

CMB lensing from Planck PR4 maps

Julien Carron,^{a,b} Mark Mirmelstein^b and Antony Lewis^b

^aUniversité de Genève, Département de Physique Théorique et CAP,
24 Quai Ansermet, CH-1211 Genève 4, Switzerland

^bDepartment of Physics & Astronomy, University of Sussex,
Brighton BN1 9QH, U.K.

E-mail: julien.carron@unige.ch, m.mirmelstein@sussex.ac.uk,
antony@cosmologist.info

Received June 24, 2022

Accepted August 9, 2022

Published September 13, 2022

Abstract. We reconstruct the Cosmic Microwave Background (CMB) lensing potential on the latest *Planck* CMB PR4 (NPIPE) maps, which include slightly more data than the 2018 PR3 release, and implement quadratic estimators using more optimal filtering. We increase the reconstruction signal to noise by almost 20%, constraining the amplitude of the CMB-marginalized lensing power spectrum in units of the *Planck* 2018 best-fit to 1.004 ± 0.024 (68% limits), which is the tightest constraint on the CMB lensing power spectrum to date. For a base Λ CDM cosmology we find $\sigma_8 \Omega_m^{0.25} = 0.599 \pm 0.016$ from CMB lensing alone in combination with weak priors and element abundance observations. Combination with baryon acoustic oscillation data gives tight 68% constraints on individual Λ CDM parameters $\sigma_8 = 0.814 \pm 0.016$, $H_0 = 68.1_{-1.1}^{+1.0} \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.313_{-0.016}^{+0.014}$. *Planck* polarized maps alone now constrain the lensing power to 7%.

Keywords: cosmological parameters from CMBR, gravitational lensing

ArXiv ePrint: [2206.07773](https://arxiv.org/abs/2206.07773)

JCAP09(2022)039

Contents

1	Introduction	1
2	Analysis	2
2.1	Data and simulations	2
2.2	Component separation	3
2.3	CMB filtering and quadratic estimators	5
2.4	κ -filtering and lensing spectrum	6
2.5	Covariance matrix	8
2.6	Point-source correction	9
2.7	Likelihoods	11
3	Results	12
3.1	GMV and TT reconstructions	12
3.2	Polarization-only reconstruction	16
3.3	ISW-lensing	18
3.4	Parameter constraints	19
4	Conclusions	22

1 Introduction

The *Planck* 2018 lensing analysis [1, hereafter PL2018] reconstructed the CMB lensing over most of the sky and derived powerful cosmological constraints, both on its own and in combination with other data. While ground-based experiments are rapidly improving the signal to noise of lensing reconstructions on small scales over parts of sky [2–5], the *Planck* data is likely to remain the only full-sky dataset for the near future, and is the only data capable of reconstructing the largest-scale lensing modes. The analysis was sub-optimal in various respects, both due to the limitations of the processed data available at the time, and to various choices made to simplify the lensing analysis. In this paper we present a new analysis of the *Planck* data, with the aim of getting a more-optimal *Planck* lensing reconstruction over most of the sky. By using a different data processing pipeline we can also hope to get a clearer picture of how robust the *Planck* lensing constraints may be to unknown systematics and modelling uncertainties.

The latest *Planck* sky maps come from the NPIPE processing pipeline [6, hereafter NPIPE or PR4], which uses $\sim 8\%$ additional measurement time from the satellite repointing manoeuvres. NPIPE also improved the low-level data processing in many respects, and provides more simulations. In addition to using PR4 data, we also introduce and test a number of improvements to the PL2018 analysis, gaining a scale-dependent improvement of up to $\sim 20\%$ in signal to noise on the lensing power spectrum measurement. These include:

- Joint inverse-variance (Wiener-)filtering of the temperature and polarization CMB maps, instead of the separate filtering performed in PL2018. This uses the expected temperature to E -polarization cross-correlation to improve recovery of the E -mode map input to the quadratic estimator (QE). For *Planck* noise levels the joint lensing reconstruction

is temperature dominated, and this translates into a modest $\sim 3\%$ reduction of the Minimum Variance (MV) bandpower errors. The resulting lensing quadratic estimator is distinct from the optimized combinations of quadratic estimators presented in early work by Okamoto and Hu [7, 8], and was named GMV (Generalized Minimum Variance) by ref. [9]. The same reference shows that the improvement expected for more sensitive experiments is larger, and is expected to reach 10% for *Simons Observatory* [9, 10].

- An additional post-processing step on the quadratic-estimated lensing convergence maps, by Wiener-filtering them according to the local reconstruction noise level on the sky. This step, introduced by ref. [11], has the virtue of weighting different parts of the sky more optimally when building the lensing power spectrum. In the noise-dominated regime, the way the QEs are constructed automatically gives a nearly-optimal inverse-noise weighting [1]. However, on scales where the signal is significant, a more uniform weighting is more optimal, and the extra filtering step gives an improved power spectrum estimate. In our case, this means a improvement of $\sim 7\%$ in GMV bandpower errors close to the peak of the lensing spectrum at $L \sim 30$.
- Accounting for the noise inhomogeneity in our baseline GMV reconstruction by using an inhomogeneous noise variance in the CMB Wiener filter. As demonstrated by [PL2018](#) (but not used in the published likelihoods), this has a large impact on the quality of the polarization-only reconstruction, for which the CMB instrument noise dominates the polarization signal on small scales. Including the inhomogeneous filtering for the GMV estimator gives a total improvement of $\sim 10\%$.
- Use of 600 end-to-end Monte-Carlo simulations (MCs), which is twice as many as was available for [PL2018](#). This allows us to reduce the errors in several terms calibrated with MCs (mean-field, MC- and RD- $N^{(0)}$ -biases and covariance matrix).
- We use a novel way to obtain a realization-dependent covariance matrix, which takes better care of the inhomogeneities of the noise across the sky. This is discussed in section 2.5.

2 Analysis

Section 2.1 introduces the data sets and section 2.2 our foreground cleaning procedure. In section 2.3 we discuss the filtering of the CMB maps and the construction of the quadratic estimators. Section 2.4 describes the additional filtering step we apply to the convergence map estimate. Section 2.5 deals with our covariance matrix estimate, and section 2.6 with the point-source correction we apply to the bandpowers. Finally, section 2.7 discuss the construction of the lensing likelihood.

2.1 Data and simulations

We use NPIPE data maps for all frequency channels in three flavours (full mission maps, as well as the ‘A’ and ‘B’ splits designed to have maximally independent noise and systematics [6]), together with the accompanying set of 600 end-to-end simulations of each map. The CMB sky signal simulation maps are the same as for PR3, and drawn from the *Planck* FFP10 cosmological model.¹ The noise components of PR3 and PR4 are however statistically independent.

¹https://github.com/carronj/plancklens/blob/master/plancklens/data/cls/FFP10_wdipole_params.ini.

A foreground component, based for each channel on the Commander sky model, is present in all simulations. On the main *Planck* High Frequency Instrument CMB channels this contains resolved bright sources, galactic dust together with some level of large-scale CIB, and zodiacal emission. The data processing is intrinsically non-linear, so a perfect separation between the CMB, noise and foregrounds component of the simulations is not possible. The simulations that we are using are therefore not completely independent. However, as discussed in more detail in section 2.2, the common component for the simulations is very small after foreground cleaning, and plays very little role in our analysis.

As in previous releases, there is a mismatch at the level of a few percent between the data and simulation power in temperature at the highest multipoles we use, $\ell \sim 2000$, presumably mainly due to residual foreground power. For this reason, before processing the maps to the lensing analysis, we add a simple Gaussian noise realization with a smooth power spectrum fit to the excess power in the foreground-cleaned simulations. This addition of power does not play an important role, except at the very low and even lensing multipoles where the mean-field is strongest (see figure 6 later on).

Besides using more data, the NPIPE processing introduced a number of modifications, or improvements, to the map-making process which are described in detail in ref. [6]. Given the complexity of the map-making process, it is difficult for us to single-out any one as being particularly critical to the lensing reconstruction. One major change in PR4 concerns large-scale polarization, where the use of template priors improves recovery of the polarized signal at large scales at the cost of a non-trivial transfer function. This is of no relevance to this paper, since low CMB multipoles do not carry much lensing signal at all, and we exclude the lowest CMB multipoles anyways.

2.2 Component separation

Previous *Planck* releases provided foreground-cleaned CMB maps from four independent methods (SMICA, SEVEM, NILC and Commander), described in detail in refs. [12, 13] for the 2015 and 2018 release respectively. However, at the time of writing the only complete set of cleaned PR4 temperature and polarization maps and simulations available were using the SEVEM component separation pipeline. For lensing reconstruction purposes from temperature, all four Planck component separation methods perform comparably well. However, we found that SEVEM performs worse on the 2018 polarized data than the baseline SMICA [14], where the template fitting procedure of SEVEM [15] is less effective than in temperature owing to the increased relative importance of the instrumental noise in polarization. For this reason, we built our own foreground-cleaned NPIPE maps, in the following way. SMICA linearly combines the frequency channels in harmonic space with a set of weights. Using the 2018 SMICA weights in temperature or polarization, with maximum multipole 2500, we simply rescaled them to account for the slight change in the effective transfer functions of each frequency channel, which we obtained empirically by cross-correlating the NPIPE frequency channel maps to the CMB FFP10 simulations input. Doing this could result in reduced suppression of foregrounds and slight suboptimality in the multipole range where both foregrounds and noise are relevant. However, we saw no evidence of a significant issue in the course of this analysis. This is not too surprising since PL2018 showed that without any cleaning, expected biases from CIB contamination are at most 2% on the CIB-bright 217GHz CMB channel without any cleaning, and down to $\sim 0.1\%$ after SMICA weighting.

To build the PR3 temperature maps, the SMICA team applied an $f_{\text{sky}} \sim 2\%$ preprocessing mask to the galactic centre and a few bright sources across the sky, and filled these pixels

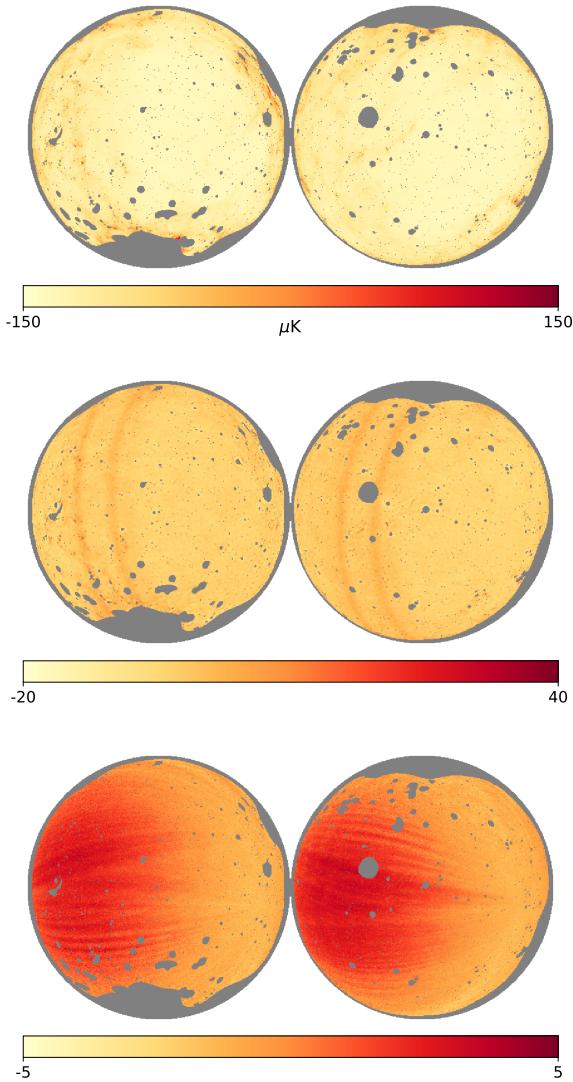


Figure 1. *Top panel:* the foreground model map input to the cleanest CMB channel 143GHz PR4 simulation maps on the lensing mask (all panels in μK , with different scales). *Middle panel:* approximate residuals after foreground cleaning using our SMICA weighting of the frequency channels, where the most visible features are a handful of bright sources leaking slightly out of the point-source mask, and zodiacal emission dust bands. *Bottom panel:* the difference of the middle panel to the PR4 simulations average, sourced by the data processing.

using diffusive inpainting prior to harmonic space weighting. If this mask is not included, the bright galactic centre produces ringing features that are distinctly visible when averaging simulations across the whole sky. We also suppress these pixels using an apodized version of the mask, but do not use an inpainting procedure on the PR4 channels. Our lensing analysis mask is then augmented to include the Boolean version of this apodized preprocessing mask. Most of the preprocessing mask is already present in the [PL2018](#) lensing mask, resulting in a negligible increase of 0.04% in masked sky fraction.

The upper panel in figure 1 shows the foreground template in the cleanest CMB channel simulations at 143GHz. Our SMICA weighting of these input foregrounds, shown on the

middle panel, gives a crude impression of the cleaning efficiency. Aside from a couple of zodiacal emission dust bands [16], the most visible residuals are a handful of sources of intensity a few tens of μK , which are the point-source component leaking a bit outside of the analysis mask. The point source mask was built for the 2015 *Planck* lensing analysis [17] based on the *Planck* catalogue of sources at 143 and 217GHz, and we find that extending the mask slightly would remove these residuals. However, they have no impact on the reconstruction (see figure 5 later on). The weighted foreground residual template varies, but only at the level of a few μK , from the simulated data-processed one as shown on the lower panel. This was obtained by taking the average of the 600 PR4 simulations, after subtracting from each simulation an approximation to their data-processed CMB component, where the isotropic effective *Planck* QuickPol beams [18, 19] were applied to the input sky. From the SMICA weights, and the pixel covariance maps produced by the NPIPE processing for each channel, we also constructed the pixel variance maps that we use in the Wiener-filtering of the combined CMB maps. We do so assuming white and independent noise across channels, and independent Stokes T, Q and U noise, following PL2018 [eqs. (15–17)]. We then rescale these maps with a constant factor such that in harmonic space this white noise value matches on the relevant (small scales) CMB multipoles approximately the empirical noise power.

2.3 CMB filtering and quadratic estimators

Given the Stokes T and $\pm_2 P \equiv Q \pm iU$ data provided on the pixelized sky, the maximum a posteriori (Wiener-filtered) modes are given by

$$\begin{pmatrix} T^{\text{WF}} \\ E^{\text{WF}} \\ B^{\text{WF}} \end{pmatrix} \equiv C^{\text{fid}} \mathcal{T}^\dagger \left[\mathcal{T} C^{\text{fid}} \mathcal{T}^\dagger + N \right]^{-1} \begin{pmatrix} T^{\text{dat}} \\ Q^{\text{dat}} \\ U^{\text{dat}} \end{pmatrix}, \quad (2.1)$$

where \mathcal{T} is our fiducial transfer function, that maps lensed T, E, B harmonic modes onto the observed Stokes maps (inclusive of the 5 arcmin beam and the HEALpix pixel window function). C^{fid} is a matrix of our fiducial lensed CMB spectra, with elements

$$[C^{\text{fid}}]_{\ell m, \ell' m'}^{XY} = \delta_{\ell \ell'} \delta_{mm'} \begin{pmatrix} C_\ell^{TT, \text{fid}} & C_\ell^{TE, \text{fid}} & 0 \\ C_\ell^{TE, \text{fid}} & C_\ell^{EE, \text{fid}} & 0 \\ 0 & 0 & C_\ell^{BB, \text{fid}} \end{pmatrix}, \quad (2.2)$$

and N is the noise matrix for the real-space Stokes data, which we assume independent between pixels, and is constructed as described in section 2.2. The fiducial spectra are those of the FFP10 cosmology, the same for the PR4 and PR3 simulation sets. The filtered maps are calculated using the conjugate-gradient method following PL2018 [20], with an improved convergence criterion putting more weight on the CMB sub-degree scales relevant for the lensing signal, resulting in faster filtering by almost a factor of two.

Foreground templates, or any set of poorly-understood modes, can be removed from the analysis by augmenting the noise matrix of eq. (2.1) and assigning them an infinite noise value, a procedure sometimes called ‘deprojection’. For example, we always deproject the temperature monopole and dipole. We also tested deprojection of the SMICA-weighted Commander foreground model, giving perfectly consistent results as shown later on.

Using these filtered CMB maps, quadratic estimators (denoted \hat{g}_{LM}^ϕ) are built in the same way as PL2018 [eqs. (3–6)]. As there, we filter CMB multipoles $2 \leq \ell \leq 2048$ but only use filtered multipoles $100 \leq \ell \leq 2048$ to construct the QEs.

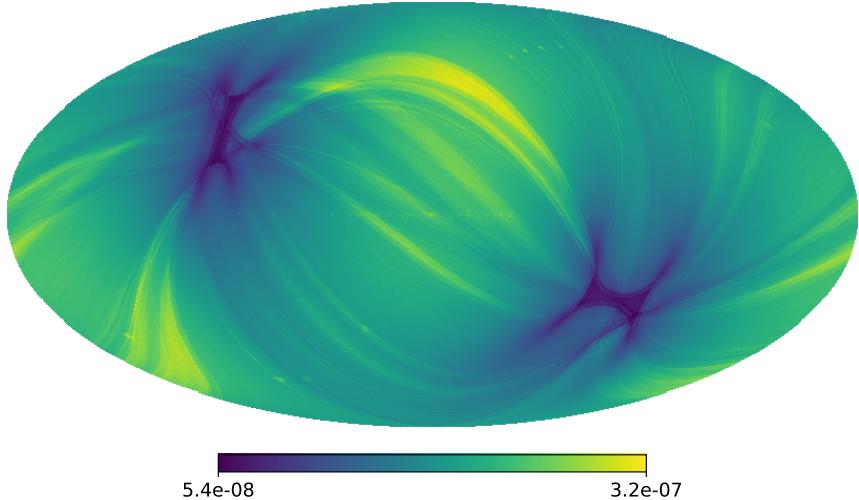


Figure 2. The large-scale effective lensing convergence reconstruction noise variance map, $N_{0,\text{eff}}^{\kappa}$, obtained by averaging the approximately flat convergence noise spectrum across the range $8 \leq L \leq 100$. After applying the same mask as for CMB filtering, and condensing this map into a simpler one with 64 homogeneous noise patches, this map is used to perform the filtering of the quadratic estimator convergence maps that are then used for our Minimum Variance estimate. This κ -filtering operation down-weights regions where the lensing reconstruction is poorest (the ecliptic equator), decreasing the bandpower errors. The overall gain remains limited, since it is CMB fluctuations, which are isotropically distributed across the sky, that dominate the reconstruction noise rather than instrumental noise.

2.4 κ -filtering and lensing spectrum

The scanning strategy of the *Planck* satellite results in many more scans of the ecliptic poles compared to the equator, resulting in inhomogeneous noise across the sky. This implies that the expected lensing reconstruction noise $N^{(0)}$ also varies depending on the sky location. The expected GMV large-scale lensing convergence reconstruction noise is shown in figure 2, spanning almost one order of magnitude.

In order to account for this, once the QEs are obtained, we filter them following the methodology which was introduced by ref. [11]. The unmasked regions of the sky are partitioned into 64 patches (based on their relative pixel CMB temperature noise level). This partitioning is done by building a histogram of the inverse noise map. The temperature and polarization noise variances in each patch (histogram bin) are then approximated as homogeneous, which is the main requirement for this filtering step. This partitioning scheme is used to construct the effective lensing reconstruction noise level of each patch, $N_{0,\text{eff}}^{\kappa}$, which is used in the filtering process. The local patch reconstruction noise values are calculated from averaging the values of $N_{0,L}^{\kappa}$ (the isotropic reconstruction noise $N_{0,L}$ for $\kappa_{LM} \equiv L(L+1)\phi_{LM}/2$) over the multipole range $8 \leq L \leq 100$. In this multipole range the reconstruction noise is relatively white. The same partitioning scheme is also used to construct an effective response map, $\mathcal{R}_{\text{eff}}^{\kappa}$, which is used to properly normalize the QEs. The values for the response map $\mathcal{R}_{\text{eff}}^{\kappa}$ are obtained in a similar way from the response \mathcal{R}_L^{κ} . Figure 2 shows the reconstruction noise map which is used in the filtering. We set $N_{0,\text{eff}}^{\kappa}(\hat{n})^{-1} = 0$ for all masked pixels. The additional filtering step is performed on the quadratic estimator convergence κ , as it is approximately local in real space, and hence is uncorrelated between patches on large scales and has approximately white noise.

Similarly to eq. (2.1), we define the Wiener-filtered estimator as

$$\hat{\kappa}^{\text{WF}} \equiv C_{\text{fid}}^{\kappa\kappa} \left[C_{\text{fid}}^{\kappa\kappa} + N_{0,\text{eff}}^{\kappa} \right]^{-1} (\mathcal{R}_{\text{eff}}^{\kappa})^{-1} \hat{g}^{\kappa}, \quad (2.3)$$

where $C_{\text{fid}}^{\kappa\kappa}$ is the fiducial theory κ power spectrum, and the QE map \hat{g}^{κ} is defined after mean field subtraction from²

$$\frac{1}{2} L(L+1) \hat{g}_{LM}^{\kappa} \equiv \left(\hat{g}_{LM}^{\phi} - \langle \hat{g}_{LM}^{\phi} \rangle_{\text{MC}} \right). \quad (2.4)$$

The bracketed matrix inverse in eq. (2.3) is the only non-trivial step, since the first and second terms in that matrix are diagonal in harmonic and pixel space respectively. This step is performed using the same conjugate gradient method and preconditioner as used for filtering the CMB maps (section 2.3), which results in fast enough convergence to the solution. Eq. (2.3), the ‘ κ -filtered estimator’, gives us a novel (yet unnormalized) lensing map whose spectrum is by construction approximately optimally weighted across the sky. We produce these maps in pairs $\hat{\kappa}_1, \hat{\kappa}_2$, with subtracted mean-field estimates in eq. (2.4) built from two non-overlapping subsets of simulations. This avoids the mean-field Monte-Carlo noise in the raw spectrum estimate, given by

$$C_L^{\hat{\kappa}\hat{\kappa}} \equiv \frac{1}{(2L+1)f_{A,L}} \sum_{M=-L}^L \hat{\kappa}_{1,LM}^{\text{WF}} \hat{\kappa}_{2,LM}^{\text{WF},*}. \quad (2.5)$$

The factor $f_{A,L}$, a scale-dependent effective sky fraction, encapsulates the effects of masking and of the κ -filtering normalization. Using the independent patch approximation, we may write [11]

$$f_{A,L} \equiv \sum_{\text{patches } p} f_p (w_L^p)^2 \quad (2.6)$$

where f_p is the sky area of patch p , and

$$w_L^p = \frac{C_{L,\text{fid}}^{\kappa\kappa}}{C_{L,\text{fid}}^{\kappa\kappa} + N_{0,\text{eff}}^{\kappa,p}} \frac{\mathcal{R}_L^{\kappa,p}}{\mathcal{R}_{\text{eff}}^{\kappa,p}}. \quad (2.7)$$

Beside the lensing spectrum signal, the raw spectrum in eq. (2.5) still contains a number of additional terms, the most relevant for *Planck* noise levels being the two of lowest orders in the lensing potential $N_L^{(0)}$ and $N_L^{(1)}$ [21, 22]. We use the now very standard realization-dependent RD- $\hat{N}_L^{(0)}$ [17, 23] and simulation-based MC- $N_L^{(1)}$ [1, 23] debiasers, which can be applied without modifications to the GMV and κ -filtered maps. At higher order, $N_L^{(3/2)}$ induced by the large-scale structure bispectrum and post-Born lensing is negligible at *Planck* noise levels [24–28], and $N_L^{(2)}$ is made negligible using the lensed CMB spectra in eq. (2.1) [21].

After bias subtraction, we apply a multiplicative Monte-Carlo correction to the bandpowers just as in PL2018, by comparing simulated reconstruction power spectra with input lensing power. In previous analyses, this residual Monte-Carlo correction is mostly sourced by the presence of the mask: masking renders the analytic isotropic normalization of the QE inaccurate close to the mask boundaries, while application of a correct anisotropic normalization appears intractable. In our case, both inhomogeneous CMB- and κ -filtering impacts

²The relation (2.4) is inverse to that of the lensing convergence κ to the potential ϕ because the unnormalized quadratic estimator \hat{g} are gradients with respect to these quantities.

the response of the QE to lensing, hence its normalization as well. Expression (2.7) and the isotropic rescaling of eq. (2.5), which originates from the approximate patch model, is of course not exact. However, all terms of the debiased spectrum scale according to the same $1/f_{A,L}$ so that a modelling error is captured automatically by the Monte-Carlo correction. The true response function only depends on the sky signal CMB spectra, which are very well constrained and whose variations always play a very subdominant role in *Planck* lensing likelihoods. As long as the correction is small, its cosmology dependence may be completely neglected. In previous releases it was at most $\sim 10\%$ on the lowest multipole bins. With the κ filtering we now find that the residual MC correction is slightly smaller (which must be seen as a coincidence rather than a virtue of the model).

We finally apply a (very small) point-source (PS) correction $\Delta\hat{C}_L^{\text{PS}}$, discussed in more detail in section 2.6. Our final bandpowers may then be written as

$$\hat{C}_L^{\kappa\kappa} \equiv C_L^{\kappa\kappa} - \text{RD-}\hat{N}_L^{(0)} - \text{MC-}\hat{N}_L^{(1)} - \Delta\hat{C}_L^{\text{PS}}. \quad (2.8)$$

We find that with κ filtering, the reconstructed bandpower variance is slightly improved around the peak of the reconstructed power spectra; not only for our GMV reconstruction (see figure 3), but also for the polarization-only reconstruction, which is much more noise-dominated. This is quantified in section 3.2. For more sensitive data from future SO and CMB-S4 experiments, the expected improvement was already demonstrated using a fiducial small flat-sky area scanning strategy in ref. [11]. Revisiting their analysis for realistic curved-sky scanning strategies from Chile similar to those of ref. [29] we find similar results, with variance improvements in the range between 5–25%. This improvement clearly depends primarily on the homogeneity level of the scan; there is greater improvement for inhomogeneous scans.

2.5 Covariance matrix

The data bandpowers are debiased with the usual realization-dependent $\text{RD-}\hat{N}_L^{(0)}$ estimate, which uses quadratic estimators that cross maps from the data with simulations. While this improves the covariance of the estimator, it is too expensive to perform the same type of realization-dependent subtraction on each and every simulation. There are at least a couple of options available:

- One option is to consider the covariance matrix obtained without (or, with realization-*independent*) debiasing. This slightly overestimates both the variance and covariance of the binned bandpowers, particularly so on small scales [30].
- Another possibility is to use a numerically cheap approximate realization-dependent debiaser: for example, PL2018 used the analytical formula for the $N_L^{(0)}$ noise on an isotropic sky with homogeneous filtering, but replacing the idealized power spectra by the empirical ones, to obtain a semi-analytic $\hat{N}_L^{(0)}$ estimate. The accuracy of this estimate is insufficient for the purpose of precision debiasing on realistic data. It is, however, quite efficient at incorporating the removal of the small-scale covariance, and, using homogeneous CMB filtering, appears good enough for a covariance matrix estimate.

Here we try to improve on the second option, by building a cheap debiaser from simulations which takes into account the filtering used on the data, and makes no assumptions of isotropy. We first discuss the case without κ -filtering: for each of 480 simulations dedicated to the covariance matrix estimate, we build an $\text{RD-}\hat{N}_L^{(0)}$ estimate using a much smaller number

(N_{RD}) of other simulations than are used to debias the data. These simulations come from the same original set, requiring the filtering of no additional maps, and we neglect the impact on the covariance of simulation subset overlap. The additional QE's ($N_{\text{RD}} \times 480$) that are required are built on the fly from the filtered maps. Since N_{RD} is small, the reconstruction noise part of the covariance is overestimated slightly. The proportionality constant is $4/N_{\text{RD}}$ as derived by [PL2018](#), appendix C. This motivates the following procedure: for a given N_{RD} , we build a correction $c_L^{N_{\text{RD}}}$ to the empirical full covariance matrix $\hat{\Sigma}^{\text{RD}}$ and rescale it as

$$\hat{\Sigma}_{LL'} \equiv \frac{\hat{\Sigma}_{LL'}^{\text{RD}}}{\sqrt{c_L^{N_{\text{RD}}} c_{L'}^{N_{\text{RD}}}}} \quad (2.9)$$

with

$$c_L \equiv 1 + \frac{4}{N_{\text{RD}}} \left(\frac{\text{MC-} \hat{N}_L^{(0)}}{C_L^{\hat{\phi}\hat{\phi}}} \right)^2, \quad (2.10)$$

where $C_L^{\hat{\phi}\hat{\phi}}$ is the mean-field subtracted raw data power. We find that even for very small N_{RD} this estimate is quite stable as a function of N_{RD} , and we used $N_{\text{RD}} \sim 16$ for our final estimates. A value of N_{RD} that is too large is both costly and not desirable: with all the simulations coming from the same set, the overlap would reintroduce the correlations that the realization-dependent debiasing removes. This covariance shares the characteristic qualitative features expected for RD- $N^{(0)}$ -debiased estimate; that is, less covariance between the spectrum estimates at different lensing multipoles, as well as lower variance at the highest lensing multipoles. With homogeneous filtering, this estimate is also numerically closer to the semi-analytic debiasing result of [PL2018](#) than without RD-debiasing. On the conservative range the size of the effect is visible but remains a few percent.

The κ -filtering operation of section 2.4 has a small effect close to the peak of the lensing power. To build a realization-dependent covariance without having to κ -filter a large number of cross-simulation maps, we simply rescale the non- κ -filtered covariance in the following manner: with r_{L_b} the ratio between the realization-independent variances in each bin b with and without κ filtering, we multiply the realization-dependent, non κ -filtered b, b' covariance by $\sqrt{r_{L_b} r_{L_{b'}}}$.

The top panel of figure 3 shows the error bars on our most precise reconstruction (red), together with variations (orange and green) and the PR3 baseline result (blue). The bottom panel shows the corresponding ratio to the PR3 errors. For most bins the improvement that we see in this work is about half from the slightly reduced noise of PR4, and half from lensing pipeline improvements.

2.6 Point-source correction

It is well known that lensing estimators can respond strongly to point-source-like signals in the CMB data [\[31, 32\]](#). To check for a residual point-source signal, we apply the point-source anisotropy estimator $S_2(\hat{n})$ to our maps. After subtraction of the expected lensing-induced trispectrum contamination to the \hat{S}_2 spectrum, we obtain a result consistent with a white noise spectrum, and non-zero at about $3\text{-}4\sigma$. We compress this spectrum further into a trispectrum shot-noise amplitude S_4 . The details of this step differ slightly from the *Planck* 2015 and 2018 analyses, since we noticed that on simulations our \hat{S}_2 spectrum estimates

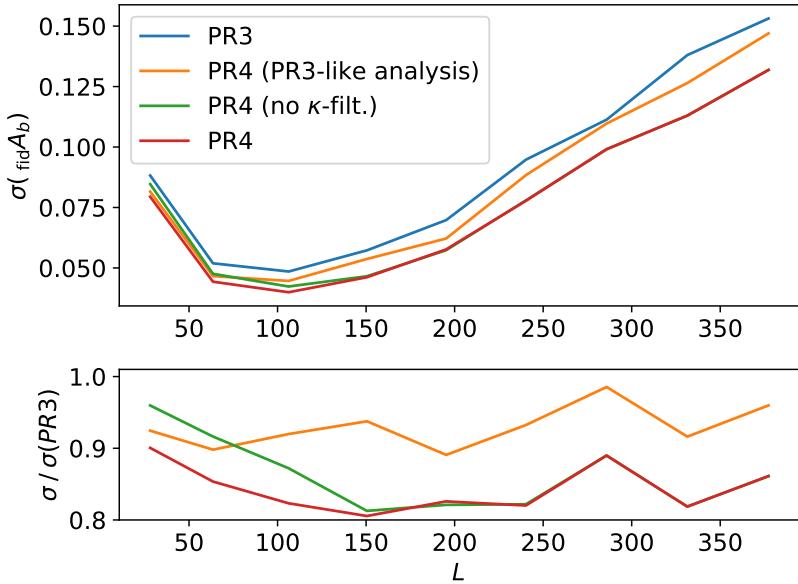


Figure 3. *Top panel:* Error bars on our reconstructed lensing power spectrum bandpower amplitudes $\text{fid}A$ relative to the FFP10 fiducial model, over the conservative range $8 \leq L \leq 400$. The inhomogeneously CMB-filtered and κ -filtered GMV bandpowers, with κ -filtering (our best result, red), and without κ filtering (green), are shown compared to the published MV [PL2018](#) reconstruction (blue). The orange line shows the results of applying the same 2018 pipeline to the NPIPE maps. *Bottom panel:* The ratio of these errors to those of the [PL2018](#) reconstruction. The reconstruction without κ -filtering, but using improved CMB-filtering with respect to PR3 (green), performs worse at low multipoles than the PR3-like analysis in agreement with purely analytic predictions. This is because using inverse reconstruction noise $N^{(0)}$ weighting for power spectrum estimation, as implied without κ -filtering, is suboptimal on large scales where the first couple of bins are now highly signal dominated. This is corrected by κ -filtering, which produces a reconstruction with (very) slightly higher $N^{(0)}$ bias but smaller error bars by effectively using more of the sky.

were highly correlated on most scales,³ with a positive cross-correlation greater than 50% at $L > 500$. Taking this into account shifts the weighting of the \hat{S}_2 spectrum towards larger scales, and provides us with a more precise estimate, reducing the error on S_4 by 30% with respect to the *Planck* analysis weighting. We measure

$$10^{13} \hat{S}_4 = 3.0 \pm 0.8 \mu\text{K}^4, \quad (2.11)$$

an estimate larger by about 30% than the result from the PR3 maps. This non-zero trispectrum affects our κ -filtered lensing power according to

$$f_{A,L} \Delta \hat{C}_L^{\text{PS}} = \hat{S}_4 \sum_p f_p (w_L^p)^2 \left(\frac{\mathcal{R}_L^{\kappa s, p}}{\mathcal{R}_L^{\kappa, p}} \right)^2 \quad (2.12)$$

where $\mathcal{R}_L^{\kappa s, p}$ is the local response of the lensing estimator to the point source field. The effect of the inhomogeneous filtering is small, the correction having very much the same shape as for more standard homogeneous filtering where all terms in the sum are identical. This correction

³This large covariance is without a realization-dependent debiaser, which could help further improve the estimate of the point-source signature, but at substantial numerical cost.

reaches a fourth of the bandpower error on the last few bins of the conservative range, and is smaller elsewhere. To account for the uncertainty in S_4 , the covariance matrix is augmented by the rank-one $\sigma_{S_4}^2 \mathcal{PS} \cdot \mathcal{PS}^t$, where

$$\mathcal{PS}_L \equiv \frac{1}{f_{A,L}} \sum_p f_p (w_L^p)^2 \left(\frac{\mathcal{R}_L^{\kappa s, p}}{\mathcal{R}_L^{\kappa, p}} \right)^2 \quad (2.13)$$

is the point-source correction template (before binning and application of the final percent-level Monte-Carlo correction) and $\sigma_{S_4}^2$ is the squared error on the trispectrum quoted above. Added in quadrature to the covariance, this is everywhere a very tiny correction to the bandpower errors: at most 0.3% on the conservative range and practically zero on the lensing power peak, but positively correlated across the lensing potential spectrum tail.

2.7 Likelihoods

We produce likelihoods for our improved bandpowers, generalizing the approach of [PL2018](#). The construction of the likelihoods for the CMB lensing bandpowers is complicated by the fact that, as well as direct dependence on $C_L^{\phi\phi}$ and via $N^{(1)}$, the estimators also carry dependence on cosmology through the lensed CMB power spectra entering $N^{(0)}$, $N^{(1)}$, the mean-field and the estimator normalization. We follow previous approaches, making use of the fact that the CMB spectra are very well constrained, which can be summarized as

- neglecting the cosmology dependence of the RD- $N^{(0)}$ debiasing (since by construction it is insensitive to leading order differences between the fiducial and true sky CMB spectra),
- linearizing the response functions and $N^{(1)}$ matrix to small changes in the CMB spectra,
- neglecting the cosmology dependence of the residual MC and PS corrections (which are small),
- and neglecting the cosmology dependence of the mean field. The dominant source of mean field is masking, which creates a strong signature directly proportional to its isotropic power spectrum even for perfectly Gaussian data. In our case, the mean field is still sizeable on the first bin of the conservative range. Unless the simulation power has been matched accurately to the measured data power the mean field sourced by the CMB signal has a cosmology dependence. By looking at the data splits this can be inferred to be between a third and half of the total. However, since the mean field is subtracted at the map level (and not from the bandpowers), the impact is quadratic in the mismatch of CMB spectra, hence completely negligible except at the very lowest multipoles (see figure 6, brown points).

This allows us to produce a fast and accurate likelihood for most cases of practical interest. In our case there is one difference to previous work: since we use inhomogeneous filtering, the responses and bias terms vary according to the location on the sky. We take this into account by approximating the sky as consisting of patches of homogeneous noise levels, so that the full linear response matrix is itself the sum of the same number of such matrices, each built for a different filtering noise level. As discussed in section 2.4 this approximation is very effective and gives a residual MC-correction (sourced predominantly by the presence

of the analysis mask) of the same size at the more standard case of homogeneous filtering (where there is a single response function and $N^{(1)}$ across the sky).

Explicitly, our likelihood is constructed as follows. Considering the decomposition of the sky in patches with roughly homogeneous filtering noise levels, we can consider that the lensing map estimate in patch p responds to the sky signal as

$$\hat{\kappa}_{LM}^{\text{WF},p}(\theta) \sim w_L^p \frac{\mathcal{R}_L^p(\theta)}{\mathcal{R}_L^p(\theta^{\text{fid}})} \kappa_{LM}^{\text{sky}}. \quad (2.14)$$

For homogeneous CMB- and no κ -filtering there is a single patch, with unit w_L^p , and we recover the [PL2018](#) likelihood construction. In this more general case we can write the linearized likelihood prediction for the observed spectrum prior to binning

$$\begin{aligned} C_L^{\kappa\kappa, \text{pred}}(\theta) &= C_L^{\kappa\kappa}(\theta) + \sum_{L'} \mathcal{M}_{LL'}^{\kappa\kappa} \left(C_{L'}^{\kappa\kappa}(\theta) - C_{L'}^{\kappa\kappa}(\theta^{\text{fid}}) \right) \\ &\quad + C_L^{\kappa\kappa}(\theta^{\text{fid}}) \sum_{\ell, XY} \mathcal{M}_{L\ell}^{XY} \left(C_{\ell}^{XY}(\theta) - C_{\ell}^{XY}(\theta^{\text{fid}}) \right) \end{aligned} \quad (2.15)$$

where the matrices involved are

$$\mathcal{M}_{L\ell}^{XY} = \sum_{\text{patches } p} f_p (\tilde{w}_L^p)^2 \left(\frac{\partial \ln (\mathcal{R}_L^p)^2 (\theta^{\text{fid}})}{\partial C_{\ell}^{XY}} + \frac{\partial N_L^{(1),p}(\theta^{\text{fid}})}{\partial C_{\ell}^{XY}} \right) \quad (2.16)$$

for $XY \in (TT, TE, EE)$, and

$$\mathcal{M}_{LL'}^{\kappa\kappa} = \sum_{\text{patches } p} f_p (\tilde{w}_L^p)^2 N_{LL'}^{(1),p}(\theta^{\text{fid}}) \quad (2.17)$$

with $\tilde{w}_L^p \equiv w_L^p / f_{A,L}$. In order to get lensing-only constraints, we make the assumption that the fluctuations in the CMB spectra in the second line of eq. (2.15) are Gaussian. We can then analytically marginalize these out, obtaining a likelihood that depends only on the cosmology's lensing power. For this we proceed just as [PL2018](#), using the smoothed PR3 `plik_lite` CMB bandpowers and covariance matrix. This leads to an increase of the lensing covariance by several percent, and to a small shift in power sourced by the small differences between the FFP10 and the *Planck* empirical spectra. We also produce a likelihood for the polarization-only reconstruction. It is of course much noisier, and for this likelihood we neglect the CMB spectra corrections altogether.

Construction of the likelihood requires calculations of many $N^{(1)}$ lensing matrices and their derivative with respect to the CMB spectra. To perform this in a more efficient way than in previous releases, especially for the GMV estimator, we have written a novel fast Fourier transform based algorithm for their calculation,⁴ in many cases orders of magnitude faster than previous methods.

3 Results

3.1 GMV and TT reconstructions

Figure 4 shows our bandpower reconstructions on the conservative range $8 \leq L \leq 400$ (green), together with the [PL2018](#) results (blue). Table 1 lists our GMV, κ -filtered bandpower values

⁴<https://github.com/NextGenCMB/lensitbiases>.

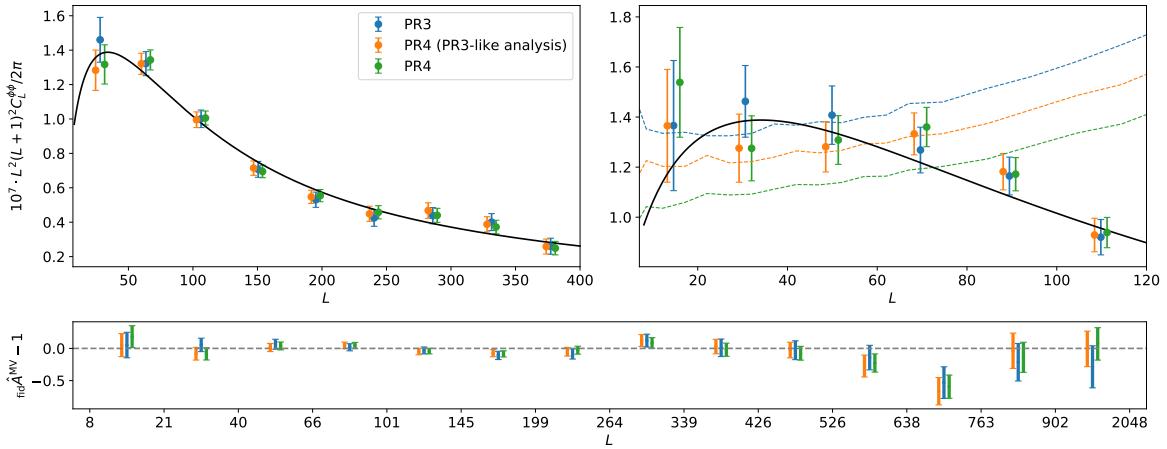


Figure 4. The left panel shows our most precise bandpower estimates over the conservative multipole range (green), together with the *Planck* 2018 (PR3) Minimum Variance bandpowers (blue). The black line is the lensing power spectrum best-fit to the *Planck* 2018 CMB spectra data (without lensing). The right panel focuses on lensing multipoles close to the Λ CDM lensing power spectrum peak, and also shows as the dotted lines the corresponding realization-dependent $N^{(0)}$ biases (without residual Monte-Carlo correction). Many more modes are now signal dominated. All bandpowers are built with the same approximate inverse-variance multipole weighting scheme given in [PL2018](#). The orange points show the results of implementing the PR3 reconstruction analysis on the NPIPE maps by using a homogeneous noise sky model as well as independent temperature to polarization filtering. The bottom panel shows the relative deviation to the fiducial FFP10 cosmology amplitude on the entire range.

which we later use to constrain parameters. To highlight the differences caused by the new PR4 data processing, figure 4 also shows our results on the new NPIPE maps with the same lensing analysis choices as 2018 (orange). The same binning scheme is used in all three cases. The reconstructions do display some differences: the right panel provides a closer view around the lensing spectrum peak, with a finer binning, where several shifts comparable to the error bars can be seen. Summary statistics and broad characteristics of the spectrum remain very similar however. Using our CMB-marginalized likelihood (see section 2.7), we fit an amplitude of the lensing spectrum relative to the Λ CDM prediction from the best-fit to the 2018 *Planck* CMB spectra (without lensing, `plikHM_TTTEEE_lowl_lowE`) [33]. On the conservative range we get

$$\text{b.f. } A_{8 \rightarrow 400}^{\text{GMV}} = 1.00 \pm 0.024, \quad (3.1)$$

with a χ^2 of 8.5 for 9 data points, with 2018 value 1.01 ± 0.026 .⁵ Both spectra show the same slight tilt with respect to the best fit, with slightly higher (lower) amplitude preferred at low (high) multipoles. Owing to the shift on the first conservative bin this tilt is only slightly less pronounced for PR4. This property is slightly more pronounced for the temperature-only reconstruction, where the two lowest conservative bins are higher than their PR4 counterparts by approximately 1σ using the same lensing pipeline.

Given the complexity of the data processing, ‘paired’ noise simulations — sharing the same random seeds between the 2018 release and NPIPE — are neither available nor planned. This makes it difficult to evaluate if the peak shifts seen can be explained solely by the data

⁵The shift in amplitude is closer to 0.003 before rounding.

L_{\min}	L_{av}	L_{\max}	MV-fid	${}_{\text{fid}}\hat{A}^{\phi^{\text{MV}}}$
8	28.1	40	1.40	0.95 ± 0.08
41	63.5	84	1.28	1.05 ± 0.05
85	106.2	129	$9.93 \cdot 10^{-1}$	1.01 ± 0.04
130	150.4	174	$7.62 \cdot 10^{-1}$	0.91 ± 0.05
175	195.1	219	$5.99 \cdot 10^{-1}$	0.92 ± 0.06
220	240.3	264	$4.84 \cdot 10^{-1}$	0.94 ± 0.08
265	285.8	309	$4.01 \cdot 10^{-1}$	1.10 ± 0.10
310	331.4	354	$3.38 \cdot 10^{-1}$	1.10 ± 0.12
355	377.3	400	$2.88 \cdot 10^{-1}$	0.86 ± 0.14

Table 1. Bandpowers on the conservative range, for our Minimum Variance (MV) baseline, including joint temperature and polarization filtering. Our polarization-only (PP) reconstruction can be found in table 2. The reconstruction uses inhomogeneous CMB filtering and the additional κ -filtering step. Numbers are given for $10^7 \cdot L^2(L+1)^2 C_L^{\phi\phi} / 2\pi$.

processing differences. We attempted to quantify this in the following approximate manner.⁶ We estimated a cross-correlation coefficient $\rho_{\ell}^{\text{noise}}$ of the noise in harmonic space from the ratio of the cross spectrum to the square roots of the auto spectra of NPIPE and 2018 SMICA maps, after subtracting the CMB signal spectrum obtained from the cross-spectrum of the NPIPE A and B splits. On small scales, we found an almost scale-independent $\rho_{\ell}^{\text{noise}} \sim 0.8$, both in temperature and polarization, and both in the SMICA maps and the main CMB frequency channels. We then constructed paired SMICA 2018 noise simulations by adding a Gaussian component to our SMICA NPIPE noise simulation, so that the new 2018 simulation has the right auto-power and cross-correlation $\rho_{\ell}^{\text{noise}}$ to the NPIPE one. Performing lensing reconstruction on the set of simulation built in this way, we can assess the significance (in this crude model) of the shifts compared to the NPIPE ones. The measured correlation value of 0.8 is smaller than expected from only 10% more data, which suggests that this estimate may be conservative. We find that the lowest bin on the peak is 3.1σ low according to this criterion for the MV reconstruction. For the TT reconstruction, the first two bins are about 2.4σ low. Other points do not stand out.

Figure 5 shows the PR3 (blue) and PR4 (orange) conservative range reconstructions from temperature only, together with several PR4 variations, always using homogeneous filtering as for PR3. The points shown include

- **fg deproj.:** The residual foreground model after SMICA weighting (see figure 1 bottom panel) is assigned infinite variance in the filtering process to project it out. This is very efficient at reducing the (already small) common component to all simulations.
- **PS ext.:** The point-source component of the lensing analysis mask is extended, increasing the masked sky fraction by 1.5%. This mask is based on resolved point sources at 143GHz and 217GHz, but a handful of bright sources can be seen leaking outside of it slightly after SMICA weighting both on the foreground model and data.

⁶We thank A. Challinor for suggesting this approach.

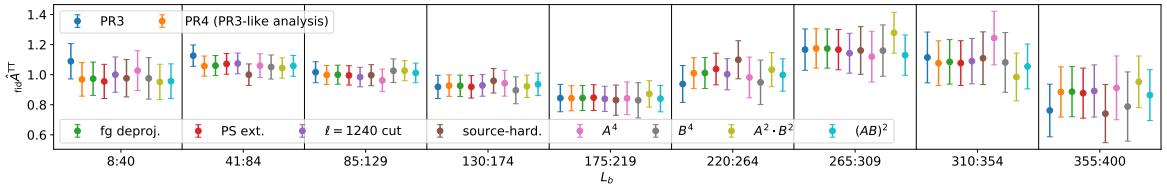


Figure 5. Lensing temperature-only bandpowers in units of the fiducial FFP10 lensing spectrum for each bin of the conservative range (without point-source subtraction), with the PR3 points in blue, and several test reconstructions on PR4 maps. All of the reconstructions used here use the same filtering and main analysis choices as PR3, with the orange points showing the PR4 result corresponding to the blue points. The filtering deprojects the residual foreground model to give the green points. The point-source component of the analysis mask is extended for the red points. The purple points exclude a range of 20 CMB multipoles centred on 1240, for reasons given in the text. The brown points show the reconstruction bias-hardened against point sources. The shifts seen in these cases are all compatible with the expected variance obtained from differencing simulations. Finally, the last four sets of points show results from the A and B splits, namely $\hat{\kappa}^{AA} \cdot \hat{\kappa}^{AA}$, $\hat{\kappa}^{BB} \cdot \hat{\kappa}^{BB}$, $\hat{\kappa}^{AA} \cdot \hat{\kappa}^{BB}$ and $\hat{\kappa}^{AB} \cdot \hat{\kappa}^{AB}$ respectively.

- **$\ell = 1240$ cut:** This excludes the CMB multipole range $1230 \leq \ell \leq 1250$ from the quadratic estimators. Motivations for this were the presence of a spike at $\ell \sim 1240$ in the otherwise smoother residual foreground model power spectrum seen on the lensing mask. This multipole also coincides with a break in the PR3 (hence by construction also our PR4) SMICA weights.
- **source-hard.:** This shows the lensing reconstruction bias-hardened against point-source contamination [31, 34]. Here a correction is made at the map level, subtracting from the lensing estimate a properly normalized estimate of a point source signature, itself obtained with a quadratic estimator. This has been shown to be very effective against contaminants that act similarly to a local lensing magnification, at the cost of some signal to noise. In this case a couple of larger shifts can be seen in the bandpowers. However, the estimators themselves differ more than in the previous cases, and these shifts are no larger than the spread expected from simulations.

The last four sets of points in figure 5 show the reconstructions from PR4 A and B data splits. We obtain lensing maps from the A and B splits and reconstruct the power from $\hat{\kappa}^{AA} \cdot \hat{\kappa}^{AA}$ (pink), $\hat{\kappa}^{BB} \cdot \hat{\kappa}^{BB}$ (grey), as well as $\hat{\kappa}^{AA} \cdot \hat{\kappa}^{BB}$ (chartreuse) and $\hat{\kappa}^{AB} \cdot \hat{\kappa}^{AB}$ (cyan). The last two have in principle vanishing contributions of the noise maps to their $N^{(0)}$ bias or mean-field power respectively. We refer to PL2018 for many more such consistency tests for PR3. We collect a couple of further points of comparison between the PR3 and PR4 results:

- $L = 2$ anomaly: the PL2018 reconstruction has a clearly anomalous-looking quadrupole. In the case of the MV reconstruction

$$\text{fid } \hat{A}_{L=2}^{\phi\phi, \text{MV}} = 22 \pm 3 \text{ (PR3)} \quad (3.2)$$

This value, and the ones below, do not include a MC correction, which would not affect this discussion. We also neglect the non-Gaussianity of the quadrupole power posterior. This high quadrupole value is completely driven by temperature, with

$$\text{fid } \hat{A}_{L=2}^{\phi\phi, \text{TT}} = 40 \pm 4 \text{ (PR3).} \quad (3.3)$$

The much noisier polarization-only reconstruction is consistent with zero. We note that the mean-field spectrum is maximal at very low multipoles and absolutely massive (in these units, the TT mean-field is $\sim 5 \cdot 10^4$ at $L = 2$, but only ~ 16 on the first conservative bin $8 \leq L \leq 40$), so that accurate reconstruction is challenging and can be ruined by small inaccuracies in the simulation modelling. On the other hand, we now see different values on the NPIPE maps. Using the same analysis as [PL2018](#) we find

$${}_{\text{fid}}\hat{A}_{L=2}^{\phi\phi,\text{MV}} = 6 \pm 2.5, \text{ (PR4, 2018-like),} \quad (3.4)$$

and

$${}_{\text{fid}}\hat{A}_{L=2}^{\phi\phi,\text{TT}} = 8 \pm 4, \text{ (PR4, 2018-like),} \quad (3.5)$$

and with our new baseline analysis

$${}_{\text{fid}}\hat{A}_{L=2}^{\phi\phi,\text{GMV}} = 1.3 \pm 1.9, \text{ (PR4, } \kappa\text{-filtered)} \quad (3.6)$$

This better-behaved temperature reconstruction result seems robust against choices of data splits and combinations, as shown in figure 6. In all cases we use the same filtering choices as [PL2018](#). We have also performed reconstruction using the full mission maps, with the SMICA-weighted Commander foreground template either de-projected in the filtering step, or subtracted from the maps prior to filtering, with very consistent results.

It is tempting to speculate from this that the NPIPE processing and simulations handle better at least some of the largest-scale non-idealities, though more concrete evidence is lacking to back this up further.

- Lensing curl spectrum from temperature: [PL2018](#) discussed and tested in some detail the presence of a possibly suspicious feature in the high- L lensing curl mode constructed from temperature. In that regard, however, the NPIPE results are very similar. Figure 7 shows the TT-curl mode spectrum, processed with the same homogeneous filtering pipeline. The reconstructions from the splits give qualitatively similar results.

3.2 Polarization-only reconstruction

Table 2 lists our updated polarization bandpowers. Compressing this into an amplitude, we find

$${}_{\text{fid}}\hat{A}_{8 \rightarrow 400}^{\phi\phi,\text{PP}} = 0.89 \pm 0.07 \text{ (PR4, } \kappa\text{-filtered),} \quad (3.7)$$

whereas the corresponding number quoted in [PL2018](#) is (with inhomogeneous CMB-filtering) 0.95 ± 0.11 . The improvement in signal to noise can be decomposed roughly as follows. The lensing reconstruction noise is proportional to two powers of the data spectra. While the instrumental noise sources only 10% of the lensing reconstruction noise for temperature, this dominates for polarization. Hence the reduced noise of the NPIPE maps has a greater impact on the polarization-only reconstruction. Using the PR4 maps in place of PR3 we see a decrease in bandpower variance of almost 25% (on the peak) to 35% (on the tail). The benefit of κ -filtering is (as expected) more subtle, but still visible on the first two bins of the conservative range, with an improvement of $\sim 8\%$. For this polarization reconstruction, we have also pre-processed the maps, by modifying the SMICA weights in order to flatten the sub-degree scales empirical noise spectra. By making the fiducial covariance model of the CMB closer to that of the data, this also slightly decreases the error bars, by about 6%.

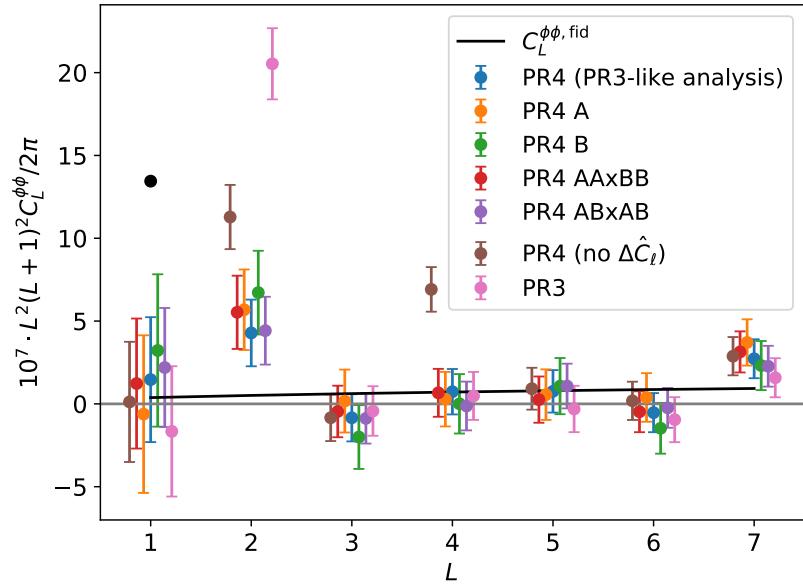


Figure 6. Lensing spectrum reconstruction from temperature on the lowest multipoles for a variety of NPIPE data splits and combinations. The comparison result for [PL2018](#) (pink) clearly shows an anomalous-looking quadrupole (all points are shown here without Monte-Carlo corrections, which do not solve this issue). The brown points show the effect of not adding the few-percent sized additional high- ℓ power to the temperature simulations to match the unresolved foreground power in the data. If the missing power can be treated as originating from a Gaussian field, the signature on the bandpowers is quadratic in the mismatch, but the mean field at the first couple of even multipoles is so large that it becomes distinctly visible. The signature is however completely negligible on all of the conservative range bins. The black dot shows the expected dipolar aberration signal from our motion with respect to the CMB frame, which appears to be consistently taken into account by the mean-field in both cases. None of these very-low lensing multipoles enter the lensing likelihood.

L_{\min}	L_{av}	L_{\max}	PP-fid	${}_{\text{fid}}\hat{A}^{\phi^{\text{PP}}}$
8	28.0	40	1.40	0.67 ± 0.23
41	63.3	84	1.28	0.94 ± 0.14
85	105.8	129	$9.95 \cdot 10^{-1}$	0.94 ± 0.15
130	149.8	174	$7.64 \cdot 10^{-1}$	0.88 ± 0.19
175	194.3	219	$6.01 \cdot 10^{-1}$	0.85 ± 0.26
220	239.3	264	$4.86 \cdot 10^{-1}$	1.18 ± 0.36
265	285.0	309	$4.02 \cdot 10^{-1}$	0.89 ± 0.56
310	330.9	354	$3.38 \cdot 10^{-1}$	0.30 ± 0.62
355	377.0	400	$2.89 \cdot 10^{-1}$	0.33 ± 0.82

Table 2. Bandpowers on the conservative range, for our polarization-only reconstruction. Our Minimum variance bandpowers can be found in table 1. This reconstruction uses inhomogeneous CMB filtering and the additional κ -filtering step. Numbers are given for $10^7 \cdot L^2(L+1)^2 C_L^{\phi\phi} / 2\pi$.

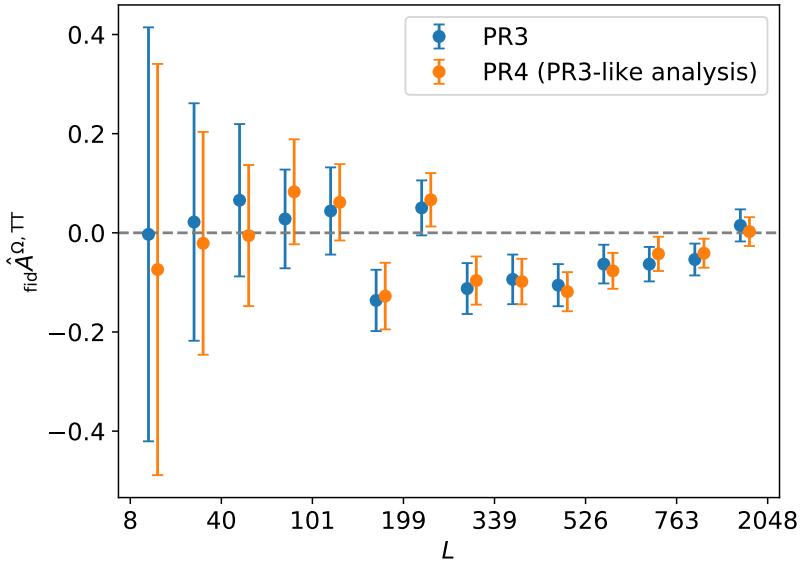


Figure 7. Lensing curl mode spectrum in units of a fiducial spectrum $L^2(L+1)^2C_L^{\Omega\Omega}/2\pi = 10^{-7}$, for the 2018 analysis and the same analysis on the new PR4 maps. The 4.3σ significance of the low feature at $264 \leq L \leq 901$ as assessed by [PL2018](#) remains virtually unchanged with the NPIPE new data processing. Reconstructions on the PR4 A and B splits give very similar results.

3.3 ISW-lensing

Cross-correlating the lensing quadratic estimator to the temperature gives the lensing ISW bispectrum [35–37], which is visible on large scales. For simplicity we evaluate this only using the simplest homogenously-filtered lensing reconstruction map. In figure 8 we show our estimate of $\hat{C}_\ell^{T\phi}$, which we obtained from the simple estimator

$$\hat{C}_L^{T\phi} = \frac{1}{(2L+1)f_{\text{sky}}} \sum_{M=-L}^L \hat{\phi}_{LM}^{\text{MV}} \hat{T}_{LM}^*, \quad (3.8)$$

where \hat{T} is our homogeneously inverse-variance Wiener-filtered temperature map T^{WF} , rescaled with the isotropic limit $C_\ell^{TT,\text{fid}}/(C_\ell^{TT,\text{fid}} + N_\ell^{TT,\text{fid}})$ of the filter in order to make it unbiased away from the mask. In principle, we could use the Wiener-filtered lensing maps here as well, but the improvement is negligible. We also apply an empirical Monte-Carlo correction to the bandpowers, calculated as the ratio of the mean $\hat{C}_L^{T\phi}$ across the simulation set to the FFP10 simulation input. This is at most a 4-5% correction in a couple of bins, and generally a tiny fraction of the error bars. The approximate Gaussian covariance formula $(C_L^{\phi\phi,\text{fid}} + N_L^{(0)})C_L^{TT,\text{fid}}/(2L+1)$ matches the empirical variance shape as a function of L , and we use it to inverse-weight $\hat{C}_L^{T\phi}$ and to build an estimate of the detection significance. We get

$$\text{fid } \hat{A}^{T\phi} = 1.01 \pm 0.25, \quad (3.9)$$

non-zero at just about 4σ . This is a bit higher than the value quoted in ref. [17] (just above 3σ), from a combination of 10% tighter error bars and a slightly higher recovered amplitude. The same analysis on the PR3 maps give $\text{fid } \hat{A}^{T\phi} = 0.94 \pm 0.30$. Finally, we note that $C_L^{E\phi}$ remains undetectable by a large margin and we see $\hat{C}_L^{E\phi}$ perfectly consistent with zero.

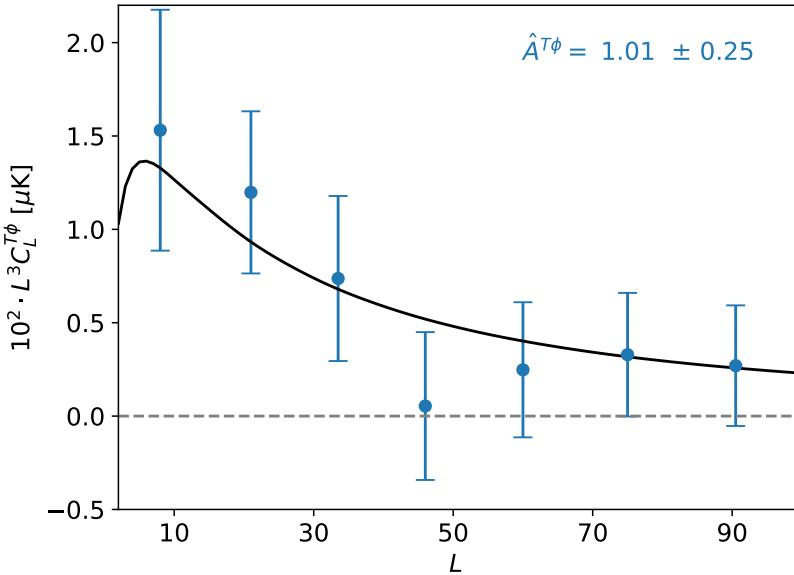


Figure 8. Estimate of the ISW lensing bispectrum $\hat{C}_L^{T\phi}$, on the multipole range $3 \leq L \leq 100$, from the cross-spectrum of the homogeneously-filtered MV NPIPE lensing reconstruction with the filtered temperature map. The black line shows the prediction in the fiducial cosmology. The detection significance is formally just above 4σ .

3.4 Parameter constraints

We use the Markov Chain Monte Carlo sampling method with Cobaya⁷ [38] to estimate cosmological parameters from the calculated likelihoods, and analyse the resulting chains using GetDist⁸ [39]. We use our baseline GMV, κ -filtered NPIPE CMB-marginalized likelihood (labelled PR4) with and without Baryon Accoustic Oscillations (BAO) likelihoods, and also the likelihood of the NPIPE analysis without applying the κ filter with PL2018 priors. We generally follow the parameter definitions and priors of PL2018, but use updated BAO likelihoods and an updated baryon density ($\Omega_b h^2$) prior based on works that were published since PL2018, as described in tables 3 and 4. The resulting base Λ CDM model parameter constraints, which we now discuss, are shown in these tables and figure 9.

As shown in the upper panel of figure 9, in the base Λ CDM model the CMB lensing-only parameter degeneracies follow a narrow tube in the space spanned by σ_8 - H_0 - Ω_m (see the detailed discussion of ref. [17] on the dependence of the lensing on cosmological parameters). This projects into a tight constraint on $\sigma_8 \Omega_m^{0.25}$, as displayed in the lower panel. Our updated bandpowers induce a slight shift upward in the parameter mean compared to PR3, and a reduction in the 68% confidence level, giving

$$\sigma_8 \Omega_m^{0.25} = 0.599 \pm 0.016 \quad (\text{PR4}), \quad (3.10)$$

where we had 0.589 ± 0.020 for PR3. The upward shift of the mean is consistent with the slightly *lower* spectrum peak seen in PR4: As can be seen in figure 3, for our noise levels the spectrum is best constrained just after the peak, so that the parameter fit is tightly constrained to the data there, in contrast to the other points. This, combined with the (more

⁷<https://github.com/CobayaSampler/cobaya>.

⁸<https://github.com/cmbant/getdist>.

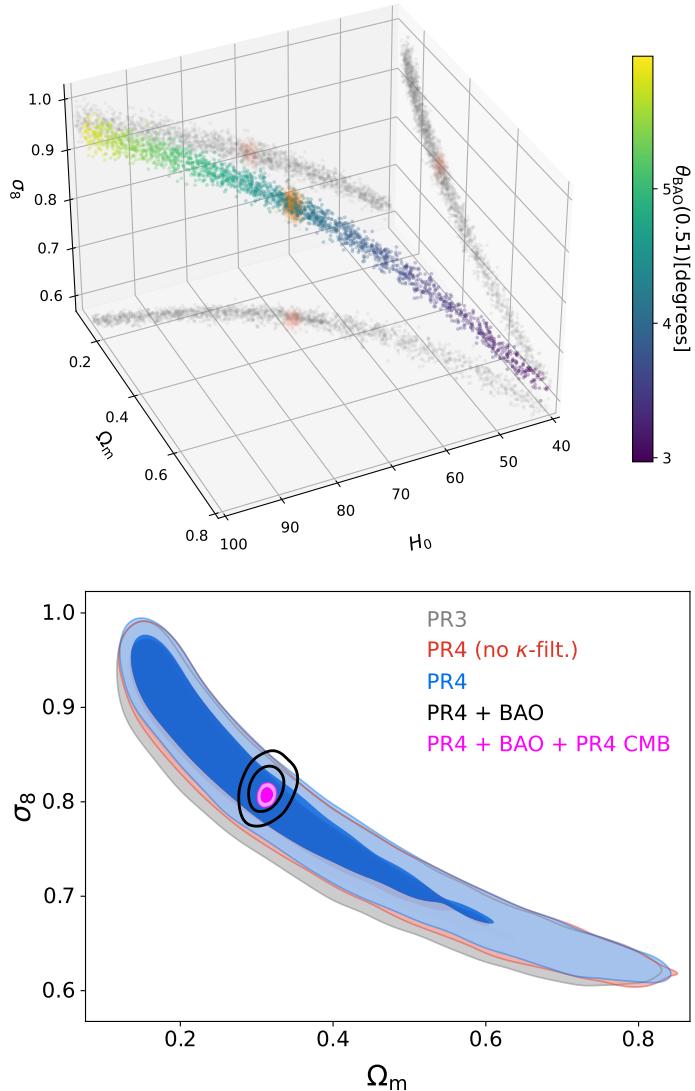


Figure 9. *Upper panel:* Posterior samples in H_0 - σ_8 - Ω_m space from our lensing-only likelihood, colour-coded by the angular size of the BAO scale $\theta_{\text{BAO}}(z) \equiv r_d/D_M(z)$ at redshift ~ 0.5 , where r_d and D_M are the comoving sound horizon and comoving distance respectively. The H_0 values are in $\text{km s}^{-1} \text{Mpc}^{-1}$. The grey points show the projection into the 2D subspaces, and the orange points the posterior samples when combined with SDSS BAO data.

Lower panel: Constraints in the σ_8 - Ω_m plane from our lensing-only likelihood, with and without κ -filtering (blue and red respectively) and in comparison to [PL2018](#) (grey). The black contours show the constraints combining our lensing likelihood with BAO data, and the magenta contours also include the recent PR4 based CMB spectra likelihood. Summary statistics are given in tables 3 and 4, and throughout the text.

intuitively, positive) correlation between the amplitude of the tail and $\sigma_8 \Omega_m^{0.25}$, produces an anti-correlation between the peak and tail. This anti-correlation is the same as that of the deflection amplitude (sourced disproportionately by the peak) to $\sigma_8 \Omega_m^{0.25}$, studied in detail in ref. [17] [see figures 8 and 9 in particular]. In order to directly confirm this, we ran a chain

Lensing-only likelihood	$\sigma_8 \Omega_m^{0.25}$
PR3	0.589 ± 0.020
PR4	0.599 ± 0.016
PR4, no κ -filtering	0.596 ± 0.017
PR4, polarization-only	0.618 ± 0.076

Table 3. Constraints on the CMB lensing parameter for different lensing datasets. The first row shows values from [PL2018](#) for comparison. We set $\Omega_b h^2 = 0.02233 \pm 0.00036$ as a tight prior on the baryon density from recent element abundances and nucleosynthesis (BBN) modelling [\[40, 41\]](#), a Gaussian prior on $n_s = 0.96 \pm 0.02$, and the Hubble constant parameter is restricted to $0.4 < h < 1$. See table 4 for results combined with BAO.

Lensing + BAO	σ_8	H_0	Ω_m
PR3	0.811 ± 0.019	$67.9_{-1.3}^{+1.2}$	$0.303_{-0.018}^{+0.016}$
PR4	0.814 ± 0.016	$68.14_{-1.10}^{+0.99}$	$0.313_{-0.016}^{+0.014}$
PR4, no κ -filtering	0.810 ± 0.016	$68.2_{-1.2}^{+1.0}$	$0.313_{-0.017}^{+0.015}$
PR4 (PR3 priors & BAO)	0.816 ± 0.017	$68.4_{-1.3}^{+1.0}$	$0.309_{-0.017}^{+0.014}$
PR4, polarization-only	$0.814_{-0.046}^{+0.051}$	$71.0_{-2.9}^{+2.1}$	$0.363_{-0.047}^{+0.036}$

Table 4. Constraints on the main lensing-related Λ CDM cosmological parameters for different lensing datasets, when combining BAO data with the lensing likelihood. The first row shows values from [PL2018](#) for comparison. The BAO likelihoods for PR4 use the 6dF Galaxy Survey measurements [\[42\]](#), the SDSS Data Release (DR) 7 measurements [\[43\]](#) and the angular and radial DR12+DR16 LRG data from eBOSS [\[44\]](#), with the exception of the fourth entry, which uses the same BAO likelihoods and priors of PR3. We use the same priors described in table 3 (except again for the fourth entry), and H_0 is in units of $\text{km s}^{-1} \text{ Mpc}^{-1}$.

with the first PR4 bin replaced by its PR3 value, which resulted in a 5 times smaller shift in $\sigma_8 \Omega_m^{0.25}$, as expected.

The points in the upper panel of figure 9 are colour-coded by the BAO angular scale at redshift $z \sim 0.5$ (the BBN prior on the baryon density that we use means that the comoving sound horizon can be predicted accurately). A measurement of the projected sound horizon from BAO measurements at low redshift therefore breaks the three-way degeneracy, without using data from CMB spectra, giving tight constraints on the parameters individually. The constraints are shown in table 4. We find a similar decrease in errors bars compared to PR3 as above, with very similar H_0 and σ_8 central values, and the entirety of the small data shift being seemingly transferred to a higher Ω_m value. However, the fourth row of the table shows

that at least part of this shift is due to the updated priors and BAO likelihoods. In the lensing + BAO chains of table 4 the estimated galaxy weak-lensing parameter $S_8 \equiv \sigma_8 (\Omega_m/0.3)^{0.5}$ central value is shifted upwards from 0.815 ± 0.036 (PR3) to 0.831 ± 0.029 (PR4), with a very similar increase seen in S_8 from PR3 to PR4 using PR3 priors & BAO.

Ref. [45] recently extended the *Planck CamSpec* high- ℓ CMB spectra likelihood to PR4, with overall consistency to PR3 and $\sim 10\%$ more constraining power compared to a comparable analysis with PR3. It also shifts the beyond Λ CDM parameters A_L and Ω_K closer to their Λ CDM values (these two parameters have a tendency to peak slightly away from Λ CDM values without the inclusion of lensing data in PR3). Combining this new likelihood with lensing and BAO (while dropping the informative priors on $\Omega_b h^2$ and n_s , and adding the PR3 low- ℓ TT and EE likelihoods), gives the magenta confidence contours in figure 9. Both σ_8 and H_0 are then constrained well below the percent level,

$$\left. \begin{array}{l} \sigma_8 = 0.8072 \pm 0.0054 \\ H_0 = 67.41 \pm 0.39 \\ \Omega_m = 0.3130 \pm 0.0053 \end{array} \right\} \text{PR4 + BAO + PR4 CMB}, \quad (3.11)$$

where H_0 is in $\text{km s}^{-1} \text{Mpc}^{-1}$. The central values are slightly shifted but consistent with the 2018 *Planck* release results, with errors that are $\sim 15\text{--}25\%$ tighter (largely driven by the 10% improvement from the PR4 high- ℓ CMB [45] and the additional sky area used by the *CamSpec* likelihood compared to the original 2018 release [46]).

We also considered extending Λ CDM to include a varying total neutrino mass, but the slightly improved lensing spectrum is not able to make much of a difference to constraints from the CMB spectra + BAO alone. We find

$$\underbrace{\sum m_\nu < 0.133 \text{ eV} \text{ (95\% CL upper bound)}}_{\text{PR4 + BAO + PR4 CMB}}. \quad (3.12)$$

With BAO and lensing only (PR4 + BAO), the constraint is obviously much weaker, finding $\sum m_\nu < 2.23 \text{ eV}$, (95% CL upper bound).

Finally, tables 3 and 4 also give the constraints using our lensing likelihood built from polarization only, showing consistency with of course larger errors. These errors are factors of 2 to 3 smaller than those achievable on this same data without inhomogeneous filtering and the improvements discussed in this work.

4 Conclusions

We have presented an updated CMB lensing reconstruction from *Planck* data, performed on the PR4 (NPIPE) maps. In part due to the lower noise in the CMB maps, and in part due to a small series of lensing pipeline improvements, error bars on the bandpowers have decreased, and constraint on the main lensing parameter has tightened to $\sigma_8 \Omega_m^{0.25} = 0.599 \pm 0.016$ from 0.589 ± 0.020 , with an upward shift in mean value caused by a slightly lower lensing spectrum peak amplitude. This upward shift in lensing parameter is not unexpected given the change in error bars [47], and its anti-correlation with the peak height for *Planck* noise levels. The shift in bandpowers at the lensing spectrum peak is more substantial (1σ). From the available simulations of PR3 and PR4, it seems difficult to assess precisely the true significance of this peak shift, but we argued it is plausibly unexpected at the 3σ level given only statistical

errors. At the lowest lensing multipoles, PR4 now gives robust results where PR3 had an anomalous-looking quadrupole.

Apart from the 1σ peak shift, our lensing results are consistent with [PL2018](#). In particular, the broad region of low lensing curl power at $264 \leq L \leq 901$ that was a potential source of worry in PR3, with a quoted significance of 2.9σ (4.2σ without consideration of look-elsewhere effects), is still very much there. As in [PL2018](#), for our likelihood analysis we restricted the bandpowers to a conservative multipole range that excludes much of this range.

We have enhanced the PR3 lensing reconstruction and likelihood pipelines to account for the inhomogeneity of the lensing reconstruction noise across the sky, and improved our joint temperature and polarization filtering scheme. These novel more-optimal techniques account for about half of our improvement in error bars. Apart from the polarization-only reconstruction, which is greatly improved but still provides weak constraints (a 7% constraint on the lensing power amplitude) compared to temperature, the impact of the analysis improvements is relatively modest for *Planck* data. It could be substantially higher for future experiments, particularly for ones with inhomogeneous scan strategies.

Our new lensing likelihoods with and without CMB-marginalization are available in a Cobaya-friendly format at https://github.com/carronj/planck_PR4_lensing. Maps and simulations are available at NERSC as indicated at this same page. Other lensing map variations can be made available upon request to the first author.

Acknowledgments

JC acknowledges support from a SNSF Eccellenza Professorial Fellowship (No. 186879), and AL support by the U.K. STFC grant ST/T000473/1. We thank Reijo Keskitalo for exchanges and help regarding the PR4 maps. We thank Kareem Marzouk and Sebastian Belkner for related collaboration and work on the PR4 maps and SMICA pipeline [48]. We thank Erik Rosenberg for sharing the PR4 `CamSpec` likelihood [45]. We thank Anthony Challinor and more generally the entire lensing working group of Simons Observatory for discussions at various stages of this work. JC thanks Louis Legrand for useful comments on the draft. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory. Some of the results in this paper have been derived using the `healpy` package [49, 50].

References

- [1] PLANCK collaboration, *Planck 2018 results. VIII. Gravitational lensing*, *Astron. Astrophys.* **641** (2020) A8 [[arXiv:1807.06210](https://arxiv.org/abs/1807.06210)] [[INSPIRE](#)].
- [2] B.D. Sherwin et al., *Two-season Atacama Cosmology Telescope polarimeter lensing power spectrum*, *Phys. Rev. D* **95** (2017) 123529 [[arXiv:1611.09753](https://arxiv.org/abs/1611.09753)] [[INSPIRE](#)].
- [3] SPT collaboration, *Constraints on Cosmological Parameters from the 500 deg² SPTpol Lensing Power Spectrum*, *Astrophys. J.* **888** (2020) 119 [[arXiv:1910.07157](https://arxiv.org/abs/1910.07157)] [[INSPIRE](#)].
- [4] W.L.K. Wu et al., *A Measurement of the Cosmic Microwave Background Lensing Potential and Power Spectrum from 500 deg² of SPTpol Temperature and Polarization Data*, *Astrophys. J.* **884** (2019) 70 [[arXiv:1905.05777](https://arxiv.org/abs/1905.05777)] [[INSPIRE](#)].
- [5] O. Darwish et al., *The Atacama Cosmology Telescope: A CMB lensing mass map over 2100 square degrees of sky and its cross-correlation with BOSS-CMASS galaxies*, *Mon. Not. Roy. Astron. Soc.* **500** (2020) 2250 [[arXiv:2004.01139](https://arxiv.org/abs/2004.01139)] [[INSPIRE](#)].

[6] PLANCK collaboration, *Planck intermediate results. LVII. Joint Planck LFI and HFI data processing*, *Astron. Astrophys.* **643** (2020) A42 [[arXiv:2007.04997](https://arxiv.org/abs/2007.04997)] [[INSPIRE](#)].

[7] T. Okamoto and W. Hu, *CMB lensing reconstruction on the full sky*, *Phys. Rev. D* **67** (2003) 083002 [[astro-ph/0301031](https://arxiv.org/abs/astro-ph/0301031)] [[INSPIRE](#)].

[8] W. Hu and T. Okamoto, *Mass reconstruction with CMB polarization*, *Astrophys. J.* **574** (2002) 566 [[astro-ph/0111606](https://arxiv.org/abs/astro-ph/0111606)] [[INSPIRE](#)].

[9] A.S. Maniyar, Y. Ali-Haïmoud, J. Carron, A. Lewis and M.S. Madhavacheril, *Quadratic estimators for CMB weak lensing*, *Phys. Rev. D* **103** (2021) 083524 [[arXiv:2101.12193](https://arxiv.org/abs/2101.12193)] [[INSPIRE](#)].

[10] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](https://arxiv.org/abs/1808.07445)] [[INSPIRE](#)].

[11] M. Mirmelstein, J. Carron and A. Lewis, *Optimal filtering for CMB lensing reconstruction*, *Phys. Rev. D* **100** (2019) 123509 [[arXiv:1909.02653](https://arxiv.org/abs/1909.02653)] [[INSPIRE](#)].

[12] PLANCK collaboration, *Planck 2015 results. IX. Diffuse component separation: CMB maps*, *Astron. Astrophys.* **594** (2016) A9 [[arXiv:1502.05956](https://arxiv.org/abs/1502.05956)] [[INSPIRE](#)].

[13] PLANCK collaboration, *Planck 2018 results. IV. Diffuse component separation*, *Astron. Astrophys.* **641** (2020) A4 [[arXiv:1807.06208](https://arxiv.org/abs/1807.06208)] [[INSPIRE](#)].

[14] J.-F. Cardoso, M. Martin, J. Delabrouille, M. Betoule and G. Patanchon, *Component separation with flexible models. Application to the separation of astrophysical emissions*, [arXiv:0803.1814](https://arxiv.org/abs/0803.1814) [[INSPIRE](#)].

[15] R. Fernandez-Cobos, P. Vielva, R.B. Barreiro and E. Martinez-Gonzalez, *Multi-resolution internal template cleaning: An application to the Wilkinson Microwave Anisotropy Probe 7-yr polarization data*, *Mon. Not. Roy. Astron. Soc.* **420** (2012) 2162 [[arXiv:1106.2016](https://arxiv.org/abs/1106.2016)] [[INSPIRE](#)].

[16] PLANCK collaboration, *Planck 2013 results. XIV. Zodiacal emission*, *Astron. Astrophys.* **571** (2014) A14 [[arXiv:1303.5074](https://arxiv.org/abs/1303.5074)] [[INSPIRE](#)].

[17] PLANCK collaboration, *Planck 2015 results. XV. Gravitational lensing*, *Astron. Astrophys.* **594** (2016) A15 [[arXiv:1502.01591](https://arxiv.org/abs/1502.01591)] [[INSPIRE](#)].

[18] E. Hivon, S. Mottet and N. Ponthieu, *QuickPol: Fast calculation of effective beam matrices for CMB polarization*, *Astron. Astrophys.* **598** (2017) A25 [[arXiv:1608.08833](https://arxiv.org/abs/1608.08833)] [[INSPIRE](#)].

[19] PLANCK collaboration, *Planck 2018 results. V. CMB power spectra and likelihoods*, *Astron. Astrophys.* **641** (2020) A5 [[arXiv:1907.12875](https://arxiv.org/abs/1907.12875)] [[INSPIRE](#)].

[20] K.M. Smith, O. Zahn and O. Dore, *Detection of Gravitational Lensing in the Cosmic Microwave Background*, *Phys. Rev. D* **76** (2007) 043510 [[arXiv:0705.3980](https://arxiv.org/abs/0705.3980)] [[INSPIRE](#)].

[21] D. Hanson, A. Challinor, G. Efstathiou and P. Bielewicz, *CMB temperature lensing power reconstruction*, *Phys. Rev. D* **83** (2011) 043005 [[arXiv:1008.4403](https://arxiv.org/abs/1008.4403)] [[INSPIRE](#)].

[22] M.H. Kesden, A. Cooray and M. Kamionkowski, *Lensing reconstruction with CMB temperature and polarization*, *Phys. Rev. D* **67** (2003) 123507 [[astro-ph/0302536](https://arxiv.org/abs/astro-ph/0302536)] [[INSPIRE](#)].

[23] SPT collaboration, *A Measurement of the Cosmic Microwave Background Gravitational Lensing Potential from 100 Square Degrees of SPTpol Data*, *Astrophys. J.* **810** (2015) 50 [[arXiv:1412.4760](https://arxiv.org/abs/1412.4760)] [[INSPIRE](#)].

[24] V. Böhm, M. Schmittfull and B.D. Sherwin, *Bias to CMB lensing measurements from the bispectrum of large-scale structure*, *Phys. Rev. D* **94** (2016) 043519 [[arXiv:1605.01392](https://arxiv.org/abs/1605.01392)] [[INSPIRE](#)].

[25] V. Böhm, B.D. Sherwin, J. Liu, J.C. Hill, M. Schmittfull and T. Namikawa, *Effect of non-Gaussian lensing deflections on CMB lensing measurements*, *Phys. Rev. D* **98** (2018) 123510 [[arXiv:1806.01157](https://arxiv.org/abs/1806.01157)] [[INSPIRE](#)].

- [26] D. Beck, G. Fabbian and J. Errard, *Lensing Reconstruction in Post-Born Cosmic Microwave Background Weak Lensing*, *Phys. Rev. D* **98** (2018) 043512 [[arXiv:1806.01216](https://arxiv.org/abs/1806.01216)] [[INSPIRE](#)].
- [27] G. Fabbian, A. Lewis and D. Beck, *CMB lensing reconstruction biases in cross-correlation with large-scale structure probes*, *JCAP* **10** (2019) 057 [[arXiv:1906.08760](https://arxiv.org/abs/1906.08760)] [[INSPIRE](#)].
- [28] G. Pratten and A. Lewis, *Impact of post-Born lensing on the CMB*, *JCAP* **08** (2016) 047 [[arXiv:1605.05662](https://arxiv.org/abs/1605.05662)] [[INSPIRE](#)].
- [29] J.R. Stevens et al., *Designs for next generation CMB survey strategies from Chile*, *Proc. SPIE Int. Soc. Opt. Eng.* **10708** (2018) 1070841 [[arXiv:1808.05131](https://arxiv.org/abs/1808.05131)] [[INSPIRE](#)].
- [30] J. Peloton, M. Schmittfull, A. Lewis, J. Carron and O. Zahn, *Full covariance of CMB and lensing reconstruction power spectra*, *Phys. Rev. D* **95** (2017) 043508 [[arXiv:1611.01446](https://arxiv.org/abs/1611.01446)] [[INSPIRE](#)].
- [31] S.J. Osborne, D. Hanson and O. Doré, *Extragalactic Foreground Contamination in Temperature-based CMB Lens Reconstruction*, *JCAP* **03** (2014) 024 [[arXiv:1310.7547](https://arxiv.org/abs/1310.7547)] [[INSPIRE](#)].
- [32] PLANCK collaboration, *Planck 2013 results. XVII. Gravitational lensing by large-scale structure*, *Astron. Astrophys.* **571** (2014) A17 [[arXiv:1303.5077](https://arxiv.org/abs/1303.5077)] [[INSPIRE](#)].
- [33] PLANCK collaboration, *Planck 2018 results. VI. Cosmological parameters*, *Astron. Astrophys.* **641** (2020) A6 [*Erratum ibid.* **652** (2021) C4] [[arXiv:1807.06209](https://arxiv.org/abs/1807.06209)] [[INSPIRE](#)].
- [34] T. Namikawa, D. Hanson and R. Takahashi, *Bias-Hardened CMB Lensing*, *Mon. Not. Roy. Astron. Soc.* **431** (2013) 609 [[arXiv:1209.0091](https://arxiv.org/abs/1209.0091)] [[INSPIRE](#)].
- [35] M. Zaldarriaga, *Lensing of the CMB: Non-Gaussian aspects*, *Phys. Rev. D* **62** (2000) 063510 [[astro-ph/9910498](https://arxiv.org/abs/astro-ph/9910498)] [[INSPIRE](#)].
- [36] W. Hu, *Weak lensing of the CMB: A harmonic approach*, *Phys. Rev. D* **62** (2000) 043007 [[astro-ph/0001303](https://arxiv.org/abs/astro-ph/0001303)] [[INSPIRE](#)].
- [37] A. Lewis, A. Challinor and D. Hanson, *The shape of the CMB lensing bispectrum*, *JCAP* **03** (2011) 018 [[arXiv:1101.2234](https://arxiv.org/abs/1101.2234)] [[INSPIRE](#)].
- [38] J. Torrado and A. Lewis, *Cobaya: Code for Bayesian Analysis of hierarchical physical models*, *JCAP* **05** (2021) 057 [[arXiv:2005.05290](https://arxiv.org/abs/2005.05290)] [[INSPIRE](#)].
- [39] A. Lewis, *GetDist: a Python package for analysing Monte Carlo samples*, [arXiv:1910.13970](https://arxiv.org/abs/1910.13970) [[INSPIRE](#)].
- [40] R.J. Cooke, M. Pettini and C.C. Steidel, *One Percent Determination of the Primordial Deuterium Abundance*, *Astrophys. J.* **855** (2018) 102 [[arXiv:1710.11129](https://arxiv.org/abs/1710.11129)] [[INSPIRE](#)].
- [41] V. Mossa et al., *The baryon density of the Universe from an improved rate of deuterium burning*, *Nature* **587** (2020) 210 [[INSPIRE](#)].
- [42] F. Beutler et al., *The 6dF Galaxy Survey: Baryon Acoustic Oscillations and the Local Hubble Constant*, *Mon. Not. Roy. Astron. Soc.* **416** (2011) 3017 [[arXiv:1106.3366](https://arxiv.org/abs/1106.3366)] [[INSPIRE](#)].
- [43] A.J. Ross, L. Samushia, C. Howlett, W.J. Percival, A. Burden and M. Manera, *The clustering of the SDSS DR7 main Galaxy sample — I. A 4 per cent distance measure at $z = 0.15$* , *Mon. Not. Roy. Astron. Soc.* **449** (2015) 835 [[arXiv:1409.3242](https://arxiv.org/abs/1409.3242)] [[INSPIRE](#)].
- [44] EBOSS collaboration, *Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory*, *Phys. Rev. D* **103** (2021) 083533 [[arXiv:2007.08991](https://arxiv.org/abs/2007.08991)] [[INSPIRE](#)].
- [45] E. Rosenberg, S. Gratton and G. Efstathiou, *CMB power spectra and cosmological parameters from Planck PR4 with CamSpec*, [arXiv:2205.10869](https://arxiv.org/abs/2205.10869) [[INSPIRE](#)].
- [46] G. Efstathiou and S. Gratton, *A Detailed Description of the CamSpec Likelihood Pipeline and a Reanalysis of the Planck High Frequency Maps*, [arXiv:1910.00483](https://arxiv.org/abs/1910.00483) [[INSPIRE](#)].

- [47] S. Gratton and A. Challinor, *Understanding parameter differences between analyses employing nested data subsets*, *Mon. Not. Roy. Astron. Soc.* **499** (2020) 3410 [[arXiv:1911.07754](https://arxiv.org/abs/1911.07754)] [[INSPIRE](#)].
- [48] K. Marzouk, A. Lewis and J. Carron, *Constraints on τ_{NL} from Planck temperature and polarization*, *JCAP* **08** (2022) 015 [[arXiv:2205.14408](https://arxiv.org/abs/2205.14408)] [[INSPIRE](#)].
- [49] A. Zonca et al., *healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python*, *J. Open Source Softw.* **4** (2019) 1298.
- [50] K.M. Górski et al., *HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759 [[astro-ph/0409513](https://arxiv.org/abs/astro-ph/0409513)] [[INSPIRE](#)].