

FROM RAW DATA TO COSMOLOGY WITH THE
ATACAMA COSMOLOGY TELESCOPE

EMILIE ROSE STORER

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
PHYSICS
ADVISER: PROFESSOR JO DUNKLEY

MARCH 2023

© Copyright by Emilie Rose Storer, 2023.

All rights reserved.

Abstract

Measurements of the cosmic microwave background (CMB) are one of the main tools used to constrain the cosmological parameters we use to describe the nature and evolution of our universe. As we study the CMB in greater detail, in addition to learning new and exciting things about the universe, we also become increasingly sensitive to various systematics arising from our instruments or our data processing techniques, so great care must be taken to ensure these do not bias our results.

In this thesis, we present the analysis of data from the Atacama Cosmology Telescope (ACT), an experiment used to measure both the temperature and polarization of the CMB. This includes work done for ACT’s recent data release (DR4) as well as progress made since then.

We describe the process of transforming raw data into maps of the sky, and then angular power spectra which are used as inputs into a likelihood to constrain cosmology. As part of this procedure, we characterize the beams of the instrument, which determine how it responds to different scales on the sky. This is critical for all science analyses, including the accurate measurement of power spectra. During the map-making step, we make additional maps with the specific purpose of checking for systematic effects which may bias our data. Also, once we’ve outlined the power spectrum pipeline, we reproduce ACT DR4 spectra using publicly available code, and in the process explore how different choices made in the analysis influence the final results.

Finally, we investigate the possibility of using the data collected by ACT during the day. This data has not been used for cosmological analyses so far due to the specific challenges it presents, mainly that the beams undergo time-dependent deformations due to heat from the Sun distorting the telescope structure. We explore different methods to mitigate and correct for these daytime beams.

Acknowledgements

Thank you, first and foremost, to Jo, for being such a fantastic advisor. Thank you for being so encouraging, and supporting me in absolutely everything, whether it be doing research, writing, traveling, giving presentations, teaching, taking classes, or serving on committees. Your enthusiasm is contagious — I always left our weekly meetings feeling reenergized, and eager to return to work.

Thank you to Suzanne Staggs, for her useful insights and positive energy. I'm grateful for her close reading of this thesis, and her many thoughtful comments and suggested improvements. The final product is now much better thanks to her.

Thank you to Lyman Page, for guiding my brief foray into instrumentation work, and for offering good advice, both scientific and otherwise.

I've greatly enjoyed being a member of the ACT collaboration. In particular, I've learned a lot from Simone Aiola, Steve Choi, Matthew Hasselfield, Marius Lungu, Adri Duivenvoorden, Thibaut Louis, and perhaps most of all, Sigurd Naess, who in addition to being an entertaining source of random “fun facts” and puzzling scenarios involving black holes, has been a very patient teacher. An extra thank you to Adri for taking over the beam pipeline!

Thank you to the various office mates I've had over the years. In particular, thank you to Zach Atkins, for so graciously sharing airborne germs when we returned to the office, and for the entertaining, occasionally maddening, riddles.¹

To the cosmologist from Peyton Hall: Thank you to Heather Prince, who ensured I saw a bit more of what NYC had to offer (other than Penn Station and the CCA), and to Andrina Nicola, for providing a listening ear over walks and coffee. Thank you to all the others who helped me (and the others from “the other side of the

¹Two ants colliding and reversing directions is *identical* to them just walking past each other!

big sculpture”) feel welcome, including Zack Li, Lachlan Lancaster, Oliver Philcox, Rodrigo Córdova Rosado, Roohi Dalal, Jahmour Givans.

I am grateful for the assistance I’ve received from support staff, including Kate Brosowski, Angela Lewis, Katherine Lamos, Janelle Jupiter, Darryl Johnson (who made sure to always greet me with a smile and an impressive amount of energy whenever I entered Jadwin near the loading docks in the mornings), and the staff at OIT.

Thank you to my father, Andy, for being such a supportive parent. To the point that he wouldn’t miss my defense for anything in the world, “even if I have to walk down there with one of those sticks [an IV pole] in my arm”. I’m glad that wasn’t necessary. Thank you for fostering my interest in science from an early age.

Merci à ma mère, Suzanne, pour m’avoir toujours supportée, même si les sciences dures, ce n’est pas tellement ton affaire. Merci d’avoir encouragé mon côté créatif, incluant mes divers projets scolaires.

Thank you to Seth, for always being there. Especially these last few months. For being annoyingly confident my thesis would turn out all right, and for understanding when I didn’t have time for unpacking moving boxes/sleeping/other fun activities because I was busy writing. I’m grateful for all the seemingly random events, starting with the phase distribution of the initial density perturbations, that led me to you.

Finally, thank you to George and Charlotte for the close companionship (although at times perhaps a bit too close for comfort). An extra thanks to Charlotte for the edits. I’m not sure what “nhmytg” means, for example, and what it has to do with daytime beam cuts, but I’m sure you were trying to help.

To my family

Contents

Abstract	iii
Acknowledgements	iv
List of Tables	xi
List of Figures	xii
1 Introduction	1
1.1 The Cosmic Microwave Background	1
1.1.1 The Early universe	2
1.1.2 Measurements	5
1.1.3 Physics of the Power Spectra	11
1.1.4 Secondary Effects	16
1.1.5 Cosmological Constraints	18
1.1.6 Foregrounds	23
1.1.7 Experiments	31
1.2 The Atacama Cosmology Telescope	40
1.2.1 The Instrument	40
1.2.2 The Data	41
1.2.3 Pipeline Overview	44
2 Maps	47
2.1 Mapmaking	47

2.1.1	Overview	47
2.1.2	Mapmaking Procedure	48
2.1.3	Noise Estimate	50
2.1.4	Point Sources	52
2.1.5	Pointing Correction	54
2.1.6	Pickup maps	55
2.1.7	Transfer Function	56
2.2	DR4 Maps	57
2.3	Map-Based Null Tests	57
2.3.1	Time Constants	59
2.3.2	Elevation of Observations	60
2.3.3	PWV	60
3	Beams	64
3.1	Introduction	65
3.2	Observations	66
3.3	Planet Map-Making	67
3.3.1	Filtering And Noise	67
3.3.2	Bias From Map-Making	69
3.4	Beam Pipeline	73
3.4.1	Map Selection And Processing	73
3.4.2	Radial Profile Fitting	76
3.4.3	Beam Window Functions	89
3.4.4	Additional Corrections	93
3.4.5	Jitter Beams	95
3.4.6	Beam Transform Covariance Matrix	100
3.5	Polarization	105
3.5.1	Main Beam	105

3.5.2	Polarized Sidelobes	113
3.5.3	Leakage Correction	118
3.6	Discussion	122
3.6.1	Beam Products	122
3.6.2	Beam Asymmetry	123
3.6.3	Conclusion	124
3.7	Appendix Materials	125
3.7.1	Shrinking Algorithm for Radial Profile Covariance Matrix . . .	125
3.7.2	Computing the Harmonic Transform of the Beam	128
3.7.3	Beam Solid Angles for Different Effective Frequencies	132
3.7.4	Transforming from $\{Q_r, U_r\}$ to $\{E, B\}$	136
3.7.5	DR3 vs DR4 Beams	139
4	Power Spectra	141
4.1	Formalism	141
4.1.1	Definitions	141
4.1.2	Estimation	144
4.2	Power Spectrum Pipeline	145
4.2.1	Enumeration of the Spectra	145
4.2.2	Masking the Foregrounds	146
4.2.3	Spatial Window Function	148
4.2.4	Ground Pickup and the Fourier-Space Filter	151
4.2.5	Fourier-Space Filter Transfer Function	153
4.2.6	Mapping Transfer Function	154
4.2.7	Calibration	155
4.2.8	Covariance Matrix	155
4.2.9	Binning and the Band Power Window Function	155
4.2.10	Coaddition	156

4.2.11	Consistency Tests	157
4.3	DR4 Results	160
4.4	Reproducing DR4 Spectra	161
4.4.1	NaWrapper	161
4.4.2	PSPipe	180
5	Daytime	182
5.1	Introduction	182
5.1.1	The Data	182
5.1.2	The Main Challenges	185
5.2	Past Work	186
5.3	Map-Making	188
5.3.1	Pointing	188
5.3.2	Beam Cuts	188
5.4	Power Spectrum Analysis	194
5.4.1	Spatial Windows Functions	194
5.4.2	Point Source Subtraction	201
5.4.3	Optimal Weighting	202
5.5	Beams	208
5.5.1	Temperature Beam	209
5.5.2	Polarization and Leakage Beam	214
5.6	Plans for Consistency and Null Tests	216
5.7	Summary of Current Status & Next Steps	218
6	Conclusion	220

List of Tables

2.1	Median Time Constants	59
2.2	Summary of the Time Constant Null Tests	60
2.3	Summary of Elevation and PWV null tests	63
3.1	Summary of Uranus observations — number used/total.	76
3.2	Instantaneous beam solid angles, gains, and FWHMs	92
3.3	Jitter values and uncertainties (arcsec ²).	101
3.4	Effective beam solid angles and uncertainties (nsr).	102
3.5	Effective frequencies [GHz].	133
3.6	Effective beam solid angles (nsr) for the CMB.	134
3.7	Effective beam solid angles (nsr) for synchrotron emission.	134
3.8	Effective beam solid angles (nsr) for dusty sources.	134
3.9	Effective beam solid angles (nsr) for the tSZ effect.	135
4.1	Summary of DR4 Region Footprint Properties	149

List of Figures

1.1	Full-sky map of the primary CMB temperature anisotropies in galactic coordinates.	5
1.2	CMB temperature power spectrum computed with CAMB.	7
1.3	Illustration of the E-mode and B-mode polarization patterns in terms of the Stokes parameters Q and U	9
1.4	CMB power spectra computed with CAMB, both with and without gravitational lensing, plotted using a log scale for the x -axis.	13
1.5	CMB power spectra computed with CAMB, both with and without gravitational lensing, plotted using a linear scale for the x -axis.	14
1.6	Brightness temperature rms as a function of frequency for different CMB foregrounds, in temperature and polarization.	26
1.7	Comparison of <i>COBE</i> , <i>WMAP</i> , and <i>Planck</i> CMB temperature maps, featuring both full-sky maps and a $10^\circ \times 10^\circ$ cutout.	33
1.8	Comparison of <i>Planck</i> , ACT+ <i>Planck</i> , and ACT-only CMB temperature maps at 150 GHz in a $3^\circ \times 3^\circ$ patch of sky, using data from ACT DR5.	34
1.9	Recent measurements of the CMB temperature anisotropy and polarization.	38
1.10	Photograph of ACT.	41
1.11	ACT detector arrays collecting data during each year of observation. .	42

1.12	Cumulative coverage map of the ACT observations from 2013 to 2016.	43
1.13	Rough outline of the ACT data analysis pipeline that transforms raw data from the telescope into constraints on cosmological parameters.	44
2.1	Intensity map of the D56 patch made using 2014 PA2 data.	50
2.2	Example of the convergence of modes in a map with the conjugate gradient algorithm.	51
2.3	Example of some larger-scale modes having converged after 3 passes of the mapmaker, each with 300 conjugate gradient steps.	52
2.4	Point source observed in the 2014 PA2 maps of the BN patch.	54
2.5	Example of a pickup map created for 2014 D56 PA2 nighttime data.	56
2.6	Rectangular cutout of the s13–s16 150 GHz maps in the D56 region.	58
2.7	Summary of the results of custom null tests.	61
2.8	Noise for different PWV splits for D56.	62
3.1	Example of a simulated planet observation used to characterize the mapping transfer function.	71
3.2	The average radial profile of the simulated Uranus observations for S16 PA2 at 150 GHz.	72
3.3	Distribution of the total number of Uranus observations that were made and ultimately became part of the final beam analysis.	74
3.4	Two maps of individual Uranus observations for S15 PA2 at 150 GHz.	75
3.5	Example of the offset fit for the radial profile of one Uranus observation for S14 PA2 at 150 GHz.	77
3.6	Measured radial beam profile for S15 PA2 at 150 GHz and the best-fit model, along with the terms in the model for the beam.	88
3.7	Maps of the main beam for S15 PA2 and PA3 at 150 GHz.	90

3.21 Ratio of the beam transforms for DR3 and DR4 at 150 GHz, for S13 and S14.	140
4.1 Graphical depiction of the spectra we use when averaging the cross-spectra for different data splits, for different scenarios.	147
4.2 Illustration of the process for obtaining the footprint for the D56 region.	150
4.3 Illustration of the process for obtaining the spatial window for use in the power spectrum analysis.	152
4.4 The Fourier-space filter transfer function, estimated from simulations and an analytic estimate, for the D56 region.	154
4.5 Summary of pair-wise consistency checks for all combinations of spectra.	159
4.6 The deep and wide spectra used in the likelihood.	162
4.7 The result of estimating the power spectrum in notebook 8 using all 4 splits of data.	164
4.8 The result of estimating the power spectrum in notebook 8 using only 2 of the 4 splits of data.	165
4.9 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when using either only 2 or all 4 splits of data, divided by the uncertainties on the DR4 spectrum.	166
4.10 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when we build the spatial window function ourselves or use the one created as part of the DR4 pipeline, divided by the uncertainties on the DR4 spectrum.	167
4.11 Difference between the spatial window function assembled in notebook 8 and the one created as part of the DR4 power spectrum pipeline. . .	169
4.12 A $1^\circ \times 1^\circ$ cutout of the fractional difference between the windows after renormalizing, centered around a point source near RA = 32° , Dec = -5° .	170

4.13 Radial profiles (centered on as masked source) for different point source mask apodizations.	170
4.14 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when using a pre-filtered map or performing the Fourier-space filtering ourselves, divided by the uncertainties on the DR4 spectrum.	172
4.15 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when using a pre-filtered map, divided by the DR4 spectrum.	172
4.16 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when we add (or not) a one-pixel offset to the Fourier-space filter, divided by the uncertainties on the DR4 spectrum.	173
4.17 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when we multiply (or not) the maps by a footprint mask before filtering them, divided by the uncertainties on the DR4 spectrum.	174
4.18 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when the weak (non-masked) point sources are added to the maps either before or after the Fourier-space filtering is performed, divided by the uncertainties on the DR4 spectrum.	175
4.19 Map used for the DR4 power spectrum analysis, after the weak point sources were added to the map but before the Fourier-space filtering was performed.	175
4.20 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when either adding the weak sources to the maps ourselves, or using the DR4 maps with the weak sources already added, divided by the uncertainties on the DR4 spectrum.	176

4.21 Difference between a (renormalized) map from the DR4 analysis with the weak point sources already added and a map to which we added the weak point sources ourselves.	176
4.22 A $1^\circ \times 1^\circ$ map cutout, centered on a bright source near RA = 32° , Dec = -5°	178
4.23 The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when either zero-padding the map or not, divided by the uncertainties on the DR4 spectrum.	180
5.1 Inverse-variance maps for PA4 at 150 GHz (coadded across all map splits) for night, day deep, and day wide.	184
5.2 Detector arrays used during each year of observation from 2017–2021. .	185
5.3 Examples of point sources seen in maps made from individual daytime TODs.	186
5.4 Footprints for the northern daytime observations from 2014–2016, along with the BN footprint from DR4.	187
5.5 Point source seen in the 2018 PA4 150 GHz maps.	189
5.6 Scanning regions for the daytime maps.	190
5.7 Histograms illustrating the number of TODs per hour in the PA5 150 GHz daytime maps, both with and without the beam cuts.	192
5.8 Histograms illustrating the number of TODs per hour, depending on the time of year, in the PA5 150 GHz daytime maps, both with and without the beam cuts.	193
5.9 Spatial window for PA5 at 150 GHz, created for the DR6 (nighttime) AA patch.	195
5.10 Spatial windows for PA4, PA5, and PA6 at 150 GHz, created for the day wide region, using the same settings as for the current nighttime analysis.	197

5.11 Selection of binary footprints for the PA4 150 GHz day wide maps, obtained using different techniques.	198
5.12 A zoom-in on the region of the inverse variance map corresponding to a hole in the binary mask.	199
5.13 Two examples of point source subtraction in both a nighttime and daytime map.	203
5.14 Mollweide projection of the daytime sky patches (split into north and south), in both equatorial and Galactic coordinates.	205
5.15 Ratio of the TT power spectrum error bars (including cosmic variance) for PA4 at 150 GHz, for various daytime scenarios vs nighttime. . . .	206
5.16 Ratio of the TT nighttime power spectrum error bars (including cosmic variance) for PA4 at 150 GHz, for different spatial weighting choices. .	208
5.17 Daytime jitter fit for a source in the day wide-north region, for PA5 at 150 GHz.	210
5.18 Preliminary distribution of the daytime jitter beam correction, for the day wide-north maps for PA5 at 150 GHz.	211
5.19 Preliminary beams for PA5 at 150 GHz, for the nighttime (DR6) and the different daytime patches.	212
5.20 Close-up of a point source in a daytime (temperature) map.	212
5.21 The day-to-night beam ratio α_b for S15 PA2 at 150 GHz, found using the power spectrum approach.	214
5.22 An illustration of the effect of the power-spectrum-based daytime beam correction.	215
5.23 Cumulative white noise depth distributions for temperature and polarization for recent ACT data releases, as well as projections for DR6. .	219

Chapter 1

Introduction

1.1 The Cosmic Microwave Background

Curious to learn more about the world surrounding them, young children often repeatedly ask “Why?” If a child fails to ever grow out of this phase, as they get older and dig deeper, inevitably they will at some point be led to the oldest light in the universe, the cosmic microwave background (CMB).¹ Having been discovered recently (as far as the timescales of physics are concerned) in 1965, the CMB is now one of the main observables studied in cosmology, the branch of science dealing with the nature and evolution of the universe. Measurements of the CMB have played an important role in helping us reach our current level of understanding of the universe. While the light itself is over 13 billion years old and our discovery of it dates to nearly 60 years ago, there is still much groundbreaking work that can be done from studying the CMB, including learning new fundamental physics.

¹ And perhaps they will even make it the topic of their Ph.D.

1.1.1 The Early universe

We observe the universe as currently expanding and cooling. Extrapolating back in time, it must once have been dense and hot. In the very early universe, when it was hot and dense, photons and baryons interacted with one another, creating a coupled photon-baryon fluid, or plasma. Thomson scattering between photons and electrons kept them in thermal equilibrium (photons had a short mean free path between interactions), and Coulomb scattering between electrons and protons acted as glue between them.

Particles of the photon-baryon fluid fell into gravitational potential wells, resulting in a compression of the fluid. The compressing fluid up-scattered photons to higher energies. This caused increased radiation pressure, which opposed the compression, resulting in rarefaction. These two phenomena of compression and rarefaction generated acoustic waves and resulted in temperature oscillations. Physically, the temperature oscillations correspond to the heating and cooling of a fluid that is being compressed and rarefied by acoustic waves (Hu et al., 1997; Hu & Dodelson, 2002).

When the universe was $\sim 380,000$ years old² (corresponding to a redshift z of ~ 1100), it cooled enough — to a temperature of ~ 3000 K — for electrons and protons to bind together, forming neutral hydrogen and helium.³ This event is known as recombination.⁴ There were then fewer electrons to scatter the photons, which allowed the photons to begin streaming freely through space (the mean free path of photons became long, of order the size of the horizon).⁵ In other words, the universe went from

²Current best estimates place the age of the universe at around 13.8 Gyr (Planck Collaboration VI, 2020).

³This temperature of ~ 3000 K may seem low, given that the ionization potential of hydrogen, 13.6 eV, corresponds to roughly 105,000 K. But it isn't! A rough temperature of $\sim 3000 - 4000$ K can be derived simply, using equilibrium theory, in particular the Saha ionization equation. More in-depth calculations can be found in Peebles (1968) and Zeldovich et al. (1968).

⁴The term recombination is somewhat misleading, since electrons and protons had not been combined before, but this term is still used for historical reasons, since it pre-dates our current understanding of the early universe.

⁵The free-streaming of photons is known as decoupling (since the photons are no longer coupled

opaque to transparent. This is when the CMB, or “primordial fireball radiation,” was released (Dicke et al., 1965). It is the oldest light in the universe, allowing us to see back to the primordial plasma. As the universe expanded, this radiation cooled, and formed what we now observe as the cosmic microwave background. The oscillations due to the counteracting forces of gravity and radiation pressure ceased when photons decoupled from baryons, and these frozen oscillations form the temperature fluctuations in the CMB, shown in Figure 1.1.

Recombination happened everywhere, but given the finite speed of light, we see it as a distant, spherical shell — the surface of last scattering.⁶

The collisions prior to recombination led to the electrons and photons being in equilibrium, so the photons had a blackbody spectrum when they set off through space. The expansion of the universe can alter this distribution only by redshifting the temperature of the radiation, so we expect this radiation to still be a blackbody spectrum today.⁷ The peak of a blackbody spectrum is related to its temperature according to Wien’s displacement law: $\lambda_{\text{peak}} \propto 1/T$.

The expansion of the universe is parametrized by a dimensionless scale factor, $a(t)$, which relates the comoving separation⁸ of two points to the physical (or proper) distance between them.⁹ By definition, the scale factor today is 1. The relation between the redshift of the light we receive from an object and the scale factor at the time the object emitted the light is $a(t) = 1/(1+z)$. The scale factor was $\sim 1/1100$

⁶Although recombination was not instantaneous — it had some finite duration.

⁷The fact that a blackbody spectrum is preserved in an expanding universe was shown in Tolman (1934) (p.171). The expansion of space stretches its peak emission to a lower temperature.

⁸The comoving distance between two points is a measure of their separation which factors out the expansion of the universe.

⁹We observe the universe to be expanding, that is, $a(t)$ to be increasing, and extrapolating back, we find that $a \rightarrow 0$ as $t \rightarrow 0$. However, we do not understand the physics at the high energies that occur as $a \rightarrow 0$, so we should not assume much knowledge about what happens as $t \rightarrow 0$.

at recombination, so the wavelength of CMB photons has been stretched by a factor of 1100 since then, meaning their energies have decreased by that factor. Hence, the photons that were then 3000 K are now $\sim 3000 \text{ K}/1100 \sim 2.72 \text{ K}$.¹⁰

So, where did the gravitational potential wells that led to the oscillations of the photon-baryon fluid come from? It is thought that quantum fluctuations were the source of primordial perturbations. These primordial perturbations have been found to be dominated by scalar, adiabatic¹¹ Gaussian fluctuations with an approximately scale-invariant power law. They were the starting point of the oscillations of the photon-baryon fluid, and the seeds for the large-scale structure we see today.¹²

In addition to scalar perturbations (i.e., perturbations in the scalar densities), one might also consider vector or tensor perturbations. Vector perturbations can be ignored here, as they decay rapidly with time (Lifshitz, 1946). Scalar and tensor perturbations evolve independently, and so they can be considered separately.

Unlike scalar perturbations, which are amplified by gravitational collapse, tensor perturbations, in the form of gravitational waves, start to decay when they enter the horizon (Lifshitz, 1946).¹³ Due to this damping, only the tensor perturbations on the largest scales (so those which entered the horizon latest) persist long enough to impact the CMB. Several models of inflation, a period of rapid expansion in the early universe, (including the simplest ones) predict primordial gravitational waves (Guth, 1981; Guth & Weinberg, 1983) (see, e.g., Kamionkowski & Kovetz (2016) for a recent

¹⁰Temperature scales as $T \propto a^{-1}$.

¹¹In the case of adiabatic perturbations, the overdensities in dark matter, baryons, photons, and neutrinos are all proportional to one another, so the ratio of energy densities across all these components is the same everywhere. Since these perturbations affect the matter content, according to Einstein's equations, they are associated with curvature perturbations.

¹²The perturbations lead to small overdense regions, which attracted both dark matter and baryons through the force of gravity. The dark matter clumps grew while the baryons underwent oscillations as part of the photon-baryon fluid. Once the baryons were no longer coupled to photons, they were free to join the growing clumps. These concentrations of matter grew to form all the structure present in the universe today.

¹³There seem to have been superhorizon scalar perturbations, and thus there might have been superhorizon tensor perturbations as well.

review). Other models for the early universe (e.g., [Steinhardt & Turok \(2002\)](#); [Ijjas & Steinhardt \(2018\)](#)) do not predict observable levels, hence they can be observationally distinguished.

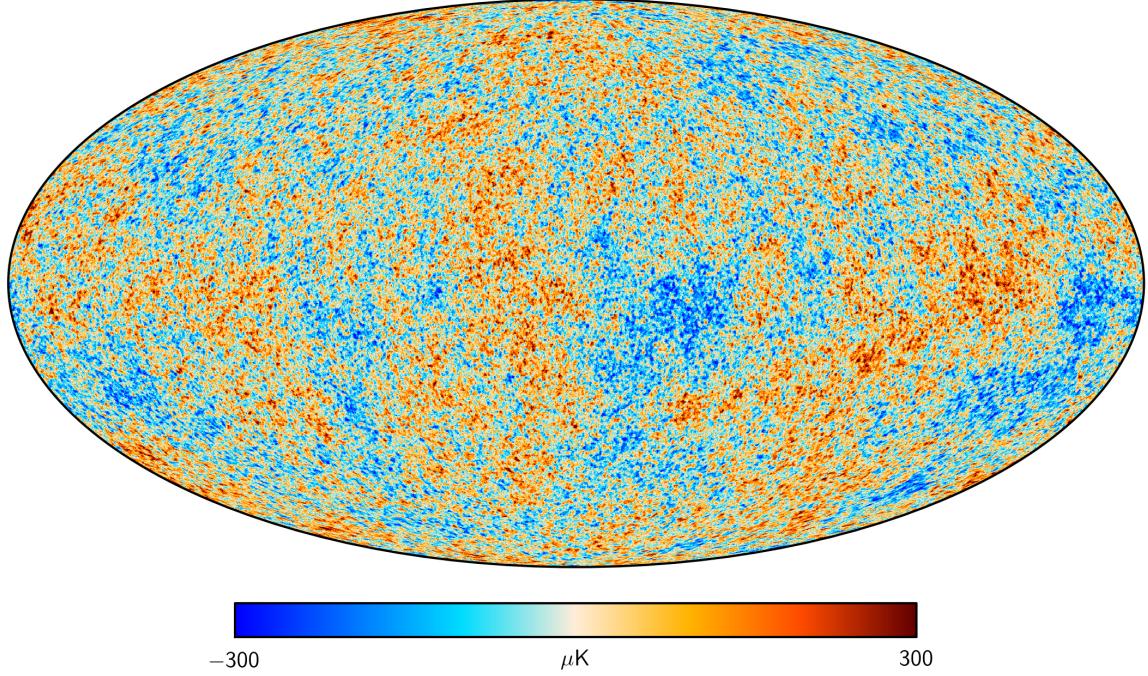


Figure 1.1: Full-sky map of the primary CMB temperature anisotropies in galactic coordinates, derived from the joint analysis of *Planck* and *WMAP* data, as presented in [Planck Collaboration I \(2016\)](#). Courtesy of ESA and The Planck Collaboration.

1.1.2 Measurements

The existence of the isotropic blackbody radiation was first predicted in the 1940s, expected to be on the order of 5 K ([Alpher & Herman, 1948](#); [Alpher & Herman, 1948](#); [Gamow, 1948](#)). This radiation was first measured in 1965, with an antenna temperature of 3.5 ± 1.0 K at 4 GHz ([Penzias & Wilson, 1965](#)). It was quickly identified as the CMB ([Dicke et al., 1965](#)). The blackbody spectrum of the CMB was confirmed by the FIRAS instrument on the *COBE* satellite in the 90s, which measured the spectrum in the range from 60 to 600 GHz, and found a temperature of 2.725 ± 0.002 K ([Mather et al., 1994](#)). This corresponds to a spectrum which peaks at a

frequency of 160.2 GHz, which corresponds to a wavelength of 1.9 mm.¹⁴ While the CMB is relatively isotropic, COBE also detected small anisotropies, of order 10^{-5} (Smoot et al., 1992). The study of these anisotropies, which encode a vast amount of information about our universe, has been the main focus of CMB experiments since then.

Individually, the anisotropies do not reveal any useful cosmological information — our models do not predict the locations of individual hot or cold spots. Rather, our interests lie in the statistics of the anisotropies.

The CMB anisotropies are well approximated as a Gaussian random field. So far, no departures from Gaussianity have been observed in the primordial CMB. In fact, primordial non-Gaussianities are very tightly constrained (Planck Collaboration IX, 2020).¹⁵ The properties of a Gaussian field are fully described by its two-point correlation function. Therefore, we study the power spectrum of CMB fluctuations, that is, the spherical harmonic transform (the equivalent of the Fourier transform for the curved sky) of the two-point correlation function. This allows us to reduce the information in a sky map to a more concise form. Further details about the formalism and computation of power spectra are presented in Section 4.1.

We study the power spectrum C_ℓ^{TT} as a function of multipole ℓ , where each ℓ is approximately related to an angular scale of $180^\circ/\ell$ on the sky. We often work with the quantity

$$D_\ell^{TT} = \frac{\ell(\ell+1)}{2\pi} C_\ell^{TT}, \quad (1.1)$$

as shown in Figure 1.2.

The spectra are generally plotted for $\ell \geq 2$, where $\ell = 2$ is the quadrupole term.

¹⁴Hence the term cosmic *microwave* background, since we see it peak in the microwave band of electromagnetic spectrum.

¹⁵While the primordial CMB is Gaussian, secondary effects, such as SZ effects and gravitational lensing, produce non-Gaussian signatures. This is also true of foregrounds.

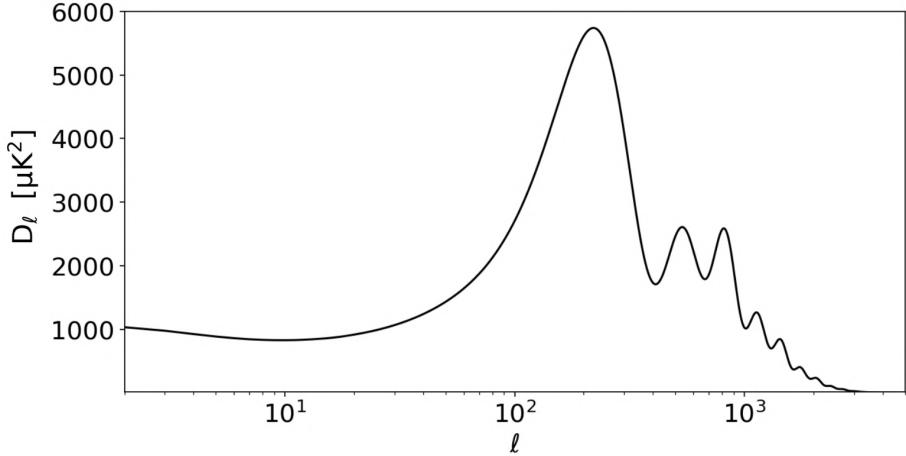


Figure 1.2: CMB temperature power spectrum computed with CAMB (Lewis & Challinor, 2011), using the best-fit Λ CDM parameters from Planck Collaboration VI (2020).

The amplitude of C_ℓ^{TT} for $\ell = 0$, the monopole, is the average temperature of the CMB across the sky, $T = 2.7255 \pm 0.0006$ K (Fixsen, 2009). The $\ell = 1$ term is the dipole, which is dominated by the Doppler shift due to the motion of the Earth with respect to the CMB. This is of amplitude 3.3621 ± 0.0010 mK (Planck Collaboration I, 2020). Both these terms are usually subtracted from the data when studying CMB anisotropies.

In addition to temperature, we also measure the polarization of the CMB. We can look at the CMB in terms of the Stokes parameters I, Q, U, V . The first parameter, I , is the intensity, or temperature, of the signal, also denoted T . Parameters Q and U are the linear polarization. Finally, V is the circular polarization, which we do not consider here, as it is expected to be negligible for the CMB (e.g., Alexander & McDonough, 2019).¹⁶

Rather than Q and U , a more natural approach is to decompose the field into a curl-free component E and a divergence-free component B , illustrated in Figure 1.3

¹⁶No circular polarization of the CMB has been detected, but upper limits have been placed (Nagy et al., 2017; Padilla et al., 2020).

(Kamionkowski et al., 1997; Zaldarriaga & Seljak, 1997).¹⁷ Unlike Q and U , the E and B polarization fields do not depend on our choice of coordinates. The polarization pattern around an E-mode cold spot ($E < 0$) is radial, whereas around an E-mode hot spot ($E > 0$) it is tangential. Meanwhile, the polarization pattern exhibits a swirling pattern in different orientations around B-mode cold and hot spots. E-modes exhibit even parity, while B-modes exhibit odd parity.

Overall, there are four power spectra we measure: the auto spectra C_ℓ^{TT} , C_ℓ^{EE} , and C_ℓ^{BB} , and the cross spectrum between temperature and E polarization, C_ℓ^{TE} . The B polarization is not expected to be correlated with temperature or E polarization in the standard model.¹⁸

Scalar perturbations give rise to both the temperature and E-mode polarization patterns in the CMB. (Scalar perturbations can only lead to polarization modes with even parity, so they can only generate E-mode polarization, not B-modes (Kamionkowski et al., 1997; Zaldarriaga & Seljak, 1997).) At the last scattering surface, Thomson scattering in the presence of local quadrupole anisotropies leads to polarization. The scattering of isotropic radiation produces no net polarization; orthogonal polarizations from orthogonal incident directions cancel out. A local quadrupole must exist in the incoming radiation from the point of view of an electron in order for the outgoing radiation after scattering to have a net polarization. At recombination, polarization-producing quadrupoles are primarily due to velocity gradients in the flow of the primordial plasma, as it flows from hot spots to cold spots and vice-versa. In the rest frame of an electron in such a flow, radiation has a quadrupolar pattern.

Unlike the temperature anisotropies, polarization is only generated by scattering. Prior to recombination, in the primordial plasma, photons and electrons were tightly

¹⁷The notation E and B is used because of the analogy with electric and magnetic fields (Zaldarriaga & Seljak, 1997).

¹⁸This is because T and E have even parity while B has odd parity, so if all interactions involving CMB photons conserve parity, $C_\ell^{TB} = C_\ell^{EB} = 0$.

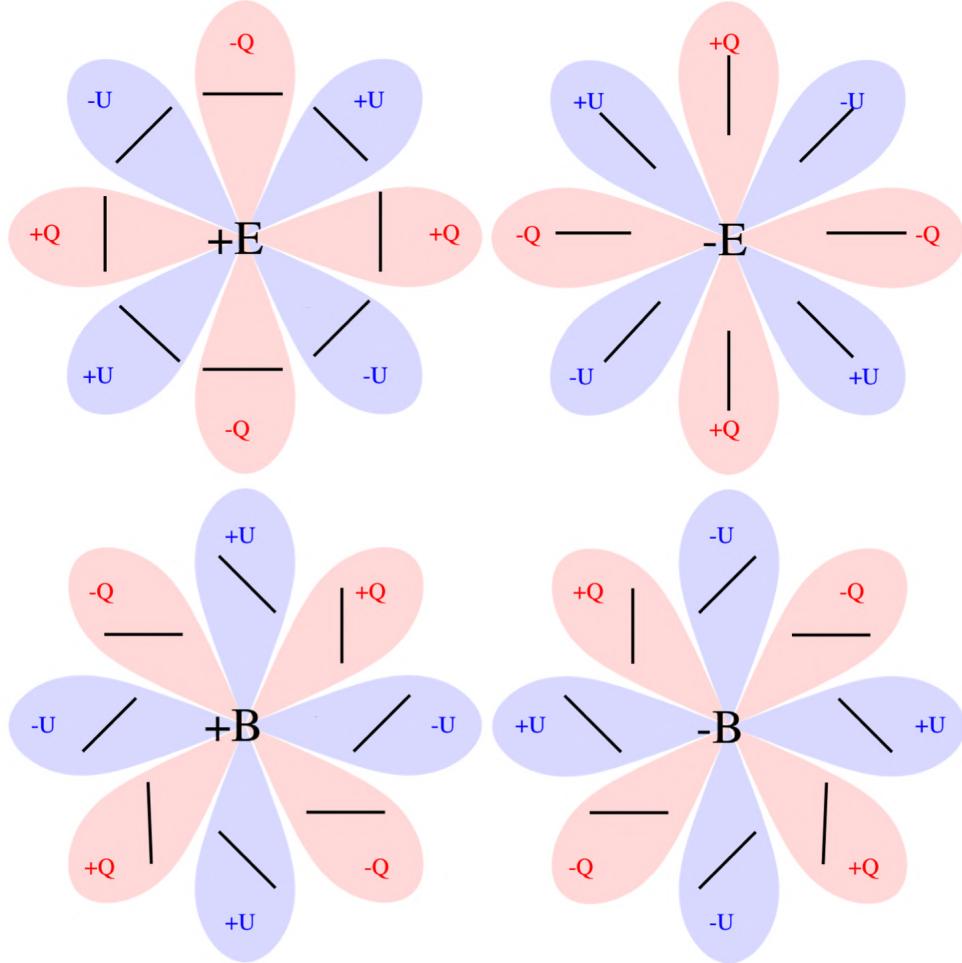


Figure 1.3: Illustration of the E-mode and B-mode polarization patterns in terms of the Stokes parameters Q (red) and U (blue). In coordinates where the top of the page is to the north, a N-S (E-W) polarization corresponds to $Q > 0$ ($Q < 0$) and a NE-SW (NW-SE) component corresponds to $U > 0$ ($U < 0$). The polarization pattern around an E-mode cold spot ($E < 0$) is radial, whereas around an E-mode hot spot ($E > 0$) it is tangential. Meanwhile, the polarization pattern exhibits a swirling pattern in different orientations around B-mode cold and hot spots. Figure courtesy of Sigurd Naess.

coupled. The many successive Thomson scatterings randomized the directions of photons. However, as recombination proceeded, the mean free path of the photons increased. This allowed photons to probe spatial variations. The difference in the fluid velocity for successive scatterings became significant. This lead to quadrupole anisotropies, polarizing the CMB photons. By the time of last scattering, a lot of the electrons had already recombined into neutral hydrogen. There were thus a reduced number of scatterers available to produce polarization. This is why the fractional polarization of the CMB is only about 10%.

Since the T and E fluctuations are both caused by the same underlying scalar perturbations, they are correlated, so their cross-spectrum, C_ℓ^{TE} , contains useful information. The peaks in the E-mode power spectrum correspond to modes at velocity extrema at recombination, so they are out of phase with the temperature peaks, which correspond to modes at displacement extrema.

While scalar perturbations do not give rise to a primordial B-mode polarization signal, tensor perturbations could,¹⁹ in the form of primordial gravitational waves. (Tensor perturbations could also give rise to temperature fluctuations and E-modes, but at a negligible level compared to what we observe.²⁰) While we have measured the lensed B-mode power spectrum (see Section 1.1.4), so far, no primordial B-modes have been detected, but upper limits have been placed. A primordial B-mode signal would be suppressed at small scales, since gravitational waves decay once they enter the horizon, so only modes close to or larger than the size of the horizon at decoupling (a few degrees) contribute significantly.

¹⁹Foregrounds (see Section 1.1.6) also give rise to B-mode polarization.

²⁰We can constrain the amplitude of tensor perturbations using the TT and/or EE spectra (e.g., [Planck Collaboration VI, 2020](#)), but these limits are not competitive with those from BB spectra (e.g., [BICEP/Keck Collaboration, 2021](#)).

1.1.3 Physics of the Power Spectra

We would like to relate the angular power spectrum of the CMB fluctuations we measure to the primordial spectrum of fluctuations and the initial conditions of the universe. To do this, we need to describe the evolution of the early universe, so how its different components interacted together and with the gravitational field. The math involves a combination of the Einstein, Boltzmann, and fluid equations. The Einstein field equations capture the relation between stress-energy perturbations and geometric spacetime perturbations, the Boltzmann equations describe the behavior of relativistic matter (photons and neutrinos), and the fluid equations outline the evolution of non-relativistic matter (baryons and dark matter). This allows us to predict the statistical properties of the CMB anisotropies from a set of initial conditions. Or conversely, we can infer initial conditions from our measurements of the CMB. Rather than review these equations at length here, we will take a more qualitative approach, focused on understanding the underlying physics and how different factors influence the CMB power spectra.

We can relate the CMB power spectra, shown in Figures 1.4 and 1.5, to the underlying physical phenomena.

For instance, the series of peaks in the D_ℓ^{TT} spectrum is due to the oscillations of the photon-baryon fluid prior to recombination. The oscillations happened at all scales, up to the size of the horizon at the end of recombination. As described above, gravity pulled the photon-baryon fluid into potential wells, then pressure countered this and pushed the fluid outward. The oscillations that were at their extrema at the time of recombination show up as the peaks. The first peak is at $\ell \approx 200$, which corresponds to a scale of roughly 1° on the sky. This corresponds to modes that entered the horizon, then collapsed to maximum compression by the time of recombination. The second peak corresponds to modes that entered the horizon, collapsed, then ex-

panded back and were at their state of maximum rarefaction when recombination occurred. The third peak corresponds to modes that went through a compression, expansion, then reach maximum compression again for recombination. And so on. So the odd-numbered peaks corresponds to compressions and the even-numbered peaks correspond to rarefaction. Baryons contribute the mass (gravity) for these oscillations, but they don't contribute to the pressure. This leads to a larger amplitude for odd-numbered peaks than for even-numbered ones.

There are several phenomena responsible for the features of the CMB anisotropies. The primary effects are summarized here:

Density perturbations are produced in high-density (low-density) regions, when radiation is compressed (rarefied) due to the coupling between matter and radiation, leading to a higher (lower) temperature. This leads to some intrinsic temperature fluctuations. When photons then decouple from the baryons and begin streaming, there are gravitational redshift (or blueshift) effects as each photon escapes the local potential. The intrinsic temperature fluctuations and the gravitational redshifts/blueshifts partially cancel out, leading to the effective temperature fluctuations. This is known as the **Sachs-Wolfe** effect (also known as the gravitational Sachs-Wolfe effect or the ordinary Sachs-Wolfe effect). For adiabatic perturbations, the gravitational redshift/blueshift effects win out over the intrinsic temperature fluctuations, so the high-density regions appear as cold spots in the CMB, and vice-versa. The D_ℓ^{TT} power spectrum is approximately flat at $\ell \lesssim 100$. This feature is called the Sachs-Wolfe plateau ([Sachs & Wolfe, 1967](#)). It is due to the gravitational redshift effects for perturbations larger than the horizon at recombination. These anisotropies have not significantly changed from the primordial power spectrum, since for scales larger than the horizon, causality dictates that the fluctuations do not evolve. (It's flat because the primordial perturbations are approximately scale-invariant.)

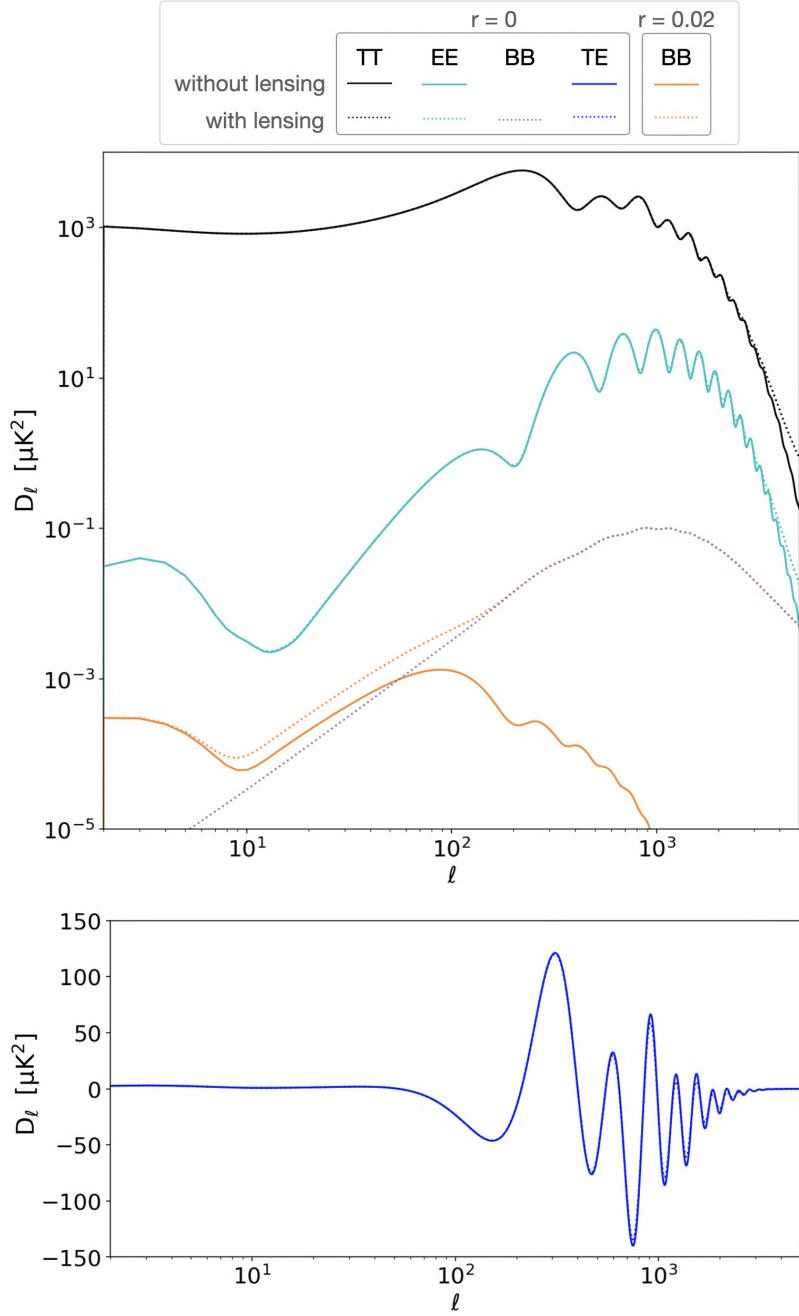


Figure 1.4: CMB power spectra computed with CAMB (Lewis & Challinor, 2011), both with (dotted lines) and without (solid lines) gravitational lensing (see Section 1.1.4), using the best-fit ΛCDM parameters from Planck Collaboration VI (2020) with tensor-to-scalar ratio $r = 0$ (see Section 1.1.5). When $r = 0$ there are no primordial B-modes, hence there is no BB spectrum to measure in the absence of lensing (no solid purple line). A BB spectrum is also plotted for a model with $r = 0.02$. The TT, EE, and TE spectra for the $r = 0.02$ model would be indistinguishable from the $r = 0$ case in this plot. A log scale is used for the x -axis to place emphasis on the low- ℓ regime.

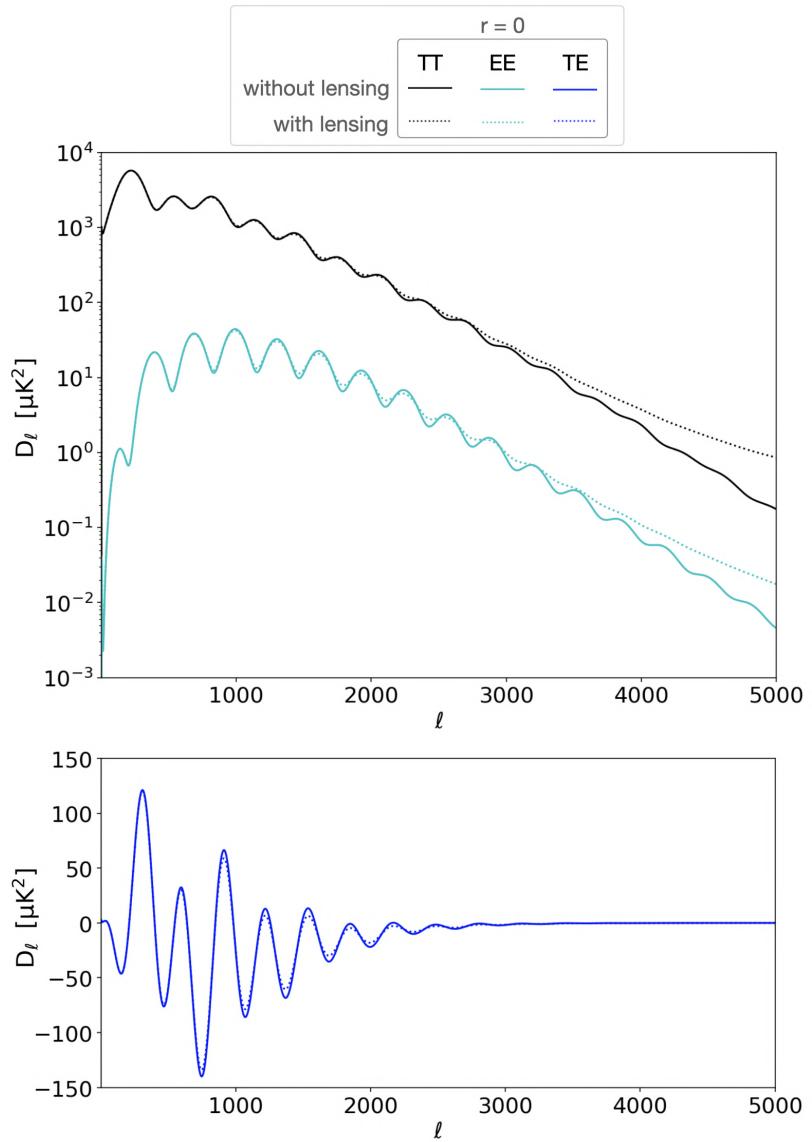


Figure 1.5: As in Figure 1.4, but with a linear scale in multipole (x -axis) to better showcase the acoustic peaks and damping tails. Due to their low amplitude, no BB spectra are plotted here.

Velocity perturbations arise from Doppler shift in frequency due to the velocity of the plasma at the time of recombination. This helps explain why the troughs in the D_ℓ^{TT} power spectrum do not extend down to zero power. They are partly filled by the Doppler effect, since they correspond to velocity extrema (the density and velocity perturbations are $\pi/2$ out of phase with each other). In temperature, the density induced anisotropies are slightly larger, hence the existence (and locations) of the peaks in the spectrum.

The falling off of the power spectrum at high ℓ is **Silk damping** (or diffusion damping) ([Silk, 1968](#)). During recombination, photons scattered off the remaining free electrons, executing a random walk. This lead to hot and cold spots mixing together on small scales, damping the CMB’s peaks. This damping is exacerbated by the fact that recombination was not instantaneous, that is, the surface of last scattering has some finite width. So the CMB photons we observe were not all emitted at the same time, hence the fluctuations are not all in phase.

The D_ℓ^{EE} spectrum exhibits similarities with the temperature spectrum, since they are both due primarily to scalar perturbations. To first order, polarization is only generated by velocity perturbations, not density. Thus, the oscillations for D_ℓ^{EE} are sharper than for D_ℓ^{TT} , and in phase opposition. The peak in polarization at $\ell \sim 100$, known as the “recombination peak,” is due to polarization from the surface of last scattering. Furthermore, polarization is also affected by Silk damping. Wavelengths much smaller than the mean free path of the photons are damped, resulting in an overall peak in D_ℓ^{EE} around the scale of the mean free path of the photons.

The D_ℓ^{TE} spectrum can be understood by examining both D_ℓ^{TT} and D_ℓ^{EE} . For instance, the changes in sign of D_ℓ^{TE} correspond to the compression and expansion extrema of D_ℓ^{TT} , when the fluid velocity vanishes. The D_ℓ^{TE} minima and maxima correspond to velocity extrema.

1.1.4 Secondary Effects

In addition to the primary effects listed above, there are several secondary mechanisms which affect the CMB we measure. These are interactions that occur along the line of sight, as the photons travel towards us.

Some time after recombination, the first stars and galaxies formed, and became a source of ionizing radiation. The gas in the universe was then reionized.²¹ Photons once again scattered off free electrons. However, by this time matter had been diluted by the expansion of the universe, so scattering due to reionization is much less significant than it was pre-recombination.

Still, since reionization, some of the photons scatter out of our line of sight, suppressing the power spectrum by a factor that depends on the optical depth to reionization. Also, some large-scale polarization signal is generated due to Thomson scattering, leading to a “reionization bump” in the E-mode and B-mode power spectra at $\ell \lesssim 10$.²² In this case, polarization is created due to electrons scattering the intrinsic CMB quadrupole.

The **integrated Sachs-Wolfe** (ISW) effect is the result of potential wells evolving with time (Sachs & Wolfe, 1967; Rees & Sciama, 1968). This is similar to the gravitational (a.k.a. non-integrated) Sachs-Wolfe effect mentioned earlier, which occurs at the surface of last scattering. In contrast, the ISW effect is produced between the surface of last scattering and today. The ISW effect is significant just after recombination (early ISW effect), and at later times (late ISW effect), when dark energy dominates.²³

²¹Unlike recombination, which was a confusing term, the word reionization here makes sense, since the electrons and protons were ionized once before (prior to “recombination”).

²²The angular scales of the reionization bump and recombination peak are mainly set by the size of the horizon at the time of each event. Since reionization occurs after recombination, the horizon is larger, so the features are larger (they occur at lower ℓ).

²³Studying the late ISW effect could give us more information about the nature of dark energy, but since it only has an effect at large scales $\ell < 10$, its usefulness is limited by cosmic variance (see Section 4.1.1).

The other two secondary phenomena primarily affect the small scales of the CMB.

The **Sunyaev-Zel'dovich** (SZ) effects are where CMB photons scatter off electrons in galaxy clusters. The thermal SZ (tSZ) effect is where CMB photons inverse Compton scattering off hot electrons (Sunyaev & Zeldovich, 1972), whereas the kinetic SZ (kSZ) effect is where CMB photons inverse Compton scattering off moving electrons (Sunyaev & Zeldovich, 1980). (The kSZ effect is typically a few percent of the tSZ effect.) In each case, the inverse Compton scattering distorts the spectrum of the CMB in the direction of galaxy clusters (so the SZ effects are *localized*). Thomson scattering in the galaxy clusters also leads to SZ-induced polarization (Sazonov & Sunyaev, 1999; Lavaux et al., 2004). For a review of the theory behind the SZ effects, see, e.g., Birkinshaw (1999).

Gravitational lensing, the bending of light due to matter overdensities in the universe,²⁴ leads to an apparent shift in the angular position of CMB photons. This results in an effective re-mapping of both temperature and polarization anisotropies (Blanchard & Schneider, 1987). This slightly smooths the peaks of the power spectrum, and shifts some power to smaller scales in the damping tail (Lewis & Challinor, 2006), as can be most clearly seen in Figure 1.5. Lensing also converts some E-mode polarization to B-mode polarization (Zaldarriaga, 2004). This is the primary contribution to the B-mode power spectrum at small scales (see Figure 1.4).

Understanding these secondary effects is essential for extracting cosmological information from the measured power spectra. They can also be of interest astrophysically.

²⁴It is mainly due to the large-scale structure between redshifts 0.5 and 3 (Blanchard & Schneider, 1987; Bernardreau, 1997; Zaldarriaga & Seljak, 1999; Lewis & Challinor, 2006).

1.1.5 Cosmological Constraints

The most widely accepted model of the universe at the moment is the Λ CDM model.

Λ refers to the cosmological constant initially introduced by Einstein into his equations in order to produce a static universe (Einstein, 1917). Once it was discovered that the universe is in fact expanding, Einstein considered the addition of the constant to be a mistake (see O’Raifeartaigh & Mitton, 2018, for an investigation of the legend of Einstein’s “biggest blunder” statement regarding the cosmological constant). Now, Λ is interpreted as the vacuum energy²⁵ responsible for the accelerated expansion of the universe, referred to as dark energy.

CDM refers to cold (non-relativistic) dark matter. Dark matter is so named because, unlike “regular” baryonic matter, it does not emit or absorb electromagnetic radiation. Instead, it only interacts gravitationally. Dark matter constitutes the majority of the matter in the universe.

The Λ CDM model contains six free parameters used to describe our universe (assuming the universe is flat)²⁶. They are as follows:

A_S , scalar amplitude: This is the amplitude of the primordial fluctuations. Changing this amplitude scales the entire power spectrum up or down (this is true for TT, TE, EE). *Planck* measures $A_S = (2.10 \pm 0.03) \times 10^{-9}$ (Planck Collaboration VI, 2020).

n_s , scalar spectral index: This is the tilt of the primordial power spectrum, which affects the tilt of the CMB power spectrum, the relative amount of power on small and large scales. A scale-invariant spectrum corresponds to $n_s = 1$. Assuming the primordial spectrum is an exact power law, *Planck* measures a slight red tilt, $n_s = 0.965 \pm 0.004$ (Planck Collaboration VI, 2020).

²⁵Energy associated with the vacuum, so for which the density remains constant even as the universe expands.

²⁶Without the assumption of a flat universe, seven parameters are needed.

τ , optical depth: The optical depth is a measure of the quantity of free electrons (due to reionization) the CMB photons travel through to reach us. Reionization affects the amplitude of the entire spectrum,²⁷ so its effect on C_ℓ^{TT} is hard to distinguish from A_S . However, as mentioned above, reionization also produces polarization on large scales, so E-mode polarization data at low ℓ can be used to constrain τ , which then allows us to constrain A_S from the amplitude. Recent measurements from *Planck* include $\tau = 0.054 \pm 0.0070$ ([Planck Collaboration VI, 2020](#)) and $\tau = 0.059 \pm 0.006$ ([Pagano et al., 2020](#)).

Ω_b , baryon density: Baryons weight the oscillations of the photon-baryon fluid prior to recombination. A change in Ω_b therefore changes the amplitude of the compressions in the oscillations relative to the peaks, thereby changing the amplitude of the odd-numbered peaks relative to the even-numbered peaks in the CMB spectrum. Increasing the baryon density increases the height of the odd-numbered peaks, and decreasing the density lowers them. *Planck* measures $\Omega_b h^2 = 0.0224 \pm 0.0001$,²⁸ which corresponds to baryons making up $\sim 5\%$ of the energy content of the universe ([Planck Collaboration VI, 2020](#)).

Ω_c , dark matter density: Dark matter affects the gravitational potentials that influence the oscillations of the photon-baryon fluid. Changing the dark matter density alters the relative strength of gravity compared to radiation pressure for the oscillations. Thus, the amplitude of the oscillations (and hence the amplitude of the acoustic peaks) is either raised or lowered. Increasing Ω_c lowers the amplitude of the peaks (and vice-versa). Similar to baryons, dark matter also makes the compressions more significant compared to the rarefactions, so increasing the density also increases the amplitude of the odd peaks relative to the even ones. *Planck* measures $\Omega_c h^2 = 0.120 \pm 0.001$, which translates to dark matter making up $\sim 26\%$ of the energy

²⁷The observed amplitude actually scales as $A_S e^{-2\tau}$.

²⁸Instead of computing Ω_b directly, we often first work with the quantity $\Omega_b h^2$, where $h = H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1})$.

content of the universe (Planck Collaboration VI, 2020). Note that this is more than 5 times greater than baryons.

θ_* , angular acoustic scale: The physical scale of the first acoustic peak corresponds to the sound horizon at recombination, r_* , so it depends on the sound speed prior to recombination. The corresponding observed angular scale depends on the physical scale, as well as the geometry of the universe (assumed to be flat in Λ CDM), and the distance to the surface of last scattering (hence the expansion history of the universe). *Planck* measures $\theta_* = (0.5965 \pm 0.0002)^\circ$, which corresponds to a physical scale of $r_* = 144.39 \pm 0.30$ Mpc (Planck Collaboration VI, 2020).

Derived parameters

We are also interested in a few derived parameters of Λ CDM. These are parameters of interesting physical significance, whose values can be computed from the six base parameters. Some parameters of particular interest are:

H_0 , Hubble constant: The Hubble constant is the value of the Hubble parameter at present time, which of course depends on the evolution of the Hubble parameter $H(t)$. The Hubble parameter is defined as $H(t) \equiv \frac{\dot{a}}{a}$ (where the dot denotes a derivate with respect to coordinate time), where $a(t)$ is the scale factor introduced earlier in the introduction. Changing the Hubble parameter would change the expansion history of the universe, therefore changing the distance to the surface of last scattering. This would alter the apparent angular scale of the acoustic peaks in the CMB power spectrum, therefore shifting the ℓ at which the peaks appear. This parameter has garnered a lot of attention in recent years, as there seems to be some tension between the values of H_0 calculated using different methods. For instance, *Planck* predicts $H_0 = 67.4 \pm 0.5$ km s $^{-1}$ Mpc $^{-1}$ within Λ CDM, whereas some measurements from local sources yield higher values, for example $H_0 = 73.04 \pm 1.04$ km s $^{-1}$ Mpc $^{-1}$ using Type 1a

Supernovae (Planck Collaboration VI, 2020; Riess et al., 2022).²⁹

Ω_Λ , dark energy density: The value of Ω_Λ , can be calculated from Ω_b and Ω_c in the standard Λ CDM model (for a flat universe, where $\Omega_{\text{tot}} = \Omega_b + \Omega_c + \Omega_\Lambda = 1$) (see, e.g., Dodelson, 2003). This also affects the expansion history of the universe and hence the multipoles at which the peaks appear in the power spectra. *Planck* finds $\Omega_\Lambda = 0.6834 \pm 0.0084$, so $\sim 68.5\%$ of the energy density of the universe is dark energy (Planck Collaboration VI, 2020).

σ_8 , late-time clustering amplitude: The late-time clustering (or fluctuation) amplitude is the rms of the density perturbations in spheres of radius $8 h^{-1}\text{Mpc}$ at redshift $z = 0$. This is another way of parametrizing the amplitude of the matter power spectrum (as opposed to the large-scale primordial amplitude A_S). It is often used when studying galaxy clusters. *Planck* finds $\sigma_8 = 0.811 \pm 0.006$ (Planck Collaboration VI, 2020).

There are also several more quantities one can compute given the standard Λ CDM parameters, including the age of the universe, and the redshifts corresponding to points in the universe’s history, such as decoupling or reionization.

Overall, Λ CDM is a remarkably simple standard model for cosmology, where the composition of the universe, the initial conditions of structure formation, and reionization history are all quantified with only six parameters (under the assumption that the universe is flat).

Additional parameters

While the Λ CDM model is a good fit to the measurements of the CMB power spectra made so far, we are always re-evaluating this, and testing different models whenever we analyze new data. One popular method of coming up with such models is

²⁹There are disagreements over whether the difference in these measurements warrants the use of the word “tension.” All around, a rather tense issue.

by extending the Λ CDM model by adding one or more free parameters (which are otherwise fixed in Λ CDM).

No convincing evidence for extensions to the base Λ CDM model exists at this time. Some of the more widely studied additional parameters include:

Ω_K , curvature: The curvature of the universe, impacts its geometry, so it changes how a physical scale translates to an observed angular scale. Thus, changing the curvature of the universe will change the apparent scale of the CMB features, shifting the multipoles ℓ of the peaks. (The physical scale is not changed.) For example, the first peak corresponds to the scale of the sound horizon at recombination. This physical scale doesn't change, but the way it translates to an observed angular scale does. In Λ CDM, Ω_K is fixed at 0, and this is consistent with current measurements. *Planck* combined with BAO measurements yields $\Omega_K = 0.001 \pm 0.002$ ([Planck Collaboration VI, 2020](#)).

r , tensor to scalar ratio: This ratio measures the power in gravitational waves (tensor modes) in the early universe, compared to the power in scalar modes. Rather than using A_T , the analogue to A_S , it is more common to use the ratio r of the tensor and scalar primordial power spectra at some pivot point k (in Mpc^{-1}), and indicate this pivot value by writing r_k . A value of $r > 0$ would mean the existence of such gravitational waves, which would lead to primordial CMB B-mode polarization. In the base Λ CDM model, $r = 0$. Primordial B-modes have not been detected by any experiments so far, so only upper limits on r exist. Recent measurements from the BICEP/Keck collaboration yield $r_{0.05} < 0.036$ at 95% confidence ([BICEP/Keck Collaboration, 2021](#)).

N_{eff} , effective number of neutrino species: Changing the effective number of relativistic species (neutrinos) at decoupling affects the evolution of perturbations by altering the expansion of the universe. Neutrinos also stream relativistically out of

density fluctuations. An increase in N_{eff} therefore suppresses and shifts the peaks in the CMB power spectra (Bashinsky & Seljak, 2004; Ma & Bertschinger, 1995). In Λ CDM, $N_{\text{eff}} = 3.046$.³⁰ Note that N_{eff} receives contributions from any particles that have the same gravitational effect as relativistic neutrinos, which is the case for any (free-streaming) radiation. Hence, N_{eff} can be used to constrain such particles as well (e.g., Baumann et al., 2016, 2018).

$\sum m_{\nu}$, sum of the masses of the neutrino species: The CMB is used to constrain $\sum m_{\nu}$ through its effect on the growth of structure, primarily by measuring the gravitational lensing of the CMB. *Planck* (combined with BAO measurements) finds $\sum m_{\nu} < 0.12$ eV (95% confidence) (Planck Collaboration VI, 2020). The (indirect) constraints on $\sum m_{\nu}$ from CMB measurements are tighter than the (direct) constraints from measurements made in laboratories. A measurement of $\sum m_{\nu}$ using CMB data should be achievable in the near future (Allison et al., 2015). This is a great example of how measurements of the CMB allow us to place constraints on fundamental physics.

There are numerous other extension being studied as well, for instance some with parameters relating to the nature of dark energy or dark matter.

1.1.6 Foregrounds

The CMB power spectra we observe depend not only on the cosmological parameters described in the previous section, but also on everything else that exists in the universe between us and the surface of last scattering. At small scales, the microwave sky is dominated by astrophysical (extragalactic) foreground emission, such as dusty and radio extragalactic point sources. At larger scales, other than the CMB, the main

³⁰Given the three neutrinos in the Standard Model of particle physics, we may naively expect $N_{\text{eff}} = 3$. However, it is slightly larger due to a small amount of entropy from electron-positron annihilation being transferred to neutrinos (as opposed to only photons) (Gnedin & Gnedin, 1998; Mangano et al., 2005; de Salas & Pastor, 2016).

contribution is from diffuse Galactic foregrounds.

When doing cosmology, these foregrounds act as contaminants, which we want to characterize and either subtract or marginalize over. One way of separating emissions from foregrounds and the CMB is to exploit their different frequency spectra, for which it helps to have observations at a range of frequencies. Another strategy is to mask them. Perhaps most common is a hybrid approach, where we mask the brightest extragalactic sources and the plane of the galaxy, and model the residual foregrounds.

There are two types of extragalactic foregrounds: radio sources and dusty, star-forming galaxies. While they emit most strongly at radio and infrared wavelengths, respectively, these sources also have significant emission at microwave frequencies. The dusty galaxies form part of the cosmic microwave background (CIB) (Puget et al., 1996; Hauser et al., 1998; Gispert et al., 2000; Lagache et al., 2005). The radio point sources emit synchrotron radiation, and the dusty sources emit thermally, so both can be modeled as following a power law emission in flux for a sufficiently narrow frequency range. Both the dusty and radio sources are clustered, but the radio source clustering power is well below the Poisson level (e.g., Sharp et al., 2010; Hall et al., 2010; Dunkley et al., 2011; Shirokoff et al., 2011; Planck Collaboration XVIII, 2011; Hajian et al., 2012; Reichardt et al., 2012; Planck Collaboration XXX, 2014; Sievers et al., 2013; Mak et al., 2017). So the CIB is modeled as a sum of both Poisson and clustered components, whereas the radio sources are modeled as Poisson power (e.g., Choi et al., 2020). The dusty sources are not significantly polarized. The radio sources are, but only slightly (e.g., Hinshaw et al., 2007; Datta et al., 2018). Conveniently, the strongest polarized sources are also among the brightest objects in temperature, and they are usually masked in our analyses. So, compared to C_ℓ^{TT} , C_ℓ^{TE} and C_ℓ^{EE} suffer from very little contamination from foregrounds at high ℓ .

A summary of the main Galactic foregrounds of interest for the CMB, in both tem-

perature and polarization, is shown in Figure 1.6. While there are a greater number of foreground components in temperature, their signal is comparable to or lower than that of the CMB. At high Galactic latitudes, CMB temperature anisotropies dominate over Galactic foregrounds (e.g., [Bennett et al., 2003](#)). While there are fewer Galactic foreground components in polarization, their amplitude relative to the CMB is greater. The emission from Galactic foregrounds is greatest along the Galactic plane, particularly near the Galactic center, but there are no regions of sky that are completely free of this contamination.

At frequencies below ~ 50 GHz, the dominant Galactic foreground in polarization is **synchrotron radiation** ([Kogut et al., 2007](#)). This radiation arises from interactions between cosmic ray electrons and the Galactic magnetic field. The electrons experience an acceleration perpendicular to their velocity, resulting in a spiral trajectory in a plane perpendicular to the magnetic field.

In theory, the polarization fraction of synchrotron radiation can be as high as 75% in a uniform magnetic field. In practice, depolarization effects and non-uniform fields lead to measurements being lower. Still, observations have found synchrotron radiation to be polarized at levels up to 40–50% (e.g., [Page et al., 2007](#); [Kogut et al., 2007](#); [Miville-Deschénes et al., 2008](#); [Dunkley et al., 2009b](#); [Planck Collaboration XXV, 2016](#); [Vidal et al., 2015](#)). This leads to a significant polarization contribution, even at high Galactic latitudes.

For electrons with a power-law energy distribution $N(E) \propto E^{-p}$, the spectrum of the synchrotron emission can be approximated as

$$I_\nu \propto B^{\frac{p+1}{2}} \nu^{\beta_s} \quad (1.2)$$

where $\beta_s = -\frac{(p+3)}{2}$ and B is the magnetic field strength ([Rybicki & Lightman, 1979](#)).³¹

³¹This includes a few assumptions, such as the electrons being well mixed, and in a uniform

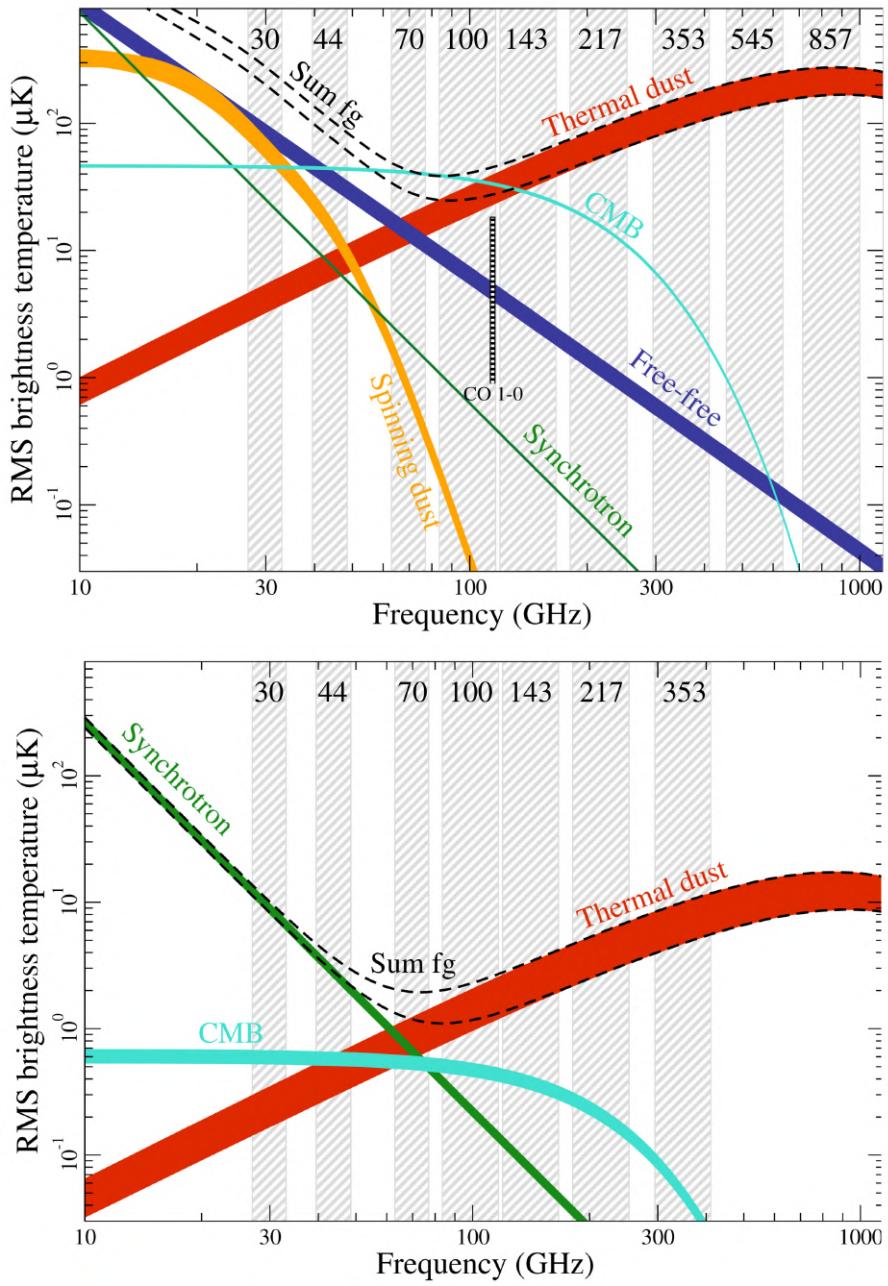


Figure 1.6: Brightness temperature rms as a function of frequency for different CMB foregrounds, in temperature (top) and polarization (bottom). For temperature (polarization), the components are smoothed to an angular resolution of 1° ($40'$), and the lower and upper edges of each line correspond to masks covering 81% (73%) and 93% (93%) of the sky, respectively. See the text for more information regarding the different foreground components. The gray shaded areas correspond to the *Planck* bands. Figures from [Planck Collaboration X \(2016\)](#), courtesy of ESA and The Planck Collaboration.

Typical values are $\beta_s \approx -2.5$ to -3.2 (e.g., [Bennett et al., 2013](#); [de Oliveira-Costa et al., 2008](#); [Dunkley et al., 2009b](#); [Fuskeland et al., 2014](#); [Kogut et al., 2007](#); [Krachmalnicoff et al., 2018](#); [Planck Collaboration X, 2016](#); [Vidal et al., 2015](#)).

There is some spatial variability to the spectrum due to its dependence on the energy distribution of the underlying electrons ([Lawson et al., 1987](#); [Platania et al., 1998](#), 2003; [Reich & Reich, 1988](#)), and there may be either steepening (due to spectral aging³²) or flattening (due to multiple components) of the spectrum ([Bennett et al., 2003](#); [Davies et al., 2006](#); [Page et al., 2007](#); [Planck Collaboration XI, 2016](#)). There is also, of course, spatial variation in the amplitude of the synchrotron emission due to variations in the magnetic field strength.

At a given frequency, synchrotron emission can also be approximately modelled as a power law, so

$$D_\ell^{\text{synch}} \propto \ell^{\alpha_s}, \quad (1.3)$$

where $\alpha_s \approx -0.7$ to -0.8 ([Planck Collaboration IV, 2020](#)).

At frequencies above ~ 70 GHz, the dominant Galactic foreground is **thermal dust emission** (see, e.g., [Ichiki, 2014](#), for a review). Non-spherical dust grains emit/absorb photons most efficiently along their longest axis, and this axis tends to be perpendicular to the local magnetic field (see, e.g., [Andersson, 2015](#); [Lazarian et al., 2015](#), for a review of alignment mechanisms). So, emission from dust grains ends up being polarized perpendicular to the Galactic magnetic field.

Dust can be modeled as a modified blackbody (sometimes referred to as a grey-body) in frequency,

$$I_\nu \propto B(\nu, T_d) \nu^{\beta_d}, \quad (1.4)$$

magnetic field.

³²As the electrons lose energy (and the higher-energy electrons lose energy more rapidly), their energy spectrum changes.

where B is the blackbody function, β_d is the index describing the frequency dependence of the dust emissivity, and T_d is the temperature of the dust. In general, typical values are $T_d \approx 19$ to 21 K and $\beta_d \approx 1$ to 2 (e.g., [Bennett et al., 2013](#); [Calzetti et al., 2000](#); [Dunkley et al., 2009b](#); [Planck Collaboration IV, 2020](#)).

The dust emission is often assumed to follow a power law as a function of multipole ℓ in D_ℓ space, so

$$D_\ell^{\text{dust}} \propto \ell^{\alpha_d}, \quad (1.5)$$

where $\alpha_d \approx -0.2$ to -0.6 ([BICEP/Keck Collaboration, 2021](#); [Planck Collaboration IV, 2020](#)).

Populations of dust may have different compositions and exist in different astrophysical environments, so their temperature and emissivity may vary spatially. And of course, different patches of sky have different levels of Galactic dust. The polarization fraction varies considerably across the sky, reaching up to $\sim 20\%$ (e.g., [Planck Collaboration XII, 2020](#)).

The three main atmospheric windows for observing the CMB are at approximately 45, 90, and 150 GHz. For the 50% of the sky with the least polarized foreground power, the ratios of the synchrotron to dust amplitudes at these frequencies are roughly 9, 0.3, and 0.02, respectively ([Choi & Page, 2015](#)). In general, the parameters for the foregrounds are fitted separately for temperature and for each polarization (e.g., [Page et al., 2007](#); [Gold et al., 2011](#); [Planck Collaboration Int. XXX, 2016](#); [Planck Collaboration X, 2016](#); [Planck Collaboration IV, 2020](#)).

Free-free emission, or thermal bremsstrahlung, is produced by scattering interactions between electrons and ions in the interstellar medium ([Rybicki & Lightman, 1979](#)). It has a very well-understood spectrum ([Dickinson et al., 2003](#); [Planck Collaboration Int. XV, 2014](#)). Free-free emission is not expected to be polarized due to the random

orientations of the scatterings. Magnetic fields could break this isotropy, but interstellar magnetic fields are too weak to produce significant polarization at microwave frequencies. Observations have indeed found free-free emission to be unpolarized, with an upper limit of a few percent (Macellari et al., 2011).

A less well understood foreground is anomalous microwave emission, or AME.³³ AME is not expected to be strongly polarized (Draine & Hensley, 2016; Hoang & Lazarian, 2016; Lazarian & Draine, 2000), and observations have placed upper limits on its polarization fraction of a few percent, falling rapidly for higher frequencies (Dickinson et al., 2011; Génova-Santos et al., 2015; Herman et al., 2022; Poidevin et al., 2019; Rubiño-Martín et al., 2012). (For a historical review of AME, see Leitch et al. (1997).) A leading candidate to explain AME is spinning dust radiation, whereby tiny, rapidly-spinning dust grains emit electric dipole radiation (Draine & Lazarian, 1998; Draine & Lazarian, 1998).³⁴ Another contribution to AME may come from larger, thermally-vibrating dust grains, which undergo fluctuations in their magnetization, resulting in magnetic dipole radiation (Draine & Lazarian, 1999). A recent review of AME research is presented in Dickinson et al. (2018).

Rotational transitions of carbon monoxide (CO) at 115, 230, and 345 GHz, were found to be a larger than anticipated source of contamination for *Planck* intensity maps (Planck HFI Core Team, 2011b). The distribution of CO line emission follows that of thermal dust emission, but with a larger concentration (so covering a smaller fraction of the sky), so for cosmological analysis, the *Planck* team masked out the corresponding areas of the sky with a Galactic mask.

Finally, for ground-based experiments, the Earth’s atmosphere also acts as a contaminant. This is mainly a concern for temperature. At millimeter wavelengths, the

³³AME is named as such because its origin is not yet fully understood.

³⁴The first paper to mention spinning dust grains as a possible source of emission (at high radio frequencies) is Erickson (1957), but detailed predictions of spinning dust as a source of AME were not made until several decades later.

atmosphere does not have significant linear polarization (Keating et al., 1998; Hanany & Rosenkranz, 2003). However, it does have circular polarization, which may be an issue if there is leakage of circular to linear polarization in the instrument (Spinelli et al., 2011; Petroff et al., 2020). Furthermore, ice crystal clouds in the upper troposphere scatter the CMB signal, leading to depolarization, and scatter back thermal radiation from the ground into linear polarization (Pietranera et al., 2007; Takakura et al., 2019).³⁵

In temperature, the emission from the atmosphere is much brighter than the CMB anisotropies. At microwave frequencies, the atmospheric emission is dominated by spectral lines of molecular oxygen at 60 and 119 GHz, and water lines at 22, 183, and 325 GHz, on top of the wings of water lines at higher frequencies. Water is more problematic than oxygen, since its inhomogeneous spatial distribution varies with time, leading to variations in the emission along the line of sight. The effect of the atmosphere is minimized by observing in so-called atmospheric windows (between the spectral lines), where the emission is lower. These main windows are: below 50 GHz, bands centered roughly around 95, 150, and 250 GHz, and also a band around 340 GHz when atmospheric conditions are particularly good.

Still, the atmosphere is one of the leading sources of noise for ground-based CMB experiments. It increases the optical loading on the detectors, amplifying their effective (photon) noise, while also introducing spatial and temporal correlations in the data. Atmospheric noise is mainly concentrated at low ℓ , so it can be suppressed by applying a high-pass filter to the data.

While modeling the atmosphere is complex, one feature that helps disentangle the CMB from atmospheric effects is that the CMB is fixed in the sky (in celestial coordinates), whereas the atmosphere is not. Since atmospheric fluctuations primarily

³⁵While this is not a concern for current experiments, for future ground-based CMB observations it will be necessary to take steps to mitigate the contamination due to clouds.

impact measurements of large angular scales (low ℓ), our ability to model these effects impacts the recovery of the CMB temperature anisotropies at these scales.

1.1.7 Experiments

Here we give a brief overview of past, current, and planned CMB experiments.

When observing the CMB from the ground, the atmosphere is one of the dominant sources of noise. There are a few strategies to deal with the atmosphere when designing a CMB experiment, which are reflected in the collection of experiments below. The advantages of ground-based telescopes are their relative ease and low cost of construction, and thus the ability to build large dishes (hence achieve higher resolution), so they are popular despite being susceptible to atmospheric contamination. To minimize the effects of the atmosphere, ground-based telescopes are placed at dry, high-elevation sites. The two most popular such sites are the Atacama Desert, in Chile, and the South Pole. The location in the Atacama Desert allows us to survey up to 70–80% of the sky, whereas the South Pole allows us to deeply observe a single small patch ($\sim 25\%$) of the sky. The atmosphere at the South Pole is also particularly stable.

While they are much more costly and logistically challenging than ground-based experiments, space-based telescopes avoid the issue of the atmosphere entirely, and can achieve full-sky coverage. However, weight restrictions limit the size of their dishes, and thus their resolution.

In terms of atmosphere, sky coverage, and resolution, an in-between option is balloon-borne telescopes, which are lifted above the lowest, densest part of the atmosphere. The downside is that these balloon-borne experiments collect data for weeks at a time at most, whereas satellites and ground-based telescopes typically observe for years.

The **Cosmic Background Explorer (COBE)** satellite made observations from 1989

to 1993 (Mather, 1993). It carried three instruments: a Diffuse Infrared Background Experiment (DIRBE) used to search for the cosmic *infrared* background (CIB) (Silverberg et al., 1993), a Differential Microwave Radiometer (DMR) used to map the CMB (Smoot et al., 1990), and a Far Infrared Absolute Spectrometer (FIRAS) used to measure the spectrum of the CMB (Mather et al., 1993). As mentioned earlier in the introduction, COBE measurements established that the CMB is a nearly perfect blackbody (Fixsen & Mather, 2002) and it mapped the anisotropies of the CMB, finding deviations in temperature of order 10^{-5} (Smoot et al., 1992).

The **Wilkinson Microwave Anisotropy Probe (WMAP)** observed from 2001 to 2010, in five frequency bands from 23 to 94 GHz (Bennett et al., 2003). Its measurements were key in establishing the Λ CDM model, which included determining that the primordial fluctuations were largely adiabatic (Spergel et al., 2003; Dunkley et al., 2009a; Bennett et al., 2013).

Planck was a European Space Agency mission which launched in 2009 and observed until 2013, in both temperature and polarization (Planck Collaboration I, 2020). It has produced the tightest constraints on cosmological parameters to date (Planck Collaboration VI, 2020). There were two instruments aboard *Planck*: the High Frequency Instrument, or HFI (Lamarre et al., 2010; Planck HFI Core Team, 2011a), and the Low Frequency Instrument, or LFI (Bersanelli et al., 2010; Mennella et al., 2011). Together, they collected data in nine frequency bands (30–857 GHz).

Maps from *COBE*, *WMAP*, and *Planck* displaying the improvements in resolution over the years are shown in Figure 1.7.

The **Atacama Cosmology Telescope (ACT)** was a 6 m telescope in Northern Chile which housed three generations of receivers: MBAC (Swetz et al., 2011), ACTPol (Thornton et al., 2016), and Advanced ACTPol (AdvACT) (Henderson et al., 2016; Ho et al., 2017; Choi et al., 2018; Li et al., 2018). Over the years, the detectors

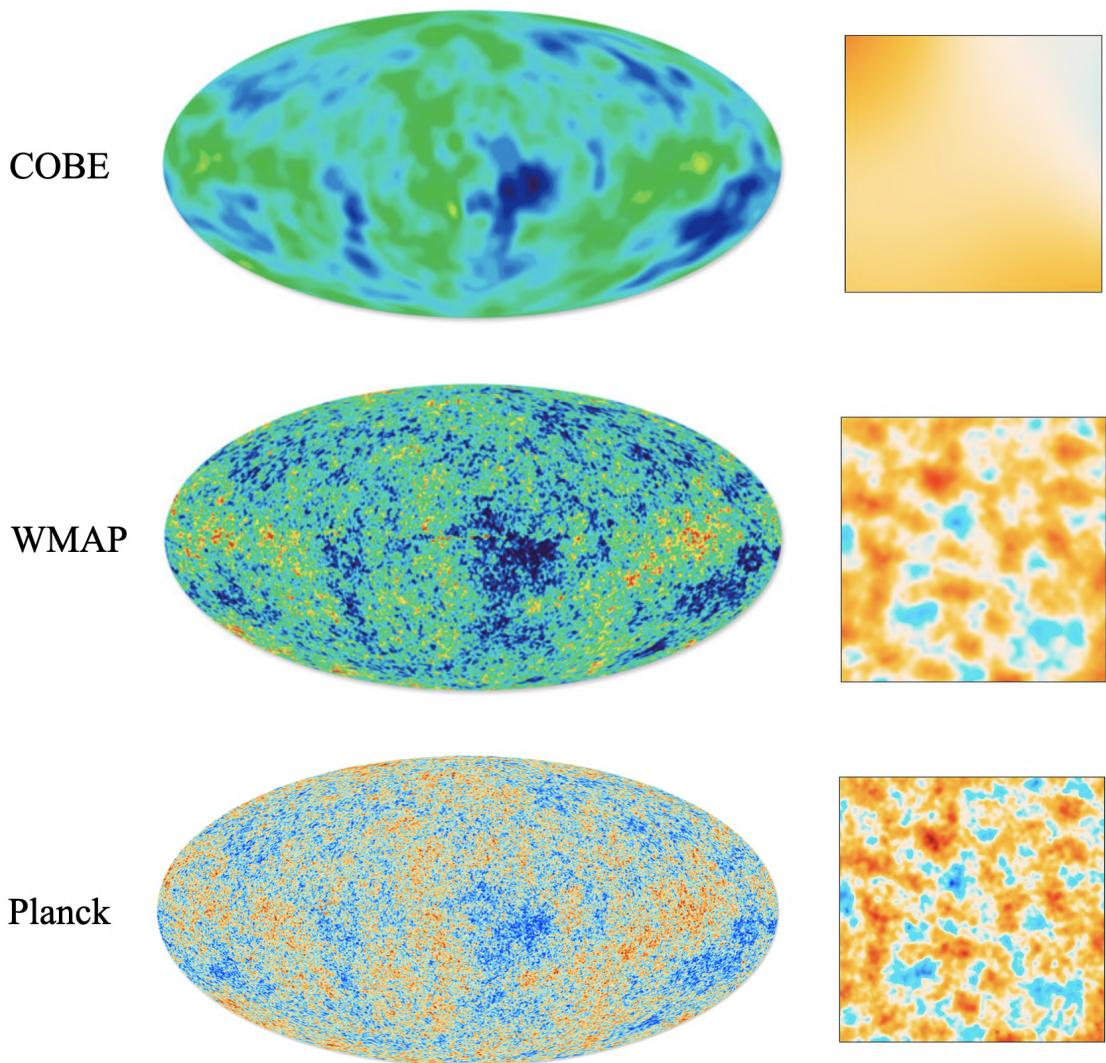


Figure 1.7: Comparison of *COBE*, *WMAP*, and *Planck* CMB temperature maps, featuring both full-sky maps (left) and a $10^\circ \times 10^\circ$ cutout (right). Image credit: NASA/COBE/DMR; NASA/WMAP Science Team; ESA and The Planck Collaboration.

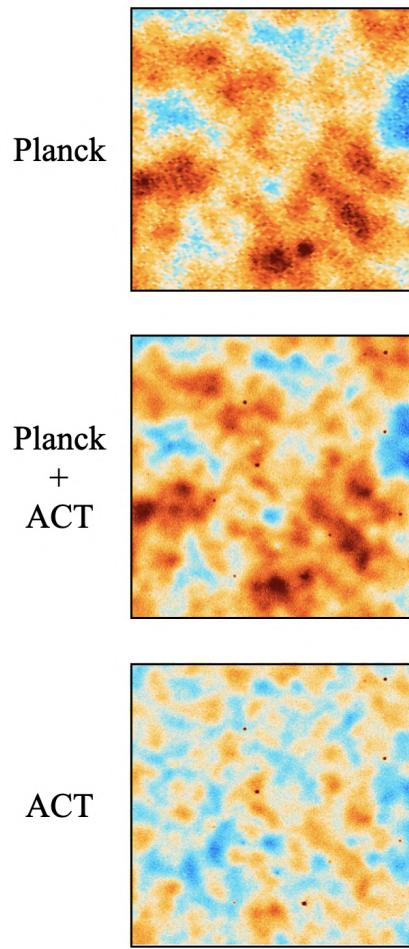


Figure 1.8: Comparison of *Planck*, *ACT+Planck*, and *ACT*-only maps at 150 GHz in a $3^\circ \times 3^\circ$ patch of sky, using data from ACT DR5. Based on a figure from [Naess et al. \(2020\)](#).

observed at a range of frequencies, from 30 to 277 GHz. From 2007 to 2022, ACT made high (arcminute) resolution maps of the CMB in both temperature and polarization, in order to measure the high- ℓ CMB anisotropies, as well as study CMB lensing and SZ clusters. ACT's several scientific accomplishments so far include making the first detection of seven acoustic peaks in the CMB power spectrum (Dunkley et al., 2011), the first detection of gravitational lensing in a CMB-only map (Das et al., 2011), the first detection of dark energy (Λ) from CMB data alone (when combined with WMAP) (Sherwin et al., 2011), the first detection of the kSZ effect (Hand et al., 2012), and the first measurement of the cross-correlation between CMB lensing and galaxy lensing (Hand et al., 2015). ACT and *Planck* maps are shown side-by-side for comparison in Figure 1.8. More details on ACT can be found in Section 1.2.

The **Cosmology Large Angular Scale Surveyor (CLASS)** is an array of four telescopes in the Atacama Desert, surveying a large fraction of the sky in bands centered at roughly 40, 90, 150, and 220 GHz (Essinger-Hileman et al., 2014). CLASS is designed to observe polarization at low ℓ (for which a large survey area is necessary), in order to measure, among other things, the reionization bump. It is the only ground-based experiment targeting these large angular scales at these frequencies. In addition to placing upper limits on r , this will allow for a cosmic-variance-limited measurement of the optical depth to the surface of last scattering, τ .

POLARBEAR (or POLARBEAR-1) is a 2.5 m telescope in the Atacama Desert of northern Chile³⁶ observing a small area of sky at 95 GHz and 150 GHz (Kermish et al., 2012). It measures solely polarization, and is focused on searching for B-mode polarization and studying gravitational lensing. POLARBEAR-2 is an upgraded receiver from POLARBEAR-1, and **Simons Array** is a project to further increase sensitivity by deploying two more telescopes with POLARBEAR-2 type receivers (Suzuki et al., 2016).

³⁶Not the South Pole, despite what its name might suggest.

BICEP and Keck Array (BK) are a series of five experiments at the South Pole: BICEP1 (or just BICEP) (Keating et al., 2003), BICEP2 (Ogburn et al., 2010), BICEP3 (Ahmed et al., 2014), Keck Array (Staniszewski et al., 2012), and BICEP Array (Hui et al., 2018). They aim to measure the polarization of the CMB (at ~ 95 and 150 GHz), in particular, B-modes.

The **South Pole Telescope (SPT)** is a 10-meter telescope located at the South Pole (Ruhl et al., 2004; Carlstrom et al., 2011). As an experiment, it has consisted of three stages: SPT-SZ (Chang et al., 2009), SPTpol (McMahon et al., 2009), and SPT-3G (Anderson et al., 2018). It makes high-resolution measurements of the CMB in temperature and polarization (at ~ 95 , 150, and 220 GHz), with the goal of measuring high- ℓ CMB fluctuations (including B-modes) as well as detecting galaxy clusters through the SZ effect.

SPIDER (Gualtieri et al., 2018) is a balloon-borne experiment designed to measure (at roughly 95, 150, and 280 GHz) the polarization of the CMB at large angular scales, targeting primordial B-modes (Ade et al., 2022).

The **Simons Observatory (SO)**, is a ground-based CMB experiment currently under construction in the Atacama Desert in Chile, with observations scheduled to begin in 2023–2024 (Galitzki et al., 2018; Simons Observatory Collaboration, 2019). It will consist of one large aperture (6 m) telescope (LAT) and six small aperture (0.5 m) telescopes (SATs). It will include 6 frequency bands from 27 to 280 GHz. The LAT will be used to measure the CMB in temperature and polarization at small angular scales, whereas the SATs will be used for larger-scale polarization measurements. The LAT will map $\sim 40\%$ of the sky at arcminute angular resolution, with an expected white noise level of $6 \mu\text{K-arcmin}$. The resulting high-resolution maps will help constrain cosmological parameters that impact the damping tail of the CMB, and map matter using gravitational lensing as well as the Sunyaev-Zel'dovich effect (both ther-

mal and kinetic). Meanwhile, the SATs will map $\sim 10\%$ of the sky to a white noise level of $2 \mu\text{K-arcmin}$, targeting large angular scales in order to measure the primordial tensor-to-scalar ratio, r .

LiteBIRD is a satellite mission which is planned to launch by the end of the decade (Hazumi et al., 2020). It will observe the polarized CMB at large angular scales in 15 frequency bands from 34 to 448 GHz, in the quest for the ever-elusive inflationary B-modes.

Further in the future, **CMB-S4** will be several times more sensitive than any existing experiment, and is expected to include telescopes both in Chile and at the South Pole, in order to exploit the unique features of both of these sites. It may consist of 18 SATs and one LAT at the South Pole, and two LATs in Chile.

Recent measurements of the CMB are shown in Figure 1.9.

The list presented here is far from exhaustive.³⁷ Over the last few decades, many more experiments have helped lay down the foundations for our present work. Today, there are several experiments (in addition to those mentioned here) working on studying the universe through the CMB. Finally, there are a variety of planned or proposed experiments for the coming years. The sampling here is merely intended to provide a quick glance at the state of affairs.

³⁷ Apologies to any readers whose favorite experiment isn't listed.

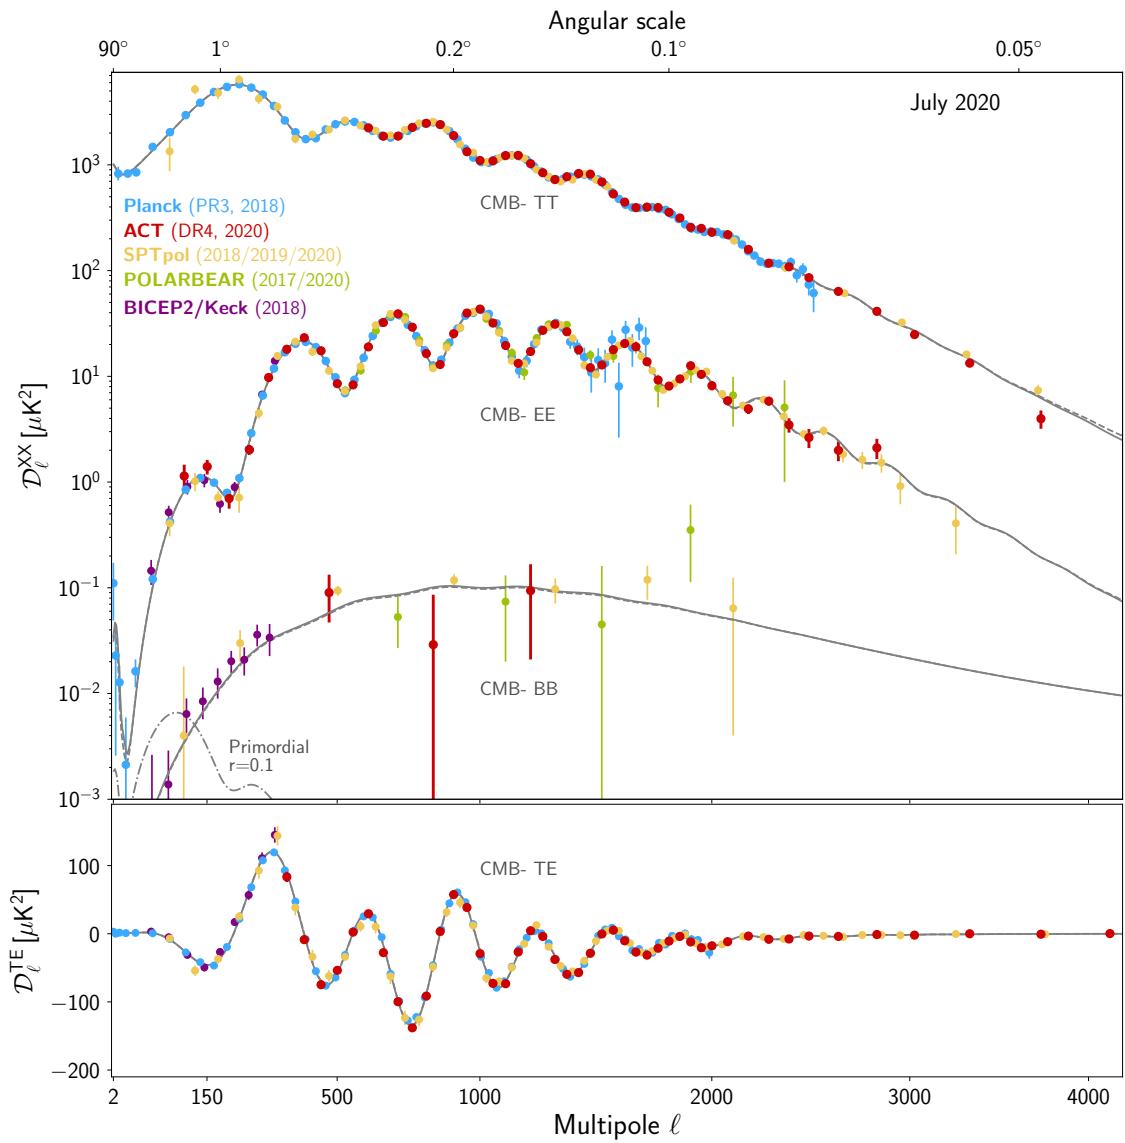


Figure 1.9: (Caption on next page.)

Figure 1.9: (Previous page.) Recent measurements of the CMB temperature anisotropy and polarization. The two models, the thin nearly overlapping gray lines, are from *Planck* (dashed line) and from ACT plus WMAP (Aiola et al. (2020), solid line). The primordial BB signal with $r = 0.1$ is also shown with the dot-dashed line. For *Planck*, we show the 2018 results (Planck Collaboration V, 2020). For ACT, we show the DR4 spectra from Aiola et al. (2020); Choi et al. (2020), including some preliminary low- ℓ EE results that were not used for cosmological analysis. For SPT we show Henning et al. (2018) for 150 GHz TT $\ell < 2000$, TE and EE, and Sayre et al. (2020) for BB. For $\ell > 2000$ we show the SPT spectrum from George et al. (2015) which has been corrected for point source emission. It is visually indistinguishable from the more precise but uncorrected spectrum in Reichardt et al. (2021). (It should be noted that with the publication of Dutcher et al. (2021) and Balkenhol et al. (2022), more recent/sensitive SPT data exist for the TT, TE, and EE spectra.) For Polarbear/Simons Array, we show EE from POLARBEAR Collaboration (2020) and BB from pipeline A in POLARBEAR Collaboration (2017). For BICEP2/Keck, we use Keck Array, BICEP2 Collaborations (2018). (It should be noted that with the recent publication of BICEP/Keck Collaboration (2021), these results are no longer current. The uncertainties on the BB spectra are now noticeably smaller than those shown here. See, for instance, Figure 13 of BICEP/Keck Collaboration (2021).) All error bars are one sigma, and points with no lower bound in TT and EE have been dropped at high ℓ . There is much more to each data set than is plotted here, for example, additional frequencies. Figure from Choi et al. (2020).

1.2 The Atacama Cosmology Telescope

1.2.1 The Instrument

The Atacama Cosmology Telescope (ACT), shown in Figure 1.10, was³⁸ a 6 m off-axis Gregorian telescope designed for millimeter wavelength observations of the CMB at arcminute resolution. Located at an altitude of 5190 m in the Atacama Desert of northern Chile, it had access to over half the sky. The high altitude and low precipitable water vapor (PWV) at this location provide excellent millimeter and submillimeter atmospheric transparency.

ACT scanned in azimuth at a few fixed elevations as the sky drifted through the field of view, in order to minimize systematic effects which arise with altitude variations. The scans were performed both in the east (where the sky is rising) and the west (where the sky is setting) so the sky could be observed at different parallactic angles. This crosslinking helps minimize systematic effects due to the scan direction in the mapmaking process and improves our ability to measure CMB modes parallel to the scan direction.

There have been three generations of ACT receivers: the millimeter bolometric array camera, MBAC, (Swetz et al., 2011) which observed at 150, 220, and 277 GHz, ACT’s first polarization-sensitive receiver, ACTPol (Thornton et al., 2016) which observed at 98 and 150 GHz, and the Advanced ACTPol (AdvACT) receiver (Henderson et al., 2016; Ho et al., 2017; Choi et al., 2018; Li et al., 2018) which in its final state was configured with detector arrays at 30, 40, 98, 150, and 220 GHz.³⁹ In addition to

³⁸Following 15 years of prolific observations, ACT has recently been retired. Decommissioning began on August 26th, 2022.

³⁹The naming convention for the ACT detector arrays may seem confusing, as it varies from paper to paper. The difficulty arises from the detector bandpasses. Each detector array is sensitive to a range of frequencies, not just one, and determining some effective central frequency is not straightforward. In this thesis, which focuses on ACTPol and AdvACT data, we will primarily deal with data from the 98 GHz and 150 GHz detectors. It should be noted that in different papers, they may instead be referred to as being at 90 or 97 GHz (instead of 98 GHz) and 148 or 149 GHz (instead of 150 GHz).

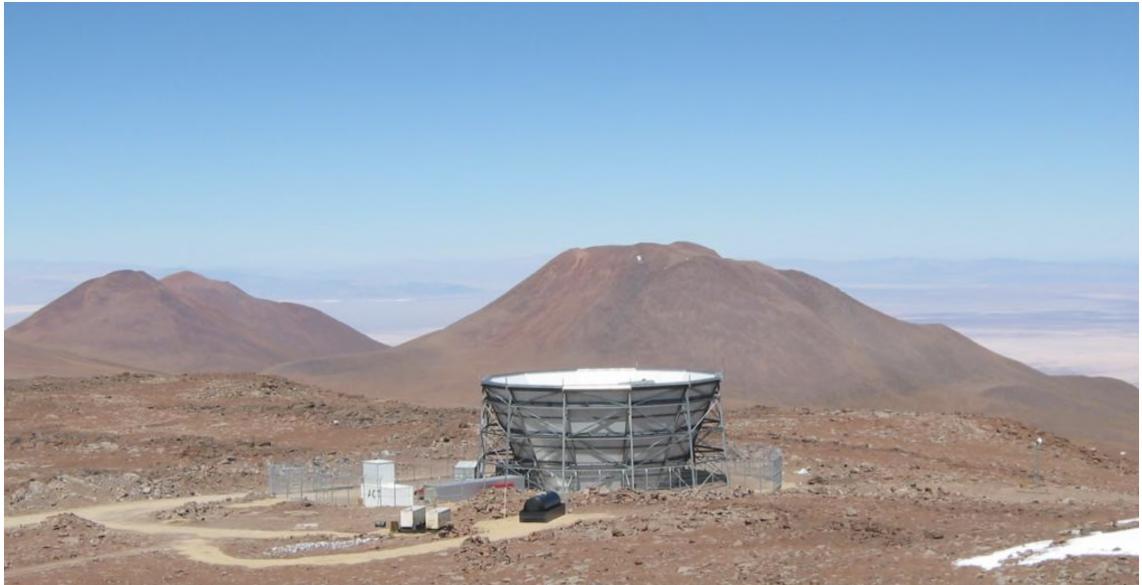


Figure 1.10: Photograph of ACT inside its stationary conical ground screen. For scale, the distance from the ground to the top of the ground screen is roughly 13 m (Swetz et al., 2011). Photo credit: Mark Devlin.

an expanded frequency coverage, AdvACT improved upon ACTPol by increasing the number of detectors, hence enhancing the instrument’s sensitivity.

The succession of polarization-sensitive detector arrays (PAs) throughout the years is illustrated in Figure 1.11. There are (at most) three separate detector arrays in the telescope at one time, each in an optics tube with its own set of filters and lenses. Beginning in 2015, all new arrays added were dichroic; that is, they simultaneously measured two polarizations.

1.2.2 The Data

The ACT data releases⁴⁰ are as follows:

DR1 consisted of MBAC data from 2008, covering a southern region (centered on RA = 60° Dec = -52.7°) at 148 GHz (e.g., Dunkley et al., 2011; Dünner et al., 2013).

⁴⁰All data are released through NASA’s LAMBDA site. <https://lambda.gsfc.nasa.gov/product/act/>

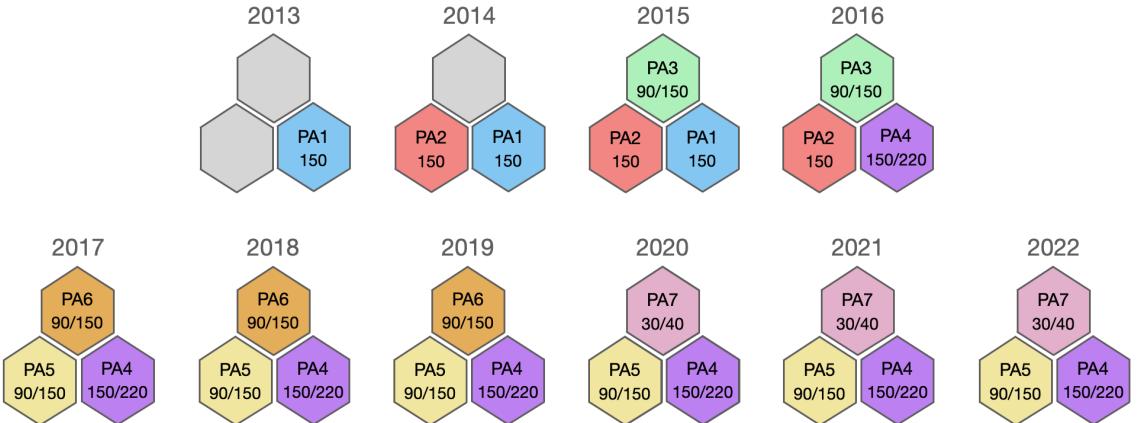


Figure 1.11: ACT polarization-sensitive detector arrays in use during each year of observation. The first detectors operated at 150 GHz. Since then, dichroic arrays have been added, contributing data in frequency bands centered roughly at 30, 40, 90, and 220 GHz. Not shown are the first generation (MBAC) arrays, since they were not sensitive to polarization.

DR2 was made up of data from 2008–2010, covering the south and the “SDSS Stripe 82” equatorial region, and included data at 217 GHz and 277 GHz in addition to 148 GHz (e.g., [Das et al., 2011](#); [Sievers et al., 2013](#); [Gralla et al., 2020](#)).

DR3 was based on ACTPol, so it included both temperature and polarization data from 2013–2014, and covered a number of regions near the equator at 150 GHz (e.g., [Naess et al., 2014](#); [Louis et al., 2017](#)).

DR4 was also based on ACTPol, covering seven regions of sky spanning roughly \sim 18,000 deg 2 , including data at 98 GHz as well as 150 GHz from 2013–2016 ([Choi et al., 2020](#); [Aiola et al., 2020](#); [Darwish et al., 2020](#); [Madhavacheril et al., 2020](#); [Namikawa et al., 2020](#); [Han et al., 2021](#); [Mallaby-Kay et al., 2021](#)). DR4 included DR3 as a subset.

DR5 included previously released MBAC and ACTPol data, as well as some preliminary data from AdvACT from 2017–2018, including some at 230 GHz ([Naess et al., 2020](#)). It consisted of high-resolution maps, containing the largest single collection of Sunyaev Zel’dovich (SZ) clusters and millimeter wave sources to date. However, due

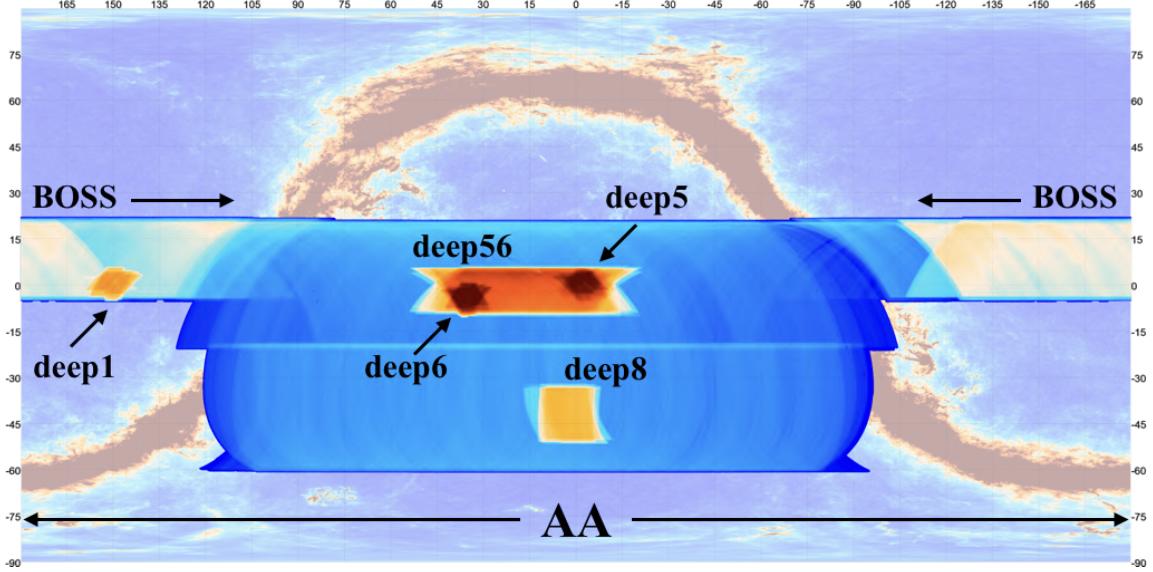


Figure 1.12: Cumulative coverage map of the ACT observations from 2013 to 2016. The observation patches have been labelled. The background is the *Planck* 353 GHz intensity map. The x -axis and y -axis represent right ascension and declination in degrees. See the text for information on the abbreviations used to describe the data in this thesis.

to their preliminary nature, these maps were not suitable for precision cosmology.

DR6 is the next ACT data release, currently in the works. It will be mentioned at various points throughout the chapters. As with all the pre-DR5 releases, DR6 will include power spectra and inferred constraints on cosmological parameters.

The main patches of sky observed by ACTPol and AdvACT at night are shown in Figure 1.12. In the early seasons small patches were observed (for example deep1, deep5, and deep6 in 2013 as part of DR3) and the patches increased in size until we began observing the “Advanced ACT” patch in 2016.⁴¹ The patch names are frequently abbreviated, so deep1, deep5, deep6, deep8, deep56, become D1, D5, D6, D8, D56, BOSS becomes BN, and Advanced ACT becomes AA. Each year of observation is referred to as a season, and we use the abbreviation S13 (or s13) for 2013, and so on.

⁴¹The scanning strategy for the daytime made the transition from deep to wide patches later, in mid-2019 (see Section 5.1.1).

1.2.3 Pipeline Overview

The data ACT collects is initially in the form of time-ordered data. That is, for each detector within an array, we have a stream of data, calibrated into mK, versus time. We split up this data into sections, each of roughly 10 minutes, and refer to each such section of data for each detector array as a TOD. The steps taken to transform these time-ordered data into constraints on cosmological parameters are illustrated in Figure 1.13, and briefly described below.

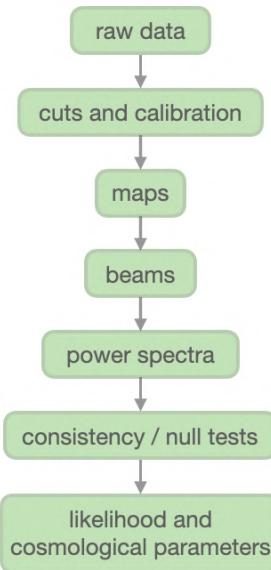


Figure 1.13: Rough outline of the ACT data analysis pipeline that transforms raw data from the telescope into constraints on cosmological parameters.

The first step taken is to assess the quality of each TOD in order to make cuts. The data samples for one detector are cut if, in general, they differ too much from the mean for the entire array; for example if there is excessive noise in the data due to electrical issues, or if the detector's gain drifts with time. The data for an entire TOD may be discarded due to issues such as bad weather, too many cut detectors within an array, problems with the telescope itself (e.g. the cryogenics), etc. The data cuts are described in more detail in [Dünner et al. \(2013\)](#).

Then, the calibration of the detectors is performed using a multitude of techniques: the relative calibration of detectors is obtained using their response to the atmosphere; the absolute calibration of the instrument is determined from planet measurements; the detector time constants are measured using planet observations and bias steps; and the polarization angles of the detectors are computed using an optical model of the telescope and lenses (Grace et al., 2014; Thornton et al., 2016).

Once data quality cuts have been made and the calibration procedures described above have been performed, the time-ordered data are projected into maps of the sky. This is the topic of Chapter 2. It is the most computationally intensive step of the analysis.

Before computing the power spectra of these maps, we must determine the telescope's beams. The beams determine how the instrument responds to different scales on the sky, and quantifying the beams (and their uncertainties) is a must in order to make precise measurements of the CMB power spectrum. The beams are covered in detail in Chapter 3.

We can then compute the power spectra, as explained in Chapter 4. Since we are making precision measurements, our data are sensitive to seemingly small changes in analysis choices, so care must be taken.

While it is tempting to immediately compute constraints on Λ CDM parameters (or another model of choice) from the power spectra, we must first check the consistency of our data. We run a series of consistency tests by comparing different power spectra, to ensure, for instance, that the spectra from two different detector arrays are in agreement. We also run map-based null tests, where we check for various systematic effects by splitting the TODs in two groups based on some parameter, for example, the atmospheric conditions, then map the two groups of TODs separately and compare the resulting spectra. These map-based tests are described in Chapter 2.

Finally, once all consistency tests have been passed, we compared our computed power spectra to the predicted spectra for different cosmological models, and thus infer constraints on the parameters.

The steps outlined here are, of course, an oversimplification, and leave out crucial components, such as the simulations used to quantify the noise in the data. But this list can serve as a useful framework upon which to build.

Chapter 2

Maps

Part of this chapter is adapted from Choi et al. (2020). My contributions to that work include splitting the data and making the maps for the map-based null tests, and writing the corresponding section of the paper. This chapter also includes some material from the associated publication Aiola et al. (2020) .

2.1 Mapmaking

2.1.1 Overview

To convert raw data into maps, the data are first divided into ten-minute time-ordered data (TOD) files and pre-processing is performed on each individual TOD.

For each set of maps (so for each season, patch,¹ detector array, and frequency) the TODs are divided into two to four subsets (depending on the patch) in a way that maximizes the crosslinking among the TODs in each subset. Each of these subsets is referred to as a split. We then map each split separately, so we end up with two to four independent maps of roughly identical coverage and depth. We use the fact

¹The patches of sky observed by ACTPol and AdvACT are shown in Figure 1.12.

that the splits have independent instrument and atmospheric noise to avoid noise bias in our computed power spectra. The maps are made with a Plate Carrée (CAR) pixelization,² with pixels of size 0.5 arcmin.

2.1.2 Mapmaking Procedure

Maps made for the primary ACT analysis are made using the mapmaking code `enki`³, written by Sigurd Naess. To make maps, we solve for the best-fit sky given the measured data and its noise. We model the time-ordered data \mathbf{d} as

$$\mathbf{d} = A\mathbf{m} + \mathbf{n} , \quad (2.1)$$

where A is the pointing matrix (it projects from map domain to time domain), \mathbf{m} is a vector whose components are the sky map pixels (for which we wish to solve), and \mathbf{n} is the noise contribution to the data, usually taken to be normally distributed with covariance matrix $N = \langle \mathbf{n}\mathbf{n}^T \rangle$.

In the simplest case, each data point is mapped to a single pixel in the map, so each row of A (corresponding to a single data sample) is composed entirely of zeros except for a single 1 in the column corresponding to that particular pixel. In reality, for our mapmaking, the model \mathbf{m} contains contributions from atmospheric noise, correlated signals in the data, and cut data samples, so the pointing matrix is not as simple. Polarization is handled by including each detector's response to the I , Q , and U Stokes parameters in the pointing matrix.

To solve for the map, we need to find the model \mathbf{m} which maximizes the likelihood

²Pre-DR4 ACT maps were made in the cylindrical equal-area (CEA) projection, but since the survey now covers a large range in declination, this would result in oversampling near the equator. The CAR projection avoids this by having pixels equi-spaced in longitude in each latitude ring, so there are the same number of pixels per ring.

³GitHub repository: <https://github.com/amaurea/enlib>

function

$$\mathcal{L} \propto \exp \left[-\frac{1}{2}(\mathbf{d} - A\mathbf{m})^T N^{-1}(\mathbf{d} - A\mathbf{m}) \right]. \quad (2.2)$$

The linear least-squares solution to this (the maximum likelihood mapmaking equation) is

$$A^T N^{-1} A \mathbf{m} = A^T N^{-1} \mathbf{d}. \quad (2.3)$$

with solution

$$\hat{\mathbf{m}} = (A^T N^{-1} A)^{-1} A^T N^{-1} \mathbf{d}. \quad (2.4)$$

Inserting the model for \mathbf{d} from Equation 2.1 into Equation 2.4 and taking $\langle \mathbf{n} \rangle = 0$ (where $\langle \rangle$ denotes the average over an infinite number of realizations), we see that $\langle \hat{\mathbf{m}} \rangle = \mathbf{m}$, that is, the solution is unbiased.

The matrix $A^T N^{-1} A$ is too big to be inverted directly, so instead we iteratively solve for \mathbf{m} using a preconditioned conjugate gradient algorithm (Shewchuk, 1994). Preconditioning involves using an approximate inverse of $A^T N^{-1} A$, which we denote P , to speed up convergence by solving the better-conditioned system

$$P A^T N^{-1} A \mathbf{m} = P A^T N^{-1} \mathbf{d}. \quad (2.5)$$

The standard mapmaking preconditioner is created using the inverse of the number of observations in each pixel.

An example of an ACT map created using this code is shown in Figure 2.1. With the conjugate gradient technique, the low signal-to-noise modes at low multipoles ($\ell \lesssim 50$) are degenerate or almost degenerate and never converge. We choose a sufficient number of steps of the conjugate gradient algorithm for the higher- ℓ modes to converge, which in this case is roughly 300. This mapmaking is the most computationally intensive step in the ACT analysis pipeline.

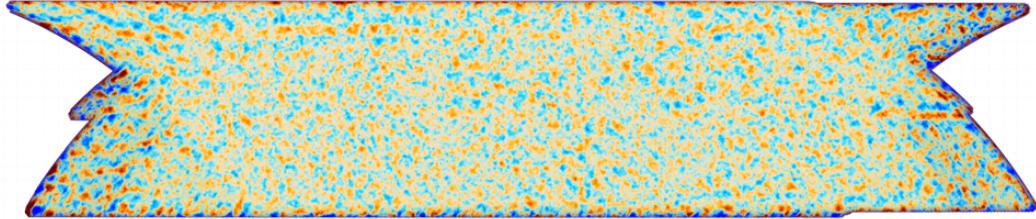


Figure 2.1: Intensity map of the D56 patch made using 2014 PA2 data. This corresponds to roughly 700 deg^2 of sky. The low- ℓ modes have been filtered out here to suppress the atmosphere, for visualization purposes.

The evolution of a map with an increasing number of conjugate gradient steps is shown in Figure 2.2. The last map in this figure, the result of 300 steps, is not quite the final product. As described in Section 2.1.3, we run through 3 passes of the mapmaker, each with 300 conjugate gradient steps. During the second and third passes of the mapmaker, the low- ℓ modes continue to converge. The result is shown in Figure 2.3.

2.1.3 Noise Estimate

The noise covariance matrix is modelled as being stationary in each 10 minute chunk of data (each TOD) and is measured from the data. It includes detector-correlated noise, including atmospheric noise.

Since the noise model is applied to the same data it is measured from, signal-induced bias may arise in the maps. When two detectors scan across the same point in the sky, the relative weights we assign to their measurements are determined based on the internal variance of their time streams. In general, due to noise in the detectors, one of them will observe a lower flux than the other. Since the detector weights are determined from the variance in the data, the detector which measures the lower flux will usually have a lower variance and thus will be given more weight when making a map. More heavily weighted low-flux measurements lead to the flux being biased low. This bias is already small since the CMB is highly subdominant to the noise in

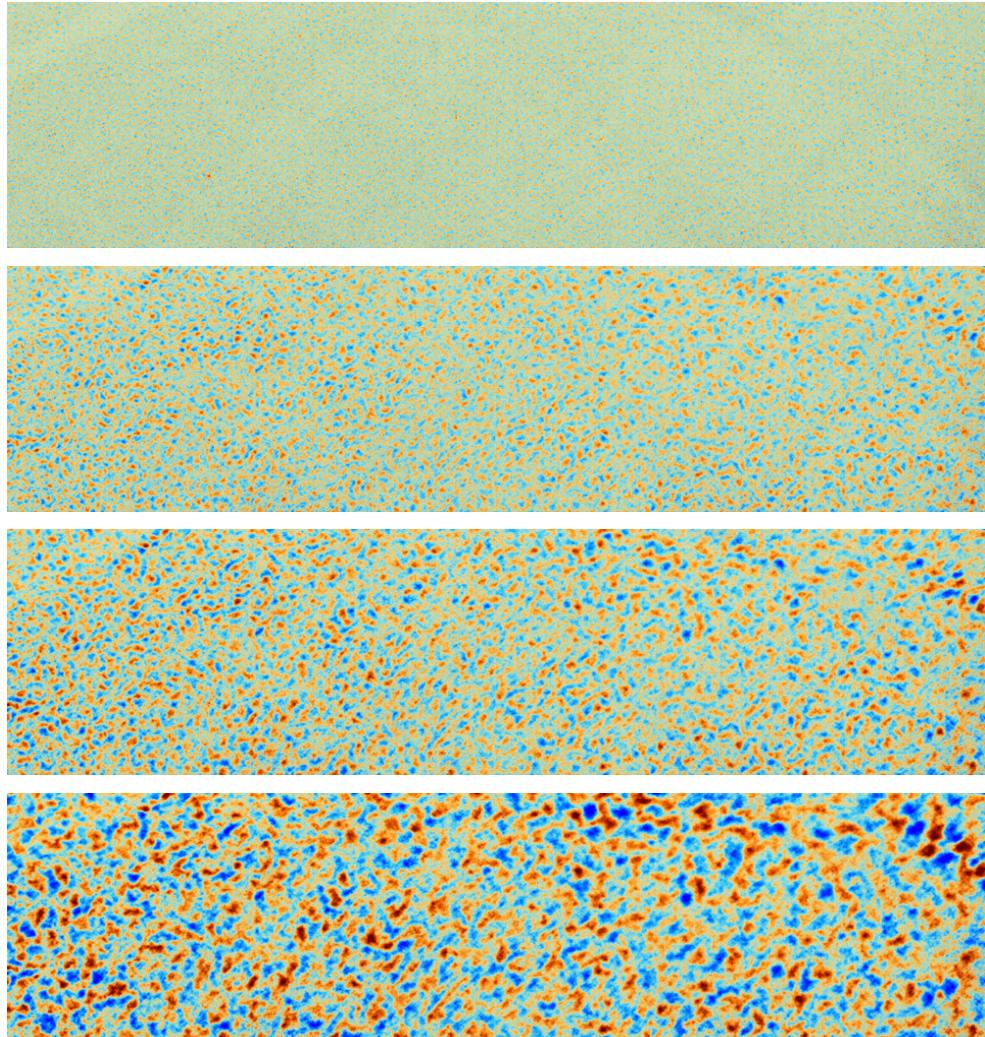


Figure 2.2: Example of the convergence of modes in a map with the conjugate gradient algorithm. These are $12^\circ \times 48^\circ$ cutouts of a map after 10, 50, 100, and 300 conjugate gradient steps, from top to bottom. The color scale spans roughly the range of $\pm 500 \mu\text{K}$. (These maps have not been corrected for the final power-spectrum-based calibration described in Section 4.2.7. Typically, the preliminary calibration at this stage is correct to within $\sim 25\%$.)

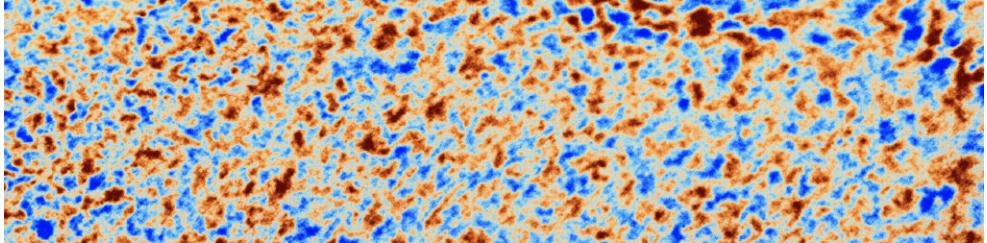


Figure 2.3: Example of some larger-scale modes having converged after 3 passes of the mapmaker, each with 300 conjugate gradient steps. This is the result of running the final map cutout from Figure 2.2 through two more passes of the mapmaker. The color scale spans roughly the range of $\pm 500 \mu\text{K}$. (This map has not been corrected for the final power-spectrum-based calibration described in Section 4.2.7. Typically, the preliminary calibration at this stage is correct to within $\sim 25\%$.)

ACT time streams, but according to simulations it is not negligible. It can be made even smaller if an estimate of the sky signal is first subtracted before the noise is estimated.

To this end, we run through 3 passes of the mapmaker. Initial maps are made for which the noise is measured from the time streams. Then these maps are subtracted from the time streams before the noise is estimated again. This improved noise estimate is then used to make new maps. Finally, these improved maps are subtracted from the time streams before the noise is estimated one final time. And this twice-improved noise estimate is used to make the final maps. At this point, the bias is undetectable. These 3-pass maps are the ones used for analyses.

2.1.4 Point Sources

Maximum likelihood mapmaking is optimal and unbiased in making CMB maps, as long as our model of the data is correct. However, in practice there are model errors, including sub-pixel structure (which leads to errors due to the sky being smooth and us modeling it as pixelated) and instrumental issues such as time-variable gain and pointing errors. In this case, the solution is biased and the error in each pixel leaks along the scanning direction by one noise correlation length. These model errors are

several orders of magnitude fainter than the signals which source them. For a smooth, low-contrast signal such as the CMB, the level of contamination is usually negligible.⁴ Non-local model errors are thus mainly a concern with a strong, compact signal such as a point source.

The most significant sources of such leakage in ACT maps are bright point sources. Due to the nature of our scanning patterns, the leakage results in an \times around the sources in the maps (with the “legs” of the \times aligned with the two scanning directions, from scanning in the east and in the west). To avoid \times es around the bright sources in the maps without cutting them, the ~ 70 brightest sources in the map receive special treatment in the mapmaker.

Pre-DR4, the sources were subtracted from the data before mapmaking and then added back to the final map. This requires knowledge of the location and amplitude of the bright point sources and does not reduce artifacts from variability. For DR4, a new method for handling the point sources has been implemented.

In theory, model errors could be eradicated by having at least one degree of freedom per sample in the map. However, in this case there would be no averaging down of the noise in the data, which would limit the usefulness of any resulting map. But there is no problem with performing such oversampling only in problematic regions such as around strong point sources, as long as these regions are small compared to the noise correlation length. In this case, the model for the time-ordered data becomes

$$\mathbf{d} = A\mathbf{m} + G\mathbf{s} + \mathbf{n} = [A \ G] [\mathbf{m} \ \mathbf{s}]^T + \mathbf{n}, \quad (2.6)$$

where \mathbf{s} are the per-sample degrees of freedom for the bright point sources, G is the

⁴While working on DR6, we have discovered that some model errors (which we are still working to understand) result in a gradual loss of power in our maps below $\ell \approx 800 - 1200$, depending on the frequency. The transfer function due to these model errors has been estimated, and found to be sufficiently close to unity for our ℓ range of greatest interest. At the moment we do not correct for the transfer function, but this is still an area of active work.

matrix mapping to the relevant samples for bright sources, and the square brackets are used to denote the concatenation of two matrices. With this notation, the equation above is of the same form as the expression $\mathbf{d} = \mathbf{A}\mathbf{m} + \mathbf{n}$ from Section 2.1.2, and so we may use the maximum likelihood mapmaking equation with the relevant substitutions. The terms \mathbf{m} and \mathbf{s} cannot be solved for separately, but that is not a problem as the total map into which the sources are projected using a white noise model, $\hat{\mathbf{m}}_{tot} = \hat{\mathbf{m}} + M_{PW}\hat{\mathbf{s}}$, is well-defined.

The effect of the source sub-sampling method can be seen in Figure 2.4. More information about different methods to avoid model errors due to point sources in maximum likelihood maps can be found in [Naess \(2019\)](#).

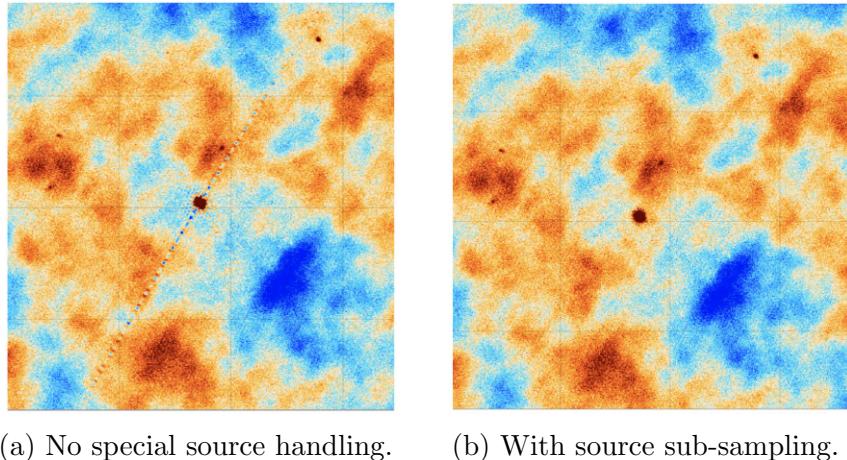


Figure 2.4: Point source observed in the 2014 PA2 maps of the BN patch (made using both daytime and nighttime data). Each map here is roughly $3.5^\circ \times 3.5^\circ$. The leakage around the point source in the map on the left-hand side resembles a \swarrow more than an \times because in this case one of the two scanning directions dominates.

2.1.5 Pointing Correction

For each season and array, a first preliminary model for the telescope pointing is constrained using observations of planets made at night. Then a correction is applied to account for temporal variations in the pointing.

First, a preliminary map is made using the nominal pointing estimate from planet observations. Then the bright point sources are located in the maps and their true positions are found by matching them to known catalogs. Assuming the beam is constant, a joint fit is then performed to the brightest sources in the time domain for a single overall pointing offset per TOD.

Pre-DR4, when a TOD did not contain a good source fit (with low enough uncertainty), the nearest neighbor with a good fit was used. Our new correction uses per-TOD pointing measurements to build a continuous model as a function of time, based on Gaussian processes. This new approach results in slightly cleaner maps of the nighttime data, but more significantly, it drastically improves the quality of the daytime maps, as shown in Section 5.3.1.

2.1.6 Pickup maps

ACT has two levels of ground screens: a screen which is fixed to the telescope and scans with it, and a 13-m high fixed ground screen surrounding the entire system (Swetz et al., 2011). Nevertheless, ground pickup still appears in the data. This ground signal can be disentangled from the sky because it has a fixed shape versus azimuth during a scan and does not rotate with parallactic angle, unlike the sky.

Pre-DR4, to remove the ground pickup, an azimuthal filter was applied to the time-ordered data. Instead, we now subtract a static pickup template per scanning pattern. This template is built by making a “pickup map,” a map of the sky in ground-centered coordinates, for all the TODs with the same scanning pattern. An example of such a map is shown in Figure 2.5. The ground signal should be common among all the TODs while the sky signal changes. This results in less large-scale subtraction than the azimuthal filter and a mapping transfer function which is closer to unity.⁵

⁵Unfortunately, despite this procedure, we still found the need to filter out the low- ℓ modes from the DR4 maps, as explained in Section 4.2.4. Due to this, and wanting to avoid introducing transfer

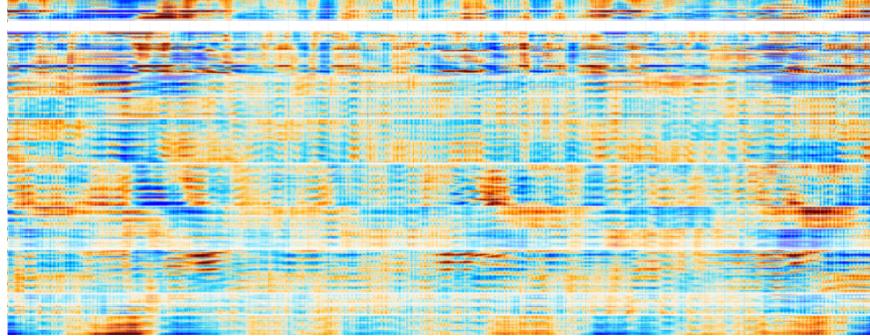


Figure 2.5: Example of a pickup map created for 2014 D56 PA2 nighttime data. The selected area here measures 43° wide by 30° tall. The color scale spans roughly the range of $\pm 500 \mu\text{K}$. (This map has not been corrected for the final power-spectrum-based calibration described in Section 4.2.7. Typically, the preliminary calibration at this stage is correct to within $\sim 25\%$.)

2.1.7 Transfer Function

A transfer function is the relationship between the output and the input of a system. In this context, we are most interested in the effect of the map-making on the power spectrum of the data.

To determine the mapping transfer function due to the pickup subtraction, first a sky signal is simulated and injected into some telescope time stream data. This combination of simulated and real data is mapped, and then the map of the real data alone is subtracted. The output map (resulting from the subtraction of the real data map from the signal injected map) is then compared to the input simulated sky signal in Fourier space. (The real data is used in order to have real noise.) The resulting transfer function is discussed further in Section 4.2.6.

functions in our analysis whenever possible (as well as due to the fact that the pickup maps add an extra step to the map-making, which is already computationally expensive), we decided to not make pickup map for DR6.

2.2 DR4 Maps

DR4 maps include data taken from 2013 to 2016, at both 98 and 150 GHz. While previous analyses used only 650 deg² of the sky, the full DR4 maps cover 17,000 deg². The upcoming DR6 maps will cover a similar area, but with much greater depth.

Example maps from DR4 are shown in Figure 2.6.

For each season/region/array/frequency there are maps for 4 data splits (except for AA which only has 2 splits), as well as a coadd map (of the data coadded across all splits). For each such split (and the coadd) there are 4 maps: **source-free maps**, which contain the I , Q , U Stokes components, **source maps**, which contain the signal from point sources, and which can simply be added to the source-free maps to produce maps of the observed sky, **inverse-variance maps**, which can be used for weighting (for power spectrum calculations, for example), and **cross-linking maps**, which are used to describe cross-linking of the data, as is explained in Section 4.2.3.

2.3 Map-Based Null Tests

We use null tests to target particular systematic effects. Specifically, we check that when the data are split roughly in half based on fast versus slow time constant, high versus low scan elevation, or high versus low precipitable water vapor (PWV, also denoted W_v), the two splits are consistent.

Performing the null tests requires making new maps.⁶ We follow the same procedure as for the primary science maps (Aiola et al., 2020). In each of the three tests, the data are split at the TOD level to maximize the systematic in question while giving roughly equal statistical weight to each subset. From the “null maps” we compute the

⁶This means the tests are computationally expensive. Making the maps needed for the suite of DR4 null tests would have taken approximately 100 years to run on a single CPU core. Thankfully, we have access to large computer clusters.

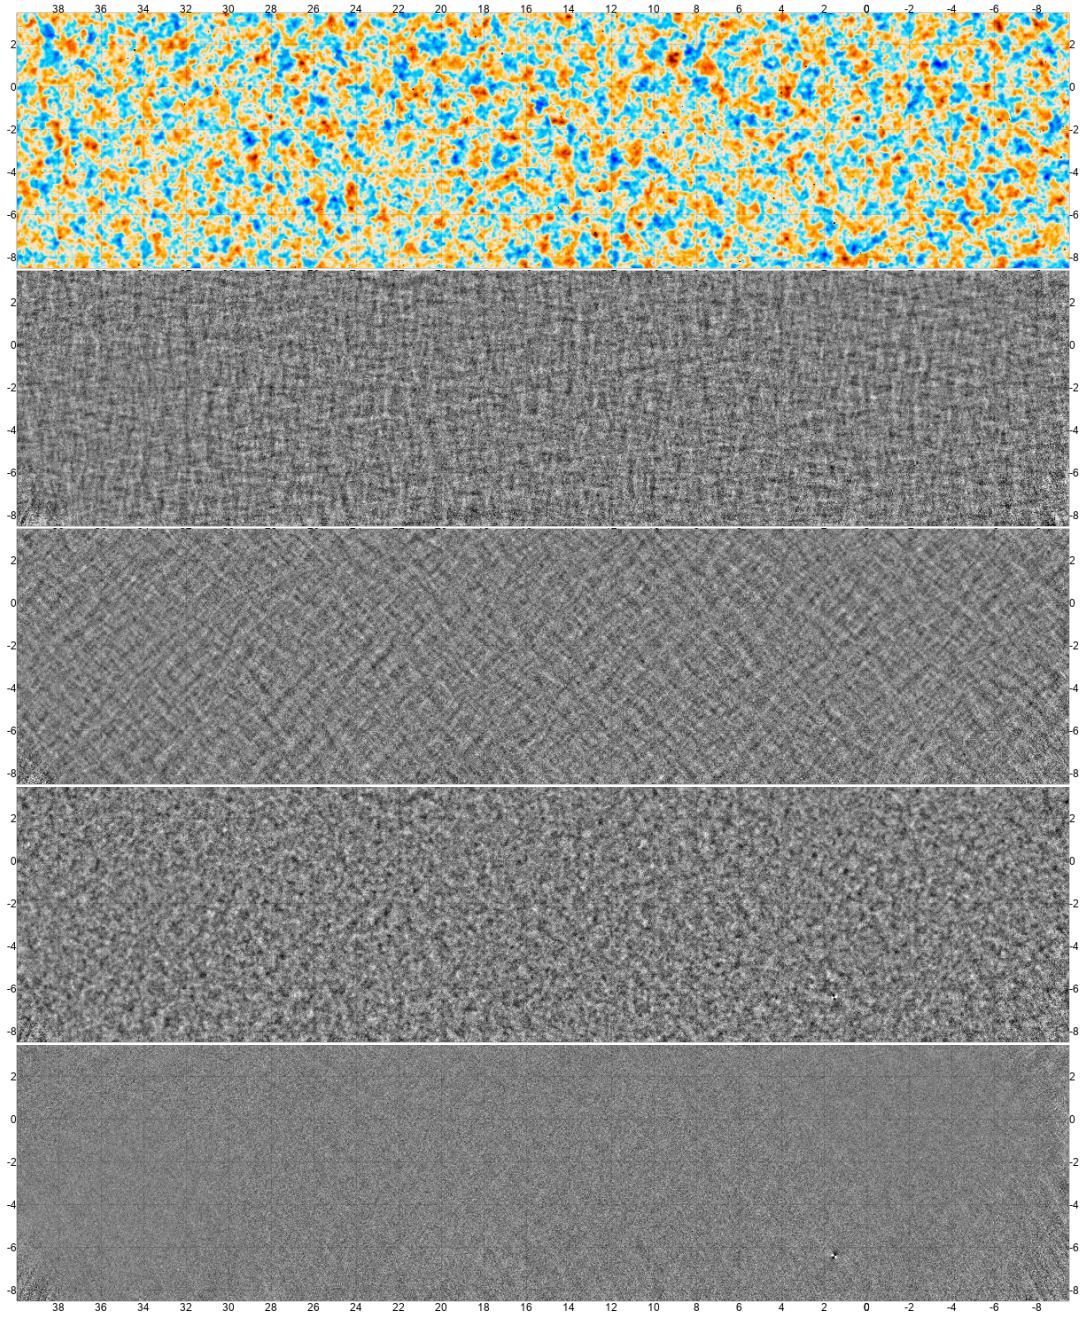


Figure 2.6: Rectangular cutout of the s13–s16 150 GHz maps in the D56 region. The selected area covers 630 deg^2 of the sky. To visually highlight the wide range of measured angular scales, modes with $|\ell| \lesssim 150$ and $|\ell_x| \lesssim 5$ have been filtered out. The top panel shows temperature fluctuations in a range of $\pm 250 \mu\text{K}$. (These maps have been corrected for the final power-spectrum-based calibration factor described in Section 4.2.7.) The remaining four black-and-white panels show (top to bottom) Q , U , E -mode, and B -mode polarization measurements in a range of $\pm 30 \mu\text{K}$. The x -axis (y -axis) shows the RA (Dec) coordinates in degrees. Figure from [Aiola et al. \(2020\)](#).

power spectrum of the difference. We then compute the χ^2 of the difference-spectrum with respect to zero. The error bars for the spectra are estimated analytically, because generating simulations for all the null tests is computationally prohibitive.

In this section, we present the map-based null tests performed for DR4.

2.3.1 Time Constants

The time response of each detector is limited by its electrothermal properties, and in the low-inductance limit can be modeled as a one-pole filter with time constant $\tau = 1/(2\pi f_{3\text{dB}})$ (Irwin & Hilton, 2005). The finite response time results in a small shift in the measured position of a point on the sky, depending on the scan direction. If not properly corrected, they can lead to a low-pass filtering of the data. The time constant null maps are designed to assess this effect.

We split the data so that “low” corresponds to τ below the median value in Table 2.1 and “high” is above. The results of the test are given in Table 2.2 and Figure 2.7. There is good consistency between the low and high detector time-constant data.

Table 2.1: Median Time Constants

Array & Band	S13	S14	S15	S16
PA1 (150 GHz)	2.1 (76)	3.9 (41)	5.4 (29)	-
PA2 (150 GHz)	-	1.9 (84)	2.3 (69)	1.7 (94)
PA3 (98 GHz)	-	-	1.2 (130)	1.1 (140)
PA3 (150 GHz)	-	-	1.1 (140)	0.98 (160)

The median time constants, τ , in ms, (along with $f_{3\text{dB}}$, in Hz), which depend on loading and base temperature. We report them for $0.5 < W_v < 1$ mm. For nominal observations, 1 Hz corresponds roughly to $\ell = 400$, thus $f_{3\text{dB}} = 35$ Hz maps to $\ell = 14,000$.

Table 2.2: Summary of the Time Constant Null Tests

Array	Frequency	χ^2/ν (PTE)
PA1	150 GHz	1285/1248 (0.23)
PA2	150 GHz	595/624 (0.79)
PA3	98 GHz	291/312 (0.80)
	150 GHz	294/312 (0.76)

For each array we report χ^2/dof (PTE, probability to exceed).

2.3.2 Elevation of Observations

Maps made from scans at low and high elevations will have different levels of ground and atmosphere contamination. The elevation split is designed to search for this contamination.

The elevation at which we split is computed separately for each region and varies between $\alpha = 37^\circ$ and $\alpha = 55^\circ$. In the BN and D8 regions, it is not possible to split the TODs into a high-elevation and a low-elevation group while maintaining enough coverage and cross-linking to make a proper map of each. In these cases, we make a map with 80% of the high-elevation data and 20% of the low-elevation data and compare this to a map with the percentages reversed. These ratios were chosen to ensure the resulting maps are sufficiently cross-linked and to ensure that the maps still allow us to test the effect of scans at different elevations.

The results of this null test are reported in Table 2.3 and shown in Figure 2.7. We see no evidence of an elevation-dependent effect.

2.3.3 PWV

The atmosphere emits and absorbs at ACT frequencies, in large part due to the excitation of the vibrational and rotational modes of water molecules. Thus, the PWV is correlated with the level of optical loading on the detectors and to the level of atmospheric fluctuations. Both the increased loading and fluctuations could bias

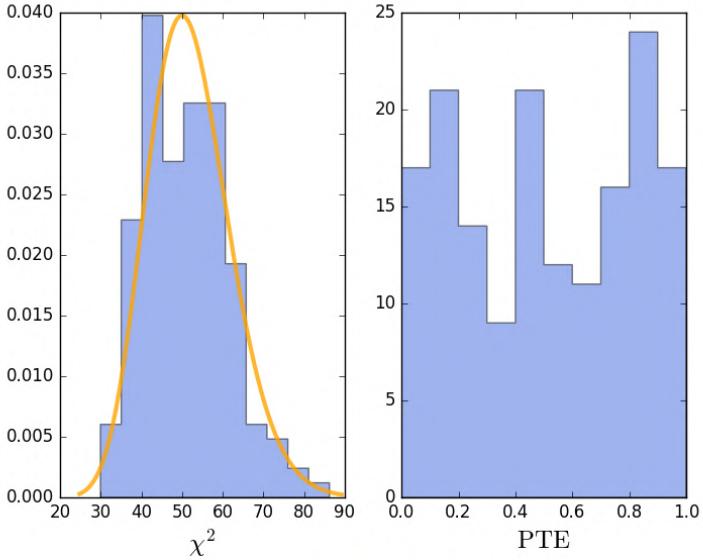


Figure 2.7: *Left:* Histogram of χ^2 from all 162 τ , PWV, and elevation null tests is shown along with the expected distribution if the null tests were uncorrelated, although the correlations will have a negligible effect on this plot. The expectation is $\chi^2 = n_{\ell,c} = 52$ corresponding to the number of ℓ bins. *Right:* Distribution of the probability to exceed for each measure of χ^2 on the left based on $n_{\ell,c} = 52$ degrees of freedom.

our maps. The low versus high PWV null test is designed to assess this possibility.

The median PWV is $W_v = 0.88$ mm, with quartile breaks at 0.63 mm and 1.36 mm. The PWV at which we split is different for each season and region. For D56 in s14, for example, TODs with $W_v < 0.80$ mm are “low” and those with $W_v > 0.80$ mm are “high.” The dividing line ranges from $W_v = 0.51$ mm to $W_v = 0.85$ mm. In general, the TODs for null tests are split so that both subsets have similar noise levels, but for PWV splits there is an additional challenge.

The noise spectra for the high-PWV and low-PWV splits are noticeably different. The $\ell \leq 4000$ noise is due to the atmosphere and $\ell \geq 4000$ noise is instrumental. So when the TODs are split such that the high- ℓ tails of the noise curves agree, as in Figure 2.8, the high-PWV split has more low- ℓ noise. Similarly, when the TODs are split such that the low- ℓ part of the curves agree, the high- ℓ parts do not. A few

cases were tested and both PWV splits described above passed our null test. For the remaining tests, the TODs were simply all split so that the high- ℓ tails of noise curves would agree.

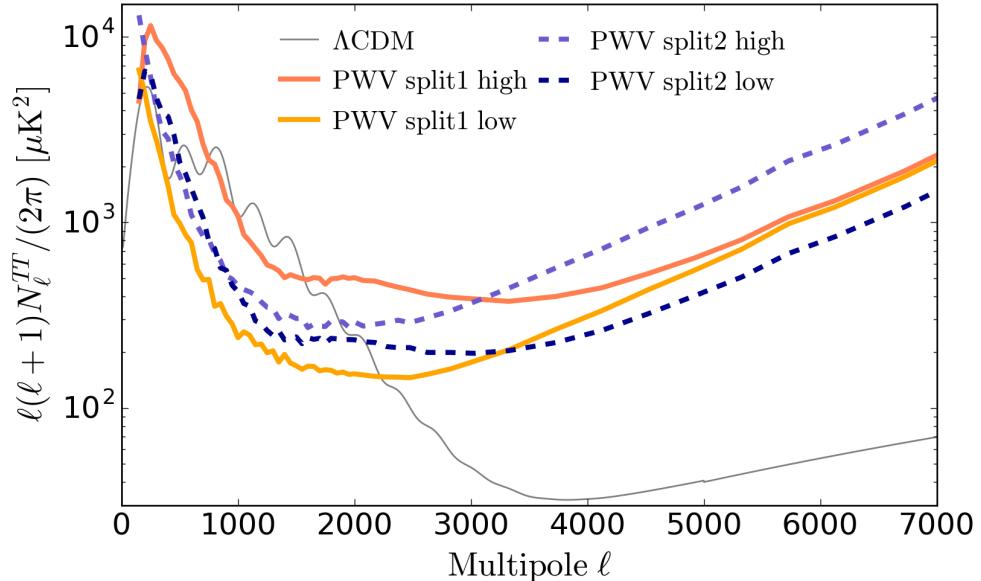


Figure 2.8: Noise for different PWV splits for D56. Solid orange and yellow (dashed purple and blue) curves show the noise power spectra from data split into observations during high and low PWV respectively, keeping the small (large) angular scale noise similar in the two splits.

We had one failure of this null test in the BN region at 98 GHz for $\ell < 600$.⁷ We are still investigating the source of this failure, but as a precaution, we eliminated the 98 GHz data for $\ell < 600$ from the cosmological analysis. This was done before unblinding. After accounting for this failure, we see no evidence of a PWV-dependent effect. The results of this null test are reported in Table 2.3 and Figure 2.7.

⁷At the time of writing this thesis, the paper [Choi et al. \(2020\)](#), in which this failure was reported, dates from over two years ago. The PTE for this null test was not reported in the paper. A brief search of our records did not yield this PTE either, so it may be lost to history.

Table 2.3: Summary of Elevation and PWV null tests

Region	PWV χ^2/ν (PTE)	Elevation χ^2/ν (PTE)
D5	-	341/312 (0.12)
D6	-	289/312 (0.82)
D56	1152/1248 (0.98)	1822/1872 (0.79)
BN	1035/936 (0.013)	1138/1248 (0.988)

We report χ^2/dof (PTE, probability to exceed) for the regions as shown. These values obtain after removing the 98 GHz data that did not pass the null test.

Chapter 3

Beams

This chapter is adapted from [Lungu et al. \(2022\)](#). I led the writing of this paper and contributed to the recent modifications to the beam analysis pipeline.

Most of the code used to analyze the beams has been reused from previous work, including [Hasselfield et al. \(2013\)](#) and [Hincks et al. \(2010\)](#). Marius Lungu made some modifications to this code, and led the beam analysis for DR4.

Improvements I contributed to the beam pipeline for this paper related to the fitting of the beam model (§3.4.2), and include: the use of the corrected Akaike information criterion (AICc) to select the optimal number of terms to fit to the beam profiles, varying the radius θ_1 at which the model for the beam changes, using MCMC to sample the posterior distribution of the parameters ℓ_{\max} , a_n , and α , in order to estimate the full covariance between all parameters. I also investigated a few other possible changes to the beam pipeline, including different ways of dealing with the offsets due to the map-maker bias when computing an average beam profile, and the possibility of fitting different exponents for the asymptotic beam behavior.

Lyman Page and Rolando Dünner (along with Marius Lungu) led the work on the scattering term for the beam (§3.4.2), Matthew Hasselfield performed the work on the

polarized sidelobes (§3.5.2), and Adri Duivenvoorden investigated the effects of the beam asymmetry (§3.6.2).

3.1 Introduction

The Atacama Cosmology Telescope (ACT) is a 6 m off-axis Gregorian telescope located at an altitude of 5190 m in the Atacama Desert of northern Chile. It is designed for millimeter wavelength observations of the cosmic microwave background (CMB) at arcminute resolution. The telescope and receiver are described in [Fowler et al. \(2007\)](#) and [Thornton et al. \(2016\)](#) respectively. This chapter describes the treatment of the ACT beams for its fourth data release, referred to as DR4. This data release includes temperature and polarization data collected by ACT between 2013 and 2016, covering seven regions of the sky spanning roughly 18,000 deg², in frequency bands centered on 98 and 150 GHz ([Thornton et al., 2016](#)). The DR4 data were obtained using two 150 GHz detector arrays, PA1 and PA2, and one dichroic detector array, PA3, operating at 98 and 150 GHz. The three array positions are not identical optically. Each year of data is referred to as a season (S13–S16) and was analyzed separately. The DR4 data products and some of their analyses are presented in [Choi et al. \(2020\)](#), [Aiola et al. \(2020\)](#), [Darwish et al. \(2020\)](#), [Han et al. \(2021\)](#), [Madhavacheril et al. \(2020\)](#), [Namikawa et al. \(2020\)](#), and [Mallaby-Kay et al. \(2021\)](#).

The beams of the telescope determine the instrument’s response to different scales on the sky. Quantifying these beams and the uncertainties on the beam measurements is one of the most challenging aspects of characterizing the instrument. An incorrect calibration of the beams would bias virtually all of the ACT science analyses, including the precision measurement of the CMB power spectrum.

The previous paper that focused on ACT beams, [Hasselfield et al. \(2013\)](#), dates from DR2, which included data through 2010. The DR3 analysis papers ([Louis et al.](#),

2017; Naess et al., 2014) referenced Hasselfield et al. (2013) and described the small changes in methodology since then. In recent years, there have been several modifications to our beam pipeline that merit further discussion. This chapter provides a comprehensive guide to the analysis of the ACT beams for DR4.

The chapter is organized as follows. In §3.2 we describe the planet observations used to measure the main portion of the ACT beams. In §3.3 we explain how these observations are made into maps. In §3.4 we describe the steps of the beam pipeline from planet maps to a model of the ACT beams and their covariance. §3.5 accounts for the polarization component of the beams. Finally, in §3.6 we discuss the publicly released beam products, assumptions made in the analysis, and future directions for ACT beam analyses. Part of the information in this chapter overlaps with that in previous ACT papers, including Aiola et al. (2020), Choi et al. (2020), and Hasselfield et al. (2013), and is included here for clarity and completeness.

3.2 Observations

Observations of sufficiently bright point-like objects are required in order to properly characterize the shape of the beams out to large angles. For ACT, planets are the best candidates for this purpose, though not all equally so. The proximity of Mercury and Venus to the Sun make them difficult to observe at night; Mars' temperature is not sufficiently constant (see, for example, Wright, 1976), so it intermittently becomes too bright; Jupiter's apparent size is too large, compared to the instrument's angular resolution, to be considered a point source (and it is too bright, as is Saturn); Saturn, though suitable for sidelobe analysis, is bright enough to induce a non-linear detector response near peak amplitude, resulting in signal suppression within the main lobe of the beam;¹ Neptune's small angular diameter relative to the beamwidth

¹Such detector non-linearity may have contributed to the $\sim 7\%$ discrepancy between the *Planck* and ACT DR2 measurements of Saturn's temperature (Planck Collaboration Int. LII, 2017; Has-

dilutes its brightness, reducing signal-to-noise. We have thus chosen to base the majority of the ACT beam analysis on observations of Uranus, which achieves adequate signal-to-noise without exceeding the dynamic range of the instrument and can be approximately treated as a point source.² As described in §3.5.2, Saturn is used to study the beam’s polarized sidelobes.

Planet observations are done by briefly interrupting the CMB survey scans, changing the azimuth at which the telescope is pointing, and scanning again at a fixed elevation. Uranus is observed roughly once per night during the observing season, but only the higher quality observations are used to characterize the beam. Before making maps of the observations, we apply the same data selection criteria as for the main maps used for cosmological analysis. For example, we discard data taken during the daytime, data taken when the weather was particularly bad, or data from detectors that do not meet criteria summarized in [Aiola et al. \(2020\)](#) and detailed in [Dünner et al. \(2013\)](#). Once maps of the observations are made, further selection criteria are applied, as described in §3.4.1.

3.3 Planet Map-Making

3.3.1 Filtering And Noise

Uranus maps were made using a special filter-and-bin map-maker built on the same framework used for building our normal maximum-likelihood maps (`enki`).³ The reason we do not use maximum-likelihood maps for planets is that these are not robust to “model errors,” as described in e.g. [Naess \(2019\)](#). The difference between

selfield et al., 2013).

²While not implemented for the current beam analysis, another possibility being studied involves using observations of Uranus to measure the central portion of the beam (where Saturn would induce a non-linear response) and observations of Saturn to measure the parts of the beam further from the peak. This way, Saturn’s brightness allows us to measure the outer parts of the beam to greater signal-to-noise than when using observations of Uranus alone.

³GitHub repository: <https://github.com/amaurea/enlib>

the true continuous sky and the pixelated model is interpreted as noise, leading to a bias which propagates along the scan direction for a noise correlation length. This is a significant effect for a bright, localized signal such as a planet.

It is possible to make maximum-likelihood models robust to these effects at significant computational and implementational cost (Piazzo, 2017), but we here choose a simpler method that exploits the compactness of the planet signal relative to the atmospheric noise correlation length to separate the two. The idea is to split the data into two regions: a target region (including the planet) which we want to map, and a reference region (the rest of the data) which is used to estimate the behavior of the atmospheric noise.⁴ The atmosphere in the target region is then estimated by interpolating the atmospheric modes from the reference region, and it is then subtracted from the target region of the map.⁵ To the extent that the signal in the target region is uncorrelated with any signal in the reference region, then this subtraction will not introduce bias on any scale in the target region. A point source like Uranus is very close to fulfilling these criteria, but in practice the outermost parts of the beam extend into the reference region, introducing a slight bias that we address in §3.3.2.

The target region must be large enough to capture all of the main beam, but small enough that the atmosphere’s behavior in the target region can be adequately predicted using only data from the reference region. We choose a 12' radius area centered on the planet as the target region, and estimate the atmosphere’s behavior here by exploiting its spatial correlation structure. Nearby detectors see similar parts of the atmosphere, so if at one moment in time some detectors are seeing the target region and some are seeing the reference region, the data from the latter are used to help determine the atmospheric modes in the former.⁶ This new interpolation tech-

⁴The atmospheric fluctuations seen by ACT are described in more detail in Morris et al. (2022).

⁵For a previous implementation of a map-making method designed to model and subtract the atmosphere from maps of planets, see Hincks et al. (2010).

⁶Ideally one would also use the temporal correlations of the atmosphere, and solve the most probable value of the atmosphere in the target region given the data in the reference region. This

nique significantly improves the noise modelling compared to the technique used in Louis et al. (2017), which used a simpler variant of filter-and-bin where the detector common mode was subtracted.

The resulting planet maps are cleaner and dominated mostly by white noise. In addition, we now inverse variance weight the detectors when making the planet maps, which was not done previously. This improved mapping has contributed to a more detailed understanding of the temperature-to-polarization leakage, which will be described in §3.5.1.

3.3.2 Bias From Map-Making

While it is approximately true that all the planet signal is contained inside the target region, a small part of it extends into the reference region due to the faint wings of the beams. This introduces a small but non-negligible bias in the target region when the atmosphere is subtracted.

To study this bias, a set of simulated beam-convolved Uranus observations are processed through the map-maker. The beam-convolved planet signal is simulated using a 2D beam model⁷ which uses the inverse Fourier transforms of two-dimensional Zernike polynomials as its basis set, similar to the 1D model of the beam described in §3.4.2. This simulated planet signal is then injected into several subsets of real time-domain data taken from CMB observations. The resulting data are then processed through the planet map-making pipeline, which includes estimating and subtracting the atmospheric noise. Both the CMB observations alone and the CMB observations with the planet signal injected are mapped separately, and then the former are subtracted from the latter. The result is a set of noise-free simulated maps of Uranus

⁷This model was developed as part of a larger, ongoing project to fit the beams in two dimensions in order to capture their asymmetry.

observations.⁸

An example of a map of a simulated Uranus observation is shown in Figure 3.1. Radial profiles are constructed by taking the azimuthal average of the input beam model and the mean output of the simulated observations. The difference between these two radial profiles represents the bias due to the atmospheric subtraction. As seen in Figure 3.2, other than small residual fluctuations near the center of the beam, the bias can be well approximated by a constant offset.⁹ In these simulations the target region for atmosphere subtraction was chosen to have a 16' radius, but the qualitative conclusions we draw are applicable to the final planet maps used for the beam analysis which have a target region with a 12' radius.¹⁰ As shown in Figure 3.2, the effect of the atmosphere removal is almost exactly as expected: the value of the input beam at the target radius (16' in this case) is uniformly subtracted within that radius to produce the output beam.

So the planet map-making bias is well approximated by a constant offset in each map, determined by the value of the Uranus signal at the edge of the target radius used for noise estimation. This offset is estimated and subtracted prior to analysis, as described in §3.4.2.

⁸In the end, it was discovered that including the atmospheric noise in the simulations does not affect the results, since the atmosphere removal is approximately linear.

⁹These small residual features only appear in the T maps when solving for the maps in T, Q, U simultaneously. When we solve only for T and set $Q = U = 0$, the fluctuations disappear. This seems to be due to an issue with the conditioning of the covariance matrix in the simulations, and is not a physical effect.

¹⁰During preliminary beam studies, maps of both Uranus and Saturn were made, with target regions both of radius 12' and 16'. Saturn was found to be too bright. In the Uranus maps with a target region of radius 16', the signal-to-noise in the radial profiles past 12' was insufficient to be of use in the analysis. Hence, the Uranus maps with a target region of radius 12' were used, in order to make the best possible prediction of the atmosphere in the region of interest.

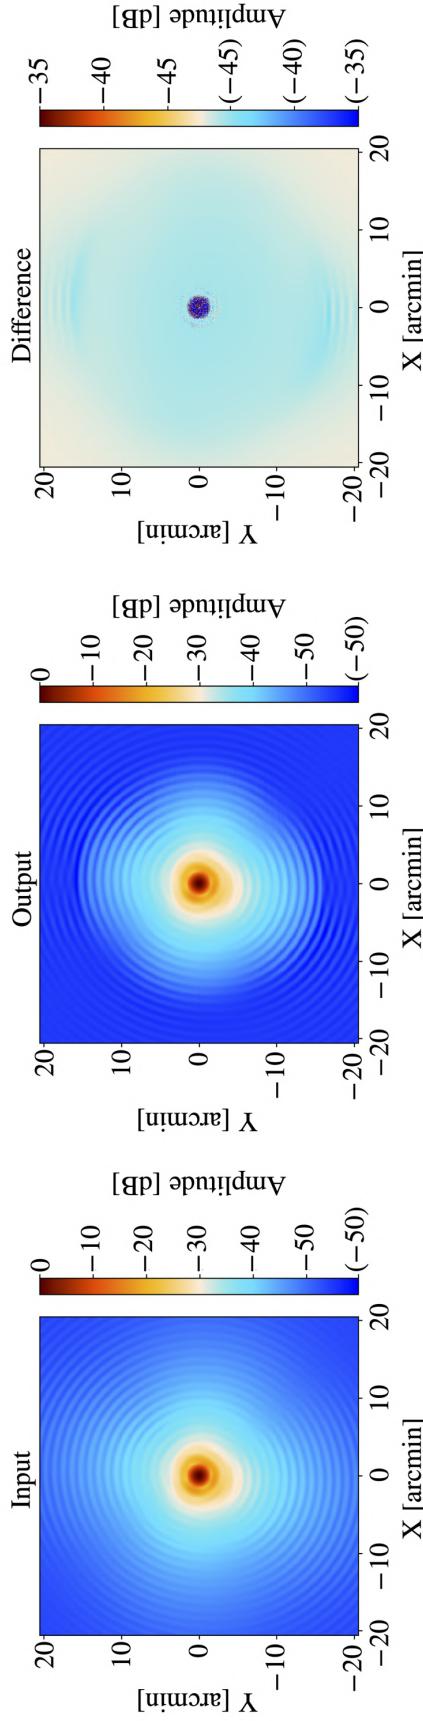


Figure 3.1: Example of a simulated planet observation used to characterize the mapping transfer function. This example is for S16 PA2 at 150 GHz. (Left) The input 2D beam model. (Center) One simulated observation. (Right) The difference between the “input” beam model (left) and the map of the “output” simulated observation (center). Note that in each case the color scale is a symmetrical log scale^a in dB (with a linear threshold of -50 dB for the input and output maps and -45 dB for the difference map) and negative values are enclosed by parentheses. Here the target region where the atmosphere was estimated and removed during the map-making process was chosen to have a radius of $16'$. The large-scale residuals seen are mostly constant within this target region of the map. The slight asymmetry in the beam (which may be safely ignored for DR4, as explained in §3.6.2) is due primarily to the positions of the telescope’s optics tubes in the focal plane.

^aA symmetrical log scale is logarithmic in both the positive and negative directions, with a small linear range around zero to avoid having values tend to infinity (https://matplotlib.org/stable/api/scale_api.html#matplotlib.scale.SymmetricalLogScale).

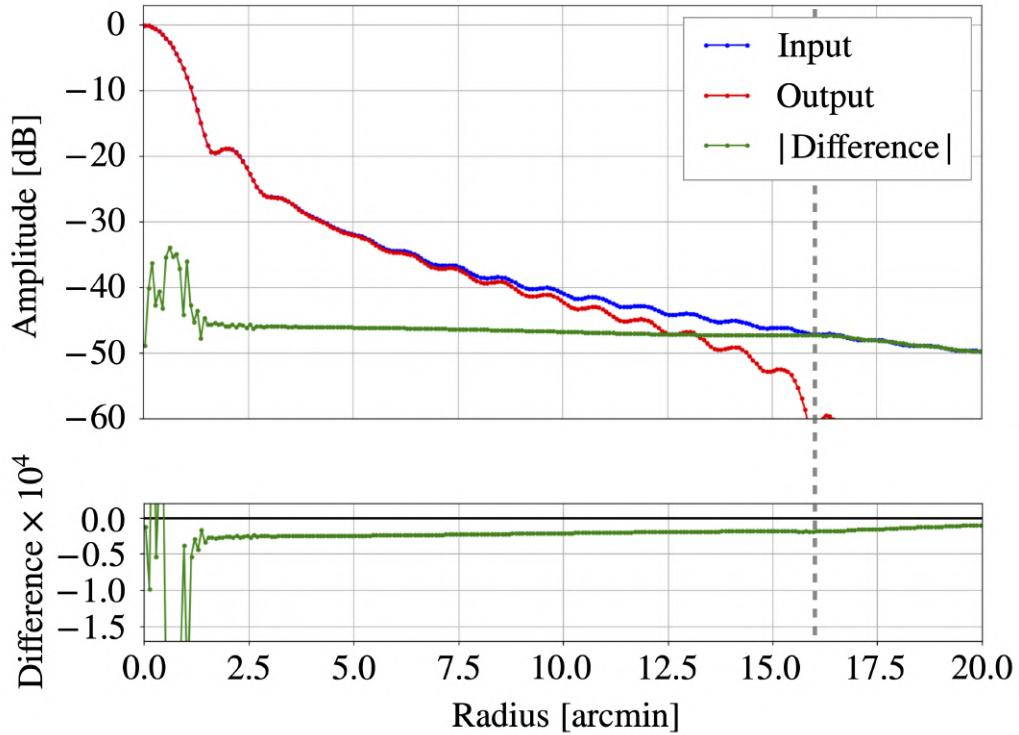


Figure 3.2: The average radial profile of the simulated Uranus observations for S16 PA2 at 150 GHz, comparing the azimuthal averages of the “input” 2D beam model and the mean map of the “output” simulated observations. The vertical dashed line indicates the $16'$ radius target region of the map within which the atmosphere was estimated and subtracted. (For the Uranus maps used to characterize the beam, a $12'$ radius target region is used, but the qualitative conclusions we draw here still apply.) Other than the small (roughly -40 to -30 dB) variations near the beam center, the difference is well approximated by a constant offset in the region from $3.5'$ to $10.0'$ where we fit for it. As described in §3.4.6, we adjust our beam covariance matrices to account for possible uncertainty due to variations in the range over which we fit the offsets, exploring the three independent ranges of $3.5' - 5.0'$, $5.0' - 7.0'$, and $7.0' - 10.0'$.

3.4 Beam Pipeline

3.4.1 Map Selection And Processing

A map of Uranus is made for each detector array, frequency, and observation.¹¹ Only a fraction of these maps is used in the final beam analysis. We select maps made from observations done at night (23–11 UTC) taken with the secondary mirror in its final focus position for the season in question. We pick the maps where there are not too many zero-weight pixels (which are indicative of poor scan coverage), where the signal-to-noise is sufficiently high (> 10), and where there are not too many horizontal stripes (determined by visual inspection of the maps). This striping in the maps is caused by large-scale variations due to residual atmosphere. The rejection of maps with too much striping is new, and is done to simplify the analysis. It results in fewer maps being selected than in [Louis et al. \(2017\)](#). For each season, array, and frequency, the number of maps discarded due to striping ranges from 3 to 42 (on average 19), which represents between 4% and 36% of the maps (on average 16%).

The number of Uranus maps used for the beam analysis versus the total number of observations is shown in Figure 3.3 and Table 3.1. For example, in 2014 (S14) we made 129 observations of Uranus. In the case of PA2, we discarded 34 observations because at the time the telescope pointing was optimized for PA1 (so Uranus was not well measured by PA2), 36 observations because the resulting maps had too much striping, and 8 because the signal-to-noise was too low. This leaves 51 Uranus observations to measure the beam. An example “good” map which was selected for the beam analysis and a “bad” map which was discarded due to too much striping are shown in Figure 3.4. In S13, a large fraction of the Uranus maps were thrown out because the observations were made in the early commissioning phase of the telescope, before

¹¹So we do not make any per-detector maps. For each detector array, frequency, and observation of Uranus, we combine the data from all the detectors that meet the quality criteria mentioned in §3.2 into one map.

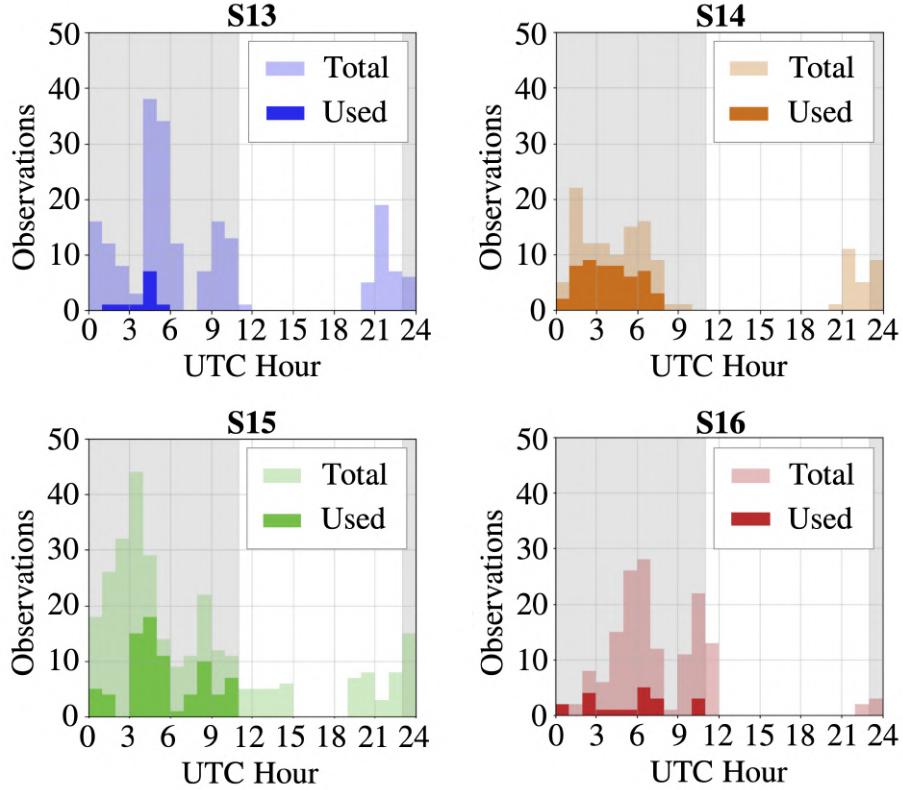


Figure 3.3: Distribution of the total number of Uranus observations that were made (faded colors) and ultimately became part of the final beam analysis (solid colors) for all arrays combined, shown by observing seasons from 2013–16. The shaded regions between 0 UTC and 11 UTC, as well as 23 and 24 UTC, demarcate the ACT nighttime dataset (note that local time at the observing site fluctuates between UTC-3 and UTC-4). See Table 3.1 for a summary of the number of observations used vs total for each detector array and season.

it had achieved its final focus. In S16, PA4 was added, and several observations of Uranus were made with the telescope pointing optimized for PA4. However, since PA4 was not included as part of the DR4 analysis, these observations were not used here, leading to a smaller fraction of the Uranus maps being selected for S16.

For the purposes of the beam measurements, we re-center each Uranus map by fitting a 2D Gaussian in the vicinity of the planet. The amplitudes from these fits are used to normalize the beam profiles to have peak values of unity. We also estimate the noise level in each map by computing the standard deviation of the pixel values outside the target region (the $12'$ radius area centered on the planet, described in §3.3.1).

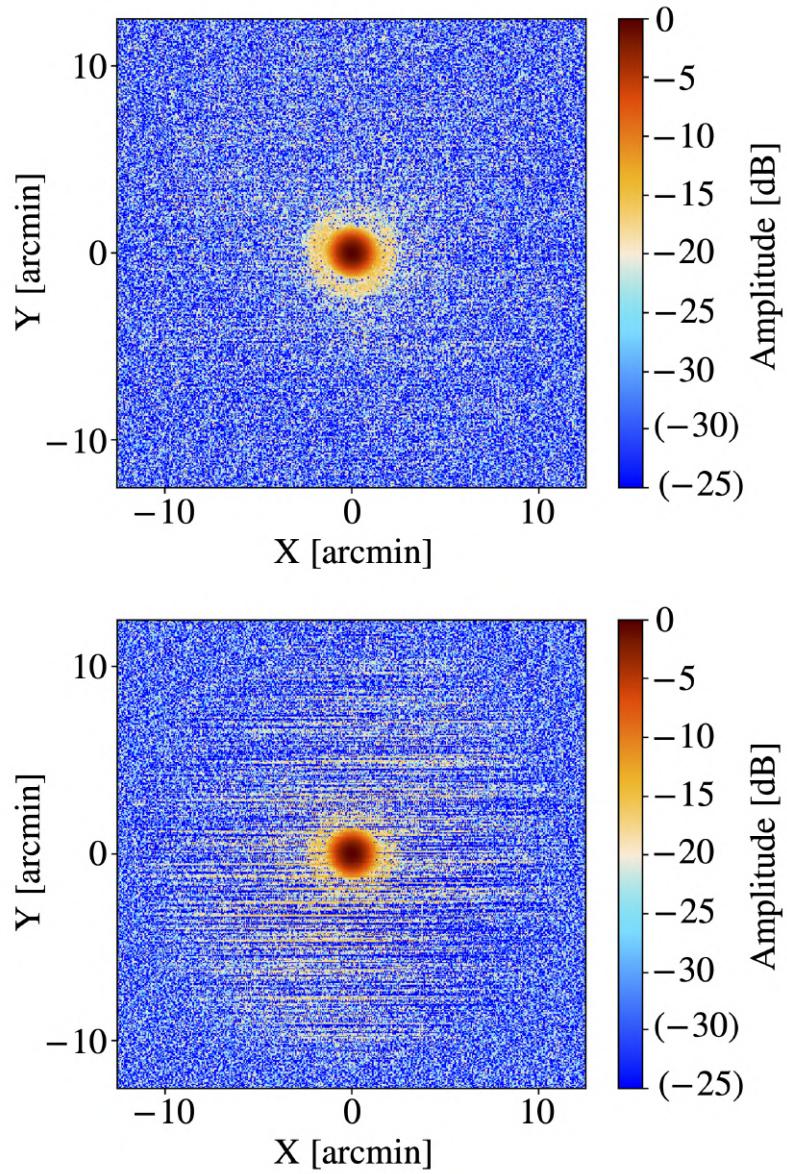


Figure 3.4: Two maps of individual Uranus observations for S15 PA2 at 150 GHz. (Top) A “good” map which was selected for the beam analysis. (Bottom) A “bad” map which was thrown out due to having too many horizontal stripes (caused by large-scale variations due to residual atmosphere remaining after the atmosphere subtraction described in §3.3.1). In each case, the color scale is a symmetrical log scale in dB, with a linear threshold of -30 dB and negative values enclosed by parentheses. Any slight asymmetry visible here may be safely ignored for DR4, as explained in §3.6.2, and is due primarily to the positions of the telescope’s optics tubes in the focal plane.

Table 3.1: Summary of Uranus observations — number used/total.

Array	Band (GHz)	Season	Used	Total
PA1	150	S13	11	197
		S14	45	129
		S15	17	133
PA2	150	S14	51	129
		S15	38	164
		S16	11	86
PA3	150	S15	8	117
		S16	6	78
PA3	98	S15	33	117
		S16	9	78

3.4.2 Radial Profile Fitting

Radial Profiles

The instantaneous ACT beams are slightly elliptical. However, when fitting the beam we work with azimuthally averaged (“symmetrized”) radial beam profiles, treating the beams as if they were circular. In our case, the cross-linking of the scans (visible in Figure 2 of [Choi et al., 2020](#)) means that the telescope beams contribute to the maps with roughly equal weight at two different orientations that are approximately 90 degrees apart for a large part of the ACT sky coverage. The effective beams are thus circularized.¹²

This symmetrizing effect works well in temperature, but it does not help with the temperature-to-polarization leakage caused by beam asymmetry. However, this effect has been quantified and may be safely ignored for DR4, as explained in §3.6.2.

The ACT beams are small enough that we use a flat sky approximation for modeling them. We denote each beam map by $B(\theta, \phi)$, where θ is the radial distance from the

¹²In reality, anisotropic noise weighting (due to ACT’s noise being correlated along the scan direction) makes the beam symmetrization from the map-maker different from what one gets by simply averaging. The overall effect is to either make the circularized beam smaller or larger than the naive prediction, depending on the direction the raw beam is elongated relative to the scanning direction. This effect is absorbed into the jitter correction described in §3.4.5.

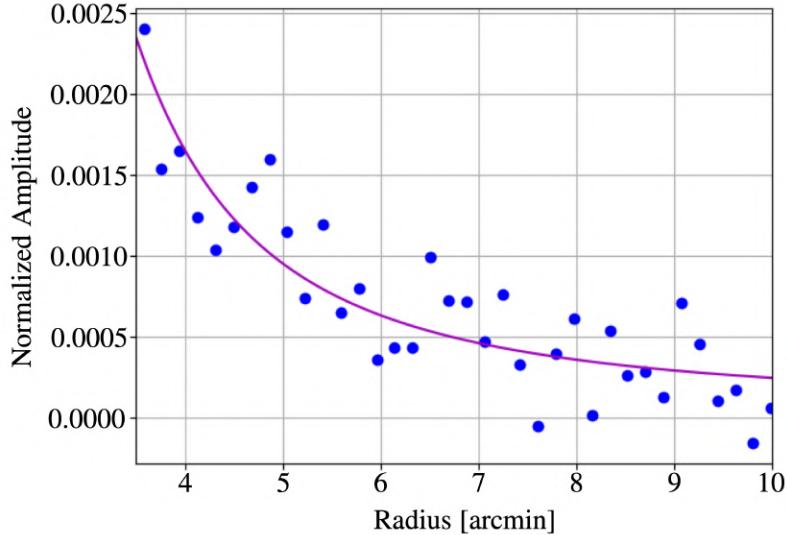


Figure 3.5: Example of the offset fit for the radial profile of one Uranus observation for S14 PA2 at 150 GHz. The points are found by taking the azimuthal average of the map for each radial bin. Uncertainties on these points are not computed. The curve is the best-fit model $\alpha/\theta^3 + \beta$, where in this case $\alpha = 0.08934$ and $\beta = 0.00012$.

beam center and ϕ is the polar angle. We set $B(0, \phi) = 1$. The symmetrized radial beam profile is then

$$B(\theta) = \frac{1}{2\pi} \int d\phi' B(\theta, \phi') . \quad (3.1)$$

Each map, with a target region of radius $12'$, is binned into a symmetrized radial profile with bins of width $11''$ out to a radius of $10'$, for a total of 55 bins. The dominant source of noise in the Uranus maps comes from variations in the atmosphere which occur at relatively large angular scales, so there is significant correlation between the bins in each radial profile.

Before combining the radial profiles into one mean profile for each season, array, and frequency, each map's profile must be corrected for the offset due to the bias induced by the planet map-maker that was described in §3.3.2. This is done by fitting the model $\alpha/\theta^3 + \beta$ (motivated in §3.4.2) plus a scattering term (described in §3.4.2) to each profile in the range $3.5' - 10.0'$ and then subtracting the measured offset, β , from each. An example of such a fit is shown in Figure 3.5.

We then compute an average radial profile independently for each season, detector array, and frequency. When averaging radial profiles, each profile is weighted by $1/N$ where N is the white noise variance estimated from the corresponding map. The averaged radial profiles extend to roughly -40 dB of the peak, or 35 dBi (40 dBi) at 150 GHz (98 GHz).

An advantage of computing each radial profile then taking the average, compared to making one average Uranus map then computing its radial profile (as done in [Hincks et al. \(2010\)](#), for example), is that it more easily allows for the estimation of the full covariance matrix for the averaged radial profile. The off-diagonal elements of this matrix are important since there is significant noise correlation between the radial bins, which propagates into the low- ℓ uncertainty of the beam window function described in [§3.4.3](#).

Radial Profile Covariance Matrix

As the averaged radial profile is computed for each season, array, and frequency, the covariance matrix between the amplitudes of the 55 radial bins is also computed in order to account for the covariant uncertainty on large angular scales. Given the small sample size, with only between 6 and 51 profile measurements used to estimate each matrix, a shrinking algorithm ([Schäfer & Strimmer, 2005](#)) is used to down-weight the off-diagonal terms of the matrix. This method is described in detail in [§3.7.1](#). In our case this shrinking procedure is necessary as a standard estimate of the covariance matrix is often so noisy it becomes ill-conditioned and cannot be inverted.

The shrinkage technique works by combining an empirical estimate of the covariance matrix \mathbf{S} (a high-dimensional estimate of the underlying covariance with little or no bias) with a model \mathbf{T} (a low-dimensional estimate which may be biased but has much smaller variance) to minimize the total mean squared error (sum of bias squared and

variance) with respect to the true underlying covariance:

$$\mathbf{C} = \widehat{\lambda}^* \mathbf{T} + (1 - \widehat{\lambda}^*) \mathbf{S} . \quad (3.2)$$

Here $\widehat{\lambda}^*$ is the parameter (often referred to as shrinkage intensity) that determines the contribution of each matrix. In our case, the covariance matrix \mathbf{S} is an unbiased empirical estimate of the covariance, the sample covariance matrix, and the model matrix \mathbf{T} is given by the diagonal of \mathbf{S} :

$$T_{ij} = \begin{cases} S_{ii} & \text{if } i = j \\ 0 & \text{if } i \neq j , \end{cases} \quad (3.3)$$

a common choice. We analytically calculate the optimal combination of the low and high dimensional estimates by determining $\widehat{\lambda}^*$ from \mathbf{S} :

$$\widehat{\lambda}^* = \frac{\sum_{i \neq j} \widehat{\text{Var}}(S_{ij})}{\sum_{i \neq j} S_{ij}^2} , \quad (3.4)$$

where $\widehat{\text{Var}}(S_{ij})$ is an estimate of the variance of each covariance matrix element. The derivation of $\widehat{\lambda}^*$ can be found in §3.7.1. For the analysis presented here, $\widehat{\lambda}^*$ ranges from 0.24 to 1.

Core Model

Since ACT's primary mirror has a diameter of $D = 6$ m, at a given wavelength λ the diffraction limited optical response of the telescope is restricted to spatial frequencies below $\ell_{\text{max}} \simeq 2\pi D/\lambda$. The Fourier transform of each beam is therefore compact on a disk. Thus, a natural choice of basis functions to model the beams in harmonic space is the Zernike polynomials, an orthonormal set of basis functions that is complete on

the unit disk.¹³

Expressed in polar coordinates ρ and ϕ in the aperture plane (where ρ is the radial distance $0 \leq \rho \leq 1$, and ϕ is the azimuthal angle $0 \leq \phi \leq 2\pi$), the Zernike polynomials may be written as

$$V_n^m(\rho, \phi) = R_n^m(\rho) e^{im\phi}, \quad (3.5)$$

where m and n are integers such that $n \geq |m|$ and $R_n^m(\rho)$ are a set of radial functions defined as

$$R_n^{\pm m}(\rho) = \begin{cases} \sum_{k=0}^{\frac{n-m}{2}} \frac{(-1)^k (n-k)!}{k! (\frac{n+m}{2} - k)! (\frac{n-m}{2} - k)!} \rho^{n-2k} & \text{for } n - |m| \text{ even} \\ 0 & \text{for } n - |m| \text{ odd.} \end{cases} \quad (3.6)$$

For an azimuthally symmetric beam, it is only necessary to consider the polynomials for which $m = 0$, which may be expressed as

$$R_{2n}^0(\rho) = P_n(2\rho^2 - 1), \quad (3.7)$$

where $P_n(x)$ are Legendre polynomials (Born & Wolf (1980)). The Fourier transform of these polynomials is

$$\tilde{R}_{2n}^0(\theta) = \int_0^1 \rho \, d\rho \, e^{-i\rho\theta} R_{2n}^0(\rho) = \frac{(-1)^n J_{2n+1}(\theta)}{\theta}, \quad (3.8)$$

where J_{2n+1} is a Bessel function of the first kind.

Motivated by this, and following previous analyses (Hasselfield et al., 2013), we adopt

¹³The Zernike polynomials are usually used to fit effects in the electric field, not in intensity, but since they are a basis set, we can use them to fit the intensity beams. So the asymptotic behavior of f_n in Equation 3.9 is $1/\theta^{1.5}$, not $1/\theta^3$ as we expect for the beams (as described in §3.4.2).

the basis functions

$$f_n(\theta\ell_{\max}) = \frac{J_{2n+1}(\theta\ell_{\max})}{\theta\ell_{\max}} \quad (3.9)$$

to fit the core of the beams in real space, where the parameter ℓ_{\max} varies the scale of the function's argument and n is a non-negative integer.

Asymptotic Behavior

On scales smaller than a few arcminutes, the noise is subdominant to the planetary signal, and so the beam profiles are well measured even by a single observation of Uranus. However, at larger angular scales, the noise quickly becomes non-negligible. Thus, the asymptotic behavior of the beams cannot be separated from the background without making some assumptions about the shape of the beams far from the core.

The illumination of the ACT optics is controlled by a cryogenic Lyot stop. The beam's Fraunhofer diffraction pattern for monochromatic radiation is described by the squared modulus of the Fourier transform of this circular aperture (Born & Wolf, 1980). The resulting intensity response with angle, or Airy pattern, is

$$A(\theta) = \left[\frac{2J_1[k(D/2)\sin\theta]}{k(D/2)\sin\theta} \right]^2, \quad (3.10)$$

where $k = 2\pi/\lambda$ is the wavenumber.¹⁴ For ACT, $k(D/2) = \pi D/\lambda$ is large, since the primary is several thousand wavelengths across. When the argument of the Bessel function is $\gg 1$ and real, as is the case for $\theta > 2'$, we can make the approximation

$$J_1(x) = \sqrt{\frac{2}{\pi x}} \cos\left(x - \frac{3\pi}{4}\right) \quad (3.11)$$

(Abramowitz & Stegun, 1972). In this regime, the Airy pattern asymptotically approaches

¹⁴Note that this corresponds to the f_0 term from the core model described in §3.4.2.

$$A(\theta) \simeq \frac{8 \cos^2[k(D/2) \sin \theta - 3\pi/4]}{\pi [k(D/2) \sin \theta]^3}, \quad (3.12)$$

such that the envelope of $A(\theta)$ falls as $1/\sin^3 \theta$. This implies that at angles larger than the “near sidelobes” (described in §3.5.2) and neglecting the effects of scattering, the beams should fall as $1/\sin^3 \theta \simeq 1/\theta^3$, since θ is small in the region considered for the beam models.

A different, more general, way to derive this asymptotic behavior of the beams comes from the geometric theory of diffraction from Keller (Keller, 1956, 1959, 1962). As detailed in Page et al. (2003), for an illuminated disk, the diffraction pattern from the rim, regardless of the interior, leads to a gain (response) of the shape

$$G(\theta) \propto \frac{1}{\sin \theta} \left(\frac{1}{\sin \theta/2} \pm \frac{1}{\cos \theta/2} \right)^2 \quad (3.13a)$$

$$\propto \frac{1}{\sin^3 \theta} (1 \pm \sin \theta), \quad (3.13b)$$

for an electric field parallel (+) or perpendicular (−) to the edge, in the far field-limit. For unpolarized light, one can simply average over both polarizations.

The beam behavior derived here assumes perfectly smooth and uniform surfaces within the telescope (i.e. perfect coherence). Nonuniformity, including mirror surface roughness, leads to diffuse scattering (i.e., decoherence of the fields), and thus a reduction in the telescope gain (Ruze, 1966). As described in the next section, an additional term is included in our beam model to account for this effect.

Scattering

The primary telescope surface can be characterized by an rms surface roughness, ϵ , and a correlation length, c . We measured these properties for ACT using photogram-

metry,¹⁵ placing 284 targets on the corners of the primary’s 71 panels where the panels join. The rms of the raw measurements was found to be approximately $30 \mu\text{m}$. However, these measurements are not a representative sample of the rms of the surface, since the targets are at the most extreme points on the corners of the panels. By fitting a polynomial surface to the measurements and looking at the residuals over the entire surface, we find an rms of $\epsilon = 20 \mu\text{m}$. Using the photogrammetry software, the uncertainty on the rms measurements is estimated to be $10 \mu\text{m}$. This does not include any misalignments due to macroscopic deformations of the telescope.

From these measurements, we can also compute the correlation length, c . The correlation function is

$$C(d) = \sum_{ij} z(\vec{r}_i)z(\vec{r}_j) = C_0 e^{-d^2/c^2}, \quad (3.14)$$

where \vec{r} is the position on the surface, z is the mean-subtracted departure from the best-fit designed surface (described in Section 3 of [Fowler et al., 2007](#)), $d = |\vec{r}_i - \vec{r}_j|$ is the separation, and the sum is over all measurement pairs with a separation d . We find that the shape of the correlation function follows the above form reasonably well. By fitting the photogrammetry measurements with Equation 3.14, for the ACT primary, we find $c = 280 \text{ mm}$. This scattering leads to another angular scale (or shoulder) in the beam response, not governed by the diameter of the Lyot stop.

For a beam $B(\theta)$ normalized to unity at $\theta = 0$, the corresponding solid angle Ω is

$$\Omega = 2\pi \int_0^\pi B(\theta) \sin \theta \, d\theta. \quad (3.15)$$

The expression for the gain due to scattering off a rough surface, $G(\theta)$, given by Equation 8 in [Ruze \(1966\)](#)¹⁶ can be re-written in terms of the unit-normalized beam

¹⁵VSTARS by GSI, <https://www.geodetic.com/v-stars/>

¹⁶There is a typo in this equation in the original paper. Inside the summation, the variance δ^2 should be raised to the n^{th} power, as written here, instead of being simply squared, as in [Ruze \(1966\)](#) ([Corkish, 1990](#)).

for an ideal reflector, $B_0(\theta)$, its corresponding solid angle, Ω_0 , the total beam, $B(\theta)$, and its corresponding solid angle, Ω , by using the relation $G(\theta) = (4\pi/\Omega)B(\theta)$. We then obtain the equation

$$B(\theta) = \frac{\Omega}{\Omega_0} B_0(\theta) e^{-\langle \delta^2 \rangle} + S(\theta), \quad (3.16)$$

with the $S(\theta)$ term given by

$$S(\theta) = \frac{\Omega}{4\pi} \left(\frac{2\pi c}{\lambda} \right)^2 e^{-\langle \delta^2 \rangle} \sum_{n=1}^{n=\infty} \frac{\langle \delta^2 \rangle^n}{n \cdot n!} e^{-(c\pi \sin(\theta)/\lambda)^2/n}, \quad (3.17)$$

where $\langle \delta^2 \rangle = (4\pi\epsilon/\lambda)^2$ is the variance in the surface phase (and ϵ and c are the surface rms and correlation length, as described earlier).

The sum converges quickly, with four terms sufficient for our purposes. We have simulated the ACT beam by taking the Fourier transform of a perturbed surface with deformations of a Gaussian shape and verified that we indeed recover the Ruze equation as it is written above.¹⁷

The first term in Equation 3.16 above is expected to decay as $1/\theta^3$, but the same does not hold true for $S(\theta)$, so the inclusion of this term in our fitting procedure is important as it affects how the beam fits are extrapolated. This scattering term $S(\theta)$, which we refer to as the Ruze beam, contains roughly 1.5% of the solid angle of the main beams at 150 GHz. The addition of this term to the beam model is new to the DR4 analysis.

¹⁷More specifically, deformations of a Gaussian shape of width (standard deviation) 250 mm were placed on a square grid of dimension 580 mm on a side to create deformations over a surface. This gave 71 squares inside a diameter of 5500 mm, the effective diameter of ACT's primary. This is a reasonable approximation given that the primary has 71 panels (Thornton et al., 2016). The amplitudes of the deformations were drawn from a normal distribution, and the overall rms of the surface was adjusted to be $\epsilon = 100 \mu\text{m}$, and the modeled surface had $c = 440 \text{ mm}$. The resulting beam followed the Ruze equation for these values of ϵ and c .

Beam Model Fitting

We separate the main beam for each season, array, and frequency into two domains: the core, where we fit the set of basis functions described in §3.4.2, and the wing, where the model is composed of a $1/\theta^3$ term in addition to the Ruze beam (where we use the measured values for ϵ and c described in the previous section). As explained in §3.4.2, a constant offset has already been subtracted from each measured beam profile before combining them into an average radial profile for each season, detector array, and frequency. The full model for each averaged beam profile can be written as

$$B(\theta) = \begin{cases} \sum_{n=0}^{n_{\text{mode}}-1} a_n f_n(\theta \ell_{\text{max}}) & \text{for } \theta \leq \theta_1 \\ \alpha/\theta^3 + S(\theta) & \text{for } \theta_1 < \theta . \end{cases} \quad (3.18)$$

The initial linear least squares fits are performed for the a_n and α for a range of values of ℓ_{max} , θ_1 , and n_{mode} (allowing n_{mode} to vary up to a maximum value of half the number of data points in the core). We first sample a range of values for the scaling parameter ℓ_{max} and then allow this non-linear parameter to vary later on. Based on the optics of the system, at 150 GHz (98 GHz) we expect to have $\ell_{\text{max}} \simeq 2\pi D/\lambda \simeq 19,000$ (12,500). We sample ℓ_{max} in integer steps from 1 to 30,000.

We fit the model out to $10'$, a value which is chosen because this is the radius out to which the beam is measured to sufficient signal-to-noise. Despite the atmosphere subtraction in the target region that extends out to $12'$, residual modes remain, so the signal-to-noise out to $12'$ is not sufficient. By only fitting out to $10'$ we ensure that the fit remains in the region where the measurement is dominated by the beam, not atmosphere.

The model is fit to each averaged radial profile, with its associated covariance matrix. We do not include continuity conditions at θ_1 , but any small discontinuities are well

below ACT’s resolution.¹⁸

In order to select the optimal number n_{mode} for each beam profile, we must strike the right balance between minimizing the χ^2 and not adding too many parameters. To do this, we use the corrected Akaike information criterion (AICc), which estimates the relative quality of models based on both their goodness of fit and their simplicity. For each choice of ℓ_{max} , θ_1 , and n_{mode} , the AICc is computed for the best-fitting model.

The uncorrected AIC is given by

$$\text{AIC} = 2k - 2 \ln(\hat{\mathcal{L}}) , \quad (3.19)$$

where k is the number of estimated parameters in the model and $\hat{\mathcal{L}}$ is the maximum value of the likelihood function for the model (Akaike, 1973, 1974). For small sample sizes, as is the case here, the AIC can exhibit a large bias. To account for this, we use the corrected criterion AICc, where

$$\text{AICc} = \text{AIC} + \frac{2k^2 + 2k}{n - k - 1} , \quad (3.20)$$

and n is the sample size (Hurvich & Tsai, 1989).

We select the values for ℓ_{max} , θ_1 , and n_{mode} corresponding to the lowest AICc.^{19,20} For the DR4 beams, the best-fit values for ℓ_{max} range from 15,604 to 17,936 (from 10,960 to 11,285) at 150 GHz (98 GHz), the best-fit values for θ_1 range from 3.58' to 7.79',

¹⁸For example, the fit shown in the top panel of Figure 3.6 may have a slight discontinuity in amplitude at θ_1 . However, as part of the beam processing, we take the harmonic transform of our beam model and then transform back to radial space to obtain the final beam profiles shown in Figure 3.8. Any small discontinuity that may have been present in the initial fit disappears due to the resolution with which we do the transform.

¹⁹In reality, this is not always precisely the case. When there exists a set of values with n_{mode} smaller by one (hence one fewer parameter) and an AICc that is not significantly different ($\Delta\text{AICc} < 2$), we actually select the model with the slightly higher AICc. We do this to avoid over-fitting the data, since in this case the AICc does not justify the addition of the extra parameter.

²⁰In earlier analyses (Louis et al., 2017; Hasselfield et al., 2013), we simply increased n_{mode} until the reduced- χ^2 fell below 1.

and the best-fit values for n_{mode} range from 9 to 13. We then use MCMC to sample the posterior distribution of the parameters ℓ_{max} , the a_n , and α , assuming uniform priors (Metropolis et al., 1953). This method produces an estimate of the parameter means and the full covariance between all the basis functions and wing parameters, including the non-linear scaling parameter ℓ_{max} .

The term $S(\theta)$ in Equation 3.18 depends on the beam’s solid angle, Ω , but the calculation of Ω depends on the beam model. To get around this issue, the beam fit is performed iteratively, first using a rough estimate for Ω , and then re-computing Ω once a beam model is obtained, and then performing the fit again. In total, we fit the beam four times, re-estimating Ω each time. By the last iteration, the change in Ω becomes undetectable, and so we have converged to a final value for the beam solid angle.

We did test whether the data prefer an asymptotic wing fit term that differs from α/θ^3 , by fitting for α/θ^b instead. While the best-fitting value for b often differs (by 10–20%) from three, it is not strongly constrained by the data. The AICc indicates that the addition of this new parameter is not warranted.

An example of the beam model fit to radial profile data is shown in Figure 3.6, for the S15 PA2 Uranus maps at 150 GHz, together with the residuals. The core functions, α/θ^3 term, and the Ruze beam are indicated in the lower panel, with the core functions used out to a radius of $\theta_1 = 5.59'$, then transitioning to the α/θ^3 plus Ruze beam model at larger radii, and using this model to extrapolate past $10'$. We measure the beam profiles down to -40 dB from the peak, which would leave a few percent of the beam’s solid angle unaccounted for if we did not use our fit to extrapolate past $10'$.

Figure 3.7 shows how the symmetrized model compares to the average Uranus map data for this case. The full set of radial profile fits for the S15 data is shown in Figure 3.8 (with some small additional corrections, as described in §3.4.4). We also

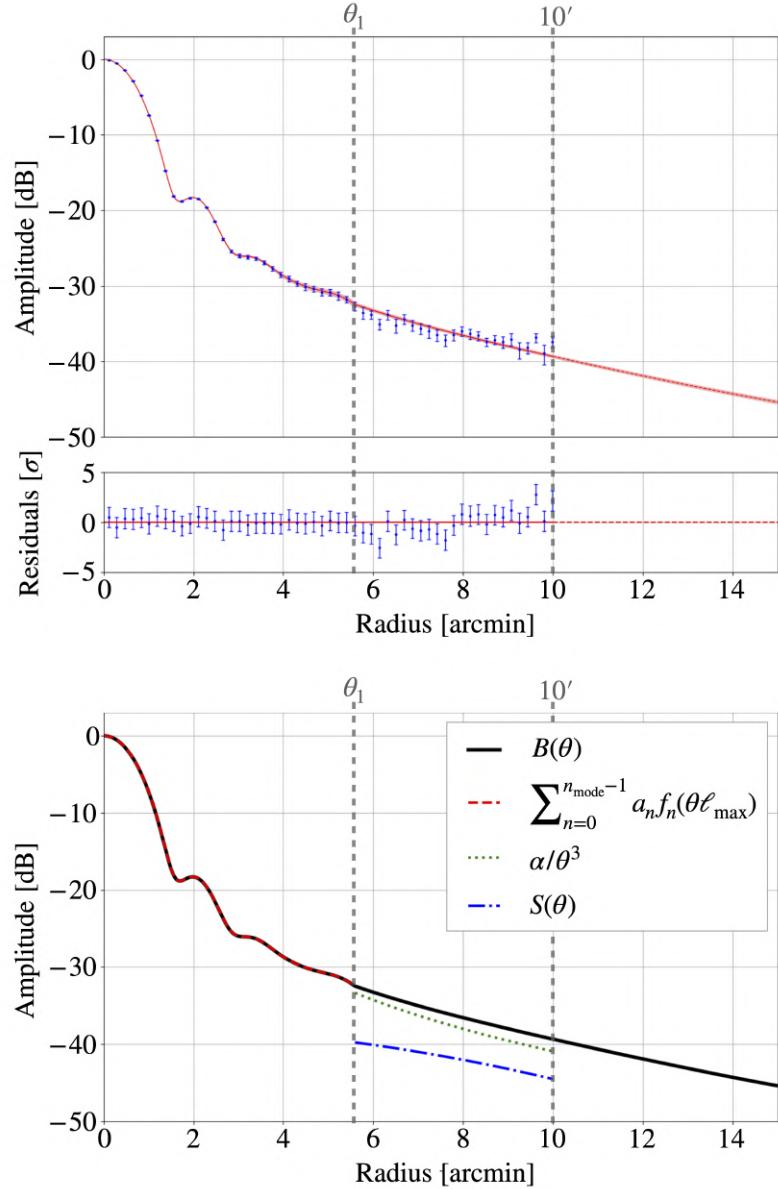


Figure 3.6: Top: Measured radial beam profile (blue) for S15 PA2 at 150 GHz and the best-fit model (red), with the (narrow) red shaded region indicating the 1σ model uncertainty bounds, and residuals shown in the middle panel. Note that the bins are correlated. Bottom: The terms in the model for the beam. Here the final best-fit value of θ_1 , where the model for the beam changes, is $5.59'$. The number of modes in the core fit (n_{mode}) is 13.

obtained model fits for the S13, S14 and S16 data that make up DR4. The solid angles, gains, and FWHMs for all seasons are reported in Table 3.2 (again, with the small corrections from §3.4.4). For reference, $1 \text{ nsr} \simeq 0.0118 \text{ arcmin}^2$.

Even though we make a high signal-to-noise measurement of each beam, the uncertainty on the solid angle is limited to $\sim 2\%$ because of the uncertainty in extrapolation, which depends on the model. The fractional uncertainties in the solid angles at 98 GHz may be larger than those at 150 GHz in part because at 98 GHz the beam is broader, so a larger fraction of the beam is affected by the uncertainty in our extrapolation.

When using the beam solid angle to calculate the flux density of a point source, the uncertainty on the measurement may be reduced by first applying a matched filter to both the beam and the source in order to remove the scales associated with the extrapolation uncertainty.

While much of this fitting method follows the approach used in Hasselfield et al. (2013) and Louis et al. (2017), notable improvements are the addition of the scattering beam term in the model, the use of MCMC sampling that includes estimating the non-linear parameter ℓ_{max} , the exploration of different radii (θ_1) out to which the core model is fit, and the use of the AICc to choose the final best-fit model.

3.4.3 Beam Window Functions

In spherical harmonic space, the beam information is encoded in the harmonic transform b_ℓ and the window function $w_\ell = b_\ell^2$, which describes the instrument's response to different multipoles, ℓ . This window function is an essential component of the DR4 power spectrum analysis in Choi et al. (2020).

The harmonic transform b_ℓ is the Legendre transform, or more accurately the Legen-

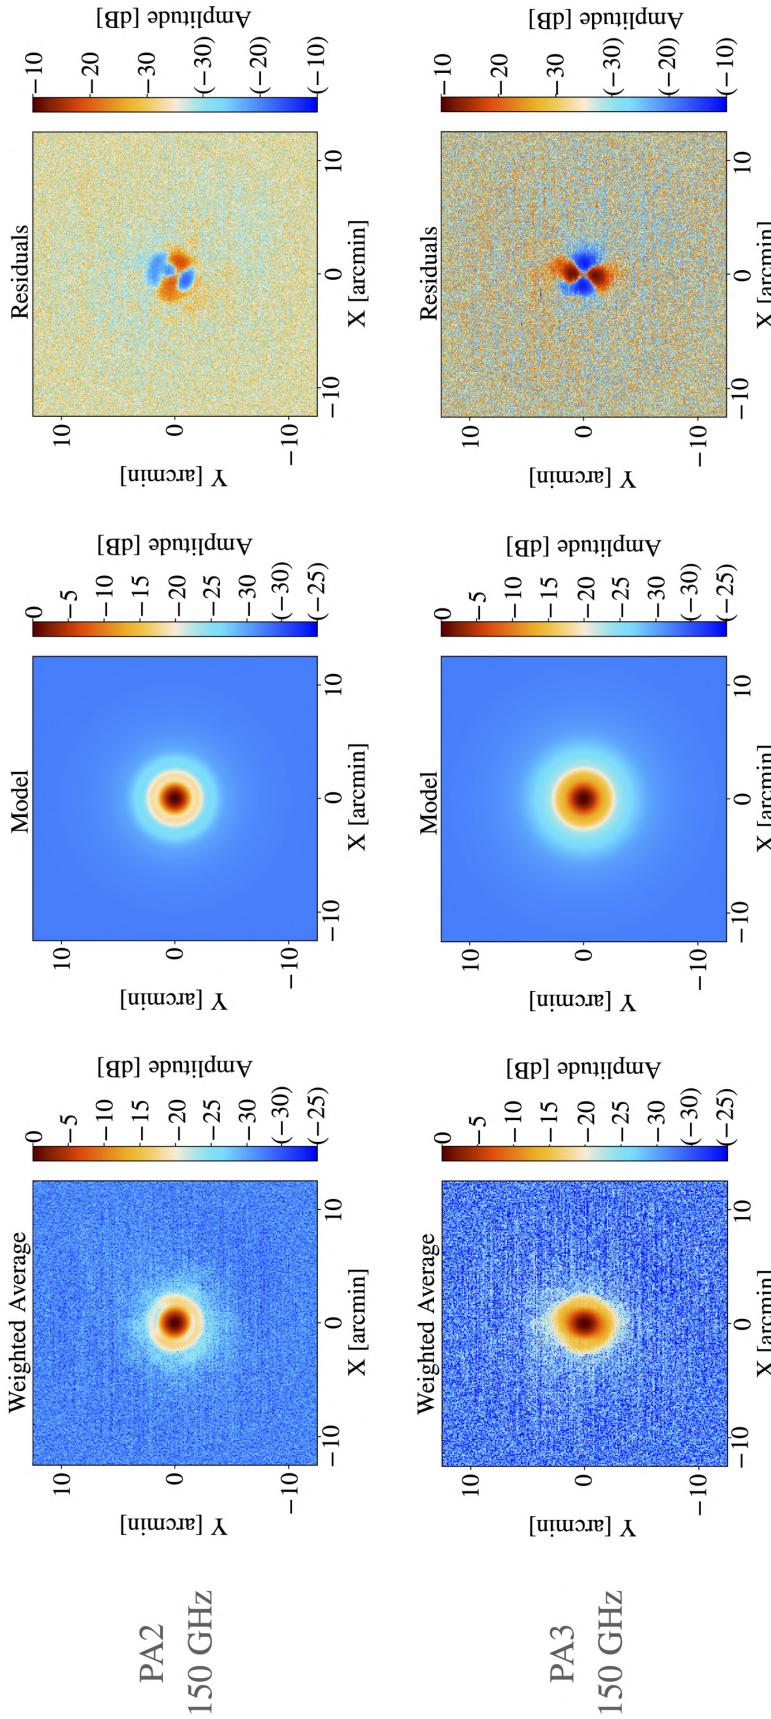


Figure 3.7: Maps of the main beam for S15 PA2 (top) and PA3 (bottom) at 150 GHz. In each case, the color scale is a symmetrical log scale in dB, with a linear threshold of -30 dB and negative values enclosed by parentheses. (Left) The weighted average of the Uranus observations used to characterize the beam. (Center) The model which is fit to the radial profile of the measured beam. (Right) The difference between the measured beam (left) and the beam model (center). Note that the color scale for the residuals is different from the other two maps. Both PA2 and PA3 are shown here, since PA2 serves as an example of a less elliptical beam, and PA3 serves as an example of a more elliptical one. The corresponding maps for S15 PA1 at 150 GHz resemble the ones shown for PA2 here, and the maps for S15 PA3 at 98 GHz are similar to the ones shown here for PA3 at 150 GHz, but broader. For each individual detector array, the residuals are fairly constant from one season to the next. As shown in Figure 3.6 for PA2, the azimuthal average of the residuals is consistent with zero, which is why the fit is successful, despite the residuals visible in the maps. These residuals are expected to have a quadrupole-like shape, since we know our beams are slightly elliptical, and the quadrupole is the dominant asymmetric azimuthal mode for an elliptical beam.

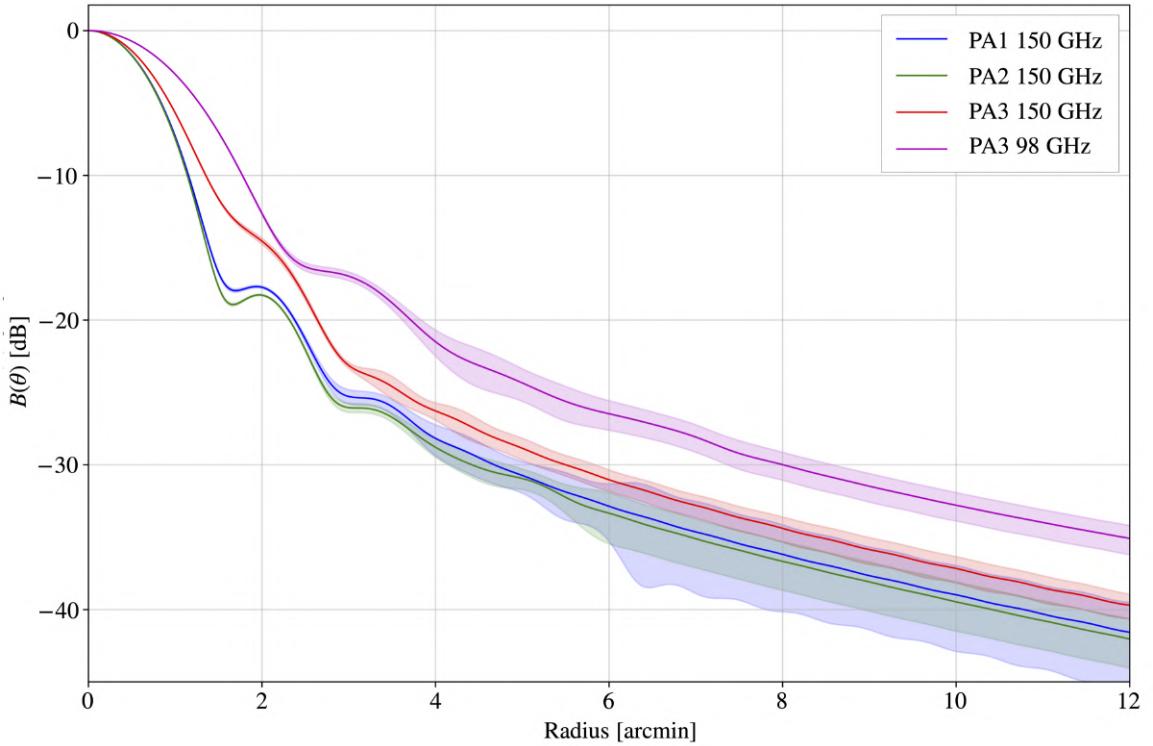


Figure 3.8: The estimated instantaneous beam profiles for S15, for the three arrays (PA1, PA2 and PA3), and for both 98 GHz and 150 GHz for PA3. The shaded bands indicate the 1σ uncertainty bounds. The uncertainties are strongly correlated between radial bins. The S13, S14 and S16 beam profiles are similar, and are included in the public data release. At 150 GHz, the beam for PA3 is quite different from those for PA1 and PA2. This is because PA3 was not focused as well, due to its position, further off-axis, compared to PA1 and PA2.

Table 3.2: Instantaneous beam solid angles, gains, and FWHMs

Array	Band (GHz)	Season	Solid Angle (nsr)	Forward Gain (dBi)	FWHM (arcmin)
PA1	150	S13	201.5 ± 3.8	77.94 ± 0.08	1.330 ± 0.001
		S14	198.5 ± 3.3	78.01 ± 0.07	1.330 ± 0.002
		S15	196.5 ± 8.8	78.06 ± 0.19	1.321 ± 0.002
PA2	150	S14	183.1 ± 3.2	78.37 ± 0.08	1.310 ± 0.001
		S15	187.8 ± 4.7	78.26 ± 0.11	1.311 ± 0.001
		S16	185.4 ± 4.8	78.31 ± 0.11	1.316 ± 0.001
PA3	150	S15	269.6 ± 5.5	76.68 ± 0.09	1.461 ± 0.002
		S16	237.5 ± 8.5	77.24 ± 0.16	1.444 ± 0.003
PA3	98	S15	503.9 ± 21.8	73.97 ± 0.19	2.001 ± 0.004
		S16	476.3 ± 21.8	74.21 ± 0.20	2.002 ± 0.004

dre polynomial transform, of the beam radial profile:

$$b_\ell = \frac{2\pi}{\Omega} \int_{-1}^1 B(\theta) P_\ell(\cos \theta) d(\cos \theta). \quad (3.21)$$

For small beams such as that of ACT, this is effectively a Fourier transform. The derivation of the Legendre transform and details about how the transform is computed are presented in §3.7.2.

We use b_ℓ instead of B_ℓ to indicate the division by Ω , which normalizes b_ℓ to unity at $\ell = 0$ (since $P_0 = 1$). $B_\ell = \Omega b_\ell$ has units of sr, whereas b_ℓ is dimensionless. We extrapolate the model beyond the fit radius of $10'$ when computing the transform. This is necessary to capture the low- ℓ part of the window function, and to account for the part of the beam solid angle that is beyond the range we fit.

A subset of the beam transforms from DR4 is shown in Figure 3.9. A similar figure in [Aiola et al. \(2020\)](#) shows window functions, b_ℓ^2 , which are used to correct the power spectra. For a given array and season, if the beam transforms for PA1 and PA2 are, respectively, b_ℓ^{PA1} and b_ℓ^{PA2} , then for the auto-power spectrum of the PA1 or PA2 maps

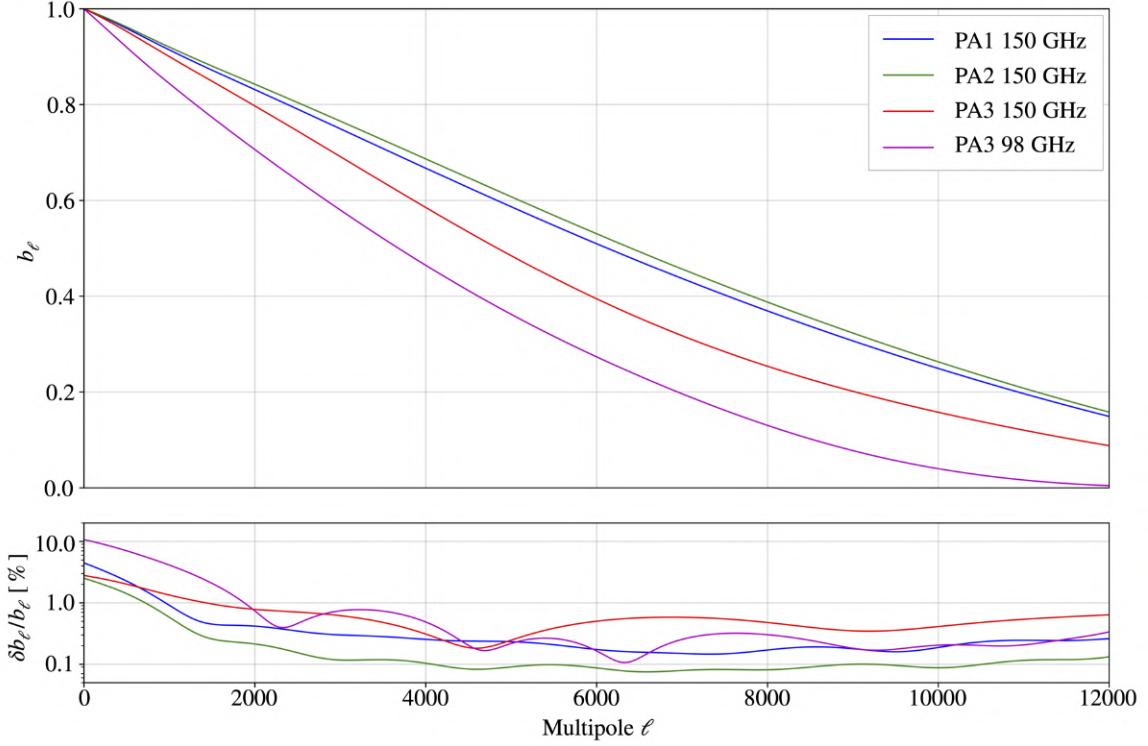


Figure 3.9: Instantaneous beam transforms and their uncertainties for the S15 data, for all the detector arrays. The uncertainties are strongly correlated between multipoles. For context when looking at this figure along with Figure 3.8, in the mapping from angle to multipole ($\ell \simeq \pi/\theta$), $1'$ corresponds to $\ell \simeq 10800$, $1.8'$ corresponds to $\ell \simeq 6000$, and $10.8'$ corresponds to $\ell \simeq 1000$.

the window functions are $(b_\ell^{\text{PA1}})^2$ and $(b_\ell^{\text{PA2}})^2$, and for the cross-spectrum of the PA1 and PA2 maps the window function is $b_\ell^{\text{PA1}} b_\ell^{\text{PA2}}$.

3.4.4 Additional Corrections

The resulting beam models and covariance matrices are an accurate description of the binned radial beam profiles, but they must be corrected for some systematic effects. Following the same approach as in Hasselfield et al. (2013), corrections are applied to account for the pixelization of the planet maps, the binning of the maps into radial annuli, Uranus' angular diameter, and the planet's effective frequency. The effect of each of these corrections is shown in Figure 3.10 for S15 PA2 at 150 GHz, as a typical example.

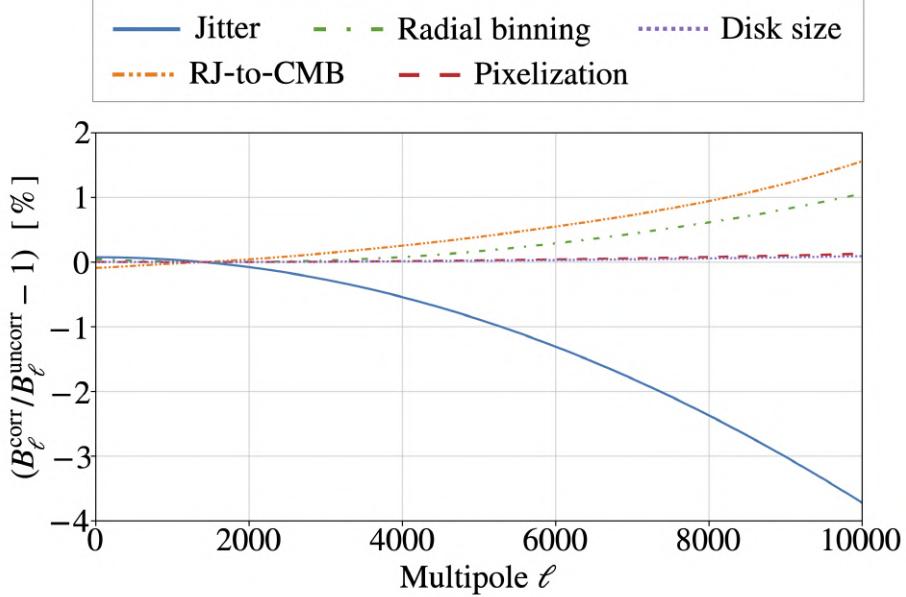


Figure 3.10: Change in the beam transform (in %) due to applying different corrections for S15 PA2 at 150 GHz. For this plot, the beam transforms have all been normalized at $\ell = 1400$. This corresponds roughly to our effective calibration scale, so any change in our beam at this multipole would be corrected by our subsequent calibration to *Planck*.

To correct for the pixelization of the planet maps, we divide the beam transform by the azimuthal average of the pixel window function, $\text{sinc}(pk_x/2\pi) \text{sinc}(pk_y/2\pi)$, where $\{k_x, k_y\}$ are the spatial frequencies and p is the pixel size in radians. This is a $\lesssim 0.1\%$ effect for $\ell < 10,000$.

Then, we estimate and correct for the transfer function induced due to binning the planet maps into radial annuli. This is done by simulating planet maps, using the best-fitting beam profile model as the input, and then estimating their radial profiles following the same procedure as with the data. Comparing the input model with the output radial profile gives an estimate for the transfer function resulting from the radial binning. This is a $\lesssim 1\%$ effect for $\ell < 10,000$.

We also correct for Uranus' angular diameter, since Uranus is large enough that it cannot be treated as a point source given the precision to which we measure the

beam. For each season, array, and frequency, we assume Uranus is a disk with radius equal to the weighted mean of the radii for all the Uranus observations contributing to the beam measurement. We then deconvolve this shape in harmonic space using the function $2J_1(\ell r)/\ell r$, where r is the radius of Uranus' disk in radians. The factor of 2 normalizes the function to unity as $\ell \rightarrow 0$. This is a $\lesssim 0.1\%$ effect for $\ell < 10,000$.

The beam at this point describes the telescope's response to a point-like source with an approximately Rayleigh-Jeans (RJ) spectrum. Near our frequencies, the temperature spectrum of Uranus goes roughly as $\nu^{-0.25}$ (Planck Collaboration Int. LII, 2017). It is sufficiently close to the RJ limit (ν^0) for our purposes. We apply a simple, first-order correction to obtain the relevant beam for the CMB blackbody spectrum using the band effective frequencies from Thornton et al. (2016).²¹ For this radiation with a band effective frequency ν_{CMB} , the beam is taken to be $B'(\ell) = B(\ell\nu_{\text{RJ}}/\nu_{\text{CMB}})$, where ν_{RJ} is the effective frequency for radiation with a Rayleigh-Jeans spectrum. This is a $\lesssim 1.5\%$ effect for $\ell < 10,000$.

The resulting beam models are referred to as the “instantaneous” beams, and it is in fact these corrected beams that are shown in Figures 3.8 and 3.9.

3.4.5 Jitter Beams

In practice, the effective beam for a given sky region, season, detector array, and frequency is broader than the instantaneous beam. This is due to combining observations taken on multiple different nights throughout each season, so the resulting effective beam for each map is not as sharp as the beam inferred from planet obser-

²¹For the foreground modeling in Choi et al. (2020), these effective frequencies were re-computed with improved passband data and upgraded code (as described in §3.7.3). Since the beam analysis for DR4 was done before this new work on the effective frequencies, the values from Thornton et al. (2016) were used. Considering the uncertainties on the passbands, the two sets of effective frequencies are consistent. In addition, given that the uncertainty on the beams is subdominant in the power spectrum analysis, whether one uses the effective frequencies from Thornton et al. (2016) or Choi et al. (2020) for the beam spectral correction does not have a significant effect on the results.

vations which have been carefully recentered before co-adding. Broadening can be caused by variations in the pointing and global alignment, as well as possible small changes in the beam over the course of a season.

As described in [Aiola et al. \(2020\)](#), ACT’s blind pointing accuracy for DR4 is comparable to the average beam FWHM. If left uncorrected, this would significantly broaden the effective beams. Instead, as was done in [Louis et al. \(2017\)](#), we correct the pointing by comparing the observed positions of bright point sources to their known catalog positions. This is done for each 10-minute section of the time-ordered data from each detector array, as described in [Aiola et al. \(2020\)](#). Instead of performing the fit in the time domain as in [Louis et al. \(2017\)](#), due to the larger data volume the fit is now performed in map-space. The resulting fit is obtained more quickly, but is slightly less accurate, leaving a larger residual variation in the beam due to pointing uncertainty.

This residual variation is captured by the effective beam, B_ℓ^{eff} , which is parametrized in terms of the instantaneous beam, B_ℓ , and a correction term that is Gaussian in ℓ , as

$$B_\ell^{\text{eff}} = B_\ell \times e^{-\ell(\ell+1)V/2}. \quad (3.22)$$

If we were to interpret the Gaussian correction term as arising purely from residual pointing errors, then V would be the residual pointing variance in square radians, which is why it is also referred to as “pointing jitter.” This variance, which in practice also includes the additional errors due to alignment and seasonal changes in the beam, is expected to be different for each sky region used in the power spectrum analysis. These different regions (Deep1, Deep5, Deep6, Deep56, Deep8, BOSS, and AdvACT) are shown in Figure 2 of [Choi et al. \(2020\)](#).

To estimate V , in each region we create a catalog of the brightest point sources that are found in the maps using a matched filter, that have a signal-to-noise of at least 10

in each map, and that are matched to catalogs of known sources. For each of these sources, we then obtain an estimate of the variance V and its associated uncertainty for each season, array, and frequency. To do this, we crop a section ($10' \times 10'$) of each map around the source and remove large-scale variations by high-pass filtering the data (by multiplying the data by $1 - G$, where G is a Gaussian filter with a FWHM of $4'$).²² We then compute the effective beam using Equation 3.22, multiply this model by the local pixel-window of the cropped map, transform it into map space, high-pass filter it, and then compute the χ^2 using an estimate of the local white noise amplitude. An example of such an individual source fit is shown in Figure 3.11.

This procedure was validated by verifying that when simulated point sources convolved with Equation 3.22 (with a known value for V) are injected into real CMB data and run through the pipeline, the input value for V is recovered.

For convenience, we exclude sources identified as galaxies, pairs of galaxies, and planetary nebulae in the known source catalogs from this fitting procedure, as these types of sources sometimes appear to ACT to be extended. This leaves approximately 20 sources (for the smaller, deep regions) to 520 (for the largest region) that are included in the fits. They are primarily quasars and radio sources.

Once all N sources have been fit using this method, we use the resulting set of pointing jitter estimates \mathbf{d} and their uncertainties $\boldsymbol{\sigma}$ to obtain an estimate for the mean and intrinsic scatter, V and σ_V , of the effective pointing jitter in each map. The likelihood for this is written as

$$\mathcal{L}(\mathbf{d}, \boldsymbol{\sigma} | V, \sigma_V) = A(\boldsymbol{\sigma}, \sigma_V) e^{-\sum_{i=1}^N (d_i - V)^2 / 2(\sigma_i^2 + \sigma_V^2)}, \quad (3.23)$$

where $A(\boldsymbol{\sigma}, \sigma_V) = \prod_i^N [2(\sigma_i^2 + \sigma_V^2)]^{-1/2}$ is a normalization factor. We estimate the pos-

²²Erratum: In [Lungu et al. \(2022\)](#), the FWHM is quoted as $8'$, but the correct value is $4'$ (which corresponds to 8 pixels of size $0.5'$).

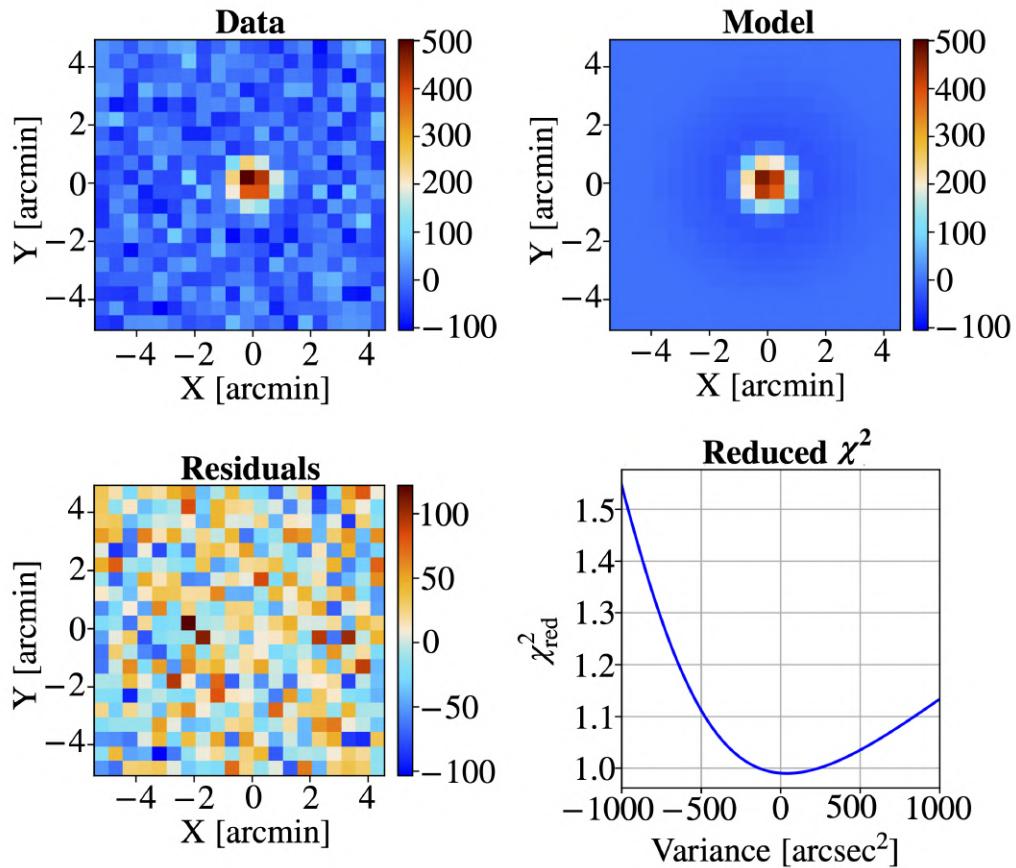


Figure 3.11: Example of the residual pointing variance, estimated by fitting an effective beam to a point source. The top panels compare the flux from the best-fitting model (right) to the data (left), with the residuals (bottom left), and the reduced χ^2 as a function of pointing variance (bottom right). In each case, the data is high-pass filtered, and the colorbars are in units of μK . The best-fit variance for this source is $41 \pm 93 \text{ arcsec}^2$.

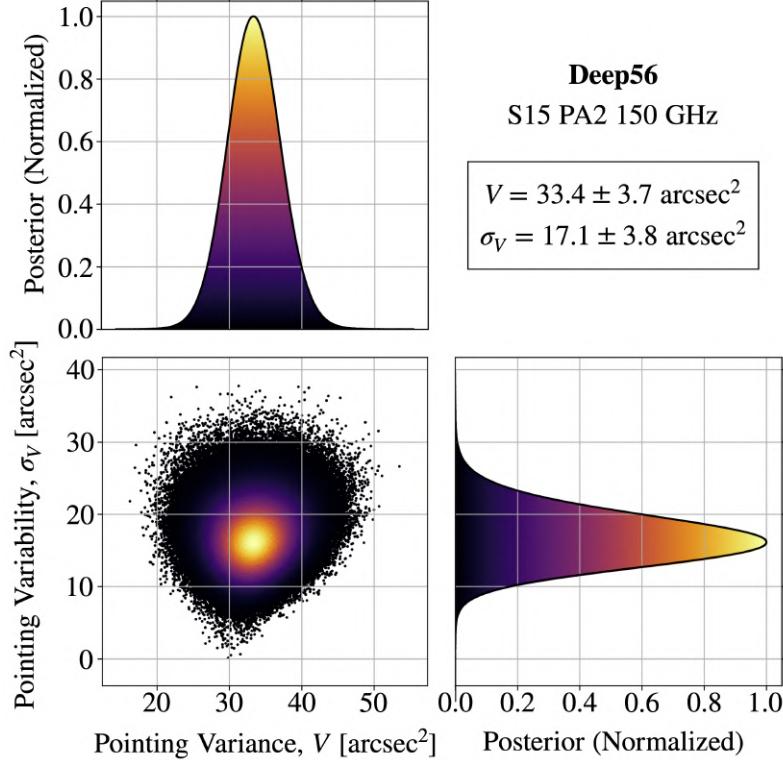


Figure 3.12: Distribution of the residual pointing variance, V , and pointing variability, σ_V , for the S15 PA2 beam at 150 GHz in the Deep56 region, resulting from fitting a pointing variance model to point sources. Notice here that the intrinsic scatter (measured with σ_V) is significantly larger than the measurement uncertainty on V . The color scale represents the normalized posterior, in the range from 0 to 1.

terior distribution for V and σ_V using MCMC with the Metropolis-Hastings algorithm (Hastings, 1970), assuming a uniform prior on V . An example of these posterior distributions for one of the maps is shown in Figure 3.12. This likelihood, which was not used for previous ACT analyses, better takes into account the intrinsic scatter in the residual pointing variance, resulting in an improved estimate.

After computing the pointing variance independently in each map, we use Equation 3.22 to obtain an effective beam, and associated covariance matrix, for each season, region, detector array, and frequency. An example of the effect of the pointing jitter correction on the beam transform is shown in Figure 3.10. The jitter values and uncertainties are given in Table 3.3 and the resulting solid angles and uncertainties

for the effective beams are given in Table 3.4.

The jitter values for PA3 at 98 GHz are significantly greater than those at 150 GHz. This empirical finding serves as a useful reminder that this “jitter” encompasses not only pointing errors, but also any changes in the beam throughout each season. These changes may be greater at 98 GHz, since that is where the beam is broader.

Some of the best-fit jitter values for the Deep8 region are negative, but considering their uncertainties, they are consistent with positive values. We allow for negative jitter values in the fits in order to account for all possible effects on the beams, not just those due to pointing. While we have not investigated whether these negative values are related to a data quality issue, it should be noted that the data from Deep8 were not used in the cosmological analysis due to the poor cross-linking in the region.

In §3.7.3 we provide the effective beam solid angles for beams at different frequencies. We applied the first-order spectral correction from §3.4.4 with effective central frequencies corresponding to synchrotron emission, dust, and the thermal Sunyaev-Zel’dovich effect. These beam solid angles differ from the CMB estimates by 0.3 to 5.3%.

3.4.6 Beam Transform Covariance Matrix

There have been several modifications since Hasselfield et al. (2013) in how we estimate uncertainties in the beam model. Previously, the non-linear scaling parameter ℓ_{\max} was varied in the fits, but the uncertainty in this parameter was not propagated to the covariance matrices. We now estimate the covariance of each beam transform, b_ℓ , by computing the Legendre transform of the sampled posterior distribution for the parameters describing the radial profile, as in Equation 3.21, and then applying the small-order corrections from §3.4.4. We then estimate the covariance matrix from this suite of ℓ -space beam transforms and the added uncertainty associated with the

Table 3.3: Jitter values and uncertainties (arcsec²).

Array	Band	Season	Deep1	Deep5	Deep6	Deep56	Deep8	BOSS	AdvACT
PA1	150 GHz	S13	47.6 ± 18.5	26.6 ± 8.9	30.3 ± 10.2	-	-	-	-
	S14	-	-	-	-	12.1 ± 4.1	-	-	-
	S15	-	-	-	-	34.9 ± 4.4	-2.0 ± 15.4	30.8 ± 3.3	-
PA2	150 GHz	S14	-	-	-	22.5 ± 3.2	-	-	-
	S15	-	-	-	-	33.4 ± 3.7	4.2 ± 6.9	13.3 ± 3.3	-
	S16	-	-	-	-	-	-	-	35.7 ± 2.4
PA3	150 GHz	S15	-	-	-	6.6 ± 5.8	-21.3 ± 24.1	65.1 ± 6.8	-
	S16	-	-	-	-	-	-	-	12.2 ± 4.5
	98 GHz	S15	-	-	-	190.9 ± 7.0	163.2 ± 18.5	246.0 ± 7.2	-
PA3	S16	-	-	-	-	-	-	-	238.2 ± 5.8

Table 3.4: Effective beam solid angles and uncertainties (nsr).

Array	Band	Season	Deep1	Deep5	Deep6	Deep56	Deep8	BOSS	AdvACT
PA1	150 GHz	S13	206.1 \pm 5.0	202.7 \pm 4.1	203.3 \pm 4.2	-	-	-	-
		S14	-	-	-	197.2 \pm 3.4	-	-	-
		S15	-	-	-	199.0 \pm 8.9	193.0 \pm 9.0	198.3 \pm 8.9	-
PA2	150 GHz	S14	-	-	-	183.9 \pm 3.3	-	-	-
		S15	-	-	-	190.3 \pm 4.8	185.7 \pm 4.8	187.1 \pm 4.7	-
		S16	-	-	-	-	-	-	188.2 \pm 4.8
PA3	150 GHz	S15	-	-	-	267.1 \pm 5.5	262.0 \pm 6.9	278.0 \pm 5.8	-
		S16	-	-	-	-	-	-	236.3 \pm 8.5
		S17	-	-	-	510.7 \pm 22.0	506.0 \pm 21.0	520.2 \pm 22.4	-
PA3	98 GHz	S16	-	-	-	-	-	-	490.3 \pm 22.4

jitter correction. These matrices are large, since the beam transforms are computed at each integer ℓ from 0 to 30,000, so we do not store them in their entirety. Rather, we decompose each matrix into independent modes (via singular value decomposition) and store the largest 10 modes. We find that this is sufficient to capture the majority of the covariance (we do not discard any modes with a singular value larger than 10^{-3} of the maximum value).

We include additional modes to account for possible uncertainty due to variations in the surface rms ϵ (which enters into the calculation of the Ruze beam, as described in §3.4.2, and was fixed to our estimate of 20 μm),²³ as well as the range over which the model for each beam offset is fit (which initially was 3.5'–10.0'). For the surface rms, the values explored are 20 μm and 30 μm . For the region over which the offsets are fit, the three independent ranges explored are 3.5'–5.0', 5.0'–7.0', and 7.0'–10.0'. We store the top 3 modes associated with these model variations. The final uncertainty for each beam is thus composed of a total of 13 modes. These beam uncertainties are later added to the data covariance matrix as part of the power spectrum analysis pipeline.

An example of the effect of these additional modes on the beam transform uncertainties is shown in Figure 3.13. The significant increase in the uncertainty at low multipoles is mainly due to the inclusion of the different fit ranges for the offsets. Previously, as described in Hasselfield et al. (2013), the uncertainties were simply doubled from their formal values to account for potential systematic variations due to different fitting ranges. Even though this earlier method lead to a smaller estimate of the beam uncertainties at low multipoles, this did not have a significant effect on our results for DR3, since the uncertainty on the beams is subdominant in the power spectrum analysis, and a low- ℓ cutoff of 500 (350) was applied to the TT (TE and

²³We do not account for possible variations in the other measured parameter for the Ruze beam, the correlation length c , since it is well constrained by our measurements.

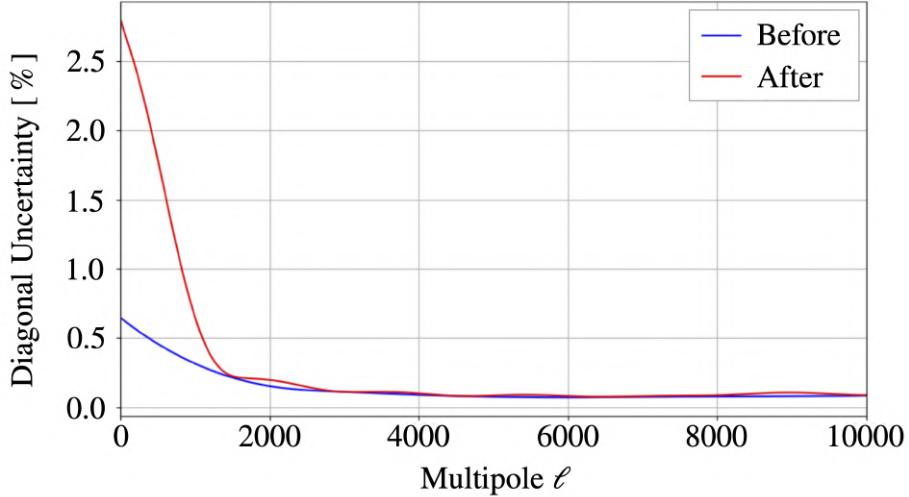


Figure 3.13: Uncertainties on the beam transform (in %) for S15 PA2 at 150 GHz, before (blue) and after (red) adding the additional modes to account for different surface rms values and different fit ranges for the beam offset levels. These uncertainties are included in Figure 3.9.

EE) data.

Another difference compared to the DR3 analysis in [Louis et al. \(2017\)](#) is the treatment of calibration uncertainty. As described in [Choi et al. \(2020\)](#), for each season, region, array, and frequency the angular power spectra from ACT are calibrated to the *Planck* temperature maps in the range $600 < \ell < 1800$. The [Louis et al. \(2017\)](#) analysis factored out the beam amplitude and uncertainty at an “effective” calibration scale (chosen in that case to be $\ell = 1400$).²⁴ However, we have changed this procedure for DR4 to reflect the fact that we are not calibrating the data at a single ℓ , but over a range of ℓ values. We now simply normalize the beam to unity at $\theta = 0$ and treat the calibration, and its uncertainty, separately in the ACT analysis. This is why the fractional beam uncertainties in Figure 3.9 are no longer smallest at $\ell = 1400$, differing from those reported for ACT DR3.

²⁴This procedure is described by Equation 11 of [Hasselfield et al. \(2013\)](#).

3.5 Polarization

3.5.1 Main Beam

Measurements of Uranus in polarization at visible and near-infrared wavelengths have shown that its disk-integrated polarization is less than 0.05% (Schmid et al., 2006). While we are not aware of any similarly precise data at millimeter wavelengths, it is expected that the relevant scattering effects in the planetary atmosphere would be much weaker, resulting in net polarization levels considerably lower than the measurement cited above. Measurements by *Planck* at 100 and 143 GHz place 68% (95%) confidence upper limits on the polarization fraction of Uranus at 2.6% (3.6%) and 1.5% (2.0%), respectively (Planck Collaboration Int. LII, 2017).

Since we do not expect Uranus to be significantly polarized in the bands observed by ACT, we interpret any polarization response measured from Uranus as being due to temperature-to-polarization (T-to-P) leakage. Although this leakage is relatively small in magnitude, the ACT DR4 data are now sensitive enough that we must account for it in our analysis. To do this, we use observations of Uranus to build an ℓ -space T-to-P leakage function for each season, array, and frequency. This measurement and correction for the leakage in the main part of our beams is new to DR4 (see also Aiola et al., 2020; Choi et al., 2020).

We begin by making maps of the Q and U Stokes parameters for the same set of Uranus observations chosen in §3.4.1 to fit the main beam. We then convert each set of $\{Q, U\}$ polarization maps to the radial Stokes parameters $\{Q_r, U_r\}$ (following the third definition of cross polarization in Ludwig, 1973) using the flat-sky approximation:

$$Q_r(\boldsymbol{\theta}) = Q(\boldsymbol{\theta}) \cos 2\phi_\theta + U(\boldsymbol{\theta}) \sin 2\phi_\theta \quad (3.24a)$$

$$U_r(\boldsymbol{\theta}) = U(\boldsymbol{\theta}) \cos 2\phi_\theta - Q(\boldsymbol{\theta}) \sin 2\phi_\theta , \quad (3.24b)$$

where $\boldsymbol{\theta} \equiv (\theta, \phi_\theta)$ are standard polar coordinates with the beam centroid as their origin and ϕ_θ increases clockwise from the positive y -axis (assuming that one uses the convention in which $+x$ points to the right and $+y$ points upward). Here Q and U follow the COSMO convention (Górski et al., 2005), whereas for the ACT maps released as part of DR4, the polarization components are defined by the IAU convention (Hamaker & Bregman, 1996). The polarization convention initially used for $\{Q, U\}$ does not matter once we have transformed to $\{Q_r, U_r\}$.

While it is not obvious how the polarized E or B beams ought to behave at larger radii, we do have some intuition about Q and U . For an unpolarized point source such as Uranus, any polarized signal will be the result of beam differences between the two axes of a polarimeter. Although the beams may be slightly different, due to e.g., differential ellipticity, we still expect each of them to decay radially as $1/\theta^3$ as they are part of the same diffraction-limited optical system. The Q and U maps are essentially just radially independent linear combinations of the various detector axes in an array, appropriately weighted by the map-maker, so the same asymptotic behavior should apply. This asymptotic behavior would then hold true for the Q_r and U_r maps as well. This was confirmed with simulations where the beams for individual polarimeter axes were constructed using an “elliptified” version of the azimuthally averaged intensity beam for a given season and array. We thus use the same basis functions in the core and the α/θ^3 term in the wing to fit the beam in polarization as we did for the main beam in §3.4.2.

Since we are ultimately interested in how leakage manifests itself in the angular power spectra, we need to translate any polarized beam models of the azimuthally averaged radial profiles of Q_r and U_r , \tilde{Q}_r and \tilde{U}_r , to an ℓ -space representation of E and B . Conveniently, \tilde{Q}_r and \tilde{U}_r have a direct correspondence to the azimuthally averaged ℓ -space E and B transforms, $\tilde{E}(\ell)$ and $\tilde{B}(\ell)$. As shown in §3.7.4, there exists a

simple relation between these components, which in the flat-sky approximation takes the form of a second-order Hankel transform:

$$\{\tilde{E}(\ell), \tilde{B}(\ell)\} = -2\pi \int \{\tilde{Q}_r(\theta), \tilde{U}_r(\theta)\} J_2(\ell\theta) \theta d\theta. \quad (3.25)$$

From the set of Uranus maps in the $\{Q_r, U_r\}$ basis, examples of which are shown in Figure 3.14, we thus construct average radial profiles for each season, array, and frequency, and we fit them in a similar way as the main beam is fit in §3.4. While for the usual temperature beam fitting pipeline we fit an offset to each individual Uranus profile before taking an average, we do not do this in polarization since we expect the mapping transfer function in that case to be zero.²⁵ The only difference in the fitting procedure here is that we do not include a scattering term in the polarized beam model, since it is not expected to matter to first order, and it is unclear how it would vary for the different detector polarizations.²⁶ Examples of the radial profile model fits in polarization are shown in Figure 3.15.

As can be seen in Figures 3.14 and 3.15, while our model is a good fit to the Q_r and U_r radial profiles, there are significant, quadrupole-like residuals in the maps when our model is subtracted. It turns out that the leakage beams are far less azimuthally symmetric than the temperature beams (as can be seen by comparing Figures 3.7 and 3.14). The treatment of this leakage will be revisited for upcoming ACT beam analyses and may be improved upon, for example, by fitting a 2D model to the polarized beam profiles in order to properly capture the asymmetry.²⁷ In the meantime, the fitting done here is sufficient. The level of residuals seen in Figure

²⁵Even if the mapping transfer function were non-zero in polarization, it would be a sub-percent-level effect in the measurement of the leakage beam, which itself is percent-level in terms of our power spectrum analysis and results, so the effect would be insignificant.

²⁶Again, this would be a sub-percent-level effect in the measurement of the percent-level leakage beams, so there would be no significant impact on our analysis and results.

²⁷Work on fitting the beams in 2D is in progress.

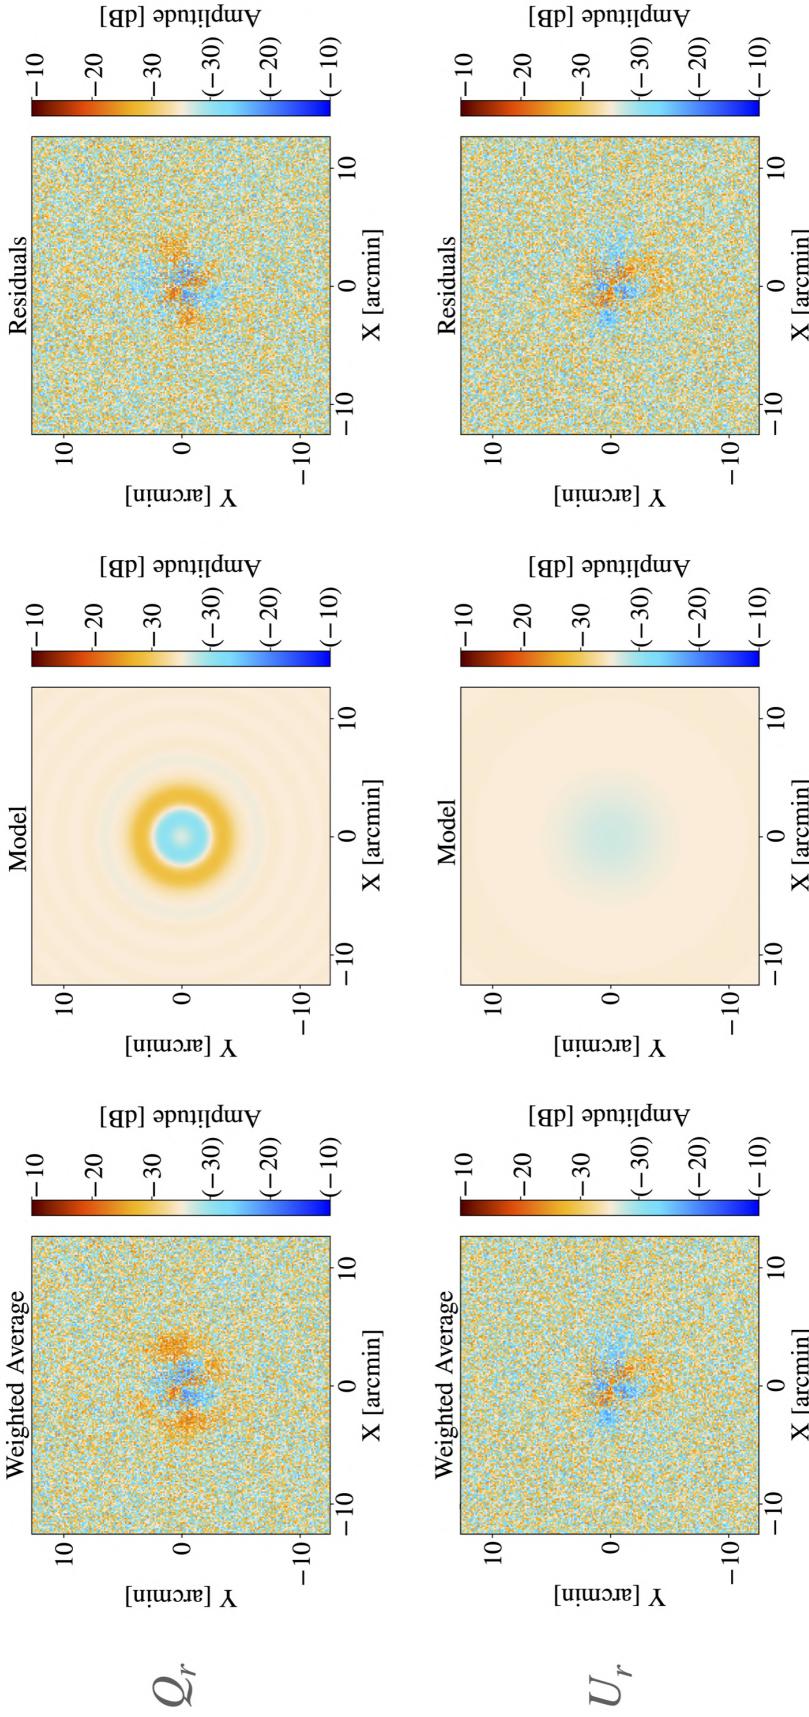


Figure 3.14: Maps of the leakage beam in Q_r (top) and U_r (bottom) for S15 PA2 at 150 GHz. In each case, the color scale is a symmetrical log scale in dB, with a linear threshold of -30 dB and negative values enclosed by parentheses. (Left) The weighted average of the Uranus observations used to characterize the beam. (Center) The model that is fit to the radial profile of the measured beam. (Right) The difference between the measured beam (left) and the beam model (center). Features in the residuals could arise due to physical effects, such as differential beam ellipticity between the two axes of a polarimeter or a polarization angle offset (Hu et al., 2003). The level of residuals seen in the last column has an insignificant effect on the results of the power spectrum analysis. As shown in Figure 3.15, the azimuthal averages of the residuals are close to zero, which is why the fits are successful.

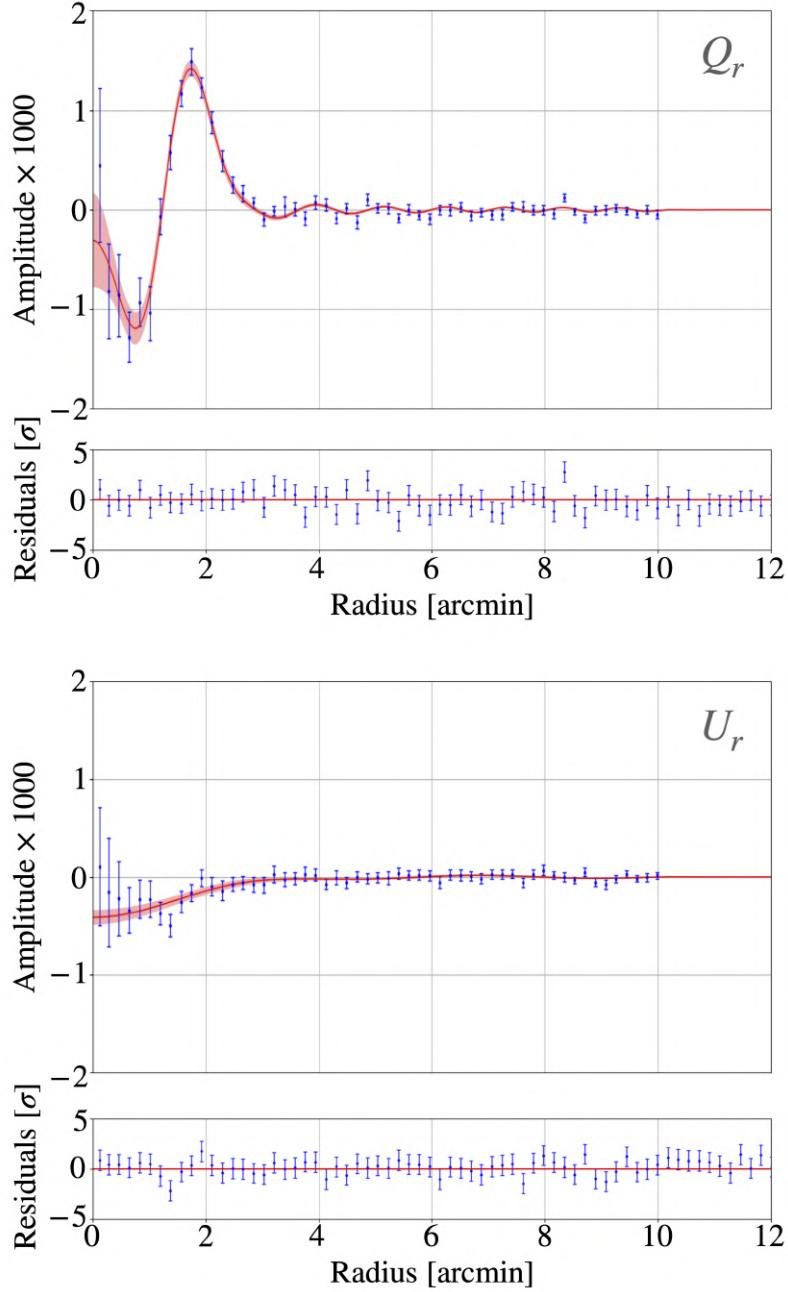


Figure 3.15: Measured radial polarization profiles (blue) in Q_r (top) and U_r (bottom) for S15 PA2 at 150 GHz and the model we fit to each profile (red), with the red shaded region indicating the 1σ model uncertainty bounds. The profiles have been normalized by the peak amplitude of the corresponding T beam (at the beam centroid, $\theta = 0$). Note that the bins are correlated.

[3.14](#) has an insignificant effect on the results of the power spectrum analysis for DR4 (see [§3.6.2](#)). The features visible in the residual maps can be expected to occur due to physical effects. For example, we simulated the T-to-P leakage for a point source due to both differential beam ellipticity between the two axes of a polarimeter and a polarization angle offset. The resulting maps of the simulated leakage beam in Q_r and U_r had strong quadrupole-like features. As shown in Figure [3.15](#), the azimuthal averages of the residuals in Q_r and U_r are close to zero, which is why the fits works well, despite the appearance of the residuals in the maps.

The Q_r and U_r radial profiles fits and their uncertainties are transformed to ℓ -space using a similar approach as for the temperature data, except using Equation [3.25](#) instead of the usual Legendre transform. The transforms for our example case are shown in Figure [3.16](#). These transforms do not include corrections for small systematic effects, as was done for the main beam fitting pipeline in [§3.4.4](#). This is because we are ultimately interested in the ratios of transfer functions for the leakage beams, in which these ℓ -space corrections cancel out. To determine the polarization leakage beams, an example of which is shown in Figure [3.17](#), we divide the E and B transforms by their corresponding (uncorrected) T transform. For PA1 and PA2 the leakage values are within 1.5% (comparable to the leakage values for SPT-3G, [Dutcher et al., 2021](#)), whereas for PA3, the leakage increases around $\ell \simeq 4000 - 6000$, and reaches over 6% (4%) at $\ell \sim 10,000$ at 150 GHz (98 GHz). We attribute the higher leakage for PA3 to an imperfectly optimized horn design in this first generation of multichroic polarimeters.

The top 10 modes from the leakage beam covariance matrices are stored. In addition, similar to the main beam analysis, two modes are added to account for variations in the model for the temperature beam.²⁸ The final leakage beam uncertainties are thus

²⁸We only consider variations in the range over which the beam offsets are fit ($3.5' - 5.0'$, $5.0' - 7.0'$, $7.0' - 10.0'$). The Ruze beam is a smaller effect, so it is (safely) ignored here.

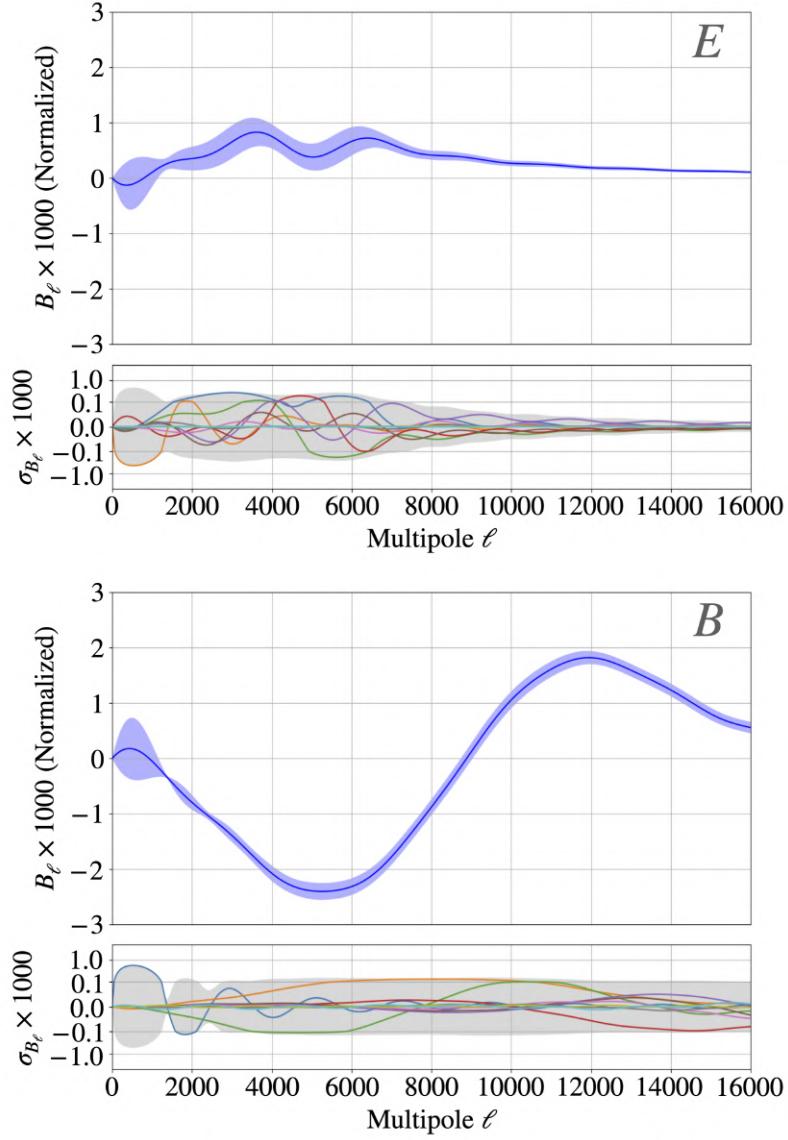


Figure 3.16: The transform of the radial beam profiles in E (top) and B (bottom) for S15 PA2 at 150 GHz. These transforms are normalized by the amplitude of the corresponding T beam transform at $\ell = 0$. The lower panels show the dominant independent modes of each transform's covariance matrix (in color) and the magnitude of the diagonals (the gray shaded regions).

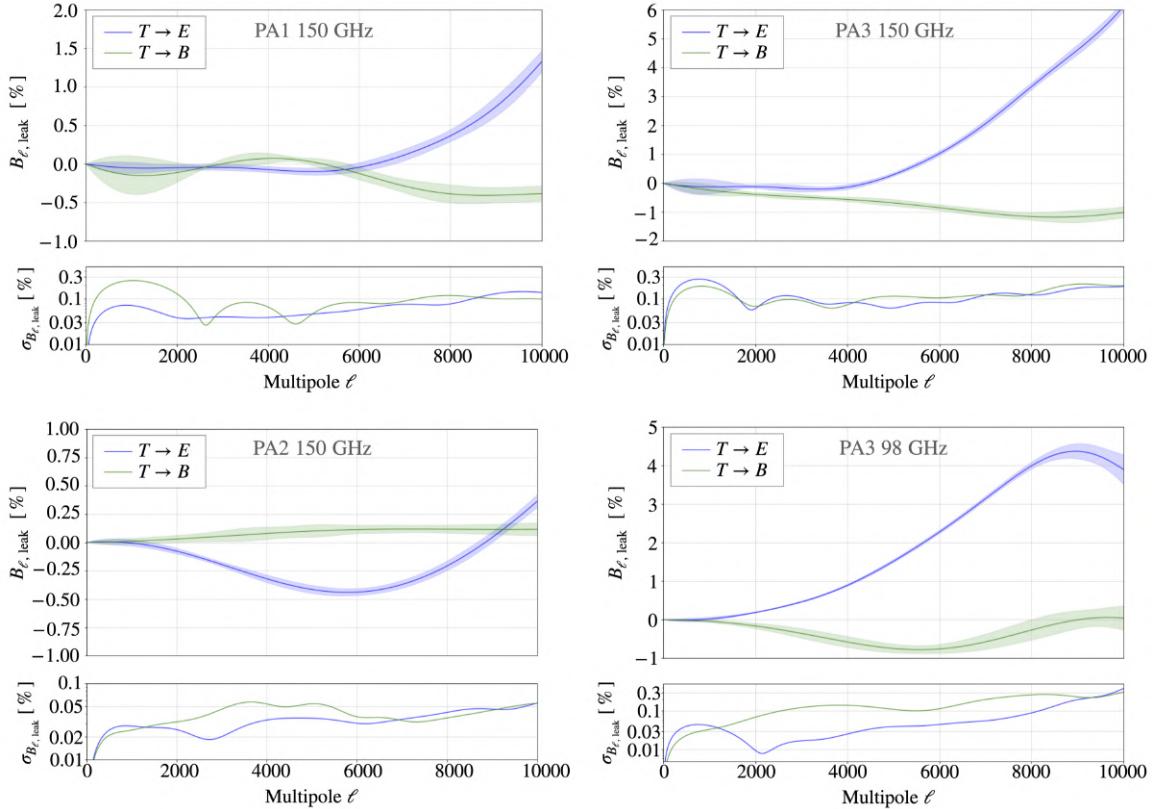


Figure 3.17: Measured temperature-to-polarization leakage in the main beams for S15 for all arrays, showing the temperature leaking into E -mode and B -mode polarization. Note the different y -axis ranges. For PA1 and PA2 the leakage values are within 1.5%, whereas for PA3, the leakage increases around $\ell \simeq 4000 - 6000$, and reaches over 6% (4%) at $\ell \sim 10,000$ at 150 GHz (98 GHz). We attribute the higher leakage for PA3 to an imperfectly optimized horn design in this first generation of multichroic polarimeters. The uncertainties (lower panels) include the adjustments for model variability.

comprised of 12 modes.

The leakage beams were used in the DR4 power spectrum likelihood, as explained in §3.5.3.

3.5.2 Polarized Sidelobes

As described in [Louis et al. \(2017\)](#) and [Aiola et al. \(2020\)](#), we detect polarized sidelobes of the main ACT beams. Although weak in amplitude, these sidelobes cause noticeable T-to-P leakage. The sidelobes for PA1 and PA2 were shown in [Louis et al. \(2017\)](#). While PA3 was not used in that analysis, it was mentioned at the time that polarized sidelobes were not detected for PA3. However, using additional observations of Saturn to conduct a more thorough analysis, we have detected sidelobes in PA3 at 150 GHz, with an amplitude roughly 10% that of the sidelobes in PA1 and PA2. The general features of the sidelobes are common to all detector arrays; they consist of a group of compact lobes, each resembling a slightly elongated image of the main beam, with approximate four-fold symmetry, strongly polarized perpendicular to the radius from the beam center (which corresponds to $-Q_r$). This results in most of the leakage due to the sidelobes being from temperature into *E*-mode polarization rather than into *B*-mode polarization. These sidelobes are stable in time. We also observe that sidelobes from Saturn are only seen if Saturn lies in a focal plane's field of view. As explained in Section 3.8 of [Aiola et al. \(2020\)](#), this is consistent with the sidelobes being due to an optical effect inside the receiver.

The sidelobes for PA1 and PA2 are at a distance of roughly 15' from the beam centroid, with an additional set visible at 30' for PA1. The sidelobes for PA3 are located approximately 30' to 40' from the beam centroid. This larger angular separation means the strongest T-to-P leakage occurs at $\ell \simeq 300$ for PA3 compared to $\ell \simeq 500$ for PA1 and PA2. Also, since the sidelobes only map to the sky when the main beam is

also in the field of view, fewer detectors are affected by each sidelobe for PA3. We do not detect any sidelobes in PA3 at 98 GHz, which implies they must either be of significantly lower amplitude than those at 150 GHz, or in a different position. For PA1, PA2, and PA3, the amount of solid angle contained in these sidelobes is roughly 1.2%, 3.4%, and 0.15%, respectively, of that in the main beam.

Preliminary studies suggest that these sidelobes are due to diffraction caused by the arrays of metal elements that make up ACT’s optical filters. We note that similar filters (Ade et al., 2006) are also used for the South Pole Telescope (SPT), POLAR-BEAR, and BICEP/Keck (Padin et al., 2008; Arnold et al., 2010; Keating et al., 2003) at varying locations in the optical path.²⁹

If the sidelobes were indeed due to diffraction from the filters, we would expect them to occur at a larger radius for lower frequencies. Based on our calculations, at 98 GHz, the sidelobes would appear starting at a distance of approximately 47' from the beam centroid. Since this is roughly the size of our field of view, the sidelobes would then map to the sky for only a small fraction of the detectors, the ones at the edges of the focal plane. This is consistent with the lack of detectable sidelobes at 98 GHz. We do not apply any sidelobe corrections to the PA3 data at 98 GHz.

As mentioned in Aiola et al. (2020), to study the sidelobes we use observations of Saturn. While Saturn’s brightness appears to induce a non-linear response near peak amplitude, it is useful for studying the relatively weak sidelobes. There are no issues due to non-linearity when observing the sidelobes. This is confirmed by the fact that the amplitude of the sidelobes seen by a detector is consistent, whether they are seen before or after the detector sees Saturn’s peak.

We apply the same treatment to the sidelobes for all detectors at 150 GHz. In short,

²⁹This issue with ACT is related to the location of the filters in the optical path and the diffraction angle produced by the filters relative to the effective view passed by the system at this location. So other experiments may or may not see such an effect, depending on the details of their optical implementation.

we model the sidelobes as a sum of polarized, spatially shifted copies of the main beam and fit the amplitudes of the T , Q , and U components of each beam instance using maps of Saturn. We then use this model to deproject the sidelobes from the time-ordered data prior to map-making. The idea is to subtract the total flux in the sidelobes, even if their shapes are not exactly zeroed out in the maps. It can be shown that this removes the low- ℓ T-to-P leakage. Maps of the sidelobes are shown in Figure 3.18, along with the T-to-P leakage functions.

For PA1 and PA2, we re-use the sidelobe models from DR3, constructed using observations of Saturn from 2014. For PA3, which was not part of DR3, we use Saturn observations from 2015 to characterize the sidelobes for DR4 in the same way as was done for DR3.

Looking more closely at how these sidelobe models are constructed, we begin with a series of observations of Saturn to which we apply the same data selection criteria as we did to Uranus, as described in §3.2. We then map the chosen observations with the `moby2`³⁰ map-maker (the same map-maker that was used for the beam analysis in [Louis et al., 2017](#)) and coadd the maps together to produce one map of Saturn for each detector array (PA1, PA2, PA3). The maps are coadded with weights based on an estimate of their white noise level determined outside the planet region.

We then visually identify regions in the maps containing compact polarized sidelobes. As can be seen in Figure 3.18, the pattern of the sidelobes in each map has approximate four-fold symmetry, with four groups of sidelobes appearing at roughly equal distances from the beam centroid. For each sidelobe we choose how many copies of the main beam should be used to model it. Stronger, elongated sidelobes are modelled by two copies of the main beam, to capture the elongation; weaker sidelobes are modelled with a single copy of the beam. Then for each of the four groups of sidelobes

³⁰GitHub repository:

<https://github.com/ACTCollaboration/moby2>

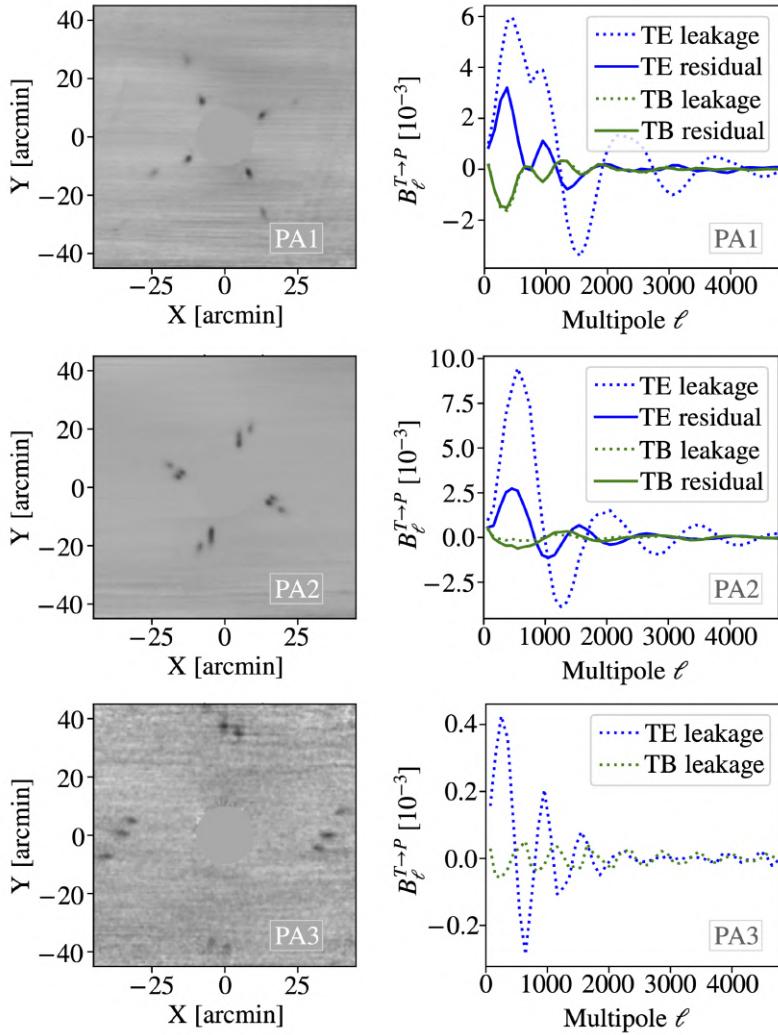


Figure 3.18: Sidelobes in the PA1, PA2, and PA3 detectors at 150 GHz. (Left) Maps of the polarized sidelobes, obtained by stacking observations of Saturn. The main part of the beam at the center of each image is masked. In each the grayscale is linear, indicating the sidelobe amplitude in the range from -0.002 (black) to $+0.001$ (white) relative to the main beam peak,^a with positive (negative) numbers corresponding to polarization parallel (perpendicular) to the radial direction. The complementary polarization leakage (corresponding to TB leakage) is smaller and not shown in the maps, but it is also estimated. (Right) The effect of the sidelobes on our measurements, expressed as beam transfer functions $B_\ell^{T \rightarrow E}$ and $B_\ell^{T \rightarrow B}$. Note that the scales on the right differ for all three arrays. The sidelobes are projected out of the time-ordered data prior to mapping, so the effective (residual) leakage is what affects our spectra. Since the sidelobes for PA3 are significantly weaker than PA1 and PA2, we do not compute residuals for PA3. No sidelobes are seen for PA3 at 98 GHz.

^aSaturn's brightness appears to compress the detector gain by a few percent. The effect can be ignored here because it is a small correction on the sidelobe leakage, which itself is a small effect.

we fit the position and amplitude of the copies of the main beam. In theory this fit could be done with any of the signals, but currently we perform this fit to a map of $P = \sqrt{Q^2 + U^2}$ because that is a bright, clear signal. With the model positions fixed, we then fit the amplitudes for each signal (T, Q_r, U_r) independently. The result is a base model for the sidelobes, but this model does not yet contain per-detector detail.

We next account for the fact that the sidelobes do not always appear in all the detectors. This is relevant to the extent that the per-detector weights in the planet maps are different from the per-detector weights in the survey maps (which are used for the CMB analysis, for example). This could possibly be a significant (up to tens of percent) effect, since the `moby2` map-maker used to make the planet maps doesn't weight the detectors by their noise, whereas the `enki` map-maker used for the survey maps does. The sidelobes occur primarily in the detectors at the periphery of the array, which is also where the noise tends to be higher, so this is a correlated effect.

To deproject the sidelobes from the per-detector time-ordered data, we need to estimate whether or not the sidelobes are seen by individual detectors, based on their position on the focal plane. The parameter we want to estimate is the radius of the “aperture”, or circle, around a detector such that if Saturn falls within the circle, the beam sidelobes are visible to that detector. To estimate this radius, we first divide the detectors into subsets. For PA3, the subsets are the three hexagonal wafers of detectors, described in [Thornton et al. \(2016\)](#). For each detector subset, we re-make maps of Saturn, and measure the total sidelobe flux in the resulting coadded map. We then compare these measurements to a model for the sidelobe flux as a function of radius, which scales with the fraction of detectors in each subset that would see each sidelobe.

As for the main beam analysis in §3.4.4, small corrections are made to the model to account for systematic effects.

The sidelobe removal described here is not perfect, so there is still residual T-to-P leakage in the ACT data due to the sidelobes. The residual TE and TB leakages, shown in Figure 3.18, are estimated by making new maps of Saturn after our sidelobe removal. Since the sidelobes for PA3 are already significantly weaker than PA1 and PA2, we do not compute residuals for PA3. We add the residuals for PA1 and PA2 to the main beam leakage for use in the power spectrum analysis, as described in the next section.

The residual sidelobe leakage in Figure 3.18 can be compared to the leakage in the main beam shown in Figure 3.17. The effect of each of these components on the final spectra (if there is no leakage correction as in §3.5.3) is shown in Figure 3.19. At low ℓ , the residual sidelobe leakage dominates, whereas the main beam leakage grows larger by $\ell \sim 2000$, and dominates at high ℓ .

3.5.3 Leakage Correction

The main beam leakage and the residual leakage from the sidelobes are included in the power spectrum likelihood (see Section 12 of Choi et al., 2020) by making use of a leakage-corrected model for the TE and EE theory spectra, computed each time the likelihood is estimated. The corrected model spectra, $T_i E'_j$ and $E_i E'_j$, are related to the input theory spectra, $T_i T_j$, $T_i E_j$ and $E_i E_j$, via

$$T_i E'_j = T_i E_j + T_i T_j \gamma_j \quad (3.26a)$$

$$E_i E'_j = E_i E_j + T_i E_j \gamma_i + T_j E_i \gamma_j + T_i T_j \gamma_i \gamma_j . \quad (3.26b)$$

Here the i and j subscripts denote different spectra and the γ factors encode the ℓ -dependent leakage (the sum of the main beam leakage from §3.5.1 and the residual sidelobe leakage from §3.5.2) and are shown in Figure 3.20. At 150 (98) GHz, the

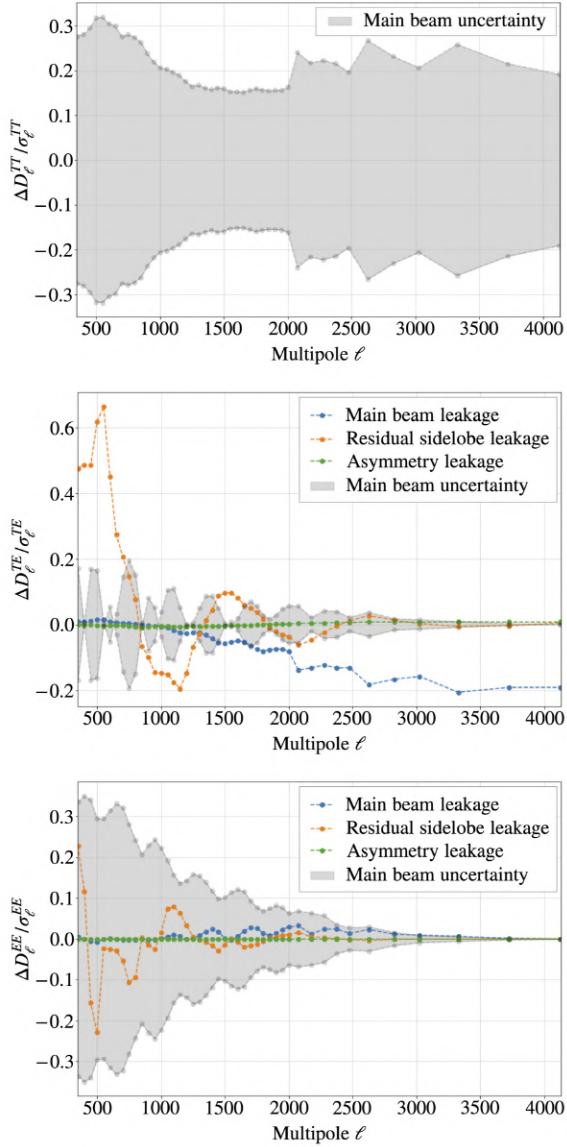


Figure 3.19: An example of the size of various effects on the inferred spectra, compared to our final uncertainties. Here we used the leakage in the main beam for S15 PA2 at 150 GHz, the residual sidelobe leakage for PA2 at 150 GHz, simulated leakage due to asymmetry for S15 PA2 at 150 GHz, and the uncertainty in the main beam for S15 PA2 deep56 at 150 GHz. The best fit Λ CDM plus foreground model for the deep 150×150 GHz spectra for ACT only from Choi et al. (2020) was used to compute the resulting change in D_ℓ^{TT} , D_ℓ^{TE} , and D_ℓ^{EE} due to these elements. These changes are then divided by the corresponding uncertainties σ_ℓ^{TT} , σ_ℓ^{TE} , and σ_ℓ^{EE} on the final binned deep spectra. At larger multipoles, eventually the uncertainty on the main beam would come to dominate. Note that we correct for the main beam leakage and the residual sidelobe leakage in the power spectrum likelihood, as explained in §3.5.3. The asymmetry leakage (simulated using an all-sky beam convolution code, as mentioned in §3.6.2) is negligible for all detector arrays. Sudden changes around $\ell = 2000$, as can be seen in the top panel, for example, are due to a change in the bin width at this scale.

amplitude of γ is never greater than 0.0035 (0.038). Around $\ell \sim 1000-3000$, this leakage correction is roughly a few-percent ($\sim 1-4\%$) effect for TE and a sub-percent ($\sim 0-0.4\%$) effect for EE .

As explained in [Choi et al. \(2020\)](#), in the power spectrum analysis we first compute a power spectrum for each season, sky region, detector array, and frequency. We then coadd over seasons and arrays to obtain one spectrum per region and frequency. Finally, the regions are divided into two groups based on the detection thresholds for point sources: deep (Deep1, Deep5, Deep6, and Deep8) and wide (BOSS and AdvACT). The spectra for the regions in these two groups are coadded, resulting in a single deep and a single wide power spectrum at each frequency.

To obtain the γ factors, we coadd the leakage beams for individual seasons and detector arrays using the same weights used to coadd the spectra, giving an effective leakage beam for both the deep and wide coadded spectra at each frequency. The uncertainties in the γ factors are incorporated in the data covariance matrix by using the main leakage beam uncertainties and the sidelobe residuals to compute another covariance matrix (similar to the main temperature beam covariance matrix) which is then added (in quadrature) to the data covariance matrix.

As described in [Aiola et al. \(2020\)](#); [Choi et al. \(2020\)](#), including this leakage correction in the likelihood reduces the residuals compared to the best-fitting Λ CDM model, but does not have a significant effect on the inferred cosmology. [Choi et al. \(2020\)](#) (in Section 12.3) also carried out a test by fitting for two scaling factors (one at 150 GHz and one at 98 GHz), that multiply the nominal values of the γ factors. The data support the baseline model, where the scaling factors are unity.

In the DR3 analysis in [Louis et al. \(2017\)](#) the residual leakage from the sidelobes was similarly treated as a systematic uncertainty in the cosmological power spectrum

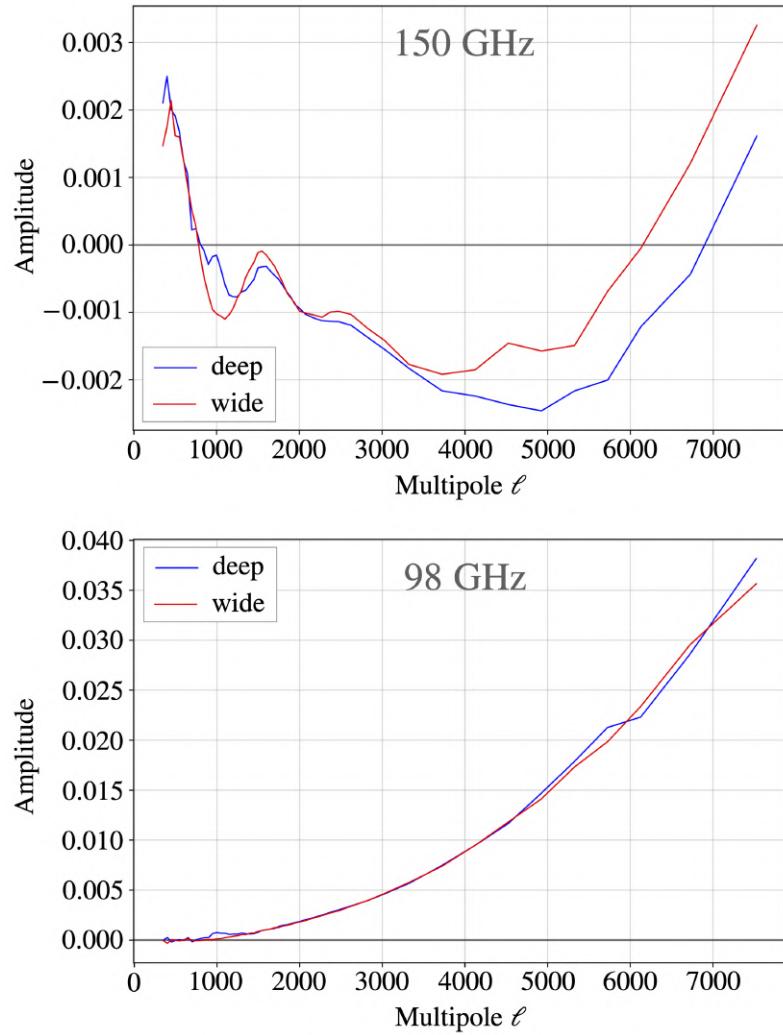


Figure 3.20: The γ factors encoding the ℓ -dependent T-to-P leakage, used in the power spectrum likelihood for both the deep and wide spectra at 150 GHz (top) and 98 GHz (bottom). The γ factors for 150 GHz include both the main beam leakage from §3.5.1 and the residual sidelobe leakage from §3.5.2. Since no sidelobes are seen at 98 GHz, those γ factors consist of only the main beam leakage.

analysis, but the correction for the main beam polarized leakage is new to DR4.³¹

3.6 Discussion

3.6.1 Beam Products

As part of DR4, several ACT data products were made publicly available on the NASA Legacy Archive for Microwave Background Data Analysis³² (LAMBDA) and at the National Energy Research Scientific Computing Center (NERSC). Details of all these data products are given in [Mallaby-Kay et al. \(2021\)](#).

The beams described in this chapter and used for the analysis in [Choi et al. \(2020\)](#) and [Aiola et al. \(2020\)](#) are included in the “ancillary products” section of the data release. For each season/region/array/frequency combination, there are multiple beam files. Both the real-space radial beam profiles and their harmonic-space transforms are provided, both for the instantaneous beams and for the beams with the jitter corrections included, and in each case are available both with and without the Rayleigh-Jeans-to-CMB spectral correction.

Note that the instantaneous beams are not region-dependent, so only the jitter-corrected beams contain region flags in their filenames. The instantaneous beams are suitable for time-domain analysis, but they should not be used when analyzing maps. The maps released as part of DR4 have not been corrected for the instrument beam. When working with the maps, one should use the jitter-corrected beams

³¹As described in [Choi et al. \(2020\)](#), for DR4 the ACT team adopted a blinding strategy in an attempt to prevent confirmation bias on the cosmological parameters. The analysis pre-unblinding used maps that deprojected the polarized sidelobes, but did not include the additional polarized leakage correction from the main beam or the sidelobe residuals. The post-unblinding analysis revealed some features in the *TE* residuals to Λ CDM that led us to a more thorough search for possible sources of systematic uncertainty in the *TE* spectrum. As a result, it was at this point that we added the correction in the likelihood for the T-to-P leakage from both the polarized sidelobe residuals, and the main beam.

³²lambda.gsfc.nasa.gov/product/act/actpol_prod_table.cfm

where the pointing variance and its uncertainty have been accounted for. The beam harmonic-space transforms can be used to correct for the beam effects in harmonic space (Bond & Efstathiou, 1987).

The TE and TB harmonic-space leakage beams for the main beam and the sidelobes and their residuals are also included in the ancillary products for DR4.

Since DR4 includes the data from DR3 as a subset, it is possible to compare some of the DR4 beams with the beams for the same season, sky region, detector array, and frequency released as part of DR3. As shown in §3.7.5, despite the differences in the analyses, the beams from DR3 and DR4 are consistent.

3.6.2 Beam Asymmetry

The asymmetry of the ACT main beams is relatively well described as elliptical, with aspect ratios varying from 1% to 20% across the different arrays, as shown in Table 1 of Choi et al. (2020). Here we give a brief summary of an investigation into this azimuthal asymmetry and the associated spurious signal.

Beam asymmetry is, to leading order, responsible for two effects. The asymmetric convolution creates anisotropy in the sky maps that distorts the shape of point sources and introduces statistical anisotropy in the inferred CMB. The magnitude of the effect is reduced by increased cross-linking: as the telescope observes a position on the sky using approximately orthogonal scan directions, the spurious signal from one scan roughly cancels with that of the other scan. The cross-linking in the ACT maps is sufficient to reduce the contamination to the power spectrum from this effect to an insignificant amount.

The second effect of beam asymmetry is to introduce T-to-P leakage. The beam asymmetry introduces a dependency in the time-ordered data on the position angle of the instrument: observations with different scan directions yield systematically

different data. For asymmetric beams with a quadrupole shape, which is the dominant asymmetric azimuthal mode of our approximately elliptical beams, the dependence on the position angle is interpreted by the map-maker as a linearly polarized sky component (Hu et al., 2003). In contrast to the first effect, the T-to-P leakage is not averaged down by cross-linking. Instead, the leakage is reduced by the instrument’s orthogonally polarized co-pointing detectors: the leakage picked up by one detector approximately cancels with the leakage picked up by its partner.³³ As mentioned in Choi et al. (2020) the residual leakage causes an additive bias to the *TE* and *TB* power spectra that is roughly constant with multipole with an amplitude that is less than 0.2σ away from zero. No attempt to remove this leakage has been made.

For the investigation described here, both effects were simulated in the time domain using an all-sky beam convolution code similar to Wandelt & Górski (2001); Reinecke et al. (2006); Prézeau & Reinecke (2010); Duivenvoorden et al. (2019).

3.6.3 Conclusion

In this chapter, we have presented the analysis of the ACT beams for DR4, which includes data from 2013–2016. Improvements to the beam pipeline include: better atmosphere subtraction for the Uranus maps, a new scattering term in the model that is fitted to the main beams, a better estimate of the uncertainty in these fits, and residual T-to-P leakage terms that are included in the ACT power spectrum likelihood. Considerable effort was spent developing a realistic model of the beams (including optical effects) and the mapping process, in order to study all elements of the analysis, including the propagation of systematic uncertainties.

The DR4 beams presented here were also used for the 2013–2016 data that were

³³This cancellation is not perfect for detector pairs with incorrect relative gain or pointing or for pairs with slightly different beams, but these effects are small compared to the leakage from detectors without a partner, which see no reduction.

part of the ACT DR5 maps (Naess et al., 2020). Finally, the pipeline presented here was used to obtain preliminary beams for the 2017–2018 Advanced ACT data included in DR5. Looking forward, as we collect and analyze more ACT data and approach the cosmic variance limit, we will become increasingly sensitive to details of the instrument beams. While some details may change³⁴, we expect to adopt similar methods as described here for analysis of the post-2016 ACT data.

3.7 Appendix Materials

3.7.1 Shrinking Algorithm for Radial Profile Covariance Matrix

In general, shrinkage is useful when estimating a covariance matrix (or any matrix, really) from a limited number of data points (Schäfer & Strimmer, 2005). This technique works by combining an empirical estimate of the covariance matrix (a high-dimensional estimate of the underlying covariance with little or no bias) with a model (a low-dimensional estimate which may be biased but has much smaller variance) to minimize the total mean squared error (sum of bias squared and variance) with respect to the true underlying covariance. This is useful when the off-diagonal elements of the covariance matrix are excessively noisy. One can analytically calculate the optimal combination of the low and high dimensional estimates, parametrized by the shrinkage intensity. A review of covariance matrix shrinkage with an example of application to cosmological analysis is given in Pope & Szapudi (2008). Shrinkage can result in a much better estimation of the covariance matrix when few measurements are available, and it does not adversely affect the covariance in the case of a large number of measurements. Since the method is computationally very simple, Pope & Szapudi (2008) suggest it should always be employed.

³⁴Future changes likely will be informed in part by the work currently being done to model the beams in 2D.

The method works as follows. Suppose we have n sets of observations and each observation yields a data vector \mathbf{x} of length p . In the case of the beams analysis for DR4, for a given season, detector array, and frequency, the number of observations n is between 6 and 51 (as listed in Table 3.1) and in each case the resulting data vector \mathbf{x} (radial profile) is of length $p = 55$. Let $x_i^{(k)}$ then represent the i^{th} element of the vector for the k^{th} observation. The estimated empirical mean of the i^{th} element across all observations is then $\bar{x}_i = (1/n) \sum_{k=1}^n x_i^{(k)}$.

If we define

$$W_{ij}^{(k)} = (x_i^{(k)} - \bar{x}_i)(x_j^{(k)} - \bar{x}_j), \quad (3.27)$$

and

$$\bar{W}_{ij} = \frac{1}{n} \sum_{k=1}^n W_{ij}^{(k)}, \quad (3.28)$$

then an unbiased, empirical estimate of the covariance, \mathbf{S} , of the data, is

$$S_{ij} = \widehat{\text{Cov}}(x_i, x_j) \quad (3.29a)$$

$$= \frac{n}{n-1} \bar{W}_{ij} \quad (3.29b)$$

$$= \frac{1}{n-1} \sum_{k=1}^n (x_i^{(k)} - \bar{x}_i)(x_j^{(k)} - \bar{x}_j). \quad (3.29c)$$

We can also compute the covariance of the elements of this covariance matrix,

$$\widehat{\text{Cov}}(S_{ij}, S_{lm}) = \frac{n}{(n-1)^3} \sum_{k=1}^n (W_{ij}^{(k)} - \bar{W}_{ij})(W_{lm}^{(k)} - \bar{W}_{lm}), \quad (3.30)$$

writing the variance of an individual covariance matrix entry as $\widehat{\text{Var}}(S_{ij}) = \widehat{\text{Cov}}(S_{ij}, S_{ij})$.

Let \mathbf{T} be the target matrix, our model with fewer (or no) free parameters. An example of such a matrix could be the identity matrix times a constant, or the target we use, which is described later, in Equation 3.33. Then we can similarly write an equation for $\widehat{\text{Cov}}(T_{ij}, S_{ij})$. The optimal shrinkage intensity, λ^* , and resulting final estimate of

the covariance matrix \mathbf{C} are then given by

$$\hat{\lambda}^* = \frac{\sum_{i,j} \widehat{\text{Var}}(S_{ij}) - \widehat{\text{Cov}}(T_{ij}, S_{ij})}{\sum_{i,j} (T_{ij} - S_{ij})^2}, \quad (3.31)$$

$$\mathbf{C} = \hat{\lambda}^* \mathbf{T} + (1 - \hat{\lambda}^*) \mathbf{S}. \quad (3.32)$$

The expression for $\hat{\lambda}^*$ above is the practical estimator suggested by [Schäfer & Strimler \(2005\)](#) in order to estimate the optimal shrinkage intensity. This is based on the analytic solution for the optimal shrinkage intensity, λ^* , introduced by [Ledoit & Wolf \(2003\)](#), which is identical to the expression above except the unbiased sample estimates $\widehat{\text{Var}}$ and $\widehat{\text{Cov}}$ are replaced by the true underlying Var and Cov. Before this analytic solution, matrix shrinkage was much less practical, since complex and computationally intensive methods were needed to find the optimal shrinkage intensity.

Looking more closely at Equation 3.31, the term $\widehat{\text{Var}}(S_{ij})$ in the numerator means that as the variances of the elements of the empirical covariance matrix decrease (so, as n becomes much larger than p) the shrinkage intensity decreases and our final estimate \mathbf{C} of the covariance matrix approaches the empirical estimate \mathbf{S} . In the case of our radial beam profile, $n \ll p$, which is precisely when shrinkage is the most important. The term $\widehat{\text{Cov}}(T_{ij}, S_{ij})$ in the numerator accounts for the fact that \mathbf{S} and \mathbf{T} are estimated from the same data, so if some elements of \mathbf{T} are equal to elements of \mathbf{S} then the two terms in the numerator cancel for those elements, and they do not affect the estimate of $\hat{\lambda}^*$. The denominator in Equation 3.31 ensures that if our choice of \mathbf{T} is very different from \mathbf{S} , then $\hat{\lambda}^*$ will be small and \mathbf{C} will be close to \mathbf{S} . A poor choice of target matrix should not therefore negatively affect \mathbf{C} .

If Equation 3.31 leads to a value of $\hat{\lambda}^*$ that is greater than one, then we set $\hat{\lambda}^* = 1$, which means our final estimate of the covariance matrix is composed of only the

target matrix. On the other hand, if the equation leads to a value of $\hat{\lambda}^*$ that is less than zero, we set $\hat{\lambda}^* = 0$, and our final estimate of the covariance matrix is equal to the empirical covariance matrix.

Several common choices for the target matrix are listed in Table 2 of [Schäfer & Strimler \(2005\)](#), along with simplified expressions for the associated shrinkage intensities $\hat{\lambda}^*$. In this analysis we use a “diagonal, unequal variance” target matrix, that is, a matrix whose diagonal elements are equal to the diagonal elements of our empirical estimate \mathbf{S} of the covariance matrix and whose off-diagonal elements are zero:

$$T_{ij} = \begin{cases} S_{ii} & \text{if } i = j \\ 0 & \text{if } i \neq j . \end{cases} \quad (3.33)$$

In this case the expression for the shrinkage intensity simplifies to

$$\hat{\lambda}^* = \frac{\sum_{i \neq j} \widehat{\text{Var}}(S_{ij})}{\sum_{i \neq j} S_{ij}^2} . \quad (3.34)$$

3.7.2 Computing the Harmonic Transform of the Beam

The spherical harmonics are functions defined on the sphere as

$$Y_l^m(\theta, \phi) = P_l^m(\cos \theta) \exp\{im\phi\} , \quad (3.35)$$

where P_l^m are the Legendre polynomials normalized for spherical harmonics. These are defined as

$$P_l^m(x) = (-1)^m \sqrt{\frac{2l+1}{4\pi}} \sqrt{\frac{(l-|m|)!}{(l+|m|)!}} (1-x^2)^{|m|/2} \frac{d^{|m|}}{dx^{|m|}} P_l(x) , \quad (3.36)$$

where the Legendre polynomials $P_l(x)$ are defined via the recurrence relation

$$P_0(x) = 1 \quad (3.37a)$$

$$P_1(x) = x \quad (3.37b)$$

$$nP_n(x) = (2n-1)xP_{n-1}(x) - (n-1)P_{n-2}(x) . \quad (3.37c)$$

The spherical harmonics are orthonormal, so they obey the relation

$$\int_0^{2\pi} \int_0^\pi Y_l^m(\theta, \phi) \overline{Y_{l'}^{m'}}(\theta, \phi) \sin \theta \, d\theta \, d\phi = \delta_{ll'} \delta_{mm'} , \quad (3.38)$$

where \bar{z} is the complex conjugate of z and δ_{ij} is the Kronecker symbol.

In a spherical harmonic transform, we compute the coefficients f_l^m used to express a function $f(\theta, \phi)$ as

$$f(\theta, \phi) = \sum_{l=0}^{\infty} \sum_{m=-l}^l f_l^m Y_l^m(\theta, \phi) . \quad (3.39)$$

The coefficients can be computed using the equation

$$f_l^m = \int_0^{2\pi} \int_0^\pi f(\theta, \phi) \overline{Y_l^m}(\theta, \phi) \sin \theta \, d\theta \, d\phi \quad (3.40a)$$

$$= \int_0^{2\pi} \int_0^\pi f(\theta, \phi) P_l^m(\cos \theta) \exp\{-im\phi\} \sin \theta \, d\theta \, d\phi . \quad (3.40b)$$

If $f(\theta, \phi)$ is independent of ϕ (as is the case for our beam), then we can write $f(\theta, \phi) = f(\theta)$ and the equation above becomes

$$f_l^m = \int_0^{2\pi} \exp\{-im\phi\} d\phi \int_0^\pi f(\theta) P_l^m(\cos \theta) \sin \theta \, d\theta . \quad (3.41)$$

The integral over ϕ then simplifies to

$$\int_0^{2\pi} e^{-im\phi} \, d\phi = 2\pi \delta_{m0} . \quad (3.42)$$

So f_l^m is only non-zero for $m = 0$, in which case we have

$$f_l^0 = 2\pi \int_0^\pi f(\theta) P_l(\cos \theta) \sin \theta \, d\theta \quad (3.43a)$$

$$= 2\pi \int_{-1}^1 f(\theta) P_l(\cos \theta) \, d\cos \theta. \quad (3.43b)$$

This is the equation for the Legendre polynomial transform, presented as a means of converting the radial beam profile $B(\theta)$ to the harmonic transform B_ℓ . However, this can be time-consuming to compute. For small beams such as ours, it is not necessary to work in the curved sky regime. We instead perform a 2D Fourier transform, which effectively becomes a Hankel transform, as shown below. The difference between the Hankel and Legendre polynomial transforms is less than a factor of 4×10^{-5} between $\ell = 0$ and $\ell = 10,000$ and the Hankel transform is much faster to compute.

Now let's consider the 2D Fourier transform of a function $f(x, y)$,

$$F(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x, y) \exp\{-i(xk_x + yk_y)\} \, dx \, dy. \quad (3.44)$$

Introducing the polar coordinates

$$x = \theta \cos \phi \quad y = \theta \sin \phi$$

$$k_x = k \cos \psi \quad k_y = k \sin \psi$$

where θ and ϕ here correspond to the θ and ϕ in spherical coordinates used throughout this chapter, we then have, in the flat sky approximation,

$$\begin{aligned} F(k \cos \psi, k \sin \psi) &\equiv \mathcal{F}(k, \psi) \\ &= \int_0^\infty \int_0^{2\pi} f(\theta, \phi) \exp\{-i\theta k(\cos \phi \cos \psi + \sin \phi \sin \psi)\} \theta \, d\theta \, d\phi. \end{aligned} \quad (3.45)$$

If our function is circularly symmetric, so independent of ϕ (as is the case for our beam model), we have $f(x, y) = f(\theta, \phi) = f(\theta)$ and the equation above becomes

$$\mathcal{F}(k, \psi) = \int_0^\infty \theta f(\theta) \int_0^{2\pi} \exp\{-i\theta k(\cos \phi \cos \psi + \sin \phi \sin \psi)\} d\theta d\phi \quad (3.46a)$$

$$= \int_0^\infty \theta f(\theta) \int_0^{2\pi} \exp\{-i\theta k \cos(\phi - \psi)\} d\theta d\phi \quad (3.46b)$$

$$= \int_0^\infty \theta f(\theta) \int_0^{2\pi} \exp\{-i\theta k \cos \alpha\} d\theta d\alpha \quad (3.46c)$$

$$= \int_0^\infty \theta f(\theta) 2 \int_0^\pi \exp\{-i\theta k \cos \alpha\} d\theta d\alpha. \quad (3.46d)$$

Using the integral representation

$$J_n(z) = \frac{(-i)^n}{\pi} \int_0^\pi \exp\{iz \cos \varphi\} \cos(n\varphi) d\varphi \quad (3.47)$$

for the Bessel functions J_n of the first kind, we have

$$J_0(z) = \frac{1}{\pi} \int_0^\pi \exp\{iz \cos \varphi\} d\varphi, \quad (3.48)$$

and so the final expression for the 2D Fourier transform of a circularly symmetric function $f(\theta)$ may be written as

$$\mathcal{F}(k) = 2\pi \int_0^\infty \theta f(\theta) J_0(-\theta k) d\theta \quad (3.49a)$$

$$= 2\pi \int_0^\infty \theta f(\theta) J_0(\theta k) d\theta, \quad (3.49b)$$

which is a Hankel transform of order zero, and where the last line follows from the identity $J_n(-z) = J_n(z)$ for integer n .

In order to compute the harmonic transform of our beam profile, we evaluate the expression above separately for the three main terms in our beam profile fit: the core term (composed of the sum of basis functions), the scattering term, and the

$1/\theta^3$ asymptotic term. The integrals for the core and scattering terms are computed numerically, but we derive an analytic expression for the integral of the $1/\theta^3$ term, shown below.

Given a fit amplitude α , the Hankel transform for the $1/\theta^3$ term may be written as

$$\mathcal{F}_{1/\theta^3}(k) = \alpha \int_0^\infty \theta \left(\frac{1}{\theta^3} \right) J_0(\theta\ell) d\theta = \alpha \int_0^\infty \frac{J_0(\theta\ell)}{\theta^2} d\theta. \quad (3.50)$$

The analytic expression we use for this integral is

$$\int \frac{J_0(\theta\ell)}{\theta^2} d\theta = \ell \left[J_1(\theta\ell) - J_0(\theta\ell) \left(\frac{\theta^2\ell^2 + 1}{\theta\ell} \right) - \frac{\pi\theta\ell}{2} \left(H_0(\theta\ell)J_1(\theta\ell) - H_1(\theta\ell)J_0(\theta\ell) \right) \right], \quad (3.51)$$

where $H_n(x)$ is the Struve function.

3.7.3 Beam Solid Angles for Different Effective Frequencies

To indicate the effect of the passbands on the beams, we tabulate the solid angles for select effective frequencies. The effective frequencies of the band centers are currently uncertain to approximately 2.4 GHz. However, since this is in part due to systematic errors in the measurements, the relative uncertainties are smaller.

Similar to the correction we make to the main beams for use with the CMB, in each case we take the beam to be $B'(\ell) = B(\ell\nu_{\text{RJ}}/\nu_{\text{S}})$, where ν_{RJ} is the effective frequency for radiation with a Rayleigh-Jeans spectrum and ν_{S} is the effective frequency for the source of interest (either CMB, synchrotron emission, dust, or the thermal Sunyaev-Zel'dovich (tSZ) effect).

As described in §3.4.4, for the beam analysis for DR4, the effective frequencies from Thornton et al. (2016) were used for the RJ-to-CMB beam spectral correction. Subsequently, the effective frequencies were re-computed for the foreground modeling for

DR4, using improved passband data and upgraded code, as detailed in Appendix D of [Choi et al. \(2020\)](#). These updated frequencies are shown here in Table 3.5 and the corresponding beam solid angles are shown in Tables 3.6, 3.7, 3.8, 3.9. Considering the uncertainties on the passbands, the frequencies from [Thornton et al. \(2016\)](#) and [Choi et al. \(2020\)](#) are consistent. In addition, given that the uncertainty on the beams is subdominant in the power spectrum analysis, which of these effective frequencies one uses for the beam spectral correction does not have a significant effect on the results. Still, the solid angles for the CMB in Table 3.6 below are slightly different from those in Table 3.4.

For any particular season/region/detector array/frequency, the uncertainty on the solid angles is 2–5%. This includes both the instantaneous beam uncertainty and the uncertainty due to the region-dependent jitter correction. As is apparent here, the derived solid angle and passband are intimately connected.

In addition, the beam and the passband are coupled, an effect which we considered for the first time in detail in [Madhavacheril et al. \(2020\)](#), and described in Appendix A of said paper. In short, the beam shape evolves as a function of frequency across the passband, and so our customary assumption of separability of these two components does not hold to high precision. This beam-bandpass coupling was modeled as part of a systematic check of the DR4 results in [Choi et al. \(2020\)](#) and was found to not have a significant effect on the inferred cosmological parameters.

Table 3.5: Effective frequencies [GHz].

Array	Band	RJ	CMB	tSZ	dust	sync
PA1	150 GHz	150.8	149.6	150.0	151.2	146.9
PA2	150 GHz	151.2	149.9	150.4	151.6	147.3
PA3	150 GHz	148.4	147.6	147.9	148.7	145.8
PA3	98 GHz	98.7	97.9	98.4	98.8	95.5

Table 3.6: Effective beam solid angles (nsr) for the CMB.

Array	Band	Season	Deep1	Deep5	Deep6	Deep56	Deep8	BOSS	AdvACT
PA1	150 GHz	S13	213	209	210	-	-	-	-
		S14	-	-	-	204	-	-	-
		S15	-	-	-	206	199	205	-
PA2	150 GHz	S14	-	-	-	190	-	-	-
		S15	-	-	-	196	192	193	-
		S16	-	-	-	-	-	-	194
PA3	150 GHz	S15	-	-	-	274	268	285	-
		S16	-	-	-	-	-	-	242
PA3	98 GHz	S15	-	-	-	546	541	556	-
		S16	-	-	-	-	-	-	524

Table 3.7: Effective beam solid angles (nsr) for synchrotron emission.

Array	Band	Season	Deep1	Deep5	Deep6	Deep56	Deep8	BOSS	AdvACT
PA1	150 GHz	S13	221	217	218	-	-	-	-
		S14	-	-	-	211	-	-	-
		S15	-	-	-	213	207	212	-
PA2	150 GHz	S14	-	-	-	196	-	-	-
		S15	-	-	-	203	198	200	-
		S16	-	-	-	-	-	-	201
PA3	150 GHz	S15	-	-	-	281	275	292	-
		S16	-	-	-	-	-	-	248
PA3	98 GHz	S15	-	-	-	575	569	585	-
		S16	-	-	-	-	-	-	552

Table 3.8: Effective beam solid angles (nsr) for dusty sources.

Array	Band	Season	Deep1	Deep5	Deep6	Deep56	Deep8	BOSS	AdvACT
PA1	150 GHz	S13	208	205	205	-	-	-	-
		S14	-	-	-	199	-	-	-
		S15	-	-	-	201	195	200	-
PA2	150 GHz	S14	-	-	-	186	-	-	-
		S15	-	-	-	192	187	189	-
		S16	-	-	-	-	-	-	190
PA3	150 GHz	S15	-	-	-	270	264	281	-
		S16	-	-	-	-	-	-	239
PA3	98 GHz	S15	-	-	-	536	531	546	-
		S16	-	-	-	-	-	-	515

Table 3.9: Effective beam solid angles (nsr) for the tSZ effect.

Array	Band	Season	Deep1	Deep5	Deep6	Deep56	Deep8	BOSS	AdvACT
PA1	150 GHz	S13	212	208	209	-	-	-	-
		S14	-	-	-	202	-	-	-
		S15	-	-	-	204	198	204	-
PA2	150 GHz	S14	-	-	-	189	-	-	-
		S15	-	-	-	195	190	192	-
		S16	-	-	-	-	-	-	193
PA3	150 GHz	S15	-	-	-	273	267	284	-
		S16	-	-	-	-	-	-	241
PA3	98 GHz	S15	-	-	-	541	536	551	-
		S16	-	-	-	-	-	-	520

3.7.4 Transforming from $\{Q_r, U_r\}$ to $\{E, B\}$

Instead of estimating the polarized ℓ -space beams directly from non-local E and B transform maps, we opt to use a locally defined map-space polarization basis for this purpose. As we show below, the fields Q_r and U_r turn out to be a convenient choice; in the flat-sky limit, they are defined in terms of local linear combinations of the usual Q and U maps as follows:

$$Q_r(\boldsymbol{\theta}) = Q(\boldsymbol{\theta}) \cos 2\phi_\theta + U(\boldsymbol{\theta}) \sin 2\phi_\theta \quad (3.52)$$

$$U_r(\boldsymbol{\theta}) = U(\boldsymbol{\theta}) \cos 2\phi_\theta - Q(\boldsymbol{\theta}) \sin 2\phi_\theta \quad (3.53)$$

where $\boldsymbol{\theta} \equiv (\theta, \phi_\theta)$ are standard polar coordinates with the beam centroid as their origin and ϕ_θ increasing clockwise from the positive y -axis (assuming that one uses the convention in which $+x$ points to the right and $+y$ points upward). Conversely, we may also define these fields in terms of their local contributions to both Q and U in the same coordinate system:

$$Q(\boldsymbol{\theta}) = Q_r(\boldsymbol{\theta}) \cos 2\phi_\theta - U_r(\boldsymbol{\theta}) \sin 2\phi_\theta \quad (3.54)$$

$$U(\boldsymbol{\theta}) = U_r(\boldsymbol{\theta}) \cos 2\phi_\theta + Q_r(\boldsymbol{\theta}) \sin 2\phi_\theta. \quad (3.55)$$

Since we are ultimately interested in how leakage manifests itself in the usual angular power spectra, we need to translate any polarized beam models of Q_r and U_r to an ℓ -space representation of E and B . As it turns out, there exists a simple relation between the azimuthally averaged versions of these components, which we derive here in the flat-sky limit. We begin with the Fourier-space expressions for E and B :

$$E(\boldsymbol{\ell}) = \hat{Q}(\boldsymbol{\ell}) \cos 2\phi_\ell + \hat{U}(\boldsymbol{\ell}) \sin 2\phi_\ell \quad (3.56)$$

$$B(\boldsymbol{\ell}) = \hat{U}(\boldsymbol{\ell}) \cos 2\phi_\ell - \hat{Q}(\boldsymbol{\ell}) \sin 2\phi_\ell \quad (3.57)$$

where $\boldsymbol{\ell} \equiv (\ell, \phi_\ell)$ is the Fourier conjugate of $\boldsymbol{\theta}$, and $\{\hat{Q}, \hat{U}\}$ are just standard Fourier transforms of $\{Q, U\}$:

$$\hat{Q}(\boldsymbol{\ell}) = \int Q(\boldsymbol{\theta}) e^{i\boldsymbol{\ell} \cdot \boldsymbol{\theta}} d\boldsymbol{\theta} = \int Q(\theta, \phi_\theta) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta \quad (3.58)$$

$$\hat{U}(\boldsymbol{\ell}) = \int U(\boldsymbol{\theta}) e^{i\boldsymbol{\ell} \cdot \boldsymbol{\theta}} d\boldsymbol{\theta} = \int U(\theta, \phi_\theta) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta. \quad (3.59)$$

Taking the azimuthal average of Equations 3.56 and 3.57, we get the one-dimensional transforms \tilde{E} and \tilde{B} :

$$\tilde{E}(\ell) = \frac{1}{2\pi} \int \hat{Q}(\boldsymbol{\ell}) \cos 2\phi_\ell d\phi_\ell + \frac{1}{2\pi} \int \hat{U}(\boldsymbol{\ell}) \sin 2\phi_\ell d\phi_\ell \quad (3.60)$$

$$\tilde{B}(\ell) = \frac{1}{2\pi} \int \hat{U}(\boldsymbol{\ell}) \cos 2\phi_\ell d\phi_\ell - \frac{1}{2\pi} \int \hat{Q}(\boldsymbol{\ell}) \sin 2\phi_\ell d\phi_\ell. \quad (3.61)$$

The expression for \tilde{E} in Equation 3.60 may be rewritten in terms of map-space Q and U with the help of Equations 3.58 and 3.59:

$$\begin{aligned} \tilde{E}(\ell) = & \frac{1}{2\pi} \int Q(\theta, \phi_\theta) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta \cos 2\phi_\ell d\phi_\ell \\ & + \frac{1}{2\pi} \int U(\theta, \phi_\theta) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta \sin 2\phi_\ell d\phi_\ell. \end{aligned} \quad (3.62)$$

Substituting for $Q(\boldsymbol{\theta})$ and $U(\boldsymbol{\theta})$ in Equation 3.62 using Equations 3.54 and 3.55, we obtain a relation between \tilde{E} and $\{Q_r, U_r\}$:

$$\begin{aligned} \tilde{E}(\ell) = & \frac{1}{2\pi} \int (Q_r(\theta, \phi_\theta) \cos 2\phi_\theta - U_r(\theta, \phi_\theta) \sin 2\phi_\theta) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta \cos 2\phi_\ell d\phi_\ell \\ & + \frac{1}{2\pi} \int (U_r(\theta, \phi_\theta) \cos 2\phi_\theta + Q_r(\theta, \phi_\theta) \sin 2\phi_\theta) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta \sin 2\phi_\ell d\phi_\ell. \end{aligned} \quad (3.63)$$

Grouping together the terms with $Q_r(\theta)$ and $U_r(\theta)$, the equation above becomes:

$$\begin{aligned}\tilde{E}(\ell) &= \frac{1}{2\pi} \int Q_r(\theta, \phi_\theta) (\cos 2\phi_\theta \cos 2\phi_\ell + \sin 2\phi_\theta \sin 2\phi_\ell) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta d\phi_\ell \\ &+ \frac{1}{2\pi} \int U_r(\theta, \phi_\theta) (\cos 2\phi_\theta \sin 2\phi_\ell - \sin 2\phi_\theta \cos 2\phi_\ell) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta d\phi_\ell.\end{aligned}\quad (3.64)$$

Then, making use of a simple trigonometric identity, Equation 3.64 may be written as:

$$\begin{aligned}\tilde{E}(\ell) &= \frac{1}{2\pi} \int Q_r(\theta, \phi_\theta) \cos 2(\phi_\theta - \phi_\ell) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta d\phi_\ell \\ &- \frac{1}{2\pi} \int U_r(\theta, \phi_\theta) \sin 2(\phi_\theta - \phi_\ell) e^{i\ell\theta \cos(\phi_\theta - \phi_\ell)} \theta d\theta d\phi_\theta d\phi_\ell.\end{aligned}\quad (3.65)$$

And making the substitution $\phi_\rho \equiv \phi_\theta - \phi_\ell$, Equation 3.65 may be expressed as:

$$\begin{aligned}\tilde{E}(\ell) &= \frac{1}{2\pi} \int Q_r(\theta, \phi_\theta) \cos 2\phi_\rho e^{i\ell\theta \cos \phi_\rho} \theta d\theta d\phi_\theta d\phi_\rho \\ &- \frac{1}{2\pi} \int U_r(\theta, \phi_\theta) \sin 2\phi_\rho e^{i\ell\theta \cos \phi_\rho} \theta d\theta d\phi_\theta d\phi_\rho.\end{aligned}\quad (3.66)$$

We are now able to write each of the two terms in the expression for \tilde{E} as three separate integrals:

$$\begin{aligned}\tilde{E}(\ell) &= \int \left(\frac{1}{2\pi} \int Q_r(\theta, \phi_\theta) d\phi_\theta \right) \left(\int \cos 2\phi_\rho e^{i\ell\theta \cos \phi_\rho} d\phi_\rho \right) \theta d\theta \\ &- \int \left(\frac{1}{2\pi} \int U_r(\theta, \phi_\theta) d\phi_\theta \right) \left(\int \sin 2\phi_\rho e^{i\ell\theta \cos \phi_\rho} d\phi_\rho \right) \theta d\theta.\end{aligned}\quad (3.67)$$

Note that the integrals of $Q_r(\theta)$ and $U_r(\theta)$ over ϕ_θ — the first set of parentheses — are simply the azimuthal averages \tilde{Q}_r and \tilde{U}_r , while the integrals over ϕ_ρ — the second set of parentheses — turn out to have simple analytic counterparts:

$$\int \cos 2\phi_\rho e^{i\ell\theta \cos \phi_\rho} d\phi_\rho = -2\pi J_2(\ell\theta)\quad (3.68)$$

$$\int \sin 2\phi_\rho e^{i\ell\theta \cos \phi_\rho} d\phi_\rho = 0. \quad (3.69)$$

So the azimuthally averaged ℓ -space E beam is simply the second-order Hankel transform of the azimuthally averaged map-space Q_r beam:

$$\tilde{E}(\ell) = -2\pi \int \tilde{Q}_r(\theta) J_2(\ell\theta) \theta d\theta. \quad (3.70)$$

One can similarly show that the same relation exists between the azimuthally averaged ℓ -space B and map-space U_r beams:

$$\tilde{B}(\ell) = -2\pi \int \tilde{U}_r(\theta) J_2(\ell\theta) \theta d\theta. \quad (3.71)$$

With Equations 3.70 and 3.71 in hand, we have a complete formalism for transforming the polarized beams using the $\{Q_r, U_r\}$ basis.

3.7.5 DR3 vs DR4 Beams

The beam transforms made publicly available for the DR3 and DR4 releases have been compared for each season, sky region, detector array, and frequency in common, as shown in Figure 3.21. Despite the changes in the analyses, the beam transforms are consistent.

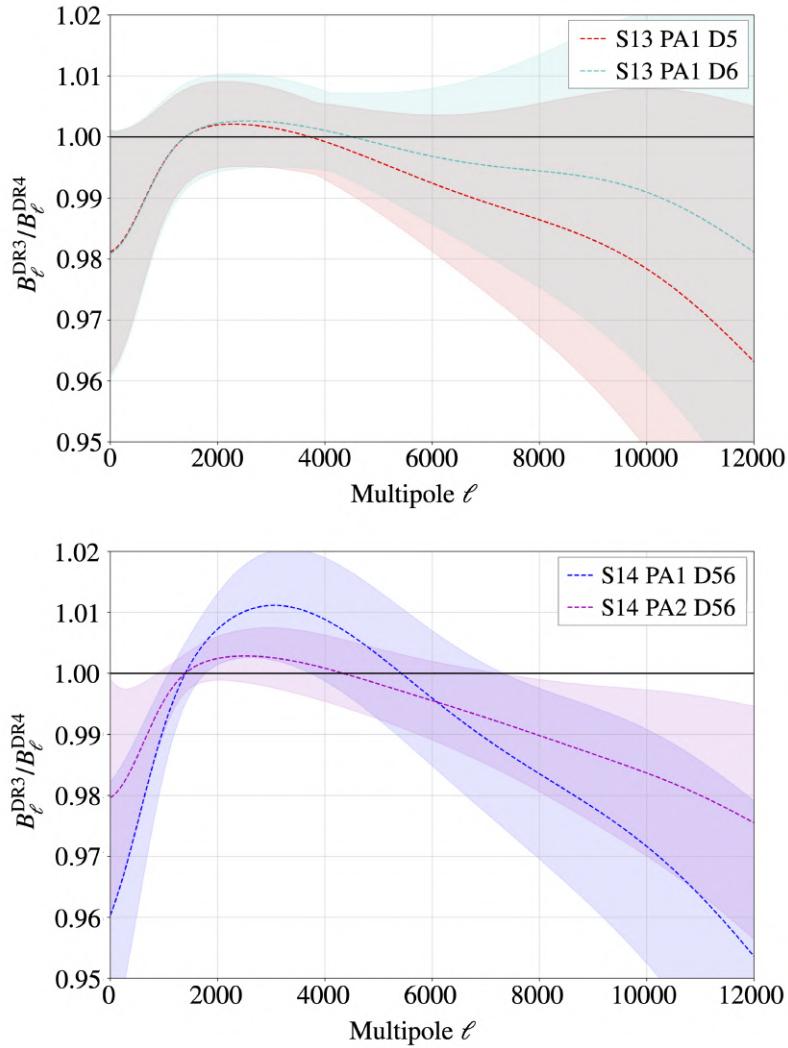


Figure 3.21: Ratio of the beam transforms for DR3 and DR4 at 150 GHz, for S13 (top) and S14 (bottom). The shaded bands indicate the 1σ uncertainty bounds, determined using the maximum uncertainty of the two transforms being compared. For these plots, the transforms have been normalized at $\ell = 1400$, which corresponds roughly to the effective calibration scale.

Chapter 4

Power Spectra

Part of this chapter is adapted from Choi et al. (2020).

In this chapter, we begin by describing the formalism for power spectrum estimation. We then describe the analysis pipeline used to compute the ACT DR4 power spectra (Choi et al., 2020). Finally, we reproduce these spectra using publicly available code, an important step for post-DR4 power spectrum analysis and community reproduction of the ACT results.

4.1 Formalism

4.1.1 Definitions

The CMB temperature fluctuations about the mean, $\frac{\delta T}{T_0}(\hat{n})$, are a Gaussian random field, and can be expressed in terms of the spherical harmonic basis functions $Y_{\ell m}$ as

$$\frac{\delta T}{T_0}(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^T Y_{\ell m}(\hat{n}) , \quad (4.1)$$

where \hat{n} is the direction in the sky, T_0 is the mean temperature ($T_0 = 2.7255 \pm 0.0006$ K according to [Fixsen \(2009\)](#)), and the expansion factors obey the expression

$$a_{\ell m}^T = \frac{1}{T_0} \int d\hat{n} \delta T(\hat{n}) Y_{\ell m}^*(\hat{n}) . \quad (4.2)$$

From the orthonormality relation

$$\int Y_{\ell m} Y_{\ell' m'}^* d^2 n = \delta_{\ell \ell'} \delta_{m m'} , \quad (4.3)$$

we can write

$$\langle a_{\ell m}^T a_{\ell' m'}^{T*} \rangle = \delta_{\ell \ell'} \delta_{m m'} C_{\ell}^{TT} . \quad (4.4)$$

This is the two-point correlation function, the variance of the $a_{\ell m}^T$, which encodes all the information for a Gaussian field.

In practice, the sky we observe is only one realization of the Gaussian random field. We only have one Universe to study, one sampling of the C_{ℓ} distribution.¹ For each ℓ , we measure $2\ell + 1$ amplitudes $a_{\ell m}$ (for $m = -\ell, -\ell + 1, \dots, 0, \dots, \ell - 1, \ell$), each from the same distribution. The variance is independent of m due to the isotropy of the universe. We then estimate the power at each ℓ by computing \hat{C}_{ℓ} averaged over the m modes,

$$\hat{C}_{\ell}^{TT} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} (a_{\ell m}^{T*} a_{\ell m}^T) . \quad (4.5)$$

The term \hat{C}_{ℓ}^{TT} represents our estimate of the temperature power spectrum C_{ℓ}^{TT} .

For several reasons, including visualization purposes,² we often work with the quantity

¹Although not explored here, methods for measuring the CMB at other locations have been proposed, e.g., [Kamionkowski et al. \(1997\)](#); [Deutsch et al. \(2018\)](#).

²When we plot $\ell(2\ell + 1)C_{\ell}/4\pi$ on a logarithmic ℓ scale (or $(2\ell + 1)C_{\ell}/4\pi$ on a linear ℓ scale), the area under the curve gives the temperature variance, or $\langle (\delta T/T_0)^2 \rangle$. For large ℓ , $D_{\ell} = \ell(\ell + 1)C_{\ell}/2\pi \approx \ell(2\ell + 1)C_{\ell}/4\pi$, and so in this context it makes sense to plot D_{ℓ} with a logarithmic ℓ scale. The reason for plotting $D_{\ell} = \ell(\ell + 1)C_{\ell}/2\pi$ (as opposed to $\ell(2\ell + 1)C_{\ell}/4\pi$) is that it leads to the flat Sachs-Wolfe plateau at low ℓ for a scale-invariant primordial power spectrum. In addition to visualization, having a flatter spectrum (hence D_{ℓ} as opposed to C_{ℓ}) is preferable for binning in ℓ .

$$D_\ell = \frac{\ell(\ell+1)}{2\pi} C_\ell . \quad (4.6)$$

For small angular scales on the sky (large multipoles ℓ), we have many values of $a_{\ell m}$, and so the statistical uncertainty on our estimate of \hat{C}_ℓ is small. However, at large angular scales (small multipoles ℓ) there is an inescapable uncertainty due to the number of modes m we can measure. This uncertainty is referred to as cosmic variance, and it scales as

$$\Delta C_\ell = \sqrt{\frac{2}{2\ell+1}} C_\ell . \quad (4.7)$$

If the fraction of sky observed is f_{sky} , then the uncertainty above increases by a factor of $\sqrt{1/f_{\text{sky}}}$, and it is usually referred to as sample variance (Scott et al., 1994).

The uncertainty on a measured power spectrum C_ℓ can be approximated as

$$\sigma^2(C_\ell) = \frac{2}{(2\ell+1)f_{\text{sky}}} \left(C_\ell + N_\ell \right)^2 , \quad (4.8)$$

where N_ℓ is the noise power spectrum (Knox, 1995). A measurement is said to be cosmic variance (or sample variance) limited when the C_ℓ term dominates over the N_ℓ term.

Analogously to temperature, we can express the polarization fields in terms of spherical harmonic basis functions as

$$E(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^E Y_{\ell m}(\hat{n}) , \quad (4.9)$$

$$B(\hat{n}) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m}^B Y_{\ell m}(\hat{n}) , \quad (4.10)$$

and their power spectra as

$$C_\ell^{EE} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} (a_{\ell m}^{E*} a_{\ell m}^E), \quad (4.11)$$

$$C_\ell^{BB} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} (a_{\ell m}^{B*} a_{\ell m}^B). \quad (4.12)$$

As mentioned in Section 1.1.2, we are also interested in the TE cross-spectrum,

$$C_\ell^{TE} = \frac{1}{2\ell+1} \sum_{m=-\ell}^{\ell} (a_{\ell m}^{T*} a_{\ell m}^E). \quad (4.13)$$

4.1.2 Estimation

The raw spherical harmonic transform (SHT), or “pseudo- C_ℓ ” spectrum of a map, is related to the true underlying power spectrum C_ℓ through the relation

$$\langle \hat{C}_\ell \rangle = \sum_{\ell'} M_{\ell\ell'} F_{\ell'} B_{\ell'}^2 \langle C_{\ell'} \rangle + \langle \hat{N}_\ell \rangle, \quad (4.14)$$

where $M_{\ell\ell'}$ is the mode-coupling matrix, F_ℓ is an ℓ -space transfer function due to filtering, B_ℓ is the beam transfer function, and \hat{N}_ℓ is the noise bias (Wandelt & Górski, 2001; Hivon et al., 2002).

The mode coupling matrix can be computed analytically, and accounts for the effects of the different masks, inverse-variance weighting, etc. when calculating the true full-sky power spectrum from the pseudo power spectrum of a map. Computing it is one of the more computationally intensive steps of the power spectrum estimation (other than the covariance matrix calculation).

4.2 Power Spectrum Pipeline

In this section, we walk through the various steps of the power spectrum analysis pipeline for DR4, as published in Choi et al. (2020). While the vast majority of what is written here will also apply to the analysis for DR6, some of the details have already changed, and some have yet to be determined.

The power spectrum code³ used for DR4 uses the now-standard curved sky pseudo- C_ℓ approach to account for the incomplete and nonuniform coverage of the sky and beam smoothing (Hivon et al., 2002; Kogut et al., 2003; Brown et al., 2005). It was tested against the power spectrum estimator code used in Louis et al. (2017) in the flat sky limit, against a suite of simulations, and against the publicly available Simons Observatory curved sky power spectrum pipeline `PSpipe`.⁴ The different codes are in excellent agreement, and the remaining difference between the curved sky codes is $< 0.01\sigma$.

4.2.1 Enumeration of the Spectra

For each set (season/region/array) of maps for DR4, the data are split temporally to have $n_d = 4$ maps each for I , Q , and U Stokes parameters. This is done so that we only compute cross-spectra and thus avoid noise bias.⁵ For the AA region (shown in Figure 1.12), due to its shallow depth, $n_d = 2$. In the same season, regions observed by different arrays have the same temporal intervals for the data splits. We compute the cross power spectrum of each pair of the data-split maps, but perform the averaging differently depending on the array and season. Specifically, a single-array power spectrum at one frequency in one season is computed from the unweighted

³The pipeline for computing the spectra was originally written for Choi & Page (2015). Its accuracy was confirmed by comparing it to an independent code from Kendrick Smith.

⁴GitHub repository: <https://github.com/simonobs/PSpipe>.

⁵If two maps with the same noise are cross-correlated, the resulting power spectrum contains the noise power. Cross spectra avoid this bias.

average of the $n_d(n_d - 1)/2$ cross data-split power spectra. For different arrays (or different frequencies within an array) in one season, we only exclude the cross spectrum between the data split maps of the same temporal period (since these maps have very similar noise) and average the $n_d^2 - n_d$ cross data-split spectra. For spectra from different seasons, we average all n_d^2 cross data-split spectra. (Since there is no temporal overlap in this case, there is no need to worry about noise bias.) This selection of $n_d(n_d - 1)/2$, $n_d^2 - n_d$, and n_d^2 spectra for averaging in different cases is depicted graphically in Figure 4.1.

In general, for power spectrum computation, the more splits the better. The limiting factor is the number of times we can split up the data while still having sufficient coverage to make reasonable maps. The larger the number of splits, the fewer data are thrown out when averaging cross data-split spectra, as visible in Figure 4.1.

4.2.2 Masking the Foregrounds

When selecting the data used to make maps, we exclude regions of contamination from the Moon and planets (as well as the Sun when looking at daytime data), so these sources of contamination need not be considered further.

Different foreground components in both intensity and polarization enter at different angular scales. For large angular scales, we apply the *Planck* “100 GHz cosmology mask” to mask regions containing large Galactic foregrounds (Planck Collaboration I, 2020) and then fit for residuals as described in Section 13 of Choi et al. (2020). At smaller angular scales, bright point sources dominate. We coadd 150 GHz maps in the deep regions (D56, D1, and D8) and find point sources with a 5σ flux greater than 15 mJy in intensity using a matched-filter source-finder. These are then masked both in the intensity and polarization maps at $5'$ ($8'$) radius and apodized beyond the mask edge with a sine function that extends over $10'$ ($15'$) at each source position for

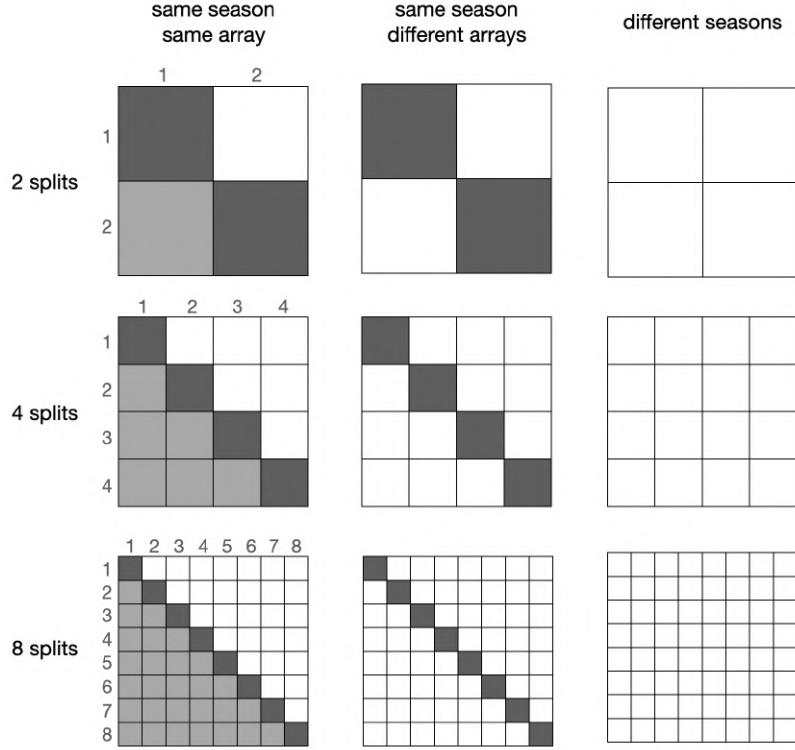


Figure 4.1: Graphical depiction of the spectra we use when averaging the cross-spectra for different data splits, for different scenarios. Each large square is split up into an $n_d \times n_d$ grid, representing the spectra that can be computed (e.g. split #1 \times split #1, split #1 \times split #2, etc.). The small white squares represent the spectra included in the average. The small dark gray squares are the spectra we exclude because they are computed with two maps from the same temporal period. The small gray squares are spectra that are redundant (they are identical to the remaining white squares).

the 150 (98) GHz maps. For the shallower and wider regions, AA and BN, we do the same but with a flux cut of 100 mJy.

As explained in [Aiola et al. \(2020\)](#), there are roughly 400 extended sources over the full region that are identified with external catalogs that are also masked. However, we do not mask Sunyaev-Zel'dovich (SZ) clusters ([Sunyaev & Zeldovich, 1972](#)) and instead include them in our foreground model.

4.2.3 Spatial Window Function

We select regions with good noise properties as follows. ACT’s constant elevation scan trajectories project into the maps as almost straight lines. When the same region is observed at different elevations or while setting as opposed to rising, scan lines are rotated with respect to the original direction and the target region is said to be cross-linked.

In general, the better the cross-linking, the better our map solutions reflect the true sky. One way to understand this is that the noise in the scanning direction of a single TOD, a roughly 10-minute stretch of time-ordered data, is large and localized in 2D Fourier space. Observations at multiple cross-linking angles improve the rotational symmetry of the noise in the Fourier plane, with the improvement related to the amount of cross-linking. We account for the degree of cross-linking in the simulations and in the spatial window as described next.

To parametrize the degree of cross-linking in a region, we make “cross-linking maps” by summing up the number of observations at each pixel by representing the projected scan angle as a polarization angle. For example, scans that project to horizontal (e.g., RA) or vertical directions (e.g., Dec) on the sky result in a $+Q_x$ or $-Q_x$ cross-linking map. Stokes I_x in this case corresponds to the usual hit-count map. We then compute the level of cross-linking from $P_x = \sqrt{Q_x^2 + U_x^2}/I_x$. $P_x = 1$ means no cross-linking (just one scan direction). From there, the smaller P_x , the better the cross-linking. We set thresholds to select regions with a minimum amount of cross-linking for each region. For instance, a threshold < 0.7 for D56 retains most of the regions observed with two orthogonal scans. For D8, which is located in a particular declination where sky rotation does not allow orthogonal scans, we investigated a threshold of 0.99 but eventually dropped the region from the cosmological analysis due to its poor cross-linking. We set the same threshold for the cross-linking maps from all seasons and

arrays for each region, set all pixels below (above) the threshold to be 1 (0).

The second step in determining the boundary of the spatial window is to threshold the noise maps in percentile to exclude the noisiest regions. Finally, we take the common boundary mask for each region (across all seasons and arrays), skip $\sim 1^\circ$ around the edge, then apodize over another $\sim 1^\circ$ around the edge with a sine function to get what we refer to as the footprint for each patch.⁶ The procedure is shown graphically in Figure 4.2. This process ensures the maps of a given region from different seasons and arrays are each weighted with the corresponding inverse variance weights, while sharing the same overall boundary. (The maps of a given region from different seasons and arrays start with slightly different boundaries due to the small pointing offsets between arrays on the telescope.)⁷ The cross-linking and inverse variance thresholds used for the different regions are listed in Table 4.1.

Table 4.1: Summary of DR4 Region Footprint Properties

	D1	D5	D6	D56	D8	BN	AA
Area (deg ²)	66	64	61	565	197	1,837	11,920
Area PS (deg ²)	23	20	20	340	120	1,400	3,600
Cross-linking threshold	0.96	0.72	0.72	0.8	0.99	0.9	—
Noise threshold	0.23	0.3	0.3	0.2	0.04	0.08	—

Notes: The top first row gives the area assuming uniform weighting out to the edge of the spatial apodization. This corresponds to the visual impression of the patches in Figure 1.12. In the second row, the area denoted “PS” is that used for power spectrum estimation after following the procedure described in this section to obtain the spatial window. The third row gives the cross-linking upper bound for selecting the spatial window. (Uniform cross-linking corresponds to an index of zero, and no cross-linking corresponds to unity.) In the bottom row, the noise threshold for selecting the spatial window is expressed as a fractional rate. For example, for D1 the 23% highest noise pixels are dropped. There are no thresholds for AA, because the regions were hand-picked by visually examining the noise and cross-linking maps.

⁶An unapodized mask would produce “ringing,” resulting in power in the spectrum of the mask, even at the highest scales.

⁷As mentioned in Section 5.4.1, for DR6 we no longer require different arrays to share the same boundary for a region, and instead define a different footprint for each detector array and frequency. This saves us from discarding data at the edge of the maps. This is much simpler to do for DR6, where we only consider one large region of sky (AA), as opposed to DR4, where there were already 13 footprints defined for different regions of sky.

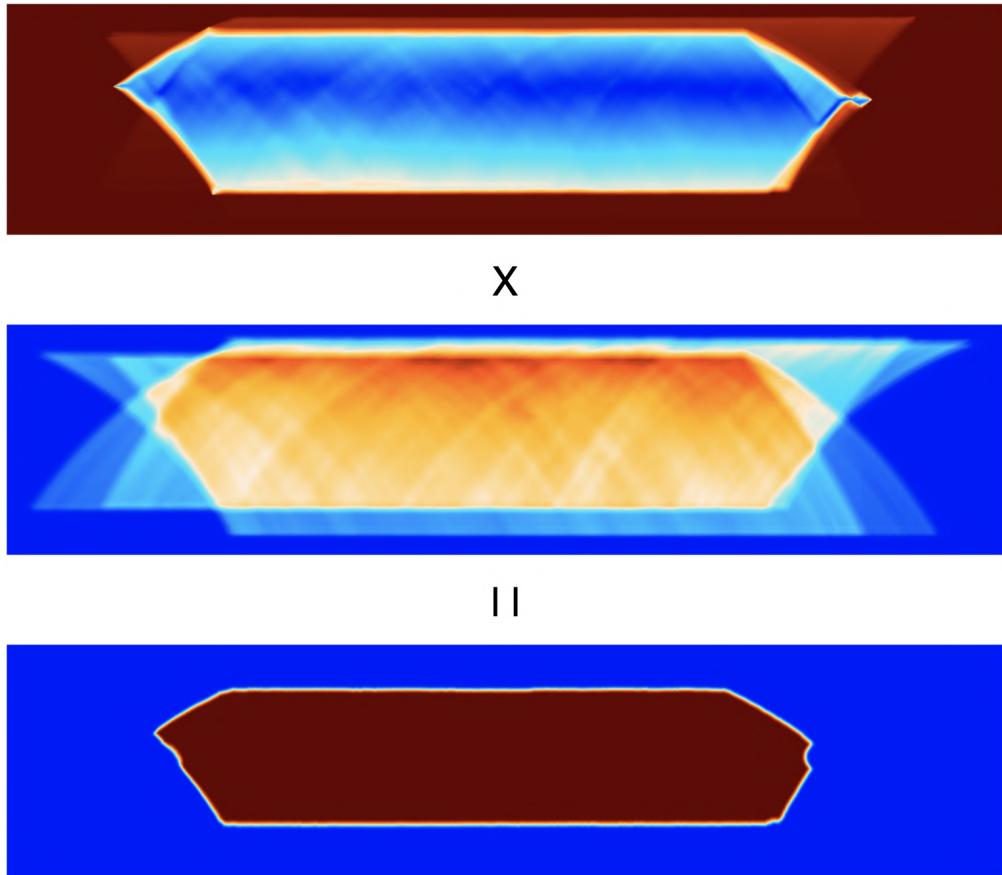


Figure 4.2: Illustration of the process for obtaining the footprint for the D56 region, for S15 PA1 data. The top panel shows the cross-linking map; the middle panel shows the normalized inverse noise variance of the coverage; and the bottom panel shows the spatial window function. The color scale spans from 0 to 1, in blue to red. After applying the cross-linking threshold to the top map and the noise threshold to the middle map, they are multiplied together to obtain the bottom map. Note that in the inverse variance map, large values (red) are most desirable, so the noise threshold cuts the blue areas, whereas in the cross-linking map, it's the opposite: blue values are best, so the cross-linking threshold cuts the red areas. See Figure 1.12 for the size and location of the region.

Once we have the footprint for a region of sky, we multiply it by the apodized point source mask and the *Planck* cosmology mask from Section 4.2.2, and the inverse variance map for the season/patch/detector array/frequency of interest (for weighting)⁸ to obtain the spatial window we will multiply with each map when computing the power spectrum. This process for obtaining the spatial window is outlined in Figure 4.3. The area of each region, both before and after applying the spatial window, is listed in Table 4.1.

4.2.4 Ground Pickup and the Fourier-Space Filter

ACT scans horizontally at different azimuths at different times of the day and year. The contamination from the ground, or other scan-synchronous pickup, is projected as constant declination stripes in the sky maps. In [Das et al. \(2011\)](#) and [Louis et al. \(2017\)](#), Fourier modes with $|\ell_x| < 90$ and $|\ell_y| < 50$ were masked to remove this ground contamination in the data.⁹ An exact mode coupling matrix was computed to account for this Fourier mask in the flat-sky power spectrum estimator code used in [Louis et al. \(2017\)](#). Because the ground contamination is projected horizontally on the equatorial coordinates (in RA direction), we continue to mask these contaminated modes in Fourier space, the space in which the modes stay localized. Then we estimate the power spectra of the filtered maps with the curved-sky code, and correct for the loss of power due to filtering with a one-dimensional transfer function determined with simulations as described in Section 4.2.5.

⁸For DR6, we apply uniform spatial weighting instead of the inverse variance weighting from DR4. More discussion about the optimal weighting to use in different situations can be found in Section 5.4.3.

⁹For the purpose of removing ground contamination, we work in the flat-sky approximation, so here ℓ_x and ℓ_y are used to denote the components of flat-sky wavevectors.

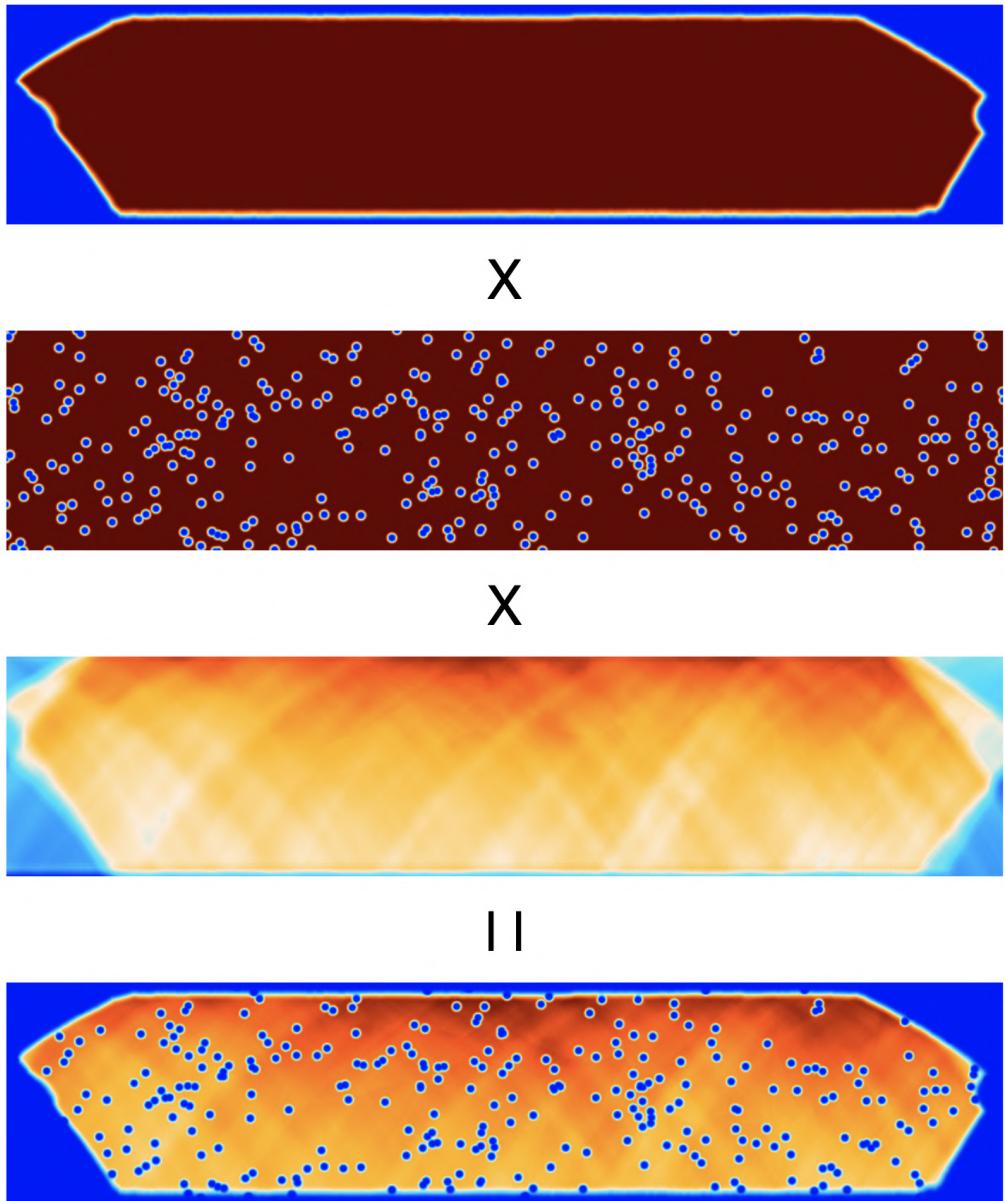


Figure 4.3: Illustration of the process for obtaining the spatial window for use in the power spectrum analysis, for S15 PA1 D56. From top to bottom, the apodized footprint for the region we are analyzing (first image) is multiplied with an apodized point source mask (second image) and a normalized inverse-variance map for weighting (third image), in order to obtain the final spatial window (fourth and final image). The color scale spans from 0 to 1, in blue to red. The *Planck* cosmology mask, not shown, is also applied (it has no overlap with the D56 region).

4.2.5 Fourier-Space Filter Transfer Function

We estimate this transfer function by comparing the power spectra of the simulated maps before and after applying the Fourier-space filter. In principle, a full transfer matrix describing the possible bin-to-bin power transfer is needed. We test the necessity of this level of complexity by examining the consistency between the transfer functions estimated with two differently shaped spectra, Λ CDM TT and EE power spectra, and find that a simple 1D implementation is sufficient for our needs. We also count the number of modes removed by our filter in the 2D Fourier plane to check the transfer function analytically.

Both the simulation-based and analytic estimates for the Fourier-space filter transfer function are shown in Figure 4.4.

We find that the transfer functions for TT and EE power spectra are consistent with each other to $< 1\%$. The transfer functions from simulations are consistent with the analytic estimate to $< 1\%$ everywhere, and for $\ell > 300$ the simulations are consistent with the analytic estimate to $< 0.5\%$.

In addition to directly acting on the TT, TE, and EE spectra, Fourier-space filtering of the Stokes Q and U maps can lead to mixing of the E and B modes.¹⁰ This is characterized by an off-diagonal E-B mixing term, estimated to be $< 1\%$ at $\ell \geq 200$ (and $< 0.1\%$ for $\ell \geq 600$). For DR4, this small mixing correction was only necessary for the sensitivity levels achieved in D56.

¹⁰We note this differs from the usual E-B mixing due to incomplete sky coverage in the pseudo- C_ℓ approach, which is analytically corrected with the mode coupling matrix.

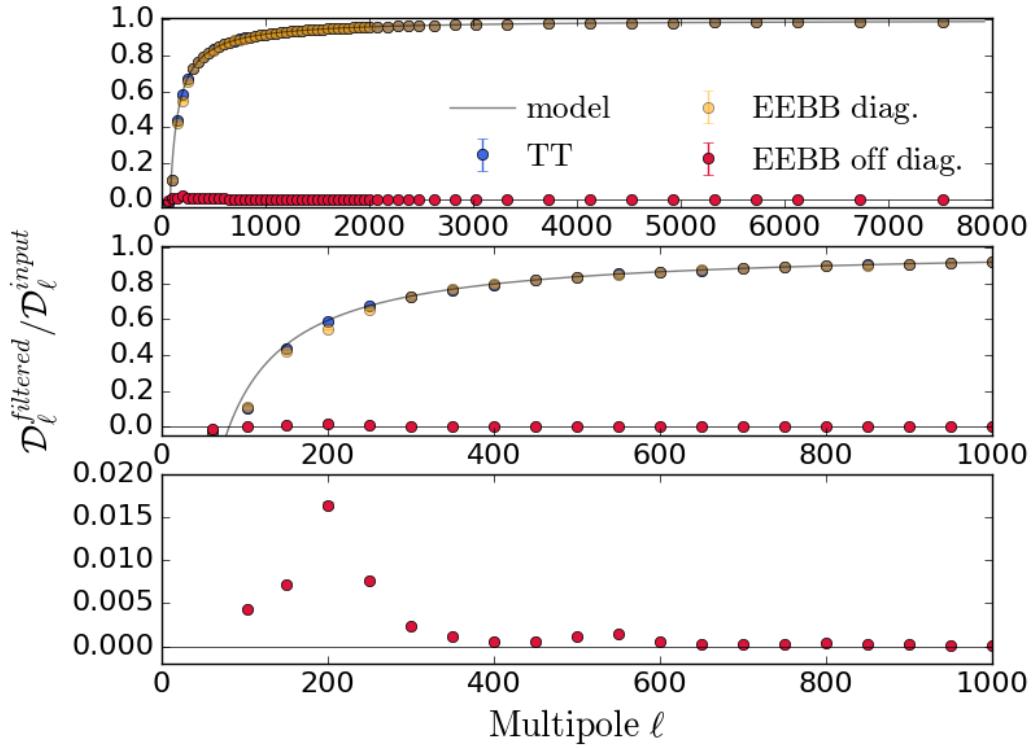


Figure 4.4: The Fourier-space filter transfer function, estimated from simulations (data points) and an analytic estimate (gray line), for the D56 region. The bottom two panels show an expanded view of the low- ℓ region from the top panel. The transfer function for TT, and the EE/BB transfer matrix elements, both diagonal and off-diagonal (which corresponds to EE/BB mixing), are computed at each ℓ bin.

4.2.6 Mapping Transfer Function

One of the attractive features of maximum likelihood mapmaking is that it produces unbiased maps. In other words, the power spectrum of an unbiased map should not need to be corrected for the mapping process. However, we add one operation to our mapmaking that does slightly bias the power spectrum. As described in Section 2.1.6, maximum-likelihood ground template maps are made in azimuth-elevation coordinates, and deprojected from the TODs. This deprojection removes the strongest ground signals and results in a flat transfer function of $\gtrsim 0.997$ for most regions at $350 < \ell < 8000$ (and smaller at $\ell < 350$, which we do not consider in this analysis). Smaller regions (D1, D5, and D6), which contribute small weight in the

total statistics of the coadded spectra, have a more complex shape with a dip to ~ 0.98 near $\ell = 4000$. Each spectrum is divided by the appropriate transfer function. The uncertainties on the transfer functions were investigated with simulations and found to be negligible ($< 0.02\sigma$).

4.2.7 Calibration

Before combining spectra, for each season/region/array/frequency, we calibrate them to the *Planck* temperature maps, weighted by the ACT spatial window (Section 4.2.3), over the range $600 < \ell < 1800$ at 98 GHz and 150 GHz. There is no separate calibration for TE or EE, but an overall polarization efficiency is allowed to vary later in the power spectrum likelihood. We calibrate through cross correlation as described in [Hajian et al. \(2011\)](#) and [Louis et al. \(2014\)](#).

4.2.8 Covariance Matrix

The covariance matrix for the power spectra is computed using noise simulations for the diagonal, pseudo-diagonal,¹¹ and “diagonal-plus-one” terms, and analytic terms for smaller corrections (see Section 6.4 of [Choi et al. \(2020\)](#)).

4.2.9 Binning and the Band Power Window Function

The power spectra are computed in bins b , within which we average over a range of ℓ s. The width of each bin is $\Delta\ell_b$, and the binned spectra are referred to as C_b .

Our power spectra are computed in 59 bands with centers spanning from $\ell = 21$ to $\ell = 7525.5$. Our cosmological analysis is based on the $n_{\ell,c} = 52$ bands from $\ell = 350.5$ to $\ell = 7525.5$ in the TE and EE spectra, and from $\ell = 600.5$ to $\ell = 7525.5$ for TT, as discussed further in Section 10 of [Choi et al. \(2020\)](#). (Here, the subscript “ c ” is

¹¹These are the diagonal terms in sub blocks of the matrix.

for “cosmology.”) The lower bound was selected as part of the blinding procedure, motivated in part by the k -space cutoff in the maps corresponding due to the Fourier-space filter described in Section 4.2.5, and in part by our experience in [Louis et al. \(2017\)](#), where the lower bound was $\ell = 500$ for TT and $\ell = 350$ for TE/EE. The maximum ℓ is determined by the signal-to-noise ratio.

The band power window functions are used to bin the theory power spectra to compare to the data. They depend on the mode coupling matrix and ℓ -space binning scheme and are slightly different for each region. Band power window functions are coadded using the power spectrum covariance matrix, which takes into account the weight variations among different regions. This coadded band power window function is used in the likelihood.

4.2.10 Coaddition

Given a vector C^{XY} (where X and Y are each either T , E , or B) encompassing several spectra (for example, spectra from different detector arrays at 150 GHz covering the same region of sky), the maximum likelihood combination of these spectra (commonly referred to as a “coadd”) C_{coadd}^{XY} can be computed with

$$C_{\text{coadd}}^{XY} = (P^T \Sigma^{-1} P)^{-1} P^T \Sigma^{-1} C^{XY}, \quad (4.15)$$

where Σ is the covariance matrix for all the spectra pairs and P is a projection matrix populated with 0s and 1s.

The covariance matrix for the coadded spectrum is then

$$\Sigma_{\text{coadd}} = (P^T \Sigma^{-1} P)^{-1}. \quad (4.16)$$

4.2.11 Consistency Tests

With multiple seasons of observations made with multiple arrays, there are many possible pair-wise combinations of data that may be tested for consistency. There are three broad classes of checks as discussed, for example, in [Louis et al. \(2019\)](#). One entails consistency between two maps, a second between the power spectra of the two maps, and a third between the power spectrum of one map and its cross spectrum with the other. We consider only the first two of these. In general, when there is a measurable signal in the maps, the consistency between maps is the most stringent test to pass. However, depending on the systematic effect, for example a multiplicative bias, the power spectrum null can be more stringent. Thus, we present both.

The power spectrum of the difference, or “null”, between maps A and B is,

$$\hat{C}_b^{A-B} = \hat{C}_b^{AA} + \hat{C}_b^{BB} - 2\hat{C}_b^{AB} , \quad (4.17)$$

where $\hat{C}_b = C_b + N_b$ is the measured quantity for example for an auto spectrum, C_b is the underlying power spectrum, and N_b is the noise. For this to be a true map-level null power spectrum, the same spatial window function needs to be used in estimating each of the power spectra on the right-hand side. We have constructed our spatial windows, as described in Section 4.2.3, to be roughly the same for the different seasons/arrays for one region, even though it required cutting otherwise good data. As a result, the same spatial window functions and maps that are used for the consistency tests are used for the cosmological analysis.

The general case for the covariance matrix of \hat{C}_b^{A-B} is worked out in [Das et al. \(2011\)](#) and [Louis et al. \(2019\)](#). When assessing the covariance, we work directly with power spectra after the n_d splits, described in Section 4.2.1, have been combined. In the limit that C_b^{AA} and C_b^{BB} cancel in Equation 4.17, in other words that the underlying

spectra and spatial windows match, and that the noise in each of the n_d splits is the same, the variance for the power spectrum of the null map for each bin b is given by:

$$\begin{aligned} (\Delta \hat{C}_b^{A-B})^2 &= \frac{1}{\nu_b} \left[\frac{2}{n_d(n_d-1)} ((N_b^{AA})^2 + (N_b^{BB})^2) \right. \\ &\quad \left. + 4 \frac{N_b^{AA} N_b^{BB}}{n_d^2} \right], \end{aligned} \quad (4.18)$$

where ν_b is the number of modes in bin b , given by $\nu_b = (2\ell + 1)\Delta\ell_b f_{\text{sky}} t_b$ (where the last term t_b , is the transfer function due to the mapping process and the Fourier-space filter), and N^{AA} is the noise in one of the n_d splits that goes into determining the spectrum of map A (and likewise N^{BB} for map B).

In contrast, the variance for the difference between power spectra is given by

$$\begin{aligned} (\Delta[\hat{C}_b^{AA} - \hat{C}_b^{BB}])^2 &= \frac{1}{\nu_b} \left[\frac{2}{n_d(n_d-1)} ((N_b^{AA})^2 + (N_b^{BB})^2) \right. \\ &\quad \left. + \frac{4}{n_d} C_b (N_b^{AA} + N_b^{BB}) \right]. \end{aligned} \quad (4.19)$$

Here again we have assumed that the power spectra of the underlying signal and spatial windows match so that $C_b = C_b^{AA} = C_b^{BB}$.

In practice, as opposed to computing Equation 4.18 directly, we compute the null power spectrum covariance with the covariance matrix Σ from Equation 4.16 and a projection matrix with elements of the form $P = (1, 1, -2)$. Similarly, for the power spectrum difference, we use $P = (1, -1)$.¹² This formalism provides a general and compact way to compute the multiple different combinations of elements of the covariance matrix that enter the consistency tests. Also, using the full covariance matrix, it can be generalized to consistency checks between power spectra of non-overlapping re-

¹²These projections are appropriate for TT and EE. The form for TE for the map spectra difference is $(1, 1, -1, -1)$.

gions as the null power spectrum error bars then contain the representative signal variance. Lastly, we correct for the different levels of point source contamination when comparing the deep and wide regions.

Once the null spectrum is found, either from individual spectra or combinations of spectra, we compute χ^2 of the $n_{\ell,c} = 52$ bins. We then compare the distribution of this χ^2 for all the difference spectra to the distribution for nulls with 500 simulations each. Figure 4.5 shows the distribution of this χ^2 for the 438 difference spectra compared to the distribution of nulls with 500 simulations for each. Additional comparisons of the data and inter-patch tests are shown in Appendix C of Choi et al. (2020). In summary, the spectra nulls within each region and the spectra nulls between regions are consistent with the nulls of 500 simulations for both types of tests.

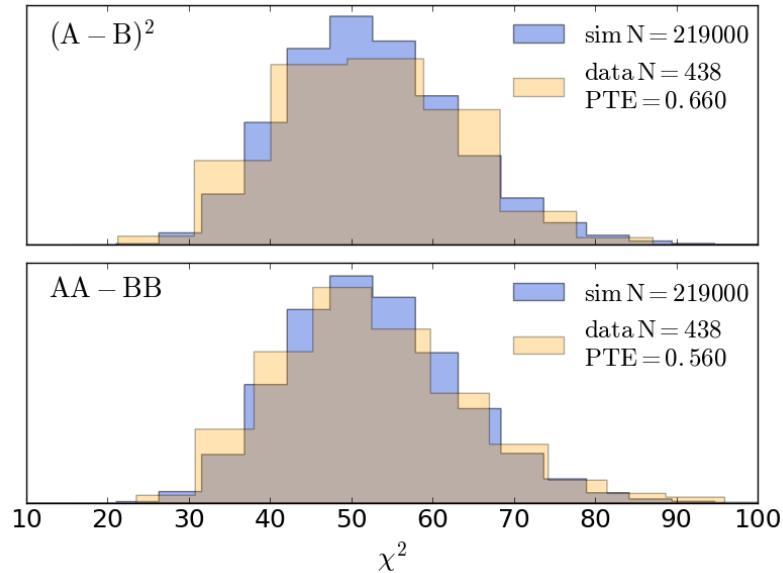


Figure 4.5: Summary of pair-wise consistency checks for all combinations of spectra from D5, D6, D56, D8, BN, and AA. D1 has one spectrum, so differences cannot be formed. The top panel shows the map-level differences, and the bottom panel shows the power spectrum differences. The expectation is $\chi^2 = n_{\ell,c} = 52$, the number of ℓ bins. The probability to exceed (PTE), the probability of obtaining a higher χ^2 than the measurement, is computed with the simulations. The histogram integrals are normalized to unity.

In a third form of consistency check, we compute χ^2 for each ℓ bin for the combination of 98, 98×150 , and 150 GHz spectra (where available) from all of D56, BN, D8, and AA. Each of these regions has more than one spectrum, so we can subtract the mean spectrum and examine the distribution around it. Because we remove the mean, residual foreground emission will be removed to some degree. This test is done using only the diagonal elements of the covariance matrix, although the cosmic variance and noise terms are appropriately separated, thus it is less rigorous than the others. For AA, we combine spectra from the five different windows into one spectrum per frequency combination. We also include D8 in this test. Then for one ℓ bin in TT, there are 20 D56, 9 BN, 9 D8, and 15 AA spectra. With 9 degrees of freedom per ℓ bin the expectation is $\chi^2 = 44$. A similar treatment for TE gives $\chi^2 = 71$. In the last step, we sum over the ℓ bins in each spectrum and find a reduced χ^2 of 1.03, 0.94, 1.03, 0.86, 0.96, 1.13 for TT, TE, EE, TB, EB, and BB respectively. This shows that the distributions about the means are well-behaved when many spectra are averaged together. The power of this test is that one checks for consistency with uncertainties of comparable size to those used in the cosmological analysis.

4.3 DR4 Results

Figure 4.6 shows the raw 98 and 150 GHz spectra that are inputs for the likelihood. These spectra are obtained by coadding all the spectra for a given frequency (98 or 150 GHz) within either all the deep or wide patches of sky. A number of features are evident. *a*) There is a power deficit in TT compared to the best-fitting model for $\ell < 600$, especially in the wide region. *b*) The agreement between 98 GHz and 150 GHz is clear. *c*) The Poisson point-source tail has a higher amplitude in the wide region. *d*) The error bars are smaller at high ℓ in the deep region when compared to the wide region, but at low ℓ the error bars for the wide region are comparatively

smaller. The latter is due to cosmic variance.

ACT’s agreement with the six-parameter Λ CDM model is comfortably within expectations.

4.4 Reproducing DR4 Spectra

In this section, we work on reproducing the DR4 spectra in a Jupyter notebook. Once we are satisfied with the results, we move on to using the `PSPipe` pipeline, which will be used for DR6.

4.4.1 `NaWrapper`

Along with the DR4 and DR5 data releases, the ACT Collaboration provided a set of Jupyter IPython notebooks to facilitate use of the data products. Both the data products and the notebooks are described in [Mallaby-Kay et al. \(2021\)](#). In the notebooks, we give examples of analyses that can be performed and recreate several key plots from the original DR4 and DR5 papers. These notebooks are available via the ACT GitHub repository.¹³

In notebook 8, we use the `NaWrapper` code base to perform a power spectrum analysis similar to that carried out for DR4 in [Choi et al. \(2020\)](#), using maps and associated data products that were made publicly available. `NaWrapper` uses the `NaMaster` code from [Alonso et al. \(2019\)](#).

We carry out this analysis in notebook 8 for the S15 deep56 PA1 data and compare the results to the equivalent spectrum computed as part of the DR4 analysis pipeline. As shown in Figures 4.7 and 4.8, the spectra from notebook 8 are generally consistent with the DR4 results, but there are visible differences. Since both analyses use the

¹³GitHub repository: https://github.com/ACTCollaboration/DR4_DR5_Notebooks

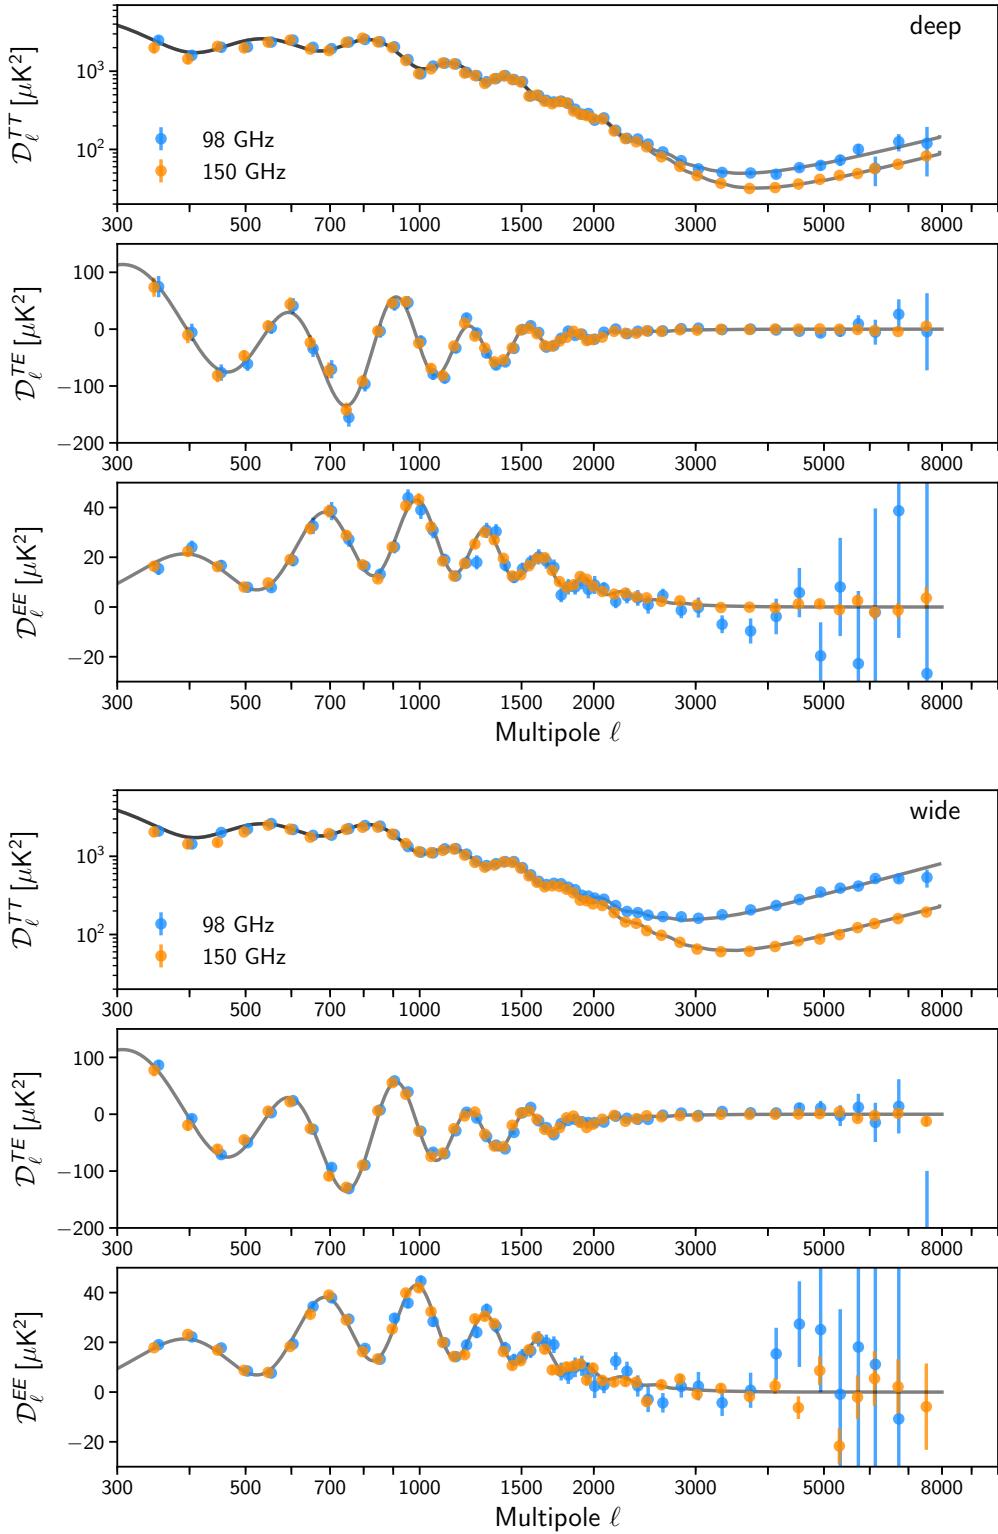


Figure 4.6: The deep (top) and wide (bottom) spectra used in the likelihood. The gray line is the best fit ΛCDM plus foreground model for ACT only. The 98 \times 150 GHz spectra are omitted for clarity.

same data, the discrepancies must arise from differences in the analyses. In order to ensure we fully understand the details of the DR4 pipeline before beginning to build the pipeline for DR6, we attempt to track down the sources of the differences.

Here, we make incremental changes to the analysis done in notebook 8 in order to better understand the source of the differences between the resulting spectrum and the DR4 results. For the remainder of Section 4.4.1, the colors and styles of lines are consistent across all plots. So, for instance, the dashed green lines in Figures 4.10, 4.14, and 4.16 all represent the same data. Note that the y -axis scales, however, change from one plot to the next. In general, as the analysis progresses and our residuals with respect to the DR4 spectra shrink, the plots zoom in around $y = 0$.

We only focus on the amplitude of the spectra here, not the size of the uncertainties. In the notebook, two options for computing uncertainties are explored: using the standard error from the six computed cross-spectra (computed with different combinations of the 4 splits) or estimating the uncertainties analytically, using the covariance matrices computed by `NaWrapper`. We do not expect these uncertainties to agree precisely with those for DR4, since DR4 included information from simulations, as well as sources of systematic errors, such as the beams, calibration, and transfer function uncertainty.

Also, while plots are only shown for TT, the results for TE and EE are similar, and when a statement is made about the spectra being in close agreement, this has been verified for polarization as well.

In this section, we identify several changes which could be made to notebook 8 in order to get the resulting spectra to agree more closely with those from DR4. These changes would all be relatively simple to implement. In some cases, however, it may be preferable to not make these modifications, as notebook 8 is intended to serve as an introduction to power spectrum estimation, and we may not want to bog down its

users with all the minutiae of the analysis at such an early stage.

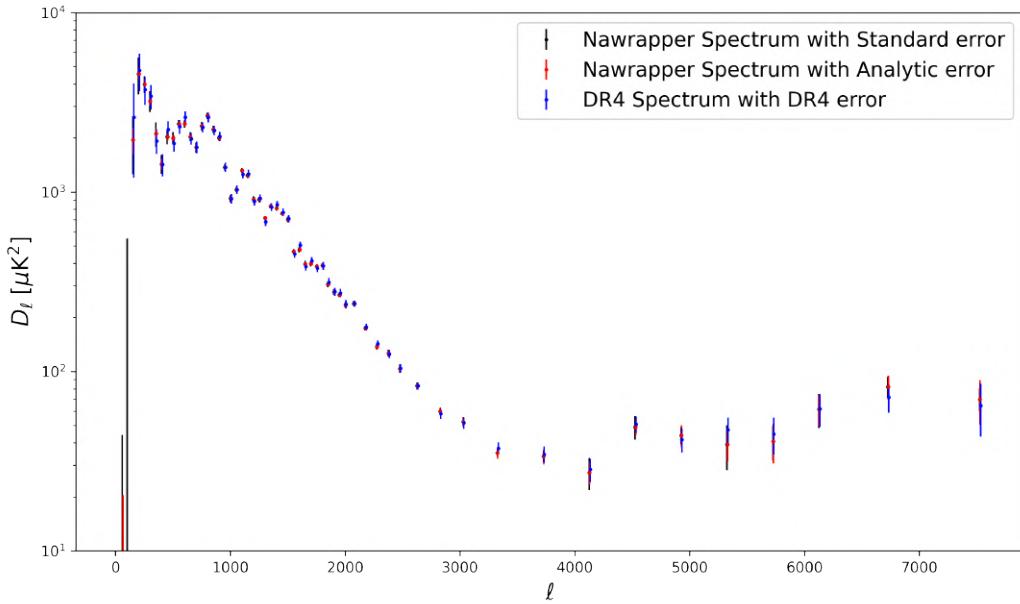


Figure 4.7: The result of estimating the power spectrum in notebook 8 using all 4 splits of data. This represents the default settings, where some of the required data products are read from disk due to computationally intensive calculations which can be difficult to run on a laptop. The difference between the two spectra shown here is plotted in Figure 4.9.

Number of Splits

When exploring the analysis choices made in notebook 8, the largest effect considered is the number of data splits analyzed. Figure 4.7 shows the result of running the notebook with its default settings. If we comment out the lines where pre-computed data products are read in and instead compute them ourselves with the default settings (which we may want to do to get a better sense of how the code works), we obtain Figure 4.8. The difference between the resulting spectra is due to the number of splits being used. The pre-computed products were generated using all 4 splits of data. However, this may be too computationally intensive to do on a laptop, so the default setting used when computing the data products in the notebook is to use only 2 out of the 4 splits. The main effect of using only half the available data is to

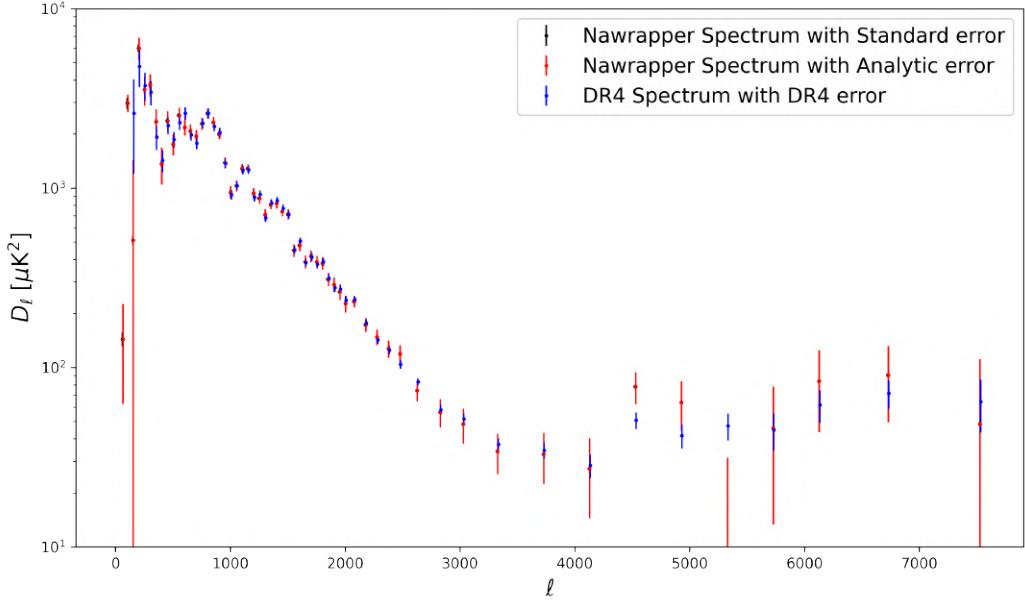


Figure 4.8: The result of estimating the power spectrum in notebook 8 using only 2 of the 4 splits of data. This option in the notebook is recommended if we want to run the calculations ourselves instead of reading in pre-computed products from disk, since running the calculations for all 4 splits may be too computationally intensive to run on a laptop. The difference between the two spectra shown here is plotted in Figure 4.9.

make the resulting spectra significantly noisier. If we change the notebook settings to use all 4 splits of data, and run it on either a powerful laptop or (ideally) submit it as a script on a computer cluster, we recover the spectrum from Figure 4.7. The residuals with respect to the official DR4 spectra are shown for both the 4-split and 2-split options in Figure 4.9.

Mask

While the 4-split spectrum is closer than the 2-split spectrum to the DR4 result, as shown in Figure 4.9, there is room for improvement.

The first analysis choice we examine is the mask, or spatial window function, we apply to the maps. The procedure used for building the DR4 spatial windows was described in Section 4.2.3. Similar to the approach for DR4, in notebook 8, we read

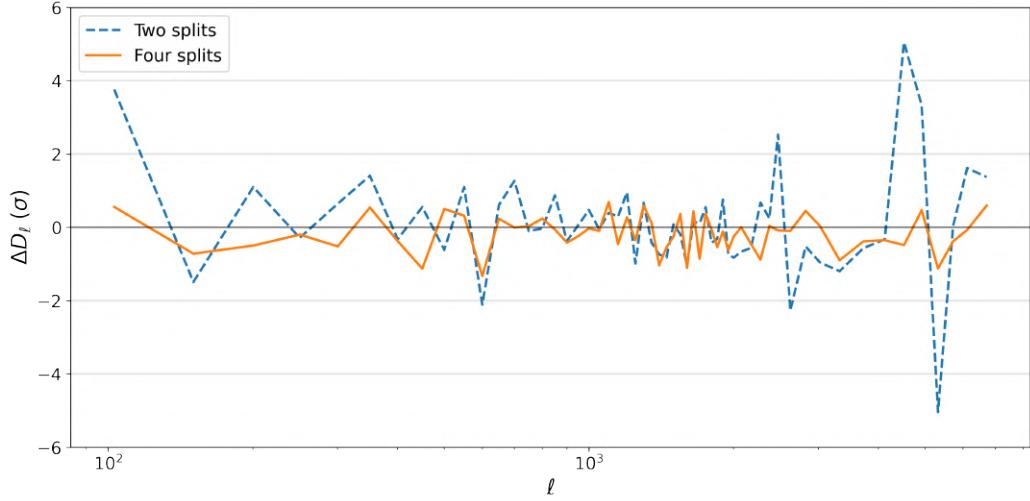


Figure 4.9: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when using either only 2 or all 4 splits of data, divided by the uncertainties on the DR4 spectrum. The variance for the 2-split spectrum is larger, since it uses half as much data as the 4-split spectrum.

in an apodized footprint outlining the patch of sky we are interested in, a point source mask, and the coadded (across all splits) inverse variance map. We apodize the source mask, then multiply all three products together (the apodized footprint, the apodized point source mask, and the inverse variance map) to form the final window function used for the analysis.

Instead of piecing together the spatial window function in this way, we can use the window function that was built as part of the DR4 analysis, although it is not one of the publicly available data products. The effect of using this pre-made mask is shown in Figure 4.10. The pre-made window is different from the one built in the notebook, as demonstrated in Figure 4.11.

The first difference is the outline of the footprint, as seen in Figure 4.11a. In the DR4 analysis, we skip 1° from the edges of the footprint when building the spatial window. This skips the edges of the survey area, where the noise is difficult to model, and skipping the edges of the area where the Fourier filtering is applied should prevent some aliasing from the filter from entering the data.

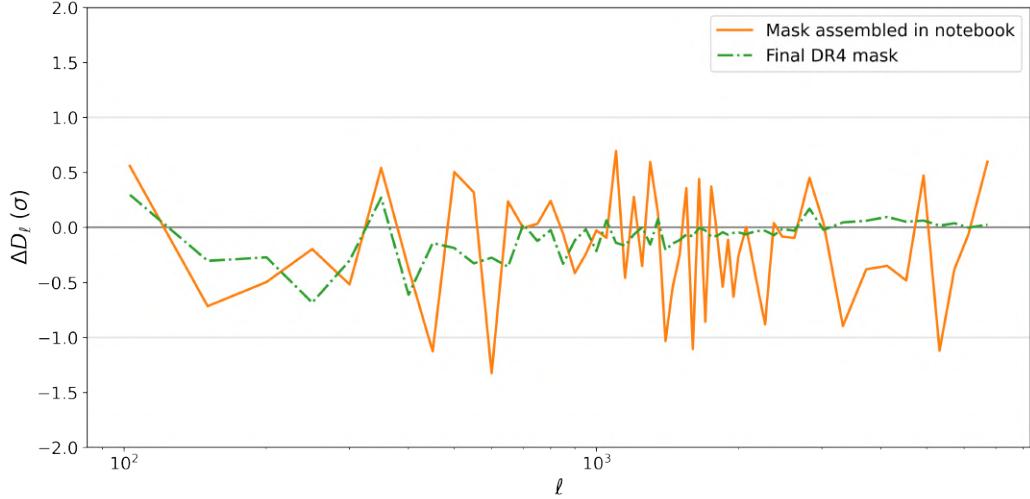


Figure 4.10: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when we build the spatial window function ourselves or use the one created as part of the DR4 pipeline, divided by the uncertainties on the DR4 spectrum.

The second difference is a different normalization for the inverse-variance-weighting portion of the masks, as is apparent in Figure 4.11b. The final inverse-variance maps released as part of DR4 have been calibrated (in $(\mu\text{K}_{\text{CMB}})^{-2}$), whereas at the stage of the analysis where the DR4 mask was built, the final calibration had not yet been applied.¹⁴ (The maps are roughly calibrated at this stage, but there is one final correction that hasn't been carried out. Although small, it does make a noticeable impact here.)

However, the ratio between the windows seen here is different from what one would infer based on the relative calibration of the sky maps. In the end, the source of this multiplicative difference between the inverse-variance maps is not too important. It does not impact the power spectra, and for the purposes of comparing the windows, we can fit and correct for the relative normalization. When this is done, as in Figure 4.11c, the masks seem to be in good agreement aside from the different footprint sizes. Upon closer examination, when the color scale is expanded in Figure 4.11d, we

¹⁴This comes up again later.

see some additional structure in the difference between the masks.

We can examine this more closely by zooming in, as in Figure 4.12. There are some widespread hash marks, indicating a difference in the inverse-variance weightings. This difference is most likely due to the inverse-variance map being smoothed in the DR4 pipeline, whereas this is not done in the notebook. Ideally, the map weight map we use should be smoothed, to avoid introducing small-scale noise in the power spectrum estimation.

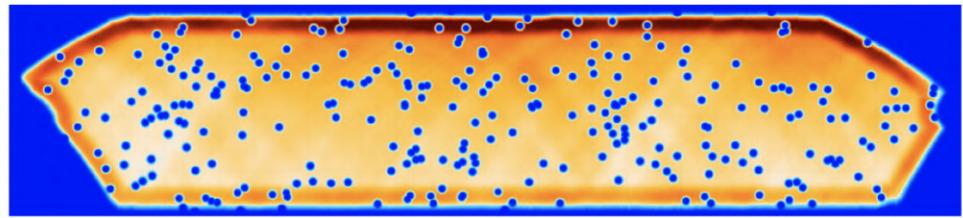
Note that in Figure 4.12 there does not seem to be a difference in the apodization around the masked point source. In notebook 8, we make use of the C^2 apodization from [Grain et al. \(2009\)](#), defined as

$$W_i = \begin{cases} -\frac{1}{2\pi} \sin\left(2\pi \frac{\delta_i}{\delta_c}\right) - \frac{\delta_i}{\delta_c} & \delta_i < \delta_c \\ 1 & \delta_i > \delta_c \end{cases} \quad (4.20)$$

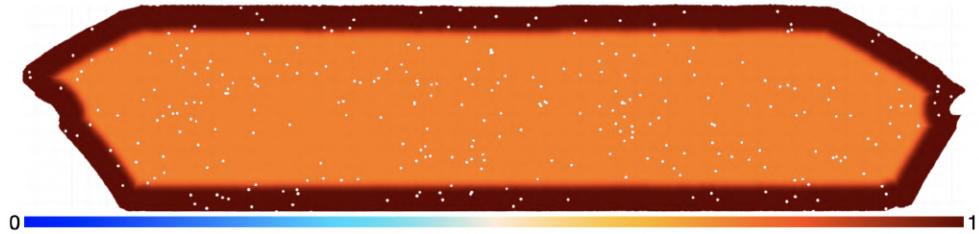
where W_i is the mask value for the i -th pixel, δ_i is the distance between the i -th pixel and the nearest zero pixel, and δ_c is the width of the apodization that smoothly interpolates between the mask value of zero surrounding the location of a point source and one elsewhere.

The radial profiles around the location of a masked source for the pre-made DR4 mask and for the point source mask we apodize ourselves, for two different choices of apodization width, are shown in Figure 4.13. A C^2 apodization of width $18'$, as is used in the notebook, is a good match to the apodization in the DR4 mask. This is despite the fact that in [Choi et al. \(2020\)](#) (and correspondingly in Section 4.2.2), the mask apodization is described as being a sine function that extends over $10'$ at each source position for the 150 GHz maps.¹⁵

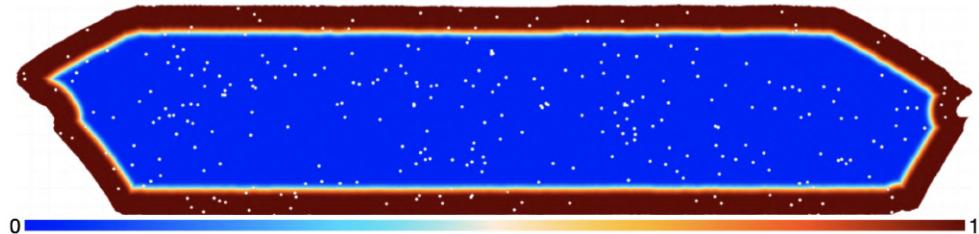
¹⁵The source of this discrepancy isn't known. One (leading) hypothesis is that it may be due to a difference in how the apodization width is defined.



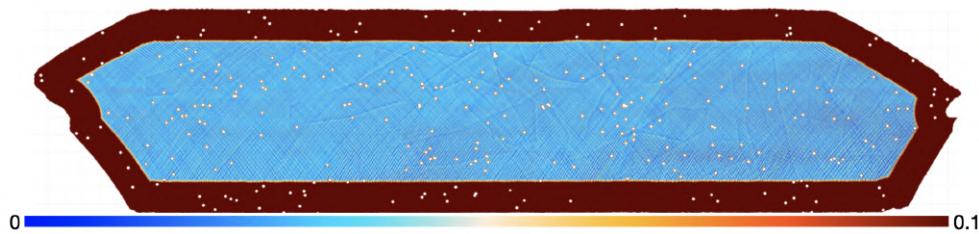
(a) The difference between the windows. (Since both windows are normalized differently, there is no meaningful color scale to be provided here.)



(b) The fractional difference between the windows, which shows that, other than a contour around the patch, the difference between the windows is primarily due to a different calibration.



(c) The fractional difference between the windows after renormalizing, as described in the text.



(d) The fractional difference between the windows after renormalizing, with a $10\times$ magnified color scale to accentuate the features in the middle of the patch.

Figure 4.11: Difference between the spatial window function assembled in notebook 8 and the one created as part of the DR4 power spectrum pipeline. In the maps above, the DR4 spatial window is subtracted from the notebook 8 window. (Hence, the red borders indicate that the DR4 mask is the more restrictive of the two.)

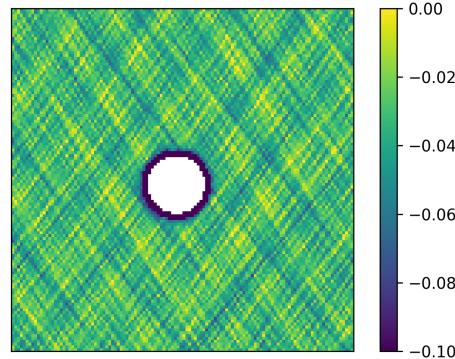


Figure 4.12: A $1^\circ \times 1^\circ$ cutout of the fractional difference between the windows after renormalizing, centered around a point source near RA = 32° , Dec = -5° .

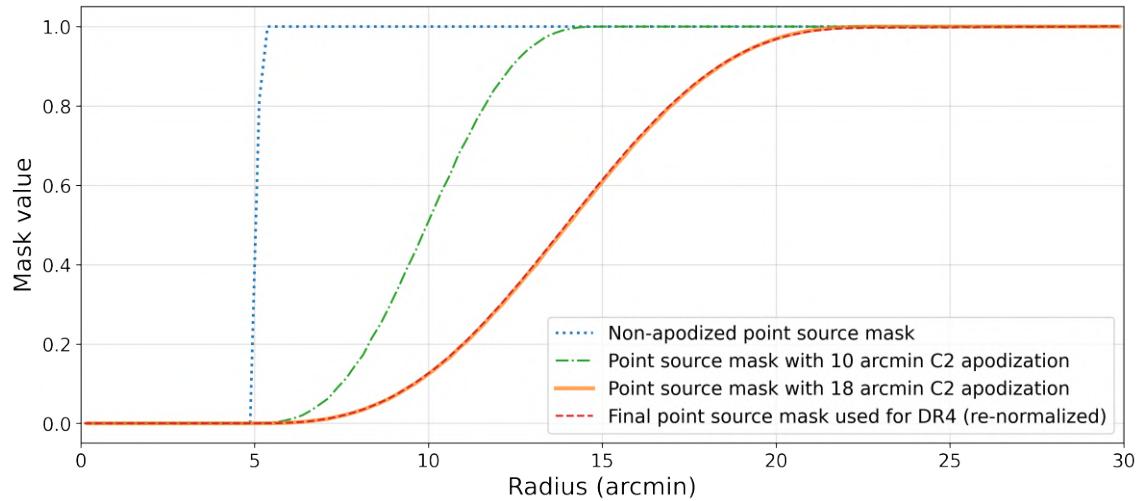


Figure 4.13: Radial profiles (centered on as masked source) for different point source mask apodizations.

At this point, we have tracked down the difference in the masks, and have a reasonable understanding of how they arise. In many cases, such as the width of the apodization function we use, there is a wide range of acceptable choices, although some may be suboptimal. Since we are trying to reproduce the DR4 results exactly, for the rest of this analysis, we will default to using the pre-made DR4 spatial windows.¹⁶

¹⁶Although an investigation into the optimality of the choices made when building the spatial windows may be relevant for future analyses.

Fourier-space Filtering

In the code used for the Fourier-space filtering for DR4, the map geometry was offset by one pixel in each dimension when defining the center pixel. Although not originally intended, this does not inherently introduce bias, but it does result in a slightly different Fourier-space filter and Fourier-space filter transfer function being applied, compared to the case with no center pixel offset. More specifically, two modes in k -space which we intended to cut (the most positive cut mode in k_x and the most negative cut mode in k_y) were not cut.

In notebook 8, we apply a Fourier-space filter without offsetting the center pixel, but we then correct for the transfer function computed for DR4, where there was a pixel offset. The transfer function for DR4 was computed with simulations, which is not something we can easily run in the notebook. There is nothing wrong with the pixel offset per se — we can apply any filter we like to the data, as long as we correct for its transfer function. The issue here is that we apply one filter (with no pixel offset) and correct for the transfer function of another (with a pixel offset). Since the difference is a mere one pixel, the effect should be small. Still, since we want to understand the source of minor discrepancies between the DR4 spectrum and the one we've computed, we consider it here. Instead of filtering the maps ourselves, we can use a pre-filtered version of the DR4 maps, although they are not publicly available.

The effect of this is shown in Figure 4.14.

If we make a version of Figure 4.14 where the differences are plotted in % rather than in σ , so plotting $(D_\ell^{\text{notebook}} - D_\ell^{\text{DR4}})/D_\ell^{\text{DR4}}$ instead of $(D_\ell^{\text{notebook}} - D_\ell^{\text{DR4}})/\sigma_\ell^{\text{DR4}}$, as is done in Figure 4.15, it becomes evident that the main source of difference between the DR4 spectrum and the spectrum we compute using the pre-filtered DR4 maps is some calibration mismatch.

The final maps released as part of DR4 are calibrated (in μK_{CMB}), but the pre-

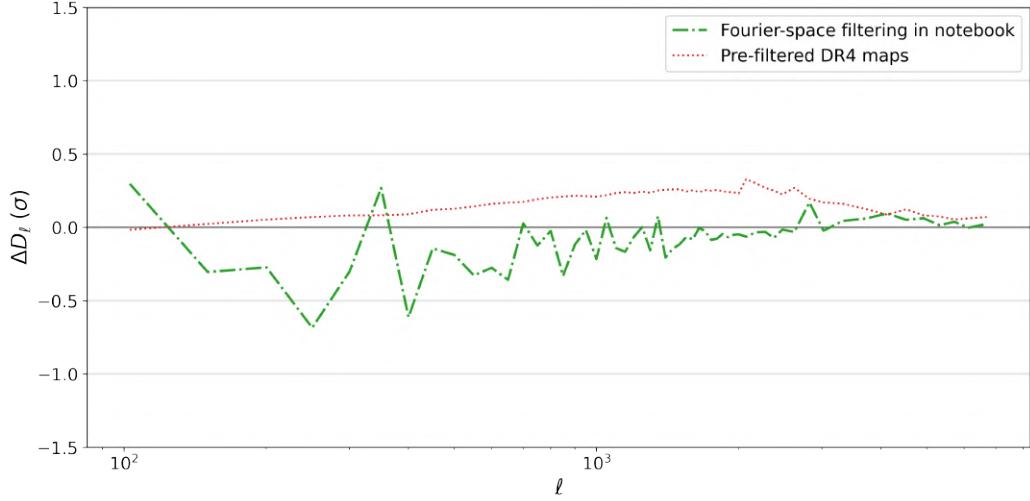


Figure 4.14: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when using a pre-filtered map or performing the Fourier-space filtering ourselves, divided by the uncertainties on the DR4 spectrum. See text for discussion.

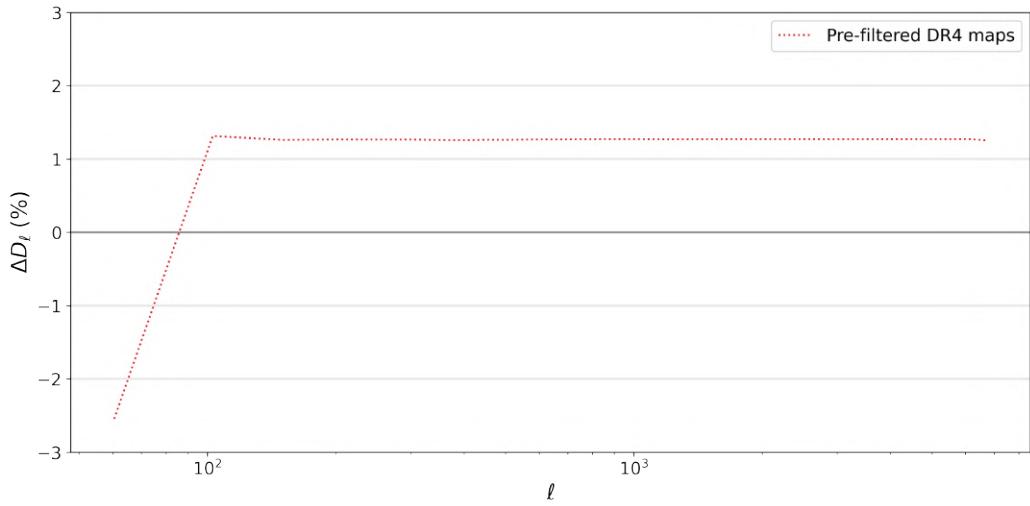


Figure 4.15: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when using a pre-filtered map, divided by the DR4 spectrum. The fact that the line above is (except for the first ℓ bin) well-described as an offset indicates that the difference between the spectra is in large part due to a calibration mismatch.

filtered maps used here, which are generated earlier on in the analysis, aren't. (They are roughly calibrated, but one final correction is missing, which is large enough to matter here.) Thankfully, it's simple to look up the relevant calibration factor (0.9874,

in this case), and apply it to the computed spectrum. With the proper calibration, the difference between the DR4 spectrum and the spectrum we compute using the pre-filtered DR4 maps becomes $< 0.002\sigma$.

Pixel Offset

Instead of using pre-filtered maps, we can Fourier-space filter the maps ourselves while also offsetting the map geometry by one pixel in each dimension, as was done for DR4.¹⁷ The result of this is shown in Figure 4.16.

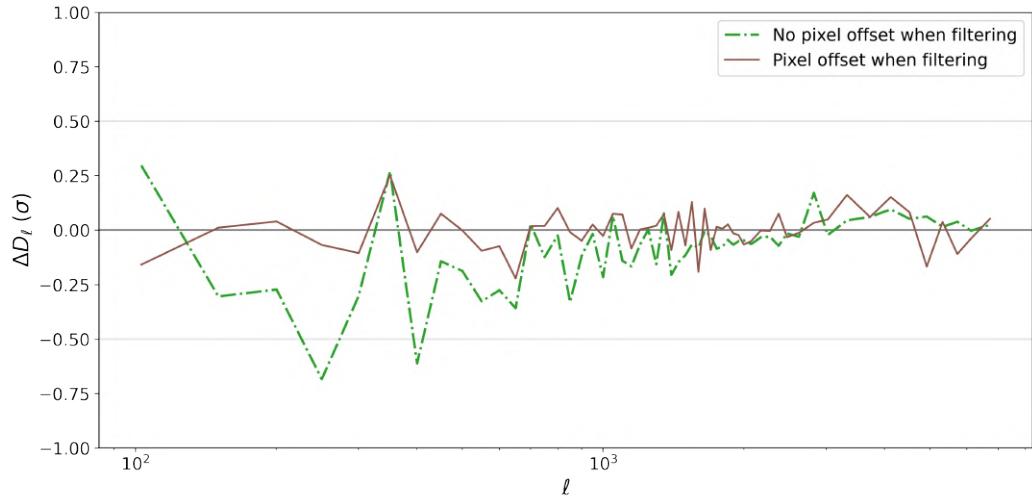


Figure 4.16: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when we add (or not) a one-pixel offset to the Fourier-space filter, divided by the uncertainties on the DR4 spectrum.

Footprint Mask Before Filtering

In the DR4 analysis, we multiply each map by a footprint mask for the region of interest before Fourier-space filtering. This masks the messy edges of the survey area. The effect of this is visible in Figure 4.17.

¹⁷There is a convenient flag in some of the map preprocessing functions in NaWrapperr in order to reproduce this effect, named “legacy_steve” after Steve Choi, who led the power spectrum analysis for DR4.

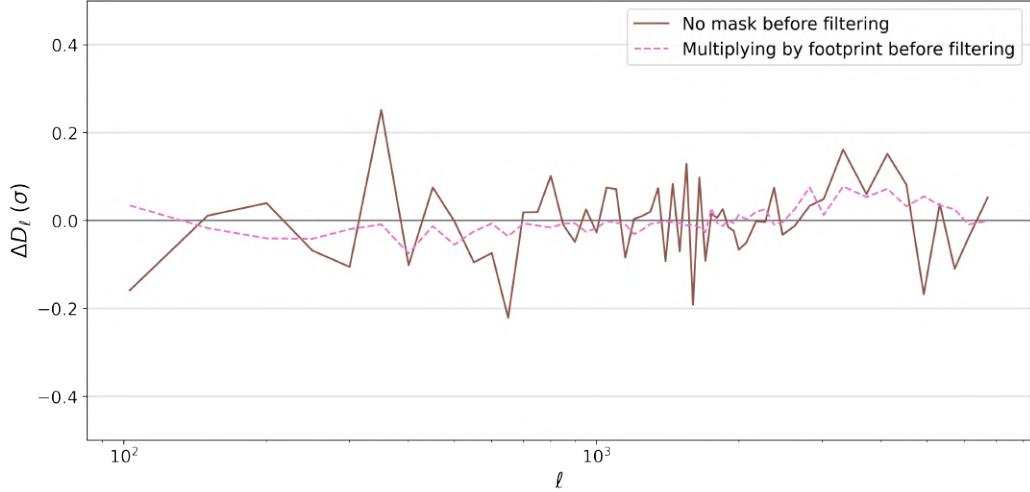


Figure 4.17: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when we multiply (or not) the maps by a footprint mask before filtering them, divided by the uncertainties on the DR4 spectrum.

Adding Weak Point Sources Back to Map Before Filtering

In notebook 8, we begin with source-free maps, filter them, then add the sources. We do this because we don't want any sharp features (in the form of bright point sources) in the maps when filtering them, as this may lead to undesirable artifacts. However, the source maps we make contain many more sources than just the bright point sources that we mask when computing the power spectra. We should be adding the weak (non-masked) sources to the maps prior to filtering them, and this is indeed what was done for the DR4 analysis. See Figure 4.18 for the change in the spectra due to adding the weak sources prior to filtering.

Details of Adding Weak Point Sources

The result above of adding the sources to the maps before filtering was not the much better agreement with the DR4 spectrum we may have been hoping for. So we study the maps in greater detail. Although they are not publicly available, we have access to the maps used for the DR4 analysis, after the weak point sources were added in, but before the Fourier-space filtering was performed. An example of such a map is

shown in Figure 4.19.

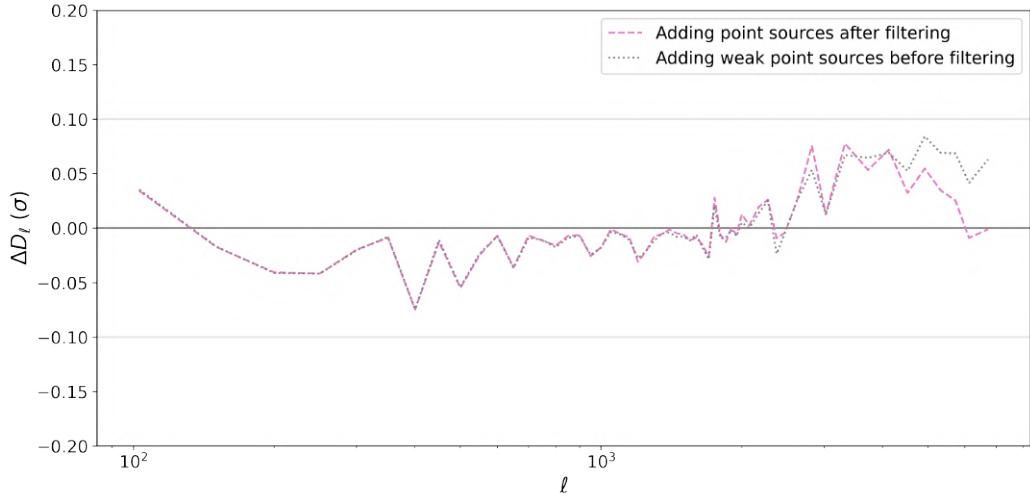


Figure 4.18: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when the weak (non-masked) point sources are added to the maps either before or after the Fourier-space filtering is performed, divided by the uncertainties on the DR4 spectrum.

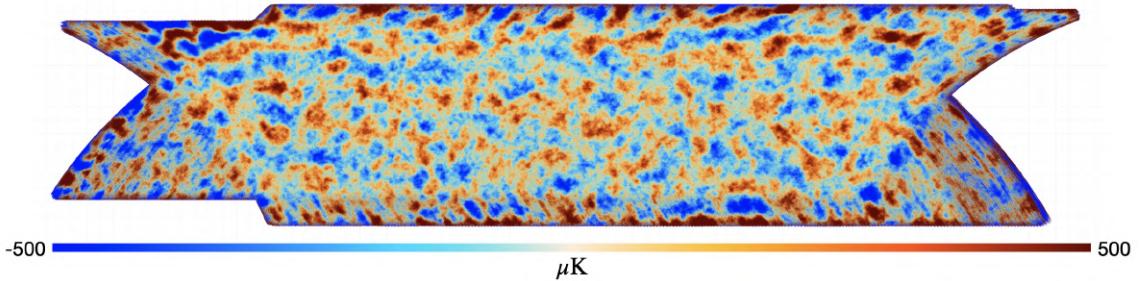


Figure 4.19: Map used for the DR4 power spectrum analysis, after the weak point sources were added to the map but before the Fourier-space filtering was performed. This map has been calibrated with the final calibration factor, so it is in units of μK_{CMB} .

Using these maps with the weak sources already added, instead of adding those sources to the maps ourselves, results in a spectrum that is in better agreement with the DR4 spectrum, as shown in Figure 4.20. Note that, as with the pre-filtered maps used previously, we needed to account for the different calibration between the final DR4 maps and the maps with the added sources, which are from an earlier stage of the analysis.

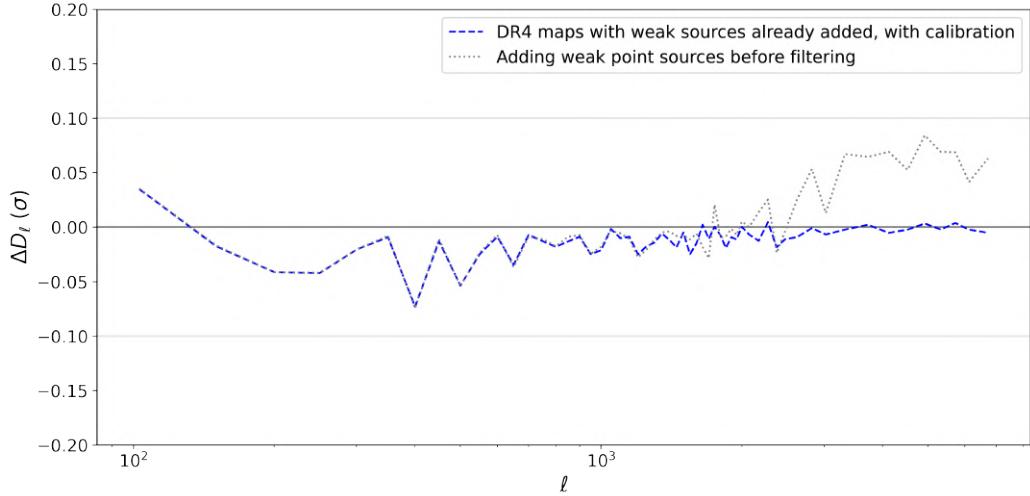


Figure 4.20: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when either adding the weak sources to the maps ourselves, or using the DR4 maps with the weak sources already added, divided by the uncertainties on the DR4 spectrum.

Where does the difference in Figure 4.20 originate from? To get to the bottom of this, we look at the difference in the maps, an example of which is shown in Figure 4.21. The difference between the maps seems to be zero in most areas, except for what appear to be the locations of point sources.

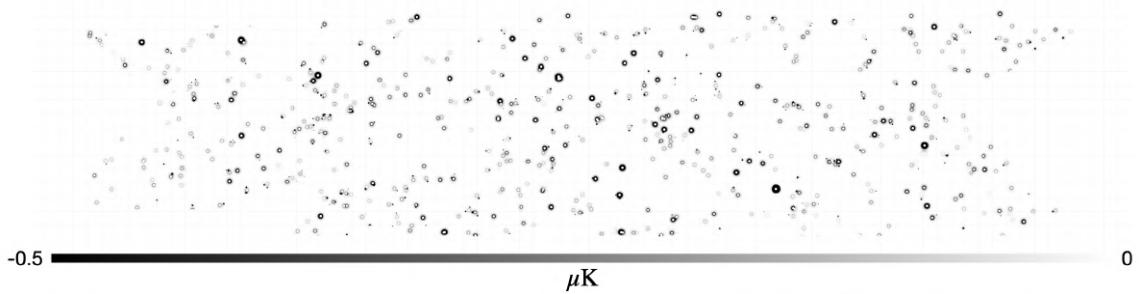


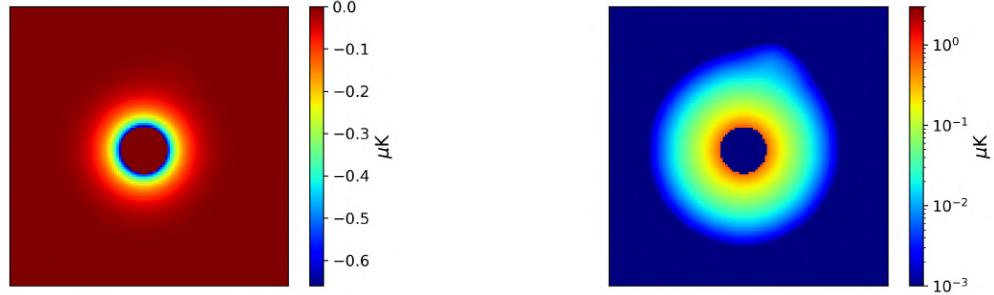
Figure 4.21: Difference between a (renormalized) map from the DR4 analysis with the weak point sources already added and a map to which we added the weak point sources ourselves.

So, digging further, we cut out a small region in the maps around a point source. The difference between the maps in this region is shown in Figure 4.22a with a linear color scale, and in Figure 4.22b with a logarithmic color scale, to more clearly see different features. Figure 4.22e shows the map of the sources for the cutout, and Figure 4.22f

shows the (non-apodized) point source mask for that area. The point source near the center of the map cutout is one of the bright sources that are masked for the power spectrum analysis, and the point source near the top edge is one of the weaker sources that are added before filtering and not masked. The source of discrepancy seems to be centered on the area around the bright, masked source.

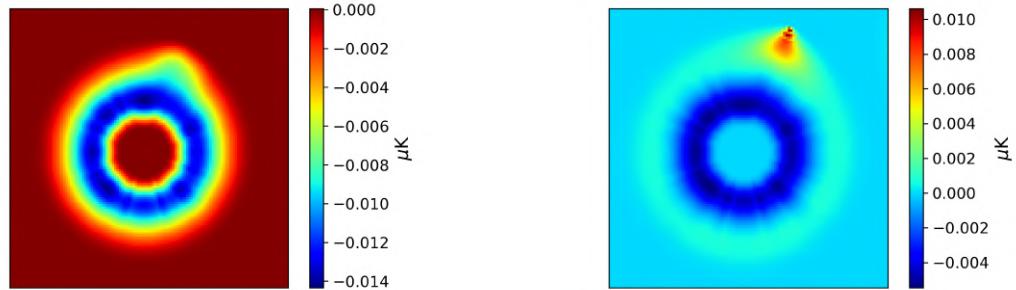
In order to make a map with weak point sources, we multiplied a source map by the non-apodized point source mask, to create a map with only the weak (non-masked) sources, then added this to the corresponding source-free map. We can try using an apodized point source mask for this instead. We try a C^2 apodization, as defined earlier, with an $18'$ radius in Figure 4.22c, and with a $24'$ radius in Figure 4.22d. The difference between the DR4 maps and the maps to which we added sources using an apodized mask is smaller in amplitude than for the maps to which the sources were added with a non-apodized point source mask. However, it also seems like the greater the apodization radius, the greater the asymmetry in the difference map around the masked source. Even in Figure 4.22b, there is structure in the direction of the weak source near the top of the map cutout. This effect is more pronounced in Figure 4.22c, and more pronounced still in Figure 4.22d. This makes sense, since the mask apodization we applied around the bright source affects the signal from the weak source, which we should not be altering.

Overall, it appears like the difference between the spectra in Figure 4.20 is due to some subtle difference, that we have not identified, in how the weak sources are added to the maps prior to filtering. The effect on the spectra is small (the difference between the two spectra in Figure 4.20 is $< 0.1\sigma$ based on the DR4 uncertainties, or $\lesssim 1.5\%$ of the amplitude of the DR4 spectra), so this detail is not a concern for analyses going forward.



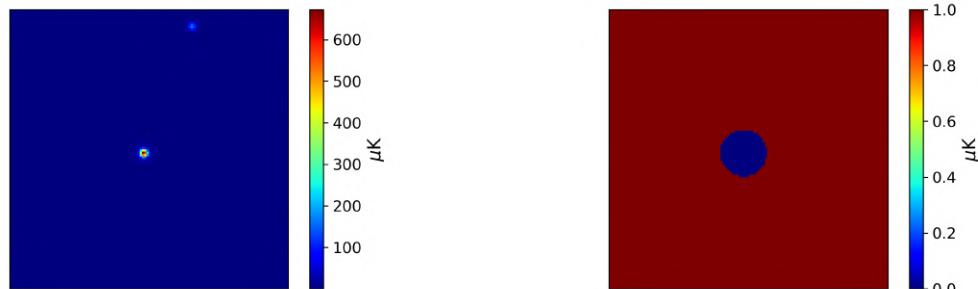
(a) Difference between a (renormalized) map from the DR4 analysis with the weak point sources already added and a map to which we added the weak point sources ourselves using a non-apodized point source mask.

(b) As in (a) but using a logarithmic color scale.



(c) As in (a) but with an $18' C^2$ apodization.

(d) As in (a) but with a $24' C^2$ apodization.



(e) Map of the sources.

(f) Non-apodized point source mask.

Figure 4.22: A $1^\circ \times 1^\circ$ map cutout, centered on a bright source near RA = 32° , Dec = -5° . Note the weaker source visible near the top of the map in panel 4.22e.

Zero-Padding

In order to speed up computations, in the DR4 pipeline, maps are zero-padded before computing power spectra. These computations involve several FFTs. FFTs of lengths that are a power of 2 are speediest, but zero-padding our maps to have dimensions which are powers of 2 in each dimension may lead to the maps being too large, either causing memory issues or suboptimal computation time. Since FFTs of lengths with low prime factors (3, 5) are only slightly longer to run than powers of 2, we zero-pad the ACT maps by changing the length of each dimension to the nearest (higher) product of positive integer powers of (2, 3, 5).

This zero-padding is not performed in notebook 8. Since zero-padding is done merely to reduce the computation time, it should not have a considerable effect on the resulting power spectrum. Indeed, the effect is not significant, but adding the zero-padding to the power spectrum analysis does make a detectable difference, as shown in Figure 4.23. This difference must arise when k-space filtering, since the spectrum obtained when using pre-filtered maps agrees with the zero-padding option here, not the “default map size” option. As with the spectrum obtained with a pre-filtered map earlier, the spectrum with zero-padding here agrees within 0.002σ with the DR4 spectrum.

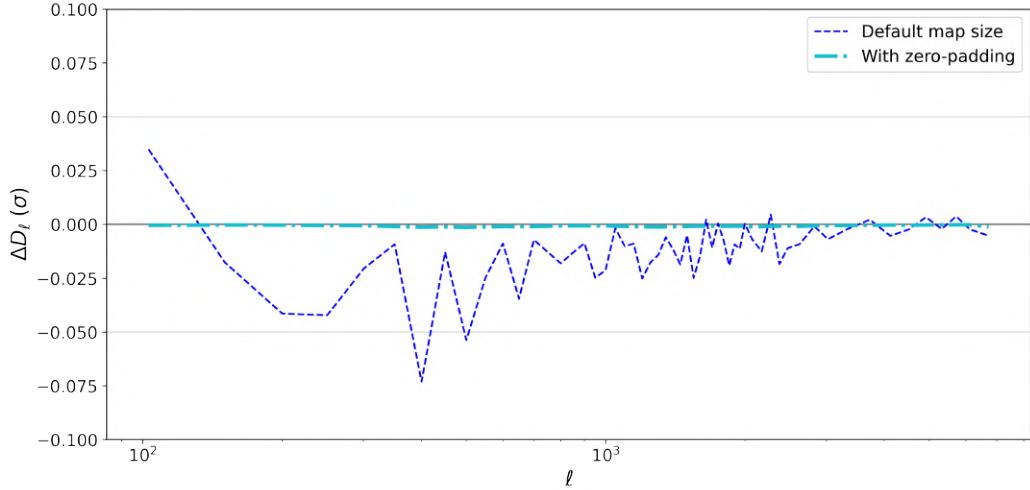


Figure 4.23: The difference between the DR4 power spectrum and the spectrum obtained in notebook 8 when either zero-padding the map or not, divided by the uncertainties on the DR4 spectrum.

In the end, we tracked down the sources of the differences between the DR4 spectra and those computed in notebook 8. This process has highlighted how different choices in the preprocessing steps prior to the power spectrum calculation can affect our results.

4.4.2 **PSPipe**

For DR6, we have transitioned to using `PSPipe`¹⁸, a public, general purpose pipeline for the computation of CMB power spectra. It was built for the analysis of the high resolution maps from the large aperture telescope of the Simons Observatory ([Simons Observatory Collaboration, 2019](#)). The `PSPipe` pipeline contains two options for codes to use for power spectrum estimation: `NaMaster` and `pspy`¹⁹ ([Louis et al., 2020](#)). Both codes have been tested against each other, and they agree to numerical precision ([Li et al., 2021](#)).

Following the extensive study of the effect of various analysis choices in Section 4.4.1,

¹⁸GitHub repository: <https://github.com/simonobs/PSPipe>

¹⁹GitHub repository: <https://github.com/simonobs/pspy>

we test power spectrum pipelines built using `NaWrapper` and `PSpipe`. While the power spectrum computations with `NaMaster` and `pspy` were already tested against each other, this is the first test of full pipelines, including all the preprocessing steps such as building window functions, Fourier-space filtering the maps, correcting for the pixel window function, adding the sources to the source-free maps, etc. With consistent inputs/settings, we found the results of the `NaWrapper` and `PSpipe` pipelines to agree within 0.001σ . Given the many intricacies of the preprocessing steps explored in the previous section, this is a significant achievement.

Chapter 5

Daytime

5.1 Introduction

5.1.1 The Data

The published ACT data analyses have focused solely on data collected during the nighttime, which we define as being from 23:00 to 11:00 UTC.¹ However, almost 50% of the TODs available for analysis were collected during the daytime (11:00 to 23:00 UTC). This chapter concerns the analysis of some of these data.

As of December 2022, the ACT team is working on the next data release, DR6. The expected baseline plan is for DR6 to include the nighttime data collected by ACT from 2017 to mid-2021, by PA4 in the 150 and 220 GHz frequency channels, and PA5 and PA6 in the 98 and 150 GHz frequency channels. The parallel daytime analysis is focused on the corresponding daytime data, that is, data collected by PA4, PA5, and PA6 from 2017 to mid-2021.

The DR6 nighttime observations spanned the same region of sky (a large patch cov-

¹Recall that the local time at the ACT site fluctuates between UTC-3 and UTC-4.

ering nearly half the sky, AA in Figure 1.12) for the entire 4.5 years. Meanwhile, the daytime scanning strategy was changed halfway through this time period. From 2017 to May 2019, the daytime observations were focused in two smaller, deep patches, which are mapped together in what we refer to as the “day deep” maps. From May 2019 to mid-2021, the daytime observations switched to the same scanning strategy as the night. The resulting maps are referred to as the “day wide” maps. Example inverse-variance maps, displaying the survey coverage and relative map depths, for night, day deep, and day wide are shown in Figure 5.1.

In past ACT analyses, each year of data was mapped separately, because the scanning strategy and the detector arrays being used changed annually. From 2017 through 2021, however, the nighttime scanning strategy stayed constant, the daytime scanning strategy changed only once (changing from deep to wide patches in mid-2019), and only one detector array was changed out (PA6 was removed in early 2020 and replaced with PA7). Therefore, all the years of nighttime observations are mapped together (with an 8-way split), the pre-mid-2019 deep daytime observations are mapped together (with an 8-way split), and the post-mid-2019 wide daytime observations are mapped together (with a 4-way split, since these maps have less depth). The day wide patch has the same sky coverage as the nighttime observations.

The detector arrays used during each year of daytime data considered here are shown in Figure 5.2. Since PA6 was removed shortly after beginning observation of the wide patches, the PA6 wide maps are quite shallow.² The left half of the southern day deep patch has shallower coverage because a larger percentage of observations there have been cut. The initial observation scans cover the patch in a roughly uniform way, but then data are cut due to poor beam quality, as described in Section 5.3.2. The beam is noticeably worse in the southern patch for $\text{RA} \gtrsim 60^\circ$.³

²This is discussed more in Section 5.4.1.

³This is before the beam cuts from Section 5.3.2.

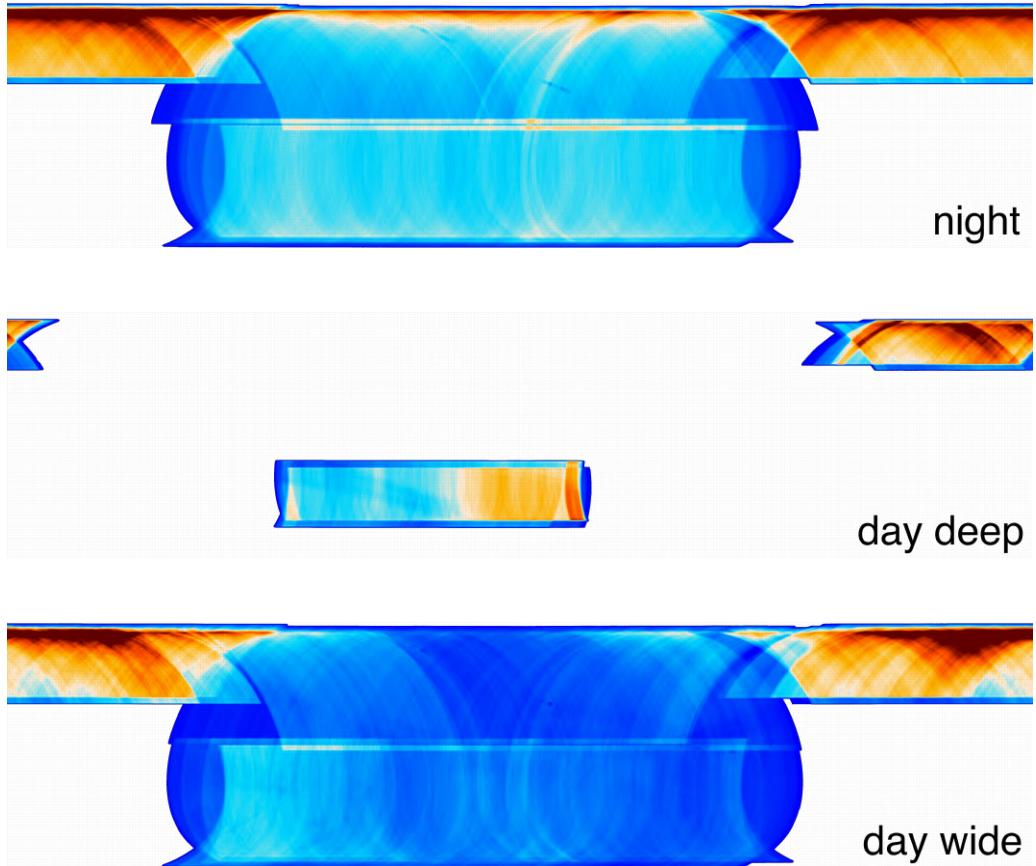


Figure 5.1: Inverse-variance maps for PA4 at 150 GHz (coadded across all map splits) for night (top), day deep (middle), and day wide (bottom). The maps span from RA = 180° (left) to RA = -180° and from Dec = -63° to Dec = 23° (top). The color scales range roughly from 5×10^{-7} to 5×10^{-4} (μK) $^{-2}$ for night and day deep, and from 5×10^{-7} to 2.5×10^{-4} (μK) $^{-2}$ for day wide. (These maps have not been corrected for the final power-spectrum-based calibration described in Section 4.2.7. Typically, the preliminary calibration at this stage is correct to within $\sim 25\%$.)

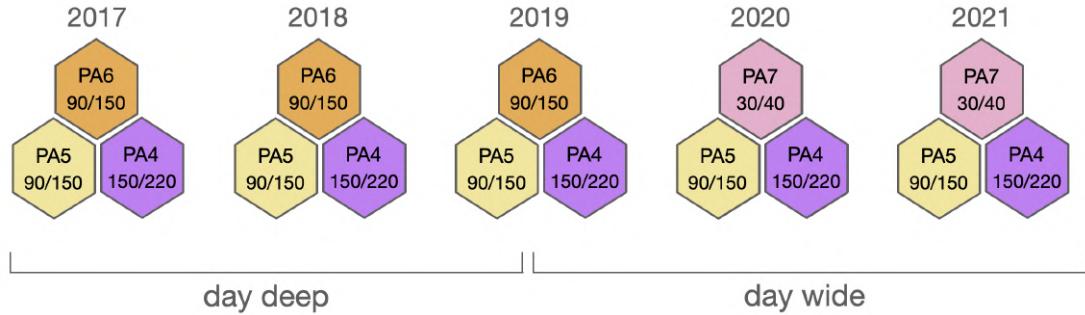
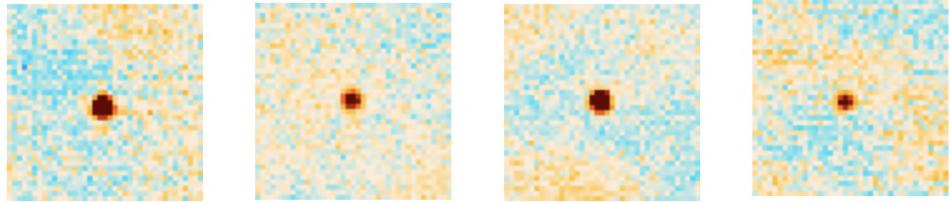


Figure 5.2: Detector arrays used during each year of observation from 2017–2021. (This is the same as the bottom half of Figure 1.11 in Section 1.2.1, but is repeated here, since that was several chapters back.) The brackets indicate the regions of sky observed during a given year.

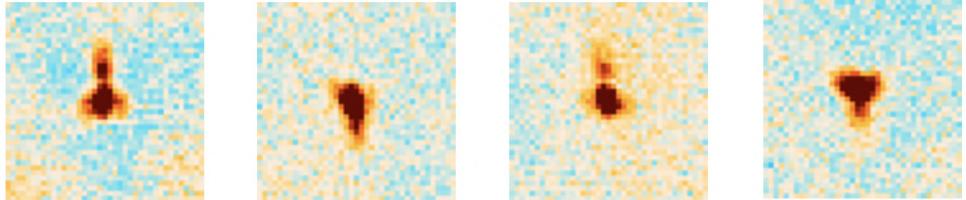
5.1.2 The Main Challenges

The daytime data are challenging to work with because the Sun’s heat causes time-dependent distortions to the telescope mirror, which leads to pointing offsets and beam deformations that change on the timescale of hours. The extent of these effects is dependent on the time of day, the temperature of the reflector, the pointing of the telescope, and possibly other factors. Attempts have been made at modelling and predicting the beam’s daytime behavior, but none has been sufficiently successful so far. They have not yet been used in cosmological analyses.

While the beam during the day looks good (azimuthally symmetric) most of the time, from time to time it can behave very badly, and the “badness” is much more than simple ellipticity, as shown in Figure 5.3 for regions around extragalactic point sources.



(a) Examples of when the beam appears to be compact and approximately radially symmetric, as expected.



(b) Examples of when the beam demonstrates pathological behavior during the day.

Figure 5.3: Examples of point sources seen in maps made from individual daytime TODs.

The approach to the daytime beams is two-fold. First, when making maps, the data with the worst, most pathological beam behaviors are excluded, as summarized in Section 5.3.2. Then, as summarized in Section 5.5, we apply a correction to account for the difference between the effective nighttime and daytime beams.

The beam cuts, combined with the pointing correction described in Section 2.1.5, lead to maps of daytime data which are indiscernible by eye to maps of nighttime data.

5.2 Past Work

The in-depth work on the daytime data began with looking at the daytime equivalent of part of DR4, using observations from 2014 to 2016. During this time, the daytime observations were made in two deep patches, one in the north and one in the south. These were very roughly in the same areas as the deep daytime patches observed from 2017 to mid-2019, but the exact footprints changed each year.

We made maps of these 2014–2016 daytime data. In doing so, we included the pointing correction described in Section 2.1.5 and 5.3.1, but did not make any beam cuts as described in Section 5.3.2. The 2014–2016 daytime maps were therefore not as high quality as the 2017–2021 maps we are working with presently.

However, the worst daytime beams were observed to occur predominantly in the south, so our efforts were focused on the northern regions, shown in Figure 5.4 along with the BN region from DR4. The BN nighttime maps were used during this analysis as well.

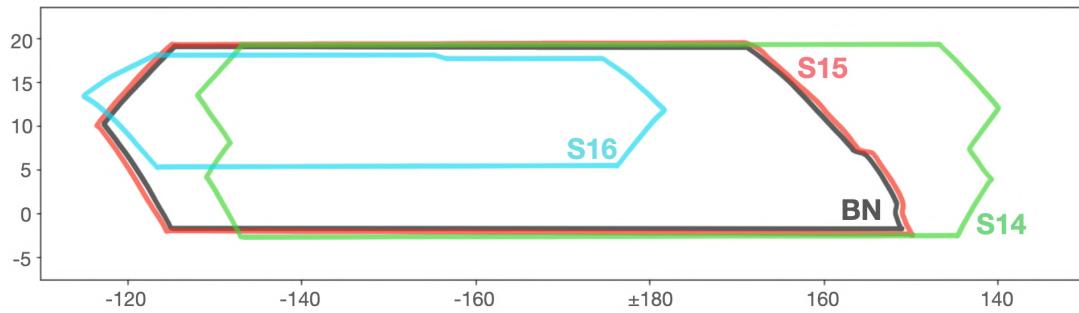


Figure 5.4: Footprints for the northern daytime observations from 2014–2016, along with the BN footprint from DR4 (see Figure 1.12) for comparison. The S15 daytime patch overlaps almost perfectly with BN. The x -axis (y -axis) shows the RA (Dec) coordinates in degrees.

We computed the 2014–2016 daytime spectra for these northern patches, assuming the same beams as for the nighttime maps of BN. Then we computed a correction factor to apply to each nighttime beam to obtain an estimate of the daytime beam. We did this using a power-spectrum-based approach, described in Section 5.5.1. Figure 5.22 in Section 5.5.1 illustrates the effect of correcting the spectra to account for these daytime beams.

During this preliminary work, the spectra were considered one by one — they were not coadded across detector arrays or seasons.⁴ This does mean we are limited by

⁴Computing all the relevant cross-season-spectra and coadding them is further complicated by all the different footprints and their different overlaps.

the signal-to-noise of the data in the northern region for one season, detector array, and frequency.

Nonetheless, we performed a significant number of tests. We tested for isotropy, by splitting each map into a number (2–8) of smaller maps, then analyzing each of the small maps separately, in order to check that the resulting power spectra and daytime beam corrections are consistent for different regions. No evidence of significant spatial variations were found. Some preliminary PWV and elevation map-based null tests were performed, as well as tests where we split the observations either by azimuth (high vs low) or time of day (beginning and end vs middle of the day). The latter tests were added because there is evidence that the daytime beams depend, in part, on the azimuth and time of day of the observations (see Section 5.3.2). These initial null tests did not reveal any issues, but we chose to terminate further efforts on the 2014–2016 data, in favor of moving on to the much larger 2017–2021 dataset.

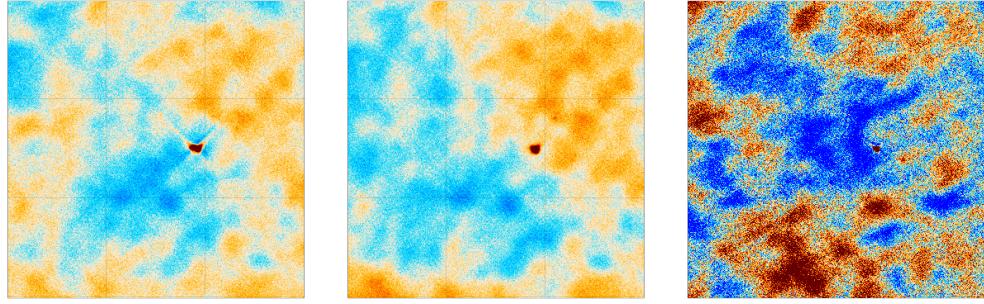
5.3 Map-Making

5.3.1 Pointing

In the map-making for the 2017–2021 ACT data, the pointing offsets are corrected using the same method applied to the nighttime data for DR4, described in Section 2.1.5. The newest pointing correction does not make a noticeable difference for the nighttime data, but significantly improves maps of the daytime data, as can be seen in Figure 5.5.

5.3.2 Beam Cuts

The beam deformations are much harder to correct for than the pointing offsets, so the approach to the beams is two-fold; data are cut where the beams are the worst



(a) Daytime map without pointing correction (b) Daytime map with pointing correction (c) Nighttime map with pointing correction

Figure 5.5: Point source seen in the 2018 PA4 150 GHz maps. The maps here are $3^\circ \times 3^\circ$. The nighttime map is noisier than the daytime ones, as it contains significantly fewer data in that region. These maps have not been filtered, so the large scale modes represent contributions from the atmosphere.

behaved, and we apply a correction to the beam profiles for the data that remain.

As part of the pointing correction⁵ described in Section 2.1.5, small thumbnail maps are produced for each moderately strong point source hit during the scans. This results in several million thumbnails, which can be stacked in order to measure the beam shape with adequate signal-to-noise. Since the telescope pointing has been observed to vary on roughly one-hour timescales, the point source thumbnails are stacked into bins of roughly one hour. A new bin is also started whenever the scanning pattern changed. This results in several thousand stacks.

For each detector array and frequency, the pixel-pixel covariance matrix of the thumbnails is computed, and then the strongest eigenmodes are used as a basis to characterize the beams. The majority of the beam-like behavior seems to be contained to the top 5 eigenmodes, with the remaining eigenmodes seeming mostly like noise, so the top 5 eigenmodes for each detector array and frequency are used as templates to fit each of the thumbnail stacks. Based on these fits, a beam “badness” statistic is defined as the fraction of the power in the top 5 eigenmodes that is not contained

⁵Both the work on the pointing correction and this work on the beam cuts were led by Sigurd Naess.

within the first eigenmode. This number is close to zero during the night, as expected, but it can reach above 10% at times during the day.

The telescope scans are carried out at four mean azimuths, of roughly -125° , -50° , $+50^\circ$, and $+125^\circ$. When the stacks are split by azimuth and then organized by time, there seem to be clearly defined regions of beam badness. There are several times when the beam is only noticeably bad for a subset of the azimuths, for instance.

The fact that there are well-defined regions of beam badness signals that it may be possible to simply cut out the daytime data with the worst beam badness. A mask is defined based on pixels with a badness of $> 4\%$, which results in cutting roughly one third of the daytime data.⁶

The scanning regions⁷ for the daytime are shown in Figure 5.6. We can study the effect of the beam cuts separately for each of these regions, since there is some evidence to suggest this might be informative.



Figure 5.6: Scanning regions for the daytime maps. These outlines represent how the data are collected —the telescope scans across these regions one at a time.

As mentioned above, it is sometimes the case that the beam for a certain period of time is only bad for a subset of the azimuths. Moreover, the scans at azimuths of

⁶The beam badness is generally lower in the 90 GHz band, and the measurements at 220 GHz are the noisiest, so we use the 150 GHz data for determining the mask.

⁷These are the regions targeted with the telescope scans, they are slightly different from the region footprints used for mapping.

roughly -50° and $+50^\circ$ are only carried out in the regions ending in “_n”, whereas the scans at azimuths of roughly -125° and $+125^\circ$ are only carried out in the regions ending in “_s”.

The effect of the beam cuts is shown in Figure 5.7, where we examine the number of TODs contributing to the maps for each daytime scanning region, both before and after the beam cuts. We see that not only are the beam badness cuts specific to certain azimuths, they also depend on the region of sky being observed. For example, considering the scans at an azimuth of $\sim 50^\circ$, several of the TODs from 12 to 16 UTC are cut from the day_wide_12h_n region, but almost no TODs from that time period are cut from the day_wide_01h_n region. We do not yet have a good understanding of why the cuts have these dependencies.

The scans are also carried out at a few discrete altitudes, but there is no appreciable difference in the beam statistics (for example in an equivalent of Figure 5.7) when splitting the scans this way.

We can further divide up the TODs by time of year, as shown in Figure 5.8. This clearly shows there are specific times and dates where the beam is bad for each scanning region. There are some interesting features to these regions of beam badness. For example, in the day_deep_12h_n region, for scan at an azimuth of roughly $+50^\circ$, most of the data from $\sim 14:00$ to $19:00$ UTC from about September to December 2018 are cut due to bad beams, but the data from the corresponding time in 2017 are not.

As mentioned in Section 5.2, in earlier work it was noticed that the daytime beams often exhibited the most pathological behavior in the southern regions of the maps, as seen by the shapes of point sources in the maps. However, this is not reflected by the beam badness statistics here. There are several possible explanations for this. For example, the percentages examined here merely report the cases of beam badness $> 4\%$. Perhaps the worst offenders, where the beam badness is $\gtrsim 10\text{--}20\%$, occur

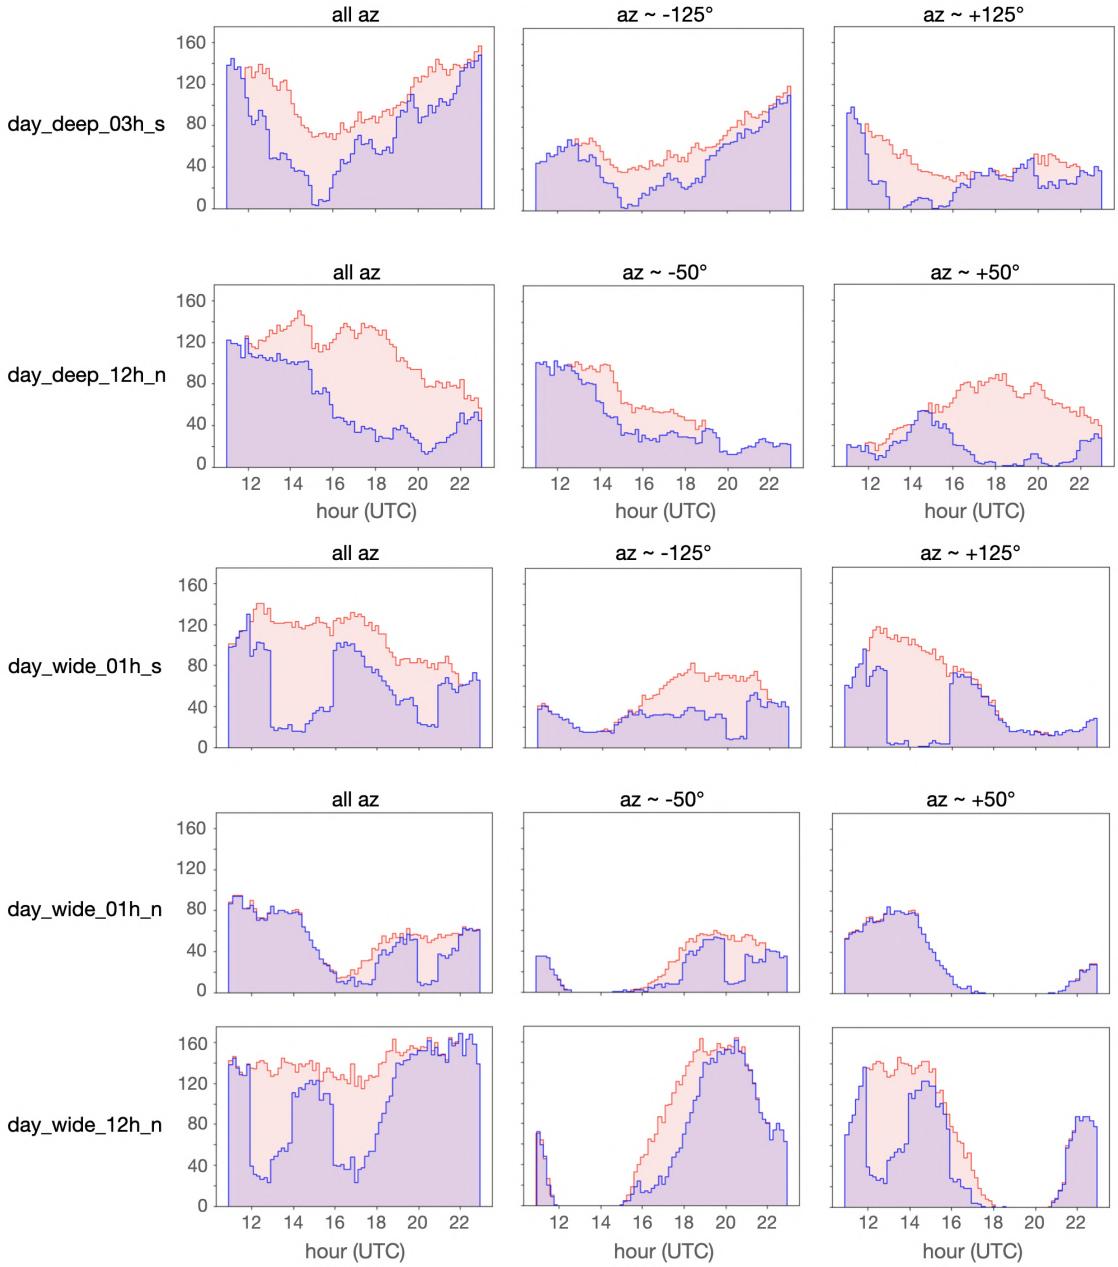


Figure 5.7: Histograms illustrating the number of TODs per hour in the PA5 150 GHz daytime maps, both with (blue) and without (red) the beam cuts. Each row corresponds to a different scanning region (these regions are shown in Figure 5.6) and each column corresponds to a different azimuth.

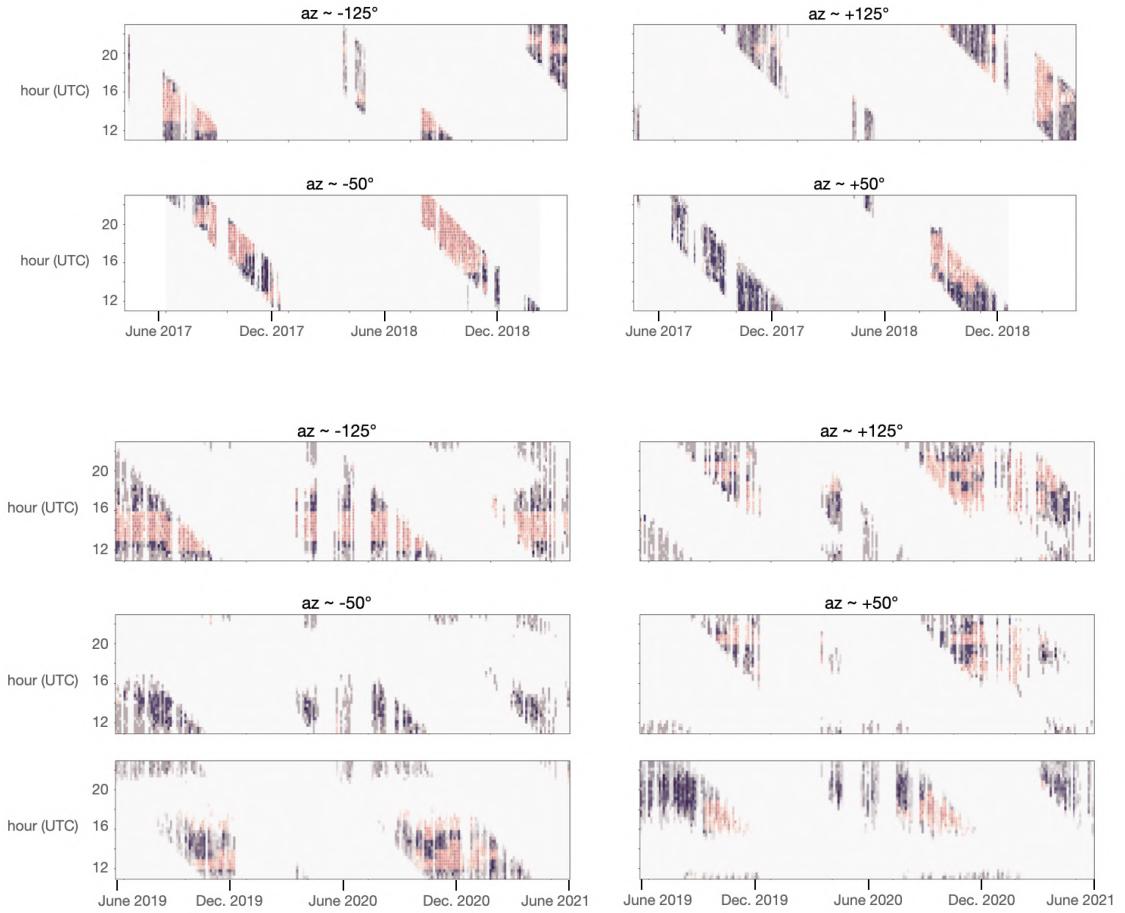


Figure 5.8: Histograms illustrating the number of TODs per hour, depending on the time of year, in the PA5 150 GHz daytime maps, both with (blue) and without (red) the beam cuts. Each column corresponds to a different azimuth and each row corresponds to a different scanning region, from top to bottom: day_deep_03h_s, day_deep_12h_n, day_wide_01h_s, day_wide_01h_n, day_wide_12h_n.

predominantly in the south, but there are more mild offenders (with a beam badness $\lesssim 10\%$) in the north.

5.4 Power Spectrum Analysis

5.4.1 Spatial Windows Functions

The procedure used to obtain the spatial window functions for DR4 was described in Section 4.2.3. To recap, in short, for DR4, for each region, the spatial window is the product of an apodized footprint, a point source mask, a Galactic mask, and an inverse variance map (for weighting). The apodized footprint is obtained by multiplying two masks, obtained by applying thresholds⁸ to the cross-linking and inverse variance maps, then skipping 1° from the edge of the resulting region, and further apodizing the edges by 1° .

For DR6, the default procedure at the moment (still in development) is similar. The spatial window is the product of an apodized footprint, a point source mask, and a Galactic mask. The default choice is to use uniform spatial weighting instead of weighting by the inverse variance map (see Section 5.4.3 for further discussion of weighting). The footprint is obtained as for DR4, masking the pixels with an inverse variance of 0 (which corresponds to no data) and/or a cross-linking value > 0.97 . However, contrary to DR4, we do not define a boundary common to all arrays. Instead, we obtain a footprint for each detector array/frequency separately. There is a pointing offset between the different arrays, so allowing each to have a different footprint avoids unnecessarily discarding data from the edges of the maps. An example of a resulting spatial window function is shown in Figure 5.9. The cross-linking threshold results in masking a strip of the maps, from a declination of roughly -43° to -26° .

⁸These thresholds are patch-dependent, and are listed in Table 2 of Choi et al. (2020).

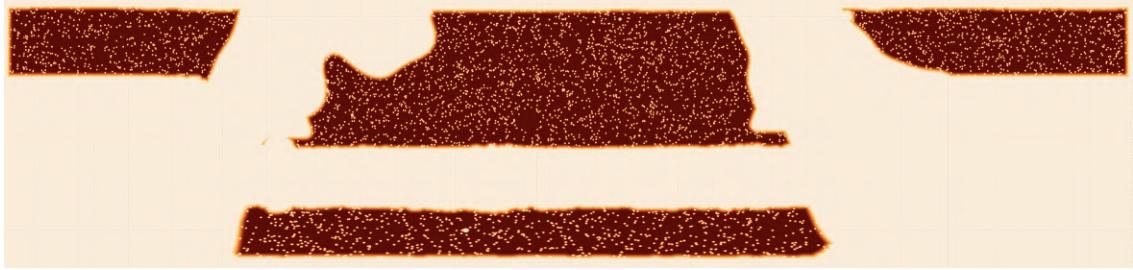


Figure 5.9: Spatial window for PA5 at 150 GHz, created for the DR6 (nighttime) AA patch.

As described in Section 1.2.1, ACT observed by scanning back and forth horizontally as the sky was either rising or setting. This allowed a given patch of sky to be observed at different parallactic angles, which leads to the cross-linking in the maps. Unfortunately, for some declinations, the sky rotations does not allow for orthogonal scans.⁹ The cross-linking in the maps is worst around $\text{Dec} \approx -35^\circ$.

For DR4, region D8 was omitted from the cosmological analysis “due to its poor cross-linking.” Issues were encountered during the analysis of these data (Choi et al., 2020), and attributed them at least in part to the lack of cross-linking in this region, a reasonable assumption, given how cross-linking is important for minimizing certain systematic effects.

However, since then, preliminary isotropy tests have found that the spectra from the poorly-crosslinked strip of the maps from $\text{Dec} = -43^\circ$ to -26° is consistent with the spectra from the rest of the maps. Analysis is still underway to fully understand the current DR6 nighttime spectra. Once this is done, we may try re-computing spectra without masking out this strip.

We find that not masking the data between $\text{Dec} = -43^\circ$ and -26° results in a $\sim 10\%$ decrease in the power spectrum uncertainties for each detector array/frequency for the nighttime and day wide maps, and a $\sim 50\%$ decrease in the power spectrum

⁹Since we observe at a few different elevations, there is no declination where we have only perfectly parallel scans, so absolutely no cross-linking (which would be a cross-linking value of 1). However, we do get uncomfortably close to this, with cross-linking values around 0.99 or higher.

uncertainties for the day deep maps (since the southern deep patch is almost entirely between these declinations).

The result of applying the current DR6 spatial window function algorithm to the day wide maps for PA4, PA5, and PA6 at 150 GHz is shown in Figure 5.10. These look noticeably different from the nighttime example. The main reason for this is that the coverage of the daytime maps is not as uniform as the nighttime ones. The beam cuts from Section 5.3.2 sometimes result in localized groups of pixels having fewer (or no) data in one (or both) scanning directions. When we mask pixels with poor cross-linking or inverse variance values, we therefore mask a number of standalone pixels throughout the maps. We then skip 1° from the edges of the mask, which means 1° away from any pixel where the mask value is set to 0. This leads to a hole with a 1° radius around what may be as little as one pixel with poor cross-linking. This is clearly undesirable. In regions, these holes merge together and mask out a significant area of the maps.

We explore different ways of building the daytime spatial windows, specifically the footprints. To keep things simple, when doing this we look at the binary footprints used before Fourier-space filtering, as shown in Figure 5.11. While only mentioned briefly in Section 4.4.1, these binary masks are similar to the footprints from Section 4.2.3. The footprints from Section 4.2.3 are used to multiply the maps when computing the power spectra, to mask the messy data near the edges of the survey region. However, prior to Fourier-space filtering the maps, we multiply them by a different, binary footprint. These binary footprints do not have apodized edges (hence the name binary — all pixels have a value of either 0 or 1), and we skip only 0.5° from their edges instead of 1° as for the final footprints. Skipping the edges of the area where the Fourier filtering is applied when computing the spectra should prevent some aliasing from the filter from entering the data.

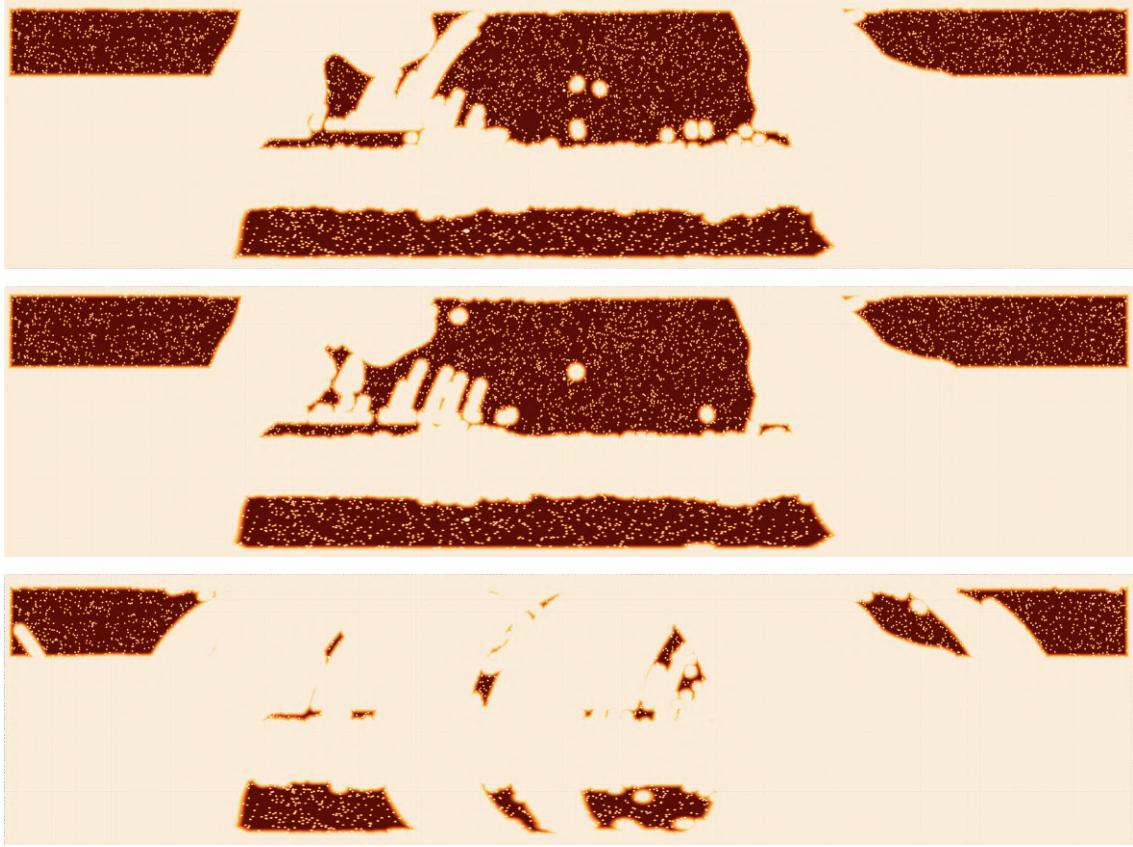


Figure 5.10: Spatial windows for PA4 (top), PA5 (middle), and PA6 (bottom) at 150 GHz, created for the day wide region, using the same settings as for the current nighttime analysis.

Since we are not concerned about a handful of pixels with poor cross-linking, we can replace the cross-linking cut by simply masking out the strip of sky from $\text{Dec} = -43^\circ$ to -26° . This removes the majority of the holes in the daytime windows.¹⁰

Since it seems that it may not be necessary to mask out the poorly cross-linked data, we create an option where it is removed altogether. After this, there are still occasionally some holes in the daytime windows. These are due to pixels for which $\text{ivar} = 0$, that is, there is no data. A zoom-in on a region with a few such pixels is shown in Figure 5.12.

¹⁰While it is not as much of an issue for the nighttime maps, there are occasionally holes in the nighttime spatial windows. Not for PA5 at 150 GHz, shown in Figure 5.9, but for PA4 at 220 GHz and PA6 at 98 and 150 GHz. These holes disappear entirely when we replace the cross-linking threshold by a declination cut.



(a) With the default settings used for the current nighttime analysis.



(b) Replacing the cross-linking cut by a declination cut.



(c) No cross-linking or declination cut.



(d) No cross-linking or declination cut, and filling the hole(s) in the mask.

Figure 5.11: Selection of binary footprints (which are applied to the maps prior to Fourier-space filtering) for the PA4 150 GHz day wide maps, obtained using different techniques.

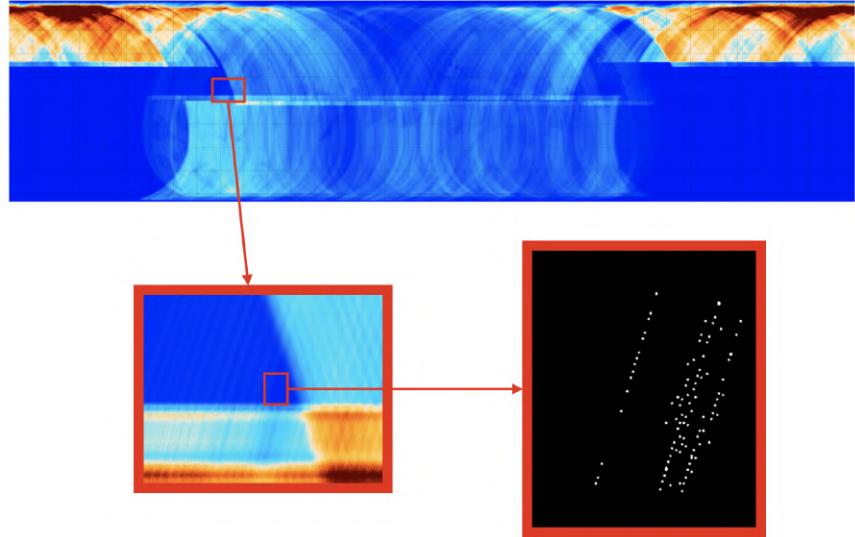


Figure 5.12: A zoom-in on the region of the inverse variance map corresponding to the hole in the binary mask in Figure 5.11c. Although the hole in the final spatial window is $\sim 2^\circ$ in diameter ($\sim 1^\circ$ in diameter in the binary mask we apply before Fourier-space filtering), it is caused by only a handful of pixels missing data in the map. In the full map (and the first level of magnification), the color scale ranges from roughly 0 to $5 \times 10^{-5} (\mu\text{K})^{-2}$. In the final zoom-in, to maximize contrast, the pixels with a value of zero are white, and all non-zero pixels are black.

A few pixels in the middle of the map with missing data are no problem.¹¹ In fact, they already happen in the standard nighttime analysis, but only in polarization, not temperature, hence they do not attract as much attention (since the masks are made using the temperature component of the inverse variance).

During the map-making process, each pixel in a map is hit (ideally) multiple times, by scans going in different directions. Each of these hits corresponds to a certain amplitude measured by the polarization-sensitive detector arrays. These measured amplitudes are used, along with the detector polarization angles and telescope pointing, to solve for values of I, Q, U for the pixel. If there is not enough information to solve for all three components, the map-maker simply attributes all the signal to I , and sets Q and U to zero. While perhaps not ideal, this happens very rarely in the maps, so it does not affect the spectra in any significant manner.

¹¹The value for $[I, Q, U]$ for that pixel is simply set to $[0, 0, 0]$.

We want to avoid masking a region with a 1° radius due to a few pixels with no data. As a reminder, our maps are made with $0.5'$ pixels, so a circle with a 1° radius contains $> 45,000$ pixels. We don't want to discard the information in all of these pixels due to a few bad apples (pixels).

So we fill in these regions of the windows and do not mask the pixels with no data when they occur in the middle of a patch rather than the edges. While they have a negligible effect on the power spectra, we can easily inpaint the handful of missing pixels with missing temperature or polarization data by taking the average of the neighboring pixels.

While a few hundred or thousand pixels in a map may be missing data, this is insignificant compared to the $\sim 50 - 240$ million pixels overall (depending on the patch of sky). In the worst of the nighttime maps, roughly 0.002% of the pixels are missing polarization data. For most of the maps, that number is $10\times$ or more lower. None of the pixels within the region of interest are missing temperature data. Most of the day deep maps have no pixels missing either all or only polarization data, but they occasionally have up to 0.008% of each. As for the day wide maps, the statistics vary from one detector array to the next. For PA4, the maps have $< 0.00005\%$ of pixels missing all data, and $\lesssim 0.001\%$ of pixels missing only the polarization data. For PA5, the maps usually have no pixels with missing data at all. For PA6, unfortunately, for some splits, over 0.5% of pixels (~ 1.5 million) within the outer boundary of the footprint have no data at all.

So for PA6, unfortunately, the hole-filling technique described above is not a good solution. The number of pixels with missing values is large enough that we cannot be confident that filling them in will have a negligible effect on the resulting power spectra. One solution would be to discard the PA6 data in the wide patch entirely.

Another solution would be to use only 2 splits for PA6 instead of 4.¹² Ideally, we would re-map the PA6 data. This would ensure the TODs would be optimally split into 2 groups. However, to avoid this computationally intensive process, we can coadd pairs of the existing maps, to obtain effectively two map splits. This results in a manageable number of pixels missing data (< 0.0002%).

5.4.2 Point Source Subtraction

Before Fourier-space filtering the maps for the power spectrum analysis as described in Section 4.2.4, we first subtract the brightest point sources. This is to avoid introducing any artifacts in the maps when filtering due to localized sharp, bright signals. When doing this, we use source-free and source-only maps. The output of the ACT map-making, however, is simply a map of the sky. In order to separate out the point sources, for each season/region/array/frequency/map split, we use our beam model to fit each point source in the map, fitting for the amplitude of each source independently. This allows us to create a map of the sources (beam model times amplitude for each source), and a source-free map, by subtracting the map of the sources from the map of the full sky.

At these (still) early stages of the daytime analysis, however, we do not yet have a final model for the daytime beams. We tried using the source maps made from nighttime data, but this yields mixed results. As shown in Figure 5.13, sometimes the nighttime source map does an adequate job of subtracting a source from the corresponding daytime map, but other times it clearly falls short. As a next step, we tried building new source maps for the daytime, using the nighttime beam models, but re-fitting for the amplitude of each source using the daytime maps. This works well enough for all sources, as can also be seen in Figure 5.13. We use these source maps

¹²We would continue using 4 splits for the other arrays, since using fewer splits results in discarding more data when computing the spectra, as explained in Section 4.2.1. But of course it is preferable to discarding all the PA6 data.

to make source-free daytime maps, then add back the weak (non-masked) sources to the maps before filtering them.

One reason why the source amplitudes might be different for the day deep, day wide, and night maps may be variability of the sources. Some of the extragalactic sources seen by ACT, active galactic nuclei (AGN) are known to have fluxes that vary from year to year (e.g., [Wright et al., 2009](#); [Chen, X. et al., 2013](#)).¹³ This could be relevant since the day deep, day wide, and night maps contain data from different time periods (2017–2019, 2019–2021, and 2017–2021, respectively).

5.4.3 Optimal Weighting

The white noise level for measurements of the power spectra can be expressed as a single number describing the root-mean-square noise per pixel, given by

$$N_\ell = \frac{4\pi f_{\text{sky}} s^2}{t_{\text{obs}} B_\ell^2}, \quad (5.1)$$

where f_{sky} is the fraction of sky observed,¹⁴ s is the instantaneous sensitivity of the experiment, t_{obs} is the observation time, and B_ℓ^2 is the beam window function ([Tegmark, 1997](#)). So, at high multipoles ℓ ($\ell \gtrsim 3000$ in this case for ACT), where the uncertainties in the spectra are dominated by noise and not cosmic variance (see Section 4.1.1), the error bars on power spectra should scale as roughly $\propto 1/t_{\text{obs}}$.

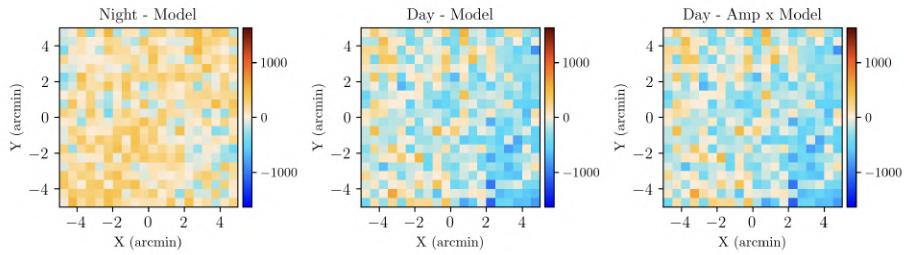
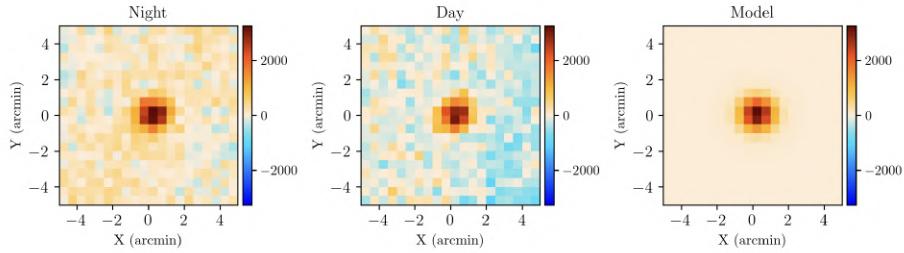
The day wide maps contain, very roughly, 1/3 the amount of data as the DR6 nighttime maps.¹⁵ So naively, one would expect the error bars on the day wide spectra to

¹³A study of the variability of AGN sources using ACT data is underway and will be the topic of a future paper.

¹⁴Or more precisely, an “effective” sky area where one accounts for the possible effects of pixel weighting ([Hivon et al., 2002](#)).

¹⁵There is just about the same amount of daytime data as nighttime data. The daytime observations are split roughly in half between day deep and day wide. Then approximately 1/3 of the data are discarded as part of the beam cuts, so only 2/3 of the data are kept. So in the end, the day wide maps contain $(1/2) \times (2/3) = 1/3$ as much data as the nighttime maps.

Source #2
 RA = 0.5500, Dec = -21.8860, Peak ratio (day/night) = 0.93, Fit amplitude = 1.01, Offset = -165.14



Source #89
 RA = 41.5000, Dec = -46.8550, Peak ratio (day/night) = 0.79, Fit amplitude = 0.65

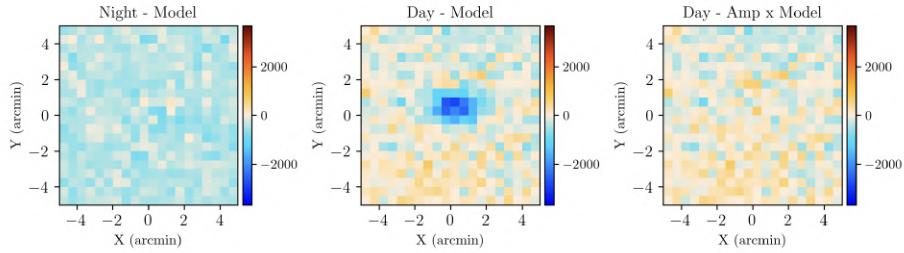
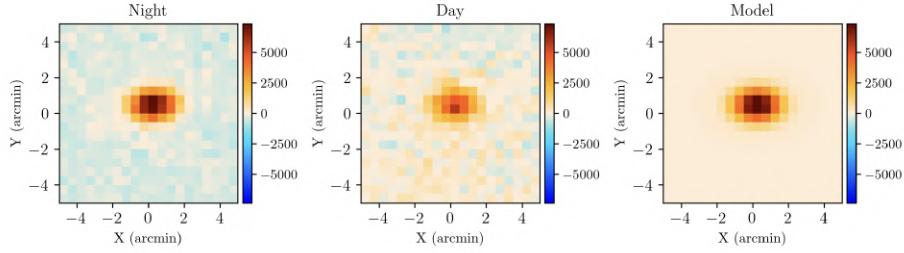


Figure 5.13: Two examples of point source subtraction in a map. The thumbnails show a point source in both a nighttime and daytime map, and the effect of subtracting the source model from each, both when using the default source model and when re-fitting the amplitude using the daytime maps. The top panel is an example of when the nighttime source model is appropriate (as-is) for subtracting the source from the daytime map, and the bottom panel is an example of when it is necessary to re-fit for the amplitude using the daytime maps.

be roughly 3 times the size of those of the DR6 spectra at high ℓ . Disappointingly, this was not the case when spectra were first computed. The initial day wide spectra, obtained using the default settings for DR6, had high- ℓ error bars 5 times larger than the nighttime ones. (For this investigation, based on the results from Section 5.4.1, we used spatial windows with no cross-linking/declination cut, and any holes due to missing pixels filled in.)

As can be seen in Figure 5.1, the inverse variance map for the wide daytime patch exhibits a greater relative variation in depth than the nighttime one. The deeper areas (to the north, on the left and right sides of the map) are relatively deeper, and the shallower areas (towards the center of the map) are relatively shallower. This significant difference in depth between portions of the maps motivated splitting the daytime wide maps into two sub-patches to be analyzed separately. While we were at it, we also split the daytime deep maps in two, since these already consisted of two distinct areas.

Based on the rationale for splitting the maps of the wide patch in two, a natural name for the wide sub-patches would be “wide shallow” (or “wide wide”) and “wide deep” (and similarly “deep shallow” (or “deep wide”) and “deep deep” for the deep sub-patches), but that would be confusing, to say the least. Instead, we name the sub-patches wide-north, wide-south, deep-north, and deep-south. These sub-patches are shown in Figure 5.14. As illustrated, the “north” and “south” labels make more sense in Galactic coordinates.¹⁶

Computing the power spectra for the day wide-north and day wide-south areas separately, then coadding these spectra (see Section 4.15) to obtain a final day wide power

¹⁶Unfortunately, we don’t avoid all possible sources of confusion, as the wide-south patch now includes the wide_01h_n (where the n stands for north) scanning region. However, this isn’t as much of a concern, since the regions used for scanning strategy purposes are rarely mentioned at the level of beam and power spectrum analysis work. The discussion of them in Section 5.3.2 was a rare exception, as the intricacies of the data at the TOD level were being explored.

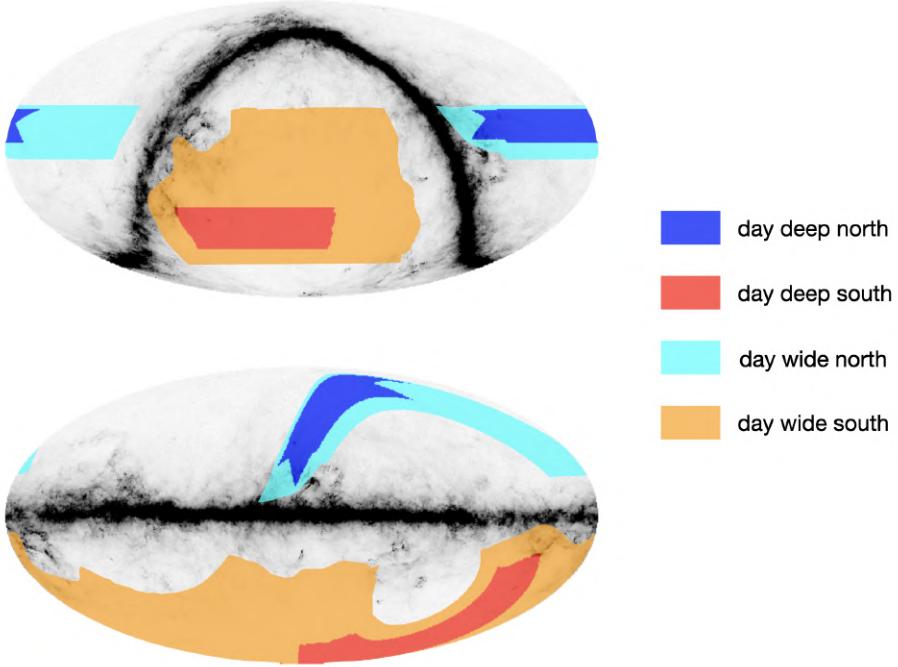


Figure 5.14: Mollweide projection of the daytime sky patches (split into north and south), in both equatorial (top) and Galactic (bottom) coordinates. The background is the *Planck* 353 GHz intensity map.

spectrum results in smaller uncertainties, as shown in Figure 5.15. The uncertainties are smallest if we use inverse variance weighting when computing the power spectra of the sub-patches, as opposed to uniform weighting. With this weighting, the error bars on the day wide spectra are indeed approximately 3 times the size of those of the DR6 nighttime spectra at high ℓ , as expected. The difference for the day deep spectra is not as pronounced, but there is still an improvement when analyzing the day deep-north and day deep-south areas separately using inverse variance weighting, as opposed to analyze day deep as a whole with uniform weighting. The day deep power spectrum error bars are roughly double the size of the DR6 ones.

The day wide and day deep spectra can be coadded, according to the procedure in Section 4.15, to obtain final daytime spectra. This involves computing the cross-patch spectra, so for deep \times wide, in order to properly account for the overlap between the

patches. This has not yet been done,¹⁷ so instead, the deep and wide spectra were roughly added together, ignoring any correlations, to get an estimate of the final error bars in Figure 5.15. The resulting high- ℓ error bars are approximately 1.6 times the size of those for the DR6 nighttime spectra. This is consistent with what we expect based solely on the amount of data that are analyzed — the day maps contain about 2/3 as much data as the nighttime maps, so based on Equation 5.1, the resulting spectra should have error bars that are roughly 1.5 times larger.

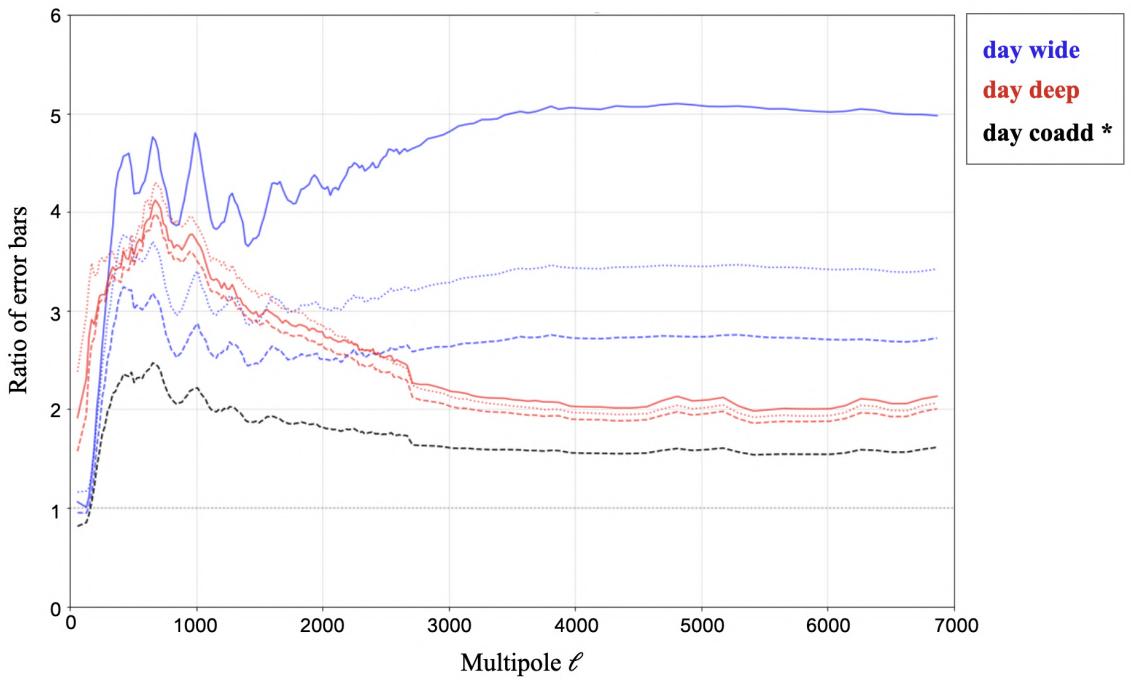


Figure 5.15: Ratio of the TT power spectrum error bars for PA4 at 150 GHz for various daytime scenarios, vs the default (preliminary) DR6 nighttime error bars. The **blue** curves are for the day wide region, and **red** for the day deep region, using uniform (solid curves) or inverse-variance (dotted curves) spatial weighting across each region, or dividing the regions into two sub-patches which are analyzed separately and then co-adding the resulting spectra (dashed curves). The two best daytime spectra (dashed curves) are preliminarily co-added, ignoring spatial correlations, and the resulting error bars are represented by the dashed **black** curve. The slight discontinuity at $\ell \sim 2700$ is due to a change in the size of the bins.

¹⁷There is nothing stopping us, we just haven't had time yet.

While Figure 5.15 only shows the error bars for TT, the conclusions we draw from looking at TE or EE instead are the same, and the ratios of uncertainties for each case are, in fact, nearly identical.

In general, in the noise dominated regime, it is most optimal to use inverse variance weighting in the pseudo- C_ℓ estimation, whereas in the signal dominated regime, equal weighting for each pixel is optimal (Efstathiou, 2004).¹⁸ It is possible that an intermediate option, somewhere between uniform and inverse-variance weighting, might result in even smaller uncertainties for the daytime power spectra. To this end, one could construct a hybrid estimator, as explained in Efstathiou (2004).

The natural¹⁹ question that arises after the work on the daytime spectra above is “What happens if we apply these different weighting schemes to the nighttime data?”

Is the approach we are using for DR6 truly the most optimal?

To answer this burning question, we split the nighttime maps into two sub-patches, the same as for the daytime wide maps, computed the spectra using either uniform or inverse variance weighting within each sub-patch, then coadded the resulting spectra. We also tried analyzing the nighttime maps as one region, but with inverse variance weighting, instead of the default uniform spatial weighting. The power spectrum error bars obtained with each of these techniques are then compared to those obtained with the default DR6 approach, as shown in Figure 5.16.

Analyzing the nighttime maps as one region with inverse variance weighting leads to larger error bars at all multipoles compared to the uniform weighting.²⁰ Meanwhile, either of the weighting schemes with two sub-patches results in a modest ($\sim 10\%$)

¹⁸In these limits, the pseudo- C_ℓ estimator is mathematically equivalent to the much more computationally expensive (and impractical) maximum likelihood estimator.

¹⁹It was natural to me, at least.

²⁰These results are for a spatial window with no cross-linking/declination cut. When we instead use a footprint where the strip of the maps with low cross-linking has been removed, the magenta curve in Figure 5.16 moves down slightly. The inverse variance weighting error bars are then $\sim 10\text{--}15\%$ larger at low ℓ compared to the uniform weighting ones, and $\sim 2\text{--}3\%$ smaller at high ℓ . So overall, uniform weighting is still preferred when the map is analyzed as one region.

improvement in the error bars at high ℓ , and little to no improvement at low ℓ . This difference is small enough that it may not be worth implementing for DR6 (since it would mean re-working an elaborate pipeline), but it is useful to keep in mind.

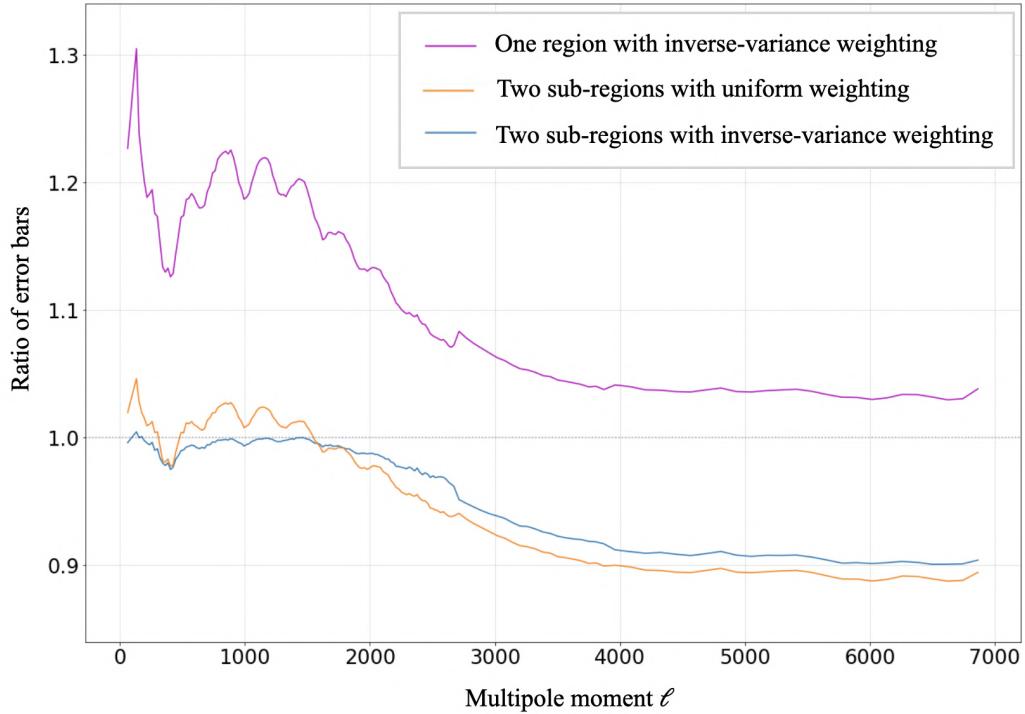


Figure 5.16: Ratio of the TT nighttime power spectrum error bars for PA4 at 150 GHz, for different spatial weighting choices. We show the ratio of error bars when splitting the maps into north and south regions, taking their spectra separately (with either uniform or inverse variance weighting), then coadding the resulting spectra, vs the default option used in DR6, where the maps are analyzed as one region with uniform weighting.

5.5 Beams

The time-variability of the beam deformations results in a possibly position-dependent effective beam in the maps. This effect is averaged down by repeated observations spread out across the sky. Our analysis relies on knowing the average effective daytime beam for each map. We plan to compute a correction to apply to the nighttime beam for each detector array/frequency/patch of daytime data in order to get an estimate

of the daytime beam. Two options for obtaining this correction are outlined below.

5.5.1 Temperature Beam

Map-Based Beam Correction

The baseline plan for obtaining daytime beams involves looking at the point sources in the maps. Using the nighttime beams as a baseline, we compute a Gaussian correction to obtain an effective beam for the daytime maps, in the same way as the nighttime jitter beams are obtained, as described in Section 3.4.5. This means model the daytime beams B_ℓ^{day} as

$$B_\ell^{\text{day}} = B_\ell^{\text{night}} \times e^{-\ell(\ell+1)V/2}, \quad (5.2)$$

where B_ℓ^{night} are the nighttime beams and V is the effective daytime jitter. This should account for the effective broadening of the beams in the daytime maps compared to the nighttime maps due to the greater variations in the beams during the day.

An example of such a daytime jitter fit is shown in Figure 5.17. It looks extremely similar to the equivalent plot of the nighttime jitter correction in Section 3.4.5.

As for the night, once the fits to the individual sources are performed, we construct a likelihood to obtain an estimate for the effective daytime jitter and its intrinsic scatter for each map, as shown in Figure 5.18.

We find *preliminary* best-fit values V for the daytime beam corrections for PA5 at 150 GHz of $\sim 115 - 145$ arcsec 2 for the different daytime sub-patches. These values represent additional pointing jitter, or effective beam smoothing, on top of the jitter correction computed for the nighttime beams. For comparison, the DR4 nighttime jitter values for 150 GHz ranged from ~ 0 to 65 arcsec 2 . Overall, the beams for the different daytime patches appear to be consistent with one another, as seen in

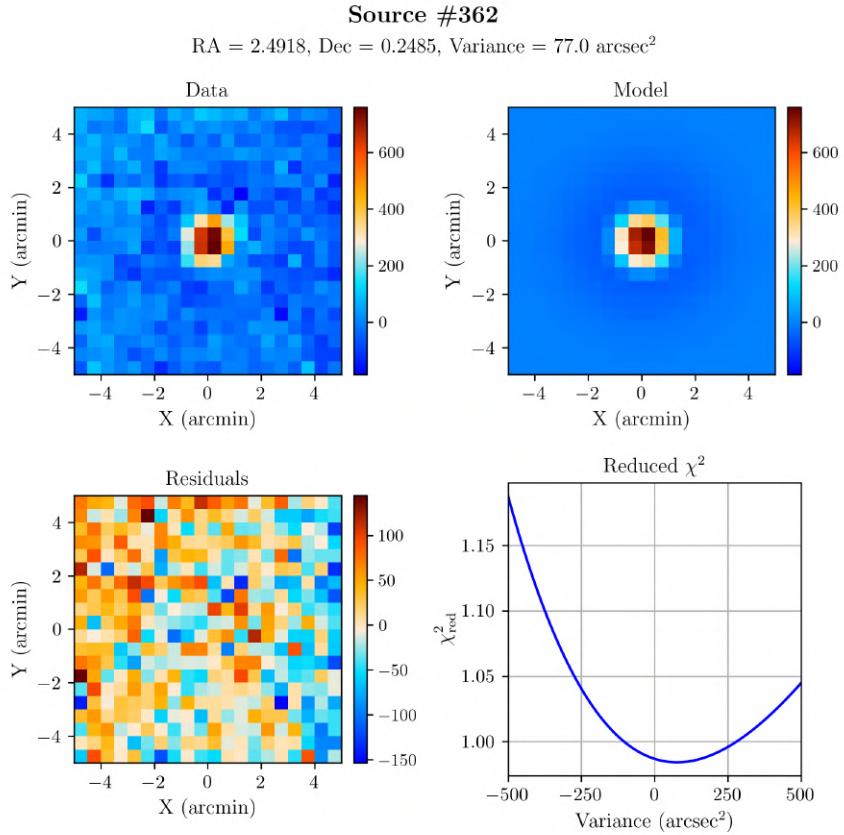


Figure 5.17: Daytime jitter fit for a source in the day wide-north region, for PA5 at 150 GHz.

Figure 5.19. This constitutes a first look at the daytime beam jitter corrections. It will be the subject of much more, in-depth future work.

This map-based correction should account for any *azimuthally symmetric* difference between the daytime and nighttime beams. However, this does not correct for any asymmetric features. There is some structure visible around the brightest point sources in daytime maps (after the bad beam cuts), at roughly 3% of the peak value, shown in Figure 5.20. This points to there being some structure in the beams that cannot be accounted for with a Gaussian correction. It remains to be seen if this will be an issue for the power spectrum analysis.

If it does prove to be a problem, one possible solution would be a similar treatment

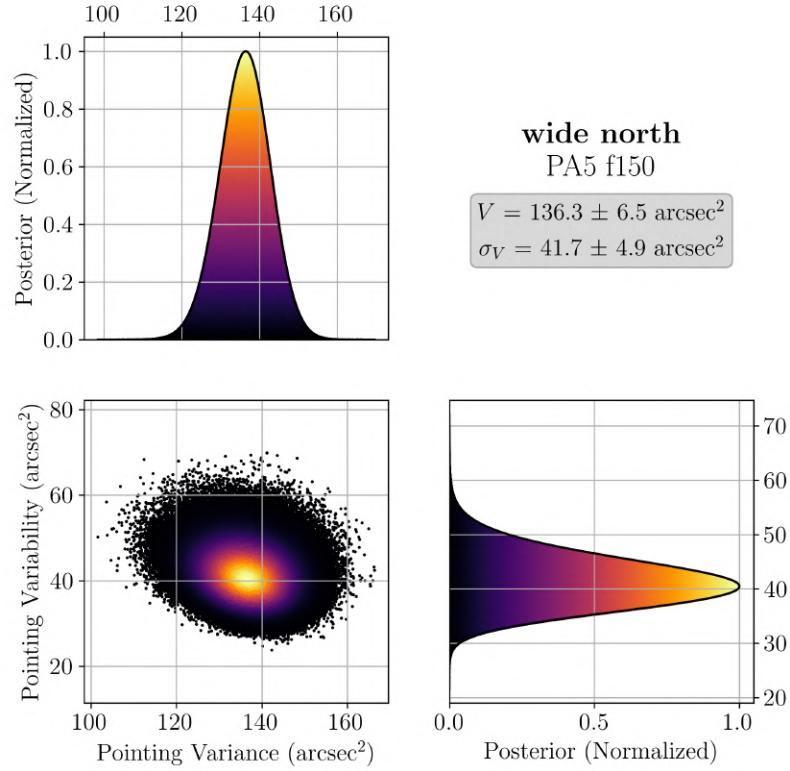


Figure 5.18: Preliminary distribution of the daytime jitter beam correction, for the day wide-north maps for PA5 at 150 GHz. The color scale represents the normalized posterior, in the range from 0 to 1.

as for the polarized sidelobes from Section 3.5.2. We could model the asymmetry as four sidelobes, each a spatially shifted copy of the main beam, and fit the amplitude of the temperature and polarization components using maps of bright point sources. We would then use this model to deproject the sidelobes from the time-ordered data before mapmaking.

Power-Spectrum-Based Beam Correction

In theory, a map-based beam correction should be able to account for any difference between the daytime and nighttime beams. If however, for some reason or other, the map-based beam corrections prove too challenging or time-consuming, we have a possible backup plan that can also act as a consistency check. It is similar to the

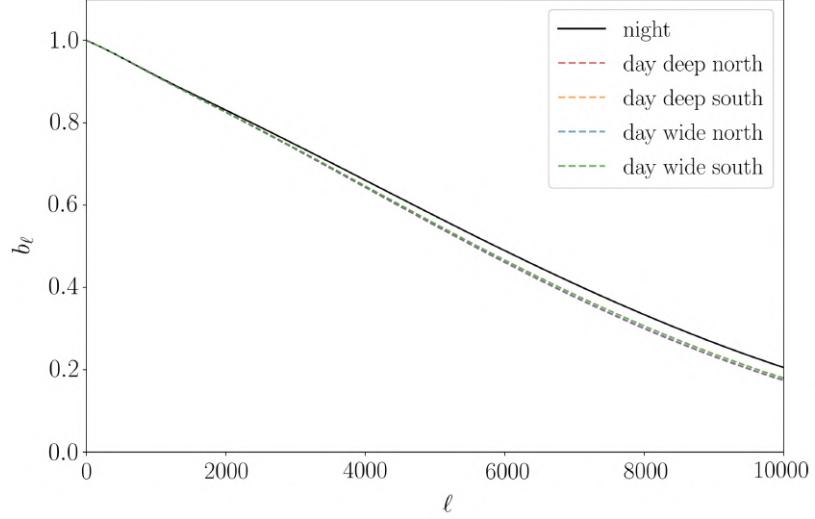


Figure 5.19: Preliminary beams for PA5 at 150 GHz, for the nighttime (DR6) and the different daytime patches. As expected, the beams are broader in the daytime. The uncertainties are not shown.

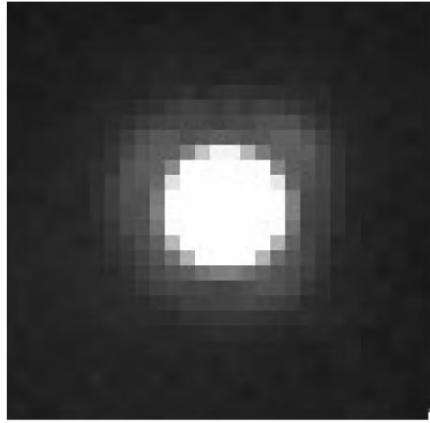


Figure 5.20: Close-up of a point source in a daytime (temperature) map. The color scale has been highly exaggerated in order to highlight the asymmetric beam structure. Still, it may be hard to see. To help guide the eye, the structure consists of four near sidelobes, distributed around the main beam in an \times pattern. Due to its appearance, this has sometimes been dubbed a “butterfly beam”.

correction applied to the daytime beams for DR5, described in [Naess et al. \(2020\)](#), although one should keep in mind that the DR5 maps were preliminary and not used for cosmological analysis.

For a given region of sky, array, and frequency, we compute the night \times night, night \times day, and day \times day temperature power spectra, assuming the nighttime beam for all calcu-

lations. (These are $\hat{C}_b^{TT, n \times n}$, $\hat{C}_b^{TT, n \times d}$, $\hat{C}_b^{TT, d \times d}$ in Equation 5.3.) When doing this, we make sure to mask the bright point sources (as we do in the regular power spectrum pipeline), in order to reduce any bias from point source variability.

In each multipole bin b , we model the 3 spectra using a day-to-night beam ratio α_b , assuming some noise-free underlying signal C_b^{TT} ,

$$\begin{aligned}\hat{C}_b^{TT, n \times n} &= C_b^{TT} + \text{noise}, \\ \hat{C}_b^{TT, n \times d} &= \alpha_b C_b^{TT} + \text{noise}, \\ \hat{C}_b^{TT, d \times d} &= \alpha_b^2 C_b^{TT} + \text{noise}.\end{aligned}\tag{5.3}$$

In each bin, we compute α_b by finding the maximum-likelihood solution to the system of equations above, marginalizing over the signal C_b^{TT} . This results in a noisy set of α_b values, to which we fit some smooth model. For [Naess et al. \(2020\)](#), this was a three parameter model: $A + (B - A) \exp\left\{-\frac{1}{2}\ell^2\sigma^2\right\}$, representing a gain error A , a high- ℓ loss of power B , and a Gaussian transition between the two. The beam ratio inferred from this smooth model is then applied to all spectra, in both temperature and polarization.

As part of the earlier work described in Section 5.2, we applied this technique to the spectra for the northern daytime regions from 2014–2016, for each season/detector array/frequency, one by one. The set of α_b values and the temperature spectra before and after correcting for the beam ratio are for one example case are shown in Figures 5.21 and 5.22. For one season/region/detector array/frequency, the polarization spectra are too noisy to see the effect of the correction. The daytime vs nighttime beam differences in Figures 5.21 and 5.22 are noticeably different from those inferred from the map-based correction in Figure 5.19. This may be due to a combination of several factors —the underlying data are from different time periods, sky regions, and detector arrays, and perhaps most importantly, no beam cuts were applied to the

2014–2016 daytime data used for Figures 5.21 and 5.22, whereas roughly 1/3 of the 2017–2021 daytime data used for Figure 5.19 were cut due to problematic beams.

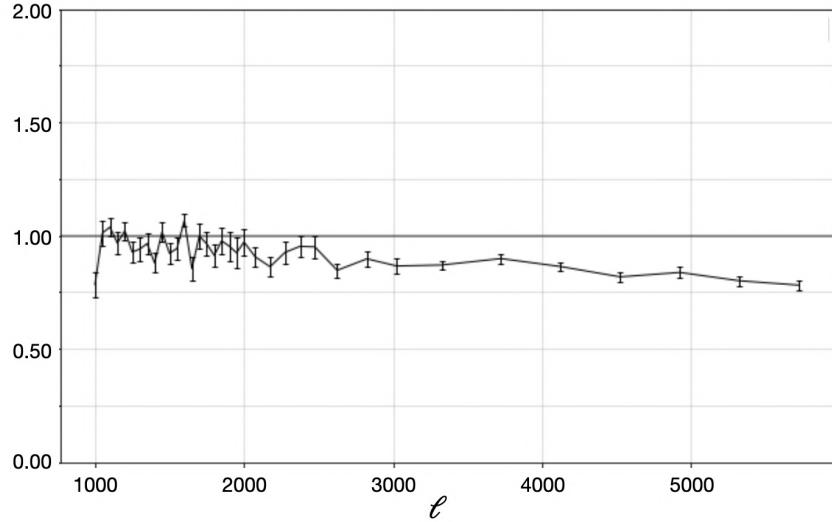


Figure 5.21: The day-to-night beam ratio α_b for S15 PA2 at 150 GHz, found using the power spectrum approach. Some smooth function is fit to these points before being used to correct the corresponding night \times day and day \times day power spectra.

5.5.2 Polarization and Leakage Beam

Whether it be map-based or power spectrum-based, the daytime beam correction will be computed using temperature data, but it will be applied to the main beam which is used for both temperature and polarization when calculating power spectra.

What about T-to-P leakage? The polarized sidelobes (Section 3.5.2) are deprojected from the time-ordered data prior to map-making. For the moment, the plan is to treat the leakage for the daytime maps the same way as for the nighttime maps. We may run some simulations, as briefly described in Section 3.6.2, to constrain the amount of T-to-P leakage that might be induced by daytime beam asymmetry if this becomes an area of concern.

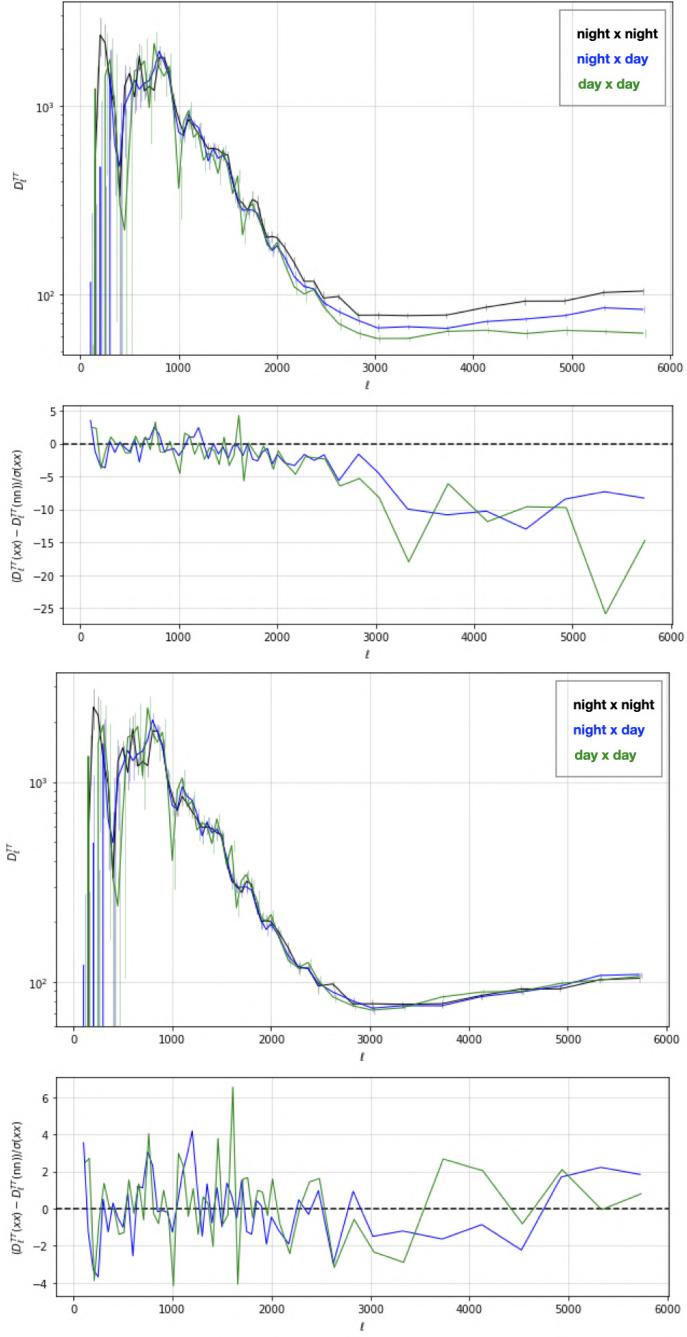


Figure 5.22: An illustration of the effect of the power-spectrum-based daytime beam correction. These plots show S15 PA2 150 GHz TT spectra, for night \times night, night \times day, and day \times day, computed in the BN region. The top (bottom) panel shows the spectra before (after) the correction for the effect of the daytime beam. Looking at the top panel, the relative loss of power on small scales (large ℓ) during the day is clearly visible.

5.6 Plans for Consistency and Null Tests

We plan to run similar consistency tests and map-based null tests as was done for DR4 (described in Sections 4.2.11 and 2.3, respectively), for both DR6 (nighttime) and the daytime analysis.

Consistency Tests

For both the DR6 nighttime and the daytime analysis, we will run a suite of consistency tests as described in Section 4.2.11. We will ensure that the spectra for different detector arrays are consistent, and for the daytime, that the spectra for different patches (and sub-patches) are consistent.²¹

Perhaps the most important consistency test for the daytime spectra will be an additional test of the consistency of the daytime and nighttime spectra. We will of course require that the daytime spectra pass all the same consistency tests as the nighttime before including them in our cosmological analysis. However, if the day vs night consistency tests are passed, this may enable certain other ACT analyses to proceed.

Map-Based Null Tests

The list of map-based nulls we will run for DR6 (nighttime) is still subject to change, but at minimum it includes PWV, elevation, and time. The PWV and elevation nulls are carried over from DR4.²² The time null test is new, and was added due to several years of data now being mapped together. For this test, the data are therefore split in half temporally (first few years vs last few years). As we further study the DR6 data, more null tests may be added to this list.

²¹We don't consider the spectra for different patches in the case of the nighttime, since it is mapped and analyzed in one patch, AA.

²²For DR4, we also ran a time constant null, as described in Section 2.3.1. This test was motivated by concerns that not properly correcting for the finite response time of each detector could lead to low-pass filtering of the data. However, no evidence of this was seen. There is no reason for this to change for DR6, and so this test will not be repeated for DR6.

For the daytime analysis, we will want to carry out all the same null tests as the night, and possibly more. The PWV and elevation null tests will, as expected, require the making of new maps. However, the time null test can be done with the existing maps, as the daytime data is already split in halves temporally —the deep patches were observed in the earlier years, and the wide patches in the later years. So the time null in this case is equivalent to an inter-patch consistency test.

In addition to those we run for the nighttime, we may want to run extra map-based null tests for the daytime, informed by our concerns about beam differences. One possibility is splitting the data by the hour at which it was taken, for example early vs late in the day, so roughly 11:00–17:00 UTC vs 17:00–23:00 UTC, or at the beginning and end vs the middle of the day, so roughly 11:00–14:00 UTC and 20:00–23:00 UTC vs 14:00–20:00 UTC. Another possibility is splitting the data based on the angle between the telescope pointing and the location of the Sun in the sky. There may be more. The main limitation is that we must be able to obtain reasonable map coverage for each null test split. For example, in the preliminary work described in Section 5.2, we found that an early vs late time split did not work for the maps being considered. Within each map, the early observations were predominantly in one half of the sky, and the late observations were predominantly in the other. So it was impossible to split the observations this way and have good coverage over the entire patch of sky being studied. For this reason, we instead ran a test where we split the TODs according to beginning and end vs middle of the day, for which it was possible to obtain maps with data spread more or less evenly over the patch.

This does naturally lead to another possible suite of tests we may want to run on the daytime maps — isotropy tests. Once again, these were done in Section 5.2, and involve computing the power spectrum of the left half vs the right half of the map, or the top half vs the bottom half. We are already partly doing this by splitting the

daytime maps into north and south sub-patches and running consistency tests on the resulting spectra. This could be further explored by splitting the sub-patches into left and right halves, or splitting the day wide-south patch into a top and bottom half. The usefulness of these tests would be limited by the uncertainties on the spectra.

5.7 Summary of Current Status & Next Steps

Some daytime power spectra have been computed, and while results so far look encouraging, no plots are shown here due to their preliminary nature, and the laundry list of caveats that accompany them.

Perhaps the largest potential issue is that both the nighttime and daytime beams are still preliminary (the daytime more so). Not all the sources of uncertainty for the nighttime beams (Section 3.4.6) have been included yet. We also have yet to account for the uncertainty associated with the daytime beam correction.

The calibration of each spectrum to *Planck*, as described in Section 4.2.7, has not yet been done. This calibration may be different for each patch/array/frequency and for night vs day.

Also, we have not begun quantifying the daytime leakage beams. For the DR4 analysis, the leakage was found to be small but non-negligible. The spectra were not corrected for this leakage directly; rather, we accounted for it in the likelihood. There is no guarantee the leakage is similarly small for the daytime, so it may have a more noticeable effect on the spectra.

While there are several details to work out, if we are able to combine DR6 night and day in an upcoming data release, we would reach an impressive survey depth, as shown in Figure 5.23. This will enable improved cosmological constraints on a variety of parameters, including the number of light relics, an early dark energy component,

and other extensions to the standard model.

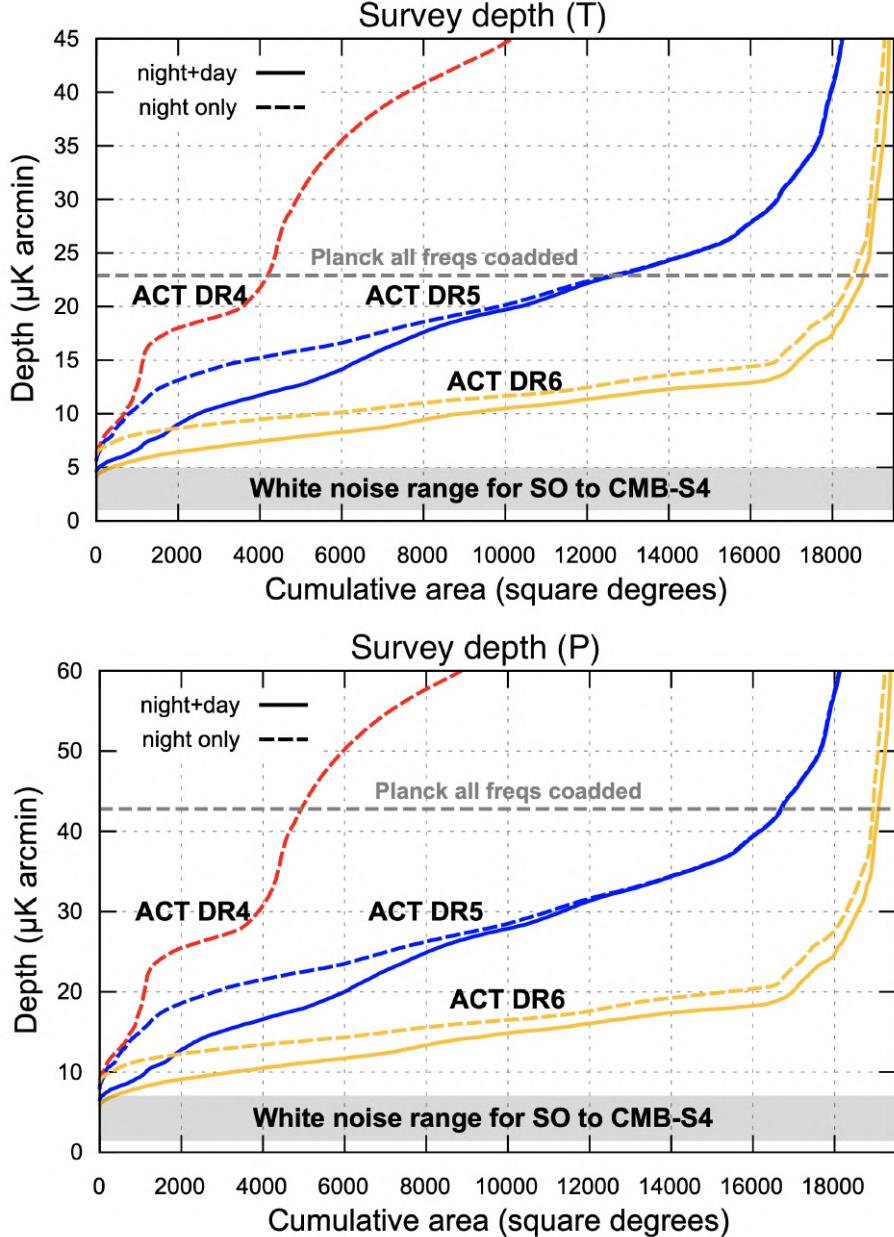


Figure 5.23: Cumulative white noise depth distributions for temperature (top) and polarization (bottom) for recent ACT data releases, as well as projections for DR6. This noise level is relevant at high multipoles ℓ , where the atmosphere is sub-dominant. Figure from Sigurd Naess.

Chapter 6

Conclusion

The cosmic microwave background is, in one sense, the final frontier. It's the farthest back we can look, the closest we can get to the origin of the Universe. The CMB photons we observe contain information not only about conditions in those early times, but about the long journey they took across the Universe to reach us. They offer us a window into the evolution of the Universe as a whole.

In this thesis, in Chapter 1 we outlined the physics of the early universe and its evolution, and gave an overview of ACT, an experiment used to study the CMB.

In Chapter 2, we described how the time-ordered data are transformed into maps of the sky, including corrections and subtle effects such as transfer functions. We also considered additional maps, made for the sole purpose of conducting null tests targeting particular systematic effects.

In Chapter 3, we quantified ACT's beams, which determine how it responds to different scales on the sky. This is an essential step in characterizing the instrument; an accurate estimate of the beams and their uncertainties is needed for all science analyses, including the precise measurement of CMB power spectra.

In Chapter 4, we described the formalism for estimating power spectra from maps, and summarized the steps of the ACT power spectrum pipeline. We reproduced the ACT spectra using publicly available code, an important step for post-DR4 power spectrum analysis and community reproduction of our results. This involved an in-depth investigation into the various analysis choices which affect the results.

Finally, in Chapter 5, we explored the data collected by ACT during the day, which to date have not been used for cosmological analyses. This is because they present additional challenges, as the Sun’s heat causes time-dependent distortions to the telescope mirror, which leads to beam deformations that change on the timescale of hours. However, if we can overcome these challenges and come to understand the daytime data well enough to pass a stringent suite of tests for systematics, then we will have significantly more data to use for a variety of science analyses, including CMB power spectra.

Overall, throughout this dissertation, we have followed the main steps in the ACT data analysis pipeline, from time-ordered data to the power spectra which encode valuable information about the parameters describing our Universe. Along the way, we scrutinized several details, including transfer functions, window functions, filtering, and masks. While at first glance this may seem fastidious, it is crucial for the level of precision cosmology we are now achieving, and to enable robust tests of subtle departures from the Λ CDM model.

While ACT is no longer observing the sky, there is still much work to be done analyzing the data we already have. And the work associated with this thesis is far from finished.

There remains important work to be done on the data analysis for ACT’s next data release, DR6, which will contain over four times the amount of data in DR4. This should result in noticeably tighter cosmological constraints. But first we must ensure

we fully understand all the subtleties of the data and its analysis, including, for example, our beams (including T-to-P leakage), and any relevant transfer functions. Crucially, we also need to pass a suite of consistency and null tests.

The same is true for the daytime data. Great progress has been made in tackling the specific challenges they present, but as for the nighttime data, our cosmological goals require precise measurements of our beams and any other factors which may impact our power spectra.

There are also data from the ACT low-frequency array, which is not included in our current analysis. Advanced ACTPol has been collecting data with this array (PA7), with bands centered at roughly 30 and 40 GHz, since 2020. This will provide leverage in removing Galactic synchrotron and extragalactic radio source contamination from our data, allowing us to create improved component-separated maps of CMB temperature and the thermal SZ effect, for example. Data at new frequencies reveals new and interesting information, but also comes with additional challenges. The beams behave differently, and so additional effort is required to characterize them. In the case of PA7, the leakage beams in particular seem different from what we are used to seeing.

Very soon, SO will begin observing, with noise levels a factor of three times lower than ACT. As these data are more sensitive, we will be even more susceptible to systematics (during both the day and the night), so new techniques will have to be developed to deal with these. Further developments could also be made to our data analysis pipeline, to avoid inducing transfer functions with our map-making and Fourier-space filtering algorithms, for instance. While it will take some work, SO promises to provide breakthrough discoveries in fundamental physics, cosmology, and astrophysics.

Bibliography

Abramowitz, M., & Stegun, I. A. 1972, Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables, tenth printing edn. (New York City: Dover), 364

Ade, P. A. R., Pisano, G., Tucker, C., & Weaver, S. 2006, [Proc. SPIE, 6275, 248](#)

Ade, P. A. R., Amiri, M., Benton, S. J., et al. 2022, [ApJ, 927, 174](#)

Ahmed, Z., Amiri, M., Benton, S. J., et al. 2014, [Proc. SPIE, 9153, 91531N](#)

Aiola, S., Calabrese, E., Maurin, L., et al. 2020, [JCAP, 2020, 047](#)

Akaike, H. 1973, Information theory and an extension of the maximum likelihood principle, in: B. N. Petrov and F. Csaki (eds), 2nd International Symposium on Information Theory (Budapest: Akadémiai Kiadó), 267-281

Akaike, H. 1974, [IEEE Trans. Automat. Contr., 19, 716](#)

Alexander, S., & McDonough, E. 2019, [Phys. Lett. B, 789, 197](#)

Allison, R., Caucal, P., Calabrese, E., Dunkley, J., & Louis, T. 2015, [Phys. Rev. D, 92, 123535](#)

Alonso, D., Sanchez, J., & and, A. S. 2019, [MNRAS, 484, 4127](#)

Alpher, R. A., & Herman, R. 1948, [Nature, 162, 774](#)

Alpher, R. A., & Herman, R. C. 1948, *Phys. Rev.*, 74, 1737

Anderson, A. J., Ade, P. A. R., Ahmed, Z., et al. 2018, *J. Low Temp. Phys.*, 193, 1057

Andersson, B. G. 2015, in *Astrophysics and Space Science Library*, Vol. 407, *Magnetic Fields in Diffuse Media*, ed. A. Lazarian, E. M. de Gouveia Dal Pino, & C. Melioli, 59

Arnold, K., Ade, P. A. R., Anthony, A. E., et al. 2010, *Proc. SPIE*, 7741, 404

Balkenhol, L., Dutcher, D., Mancini, A. S., et al. 2022, [arXiv:2212.05642](https://arxiv.org/abs/2212.05642)

Bashinsky, S., & Seljak, U. 2004, *Phys. Rev. D*, 69

Baumann, D., Green, D., Meyers, J., & Wallisch, B. 2016, *JCAP*, 2016, 007

Baumann, D., Green, D., & Wallisch, B. 2018, *JCAP*, 2018, 029

Bennett, C. L., Bay, M., Halpern, M., et al. 2003, *ApJ*, 583, 1

Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, 208, 20

Bernardeau, F. 1997, *A&A*, 324, 15

Bersanelli, M., Mandolesi, N., Butler, R. C., et al. 2010, *A&A*, 520, A4

BICEP/Keck Collaboration. 2021, *Phys. Rev. Lett.*, 127, 151301

Birkinshaw, M. 1999, *Physics Reports*, 310, 97

Blanchard, A., & Schneider, J. 1987, *A&A*, 184, 1

Bond, J. R., & Efstathiou, G. 1987, *MNRAS*, 226, 655

Born, M., & Wolf, E. 1980, *Principles of Optics*, 6th edn. (Oxford, England: Pergamon Press)

Brown, M. L., Castro, P. G., & Taylor, A. N. 2005, [MNRAS](#), **360**, 1262

Calzetti, D., Armus, L., Bohlin, R. C., et al. 2000, [ApJ](#), **533**, 682

Carlstrom, J. E., Ade, P. A. R., Aird, K. A., et al. 2011, [PASP](#), **123**, 568

Chang, C. L., Ade, P. A. R., Aird, K. A., et al. 2009, [AIP Conf. Proc.](#), **1185**, 475

Chen, X., Rachen, J. P., López-Caniego, M., et al. 2013, [A&A](#), **553**, A107

Choi, S. K., & Page, L. A. 2015, [JCAP](#), **2015**, 020

Choi, S. K., Austermann, J., Beall, J. A., et al. 2018, [J. Low Temp. Phys.](#), **193**, 267

Choi, S. K., Hasselfield, M., Ho, S.-P. P., et al. 2020, [JCAP](#), **2020**, 045

Corkish, R. 1990, [IEEE Antennas Propag. Mag.](#), **32**, 6

Darwish, O., Madhavacheril, M. S., Sherwin, B. D., et al. 2020, [MNRAS](#), **500**, 2250

Das, S., Marriage, T. A., Ade, P. A. R., et al. 2011, [ApJ](#), **729**, 62

Datta, R., Aiola, S., Choi, S. K., et al. 2018, [MNRAS](#), **486**, 5239

Davies, R. D., Dickinson, C., Banday, A. J., et al. 2006, [MNRAS](#), **370**, 1125

de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, [MNRAS](#), **388**, 247

de Salas, P. F., & Pastor, S. 2016, [JCAP](#), **2016**, 051

Deutsch, A.-S., Dimastrogiovanni, E., Johnson, M. C., Münchmeyer, M., & Terrana, A. 2018, [Phys. Rev. D](#), **98**, 123501

Dicke, R. H., Peebles, P. J. E., Roll, P. G., & Wilkinson, D. T. 1965, [ApJ](#), **142**, 414

Dickinson, C., Davies, R. D., & Davis, R. J. 2003, [MNRAS](#), **341**, 369

Dickinson, C., Peel, M., & Vidal, M. 2011, [MNRAS Lett.](#), **418**, L35

Dickinson, C., Ali-Haïmoud, Y., Barr, A., et al. 2018, *New Astron. Rev.*, 80, 1

Dodelson, S. 2003, *Modern Cosmology* (Amsterdam: Academic Press)

Draine, B. T., & Hensley, B. S. 2016, *ApJ*, 831, 59

Draine, B. T., & Lazarian, A. 1998, *ApJ*, 494, L19

Draine, B. T., & Lazarian, A. 1998, *ApJ*, 508, 157

—. 1999, *ApJ*, 512, 740

Duivenvoorden, A. J., Gudmundsson, J. E., & Rahlin, A. S. 2019, *MNRAS*, 486, 5448

Dunkley, J., Komatsu, E., Nolta, M. R., et al. 2009a, *ApJS*, 180, 306

Dunkley, J., Amblard, A., Baccigalupi, C., et al. 2009b, *AIP Conf. Proc.*, 1141, 222

Dunkley, J., Hlozek, R., Sievers, J., et al. 2011, *ApJ*, 739, 52

Dutcher, D., Balkenhol, L., Ade, P., et al. 2021, *Phys. Rev. D*, 104

Dünner, R., Hasselfield, M., Marriage, T. A., et al. 2013, *ApJ*, 762, 10

Efstathiou, G. 2004, *MNRAS*, 349, 603

Einstein, A. 1917, *Sitzungsberichte der König. Preuss. Akad.*, 142

Erickson, W. C. 1957, *ApJ*, 126, 480

Essinger-Hileman, T., Ali, A., Amiri, M., et al. 2014, *Proc. SPIE*, 9153

Fixsen, D. J. 2009, *ApJ*, 707, 916

Fixsen, D. J., & Mather, J. C. 2002, *ApJ*, 581, 817

Fowler, J. W., Niemack, M. D., Dicker, S. R., et al. 2007, *Appl. Opt.*, 46, 3444

Fuskeland, U., Wehus, I. K., Eriksen, H. K., & Naess, S. K. 2014, *ApJ*, 790, 104

Galitzki, N., Baildon, T., Barron, D., et al. 2018, *Proc. SPIE*, 10708, 1070804

Gamow, G. 1948, *Nature*, 162, 680

George, E. M., Reichardt, C. L., Aird, K. A., et al. 2015, *ApJ*, 799, 177

Gispert, R., Lagache, G., & Puget, J. L. 2000, *A&A*, 360, 1

Gnedin, N. Y., & Gnedin, O. Y. 1998, *ApJ*, 509, 11

Gold, B., Odegard, N., Weiland, J. L., et al. 2011, *ApJS*, 192, 15

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759

Grace, E., Beall, J., Bond, J. R., et al. 2014, *Proc. SPIE*, 9153, 915310

Grain, J., Tristram, M., & Stompor, R. 2009, *Phys. Rev. D*, 79

Gralla, M. B., Marriage, T. A., Addison, G., et al. 2020, *ApJ*, 893, 104

Gualtieri, R., Filippini, J. P., Ade, P. A. R., et al. 2018, *J. Low Temp. Phys.*, 193, 1112

Guth, A. H. 1981, *Phys. Rev. D*, 23, 347

Guth, A. H., & Weinberg, E. J. 1983, *Nucl. Phys. B*, 212, 321

Génova-Santos, R., Martín, J. A. R., Rebolo, R., et al. 2015, *MNRAS*, 452, 4169

Hajian, A., Acquaviva, V., Ade, P. A. R., et al. 2011, *ApJ*, 740, 86

Hajian, A., Viero, M. P., Addison, G., et al. 2012, *ApJ*, 744, 40

Hall, N. R., Keisler, R., Knox, L., et al. 2010, *ApJ*, 718, 632

Hamaker, J. P., & Bregman, J. D. 1996, *A&AS*, 117, 161

Han, D., Sehgal, N., MacInnis, A., et al. 2021, *JCAP*, 2021, 031

Hanany, S., & Rosenkranz, P. 2003, *New Astron. Rev.*, 47, 1159

Hand, N., Addison, G. E., Aubourg, E., et al. 2012, *Phys. Rev. Lett.*, 109, 041101

Hand, N., Leauthaud, A., Das, S., et al. 2015, *Phys. Rev. D*, 91, 062001

Hasselfield, M., Moodley, K., Bond, J. R., et al. 2013, *ApJS*, 209, 17

Hastings, W. K. 1970, *Biometrika*, 57, 97

Hauser, M. G., Arendt, R. G., Kelsall, T., et al. 1998, *ApJ*, 508, 25

Hazumi, M., Ade, P. A., Adler, A., et al. 2020, *Proc. SPIE*, 11443, 114432F

Henderson, S. W., Allison, R., Austermann, J., et al. 2016, *J. Low Temp. Phys.*, 184, 772

Henning, J. W., Sayre, J. T., Reichardt, C. L., et al. 2018, *ApJ*, 852, 97

Herman, D., Hensley, B., Andersen, K. J., et al. 2022, *A&A*

Hincks, A. D., Acquaviva, V., Ade, P. A. R., et al. 2010, *ApJS*, 191, 423

Hinshaw, G., Nolta, M. R., Bennett, C. L., et al. 2007, *ApJS*, 170, 288

Hivon, E., Gorski, K. M., Netterfield, C. B., et al. 2002, *ApJ*, 567, 2

Ho, S.-P. P., Austermann, J. A., Beall, J. A., et al. 2017, *Proc. SPIE*, 9914, 991418

Hoang, T., & Lazarian, A. 2016, *ApJ*, 821, 91

Hu, W., & Dodelson, S. 2002, *ARA&A*, 40, 171

Hu, W., Hedman, M. M., & Zaldarriaga, M. 2003, *Phys. Rev. D*, 67

Hu, W., Sugiyama, N., & Silk, J. 1997, *Nature*, 386, 37

Hui, H., Ade, P. A. R., Ahmed, Z., et al. 2018, *Proc. SPIE*, 10708, 1070807

Hurvich, C. M., & Tsai, C.-L. 1989, *Biometrika*, 76, 297

Ichiki, K. 2014, *Prog. Theor. Exp. Phys.*, 2014, 06B109

Ijjas, A., & Steinhardt, P. J. 2018, *Class. Quantum Gravity*, 35, 135004

Irwin, K., & Hilton, G. 2005, Transition-Edge Sensors, ed. C. Enss (Berlin, Heidelberg: Springer Berlin Heidelberg), 63

Kamionkowski, M., Kosowsky, A., & Stebbins, A. 1997, *Phys. Rev. D*, 55, 7368

Kamionkowski, M., & Kovetz, E. D. 2016, *ARA&A*, 54, 227

Keating, B., Timbie, P., Polnarev, A., & Steinberger, J. 1998, *ApJ*, 495, 580

Keating, B. G., Ade, P. A. R., Bock, J. J., et al. 2003, *Proc. SPIE*, 4843, 284

Keck Array, BICEP2 Collaborations. 2018, *Phys. Rev. Lett.*, 121

Keller, J. 1956, *IEEE Trans. Antennas Propag.*, 4, 312

—. 1959, *J. of Appl. Phys.*, 30, 1452

—. 1962, *J. Opt. Soc. Am.*, 52, 116

Kermish, Z. D., Ade, P., Anthony, A., et al. 2012, *Proc. SPIE*, 8452, 84521C

Knox, L. 1995, *Phys. Rev. D*, 52, 4307

Kogut, A., Spergel, D. N., Barnes, C., et al. 2003, *ApJS*, 148, 161

Kogut, A., Dunkley, J., Bennett, C. L., et al. 2007, *ApJ*, 665, 355

Krachmalnicoff, N., Carretti, E., Baccigalupi, C., et al. 2018, *A&A*, 618, A166

Lagache, G., Puget, J.-L., & Dole, H. 2005, *ARA&A*, 43, 727

Lamarre, J., Puget, J., Ade, P. A. R., et al. 2010, *A&A*, 520, A9

Lavaux, G., Diego, J. M., Mathis, H., & Silk, J. 2004, *MNRAS*, 347, 729

Lawson, K. D., Mayer, C. J., Osborne, J. L., & Parkinson, M. L. 1987, *MNRAS*, 225, 307

Lazarian, A., Andersson, B. G., & Hoang, T. 2015, in *Polarimetry of Stars and Planetary Systems*, 81

Lazarian, A., & Draine, B. T. 2000, *ApJL*, 536, L15

Ledoit, O., & Wolf, M. 2003, *J. Empir. Finance*, 10, 603

Leitch, E. M., Readhead, A. C. S., Pearson, T. J., & Myers, S. T. 1997, *ApJ*, 486, L23

Lewis, A., & Challinor, A. 2006, *Phys. Rep.*, 429, 1

Lewis, A., & Challinor, A. 2011, CAMB: Code for Anisotropies in the Microwave Background, Astrophysics Source Code Library, record ascl:1102.026

Li, Y., Austermann, J. E., Beall, J. A., et al. 2018, *Proc. SPIE*, 10708, 107080A

Li, Z., Louis, T., Calabrese, E., et al. 2021, *arXiv e-prints*, arXiv:2112.13839

Lifshitz, E. M. 1946, *Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki*, 16, 587

Louis, T., Garrido, X., Soussana, A., et al. 2019, *Phys. Rev. D*, 100, 023518

Louis, T., Naess, S., Garrido, X., & Challinor, A. 2020, *Phys. Rev. D*, 102, 123538

Louis, T., Addison, G. E., Hasselfield, M., et al. 2014, *JCAP*, 2014, 016

Louis, T., Grace, E., Hasselfield, M., et al. 2017, *JCAP*, 2017, 031

Ludwig, A. 1973, *IEEE Trans. Antennas Propag.*, 21, 116

Lungu, M., Storer, E. R., Hasselfield, M., et al. 2022, *JCAP*, 2022, 044

Ma, C.-P., & Bertschinger, E. 1995, *ApJ*, 455, 7

Macellari, N., Pierpaoli, E., Dickinson, C., & Vaillancourt, J. E. 2011, *MNRAS*, 418, 888

Madhavacheril, M. S., Hill, J. C., Næss, S., et al. 2020, *Phys. Rev. D*, 102

Mak, D. S. Y., Challinor, A., Efstathiou, G., & Lagache, G. 2017, *MNRAS*, 466, 286

Mallaby-Kay, M., Atkins, Z., Aiola, S., et al. 2021, *ApJS*, 255, 11

Mangano, G., Miele, G., Pastor, S., et al. 2005, *Nucl Phys B*, 729, 221

Mather, J. C. 1993, *Proc. SPIE*, 2019, 146

Mather, J. C., Fixsen, D. J., & Shafer, R. A. 1993, *Proc. SPIE*

Mather, J. C., Cheng, E. S., Cottingham, D. A., et al. 1994, *ApJ*, 420, 439

McMahon, J. J., Aird, K. A., Benson, B. A., et al. 2009, *AIP Conf. Proc.*, 1185, 511

Mennella, A., Butler, R. C., Curto, A., et al. 2011, *A&A*, 536, A3

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, *J. Chem. Phys.*, 21, 1087

Miville-Deschénes, M. A., Ysard, N., Lavabre, A., et al. 2008, *A&A*, 490, 1093

Morris, T. W., Bustos, R., Calabrese, E., et al. 2022, *Phys. Rev. D*, 105

Naess, S., Hasselfield, M., McMahon, J., et al. 2014, *JCAP*, 2014, 007

Naess, S., Aiola, S., Austermann, J. E., et al. 2020, *JCAP*, 2020, 046–046

Naess, S. K. 2019, *JCAP*, 2019, 060–060

Nagy, J. M., Ade, P. A. R., Amiri, M., et al. 2017, *ApJ*, 844, 151

Namikawa, T., Guan, Y., Darwish, O., et al. 2020, *Phys. Rev. D*, 101

Ogburn, R. W., I., Ade, P. A. R., Aikin, R. W., et al. 2010, *Proc. SPIE*, 7741, 77411G

O’Raifeartaigh, C., & Mitton, S. 2018, *Phys. Perspect.*, 20, 318

Padilla, I. L., Eimer, J. R., Li, Y., et al. 2020, *ApJ*, 889, 105

Padin, S., Staniszewski, Z., Keisler, R., et al. 2008, *Appl. Opt.*, 47, 4418

Pagano, L., Delouis, J.-M., Mottet, S., Puget, J.-L., & Vibert, L. 2020, *A&A*, 635, A99

Page, L., Jackson, C., Barnes, C., et al. 2003, *ApJ*, 585, 566–586

Page, L., Hinshaw, G., Komatsu, E., et al. 2007, *ApJS*, 170, 335

Peebles, P. J. E. 1968, *ApJ*, 153, 1

Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, 142, 419

Petroff, M. A., Eimer, J. R., Harrington, K., et al. 2020, *ApJ*, 889, 120

Piazzo, L. 2017, *IEEE Trans. Image Process.*, 26, 5232

Pietranera, L., Buehler, S. A., Calisse, P. G., et al. 2007, *MNRAS*, 376, 645

Planck HFI Core Team. 2011a, *A&A*, 536, A4

—. 2011b, *A&A*, 536, A6

Planck Collaboration XVIII. 2011, *A&A*, 536, A18

Planck Collaboration XXX. 2014, *A&A*, 571, A30

Planck Collaboration I. 2016, *A&A*, 594, A1

Planck Collaboration X. 2016, *A&A*, 594, A10

Planck Collaboration XI. 2016, [A&A](#), 594, A11

Planck Collaboration XXV. 2016, [A&A](#), 594, A25

Planck Collaboration I. 2020, [A&A](#), 641, A1

Planck Collaboration IV. 2020, [A&A](#), 641, A4

Planck Collaboration V. 2020, [A&A](#), 641, A5

Planck Collaboration VI. 2020, [A&A](#), 641, A6

Planck Collaboration IX. 2020, [A&A](#), 641, A9

Planck Collaboration XII. 2020, [A&A](#), 641, A12

Planck Collaboration Int. XV. 2014, [A&A](#), 565, A103

Planck Collaboration Int. XXX. 2016, [A&A](#), 586, A133

Planck Collaboration Int. LII. 2017, [A&A](#), 607, A122

Platania, P., Bensadoun, M., Bersanelli, M., et al. 1998, [ApJ](#), 505, 473

Platania, P., Burigana, C., Maino, D., et al. 2003, [A&A](#), 410, 847

Poidevin, F., Rubiño-Martín, J. A., Dickinson, C., et al. 2019, [MNRAS](#), 486, 462

POLARBEAR Collaboration. 2017, [ApJ](#), 848, 121

—. 2020, [ApJ](#), 904, 65

Pope, A. C., & Szapudi, I. 2008, [MNRAS](#), 389, 766–774

Prézeau, G., & Reinecke, M. 2010, [ApJS](#), 190, 267

Puget, J. L., Abergel, A., Bernard, J. P., et al. 1996, [A&A](#), 308, L5

Rees, M. J., & Sciama, D. W. 1968, [Nature](#), 217, 511

Reich, P., & Reich, W. 1988, A&A, 196, 211

Reichardt, C. L., Shaw, L., Zahn, O., et al. 2012, [ApJ, 755, 70](#)

Reichardt, C. L., Patil, S., Ade, P. A. R., et al. 2021, [ApJ, 908, 199](#)

Reinecke, M., Dolag, K., Hell, R., Bartelmann, M., & Enßlin, T. A. 2006, [A&A, 445, 373](#)

Riess, A. G., Yuan, W., Macri, L. M., et al. 2022, [ApJL, 934, L7](#)

Rubiño-Martín, J. A., López-Caraballo, C., Génova-Santos, R., & Rebolo, R. 2012, [Adv. Astron., 2012](#)

Ruhl, J., Ade, P. A. R., Carlstrom, J. E., et al. 2004, [Proc. SPIE, 5498, 11](#)

Ruze, J. 1966, Proc. IEEE, 54, 633

Rybicki, G. B., & Lightman, A. P. 1979, Radiative processes in astrophysics

Sachs, R. K., & Wolfe, A. M. 1967, [ApJ, 147, 73](#)

Sayre, J., Reichardt, C., Henning, J., et al. 2020, [Phys. Rev. D, 101](#)

Sazonov, S. Y., & Sunyaev, R. A. 1999, [MNRAS, 310, 765](#)

Schmid, H. M., Joos, F., & Tschan, D. 2006, [A&A, 452, 657](#)

Schäfer, J., & Strimmer, K. 2005, [Stat. Appl. Genet. Mol. Biol., 4](#)

Scott, D., Srednicki, M., & White, M. 1994, [ApJL, 421, L5](#)

Sharp, M. K., Marrone, D. P., Carlstrom, J. E., et al. 2010, [ApJ, 713, 82](#)

Sherwin, B. D., Dunkley, J., Das, S., et al. 2011, [Phys. Rev. Lett., 107, 021302](#)

Shewchuk, J. R. 1994, An Introduction to the Conjugate Gradient Method Without the Agonizing Pain

Shirokoff, E., Reichardt, C. L., Shaw, L., et al. 2011, *ApJ*, 736, 61

Sievers, J. L., Hlozek, R. A., Nolta, M. R., et al. 2013, *JCAP*, 2013, 060

Silk, J. 1968, *ApJ*, 151, 459

Silverberg, R. F., Hauser, M. G., Boggess, N. W., et al. 1993, *Proc. SPIE*, 2019, 180

Simons Observatory Collaboration. 2019, *JCAP*, 2019, 056

Smoot, G., Bennett, C., Weber, R., et al. 1990, *ApJ*, 360, 685

Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJL*, 396, L1

Spergel, D. N., Verde, L., Peiris, H. V., et al. 2003, *ApJS*, 148, 175

Spinelli, S., Fabbian, G., Tartari, A., Zannoni, M., & Gervasi, M. 2011, *MNRAS*, 414, 3272

Staniszewski, Z., Aikin, R. W., Amiri, M., et al. 2012, *J. Low Temp. Phys.*, 167, 827

Steinhardt, P. J., & Turok, N. 2002, *Science*, 296, 1436

Sunyaev, R. A., & Zeldovich, Y. B. 1972, *Comments Astrophys. Space Phys.*, 4, 173

—. 1980, *MNRAS*, 190, 413

Suzuki, A., Ade, P., Akiba, Y., et al. 2016, *J. Low Temp. Phys.*, 184, 805

Swetz, D. S., Ade, P. A. R., Amiri, M., et al. 2011, *ApJS*, 194, 41

Takakura, S., Aguilar-Faúndez, M. A. O., Akiba, Y., et al. 2019, *ApJ*, 870, 102

Tegmark, M. 1997, *Phys. Rev. D*, 56, 4514

Thornton, R. J., Ade, P. A. R., Aiola, S., et al. 2016, *ApJS*, 227, 21

Tolman, R. C. 1934, *Relativity, Thermodynamics, and Cosmology* (Clarendon Press)

Vidal, M., Dickinson, C., Davies, R. D., & Leahy, J. P. 2015, *MNRAS*, **452**, 656

Wandelt, B. D., & Górska, K. M. 2001, *Phys. Rev. D*, **63**, 123002

Wright, E. L. 1976, *ApJ*, **210**, 250

Wright, E. L., Chen, X., Odegard, N., et al. 2009, *ApJS*, **180**, 283

Zaldarriaga, M. 2004, in *Measuring and Modeling the Universe*, ed. W. L. Freedman, Vol. 2, 309

Zaldarriaga, M., & Seljak, U. 1999, *Phys. Rev. D*, **59**, 123507

Zaldarriaga, M., & Seljak, U. c. v. 1997, *Phys. Rev. D*, **55**, 1830

Zeldovich, Y. B., Kurt, V. G., & Syunyaev, R. A. 1968, *Zhurnal Eksperimentalnoi i Teoreticheskoi Fiziki*, **55**, 278