


BEYONDPLANCK

IX. Bandpass and beam leakage corrections

T. L. Svalheim¹ , A. Zonca², K. J. Andersen¹, R. Aurlen¹, R. Banerji¹, M. Bersanelli^{3,4,5}, S. Bertocco⁶, M. Brilenkov¹, M. Carbone⁷, L. P. L. Colombo³, H. K. Eriksen¹, M. K. Foss¹, C. Franceschet^{3,5}, U. Fuskeland¹, S. Galeotta⁶, M. Galloway¹, S. Gerakakis⁷, E. Gjerløw¹, B. Hensley⁸, D. Herman¹, M. Iacobellis⁷, M. Ieronymaki⁷, H. T. Ihle¹, J. B. Jewell⁹, A. Karakci¹, E. Keihänen^{10,11}, R. Keskitalo¹², G. Maggio⁶, D. Maino^{3,4,5}, M. Maris⁶, S. Paradiso^{3,5}, B. Partridge¹⁴, M. Reinecke¹⁵, A.-S. Suur-Uski^{10,11}, D. Tavagnacco^{6,13}, H. Thommesen¹, D. J. Watts¹, I. K. Wehus¹, and A. Zacchei⁶

¹ Institute of Theoretical Astrophysics, University of Oslo, Blindern, Oslo, Norway
e-mail: t.l.svalheim@astro.uio.no

² San Diego Supercomputer Center, University of California, San Diego, USA

³ Dipartimento di Fisica, Università degli Studi di Milano, Via Celoria 16, Milano, Italy

⁴ INAF/IASF Milano, Via E. Bassini 15, Milano, Italy

⁵ INFN, Sezione di Milano, Via Celoria 16, Milano, Italy

⁶ INAF – Osservatorio Astronomico di Trieste, Via G.B. Tiepolo 11, Trieste, Italy

⁷ Planetek Hellas, Leoforos Kifisias 44, Marousi 151 25, Greece

⁸ Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA

¹⁰ Department of Physics, University of Helsinki, Gustaf Hällströmin katu 2, Helsinki, Finland

¹¹ Helsinki Institute of Physics, University of Helsinki, Gustaf Hällströmin katu 2, Helsinki, Finland

¹² Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

¹³ Haverford College Astronomy Department, 370 Lancaster Avenue, Haverford, PA, USA

¹⁴ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

¹⁵ Dipartimento di Fisica, Università degli Studi di Trieste, Via A. Valerio 2, Trieste, Italy

Received 10 January 2022 / Accepted 25 February 2022

ABSTRACT

We discuss the treatment of bandpass and beam leakage corrections in the Bayesian BEYONDPLANCK cosmic microwave background (CMB) analysis pipeline as applied to the *Planck* LFI measurements. As a preparatory step, we first applied three corrections to the nominal LFI bandpass profiles, including the removal of a known systematic effect in the ground measuring equipment at 61 GHz, along with a smoothing of standing wave ripples and edge regularization. The main net impact of these modifications is an overall shift in the 70 GHz bandpass of +0.6 GHz. We argue that any analysis of LFI data products, either from *Planck* or BEYONDPLANCK, should use these new bandpasses. In addition, we fit a single free bandpass parameter for each radiometer of the form $\Delta_i = \Delta_0 + \delta_i$, where Δ_0 represents an absolute frequency shift per frequency band and δ_i is a relative shift per detector. The absolute correction is only fitted at 30 GHz, with a full χ^2 -based likelihood, resulting in a correction of $\Delta_{30} = 0.24 \pm 0.03$ GHz. The relative corrections were fitted using a spurious map approach that is fundamentally similar to the method pioneered by the WMAP team, but excluding the introduction of many additional degrees of freedom. All the bandpass parameters were sampled using a standard Metropolis sampler within the main BEYONDPLANCK Gibbs chain and the bandpass uncertainties were thus propagated to all other data products in the analysis. In summary, we find that our bandpass model significantly reduces leakage effects. For beam leakage corrections, we adopted the official *Planck* LFI beam estimates without any additional degrees of freedom and we only marginalized over the underlying sky model. We note that this is the first time that leakage from beam mismatch has been included for *Planck* LFI maps.

Key words. cosmic background radiation

1. Introduction

Cosmic microwave background (CMB), initially discovered by Penzias & Wilson (1965), now stands as one of the most important sources of information in cosmology. The most recent full-sky measurements of this signal were made by the *Planck* satellite (Planck Collaboration I 2020) from an orbit around the second Sun–Earth Lagrange point between 2009 and 2013, using two complementary instruments to observe the sky in nine frequency bands between 30 and 857 GHz. These measurements have put strong constraints on a wide range of both cosmological parameters and physical phe-

nomena, forming one of the cornerstones of contemporary cosmology.

Although the official *Planck* data processing ended in 2020 (Planck Collaboration I 2020; Planck Collaboration Int. LVII 2020), several open questions regarding low-level instrumental effects in *Planck* remained unanswered at that time. Addressing these concerns for the Low Frequency Instrument (LFI; Planck Collaboration II 2020) is one of the main motivations devised for the BEYONDPLANCK project BeyondPlanck Collaboration (2023). The BEYONDPLANCK machinery is unique in that it processes raw time-ordered data (TOD) into final cosmological and astrophysical results

within one single integrated end-to-end Bayesian analysis framework. Its computational engine is Commander (Eriksen et al. 2004, 2008; Galloway et al. 2023), which was originally developed for *Planck* component separation purposes (Planck Collaboration XII 2014; Planck Collaboration X 2016; Planck Collaboration IV 2020) and which uses Gibbs sampling (Geman & Geman 1984) to draw samples from a large global posterior distribution. The BEYONDPLANCK project has generalized this code to also account for low-level data processing and mapmaking, and thereby integrated the full analysis pipeline into a self-consistent Bayesian framework. For a full description of the BEYONDPLANCK project, we refer the interested reader to BeyondPlanck Collaboration (2023).

A defining theme for the BEYONDPLANCK approach is a detailed statistical exploration of the interplay between instrumental effects and astrophysical foregrounds. The area that is particularly relevant in this respect is related to spectral responses. Specifically, LFI comprised 22 radiometer chain assemblies (RCAs) grouped in three bands with nominal frequencies of 30, 44, and 70 GHz. Within a band, each RCA has a slightly different spectral response (or bandpass profile) and center frequency compared to the others. The *Planck* LFI spectral responses used in the official analysis were provided as part of the 2018 data release¹. These response functions were measured on the ground prior to launch, as described in Zonca et al. (2009). However, laboratory measurements of bandpass profiles are generally a highly non-trivial task, as even small environmental variations and interferences may affect the results. Several systematic issues were identified during the LFI testing campaign and some of these were left uncorrected in the final LFI products, even though approaches for possible improvements had been discussed. In this paper, we finally implement these corrections and make the resulting bandpass profiles publicly available². We also recommend that any future analysis of the *Planck* LFI data should use the new bandpasses, as the effects are non-negligible and lead to an improved internal consistency.

Even with perfect knowledge of the instrument, there will be artefacts in the final frequency and component maps when using a multi-detector mapmaking algorithm (which includes those employed by *Planck*; Planck Collaboration II 2020; Planck Collaboration III 2020) if the detectors have different sensitivities, unless these differences are properly accounted for (e.g., Page et al. 2007; Planck Collaboration Int. LVII 2020; Delouis et al. 2019). This applies both to differences in bandpass and beam profiles. In this paper, we discuss how such corrections are applied within the BEYONDPLANCK framework, both in terms of how to correct deterministically for an assumed known response functions, and how to account for uncertainties in the response functions themselves. We note that since the effect of the bandpass is global and influences both gain and foreground results, sampling corrections are both phenomenologically challenging and computationally demanding.

We note that for the white noise levels typical of WMAP and *Planck*, previous static leakage corrections suppress residual effects that are well below the levels relevant for cosmological interpretation. However, future CMB experiments will target primordial gravitational waves (e.g., Kamionkowski & Kovetz 2016 and references therein) and, more, specifically the tensor-to-scalar ratio, r . The current measurements from BICEP2/Keck and *Planck* constrain this parameter to $r < 0.036$ (Ade et al. 2021) and $r < 0.032$ (Tristram et al. 2022), respectively, and

these low levels correspond to signals that are only a few hundred nanokelvin on the sky. To reach levels of $r \lesssim 10^{-3}$ – 10^{-4} , highly accurate modeling of bandpass and beam leakage effects will be critically important and we believe that an integrated method of the kind shown in this paper will be required.

The rest of this paper is organized as follows. We start with a discussion of the LFI bandpass pre-processing steps in Sect. 2. We then present our own algorithms in Sect. 3 and show how these directly build on and generalize previous efforts. The results from the main Markov chain analysis are reported in Sect. 4. We present our conclusions in Sect. 5.

2. Bandpass pre-processing

We start our discussion by reviewing the bandpass profiles provided by the official *Planck* LFI data processing center (DPC) (Planck Collaboration II 2014), as shown in Fig. 1 for each radiometer. As summarized by Zonca et al. (2009), these functions were characterized through two complementary approaches. The first was aimed at modeling and characterizing each optical element individually and then combining these into a complete model. The second approach was carried out through integrated cryogenic system tests for LFI by Thales Alenia Space in Vimodrone for the 30 and 44 GHz channels, as well as by DA-Design Ylinen in Finland for the 70 GHz channel.

While the results from the two approaches generally agreed, there were also notable differences, some of which may be identified by eye in Fig. 1. Perhaps the most striking feature is standing waves, which result in high-frequency ripples across the bandpass. While some of these may be due to real features in RCA itself, others may be caused by standing waves excited by the testing equipment. One particularly notable example of this is the 70 GHz channel, for which the input load was placed directly in front of the feed horn during the test, resulting in strong standing waves between the two components. At the same time, the precise phase of the ripples is sensitive to environmental properties and can change depending, for instance, on the ambient temperature. As such, it is non-trivial to assign a physical reality to these ripples as far as the real measurements are concerned. Zonca et al. (2009) therefore suggested that these should be removed through low-pass filtering before cosmological data analysis on the grounds that the independent end-to-end test showed general agreement within the range of reliable results, but small-scale ripples were not aligned among the two; this is illustrated in Fig. 2, which is a direct reproduction of Fig. 5 from the original paper. Smoothing the ripples was, however, never actually implemented during the official *Planck* analysis and, thus, this process is a feature of the current paper. Technically speaking, we implement this filter in logarithmic space using a three-pole low pass filter at 0.2 times the Nyquist frequency. We note, however, that the specific details regarding the filter are not critically important.

A second known artefact induced by the testing equipment is the excess seen in the 70 GHz bandpasses below 61 GHz, which is due to a known systematic effect in all gain measurements of the backend module (BEM); it is thus caused by the test equipment – and not the radiometers themselves (Zonca et al. 2009). This feature should simply be removed and the 70 GHz bandpass should be limited to 61–80 GHz.

More generally, the edges of the bandpasses are not well characterized, and standing wave effects also tend to be relatively more significant near the edges. The bandpass cutoffs should therefore be regularized through smooth apodization. In this paper, we implement this by calculating the derivative of

¹ <https://pla.esac.esa.int/>

² https://beyondplanck.science/products/files_v2/

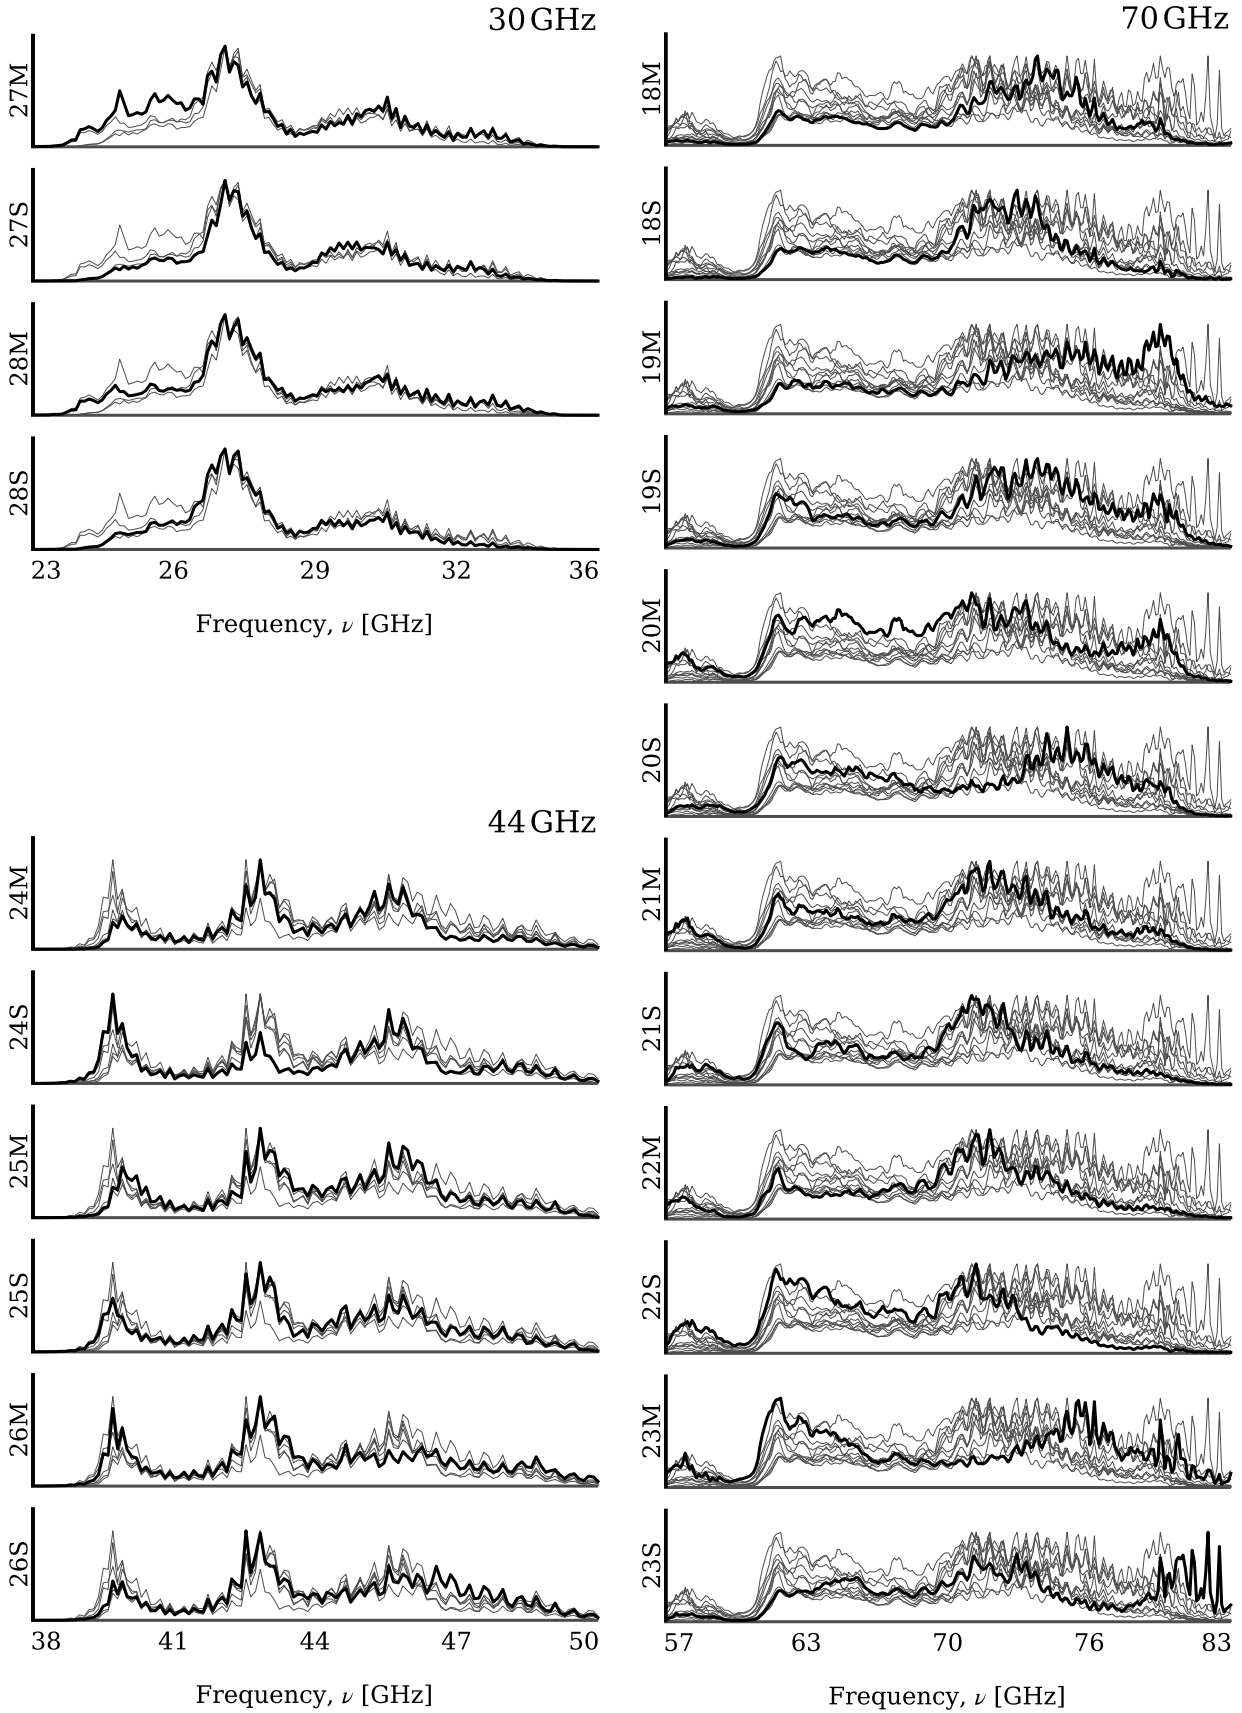


Fig. 1. Normalized radiometer bandpass profiles for LFI 30 GHz (*top left*), 44 GHz (*bottom left*), and 70 GHz (*right*) with a resolution of 100 MHz. Within each feed horn are two orthogonal linearly polarized components, denoted “M” and “S” for main- and side-arm radiometers, for measuring Q and U polarization. Individual bandpasses are highlighted in bold and the remaining are plotted behind for ease of comparison.

Table 1. Effective center frequencies, ν_{eff} , and unit conversion factors, U_c , for nominal and corrected bandpass profiles.

Quantity	Unit	Band	Nominal	Corrected	Change [%]
ν_{eff}	GHz	30	28.756	28.747	-0.03
		44	44.121	44.148	0.06
		70	70.354	70.795	0.63
U_c	K_{CMB}/K_{RJ}	30	1.0217	1.0217	0.00
		44	1.0515	1.0515	0.00
		70	1.1353	1.1369	0.14
U_c	$\text{MJy sr}^{-1}/K_{\text{CMB}}$	30	24.269	24.254	-0.06
		44	55.723	55.834	0.19
		70	129.22	132.45	2.49

the low-pass filtered profiles near the edges and extrapolating smoothly to zero. For the 70 GHz channel, the derivative is evaluated at 61 GHz to remove the BEM artefact discussed above.

Figure 3 shows a comparison of the raw and corrected bandpasses for each radiometer. Here, we see that the main change applied to the 30 GHz radiometers is the low-pass filter, as the edges were already quite well characterized in the original measurements. For the 44 GHz band, the most prominent correction is the high frequency cutoff. In particular, extrapolating the behavior for horn 24 is not trivial, as the profiles do not show a convincing downward trend before truncation. However, the profiles for both horn 25 and 26 suggest that the power drop-off is most likely just outside the measured range.

Edge trimming is also the dominant effect for the 70 GHz channels, both at low and high frequencies. The substantial amplitude and width of the 61 GHz spike feature implies that even the center frequency of the 70 GHz channel will change notably for this channel. This is quantified in Table 1 (see next section for exact definitions of each quantity). Here, we see that the 70 GHz center frequency increases by 0.44 GHz, or equivalently, by 0.6%, which translates into a difference in the conversion factor between flux density and thermodynamic temperature units, U_c , of 2.5%. The impact of these corrections in terms of the data quality and astrophysical products is considered in Sect. 4, where we find that the new profiles require smaller corrections during leakage removal. The Python script that applies the various corrections is available online³.

3. Methodology

The main goal of the BEYONDPLANCK project is to perform an end-to-end Bayesian analysis of the *Planck* LFI data. To do so, we start by writing down a physically motivated parametric model which can be fitted to our calibrated time-ordered data, \mathbf{d} , for radiometer, j :

$$\mathbf{d}_j = \mathbf{P}_j \mathbf{B}_j \int s(\nu) \tau_j(\nu) d\nu + \mathbf{n}_j. \quad (1)$$

Here, \mathbf{P} is a pointing matrix that maps the sky signal s into time domain; \mathbf{B} is the beam profile; and $\tau(\nu)$ is the bandpass. The sky signal as observed by radiometer j may be written in the following form,

$$s^j = \sum_{c=1}^{N_{\text{comp}}} \mathbf{a}_c \left[U_j \int f_c(\nu; \beta) \tau_j(\nu) d\nu \right] \equiv \sum_{c=1}^{N_{\text{comp}}} \mathbf{M}_{jc} \mathbf{a}_c = \mathbf{M}_j \mathbf{a}, \quad (2)$$

³ https://github.com/trygvels/BeyondPlanck_LFI_bandpass




Fig. 2. Bandpass profile for a 30 GHz radiometer measured by the standard approach (LARFM), versus the end-to-end cryogenic testing. Figure from Zonca et al. (2009).

where the sum runs over all relevant astrophysical sky components, c , each with its own unit conversion factor, U_j , that converts from its own intrinsic unit to the common brightness temperature unit adopted for s . As shown by (Planck Collaboration IX 2014), this may be written as

$$U_{ij} = \frac{\int \tau(\nu) \frac{dI_c}{dX_i} d\nu}{\int \tau(\nu) \frac{dI_c}{dX_j} d\nu}, \quad (3)$$

where dI_c/dX_i is the intensity derivative expressed in the given unit convention.

As defined by Eq. (2), \mathbf{M}_{cj} is referred to as the “mixing matrix,” and this matrix scales the set of component amplitudes \mathbf{a} to arbitrary frequencies using the spectral energy distribution, f , of each component, which may be described by some set of spectral parameters β . Using this notation, Eq. (1) can be written in the following slightly more compact form:

$$\mathbf{d}_j = \mathbf{P}_j \mathbf{B}_j \mathbf{M}_j \mathbf{a} + \mathbf{n}_j. \quad (4)$$

For a full discussion of the BEYONDPLANCK data model and notation, we refer to BeyondPlanck Collaboration (2023) and references therein.

In this paper, we are particularly interested in how beam and bandpass differences between detectors create leakage effects and how we can correct for these. The first complication in that respect is related to binning the time-ordered data, \mathbf{d} , into the Stokes parameter maps, \mathbf{m}_ν , using the mapmaking equation:

$$\left(\sum_{j \in \nu} \mathbf{P}_j^t \mathbf{N}_j^{-1} \mathbf{P}_j \right) \mathbf{m}_\nu = \sum_j \mathbf{P}_j^t \mathbf{N}_j^{-1} \mathbf{d}_j, \quad (5)$$

where we define $\mathbf{N}_j = \langle \mathbf{nn}^t \rangle$ as the (time-domain) noise covariance matrix for detector j .

This equation implicitly assumes that all radiometers, j , observe the same sky, s . However, although the various detectors actually do observe the same sky, their “response” to the spectral distribution of the different foreground components varies, because the bandpass and beam profiles differ, as described by

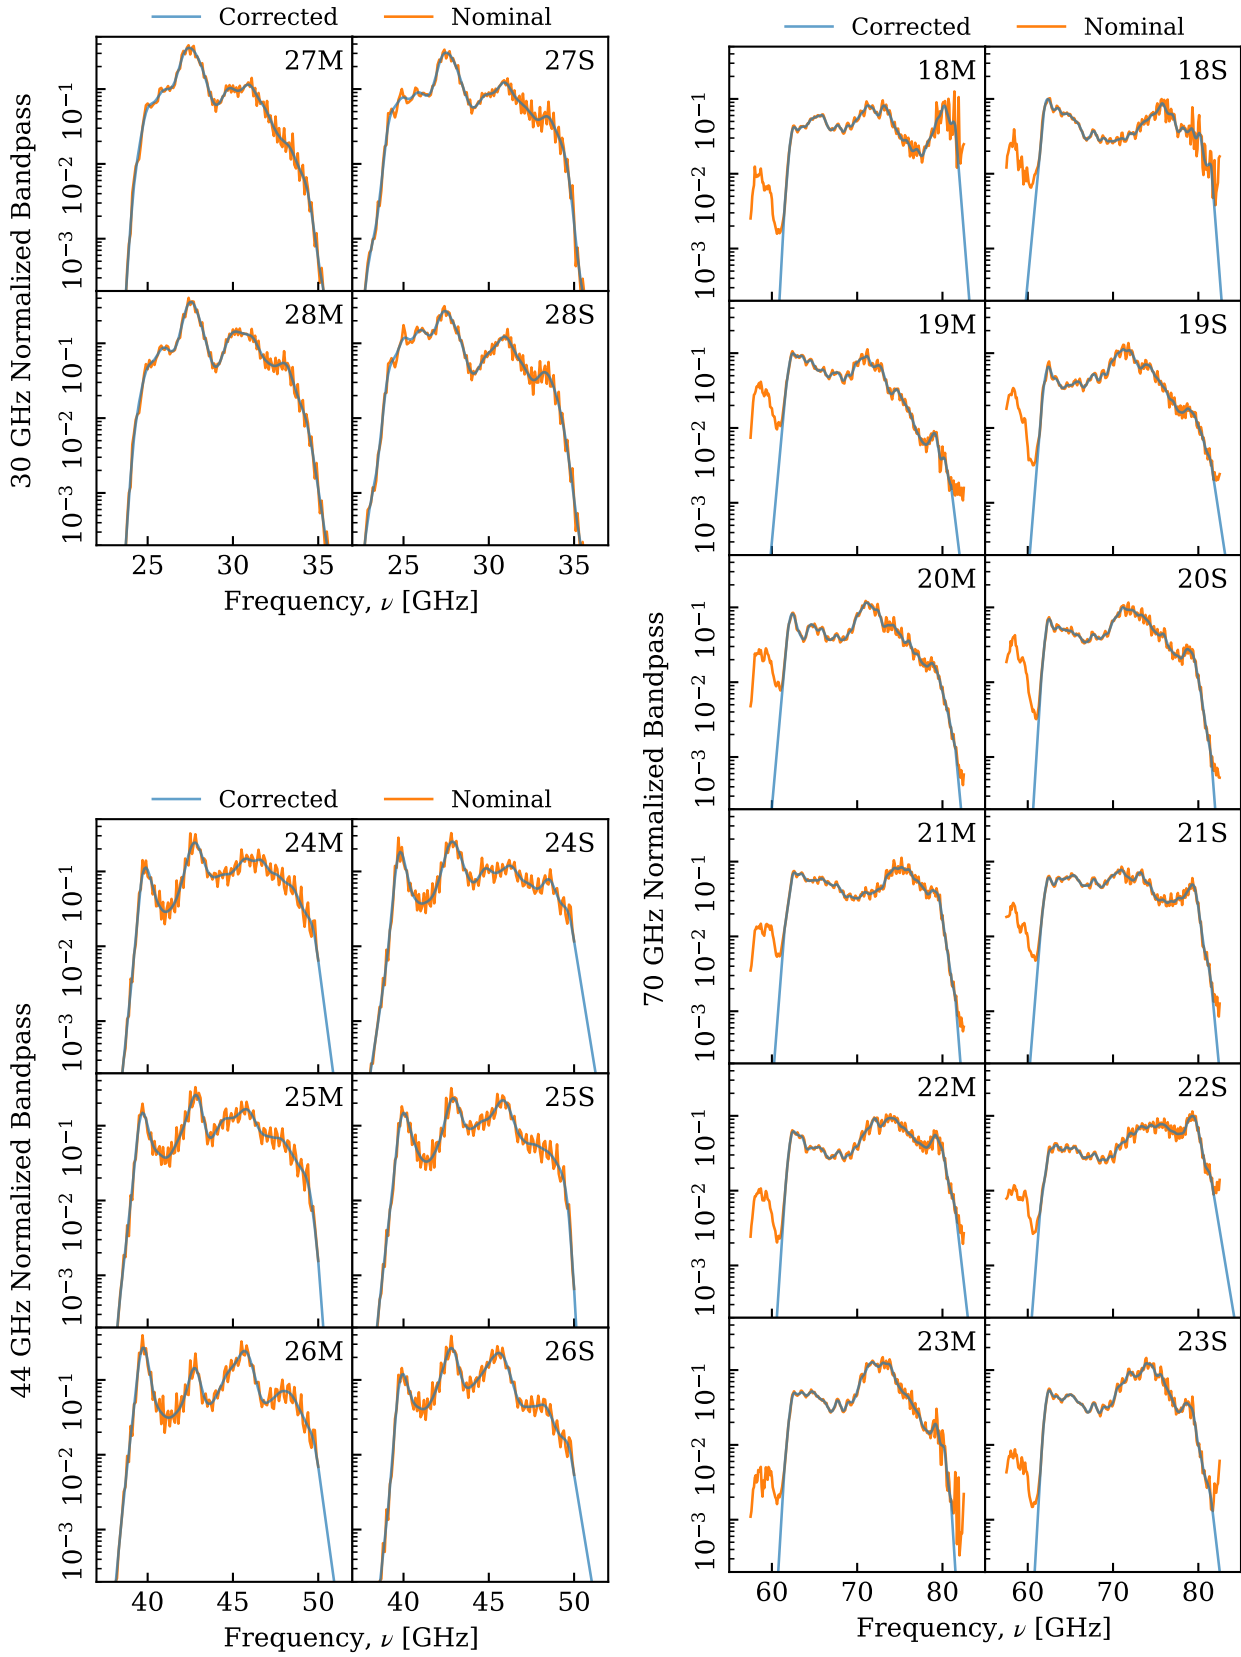


Fig. 3. Nominal and corrected bandpass profiles for LFI 30 GHz (top left), 44 GHz (bottom left), and 70 GHz (right) radiometers. Orange curves show raw bandpass profiles as provided by the *Planck* LFI DPC, and blue curves show the corrected profiles discussed in Sect. 2.

Eq. (1). When co-adding these different measurements into a sky map, any difference from the mean will be interpreted as noise by the mapmaker and, accordingly, it would be distributed between the various Stokes parameters according to the local scanning and noise levels at any given time. Furthermore, since these discrepancies are not stochastic but, rather, deterministically predictable by the bandpass and beam profiles, they do not average down over time. They therefore induce systematic errors in the final maps that are directly correlated with the sky signal itself. These effects are particularly significant for polarization analysis, where different bands must be combined to extract the very low Q and U signals. Thus, these errors are particularly worrisome for cosmological and astrophysical analyses.

We define three different effects caused by bandpass and beam errors. The first effect is referred to as the “bandpass mismatch,” and this describes the deterministic effect discussed above, namely, that different bandpass profiles create spurious leakage during multi-detector mapmaking, creating what is often referred to as “temperature-to-polarization leakage.” However, this name is somewhat of a misnomer, since all the Stokes parameters are formally coupled (see, e.g., [Planck Collaboration Int. LVII 2020](#)). Still, the effect is relatively much more important for polarization than for temperature because of its far lower signal-to-noise ratio. Indeed, this effect is the single strongest instrumental polarization contaminant for the LFI 30 GHz channel (see, e.g., Fig. 16 in [BeyondPlanck Collaboration 2023](#)). However, it is also entirely predictable and could, in principle, be removed to machine accuracy if both the astrophysical sky and the detector bandpasses were perfectly known.

However, as discussed above, the bandpasses are by no means perfectly known and those uncertainties create additional leakage that is not deterministically correctable. Furthermore, since all radiometer bandpasses within a band are uncertain, the combined co-added frequency bandpass is also uncertain. And this uncertainty produces an artefact when incorrectly translating the foreground sky model to the observed signal through the mixing matrix in Eq. (2). We refer to the effect caused by bandpass uncertainties as “bandpass errors,” and we attempt to minimize and marginalize over these by parameterizing the bandpasses. We then fit the associated free parameters as part of the main Gibbs sampling process.

The third and final leakage effect considered here is “beam mismatch,” which is a deterministic leakage effect arising from differences between the main beams of the radiometers in a given frequency channel. These are in principle similar to the bandpass mismatch effect and give rise to small-scale artefacts for main beam variations and large-scale artefacts for sidelobe variations. In this paper, we only account for static FWHM differences in the main beam between radiometers, but not asymmetric beams (these will be discussed in future publications, along with “beam errors,” i.e., uncertainties in the actual beam profiles). For a real-world example of the importance of detector sidelobe differences, we refer to the WMAP re-analysis reported by [Watts et al. \(2023\)](#).

3.1. Leakage corrections

As discussed above, bandpass mismatch artefacts arise because Eq. (5) assumes that all radiometers measure the same signal in any given pixel of the sky; however, in reality, they do not because their different bandpass and beam profiles couple differently with the SED of foreground emissions. To account for these

differences, we define the following leakage correction term for each radiometer:

$$\delta s_{j,t}^{\text{leak}} = \mathbf{P}_{tp}^j \mathbf{B}_{pp'}^j \left(s_{jp'}^{\text{sky}} - \langle s_{jp'}^{\text{sky}} \rangle \right). \quad (6)$$

Here, s_j^{sky} is a model of the sky as actually seen by radiometer j at pixel p , taking into account its specific bandpass and beam profile, and angled brackets denote an average over all radiometers evaluated pixel-by-pixel and t denotes a single time sample. We note that this leakage term explicitly accounts for both bandpass and beam differences via s_j^{sky} and $\mathbf{B}_{pp'}^j$.

In order to create a leakage-cleaned frequency map, we simply subtract this leakage term from the calibrated data prior to mapmaking:

$$\left(\sum_{j \in \nu} \mathbf{P}_j^t \mathbf{N}_j^{-1} \mathbf{P}_j \right) \mathbf{m}_\nu = \sum_{j \in \nu} \mathbf{P}_j^t \mathbf{N}_j^{-1} (\mathbf{d}_j - \delta \mathbf{s}_j^{\text{leak}}). \quad (7)$$

In this equation, the right-hand side corresponds to a stationary sky signal in which all radiometers see the same effective sky signal, defined by the mean over all detectors, while the mean itself is not affected by the leakage correction, since $\delta \mathbf{s}_j^{\text{leak}}$ sums to zero by construction.

These corrections are conceptually similar to those applied by the *Planck* LFI DPC ([Planck Collaboration II 2016, 2020](#)) and *Planck* DR4 ([Planck Collaboration Int. LVII 2020](#)) pipelines, although implementation-wise they differ significantly from both. Firstly, neither of the two previous pipelines applies any beam leakage correction, while in the current work we account for the different beam FWHMs for each radiometer ([Planck Collaboration III 2016](#)) when evaluating Eq. (6). Secondly, while both of the previous pipelines use a linear approximation to the component SEDs to evaluate the bandpass effects, we evaluate the full integral in Eq. (2) for each case, as described in [Galloway et al. \(2023\)](#). Thirdly, while the DPC pipelines makes separate maps for the raw TOD and the leakage correction, and subtract the latter as a post-processing step, we applied the corrections directly in time-domain prior to mapmaking, as in *Planck* DR4. Fourthly, while the DPC approach does not make any direct adjustments to the bandpass profiles themselves, we fit a parametric function for each of these, as described in the next section as part of the larger sampling framework, and thereby achieve a better overall fit. The *Planck* DR4 achieves a similar goal by fitting separate sky signal template terms for the mean and derivative of the foreground SED evaluated at the center frequency as part of their generalized mapmaker ([Planck Collaboration Int. LVII 2020](#)). Finally, neither of the previous pipelines allow for uncertainties in the astrophysical sky model itself, but rather adopts a Commander sky model as a static conditional input.

Figure 4 shows as comparison between the uncorrected (top row) and leakage corrected (second row) LFI 30 GHz map, as well as the individual contributions from bandpass leakage corrections (third row) and beam leakage corrections (bottom row). Comparing these maps, we immediately see that the bandpass leakage corrections are highly significant along the Galactic plane, with corrections on the order of unity. Indeed, the corrections are sufficiently large that the central Galactic plane switches sign in both Stokes Q and U . This makes any estimates of spectral parameters highly dependent on the leakage corrections, and joint estimation of both leakage corrections and spectral parameters is essential; for a specific discussion of Bayesian estimates of the spectral index

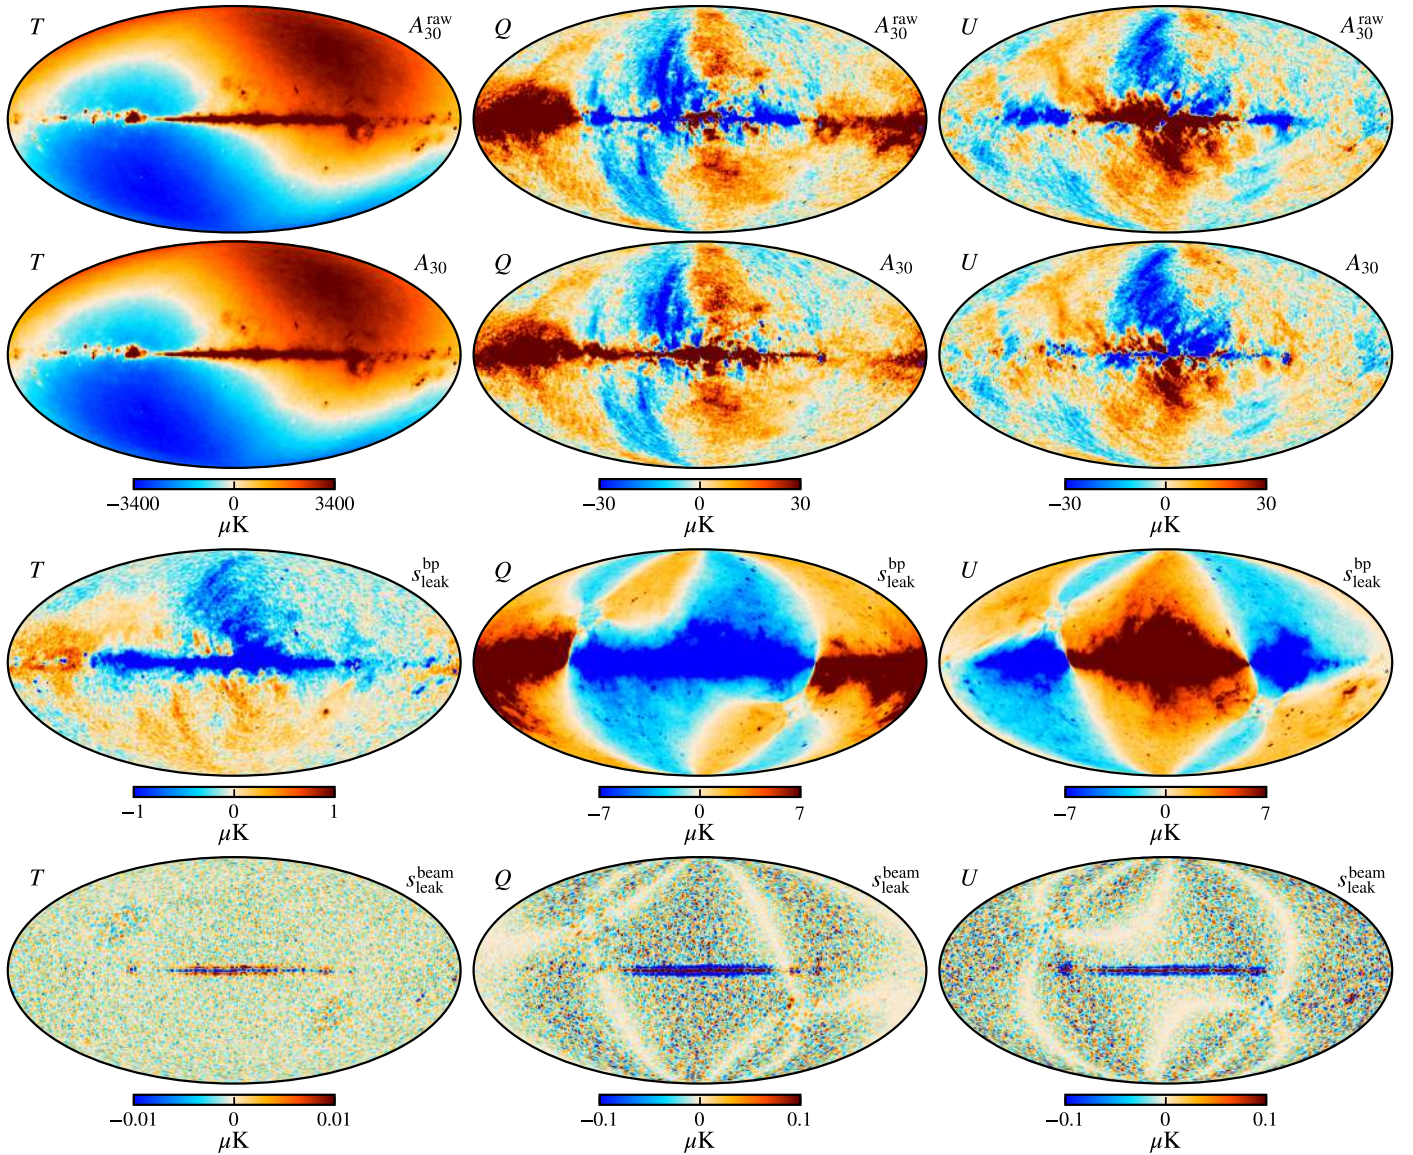


Fig. 4. Comparison of the raw (*top row*) and leakage corrected (*second row*) *Planck* 30 GHz channel. The *third and fourth rows* show the individual bandpass and beam leakage corrections, respectively. Columns shows Stokes T , Q , and U parameters. All maps have been smoothed to an angular resolution of 1° FWHM, except for the bottom row, which has been smoothed to 2° FWHM.

of polarized synchrotron emission (see [Svalheim et al. 2023](#)). Likewise, the bandpass leakage corrections at high Galactic latitudes are dominated by the foreground monopoles, most notably those of synchrotron, AME, and free-free emission; thus, consistent and joint estimation of monopoles and leakage corrections is therefore essential for large-scale CMB extraction. A novel feature of the BEYONDPLANCK processing is component-based monopole determination, which allows more easily self-consistent zero-level estimation across frequencies; for a discussion of this approach, see [Andersen et al. \(2023\)](#).

Comparing the bottom two rows of Fig. 4, it is visually obvious that the bandpass corrections are generally much more important than the beam errors on large angular scales. However, the beam corrections can also be important for CMB power spectrum analysis on smaller angular scales. This is illustrated in Fig. 5, which shows a $30^\circ \times 30^\circ$ zoom-in of the 44 GHz leakage correction map (top right panel) near the galactic south pole. This effect is particularly strong for the 44 GHz channel, because one of its three feeds (horn 24) has a significantly smaller beam

width ($24'$ FWHM) than the other two (horns 25 and 26; $30'$ FWHM), as reported by [Planck Collaboration IV \(2016\)](#).

In this plot, we see that the spurious beam-induced small-scale polarization fluctuations are typically at a level of $1\text{--}2\ \mu\text{K}$. In addition, these fluctuations are primarily seeded by CMB temperature fluctuations, as seen by comparing the top two panels. The in-set circle provides a visual guide that makes it easy to identify correlations by eye. These spurious temperature-to-polarization leakage fluctuations induce spurious small-scale polarization modes. The bottom panel in Fig. 5 compares the angular power spectrum of the 44 GHz channel with the *Planck* 2018 best-fit ΛCDM EE spectrum, as well as a scaled version of the TT spectrum. The fact is that not only can uncorrected beam leakage correction confuse EE and BB measurements, but they can also obviously be a significant contaminant for TE , TB , and EB correlations, and thereby contaminate constraints on both standard and non-standard physics, for instance, gravitational lensing (e.g., [Planck Collaboration Int. XLI 2016](#)) or birefringence (e.g., [Minami & Komatsu 2020](#)).




Fig. 5. *Top left:* Gnomonic projection of a $30^\circ \times 30^\circ$ area centered on $(l, b) = (0^\circ, -70^\circ)$ in the BEYONDPLANCK posterior mean CMB temperature map (Colombo et al. 2023). *Top right:* BEYONDPLANCK 44 GHz Q-leakage map in the same area. *Bottom:* Angular polarization power spectra of the masked 44 GHz leakage map, compared with the *Planck* best-fit Λ CDM TT (scaled by 10^{-3}) and EE spectra.

3.2. Monte Carlo sampling of parametric bandpass models

Since each radiometer bandpass is measured with non-negligible uncertainties, our leakage correction model as defined by Eq. (6) is not perfect. For instance (as discussed in Sect. 2), the modifications made to the 70 GHz bandpasses in this paper changes its overall center frequency by 0.6%, and this is likely to have a non-trivial impact on foreground estimates derived from these data, whereas relative measurement errors between individual detectors will create bandpass leakage (as discussed above).

It is therefore important to parameterize the bandpasses themselves, fit the free parameters to the data, and marginalize over the corresponding uncertainties. Obviously, bandpasses have an indefinite number of degrees of freedom, being essentially free functions of frequency, and flight data are far from sufficient to constrain these frequency-by-frequency. It is therefore customary to base parametric models on the laboratory measurements, τ_0 , and only apply relatively mild corrections to these. For instance, Planck Collaboration X (2016) introduced a simple linear shift model:

$$\tau(\nu) = \tau_0(\nu + \Delta_{\text{bp}}), \quad (8)$$

where Δ_{bp} denotes a linear shift in frequency space. Another possible model is a power-law tilt model:

$$\tau(\nu) = \tau_0(\nu) \left(\frac{\nu}{\nu_c} \right)^n, \quad (9)$$

where ν_c is the center frequency, and n is a spectral index. We have implemented support for both models in our codes, but we focus only on the linear shift model as introduced in Planck Collaboration IX (2016) in the following. Alternative models will be explored further in future work.

For the linear shift model, we split the total shift for radiometer j into two components,

$$\Delta_{\text{bp}}^j = \bar{\Delta}_{\text{bp}} + \delta_{\text{bp}}^j. \quad (10)$$

Here, $\bar{\Delta}_{\text{bp}}$ corresponds to an absolute frequency shift for the overall co-added frequency band, while δ_{bp}^j is a relative frequency shift for radiometer j only, with the constraint that $\sum_j \delta_{\text{bp}}^j = 0$.

Absolute and relative bandpass corrections generally have very different impacts on the final sky maps. Intuitively, an absolute frequency shift can be interpreted as a “foreground-only calibration change,” in the sense that each foreground component becomes either weaker or brighter in the given frequency channel. The archetypal signature of an absolute bandpass error is that the residual map ($\mathbf{d}_\nu - \mathbf{s}_\nu$) shows an imprint of the Galactic dipole, but there is no corresponding imprint of a residual CMB dipole; if there are both a CMB dipole and a Galactic plane imprint, then the problem is an absolute gain error. In contrast, relative bandpass errors primarily lead to temperature-to-polarization leakage, as visualized in Fig. 4, with a pattern defined by the Galactic intensity foregrounds modulated by the scanning strategy and polarization angle of the experiment in question.

Based on these observations, we employed different sampling techniques for $\bar{\Delta}_{\text{bp}}$ and δ_{bp}^j , each with its own likelihood and proposal density. In both cases, however, we employed a standard Metropolis MCMC sampler with a tuned covariance matrix (see Appendix A in BeyondPlanck Collaboration 2023), and we sampled the absolute and relative corrections separately and alternately in each main Gibbs step; this interleaved sampling is exclusively for the convenience of implementation, as the code becomes simpler by considering one type of corrections at any given time.

3.3. Absolute bandpass sampling

To derive a likelihood for the $\bar{\Delta}_{\text{bp}}$ parameter, we note that this parameter conditionally only affects the sky signal model through the mixing matrix in Eq. (2)⁴, and we may therefore form the following data-minus-signal residual,

$$\mathbf{r}_\nu = \mathbf{d}_\nu - \mathbf{s}_\nu(\bar{\Delta}_{\text{bp}}), \quad (11)$$

where

$$\mathbf{s}_\nu = \sum_c \mathbf{M}_{c_j}(\beta, \Delta_{\text{bp}}^j) \mathbf{a}^c. \quad (12)$$

Under the assumption that \mathbf{r}_ν may be modeled as Gaussian noise, as per Eq. (1), this results in a log-likelihood of the following form:

$$-2 \ln \mathcal{L}(\bar{\Delta}_{\text{bp}}) = (\mathbf{d}_\nu - \mathbf{s}_\nu(\bar{\Delta}_{\text{bp}}))^T \mathbf{N}_\nu^{-1} (\mathbf{d}_\nu - \mathbf{s}_\nu(\bar{\Delta}_{\text{bp}})) = \chi^2(\bar{\Delta}_{\text{bp}}), \quad (13)$$

up to an irrelevant constant. Absolute bandpass sampling is thus equivalent to a standard χ^2 fit, and we may sample from the corresponding conditional distribution through standard Metropolis

⁴ In principle, the bandpass also affects the beam calculations, but since we do not apply any stochastic beam corrections in this paper, we ignore this effect for now.




Fig. 6. Mask used for bandpass correction sampling. The accepted sky fraction is $f_{\text{sky}} = 0.953$.

MCMC, with an acceptance rate of

$$a = \min(1, e^{-\frac{1}{2}(\chi_{i+1}^2 - \chi_i^2)}). \quad (14)$$

In practice, we applied a Galactic mask in these evaluations by setting $N_p^{-1} = 0$ for masked pixels. However, since the calibration signal in question is precisely the Galactic plane, it is desirable to include as much of the sky as possible and we only excluded the central parts of the Milky Way (within which foreground modeling is very complicated), as well as particularly bright point sources such as Tau-A. The actual mask used is shown as a black region in Fig. 6 and it excludes 4.7% of the sky.

Finally, following [Planck Collaboration X \(2016\)](#), we only estimated an absolute correction for the LFI 30 GHz channel. Allowing all channel corrections to be fitted freely is impossible, as this would result in perfect degeneracies between the bandpass parameters and the foreground SED parameters. For this study in particular, the decision to omit high-frequency Planck measurements severely limits thermal dust spectral index constraints, which could potentially alleviate the degeneracy. The motivation for fitting the 30 GHz channel is simply that this parameter turns out to be sufficiently non-degenerate to allow a robust fit, while at the same time non-negligible residuals occur when it is not fitted. The 30 GHz bandpass parameter is in practice constrained by its nearest WMAP channel, namely the *Ka*-band at 33 GHz.

3.4. Relative bandpass sampling

For the relative bandpass corrections, δ_{bp}^j , we adopt an alternative likelihood that is inspired by the spurious map approach pioneered by [Page et al. \(2007\)](#) for CMB polarization purposes. The motivation for this is that relative bandpass corrections primarily have a real-world impact for polarization, while a full and formally statistically correct χ^2 fit, as defined by Eq. (13), is vastly dominated by the high signal-to-noise ratio (S/N) of the intensity signal. In practice, a global fit would therefore tend to use the extra degrees of freedom to fit intensity foreground SED errors, rather than large-scale polarization artefacts.

As noted by [Page et al. \(2007\)](#), relative bandpass errors cause intensity-to-polarization leakage with a very specific observational signature, namely, that the spurious polarization signal does not depend on the polarization angle orientation of a given detector, but only on its bandpass properties. They use this to define an additional correction map for each radiometer, S_j , that corresponds to the difference between the intensity signal seen

by detector, j , and the mean over all radiometers. Thus, the time-domain signal measured by detector, j , may be written in the form:

$$s_j = T + Q \cos 2\psi_j + U \sin 2\psi_j + \sum_{i=1}^{N_{\text{det}}-1} S_i \delta_{ij}, \quad (15)$$

and the pointing matrix \mathbf{P} in Eq. (1) may be modified accordingly.

In general, one can only solve for $N - 1$ spurious sky signals, where N is the number of radiometers in a band, to avoid a perfect degeneracy with the mean intensity signal. For simplicity, we therefore consider a minimal case with $N = 2$. In this case, the single-pixel mapmaking equation reads:

$$\begin{bmatrix} 1 & \cos 2\psi & \sin 2\psi & \delta_{1j} \\ \cos 2\psi & \cos^2 2\psi & \cos 2\psi \sin 2\psi & \cos 2\psi \delta_{1j} \\ \sin 2\psi & \sin 2\psi \cos 2\psi & \sin^2 2\psi & \sin 2\psi \delta_{1j} \\ \delta_{1j} & \cos 2\psi \delta_{1i} & \sin 2\psi \delta_{1j} & \delta_{1j} \end{bmatrix} \begin{bmatrix} T \\ Q \\ U \\ S_1 \end{bmatrix} = \begin{bmatrix} d \\ d \cos 2\psi \\ d \sin 2\psi \\ d \delta_{1j} \end{bmatrix}. \quad (16)$$

For the tightly interconnected WMAP scanning strategy, this equation may be solved pixel-by-pixel without inducing a prohibitive noise penalty and as a result, [Page et al. \(2007\)](#) simply chose to deliver polarization sky maps that are explicitly marginalized over S . However, this is not possible for the *Planck* scanning strategy, for which the polarization angle of a given detector only varies by a few tens of degrees over large areas of the sky ([Planck Collaboration I 2014](#)). For these, the condition number of the matrix in Eq. (16) leads to a massive noise increase, to the point that the map becomes unusable for astrophysical and cosmological analyses.

However, even though the noise per pixel is excessive, the aggregated signal-to-noise ratio in S_j across the full sky is still high. In this paper, we propose using the spurious map approach instead to fit the small number of relative bandpass shifts through the following goodness-of-fit quantity:

$$\chi^2 = \sum_{j=1}^{N_{\text{det}}-1} \sum_p \left(\frac{S_j(p)}{\sigma_j(p)} \right)^2. \quad (17)$$

Here, $\sigma_j(p)$, is the uncertainty arising from the *IQUS* solution above, which is defined as the diagonal element of the inverse coupling matrix in Eq. (16). The rest of the algorithm is identical to that described in the previous section, just with a different χ^2 expression in Eq. (14).

Intuitively, this approach combines the *Planck* idea of using a parametric foreground model to perform relative bandpass corrections with the spurious map approach from WMAP with the goal of optimizing the free parameters using only polarization information. As such, the algorithm significantly reduces temperature-to-polarization leakage through a very small number of additional degrees of freedom and a negligible noise boost.

4. Results

4.1. Posterior summary

We are now finally ready to present the results from the above algorithms as applied within the BEYONDPLANCK analysis framework, and we start by inspecting the resulting




Fig. 7. Markov chain trace plots for bandpass and spectral index parameters. The two chains are indicated with different opacities.

Markov chains and their internal correlations. As described by [BeyondPlanck Collaboration \(2023\)](#), two independent Monte Carlo chains are produced in the main BEYONDPLANCK processing, each resulting in 750 samples. Figure 7 shows a representative subset of these, where in addition to the individual detector bandpass shifts, we also include the AME ν_p parameter in temperature, as well as the spectral index in polarization for both thermal dust and synchrotron, β_d and β_s , respectively. For an in-

depth discussion on these results, we refer the interested reader to [Svalheim et al. \(2023\)](#).

Overall, we see that the correlation length is substantial in these chains, and it is clear that several of the Metropolis step lengths would benefit from further optimization in a future run. Indeed, using the current chains to tune a Metropolis proposal matrix for a future iteration of the analysis appears to be promising. However, there is generally no doubt that bandpass corrections are among the most difficult parameters to sample in the entire BEYONDPLANCK data model, because bandpass corrections are global and they do not only affect leakage corrections, but also the unit conversions, thereby impacting the entire foreground model. This global impact leads to a long Markov chain correlation length.

In Fig. 8, we plot Pearson’s correlation coefficient, p , for various parameter pairs. Overall, we see that the bandpass parameters are most strongly correlated with the AME peak frequency, ν_p , in intensity, and the synchrotron and thermal dust spectral indices, β_s and β_d , in polarization for which correlations around $|p| \approx 0.5$ are observed. An in-depth discussion on temperature foreground degeneracies can be found in [Andersen et al. \(2023\)](#). We also note that this plot serves as a powerful reminder of the usefulness of global cross-experiment analysis, as improved constraints on either ν_p , β_s or β_d from, for instance, *Planck* HFI ([Planck Collaboration III 2020](#)), C-BASS ([King et al. 2010](#)), or QUIJOTE ([Génona-Santos et al. 2015](#)) would translate directly into improved constraints on the *Planck* LFI bandpasses and, therefore, better maps overall. Generally speaking, older data sets may almost always be improved when new experiments become available.

The average properties of the Markov chains are listed in Table 2, with the posterior mean and standard deviation for each bandpass correction parameter, both for co-added frequency channels and individual radiometers, and the same information is visualized in Fig. 9.

Figure 10 compares the BEYONDPLANCK posterior mean leakage maps with the corresponding *Planck* 2018 LFI leakage maps for all three LFI channels. As discussed in Sect. 3.1, there are several algorithmic differences between these two pipelines. Firstly, the *Planck* 2018 approach does not account for beam mismatch leakage and, thus, we see significantly fewer small-scale fluctuations in these maps compared to BEYONDPLANCK; this is particularly striking in the 44 GHz channel. Secondly, we note that the *Planck* 2018 correction maps exhibit significantly weaker large-scale corrections at high Galactic latitudes compared to BEYONDPLANCK. We have traced this effect down to a difference in the net foreground monopole in the two approaches and, in particular, we have found that the BEYONDPLANCK 30 GHz leakage map takes a very similar shape if we add an artificial negative monopole of $-200 \mu\text{K}$ to the AME amplitude map. In this respect, we note that the *Planck* 2018 pipeline uses the Madam map-making code ([Keihänen et al. 2005](#)), which explicitly removes all monopole terms from the co-added map ([Planck Collaboration II 2020](#)). At the same time, the Commander-based foreground maps used to generate the correction templates do include a physical estimate of the monopoles. To account for this issue, the DPC processing generated a new set of correction maps using the *Planck* levelS simulation package ([Reinecke et al. 2006](#)), creating TODs from the Commander leakage map and correction factors. The next step was to bin these into maps using the same Madam map-making code, resulting in correction maps without the monopole term from which the final correction templates were generated. In retrospect, it appears that this rather involved pipeline did underestimate the




Fig. 8. Pearson's correlations between bandpass correction and other parameters in the global BEYONDPLANCK model as evaluated directly from the Markov chain.

monopole contribution in the final maps, and this serves as a useful reminder of an important benefit of an integrated end-to-end pipeline: passing data objects from one operation to the next in a self-consistent manner becomes much more transparent when all parts of the code use the same data model.

4.2. Comparison of nominal and corrected LFI bandpasses

As discussed in Sect. 2, we applied a number of pre-processing steps to the *Planck* LFI bandpasses before performing the main BEYONDPLANCK end-to-end analysis. In this section, we con-

Table 2. Bandpass correction posterior means and standard deviations for all *Planck* LFI frequency bands and individual radiometers.

Radiometer	Δ_{bp} [GHz]
LFI 30 GHz	
27M	0.24 ± 0.03
27S	0.28 ± 0.03
28M	0.47 ± 0.03
28S	0.08 ± 0.03
28S	0.13 ± 0.03
LFI 44 GHz	
24M	-0.55 ± 0.02
24S	0.16 ± 0.02
25M	0.28 ± 0.03
25S	0.07 ± 0.02
26M	0.09 ± 0.01
26S	-0.05 ± 0.01
LFI 70 GHz	
18M	-0.87 ± 0.17
18S	0.14 ± 0.12
19M	0.73 ± 0.16
19S	0.59 ± 0.12
20M	-0.27 ± 0.12
20S	0.20 ± 0.22
21M	0.06 ± 0.20
21S	-0.71 ± 0.08
22M	0.36 ± 0.32
22S	-0.32 ± 0.16
23M	0.13 ± 0.13
23S	-0.05 ± 0.05

Notes. We note that the absolute bandpass corrections at 44 and 70 GHz are fixed to zero.




Fig. 9. Estimated bandpass corrections for each LFI radiometer. Error bars indicate $\pm 3\sigma$ uncertainties. These parameters are constrained to have vanishing mean within each frequency (except for 30 GHz, which has an equal absolute shift to all radiometers as well), and are as such strongly correlated.

sider the impact of these changes in terms of some key data products that highlight their effects.

The first quantity we considered is the spurious map, S_j , as defined by Eq. (16). These are shown for the 30 GHz channel in Fig. 11 for three different analysis configurations, all in the form of S_j/σ_j , masked by the bandpass processing mask and smoothed to 3° FWHM. The top row shows the three 30 GHz spurious maps that result if we attempt to produce maps with no bandpass corrections at all and simply take the publicly available

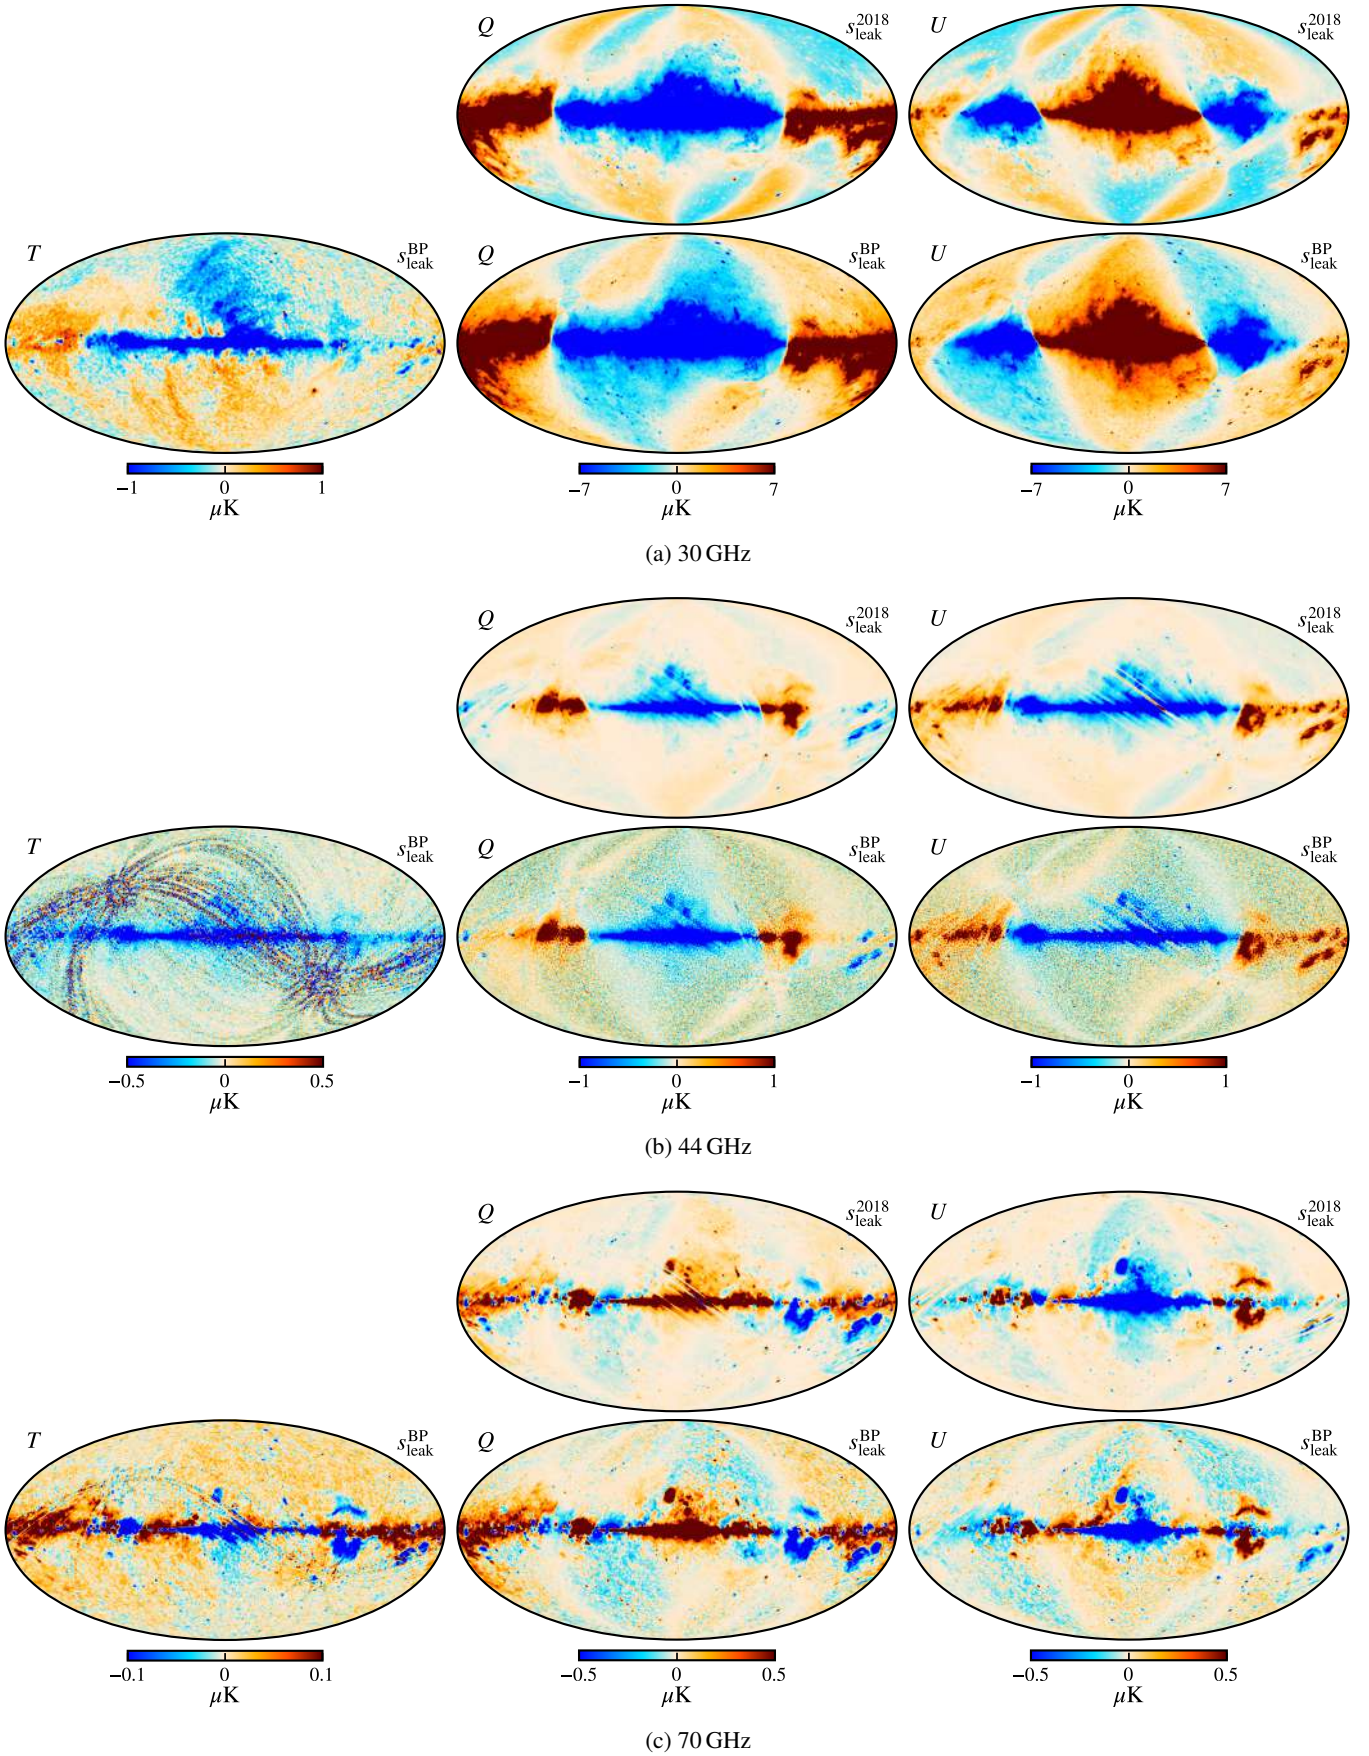


Fig. 10. Full leakage correction maps for the 30 GHz (a), 44 GHz (b) and 70 GHz (c) bands in Stokes T , Q , and U from the official *Planck* 2018 processing (top) and BEYONDPLANCK (bottom). All maps have been smoothed to an angular resolution of 1° FWHM.

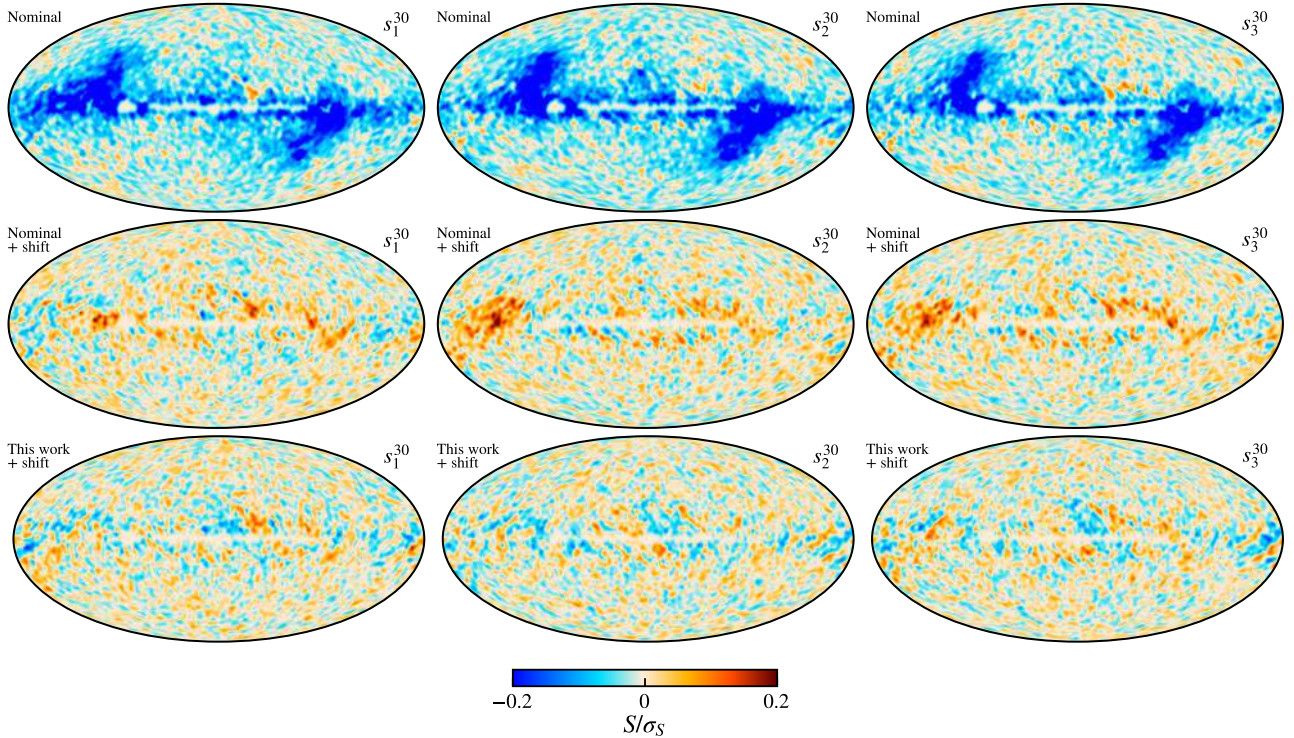


Fig. 11. Spurious maps for the 30 GHz band with nominal profiles without corrections (*top*), nominal profiles with corrections (*middle*) and new profiles with corrections (*bottom*) using the processing mask showed in Fig. 6 ($f_{\text{sky}} = 0.953$), and smoothed with a gaussian beam of 3° FWHM. This quantity is unitless, as it is defined as S/σ_s .




Fig. 12. Estimated bandpass corrections for the 70 GHz radiometers. These parameters are constrained to have vanishing mean and, as such, they are strongly correlated.

profiles at face value. Here, we see coherent structures that are clearly inconsistent with noise and morphologically associated with the Galactic plane and the so-called “*Planck* Deep Fields” (centered on the Ecliptic poles). In these regions, the S/N is vastly higher than in the Ecliptic plane because of the dense polarization angle sampling of the *Planck* scanning strategy. The second row shows a similar case, using the nominal LFI bandpasses, but this time allowing a frequency shift per radiometer, as described in Sect. 3. The residuals are clearly reduced, indicating that the fitting algorithm does work as expected. Finally, the bottom row shows the same, but now using the corrected LFI bandpasses from Sect. 2 as input. At this point, the residuals are significantly closer to randomly distributed, with only small residuals appearing near the Galactic center.

Figure 12 shows a comparison of posterior mean and standard deviations for the 70 GHz bandpass corrections when using the nominal (orange points) and corrected (blue points) profiles. Here, we see that the nominal profiles generally require much larger correction factors than the corrected profiles. In particular, 18M sees $\Delta_{\text{bp}} \sim 1.3$ GHz using the nominal profiles as opposed to ~ 0.8 GHz for the new ones. This is not surprising, considering that this profile exhibits some of the strongest systematic artefacts in Fig. 1.

5. Conclusions

In this paper, we discuss how bandpass and beam leakage effects may be mitigated in a Bayesian CMB analysis pipeline and we summarize the results from the BEYONDPLANCK pipeline. As a preparatory step, we have also provided a set of corrected LFI bandpass profiles that correct for known systematic features dating back to the ground calibration phase and previously reported in the literature. The most notable result from these corrections is that the overall 70 GHz center frequency is shifted up by 0.6%, and simple map-based comparisons show that differences in the resulting frequency maps are on the order of a couple of microkelvins.

We argue that the proposed algorithms are substantially simpler to implement compared to previous methods, thanks to the tight integration between astrophysical component separation and low-level mapmaking; all the required components are already available from different parts of the BEYONDPLANCK framework. As a result, the leakage correction may actually be defined in terms of two or three very simple equations, and the practical code implementation amounts to a few hundred lines of code. We also note that beam leakage corrections are trivial to implement, simply by accounting for the different detector beam responses when scanning the model sky with the pointing operator.

We believe that the importance of these methods will become increasingly critical for next-generation experiments. While the current *Planck* LFI polarization observations are intrinsically noise-dominated, and the various corrections discussed in this paper are relatively minor compared to the overall noise level, the same will not hold true for future *B*-mode experiments such as LiteBIRD or CMB-S4. For these, establishing highly accurate bandpass and beam leakage corrections will be absolutely essential in order to reach the required accuracy in terms of nanokelvin.

Acknowledgements. We thank Prof. Pedro Ferreira and Dr. Charles Lawrence for useful suggestions, comments and discussions. We also thank the entire *Planck* and WMAP teams for invaluable support and discussions, and for their dedicated efforts through several decades without which this work would not be possible. The current work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement numbers 776282 (COMPET-4; BEYONDPLANCK), 772253 (ERC; BITS2COSMOLOGY), and 819478 (ERC; COSMOGLOBE). In addition, the collaboration acknowledges support from ESA; ASI and INAF (Italy); NASA and DoE (USA); Tekes, Academy of Finland (grant no. 295113), CSC, and Magnus Ehrnrooth foundation (Finland); RCN (Norway; grant nos. 263011, 274990); and PRACE (EU).

References

- Ade, P. A. R., Ahmed, Z., Amiri, M., et al. 2021, *Phys. Rev. Lett.*, **127**, 151301
- Andersen, K. J., Herman, D., Aurlien, R., et al. 2023, *A&A*, **675**, A13 (BeyondPlanck SI)
- BeyondPlanck Collaboration (Andersen, K. J., et al.) 2023, *A&A*, **675**, A1 (BeyondPlanck SI)
- Colombo, L. P. L., Eskilt, J. R., Paradiso, S., et al. 2023, *A&A*, **675**, A11 (BeyondPlanck SI)
- Delouis, J. M., Pagano, L., Mottet, S., Puget, J. L., & Vibert, L. 2019, *A&A*, **629**, A38
- Eriksen, H. K., O'Dwyer, I. J., Jewell, J. B., et al. 2004, *ApJS*, **155**, 227
- Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, **676**, 10
- Galloway, M., Andersen, K. J., Aurlien, R., et al. 2023, *A&A*, **675**, A3 (BeyondPlanck SI)
- Geman, S., & Geman, D. 1984, *IEEE Trans. Pattern Anal. Mach. Intell.*, **6**, 721
- Génova-Santos, R., Rubiño-Martín, J. A., Rebolo, R., et al. 2015, *MNRAS*, **452**, 4169
- Kamionkowski, M., & Kovetz, E. D. 2016, *ARA&A*, **54**, 227
- Keihänen, E., Kurki-Suonio, H., & Poutanen, T. 2005, *MNRAS*, **360**, 390
- King, O. G., Copley, C., Davies, R., et al. 2010, *SPIE Conf. Ser.*, **7741**, 1
- Minami, Y., & Komatsu, E. 2020, *Phys. Rev. Lett.*, **125**
- Page, L., Hinshaw, G., Komatsu, E., et al. 2007, *ApJS*, **170**, 335
- Penzias, A. A., & Wilson, R. W. 1965, *ApJ*, **142**, 419
- Planck Collaboration I. 2014, *A&A*, **571**, A1
- Planck Collaboration II. 2014, *A&A*, **571**, A2
- Planck Collaboration IX. 2014, *A&A*, **571**, A9
- Planck Collaboration XII. 2014, *A&A*, **571**, A12
- Planck Collaboration II. 2016, *A&A*, **594**, A2
- Planck Collaboration III. 2016, *A&A*, **594**, A3
- Planck Collaboration IV. 2016, *A&A*, **594**, A4
- Planck Collaboration IX. 2016, *A&A*, **594**, A9
- Planck Collaboration X. 2016, *A&A*, **594**, A10
- Planck Collaboration I. 2020, *A&A*, **641**, A1
- Planck Collaboration II. 2020, *A&A*, **641**, A2
- Planck Collaboration III. 2020, *A&A*, **641**, A3
- Planck Collaboration IV. 2020, *A&A*, **641**, A4
- Planck Collaboration Int. XLI. 2016, *A&A*, **596**, A102
- Planck Collaboration Int. LVII. 2020, *A&A*, **643**, A42
- Reinecke, M., Dolag, K., Hell, R., Bartelmann, M., & Enßlin, T. A. 2006, *A&A*, **445**, 373
- Svalheim, T. L., Andersen, K. J., Aurlien, R., et al. 2023, *A&A*, **675**, A14 (BeyondPlanck SI)
- Tristram, M., Banday, A. J., Górski, K. M., et al. 2022, *Phys. Rev. D*, **105**, 083524
- Watts, D. J., Galloway, M., Ihle, H. T., et al. 2023, *A&A*, **675**, A16 (BeyondPlanck SI)
- Zonca, A., Franceschet, C., Battaglia, P., et al. 2009, *J. Instrum.*, **4**, 2010