Choice of Planck CMB likelihood in cosmological analyses

Hidde T. Jense,[1,*] Marc Viña,[1] Erminia Calabrese,[1] and J. Colin Hill[2]

[1]*School of Physics and Astronomy, Cardiff University, The Parade, CF24 3AA Cardiff, Wales, United Kingdom*

[2]*Department of Physics, Columbia University, New York, New York 10027, USA*

(Received 14 November 2025; accepted 26 January 2026; published 26 February 2026)

We compare cosmological parameters from different *Planck* sky maps and likelihood pipelines, assessing the robustness of cosmological results with respect to the choice of the latest *Planck* maps-likelihood combination. We show that, for the *Planck* multipole range retained in combination with ground-based observations, different products give very similar cosmological solutions; small remaining differences are reduced by the addition of other CMB datasets to *Planck*. In particular, constraints on extended cosmological models benefit from the addition of small-scale power from ground-based experiments and are completely insensitive to the choice of *Planck* maps and likelihood. For this work, we derive and release a nuisance-marginalized dataset and `CamSpec-NPIPE-lite` likelihood for the *Planck* NPIPE data injected into the `CamSpec` likelihood—which are usually used to obtain the reference *Planck* PR4 cosmology. Using the extracted CMB spectra, we show that the additional constraining power for cosmology is coming from polarization at all scales and from temperature at multipoles above 1500 when going from PR3 to PR4. We also show that full marginalization over the `CamSpec` foreground nuisance parameters can impact parameter inference and model selections when truncating some scales; our new likelihood enables correct combinations with other CMB datasets.

DOI: 10.1103/ff9j-skyj

I. INTRODUCTION

The cosmic microwave background (CMB) has been a primary source of constraints on cosmological models for decades. Until recently, the most stringent limits on the parameters of the standard cosmological model, ΛCDM, and its extensions were set using observations of CMB temperature and polarization anisotropies from the *Planck* satellite mission [1–5]. Constraints on cosmological parameters similar to or better than those obtained with *Planck* are now becoming possible with data from the ground-based Atacama Cosmology Telescope [ACT [6–8]] and South Pole Telescope [SPT [9]], exploiting information encoded in the CMB at intermediate and small scales—multipoles between $400 \lesssim \ell \lesssim 6500$. However, *Planck* remains the dominant dataset in tracing the behavior of the very large scales, at $2 < \ell < 1600$. This range of observations from *Planck* is currently used in combination with ACT and SPT to derive new leading cosmological results and will continue to supplement large-scale data to future ground-based experiments until a new CMB satellite becomes operational.

The *Planck* products available to the community for cosmological exploitation are based on two sets of sky maps. The "Legacy" maps [1] are processed through the `Plik` likelihood described in [2] (PL20 hereafter) with cosmology results reported in [3] (PC20 hereafter). The Legacy maps are also analyzed in [10] (EG21 hereafter) with the alternative `CamSpec` likelihood pipeline. The "NPIPE" maps [11] were generated after introductions of several data-processing improvements in the handling of large-scale systematics subsequent to the final *Planck* legacy release. The reduction in systematics allowed a small increase in the sky fraction used for cosmology, and revised cosmological limits from these maps have been presented in [4] (R22 hereafter) with the `CamSpec` likelihood pipeline and in [5] with the `HiLLiPoP` likelihood.

The two sets of maps define the *Planck* release names, with Legacy corresponding to the *Planck* data Release 3, PR3, and NPIPE maps identified as PR4. However, the reference cosmology for each data release relies also on the choice of the likelihood. PR3 cosmology is then usually referring to the Legacy maps processed with `Plik`, and PR4 cosmology usually points to NPIPE maps processed with `CamSpec`. A summary of these products is presented in Table I.

*Contact author: jenseh@cardiff.ac.uk

TABLE I. Summary of *Planck* products available for cosmological exploitation, making explicit the choice of sky maps and likelihood pipelines. PR3 cosmology is usually drawn from the final legacy release of *Planck*, using the Legacy maps processed with the `Plik` likelihood; PR4 cosmology is usually drawn from the analysis of NPIPE maps processed with the `CamSpec` likelihood.

Planck products	Likelihood	Maps	Reference	Known as
`Plik-Legacy`	`Plik`	Legacy	*Planck* Collaboration [1–3]	PR3 reference cosmology
`CamSpec-Legacy`	`CamSpec`	Legacy	Efstathiou & Gratton [10]	
`CamSpec-NPIPE`	`CamSpec`	`NPIPE`	Rosenberg *et al.* [4]	PR4 reference cosmology
`HiLLiPoP-NPIPE`	`HiLLiPoP`	`NPIPE`	Tristram *et al.* [5]	

In this paper, we compare and discuss differences between the reference PR3 and PR4 cosmological results, using the Legacy maps in input to the `Plik` likelihood code (hereafter labeled `Plik-Legacy`) and the NPIPE maps in input to the `CamSpec` likelihood (hereafter labeled `CamSpec-NPIPE`). While both sets of maps have been processed with the `CamSpec` likelihood, the details of the likelihood method applied to Legacy and NPIPE maps in `CamSpec` differ (see EG21 and R22), and a `Plik-NPIPE` likelihood is not available. When using the full dataset, the cosmologies retrieved from the `Plik-Legacy` and `CamSpec-NPIPE` likelihoods are similar but show order of 1σ movements in cosmological parameters, as discussed in R22. Some of these shifts are due to different choices in data characterization made between `Plik` and `CamSpec` and present even when using the same Legacy maps—see discussion in PL20, PC20, and EG21. Additionally, `CamSpec-NPIPE` provides more stringent constraints on some cosmological parameters, due to the increased sky fraction retained in the NPIPE maps (R22) corresponding to ~10% more data. Here we extract and use the CMB signal-to-noise in temperature and polarization as a metric to show where the additional constraining power for cosmology in `CamSpec-NPIPE` is localized. We also add to previous comparisons, focusing in particular on the restricted *Planck* multipole range used in combined analyses with ACT and SPT—to, e.g., form the "`P-ACT`" combination of [7] and "CMB-SPA" combination of [9]—to assess the robustness of new, leading constraints on cosmological models with respect to the choice of *Planck* likelihoods and maps.

To perform these studies, we work with foreground and nuisance-marginalized likelihoods, deriving and releasing a new `CamSpec-NPIPE-lite` likelihood. We show that attention to marginalization over nuisances is required when truncating the `CamSpec-NPIPE` likelihood and is necessary for likelihood combinations.

The paper is organized as follows. In Sec. II, we summarize the main aspects of the `Plik` and `CamSpec` likelihoods and present a new, foreground-marginalized `CamSpec-NPIPE` likelihood; Sec. III compares and discusses the constraints on cosmological parameters from various likelihoods. Finally, in Sec. IV, we derive and discuss cosmological constraints from the combination of the various *Planck* likelihoods with small-scale measurements from ACT DR6. We conclude in Sec. V.

II. *PLANCK* LIKELIHOODS

A. Likelihood data and model

Both `Plik` and `CamSpec` likelihoods take in input versions of cross-frequency angular power spectra from the *Planck*'s High Frequency Instrument (HFI), together with a covariance matrix capturing the instrumental noise properties of the multi-frequency maps. The two likelihoods differ in the amount and format of data, as well as in the detailed modeling of the spectra—as explained in [12], PL20, EG21, and R22. In this work, we do not alter any preprocessing and modeling done before or at likelihood level. To compare and discuss cosmological results between different likelihood packages, we rely on the underlying CMB-only power in the data. To extract this, we simply need to add an additional processing step of the multi-frequency data and operate on the foregrounds and systematics components accounted for in the likelihood. In what follows, we then summarize only the main aspects relevant for this work and refer the reader to PL20, EG21, and R22 for the full details.

All likelihood evaluations in this work are performed within the Simons Observatory[1] `MFLike` framework,[2] which provides consistent handling of input spectra, covariances, and marginalization procedures. A version of `Plik-Legacy` in this framework was presented in [13][3]; the equivalent for `CamSpec-NPIPE` is built in this work [14].

1. `Plik-Legacy`

The `Plik` likelihood uses a total of 16 auto- and cross-spectra in TT, TE, and EE constructed from T and E Legacy half-mission maps at frequencies 100, 143, and 217 GHz. The spectra only cover the range of multipoles over which

[1]https://simonsobservatory.org/.
[2]https://github.com/simonsobs/LAT_MFLike.
[3]Available at https://github.com/simonsobs/LAT_MFLike/tree/mflike-plik.

the likelihood of the CMB power spectra can be well-approximated as Gaussian and span multipoles $30 < \ell < 2508$ in temperature and $30 < \ell < 1996$ in polarization—the upper and lower multipoles are different across spectra, and the number of cross-frequency spectra is different between temperature and TE/EE (PL20). The spectra are binned for a total of 613 datapoints in the full temperature plus polarization dataset.

Before computing spectra, the maps are masked to clean for Galactic dust and bright point sources in temperature, and only Galactic dust in polarization. This leaves the likelihood to include Galactic and extragalactic foreground emission. The model of foregrounds in temperature scales across frequencies templates for the thermal and kinetic Sunyaev–Zel'dovich effects (tSZ and kSZ), the cosmic infrared background (CIB), and tSZxCIB correlations, and solves for unresolved point sources as a different varying shot noise term in each cross-spectrum. Thermal dust emission from the Galaxy is estimated by cross-correlating each map individually with the 545 GHz map, building a template for each cross-spectrum that takes into account the correct sky fraction and frequency-dependent point source mask that differs per map. The resulting template is then only left to scale via an amplitude, which is constrained with a Gaussian prior during parameter inference. In polarization, extragalactic emission is limited to a point source contribution, which is negligible for the sensitivity of *Planck*. Polarized dust emission is estimated similarly to the thermal dust emission, except that the cross-correlation is done with the 353 GHz channel (which is the highest *Planck* frequency channel that measures polarization), and the template used is a simple power law with a fixed spectral index and an amplitude that is free to vary only for the TE cross-spectrum. The EE spectra are not very sensitive to the value of the dust amplitude, and these are left fixed to the value found from the cross-correlation with the 353 GHz channel.

Residual systematics such as beam leakage, subpixel effects, and correlated noise corrections are also accounted for at likelihood level with fixed frequency-dependent templates.

The `Plik` likelihood function comes down to a Gaussian likelihood:

$$-2\ln\mathcal{L}(\theta) = (\mathcal{D}_b^{\rm th}(\theta) - \mathcal{D}_b^{\rm d})\mathbf{\Sigma}^{-1}(\mathcal{D}_b^{\rm th}(\theta) - \mathcal{D}_b^{\rm d})^T, \quad (1)$$

where for each cross-spectrum XY covering TT, TE, and EE, $\mathcal{D}_b^{\rm d}$ is the binned data vector, $\mathbf{\Sigma}$ is the covariance matrix, and $\mathcal{D}_b^{\rm th}(\theta)$ is the binned theory model defined by some parameters θ, built from the components

$$\mathcal{D}_b^{{\rm th},XY}(\theta) = \sum_\ell \mathbf{w}_b^\ell c_{\nu_1}^X c_{\nu_2}^Y \Big(\mathcal{D}_\ell^{{\rm CMB},XY}(\theta_1) + \mathcal{D}_{\ell,\nu_1,\nu_2}^{{\rm fg},XY}(\theta_2) + \mathcal{D}_\ell^{{\rm sys},XY}\Big), \quad (2)$$

where $\mathbf{w}_b^\ell$ is the binning matrix used by `Plik`, $\mathcal{D}_\ell^{\rm CMB}$ is the CMB theory model, $\mathcal{D}_{\ell,\nu_1,\nu_2}^{\rm fg}$ is the foreground model at cross-frequency $\nu_1 \times \nu_2$, and $\mathcal{D}_\ell^{\rm sys}$ is the residual systematic correction. The final data vector contains binned data, where each bin b is specific to a cross-spectrum XY, frequencies ν_1, ν_2, and ℓ range. Here, we divide the full parameter vector θ between the cosmological parameters θ_1 and a set of secondary emission parameters θ_2 accounting for foregrounds. The components of the theory vector are corrected by factors $c_{\nu_1}^X$, $c_{\nu_2}^Y$, which contain all relevant calibration and polarization efficiency factors for each cross-spectrum—with some of them left as free parameters in the likelihoods and others fixed to measured quantities. `Plik-Legacy` defines the calibration relative to the 143 GHz temperature map, leaving five free parameters: two for the calibration of the 100 and 217 GHz maps relative to this, and three for the polarization efficiencies of the polarization maps. The 143 GHz map is calibrated off the dipole measurement, and the uncertainty in this measurement is captured in a total calibration parameter $A_{\rm Planck}$, which linearly scales all spectra. The full posterior mode also contains contributions from parameter priors, most of which are wide, uniform priors with little information for parameters that are well-constrained by the data. Gaussian priors are imposed on the Galactic dust emission amplitudes from the measurement at 545 GHz, and the overall calibration factor $A_{\rm Planck} = (100 \pm 0.25)\%$ propagating the uncertainty in the dipole calibration.

The official `Plik-Legacy` product is also released in a compressed form (PL20), known as "`Plik-lite`" and labeled throughout here as `Plik-Legacy-lite`, which marginalizes over foreground and calibration parameters and systematics templates to produce a reduced data vector that retains the CMB-only cosmological constraining power while improving computational efficiency.

2. *CamSpec-NPIPE*

The `CamSpec-NPIPE` likelihood is constructed from five unbinned angular power spectra: three TT cross-spectra (143×143, 143×217, and 217×217 GHz), and coadded TE and EE spectra built from the `NPIPE` maps. The 143×143 GHz TT spectrum as well as the TE and EE spectra span the multipole range $30 < \ell < 2000$, while the 143×217 and 217×217 GHz TT spectra cover $500 < \ell < 2500$. The polarization spectra are formed by inverse-noise-variance weighting of all cross-spectra between 100, 143, and 217 GHz; and only cross-split spectra are used in the construction of the likelihood, avoiding the noise bias associated with auto-spectra. This means any likelihood spectrum is built up only of $\mathrm{A} \times \mathrm{B}$ spectra, where A and B are the different detector set splits chosen by `NPIPE`.

One of the main differences with `Plik-Legacy`—beyond the use of different sets of maps—is that significant foreground removal is done at map level and not modeled in the likelihood. This precleaning includes both masking and the use of additional frequency channels to suppress foreground contamination. For example, a relatively conservative Galactic mask is applied to the NPIPE maps, especially at low Galactic latitudes, to reduce contamination from diffuse foregrounds. This is sufficient to render postcleaning dust residuals negligible, even at the most contaminated frequency used in this likelihood (e.g., at 217 GHz). In temperature, maps are cleaned by template subtraction using higher-frequency maps dominated by thermal dust emission, and by masking bright sources. In polarization, only masking is applied. These steps reduce the residual foreground contamination to a level where explicit modeling in polarization becomes unnecessary and the spectra can be coadded across frequencies to inject in the likelihood a single CMB-only spectrum.

In temperature, residual extragalactic foreground power is still present but can no longer be modeled with specific astrophysical components (such as tSZ, CIB, etc.) because the precleaning steps have coupled the residual emission associated to different contributions. The frequency-dependent residual power is then modeled with a simple power-law form:

$$\mathcal{D}^{\mathrm{fg,TT}}_{\ell,\nu_1,\nu_2}(A_{\nu_1,\nu_2},\gamma_{\nu_1,\nu_2}) = A_{\nu_1,\nu_2}\left(\frac{\ell}{\ell_0}\right)^{\gamma_{\nu_1,\nu_2}}, \tag{3}$$

with nuisance amplitude parameters A_{ν_1,ν_2} and power-law index γ_{ν_1,ν_2} for each frequency pair. The pivot scale is fixed at $\ell_0 = 1500$. This foreground term is added to the CMB theory before correcting the full model for remaining uncertainties in calibration:

$$\mathcal{D}^{\mathrm{th,,XY}}_{\ell,\nu_1,\nu_2}(\theta) = \frac{1}{C^{XY}}\left(\mathcal{D}^{\mathrm{CMB,XY}}_{\ell}(\theta_1) + \mathcal{D}^{\mathrm{fg,XY}}_{\ell,\nu_1,\nu_2}(\theta_2)\right). \tag{4}$$

`CamSpec-NPIPE` does this with three parameters capturing small residual mismatches in gain and polarization efficiency: C^{TE} and C^{EE} allow for rescaling of the TE and EE spectra, respectively, and A_{Planck}, a global calibration amplitude, is applied to all spectra. Like for `Plik`, this model is compared to the data in a Gaussian likelihood.

We implement all components of the `CamSpec-NPIPE` likelihood in the `MFLike` framework and check that our implementation reproduces the results of R22 as shown in later sections. From this, we derive and release the equivalent of `Plik-Legacy-lite`, `CamSpec-NPIPE-lite`, as described in the next section.

B. A foreground-marginalized `CamSpec-NPIPE` likelihood

We employ the Gibbs sampling method developed before for ACT, *Planck*, and SPT [2,7,9,12,15–17] to extract the CMB-only power from the multi-frequency `CamSpec-NPIPE` dataset.

For the theory vector of Eq. (4), the Gibbs sampling method builds a Gaussian estimator of the CMB-only bandpowers $\mathcal{D}^{\mathrm{CMB,XY}}_{\ell}$, marginalized over the secondary parameters θ_2, sampled independently of any cosmological model with parameters θ_1. Gibbs sampling allows us to sample $\mathcal{D}^{\mathrm{CMB}}_{\ell}$ and θ_2 even if we do not have a good method of jointly sampling $(\mathcal{D}^{\mathrm{CMB}}_{\ell},\theta_2)$. In this case, we can easily perform a Gaussian sampling over the CMB bandpowers $\mathcal{D}^{\mathrm{CMB}}_{\ell}|\theta_2$, and we can use the Metropolis-Hastings algorithm to sample the nuisance parameters $\theta_2|\mathcal{D}^{\mathrm{CMB}}_{\ell}$. To do this, we construct the *mapping matrix* $\mathbf{M}$, a rectangular matrix with $M_{ij} = 1$ if the ith entry of the CMB vector should contain the bandpowers contained in the jth entry of the data vector, and 0 otherwise. If we have a sample $\mathcal{D}^{\mathrm{CMB}}_{\ell}$, our full model vector becomes

$$\mathcal{D}_{\ell,\nu_1,\nu_2}(\theta) = \mathbf{M}\mathcal{D}^{\mathrm{CMB}}_{\ell} + \mathcal{D}^{\mathrm{sec}}_{\ell,\nu_1,\nu_2}(\theta_2), \tag{5}$$

where instead of parametrizing $\mathcal{D}^{\mathrm{CMB}}_{\ell}(\theta_1)$ with some Einstein-Boltzmann code based on cosmological parameters θ_1, we now use our Gaussian sample for the CMB bandpowers. Then, our full likelihood becomes

$$\begin{aligned}-2\ln\mathcal{L}(\mathcal{D}^{\mathrm{CMB}}_{\ell},\theta_2) = {}&(\mathbf{M}\mathcal{D}^{\mathrm{CMB}}_{\ell} + \mathcal{D}^{\mathrm{sec}}_{\ell,\nu_1,\nu_2}(\theta_2) - \hat{\mathcal{D}}_{\ell,\nu_1,\nu_2})\\ &\mathbf{\Sigma}^{-1}(\mathbf{M}\mathcal{D}^{\mathrm{CMB}}_{\ell} + \mathcal{D}^{\mathrm{sec}}_{\ell,\nu_1,\nu_2}(\theta_2) - \hat{\mathcal{D}}_{\ell,\nu_1,\nu_2})^T,\end{aligned} \tag{6}$$

with $\hat{\mathcal{D}}_{\ell,\nu_1,\nu_2}$ being our data vector.

For fixed nuisances θ_2, the conditional distribution of the CMB bandpowers, $\mathbb{P}(\mathcal{D}^{\mathrm{CMB}}_{\ell}|\theta_2,\hat{\mathcal{D}}_{\ell,\nu_1,\nu_2})$ is given by

$$\begin{aligned}&-2\ln\mathbb{P}(\mathcal{D}^{\mathrm{CMB}}_{\ell}|\theta_2,\hat{\mathcal{D}}_{\ell,\nu_1,\nu_2})\\ &= (\mathcal{D}^{\mathrm{CMB}}_{\ell} - \hat{\mathcal{D}}^{\mathrm{res}}_{\ell})^T\mathbf{Q}^{-1}(\mathcal{D}^{\mathrm{CMB}}_{\ell} - \hat{\mathcal{D}}^{\mathrm{res}}_{\ell}),\end{aligned} \tag{7}$$

where $\hat{\mathcal{D}}^{\mathrm{res}}_{\ell}$ is the residual of the data minus the nuisance components, averaged across all cross-spectra, and $\mathbf{Q}$ is the covariance of the CMB bandpowers. We can find these by taking the derivative with respect to the CMB bandpowers, assuming this is a Gaussian distribution with uniform flat priors, giving us

$$\mathbf{Q} = \mathbf{M}^T\mathbf{\Sigma}^{-1}\mathbf{M} \tag{8}$$

and

$$\hat{\mathcal{D}}^{\mathrm{res}}_{\ell} = \mathbf{Q}^{-1}[\mathbf{M}^T\mathbf{\Sigma}^{-1}(\hat{\mathcal{D}}_{\ell,\nu_1,\nu_2} - \mathcal{D}^{\mathrm{sec}}_{\ell,\nu_1,\nu_2}(\theta_2))]. \tag{9}$$

We can draw a random sample from this Gaussian distribution by taking the Cholesky decomposition of the

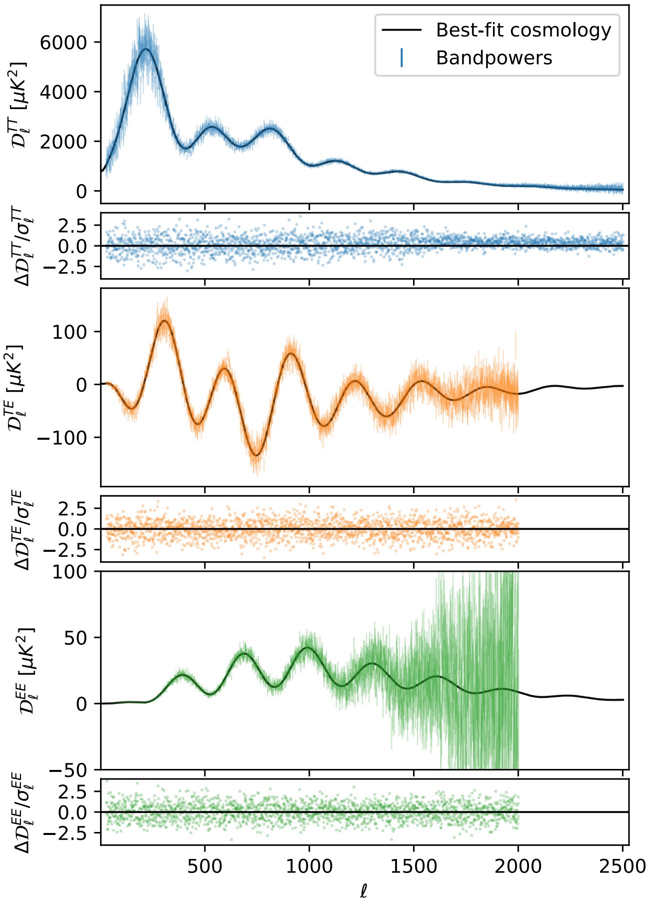

FIG. 1. CMB foreground-marginalized `CamSpec-NPIPE-lite` spectra. The TE and EE spectra appear normal distributed around the best-fitting model, as does the TT spectrum below $\ell < 1500$. At $\ell > 1500$, the TT data points are highly correlated and the full behaviour of the data is captured in the covariance matrix.

covariance matrix, $\mathbf{Q} = \mathbf{L}\mathbf{L}^T$, and drawing a random vector

$$\mathcal{D}_\ell^{\rm CMB} = \hat{\mathcal{D}}_\ell^{\rm res} + \mathbf{L}y, \tag{10}$$

where $y \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ is a vector drawn from a standard normal distribution. This method means that the vector $\mathcal{D}_\ell^{\rm CMB}$ is drawn from a multivariate normal distribution with a mean and variance imposed by the data vector and covariance matrix.

We can now estimate the nuisance-marginalized CMB bandpowers by Gibbs sampling the conditional distributions $\mathbb{P}(\mathcal{D}_\ell^{\rm CMB}|\theta_2, \hat{\mathcal{D}}_{\ell,\nu_1,\nu_2})$ and $\mathbb{P}(\theta_2|\mathcal{D}_\ell^{\rm CMB}, \hat{\mathcal{D}}_{\ell,\nu_1,\nu_2})$. Remaining for the nuisance parameters are the six temperature residual foreground parameters, three amplitudes A_{ν_1,ν_2}, and three power law indices γ_{ν_1,ν_2}, one pair for each of the three TT cross-spectra 143×143, 143×217, and 217×217. We keep the three calibration parameters $A_{\rm Planck}$, $C^{\rm TE}$, and $C^{\rm EE}$ fixed to unity in the foreground marginalization, and leave them for the likelihood to

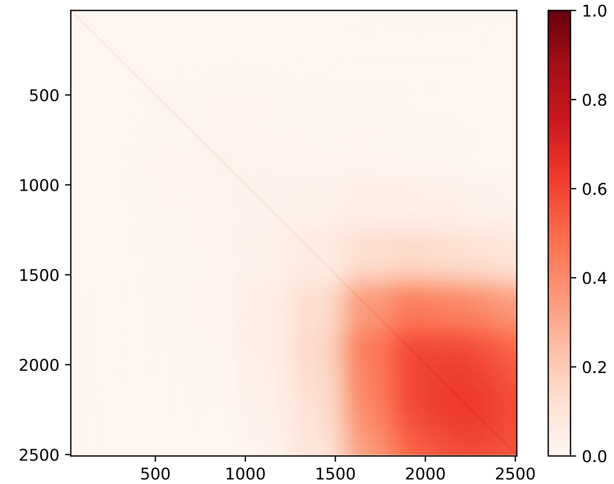

FIG. 2. TTTT subblock of the correlation matrix. The marginalization over foregrounds captures a lot of information in the $\ell > 1500$ part of the covariance, leading to high correlations between different bins.

measure, as they are fully degenerate with the amplitude of the CMB bandpowers. It is possible to modify the mapping matrix to include any multiplicative calibration parameters and marginalize over them, using a prior to constrain them to avoid the full amplitude degeneracy. We opted however not to do so, as to leave the option for external datasets to jointly marginalize over the *Planck* dipole calibration measurement.

We show the extracted power spectra from the marginalization procedure in Fig. 1. The CMB spectra are shown alongside the best-fitting ΛCDM model; a deviation from the best-fit cosmology of the TT bandpowers is visible at $\ell > 1500$ and caused by the foreground-induced correlation between the data points which is fully captured in the CMB covariance matrix [similarly to the ACT DR6 small-scale temperature, see Fig. 34 of [7]]. We show the TTTT block of the correlation matrix in Fig. 2; the remaining blocks are mostly zero or diagonal, as there is no residual TE/EE foregrounds, nor is there much important cross-spectrum covariance.

In previous applications, this marginalization procedure relied on the possibility to disentangle (most) frequency-dependent effects from the black-body CMB signal—i.e., the amplitude of a given foreground component would increase or decrease with frequency while the CMB amplitude would remain constant. In the `CamSpec-NPIPE` model, this is no longer possible because the foreground residual power is distinct across frequencies. We find that the method still works and is able to lock the CMB signal across spectra, but we lose constraining power on the foreground nuisance parameters with respect to the full likelihood. As shown in Fig. 3, we recover consistent but broader posteriors for the foregrounds (similarly to the Poisson source terms in the case of `Plik-Legacy`, PL20). Figure 3 also shows that, different from the full

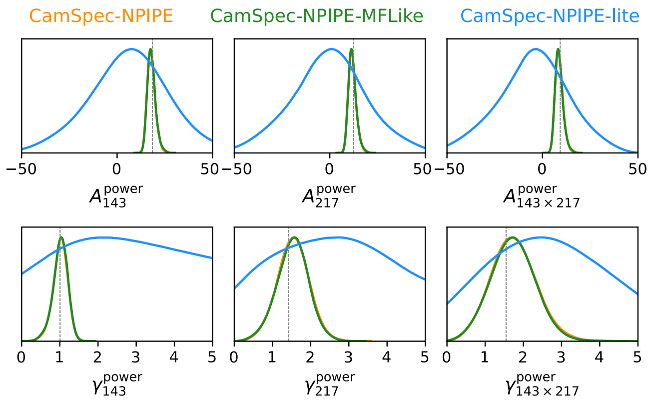

FIG. 3. Posterior distributions of the foreground nuisance parameters from the multi-frequency likelihood (orange for the original `CamSpec-NPIPE` likelihood and green for our implementation in `MFLike`), and from the foreground marginalization procedure (blue). Due to the fact that each foreground appears in only one spectrum each, they are fully degenerate in the foreground marginalization procedure, and thus highly unconstrained. Orange and green posteriors are obtained with a fit to the full *Planck* TT/TE/EE + lowT + lowE dataset but receive contribution only from the high-ℓ `CamSpec-NPIPE` likelihood. The vertical lines indicate the best-fitting values from the `CamSpec-NPIPE` chain (original likelihood, shown in orange).

likelihood, for the foreground marginalization procedure, we let the amplitude of parameters go negative. We found that opening up the amplitude parameters like this, while not necessarily physical, improved the agreement between the two likelihoods. We explain this by the fact that we model the likelihood as a Gaussian, and without the frequency information the amplitudes will broaden in such a way that imposing a positivity prior truncates one tail of the posterior mode.

After marginalization, we take the resulting TT, TE, and EE samples and compute the multivariate Gaussian mean $\hat{\mathcal{D}}_\ell$ and covariance $\mathbf{\Sigma}$ of these samples. We built a simple Gaussian likelihood, `CamSpec-NPIPE-lite` [18], as

$$-2\ln\mathcal{L}(\mathcal{D}_\ell^{\rm CMB}, c^{XY}) = (\mathcal{D}_\ell^{\rm CMB}/c^{XY} - \hat{\mathcal{D}}_\ell)\,\mathbf{\Sigma}^{-1}(\mathcal{D}_\ell^{\rm CMB}/c^{XY} - \hat{\mathcal{D}}_\ell)^T, \quad (11)$$

where c^{XY} is a set of calibration parameters for the XY spectra: $c^{TT} = A_{\rm Planck}{}^2$, $c^{\rm TE} = {\rm cal}^{\rm TE} A_{\rm Planck}{}^2$, $c^{\rm EE} = {\rm cal}^{\rm EE} A_{\rm Planck}{}^2$. Thus, the entire likelihood only depends on three remaining parameters $\{A_{\rm Planck}, {\rm cal}^{\rm TE}, {\rm cal}^{\rm EE}\}$ that need to be included in likelihood evaluations.

III. COSMOLOGICAL CONSTRAINTS

A. Parameter run setup

We implement our likelihoods in Cobaya [19,20], and we use it to perform cosmological inference with Monte Carlo Markov chain runs. We use the Einstein-Boltzmann solver code CAMB for generating the lensed CMB power spectrum with the accuracy settings described in Calabrese *et al.* [6], even when only using *Planck* data, for consistency. We run our chains to a convergence criterion of $R - 1 < 0.01$, discarding 30% of our samples as a burn-in fraction.

As a baseline comparison, we derive constraints on the ΛCDM model parameters from the official `Plik-Legacy` and `CamSpec-NPIPE` likelihoods (as implemented in Cobaya), from our `MFLike` implementations, as well as our `CamSpec-NPIPE-lite` likelihood. We vary the standard set of six base ΛCDM parameters: the baryon and cold dark matter densities, $\Omega_b h^2$ and $\Omega_c h^2$, the amplitude and spectral index of primordial curvature scalar perturbations, A_s and n_s defined at a pivot scale $k_0 = 0.05\ {\rm Mpc}^{-1}$, an approximation of the angular scale at recombination, $\theta_{\rm MC}$, and the optical depth to reionization, $\tau_{\rm reio}$. We assume flatness and the existence of massive neutrinos with a total fixed mass of 0.06 eV. We also test select extension models, including variations in the lensing peak smearing $\Lambda{\rm CDM} + A_L$ [with $A_L = 1$ in ΛCDM [21]], in the effective number of relativistic species $\Lambda{\rm CDM} + N_{\rm eff}$ [$N_{\rm eff} = 3.044$ in ΛCDM [22,23]], and in the running of the spectral index of primordial scalar curvature perturbations $\Lambda{\rm CDM} + {\rm d}n_s/{\rm d}\ln k$ [${\rm d}n_s/{\rm d}\ln k = 0$ in ΛCDM [24]].

We follow the naming convention for data combinations used by PL20 and Louis *et al.* [7].

(i) We refer to *Planck* TT/TE/EE when using the data combination of TT, TE, and EE in the multipole range $30 < \ell < 2508$ from either the `Plik-Legacy` or `CamSpec-NPIPE` likelihood. To these, we always add the large-scale temperature data of the `Commander` likelihood from PL20 at $\ell < 30$ and denote it with "lowT." We further add the EE large-scale polarization data from the *Planck* `Sroll2` likelihood [25] with "`Sroll2`"; we sometimes replace this full likelihood with an equivalent Gaussian prior $\tau_{\rm reio} = (5.44 \pm 0.73) \times 10^{-2}$ for quicker evaluations and label this "lowE."[4]

(ii) In some cases, we investigate the constraints from the *Planck* largest scales only—the so-called "cut" data ranges, with the high-ℓ multipole range chosen by [7] of $\ell < 1000/600/600$ in TT/TE/EE. We explicitly call these truncated likelihoods `Plik-Legacy-lite-cut` and `CamSpec-NPIPE-lite-cut`, referring to the full ranges otherwise.

(iii) We combine the two *Planck* `lite-cut` likelihoods with the small-scale data from ACT DR6. In these cases, we use the nomenclature from L25, calling

[4] As shown in other ACT and SPT works, this prior is enough to break the $A_s - \tau_{\rm reio}$ degeneracy.

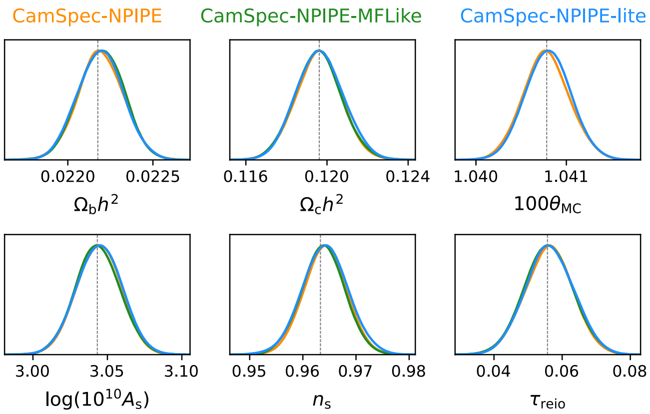

FIG. 4. Posteriors of the ΛCDM cosmological parameters recovered from different implementations of the CamSpec-NPIPE likelihood. We find very good agreement in all five free parameters (we infer the optical depth $\tau_{\rm reio}$ through a prior), and see only a slight widening of 9% on the constraint on n_s from the lite likelihood due to the foreground marginalization. All posteriors are obtained with the full *Planck* TT/TE/EE + lowT + lowE dataset. The vertical lines indicate the best-fitting values from the CamSpec-NPIPE chain (original likelihood, shown in orange).

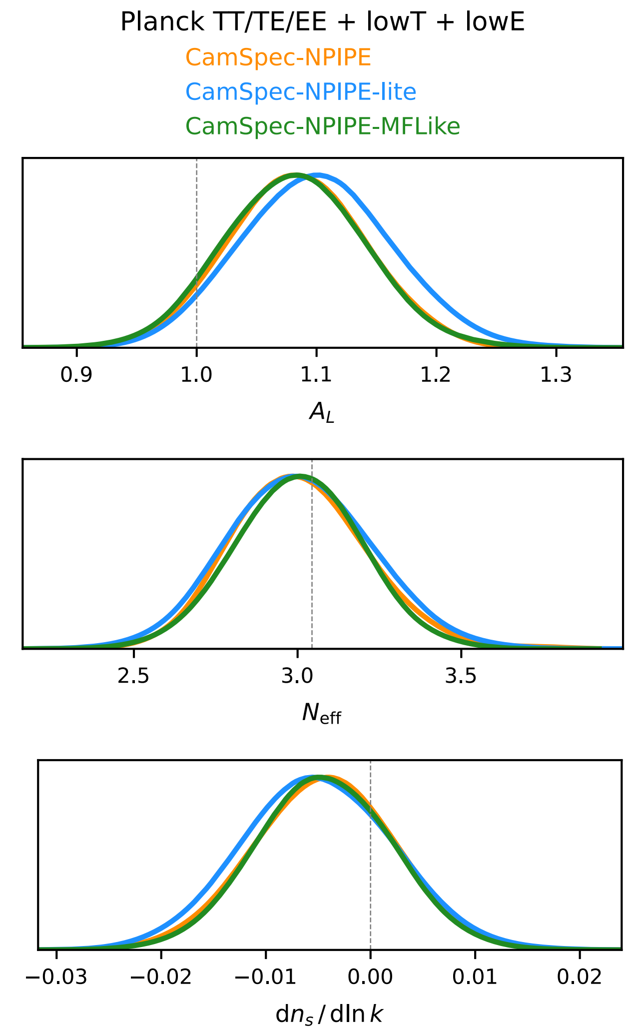

FIG. 5. Comparison of the constraints on select single-parameter extensions to ΛCDM as obtained from the full *Planck* TT/TE/EE + lowT + lowE dataset using the original CamSpec-NPIPE likelihood (orange), our implementation in MFLike (green), and our CamSpec-NPIPE-lite likelihood (blue). We show the constraints on the lensing peak smearing amplitude A_L (top panel), the effective number of relativistic species $N_{\rm eff}$ (middle panel), and the running of the spectral index $dn_s/d\ln k$ (bottom panel). We observe that foreground marginalization has little effect on the constraints of extension models. The constraints from the lite likelihood are consistent and at most 5–10% wider than those from the full likelihoods. The vertical lines show the expected values of these parameters in ΛCDM.

P-ACT the combination of either *Planck* lite-cut likelihood with the ACT DR6-lite likelihood, as well as the large-scale temperature and polarization measurements from Commander and Sroll2. We explicitly differentiate between P-ACT (Plik-Legacy) and P-ACT (CamSpec-NPIPE).

B. Parameter results

1. CamSpec-NPIPE-lite vs full multi-frequency

To validate our implementation of the CamSpec likelihood, we perform cosmological parameter inference with two versions of the full multi-frequency CamSpec-NPIPE likelihood: the original version from R22 and our MFLike implementation, finding perfect agreement both at parameter posterior level and χ^2 evaluation. We also compare the results from these full likelihoods to those derived from our compressed CamSpec-NPIPE-lite likelihood. The constraints on cosmological parameters are shown in Fig. 4. We show an excellent recovery of all five cosmological parameters,[5] with only a minor ($\lesssim 10\%$) widening of the distribution of the primordial power spectrum tilt n_s from the lite likelihood.

The constraints on the parameters of the three extended models tested here are compared between the full and lite likelihoods in Fig. 5, again finding good agreement between the different results. We recover the exact same posteriors between our CamSpec-NPIPE-MFLike and the official CamSpec-NPIPE implementations, and good agreement between those and the parameter posteriors obtained when using CamSpec-NPIPE-lite. Also in this case we observe a minor broadening of parameter uncertainties (5–10%) from the lite likelihood due to the foreground marginalization step. For A_L, we note that the broadening is more enhanced on the high-end tail of the

[5] The optical depth $\tau_{\rm reio}$ is measured in all cases by Sroll2.

TABLE II. Comparison of best-fit ΛCDM solutions obtained with different *Planck* likelihoods fitted to either Plik-Legacy or CamSpec-NPIPE datasets (shortened here into PL and CN, respectively). The rows give the χ^2 value of a given high-ℓ dataset for the best-fitting cosmology from each column, with the number in brackets referring to $\Delta\chi^2$ with respect to the self best-fitting curve of each dataset. There is good agreement between the two *Planck* products and remaining differences decrease for a cut multipole range or when combined with other CMB data. Similar to R22, to calculate χ^2, we minimize for the calibration parameters, while keeping the cosmological parameters fixed.

	Theory					
	Planck TT/TE/EE + lowT + lowE		*Planck*-cut TT/TE/EE + lowT + lowE		P-ACT	
	PL-lite	CN-lite	PL-lite-cut	CN-lite-cut	PL-lite-cut	CN-lite-cut
High-ℓ dataset						
PL-lite	583.27	587.10 (3.83)				
CN-lite	6745.63 (2.88)	6742.75				
PL-lite-cut			216.41	218.17 (1.76)		
CN-lite-cut			2332.46 (1.75)	2330.71		
P-ACT (PL-lite-cut)					378.73	379.47 (0.74)
P-ACT (CN-lite-cut)					2499.18 (0.39)	2498.79

posterior, the full two-dimensional posterior shows that this is down to a one-sided broadening of parameters along the $A_L - \Omega_b h^2$ degeneracy line, instead of a shift in central value.

2. Comparison with Plik-Legacy

Differences between Plik-Legacy and CamSpec-NPIPE cosmologies are expected due to small differences in the amount of data used and many differences in data characterization; these have been discussed at length in PL20, EG21, and R22. To summarize, parameters in ΛCDM are consistent but with the mean of the posteriors shifting by fractions of σ's and in the case of $\Omega_b h^2$ by 1σ (see Fig. 7). Following R22, to give a quantitative comparison, we confront either best-fitting cosmology to either dataset. This is shown in Table II. We find $\Delta\chi^2 = 2.88$ between the CamSpec-NPIPE-lite and Plik-Legacy-lite best-fit ΛCDM theories when compared to the CamSpec-NPIPE-lite data, and $\Delta\chi^2 = 3.83$ when compared to the Plik-Legacy-lite data, showing good agreement between the two cosmological solutions. For these calculations, we keep the cosmological parameters fixed to the best-fit values but minimize for calibrations.

The small increase in data going from the Legacy maps to NPIPE and the precleaning steps allowing to retain more sky fraction in CamSpec-NPIPE translates into tighter limits for both ΛCDM and extended model parameters. To localize where in spectrum space the additional constraining power is coming from, we show a comparison of the Plik-Legacy-lite and CamSpec-NPIPE-lite per-ℓ constraining power in Fig. 6, calculating the signal-to-noise ratio per bin in the spectra, corrected for the different bin sizes and the slight difference in the treatment of polarization efficiencies. We see that the CamSpec-NPIPE-lite data has better constraining power in TE and EE, as well as in TT at $\ell > 1500$. Better polarization and improved small scales in temperature are particularly important for measuring n_s and extended model parameters.

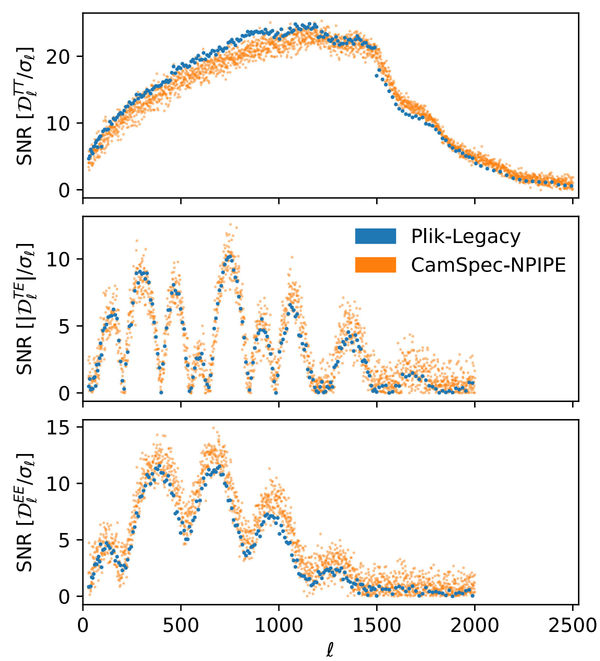

FIG. 6. Comparison of the signal-to-noise ratios $\mathcal{D}_\ell/\sigma_\ell$ of Plik-Legacy (blue) and CamSpec-NPIPE (orange) after foreground-marginalization. The Plik data points contain an additional factor of $\sqrt{\Delta\ell}$ to correct for the bin widths of the data, while the NPIPE is plotted with an additional 1% error margin since the likelihood marginalizes over additional TE and EE calibration factors. The constraining power in TT from Plik-Legacy is higher at $\ell < 1200$ due to the inclusion of the 100×100 spectrum, which is absent in CamSpec-NPIPE.

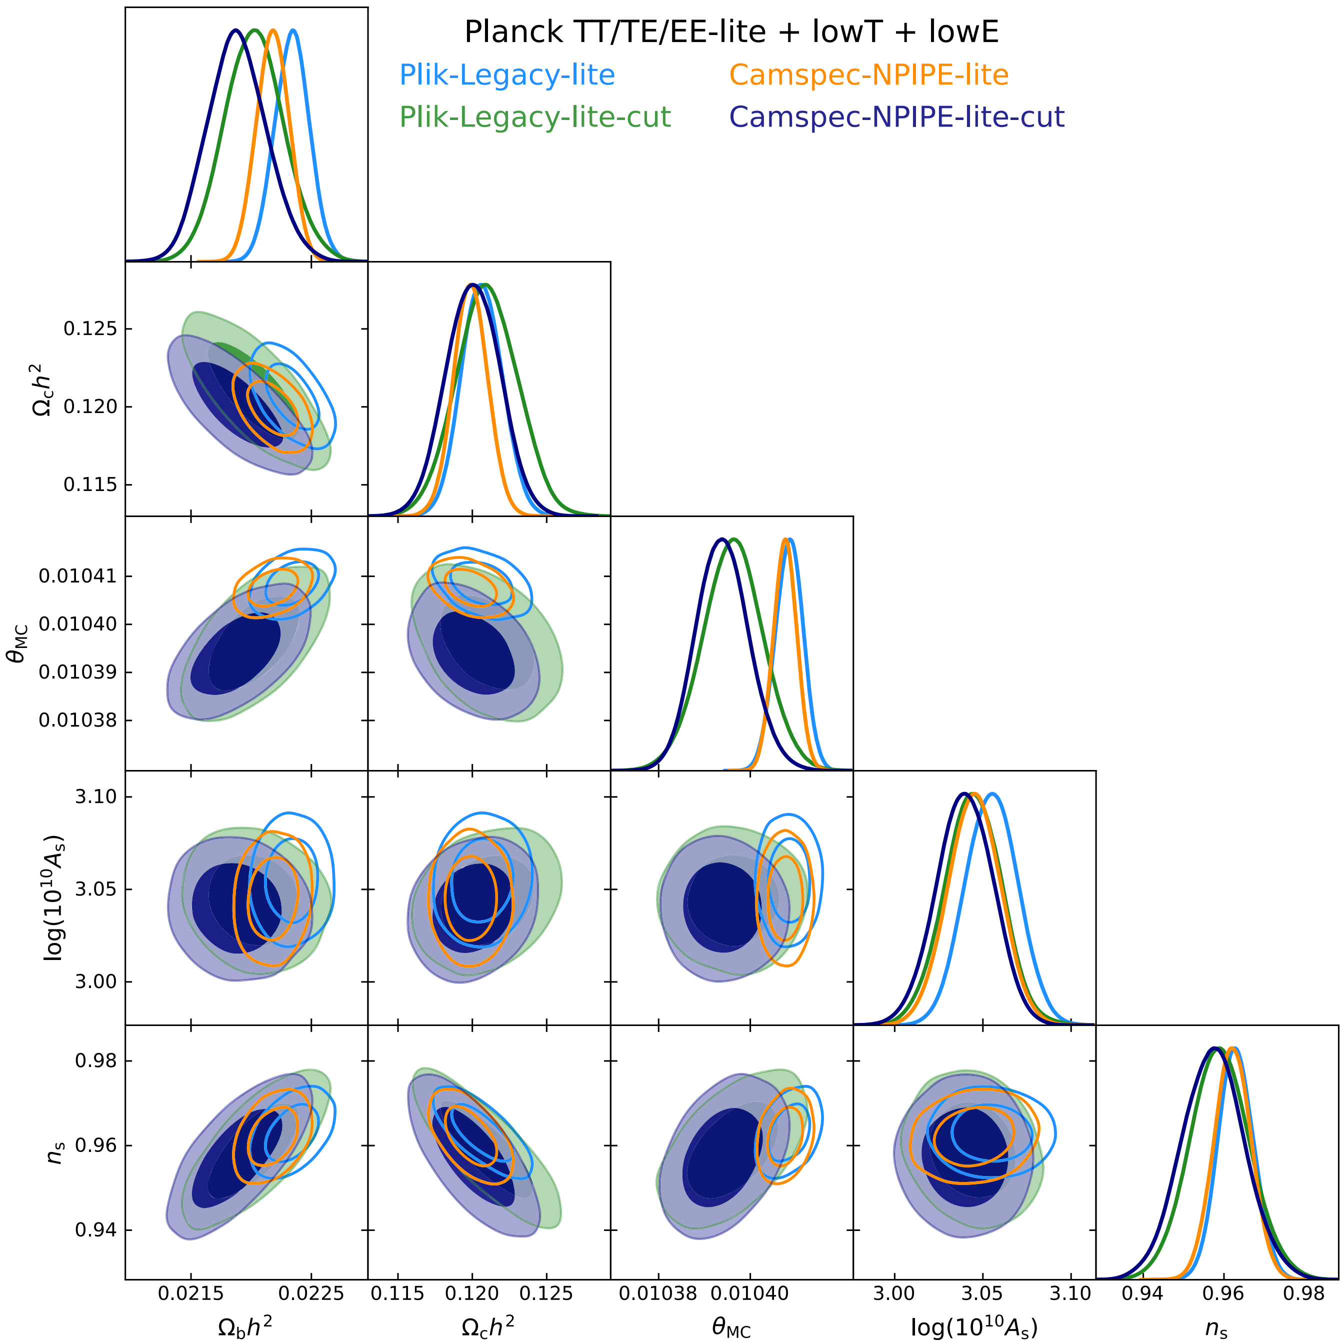

FIG. 7. Comparison of ΛCDM parameters as obtained from `Plik-Legacy` and `CamSpec-NPIPE` likelihoods, using the full dataset (blue and orange, respectively) and a restricted range keeping only the large scales with $\ell < 1000/600/600$ in TT/TE/EE (green and navy, respectively). The results from the full datasets (open contours) do not overlap as much in parameter space as the cut likelihoods (filled contours). The numerical values of these constraints are given in Table III.

3. Parameters from a restricted multipole range

We now test a restricted multipole range, cutting the *Planck* power spectra in TT/TE/EE at $\ell < 1000/600/600$, which is the point at which ground-based experiments currently probe power spectra to higher precision than *Planck*. These criteria were first set out in Louis *et al.* [7] for ACT DR6 and later reused in Camphuis *et al.* [9] for the combination of *Planck*, ACT, and SPT-3G.

We show the constraints on ΛCDM parameters in Fig. 7 and their numerical values in Table III, comparing the `CamSpec-NPIPE` and `CamSpec-NPIPE-cut` constraints versus the equivalent for `Plik`. The restricted multipole range gives very similar parameters between the two likelihoods, improving on the agreement from the full dataset. The largest shift is of $\sim 0.4\sigma$ in $\Omega_b h^2$, and $\Delta\chi^2$ with respect to the fit to the other dataset is reduced by ~ 1–2 points (see Table 2).

TABLE III. Constraints on the ΛCDM basic varying parameters from `Plik-Legacy-lite` and `CamSpec-NPIPE-lite`, comparing the full (left) and restricted multipole (right) ranges. These numerical values correspond to the results shown in Fig. 7.

	`Plik-Legacy-lite`	`CamSpec-NPIPE-lite`	`Plik-Legacy-lite-cut`	`CamSpec-NPIPE-lite-cut`
$\Omega_b h^2$	$(2.237 \pm 0.015) \times 10^{-2}$	$(2.218 \pm 0.014) \times 10^{-2}$	$(2.203 \pm 0.027) \times 10^{-2}$	$(2.188 \pm 0.025) \times 10^{-2}$
$\Omega_c h^2$	$(12.01 \pm 0.14) \times 10^{-2}$	$(11.99 \pm 0.12) \times 10^{-2}$	$(12.07 \pm 0.22) \times 10^{-2}$	$(12.00 \pm 0.18) \times 10^{-2}$
θ_{MC}	$(104.087 \pm 0.031) \times 10^{-4}$	$(104.077 \pm 0.026) \times 10^{-4}$	$(103.952 \pm 0.070) \times 10^{-4}$	$(103.944 \pm 0.057) \times 10^{-4}$
n_s	0.9651 ± 0.0043	0.9622 ± 0.0046	0.9594 ± 0.0078	0.9575 ± 0.0077
$\log(10^{10} A_s)$	3.054 ± 0.014	3.046 ± 0.016	3.043 ± 0.016	3.040 ± 0.016

IV. COMBINING PLANCK WITH OTHER CMB DATASETS

One motivation for this work is to generate a `CamSpec-NPIPE` product that can be easily and correctly combined at likelihood level with other CMB datasets. The foreground marginalization performed here allows one to properly treat the foreground contribution in the large-scale `NPIPE` data when combining the cut likelihood with a different small-scale measurement, for example, from ACT DR6.

With this product now available, we derive here `P-ACT` cosmological constraints for different *Planck* likelihoods in combination with ACT DR6 as defined in Sec. III A, to check if the specific choice of *Planck* is important in these combined results. Throughout this section, we only combine `lite` likelihoods.

A. Constraints on ΛCDM

We show a comparison between the constraints obtained with `Plik-Legacy` and `CamSpec-NPIPE` in combination with small-scale ACT DR6 data in Fig. 8.

Overall, we see that the agreement between `Plik-Legacy` and `CamSpec-NPIPE` is maximum when combined with ACT DR6, reducing their difference to $\Delta\chi^2 = 0.4$–0.7 in the `P-ACT` combination. This is partially due to having cut the *Planck* likelihoods and partially to the additional constraining power from ACT. The 1.0σ difference in $\Omega_b h^2$ between the two full *Planck* datasets goes down to 0.4σ for the *Planck* cut cases and then 0.15σ when combining with ACT DR6. In the `P-ACT` combination, the largest parameter difference is a 0.7σ difference in $\Omega_c h^2$. We compare the numerical values of the ΛCDM parameter constraints in Table IV.

B. Constraints on extension models

We now compare the constraints on the extended models $\Lambda\mathrm{CDM} + A_L$ and $\Lambda\mathrm{CDM} + N_{\mathrm{eff}}$. We expect that the A_L extension carries significant weight from the *Planck* measurement of large angular scales, while the N_{eff} extension should be dominated by the ACT measurement of small scales. As such, these two extensions are good proxies for the effects of combining the two different datasets. The marginal constraints on the extension parameters for the different *Planck* likelihoods and their combinations with ACT are shown in Fig. 9.

We see that the combination of *Planck* with ACT removes completely the effect of the choice of *Planck* likelihood on the constraint of the extension parameter. Both A_L and N_{eff} are recovered to the same value, with minimal differences in other parameters.

We extend this test to run the CMB-SPA combination of [9], and we find the same conclusions: the results are insensitive to the choice of *Planck*.

C. Evidence for evolving dark energy

The recent DESI DR2 results have generated a lot of interest in the prospects of dynamical dark energy, with a tentative statistical preference for $w_0 w_a$ over ΛCDM [26,27] where w_0 and w_a parametrize the present-day value of the dark energy equation of state and its variation with time, respectively. In ΛCDM $w_0 = -1, w_a = 0$ [28,29]. In light of the vast range of dataset combinations that have been explored for this model, and their difference in statistical preference [see e.g., Ref. [30]], we test here both the use of `CamSpec-NPIPE` instead of `Plik-Legacy` and the effects of foreground marginalization for `CamSpec-NPIPE` in the combination `P-ACT-LBS`[6] on the $w_0 w_a$ constraints and statistical preference over ΛCDM. In deriving best-fit points in likelihood space, we use the `BOBYQA` minimizer [31–33].

After foreground marginalization, we use Wilks' theorem to find that the statistical preference for $w_0 w_a$ is 2.5σ when using `Plik-Legacy-lite-cut` data, or 2.4σ when using the foreground-marginalized `CamSpec-NPIPE-lite-cut` data. With a simple truncation of the full multi-frequency likelihood—leaving the foreground unconstrained—for `CamSpec-NPIPE-cut`, we find 2.3σ. We conclude that the `P-ACT` evidence for $w_0 w_a$ does not depend on the choice of *Planck* (see also Fig. 10) and that the treatment of the nuisance parameters changes

[6]`P-ACT-LBS` is defined as in [6] as the combination of large-scale CMB measurements from *Planck* `Commander` for TT, `Sroll2` for EE, *Planck* for the intermediate scales $30 < \ell < 1000/600/600$ in TT/TE/EE, ACT DR6 for the range $600 < \ell < 6500$, the *Planck* NPIPE + ACT DR6 lensing measurement, BAO measurements from DESI DR2, and supernovae from the Pantheon + sample.

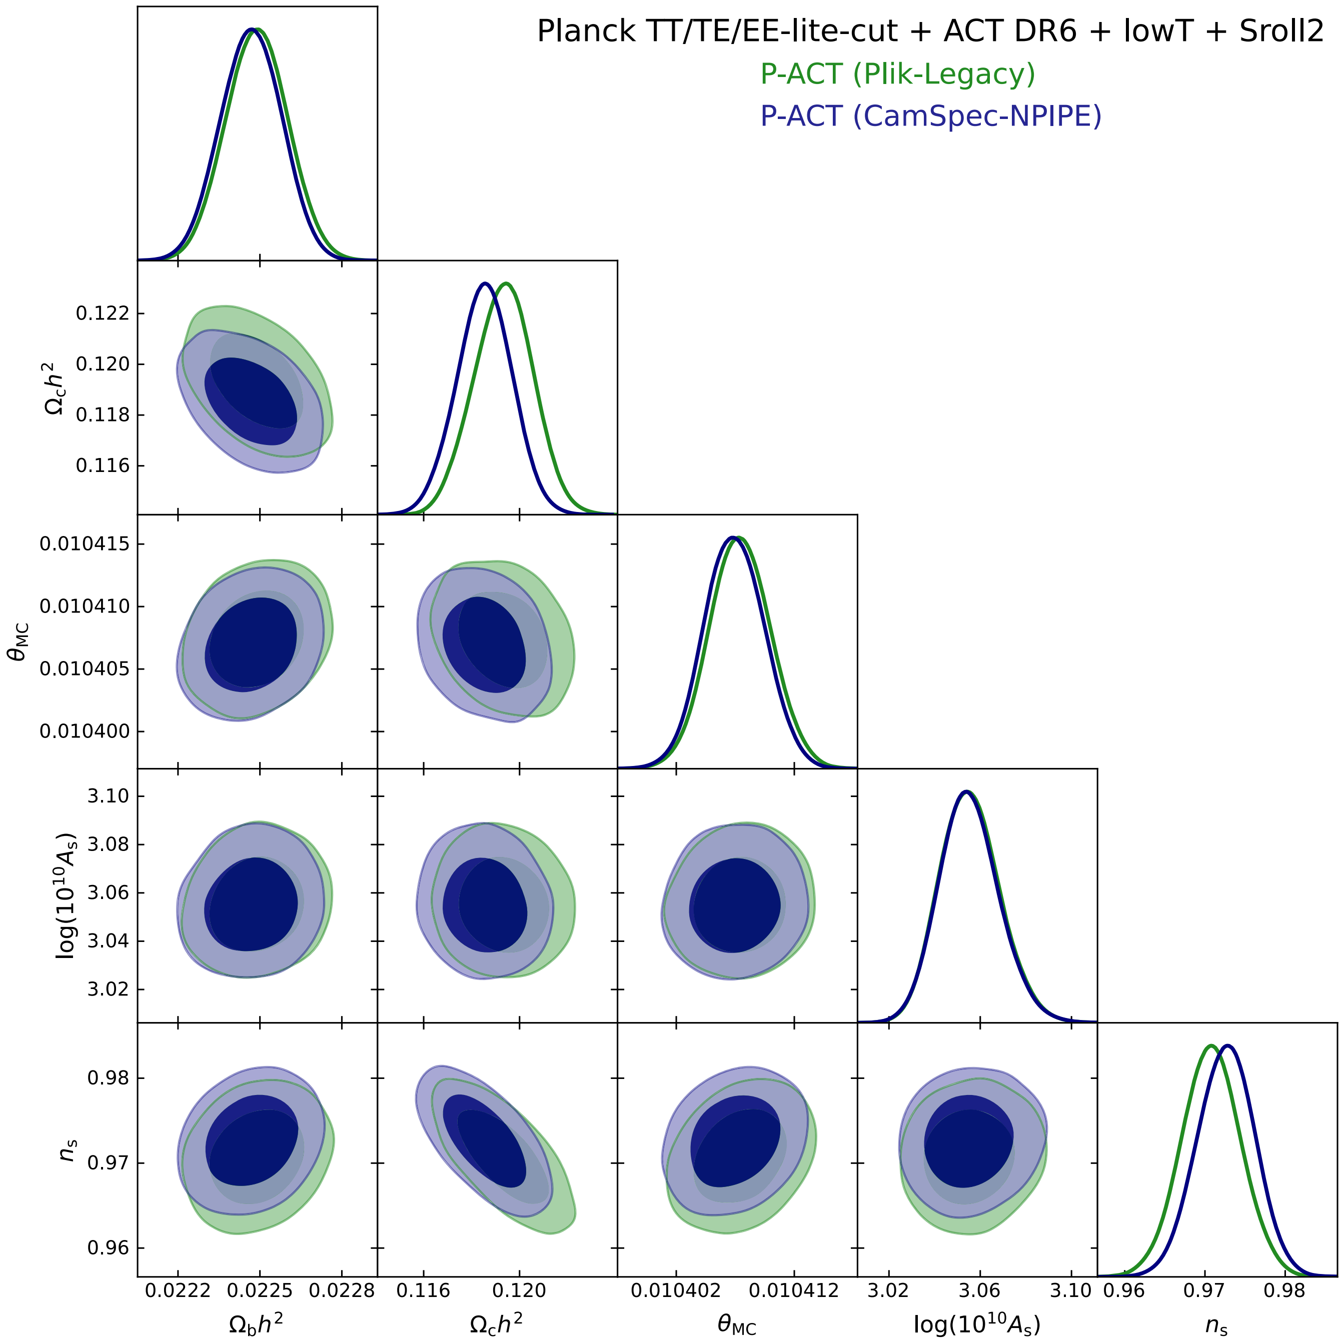

FIG. 8. ΛCDM constraints from the `P-ACT` combination with either `Plik-Legacy` (green) or `CamSpec-NPIPE` (dark blue) in all five cosmological parameters. We find excellent agreement in between either choice in the `P-ACT` combination. Compared to Fig. 7, we notice that there is less difference in most parameters.

this number by $\sim 0.1\sigma$. We note that [30] found the same preference when using the `P-ACT` combination with `Plik-Legacy`; their combination with `CamSpec-NPIPE` gives a 2.8σ preference for $w_0 w_a$ but assumes a substantially different dataset in both ℓ cuts ($\ell < 2000/1000/1000$ in TT/TE/EE) and the choice of low-ℓ EE likelihood and therefore it is not directly comparable with the `P-ACT` results obtained with `Plik-Legacy`.

We summarize these numbers in Table V and note that they fluctuate at the level of $0.2-0.3\sigma$ depending on choices of low-ℓ EE likelihood, minimizer used for deriving the best-fit points, and accuracy settings in theory calculations. The threshold of 3σ used to claim tentative evidence is therefore conditioned to several assumptions made in the analysis, including the treatment of foregrounds.

TABLE IV. Constraints on the ΛCDM basic varying (top) and key derived (bottom) parameters from `P-ACT`, comparing using `Plik-Legacy-lite` and `CamSpec-NPIPE-lite` in combination with ACT.

	`P-ACT (Plik-Legacy)`	`P-ACT (CamSpec-NPIPE)`
$\Omega_b h^2$	$(2.249 \pm 0.011) \times 10^{-2}$	$(2.247 \pm 0.011) \times 10^{-2}$
$\Omega_c h^2$	$(11.93 \pm 0.12) \times 10^{-2}$	$(11.85 \pm 0.12) \times 10^{-2}$
θ_{MC}	$(104.074 \pm 0.026) \times 10^{-4}$	$(104.069 \pm 0.025) \times 10^{-4}$
n_s	0.9708 ± 0.0037	0.9726 ± 0.0036
$\log(10^{10} A_s)$	$3.055^{+0.012}_{-0.014}$	$3.055^{+0.012}_{-0.014}$
H_0	67.62 ± 0.51	67.86 ± 0.48
σ_8	0.8146 ± 0.0063	0.8125 ± 0.0062
Ω_m	$(31.03 \pm 0.72) \times 10^{-2}$	$(30.63 \pm 0.67) \times 10^{-2}$
$A_s e^{-2\tau}$	$(1.883 \pm 0.012) \times 10^{-9}$	$(1.879 \pm 0.013) \times 10^{-9}$

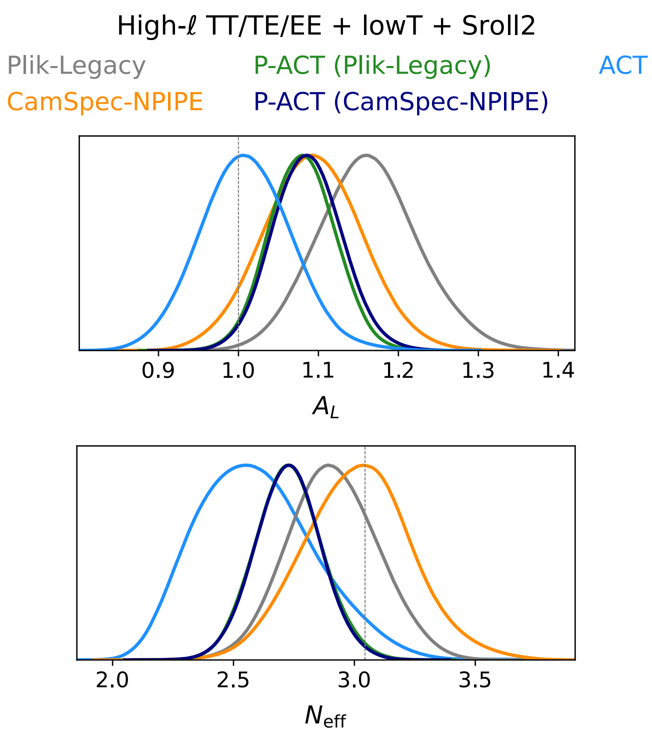

FIG. 9. Comparisons of the ACT, *Planck*, and `P-ACT` constraints on A_L and N_{eff} for the `Plik-Legacy` and `CamSpec-NPIPE` likelihoods. In both cases, the `P-ACT` combination gives identical results for either *Planck* likelihood.

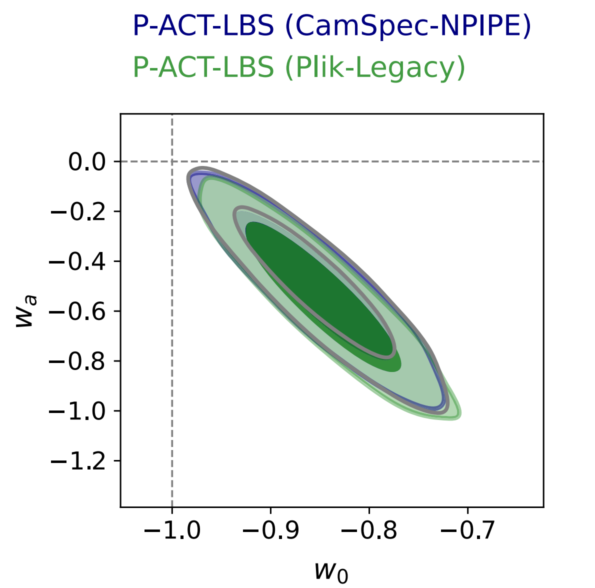

FIG. 10. Constraints on $w_0 w_a$ from `P-ACT-LBS` using either `Plik-Legacy` (green) or `CamSpec-NPIPE` (dark blue). The dashed lines denote ΛCDM at $w_0 = -1, w_a = 0$. The empty gray contour uses the same data combination but with the non-marginalized `CamSpec-NPIPE` likelihood. The contours are insensitive to the choice of *Planck*.

TABLE V. Summary of the statistics for $w_0 w_a$ over ΛCDM from the `P-ACT-LBS` combination using different *Planck* likelihoods. The $\Delta\chi^2$ column lists the $\chi^2_{w_0 w_a} - \chi^2_{\Lambda\mathrm{CDM}}$ difference at the Maximum A Posteriori point, and the preference represents the frequentist significance from this χ^2 difference.

`P-ACT-LBS` for *Planck* likelihood	$\Delta\chi^2$	Preference
`Plik-Legacy-lite-cut`	−8.8	2.5σ
`CamSpec-NPIPE-lite-cut`	−8.1	2.4σ
`CamSpec-NPIPE-cut`	−7.9	2.3σ

V. CONCLUSIONS

We have summarized differences in cosmological results resulting from different combinations of *Planck* sky maps and likelihood pipelines. To ease the comparison and to build a product that can be correctly combined with other CMB datasets at likelihood level, we have derived a foreground-marginalized `lite` version of the `CamSpec-NPIPE` likelihood. We use this new `lite` likelihood to localize in CMB spectrum space the additional constraining power present in `CamSpec-NPIPE` compared to `Plik-Legacy` and then perform cosmological analysis with a restricted multipole range as now commonly done by ground-based experiments. We show that limits on cosmological parameters, in particular in extended models, are completely insensitive to the choice of *Planck* maps/likelihood when keeping the *Planck* large scales in combination with small scales from ACT DR6.

ACKNOWLEDGMENTS

We are grateful to ACT collaborators for useful discussions. We thank Ian Harrison and Serena Giardiello for useful feedback. H. T. J., M. V., and E. C. acknowledge support from the Horizon 2020 ERC Starting Grant (Grant Agreement No. 849169). J. C. H. acknowledges support from NSF Grant No. AST-2108536, the Sloan Foundation, and the Simons Foundation.

DATA AVAILABILITY

The data that support the findings of this article are openly available [14,18].

[1] Planck Collaboration, Planck 2018 results. I. Overview and the cosmological legacy of Planck, Astron. Astrophys. **641**, A1 (2020).

[2] Planck Collaboration, Planck 2018 results. V. CMB power spectra and likelihoods, Astron. Astrophys. **641**, A5 (2020).

[3] Planck Collaboration, Planck 2018 results. VI. Cosmological parameters, Astron. Astrophys. **641**, A6 (2020); **652**, C4 (E) (2021).

[4] E. Rosenberg, S. Gratton, and G. Efstathiou, CMB power spectra and cosmological parameters from Planck PR4 with `CamSpec`, Mon. Not. R. Astron. Soc. **517**, 4620 (2022).

[5] M. Tristram *et al.*, Cosmological parameters derived from the final Planck data release (PR4), Astron. Astrophys. **682**, A37 (2024).

[6] E. Calabrese, J. C. Hill, H. Jense, A. La Posta *et al.*, The Atacama Cosmology Telescope: DR6 constraints on extended cosmological models, J. Cosmol. Astropart. Phys. 11 (2025) 063.

[7] T. Louis, A. La Posta, Z. Atkins, H. Jense *et al.*, The Atacama Cosmology Telescope: DR6 power spectra, likelihoods and ΛCDM parameters, J. Cosmol. Astropart. Phys. 11 (2025) 062.

[8] S. Naess, Y. Guan, A. Duivenvoorden, M. Hasselfield *et al.*, The Atacama Cosmology Telescope: DR6 maps, J. Cosmol. Astropart. Phys. 11 (2025) 061.

[9] E. Camphuis *et al.*, SPT-3G D1: CMB temperature and polarization power spectra and cosmology from 2019 and 2020 observations of the SPT-3G main field, arXiv:2506.20707.

[10] G. Efstathiou and S. Gratton, A detailed description of the `CamSpec` likelihood pipeline and a reanalysis of the Planck high frequency maps, Open J. Astrophys. **4**, 8 (2021).

[11] Planck Collaboration *et al.*, Planck intermediate results. LVII. Joint Planck LFI and HFI data processing, Astron. Astrophys. **643**, A42 (2020).

[12] Planck Collaboration, Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters, Astron. Astrophys. **594**, A11 (2016).

[13] Z. Li *et al.*, The Simons Observatory: A new open-source power spectrum pipeline applied to the `Planck Legacy` data, J. Cosmol. Astropart. Phys. 09 (2023) 048.

[14] M. Viña Bertran, `CamSpec-NPIPE-MFLike`, https://github.com/MarcVB93/mflike-camspec (2025).

[15] L. Balkenhol, Compressed "CMB-lite" likelihoods using automatic differentiation, Open J. Astrophys. **8**, 17 (2025).

[16] S. K. Choi *et al.*, The Atacama Cosmology Telescope: A measurement of the cosmic microwave background power spectra at 98 and 150 GHz, J. Cosmol. Astropart. Phys., 12 (2020) 045.

[17] J. Dunkley *et al.*, The Atacama Cosmology Telescope: Likelihood for small-scale CMB data, J. Cosmol. Astropart. Phys. 07 (2013) 025.

[18] H. T. Jense, `CamSpec-NPIPE-lite`, https://github.com/HTJense/camspec_npipe-lite (2025).

[19] J. Torrado and A. Lewis, Cobaya: Bayesian analysis in cosmology, Astrophysics Source Code Library (2019), https://ascl.net/1910.019.

[20] J. Torrado and A. Lewis, Cobaya: Code for Bayesian analysis of hierarchical physical models, J. Cosmol. Astropart. Phys. 05 (2021) 057.

[21] E. Calabrese, A. Slosar, A. Melchiorri, G. F. Smoot, and O. Zahn, Cosmic microwave weak lensing data as a test for the dark universe, Phys. Rev. D **77**, 123531 (2008).

[22] J. J. Bennett, G. Buldgen, P. F. de Salas, M. Drewes, S. Gariazzo, S. Pastor, and Y. Y. Y. Wong, Towards a precision calculation of the effective number of neutrinos $N_{\rm eff}$ in the standard model. Part II. Neutrino decoupling in the presence of flavour oscillations and finite-temperature QED, J. Cosmol. Astropart. Phys. 04 (2021) 073.

[23] M. Drewes, Y. Georis, M. Klasen, L. P. Wiggering, and Y. Y. Y. Wong, Towards a precision calculation of N_{eff} in the standard model. Part III. Improved estimate of NLO contributions to the collision integral, J. Cosmol. Astropart. Phys. 06 (2024) 032.

[24] A. Kosowsky and M. S. Turner, CBR anisotropy and the running of the scalar spectral index, Phys. Rev. D **52**, R1739 (1995).

[25] L. Pagano, J. M. Delouis, S. Mottet, J. L. Puget, and L. Vibert, Reionization optical depth determination from Planck HFI data with ten percent accuracy, Astron. Astrophys. **635**, A99 (2020).

[26] DESI Collaboration, DESI DR2 results I: Baryon acoustic oscillations from the Lyman alpha forest, Phys. Rev. D **112**, 083514 (2025).

[27] DESI Collaboration, DESI DR2 results II: Measurements of baryon acoustic oscillations and cosmological constraints, Phys. Rev. D **112**, 083515 (2025).

[28] M. Chevallier and D. Polarski, Accelerating universes with scaling dark matter, Int. J. Mod. Phys. D **10**, 213 (2001).

[29] E. V. Linder, Exploring the expansion history of the universe, Phys. Rev. Lett. **90**, 091301 (2003).

[30] DESI Collaboration, Cosmological implications of DESI DR2 BAO measurements in light of the latest ACT DR6 CMB data, Phys. Rev. D **112**, 083529 (2025).

[31] C. Cartis, J. Fiala, B. Marteau, and L. Roberts, Improving the flexibility and robustness of model-based derivative-free optimization solvers, arXiv:1804.00154.

[32] M. J. D. Powell, The BOBYQA algorithm for bound constrained optimization without derivatives, DAMTP, University of Cambridge, https://www.damtp.cam.ac.uk/user/na/NA_papers/NA2009_06.pdf.

[33] C. Cartis, L. Roberts, and O. Sheridan-Methven, Escaping local minima with local derivative-free methods: A numerical investigation, Optimization **71**, 2343 (2021).