CMB LENSING MEASUREMENTS WITH TWO YEARS OF DATA FROM THE SPT-3G SURVEY

BY

CAIL DALEY

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Astronomy
in the Graduate College of the
University of Illinois Urbana-Champaign, 2024

Urbana, Illinois

Doctoral Committee:

Professor Gil Holder, Chair and Director of Research
Professor Brian Fields
Associate Professor Jeffrey Filippini
Professor Paul Ricker

Abstract

The South Pole Telescope (SPT) is a 10 meter offset Gregorian telescope designed to observe the Cosmic Microwave Background (CMB) and astrophysical foregrounds at millimeter wavelengths. Radiation from the CMB is gravitationally lensed by large scale structure as it travels through the cosmos, complicating the analysis of the primordial CMB but also embedding information about the evolution and structure of the universe at more recent times. Measurements of the lensing potential can be used to constrain cosmological parameters (e.g. H_0, Ω_m, σ_8, and the sum of neutrino masses) and to remove lensing contamination from CMB B-mode polarization maps in the search for primordial gravitational waves. The SPT-3G experiment surveys $1500\,\mathrm{deg}^2$ of the southern sky with $\sim 1'$ angular resolution, allowing for high signal-to-noise measurements of the lensing potential out to sub-degree scales that can probe the nonlinear regime of the lensing power spectrum.

In this thesis I present CMB lensing measurements made with data from the 2019 and 2020 observing seasons of the SPT-3G experiment. With per-pixel noise levels that are a factor of 3-4 lower than recent *Planck* and Atacama Cosmology Telescope lensing measurements, the SPT-3G lensing maps are the deepest CMB lensing maps ever made. While the analysis is ongoing, preliminary results indicate that we will measure the lensing power spectrum with a combined signal-to-noise of nearly 40σ, resulting in a 2-3% constraint on the amplitude of the lensing power spectrum A_ϕ. We expect to measure the Hubble constant H_0 to similar precision as other state of the art CMB lensing measurements, and expect to place the tightest CMB lensing constraints to date on the structure growth parameter S_8. The SPT-3G

lensing maps will also be used to delens CMB polarization maps in the search for primordial gravitational waves in collaboration with BICEP/*Keck*, and will be excellent for cross-correlations especially given the high signal to noise of the polarization-based reconstruction that is robust to foreground contamination.

To Micha.

Acknowledgments

I would like to begin by expressing my deepest gratitude to my advisor, Gil Holder, for encouraging me, challenging me, and teaching me so much about cosmology. Gil has shown me a playful approach to science that still demands the highest standards of work, and I can only hope to balance these qualities in my own path through life.

I am grateful to my committee members, Profs. Brian Fields, Jeff Filippini, and Paul Ricker, for their guidance and support throughout my time at the University of Illinois. While I never had the pleasure of taking a class with Jeff, Brian and Paul taught my favorite classes in graduate school I was lucky to have had such dedicated instructors. I would also like to thank Prof. Leslie Looney for welcoming me into his home for Thanksgiving dinner when I was a first-year graduate student, and for his patient support as department chair.

Joining the South Pole Telescope (SPT) collaboration was the best decision I made in graduate school, and there are so many people within the collaboration that I would like to thank. My faculty mentor Tom Crawford has been an incredible source of wisdom on both life and signal processing, and I have deep respect for his dedication to supporting graduate students in SPT. I would also like to thank Tom and Wai Ling Kimmy Wu for their support in my postdoctoral job search, especially for writing me letters of recommendation. I am very grateful to Kimmy, Lloyd Knox, and Marius Millea for hosting me at their home institutions and giving me opportunities to share my work in the summer of 2022.

It has been a great honor to work with and befriend so many postdocs and graduate students in the SPT. I don't think I'll ever be able to express the depths of my gratitude to Srinivasan Raghunathan and Yuuki Omori, who have become informal mentors over the past few years and have shown me incredible kindness and patience. I will never forget the time they devoted to helping me with both science and big life decisions. Sasha Rahlin was always available to help me with software-related questions, and I learned a great deal while working with her to improve the SPT-3G software library. I have had also had the pleasure of working closely with many wonderful graduate students (some now postdocs) in SPT, including Allen Foster, Chris Tandoi, Chunyu Lu, Etienne Camphuis, Kedar Phadke, Melanie Archipley, Sam Guns, Wei Quan, and Zhaodi Pan.

I would especially like to thank Chris Tandoi and Melanie Archipley for their camaraderie and spirited debates both within SPT and at the University of Illinois; they have been some of my closest friends in graduate school. And I am lucky to have been part of such a great cohort of graduate students at the University of Illinois: Alex Gagliano, Amelia Mangian, Cassidy Wagner, Colin Burke, Jianyang Fu, Melanie Archipley, and Patrick Aleo. I will always treasure the great times we had in the first few years of graduate school.

I would like to thank all my housemates at Vinehouse over the years; I couldn't have asked for a better group of people to share my time in Urbana with.

I would like to thank my parents Dan and Lynn, and my siblings David, Kristen, and Rae, for their love and and support even if they don't understand what I do all day.

And finally, more than anyone else I would like thank my partner Micha, who has undoubtedly been the greatest gift of the last six years. I would not be half the person I am today without her unwavering support and unconditional love, and I cannot wait to keep growing together in the next chapter of our lives.

This work was funded by a South Pole Telescope grant from the National Science Foundation with award number AWD0100221.

I would like to thank Brand and Brand and Monica Fortner for their generous financial support of my graduate education and research.

This work was partially supported by the Center for AstroPhysical Surveys (CAPS) at the National Center for Supercomputing Applications (NCSA), University of Illinois Urbana-Champaign.

Some of this work was done using services provided by the OSG Consortium (Pordes et al., 2007; Sfiligoi et al., 2009; OSG, 2006, 2015), which is supported by the National Science Foundation awards #2030508 and #1836650.

Portions of Chapter 3 were published in the Astrophysical Journal with the title "Optimal Cosmic Microwave Background Lensing Reconstruction and Parameter Estimation with SPTpol Data" (Millea et al., 2021). I was second author on this work, and all the text in this chapter is my own.

Table of contents

List of Figures . . . . **ix**

Chapter 1 Introduction **1**
- 1.1 The Cosmic Microwave Background 2
- 1.2 The ΛCDM Model & Cosmological Parameters 6
- 1.3 The South Pole Telescope 8
- 1.4 The Flat-Sky Approximation and Map Projections 15
- 1.5 Thesis Overview 21

Chapter 2 Gravitational Lensing of the CMB **23**
- 2.1 An Introduction to Weak Lensing 24
- 2.2 Weak Lensing of the CMB 28
- 2.3 CMB Lensing Science Cases 31
- 2.4 Reconstructing the Lensing Potential 34

Chapter 3 SPTpol Lensing **51**

Chapter 4 SPT-3G Lensing **57**
- 4.1 Analysis Overview and Motivation 59
- 4.2 Data 63
- 4.3 Simulations 74
- 4.4 Map Processing 78
- 4.5 Lensing Map Reconstruction 87
- 4.6 Power Spectrum Estimation 93
- 4.7 Null Tests and Consistency Checks 100
- 4.8 Obfuscated Data Bandpowers 110
- 4.9 Discussion and Future Work 111

Chapter 5 Summary and Conclusions **117**

Bibliography **121**

List of Figures

1.1 Temperature anisotropies of the cosmic microwave background as observed by the *Planck* satellite. 3
1.2 The South Pole and BICEP Telescopes at the Amundsen-Scott Station, with the Milky Way and Aurora Australis visible in the background. 9
1.3 The survey footprints of the SPT fields used in this thesis. 10
1.4 The SPT-3G footprint in two different flat-sky projections. 17

2.1 A schematic of a weak gravitational lensing scenario. 26

3.1 SPTpol CMB polarization and QE lensing convergence maps from the SPTpol $100\,\mathrm{deg}^2$ ultradeep analysis. 52
3.2 Ingredients of the inverse variance filter used in the SPTpol $100\,\mathrm{deg}^2$ ultradeep analysis. 53
3.3 Components of the SPTpol $100\,\mathrm{deg}^2$ ultradeep lensing power spectrum pipeline and final debiased bandpowers. 55
3.4 Comparison of A_ϕ constraints from the Bayesian and QE pipelines the SPTpol $100\,\mathrm{deg}^2$ ultradeep analysis. 56

4.1 Forecasted constraints on ΛCDM parameters from the full-depth SPT-3G survey. 58
4.2 Lensing reconstruction noise levels from the SPT-3G 2019-2020 dataset. . . 60
4.3 Example of point source filtering wings. 65
4.4 SPT-3G beams. 67
4.5 Cross-spectrum transfer function and ratio between cross- and auto-spectrum transfer functions for the SPT-3G 2019+2020 flat-sky lensing analysis. . . . 79
4.6 Minimum variance linear frequency combination weights for the SPT-3G 2019+2020 flat-sky lensing analysis. 80
4.7 Transfer function-deconvolved noise power spectra for the SPT-3G 2019+2020 flat-sky lensing analysis. 82
4.8 Combined power spectra of the individual frequency maps and their minimum variance linear combination for the SPT-3G 2019+2020 flat-sky lensing analysis. 83
4.9 Downsampled and inpainted SPT-3G 2019+2020 T, Q, and U maps. 85
4.10 Wiener filtered SPT-3G 2019+2020 T, E, and B maps. 86
4.11 Power spectra of the SPT-3G 2019+2020 downsampled and Wiener filtered maps. 87

4.12 Convergence maps for the SPT-3G 2019+2020 flat-sky lensing analysis. . . 89
4.13 Mean-fields for the SPT-3G 2019+2020 flat-sky lensing analysis. 91
4.14 Lensing map noise and MC response correction for the SPT-3G 2019+2020 flat-sky lensing analysis. 92
4.15 Cross-correlation coefficients for the SPT-3G 2019+2020 flat-sky lensing analysis. 94
4.16 Minimum variance simulation bandpowers for the 2019+2020 SPT-3G lensing analysis. 97
4.17 Demonstration of the improved off-diagonal $\mathbb{N}_{\mathbf{L}}^{0,\mathrm{RDSA}}$ bandpower covariance. 98
4.18 Lensed simulation pipeline test for the SPT-3G 2019+2020 flat-sky lensing analysis. 102
4.19 Unlensed simulation pipeline test for the SPT-3G 2019+2020 flat-sky lensing analysis. 106
4.20 Characterization of foreground biases in the SPT-3G 2019+2020 flat-sky lensing analysis with the Agora simulation. 108
4.21 Obfuscated data bandpowers for the SPT-3G 2019+2020 flat-sky lensing analysis. 111
4.22 Forecasted parameter constraints from the 2019+2020 SPT-3G lensing analyses compared to other early-universe probes. 112
4.23 Forecasted parameter constraints and degeneracies for the 2019+2020 SPT-3G lensing analyses. 114
4.24 SPT-3G 2019+2020 and *Planck* lensing maps stacked on regions of varying quasar density. 115

Chapter 1

Introduction

The Cosmic Microwave Background (CMB) is an incredibly rich cosmological dataset, the earliest and most distant electromagnetic source we can observe. If we compare the 13.8 Gyr age of the universe to a human lifetime, the CMB is analogous to a day-old baby picture of the universe. Photons from the CMB have been traveling towards us since only 400,000 years after the beginning of the universe, which is thought to have been characterized by a period of exponential expansion known as inflation allowing the universe to be flat and homogeneous on observable scales (Guth, 1981; Linde, 1982; Albrecht & Steinhardt, 1982). While we do not have direct observational access to electromagnetic information prior to the epoch of recombination when the universe transitioned from optically thick to optically thin, it is widely held that quantum fluctuations in the initial moments of the universe were blown up to macroscopic scales during inflation (Lyth & Riotto, 1999; Baumann et al., 2009). These quantum fluctuations seeded density fluctuations that have evolved under the influence of gravity into the large-scale structure of the universe we see today including cosmic voids, filaments, galaxy clusters, and galaxies (Peebles, 1967; Zel'dovich, 1970; Hawking, 1982). Corresponding temperature fluctuations were frozen into the radiation field of the universe during the period of recombination, when the universe cooled enough for electrons and protons to combine into neutral hydrogen atoms and photons were able to free-stream through the universe (Peebles, 1968; Zel'dovich, 1970).

Fluctuations in the radiation field during recombination can be detected today as tiny perturbations in the CMB, which is a nearly ideal blackbody emitter peaking in the millimeter with an effective temperature of $2.7260 \pm 0.0013\,\mathrm{K}$ (Fixsen, 2009). Characterizing the primary anisotropies of the CMB can shed light on the dynamics of inflation and constrain the parameters of the concordance ΛCDM cosmological model with remarkable precision (Baumann et al., 2009; Challinor, 2013). CMB photons interact with matter as they travel through the cosmos, complicating the analysis of the primordial CMB but also embedding more information about the evolution and structure of the universe at more recent times. Finally, astrophysical objects can stand out against the backdrop of the CMB, allowing the discovery and characterization of galaxies, galaxies clusters, and transient events at millimeter wavelengths.

In the remainder of this chapter I will provide a brief overview of the CMB, the ΛCDM model of cosmology, the South Pole Telescope, and the flat-sky approximation and map projections used in this thesis.

1.1 The Cosmic Microwave Background

The CMB is a snapshot of the universe during the epoch of recombination when the primordial plasma became optically thin and photons began to free-stream through the universe. Before recombination ($z \lesssim 1100$), the universe was optically thick with photons continuously undergoing Thomson scattering due to the high density of free electrons. A combination of Thomson scattering and Coulomb scattering between free protons and electrons tightly coupled photons and baryons in the hot plasma, giving rise to a photon-baryon fluid that experienced acoustic oscillations driven by the competing forces of gravity and radiation pressure (Hu & Dodelson, 2002). As temperatures reached $\sim 3000\,\mathrm{K}$, neutral hydrogen atoms gradually condensed out of protons and electrons in the plasma; with fewer free electrons to

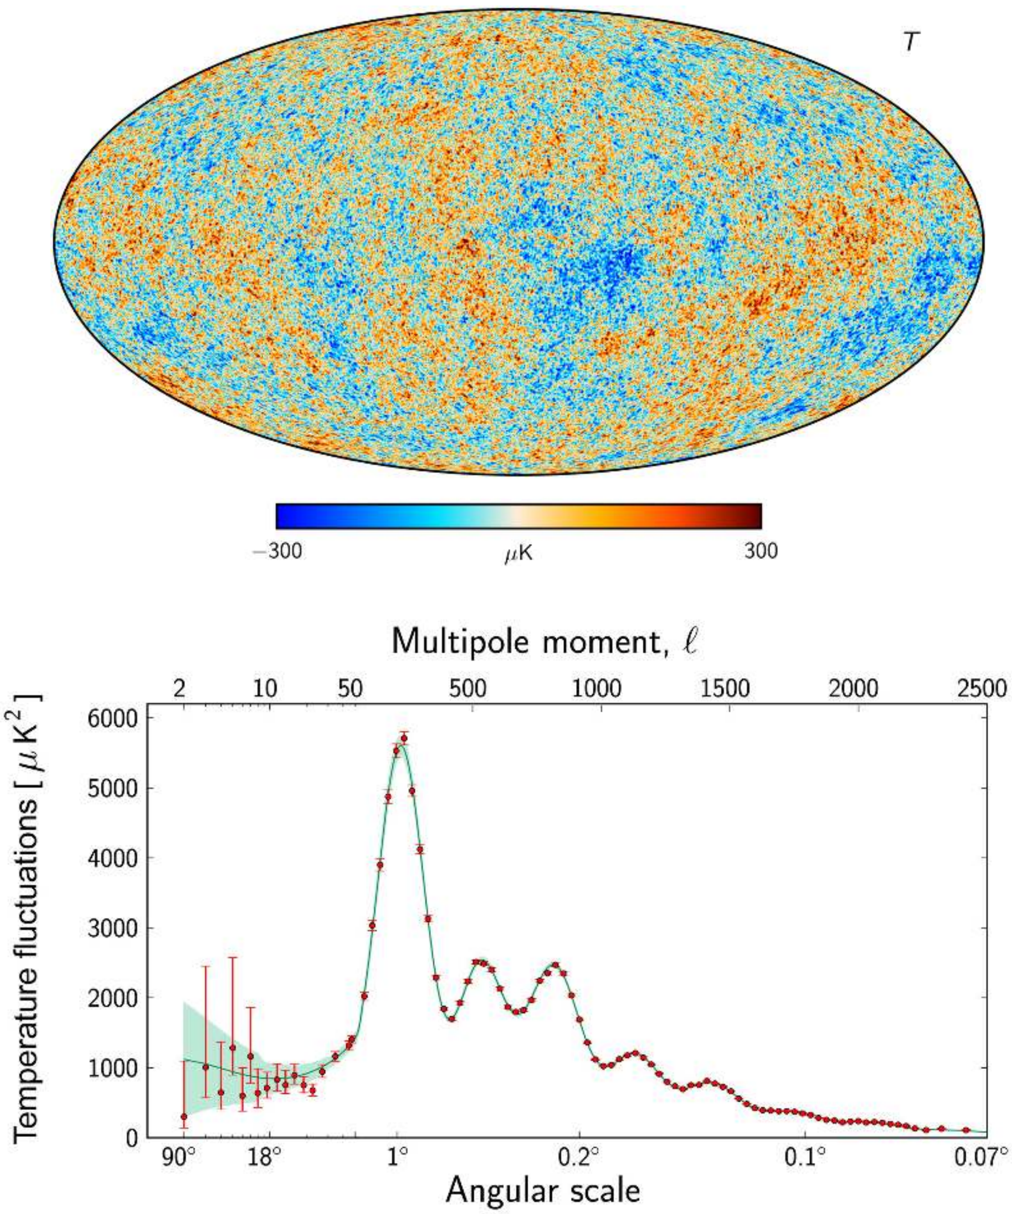

Figure 1.1: Temperature anisotropies of the cosmic microwave background as observed by the *Planck* satellite. Top: *map* of the CMB with the monopole and dipole subtracted in units of temperature. The map is estimated from multi-frequency all-sky maps with the `COMMANDER` algorithm (Planck Collaboration et al., 2020a), and the Galactic plane has been filled in with a Gaussian-constrained realization of the CMB. Bottom: *power spectrum* showing the fluctuations in the above map as a function of angular scale; acoustic oscillations that were frozen into the density and radiation field of the early universe are visible as peaks and troughs in the power spectrum and encode a huge amount of information about the universe. The red data points are *Planck* measurements while the green curve is the best-fit 6 parameter ΛCDM model, demonstrating the remarkable success of the model in explaining the CMB. Credit: ESA and the *Planck* Collaboration.

interact with, photons decoupled from baryons and radiation pressure was no longer able to support the acoustic oscillations, freezing them into the density field of the early universe. Photons that had been trapped in the plasma began to propagate freely from the "surface of last scattering," creating an all-sky thermal background that has since been redshifted from the near-infrared to the millimeter.

The CMB was first discovered by Penzias and Wilson in 1965, who were awarded the Nobel Prize in Physics in 1978 for their discovery (Penzias & Wilson, 1965). It is the most perfect blackbody observed in nature, and is remarkably isotropic once the dipole from our motion relative to the CMB rest frame is removed. However, tiny anisotropies on the order of one part in 10^4 (Figure 1.1) have been measured with increasing precision by space and balloon-borne missions such as the Cosmic Background Explorer (COBE; Smoot et al., 1992), BOOMERanG experiment (de Bernardis et al., 2000), Wilkinson Microwave Anisotropy Probe (WMAP; Bennett et al., 2003), *Planck* observatory (Planck Collaboration et al., 2014a), and SPIDER telescope (Ade et al., 2022a). Ground-based experiments including the DASI interferometer (Kovac et al., 2002), BICEP/*Keck* experiment (Chiang et al., 2010), South Pole Telescope (SPT; Lueker et al., 2010) and Atacama Cosmology Telescope (ACT; Fowler et al., 2010) have also made significant contributions to our understanding of CMB anisotropies. These anisotropies are the imprint of over- and under-densities in the primordial plasma, and their power spectrum (describing the amplitude of fluctuations as a function of angular scale) is a strong probe of cosmological parameters. At recombination fluctuations are expected to be Gaussian, so different angular scales are completely uncorrelated and the power spectrum captures all statistical information present in the primordial CMB (Lewis & Challinor, 2006). Measurements of the CMB power spectrum indicate that we live in a spatially flat universe dominated by dark matter and dark energy (Hu & Dodelson, 2002).

The CMB is also polarized at the $\sim$10% level due to local quadrupole moments in the radiation field at the time of recombination when the universe was transitioning from optically thick to optically thin. In the process of Thomson scattering an electron oscillates along with the electric field of incident radiation, in turn emitting radiation with an intensity that peaks in the direction orthogonal to the incident electric field. The scattered radiation will then have a linear polarization parallel to the incoming electric field. If the local radiation field at the surface of last scattering is isotropic, an equal amount of photons will scatter from each direction, and the net polarization will be zero. However if an electron experiences a local quadrupole such that the intensity of incident photons separated by $90°$ is imbalanced, a net linear polarization will be generated (Hu & White, 1997). In the optically thick limit of the primordial plasma, frequent Thomson scattering smooths out quadrupolar features in the radiation field so linear CMB polarization can only be generated when the local optical depth is $\tau \lesssim 1$. Thus linear polarization in the CMB is produced only as the scattering events that generate it die out (Hu & Dodelson, 2002).

The CMB can be observed in total intensity T, which probes the effective temperature field at recombination, or in polarization via the Stokes Q and U parameters for a total of three distinct observable fields. Circular Stokes V polarization from vortical motions in the primordial photon-baryon fluid is expected to vanish by the time of recombination due to the decay of vector modes with the expansion of the universe (Hu & White, 1997; Kamionkowski et al., 1997). It is often more convenient to decompose the spin-two Q and U polarization field into a curl-free E component and divergence-free B component. Scalar density perturbations in the primordial plasma are only capable of generating E-modes, so any B-mode power in the primordial CMB is expected to come from vector perturbations (which are expected to be damped by inflation) or tensor perturbations sourced by primordial gravitational waves from inflation (Seljak & Zaldarriaga, 1997; Kamionkowski et al., 1997; Zaldarriaga, 2001). A detection of such gravitational waves, whose amplitude is parameterized by the tensor-to-scalar ratio r, would confirm one of the main predictions of inflationary theory. Certain families of cosmological models allow r to be arbitrary small, but some theories suggest $r \gtrsim 10^{-3}$, a level which could be detectable with current and next-generation CMB telescopes (Kamionkowski & Kovetz, 2016).

Thus far we have only considered the *primary* anisotropies of the CMB, which provide information about the universe at and prior to the surface of last scattering. But CMB photons also interact with matter in the universe at later times as they propagate from the surface of last scattering, imprinting information about the growth of structure and reionization history of the universe as well as baryonic astrophysics. These interactions can change the direction, intensity, or frequency of CMB photons and generate *secondary* CMB anisotropies. Secondary anisotropies can be broadly divided into gravitational interactions and electron scattering interactions (Aghanim et al., 2008). Gravitational interactions include lensing which deflects CMB photons but does not change their frequency (see Chapter 2) and the integrated Sachs-Wolfe, Rees-Sciama, and moving-lens effects which all involve the gravitational redshift of CMB photons as they travel through time-varying gravitational potentials (Lewis & Challinor, 2006; Sachs & Wolfe, 1967; Rees & Sciama, 1968; Birkinshaw & Gull, 1983). Electron scattering interactions include the thermal and kinetic Sunyaev-Zel'dovich effects (tSZ and kSZ) involving inverse Compton scattering off hot electrons in galaxy clusters; the tSZ effect boosts the frequency of scattered photons while the kSZ signal is spectrally indistinguishable from primary CMB fluctuations (Sunyaev & Zeldovich, 1970, 1972; Birkinshaw, 1999). CMB secondary anisotropies are a rich source of information about the universe at later times.

1.2 The ΛCDM Model & Cosmological Parameters

Cosmological observations over the last several decades have led to the acceptance of a simple six-parameter ΛCDM concordance model capable of predicting a wide range of observable properties of the universe. The ΛCDM model assumes a spatially flat universe composed of matter (baryons, non-relativistic neutrinos, and cold dark matter–CDM) responsible for the growth of structure in the universe, radiation (photons and relativistic neutrinos), and dark energy parameterized by a constant Λ responsible for the accelerating expansion of the universe (Perivolaropoulos & Skara, 2022). It also assumes that on cosmological scales, the universe is statistically isotropic and homogeneous, and that the behavior of gravity can be described by Albert Einstein's theory of General Relativity (Einstein, 1916).

The ΛCDM model has six free parameters: the baryon density $\Omega_b h^2$, the cold dark matter density $\Omega_c h^2$, the Hubble constant H_0, the optical depth to reionization τ, and the amplitude A_s and spectral index n_s of the primordial scalar power spectrum. Here $h = H_0/100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ is the dimensionless Hubble constant, and $\Omega_i = \rho_i/\rho_{\mathrm{crit}}$ is the density parameter of a given species relative to the critical density $\rho_{\mathrm{crit}} = 3H_0^2/8\pi G$. The angular size of the sound horizon at recombination θ_s is a derived parameter that can be useful when fitting to CMB data. Similarly σ_8 is a derived parameter describing the present-day amplitude of matter density fluctuations on $8h^{-1}$ Mpc scales:

$$\sigma_8 = \frac{1}{2\pi^2} \int \mathrm{d}k\, k^2 P(k) W(kR_8)^2, \tag{1.1}$$

where $P(k)$ is the matter power spectrum and $W(kR_8)$ is the Fourier transform of a three-dimensional top hat filter with a radius of $8h^{-1}$ Mpc. Closely related to σ_8 is the structure growth parameter $S_8 = \sigma_8(\Omega_m/0.3)^{0.5}$ which is less dependent on assumptions about the matter density and allows for more straightforward comparisons between early- and late-time measurements.

The term "concordance" is often used to describe the ΛCDM model because it can predict numerous observed properties of our universe including the statistics of the CMB (Figure 1.1; Planck Collaboration et al., 2020b) and large scale structure (Bernardeau et al., 2002; Planelles et al., 2015), the abundances of light nuclei (Cyburt et al., 2016; Fields et al., 2020), and the expansion history of the universe (Riess et al., 1998; Perlmutter et al., 1999; Freedman et al., 2001). Although the ΛCDM model has been incredibly successful in explaining the behavior of our universe with a minimal number of free parameters, theoretical and observational difficulties with the model remain. Theoretical predictions of the easilycalculable contributions to the vacuum energy density are larger than the observed value by as much as 120 orders of magnitude (Weinberg, 2000), requiring the unknown terms in the vacuum energy density to cancel to an incredible degree of precision; on the other hand the coincidence between the observed present-day vacuum and matter energy densities is puzzling given their extremely different evolutionary histories (Perivolaropoulos & Skara, 2022). On the observational side, discrepancies or "tensions" between largely independent probes of ΛCDM parameters have been growing in recent years as the precision of cosmological measurements has increased. Two of the most (in)famous of these tensions concern the Hubble constant H_0 governing the expansion rate of the universe and the structure growth parameter S_8 related to the linear-theory amplitude of matter density fluctuations. In both cases the discrepancies lie between combinations of early-universe probes such as the CMB, baryonic acoustic oscillations (BAO; DESI Collaboration et al., 2024), and elemental abundances vs. late-time probes including galaxy lensing (shear; Dark Energy Survey and Kilo-Degree Survey Collaboration et al., 2023), galaxy clustering (Abbott et al., 2022), and type-Ia supernovae (Riess et al., 2019; Wong et al., 2020). Tensions could be due to sampling variance, new physics, unknown systematic errors in either the early- or late-universe datasets, or a combination thereof.

1.3 The South Pole Telescope

The South Pole Telescope is a 10 m offset Gregorian telescope with $\sim 1'$ resolution designed to observe the CMB and astrophysical foregrounds such as galaxies and galaxy clusters at millimeter wavelengths (Carlstrom et al., 2011). Located at the Amundsen-Scott South Pole Station in Antarctica at an elevation of over 9000 feet (Figure 1.2), the telescope benefits from perhaps the best millimeter-wave observing conditions on Earth, with atmospheric brightness fluctuations more than an order of magnitude smaller than other observing sites (Bussmann et al., 2005). The SPT saw first light in 2007, and took data for the 2500 deg^2 SPT-SZ survey until 2012. The original 966-detector camera was then replaced with the polarization-sensitive

Figure 1.2: The South Pole Telescope (center left) and BICEP telescope (center) at the Amundsen-Scott South Pole Station, with the Milky Way and Aurora Australis visible in the background. The SPT and BICEP/*Keck* experiments produce complimentary CMB datasets that can be combined to search for primordial gravitational waves. Photo by SPT winter over Geoffrey Chen.

SPTpol receiver boasting 1536 detectors (Austermann et al., 2012) which took observations between 2012 and 2016. The SPTpol survey observed a 500 deg^2 field with a nested 100 deg^2 "ultradeep" field (Figure 1.3, green) as well as an extended 2770 deg^2 extended survey to detect galaxy clusters via the SZ effect (Bleem et al., 2020). The telescope was outfitted with the third-generation ~16 000 detector SPT-3G camera in 2017 and underwent an engineering year in 2018 before beginning a five-year, 1500 deg^2 survey of the southern sky in 2019 (Figure 1.3, purple; Benson et al., 2014; Sobrin et al., 2022). The first five years of the main SPT-3G survey have achieved white noise levels of 3.0 μK-arcmin, 2.5 μK-arcmin and 8.9 μK-arcmin at 90 GHz, 150 GHz and 220 GHz in temperature, with polarization noise levels higher by a factor of $\sqrt{2}$; a minimum-variance linear combination of the frequency bands reduces the white noise level to 1.9 μK-arcmin (Prabhu et al., 2024).

The main 1500 deg^2 field is observed at a ~daily cadence for eight months each year (April-November), but during the austral summer the Sun rises above the horizon and is close enough to the winter field to be picked up by the telescope's sidelobes. To avoid sun contamination SPT-3G switches to observations of three "summer" fields totaling 2650 deg^2 for the remaining four months of the year. The white noise levels of the summer fields integrated over the fiducial 5-year SPT-3G survey is roughly 3-4 times higher than the

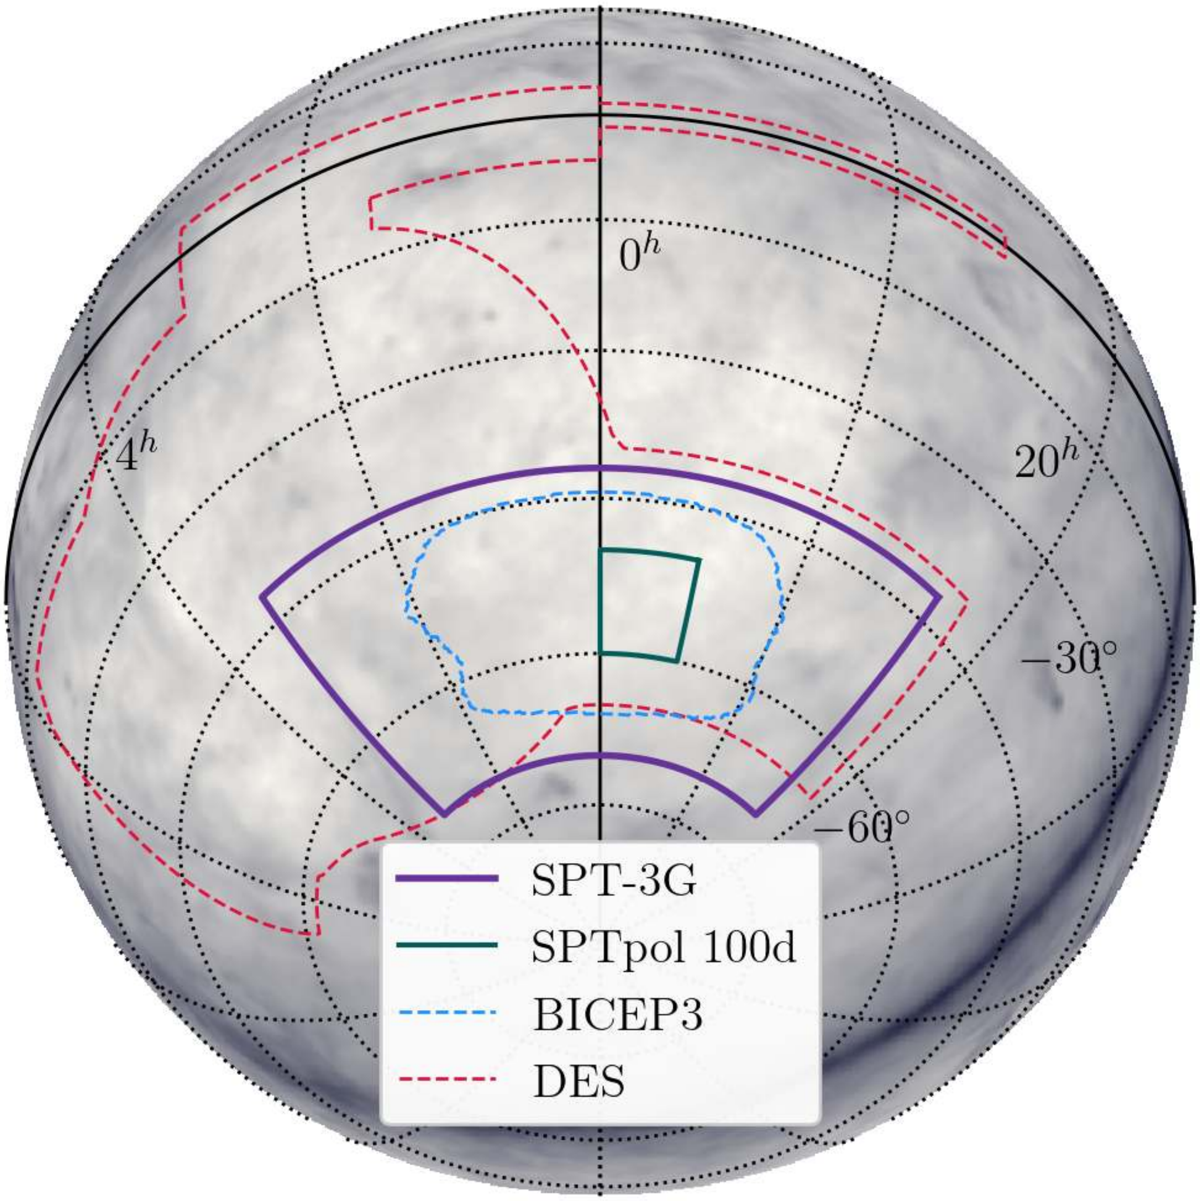

Figure 1.3: The survey footprints of the SPT fields used in this thesis overlaid on thermal dust emission from Planck Collaboration et al. (2015a) and compared to the footprints of BICEP3 and the Dark Energy Survey (DES).

corresponding main field noise levels. The constraints on cosmological parameters from the combined main and summer fields are now dominated by sample variance, so longer integration times do not significantly improve our constraints. As a result SPT-3G has begun new "wide" field observations totaling $6000\,\text{deg}^2$ in the 2024 season, covering all parts of the sky with minimal Galactic contamination feasible to observe from the South Pole before returning to the main $1500\,\text{deg}^2$ field for the 2025 and 2026 seasons. The extended field observations are forecasted to achieve noise levels of $14\,\mu$K-arcmin, $12\,\mu$K-arcmin and $42\,\mu$K-arcmin at 90 GHz, 150 GHz and 220 GHz respectively, and will bring the total SPT-3G survey area to over $10\,000\,\text{deg}^2$ (Prabhu et al., 2024).

The SPT provides complementary datasets to three of the other major CMB experiments of the past ∼decade (*Planck*, ACT, and BICEP/*Keck*) which we briefly describe below. The *Planck* satellite was launched in 2009 and took data continuously until its decommissioning in 2013, and was primarily designed to measure the temperature and polarization anisotropies of the CMB (Planck Collaboration et al., 2014a). As a space-based mission *Planck* did not have to contend with atmospheric noise, and was able to survey the entire sky although the galactic plane was masked for cosmology analyses leaving a sky fraction of e.g. $f_{\text{sky}} = 0.671$ for the 2018 lensing analysis (Planck Collaboration et al., 2020c). *Planck* was sensitive to frequencies between 20 GHz and 1000 GHz, enabling the production of component-separated CMB maps with reduced contamination from foregrounds with a different frequency dependence than the CMB (Planck Collaboration et al., 2020a). The `SMICA` component-separated CMB maps used in the 2018 lensing analysis achieved a resolution of $5'$ and median white noise levels of $27\,\mu$K-arcmin and $52\,\mu$K-arcmin in temperature and polarization respectively (Planck Collaboration et al., 2020c). Planck remains the gold standard for CMB anisotropy measurements on large scales and places the tightest CMB constraints on ΛCDM parameters thanks to its all-sky coverage. In comparison, SPT-3G's improved resolution and noise levels (by factors of ∼5 and ∼10 respectively) push to new frontiers on small scales especially in the cases of CMB polarization and lensing.

The Atacama Cosmology Telescope (ACT; Fowler et al., 2007; Thornton et al., 2016; Henderson et al., 2016) is a 6 m telescope located in the Atacama Desert in Chile that has been observing the CMB since 2008. The SPT and ACT are similar in many ways, each with three generations of receivers installed on comparable timescales sharing roughly the same frequency coverage. The most important difference between the two experiments is their location: ACT's higher latitude allows it to survey a much larger fraction of the sky, but the observing conditions in Chile lead to higher noise levels and more challenging systematics. Atmospheric noise brightness fluctuations are larger at the Atacama site by a factor of 30 (Bussmann et al., 2005), and the daily appearance of the sun reduces the stability of observing conditions. While SPT-3G may perform uninterrupted observations of the main field for 8 months each year while the Sun is at or below the horizon, ACT's beam can be deformed the Sun in daytime and the recent ACT DR6 lensing results did not make use of daytime data due as more extensive efforts were needed to control systematics in the daytime data (Madhavacheril et al., 2024; Qu et al., 2024). The early science-grade CMB maps used in the ACT DR6 lensing analysis achieve a spatially-varying white noise level between 8 μK-arcmin and 15 μK-arcmin over 9400 deg^2. In conclusion, ACT and the SPT-3G main survey provide complimentary datasets, with ACT covering a much larger fraction of the sky (23% of the sky vs. 3.5% when excluding the galaxy) and SPT-3G achieving lower noise levels (1.9 μK-arcmin vs. 8 μK-arcmin to 15 μK-arcmin).

The BICEP/*Keck* experiment was specifically designed to detect the B-mode signature of primordial gravitational waves from inflation, and has achieving progressively lower noise levels by observing a ~400 deg^2 to 600 deg^2 patch of sky with increasingly sensitive receivers since 2010 (BICEP2 Collaboration et al., 2014; BICEP2 and Keck Array Collaborations et al., 2015; Ade et al., 2022b). Located next to the SPT at the Amundsen-Scott South Pole Station (Figure 1.2) and observing a similar patch of sky (Figure 1.3), BICEP/*Keck* also benefits from the excellent observing conditions at the South Pole and has achieved noise levels of 2.8 μK-arcmin, 2.8 μK-arcmin and 8.8 μK-arcmin at 95 GHz, 150 GHz and 220 GHz in the latest combined dataset (Ade et al., 2021). Since inflation is predicted to generate a B-mode signal on degree scales, BICEP/*Keck* instruments are small-aperture telescopes designed to be maximally sensitive to these large angular scales, with excellent control of systematics. The latest BICEP/*Keck* results have placed the tightest constraints on primordial gravitational waves to date, setting a 95% upper limit $r < 0.0036$ on the tensor to scalar ratio and achieving an uncertainty of $\sigma(r) = 0.009$ (Ade et al., 2021).

The uncertainty on r is now dominated by sample variance from lensing B-modes, and removal of this contamination via "delensing" of BICEP/*Keck*'s B-mode maps is required to significantly tighten the constraints on r (BICEP/Keck Collaboration et al., 2021). However BICEP/*Keck*'s small-aperture telescopes have relatively low angular resolution (for example the resolution of the BICEP3 instrument is $\sim 21'$; Ade et al., 2022b) and thus BICEP/*Keck* data does not cover the range of angular scales needed to reconstruct the lensing potential with the signal to noise needed for delensing. Lensing reconstruction is performed by picking out correlations between pairs of spherical harmonics on different scales, so pushing to smaller spatial scales increases the total number of pairs than can be used. SPT on the other hand has relatively high resolution for a CMB telescope, making it particularly well suited to mapping the CMB lensing potential although it does not have the sensitivity and control of systematics on large scales required to detect primordial gravitational waves. Thus the overlapping intermediate-resolution BICEP/*Keck* and high-resolution SPT-3G datasets are quite complementary in this respect, and the two experiments have agreed to share data and collaborate in an effort to use SPT-3G lensing maps to delens combined BICEP/*Keck* and SPT-3G polarization maps. An initial demonstration of this approach with *Planck*, BICEP/*Keck*, and SPTpol data resulted in a reduction of the uncertainty on r of 10%, and predicted that the uncertainty would be reduced by a factor of 2.5 with the inclusion of SPT-3G data (BICEP/Keck Collaboration et al., 2021)

SPT observations have enabled a broad range of cosmological and astrophysical analyses. The SPT has measured power spectra of the primary CMB anisotropies first in temperature

(Keisler et al., 2011) and then in polarization (Crites et al., 2015; Henning et al., 2017; Dutcher et al., 2021). Measurements of secondary CMB anisotropy power spectra have also been made with increasing precision (Lueker et al., 2010; Reichardt et al., 2021). These secondary power spectra include contributions from emissive millimeter sources including resolved radio sources and dusty galaxies as well as the unresolved cosmic infrared background (CIB; see Kashlinsky et al., 2018, for a review). The SPT has produced catalogs of these sources (Vieira et al., 2010; Everett et al., 2020) which include fascinating high-redshift galaxies magnified by strong lensing (e.g. Vieira et al., 2013). Individual galaxy clusters also appear as decrements in the SPT 90 GHz and 150 GHz bands due to the SZ effect; SPT performed the first discovery of galaxy clusters using the SZ effect (Staniszewski et al., 2009), and has since released catalogs of these clusters (Plagge et al., 2010; Bleem et al., 2020). The SPT is very well suited to measure gravitational lensing of the CMB due to its low noise and high angular resolution, and has reported a suite of successful CMB lensing measurements (van Engelen et al., 2012; Story et al., 2015; Omori et al., 2017; Wu et al., 2019; Millea et al., 2021; Pan et al., 2023). The first detection of lensing B-modes was performed with SPT (Hanson et al., 2013), and estimates of the lensing potential have been used to delens polarization maps as mentioned above (Manzotti et al., 2017; BICEP/Keck Collaboration et al., 2021). Cross-correlations between CMB lensing maps and other tracers of large scale structure have been used to constrain both cosmological and astrophysical parameters (Bleem et al., 2012; Holder et al., 2013; Geach et al., 2013). The main SPT fields almost entirely overlap with the Dark Energy Survey (DES; Figure 1.3), allowing for joint analyses of the auto- and cross-spectra of CMB lensing, galaxy lensing, and galaxy clustering (5×2pt and 6×2 pt; Baxter et al., 2016; Omori et al., 2018; Chang et al., 2023; Abbott et al., 2023). Finally, SPT-3G's ~daily observations of the main field have allowed for time-domain studies including transient searches (Guns et al., 2021), asteroid measurements (Chichura et al., 2022), AGN variability studies (Hood et al., 2023), and stellar flare catalogs (Tandoi et al., 2024).

1.4 The Flat-Sky Approximation and Map Projections

The CMB is an all-sky spherical dataset, and a host of software packages have been developed to store and analyze CMB data directly on the sphere (notably HEALPix, Górski et al., 2005). We often use equatorial coordinates right ascension (R.A.) and declination (Decl.) represented on the sphere as (θ, ϕ), where θ is the angular distance northward perpendicular to the celestial equator and ϕ is the angular distance eastward perpendicular to the March equinox. However for ground-based experiments with relatively small sky fractions, it is often more convenient to analyze data in the "flat-sky" approximation by projecting a spherical map with spherical coordinates (θ, ϕ) to a (x, y) grid that treats the observed patch of sky as a Euclidean plane. This allows computationally expensive spherical harmonic transforms (SHTs) to be replaced with highly-optimized fast Fourier transforms (FFTs), and may dramatically reduce the map memory footprint for ground-based experiments that observe only a fraction of the sky. Fourier modes can also be conceptually easier to work with than spherical harmonics.

We can relate a spin-0 field $X(\theta, \phi)$ to its spherical harmonic transform $X_{\ell m}$ with

$$
\begin{aligned}
X(\theta, \phi) &= \sum_{\ell=0}^{\ell=\infty} \sum_{m=-\ell}^{m=\ell} X_{\ell m} Y_{\ell m}(\theta, \phi) \\
X_{\ell m} &= \int \mathrm{d}\Omega \, X(\theta, \phi) Y^*_{\ell m}(\theta, \phi),
\end{aligned}
$$

where $Y_{\ell m, \hat{\mathbf{n}}}$ are the spherical harmonic functions and $\mathrm{d}\Omega = \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi$ is the differential solid angle. We take the flat-sky limit by assuming that $\ell \gg 1$ and all photon trajectories are perpendicular to the flat-sky plane, replacing spherical (θ, ϕ) with planar $\hat{\mathbf{n}} = (x, y)$, and adopting a plane-wave harmonic basis $\boldsymbol{\ell} = (\ell_x, \ell_y)$. Then the flat-sky map $X_{\hat{\mathbf{n}}}$ can be related to its Fourier transform $X_{\boldsymbol{\ell}}$ with

$$X_{\boldsymbol{\ell}} = \int \mathrm{d}^2\hat{\mathbf{n}}\, X_{\hat{\mathbf{n}}} e^{-i\boldsymbol{\ell}\cdot\hat{\mathbf{n}}}, \tag{1.2}$$

$$X_{\hat{\mathbf{n}}} = \int \frac{\mathrm{d}^2\boldsymbol{\ell}}{(2\pi)^2} X_{\boldsymbol{\ell}} e^{i\boldsymbol{\ell}\cdot\hat{\mathbf{n}}}. \tag{1.3}$$

We will denote the power spectrum of a field X as the Fourier-space product of the field with its complex conjugate, $C_\ell = \langle X_{\boldsymbol{\ell}} X^*_{\boldsymbol{\ell}} \rangle$. We will also use bolded subscripts e.g. $X_{\boldsymbol{\ell}}$ to denote two-dimensional Fourier transforms.

The flat-sky approximation requires the use of a specific projection that maps spherical (θ, ϕ) to planar (x, y) coordinates. The field of cartography has produced numerous projections over the centuries that preserve different properties of the sphere such as area, distance, or angle. Each projection has its own advantages and drawbacks, and a specific projection should be chosen with the applications of the map in mind. SPT-3G usually uses the Lambert Zenithal Equal-Area projection (ProjZEA), which was first published by French-Swiss polymath Johann Heinrich Lambert (Lambert, 1772):

$$x = k \cos\theta \sin(\phi - \phi_0), \tag{1.4}$$

$$y = k(\sin\theta\cos\theta_0 - \cos\theta\sin\theta_0\cos(\phi - \phi_0)), \tag{1.5}$$

where (θ_0, ϕ_0) denotes the center of the projection and

$$k = \sqrt{\frac{2}{1 + \sin\theta\sin\theta_0 + \cos\theta\cos\theta_0\cos(\phi - \phi_0)}}. \tag{1.6}$$

As an equal-area projection each ProjZEA pixel contains the same solid angle on the sphere, which is useful for power spectrum analyses as it mitigates mode-mixing between different angular scales; non-equal-area projections tend to broaden the acoustic peaks of the CMB power spectrum. However, equal-area projections necessarily do not preserve angular

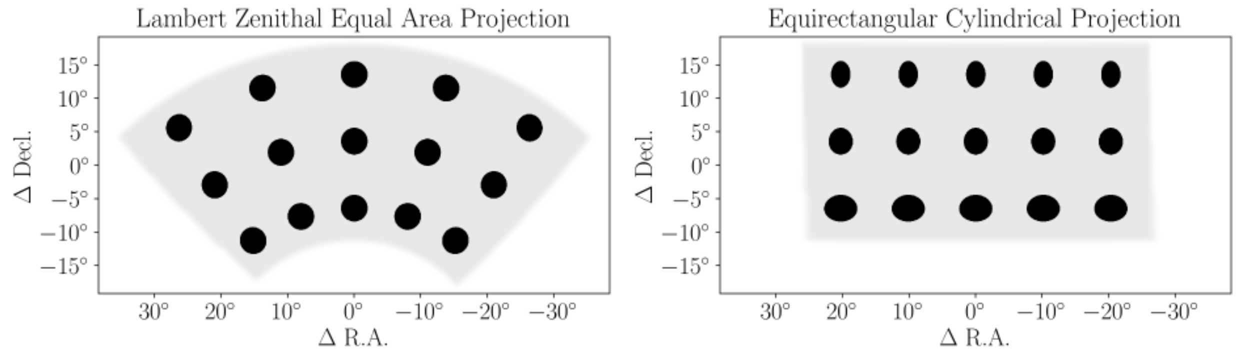

Figure 1.4: The SPT-3G footprint (gray) in two different flat-sky projections, with disks defined in right ascension and declination (black) superimposed to show distortion of each projection. The disks are spaced along three rows and five columns of constant right ascension and declination respectively. In the Lambert Zenithal Equal-Area projection commonly used by SPT-3G (ProjZEA; left) the disks remain circular and equal-area at all locations in the map, but lines of constant right ascension and declination do not align with the x and y axes. In the Equirectangular Cylindrical projection commonly used by BICEP/*Keck* (right), lines of constant right ascension and declination do align with the x and y axes, but the disks reveal a declination-dependent distortion and do not contain the same areas.

relationships on the sphere (i.e. they are not conformal). As a result lines of constant declination do not align with the y-axis except at the center of the projection and are not parallel to each other, whereas lines of constant right ascension become curved in ProjZEA imparting a characteristic "windshield-wiper" appearance when large rectangular patches of θ, ϕ are projected (Figure 1.4). The flat-sky approximation becomes worse as the sky area is increased, and the $1500\,\mathrm{deg}^2$ SPT-3G survey footprint is large enough that the distortion of angles in ProjZEA is quite significant.

ProjZEA is better than many other projections for one-dimensional power spectrum analyses as it induces minimal mode mixing between different angular scales $|\boldsymbol{\ell}| \neq |\boldsymbol{\ell}'|$, but it does mix modes with harmonic multipole ℓ travelling purely in the North or East directions into multiple flat-sky modes with $|\boldsymbol{\ell}'| = \ell$ but varying $\tan\psi_{\boldsymbol{\ell}'} = \ell'_y/\ell'_x$. This has significant consequences for SPT data since the telescope performs constant-declination scans across the sky, meaning that the filtering we apply to these one-dimensional scans only acts in the East-West direction. In cylindrical projections or for very small patches of sky in ProjZEA this filtering cleanly projects onto ℓ_x modes, but for larger patches the filtering "splays" over a wide range of ℓ_x and ℓ_y modes. In the SPT-3G patch for example, the scan direction on the left, center, and right parts of the map corresponds to $\ell_y = \ell_x, \ell_y = 0$, and $\ell_y = -\ell_x$ respectively. This makes modeling and correcting for scan-direction filtering effects in ProjZEA Fourier-space much more difficult, as we will see in later chapters.

1.4.1 Considerations for Polarized Maps

The spin-2 Q and U polarization fields observed by CMB experiments can be decomposed into E and B fields by applying a local rotation in Fourier space,

$$\begin{bmatrix} E_{\boldsymbol{\ell}} \\ B_{\boldsymbol{\ell}} \end{bmatrix} = \begin{bmatrix} \cos 2\psi_{\boldsymbol{\ell}} & \sin 2\psi_{\boldsymbol{\ell}} \\ -\sin 2\psi_{\boldsymbol{\ell}} & \cos 2\psi_{\boldsymbol{\ell}} \end{bmatrix} \begin{bmatrix} Q_{\boldsymbol{\ell}} \\ U_{\boldsymbol{\ell}} \end{bmatrix} \tag{1.7}$$

where $\tan \psi_{\boldsymbol{\ell}} = \ell_y/\ell_x$ is the slope of a given mode $\boldsymbol{\ell}$ in the Fourier plane (Hirata & Seljak, 2003a). E and B fields are often described as curl-free and divergence-free analogous to electric and magnetic fields, however this analogy is not exact as Q and U make up a spin-two headless vector field rather as opposed to the a spin-one vector fields relevant in electromagnetism. The EB basis allows us to describe the spin-2 polarization field in terms of a scalar E field and a pseudo-scalar B field. If we perform a few trigonometric substitutions and apply the above rotation matrix to Q and U, we find:

$$\boldsymbol{\ell}^2 E_{\boldsymbol{\ell}} = \left(\ell_x^2 - \ell_y^2\right) Q_{\boldsymbol{\ell}} + 2\ell_x \ell_y U_{\boldsymbol{\ell}} \tag{1.8}$$

$$\boldsymbol{\ell}^2 B_{\boldsymbol{\ell}} = -2\ell_x \ell_y Q_{\boldsymbol{\ell}} + \left(\ell_x^2 - \ell_y^2\right) U_{\boldsymbol{\ell}}. \tag{1.9}$$

Recalling the Fourier derivative property $\nabla X_{\boldsymbol{\ell}} = -i\boldsymbol{\ell} X_{\boldsymbol{\ell}}$, we see that we can form fields $\chi_E = \boldsymbol{\ell}^2 E_{\boldsymbol{\ell}}$ and $\chi_B = \boldsymbol{\ell}^2 B_{\boldsymbol{\ell}}$ depending only on linear combinations of the second partial derivatives of Q and U. Calculating E and B from pixelized partial-sky Q and U maps is a delicate process due to the interaction between the the boundaries of the observed region and the spatial derivatives of Q and U; the EB decomposition is not well-defined for a partial-sky map and can produce "ambiguous" modes that contribute to both E and B. Several "purified" estimators that leverage χ_E and χ_B have been developed to separate out these ambiguous modes that can otherwise cause $E \rightleftharpoons B$ mixing (Bunn et al., 2003; Smith & Zaldarriaga, 2007)

Even when using the simple EB estimator given in Equation (1.7), care must be taken to preserve the mathematical properties of the input maps. $Q_{\hat{\mathbf{n}}}, U_{\hat{\mathbf{n}}}, E_{\hat{\mathbf{n}}}, B_{\hat{\mathbf{n}}}$ are real-valued maps, so their Fourier transforms must obey the Hermitian property $X_{-\boldsymbol{\ell}} = X^*_{\boldsymbol{\ell}}$, i.e. they must be equal to their conjugate transpose. However, Equation (1.7) does not preserve this property when applied naively to FFTs of even-sided pixelized maps. It turns out that the Hermitian property is broken by the $\sin 2\psi_{\boldsymbol{\ell}} = 2\ell_x \ell_y / \boldsymbol{\ell}^2$ term, which may be initially surprising as $2\ell_x \ell_y / \boldsymbol{\ell}^2$ should obey the Hermitian property. But things become clearer if we think of this term as a product of partial derivatives in the $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ directions and consider a note on FFT-based differentiation (Johnson, 2011) by the co-creator of the FFTW library (Frigo & Johnson, 2005). Johnson shows that the Nyquist frequency of the derivative of a function with an even number of samples must be zero to preserve the function's symmetries. Specifically, a non-zero imaginary Nyquist component violates the Hermitian property, while a non-zero real Nyquist component violates the property that a derivative of an even function must be odd and vice versa. A naive FFT-based derivative $-i\ell_x X_{\boldsymbol{\ell}}$ will produce a non-zero Nyquist coefficient, and Johnson argues that this component should be manually nulled to preserve the input function's symmetries. As a consequence, applying the symmetry-preserving first derivative *twice* to a function with an even number of samples is *not* equivalent to the second derivative of the function, since the Nyquist component is not required to be zero in the latter case. Returning to the application of $QU \rightleftharpoons EB$ transformations in CMB data analysis, nulling the Nyquist component of the $\sin 2\psi_{\boldsymbol{\ell}}$ is not an ideal solution because it loses information and makes it impossible to recover the input QU maps from the EB maps. Instead we manually enforce the symmetry on the $\sin 2\psi_{\boldsymbol{\ell}}$ term by multiplying the Nyquist coefficients by -1, which does not lose any information and preserves the Hermitian properties of the maps at the cost of giving an incorrect Nyquist component in the EB maps. I led an effort to understand and solve these issues in SPT-3G's software library, which had existed for many years. In practice this correction not terribly important since modes near Nyquist frequency are often lowpassed in CMB analyses to avoid aliasing, but this finding serves as a striking reminder that rigorously calculating derivatives of pixelized maps is not as simple as continuous Fourier analysis would suggest and can break the conjugate transpose symmetries of FFTs if not done carefully.

Even more care must be taken when estimating EB from QU maps in non-cylindrical projections, which induce a spatially-varying *effective* rotation such that sky North and the projection y-axis are not in general aligned with each other. Since the orientation of Q and U are defined relative to cardinal directions on the sphere, this spatially-varying rotation changes the alignment of the Q $+$ and U $\times$ patterns of curl-free polarization map. If not corrected for, this spatially varying rotation introduces curl to the projected polarization maps that leaks E into B and vice-versa.

When dealing with polarized QU data, in most cases one wants to undo the rotation induced by projection, or "flatten" the QU field to undo this leakage. A spin-2 QU polarization field on the sphere is transformed under a local rotation ξ as

$$\left(Q_{\hat{\mathbf{n}}}^{\mathrm{rot}} + iU_{\hat{\mathbf{n}}}^{\mathrm{rot}}\right) = (Q_{\hat{\mathbf{n}}} + iU_{\hat{\mathbf{n}}})e^{-i2\xi}, \tag{1.10}$$

rotating Q into U and vice versa. The rotation itself is easily accomplished using the equation above, but calculating the appropriate rotation angle to undo the effective rotation of non-cylindrical projections is more delicate. One could define the rotation angle as the angle between sky North $\hat{\boldsymbol{\theta}}$ and flat-sky "up" $\hat{\boldsymbol{y}}$, or as the angle between sky East $\hat{\boldsymbol{\phi}}$ and flat-sky "right" $\hat{\boldsymbol{x}}$. For angle-preserving conformal projections these angles are equivalent, and the undoing of the effective rotation can be exact to numerical precision. However non-conformal projections such as ProjZEA do not preserve angular relationships on the sphere; for example, celestial North and East are only separated by 90° at the center of the projection. Thus we use an equal combination of the North vs. "up" and East vs. "right" estimates of the rotation angle:

$$2\xi = \mathrm{atan2}\left(\frac{\partial\theta/\partial y}{\partial\theta/\partial x}\right) + \mathrm{atan2}\left(-\frac{\partial\psi/\partial x}{\partial\psi/\partial y}\right). \tag{1.11}$$

In practice we calculate these derivatives by applying finite difference methods to pixelized maps of right ascension and declination, but automatic differentiation methods that have been popular in recent years could also be applied to Equations (1.4) and (1.5) to calculate the gradients more precisely.

To summarize, all non-cylindrical projections introduce a spatially-varying rotation that mixes Q into U and thus E into B; this can be corrected by locally "flattening" the QU polarization field with Equation (1.10). However a well-defined rotation angle does not exist for non-conformal projections such as ProjZEA, and a residual curl will remain in the projected maps with an amplitude that grows with distance from the projection center. More advanced methods to address this residual curl can be imagined, such as nulling any B-mode signal in a curved-sky pixelization of the maps, projecting, and then constructing a leakage template from the resulting leaked flat-sky B-modes. We do not employ such methods in this thesis, however.

1.5 Thesis Overview

Since joining the SPT-3G collaboration in the spring of my first year of graduate school, I have had the pleasure of participating in a wide range of projects dealing with high-level science, data analysis, and software development. I was second author on the first application of joint Bayesian parameter inference and optimal lensing reconstruction to CMB data (Millea et al., 2021), and third author on the first SPT-3G transient search (Guns et al., 2021). SPT-3G transients will not be discussed in this thesis but comprised a major part of my work in SPT during the early part of my graduate career, and I was honored to receive the AAS

Chambliss Astronomy Achievement Student Award for my presentation of SPT-3G results on millimeter-wavelength flare stars at the 237th AAS meeting in January 2021. For the last three and a half years I have been leading a CMB lensing analysis with data from the 2019 and 2020 SPT-3G observing seasons which is now in its final stages. The CMB maps I made as part of this analysis are being used in several other ongoing SPT-3G projects including lensing, delensing, power spectrum, and birefringence analyses. During my time in SPT-3G I also been one of the most active graduate student contributors to the SPT-3G software repository, with nearly 120 commits amounting to 25 000 and 12 000 lines of code added and deleted respectively.

In the following pages I will describe my research on gravitational lensing of the CMB. My dissertation mainly focuses on the reconstruction of the most sensitive CMB lensing maps made to date with two years of SPT-3G data, but also includes results from a SPTpol lensing analysis. In Chapter 2 I introduce gravitational lensing of the CMB and the basics of weak lensing before describing the standard quadratic estimator used to reconstruct CMB lensing maps. In Chapter 3 I summarize a CMB lensing analysis (Millea et al., 2021) which used $100\,\mathrm{deg}^2$ of SPTpol data to compare the standard quadratic estimator to a novel Bayesian method. Chapter 4 describes the SPT-3G 2019+2020 flat-sky quadratic estimator analysis that is the main subject of my thesis; some of the information presented in this chapter is not yet public. I conclude in Chapter 5 with a summary of my work and a discussion of future directions in the field of CMB lensing.

Chapter 2

Gravitational Lensing of the CMB

Photons from the primordial CMB are deflected by over- and under-densities in the matter distribution as they propagate across the universe, changing the statistics of the observed CMB anisotropies and encoding information about cosmic structure integrated over a broad range of redshifts. Given the quantity and depth of large scale structure potential wells between us and the CMB, one can expect photons to experience $\sim 2'$ deflections on average; on the other hand, the average size of the potential wells along our line of sight is $\sim 2°$ so deflections will be coherent on degree scales (Lewis & Challinor, 2006). Hot and cold spots will thus be magnified in some patches of the lensed sky and shrunk by the same amount in others, mixing a single angular scale ℓ in the unlensed CMB into a range of scales in the lensed CMB.

Mixing of scales from lensing affects both the power spectrum of the lensed CMB and its covariance. The acoustic peaks and troughs of the primary CMB are smoothed since power at the peaks is transferred towards the troughs and vice versa. The CMB has a red spectrum beyond the first few acoustic peaks with larger fluctuations on large scales, so power is also preferentially shifted power to smaller scales. Finally, this mixing of scales breaks the Gaussian spatial statistics of the primary CMB. In the unlensed CMB, different angular modes are uncorrelated and the power spectrum covariance is diagonal: $\langle X_{\boldsymbol{\ell}} X_{\boldsymbol{\ell}'} \rangle = C_\ell \delta_{\ell,\ell'}$

where $\delta_{\ell,\ell'}$ is the Dirac delta function. This no longer holds for the lensed CMB, where the off-diagonals of the covariance encode the amount of mixing between different scales from lensing. There are yet further implications for lensing of CMB polarization if we recall the second-derivative relation between QU and EB discussed in Chapter 1.4.1; mixing of angular scales in Q and U maps will couple with the derivatives of the $QU \rightarrow EB$ basis transformation and mix E into B, drowning out any primordial gravitational wave signal from inflation.

This chapter is organized as follows. I begin by introducing the basics of weak gravitational lensing in Chapter 2.1 before applying these concepts to the context of CMB lensing in Chapter 2.2. Chapter 2.3 explores the science cases of CMB lensing reconstruction. I conclude in Chapter 2.4 by discussing the practical aspects of lensing reconstruction with the quadratic estimator that will be employed in Chapters 3 and 4.

2.1 An Introduction to Weak Lensing

We begin with a brief historical review of gravitational lensing based on Schneider et al. (1992). The idea that light could be affected by an inhomogeneous gravitational field has existed since at least the time of Issac Newton, who wondered in the 1704 edition of *Opticks*: "Do not Bodies act upon Light at a distance, and by their action bend its Rays; and is not this action strongest at the least distance?" This idea was taken further by the British scientist Henry Cavendish and French polymath Pierre-Simon Laplace, who independently realized towards the end of the 18th century that the escape velocity of a test mass from a sufficiently massive and compact object would exceed the speed of light. These ideas preceded the development of a general relativistic theory of black holes by over one hundred years.

A prediction for the deflection of light by a massive object was first published by Johann Georg von Soldner in 1801, who calculated that a ray of light passing a distance r from an object of mass M would be deflected by an angle $\hat{\alpha} = 2GM/c^2 r$ and noted that this evaluates to $\hat{\alpha} = 0.87''$ for light passing just above the surface of the Sun. This prediction, arrived at independently by Einstein in 1911, is a factor of two smaller than the correct general relativistic calculation (also derived by Einstein) because it neglects the curvature of spacetime (Lewis & Challinor, 2006; Meneghetti, 2021). The first experimental confirmation of gravitational lensing, and of the general relativistic prediction, came when Frank Dyson and Arthur Eddington observed the deflection of starlight by the Sun during the 1919 solar eclipse (Dyson et al., 1920)

In general relativity, photons follow null geodesics in a curved spacetime; however most astrophysical contexts including CMB lensing allow a simpler linearized treatment (Bartelmann & Schneider, 2001; Lewis & Challinor, 2006). In the presence of first-order density perturbations given by a Newtonian potential Φ, a light ray experiences an effective index of refraction $\eta = 1 - 2\Phi/c^2$. Following Fermat's principle that light rays travel along paths that minimize the travel time (or more precisely, along paths of stationary time) we can write the deflection angle along such a path $\mathrm{d}r$ in physical coordinates as

$$\hat{\boldsymbol{\alpha}}_{\mathrm{tot}} = -\frac{2}{c^2} \int \mathrm{d}r \, \boldsymbol{\nabla}_{\perp} \Phi \tag{2.1}$$

where $\boldsymbol{\nabla}_{\perp}$ is the gradient *perpendicular* to the photon path (Kilbinger, 2015).

The deflection angle as defined above can be difficult to work work with as it involves a full integral along the deflected photon path, amounting to a full ray-tracing calculation. Instead the Born approximation is often used, which assumes that the deflection angle is small and thus the gradient of the potential can be evaluated along the *undeflected* path. Weak lensing is often used to mean that the deflection angles are small, and the Born approximation is valid for arcminute-scale deflections expected in CMB lensing (Lewis & Challinor, 2006). We will employ the thin screen approximation in which all light is emitted from a single source plane and the size of the lens is small relative to the distances between the observer, lens, and source (Bartelmann & Schneider, 2001; Meneghetti, 2021). In this case we can also treat

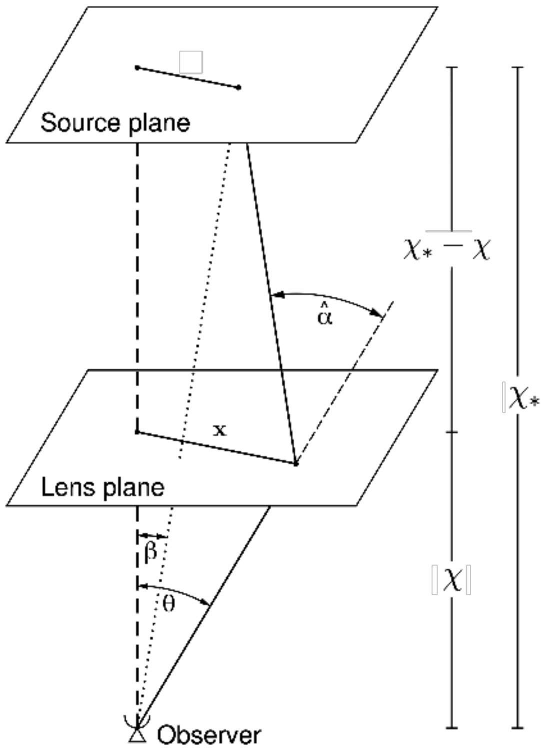

Figure 2.1: A schematic of a weak gravitational lensing scenario in the thin screen approximation, adapted from Bartelmann & Schneider (2001). An observer is positioned at comoving distances χ from the lens plane and χ_* from the source plane. A light ray emitted from a direction $\boldsymbol{\beta}$ in the source plane is deflected by an angle $\hat{\boldsymbol{\alpha}}$ at the lens plane, and the deflected ray is observed at a direction $\boldsymbol{\theta}$ by the observer.

the lens as a single plane, and assume that the photon trajectory is composed of straight-line segments with a single deflection at the lens plane. A typical lensing scenario given these assumptions is shown in Figure 2.1.

The deflection angle in Equation (2.1) gives the deflection experienced at the *lens* plane, rather than the change between the original and deflected directions $\boldsymbol{\beta}$ and $\boldsymbol{\theta}$ experienced by an observer. When accounting for the full geometry of the source-lens-observer system, it is easier to work with comoving coordinates that are unaffected by the expansion of the universe. Consider two points with a comoving separation $\mathbf{x}$ in the lens plane located at a comoving distance χ from the observer. If these points have an an angular separation $\boldsymbol{\theta}$ from the observer's perspective, we can relate the apparent angular separation to their comoving separation by $\mathbf{x} = \chi\boldsymbol{\theta}$ in a flat universe. This allows us us to write the angular deflection at the lens plane as (Kilbinger, 2015)

$$\hat{\boldsymbol{\alpha}} = -\frac{2}{c^2}\boldsymbol{\nabla}_{\perp}\Phi(\mathbf{x}, \chi), \tag{2.2}$$

and relate the deflected and undeflected comoving positions on the *source* plane at distance χ_* from the observer as

$$\chi_*\boldsymbol{\theta} = \chi_*\boldsymbol{\beta} - (\chi_* - \chi)\hat{\boldsymbol{\alpha}}. \tag{2.3}$$

This is the **lens equation**, which can be written purely in terms of angular separations as

$$\boldsymbol{\beta} = \boldsymbol{\theta} + \boldsymbol{\alpha}, \tag{2.4}$$

where $\boldsymbol{\alpha}$ is the *scaled* deflection angle

$$\boldsymbol{\alpha} = \frac{\chi_* - \chi}{\chi_*}\hat{\boldsymbol{\alpha}}. \tag{2.5}$$

Finally, we can use the Born approximation to calculate an effective integrated lensing potential ϕ that incorporates the effect of an arbitrary number of lens planes along the line of sight (Lewis & Challinor, 2006; Kilbinger, 2015; Meneghetti, 2021):

$$\phi(\boldsymbol{\theta}) = -\frac{2}{c^2}\int_0^{\chi_*} \mathrm{d}\chi \, \frac{\chi_* - \chi}{\chi\chi_*}\Phi(\boldsymbol{\theta}\chi, \chi). \tag{2.6}$$

The projected lensing potential is related to the net scaled deflection angle by $\boldsymbol{\alpha} = \boldsymbol{\nabla}\phi$. We will also introduce the convergence $\kappa = -\nabla^2\phi/2$ which can be interpreted as both the projected surface density along the line of sight and the isotropic increase or decrease in the size of a lensed source (Kilbinger, 2015).

2.2 Weak Lensing of the CMB

CMB photons have been propagating across the universe for nearly the entirety of its existence, so one might ask if the approximation of all large scale structure as a single infinitesimally thin plane accurately captures the effect of gravitational lensing on the CMB. Recombination is approximately instantaneous compared to the lifetime of the universe so the source plane approximation is reasonable, but any line of sight to the CMB intersects on the order of 50 potential wells (Lewis & Challinor, 2006). On large scales where the cosmic structure is well described by linear perturbation theory, both the Newtonian potential Φ and the lensing potential ϕ are small and the Born approximation remains accurate. However as matter perturbations evolve under the influence of gravity, non-linear structures emerge first on small scales and grow hierarchically to progressively larger scales (Planelles et al., 2015). As a result the lensing potential becomes non-Gaussian on increasingly large scales with time, and the Born approximation breaks down. Fortunately even the largest non-Gaussian scales are small relative to the degree-scale peak of the deflection angle power spectrum, and are only expected to bias temperature-based lensing measurements at $\sim$1% level (Lewis & Challinor, 2006; Böhm et al., 2018). The thin screen and Born approximations are thus well-justified for CMB lensing.

We will now discuss the effects of lensing on the statistics of CMB in greater detail and motivate the use of statistical estimators to reconstruct the lensing potential. To do so we will return to the flat-sky approximation and consider a CMB map $X(\hat{\mathbf{n}}) \in \{T, Q, U\}$ and its Fourier transform $X_{\boldsymbol{\ell}}$. Weak lensing deflects CMB photons by an angle $\boldsymbol{\alpha} = \boldsymbol{\nabla}\phi$, so we can relate the lensed field X to the unlensed field X^U by

$$X(\hat{\mathbf{n}}) = X^U(\hat{\mathbf{n}} + \boldsymbol{\nabla}\phi) \tag{2.7}$$

If we perform a first-order Taylor expansion of this expression, we find

$$X_{\hat{\mathbf{n}}} \approx X^U_{\hat{\mathbf{n}}} + \nabla_i \phi_{\hat{\mathbf{n}}} \nabla^i X^U_{\hat{\mathbf{n}}}. \tag{2.8}$$

Taking the Fourier transform of both sides, the second term becomes a convolution between gradients of the potential and the unlensed CMB (Hu, 2000),

$$X_{\boldsymbol{\ell}} \approx X^U_{\boldsymbol{\ell}} + (\nabla_i \phi_{\mathbf{L}} * \nabla^i X^U_{\boldsymbol{\ell}}) \tag{2.9}$$

$$\approx X^U_{\boldsymbol{\ell}} + \int \mathrm{d}\boldsymbol{\ell}' \; \nabla_i \phi_{\mathbf{L}} \nabla^i X^U_{\boldsymbol{\ell}}, \tag{2.10}$$

where we have introduced the variable $\mathbf{L}$ to denote the Fourier modes of the lensing potential. Throughout the rest of this thesis we will use $\boldsymbol{\ell}$ for CMB modes and $\mathbf{L}$ for lensing modes. The lensing modes $\mathbf{L}$ are related to the convolution variable $\boldsymbol{\ell}'$ in Equation (2.10) by $\mathbf{L} = \boldsymbol{\ell} - \boldsymbol{\ell}'$; this illustrates how lensing induces off-diagonal correlations between different angular scales in the lensed CMB separated by $\mathbf{L}$. If we use the Fourier derivative property $\nabla X_{\boldsymbol{\ell}} = -i\boldsymbol{\ell} X_{\boldsymbol{\ell}}$ we can rewrite Equation (2.10) as

$$X_{\boldsymbol{\ell}} \approx X^U_{\boldsymbol{\ell}} + \int \mathrm{d}\boldsymbol{\ell}' \; (-\boldsymbol{\ell} \cdot \mathbf{L} \phi_{\mathbf{L}}) X^U_{\boldsymbol{\ell}}, \tag{2.11}$$

and simplify the integrand to a product of the unlensed field with a lensing kernel $W_{\boldsymbol{\ell},\mathbf{L}}$:

$$X_{\boldsymbol{\ell}} \approx X^U_{\boldsymbol{\ell}} + \int \mathrm{d}\boldsymbol{\ell}' \, W_{\boldsymbol{\ell},\mathbf{L}} X^U_{\boldsymbol{\ell}}. \tag{2.12}$$

The second term in Equation (2.12) gives us the expression for the lensing signal $\delta X_{\boldsymbol{\ell}} = X_{\boldsymbol{\ell}} - X^U_{\boldsymbol{\ell}}$ to first order in ϕ.

Lensing remaps the Stokes $T_{\hat{\mathbf{n}}}, Q_{\hat{\mathbf{n}}}, U_{\hat{\mathbf{n}}}$ maps in an identical fashion; however the lensing signal $\delta X_{\boldsymbol{\ell}}$ will not be the same in temperature vs. polarization since the polarization fields have different power spectra and thus different couplings to the lensing kernel. The effect of lensing on the curl- and divergence-free polarization basis is less straightforward; we can define the E and B lensing signal by combining Equation (2.12) for Q and U with the $QU \to EB$ basis transformation given in Equation (1.7) to obtain

$$\delta E_{\boldsymbol{\ell}} = \int \mathrm{d}\boldsymbol{\ell}'\, W_{\boldsymbol{\ell},\mathbf{L}} \big[E^U_{\boldsymbol{\ell}'} \cos 2\psi_{\boldsymbol{\ell},\boldsymbol{\ell}'} - B^U_{\boldsymbol{\ell}'} \sin 2\psi_{\boldsymbol{\ell},\boldsymbol{\ell}'}\big] \tag{2.13}$$

$$\delta B_{\boldsymbol{\ell}} = \int \mathrm{d}\boldsymbol{\ell}'\, W_{\boldsymbol{\ell},\mathbf{L}} \big[B^U_{\boldsymbol{\ell}'} \cos 2\psi_{\boldsymbol{\ell},\boldsymbol{\ell}'} + E^U_{\boldsymbol{\ell}'} \sin 2\psi_{\boldsymbol{\ell},\boldsymbol{\ell}'}\big], \tag{2.14}$$

where $\psi_{\boldsymbol{\ell},\boldsymbol{\ell}'} = \psi_{\boldsymbol{\ell}'} - \psi_{\boldsymbol{\ell}}$ (Hu & Okamoto, 2002). Hence lensing mixes E-modes and B-modes, adding power to the B-mode spectrum and making the search for primordial gravitational waves more difficult. Since any B-mode signal in the primordial CMB is expected to be several orders of magnitude smaller than the E-mode signal, lensing mainly mixes E into B.

The framing of lensing as Fourier-space convolution between a CMB field and a lensing kernel $W_{\boldsymbol{\ell},\mathbf{L}} = -\boldsymbol{\ell} \cdot \mathbf{L}\phi_L$ suggests a way to reconstruct the lensing potential by performing another Fourier-space convolution, this time between two lensed CMB fields. This the key idea behind the *quadratic estimator*, which estimates each mode $\mathbf{L}$ in the lensing reconstruction by performing a weighted average over pairs of modes in the CMB maps. In map space, we can think of this as picking out the correlated (de)magnifications of CMB anisotropies in the neighborhood of a given degree-scale lensing clump (Lewis & Challinor, 2006). The quadratic estimator was first derived by Hu & Okamoto (2002) and has become the standard tool for CMB lensing reconstruction, used by *Planck*, ACT, and SPT to map the matter distribution of the universe with great success. In the following section we will touch on some of the science applications of reconstructed lensing maps before defining the quadratic estimator in more detail.

2.3 CMB Lensing Science Cases

CMB lensing maps trace the integrated matter distribution of the universe between us and the surface of last scattering. The CMB lensing redshift kernel $W(z)$ describes the relative contributions of different redshifts to the projected matter density map κ, and extends to much higher redshifts than other tracers of large scale structure. As such CMB lensing is one of the best tools for understanding the cosmic structure at early times, complementing information from the primary CMB at recombination and galaxy surveys at late times. Since CMB lensing measures structure on large scales and at high redshifts where non-linear growth has not yet begun, it is also relatively insensitive to astrophysical biases than can affect galaxy surveys (Abazajian et al., 2016; Qu et al., 2024). This is especially true for polarization-only lensing reconstructions which are less sensitive to foreground biases. We discuss three major applications of CMB lensing measurements below: cosmological parameters from the power spectrum of the lensing maps, cross-correlations with other large scale structure tracers, and delensing.

Lensing Power Spectrum and Cosmological Parameters. The lensing power spectrum can constrain ΛCDM parameters independent of the primary CMB and also provides a complementary cross-check to primary CMB parameter estimates. Combining CMB lensing measurements with primary CMB power spectra accounts for the lensing-induced smoothing of the acoustic peaks at the likelihood level, allowing for a more precise measurement of the angular size of the sound horizon at recombination and the Hubble constant (see Hu et al., 2001, for a discussion of the relation between the CMB acoustic peaks and cosmological parameters). CMB lensing constrains a narrow band in σ_8-H_0-Ω_m space, and this degeneracy can be broken by combining with other datasets that trace the expansion history of the universe such as baryonic acoustic oscillations (Planck Collaboration et al., 2016a; Madhavacheril et al., 2024). CMB lensing can also inform extensions to the concordance ΛCDM model that modify the mean curvature or structure formation history of the universe (Pan et al., 2023; Madhavacheril et al., 2024). The matter power spectrum is damped on small scales by massive neutrinos, so lensing can dramatically improve estimates of the sum of the neutrino masses when compared to primary CMB measurements that are only mildly sensitive to this parameter (Kaplinghat et al., 2003; Lesgourgues & Pastor, 2006) Finally, CMB lensing constraints on structure growth parameter can be compared to independent constraints from galaxy surveys to shed light on the S_8 tension discussed in Chapter 1.2.

Cross-correlations. CMB lensing probes cosmological parameters and traces the matter distribution of the universe at intermediate redshifts, while optical galaxy surveys contain information from the lower-redshift universe. Cross-correlation studies between these datasets are a powerful tool that provide a more complete picture of cosmic evolution and break degeneracies between cosmological parameters. Since measurements of structure from CMB lensing and galaxy surveys use different types of input data, wavelengths, and estimators, sensitivity to systematic biases is often reduced in cross-correlation analyses. Along these lines, cross-correlations can be used to calibrate shear bias terms in galaxy weak lensing measurements (Vallinotto, 2012; Schaan et al., 2017). Cross-correlations can also inform ΛCDM model tensions between high-redshift and local-universe datasets (Kim et al., 2024). A common approach is to form combinations of the two-point functions between galaxy clustering, galaxy lensing, and CMB lensing (e.g. 5×2pt, 6×2pt, as performed in Abbott et al., 2023). Cross-correlations can also be used for tomographic studies, probing structure growth or non-Gaussianities in the matter power spectrum as a function of redshift (Zhang et al., 2023; Kim et al., 2024).

Delensing. Both the unlensed CMB and the CMB lensing potential are expected to be Gaussian, and thus fully described by their two-point power spectra. However, a measurement of the *lensed* CMB power spectrum and the lensing power spectrum is not sufficient to reconstruct all information in the primary CMB, even if the measurement errors are arbitrarily small. To understand this consider the "Knox formula" describing the variance of an unbiased estimate (e.g. $\langle C_\ell^{\mathrm{est}}\rangle = C_\ell^{\mathrm{true}}$) of a map's power spectrum (Knox, 1995),

$$\mathrm{var}\left(C_\ell^{\mathrm{est}}\right) = \left\langle \left(C_\ell^{\mathrm{est}} - C_\ell^{\mathrm{true}}\right)^2 \right\rangle = \frac{2}{2\ell+1}\left(C_\ell^{\mathrm{true}} + n^2 e^{\ell^2\sigma_b^2}\right)^2, \tag{2.15}$$

where n is a white noise level in e.g μK-arcmin and σ_b is the full width at half maximum (FWHM) of a Gaussian instrumental beam. Even if the white noise goes zero, the variance of a measurement on a given angular scale will be $2C_\ell^{\mathrm{true}}/(2\ell+1)$ and will increase with the amplitude of the true signal. This is known as sample variance (or cosmic variance for a full-sky measurement) and provides a fundamental limit on our ability to determine the "true" power spectrum of a map especially on large scales.

The lensed CMB power spectrum receives contributions from both the unlensed power spectrum and the lensing power spectrum, increasing the amount of sample variance in the measurement. However the sample variance from lensing can be reduced by undeflecting or *delensing* CMB maps with the observed realization of the lensing potential, thus tightening the error bars on the primary CMB power spectrum and improving constraints on parameters. As discussed earlier, this is especially important in the search for primordial gravitational waves where the lensing B-mode power spectrum is orders of magnitude larger than the predicted primordial signal. B-mode delensing is predicted to reduce the uncertainty on r by a factor of 2.5 with combined BICEP/*Keck* and SPT-3G data (BICEP/Keck Collaboration et al., 2021), and by more than an order of magnitude for the proposed CMB-S4 experiment[1] (Abazajian et al., 2016). Delensing can also be applied to T and E maps, shedding light on the primordial helium fraction, effective number of neutrino species, early dark energy

[1]The CMB-S4 project as proposed would include telescopes at both the South Pole and in the Chilean Atacama Desert, effectively combining and scaling up the SPT and ACT experiments. CMB-S4 was ranked as a top priority in the 2020 Decadal Survey and 2023 Particle Physics Project Prioritization Panel, but in May of this year the National Science Foundation announced that the CMB-S4 project would not be moved to the next design/funding stage due to a need to prioritize the recapitalization of critical infrastructure at the South Pole. While this decision was only driven by the South Pole component of the project and CMB-S4 may continue with a Chile-only plan, primordial gravitational wave science can only be done with the incredible observing conditions at the South Pole. This reinforces the importance of the joint BICEP/*Keck* and SPT-3G delensing program, which is not expected to be affected by the NSF decision.

models, and the physics of inflation (Green et al., 2017; Ange & Meyers, 2023; Trendafilova et al., 2024).

2.4 Reconstructing the Lensing Potential

Since the same lensing kernel $W_{\boldsymbol{\ell},\mathbf{L}}$ discussed in Chapter 2.2 is applied to all CMB fields, one can pick out this kernel by performing a weighted average over pairs of Fourier modes from different lensed fields. This is the idea behind the quadratic estimator (QE; Hu & Okamoto, 2002), which calculates the lensing potential by convolving pairs of inverse-variance filtered maps $\overline{X}_{\boldsymbol{\ell}}$ and $\overline{Y}_{\boldsymbol{\ell}}$ with a kernel $W^{XY}_{\boldsymbol{\ell},\mathbf{L}-\boldsymbol{\ell}}$ formed from theory power spectra,

$$\overline{\phi}^{XY}_{\mathbf{L}} = \int \mathrm{d}^2\boldsymbol{\ell}\, \overline{X}_{\boldsymbol{\ell}} \overline{Y}^*_{\boldsymbol{\ell}-\mathbf{L}} W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}. \tag{2.16}$$

where $XY \in \{TT, EE, TE, EB, TB\}$. This acts as a convolution of filtered maps in Fourier space, or a multiplication of the filtered maps in real space. The kernel $W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}$ can be thought of as a matched filter that maximizes the response of the quadratic estimator to the lensing-induced correlations in the inverse-variance filtered CMB maps.

Evaluating $\overline{\phi}^{XY}_{\mathbf{L}}$ can be quite costly and scales as $\mathcal{O}(\ell^6_{\max})$ if done as a convolution in Fourier space, since $\overline{\phi}_{\mathbf{L}}$, $\overline{X}_{\boldsymbol{\ell}}$, and $\overline{Y}_{\boldsymbol{\ell}}$ each have $\ell^2_{\max}$ Fourier pixels. However, $W_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}$ can be separated into a sum over N terms where each term W_i is factored as $W_i = J_i(\boldsymbol{\ell}) \cdot K_i(\boldsymbol{\ell}-\mathbf{L}) \cdot L_i(\mathbf{L})$. In this formulation, since $J_i(\boldsymbol{\ell})$ is only a function of the scales in map $\overline{X}$ and $K_i(\boldsymbol{\ell}-\mathbf{L})$ is only a function of the scales in map $\overline{Y}$, these components can be applied to the maps separately. Thus, instead of convolving terms $J_i(\boldsymbol{\ell})\overline{X}_{\boldsymbol{\ell}}$ and $K_i(\boldsymbol{\ell}-\mathbf{L})\overline{Y}_{\boldsymbol{\ell}-\mathbf{L}}$ in harmonic space, they can be inverse-Fourier transformed, multiplied in map space, and then Fourier transformed back to harmonic space before applying the $L_i(\mathbf{L})$ component of the kernel. Replacing parts of the convolution with Fourier transforms in this way improves cost from $\mathcal{O}(\ell^6_{\max})$ to $\mathcal{O}(N_i \cdot \ell^2_{\max} \cdot \log \ell_{\max})$, where N_i is the number of terms in the sum. This separation of the kernel also helps build an intuitive understanding of the quadratic estimator: it is the filtered product of two CMB maps with yet different filtering applied to each map.

In the remainder of this chapter I will introduce the ingredients of quadratic estimator lensing map and power spectrum reconstruction which will be used in Chapters 3 and 4. First, input CMB maps for an estimator XY must be filtered to maximize the signal to noise of the lensing estimator. The initial quadratic estimate of the lensing map is biased by a mean field, which must be subtracted at the map level, and the lensing map must also be properly normalized. Once the lensing *map* has been corrected we can estimate its power spectrum, which is also biased and must have several noise terms subtracted. The normalization, mean field bias, and power spectrum noise terms are all estimated from a suite of lensed sky simulations. The resulting debiased power spectrum can then be used to constrain cosmological parameters.

2.4.1 Inverse Variance Filtering

Maps of the millimeter sky from CMB experiments are not a direct measurement of the lensed CMB, and contain a host of other signals and effects including foregrounds, noise, instrumental beam, and map filtering. To obtain an unbiased estimate of the lensed CMB and maximize the signal-to-noise of the lensing reconstruction, it is desirable to inverse-variance filter the CMB maps that are fed to the quadratic estimator. The inverse-variance filter implementation used by SPT corrects for observational and filtering effects and then (in the simplest case) weights the map by inverse of the combined signal and noise variance, $\mathbb{C}^{-1} = (\mathbb{S} + \mathbb{N})^{-1}$. First we will consider the components that contribute to a single binned TQU pixel $d_{\hat{\mathbf{n}}}$ in a SPT map. The harmonic-space sky signal $X_{\boldsymbol{\ell}} \in \{T_{\boldsymbol{\ell}}, E_{\boldsymbol{\ell}}, B_{\boldsymbol{\ell}}\}$ is related to the TQU map values by the observation matrix operator $\mathbb{P}_{\boldsymbol{\ell}\to\hat{\mathbf{n}}}$ which includes beam, filtering, and pixelization effects, the $EB \to QU$ transformation, and an inverse Fourier transform to map space. We will use blackboard bold font (e.g. $\mathbb{B}$) to denote two-dimensional operators as well as Fourier-space objects that do not correspond to a map when inverse Fourier transformed, for example a covariance or transfer function. The data pixel values can then be modeled as

$$d_{\hat{\mathbf{n}}} = \mathbb{P}_{\boldsymbol{\ell}\to\hat{\mathbf{n}}} X_{\boldsymbol{\ell}} + \mathbb{P}_{\boldsymbol{\ell}\to\hat{\mathbf{n}}} N_{\boldsymbol{\ell}} + n_{\hat{\mathbf{n}}}, \tag{2.17}$$

where $N_{\boldsymbol{\ell}}$ is Fourier-space atmospheric noise and $n_{\hat{\mathbf{n}}}$ is pixel-space detector white noise assumed to be uncorrelated between pixels. The atmospheric noise is beamed and filtered in the same way as the true sky signal, whereas the detector noise is not beamed, but is still filtered and pixelized. However in our data model, we approximate the noise as unaffected by the observation matrix operator $\mathbb{P}_{\boldsymbol{\ell}\to\hat{\mathbf{n}}}$; while this is not strictly true and causes us to overestimate our noise level, doing so will only slightly reduce the optimality of the filter and does not bias the filtered maps or lensing reconstruction.

The pixel-space white noise $n_{\hat{\mathbf{n}}}$ captures spatially varying properties of the SPT maps such as the survey mask (beyond which the $n_{\hat{\mathbf{n}}}$ can be thought of as infinite) whereas the other two terms of $d_{\hat{\mathbf{n}}}$ are assumed to be diagonal in Fourier space. As a result, the inverse variance filter for $d_{\hat{\mathbf{n}}}$ is not diagonal in either pixel or Fourier space, and must be solved iteratively using a conjugate gradient solver. Before introducing the full filter, it can be instructive to consider the purely Fourier-diagonal approximation to the filter, which is also necessary for the lensing estimator normalization discussed later. If we replace the spatially varying white noise with a constant noise level n_ℓ, and the observation matrix operator $\mathbb{P}_{\boldsymbol{\ell}\to\hat{\mathbf{n}}}$ with a transfer function $\mathbb{T}_{\boldsymbol{\ell}}$, we can write the simplified Fourier-diagonal data model as

$$d_{\boldsymbol{\ell}} = \mathbb{T}_{\boldsymbol{\ell}} X_{\boldsymbol{\ell}} + \mathbb{T}_{\boldsymbol{\ell}} N_{\boldsymbol{\ell}} + n_\ell. \tag{2.18}$$

Since the transfer function (and the total observation matrix) operate on the sky signal $X_{\boldsymbol{\ell}}$ and atmospheric noise $N_{\boldsymbol{\ell}}$ in the same way, we need not differentiate between them in the construction of this filter and can define the total sky variance $\mathbb{S}_{\boldsymbol{\ell}}$ as the power spectrum of $X_{\boldsymbol{\ell}} + N_{\boldsymbol{\ell}}$. To obtain the inverse-variance filtered sky $\overline{X}_{\boldsymbol{\ell}}$ we can deconvolve the transfer function from the data, and then divide by the sum of the on-sky variance and transfer-function-deconvolved white noise, $\mathbb{C}_{\boldsymbol{\ell}} = \mathbb{S}_{\boldsymbol{\ell}} + \mathbb{T}_{\boldsymbol{\ell}}^{-2} n$:

$$\overline{X}_{\boldsymbol{\ell},\text{diagonal}} = \frac{\mathbb{T}_{\boldsymbol{\ell}}^{-1} d_{\boldsymbol{\ell}}}{\mathbb{S}_{\boldsymbol{\ell}} + \mathbb{T}_{\boldsymbol{\ell}}^{-2} n} \tag{2.19}$$

$$= \mathbb{C}_{\boldsymbol{\ell}}^{-1} \mathbb{T}_{\boldsymbol{\ell}}^{-1} d_{\boldsymbol{\ell}}. \tag{2.20}$$

This makes it easy to see that the Fourier-diagonal inverse-variance filter $\mathbb{F}_{\boldsymbol{\ell}}$ is simply the inverse of the transfer function and total effective variance $\mathbb{C}_{\boldsymbol{\ell}}$:

$$\mathbb{F}_{\boldsymbol{\ell}} = \mathbb{C}_{\boldsymbol{\ell}}^{-1} \mathbb{T}_{\boldsymbol{\ell}}^{-1}. \tag{2.21}$$

An optimal signal-to-noise filtered map can then be constructed by multiplying by the desired *signal* variance, for example that of the lensed CMB. This is commonly known as a Wiener filter after American mathematician and computer scientist Norbert Wiener, who derived this filter in a classified 1942 document aiming to improve wartime radar communication (Wiener, 1949) although an independent derivation of this filter was published in 1941 by Soviet mathematician Andrey Kolmogorov (Kolmogorov, 1941). The Wiener-filtered map will almost perfectly reconstruct the signal in the signal-dominated regime, and be down-weighted in the noise dominated regime.

In the presence of spatially varying noise, the inverse variance filter is more complicated and must be solved iteratively. We solve the following equation for the inverse variance filtered field $\overline{X}$:

$$\left[\mathbb{C}^{-1} + \mathbb{P}^{\dagger} n \mathbb{P}\right] \mathbb{C}\overline{X} = \mathbb{P}^{\dagger} n^{-1} d \tag{2.22}$$

where we have dropped the Fourier and pixel-space subscripts for readability. Note that we are now dividing out the spatially varying noise level from the data as well as undoing the filtering, pixelization, and $EB \rightarrow QU$ operations captured in $\mathbb{P}$. The operators acting on $\overline{X}$ are difficult to invert as they mix pixel- and Fourier-space operations and thus are not diagonal in any basis, so we solve this equation iteratively using a conjugate gradient method.

In practice, both the atmospheric noise $\mathbb{N}_{\boldsymbol{\ell}}$ and detector white noise $n_{\hat{\mathbf{n}}}$ are determined from "signflip" noise realizations estimated by combining different observations in a way that cancels out any constant signal on the sky. This is done by randomly permuting the full set of observations, splitting the permuted observations into two bundles with approximately equal weights, and then subtracting the two bundles to form a noise-only difference map. This process can be repeated many times to form many nearly-independent noise realizations which contain all the filtering and pixelization effects discussed above. The scalar white noise level is estimated as the mean power at high ℓ to avoid contributions from the $1/f$ atmospheric noise that dominates on large scales; the inverse of this noise level is then multiplied by the boundary mask to capture the survey geometry and form n in Equation (2.22). The white noise level is then subtracted off from the noise power spectrum, leaving only the ℓ-dependent atmospheric noise, which is then corrected for the by the transfer function to obtain the atmospheric component of $\mathbb{C}$. The remaining components of $\mathbb{C}$ are the lensed CMB and foreground power spectra.

2.4.2 Lensing Map Reconstruction

The raw quadratic estimate $\bar{\phi}^{XY}_{\mathbf{L}}$ in Equation (2.16) is an inverse-variance weighted estimate of the lensing potential and must be properly scaled by an estimate of the estimator variance to obtain a useful lensing map. The variance is usually calculated as the inverse of an analytic response function $\mathbb{R}^{XY,\text{Analytic}}_{\mathbf{L}}$ which is then used to normalize the estimator. The analytic response takes as input the Fourier-diagonal approximation of the inverse-variance filter $\mathbb{F}_{\boldsymbol{\ell}}$ given in Equation (2.21):

$$\mathbb{R}^{XY,\text{Analytic}}_{\mathbf{L}} = \int \mathrm{d}^2\boldsymbol{\ell}\, W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}} W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}} \mathbb{F}^{X}_{\boldsymbol{\ell}} \mathbb{F}^{Y}_{\boldsymbol{\ell}-\mathbf{L}}. \tag{2.23}$$

This analytic response is similar to the quadratic estimate of Equation (2.16), but with an extra factor of the QE kernel $W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}$ and the filtered maps replaced by the diagonal approximation filter to the filter $\mathbb{F}_{\boldsymbol{\ell}}$, yielding the inverse of the estimator variance.

The Fourier-diagonal approximation to the inverse-variance filter $\mathbb{F}_{\boldsymbol{\ell}}$ is particularly inaccurate for the $1500\,\mathrm{deg}^2$ SPT-3G patch given flat-sky projection effects in ProjZEA. As a result, we instead calculate a realization-dependent semi-analytic (RDSA) response function by replacing $\mathbb{F}^X_{\boldsymbol{\ell}}, \mathbb{F}^Y_{\boldsymbol{\ell}}$ with the power spectra of the inverse-variance filtered maps $\overline{X}_{\boldsymbol{\ell},i}, \overline{Y}_{\boldsymbol{\ell},i}$ for a given simulation realization i:

$$\mathbb{R}^{XY,\mathrm{RDSA}}_{\mathbf{L},i} = \int \mathrm{d}^2\boldsymbol{\ell}\, W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}} W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}} \left\langle \overline{X}_{\boldsymbol{\ell},i}\overline{X}^*_{\boldsymbol{\ell},i} \right\rangle \left\langle \overline{Y}_{\boldsymbol{\ell}-\mathbf{L},i}\overline{Y}^*_{\boldsymbol{\ell}-\mathbf{L},i} \right\rangle . \tag{2.24}$$

We then calculate a *mean* semi-analytic response by averaging over simulations:

$$\mathbb{R}^{XY,\mathrm{SA}}_{\mathbf{L}} = \left\langle \mathbb{R}^{XY,\mathrm{RDSA}}_{\mathbf{L},i} \right\rangle . \tag{2.25}$$

We choose to do this in two steps rather than calculating $\mathbb{R}^{\mathrm{SA}}_{\mathbf{L}}$ directly from the ensemble-averaged power spectra of the inverse-variance filtered maps because the realization-dependent $\mathbb{R}^{\mathrm{RDSA}}_{\mathbf{L},i}$ will be useful in the power spectrum estimation step described later. The semi-analytic approach described above is slightly different than the analytic approach employed by other quadratic estimator analyses, but we find that it does a better job of capturing projection effects in the pipeline and slightly reduces the lensing power spectrum error bars.

It is more convenient and intuitive to work with the lensing convergence $\kappa = -\nabla^2\phi/2$ which traces the integrated mass along the line of sight; we can easily convert from ϕ to κ in Fourier space by applying a factor of $L(L+1)/2$ to the lensing map. Lensing reconstructions can be formed from many combinations of CMB fields ($TT, TE, ET, TB, BT, EE, EB, BE$; there are two different estimators XY and YX for a given set of input maps), and it is desirable to combine these estimators to form a minimum variance estimate of the lensing map. Since each estimator $\overline{\phi}^{XY}_{\mathbf{L}}$ is inverse-variance filtered before the response is applied, we can simply sum the estimators and normalize by the sum of the responses. Putting all of this together, we have the following expression for the normalized minimum-variance[2] map:

$$\tilde{\kappa}_{\mathbf{L}} = \frac{L(L+1)}{2} \frac{\sum_{XY} \overline{\phi}^{XY}}{\sum_{XY} \mathbb{R}_{\mathbf{L}}^{XY,\mathrm{SA}}}. \tag{2.26}$$

Typically we produce two minimum-variance estimators,

$$\tilde{\kappa}_{\mathbf{L}}^{\mathrm{MV}} \in \{TT, TE, ET, TB, BT, EE, EB, BE\}, \tag{2.27}$$

$$\tilde{\kappa}_{\mathbf{L}}^{\mathrm{POL}} \in \{EE, EB, BE\}. \tag{2.28}$$

The minimum-variance polarization estimator tends to be less sensitive to foregrounds since polarized foregrounds are expected to be negligible on the scales relevant to lensing in the SPT patch, and can produce a cleaner lensing map without relying on foreground modeling or bias-mitigation techniques.

The convergence map, while normalized, is still biased by the lensing estimator's sensitivity to statistical anisotropies caused by effects other than lensing, including the survey mask and inhomogeneous map noise. We can correct for this "mean-field" bias provided realistic simulations by averaging together a suite of lensing reconstructions, where each reconstruction is formed from a different realization of the CMB and lensing potential:

$$\overline{\kappa}_{\mathbf{L}}^{\mathrm{MF}} = \langle \tilde{\kappa}_{\mathbf{L}} \rangle. \tag{2.29}$$

Each lensing reconstruction is a combination of the true lensing signal, reconstruction noise and the mean-field; only the latter remains constant across realizations. As a result, the true lensing signal and the reconstruction noise average down over many realizations as $1/\sqrt{N_{\mathrm{sims}}}$ in the map, allowing the mean-field to be precisely estimated given a sufficient

[2]Technically, this is only the inverse-variance combination and is not formally minimum-variance as it ignores TE correlations; Maniyar et al. (2021) present a more optimal global-minimum-variance (GMV) estimator properly handles the TE case. Nevertheless we will refer to the inverse-variance combination as the minimum-variance estimator here for simplicity.

number of simulations. The accuracy of the mean-field can be improved for a fixed number of simulations by subtracting the true convergence $\kappa_{\mathbf{L}}^{\rm Truth}$ from each reconstruction, leaving only the reconstruction noise as a source of uncertainty in the mean field (assuming the sims capture all the effects in the data):

$$\tilde{\kappa}_{\mathbf{L}}^{\rm MF} = \left\langle \tilde{\kappa}_{\mathbf{L}} - \kappa_{\mathbf{L}}^{\rm Truth} \right\rangle . \tag{2.30}$$

When subtracting the mean-field from the reconstruction of a given realization $\tilde{\kappa}_{\mathbf{L},i}$, one must be careful to not include that realization in the mean-field average. Some lensing analyses split the N available simulations into two batches—one batch for mean field estimation and one batch for pipeline calibration and validation—but doing increases the uncertainty in the mean field estimate and reduces the analysis' sensitivity to small biases in the pipeline. An alternative approach is to subtract off realization i's contribution to the mean field:

$$\tilde{\kappa}_{\mathbf{L},i}^{\rm MF} = \frac{N}{N-1}\tilde{\kappa}_{\mathbf{L}}^{\rm MF} - \tilde{\kappa}_{\mathbf{L},i} + \kappa_{\mathbf{L},i}^{\rm Truth}. \tag{2.31}$$

This allows the analysis to use all N simulations for calibration and validation, and $N-1$ simulations for mean field estimation.

After subtracting the mean field, the lensing map will be unbiased, but the analytic response calculation is based on a number of simplifying assumptions and is unlikely to capture all sources of variance in the estimator when applied to realistic maps. As a result, we apply an additional multiplicative Monte Carlo (MC) correction, estimated by cross-correlating the mean-field subtracted lensing reconstruction $\tilde{\kappa}_{\mathbf{L}}^{\rm MFsub}$ with the input convergence maps:

$$\mathbb{R}_{\mathbf{L}}^{\rm MC} = \frac{\left\langle \tilde{\kappa}_{\mathbf{L}}^{\rm MFsub} \kappa_{\mathbf{L}}^{\rm Truth*} \right\rangle}{\left\langle \kappa_{\mathbf{L}}^{\rm Truth} \kappa_{\mathbf{L}}^{\rm Truth*} \right\rangle}. \tag{2.32}$$

Having defined the mean field and MC corrections, we can now construct an unbiased and properly normalized lensing map

$$\hat{\kappa}_{\mathbf{L}} = \frac{\tilde{\kappa}_{\mathbf{L}} - \tilde{\kappa}_{\mathbf{L}}^{\mathrm{MF}}}{\mathbb{R}_{\mathbf{L}}^{\mathrm{MC}}}. \tag{2.33}$$

In summary, the steps in this process are:

1. evaluate the raw QE on inverse-variance filtered maps,

2. normalize by the analytic response,

3. convert from ϕ to κ,

4. form a minimum-variance combination of the estimators,

5. remove the mean-field bias,

6. apply a multiplicative MC correction.

2.4.3 Foreground Bias-Hardening

CMB temperature maps contain signals from a variety of extragalactic foregrounds that can bias lensing reconstructions including point sources, thermal and kinetic Sunyaev Zel'dovich effects (tSZ and kSZ), and the cosmic infrared background (van Engelen et al., 2013; Osborne et al., 2013; Ferraro & Hill, 2018; Baxter et al., 2019; Omori, 2024). Diffuse galactic foregrounds may also be present, but are expected to be negligible in the SPT-3G patch on the scales relevant for lensing. Non-Gaussian foregrounds can be picked up by the lensing estimator and introduce systematic biases to the lensing map. Since these foregrounds trace large scale structure in the universe, the bias they contribute to the lensing reconstruction will be partially correlated with the true lensing convergence. This can in turn can affect the results of joint cross-correlation analyses of galaxy clustering, galaxy lensing, and CMB lensing (3×2 point; Baxter et al., 2019). Furthermore, the reconstructed CMB lensing power spectrum will receive contributions from both the trispectrum of the foregrounds and the bispectrum between foregrounds and the lensing convergence.

Various methods can be used to mitigate these biases, including inpainting or masking of points sources and galaxy clusters (van Engelen et al., 2012), and multi-frequency cleaning that leverages the spectral dependence of different foregrounds (Cardoso et al., 2008; Madhavacheril & Hill, 2018; Darwish et al., 2021). While these methods can greatly reduce the amount of foreground contamination in the input maps, they come with a reduction of useable sky area in the case of masking/inpainting, and may not completely null the foregrounds in the case of multi-frequency cleaning. Another approach, which can be combined with the previously mentioned techniques, is to modify the lensing estimator to be less sensitive to foregrounds. This technique, known as bias-hardening, was first introduced as a method to mitigate the mean-field bias in Namikawa et al. (2013) and was extended to foreground biases in Osborne et al. (2013). The basic idea of the bias-hardened lensing estimator is to construct a new estimator $\tilde{s}^{TT}_{\mathbf{L}}$ sensitive to foregrounds modeled as a collection of sources with identical profiles, and project the configurations sensitive to that profile out of the standard quadratic estimator. In the case of point sources, the foreground estimator amounts to simply squaring the input map.

The standard convergence and foreground quadratic estimators each acquire a bias from the other as they are both sensitive to the same configurations $\boldsymbol{\ell}, \boldsymbol{\ell}'$. Here we follow the implementation of Sailer et al. (2020), where this bias to the true fields $\kappa_{\mathbf{L}}$ and $s_{\mathbf{L}}$ is given by

$$\begin{pmatrix} \left\langle \tilde{\kappa}^{TT}_{\mathbf{L}} \right\rangle \\ \left\langle \tilde{s}^{TT}_{\mathbf{L}} \right\rangle \end{pmatrix} = \begin{pmatrix} 1 & \mathbb{R}^{TT,\kappa\kappa}_{\mathbf{L}} \mathbb{R}^{TT,\kappa s}_{\mathbf{L}} \\ \mathbb{R}^{TT,ss}_{\mathbf{L}} \mathbb{R}^{TT,\kappa s}_{\mathbf{L}} & 1 \end{pmatrix} \begin{pmatrix} \kappa_{\mathbf{L}} \\ s_{\mathbf{L}} \end{pmatrix} \tag{2.34}$$

where $\mathbb{R}^{TT,\kappa\kappa}_{\mathbf{L}}$ and $\mathbb{R}^{TT,ss}_{\mathbf{L}}$ are the standard analytic response functions given each estimator's weights, and $\mathbb{R}_{\mathbf{L}}^{TT,\kappa s}$ is the cross-response between the convergence and foreground estimators:

$$\mathbb{R}_{\mathbf{L}}^{TT,\kappa s} = \int \mathrm{d}^2\boldsymbol{\ell}\, W_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}^{TT,\kappa} W_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}^{TT,s} \mathbb{F}_{\boldsymbol{\ell}}^{T} \mathbb{F}_{\boldsymbol{\ell}-\mathbf{L}}^{T}. \tag{2.35}$$

Equation (2.34) can then be solved for the bias-hardened estimators $\hat{\kappa}_{\mathbf{L}}^{TT,\mathrm{BH}}$ and $\hat{s}_{\mathbf{L}}^{TT,\mathrm{BH}}$:

$$\begin{pmatrix} \hat{\kappa}_{\mathbf{L}}^{TT,\mathrm{BH}} \\ \hat{s}_{\mathbf{L}}^{TT,\mathrm{BH}} \end{pmatrix} = \begin{pmatrix} 1 & \mathbb{R}_{\mathbf{L}}^{TT,\kappa\kappa}\mathbb{R}_{\mathbf{L}}^{TT,\kappa s} \\ \mathbb{R}_{\mathbf{L}}^{TT,ss}\mathbb{R}_{\mathbf{L}}^{TT,\kappa s} & 1 \end{pmatrix}^{-1} \begin{pmatrix} \tilde{\kappa}_{\mathbf{L}}^{TT} \\ \tilde{s}_{\mathbf{L}}^{TT} \end{pmatrix}. \tag{2.36}$$

The bias-hardened estimator does have slightly increased noise relative to the standard quadratic estimator, since the standard quadratic estimator weights are designed to minimize variance. The bias-hardened noise $\mathbb{N}_{\mathbf{L}}^{TT,\mathrm{BH}}$ is related to the standard quadratic estimator noise $\mathbb{N}_{\mathbf{L}}^{TT}$ by the determinant of the κ-s mixing matrix given in Equation (2.34):

$$\mathbb{N}_{\mathbf{L}}^{TT,\mathrm{BH}} = \frac{\mathbb{N}_{\mathbf{L}}^{TT}}{1 - \mathbb{R}_{\mathbf{L}}^{TT,\kappa\kappa}\mathbb{R}_{\mathbf{L}}^{TT,ss}\left(\mathbb{R}_{\mathbf{L}}^{TT,\kappa s}\right)^2} \tag{2.37}$$

We construct a source-hardened version of the minimum-variance estimate given in Equation (2.23) by subtracting off the standard TT estimator and replacing it with the bias-hardened TT estimator following Qu et al. (2024):

$$\hat{\kappa}_{\mathbf{L}}^{MV,\mathrm{BH}} = \frac{\mathbb{R}_{\mathbf{L}}^{MV,\kappa\kappa}\hat{\kappa}_{\mathbf{L}}^{MV} - \mathbb{R}_{\mathbf{L}}^{TT,\kappa\kappa}\hat{\kappa}_{\mathbf{L}}^{TT} + \mathbb{R}_{\mathbf{L}}^{TT,\kappa\kappa}\hat{\kappa}_{\mathbf{L}}^{TT,BH}}{\mathbb{R}_{\mathbf{L}}^{MV,\kappa\kappa}} \tag{2.38}$$

Any profile can be hardened against with this method; we use a simple point source profile for the SPT-3G lensing analysis. Since the profile is convolved with the instrumental beam, point-source hardening also hardens against cluster profiles to some extent.

2.4.4 Lensing Power Spectrum Estimation

In order to relate the lensing map $\hat{\kappa}_{\mathbf{L}}$ to cosmological parameters, we take its auto-spectrum and form biased lensing power spectra,

$$C_{\mathbf{L}}^{\hat{\kappa}\hat{\kappa}} = f_{\text{mask},4}^{-1} \left\langle \hat{\kappa}_{\mathbf{L}} \hat{\kappa}_{\mathbf{L}}^{*} \right\rangle . \tag{2.39}$$

Here $f_{\text{mask},4}$ corrects for the masked sky fraction, and is calculated by raising the survey mask to the fourth power and taking its mean. This differs from the standard CMB power spectrum sky fraction normalization $f_{\text{mask},2}$ which only squares the mask and takes the mean, because the quadratic estimate is formed from the *product* of two filtered maps.

Because the quadratic estimate is a 2-point function of the input CMB maps, its power spectrum is a 4-point function, or trispectrum, of the CMB. The CMB trispectrum contains both *disconnected* terms that can be expressed in terms of products of 2-point functions and *connected* terms that are pure 4-point functions. While the CMB lensing signal contributes a connected term to the trispectrum, there remain both connected and disconnected noise terms that do not contain any lensing signal but still contribute to the variance of the quadratic estimator. The two lowest-order terms are known as $\mathbb{N}_{\mathbf{L}}^{0}$ and $\mathbb{N}_{\mathbf{L}}^{1}$, so named because they are zeroth- and first-order in the lensing potential respectively (Kesden et al., 2003). A $\mathbb{N}_{\mathbf{L}}^{2}$ bias also exists if unlensed CMB power spectra are used in the quadratic estimator as initially derived by Okamoto & Hu (2003), but this bias can be easily eliminated by using the lensed CMB power spectra instead (Hanson et al., 2011; Lewis et al., 2011; Anderes, 2013). Under the assumption that the CMB lensing potential is Gaussian, the next-highest order term scales as $\mathcal{O}([C_L^{\phi\phi}]^3)$ and may be discounted; however if the lensing potential is allowed to be non-Gaussian and/or post-Born approximations are included, a $\mathbb{N}_{\mathbf{L}}^{3/2}$ bias appears (Böhm et al., 2016; Pratten & Lewis, 2016). This bias can approach 1 % of the lensing power spectrum on large scales in the TT estimator, but is negligible for polarization estimators and can be safely ignored given the relative signal to noise from temperature and polarization estimators for SPT-3G data (Böhm et al., 2018; Beck et al., 2018). The $\mathbb{N}^0_{\mathbf{L}}$ and $\mathbb{N}^1_{\mathbf{L}}$ terms do contribute significantly to the variance of the quadratic estimators, and thus bias the reconstructed lensing power spectrum if not properly accounted for; I will outline the procedure for debiasing the lensing power spectrum below.

The disconnected $\mathbb{N}^0_{\mathbf{L}}$ bias arises from chance 2-point (Gaussian) correlations between Gaussian CMB, foregrounds and noise fluctuations in the pairs of input maps; this bias is zeroth-order in the lensing potential and would exist even in the trispectrum of the unlensed CMB, although it has an implicit dependence on the lensing potential via the lensed CMB power spectrum (Kesden et al., 2003). In ideal cases the $\mathbb{N}^0_{\mathbf{L}}$ bias is the inverse of the quantity predicted by the analytic response calculation in Equation (2.23); however, in a similar spirit to the multiplicative MC bias, we estimate the additive $\mathbb{N}^0_{\mathbf{L}}$ bias from simulations that capture the full complexity of the data. We do so by cross-correlating two lensing maps formed from different CMB and lensing potential realizations, ensuring that all power in the cross-spectrum of these maps is due to chance Gaussian correlations rather than from lensing. To make the notation more readable, we define (U_1, V_2) to be a lensing reconstruction formed from different CMB realizations U and V lensed by different lensing potentials 1 and 2. Then the cross-spectrum of these lensing reconstructions is written as $C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1, V_2)\ (V_2, U_1)]$. Using this notation, the general form of the $\mathbb{N}^0_{\mathbf{L}}$ bias is

$$\mathbb{N}^0_{\mathbf{L}} = \Big\langle + C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1, V_2)\ (U_1, V_2)] \\ + C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1, V_2)\ (V_2, U_1)] \Big\rangle_{U,V}. \tag{2.40}$$

For individual asymmetric estimators e.g. TE, TB, EB, the two terms above differ, but for TT, EE and the minimum variance combinations, symmetry is preserved and the two terms are identical. As a result, in practice we calculate $\mathbb{N}^0_{\mathbf{L}} = 2\langle C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1, V_2)\ (U_1, V_2)]\rangle_{U,V}$ for the minimum-variance estimator to reduce the computational cost of the bias estimation.

While the $\mathbb{N}^0_{\mathbf{L}}$ estimated from simulations will likely be a better estimate of the disconnected noise bias than the analytic response, it relies on perfect agreement between simulations and data, which is difficult to achieve in practice. To relax this requirement, one can instead use a "realization-dependent" estimate $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}}$ (Namikawa et al., 2013), which replaces a simulation with the data D in one leg of each quadratic estimate, and then subtracts the standard $\mathbb{N}^{0}_{\mathbf{L}}$ bias:

$$\begin{aligned}\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}} = \Big\langle & + C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(D,U_1)\ (D,U_1)] \\ & + C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,D)\ (U_1,D)] \\ & + C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,D)\ (D,U_1)] \\ & + C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(D,U_1)\ (U_1,D)] \\ & - C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_2)\ (U_1,V_2)] \\ & - C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_2)\ (V_2,U_1)] \Big\rangle_{U,V}. \end{aligned} \tag{2.41}$$

Again the symmetry of the minimum-variance estimators allows us to simplify this expression to $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}} = 4\left\langle C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(D,U_1)\ (D,U_1)]\right\rangle_U - \mathbb{N}^{0}_{\mathbf{L}}$, reducing the computational cost of the $\mathbb{N}^{0\mathrm{RD}}_{\mathbf{L}}$ bias by a factor of 6. This realization-dependent calculation mitigates any discrepancies between the simulations and data, including from assumed cosmology, foreground models, and sample variance of the data. In addition, the realization-dependent approach greatly reduces the covariance of the lensing power spectrum reconstruction for certain estimators especially on small scales, effectively reducing the error bars on the bandpowers (Hanson et al., 2011; Nguyên et al., 2019). $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}}$ can also reduce the correlated error in joint estimates of the lensing potential amplitude from lensing power spectrum and CMB temperature power spectrum peak-smoothing (Schmittfull et al., 2013).

As discussed above the $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}}$ calculation can reduce the covariance of the estimated data lensing power spectrum. It is desirable to apply this calculation to simulations as well in order to properly capture this reduction in the covariance, but doing so is very computationally expensive since each lensing reconstruction must be cross-correlated with the reconstructions of all other simulations, leading to an $\mathcal{O}(N_{\rm sims}^2)$ scaling compared to the $\mathcal{O}(N_{\rm sims})$ scaling of the other bias calculations. A compromise is to use a semi-analytic estimate of the $\mathbb{N}_{\mathbf{L}}^{\rm RD}$ bias (Planck Collaboration et al., 2016a; Qu et al., 2024) that only considers Fourier-diagonal (power spectrum) information in the simulations and ignores the full set of off-diagonals used in the quadratic estimator calculation. To do this we take the realization-dependent semi-analytic response $\mathbb{R}_{\mathbf{L},i}^{\rm RDSA}$ discussed in the previous section and normalize it by the mean semi-analytic response $\mathbb{R}_{\mathbf{L}}^{\rm SA}$ to get a quantity that behaves like the noise of the estimator:

$$\mathbb{N}_{\mathbf{L},i}^{0,\rm RDSA} = \frac{\mathbb{R}_{\mathbf{L},i}^{\rm RDSA}}{f_{\rm mask,2}^2 (\mathbb{R}_{\mathbf{L}}^{\rm SA})^2}, \tag{2.42}$$

where we have normalized by the 2-point mask correction since the semi-analytic response is calculated from the power spectra of the filtered maps. The semi-analytic $\mathbb{N}_{\mathbf{L},i}^{0,\rm RDSA}$ does not capture the full set of chance correlations as the full $\mathbb{N}_{\mathbf{L}}^{0}$ or $\mathbb{N}_{\mathbf{L}}^{0,\rm RD}$ calculations do, and is a less accurate estimate of the disconnected noise bias as a result. To mitigate this, we do not directly subtract $\mathbb{N}_{\mathbf{L},i}^{0,\rm RDSA}$ from each lensing reconstruction's power spectrum but instead subtract the standard $\mathbb{N}_{\mathbf{L}}^{0}$ combined with a residual defined as that simulation's $\mathbb{N}_{\mathbf{L}}^{0,\rm RDSA}$ estimate minus the mean $\mathbb{N}_{\mathbf{L}}^{0,\rm RDSA}$:

$$\Delta\mathbb{N}_{\mathbf{L},i}^{0,\rm RDSA} = \mathbb{N}_{\mathbf{L},i}^{0,\rm RDSA} - \langle \mathbb{N}_{\mathbf{L}}^{0,\rm RDSA} \rangle. \tag{2.43}$$

This reduces the mis-estimation of the overall size of the $\mathbb{N}_{\mathbf{L}}^{0,\rm RDSA}$ bias while still accounting from sim-to-sim scatter and thus reducing the covariance of the simulated lensing power spectrum.

The final significant bias to the lensing power spectrum is the connected $\mathbb{N}_{\mathbf{L}}^{1}$ term in the trispectrum, which is sourced by chance correlations between the CMB and the lensing potential. The $\mathbb{N}_{\mathbf{L}}^{1}$ bias is linear in the lensing potential and its analytic expression involves an integral over the lensing power spectrum; as such the $\mathbb{N}_{\mathbf{L}}^{1}$ noise in a given multipole $\mathbf{L}$ is sourced from the lensing signal at other multipoles $\mathbf{L}' \neq \mathbf{L}$ (Kesden et al., 2003). This bias can be estimated by cross-correlating lensing maps formed from *different* CMB realizations lensed by the *same* lensing potential:

$$\begin{aligned}
\mathbb{N}^1_{\mathbf{L}} = \Big\langle &+ C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_1)\ (U_1,V_1)] \\
&+ C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_1)\ (V_1,U_1)] \\
&- C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_2)\ (U_1,V_2)] \\
&- C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_2)\ (V_2,U_1)] \Big\rangle_{U,V},
\end{aligned} \tag{2.44}$$

where the final two lines are equivalent to the $\mathbb{N}^0_{\mathbf{L}}$ bias calculation. For symmetric minimum-variance combinations the first two terms are again identical, and the $\mathbb{N}^1_{\mathbf{L}}$ bias can be calculated with fewer evaluations as $\mathbb{N}^1_{\mathbf{L}} = 2\langle C^{\hat{\kappa}\hat{\kappa}}_{\mathbf{L}}[(U_1,V_1)\ (U_1,V_1)]\rangle_{U,V} - \mathbb{N}^0_{\mathbf{L}}$. Equation (2.44) shows that the $\mathbb{N}^1_{\mathbf{L}}$ bias is independent of the input CMB map noise and foregrounds since the $\mathbb{N}^0_{\mathbf{L}}$ bias differences out; as such one can obtain faster convergence if the $\mathbb{N}^1_{\mathbf{L}}$ bias is estimated from lensed CMB-only simulations. The $\mathbb{N}^1_{\mathbf{L}}$ bias is only correct for the exact lensing power spectrum used in the simulations, and introduces a cosmological model dependence that must be marginalized over when performing parameter estimation.

In summary, the quadratic estimator lensing power spectrum is biased by one disconnected noise term $\mathbb{N}^0_{\mathbf{L}}$. A more accurate realization-dependent estimate of the disconnected term $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}}$ is used for the data lensing power spectrum, while a semi-analytic estimate of the $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L},i}$ is used for the simulations; both reduce the covariance of the respective data and simulation lensing power spectra. A smaller connected $\mathbb{N}^1_{\mathbf{L}}$ bias linear in the lensing potential is also accounted for. Putting this all together we have the following expressions for the data and simulation power spectra estimates:

$$\widehat{C}^{\kappa\kappa}_{L,\,\mathrm{data}} = C^{\hat{\kappa}\hat{\kappa}}_{L} - \mathbb{N}^{0,\mathrm{RD}}_{L} - \mathbb{N}^1_{L}. \tag{2.45}$$

$$\widehat{C}^{\kappa\kappa}_{L,\,\mathrm{sim}} = C^{\hat{\kappa}\hat{\kappa}}_{L} - \mathbb{N}^0_{L} - \Delta\mathbb{N}^{0,\mathrm{RDSA}}_{L} - \mathbb{N}^1_{L}. \tag{2.46}$$

One final note is that due to non-idealities in the pipeline, the lensing power spectrum may not be perfectly debiased even after these corrections. Especially with a large number of simulations, arbitrarily small biases can be detected. A possible source of residual bias is masking of point sources and clusters (Wu et al., 2019; Qu et al., 2024). Mode-coupling from masking can reduce the cross correlation with the input convergence in the MC response correction $\mathbb{R}^{\mathrm{MC}}_{\mathbf{L}}$ (Equation (2.32)); for example in the presence of flat mode mixing where 1% of each lensing mode mixes into each other mode equally, the cross-correlation with the input convergence will be low by a roughly the same amount. This will bias the estimated lensing power spectrum $\widehat{C}^{\kappa\kappa}_{L,\mathrm{data}}$ high by 2% via the MC response correction. There are various way to correct for this; Wu et al. (2019) apply a final $\sim$5% multiplicative correction f_{PS}, whereas Qu et al. (2024) introduce a final additive L-dependent bias term.

Chapter 3

SPTpol Lensing

Until recently all published lensing measurements have used the quadratic estimator, which is nearly minimum-variance at the instrumental noise levels of previous SPTpol $100\,\mathrm{deg}^2$ (Wu et al., 2019) and $500\,\mathrm{deg}^2$ (Story et al., 2015) analyses. The quadratic estimator is a frequentist approach, producing a point estimate of the lensing potential given a set of simulations generated with fiducial parameter values. However the quadratic estimator becomes suboptimal as the polarization noise level falls below the $\sim 5\,\mu$K-arcmin effective B-mode noise floor from lensing. This was realized soon after the quadratic estimator was first proposed (Hirata & Seljak, 2003a,b; Seljak & Hirata, 2004), and Hirata & Seljak (2003b) first demonstrated an improved estimator based on maximizing the Bayesian CMB lensing posterior. Carron & Lewis (2017) revisited the construction of more optimal estimators with a maximum a posteriori (MAP) estimator based on iterative quadratic estimates applicable to realistic CMB datasets. The POLARBEAR collaboration (Adachi et al., 2020) first demonstrated the use of these techniques on real data, using the iterative MAP ϕ estimate to delens B-mode polarization maps.

While MAP approaches are more optimal than the traditional quadratic estimator, it is also more challenging to estimate and interpret their power spectrum; like the quadratic estimator they require normalization and noise bias subtraction, but analytic calculations of

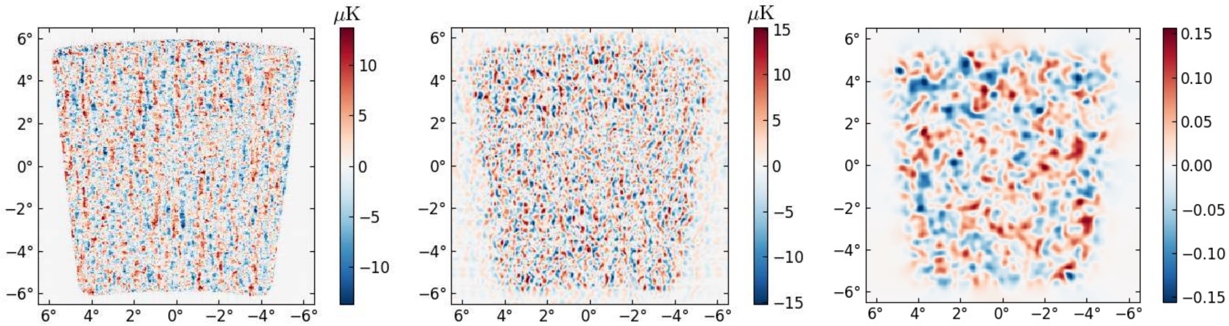

Figure 3.1: SPTpol CMB polarization and QE lensing convergence maps from the SPTpol $100\,\mathrm{deg}^2$ ultradeep analysis. The observed CMB Q map is shown in **(a)**, with the directional nature of the Stokes vector clearly visible. The Q and U maps are combined via Equation (1.7) to produce E and B maps, which are subsequently inverse-variance filtered; the Wiener filtered E map $C_\ell^{EE}\bar{E}_{\hat{\mathbf{n}}}$ is shown in **(b)**. Pairs of filtered maps are fed to the quadratic estimator to create the minimum-variance estimate of the CMB convergence $\kappa \equiv -\nabla^2\phi/2$ shown in **(c)**. The convergence map has been smoothed with a $20'$ Gaussian beam to remove noise-dominated scales in the reconstruction.

these biases are not available. An alternative approach is to use Bayesian inference to directly estimate cosmological parameters, thus avoiding the need to normalize and debias the power spectrum of the MAP estimate. This approach has been developed in Anderes et al. (2015); Millea et al. (2019, 2020), and is applied to real data for the first time here, published as Millea et al. (2021). We simultaneously perform a Bayesian reconstruction of the lensing potential and estimate cosmological parameters A_ϕ and A_L, as well as systematic parameters including the polarization rotation angle, temperature-to-calibration leakage coefficients, and beam eigenmodes. This marks the first time cosmological parameters have been estimated from optimal lensing reconstructions with real data. The Bayesian method is compared to the standard quadratic estimator to build confidence in this new method and demonstrate the improvements that can be attained with optimal lensing reconstruction. I was second author on Millea et al. (2021), where I was primarily responsible for leading the quadratic estimator analysis.

To achieve the low noise levels at which optimal reconstruction methods start to show improvements over the standard quadratic estimator, we combine overlapping CMB maps from the SPTpol $100\,\mathrm{deg}^2$ (Story et al., 2015, Figure 1.3) and $500\,\mathrm{deg}^2$ (Wu et al., 2019)

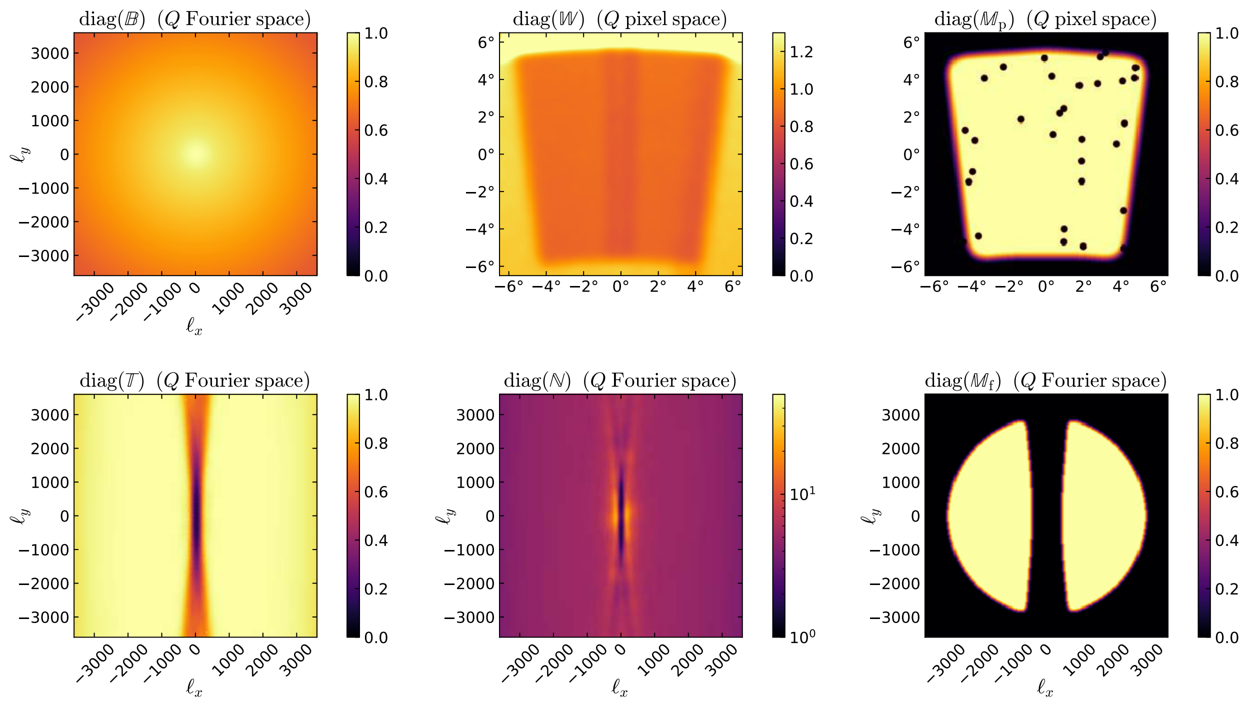

Figure 3.2: Ingredients of the inverse variance filter used in the SPTpol 100 deg^2 ultradeep analysis. The first row shows the beam $\mathbb{B}$, pixel weights $\mathbb{W}$, and pixel mask $\mathbb{M}_p$. The second row shows the transfer function $\mathbb{T}$, noise power spectrum $\mathbb{N}$, and Fourier mask $\mathbb{M}_f$. Projection effects are minimal in the $\sim 10^\circ \times 10^\circ$ patch so the scan-direction timestream filtering cleanly projects onto the $\hat{\mathbf{x}}$ direction, producing a localized trench along the ℓ_y axis in Fourier space.

lensing analyses. This allows us to create "ultradeep" 100 deg^2 polarization maps with 6 μK-arcmin noise levels (Figure 3.1a). Maps are made in the ProjZEA projection with relatively large $3'$ pixels to reduce the computational cost of the Bayesian estimator; we only considers scales below $\ell < 3000$ in the CMB maps when performing lensing reconstruction.

The CMB maps are inverse variance filtered with Equation (2.22) with the resulting Wiener filtered E map shown in Figure 3.1b. The map- and pixel-space ingredients to the filter are shown in Figure 3.2. We take this opportunity to build intuition for projection effects in ProjZEA, which are relatively small for the SPTpol 100 deg^2 patch of sky. The patch is defined as a square in right ascension and declination, but becomes slightly trapezoidal in ProjZEA; the scan-direction timestream filtering along lines of constant declination projects onto the $\hat{\mathbf{x}}$ direction in ProjZEA. In Fourier space, the timestream filtering appears as a trough along the ℓ_x axis since the highpass component of the filter removes large-scale ℓ_x

modes. However even for this small patch the scan direction in ProjZEA varies slightly with right ascension, introducing a mild "splaying" into the transfer function caused by the averaging together of the different scan directions. This effect is much larger for the 1500 deg^2 SPT-3G field as discussed in the following chapter.

We evaluate the quadratic estimator and normalize by the analytic response before correcting for the mean-field and MC response estimated from simulations. The reconstructed convergence map is shown in Figure 3.1c, where we have smoothed the map with a $20'$ Gaussian beam to remove noise-dominated scales in the reconstruction. In order to relate the lensing map $\hat{\phi}_{\mathbf{L}}$ to cosmological parameters such as A_ϕ, we take its auto spectrum and form biased lensing power spectra $C_L^{\hat{\kappa}\hat{\kappa}}$. We then subtract the $\mathbb{N}_L^{0,\mathrm{RD}}$ and $\mathbb{N}_L^1$ biases estimated from simulations, and apply a final percent-level multiplicative MC correction f_{PS} to correct for any remaining mode-coupling effects from point source masking. Figure 3.3 contains the main results from the quadratic estimator analysis. The top panel shows the initial noise-biased power spectrum of the convergence map $C_L^{\hat{\kappa}\hat{\kappa}}$ as well as the noise bias terms $\mathbb{N}_L^{0,\mathrm{RD}}$ and $\mathbb{N}_L^1$. Also shown is the cross spectrum between input and reconstructed convergence maps from simulations used to calculate MC response correction $\mathbb{R}_{|\mathbf{L}|}^{\mathrm{MC}}$. In the the bottom panel of Figure 3.3 we present the final debiased bandpowers $\widehat{C}_L^{\phi\phi}$, showing excellent agreement with our suite of simulations and the best-fit *Planck* cosmology. We find the quadratic estimate of the lensing amplitude to be $A_\phi = 1.00 \pm 0.15$.

The Bayesian lensing approach involves a Monte Carlo Markov Chain (MCMC) procedure that yields samples from the posterior distribution of the CMB lensing potential and cosmological parameters. This method is optimal at any noise level and fully utilizes all orders of information in the data. The Bayesian analysis in this study is used to jointly infer the parameters A_ϕ and A_L, providing a self-consistent model of the data that includes both parameters. This joint inference allows us to account for the correlation between the reconstructed lensing potential and the delensed CMB power spectra. The parameter A_ϕ determines the overall amplitude of the lensing power spectrum, while A_L quantifies the

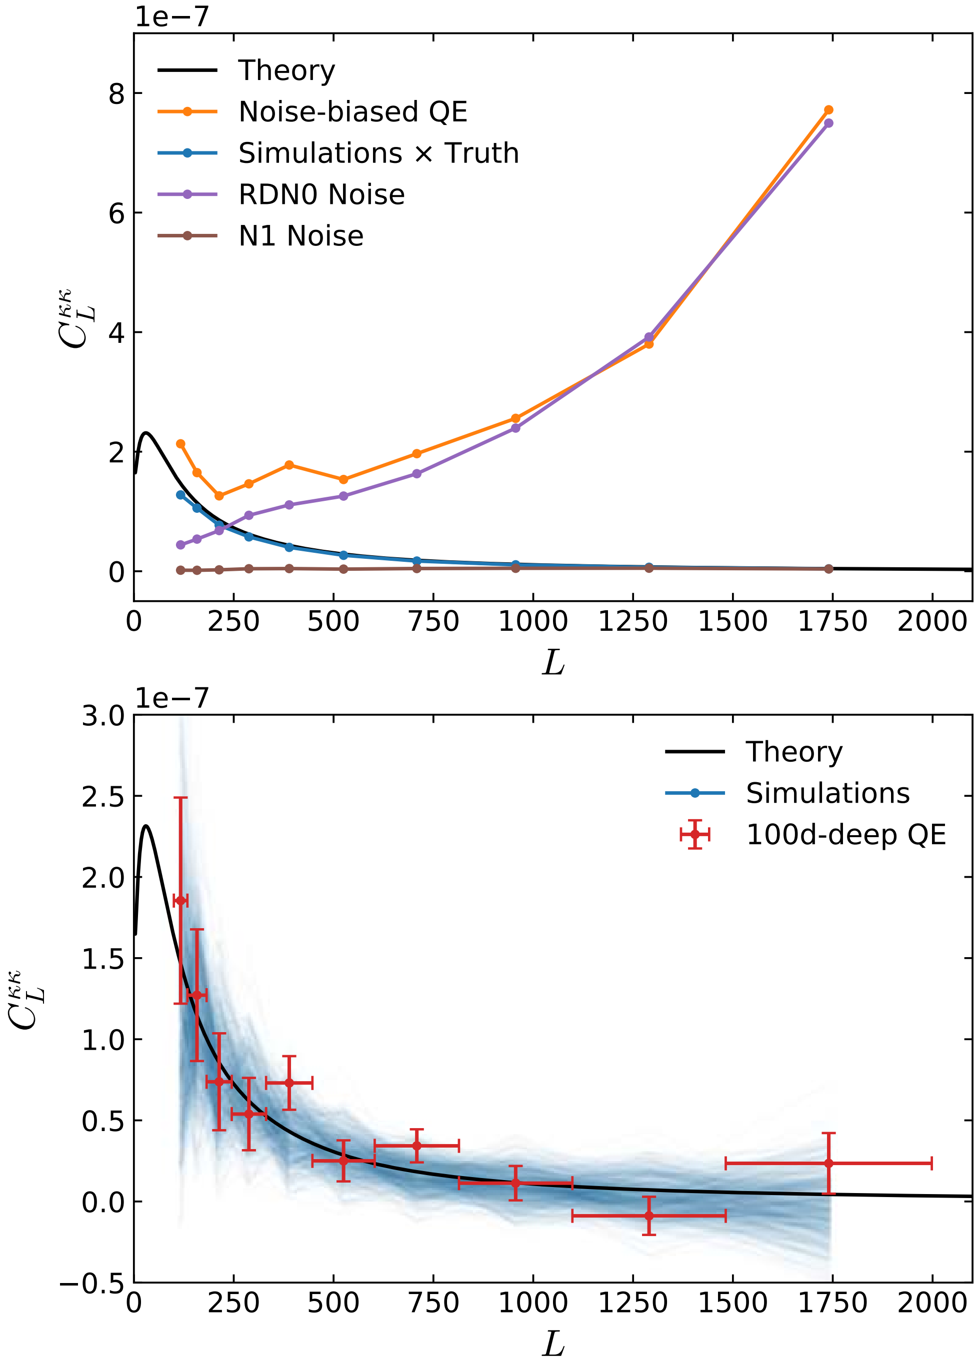

Figure 3.3: Components of the SPTpol 100 deg^2 ultradeep lensing power spectrum pipeline and final debiased bandpowers. The top panel shows the noise-biased spectrum $C_L^{\hat{\kappa}\hat{\kappa}}$ and estimates of the additive bias terms $\mathbb{N}_L^{0,\mathrm{RD}}$ and $\mathbb{N}_L^1$. In the bottom panel we show the debiased data bandpowers compared to theory and simulations.

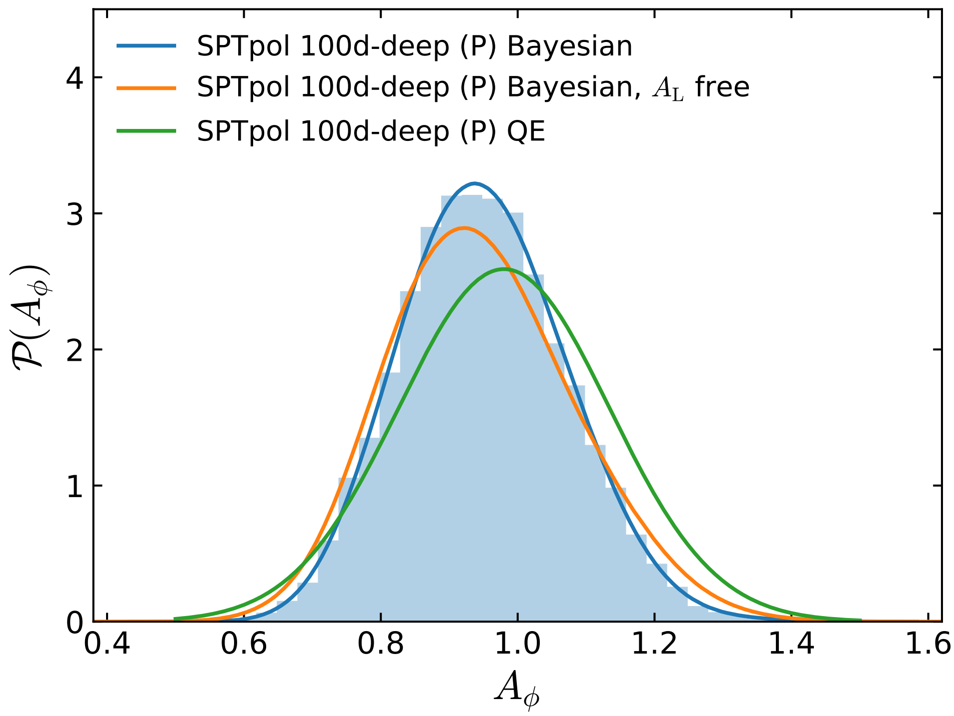

Figure 3.4: Comparison of A_ϕ constraints from the Bayesian and QE pipelines the SPTpol $100\,\mathrm{deg}^2$ ultradeep analysis. While the two constraints are consistent with each other, the Bayesian method achieves a 17% tighter constraint on A_ϕ compared to the QE method.

peak-smoothing effect from lensing on the *CMB* power spectrum; this allows a separation of the 2-point and 4-point effects of CMB lensing.

We compare the constraints on A_ϕ from the Bayesian and QE methods in Figure 3.4; the Bayesian pipeline obtains $A_\phi = 0.95 \pm 0.12$, consistent with the quadratic estimator but with 26% tighter error bars. However the Bayesian method also uses CMB power spectrum information to constrain A_L which in turns tightens the constraints on A_ϕ. Ignoring the 2-point CMB constraints we find that the Bayesian pipeline gives a 17% tighter constraint on A_ϕ compared to the QE method, demonstrating the improvement achievable with optimal lensing reconstructions from deep polarization maps. This experience leading a quadratic estimator analysis on a relatively small and simple dataset was an excellent preparation for the SPT-3G lensing analysis discussed in the next chapter.

Chapter 4

SPT-3G Lensing

The SPT-3G experiment is extremely well suited for lensing measurements thanks to its low noise levels and high angular resolution compared to *Planck* as well as its larger sky coverage compared to SPTpol. This combination of factors allows a larger number of CMB multipoles $\boldsymbol{\ell}$ to be measured with high signal-to-noise, in turn increasing the number of high signal-to-noise modes that can contribute to a given lensing multipole $\mathbf{L} = \boldsymbol{\ell} - \boldsymbol{\ell}'$. Lensing measurements from the full-depth SPT-3G dataset will break degeneracies between cosmological parameters inferred from primary CMB measurements, and improve constraints on the Hubble constant H_0 by a factor of two compared to *Planck* (Figure 4.1; see also Prabhu et al., 2024). In addition to placing competitive constraints on ΛCDM parameters, SPT-3G lensing maps will be excellent for cross-correlations with other tracers of large-scale structure, especially because of SPT-3G's ability to make deep polarization-only lensing maps that will be robust against foreground-induced biases which can significantly affect temperature-based lensing reconstructions. Finally, SPT-3G lensing maps will be instrumental in removing the contribution from lensing-induced B-modes in the joint BICEP-*Keck*/SPT search for primordial gravitational waves, as sample variance from lensing has already become the largest source of uncertainty in this search (BICEP2 Collaboration et al., 2018; BICEP/Keck Collaboration et al., 2021).

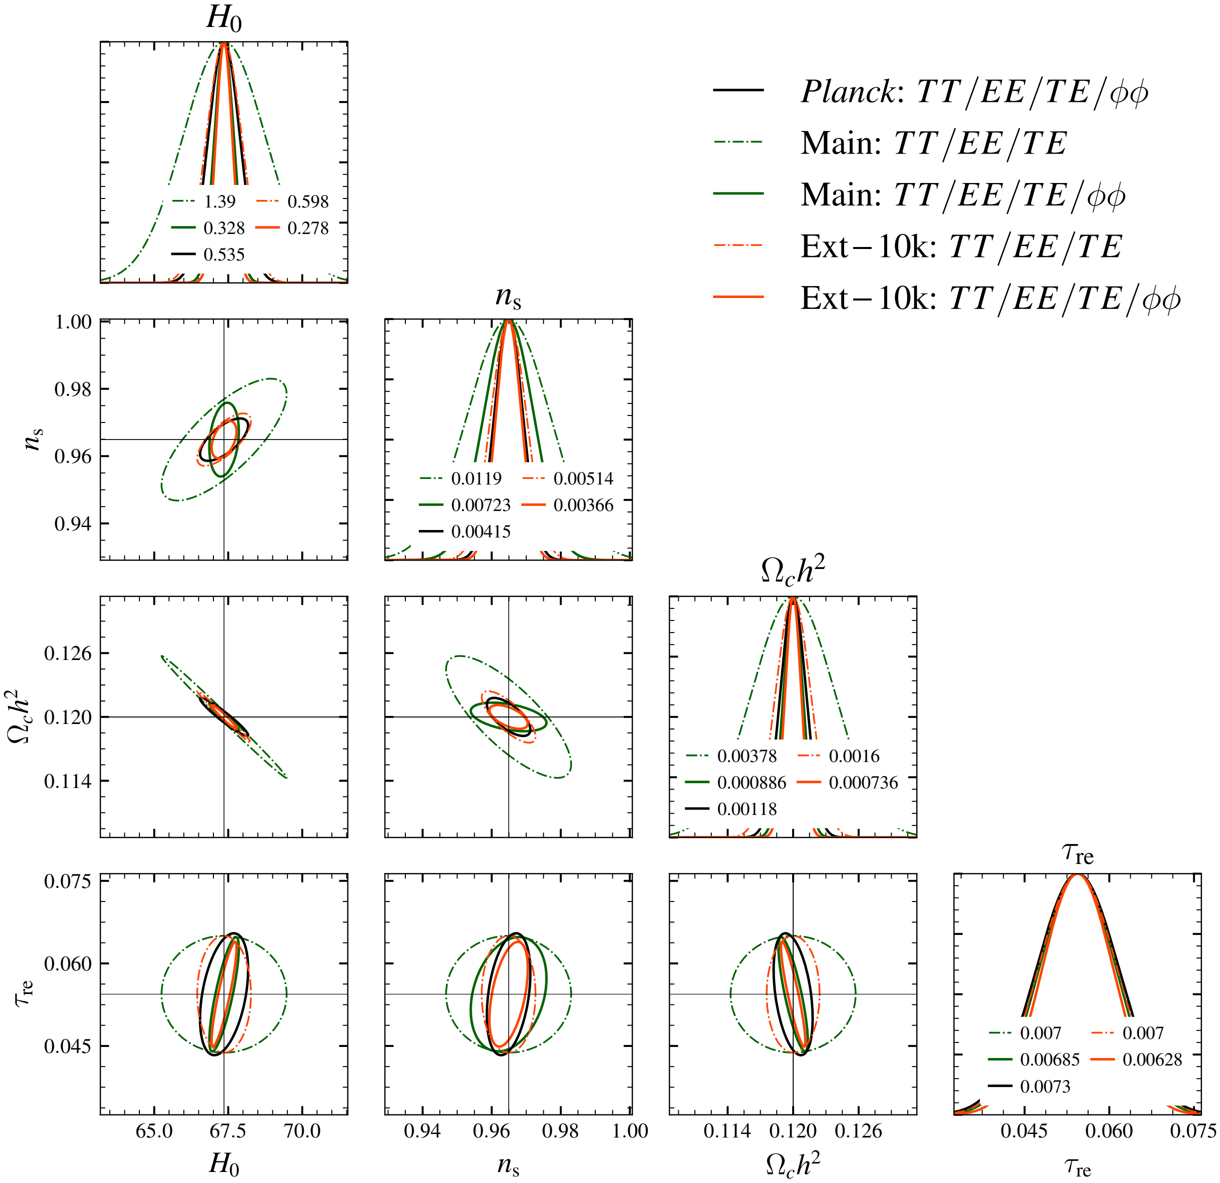

Figure 4.1: Forecasted constraints on ΛCDM parameters from the full-depth SPT-3G survey. The profiles and contours represent marginalized 1-dimensional and 2-dimensional 68% contours for the Hubble constant H_0, scalar spectral index n_s, cold dark matter density $\Omega_c h^2$, and reionization optical depth $\tau_{\rm re}$. Constraints from the 1500 deg^2 main survey are shown in green while the extended 10 000 deg^2 survey is shown in red; the solid lines show the constraints from the primary CMB alone while the dashed lines show the improvement when lensing information is added. Combined *Planck* constraints from the primary CMB and lensing are shown in black. SPT-3G parameter constraints from the full-depth main survey will provide a consistency check on measurements from *Planck*; the inclusion of lensing information also breaks degeneracies between parameters and will allow us to set even tighter constraints on the Hubble constant than *Planck*. Figure from the SPT-3G forecasting paper (Prabhu et al., 2024).

The first lensing measurement made with SPT-3G was presented in Pan et al. (2023), and found a lensing amplitude $A_\phi = 1.020 \pm 0.060$ relative to a fiducial *Planck* 2018 ΛCDM cosmology. This measurement used temperature-only 95 GHz and 150 GHz data from half of the SPT-3G focal plane over half of the 2018 observing season due to problems with the telescope drive system and receiver, and was subject to excess low-frequency frequency noise in the CMB observations from detector wafer temperature drifts. Despite these limitations Pan et al. (2023) were able to measure the lensing power spectrum with per-mode signal-to-noise greater than unity out to $L \lesssim 100$. Constraints on H_0 and Ω_m were ∼30% tighter than previous SPT lensing measurements (Omori et al., 2017; Wu et al., 2019), and competitive with measurements from the *Planck* and ACT experiments (Planck Collaboration et al., 2020c; Madhavacheril et al., 2024).

4.1 Analysis Overview and Motivation

The next SPT-3G lensing analysis, which is the main topic of this thesis, uses data from the 2019 and 2020 observing seasons. Due to the reduced data volume used in the 2018 analysis, the 2019+2020 dataset represents a factor of ∼10 increase in data volume, with ProjZEA map noise levels of 5.3 μK-arcmin, 4.4 μK-arcmin and 16 μK-arcmin in temperature and 8.2 μK-arcmin, 6.6 μK-arcmin and 26 μK-arcmin in polarization at 90 GHz, 150 GHz and 220 GHz respectively. As a result, the 2019+2020 analysis will produce the deepest CMB lensing maps ever made per-pixel, with reconstruction noise reduced by a factor of ∼three compared to previous measurements (Figure 4.2) and a resulting ∼ 2% measurement of the amplitude of the lensing power spectrum A_ϕ.

The combination of sky coverage and map depth in the 2019+2020 dataset brings new opportunities and challenges that motivate different approaches to lensing reconstruction than the standard flat-sky quadratic estimator that was used in the previous SPTpol and SPT-3G lensing analyses of Story et al. (2015), Wu et al. (2019), Millea et al. (2021), and Pan et al.

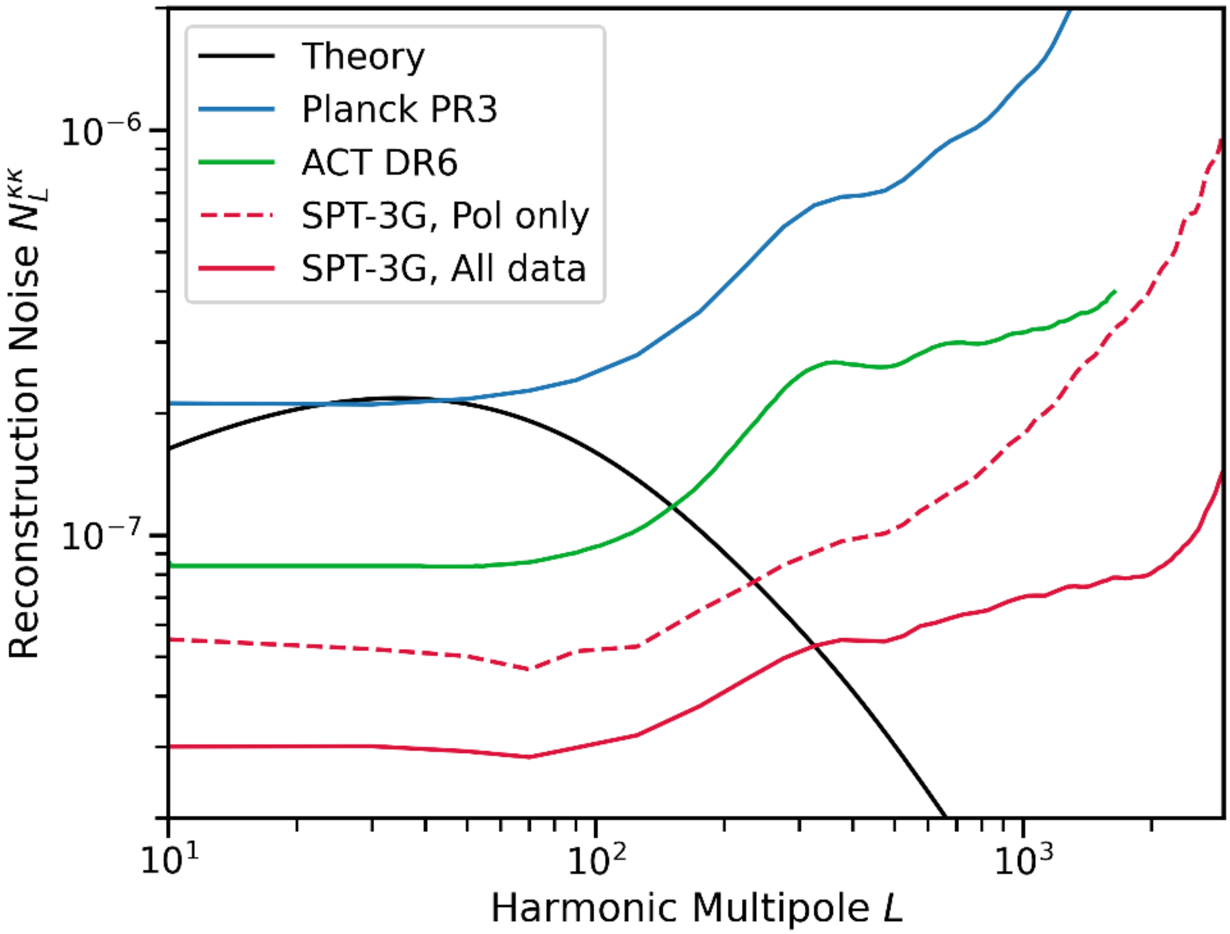

Figure 4.2: Per-mode lensing reconstruction noise levels from the SPT-3G 2019-2020 flat-sky quadratic estimator analysis as a function of lensing multipole L, compared to noise levels from other experiments and the expected lensing power spectrum. The best-fit *Planck* 2018 ΛCDM lensing power spectrum is shown in black, and the Planck 2018 lensing reconstruction noise levels are shown in blue (Planck Collaboration et al., 2020c). Reconstruction noise levels from the recent ACT measurement of Madhavacheril et al. (2024) are shown in green. SPT-3G 2019+2020 minimum variance and polarization-only noise levels are shown with red solid and dashed lines respectively. The 2019+2020 SPT-3G minimum-variance reconstruction has a per-mode signal-to-noise greater than unity out to $L \lesssim 330$, is nearly four (three) times deeper than the Planck (ACT) measurement; the SPT-3G lensing maps will be the deepest ever made per-pixel. Roughly half the SPT-3G 2019+2020 signal-to-noise comes from polarization, allowing the production of clean polarization-only lensing maps robust to foreground biases.

(2023). On one hand the flat-sky approximation begins to break down for the 1500 $\deg^2$ winter field as discussed in Chapter 1.4, leading to suboptimalities in $QU \rightarrow EB$ transformation, transfer function estimation, and inverse-variance filtering, all of which increase the variance of the lensing reconstruction. This motivates the use of a curved-sky lensing estimator leveraging the HEALPix projection (Górski et al., 2005) and spherical harmonic transforms as was done in e.g. Omori et al. (2018) and Planck Collaboration et al. (2020c). Curved-sky lensing estimators will be critical for future lensing analyses using the SPT-3G summer and wide fields, which will bring the total sky coverage of the SPT-3G experiment to over 10 000 $\deg^2$. On the other hand, the polarization noise levels of the 2019+2020 dataset are beginning to approach the $\sim 5\,\mu$K-arcmin lensing B-mode noise floor where the quadratic estimator becomes suboptimal as discussed in Chapter 3. This motivates the use of Bayesian methods which can extract all higher-order information from the data as demonstrated on real data for the first time in Millea et al. (2021). Given these considerations, three independent analysis pipelines are used to make lensing measurements with the 2019+2020 SPT-3G dataset: a flat-sky quadratic estimator pipeline, a curved-sky quadratic estimator pipeline, and a flat-sky Bayesian pipeline based on the MUSE algorithm (Millea & Seljak, 2021). This three-pipeline approach brings the additional advantage of allowing rigorous cross-checks of the results, confirming the robustness of the measurements to different analysis choices and estimators. These pipelines are briefly summarized below:

- *Flat-sky quadratic estimator:* **I am leading the flat-sky quadratic estimator pipeline,** which is similar to the pipeline used in the 2018 SPT-3G lensing analysis (Pan et al., 2023) but includes a range of improvements. The 2019+2020 flat-sky pipeline adds polarization and 220 GHz data, as well as implementing a minimum variance combination of the frequency bands and foreground bias-hardening for the first time in a flat-sky SPT lensing analysis. I have also refactored the large majority of the codebase in an effort to improve readability, accuracy, modularity, and performance. The advantages of the flat-sky quadratic estimator pipeline include its speed and modularity relative to the other two pipelines, and the implementation of bias-hardened estimators not available in the other two pipelines. However, it produces roughly 15% larger error bars on the lensing power spectrum due to projection issues and the flat-sky approximation, as well as its use of the suboptimal quadratic estimator (SQE) over the global minimum variance estimator used by the curved-sky pipeline and the MUSE algorithm used by the Bayesian pipeline.

- *Curved-sky quadratic estimator:* The curved-sky pipeline has lower noise and tighter error bars than the flat-sky pipe because it does not suffer from projection effects, and it includes an implementation of the global minimum variance estimator (GMV; Maniyar et al., 2021). The GMV estimator takes into account the TE correlations in both the inverse-variance filter and quadratic estimator steps, and reduces the noise of the estimator by roughly 10% for third-generation CMB experiments like SPT-3G and SO. The curved-sky GMV estimator is somewhat slower than the flat-sky estimator due to the use of spherical harmonic transforms over fast Fourier transforms and the extra computational cost of the GMV estimator; it is also somewhat more difficult to evaluate custom combinations of estimators with the GMV.

- *Bayesian pipeline (MUSE):* This pipeline uses the MUSE algorithm (Millea & Seljak, 2021) which is an evolution of the optimal sampling-based approach used in Millea et al. (2021). Hamiltonian Monte Carlo methods were found to be too computationally expensive for the $1500\,\mathrm{deg}^2$ SPT-3G field, so the MUSE algorithm was developed as a near-optimal fast alternative that performs an approximate marginalization over latent parameter space using a suite of forward-model simulations of the data. The MUSE pipeline has the advantage of jointly estimating the unlensed CMB power spectrum *and* the lensing power spectrum, meaning that it is both a primary CMB and lensing analysis. The MUSE pipeline only uses polarization data in order to simplify the scope and complexity of this first application of the algorithm to real data; as a result it provides tighter constraints than the polarization-only quadratic estimators, but has less signal-to-noise than the all-data quadratic estimators.

The remainder of this chapter is organized as follows. Chapter 4.2 introduces the SPT-3G 2019+2020 dataset and describes the mapmaking, beam estimates, and calibration procedures used to produce the CMB maps used as input to the quadratic estimator. The simulations used to normalize and debias the lensing maps and power spectrum and characterize foreground biases are introduced in Chapter 4.3. Chapter 4.4 describes the processing steps applied to the CMB maps including a minimum variance linear combination of the frequency bands, inpainting of point sources and galaxy clusters, and anti-aliasing techniques. The lensing reconstruction process is described in Chapter 4.5 including inverse variance filtering and the normalization and debiasing of the lensing maps; the lensing maps themselves are presented here. Power spectrum estimation is discussed in Chapter 4.6, including the application of an obfuscation factor to avoid confirmation bias. Chapter 4.7 presents the tests and consistency checks that must be performed before finalizing the analysis choices, removing the obfuscation factor, and proceeding to cosmological parameter estimation. These tests and checks are currently in progress, so we show the obfuscated lensing power spectra in Chapter 4.8. Finally, Chapter 4.9 concludes with a discussion of future work and the implications of the SPT-3G lensing analysis for cosmology.

4.2 Data

For this lensing analysis we use data from SPT-3G's 2019 and 2020 winter observing seasons. The $1500\,\mathrm{deg}^2$ SPT-3G winter footprint extends from $-42°$ to $-70°$ in declination (Dec) and $-50°$ to $50°$ in right ascension (RA); the footprint is divided into four subfields with declination centers of $-44.^{\circ}75$, $-52.^{\circ}25$, $-59.^{\circ}75$, and $-67.^{\circ}25$ to ensure consistent detector response and linearity across the patch. Only one subfield is mapped in a given observation, which lasts roughly two hours; more than 3000 subfield observations were taken during the 2019-2020 observing season. The subfields are observed 1034, 902, 772, and 578 times respectively across the two seasons for a total of 3286 observations; the number of observations of each subfield is chosen to produce approximately uniform map noise levels across subfields. During a two-hour observation, the SPT performs 72 constant-elevation (and thus constant-declination) scans across a given subfield; the scan speed is roughly $1\,\mathrm{deg\,s^{-1}}$, and the telescope takes an elevation step of 12.5 arcmin between scans. Each of the $\sim$16 000 detectors in the SPT-3G camera sample the sky as the telescope scans, producing one-dimensional "timestreams" that record flux as a function of right ascension. The timestreams are originally sampled at 152.6 Hz, but for this analysis we use downsampled timestreams with a sample rate of 76.3 Hz which corresponds to on-sky sample rates between $0.3\,\mathrm{arcmin\,s^{-1}}$ and $0.6\,\mathrm{arcmin\,s^{-1}}$ depending on declination. The Nyquist frequency of the downsampled timestreams is $\ell \geq 18\,000$, well beyond the angular scales of interest for this analysis.

4.2.1 Mapmaking

We filter the aforementioned timestreams and bin their samples together to make a pixelated map of the sky in a process known as mapmaking. The timestreams are low-passed using a filter of the form $e^{(-\ell/\ell_0)^6}$ with $\ell_0 \sim 13\,000$ to remove small-scale noise that could be aliased to larger scales when binning the timestreams into pixels. In addition, we high-pass filter the timestreams to remove large-scale instrumental and atmospheric noise by subtracting a 19th-order Legendre polynomial fit and then projecting out Fourier modes at $\ell \leq 300$. This two-step high-pass ensures that large-scale features are removed even if they are not well described in Fourier space.

The filtering described above can produce undesired periodic features or "ringing" in the timestreams around the locations of point sources Figure 4.3. To avoid these features, we mask point sources and clusters when fitting the filtering templates by assigning zero weight to all samples within a given radius of the source. This masking only applies to the filtering templates, and the template is subtracted from all samples, including those

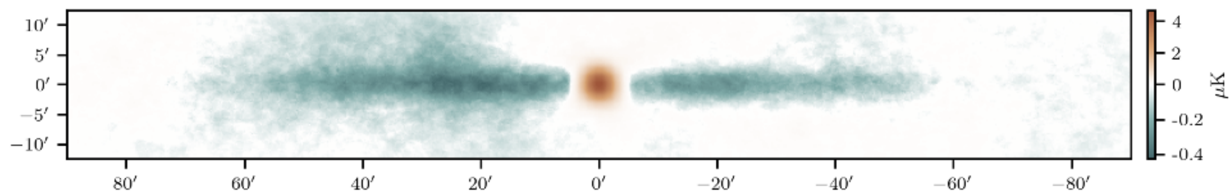

Figure 4.3: Example of point source filtering wings at 150 GHz in one region of the SPT-3G winter field using the Samson-Flamsteed or sinusoidal projection. The heatmap shows a median stack of point sources with a signal-to-noise > 5 and a flux < 30 mJy. To capture the dynamic range of stack, positive and negative values employ linear color stretches with different slopes; the median amplitude of the filtering wings is roughly 10 % that of the sources. Data provided by Kayla Kornoelje and Lindsey Bleem.

that were masked when constructing the template. This procedure reduces artificial ringing around point sources and clusters, but does not mask the sources themselves in the final map. However, these "masked" regions have different filtering than the rest of the map and cannot be used for cosmological analyses.

We mask all emissive sources with 150 GHz flux > 6 mJy and all galaxy clusters with a signal-to-noise > 10 for a total of 2655 masked objects. The source and cluster catalogs used are preliminary versions of the SPT-3G catalogs that will be published at a later date in Archipley et al. (in prep). A circular mask is used in all cases; for point sources, the mask radius is defined on a per-source basis as the distance at which the beam-convolved source signal-to-noise falls below 1, e.g. where the source flux intersects with the noise floor. Clusters are also masked out to a radius where the beam-convolved signal-to-noise falls below 1, but we combine the beam with a beta profile (Cavaliere & Fusco-Femiano, 1976, 1978; Plagge et al., 2010) to account for the angular extent of the cluster. The function form of the beta profile is $f(\theta) \sim (1 + (\theta/\theta_c)^2)^{-1}$, where θ_c is the best-fit core radius of the cluster. Thermal emission from weather balloons launched at the South Pole by the Antarctic Meteorological Research Center and National Oceanic and Atmospheric Administration can dominate the CMB emission in single-observations maps, so we also mask the the balloon flight paths in the timestreams.

After filtering the timestreams, we bin the samples into $0.5625'$ flat-sky pixels using an inverse-variance weighting scheme. The ProjZEA projection is centered on a right ascension of $0°$ and a declination of $-59.55°$. The pixel size was chosen so that $1500\,\mathrm{deg}^2$ patch fits in a 4096×8192 ProjZEA map, as the efficiency of FFTs is maximized when the map size is a power of two. Weights are determined from the detector noise in the range $320 < \ell < 4000$. This downweights high-noise detectors and scans within a single observation, and observations on days with poor observing conditions will contribute less to the final map. The binning operation picks up one factor of the output pixelization's pixel window function, $\mathrm{PWF}(0.5625')$.

I ran the mapmaking pipeline on all 3286 observations from the 2019 and 2020 winter observing seasons at 90 GHz, 150 GHz and 220 GHz individually. The resulting flat-sky and HEALPix temperature and polarization maps are used by many ongoing SPT-3G projects, including lensing, delensing and birefringence analyses. Each mapmaking job takes approximately 2 hours to run, and occupies approximately 5 GB of memory. Three versions of the maps were produced as we converged on the final mapmaking choices such as timestream filtering parameters and point source masking thresholds, for a total of 60 000 computing hours. Mapmaking jobs were run in parallel on the Open Science Grid high-throughput computing pool (Pordes et al., 2007; Sfiligoi et al., 2009; OSG, 2006, 2015).

After maps have been made for each observation in the 2019 and 2020 winter observing seasons, we sum these maps into a full-depth "coadd" with minimal data cuts. We exclude all observations before March 22nd and after November 30th to avoid sun contamination, and cut $\sim$5 observations that produce empty maps.

In addition to full-depth coadds, we also construct 500 nearly independent noise realizations by summing the observations with random signflips applied, cancelling the signal but not the noise. This allows us to add noise with the correct statistical properties to our simulations. We estimate that the noise realizations are correlated with each other at the percent level due to the finite number of observations in the dataset, but this is not expected to significantly impact the lensing analysis.

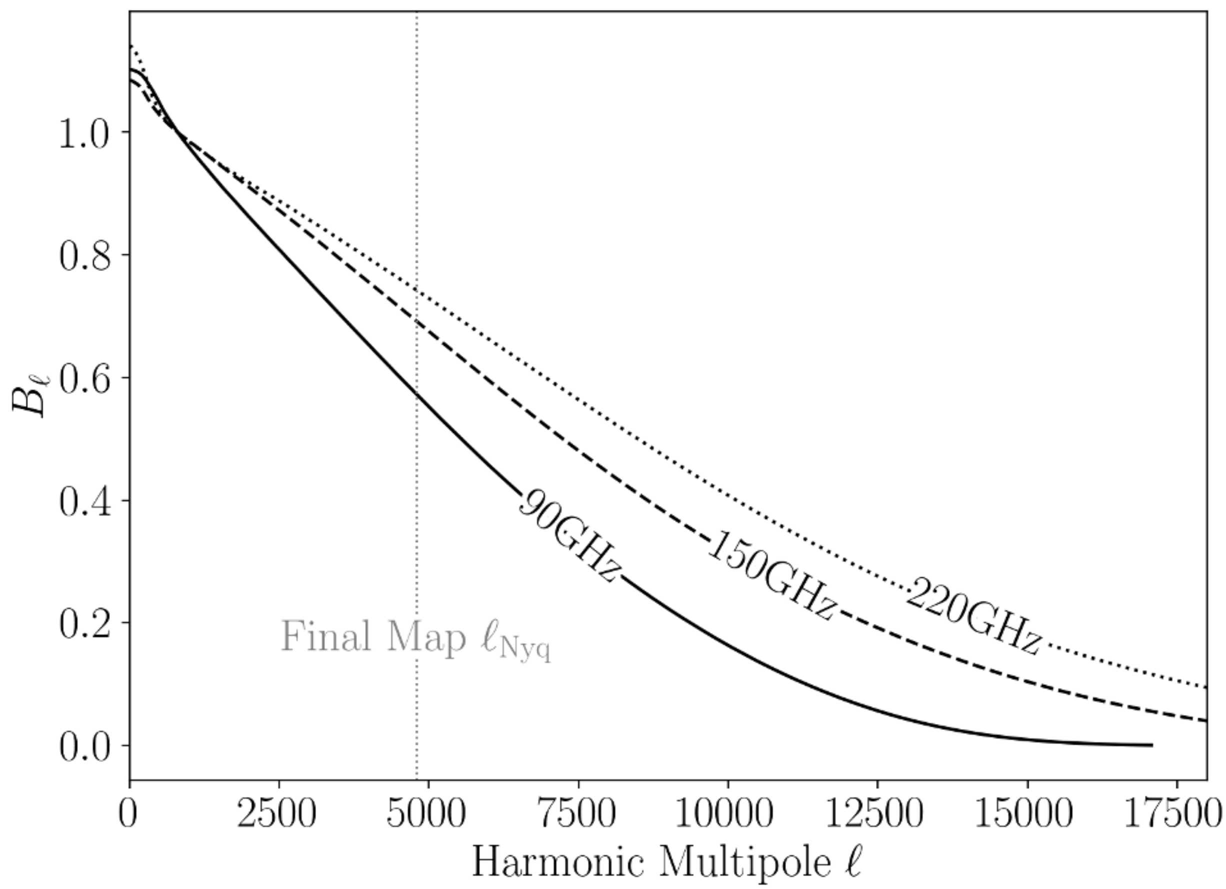

Figure 4.4: Preliminary SPT-3G beams at 90 GHz, 150 GHz and 220 GHz used in this analysis; final versions will be presented in Huang et al. (in prep). The beams are estimated from composite observations of AGN and Saturn, and are used to characterize the point source response of the instrument. The angular resolution of the SPT-3G instrument increases with frequency, with beam full-width half-maxima (FWHM) of $1.57'$, $1.17'$ and $1.04'$ at 90 GHz, 150 GHz and 220 GHz respectively (Sobrin et al., 2022). The beam response is $\gtrsim 0.6$ in the range of angular scales considered in this analysis.

4.2.2 Beams

The SPT-3G beams describe the point source response of the instrument as a function of angular scale, and are a measure of the resolution of the instrument. The beams can be roughly approximated as Gaussian, with full-width half-maxima (FWHM) of $1.57'$, $1.17'$ and $1.04'$ at 90 GHz, 150 GHz and 220 GHz respectively (Sobrin et al., 2022). However, extremely precise characterization of the beam is necessary for a wide range of SPT-3G analyses. The SPT-3G beams are estimated in a similar manner as in SPT-SZ (Keisler et al., 2011), using composite observations of bright point-like AGN and planets. This analysis, which I will briefly summarize, will be published in a forthcoming collaboration paper (Huang et al., in prep).

Observations of AGN and Saturn must be stitched together to produce an accurate estimate of the beam on all scales; this is because the AGN are not bright enough to capture the faint beam sidelobes at longer angular separations, while Saturn is spatially extended and saturates the detectors when viewed directly. To characterize the inner beam, 8 AGN at different locations in the SPT-3G field are used; these AGN maps are normalized and corrected for the effects of declination-varying time constants and CMB contamination. Characterization of the outer beam with Saturn is difficult because the detectors saturate when Saturn is viewed directly, and take time to return to linear responsivity As a result, only approaching scans of Saturn can be used, and only beyond a certain radius. This gives an inner radius beyond which Saturn is not saturated, and an outer radius with which the AGN have high enough signal-to-noise. The AGN and planet maps are stitched within this annulus which has inner radii of $2'$, $1.5'$ and $1'$ at 90 GHz, 150 GHz and 220 GHz, and a width of $1.5'$ at all bands. The beam profile is then estimated from these stitched maps; the beam covariance is still being finalized, but statistical covariance is estimated from the scatter among the 8 AGN, and systematic covariance from factors such as the stitching radius and CMB background subtraction will also be included. The beams used in this analysis are showed in Figure 4.4.

4.2.3 Calibration

The true or absolute brightness of the CMB is difficult to measure due to the interplay between instrumental and atmospheric effects, both of which can vary with time. Both relative and absolute calibration are performed on SPT-3G data. Schematically, relative calibration ensures that individual detectors and observations are consistent with each other, while absolute calibration with *Planck* ensures that the coadded maps are faithful to the true

CMB. *Planck*'s absolute calibration is of excellent quality thanks in large part to the fact that it is a space-based mission, which imparts two main advantages. *Planck* does not suffer from atmospheric effects, and more importantly the orbital modulation of the CMB dipole can be combined with the satellite's exquisitely characterized orbital motion and precise measurements of the absolute CMB temperature (Fixsen, 2009) to calibrate the on-board instruments (Planck Collaboration et al., 2015b, 2016b).

SPT-3G's relative calibration is described in Sobrin et al. (2022), but I will briefly summarize the process here. Because the detector properties change with weather and pointing elevation, relative calibration observations must be performed regularly and at multiple declinations. The Galactic HII regions RCW38 and MAT5A (Coble et al., 2003) are observed every 8 days in a "dense" mode such that every detector scans across the source; MAT51 is used to calibrate the upper two subfields, while RCW38 is used for the lower two. Each detector's map of a HII region is fit to a band-averaged template, and this fit is compared to the known flux of the source to determine each detector's calibration in CMB temperature units. However detector calibration can vary on shorter timescales than the 8-day cadence of HII region observations, so two additional corrections are applied to ensure that each observation's calibration is consistent. Before each HII region and CMB field observation, the detector response to a thermal source located behind the secondary mirror is measured, and the CMB field observations are corrected by the ratio of these responses. In addition, to capture variation in atmospheric transmission due to weather, the HII regions are briefly re-observed before each CMB field observation in a "sparse" mode. Because each detector does not see the source in this sparse mode, the comparison between sparse and dense HII calibration observations is performed at the coadded map level, with the relative calibration defined as the ratio of the integrated flux within a $4' \times 4'$ box around the source. The relative calibration procedure is largely automated as part of SPT-3G autoprocessing pipeline, and the calibration numbers are bundled with each observation's metadata such that individual observation maps produced by the mapmaker are properly calibrated relative to each other without any special input form the person making the maps.

Absolute calibration relative to *Planck* is performed on the full-depth coadded maps, and is done in four steps:

1. Temperature calibration
2. Temperature-to-polarization leakage removal
3. QU rotation
4. Polarization calibration

Temperature calibration is performed by comparing the SPT-3G data to *Planck* frequency maps passed through the SPT-3G observing pipeline. This external comparison is done at by cross-correlating the SPT-3G 150 GHz and *Planck* 143 GHz maps; we then internally calibrate the 90 GHz and 220 GHz SPT maps by cross-correlating with the calibrated 150 GHz SPT map. Because the SPT maps have higher signal-to-noise than *Planck* in the $\ell \in [600, 1200]$ range used for calibration, the internal calibration produces more precise calibration factors for 90 GHz and 220 GHz compared to calibrating these bands to the corresponding *Planck* bands directly. Both the external and internal calibration factors are defined as the ratio of the cross-spectrum of the full-depth *Planck* and SPT-3G maps to the cross-spectrum of two half-depth SPT-3G maps. For example, at 90 GHz:

$$\begin{aligned} T^{150}_{\text{cal, external}} &= \frac{\text{SPT}^{150}_{\text{full}} \times Planck^{143}_{\text{full}}}{\text{SPT}^{150}_{\text{half1}} \times \text{SPT}^{150}_{\text{half2}}}, \\ T^{90}_{\text{cal, internal}} &= \frac{\text{SPT}^{150}_{\text{full}} \times Planck^{143}_{\text{full}}}{\text{SPT}^{150}_{\text{half1}} \times \text{SPT}^{90}_{\text{half2}}}, \\ T^{220}_{\text{cal, internal}} &= \frac{\text{SPT}^{150}_{\text{full}} \times Planck^{143}_{\text{full}}}{\text{SPT}^{150}_{\text{half1}} \times \text{SPT}^{220}_{\text{half2}}}, \end{aligned} \tag{4.1}$$

This cross-correlation approach avoids contributions from the noise of either experiment. The internal and external calibrations are combined to produce final temperature calibration factors for each band and subfield, as presented in Table 4.2.3.

	90 GHz	150 GHz	220 GHz
Subfield-averaged			
$\Delta\psi$	0.008 06	0.006 75	−0.005 63
R_P	1.073	1.083	1.189
ra0hdec-44.75			
R_T	1.064 ± 0.004	1.013 ± 0.004	1.011 ± 0.007
$\epsilon^{Q,TT}$	0.0025 ± 0.0006	0.0026 ± 0.0006	0.0026 ± 0.0010
$\epsilon^{Q,TQ}$	0.83 ± 0.08	0.90 ± 0.08	0.85 ± 0.12
$\epsilon^{U,TT}$	0.0054 ± 0.0005	0.0072 ± 0.0004	0.0081 ± 0.0008
ra0hdec-52.25			
R_T	1.069 ± 0.003	1.033 ± 0.003	1.003 ± 0.005
$\epsilon^{Q,TT}$	0.0031 ± 0.0007	0.0030 ± 0.0006	0.0038 ± 0.0010
$\epsilon^{Q,TQ}$	0.91 ± 0.09	0.95 ± 0.09	0.77 ± 0.13
$\epsilon^{U,TT}$	0.0060 ± 0.0005	0.0070 ± 0.0004	0.0065 ± 0.0009
ra0hdec-59.75			
R_T	1.070 ± 0.004	1.001 ± 0.004	1.011 ± 0.005
$\epsilon^{Q,TT}$	0.0076 ± 0.0007	0.0094 ± 0.0006	0.0186 ± 0.0009
$\epsilon^{Q,TQ}$	0.79 ± 0.09	0.80 ± 0.09	0.68 ± 0.13
$\epsilon^{U,TT}$	0.0089 ± 0.0005	0.0130 ± 0.0005	0.0111 ± 0.0009
ra0hdec-67.25			
R_T	1.073 ± 0.007	1.014 ± 0.008	1.017 ± 0.009
$\epsilon^{Q,TT}$	0.0065 ± 0.0008	0.0088 ± 0.0008	0.0181 ± 0.0011
$\epsilon^{Q,TQ}$	0.88 ± 0.11	1.00 ± 0.11	0.95 ± 0.15
$\epsilon^{U,TT}$	0.0087 ± 0.0006	0.0118 ± 0.0006	0.0132 ± 0.0010

Table 4.1: Calibration parameters for the SPT-3G 2019+2020 maps at each observing frequency. For each subfield, we apply an absolute temperature calibration factor R_T and monopole T-to-P leakage coefficients $\epsilon^{Q,TT}$, $\epsilon^{Q,TQ}$, $\epsilon^{U,TT}$. We also apply a global polarization rotation angle $\Delta\psi$ and a relative polarization calibration factor R_P to all subfields.

Following temperature calibration, we jointly estimate the temperature-to-polarization (T-to-P) monopole leakage coefficients and the polarization rotation angle. T-to-P leakage can be caused by a variety of sources. Gain mismatch between detector pairs leaks a scaled copy of the temperature map into polarization ("monopole" leakage), whereas differential detector pointing and beam ellipticity introduce copies of the first and second derivatives of the temperature map ("dipole" and "quadrupole" leakage; Hu et al., 2003; Henning et al., 2017). We find no evidence of dipole or quadrupole leakage in the SPT-3G data, so we only correct for monopole leakage. A non-zero polarization rotation angle can be caused by slight misalignments in the orientation of the Q and U detectors; this effect will mix Q $\rightleftharpoons$ U, creating spurious EB and TB (via the TE) correlations which are otherwise predicted to be zero by the standard model (Keating et al., 2012). We correct for the monopole leakage coefficients by subtracting a template of the temperature map from the QU maps, and we correct for the polarization rotation angle by rotating the QU maps such that the EB and TB correlations are nulled. We need not consider $E \rightleftharpoons B$ mixing from gravitational lensing or partial-sky masking, because these effects only mix modes at different multipoles, i.e. $\langle E_\ell B_{\ell'} \rangle$ from lensing or masking is only non-zero when $\ell \neq \ell'$ (Keating et al., 2012).

Because T-to-P leakage affects the polarization angle estimate and vice versa, these two calibration parameters must be jointly estimated in an iterative manner. We first estimate the monopole leakage coefficients by fitting a template to the C_ℓ^{TQ} and C_ℓ^{TU} cross-spectra using auto-spectra from the data and cross-spectra from simulations that have no monopole leakage:

$$C_{\ell,\mathrm{template}}^{TQ} = \epsilon^{Q,TT} C_{\ell,\mathrm{data}}^{TT} + \epsilon^{Q,TQ} C_{\ell,\mathrm{sim}}^{TQ}, \tag{4.2}$$

$$C_{\ell,\mathrm{template}}^{TU} = \epsilon^{Q,TT} C_{\ell,\mathrm{data}}^{TT} + \epsilon^{U,TU} C_{\ell,\mathrm{sim}}^{TU}. \tag{4.3}$$

The templates nominally depend on four free parameters ($\epsilon^{Q,TT},\ \epsilon^{Q,TQ},\ \epsilon^{U,TT},\ \epsilon^{U,TU}$), but the on-sky two-dimensional $C_{\boldsymbol{\ell}}^{TQ}$ and $C_{\boldsymbol{\ell}}^{TU}$ cross-spectra have alternating signs in each quadrant of the Fourier plane and azimuthally average to zero. However our scan-direction timestream filtering (|) aligns with the North-South maxima of the $Q(+)$ patterns but the null crossings of the $U(\times)$ patterns, leading to a non-zero azimuthal average for $C_{\boldsymbol{\ell}}^{TQ}$ while preserving the zero average for $C_{\boldsymbol{\ell}}^{TU}$. As a result C_{ℓ}^{TU} will be zero and we only need to fit for $\epsilon^{Q,TT}$, $\epsilon^{Q,TQ}$, $\epsilon^{U,TT}$. For the polarization rotation angle $\Delta\psi$, the observed Q and U are related to the primordial $\tilde{Q}$ and $\tilde{U}$ by

$$(Q + iU) = (\tilde{Q} + i\tilde{U})e^{i2\Delta\psi}, \tag{4.4}$$

which introduces artificial TB and EB correlations:

$$C_{\ell}^{TB} = -\sin(2\Delta\psi)\tilde{C_{\ell}}^{TE} \tag{4.5}$$

$$C_{\ell}^{EB} = \frac{1}{2}\sin(4\Delta\psi)(\tilde{C_{\ell}}^{BB} - \tilde{C_{\ell}}^{EE}). \tag{4.6}$$

This fit is performed by generating 100 $\left(\epsilon^{Q,TT}, \epsilon^{Q,TQ}, \epsilon^{U,TT}, \Delta\psi\right)$ Latin hypercube training points; these parameters are applied to the data maps (first T-to-P deprojection, then polarization rotation), and the resulting TB and EB data spectra are computed. These 100 spectra are used to train a `CosmoPower` (Mancini et al., 2021) emulator, which can then predict the data spectra for any set of calibration parameters. An MCMC chain is run to find the best fit parameters, which are applied to the data maps. This process is repeated once with the order of the T-to-P deprojection and polarization rotation reversed, producing the final calibration numbers shown in Table 4.2.3. While these calibration parameters are estimated individually for each subfield, we have no reason to expect the polarization angle rotation $\Delta\psi$ to vary with declination, and use a single value for the entire winter field.

The final step is the absolute polarization calibration, which is performed on the Q and U maps in a manner analogous to the absolute temperature calibration described in Equation (4.1). We combine the Q and U coefficients for each subfield to produce a single absolute calibration factor for each band. We then correct these numbers by an additional factor of $\sqrt{0.966}$ to match the final recalibration factor of applied to the 143 GHz *Planck* bandpowers (Equation 45, Planck Collaboration et al., 2019).

4.3 Simulations

In addition to making maps of the real data, we construct a suite of "mock" observations from simulated skies in order to characterize our mapmaking filtering and calculate lensing corrections including the mean-field MC response, and noise bias terms described in Chapter 2. We also use the mock observations to test for biases in our pipeline and estimate the lensing bandpower covariance. We generate three types of simulations: 10 non-Gaussian Agora simulations, 650 Gaussian lensed simulations, and 10 unlensed simulations.

1. *Non-Gaussian Agora simulations for foreground bias characterization.* As discussed in Chapter 2, non-Gaussian foregrounds can bias CMB lensing reconstructions and it is important to characterize the size of these biases in our lensing power spectrum. To do so we use ten SPT-3G size patches from the Agora simulation (Omori, 2024), a suite of extragalactic skies generated using N-body simulations that include tSZ, kSZ, and radio source components correlated with the ray-traced gravitational lensing signal. The various foreground components are fit to the calibrated SPT-3G data power spectra after masking >6 mJy sources $> 10\sigma$ (roughly corresponding to $M_{500c} = 2 \times 10^{14}\,\mathrm{M}_\odot$), such that the power spectrum of the Agora simulations precisely matches the power spectrum of our data maps at 95 GHz, 150 GHz and 220 GHz.

2. *Gaussian Lensed Simulations.* The Gaussian realizations consist of Gaussian realizations of the unlensed CMB, which are then lensed by a Gaussian realization of the lensing potential using `Lenspix` (Lewis, 2005, 2011). The Gaussian realizations of the unlensed CMB and lensing potential are generated using the Planck 2018 `TT,TE,EE+lowE+lensing` best-fit power spectra (Planck Collaboration et al., 2020b;

Planck Collaboration, 2018). Correlated cross-frequency Gaussian foreground power based on the AGORA foreground components discussed above are then added such that the power spectrum of the Gaussian lensed simulations matches both AGORA and the data. We construct 650 simulations in this manner that are split into two groups: 500 `cmb1phi1` realizations and 150 `cmb2phi1` realizations. The 150 `cmb2phi1` realizations share the same lensing potential as the first 150 `cmb1phi1` realizations but have different CMB realizations, and are used to estimate the $\mathbb{N}^1_{\mathbf{L}}$ bias described in Chapter 2.

3. *Unlensed Simulations.* We also construct ten unlensed simulations to test that our lensing pipeline does not spuriously detect a lensing signal in the absence of lensing and does not suffer from additive biases. The unlensed simulations are generated in a similar manner to the Gaussian lensed simulations, but the Gaussian CMB realizations are instead generated with the best-fit *lensed* Planck 2018 power spectra, to ensure the unlensed maps have the proper two-point statistics. We will be generating a larger number of unlensed simulations to properly characterize the covariance of the unlensed pipeline test described in Chapter 4.7.1.

In all cases the simulated skies input to mock-observing are NSIDE 8192 HEALPix maps (Górski et al., 2005). Mock observing is nearly identical to the data mapmaking described in Chapter 4.2.1, with the mock timestreams generated from the input HEALPix maps using bilinear interpolation. This interpolation from input map to timestreams picks up two factors of the NSIDE 8192 HEALPix pixel window function $\mathrm{PWF}(8192)^2$; we deconvolve this from the mock observations as described in the following section. We do not inject point sources into the input skies, but we do mask the locations of point sources in the timestreams in the same manner as the data.

A mock observation is run for a given input map realization, subfield observation, and band, and takes a similar amount of time as a data observation job (2 hours). We coadd the suite of individual subfield mock-observations into a single coadd per input sky realization and band. It is not feasible to run mock observation jobs for all combinations of 670 input map realizations $\times$ 3 frequencies $\times$ 3286 subfield observations, so instead we mock observe 10% of the subfield observations for each input map realization and frequency. I was part of an effort to characterize and optimize the performance of the mock observation pipeline; the improvements included a more efficient timestream-to-pixel binning algorithm, sparse storage of input maps in memory, and elimination of redundant pixel weight calculations. This effort reduced memory requirements from 17 GB to 6 GB, disk space usage by a factor of 3, job runtime by 15%, and increased the number of jobs that could be run in parallel on the Open Science Grid by a factor of 2.

In addition to using the mock observations to calibrate and test our quadratic estimator pipeline, we also use them to estimate a transfer function $\mathbb{T}_{\boldsymbol{\ell}}$ for each band that relates the Fourier transform of the on-sky power to the Fourier transform of the pixelized map for both simulations and data. The transfer function effectively approximates the mapmaking process as a Fourier-diagonal linear operator and encodes effects such as beams, timestream filtering and pixel window functions: $X^{\mathrm{mock}}_{\boldsymbol{\ell}} = \mathbb{T}_{\boldsymbol{\ell}} X^{\mathrm{input}}_{\boldsymbol{\ell}}$. Due to flat-sky projection effects for the $1500\,\mathrm{deg}^2$ patch discussed in Chapter 1.4 the timestream filtering is not diagonal in Fourier space, increasing the degree to which the inverse variance and lensing reconstruction are suboptimal. This is largely due to the fact that the scan-direction timestream filtering that dominates the anisotropic part of the transfer function contributes to all modes with $|\ell_y| > \ell_x$ due to mode-mixing; for a small patch of sky like the $100\,\mathrm{deg}^2$ SPTpol patch discussed in Chapter 3 the timestream filtering instead nulls modes with $\ell_x < 300$ in a relatively clean manner.

There are two ways to define the transfer function, depending on how one wants to address the mode-mixing induced by the ProjZEA projection. If we are only interested in the amount of *correlated* signal transferred from a given mode in the input map to the mock observation, the transfer function should depend on the cross spectrum between the two maps. On the other hand if we are instead interested in the relative *total* power, the transfer function should depend on the auto-spectrum of the mock observation instead. This leads to the following two transfer function definitions:

$$\mathbb{T}_{\boldsymbol{\ell}}^{\text{cross}} = \frac{\left\langle X_{\boldsymbol{\ell}}^{\text{mock}} X_{\boldsymbol{\ell}}^{\text{input}*} \right\rangle}{\left\langle X_{\boldsymbol{\ell}}^{\text{input}} X_{\boldsymbol{\ell}}^{\text{input}*} \right\rangle}, \tag{4.7}$$

$$\mathbb{T}_{\boldsymbol{\ell}}^{\text{auto}} = \left(\frac{\left\langle X_{\boldsymbol{\ell}}^{\text{mock}} X_{\boldsymbol{\ell}}^{\text{mock}*} \right\rangle}{\left\langle X_{\boldsymbol{\ell}}^{\text{input}} X_{\boldsymbol{\ell}}^{\text{input}*} \right\rangle} \right)^{1/2}. \tag{4.8}$$

with $\langle \ldots \rangle$ signifying an average over realizations. We show $T_{\boldsymbol{\ell}}^{\text{cross}}$ and the ratio $T_{\boldsymbol{\ell}}^{\text{cross}}/T_{\boldsymbol{\ell}}^{\text{auto}}$ in Figure 4.5 and note that the two estimates differ by roughly 10% in modes that are hit by the timestream lowpass filter in some parts of the map but not others. Since one of the steps in the inverse-variance filtering described in Chapter 2.4.1 is deconvolving the transfer function from the data, this 10% propagates directly into the inverse-variance filtered maps and thus into the quadratic estimate and its normalization.

The choice of transfer function has implications for the inverse-variance filtering and lensing reconstruction normalization discussed in Chapters 2.4.1 and 2.4.2. We find that with when using $\mathbb{T}_{\boldsymbol{\ell}}^{\text{auto}}$, the power spectrum of the Wiener filtered maps $C_\ell^2 \left\langle \overline{X}_{\boldsymbol{\ell}} \overline{X}_{\boldsymbol{\ell}}^* \right\rangle$ agrees perfectly with theory, whereas using $\mathbb{T}_{\boldsymbol{\ell}}^{\text{cross}}$ gives a Wiener filtered power spectrum that exceeds theory. This is because the $\mathbb{T}_{\boldsymbol{\ell}}^{\text{cross}}$ Wiener filtered maps have the correct amount of *correlated* signal, but the uncorrelated power from ProjZEA mode-mixing in the power spectrum is boosted by the transfer function deconvolution giving too much *total* power in the maps. We find that our final lensing bandpowers are not biased in any way by the choice of transfer function, but using $\mathbb{T}_{\boldsymbol{\ell}}^{\text{auto}}$ slightly reduces the lensing map noise and power spectrum error bars. As such we choose to use $\mathbb{T}_{\boldsymbol{\ell}}^{\text{auto}}$ for the remainder of the analysis.

One complication to using $\mathbb{T}_{\boldsymbol{\ell}}^{\text{auto}}$ is that the T, E, and B transfer functions are not consistent with each other as is the case for $\mathbb{T}_{\boldsymbol{\ell}}^{\text{cross}}$. This has been seen in multiple curved-sky pipelines and thus cannot be related to the ProjZEA-specific mode-mixing discussed above. We suspect this discrepancy to be caused by mode-mixing in the mock observing process, perhaps from incomplete pixel coverage related to only using 10% of the subfield observations for mock observing. We choose to use the T transfer function for E and B to avoid propagating these poorly understood effects into the lensing reconstruction, but acknowledge that this is an imperfect solution and we hope to gain a better understanding of the origin of this discrepancy for future analyses.

4.4 Map Processing

Three main processing steps are applied to the data and mock-observed coadds before the inverse-variance filtering. The individual frequency maps are combined into a minimum-variance estimate of the CMB, point sources and clusters are inpainted, and the maps are anti-aliased and downsampled. These steps are applied to both the data and mock-observed maps.

We form a linear combination of the frequency bands to produce the best possible estimate of the CMB by minimizing the variance contributions from foregrounds and noise. Also known as component separation (Cardoso et al., 2008; Planck Collaboration et al., 2014b), linear combination techniques leverage the different frequency dependence of various signals in the data. Here we follow Prabhu et al. (2024) and choose a minimum-variance combination that isolates signals with the same frequency dependence as the CMB, but one can also choose to null specific foreground contributions such as tSZ or CIB (Raghunathan & Omori, 2023). We form the minimum variance combination as a weighted sum of the three frequency bands:

$$X_{\boldsymbol{\ell}}^{\mathrm{MV}} = \sum_i \mathbb{W}_{\boldsymbol{\ell}}^i X_{\boldsymbol{\ell}}^i, \tag{4.9}$$

where the length-3 vector of weights obeys $\sum_i \mathbb{W}_{\boldsymbol{\ell}}^i = 1$ and is calculated as

$$\mathbb{W}_{\boldsymbol{\ell}} = \frac{\mathbb{C}_{\boldsymbol{\ell}}^{-1}\mathbb{A}_s}{\mathbb{A}_s^{\dagger}\mathbb{C}_{\boldsymbol{\ell}}^{-1}\mathbb{A}_s}. \tag{4.10}$$

Here $\mathbb{C}_{\boldsymbol{\ell}}$ is a 3×3 covariance matrix of the foregrounds and noise at a given $\boldsymbol{\ell}$ and $\mathbb{A}_s = [1\ 1\ 1]$ is

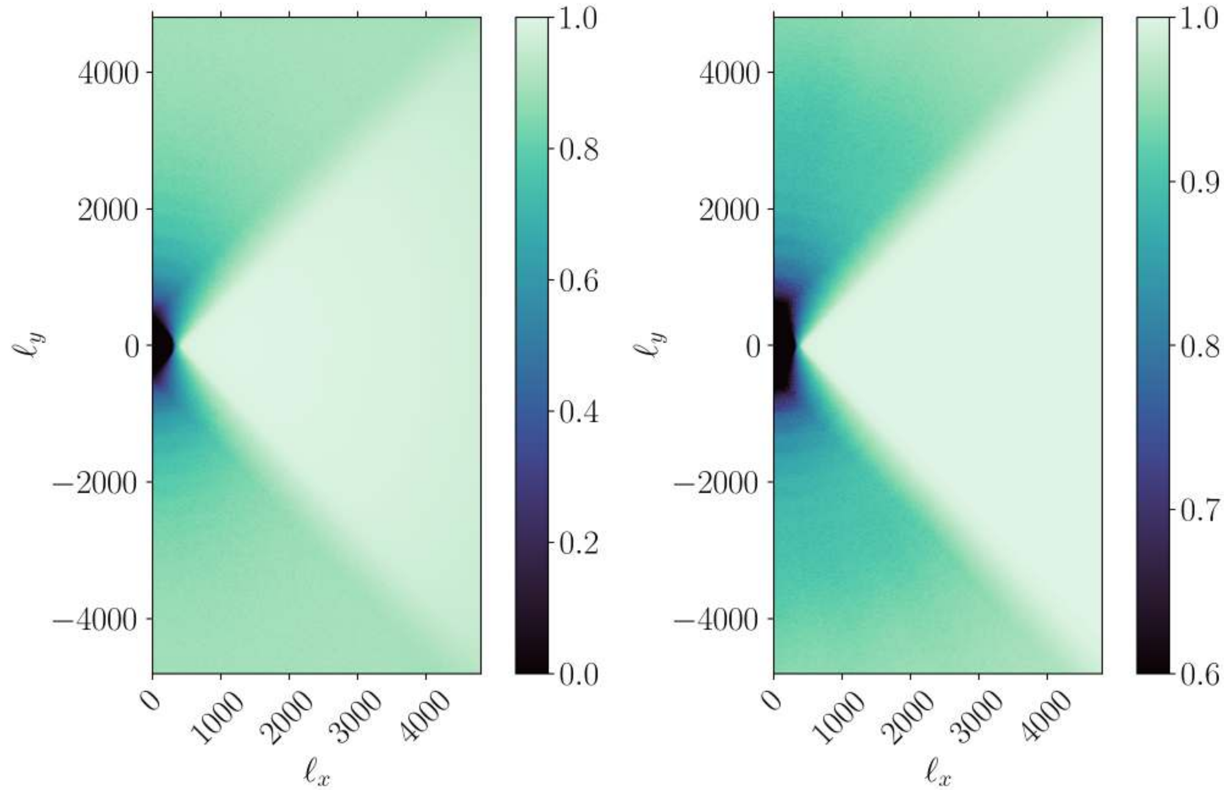

Figure 4.5: Cross-spectrum transfer function (left) and ratio between cross- and auto-spectrum transfer functions (right) for the SPT-3G 2019+2020 flat-sky lensing analysis. The low-amplitude region near the origin of the cross-spectrum transfer function (left) corresponds to modes that are nulled by the timestream highpass filter at all locations in the map. The higher-amplitude wedge where $-\ell_x < \ell_y < \ell_x$ includes modes that are not hit by the timestream lowpass filter anywhere in the map, whereas the intermediate-amplitude region outside this wedge includes modes that are hit by the lowpass filter in some but not all regions of the map. The ratio between the cross- and auto-spectrum transfer functions (right) shows the degree to which the approximation of the transfer function as Fourier-diagonal is breaks down–the two estimates are identical within the main wedge of the transfer function, but outside this wedge at least 10% of the power is uncorrelated with the input, indicating the level of mode-mixing induced by the ProjZEA projection.

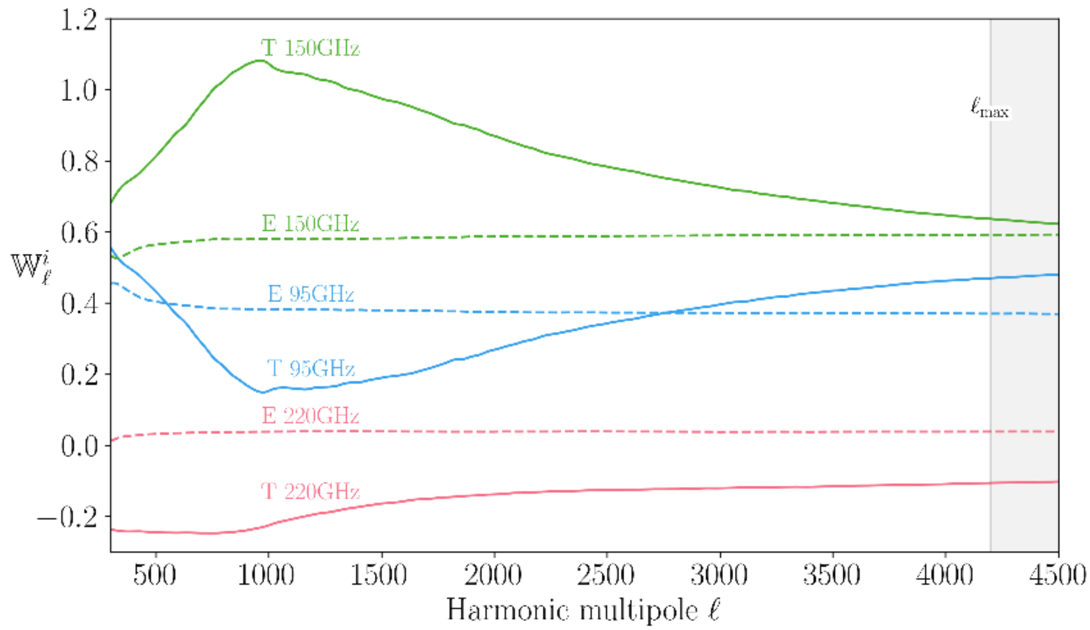

Figure 4.6: Minimum variance linear frequency combination weights as a function of harmonic multipole. In temperature the weights are tuned to minimize noise and foregrounds that are correlated across frequencies, with negative 220 GHz weights to suppress the correlated atmospheric $1/f$ noise and CIB. In polarization, the weights are featureless and correspond to an inverse-variance combination of the per-frequency white noise levels. Scales that are lowpass-filtered in the downsampling step are shaded gray.

the frequency response of the CMB. The linear combination technique also produces estimates of the residual noise and foreground power $\left(\mathbb{A}_s^\dagger \mathbb{C}_{\boldsymbol{\ell}}^{-1} \mathbb{A}_s\right)^{-1}$ in the minimum-variance map, which we later use as inputs to the inverse-variance filtering. We show the one-dimensional weights $\mathbb{W}^i_{\boldsymbol{\ell}}$ in Figure 4.6. The 220GHz band is effectively used as noise- and foreground-subtraction template in temperature, since both the correlated atmospheric $1/f$ noise and the CIB are strongest at 220 GHz (Reichardt et al., 2021).

The noise covariance is estimated from cross-spectra of the signflip noise realizations. We show the per-frequency and minimum variance noise power spectra in Figure 4.7, with correlated atmospheric $1/f$ noise clearly visible on large scales in temperature. We achieve minimum-variance white-noise levels of 4.1 μK-arcmin in temperature and 5.0 μK-arcmin in polarization; we note that the minimum variance combination breaks the typical $\sqrt{2}$ relation between temperature and polarization noise levels. This is because the the temperature combination jointly minimizes noise *and* foregrounds which are both correlated across frequencies, whereas in polarization the foregrounds are negligible and the combination is close to a pure inverse-variance combination of white noise levels.

The foreground covariance is generated from fits to the data auto- and cross-spectra as discussed in Chapter 4.3. Any mis-modeling of the foreground covariance will make the linear combination sub-optimal, but will not introduce a bias. We show the combined power spectra of the individual frequency maps and their minimum variance combination in Figure 4.8, as well as the foreground and noise components. The minimum variance combination is consistent with the individual frequency maps on large scales, but has reduced noise and foreground power on small scales.

Point sources and clusters are masked in timestream *filtering* during the mapmaking process, but the sources and clusters are still present in the map, and the masked regions have a different transfer function than the rest of the map. Masking these locations in the CMB complicates the reconstruction mean-field, so instead we inpaint sources and clusters to avoid contaminating the lensing reconstruction. Inpainting involves filling in the masked regions with a Gaussian constrained realization (Hoffman & Ribak, 1991; Benoit-Lévy et al., 2013; Raghunathan et al., 2019) of the CMB, foregrounds, and noise, such that the inpainted regions have the same two-point statistics as the rest of the map. Inpainting for a given source is performed by defining a circular region X_1 to be inpainted with radius R_1 equal to the timestream masking radius, and a surrounding annular region X_2 extending from R_1 to $R_1 + 2'$. A Gaussian simulation $X^{\rm sim}$ with the same power spectrum as the data $X^{\rm data}$ is generated, and Gaussian constrained realization of the inner region is calculated as

$$X_1^{\rm inpainted} = X_1^{\rm sim} + \mathbf{C}_{12}\mathbf{C}_{22}^{-1}\big(X_1^{\rm data} - X_1^{\rm sim}\big), \tag{4.11}$$

where $\mathbf{C}_{12}$ is the pixel-space covariance between the inner and outer regions, and $\mathbf{C}_{22}$ is the pixel-space covariance of the outer region estimated directly from the data. We note that

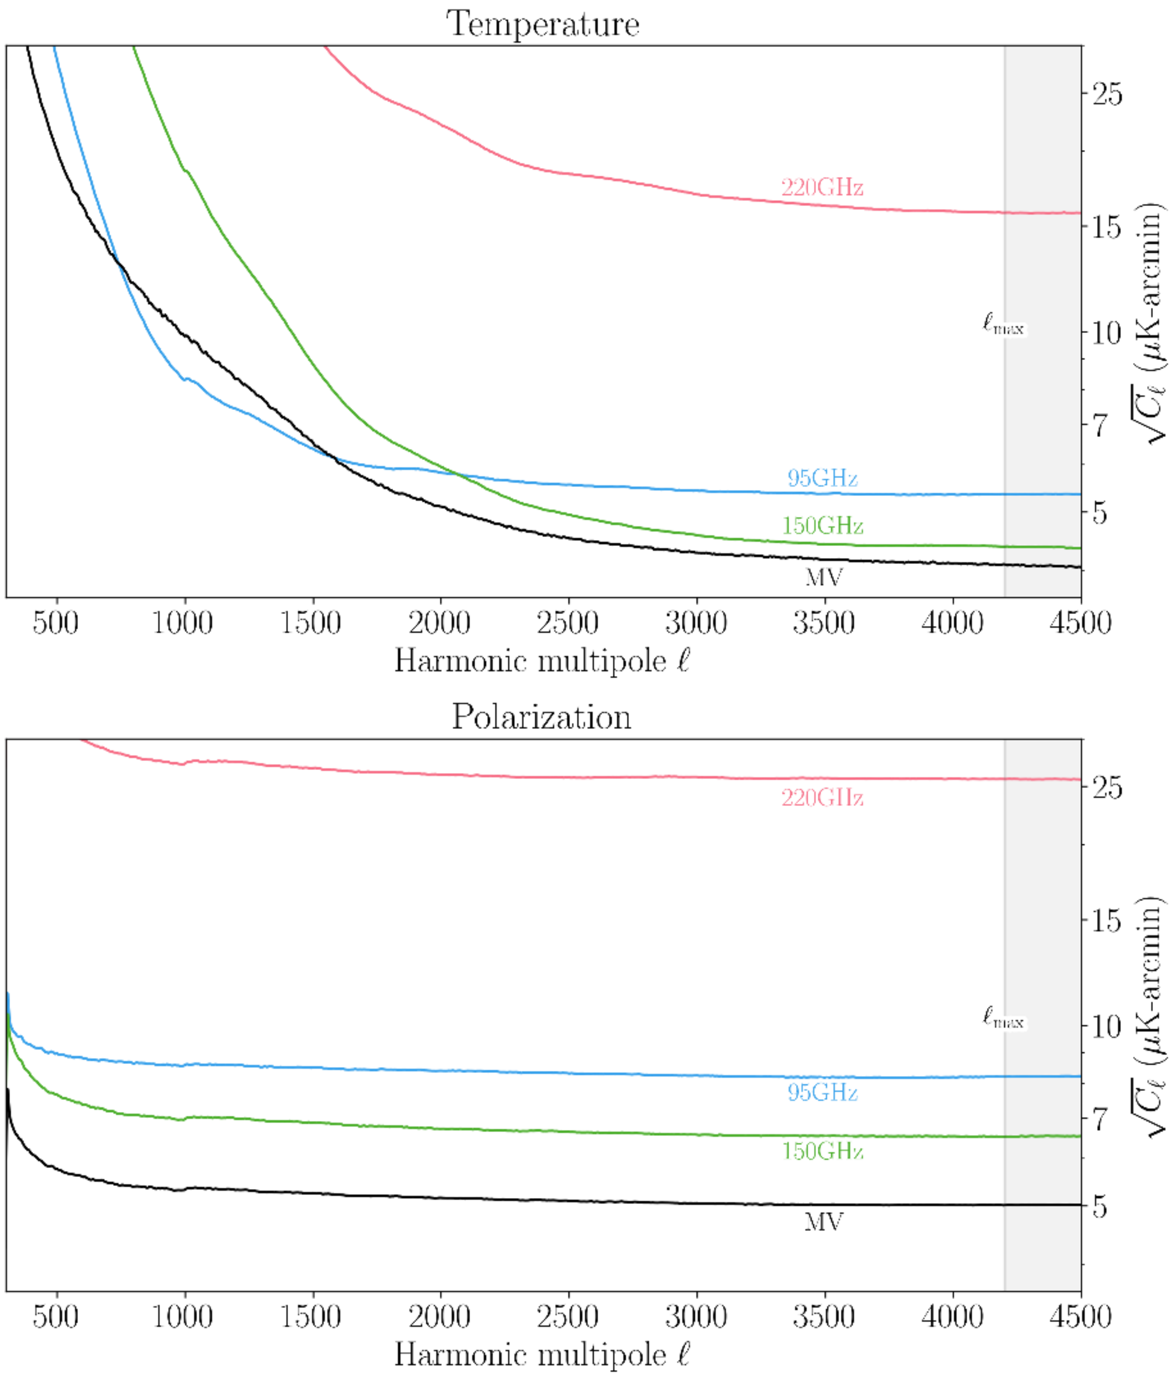

Figure 4.7: Transfer function-deconvolved noise power spectra of the individual frequency maps and their minimum variance (MV) linear combination as a function of harmonic multipole, in units of μK-arcmin. Atmospheric $1/f$ noise is clearly visible at large scales in temperature, while polarization is almost entirely composed of white noise. The cross-frequency spectra are not shown here, but the atmospheric noise is correlated across frequencies (particularly between 220 GHz and 150 GHz); these correlations are leveraged to suppress the noise in the minimum variance combination. Scales that are lowpass-filtered in the downsampling step are shaded gray.

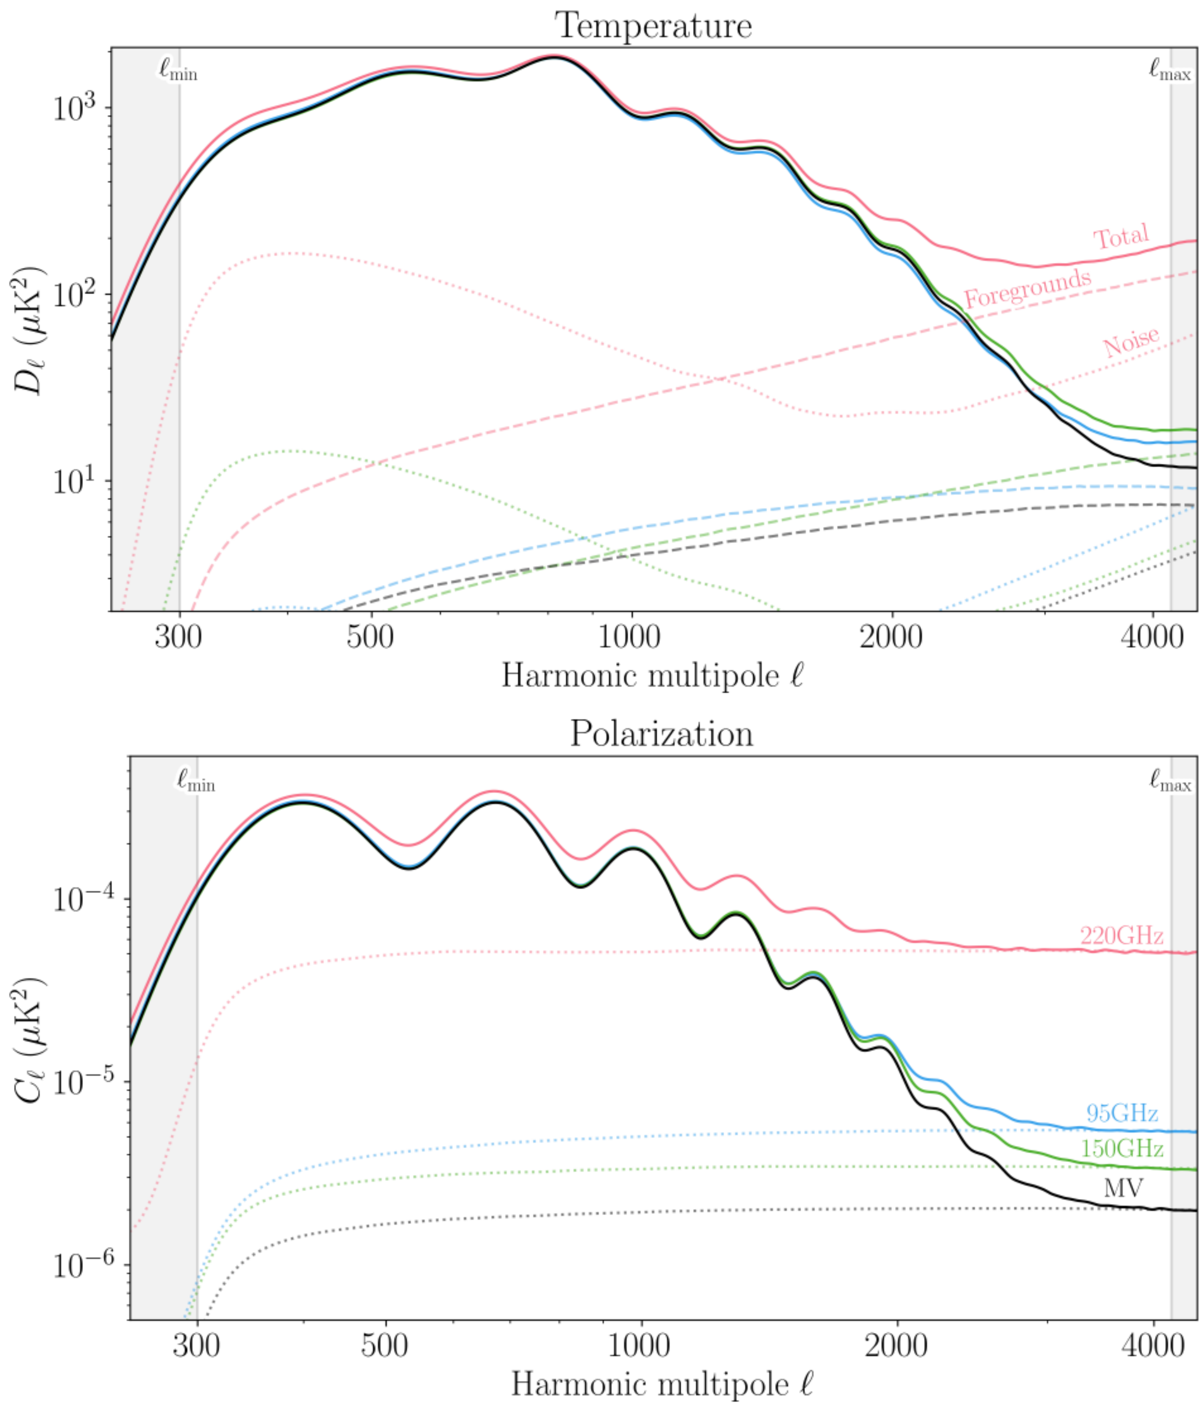

Figure 4.8: Combined power spectra of the individual frequency maps and their minimum variance (MV) linear combination as a function of harmonic multipole. 95 GHz, 150 GHz and 220 GHz are shown blue, green and pink respectively, while the minimum variance combination is shown in black. The total, foreground, and noise power are shown in solid, dashed, and dotted lines respectively. Scales that are lowpass-filtered in the downsampling step are shaded gray. While polarized extragalactic foregrounds are included in our foreground modeling, they are orders of magnitude below the signal and are not visible in the lower plot.

lensing information is not preserved, since the inpainted estimate is Gaussian; we mask the majority of inpainted regions in the lensing reconstruction.

We choose to inpaint the minimum variance combination to avoid jointly inpainting the individual frequency maps; however, the ℓ-dependence of the linear combination weights tends to spread around point source power, so we perform an initial approximate inpainting step to remove point sources from the individual frequency maps before forming the minimum variance combination. The approximate inpainting is performed by masking sources, smoothing a copy of the masked map, and then replacing the masked regions with the smoothed map. This process is performed iteratively by smoothing on different scales (to capture different Fourier modes in the map) until the masked regions are filled in. After this initial approximate inpainting step, we form the minimum variance combination and perform full Gaussian constrained inpainting on the combined map.

The final processing step is to anti-alias the map and downsample them to a lower resolution. The maps are initially made at $0.5625'$ resolution; this is beyond what is needed for lensing reconstruction since small scales in temperature map are contaminated by foregrounds and small scales in polarization are noise-dominated. We downsample the maps by a factor of 4 to $2.25'$ resolution, first applying a lowpass filter to remove power close to the Nyquist frequency of the downsampled map. The lowpass filter slowly tapers to zero between $\ell = 4300$ and $\ell = 4800$ (the Nyquist frequency of the downsampled map) using a Tukey window function. For mock observations, we also deconvolve the factor of $\mathrm{PWF}(8192)^2$ picked up when interpolating the input maps into mock timestreams at this step, since these factors of the HEALpix pixel window function are not present in the data maps. To downsample the maps, we average together 4×4 blocks of pixels in the original map to form a single pixel in the downsampled map. We show the inpainted and downsampled T, Q, and U maps in Figure 4.9.

This binning operation picks up the pixel window function of the output $2.25'$ map, which we apply to the residual noise and foreground power spectra used in the inverse-variance

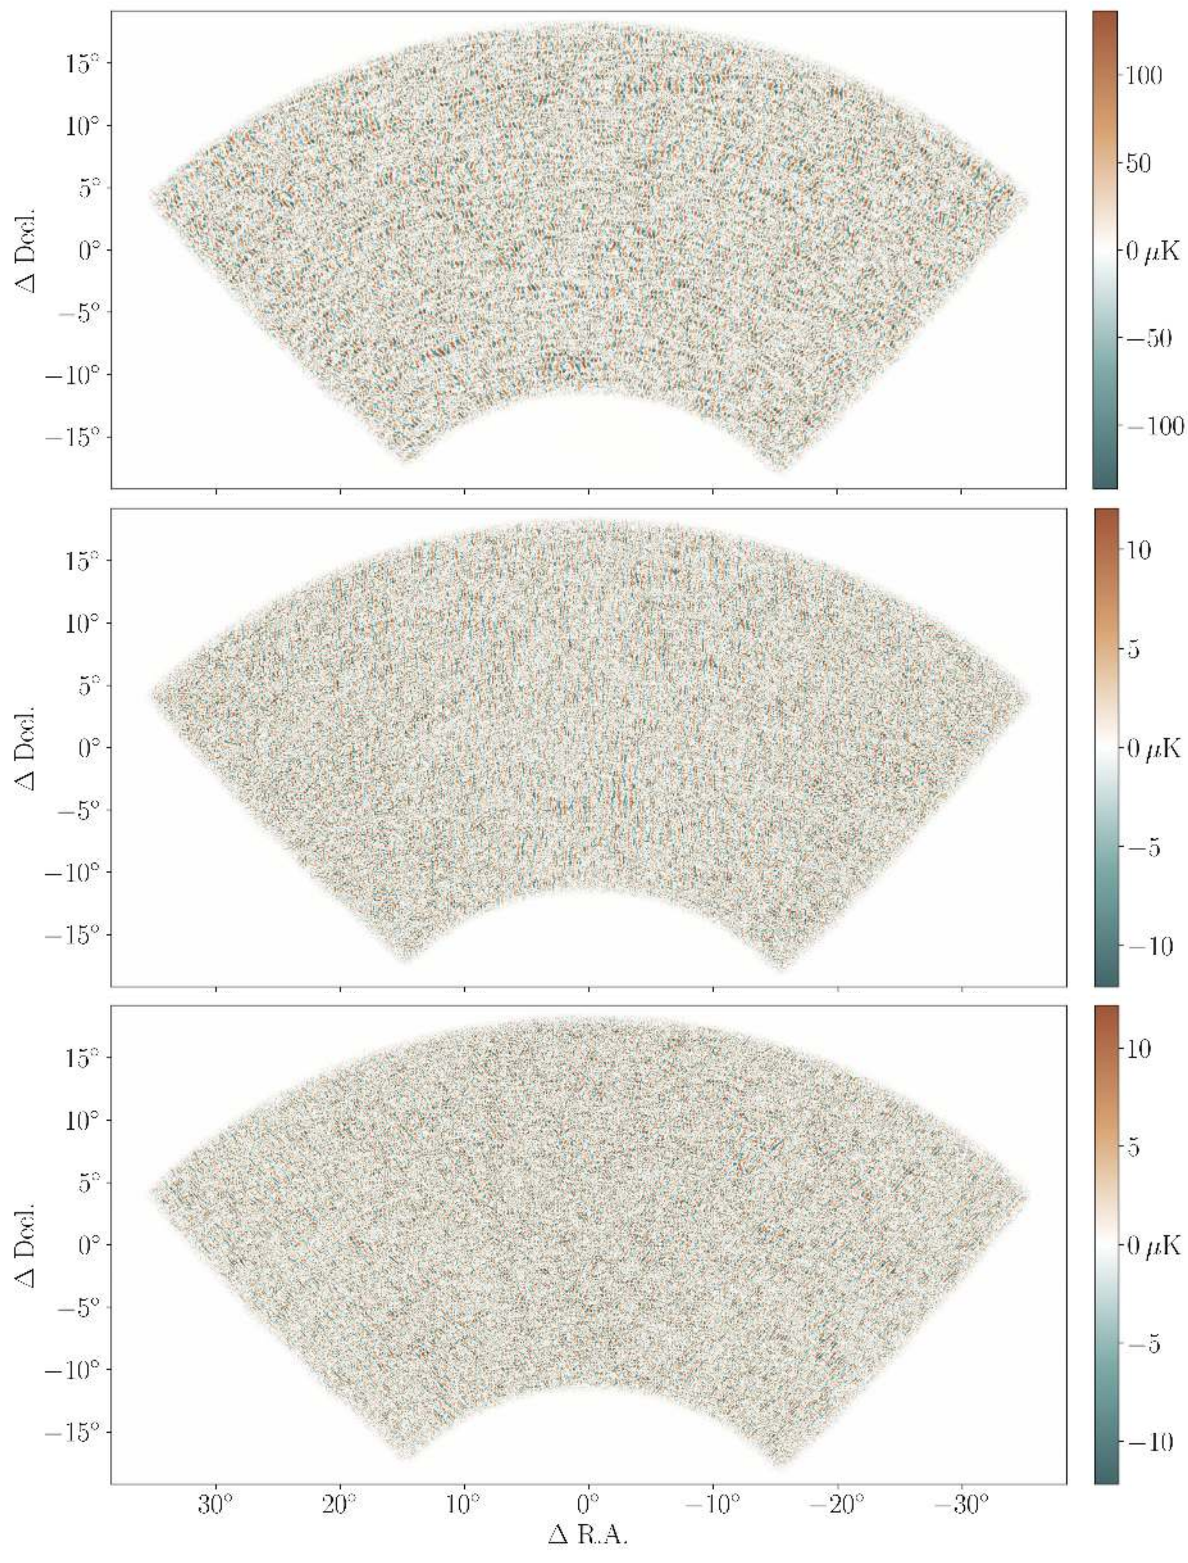

Figure 4.9: Downsampled and inpainted SPT-3G 2019+2020 T, Q, and U maps from top to bottom. The polarization maps have not been flattened at this stage, so the QU patterns are still oriented relative to the cardinal directions.

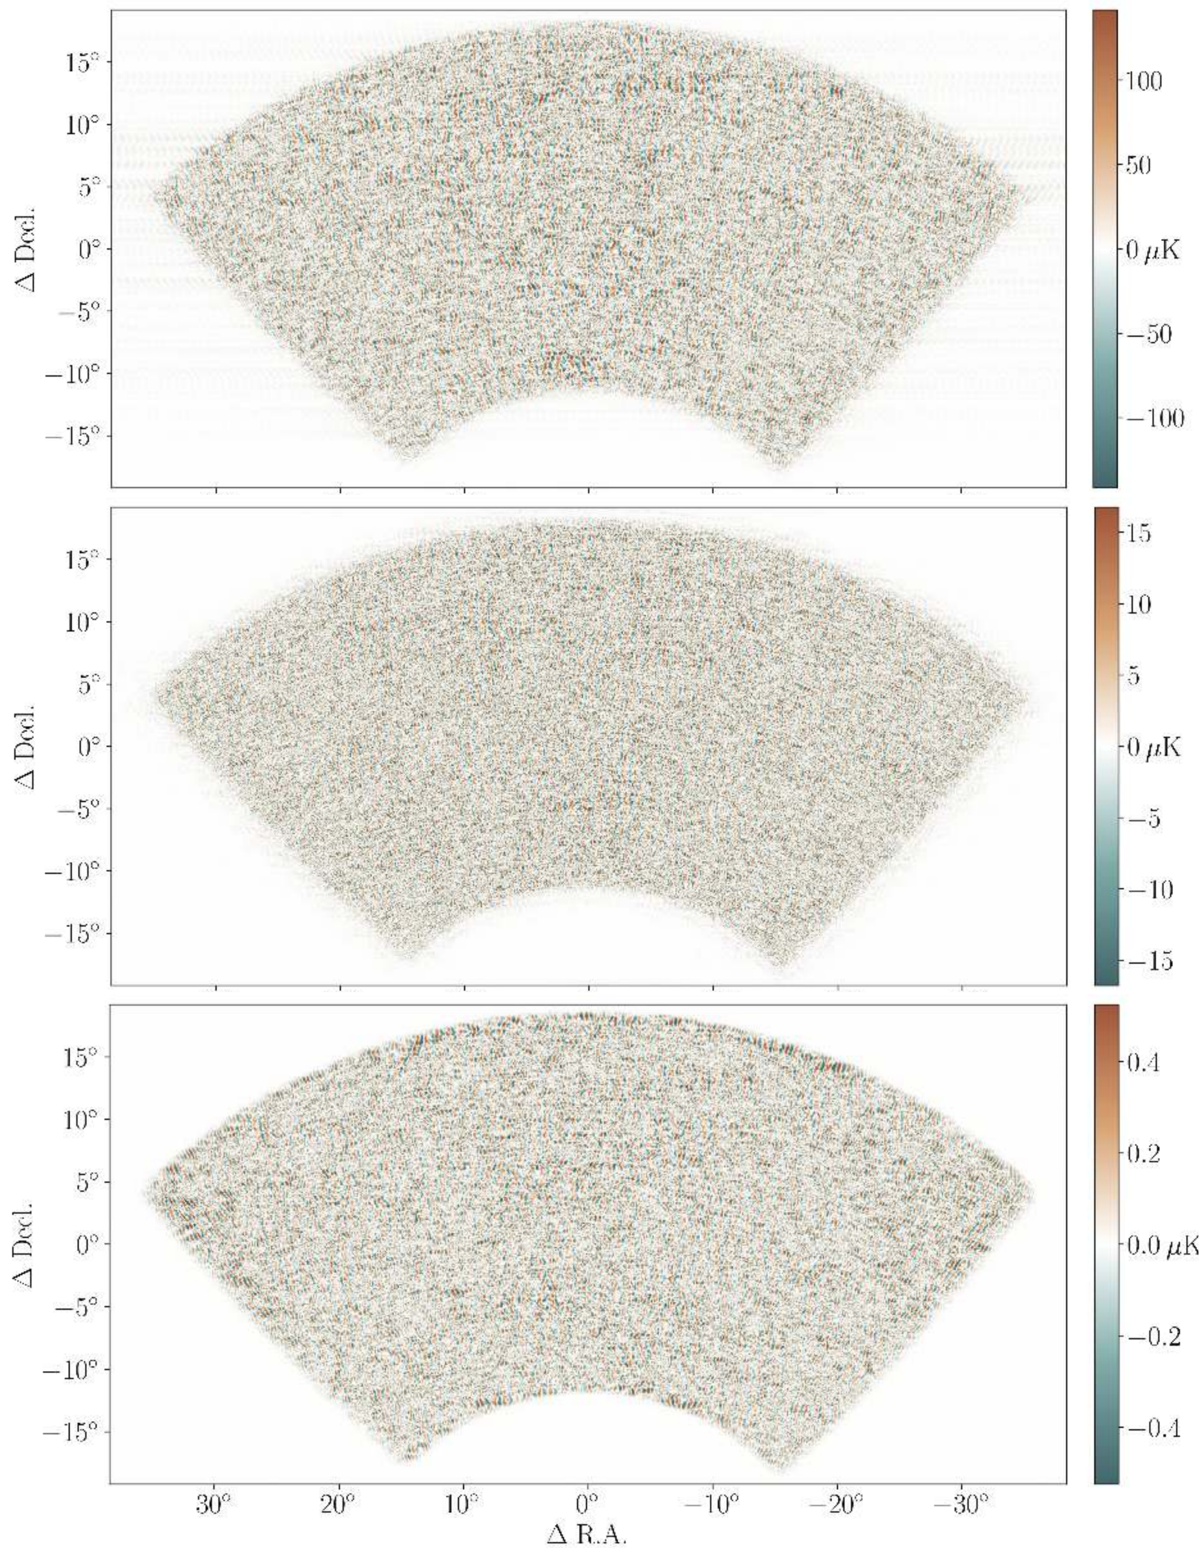

Figure 4.10: Wiener filtered SPT-3G 2019+2020 T, E, and B maps from top to bottom. Indications of $E \to B$ leakage are evident at the mask boundary in B, but we find no indication that this biases the lensing reconstruction.

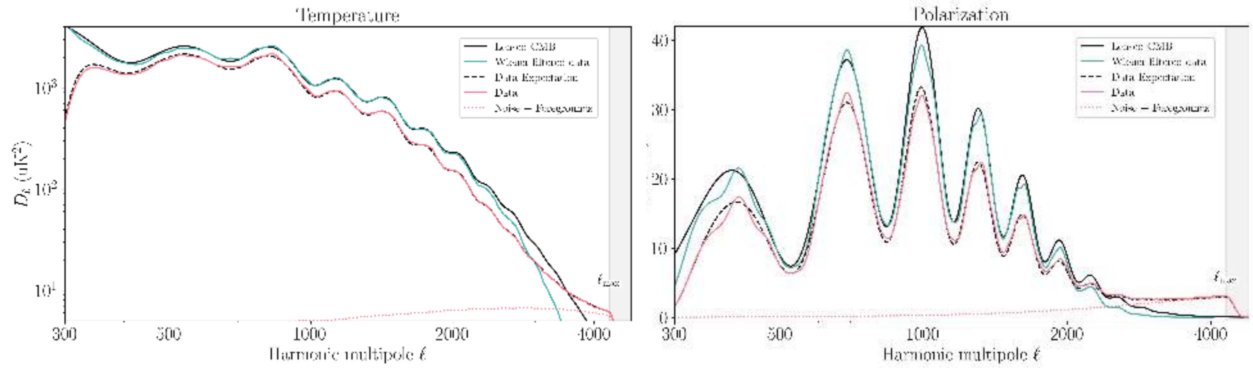

Figure 4.11: Power spectra of the SPT-3G 2019+2020 downsampled and Wiener filtered T and E maps as a function of harmonic multipole ℓ. We find good agreement between the downsampled power spectra and our expectation of the lensed CMB, foregrounds, noise, and effective transfer function. As a signal-to-noise weighted quantity, the Wiener filtered power spectra agree with the theoretical lensed CMB power spectra on signal-dominated scales, but are downweighted on small scales where noise and foregrounds are larger.

filtering for consistency. At this stage we also calculate a final effective transfer function for the anti-aliased and downsampled minimum-variance map which incorporates all effects that transform an idealized sky into our final map:

$$\mathbb{T}^{\mathrm{MV}}_{\boldsymbol{\ell}} = \mathrm{PWF}_\ell(2.25') \cdot F^{\mathrm{AA}}_\ell \cdot \sum_i \mathbb{W}^i_{\boldsymbol{\ell}} \mathbb{T}^i_{\boldsymbol{\ell}} B^i_\ell, \tag{4.12}$$

where $\mathrm{PWF}_\ell(2.25')$ is the output map pixel window function, F^{AA}_ℓ is the anti-aliasing filter, B^i_ℓ is the per-frequency beam, and $\mathbb{T}^i_{\boldsymbol{\ell}}$ is the per-frequency transfer function. We show the power spectra of the downsampled data maps compared to our expectation of the lensed CMB, foregrounds, noise, and effective transfer function in Figure 4.11. We find good agreement between data and expectation.

4.5 Lensing Map Reconstruction

The first step in the lensing reconstruction process is to deconvolve the effective transfer function and inverse-variance filter the downsampled T, Q, and U maps as described in Chapter 2.4.1. In addition to the maps themselves, the filter takes as input lensed CMB power spectra, the effective transfer function $\mathbb{T}_{\boldsymbol{\ell}}^{\mathrm{MV}}$, two-dimensional transfer function-convolved atmospheric noise and foregound power spectra $\mathbb{N}_{\boldsymbol{\ell}}^{\mathrm{Atm}}$ and $\mathbb{S}_{\boldsymbol{\ell}}^{\mathrm{FG}}$, and a pixel-space white noise level n. We estimate the pixel-space white noise as the mean square root of the total noise power spectrum $\mathbb{N}_{\boldsymbol{\ell}}$ in the range $\ell \in [4000, 4200]$, and then define the residual atmospheric noise as $\mathbb{N}_{\boldsymbol{\ell}}^{\mathrm{Atm}} = \mathbb{N}_{\boldsymbol{\ell}} - n^2$. The white noise level is combined with the survey mask to create spatially-varying estimate of the map variance.

We iteratively solve Equation (2.22) with a conjugate gradient method until either the dot product of the residuals has decreased by a factor of 10^8 relative to the first iteration, or 100 iterations have been performed. We mask the $\ell_x < 300, |\ell_y| < 1000$ region where the timestream highpass filter is relatively well localized (Figure 4.5) as well as $\ell > 4000$ in the solver; these regions are also ignored by the quadratic estimator. We show the Wiener filtered data maps $C_\ell \overline{X}_{\boldsymbol{\ell}}$ in Figure 4.10 and power spectra $C_\ell^2 \left\langle \overline{X}_{\boldsymbol{\ell}} \overline{X}_{\boldsymbol{\ell}}^* \right\rangle$ in Figure 4.11, where $\overline{X}_{\boldsymbol{\ell}}$ is the inverse variance filtered map. In the signal-dominated regime, the Wiener filtered power spectra are consistent with the theoretical lensed CMB power spectra used in the simulations which gives confidence that our simulations agree well with the data and will allow us to accurately estimate the lensing normalization, mean-field, and bias terms.

We evaluate the quadratic estimator, normalize the lensing map, and remove the mean-field bias as described in Chapter 2.4. The temperature, polarization, and minimum variance convergence maps are shown in Figure 4.12.

We estimate the mean-field, shown in Figure 4.13, from the suite of 500 Gaussian lensed simulations. The mean-field corrects for structure in the raw lensing map that remains constant across realizations of the lensed CMB; in addition to large-scale contributions from the survey mask, there is a noticeable mean-field at the location of point-sources and galaxy clusters where the timestream filtering is masked in mapmaking. While there are no sources injected into the simulations, these regions are still masked during mock-observing. In both data and simulations, the timestream filtering mask couples with the CMB and introduces small artifacts into the map, adding non-Gaussian structure to the CMB maps at the location

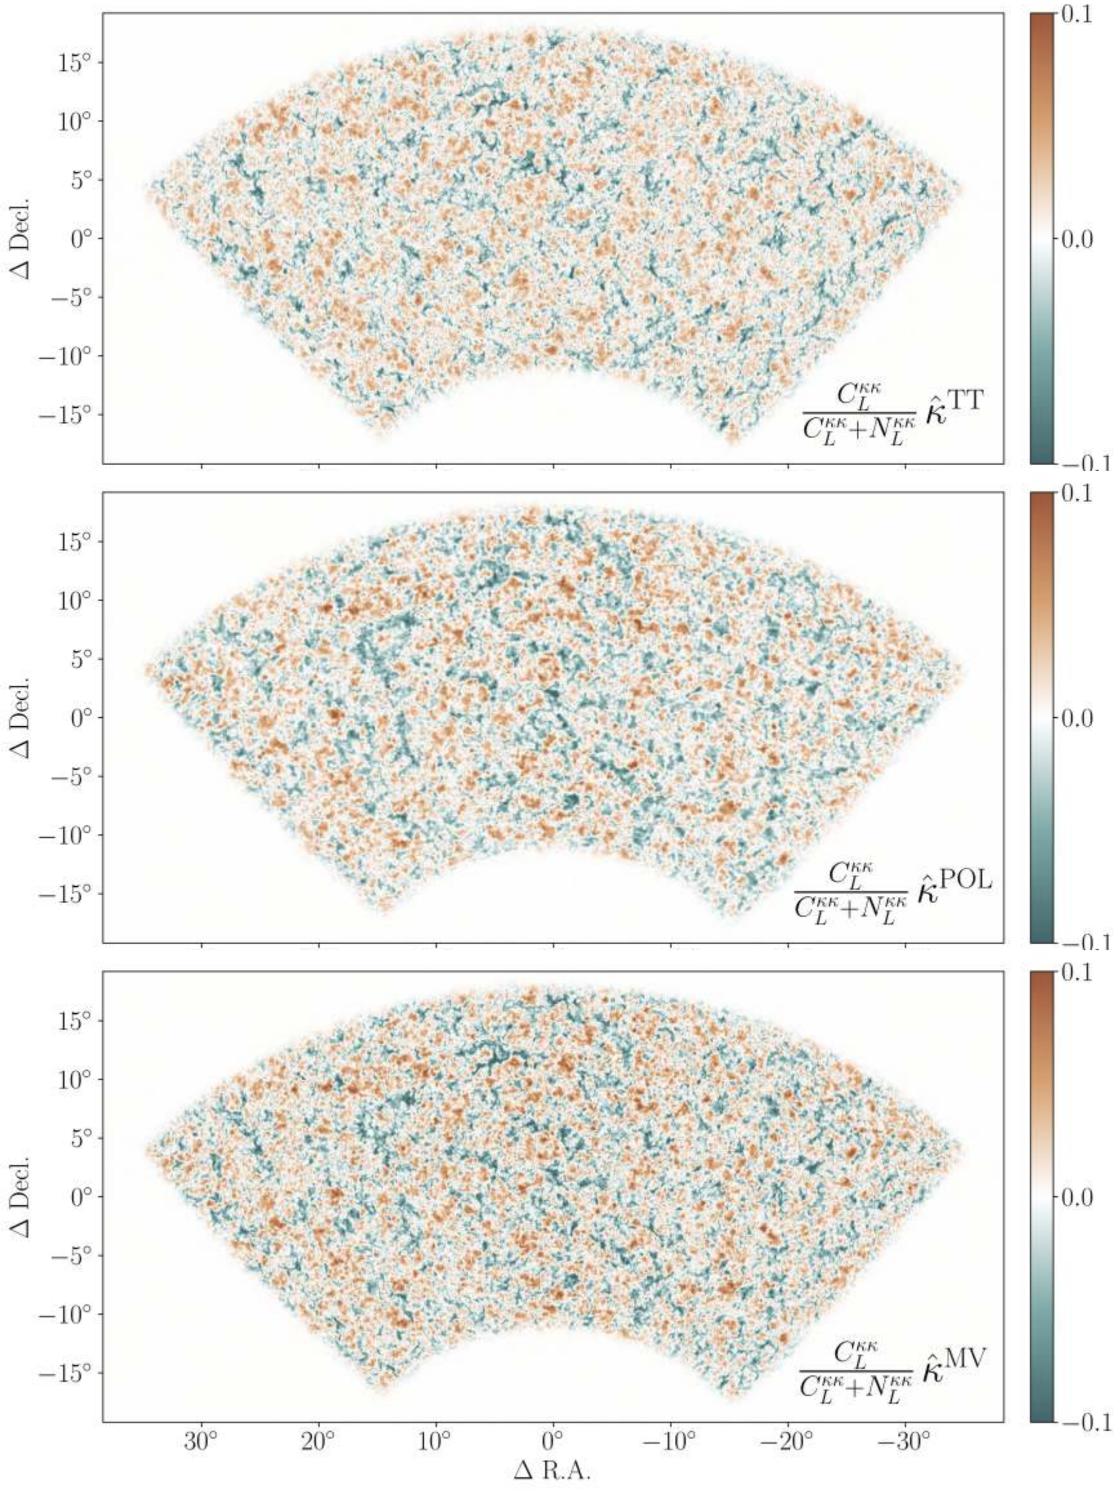

Figure 4.12: Convergence maps for the SPT-3G 2019+2020 flat-sky lensing analysis. The temperature $\hat{\kappa}^{\mathrm{TT}}_{\mathbf{L}}$, polarization $\hat{\kappa}^{\mathrm{POL}}_{\mathbf{L}}$, and minimum variance $\hat{\kappa}^{\mathrm{MV}}_{\mathbf{L}}$ maps are shown with a simple Fourier-space Wiener filter applied to highlight the signal-dominated scales and relative signal-to-noise of the estimators. The temperature estimator has better signal-to-noise on small scales, while the polarization estimator better signal-to-noise on large scales.

of point sources and thus contributing to the mean-field. However, because the coupling depends on the specific *realization* of the CMB, the mean-field will not perfectly remove the artifacts from any given realization. To minimize the impact of these artifacts on the lensing reconstruction, we mask point source locations when plotting lensing maps and estimating lensing power spectra.

We also estimate an Monte Carlo (MC) correction $\mathbb{R}^{\mathrm{MC}}_{\mathbf{L}}$ to the quadratic estimator normalization using the 500 Gaussian lensed simulations. The ProjZEA transfer function is highly anisotropic, and this propagates through to the the MC response and the lensing map noise (Figure 4.14). As a result we estimate all normalization and noise bias terms in two dimensions. However the raw two-dimensional MC response is too noisy to be used directly, and simple smoothing is not sufficient because the response exhibits structure on a wide range of scales. Instead an fit a multi-scale azimuthally-varying model to the MC response. Azimuthal variation is modelled with multipole-like weights $\mathbb{W}^{n,c}_{\mathbf{L}} = \cos(n\psi_{\mathbf{L}})$ and $\mathbb{W}^{n,s}_{\mathbf{L}} = \sin(n\psi_{\mathbf{L}})$, where $\psi_{\mathbf{L}} = \mathrm{atan2}(L_y/L_x)$ is the angle of a given mode in Fourier space. These weights are applied to the raw MC response and azimuthally averaged to produce a one-dimensional vector, to which we fit a multi-scale Gaussian process. The one-dimensional fit is then applied back to $\mathbb{W}^{n}_{\mathbf{L}}$ to produce a 2-dimensional model. We do the Gaussian process fit in one dimension because it is not computationally feasible to fit a Gaussian process to the full two-dimensional MC response. This fit is performed iteratively for both sines and cosines with $n \in [1, 5]$ until the model residuals are consistent with noise. We acknowledge that this is an awkward approach, but the resulting model is a reasonable match to the noisy MC response (Figure 4.14) and there is no apparent structure in the final residuals.

Before continuing to power spectrum estimation, we perform a basic check on our maps by cross-correlating the reconstructions with the input convergence (for simulations) and an external tracer of large-scale structure (for data). We choose to cross-correlate with the 545 GHZ CIB map from Lenz et al. (2019), and also cross-correlate the minimum variance Planck Collaboration et al. (2020c) lensing map with the same CIB map. To

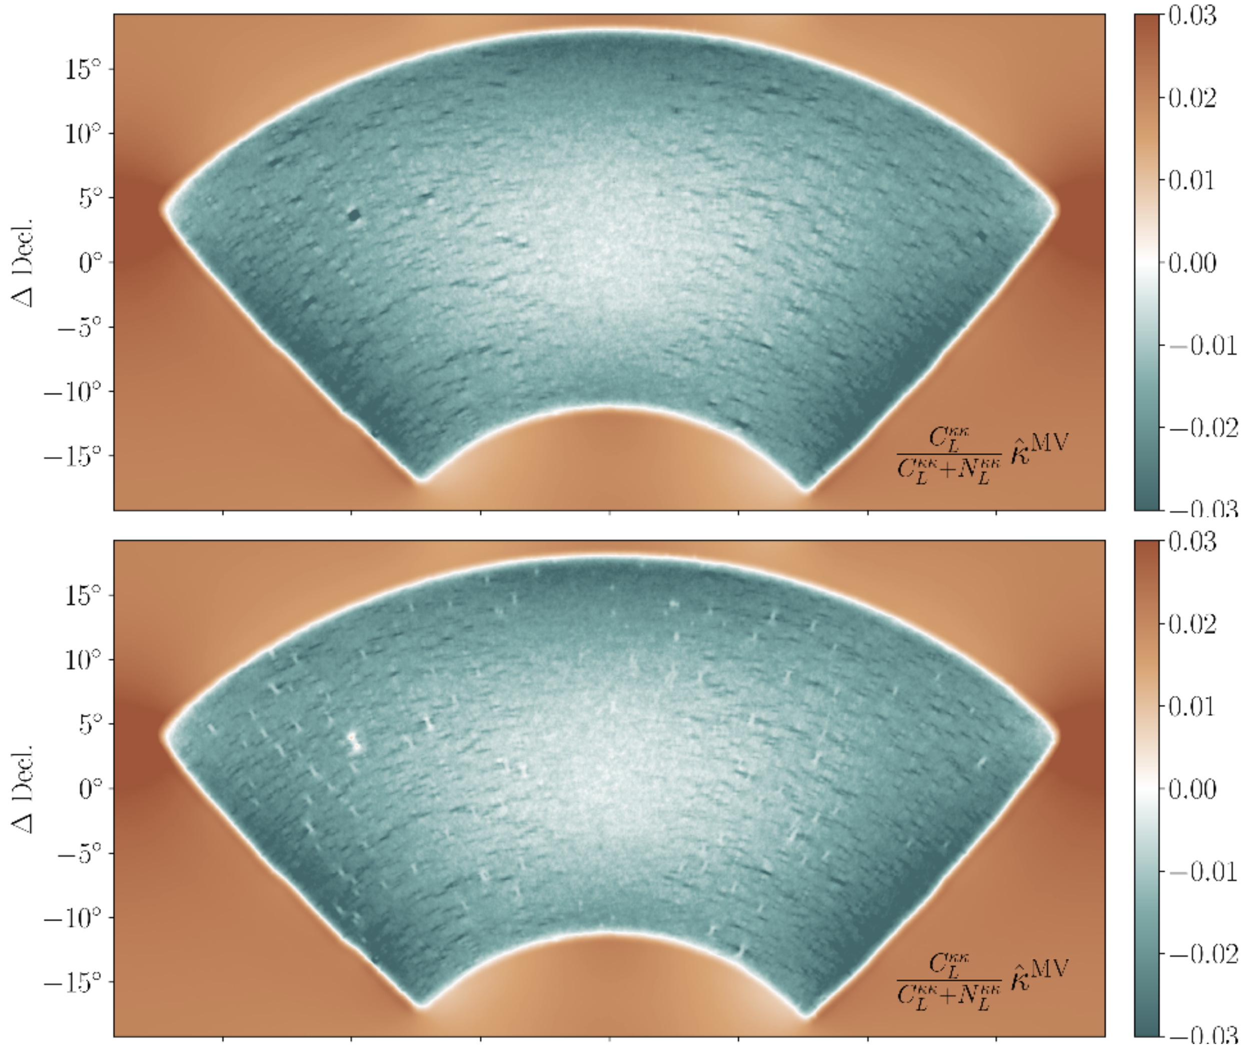

Figure 4.13: Mean-fields for the SPT-3G 2019+2020 flat-sky lensing analysis. The minimum-variance mean-field (top) reveals large-scale structure in the raw lensing map from the survey mask as well as localized artifacts from point source masking. Since these artifacts are caused by a coupling between the CMB and timestream filtering, they extend in the scan direction. In the center of the map, the $\ell_x < 300$ mask applied in the inverse-variance filtering step largely removes these extended artifacts and they do not propagate into the lensing map. However towards the edges of the map, where the scan direction does not align with $\hat{\mathbf{x}}$ in ProjZEA, the artifacts are not suppressed and contribute to the lensing map. To minimize the impact of these artifacts, we mask the locations of point sources and galaxy clusters when plotting lensing maps and estimating lensing power spectra (bottom).

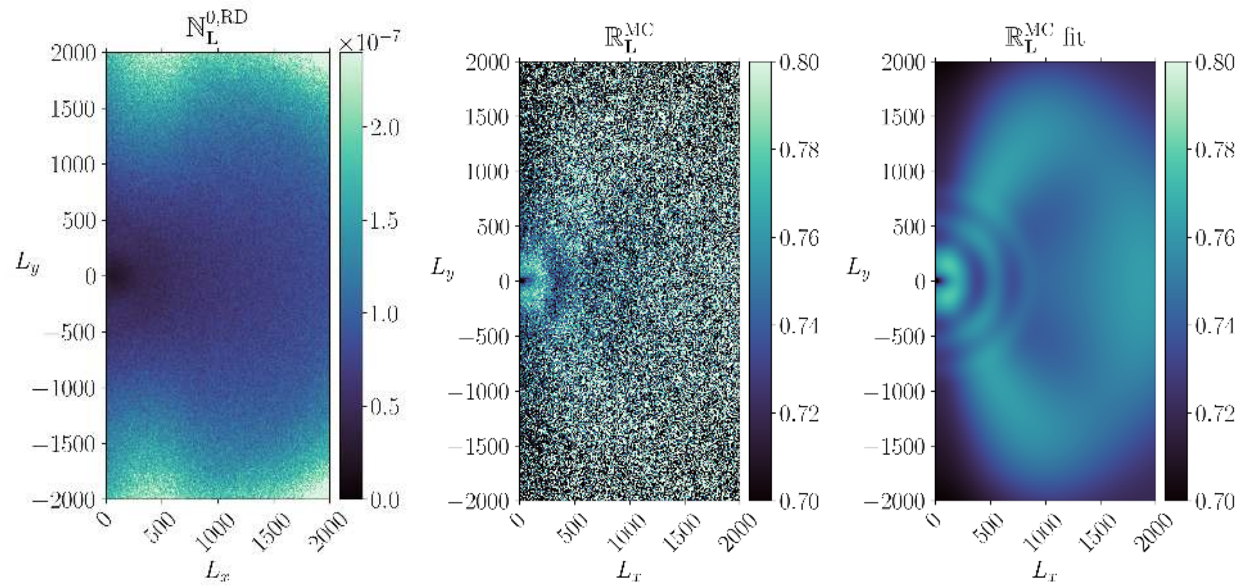

Figure 4.14: Lensing map noise and MC response correction for the SPT-3G 2019+2020 flat-sky lensing analysis in Fourier space. The realization-dependent noise $\mathbb{N}_{\boldsymbol{\ell}}^{0,\mathrm{RD}}$ (**left**) highly anisotropic, varying by a factor of two in the $\hat{\mathbf{x}}$ vs. $\hat{\mathbf{y}}$ directions; as a result we calculate all normalization and noise bias terms in two dimension. However the raw two-dimensional MC response $\mathbb{R}_{\boldsymbol{\ell}}^{\mathrm{MC}}$ (**center**) is too noisy to be used directly even when estimated from 500 simulations. Instead, we fit a multi-scale Gaussian process model to the raw MC response (**right**) to improve the stability of the lensing power spectrum estimation.

avoid confirmation bias when comparing with *Planck*, we only calculate the cross-correlation coefficient $\rho_L = C_L^{\hat{\kappa}\mathrm{CIB}}/\sqrt{C_L^{\hat{\kappa}\hat{\kappa}} C_L^{\mathrm{CIB\,CIB}}}$ which is insensitive to the overall amplitude of the signal. The SPT-3G, Planck, and CIB maps also have different analysis masks, making it even more difficult to make precise conclusions about the final lensing power spectrum from this comparison. Cross-correlation coefficients are shown in Figure 4.15. For the simulations we achieve a correlation coefficient of over 90% on large scales where the reconstruction noise is low, with over half of the signal coming from polarization. When cross-correlating the SPT-3G and *Planck* temperature reconstructions with the CIB, we find that the SPT-3G reconstruction has a ∼40% higher correlation coefficient compared to *Planck* on large scales thanks to SPT-3G's lower reconstruction noise. *Planck* lensing is temperature-dominated while SPT-3G lensing is polarization-dominated, so the improvement is even more impressive for the minimum-variance estimator.

We can also take this opportunity to check the efficacy of the bias-hardened estimator, as any CIB signal in the CMB temperature maps will bias the lensing reconstruction. The CIB bias is *anti-correlated* with the underlying matter distribution because point-like sources create negative features in a reconstructed convergence map (e.g. Figure 4.13), partially cancelling out the positive contributions from large scale structure traced by CIB galaxies. Thus the CIB can bias the temperature lensing power spectrum *low* by ∼3% below L of 1500, although both the amplitude and L-dependence of this bias depends on CMB map noise levels, point source masking thresholds, and the ℓ range used in the reconstruction (van Engelen et al., 2013). We find that bias-hardening increases the correlation coefficient by several percent on large scales, indicating some level of CIB bias in the standard temperature estimator. Both bias-hardening and the larger signal to noise in polarization will help to mitigate this bias.

4.6 Power Spectrum Estimation

Power spectrum estimation largely follows the methodology described in Chapter 2.4.4. We use 500 Gaussian lensed simulations to estimate the disconnected $\mathbb{N}^0_{\mathbf{L}}$ bias, and 150 pairs of simulations with different CMB realizations lensed by the same ϕ realization to estimate the connected $\mathbb{N}^1_{\mathbf{L}}$ bias. We estimate the more accurate realization-dependent disconnected bias $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}}$ by replacing a simulation with the data in one leg of the quadratic estimator. The semi-analytic version of this bias $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L}}$ is calculated with from inverse variance filtered simulation power spectra and is used to reduce the off-diagonal terms in the bandpower covariance matrix. This leaves us with noise bias-subtracted power spectra $\widehat{C}^{\kappa\kappa}_{\mathbf{L},\mathrm{data}}$ and $\widehat{C}^{\kappa\kappa}_{\mathbf{L},\mathrm{sim}}$.

When calculating a power spectrum, we azimuthally average the two-dimensional power spectrum $\langle \hat{\kappa}_{\mathbf{L}} \hat{\kappa}^*_{\mathbf{L}} \rangle$ and bin into one-dimensional bandpowers $C^{\kappa\kappa}_L$. Since the lensing noise is highly anisotropic, we use a two-dimensional signal to *variance* weighting so that for a given

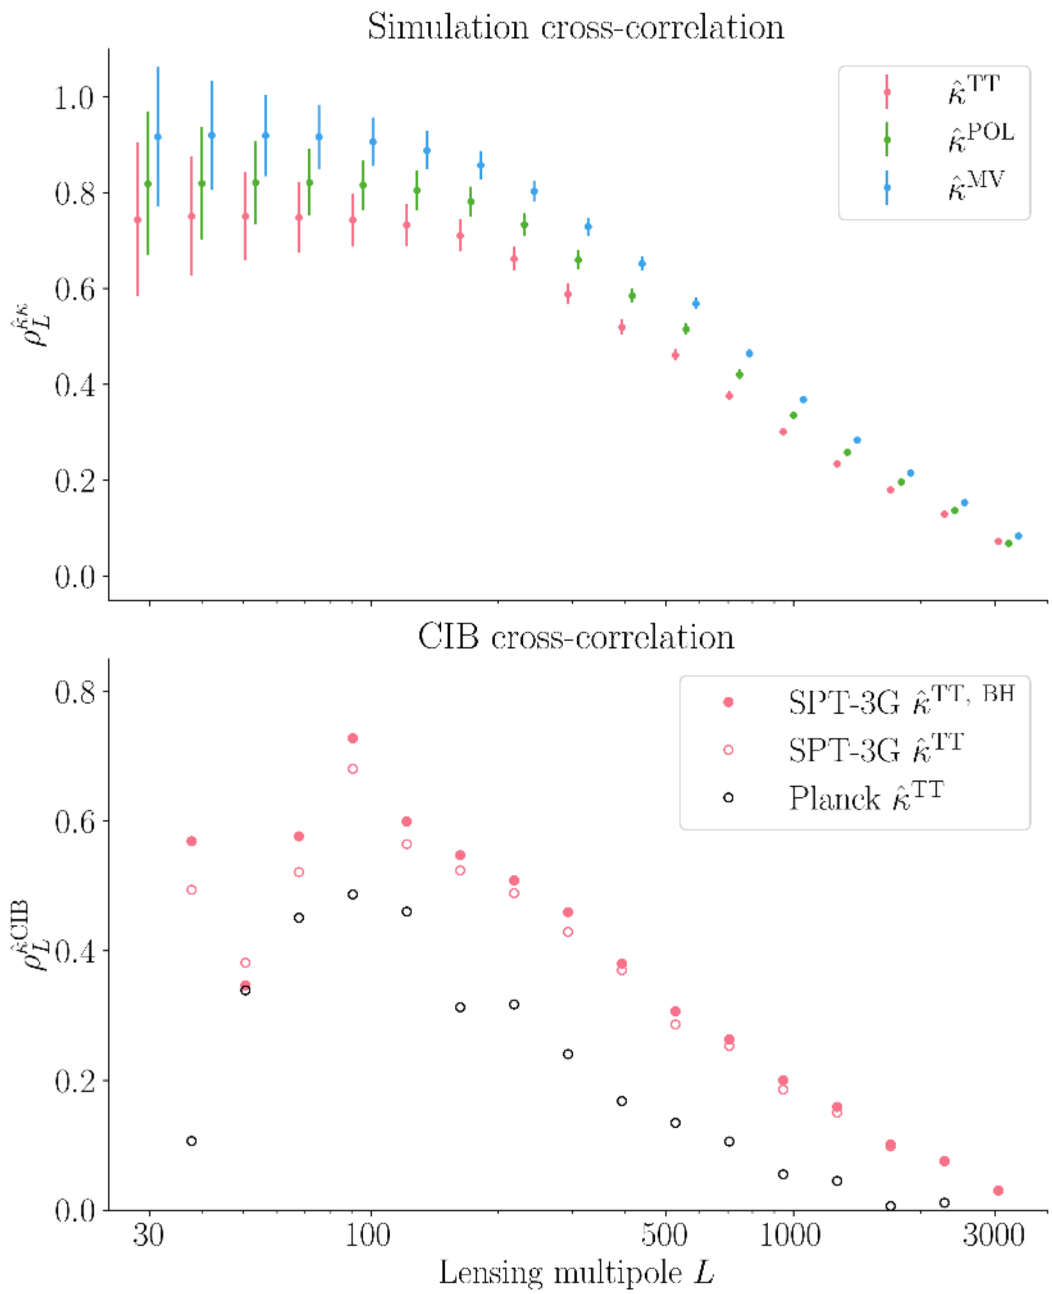

Figure 4.15: Cross-correlation coefficients for the SPT-3G 2019+2020 flat-sky lensing analysis as a function of lensing multipole L. **Left:** mean correlation coefficients between the reconstructed and input convergence maps from simulations. The correlation coefficient is reduced on large scales where the reconstruction noise is higher. Most of the signal to noise in the minimum variance estimator comes from polarization. **Right:** correlation coefficients for SPT-3G (pink) and Planck (black) temperature lensing maps. The SPT-3G lensing reconstruction noise is much lower than Planck's, leading to an improved correlation coefficient. We see that the bias-hardened estimator (closed circles) increases the correlation coefficient by several percent below L of 1500, consistent with a negative CIB bias in the standard temperature estimator.

bin b,

$$C_b^{\kappa\kappa} = \frac{\sum_{\mathbf{L}\in b} \mathbb{W}_{\mathbf{L}} \left\langle \hat{\kappa}_{\mathbf{L}} \hat{\kappa}^*_{\mathbf{L}} \right\rangle}{\sum_{\mathbf{L}\in b} \mathbb{W}_{\mathbf{L}}}, \tag{4.13}$$

where the weights are defined as

$$\mathbb{W}_{\mathbf{L}} = \frac{C_L^{\kappa\kappa,\mathrm{theory}}}{\mathrm{var}(\widehat{C}^{\kappa\kappa}_{\mathbf{L},\mathrm{sim}})}. \tag{4.14}$$

This weighting maximizes sensitivity to deviations from the fiducial theory power spectrum (Planck Collaboration et al., 2014c), and improves the signal-to-noise of our bandpowers by 25% at low L. In our fiducial setup we use 15 logarithmically spaced bins in the range $L \in [24, 1948]$, with two higher-L bins that are only used for diagnostic purposes.

We mask the locations of point sources and galaxy clusters in the lensing map as the point source filtering artifacts at these locations can bias the lensing power spectrum. However, masking all 2655 locations leads to a significant loss of sky area and a large amount of mode-coupling. While codes such as `PolSpice` (Challinor et al., 2011) and `NaMaster` (Alonso et al., 2019) can correct for mode-coupling from a complicated mask, neither code has the ability to perform two-dimensional weighting in flat-sky which is critical for our analysis as discussed above. Furthermore, the smallest point masks applied in mapmaking are only $3'$ across, barely larger than a single pixel in the downsampled ProjZEA map; an appropriate apodization of these tiny masks given our $2.25'$ pixel size would mask a significant fraction of the $1500\,\mathrm{deg}^2$ survey area. Thus we choose to only mask sources and clusters in the reconstruction that have a mask larger than $12'$ across. This effectively masks the 10% of sources and clusters with the highest signal-to-noise, and reduces the useable sky area by only 0.6%. We find that this approach reduces mode-coupling to manageable levels, and that masking more sources only reduces the excess power from point source artifacts by fractions of a percent.

One additional small bias emerges because we choose to use all 500 simulation to estimate a single mean field. The $\mathbb{N}^0_{\mathbf{L}}$ and $\mathbb{N}^1_{\mathbf{L}}$ present in the individual lensing reconstructions average down by a factor of $\sqrt{500}$ in the mean field map, but are not negligible compared to the lensing signal at small scales. This residual noise is injected into the mean-field subtracted map, and contributes a small $(\mathbb{N}^0_{\mathbf{L}} + \mathbb{N}^1_{\mathbf{L}})/500$ bias to the auto-spectrum. Previous SPT analyses have chosen to split the simulations in half and estimate the lensing power spectrum using two independent mean fields, effectively taking a cross-spectrum between maps with independent mean-field noise to avoid this bias (Omori et al., 2017; Wu et al., 2019; Pan et al., 2023). We find that the cross-spectrum mean-field approach does not change the level of bias in our pipeline, but does increase the power spectrum error bars at high L. We instead choose to calculate and subtract off the power spectrum of the averaged-down noise terms present in the mean-field, allowing us to use all simulations in the mean field estimate and reducing the *map-level* noise in the mean-field by $\sqrt{2}$.

We show the mean of 500 minimum variance simulation bandpowers $\left\langle \widehat{C}^{\kappa\kappa}_{L,\mathrm{sim}} \right\rangle$ in Figure 4.16. The realization-dependent data noise $\mathbb{N}^{0,\mathrm{RD}}_{\mathbf{L}}$ crosses the theory lensing spectrum at $L \sim 330$, indicating that our measurement will be signal-dominated well beyond the peak of the lensing power spectrum. This is a significant improvement over the highest signal-to-noise per mode measurements with *Planck* (Carron et al., 2022), SPT (Wu et al., 2019), and ACT (Madhavacheril et al., 2024; Qu et al., 2024) data, which are signal-dominated to $L \sim 75$, 150, and, 180 respectively. The power spectrum of the mean-field is larger than the signal below $L \sim 35$, so the the lowest L bin may unreliable if the mean-field is not estimated with sufficient accuracy. However the SPT-3G mean field is much smaller than in recent lensing measurements from SPT-3G and ACT, allowing us to push to lower L. We find that a $\sim$2% multiplicative bias remains in the noise bias-subtracted simulation bandpowers; we correct the data bandpowers by the inverse of this bias. We discuss potential sources of and solutions to this bias in Chapter 4.7.1.

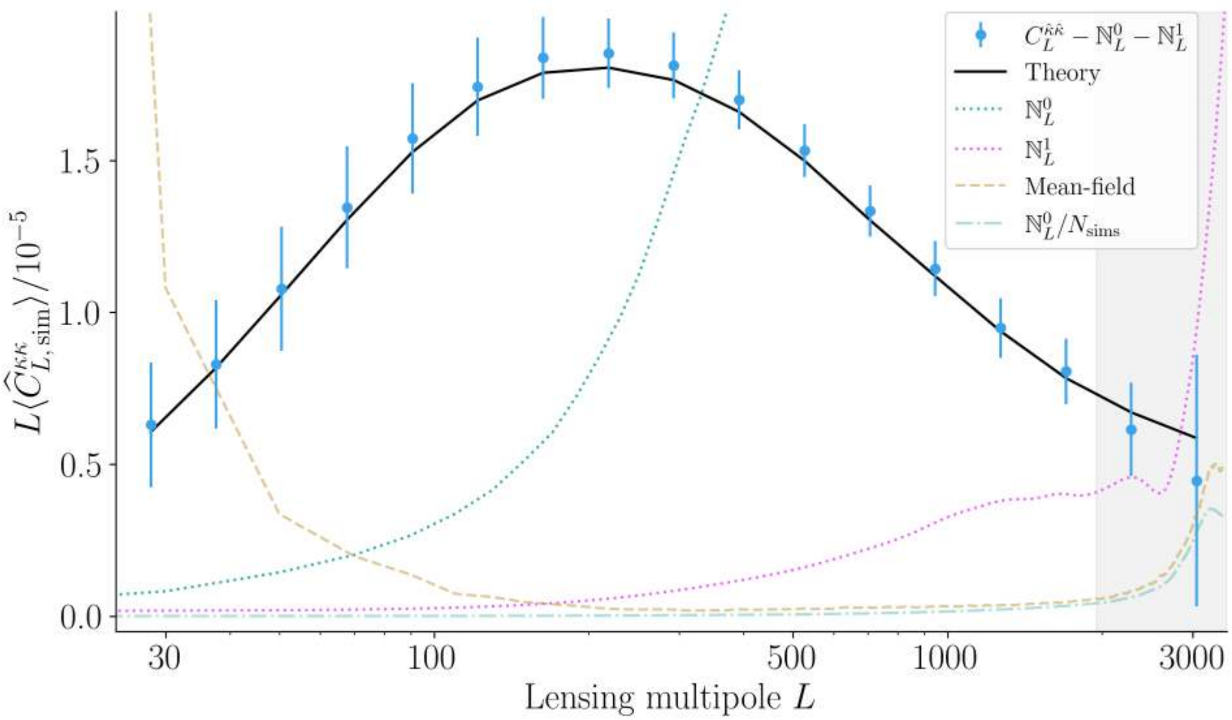

Figure 4.16: Mean of 500 minimum variance simulation bandpowers for the 2019+2020 SPT-3G lensing analysis as a function of lensing multipole L. The $\mathbb{N}_{\mathbf{L}}^{0,\rm RDSA}$- and $\mathbb{N}_{\mathbf{L}}^{1}$-subtracted bandpowers (blue) are compared to the theory lensing spectrum (black). The disconnected bias for data $\mathbb{N}_{\mathbf{L}}^{0,\rm RD}$ (dotted green) is very large at high L, and is nearly identical to its simulation analog $\mathbb{N}_{\mathbf{L}}^{0,\rm RDSA}$ indicating good agreement between the data and simulation CMB maps. The $\mathbb{N}_{\mathbf{L}}^{1}$ connected bias (dotted pink) is much smaller but is comparable to the signal at high L. The mean-field (dashed orange) has structure at both low and high L. The low L behavior is largely caused by the survey mask, while the high L noise is mostly from $\mathbb{N}_{\mathbf{L}}^{0}$ in the individual reconstructions used to form the mean-field. However the $\mathbb{N}_{\mathbf{L}}^{0}/N_{\rm sims}$ contribution (dash-dotted green) does not explain the entirety of the small-scale structure in the mean-field; the remaining mean-field power is likely related to point source masking artifacts. After subtracting the disconnected and connected bias terms and mean-field noise, we find a $\sim$2% multiplicative bias in the simulation bandpowers; we discuss this bias more in Chapter 4.7.1.

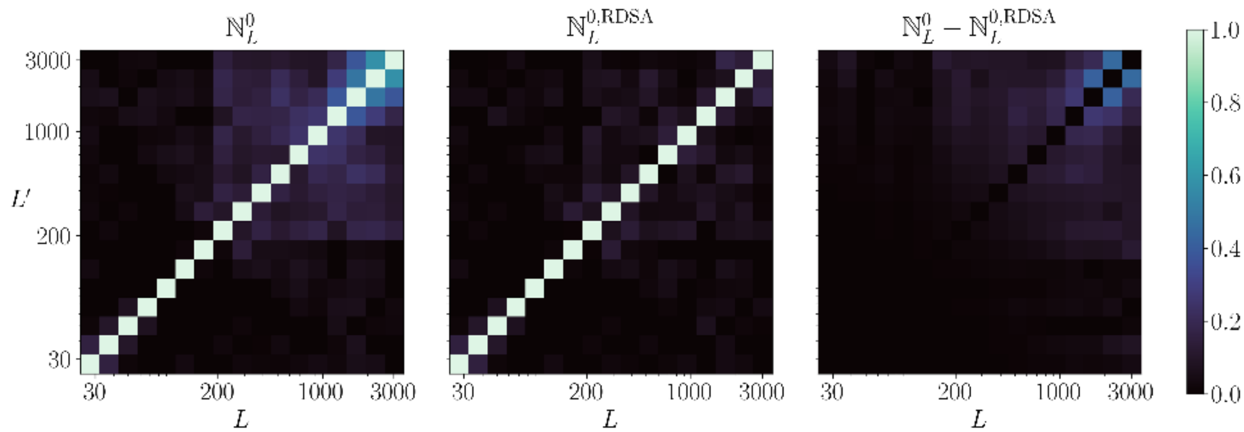

Figure 4.17: Demonstration of the improved off-diagonal $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L}}$ bandpower covariance for the minimum variance estimator. L bins above $L = 200$ are correlated at the ∼40% level in the standard $\mathbb{N}^0_{\mathbf{L}}$ noise covariance, but off-diagonal terms are reduced to ∼10-15% with $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L}}$. We note that mode-coupling may still affect the $\mathbb{R}^{\mathrm{MC}}_{\mathbf{L}}$ calculation as discussed in Chapter 4.7.1.

4.6.1 Bandpower Covariance

We estimate the covariance from the set of $N_{\mathrm{sims}} = 500$ bandpowers reconstructed from simulations and debiased using $N_L^{(0),\mathrm{RDSA}}$ to ensure that the simulation bandpowers have the same covariance as the $N_L^{(0),\mathrm{RD}}$-debiased data bandpowers as discussed in Chapter 2.4.4. The $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L}}$ correction greatly reduces correlation between bandpowers at high L (Figure 4.17) We correct the inverse of the covariance matrix by the Hartlap correction (Hartlap et al., 2007) to account for the finite number of simulations used to estimate the covariance:

$$\alpha_{\mathrm{cov}} = \frac{N_{sims} - N_{\mathrm{bin}} - 2}{N_{\mathrm{sims}} - 1}, \tag{4.15}$$

where $N_{\mathrm{bin}} = 15$ is the fiducial number of bandpowers.

4.6.2 Obfuscation of Results

When making precision measurements of cosmological parameters, it is vital to minimize the bias to these measurements from both systematics in the data and human decisions.

This can be difficult to balance, because unexpected results driven by new physics could be attributed by humans to systematics in the data, or vice versa. The value of a measurement is affected by numerous analysis choices, and researchers may (unintentionally) be susceptible to confirmation bias towards results consistent with theory or previous experiments. For example, in the instance of this lensing analysis, researchers might be tempted to make analysis choices that push the $C_L^{\kappa\kappa}$ bandpowers towards consistency with state of the art *Planck* and ACT bandpowers, or the ΛCDM cosmological model in general. On the other hand, researchers may initially find results consistent with their prior expectations, and may not be motivated to search further for systematics. To avoid compromising the quality of the final results because of these considerations, researchers may choose to hide the final results of an analysis (in this case lensing bandpowers and resulting cosmological parameters) until all analysis choices have been finalized. This practice, which I will refer to as "obfuscation"[1] has been well-established in the particle physics community for decades (Arisaka et al., 1993) and is becoming increasingly common in the observational cosmology community (Muir et al., 2020; Brieden et al., 2020).

Until we have finalized our analysis choices and passes the tests described in Section 4.7, we do not plot or read the debiased lensing bandpowers without first applying an obfuscating factor. When we do plot the obfuscated lensing bandpowers, we remove all y-axis markers and do not overlay theory expectations or measurements from other experiments. We do not apply any obfuscation at the level of CMB bandpowers, since we require good agreement between our CMB data and simulations (which are constructed from theory spectra) in order to properly estimate the mean field, MC corrections, and noise bias terms. We obfuscate our lensing bandpowers by applying a multiplicative factor $f(L_i) = A + B_i$, where $A \sim N(1, 0.05^2)$ and $B_i \sim N(0, 0.05^2)$ are drawn from normal distributions. The standard deviation of the overall scaling A was chosen to roughly match the size of potential bias from foregrounds in

[1]The term "blinding" is most commonly used to describe this practice, but some collaborations are trying to move away form this term on the account that it can be ableist. Even though blinding can be considered a positive term in this context, similar to the phrase "justice is blind," it still associates blindness with an reduced ability to fully know the world (Ades, 2020).

the minimum variance estimator, while the standard deviation of the L-dependent component was chosen to roughly match the uncertainty in each bin. This allows us to see catastrophic issues in the obfuscated data bandpowers without being able to tweak the bandpowers towards any expected value at the 5-10% level. To avoid accidentally looking at the raw data or the the obfuscating factor, we store the obfuscating coefficients in a non-human-readable file and write a function that applies the coefficients to the data vector each time it is read from disk.

4.7 Null Tests and Consistency Checks

We perform a suite of pipeline tests, data consistency tests, and characterization checks to ensure that our pipeline is sufficiently unbiased, that the data is robust to analysis choices, and to characterize systematic biases in our measurement of the lensing power spectrum. Our fiducial CMB multipole range used in the lensing reconstruction is

$$\ell_{x,\mathrm{min}} = 300, \qquad \ell_{y,\mathrm{min}} = 1000, \qquad \ell^{T}_{\mathrm{max}} = 3500, \qquad \ell^{P}_{\mathrm{max}} = 4000.$$

We use a lower ℓ_{max} in temperature to reduce our sensitivity to foreground biases. For our baseline lensing power spectrum we use lensing multipoles between 24 and 1948, divided into 15 non-overlapping bandpower bins with edges

$$L_{\mathrm{bin\ edges}} = [24, 32, 43, 58, 78, 104, 139, 187, 250, 336, 450, 603, 808, 1084, 1453, 1948].$$

We discard the lowest Ls < 24 as they are dominated by the mean field which may not be exactly the same for simulations and data. Even the first fiducial bin may be overly sensitive to the accuracy of the mean field estimate, and in the future we will assess the impact of the mean field on the lowest L bins with the data consistency tests described below. Conversely, we discard the highest Ls because they are dominated by the combined noise bias terms which have a very blue spectrum and cannot be estimated with infinite precision; these scales are also more sensitive to foregrounds, and add little information for cosmological parameter estimation.

The majority of the tests we perform amount to constructing a null spectrum, calculating a χ^2 from 0 for the null spectrum, and then calculating a probability to exceed (PTE) from this χ^2 given degrees of freedom equal to the number of bandpowers. For each test we run the pipeline on all 500 simulations with the given configuration of analysis choices (if different from fiducial), and construct a bandpower covariance matrix specific to the test from the resulting suite of bandpowers. We calculate the χ^2 from zero as

$$\chi_0^2 = \mathbf{d}^T \mathbf{C}^{-1} \mathbf{d}, \tag{4.16}$$

where $\mathbf{d}$ is the null spectrum, and $\mathbf{C}$ is the bandpower covariance matrix. Unless specified otherwise, we perform each test with three estimators: temperature-only (TT), polarization-only (POL), and minimum-variance (MV). Because the minimum-variance estimator is simply a combination of the temperature and polarization estimators, we do not consider the minimum variance estimator to be an independent test and take the number of independent tests to be $N_{\text{tests}} = 2$ when calculating PTE $> 0.05/N_{\text{tests}}$ thresholds unless stated otherwise.

4.7.1 Pipeline Tests

We perform pipeline two tests, one on lensed simulations and one on unlensed simulations, to ensure that we do not suffer from significant multiplicative or additive biases.

Lensed Simulation Pipeline Test. We first run our pipeline on all 500 simulations with our baseline configuration to ensure that our pipeline is unbiased. For this test we require that the *per-bin* mean of all 500 bandpowers constructed from simulations is within $0.2\sigma_i$ of the input $C_{L,i}^{\kappa\kappa,\text{theory}}$, where σ_i is the ith diagonal element of the bandpower covariance matrix. In other words, we require that the bias is less than 20% of the statistical uncertainty in each bin. We compute a multiplicative correction factor $C_{L,i}^{\kappa\kappa,\text{theory}}/\hat{C}_{L,i}^{\kappa\kappa}$ that we apply to all future

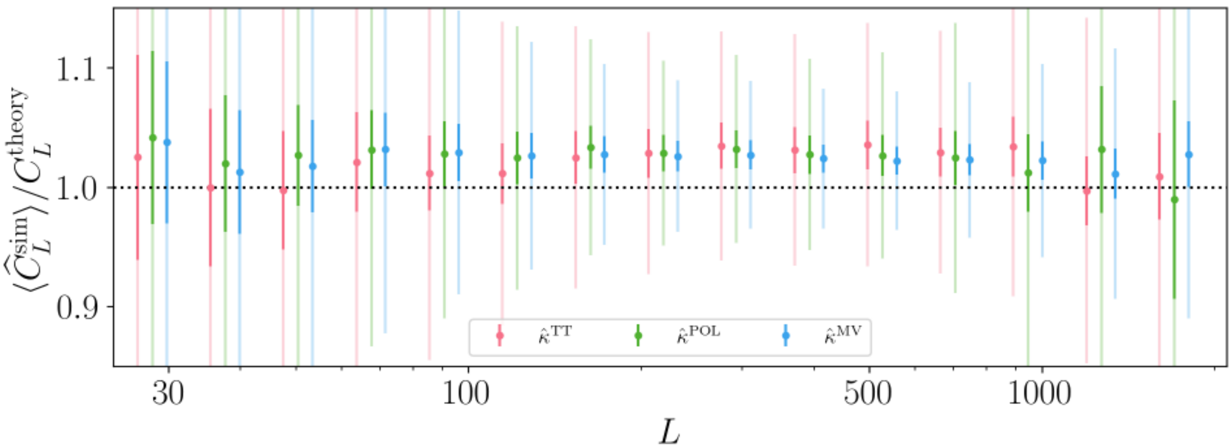

Figure 4.18: Lensed simulation pipeline test assessing the relative bias in the mean noise-subtracted bandpowers for the temperature, polarization, and minimum variance estimators as a function of lensing multipole L. The lighter error bars show the 1σ uncertainty for a single realization of the sky, while the darker error bars show the 0.2σ threshold we require to pass this test. We find a constant $\sim$2% bias across the full L range for all estimators, amounting to a $\sim 0.4\sigma$ bias in the highest signal to noise bins. Consequently we do not currently pass this test; we discuss potential solutions below.

bandpowers to correct for any $< 0.2\sigma_i$ bias that may be revealed by this test.

We show the results of the lensed simulation pipeline test in Figure 4.16. After subtracting the noise terms from the lensing power spectrum, this test reveals a flat $\sim$2% bias for all estimators corresponding to $\sim$40% of the statistical uncertainty in the highest signal to noise bins. As a result we do not pass this test and are working to understand and reduce the bias to an acceptable level. We suspect this bias may be related to the masking of point source location in the mapmaking timestream filtering, which could manifest in several ways.

The timestream point source masking couples with the CMB in the region around the point source mask in a way that creates correlations between different modes in the map, which is in turn picked up by the lensing estimator. While on average this introduces negative features in the reconstructed convergence map (e.g. Figure 4.13), the coupling depends on the specific realization of the CMB and thus the mean-field will not perfectly remove the artifacts from any given realization. This also means that the contribution to the lensing map auto-spectrum from the remaining artifacts will not be perfectly removed by the $\mathbb{N}^0_{\mathbf{L}}$ bias, which is is formed from lensing maps that have a different CMB realization in each leg of the quadratic estimator. As discussed in Chapter 2.4.4 we mask only 10% of source locations that have the largest timestream masking radius, which is a good proxy for the amplitude of the filtering artifacts. This threshold was chosen as a compromise between minimizing bias from filtering artifacts, mode-coupling from complicated mask structure, and loss of sky area. However we may revisit this decision and investigate more aggressive masking thresholds, or try Gaussian constrained realization inpainting to remove the artifacts from the lensing map without worsening mode-coupling from the mask. Running the lensing pipeline on mock observations generated without masking of point source locations would also be a useful test. If filtering artifacts are indeed responsible for the bias they will be present in the data as well, and we can add a new bias estimated from the difference between the lensing power spectrum from masked and unmasked simulations to our standard pipeline.

Mode-coupling (from masking, projection effects, or other sources) can itself introduce a bias in the lensing power spectrum. Investigations of mode-coupling from masking point source locations indicates that power is generally transferred from low to high L modes; this would tend to reduce to the MC response $\mathbb{R}^{\mathrm{MC}}_{\mathbf{L}}$ which is effectively a *cross-spectrum* transfer function for the lensing map. As discussed in 4.3, the *CMB* cross-spectrum transfer function $\mathbb{T}^{\mathrm{cross}}_{\mathbf{L}}$ tends to bias the CMB auto-spectrum high when it is deconvolved from the maps; similarly, one could conceive the $\mathbb{R}^{\mathrm{MC}}_{\mathbf{L}}$ biasing the lensing power spectrum high when it is deconvolved from the lensing maps. Even a 1% bias to $\mathbb{R}^{\mathrm{MC}}_{\mathbf{L}}$ could explain the 2% bias we see in the lensing power spectrum.

In summary, we are actively working to understand and reduce the 2% bias we see in the lensed simulation pipeline test. Such biases are quite common and have been seen in *Planck* ($\lesssim$10%; Planck Collaboration et al., 2020c; Carron et al., 2022), SPTpol ($\sim$5%; Wu et al., 2019), and ACT (mentioned but not quantified; Madhavacheril et al., 2024; Qu et al., 2024) lensing analyses. In order to proceed with the other pipeline and data tests while working on this bias, we correct the unlensed, Agora, and data bandpowers by the per-bin bias factor.

Motivated by the lack of bias in the unlensed simulation pipeline test described below, we treat the factor as multiplicative rather than additive.

Unlensed Simulation Pipeline Test. Next we confirm that our pipeline returns a null spectrum when evaluated on unlensed simulations. We perform lensing reconstruction and power spectrum estimation on 10 unlensed simulations that have the same power spectrum as the lensed simulations, and estimate the bandpower covariance directly from the unlensed simulations to avoid additional sample variance from lensing present in the lensed simulations. For this test we have two requirements. First, we require that the mean of the ten unlensed reconstruction power spectra has a PTE $> 0.05/N_{\mathrm{tests}}$. Second, we require that the power spectrum for *each* realization is consistent with the null hypothesis of no lensing by confirming that the PTEs of the ten simulations are uniformly distributed and do not cluster towards a given value. We take this approach rather than setting a PTE requirement for each individual simulation because given enough simulations, any given PTE threshold will eventually be exceeded by random chance.

With only ten unlensed simulations, the estimate of the bandpower covariance matrix will be singular and we cannot perform the full test as described above until more simulations are generated. In the meantime, we use the $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L}}$ covariance which has no dependence on the lensing signal. This is conservative because $\mathbb{N}^{0,\mathrm{RDSA}}_{\mathbf{L}}$ estimated from lensed simulations will receive contributions from lensing sample variance and will thus *overestimate* the unlensed covariance. We show the results of this test in Figure 4.19. None of the estimators shown any sign of bias, although the mean of the ten simulations reveals an oscillatory pattern around zero in the temperature estimator at low L where the estimator has the lowest signal to noise. This pattern may be due to a combination of large error bars and small number statistics; the PTEs of the individual simulations are not clustered around any particular value, although there is a slight preference for higher PTEs due to the overestimation of the unlensed covariance. This test currently passes but we will revisit with more simulations in the near future. We use the lack of bias in the unlensed simulations as indication that the lensed simulation bias is multiplicative rather than additive.

4.7.2 Characterization of Foreground Biases with Agora

Foregrounds can bias our temperature and minimum variance lensing reconstructions, and it is important to get a sense of the magnitude and scale dependence (both in the CMB maps and the lensing reconstruction) of these biases before finalizing our analysis choices and removing the obfuscation factor. We use the Agora Omori (2024) simulations to characterize the impact of foregrounds on our lensing power spectrum and the efficacy of various mitigation strategies such as restricting the temperature $\ell_{\max}$ or bias-hardening. We have ten $1500\,\mathrm{deg}^2$ patches of sky cut out from the single Agora sky realization, and estimate the lensing power spectrum for each patch. We do not have strict pass/fail criteria for this test, but rather use the results to inform our analysis choices and assure ourselves that our parameter constraints are not extremely sensitive to foregrounds. Of course the Agora simulations may not capture all the complexities of the real sky, and we will also use the data consistency tests described below to ensure that our data is robust to analysis choices and not biased by systematics. We also marginalize over foregrounds when estimating cosmological parameters, further reducing the sensitivity of our final constraints to foreground biases.

We show the results of the Agora foreground characterization in Figure 4.20. Interpretation of these results is complicated by the fact that some of the Agora patches used here are known to have a $\sim$25% low lensing amplitude due to sample variance at low L. As a result the mean lensing signal of the ten patches may also be low, which is borne out by the fact that even the foreground-insensitive polarization estimator is biased low at low L. However we can still assess the sensitivity to foregrounds by comparing the standard temperature and minimum variance estimators to the polarization and bias-hardened estimators. We find evidence for a negative foreground bias at low L in both the temperature and minimum variance estimators as well as a positive bias at high L in the temperature estimator. The PTEs of the standard

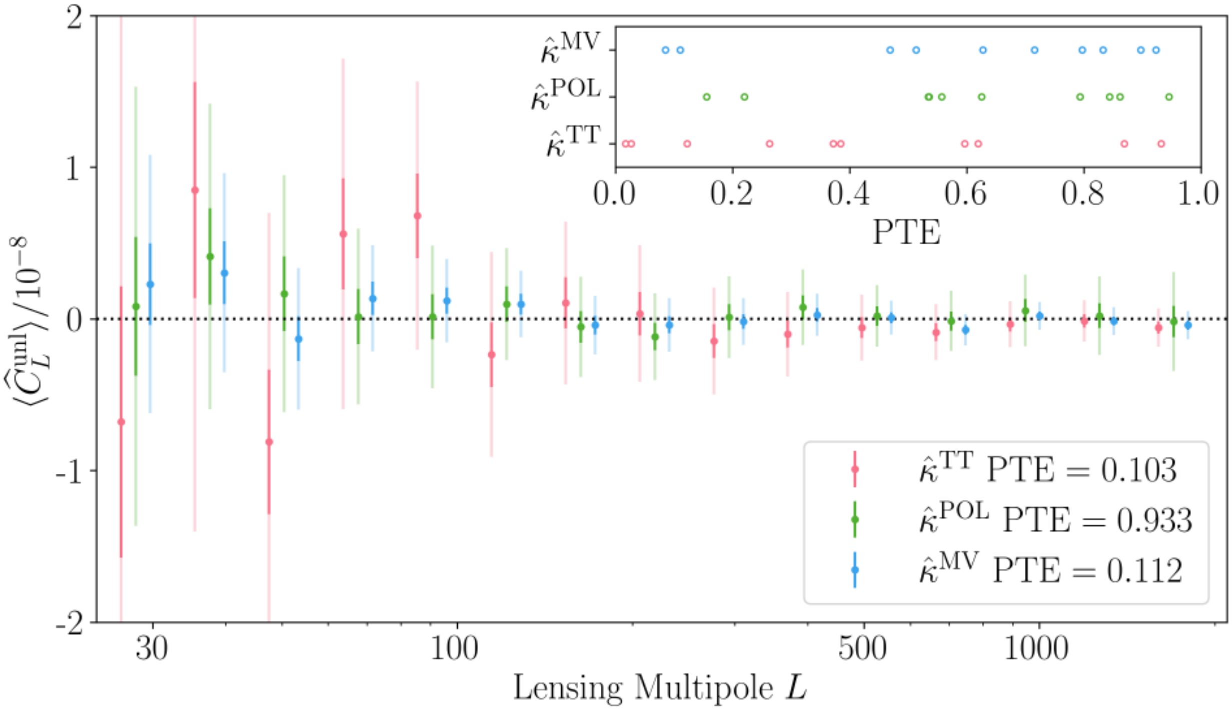

Figure 4.19: Unlensed simulation pipeline test assessing the level of absolute bias in the temperature, polarization, and minimum variance estimators as a function of lensing multipole L when our pipeline is run on maps with no lensing signal. The points in the main plot show the mean bandpowers across the ten available unlensed simulations. The error bars are calculated from the standard deviation of $\mathbb{N}_{\mathbf{L}}^{0,\mathrm{RDSA}}$ estimated from the Gaussian lensed simulations, with lighter error bars showing the 1σ uncertainty for a single realization and the darker error bars showing the uncertainty on the mean. The inset in the top right shows the distribution PTEs for the ten simulations. While there is no evidence of a systematic offset and the distribution of PTEs does not cluster around any particular value, the temperature estimator shows an oscillatory pattern around zero at low L. This may be due to a combination of large error bars and small number statistics, and we will revisit this test with more simulations in the near future.

temperature and minimum variance estimators are extremely low, while the bias-hardened versions raise the PTEs above our PTE $= 0.05/N_{\text{tests}}$ rule of thumb. This indicates that the bias-hardened estimators will be critical for reducing foreground biases in our final analysis, and reducing the temperature $\ell^T_{\max}$ may also be important. We will perform more tests with varying $\ell^T_{\max}$ to understand the scale-dependence of the foreground biases.

4.7.3 Data Consistency Tests

We perform three types of data consistency tests to ensure that our data is robust to analysis choices and that our lensing bandpowers are not biased by systematics. For each type of test and each configuration of analysis choices, we run the test on the data, construct a null spectrum, and then calculate a χ^2_0 and PTE using a covariance matrix estimated by running the same test on the lensed simulations. We are currently focusing on addressing the bias in the lensed simulation pipeline test, and will revisit the data consistency tests once we have a better understanding of the bias; nevertheless we briefly describe the tests we plan to perform below.

Curl Test. The general form of the lensing deflection field $\boldsymbol{\alpha}_{\hat{\mathbf{n}}}$ can be separated into a curl-free lensing potential $\phi_{\hat{\mathbf{n}}}$ and a divergence-free curl potential $\Omega_{\hat{\mathbf{n}}}$ such that in the flat-sky approximation, $\boldsymbol{\alpha}_{\hat{\mathbf{n}}} = \boldsymbol{\nabla}\phi_{\hat{\mathbf{n}}} + \boldsymbol{\nabla} \times \Omega_{\hat{\mathbf{n}}}$ (Hirata & Seljak, 2003a; Namikawa et al., 2012; Robertson & Lewis, 2023). The contribution to the deflection angle from the curl potential can also be estimated with the quadratic estimator by modifying the kernel $W^{XY}_{\boldsymbol{\ell},\boldsymbol{\ell}-\mathbf{L}}$ in Equation (2.16); conceptually the curl kernel depends on $\boldsymbol{\ell} \times \mathbf{L}$ rather than $\boldsymbol{\ell} \cdot \mathbf{L}$. While $\Omega_{\hat{\mathbf{n}}}$ is zero in the Born approximation, interactions between multiple lens planes can generate a curl signal. The resulting curl auto-spectrum is expected to be undetectable with SPT-3G, but the post-Born curl signal may be detectable with future experiments via cross-correlations with a curl template formed from multi-redshift tracers of the matter density (Pratten & Lewis, 2016; Robertson & Lewis, 2023). On the other hand, non-Gaussian systematics and

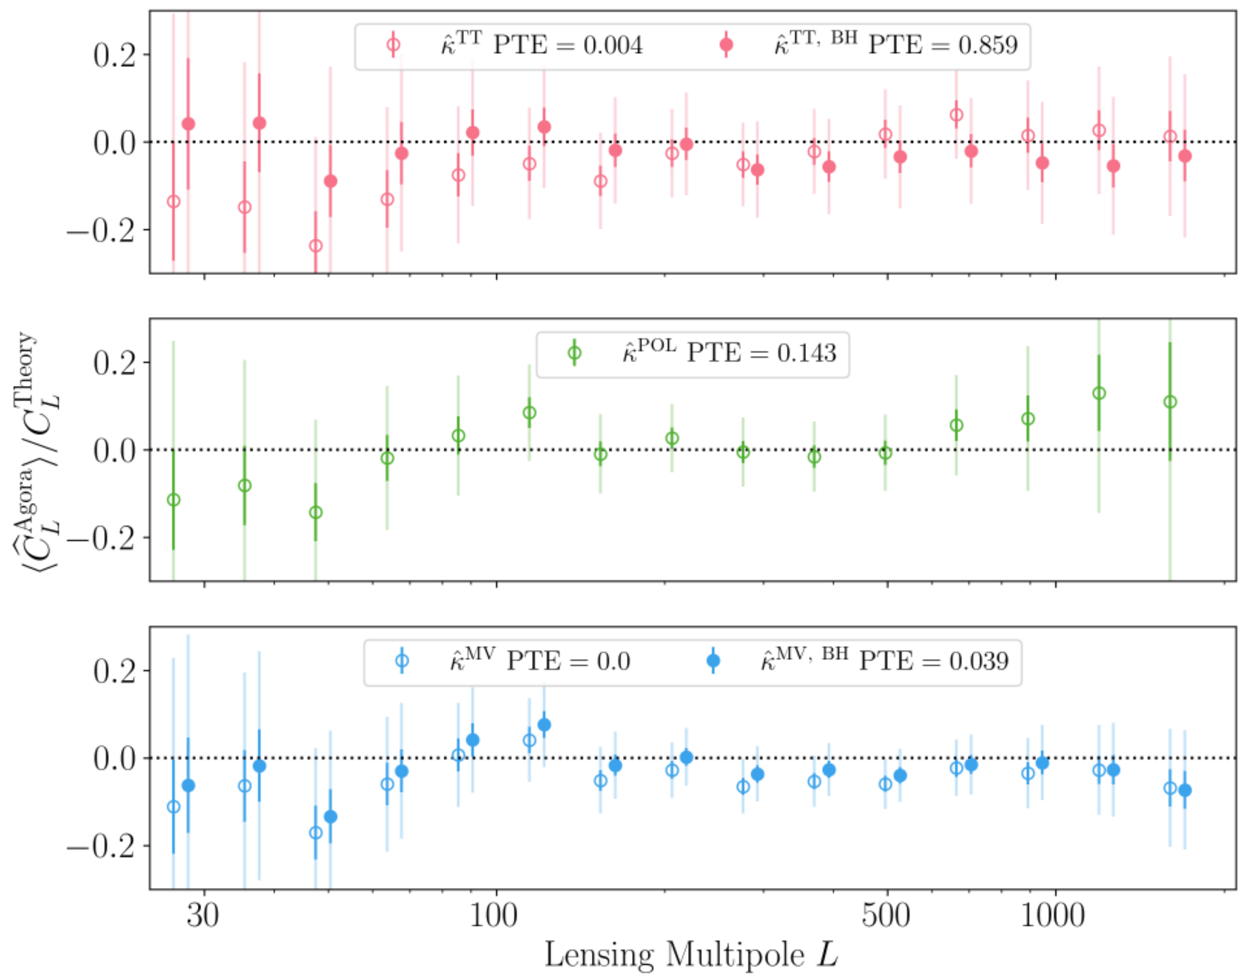

Figure 4.20: Characterization of foreground biases in the SPT-3G 2019+2020 flat-sky lensing analysis with the Agora simulation. We show the mean temperature, polarization, and minimum variance lensing power spectra for ten 1500 deg^2 patches of the Agora sky realization as a function of lensing multipole L. For the temperature and minimum variance combinations, we also show the bias-hardened estimators which reduce foreground biases. The error bars are calculated from the standard deviation of the Gaussian lensed simulations, with lighter error bars showing the 1σ uncertainty for a single realization and the darker error bars showing the uncertainty on the mean. Interpretation of these results is complicated by the fact that some of the Agora patches used here are known to have a low lensing amplitude as discussed in the text. Nevertheless we note that the standard temperature and minimum variance estimators have extremely low PTEs, while the bias-hardened versions have PTEs above our PTE$= 0.05/N_{\rm tests}$ rule of thumb.

foregrounds could also generate a curl-like signal, making the curl test a useful check for systematics in our data and a further validation of our pipeline (Cooray et al., 2005). For the curl test we use the estimators described in Namikawa et al. (2012). We do not apply an MC correction as the curl signal in our simulations is zero, but perform bias subtraction in the same manner as for the convergence estimators as the N_L^0 and N_L^1 biases also contribute to the quadratic estimator curl trispectrum (van Engelen et al., 2012, Appendix A). We only consider the minimum-variance curl estimator, so the PTE threshold is simply 0.05.

Consistency Across Estimators. We test the consistency between independent groupings of estimators to ensure there are not problems with specific estimators and that we are not sensitive to foreground biases. We perform three different tests:

$$C_L^{\mathrm{TT}} - C_L^{\mathrm{POL}}, \qquad C_L^{\mathrm{MV-TT}} - C_L^{\mathrm{TT}}, \qquad C_L^{\mathrm{TTTEEE}} - C_L^{\mathrm{TBEB}}.$$

The first two tests in particular are sensitive to foreground biases; we also look at χ_0-statistics in addition to χ_0^2 as the former is more sensitive to one-sided biases that one might expect from foregrounds. Since all three tests amount to different splits of the minimum-variance estimator, we take the number of independent tests to be $N_{\mathrm{tests}} = 1$ and require PTE > 0.05 for each test.

Stability Given CMB Multipole Range Used. We vary the CMB multipole range used in the lensing reconstruction to ensure that our results are robust to the choice of $\ell_{x\mathrm{min}}$, $\ell_{y\mathrm{min}}$, ℓ_{max}^T, ℓ_{max}^P and that subsets of the data are consistent with each other. Varying ℓ_{max}^T tests for biases from foregrounds such as CIB and tSZ, as these biases tend to increase with ℓ_{max}. We vary ℓ_{max}^P to test for the effects of high-ℓ systematics in our polarization data that have been hinted at by comparisons between mock and data EE power spectra.[2] On the other hand, varying $\ell_{x\mathrm{min}}$ and $\ell_{y\mathrm{min}}$ tests for issues related to inverse-variance filtering near the $\ell = 300$

[2]This is not public information.

timestream lowpass. We vary $\ell_{x\min}$, $\ell_{y\min}$, $\ell^T_{\max}$, $\ell^P_{\max}$ separately, calculating a null spectrum between each configuration and the fiducial configuration for both data and simulations to properly estimate the covariance of each null spectrum. The multipole variations are listed below, with the fiducial values in bold:

$$\begin{aligned}
\ell_{x\min} &\in \{\mathbf{300}, 350, 400\}, \\
\ell_{y\min} &\in \{700, \mathbf{1000}, 1300\}, \\
\ell^T_{\max} &\in \{3000, \mathbf{3500}, 4000\}, \\
\ell^P_{\max} &\in \{3000, 3500, \mathbf{4000}, 4500\}.
\end{aligned}$$

This amounts to 7 independent tests for TT and 8 for POL, for a total of 15 independent tests when calculating PTE $> 0.05/N_{\text{tests}}$ thresholds.

The tests describe above will allow us to finalize our analysis choices and build confidence in the robustness of our measurement before removing the obfuscation factor from the data.

4.8 Obfuscated Data Bandpowers

Although we have not yet finalized our analysis choices and removed the obfuscation factor, our obfuscation policy allows us to plot the data bandpowers with the factor applied. Motivated by the sensitivity to foreground bias and improvement from bias-hardening demonstrated by the Agora checks, we show both the standard and bias-hardened minimum variance bandpowers in Figure 4.21. The obfuscated bandpowers give a $\sqrt{\chi^2_0}$ from zero of $\sim$39, suggesting a measurement of the lensing power spectrum with a total signal to noise of nearly 40σ. This is competitive with the highest signal to noise measurements from Planck (40-42σ, Planck Collaboration et al., 2020c; Carron et al., 2022) and ACT (43σ, Madhavacheril et al., 2024; Qu et al., 2024) over only 3.5% of the sky.

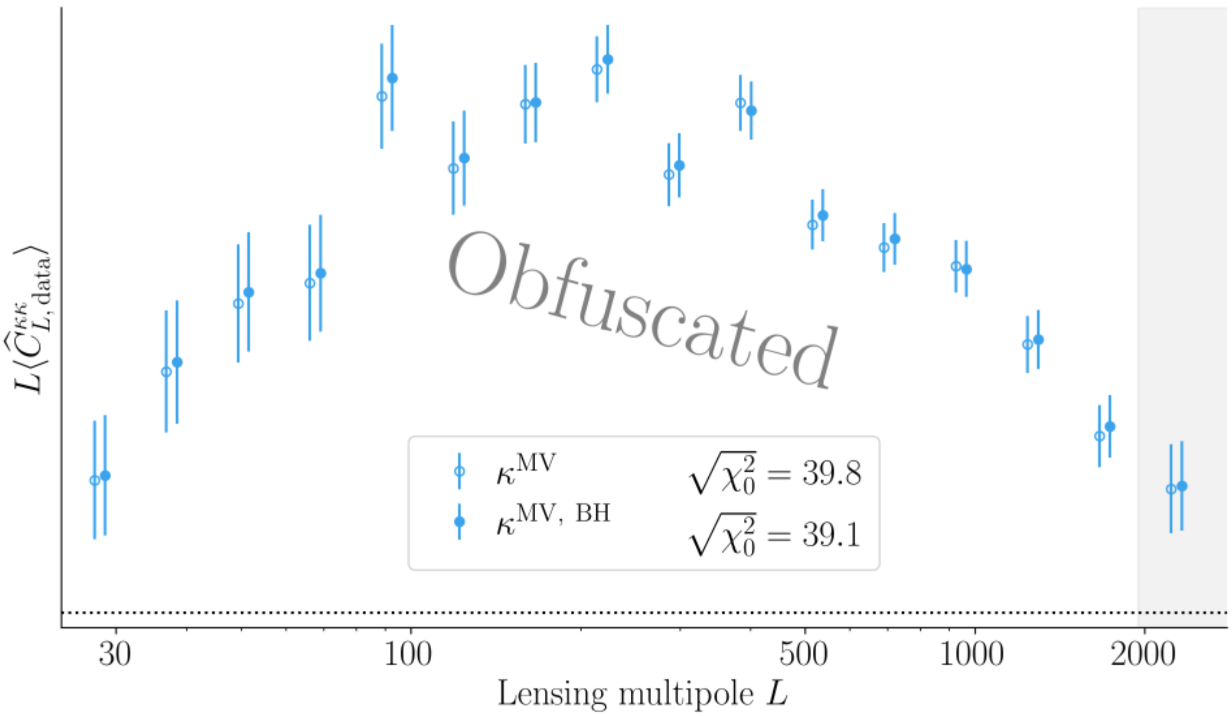

Figure 4.21: Obfuscated data bandpowers for the SPT-3G 2019+2020 flat-sky lensing analysis as a function of harmonic multipole L. We show the standard and bias hardened minimum variance bandpowers with error bars estimated from the Gaussian lensed simulations. Consistent with our obfuscation policy we do not show the y-axis or compare to theory expectations or other experiments. We expect the application of the obfuscating factor to increase the scatter by roughly $\sqrt{2}$. We find a $\sqrt{\chi_0^2}$ from zero of $\sim$39 for the bias hardened estimator, competitive with recent *Planck* and ACT measurements. Consistent with the characterization of foreground biases with the Agora simulations, we see that bias-hardening increases the lensing power over a broad range of scales, although this is counteracted by the larger bias-hardened error bars in the $\sqrt{\chi_0^2}$ calculation.

4.9 Discussion and Future Work

The three 2019+2020 lensing analyses are currently all nearing completion, with the quadratic estimator pipelines working on the passing of consistency and null tests described in Chapter 4.7, and the Bayesian pipeline focusing on the interpretation and validation of their results. The collaboration is aiming for a September 1st internal deadline for the completion of tests and the removal of the obfuscation factor for the two quadratic estimator analyses as well as a primary CMB power spectrum analysis using the same data. Many exciting avenues will

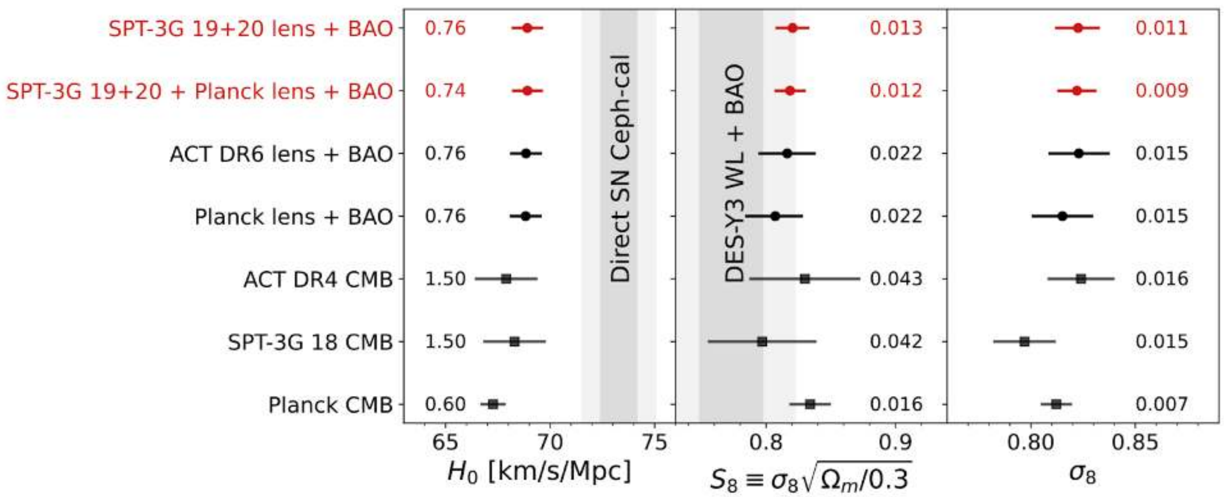

Figure 4.22: Forecasted parameter constraints from the 2019+2020 SPT-3G lensing analyses compared to other early-universe probes combined with DESI BAO constraints (DESI Collaboration et al., 2024). Figure kindly provided by Federico Bianchini.

open up after the completion of the quadratic estimator pipeline tests and the subsequent removal of the obfuscation factor described in Section 4.6.2, including rigorous comparison of the results of each pipeline, parameter estimation, B-mode delensing, and cross-correlation with external tracers of large scale structure. Pipeline comparisons and parameter estimation will be included in the initial suite of publications from the 2019+2020 data, while delensing and cross-correlations are longer-term standalone projects that will be pursued in the coming years. The SPT collaboration is still in the process of finalizing the division of papers for the 2019+2020 primary CMB and lensing analyses, but the current plan is for me to lead a paper comparing the three lensing pipelines in order to establish confidence in the results and to clarify the differences between pipelines and the impact of different analysis choices of each pipeline on the lensing bandpowers and perhaps even cosmological constraints. This paper will also be an excellent opportunity to build intuition between components of the quadratic estimator and the MUSE framework, for example how the connected $\mathbb{N}^1_L$ quadratic estimator bias term is accounted for in a Bayesian context.

Experts in cosmological parameter estimation within the collaboration have already begun to explore constraints with simulated data vectors using the quadratic estimator covariance, including marginalization over instrumental systematics and foregrounds, and cosmology-dependence of the quadratic estimator MC and $\mathbb{N}_L^{(1)}$ corrections. Cosmological parameter constraints from the 2019+2020 lensing analyses will be state of the art, and when combined with baryonic acoustic oscillations measurements we will be able to add new perspectives to the H_0 and S_8 tensions between early- and late-time measurements of the rate of expansion of the universe and the amplitude of matter fluctuations (Figure 4.22). SPT-3G 2-19+2020 constraints will also have slightly different degeneracies in parameter space compared to recent lensing measurements from Planck (Figure 4.23), providing an important check on the robustness of the ΛCDM model and allowing explorations of model extensions.

The SPT-3G 2019+2020 lensing maps will also be crucial in the joint effort between BICEP/*Keck* to reduce B-mode lensing sample variance in the search for primordial gravitational waves. A joint set of simulations for the two experiments is currently being generated, and with the SPT-3G lensing maps in hand we will soon be able to begin the delensing process. Work on constructing a combined low- and mid-ℓ set of maps to be delensed with the SPT-3G lensing maps is already underway.

Finally, the SPT-3G 2019+2020 lensing maps can be used for cross correlations with other tracers of large scale structure. The incredible sensitivity of SPT-3G's polarization lensing maps will allow cross-correlations with reduced sensitivity to foregrounds. Cross-correlations with Dark Energy Survey (DES) galaxy clustering and weak lensing maps (5×2pt, 6×2 Omori et al., 2018; Chang et al., 2023; Abbott et al., 2023) can begun immediately, and the initial data release from the Euclid satellite will partially overlap with the SPT-3G main field. Combining measurements from Euclid mission with CMB data can shrink error bars on cosmological parameters by as much as a factor of 10 compared to Euclid-only results (Euclid Collaboration et al., 2022). On a slightly longer timescale, the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) will overlap with the SPT-3G main field and will provide a wealth of data for cross-correlations with SPT-3G lensing maps.

I have begun exploring a cross-correlation of the SPT-3G lensing maps with a catalog of

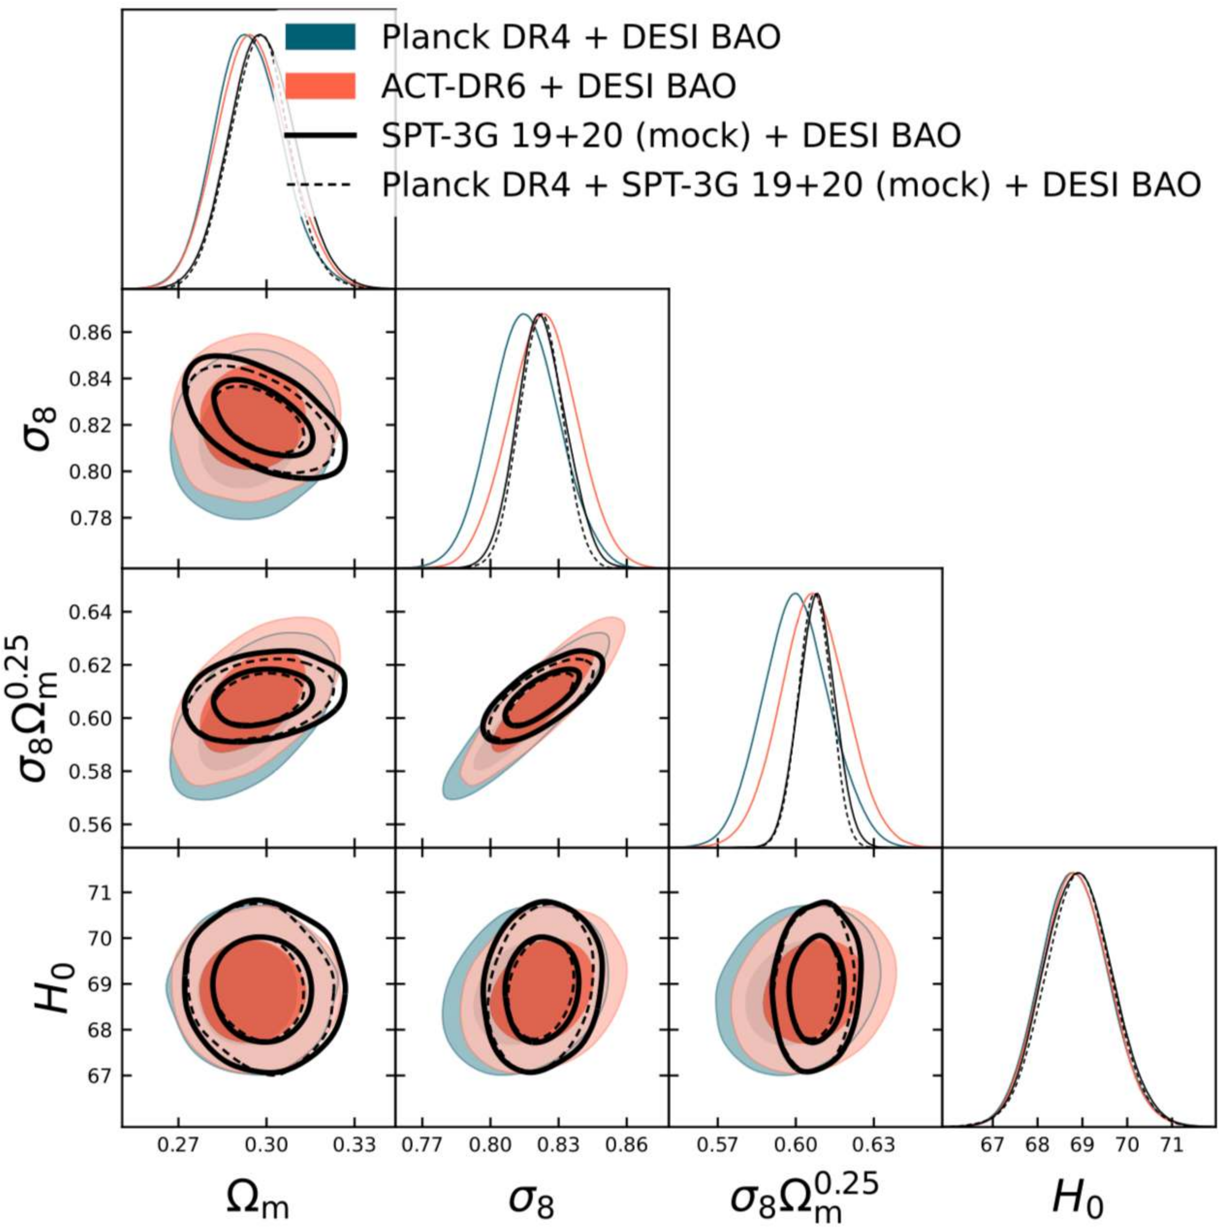

Figure 4.23: Forecasted parameter constraints and degeneracies for the 2019+2020 SPT-3G lensing analyses compared to lensing measurements from other experiments combined with DESI BAO constraints (DESI Collaboration et al., 2024). Figure kindly provided by Federico Bianchini.

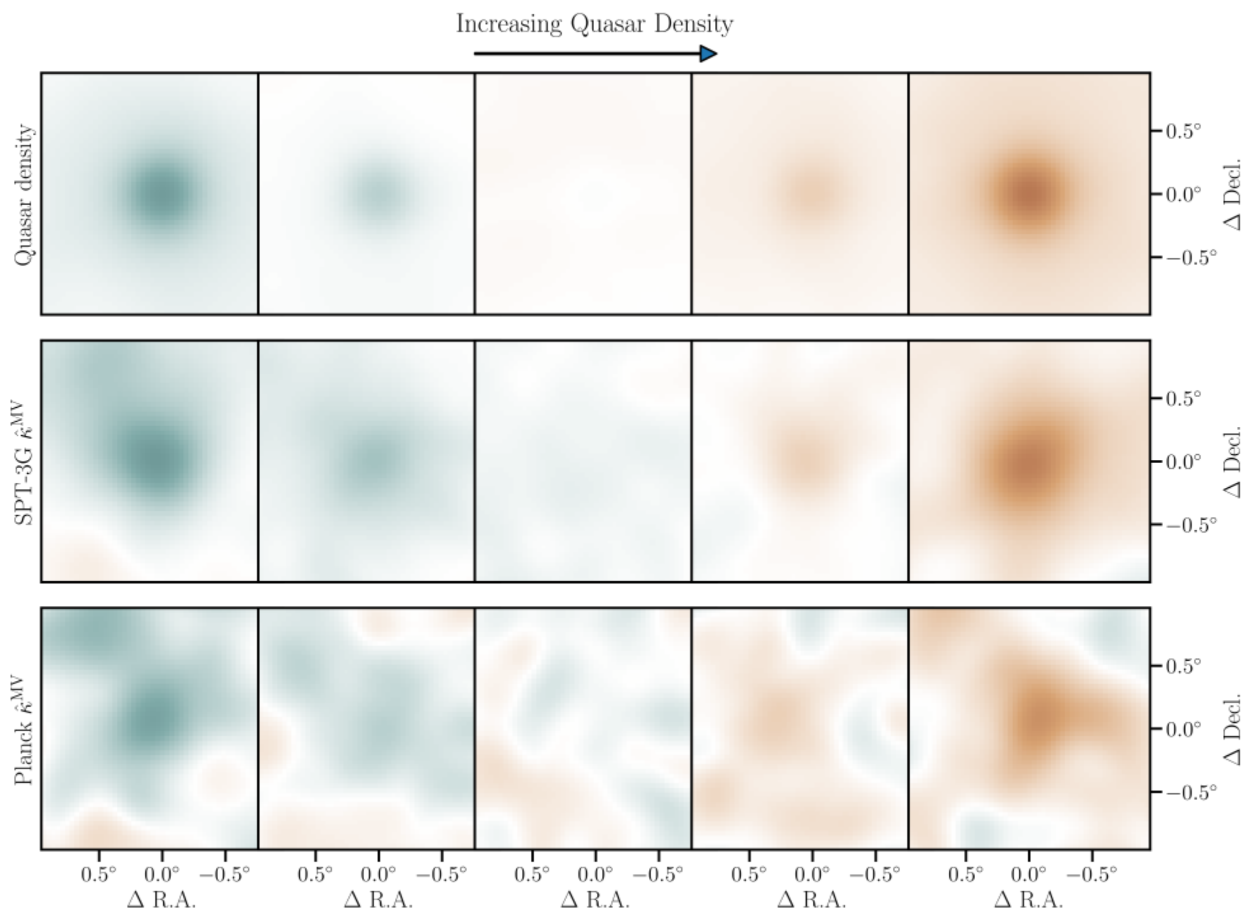

Figure 4.24: SPT-3G 2019+2020 and *Planck* lensing maps stacked on regions of varying quasar density using the Yang & Shen (2022) DES quasar catalog. The lensing signal is strongly correlated with the quasar density, and the SPT-3G stacks trace the quasar density with much greater signal to noise than the *Planck* stacks thanks to SPT-3G's low lensing noise.

DES quasars from UIUC's very own Yang & Shen (2022). Taking inspiration from Geach et al. (2013), I have stacked the SPT-3G lensing lensing maps on regions of different quasar density and demonstrated a strong correlation between the lensing signal and quasar density Figure 4.24. I plan to continue this investigation by performing cross-correlations in Fourier space and by adding multiple tracers at different redshift to facilitate tomographic studies. Tomographic cross-correlations can be used to project out the low-redshift component of the lensing signal, reducing late-time nonlinear and baryonic effects and allowing searches for non-Gaussianity in the lensing signal at high redshifts (Zhang et al., 2023). Lensing maps cleaned of low-redshift components in this way can also be used to reduce biases on neutrino mass measurements (McCarthy et al., 2021). I am also interested in cross-correlating SPT and ACT lensing maps with each other to assess the consistency of the two datasets and explore the effects of foreground biases.

The SPT-3G 2019+2020 lensing results will be a major contribution to the field of observational cosmology, and will open up many exciting avenues for future research including parameter estimation, delensing, and cross-correlations.

Chapter 5

Summary and Conclusions

This thesis has focused on gravitational lensing of the cosmic microwave background (CMB) using multiple datasets from the South Pole Telescope (SPT).

Chapter 1 begins with an introduction to the CMB and the concordance ΛCDM cosmological model, as well as some of the tensions between early- and late-time cosmological measurements that motivate extensions to the standard model. The South Pole Telescope and its three generations of SPT-SZ, SPTpol, and SPT-3G receivers are introduced, and SPT is put in broader context with respect to other CMB experiments. Many of SPT's science cases are highlighted, including measurements of the primary CMB power spectrum, CMB secondaries, and studies of millimeter point sources and galaxy clusters. The chapter concludes with a discussion of the flat-sky approximation and map projections used in this thesis, particularly implications for polarization maps.

Chapter 2 opens with a discussion of weak lensing in general and CMB lensing in particular. CMB lensing reconstruction is motivated with a summary of science applications before the formalism of the Hu & Okamoto (2002) quadratic estimator reconstruction technique is introduced. Ingredients of the quadratic estimator and lensing power spectrum are discussed in detail including the normalization and debiasing of both the lensing map and power spectrum.

Chapter 3 summarizes the results from (Millea et al., 2021), which was the first demonstration of joint optimal lensing reconstruction and cosmological parameter estimation with real CMB data. The analysis used $100\,\mathrm{deg}^2$ of exquisitely deep SPTpol data with noise levels that approach the lensing B-mode noise floor where the quadratic estimator becomes suboptimal, motivating the use of an optimal Bayesian method that leverages all orders of the lensing signal. The Bayesian method is compared to the standard quadratic estimator to build confidence in the new method and demonstrate the improvements that can be attained with optimal lensing reconstruction. The Bayesian and quadratic estimator results were consistent with each other as well as with previous SPT analyses that used a subset of the data, but the Bayesian method achieved 17% tighter constraints on the amplitude of the lensing power spectrum A_ϕ compared to the quadratic estimator pipeline. I led the quadratic estimator pipeline and was second author on the paper; this was my first scientific publication in graduate school. Leading the quadratic estimator on a relatively small and simple dataset was an excellent preparation for the SPT-3G lensing analysis that I am currently leading.

Chapter 4 describes the ongoing SPT-3G 2019+2020 lensing analysis which is the primary focus of my thesis. The CMB lensing maps produced by this analysis will be the deepest ever made, with per-mode reconstruction noise levels lower than recent *Planck* and ACT lensing measurements by a factor of 3-4. The chapter opens with a motivation of the analysis before describing the mapmaking, simulation, and map processing steps that are performed before lensing reconstruction. Lensing reconstruction is described in detail, including the inverse variance filtering, normalization, and debiasing of the lensing maps. After the lensing maps are presented, lensing power spectrum estimation is discussed and the use of an obfuscation factor to avoid confirmation bias is introduced. Tests and consistency checks that must be performed before finalizing the analysis choices are then discussed, and since these tests are ongoing, only obfuscated lensing bandpowers are presented. The chapter closes with a discussion of future work and the scientific impact of the results.

The SPT-3G 2019+2020 lensing analysis is expected to measure the lensing power spectrum with a signal-to-noise of nearly 40σ, with a 2-3% constraint on the amplitude of the lensing power spectrum A_ϕ. This is competitive with the best constraints from *Planck* and ACT, and will place powerful independent constraints on ΛCDM parameters and extensions to the standard model. In particular, we will be able to measure the Hubble constant H_0 with a precision similar to that of the best *Planck* and ACT lensing measurements, and place the tightest CMB lensing constraints to date on structure growth parameter S_8. As a result the SPT-3G 2019+2020 parameter constraints will shed new light on tensions between estimates of H_0 and S_8 from early-universe vs. late-time measurements. The SPT-3G lensing maps will also be used to delens CMB polarization maps in the search for primordial gravitational waves in collaboration with BICEP/*Keck*. These results will be critical for placing the tightest possible constraints on the tensor-to-scalar ratio r, especially given the uncertainty surrounding the future of the South Pole component of the proposed CMB-S4 experiment.

The SPT-3G lensing maps will also be excellent for cross-correlations, especially given the high signal to noise of the polarization reconstruction that is relatively free from foreground contamination. Cross-correlations with optical surveys including DES, Euclid, and eventually LSST will allow us to probe the growth of structure over cosmic time, break degeneracies between cosmological parameters, and calibrate weak lensing shear measurements. We can also use cross-correlations to produce clean high-redshift maps of cosmic structure useful for constraining parameters that govern structure formation in the universe, including the sum of neutrino masses.

I am particularly excited about cross-correlation prospects with Euclid as I will be beginning a postdoctoral position this fall at CEA Paris-Saclay where I will be working with Euclid weak lensing data. I will be joining the TOSCA[1] project to explore new weak lensing statistics and the synergies between optical and radio radio experiments, specifically focusing on weak lensing measurements from Euclid and the Square Kilometer Array Observatory. This position comes with a generous allocation of time for non-TOSCA projects, and I

[1] https://tosca.cosmostat.org

look forward to working with members of the Euclid and SPT collaborations to perform cross-correlations with the latest SPT-3G lensing maps. I am very lucky to have had the privilege of learning so much about CMB lensing at the University of Illinois and in the South Pole Telescope collaboration, and I am excited to extend this experience to the fields of optical and radio weak lensing in the coming years.

Bibliography

Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, arXiv e-prints, arXiv:1610.02743. https://arxiv.org/abs/1610.02743

Abbott, T. M. C., Aguena, M., Alarcon, A., et al. 2022, Phys. Rev. D, 105, 023520, doi: 10.1103/PhysRevD.105.023520

—. 2023, Phys. Rev. D, 107, 023531, doi: 10.1103/PhysRevD.107.023531

Adachi, S., Aguilar Faúndez, M. A. O., Akiba, Y., et al. 2020, Phys. Rev. Lett., 124, 131301, doi: 10.1103/PhysRevLett.124.131301

Ade, P. A. R., Ahmed, Z., Amiri, M., et al. 2021, Phys. Rev. Lett., 127, 151301, doi: 10.1103/PhysRevLett.127.151301

Ade, P. A. R., Amiri, M., Benton, S. J., et al. 2022a, ApJ, 927, 174, doi: 10.3847/1538-4357/ac20df

Ade, P. A. R., Ahmed, Z., Amiri, M., et al. 2022b, ApJ, 927, 77, doi: 10.3847/1538-4357/ac4886

Ades, R. 2020, An End to Blind Review?, American Philosophical Association Blog. https://blog.apaonline.org/2020/02/20/an-end-to-blind-review/?amp#top

Aghanim, N., Majumdar, S., & Silk, J. 2008, Reports on Progress in Physics, 71, 066902, doi: 10.1088/0034-4885/71/6/066902

Albrecht, A., & Steinhardt, P. J. 1982, Phys. Rev. Lett., 48, 1220, doi: 10.1103/PhysRevLett.48.1220

Alonso, D., Sanchez, J., Slosar, A., & LSST Dark Energy Science Collaboration. 2019, MNRAS, 484, 4127, doi: 10.1093/mnras/stz093

Anderes, E. 2013, Phys. Rev. D, 88, 083517, doi: 10.1103/PhysRevD.88.083517

Anderes, E., Wandelt, B. D., & Lavaux, G. 2015, ApJ, 808, 152, doi: 10.1088/0004-637X/808/2/152

Ange, J., & Meyers, J. 2023, J. Cosmology Astropart. Phys., 2023, 045, doi: 10.1088/1475-7516/2023/10/045

Arisaka, K., Auerbach, L. B., Axelrod, S., et al. 1993, Phys. Rev. Lett., 70, 1049, doi: 10.1103/PhysRevLett.70.1049

Austermann, J. E., Aird, K. A., Beall, J. A., et al. 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8452, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VI, ed. W. S. Holland & J. Zmuidzinas, 84521E, doi: 10.1117/12.927286

Bartelmann, M., & Schneider, P. 2001, Phys. Rep., 340, 291, doi: 10.1016/S0370-1573(00)00082-X

Baumann, D., Jackson, M. G., Adshead, P., et al. 2009, in American Institute of Physics Conference Series, Vol. 1141, CMB Polarization Workshop: Theory and Foregrounds: CMBPol Mission Concept Study, ed. S. Dodelson, D. Baumann, A. Cooray, J. Dunkley, A. Fraisse, M. G. Jackson, A. Kogut, L. Krauss, M. Zaldarriaga, & K. Smith (AIP), 10–120, doi: 10.1063/1.3160885

Baxter, E., Clampitt, J., Giannantonio, T., et al. 2016, MNRAS, 461, 4099, doi: 10.1093/mnras/stw1584

Baxter, E. J., Omori, Y., Chang, C., et al. 2019, Phys. Rev. D, 99, 023508, doi: 10.1103/PhysRevD.99.023508

Beck, D., Fabbian, G., & Errard, J. 2018, Phys. Rev. D, 98, 043512, doi: 10.1103/PhysRevD.98.043512

Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, ApJS, 148, 1, doi: 10.1086/377253

Benoit-Lévy, A., Déchelette, T., Benabed, K., et al. 2013, A&A, 555, A37, doi: 10.1051/0004-6361/201321048

Benson, B. A., Ade, P. A. R., Ahmed, Z., et al. 2014, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9153, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy VII, 91531P, doi: 10.1117/12.2057305

Bernardeau, F., Colombi, S., Gaztañaga, E., & Scoccimarro, R. 2002, Phys. Rep., 367, 1, doi: 10.1016/S0370-1573(02)00135-7

BICEP2 and Keck Array Collaborations, Ade, P. A. R., Ahmed, Z., et al. 2015, ApJ, 811, 126, doi: 10.1088/0004-637X/811/2/126

BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2014, ApJ, 792, 62, doi: 10.1088/0004-637X/792/1/62

BICEP2 Collaboration, Keck Array Collaboration, Ade, P. A. R., et al. 2018, Phys. Rev. Lett., 121, 221301, doi: 10.1103/PhysRevLett.121.221301

BICEP/Keck Collaboration, SPTpol Collaboration, Ade, P. A. R., et al. 2021, Phys. Rev. D, 103, 022004, doi: 10.1103/PhysRevD.103.022004

Birkinshaw, M. 1999, Phys. Rep., 310, 97, doi: 10.1016/S0370-1573(98)00080-5

Birkinshaw, M., & Gull, S. F. 1983, Nature, 302, 315, doi: 10.1038/302315a0

Bleem, L. E., van Engelen, A., Holder, G. P., et al. 2012, ApJ, 753, L9, doi: 10.1088/2041-8205/753/1/L9

Bleem, L. E., Bocquet, S., Stalder, B., et al. 2020, ApJS, 247, 25, doi: 10.3847/1538-4365/ab6993

Böhm, V., Schmittfull, M., & Sherwin, B. D. 2016, Phys. Rev. D, 94, 043519, doi: 10.1103/PhysRevD.94.043519

Böhm, V., Sherwin, B. D., Liu, J., et al. 2018, Phys. Rev. D, 98, 123510, doi: 10.1103/PhysRevD.98.123510

Brieden, S., Gil-Marín, H., Verde, L., & Bernal, J. L. 2020, J. Cosmology Astropart. Phys., 2020, 052, doi: 10.1088/1475-7516/2020/09/052

Bunn, E. F., Zaldarriaga, M., Tegmark, M., & de Oliveira-Costa, A. 2003, Phys. Rev. D, 67, 023501, doi: 10.1103/PhysRevD.67.023501

Bussmann, R. S., Holzapfel, W. L., & Kuo, C. L. 2005, ApJ, 622, 1343, doi: 10.1086/427935

Cardoso, J.-F., Martin, M., Delabrouille, J., Betoule, M., & Patanchon, G. 2008, arXiv e-prints, arXiv:0803.1814, doi: 10.48550/arXiv.0803.1814

Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, PASP, 123, 568, doi: 10.1086/659879

Carron, J., & Lewis, A. 2017, Phys. Rev. D, 96, 063510, doi: 10.1103/PhysRevD.96.063510

Carron, J., Mirmelstein, M., & Lewis, A. 2022, J. Cosmology Astropart. Phys., 2022, 039, doi: 10.1088/1475-7516/2022/09/039

Cavaliere, A., & Fusco-Femiano, R. 1976, A&A, 49, 137

—. 1978, A&A, 70, 677

Challinor, A. 2013, in Astrophysics from Antarctica, ed. M. G. Burton, X. Cui, & N. F. H. Tothill, Vol. 288, 42–52, doi: 10.1017/S1743921312016663

Challinor, A., Chon, G., Colombi, S., et al. 2011, PolSpice: Spatially Inhomogeneous Correlation Estimator for Temperature and Polarisation, Astrophysics Source Code Library, record ascl:1109.005

Chang, C., Omori, Y., Baxter, E. J., et al. 2023, Phys. Rev. D, 107, 023530, doi: 10.1103/PhysRevD.107.023530

Chiang, H. C., Ade, P. A. R., Barkats, D., et al. 2010, ApJ, 711, 1123, doi: 10.1088/0004-637X/711/2/1123

Chichura, P. M., Foster, A., Patel, C., et al. 2022, ApJ, 936, 173, doi: 10.3847/1538-4357/ac89ec

Coble, K., Ade, P. A. R., Bock, J. J., et al. 2003, arXiv e-prints, astro, doi: 10.48550/arXiv.astro-ph/0301599

Cooray, A., Kamionkowski, M., & Caldwell, R. R. 2005, Phys. Rev. D, 71, 123527, doi: 10.1103/PhysRevD.71.123527

Crites, A. T., Henning, J. W., Ade, P. A. R., et al. 2015, ApJ, 805, 36, doi: 10.1088/0004-637X/805/1/36

Cyburt, R. H., Fields, B. D., Olive, K. A., & Yeh, T.-H. 2016, Reviews of Modern Physics, 88, 015004, doi: 10.1103/RevModPhys.88.015004

Dark Energy Survey and Kilo-Degree Survey Collaboration, Abbott, T. M. C., Aguena, M., et al. 2023, The Open Journal of Astrophysics, 6, 36, doi: 10.21105/astro.2305.17173

Darwish, O., Madhavacheril, M. S., Sherwin, B. D., et al. 2021, MNRAS, 500, 2250, doi: 10.1093/mnras/staa3438 de Bernardis, P., Ade, P. A. R., Bock, J. J., et al. 2000, Nature, 404, 955, doi: 10.1038/35010035

DESI Collaboration, Adame, A. G., Aguilar, J., et al. 2024, arXiv e-prints, arXiv:2404.03002, doi: 10.48550/arXiv.2404.03002

Dutcher, D., Balkenhol, L., Ade, P. A. R., et al. 2021, Phys. Rev. D, 104, 022003, doi: 10.1103/PhysRevD.104.022003

Dyson, F. W., Eddington, A. S., & Davidson, C. 1920, Philosophical Transactions of the Royal Society of London Series A, 220, 291, doi: 10.1098/rsta.1920.0009

Einstein, A. 1916, Annalen der Physik, 354, 769, doi: 10.1002/andp.19163540702

Euclid Collaboration, Ilić, S., Aghanim, N., et al. 2022, A&A, 657, A91, doi: 10.1051/0004-6361/202141556

Everett, W. B., Zhang, L., Crawford, T. M., et al. 2020, ApJ, 900, 55, doi: 10.3847/1538-4357/ab9df7

Ferraro, S., & Hill, J. C. 2018, Phys. Rev. D, 97, 023512, doi: 10.1103/PhysRevD.97.023512

Fields, B. D., Olive, K. A., Yeh, T.-H., & Young, C. 2020, J. Cosmology Astropart. Phys., 2020, 010, doi: 10.1088/1475-7516/2020/03/010

Fixsen, D. J. 2009, ApJ, 707, 916, doi: 10.1088/0004-637X/707/2/916

Fowler, J. W., Niemack, M. D., Dicker, S. R., et al. 2007, Appl. Opt., 46, 3444, doi: 10.1364/AO.46.003444

Fowler, J. W., Acquaviva, V., Ade, P. A. R., et al. 2010, ApJ, 722, 1148, doi: 10.1088/0004-637X/722/2/1148

Freedman, W. L., Madore, B. F., Gibson, B. K., et al. 2001, ApJ, 553, 47, doi: 10.1086/320638

Frigo, M., & Johnson, S. G. 2005, Proceedings of the IEEE, 93, 216

Geach, J. E., Hickox, R. C., Bleem, L. E., et al. 2013, ApJ, 776, L41, doi: 10.1088/2041-8205/776/2/l41

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, ApJ, 622, 759, doi: 10.1086/427976

Green, D., Meyers, J., & van Engelen, A. 2017, J. Cosmology Astropart. Phys., 2017, 005, doi: 10.1088/1475-7516/2017/12/005

Guns, S., Foster, A., Daley, C., et al. 2021, ApJ, 916, 98, doi: 10.3847/1538-4357/ac06a3

Guth, A. H. 1981, Phys. Rev. D, 23, 347, doi: 10.1103/PhysRevD.23.347

Hanson, D., Challinor, A., Efstathiou, G., & Bielewicz, P. 2011, Phys. Rev. D, 83, 043005, doi: 10.1103/PhysRevD.83.043005

Hanson, D., Hoover, S., Crites, A., et al. 2013, Phys. Rev. Lett., 111, 141301, doi: 10.1103/PhysRevLett.111.141301

Hartlap, J., Simon, P., & Schneider, P. 2007, A&A, 464, 399, doi: 10.1051/0004-6361:20066170

Hawking, S. W. 1982, Physics Letters B, 115, 295, doi: 10.1016/0370-2693(82)90373-2

Henderson, S. W., Allison, R., Austermann, J., et al. 2016, Journal of Low Temperature Physics, 184, 772, doi: 10.1007/s10909-016-1575-z

Henning, J. W., Sayre, J. T., Reichardt, C. L., et al. 2017, ApJ, 852, 97, doi: 10.3847/1538-4357/aa9ff4

Hirata, C. M., & Seljak, U. 2003a, Phys. Rev. D, 68, 083002, doi: 10.1103/PhysRevD.68.083002

—. 2003b, Phys. Rev. D, 67, 043001, doi: 10.1103/PhysRevD.67.043001

Hoffman, Y., & Ribak, E. 1991, ApJ, 380, L5, doi: 10.1086/186160

Holder, G. P., Viero, M. P., Zahn, O., et al. 2013, ApJ, 771, L16, doi: 10.1088/2041-8205/771/1/L16

Hood, J. C., I., Simpson, A., McDaniel, A., et al. 2023, ApJ, 945, L23, doi: 10.3847/2041-8213/acbf45

Hu, W. 2000, Phys. Rev. D, 62, 043007, doi: 10.1103/PhysRevD.62.043007

Hu, W., & Dodelson, S. 2002, ARA&A, 40, 171, doi: 10.1146/annurev.astro.40.060401.093926

Hu, W., Fukugita, M., Zaldarriaga, M., & Tegmark, M. 2001, ApJ, 549, 669, doi: 10.1086/319449

Hu, W., Hedman, M. M., & Zaldarriaga, M. 2003, Phys. Rev. D, 67, 043004, doi: 10.1103/PhysRevD.67.043004

Hu, W., & Okamoto, T. 2002, ApJ, 574, 566, doi: 10.1086/341110

Hu, W., & White, M. 1997, New A, 2, 323, doi: 10.1016/S1384-1076(97)00022-5

Johnson, S. G. 2011, Notes on FFT-based differentiation, MIT Applied Mathematics. https://api.semanticscholar.org/CorpusID:1087476

Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, Phys. Rev. Lett., 78, 2058, doi: 10.1103/PhysRevLett.78.2058

Kamionkowski, M., & Kovetz, E. D. 2016, ARA&A, 54, 227, doi: 10.1146/annurev-astro-081915-023433

Kaplinghat, M., Knox, L., & Song, Y.-S. 2003, Phys. Rev. Lett., 91, 241301, doi: 10.1103/PhysRevLett.91.241301

Kashlinsky, A., Arendt, R. G., Atrio-Barandela, F., et al. 2018, Reviews of Modern Physics, 90, 025006, doi: 10.1103/RevModPhys.90.025006

Keating, B., Shimon, M., & Yadav, A. 2012, ApJ, 762, L23, doi: 10.1088/2041-8205/762/2/l23

Keisler, R., Reichardt, C. L., Aird, K. A., et al. 2011, ApJ, 743, 28, doi: 10.1088/0004-637X/743/1/28

Kesden, M., Cooray, A., & Kamionkowski, M. 2003, Phys. Rev. D, 67, 123507, doi: 10.1103/PhysRevD.67.123507

Kilbinger, M. 2015, Reports on Progress in Physics, 78, 086901, doi: 10.1088/0034-4885/78/8/086901

Kim, J., Sailer, N., Madhavacheril, M. S., et al. 2024, arXiv e-prints, arXiv:2407.04606. https://arxiv.org/abs/2407.04606

Knox, L. 1995, Phys. Rev. D, 52, 4307, doi: 10.1103/PhysRevD.52.4307

Kolmogorov, A. N. 1941, Moscow University Bulletin, 2, 1

Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, Nature, 420, 772, doi: 10.1038/nature01269

Lambert, J. H. 1772, Anmerkungen und Zusätze zur Entwerfung der Land- und Himmelscharten, ed. A. Wangerin (University of Michigan Library Digital Collections). https://name.umdl.umich.edu/ABR2581.0001.001

Lenz, D., Doré, O., & Lagache, G. 2019, ApJ, 883, 75, doi: 10.3847/1538-4357/ab3c2b

Lesgourgues, J., & Pastor, S. 2006, Phys. Rep., 429, 307, doi: 10.1016/j.physrep.2006.04.001

Lewis, A. 2005, Phys. Rev. D, 71, 083008, doi: 10.1103/PhysRevD.71.083008

—. 2011, LensPix: Fast MPI full sky transforms for HEALPix, Astrophysics Source Code Library, record ascl:1102.025

Lewis, A., & Challinor, A. 2006, Phys. Rep., 429, 1, doi: 10.1016/j.physrep.2006.03.002

Lewis, A., Challinor, A., & Hanson, D. 2011, J. Cosmology Astropart. Phys., 2011, 018, doi: 10.1088/1475-7516/2011/03/018

Linde, A. D. 1982, Physics Letters B, 108, 389, doi: 10.1016/0370-2693(82)91219-9

Lueker, M., Reichardt, C. L., Schaffer, K. K., et al. 2010, ApJ, 719, 1045, doi: 10.1088/0004-637X/719/2/1045

Lyth, D. H. D. H., & Riotto, A. A. 1999, Phys. Rep., 314, 1, doi: 10.1016/S0370-1573(98)00128-8

Madhavacheril, M. S., & Hill, J. C. 2018, Phys. Rev. D, 98, 023534, doi: 10.1103/PhysRevD.98.023534

Madhavacheril, M. S., Qu, F. J., Sherwin, B. D., et al. 2024, ApJ, 962, 113, doi: 10.3847/1538-4357/acff5f

Mancini, A. S., Piras, D., Alsing, J., Joachimi, B., & Hobson, M. P. 2021, MNRAS, 511, 1771, doi: 10.1093/mnras/stac064

Maniyar, A. S., Ali-Haïmoud, Y., Carron, J., Lewis, A., & Madhavacheril, M. S. 2021, Phys. Rev. D, 103, 083524, doi: 10.1103/PhysRevD.103.083524

Manzotti, A., Story, K. T., Wu, W. L. K., et al. 2017, ApJ, 846, 45, doi: 10.3847/1538-4357/aa82bb

McCarthy, F., Foreman, S., & van Engelen, A. 2021, Phys. Rev. D, 103, 103538, doi: 10.1103/PhysRevD.103.103538

Meneghetti, M. 2021, Introduction to Gravitational Lensing: With Python Examples (Springer International Publishing), doi: 10.1007/978-3-030-73582-1

Millea, M., Anderes, E., & Wandelt, B. D. 2019, Phys. Rev. D, 100, 023509, doi: 10.1103/PhysRevD.100.023509

—. 2020, Phys. Rev. D, 102, 123542, doi: 10.1103/PhysRevD.102.123542

Millea, M., & Seljak, U. 2021, Phys. Rev. D, 105, 103531, doi: 10.1103/physrevd.105.103531

Millea, M., Daley, C. M., Chou, T. L., et al. 2021, ApJ, 922, 259, doi: 10.3847/1538-4357/ac02bb

Muir, J., Bernstein, G. M., Huterer, D., et al. 2020, MNRAS, 494, 4454, doi: 10.1093/mnras/staa965

Namikawa, T., Hanson, D., & Takahashi, R. 2013, MNRAS, 431, 609, doi: 10.1093/mnras/stt195

Namikawa, T., Yamauchi, D., & Taruya, A. 2012, J. Cosmology Astropart. Phys., 2012, 007, doi: 10.1088/1475-7516/2012/01/007

Nguyên, H. N., Sehgal, N., & Madhavacheril, M. S. 2019, Phys. Rev. D, 99, 023502, doi: 10.1103/PhysRevD.99.023502

Okamoto, T., & Hu, W. 2003, Phys. Rev. D, 67, 083002, doi: 10.1103/physrevd.67.083002

Omori, Y. 2024, MNRAS, 530, 5030, doi: 10.1093/mnras/stae1031

Omori, Y., Chown, R., Simard, G., et al. 2017, ApJ, 849, 124, doi: 10.3847/1538-4357/aa8d1d

Omori, Y., Giannantonio, T., Porredon, A., et al. 2018, Phys. Rev. D, 100, 043501, doi: 10.1103/physrevd.100.043501

Osborne, S. J., Hanson, D., & Doré, O. 2013, J. Cosmology Astropart. Phys., 2014, 024, doi: 10.1088/1475-7516/2014/03/024

OSG. 2006, OSPool, OSG, doi: 10.21231/906P-4D78

—. 2015, Open Science Data Federation, OSG, doi: 10.21231/0KVZ-VE57

Pan, Z., Bianchini, F., Wu, W. L. K., et al. 2023, Phys. Rev. D, 108, 122005, doi: 10.1103/PhysRevD.108.122005

Peebles, P. J. E. 1967, ApJ, 147, 859, doi: 10.1086/149077

—. 1968, ApJ, 153, 1, doi: 10.1086/149628

Penzias, A. A., & Wilson, R. W. 1965, ApJ, 142, 419, doi: 10.1086/148307

Perivolaropoulos, L., & Skara, F. 2022, New A Rev., 95, 101659, doi: 10.1016/j.newar.2022.101659

Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, ApJ, 517, 565, doi: 10.1086/307221

Plagge, T., Benson, B. A., Ade, P. A. R., et al. 2010, ApJ, 716, 1118, doi: 10.1088/0004-637X/716/2/1118

Planck Collaboration. 2018, Planck Legacy Archive. https://www.cosmos.esa.int/web/planck/pla

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014a, A&A, 571, A1, doi: 10.1051/0004-6361/201321529

—. 2014b, A&A, 571, A12, doi: 10.1051/0004-6361/201321580

—. 2014c, A&A, 571, A17, doi: 10.1051/0004-6361/201321543

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2015a, A&A, 594, A10, doi: 10.1051/0004-6361/201525967

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2015b, A&A, 594, A5, doi: 10.1051/0004-6361/201526632

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016a, A&A, 594, A15, doi: 10.1051/0004-6361/201525941

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016b, A&A, 594, A8, doi: 10.1051/0004-6361/201525820

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2019, A&A, 641, A5, doi: 10.1051/0004-6361/201936386

Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020a, A&A, 641, A4, doi: 10.1051/0004-6361/201833881

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020b, A&A, 641, A6, doi: 10.1051/0004-6361/201833910

—. 2020c, A&A, 641, A8, doi: 10.1051/0004-6361/201833886

Planelles, S., Schleicher, D. R. G., & Bykov, A. M. 2015, Space Sci. Rev., 188, 93, doi: 10.1007/s11214-014-0045-7

Pordes, R., Petravick, D., Kramer, B., et al. 2007, in 78, Vol. 78, J. Phys. Conf. Ser., 012057, doi: 10.1088/1742-6596/78/1/012057

Prabhu, K., Raghunathan, S., Millea, M., et al. 2024, arXiv e-prints, arXiv:2403.17925, doi: 10.48550/arXiv.2403.17925

Pratten, G., & Lewis, A. 2016, J. Cosmology Astropart. Phys., 2016, 047, doi: 10.1088/1475-7516/2016/08/047

Qu, F. J., Sherwin, B. D., Madhavacheril, M. S., et al. 2024, ApJ, 962, 112, doi: 10.3847/1538-4357/acfe06

Raghunathan, S., Holder, G. P., Bartlett, J. G., et al. 2019, J. Cosmology Astropart. Phys., 2019, 037, doi: 10.1088/1475-7516/2019/11/037

Raghunathan, S., & Omori, Y. 2023, ApJ, 954, 83, doi: 10.3847/1538-4357/ace0c6

Rees, M. J., & Sciama, D. W. 1968, Nature, 217, 511, doi: 10.1038/217511a0

Reichardt, C. L., Patil, S., Ade, P. A. R., et al. 2021, ApJ, 908, 199, doi: 10.3847/1538-4357/abd407

Riess, A. G., Casertano, S., Yuan, W., Macri, L. M., & Scolnic, D. 2019, ApJ, 876, 85, doi: 10.3847/1538-4357/ab1422

Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, AJ, 116, 1009, doi: 10.1086/300499

Robertson, M., & Lewis, A. 2023, J. Cosmology Astropart. Phys., 2023, 048, doi: 10.1088/1475-7516/2023/08/048

Sachs, R. K., & Wolfe, A. M. 1967, ApJ, 147, 73, doi: 10.1086/148982

Sailer, N., Schaan, E., & Ferraro, S. 2020, Phys. Rev. D, 102, 063517, doi: 10.1103/PhysRevD.102.063517

Schaan, E., Krause, E., Eifler, T., et al. 2017, Phys. Rev. D, 95, 123512, doi: 10.1103/PhysRevD.95.123512

Schmittfull, M. M., Challinor, A., Hanson, D., & Lewis, A. 2013, Phys. Rev. D, 88, 063012, doi: 10.1103/PhysRevD.88.063012

Schneider, P., Ehlers, J., & Falco, E. E. 1992, Gravitational Lenses, doi: 10.1007/978-3-662-03758-4

Seljak, U., & Hirata, C. M. 2004, Phys. Rev. D, 69, 043005, doi: 10.1103/PhysRevD.69.043005

Seljak, U., & Zaldarriaga, M. 1997, Phys. Rev. Lett., 78, 2054, doi: 10.1103/PhysRevLett.78.2054

Sfiligoi, I., Bradley, D. C., Holzman, B., et al. 2009, in 2, Vol. 2, 2009 WRI World Congress on Computer Science and Information Engineering, 428–432, doi: 10.1109/CSIE.2009.950

Smith, K. M., & Zaldarriaga, M. 2007, Phys. Rev. D, 76, 043001, doi: 10.1103/PhysRevD.76.043001

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, ApJ, 396, L1, doi: 10.1086/186504

Sobrin, J. A., Anderson, A. J., Bender, A. N., et al. 2022, ApJS, 258, 42, doi: 10.3847/1538-4365/ac374f

Staniszewski, Z., Ade, P. A. R., Aird, K. A., et al. 2009, ApJ, 701, 32, doi: 10.1088/0004-637X/701/1/32

Story, K. T., Hanson, D., Ade, P. A. R., et al. 2015, ApJ, 810, 50, doi: 10.1088/0004-637X/810/1/50

Sunyaev, R. A., & Zeldovich, Y. B. 1970, Ap&SS, 7, 3, doi: 10.1007/BF00653471

—. 1972, Comments on Astrophysics and Space Physics, 4, 173

Tandoi, C., Guns, S., Foster, A., et al. 2024, arXiv e-prints, arXiv:2401.13525, doi: 10.48550/arXiv.2401.13525

Thornton, R. J., Ade, P. A. R., Aiola, S., et al. 2016, ApJS, 227, 21, doi: 10.3847/1538-4365/227/2/21

Trendafilova, C., Hotinli, S. C., & Meyers, J. 2024, J. Cosmology Astropart. Phys., 2024, 017, doi: 10.1088/1475-7516/2024/06/017

Vallinotto, A. 2012, ApJ, 759, 32, doi: 10.1088/0004-637X/759/1/32 van Engelen, A., Bhattacharya, S., Sehgal, N., et al. 2013, ApJ, 786, 13, doi: 10.1088/0004-637x/786/1/13

van Engelen, A., Keisler, R., Zahn, O., et al. 2012, ApJ, 756, 142, doi: 10.1088/0004-637x/756/2/142

Vieira, J. D., Crawford, T. M., Switzer, E. R., et al. 2010, ApJ, 719, 763, doi: 10.1088/0004-637X/719/1/763

Vieira, J. D., Marrone, D. P., Chapman, S. C., et al. 2013, Nature, 495, 344, doi: 10.1038/nature12001

Weinberg, S. 2000, arXiv e-prints, astro, doi: 10.48550/arXiv.astro-ph/0005265

Wiener, N. 1949, Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications (The MIT Press), doi: 10.7551/mitpress/2946.001.0001

Wong, K. C., Suyu, S. H., Chen, G. C. F., et al. 2020, MNRAS, 498, 1420, doi: 10.1093/mnras/stz3094

Wu, W. L. K., Mocanu, L. M., Ade, P. A. R., et al. 2019, ApJ, 884, 70, doi: 10.3847/1538-4357/ab4186

Yang, Q., & Shen, Y. 2022, ApJS, 264, 9, doi: 10.3847/1538-4365/ac9ea8

Zaldarriaga, M. 2001, Phys. Rev. D, 64, 103001, doi: 10.1103/PhysRevD.64.103001

Zel'dovich, Y. B. 1970, A&A, 5, 84

Zhang, Z. J., Omori, Y., & Chang, C. 2023, MNRAS, 524, 6392, doi: 10.1093/mnras/stad2330