrum level and split our galaxy sample g into n non-overlapping redshift bins, denoted by $g_1, \cdots, g_n$, so that our averaged data vector is $\boldsymbol{\mu}_\ell^T = (C_\ell^{\kappa\kappa}, C_\ell^{\kappa g_1}, \cdots, C_\ell^{\kappa g_n}, C_\ell^{g_1 g_1}, \cdots, C_\ell^{g_n g_n})$. The corresponding Fisher matrix is

$$F_{ij} \simeq \sum_{\ell=\ell_{\rm min}}^{\ell_{\rm max}} \frac{\partial \boldsymbol{\mu}_\ell^T}{\partial \theta_i} \boldsymbol{C}_\ell^{-1} \frac{\partial \boldsymbol{\mu}_\ell}{\partial \theta_j}, \tag{10.17}$$

with the elements of the covariance $\boldsymbol{C}_\ell$ given by[5]

$$(\boldsymbol{C}_\ell)_{XY} = \frac{2}{2\ell+1} \times \begin{cases} \frac{1}{f_{\rm sky}^{\rm CMB}} (C_\ell^{\kappa\kappa} + N_\ell^\kappa)^2 & X = Y = \kappa\kappa \\ \frac{1}{2 f_{\rm sky}^{\cap}} \left[(C_\ell^{\kappa\kappa} + N_\ell^\kappa) C_\ell^{g_i g_i} \delta_{ij}^K + C_\ell^{\kappa g_i} C_\ell^{\kappa g_j} \right] & X = \kappa g_i,\ Y = \kappa g_j \\ \frac{1}{f_{\rm sky}^{\rm LSS}} (C_\ell^{g_i g_i})^2 \delta_{ij}^K & X = g_i g_i,\ Y = g_j g_j \\ \frac{1}{f_{\rm sky}^{\rm CMB}} (C_\ell^{\kappa\kappa} + N_\ell^\kappa) C_\ell^{\kappa g_i} & X = \kappa\kappa,\ Y = \kappa g_i \\ \frac{f_{\rm sky}^{\cap}}{f_{\rm sky}^{\rm LSS} f_{\rm sky}^{\rm CMB}} (C_\ell^{\kappa g_i})^2 & X = \kappa\kappa,\ Y = g_i g_i \\ \frac{1}{f_{\rm sky}^{\rm LSS}} C_\ell^{g_i g_i} C_\ell^{\kappa g_i} \delta_{ij}^K & X = g_i g_i,\ Y = \kappa g_j \end{cases} \tag{10.18}$$

where $f_{\rm sky}^{\rm LSS}$, $f_{\rm sky}^{\rm CMB}$ and $f_{\rm sky}^{\cap}$ are the sky fractions of the LSS survey, CMB lensing map, and overlap of the two respectively. Throughout we assume $f_{\rm sky}^{\rm CMB} = 0.4$, which is consistent with an SO-like experiment, and assume full spatial overlap of the CMB and LSS surveys for cross-correlations: $f_{\rm sky}^{\cap} = \min(f_{\rm sky}^{\rm LSS}, f_{\rm sky}^{\rm CMB})$. In all forecasts we take $\ell_{\rm min} = 30$. To remove information from modes that are too small to be accurately modeled using PT, we multiply the appropriate elements of the covariance matrix ($C_\ell^{g_i g_i}$ and $C_\ell^{\kappa g_i}$) by 10^{20} whenever $\ell > k_{\rm nl}(z_i)\chi(z_i)$, where z_i is central redshift in the i'th redshift bin.

Due to projection effects, the lensing auto-power spectrum receives substantial contributions from small scales, beyond the reach of perturbation theory. For such modes, baryonic effects cannot be neglected either, further complicating the model predictions. To be conservative we impose a cut ($\ell < 500$) on the lensing angular modes included in our forecasts, but note that this is active area of research and improvements are likely to possible by the time the data arrive [236, 106, 391, 148]. For $\ell < 500$, baryonic feedback contributes at the subpercent level to the convergence power spectrum [149], while at $\ell = 500$, only $\sim 15\%$ of the power comes from $k > 0.2\,h\,{\rm Mpc}^{-1}$.

[5] As in §10.3, we include the shot noise in our fiducial $C_\ell^{g_i g_i}$, so it does not explicitly appear in our covariance. The $f_{\rm sky}$ factors are determined using the following argument. Suppose that $\hat{C}_\ell^X = (N^X)^{-1} \sum_{m \in S^X} |a_{\ell m}^X|^2$ is some unbiased estimator for C_ℓ^X that sums over some m's in a set S^X, and N^X is the number of m's in S^X. It is then straightforward to show that ${\rm Cov}[\hat{C}^X, \hat{C}^Y] \propto N^\cap / N^X N^Y$, where $N^\cap$ is the number of m's in the intersection of S^X and S^Y. Thinking of N^X as $(2\ell+1) f_{\rm sky}^X$, we obtain the factors of $f_{\rm sky}$ in Eq. (10.18).

Parameter	Definition	Fiducial value
h	Hubble parameter: $H_0/100$ km/s/Mpc	0.677
$\ln(A_s)$	ln(primordial amplitude)	-19.97
n_s	spectral index	0.96824
ω_c	fractional dark matter density: $\Omega_c h^2$	0.11923
ω_b	fractional baryon density: $\Omega_b h^2$	0.02247
τ	optical depth to reionization	0.0568
b	linear Eulerian bias: $\delta_g \supset b\delta_{cb}$	survey dependent
b_2	quadratic Eulerian bias: $\delta_g \supset \frac{1}{2}b_2\delta_{cb}^2$	$(8/21) \times (b-1)$
b_s	shear bias: $\delta_g \supset b_s s^2$	$-(2/7) \times (b-1)$
α_0	$P_{gg} \supset \alpha_0 (k/k_*)^2 P_{cb}^{\rm Zel}$	$1.22 + 0.24\, b^2\, (z - 5.96)$
α_{2n}	$P_{gg} \supset \alpha_{2n}\mu^{2n}(k/k_*)^2 P_{cb}^{\rm Zel}$, $n = 1, 2$	0
N_0	Poisson white noise: $P_{gg} \supset N_0$	survey dependent
N_2	FoG-like contribution: $P_{gg} \supset N_2(k\mu)^2$	survey dependent
N_4	FoG-like contribution: $P_{gg} \supset N_4(k\mu)^4$	0
T_b	21-cm brightness temperature	Eq. (10.9)
α_x	$P_{gm} \supset \alpha_x (k/k_*)^2 P_{cb}^{\rm Zel}$	0

Table 10.2: A list of the free parameters in our base model, which are marginalized over in all forecasts unless explicitly stated otherwise. In the first set of rows we list the ΛCDM parameters and their fiducial values. In the second set of rows we list the nuisance terms that we use to model redshift-space clustering. The last row shows the counterterm introduced when including lensing cross-correlations. Here δ_g is the non-linear density contrast of the tracer with power spectrum P_{gg}, δ_{cb} is the non-linear density contrast of CDM+baryons, $P_{cb}^{\rm Zel}$ is the power spectrum of the CDM+baryons within the Zeldovich approximation, $s_{ij} \equiv (\partial_i\partial_j/\partial^2 - \delta_{ij}/3)\delta_{cb}$ is the shear field, $s^2 \equiv s_{ij}s_{ij}$, and $k_* = 1\,h\,{\rm Mpc}^{-1}$.

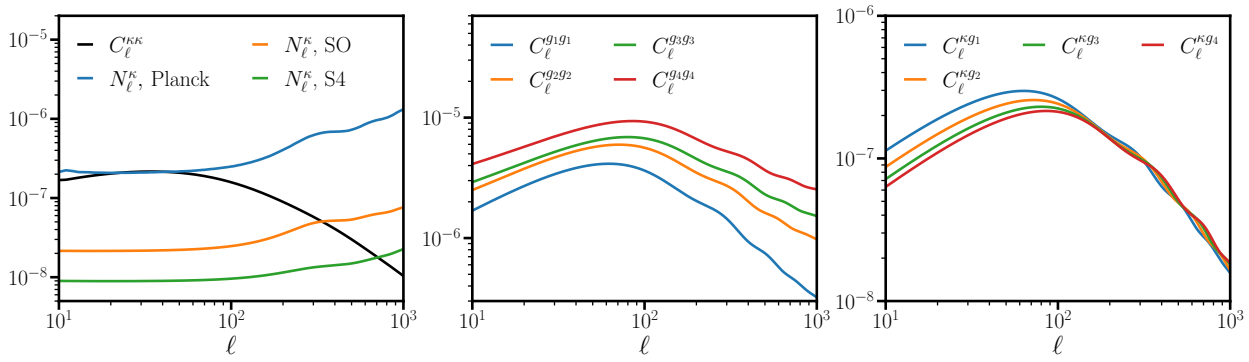


Figure 10.6: *Left:* Lensing noise N_ℓ^κ for Planck, Simons Observatory and CMB-S4. The convergence power spectrum is shown in black. *Middle:* The galaxy auto-correlation $C_\ell^{g_ig_i}$ for MegaMapper, where g_i is the LBG sample in the $i^{\rm th}$ redshift bin: $z \in [2.0, 2.8], [2.8, 3.5], [3.5, 4.3]$, or $[4.3, 5.0]$. *Right:* The convergence-galaxy cross-spectra in each redshift bin.

Combined forecasts

In our forecasts we treat the CMB surveys described in §10.2 as priors on ΛCDM, M_ν, $N_{\rm eff}$ and Ω_k. Since our CMB priors are calculated using *unlensed* power spectra, we neglect any covariances between these priors and lensing or full shape data. Thus to combine a CMB prior with a LSS Fisher matrix $\boldsymbol{F}_{\rm LSS}$ with basis $\{\Lambda{\rm CDM}, M_\nu, N_{\rm eff}, \Omega_k, \cdots\}$, we simply add the CMB Fisher matrix to the appropriate submatrix of $\boldsymbol{F}_{\rm LSS}$.

When combining lensing[6] $(C_\ell^{\kappa\kappa}, C_\ell^{\kappa g_i}, C_\ell^{g_ig_i})$ and full shape clustering data $(P_{gg}(\boldsymbol{k}))$ from the same survey, we impose a $k_\parallel > 10^{-3}\,h\,{\rm Mpc}^{-1}$ cut on the full shape Fisher matrix. Since lensing probes modes with $k_\parallel \sim 0$, this cut nulls any covariance between the two observables in linear theory. We therefore neglect any covariance between the lensing and full shape data, and simply add the respective Fisher matrices when showing combined constraints.

We also neglect any (negligibly small) covariances between two non-overlapping LSS surveys when combining their full shape data. The same cannot be done when combining lensing data from two surveys, however, since the covariance between $C_\ell^{\kappa g_i}$ and $C_\ell^{\kappa g_j}$ is nonzero even when the tracers g_i and g_j are at different redshifts. To avoid this complication we never combine lensing data from two different surveys − when combining full shape information from two spectroscopic surveys with lensing data, we only include the lensing data from the higher redshift survey. Since the high-z survey (e.g. MegaMapper) typically dominates the constraining power, neglecting the lensing-galaxy cross-correlations of the low-z survey (e.g. DESI) has a negligible impact on our results. Since the foreground wedge (§10.3) restricts access to modes with $k_\parallel \sim 0$, we never include 21-cm × CMB lensing in our results. Thus

[6]From here on out we abbreviate $(C_\ell^{\kappa\kappa}, C_\ell^{\kappa g_i}, C_\ell^{g_ig_i})$ as "lensing data", "+ lensing" or simply "lensing", unless stated otherwise.

when adding lensing to (low-z spectroscopic survey) + (21-cm) clustering data, we only include the lensing data from the low-z survey.

The full set of free parameters in our base model is given in Table 10.2. In all forecasts that include full shape information we marginalize over $\{\Lambda\text{CDM}, b, b_2, b_s, \alpha_0, \alpha_2, \alpha_4, N_0, N_2, N_4\}$, while in all forecasts that include lensing we marginalize over $\{\Lambda\text{CDM}, b, b_2, b_s, \alpha_0, N_0, \alpha_x\}$. We marginalize over the 21-cm brightness temperature T_b in any forecast that includes 21-cm full shape data. We treat all time-dependent nuisance terms (biases, counterterms, stochastic terms, T_b) in different redshift bins as separate, independent variables which are individually marginalized over. We do the same for nuisance terms of separate surveys when showing combined constraints.

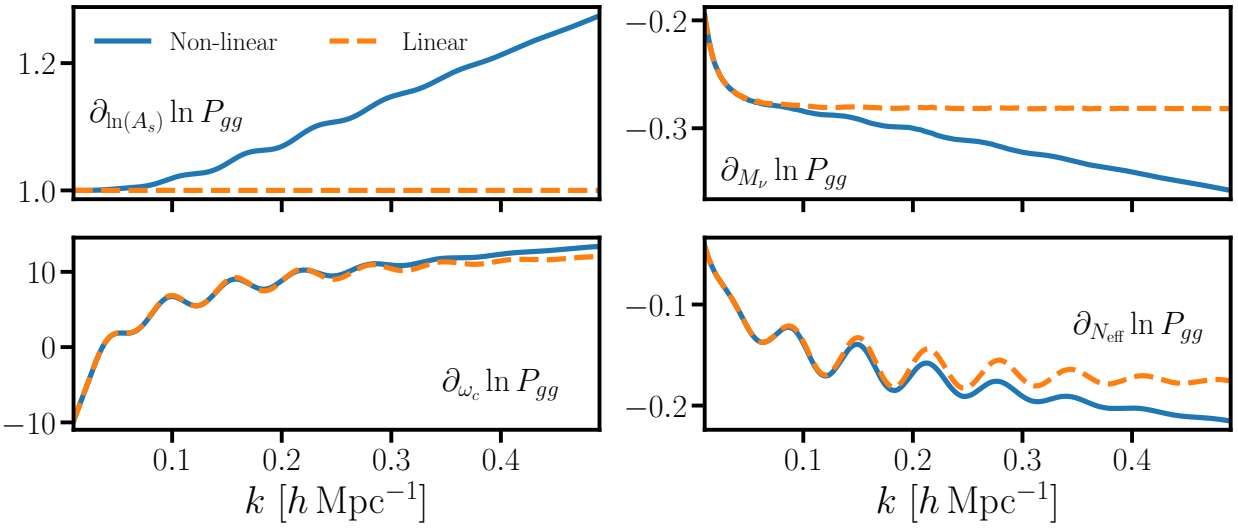


Figure 10.7: Some exemplary logarithmic derivatives at $z = 2.5$, $\mu = 0$. The solid curves are calculated using Lagrangian perturbation theory, the dashed curves are calculated using linear theory. These derivatives do not account for the A-P (§10.3) effect.

Derivatives

With the exception of derivatives with respect to the three stochastic terms (N_0, N_2, N_4) and the 21-cm brightness temperature, all of our derivatives are calculated numerically. For parameters whose fiducial value, $\theta_{\rm fid}$, is within the physically allowed region we use a five-point, two-sided derivative with relative step size $\delta\theta/\theta_{\rm fid}$, which takes the form:

$$\left.\frac{\partial P}{\partial \theta}\right|_{\theta_{\rm fid}} = \frac{-P_{+2} + 8P_{+1} - 8P_{-1} + P_{-2}}{12\,\delta\theta} + \mathcal{O}\left(\delta\theta^4\right). \tag{10.19}$$

where $P_{\pm n} \equiv P(\theta_{\rm fid} \pm n\delta\theta)$. When the parameter has its fiducial value at a boundary of the allowed region we instead use a five-point, one-sided derivative:

$$\left.\frac{\partial P}{\partial \theta}\right|_{\theta_{\rm fid}} = \frac{-3P_{+4} + 16P_{+3} - 36P_{+2} + 48P_{+1} - 25P_{+0}}{12\,\delta\theta} + \mathcal{O}\left(\delta\theta^4\right). \tag{10.20}$$

We take $\delta\theta/\theta_{\rm fid} = 0.01$ for most parameters. When the fiducial value of the parameter is zero, we assume $\delta\theta = 0.01$ by default. τ, M_ν and $A_{\rm lin}$ (see §10.4) require larger (or smaller) step sizes for numerical convergence: $\delta\tau = 0.3\tau$, $\delta M_\nu = 0.05 M_\nu$ and $\delta A_{\rm lin} = 0.002$ respectively. For these choices the convergence of the derivatives is better than 0.5% for all parameters. We show some illustrative examples in Fig. 10.7.

When taking derivatives with respect to M_ν, we fix $m_{\nu_1} = 0$ to avoid numerical error, and take steps of the form $m_{\nu_2}, m_{\nu_3} \to m_{\nu_2} + \delta M_\nu/2, m_{\nu_3} + \delta M_\nu/2$ to preserve the atmospheric mass splitting.

Alcock-Paczynski effect

Power spectra, being dimensionfull quantities expressed in terms of physical distances, are usually computed from observations by assuming a "fiducial model" to convert redshifts and angular positions to 3D comoving coordinates. This fixed fiducial cosmology must be accounted for when comparing the observations to a theoretical model that predicts a different conversion to comoving coordinates. The distortion this introduces is known as the Alcock-Paczynski (A-P) effect [30]. The wavevector measured using the fiducial cosmology $\boldsymbol{k}_{\rm fid}$ is related to the true wavevector $\boldsymbol{k}$ by:

$$\boldsymbol{k}_{\rm fid} = \boldsymbol{k}_\parallel \alpha_\parallel(z) + \boldsymbol{k}_\perp \alpha_\perp(z). \tag{10.21}$$

where $\alpha_\parallel(z) = H(z)_{\rm fid}/H(z)_{\rm true}$ and $\alpha_\perp(z) = D_A(z)_{\rm true}/D_A(z)_{\rm fid}$. The correlation function, being a dimensionless quantity, is left unchanged by this rescaling. Thus $P_{gg}^{\rm obs}(\boldsymbol{k}_{\rm fid}, z) d^3\boldsymbol{k}_{\rm fid} = P_{gg}^{\rm true}(\boldsymbol{k}, z) d^3\boldsymbol{k}$, which gives [548]:

$$P_{gg}^{\rm obs}(\boldsymbol{k}_{\rm fid}, z) = \alpha_\parallel^{-1}(z)\alpha_\perp^{-2}(z) P_{gg}^{\rm true}(\boldsymbol{k}, z). \tag{10.22}$$

This effect is included in our derivatives via the chain rule:

$$\begin{aligned}\frac{\partial P}{\partial \theta} \to \frac{\partial P}{\partial \theta} &- \underbrace{\left[\frac{\partial \alpha_\parallel}{\partial \theta} + 2\frac{\partial \alpha_\perp}{\partial \theta}\right]}_{\text{volume rescaling}} P \\ &- \underbrace{\mu(1-\mu^2)\left[\frac{\partial \alpha_\parallel}{\partial \theta} - \frac{\partial \alpha_\perp}{\partial \theta}\right]}_{-\partial_\theta \mu} \frac{dP}{d\mu} - \underbrace{k\left[\frac{\partial \alpha_\parallel}{\partial \theta}\mu^2 + \frac{\partial \alpha_\perp}{\partial \theta}(1-\mu^2)\right]}_{-\partial_\theta k} \frac{dP}{dk}.\end{aligned} \tag{10.23}$$

BAO modeling

One of the major scientific goals of future redshift surveys is the mapping of the expansion history of the Universe using the baryon acoustic oscillation (BAO) technique [617]. The BAO features in the power spectrum form a particularly robust standard ruler, and can benefit from a technique known as "reconstruction" [209, 450], that improves the distance fidelity. Since BAO measurements are analyzed differently than many of the other parameters we've discussed, we forecast constraints on the distance scale differently, as we now discuss.

We hold the shape of the linear theory power spectrum fixed at a fiducial value and compute the post-reconstruction power spectrum ($P_{\rm recon}$) with the "Rec-Sym" convention within the Zeldovich approximation [139]. We use the same quadratic bias model of §10.3, with fiducial bias values listed in Table 10.2, and approximate the reconstructed power spectrum by its first three non-zero multipoles. We include the A-P parameters, $\alpha_\parallel$ and $\alpha_\perp$, that are defined to include a factor of the sound horizon at the drag epoch (r_d), and marginalize over a linear bias and 15 "broad band" polynomials as is usually done experimentally (e.g. refs. [28, 205]):

$$P_{\rm obs}(k_\perp, k_\parallel; \boldsymbol{b}) = \alpha_\parallel^{-1}\alpha_\perp^{-2}\, P_{\rm recon}(k_\perp/\alpha_\perp, k_\parallel/\alpha_\parallel, \boldsymbol{b}) + \sum_{n=0}^{4}\sum_{m=0}^{2} c_{nm} k^n \mu^{2m} \tag{10.24}$$

with $\mu = k_\parallel/k$ and $\boldsymbol{b} = (b, b_2, b_s)$. We set the fiducial value of c_{00} coefficient to the Poisson shot noise $1/\bar{n}_{\rm gal}$, and the fiducial values of all higher order coefficients to zero.

The Fisher matrix for the reconstructed power spectrum takes the same functional form as Eq. (10.10). In our forecasts (§10.4) we interpret the marginalized errors on the tangential and radial distance measures as $\sigma_{\alpha_\perp} D_A/r_d$ and $\sigma_{\alpha_\parallel} H r_d$ respectively. In Appendix G.3 we compare this approach to the traditional method of ref. [547].

Redshift uncertainties

An error $\sigma_z/(1+z)$ in the measurement of the redshift of an object corresponds to an error in the radial distance to said object $\sigma_\chi = \sigma_z d\chi/dz = [c(1+z)/H(z)]\sigma_z/(1+z)$. As a result of this error, the observed density field at Euleurian position $\boldsymbol{x}$ is related to the true density field by a simple translation $\delta^{\rm obs}(\boldsymbol{x}) = \delta(\boldsymbol{x} - \epsilon(\boldsymbol{x})\hat{\boldsymbol{n}})$, where $\hat{\boldsymbol{n}}$ is the direction along the line of sight, and $\epsilon(\boldsymbol{x})$ is a stochastic variable which we assume is drawn from a Gaussian distribution with mean zero and standard deviation σ_χ. Thus on average, the observed density field is damped in the $\hat{\boldsymbol{n}}$ direction:

$$\begin{aligned}\langle \delta^{\rm obs}(\boldsymbol{k})\rangle_\epsilon &= \int d^3\boldsymbol{y} e^{-i\boldsymbol{k}\cdot\boldsymbol{y}}\delta(\boldsymbol{y}) \frac{1}{\sqrt{2\pi}\sigma_\chi}\int d[\epsilon(\boldsymbol{x})] e^{-\epsilon(\boldsymbol{x})^2/2\sigma_\chi^2} e^{-ik_\parallel \epsilon(\boldsymbol{x})} \\ &= e^{-k_\parallel^2\sigma_\chi^2/2}\delta(\boldsymbol{k}),\end{aligned} \tag{10.25}$$

which in turn damps the observed power spectrum by a factor of $\exp[-k^2\mu^2\sigma_\chi^2]$. In our forecasts we assume that the redshift uncertainty can be well characterized and ignore this

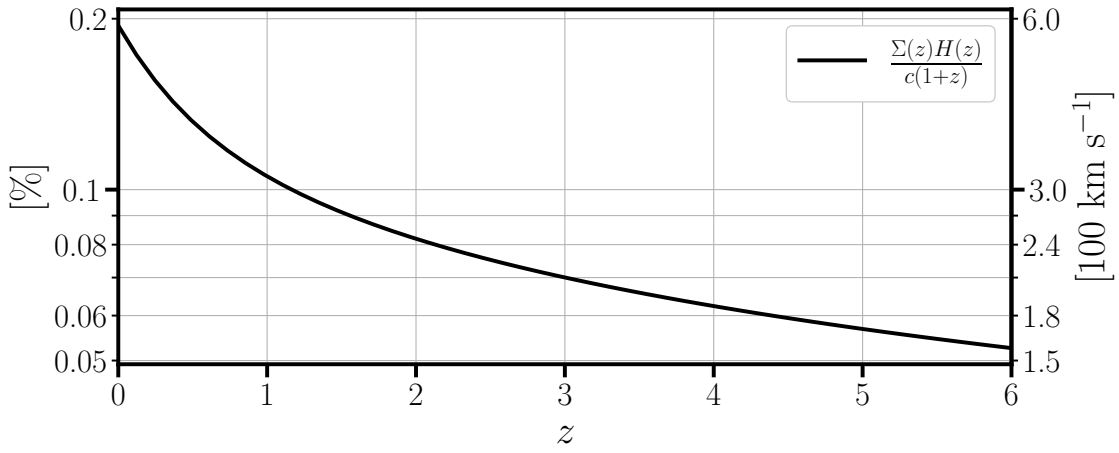


Figure 10.8: $\Sigma(z)H(z)/[c(1+z)]$ as a function of redshift. If the redshift uncertainties $\sigma_z/(1+z)$ are near or above this curve, then they will suppress modes that could carry valuable cosmological information. The left y-axis shows the requirement on $\sigma_z/(1+z)$ in percent, while the right y-axis converts that to a velocity scale, in $100\,\mathrm{km\,s^{-1}}$.

extra factor in the derivatives. Therefore, at the Fisher matrix level (Eqn. (10.10)), we can in practice absorb the effect of redshift uncertainties in the noise, so that $N_0 \to N_0 \exp[k^2\mu^2\sigma_\chi^2]$.

These uncertainties become increasingly important as an experiment pushes to higher redshifts. If we want to accurately reconstruct all modes with $k < k_{\rm nl}$, we require that the uncertainty in radial length measurements is smaller than the corresponding wavelength at the non-linear scale: $\sigma_\chi < k_{\rm nl}^{-1} = \Sigma(z)$. This in turn implies $\sigma_z/(1+z) < \Sigma(z)H(z)/[c(1+z)]$. The RHS of this inequality is plotted in Fig. 10.8. Obtaining small redshift uncertainties at high redshift is challenging, because the tracers are intrinsically fainter, but also because some of the broader and easier lines to measure provide intrinsically less accurate redshifts (due to asymmetric profiles, or because they are affected by winds and outflows for example) [549]. While it may be possible to use the Lyα line profile to recover the systemic velocity to $\lesssim 100\,\mathrm{km/s}$ [599] — which would be more than sufficient for our science cases — how well this works on large samples with achievable integration times remains an open question. We delay discussing the impact of redshift uncertainties on our forecasts to §10.5.

10.4 Results

In sections 10.2 and 10.3 we outlined the surveys, samples and methods used to create our forecasts. In this section we summarize our results, estimating the constraints on base ΛCDM, curvature, neutrino mass, relativistic species, primordial features, primordial non-Gaussianity, dynamical dark energy, and gravitational slip.

Because we are interested in the performances of proposed future experiments in the

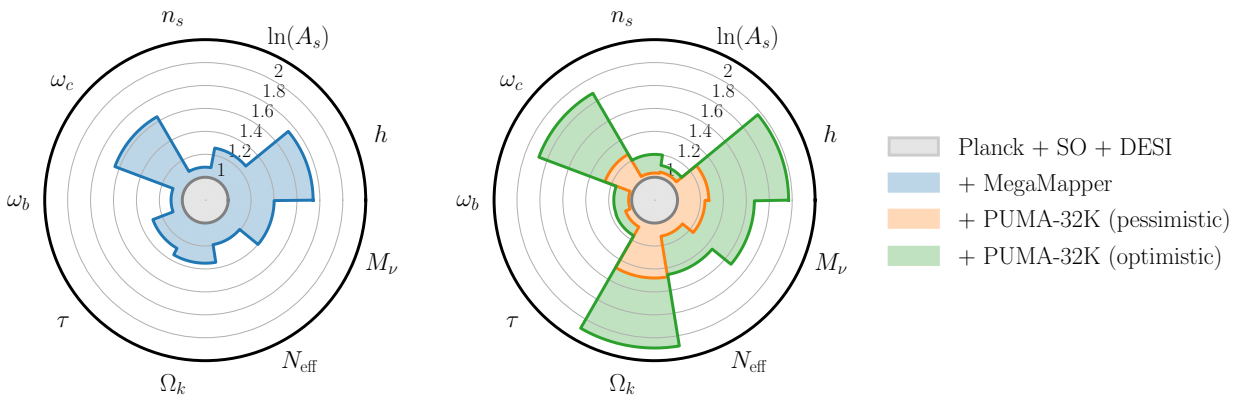


Figure 10.9: For each parameter X on the perimeter, the corresponding radius is the ratio of the uncertainty σ_X from Planck+SO+DESI (PSD) alone to the uncertainty from Planck+SO+DESI+high redshift ($z > 2$) data. The constraints on ΛCDM are calculated assuming our base model (see §10.3) while each extension to our base model (M_ν, $N_{\rm eff}$, Ω_k) is treated as a single parameter extension. These constraints include both full shape and lensing data (see §10.3 for our conventions when combining surveys). For PSD+MegaMapper we include MegaMapper×SO lensing, while for all other configurations DESI×SO lensing is included. In green (orange) we show constraints for optimistic (pessimistic) foreground assumptions (§10.3).

context of current state of the art surveys, we compute the Fisher information for DESI and Euclid as well. This way all current and future surveys are studied using the same perturbation theory scheme and share the same model assumptions, allowing for a fairer comparison. In light of existing tensions between LSS datasets and Planck (e.g. the H_0 and S_8 tensions), we also present forecasted uncertainties on relevant parameters from individual surveys without any CMB external information. In doing so we obtain an estimate for the expected level of agreement (or disagreement) between disjoint LSS and CMB analyses.

We note that by taking ΛCDM as our fiducial cosmology, we are implicitly assuming that these tensions are either statistical flukes or due to systematic errors which will ultimately be resolved in the future. More generally, we implicitly assume that our base model (or our base model + the extensions considered in the following subsections) is capable of accurately fitting future LSS and CMB data to the required precision. Should future tensions arise, or should current tensions worsen, it may be necessary to introduce extensions to this base model to accurately fit the data. Marginalizing over the free parameters associated with these extensions would inevitably result in larger errors than those derived assuming our base model.

Such extensions could arise from new physics (e.g. early dark energy), new nuisance parameters which model systematic effects (e.g. relative velocities), or from freeing up parameters which had previously been held fixed. The latter is explored in detail in §10.4.

We note that we do not explore extending the base model to include the running of the spectral index α_s. We expect a detection of the running to be well beyond the reach of the surveys considered in this chapter[7], and it is therefore unlikely that one would need to extend the base model to fit for the running.

As partially summarized in Fig. 10.9, high redshift surveys would yield significant gains in constraining power over current LSS and CMB surveys. For example, the addition of PUMA to the combination of Planck, SO and DESI tightens the constraints on ω_c, h and Ω_k by a factor of ~ 2, while the constraints on M_ν improve by more than 60%. The addition of MegaMapper also yields to a large reduction of the uncertainties. Other extensions of the base model, such as primordial features (§10.4) and EDE (§10.4), show similar improvements, which we explore in detail in the following subsections.

Distance measures

The first results we show concern the uncertainty on the distance measures obtained from measurements of the BAO feature (§10.3). Fig. 10.10 shows the relative error on $\alpha_\perp$ and $\alpha_\parallel$, which can be interpreted as relative errors on D_A/r_d and $r_d H$, for different spectroscopic and 21-cm surveys. Near-term surveys (DESI, Euclid and HIRAX) could measure the angular diameter distance to sub-percent precision all the way to $z \sim 2$, and the Hubble parameter to percent precision in the same redshift range. For 21-cm instruments we always show a shaded region encompassing optimistic and pessimistic foreground assumptions. We find that an MSE-like survey, shown in green, nicely extends the current spectroscopic program to higher redshifts, achieving the same accuracy in the redshift range $2 \leq z \leq 4$ as DESI and Euclid at low redshift. An instrument like MegaMapper would further improve these constraints, yielding 0.5% measurements of both D_A/r_d and $r_d H$ at $z \lesssim 3.5$ and 1% measurements at $z \sim 5$. A Stage-II 21-cm experiment like PUMA-32K, shown in brown, could in principle outperform spectroscopic surveys, providing almost cosmic variance limited measurements of the Hubble parameter out to $z \sim 6$. For the measurement of distances perpendicular to the line of sight, 21-cm instruments are severely limited by the foreground contamination, but even in our pessimistic scenario PUMA-32K would reduce the uncertainties on D_A/r_d by roughly 30% compared to MegaMapper.

We would like to stress that for 21-cm instruments it is not yet clear how to perform BAO reconstruction and therefore enhance the constraints on the distance parameters. In auto-correlation experiments new approaches have been considered [444], but no practical implementation has been proposed for interferometers. On the other hand we notice that the main purpose of BAO reconstruction is to undo large scale flows and sharpen the position of the BAO feature in configuration space. The effect of such large scale motions is smaller at high redshifts, and becomes irrelevant at $z \gtrsim 3$, after which we expect BAO reconstruction to yield diminishing returns.

[7] We find $\sigma(\alpha_s) = 0.005$ from our fiducial LiteBIRD + CMB-S4 + DESI + MegaMapper + lensing survey, while the expected value for the running is ~ 0.001.

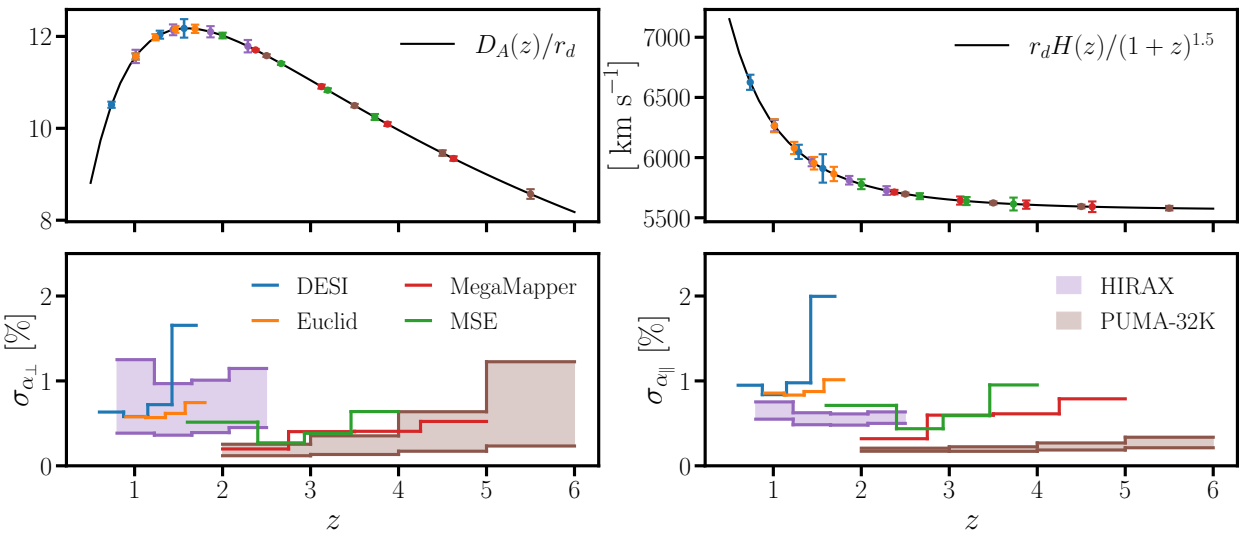


Figure 10.10: Error on the A-P parameters from the reconstructed power spectrum, which can be interpreted as relative errors on $D_A(z)/r_d$ and $r_dH(z)$ respectively. The boundaries of the shaded regions denote optimistic/pessimistic foreground assumptions for the 21-cm surveys. In the top panels we show the errorbars for the optimistic case.

Finally, it is possible to combine BAO information with the full shape analysis presented in the next few sections [468]. This would however require a model for the covariance between the two measurements, which is currently best obtained using mock catalogs. For this reason we do not include any prior on the distance measures in the forecasts discussed in the next sections.

Base ΛCDM

Historically, the standard (ΛCDM) model has been extremely successful in modeling cosmological observables, ranging from the Universe's expansion history to the statistics of CMB anisotropies and the evolution of galaxy clustering (e.g. Fig. 10.3). Future LSS and CMB surveys will significantly improve over current constraints, providing the most stringent test of ΛCDM to date.

Within the next $\sim$ 5 years, our forecasts indicate that the combination of DESI and Planck will improve upon current CMB measurements of the Hubble constant and ω_c by more than a factor of 2. By themselves, DESI and Euclid will provide similar constraints to Planck on several ΛCDM parameters, but LSS and CMB data combined are larger than the sum of their parts. In Table 10.3 we also show how the inclusion of CMB lensing (and its cross-correlation with galaxies) improves the constraints. For Planck lensing and current or upcoming surveys like DESI and Euclid, the addition of lensing does not yield significant improvements. Given the time frame of future LSS surveys, we expect to (in principle) be

Configuration	$10^4 h$	$10^4 \omega_c$	$10^3 \tau$	$10^3 \Omega_k$	$10^3 N_{\rm eff}$	M_ν [meV]
Planck	62.14	13.66	7.87	17.19	187.09	83.39
DESI	63.83	42.29	···	20.38	599.83	311.98
Euclid	58.02	35.55	···	17.0	476.77	245.8
MSE	49.84	21.36	···	10.79	214.09	129.18
MegaMapper	45.75	19.42	···	8.01	173.57	104.58
HIRAX	48.9/98.2	23.2/42.5	···	10.7/29.4	267.5/707.0	157.0/516.9
PUMA-5K	39.4/64.9	14.6/16.8	···	5.5/12.1	101.9/204.1	67.4/154.2
PUMA-32K	32.8/49.7	12.2/14.3	···	3.8/8.8	81.4/150.7	52.4/116.1
Planck + DESI	31.83	7.28	5.86	1.24	112.8	41.49
+ lensing	28.72	6.58	5.17	1.19	107.32	38.01
Planck+SO+DESI (PSD)	23.42	6.13	5.83	0.68	43.93	39.76
+ lensing	18.68	4.9	4.2	0.68	42.84	35.39
Planck + SO + Euclid	23.04	6.05	5.64	0.66	43.64	37.17
+ lensing	17.23	4.54	3.97	0.65	42.62	33.67
Planck + SO + MSE	17.63	4.76	5.5	0.58	41.51	31.57
+ lensing	12.47	3.41	3.46	0.57	39.27	29.2
PSD + MegaMapper	13.77	3.77	4.67	0.51	38.1	26.33
+ lensing	10.65	2.96	3.23	0.5	35.91	25.29
PSD + HIRAX	16.9/22.4	4.4/5.9	5.8/5.8	0.6/0.6	41.1/43.6	34.1/38.9
+ lensing	14.9/18.2	3.9/4.8	3.9/4.2	0.6/0.6	40.3/42.6	30.0/34.5
PSD + PUMA-5K	12.9/19.7	3.5/5.2	5.8/5.8	0.4/0.5	34.2/41.5	27.1/35.3
+ lensing	11.9/16.7	3.2/4.4	3.7/4.0	0.4/0.5	33.8/40.6	25.2/31.6
PSD + PUMA-32K	10.0/16.6	2.8/4.4	5.7/5.8	0.3/0.5	29.6/39.0	22.4/31.6
+ lensing	9.6/14.7	2.6/3.9	3.6/3.9	0.3/0.5	29.4/38.3	21.2/28.5
LiteBIRD+S4+DESI (LSD)	21.39	5.74	2.02	0.34	24.57	34.78
+ lensing	12.86	3.41	1.93	0.34	24.02	24.38
LSD + MegaMapper	12.76	3.52	1.94	0.31	22.63	20.02
+ lensing	8.86	2.43	1.74	0.31	21.62	14.92
LSD + HIRAX	15.9/20.6	4.3/5.5	2.0/2.0	0.3/0.3	23.7/24.5	30.2/34.1
+ lensing	11.3/12.7	3.0/3.4	1.9/1.9	0.3/0.3	23.3/24.0	19.5/23.6
LSD + PUMA-5K	12.4/18.4	3.4/4.9	2.0/2.0	0.3/0.3	21.3/24.1	24.7/31.1
+ lensing	9.7/12.1	2.6/3.2	1.8/1.9	0.3/0.3	21.0/23.6	17.4/21.4

LSD + PUMA-32K	9.7/15.8	2.7/4.2	2.0/2.0	0.3/0.3	19.4/23.4	20.3/28.0
+ lensing	8.3/11.3	2.2/3.0	1.8/1.9	0.3/0.3	19.3/23.0	14.5/18.9
Everything bagel	6.95	1.94	1.69	0.24	18.36	12.24

Table 10.3: 1σ uncertainties on selected ΛCDM parameters and extensions from a combination of primary CMB, large-scale clustering ($P_{gg}(\boldsymbol{k})$), and CMB×LSS lensing data. For all configurations that include 21-cm data, we give constraints for optimistic/pessimistic foreground assumptions. The Planck constraints in the first row are the diagonal components of the measured TT,TE,EE+lowE covariance [475]. These would be the true constraints if the Planck likelihood were a perfect Gaussian, and agree with the constraints quoted in ref. [475] to within 5% (with the exception of M_ν, where non-Gaussianities in the likelihood increase the true σ_{M_ν} to 130 meV). Everywhere else in this chapter, Planck refers to the CMB prior described in §10.2. In rows $2-8$ (LSS alone) we include a conservative $\sigma(\omega_b) = 0.0005$ prior from BBN [475]. PSD is short for Planck + SO + DESI, while LSD is short for LiteBIRD + S4 + DESI. "Everything bagel" is short for LSD + PUMA-32K (with optimistic foreground assumptions) + MegaMapper + MegaMapper×S4 lensing, and neglects any correlations between 21-cm and LBG clustering.

able to combine them with Planck+SO+DESI (PSD) data. The addition of MegaMapper to PSD improves the constraint on the Hubble constant by almost a factor of 2, and by $\sim 40\%$ on ω_c. We find that cross-correlations with CMB lensing maps return larger gains for a MegaMapper-like survey, presumably due to the increased overlap with the CMB lensing kernel. For 21-cm surveys, foregrounds will likely make the cross correlation with CMB lensing maps not possible. Thus, in Table 10.3, the "+lensing" rows below any PSD+21 cm surveys only include the cross-correlation of DESI with SO (or S4) lensing maps. Nonetheless PSD+PUMA-5K or -32K could yield comparable constraints to PSD+MegaMapper on most ΛCDM parameters. In the future, density reconstruction methods could help 21-cm surveys to recover the modes lost in the foreground wedge, allowing for a non-zero correlation with CMB lensing maps [415, 414].

On a similar timescale to future LSS surveys, CMB experiments like LiteBIRD and CMB-S4 will also be taking data. It is therefore interesting to consider their combination with DESI as the baseline, which we dub LSD. As shown in Table 10.3, clustering data at high redshifts enable significant improvements in ΛCDM constraints over a nearly cosmic variance limited CMB survey. The uncertainties on h and ω_c from LSD only are $> 50\%$ larger than those from LSD+MegaMapper. Improvements from cross-correlations with CMB lensing are also of the same order. We note that constraints from LSD+PUMA-32K are comparable to those from LSD+MegaMapper, despite 21-cm surveys having a lower effective shot noise. This is ultimately due to degeneracies between the brightness temperature T_b, A_s and the linear bias b, which in turn worsen the constraints on any cosmological parameters that the latter are degenerate with.

Structure growth

Future high-redshift surveys will enable 1% measurements of the amplitude of structure growth $\sigma_8(z)$ out to $z \sim 5$. We consider two different methodologies for structure growth measurements: one in which σ_8 is treated as a derived parameter from ΛCDM, and another, more cosmology-agnostic approach in which the shape of the power spectrum is held fixed.

In the derived-from-ΛCDM procedure we calculate the Fisher matrix from full shape clustering data with basis $\{\theta_i\} = \{\Lambda\text{CDM}, b, b_2, b_s, N_0, N_2, N_4, \alpha_0, \alpha_2, \alpha_4\}$ in each redshift bin. Our choice of ΛCDM parameters is listed in Table 10.2 (τ is held fixed here, since no primary CMB information is included). We then transform to a basis that includes $\sigma_8(z)$, where z is the central redshift of some redshift bin. The Fisher matrix transforms as a two-form under a change of basis: $\boldsymbol{F} \to \boldsymbol{R}^T \boldsymbol{F} \boldsymbol{R}$, where

$$\boldsymbol{R}^{-1} = \begin{pmatrix} \frac{\partial \sigma_8(z)}{\partial \theta_1} & \frac{\partial \sigma_8(z)}{\partial \theta_2} & \cdots & \frac{\partial \sigma_8(z)}{\partial \theta_n} \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix}. \tag{10.26}$$

We then add a conservative $\sigma(\omega_b) = 0.0005$ [475] prior from BBN, and invert this matrix to yield the marginalized constraint on $\sigma_8(z)$. Thus the measurements of $\sigma_8(z)$ are local, derived only from data within the redshift bin of interest, apart from the shared ω_b prior. As shown in the left panel of Fig. 10.11, DESI and Euclid will measure $\sigma_8(z)$ to $3-7\%$ across the range $0.6 < z < 2$, while MSE and MegaMapper will extend these measurements out to $z = 5$ with less than 5% error. Due to a degeneracy between σ_8 and T_b, the constraints from HIRAX are more modest than those from DESI or Euclid. Since σ_8 and T_b are perfectly degenerate in linear theory, 21-cm measurements of σ_8 heavily depend on how accurately one can model and measure non-linear effects [124]. While PUMA also suffers from the same $\sigma_8 - T_b$ degeneracy, it has a much lower shot noise than HIRAX, and is thus more sensitive to non-linear scales that differentiate the effects of the two parameters, enabling $5-10\%$ measurements across its redshift range. We find that PUMA-32K (with a BBN prior on ω_b) can measure $T_b(z)$ to the same relative accuracy as $\sigma_8(z)$, or $5-10\%$ depending on the redshift bin. Thus a $\approx$ 1% prior on $T_b(z)$ would be required for 21-cm measurements of $\sigma_8(z)$ to be competitive with MegaMapper, which is likely beyond the reach of astrophysical determinations from high-column-density systems.

In Fig. 10.11 we also show the improvement from adding SO×LSS lensing to full shape data. When adding lensing to this figure we only include the cross-correlation, and not the convergence power spectrum. This ensures that the constraints on $\sigma_8(z)$ remain local apart from the ω_b prior. Thus these constraints make no assumptions about the expansion history outside of the redshift bin of interest. As was true in the previous section, lensing cross-correlations offer smaller improvements for DESI and Euclid, while for MegaMapper, the improvement can be more than a factor of 2. With lensing cross-correlations, MegaMapper can measure $\sigma_8(z)$ with a 1% error across its entire redshift range.

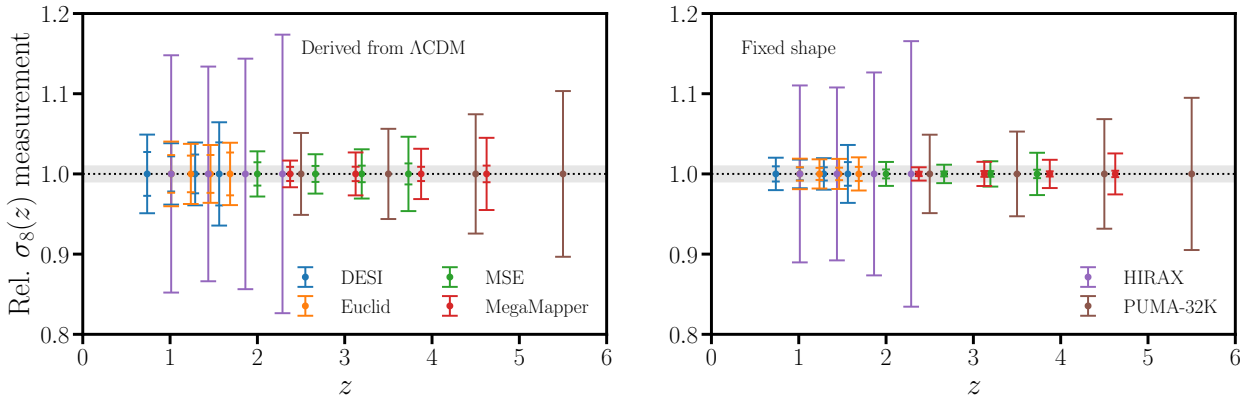


Figure 10.11: *Left:* Errors on $\sigma_8(z)$ when derived from ΛCDM, assuming a $\sigma(\omega_b) = 0.0005$ prior from BBN. *Right:* Errors on $\sigma_8(z)$ from full shape data when calculated using the fixed shape procedure, without any external priors. The errorbars with smaller caps include cross-correlations with SO lensing, but do not include any information from the convergence auto-spectrum. All 21-cm measurements assume optimistic foregrounds. The gray bands extend from 1 ± 0.01.

While the derived-from-ΛCDM constraints are only valid within the ΛCDM framework, fixed shape constraints assume very little about the underlying cosmology. For any cosmology whose free parameters include an amplitude of the matter power spectrum, one can obtain a fixed shape measurement by fitting to the amplitude with all other parameters fixed at some fiducial value. Thus fixed-shape forecasts give a sense of a "best-case" scenario, since the error on the amplitude from the fixed shape procedure is always smaller than the error derived from the full cosmological model. In our case we fix $\{h, \omega_c, \omega_b, n_s, \tau\}$ at their fiducial ΛCDM values, and treat the amplitude of the (linear) matter power spectrum as an independent variable in each redshift bin. In each bin we marginalize over the nuisance terms listed in Table 10.2. With this procedure, we find that DESI and Euclid yield 2% measurements across $0.6 < z < 2$, while MSE and MegaMapper extend these measurements out to $z = 5$ with comparable uncertainty. When cross-correlations with SO lensing are included, these errors decrease to sub-percent values across the entire redshift range covered by spectroscopic surveys ($0.6 < z < 5$). Again, we see that 21-cm measurements suffer from a $\sigma_8 - T_b$ degeneracy, resulting in much more modest $\mathcal{O}(10\%)$ constraints.

We find that fixed shape measurements from spectroscopic surveys have 2 to 5 times smaller error than when derived from ΛCDM (without lensing cross-correlations). For 21-cm surveys, however, the two methods yield similar constraints, presumably due to the $\sigma_8 - T_b$ degeneracy dominating over any other degeneracy between σ_8 and the remaining ΛCDM parameters, such as τ.

Massive neutrinos

One of the best-motivated extensions beyond the ΛCDM model described above is allowing the neutrino masses to be non-minimal [646]. Measurements of solar and atmospheric neutrino oscillations probe $m_{\nu_2}^2 - m_{\nu_1}^2$ and $m_{\nu_2}^2 - m_{\nu_3}^2$ respectively, which together imply a minimum value for the sum of the neutrino masses $M_\nu = \sum_i m_{\nu_i} > 60\,\mathrm{meV}$. Oscillation measurements have yet to determine if neutrinos follow a normal ($m_{\nu_2} < m_{\nu_3}$) or inverted ($m_{\nu_2} > m_{\nu_3}$) mass hierarchy, which is one of the major goals of upcoming oscillation measurements and neutrino-less double beta decay experiments. Cosmological constraints of M_ν are complementary to these experiments. CMB and BAO data currently provide our tightest upper limits on M_ν [153], and future surveys should tighten these constraints significantly.

Neutrinos are relativistic at early times, washing out their primordial fluctuations at length scales smaller than the horizon. However, once the neutrino fluid cools and becomes non-relativistic (occurring at a scale factor $a_{\rm nr}$), so that its energy density is dominated by the neutrino mass M_ν, neutrinos cluster in the same way as the CDM+baryon fluid. The comoving length scale associated with this transition is $k_{\rm nr} = a_{\rm nr}H(a_{\rm nr}) \sim 10^{-3}h\mathrm{Mpc}^{-1}$ [1]. Thus, another way of saying the above is that neutrinos only cluster at large scales ($k \ll k_{\rm nr}$). At small scales, the neutrino energy density is nearly spatially homogeneous, contributing to the Hubble drag associated with the evolution of the CDM+baryon fluid, but not to the gravitational potential, which drives the clustering. Thus a nonzero neutrino mass slows the growth of structure, damping the matter power spectrum at scales smaller than the horizon at $a_{\rm nr}$.

Since this effect is largest at small scales and at late times, high redshift LSS surveys − which will measure high k modes in the late-time Universe − are ideally situated for M_ν measurements. As shown in Table 10.3, we see that $\sigma(M_\nu)$ drops dramatically with the addition of high redshift data: the constraint from PSD is 35 (40) meV with(out) lensing, which drops to 25 (26) meV for PSD+MegaMapper, or 21 (22) meV for PSD+PUMA-32K. As shown in Fig. 10.12, these results depend significantly on the $k_{\rm max}$ used in the forecast. The constraints change by nearly a factor of two for a LSD+MegaMapper-like survey between $0.2 < k_{\rm max}/k_{\rm nl} < 1.0$. Likewise, the 21-cm foreground assumptions have an equally significant effect on the forecasted neutrino constraints, impacting the results by up to 50% for the highest values of $k_{\rm max}$ for PSD+PUMA-32K[8]. However, even with a more conservative cutoff of $k_{\rm max}/k_{\rm nl} = 0.5$, and with mildly conservative foreground assumptions, a next generation CMB survey coupled to either MegaMapper or PUMA-32K could comfortably achieve $\sigma(M_\nu) \lesssim 20$ meV.

Currently in the literature there is a significant bifurcation in M_ν forecasts for the upcoming Euclid satellite. Ref. [104] found 34 meV for Planck+SO+Euclid+Euclid×SO lensing using a PT-based method in which nine nuisance terms were marginalized over, including four bias parameters (b, b_2, b_s, b_3), three counterterms (α_0, α_2, α_4) and two stochastic contri-

[8]For parameters with a physical priors, *e.g.* $M_\nu > 0$, the Fisher matrix approach does not necessarily return trustworthy error bars if the fiducial value chosen is close to the physical boundary. For this reason we caution to reader about the interpretation of our results when $\sigma(M_\nu)$ is large.

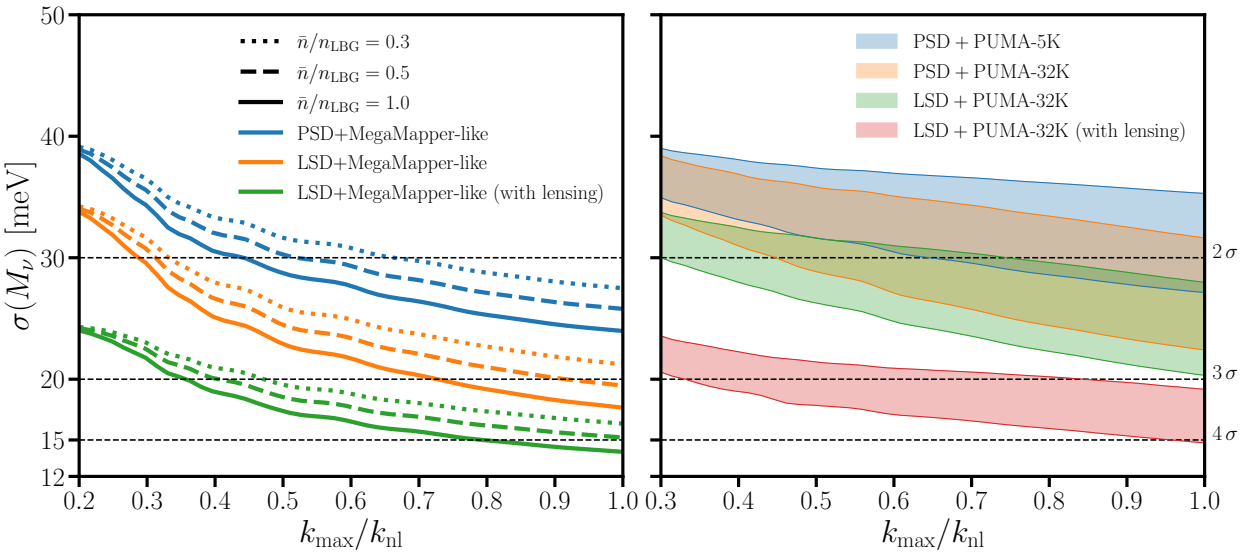


Figure 10.12: Constraints on M_ν as a function of $k_{\rm max}/k_{\rm nl}$ from the primary CMB and full shape data. The curves that are labeled "(with lensing)" include information from DESI×S4 lensing ($\kappa\kappa$, κg, gg). In both plots we set $k_{\rm max} = k_{\rm nl}$ for DESI, and only vary $k_{\rm max}/k_{\rm nl}$ for the high redshift survey of interest. *Left:* Constraints for a MegaMapper-like survey that observes a redshift-independent fraction $\bar{n}/n_{\rm LBG}$ of LBGs from $2 < z < 5$ with $f_{\rm sky} = 0.34$. *Right*: Constraints for PUMA. The edges of the shaded regions represent optimistic and pessimistic foreground assumptions.

butions (N_0, N_2). When we fix N_4 and match their $k_{\rm max} = 0.2\,h\,{\rm Mpc}^{-1}$ cut we find 40 meV for Planck+SO+Euclid+Euclid×SO lensing. A $\sim 15\%$ discrepancy is entirely reasonable given that their Euclid sample spans a broader redshift range ($0.6 < z < 2.1$) than ours ($0.9 < z < 1.8$), uses slightly different number densities and biases, and also marginalizes over a b_3 term. By contrast, ref. [147] finds 17 meV for Planck+Euclid+lensing (PS only) using a PT-based approach in which seven nuisance parameters were marginalized over, including three bias terms, three counterterms, and the shot noise. When we fix N_2 and N_4 and set $k_{\rm max} = 0.5\,h\,{\rm Mpc}^{-1}$, so that our set of nuisance terms and $k_{\rm max}$ match ref. [147], we find 30 meV for Planck+Euclid+lensing+BAO. If we fix τ our constraint drops to 20 meV, in good agreement with ref. [147]. The residual difference could be due to the fact that the authors in [147] employ a full MCMC approach as opposed to the Fisher matrix presented in this chapter.

Similar results for the cross-correlation of galaxies with CMB lensing were also recently obtained by [142], who consider a non-linear bias expansion similar to ours. Our results are largely consistent with recent forecasts for upcoming high redshift surveys. Ref. [228] found 26 meV for PSD+MegaMapper[9], while we find 25 meV. Recent PUMA forecasts [42] show

[9]This includes BAO only and no lensing, uses a different $k_{\rm max}$, and does not marginalize over non-linear

M_ν improving from 38 meV with Planck+LSST+DESI to 20 meV with CMB-S4+PUMA-32K, while we get 40 meV for Planck+SO+DESI and 20 meV for LSD+PUMA-32K.

We note that the inclusion of a future cosmic-variance limited measurement of τ from large scale CMB polarization such as from the proposed LiteBIRD satellite provides significant improvements on M_ν by removing the $\tau - A_s$ degeneracy present in the primary CMB. The very large-scale polarization signal is challenging to measure from the ground, potentially requiring a future satellite mission[10]. However, small scale information in both the CMB temperature (through fluctuations in the kinematic Sunyaev-Zel'dovich effect [230, 36]), CMB lensing [402], or the 21cm signal from reionization [355, 92] can potentially constrain τ to a similar level as LiteBIRD ($\sigma(\tau) \approx 0.002$), if astrophysical systematics can be kept under control. Including such a prior on τ reduces the PSD+MegaMapper neutrino mass constraint from 25 meV to 16 meV (including lensing) while for PSD+PUMA-32K it drops from 21 meV to 16 meV. As shown in Table 10.3 and Fig. 10.20, these improvements are similar to those obtained by upgrading from PSD to LSD, suggesting that the improvement from these future CMB data is almost entirely dominated by the better τ measurement.

Finally we note that with a LiteBIRD-like prior on τ and $k_{\rm max}/k_{\rm nl} = 1$, MegaMapper (which corresponds to $\bar{n}/n_{\rm LBG} \approx 0.5$) is primarily limited by its shot noise, and that small improvements are possible with a higher LBG target efficiency. Increasing the number density to $\bar{n}/n_{\rm LBG} = 1$ decreases the LSD + MegaMapper (without lensing) constraint from 20 to 18 meV.

Light relics

Measurements of the total amount of radiation in the Universe are a powerful way to search for new physics. Any light particle that was once in thermal equilibrium with the Standard Model, and hence had a cosmological abundance, leads to well known modifications in the power spectrum of the CMB and LSS. These effects are commonly parameterized through $N_{\rm eff}$, defined such that

$$\rho_r = \frac{\pi^2}{15}\left[1 + \frac{7}{8}\left(\frac{4}{11}\right)^{4/3} N_{\rm eff}\right] T_\gamma^4 \tag{10.27}$$

during the radiation dominated era, with $N_{\rm eff} = 3.046$ [381] in the Standard Model (or $N_{\rm eff} = 3$ in the limit of instantaneous neutrino decoupling). The presence of new light relics would shift this value to $N_{\rm eff} = 3.046 + \Delta N_{\rm eff}$, where $\Delta N_{\rm eff}$ depends on the number of independent spin states and decoupling temperature(s). By considering decoupling temperatures that are significantly larger than the top quark mass, one obtains lower limits to $\Delta N_{\rm eff}$ for broad classes of particles: $\Delta N_{\rm eff} > 0.027$ for scalars, 0.047 for Weyl fermions, and 0.054 for vector bosons [1]. In many examples of physics beyond the Standard Model, these new states are

terms, so the comparison is not entirely fair.

[10]The CLASS experiment [217] aims to measure these very large-scale polarization fluctuations from the ground, and depending on performance, it may lead to competitive constraints on τ compared to a future space mission.

too weakly coupled to be probed by facilities on Earth, therefore highlighting the synergy between cosmological probes and traditional particle physics experiments.

As shown in Fig. 10.7, the presence of extra radiation produces a change in the position of the BAO peaks and a damping of the matter power spectrum on small scales [62]. The density perturbations of the new degrees of freedom also produce a phase shift in the oscillation scale of the baryon-photon fluid before recombination [62]. In the CMB power spectrum, the small-scale suppression due to extra relativistic species is largely determined by a shift in the damping scale k_d, which is sourced by a change in the early time Hubble parameter $H(z)$ through the change in ρ_r. As such, $N_{\rm eff}$ is rather degenerate with the primordial helium abundance Y_p in CMB data, and marginalizing over Y_p can have a significant effect on $N_{\rm eff}$ constraints, weakening them by up to a factor of three[11] [69]. Note that this degradation is present despite the $N_{\rm eff}$-induced phase shift in the CMB peaks location which is independent of Y_p and therefore partly helps break this degeneracy, suggesting that the phase-shift alone is not playing a major role.

By contrast, the growth of dark matter density fluctuations on scales probed by LSS experiments is sensitive to the modified expansion history introduced by new free streaming particles, and it is not affected by Silk damping. It is therefore largely insensitive to Y_p, which is thus not included in our analysis. The BAO phase shift is also not the major source of information on $N_{\rm eff}$ in galaxy clustering data [70], whose constraining power is mostly driven by the change in the full shape of the power spectrum.

In our forecasts we marginalize over ΛCDM and the nuisance terms listed in Table 10.2. Our results are summarized in Table 10.3 and Fig. 10.13. Near term LSS surveys (DESI, Euclid, MSE) will improve over the SO goal of $\sigma(N_{\rm eff}) = 0.05$ by only $\sim 15\%$. The addition of high redshift data reduces this uncertainty to 0.03, a 40% improvement over Stage-III CMB measurements alone. With a Stage-IV CMB survey these constraints drop to 0.019, or 0.018 when a combination of high redshift LBG and 21-cm data are included. A measurement with uncertainty $\sigma(N_{\rm eff}) = 0.018$ could detect light scalars (vectors) at 1.5σ (3σ) significance, regardless of their freeze-out temperature, and would be capable of detecting (or ruling out) all light relics with a freeze-out temperature smaller than the QCD phase transition (~ 0.2 GeV, or $\Delta N_{\rm eff} \sim 0.1$) at $> 5\sigma$ significance [1].

As shown in Fig. 10.13 the constraints from a PUMA-NK survey plateau with detector number N at around $N = 32$, suggesting that PUMA-32K constraints are limited by the foreground removal of modes inside the wedge. From Table 10.3, we see that both MegaMapper and PUMA-32K achieve similar constraints (~ 0.02 with CMB-S4). Thus MegaMapper is also sufficient for achieving a nearly cosmic-variance limited $N_{\rm eff}$ measurement from LSS. As one would expect, from Fig. 10.13 we see that these constraints significantly depend on the $k_{\rm max}$ assumed in the forecast, changing at the $\sim 25\%$ level between $0.1 < k_{\rm max}/k_{\rm nl} < 1$. From

[11] In our forecasts Y_p is treated as a derived parameter from the BBN consistency relation by default (see later in this section). The BBN calculation assumes a known value for the neutron lifetime and no other physics beyond the Standard Model. However, there is a current $4\,\sigma$ discrepancy between beam [439] and bottle [458] measurements of this lifetime. Thus it may be necessary to marginalize over Y_p to obtain an unbiased measurement of $N_{\rm eff}$.

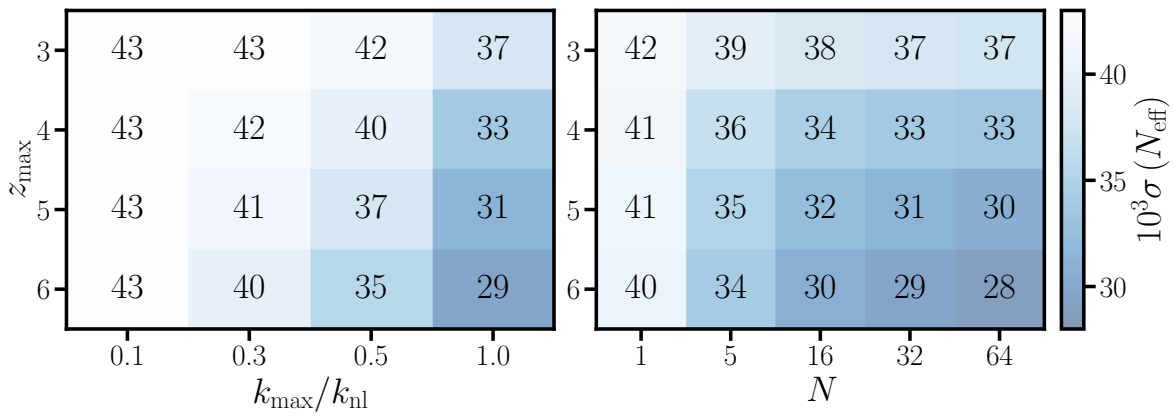


Figure 10.13: $N_{\rm eff}$ constraints from Planck+SO (primary CMB) +DESI+PUMA-NK (full shape) for variable dish number N and $z_{\max}$. No lensing data is included in this figure. In the right plot we fix $k_{\max}/k_{\rm nl} = 1$, while in the left plot we fix $N = 32$. For reference, the constraint from Planck+SO (primary CMB) +DESI (full shape) is 0.044.

Fig. 10.13 we also see that decreasing PUMA's redshift range to $2 < z < 4$, which halves the survey volume, only increases the $N_{\rm eff}$ constraint by $\sim 10\%$, which implies that most of the information is coming from the breaking of the degeneracies between cosmological parameters obtained by combining CMB with LSS data.

Our results are in good agreement with ref. [42], which show $N_{\rm eff}$ constraints dropping from 0.026 with CMB-S4 to 0.013 with CMB-S4 + PUMA-32K. Without lensing, our constraints are 0.019 for LSD + PUMA-32K. This apparent discrepancy can be entirely explained by differences in nuisance parameters. When non-linear nuisance parameters $(b_2, b_s, N_2, N_4, \alpha_2, \alpha_4)$ are held fixed, we find $\sigma(N_{\rm eff}) \sim 0.012$, which is in excellent agreement with ref. [42]. This suggests that improved models of galaxy formation, which could yield tight priors on non-linear nuisance parameters, would significantly improve a LSS measurement of $N_{\rm eff}$, with the "extra information" primarily coming from small scales.

In our fiducial analysis we have fixed Y_p as a function of $N_{\rm eff}$ using the Big Bang Nucleosynthesis (BBN) consistency relation [553, 325, 295], rather than varying it independently. However, as an important check of the Standard Model of particle physics and cosmology, we can also treat them as independent quantities. In this case, marginalizing over Y_p can significantly degrade the constraints on $N_{\rm eff}$ from CMB surveys. For example, the constraints from Planck + SO (LiteBIRD + S4) increase from 0.047 (0.025) to 0.116 (0.077) with Y_p marginalized over. When clustering data from high redshift surveys are included, this degradation is reduced by $\sim 50\%$. For example, the constraint from PSD + PUMA-32K (LSD + PUMA-32K) only increases from 0.030 (0.019) to 0.048 (0.041). Since LSS is insensitive to Y_p, the resulting constraints with Y_p marginalized over are primarily driven by LSS. These constraints are comparable to the SO goal (0.05), suggesting that future high-redshift LSS surveys will have a similar constraining power to a Stage-III CMB survey, but with the added

benefit of being immune to Y_p.

Primordial non-Gaussianity

In the simplest inflationary models one expands the action of a single scalar field minimally coupled to gravity about a spatially homogeneous solution. The leading order corrections to the action are quadratic in the small fluctuations about this solution. Thus these fluctuations are drawn from a Gaussian distribution to lowest order. Higher order terms in the action, arising from interactions of the inflaton with itself or other fields, result in trace amounts of Primordial non-Gaussianity (PNG). Of particular interest are local PNG, which vanish in the squeezed limit for slow-roll single field models of inflation [379, 157].

The non-Gaussian contributions to the primordial potential Φ are commonly modeled as a perturbative expansion in a Gaussian random field ϕ [531]:

$$\Phi(\boldsymbol{x}) = \phi(\boldsymbol{x}) + f_{\rm NL} \int d^3y\, d^3z\; W(\boldsymbol{y},\boldsymbol{z})\phi(\boldsymbol{x}+\boldsymbol{y})\phi(\boldsymbol{x}+\boldsymbol{z}) + \mathcal{O}(\phi^3), \tag{10.28}$$

where $W(\boldsymbol{y},\boldsymbol{z})$ is some kernel and $f_{\rm NL}$ controls the amplitude of the non-Gaussianity. In this chapter we specify to the local type of PNG in which $W(\boldsymbol{y},\boldsymbol{z}) = \delta^D(\boldsymbol{y})\delta^D(\boldsymbol{z})$. Currently the tightest constraints on the amplitude of local PNG $f_{\rm NL}^{\rm Loc} = -0.9 \pm 5.1$ come from the Planck satellite [153], which are within a factor of 3 of a cosmic-variance limited CMB survey [197]. LSS surveys are not yet competitive with the CMB, with current constraints around $\sigma(f_{\rm NL}^{\rm Loc}) \sim 25$ [120]. Pushing beyond $\sigma(f_{\rm NL}^{\rm Loc}) \sim 1$ requires the addition of future 3D LSS data, which has access to significantly more modes than the 2D CMB on large scales.

In addition to modifying the Universe's initial conditions (e.g. producing a non-zero bispectrum) PNG modifies the non-linear evolution of structure growth (producing new mode-couplings), which can in principle be modeled using standard perturbative techniques. However, the development of a PT-based code which self-consistently models non-linearities (in redshift space) in the presence of PNG is beyond the scope of this thesis. We instead focus on PNG's effect on the linear bias, and restrict our forecasts to $k_{\rm max} < 0.06\ h{\rm Mpc}^{-1}$ where linear theory is a good approximation for the high redshifts we consider. As seen in Eq. 10.28, PNG couples short wavelength modes of Φ to long wavelength modes of ϕ. This coupling induces an additional scale-dependent bias ($b \to b + \Delta b$) which takes the form [556]:

$$\Delta b \equiv \frac{3\Omega_m H_0^2}{2k^2 D(z) T(k)} b_\phi f_{\rm NL}^{\rm Loc} = \frac{3\Omega_m \delta_c H_0^2}{k^2 D(z) T(k)} (b-1) f_{\rm NL}^{\rm Loc}, \tag{10.29}$$

where $D(z)$ is normalized to a in the matter dominated era and $\delta_c = 1.686$. In the above equation we made the assumption that $b_\phi = 2\delta_c(b-1)$ [556, 86], which does not necessarily hold for the galaxies that will be observed by current and future LSS surveys [87, 60, 419], and it cautions against combining the constraints from different instruments.

Shown in Fig. 10.14 are our forecasted constraints on $f_{\rm NL}^{\rm Loc}$ from a variety of upcoming or proposed LSS surveys, using the power spectrum only. Since all information about $f_{\rm NL}^{\rm Loc}$

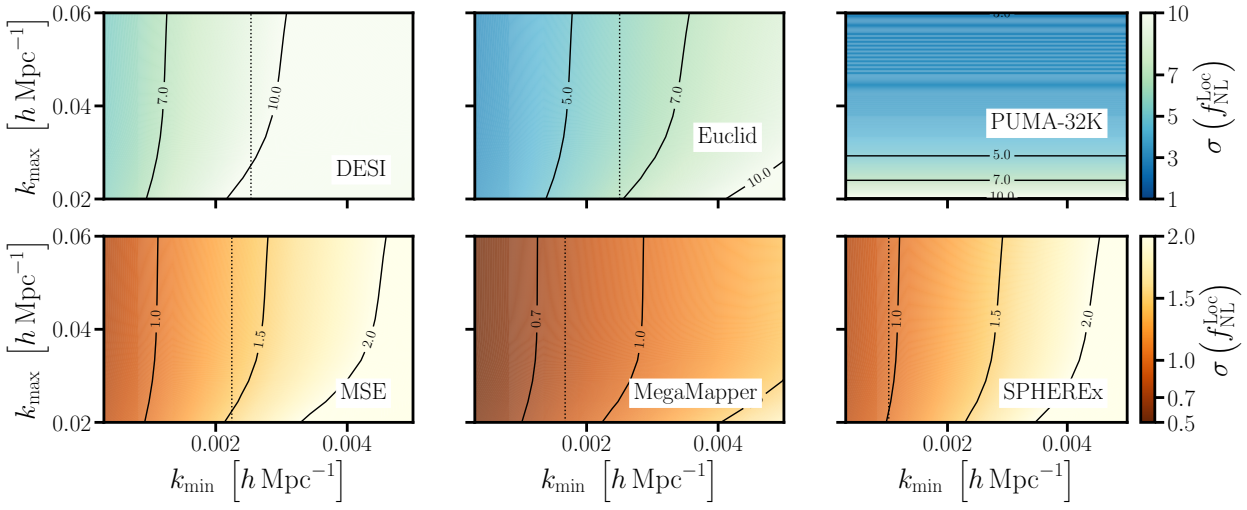


Figure 10.14: Constraints on local primordial non-Gaussianity as a function of $k_{\rm min}$ and $k_{\rm max}$, using the full shape power spectrum without any CMB priors. For SPHEREx we set $\sigma_z/(1+z) = 0.05$, while for all other surveys we set $\sigma_z = 0$. The vertical lines are the lowest possible $k_{\rm min}$ for each survey: $2\pi/V^{1/3}$, where V is the volume of the survey (including sky coverage).

comes from low k where linear theory is an excellent approximation, in our forecasts we only marginalize over ΛCDM, b and T_b (for 21-cm surveys[12]). For all surveys we include a prior on ΛCDM from Planck. Even with just the power spectrum, we see that near-term LSS surveys (such as Euclid) have the capability of achieving similar constraints on PNG as Planck, while a future high-redshift LBG survey has the potential to achieve $\sigma(f_{\rm NL}^{\rm Loc}) \sim 0.8$, more than a factor of 2 improvement over a cosmic variance limited CMB survey. Similar constraints from the power spectrum should be obtainable with the SPHEREx mission [197].

Here we have only focused on local PNG from the power spectrum, making use of the $f_{\rm NL}^{\rm Loc}$ induced scale-dependence of linear bias. We note that adding bispectrum or higher point function information can potentially help [488, 319, 419]. Extracting bispectrum information will be crucial for constraining shapes of PNG beyond the local type considered here (such as equilateral or orthogonal, for example), where the effects on the power spectrum are largely degenerate with non linear evolution [249].

Finally we note that the signatures of a non-zero $f_{\rm NL}^{\rm Loc}$ appear on very large scales, which are often most affected by systematics (especially angular measurements) [322, 487, 335, 23, 505, 615]. In Section 10.5 we explore deprojecting angular modes by introducing a minimum $k_\parallel$ cut. Alternatively, one can use the fact that most galaxy imaging systematics [304] are uncorrelated with CMB lensing maps, and measure the scale dependent bias by cross-correlating the galaxy positions with CMB lensing. Ref. [534] found that $\sigma(f_{\rm NL}^{\rm Loc}) \lesssim 1$ can be

[12]See ref. [161] for a detailed discussion of degeneracies between 21-cm foregrounds and $f_{\rm NL}^{\rm Loc}$.

achieved by future surveys using this technique, thus providing an important independent check on systematics.

Primordial features

Current observations of the CMB are consistent with the primordial potential having a nearly scale-invariant power spectrum. This is a generic prediction of single-field slow-roll inflation, sourced by the (near) flatness of the inflationary potential [68]. However, attempts to connect this simple picture of inflation with more fundamental ultra-violet scenarios often result in deviation from scale invariance on some scales [557]. In this chapter we focus on a subclass of these features (sharp features), which are commonly modeled as a step in the inflationary potential [17]. This step produces a scale-dependent oscillatory feature in the primordial power spectrum, which we model as a sinusoidal modulation of the linear power spectrum: $P_m(k) \to [1 + A_{\rm lin} \sin(\omega_{\rm lin} k + \phi_{\rm lin})] P_m(k)$ [55, 83, 476].

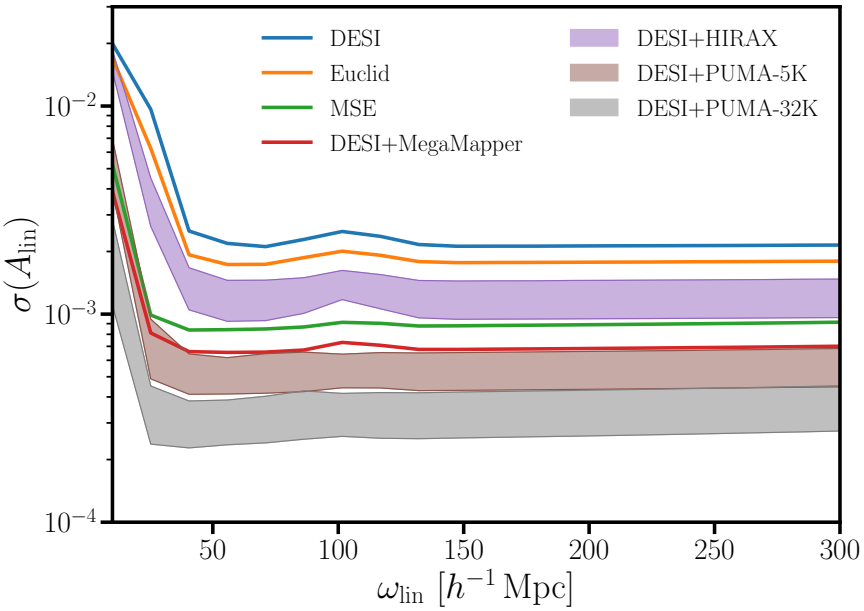


Figure 10.15: Constraints on $A_{\rm lin}$ for various fiducial values of $\omega_{\rm lin}$, assuming $\phi_{\rm lin} = \pi/2$. For all surveys we include a Planck+SO prior on ΛCDM. The slight "bump" in the constraints at $\omega_{\rm lin} \sim 100\ h^{-1}$Mpc is due to the degeneracy with the BAO frequency.

In our forecasts we fix both $\omega_{\rm lin}$ and $\phi_{\rm lin}$, and find the corresponding uncertainty on $A_{\rm lin}$ after marginalizing over the parameters in Table 10.2. Our results are shown in Fig. 10.15 for $10\ h^{-1}{\rm Mpc} < \omega_{\rm lin} < 300\ h^{-1}{\rm Mpc}$, assuming $k_{\rm max} = k_{\rm nl}$[13]. We find that near-term

[13]Since the high frequency oscillatory feature induced by primordial features is immune to small-scale non-linearities [83] up to an overall damping of their amplitude [56], one could in principle extend the $k_{\rm max}$

spectroscopic surveys will be capable of achieving $\sigma(A_{\rm lin}) \sim 2 \times 10^{-3}$, while HIRAX has the capability to improve upon these constraints by a factor of two. Order of magnitude improvement are possible by extending out to high redshift, with an optimistic configuration of PUMA-32K achieving $\sigma(A_{\rm lin}) \sim 2 \times 10^{-4}$.

These results are largely consistent with ref. [83], who found $\sigma(A_{\rm lin}) \simeq 2 \times 10^{-3}, 1.5 \times 10^{-3}$, and 10^{-4} for DESI, Euclid, and a cosmic variance limited LSS survey (out to $z = 6$) respectively. Our results for DESI and Euclid are within $\sim 25\%$ of these values, which is entirely reasonable given that ref. [83] uses a different forecasting approach in which $P_{gg}^{\rm w}(k)/P_{gg}^{\rm nw}(k)$ is the observable of interest. This suggests that PUMA-32K (with optimistic foreground assumptions) is within a factor of two of a cosmic variance limited survey. We note that the slight "bump" in our constraints around $\omega_{\rm lin} \sim 100\ h^{-1}$Mpc is smaller than the bump found by ref. [83], which is due to a degeneracy with the BAO frequency. This suggests that the full shape information is sufficient to break the degeneracy between primordial features and the BAO feature.

Dynamical dark energy

Models of dark energy are typically invoked to explain the observed accelerated expansion of the Universe at the present epoch. Current observations are consistent with the existence of a cosmological constant Λ. In this section we quantify the sensitivity of future surveys to departures from the Λ-model, using both a "model-independent" approach[14], and through constraints on a specific model of dynamical dark energy.

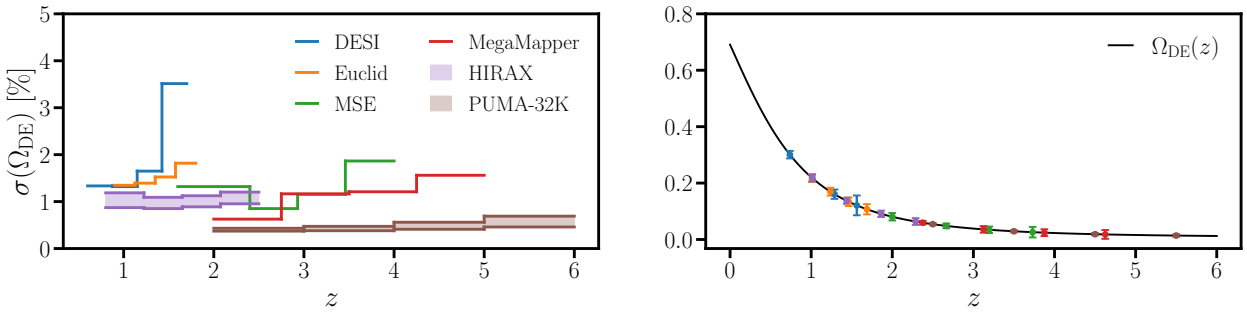


Figure 10.16: The absolute error on the dark energy fraction ($\Omega_{\rm DE}$) achievable by current and future LSS surveys.

used in an analysis beyond the nonlinear scale. In practice, however, we find that both the MegaMapper and PUMA-32K constraints improve by at most 2% when $k_{\rm max}$ is increased from $k_{\rm nl}$ to $1.4\,k_{\rm nl}$. For MegaMapper this is mostly due to the large value of the shot noise, which limits the ability of reconstructing oscillatory features on small scales.

[14]Ref. [354] performs a similar analysis to estimate constraints on $\Omega_{\rm DE}(z)$, focusing on the redshift interval $1.5 < z < 2.5$.

Over the redshift range covered by our observations (typically up to $z \sim 5-6$), direct expansion and growth measurements can constrain any dynamical dark energy model by quantifying the fraction of dark energy in a number of redshift bins as we now explain. At redshifts $z < 6$ the Universe is dominated by nonrelativstic matter and dark energy, so that the first Friedmann equation is approximately given by $\Omega_{\rm DE}(z)+\Omega_m(z) = 1$, where $\Omega_m(z) = \rho_m(z)/\rho_{\rm crit}(z) = \omega_m(1+z)^3/h^2(z)$ and $h(z)$ is defined through $H(z) = 100\,h(z)$ km/s/Mpc. We estimate the error on the dark energy fraction via standard error propagation, and neglect the covariance between $h(z)$ and ω_m, so that $\sigma^2_{\Omega_{\rm DE}(z)} = \Omega^2_m(z)(4[\sigma_{h(z)}/h(z)]^2 + [\sigma_{\omega_m}/\omega_m]^2)$. We take the relative error on $h(z)$ from the radial BAO measurements ($\sigma_{h(z)}/h(z) = \sigma_{\alpha_\parallel(z)}$), which are described in sections 10.3 and 10.4. We assume that the measurement of ω_m is driven primarily from CMB measurements, and take $\sigma_{\omega_m} = 2.4 \times 10^{-4}$ for all surveys, motivated by the constraints achievable by LSD + MegaMapper (see Table 10.3). Our results are shown in Fig. 10.16. We find that MegaMapper will be able to constrain the dark energy fraction to better than 2% across its redshift range, while PUMA-32K could improve these measurements to the sub-percent level.

While the simplest models have $\Omega_{\rm DE}$ drop rapidly at higher redshift, so that dark energy is a late-time phenomenon, there is little guidance from theory as to how $\rho_{\rm DE}$ should evolve, and the presence of an inflationary epoch in the very early Universe calls into question whether there should be only one late-time epoch of dark energy domination. Recently, early dark energy (EDE) has been invoked to decrease the CMB sound horizon via an increase in the expansion rate before recombination. Doing so would increase the inferred value of H_0 from CMB observations, making EDE a possible candidate for resolving the H_0 tension [207]. Current constraints from a combination of CMB, SNe and LSS data tend to disfavor this explanation [288, 307, 163, 324], with much of the constraint coming from EDE's effects on the growth of large-scale structure.

This effect is much more general than the above-mentioned example however, and future LSS surveys can surprisingly put very tight constraints on EDE models at redshifts much higher than the range covered by the observations, thanks to the exquisite precision with which they measure the shape of the power spectrum. Since the growth of structure is damped by Hubble expansion, long periods where unclustered species dominate the expansion lead to suppression of large-scale structure that can be detected by comparing early- and late-time measures of the fluctuation amplitude. A precise measurement of the power spectrum shape can also be used to place constraints on (or detect) short-term deviations from matter or radiation domination, since such periods will change the shape of the power spectrum due to differential growth of modes. An accurate measurement of the power spectrum constrains deviations in the expansion history over a broad range of redshifts [137].

We give an illustrative example in Fig. 10.17, which considers the EDE scenario of ref. [484], in which a light scalar ϕ with potential $V(\phi) \propto [1-\cos(\phi/f)]^3$ is initially displaced from its minimum. The field is parameterized by its maximum energy contribution, $f_{\rm EDE} \equiv \max[\rho_{\rm EDE}(z)/\rho_{\rm tot}(z)]$, which occurs at redshift z_c, and the initial value of the field $\theta_i \equiv \phi_i/f$. Values of $z_c \gtrsim z_\star$, where $z_\star \sim 1100$ is recombination, would help to relieve

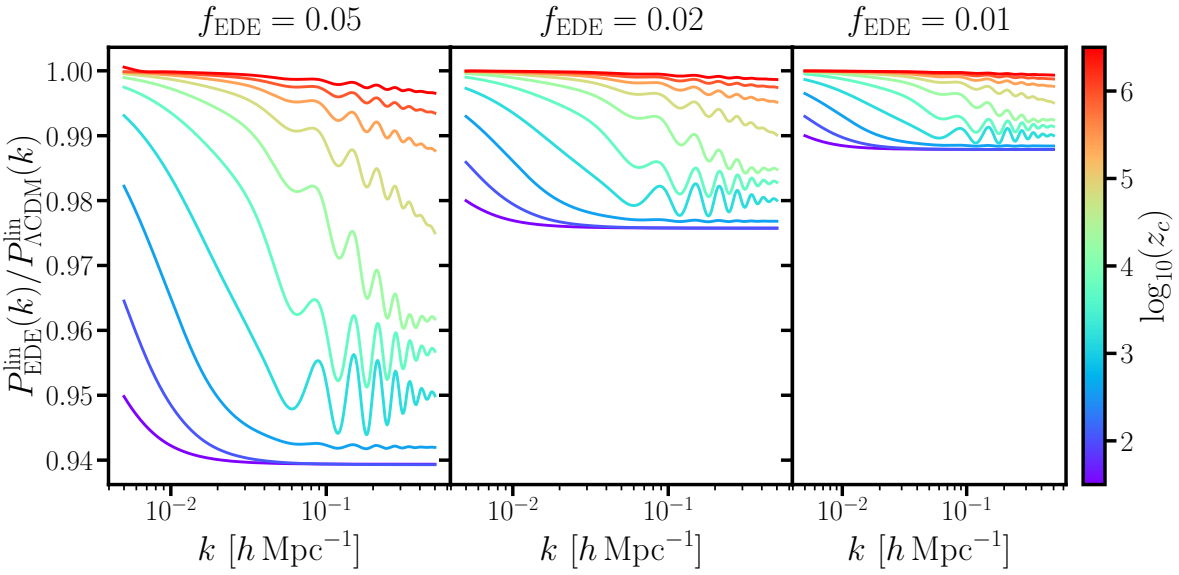


Figure 10.17: Change in the linear matter power spectrum (at $z = 2$) when including Early Dark Energy (EDE) for various values of $f_{\rm EDE}$ and $1.5 < \log_{10}(z_c) < 6.5$. $\theta_i = 2.83$, and all standard cosmological parameters are held at their fiducial ΛCDM values. For $z_c \lesssim 30$, the ratio $P^{\rm lin}_{\rm EDE}(k)/P^{\rm lin}_{\Lambda{\rm CDM}}(k)$ is roughly given by $1 - f_{\rm EDE}$.

the Hubble tension, but we show that future large-scale structure surveys can probe even small values of $f_{\rm EDE}$ over a very broad range of redshifts. Using the publicly available code `CLASS_EDE` [288] to produce the initial power spectra and `velocileptors` to treat the damping of features from non-linear evolution and biasing [137], we forecast constraints on $f_{\rm EDE}$ for values of z_c ranging from ~ 30 to 10^6. In all forecasts z_c and $\theta_i = 2.83$ are held fixed, and we marginalize over the parameters listed in Table 10.2. The results are shown in Fig. 10.18.

We see that near-term spectroscopic surveys are capable of constraining the amplitude of early dark energy at the few percent level out to redshift 10^4, while PUMA would be capable of achieving subpercent level constraints out to redshift 10^5. These are order of magnitude improvements over the $\mathcal{O}(10\%)$ constraints found by refs. [560, 288] from a combination of current BOSS and Planck data. As shown in Fig. 10.17, EDE that peaks prior to recombination produces a phase shift in the BAO feature. From Fig. 10.17 we see that this effect peaks at around $z_c \sim 10^{3.5}$, which sensibly corresponds to the minimum in Fig. 10.18. However, EDE that peaks after recombination induces no phase shift in the BAO signature. In this case, EDE damps the linear power spectrum in a scale-independent way (for $k > 0.01\ h{\rm Mpc}^{-1}$) that is nearly degenerate with A_s, b and T_b. Thus all of the information about post-recombination EDE comes from low k. This puts 21-cm observations at a disadvantage, as low k modes are excluded by the foreground wedge. We see this effect in Fig. 10.18. PUMA-32K offers no improvement over DESI for $z_c < z_\star$, while MegaMapper improves over DESI by more than a factor of two. We note that constraints in the $z_c < z_\star$

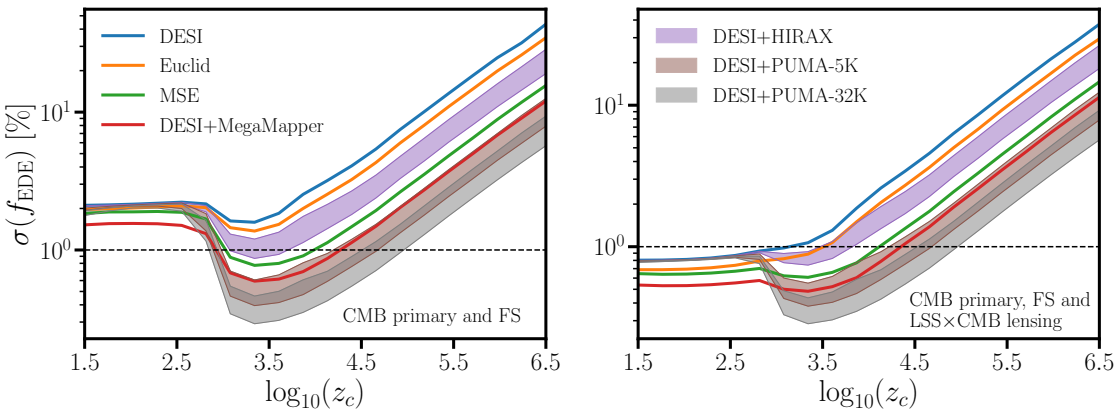


Figure 10.18: Constraints on the maximum amplitude of early dark energy ($f_{\rm EDE}$) as a function of the time at which EDE peaks z_c, assuming $\theta_i = 2.83$. We include a Planck+SO prior on ΛCDM for all experiments. In the left panel we show constraints from full shape (FS) measurements only, while in the right panel we include a prior on ΛCDM and the nuisance parameters in Table 10.2 from SO lensing and cross-correlations with the respective galaxy surveys, following the convention in §10.3.

regime benefit significantly from the addition of lensing cross-correlations (MegaMapper constraints decrease by more than a factor of ~ 2), which soften the $f_{\rm EDE} - b$ degeneracy.

Modified gravity

Modifications to GR have been proposed as an alternative explanation for the observed accelerated expansion of the Universe at low redshift [311, 314, 37, 557]. The phenomenology of such models can be very rich, but a common ingredient is that the responses of the Newtonian potential Ψ and spatial curvature Φ to matter and energy (which are equal in the low redshift Universe in GR) can differ.

To estimate the sensitivity of future surveys to deviations from GR we examine the gravitational slip parameter $\gamma = \Phi/\Psi$ [151, 627, 315, 314, 37, 557]. We follow the method of ref. [131], treating γ as a scale-independent factor, and consider a scenario in which Poisson's equation for Φ is modified but Ψ is left unchanged. This amounts to making the substitution $\delta_\kappa = c_\kappa \delta_m$ in Eq. 10.14, where $c_\kappa = (1+\gamma)/2$, since lensing probes the sum of the two potentials $\Phi + \Psi$. We approximate c_κ as a piecewise constant linear Eulerian bias, whose value in each redshift bin is treated as an independent parameter.

In Fig. 10.19 we show constraints on γ_i, where γ_i is the value of the gravitational slip in the $i^{\rm th}$ redshift bin. In these constraints we include LSS×CMB lensing, full shape clustering data, the primary CMB, and marginalize over all parameters listed in Table 10.2. The direct constraints on γ_i come only from $C_\ell^{\kappa\kappa}$ and $C_\ell^{\kappa g_i}$, while the full shape and primary CMB data serve as priors on ΛCDM and the nuisance parameters. Thus the constraints in Fig. 10.19

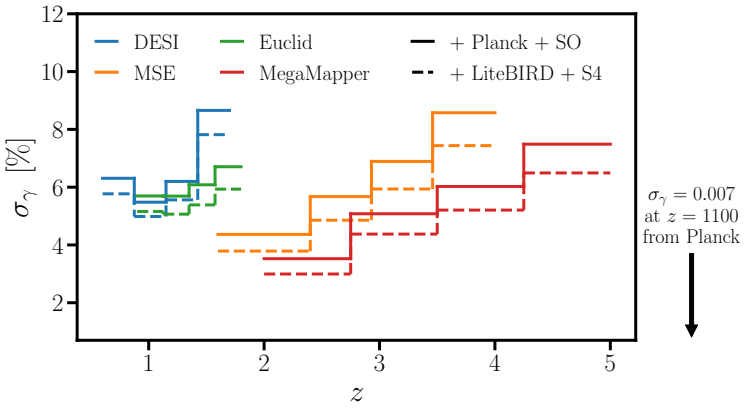


Figure 10.19: Error on the gravitational slip parameter γ from LSS×CMB lensing cross-correlations, LSS full shape data, and the primary CMB. For comparison, we show the constraint on γ from Planck [473] at $z \approx 1100$ (recombination) to the right of the plot.

should not be interpreted as independent measurements. Instead, they should be interpreted as a survey's sensitivity to γ_i as measured by the full dataset (lensing cross-correlations and full shape clustering across the entire redshift range of the survey), assuming ΛCDM as the fiducial cosmology.

From Fig. 10.19 we see that spectroscopic surveys are capable of measuring $\gamma(z)$ to better than 10% out to redshift 6 with current CMB data. The improvement from upgrading to LiteBIRD + CMB-S4 from Planck + SO is $\sim 15\%$ for all LSS surveys at all redshifts of interest. We see that MegaMapper improves over MSE by $20-40\%$ across $2 < z < 4$. These results are in rough agreement with ref. [131], who considered an LBG survey covering 1000 square degrees and found $0.02 < \sigma_\gamma < 0.20$ for $2 < z < 5$ after marginalizing over σ_8 and $b_{\rm LBG}$. If we only marginalize over $\ln(A_s)$ and $b_{\rm LBG}$, and reduce MegaMapper's sky coverage to $f_{\rm sky} = 0.024$, mimicking the approach of ref. [131], we find $0.09 < \sigma_\gamma < 0.23$ for $2 < z < 5$. The disagreement at low redshift is likely due to different binning schemes (we use four redshift bins, while ref. [131] uses three), and differences in the assumed LBG number densities.

Extended base model and multi-extension fits

Our discussion so far focused on one parameter extensions of our base ΛCDM model. We now investigate how our constraints degrade as we extend our base model by a single parameter, or as we simultaneously vary multiple extensions.

We first consider the constraints on the ΛCDM parameters when the base model is extended by a single parameter. In Table 10.4 we show the percent increase to PSD + MegaMapper (using full shape data only, without lensing) constraints as we extend the base model to include M_ν, $N_{\rm eff}$, Ω_k or $f_{\rm EDE}$ (with $\log_{10}(z_c) = 3.61$). In particular, we find that

constraints on h, ω_c and ω_b worsen by roughly a factor of 2 when $N_{\rm eff}$ is allowed to vary. When the base model is extended to include M_ν or Ω_k, all ΛCDM constraints increase by less than 40%. Allowing $f_{\rm EDE}$ to vary, however, has a negligible (few percent) impact on ΛCDM constraints.

In Table 10.4 we also consider the impact of extending the base model on several one-parameter extensions. In particular, we find that $\{M_\nu, N_{\rm eff}, \Omega_k\}$ are not internally degenerate – allowing any two of these three extensions to vary simultaneously worsens the constraints by less than 17%. We find a mild degeneracy between $f_{\rm EDE}$ and Ω_k; varying the two simultaneously worsens the constraints on both parameters by $\approx 30\%$.

Extension	h	$\ln(A_s)$	n_s	ω_c	ω_b	τ	M_ν	$N_{\rm eff}$	Ω_k	$f_{\rm EDE}$
M_ν	39.7	29.5	2.7	26.6	0.0	35.3	–	11.5	12.5	16.2
$N_{\rm eff}$	138.0	0.0	46.0	67.0	73.0	1.0	11.5	–	0.3	8.5
Ω_k	23.4	1.3	12.4	32.6	0.0	3.4	12.5	0.3	–	31.8
$f_{\rm EDE}$	0.4	1.6	4.4	0.8	3.4	2.5	16.2	8.5	31.8	–

Table 10.4: Percent increase to PSD + MegaMapper constraints on ΛCDM (columns 2-7) and several one-parameter extensions (columns 8-11) when the base model is extended by a single parameter. The constraint on $f_{\rm EDE}$ is for fixed $\log_{10}(z_c) = 3.61$ and $\theta_i = 2.83$.

In Fig. 10.20 we consider changes in our constraints as we vary more than one extension simultaneously. The blue bars in the left panels of Fig. 10.20 shows the $1\,\sigma$ uncertainty on the ΛCDM parameters from PSD, which we compare to PSD + PUMA-32K, shown in orange (the analogous figure for MegaMapper looks nearly identical). Allowing for multiple extensions of the baseline model results in an overall degradation of the constraints, as shown by the different shading of the blue and orange bars. When the neutrino mass is marginalized over, this effect is the strongest on h, A_s and τ. This trend increases when simultaneously fitting for $N_{\rm eff}$ and Ω_k. The net improvement over PSD brought by future surveys is also reduced as more extensions are added to the base ΛCDM model, indicating that the set of observables we consider is not powerful enough to break the new degeneracies among parameters. Fitting to more than one extension to ΛCDM simultaneously without significant degradation to ΛCDM requires next generation space-based CMB surveys, such as LiteBIRD, or other measurements which are capable of constraining the primordial amplitude and optical depth to extremely high precision. These constraints are shown by the red and green bars respectively. As seen in Fig. 10.20, the constraints on τ and A_s from next generation CMB surveys are negligibly impacted when simultaneously fitting for multiple extensions. Thus no additional penalty is incurred to the remaining four ΛCDM parameters from the would-be worsened degeneracies with τ and A_s.

On the other hand, the constraints to our selected extensions to ΛCDM (right panel of Fig. 10.20) remain relatively unchanged while fitting to more than one extension. Just as

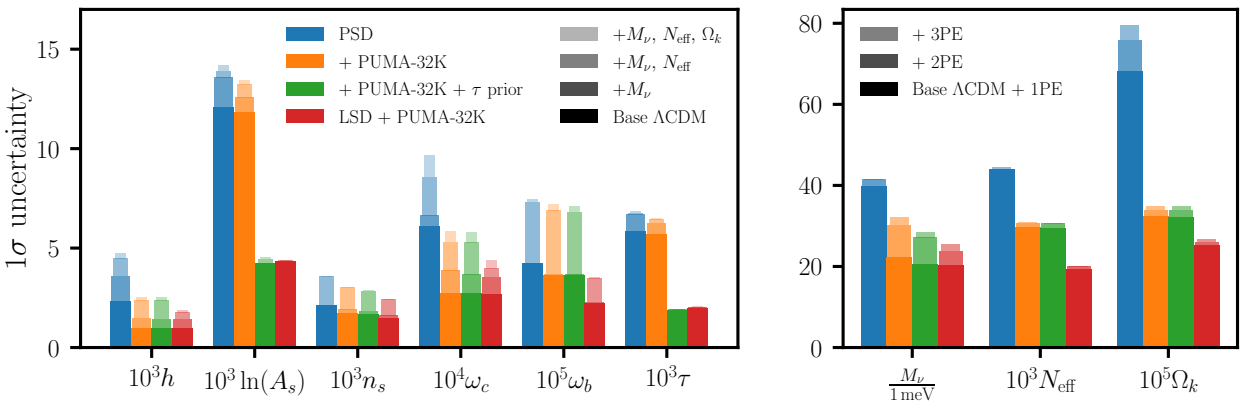


Figure 10.20: Constraints on ΛCDM and extensions from four different combinations of experiments, using primary CMB and clustering data only. We include a $\sigma(\tau) = 0.002$ prior for the green bars. *Left*: The shading/width denotes which parameters are allowed to vary when fitting for ΛCDM (e.g. for Base ΛCDM+M_ν we keep $N_{\rm eff}$ and Ω_k fixed). *Right*: The shading denotes the number N of Parameter Extensions (NPE) that are simultaneously allowed to vary in our fits, starting with the relevant parameter on the x-axis and continuing along the cycle $(M_\nu, N_{\rm eff}, \Omega_k)$. E.g. for Ω_k, 2PE is equivalent to Ω_k, M_ν.

before (Table 10.4), this suggests that $\{M_\nu, N_{\rm eff}, \Omega_k\}$ are not internally degenerate.

Linear vs. non-linear modeling

Currently, the vast majority of forecasts assume linear theory and scale-independent bias (though see [147, 104, 142] for notable exceptions), and model the effects of FoG-like contributions with a specific functional form (typically a Gaussian or Lorentzian). That is, the typical linear model (LM) for the redshift-space galaxy power spectrum takes the form [624]:

$$P_{gg}^{\rm LM}(\boldsymbol{k}) = \frac{(b + f\mu^2)^2}{1 + \sigma_v^2 (k\mu)^2} P^{\rm lin}(k) + N_0. \tag{10.30}$$

To explore the difference between this simplified linear model and our fiducial analysis, we create forecasts for DESI, MegaMapper and PUMA-32K (with optimistic foregrounds) following the procedure of §10.3 with Eq. (10.30) replacing the full model (FM), given by Eq. (10.7). A comparison between the constraints on ΛCDM and a selection of single parameter extensions is shown in Fig. 10.21. From Fig. 10.21 we see that the linear model generically predicts overly optimistic constraints: for both PSD and PSD + MegaMapper, the constraints using the full model are $20 - 40\%$ larger than those using the linear model for most parameters, with n_s being a notable exception, where the linear and full models disagree by factors of 3 and 6 for the respective survey combinations. For PSD + PUMA-32K the disagreement is significantly worse: $\sigma(M_\nu)$ and $\sigma(N_{\rm eff})$ from linear theory are overly optimistic by a factor of ~ 2, while the constraints on n_s disagree by a factor of 12.

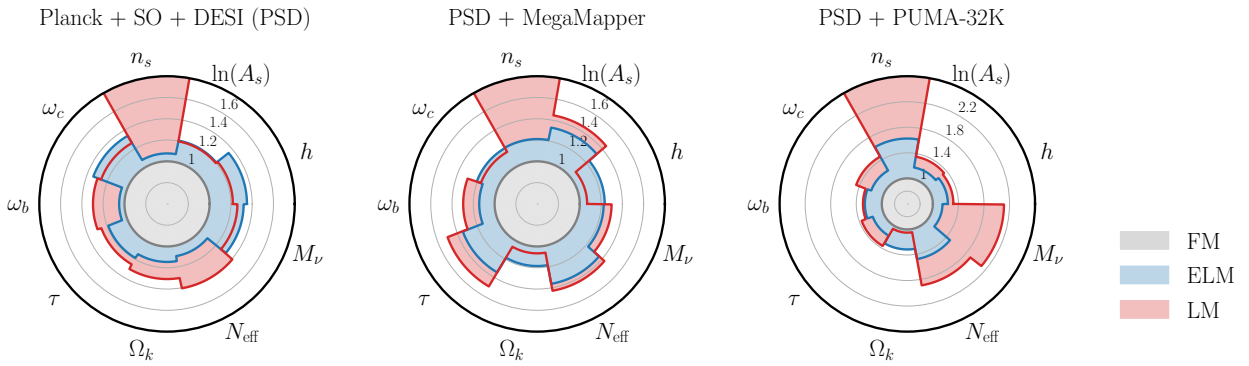


Figure 10.21: Constraints on ΛCDM and several one-parameter extensions when using the full model (FM), relative to the linear model (LM, red) or extended linear model (ELM, blue). In this plot we only include the primary CMB and full shape clustering data. Note that the largest radius in the two leftmost panels is 1.8, but 2.4 in the right panel. For the linear model, the values for n_s extend well beyond these limits: from left to right the ratios are 3.1, 6.8, and 11.6.

Leaving aside the fact that linear theory is clearly not sufficient for an unbiased estimate of cosmological parameters in the galaxy power spectrum, linear theory results in inaccurate forecasts for several reasons. Perhaps the most obvious is that the approximation of linear dynamics breaks down at intermediate to high k. Two beyond-linear effects are particularly important: the broadening of the BAO peak due to non-linear evolution and decorrelation of the evolved field with the initial conditions [85, 210, 388, 387, 160, 449, 442, 117, 579, 394, 623, 544, 533, 50, 609, 398, 98, 545]. In Fig. 10.7 we show a few logarithmic derivatives of the tracer power spectrum using both linear and non-linear theory at $z = 2.25$. The non-linear scale at this redshift is roughly $0.46\,h\,\mathrm{Mpc}^{-1}$, and we see significant discrepancies (more than a few percent) by $k \sim 0.2\,h\,\mathrm{Mpc}^{-1} < 0.5\,k_{\rm nl}$. In addition to linear theory breaking down on small scales, linear forecasting is inaccurate since the true angular structure is more complex than the $b + f\mu^2$ form allows, the bias model has too few degrees of freedom, linear forecasting ignores contributions from small scale physics (i.e. counterterms) and the FoG model is overly simplified.

The linear model can be trivially extended to address the latter two shortfalls: the effects of small-scale physics can be mimicked with the counterterm-like contribution $\mu^{2n}k^2P^{\rm lin}$, and the functional dependence of the FoG-like contribution can be relaxed via a Taylor expansion. Doing so results in the "extended linear model" (ELM):

$$P_{gg}^{\rm ELM} = (b + f\mu^2)^2 P^{\rm lin} + (\alpha_0 + \alpha_2\mu^2 + \alpha_4\mu^4)(k/k_*)^2 P^{\rm lin} + N_0 + N_2(\mu k)^2 + N_4(\mu k)^4, \quad (10.31)$$

where we have suppressed the explicit $\boldsymbol{k}$ dependence. A comparison of the ELM with the FM is shown in Fig. 10.21. For PSD (PSD + MegaMapper) we find a minor improvement

in the ELM over the LM, with the mean disagreement[15] with the full model decreasing from 28% (28%) to 21% (25%). Most notable is the improved ELM prediction of $\sigma(n_s)$, which agrees with the full model to within 20% for both PSD and PSD + MegaMapper. For PSD + PUMA-32K, we find that the mean disagreement with the full model decreases from 44% to 25%, which is primarily driven by the significantly more accurate ELM predictions of $\sigma(M_\nu)$ and $\sigma(N_{\rm eff})$, disagreeing by 20% and 40% with the full model, as opposed to a factor of 2. Likewise, the disagreement for $\sigma(n_s)$ decreases from a factor of 12 to 50%.

10.5 Observational challenges

In the previous section we neglected or delayed discussing the effects of several observational challenges that would arise in a more realistic treatment of future surveys. In this section we quantify some of these effects, summarizing the importance of the 21-cm foreground wedge, examining redshift uncertainties, the impact of a shallower u-band depth for LBG surveys, and the impact of a $k_\parallel$ cut on constraints of primordial non-Gaussianity.

Throughout §10.4 we presented both optimistic and pessimistic forecasts for 21-cm surveys, representing the ability to mitigate foreground contamination in modes with low $k_\parallel$. Here we summarize the main results. Foreground assumptions can impact the precision of distance measurements perpendicular to the line of sight by more than a factor of 2, while measurements parallel to the line of sight are negligibly impacted. For PSD + PUMA-32K, the constraints on h, ω_c, $A_{\rm lin}$ and Ω_k vary by more than 50% depending on the foreground assumptions, while the constraints on $N_{\rm eff}$ and M_ν vary by more than 30%. For EDE the constraints can vary by more than a factor of 2. Finally, while the 21-cm constraints of $f_{\rm NL}^{\rm Loc}$ under optimistic foreground assumptions will not be competitive with a high-z galaxy survey, pessimistic foreground assumptions would inhibit a 21-cm measurement of $f_{\rm NL}^{\rm Loc}$ from scale-dependent bias altogether.

As previously discussed in §10.3, at high redshifts it becomes increasingly difficult to obtain small redshift errors, while the effects of redshift errors become increasingly more important due to the non-linear scale shifting to higher k. In Fig. 10.22 we show the impact of redshift errors on the constraints of ΛCDM and several one-parameter extensions from MegaMapper full shape data. We find that redshift errors negligibly impact our results for $\sigma_z/(1+z) < 10^{-3}$, only degrading the constraints at the few percent level. However, when $\sigma_z/(1+z)$ nears 0.3% we begin to see non-negligible ($\sim 25\%$) effects on some parameters, most notably M_ν, $f_{\rm EDE}$ and $A_{\rm lin}$. As shown in the right pannel of Fig. 10.22, including redshift errors is roughly equivalent to enforcing a strict $k_{\rm max}/k_{\rm nl}$ cut, with $\sigma_z/(1+z) = 0.01$ (0.003) corresponding to roughly $k_{\rm max}/k_{\rm nl} = 0.3$ (0.6) for our fiducial MegaMapper survey.

To estimate the impact of a lower u-band depth on MegaMapper constraints, we consider the idealized MegaMapper-like survey described in §10.2, which we split into four redshift bins spanning $2 < z < 5$. To mimic the effect of a lower u-band depth, we vary $m_{\rm UV}$ in

[15] The "mean disagreement" $\equiv \sum_\theta [\sigma^{\rm FM}(\theta)/\sigma^{\rm ELM}(\theta) - 1]/8$, where $\theta \in \{h, \ln(A_s), \omega_c, \omega_b, \tau, \Omega_k, N_{\rm eff}, M_\nu\}$.

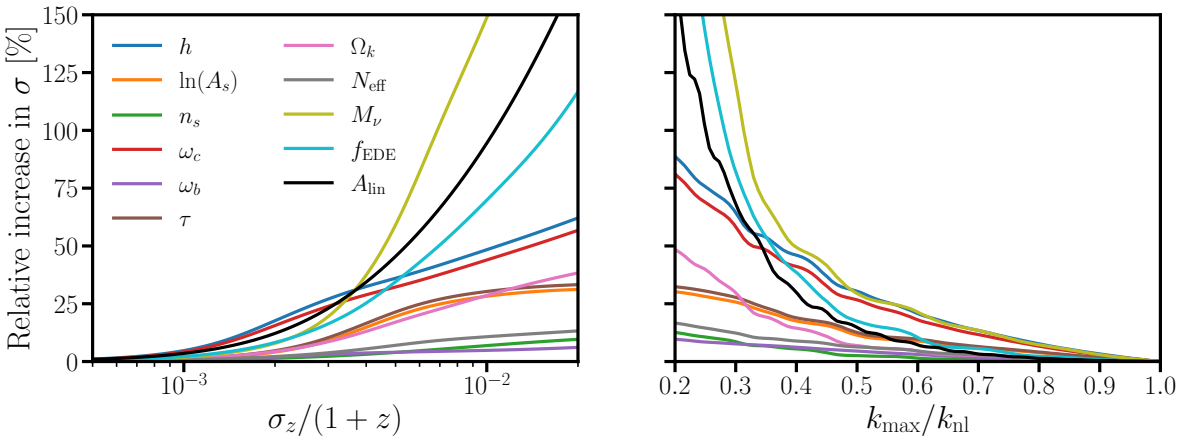


Figure 10.22: *Left*: Constraints on ΛCDM and several one-parameter extensions from MegaMapper full shape data vs redshift uncertainty, with $k_{\rm max}/k_{\rm nl} = 1$. The constraints are measured relative to $\sigma_z = 0$. *Right*: Change in constraints as a function of $k_{\rm max}/k_{\rm nl}$, relative to $k_{\rm max}/k_{\rm nl} = 1$. In both plots we include a prior on $\{\Lambda{\rm CDM}, N_{\rm eff}, \Omega_k, M_\nu\}$ from Planck + SO. The constraint on $f_{\rm EDE}$ is with fixed $\log_{10}(z_c) = 3.61$, while the constraint on $A_{\rm lin}$ is with fixed $\omega_{\rm lin} = 147\ h^{-1}\,{\rm Mpc}$.

the two central redshift bins (spanning $2.75 < z < 4.25$) between 23 and 24.5. We leave $m_{\rm UV} = 24.5$ in the lowest and highest redshift bins unchanged, and assume a mean 50% target efficiency ($\bar{n}/n_{\rm LBG} = 0.5$) in all redshift bins, regardless of $m_{\rm UV}$. As shown in the right panel in Fig. 10.23, the constraints from Planck + SO + MegaMapper on ΛCDM and common one parameter extensions ($N_{\rm eff}$, M_ν, Ω_k) increase by less than $\approx 15\%$ as the u-band depth is decreased from 24.5 to 23, while the constraints on $f_{\rm EDE}$ (at fixed $\log_{10}(z_c) = 3.61$) and $A_{\rm lin}$ (at fixed $\omega_{\rm lin} = 147\ h^{-1}\,{\rm Mpc}$) increase by less than 25%.

We also perform a simplistic area vs depth trade-off study for the same MegaMapper-like survey, which we discuss here. For fixed observing time, we assume the sky coverage and the change to the limiting apparent magnitude scale as $f_{\rm sky} \propto t^{-1}$ and $\delta m_{\rm UV} \propto 2.5\log_{10}(\sqrt{t})$ respectively, which should be a very good approximation given that MegaMapper will be deep into the sky-noise limit. The fiducial survey corresponds to $t = 1$. In principle, changing the limiting magnitude changes the bias of the LBG sample. We account for this effect by rescaling $n(m = 24.5, z) \to n(m, z)b^2(m, z)/b^2(m = 24.5, z)$ in our covariance.

As shown in the left panel of Fig. 10.23, the fiducial MegaMapper sky coverage ($f_{\rm sky} = 0.34$) is well optimized for the parameters considered in this chapter. In particular, we find that the fiducial MegaMapper sky coverage yields constraints that are within 5% of their optimal values.

Finally we consider the impact of a $k_\parallel$ cut, which could be imposed to remove systematic errors, on constraints of local primordial non-Gaussianity. For our constraints we follow the method described in §10.4 with $k_{\rm min} = 0.001\ h\,{\rm Mpc}^{-1}$, and remove modes whose component

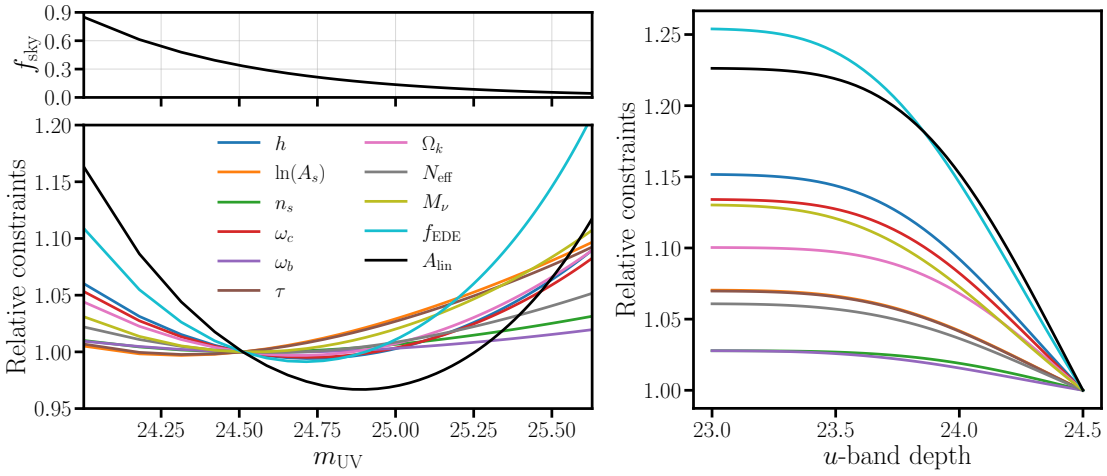


Figure 10.23: *Left:* Relative improvement to Planck + SO + MegaMapper-like parameters as (MegaMapper) area is traded for depth, for fixed observing time. We assume that $m_{\rm UV}(t) = 24.5 + 2.5\log_{10}(\sqrt{t})$ and $f_{\rm sky}(t) = 0.34/t$, and scan through $0.4 < t < 8$. *Right:* Degradation to Planck + SO + MegaMapper-like constraints with a shallower u-band depth. Here the MegaMapper-like survey has a default 24.5 limiting magnitude, and detects 50% of the LBGs for a given $m_{\rm UV}$.

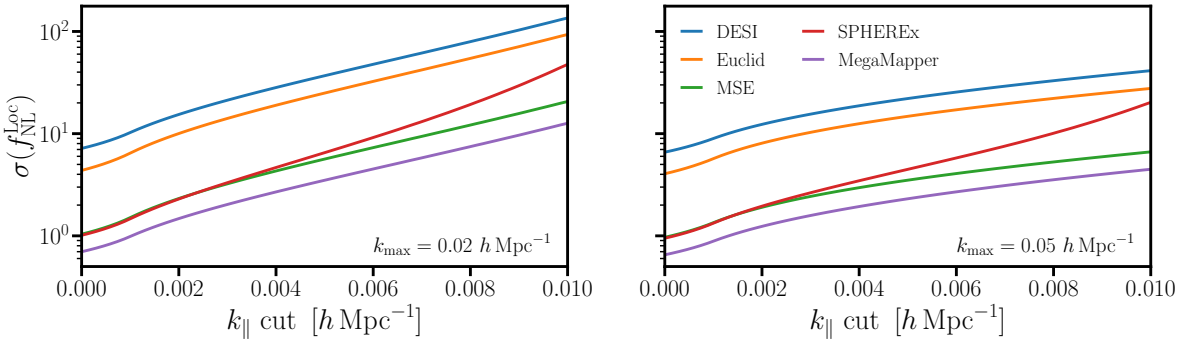


Figure 10.24: Constraints on the amplitude of local primordial non-Gaussianity with a cut on $k_\parallel$ modes. In the left panel we set $k_{\rm max} = 0.02\ h\,{\rm Mpc}^{-1}$, while in the right panel we set $k_{\rm max} = 0.05\ h\,{\rm Mpc}^{-1}$. In both panels we take $k_{\rm min} = 0.001\ h\,{\rm Mpc}^{-1}$

along the line of sight is smaller than some $k_\parallel$ cut, which we vary between 0 and 0.01 $h\,{\rm Mpc}^{-1}$. Our results are illustrated in Fig. 10.24. With a conservative $k_{\rm max} = 0.02\ h\,{\rm Mpc}^{-1}$, we find that MegaMapper constraints degrade from $\sigma(f_{\rm NL}^{\rm Loc}) = 0.7$ to $\sigma(f_{\rm NL}^{\rm Loc}) = 10$ as the $k_\parallel$ cut is increased from 0 to 0.01 $h\,{\rm Mpc}^{-1}$. For a more optimistic $k_{\rm max} = 0.05\ h\,{\rm Mpc}^{-1}$ these constraints degrade from 0.7 to 4. In addition, we find that the constraints from SPHEREx degrade more rapidly with a cut on $k_\parallel$ than for spectroscopic surveys due to its redshift errors

(as in §10.4, we assume $\sigma_z/(1+z) = 0.05$). This emphasizes the importance of spectroscopy for diagnosing systematic errors without blowing up the error on $f_{\rm NL}^{\rm Loc}$.

10.6 Conclusions

Advances in experimental techniques make it possible to map the high redshift Universe at high fidelity in the near future. This will increase the observed volume by many-fold, while providing unprecedented access to very large scales, allowing the reconstruction of an order of magnitude more "primordial modes". Such surveys would improve measurements of the expansion history and spatial curvature of the Universe, primordial non-Gaussianity, features in the power spectrum, models of inflation, dark energy and modified gravity, light relics and neutrino masses. Recently developed theoretical models, well suited to the quasi-linear clustering of the early Universe, provide a firm basis for forecasting the science reach of such surveys including such complexities are non-linear evolution, redshift-space distortions and non-linear bias. In this chapter we have presented such forecasts for a handful of proposed experiments (§10.2) that illustrate the potential of next-generation large-scale structure surveys to inform our theories of fundamental physics and discussed some of the limitations and synergies with other probes. We used a Fisher matrix formalism (§10.3) to forecast errors on cosmological parameters from the experiments individually and in combination with other probes, such as primary CMB anisotropies and CMB lensing. Baryon acoustic oscillations (BAO) are treated as a special case, as are methods for determining the amplitude of clustering using a fixed shape of the power spectrum (e.g. refs. [28, 205]).

Section 10.4 shows that such surveys could map the expansion history and spatial curvature of the Universe, in a relatively model independent fashion and with high precision, over more than 90% of cosmic time ($z < 6$). Fig. 10.10 shows that MegaMapper could provide sub-percent measures of the angular diameter distance to $z \approx 5$ while a PUMA-like experiment could reduce the errors by a further factor of 2 under optimistic assumptions for foregrounds. The line-of-sight distance, i.e. the expansion rate, can be determined with similar accuracy with a PUMA-like experiment providing $< 0.5\%$ precision to $z \approx 6$ for a wide range of foreground assumptions.

Section 10.4 shows that future high redshift surveys will enable percent-level measurements of structure growth out to redshift 5. Fig. 10.11 shows that PUMA could achieve $5-10\%$ measurements of $\sigma_8(z)$ across $2 < z < 6$, while MegaMapper could reduce these errors to less than 1% across $2 < z < 5$. 21-cm surveys are severely limited by a degeneracy with the mean 21-cm brightness temperature, while high redshift spectroscopic surveys benefit greatly from the addition of CMB lensing cross-correlations, which decrease the error on $\sigma_8(z)$ by a factor of 2.

The ΛCDM model with minimal-mass neutrinos provides a highly predictive and rigid model that can explain the majority of cosmological observations [473, 205] and thus makes a natural target for forecasts (§10.4). Table 10.3 and Figs. 10.9 and 10.20 present a summary of our findings. Upcoming surveys like DESI and Euclid, in combination with Planck, will

improve current measurements of ω_c and h by a factor of two, but will largely exhaust the information available from the low redshift power spectrum. Adding high-z redshift measurements allows further improvements by more than a factor of 3, exceeding the reach of even ambitious next-generation CMB experiments. The improvement in the determinant of the Fisher matrix for base ΛCDM in going from PSD to PSD + MegaMapper (PSD + PUMA-32K) is a factor of 2×10^4 (7×10^4).

Massive neutrinos provide one of the best motivated extensions of the minimal ΛCDM model (§10.4), and large-scale structure provides the tightest constraints on the sum of the neutrino masses [646]. Table 10.3 shows that constraints on M_ν tighten significantly with the inclusion of high-redshift LSS data. While the forecasts depend sensitively on the smallest scales that can be used (§10.4; Fig. 10.12) the next-generation of CMB surveys coupled with a next-generation, high-redshift LSS survey should comfortably achieve $\sigma(M_\nu) = 20$ meV. This could be further improved by a better measurement of τ (from the CMB or other means) and the inclusion of CMB lensing, which aids in degeneracy breaking (see Fig. G.3).

Light relics (§10.4) suppress structure on small scales and thus can also be well constrained by high-redshift large-scale structure surveys (Table 10.3 and Fig. 10.13). We find that both MegaMapper and PUMA are nearly cosmic variance limited for this measurement and capable of achieving $\sigma(N_{\rm eff}) \simeq 0.02$ in combination with future CMB experiments (depending upon the $k_{\rm max}$ assumed: Table 10.3 and Fig. 10.13). Such a measurement could detect light vectors at 3σ, regardless of their freeze-out temperature, and would be capable of detecting (or ruling out) all light relics with a freeze-out temperature smaller than the QCD phase transition (~ 0.2 GeV) at 5σ significance. Interestingly we find that a proper marginalization over nuisance parameters (such as non-linear biases and counterterms) is important for forecasting $\sigma(N_{\rm eff})$, with constraints underestimated by $\simeq 40\%$ when these parameters are held fixed (as is often done).

One window into the very early Universe and the physics of inflation is the study of primordial non-Gaussianity (§10.4). The Planck constraints $f_{\rm NL}^{\rm loc} = -0.9 \pm 5.1$ are within a factor of 3 of the cosmic variance limit, so pushing to $\sigma(f_{\rm NL}^{\rm loc}) \lesssim 1$ requires LSS data. We find a high-z LSS survey would be capable of a 5-fold improvement over a cosmic variance limited CMB survey, with $\sigma(f_{\rm NL}^{\rm loc}) \sim 0.8$, just from the power spectrum alone. Further improvements could be possible by including information from the bispectrum or from including CMB lensing.

A second window into the physics of inflation from large-scale structure surveys is to search for primordial features (§10.4) in the power spectrum. The space of possible features is large, so as an example we show that a future LSS survey could constrain the amplitude of a sinusoidal modulation of $P(k)$ to better than 0.1% over a broad range of frequencies (Fig. 10.15). A precise measurement of the power spectrum shape can also be used to place constraints on (or detect) short-term deviations from matter or radiation domination, since such periods will change the shape of the power spectrum due to differential growth of modes. Thus the shape of the matter power spectrum has the potential to detect sub-percent fractional amounts of Early Dark Energy to $z \sim 10^5$ (§10.4), probing Dark Energy all the way to when the Universe was a only a few years old (Fig. 10.18).

The precision of these measurements, combined with CMB observations, also has the promise of tightening constraints on the gravitational slip, which is equal to one in General Relativity but generically departs from unity in other models. For the simplest case of a scale-independent slip, we show that future high redshift experiments could constrain γ to a few percent (Fig. 10.19) with the tightest constraints coming from MegaMapper near $z \approx 2-3$.

We find that inclusion of models that go beyond constant bias and linear theory is important for accurately assessing the potential of future experiments to constrain parameters. Neglecting the complexities of bias or fixing too many nuisance parameters can artificially inflate the precision of parameter measurements by a factor of up to 12 in ΛCDM models (§10.4 and Fig. 10.21).

At higher redshift the non-linear scale shifts to higher k, and the impact of fingers of god on the redshift-space clustering also declines. To make use of the extra information available along the line of sight requires an increase in redshift precision (§10.3; Figs. 10.5 and 10.8). Achieving this precision observationally may be challenging if broad lines, or lines with significant radiative transfer effects, are used. In §10.5 we discuss the impact of redshift uncertainties on a MegaMapper-like survey, and find that our constraints degrade by less than 15% for $\sigma_z/(1+z) < 0.002$, which corresponds to a systematic velocity of 600 km s^{-1}.

Studies of linear or quasi-linear large-scale structure with redshift surveys and the CMB currently provide our tightest constraints on cosmology and fundamental physics. Pushing the redshift and volume frontier will provide guaranteed, significant improvements in the state-of-the-art in a manner that can be reliably forecast and optimized. This chapter presents a first step in this direction, considering a number of proposed experiments that would map the Universe in the epochs before cosmic noon.

Acknowledgments

We would like to thank Arjun Dey and David Schlegel for helpful discussions on redshift errors for high redshift galaxies and Kyle Dawson for conversations and calculations regarding SpecTel. We also thank Shi-Fan Chen, Mikhail Ivanov, Marko Simonovic, Juna Kollmeier, Eric Linder, Emmanuel Schaan, Uroš Seljak, Leonardo Senatore, Anže Slosar and Michael Wilson for useful discussions during the preparation of this manuscript. N.S. thanks Haruki Ebina for pointing out errors in our galaxy−CMB lensing covariance (the $f_{\rm sky}$ factors in Eq. 10.18) in a previous draft. N.S. is supported by the NSF. S.F. is supported by the Physics Division of Lawrence Berkeley National Laboratory. M.W. is supported by the DOE and the NSF. This research has made use of NASA's Astrophysics Data System and the arXiv preprint server. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

Chapter 11

Measuring the optical depth from the epoch of reionization 21cm signal

This chapter was originally published as

> Noah Sailer, Shi-Fan Chen, and Martin White. "Optical depth to reionization from perturbative 21 cm clustering". In: *JCAP* 10 (2022), p. 007. DOI: 10.1088/1475-7516/2022/10/007. arXiv: 2205.11504 [astro-ph.CO]

The optical depth τ is the least well determined parameter in the standard model of cosmology, and one whose precise value is important for both understanding reionization and for inferring fundamental physics from cosmological measurements. We forecast how well future epoch of reionization experiments could constraint τ using a symmetries-based bias expansion that highlights the special role played by anisotropies in the power spectrum on large scales. Given a parametric model for the ionization evolution inspired by the physical behavior of more detailed reionization simulations, we find that future 21cm experiments could place tight constraints on the timing and duration of reionization and hence constraints on τ that are competitive with proposed, space-based CMB missions provided they can measure $k \approx 0.1\,h\,\mathrm{Mpc}^{-1}$ with a clean foreground wedge across redshifts spanning the most active periods of reionization, corresponding to ionization fractions $0.2 \lesssim x \lesssim 0.8$. Significantly improving upon existing CMB-based measurements with next-generation 21cm surveys would require substantially longer observations (~ 5 years) than standard $\mathcal{O}(1000$ hour) integration times. Precise measurements of smaller scales will not improve constraints on τ until a better understanding of the astrophysics of reionization is achieved. In the presence of noise and foregrounds even future 21cm experiments will struggle to constrain τ if the ionization evolution deviates significantly from simple parametric forms.

11.1 Introduction

Within the Standard (ΛCDM) Model of Cosmology the optical depth to reionization τ is the least well-determined of the standard parameters [475]. Though it is an almost purely astrophysical parameter its uncertainty has a large impact on inferences about fundamental physics because, when fitting cosmic microwave background (CMB) data, it is strongly degenerate with parameters like the amplitude of primordial fluctuations (A_s, σ_8) and the sum of the neutrino masses M_ν. In addition to hindering higher-precision measurements of these parameters, τ derives from the Thomson scattering of free electrons and is thus a probe of the history of cosmic reionization [75, 405, 398]. Beyond suppressing the amplitude of CMB temperature anisotropies this Thomson scattering results in a range of physical phenomena in the CMB including large-scale (E-mode) polarization [473] and the kinetic Sunyaev-Zeldovich (kSZ) effect [574, 202, 230, 36] that can be used to measure τ. In addition, probes of the Epoch of Reionization (EoR) have the potential to directly constrain the global ionization history and therefore the abundance of free electrons that source τ [59, 400, 355, 92, 260, 257]. Our goal in this paper is to investigate this final method using recent advances combining cosmological perturbation theory with techniques based on effective field theories.

There have been numerous approaches to modeling the 21cm signal from the EoR, ranging from detailed radiative transfer simulations coupled to hydrodynamics [459, 318] to semi-analytic [428] and hybrid [590] models. The majority are based on N-body simulations. By contrast, our approach to modeling 21cm observations will be based upon cosmological perturbation theory with a symmetries-based bias expansion. This approach is now quite commonplace in large-scale structure studies, but differs in some key respects from the standard approaches within the reionization field (see references above for examples of analytic or semi-numerical models and refs. [291, 399, 489] for notable examples similar in spirit to our approach). No matter how complex the physical processes leading to ionized regions, as long as they obey certain symmetries (e.g. rotational invariance and the equivalence principle) and are quasi-local we can write δ_{21} as an expansion in powers of the density field, tidal field and low-order derivatives [395, 189]. This expansion will be valid on large scales, in our case on scales larger than the typical size of an ionized bubble [399, 489]. The impact of small-scale physics not explicitly included in the model is then handled by correction terms (i.e. counterterms), which are also restricted by symmetry considerations. By marginalizing over the bias parameters and counterterms we can focus on the cosmological information contained in large scale modes while efficiently removing sensitivity to small-scale physics. Since our expansion is based entirely on symmetry, it makes minimal assumptions about that physics. Conversely, information about the sources and morphology of reionization is encapsulated in the small-scale modes and bias coefficients.

Given that perturbation theory and the bias expansion marginalize away much of the "astrophysics" of the 21cm signal it is worth asking whether it is still possible to recover useful information about reionization within such a framework, or whether the study of reionization necessarily requires detailed modeling of the small-scale physics of the EoR (e.g. the composition of ionizing sources). For example, while the overall amplitude of the

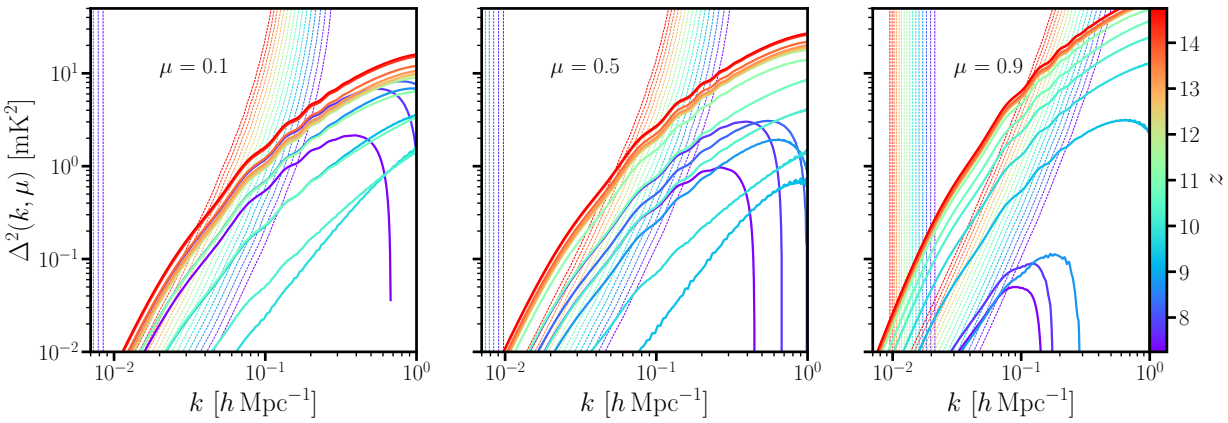


Figure 11.1: The dimensionless 21cm power spectrum ($\Delta^2 = k^3P/2\pi^2$) as a function of redshift for our fiducial model (solid lines, see §11.4) and the thermal noise per mode (steeply rising, dashed lines) for the case of 1000 14m dishes with a 10 year integration time (first column of Table 11.1). The step in thermal noise (vertical lines) at low-k is due to a absence of baselines capable of measuring modes with $k_\perp\chi/2\pi < D/\lambda_{21}^{\rm obs}$, where D is the diameter of the dishes. The panels show $\mu = 0.1$ (left), $\mu = 0.5$ (middle) and $\mu = 0.9$ (right) with colors indicating redshift. Note that the spectrum is quite anisotropic (the amplitudes vary significantly with μ, particularly at lower redshifts), with the shape encoding valuable information about the cosmology and reionization history.

observed 21cm power spectrum is proportional to the brightness temperature $\overline{T}_b$ and can therefore be used to predict the electron fraction (§11.2), even at linear order in perturbation theory the power spectrum will also be proportional to the square of the linear bias b^2 whose value cannot be determined absent detailed knowledge of small-scale behavior of neutral hydrogen, the host halos of ionizing sources etc., leading to an exact degeneracy with the electron fraction.

However, while many potential measurements of the global reionization history, e.g. the overall amplitude of the 21cm signal, are polluted by small-scale physics, the presence of fundamental symmetries implies that others are protected. For example, the equivalence principle dictates that all matter behaves identically under the force of gravity, such that the peculiar velocity of neutral hydrogen on large scales must be the same[1] as that of matter and other tracers like galaxies. Since 21cm surveys infer line-of-sight distances in terms of observed frequencies these peculiar velocities are then encoded in anisotropies in the observed clustering due to redshift-space distortions (RSD), where an object at $\boldsymbol{x}$ appears to be at redshift-space position $\boldsymbol{s} = \boldsymbol{x} + \boldsymbol{u}$, with $\boldsymbol{u} = \hat{\boldsymbol{n}}\ (\hat{\boldsymbol{n}} \cdot \boldsymbol{v}_{\rm pec})/\mathcal{H}$. Within linear theory the redshift-space power spectrum is given by [316]

$$P_{\rm obs}(k) = \overline{T}_b^2(b + f\mu^2)^2 P_{\rm lin}(k), \quad \mu = \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{k}} \tag{11.1}$$

[1]The relative velocity between baryons and dark matter is discussed in §11.5.

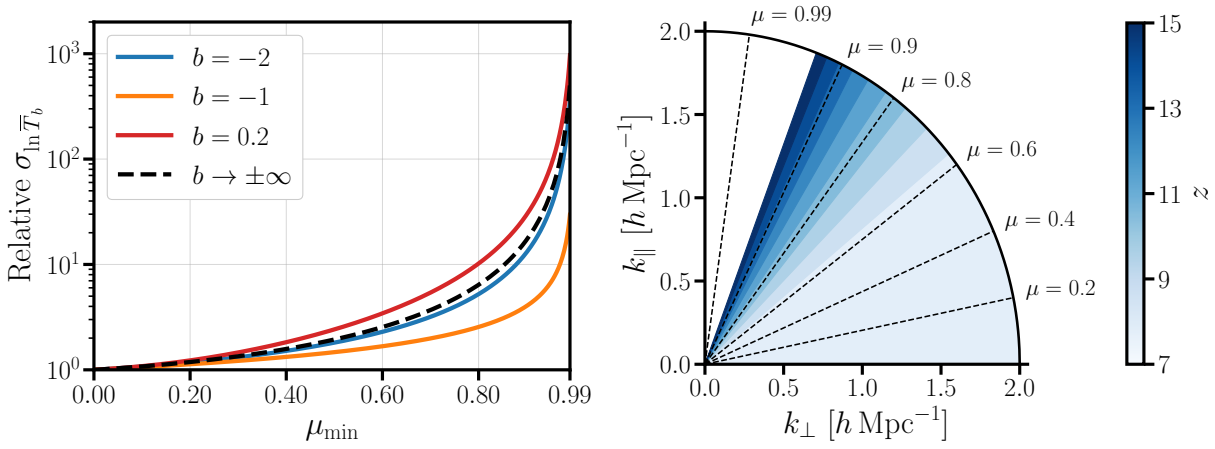


Figure 11.2: *Left:* Error on $\ln \overline{T}_b$ (normalized to $\mu_{\rm min} = 0$) after marginalizing over the linear bias, as a function of $\mu_{\rm min}$, following the simplified scenario outlined in §11.1. The colors denote the assumed fiducial value of the linear bias. *Right:* Foreground wedge for our fiducial scenario ($N_w = 3$; Eq. 11.18) as a function of redshift, assuming 14 m diameter dishes (see §11.4 for a more detailed discussion of the foreground wedge).

where $\hat{\boldsymbol{n}}$ is the line of sight unit vector and $f(z) = d\ln D/d\ln a$ is the growth rate, with $f \simeq 1$ to sub-percent precision at the matter dominated redshifts where reionization occurs (see Fig. 11.1). As the growth rate and $P_{\rm lin}$ are purely a function of cosmology, the size of the anisotropy (μ^2) on large scales is then uncontaminated by small-scale physics and for any given cosmology can be used to infer the brightness temperature $\overline{T}_b$, unlike the isotropic piece proportional to b^2 [59, 400]. In addition to the linear bias, $\overline{T}_b$ is perfectly degenerate in linear theory with the amplitude of matter fluctuations σ_8 within a single redshift bin. This degeneracy can be broken[2] using the different time-evolution of $\sigma_8(z) \propto (1+z)^{-1}$ and $\overline{T}_b(z) \propto (1+z)^{1/2}[1-x(z)]$, or by combing 21cm observations with primary CMB fluctuations (which constrain $\sigma_8 e^{-\tau}$).

The left panel of Figure 11.2 shows our relative ability to measure the brightness temperature as a function of the minimum angular scale $\mu_{\rm min}$ in an idealized setup[3]. Clearly,

[2] At lower redshifts non-linear evolution of large-scale structure can be used to break this degeneracy [124, 512], but at the early times relevant for EoR such effects are very small on the scales that can be reliably modeled.

[3] Specifically, we consider a scenario where the survey is noiseless and the signal is given by linear theory, and the treat the linear power spectrum $P_{\rm lin}$ as fixed. Our ability to constraint $\overline{T}_b$ is then independent of the linear power spectrum shape and the Fisher matrix is given by

$$F_{ij} \propto \int_{\mu_{\rm min}}^{1} d\mu \left[\partial_i (b+\mu^2)^2 \overline{T}_b^2\right]\left[\partial_j (b+\mu^2)^2 \overline{T}_b^2\right] \Big/ (b+\mu^2)^4 \overline{T}_b^4,$$

where $i \in \{\overline{T}_b, b\}$.

since our ability to constrain $\overline{T}_b$ depends on resolving the anisotropy due to RSD, reducing the angular range over which we fit results in dramatic reductions in constraining power (see also ref. [479]). This is particularly relevant for 21cm interferometers wherein chromaticity in instrument leads to foreground leakage into a "wedge" below some $\mu_{\rm min}$ that renders those modes unusable (§11.4). The wedge angle corresponding to our fiducial setup is shown in the right panel of Figure 11.2.

Our goal in this chapter will be to investigate the feasibility of using the EoR 21cm signal in redshift space to constrain the global reionization history within a perturbation theory framework. The remainder of this chapter is organized as follows: Section 11.2 discusses the relationship between the optical depth and reionization as probed by 21cm measurements. This section describes our parameterizations of the ionization fraction evolution. Section 11.3 gives a quick overview of Lagrangian perturbation theory and the symmetries-based bias expansion that we employ throughout this chapter. Section 11.4 discusses the Fisher matrix formalism that we employ in our forecasts, and the fiducial models and scale cuts that we investigate. We conclude in Section 11.5 with a discussion of implications from our forecasts.

11.2 Extracting the optical depth from 21cm clustering

In this section we describe our method for constraining τ from a measurement of 21cm clustering during the reionization era. We start by identifying the physical quantities that we need to constrain in order to infer τ and then describe how these are related to the 21cm power spectrum before discussing how we choose to parameterize the evolving ionization fraction.

The optical depth as a derived parameter

We begin by reviewing how the optical depth can be computed given a cosmological background and a reionization history (i.e. the evolution of the ionization fraction) [355]. The mean CMB optical depth is by definition $\tau = \sigma_T \int_0^{z_{\rm rec}} dz\, \bar{n}_e(z) dr/dz$, where σ_T is the Thompson cross-section, $\bar{n}_e(z)$ is the volume-averaged free electron (proper) number density, $z_{\rm rec} \simeq 1100$ is redshift of last scattering, and dr/dz is the line-of-sight (proper) distance per unit redshift. Throughout we assume that the Universe is flat, in which case $dr/dz = c/(1+z)H(z)$, and use overlines to denote *volume* averages. We assume that free electrons are sourced solely from hydrogen and (singly ionized) helium, neglecting the free electrons from doubly-ionized helium[4] and additional isotopes of H and He, as well as

[4]Ref. [355] performs a back-of-the-envelope calculation to show that doubly-ionized helium contributes to $\sim 1.4\%$ of the total optical depth. Assuming that the best fit $\tau = 0.0561$ from Planck is the truth, this amounts to $\tau_{\rm HeIII} \sim 0.0008$.

heavier elements. We further assume that hydrogen and helium have identical ionization fractions: $x_{\rm HII} \equiv x_{\rm HeII}$. This is a standard approximation in the literature [373, 150, 355], due to the similar ionization potentials (13.6 vs 24.6 eV) of the two elements and the fact that neutral H efficiently shields the He from ionizing radiation. Given these assumptions, the free electron number density can be rewritten as $n_e(\boldsymbol{x}) = x_{\rm HII}(\boldsymbol{x})\rho_b(\boldsymbol{x})/\mu m_{\rm H}$, where ρ_b is the baryon density, $\mu = 1 + (m_{\rm He}/m_{\rm H} - 1)Y_p/4$ is the mean molecular weight and $Y_p = 4n_{\rm He}/n_b$ is the primordial helium abundance. It is standard to rewrite the baryon density as $\rho_b(\boldsymbol{x}) \equiv \overline{\rho}_b(z)[1 + \delta_b(\boldsymbol{x})]$, where $\overline{\rho}_b(z) = 3H_0^2\Omega_b(1+z)^3/8\pi G$ is the mean baryon density, and δ_b is the density contrast. With these definitions and assumptions in hand, the optical depth can be rewritten as [355]

$$\begin{aligned}\tau &= \frac{3H_0^2\Omega_b\sigma_T c}{8\pi G\mu m_{\rm H}}\int_0^{z_{\rm rec}} dz\,\frac{(1+z)^2}{H(z)}\overline{x_{\rm HII}(1+\delta_b)}(z) \\ &= 0.0691\frac{\omega_b}{h\mu}\int_0^{z_{\rm rec}} dz\,\frac{(1+z)^2}{E(z)}x(z),\end{aligned} \tag{11.2}$$

where $\omega_b = \Omega_b h^2$, $H_0 \equiv 100\,h\,$km/s/Mpc, $E(z) = H(z)/H_0$, and we've defined the *mass-weighted* average ionization fraction $x(z) \equiv \overline{x_{\rm HII}(1+\delta_b)}(z)$. In going from the first to second line of Eq. (11.2) we have substituted the values of quantities which are known to well beyond the precision of our τ measurement. In particular we see that τ may be derived from the baryon fraction ω_b, the expansion history $E(z)$, the primordial helium abundance $\mu = 1 + 0.7428\,Y_p$, and the "reionization history" $x(z)$.

Ionization fraction from 21cm clustering

The 21cm brightness temperature, defined as the difference between the observed 21cm brightness and that of the CMB, at redshift-space position $\boldsymbol{s}$ is given by [383, 239]:

$$\begin{aligned}T_b(\boldsymbol{s}, z) &= \frac{3}{16}\frac{\hbar c^3 A_{10}}{k_B\nu_{21}^2}\frac{[1 - x_{\rm HII}(\boldsymbol{x}, z)]n_{\rm H}(\boldsymbol{x}, z)}{(1+z)(dv_\parallel/dr_\parallel)}\left[1 - \frac{T_{\rm CMB}(z)}{T_s(\boldsymbol{x}, z)}\right] \\ &\equiv \overline{T}_b(z)[1 + \delta_{21}(\boldsymbol{s}, z)]\end{aligned} \tag{11.3}$$

where $\boldsymbol{x}$ is the corresponding position in real space, $\langle\delta_{21}\rangle = 0$, $A_{10} = 2.85\times10^{-15}$ Hz is the spontaneous emission coefficient of the 21cm transition, $\nu_{21} = 1.42$ GHz is the rest frame 21cm transition frequency and T_s is the spin temperature, defined so that the ratio of the number of excited (spin aligned) hydrogen atoms to the number of hydrogen atoms in the ground state (spins anti-aligned) is given by $3\exp[-2\pi\hbar\nu_{21}/k_B T_s]$[5]. The derivative $dv_\parallel/dr_\parallel$ is the (proper) gradient of the (proper) velocity along the line of sight and sources the

[5]This expression follows from the more general form

$$T_b(\mathbf{s}, z) = \frac{T_s - T_{\rm CMB}}{1+z}(1 - e^{-\tau_{21}(\mathbf{s})}) \tag{11.4}$$

redshift space anisotropy in the observed signal; dividing a real-space density by this factor yields the redshift-space density. We assume that the spin temperature dominates over the CMB ($T_s \gg T_{\rm CMB}$), and discuss the limitations of this approximation in §11.4. Under these assumptions, the mean 21cm brightness temperature takes the form:

$$\overline{T}_b(z) = 189.1[1 - x(z)]\frac{(1+z)^2}{E(z)}\ \frac{\omega_b}{h\mu}\left(1 - \frac{Y_p}{4}\right) \quad \text{mK}, \tag{11.5}$$

where we have substituted the values of known quantities.

A 21cm interferometer directly measures the Fourier modes $T_b(\boldsymbol{k})$, from which one measures the power spectrum:

$$P_{\rm obs}(\boldsymbol{k}, z) = \overline{T}_b^2(z) P_{21}(\boldsymbol{k}, z) + P_N(\boldsymbol{k}, z), \tag{11.6}$$

where $\langle \delta_{21}(\boldsymbol{k}) \delta_{21}(\boldsymbol{k}') \rangle = (2\pi)^3 \delta^D(\boldsymbol{k} + \boldsymbol{k}') P_{21}(\boldsymbol{k})$ and $P_N(\boldsymbol{k}, z)$ is the thermal noise of the instrument. In particular, we note that the mean 21cm brightness temperature controls the overall amplitude of the observed (redshift space) power spectrum of the 21cm signal. This mean temperature is in turn proportional to the neutral fraction $1 - x$. Thus, by measuring the amplitude of the 21cm power spectrum one can in turn infer a measurement of the ionization fraction x.

Parametric reionization history

The presence of a series of acoustic peaks in the angular power spectra of the CMB indicates that the Universe was dense and ionized at early times and then underwent a rapid transition to being (largely) neutral at $z \simeq 1100$. This neutral period lasted for a significant time before the first astrophysical sources (in currently popular models, massive stars in relatively low mass early galaxies) reionized the Universe. Lower limits on the redshift of reionization come from the lack of complete absorption of UV photons from high redshift QSOs [262]. While significant questions remain, there are numerous observational constraints on the process of reionization [75, 405, 398]. Current observations point towards a "late and fast" reionization period, though with considerable uncertainty [644].

Rather than use a simulation-based model to infer the evolution of the ionization fraction from a set of physical parameters (star formation rates, escape fractions, baryon content, etc.), we will instead take a more astrophysics-agnostic approach and adopt an analytic parameterization of $x(z)$. Ref. [590] found that the Weibull function, which has three free parameters:

$$x(z) = \text{Exp}\left[-\max\left(\frac{z-a}{b}, 0\right)^c\right], \tag{11.7}$$

where τ_{21} is the 21cm optical depth, which in the small τ_{21} limit gives $1 - e^{-\tau_{21}} \approx \tau_{21}$. Since the 21cm signal comes from the diffuse, optically-thin IGM this approximation is expected to be a good one, but it is worth noting that in the opposite, optically thick limit the nonlinearity of the exponential mapping tends to contaminate the linear redshift-space signal critical to obtaining τ constraints. This is due to the fact that the exponential mapping breaks the number conservation symmetry that protects the Kaiser form of RSD [138].

is capable of accurately fitting a wide range of reionization scenarios. We will adopt this Weibull form as the default model in our forecasts. A comparison of the best-fit Weibull function with our fiducial reionization history is shown in the left panel of Fig. 11.3.

In §11.4 we will consider how freeing up the functional form of $x(z)$ impacts our constraints. To examine this we will model $x(z)$ with Lagrange interpolating functions, as proposed by ref. [588]. The parameters of the model are a set of mass-weighted ionization fractions $\{x(z_i)\}$ for a chosen (fixed) set of redshifts $\{z_i\}$. These points are smoothly connected via:

$$x(z) = \min\left[\exp\left(\sum_{i=1}^{N} p_i(z)\right), 1\right] \quad \text{where} \quad p_i(z) \equiv \ln(x(z_i)) \prod_{j\neq i} \frac{\ln(1+z) - \ln(1+z_j)}{\ln(1+z_i) - \ln(1+z_j)}. \tag{11.8}$$

We consider how our constraints vary as a function of the number of free parameters in §11.4.

11.3 The 21cm clustering signal in Lagrangian perturbation theory

To robustly account for scale-dependent bias and structure formation in the quasi-linear regime we will adopt the formalism of Lagrangian Perturbation Theory (LPT). The literature on perturbative modeling of large-scale structure is vast and we will not try to give a comprehensive overview of the subject or rederive results; rather, our goal is to give a reasonably self-contained summary and highlight important physical points distinguishing the 21cm signal from the more standard treatment of galaxies. A review of perturbation theory can be found in refs. [82, 194], and of bias in ref. [189].

Within LPT cosmological structure formation is modeled as the gravitational dynamics of an effective fluid whose short-wavelength (UV) physics are integrated out, with corresponding effects on large-scale structure parametrized by effective-theory counterterms and stocastic contributions [607, 482]. Unlike in Eulerian effective theories of large-scale structure (see e.g. [67, 118]), which model the resulting fluid via its densities and velocity potentials, within the Lagrangian picture of fluid dynamics one traces the trajectories of infinitesimal fluid elements labeled by their initial (Lagrangian) positions $\boldsymbol{q}$. The dynamics are encoded in a displacement vector $\boldsymbol{\Psi}(\boldsymbol{q}, \eta)$ defined such that the Eulerian (comoving) positions $\boldsymbol{x}$ of the fluid elements at a conformal time η is $\boldsymbol{x}(\boldsymbol{q}, \eta) = \boldsymbol{q} + \boldsymbol{\Psi}(\boldsymbol{q}, \eta)$, and where the displacement vector is sourced by the gravitational potential to follow $\ddot{\boldsymbol{\Psi}}(\boldsymbol{q}) + \mathcal{H}\dot{\boldsymbol{\Psi}}(\boldsymbol{q}) = -\nabla_{\boldsymbol{x}}\Phi(\boldsymbol{x})$, with $\mathcal{H}$ defined as the conformal Hubble parameter. The gravitational potential in turn follows Poisson's equation, with the matter (over)density obtained via number conservation:

$$1 + \delta_{\rm m}(\boldsymbol{x}, \eta) = \int d^3\boldsymbol{q}\ \delta^D(\boldsymbol{x} - \boldsymbol{q} - \boldsymbol{\Psi}(\boldsymbol{q}, \eta)). \tag{11.9}$$

The velocity field is given straightforwardly as the time derivative of the displacement.

In general, cosmological observables like the 21cm signal do not directly probe the matter field described by the displacements $\boldsymbol{\Psi}$ but rather are determined by complex small-scale astrophysical phenomena and their response to the large-scale matter distribution over cosmic time [189]. For example, whether a galaxy exists at a spacetime point $(\boldsymbol{x}, \eta)$ will generically depend on the matter density, velocity gradients and tidal fields in a neighborhood about its trajectory $\boldsymbol{q} + \boldsymbol{\Psi}(\boldsymbol{q}, \eta)$ approximately the size of its Lagrangian halo radius R_h. While the size of the nonlocality in the galaxy bias expansion is governed primarily by R_h (since the dominant physical effect is the nonlinear gravitational collapse of the Lagrangian patch into a halo), the 21cm signal during reionization is sensitive to other, stronger nonlocalities such as the mean free path $\lambda_{\rm mfp}$ of ionizing photons, which unlike nonrelativistic matter can travel significant ($\sim H^{-1}$) distances in a Hubble time in the absence of absorption. In this case the size of nonlocalities is roughly set by the size of the ionized bubbles during reionization, with smaller bubbles earlier on in the process corresponding to our ability to perturbatively model the signal over a wider range of scales [399].

Within the (Lagrangian) effective-theory framework [395, 543, 47, 606, 136] we can write the quasi-local dependence of the 21cm signal at $\boldsymbol{x}$ on neighboring points along its trajectory — that is, nearby in $\boldsymbol{q}$ — as an expansion order-by-order in the initial conditions at $\boldsymbol{q}$

$$F(\boldsymbol{q}) = 1 + b_1^{\rm L}\delta(\boldsymbol{q}) + \frac{1}{2}b_2^{\rm L}\big(\delta^2(\boldsymbol{q}) - \langle\delta^2\rangle\big) + b_s^{\rm L}\big(s^2(\boldsymbol{q}) - \langle s^2\rangle\big) + \cdots \\ + R_*^2\nabla^2\delta(\boldsymbol{q}) + \cdots + \epsilon(\boldsymbol{q}). \tag{11.10}$$

This functional is then dynamically advected to the real-space position $\boldsymbol{x}$ [387]:

$$1 + \delta_{21}(\boldsymbol{x}) = \int d^3\boldsymbol{q}\, F(\boldsymbol{q})\, \delta^D(\boldsymbol{x} - \boldsymbol{q} - \boldsymbol{\Psi}(\boldsymbol{q})) \tag{11.11}$$

in the same way as the matter field (Eq. 11.9).[6] In this way the 21cm clustering signal is captured in LPT via a combination of dynamics ($\boldsymbol{\Psi}$) and dependence on the initial conditions $F(\boldsymbol{q})$.

The form of the functional $F(\boldsymbol{q})$ is restricted by symmetries and the equivalence principle to depend only on scalar combinations of the density, velocity gradients and the tidal field. Within perturbation theory these quantities can all be written in terms of the initial density field, i.e. to linear order scalars formed from $\partial_i v_j$ and $\partial_i\partial_j\Phi \sim (k_i k_j/k^2)\delta$. To second order, contributions to the bias expansion can be summarized by its dependence on δ^2 and $s^2 = s_{ij}s_{ij}$, where $s_{ij} \equiv (\partial_i\partial_i/\partial^2 - \delta_{ij}/3)\,\delta$ is the shear field. Beyond the local value of these initial conditions along the trajectory starting at $\boldsymbol{q}$, the existence of nonlocalities due to e.g. ionized bubbles can be handled by including derivative corrections such as $\nabla^2\delta$, i.e. performing a

[6] Fluctuations in the 21cm signal can be sourced by fluctuations in either the baryon density δ_b or the electron fraction δ_x and treating these separately is common in 21cm studies. However, since it can equally be thought of as fluctuations in the density of neutral hydrogen, in this chapter we directly model this product using a single bias expansion instead of separately. This avoids needing to model correlated fluctuations in δ_b and δ_x when evaluating $\delta_b\delta_x$, which in perturbation theory language is known as a "composite operator".

Taylor series expansion of δ around $\boldsymbol{q}$. These corrections become increasingly prominent as we probe scales closer to the bubble size. Depending on how small in scale we need to fit to it may be necessary to include terms higher order in density and derivatives, though it is important to stress that our perturbative treatment of $F(\boldsymbol{q})$ as a quasi-local functional will necessarily break down on scales smaller than the bubble size where the physics of reionization are nonlocal.

The last term in Eq. 11.10 is the residual, $\epsilon(\boldsymbol{q})$, which represents stochastic small-scale contributions to the clustering of the 21cm signal uncorrelated with the large-scale modes we will be interested in. To lowest order these modes only correlate in close proximity $\langle\epsilon(\boldsymbol{q})\epsilon(\boldsymbol{q}')\rangle \sim \delta^D(\boldsymbol{q}-\boldsymbol{q}')$ or, equivalently, have a flat spectrum in Fourier space: $P_\epsilon(k) = \text{const.}$

Finally, we need to account for contributions to the *observed* densities due to peculiar velocities in the form of redshift-space distortions (RSD). Specifically, since line-of-sight distances are inferred from redshifts they incur an additional Doppler contribution from the peculiar velocity $\boldsymbol{u} = \hat{\boldsymbol{n}}\,(\hat{\boldsymbol{n}}\cdot\boldsymbol{v}_{\rm pec})/\mathcal{H}$, such that the 21cm signal in redshift-space is given by [388, 610, 134]

$$1+\delta_{21}(\boldsymbol{s}) = \int d^3\boldsymbol{q}\, F(\boldsymbol{q})\,\delta^D(\boldsymbol{s}-\boldsymbol{q}-\boldsymbol{\Psi}_s(\boldsymbol{q})) \quad ; \quad \boldsymbol{\Psi}_s(\boldsymbol{q}) = \boldsymbol{\Psi}+\boldsymbol{u}. \tag{11.12}$$

In the limit where the real-to-redshift-space mapping is unique ($\boldsymbol{s}\leftrightarrow\boldsymbol{x}\leftrightarrow\boldsymbol{q}$) we can integrate over the delta function to recover the familiar expression $\delta(\boldsymbol{s}) = aH/(dv_\parallel/dr)\delta(\boldsymbol{x})$, where $v_\parallel$ is the line of sight velocity including Hubble flow [239]. Within the Einstein-de Sitter approximation, which is an excellent approximation at the redshifts where reionization occurs, the $m^{\rm th}$ order displacement obeys the simple relation

$$\boldsymbol{\Psi}_s^{(m)}(\boldsymbol{q},\eta) = \boldsymbol{\Psi}^{(m)} + mf\hat{\boldsymbol{n}}\,\left(\hat{\boldsymbol{n}}\cdot\boldsymbol{\Psi}^{(m)}\right). \tag{11.13}$$

with $f\simeq 1$, though it is important to note that, like the real-space densities, the velocities too require effective-theory corrections to properly account for their dependence on small-scale physics [136]. To lowest order, the above calculation recovers the familiar Kaiser [316] form:

$$1+\delta_{21} = \frac{1+b_1^L\delta(\boldsymbol{q})}{1+(aH)^{-1}\,dv_\parallel^{\rm pec}/dr}$$

and, using that $v_{{\rm pec},i}/(aH) = -f\nabla_i^{-1}\delta$ we have in Fourier space

$$\delta_{21}(\boldsymbol{k}) = \left([1+b_1^L]+f\mu^2\right)\delta(\boldsymbol{k}), \quad \mu = \hat{\boldsymbol{n}}\cdot\hat{k}$$

such that the linear theory power spectrum is given by $(b+f\mu^2)^2P_{\rm lin}(k)$ [316], where $b = 1+b_1^L$. This is the familiar result used in the introduction.

11.4 Forecasts

Now we turn to our forecasts of how well a (future) 21cm interferometer could determine τ by measuring the evolving 21cm power spectrum. Our forecasts use the Fisher matrix

formalism, with an extensive set of nuisance terms to handle the complexities of the reionization process. These forecasts allow us to determine the key observations that would allow a robust extraction of τ and highlight critical assumptions. We begin by defining our fiducial cosmology, reionization history and 21cm clustering model before describing the range of scales and redshifts used in our fits.

Model and fiducial scenario

In this chapter we use `CLASS` [97] to compute the linear CDM+baryon power spectrum, and `velocileptors`[7] [136] to model non-linearities and redshift-space distortions (see §11.3) to one-loop order within LPT:

$$P_{21}(\boldsymbol{k}) = P_{21}^{\text{1-loop}}(\boldsymbol{k}) + \left(\alpha_0 + \alpha_2\mu^2 + \alpha_4\mu^4\right)k^2 P_{cb}^{\text{Zel}}(\boldsymbol{k}) + N_0 + N_2(\mu k)^2 + N_4(\mu k)^4, \quad (11.14)$$

where N_2 and N_4 encode the effects of small-scale velocities, N_0 is the shot noise, P_{cb}^{Zel} is the CDM+baryon power spectrum in the Zel'dovich approximation, and $\mu = \hat{\boldsymbol{k}} \cdot \hat{\boldsymbol{n}}$ is the cosine of the angle between the wavevector and the line of sight. This model includes a bias expansion to quadratic order in the linear CDM+baryon density contrast (i.e. the first line of Eq. 11.10). For further details on the calculations see refs. [136, 134].

We use the `FishLSS`[8] [512] code to compute the Fisher information for the full-shape 21cm power spectrum $P_{\text{obs}}(\boldsymbol{k})$ in the basis:

$$(A_s, n_s, h, \omega_c, \omega_b, M_\nu) \cup \{\ln \overline{T}_b, b, b_2, b_s, \alpha_0, \alpha_2, \alpha_4, N_0, N_2, N_4\}_z, \quad (11.15)$$

where M_ν is the sum of the neutrino masses, which we hold fixed unless otherwise stated, and the biases here (b, b_2, b_s) are in the Eulerian frame. We treat all time-dependent nuisance terms (the right-hand set in Eq. 11.15; the z subscript indicates time-dependence) in different redshift bins as separate, independent variables which are individually varied. We direct the reader to ref. [512] for additional details regarding the numerical computation of derivatives and evaluation of the full shape Fisher matrix, as well as information regarding the CMB Fisher matrices used below. Unless otherwise stated we will include a prior on ΛCDM+M_ν from Planck primary CMB measurements in all of our forecasts.

As discussed in §11.2 we will parameterize the reionization history with set of parameters $\{x_i\}$. Using this parametrization we can trade the free parameters $\{\ln \overline{T}_b\}_z$ for the set $\{x_i\}$ via the chain rule: $\partial_{x_i} P_{\text{obs}}(z_j) = (\partial_{x_i} \ln \overline{T}_b)\partial_{\ln \overline{T}_b} P_{\text{obs}}(z_j)$, since the value of the ionization fraction has little impact on the evolution of large-scale matter fluctuations and primarily impacts P_{obs} through an amplitude change. In our fiducial scenario we use the Weilbull form ($\{x_i\} = \{a, b, c\}$), in which case our fiducial full set of free parameters is given by

$$\underbrace{(A_s, n_s, h, \omega_c, \omega_b, a, b, c)}_{\equiv \boldsymbol{\theta}} \cup \{b, b_2, b_s, \alpha_0, \alpha_2, \alpha_4, N_0, N_2, N_4\}_z. \quad (11.16)$$

[7] https://github.com/sfschen/velocileptors
[8] https://github.com/NoahSailer/FishLSS

In this scenario our total number of free parameters is $8 + 9 \times$ (number of redshift bins). Once we have the Fisher matrix in the basis Eq. (11.16), we then average over the nuisance parameters $\{b, b_2, b_s, \alpha_0, \alpha_2, \alpha_4, N_0, N_2, N_4\}_z$ to obtain the nuisance-averaged 8×8 Fisher matrix $\boldsymbol{F}$, and use Eq. (11.2) to propagate the errors on $\boldsymbol{\theta}$ to the optical depth: $\sigma_\tau^2 = \partial_{\boldsymbol{\theta}}\tau \cdot \boldsymbol{F}^{-1} \cdot \partial_{\boldsymbol{\theta}}\tau$.

We base our fiducial reionization history off of the SCORCH [199] simulations, which generalize the radiation-hydrodynamic code `Radhydro` [591, 589] to include varying radiation escape fractions: $f_{\rm esc}(z) = f_8[(1+z)/9]^{a_8}$. We choose to adopt the redshift-independent scenario ($a_8 = 0$) as our fiducial model for $x(z)$, which is shown in the left panel of Fig. 11.3. This scenario predicts $\tau = 0.0639$ when all of the cosmological parameters are fixed at their best-fit Planck values [473], which is within 1.1σ of the Planck measurement $\tau = 0.0561 \pm 0.0071$. In Fig. 11.3 we also show the best fit Weibull function ($a = 6.62$, $b = 1.85$, $c = 1.14$), which agrees with the simulation output to within $|\delta x| \leq 0.02$, resulting in a very small systematic shift in the optical depth: $|\delta\tau| = 0.0005$. Additionally Fig. 11.3 shows the cumulative contributions to τ from redshifts $< z$ (normalized to the total τ). Note that ≈ 0.5 of the total τ comes from redshifts below 5 where the Universe is fully ionized (and hence has no uncertainty from ignorance of x). Most of the rest comes from $z < 11$.

For the clustering we adopt the evolution of b measured from simulations[9] in ref. [399] and set our fiducial counterterm $\alpha_0(z)$ from the effective bubble radius evolution $R_{\rm eff}(z)$ in the same work (Fig. 11.3). These biases are consistent with those measured by ref. [489], however other theoretical calculations have found different values [632]. The effective bubble radius[10] is defined such that $\delta_{21} \supset b(1 - R_{\rm eff}^2 k^2/3)\delta$, so that the power spectrum $P_{21} \supset -2b^2 R_{\rm eff}^2 k^2 P/3 + \cdots$, and so we identify $\alpha_0 = -2b^2 R_{\rm eff}^2/3$. For the higher order biases, b_2 and b_s, we assume vanishing Lagrangian bias corresponding to the coevolution values for Eulerian biases and we further take $N_0 = N_2 = N_4 = \alpha_2 = \alpha_4 = 0$. Since these terms are relevant primarily at small scales and in the absence of receiver noise, the precise values are not important. We have explicitly checked that changing the fiducial value of b_2 from its coevolution value to $b_2(z) = -2$, a value that roughly agrees with that measured by ref. [399] from simulations, impacts our forecasts (for the idealized scenario of §11.4) by less than 7%.

In Fig. 11.3 we note that the linear bias starts off positive and eventually becomes negative while $R_{\rm eff}$ grows towards the end of reionization. The growth of the bubbles is self-explanatory. Ref. [399] gives a simple explanation for the evolution of the bias assuming that the ionized bubbles trace the locations of the sources in proportion to their ionizing luminosity and that recombinations and light travel delays can be neglected. In this model a large region containing more sources will have more ionization so $\delta_{\rm HII} = \delta_S + \delta$, where δ_S is the source overdensity field. Writing $\delta = (1-x)\delta_{\rm HI} + x\delta_{\rm HII}$ we note that the first term on

[9] These values are taken from the left-hand panel of Fig. 7. We fit a Weibull function to the reionization history of ref. [399], finding $(a, b, c) = (6.73, 1.0, 0.92)$, and use this to appropriately translate these curves to our fiducial reionization history (assuming that $b(x)$ and $R_{\rm eff}^2(x)$ are fixed).

[10] The measured values of $R_{\rm eff}^2$ are poorly constrained at early times [399], so the negative values of $R_{\rm eff}^2$ may just be a numerical artifact. This has a negligible impact on our constraints since in this regime the hard $k_{\rm max}^h$ cut dominates over the bubble size such that $\alpha_0 k^2 \ll 1$.

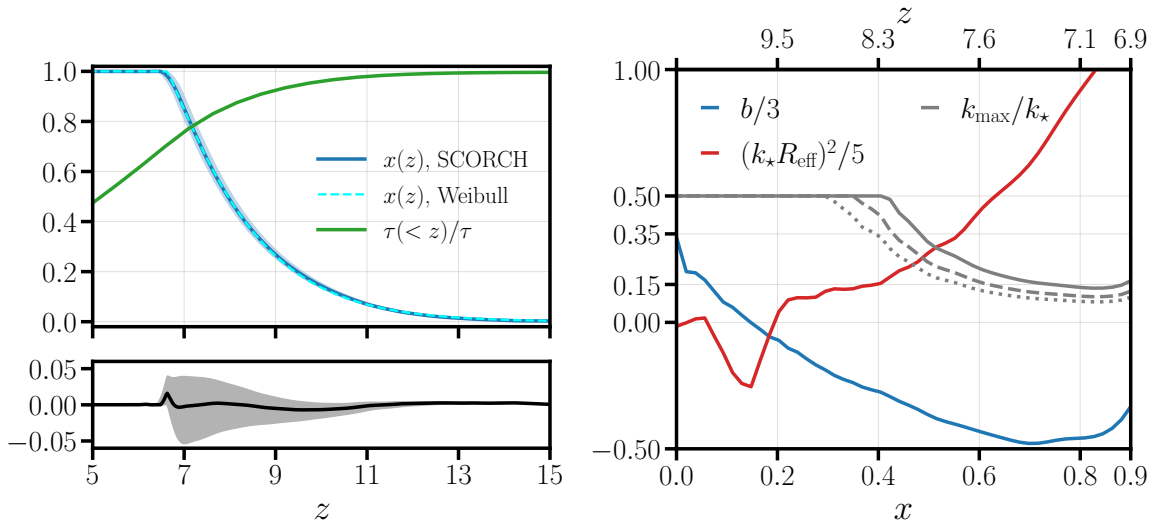


Figure 11.3: *Left*: In the top panel we show our fiducial reionization history (blue) and the best fit Weibull function (cyan). The bottom panel shows the residual $|\delta x| < 0.02$. This residual results in a systematic error of $|\delta\tau| = 0.0005$. The edges of shaded bands represent the δz by which one can shift the reionization history and change the optical depth by $|\delta\tau| = 0.001$. In green we show the (relative) cumulative contribution to the optical depth for redshifts $< z$. *Right*: The fiducial linear Eulerian bias b, the effective bubble size $R_{\rm eff}$ and a representative evolution of $k_{\rm max}$ (see Eq. (11.17)), where $k_\star = 1\,h\,{\rm Mpc}^{-1}$. The grey solid, dashed and dotted lines correspond to $X_\alpha = 3$, 4, 5 respectively. Such limits ensure that for $k < k_{\rm max}$ the corrections are perturbative and the higher-order counterterms are negligible.

the rhs is δ_{21} while the last is $x(\delta + \delta_S)$ and thus $\delta_{21} = (1-x)\delta - x\delta_S = (1 - x[1 + b_S])\delta$. Assuming the sources of reionization are highly clustered we expect b to be negative once reionization gets underway.

Idealistic survey

We first examine an idealized noiseless survey which covers half the sky over the redshift range $7 < z < 15$, which we will split into 16 redshift bins[11] with $\Delta z = 0.5$. To investigate the sensitivity of these measurements in the $\boldsymbol{k} - z$ parameter space, we will first consider redshift-independent scale cuts, with our fiducial cuts being $k_\parallel^{\rm min} = 0.05\ h\,{\rm Mpc}^{-1}$, $k_{\rm max} = 0.5\ h\,{\rm Mpc}^{-1}$ and $\mu_{\rm min} = 0.7$. A more realistic treatment which accounts for the redshift dependence of the foreground wedge and $k_{\rm max}$ will be considered in §11.4. For this idealized scenario, and using the best-fit Weibull function to parameterize the reionization history, we find a fiducial

[11] We have checked that doubling the number of redshifs bins (so that $\Delta z = 0.25$) does not significantly impact our forecasts, suggesting that $\Delta z = 0.5$ is sufficient for convergence.

error on the optical depth $\sigma_\tau = 0.00016$ with a Planck prior on ΛCDM, and 0.00026 without (21cm clustering only). This illustrates that if reionization proceeds smoothly in redshift and can be parameterized by a few parameters then there is more than enough signal in the clustering to infer τ with very high accuracy.

Exploring this further, Fig. 11.4 shows how the errors depend upon the range of scales and redshifts probed. As Fig. 11.4 makes clear, the constraints are more sensitive to an experiment's ability to resolve large scales (i.e. low k modes) and obtain sufficient dynamic range in μ than to the range of redshifts it probes or its ability to probe very small scales. Among other things this suggests that foreground subtraction and instrument calibration will matter as much for measuring τ as raw instrumental sensitivity. Quantitatively we see that the constraints on τ degrade significantly if the experiment cannot measure $k < 0.2\,h\,\mathrm{Mpc}^{-1}$, but it is only necessary to measure well modes up to $k \simeq 0.2\,h\,\mathrm{Mpc}^{-1}$ to get almost all of the constraining power in the data. This is because the smaller-scale modes mostly serve to constrain astrophysical parameters. All of these modes are linear, from the perspective of dynamics.

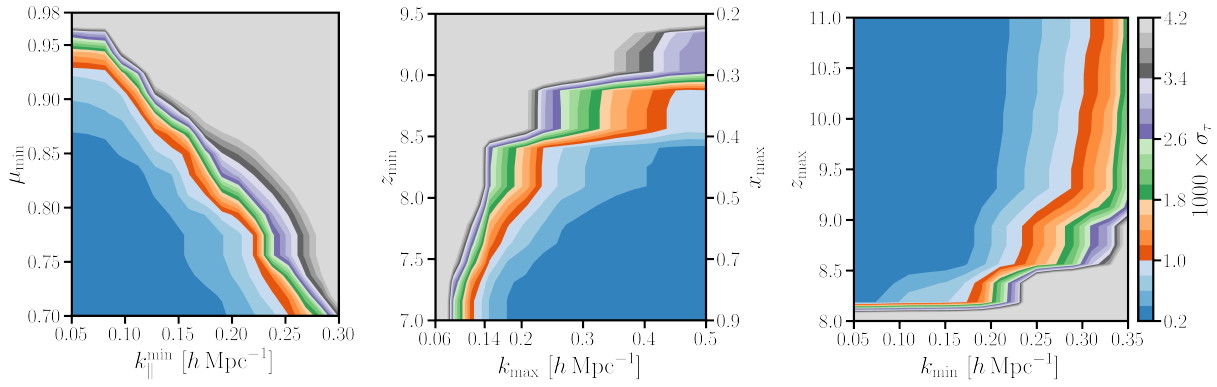


Figure 11.4: Errors on the optical depth (times a thousand) for a noiseless survey covering $7 < z < 15$ over half the sky. The survey has been split into 16 redshift bins with $\Delta z = 0.5$. The fiducial scale cuts are $k_{\max} = 0.5\ h\,\mathrm{Mpc}^{-1}$, $k_\parallel^{\min} = 0.05\ h\,\mathrm{Mpc}^{-1}$ and $\mu_{\min} = 0.7$, and we include a prior on ΛCDM from Planck primary CMB measurements.

Much of the integral responsible for τ comes from 'intermediate' redshifts, so the survey does not need to probe very high or low redshifts in order to constrain τ very well. This is also illustrated in Fig. 11.4 where we see that the errors on τ remain small for $z_{\min} < 8$ and $z_{\max} > 9.5$ (in our fiducial reionization history) corresponding to coverage of ionization fractions $0.2 \lesssim x(z) \lesssim 0.5$. Only once the survey is unable to probe those epochs does the inferred error on τ rise dramatically. Varying the Weibull parameter b, which affects the duration of reionization, in our fiducial model within upper limits on the duration of reionization from the CMB ($b < 2.8$) changes our forecasts for σ_τ by $< 25\%$. We expect that $T_S \gg T_{\mathrm{cmb}}$ for $z \lesssim 10$ [75, 405, 398], justifying our neglect of spin temperature fluctuations in our fiducial model. If reionization is particularly extended or has a non-trivial high z

component this assumption would need to be revisited. Such scenarios are disfavored, but if the spin temperature can fluctuate while x is significantly different from zero then the connection between 21cm clustering and τ would be much more complex.

Finally, to investigate how sensitive our constraints are to the number of free variables in our parameterization of the reionization history, we consider swapping the Weibull function for a Lagrangian interpolating function (see §11.2) with N free parameters, which we choose to be the ionization fractions $\{x(z_i)\}$ at redshifts $z_i = 7, 7+6/N, \cdots, 13$. For $N = 3$, 5, 7, 9 we find $1000 \times \sigma_\tau = 0.20$, 0.25, 1.95, 4.47 for our set of fiducial scale cuts. The scaling of the error with N gets significantly worse with more realistic scale and temporal cuts: for example, after enforcing $z_{\rm max} = 10$ and adding a redshift-dependent foreground wedge ($N_w = 3$, see §11.4) the error increases from 0.4 to 6.1 for $N = 3$, 5. It is worth noting that not all of the loss of constraining power with increasing N is necessarily physical: astrophysical processes contributing to reionization typically evolve on long (fraction of Hubble) time scales, leading to constraints on the variation of $x(z)$ between interpolating points. Introducing physically motivated smoothness priors to parametrically fit e.g. the expansion history have recently been studied in ref. [495], and we leave it to future work to develop these techniques in the context of reionization.

More realistic scenarios

Having determined what range of scales and redshifts are most important for constraining τ we now consider more realistic scenarios in which we add instrument noise, foregrounds and redshift evolution. Our purpose is not to model any particular 21cm survey in detail, but rather to explore the impact that various observational choices have on the inferred constraints on τ. The current generation of 21cm interferometers does not have sufficient sensitivity to improve upon existing measurements of τ, so we shall consider hypothetical next-generation instruments with integration times on the order of years, consistent with proposed post-reionization 21cm surveys [42], as opposed to the more standard 1000 hours[12].

Our fiducial instrument will be a hexagonally packed array of 10^3, non-tracking, 14 m dishes. This is a scaled up version of the Hydrogen Epoch of Reionization Array (HERA; [174]). We compute the baseline distribution using `21cmSense`[13] [481, 480]. The thermal noise power is based on Appendix D of ref. [558]. At high redshifts we are sky noise dominated and since most of our constraining power on τ comes from intermediate k the baseline distribution is of secondary importance and we are mostly sensitive to the total collecting area (and the integration time). HERA frequencies range from 50 to 250 MHz, which translates to $4.8 < z < 27.5$ and we shall assume a future array would cover a comparable frequency

[12] Reducing the integration time for our fiducial SKA-like survey (see the SKA-like column in Table 11.1) from 5 years to 1000 hours increases the $1\,\sigma$ error on τ from 0.0009 to 0.0067, which is a marginal improvement over Planck alone ($\sigma_\tau = 0.0071$). For the 1000 dish scenario, reducing the integration time to 1000 hours results in no improvement over Planck in all of our scenarios.

[13] https://github.com/steven-murray/21cmSense

range. As we have seen, such a wide redshift range is more than sufficient to provide tight constraints on τ.

The range of wavenumbers transverse to the line of sight that can be probed by an interferometer depends upon the distribution of baseline separations [584]. Since the most constraining modes for our purposes are the low k modes we are primarily interested in the number of short baselines. If we take a minimum dish separation of 14 m for our fiducial cosmology $k_\perp^{\rm min} = 6\times10^{-3} h\,{\rm Mpc}^{-1}$ at $z=10$ and $4\times10^{-3} h\,{\rm Mpc}^{-1}$ at $z=15$. The baselines that probe $k_\perp \sim 0.1\,h\,{\rm Mpc}^{-1}$ are 240 m at $z=10$ and 380 m at $z=15$. All of these baselines would be well probed by next-generation experiments.

The minimum wavenumber that can be probed in the line of sight direction depends upon the ability of the experiment to subtract the foregrounds that are many orders of magnitude brighter than the signal, but spectrally smooth. Methods for foreground subtraction are rapidly improving, so we shall simply explore how our constraint on τ depends upon $k_\parallel^{\rm min}$, which inherently limits the accessible μ-range. For our fiducial scenario we will take $k_\parallel^{\rm min} = 0.1\,h\,{\rm Mpc}^{-1}$. The maximum observationally accessible wavenumber is normally set by thermal noise and the lack of many widely separated baselines. However since we are attempting to focus on perturbative scales this will not be the only limitation of interest to us. Rather our $k_{\rm max}$ will (also) be set by the scale at which our bias expansion becomes inapplicable, which is mostly set by the evolution of the "effective" bubble size. We shall take $k_{\rm max}$ to be the smaller of $k^h_{\rm max} = 0.5\,h\,{\rm Mpc}^{-1}$ and the scale at which our lowest-order counterterm is X_α^{-2} of the linear power spectrum:

$$k_{\rm max}(z) = \min\left[\frac{1}{X_\alpha\sqrt{|\alpha_0(z)|}}\,,\,k^h_{\rm max}\right] \quad \text{where} \quad \alpha_0(z) = -\frac{2}{3}b^2(z)R^2_{\rm eff}(z) \tag{11.17}$$

and consider $X_\alpha = 3$ and 4. In principle the validity of the bias expansion in real space is set strictly by the nonlocality scale $R_{\rm eff}$ but we have adopted $1/\sqrt{\alpha_0}$ to serve as a catch-all, including both the validity of the bias expansion and the significance of the linear RSD signal from which most of our constraint derives; we have checked that this choice does not significantly impact our σ_τ forecast.

Finally, contamination from systematic errors in the inherently chromatic response of an interferometer, interacting with astrophysical sources that are many orders of magnitude brighter than the signals of interest, leads to contamination within a "foreground wedge" [12, 243]. A particularly lucid and pedagogical introduction to the physical origin of the wedge can be found in ref. [420]. Given the rapidly evolving state of 21cm analysis, we shall take a simple model of the wedge and explore how variations in performance will affect our inference. In particular we follow ref. [42] in modeling the foreground wedge as allowing only modes with

$$k_\parallel > \max\left[k_\perp\frac{\chi(z)H(z)}{c(1+z)}\sin(\theta_w(z))\,,\,k_\parallel^{\rm min}\right] \quad \text{where} \quad \theta_w(z) = N_w\frac{1.22}{2\sqrt{0.7}}\frac{\lambda_{\rm obs}}{D_{\rm phys}} \tag{11.18}$$

where $\lambda_{\rm obs} = (1+z)\,21\,{\rm cm}$ is the observed wavelength, $D_{\rm phys}$ is the physical diameter of the dish, $k_\parallel = k\mu$ and $k_\perp = k\sqrt{1-\mu^2}$. For our fiducial scenario we will take $N_w = 3$.

Fig. 11.2 shows the foreground wedge as a function of redshift for $N_w = 3$ and $D_{\rm phys} = 14\,{\rm m}$. The minimum μ uncontaminated by foregrounds rises from $\simeq 0.5$ at $z = 6$ to $\simeq 0.9$ at $z = 15$. For comparison the minimum μ kept in the current analysis of HERA data is $\simeq 0.98$ [12, 489], which would render measurement of RSD ineffective. By contrast SKA predicts an observable window of $\mu > 0.67$. Clearly considerable effort remains to be done on instrument calibration and foreground removal to enable our forecasts.

Scenario	1000 dishes								SKA-like	Planck	PSD
N_w	3	3	3	3	3	3	3	**6**	3	–	–
X_α	3	3	3	3	3	3	**4**	3	3	–	–
$k^h_{\rm max}$ $[h\,{\rm Mpc}^{-1}]$	0.5	0.5	0.5	0.5	0.5	**0.3**	0.5	0.5	0.5	–	–
$k^{\rm min}_\parallel$ $[h\,{\rm Mpc}^{-1}]$	0.1	0.1	0.1	0.1	**0.2**	0.1	0.1	0.1	0.1	–	–
$z_{\rm max}$	12	12	12	**10**	12	12	12	12	**10**	–	–
$z_{\rm min}$	7	7	**9**	7	7	7	7	7	7	–	–
$t_{\rm int}$ [years]	**10**	5	5	5	5	5	5	5	5	–	–
$1000 \times \sigma_\tau$	1.9	3.0	7.0	3.8	6.9	3.8	3.9	6.3	0.9	7.1	5.8
σ_{M_ν} [meV]	30	33	39	33	39	33	34	38	26	83	40

Table 11.1: Forecasts for the uncertainty on the optical depth and the sum of the neutrino masses for different experimental configurations. All forecasts on the optical depth include a prior from Planck and assume $f_{\rm sky} = 0.5$. The neutrino mass constraints all include Planck, SO and the emission line galaxy sample of DESI (PSD) following ref. [512]. The baseline constraints for these are shown in the last two columns. Columns 2-9 show the impact of varying each parameter one at a time from our fiducial to a more pessimistic scenario. As a more futuristic example we include an SKA-like forecast, in which the collecting area is increased to $1\,{\rm km}^2$, which increases the sensitivity significantly over our fiducial scenario.

Our main results are summarized in Table 11.1, which presents forecasts for the uncertainty on the optical depth and the sum of the neutrino masses for different experimental configurations. The 21cm-based forecasts are compared to the existing measurement precision of Planck [473, 475] and to a forecast for how well τ and M_ν could be constrained by a combination of future CMB and galaxy redshift surveys (Planck, Simons Observatory and the emission line galaxy sample of DESI, indicated as PSD [512]). Columns 2-9 show the impact of varying experimental configurations, one element at a time, at fixed array size (10^3 dishes) while the column labeled SKA-like shows the improvement afforded by a much larger array (with $1\,{\rm km}^2$ collecting area), which reduces the thermal noise by roughly a factor of 5. We see that $z_{\rm min}$, $k^{\rm min}_\parallel$ and N_w have the largest impact on σ_τ, as expected from our earlier discussion. Once $z_{\rm min} > 9$ we are no longer measuring the range of redshifts (or x)

that contribute the most to τ. For $k_{\parallel}^{\rm min} > 0.1\,h\,{\rm Mpc}^{-1}$ we become sensitive primarily to the astrophysics controlling bubble sizes which our method explicitly marginalizes over. When $k_{\rm min} \simeq 0.2\,h\,{\rm Mpc}^{-1}$ our $k_{\rm max} \leq k_{\rm min}$ for $x > 0.5$ and we are unable to probe contributions to τ from $z < 8$. For larger N_w the range in μ is too restricted to be able to separate the modes along the line of sight from those transverse to it, and thus we cannot break the degeneracy with bias. By contrast we are relatively insensitive to the smallest scales we model – changes to X_α and $k^h_{\rm max}$ change σ_τ only modestly from our baseline configuration. Column 2 and the SKA-like column shows the constraints improve non-negligibly for lower thermal noise levels, indicating that our fiducial array is far from sample-variance limited. The constraints from the 10 yr survey with our fiducial array would have the same precision as the forecasts for σ_τ from the proposed LiteBIRD satellite mission [569]. In contrast to the results for a noiseless survey (§11.4, or SKA), here we find that a noisy survey is moderately sensitive to the maximum redshift: decreasing $z_{\rm max}$ from 12 to 10 to avoid the impact of spin-temperature fluctuations can increase the uncertainty on the optical depth by $\sim 25\%$.

11.5 Discussion and conclusions

The optical depth to reionization τ is one of the key parameters in the 6-parameter ΛCDM model of cosmology [473]. Holding the amplitude of CMB fluctuation fixed, τ controls the amplitude of fluctuations in the Universe. Currently the tightest constraints on this parameter come from observations of CMB polarization, but in the future 21cm experiments have the potential to dramatically improve upon our constraints. We discuss an approach to constraining τ built around ideas very familiar in large-scale structure studies but less often used in reionization. We emphasize the role of large scale modes, redshift-space distortions and a symmetries-based bias model that allows a principled but efficient marginalization over astrophysical parameters. The split encoded in this formalism means we can focus on the cosmological information contained in large scale modes while efficiently removing sensitivity to small-scale physics (and since our expansion is based entirely on symmetry, it makes minimal assumptions about that physics). Conversely, information about the sources and morphology of reionization is encapsulated in the small-scale modes and bias coefficients, with the large-scale modes being set largely by the timing of reionization. This approach to analyzing 21cm interferometer data has implications for the range of scales and redshifts that contain most of the cosmological signal, and hence on the most impactful observational settings.

Given sufficient dynamic range in scale, and the ability to measure the anisotropy in the clustering that is imprinted by peculiar velocities, future 21cm experiments should be able to place constraints on τ that are much tighter than constraints from existing CMB experiments or even future galaxy redshift surveys and competitive with proposed, future space-based CMB anisotropy missions. In order to make these inferences robust to uncertainties in astrophysical models, a premium should be placed on measuring large scales and on cleaning the foreground wedge as much as possible. By contrast very high angular resolution is

less helpful. We find that much of the constraining power to τ is lost once $k_{\parallel}^{\rm min}$ rises to $\approx 0.2\,h\,{\rm Mpc}^{-1}$ or if the foreground wedge extends past $\approx 3\times$ the primary beam. The experiment must also measure well the redshifts for which reionization is most active, i.e. those with $0.2 < x < 0.8$.

We have performed our forecasts within the context of Lagrangian perturbation theory, since this has a natural connection to the excursion set formalism and the theory of peaks and easily handles redshift-space distortions. At low redshifts the Lagrangian bias expansion has been shown to perform better for halos and galaxies than the equivalent Eulerian formalism [415], however which is preferable for reionization is currently unknown. The Lagrangian and Eulerian approaches to perturbation theory both predict that the dynamics on the scales and redshifts relevant for determining τ are very close to linear, with small corrections that can be very well described by either formalism. The primary limitation is not the non-linear scale but the scale of non-locality (due to ionized bubbles) in the bias expansion. We have treated this as a fundamental limitation in our forecasts and treated the bias parameters as completely free; however, future information about the process of reionization from a combination of simulations and observations could allow us to place priors on the bias coefficients. Such priors would allow information to be extracted from smaller scales than we have considered. In this regard we encourage simulators (and observers) to provide information on the anisotropy of the power spectrum when publishing their models.

Our forecasts assume that reionization proceeds smoothly and in a manner that can be described as "late and fast". This is consistent with the current observations, but our understanding of the process of reionization is highly incomplete. Since the low redshift Universe is known to be highly ionized [222], the main cause for concern would be multiple or partial ionizations at high z. Such episodes would leave a clear signal in P_{21} that would be at odds with our default scenario, so would not lead to inference of a highly precise but inaccurate τ unless $x(z)$ and the spin temperature evolution somehow conspired. For this reason we have not considered such scenarios.

Our forecasts suggest that a scaled up version of HERA has the potential to measure the optical depth with statistical uncertainty $\sigma_\tau = 0.002 - 0.004$, while a SKA-like survey could achieve $\sigma_\tau = 0.0009$ (assuming a $5-10$ year integration time, we find no improvement over Planck with only $\sim$ 1000 hours). We note that this improvement in sensitivity is more modest than models which claim to accurately predict and derive τ-constraints from small-scale 21cm fluctuations (e.g. refs. [355, 257, 260], which use `21cmFAST` as the theory model). However, these more ambitious models run an increased risk of returning biased results should their predictions differ significantly from the true small-scale 21cm signal. By contrast, our approach of modeling the 21cm intensity field using a symmetries-based bias expansion has been chosen to mitigate astrophysical uncertainties, however, several notable systematics could remain. First, any error in the best-fit Weibull function to the true reionization history propagates as a systematic error to the optical depth measurement. This systematic error can be reduced by freeing up the functional form of the $x(z)$ parametrization, however, doing so (§11.4) comes with a significant noise penalty when a realistic foreground

wedge and $z_{\rm max}$ are accounted for. Second, we have neglected the contributions from doubly reionized helium, which is expected to contribute $\tau_{\rm He} \sim 0.0008$ to the optical depth. Even for an SKA-like survey the second ionization of helium would need to be uncertain by $\mathcal{O}(1)$ to give significant errors. Additionally, we have neglected spin temperature fluctuations and uncertainties in the primordial helium abundance. We have shown in §11.4 how our constraints vary with $z_{\rm max}$, which can be lowered to remove any spin-temperature induced systematics, while within our formalism it would be straightforward to marginalize over any uncertainty in Y_p[14].

Finally, in our discussion of perturbation theory we have made the implicit assumption that all matter (dark matter, baryons and neutrinos) in the Universe can be treated as a single fluid. In the case of neutrinos it is known [121, 72] that galaxies primarily trace the cold dark matter and baryons; since the bias coefficients for the 21cm signal primarily reflect those of neutral hydrogen and ionizing sources the same approximation likely holds in this context as well. Residual differences post-recombination in the baryon (b) and dark matter (c) fields can lead to additional bias contributions proportional to relative densities and velocities [633, 595, 530, 130]. These contributions are either stationary or decay with time, and thus tend to be extremely subdominant at low redshifts, but could conceivably be a non-negligble effect at the EoR and, at k beyond the acoustic scale, degenerate with the linear bias. However, we note that linear theory contributions to the observed clustering due to the relative velocity v_{bc} should be significantly below 1% of the total matter velocity contributions even at high redshifts and these terms should contribute negligibly to redshift-space distortions. Since our reionization measurement depends mostly on the observed anisotropy it should be quite immune to these additional biases.

Acknowledgments

We would like to thank Avery Meiksin for helpful discussions on Helium reionization, and Joshua Dillon for helpful discussion on thermal noise calculation. S.C. thanks Nickolas Kokron for useful discussions on the form of the 21cm signal. N.S. is supported by the Physics Division of Lawrence Berkeley National Laboratory. S.C. and M.W. are supported by the DOE. This research has made use of NASA's Astrophysics Data System and the arXiv preprint server. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

Note added. While this chapter was being finalized, ref. [489] appeared which takes a highly related approach to modeling 21cm perturbations using a large-scale-structure-

[14]Given a prior $\sigma_{Y_p} = 0.2437 \pm 0.008$ from Planck [473] and direct measurements of helium abundance, uncertainties in Y_p would result in $\sigma_\tau \sim 0.0003$. We therefore do not expect varying Y_p to significantly impact our results when realistic noise is included.

inspired bias modeling. Our Lagrangian formalism can be exactly mapped into their Eulerian one [136] so the approaches are highly synoptic.

Appendix A

Geometric foreground mitigation I

A.1 Foreground power spectra

Our processing of the Sehgal simulations follows [524]. After masking the sources detected in the foreground maps at 5 mJy, which are detected at about 5σ given the noise level assumed, the power spectrum of each foreground at 148 GHz is shown in Fig. A.1.

A.2 Hybrid estimator

The hybrid estimator is defined to be

$$\hat{\kappa}^{\rm H}_L = N^{\rm H}_L \left(\hat{\kappa}^{\rm QE}_L / N^{\rm QE}_L + \hat{\kappa}^{\rm S}_L / N^{\rm S}_L \right) \tag{A.1}$$

where the standard estimator $\hat{\kappa}^{\rm QE}_L$ is calculated using $\ell_{\rm min,T} \leq \ell \leq 2000$ and the shear estimator $\hat{\kappa}^{\rm S}_L$ is calculated using $2000 \leq \ell \leq \ell_{\rm max,T}$. The noise of the hybrid estimator is simply

$$N^{\rm H}_L = \frac{1}{1/N^{\rm QE}_L + 1/N^{\rm S}_L}.$$

assuming that shear and QE are uncorrelated. The bias B_L to the lensing signal is defined by

$$C^{\rm QE}_L = C^{\kappa}_L + N^{\rm QE}_L + B^{\rm QE}_L \tag{A.2}$$

$$C^{\rm S}_L = C^{\kappa}_L + N^{\rm S}_L + B^{\rm S}_L. \tag{A.3}$$

Assuming that the noises of $\hat{\kappa}^{\rm QE}_L$ and $\hat{\kappa}^{\rm S}_L$ are uncorrelated, the bias to the Hybrid estimator is

$$B^{\rm H}_L = \left(N^{\rm H}_L\right)^2 \left(\frac{B^{\rm QE}_L}{\left(N^{\rm QE}\right)^2} + \frac{B^{\rm S}_L}{\left(N^{\rm S}\right)^2} \right). \tag{A.4}$$

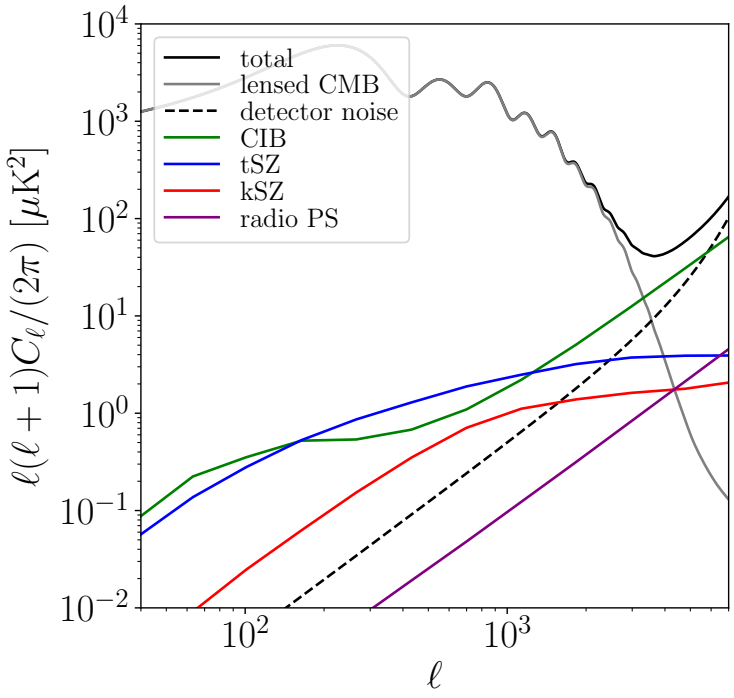


Figure A.1: Foreground power spectra at 148 GHz after being masked with a match-filter with a flux cut of 5 mJy and a mask patch radius of 3'. We assume a 6 μK-arcmin noise, and a 1.4' beam.

Here $B_L^{\rm QE}$ is the bias when the standard quadratic estimator is calculated using $\ell_{\rm min,T} \leq \ell \leq 2000$. We approximate $B_L^{\rm S}$ as the bias from Shear when calculated with $\ell_{\rm min,T} \leq \ell \leq \ell_{\rm max,T}$. This approximation should be OK since the contribution from $\ell_{\rm min,T} \leq \ell \leq 2000$ is negligible. The bias to the hybrid estimator's cross-correlation with LSST is calculated by inverse noise weighting the biases to QE and Shear.

A.3 Derivation of Eq. (2.20)

We start with

$$\boldsymbol{M_L} = \begin{pmatrix} 1 & N_{\boldsymbol{L}}^{\kappa}\mathcal{R}_{\boldsymbol{L}}^{\kappa,s_1} & \cdots & N_{\boldsymbol{L}}^{\kappa}\mathcal{R}_{\boldsymbol{L}}^{\kappa,s_n} \\ N_{\boldsymbol{L}}^{s_1}\mathcal{R}_{\boldsymbol{L}}^{s_1,\kappa} & 1 & \cdots & N_{\boldsymbol{L}}^{s_1}\mathcal{R}_{\boldsymbol{L}}^{s_1,s_n} \\ \vdots & \vdots & \vdots & \vdots \\ N_{\boldsymbol{L}}^{s_n}\mathcal{R}_{\boldsymbol{L}}^{s_n,\kappa} & N_{\boldsymbol{L}}^{s_n}\mathcal{R}_{\boldsymbol{L}}^{s_n,s_1} & \cdots & 1 \end{pmatrix} \equiv \begin{pmatrix} 1 & \boldsymbol{w}_{\boldsymbol{L}}^T \\ \boldsymbol{v_L} & \boldsymbol{A_L} \end{pmatrix}, \tag{A.5}$$

where $\boldsymbol{A_L}$ is a $n \times n$ matrix and $\boldsymbol{w_L}, \boldsymbol{v_L}$ are n-dimensional column vectors. By splitting the matrix into blocks, we can express its inverse and determinant as

$$\boldsymbol{M_L^{-1}} = \begin{pmatrix} (1 - \boldsymbol{w_L^T A_L^{-1} v_L})^{-1} & -(1 - \boldsymbol{w_L^T A_L^{-1} v_L})^{-1} \boldsymbol{w_L^T A_L^{-1}} \\ -\boldsymbol{A_L^{-1} v_L}(1 - \boldsymbol{w_L^T A_L^{-1} v_L})^{-1} & \boldsymbol{A_L^{-1}} + \boldsymbol{A_L^{-1} v_L}(1 - \boldsymbol{w_L^T A_L^{-1} v_L})^{-1} \boldsymbol{w_L^T A_L^{-1}} \end{pmatrix} \tag{A.6}$$

$$\det(\boldsymbol{M_L}) = \det(\boldsymbol{A_L})(1 - \boldsymbol{w_L^T A_L^{-1} v_L}). \tag{A.7}$$

Using this notation, the bias-hardened lensing estimator takes the form

$$\hat{\kappa}_{\boldsymbol{L}}^{\rm BH} = \frac{\hat{\kappa}_{\boldsymbol{L}} - \boldsymbol{w_L^T A_L^{-1} \hat{s}_L}}{1 - \boldsymbol{w_L^T A_L^{-1} v_L}}, \tag{A.8}$$

where $\hat{\boldsymbol{s}}_{\boldsymbol{L}} \equiv (\hat{s}_{1,\boldsymbol{L}}, \hat{s}_{2,\boldsymbol{L}}, \cdots, \hat{s}_{n,\boldsymbol{L}})^T$. The variance of the bias-hardened lensing estimator is simply

$$\begin{aligned} \langle \hat{\kappa}_{\boldsymbol{L}}^{\rm BH} \hat{\kappa}_{-\boldsymbol{L}}^{\rm BH} \rangle = (1 - \boldsymbol{w_L^T A_L^{-1} v_L})^{-2} \Bigg[& \langle \hat{\kappa}_{\boldsymbol{L}} \hat{\kappa}_{-\boldsymbol{L}} \rangle + \sum_{ij} \sum_{ab} (\boldsymbol{w_L})_i (\boldsymbol{A_L^{-1}})_{ij} (\boldsymbol{w_L})_a (\boldsymbol{A_L^{-1}})_{ab} \langle \hat{s}_{j,-\boldsymbol{L}} \hat{s}_{b,\boldsymbol{L}} \rangle \\ & - 2 \sum_{ij} (\boldsymbol{w_L})_i (\boldsymbol{A_L^{-1}})_{ij} {\rm Re} \langle \hat{\kappa}_{-\boldsymbol{L}} \hat{s}_{j,\boldsymbol{L}} \rangle \Bigg], \end{aligned} \tag{A.9}$$

from which we get the noise

$$\begin{aligned} N_{\boldsymbol{L}}^{\kappa^{\rm BH}} = (1 - \boldsymbol{w_L^T A_L^{-1} v_L})^{-2} \Bigg[& N_{\boldsymbol{L}}^{\kappa} + \sum_{ij} (\boldsymbol{w_L})_i (\boldsymbol{A_L^{-1}})_{ij} N_{\boldsymbol{L}}^{s_j} \sum_{ab} (\boldsymbol{w_L})_a (\boldsymbol{A_L^{-1}})_{ab} N_{\boldsymbol{L}}^{s_b} \mathcal{R}_{\boldsymbol{L}}^{s_b, s_j} \\ & - 2 N_{\boldsymbol{L}}^{\kappa} \sum_{ij} (\boldsymbol{w_L})_i (\boldsymbol{A_L^{-1}})_{ij} N_{\boldsymbol{L}}^{s_j} \mathcal{R}_{\boldsymbol{L}}^{\kappa, s_j} \Bigg]. \end{aligned} \tag{A.10}$$

In the past two equations we've assumed that the noises and responses are even functions of $\boldsymbol{L}$, which is true so long that the sources are modeled as halos with profiles $u_{\boldsymbol{L}}$. Let's focus on the second term of the RHS of the equation above. Note that $N_{\boldsymbol{L}}^{s_b} \mathcal{R}_{\boldsymbol{L}}^{s_b, s_j}$ is the b'th component of the j'th column vector of $\boldsymbol{A_L}$. Therefore

$$\sum_b (\boldsymbol{A_L^{-1}})_{ab} N_{\boldsymbol{L}}^{s_b} \mathcal{R}_{\boldsymbol{L}}^{s_b, s_j} = \delta_{aj}^K. \tag{A.11}$$

From this we find

$$\sum_{ab} (\boldsymbol{w_L})_a (\boldsymbol{A_L^{-1}})_{ab} N_{\boldsymbol{L}}^{s_b} \mathcal{R}_{\boldsymbol{L}}^{s_b, s_j} = \sum_a (\boldsymbol{w_L})_a \sum_b (\boldsymbol{A_L^{-1}})_{ab} N_{\boldsymbol{L}}^{s_b} \mathcal{R}_{\boldsymbol{L}}^{s_b, s_j} = \sum_a (\boldsymbol{w_L})_a \delta_{aj}^K = (\boldsymbol{w_L})_j. \tag{A.12}$$

Recall that the j'th component of $\boldsymbol{w_L}$ is just $N^{\kappa}_{\boldsymbol{L}}\mathcal{R}^{\kappa,s_j}_{\boldsymbol{L}}$. Plugging this result back into our expression for the noise gives

$$\begin{aligned} N^{\kappa^{\rm BH}}_{\boldsymbol{L}} &= (1-\boldsymbol{w}^T_{\boldsymbol{L}}\boldsymbol{A}^{-1}_{\boldsymbol{L}}\boldsymbol{v}_{\boldsymbol{L}})^{-2}\left[1-\sum_{ij}(\boldsymbol{w}_{\boldsymbol{L}})_i(\boldsymbol{A}^{-1}_{\boldsymbol{L}})_{ij}\underbrace{N^{s_j}_{\boldsymbol{L}}\mathcal{R}^{\kappa,s_j}_{\boldsymbol{L}}}_{=(\boldsymbol{v}_{\boldsymbol{L}})_j}\right]N^{\kappa}_{\boldsymbol{L}} \\ &= (1-\boldsymbol{w}^T_{\boldsymbol{L}}\boldsymbol{A}^{-1}_{\boldsymbol{L}}\boldsymbol{v}_{\boldsymbol{L}})^{-1}N^{\kappa}_{\boldsymbol{L}} \\ &= \frac{\det(\boldsymbol{A}_{\boldsymbol{L}})}{\det(\boldsymbol{M}_{\boldsymbol{L}})}N^{\kappa}_{\boldsymbol{L}}. \end{aligned} \tag{A.13}$$

A.4 tSZ-like profile

If the tSZ foreground map is modeled as $s_{\boldsymbol{\ell}} = \sum_i s_i u_{\boldsymbol{\ell}}$, then the tSZ power spectrum is proportional to $|u_{\boldsymbol{\ell}}|^2$. Recall that the bias-hardened estimator is insensitive to the normalization of the profile. Therefore we take the square root of the tSZ power spectrum (measured from the simulations) as our tSZ-like profile, which is plotted in Fig. A.2. In practice, determining the exact profile from the measured tSZ power spectrum may be challenging. However, we have checked that the residual biases from tSZ and other foregrounds vary slowly when the assumed profile is changed, such that a precise knowledge of the profile is not needed. This is encouraging, and suggests that this method will be useful in practice.

A.5 Noise price of bias-hardening increases with profile size

We found that the noise price for hardening against a tSZ-like profile is larger than hardening against point sources. To gain an intuition for why this is the case, we plot the noise when hardening against a single Gaussian profile $u_{\boldsymbol{\ell}} = e^{-\sigma^2\ell^2/2}$ for different values of σ in Fig. A.3. We find that the noise increases with σ; that is, the larger the profile the larger the cost in noise. As shown in Fig. 2.1, in the regime where the response is small the noise cost goes as $\mathcal{R}^2_L$. The response $\mathcal{R}_L \propto \int d^2\boldsymbol{\ell} f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}/C^{\rm tot}_{\ell}C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}$ is larger when the linear response to sources $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = e^{\sigma^2L^2/2}e^{-\sigma^2\ell^2/2}e^{-\sigma^2|\boldsymbol{L}-\boldsymbol{\ell}|^2/2}$ looks more like $f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$. For point sources $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ is flat, however, as the size of the profile gets larger $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ looks slightly more similar to the linear response to lensing, making it more difficult for the source-hardened estimator to remove the sources, and results in a higher noise price.

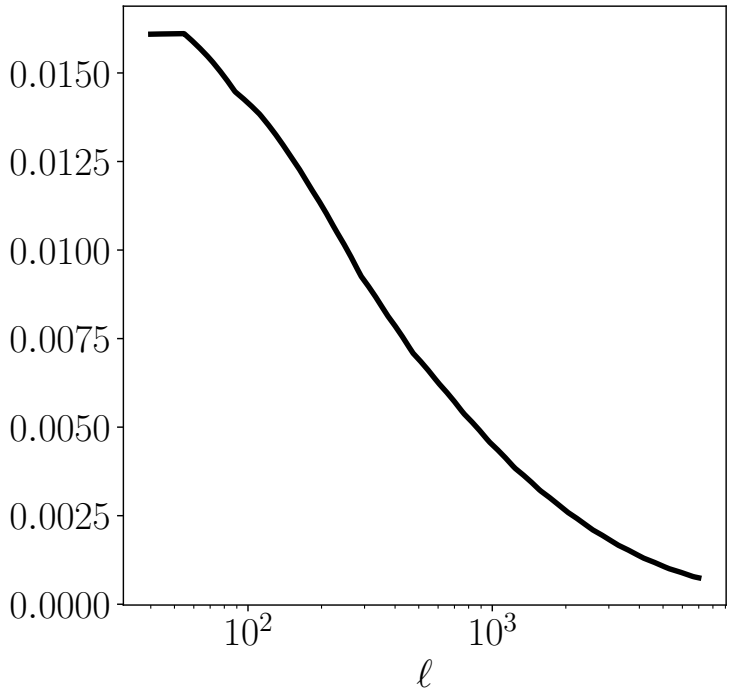


Figure A.2: tSZ-like profile $u_{\boldsymbol{\ell}}$. For $\ell \lesssim 60$ we approximate the profile as a flat curve.

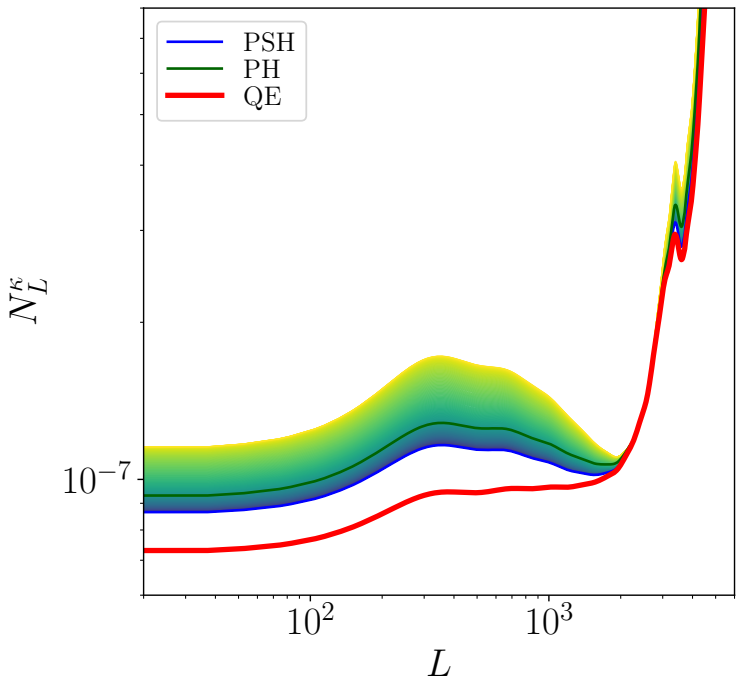


Figure A.3: **Shaded (blue to yellow) region:** Noise when bias hardening against a single Gaussian profile $f^s_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = e^{\sigma^2 L^2/2} e^{-\sigma^2 \ell^2/2} e^{-\sigma^2 |\boldsymbol{L}-\boldsymbol{\ell}|^2/2}$ for $\sigma \in [0', 2.4']$. σ is monotonically increasing from blue to yellow. Note that $\sigma = 0'$ corresponds to hardening against point sources (PSH). **Red:** Noise of the standard QE.

A.6 Noise and cross-correlation checks

As a simple pipeline check we verify our calculation of the noise and response when bias-hardening against point sources (PSH). These checks are shown in the plots below.

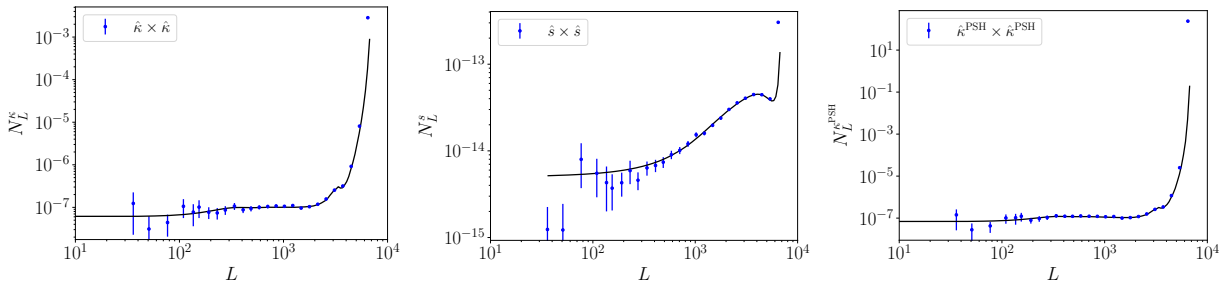


Figure A.4: Verification of the noises of each estimator. The auto-correlations (blue) of the estimators $\hat{\kappa}, \hat{s},$ and $\hat{\kappa}^{\rm PSH}$ when run on a Gaussian random field with power spectrum $C_L^{\rm tot}$. In black are the theory curves.

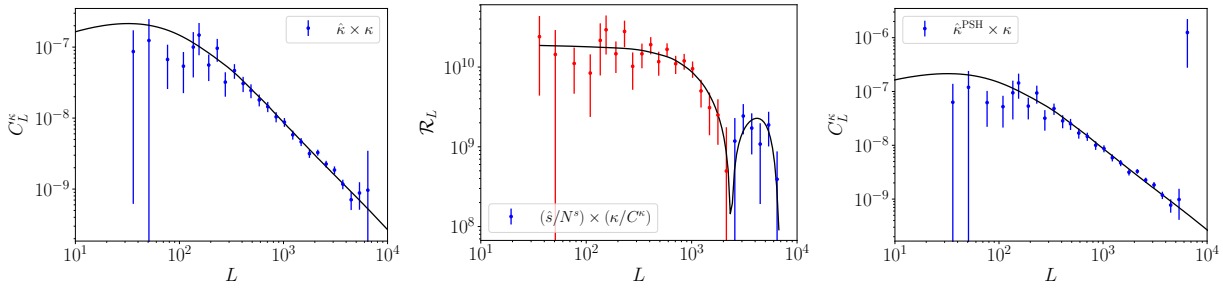


Figure A.5: Cross-correlating the true convergence κ with $\hat{\kappa}$, $\hat{s}$, and $\hat{\kappa}^{\rm BH}$ when run on a lensed CMB map. We see that both the standard QE $\hat{\kappa}$ and the point source hardened estimator $\hat{\kappa}^{\rm PSH}$ recover the lensing signal C_L^κ when correlated with the true convergence. The cross correlation of the point source estimator $\hat{s}$ with the true lensing convergence gives $N_L^s C_L^\kappa \mathcal{R}_L$, which is verified in the middle plot. The blue points are positive, whereas the red points are negative. The response has a zero crossing at $L \sim 2500$.

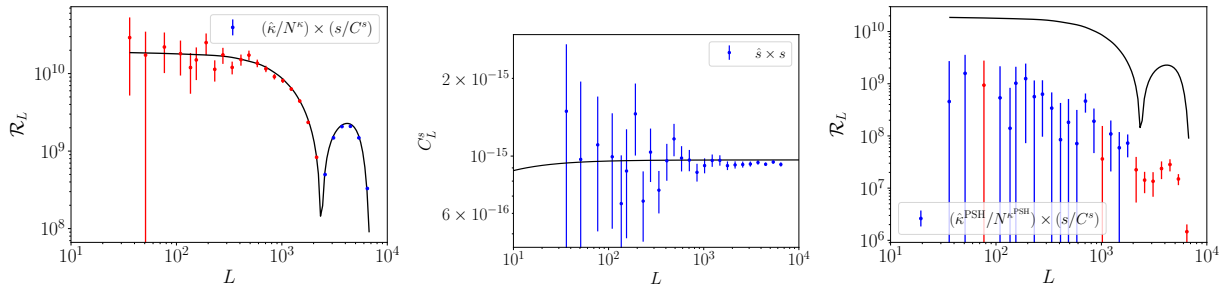


Figure A.6: Cross-correlating the true source map s with $\hat{\kappa}$, $\hat{s}$, and $\hat{\kappa}^{\rm PSH}$ when run on the true point source map s. The cross-correlation of $\hat{\kappa}$ with s gives $N_L^{\kappa} C_L^s \mathcal{R}_L$, which is verified in the plot on the left. $\hat{s}$ recovers the source signal C_L^s when cross correlated with the true source map. Since the point-source hardened estimator is designed to have zero response to point sources, the cross-correlation of $\hat{\kappa}^{\rm PSH}$ and s should be zero, which is consistent with the result found in the rightmost plot.

Appendix B

Geometric foreground mitigation II

B.1 Expansions in the large-lens limit

Any function of $|\boldsymbol{L}-\boldsymbol{\ell}|$ can be Taylor expanded in powers of L/ℓ to give:

$$u_{|\boldsymbol{L}-\boldsymbol{\ell}|} = u_\ell - L u'_\ell \cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{1}{4}\left(\frac{L}{\ell}\right)^2 \left[D^+_\ell u_\ell + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) D^-_\ell u_\ell\right] \tag{B.1}$$

with corrections going as $\mathcal{O}((L/\ell)^3)$, where in Eq. (B.1) we have defined the operators

$$D^\pm_\ell \equiv \ell^2 \frac{d^2}{d\ell^2} \pm \frac{d}{d\ln\ell}. \tag{B.2}$$

We can similarly Taylor expand the following geometric factors to quadratic order:

$$\begin{aligned}
\cos(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}) &= 1 - \left(\frac{L}{\ell}\right)^2 + \left(\frac{L}{\ell}\right)^2 \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) \\
\sin(2\theta_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}) &= 2\frac{L}{\ell}\sin(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \left(\frac{L}{\ell}\right)^2 \sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) \\
\cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}}) &= \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{L}{\ell}\left[\cos(3\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - \cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right] + \left(\frac{L}{\ell}\right)^2 \left[\cos(4\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right] \\
\sin(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}}) &= \sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{L}{\ell}\left[\sin(3\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - \sin(\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right] + \left(\frac{L}{\ell}\right)^2 \left[\sin(4\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - \sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right]
\end{aligned} \tag{B.3}$$

Using these expressions one can Taylor expand the linear responses in Eq. (3.4):

$$
\begin{aligned}
f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= C^{TT}_\ell \left[\frac{d\ln \ell^2 C^{TT}_\ell}{d\ln\ell} + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\frac{d\ln C^{TT}_\ell}{d\ln\ell}\right] \\
&\quad - \frac{1}{2}\frac{L}{\ell}\left[D^+_\ell C^{TT}_\ell + 4\frac{dC^{TT}_\ell}{d\ln\ell}\cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})D^-_\ell C^{TT}_\ell\right] \\
f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= C^{TE}_\ell \left[\frac{d\ln \ell^2 |C^{TE}_\ell|}{d\ln\ell} + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\frac{d\ln C^{TE}_\ell}{d\ln\ell}\right] \\
&\quad - \frac{1}{2}\frac{L}{\ell}\left[\big(D^+_\ell + 4\big)C^{TE}_\ell + 4\frac{dC^{TE}_\ell}{d\ln\ell}\cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\big(D^-_\ell - 4\big)C^{TE}_\ell\right] \\
f^{TB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= 2\sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})C^{TE}_\ell + \frac{L}{\ell}C^{TE}_\ell\left[\sin(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \sin(3\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right] \\
f^{EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= C^{EE}_\ell \left[\frac{d\ln \ell^2 C^{EE}_\ell}{d\ln\ell} + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\frac{d\ln C^{EE}_\ell}{d\ln\ell}\right] \\
&\quad - \frac{1}{2}\frac{L}{\ell}\left[D^+_\ell C^{EE}_\ell + 4\frac{dC^{EE}_\ell}{d\ln\ell}\cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})D^-_\ell C^{EE}_\ell\right] \\
f^{EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= 2\sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})C^{EE}_\ell + \frac{L}{\ell}C^{EE}_\ell\left[\sin(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \sin(3\theta_{\boldsymbol{L},\boldsymbol{\ell}}),\right]
\end{aligned}
\tag{B.4}
$$

with corrections going as $\mathcal{O}((L/\ell)^2)$. The colors here denote terms which are orthogonal to the relevant foreground linear response functions to $\mathcal{O}((L/\ell)^2)$ (blue), and those which aren't (red). That is, within the LLS approach we wish to retain sensitivity to the (clean) blue terms, while killing sensitivity to the red.

Likewise, one can Taylor expand the angular dependent pieces of the foreground linear

responses (Table 3.1) to give:

$$
\begin{aligned}
g^{TT,T}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \mathcal{A}\frac{u_\ell}{u_L}\left\{u_\ell - Lu'_\ell\cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{1}{4}\left(\frac{L}{\ell}\right)^2\left[D^+_\ell + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})D^-_\ell\right]u_\ell\right\} \\
g^{TE,E}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \mathcal{B}\frac{u_\ell}{u_L}\left\{u_\ell\cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{L}{\ell}\left[(u_\ell - \ell u'_\ell/2)\cos(3\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - (u_\ell + \ell u'_\ell/2)\cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right]\right. \\
&\quad \left. + \frac{1}{4}\left(\frac{L}{\ell}\right)^2\left[D^-_\ell u_\ell/2 + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})(D^+_\ell u_\ell - 4u_\ell) + (4u_\ell + D^-_\ell u_\ell/2)\cos(4\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right]\right\} \\
g^{TE,B}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= -\mathcal{B}\frac{u_\ell}{u_L}\left\{u_\ell\sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{L}{\ell}\left[(u_\ell - \ell u'_\ell/2)\sin(3\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - (u_\ell + \ell u'_\ell/2)\sin(\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right]\right. \\
&\quad \left. + \frac{1}{4}\left(\frac{L}{\ell}\right)^2\left[\sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})(D^+_\ell u_\ell - 4u_\ell) + \sin(4\theta_{\boldsymbol{L},\boldsymbol{\ell}})(4u_\ell + D^-_\ell u_\ell/2)\right]\right\} \\
g^{EE,T}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \mathcal{C}\frac{u_\ell}{u_L}\left\{u_\ell - Lu'_\ell\cos(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \frac{1}{4}\left(\frac{L}{\ell}\right)^2\left[(D^+_\ell - 4)u_\ell + \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})(D^-_\ell + 4)u_\ell\right]\right\} \\
g^{EB,T}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \mathcal{C}\frac{u_\ell}{u_L}\left\{2\frac{L}{\ell}u_\ell\sin(\theta_{\boldsymbol{L},\boldsymbol{\ell}}) - \left(\frac{L}{\ell}\right)^2\left[\ell u'_\ell + u_\ell\right]\sin(2\theta_{\boldsymbol{L},\boldsymbol{\ell}})\right\},
\end{aligned}
\tag{B.5}
$$

while the remaining two responses satisfy $g^{TB,E}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = -g^{TE,B}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ and $g^{TB,B}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = g^{TE,E}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$. Terms in $\{\cdots\}$'s are accurate to $\mathcal{O}((L/\ell)^3)$.

B.2 Explicit expressions for the global minimum variance weights

Let $\boldsymbol{F}^{\rm GMV}$ denote the minimum variance set of weights (Eq. (3.10)) for a quadratic estimator taking the form of Eq. (3.8), where the index i runs over T, E, B at a single frequency:

$$
\boldsymbol{F}^{\rm GMV}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \begin{pmatrix} F^{{\rm GMV},TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} & F^{{\rm GMV},TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} & F^{{\rm GMV},TB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \\ F^{{\rm GMV},TE}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}} & F^{{\rm GMV},EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} & F^{{\rm GMV},EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \\ F^{{\rm GMV},TB}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}} & F^{{\rm GMV},EB}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}} & 0 \end{pmatrix}. \tag{B.6}
$$

Assuming $\tilde{C}^{TB} = \tilde{C}^{EB} = 0$, we find:

$$
\begin{aligned}
F^{\mathrm{GMV},TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{\tilde{C}^{EE}_{\ell}\tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} + \tilde{C}^{TE}_{\ell}\tilde{C}^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{EE}_{\ell}\tilde{C}^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{TE}_{\ell}\tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TE}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}}}{2\mathcal{D}_{\ell}\mathcal{D}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\
F^{\mathrm{GMV},TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{\tilde{C}^{EE}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} + \tilde{C}^{TE}_{\ell}\tilde{C}^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TE}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}} - \tilde{C}^{TE}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{EE}_{\ell}\tilde{C}^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2\mathcal{D}_{\ell}\mathcal{D}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\
F^{\mathrm{GMV},TB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{\tilde{C}^{EE}_{\ell}f^{TB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{TE}_{\ell}f^{EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2\mathcal{D}_{\ell}\tilde{C}^{BB}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\
F^{\mathrm{GMV},EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{\tilde{C}^{TE}_{\ell}\tilde{C}^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} + \tilde{C}^{TT}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{EE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{TE}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{TT}_{\ell}\tilde{C}^{TE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}f^{TE}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}}}{2\mathcal{D}_{\ell}\mathcal{D}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\
F^{\mathrm{GMV},EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{\tilde{C}^{TT}_{\ell}f^{EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \tilde{C}^{TE}_{\ell}f^{TB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2\mathcal{D}_{\ell}\tilde{C}^{BB}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}
\end{aligned}
\tag{B.7}
$$

where we have defined $\mathcal{D}_{\ell} \equiv \tilde{C}^{TT}_{\ell}\tilde{C}^{EE}_{\ell} - (\tilde{C}^{TE}_{\ell})^2$. We note that our convention of summing over both permutations of pairs, e.g. both TE and ET, differs from that used in [382]. As a result our TE, TB and EB weights differ by a conventional factor of 2. Additionally we find that the last term of $F^{\mathrm{GMV},TE}$ differs from that in [382], where the arguments $\boldsymbol{\ell}$ and $\boldsymbol{L}-\boldsymbol{\ell}$ have been flipped.

B.3 Cleaning combined estimators

Here we explore the feasibility of cleaning the GMV estimator via bias hardening and large-scale suppression. We find that for both cases one should in practice take an optimal linear combination of the relevant single-pair estimators, rather than to harden (or suppress) a GMV estimator.

Bias hardening

In section 3.3 we found that the optimal temperature-only bias hardened estimator is insensitive to the amplitude $\mathcal{A}$ of the foregrounds. Here we show that this is always be the case when bias hardening a single-pair lensing estimator.

Proof: In general, the optimal bias hardened weights for an estimator constructed from a single pair X, Y takes the (schematic) form:

$$
F^{XY} + \alpha^{T} G^{XY,T} + \alpha^{E} G^{XY,E} + \alpha^{B} G^{XY,B}, \tag{B.8}
$$

where F^{XY} is the MV lensing weight and $G^{XY,Z}$ is the MV weight for the foreground in the Z'th map. WLOG we will assume that the weights $G^{XY,Z} \propto g^{XY,Z}$ have been rescaled to remove the unknown prefactor ($\mathcal{A}$, $\mathcal{B}$ or $\mathcal{C}$) in the foreground response g, which can be

thought of as being absorbed into coefficients $(\alpha^T, \alpha^E, \alpha^B)$. These coefficients are chosen so that the bias hardened estimator has zero response to s^T, s^E and s^B. This is enforced by integrating Eq. (B.8) against $g^{XY,T}$, and forcing this integral to vanish (and similarly for $g^{XY,E}$ and $g^{XY,B}$). Since these integral constraints are set to zero, we are always free to divide out the unknown prefactor introduced by the foreground response $g^{XY,Z}$ in all three constraint equations. We can then solve for the coefficients $(\alpha^T, \alpha^E, \alpha^B)$ without any knowledge of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$. Thus Eq. (B.8) is independent of $(\mathcal{A}, \mathcal{B}, \mathcal{C})$. □

This is no longer true, however, when bias hardening the GMV estimator. In this case, the optimal bias hardened weights take the same form as Eq. (B.8), but with F^{XY} replaced by the GMV lensing weights, and similarly for the foreground weights (e.g. the GMV weights for s^T takes the same form as Eq. (3.10), but with f^{ij} replaced by $g^{XY,T}$, see Appendix B.4). From Table 3.1 note that the GMV weights for s^E and s^B are proportional to $\mathcal{B}$. However, the GMV weights for s^T contain a linear combination of $\mathcal{A}$ and $\mathcal{C}$. Consequentially, one cannot bias harden the GMV κ estimator without assuming the ratio $\mathcal{C}/\mathcal{A} \simeq \langle p_i^2 \rangle / 2$.

To avoid assuming this ratio, one can simply take an optimal linear combination of the single-pair bias hardened estimators, which takes the form $\sum_i w_{\boldsymbol{L}}^i \hat{\kappa}_{\boldsymbol{L}}^{\mathrm{BH},i}$, where $i \in \{TT, TE, TB, EE, EB\}$ and $w_i = \sum_j C_{ji}^{-1} / \sum_{ij} C_{ij}^{-1}$, where C_{ij} is the covariance between $\hat{\kappa}^{\mathrm{BH},i}$ and $\hat{\kappa}^{\mathrm{BH},j}$:

$$C_{\boldsymbol{L}}^{XY,MN} = N_{\boldsymbol{L}}^{XY} N_{\boldsymbol{L}}^{MN} \int_{\boldsymbol{\ell}} F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{XY} \left[F_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}^{MN} \tilde{C}_{\ell}^{XM} \tilde{C}_{|\boldsymbol{L}-\boldsymbol{\ell}|}^{YN} + F_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}}^{MN} \tilde{C}_{\ell}^{XN} \tilde{C}_{|\boldsymbol{L}-\boldsymbol{\ell}|}^{YM} \right]. \tag{B.9}$$

The noise of this optimal linear combination is given by $1/\sum_{ij} C_{ij}^{-1}$.

This simple linear combination of bias hardened estimators is very close to optimal (while remaining practical) for the following reason. The GMV estimator can be separated into two statistically independent GMV estimators constructed from $\{TT, TE, EE\}$ and $\{TB, EB\}$ [382]. Following the same argument as above one cannot bias harden the former without assuming the ratio $\mathcal{C}/\mathcal{A}$. To get around this, one can further decompose the set into $\{TT, TE\}$ and $\{EE\}$. In practice the GMV estimator constructed from $\{TT, TE\}$ is not FFT-able. To make this estimator FFT-able one can neglect the cross-correlation $\tilde{C}^{TE}$, however, in this approximation the $\{TT, TE\}$ GMV estimator reduces to the optimal linear combination of $\hat{\kappa}^{TT}$ and $\hat{\kappa}^{EE}$. Thus, to bias harden the $\{TT, TE, EE\}$ GMV estimator in an FFT-able way without assuming $\mathcal{C}/\mathcal{A}$ one must fully decompose $\{TT, TE, EE\}$ into an optimal linear combination of TT, TE and EE.

While one can in principle bias harden the GMV estimator constructed from $\{TB, EB\}$ in practice the gains for doing so (as opposed to an optimal linear combination) are expected to be minimal. We have checked that the reconstruction noise for the GMV constructed from $\{TB, EB\}$ and the optimal linear combination of TB and EB differ by less than half a percent for $L < 5000$.

Large-lens suppression

One can generalize the single-pair LLS estimator to a global form in a straightforward way. However, as we show below, we find that an optimal linear combination of single-pair LLS estimators is practically more useful.

From Eqs. (3.33) and (3.34) a natural ansatz for the global LLS weights is:

$$\boldsymbol{F}^{\rm LLS}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \begin{pmatrix} A_{L,\ell}\cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) & B_{L,\ell} & 0 \\ B_{L,|\boldsymbol{L}-\boldsymbol{\ell}|} & C_{L,\ell}\cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}}) & F^{{\rm MV},EB}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \\ 0 & F^{{\rm MV},EB}_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}} & 0 \end{pmatrix}, \tag{B.10}$$

with corrections going as $\mathcal{O}(L/\ell)$, where the matrix is in the basis $\{T, E, B\}$. The coefficients (A, B, C) should be chosen to minimize Eq. (3.9) in the large-lens limit:

$$\begin{aligned}\mathcal{N}^\kappa_{\boldsymbol{0}} = 2(N^\kappa_{\boldsymbol{0}})^2 \int \ell\, d\ell \Big[\frac{1}{2}(\tilde{C}^{TT}_\ell)^2 A^2_{0,\ell} + \frac{1}{2}(\tilde{C}^{EE}_\ell)^2 C^2_{0,\ell} + 2\big((\tilde{C}^{TE}_\ell)^2 + \tilde{C}^{TT}_\ell \tilde{C}^{EE}_\ell\big) B^2_{0,\ell} \\ + (\tilde{C}^{TE}_\ell)^2 A_{0,\ell} C_{0,\ell}\Big] + 2(N^\kappa_{\boldsymbol{0}})^2 \int_{\boldsymbol{\ell}} 2\tilde{C}^{EE}_\ell \tilde{C}^{BB}_\ell (F^{{\rm MV},EB}_{\boldsymbol{\ell},-\boldsymbol{\ell}})^2,\end{aligned} \tag{B.11}$$

where

$$(N^\kappa_{\boldsymbol{0}})^{-1} = \int_{\boldsymbol{\ell}} 2 f^{EB}_{\boldsymbol{\ell},-\boldsymbol{\ell}} F^{{\rm MV},EB}_{\boldsymbol{\ell},-\boldsymbol{\ell}} + \int \ell\, d\ell \left[\frac{1}{2}\frac{dC^{EE}_\ell}{d\ln\ell} C_{0,\ell} + 2C^{TE}_\ell \frac{d\ln\ell^2 C^{TE}_\ell}{d\ln\ell} B_{0,\ell} + \frac{1}{2}\frac{dC^{TT}_\ell}{d\ln\ell} A_{0,\ell}.\right] \tag{B.12}$$

By varying A, B, C individually and setting $\delta\mathcal{N}^\kappa = 0$ we find:

$$\begin{aligned} A_{L,\ell} &= \frac{(\tilde{C}^{EE}_\ell)^2 \frac{dC^{TT}_\ell}{d\ln\ell} - (\tilde{C}^{TE}_\ell)^2 \frac{dC^{EE}_\ell}{d\ln\ell}}{2\left[(\tilde{C}^{TT}_\ell)^2(\tilde{C}^{EE}_\ell)^2 - (\tilde{C}^{TE}_\ell)^4\right]} \qquad C_{L,\ell} = \frac{(\tilde{C}^{TT}_\ell)^2 \frac{dC^{EE}_\ell}{d\ln\ell} - (\tilde{C}^{TE}_\ell)^2 \frac{dC^{TT}_\ell}{d\ln\ell}}{2\left[(\tilde{C}^{TT}_\ell)^2(\tilde{C}^{EE}_\ell)^2 - (\tilde{C}^{TE}_\ell)^4\right]} \\ B_{L,\ell} &= \frac{C^{TE}_\ell \frac{d\ln\ell^2|C^{TE}_\ell|}{d\ln\ell}}{2\left[\tilde{C}^{TT}_\ell \tilde{C}^{EE}_\ell + (\tilde{C}^{TE}_\ell)^2\right]}. \end{aligned} \tag{B.13}$$

up to $\mathcal{O}(L/\ell)$ corrections.

As shown in Fig. B.1, the global LLS estimator (GLLS) negligibly improves upon the optimal linear combination of single-pair LLS estimators at large scales. At small scales where the $L/\ell \ll 1$ approximation breaks down the GLLS estimator is no longer guaranteed to be optimal. In fact we find that the optimal linear combination of LLS estimators outperforms (or is nearly identical to) the GLLS for all relevant scales.

B.4 Bias hardening as constrained minimization

Here we will derive an optimal bias hardened estimator using Lagrange multipliers. Let's suppose that we have a foreground s^i that we wish to harden a generic quadratic estimator

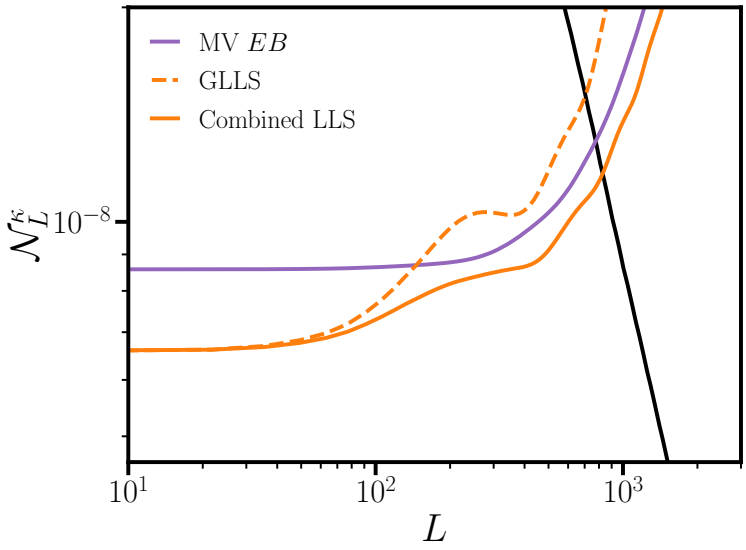


Figure B.1: Lensing reconstruction noise for the optimal linear combination of single-pair LLS estimators (solid orange) and for the Global LLS estimator (GLLS, dashed orange). For reference we show the reconstruction noise for the MV EB estimator (purple) and the convergence power spectrum (black).

(Eq. (3.8)) against, with the relevant linear responses given by $g^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$. To construct an optimal bias hardened estimator, we choose our lensing weights F^{ij} to minimize the noise of our lensing estimator under the constraint of having no response to the foreground. That is, we wish to minimize the Lagrangian:

$$\mathcal{L}_{\boldsymbol{L}}[F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}},\lambda^i_{\boldsymbol{L}}] = \mathcal{N}^\kappa_{\boldsymbol{L}} + 2\sum_i \lambda^i_{\boldsymbol{L}} \int_{\boldsymbol{\ell}} F^{jk}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} g^{jk,i}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tag{B.14}$$

where $\lambda^i_{\boldsymbol{L}}$ are some Lagrange multipliers. We first vary the weights, which gives

$$\delta\mathcal{L}_{\boldsymbol{L}} = 4(N^\kappa_{\boldsymbol{L}})^2 \sum_{mn} \int_{\boldsymbol{\ell}} \left[\sum_{ij} F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tilde{C}^{im}_\ell \tilde{C}^{jn}_{|\boldsymbol{L}-\boldsymbol{\ell}|} - \frac{1}{2}\frac{\mathcal{N}^\kappa_{\boldsymbol{L}}}{N^\kappa_{\boldsymbol{L}}} f^{mn}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} + \frac{1}{2}\sum_k \sigma^k_{\boldsymbol{L}} g^{mn,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \right] \delta F^{mn}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \tag{B.15}$$

where $\sigma^k_{\boldsymbol{L}} = \lambda^k_{\boldsymbol{L}}/(N^\kappa_{\boldsymbol{L}})^2$ are a renormalized set of Lagrange multipliers. After setting $\delta\mathcal{L}_{\boldsymbol{L}} = 0$ and fixing the normalization of the weights ($\mathcal{N}^\kappa_{\boldsymbol{L}} = N^\kappa_{\boldsymbol{L}}$), we find that the optimal zero-response weights satisfy

$$\begin{aligned} F^{ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} &= \frac{1}{2}\sum_{mn} \left(\tilde{\boldsymbol{C}}^{-1}_\ell\right)^{im} \left(\tilde{\boldsymbol{C}}^{-1}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right)^{jn} \left[f^{mn}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \sum_k \sigma^k_{\boldsymbol{L}} g^{mn,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \right] \\ &= F^{\mathrm{GMV},ij}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - \sum_k \sigma^k_{\boldsymbol{L}} G^{ij,k}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \end{aligned} \tag{B.16}$$

where $F^{\mathrm{GMV},ij}$ are the weights for the GMV κ estimator and $G^{ij,k}$ are the GMV weights for s^k. That is, the true minimum variance estimator with zero response to the foregrounds is some linear combination of the GMV estimators for κ and s^i. The linear combination which nulls the response of $\hat{\kappa}$ to s^i is unique, given by Eq. (3.15).

This approach can in principle be extended to null not only the field-level bias (and hence the primary and tripectrum biases, see Section 3.6), but also the secondary bias. Consider the temperature-only case for simplicity. In this scenario the secondary bias is proportional to:

$$\text{Secondary} \propto \int_{\boldsymbol{\ell}\boldsymbol{\ell}'} F^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} F^{TT}_{\boldsymbol{\ell}',-\boldsymbol{L}-\boldsymbol{\ell}'} f^{TT}_{\boldsymbol{\ell},-\boldsymbol{L}-\boldsymbol{\ell}'} g^{TT,T}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}}. \tag{B.17}$$

We null the secondary by adding a second term to the Lagrangian, yielding:

$$F^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{f^{TT}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2\tilde{C}^{TT}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} + \sigma_{\boldsymbol{L}} \frac{g^{TT,T}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2\tilde{C}^{TT}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} + \frac{\lambda_{\boldsymbol{L}}}{2\tilde{C}^{TT}_{\ell}\tilde{C}^{TT}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \int_{\boldsymbol{\ell}'} F^{TT}_{\boldsymbol{\ell}',-\boldsymbol{L}-\boldsymbol{\ell}'} f^{TT}_{\boldsymbol{\ell},-\boldsymbol{L}-\boldsymbol{\ell}'} g^{TT,T}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}} \tag{B.18}$$

where σ and λ are chosen to null the field-level and secondary biases respectively. In practice one would have to solve for these weights numerically, which could potentially be accomplished in an iterative fashion with the bias-hardened weights as an initial guess.

B.5 Derivation of polarization-based LLS estimators

In a similar manner to 3.4 we expand our lensing weights into multipoles:

$$F^{XY}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \sum_{m=0}^{\infty} A^{(m)}_{L,\ell} \cos(m\theta_{\boldsymbol{L},\boldsymbol{\ell}}) + \sum_{m=1}^{\infty} B^{(m)}_{L,\ell} \sin(m\theta_{\boldsymbol{L},\boldsymbol{\ell}}), \tag{B.19}$$

The noise of the lensing estimator estimator in the large lens limit is then

$$\mathcal{N}_0 \propto \frac{\int \ell\, d\ell \tilde{C}^{XX}_{\ell}\tilde{C}^{YY}_{\ell} \left[2\left(A^{(0)}_{0,\ell}\right)^2 + \sum_{m=1}^{\infty}\left(A^{(m)}_{0,\ell}\right)^2 + \sum_{m=1}^{\infty}\left(B^{(m)}_{0,\ell}\right)^2\right]}{\left[\sum_{m=0}^{\infty} A^{(m)}_{0,\ell} f^{XY,m}_{c,\ell} + \sum_{m=1}^{\infty} B^{(m)}_{0,\ell} f^{XY,m}_{s,\ell}\right]^2}. \tag{B.20}$$

where we have neglected the cross-correlation $\tilde{C}^{XY}$, and we have defined

$$f^{XY,m}_{c,\ell} = \int \frac{d\theta_{\boldsymbol{L},\boldsymbol{\ell}}}{2\pi} \cos(m\theta_{\boldsymbol{L},\boldsymbol{\ell}}) f^{XY}_{\boldsymbol{\ell},-\boldsymbol{\ell}} \qquad f^{XY,m}_{s,\ell} = \int \frac{d\theta_{\boldsymbol{L},\boldsymbol{\ell}}}{2\pi} \sin(m\theta_{\boldsymbol{L},\boldsymbol{\ell}}) f^{XY}_{\boldsymbol{\ell},-\boldsymbol{\ell}}. \tag{B.21}$$

Varying $A^{(m)}$, $B^{(m)}$ individually and setting $\delta\mathcal{N}_0 = 0$ then gives

$$A^{(0)}_{L,\ell} \propto \frac{f^{XY,0}_{c,\ell}}{2\tilde{C}^{XX}_{\ell}\tilde{C}^{YY}_{\ell}} \quad A^{(m)}_{L,\ell} \propto \frac{f^{XY,m}_{c,\ell}}{\tilde{C}^{XX}_{\ell}\tilde{C}^{YY}_{\ell}} \quad B^{(m)}_{L,\ell} \propto \frac{f^{XY,m}_{s,\ell}}{\tilde{C}^{XX}_{\ell}\tilde{C}^{YY}_{\ell}}, \tag{B.22}$$

up to $\mathcal{O}(L/\ell)$, where the (common) proportionality constant is any function of L. That is, we can read off the optimal values of $A^{(m)}, B^{(m)}$ directly from the multipole expansion of the linear response f^{XY} in the large-lens limit. To suppress the response to foregrounds we will manually set the coefficients which are sensitive to foregrounds to zero. We do so below for each single-pair estimator, making frequent use of Eqs. (B.4) and (B.5).

TE: Note that $g^{TE,E}$ contains a $\cos(2\theta)$ term which is $\mathcal{O}(1)$, while $g^{TE,B}$ has a $\mathcal{O}(1)$ $\sin(2\theta)$ term. We therefore manually set $A^{(2)} = B^{(2)} = 0$ to suppress responses to s^E and s^B. Both $g^{TE,E}$ and $g^{TE,E}$ only have a monopole at $\mathcal{O}((L/\ell)^2)$. Thus we can retain the monopole $A^{(0)}$, whose optimal value takes the form

$$A^{(0)}_{L,\ell} = \frac{C^{TE}_\ell}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_\ell} \frac{d\ln \ell^2 |C^{TE}_\ell|}{d\ln \ell}, \tag{B.23}$$

where we have made us of Eqs. (B.22) and (B.4).

TB: Likewise $g^{TB,E}$ contains a $\mathcal{O}(1)$ $\sin(2\theta)$ term while $g^{TE,B}$ has a $\mathcal{O}(1)$ $\cos(2\theta)$ term. We again set $A^{(2)} = B^{(2)} = 0$ to suppress responses to s^E and s^B. Both $g^{TE,E}$ and $g^{TE,E}$ only have a monopole at $\mathcal{O}((L/\ell)^2)$. Thus we can retain the monopole $A^{(0)}$, however, the response f^{TB} only has a $\sin(2\theta)$ term to leading order, and so we set $A^{(0)} = 0$. There is no LLS TB estimator to leading order.

EE: This case is identical to TT. We only retain the $\cos(2\theta)$ term.

EB: Note that $F^{\mathrm{MV},EB} = f^{EB}/\tilde{C}^{EE}\tilde{C}^{BB}$ only has a $\sin(2\theta)$ term at $\mathcal{O}(1)$, while at $\mathcal{O}(L/\ell)$ it has a $\cos(\theta)$, $\sin(\theta)$ and $\sin(3\theta)$ term. Note also that $g^{EB,T}$ only has a $\sin(\theta)$ term at $\mathcal{O}(L/\ell)$. Thus the foreground response of the MV EB estimator is already suppressed by $\mathcal{O}((L/\ell)^2)$.

B.6 Worked example: explicitly hardening the TE estimator

From Table 3.1 we see that the TE has a non-zero linear response to s^E and s^B in addition to lensing. Thus to bias harden $\hat{\kappa}$ against s^E and s^B we must construct (optimal) estimators for κ, s^E and s^B, which take the form:

$$\begin{aligned}
\hat{\kappa}_{\boldsymbol{L}} &= N^\kappa_{\boldsymbol{L}} \int_{\boldsymbol{\ell}} \frac{f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \tilde{T}_{\boldsymbol{\ell}} \tilde{E}_{\boldsymbol{L}-\boldsymbol{\ell}} \quad \text{where} \quad (N^\kappa_{\boldsymbol{L}})^{-1} = \int_{\boldsymbol{\ell}} \frac{(f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}})^2}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\
\hat{s}^E_{\boldsymbol{L}} &= N^E_{\boldsymbol{L}} \int_{\boldsymbol{\ell}} \frac{\cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \tilde{T}_{\boldsymbol{\ell}} \tilde{E}_{\boldsymbol{L}-\boldsymbol{\ell}} \quad \text{where} \quad (N^E_{\boldsymbol{L}})^{-1} = \int_{\boldsymbol{\ell}} \frac{\cos^2(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\
\hat{s}^B_{\boldsymbol{L}} &= N^B_{\boldsymbol{L}} \int_{\boldsymbol{\ell}} \frac{\sin(2\theta_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{L}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \tilde{T}_{\boldsymbol{\ell}} \tilde{E}_{\boldsymbol{L}-\boldsymbol{\ell}} \quad \text{where} \quad (N^B_{\boldsymbol{L}})^{-1} = \int_{\boldsymbol{\ell}} \frac{\sin^2(2\theta_{\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{L}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}},
\end{aligned} \tag{B.24}$$

where $\tilde{C}^{XX}_\ell$ is the total power spectrum, including the lensed field $\tilde{X}$ as well as instrumental noise and foregrounds. Here we have neglected the cross-correlation between T and E in

the minimum variance weights to make the estimators FFT-able, and we have assumed that the foreground is composed of point sources ($u_\ell = 1$). Note that we have neglected the normalization $\mathcal{B}$ of the foregrounds. This makes $\hat{s}^E$ an estimator of $\mathcal{B}\, s^E$, and similarly for $\hat{s}^B$. As we show in Appendix B.3, the bias hardened TE lensing estimator is insensitive to $\mathcal{B}$, which we will also show below explicitly.

The field-level biases to the minimum variance estimators above take the form:

$$\begin{pmatrix} \langle \hat{\kappa}_{\boldsymbol{L}} \rangle' \\ \langle \hat{s}^E_{\boldsymbol{L}} \rangle' \\ \langle \hat{s}^B_{\boldsymbol{L}} \rangle' \end{pmatrix} = \begin{pmatrix} 1 & \mathcal{R}^{\kappa E}_{\boldsymbol{L}} & \mathcal{R}^{\kappa B}_{\boldsymbol{L}} \\ \mathcal{R}^{E\kappa}_{\boldsymbol{L}} & 1 & \mathcal{R}^{EB}_{\boldsymbol{L}} \\ \mathcal{R}^{B\kappa}_{\boldsymbol{L}} & \mathcal{R}^{BE}_{\boldsymbol{L}} & 1 \end{pmatrix} \begin{pmatrix} \kappa_{\boldsymbol{L}} \\ \mathcal{B}\, s^E_{\boldsymbol{L}} \\ \mathcal{B}\, s^B_{\boldsymbol{L}} \end{pmatrix}, \tag{B.25}$$

where we have defined six response functions:

$$\begin{aligned} \mathcal{R}^{\kappa E}_{\boldsymbol{L}} &= N^\kappa_{\boldsymbol{\ell}} \int_{\boldsymbol{\ell}} \frac{f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} & \mathcal{R}^{E\kappa}_{\boldsymbol{L}} &= N^E_{\boldsymbol{\ell}} \int_{\boldsymbol{\ell}} \frac{f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\ \mathcal{R}^{\kappa B}_{\boldsymbol{L}} &= N^\kappa_{\boldsymbol{\ell}} \int_{\boldsymbol{\ell}} \frac{f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \sin(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} & \mathcal{R}^{B\kappa}_{\boldsymbol{L}} &= N^B_{\boldsymbol{\ell}} \int_{\boldsymbol{\ell}} \frac{f^{TE}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} \sin(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \\ \mathcal{R}^{EB}_{\boldsymbol{L}} &= N^E_{\boldsymbol{\ell}} \int_{\boldsymbol{\ell}} \frac{\cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}}) \sin(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} & \mathcal{R}^{BE}_{\boldsymbol{L}} &= N^B_{\boldsymbol{\ell}} \int_{\boldsymbol{\ell}} \frac{\cos(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}}) \sin(2\theta_{\boldsymbol{L},\boldsymbol{L}-\boldsymbol{\ell}})}{\tilde{C}^{TT}_\ell \tilde{C}^{EE}_{|\boldsymbol{L}-\boldsymbol{\ell}|}}. \end{aligned} \tag{B.26}$$

The bias hardened estimators are simply:

$$\begin{pmatrix} \hat{\kappa}^{\rm BH}_{\boldsymbol{L}} \\ \hat{s}^{{\rm BH},E}_{\boldsymbol{L}} \\ \hat{s}^{{\rm BH},B}_{\boldsymbol{L}} \end{pmatrix} = \begin{pmatrix} 1 & \mathcal{R}^{\kappa E}_{\boldsymbol{L}} & \mathcal{R}^{\kappa B}_{\boldsymbol{L}} \\ \mathcal{R}^{E\kappa}_{\boldsymbol{L}} & 1 & \mathcal{R}^{EB}_{\boldsymbol{L}} \\ \mathcal{R}^{B\kappa}_{\boldsymbol{L}} & \mathcal{R}^{BE}_{\boldsymbol{L}} & 1 \end{pmatrix}^{-1} \begin{pmatrix} \hat{\kappa}_{\boldsymbol{L}} \\ \hat{s}^E_{\boldsymbol{L}} \\ \hat{s}^B_{\boldsymbol{L}} \end{pmatrix}. \tag{B.27}$$

B.7 Expected biases to Simons Observatory

In Fig. B.2 we show the expected sensitivity to the lensing amplitude, and the bias from radio point sources, for Simons Observatory-like sensitivity. The assumptions regarding the experimental setup and simulations are identical to those in the main text, but with $\Delta_T = 6$ μK-arcmin instead of $\Delta_T = 1$ μK-arcmin. We see that for polarization-based reconstruction the point source-induced bias is negligible.

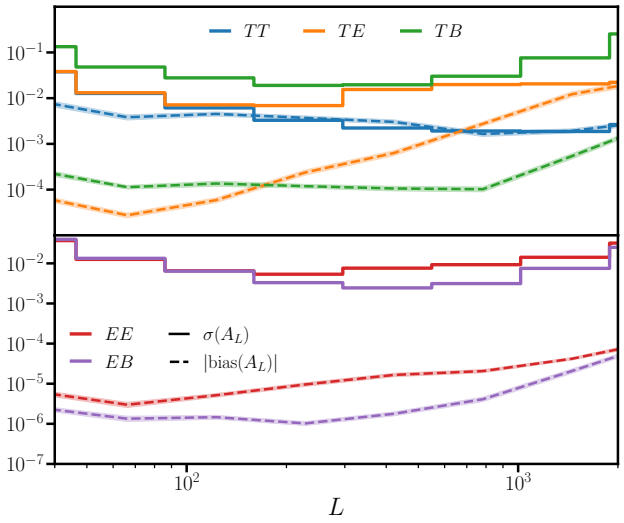


Figure B.2: The same as the top two panels of Fig. 3.4, but with Simons Observatory-like sensitivity instead of CMB-S4.

Appendix C

Optimal multifrequency weighting

C.1 Cancellations

As mentioned in §4.5, the definition of bias($A_{\rm lens}$) permits undesirable cancellations in biases across lensing multipoles. A simple method for avoiding these cancellations is to replace bias(C_L^κ) with |bias(C_L^κ)| in Eq. (4.17). Fig. C.1 shows how our results change with this substitution. In particular for the profile hardened estimator, we find that one has to deproject out to $t \sim 0.4$ (as opposed to $t \sim 0.2$) in order to achieve bias/noise $< 1/2$.

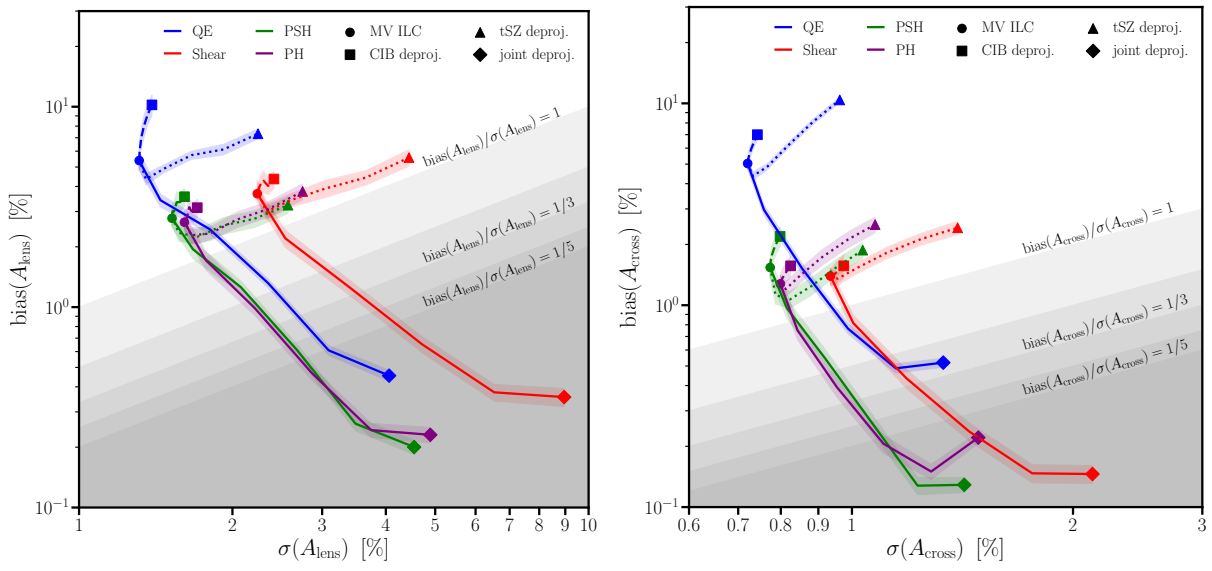


Figure C.1: The same as Fig. 4.3, but with |bias(C_L^κ)| replacing bias(C_L^κ) in Eq. (4.17), and similarly for the cross-correlation with the mock LSST sample.

In addition to the aforementioned cancellations, for any given foreground, the primary, secondary and trispectrum terms may also cancel separately due to the linear combination

of frequencies. While the distinction of primary, secondary and trispectrum terms is somewhat arbitrary, this cancellation relies on the relative sizes of the foreground bispectrum and trispectrum to be correct in the simulation used. Additionally, at any lensing multipole, the lensing biases from two distinct foregrounds may have different signs and partially cancel. While one can in principle take advantage of these cancellations, doing so relies on the simulations having the correct relative amplitude of each foreground, including their correlations. We will explore the sensitivity of our optimized ILC to these results in a future work, which will more realistically include polarization, in addition to temperature data.

C.2 Proof that ILC minimizes the lensing noise for all QEs

In this Appendix, we show that the linear combination of frequency maps which minimizes the lensing noise is the standard ILC, for all the estimators we consider. Specifically, we show the following: if the spatial weights $F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ of the estimator $\hat{\kappa}$ depend on the frequency weights $\boldsymbol{w}_\ell$ only through the map noise $C^{\rm tot}_\ell$, then the frequency weights $\boldsymbol{w}_\ell$ that extremize $\mathcal{N}_L[\boldsymbol{w}]$ are the ILC weights. To minimize the noise subject to the constraint $\boldsymbol{w}_\ell^T\mathbf{1}=1$ we set

$$\begin{aligned}\frac{\delta\mathcal{N}_L[\boldsymbol{w}]}{\delta\boldsymbol{w}_\ell} &= 4N_L^2F^2_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\left[C^{\rm tot}_\ell + C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}\right]\boldsymbol{C}_\ell\boldsymbol{w}_\ell + 4N_L^2\int\frac{d^2\boldsymbol{\ell}'}{(2\pi)^2}F_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}\frac{\delta F_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}}{\delta\boldsymbol{w}_\ell}C^{\rm tot}_{\ell'}C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}'|}\\ &\quad - 2N_L\mathcal{N}_L[\boldsymbol{w}]\int\frac{d^2\boldsymbol{\ell}'}{(2\pi)^2}\frac{\delta F_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}}{\delta\boldsymbol{w}_\ell}f^{\kappa}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}\\ &\quad = \lambda_{\boldsymbol{\ell},\boldsymbol{L}}\mathbf{1},\end{aligned} \tag{C.1}$$

where $\lambda_{\boldsymbol{\ell},\boldsymbol{L}}$ is some Lagrange multiplier. If the spatial weights $F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}$ only depend on frequency weights $\boldsymbol{w}_\ell$ through $C^{\rm tot}_\ell$ and $C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}$, then

$$\begin{aligned}\frac{\delta F^{\kappa}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}}{\delta\boldsymbol{w}_\ell} &= \frac{\partial F^{\kappa}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}}{\partial C^{\rm tot}_{\ell'}}\frac{\delta C^{\rm tot}_{\ell'}}{\delta\boldsymbol{w}_\ell} + \frac{\partial F^{\kappa}_{\boldsymbol{\ell}',\boldsymbol{L}-\boldsymbol{\ell}'}}{\partial C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}'|}}\frac{\delta C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}'|}}{\delta\boldsymbol{w}_\ell}\\ &= (2\pi)^2\frac{\partial F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{\partial C^{\rm tot}_\ell}\left[\delta^D_{\boldsymbol{\ell}-\boldsymbol{\ell}'} + \delta^D_{\boldsymbol{\ell}-\boldsymbol{L}+\boldsymbol{\ell}'}\right]\boldsymbol{C}_\ell\boldsymbol{w}_\ell\\ &\equiv \mathcal{A}_{\boldsymbol{\ell},\boldsymbol{\ell}',\boldsymbol{L}}\boldsymbol{C}_\ell\boldsymbol{w}_\ell\end{aligned} \tag{C.2}$$

where $\mathcal{A}_{\boldsymbol{\ell},\boldsymbol{\ell}',\boldsymbol{L}}$ is some scalar function. Plugging this in to our equation for $\delta\mathcal{N}_L[\boldsymbol{w}]/\delta\boldsymbol{w}_\ell$ gives

$$\frac{\delta\mathcal{N}_L[\boldsymbol{w}]}{\delta\boldsymbol{w}_\ell} = \mathcal{B}_{\boldsymbol{\ell},\boldsymbol{L}}\boldsymbol{C}_\ell\boldsymbol{w}_\ell = \lambda_{\boldsymbol{\ell},\boldsymbol{L}}\mathbf{1} \tag{C.3}$$

for some scalar function $\mathcal{B}_{\boldsymbol{\ell},\boldsymbol{L}}$. Solving this equation for $\boldsymbol{w}_\ell$ gives $\boldsymbol{w}_\ell = \lambda_{\boldsymbol{\ell},\boldsymbol{L}}\boldsymbol{C}_\ell^{-1}\mathbf{1}/\mathcal{B}_{\boldsymbol{\ell},\boldsymbol{L}}$. The Lagrange multiplier is chosen so the weights sum to unity, which yields the standard ILC weights, given by Eq. (4.7).

Note that the standard QE, bias-hardened and shear estimators have respective weights:

$$F^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\rm tot}_{\ell}C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \quad , \quad F^{\kappa^{\rm BH}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \frac{f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} - N^{\kappa}_{L}\mathcal{R}_{L}f^{s}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}}{2C^{\rm tot}_{\ell}C^{\rm tot}_{|\boldsymbol{L}-\boldsymbol{\ell}|}} \quad ,$$
$$F^{\rm shear}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}} = \cos(2\theta_{\boldsymbol{L},\boldsymbol{\ell}-\boldsymbol{L}/2})\frac{C^{0}_{\ell}}{2(C^{\rm tot}_{\ell})^{2}}\frac{d\ln C^{0}_{\ell}}{d\ln\ell}, \tag{C.4}$$

which only depend on $\boldsymbol{w}_\ell$ through $C^{\rm tot}_\ell$. Thus their noises are minimized by the ILC weights.

C.3 Attempts at numerically minimizing the loss functions: approximations

Finding the temperature weights which minimize the desired combination of lensing noise and bias is a difficult problem. Indeed, the space of parameters to optimize is large, since the temperature weights are free functions of scale and frequency. For each choice of temperature weights, one needs to perform multiple lensing reconstructions on the foreground simulations to estimate the primary, secondary and trispectrum biases. This step is computationally expensive. This needs to be repeated for each choice of estimator (QE, shear, PSH, PH), and at each step when exploring the space of temperature weights. Below, we describe approximations which dramatically simplify the problem, and point out their inaccuracies.

One may assume that each foreground is determined by a single spatial template, whose amplitude is rescaled at each frequency (Approximation 1). This approximation is exact for the CMB and for the tSZ, ignoring relativistic corrections. It is also excellent for the CIB, whose maps at different frequencies are highly correlated [378, 477].

To simplify further, we note that most of the CMB lensing signal to noise comes from temperature multipoles around $\ell = 3000$. As a result, most of the foreground biases also comes from these multipoles. At these multipoles, the ILC weights are slowly varying. We may therefore neglect the scale dependence of the ILC weights $\boldsymbol{w}^{\rm ILC}$, and look for scale-independent optimal weights $\boldsymbol{w}$ (Approximation 2). This approximation dramatically reduces the dimensionality of the minimization problem, making it much more tractable. In this approximation, the primary $\mathcal{P}^{s}_{0,L}$, secondary $\mathcal{S}^{s}_{0,L}$ and trispectrum $\mathcal{T}^{s}_{0,L}$ biases to $A_{\rm lens}$ only need to be computed at one frequency (e.g., 150 GHz). They can then simply be rescaled to any other frequency. Furthermore, because the temperature weights are scale-independent, the bispectrum and trispectrum biases only need to be computed from simulations once. They can then be rescaled for any new choice of temperature weights. In short, approximations 1 and 2 allow us to compute the foreground biases from simulation only once, rather than having to do it for every frequency and for every choice of weights $\boldsymbol{w}_\ell$.

Finally, we expect the optimal multifrequency combination will be close to the standard ILC. We indeed show this to be true for the profile hardened estimator when exploring part of the parameter space, going from ILC to some form of deprojection. This motivates us to Taylor expand the noise about the ILC weights to second order in $\boldsymbol{\epsilon}$ (Approximation

3): $\boldsymbol{w} = \boldsymbol{w}^{\rm ILC}_{\ell=3000} + \boldsymbol{\epsilon}$. This approximation dramatically speeds up the noise computation during the optimization step, and accurately predicts the variance and bias of $A_{\rm lens}$ except at 27 GHz. A simple method for getting around this is to set the 27 GHz weight to zero, in which case the Taylor expansion agrees with the true noise to the subpercent level for the $\boldsymbol{\epsilon}$'s relevant for achieving bias/noise $\sim 1/2$.

In this simplified case, the scaling of the foreground biases with the frequency weights is simply:

$$\text{bias}(C_L^\kappa)[\boldsymbol{w}] = \sum_s \left[\mathcal{P}^s_{0,L} + \mathcal{S}^s_{0,L}\right] \left(\frac{\boldsymbol{A}_s \cdot \boldsymbol{w}}{A_{s,150}}\right)^2 + \sum_s \mathcal{T}^s_{0,L} \left(\frac{\boldsymbol{A}_s \cdot \boldsymbol{w}}{A_{s,150}}\right)^4 \tag{C.5}$$

where $A_{s,150}$ is the scale-factor for the 150 GHz channel. In practice, this expression can be easily generalized to include the lensing bias from correlated foregrounds, e.g., CIB and tSZ. In summary, we are thus looking for a handful of parameters $(\epsilon_\nu)_{\nu=1\ldots5}$, which minimize the sum of squared variance and squared bias on $A_{\rm lens}$.

However, we found approximations 2 and 3 not to be sufficiently accurate. Assuming the weights to be scale-independent leads to a significant misestimation of the variance and bias of $A_{\rm lens}$, at the $\mathcal{O}(1)$ level.

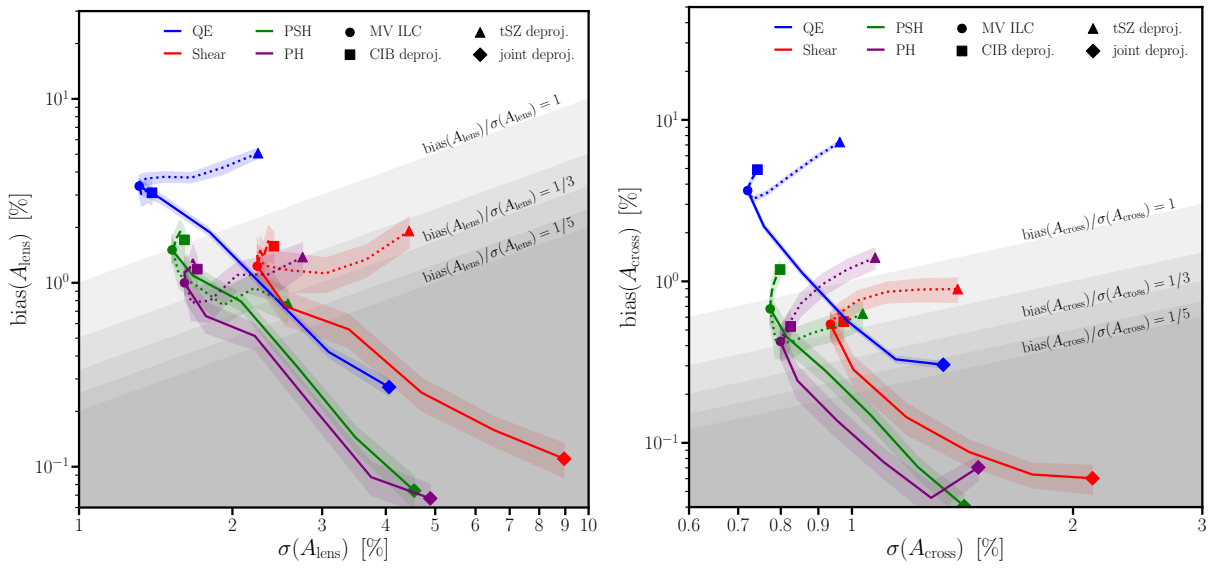


Figure C.2: The same as Fig. 4.3, using the more agressive masking technique described in App. C.4. When calculating σ, we neglect any differences in sky coverage due to more of the sky being masked (8% of the sky as opposed to 3% with the fiducial mask used in the main text), as this is a negligible effect.

C.4 Dependence of the lensing biases on the masking

We explore a more aggressive masking approach by running the matched filter at all the SO frequencies, masking only the sources (using disks with a 2 arcmin radius) that are detected with 5σ significance, and using the union of the masks at each frequency. This allows us to find much more IR (radio) sources, which are brighter at higher (lower) frequencies. In an analysis on real data, this may make it more delicate to estimate the normalization and mean field of the lensing estimators [336], which has not been taken into account in our lensing biases calculations.

We note that in practice, masking changes the observed map covariance $\boldsymbol{C}_\ell$, which in turn will influence the form of the ILC weights. For simplicity we have ignored this effect, and thus use the same ILC weights for the more aggressive masking approach as we did in the main text, which are derived using the theory foreground spectra from [200].

The resulting biases from this masking technique are shown in Fig. C.2. In cross-correlation, the aggressive masking reduces the foreground bias for all estimators, as expected. In particular, at ILC ($t = 0$) the biases to the cross-correlation for PH, Shear, PSH, and QE are reduced by $\sim 30\%$, 40%, 10%, and 60% respectively, relative to the biases calculated using the single mask at 150 GHz. In auto-correlation, the biases for PH, Shear and PSH are similarly reduced using the more aggressive masking technique. However, the bias to the QE at ILC increases by 50%. This can be understood as follows. In auto-correlation, both the primary and trispectrum terms contribute to the foreground bias. The aggressive masking reduces both, but reduces the trispectrum much more. For the QE, a "lucky" cancellation allows the dominant negative primary bias (at low lensing-L) to cancel the dominant positive trispectrum bias (at high lensing-L). By reducing the trispectrum term more than the primary term, the aggressive masking spoils this cancellation, thus enhancing the QE bias in auto-correlation.

C.5 Dependence of the lensing biases on $\ell_{\mathrm{max},T}$

In this section we explore how our results change with $\ell_{\mathrm{max},T}$. As shown in Fig. C.3, we still find that profile hardening outperforms the remaining estimators for any bias/noise cutoff and for any $\ell_{\mathrm{max},T}$. We also find that small improvements are possible by pushing to a higher $\ell_{\mathrm{max},T}$ with more joint deprojection. For the auto-correlation, the bias to the profile hardened estimator can be reduced below half of the statistical uncertainty at a 30% noise penalty by choosing $\ell_{\mathrm{max},T} = 4000$ and $t \sim 0.4$ (as opposed to a 50% penalty with $\ell_{\mathrm{max},T} = 3000$ and $t \sim 0.2$), relative to the standard QE at ILC with $\ell_{\mathrm{max},T} = 3000$. For the cross-correlation, the noise penalty for achieving bias/noise $< 1/2$ with the profile hardened estimator can be reduced from 15% to 10% by similarly increasing $\ell_{\mathrm{max},T}$ and t.

We note that for sufficiently high t, the biases and noises become roughly insensitive to the value of $\ell_{\mathrm{max},T}$. This is due to the map noise blowing up at high-ℓ as the weights are pushed to joint deprojection, as shown in Fig. 4.1. This in turn down-weights high-ℓ modes

and acts as an effective cutoff with $\ell_{\text{max eff}} < \ell_{\text{max},T}$. Instead of partial joint deprojection (Eq. (4.19)), one could naively use the standard ILC weights and reduce $\ell_{\text{max},T}$ to remove these high-ℓ modes, and hence the bias, in a potentially a more optimal way. However, in practice we find it more optimal to use a high $\ell_{\text{max},T}$ with joint deprojection than a lower $\ell_{\text{max},T}$ with the standard ILC weights. This is illustrated in Fig. C.3, where we plot the bias and noise at ILC for $\ell_{\text{max},T} = 1500$, 2000, and 2500 (individual circles). While lowering the $\ell_{\text{max},T}$ decreases the bias (as expected), the noise cost for doing so is much larger than increasing t with a higher $\ell_{\text{max},T}$ (as illustrated by the three lines for each estimator).

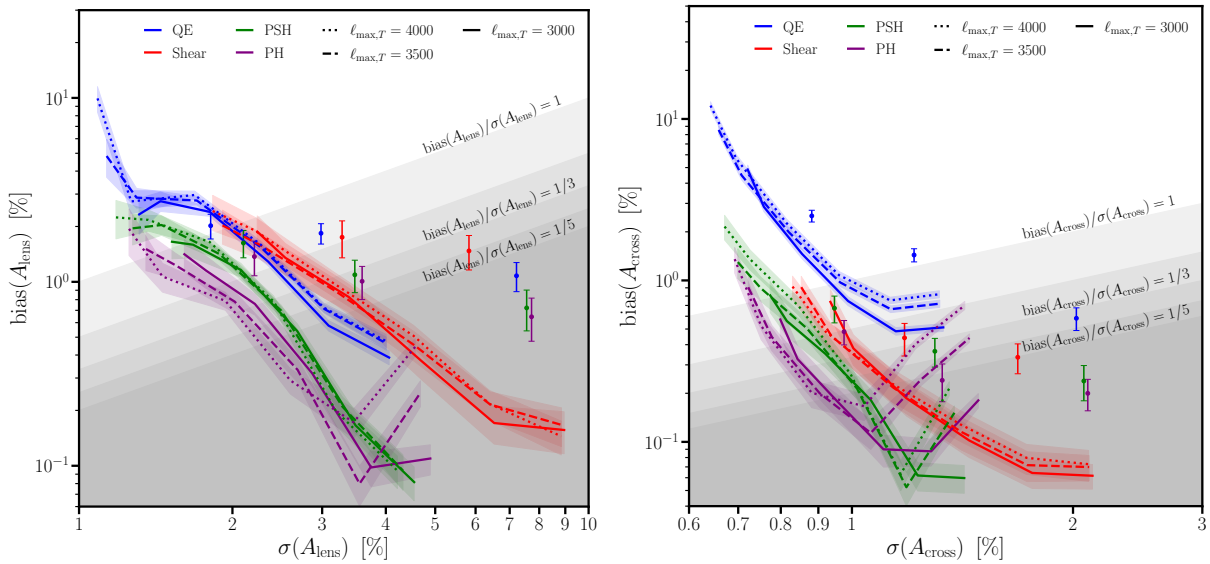


Figure C.3: Noise vs bias for several $\ell_{\text{max},T}$'s. For $\ell_{\text{max},T} = 3000$, 3500, and 4000 we plot the line going from the standard ILC to joint deprojection, with the linestyle denoting the $\ell_{\text{max},T}$. For $\ell_{\text{max},T} = 1500$, 2000, and 2500 we plot the noise and bias at the standard ILC for each estimator (denoted by the circles). Note that for $\ell_{\text{max},T} = 1500$, the Shear estimator's noise is larger than 10% (3%) for the auto- (cross-) correlation, extending beyond the domain of the plot.

Appendix D

Tomographic structure growth I

D.1 Data and code availability

The code base used to estimate ancillary power spectra and covariances, perform mock tests and cosmological analyses is publicly available (`MaPar`). The fiducial bandpowers, window functions and covariance used in our fiducial analysis are available at this URL, along with the chains and best-fit prediction for a subset of the analyses presented here. The `Aemulus` emulator weights will be made available upon the publication of an upcoming DES-DESI cross-correlation analysis [129]. The configuration for our baseline analysis is given by `fiducial_noBAO_pr4-dr6.yaml`, located in the `MaPar/yamls/` directory along with several other alternative analysis configurations. The fiducial redshift distributions are located in the `MaPar/data/dNdzs/` directory, while the fiducial power spectra used for covariance estimation can be found in `MaPar/spectra/fiducial/`. See the `README` for further information.

D.2 Signal to noise

For a power spectrum C_ℓ^{ab} with measurement error $\sigma(C_\ell^{ab})$ the signal-to-noise ratio can be expressed as

$$\begin{aligned}\mathrm{SNR}^2 &= \sum_\ell \frac{(C_\ell^{ab})^2}{\sigma^2(C_\ell^{ab})} \\ &= \int dz\, d\ln k \left[\frac{W^a(\chi)W^b(\chi)}{\chi^2 H(z)} k P_{ab}(k) \sum_\ell \frac{1}{\ell\, \sigma^2(C_\ell^{ab})} W^a(\ell k^{-1})\, W^b(\ell k^{-1})\, P_{ab}(\ell\chi^{-1})\right].\end{aligned} \tag{D.1}$$

where we used

$$C_\ell^{ab} \simeq \int d\chi \frac{W^a(\chi)W^b(\chi)}{\chi^2} P_{ab}(\ell\chi^{-1}) = \ell^{-1} \int dk\, W^a(k^{-1}l)\, W^b(k^{-1}l)\, P_{ab}(k) \tag{D.2}$$

to go from the first to the second line. In Fig. 5.1 we plot SNR(k, z), which is defined as the square root of the quantity in brackets in Eq. (D.1).

D.3 Parameter-based comparison of analysis choices

Our methodology for estimating ancillary power spectra and covariances is summarized in §5.3. There are several differences between the methods adopted here and in the companion paper: (1) the companion paper [320] adopts a slightly different binning for high-ℓ bandpowers ($\ell > 971$) that has minor impacts on the numerical computation of the window functions and bandpowers for $\ell < 600$, (2) the covariance matrices computed in [320] are either estimated numerically or estimated as an analytic-numerical hybrid, whereas the ancillary covariance matrices used in this chapter are all estimated analytically, (3) we always remask the CMB lensing map when computing pseudo-C_ℓ's whereas [320] remasks the PR4 map but not the DR6 map, and (4) we adopt a slightly different approach for computing the "normalization" correction, which we now describe. We refer the reader to the companion paper [320] for a description of the transfer function for computing this correction.

Mode-couplings arising from e.g. masking or anisotropic filtering of CMB maps are not properly forward modeled in CMB lensing cross-correlations measurements (which are integrals of squeezed limit bispectra) by the MASTER algorithm. To approximately account for these effects, the current state-of-the-art (see [515] for an improved approximation) is to multiply a CMB lensing cross-correlation measurement by a Monte-Carlo "normalization correction" estimated from a set of simulated CMB lensing reconstructions, which we compute as

$$(\text{MC correction})_L = \frac{\sum_\ell W_{L\ell} \sum_{i=1}^{N_{\rm sim}} \sum_{m=-\ell}^{\ell} \{M^\kappa \kappa^i\}_{\ell m} \{M^g \kappa^i\}^*_{\ell m}}{\sum_\ell W_{L\ell} \sum_{i=1}^{N_{\rm sim}} \sum_{m=-\ell}^{\ell} \{M^\kappa \hat{\kappa}^i\}_{\ell m} \{M^g \kappa^i\}^*_{\ell m}} \tag{D.3}$$

where M^κ (M^g) is the CMB lensing (galaxy) mask, $\hat{\kappa}^i$ (κ^i) is the reconstructed (input) CMB lensing convergence from the i'th CMB lensing reconstruction simulation, $\{AB\}_{\ell m} \equiv \int d^2\hat{\boldsymbol{n}} Y^*_{\ell m}(\hat{\boldsymbol{n}}) A(\hat{\boldsymbol{n}}) B(\hat{\boldsymbol{n}})$, and $W_{L\ell}$ is the window function corresponding to the bandpower L. We show the normalization corrections used in our ancillary cross-correlation measurements in Fig. D.1.

The ancillary power spectra and covariances are used in the main text to perform systematics checks and to reanalyze the cross-correlation with *Planck* PR3, while for our fiducial analysis we use the power spectra measurements and hybrid covariance from the companion paper [320]. As a cross-check, we independently measured the galaxy auto-correlations, the PR4 and DR6 cross-correlations and estimated their associated covariances following §5.3 and the discussion above. In Fig. D.2 we verify at the parameter-level the consistency of these measurements for the linear theory PR4+DR6 analysis. In particular, we find find that the S_8 constraints are identical to three decimal places.

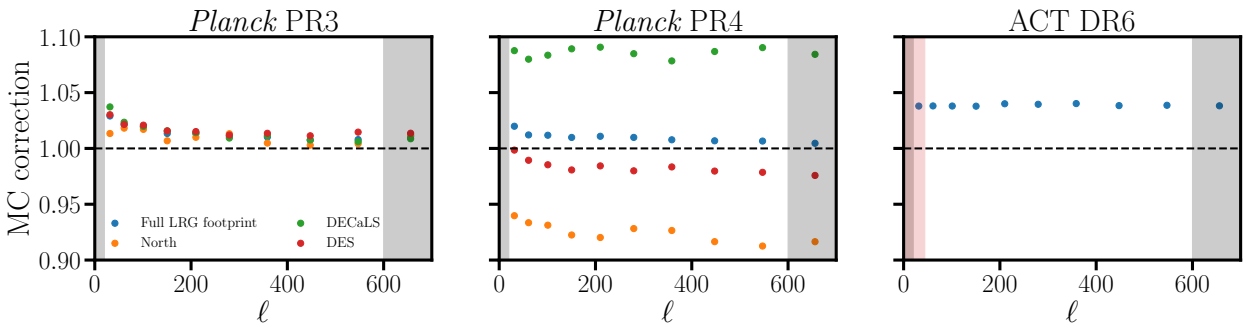


Figure D.1: Monte-Carlo normalization corrections for the baseline PR3 (left), PR4 (middle) and DR6 (right) CMB lensing reconstructions for different LRG masks. For PR3 and PR4 we show the normalization corrections for the full LRG footprint (blue) and the intersection of the LRG footprint with the Northern (orange), DECaLS (green) and DES (red) imaging regions. Note in particular the large $\mathcal{O}(10\%)$ variations for PR4, which arise from the optimal anisotropic filtering adopted by [119].

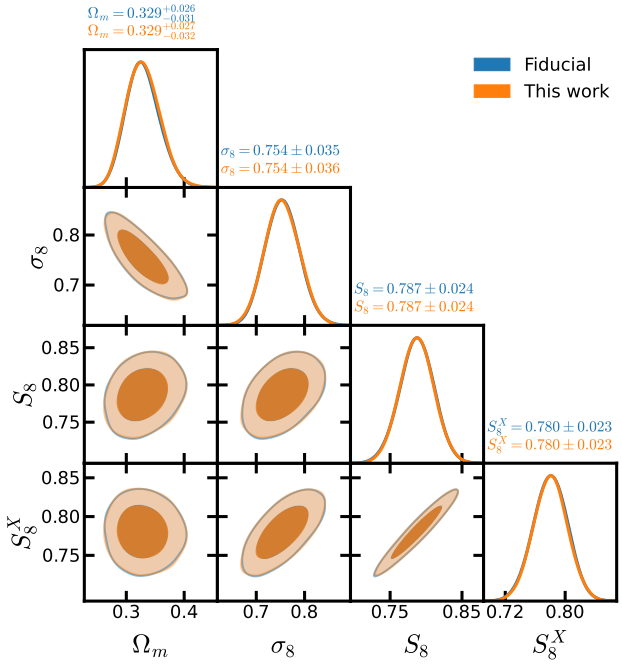


Figure D.2: Linear theory constraints from PR4+DR6 using our fiducial measurements, window functions, and hybrid covariance vs those obtained following §5.3 and Appendix D.3.

D.4 Analytic marginalization and Bayesian PTEs

Analytic marginalization: Suppose that our theory prediction $\boldsymbol{t}(\boldsymbol{\theta},\boldsymbol{\phi}) \equiv \boldsymbol{A}(\boldsymbol{\theta}) + \sum_i \phi_i \boldsymbol{B}_i(\boldsymbol{\theta})$ depends generically on a set of parameters $\boldsymbol{\theta}$ and linearly on a set of parameters $\boldsymbol{\phi}$. We assume that the likelihood and priors on linear parameters are Gaussian, in which case the posterior is

$$P(\boldsymbol{\theta},\boldsymbol{\phi}|\boldsymbol{d}) \propto e^{-\frac{1}{2}(\boldsymbol{t}-\boldsymbol{d})^T \boldsymbol{C}_d^{-1}(\boldsymbol{t}-\boldsymbol{d})}\ e^{-\frac{1}{2}(\boldsymbol{\phi}-\boldsymbol{\mu})^T \boldsymbol{C}_\phi^{-1}(\boldsymbol{\phi}-\boldsymbol{\mu})}\ P(\boldsymbol{\theta}), \tag{D.4}$$

where $\boldsymbol{d}$ is the data, $\boldsymbol{C}_d$ is the data covariance and $\boldsymbol{\mu}$ ($\boldsymbol{C}_\phi$) is the mean (covariance) of the prior on the linear parameters. Defining $\boldsymbol{\delta} \equiv \boldsymbol{A} - \boldsymbol{d}$, $(\boldsymbol{M}^{-1})_{ij} \equiv \boldsymbol{B}_i^T \boldsymbol{C}_d^{-1} \boldsymbol{B}_j + (\boldsymbol{C}_\phi^{-1})_{ij}$, $V_i \equiv \boldsymbol{B}_i^T \boldsymbol{C}_d^{-1}\boldsymbol{\delta} - (\boldsymbol{C}_\phi^{-1}\boldsymbol{\mu})_i$ and $\bar{\boldsymbol{\phi}} \equiv -\boldsymbol{M}\boldsymbol{V}$, we may re-express the posterior as

$$P(\boldsymbol{\theta},\boldsymbol{\phi}|\boldsymbol{d}) \propto e^{-\frac{1}{2}\boldsymbol{\delta}^T \boldsymbol{C}_d^{-1}\boldsymbol{\delta}}\ e^{\frac{1}{2}\boldsymbol{V}^T\boldsymbol{M}\boldsymbol{V}}\ e^{-\frac{1}{2}\left(\boldsymbol{\phi}-\bar{\boldsymbol{\phi}}\right)^T \boldsymbol{M}^{-1}\left(\boldsymbol{\phi}-\bar{\boldsymbol{\phi}}\right)}\ P(\boldsymbol{\theta}), \tag{D.5}$$

where we have absorbed all terms independent of $(\boldsymbol{\theta},\boldsymbol{\phi})$ into a proportionality constant. We perform the (Gaussian) integral over the linear parameters analytically to get the marginal posterior

$$P_m(\boldsymbol{\theta}|\boldsymbol{d}) \equiv \int d\boldsymbol{\phi}\ P(\boldsymbol{\theta},\boldsymbol{\phi}|\boldsymbol{d}) \propto e^{-\frac{1}{2}\boldsymbol{\delta}^T \boldsymbol{C}_d^{-1}\boldsymbol{\delta}}\ e^{\frac{1}{2}\boldsymbol{V}^T\boldsymbol{M}\boldsymbol{V}}\ \sqrt{\det(\boldsymbol{M})}\ P(\boldsymbol{\theta}). \tag{D.6}$$

The linear parameters in our fiducial analysis include shot noise and counterterms, which we marginalize over analytically using Eq. (D.6). When maximizing the posterior we set the linear parameters to their best fit values ($\boldsymbol{\phi}_{\text{best-fit}} = \bar{\boldsymbol{\phi}} = -\boldsymbol{M}V$) for each set of generic (non-linear) parameters $\boldsymbol{\theta}$. As an explicit check of our numerical implementation of Eq. (D.6) we consider fitting to the z_3 model prediction (see §5.5) with and without analytic marginalization, where for the latter we directly sample shot noise and counterterms. We adopt the fiducial HEFT model (§5.4) with a PR4-like covariance using the priors listed in Table 5.3 with the following exceptions: we place a uniform $\mathcal{U}(-5,5)$ prior on b_s and sample α_x directly (instead of ϵ) with a wide $\mathcal{N}(0,50)$ prior. These results are summarized in Fig. D.3. We find excellent agreement between the posteriors of non-linear parameters.

Bayesian PTE: The Bayesian probability to exceed (PTE) is defined as the probability that a random data realization ($\boldsymbol{d}^{\text{rand}}$) will have a larger average[1] χ^2 than the observed data ($\boldsymbol{d}$) [246]:

$$\begin{aligned}
\text{PTE} &\equiv \text{Prob}\left(\chi^2(\boldsymbol{\theta},\boldsymbol{\phi},\boldsymbol{d}^{\text{rand}}) \geq \chi^2(\boldsymbol{\theta},\boldsymbol{\phi},\boldsymbol{d})\,|\,\boldsymbol{d}\right) \\
&= \int d\boldsymbol{\theta}\ d\boldsymbol{\phi}\ d\boldsymbol{d}^{\text{rand}}\ H\left(\chi^2(\boldsymbol{\theta},\boldsymbol{\phi},\boldsymbol{d}^{\text{rand}}) - \chi^2(\boldsymbol{\theta},\boldsymbol{\phi},\boldsymbol{d})\right) P(\boldsymbol{\theta},\boldsymbol{\phi},\boldsymbol{d}^{\text{rand}}|\boldsymbol{d}) \\
&\text{with}\quad \chi^2(\boldsymbol{\theta},\boldsymbol{\phi},\boldsymbol{d}) = \left(\boldsymbol{t}(\boldsymbol{\theta},\boldsymbol{\phi}) - \boldsymbol{d}\right)^T \boldsymbol{C}_d^{-1} \left(\boldsymbol{t}(\boldsymbol{\theta},\boldsymbol{\phi}) - \boldsymbol{d}\right)
\end{aligned} \tag{D.7}$$

[1] Bayesian PTEs can be defined for any test-statistic (e.g. χ^4). We take the test statistic to be the "standard" χ^2.

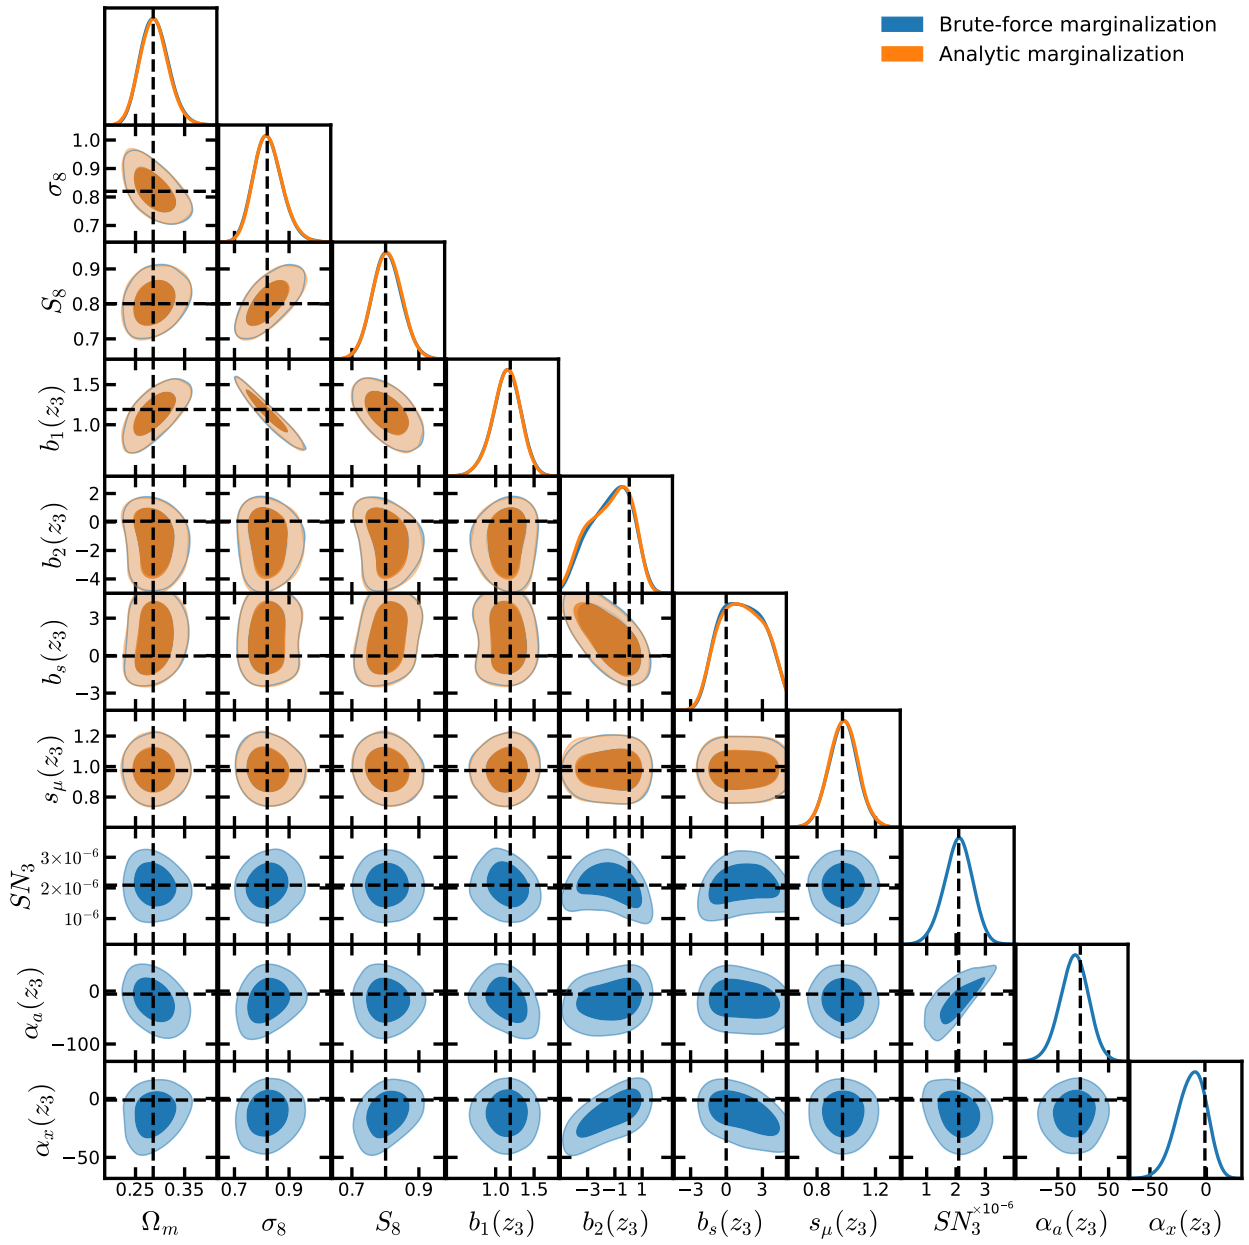


Figure D.3: A comparison of analytic (blue) and brute-force marginalization (orange) over shot noise and counterterms. The true values are indicated by black dashed lines. We fit to the z_3 model prediction (see §5.5) using our fiducial HEFT model and a PR4-like covariance. In this figure we have adopted wider priors for b_s and α_x than used in our fiducial analysis (see text of §D.4).

where $H(x)$ is the Heaviside step function and $P(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d}^{\rm rand}|\boldsymbol{d}) = P(\boldsymbol{d}^{\rm rand}|\boldsymbol{\theta}, \boldsymbol{\phi})P(\boldsymbol{\theta}, \boldsymbol{\phi}|\boldsymbol{d})$ is

the joint distribution of the random data realization and model parameters $(\boldsymbol{\theta}, \boldsymbol{\phi})$ given a measurement of the data $(\boldsymbol{d})$.

Given that the likelihood $P(\boldsymbol{d}^{\rm rand}|\boldsymbol{\theta}, \boldsymbol{\phi})$ is a Gaussian with mean $\boldsymbol{t}(\boldsymbol{\theta}, \boldsymbol{\phi})$ and (presumed known) covariance $\boldsymbol{C}_d$, the quantity $\chi^2(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d}^{\rm rand})$ is drawn from a χ^2 distribution with N_d (number of data points) degrees of freedom for a fixed set of $(\boldsymbol{\theta}, \boldsymbol{\phi})$. The probability then that this quantity exceeds the measured value $\chi^2(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d})$ is given by the PTE $(= 1 - \text{CDF})$ of a χ^2 distribution with N_d degrees of freedom:

$$\int d\boldsymbol{d}^{\rm rand}\ H\big(\chi^2(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d}^{\rm rand}) - \chi^2(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d})\big) P(\boldsymbol{d}^{\rm rand}|\boldsymbol{\theta}, \boldsymbol{\phi}) = 1 - \frac{1}{\Gamma(N_d/2)}\gamma\left(\frac{N_d}{2}, \frac{\chi^2(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d})}{2}\right) \tag{D.8}$$

where $\gamma(s, x) = \int_0^x t^{s-1}e^{-t}dt$ is the incomplete gamma function and $\Gamma(s) = \lim_{x\to\infty}\gamma(s, x)$.

Plugging this result in to the definition of the Bayesian PTE we have:

$$\text{PTE} = \int d\boldsymbol{\theta}\ d\boldsymbol{\phi}\left[1 - \frac{1}{\Gamma(N_d/2)}\gamma\left(\frac{N_d}{2}, \frac{\chi^2(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{d})}{2}\right)\right] P(\boldsymbol{\theta}, \boldsymbol{\phi}|\boldsymbol{d}). \tag{D.9}$$

The remaining integrals over the model parameters are performed numerically. Analytic marginalization complicates the numerical integration slightly. For each step in the chain (fixed $\boldsymbol{\theta}$) we integrate over the linear parameters $\boldsymbol{\phi}$ via Monte-Carlo, where each $\boldsymbol{\phi}$ is drawn from a Gaussian distribution with mean $\bar{\boldsymbol{\phi}}$ and covariance $\boldsymbol{M}$ (see Eq. D.5). We integrate over non-linear parameters by averaging the PTEs (averaged over $\boldsymbol{\phi}$) from random elements of the chain.

D.5 Supplemental constraints

We show the full set of parameter constraints from our baseline PR4+DR6 fit in Fig. D.4.

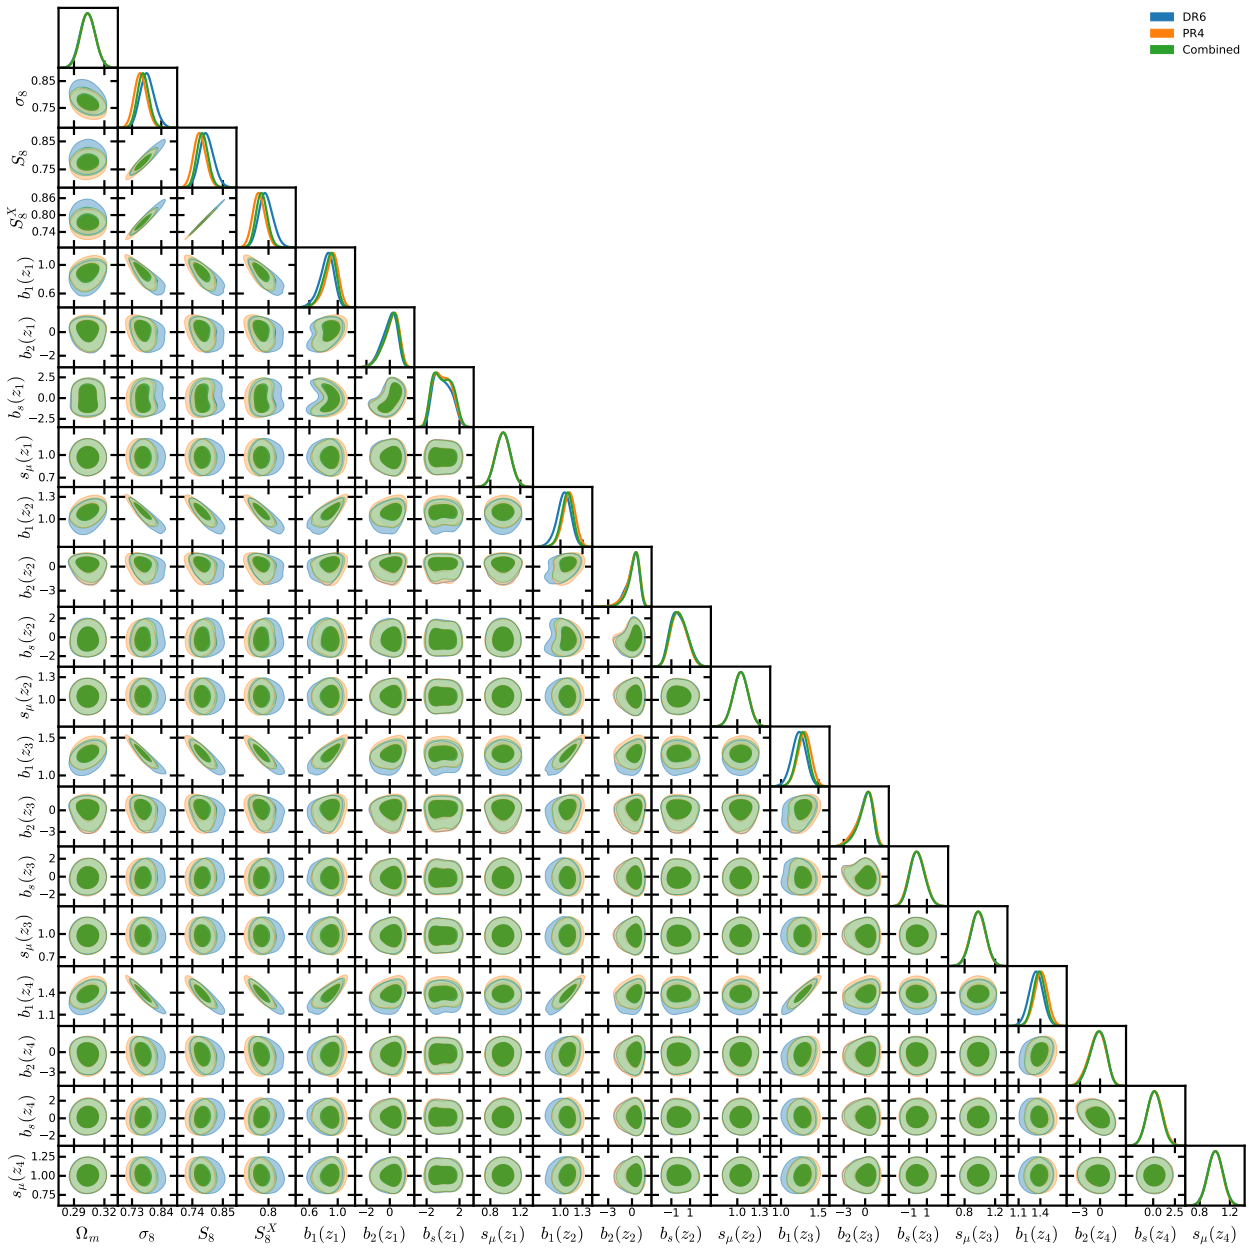


Figure D.4: Parameter constraints from the baseline PR4+DR6 fit.

Appendix E

Baryon distribution I

E.1 Response functions of more general stacked estimators

We define a more general stacked patchy screening estimator as weighted average of a filtered CMB map centered on a set of galaxy positions:

$$\widehat{\mathcal{T}}(\boldsymbol{r}) \equiv -\frac{1}{N_g\langle \mathcal{W}[T_l]T_l\rangle}\sum_{i=1}^{N_g}\mathcal{W}[T_l(\xi\boldsymbol{r}+\boldsymbol{r}_i)]T_s(\boldsymbol{r}+\boldsymbol{r}_i), \tag{E.1}$$

where $\boldsymbol{r}_i$ is the position of the i'th galaxy, $T_l(\boldsymbol{r}) = \int_{\boldsymbol{\ell}} e^{i\boldsymbol{\ell}\cdot\boldsymbol{r}}W_\ell^l T_{\boldsymbol{\ell}}$ is a filtered CMB map with isotropic (low pass) filter W_ℓ^l and likewise for $T_s(\boldsymbol{r})$, while $\mathcal{W}[T]$ is a general weighting function. As in the main text, ξ is a number specifying where the long map should be evaluated relative to the short map. Provided that the weighting function has a one-dimensional Fourier transform $\mathcal{W}_\omega$ we can rewrite

$$\widehat{\mathcal{T}}(\boldsymbol{r}) = -\frac{1}{N_g\langle \mathcal{W}[T_l]T_l\rangle}\sum_{i=1}^{N_g}\int\frac{d\omega}{2\pi}\mathcal{W}_\omega e^{i\omega T_l(\xi\boldsymbol{r}+\boldsymbol{r}_i)}T_s(\boldsymbol{r}+\boldsymbol{r}_i). \tag{E.2}$$

To compute the response of $\widehat{\mathcal{T}}(\boldsymbol{r})$ to lensing we average over primary CMB realizations in the presence of a fixed late time structure realization, which we denote as $\langle\widehat{\mathcal{T}}(\boldsymbol{r})\rangle'$. We Taylor expand the exponential and repeatedly Fourier transform to find

$$\left\langle e^{i\omega T_l(\xi\boldsymbol{r}+\boldsymbol{r}_i)}T_s(\boldsymbol{r}+\boldsymbol{r}_i)\right\rangle' = \sum_{n=0}^{\infty}\frac{(i\omega)^{2n+1}}{(2n+1)!}\int_{\boldsymbol{\ell}\boldsymbol{\ell}_1\cdots\boldsymbol{\ell}_{2n+1}} e^{i\boldsymbol{\ell}\cdot(\boldsymbol{r}+\boldsymbol{r}_i)}e^{i(\boldsymbol{\ell}_1+\cdots+\boldsymbol{\ell}_{2n+1})\cdot(\xi\boldsymbol{r}+\boldsymbol{r}_i)}\langle T_{\boldsymbol{\ell}}^s T_{\boldsymbol{\ell}_1}^l\cdots T_{\boldsymbol{\ell}_{2n+1}}^l\rangle', \tag{E.3}$$

where we neglect terms with an odd number of T's, as these terms vanish in the partial average. Note that for a fixed lensing realization the lensed CMB $T_{\boldsymbol{\ell}}$ is still a Gaussian random field, just not a statistically isotropic one, so we can use Wick's theorem to reduce $\langle T_{\boldsymbol{\ell}}T_{\boldsymbol{\ell}_1}\cdots T_{\boldsymbol{\ell}_{2n+1}}\rangle'$ to something like $\sim\langle T^2\rangle^{n+1}$. Note that only terms that are fully symmetric

in the last $2n+1$ indices contribute to Eq. (E.3). The symmetric contribution to the $2n+2$ point function is given by

$$\begin{aligned}\langle T_{\boldsymbol{\ell}} T_{(\boldsymbol{\ell}_1} \cdots T_{\boldsymbol{\ell}_{2n+1})}\rangle' &= (2n+1)\langle T_{\boldsymbol{\ell}} T_{(\boldsymbol{\ell}_1)}\rangle' \langle T_{\boldsymbol{\ell}_2} \cdots T_{\boldsymbol{\ell}_{2n+1})}\rangle' \\ &= \frac{(2n+1)!}{2^n n!}\langle T_{\boldsymbol{\ell}} T_{(\boldsymbol{\ell}_1)}\rangle' \langle T_{\boldsymbol{\ell}_2} T_{\boldsymbol{\ell}_3}\rangle' \cdots \langle T_{\boldsymbol{\ell}_{2n}} T_{\boldsymbol{\ell}_{2n+1})}\rangle'.\end{aligned} \tag{E.4}$$

Substituting this in to Eq. (E.3) and resumming gives

$$\left\langle e^{i\omega T_l(\xi \boldsymbol{r}+\boldsymbol{r}_i)} T_s(\boldsymbol{r}+\boldsymbol{r}_i)\right\rangle' = i\omega \int_{\boldsymbol{\ell}\boldsymbol{\ell}'} e^{i(\boldsymbol{\ell}+\boldsymbol{\ell}')\cdot\boldsymbol{r}_i} e^{i(\boldsymbol{\ell}+\xi\boldsymbol{\ell}')\cdot\boldsymbol{r}} \langle T^s_{\boldsymbol{\ell}} T^l_{\boldsymbol{\ell}'}\rangle' \,\mathrm{Exp}\left[-\tfrac{\omega^2}{2}\int_{\boldsymbol{\ell}\boldsymbol{\ell}'} e^{i(\boldsymbol{\ell}+\boldsymbol{\ell}')\cdot(\xi\boldsymbol{r}+\boldsymbol{r}_i)}\langle T^l_{\boldsymbol{\ell}} T^l_{\boldsymbol{\ell}'}\rangle'\right]. \tag{E.5}$$

A nearly identical calculation can be carried out for the denominator:

$$\langle \mathcal{W}[T_l] T_l\rangle = \int d\omega \mathcal{W}_\omega \langle e^{i\omega T_l} T_l\rangle / 2\pi.$$

To compute $\langle e^{i\omega T_l} T_l\rangle$, take Eq. (E.5) and replace T^s with T^l, the partial average $\langle\cdots\rangle'$ with a full ensemble average $\langle\cdots\rangle$, and set $\xi = 1$. The integrals inside and outside the exponential are identical and simplify to $\langle T_l^2\rangle$, giving

$$\langle \mathcal{W}[T_l(\boldsymbol{r})] T_l(\boldsymbol{r})\rangle = \langle T_l^2\rangle \int \frac{d\omega}{2\pi} e^{-\frac{\omega^2}{2}\langle T_l^2\rangle} i\omega \mathcal{W}_\omega. \tag{E.6}$$

Putting this all together gives

$$\langle \widehat{\mathcal{T}}(\boldsymbol{r})\rangle' = -\tfrac{1}{N_g\langle T_l^2\rangle}\sum_{i=1}^{N_g}\int_{\boldsymbol{\ell}\boldsymbol{\ell}'} e^{i(\boldsymbol{\ell}+\boldsymbol{\ell}')\cdot\boldsymbol{r}_i} e^{i(\boldsymbol{\ell}+\xi\boldsymbol{\ell}')\cdot\boldsymbol{r}} \langle T^s_{\boldsymbol{\ell}} T^l_{\boldsymbol{\ell}'}\rangle' \frac{\int d\omega\, \omega\, \mathcal{W}_\omega \,\mathrm{Exp}\left[-\frac{\omega^2}{2}\int_{\boldsymbol{\ell}\boldsymbol{\ell}'} e^{i(\boldsymbol{\ell}+\boldsymbol{\ell}')\cdot(\xi\boldsymbol{r}+\boldsymbol{r}_i)}\langle T^l_{\boldsymbol{\ell}} T^l_{\boldsymbol{\ell}'}\rangle'\right]}{\int d\omega\, \omega\, \mathcal{W}_\omega \,\mathrm{Exp}\left[-\frac{\omega^2}{2}\langle T_l^2\rangle\right]}. \tag{E.7}$$

Perturbative expansion

To compute the response of $\widehat{\mathcal{T}}$ to lensing at a given order, we Taylor expand $\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle'$ around a fixed lensing realization

$$\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle' = (2\pi)^2\delta^D_{\boldsymbol{L}} C^0_\ell + f^\kappa_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}\,\kappa_{\boldsymbol{L}} + \sum_{n=1}^{\infty}\int_{\boldsymbol{\ell}_1\cdots\boldsymbol{\ell}_n} f^{\kappa^{n+1}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}_1,\cdots,\boldsymbol{\ell}_n}\left(\kappa_{\boldsymbol{\ell}_1}\cdots\kappa_{\boldsymbol{\ell}_n}\kappa_{\boldsymbol{L}-\boldsymbol{\ell}_{1n}}\right), \tag{E.8}$$

and collect all terms in Eq. (E.7) at a given order in κ, where $f^{\kappa^{n+1}}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}_1,\cdots,\boldsymbol{\ell}_n}$ defines the n'th order response function and we've adopted the shorthand $\boldsymbol{\ell}_{1n} \equiv \sum_{i=1}^n \boldsymbol{\ell}_i$. The higher order

response functions (f^{κ^n} for $n > 1$) can be found by Taylor expanding the the lensed CMB in powers of the convergence

$$T_{\boldsymbol{\ell}} = T^0_{\boldsymbol{\ell}} + \sum_{n=1}^{\infty} \int_{\boldsymbol{\ell}_1 \cdots \boldsymbol{\ell}_n} K^{(n)}_{\boldsymbol{\ell},\boldsymbol{\ell}_1,\cdots,\boldsymbol{\ell}_n} \kappa_{\boldsymbol{\ell}_1} \cdots \kappa_{\boldsymbol{\ell}_n} T^0_{\boldsymbol{\ell}-\boldsymbol{\ell}_{1n}} \quad \text{where} \quad K^{(n)}_{\boldsymbol{\ell},\boldsymbol{\ell}_1,\cdots,\boldsymbol{\ell}_n} = \frac{(-2)^n}{n!} \prod_{i=1}^{n} \frac{\boldsymbol{\ell}_i}{\ell_i^2} \cdot (\boldsymbol{\ell} - \boldsymbol{\ell}_{1n}). \tag{E.9}$$

For example, the quadratic response function f^{κ^2} is given by

$$\begin{aligned} f^{\kappa^2}_{\boldsymbol{\ell}_1,\boldsymbol{\ell}_2,\boldsymbol{\ell}_3} &= K^{(2)}_{\boldsymbol{\ell}_1,\boldsymbol{\ell}_1+\boldsymbol{\ell}_2-\boldsymbol{\ell}_3,\boldsymbol{\ell}_3} C^0_{\ell_2} + \frac{1}{2} K^{(1)}_{\boldsymbol{\ell}_1,\boldsymbol{\ell}_3} K^{(1)}_{\boldsymbol{\ell}_2,\boldsymbol{\ell}_1+\boldsymbol{\ell}_2-\boldsymbol{\ell}_3} C^0_{|\boldsymbol{\ell}_1-\boldsymbol{\ell}_3|} + (\boldsymbol{\ell}_1 \leftrightarrow \boldsymbol{\ell}_2) \\ f^{\kappa^2}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}'} &= \frac{2\ell'_i (\boldsymbol{L}-\boldsymbol{\ell}')_j}{\ell'^2 |\boldsymbol{L}-\boldsymbol{\ell}'|^2} \left[\ell_i \ell_j C^0_\ell - (\boldsymbol{\ell}-\boldsymbol{\ell}')_i (\boldsymbol{\ell}-\boldsymbol{\ell}')_j C^0_{|\boldsymbol{\ell}-\boldsymbol{\ell}'|} \right] + (\boldsymbol{\ell} \leftrightarrow \boldsymbol{L}-\boldsymbol{\ell}) \end{aligned} \tag{E.10}$$

where C^0_ℓ is the power spectrum of the unlensed CMB and summation over i, j is implicit.

- **Linear order:** Expanding Eq. (E.7) up to linear order in κ gives

$$\big\langle \widehat{\mathcal{T}}(\boldsymbol{r}) \big\rangle' = \overline{\mathcal{T}}(r) + \frac{1}{N_g} \sum_{i=1}^{N_g} \int_{\boldsymbol{L}} e^{i\boldsymbol{L}\cdot\boldsymbol{r}_i} \, \mathcal{R}^{\mathcal{T}\kappa}_{\boldsymbol{L}}(\boldsymbol{r}) \, \kappa_{\boldsymbol{L}} \tag{E.11}$$

  where $\overline{\mathcal{T}}(r) = -\big\langle T_l(\xi\boldsymbol{r}) T_s(\boldsymbol{r}) \big\rangle / \langle T_l^2 \rangle$ is the (full-sky) mean field. The linear response function is

$$\mathcal{R}^{\mathcal{T}\kappa}_{\boldsymbol{L}}(\boldsymbol{r}) = \frac{e^{i\xi\boldsymbol{L}\cdot\boldsymbol{r}}}{\langle T_l^2 \rangle} \int_{\boldsymbol{\ell}} \Bigg[\underbrace{\frac{\overline{\mathcal{T}}(r)}{2} \frac{d \ln \mathcal{Z}[T_l^{\rm rms}]}{d \ln T_l^{\rm rms}} W^l_{|\boldsymbol{L}-\boldsymbol{\ell}|} W^l_{\ell}}_{\text{vanishes in practice}} - e^{i(1-\xi)(\boldsymbol{L}-\boldsymbol{\ell})\cdot\boldsymbol{r}} W^s_{|\boldsymbol{L}-\boldsymbol{\ell}|} W^l_{\ell} \Bigg] f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}}, \tag{E.12}$$

  where we defined $\mathcal{Z}[T] \equiv \int d\omega \, \omega \, \mathcal{W}_\omega \, \mathrm{Exp}[-(\omega T)^2/2]$ and $T_l^{\rm rms} = \sqrt{\langle T_l^2 \rangle}$. In practice we assume the long and short filters are disjoint (i.e. $W^l_\ell W^s_\ell = 0$) so that the full-sky mean field vanishes ($\overline{\mathcal{T}} = 0$). In this case the response of $\widehat{\mathcal{T}}$ to lensing is independent of the weight function $\mathcal{W}[T]$. Note that this derivation is agnostic to the functional form of the lensing response function (f^κ). It immediately follows that the general stacked estimator's (Eq. E.1) linear response to any extragalactic foreground is independent of the weighting function provided that the full-sky mean field vanishes.

  Note that in principle the linear response is a function of the vectors $(\boldsymbol{r}, \boldsymbol{L})$. This is relevant for oriented stacking, however, in this chapter we assume (uniform) random orientations, in which case the response can be replaced with its angle average, making it a function of the scalars (r, L).

  We note that the linear order term can be isolated from simulations by lensing either the short or long map to linear order, leaving the remaining map unlensed, and averaging over the resulting two profiles.

- **Quadratic order:** The quadratic contribution to Eq. (E.7) can be written as

$$\big\langle \widehat{\mathcal{T}}(\boldsymbol{r})\big\rangle' \supset \frac{1}{N_g}\sum_{i=1}^{N_g}\int_{\boldsymbol{L}} e^{i\boldsymbol{L}\cdot\boldsymbol{r}_i}\int_{\boldsymbol{\ell}} \mathcal{R}^{\mathcal{T}\kappa^2}_{\boldsymbol{L},\boldsymbol{\ell}}(\boldsymbol{r})\;\kappa_{\boldsymbol{\ell}}\,\kappa_{\boldsymbol{L}-\boldsymbol{\ell}} \tag{E.13}$$

where the quadratic response function is

$$\begin{aligned}\mathcal{R}^{\mathcal{T}\kappa^2}_{\boldsymbol{L},\boldsymbol{\ell}'}(\boldsymbol{r}) = \frac{e^{i\xi\boldsymbol{L}\cdot\boldsymbol{r}}}{\langle T_l^2\rangle}\int_{\boldsymbol{\ell}}\Bigg[&\frac{1}{2}\frac{d\ln\mathcal{Z}[T_l^{\rm rms}]}{d\ln T_l^{\rm rms}}\mathcal{R}^{\mathcal{T}\kappa}_{\boldsymbol{\ell}'}(\boldsymbol{r})\,W^l_{|\boldsymbol{L}-\boldsymbol{\ell}'-\boldsymbol{\ell}|}\,W^l_{\ell}\,f^{\kappa}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell}'-\boldsymbol{\ell}}\\ &- e^{i(1-\xi)(\boldsymbol{L}-\boldsymbol{\ell})\cdot\boldsymbol{r}}\,W^s_{|\boldsymbol{L}-\boldsymbol{\ell}|}\,W^l_{\ell}\,f^{\kappa^2}_{\boldsymbol{\ell},\boldsymbol{L}-\boldsymbol{\ell},\boldsymbol{\ell}'}\Bigg],\end{aligned} \tag{E.14}$$

assuming that the mean field vanishes ($\overline{\mathcal{T}}=0$).

E.2 Fourier transform of the signed and thresholded estimator

Consider the signed and thresholded estimator $\mathcal{W}[T] = \mathrm{Sign}[T]H(|T|-T_c)$ where $H(x)$ is the Heaviside step function and T_c is a cutoff temperature [536, 154]. To compute the Fourier transform $\mathcal{W}_\omega = \int_{-\infty}^{\infty} dT\; e^{-i\omega T}\,\mathcal{W}[T]$ we first note that we may write $\mathcal{W}[T] = \lim_{k\to+\infty}\mathcal{W}[k,T]$ where

$$\mathcal{W}[k,T] = \frac{\tanh(k\,T)}{2}\Big(1+\tanh\Big(k\,T\,\tanh(k\,T)-k\,T_c\Big)\Big). \tag{E.15}$$

The function $\mathcal{W}[k,T]$ is meromorphic in T for any finite k. We plot $|\mathcal{W}[k,T]|$ in Fig. E.1 for $k\,T_c = 0.5$. To ensure a convergent result for $\mathcal{W}_\omega$ we bend the contour slightly so that T acquires a negative imaginary component as $|T|\to\infty$. This is illustrated in Fig. E.1, where we take $T\to(1\pm i\epsilon)T$ to shift the contour from the solid to dashed white line. We are always free to make ϵ small enough to dodge any poles for a finite k, leaving us with (for $\omega>0$)

$$\begin{aligned}\mathcal{W}_\omega &= \lim_{k\to+\infty}\lim_{\epsilon\to0^+}\left[\int_{T_c}^{\infty} dT\; e^{-(i+\epsilon)\omega T} - \int_{-\infty}^{-T_c} dT\; e^{-(i-\epsilon)\omega T} + \mathcal{O}(k^{-1})\right]\\ &= \lim_{\epsilon\to0^+}\frac{1}{\omega}\left[\frac{e^{-(i+\epsilon)\omega T_c}}{i+\epsilon}+\frac{e^{(i-\epsilon)\omega T_c}}{i-\epsilon}\right]\\ &= \frac{2}{i\omega}\cos(\omega T_c).\end{aligned} \tag{E.16}$$

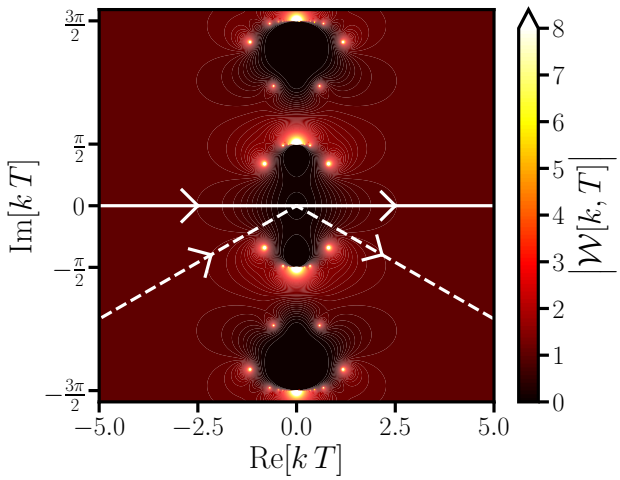


Figure E.1: The pole structure of $\mathcal{W}[k,T]$ for $k\,T_c = 0.5$.

E.3 Mean field bias

Just as for CMB lensing, the stacked patchy screening estimator acquires a non-zero mean when the galaxy and CMB data have partial sky coverage (even if the long and short filters are disjoint). This "mean field" must be subtracted from Eq. (7.5) to obtain an unbiased estimate of the gas profile. The mean field bias can be estimated by simply replacing the galaxy positions with a set of random positions uniformly dispersed within the galaxy mask. To rigorously show this, note that we can rewrite Eq. (7.5) as

$$\widehat{\mathcal{T}}(\boldsymbol{r}) = \frac{-1}{\langle T_l^2\rangle}\int_{\boldsymbol{L}\boldsymbol{\ell}}\left[\frac{1}{N_g}\sum_{i=1}^{N_g} e^{i\boldsymbol{L}\cdot\boldsymbol{r}_i}\right] e^{i(\boldsymbol{L}-(1-\xi)\boldsymbol{\ell})\cdot\boldsymbol{r}} W^s_{|\boldsymbol{L}-\boldsymbol{\ell}|} W^l_\ell T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}. \tag{E.17}$$

The sum over galaxy positions only includes positions within the galaxy mask $M(\boldsymbol{r})$, and is related to the masked density contrast $\tilde{g}$ via

$$\tilde{g}(\boldsymbol{r}) = \frac{A_M}{N_g}\sum_{i=1}^{N_g}\delta^D(\boldsymbol{r}-\boldsymbol{r}_i) - M(\boldsymbol{r}), \tag{E.18}$$

where $A_M = \int d^2\boldsymbol{r} M(\boldsymbol{r})$ is the area within the mask M. Fourier transforming this expression and substituting it in to Eq. (E.17) gives

$$\langle\widehat{\mathcal{T}}(\boldsymbol{r})\rangle = -\frac{1}{A_M\langle T_l^2\rangle}\int_{\boldsymbol{L}\boldsymbol{\ell}} e^{i(\boldsymbol{L}-(1-\xi)\boldsymbol{\ell})\cdot\boldsymbol{r}} W^s_{|\boldsymbol{L}-\boldsymbol{\ell}|} W^l_\ell \Big[\langle\tilde{g}_{-\boldsymbol{L}} T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle + M_{-\boldsymbol{L}}\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle\Big]. \tag{E.19}$$

The first term is the signal while the second term is the mean field bias. Even if the long and short filters are disjoint ($W^l_\ell W^s_\ell = 0$), the mean field bias does not vanish provided that

the galaxy *and* CMB maps have partial sky coverage (in which case $M_{-\boldsymbol{L}}\langle T_{\boldsymbol{\ell}} T_{\boldsymbol{L}-\boldsymbol{\ell}}\rangle \neq 0$ for $L > 0$). To isolate the second term one can choose a set of positions $\{\boldsymbol{r}_i\}$ from a distribution that is independent of the CMB data satisfying $\langle \tilde{g} \rangle = 0$ (e.g. uniform-area sampling within the masked region).

Appendix F

Baryon distribution II

F.1 Simulation test of patchy screening and the effect of CMB lensing

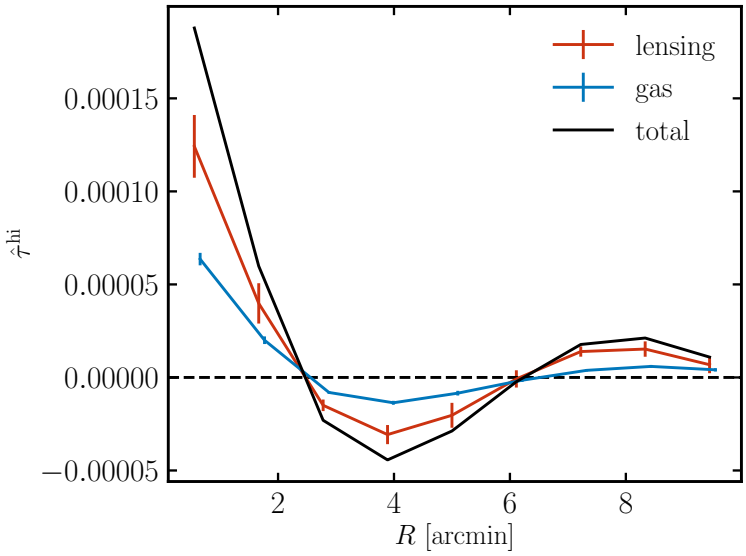


Figure F.1: Simulated high-pass-filtered profiles of the 'patchy screening' effect as implemented in Eq. 8.13. Maps are generated by applying the 'transfer function method' to the AbacusSummit huge box density maps between $z = 0.45$ and 1 to match as closely as possible the LRG sample redshift support (see the text in App. F.1 for more details). Stacks are performed on all halos in that redshift range above $10^{13.07}$ such that the mean halo mass matches that of the LRG sample, $10^{13.3}$ in units of $M_\odot/h$.

In addition to the lensed maps, we also create 'patchy screening'-only maps by generating

a random Gaussian relization from the lensed C_ℓ and then multiplying the thus-obtained CMB temperature map by $\exp[-\tau]$. To get the τ map on the full sky, we adopt the transfer function technique presented in [455, 359], for which the authors show that multiplying the matter field of an N-body simulation by the transfer function:

$$T(k) = \sqrt{\frac{P_{\tau\tau}}{P_{\rm mm}}} \tag{F.1}$$

in Fourier space leads to a new field $\hat{\tau}$ that matches the true τ field of a hydro simulation at the field level down to very small scales. As noted previously [265], Illustris at $z = 0.5$ provides a good fit to the optical depth profiles observed for the LRGs. Thus, we adopt the transfer function from the Illustris simulation when constructing the theory model used in Eq. 8.18.

Next since we work with the HEALPix density sheets rather than the 3D density fields [253], we convolve each sheet with a transfer function in harmonic space evaluated as $T(\ell = k\chi)$, where χ is the comoving distance to the sheet. This finding holds across different feedback models and can be quantified via the cross-correlation coefficient, $r(k)$, between τ and $\hat{\tau}$, which remains above 0.95 until $k \approx 1.25\ h/\mathrm{Mpc}$ for Illustris [437], used in this study.

The result of the comparison is shown in Fig. F.1. We see that the contribution from the gas portion of the signal, i.e. τ, makes up about a third of the signal, with the lensing contributing about two thirds. The amplitude of the signal is likely somewhat exaggerated due to the fact we are not accounting for miscentering in this test performed on halos, and on very small scales ($\sim$1 arcmin), the transfer function method breaks down, as the cross-correlation coefficient between the dark matter field and the gas field in the Illustris simulation becomes small, i.e. $r(k) = P_{\tau m}/\sqrt{P_{\tau\tau}P_{mm}} \ll 1$.

F.2 Comparison of the signal-to-noise

In Fig. F.2, we compare the signal-to-noise per radial bin for our measurements using ACT and DESI data of the 'patchy screening' signal, resulting from τ and CMB lensing, and the kSZ effect, resulting from τ after reconstructing the line-of-sight velocity field. We see that the majority of the signal-to-noise for the high-pass filtered quantities comes at small aperture and decreases towards large ones, where part of the signal has been filtered out. Similarly, the CAP-filtered quantity does not have a lot of sensitivity at large apertures, where the noise is dominated by the primary CMB. Most of its sensitivity comes at small apertures (2-4 arcmin), with the first radial bin being dominated by the ACT beam noise. We stress that the CAP filter cannot be applied to the 'patchy screening' estimator due to leakage of the large-scale CMB fluctuations.

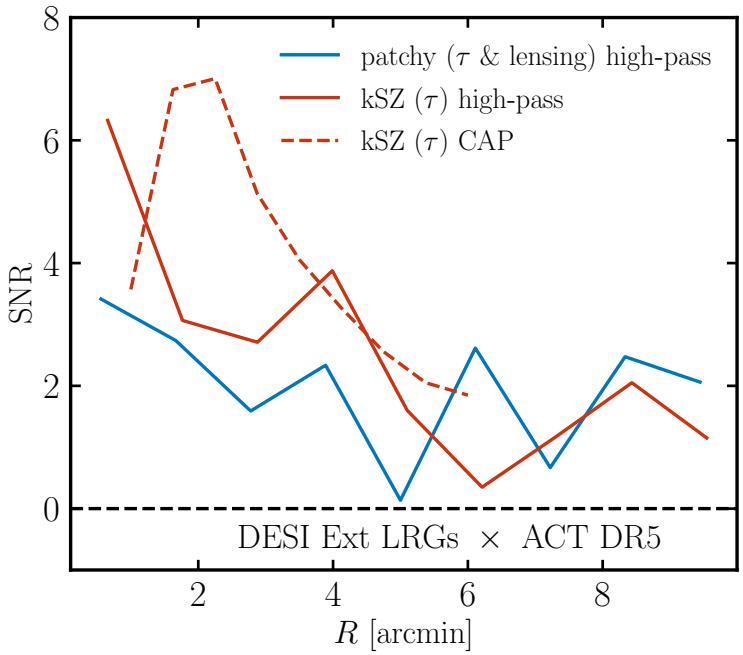


Figure F.2: Comparison of the signal-to-noise per radial bin for our measurements using ACT and DESI data of the 'patchy screening' signal, resulting from τ and lensing, and the kSZ effect, resulting from τ after reconstructing the line-of-sight velocity field.

Appendix G

Cosmology at high redshift

G.1 Primordial figure of merit ($N_{\rm modes}$)

Here we describe a new figure of merit for "counting reconstructable linear modes." This is similar to the effective volume, but integrated to a $k_{\rm max}$ that incorporates the decorrelation of the observed nonlinear field with the initial conditions at high k, and hence loss of primordial information on small scales. This loss of information can be modeled as a Gaussian "propagator" $G(\boldsymbol{k})$. That is, we split the observed nonlinear overdensities $\delta_F(\boldsymbol{k}) = G(\boldsymbol{k})\delta_L(\boldsymbol{k}) + d(\boldsymbol{k})$ into two pieces: one which is correlated with the linear overdensities, $\delta_L(\boldsymbol{k})$, and one which isn't $\langle \delta_L(\boldsymbol{k}) d(\boldsymbol{k}') \rangle = 0$. The Gaussian propogator is given by:

$$G(k,\mu) \equiv \frac{\langle \delta_F \delta_L \rangle}{\langle \delta_L \delta_L \rangle} \simeq \exp\left[-\frac{1}{2} \left(k_\perp^2 + k_\parallel^2 \{1+f\}^2 \right) \Sigma^2 \right], \tag{G.1}$$

where Σ is set by the rms displacement within the Zel'dovich approximation, as in Eq. (10.12).

In the sample variance limit, the relative uncertainty of the linear power spectrum's amplitude A is given by $\sqrt{2/N_{\rm modes}}$. We can calculate this uncertainty from the Fisher matrix:

$$F_{ii} = \frac{f_{\rm sky}}{2} \int_{z_{\rm min}}^{z_{\rm max}} dz \frac{dV}{dz} \int_{k-{\rm wedge}(z)}^{\infty} \frac{d^3\boldsymbol{k}}{(2\pi)^3} \left(\frac{\partial_{\theta_i} P_F(\boldsymbol{k},z)}{P_F(\boldsymbol{k},z) + N(z)} \right)^2, \tag{G.2}$$

where $k-{\rm wedge}(z)$ defines the lower limit of integration. The uncertainty on the amplitude is $\delta A/A = \delta P_L(\boldsymbol{k})/P_L(\boldsymbol{k}) = 1/\sqrt{F_{AA}}$, so that $N_{\rm modes} = 2F_{AA}$. This defines the figure of merit $N_{\rm modes}$.

Let $\hat{A} = P_L(\boldsymbol{k})/\hat{P}_L(\boldsymbol{k})$ be an estimator for the amplitude of the linear power spectrum $P_L(\boldsymbol{k}) = (b + f\mu^2)^2 P_m(k)$, where $\hat{P}_L(\boldsymbol{k})$ is any (unbiased) estimator for the linear galaxy power spectrum in redshift space, and $P_m(k)$ is the linear matter power spectrum. Since $P_F(\boldsymbol{k}) = G^2(\boldsymbol{k}) P_L(\boldsymbol{k}) + P_d(\boldsymbol{k})$ does not explicitly depend on $\hat{P}_L(\boldsymbol{k})$, the derivative of the

non-linear power spectrum with respect to $\hat{A}$ is simply[1] $G^2(\boldsymbol{k})\hat{P}_L(\boldsymbol{k})$. This derivative should be evaluated at the fiducial cosmology, so $\partial_A P_F(\boldsymbol{k}) = G^2(\boldsymbol{k})P_L(\boldsymbol{k})$. Putting this together gives:

$$N_{\rm modes} = f_{\rm sky}\int_{z_{\rm min}}^{z_{\rm max}} dz \frac{dV}{dz}\int_{k-{\rm wedge}(z)}^{\infty}\frac{d^3\boldsymbol{k}}{(2\pi)^3}\left(\frac{G^2(\boldsymbol{k},z)P_L(\boldsymbol{k},z)}{P_F(\boldsymbol{k},z)+N(z)}\right)^2. \tag{G.3}$$

The calculation of the CMB equivalent of $N_{\rm modes}$ is almost identical. We include the TT, TE, and EE power spectra $\boldsymbol{\mu}_\ell = \left(C_\ell^{\rm TT}, C_\ell^{\rm TE}, C_\ell^{\rm EE}\right)^T$ in our Fisher forecast, which has covariance matrix:

$$\boldsymbol{C}_{\ell,\ell'} = \frac{2\delta^K_{\ell,\ell'}}{(2\ell+1)f_{\rm sky}}\begin{pmatrix} \left(C_\ell^{\rm TT}+N_\ell\right)^2 & \left(C_\ell^{\rm TT}+N_\ell\right)C_\ell^{\rm TE} & \left(C_\ell^{\rm TE}\right)^2 \\ \left(C_\ell^{\rm TT}+N_\ell\right)C_\ell^{\rm TE} & \frac{1}{2}\left[\left(C_\ell^{\rm TE}\right)^2+\left(C_\ell^{\rm TT}+N_\ell\right)\left(C_\ell^{\rm EE}+2N_\ell\right)\right] & \left(C_\ell^{\rm EE}+2N_\ell\right)C_\ell^{\rm TE} \\ \left(C_\ell^{\rm TE}\right)^2 & \left(C_\ell^{\rm EE}+2N_\ell\right)C_\ell^{\rm TE} & \left(C_\ell^{\rm EE}+2N_\ell\right)^2 \end{pmatrix}, \tag{G.4}$$

where

$$N_\ell = \Delta_T^2 \exp\left(\frac{\theta^2_{\rm FWHM}\ell^2}{8\ln 2}\right) \tag{G.5}$$

accounts for the instrumental noise. Here Δ_T is the noise level of the experiment and $\theta_{\rm FWHM}$ is the full-width at half-maximum (FWHM) of the beam in radians. Given $\boldsymbol{\mu}_\ell$ and $\boldsymbol{C}_{\ell,\ell'}$, the Fisher matrix is simply:

$$F_{ij} = \frac{1}{2}\sum_{\ell,\ell_1,\ell_2,\ell_3}{\rm Tr}\left[\boldsymbol{C}^{-1}_{\ell,\ell_1}\frac{\partial \boldsymbol{C}_{\ell_1,\ell_2}}{\partial p_i}\boldsymbol{C}^{-1}_{\ell_2,\ell_3}\frac{\partial \boldsymbol{C}_{\ell_3,\ell}}{\partial p_j}\right] + \sum_{\ell,\ell'}\frac{\partial \boldsymbol{\mu}_\ell^T}{\partial p_i}\boldsymbol{C}^{-1}_{\ell,\ell'}\frac{\partial \boldsymbol{\mu}_{\ell'}}{\partial p_j}. \tag{G.6}$$

The second term dominates since the summand is weighted by the number of modes, whereas in the first term all factors of $2\ell+1$ cancel. Just as before, we take $N_{\rm modes} = 2F_{\rm AA}$:

$$N_{\rm modes} \simeq 2\sum_{\ell,\ell'}\boldsymbol{\mu}_\ell^T\boldsymbol{C}^{-1}_{\ell,\ell'}\boldsymbol{\mu}_{\ell'}. \tag{G.7}$$

Due to the bias from foreground contamination, $C_\ell^{\rm TT}$ will realistically be measurable out to $\ell_{\rm max} \simeq 3000$. We enforce this in our calculation of $N_{\rm modes}$ by multiplying $C_\ell^{\rm TT}$ by 10^6 in our covariance matrix when $\ell > 3000$.

Our definition of $N_{\rm modes}$ can be easily generalized to represent a survey's ability to measure the amplitude of the linear matter power spectrum (at $z = 0$) in a specific k-bin. Following the same argument as before, let $\hat{A}(k) = P_m(k)/\hat{P}_m(k)$ be an estimator for the amplitude, where $P_m(k)$ is the linear matter power spectrum at $z = 0$. Then $\partial_{\hat{A}(k')}P_F(\boldsymbol{k},z) = G^2(\boldsymbol{k},z)T^2(k,z)(b+f\mu^2)^2\hat{P}_m(k)\,\Theta(\Delta k-|k-k'|)$, where $T(k,z)$ is the transfer function, $\Theta(x)$ is the Heaviside step function, and Δk is the width of the k-bin. When evaluated at the fiducial cosmology, this yields $\partial_{\hat{A}(k')}P_F(\boldsymbol{k},z) = G^2(\boldsymbol{k},z)P_L(\boldsymbol{k},z)\Theta(\Delta k-|k-k'|)$.

[1] This is not true in a strict mathematical sense. In reality the decorrelated piece $P_d(\boldsymbol{k})$ is still dependent on the initial conditions, although the manner in which it depends on $P_L(\boldsymbol{k})$ in highly non-trivial and subdominant, so we neglect it in our derivative.

Just as before, we interpret $2F_{A(k)A(k)}$ as the effective number of sample variance limited modes that measure the amplitude:

$$N_{\rm modes}(k) = f_{\rm sky}\int_{z_{\rm min}}^{z_{\rm max}} dz \frac{dV}{dz}\frac{k^2\Delta k}{2\pi^2}\int_0^1 d\mu \left(\frac{G^2(\boldsymbol{k},z)P_L(\boldsymbol{k},z)}{P_F(\boldsymbol{k},z)+N(z)}\right)^2. \tag{G.8}$$

We make use this expression in Fig. 10.3, where the error in each bin is $P_m(k)\sqrt{2/N_{\rm modes}(k)}$.

One note on the interpretation of $N_{\rm modes}$ is in order here: we have defined $N_{\rm modes}$ to be the number of modes recoverable that are correlated with the initial conditions. Therefore it is a figure of merit for how well we can constrain the primordial potential or other statistical properties of the initial conditions. $N_{\rm modes}$ as defined here does not take into account the sensitivity of a given observable to the parameters of interest, so it should not be taken as the only comparison metric, especially between high and low-redshift probes (such as CMB vs LSS). For example, it is very possible for an LSS experiment with lower $N_{\rm modes}$ than the CMB to be a lot more sensitive to Dark Energy or neutrino masses, since it performs measurements on a range of redshifts where they are dynamically important.

For this reason, while a direct comparison between CMB and LSS experiments based on $N_{\rm modes}$ alone might be misleading, we believe that it provides a useful metric to compare different survey designs covering a similar redshift range. Since it represents a proxy for "primordial information" available, survey optimization aimed at maximizing it can ensure better sensitivity to the physics of the early Universe.

G.2 Covariance approximations

In the main text we have followed the "usual" approach of approximating the covariance matrix by the Gaussian (disconnected) piece. In Eq. (10.10) we also assume that the survey has a uniform number density in each redshift bin. As emphasized by ref. [611], a more realistic approach would include the impact of window functions, galaxy weights, etc. Inclusion of those effects leads to a covariance with the same functional form

$$C \propto (P^2 + 2P\bar{n}^{-1} + \bar{n}^{-2})/V \tag{G.9}$$

with the substitutions

$$V \to I_{22}^2/I_{44}, \quad \bar{n}^{-1} \to I_{34}/I_{44}, \quad \bar{n}^{-2} \to I_{24}/I_{44}, \tag{G.10}$$

where the I_{ij}'s are integrals over the number density [611]:

$$I_{ij} = \int d^3x\ \bar{n}^i(\mathbf{x})w^j(\mathbf{x}) \qquad \text{with} \quad w(\mathbf{x}) = \frac{1}{1+\bar{n}(\mathbf{x})P_{\rm fid}} \tag{G.11}$$

and $P_{\rm fid}$ a fiducial level of power. Ref. [611] argues that for "realistic" surveys, the ratios of I_{ij} can be numerically quite different than a typical "effective volume" or "constant number density" approximation might lead you to believe. We find that these corrections are

negligible for the majority of our forecasts, as shown in Fig. G.1. The corrections are larger for surveys with smaller number densities. For DESI, the volume correction in the highest redshift bin is 25%, while the number density correction in the lowest redshift bin is 30%. For all other surveys, the correction to the volume is less that 13%, and the correction to the number density is less than 12%. We thus expect these corrections to have a small ($\lesssim 10\%$) impact on our parameter constraints.

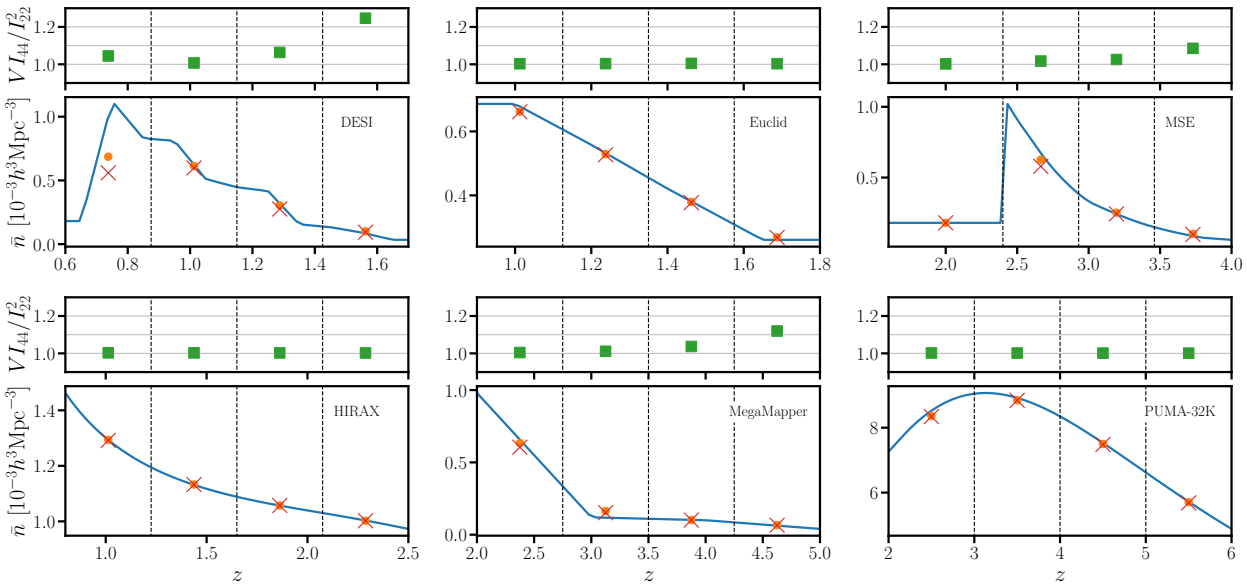


Figure G.1: Corrections to the number densities and volume from the effects considered by ref. [611] for six of the surveys considered in this chapter, when $P_{\rm fid} = 10000\ h^{-3}\mathrm{Mpc}^3$. In blue is the "true" number density $\bar{n}(z)$ of each survey. In orange we plot I_{44}/I_{34} at the center of each redshift bin. In red is $\sqrt{I_{44}/I_{24}}$. In green we plot the ratio VI_{44}/I_{22}^2.

G.3 BAO forecasting

In §10.3 we outlined our procedure for inferring distance constraints from the BAO feature in the power spectrum. In this section we compare our method to the standard procedure[2] of ref. [547] or [186, 235]. In the standard procedure it is assumed that the BAO signal is exponentially damped, but to avoid any information coming from the damping these parameters are not varied in the Fisher matrix. Rather the A-P parameters, $\alpha_\parallel$ and $\alpha_\perp$ are allowed to vary so as to "shift" the BAO template. A scale-independent bias is assumed.

[2]Implemented using: https://www.cfa.harvard.edu/deisenst/acousticpeak/bao_forecast.c

For an isotropic shift, s, the Fisher matrix would look schematically like [547]

$$F_{ss} \approx \frac{V_{\rm surv}}{2} \int \frac{d^3k}{(2\pi)^3} \left(\frac{\partial P_{\rm osc} e^{-k^2\Sigma^2/2}}{\partial \ln s} \right)^2 \frac{1}{(P_L + \bar{n}^{-1})^2} \tag{G.12}$$

where in the numerator $P_{\rm osc}$ represents only the oscillatory, or BAO, piece of the power spectrum and the exponential indicates the damping due to non-linear structure formation. The 2D model is only slightly more complex, with different damping across and along the line of sight and redshift-space distortions included assuming linear theory. To include the effects of reconstruction the damping parameters (Σ) are reduced by ≈ 0.5. By contrast our method uses the Zeldovich-based model of ref. [134]. This model automatically accounts for the differences in damping of the different components, more complex bias, non-linear redshift-space distortions and the manner in which reconstruction reduces the damping. By using a forward model of the reconstruced power spectrum, our forecast is of the standard Fisher form wherein a parameter error is forecast through the theory derivative rather than requiring any modification of the procedure.

However, our method is actually quite close (within $\sim 15\%$) to the standard approach in practice, as shown in Fig. G.2. The damping of the BAO signal by non-linear evolution is very close to Gaussian, as assumed in the standard approach, and while reconstruction affects each of the damping terms differently along and across the line of sight the major impact is to reduce the scale of the damping by approximately a factor of 2 (indeed it was a perturbative analysis of BAO reconstruction that suggested this approximation).

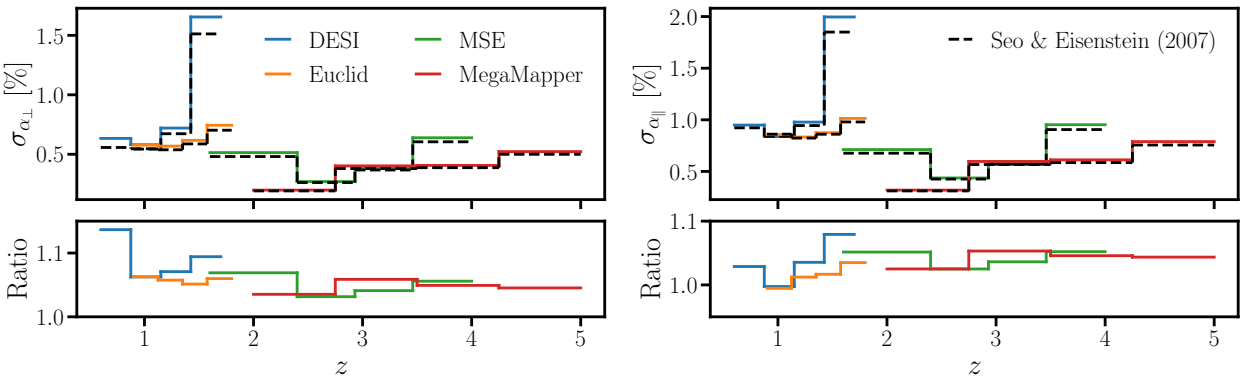


Figure G.2: In the top two pannels we compare our distance constraints (solid), calculated using the method of §10.3, to those found using the standard procedure of ref. [547] (dashed). The ratio of the two methods is shown in the bottom two panels.

G.4 Lensing auto- vs. cross-correlation

Here we comment on the information content of the CMB lensing auto power spectrum, vs the cross-correlation with galaxies. Given a spectroscopic survey with sufficiently high number densities that covers the majority of the redshift range where the CMB lensing kernel is large, one can in principle reconstruct $C_\ell^{\kappa\kappa}$ from the galaxy auto-correlation and galaxy-convergence cross-correlation. This is easily seen in linear theory for a noiseless galaxy survey, so that the number density $n \to \infty$. In this limit we can make the redshift bins arbitrarily small, and approximate the integral expression for $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$ (Eq. (10.14)) at a single comoving distance χ_i. Under these approximations, $C_\ell^{\kappa\kappa}$ is related to $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$ via

$$C_\ell^{\kappa\kappa} = \sum_i \left(C_\ell^{\kappa g_i}\right)^2 / C_\ell^{g_i g_i}, \tag{G.13}$$

where i runs over redshift bins covering $0 < z < z_\star$. Thus $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$ carries at least the same information as $C_\ell^{\kappa\kappa}$ in this idealized case. Since $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$ can in principle be measured to a higher signal-to-noise than $C_\ell^{\kappa\kappa}$, and can be used to measure time evolution, one might naively expect that parameter constraints are driven by $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$. However, in our forecasts we find this not to be the case, as illustrated in Fig. G.3. The contours from an SO measurement of the convergence auto-correlation (red) are significantly tighter than the contours from galaxy auto- and cross-correlations, as measured by SO and our idealized PUMA-32K survey (light blue).

There are multiple ways in which this simplistic argument goes wrong: (1) the assumption that $C_\ell^{\kappa g_i}$, $C_\ell^{g_i g_i}$ can be measured to a high signal-to-noise ratio depends on the details of the survey, (2) a redshift survey might not cover the entire region where the lensing kernel is large, and (3) forecasts that include galaxy auto- and cross-correlations marginalize over galaxy biases, counterterms, and a Poisson-like stochastic contribution; however, including $C_\ell^{\kappa\kappa}$ in a forecast introduces no new nuisance parameters. This is by no means an exhaustive list, since Eq. (G.13) is no longer true for a survey with finite noise, redshift bins of finite extent, when magnification bias is included, or when modeling the tracer beyond linear theory, etc.

To test which of these effects causes the largest degradation to constraints from $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$, we consider idealized versions of the HIRAX and PUMA-32K surveys which have no foreground wedge and are assumed to measure T_b to infinite precision (so that it is held fixed), dubbed HIRAX-ideal and PUMA-ideal, enabling cross-correlations with CMB lensing. We take HIRAX-ideal and PUMA-ideal as representative (toy-model) low noise surveys at high redshift. As illustrated in Fig. G.3, we see a negligible improvement to parameter constraints when upgrading from HIRAX-ideal (grey) to PUMA-ideal (light blue). Thus (1) is not the problem, and PUMA-ideal has sufficiently low noise levels for high fidelity measurements of $C_\ell^{\kappa g_i}$ and $C_\ell^{g_i g_i}$. We also see that extending PUMA-ideal to $z = 0.2$ (light vs. dark blue) has a small impact on the constraints, suggesting that PUMA's fiducial redshift range ($2 < z < 6$) sufficiently covers the region where the CMB lensing kernel is large, eliminating (2). Instead, the primary culprit is (3). From Fig. G.3 we see that the covariance from the

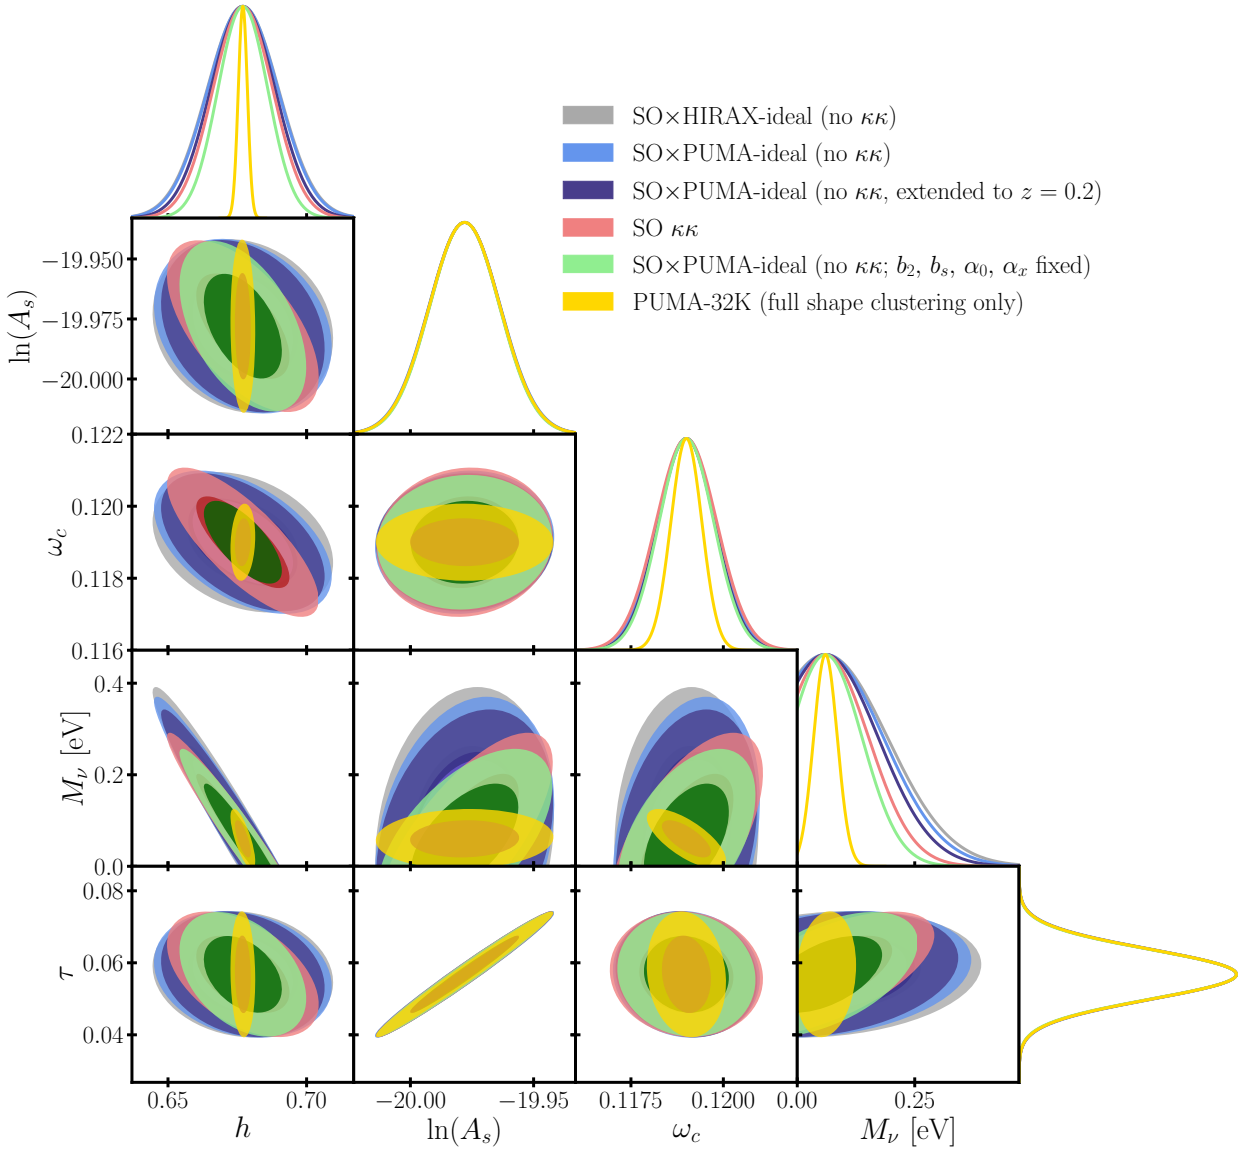


Figure G.3: Constraints from cross-correlations of SO lensing with HIRAX-ideal (greay), PUMA-ideal (light blue), a version of PUMA-ideal that extends down to $z = 0.2$ (dark blue), and PUMA-ideal with non-linear nuisance terms held fixed (green). In red we plot the constraints from the SO lensing convergence autospectrum. To illustrate how the degeneracies from lensing and full shape measurements differ, we plot the constraints from PUMA-32K full shape data in yellow. All contours share the same Planck + SO prior on ΛCDM+M_ν.

convergence auto-correlation (red) is nearly identical to the covariance from the galaxy auto- and cross-correlation (green) when the nuisance terms are fixed.

We should stress that the value of cross-correlations extends considerably beyond constraining standard cosmological parameters. While $C_\ell^{\kappa\kappa}$ depends on the integrated matter fluctuations over a wide redshift range, as discussed in Sec. 10.4, they allow for a tomographic reconstruction of the amplitude of structure $\sigma_8(z)$ that only depends on the local amplitude. This can provide important model-independent tests beyond the rigidity of ΛCDM, considerably opening up the discovery space.

G.5 Empirical checks

As discussed in §10.3, we only integrate over modes where the N_2 term ($N_0\sigma_v^2(k\mu)^2$) is less than 20% of the fiducial power. This "N_2 cut" is only relevant for MSE and MegaMapper, whose number densities are sufficiently low enough for the N_2 term to contribute a non-negligible amount to the fiducial power for $k < k_{\rm nl}$. As we show in Fig. G.4, our results are largely insensitive to the choice of N_2 cutoff. For any cutoff above 4% the constraints from Planck + SO + MegaMapper change by at most 10%.

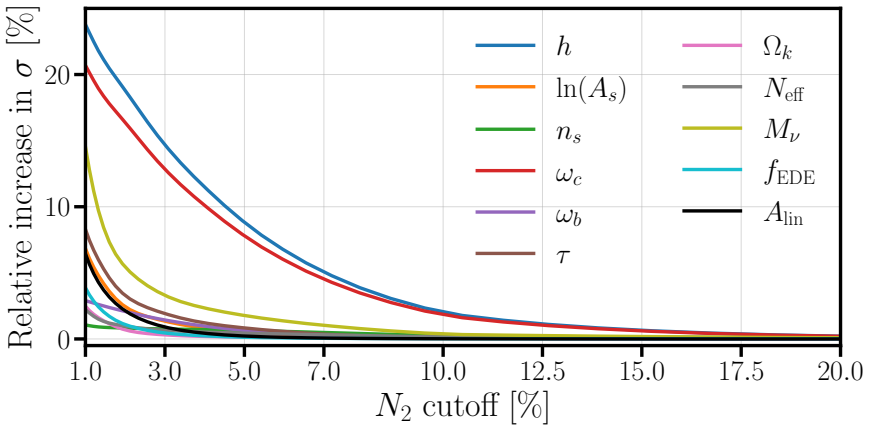


Figure G.4: Planck + SO + MegaMapper (full shape only) constraints as the N_2 cutoff is varied. This is with $\log_{10}(z_c) = 3.61$ and $\omega_{\rm lin} = 147\ h^{-1}\,$Mpc.

To quantify the importance of nuisance parameters' fiducial values, we consider a "MiniMapper" survey that observes LBGs in a single redshift bin covering $2 < z < 3$ with $f_{\rm sky} = 0.34$. We assume that MiniMapper has 100% efficiency for a limiting magnitude $m_{\rm UV} = 24.5$. The default values of MiniMapper's nuisance terms (following the conventions in §10.3) are listed in the first row of Table G.1. The constraints from Planck + MiniMapper on $\{h, \omega_c, \tau, \Omega_k, N_{\rm eff}, M_\nu\}$ with these default values are shown in the last row of Table G.2. The remaining rows of Table G.2 show the constraints when the fiducial value of each of the seven nuisance terms is individually varied, with the new fiducial values listed in the second row of Table G.1. We find that varying the values of $\{\alpha_0, \alpha_2, N_2, \alpha_4\}$ has a negligible ($< 6\%$) impact on our results, while varying the fiducial values of $\{b_2, b_s\}$ has less than a

Parameter	b_2	b_s	α_0	α_2	α_4	N_2	N_4
Default value	0.93	-0.70	-8.65	0	0	-728	0
New value	1.8	-1.4	0	5	5	0	318

Table G.1: Default and new fiducial values for the nuisance parameters of the MiniMapper survey.

Parameter changed	$10^4 h$	$10^4 \omega_c$	$10^3 \tau$	$10^3 \Omega_k$	$10^3 N_{\rm eff}$	M_ν [meV]
b_2	25.1	5.82	6.12	0.71	72.35	34.02
b_s	22.77	5.42	4.96	0.69	77.17	28.23
α_0	19.58	4.79	5.35	0.69	81.02	31.96
α_2	19.4	4.75	5.34	0.69	80.38	31.85
α_4	19.23	4.71	5.34	0.69	79.88	31.83
N_2	19.75	4.83	5.32	0.69	81.35	31.73
N_4	19.69	4.82	5.31	0.69	81.23	31.7
Default	19.68	4.82	5.31	0.69	81.23	31.7

Table G.2: Constraints on selected ΛCDM parameters and extensions for Planck + MiniMapper. In each row we change the relevant nuisance parameter's fiducial value to the new value listed in Table G.1. The fiducial values of the remaining nuisance parameters are left unchanged. In the last row we show the constraints when all of the nuisance terms fiducial values' are set to the default values in Table G.1.

27% impact. We should be able to refine our fiducial values as we learn more about the clustering of these high redshift samples.

We have also checked the sensitivity of Early Dark Energy constraints to the fiducial value of $f_{\rm EDE}$. We find that the constraint on $f_{\rm EDE}$ (for $\log_{10}(z_c) = 3.5$ and $\theta_i = 2.83$) from Planck + MiniMapper changes by less than 3% when the fiducial value of $f_{\rm EDE}$ is increased from 0 to 0.02.

G.6 Python implementation

The Python code used to create these forecasts (`FishLSS`) is publicly available[3]. We include an example Jupyter notebook that walks the user through using the code in detail, and

[3] https://github.com/NoahSailer/FishLSS

provide a brief overview of the organizational structure below.

The primary module is `fisherForecast.py`, which takes a cosmology (an instance of `CLASS`) and experiment (an instance of `experiment.py`) as inputs, and creates a Fisher forecast for a user-specified set of parameters and observables. `fisherForecast.py` contains functions for calculating derivatives of power spectra, computing the 21-cm foreground wedge (under optimistic or pessimistic assumptions), computing Fisher matrices for both 3D and 2D clustering measurements, combining the Fisher matrices from two uncorrelated measurements, and calculating both $N_{\rm modes}$ and $N_{\rm modes}(k)$. `fisherForecast.py` has three helper modules – `experiment.py`, `twoPoint.py`, and `twoPointNoise.py` – which we describe below.

`experiment.py` specifies the experimental configuration of the survey, including the redshift range, redshift bins, linear bias, nonlinear biases (optional), redshift errors, velocity dispersion for FoG-like contributions, and sky coverage. For 21-cm surveys, additional details of the interferometric array are stored in `experiment.py`, such as the number of detectors, fill factor and observing time. We include a set of flags (`LBG`, `HI`, `Halpha`, `ELG`, `Euclid`, `MSE`) which may be switched to `True` should the user wish to use the default biases and number densities of a specific survey described in §10.2.

Power spectra are computed with `twoPoint.py`. In particular, `twoPoint.py` has several functions for computing the fiducial linear biases for the surveys and samples considered in this chapter, as well as functions for computing both the matter and galaxy power spectra in redshift space ($P_{mm}(\boldsymbol{k})$, $P_{gg}(\boldsymbol{k})$); the real space galaxy-matter cross spectrum ($P_{gm}(k)$); the CMB lensing auto-spectrum, cross-correlation with galaxies, and the projected galaxy auto-spectrum ($C_\ell^{\kappa\kappa}$, $C_\ell^{\kappa g}$, C_ℓ^{gg}); and the reconstructed power spectrum. These functions require the installation of both `CLASS`[4] [97] and `velocileptors`[5] [131], which in turn requires `pyFFTW`[6].

We calculate covariances (at the $\boldsymbol{k}$ or ℓ-level) with `twoPointNoise.py`. In particular, `twoPointNoise.py` contains functions for computing the Poisson shot noise of each of the samples considered in this chapter, as well as functions for computing Eqs. (10.10) and (10.18).

While we do not include our code for computing the CMB priors, the Fisher matrices used for Planck, SO, S4, and LiteBIRD can be found in the `input` folder.

[4] https://github.com/lesgourg/class_public
[5] https://github.com/sfschen/velocileptors
[6] https://hgomersall.github.io/pyFFTW/

Bibliography

[1] K. N. Abazajian et al. "CMB-S4 Science Book, First Edition". In: *ArXiv e-prints* (Oct. 2016). arXiv: 1610.02743.

[2] Kevork Abazajian et al. "Snowmass 2021 CMB-S4 White Paper". In: *arXiv e-prints* (Mar. 2022), arXiv:2203.08024. arXiv: 2203.08024 [astro-ph.CO].

[3] Kevork Abazajian et al. "CMB-S4 Science Case, Reference Design, and Project Plan". In: *arXiv e-prints*, arXiv:1907.04473 (July 2019), arXiv:1907.04473. arXiv: 1907.04473 [astro-ph.IM].

[4] T. Abbott et al. "The Dark Energy Survey: more than dark energy – an overview". In: *Mon. Not. Roy. Astron. Soc.* 460.2 (2016), pp. 1270–1299. DOI: 10.1093/mnras/stw641. arXiv: 1601.00329 [astro-ph.CO].

[5] T. M. C. Abbott et al. "Dark Energy Survey Year 1 Results: Cosmological constraints from cluster abundances and weak lensing". In: *Phys. Rev. D* 102.2 (2020), p. 023509. DOI: 10.1103/PhysRevD.102.023509. arXiv: 2002.11124 [astro-ph.CO].

[6] T. M. C. Abbott et al. "Dark Energy Survey Year 3 results: Cosmological constraints from galaxy clustering and weak lensing". In: *Phys. Rev. D* 105.2 (2022), p. 023520. DOI: 10.1103/PhysRevD.105.023520. arXiv: 2105.13549 [astro-ph.CO].

[7] T. M. C. Abbott et al. "DES Y3 + KiDS-1000: Consistent cosmology combining cosmic shear surveys". In: *Open J. Astrophys.* 6 (2023), p. 2305.17173. DOI: 10.21105/astro.2305.17173. arXiv: 2305.17173 [astro-ph.CO].

[8] T. M. C. Abbott et al. "Joint analysis of Dark Energy Survey Year 3 data and CMB lensing from SPT and Planck. III. Combined cosmological constraints". In: *Phys. Rev. D* 107.2 (2023), p. 023531. DOI: 10.1103/PhysRevD.107.023531. arXiv: 2206.10824 [astro-ph.CO].

[9] T. M. C. Abbott et al. "The Dark Energy Survey: Cosmology Results with ∼1500 New High-redshift Type Ia Supernovae Using the Full 5 yr Data Set". In: *Astrophys. J. Lett.* 973.1 (2024), p. L14. DOI: 10.3847/2041-8213/ad6f9f. arXiv: 2401.02929 [astro-ph.CO].

[10] Elcio Abdalla et al. "Cosmology intertwined: A review of the particle physics, astrophysics, and cosmology associated with the cosmological tensions and anomalies". In: *JHEAp* 34 (2022), pp. 49–211. DOI: 10.1016/j.jheap.2022.04.002. arXiv: 2203.06142 [astro-ph.CO].

[11] F. B. Abdalla et al. "Cosmology from HI galaxy surveys with the SKA". In: *Advancing Astrophysics with the Square Kilometre Array (AASKA14)*, 17 (Apr. 2015), p. 17. arXiv: 1501.04035.

[12] Zara Abdurashidova et al. "First Results from HERA Phase I: Upper Limits on the Epoch of Reionization 21 cm Power Spectrum". In: *Astrophys. J.* 925.2, 221 (Feb. 2022), p. 221. DOI: 10.3847/1538-4357/ac1c78. arXiv: 2108.02263 [astro-ph.CO].

[13] Y. Sultan Abylkairov, Omar Darwish, J. Colin Hill, and Blake D. Sherwin. "Partially constrained internal linear combination: A method for low-noise CMB foreground mitigation". In: *Physical Review D* 103.10 (May 2021). ISSN: 2470-0029. DOI: 10.1103/physrevd.103.103510. URL: http://dx.doi.org/10.1103/PhysRevD.103.103510.

[14] R. Adam et al. "Planck intermediate results. XLVII. Planck constraints on reionization history". In: *Astron. Astrophys.* 596 (2016), A108. DOI: 10.1051/0004-6361/201628897. arXiv: 1605.03507 [astro-ph.CO].

[15] A. G. Adame et al. "DESI 2024 VI: cosmological constraints from the measurements of baryon acoustic oscillations". In: *JCAP* 02 (2025), p. 021. DOI: 10.1088/1475-7516/2025/02/021. arXiv: 2404.03002 [astro-ph.CO].

[16] G. Adame et al. "The Early Data Release of the Dark Energy Spectroscopic Instrument". In: (June 2023). DOI: 10.5281/zenodo.7964161. arXiv: 2306.06308 [astro-ph.CO].

[17] Jennifer Adams, Bevan Cresswell, and Richard Easther. "Inflationary perturbations from a potential with a step". In: *Physical Review D* 64.12 (Nov. 2001). ISSN: 1089-4918. DOI: 10.1103/physrevd.64.123514. URL: http://dx.doi.org/10.1103/PhysRevD.64.123514.

[18] P. A. R. Ade et al. "BICEP / Keck XVII: Line of Sight Distortion Analysis: Estimates of Gravitational Lensing, Anisotropic Cosmic Birefringence, Patchy Reionization, and Systematic Errors". In: *arXiv e-prints*, arXiv:2210.08038 (Oct. 2022), arXiv:2210.08038. arXiv: 2210.08038 [astro-ph.CO].

[19] P. A. R. Ade et al. "Planck 2015 results. XVI. Isotropy and statistics of the CMB". In: *Astron. Astrophys.* 594 (2016), A16. DOI: 10.1051/0004-6361/201526681. arXiv: 1506.07135 [astro-ph.CO].

[20] P. A. R. Ade et al. "Planck 2015 results. XXIV. Cosmology from Sunyaev-Zeldovich cluster counts". In: *Astron. Astrophys.* 594 (2016), A24. DOI: 10.1051/0004-6361/201525833. arXiv: 1502.01597 [astro-ph.CO].

[21] P. A. R. Ade et al. "Planck 2013 results. XVII. Gravitational lensing by large-scale structure". In: *Astronomy & Astrophysics* 571 (Oct. 2014), A17. ISSN: 1432-0746. DOI: 10.1051/0004-6361/201321543. URL: http://dx.doi.org/10.1051/0004-6361/201321543.

[22] Peter Ade et al. "The Simons Observatory: science goals and forecasts". In: *J. Cosm. Astrop. Phys.* 2019.2, 056 (Feb. 2019), p. 056. DOI: 10.1088/1475-7516/2019/02/056. arXiv: 1808.07445 [astro-ph.CO].

[23] Nishant Agarwal et al. "Characterizing unknown systematics in large scale structure surveys". In: *JCAP* 04 (2014), p. 007. DOI: 10.1088/1475-7516/2014/04/007. arXiv: 1309.2954 [astro-ph.CO].

[24] N. Aghanim et al. "Planck 2018 results. III. High Frequency Instrument data processing and frequency maps". In: *Astron. Astrophys.* 641 (2020), A3. DOI: 10.1051/0004-6361/201832909. arXiv: 1807.06207 [astro-ph.CO].

[25] N. Aghanim et al. "Planck 2018 results. VIII. Gravitational lensing". In: *Astron. Astrophys.* 641 (2020), A8. DOI: 10.1051/0004-6361/201833886. arXiv: 1807.06210 [astro-ph.CO].

[26] Simone Aiola et al. "The Atacama Cosmology Telescope: DR4 Maps and Cosmological Parameters". In: *JCAP* 12 (2020), p. 047. DOI: 10.1088/1475-7516/2020/12/047. arXiv: 2007.07288 [astro-ph.CO].

[27] Y. Akrami et al. "Planck 2018 results. IV. Diffuse component separation". In: *Astron. Astrophys.* 641 (2020), A4. DOI: 10.1051/0004-6361/201833881. arXiv: 1807.06208 [astro-ph.CO].

[28] S. Alam et al. "The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: cosmological analysis of the DR12 galaxy sample". In: *Mon. Not. R. Astron. Soc.* 470 (Sept. 2017), pp. 2617–2652. DOI: 10.1093/mnras/stx721. arXiv: 1607.03155.

[29] Shadab Alam et al. "The Eleventh and Twelfth Data Releases of the Sloan Digital Sky Survey: Final Data from SDSS-III". In: *Astrophys. J. Supp.* 219.1, 12 (July 2015), p. 12. DOI: 10.1088/0067-0049/219/1/12. arXiv: 1501.00963 [astro-ph.IM].

[30] C. Alcock and B. Paczynski. "An evolution free test for non-zero cosmological constant". In: *Nature* 281 (Oct. 1979), p. 358. DOI: 10.1038/281358a0.

[31] Itamar J. Allali, Praniti Singh, JiJi Fan, and Lingfeng Li. "Reionization and the Hubble Constant: Correlations in the Cosmic Microwave Background". In: *arXiv e-prints*, arXiv:2503.05691 (Mar. 2025), arXiv:2503.05691. DOI: 10.48550/arXiv.2503.05691. arXiv: 2503.05691 [astro-ph.CO].

[32] R. Allison, P. Caucal, E. Calabrese, J. Dunkley, and T. Louis. "Towards a cosmological neutrino mass detection". In: *Phys. Rev. D* 92.12, 123535 (Dec. 2015), p. 123535. DOI: 10.1103/PhysRevD.92.123535. arXiv: 1509.07471 [astro-ph.CO].

[33] David Alonso et al. "The LSST Dark Energy Science Collaboration (DESC) Science Requirements Document". In: (Sept. 2018). arXiv: 1809.01669 [astro-ph.CO].

[34] David Alonso, Philip Bull, Pedro G. Ferreira, Roy Maartens, and M. Santos. "Ultra large-scale cosmology in next-generation experiments with single tracers". In: *Astrophys. J.* 814.2 (2015), p. 145. DOI: 10.1088/0004-637X/814/2/145. arXiv: 1505.07596 [astro-ph.CO].

[35] David Alonso, Javier Sanchez, and Anže Slosar. "A unified pseudo-C_ℓ framework". In: *Mon. Not. Roy. Astron. Soc.* 484.3 (2019), pp. 4127–4151. DOI: 10.1093/mnras/stz093. arXiv: 1809.09603 [astro-ph.CO].

[36] Marcelo A. Alvarez, Simone Ferraro, J. Colin Hill, Renée Hložek, and Margaret Ikape. "Mitigating the optical depth degeneracy using the kinematic Sunyaev-Zel'dovich effect with CMB-S4". In: *Phys. Rev. D* 103.6 (2021), p. 063518. DOI: 10.1103/PhysRevD.103.063518. arXiv: 2006.06594 [astro-ph.CO].

[37] Luca Amendola et al. "Cosmology and fundamental physics with the Euclid satellite". In: *Living Reviews in Relativity* 21.1, 2 (Apr. 2018), p. 2. DOI: 10.1007/s41114-017-0010-3. arXiv: 1606.00180 [astro-ph.CO].

[38] Stefania Amodeo et al. "Atacama Cosmology Telescope: Modeling the gas thermodynamics in BOSS CMASS galaxies from kinematic and thermal Sunyaev-Zel'dovich measurements". In: *Phys. Rev. D* 103.6, 063514 (Mar. 2021), p. 063514. DOI: 10.1103/PhysRevD.103.063514. arXiv: 2009.05558 [astro-ph.CO].

[39] A. Amon et al. "Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to data calibration". In: *Phys. Rev. D* 105.2 (2022), p. 023514. DOI: 10.1103/PhysRevD.105.023514. arXiv: 2105.13543 [astro-ph.CO].

[40] Alexandra Amon and George Efstathiou. "A non-linear solution to the S_8 tension?" In: *Mon. Not. R. Astron. Soc.* 516.4 (Nov. 2022), pp. 5355–5366. DOI: 10.1093/mnras/stac2429. arXiv: 2206.11794 [astro-ph.CO].

[41] K. J. Andersen et al. "BeyondPlanck I. Global Bayesian analysis of the Planck Low Frequency Instrument data". In: *BeyondPlanck Release Conference*. Nov. 2020. DOI: 10.1051/0004-6361/202244953. arXiv: 2011.05609 [astro-ph.CO].

[42] Réza Ansari et al. "Inflation and Early Dark Energy with a Stage II Hydrogen Intensity Mapping experiment". In: (Oct. 2018). arXiv: 1810.09572 [astro-ph.CO].

[43] Maria Archidiacono, Emanuele Castorina, Diego Redigolo, and Ennio Salvioni. "Unveiling dark fifth forces with linear cosmology". In: *JCAP* 10 (2022), p. 074. DOI: 10.1088/1475-7516/2022/10/074. arXiv: 2204.08484 [astro-ph.CO].

[44] Maria Archidiacono, Deanna C. Hooper, Riccardo Murgia, Sebastian Bohr, Julien Lesgourgues, and Matteo Viel. "Constraining Dark Matter-Dark Radiation interactions with CMB, BAO, and Lyman-α". In: *JCAP* 10 (2019), p. 055. DOI: 10.1088/1475-7516/2019/10/055. arXiv: 1907.01496 [astro-ph.CO].

[45] Giovanni Aricò, Raul E. Angulo, Matteo Zennaro, Sergio Contreras, Angela Chen, and Carlos Hernández-Monteagudo. "DES Y3 cosmic shear down to small scales: Constraints on cosmology and baryons". In: *Astron. Astrophys.* 678, A109 (Oct. 2023), A109. DOI: 10.1051/0004-6361/202346539. arXiv: 2303.05537 [astro-ph.CO].

[46] Mahnaz Asghari, Jose Beltrán Jiménez, Shahram Khosravi, and David F. Mota. "On structure formation from a small-scales-interacting dark sector". In: *JCAP* 04 (2019), p. 042. DOI: 10.1088/1475-7516/2019/04/042. arXiv: 1902.05532 [astro-ph.CO].

[47] V. Assassi, D. Baumann, D. Green, and M. Zaldarriaga. "Renormalized halo bias". In: *J. Cosm. Astrop. Phys.* 8, 056 (Aug. 2014), p. 056. DOI: 10.1088/1475-7516/2014/08/056. arXiv: 1402.5916.

[48] Alejandro Aviles and Arka Banerjee. "A Lagrangian Perturbation Theory in the presence of massive neutrinos". In: *JCAP* 10 (2020), p. 034. DOI: 10.1088/1475-7516/2020/10/034. arXiv: 2007.06508 [astro-ph.CO].

[49] Micaela B. Bagley et al. "HST Grism-derived Forecasts for Future Galaxy Redshift Surveys". In: *arXiv e-prints*, arXiv:2006.03602 (June 2020), arXiv:2006.03602. arXiv: 2006.03602 [astro-ph.GA].

[50] Tobias Baldauf, Mehrdad Mirbabayi, Marko Simonović, and Matias Zaldarriaga. "Equivalence principle and the baryon acoustic peak". In: *Phys. Rev. D* 92.4, 043514 (Aug. 2015), p. 043514. DOI: 10.1103/PhysRevD.92.043514. arXiv: 1504.04366 [astro-ph.CO].

[51] Tobias Baldauf, Uros Seljak, Vincent Desjacques, and Patrick McDonald. "Evidence for quadratic tidal tensor bias from the halo bispectrum". In: *Phys. Rev. D* 86 (8 Oct. 2012), p. 083540. DOI: 10.1103/PhysRevD.86.083540. URL: https://link.aps.org/doi/10.1103/PhysRevD.86.083540.

[52] Antón Baleato Lizancos and Simone Ferraro. "Impact of extragalactic foregrounds on internal delensing of the CMB B-mode polarization". In: *Phys. Rev. D* 106.6 (2022), p. 063534. DOI: 10.1103/PhysRevD.106.063534. arXiv: 2205.09000 [astro-ph.CO].

[53] Antón Baleato Lizancos and Martin White. "Harmonic analysis of discrete tracers of large-scale structure". In: (Dec. 2023). arXiv: 2312.12285 [astro-ph.CO].

[54] Antón Baleato Lizancos and Martin White. "The impact of anisotropic redshift distributions on angular clustering". In: *JCAP* 07 (2023), p. 044. DOI: 10.1088/1475-7516/2023/07/044. arXiv: 2305.15406 [astro-ph.CO].

[55] Mario Ballardini, Fabio Finelli, Cosimo Fedeli, and Lauro Moscardini. "Probing primordial features with future galaxy surveys". In: *JCAP* 10 (2016). [Erratum: JCAP 04, E01 (2018)], p. 041. DOI: 10.1088/1475-7516/2016/10/041. arXiv: 1606.03747 [astro-ph.CO].

[56] Mario Ballardini, Riccardo Murgia, Marco Baldi, Fabio Finelli, and Matteo Viel. "Non-linear damping of superimposed primordial oscillations on the matter power spectrum in galaxy surveys". In: *JCAP* 04.04 (2020), p. 030. DOI: 10.1088/1475-7516/2020/04/030. arXiv: 1912.12499 [astro-ph.CO].

[57] K. Bandura et al. "Canadian Hydrogen Intensity Mapping Experiment (CHIME) pathfinder". In: *Ground-based and Airborne Telescopes V*. Vol. 9145. Proc. SPIE. July 2014, 914522, p. 914522. DOI: 10.1117/12.2054950. arXiv: 1406.2288 [astro-ph.IM].

[58] Danial Baradaran, Boryana Hadzhiyska, Martin J. White, and Noah Sailer. "Predicting the 21-cm field with a hybrid effective field theory approach". In: *Phys. Rev. D* 110.10 (2024), p. 103517. DOI: 10.1103/PhysRevD.110.103517. arXiv: 2406.13079 [astro-ph.CO].

[59] Rennan Barkana and Abraham Loeb. "A Method for Separating the Physics from the Astrophysics of High-Redshift 21 Centimeter Fluctuations". In: *Astrophys. J. Lett.* 624.2 (May 2005), pp. L65–L68. DOI: 10.1086/430599. arXiv: astro-ph/0409572 [astro-ph].

[60] Alexandre Barreira. "On the impact of galaxy bias uncertainties on primordial non-Gaussianity constraints". In: *JCAP* 12 (2020), p. 031. DOI: 10.1088/1475-7516/2020/12/031. arXiv: 2009.06622 [astro-ph.CO].

[61] Matthias Bartelmann and Peter Schneider. "Weak gravitational lensing". In: *Physics Reports* 340.4 (2001), pp. 291–472. ISSN: 0370-1573. DOI: https://doi.org/10.1016/S0370-1573(00)00082-X. URL: https://www.sciencedirect.com/science/article/pii/S037015730000082X.

[62] Sergei Bashinsky and Uros Seljak. "Neutrino perturbations in CMB anisotropy and matter clustering". In: *Phys. Rev. D* 69 (2004), p. 083002. DOI: 10.1103/PhysRevD.69.083002. arXiv: astro-ph/0310198.

[63] N. Battaglia, J. R. Bond, C. Pfrommer, and J. L. Sievers. "On the Cluster Physics of Sunyaev-Zel'dovich and X-Ray Surveys. I. The Influence of Feedback, Non-thermal Pressure, and Cluster Shapes on Y-M Scaling Relations". In: *Astrophys. J.* 758.2, 74 (Oct. 2012), p. 74. DOI: 10.1088/0004-637X/758/2/74. arXiv: 1109.3709 [astro-ph.CO].

[64] N. Battaglia, J. R. Bond, C. Pfrommer, J. L. Sievers, and D. Sijacki. "Simulations of the Sunyaev-Zel'dovich Power Spectrum with Active Galactic Nucleus Feedback". In: *Astrophys. J.* 725.1 (Dec. 2010), pp. 91–99. DOI: 10.1088/0004-637X/725/1/91. arXiv: 1003.4256 [astro-ph.CO].

[65] N. Battaglia, A. Natarajan, H. Trac, R. Cen, and A. Loeb. "Reionization on Large Scales. III. Predictions for Low-l Cosmic Microwave Background Polarization and High-l Kinetic Sunyaev-Zel'dovich Observables". In: *Astrophys. J.* 776.2, 83 (Oct. 2013), p. 83. DOI: 10.1088/0004-637X/776/2/83. arXiv: 1211.2832 [astro-ph.CO].

[66] Nicholas Battaglia, Simone Ferraro, Emmanuel Schaan, and David N. Spergel. "Future constraints on halo thermodynamics from combined Sunyaev-Zel'dovich measurements". In: *J. Cosm. Astrop. Phys.* 2017.11, 040 (Nov. 2017), p. 040. DOI: 10.1088/1475-7516/2017/11/040. arXiv: 1705.05881 [astro-ph.CO].

[67] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga. "Cosmological non-linearities as an effective fluid". In: *J. Cosm. Astrop. Phys.* 7, 051 (July 2012), p. 51. DOI: 10.1088/1475-7516/2012/07/051. arXiv: 1004.2488 [astro-ph.CO].

[68] Daniel Baumann. "TASI Lectures on Inflation". In: *arXiv e-prints*, arXiv:0907.5424 (July 2009), arXiv:0907.5424. arXiv: 0907.5424 [hep-th].

[69] Daniel Baumann, Daniel Green, Joel Meyers, and Benjamin Wallisch. "Phases of New Physics in the CMB". In: *JCAP* 01 (2016), p. 007. DOI: 10.1088/1475-7516/2016/01/007. arXiv: 1508.06342 [astro-ph.CO].

[70] Daniel Baumann, Daniel Green, and Matias Zaldarriaga. "Phases of New Physics in the BAO Spectrum". In: *JCAP* 11 (2017), p. 007. DOI: 10.1088/1475-7516/2017/11/007. arXiv: 1703.00894 [astro-ph.CO].

[71] E. J. Baxter et al. "Dark Energy Survey Year 1 results: Methodology and projections for joint analysis of galaxy clustering, galaxy lensing, and CMB lensing two-point functions". In: *Phys. Rev. D* 99.2, 023508 (Jan. 2019), p. 023508. DOI: 10.1103/PhysRevD.99.023508. arXiv: 1802.05257 [astro-ph.CO].

[72] Adrian E. Bayer, Arka Banerjee, and Uroš Seljak. "Beware of Fake νs: The Effect of Massive Neutrinos on the Non-Linear Evolution of Cosmic Structure". In: *arXiv e-prints*, arXiv:2108.04215 (Aug. 2021), arXiv:2108.04215. arXiv: 2108.04215 [astro-ph.CO].

[73] Dominic Beck, Josquin Errard, and Radek Stompor. "Impact of Polarized Galactic Foreground Emission on CMB Lensing Reconstruction and Delensing of B-Modes". In: *JCAP* 06 (2020), p. 030. DOI: 10.1088/1475-7516/2020/06/030. arXiv: 2001.02641 [astro-ph.CO].

[74] Dominic Beck, Giulio Fabbian, and Josquin Errard. "Lensing Reconstruction in Post-Born Cosmic Microwave Background Weak Lensing". In: *Phys. Rev. D* 98.4 (2018), p. 043512. DOI: 10.1103/PhysRevD.98.043512. arXiv: 1806.01216 [astro-ph.CO].

[75] George D. Becker, James S. Bolton, and Adam Lidz. "Reionisation and High-Redshift Galaxies: The View from Quasar Absorption Lines". In: *Publ. Astron. Soc. Australia* 32, e045 (Dec. 2015), e045. DOI: 10.1017/pasa.2015.45. arXiv: 1510.03368 [astro-ph.CO].

[76] Niklas Becker, Deanna C. Hooper, Felix Kahlhoefer, Julien Lesgourgues, and Nils Schöneberg. "Cosmological constraints on multi-interacting dark matter". In: *JCAP* 02 (2021), p. 019. DOI: 10.1088/1475-7516/2021/02/019. arXiv: 2010.04074 [astro-ph.CO].

[77] K. M. Belotsky, A. A. Kirillov, N. O. Nazarova, and S. G. Rubin. "Reionization effect enhancement due to primordial black holes". In: *Int. J. Mod. Phys. D* 26.09 (2017), p. 1750102. DOI: 10.1142/S0218271817501024. arXiv: 1702.06338 [astro-ph.CO].

[78] Roger de Belsunce, Steven Gratton, William Coulton, and George Efstathiou. "Inference of the optical depth to reionization from low multipole temperature and polarization Planck data". In: *Mon. Not. Roy. Astron. Soc.* 507.1 (2021), pp. 1072–1091. DOI: 10.1093/mnras/stab2215. arXiv: 2103.14378 [astro-ph.CO].

[79] C. L. Bennett et al. "Nine-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results". In: *Astrophys. J. Suppl.* 208 (2013), p. 20. DOI: 10.1088/0067-0049/208/2/20. arXiv: 1212.5225 [astro-ph.CO].

[80] B. A. Benson et al. "SPT-3G: a next-generation cosmic microwave background polarization experiment on the South Pole telescope". In: Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 9153, 91531P (July 2014), 91531P. DOI: 10.1117/12.2057305. arXiv: 1407.2973 [astro-ph.IM].

[81] José Ramón Bermejo-Climent, Mario Ballardini, Fabio Finelli, Daniela Paoletti, Roy Maartens, José Alberto Rubiño-Martín, and Luca Valenziano. "Cosmological parameter forecasts by a joint 2D tomographic approach to CMB and galaxy clustering". In: *Phys. Rev. D* 103.10 (2021), p. 103502. DOI: 10.1103/PhysRevD.103.103502. arXiv: 2106.05267 [astro-ph.CO].

[82] F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro. "Large-scale structure of the Universe and cosmological perturbation theory". In: *Phys. Rep.* 367 (Sept. 2002), pp. 1–248. DOI: 10.1016/S0370-1573(02)00135-7. eprint: astro-ph/0112551.

[83] Florian Beutler, Matteo Biagetti, Daniel Green, Anže Slosar, and Benjamin Wallisch. "Primordial features from linear to nonlinear scales". In: *Physical Review Research* 1.3 (Dec. 2019). ISSN: 2643-1564. DOI: 10.1103/physrevresearch.1.033209. URL: http://dx.doi.org/10.1103/PhysRevResearch.1.033209.

[84] Florian Beutler et al. "The 6dF Galaxy Survey: baryon acoustic oscillations and the local Hubble constant". In: *Mon. Not. R. Astron. Soc.* 416.4 (Oct. 2011), pp. 3017–3032. DOI: 10.1111/j.1365-2966.2011.19250.x. arXiv: 1106.3366 [astro-ph.CO].

[85] S. Bharadwaj. "The Evolution of Correlation Functions in the Zeldovich Approximation and Its Implications for the Validity of Perturbation Theory". In: *Astrophys. J.* 472 (Nov. 1996), pp. 1–+. DOI: 10.1086/178036. eprint: arXiv:astro-ph/9606121.

[86] Matteo Biagetti. "The Hunt for Primordial Interactions in the Large Scale Structures of the Universe". In: *Galaxies* 7.3 (2019), p. 71. DOI: 10.3390/galaxies7030071. arXiv: 1906.12244 [astro-ph.CO].

[87] Matteo Biagetti, Titouan Lazeyras, Tobias Baldauf, Vincent Desjacques, and Fabian Schmidt. "Verifying the consistency relation for the scale-dependent bias from local primordial non-Gaussianity". In: *Mon. Not. Roy. Astron. Soc.* 468.3 (2017), pp. 3277–3288. DOI: 10.1093/mnras/stx714. arXiv: 1611.04901 [astro-ph.CO].

[88] F. Bianchini et al. "Constraints on Cosmological Parameters from the 500 deg^2 SPTpol Lensing Power Spectrum". In: *Astrophys. J.* 888 (2020), p. 119. DOI: 10.3847/1538-4357/ab6082. arXiv: 1910.07157 [astro-ph.CO].

[89] Federico Bianchini and Abhishek S. Maniyar. "The Cosmic Microwave Background – Secondary Anisotropies". In: (Jan. 2025). arXiv: 2501.13913 [astro-ph.CO].

[90] Federico Bianchini and Marius Millea. "Inference of gravitational lensing and patchy reionization with future CMB data". In: *Phys. Rev. D* 107.4 (2023), p. 043521. DOI: 10.1103/PhysRevD.107.043521. arXiv: 2210.10893 [astro-ph.CO].

[91] L. Bigwood et al. "Weak lensing combined with the kinetic Sunyaev-Zel'dovich effect: a study of baryonic feedback". In: *Mon. Not. R. Astron. Soc.* 534.1 (Oct. 2024), pp. 655–682. DOI: 10.1093/mnras/stae2100. arXiv: 2404.06098 [astro-ph.CO].

[92] Tashalee S. Billings, Paul La Plante, and James E. Aguirre. *Extracting the Optical Depth to Reionization τ from 21 cm Data Using Machine Learning Techniques*. 2021. arXiv: 2103.14563 [astro-ph.CO].

[93] M. Birkinshaw. "The Sunyaev-Zel'dovich effect". In: *Phys. Rep.* 310.2-3 (Mar. 1999), pp. 97–195. DOI: 10.1016/S0370-1573(98)00080-5. arXiv: astro-ph/9808050 [astro-ph].

[94] A. Blanchard and J. Schneider. "Gravitational lensing effect on the fluctuations of the cosmic background radiation". In: *Astron. Astrophys.* 184.1-2 (Oct. 1987), pp. 1–6.

[95] A. Blanchard et al. "Euclid preparation". In: *å* 642 (Oct. 2020), A191. ISSN: 1432-0746. DOI: 10.1051/0004-6361/202038071. arXiv: 1910.09273. URL: http://dx.doi.org/10.1051/0004-6361/202038071.

[96] Alain Blanchard et al. "Gravitation And the Universe from large Scale-Structures: The GAUSS mission concept. Mapping the cosmic web up to the reionization era". In: (Feb. 2021). DOI: 10.1007/s10686-021-09717-9. arXiv: 2102.03931 [astro-ph.CO].

[97] D. Blas, J. Lesgourgues, and T. Tram. "The Cosmic Linear Anisotropy Solving System (CLASS). Part II: Approximation schemes". In: *J. Cosm. Astrop. Phys.* 7, 034 (July 2011), p. 034. DOI: 10.1088/1475-7516/2011/07/034. arXiv: 1104.2933.

[98] Diego Blas, Mathias Garny, Mikhail M. Ivanov, and Sergey Sibiryakov. "Time-sliced perturbation theory II: baryon acoustic oscillations and infrared resummation". In: *J. Cosm. Astrop. Phys.* 2016.7, 028 (July 2016), p. 028. DOI: 10.1088/1475-7516/2016/07/028. arXiv: 1605.02149 [astro-ph.CO].

[99] Jonathan Blazek, Niall MacCrann, M. A. Troxel, and Xiao Fang. "Beyond linear galaxy alignments". In: *Phys. Rev. D* 100.10 (2019), p. 103506. DOI: 10.1103/PhysRevD.100.103506. arXiv: 1708.09247 [astro-ph.CO].

[100] S. Bocquet et al. "SPT Clusters with DES and HST Weak Lensing. II. Cosmological Constraints from the Abundance of Massive Halos". In: *arXiv e-prints*, arXiv:2401.02075 (Jan. 2024), arXiv:2401.02075. DOI: 10.48550/arXiv.2401.02075. arXiv: 2401.02075 [astro-ph.CO].

[101] Vanessa Böhm, Chirag Modi, and Emanuele Castorina. "Lensing corrections on galaxy-lensing cross correlations and galaxy-galaxy auto correlations". In: *JCAP* 03 (2020), p. 045. DOI: 10.1088/1475-7516/2020/03/045. arXiv: 1910.06722 [astro-ph.CO].

[102] Boris Bolliet, J. Colin Hill, Simone Ferraro, Aleksandra Kusiak, and Alex Krolewski. "Projected-field kinetic Sunyaev-Zel'dovich Cross-correlations: halo model and forecasts". In: *J. Cosm. Astrop. Phys.* 2023.3, 039 (Mar. 2023), p. 039. DOI: 10.1088/1475-7516/2023/03/039. arXiv: 2208.07847 [astro-ph.CO].

[103] J. R. Bond, G. Efstathiou, and J. Silk. "Massive Neutrinos and the Large-Scale Structure of the Universe". In: *Phys. Rev. Lett.* 45 (24 Dec. 1980), pp. 1980–1984. DOI: 10.1103/PhysRevLett.45.1980. URL: https://link.aps.org/doi/10.1103/PhysRevLett.45.1980.

[104] Aoife Boyle and Fabian Schmidt. "Neutrino mass constraints beyond linear order: cosmology dependence and systematic biases". In: *arXiv e-prints*, arXiv:2011.10594 (Nov. 2020), arXiv:2011.10594. arXiv: 2011.10594 [astro-ph.CO].

[105] Diogo Bragança, Yaniv Donath, Leonardo Senatore, and Henry Zheng. "Peeking into the next decade in Large-Scale Structure Cosmology with its Effective Field Theory". In: (July 2023). arXiv: 2307.04992 [astro-ph.CO].

[106] Diogo P. L. Bragança, Matthew Lewandowski, David Sekera, Leonardo Senatore, and Raphael Sgier. "Baryonic effects in the Effective Field Theory of Large-Scale Structure and an analytic recipe for lensing in CMB-S4". In: (Oct. 2020). arXiv: 2010.02929 [astro-ph.CO].

[107] Thejs Brinckmann, Deanna C. Hooper, Maria Archidiacono, Julien Lesgourgues, and Tim Sprenger. "The promising future of a robust cosmological neutrino mass measurement". In: *JCAP* 01 (2019), p. 059. DOI: 10.1088/1475-7516/2019/01/059. arXiv: 1808.05955 [astro-ph.CO].

[108] Stephen P. Brooks and Andrew Gelman. "General Methods for Monitoring Convergence of Iterative Simulations". In: *Journal of Computational and Graphical Statistics* 7.4 (1998), pp. 434–455. DOI: 10.1080/10618600.1998.10474787. eprint: https://www.tandfonline.com/doi/pdf/10.1080/10618600.1998.10474787. URL: https://www.tandfonline.com/doi/abs/10.1080/10618600.1998.10474787.

[109] Dillon Brout et al. "The Pantheon+ Analysis: Cosmological Constraints". In: *Astrophys. J.* 938.2 (2022), p. 110. DOI: 10.3847/1538-4357/ac8e04. arXiv: 2202.04077 [astro-ph.CO].

[110] Manuel A. Buen-Abad, Gustavo Marques-Tavares, and Martin Schmaltz. "Non-Abelian dark matter and dark radiation". In: *Phys. Rev. D* 92.2 (2015), p. 023531. DOI: 10.1103/PhysRevD.92.023531. arXiv: 1505.03542 [hep-ph].

[111] Manuel A. Buen-Abad, Martin Schmaltz, Julien Lesgourgues, and Thejs Brinckmann. "Interacting Dark Sector and Precision Cosmology". In: *JCAP* 01 (2018), p. 008. DOI: 10.1088/1475-7516/2018/01/008. arXiv: 1708.09406 [astro-ph.CO].

[112] Philip Bull, Pedro G. Ferreira, Prina Patel, and Mario G. Santos. "Late-time cosmology with 21cm intensity mapping experiments". In: *Astrophys. J.* 803.1 (2015), p. 21. DOI: 10.1088/0004-637X/803/1/21. arXiv: 1405.1452 [astro-ph.CO].

[113] James S. Bullock and Michael Boylan-Kolchin. "Small-Scale Challenges to the ΛCDM Paradigm". In: *Annu. Rev. Astron. Astrophys.* 55.1 (Aug. 2017), pp. 343–387. DOI: 10.1146/annurev-astro-091916-055313. arXiv: 1707.04256 [astro-ph.CO].

[114] Hongbo Cai, Mathew S. Madhavacheril, J. Colin Hill, and Arthur Kosowsky. "Bias to cosmic microwave background lensing reconstruction from the kinematic Sunyaev-Zel'dovich effect at reionization". In: *Phys. Rev. D* 105.4 (2022), p. 043516. DOI: 10.1103/PhysRevD.105.043516. arXiv: 2111.01944 [astro-ph.CO].

[115] Erminia Calabrese et al. "The Atacama Cosmology Telescope: DR6 Constraints on Extended Cosmological Models". In: *arXiv e-prints* (Mar. 2025). arXiv: 2503.14454 [astro-ph.CO].

[116] V. Calafut et al. "The Atacama Cosmology Telescope: Detection of the pairwise kinematic Sunyaev-Zel'dovich effect with SDSS DR15 galaxies". In: *Phys. Rev. D* 104.4, 043502 (Aug. 2021), p. 043502. DOI: 10.1103/PhysRevD.104.043502. arXiv: 2101.08374 [astro-ph.CO].

[117] J. Carlson, B. Reid, and M. White. "Convolution Lagrangian perturbation theory for biased tracers". In: *Mon. Not. R. Astron. Soc.* 429 (Feb. 2013), pp. 1674–1685. DOI: 10.1093/mnras/sts457. arXiv: 1209.0780.

[118] J. J. M. Carrasco, M. P. Hertzberg, and L. Senatore. "The effective field theory of cosmological large scale structures". In: *Journal of High Energy Physics* 9 (Sept. 2012), p. 82. DOI: 10.1007/JHEP09(2012)082. arXiv: 1206.2926 [astro-ph.CO].

[119] Julien Carron, Mark Mirmelstein, and Antony Lewis. "CMB lensing from Planck PR4 maps". In: *JCAP* 09 (2022), p. 039. DOI: 10.1088/1475-7516/2022/09/039. arXiv: 2206.07773 [astro-ph.CO].

[120] Emanuele Castorina et al. "Redshift-weighted constraints on primordial non-Gaussianity from the clustering of the eBOSS DR14 quasars in Fourier space". In: *JCAP* 09 (2019), p. 010. DOI: 10.1088/1475-7516/2019/09/010. arXiv: 1904.08859 [astro-ph.CO].

[121] Emanuele Castorina, Carmelita Carbone, Julien Bel, Emiliano Sefusatti, and Klaus Dolag. "DEMNUni: The clustering of large-scale structures in the presence of massive neutrinos". In: *JCAP* 07 (2015), p. 043. DOI: 10.1088/1475-7516/2015/07/043. arXiv: 1505.07148 [astro-ph.CO].

[122] Emanuele Castorina, Emiliano Sefusatti, Ravi K. Sheth, Francisco Villaescusa-Navarro, and Matteo Viel. "Cosmology with massive neutrinos II: on the universality of the halo mass function and bias". In: *JCAP* 02 (2014), p. 049. DOI: 10.1088/1475-7516/2014/02/049. arXiv: 1311.1212 [astro-ph.CO].

[123] Emanuele Castorina and Francisco Villaescusa-Navarro. "On the spatial distribution of neutral hydrogen in the Universe: bias and shot-noise of the HI power spectrum". In: *Monthly Notices of the Royal Astronomical Society* 471.2 (June 2017), pp. 1788–1796. ISSN: 1365-2966. DOI: 10.1093/mnras/stx1599. URL: http://dx.doi.org/10.1093/mnras/stx1599.

[124] Emanuele Castorina and Martin White. "Measuring the growth of structure with intensity mapping surveys". In: *J. Cosm. Astrop. Phys.* 2019.6, 025 (June 2019), p. 025. DOI: 10.1088/1475-7516/2019/06/025. arXiv: 1902.07147 [astro-ph.CO].

[125] Emanuele Castorina et al. "Packed Ultra-wideband Mapping Array (PUMA): Astro2020 RFI Response". In: *arXiv e-prints*, arXiv:2002.05072 (Feb. 2020), arXiv:2002.05072. arXiv: 2002.05072 [astro-ph.IM].

[126] Solène Chabanier, Marius Millea, and Nathalie Palanque-Delabrouille. "Matter power spectrum: from Lyα forest to CMB scales". In: *Mon. Not. Roy. Astron. Soc.* 489.2 (2019), pp. 2247–2253. DOI: 10.1093/mnras/stz2310. arXiv: 1905.08103 [astro-ph.CO].

[127] K. C. Chan, R. Scoccimarro, and R. K. Sheth. "Gravity and large-scale nonlocal bias". In: *Phys. Rev. D* 85.8, 083509 (Apr. 2012), p. 083509. DOI: 10.1103/PhysRevD.85.083509. arXiv: 1201.3614 [astro-ph.CO].

[128] C. Chang et al. "Joint analysis of Dark Energy Survey Year 3 data and CMB lensing from SPT and Planck. II. Cross-correlation measurements and cosmological constraints". In: *Phys. Rev. D* 107.2 (2023), p. 023530. DOI: 10.1103/PhysRevD.107.023530. arXiv: 2203.12440 [astro-ph.CO].

[129] S. Chen et al. "Not all lensing is low: An analysis of DESI×DES using the Lagrangian Effective Theory of LSS". In: *arXiv e-prints*, arXiv:2407.04795 (July 2024), arXiv:2407.04795. DOI: 10.48550/arXiv.2407.04795. arXiv: 2407.04795 [astro-ph.CO].

[130] Shi-Fan Chen, Emanuele Castorina, and Martin White. "Biased tracers of two fluids in the Lagrangian picture". In: *J. Cosm. Astrop. Phys.* 2019.6, 006 (June 2019), p. 006. DOI: 10.1088/1475-7516/2019/06/006. arXiv: 1903.00437 [astro-ph.CO].

[131] Shi-Fan Chen, Emanuele Castorina, Martin White, and Anže Slosar. "Synergies between radio, optical and microwave observations at high redshift". In: *J. Cosm. Astrop. Phys.* 2019.7, 023 (July 2019), p. 023. DOI: 10.1088/1475-7516/2019/07/023. arXiv: 1810.00911 [astro-ph.CO].

[132] Shi-Fan Chen, Mikhail M. Ivanov, Oliver H. E. Philcox, and Lukas Wenzl. "Suppression without Thawing: Constraining Structure Formation and Dark Energy with Galaxy Clustering". In: *Phys. Rev. Lett.* 133.23, 231001 (Dec. 2024), p. 231001. DOI: 10.1103/PhysRevLett.133.231001. arXiv: 2406.13388 [astro-ph.CO].

[133] Shi-Fan Chen and Nickolas Kokron. "A Lagrangian theory for galaxy shape statistics". In: *JCAP* 01 (2024), p. 027. DOI: 10.1088/1475-7516/2024/01/027. arXiv: 2309.16761 [astro-ph.CO].

[134] Shi-Fan Chen, Zvonimir Vlah, Emanuele Castorina, and Martin White. "Redshift-Space Distortions in Lagrangian Perturbation Theory". In: *arXiv e-prints*, arXiv:2012.04636 (Dec. 2020), arXiv:2012.04636. arXiv: 2012.04636 [astro-ph.CO].

[135] Shi-Fan Chen, Zvonimir Vlah, and Martin White. "A new analysis of galaxy 2-point functions in the BOSS survey, including full-shape information and post-reconstruction BAO". In: *JCAP* 02.02 (2022), p. 008. DOI: 10.1088/1475-7516/2022/02/008. arXiv: 2110.05530 [astro-ph.CO].

[136] Shi-Fan Chen, Zvonimir Vlah, and Martin White. "Consistent Modeling of Velocity Statistics and Redshift-Space Distortions in One-Loop Perturbation Theory". In: *JCAP* 07 (2020), p. 062. DOI: 10.1088/1475-7516/2020/07/062. arXiv: 2005.00523 [astro-ph.CO].

[137] Shi-Fan Chen, Zvonimir Vlah, and Martin White. "Modeling features in the redshift-space halo power spectrum with perturbation theory". In: *arXiv e-prints*, arXiv:2007.00704 (July 2020), arXiv:2007.00704. arXiv: 2007.00704 [astro-ph.CO].

[138] Shi-Fan Chen, Zvonimir Vlah, and Martin White. "The Lyα forest flux correlation function: a perturbation theory perspective". In: *J. Cosm. Astrop. Phys.* 2021.5, 053 (May 2021), p. 053. DOI: 10.1088/1475-7516/2021/05/053. arXiv: 2103.13498 [astro-ph.CO].

[139] Shi-Fan Chen, Zvonimir Vlah, and Martin White. "The reconstructed power spectrum in the Zeldovich approximation". In: *J. Cosm. Astrop. Phys.* 2019.9, 017 (Sept. 2019), p. 017. DOI: 10.1088/1475-7516/2019/09/017. arXiv: 1907.00043 [astro-ph.CO].

[140] Shi-Fan Chen, Martin White, Joseph DeRose, and Nickolas Kokron. "Cosmological analysis of three-dimensional BOSS galaxy clustering and Planck CMB lensing cross correlations via Lagrangian perturbation theory". In: *JCAP* 07.07 (2022), p. 041. DOI: 10.1088/1475-7516/2022/07/041. arXiv: 2204.10392 [astro-ph.CO].

[141] Shi-Fan Stephen Chen. "Perturbation Theory Models for Precision Cosmology with Large-Scale Structure Surveys". PhD thesis. UC, Berkeley (main), 2022.

[142] Shu-Fan Chen, Hayden Lee, and Cora Dvorkin. *Precise and Accurate Cosmology with CMBxLSS Power Spectra and Bispectra*. 2021. arXiv: 2103.01229 [astro-ph.CO].

[143] X. Chen. "The Tianlai Project: a 21CM Cosmology Experiment". In: *International Journal of Modern Physics Conference Series*. Vol. 12. International Journal of Modern Physics Conference Series. Mar. 2012, pp. 256–263. DOI: 10.1142/S2010194512006459. arXiv: 1212.6278 [astro-ph.IM].

[144] Michel Chevallier and David Polarski. "Accelerating universes with scaling dark matter". In: *Int. J. Mod. Phys. D* 10 (2001), pp. 213–224. DOI: 10.1142/S0218271801000822. arXiv: gr-qc/0009008.

[145] P. R. V. Chintalapati, G. Gutierrez, and M. H. L. S. Wang. "A systematic study of projection biases in the Weak Lensing analysis of cosmic shear and the combination of galaxy clustering and galaxy-galaxy lensing". In: (Mar. 2024). arXiv: 2403.08101 [astro-ph.CO].

[146] Nora Elisa Chisari et al. "Core Cosmology Library: Precision Cosmological Predictions for LSST". In: *Astrophys. J. Suppl.* 242.1 (2019), p. 2. DOI: 10.3847/1538-4365/ab1658. arXiv: 1812.05995 [astro-ph.CO].

[147] Anton Chudaykin and Mikhail M. Ivanov. "Measuring neutrino masses with large-scale structure: Euclid forecast with controlled theoretical error". In: *J. Cosm. Astrop. Phys.* 2019.11, 034 (Nov. 2019), p. 034. DOI: 10.1088/1475-7516/2019/11/034. arXiv: 1907.06666 [astro-ph.CO].

[148] Eegene Chung, Simon Foreman, and Alexander van Engelen. "Baryonic effects on CMB lensing and neutrino mass constraints". In: *Physical Review D* 101.6 (Mar. 2020). ISSN: 2470-0029. DOI: 10.1103/physrevd.101.063534. URL: http://dx.doi.org/10.1103/PhysRevD.101.063534.

[149] Eegene Chung, Simon Foreman, and Alexander van Engelen. "Baryonic effects on CMB lensing and neutrino mass constraints". In: *Phys. Rev. D* 101.6 (2020). [Erratum: Phys.Rev.D 102, 109903 (2020)], p. 063534. DOI: 10.1103/PhysRevD.101.063534. arXiv: 1910.09565 [astro-ph.CO].

[150] Sebastien Clesse, Laura Lopez-Honorez, Christophe Ringeval, Hiroyuki Tashiro, and Michel H. G. Tytgat. "Background reionization history from omniscopes". In: *Phys. Rev. D* 86 (2012), p. 123506. DOI: 10.1103/PhysRevD.86.123506. arXiv: 1208.4277 [astro-ph.CO].

[151] Timothy Clifton, Pedro G. Ferreira, Antonio Padilla, and Constantinos Skordis. "Modified gravity and cosmology". In: *Physics Reports* 513.1-3 (Mar. 2012), pp. 1–189. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2012.01.001. URL: http://dx.doi.org/10.1016/j.physrep.2012.01.001.

[152] Thomas Colas, Guido D'amico, Leonardo Senatore, Pierre Zhang, and Florian Beutler. "Efficient Cosmological Analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure". In: *JCAP* 06 (2020), p. 001. DOI: 10.1088/1475-7516/2020/06/001. arXiv: 1909.07951 [astro-ph.CO].

[153] Planck Collaboration et al. *Planck 2018 results. IX. Constraints on primordial non-Gaussianity*. 2019. arXiv: 1905.05697 [astro-ph.CO].

[154] William R. Coulton et al. "The Atacama Cosmology Telescope: A search for late-time anisotropic screening of the Cosmic Microwave Background". In: (Jan. 2024). arXiv: 2401.13033 [astro-ph.CO].

[155] William R. Coulton et al. "The Atacama Cosmology Telescope: High-resolution component-separated maps across one-third of the sky". In: *arXiv e-prints*, arXiv:2307.01258 (July 2023), arXiv:2307.01258. DOI: 10.48550/arXiv.2307.01258. arXiv: 2307.01258 [astro-ph.CO].

[156] Nathaniel Craig, Daniel Green, Joel Meyers, and Surjeet Rajendran. "No νs is Good News". In: *JHEP* 09 (2024), p. 097. DOI: 10.1007/JHEP09(2024)097. arXiv: 2405.00836 [astro-ph.CO].

[157] Paolo Creminelli and Matias Zaldarriaga. "Single field consistency relation for the 3-point function". In: *JCAP* 10 (2004), p. 006. DOI: 10.1088/1475-7516/2004/10/006. arXiv: astro-ph/0407059.

[158] Devin Crichton et al. "Evidence for the thermal Sunyaev-Zel'dovich effect associated with quasar feedback". In: *Mon. Not. R. Astron. Soc.* 458.2 (May 2016), pp. 1478–1492. DOI: 10.1093/mnras/stw344. arXiv: 1510.05656 [astro-ph.GA].

[159] Neil H. M. Crighton et al. "The neutral hydrogen cosmological mass density at z = 5". In: *Mon. Not. Roy. Astron. Soc.* 452.1 (2015), pp. 217–234. DOI: 10.1093/mnras/stv1182. arXiv: 1506.02037 [astro-ph.CO].

[160] Martín Crocce and Román Scoccimarro. "Nonlinear evolution of baryon acoustic oscillations". In: *Phys. Rev. D* 77.2, 023533 (Jan. 2008), p. 023533. DOI: 10.1103/PhysRevD.77.023533. arXiv: 0704.2783 [astro-ph].

[161] Steven Cunnington, Stefano Camera, and Alkistis Pourtsidou. "The degeneracy between primordial non-Gaussianity and foregrounds in 21 cm intensity mapping experiments". In: *Mon. Not. Roy. Astron. Soc.* 499.3 (2020), pp. 4054–4067. DOI: 10.1093/mnras/staa2986. arXiv: 2007.12126 [astro-ph.CO].

[162] Guido D'Amico, Jérôme Gleyzes, Nickolas Kokron, Katarina Markovic, Leonardo Senatore, Pierre Zhang, Florian Beutler, and Héctor Gil-Marín. "The Cosmological Analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure". In: *JCAP* 05 (2020), p. 005. DOI: 10.1088/1475-7516/2020/05/005. arXiv: 1909.05271 [astro-ph.CO].

[163] Guido D'Amico, Leonardo Senatore, Pierre Zhang, and Henry Zheng. "The Hubble Tension in Light of the Full-Shape Analysis of Large-Scale Structure Data". In: *arXiv e-prints*, arXiv:2006.12420 (June 2020), arXiv:2006.12420. arXiv: 2006.12420 [astro-ph.CO].

[164] Roohi Dalal et al. "Hyper Suprime-Cam Year 3 results: Cosmology from cosmic shear power spectra". In: *Phys. Rev. D* 108.12 (2023), p. 123519. DOI: 10.1103/PhysRevD.108.123519. arXiv: 2304.00701 [astro-ph.CO].

[165] Omar Darwish et al. "The Atacama Cosmology Telescope: A CMB lensing mass map over 2100 square degrees of sky and its cross-correlation with BOSS-CMASS galaxies". In: *Mon. Not. Roy. Astron. Soc.* 500.2 (2020), pp. 2250–2263. DOI: 10.1093/mnras/staa3438. arXiv: 2004.01139 [astro-ph.CO].

[166] Omar Darwish, Blake D. Sherwin, Noah Sailer, Emmanuel Schaan, and Simone Ferraro. "Optimizing foreground mitigation for CMB lensing with combined multifrequency and geometric methods". In: *arXiv e-prints*, arXiv:2111.00462 (Oct. 2021), arXiv:2111.00462. arXiv: 2111.00462 [astro-ph.CO].

[167] Sudeep Das, Amir Hajian, and David N. Spergel. "Efficient Power Spectrum Estimation for High Resolution CMB Maps". In: *Phys. Rev. D* 79 (2009), p. 083008. DOI: 10.1103/PhysRevD.79.083008. arXiv: 0809.1092 [astro-ph].

[168] Rahul Datta et al. "The Atacama Cosmology Telescope: Two-season ACTPol Extragalactic Point Sources and their Polarization properties". In: *Mon. Not. Roy. Astron. Soc.* 486.4 (2019), pp. 5239–5262. DOI: 10.1093/mnras/sty2934. arXiv: 1811.01854 [astro-ph.CO].

[169] Kyle Dawson, Josh Frieman, Katrin Heitmann, Bhuvnesh Jain, Steve Kahn, Rachel Mandelbaum, Saul Perlmutter, and Anže Slosar. "Cosmic Visions Dark Energy: Small Projects Portfolio". In: (Feb. 2018). arXiv: 1802.07216 [astro-ph.CO].

[170] Kyle S. Dawson et al. "The Baryon Oscillation Spectroscopic Survey of SDSS-III". In: *Astron. J.* 145.1, 10 (Jan. 2013), p. 10. DOI: 10.1088/0004-6256/145/1/10. arXiv: 1208.0022 [astro-ph.CO].

[171] Roger de Belsunce, Steven Gratton, and George Efstathiou. "Testing for spectral index variations in polarized CMB foregrounds". In: *Mon. Not. R. Astron. Soc.* 517.2 (Dec. 2022), pp. 2855–2866. DOI: 10.1093/mnras/stac2825. arXiv: 2205.13968 [astro-ph.CO].

[172] F. De Bernardis et al. "Detection of the pairwise kinematic Sunyaev-Zel'dovich effect with BOSS DR11 and the Atacama Cosmology Telescope". In: *J. Cosm. Astrop. Phys.* 2017.3, 008 (Mar. 2017), p. 008. DOI: 10.1088/1475-7516/2017/03/008. arXiv: 1607.02139 [astro-ph.CO].

[173] Anna de Graaff, Yan-Chuan Cai, Catherine Heymans, and John A. Peacock. "Probing the missing baryons with the Sunyaev-Zel'dovich effect from filaments". In: *Astron. Astrophys.* 624, A48 (Apr. 2019), A48. DOI: 10.1051/0004-6361/201935159. arXiv: 1709.10378 [astro-ph.CO].

[174] David R. DeBoer et al. "Hydrogen Epoch of Reionization Array (HERA)". In: *Publ. Astron. Soc. Pac.* 129.974 (2017), p. 045001. DOI: 10.1088/1538-3873/129/974/045001. arXiv: 1606.07473 [astro-ph.IM].

[175] J. Delabrouille and J. -F. Cardoso. "Diffuse source separation in CMB observations". In: *arXiv e-prints*, astro-ph/0702198 (Feb. 2007), astro–ph/0702198. arXiv: astro-ph/0702198 [astro-ph].

[176] J. M. Delouis, L. Pagano, S. Mottet, J. L. Puget, and L. Vibert. "SRoll2: an improved mapmaking approach to reduce large-scale systematic effects in the Planck High Frequency Instrument legacy maps". In: *Astron. Astrophys.* 629, A38 (Sept. 2019), A38. DOI: 10.1051/0004-6361/201834882. arXiv: 1901.11386 [astro-ph.CO].

[177] J. DeRose et al. "Dark Energy Survey Year 3 results: Cosmology from combined galaxy clustering and lensing validation on cosmological simulations". In: *Phys. Rev. D* 105.12 (2022), p. 123520. DOI: 10.1103/PhysRevD.105.123520. arXiv: 2105.13547 [astro-ph.CO].

[178] J. DeRose et al. "The Buzzard Flock: Dark Energy Survey Synthetic Sky Catalogs". In: (Jan. 2019). arXiv: 1901.02401 [astro-ph.CO].

[179] Joeseph DeRose et al. "Mocking DESI cross correlations using the Buzzard simulations, in preparation". In: (2023). arXiv: 24MM.XXXXX [astro-ph.CO].

[180] Joseph DeRose, Matthew R. Becker, and Risa H. Wechsler. "Modeling Redshift-space Clustering with Abundance Matching". In: *Astrophys. J.* 940.1 (2022), p. 13. DOI: 10.3847/1538-4357/ac9968. arXiv: 2105.12104 [astro-ph.CO].

[181] Joseph DeRose et al. "Aemulus ν: precise predictions for matter and biased tracer power spectra in the presence of neutrinos". In: *JCAP* 07 (2023), p. 054. DOI: 10.1088/1475-7516/2023/07/054. arXiv: 2303.09762 [astro-ph.CO].

[182] DESI Collaboration. "Evidence for evolving dark energy from DESI". In: (2023). arXiv: 2307.67890 [astro-ph.CO].

[183] DESI Collaboration et al. "DESI 2024 V: Full-Shape Galaxy Clustering from Galaxies and Quasars". In: *arXiv e-prints*, arXiv:2411.12021 (Nov. 2024), arXiv:2411.12021. DOI: 10.48550/arXiv.2411.12021. arXiv: 2411.12021 [astro-ph.CO].

[184] DESI Collaboration et al. "DESI DR2 Results II: Measurements of Baryon Acoustic Oscillations and Cosmological Constraints". In: *arXiv e-prints*, arXiv:2503.14738 (Mar. 2025), arXiv:2503.14738. arXiv: 2503.14738 [astro-ph.CO].

[185] DESI Collaboration et al. "Overview of the Instrumentation for the Dark Energy Spectroscopic Instrument". In: *Astron. J.* 164.5, 207 (Nov. 2022), p. 207. DOI: 10.3847/1538-3881/ac882b. arXiv: 2205.10939 [astro-ph.IM].

[186] DESI Collaboration et al. "The DESI Experiment Part I: Science,Targeting, and Survey Design". In: *arXiv e-prints*, arXiv:1611.00036 (Oct. 2016), arXiv:1611.00036. DOI: 10.48550/arXiv.1611.00036. arXiv: 1611.00036 [astro-ph.IM].

[187] DESI Collaboration et al. "The DESI Experiment Part II: Instrument Design". In: *arXiv e-prints*, arXiv:1611.00037 (Oct. 2016), arXiv:1611.00037. DOI: 10.48550/arXiv.1611.00037. arXiv: 1611.00037 [astro-ph.IM].

[188] DESI Collaboration et al. "Validation of the Scientific Program for the Dark Energy Spectroscopic Instrument". In: *Astron. J.* 167.2, 62 (Feb. 2024), p. 62. DOI: 10.3847/1538-3881/ad0b08. arXiv: 2306.06307 [astro-ph.CO].

[189] V. Desjacques, D. Jeong, and F. Schmidt. "Large-Scale Galaxy Bias". In: *ArXiv e-prints* (Nov. 2016). arXiv: 1611.09787.

[190] Anne-Sylvie Deutsch, Matthew C. Johnson, Moritz Münchmeyer, and Alexandra Terrana. "Polarized Sunyaev Zel'dovich tomography". In: *JCAP* 04 (2018), p. 034. DOI: 10.1088/1475-7516/2018/04/034. arXiv: 1705.08907 [astro-ph.CO].

[191] Arjun Dey et al. "Overview of the DESI Legacy Imaging Surveys". In: *Astron. J.* 157.5, 168 (May 2019), p. 168. DOI: 10.3847/1538-3881/ab089d. arXiv: 1804.08657 [astro-ph.IM].

[192] Eleonora Di Valentino et al. "Cosmology Intertwined II: The Hubble Constant Tension". In: *Astroparticle Physics* 131 (2021), p. 102605. DOI: 10.1016/j.astropartphys.2021.102605. arXiv: 2008.11284.

[193] Eleonora Di Valentino, Alessandro Melchiorri, Olga Mena, and Sunny Vagnozzi. "Nonminimal dark sector physics and cosmological tensions". In: *Phys. Rev. D* 101.6 (2020), p. 063502. DOI: 10.1103/PhysRevD.101.063502. arXiv: 1910.09853 [astro-ph.CO].

[194] Scott Dodelson. *Modern Cosmology*. 2003.

[195] Scott Dodelson and Jay M. Jubas. "Reionization and Its Imprint on the Cosmic Microwave Background". In: *Astrophys. J.* 439 (Feb. 1995), p. 503. DOI: 10.1086/175191. arXiv: astro-ph/9308019 [astro-ph].

[196] Jamie Donald-McCann, Rafaela Gsponer, Ruiyang Zhao, Kazuya Koyama, and Florian Beutler. "Analysis of unified galaxy power spectrum multipole measurements". In: *Mon. Not. R. Astron. Soc.* 526.3 (Dec. 2023), pp. 3461–3481. DOI: 10.1093/mnras/stad2957. arXiv: 2307.07475 [astro-ph.CO].

[197] Olivier Doré et al. "Cosmology with the SPHEREX All-Sky Spectral Survey". In: *arXiv e-prints*, arXiv:1412.4872 (Dec. 2014), arXiv:1412.4872. arXiv: 1412.4872 [astro-ph.CO].

[198] Olivier Doré et al. "Science Impacts of the SPHEREx All-Sky Optical to Near-Infrared Spectral Survey: Report of a Community Workshop Examining Extragalactic, Galactic, Stellar and Planetary Science". In: *arXiv e-prints*, arXiv:1606.07039 (June 2016), arXiv:1606.07039. arXiv: 1606.07039 [astro-ph.CO].

[199] Aristide Doussot, Hy Trac, and Renyue Cen. "SCORCH. II. Radiation-hydrodynamic Simulations of Reionization with Varying Radiation Escape Fractions". In: *Astrophys. J.* 870.1 (2019), p. 18. DOI: 10.3847/1538-4357/aaef75. arXiv: 1712.04464 [astro-ph.CO].

[200] J. Dunkley et al. "The Atacama Cosmology Telescope: likelihood for small-scale CMB data". In: *JCAP* 07 (2013), p. 025. DOI: 10.1088/1475-7516/2013/07/025. arXiv: 1301.0776 [astro-ph.CO].

[201] Cora Dvorkin, Wayne Hu, and Kendrick M. Smith. "B-mode CMB polarization from patchy screening during reionization". In: *Phys. Rev. D* 79.10, 107302 (May 2009), p. 107302. DOI: 10.1103/PhysRevD.79.107302. arXiv: 0902.4413 [astro-ph.CO].

[202] Cora Dvorkin and Kendrick M. Smith. "Reconstructing Patchy Reionization from the Cosmic Microwave Background". In: *Phys. Rev. D* 79 (2009), p. 043003. DOI: 10.1103/PhysRevD.79.043003. arXiv: 0812.1566 [astro-ph].

[203] Cora Dvorkin et al. "Neutrino Mass from Cosmology: Probing Physics Beyond the Standard Model". In: *Bull. Am. Astron. Soc.* 51.3, 64 (May 2019), p. 64. arXiv: 1903.03689 [astro-ph.CO].

[204] Haruki Ebina and Martin White. "Cosmology before noon with multiple galaxy populations". In: (Jan. 2024). arXiv: 2401.13166 [astro-ph.CO].

[205] eBOSS Collaboration et al. "The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological Implications from two Decades of Spectroscopic Surveys at the Apache Point observatory". In: *arXiv e-prints*, arXiv:2007.08991 (July 2020), arXiv:2007.08991. arXiv: 2007.08991 [astro-ph.CO].

[206] D. Eckert et al. "Non-thermal pressure support in X-COP galaxy clusters". In: *Astron. Astrophys.* 621, A40 (Jan. 2019), A40. DOI: 10.1051/0004-6361/201833324. arXiv: 1805.00034 [astro-ph.CO].

[207] G. Efstathiou. "A Lockdown Perspective on the Hubble Tension (with comments from the SH0ES team)". In: *arXiv e-prints*, arXiv:2007.10716 (July 2020), arXiv:2007.10716. arXiv: 2007.10716 [astro-ph.CO].

[208] George Efstathiou and Steven Gratton. "A Detailed Description of the CamSpec Likelihood Pipeline and a Reanalysis of the Planck High Frequency Maps". In: *arXiv e-prints* (Oct. 2019). DOI: 10.21105/astro.1910.00483. arXiv: 1910.00483 [astro-ph.CO].

[209] D. J. Eisenstein, H.-J. Seo, E. Sirko, and D. N. Spergel. "Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak". In: *Astrophys. J.* 664 (Aug. 2007), pp. 675–679. DOI: 10.1086/518712. eprint: astro-ph/0604362.

[210] D. J. Eisenstein, H.-J. Seo, and M. White. "On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter". In: *Astrophys. J.* 664 (Aug. 2007), pp. 660–674. DOI: 10.1086/518755. eprint: astro-ph/0604361.

[211] W. Elbers et al. "Constraints on Neutrino Physics from DESI DR2 BAO and DR1 Full Shape". In: *arXiv e-prints*, arXiv:2503.14744 (Mar. 2025), arXiv:2503.14744. DOI: 10.48550/arXiv.2503.14744. arXiv: 2503.14744 [astro-ph.CO].

[212] Willem Elbers, Carlos S. Frenk, Adrian Jenkins, Baojiu Li, and Silvia Pascoli. "An optimal non-linear method for simulating relic neutrinos". In: *Mon. Not. Roy. Astron. Soc.* 507.2 (2021), pp. 2614–2631. DOI: 10.1093/mnras/stab2260. arXiv: 2010.07321 [astro-ph.CO].

[213] Willem Elbers, Carlos S. Frenk, Adrian Jenkins, Baojiu Li, and Silvia Pascoli. "Higher order initial conditions with massive neutrinos". In: *Mon. Not. Roy. Astron. Soc.* 516.3 (2022), pp. 3821–3836. DOI: 10.1093/mnras/stac2365. arXiv: 2202.00670 [astro-ph.CO].

[214] Willem Elbers, Carlos S. Frenk, Adrian Jenkins, Baojiu Li, and Silvia Pascoli. "Negative neutrino masses as a mirage of dark energy". In: *Phys. Rev. D* 111.6 (2025), p. 063534. DOI: 10.1103/PhysRevD.111.063534. arXiv: 2407.10965 [astro-ph.CO].

[215] Richard Ellis and Kyle Dawson. "SpecTel: A 10-12 meter class Spectroscopic Survey Telescope". In: *Bull. Am. Astron. Soc.* Vol. 51. Sept. 2019, 45, p. 45. arXiv: 1907.06797 [astro-ph.IM].

[216] H. K. Eriksen, A. J. Banday, K. M. Gorski, and P. B. Lilje. "On Foreground Removal from theWilkinson Microwave Anisotropy ProbeData by an Internal Linear Combination Method: Limitations and Implications". In: *The Astrophysical Journal* 612.2 (Sept. 2004), pp. 633–646. ISSN: 1538-4357. DOI: 10.1086/422807. URL: http://dx.doi.org/10.1086/422807.

[217] Thomas Essinger-Hileman et al. "CLASS: the cosmology large angular scale surveyor". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII.* Ed. by Wayne S. Holland and Jonas Zmuidzinas. Vol. 9153. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. July 2014, 91531I, p. 91531I. DOI: 10.1117/12.2056701. arXiv: 1408.4788 [astro-ph.IM].

[218] Ivan Esteban, M. C. Gonzalez-Garcia, Michele Maltoni, Ivan Martinez-Soler, João Paulo Pinheiro, and Thomas Schwetz. "NuFit-6.0: updated global analysis of three-flavor neutrino oscillations". In: *JHEP* 12 (2024), p. 216. DOI: 10.1007/JHEP12(2024)216. arXiv: 2410.05380 [hep-ph].

[219] Giulio Fabbian, Matteo Calabrese, and Carmelita Carbone. "CMB weak-lensing beyond the Born approximation: a numerical approach". In: *JCAP* 02 (2018), p. 050. DOI: 10.1088/1475-7516/2018/02/050. arXiv: 1702.03317 [astro-ph.CO].

[220] Giulio Fabbian, Julien Carron, Antony Lewis, and Margherita Lembo. "Lensed CMB power spectrum biases from masking extragalactic sources". In: *Physical Review D* 103.4 (Feb. 2021). ISSN: 2470-0029. DOI: 10.1103/physrevd.103.043535. URL: http://dx.doi.org/10.1103/PhysRevD.103.043535.

[221] Giulio Fabbian, Antony Lewis, and Dominic Beck. "CMB lensing reconstruction biases in cross-correlation with large-scale structure probes". In: *JCAP* 10 (2019), p. 057. DOI: 10.1088/1475-7516/2019/10/057. arXiv: 1906.08760 [astro-ph.CO].

[222] Xiaohui Fan et al. "Constraining the Evolution of the Ionizing Background and the Epoch of Reionization with z~6 Quasars. II. A Sample of 19 Quasars". In: *Astron. J.* 132.1 (July 2006), pp. 117–136. DOI: 10.1086/504836. arXiv: astro-ph/0512082 [astro-ph].

[223] Gerrit S. Farren et al. "The Atacama Cosmology Telescope: Cosmology from Cross-correlations of unWISE Galaxies and ACT DR6 CMB Lensing". In: *Astrophys. J.* 966.2, 157 (May 2024), p. 157. DOI: 10.3847/1538-4357/ad31a5. arXiv: 2309.05659 [astro-ph.CO].

[224] Chang Feng and Gilbert Holder. "Polarization of the Cosmic Infrared Background Fluctuations". In: *Astrophys. J.* 897.2 (2020), p. 140. DOI: 10.3847/1538-4357/ab9013. arXiv: 1907.12085 [astro-ph.CO].

[225] Chang Feng and Gilbert Holder. "Searching for patchy reionization from cosmic microwave background with hybrid quadratic estimators". In: *Phys. Rev. D* 99.12, 123502 (June 2019), p. 123502. DOI: 10.1103/PhysRevD.99.123502. arXiv: 1808.01592 [astro-ph.CO].

[226] Yu Feng, Simeon Bird, Lauren Anderson, Andreu Font-Ribera, and Chris Pedersen. *MP-Gadget/MP-Gadget: A tag for getting a DOI*. Version FirstDOI. Oct. 2018. DOI: 10.5281/zenodo.1451799. URL: https://doi.org/10.5281/zenodo.1451799.

[227] S. Ferraro and J. C. Hill. "Bias to CMB Lensing Reconstruction from Temperature Anisotropies due to Large-Scale Galaxy Motions". In: *ArXiv e-prints* (May 2017). arXiv: 1705.06751.

[228] Simone Ferraro et al. "Inflation and Dark Energy from Spectroscopy at $z > 2$". In: (Mar. 2019). arXiv: 1903.09208 [astro-ph.CO].

[229] Simone Ferraro, J. Colin Hill, Nick Battaglia, Jia Liu, and David N. Spergel. "Kinematic Sunyaev-Zel'dovich effect with projected fields. II. Prospects, challenges, and comparison with simulations". In: *Phys. Rev. D* 94.12, 123526 (Dec. 2016), p. 123526. DOI: 10.1103/PhysRevD.94.123526. arXiv: 1605.02722 [astro-ph.CO].

[230] Simone Ferraro and Kendrick M. Smith. "Characterizing the epoch of reionization with the small-scale CMB: Constraints on the optical depth and duration". In: *Phys. Rev. D* 98.12 (2018), p. 123519. DOI: 10.1103/PhysRevD.98.123519. arXiv: 1803.07036 [astro-ph.CO].

[231] Christian Fidler, Nils Sujata, and Maria Archidiacono. "Relativistic bias in neutrino cosmologies". In: *JCAP* 06 (2019), p. 035. DOI: 10.1088/1475-7516/2019/06/035. arXiv: 1812.09266 [astro-ph.CO].

[232] George B. Field. "An Attempt to Observe Neutral Hydrogen Between the Galaxies." In: *Astrophys. J.* 129 (May 1959), p. 525. DOI: 10.1086/146652.

[233] Karl B. Fisher, Caleb A. Scharf, and Ofer Lahav. "A Spherical Harmonic Approach to Redshift Distortion and a Measurement of Omega from the 1.2 Jy IRAS Redshift Survey". In: *Mon. Not. Roy. Astron. Soc.* 266 (1994), pp. 219–226. DOI: 10.1093/mnras/266.1.219. arXiv: astro-ph/9309027.

[234] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright. "The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set". In: *Astrophys. J.* 473 (Dec. 1996), p. 576. DOI: 10.1086/178173. arXiv: astro-ph/9605054 [astro-ph].

[235] Andreu Font-Ribera, Patrick McDonald, Nick Mostek, Beth A. Reid, Hee-Jong Seo, and An Slosar. "DESI and other dark energy experiments in the era of neutrino mass measurements". In: *JCAP* 05 (2014), p. 023. DOI: 10.1088/1475-7516/2014/05/023. arXiv: 1308.4164 [astro-ph.CO].

[236] S. Foreman and L. Senatore. "The EFT of Large Scale Structures at all redshifts: analytical predictions for lensing". In: *J. Cosm. Astrop. Phys.* 4, 033 (Apr. 2016), p. 033. DOI: 10.1088/1475-7516/2016/04/033. arXiv: 1503.01775.

[237] Simon Foreman, P. Daniel Meerburg, Alexander van Engelen, and Joel Meyers. "Lensing reconstruction from line intensity maps: the impact of gravitational nonlinearity". In: *J. Cosm. Astrop. Phys.* 2018.7, 046 (July 2018), p. 046. DOI: 10.1088/1475-7516/2018/07/046. arXiv: 1803.04975 [astro-ph.CO].

[238] Masataka Fukugita and P. J. E. Peebles. "The Cosmic Energy Inventory". In: *Astrophys. J.* 616.2 (Dec. 2004), pp. 643–668. DOI: 10.1086/425155. arXiv: astro-ph/0406095 [astro-ph].

[239] Steven Furlanetto, S. Peng Oh, and Frank Briggs. "Cosmology at Low Frequencies: The 21 cm Transition and the High-Redshift Universe". In: *Phys. Rept.* 433 (2006), pp. 181–301. DOI: 10.1016/j.physrep.2006.08.002. arXiv: astro-ph/0608032.

[240] Gaia Collaboration et al. "Gaia Data Release 2. Summary of the contents and survey properties". In: *Astron. Astrophys.* 616, A1 (Aug. 2018), A1. DOI: 10.1051/0004-6361/201833051. arXiv: 1804.09365 [astro-ph.GA].

[241] Gaia Collaboration et al. "Gaia Early Data Release 3. Summary of the contents and survey properties (Corrigendum)". In: *Astron. Astrophys.* 650, C3 (June 2021), p. C3. DOI: 10.1051/0004-6361/202039657e.

[242] N. Galitzki et al. "The Simons Observatory: instrument overview". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Vol. 10708. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series. July 2018, 1070804, p. 1070804. DOI: 10.1117/12.2312985. arXiv: 1808.04493 [astro-ph.IM].

[243] H. Gan et al. "Statistical analysis of the causes of excess variance in the 21 cm signal power spectra obtained with the Low-Frequency Array". In: *arXiv e-prints*, arXiv:2203.02345 (Mar. 2022), arXiv:2203.02345. arXiv: 2203.02345 [astro-ph.CO].

[244] Carlos García-García, Jaime Ruiz Zapatero, David Alonso, Emilio Bellini, Pedro G. Ferreira, Eva-Maria Mueller, Andrina Nicola, and Pilar Ruiz-Lapuente. "The growth of density perturbations in the last ∼10 billion years from tomographic large-scale structure data". In: *JCAP* 10 (2021), p. 030. DOI: 10.1088/1475-7516/2021/10/030. arXiv: 2105.12108 [astro-ph.CO].

[245] Andrew Gelman and Donald B. Rubin. "Inference from Iterative Simulation Using Multiple Sequences". In: *Statistical Science* 7.4 (1992), pp. 457–472. DOI: 10.1214/ss/1177011136. URL: https://doi.org/10.1214/ss/1177011136.

[246] Andrew B. Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Boca Ratan, Florida: Chapman and Hall/CRC, 1995.

[247] V. Ghirardini et al. "The SRG/eROSITA All-Sky Survey: Cosmology Constraints from Cluster Abundances in the Western Galactic Hemisphere". In: *arXiv e-prints*, arXiv:2402.08458 (Feb. 2024), arXiv:2402.08458. DOI: 10.48550/arXiv.2402.08458. arXiv: 2402.08458 [astro-ph.CO].

[248] William Giarè, Eleonora Di Valentino, and Alessandro Melchiorri. "Measuring the reionization optical depth without large-scale CMB polarization". In: *Phys. Rev. D* 109.10, 103519 (May 2024), p. 103519. DOI: 10.1103/PhysRevD.109.103519. arXiv: 2312.06482 [astro-ph.CO].

[249] Jérôme Gleyzes, Roland de Putter, Daniel Green, and Olivier Doré. "Biasing and the search for primordial non-Gaussianity beyond the local type". In: *Journal of Cosmology and Astroparticle Physics* 2017.04 (Apr. 2017), pp. 002–002. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2017/04/002. URL: http://dx.doi.org/10.1088/1475-7516/2017/04/002.

[250] Vera Gluscevic, Marc Kamionkowski, and Duncan Hanson. "Patchy screening of the cosmic microwave background by inhomogeneous reionization". In: *Phys. Rev. D* 87.4, 047303 (Feb. 2013), p. 047303. DOI: 10.1103/PhysRevD.87.047303. arXiv: 1210.5507 [astro-ph.CO].

[251] Samuel Goldstein, J. Colin Hill, Vid Iršič, and Blake D. Sherwin. "Canonical Hubble-Tension-Resolving Early Dark Energy Cosmologies Are Inconsistent with the Lyman-α Forest". In: *Phys. Rev. Lett.* 131.20 (2023), p. 201001. DOI: 10.1103/PhysRevLett.131.201001. arXiv: 2303.00746 [astro-ph.CO].

[252] K. M. Górski, 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". In: *Astrophys. J.* 622 (2005), pp. 759–771. DOI: 10.1086/427976. arXiv: astro-ph/0409513.

[253] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. "HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere". In: *Astrophys. J.* 622 (Apr. 2005), pp. 759–771. DOI: 10.1086/427976. eprint: arXiv:astro-ph/0409513.

[254] Steven Gratton and Anthony Challinor. "Understanding parameter differences between analyses employing nested data subsets". In: *Mon. Not. R. Astron. Soc.* 499.3 (Dec. 2020), pp. 3410–3416. DOI: 10.1093/mnras/staa2996. arXiv: 1911.07754 [astro-ph.IM].

[255] Daniel Green and Joel Meyers. "The Cosmological Preference for Negative Neutrino Mass". In: (July 2024). arXiv: 2407.07878 [astro-ph.CO].

[256] Daniel Green et al. "Messengers from the Early Universe: Cosmic Neutrinos and Other Light Relics". In: *Bull. Am. Astron. Soc.* 51.3, 159 (May 2019), p. 159. arXiv: 1903.04763 [astro-ph.CO].

[257] Bradley Greig et al. "Interpreting LOFAR 21-cm signal upper limits at z ≈ 9.1 in the context of high-z galaxy and reionization observations". In: *Mon. Not. Roy. Astron. Soc.* 501.1 (2021), pp. 1–13. DOI: 10.1093/mnras/staa3593. arXiv: 2006.03203 [astro-ph.CO].

[258] Bradley Greig, Andrei Mesinger, and Eduardo Bañados. "Constraints on reionization from the z = 7.5 QSO ULASJ1342+0928". In: *Mon. Not. Roy. Astron. Soc.* 484.4 (2019), pp. 5094–5101. DOI: 10.1093/mnras/stz230. arXiv: 1807.01593 [astro-ph.CO].

[259] Bradley Greig, Andrei Mesinger, Zoltán Haiman, and Robert A. Simcoe. "Are we witnessing the epoch of reionization at z=7.1 from the spectrum of J1120+0641?" In: *Mon. Not. Roy. Astron. Soc.* 466.4 (2017), pp. 4239–4249. DOI: 10.1093/mnras/stw3351. arXiv: 1606.00441 [astro-ph.CO].

[260] Bradley Greig, Cathryn M. Trott, Nichole Barry, Simon J. Mutch, Bart Pindor, Rachel L. Webster, and J. Stuart B. Wyithe. "Exploring reionization and high-z galaxy observables with recent multiredshift MWA upper limits on the 21-cm signal". In: *Mon. Not. Roy. Astron. Soc.* 500.4 (2020), pp. 5322–5335. DOI: 10.1093/mnras/staa3494. arXiv: 2008.02639 [astro-ph.CO].

[261] Andrei Gruzinov and Wayne Hu. "Secondary Cosmic Microwave Background Anisotropies in a Universe Reionized in Patches". In: *Astrophys. J.* 508.2 (Dec. 1998), pp. 435–439. DOI: 10.1086/306432. arXiv: astro-ph/9803188 [astro-ph].

[262] James E. Gunn and Bruce A. Peterson. "On the Density of Neutral Hydrogen in Intergalactic Space." In: *Astrophys. J.* 142 (Nov. 1965), pp. 1633–1636. DOI: 10.1086/148444.

[263] N. Gupta et al. "Fractional polarization of extragalactic sources in the 500 deg2 SPTpol survey". In: *Mon. Not. Roy. Astron. Soc.* 490.4 (2019), pp. 5712–5721. DOI: 10.1093/mnras/stz2905. arXiv: 1907.02156 [astro-ph.CO].

[264] J. Guy et al. "The Spectroscopic Data Processing Pipeline for the Dark Energy Spectroscopic Instrument". In: *Astron. J.* 165.4, 144 (Apr. 2023), p. 144. DOI: 10.3847/1538-3881/acb212. arXiv: 2209.14482 [astro-ph.IM].

[265] B. Hadzhiyska et al. "Evidence for large baryonic feedback at low and intermediate redshifts from kinematic Sunyaev-Zel'dovich observations with ACT and DESI photometric galaxies". In: (July 2024). arXiv: 2407.07152 [astro-ph.CO].

[266] Boryana Hadzhiyska, Simone Ferraro, Bernardita Ried Guachalla, and Emmanuel Schaan. "Velocity reconstruction in the era of DESI and Rubin/LSST. II. Realistic samples on the light cone". In: *Phys. Rev. D* 109.10, 103534 (May 2024), p. 103534. DOI: 10.1103/PhysRevD.109.103534. arXiv: 2312.12434 [astro-ph.CO].

[267] Boryana Hadzhiyska, Simone Ferraro, and Rongpu Zhou. "Tracing cosmic gas in filaments and halos: Low-redshift insights from the kinematic Sunyaev-Zel'dovich effect". In: *Phys. Rev. D* 111.2, 023534 (Jan. 2025), p. 023534. DOI: 10.1103/PhysRevD.111.023534. arXiv: 2412.03631 [astro-ph.CO].

[268] Boryana Hadzhiyska, Carlos García-García, David Alonso, Andrina Nicola, and Anže Slosar. "Hefty enhancement of cosmological constraints from the DES Y1 data using a hybrid effective field theory approach to galaxy bias". In: *J. Cosm. Astrop. Phys.* 2021.9, 020 (Sept. 2021), p. 020. DOI: 10.1088/1475-7516/2021/09/020. arXiv: 2103.09820 [astro-ph.CO].

[269] Boryana Hadzhiyska, Lehman H. Garrison, Daniel Eisenstein, and Sownak Bose. "The halo light-cone catalogues of ABACUSSUMMIT". In: *Mon. Not. R. Astron. Soc.* 509.2 (Jan. 2022), pp. 2194–2208. DOI: 10.1093/mnras/stab3066. arXiv: 2110.11413 [astro-ph.CO].

[270] Boryana Hadzhiyska, Noah Sailer, and Simone Ferraro. "Mapping the gas density with the kinematic Sunyaev-Zel'dovich and patchy screening effects: a self-consistent comparison". In: *arXiv e-prints* (June 2025). arXiv: 2506.17379 [astro-ph.CO].

[271] Boryana Hadzhiyska, Kevin Wolz, Susanna Azzoni, David Alonso, Carlos García-García, Jaime Ruiz-Zapatero, and Anže Slosar. "Cosmology with 6 parameters in the Stage-IV era: efficient marginalisation over nuisance parameters". In: *The Open Journal of Astrophysics* 6, 23 (July 2023), p. 23. DOI: 10.21105/astro.2301.11895. arXiv: 2301.11895 [astro-ph.CO].

[272] Boryana Hadzhiyska et al. "Interpreting Sunyaev-Zel'dovich observations with MillenniumTNG: mass and environment scaling relations". In: *Mon. Not. R. Astron. Soc.* 526.1 (Nov. 2023), pp. 369–382. DOI: 10.1093/mnras/stad2751. arXiv: 2305.00992 [astro-ph.CO].

[273] Boryana Hadzhiyska et al. "Synthetic light-cone catalogues of modern redshift and weak lensing surveys waith ABACUSSUMMIT". In: *Mon. Not. R. Astron. Soc.* 525.3 (Nov. 2023), pp. 4367–4387. DOI: 10.1093/mnras/stad2563. arXiv: 2305.11935 [astro-ph.CO].

[274] ChangHoon Hahn et al. "The DESI Bright Galaxy Survey: Final Target Selection, Design, and Validation". In: *Astron. J.* 165.6 (2023), p. 253. DOI: 10.3847/1538-3881/accff8. arXiv: 2208.08512 [astro-ph.CO].

[275] Oliver Hahn, Cornelius Rampf, and Cora Uhlemann. "Higher order initial conditions for mixed baryon–CDM simulations". In: *Mon. Not. Roy. Astron. Soc.* 503.1 (2021), pp. 426–445. DOI: 10.1093/mnras/staa3773. arXiv: 2008.09124 [astro-ph.CO].

[276] Shaul Hanany et al. "PICO: Probe of Inflation and Cosmic Origins". In: (Mar. 2019). arXiv: 1902.10541 [astro-ph.IM].

[277] Nick Hand et al. "Evidence of Galaxy Cluster Motions with the Kinematic Sunyaev-Zel'dovich Effect". In: *Phys. Rev. Lett.* 109.4, 041101 (July 2012), p. 041101. DOI: 10.1103/PhysRevLett.109.041101. arXiv: 1203.4219 [astro-ph.CO].

[278] Qianjun Hang, Shadab Alam, John A. Peacock, and Yan-Chuan Cai. "Galaxy clustering in the DESI Legacy Survey and its imprint on the CMB". In: *Mon. Not. Roy. Astron. Soc.* 501.1 (2021), pp. 1481–1498. DOI: 10.1093/mnras/staa3738. arXiv: 2010.00466 [astro-ph.CO].

[279] Duncan Hanson, Anthony Challinor, and Antony Lewis. "Weak lensing of the CMB". In: *General Relativity and Gravitation* 42.9 (Sept. 2010), pp. 2197–2218. DOI: 10.1007/s10714-010-1036-y. arXiv: 0911.0612 [astro-ph.CO].

[280] Charles R Harris et al. "Array programming with NumPy". In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.

[281] Masashi Hazumi and et al. "LiteBIRD: A Satellite for the Studies of B-Mode Polarization and Inflation from Cosmic Background Radiation Detection". In: *Journal of Low Temperature Physics* 194 (2019), pp. 443–452.

[282] Chen Heinrich and Wayne Hu. "Reionization effective likelihood from Planck 2018 data". In: *Phys. Rev. D* 104.6 (2021), p. 063505. DOI: 10.1103/PhysRevD.104.063505. arXiv: 2104.13998 [astro-ph.CO].

[283] S. W. Henderson et al. "Advanced ACTPol Cryogenic Detector Arrays and Readout". In: *Journal of Low Temperature Physics* 184.3-4 (Aug. 2016), pp. 772–779. DOI: 10.1007/s10909-016-1575-z. arXiv: 1510.02809 [astro-ph.IM].

[284] J. W. Henning et al. "Measurements of the Temperature and E-Mode Polarization of the CMB from 500 Square Degrees of SPTpol Data". In: *Astrophys. J.* 852.2 (2018), p. 97. DOI: 10.3847/1538-4357/aa9ff4. arXiv: 1707.09353 [astro-ph.CO].

[285] Catherine Heymans et al. "CFHTLenS tomographic weak lensing cosmological parameter constraints: Mitigating the impact of intrinsic galaxy alignments". In: *Mon. Not. Roy. Astron. Soc.* 432 (2013), p. 2433. DOI: 10.1093/mnras/stt601. arXiv: 1303.1808 [astro-ph.CO].

[286] Catherine Heymans et al. "KiDS-1000 Cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints". In: *Astron. Astrophys.* 646 (2021), A140. DOI: 10.1051/0004-6361/202039063. arXiv: 2007.15632 [astro-ph.CO].

[287] J. Colin Hill, Simone Ferraro, Nick Battaglia, Jia Liu, and David N. Spergel. "Kinematic Sunyaev-Zel'dovich Effect with Projected Fields: A Novel Probe of the Baryon Distribution with Planck, WMAP, and WISE Data". In: *Phys. Rev. Lett.* 117.5, 051301 (July 2016), p. 051301. DOI: 10.1103/PhysRevLett.117.051301. arXiv: 1603.01608 [astro-ph.CO].

[288] J. Colin Hill, Evan McDonough, Michael W. Toomey, and Stephon Alexander. "Early Dark Energy Does Not Restore Cosmological Concordance". In: *arXiv e-prints*, arXiv:2003.07355 (Mar. 2020), arXiv:2003.07355. arXiv: 2003.07355 [astro-ph.CO].

[289] Christopher M. Hirata and Uros Seljak. "Reconstruction of lensing from the cosmic microwave background polarization". In: *Phys. Rev. D* 68 (2003), p. 083002. DOI: 10.1103/PhysRevD.68.083002. arXiv: astro-ph/0306354.

[290] E. Hivon, K. M. Gorski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen. "Master of the cosmic microwave background anisotropy power spectrum: a fast method for statistical analysis of large and complex cosmic microwave background data sets". In: *Astrophys. J.* 567 (2002), p. 2. DOI: 10.1086/338126. arXiv: astro-ph/0105302.

[291] Kai Hoffmann, Yi Mao, Jiachuan Xu, Houjun Mo, and Benjamin D. Wandelt. "Signatures of cosmic reionization on the 21-cm two- and three-point correlation function I: quadratic bias modelling". In: *Mon. Not. R. Astron. Soc.* 487.3 (Aug. 2019), pp. 3050–3068. DOI: 10.1093/mnras/stz1472. arXiv: 1802.02578 [astro-ph.CO].

[292] Benjamin Horowitz, Simone Ferraro, and Blake D. Sherwin. "Reconstructing Small Scale Lenses from the Cosmic Microwave Background Temperature Fluctuations". In: *Mon. Not. Roy. Astron. Soc.* 485.3 (2019), pp. 3919–3929. DOI: 10.1093/mnras/stz566. arXiv: 1710.10236 [astro-ph.CO].

[293] Benjamin Horowitz and Uroš Seljak. "Cosmological constraints from thermal Sunyaev–Zeldovich power spectrum revisited". In: *Mon. Not. Roy. Astron. Soc.* 469.1 (2017), pp. 394–400. DOI: 10.1093/mnras/stx766. arXiv: 1609.01850 [astro-ph.CO].

[294] Jiamin Hou et al. "The Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: BAO and RSD measurements from anisotropic clustering analysis of the Quasar Sample in configuration space between redshift 0.8 and 2.2". In: *Mon. Not. Roy. Astron. Soc.* 500.1 (2020), pp. 1201–1221. DOI: 10.1093/mnras/staa3234. arXiv: 2007.08998 [astro-ph.CO].

[295] Zhen Hou, Ryan Keisler, Lloyd Knox, Marius Millea, and Christian Reichardt. "How Massless Neutrinos Affect the Cosmic Microwave Background Damping Tail". In: *Phys. Rev. D* 87 (2013), p. 083008. DOI: 10.1103/PhysRevD.87.083008. arXiv: 1104.2333 [astro-ph.CO].

[296] Cullan Howlett, Antony Lewis, Alex Hall, and Anthony Challinor. "CMB power spectrum parameter degeneracies in the era of precision cosmology". In: *Journal of Cosmology and Astroparticle Physics* 2012.04 (Apr. 2012), p. 027. DOI: 10.1088/1475-7516/2012/04/027. URL: https://dx.doi.org/10.1088/1475-7516/2012/04/027.

[297] Cullan Howlett, Khaled Said, John R. Lucey, Matthew Colless, Fei Qin, Yan Lai, R. Brent Tully, and Tamara M. Davis. "The sloan digital sky survey peculiar velocity catalogue". In: *Mon. Not. Roy. Astron. Soc.* 515.1 (2022), pp. 953–976. DOI: 10.1093/mnras/stac1681. arXiv: 2201.03112 [astro-ph.CO].

[298] Wayne Hu, Simon DeDeo, and Chris Vale. "Cluster mass estimators from CMB temperature and polarization lensing". In: *New Journal of Physics* 9.12 (Dec. 2007), p. 441. DOI: 10.1088/1367-2630/9/12/441. arXiv: astro-ph/0701276 [astro-ph].

[299] Wayne Hu, Daniel J. Eisenstein, and Max Tegmark. "Weighing neutrinos with galaxy surveys". In: *Phys. Rev. Lett.* 80 (1998), pp. 5255–5258. DOI: 10.1103/PhysRevLett.80.5255. arXiv: astro-ph/9712057.

[300] Wayne Hu and Takemi Okamoto. "Mass Reconstruction with Cosmic Microwave Background Polarization". In: *Astrophys. J.* 574.2 (Aug. 2002), pp. 566–574. DOI: 10.1086/341110. arXiv: astro-ph/0111606 [astro-ph].

[301] Wayne Hu, Douglas Scott, and Joseph Silk. "Reionization and cosmic microwave background distortions: A complete treatment of second-order Compton scattering". In: *Phys. Rev. D* 49.2 (Jan. 1994), pp. 648–670. DOI: 10.1103/PhysRevD.49.648. arXiv: astro-ph/9305038 [astro-ph].

[302] Edwin Hubble. "A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae". In: *Proceedings of the National Academy of Sciences* 15.3 (1929), pp. 168–173. DOI: 10.1073/pnas.15.3.168.

[303] John D. Hunter. "Matplotlib: A 2D Graphics Environment". In: *Computing in Science & Engineering* 9.3 (2007), pp. 90–95. DOI: 10.1109/MCSE.2007.55.

[304] Dragan Huterer, Carlos E. Cunha, and Wenjuan Fang. "Calibration errors unleashed: effects on cosmological parameters and requirements for large-scale structure surveys". In: *Mon. Not. R. Astron. Soc.* 432.4 (July 2013), pp. 2945–2961. DOI: 10.1093/mnras/stt653. arXiv: 1211.1015 [astro-ph.CO].

[305] S. Ilić et al. "*Euclid* preparation: XV. Forecasting cosmological constraints for the *Euclid* and CMB joint analysis". In: (June 2021). arXiv: 2106.08346 [astro-ph.CO].

[306] S. Ilić et al. "Reconstructing the epoch of reionisation with Planck PR4". In: *arXiv e-prints*, arXiv:2504.13254 (Apr. 2025), arXiv:2504.13254. arXiv: 2504.13254 [astro-ph.CO].

[307] Mikhail M. Ivanov, Evan McDonough, J. Colin Hill, Marko Simonović, Michael W. Toomey, Stephon Alexand er, and Matias Zaldarriaga. "Constraining Early Dark Energy with Large-Scale Structure". In: *arXiv e-prints*, arXiv:2006.11235 (June 2020), arXiv:2006.11235. arXiv: 2006.11235 [astro-ph.CO].

[308] Mikhail M. Ivanov, Oliver H. E. Philcox, Marko Simonović, Matias Zaldarriaga, Takahiro Nischimichi, and Masahiro Takada. "Cosmological constraints without nonlinear redshift-space distortions". In: *Phys. Rev. D* 105.4 (2022), p. 043531. DOI: 10.1103/PhysRevD.105.043531. arXiv: 2110.00006 [astro-ph.CO].

[309] Mikhail M. Ivanov, Marko Simonović, and Matias Zaldarriaga. "Cosmological Parameters from the BOSS Galaxy Power Spectrum". In: *JCAP* 05 (2020), p. 042. DOI: 10.1088/1475-7516/2020/05/042. arXiv: 1909.05277 [astro-ph.CO].

[310] Željko Ivezić et al. "LSST: From Science Drivers to Reference Design and Anticipated Data Products". In: *Astrophys. J.* 873.2, 111 (Mar. 2019), p. 111. DOI: 10.3847/1538-4357/ab042c. arXiv: 0805.2366 [astro-ph].

[311] Bhuvnesh Jain and Justin Khoury. "Cosmological tests of gravity". In: *Annals of Physics* 325.7 (July 2010), pp. 1479–1516. DOI: 10.1016/j.aop.2010.04.002. arXiv: 1004.3294 [astro-ph.CO].

[312] Divesh Jain, Tirthankar Roy Choudhury, Srinivasan Raghunathan, and Suvodip Mukherjee. "Probing the physics of reionization using kinematic Sunyaev–Zeldovich power spectrum from current and upcoming cosmic microwave background surveys". In: *Mon. Not. Roy. Astron. Soc.* 530.1 (2024), pp. 35–51. DOI: 10.1093/mnras/stae748. arXiv: 2311.00315 [astro-ph.CO].

[313] Harold Jeffreys. "An Invariant Form for the Prior Probability in Estimation Problems". In: *Proceedings of the Royal Society of London Series A* 186.1007 (Sept. 1946), pp. 453–461. DOI: 10.1098/rspa.1946.0056.

[314] A. Joyce, L. Lombriser, and F. Schmidt. "Dark Energy Versus Modified Gravity". In: *Annual Review of Nuclear and Particle Science* 66 (Oct. 2016), pp. 95–122. DOI: 10.1146/annurev-nucl-102115-044553. arXiv: 1601.06133.

[315] Austin Joyce, Bhuvnesh Jain, Justin Khoury, and Mark Trodden. "Beyond the cosmological standard model". In: *Phys. Rep.* 568 (Mar. 2015), pp. 1–98. DOI: 10.1016/j.physrep.2014.12.002. arXiv: 1407.0059 [astro-ph.CO].

[316] N. Kaiser. "Clustering in real space and in redshift space". In: *Mon. Not. R. Astron. Soc.* 227 (July 1987), pp. 1–21. DOI: 10.1093/mnras/227.1.1.

[317] Nick Kaiser. "Weak Gravitational Lensing of Distant Galaxies". In: *Astrophys. J.* 388 (Apr. 1992), p. 272. DOI: 10.1086/171151.

[318] R. Kannan, E. Garaldi, A. Smith, R. Pakmor, V. Springel, M. Vogelsberger, and L. Hernquist. "Introducing the THESAN project: radiation-magnetohydrodynamic simulations of the epoch of reionization". In: *Mon. Not. R. Astron. Soc.* 511.3 (Apr. 2022), pp. 4005–4030. DOI: 10.1093/mnras/stab3710. arXiv: 2110.00584 [astro-ph.GA].

[319] Dionysios Karagiannis, Andrei Lazanu, Michele Liguori, Alvise Raccanelli, Nicola Bartolo, and Licia Verde. "Constraining primordial non-Gaussianity with bispectrum and power spectrum from upcoming optical and radio surveys". In: *Mon. Not. Roy. Astron. Soc.* 478.1 (2018), pp. 1341–1376. DOI: 10.1093/mnras/sty1029. arXiv: 1801.09280 [astro-ph.CO].

[320] Joshua Kim et al. "The Atacama Cosmology Telescope DR6 and DESI: structure formation over cosmic time with a measurement of the cross-correlation of CMB lensing and luminous red galaxies". In: *J. Cosm. Astrop. Phys.* 2024.12, 022 (Dec. 2024), p. 022. DOI: 10.1088/1475-7516/2024/12/022. arXiv: 2407.04606 [astro-ph.CO].

[321] Soma King and Philip Lubin. "Circular polarization of the CMB: Foregrounds and detection prospects". In: *Phys. Rev. D* 94.2 (2016), p. 023501. DOI: 10.1103/PhysRevD.94.023501. arXiv: 1606.04112 [astro-ph.CO].

[322] Ellie Kitanidis et al. "Imaging Systematics and Clustering of DESI Main Targets". In: *Mon. Not. Roy. Astron. Soc.* 496.2 (2020), pp. 2262–2291. DOI: 10.1093/mnras/staa1621. arXiv: 1911.05714 [astro-ph.CO].

[323] Ellie Kitanidis and Martin White. "Cross-Correlation of Planck CMB Lensing with DESI-Like LRGs". In: *Mon. Not. Roy. Astron. Soc.* 501.4 (2021), pp. 6181–6198. DOI: 10.1093/mnras/staa3927. arXiv: 2010.04698 [astro-ph.CO].

[324] Anatoly Klypin et al. "Clustering and Halo Abundances in Early Dark Energy Cosmological Models". In: *arXiv e-prints*, arXiv:2006.14910 (June 2020), arXiv:2006.14910. arXiv: 2006.14910 [astro-ph.CO].

[325] James P. Kneller and Gary Steigman. "BBN for pedestrians". In: *New J. Phys.* 6 (2004), p. 117. DOI: 10.1088/1367-2630/6/1/117. arXiv: astro-ph/0406320.

[326] Yosuke Kobayashi, Takahiro Nishimichi, Masahiro Takada, and Hironao Miyatake. "Full-shape cosmology analysis of the SDSS-III BOSS galaxy power spectrum using an emulator-based halo model: A 5% determination of $\sigma 8$". In: *Phys. Rev. D* 105.8 (2022), p. 083517. DOI: 10.1103/PhysRevD.105.083517. arXiv: 2110.06969 [astro-ph.CO].

[327] Nickolas Kokron, Joseph DeRose, Shi-Fan Chen, Martin White, and Risa H. Wechsler. "The cosmology dependence of galaxy clustering and lensing from a hybrid N-body-perturbation theory model". In: *Mon. Not. R. Astron. Soc.* 505.1 (July 2021), pp. 1422–1440. DOI: 10.1093/mnras/stab1358. arXiv: 2101.11014 [astro-ph.CO].

[328] Krause, E. and Hirata, C. M. "Weak lensing power spectra for precision cosmology - Multiple-deflection, reduced shear, and lensing bias corrections". In: *A&A* 523 (2010), A28. DOI: 10.1051/0004-6361/200913524. URL: https://doi.org/10.1051/0004-6361/200913524.

[329] Alex Krolewski, Simone Ferraro, Edward F. Schlafly, and Martin White. "unWISE tomography of Planck CMB lensing". In: *JCAP* 05 (2020), p. 047. DOI: 10.1088/1475-7516/2020/05/047. arXiv: 1909.07412 [astro-ph.CO].

[330] Alex Krolewski, Simone Ferraro, and Martin White. "Cosmological constraints from unWISE and Planck CMB lensing tomography". In: *JCAP* 12.12 (2021), p. 028. DOI: 10.1088/1475-7516/2021/12/028. arXiv: 2105.03421 [astro-ph.CO].

[331] Suresh Kumar and Rafael C. Nunes. "Observational constraints on dark matter–dark energy scattering cross section". In: *Eur. Phys. J. C* 77.11 (2017), p. 734. DOI: 10.1140/epjc/s10052-017-5334-3. arXiv: 1709.02384 [astro-ph.CO].

[332] Aleksandra Kusiak, Boris Bolliet, Simone Ferraro, J. Colin Hill, and Alex Krolewski. "Constraining the baryon abundance with the kinematic Sunyaev-Zel'dovich effect: Projected-field detection using P l a n c k , W M A P , and u n W I S E". In: *Phys. Rev. D* 104.4, 043518 (Aug. 2021), p. 043518. DOI: 10.1103/PhysRevD.104.043518. arXiv: 2102.01068 [astro-ph.CO].

[333] Alex Laguë, Fiona McCarthy, Mathew Madhavacheril, J. Colin Hill, and Frank J. Qu. "Constraints on dark matter-dark energy scattering from ACT DR6 CMB lensing measurements". In: *Phys. Rev. D* 110.2, 023536 (July 2024), p. 023536. DOI: 10.1103/PhysRevD.110.023536. arXiv: 2402.08149 [astro-ph.CO].

[334] R. Laureijs and et al. "Euclid Definition Study Report". In: *arXiv e-prints* (2011). eprint: 1110.3193.

[335] Boris Leistedt, Hiranya V. Peiris, and Nina Roth. "Constraints on Primordial Non-Gaussianity from 800 000 Photometric Quasars". In: *Phys. Rev. Lett.* 113.22 (2014), p. 221301. DOI: 10.1103/PhysRevLett.113.221301. arXiv: 1405.4315 [astro-ph.CO].

[336] Margherita Lembo, Giulio Fabbian, Julien Carron, and Antony Lewis. "CMB lensing reconstruction biases from masking extragalactic sources". In: *arXiv e-prints*, arXiv:2109.13911 (Sept. 2021), arXiv:2109.13911. arXiv: 2109.13911 [astro-ph.CO].

[337] Julien Lesgourgues, Gianpiero Mangano, Gennaro Miele, and Sergio Pastor. *Neutrino Cosmology*. Cambridge University Press, Feb. 2013. ISBN: 978-1-108-70501-1, 978-1-139-60341-6.

[338] Julien Lesgourgues and Sergio Pastor. "Massive neutrinos and cosmology". In: *Phys. Rept.* 429 (2006), pp. 307–379. DOI: 10.1016/j.physrep.2006.04.001. arXiv: astro-ph/0603494.

[339] Julien Lesgourgues and Sergio Pastor. "Neutrino mass from Cosmology". In: *Adv. High Energy Phys.* 2012 (2012), p. 608515. DOI: 10.1155/2012/608515. arXiv: 1212.6154 [hep-ph].

[340] Julien Lesgourgues and Thomas Tram. "Cosmological neutrino mass detection: The best probe of neutrino lifetime". In: *Journal of Cosmology and Astroparticle Physics* 2013.09 (2013), p. 032. DOI: 10.1088/1475-7516/2013/09/032. arXiv: 1307.0637.

[341] Michael Levi et al. "The DESI Experiment, a whitepaper for Snowmass 2013". In: *arXiv e-prints*, arXiv:1308.0847 (Aug. 2013), arXiv:1308.0847. DOI: 10.48550/arXiv.1308.0847. arXiv: 1308.0847 [astro-ph.CO].

[342] A LEWIS and A CHALLINOR. "Weak gravitational lensing of the CMB". In: *Physics Reports* 429.1 (June 2006), pp. 1–65. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2006.03.002. URL: http://dx.doi.org/10.1016/j.physrep.2006.03.002.

[343] Antony Lewis. "Efficient sampling of fast and slow cosmological parameters". In: *Phys. Rev. D* 87.10 (2013), p. 103529. DOI: 10.1103/PhysRevD.87.103529. arXiv: 1304.4473 [astro-ph.CO].

[344] Antony Lewis. "GetDist: a Python package for analysing Monte Carlo samples". In: *arXiv e-prints*, arXiv:1910.13970 (Oct. 2019), arXiv:1910.13970. DOI: 10.48550/arXiv.1910.13970. arXiv: 1910.13970 [astro-ph.IM].

[345] Antony Lewis and Sarah Bridle. "Cosmological parameters from CMB and other data: A Monte Carlo approach". In: *Phys. Rev. D* 66 (2002), p. 103511. DOI: 10.1103/PhysRevD.66.103511. arXiv: astro-ph/0205436.

[346] Antony Lewis, Anthony Challinor, and Anthony Lasenby. "Efficient Computation of Cosmic Microwave Background Anisotropies in Closed Friedmann-Robertson-Walker Models". In: *The Astrophysical Journal* 538.2 (Aug. 2000), p. 473. DOI: 10.1086/309179. URL: https://dx.doi.org/10.1086/309179.

[347] Ming Li, Raul E. Angulo, Simon D. M. White, and Jens Jasche. "Matched filter optimization of kSZ measurements with a reconstructed cosmological flow field". In: *Mon. Not. Roy. Astron. Soc.* 443.3 (2014), pp. 2311–2326. DOI: 10.1093/mnras/stu1224. arXiv: 1404.0007 [astro-ph.CO].

[348] Shaohong Li, Yi Zheng, Ziyang Chen, Haojie Xu, and Xiaohu Yang. "Detection of Pairwise Kinetic Sunyaev–Zel'dovich Effect with DESI Galaxy Groups and Planck in Fourier Space". In: *Astrophys. J. Supp.* 271.1, 30 (Mar. 2024), p. 30. DOI: 10.3847/1538-4365/ad1bd8. arXiv: 2401.03507 [astro-ph.CO].

[349] Xiangchong Li et al. "Hyper Suprime-Cam Year 3 results: Cosmology from cosmic shear two-point correlation functions". In: *Phys. Rev. D* 108.12 (2023), p. 123518. DOI: 10.1103/PhysRevD.108.123518. arXiv: 2304.00702 [astro-ph.CO].

[350] Yunyang Li et al. "A Measurement of the Largest-Scale CMB E-mode Polarization with CLASS". In: *arxiv e-prints* (Jan. 2025). arXiv: 2501.11904 [astro-ph.CO].

[351] Zack Li, Giuseppe Puglisi, Mathew S. Madhavacheril, and Marcelo A. Alvarez. "Simulated catalogs and maps of radio galaxies at millimeter wavelengths in Websky". In: *JCAP* 08.08 (2022), p. 029. DOI: 10.1088/1475-7516/2022/08/029. arXiv: 2110.15357 [astro-ph.GA].

[352] D. Nelson Limber. "The Analysis of Counts of the Extragalactic Nebulae in Terms of a Fluctuating Density Field." In: *Astrophys. J.* 117 (Jan. 1953), p. 134. DOI: 10.1086/145672.

[353] Eric V. Linder. "Exploring the expansion history of the universe". In: *Phys. Rev. Lett.* 90 (2003), p. 091301. DOI: 10.1103/PhysRevLett.90.091301. arXiv: astro-ph/0208512.

[354] Eric V. Linder. "The Rise of Dark Energy". In: (June 2021). arXiv: 2106.09581 [astro-ph.CO].

[355] Adrian Liu, Jonathan R. Pritchard, Rupert Allison, Aaron R. Parsons, Uroš Seljak, and Blake D. Sherwin. "Eliminating the optical depth nuisance from the CMB with 21 cm cosmology". In: *Phys. Rev. D* 93.4 (2016), p. 043013. DOI: 10.1103/PhysRevD.93.043013. arXiv: 1509.08463 [astro-ph.CO].

[356] Adrian Liu and J. Richard Shaw. "Data Analysis for Precision 21 cm Cosmology". In: *Publ. Astron. Soc. Pacific* 132.1012, 062001 (June 2020), p. 062001. DOI: 10.1088/1538-3873/ab5bfd. arXiv: 1907.08211 [astro-ph.IM].

[357] Hongwan Liu, Tracy R. Slatyer, and Jesús Zavala. "Contributions to cosmic reionization from dark matter annihilation and decay". In: *Phys. Rev. D* 94.6 (2016), p. 063507. DOI: 10.1103/PhysRevD.94.063507. arXiv: 1604.02457 [astro-ph.CO].

[358] Jia Liu and J. Colin Hill. "Cross-correlation of Planck CMB lensing and CFHTLenS galaxy weak lensing maps". In: *Physical Review D* 92.6 (Sept. 2015). ISSN: 1550-2368. DOI: 10.1103/physrevd.92.063517. URL: http://dx.doi.org/10.1103/PhysRevD.92.063517.

[359] R. Henry Liu et al. "Comparison of Gaussian Process Transfer Function and Hybrid Effective Field Theory Methods for Baryon Pasting on Simulations". In: *in preparation* (2024).

[360] M. Lokken et al. "Superclustering with the Atacama Cosmology Telescope and Dark Energy Survey: II. Anisotropic large-scale coherence in hot gas, galaxies, and dark matter". In: (Sept. 2024). arXiv: 2409.04535 [astro-ph.CO].

[361] M. Lokken et al. "Superclustering with the Atacama Cosmology Telescope and Dark Energy Survey. I. Evidence for Thermal Energy Anisotropy Using Oriented Stacking". In: *Astrophys. J.* 933.2 (2022), p. 134. DOI: 10.3847/1538-4357/ac7043. arXiv: 2107.05523 [astro-ph.CO].

[362] Thibaut Louis et al. "The Atacama Cosmology Telescope: DR6 Power Spectra, Likelihoods and ΛCDM Parameters". In: (Mar. 2025). arXiv: 2503.14452 [astro-ph.CO].

[363] A. Loureiro et al. "KiDS and Euclid: Cosmological implications of a pseudo angular power spectrum analysis of KiDS-1000 cosmic shear tomography". In: *Astron. Astrophys.* 665 (2022), A56. DOI: 10.1051/0004-6361/202142481. arXiv: 2110.06947 [astro-ph.CO].

[364] Marilena LoVerde. "Halo bias in mixed dark matter cosmologies". In: *Phys. Rev. D* 90.8 (2014), p. 083530. DOI: 10.1103/PhysRevD.90.083530. arXiv: 1405.4855 [astro-ph.CO].

[365] Marilena LoVerde and Niayesh Afshordi. "Extended Limber Approximation". In: *Phys. Rev. D* 78 (2008), p. 123506. DOI: 10.1103/PhysRevD.78.123506. arXiv: 0809.5112 [astro-ph].

[366] Marilena LoVerde and Niayesh Afshordi. "Extended Limber approximation". In: *Physical Review D* 78.12 (Dec. 2008). ISSN: 1550-2368. DOI: 10.1103/physrevd.78.123506. URL: http://dx.doi.org/10.1103/PhysRevD.78.123506.

[367] Marilena Loverde and Zachary J. Weiner. "Massive neutrinos and cosmic composition". In: *JCAP* 12 (2024), p. 048. DOI: 10.1088/1475-7516/2024/12/048. arXiv: 2410.00090 [astro-ph.CO].

[368] LSST Science Collaboration et al. "LSST Science Book, Version 2.0". In: *arXiv e-prints*, arXiv:0912.0201 (Dec. 2009), arXiv:0912.0201. arXiv: 0912.0201 [astro-ph.IM].

[369] Matteo Lucca. "Dark energy–dark matter interactions as a solution to the S8 tension". In: *Phys. Dark Univ.* 34 (2021), p. 100899. DOI: 10.1016/j.dark.2021.100899. arXiv: 2105.09249 [astro-ph.CO].

[370] Gabriel P. Lynch and Lloyd Knox. "What's the matter with Σm_ν?" In: *arXiv e-prints* (Mar. 2025). arXiv: 2503.14470 [astro-ph.CO].

[371] Niall MacCrann et al. "The Atacama Cosmology Telescope: reionization kSZ trispectrum methodology and limits". In: *Mon. Not. Roy. Astron. Soc.* 532.4 (2024), pp. 4247–4260. DOI: 10.1093/mnras/stae1746. arXiv: 2405.01188 [astro-ph.CO].

[372] Niall MacCrann et al. "The Atacama Cosmology Telescope: Mitigating the Impact of Extragalactic Foregrounds for the DR6 Cosmic Microwave Background Lensing Analysis". In: *Astrophys. J.* 966.1, 138 (May 2024), p. 138. DOI: 10.3847/1538-4357/ad2610. arXiv: 2304.05196 [astro-ph.CO].

[373] Piero Madau, Avery Meiksin, and Martin J. Rees. "21 Centimeter Tomography of the Intergalactic Medium at High Redshift". In: *Astrophys. J.* 475.2 (Feb. 1997), pp. 429–444. DOI: 10.1086/303549. arXiv: astro-ph/9608010 [astro-ph].

[374] Mathew S. Madhavacheril and J. Colin Hill. "Mitigating foreground biases in CMB lensing reconstruction using cleaned gradi ents". In: *Phys. Rev. D* 98.2, 023534 (July 2018), p. 023534. DOI: 10.1103/PhysRevD.98.023534. arXiv: 1802.08230 [astro-ph.CO].

[375] Mathew S. Madhavacheril, Kendrick M. Smith, Blake D. Sherwin, and Sigurd Naess. "CMB lensing power spectrum estimation without instrument noise bias". In: (Nov. 2020). DOI: 10.1088/1475-7516/2021/05/028. arXiv: 2011.02475 [astro-ph.CO].

[376] Mathew S. Madhavacheril et al. "Atacama Cosmology Telescope: Component-separated maps of CMB temperature and the thermal Sunyaev-Zel'dovich effect". In: *Physical Review D* 102.2 (July 2020). ISSN: 2470-0029. DOI: 10.1103/physrevd.102.023534. URL: http://dx.doi.org/10.1103/PhysRevD.102.023534.

[377] Mathew S. Madhavacheril et al. "The Atacama Cosmology Telescope: DR6 Gravitational Lensing Map and Cosmological Parameters". In: *Astrophys. J.* 962.2, 113 (Feb. 2024), p. 113. DOI: 10.3847/1538-4357/acff5f. arXiv: 2304.05203 [astro-ph.CO].

[378] Daisy Suet Ying Mak, Anthony Challinor, George Efstathiou, and Guilaine Lagache. "Measurement of CIB power spectra over large sky areas from Planck HFI maps". In: *Mon. Not. R. Astron. Soc.* 466.1 (Apr. 2017), pp. 286–319. DOI: 10.1093/mnras/stw3112. arXiv: 1609.08942 [astro-ph.CO].

[379] Juan Maldacena. "Non-gaussian features of primordial fluctuations in single field inflationary models". In: *Journal of High Energy Physics* 2003.05 (May 2003), pp. 013–013. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2003/05/013. URL: http://dx.doi.org/10.1088/1126-6708/2003/05/013.

[380] M. Mallaby-Kay et al. "Kinematic Sunyaev-Zel'dovich effect with ACT, DES, and BOSS: A novel hybrid estimator". In: *Phys. Rev. D* 108.2, 023516 (July 2023), p. 023516. DOI: 10.1103/PhysRevD.108.023516. arXiv: 2305.06792 [astro-ph.CO].

[381] Gianpiero Mangano, Gennaro Miele, Sergio Pastor, Teguayco Pinto, Ofelia Pisanti, and Pasquale D. Serpico. "Relic neutrino decoupling including flavour oscillations". In: *Nuclear Physics B* 729.1-2 (Nov. 2005), pp. 221–234. ISSN: 0550-3213. DOI: 10.1016/j.nuclphysb.2005.09.041. URL: http://dx.doi.org/10.1016/j.nuclphysb.2005.09.041.

[382] Abhishek S. Maniyar, Yacine Ali-Haımoud, Julien Carron, Antony Lewis, and Mathew S. Madhavacheril. "Quadratic estimators for CMB weak lensing". In: *Phys. Rev. D* 103.8 (2021), p. 083524. DOI: 10.1103/PhysRevD.103.083524. arXiv: 2101.12193 [astro-ph.CO].

[383] Yi Mao, Paul R. Shapiro, Garrelt Mellema, Ilian T. Iliev, Jun Koda, and Kyungjin Ahn. "Redshift-space distortion of the 21-cm background from the epoch of reionization - I. Methodology re-examined". In: *Mon. Not. R. Astron. Soc.* 422.2 (May 2012), pp. 926–954. DOI: 10.1111/j.1365-2966.2012.20471.x. arXiv: 1104.2094 [astro-ph.CO].

[384] Giovanni Marozzi, Giuseppe Fanizza, Enea Di Dio, and Ruth Durrer. "CMB-lensing beyond the Born approximation". In: *JCAP* 09 (2016), p. 028. DOI: 10.1088/1475-7516/2016/09/028. arXiv: 1605.08761 [astro-ph.CO].

[385] G. A. Marques et al. "Cosmological constraints from the tomography of DES-Y3 galaxies with CMB lensing from ACT DR4". In: *JCAP* 01 (2024), p. 033. DOI: 10.1088/1475-7516/2024/01/033. arXiv: 2306.17268 [astro-ph.CO].

[386] Charlotte A. Mason et al. "Inferences on the timeline of reionization at z $\sim$ 8 from the KMOS Lens-Amplified Spectroscopic Survey". In: *Mon. Not. Roy. Astron. Soc.* 485.3 (2019), pp. 3947–3969. DOI: 10.1093/mnras/stz632. arXiv: 1901.11045 [astro-ph.CO].

[387] T. Matsubara. "Nonlinear perturbation theory with halo bias and redshift-space distortions via the Lagrangian picture". In: *Phys. Rev. D* 78.8, 083519 (Oct. 2008), p. 083519. DOI: 10.1103/PhysRevD.78.083519. arXiv: 0807.1733.

[388] T. Matsubara. "Resumming cosmological perturbations via the Lagrangian picture: One-loop results in real space and in redshift space". In: *Phys. Rev. D* 77.6, 063530 (Mar. 2008), p. 063530. DOI: 10.1103/PhysRevD.77.063530. arXiv: 0711.2521.

[389] M. Maus et al. "An analysis of parameter compression and full-modeling techniques with Velocileptors for DESI 2024 and beyond". In: *arXiv e-prints*, arXiv:2404.07312 (Apr. 2024), arXiv:2404.07312. DOI: 10.48550/arXiv.2404.07312. arXiv: 2404.07312 [astro-ph.CO].

[390] Mark Maus et al. "A joint analysis of RSD and galaxy $\times$ CMB lensing cross-correlations with DESI DR1 LRG galaxies". In: (2025). arXiv: 25MM.XXXXX [astro-ph.CO].

[391] Fiona McCarthy, Simon Foreman, and Alexander van Engelen. "Avoiding baryonic feedback effects on neutrino mass measurements from CMB lensing". In: (Nov. 2020). arXiv: 2011.06582 [astro-ph.CO].

[392] Fiona McCarthy, J. Colin Hill, and Mathew S. Madhavacheril. "Baryonic feedback biases on fundamental physics from lensed CMB power spectra". In: (Mar. 2021). arXiv: 2103.05582 [astro-ph.CO].

[393] Ian G. McCarthy et al. "FLAMINGO: combining kinetic SZ effect and galaxy–galaxy lensing measurements to gauge the impact of feedback on large-scale structure". In: *Mon. Not. R. Astron. Soc.* 540.1 (June 2025), pp. 143–163. DOI: 10.1093/mnras/staf731. arXiv: 2410.19905 [astro-ph.CO].

[394] N. McCullagh and A. S. Szalay. "Nonlinear Behavior of Baryon Acoustic Oscillations from the Zel'dovich Approximation Using a Non-Fourier Perturbation Approach". In: *Astrophys. J.* 752, 21 (June 2012), p. 21. DOI: 10.1088/0004-637X/752/1/21. arXiv: 1202.1306 [astro-ph.CO].

[395] Patrick McDonald and Arabindo Roy. "Clustering of dark matter tracers: generalizing bias for the coming era of precision LSS". In: *J. Cosm. Astrop. Phys.* 2009.8, 020 (Aug. 2009), p. 020. DOI: 10.1088/1475-7516/2009/08/020. arXiv: 0902.0991 [astro-ph.CO].

[396] Ian McGreer, Andrei Mesinger, and Valentina D'Odorico. "Model-independent evidence in favour of an end to reionization by $z \approx 6$". In: *Mon. Not. Roy. Astron. Soc.* 447.1 (2015), pp. 499–505. DOI: 10.1093/mnras/stu2449. arXiv: 1411.5375 [astro-ph.CO].

[397] B. R. McNamara and P. E. J. Nulsen. "Heating Hot Atmospheres with Active Galactic Nuclei". In: *Annual Review of Astronomy and Astrophysics* 45.1 (Sept. 2007), pp. 117–175. DOI: 10.1146/annurev.astro.45.051806.110625. arXiv: 0709.2152 [astro-ph].

[398] M. McQuinn and M. White. "Cosmological perturbation theory in 1+1 dimensions". In: *J. Cosm. Astrop. Phys.* 1, 043 (Jan. 2016), p. 043. DOI: 10.1088/1475-7516/2016/01/043. arXiv: 1502.07389.

[399] Matthew McQuinn and Anson D'Aloisio. "The observable 21cm signal from reionization may be perturbative". In: *JCAP* 10 (2018), p. 016. DOI: 10.1088/1475-7516/2018/10/016. arXiv: 1806.08372 [astro-ph.CO].

[400] Matthew McQuinn, Oliver Zahn, Matias Zaldarriaga, Lars Hernquist, and Steven R. Furlanetto. "Cosmological Parameter Estimation Using 21 cm Radiation from the Epoch of Reionization". In: *Astrophys. J.* 653.2 (Dec. 2006), pp. 815–834. DOI: 10.1086/505167. arXiv: astro-ph/0512263 [astro-ph].

[401] Alexander Mead and Licia Verde. "Including beyond-linear halo bias in halo models". In: *arXiv e-prints*, arXiv:2011.08858 (Nov. 2020), arXiv:2011.08858. arXiv: 2011.08858 [astro-ph.CO].

[402] P. Daniel Meerburg, Joel Meyers, Kendrick M. Smith, and Alexander van Engelen. "Reconstructing CMB fluctuations and the mean reionization optical depth". In: *Phys. Rev. D* 95.12 (2017), p. 123538. DOI: 10.1103/PhysRevD.95.123538. arXiv: 1701.06992 [astro-ph.CO].

[403] Pieter Daniel Meerburg et al. "Primordial Non-Gaussianity". In: *Bull. Am. Astron. Soc.* 51.3, 107 (May 2019), p. 107. arXiv: 1903.04409 [astro-ph.CO].

[404] A. M. Meisner, D. Lang, E. F. Schlafly, and D. J. Schlegel. "unWISE Coadds: The Five-year Data Set". In: *Publ. Astron. Soc. Pacific* 131.1006 (Dec. 2019), p. 124504. DOI: 10.1088/1538-3873/ab3df4. arXiv: 1909.05444 [astro-ph.IM].

[405] Andrei Mesinger. "Understanding the Epoch of Cosmic Reionization". In: *Astrophysics and Space Science Library* 423 (Jan. 2016). DOI: 10.1007/978-3-319-21957-8.

[406] Michaël Michaux, Oliver Hahn, Cornelius Rampf, and Raul E. Angulo. "Accurate initial conditions for cosmological N-body simulations: Minimizing truncation and discreteness errors". In: *Mon. Not. Roy. Astron. Soc.* 500.1 (2020), pp. 663–683. DOI: 10.1093/mnras/staa3149. arXiv: 2008.09588 [astro-ph.CO].

[407] Marius Millea and François Bouchet. "Cosmic Microwave Background Constraints in Light of Priors Over Reionization Histories". In: *Astron. Astrophys.* 617 (2018), A96. DOI: 10.1051/0004-6361/201833288. arXiv: 1804.08476 [astro-ph.CO].

[408] Marius Millea and Uros Seljak. "Marginal unbiased score expansion and application to CMB lensing". In: *Phys. Rev. D* 105.10 (2022), p. 103531. DOI: 10.1103/PhysRevD.105.103531. arXiv: 2112.09354 [astro-ph.CO].

[409] Timothy N. Miller et al. "The Optical Corrector for the Dark Energy Spectroscopic Instrument". In: *arXiv e-prints*, arXiv:2306.06310 (June 2023), arXiv:2306.06310. DOI: 10.48550/arXiv.2306.06310. arXiv: 2306.06310 [astro-ph.IM].

[410] Nishant Mishra and Emmanuel Schaan. "Bias to CMB lensing from lensed foregrounds". In: *Phys. Rev. D* 100.12, 123504 (Dec. 2019), p. 123504. DOI: 10.1103/PhysRevD.100.123504. arXiv: 1908.08057 [astro-ph.CO].

[411] Hironao Miyatake et al. "Hyper Suprime-Cam Year 3 results: Cosmology from galaxy clustering and weak lensing with HSC and SDSS using the emulator based halo model". In: *Phys. Rev. D* 108.12 (2023), p. 123517. DOI: 10.1103/PhysRevD.108.123517. arXiv: 2304.00704 [astro-ph.CO].

[412] C. Modi, E. Castorina, and U. Seljak. "Halo bias in Lagrangian Space: Estimators and theoretical predictions". In: *ArXiv e-prints* (Dec. 2016). arXiv: 1612.01621.

[413] Chirag Modi, Shi-Fan Chen, and Martin White. "Simulations and symmetries". In: *Mon. Not. Roy. Astron. Soc.* 492.4 (2020), pp. 5754–5763. DOI: 10.1093/mnras/staa251. arXiv: 1910.07097 [astro-ph.CO].

[414] Chirag Modi, Martin White, Emanuele Castorina, and Anže Slosar. "Mind the gap: the power of combining photometric surveys with intensity mapping". In: (Feb. 2021). arXiv: 2102.08116 [astro-ph.CO].

[415] Chirag Modi, Martin White, Anže Slosar, and Emanuele Castorina. "Reconstructing large-scale structure with neutral hydrogen surveys". In: *J. Cosm. Astrop. Phys.* 2019.11, 023 (Nov. 2019), p. 023. DOI: 10.1088/1475-7516/2019/11/023. arXiv: 1907.02330 [astro-ph.CO].

[416] Chirag Modi, Martin White, and Zvonimir Vlah. "Modeling CMB lensing cross correlations with CLEFT". In: *Journal of Cosmology and Astroparticle Physics* 2017.08 (Aug. 2017), pp. 009–009. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2017/08/009. URL: http://dx.doi.org/10.1088/1475-7516/2017/08/009.

[417] R. Moessner, B. Jain, and J. V. Villumsen. "The effect of weak lensing on the angular correlation function of faint galaxies". In: *Mon. Not. Roy. Astron. Soc.* 294 (1998), p. 291. DOI: 10.1046/j.1365-8711.1998.01225.x. arXiv: astro-ph/9708271.

[418] Paulo Montero-Camacho and Christopher M. Hirata. "Exploring circular polarization in the CMB due to conventional sources of cosmic birefringence". In: *JCAP* 08 (2018), p. 040. DOI: 10.1088/1475-7516/2018/08/040. arXiv: 1803.04505 [astro-ph.CO].

[419] Azadeh Moradinezhad Dizgah, Matteo Biagetti, Emiliano Sefusatti, Vincent Desjacques, and Jorge Noreña. "Primordial Non-Gaussianity from Biased Tracers: Likelihood Analysis of Real-Space Power Spectrum and Bispectrum". In: (Oct. 2020). arXiv: 2010.14523 [astro-ph.CO].

[420] Miguel F. Morales, Bryna Hazelton, Ian Sullivan, and Adam Beardsley. "Four Fundamental Foreground Power Spectrum Shapes for 21 cm Cosmology Observations". In: *Astrophys. J.* 752.2, 137 (June 2012), p. 137. DOI: 10.1088/0004-637X/752/2/137. arXiv: 1202.3830 [astro-ph.IM].

[421] Surhud More et al. "Hyper Suprime-Cam Year 3 results: Measurements of clustering of SDSS-BOSS galaxies, galaxy-galaxy lensing, and cosmic shear". In: *Phys. Rev. D* 108.12 (2023), p. 123520. DOI: 10.1103/PhysRevD.108.123520. arXiv: 2304.00703 [astro-ph.CO].

[422] Emily Moser et al. "The Circumgalactic Medium from the CAMELS Simulations: Forecasting Constraints on Feedback Processes from Future Sunyaev-Zeldovich Observations". In: *Astrophys. J.* 933.2, 133 (July 2022), p. 133. DOI: 10.3847/1538-4357/ac70c6. arXiv: 2201.02708 [astro-ph.CO].

[423] Tony Mroczkowski et al. "Astrophysics with the Spatially and Spectrally Resolved Sunyaev-Zeldovich Effects. A Millimetre/Submillimetre Probe of the Warm and Hot Universe". In: *Space Sci. Rev.* 215.1, 17 (Feb. 2019), p. 17. DOI: 10.1007/s11214-019-0581-2. arXiv: 1811.02310 [astro-ph.CO].

[424] Viatcheslav Mukhanov. *Physical Foundations of Cosmology*. Cambridge, UK: Cambridge University Press, 2005. ISBN: 978-0-521-80479-0.

[425] Moritz Münchmeyer, Mathew S. Madhavacheril, Simone Ferraro, Matthew C. Johnson, and Kendrick M. Smith. "Constraining local non-Gaussianities with kinetic Sunyaev-Zel'dovich tomography". In: *Phys. Rev. D* 100.8, 083508 (Oct. 2019), p. 083508. DOI: 10.1103/PhysRevD.100.083508. arXiv: 1810.13424 [astro-ph.CO].

[426] Julian B. Muñoz and Cora Dvorkin. "Efficient Computation of Galaxy Bias with Neutrinos and Other Relics". In: *Phys. Rev. D* 98.4 (2018), p. 043503. DOI: 10.1103/PhysRevD.98.043503. arXiv: 1805.11623 [astro-ph.CO].

[427] Julian B. Muñoz, Jordan Mirocha, John Chisholm, Steven R. Furlanetto, and Charlotte Mason. "Reionization after JWST: a photon budget crisis?" In: *Mon. Not. Roy. Astron. Soc.* 535.1 (2024), pp. L37–L43. DOI: 10.1093/mnrasl/slae086. arXiv: 2404.07250 [astro-ph.CO].

[428] Julian B. Muñoz, Yuxiang Qin, Andrei Mesinger, Steven G. Murray, Bradley Greig, and Charlotte Mason. "The Impact of the First Galaxies on Cosmic Dawn and Reionization". In: *arXiv e-prints*, arXiv:2110.13919 (Oct. 2021), arXiv:2110.13919. arXiv: 2110.13919 [astro-ph.CO].

[429] Adam D. Myers et al. "The Target-selection Pipeline for the Dark Energy Spectroscopic Instrument". In: *Astron. J.* 165.2 (2023), p. 50. DOI: 10.3847/1538-3881/aca5f9. arXiv: 2208.08518 [astro-ph.IM].

[430] Sigurd Naess et al. "The Atacama Cosmology Telescope: DR6 Maps". In: *arXiv e-prints* (Mar. 2025). arXiv: 2503.14451 [astro-ph.CO].

[431] Sigurd Naess et al. "The Atacama Cosmology Telescope: arcminute-resolution maps of 18 000 square degrees of the microwave sky from ACT 2008-2018 data combined with Planck". In: *J. Cosm. Astrop. Phys.* 2020.12, 046 (Dec. 2020), p. 046. DOI: 10.1088/1475-7516/2020/12/046. arXiv: 2007.07290 [astro-ph.IM].

[432] Daisuke Nagai, Andrey V. Kravtsov, and Alexey Vikhlinin. "Effects of Galaxy Formation on Thermodynamics of the Intracluster Medium". In: *Astrophys. J.* 668.1 (Oct. 2007), pp. 1–14. DOI: 10.1086/521328. arXiv: astro-ph/0703661 [astro-ph].

[433] Minami Nakane et al. "Lyα Emission at z = 7–13: Clear Evolution of Lyα Equivalent Width Indicating a Late Cosmic Reionization History". In: *Astrophys. J.* 967.1, 28 (May 2024), p. 28. DOI: 10.3847/1538-4357/ad38c2. arXiv: 2312.06804 [astro-ph.GA].

[434] Toshiya Namikawa, Duncan Hanson, and Ryuichi Takahashi. "Bias-hardened CMB lensing". In: *Mon. Not. R. Astron. Soc.* 431.1 (May 2013), pp. 609–620. DOI: 10.1093/mnras/stt195. arXiv: 1209.0091 [astro-ph.CO].

[435] Toshiya Namikawa, Anirban Roy, Blake D. Sherwin, Nicholas Battaglia, and David N. Spergel. "Constraining reionization with the first measurement of the cross-correlation between the CMB optical-depth fluctuations and the Compton y-map". In: *Phys. Rev. D* 104.6 (2021), p. 063514. DOI: 10.1103/PhysRevD.104.063514. arXiv: 2102.00975 [astro-ph.CO].

[436] Toshiya Namikawa and Ryuichi Takahashi. "Bias-Hardened CMB Lensing with Polarization". In: *Mon. Not. Roy. Astron. Soc.* 438.2 (2014), pp. 1507–1517. DOI: 10.1093/mnras/stt2290. arXiv: 1310.2372 [astro-ph.CO].

[437] D. Nelson et al. "The illustris simulation: Public data release". In: *Astronomy and Computing* 13 (Nov. 2015), pp. 12–37. DOI: 10.1016/j.ascom.2015.09.003. arXiv: 1504.00362 [astro-ph.CO].

[438] L. B. Newburgh et al. "HIRAX: A Probe of Dark Energy and Radio Transients". In: *Proc. SPIE Int. Soc. Opt. Eng.* 9906 (2016). Ed. by Helen J. Hall, Roberto Gilmozzi, and Heather K. Marshall, p. 99065X. DOI: 10.1117/12.2234286. arXiv: 1607.02059 [astro-ph.IM].

[439] J. S. Nico et al. "Measurement of the neutron lifetime by counting trapped protons in a cold neutron beam". In: *Phys. Rev. C* 71 (2005), p. 055502. DOI: 10.1103/PhysRevC.71.055502. arXiv: nucl-ex/0411041.

[440] Andrina Nicola et al. "Galaxy bias in the era of LSST: perturbative bias expansions". In: *J. Cosm. Astrop. Phys.* 2024.2, 015 (Feb. 2024), p. 015. DOI: 10.1088/1475-7516/2024/02/015. arXiv: 2307.03226 [astro-ph.CO].

[441] Ivan Nikolić, Andrei Mesinger, Yuxiang Qin, and Adélie Gorce. "Inferring reionization and galaxy properties from the patchy kinetic Sunyaev–Zel'dovich signal". In: *Mon. Not. Roy. Astron. Soc.* 526.2 (2023), pp. 3170–3183. DOI: 10.1093/mnras/stad2961. arXiv: 2307.01265 [astro-ph.CO].

[442] Y. Noh, M. White, and N. Padmanabhan. "Reconstructing baryon oscillations". In: *Phys. Rev. D* 80.12, 123501 (Dec. 2009), p. 123501. DOI: 10.1103/PhysRevD.80.123501. arXiv: 0909.1802.

[443] Andrej Obuljen, Emanuele Castorina, Francisco Villaescusa-Navarro, and Matteo Viel. "High-redshift post-reionization cosmology with 21cm intensity mapping". In: *JCAP* 05 (2018), p. 004. DOI: 10.1088/1475-7516/2018/05/004. arXiv: 1709.07893 [astro-ph.CO].

[444] Andrej Obuljen, Francisco Villaescusa-Navarro, Emanuele Castorina, and Matteo Viel. "Baryon Acoustic Oscillations reconstruction with pixels". In: *JCAP* 09 (2017), p. 012. DOI: 10.1088/1475-7516/2017/09/012. arXiv: 1610.05768 [astro-ph.CO].

[445] Takemi Okamoto and Wayne Hu. "CMB lensing reconstruction on the full sky". In: *Phys. Rev. D* 67 (2003), p. 083002. DOI: 10.1103/PhysRevD.67.083002. arXiv: astro-ph/0301031.

[446] Stephen J. Osborne, Duncan Hanson, and Olivier Doré. "Extragalactic foreground contamination in temperature-based CMB lens reconstruction". In: *J. Cosm. Astrop. Phys.* 2014.3, 024 (Mar. 2014), p. 024. DOI: 10.1088/1475-7516/2014/03/024. arXiv: 1310.7547 [astro-ph.CO].

[447] J. P. Ostriker and E. T. Vishniac. "Generation of Microwave Background Fluctuations from Nonlinear Perturbations at the ERA of Galaxy Formation". In: *Astrophys. J. Lett.* 306 (July 1986), p. L51. DOI: 10.1086/184704.

[448] Jeremiah P. Ostriker, Paul Bode, and Arif Babul. "A Simple and Accurate Model for Intracluster Gas". In: *Astrophys. J.* 634.2 (Dec. 2005), pp. 964–976. DOI: 10.1086/497122. arXiv: astro-ph/0504334 [astro-ph].

[449] N. Padmanabhan, M. White, and J. D. Cohn. "Reconstructing baryon oscillations: A Lagrangian theory perspective". In: *Phys. Rev. D* 79.6, 063523 (Mar. 2009), p. 063523. DOI: 10.1103/PhysRevD.79.063523. arXiv: 0812.2905.

[450] N. Padmanabhan, X. Xu, D. J. Eisenstein, R. Scalzo, A. J. Cuesta, K. T. Mehta, and E. Kazin. "A 2 per cent distance to z = 0.35 by reconstructing baryon acoustic oscillations - I. Methods and application to the Sloan Digital Sky Survey". In: *Mon. Not. R. Astron. Soc.* 427 (Dec. 2012), pp. 2132–2145. DOI: 10.1111/j.1365-2966.2012.21888.x. arXiv: 1202.0090.

[451] Nikhil Padmanabhan et al. "The clustering of luminous red galaxies in the Sloan Digital Sky Survey imaging data". In: *Mon. Not. R. Astron. Soc.* 378.3 (July 2007), pp. 852–872. DOI: 10.1111/j.1365-2966.2007.11593.x. arXiv: astro-ph/0605302 [astro-ph].

[452] L. Pagano, J. -M. Delouis, S. Mottet, J. -L. Puget, and L. Vibert. "Reionization optical depth determination from Planck HFI data with ten percent accuracy". In: *Astron. Astrophys.* 635 (2020), A99. DOI: 10.1051/0004-6361/201936630. arXiv: 1908.09856 [astro-ph.CO].

[453] Z. Pan et al. "Measurement of gravitational lensing of the cosmic microwave background using SPT-3G 2018 data". In: *Phys. Rev. D* 108.12 (2023), p. 122005. DOI: 10.1103/PhysRevD.108.122005. arXiv: 2308.11608 [astro-ph.CO].

[454] S. Pandey et al. "Cross-correlation of Dark Energy Survey Year 3 lensing data with ACT and P l a n c k thermal Sunyaev-Zel'dovich effect observations. II. Modeling and constraints on halo pressure profiles". In: *Phys. Rev. D* 105.12, 123526 (June 2022), p. 123526. DOI: 10.1103/PhysRevD.105.123526. arXiv: 2108.01601 [astro-ph.CO].

[455] Hyunbae Park, Marcelo A. Alvarez, and J. Richard Bond. "The Impact of Baryonic Physics on the Kinetic Sunyaev-Zel'dovich Effect". In: *Astrophys. J.* 853.2, 121 (Feb. 2018), p. 121. DOI: 10.3847/1538-4357/aaa0da. arXiv: 1710.02792 [astro-ph.CO].

[456] Hyunbae Park, Paul R. Shapiro, Eiichiro Komatsu, Ilian T. Iliev, Kyungjin Ahn, and Garrelt Mellema. "The Kinetic Sunyaev-Zel'dovich effect as a probe of the physics of cosmic reionization: the effect of self-regulated reionization". In: *Astrophys. J.* 769 (2013), p. 93. DOI: 10.1088/0004-637X/769/2/93. arXiv: 1301.3607 [astro-ph.CO].

[457] Sanjaykumar Patil, Srinivasan Raghunathan, and Christian L. Reichardt. "Suppressing the Thermal SZ-induced Variance in CMB-cluster Lensing Estimators". In: *The Astrophysical Journal* 888.1 (Dec. 2019), p. 9. DOI: 10.3847/1538-4357/ab55dd. URL: https://doi.org/10.3847/1538-4357/ab55dd.

[458] R. W. Pattie Jr. et al. "Measurement of the neutron lifetime using a magneto-gravitational trap and in situ detection". In: *Science* 360.6389 (2018), pp. 627–632. DOI: 10.1126/science.aan8895. arXiv: 1707.01817 [nucl-ex].

[459] Andreas H. Pawlik, Alireza Rahmati, Joop Schaye, Myoungwon Jeon, and Claudio Dalla Vecchia. "The Aurora radiation-hydrodynamical simulations of reionization: calibration and first results". In: *Mon. Not. R. Astron. Soc.* 466.1 (Apr. 2017), pp. 960–973. DOI: 10.1093/mnras/stw2869. arXiv: 1603.00034 [astro-ph.GA].

[460] John A. Peacock. *Cosmological Physics*. 1999.

[461] Marcos Pellejero-Ibanez, Raul E. Angulo, Matteo Zennaro, Jens Stuecker, Sergio Contreras, Giovanni Arico, and Francisco Maion. "The bacco simulation project: bacco hybrid Lagrangian bias expansion model in redshift space". In: *Mon. Not. Roy. Astron. Soc.* 520.3 (2023), pp. 3725–3741. DOI: 10.1093/mnras/stad368. arXiv: 2207.06437 [astro-ph.CO].

[462] Will J. Percival et al. "Cosmology with the MaunaKea Spectroscopic Explorer". In: (Mar. 2019). arXiv: 1903.03158 [astro-ph.CO].

[463] Will J. Percival et al. "Parameter constraints for flat cosmologies from CMB and 2dFGRS power spectra". In: *Mon. Not. Roy. Astron. Soc.* 337 (2002), p. 1068. DOI: 10.1046/j.1365-8711.2002.06001.x. arXiv: astro-ph/0206256.

[464] S. Perlmutter et al. "Measurements of Ω and Λ from 42 High-Redshift Supernovae". In: *The Astrophysical Journal* 517.2 (1999), pp. 565–586. DOI: 10.1086/307221.

[465] Fred M. Persi, David N. Spergel, Renyue Cen, and J. P. Ostriker. "Hot Gas in Superclusters and Microwave Background Distortions". In: *Astrophys. J.* 442 (Mar. 1995), p. 1. DOI: 10.1086/175416.

[466] Massimo Persic and Paolo Salucci. "The baryon content of the universe". In: *Mon. Not. R. Astron. Soc.* 258.1 (Sept. 1992), 14P–18P. DOI: 10.1093/mnras/258.1.14P. arXiv: astro-ph/0502178 [astro-ph].

[467] A. Pezzotta et al. "Euclid preparation. TBD. Galaxy power spectrum modelling in real space". In: (Dec. 2023). arXiv: 2312.00679 [astro-ph.CO].

[468] Oliver H. E. Philcox, Mikhail M. Ivanov, Marko Simonović, and Matias Zaldarriaga. "Combining Full-Shape and BAO Analyses of Galaxy Power Spectra: A 1.6\% CMB-independent constraint on H_0". In: *JCAP* 05 (2020), p. 032. DOI: 10.1088/1475-7516/2020/05/032. arXiv: 2002.04035 [astro-ph.CO].

[469] Giulia Piccirilli, Giulio Fabbian, David Alonso, Kate Storey-Fisher, Julien Carron, Antony Lewis, and Carlos García-García. "Growth history and quasar bias evolution at z ¡ 3 from Quaia". In: *J. Cosm. Astrop. Phys.* 2024.6, 012 (June 2024), p. 012. DOI: 10.1088/1475-7516/2024/06/012. arXiv: 2402.05761 [astro-ph.CO].

[470] Annalisa Pillepich et al. "Simulating galaxy formation with the IllustrisTNG model". In: *Mon. Not. R. Astron. Soc.* 473.3 (Jan. 2018), pp. 4077–4106. DOI: 10.1093/mnras/stx2656. arXiv: 1703.02970 [astro-ph.GA].

[471] Planck Collaboration et al. "Planck 2015 results. I. Overview of products and scientific results". In: *Astron. Astrophys.* 594, A1 (Sept. 2016), A1. DOI: 10.1051/0004-6361/201527101. arXiv: 1502.01582 [astro-ph.CO].

[472] Planck Collaboration et al. "Planck 2015 results. XX. Constraints on inflation". In: *Astron. Astrophys.* 594, A20 (Sept. 2016), A20. DOI: 10.1051/0004-6361/201525898. arXiv: 1502.02114 [astro-ph.CO].

[473] Planck Collaboration et al. "Planck 2018 results. I. Overview and the cosmological legacy of Planck". In: *arXiv e-prints*, arXiv:1807.06205 (July 2018), arXiv:1807.06205. arXiv: 1807.06205 [astro-ph.CO].

[474] Planck Collaboration et al. "Planck 2018 results. V. CMB power spectra and likelihoods". In: *Astron. Astrophys.* 641, A5 (Sept. 2020), A5. DOI: 10.1051/0004-6361/201936386. arXiv: 1907.12875 [astro-ph.CO].

[475] Planck Collaboration et al. "Planck 2018 results. VI. Cosmological parameters". In: *arXiv e-prints*, arXiv:1807.06209 (July 2018), arXiv:1807.06209. arXiv: 1807.06209 [astro-ph.CO].

[476] Planck Collaboration et al. "Planck 2018 results. X. Constraints on inflation". In: *arXiv e-prints*, arXiv:1807.06211 (July 2018), arXiv:1807.06211. arXiv: 1807.06211 [astro-ph.CO].

[477] Planck Collaboration et al. "Planck early results. XVIII. The power spectrum of cosmic infrared background anisotropies". In: *Astron. Astrophys.* 536, A18 (Dec. 2011), A18. DOI: 10.1051/0004-6361/201116461. arXiv: 1101.2028 [astro-ph.CO].

[478] Planck Collaboration et al. "Planck intermediate results. LVII. Joint Planck LFI and HFI data processing". In: *Astron. Astrophys.* 643, A42 (Nov. 2020), A42. DOI: 10.1051/0004-6361/202038073. arXiv: 2007.04997 [astro-ph.CO].

[479] Jonathan C. Pober. "The impact of foregrounds on redshift space distortion measurements with the highly redshifted 21-cm line". In: *Mon. Not. R. Astron. Soc.* 447.2 (Feb. 2015), pp. 1705–1712. DOI: 10.1093/mnras/stu2575. arXiv: 1411.2050 [astro-ph.CO].

[480] Jonathan C. Pober et al. "What Next-Generation 21 cm Power Spectrum Measurements Can Teach Us About the Epoch of Reionization". In: *Astrophys. J.* 782 (2014), p. 66. DOI: 10.1088/0004-637X/782/2/66. arXiv: 1310.7031 [astro-ph.CO].

[481] Jonathan C. Pober et al. "The Baryon Acoustic Oscillation Broadband and Broadbeam Array: Design Overview and Sensitivity Forecasts". In: *Astron. J.* 145.3, 65 (Mar. 2013), p. 65. DOI: 10.1088/0004-6256/145/3/65. arXiv: 1210.2413 [astro-ph.CO].

[482] R. A. Porto, L. Senatore, and M. Zaldarriaga. "The Lagrangian-space Effective Field Theory of large scale structures". In: *J. Cosm. Astrop. Phys.* 5, 022 (May 2014), p. 022. DOI: 10.1088/1475-7516/2014/05/022. arXiv: 1311.2168.

[483] Vivian Poulin, José Luis Bernal, Ely D. Kovetz, and Marc Kamionkowski. "Sigma-8 tension is a drag". In: *Phys. Rev. D* 107.12 (2023), p. 123538. DOI: 10.1103/PhysRevD.107.123538. arXiv: 2209.06217 [astro-ph.CO].

[484] Vivian Poulin, Tristan L. Smith, Tanvi Karwal, and Marc Kamionkowski. "Early Dark Energy can Resolve the Hubble Tension". In: *Physical Review Letters* 122.22 (June 2019). ISSN: 1079-7114. DOI: 10.1103/physrevlett.122.221301. URL: http://dx.doi.org/10.1103/PhysRevLett.122.221301.

[485] Geraint Pratten and Antony Lewis. "Impact of post-Born lensing on the CMB". In: *JCAP* 08 (2016), p. 047. DOI: 10.1088/1475-7516/2016/08/047. arXiv: 1605.05662 [astro-ph.CO].

[486] Heather Prince, Kavilan Moodley, Jethro Ridl, and Martin Bucher. "Real space lensing reconstruction using cosmic microwave background polarization". In: *J. Cosm. Astrop. Phys.* 2018.1, 034 (Jan. 2018), p. 034. DOI: 10.1088/1475-7516/2018/01/034. arXiv: 1709.02227 [astro-ph.CO].

[487] Anthony R. Pullen and Christopher M. Hirata. "Systematic Effects in Large-Scale Angular Power Spectra of Photometric Quasars and Implications for Constraining Primordial Non-Gaussianity". In: *Publ. Astron. Soc. Pacific* 125.928 (June 2013), p. 705. DOI: 10.1086/671189. arXiv: 1212.4500 [astro-ph.CO].

[488] Roland de Putter. "Primordial physics from large-scale structure beyond the power spectrum". In: (Feb. 2018). arXiv: 1802.06762 [astro-ph.CO].

[489] Wenzer Qin, Katelin Schutz, Aaron Smith, Enrico Garaldi, Rahul Kannan, Tracy R. Slatyer, and Mark Vogelsberger. "An Effective Bias Expansion for 21 cm Cosmology in Redshift Space". In: (May 2022). arXiv: 2205.06270 [astro-ph.CO].

[490] Frank J. Qu et al. "The Atacama Cosmology Telescope DR6 and DESI: Structure growth measurements from the cross-correlation of DESI Legacy Imaging galaxies and CMB lensing from ACT DR6 and Planck PR4". In: (Oct. 2024). arXiv: 2410.10808 [astro-ph.CO].

[491] Frank J. Qu, Anthony Challinor, and Blake D. Sherwin. "CMB lensing with shear-only reconstruction on the full sky". In: *arXiv e-prints* (Aug. 2022), arXiv:2208.14988. arXiv: 2208.14988 [astro-ph.CO].

[492] Frank J. Qu, Marius Millea, and Emmanuel Schaan. "Impact & Mitigation of Polarized Extragalactic Foregrounds on Bayesian Cosmic Microwave Background Lensing". In: (June 2024). arXiv: 2406.15351 [astro-ph.CO].

[493] Frank J. Qu et al. "The Atacama Cosmology Telescope: A Measurement of the DR6 CMB Lensing Power Spectrum and Its Implications for Structure Growth". In: *Astrophys. J.* 962.2, 112 (Feb. 2024), p. 112. DOI: 10.3847/1538-4357/acfe06. arXiv: 2304.05202 [astro-ph.CO].

[494] S. Raghunathan et al. "First Constraints on the Epoch of Reionization Using the Non-Gaussianity of the Kinematic Sunyaev-Zel'dovich Effect from the South Pole Telescope and Herschel-SPIRE Observations". In: *Phys. Rev. Lett.* 133.12 (2024), p. 121004. DOI: 10.1103/PhysRevLett.133.121004. arXiv: 2403.02337 [astro-ph.CO].

[495] Marco Raveri et al. "A joint reconstruction of dark energy and modified growth evolution". In: *arXiv e-prints*, arXiv:2107.12990 (July 2021), arXiv:2107.12990. arXiv: 2107.12990 [astro-ph.CO].

[496] C. L. Reichardt et al. "An Improved Measurement of the Secondary Cosmic Microwave Background Anisotropies from the SPT-SZ + SPTpol Surveys". In: *Astrophys. J.* 908.2 (2021), p. 199. DOI: 10.3847/1538-4357/abd407. arXiv: 2002.06197 [astro-ph.CO].

[497] Mathieu Remazeilles, Jacques Delabrouille, and Jean-François Cardoso. "CMB and SZ effect separation with constrained Internal Linear Combinations". In: *Monthly Notices of the Royal Astronomical Society* 410.4 (Nov. 2011), pp. 2481–2487. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2010.17624.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2010.17624.x.

[498] Massimo Ricotti, Jeremiah P. Ostriker, and Katherine J. Mack. "Effect of Primordial Black Holes on the Cosmic Microwave Background and Cosmological Parameter Estimates". In: *Astrophys. J.* 680 (2008), p. 829. DOI: 10.1086/587831. arXiv: 0709.0524 [astro-ph].

[499] Bernardita Ried Guachalla, Emmanuel Schaan, Boryana Hadzhiyska, and Simone Ferraro. "Velocity reconstruction in the era of DESI and Rubin/LSST. I. Exploring spectroscopic, photometric, and hybrid samples". In: *Phys. Rev. D* 109.10, 103533 (May 2024), p. 103533. DOI: 10.1103/PhysRevD.109.103533. arXiv: 2312.12435 [astro-ph.CO].

[500] Adam G. Riess and et al. "Cosmic Distances Calibrated to 1% Precision with Gaia EDR3 Parallaxes and Hubble Space Telescope Photometry of 75 Cepheids". In: *Astrophys. J. Lett.* 908.1 (2021), p. L6. DOI: 10.3847/2041-8213/ac2f7c. arXiv: 2112.04510 [astro-ph.CO].

[501] Adam G. Riess et al. "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant". In: *The Astronomical Journal* 116.3 (1998), pp. 1009–1038. DOI: 10.1086/300499.

[502] Naomi Clare Robertson et al. "Strong detection of the CMB lensing and galaxy weak lensing cross-correlation from ACT-DR4, Planck Legacy, and KiDS-1000". In: *Astron. Astrophys.* 649 (2021), A146. DOI: 10.1051/0004-6361/202039975. arXiv: 2011.11613 [astro-ph.CO].

[503] Erik Rosenberg, Steven Gratton, and George Efstathiou. "CMB power spectra and cosmological parameters from Planck PR4 with CamSpec". In: *Mon. Not. Roy. Astron. Soc.* 517.3 (2022), pp. 4620–4636. DOI: 10.1093/mnras/stac2744. arXiv: 2205.10869 [astro-ph.CO].

[504] Ashley J. Ross, Lado Samushia, Cullan Howlett, Will J. Percival, Angela Burden, and Marc Manera. "The clustering of the SDSS DR7 main Galaxy sample – I. A 4 per cent distance measure at $z = 0.15$". In: *Mon. Not. Roy. Astron. Soc.* 449.1 (2015), pp. 835–847. DOI: 10.1093/mnras/stv154. arXiv: 1409.3242 [astro-ph.CO].

[505] Ashley J. Ross et al. "The clustering of galaxies in the SDSS-III DR9 Baryon Oscillation Spectroscopic Survey: constraints on primordial non-Gaussianity". In: *Mon. Not. R. Astron. Soc.* 428.2 (Jan. 2013), pp. 1116–1127. DOI: 10.1093/mnras/sts094. arXiv: 1208.1491 [astro-ph.CO].

[506] A. Roy, A. Lapi, D. Spergel, and C. Baccigalupi. "Observing patchy reionization with future CMB polarization experiments". In: *J. Cosm. Astrop. Phys.* 2018.5, 014 (May 2018), p. 014. DOI: 10.1088/1475-7516/2018/05/014. arXiv: 1801.02393 [astro-ph.CO].

[507] Anirban Roy, Alexander van Engelen, Vera Gluscevic, and Nicholas Battaglia. "Probing the Circumgalactic Medium with Cosmic Microwave Background Polarization Statistical Anisotropy". In: *Astrophys. J.* 951.1 (2023), p. 50. DOI: 10.3847/1538-4357/acd194. arXiv: 2201.05076 [astro-ph.CO].

[508] John J. Ruan, Matthew McQuinn, and Scott F. Anderson. "Detection of Quasar Feedback from the Thermal Sunyaev-Zel'dovich Effect in Planck". In: *Astrophys. J.* 802.2, 135 (Apr. 2015), p. 135. DOI: 10.1088/0004-637X/802/2/135. arXiv: 1502.01723 [astro-ph.CO].

[509] David Rubin et al. "Union Through UNITY: Cosmology with 2,000 SNe Using a Unified Bayesian Framework". In: *arXiv e-prints* (Nov. 2023). arXiv: 2311.12098 [astro-ph.CO].

[510] Sayan Saha, Louis Legrand, and Julien Carron. "Cluster profiles from beyond-the-QE CMB lensing mass maps". In: *J. Cosm. Astrop. Phys.* 2024.1, 024 (Jan. 2024), p. 024. DOI: 10.1088/1475-7516/2024/01/024. arXiv: 2307.11711 [astro-ph.CO].

[511] Noah Sailer et al. "Cosmological constraints from the cross-correlation of DESI Luminous Red Galaxies with CMB lensing from Planck PR4 and ACT DR6". In: *JCAP* 06 (2025), p. 008. DOI: 10.1088/1475-7516/2025/06/008. arXiv: 2407.04607 [astro-ph.CO].

[512] Noah Sailer, Emanuele Castorina, Simone Ferraro, and Martin White. "Cosmology at high redshift — a probe of fundamental physics". In: *JCAP* 12.12 (2021), p. 049. DOI: 10.1088/1475-7516/2021/12/049. arXiv: 2106.09713 [astro-ph.CO].

[513] Noah Sailer, Shi-Fan Chen, and Martin White. "Optical depth to reionization from perturbative 21 cm clustering". In: *JCAP* 10 (2022), p. 007. DOI: 10.1088/1475-7516/2022/10/007. arXiv: 2205.11504 [astro-ph.CO].

[514] Noah Sailer, Joseph DeRose, Simone Ferraro, Shi-Fan Chen, Rongpu Zhou, Martin White, Joshua Kim, and Mathew Madhavacheril. "Evolution of structure growth during dark energy domination: Insights from the cross-correlation of DESI galaxies with CMB lensing and galaxy magnification". In: *Phys. Rev. D* 111.10, 103540 (May 2025), p. 103540. DOI: 10.1103/27rg-tq8z. arXiv: 2503.24385 [astro-ph.CO].

[515] Noah Sailer, Gerrit Farren, Simone Ferraro, and Martin White. "A cookbook for CMB lensing cross-correlations". In: (2024). arXiv: 24MM.XXXXX [astro-ph.CO].

[516] Noah Sailer, Gerrit S. Farren, Simone Ferraro, and Martin White. "Dispuτable: the high cost of a low optical depth". In: *arXiv e-prints* (Apr. 2025). arXiv: 2504.16932 [astro-ph.CO].

[517] Noah Sailer, Simone Ferraro, and Emmanuel Schaan. "Foreground-immune CMB lensing reconstruction with polarization". In: *Phys. Rev. D* 107.2 (2023), p. 023504. DOI: 10.1103/PhysRevD.107.023504. arXiv: 2211.03786 [astro-ph.CO].

[518] Noah Sailer, Boryana Hadzhiyska, and Simone Ferraro. "Bias hardened estimators of patchy screening profiles". In: *arXiv e-prints* (June 2025). arXiv: 2506.17217 [astro-ph.CO].

[519] Noah Sailer, Emmanuel Schaan, and Simone Ferraro. "Lower bias, lower noise CMB lensing with foreground-hardened estimators". In: *Phys. Rev. D* 102.6 (2020), p. 063517. DOI: 10.1103/PhysRevD.102.063517. arXiv: 2007.04325 [astro-ph.CO].

[520] Noah Sailer, Emmanuel Schaan, Simone Ferraro, Omar Darwish, and Blake Sherwin. "Optimal multifrequency weighting for CMB lensing". In: *Phys. Rev. D* 104.12 (2021), p. 123514. DOI: 10.1103/PhysRevD.104.123514. arXiv: 2108.01663 [astro-ph.CO].

[521] Andrés N. Salcedo, Hao-Yi Wu, Eduardo Rozo, David H. Weinberg, Chun-Hao To, Tomomi Sunayama, and Andy Lee. "Dark Energy Survey Year 1 Clusters are Consistent with Planck". In: *arXiv e-prints*, arXiv:2310.03944 (Oct. 2023), arXiv:2310.03944. DOI: 10.48550/arXiv.2310.03944. arXiv: 2310.03944 [astro-ph.CO].

[522] Christoph Saulder et al. "Target selection for the DESI Peculiar Velocity Survey". In: *Mon. Not. Roy. Astron. Soc.* 525.1 (2023), pp. 1106–1125. DOI: 10.1093/mnras/stad2200. arXiv: 2302.13760 [astro-ph.CO].

[523] Emmanuel Schaan, Simone Ferraro, et al. "Evidence for the kinematic Sunyaev-Zel'dovich effect with the Atacama Cosmology Telescope and velocity reconstruction from the Baryon Oscillation Spectroscopic Survey". In: *Phys. Rev. D* 93.8 (2016), p. 082002. DOI: 10.1103/PhysRevD.93.082002. arXiv: 1510.06442 [astro-ph.CO].

[524] Emmanuel Schaan and Simone Ferraro. "Foreground-Immune Cosmic Microwave Background Lensing with Shear-Only Reconstruction". In: *Phys. Rev. Lett.* 122.18, 181301 (May 2019), p. 181301. DOI: 10.1103/PhysRevLett.122.181301. arXiv: 1804.06403 [astro-ph.CO].

[525] Emmanuel Schaan et al. "Atacama Cosmology Telescope: Combined kinematic and thermal Sunyaev-Zel'dovich measurements from BOSS CMASS and LOWZ halos". In: *Phys. Rev. D* 103.6, 063513 (Mar. 2021), p. 063513. DOI: 10.1103/PhysRevD.103.063513. arXiv: 2009.05557 [astro-ph.CO].

[526] Edward F. Schlafly, Aaron M. Meisner, and Gregory M. Green. "The unWISE Catalog: Two Billion Infrared Sources from Five Years of WISE Imaging". In: *Astrophys. J. Supp.* 240.2, 30 (Feb. 2019), p. 30. DOI: 10.3847/1538-4365/aafbea. arXiv: 1901.03337 [astro-ph.IM].

[527] Edward F. Schlafly et al. "Survey Operations for the Dark Energy Spectroscopic Instrument". In: *Astron. J.* 166.6, 259 (Dec. 2023), p. 259. DOI: 10.3847/1538-3881/ad0832. arXiv: 2306.06309 [astro-ph.CO].

[528] David Schlegel, Juna A. Kollmeier, and Simone Ferraro. "The MegaMapper: a $z > 2$ spectroscopic instrument for the study of Inflation and Dark Energy". In: *Bull. Am. Astron. Soc.* Vol. 51. Sept. 2019, 229, p. 229. arXiv: 1907.11171 [astro-ph.IM].

[529] David J. Schlegel, Douglas P. Finkbeiner, and Marc Davis. "Maps of dust IR emission for use in estimation of reddening and CMBR foregrounds". In: *Astrophys. J.* 500 (1998), p. 525. DOI: 10.1086/305772. arXiv: astro-ph/9710327.

[530] F. Schmidt. "Effect of relative velocity and density perturbations between baryons and dark matter on the clustering of galaxies". In: *Phys. Rev. D* 94.6, 063508 (Sept. 2016), p. 063508. DOI: 10.1103/PhysRevD.94.063508. arXiv: 1602.09059.

[531] Fabian Schmidt and Marc Kamionkowski. "Halo clustering with nonlocal non-Gaussianity". In: *Physical Review D* 82.10 (Nov. 2010). ISSN: 1550-2368. DOI: 10.1103/physrevd.82.103002. URL: http://dx.doi.org/10.1103/PhysRevD.82.103002.

[532] M. Schmittfull, T. Baldauf, and M. Zaldarriaga. "Iterative initial condition reconstruction". In: *Phys. Rev. D* 96.2, 023505 (July 2017), p. 023505. DOI: 10.1103/PhysRevD.96.023505. arXiv: 1704.06634.

[533] Marcel Schmittfull, Yu Feng, Florian Beutler, Blake Sherwin, and Man Yat Chu. "Eulerian BAO reconstructions and N -point statistics". In: *Phys. Rev. D* 92.12, 123522 (Dec. 2015), p. 123522. DOI: 10.1103/PhysRevD.92.123522. arXiv: 1508.06972 [astro-ph.CO].

[534] Marcel Schmittfull and Uros Seljak. "Parameter constraints from cross-correlation of CMB lensing with galaxy clustering". In: *Phys. Rev. D* 97.12, 123540 (June 2018), p. 123540. DOI: 10.1103/PhysRevD.97.123540. arXiv: 1710.09465 [astro-ph.CO].

[535] Denise M. Schmitz, Christopher M. Hirata, Jonathan Blazek, and Elisabeth Krause. "Time evolution of intrinsic alignments of galaxies". In: *JCAP* 07 (2018), p. 030. DOI: 10.1088/1475-7516/2018/07/030. arXiv: 1805.02649 [astro-ph.CO].

[536] Theo Schutt, Abhishek S. Maniyar, Emmanuel Schaan, William R. Coulton, and Nishant Mishra. "New temperature inversion estimator to detect CMB patchy screening by large-scale structure". In: *Phys. Rev. D* 109.10 (2024), p. 103539. DOI: 10.1103/PhysRevD.109.103539. arXiv: 2401.13040 [astro-ph.CO].

[537] D. M. Scolnic et al. "The Complete Light-curve Sample of Spectroscopically Confirmed SNe Ia from Pan-STARRS1 and Cosmological Constraints from the Combined Pantheon Sample". In: *Astrophys. J.* 859.2 (2018), p. 101. DOI: 10.3847/1538-4357/aab9bb. arXiv: 1710.00845 [astro-ph.CO].

[538] L. F. Secco et al. "Dark Energy Survey Year 3 results: Cosmology from cosmic shear and robustness to modeling uncertainty". In: *Phys. Rev. D* 105.2 (2022), p. 023515. DOI: 10.1103/PhysRevD.105.023515. arXiv: 2105.13544 [astro-ph.CO].

[539] Emiliano Sefusatti, Martin Crocce, Roman Scoccimarro, and Hugh Couchman. "Accurate Estimators of Correlation Functions in Fourier Space". In: *Mon. Not. Roy. Astron. Soc.* 460.4 (2016), pp. 3624–3636. DOI: 10.1093/mnras/stw1229. arXiv: 1512.07295 [astro-ph.CO].

[540] Neelima Sehgal et al. "CMB-HD: An Ultra-Deep, High-Resolution Millimeter-Wave Survey Over Half the Sky". In: (June 2019). arXiv: 1906.10134 [astro-ph.CO].

[541] Neelima Sehgal, Paul Bode, Sudeep Das, Carlos Hernand ez-Monteagudo, Kevin Huffenberger, Yen-Ting Lin, Jeremiah P. Ostriker, and Hy Trac. "Simulations of the Microwave Sky". In: *Astrophys. J.* 709.2 (Feb. 2010), pp. 920–936. DOI: 10.1088/0004-637X/709/2/920. arXiv: 0908.0540 [astro-ph.CO].

[542] Neelima Sehgal et al. *CMB-HD: Astro2020 RFI Response*. 2020. arXiv: 2002.12714 [astro-ph.CO].

[543] L. Senatore. "Bias in the effective field theory of large scale structures". In: *J. Cosm. Astrop. Phys.* 11, 007 (Nov. 2015), p. 007. DOI: 10.1088/1475-7516/2015/11/007. arXiv: 1406.7843.

[544] L. Senatore and M. Zaldarriaga. "The IR-resummed Effective Field Theory of Large Scale Structures". In: *J. Cosm. Astrop. Phys.* 2, 013 (Feb. 2015), p. 13. DOI: 10.1088/1475-7516/2015/02/013. arXiv: 1404.5954.

[545] H.-J. Seo, F. Beutler, A. J. Ross, and S. Saito. "Modeling the reconstructed BAO in Fourier space". In: *Mon. Not. R. Astron. Soc.* 460 (Aug. 2016), pp. 2453–2471. DOI: 10.1093/mnras/stw1138. arXiv: 1511.00663.

[546] Hee-Jong Seo, Scott Dodelson, John Marriner, Dave Mcginnis, Albert Stebbins, Chris Stoughton, and Alberto Vallinotto. "A GROUND-BASED 21 cm BARYON ACOUSTIC OSCILLATION SURVEY". In: *The Astrophysical Journal* 721.1 (Aug. 2010), pp. 164–173. DOI: 10.1088/0004-637x/721/1/164. URL: https://doi.org/10.1088/0004-637x/721/1/164.

[547] Hee-Jong Seo and Daniel J. Eisenstein. "Improved Forecasts for the Baryon Acoustic Oscillations and Cosmological Distance Scale". In: *The Astrophysical Journal* 665.1 (Aug. 2007), pp. 14–24. ISSN: 1538-4357. DOI: 10.1086/519549. URL: http://dx.doi.org/10.1086/519549.

[548] Hee-Jong Seo and Daniel J. Eisenstein. "Probing Dark Energy with Baryonic Acoustic Oscillations from Future Large Galaxy Redshift Surveys". In: *The Astrophysical Journal* 598.2 (Dec. 2003), pp. 720–740. ISSN: 1538-4357. DOI: 10.1086/379122. URL: http://dx.doi.org/10.1086/379122.

[549] Alice E. Shapley, Charles C. Steidel, Max Pettini, and Kurt L. Adelberger. "Rest-Frame Ultraviolet Spectra of z~3 Lyman Break Galaxies". In: *Astrophys. J.* 588.1 (May 2003), pp. 65–89. DOI: 10.1086/373922. arXiv: astro-ph/0301230 [astro-ph].

[550] R. K. Sheth, H. J. Mo, and G. Tormen. "Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes". In: *Monthly Notices of the Royal Astronomical Society* 323.1 (May 2001), pp. 1–12. ISSN: 1365-2966. DOI: 10.1046/j.1365-8711.2001.04006.x. URL: http://dx.doi.org/10.1046/j.1365-8711.2001.04006.x.

[551] J. Michael Shull, Britton D. Smith, and Charles W. Danforth. "The Baryon Census in a Multiphase Intergalactic Medium: 30% of the Baryons May Still be Missing". In: *Astrophys. J.* 759.1, 23 (Nov. 2012), p. 23. DOI: 10.1088/0004-637X/759/1/23. arXiv: 1112.2706 [astro-ph.CO].

[552] Joseph Harry Silber et al. "The Robotic Multiobject Focal Plane System of the Dark Energy Spectroscopic Instrument (DESI)". In: *Astron. J.* 165.1, 9 (Jan. 2023), p. 9. DOI: 10.3847/1538-3881/ac9ab1. arXiv: 2205.09014 [astro-ph.IM].

[553] Vimal Simha and Gary Steigman. "Constraining The Early-Universe Baryon Density And Expansion Rate". In: *JCAP* 06 (2008), p. 016. DOI: 10.1088/1475-7516/2008/06/016. arXiv: 0803.3465 [astro-ph].

[554] C. Skordis, A. Pourtsidou, and E. J. Copeland. "Parametrized post-Friedmannian framework for interacting dark energy theories". In: *Phys. Rev. D* 91.8 (2015), p. 083537. DOI: 10.1103/PhysRevD.91.083537. arXiv: 1502.07297 [astro-ph.CO].

[555] Anze Slosar, Xingang Chen, Cora Dvorkin, Daan Meerburg, Benjamin Wallisch, Daniel Green, and Eva Silverstein. "Scratches from the Past: Inflationary Archaeology through Features in the Power Spectrum of Primordial Fluctuations". In: *Bull. Am. Astron. Soc.* 51.3, 98 (May 2019), p. 98. arXiv: 1903.09883 [astro-ph.CO].

[556] Anze Slosar, Christopher Hirata, Uros Seljak, Shirley Ho, and Nikhil Padmanabhan. "Constraints on local primordial non-Gaussianity from large scale structure". In: *JCAP* 08 (2008), p. 031. DOI: 10.1088/1475-7516/2008/08/031. arXiv: 0805.3580 [astro-ph].

[557] Anze Slosar, Rachel Mandelbaum, and Daniel Eisenstein. "Dark Energy and Modified Gravity". In: *Bull. Am. Astron. Soc.* 51.3, 97 (May 2019), p. 97. arXiv: 1903.12016 [astro-ph.CO].

[558] Anze Slosar et al. "Packed Ultra-wideband Mapping Array (PUMA): A Radio Telescope for Cosmology and Transients". In: *Bull. Am. Astron. Soc.* Vol. 51. Sept. 2019, 53, p. 53. arXiv: 1907.12559 [astro-ph.IM].

[559] Kendrick M. Smith and Simone Ferraro. "Detecting Patchy Reionization in the Cosmic Microwave Background". In: *Phys. Rev. Lett.* 119.2 (2017), p. 021301. DOI: 10.1103/PhysRevLett.119.021301. arXiv: 1607.01769 [astro-ph.CO].

[560] Tristan L. Smith, Vivian Poulin, José Luis Bernal, Kimberly K. Boddy, Marc Kamionkowski, and Riccardo Murgia. *Early dark energy is not excluded by current large-scale structure data*. 2020. arXiv: 2009.10740 [astro-ph.CO].

[561] Emanuele Sobacchi and Andrei Mesinger. "The clustering of Lyman α emitters at z $\approx$ 7: implications for reionization and host halo masses". In: *Mon. Not. Roy. Astron. Soc.* 453.2 (2015), pp. 1843–1854. DOI: 10.1093/mnras/stv1751. arXiv: 1505.02787 [astro-ph.CO].

[562] B. Soergel et al. "Detection of the kinematic Sunyaev-Zel'dovich effect with DES Year 1 and SPT". In: *Mon. Not. R. Astron. Soc.* 461.3 (Sept. 2016), pp. 3172–3193. DOI: 10.1093/mnras/stw1455. arXiv: 1603.03904 [astro-ph.CO].

[563] D. Spergel et al. "Wide-Field InfrarRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report". In: *arXiv e-prints*, arXiv:1503.03757 (Mar. 2015), arXiv:1503.03757. arXiv: 1503.03757 [astro-ph.IM].

[564] Volker Springel. "The Cosmological simulation code GADGET-2". In: *Mon. Not. Roy. Astron. Soc.* 364 (2005), pp. 1105–1134. DOI: 10.1111/j.1365-2966.2005.09655.x. arXiv: astro-ph/0505010.

[565] Benjamin E. Stahl, Thomas de Jaeger, Supranta S. Boruah, WeiKang Zheng, Alexei V. Filippenko, and Michael J. Hudson. "Peculiar-velocity cosmology with Types Ia and II supernovae". In: *Mon. Not. Roy. Astron. Soc.* 505.2 (2021), pp. 2349–2360. DOI: 10.1093/mnras/stab1446. arXiv: 2105.05185 [astro-ph.CO].

[566] George Stein, Marcelo A. Alvarez, J. Richard Bond, Alexander van Engelen, and Nicholas Battaglia. "The Websky Extragalactic CMB Simulations". In: *JCAP* 10 (2020), p. 012. DOI: 10.1088/1475-7516/2020/10/012. arXiv: 2001.08787 [astro-ph.CO].

[567] Benjamin Stölzner et al. "KiDS-Legacy: Consistency of cosmic shear measurements and joint cosmological constraints with external probes". In: (Mar. 2025). arXiv: 2503.19442 [astro-ph.CO].

[568] Meng Su, Amit P. S. Yadav, Matthew McQuinn, Jaiyul Yoo, and Matias Zaldarriaga. "An Improved Forecast of Patchy Reionization Reconstruction with CMB". In: *arXiv e-prints*, arXiv:1106.4313 (June 2011), arXiv:1106.4313. DOI: 10.48550/arXiv.1106.4313. arXiv: 1106.4313 [astro-ph.CO].

[569] H. Sugai et al. "Updated Design of the CMB Polarization Experiment Satellite LiteBIRD". In: *Journal of Low Temperature Physics* 199.3-4 (Jan. 2020), pp. 1107–1117. ISSN: 1573-7357. DOI: 10.1007/s10909-019-02329-w. URL: http://dx.doi.org/10.1007/s10909-019-02329-w.

[570] Naonori S. Sugiyama, Teppei Okumura, and David N. Spergel. "A direct measure of free electron gas via the kinematic Sunyaev-Zel'dovich effect in Fourier-space analysis". In: *Mon. Not. R. Astron. Soc.* 475.3 (Apr. 2018), pp. 3764–3785. DOI: 10.1093/mnras/stx3362. arXiv: 1705.07449 [astro-ph.CO].

[571] Sunao Sugiyama et al. "Hyper Suprime-Cam Year 3 results: Cosmology from galaxy clustering and weak lensing with HSC and SDSS using the minimal bias model". In: *Phys. Rev. D* 108.12 (2023), p. 123521. DOI: 10.1103/PhysRevD.108.123521. arXiv: 2304.00705 [astro-ph.CO].

[572] R. A. Sunyaev and Ia. B. Zeldovich. "Intergalactic Gas in Clusters of Galaxies, the Microwave Background, and Cosmology". In: *Astrophys. Space Phys. Res.* 1 (Jan. 1981), p. 1.

[573] R. A. Sunyaev and Ya. B. Zeldovich. "The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies". In: *Comments on Astrophysics and Space Physics* 4 (Nov. 1972), p. 173.

[574] R. A. Sunyaev and Ya. B. Zeldovich. "The velocity of clusters of galaxies relative to the microwave background - The possibility of its measurement." In: *Mon. Not. R. Astron. Soc.* 190 (Feb. 1980), pp. 413–420. DOI: 10.1093/mnras/190.3.413.

[575] Masahiro Takada and Olivier Dore. "Geometrical Constraint on Curvature with BAO experiments". In: *Phys. Rev. D* 92.12 (2015), p. 123518. DOI: 10.1103/PhysRevD.92.123518. arXiv: 1508.02469 [astro-ph.CO].

[576] R. Takahashi, M. Sato, T. Nishimichi, A. Taruya, and M. Oguri. "Revising the Halofit Model for the Nonlinear Matter Power Spectrum". In: *Astrophys. J.* 761, 152 (Dec. 2012), p. 152. DOI: 10.1088/0004-637X/761/2/152. arXiv: 1208.2701.

[577] Mengtao Tang, Daniel P. Stark, Michael W. Topping, Charlotte Mason, and Richard S. Ellis. "JWST/NIRSpec Observations of Lyman α Emission in Star-forming Galaxies at $6.5 \lesssim z \lesssim 13$". In: *Astrophys. J.* 975.2, 208 (Nov. 2024), p. 208. DOI: 10.3847/1538-4357/ad7eb7. arXiv: 2408.01507 [astro-ph.GA].

[578] Hideki Tanimura, Saleem Zaroubi, and Nabila Aghanim. "Direct detection of the kinetic Sunyaev-Zel'dovich effect in galaxy clusters". In: *Astron. Astrophys.* 645, A112 (Jan. 2021), A112. DOI: 10.1051/0004-6361/202038846. arXiv: 2007.02952 [astro-ph.CO].

[579] S. Tassev and M. Zaldarriaga. "Towards an optimal reconstruction of baryon oscillations". In: *J. Cosm. Astrop. Phys.* 10, 006 (Oct. 2012), p. 006. DOI: 10.1088/1475-7516/2012/10/006. arXiv: 1203.6066.

[580] Brook Taylor. *Methodus Incrementorum Directa et Inversa.* Latin. London: Typis Gul. Innys, 1715.

[581] Max Tegmark, Angélica de Oliveira-Costa, and Andrew J. S. Hamilton. "High resolution foreground cleaned CMB map from WMAP". In: *Physical Review D* 68.12 (Dec. 2003). ISSN: 1089-4918. DOI: 10.1103/physrevd.68.123523. URL: http://dx.doi.org/10.1103/PhysRevD.68.123523.

[582] The CMB-HD Collaboration et al. "Snowmass2021 CMB-HD White Paper". In: *arXiv e-prints*, arXiv:2203.05728 (Mar. 2022), arXiv:2203.05728. DOI: 10.48550/arXiv.2203.05728. arXiv: 2203.05728 [astro-ph.CO].

[583] The MSE Science Team et al. "The Detailed Science Case for the Maunakea Spectroscopic Explorer, 2019 edition". In: *arXiv e-prints*, arXiv:1904.04907 (Apr. 2019), arXiv:1904.04907. arXiv: 1904.04907 [astro-ph.IM].

[584] A. R. Thompson, J. M. Moran, and G. W. Swenson Jr. *Interferometry and Synthesis in Radio Astronomy, 3rd Edition*. 2017.

[585] R. J. Thornton et al. "The Atacama Cosmology Telescope: The Polarization-sensitive ACTPol Instrument". In: *Astrophys. J. Supp.* 227.2, 21 (Dec. 2016), p. 21. DOI: 10.3847/1538-4365/227/2/21. arXiv: 1605.06569 [astro-ph.IM].

[586] Jesus Torrado and Antony Lewis. "Cobaya: Code for Bayesian Analysis of hierarchical physical models". In: *JCAP* 05 (2021), p. 057. DOI: 10.1088/1475-7516/2021/05/057. arXiv: 2005.05290 [astro-ph.IM].

[587] Jesús Torrado and Antony Lewis. *Cobaya: Bayesian analysis in cosmology*. Astrophysics Source Code Library, record ascl:1910.019. Oct. 2019.

[588] Hy Trac. "Parametrizing the Reionization History with the Redshift Midpoint, Duration, and Asymmetry". In: *Astrophys. J. Lett.* 858.2 (2018), p. L11. DOI: 10.3847/2041-8213/aabff0. arXiv: 1804.00672 [astro-ph.CO].

[589] Hy Trac and Renyue Cen. "Radiative transfer simulations of cosmic reionization. 1. Methodology and initial results". In: *Astrophys. J.* 671 (2007), p. 1. DOI: 10.1086/522566. arXiv: astro-ph/0612406.

[590] Hy Trac, Nianyi Chen, Ian Holst, Marcelo A. Alvarez, and Renyue Cen. "AMBER: A Semi-numerical Abundance Matching Box for the Epoch of Reionization". In: *Astrophys. J.* 927.2 (2022), p. 186. DOI: 10.3847/1538-4357/ac5116. arXiv: 2109.10375 [astro-ph.CO].

[591] Hy Trac and Ue-Li Pen. "Out - of - core hydrodynamic simulations for cosmological applications". In: *New Astron.* 11 (2006), pp. 273–286. DOI: 10.1016/j.newast.2005.08.003. arXiv: astro-ph/0402444.

[592] M. Tristram et al. "Cosmological parameters derived from the final (PR4) Planck data release". In: *Astron. Astrophys.* 682 (2024), A37. DOI: 10.1051/0004-6361/202348015. arXiv: 2309.10034 [astro-ph.CO].

[593] M. Tristram et al. "Planck constraints on the tensor-to-scalar ratio". In: *Astron. Astrophys.* 647, A128 (Mar. 2021), A128. DOI: 10.1051/0004-6361/202039585. arXiv: 2010.01139 [astro-ph.CO].

[594] Tilman Tröster et al. "Cosmology from large-scale structure: Constraining ΛCDM with BOSS". In: *Astron. Astrophys.* 633 (2020), p. L10. DOI: 10.1051/0004-6361/201936772. arXiv: 1909.11006 [astro-ph.CO].

[595] Dmitriy Tseliakhovich and Christopher Hirata. "Relative velocity of dark matter and baryonic fluids and the formation of the first structures". In: *Phys. Rev. D* 82.8, 083520 (Oct. 2010), p. 083520. DOI: 10.1103/PhysRevD.82.083520. arXiv: 1005.2416 [astro-ph.CO].

[596] Sunny Vagnozzi, Thejs Brinckmann, Maria Archidiacono, Katherine Freese, Martina Gerbino, Julien Lesgourgues, and Tim Sprenger. "Bias due to neutrinos must not uncorrect'd go". In: *JCAP* 09 (2018), p. 001. DOI: 10.1088/1475-7516/2018/09/001. arXiv: 1807.04672 [astro-ph.CO].

[597] A. van Engelen, S. Bhattacharya, N. Sehgal, G. P. Holder, O. Zahn, and D. Nagai. "CMB Lensing Power Spectrum Biases from Galaxies and Clusters Using High-angular Resolution Temperature Maps". In: *Astrophys. J.* 786.1, 13 (May 2014), p. 13. DOI: 10.1088/0004-637X/786/1/13. arXiv: 1310.7023 [astro-ph.CO].

[598] Keith Vanderlinde et al. "The Canadian Hydrogen Observatory and Radio-transient Detector (CHORD)". In: *Canadian Long Range Plan for Astronomy and Astrophysics White Papers*. Vol. 2020. Oct. 2019, 28, p. 28. DOI: 10.5281/zenodo.3765414. arXiv: 1911.01777 [astro-ph.IM].

[599] A. Verhamme et al. "Recovering the systemic redshift of galaxies from their Lyman alpha line profile". In: *Mon. Not. R. Astron. Soc.* 478.1 (July 2018), pp. L60–L65. DOI: 10.1093/mnrasl/sly058. arXiv: 1804.01883 [astro-ph.GA].

[600] Francisco Villaescusa-Navarro, Arka Banerjee, Neal Dalal, Emanuele Castorina, Roman Scoccimarro, Raul Angulo, and David N. Spergel. "The imprint of neutrinos on clustering in redshift-space". In: *Astrophys. J.* 861.1 (2018), p. 53. DOI: 10.3847/1538-4357/aac6bf. arXiv: 1708.01154 [astro-ph.CO].

[601] Francisco Villaescusa-Navarro, Federico Marulli, Matteo Viel, Enzo Branchini, Emanuele Castorina, Emiliano Sefusatti, and Shun Saito. "Cosmology with massive neutrinos I: towards a realistic modeling of the relation between matter, haloes and galaxies". In: *JCAP* 03 (2014), p. 011. DOI: 10.1088/1475-7516/2014/03/011. arXiv: 1311.0866 [astro-ph.CO].

[602] Francisco Villaescusa-Navarro et al. "Ingredients for 21 cm Intensity Mapping". In: *Astrophys. J.* 866.2, 135 (Oct. 2018), p. 135. DOI: 10.3847/1538-4357/aadba0. arXiv: 1804.09180 [astro-ph.CO].

[603] Jens Verner Villumsen. "Clustering of faint galaxies: omega (Theta), induced by weak gravitational lensing". In: (Dec. 1995). arXiv: astro-ph/9512001.

[604] Pauli Virtanen et al. "SciPy 1.0: fundamental algorithms for scientific computing in Python". In: *Nat. Methods* 17.3 (Mar. 2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.

[605] Ethan T. Vishniac. "Reionization and Small-Scale Fluctuations in the Microwave Background". In: *Astrophys. J.* 322 (Nov. 1987), p. 597. DOI: 10.1086/165755.

[606] Z. Vlah, E. Castorina, and M. White. "The Gaussian streaming model and convolution Lagrangian effective field theory". In: *J. Cosm. Astrop. Phys.* 12, 007 (Dec. 2016), p. 007. DOI: 10.1088/1475-7516/2016/12/007. arXiv: 1609.02908.

[607] Z. Vlah, M. White, and A. Aviles. "A Lagrangian effective field theory". In: *J. Cosm. Astrop. Phys.* 9, 014 (Sept. 2015), p. 014. DOI: 10.1088/1475-7516/2015/09/014. arXiv: 1506.05264.

[608] Zvonimir Vlah, Nora Elisa Chisari, and Fabian Schmidt. "An EFT description of galaxy intrinsic alignments". In: *JCAP* 01 (2020), p. 025. DOI: 10.1088/1475-7516/2020/01/025. arXiv: 1910.08085 [astro-ph.CO].

[609] Zvonimir Vlah, Uroš Seljak, Man Yat Chu, and Yu Feng. "Perturbation theory, effective field theory, and oscillations in the power spectrum". In: *J. Cosm. Astrop. Phys.* 2016.3, 057 (Mar. 2016), p. 057. DOI: 10.1088/1475-7516/2016/03/057. arXiv: 1509.02120 [astro-ph.CO].

[610] Zvonimir Vlah and Martin White. "Exploring redshift-space distortions in large-scale structure". In: *J. Cosm. Astrop. Phys.* 2019.3, 007 (Mar. 2019), p. 007. DOI: 10.1088/1475-7516/2019/03/007. arXiv: 1812.02775 [astro-ph.CO].

[611] Digvijay Wadekar, Mikhail M. Ivanov, and Roman Scoccimarro. "Cosmological constraints from BOSS with analytic covariance matrices". In: *Phys. Rev. D* 102.12, 123521 (Dec. 2020), p. 123521. DOI: 10.1103/PhysRevD.102.123521. arXiv: 2009.00622 [astro-ph.CO].

[612] Terry P. Walker, Gary Steigman, David N. Schramm, Keith A. Olive, and Ho-Shik Kang. "Primordial Nucleosynthesis Redux". In: *Astrophys. J.* 376 (July 1991), p. 51. DOI: 10.1086/170255.

[613] Feige Wang et al. "A Significantly Neutral Intergalactic Medium Around the Luminous z = 7 Quasar J0252–0503". In: *Astrophys. J.* 896.1 (2020), p. 23. DOI: 10.3847/1538-4357/ab8c45. arXiv: 2004.10877 [astro-ph.GA].

[614] D. J. Watts, G. E. Addison, C. L. Bennett, and J. L. Weiland. "Beyond optical depth: Future determination of ionization history from the CMB". In: *arxiv e-prints* (Oct. 2019). DOI: 10.3847/1538-4357/ab5fd5. arXiv: 1910.00590 [astro-ph.CO].

[615] Noah Weaverdyck and Dragan Huterer. "Mitigating contamination in LSS surveys: a comparison of methods". In: *Mon. Not. Roy. Astron. Soc.* 503.4 (2021), pp. 5061–5084. DOI: 10.1093/mnras/stab709. arXiv: 2007.14499 [astro-ph.CO].

[616] Risa H. Wechsler, Joseph DeRose, Michael T. Busha, Matthew R. Becker, Eli Rykoff, and August Evrard. "ADDGALS: Simulated Sky Catalogs for Wide Field Galaxy Surveys". In: *Astrophys. J.* 931.2 (2022), p. 145. DOI: 10.3847/1538-4357/ac5b0a. arXiv: 2105.12105 [astro-ph.CO].

[617] D. H. Weinberg, M. J. Mortonson, D. J. Eisenstein, C. Hirata, A. G. Riess, and E. Rozo. "Observational probes of cosmic acceleration". In: *Phys. Rep.* 530 (Sept. 2013), pp. 87–255. DOI: 10.1016/j.physrep.2013.05.001. arXiv: 1201.2434.

[618] Steven Weinberg. *Cosmology*. Oxford, UK: Oxford University Press, 2008. ISBN: 978-0-19-852682-7.

[619] Rainer Weinberger et al. "Simulating galaxy formation with black hole driven thermal and kinetic feedback". In: *Mon. Not. R. Astron. Soc.* 465.3 (Mar. 2017), pp. 3291–3308. DOI: 10.1093/mnras/stw2944. arXiv: 1607.03486 [astro-ph.GA].

[620] Lukas Wenzl et al. "Atacama Cosmology Telescope: DR6 gravitational lensing and SDSS BOSS cross-correlation measurement and constraints on gravity with the EG statistic". In: *Phys. Rev. D* 111.4 (2025), p. 043535. DOI: 10.1103/PhysRevD.111.043535. arXiv: 2405.12795 [astro-ph.CO].

[621] Lukas Wenzl et al. "Constraining gravity with a new precision E_G estimator using Planck + SDSS BOSS data". In: *Phys. Rev. D* 109.8, 083540 (Apr. 2024), p. 083540. DOI: 10.1103/PhysRevD.109.083540. arXiv: 2401.12971 [astro-ph.CO].

[622] M. White. "Reconstruction within the Zeldovich approximation". In: *Mon. Not. R. Astron. Soc.* 450 (July 2015), pp. 3822–3828. DOI: 10.1093/mnras/stv842. arXiv: 1504.03677.

[623] M. White. "The Zel'dovich approximation". In: *Mon. Not. R. Astron. Soc.* 439 (Apr. 2014), pp. 3630–3640. DOI: 10.1093/mnras/stu209. arXiv: 1401.5466.

[624] M. White, B. Reid, C.-H. Chuang, J. L. Tinker, C. K. McBride, F. Prada, and L. Samushia. "Tests of redshift-space distortions models in configuration space for the analysis of the BOSS final data release". In: *Mon. Not. R. Astron. Soc.* 447 (Feb. 2015), pp. 234–245. DOI: 10.1093/mnras/stu2460. arXiv: 1408.5435.

[625] Martin White et al. "Cosmological constraints from the tomographic cross-correlation of DESI Luminous Red Galaxies and Planck CMB lensing". In: *JCAP* 02.02 (2022), p. 007. DOI: 10.1088/1475-7516/2022/02/007. arXiv: 2111.09898 [astro-ph.CO].

[626] Lily R. Whitler, Charlotte A. Mason, Keven Ren, Mark Dijkstra, Andrei Mesinger, Laura Pentericci, Michele Trenti, and Tommaso Treu. "The impact of scatter in the galaxy UV luminosity to halo mass relation on Ly α visibility during the epoch of reionization". In: *Mon. Not. Roy. Astron. Soc.* 495.4 (2020), pp. 3602–3613. DOI: 10.1093/mnras/staa1178. arXiv: 1911.03499 [astro-ph.CO].

[627] Clifford M. Will. "The Confrontation between General Relativity and Experiment". In: *Living Reviews in Relativity* 17.1, 4 (Dec. 2014), p. 4. DOI: 10.12942/lrr-2014-4. arXiv: 1403.7377 [gr-qc].

[628] M. J. Wilson and Martin White. "Cosmology with dropout selection: straw-man surveys & CMB lensing". In: *J. Cosm. Astrop. Phys.* 2019.10, 015 (Oct. 2019), p. 015. DOI: 10.1088/1475-7516/2019/10/015. arXiv: 1904.13378 [astro-ph.CO].

[629] Edward L. Wright et al. "The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance". In: *Astron. J.* 140.6 (Dec. 2010), pp. 1868–1881. DOI: 10.1088/0004-6256/140/6/1868. arXiv: 1008.0031 [astro-ph.IM].

[630] Jiachuan Xu et al. "Constraining baryonic physics with DES Y1 and Planck data: Combining galaxy clustering, weak lensing, and CMB lensing". In: *Phys. Rev. D* 110.6, 063532 (Sept. 2024), p. 063532. DOI: 10.1103/PhysRevD.110.063532. arXiv: 2311.08047 [astro-ph.CO].

[631] Kun Xu, Y. P. Jing, Hongyu Gao, Xiaolin Luo, and Ming Li. "Accurate Measurement of the Lensing Magnification by BOSS CMASS Galaxies and Its Implications for Cosmology and Dark Matter". In: *Astrophys. J.* 973.2 (2024), p. 102. DOI: 10.3847/1538-4357/ad6156. arXiv: 2405.16484 [astro-ph.CO].

[632] Wenxiao Xu, Yidong Xu, Bin Yue, Ilian T. Iliev, Hy Trac, Liang Gao, and Xuelei Chen. "The HI bias during the Epoch of Reionization". In: *Mon. Not. Roy. Astron. Soc.* 490.4 (2019), pp. 5739–5748. DOI: 10.1093/mnras/stz2926. arXiv: 1907.05098 [astro-ph.CO].

[633] Jaiyul Yoo, Neal Dalal, and Uroš Seljak. "Supersonic relative velocity effect on the baryonic acoustic oscillation measurements". In: *J. Cosm. Astrop. Phys.* 2011.7, 018 (July 2011), p. 018. DOI: 10.1088/1475-7516/2011/07/018. arXiv: 1105.3732 [astro-ph.CO].

[634] Byeonghee Yu, Robert Z. Knight, Blake D. Sherwin, Simone Ferraro, Lloyd Knox, and Marcel Schmittfull. "Towards Neutrino Mass from Cosmology without Optical Depth Information". In: (Sept. 2018). arXiv: 1809.02120 [astro-ph.CO].

[635] Sihan Yuan et al. "The DESI one-per cent survey: exploring the halo occupation distribution of luminous red galaxies and quasi-stellar objects with ABACUSSUMMIT". In: *Mon. Not. R. Astron. Soc.* 530.1 (May 2024), pp. 947–965. DOI: 10.1093/mnras/stae359. arXiv: 2306.06314 [astro-ph.CO].

[636] Matias Zaldarriaga and Uros Seljak. "An all sky analysis of polarization in the microwave background". In: *Phys. Rev. D* 55 (1997), pp. 1830–1840. DOI: 10.1103/PhysRevD.55.1830. arXiv: astro-ph/9609170.

[637] Matteo Zennaro, Raul E. Angulo, Sergio Contreras, Marcos Pellejero-Ibáñez, and Francisco Maion. "Priors on Lagrangian bias parameters from galaxy formation modelling". In: *Mon. Not. Roy. Astron. Soc.* 514.4 (2022), pp. 5443–5456. DOI: 10.1093/mnras/stac1673. arXiv: 2110.05408 [astro-ph.CO].

[638] A. Zenteno et al. "A joint SZ-X-ray-optical analysis of the dynamical state of 288 massive galaxy clusters". In: *Mon. Not. R. Astron. Soc.* 495.1 (June 2020), pp. 705–725. DOI: 10.1093/mnras/staa1157. arXiv: 2004.01721 [astro-ph.GA].

[639] Pierre Zhang, Guido D'Amico, Leonardo Senatore, Cheng Zhao, and Yifu Cai. "BOSS Correlation Function analysis from the Effective Field Theory of Large-Scale Structure". In: *JCAP* 02.02 (2022), p. 036. DOI: 10.1088/1475-7516/2022/02/036. arXiv: 2110.07539 [astro-ph.CO].

[640] Rongpu Zhou et al. "Target Selection and Validation of DESI Luminous Red Galaxies". In: *Astron. J.* 165.2 (2023), p. 58. DOI: 10.3847/1538-3881/aca5fb. arXiv: 2208.08515 [astro-ph.CO].

[641] Rongpu Zhou et al. "The Clustering of DESI-like Luminous Red Galaxies Using Photometric Redshifts". In: *Mon. Not. Roy. Astron. Soc.* 501.3 (2021), pp. 3309–3331. DOI: 10.1093/mnras/staa3764. arXiv: 2001.06018 [astro-ph.CO].

[642] Rongpu Zhou, Arjun Dey, Dustin Lang, John Moustakas, Edward F. Schlafly, and David J. Schlegel. "Comparing the Photometric Calibration of DESI Imaging and Gaia Synthetic Photometry". In: *Research Notes of the American Astronomical Society* 7.5, 105 (May 2023), p. 105. DOI: 10.3847/2515-5172/acd7ef. arXiv: 2305.16630 [astro-ph.IM].

[643] Rongpu Zhou et al. "DESI luminous red galaxy samples for cross-correlations". In: *J. Cosm. Astrop. Phys.* 2023.11, 097 (Nov. 2023), p. 097. DOI: 10.1088/1475-7516/2023/11/097. arXiv: 2309.06443 [astro-ph.CO].

[644] Yongda Zhu et al. "Long Dark Gaps in the Lyβ Forest at $z < 6$: Evidence of Ultra Late Reionization from XQR-30 Spectra". In: (May 2022). arXiv: 2205.04569 [astro-ph.CO].

[645] Hu Zou et al. "Project Overview of the Beijing-Arizona Sky Survey". In: *Publ. Astron. Soc. Pacific* 129.976 (June 2017), p. 064101. DOI: 10.1088/1538-3873/aa65ba. arXiv: 1702.03653 [astro-ph.GA].

[646] P. A. Zyla et al. "Review of Particle Physics". In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.