

https://dspace.org/

The Atacama Cosmology Telescope Mitigating the Impact of Extragalactic Foregrounds for the DR6 Cosmic Microwave Background Lensing Analysis

Niall MacCrann; Blake D. Sherwin; Frank J. Qu; Toshiya Namikawa; Matthew Hilton ...

IOP PUBLISHING LTD
https://hdl.handle.net/10539/40846
WOS

Downloaded from a DSpace repository. https://dspace.org/

The Astrophysical Journal, 966:138 (17pp), 2024 May 1 https://doi.org/10.3847/1538-4357/ad2610
© 2024. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

The Atacama Cosmology Telescope: Mitigating the Impact of Extragalactic Foregrounds for the DR6 Cosmic Microwave Background Lensing Analysis

Niall MacCrann[1,2], Blake D. Sherwin[1,2], Frank J. Qu[1], Toshiya Namikawa[3], Mathew S. Madhavacheril[4], Irene Abril-Cabezas[1], Rui An[5], Jason E. Austermann[6], Nicholas Battaglia[7], Elia S. Battistelli[8], James A. Beall[6], Boris Bolliet[1], J. Richard Bond[9], Hongbo Cai[10], Erminia Calabrese[11], William R. Coulton[12], Omar Darwish[13], Shannon M. Duff[6], Adriaan J. Duivenvoorden[12,14], Jo Dunkley[14,15], Gerrit S. Farren[1], Simone Ferraro[16], Joseph E. Golec[17,18], Yilun Guan[19], Dongwon Han[1], Carlos Hervías-Caimapo[20], J. Colin Hill[21], Matt Hilton[22], Renée Hložek[23,24,25], Johannes Hubmayr[6], Joshua Kim[4], Zack Li[9], Arthur Kosowsky[10], Thibaut Louis[26], Jeff McMahon[17,18,27,28], Gabriela A. Marques[29], Kavilan Moodley[30], Sigurd Naess[31], Michael D. Niemack[7,32], Lyman Page[14], Bruce Partridge[33], Emmanuel Schaan[34,35], Neelima Sehgal[36], Cristóbal Sifón[37], Edward J. Wollack[38], Maria Salatino[39], Joel N. Ullom[6], Jeff Van Lanen[6], Alexander Van Engelen[40], and Lukas Wenzl[7]

[1] DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 OWA, UK; nm746@cam.ac.uk
[2] Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge CB3 0HA, UK
[3] Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwa, 277-8583, Japan
[4] Department of Physics and Astronomy, University of Pennsylvania, 209 South 33rd Street, Philadelphia, PA 19104, USA
[5] Department of Physics and Astronomy, University of Southern California, Los Angeles, CA 90089, USA
[6] NIST, Quantum Sensors Group, 325 Broadway, Boulder, CO 80305, USA
[7] Department of Astronomy, Cornell University, Ithaca, NY 14853, USA
[8] Sapienza University of Rome, Physics Department, Piazzale Aldo Moro 5—I-00185, Rome, Italy
[9] Canadian Institute for Theoretical Astrophysics, University of Toronto, Toronto, ON M5S 3H8, Canada
[10] Department of Physics and Astronomy, University of Pittsburgh, Pittsburgh, PA 15260, USA
[11] School of Physics and Astronomy, Cardiff University, The Parade, Cardiff, Wales CF24 3AA, UK
[12] Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA
[13] Université de Genève, Département de Physique Théorique and Centre for Astroparticle Physics, 24 quai Ernest-Ansermet, CH-1211 Genève 4, Switzerland
[14] Joseph Henry Laboratories of Physics, Jadwin Hall, Princeton University, Princeton, NJ 08544, USA
[15] Department of Astrophysical Sciences, Peyton Hall, Princeton University, Princeton, NJ 08544, USA
[16] Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, USA
[17] Department of Physics, University of Chicago, 5720 South Ellis Avenue, Chicago, IL 60637, USA
[18] Kavli Institute for Cosmological Physics, University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637, USA
[19] Dunlap Institute for Astronomy and Astrophysics, University of Toronto, 50 St. George Street, Toronto, ON M5S 3H4, Canada
[20] Instituto de Astrofísica and Centro de Astro-Ingeniería, Facultad de Física, Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, 7820436 Macul, Santiago, Chile
[21] Department of Physics, Columbia University, 538 West 120th Street, New York, NY 10027, USA
[22] Wits Centre for Astrophysics, School of Physics, University of the Witwatersrand, Private Bag 3, Johannesburg 2050, South Africa
[23] David A. Dunlap Department of Astronomy & Astrophysics, 50 St George Street, Toronto, ON M5S 3H4, Canada
[24] Dunlap Institute for Astronomy & Astrophysics, 50 St George Street, Toronto, ON M5S 3H4, Canada
[25] Astrophysics Research Centre, University of KwaZulu-Natal, Westville Campus, Durban 4041, South Africa
[26] Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France
[27] Department of Astronomy and Astrophysics, University of Chicago, 5640 S. Ellis Avenue, Chicago, IL 60637, USA
[28] Enrico Fermi Institute, University of Chicago, Chicago, IL 60637, USA
[29] Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510, USA
[30] Astrophysics Research Centre, School of Mathematics, Statistics and Computer Science, University of KwaZulu-Natal, Durban 4001, South Africa
[31] Institute of Theoretical Astrophysics, University of Oslo, Norway
[32] Department of Physics, Cornell University, Ithaca, NY 14853, USA
[33] Department of Physics and Astronomy, Haverford College, Haverford, PA 19041, USA
[34] SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
[35] Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Stanford, CA 94305, USA
[36] Physics and Astronomy Department, Stony Brook University, Stony Brook, NY 11794, USA
[37] Instituto de Física, Pontificia Universidad Católica de Valparaíso, Casilla 4059, Valparaíso, Chile
[38] NASA Goddard Space Flight Center, 8800 Greenbelt Road., Greenbelt, MD 20771, USA
[39] Physics Department, Stanford University, 382 via Pueblo, Stanford, CA 94305, USA
[40] School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA

Received 2023 April 19; revised 2023 June 22; accepted 2023 July 6; published 2024 April 30

Abstract

We investigate the impact and mitigation of extragalactic foregrounds for the cosmic microwave background (CMB) lensing power spectrum analysis of Atacama Cosmology Telescope (ACT) data release 6 (DR6) data. Two independent microwave sky simulations are used to test a range of mitigation strategies. We demonstrate that finding and then subtracting point sources, finding and then subtracting models of clusters, and using a profile bias-hardened lensing estimator together reduce the fractional biases to well below statistical uncertainties, with the

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

inferred lensing amplitude, $A_{\rm lens}$, biased by less than 0.2σ. We also show that another method where a model for the cosmic infrared background (CIB) contribution is deprojected and high-frequency data from Planck is included has similar performance. Other frequency-cleaned options do not perform as well, either incurring a large noise cost or resulting in biased recovery of the lensing spectrum. In addition to these simulation-based tests, we also present null tests on the ACT DR6 data for sensitivity of our lensing spectrum estimation to differences in foreground levels between the two ACT frequencies used, while nulling the CMB lensing signal. These tests pass whether the nulling is performed at the map or bandpower level. The CIB-deprojected measurement performed on the DR6 data is consistent with our baseline measurement, implying that contamination from the CIB is unlikely to significantly bias the DR6 lensing spectrum. This collection of tests gives confidence that the ACT DR6 lensing measurements and cosmological constraints presented in companion papers to this work are robust to extragalactic foregrounds.

Unified Astronomy Thesaurus concepts: Cosmology (343); Large-scale structure of the universe (902); Weak gravitational lensing (1797); Cosmic microwave background radiation (322)

1. Introduction

Gravitational lensing provides a relatively direct method of probing the matter distribution in the universe, an otherwise challenging task given the domination of dark matter over visible matter. By measuring lensing statistics, we can therefore constrain the properties of dark energy and neutrinos by modeling their impact on the statistics of the matter distribution across a wide range in redshift. The cosmic microwave background (CMB) is a useful source of photons for lensing since its statistics in the absence of lensing are close to Gaussian and isotropic. Furthermore, we know its redshift precisely, which is required for inference since the strength of lensing depends on source redshift. The deflection of CMB photons by lensing, induced by a given realization of the matter field, breaks the statistical isotropy of the CMB, generating an off-diagonal covariance between Fourier modes. This can be used to construct quadratic estimators for the lensing convergence, which is closely related to the projected matter density field (see Lewis & Challinor 2006 for a review of CMB lensing).

One of the main challenges in robust CMB lensing estimation comes from the presence of other *secondary* anisotropies, i.e., physical processes between us and the surface of last scattering that generate millimeter radiation (such as emission from dusty galaxies and radio sources) or scatter the CMB photons, such as the thermal and kinematic Sunyaev–Zel'dovich (henceforth tSZ and kSZ, respectively) effects. These sources of statistical anisotropy can bias the lensing reconstruction. Fortunately, various methods have been developed that mitigate these effects.

First, the frequency dependence of these foregrounds can be exploited to project out a given spectral energy distribution (SED) or, equivalently, form linear combinations of individual frequency maps that null a given SED (e.g., Remazeilles et al. 2011; Madhavacheril & Hill 2018). High signal-to-noise ratio (S/N) point sources and clusters can be detected using a matched-filter approach (e.g., Haehnelt & Tegmark 1996; Staniszewski et al. 2009), for which models can be fitted and subtracted, or regions around these detections can be masked or "inpainted" (e.g., Bucher & Louis 2012). All of these methods aim to remove the contaminating contributions of foregrounds to the maps entering the lensing reconstruction.

Second, there are methods developed for lensing estimation specifically that amend the usual quadratic estimator for lensing reconstruction to reduce sensitivity to foregrounds. The quadratic estimator reconstructs lensing potential modes, $\phi(\boldsymbol{L})$, by averaging over pairs of temperature modes, $T(\boldsymbol{l})T(\boldsymbol{L}-\boldsymbol{l})$. *Bias hardening* (Namikawa & Takahashi 2014; Osborne et al. 2014; Sailer et al. 2020) involves amending the estimator such that it has zero response to the mode coupling generated by, for example, Poisson-distributed sources. Finally, the *shear* estimator has recently been developed by Schaan & Ferraro (2019) (and generalized to the curved sky by Qu et al. 2023). This estimator uses only the quadrupolar contribution to the coupling between observed CMB Fourier modes induced by lensing, which is largely unaffected by extragalactic foregrounds (Schaan & Ferraro 2019).

We note here that extragalactic rather than Galactic foregrounds are the focus of this paper. Galactic dust contamination falls sharply at small scales (high l), and the accompanying lensing power spectrum estimation in Qu et al. (2024) uses CMB modes at $l > 600$ only, which should ensure minimal contamination from Galactic foregrounds (Challinor et al. 2018; Beck et al. 2020). Qu et al. (2024) also demonstrate that the measured lensing power spectrum is stable to using an even more conservative $l_{\rm min}$. Furthermore, unlike extragalactic foregrounds, we can use the large-scale anisotropy of Galactic foregrounds (they are much higher in amplitude close to the Galactic plane) to test sensitivity. If the lensing power spectrum measurement was significantly biased by Galactic foregrounds, that bias would be highly sensitive to the strictness of the Galactic mask used; however, Qu et al. (2024) demonstrate that the measured lensing power spectrum is very stable to using a more conservative Galactic mask than the baseline choice.

In recent years the Planck satellite has provided the data for the state-of-the-art CMB lensing reconstruction and (auto-) power spectrum measurements (Planck Collaboration et al. 2020a; Carron et al. 2022), building on initial detections of CMB lensing cross-correlations with Wilkinson Microwave Anisotropy Probe satellite data (Smith et al. 2007; Hirata et al. 2008), and the lensing (auto) power spectrum by the ground-based Atacama Cosmology Telescope[41] (ACT) and the South Pole Telescope[42] (SPT) (Das et al. 2011; van Engelen et al. 2012; see, e.g., Omori et al. 2017; Sherwin et al. 2017; Wu et al. 2019; Bianchini et al. 2020; Millea et al. 2021 for subsequent analyses from these experiments). Much larger telescopes and detector arrays are feasible for ground-based experiments, allowing for higher-resolution, lower-noise observations. With enhanced instrumentation, ACT and SPT data can now achieve comparable statistical power to Planck.

[41] https://act.princeton.edu/
[42] https://pole.uchicago.edu/public/Home.html

Meanwhile, the upcoming Simons Observatory[43] SO and the planned CMB-S4[44] (Abazajian et al. 2016) will further increase sensitivity, especially in polarization.

However, these data come with challenges in controlling foregrounds. First, the higher CMB l-modes (i.e., smaller angular scales) accessible with these experiments are more contaminated by certain extragalactic foregrounds, especially dusty galaxies and the SZ effects (see, e.g., Osborne et al. 2014; van Engelen et al. 2014). Ground-based experiments also have reduced frequency coverage at these high CMB multipoles compared to that leveraged by Planck to perform multifrequency cleaning (Planck Collaboration et al. 2020c); for this analysis we use only f090 (77–112 GHz) and f150 (124–172 GHz) frequency bands from ACT.

The aforementioned analyses of previous ground-based data (e.g., Wu et al. 2019; Sherwin et al. 2017) achieved lensing power spectrum precision of $\gtrsim$5%, and thus the demands on the systematic uncertainties were lesser than for the ACT data release 6 (DR6) data (which achieve 2% precision), meaning that potential residual foreground biases at the few percent level (generally after masking or inpainting sources and SZ clusters) were not catastrophic. This lesser precision enabled, e.g., Sherwin et al. (2017) to marginalize over simulation-informed foreground contamination templates with 3% (on $A_{\rm lens}$) systematic uncertainty. With ACT DR6 data, this approach would significantly inflate the cosmological parameter uncertainties, so more robust quantification and mitigation (e.g., the use of bias hardening) of foreground biases are required.

We argue in the following that we can effectively mitigate such extragalactic foreground biases in ACT DR6 data, enabling robust lensing power spectrum measurements and cosmological inference, as described in companion papers Qu et al. (2024) and Madhavacheril et al. (2024). In Section 2, we briefly cover relevant formalisms for lensing quadratic estimators (including bias-hardening and frequency-cleaned estimators) and for quantifying biases from extragalactic foregrounds. In Section 3, we describe the ACT DR6 data and give a brief overview of the DR6 lensing analysis of Qu et al. (2024) to provide context for the subsequent foreground bias estimates. In Section 4, we describe the microwave-sky simulations used to test our mitigation strategies, and in Section 5 we present predicted biases from these simulations. In Section 6, we present null tests performed on the DR6 data, which test for sensitivity to the difference in foreground biases for the two ACT frequency channels used, and present consistency of a cosmic infrared background (CIB) deprojected analysis with our baseline analysis. We present our conclusions in Section 7.

2. Formalism and Methods

We briefly describe the formalism associated with the (bias-hardened) quadratic estimators and multifrequency cleaning approaches used in this work. Note that we focus on the CMB temperature field, since we assume that polarization is not affected by extragalactic foregrounds—certainly we do not expect fields such as the tSZ effect and CIB to be polarized at a significant level for our observations, and any bright polarized sources in the DR6 data are subtracted or masked and inpainted (Qu et al. 2024; S. Naess et al. 2024, in preparation).

2.1. Quadratic Estimator

We use throughout the quadratic estimator (Hu & Okamoto 2002), using the curved-sky calculations implemented in FALAFEL[45] and the normalization and reconstruction noise calculations implemented in TEMPURA.[46] For simplicity, we will present in the text the flat-sky equations, where the quadratic estimator for Fourier mode L of the lensing potential, $\phi_{\boldsymbol{L}}$, is given by (e.g., Hanson et al. 2011)

$$\hat{\phi}_{\boldsymbol{L}} = \frac{1}{2} A_L^{\phi} \int \frac{d^2\boldsymbol{l}}{(2\pi)^2} \frac{f_{L,l} T_l T_{L-l}}{C_l^{\rm tot} C_{|L-l|}^{\rm tot}}, \tag{1}$$

with the normalization factor A_L^{ϕ} given by

$$A_L^{\phi} = \left[\int \frac{d^2\boldsymbol{l}}{(2\pi)^2} \frac{f_{L,l}^2}{2 C_l^{\rm tot} C_{|L-l|}^{\rm tot}} \right]^{-1} \tag{2}$$

and

$$f_{L,l} = \tilde{C}_l \boldsymbol{l} \cdot \boldsymbol{L} + \tilde{C}_{|L-l|} \boldsymbol{L} \cdot |\boldsymbol{L} - \boldsymbol{l}|. \tag{3}$$

Here $\tilde{C}_l$ is the lensed CMB temperature power spectrum, while $C_l^{\rm tot}$ is the total observed CMB power spectrum (i.e., including noise). It is also common to work with the convergence, κ (and its power spectrum $C_L^{\kappa\kappa}$), instead of the lensing potential, which is simply related to ϕ via

$$\kappa_{\boldsymbol{L}} = L(L+1)\phi_{\boldsymbol{L}}/2. \tag{4}$$

The power spectrum of the reconstruction has what is known as an N^0 bias, which is the expectation for the quadratic estimator operating on a Gaussian, statistically isotropic field with the same power spectrum as the lensed CMB (i.e., in the absence of the statistcal anisotropy induced by lensing). For the basic, temperature-only, quadratic estimator described here, N^0 is equal to A_L^{ϕ}. N^0 can be calculated analytically or from simulations and subtracted off. The realization-dependent N^0 (RDN0) method of Hanson et al. (2011) and Namikawa et al. (2013) minimizes the impact of small differences between the CMB power spectrum assumed in the N^0 calculation and the true CMB power spectrum in the data.

2.2. Foreground Biases

Foregrounds perturb the CMB temperature, $T_{\rm CMB}$, as

$$T(\boldsymbol{l}) = T_{\rm CMB}(\boldsymbol{l}) + T^{\rm fg}(\boldsymbol{l}). \tag{5}$$

Denoting a quadratic estimator (see Equation (1)) on two temperature maps A and B as $Q(T^A, T^B)$, the contamination of $C_L^{\hat{\phi}\hat{\phi}}$ is (e.g., Osborne et al. 2014; van Engelen et al. 2014)

$$\begin{aligned} \Delta C_L^{\hat{\phi}\hat{\phi}} = & 2\langle Q(T^{\rm fg}, T^{\rm fg})\phi \rangle_L \\ & + 4\langle Q(T^{\rm fg}, T_{\rm CMB}) Q(T^{\rm fg}, T_{\rm CMB}) \rangle_L \\ & + \langle Q(T^{\rm fg}, T^{\rm fg}) Q(T^{\rm fg}, T^{\rm fg}) \rangle_L, \end{aligned} \tag{6}$$

where we have made the subtitution $Q(T_{\rm CMB}, T_{\rm CMB}) = \phi$ in the first line. The first and second terms are known as the *primary*

[43] https://simonsobservatory.org/
[44] https://cmb-s4.org/
[45] https://github.com/simonsobs/falafel
[46] https://github.com/simonsobs/tempura

and *secondary bispectrum* terms, respectively, since they depend on bispectra involving $T^{\rm fg}$ and ϕ (in the secondary bispectrum case the ϕ-dependent term arises from the presence of one $T_{\rm CMB}$ in each of the quadratic estimators). The third term is known as the trispectrum contribution, since it depends only on the trispectrum of $T^{\rm fg}$. Given a simulation of $T^{\rm fg}$ and ϕ, we can estimate these terms individually, and without noise from the primary CMB, following Schaan & Ferraro (2019) and briefly summarized here: The primary bispectrum term is computed directly as in the first line of Equation (6), using the true ϕ-field for the simulation. Noise on the secondary bispectrum estimate is reduced greatly by using the difference $\langle Q(T^{\rm fg}, T_{\rm CMB})Q(T^{\rm fg}, T_{\rm CMB})\rangle_L - \langle Q(T^{\rm fg}, T'_{\rm CMB})Q(T^{\rm fg}, T'_{\rm CMB})\rangle_L$, where $T'_{\rm CMB}$ is formed from the same unlensed CMB as $T_{\rm CMB}$, but lensed by an independent κ field. The estimator for the trispectrum term in Equation (6) has an "N^0 bias" (the disconnected trispectrum) that we subtract;[47] this is given by

$$N_L^{0,\rm fg} = (A_L^\phi)^2 \int \frac{d^2\boldsymbol{l}}{(2\pi)^2} \frac{f_{L,l}^2 C_l^{\rm fg} C_{|\boldsymbol{L}-\boldsymbol{l}|}^{\rm fg}}{2(C_l^{\rm tot} C_{|\boldsymbol{L}-\boldsymbol{l}|}^{\rm tot})^2}. \quad (7)$$

Note that the integral here is the same as for A_L^ϕ, but replacing $C_l^{\rm tot}$ with $(C_l^{\rm tot})^2/C_l^{\rm fg}$.

2.3. Bias-hardened Estimators

In general, a bias-hardened estimator for a field $x_{\boldsymbol{L}}$ in the presence of a contaminant $y_{\boldsymbol{L}}$ is (Namikawa et al. 2013; Osborne et al. 2014; Sailer et al. 2020)

$$x_{\boldsymbol{L}} = \frac{\hat{x}_{\boldsymbol{L}} - A_L^x R_L^{xy} \hat{y}_{\boldsymbol{L}}}{1 - A_L^x N_L^y (R_L^{xy})^2}, \quad (8)$$

where $\hat{x}_{\boldsymbol{L}}$ is the nonhardened estimator, given by

$$\hat{x}_{\boldsymbol{L}} = \frac{1}{2} A_L^x \int \frac{d^2\boldsymbol{l}}{(2\pi)^2} \frac{f_{L,l}^x T_l T_{L-l}}{C_l^{\rm tot} C_{|\boldsymbol{L}-\boldsymbol{l}|}^{\rm tot}}, \quad (9)$$

and $\hat{y}_{\boldsymbol{L}}$ is defined analogously. R^{xy} is the response of the estimator for field x to the presence of field y and is given by

$$R_L^{xy} = \int \frac{d^2\boldsymbol{l}}{(2\pi)^2} \frac{f_{L,l}^y f_{L,l}^x}{2C_l^{\rm tot} C_{|\boldsymbol{L}-\boldsymbol{l}|}^{\rm tot}}. \quad (10)$$

A_L^x is the estimator normalization, given by $1/R_L^{xx}$. Here f^x and f^y are functions that describe the mode coupling induced in the observed CMB by the fields x and y. For lensing, this is given by Equation (3), while for point sources it is a constant. Sailer et al. (2020) showed that this method can also harden against Poisson-distributed sources with some (Fourier space) radial profile $u(\boldsymbol{l})$ (with unknown amplitude), in this case,

$$f_{l,L-l} = \frac{u(\boldsymbol{l})u(\boldsymbol{L}-\boldsymbol{l})}{u(\boldsymbol{L})}. \quad (11)$$

Sailer et al. (2023) further demonstrate that good performance can be achieved even with an imperfect choice profile (e.g., CIB is mitigated even when using a profile chosen for tSZ).

[47] Note that in a real data analysis this would be accounted for by the RDN0 correction described above, in the usual case where, after some foreground mitigation, foreground power is subdominant to CMB power.

Since we are effectively deprojecting modes that contain some information on lensing, there is a noise cost to doing bias hardening (for the DR6 lensing analysis presented in Qu et al. 2024 it is ∼10%), with the N^0 of the bias-hardened estimator given by

$$N^{0,x,\rm BH} = N^{0,x}/(1 - N^{0,x}N^{0,y}(R^{xy})^2). \quad (12)$$

When computing the foreground trispectrum contamination, we again need to subtract an N^0 contribution (via the last term in Equation (6)), given by

$$N^{0,\rm fg,BH} = \frac{N^{0,\rm fg} + (A^x)^2(R^{xy}A^y - 2R^{xy}A^y R_{\rm fg}^{xy})}{[1 - A^x A^y (R^{xy})^2]^2}, \quad (13)$$

where $N^{0,\rm fg}$ is given by Equation (7) (and is also equal to $(A_L^x)^2/R_L^{xx,\rm fg}$),

$$R^{xy,\rm fg} = \int \frac{d^2\boldsymbol{l}}{(2\pi)^2} \frac{f_{L,l}^y f_{L,l}^x C_l^{\rm fg} C_{|\boldsymbol{L}-\boldsymbol{l}|}^{\rm fg}}{2C_l^{\rm tot} C_{|\boldsymbol{L}-\boldsymbol{l}|}^{\rm tot}}, \quad (14)$$

and $C_l^{\rm fg}$ is the power spectrum of the foreground-only temperature field.

2.4. Frequency-cleaned, Asymmetric and Symmetrized Estimators

Even for the temperature-only case, we need not use the same two maps in our quadratic estimator for the lensing potential; we can also use two temperature maps that have had different levels of foreground cleaning applied, which can result in a better bias−variance trade-off than using the same (noisy) frequency-cleaned temperature map in both legs of the quadratic estimator. Indeed, if just one of the maps has zero foreground contamination, we might expect the quadratic estimator to be unbiased by foregrounds in cross-correlation with the true convergence field or some tracer of it (see Hu et al. 2007; Madhavacheril & Hill 2018; Darwish et al. 2021). However, the secondary bispectrum contribution to the auto-power spectrum of the reconstruction is not removed in general.

The standard quadratic estimator in Equation (1) is not symmetric when two different temperature maps are used, i.e., $\phi_{XY} = Q(X, Y)$ is not in general equal to $\phi_{YX} = Q(Y, X)$, for $X \neq Y$. Darwish et al. (2023) therefore propose using a minimum-variance linear combination of the two asymmetric estimators:

$$\phi_{\rm sym}(\boldsymbol{L}) = \boldsymbol{W}(\boldsymbol{L}) \begin{pmatrix} \phi_{XY}(\boldsymbol{L}) \\ \phi_{YX}(\boldsymbol{L}) \end{pmatrix}, \quad (15)$$

where $\boldsymbol{W}(\boldsymbol{L})$ is some weight matrix. We define the matrix

$$N^0(\boldsymbol{L}) = \begin{pmatrix} N^{0,XYXY}(\boldsymbol{L}) & N^{0,XYYX}(\boldsymbol{L}) \\ N^{0,YXXY}(\boldsymbol{L}) & N^{0,YXYX}(\boldsymbol{L}) \end{pmatrix}, \quad (16)$$

where $N^{0,ABCD}(\boldsymbol{L})$ is the variance of the lensing reconstruction power spectrum for the general case of the cross-correlation of lensing reconstructions ϕ_{AB} and ϕ_{CD}, derived from CMB temperature maps A, B, C, D (see end of Section 2.1 for discussion of N^0).

Then, the normalized minimum variance $\boldsymbol{W}(\boldsymbol{L})$ is given by

$$W(\boldsymbol{L}) = \frac{(N^0)^{-1}}{|N^0|(N^{0,XYXY} - N^{0,YXYX})}, \quad (17)$$

where we have used that $N^{0,XYYX} = N^{0,YXXY}$. We shall refer to this as the D21 estimator.

3. ACT DR6 Data and Lensing Power Spectrum Analysis

We briefly describe the ACT DR6 data and how they are used for the lensing power spectrum analysis described in Qu et al. (2024) and Madhavacheril et al. (2024), including summarizing the baseline foreground mitigation, which will provide useful context for our description of the simulation processing in Section 4.2.

ACT DR6 includes maximum likelihood maps at 0.′5 pixel scale. The lensing power spectrum measurement (Qu et al. 2024) that this paper aims to support uses the first science-grade version of the ACT DR6 maps, labeled dr6.01 and generated from observations performed between 2017 May and 2021 June.[48] For the lensing analysis, we use the f090 band and f150 band data (with central frequencies ∼97 and ∼150 GHz, respectively). For each frequency, for temperature (T) and polarization data (Q and U), an optimal coadd (including data taken during the daytime not used in the lensing analysis) is generated, and a point-source catalog is generated using a matched-filter algorithm. These point-source catalogs are used to subtract point-source models from each single-array map[49] before they enter a weighted coaddition step in harmonic space, with weights given by the inverse noise power spectrum of a given map (see Section 4.6 of Qu et al. 2024 for further details). In addition, for the temperature data in our baseline approach, we subtract a map that models the tSZ contributions from detected clusters, which is a superposition of the best-fit templates inferred by the matched-filter cluster-finding code NEMO[50] (Hilton et al. 2021), for S/N > 5 detections (see Hilton et al. 2021; Section 2.1.3 of Qu et al. 2024).

The lensing reconstruction and power spectrum estimation are performed using the "cross-correlation-only" estimator of Madhavacheril et al. (2021), with our implementation using maps constructed from four independent splits of the DR6 data. The final lensing power spectrum estimate is constructed by combining the various different estimators one can form from using a different one of the four independent data splits in each of the four legs of the lensing power spectrum estimator (see Section 4.8.1 of Qu et al. 2024). We use a *"minimum-variance"*[51] (henceforth MV) lensing reconstruction that combines temperature and polarization information via a weighted sum of the reconstructions from all pairs $XY \in [TT, TE, TB, EE, EB]$:

$$\hat{\phi}^{\rm MV}(\boldsymbol{L}) = \sum_{XY} w_{XY} \hat{\phi}^{AB}(\boldsymbol{L}), \tag{18}$$

with weights given by the inverse of the normalization (Equation (2) for the TT case).

[48] Since these maps were generated, the ACT team has made some refinements to the mapmaking that improve the large-scale transfer function and polarization noise levels and include data taken in 2022. The team expects to use a second version of the maps for the DR6 public data release and for further lensing analyses.

[49] That is, each map corresponding to a portion of the data observed by a given detector array at a certain frequency.

[50] https://nemo-sz.readthedocs.io/en/latest/

[51] This is in quotes because our implementation does not take into account all cross-correlations between pairs AB in Equation (22); see Qu et al. (2024) for more details.

Through the tests described in this paper, we arrive at a set of foreground mitigation analysis choices for the baseline ACT DR6 lensing analysis, which we summarize briefly here:

1. Models for S/N > 4 point sources are subtracted from each map entering the coadd. This helps mitigate the impact of radio point sources and the brighter among the dusty galaxies that constitute the CIB.
2. A cluster tSZ model map, generated from cluster candidates identified by the NEMO code, is subtracted from each frequency map (see Section 4.2 for more details).
3. A bias-hardened estimator is used for the TT contribution to the $\hat{\phi}^{\rm MV}(\boldsymbol{L})$ estimate, with profile proportional to $(C_l^{\rm tSZ})^{0.5}$, where $C_l^{\rm tSZ}$ is the tSZ power spectrum, which we estimate from the WEBSKY simulations (see Section 2.3 for discussion of bias hardening). We will refer to this estimator as the *profile-hardened estimator*.

4. Simulations and Processing

In this section, we use two extragalactic microwave-sky simulations (described in Section 4.1), to make predictions for biases caused by contamination due to extragalactic foregrounds to the estimated lensing power spectrum.

4.1. Microwave-sky Simulations

Given the multiple complex astrophysical processes that distort or contaminate the observed CMB, we start by predicting biases in the estimated lensing power spectrum from cosmological simulations that attempt to model the nonlinear, correlated fields that constitute extragalactic foregrounds. We use two such simulations: the WEBSKY simulations (Stein et al. 2020), and the simulations from Sehgal et al. (2010), which we will refer to as the S10 simulations.

The WEBSKY simulation is built on a full-sky dark matter halo catalog based on a 600 $(\mathrm{Gpc}/h)^3$ mock matter field generated using the fast "mass-Peak Patch" algorithm (Stein et al. 2019). CIB is generated using a halo model with parameters fit to Herschel power spectrum measurements (see Shang et al. 2012; Viero et al. 2013 for details). Gas is distributed around halos by assuming spherically symmetric, parametric radial distributions from Battaglia et al. (2012), which are then used to model the kSZ effect and the Compton-y field. Maps of radio emission from galaxies are also generated using a halo model, as described in Li et al. (2022).

The S10 simulation is built on a 1 h^{-1} Gpc N-body simulation, with lensing quantities calculated via ray-tracing. Gas is added to the dark matter halos without assuming spherical symmetry (following Bode et al. 2007), allowing for additional complexity in the SZ observables relative to WEBSKY. Radio sources and infrared galaxies are added using halo occupation distributions tuned to a variety of multi-wavelength observations (see Sehgal et al. 2010 for details).

4.2. Simulated ACT DR6 Maps

In order to generate realistic estimates of the contamination of the ACT DR6 lensing power spectrum due to extragalactic foregrounds, we need to perform some preprocessing of the microwave-sky simulations described in Section 4.1, such that they have some of the same observational properties as the real ACT data.

1. We generate (in differential CMB temperature units) a total foreground map (for each frequency) by summing the contributions from the tSZ, kSZ, CIB, and radio point sources. We neglect the dependence of the foreground components across the ACT passbands (since we only have simulation outputs at a small number of discrete single frequencies) and simply choose the simulated frequency closest in frequency to 90 and 150 GHz (93 and 145 GHz for WEBSKY, 90 and 148 GHz for S10). This is a reasonable approach given that there is theoretical uncertainty in the foreground components studied here that is larger than the $\mathcal{O}(10\%)$ errors induced by neglecting the passbands. Fortunately, as we will show, the residual biases (as a fraction of the lensing signal) after foreground mitigation are at the percent level; hence, the dependence across passbands is only relevant at the $\mathcal{O}(0.1\%)$, well below our current statistical precision.
2. We then add a realization of the CMB, lensed by the appropriate κ field.
3. We then convolve with a Gaussian beam appropriate for the frequency (with FWHM $= 2\farcm2$ and $1\farcm4$ for 90 and 150 GHz, respectively), convert the maps to the *plate carrée* cylindrical (CAR) pixelization used for ACT data and pipelines, and add a random realization of the instrumental noise, approximated as being Gaussian with the spatially varying inverse variance estimated for the DR6 day+night coadded data[52] (see S. Naess et al. 2024, in preparation).
4. We input these simulated maps to the matched-filter source- and cluster-finding algorithms implemented by NEMO. For each frequency, we initially run NEMO in point-source-finding mode with an S/N > 4 threshold, which outputs a point-source catalog that can be used to generate a point-source model map.
5. We then run NEMO in cluster-finding mode, jointly on the 90 and 150 GHz maps. The matched-filter templates are a range of cluster model templates, in this case based on the universal pressure profile of Arnaud et al. (2010).[53] A total of 15 filters with different angular sizes are constructed by varying the mass and redshift of the cluster model (see Hilton et al. 2021 for details). The point-source catalogs from the previous steps are used to mask point sources during the cluster finding. This outputs a catalog of candidate clusters with estimated S/N and best-fit cluster model template (with the best-fit redshift and halo mass simply those of the best-fit template).

For the baseline DR6 lensing analysis, models for both the point-source and tSZ cluster contribution are subtracted at the map level. We therefore perform this model subtraction on the foreground-only maps that are required for the foreground bias estimator used here (described in Section 2.2).

We present in the appendices two variations on this procedure to address concerns about its accuracy. First, for both simulations, we have limited volume, so the bias estimates we obtain may contain significant statistical uncertainty due to the limited area of the foreground-only maps used, especially after applying the $f_{\rm sky} \sim 0.3$ DR6 analysis mask. In Appendix C we present foreground bias estimates from a version of the WEBSKY simulations rotated such that a nonoverlapping region of the simulation populates the DR6 mask. The foreground biases are consistent, implying that cosmic variance is not an important contributor to uncertainty in the foreground bias estimates presented here.

Second, we note that the noise model assumed above is rather simplified; while it probably includes the small-scale noise important for cluster and source finding fairly well (and note that it is only at these steps in our foreground bias estimation that the noise fields are used), it does lack larger-scale noise correlations due to the atmosphere that may be important, e.g., when it comes to detecting large, low-redshift clusters.[54] We demonstrate in Appendix B that our bias results are insensitive to using a noise model that includes large-scale correlations.

We note here also that while cluster and point-source model subtraction should mitigate the tSZ effect and CIB, respectively, we do not implement a specific mitigation scheme for the kSZ effect. However, it is included in our simulations, so it will contribute to the foreground bias estimates presented here, and it is not expected to generate a significant bias at ACT DR6 accuracy levels (Ferraro & Hill 2018; Cai et al. 2022).

4.3. Adding Planck High-frequency Data and Harmonic ILC

In addition to our baseline approach, which uses only ACT maps to estimate the lensing (specifically the f090 and f150 data), we also investigate a frequency cleaning approach where we include high-frequency data from Planck (353 and 545 GHz), where the CIB is brighter. To reduce the information on the primary CMB coming from Planck (i.e., to make this a lensing measurement that is mostly independent of Planck CMB information), we do not include lower-frequency Planck data.

To simulate the Planck data, we generate the total foreground-only maps at the Planck frequencies and apply a point-source flux threshold of 304 and 555 mJy to the 353 and 545 GHz channels, respectively (taking these thresholds from Table 1 of Planck Collaboration et al. 2016). We apply the ACT DR6 analysis mask to isolate the same sky area for the ACT and Planck data. Note that for the S10 simulations no 545 GHz data were generated; therefore, we simply scale the 353 GHz CIB simulation using an SED (Madhavacheril et al. 2020) defined as

$$f_{\rm CIB}(\nu) \propto \frac{\nu^{3+\beta}}{\exp h\nu/k_{\rm B}T_{\rm CIB} - 1}\left(\frac{dB(\nu, T)}{dT}\bigg|_{T_{\rm CMB}}\right)^{-1}, \quad (19)$$

where $\beta = 1.2$, $T_{\rm CIB} = 24$ K, and $B(\nu, T)$ is the Planck function. Note that since the CIB does not exactly follow such an l-independent spectrum, the use of this rescaling could

[52] Note that this is the noise level of the maps used for source and cluster finding, while the DR6 lensing analysis uses only nighttime observations, which have higher noise. However, the method used here for calculating foreground biases does not require simulating the latter noise level.

[53] We note here that this is not the same theoretical model for the tSZ profile used in constructing the WEBSKY simulations (which used Battaglia et al. 2012); the fact that we find that tSZ biases in WEBSKY are well under control encourages us that our foreground biases are not very sensitive to the theoretical cluster profile assumed when detecting and subtracting models for clusters.

[54] There has been extensive recent work on realistic noise simulations for single-array maps by Atkins et al. (2023); however, we did not have access to realistic noise simulations for the deep coadds used for source and cluster detection during this work. Hence, we perform the sensitivity test outlined in Appendix B.

underestimate foreground biases in tests of CIB deprojection on the S10 simulations; however, the WEBSKY simulations do have a more realistic 545 GHz channel. As for the other foreground components, the tSZ effect can be accurately scaled to 545 GHz with SED

$$f_{\rm tSZ} \propto \frac{h\nu}{k_{\rm B}T_{\rm CMB}}\frac{e^{h\nu/k_{\rm B}T_{\rm CMB}}+1}{e^{h\nu/k_{\rm B}T_{\rm CMB}}-1} - 4, \quad (20)$$

where kSZ is frequency independent in CMB units, and radio point sources can be assumed to have negligible flux contribution at 545 GHz (Dunkley et al. 2013).

We produce a harmonic-space *internal linear combination* (ILC; see, e.g., Remazeilles et al. 2011; Madhavacheril et al. 2020) of the maps that minimizes the variance of the output maps and additionally deprojects one or more foreground components by assuming an l-independent SED for those foreground components.[55] We use the same point-source and tSZ cluster-subtracted ACT maps as in our baseline analysis, while the Planck maps are only point source subtracted. We investigate below the deprojection of the CIB and/or the tSZ effect, with SEDs given by Equations (19) and (20), respectively. Again, we note that the assumption of an l-independent SED is only approximate in the case of the CIB. Minimizing the variance requires providing total power and cross-spectra for all the input maps; we construct these for frequencies i and j as

$$C^{\rm total}_{l,ij} = C^{\rm cmb}_l + C^{\rm fg}_{l,ij} + \delta_{ij}N_{l,i}, \quad (21)$$

where $C^{\rm fg}_{l,ij}$ is the foreground (cross-)spectrum estimated from the simulation foreground-only maps and $N_{l,\rm i}$ is the noise-only power spectrum, estimated from ACT DR6 noise simulations and the Planck noise-only simulations[56] provided with the NPIPE maps (Planck Collaboration et al. 2020b). Note that since we do not use signal power spectra measured directly from simulated maps containing lensed CMB signal, we will not incur any *ILC bias* (Delabrouille et al. 2009), which comes from the down-weighting of modes with high signal variance that can occur in an ILC. When generating ILC maps on the real DR6 data, we do use power spectra measured directly from the data, but we apply aggressive smoothing to avoid this ILC bias.[57]

5. Simulation Results: Foreground Biases to the Lensing Power Spectrum

Below we present our simulation-derived predicted biases due to extragalactic foregrounds, $\Delta C_L^{\hat\phi\hat\phi}$. We calculate this following the methodology of Schaan & Ferraro (2019), who outline how each of the terms primary bispectrum, secondary bispectrum, and trispectrum can be estimated from simulations without noise due to the CMB or instrument noise realization. The total bias is then simply the sum of these terms (see Equation (6)).

We estimate foreground biases for the temperature-only case, as well as for our MV estimator, which is the power spectrum of $\hat\phi^{\rm MV}(\boldsymbol{L})$, a lensing reconstruction that uses both temperature and polarization data, given by a weighted sum of the reconstructions from all pairs $AB \in [TT, TE, TB, EE, EB]$:

$$\hat\phi^{\rm MV}(\boldsymbol{L}) = \sum_{AB} w_{AB}\hat\phi^{AB}(\boldsymbol{L}) \quad (22)$$

(see Qu et al. 2024 for the exact details of these weights in the bias-hardened case).

We weight each $\hat\phi^{AB}(\boldsymbol{L})$ by the inverse of the normalization (which is equivalent to its response to lensing). The bias on the MV lensing power spectrum estimate is then (Darwish et al. 2021)

$$\begin{aligned}\Delta C_L^{\hat\phi,\rm MV} = [&(w_L^{TT})^2\Delta C_L^{\hat\phi,TTTT} \\ &+ \frac{1}{2}\sum_{XY\in(EE,BB,TE)} w_L^{TT}w_L^{XY}P_L^{TT} \\ &+ w_L^{TT}w_L^{TE}S_L^{TTTE} + w_L^{TE}w_L^{TE}S_L^{TETE} \\ &+ w_L^{TT}w_L^{TB}S_L^{TTTB} + w_L^{TB}w_L^{TB}S_L^{TBTB} \\ &+ w_L^{TE}w_L^{TB}S_L^{TETB} + w_L^{TB}w_L^{TB}S_L^{TBTB}], \end{aligned} \quad (23)$$

where $\Delta C_L^{\hat\phi,TTTT}$ is the bias on the temperature-only estimator. P_L^{TT} is the primary bispectrum bias on the temperature-only estimator and is the dominant additional bias term for the MV estimator, coming from the correlation of foreground temperature with the lensing convergence reconstructed from polarization, $\langle Q(T^{\rm fg}, T^{\rm fg})\hat\kappa^{\rm pol}\rangle$. There are additional secondary bispectrum terms S^{ABCD} of the form

$$\langle Q(A^{\rm CMB}, B^{\rm fg})Q(C^{\rm CMB}, D^{\rm fg})\rangle, \quad (24)$$

with $A, B, C, D \in [T, E, B]$. Even the largest of these, S^{TTTE}, is negligible for our analysis.

We also show the bias $\Delta A_{\rm lens}$ in the inferred lensing power spectrum amplitude $A_{\rm lens}$, which we approximate as

$$\Delta A_{\rm lens} = \frac{\sum_L \sigma_L^{-2}\Delta C_L^{\hat\phi\hat\phi}C_L^{\phi\phi}}{\sum_L (C_L^{\phi\phi}/\sigma_L)^2} \quad (25)$$

for the case of a diagonal covariance matrix on $C_L^{\hat\phi\hat\phi}$ with diagonal elements σ_L and true signal $C_L^{\phi\phi}$. The uncertainty on $A_{\rm lens}$ is given by $(\sum_L (C_L^{\phi\phi}/\sigma_L)^2)^{-0.5}$, so

$$\frac{\Delta A_{\rm lens}}{\sigma(A_{\rm lens})} = \frac{\sum_L \sigma_L^{-2}\Delta C_L^{\hat\phi\hat\phi}C_L^{\phi\phi}}{\sqrt{\sum_L (C_L^{\phi\phi}/\sigma_L)^2}}. \quad (26)$$

In this work, we calculate $\sigma(A_{\rm lens})$ based on an approximate covariance (assuming a reconstruction noise based on the analytic full-sky N^0, scaled appropriately for the ACT sky fraction); this allows us to quickly recompute the covariance for the various different mitigation strategies explored here. This results in an underestimation of $\sigma(A_{\rm lens})$ by $\sim$10% for the temperature-only case and $\sim$5% for the MV case, relative to the more accurate uncertainty for the DR6 data, recovered using simulations in Qu et al. (2024), which we will call $\sigma^{\rm data}(A_{\rm lens})$. Given computational limitations, we do not have $\sigma^{\rm data}(A_{\rm lens})$ for most of the analysis variations tested here. Therefore, when we report $A_{\rm lens}$ uncertainties in Table 1, rather than quoting these

[55] A more advanced ILC analysis of the ACT DR6 maps, using needlet methods, is also underway and will be reported in W. Coulton et al. (2024, in preparation).
[56] See https://wiki.cosmos.esa.int/planck-legacy-archive/index.php/NPIPE_Introduction for documentation.
[57] See W. Coulton et al. (2024, in preparation) for further investigation of and mitigation methods for ILC bias.

Table 1
The Bias and Uncertainty on the Inferred Lensing Power Spectrum Amplitude, $A_{\rm lens}$, Predicted from the WEBSKY and S10 Simulations, for the *TT* and MV Estimators, for a Range of Analysis Variations

Simulation	Analysis Version	$\Delta A_{\rm lens}^{TT} \times 100$	$\sigma^{TT}(A_{\rm lens}) \times 100$	$\Delta A_{\rm lens}^{\rm MV} \times 100$	$\sigma^{\rm MV}(A_{\rm lens}) \times 100$
websky	baseline	−1.2	3.7	−0.42	2.3
websky	CIB-deproj incl. Planck	−0.83	3.8	−0.32	2.3
S10	baseline	−0.24	3.7	−0.088	2.3
S10	CIB-deproj incl. Planck	−0.41	3.8	−0.13	2.3
websky	tSZ-deproj + PSH	−1.9	32	−0.12	4.2
websky	D21-tSZ-deproj + PSH	1.4	7.4	0.56	2.3
S10	tSZ-deproj + PSH	0.19	25	−0.024	4.1
S10	D21-tSZ-deproj + PSH	1.2	6.7	0.45	2.3
websky	CIB-deproj incl. Planck + PRH	−0.81	3.8	−0.31	2.3
websky	D21-CIB-deproj incl. Planck + PRH	0.61	3.7	0.2	2.3
S10	CIB-deproj incl. Planck + PRH	−0.49	3.8	−0.17	2.3
S10	D21-CIB-deproj incl. Planck + PRH	0.54	3.7	0.18	2.3
websky	tSZ and CIB-deproj incl. Planck + PRH	8.7	12	0.53	3.6
websky	D21 tSZ and CIB-deproj incl. Planck + PRH	0.63	6.2	0.22	2.3
S10	tSZ and CIB-deproj incl. Planck + PRH	1.5	11	0.098	3.5
S10	D21 tSZ and CIB-deproj incl. Planck + PRH	0.24	5.9	0.09	2.3

Note. Note that in the unbiased case $A_{\rm lens} = 1$, so $\Delta A_{\rm lens}$ constitutes a fractional bias. "PSH" and "PRH" indicated point-source hardeneing and profile hardening, respectively (see Section 2.3 for details).

underestimated uncertainties, we scale them by the ratio $\sigma^{\rm data}(A_{\rm lens})/\sigma(A_{\rm lens})$ for the baseline analysis, for which we have both these approximate uncertainties and the accurate simulation-based estimates.

5.1. Baseline Analysis

In the baseline power spectrum analysis of Qu et al. (2024) and Madhavacheril et al. (2024), we perform lensing reconstruction on a weighted coadd of the f090 and f150 maps, with each frequency i weighted by the inverse of its one-dimensional harmonic-space noise power spectrum N_l^i as estimated from simulations. The inverse noise filter used in the quadratic estimator is diagonal in harmonic space and uses the total power spectrum $C_l^{\rm tot} = C_l^{\rm cmb} + N_l^{\rm coadd}$, the sum of a fiducial CMB power spectrum and the coadd noise power spectrum.

We perform these same steps on the simulated foreground-only maps used to estimate foreground biases here, i.e., we use the same l-dependent weights to coadd the f090 and f150 maps and the same filtering of the maps entering the quadratic estimator (rather than using the total power spectrum of the foreground-only maps themselves, e.g., since the aim here is to use the same weighting of modes as used in the real data reconstruction).

The solid lines in Figure 1 show the predicted bias to the lensing reconstruction power spectrum from temperature data only (left panel) and for the MV estimator (right panel), as a fraction of the expected $C_L^{\kappa\kappa}$ signal, for the baseline method. For both the WEBSKY (purple lines) and S10 (green lines) simulations, the absolute fractional biases are within 2% up to $L \approx 500$, which is where most of the S/N of the DR6 measurement will come from (the solid light-gray line indicates the cumulative S/N as a function of the maximum L included). For WEBSKY the fractional bias does start to exceed that level at higher L and supports our pre-unblinding decision to limit the baseline analysis to $L \leqq 763$. For guidance, the dotted gray line indicates the DR6 1σ uncertainty divided by 10, indicating that biases for each bandpower are mostly below 0.1σ, except for WEBSKY at $L > 500$, where they are still well below 0.5σ.

We emphasize that the result for the MV estimator is most important, since that is what we use for cosmological inference in Qu et al. (2024) and Madhavacheril et al. (2024). However, it is encouraging that biases in the temperature-only estimator are also small, since it allows us to use the consistency of the MV and temperature-only measurements as a test of other systematics that affect temperature and polarization differently.

To more concisely quantify the cumulative impact of these scale-dependent fractional biases, Figure 2 shows the bias in inferred $A_{\rm lens}$ as a function of the maximum L used, $L_{\rm max}$. For our baseline range of $L_{\rm max} = 763$, we report in Table 1 and the legend to Figure 1 the bias in inferred $A_{\rm lens}$, which is well below 1σ for both simulations, for both the temperature-only (-0.31σ for WEBSKY and -0.06σ for S10) and MV cases (-0.18σ for WEBSKY and -0.04σ for S10).

When including higher-L scales, the predicted absolute biases are still modest, not exceeding 0.3σ when including up to $L_{\rm max} = 1300$. We therefore believe that it is reasonable for the lensing power spectrum analyses in Qu et al. (2024) and Madhavacheril et al. (2024) to also consider the extended range, $L_{\rm max} = 1300$.

To provide insight into the impact of each of our mitigation strategies, Figure 3 shows the foreground biases predicted from WEBSKY when no mitigation is used (purple solid lines), when cluster model subtraction is included (green solid lines), when source subtraction is included (orange solid lines), and when both cluster models and sources are subtracted (black solid lines). For each of these cases we also show as dashed lines the case where the profile-hardened estimator is used, as in our baseline analysis. Without profile hardening, both cluster and model subtraction are required to get the biases down to the few percent level. Using profile hardening in addition enables us to reduces foreground bias of $A_{\rm lens}$ to below 1%.

One might be concerned that the profile assumed for profile hardening is derived from the WEBSKY simulation tSZ field,

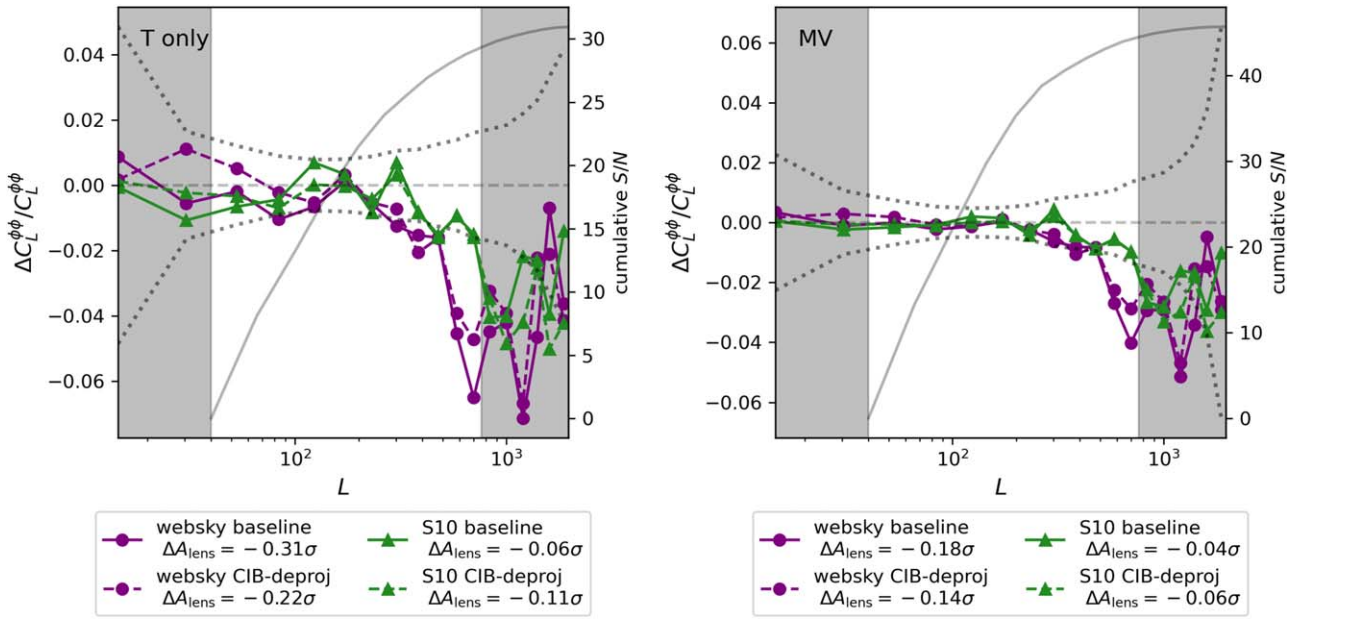


Figure 1. Fractional bias due to extragalactic foregrounds to the estimated CMB lensing power spectrum, for an ACT DR6-like analysis. Left panel: the bias for the temperature-only power spectrum. Right panel: the bias for the MV power spectrum (see Equation (23)), which is the measurement used for cosmological inference in Qu et al. (2024) and Madhavacheril et al. (2024). In both panels, purple circles with solid connecting lines indicate biases estimated from the WEBSKY (Stein et al. 2020) simulations, and green triangles with solid connecting lines indicate biases estimated from the S10 (Sehgal et al. 2010) simulations. Dashed lines indicate the foreground biases for the CIB-deprojected analysis variation described in Section 5.2. The gray dotted line indicates the 1σ uncertainty of the DR6 bandpower measurement divided by 10, and the solid gray line indicates the cumulative S/N when only scales up to L are included. The gray shaded regions indicate scales not used in the cosmological inference, as described in Qu et al. (2024). For all cases, we quote in the legend the total bias in inferred lensing power spectrum amplitude $A_{\rm lens}$ in units of the DR6 1σ uncertainty on that quantity. Dashed gray lines indicate zero bias.

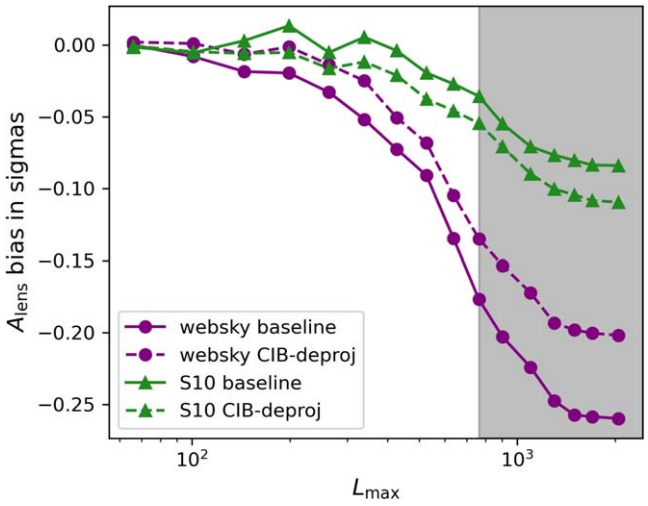


Figure 2. The key result of our simulation tests—the bias in inferred lensing power spectrum, $A_{\rm lens}$, in units of the 1σ uncertainty, as a function of the maximum scale, $L_{\rm max}$ (for the MV estimator). Purple circles and solid (dashed) lines show the prediction from WEBSKY for the baseline (CIB-deprojected analysis), while green triangles show the predictions for the S10 simulations. The gray shaded regions show ranges not included in the baseline cosmology analysis in Qu et al. (2024) and Madhavacheril et al. (2024).

and thus we may overestimate the effectiveness of profile hardening, relative to the performance on the data (where we also use the same profile derived from WEBSKY, which may not match the data as perfectly). We find, however, that the performance of profile hardening is very insensitive to the precise profile assumed; for example, when using a profile based on the S10 tSZ power spectrum instead, we find that the predicted $A_{\rm lens}$ for WEBSKY differs by only 0.3%.

5.2. CIB-deprojected Analysis

In addition to the baseline analysis, we perform the lensing reconstruction on CIB-deprojected maps, using the methodology described in Section 4.3. We include the Planck 353 and 545 GHz channels, in which the CIB has much higher amplitude than at ACT frequencies. In approximate terms, these high-frequency channels provide maps of the CIB that are "subtracted" out by the CIB deprojection, while providing very little information on the CMB. Thus, this analysis is still largely independent of Planck CMB information (it is to maintain this that we do not use the Planck 217 GHz channel). Note that for the ACT maps point sources and tSZ clusters are still modeled and subtracted, as in our baseline approach, while for the Planck data only point sources are subtracted. We still also use the same bias-hardened quadratic estimator as in our baseline approach.

The dashed lines in Figure 1 show the fractional biases to $C_L^{\hat{\kappa}\hat{\kappa}}$ predicted from the WEBSKY (purple dashed) and S10 (green dashed) simulations. For both simulations the predicted biases are quite similar to the baseline case, suggesting that this CIB deprojection approach is also a useful option to use on the DR6 data.

5.3. Other Options

We report in Table 1 statistics on the $A_{\rm lens}$ bias for various other mitigation strategies, summarized below:

1. tSZ deprojection with only the two ACT channels used here (f090 and f150) greatly increases the reconstruction noise (by a factor of $\sim$10 for the temperature-only estimator). This is because the tSZ amplitude at 150 GHz is roughly half that at 90 GHz; therefore, to null the tSZ

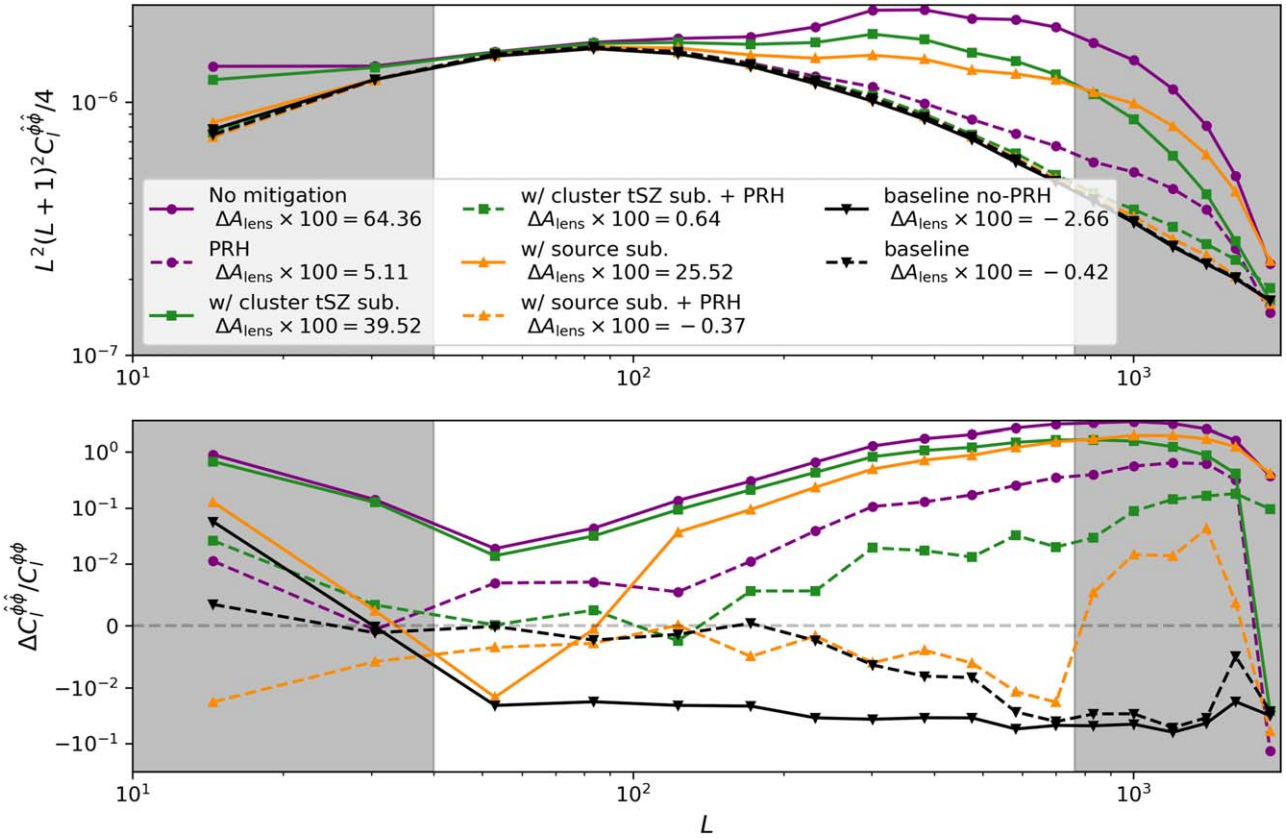


Figure 3. A demonstration of our three main mitigation strategies at work: the total recovered lensing power spectrum (top panel), and its fractional bias with respect to the truth (bottom panel), for the temperature-only case, as predicted from the WEBSKY simulation. Purple lines labeled "no mitigation" show the case where no foreground mitigation is applied. Green (orange) lines have cluster (point-source) models subtracted, while black solid lines have both subtracted. Dashed lines indicated the use of a profile-hardened lensing estimator (labeled "PRH"; see Sections 2.3 and 3 for details). In the legend we also note the fractional bias in lensing power spectrum amplitude, $\Delta A_{\rm lens}$, in percent. Dashed gray lines indicate zero bias. The gray shaded regions indicate scales not used in the cosmological inference, as described in Qu et al. (2024).

effect requires weighting the noisier 150 GHz with roughly twice the weight as the 90 GHz data. In addition, tSZ deprojection upweights the CIB (which is stronger at 150GHz), resulting in $A_{\rm lens}$ biases at the percent level. For these reasons, simply using tSZ-deprojected maps, with only ACT channels, in both legs of the quadratic estimator is not a viable option. The noise cost is reduced for the D21 estimator but is still a factor of $\sim$2 for the temperature estimator. The biases are slightly smaller for these tSZ-deprojected cases when performing point-source hardening (indicated in Table 1 by "PSH") rather than profile hardening (indicated in Table 1 by "PRH"), presumably since it is better suited to the dusty galaxies responsible for the CIB. When performing tSZ deprojection and not also deprojecting the CIB, including high-frequency data from Planck is not useful because the CIB contamination from these high frequencies becomes very large. Hence, we do not show results for that option here.

2. CIB deprojection using only ACT channels similarly has a large noise cost; including Planck channels solves this by effectively providing a relatively high S/N CIB map to subtract (as described in Section 5.2). As well as our CIB deprojection option, where both temperature maps in the quadratic estimator are cleaned, we apply the D21 estimator for this case. For this case, labeled "websky/S10 D21-CIB-deproj incl. Planck + PRH," percent-level $A_{\rm lens}$ biases remain. On inspecting the contributions to this bias, we find that this is due to an increased trispectrum term relative to the fully cleaned estimator (the primary and secondary terms are approximately unchanged). Since we would not expect the tSZ or CIB trispectra to increase for the D21 estimator relative to the fully CIB-cleaned estimator, this is likely due to the presence of terms of the form $\langle Q(T^{\rm CIB}T^{\rm tSZ})Q(T^{\rm CIB}T^{\rm tSZ})\rangle$ for the D21 estimator (this term is approximately zero for the usual CIB-deprojected estimator where all legs have an approximate CIB spectrum deprojected).
3. When including the Planck high-frequency data, we have sufficient degrees of freedom to deproject both tSZ and CIB. However, there is a large noise cost to doing so, resulting in an $A_{\rm lens}$ uncertainty $\sim$5 times larger than that of the baseline analysis (and $\sim$2 times larger for the D21 estimator) for the temperature-only estimator.

6. Data Nulls

Having settled on our mitigation strategies based on the predicted biases from WEBSKY and S10 simulations (Section 5), we now turn to the DR6 data to further validate the performance of these strategies. In Section 6.1 we present null tests involving differences of ACT single-frequency information (in which the CMB lensing signal is nulled), then in Section 6.2 we compare the DR6 bandpowers estimated from a CIB-deprojected map to our baseline DR6 result.

6.1. Frequency Difference Tests

The two ACT channels we have used have different sensitivities to foregrounds, in particular, the tSZ effect has a higher amplitude at 90 GHz than at 150 GHz, and the opposite is true for the CIB. We can use differences between the single-frequency data to form null tests, since the lensed CMB signal is nulled in these differences, while foregrounds are not. If our mitigation strategies are working well, however, we will still find that our lensing estimators applied to these foreground-only maps return null signals. We consider three such null tests in the following that have somewhat different sensitivity to the different terms in the foreground expansion of Equation (6).

6.1.1. Null Map Auto Spectrum

We perform reconstruction on the difference of the individual frequency maps (using temperature data only to maximize sensitivity to extragalactic foregrounds) and take the power spectrum:

$$C_L^{\rm null,1} = \langle Q(T^{90} - T^{150}, T^{90} - T^{150}) \times Q(T^{90} - T^{150}, T^{90} - T^{150})\rangle. \quad (27)$$

Since CMB signal is nulled in the input maps to this reconstruction, this measurement is insensitive to the bispectrum terms and depends only on the trispectrum of the frequency difference map $T^{90} - T^{150} = T^{\rm fg,90} - T^{\rm fg,150}$, where $T^{{\rm fg},i}$ is the foreground contribution to frequency i. The top panel of Figure 4 shows this measurement on DR6 data, showing a null signal, as well as the predictions from the WEBSKY and S10 simulations. The solid gray line indicates the $C_L^{\kappa\kappa}$ theory prediction divided by 10, so any foreground trispectrum hiding beneath the noise here is well below the true lensing signal.

6.1.2. Null Map × $\hat{\kappa}^{\rm MV}$ Spectrum

We cross-correlated the reconstruction based on the frequency difference map with the baseline (i.e., nonnulled) reconstruction, $\hat{\kappa}^{\rm MV}$:

$$C_L^{\rm null,2} = \langle Q(T^{90} - T^{150}, T^{90} - T^{150})\hat{\kappa}^{\rm MV}\rangle. \quad (28)$$

If $\hat{\kappa}^{\rm MV}$ were the true κ, this measurement would be sensitive only to the difference in the primary bispectra contributions for the two frequencies. In fact, since $\hat{\kappa}^{\rm MV}$ will have small foreground biases, there will also be a trispectrum contribution present of the form

$$\langle Q(T^{\rm fg,90} - T^{\rm fg,150}, T^{\rm fg,90} - T^{\rm fg,150}) \times Q(T^{\rm fg,coadd}, T^{\rm fg,coadd})\rangle. \quad (29)$$

The middle panel of Figure 4 shows this measurement on DR6 data, as well as the predictions from the WEBSKY and S10 simulations, showing a null signal. Again, the solid gray line indicates the $C_L^{\kappa\kappa}$ theory prediction divided by 10.

6.1.3. Bandpower Frequency Difference

We take the difference of the auto-spectra of reconstructions performed on single-frequency maps, i.e.,

$$C_L^{\rm null,3} = C_L^{\hat{\kappa}\hat{\kappa},90\ {\rm GHz}} - C_L^{\hat{\kappa}\hat{\kappa},150\ {\rm GHz}}. \quad (30)$$

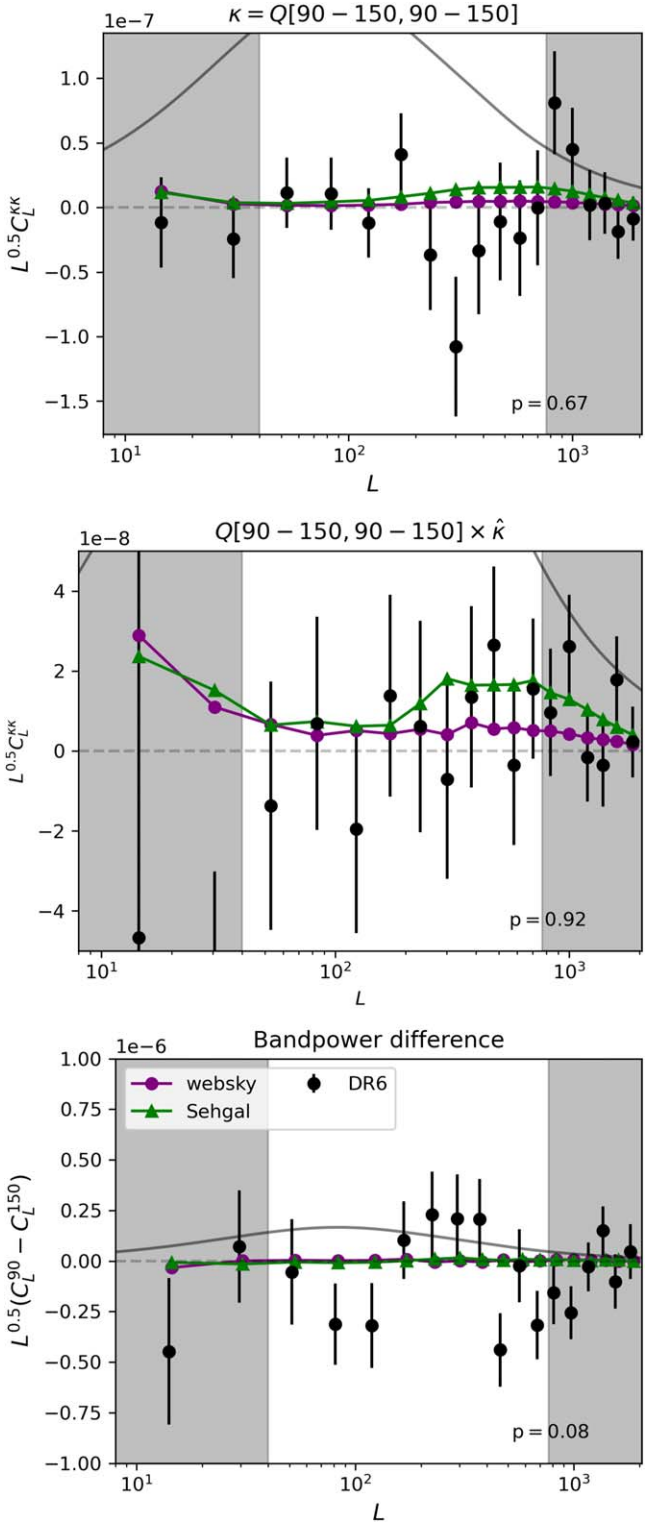


Figure 4. Frequency difference null tests, designed to test the sensitivity of our lensing power spectrum estimation to differences in the foreground contributions to the f090 and f150 temperature data. In each panel, the CMB lensing signal is nulled in a different way; see Sections 6.1.1–6.1.3 for details. All the null tests perform satisfactorily, i.e., the p-values are $\geqslant$0.05. In each panel, the DR6 data are shown as black circles, and the simulation predictions for WEBSKY and S10 are shown as purple circles and green triangles, respectively. Solid gray lines indicate the theoretical lensing power spectrum multiplied by 0.1, and dashed gray lines indicate zero bias. The gray shaded regions indicate scales not used in the cosmological inference, as described in Qu et al. (2024).

This null is sensitive to all three contributions (i.e., primary bispectrum, secondary bispectrum, and trispectrum) in Equation (6), where one would substitute $T^{\rm fg} = T^{\rm fg,90} - T^{\rm fg,150}$ to model the result of this measurement.

The bottom panel of Figure 4 shows this measurement on DR6 data, as well as the predictions from the WEBSKY and S10 simulations, showing a null signal. This test is noisier than the first two, since lensed CMB is not nulled at the map level. We note that one could form other null measurements, for example, those of the form

$$\langle Q(T^{90} - T^{150}, T^{\rm coadd}) Q(T^{90} - T^{150}, T^{\rm coadd}) \rangle, \qquad (31)$$

in order to target and disentangle the secondary bispectrum contribution, but given that we find this to be very small in simulations, we leave such exercises for future work.

6.2. Consistency of CIB-deprojected Analysis

As described in Section 5, the CIB-deprojected version of our analysis performs well on simulations, with predicted systematic bias in the lensing amplitude well below our statistical uncertainties. This prediction of course depends on the simulations, which may contain some inaccuracies in modeling extragalactic foreground components.[58] It is therefore useful to test for consistency of the CIB-deprojected and baseline analysis on the DR6 data. We generate CIB-deprojected temperature maps by combining the DR6 data with Planck 353 and 545 GHz data, using the same procedure as described in Section 4.3. As explained in Qu et al. (2024), we additionally remove a small area at the edge of our mask that has strong features in the high-frequency Planck data due to Galactic dust. We generated 600 simulations of these maps, using the Planck NPIPE noise simulations provided by Planck Collaboration et al. (2020b); these are used for the N^0 and mean-field bias corrections (see Qu et al. 2024 for details).

The deprojected temperature maps are then used in our lensing reconstruction (including profile hardening as in the baseline analysis) and lensing power spectrum estimation (in combination with the same polarization data as are used in the baseline analysis).

Figure 5 compares this CIB-deprojected measurement to the baseline measurement (which does not perform frequency cleaning), finding no evidence for inconsistency. This result implies that CIB contamination is unlikely to be significant in our baseline analysis.

We note that we do not perform an equivalent consistency test for tSZ deprojection; as described in Section 5.3, tSZ deprojection incurs a very large noise cost when including only ACT data and can generate very large biases due to boosting the CIB when including higher frequencies from Planck. While joint tSZ and CIB deprojection could address this large bias, we do not find this to be an effective strategy in our simulation tests since it incurs large noise costs and nonnegligible biases.

7. Conclusions

Extragalactic foregrounds are a potentially significant source of systematic bias in CMB lensing estimation, especially for temperature-dominated current data sets such as ACT. We have

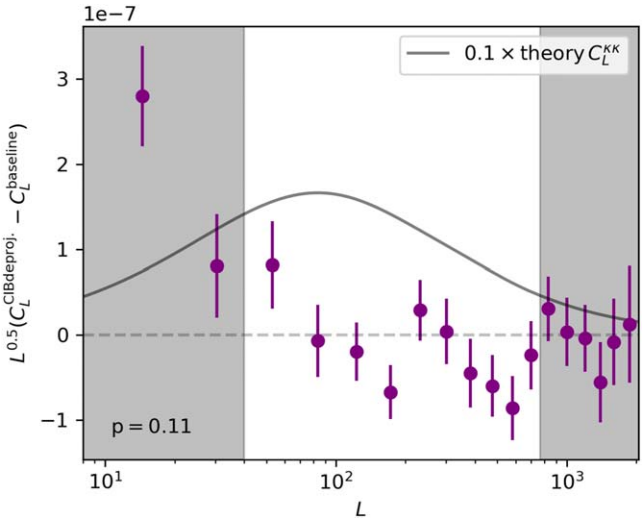


Figure 5. The difference in the ACT DR6 lensing power spectrum estimated from CIB-deprojected maps (as described in Section 6.2) and the baseline measurement. Gray shaded regions indicated L-ranges excluded from the baseline analysis. The p-value = 0.11 for the null result (including only data points within the baseline analysis range). The dashed gray line indicates zero bias, while the solid gray line is the prediced lensing power spectrum divided by 10. The gray shaded regions indicate scales not used in the cosmological inference, as described in Qu et al. (2024).

argued that the mitigation strategies implemented for the DR6 lensing power spectrum analysis (see Madhavacheril et al. 2024; Qu et al. 2024) ensure negligible bias due to extragalactic foregrounds. These mitigation strategies are as follows: (i) finding (using a matched-filter algorithm) and subtracting models for S/N > 4 point sources to remove contamination from radio sources and dusty galaxies (or the CIB), (ii) finding (using a matched-filter algorithm) and subtracting models of galaxy clusters to remove tSZ contamination, and (iii) using a profile bias-hardened quadratic estimator for the lensing reconstruction.

We show first that on two sets of microwave-sky simulations, WEBSKY (Stein et al. 2020) and S10 (Sehgal et al. 2010), the predicted level of bias to the estimated CMB lensing power spectrum is well below our statistical uncertainties. For the baseline analysis, with the MV estimator, the size of the fractional bias is below 1% for most of the fiducial range of scales, L, used; the bias to the inferred lensing amplitude, $A_{\rm lens}$, is below 0.2σ. When extending to higher L, foreground biases become more significant, but the bias to $A_{\rm lens}$ remains at below 0.3σ for $L_{\rm max} = 1300$.

We note here that, in relative terms, the residual biases for the two simulations are somewhat different—for the baseline MV case, we find $\Delta A_{\rm lens} = -0.42$ for WEBSKY and $\Delta A_{\rm lens} = -0.09$ for S10. These residual biases arise from the presence of several sub-percent-level residual biases, which can take either sign (e.g., CIBxtSZ terms can be negative), for each of which there is some theoretical uncertainty. Given this, we argue that it is not that surprising to see differences of 0.5% in $A_{\rm lens}$ biases between the simulations, arising from the sum of several small but uncertain contributions.

While the two simulations use somewhat different methods for implementing astrophysics and they both imply that any residual foreground biases are comfortably within our requirements, one may still be concerned that these two simulations are not an adequate set for validating the DR6 lensing measurement. We cannot rule out the possibility that some

[58] Note that these simulations are typically tuned to real data at the power spectrum level, while the extragalactic foreground biases here depend on higher-order statistics of the foreground fields.

other hypothetical simulation would have larger biases, e.g., if for whatever reason the tSZ effect or CIB was much more non-Gaussian than predicted by these simulations but still had the same power spectrum as websky and S10 (since these power spectra were tuned to observations). Some comparisons of higher-order (than power spectrum) statistics between simulations and data have been made in the literature (e.g., Feng & Holder 2019), but certainly further work on iteratively improving microwave-sky simulations as new data arrive should be a priority going forward. This motivates the data-based tests that we present in Section 6.

We present null tests performed on the DR6 data that leverage the frequency dependence of extragalactic foregrounds and thus do not depend on having realistic microwave-sky simulations. We investigate three "lensing" power spectra where the CMB lensing signal is nulled by differencing the f090 and f150 data at both the map level and the bandpower level, exploiting different sensitivities to the primary bispectrum, secondary bispectrum, and trispectrum foreground components. All of these null tests pass (with p-value $\geqslant 0.05$).

Finally, we demonstrate that using CIB-deprojected maps in our lensing estimation produces lensing power spectrum bandpowers that are consistent with our baseline measurements, implying that CIB contamination is not likely to be a significant contaminant in the DR6 measurement. We note here a further test presented in our companion paper (Qu et al. 2024), which is the consistency with the baseline measurement of the shear estimator of Schaan & Ferraro (2019) and Qu et al. (2023); this estimator uses only the quadrupolar contribution to the CMB mode coupling induced by lensing. While the lensing power spectrum measurement with the shear estimator is somewhat noisier than the baseline measurement, it is very encouraging that Qu et al. (2023) find that the difference in the bandpowers is consistent with zero, with $\Delta A_{\rm lens} = 0.01 \pm 0.05$.

It is worth commenting here on the use of the DR6 lensing reconstruction maps for cross-correlation studies. The foreground bias estimates for the lensing auto-spectrum provided here are not directly applicable to cross-correlations of the lensing reconstructions with, e.g., maps of galaxy overdensity. These cross-correlations are impacted by biases analogous to the primary bispectrum bias described in Section 2.2, due to the correlation between galaxy overdensity and CMB foregrounds that also trace the large-scale structure, especially the CIB and tSZ. The size of the contamination will depend on the specific tracer sample used for cross-correlation, but we do expect the mitigation strategies used here to also be very effective for these cross-correlations, as will be demonstrated for the case of unWISE galaxies in G. Farren et al. (2024, in preparation) and CMASS galaxies in L. Wenzl et al. (2024, in preparation).

While polarization data will become increasingly important for upcoming SO lensing analyses, much of the S/N will still depend on CMB temperature data, so careful treatment of extragalactic foregrounds will be required. With additional frequency channels at high resolution, as will be provided by SO, deprojecting both tSZ and CIB could be more fruitful, including, for example, partial deprojection or composite approaches explored in Abylkairov et al. (2021), Sailer et al. (2021), and Darwish et al. (2023). While deeper upcoming data from, e.g., SO will demand greater control of foreground biases (given the reduced statistical uncertainties), it will also allow fainter point sources, dusty galaxies, and clusters to be detected and modeled or masked, although care must be taken not to introduce selection biases by preferentially masking higher convergence regions of the sky (Lembo et al. 2022).

Acknowledgments

Support for ACT was provided by the U.S. National Science Foundation through awards AST-0408698, AST-0965625, and AST-1440226 for the ACT project, as well as awards PHY-0355328, PHY-0855887, and PHY-1214379. Funding was also provided by Princeton University, the University of Pennsylvania, and a Canada Foundation for Innovation (CFI) award to UBC. ACT operated in the Parque Astronómico Atacama in northern Chile under the auspices of the Agencia Nacional de Investigación y Desarrollo (ANID). The development of multichroic detectors and lenses was supported by NASA grants NNX13AE56G and NNX14AB58G. Detector research at NIST was supported by the NIST Innovations in Measurement Science program.

Computing for ACT was performed using the Princeton Research Computing resources at Princeton University, the National Energy Research Scientific Computing Center (NERSC), and the Niagara supercomputer at the SciNet HPC Consortium.

Some computations were performed on the Niagara supercomputer at the SciNet HPC Consortium and the Symmetry cluster at the Perimeter Institute. SciNet is funded by the CFI under the auspices of Compute Canada, the Government of Ontario, the Ontario Research Fund—Research Excellence, and the University of Toronto.

N.M., B.D.S., F.Q., B.B., I.A.C., G.S.F., and D.H. acknowledge support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 851274). B.D.S. further acknowledges support from an STFC Ernest Rutherford Fellowship. M.M. acknowledges support from NASA grant 21-ATP21-0145. E.C. acknowledges support from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 849169). K.M. acknowledges support from the National Research Foundation of South Africa. O.D. acknowledges support from an SNSF Eccellenza Professorial Fellowship (No. 186879). C.S. acknowledges support from the Agencia Nacional de Investigación y Desarrollo (ANID) through FONDECYT grant No. 11191125 and BASAL project FB210003. I.A.C. acknowledges support from Fundación Mauricio y Carlota Botton. M.H. acknowledges support from the National Research Foundation of South Africa (grant No. 137975). S.N. was supported by a grant from the Simons Foundation (CCA 918271, PBL). N.S. acknowledges support from NSF grant No. AST-1907657. J.R.B. acknowledges support from NSERC and CIFAR and the Canadian Digital Alliance. G.A.M. is part of Fermi Research Alliance LLC under contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. O.D. acknowledges support from an SNSF Eccellenza Professorial Fellowship (No. 186879). J.C.H. acknowledges support from NSF grant AST-2108536, NASA grants 21-ATP21-0129 and 22-ADAP22-0145, D.O.E. grant DE-SC00233966, the Sloan Foundation, and the Simons Foundation. T.N. acknowledges support from JSPS KAKENHI (grant Nos. JP20H05859 and JP22K03682) and World Premier International Research Center Initiative (WPI), MEXT, Japan.

Appendix A
Contribution to Foreground Bias from Individual Foreground Components

We show in Figure 6 the foreground biases for each individual extragalactic foreground component, for WEBSKY (left panel) and S10 (right panel). These are estimated by rerunning the simulation processing described in Section 4.2, but in each case including only a single foreground component in the maps. We also show the total bias as the black circular markers and dashed lines. Note that this is not simply a sum of the individual components since there are additional terms owing to the correlation between the foreground components (e.g., CIB is quite correlated with tSZ).

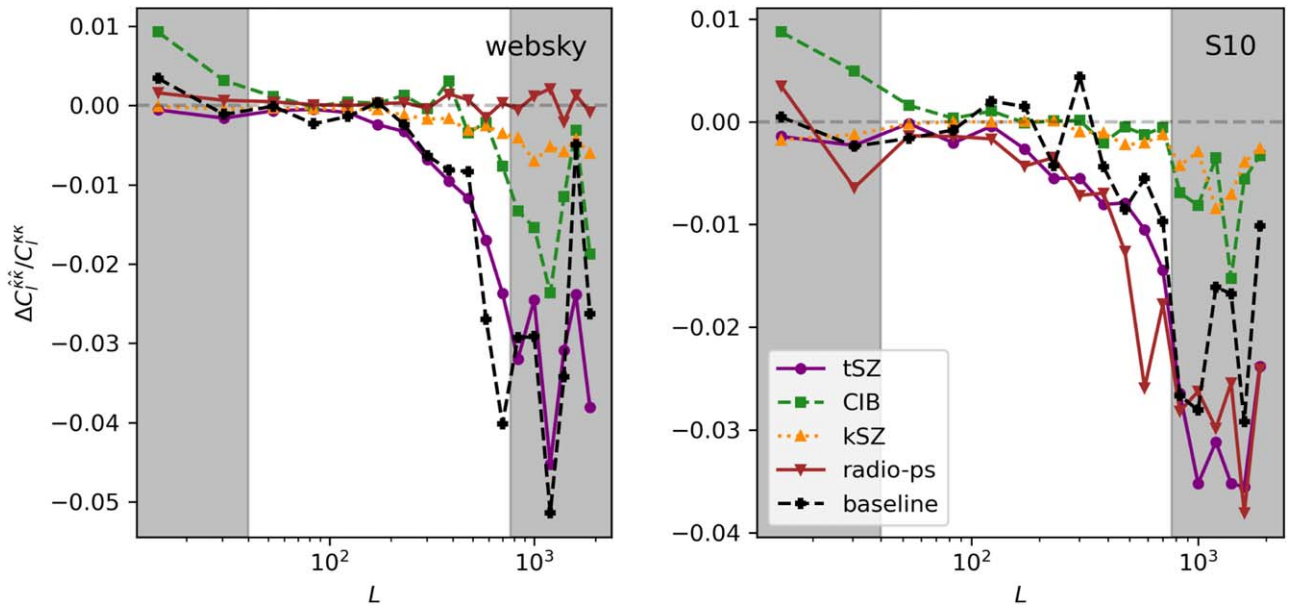


Figure 6. Contributions to the lensing power spectrum bias from individual foreground components: the tSZ effect, the CIB, the kSZ effect, and radio sources ("radio-ps"), as predicted from the WEBSKY (left panel) and S10 (right panel) simulations. The total bias is labeled "baseline" and is not equal to the sum of the individual contributions, since there are also cross-terms present owing to correlations on the sky between the different foregrounds (e.g., CIB is correlated with tSZ).

Appendix B
Results with 1/*f* Modulated Noise

Our simulation-based foreground bias estimates depend on the effectiveness of point-source and cluster detection, which in turn depends on the properties of the noise added to the simulated foreground maps. Above we use a simple, local variance $1/\texttt{ivar}$ model for the map noise, where `ivar` is the inverse variance map estimated for the DR6 coadd data). We test here the inclusion of additional large-scale correlations by generating a simulated noise map as a Gaussian random field drawn from a power spectrum $N_l = (1 + (l/l_{\rm knee})^{\alpha})$, which is then multiplied by $\sqrt{1/\texttt{ivar}}$. On large scales ($l < l_{\rm knee}$), this introduces correlated noise that resembles that expected as a result of the atmosphere, while still achieving the correct pixel variance at small scales ($l \gg l_{\rm knee}$). It is found to be a good fit to ACT data in Naess et al. (2020), from where we take the parameter values $l_{\rm knee} = (2000, 3000)$ and $\alpha = (-3, -3)$ for the (90, 150) GHz channels.

As shown in Figure 7, we do find that the results of the NEMO cluster-finding code are somewhat sensitive to this change in noise model, with fewer clusters found at low redshift (or, at least, assigned low-redshift best-fit templates) and somewhat fewer clusters detected in total. This is expected —the increased noise on large scales reduces the effectiveness in detecting the larger (in angular size) low-redshift clusters (see, e.g., discussion in Hilton et al. 2021). Nonetheless, there is little impact on the resulting foreground biases predicted for the lensing power spectrum, with negligible change to the bias in the infereed $A_{\rm lens}$ (see Figure 8).

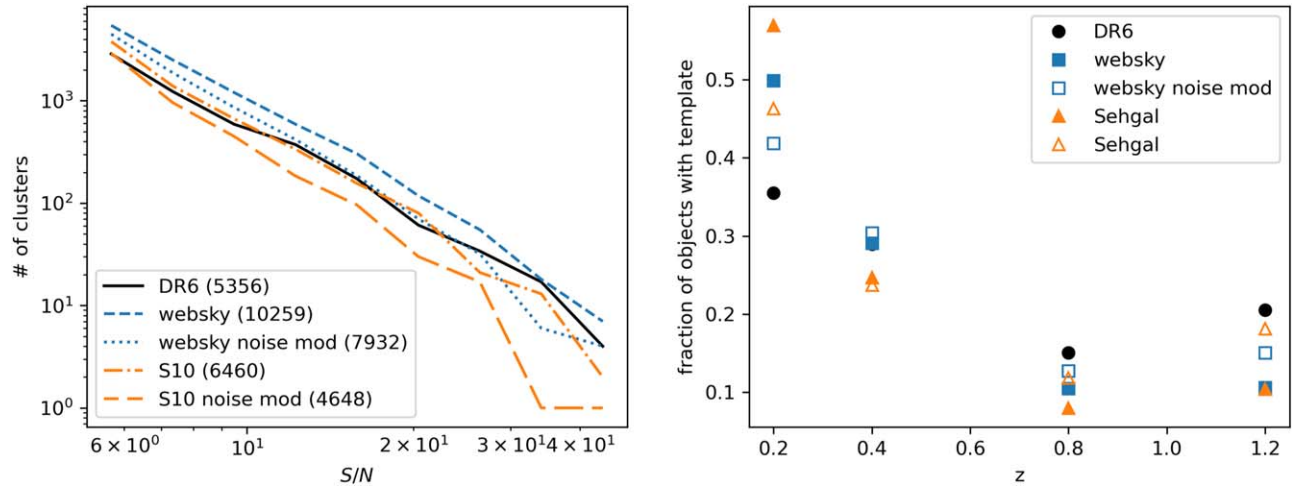


Figure 7. Left: the number of S/N > 5 cluster candidates detected by NEMO as a function of S/N, for the DR6 data, and the WEBSKY and S10 simulations. As well as our fiducial noise model, we show the number counts assuming the modulated $1/f$ noise model described in Appendix B. The total number of S/N > 5 cluster candidates for each case is noted in the legend. Right: fraction of candidate clusters as a function of best-fit template redshift.

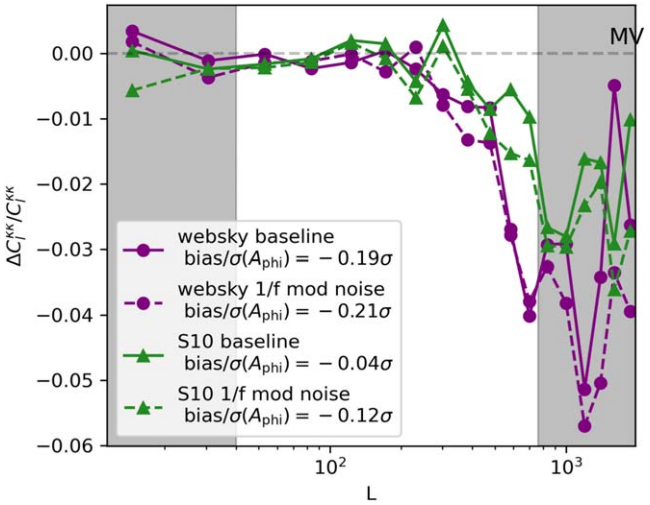


Figure 8. The dashed lines show predicted biases to the MV lensing reconstruction power spectrum, for the WEBSKY and S10 simulations, when assuming the $1/f$ modulated noise model described in Appendix B rather than our fiducial noise model (solid lines). In the legend we note the bias in inferred lensing amplitude as a fraction of the 1σ uncertainty for the ACT DR6 lensing analysis. The gray shaded regions indicate scales not used in the cosmological inference, as described in Qu et al. (2024). Dashed gray lines indicate zero bias.

Appendix C
Uncertainty on Bias Estimates due to Finite Simulation Volume

In Section 5 we present predictions of the foreground biases to $C_L^{\hat{\kappa}\hat{\kappa}}$ based on simulations. Unlike lensing estimation from a normal CMB map, these bias predictions are not affected by instrumental noise or noise on the CMB power spectrum. However, there does exist some uncertainty associated with the finite volume of simulation from which they are estimated. For both the WEBSKY and S10 simulations one full sky is available, and in order to ensure realistic noise properties for cluster and source finding, we further apply the ACT DR6 mask. We generate a close-to-independent (within the ACT mask) realization of the WEBSKY simulation by rotating the simulation maps such that the sky area that enters the ACT DR6 mask does not overlap with that initially entering the ACT mask. A simple rotation by 90° around the y-axis, implemented using pixell's[59] `rotate_alm` function with angle arguments `(0., -np.pi/2, 0.)`, generates a map that has negligible overlap with the original area allowed by the mask.

The green dashed line in Figure 9 shows the foreground-induced bias for the MV case, with the bias for the rotated case well within requirements (an $A_{\rm lens}$ bias of -0.11σ) and similar to the baseline (unrotated case), implying that cosmic variance is not a significant source of uncertainty in our foreground bias predictions.

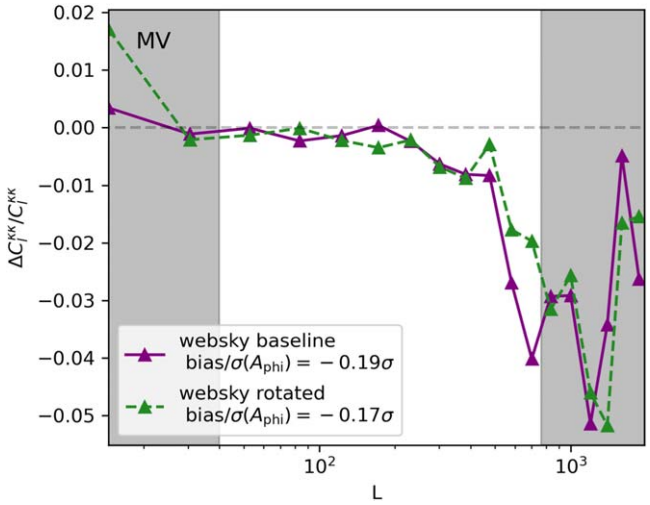


Figure 9. The dashed lines show predicted biases to the lensing reconstruction power spectrum for the rotated (as described in Appendix C) WEBSKY simulations, compared to the baseline case (solid line). The gray shaded regions indicate scales not used in the cosmological inference, as described in Qu et al. (2024).

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

ORCID iDs

Niall MacCrann https://orcid.org/0000-0002-8998-3909
Irene Abril-Cabezas https://orcid.org/0000-0003-3230-4589
Jason E. Austermann https://orcid.org/0000-0002-6338-0069
Nicholas Battaglia https://orcid.org/0000-0001-5846-0411
Elia S. Battistelli https://orcid.org/0000-0001-5210-7625
William R. Coulton https://orcid.org/0000-0002-1297-3673
Adriaan J. Duivenvoorden https://orcid.org/0000-0003-2856-2382
Jo Dunkley https://orcid.org/0000-0002-7450-2586
Gerrit S. Farren https://orcid.org/0000-0001-5704-1127
Simone Ferraro https://orcid.org/0000-0003-4992-7854
Yilun Guan https://orcid.org/0000-0002-1697-3080
Matt Hilton https://orcid.org/0000-0002-8490-8117
Arthur Kosowsky https://orcid.org/0000-0002-3734-331X
Gabriela A. Marques https://orcid.org/0000-0002-8571-8876
Kavilan Moodley https://orcid.org/0000-0001-6606-7142
Sigurd Naess https://orcid.org/0000-0002-4478-7111
Lyman Page https://orcid.org/0000-0002-9828-3525
Emmanuel Schaan https://orcid.org/0000-0002-4619-8927
Neelima Sehgal https://orcid.org/0000-0002-9674-4527
Edward J. Wollack https://orcid.org/0000-0002-7567-4451
Lukas Wenzl https://orcid.org/0000-0001-5245-2058

References

Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016,
Abylkairov, Y. S., Darwish, O., Hill, J. C., & Sherwin, B. D. 2021, PhRvD, 103, 103510
Arnaud, M., Pratt, G. W., Piffaretti, R., et al. 2010, A&A, 517, A92
Atkins, Z., Duivenvoorden, A. J., Coulton, W. R., et al. 2023, JCAP, 2023, 071
Battaglia, N., Bond, J. R., Pfrommer, C., & Sievers, J. L. 2012, ApJ, 758, 75
Beck, D., Errard, J., & Stompor, R. 2020, JCAP, 2020, 030
Bianchini, F., Wu, W. L. K., Ade, P. A. R., et al. 2020, ApJ, 888, 119
Bode, P., Ostriker, J. P., Weller, J., & Shaw, L. 2007, ApJ, 663, 139
Bucher, M., & Louis, T. 2012, MNRAS, 424, 1694
Cai, H., Madhavacheril, M. S., Hill, J. C., & Kosowsky, A. 2022, PhRvD, 105, 043516
Carron, J., Mirmelstein, M., & Lewis, A. 2022, JCAP, 2022, 039
Challinor, A., Allison, R., Carron, J., et al. 2018, JCAP, 2018, 018
Darwish, O., Madhavacheril, M. S., Sherwin, B. D., et al. 2021, MNRAS, 500, 2250
Darwish, O., Sherwin, B. D., Sailer, N., Schaan, E., & Ferraro, S. 2023, PhRvD, 107, 043519
Das, S., Sherwin, B. D., Aguirre, P., et al. 2011, PhRvL, 107, 021301
Delabrouille, J., Cardoso, J. F., Le Jeune, M., et al. 2009, A&A, 493, 835
Dunkley, J., Calabrese, E., Sievers, J., et al. 2013, JCAP, 2013, 025
Feng, C., & Holder, G. 2019, arXiv:1905.02084
Ferraro, S., & Hill, J. C. 2018, PhRvD, 97, 023512
Haehnelt, M. G., & Tegmark, M. 1996, MNRAS, 279, 545
Hanson, D., Challinor, A., Efstathiou, G., & Bielewicz, P. 2011, PhRvD, 83, 043005
Hilton, M., Sifón, C., Naess, S., et al. 2021, ApJS, 253, 3
Hirata, C. M., Ho, S., Padmanabhan, N., Seljak, U., & Bahcall, N. A. 2008, PhRvD, 78, 043520
Hu, W., DeDeo, S., & Vale, C. 2007, NJPh, 9, 441
Hu, W., & Okamoto, T. 2002, ApJ, 574, 566
Lembo, M., Fabbian, G., Carron, J., & Lewis, A. 2022, PhRvD, 106, 023525
Lewis, A., & Challinor, A. 2006, PhR, 429, 1
Li, Z., Puglisi, G., Madhavacheril, M. S., & Alvarez, M. A. 2022, JCAP, 2022, 029
Madhavacheril, M. S., & Hill, J. C. 2018, PhRvD, 98, 023534
Madhavacheril, M. S., Hill, J. C., Næss, S., et al. 2020, PhRvD, 102, 023534
Madhavacheril, M. S., Smith, K. M., Sherwin, B. D., & Naess, S. 2021, JCAP, 2021, 028
Madhavacheril, M., Qu, F., Sherwin, B., et al. 2024, ApJ, 962, 113
Millea, M., Daley, C. M., Chou, T. L., et al. 2021, ApJ, 922, 259
Naess, S., Aiola, S., Austermann, J. E., et al. 2020, JCAP, 2020, 046
Namikawa, T., Hanson, D., & Takahashi, R. 2013, MNRAS, 431, 609
Namikawa, T., & Takahashi, R. 2014, MNRAS, 438, 1507
Omori, Y., Chown, R., Simard, G., et al. 2017, ApJ, 849, 124
Osborne, S. J., Hanson, D., & Doré, O. 2014, JCAP, 2014, 024
Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, A&A, 594, A26
Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020a, A&A, 641, A8
Planck Collaboration, Akrami, Y., Andersen, K. J., et al. 2020b, A&A, 643, A42
Planck Collaboration, Akrami, Y., Ashdown, M., et al. 2020c, A&A, 641, A4
Qu, F. J., Challinor, A., & Sherwin, B. D. 2023, PhRvD, 108, 063518
Qu, F., Sherwin, B., Madhavacheril, M., et al. 2024, ApJ, 962, 112
Remazeilles, M., Delabrouille, J., & Cardoso, J.-F. 2011, MNRAS, 410, 2481
Sailer, N., Ferraro, S., & Schaan, E. 2023, PhRvD, 107, 023504
Sailer, N., Schaan, E., & Ferraro, S. 2020, PhRvD, 102, 063517
Sailer, N., Schaan, E., Ferraro, S., Darwish, O., & Sherwin, B. 2021, PhRvD, 104, 123514
Schaan, E., & Ferraro, S. 2019, PhRvL, 122, 181301
Sehgal, N., Bode, P., Das, S., et al. 2010, ApJ, 709, 920
Shang, C., Haiman, Z., Knox, L., & Oh, S. P. 2012, MNRAS, 421, 2832
Sherwin, B. D., van Engelen, A., Sehgal, N., et al. 2017, PhRvD, 95, 123529
Smith, K. M., Zahn, O., & Doré, O. 2007, PhRvD, 76, 043510
Staniszewski, Z., Ade, P. A. R., Aird, K. A., et al. 2009, ApJ, 701, 32
Stein, G., Alvarez, M. A., & Bond, J. R. 2019, MNRAS, 483, 2236
Stein, G., Alvarez, M. A., Bond, J. R., van Engelen, A., & Battaglia, N. 2020, JCAP, 2020, 012
van Engelen, A., Bhattacharya, S., Sehgal, N., et al. 2014, ApJ, 786, 13
van Engelen, A., Keisler, R., Zahn, O., et al. 2012, ApJ, 756, 142
Viero, M. P., Wang, L., Zemcov, M., et al. 2013, ApJ, 772, 77
Wu, W. L. K., Mocanu, L. M., Ade, P. A. R., et al. 2019, ApJ, 884, 70