

Planck intermediate results

LVII. Joint Planck LFI and HFI data processing

Planck Collaboration: Y. Akrami^{11,49,50}, K. J. Andersen⁵⁰, M. Ashdown^{56,3}, C. Baccigalupi⁶⁵, M. Ballardini^{16,33}, A. J. Banday^{78,6}, R. B. Barreiro⁵², N. Bartolo^{20,53}, S. Basak⁷¹, K. Benabed^{48,73}, J.-P. Bernard^{78,6}, M. Bersanelli^{23,38}, P. Bielewicz^{64,65}, J. R. Bond⁵, J. Borrill^{9,76}, C. Burigana^{37,21,40}, R. C. Butler³³, E. Calabrese⁶⁸, B. Casaponsa⁵², H. C. Chiang^{18,4}, L. P. L. Colombo²³, C. Combet⁵⁸, B. P. Crill^{54,8}, F. Cuttaia³³, P. de Bernardis²², A. de Rosa³³, G. de Zotti³⁴, J. Delabrouille¹, E. Di Valentino⁵⁵, J. M. Diego⁵², O. Doré^{54,8}, M. Douspis⁴⁷, X. Dupac²⁶, H. K. Eriksen⁵⁰, R. Fernandez-Cobos⁵², F. Finelli^{33,40}, M. Frailis³⁵, A. A. Fraisse¹⁸, E. Franceschi³³, A. Frolov⁷², S. Galeotta³⁵, S. Galli^{48,73}, K. Ganga¹, M. Gerbino²⁹, T. Ghosh^{67,7}, J. González-Nuevo¹³, K. M. Górski^{54,79}, A. Gruppuso^{33,40}, J. E. Gudmundsson^{77,18}, W. Handley^{56,3}, G. Helou⁸, D. Herranz⁵², S. R. Hildebrandt^{54,8}, E. Hivon^{48,73}, Z. Huang⁶⁹, A. H. Jaffe⁴⁵, W. C. Jones¹⁸, E. Keihänen¹⁷, R. Keskitalo^{9,*}, K. Kiiveri^{17,31}, J. Kim⁶², T. S. Kisner⁶⁰, N. Krachmalnicoff⁶⁵, M. Kunz^{10,47,2}, H. Kurki-Suonio^{17,31}, A. Lasenby^{3,56}, M. Lattanzi^{41,21}, C. R. Lawrence⁵⁴, M. Le Jeune¹, F. Levrier⁷⁴, M. Liguori^{20,53}, P. B. Lilje⁵⁰, M. Lilley^{48,73}, V. Lindholm^{17,31}, M. López-Caniego²⁶, P. M. Lubin¹⁹, J. F. Macías-Pérez⁵⁸, D. Maino^{23,38,42}, N. Mandolesi^{33,21}, A. Marcos-Caballero⁵², M. Maris³⁵, P. G. Martin⁵, E. Martínez-González⁵², S. Matarrese^{20,53,28}, N. Mauri⁴⁰, J. D. McEwen⁶³, P. R. Meinhold¹⁹, A. Mennella^{23,38}, M. Migliaccio^{25,43}, S. Mitra^{21,33,41}, L. Montier^{78,6}, G. Morgante³³, A. Moss⁷⁰, P. Natoli^{21,75,41}, D. Paoletti^{33,40}, B. Partridge³⁰, G. Patanchon¹, D. Pearson⁵⁴, T. J. Pearson^{8,46}, F. Perrotta⁶⁵, F. Piacentini²², G. Polenta⁷⁵, J. P. Rachen¹⁴, M. Reinecke⁶², M. Remazeilles⁵⁵, A. Renzi⁵³, G. Rocha^{54,8}, C. Rosset¹, G. Roudier^{1,74,54}, J. A. Rubiño-Martín^{51,12}, B. Ruiz-Granados^{51,12}, L. Salvati^{32,36}, M. Savelainen^{17,31,61}, D. Scott¹⁵, C. Sirignano^{20,53}, G. Sirri⁴⁰, L. D. Spencer⁶⁸, A.-S. Suur-Uski^{17,31}, L. T. Svalheim⁵⁰, J. A. Tauber²⁷, D. Tavagnacco^{35,24}, M. Tenti³⁹, L. Terenzi³³, H. Thommesen⁵⁰, L. Toffolatti^{13,33}, M. Tomasi^{23,38}, M. Tristram⁵⁷, T. Trombetti^{37,41}, J. Valiviita^{17,31}, B. Van Tent⁵⁹, P. Vielva⁵², F. Villa³³, N. Vittorio²⁵, B. D. Wandelt^{48,73}, I. K. Wehus⁵⁰, A. Zacchei³⁵, and A. Zonca⁶⁶

(Affiliations can be found after the references)

Received 2 April 2020 / Accepted 1 July 2020

ABSTRACT

We present the NPIPE processing pipeline, which produces calibrated frequency maps in temperature and polarization from data from the *Planck* Low Frequency Instrument (LFI) and High Frequency Instrument (HFI) using high-performance computers. NPIPE represents a natural evolution of previous *Planck* analysis efforts, and combines some of the most powerful features of the separate LFI and HFI analysis pipelines. For example, following the LFI 2018 processing procedure, NPIPE uses foreground polarization priors during the calibration stage in order to break scanning-induced degeneracies. Similarly, NPIPE employs the HFI 2018 time-domain processing methodology to correct for bandpass mismatch at all frequencies. In addition, NPIPE introduces several improvements, including, but not limited to: inclusion of the 8% of data collected during repointing manoeuvres; smoothing of the LFI reference load data streams; in-flight estimation of detector polarization parameters; and construction of maximally independent detector-set split maps. For component-separation purposes, important improvements include: maps that retain the CMB Solar dipole, allowing for high-precision relative calibration in higher-level analyses; well-defined single-detector maps, allowing for robust CO extraction; and HFI temperature maps between 217 and 857 GHz that are binned into 0.9 pixels ($N_{\text{side}} = 4096$), ensuring that the full angular information in the data is represented in the maps even at the highest *Planck* resolutions. The net effect of these improvements is lower levels of noise and systematics in both frequency and component maps at essentially all angular scales, as well as notably improved internal consistency between the various frequency channels. Based on the NPIPE maps, we present the first estimate of the Solar dipole determined through component separation across all nine *Planck* frequencies. The amplitude is $(3366.6 \pm 2.7) \mu\text{K}$, consistent with, albeit slightly higher than, earlier estimates. From the large-scale polarization data, we derive an updated estimate of the optical depth of reionization of $\tau = 0.051 \pm 0.006$, which appears robust with respect to data and sky cuts. There are 600 complete signal, noise and systematics simulations of the full-frequency and detector-set maps. As a *Planck* first, these simulations include full time-domain processing of the beam-convolved CMB anisotropies. The release of NPIPE maps and simulations is accompanied with a complete suite of raw and processed time-ordered data and the software, scripts, auxiliary data, and parameter files needed to improve further on the analysis and to run matching simulations.

Key words. cosmic background radiation – cosmology: observations – cosmological parameters – Galaxy: general – methods: data analysis

1. Introduction

This paper, the last by the Planck Collaboration, describes a *Planck*¹ data-processing pipeline called NPIPE and the calibrated data, maps, simulations, and other data products that

* Corresponding author: R. Keskitalo,
e-mail: rtkeskitalo@lbl.gov

¹ *Planck* (<http://www.esa.int/Planck>) is a project of the European Space Agency (ESA) with instruments provided by two scientific consortia funded by ESA member states and led by Principal Investigators from France and Italy, telescope reflectors provided through a collaboration between ESA and a scientific consortium led and funded by Denmark, and additional contributions from NASA (USA).

it produces. NPIPE represents a natural evolution of the previous *Planck* analysis efforts made within the *Planck* data processing centres (DPCs) in Paris and Trieste, but is uniquely designed to analyse data from both the Low Frequency Instrument (LFI) and the High Frequency Instrument (HFI) within the same framework. Furthermore, NPIPE is implemented to execute efficiently on massively parallel high-performance computing (HPC) systems. Indeed, it was both developed and run on the HPC facilities hosted by the National Energy Research Scientific Computing Center (NERSC), motivating the “N” in NPIPE. The main technical design criteria for NPIPE are therefore:

- the ability to handle both LFI and HFI data;
- efficient execution on HPC systems;
- adaptability to evolving HPC architectures;
- *Planck* data access based on Exchange File Format files rather than databases;
- minimal I/O during processing.

Running in a massively parallel HPC environment, NPIPE overcomes some of the practical difficulties and limitations imposed on the pipelines operated by the *Planck* DPCs. With the above design requirements in mind, NPIPE developed incrementally, with the first step being a prototyping platform for testing new implementations of existing HFI preprocessing modules. Early success in 4 K line removal and inclusion of the data taken during repointing manoeuvres (excluded in the DPC processing) pre-saged the development of a full data-processing pipeline from raw time-ordered data (TOD) all the way to maps. Since most of the DPC pipeline modules were deeply integrated into the database-driven architecture at the DPC, this led to fresh implementations of a majority of the processing modules. The capability to handle LFI data was added while studying reduction of the high-frequency noise in the LFI TOD through improved decorrelation of the $1/f$ noise.

Having a single data-processing pipeline that could handle both LFI and HFI data led to significant improvements in the calibration procedures, which combine the strengths of both LFI and HFI calibration. Specifically, the NPIPE calibration procedure is a synthesis of the high S/N approach that works well for the HFI 353 GHz channel and the noise-limited approach that is critical to LFI 70 GHz calibration. Numerous smaller improvements to the data-processing modules were motivated by experiences with the nine *Planck* frequencies and their diverse features. The LFI- and HFI-DPC pipelines evolved in the post-launch period, often in response to instrument-specific effects that emerged as the calibration accuracy improved. NPIPE builds on knowledge accumulated through the years, and takes a global approach, aiming at coherent treatment of the full *Planck* data set.

In what follows, we detail NPIPE processing and a full data set of *Planck* maps that result from it. These NPIPE data maps are supported by a suite of simulated maps that capture the relevant noise and systematic errors present in the data. We also compare the NPIPE results to the two previous public releases of *Planck* temperature and polarization maps, namely the second data release in 2015 (“PR2”; [Planck Collaboration I 2016](#)) and the third release in 2018 (“PR3”; [Planck Collaboration I 2020](#)). Comparison with both earlier releases allows the reader to assess the magnitude of the differences we find.

To characterize the new maps, we consider a few important applications for which component separation is essential. Indeed, robustly supporting component-separation applications ranks among the primary motivations for the NPIPE products, and several features have been added specifically to meet the requirements of future astrophysical analysis. Among these are: maps that retain the cosmic microwave background (CMB) Solar dipole, and thereby allow for joint component separation and high-precision relative calibration; robust single-detector temperature-only sky maps that allow for fine-grained CO extraction, and thereby weaker degeneracies with respect to CMB, thermal dust, and free-free emission; and HFI maps between 217 and 857 GHz pixelized at a HEALPIX² ([Górski et al.](#)

2005) resolution of $N_{\text{side}} = 4096$, corresponding to a pixel size of 0.88, which ensures properly signal-bandwidth-limited (not limited by pixel size) sampling of the highest resolution *Planck* beams of 5’.

Taking advantage of the new NPIPE maps, we present the first *Planck*-only astrophysical sky model that simultaneously constrains the CMB Solar dipole, higher-order CMB temperature fluctuations, low-frequency foregrounds, thermal dust, and CO line emission. This joint analysis has at least three major advantages compared to previous analyses. First, it results in a global LFI+HFI CMB Solar dipole estimate with minimal foreground contamination. Second, it eliminates the need for estimating dipole residuals per frequency band during component separation, and thereby reduces large-scale degeneracies in the astrophysical component maps. Third, it allows for high-precision relative, inter-frequency gain calibration during component separation. The resulting astrophysical components are compared to corresponding previous releases, and are found to be consistent with earlier maps, but exhibiting lower noise and systematics. Similarly, the derived CMB dipole parameters are also consistent with previous estimates, although the new result has a slightly higher amplitude. Properly taking into account uncertainties from higher-order CMB fluctuations (Thommesen et al., in prep.), the NPIPE results are consistent with the *Planck* 2018 constraints at the $1-2\sigma$ level.

Finally, we constrain the optical depth of reionization τ with the large-scale polarization NPIPE maps. Due to the low levels of systematic residuals in these maps, we observe good consistency among results derived from different frequency maps, and the resulting large-scale angular CMB power spectrum and cosmological parameters appear robust with respect to both data and sky cuts.

The release of NPIPE maps and simulations is accompanied with the full software suite and input data so that the results can be reproduced, and even improved, given sufficient computational resources. The rest of the paper is organized as follows. In Sect. 2 we summarize the low-level NPIPE data processing, highlighting in particular where NPIPE differs from the DPC processing. In Sect. 3 we comment on the destriping mapmaking algorithm employed by NPIPE. In Sect. 4 we present and characterize the NPIPE frequency and detector maps, and in Sect. 5 we describe the associated end-to-end simulations. In Sect. 6 we compare NPIPE results with those of previous data releases. In Sect. 7 we discuss the astrophysical component maps derived from the NPIPE-only data set, before considering the CMB Solar dipole in Sect. 8 and the optical depth of reionization in Sect. 9. Conclusions are presented in Sect. 10, and algorithmic details are provided in the appendices.

2. Data processing

NPIPE re-implements most of the Level 2 data processing performed at the *Planck* DPCs ([Planck Collaboration II 2019](#); [Planck Collaboration III 2020](#)), while introducing a number of detailed changes that will be described later in this section. We divide our processing into two main steps.

Local *preprocessing* covers all the steps that can take place at the single-detector³ level without projecting the signal onto

² <https://healpix.sourceforge.io>. The square roots of the uniform pixel areas at $N_{\text{side}} = 1024, 2048, \text{ and } 4096$ are 3.4, 1.7, and 0.88, respectively.

³ In this paper, a “detector” is understood to be an LFI radiometer or an HFI bolometer. Each LFI horn and each polarized HFI horn feeds two linearly polarized detectors aligned in orthogonal directions. Note that each LFI radiometer actually comprises two detector diodes that are co-added to optimize rejection of coherent fluctuations.

maps. Furthermore, every preprocessing step (except for the mission averaging of the LFI low-pass filter and spike templates) can be performed at the single-pointing-period⁴ level (see Sect. 2.3.9) without reference to data beyond it.

Global *reprocessing* encompasses single- and multi-detector operations that require much or all of the data to be in memory. Examples of systematics handled during reprocessing are gain fluctuations and residuals from the low-frequency bolometer transfer function.

The data are written to disk after preprocessing and reprocessing to flexibly allow development of these two major steps. Furthermore, mapmaking of the reprocessed TOD can be separately optimized after the data are processed.

The parts of NPIPE data processing that continue to employ DPC-derived products are the following.

- LFI analogue-to-digital converter (ADC) correction. We apply the DPC-provided correction profiles.
- HFI Initial ADC correction measurement. We start from the DPC-provided correction profiles.
- HFI bolometer transfer-function measurement. We deconvolve the DPC-measured transfer function.
- Beam and focal plane geometry measurements. Such changes can be implemented in NPIPE, but the data set described and released here retains the DPC values. Neither are expected to differ from PR3, since we have not updated the bolometer transfer functions.

2.1. Differences between NPIPE and the PR3 data processing

We list here the most important differences between NPIPE and PR3 data processing, with section numbers for further details.

- We use a hybrid calibration scheme to work around a degeneracy in the *Planck* calibration process. The 30 and 353 GHz data are first calibrated along with the polarized sky much like PR3 HFI is calibrated. We then mimic the PR3 LFI calibration and calibrate the intermediate CMB frequencies (44–217 GHz), using an approximation that the sky polarization can be captured with templates derived from the foreground frequencies (Sect. 2.4.13). The calibration for the CMB frequencies is thus based on temperature-only sky and polarized foreground templates (see Sect. 2.4.13). The use of this polarization prior significantly reduces the uncertainty in gain and the large-scale polarization systematics, but also requires a measurable and correctable transfer function affecting large-scale CMB polarization (described in Sect. 4.3).
- We include the *Planck* repointing manoeuvre data in all of our processing, leading to a roughly 8% increase in total integration time (2.3.1).
- We retain the Solar dipole in the frequency maps for component separation.
- We include in our pointing solution the latest star-camera distortion corrections that were developed for *Herschel* data processing (Sect. 2.2.2).
- We correct all LFI radiometers for 1 Hz housekeeping spikes, rather than just the 44 GHz ones (Sect. 2.3.5).

⁴ A pointing period or ring consists of the time to orient the spacecraft spin axis with three precise thruster burns (the *repointing manoeuvre*) and the continuous science scan that follows the repointing. Each repointing manoeuvre lasted about four minutes, with the majority of time spent passively slewing the spacecraft spin axis into the new orientation. The spin axis was adjusted by 2' during a standard repointing. The subsequent science scan lasted between 35 and 75 min.

- We fit for more 4 K cooler lines, and use the global signal estimate (Sect. 2.3.2) for signal removal (Sect. 2.3.5).

- We modify the LFI sky-load differencing to include a low-pass filter that reduces the injection of uncorrelated noise *into* the differenced signal (Sect. 2.3.6).

- We use our own signal estimate during HFI glitch detection and removal (Sect. 2.3.4).

- We extend the HFI glitch flags during transfer function deconvolution, to avoid leaking constrained-realization power from the gaps into the surrounding data. This leads to less small-scale noise and lower noise correlations between half-rings⁵ (Sect. 2.3.9).

- We fit for more HFI transfer-function harmonics, to better address the odd-even survey differences (Sect. 2.4.5).

- We subtract only the seasonally varying part of zodiacal emission, in order not to bias the dust component in the maps (Sect. 2.4.6).

- We provide polarization-corrected, single-detector, and single-horn temperature maps at all frequencies (Sects. 4.1.3 and 4.1.4).

- We correct the HFI polarization angles and efficiencies, to address polarized signal-like residuals in single-detector maps (Sect. 2.4.15).

- We bin our single-detector maps from 217 to 857 GHz at $N_{\text{side}} = 4096$, to fully sample the narrow beam.

- We make detector-set maps for cross correlation analysis and systematics-level studies that are fully independently processed (Sect. 4.1.5).

- We provide a consistent, low-resolution data set, with estimates of pixel-pixel noise covariance across all *Planck* frequencies (Sect. 4.2).

- We fit for signal distortion to address second-order ADC nonlinearity in the HFI CMB channels from 100 to 217 GHz (Sect. 2.4.2).

- We extend to LFI data the HFI approach of correcting for bandpass mismatch in the time domain (Sect. 2.4.7).

- We destripe all data with Madam, using extremely short 167 ms baselines (Sect. 3).

- We calibrate with apodized masks that allow us to access most of the sky while down-weighting Galactic regions where even small issues in the signal model can cause problems with calibration (Sect. 2.4.12).

2.2. Inputs

The inputs to NPIPE are the raw, digitized data as they were transmitted from the spacecraft and decompressed by the DPC Level 1 processing (Planck Collaboration II 2014; Planck Collaboration VI 2014). The HFI TOD are initially corrected with the same ADC nonlinearity profiles as PR3 data. The data are stored in *Planck* exchange file format FITS files that we have extended to accommodate raw, undifferenced LFI timelines.

Each LFI horn feeds two polarization-orthogonal radiometers (“detectors” in this paper, see footnote 3). Each radiometer comprises two diodes that alternate between observing the sky and a 4 K reference load. In total, each LFI horn provides four sky timelines and four reference timelines. From the 11 LFI horns there are then 88 discrete timelines that are preprocessed into 22 detector timelines.

For each polarized HFI horn there are two polarization-sensitive bolometers (PSBs labelled “a” and “b”) at orthogonal

⁵ Half-ring data sets are built from the first and second halves of the pointing periods.

angles of polarization sensitivity. Some of the HFI horns instead house nearly unpolarized “spiderweb” bolometers (SWBs). There are 36 HFI horns in total, 16 polarized horns and 20 nearly unpolarized. In addition, two dark bolometers measured the thermal baseline on the focal plane. With two of the SWBs irrecoverably compromised by random telegraphic signal, we have 52 bolometer timelines that are preprocessed into 50 optical timelines.

2.2.1. Time span

HFI science operations spanned the period between 12 August 2009 and 15 January 2012, comprising almost five sky surveys⁶, 885 operational days (29 months), and 27 000 pointing periods. We follow [Planck Collaboration III \(2020\)](#) and exclude the last 22 days (956 pointing periods) of the HFI science data due to large drifts in the thermal baseline. This leaves the fifth HFI sky survey only 80% complete. The drifts were considered harmful both due to the correlated noise fluctuations they induce and due to the challenges the changes in the baseline introduce to ADC nonlinearity correction.

LFI science operations continued until 10 October 2013, comprising 8.3 sky surveys, 1 513 operational days (49.6 months), and 46 000 pointing periods. The LFI maps in PR3 are based on the eight completed sky surveys ([Planck Collaboration II 2019](#)). NPIPE maps include the additional 60 days of the ninth sky survey.

2.2.2. Pointing

The boresight pointing in NPIPE is based on the same raw attitude information as PR3. We processed the raw attitude quaternions to include a star-camera field-of-view distortion correction that was developed for *Herschel* and updates some guide-star positions necessitated by the delay in the *Planck* launch ([Tuttlebee 2013](#)). The corrections are only a few arc seconds at most, but are systematic rather than statistical.

Reprocessing the attitude history allowed us to extend the low-pass filtering used to suppress high-frequency noise in the reconstructed pointing. In PR3, the filtering only covered the stable science scans, leaving the repointing manoeuvres considerably noisier. In NPIPE the same low-pass filter is used on both science scans and repointing manoeuvres, better supporting our use of the repointing manoeuvre data in mapmaking.

NPIPE uses the same PTCOR correction ([Planck Collaboration I 2016](#)) as PR3 for the thermal deformation of the angle between the boresight and the star camera. For convenience, we apply the correction directly to the attitude history files (AHFs), thus removing the need to apply the correction every time the pointing is read from the AHFs.

2.3. Local preprocessing

We present flow charts of the preprocessing steps for LFI in Fig. 1, dark HFI bolometers in Fig. 2, and optical HFI bolometers in Fig. 3. The flow charts are annotated with the appropriate section numbers.

⁶ The scan strategy was such that *Planck* observed about 95% of the sky over the course of six months, but the entire sky exactly twice every 12 months. The approximately six-month periods are called “surveys”. See Sect. 4 of [Planck Collaboration I \(2014\)](#), and particularly Table 1, for precise definitions.

2.3.1. Repointing manoeuvres

NPIPE was designed to retain the data from the times between fixed pointings of the spin axis that were previously neglected. While these repointing manoeuvres are marked by changes in the rotation of the spacecraft, these changes alter the angular momentum as little as possible to conserve propellant. This means that the scanning rate on the sky does not change and the effective beam is unaffected. The DPC pipelines omitted these data for a number of reasons. Firstly, they lack the repetitive overlapping scans that allow effective compression of the data onto ring maps. Secondly, the pointing solution is less accurate. Lastly, there was some concern that the thermal environment onboard the spacecraft was less stable during and immediately after thruster burns. With the benefit of hindsight, it is now clear that none of these concerns is very serious.

NPIPE deals with the lack of repetitive scans by using a global signal estimate (Sect. 2.3.2), and only using ring-compressed data for calibration and other template fitting. Our pointing solution for the repointings is more stable than the earlier versions, due to our use of a low-pass filter to suppress small-timescale fluctuations (Sect. 2.2.2). The concern for thermal fluctuations is greatly alleviated by our application of horn-symmetric detector weights (Sect. 2.4.11), leading to effective cancellation of thermal common modes within polarized horns. Tests show that the repointing-period data are perfectly usable (Sect. 4.4).

2.3.2. Global signal estimate

Many of the preprocessing modules require an estimate of the detected sky signal. In particular, the HFI-DPC modules make extensive use of a phase-binned⁷ signal estimate, where data from a single pointing period and a single detector are used to estimate the periodic signal present in the data. This binning is made possible by the repetitive *Planck* scanning strategy, which repeats the same scanning circle 39–65 times at a fixed rate of one revolution per minute. Binning the data according to the spacecraft spin phase produces an unbiased estimate of the signal, but the phase-binned data contain a significant amount of noise that is made scan-synchronous by the binning. Moreover, the phase-binned estimate is not applicable to the 4 min repointing manoeuvres at the beginnings of each science scan.

Except in glitch removal, NPIPE uses an alternative that we call the “global signal estimate”. We sample a full-mission, single-detector temperature map from a previous iteration of NPIPE, together with polarization from the full set of detectors within each frequency channel. Using the last run of the pipeline as input to preprocessing is justified by the fact that the relevant, high S/N modes of the maps converged early in the development process⁸. This map already contains the Solar dipole. We further add an estimate of the orbital dipole based on the known spacecraft velocity. The time-dependent orbital dipole is not included in the map. For HFI, this global estimate of the sky signal is

⁷ Each pointing period consists of tens of repeating scanning circles. Individual detector data for one period form a one-dimensional data set that can be indexed and binned according to the spin phase of the spacecraft.

⁸ In fact, the large-scale temperature signal has remained essentially unchanged since Public Release 1 in 2013 ([Planck Collaboration I 2014](#)) except for small overall calibration adjustment. One telling measure of the robustness of the temperature map is the small component separation residual discussed in Sect. 7.

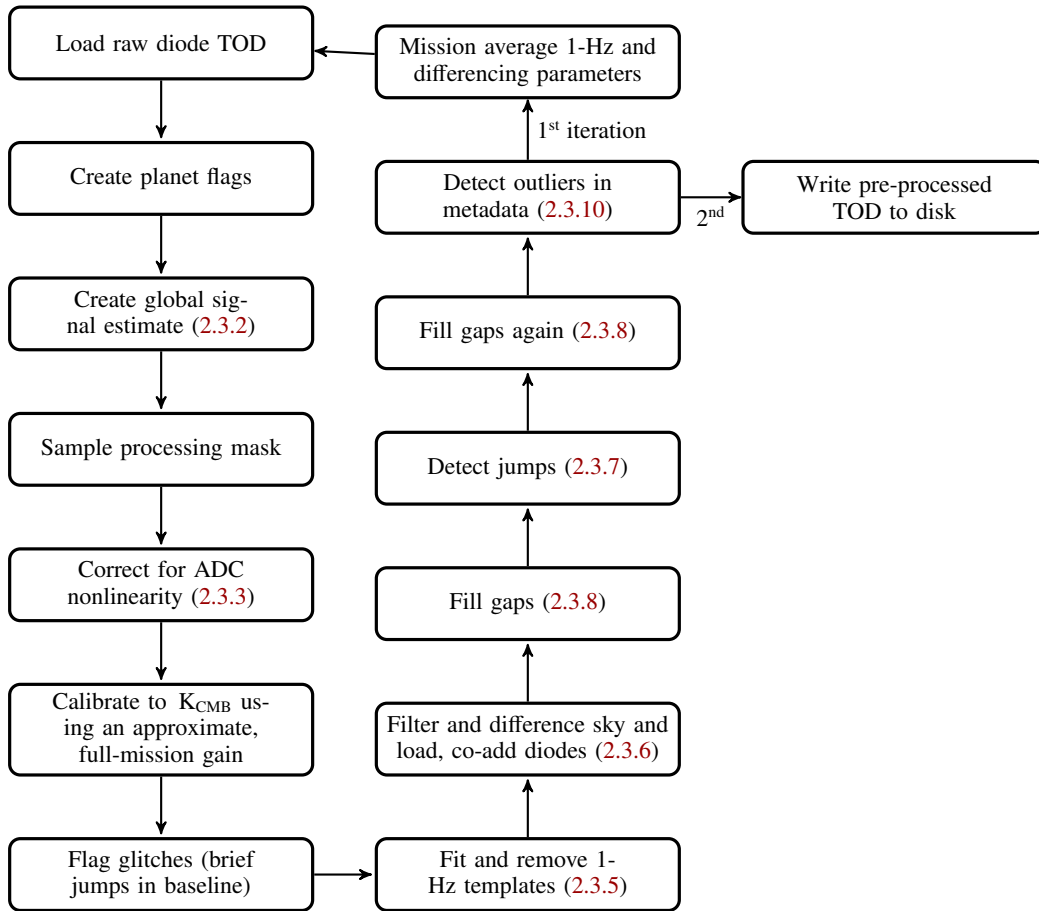


Fig. 1. NPIPE preprocessing flow chart for LFI, indicating the sections in which major steps are discussed. The main loop runs twice, first fitting the 1 Hz spike templates and low-pass filter parameters independently on each pointing period, then mission-averaging the template amplitudes and filter parameters and running the entire processing pipeline again with these fixed values.

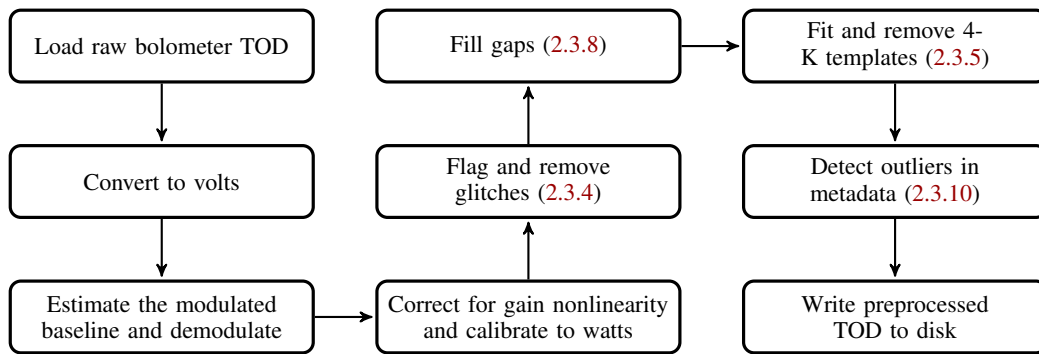


Fig. 2. NPIPE preprocessing flow chart for dark HFI bolometers.

then convolved with the estimated bolometer transfer function in order to produce the required estimate of the sky signal.

2.3.3. ADC nonlinearity correction

Both *Planck* instruments experience detectable levels of nonlinearity in their analogue-to-digital conversion (ADC). The level of nonlinearity in LFI data is much less than in HFI, and is straightforward to correct. NPIPE applies the LFI DPC correction profiles as part of our preprocessing. Details of the LFI nonlinearity and the measured correction can be found in [Planck Collaboration II \(2019\)](#).

For HFI nonlinearity in the digitization was found to be the most important systematic error affecting large-scale polarization (see discussion in [Planck Collaboration III 2020](#); [Delouis et al. 2019](#)). The first-order manifestation of the nonlinearity is apparent gain changes that correlate with changes in the level of signal in the detector. Secondly we observe distortions of the signal itself.

After the end of HFI observations, a two-year campaign was carried out during the LFI extended mission to gather statistical information about the ADC effects. The problem is not fully tractable, since the digitized data transmitted to Earth are down-sampled onboard the spacecraft. Nevertheless, the derived ADC

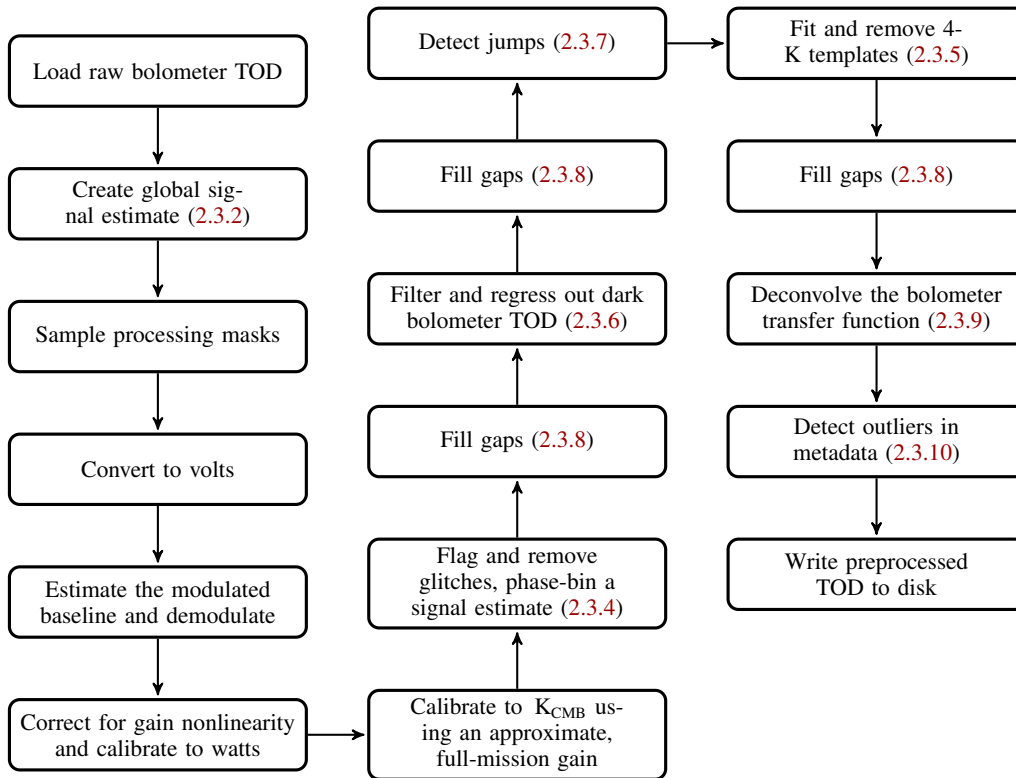


Fig. 3. NPIPE preprocessing flow chart for optical HFI bolometers.

correction profiles greatly improve the effective gain stability in the data, as well as consistency between the odd- and even-parity data sets split from the square-wave-modulated data (Planck Collaboration VII 2016). Even after applying these ADC correction profiles, however, residual ADC nonlinearity (ADCNL) remains the primary systematic error in the large-scale HFI polarization data (Planck Collaboration Int. XLVI 2016). NPIPE begins with the same initial ADCNL-corrected TOD as PR3, but uses a different model to measure and remove the residual ADCNL (see Sect. 2.4.2).

2.3.4. Glitch removal

Cosmic-ray glitches in the HFI data (see Planck Collaboration X 2014) are detected, fitted, and removed in NPIPE using the same despiking software that runs in the HFI-DPC pipeline (Planck HFI Core Team (2011), Planck Collaboration VI (2014)). However, NPIPE does its own signal removal and runs despiking in the “dark” mode developed for processing the dark bolometer data.

For the signal estimate, we reorder the signal by the satellite spin phase and fit a sixth-order polynomial in a sliding window of three bin widths at a time, keeping the polynomial fit in the centre bin. Each bin is as wide as the average full width at half maximum (FWHM) of the beam at that frequency. The DPC signal estimate uses 1.5 bins regardless of the optical beam width.

We also experimented with using the global signal estimate in place of the phase-binned signal estimate. This has the advantage that the signal estimate does not include noise from the pointing period that is being analysed. However, since the signal estimate and the resulting glitch residuals touch virtually every datum in HFI, it was safer to use the local signal estimate instead of potentially biasing the entire data set with errors in the global signal estimate.

Using a phase-binned signal estimate introduces a noise bias in the glitch fits as the fitted glitch templates attempt to accommodate some of the noise in the signal estimate. This bias was evident in the noise estimates derived from half-ring difference maps, which were about a percent short of the full-data noise due to correlated noise cancellation (Planck Collaboration XII 2016). Our FWHM-wide binning scheme and the extra smoothing from the polynomial fits limits the amount of noise in the signal estimate and the amount of half-ring bias. Nevertheless, users of NPIPE products are advised to be cognizant of the fact that the small-scale noise in the half-ring maps is correlated at a low level due to correlated errors in the glitch residuals.

We have raised the glitch detection threshold to reflect improvements in the signal estimate and to reduce the rate of false positives. Our short-baseline destriping (see Sect. 3) is also better equipped to remove residual glitch tails than the ring off-set model applied in DPC processing. The net effect of all the changes in glitch detection is to reduce the overall glitch detection rates to an average of 8% in each detector, about half that of the DPC pipeline. While this may seem like a large difference, the energy distribution of the glitches is such that the affected glitch population is mostly right at the detection limit. In Sect. 6 we show that the resulting NPIPE maps have lower noise than their 2018 counterparts, despite allowing more glitches through, and even after correcting for the added integration time.

2.3.5. Removal of frequency spikes

The drive electronics for the 4 K cooler cause interference in the HFI data at many harmonics. We follow the DPC approach in fitting time-domain sine and cosine templates at the harmonic frequencies. The DPC pipeline fitted for nine harmonics, namely, 10, 16.7, 20, 30, 40, 50, 60, 70, and 80 Hz. We have




Fig. 4. Interference from the drive electronics of the 4 K cooler in the HFI detectors. *Left:* measured 4 K line amplitudes for the most intense lines in bolometer 143-3a. *Right:* new 16.03 Hz line template amplitudes for bolometer 100-1b. For unknown reasons, the line is only detected intermittently around day 300. For reference, the CMB temperature fluctuations typically fit within $\pm 250 \mu\text{K}$ and polarization fluctuations within $\pm 2 \mu\text{K}$; it is therefore clear that the cooler lines must be modelled and removed to high precision. With the *Planck* scan rate being 6° s^{-1} , the 10 Hz line corresponds to an angular scale of $36'$ and multipole $\ell \approx 600$. For higher-order harmonics of 10 Hz, divide the angular scale and multiply the multipole with the appropriate factor.

added a further six harmonics that are transiently detected in some of the bolometers, namely 16.0, 25.0, 43.4, 46.2, 47.6, and 56.8 Hz.

The line phases and amplitudes in the data are observed to change over timescales of minutes, although the source of the variation is not understood. Fitting for the lines at the pointing period level leaves detectable residuals. We have designated six harmonics as “intense”, meaning that their S/N allows for shorter fitting period (5 min). These intense lines are (in approximate order of intensity) at 70, 30, 10, 50, 16.7, and 20 Hz. The 5 min fitting period is a compromise between leaving line residuals and applying a notch filter to the instrumental noise. Figure 4 shows the measured line amplitude for some of the lines in a particularly badly-affected bolometer, 143-3a.

Our pipeline uses the global signal estimate rather than a phase-binned signal estimate for spike removal. The use of this approach avoids a resonant ring issue where the spin rate in some of the pointing periods causes the harmonic templates to be synchronous with the sky signal.

LFI data show frequency spikes at 1 Hz that are understood to originate in the housekeeping electronics. To deal with these, we employ the same methods as with HFI, constructing time-domain sine and cosine templates at 1 Hz and all of its harmonics (and their aliased counterparts) up to the Nyquist frequency. With the LFI sampling frequencies, this amounts to 47, 68, and 116 individual fitting frequencies at 30, 44, and 70 GHz respectively. These template amplitudes were found to be stable (no model exists to predict their amplitude), so we averaged their amplitudes across the mission, then subtracted them from the TOD. While the LFI-DPC processing only corrected for the 1 Hz spikes in the most-affected 44 GHz data, NPIPE fits and corrects for the electronic interference in all of the LFI detectors. Expanding the correction to all of the detectors is particularly important because the low-pass-filtered load signal does not cancel the interference as efficiently as an unfiltered signal. We show co-added line templates for all LFI diodes in Fig. 5.

The LFI-DPC approach to building the 1 Hz spike template differs from that of NPIPE. The DPC template was built by

binning diode data into a 1 s-long template according to the phase within the 1 s period. NPIPE instead considers harmonic templates⁹ to model the parasitic oscillation. In the limit of excellent S/N, the two approaches are equivalent, since the DPC template can be binned arbitrarily finely. For shorter fitting periods and lower S/N, the harmonic templates are a more economical decomposition of the parasitic effect, and allow for a more robust fit.

2.3.6. Sky-load differencing and thermal decorrelation

Both LFI and HFI perform decorrelation to reduce time-correlated ($1/f$) noise fluctuations in the data. The LFI $1/f$ templates are the reference-load timestreams, one for each diode. The HFI thermal templates come from the two dark bolometers.

Traditionally the LFI sky-load differencing has relied on a scaling factor, R :

$$\text{difference} = \text{sky} - R \times \text{load}. \quad (1)$$

For PR3, the LFI DPC calculated the scaling factor for each operational day. These least squares estimates of R balance two opposing effects:

1. low-frequency $1/f$ fluctuations in the load timestreams are highly correlated with the fluctuations in the sky timestream;
2. high-frequency noise in the load timestream is essentially uncorrelated with the sky timestream.

A value of R calculated in this way suppresses the $1/f$ noise, while actually slightly *increasing* the high-frequency noise in the differenced signal.

In NPIPE we have attempted to suppress the injection of uncorrelated high frequency noise into the differenced signal. In

⁹ “Harmonic” in this context means sine and cosine templates needed to represent 1 s periodic signal, accounting for potential aliasing from downsampling the signal. The frequencies of the templates are of the form $nf + kf_r$, where n is a non-negative integer, f is the parasitic frequency (1 Hz), $k \in \{-1, 0, 1\}$ and f_r is the floating-point modulo of sampling frequency, f_s , and f .




Fig. 5. LFI spike templates built from the harmonic templates fitted to each of the 44 timelines. Each of the coloured curves represents a different diode. The CMB temperature fluctuations have a maximum amplitude of about $250 \mu\text{K}_{\text{CMB}}$ and the polarization fluctuations about $2 \mu\text{K}_{\text{CMB}}$. The 1 Hz features would be seen at approximately $\ell = 60$.

a somewhat simplified model, the sky and load each have an independent white noise component, σ_{sky} and σ_{load} , and share a fully correlated $1/f$ component with power $P_{\text{corr}}(f)$. The corresponding power spectral densities (PSD) are

$$P_{\text{sky}}(f) = \sigma_{\text{sky}}^2 + R_0^2 P_{\text{corr}}(f), \quad (2)$$

$$P_{\text{load}}(f) = \sigma_{\text{load}}^2 + P_{\text{corr}}(f), \quad (3)$$

$$P_{\text{corr}}(f) = \sigma^2 \left[1 + \left(\frac{f}{f_{\text{knee}}} \right)^\alpha \right], \quad (4)$$

where R_0 is a scaling factor that eliminates the correlated component in the sky–load difference in Eq. (1). Note that we are adopting a notation where σ^2 has PSD units. Using an arbitrary scaling factor, the PSD of the differenced timestream gives

$$P_{\text{diff}}(f) = \sigma_{\text{sky}}^2 + R^2 \sigma_{\text{load}}^2 + (R_0 - R)^2 P_{\text{corr}}(f). \quad (5)$$

We can thus derive a frequency-dependent scaling factor, $R(f)$, that minimizes the noise power in the differenced signal:

$$R(f) = R_0 \frac{P_{\text{corr}}(f)}{P_{\text{corr}}(f) + \sigma_{\text{load}}^2}. \quad (6)$$

In this simplified model, Eq. (6) defines an optimal low-pass filter that minimizes noise power in the differenced signal. We show examples of the sky and load PSDs and the differenced PSD in Fig. 6, and compare the low-pass-filtered approach to a direct scaling of the load signal. The two examples show that when the uncorrelated instrumental noise is comparable to the correlated $1/f$ noise, low-pass filtering can substantially improve both the low and high frequency noise in the differenced signal.

As a refinement to the model, we consider that a radiometer comprises two diodes, each with sky and load, and the two diodes are also correlated with each other, complicating the situation. We use Eq. (6) to inform us of a parametric representation of a low-pass filter with only three free parameters:

$$F^{\text{lp}}(f; R, \sigma, \alpha) = R^2 \frac{f^\alpha}{f^\alpha + \sigma^2} \quad (7)$$

and we use nonlinear minimization to optimize the low-pass filter for both diodes and their relative weights that define the co-added radiometer signal, \mathbf{d} :

$$\mathbf{d} = w \cdot (s_0 - F^{\text{lp}}(I_0; R_0, \sigma_0, \alpha_0)) + (1 - w) \cdot (s_1 - F^{\text{lp}}(I_1; R_1, \sigma_1, \alpha_1)). \quad (8)$$

Here “ s ” and “ I ” represent the sky and the load, respectively, and indices “0” and “1” identify the two diodes in the radiometer. The fit is performed while marginalizing over a sky-signal estimate.

The best-fit parameters ($w, R_0, \sigma_0, \alpha_0, R_1, \sigma_1, \alpha_1$) display a fair amount of scatter from pointing period to pointing period. However, their sample medians are very stable, suggesting that the scatter is driven by noise. We decided to run the main LFI preprocessing loop twice. The first iteration is performed to measure the best-fit, low-pass filters (Eq. (7)), as well as to construct 1 Hz spike templates (Sect. 2.3.5) for every pointing period. Then the parameters are averaged across the entire mission and the processing is repeated, keeping the filter and spike parameters fixed.

2.3.7. Jump correction

The *Planck* data also contain “jumps”, which are sharp changes in the baseline level that persist (unlike spikes). The NPIPE jump-correction module uses a matched filter to detect baseline changes in the TOD. We set the width of the filter to approximately 4 min of data. The width was chosen as a compromise between S/N (wider is better) and $1/f$ noise fluctuations (narrower is better).

We use the global signal estimate for signal removal, avoiding obvious issues that arise from building a signal estimate from data that contain jumps we are trying to detect. Once a jump is detected, the filtered signal provides an estimate of the jump size, which is then corrected for. Since both the exact position and size of the jump are subject to uncertainty, we flag one filter width (4 min) of data around the detected and corrected jump. This approach differs from the implementation used in the HFI-DPC pipeline, which identified pointing periods with jumps by examining the cumulative sum of samples.

The total number of detected jumps varies by detector, but on the average we identify, correct, and flag about a thousand jumps in each detector, once pointing periods with outlier statistics (see Sect. 2.3.10) have been discarded.

2.3.8. Gap filling

As discussed in Sects. 2.3.4 and 2.3.7, the two instruments are subject to glitches and jumps of different origin, and both

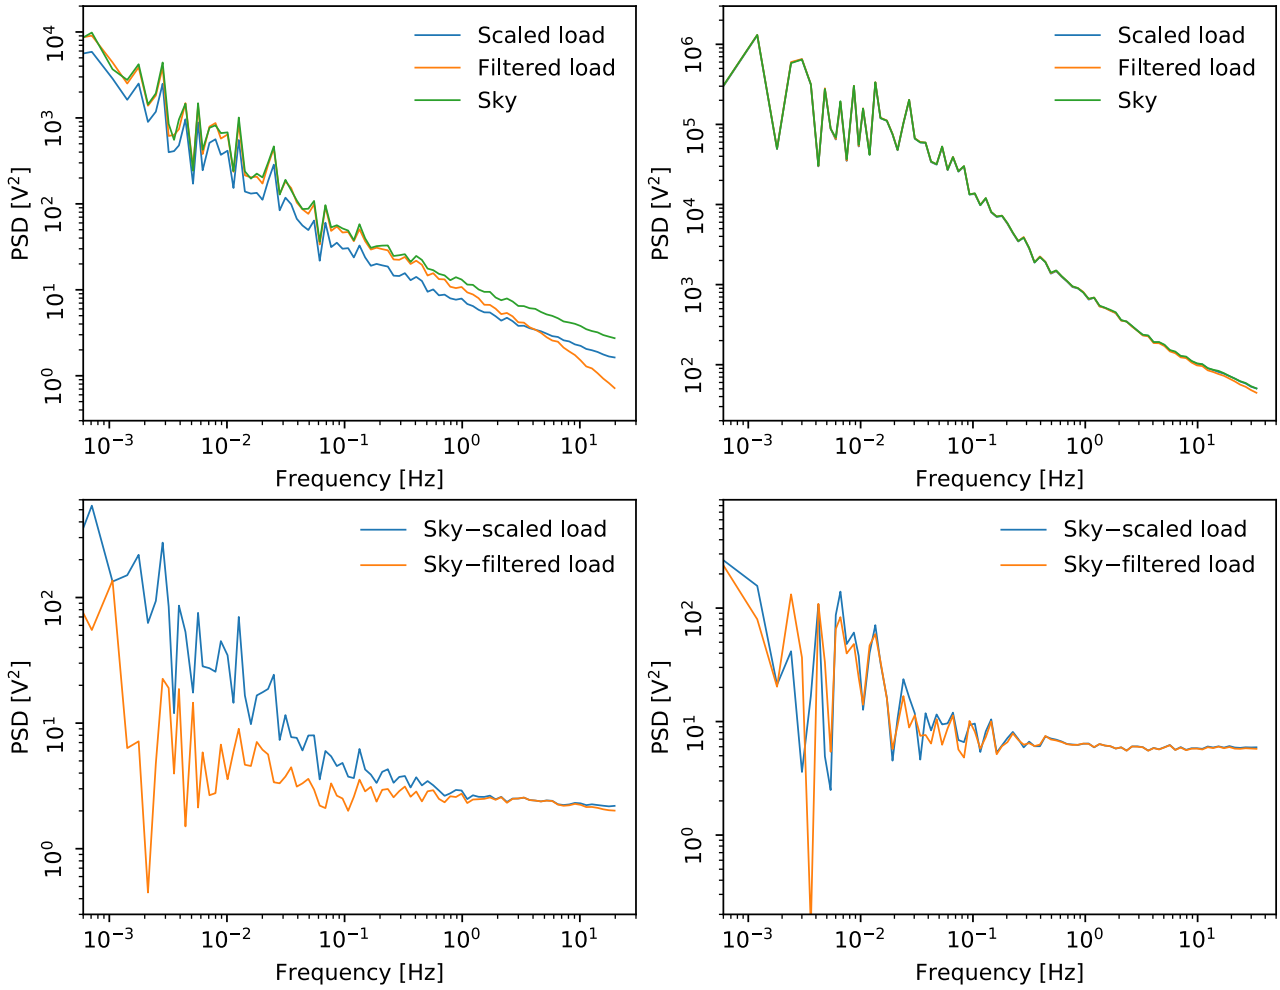


Fig. 6. NPIPE applies a low-pass filter to the LFI load signal before decorrelation. In PR3, a single scaling coefficient was fitted each operational day. We demonstrate the two approaches for one pointing period on two diodes: LFI2500 (*left*), with a relatively low amount of $1/f$ noise; and LFI1800 (*right*), which is dominated by $1/f$ at all frequencies. *Top*: power spectral densities of sky and best-fit load signals. *Bottom*: power spectral densities of sky-load differenced signals. The PSDs were binned into 300 logarithmically-spaced bins for clarity. The scaled load template leaves more noise at both high and low frequencies than a filtered version, where the uncorrelated power is suppressed. When the $1/f$ fluctuations dominate across the spectral domain, the scaling and filtering approaches achieve effectively the same result.

require removal of sections of data. To facilitate the use of Fourier techniques (e.g., bolometer transfer function deconvolution, Sect. 2.3.9), the gaps need to be filled with a constrained realization.

NPIPE deals with gap filling using the same basic approach for all *Planck* channels. For LFI, since it has relatively few gaps, the signal part of the constrained realization comes from the global signal estimate (Sect. 2.3.2), while for HFI it comes from the phase-binned signal estimate. The decision to use the (noisy) phase-binned signal estimate instead of the global signal estimate was made out of an abundance of caution: even small but systematic errors in the signal estimate used to fill 10–20% of the data before applying a filter could potentially lead to detectable bias. The $1/f$ noise fluctuations are matched using a 5th order polynomial fit to the signal-subtracted data, and the gap-filling procedure is completed with a white-noise realization.

2.3.9. Deconvolution of the bolometer time response

The time response of the HFI bolometers gives the relation between the optical signal incident on the bolometers and the output of the readout electronics, characterized by a gain, and a

time shift, dependent on the temporal frequency of the incoming signal (Planck Collaboration VII 2016). It is described by a linear complex transfer function in the frequency domain, called the “time transfer function”, which must be deconvolved from the data. NPIPE performs this deconvolution at the end of preprocessing. The gap-filled data are Fourier transformed, divided by the measured complex transfer function, and transformed signal back into the real domain. The complex transfer functions are identical to those used in PR3, and described in Planck Collaboration VII (2016).

We discovered that the application of the transfer function to glitch-removed and gap-filled data was picking up a considerable amount of power from the gap-filled samples and mixing it with the science data in the unflagged samples. To suppress this effect, our deconvolution module transforms a delta function signal and measures a window over which more than 10% of the deconvolved power can be attributed to the constrained realization present in the gaps. These samples are flagged in addition to the usual quality flags. The issue is illustrated in Fig. 7. The effect was not corrected for in PR3, contributing a small amount of sky-synchronous noise (especially at 100 GHz), responsible for several percent of the small-scale noise variance in the maps.




Fig. 7. Demonstration of the leakage of flagged signal onto science data. We created a zero (null) input signal (heavy blue line) with a gap between sample indices 0 and 10. We then filled the gap with an assumed signal = 1.0, then deconvolved using the 100-2b transfer function. Three samples immediately preceding the gap are significantly compromised by the signal in the gap. Such samples are dropped in the NPIPE analysis.

2.3.10. Outlier pointing-period detection

NPIPE accumulates statistics of the data for each pointing period during the preprocessing. Once all of the pointing periods are processed, the metadata are analysed for outlier periods, which are subsequently flagged completely. We used glitch rate, apparent gain, and noise rms to test for outliers by looking for $>5\sigma$ deviations in the statistics after removing a running average. Rings with more than three such deviations were flagged as outliers. Table 1 shows the fraction of outlier pointing periods at each *Planck* frequency. It also shows the total fraction of pointing periods discarded after we also reject unusable data from the sorption-cooler switchover at the end of first year of operations, and data acquired during a 10 day elevated spin-rate campaign.

2.4. Global reprocessing

NPIPE applies a number of templates to correct the data using a generalized destriper, much like the SRO11 method described in [Planck Collaboration Int. XLVI \(2016\)](#) for HFI large-scale polarization analysis. Our data model can be cast in the standard destripping form:

$$\mathbf{d} = \mathbf{P}\mathbf{m} + \mathbf{F}\mathbf{a} + \mathbf{n}, \quad (9)$$

where (multi-) detector time-ordered data \mathbf{d} are a linear combination of the sought-for sky map \mathbf{m} , sampled into the time domain with the pointing matrix \mathbf{P} , and an arbitrary number of time-domain templates presented as columns of the sparse template matrix \mathbf{F} and their amplitudes \mathbf{a} . When $1/f$ noise offsets are included in \mathbf{F} , the noise term, \mathbf{n} , is approximated as white noise.

If we choose not to impose a prior on the distribution of the template amplitudes, \mathbf{a} , we can marginalize over the sky map, \mathbf{m} , and solve for the template amplitudes ([Keihänen et al. 2004](#)):

$$\mathbf{a} = (\mathbf{F}^T \mathbf{N}^{-1} \mathbf{Z} \mathbf{F})^{-1} \mathbf{F}^T \mathbf{N}^{-1} \mathbf{Z} \mathbf{d}, \quad (10)$$

where

$$\mathbf{Z} = \mathbf{I} - \mathbf{P}(\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1}. \quad (11)$$

Traditionally, \mathbf{F} has comprised the baseline offset templates of a step function model of the $1/f$ noise¹⁰ – one column for each baseline period (anywhere between 1 s and 1 h) and every detector. Here, we add additional columns to \mathbf{F} for:

- linearized gain fluctuations (discussed in Sect. 2.4.1);
- HFI signal distortion (Sect. 2.4.2);
- orbital dipole (Sect. 2.4.3);
- far-sidelobe pickup (Sect. 2.4.4);
- HFI transfer-function residuals (Sect. 2.4.5);
- 6-component zodiacal light model (Sect. 2.4.6);
- bandpass mismatch (Sect. 2.4.7); and
- foreground polarization (Sect. 2.4.13).

Notice that there is no separate template for the Solar dipole; we consider it as an integral part of the sky, \mathbf{m} . For this reason, NPIPE calibration is not impacted by uncertainty in the Solar dipole estimate.

Polarization plays a key part in the degeneracy between the template amplitudes, \mathbf{a} , and the sky map \mathbf{m} in Eq. (9). Especially at the CMB frequencies, where the foreground polarization is fainter, *Planck* scanning strategy allows for apparent gain fluctuations that are compensated with scan-synchronous polarization residuals. LFI calibration in the 2015 and 2018 releases addressed this degeneracy by fixing the polarized part of \mathbf{m} to a polarized sky model derived using *Commander* ([Eriksen et al. 2004](#)). No such prior was used in the HFI 2015 or 2018 processing. We adopt the LFI approach here and also use a polarization prior, but rather than requiring repeated iterations between NPIPE and *Commander*, we sample smoothed NPIPE 30, 217, and 353 GHz maps as polarization templates in \mathbf{F} . When fitting for these polarization templates, the sky map \mathbf{m} is unpolarized, effectively preventing the generation of spurious gain fluctuations that would be compensated by polarization errors. This approach leads to significant improvement in the large-scale polarization systematics in the HFI maps, particularly at 100 GHz. Once time-dependent gain fluctuations and other time-varying signals are corrected for, the last reprocessing iteration freezes the gains and then solves for bandpass mismatch over a polarized sky, without fitting for polarization templates. This final fitting step also fits over a larger fraction of the sky than the previous iterations (Sect. 2.4.12).

NPIPE reprocessing solves for the template amplitudes, \mathbf{a} , in Eq. (9) and produces a template-cleaned TOD: $\mathbf{d}_{\text{clean}} = \mathbf{d} - \mathbf{F}\mathbf{a}$. Solving for the amplitudes is made nonlinear by the fact that the gain fluctuation template in \mathbf{F} is derived from the sky map, \mathbf{m} . Notionally, one could explicitly write out \mathbf{m} in \mathbf{F} in Eq. (9) and solve for the amplitudes using nonlinear minimization techniques. In practice, this is unfeasible owing to the dimensionality of the problem. Instead, we iterate between solving for the amplitudes, \mathbf{a} , cleaning the TOD and updating \mathbf{m} and \mathbf{F} . The iterations also enable speeding up the template generation and fitting by compressing the data onto phase-binned HEALPix rings¹¹,

¹⁰ We use the term $1/f$ -noise for instrument noise that has a power spectral density (PSD) that can be approximated with a power law: $P(f) \propto f^\alpha$, where f is the frequency and α is the slope of the spectrum. The approximation is valid up to a so-called knee frequency, where f^α transitions into a high frequency noise plateau and other features. See Fig. 27 for examples.

¹¹ Each pointing period is binned onto a subset of HEALPix pixels at the destripping resolution. We must exclude the repointing manoeuvres in doing so, since they break the repetitive scanning pattern we are leveraging.

Table 1. Discarded and quality-flagged data.

Frequency [GHz]	Outlier rings ^(a) [%]	Discarded rings ^(b) [%]	Discarded data ^(c) [%]	Relative total samples ^(d)	Relative masked samples ^(e)
30	0.71	1.88	2.62	1.119	1.095
44	0.57	1.76	2.19	1.120	1.098
70	0.71	1.89	2.84	1.114	1.091
100	0.97	2.17	20.00	1.167	1.163
143	1.16	2.36	21.53	1.167	1.160
217	0.95	2.16	20.98	1.221	1.209
353	0.90	2.08	21.78	1.179	1.174
545	0.89	2.06	15.93	1.068	1.066
857	0.93	2.11	17.76	1.019	1.015

Notes. ^(a)Rings with anomalous glitch rate, noise level, gain, or numerous jumps. A percentage of rings is not necessarily the same as a percentage of integration time, the ring length varies seasonally by almost a factor of 2. ^(b)Total fraction of rings discarded because of outlier statistics and spacecraft events such as the sorption-cooler switchover, elevated spin-rate campaign and LFI reboot on day 1085 since launch. ^(c)Total fraction of samples flagged because of pointing, jumps, cosmic-ray glitches, outlier pointing periods and spacecraft events. ^(d)Ratio of total integration time in the full-frequency maps, in the sense (NPIPE /2018). ^(e)Ratio of total integration time in the full-frequency maps, in the sense (NPIPE /2018), over the 50% of the sky with the least foregrounds. The distinction matters for LFI, because the ninth survey – included only in NPIPE – lies almost entirely in the Galactic plane.

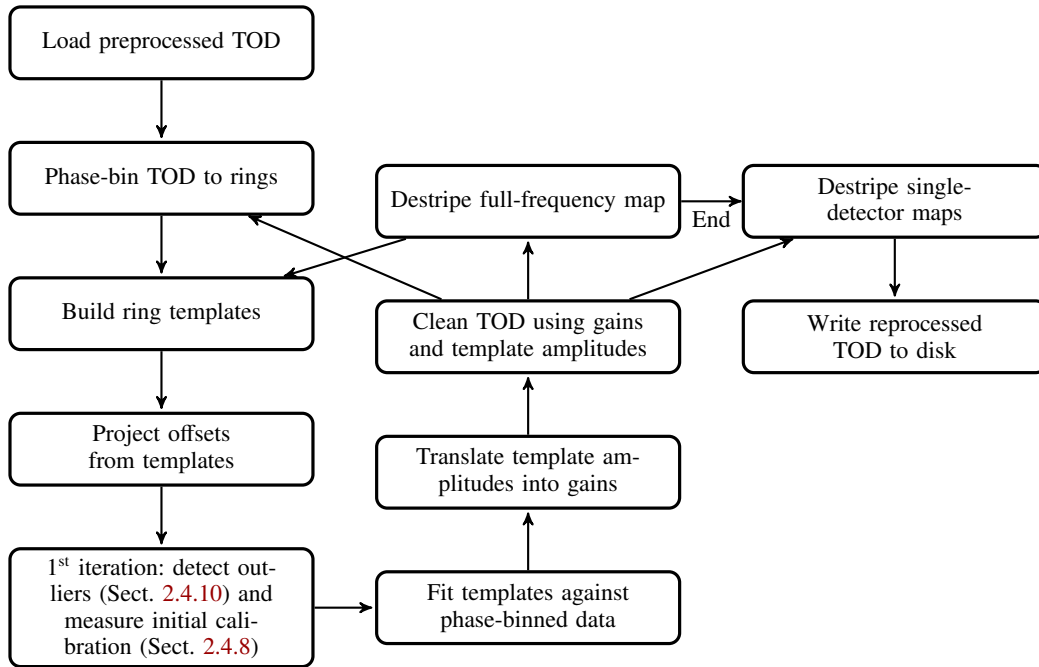


Fig. 8. NPIPE reprocessing flow chart. The *first two columns* contain the main iteration loop where the most recent frequency map is sampled into a new calibration template and the templates are iteratively fitted and subtracted from the TOD. Reprocessing finishes with the *third column* where single-detector data are de-polarized using the full-frequency map.

while retaining the fully sampled data for cleaning and mapmaking. NPIPE reprocessing iterations follow these steps:

1. compress TOD onto HEALPix rings;
2. solve for maximum likelihood template amplitudes \mathbf{a} in Eq. (9) using compressed data;
3. clean TOD using the amplitudes in \mathbf{a} ; and
4. update the gain template by running Madam on the cleaned TOD.

The recovered template amplitudes rapidly converge to zero as the cleaned data become consistent with the frequency estimate of the sky, \mathbf{m} . In practice we only need three iterations for the residual errors to become negligible.

Figure 8 shows the flow chart of NPIPE reprocessing. We will now discuss how we build each of the time-domain templates in F.

2.4.1. Gain fluctuations

Calibrating the detector data is a fundamentally nonlinear problem. We linearize the calibration problem by fitting a gain fluctuation template that measures deviations from the average gain for a given frequency. The amplitude of the template, δg , provides us with a gain correction that we may apply to the detector in question:

$$\mathbf{d}' = \frac{1}{1 + \delta g} \mathbf{d}. \quad (12)$$

Our gain template is the most recent iteration map downgraded to low resolution to improve S/N. The sky map already includes the large Solar dipole. Once the template is sampled, we add to it all of the time-domain templates that we have corrected for

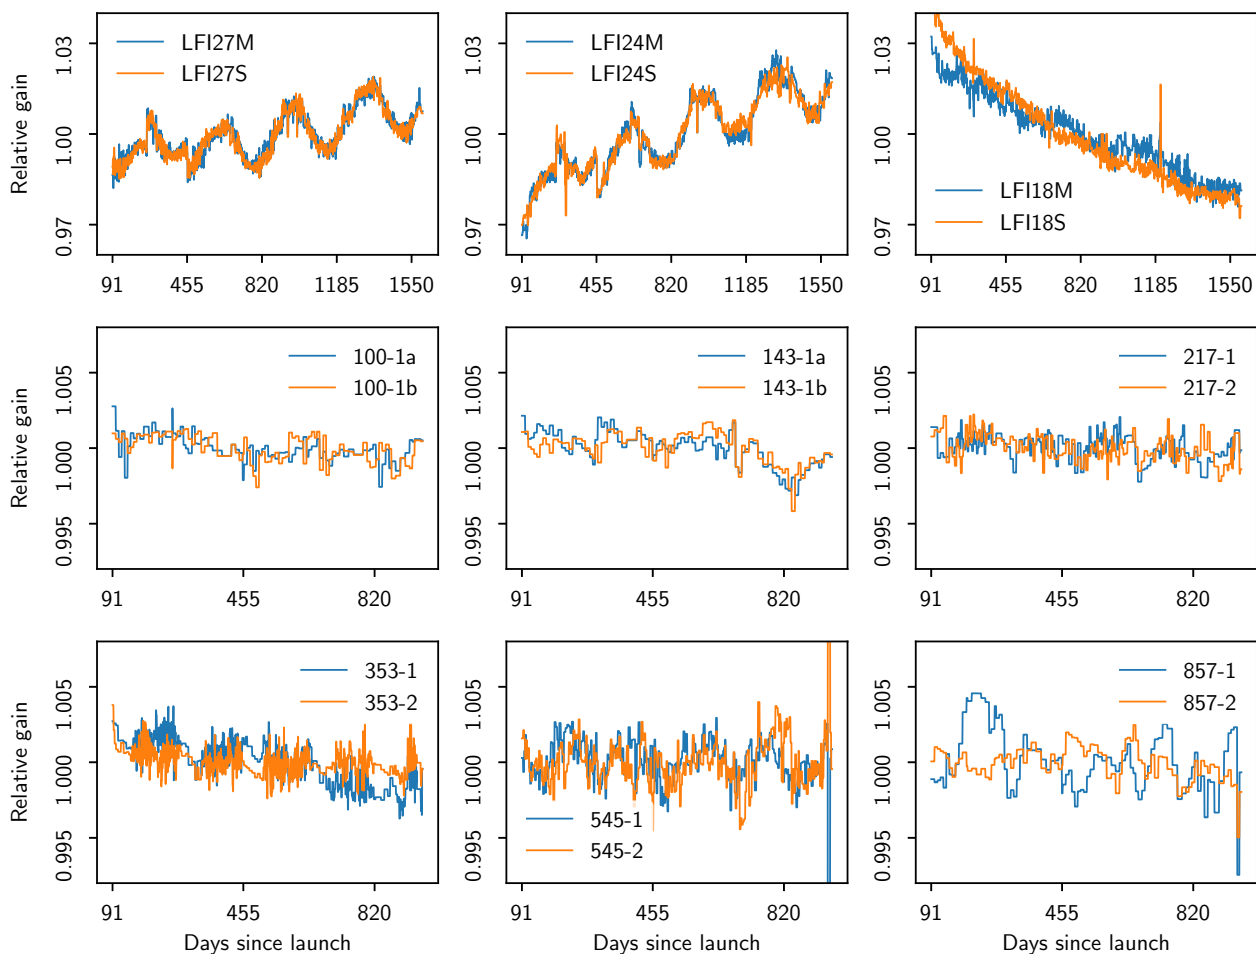


Fig. 9. Sampling of measured gain fluctuations at all *Planck* frequencies. The *top* row shows 30, 44, and 70 GHz, *from left to right*, for representative detectors, with “M” and “S” referring to “main” and “side” (see [Planck Collaboration III 2020](#)). The *second and third* rows show HFI frequencies, indicated in the detector names. The vertical lines indicate complete observing years. Seasonal effects due to Solar distance are apparent in the 30 and 44 GHz gains, as well as discontinuities from switching to the redundant 20 K sorption-cooler system on day 455 and changing the transponder setting on day 270.

in the cleaned TOD. That means adding the orbital dipole, far sidelobes, bandpass mismatch, HFI transfer-function residuals, and zodiacal light. It is true that a gain template sampled from the same sky map we are trying to estimate will have noise that is correlated with the detector TOD, but the effect is negligible because of the overwhelming number of samples accumulated into any sky pixel.

The gain template is split into disjoint chunks of time to trace changes in the gain. Each gain step is long enough to amount to roughly equal S/N, except for the few steps that are broken by known discontinuities or run into a preset maximum step length of 300 pointing periods (about 10 days). We have optimized the S/N threshold for each frequency to allow adequate tracking of gain fluctuations without excessively contributing to the map noise through noise-driven gain errors. We show the measured gain fluctuations for a subset of *Planck* detectors in Fig. 9. Unlike the LFI-DPC gain solution, the NPIPE gains are not filtered at all. Details of the filtering required to render LFI-DPC ring gains usable can be found in [Planck Collaboration V \(2016\)](#).

2.4.2. Residual HFI ADC nonlinearity

In previous *Planck* releases, residual HFI ADCNL was approximately addressed by correcting for apparent gain fluctuations

([Planck Collaboration VIII 2016](#)). Observed changes in the effective detector gain may be combined with the approximate level of the input signal to create an effective model of the residual ADC nonlinearity. This approach was shown to be effective in reducing large-scale polarization systematics in [Planck Collaboration III \(2020\)](#). We refer to this model of the ADCNL causing multiplicative, gain-like fluctuations in the TOD as the “linear gain model” of ADCNL.

Despite the success of the linear gain model in suppressing the nonlinearity errors, ADCNL was still the dominant source of large-scale polarization uncertainty in [Planck Collaboration III \(2020\)](#), prompting the use of inter-frequency cross-spectral methods for estimating the re-ionization optical depth from the maps ([Planck Collaboration Int. XLVI 2016](#)).

NPIPE makes two important updates to the handling of residual ADCNL. Firstly, we do not assume the apparent gain fluctuations to be exclusively sourced by the ADCNL. In this way, our gain steps are free to address both apparent and real gain fluctuations. Secondly, we add another calibration template that has all values above the median signal nulled. The function of this second “distortion” template is to enable us to adapt to signal distortions that arise from a different effective gain being imposed at opposite ends of the signal range. This is a natural extension of the linear gain model when the signal covers a wide dynamic range and the extremes of the signal probe different registers of

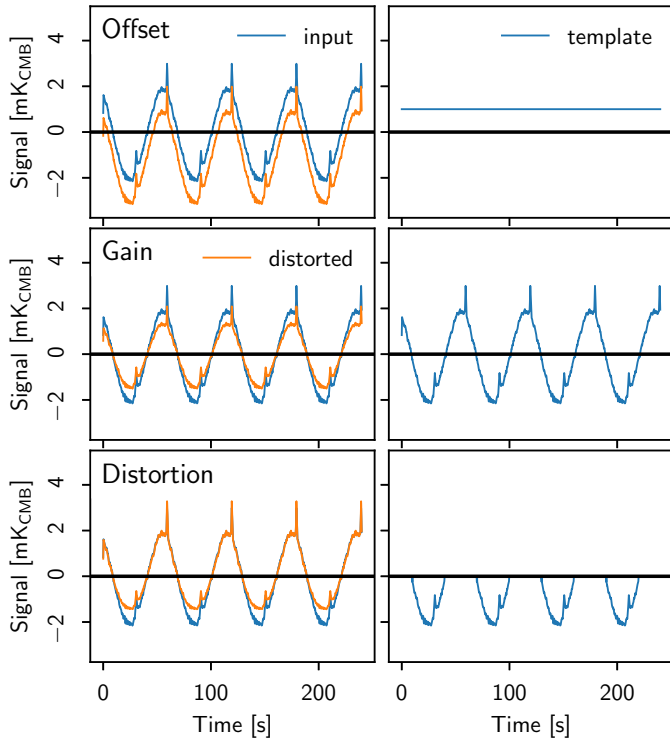


Fig. 10. Three degrees of signal distortion due to HFI ADC nonlinearity (*left panels*) and the NPIPE templates used to fit for them (*right panels*). The blue curves give an interval of data from a specific bolometer. The main variations seen come from the detector scanning across the dipole and the Galactic plane, including the ADCNL effects. The orange curves show the steps taken to fit the distortions. *Top*: (harmless) signal offset is captured by destriping baselines (i.e., moving the orange curve upwards). *Middle*: linear gain fluctuations are addressed by calibration. *Bottom*: linear gain varies as a function of signal level and requires a separate gain template for the opposite ends of the signal range. This correction is unique to NPIPE.

Table 2. HFI transfer-function bins.

Band	First harmonic ^(a)	f_{\min} [Hz]
1	1	0.025
2	2	0.042
3	4	0.075
4	6	0.108
5	9	0.158
6	13	0.225
7	18	0.308
8	25	0.425
9	34	0.575
10	46	0.775
11	61	1.025
12	81	1.358
13	106	1.775
14	139	2.325
15	181	3.025
16	235	3.925

Notes. ^(a)The satellite spin rate, 16.7 mHz.

the ADC. We demonstrate the first three leading orders of signal error due to ADCNL in Fig. 10. Any fluctuation in the TOD that can be fitted with the gain and distortion templates, could also be fitted with two distortion templates having the opposite signal halves nulled. Arranging the templates into gain- and distortion-only parts simplifies the interpretation. The template

corrections are applied to the TOD by multiplying the median-subtracted signal and half-signal (samples that have values below the median) with the appropriate factors. We show examples of the fitted distortion template amplitudes in Fig. 11.

It is worth emphasizing that the NPIPE model for the ADCNL is not a precise model that derives the correction from known inputs. Such a model is, in fact, impossible to construct, due to the down-sampling that occurred onboard the spacecraft. Instead, we simply fit piece-wise stationary time-domain templates to the mismatch between detectors, in an effort to reduce temperature-to-polarization leakage that would otherwise dominate the large-scale polarization. The shape of the templates is well motivated by our understanding of the ADCNL. The language of linear gain drifts is useful in describing the signal effects, but there is no way of disentangling actual changes in bolometer gain (expected to be of the order of 10^{-5}) from apparent drifts sourced by ADCNL.

2.4.3. Orbital dipole

The motion of *Planck* around the Sun causes a seasonal Doppler modulation of the CMB monopole. The spacecraft position is known to an accuracy better than 1 km and the velocity better than 1 cm s^{-1} (Godard et al. 2009). The CMB monopole is $2.72548 \text{ K} \pm 0.57 \text{ mK}$ (Fixsen 2009). Because both the spacecraft motion and the CMB monopole are extremely well known, this $\sim 300 \mu\text{K}_{\text{CMB}}$ dipole¹² signature is our best absolute calibrator. We fit for an orbital dipole template in each detector TOD and find a per-frequency calibration coefficient that brings the average orbital dipole template to the expected level. We do not require that all detectors observe an identical dipole, since far sidelobe effects alter the perceived amplitude at the individual horn level. We do, however, subtract the same orbital dipole from both detectors in a horn. Far sidelobe effects are approximately corrected by convolving the orbital dipole template up to the second-order quadrupole (sometimes called the “kinematic quadrupole”) with a GRASP¹³ estimate of the far sidelobes.

The orbital dipole fits provide for absolute calibration for all frequency channels except for 857 GHz, where NPIPE uses the same planetary calibration as in PR3.

2.4.4. Far-sidelobe pickup

We calculate timeline estimates of the far-sidelobe (FSL) pickup by convolving PR2 full-frequency maps (with a Solar dipole added) with a model of the beam. We considered updating the convolved timelines from PR3 or NPIPE temperature maps, but the differences would have been negligible.

To perform this convolution, we used available GRASP estimates (which are monochromatic) of the 4π detector beams. Full spectral treatment of the frequency-dependent FSL would have amounted to a percent-level correction to a 10^{-3} level systematic. The estimates are only available for the 30–353 GHz channels, so we used 353 GHz beams as proxies for 545 and 857 GHz channels. Except at the submillimetre frequencies, the FSL template is too faint and degenerate to include in the fitting. We simply subtract it and correct the amplitude of the subtracted

¹² The range of the orbital dipole signature depends seasonally on the phase of the 7.5° spin axis precession and the position of each detector on the focal plane.

¹³ The GRASP software was developed by TICRA (Copenhagen, DK) for analysing general reflector antennas (<https://www.ticra.com/>).




Fig. 11. Sampling of measured signal distortions (cf. Sect. 2.4.2) at 100, 143, and 217 GHz. These are the only frequencies where the distortion template is fitted. The distortion templates are fit at the same time steps as the gain templates in Fig. 9.

Table 3. Zodiacal template emissivities.

Frequency [GHz]	λ [μm]	Cloud	Band 1	Band 2	Band 3	Blob	Ring
857	350	0.285	0.53	0.274	0.833	0.317	0.782
545	550	0.182	0.973	0.32	1.37	0.214	0.438
353	849	0.105	0.954	0.207	1.00	0.0883	0.449
217	1382	0.0391	0.823	0.0999	0.685	0.0211	0.332
143	2096	0.0106	0.777	0.0693	0.494	0.0207	0.243
100	2998	0.0124	0.68	0.0779	0.523	-0.0327	-0.0282

Notes. The statistical uncertainties for these amplitudes are found to be negligible. See [Planck Collaboration XIV \(2014\)](#) for an explanation of the components. Note that NPIPE only fits and corrects for the time-varying part of the Zodiacal emission model.

template as our estimate of the gain fluctuations improves. For 545 and 857 GHz, the FSL template is successfully fitted with the individual SWBs, yielding roughly consistent best-fit amplitudes between the detectors.

2.4.5. HFI transfer-function residuals

The time response of the bolometers to sky signals is complicated. HFI uses planet observations and stacked cosmic-ray glitches to infer the shape of the bolometer transfer function and deconvolve it. These calibrators are excellent for probing the intermediate- and high-frequency components of the transfer function, but leave an uncertainty about the relative values of the low- and high-frequency components. The most striking effect is a relative change in phase between the Solar dipole and signals on much smaller angular scales. We construct templates that correspond to the real and imaginary parts of a binned residual transfer function, fit the templates, construct the transfer function from the fit amplitudes, and deconvolve it from the signal. The real part of the transfer function corresponds to a small, scale-dependent change in calibration, and can be fit reliably only at 353 GHz and above. The imaginary part introduces a phase shift between different scales, and is fit at all HFI frequencies.

We quantize the transfer-function bins to contain one or more harmonics of the 16.7 mHz spin frequency, and exponentially increase the bin width according to

$$w_b = \text{floor}(e^{b/4})/60 \text{ s}, \quad (13)$$

where b is the band index, w_b is the width of the band in Hz, the first band begins at 0.5/60 s, and the “floor” function returns the largest integer value that is smaller than its argument. The last band is extended all the way to the Nyquist frequency. To avoid a degeneracy with overall calibration, we omit the first band of

the real transfer-function templates. We find that we can reliably fit up to 16 bands at 353 GHz and above.

In the HFI 2018 processing, the empirical transfer function was fitted in only four bins. A pattern of “zebra stripes” was identified in the 353 GHz odd–even survey difference maps and shown to be caused by the inability of the coarse-grained transfer-function model to track the actual residual. The odd–even survey differences for both methods are described further in Sect. 6.3.1.

2.4.6. Zodiacal emission

Interplanetary dust in the Solar System forms a cloud through which we observe the CMB. Thermal emission from zodiacal dust is visible at high *Planck* frequencies ([Planck Collaboration XIV 2014](#)), and has characteristic seasonal variations as the satellite moves with respect to the cloud. The seasonal variations are of concern for our template fitting, since they might bias the gains and noise offsets. We use the six geometric components described by [Kelsall et al. \(1998\)](#) to construct six independent templates that capture the seasonal variations. To preserve the mean zodiacal emission in the signal, our templates are the differences of expected emission between the current spacecraft position and another position 180° along Earth’s orbit. The opposite position is found by inverting the heliocentric position vector. Furthermore, to boost the S/N we use the same template amplitude for all detectors in a frequency channel.

Our recovered zodiacal template amplitudes are roughly in agreement with the 2018 results ([Planck Collaboration III 2020](#)), but offer no new information, owing to the removal of the common mode and the emphasis on cleaning the data of seasonal effects. We defer the actual zodiacal emission removal to the component-separation process, where multi-frequency

information offers better means of cleaning the frequencies where the emission is faint. The recovered template emissivities are listed in Table 3.

One may assess the success of the zodiacal emission removal by examining the difference between maps made from odd and even surveys. Comparison of NPIPE and PR3 odd–even survey differences (see Sect. 6.3.1) suggests that NPIPE processing achieves at least the same level of zodiacal residuals as the 2018 processing; however, it is hard to quantify the performance because of other features in the same survey-difference maps.

2.4.7. Bandpass mismatch

Planck detectors within a frequency band have differing bandpasses, causing them to see foreground components at different intensities. In theory, one can use a frequency-dependent model of the foreground emission, coupled with a set of measured detector bandpasses, to predict the amount of bandpass mismatch and derive correction maps. In practice, uncertainties in both the measured bandpasses and the sky model require fitting the corrections directly against the data. For all continuum-emission foregrounds (synchrotron, free-free, anomalous microwave emission, and dust), we use a high-resolution *Commander* sky model to predict the frequency derivative of the foreground emission, pixel-by-pixel, at the centre of the frequency band. The fitting amplitude directly gives us an estimate of the relative difference in centre frequency between the detectors. We also consider narrow-line emission from CO regions, which are fitted with a separate template. The bandpass templates are *HEALPix* maps that we bilinearly interpolate to the detector sample positions.

For added flexibility, we choose to include some of the foreground components as separate templates. This compromises any direct physical interpretation of the amplitude of the bandpass-mismatch template (i.e., difference in centre frequency), but leads to better agreement between full-frequency and single-detector maps. For the 30–70 GHz channels, these extra templates are the frequency derivatives of the *Commander* anomalous microwave emission (AME) and dust frequency components, chosen for their similarity to maps of the full-frequency versus single-detector mismatch. For 70 GHz we also fit for an HCN ($J = 1 \rightarrow 0$, 88.63 GHz) emission-line template. For 100–857 GHz, we use the frequency derivative of the free-free component instead of AME or dust.

2.4.8. Initial calibration

In order to speed up convergence, we perform a linear regression of the raw TOD against the 20% of the sky nearest to the dipole extrema. The target is the 2015 *Planck* dipole, which we use as a crude initial calibration. The fit provides an initial, fixed, calibration coefficient for each detector; these are overridden by the orbital dipole fit and relative gain corrections during the first reprocessing iteration.

2.4.9. Submillimetre processing

Planck 545 and 857 GHz frequency channels have unique processing requirements. Foregrounds dominate across the sky, wide bandpasses cover a number of potential emission lines, and the detectors, although nominally polarization-independent, are in fact weakly sensitive to polarization, with polarization efficiencies measured to be around 6%. In contrast, the HFI PSBs

have values around 90%. For 545 GHz, we marginalize over the NPIPE 30, 217, and 353 GHz polarization maps, smoothed to 60', to project out detector mismatch due to polarization. For 857 GHz, we found that this approach does not improve detector agreement, and chose to overlook the small polarization sensitivity in reprocessing.

2.4.10. Outlier detection

NPIPE reprocessing searches for outlier pointing periods in the data at two stages. First, a submatrix of F is fitted against the TOD, one detector and one pointing period at a time. Pointing periods showing anomalous fitting coefficients are flagged. Second, the destriper estimates the initial residual for every detector and every pointing period by binning, resampling, and subtracting the input TOD. Again, anomalous periods are flagged. These steps result in fewer than ten pointing periods permanently flagged for each detector. The total numbers of additional discarded pointing periods (across all detectors) at each frequency is 23, 25, 65, 9, 16, 18, 14, 0, and 11 for 30 GHz through 857 GHz, respectively.

2.4.11. Horn symmetrization

Each pair of polarization-orthogonal detectors that shares a horn forms a single polarization-sensitive unit, with nearly identical optical properties. We enforce this by:

- applying a joint set of quality flags to both detectors and dropping the same outlier data for each;
- using a common noise weight for both detectors;
- using the same pointing;
- using the same FSL template;
- using the same zodiacal-light templates and template amplitudes.

The purpose of the symmetrization is to limit the temperature-to-polarization (“T-to-P”) leakage. Since the detectors in a horn share the same optical feed, the difference in samples collected simultaneously from the two is free from most of the optical mismatch that would otherwise cause leakage. In fact, our symmetrization scheme has the same effect as differencing the TOD from the two detectors, then using the resulting difference to solve for sky polarization.

One potential source of T-to-P leakage is the sub-pixel structure that includes sharp foreground features as well as the dipole gradient. Samples drawn from various corners of the pixel may indicate different signal levels, and if those observations are associated with different observing angles, the difference can be interpreted as polarization. To the extent that the polarized pairs of detectors are optically aligned, this source of leakage is blocked by our symmetrization procedures.

We have verified that the symmetrization reduces T-to-P leakage in the Galactic plane and in the vicinity of compact sources, where beam effects are pronounced. While the effect is small, it was nevertheless considered strong enough to merit the increase in noise from symmetrizing the flags and using slightly less optimal noise weights for the detectors.

2.4.12. Last iteration and processing masks

Each iteration of the reprocessing uses a processing mask to avoid issues arising from sub-pixel structure, variable unresolved (“point”) sources, and small incompatibilities in our bandpass-mismatch templates. The same mask is not necessarily optimal

for all of the templates. In particular, the gains require only the cleanest and most reliable parts of the sky to be considered, while the bandpass corrections can greatly benefit from including as much as possible of the intense Galactic emission. To best serve both applications, we mask more of the sky while we iterate over the calibration, then freeze the gain solution and run the last iteration with more of the Galaxy exposed.

Point sources, especially variable radio sources, and other steep signal gradients on the sky may cause an error in the fitted template amplitudes. We use processing masks to avoid regions of the sky where the signal model is inaccurate.

NPIPE reprocessing uses apodized processing masks. This means that it is possible to down-weight regions of the sky without discarding them completely when solving for the template amplitudes. This approach is particularly useful for measuring the gain fluctuations both at the dipole extrema and over intense Galactic emission. The use of apodized processing masks is a new development in TOD processing, but has been used extensively in map domain analysis.

We build our apodized masks from an earlier iteration of NPIPE temperature maps. First, we subtract PR2 dipole and CMB map from each frequency map, and then find the scaling coefficients that suppress the foreground amplitude in each pixel below a fixed threshold. These thresholds are not rigorously optimized, but rather are chosen to visually suppress regions of the sky where single-detector maps show notable disagreement with the full-frequency map. Further adjustment is performed when the gain solutions show sharp and recurring seasonal trends. Finally, point sources are added to the masks from the second *Planck* compact source catalogue ([Planck Collaboration XXVI 2016](#)).

We present our processing masks in Fig. 12. The final short-baseline destriping uses binary masks, but we are able to use a large sky fraction due to the time-domain bandpass-mismatch corrections.

2.4.13. Polarization estimates

The full polarized solution of the template amplitudes is degenerate at CMB frequencies and subject to large noise uncertainty. Without external constraints, the template fitting converges to a self-consistent but sub-optimal solution, with large-scale polarization power leaking in from the dipole (see Sect. 6.3.1). These issues can be remedied by performing the template fits on a temperature-only sky. Polarization cannot be overlooked without biasing the solution, but it is possible to marginalize over the sky polarization by adding additional templates that model the foreground polarization to the template matrix. To this end, the 217 GHz reprocessing fits for 30 and 353 GHz polarization templates smoothed to 1° , while 44–143 GHz reprocessing fits for 30, 217, and 353 GHz. For each template, we only allow for a single template amplitude across all detectors. This improves the S/N at the cost of ignoring the polarized bandpass mismatch. We do not explicitly fit for bandpass mismatch in polarization, but the sky model does predict a percent level correction based on the central frequencies fitted with temperature which is included when applying the bandpass correction.

To construct the time-domain templates, we smooth each foreground frequency map with a 1° -FWHM Gaussian beam and sample according to the detector polarization parameters and orientation.

2.4.14. Degeneracies

Including sampled sky maps as templates in Eq. (9) creates a degeneracy. For instance, the same bandpass-mismatch template is sampled for every detector, and we can shift power between the template amplitudes in \mathbf{a} and the sky map, \mathbf{m} . Such degeneracies prevents direct solution of Eq. (10); instead, we perform a conjugate gradient (CG) iteration for \mathbf{a} that minimizes the residual in

$$(\mathbf{F}^T \mathbf{N}^{-1} \mathbf{Z} \mathbf{F}) \mathbf{a} = \mathbf{F}^T \mathbf{N}^{-1} \mathbf{Z} \mathbf{d}. \quad (14)$$

Once the CG iteration converges, we adjust the template amplitudes, \mathbf{a} , as follows:

- the bandpass-mismatch correction for each frequency must average to zero;
- relative gain corrections for each frequency must average to zero; and
- the average of all noise offsets must equal zero.

These adjusted template amplitudes are then used to correct the TOD. An alternative (but equivalent) approach would have been to modify the destriping equation itself to include these priors.

The matter of degenerate modes is a common topic in destriping. They arise due to the projection matrix, \mathbf{Z} , which removes the sky-synchronous part of the signal. Even in the case of traditional destriping that only fits for noise offsets, an overall offset of the solved map is set by the initial guess rather than being solved from the data. Beyond that, any sky-synchronous mode in the templates is invisible to the destriper.

The unadjusted average of all of the gain fluctuations tends to be large (of order unity rather than zero) and matched by an opposite average value in the orbital dipole template. This simply reflects a degeneracy between the two templates, because the gain template includes the orbital dipole. To deal with this, we subtract the frequency average of the gain fluctuations and adjust the orbital dipole amplitudes accordingly to keep the level of total orbital dipole unchanged.

There is a degeneracy between the templates used to model foreground polarization and the templates used to capture bandpass mismatch. This degeneracy requires no active mitigation due to the way bandpass-mismatch correction is only applied in the final template fitting step that does not involve polarization templates. For more discussion, see Appendix H.

2.4.15. Polarization parameters

During the NPIPE development a number of diagnostics were produced to assess the performance of the template corrections. One very useful diagnostic is a set of polarization-corrected, single-detector maps reflecting the final template amplitudes. Any mismatch between the full-frequency temperature map and these single detector maps can be considered a residual and an indication of a deficiency in the data model or the fitting procedure. Once other sources of residuals have been addressed, the remainder, especially for some of the 353 GHz bolometers, exhibits a residual signal that closely resembles maps of polarized sky emission. We suggest that these residuals arise from small errors in the polarization efficiencies or angles that we adopted from the DPCs.

We derive two time-domain polarization templates by scanning the full-frequency polarization map, one with the usual Q and U weights, and another using the polarization angle derivative of the Q and U weights. We then bin these two

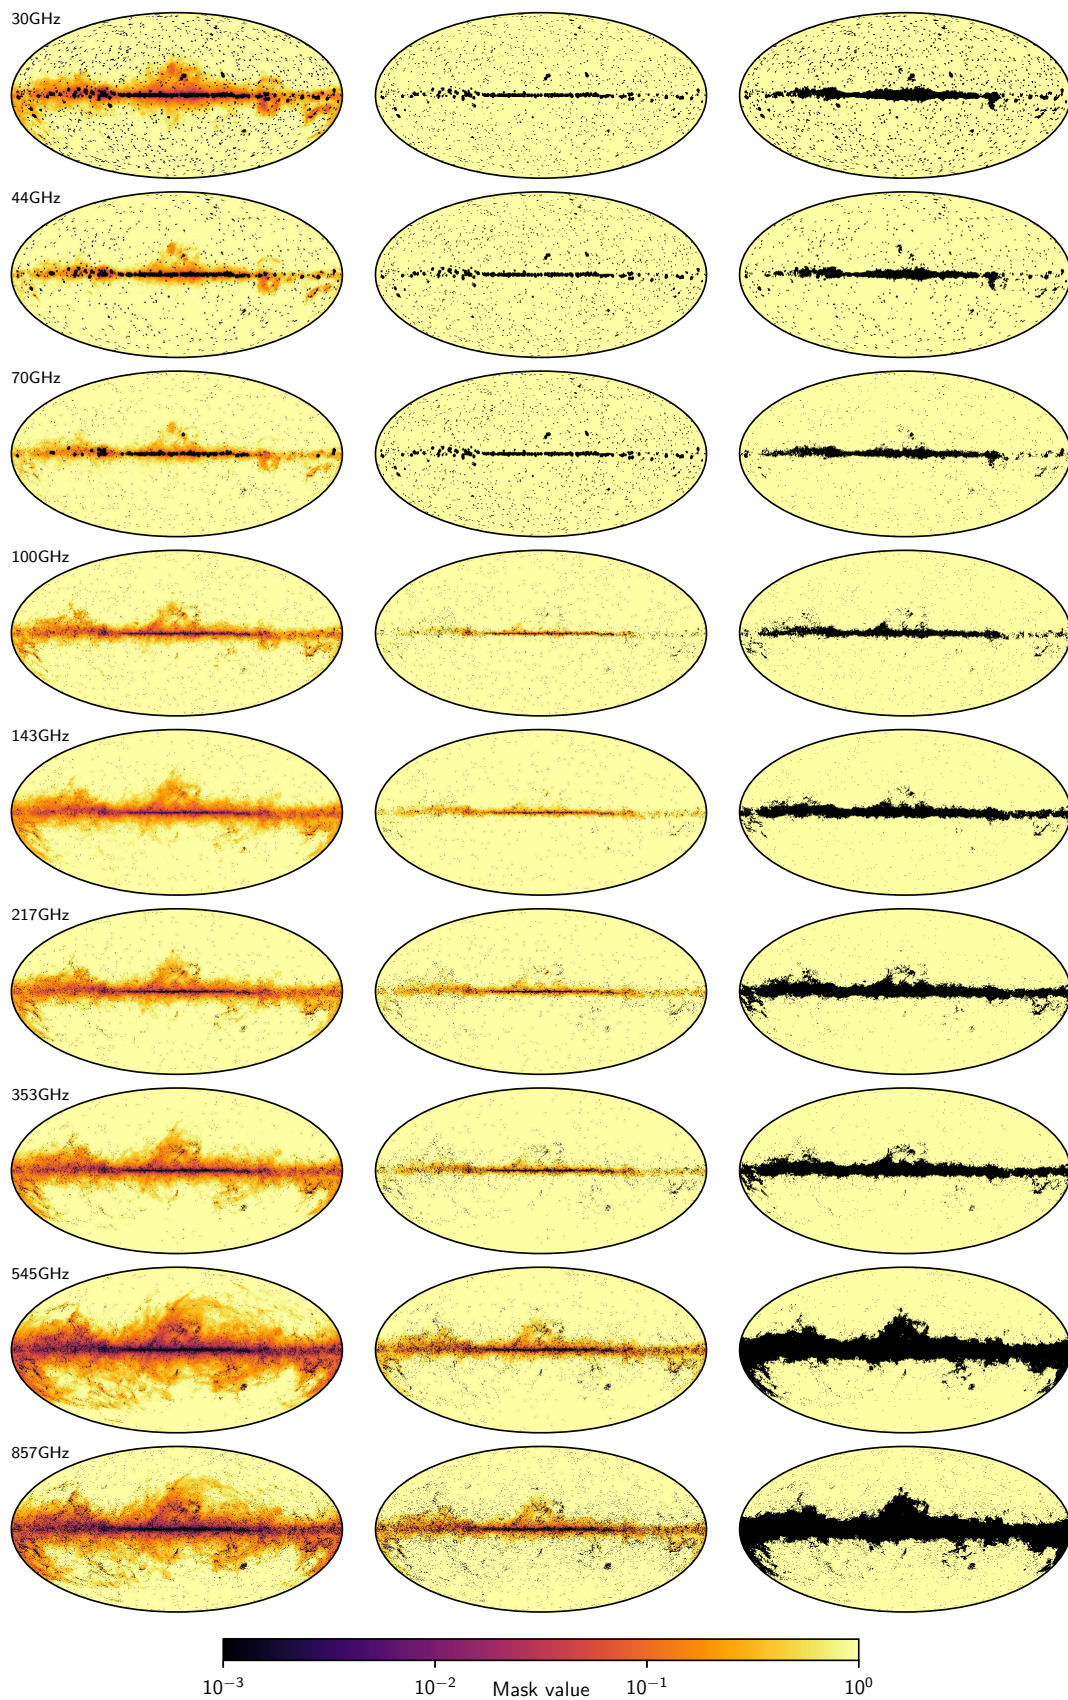


Fig. 12. Sharp and apodized processing masks used in NPIPE processing. The three columns are for the calibration (*left*), bandpass correction (*middle*), and destriping (*right*) masks. The destriping masks are explicitly binary because Madam (which we use for this step) does not support floating-point masks.

Table 4. NPIPE polarization efficiencies and angles.

Bolometer	Pol. efficiency [%]	Δ [%]	Pol. angle [deg]	Δ [deg]
100-1a	94.35 ± 1.03	−0.36	23.32 ± 0.28	+0.22
100-1b	96.98 ± 0.87	+2.67	111.93 ± 0.24	+0.13
100-2a	96.03 ± 0.45	−0.15	44.26 ± 0.18	−0.84
100-2b	89.74 ± 0.45	−0.50	134.57 ± 0.21	+0.27
100-3a	89.56 ± 0.71	−0.53	179.33 ± 0.14	−0.87
100-3b	93.23 ± 1.06	−0.22	90.29 ± 0.20	+0.19
100-4a	93.88 ± 0.63	−1.83	156.96 ± 0.16	+0.46
100-4b	93.02 ± 0.69	+0.75	67.95 ± 0.17	−0.06
143-1a	85.55 ± 0.20	+2.32	45.74 ± 0.09	+0.24
143-1b	86.43 ± 0.34	+1.84	137.66 ± 0.12	−0.04
143-2a	85.35 ± 0.28	−2.09	45.19 ± 0.11	−0.21
143-2b	88.97 ± 0.20	−0.21	135.49 ± 0.09	+0.29
143-3a	86.03 ± 0.45	+2.13	0.08 ± 0.09	+0.38
143-3b	88.69 ± 0.35	−1.25	92.74 ± 0.09	−0.36
143-4a	91.75 ± 0.40	−1.36	0.61 ± 0.08	−0.29
143-4b	93.72 ± 0.39	+0.88	89.61 ± 0.09	+0.31
143-5	7.52 ± 0.30	+1.25	68.95 ± 1.03	+3.25
143-6	3.87 ± 0.38	−0.30	69.71 ± 2.06	−0.89
143-7	3.06 ± 0.32	+1.74	95.24 ± 2.96	−7.56
217-1	2.34 ± 0.20	−1.50	97.11 ± 2.18	−1.29
217-2	3.33 ± 0.20	+1.34	95.87 ± 1.57	+13.37
217-3	1.23 ± 0.15	−2.72	163.02 ± 2.27	−7.88
217-4	3.78 ± 0.14	−0.61	121.20 ± 1.56	+1.20
217-5a	96.14 ± 0.12	+1.13	47.08 ± 0.04	+0.08
217-5b	97.35 ± 0.15	+2.15	135.88 ± 0.08	−0.22
217-6a	90.12 ± 0.09	−4.82	46.01 ± 0.06	−0.09
217-6b	93.91 ± 0.09	−1.48	135.97 ± 0.04	+0.07
217-7a	94.09 ± 0.15	+0.04	179.43 ± 0.02	−0.17
217-7b	95.15 ± 0.14	+1.48	90.53 ± 0.03	+0.03
217-8a	96.76 ± 0.20	+2.55	1.18 ± 0.04	+0.88
217-8b	95.80 ± 0.20	+1.68	89.90 ± 0.05	−0.80
353-1	0.23 ± 0.13	−2.97	101.74 ± 18.81	−1.36
353-2	7.49 ± 0.05	+2.78	120.68 ± 0.63	+6.08
353-3a	89.93 ± 0.09	+0.95	45.70 ± 0.06	+0.20
353-3b	90.90 ± 0.11	−1.17	133.64 ± 0.07	−0.26
353-4a	87.35 ± 0.13	+0.14	45.67 ± 0.05	−0.03
353-4b	89.85 ± 0.08	−1.74	135.52 ± 0.06	−0.08
353-5a	85.44 ± 0.11	+0.75	177.87 ± 0.02	+0.17
353-5b	90.95 ± 0.11	+3.35	89.43 ± 0.04	−0.17
353-6a	84.99 ± 0.58	−2.55	0.22 ± 0.12	+0.52
353-6b	92.56 ± 0.66	+3.79	88.86 ± 0.09	−0.64
353-7	5.31 ± 0.12	−2.68	126.81 ± 0.70	+5.31
353-8	9.88 ± 0.08	+2.06	139.30 ± 0.35	+6.30

Notes. We report corrections with respect to the ground-measured values used in PR3. In Rosset et al. (2010) the statistical uncertainties for polarization efficiency are estimated to be 0.1–0.6% for PSBs and SWBs operating between 100 and 353 GHz. Polarization angle systematic uncertainties for the PSBs were estimated by Rosset et al. (2010) to be 0°9 (not including a 0°3 common error); uncertainties for the SWBs were estimated to be up to 5°5. For NPIPE, statistical uncertainties in the PSB polarization angles are small, while those for the SWBs are up to 2°3. The NPIPE uncertainties include statistical and systematic errors estimated using 60 simulations.

templates using the same intensity weights as with the polarization-corrected, single-detector TOD. Fitting the resulting two polarization templates against the single-detector residual map allows us to measure a correction to the polarization efficiency and polarization angle that we incorporate into a refined version of the instrument model. The template fits are limited by other residuals in the single-detector maps, so their

uncertainties are systematic, not statistical, and are difficult to assess without simulations.

We have been able to reduce the single-detector/full-frequency residuals only from 100 to 353 GHz (as shown in Table 4). The LFI residual maps do not show a residual consistent with an error in the polarization parameters, either due to lack of such error or due to higher noise content.

One might argue that the binned full-mission maps are not appropriate for fitting the polarization parameters because they do not capture the fact that the detectors are rotated with respect to the pixels between different surveys. We tried fitting for the polarized templates during destripping (similar to other templates), but found that the change in position angle over most of the sky is insufficient and the template fits are unreliable. This is why the polarization parameter corrections are based on map domain fits.

3. Destripping

If the relevant systematics are sufficiently suppressed, low-frequency noise fluctuations in each detector begin to dominate large-scale uncertainties in the final maps. In an effort to extract maximal information from the data, we have adopted an aggressive approach to destripping, fitting each detector with 167 ms baseline offsets that correspond to 1° steps on the sky. These short baselines are used to capture as much of the low-frequency instrumental noise as is statistically possible, but they also adapt to other large-scale residuals such as ADC nonlinearities and gain fluctuations. The baseline-offset solution is regularized by a prior on the baseline distribution (Keihänen et al. 2005, 2010), which is derived from the detector-noise power spectral densities (PSDs; see Sect. 5.1).

We tested the efficacy of the short baselines by destripping and projecting dark-bolometer data from the two dark bolometers as if they were bolometers 857-1 and 857-3. Figure 13 shows the power spectra of maps destripped with various baseline lengths. Using the dark bolometer data instead of synthetic TOD has the advantage that they include realistic glitch and 4 K line residuals.

In choosing the destripping resolution, we have balanced S/N, which improves with lower destripping resolution, with signal error, which is best when there is no detectable sub-pixel structure. Our destripping resolution is $N_{\text{side}} = 512$ for all LFI frequencies, $N_{\text{side}} = 1024$ for 100–353 GHz, and $N_{\text{side}} = 2048$ for 545 and 857 GHz. These resolutions allow for modest sub-pixel signal gradients outside the destripping mask, but do not lead to temperature-to-polarization leakage due to the horn symmetrization (Sect. 2.4.11).

Solving for short baselines also serves a secondary purpose in allowing us to apply the Madam approach for estimating the residual noise in the pixel-to-pixel covariance in a low-resolution version of the NPIPE data set. In Madam, the noise is broken down into baseline offsets and white noise. When the baselines are too long to capture the $1/f$ noise fluctuations, Madam can only accommodate the extra power in the white noise component of the signal model, leading to mismatch between the noise matrix and the actual noise in the maps.

4. NPIPE data set

We now describe the contents of the NPIPE data release, highlighting differences between NPIPE and other *Planck* public releases that may be relevant to the user. For information regarding the dissemination of the maps and the related software, see Appendix A.




Fig. 13. Power spectra of destriped dark-bolometer maps using various baseline lengths. The “full” case on the *left panel* shows the pseudo- C_ℓ power spectrum over the entire map, including the Galactic plane that was masked out while solving for the baseline offsets. *Right panel*: power spectrum taken over the sky pixels that were considered when solving for the baselines. Shorter baselines clearly suppress noise above $\ell = 20$. The 167 ms case shows an elevated residual at $\ell < 100$ inside the destriping mask, likely due to the filter’s inability to constrain the solution. The residual at $\ell < 20$ is dominated by the irreducible uncertainty set by the *Planck* scan strategy, and cannot be improved without external information.

4.1. Maps

All NPIPE maps are calibrated to thermodynamic temperature units in kelvins (K_{CMB}). The 857 GHz calibration is the same as in PR3, but is converted into K_{CMB} temperature units using the measured bandpasses. Note that NPIPE maps include the Solar dipole; it may be removed using the parameters supplied in Sect. 8.

The conversion factor from thermodynamic temperature to flux density units is $58.04 \text{ MJy sr}^{-1} / K_{\text{CMB}}$ for the 545 GHz maps and $2.27 \text{ MJy sr}^{-1} / K_{\text{CMB}}$ for the 857 GHz maps (see [Planck Collaboration IX 2014](#)).

4.1.1. Frequency maps

We present temperature maps at all nine *Planck* frequencies in Fig. 14, and with the Solar dipole removed in Fig. 15. The zero levels of the maps are adjusted for plotting by evaluating the monopole outside a Galactic mask. The mapmaking procedure leaves the true monopole undetermined. Stokes Q and U polarization maps and the polarization amplitude at the seven polarized frequencies are shown in Fig. 16. Polarization maps smoothed to 3° are shown in Fig. 17.

4.1.2. Half-ring maps

The repetitive *Planck* scanning strategy (each scanning circle is repeated 30–75 times) allows us to build subset maps that have effectively identical systematics but independent instrument noise. We split each pointing period into two half-rings, and assign the halves to separate subsets, which we call “HR1” and “HR2”. The half-ring difference (i.e., HR1–HR2) is a useful measure of the instrumental noise, but should not be assumed to be unbiased for large-scale cross-spectral analysis. The half-ring maps share all the calibration, bandpass mismatch, and other template residuals.

The noise in the HFI half-ring maps contains a small amount of correlated error between the two ring halves (cf. Sect. 2.3.4

and Fig. 18). This error comes from a noisy signal estimate used in the removal of the glitches.

4.1.3. Single-detector maps

NPIPE builds unpolarized single-detector maps from the reprocessed TOD by subtracting an estimate of the polarization response from the full-frequency map. The single-detector maps are destriped with Madam independently, using the same destriping parameters as for the full-frequency data. The maps are provided with and without bandpass-mismatch corrections. Mismatch-corrected maps provide a useful diagnostic for reprocessing performance, while the uncorrected maps are used as inputs to temperature component separation. This provides component-separation algorithms with samples of the sky for each detector bandpass, which is necessary, for example, to measure the CO line emission. Sampling and subtracting the polarization signal from the full-frequency map injects a small amount of correlated noise into the single-detector maps. This noise is attenuated by the fact that we downgrade the frequency map to the same resolution as our bandpass-mismatch templates, i.e., $N_{\text{side}} = 512$ for 30–100 GHz and $N_{\text{side}} = 1024$ for 143–353 GHz. Most of this correlated noise is cancelled in the single-horn maps (Sect. 4.1.4), so applications that are particularly sensitive to small-scale noise correlations are best served using the horn maps instead.

4.1.4. Single-horn maps

In addition to polarization-corrected single-detector maps, we also provide unpolarized single-horn maps that naturally cancel most if not all of the polarization. These maps are linear combinations of the polarization-orthogonal detector data in the polarized *Planck* horns, with the weights chosen to maximally cancel polarization in the averaged map. The process also eliminates noise present in the resampled polarization estimate.

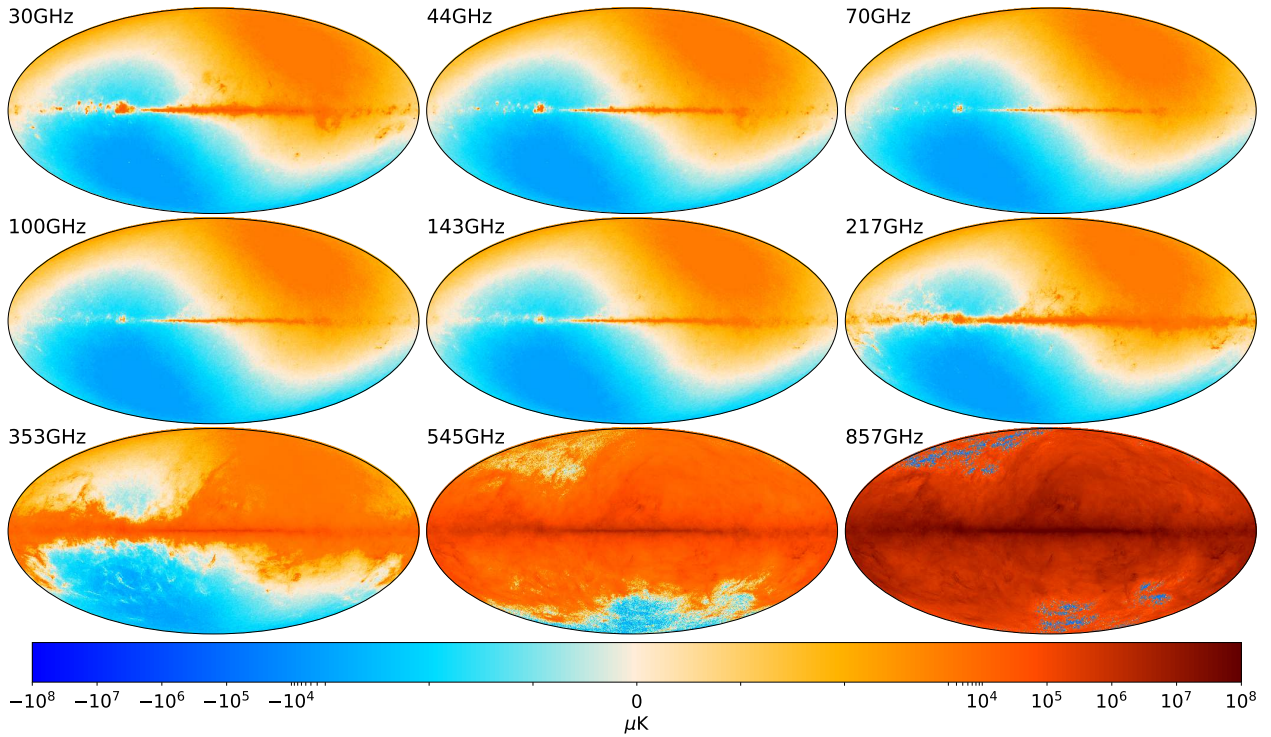


Fig. 14. NPIPE temperature maps, including the Solar dipole. The scaling is linear between -3 and 3 mK.

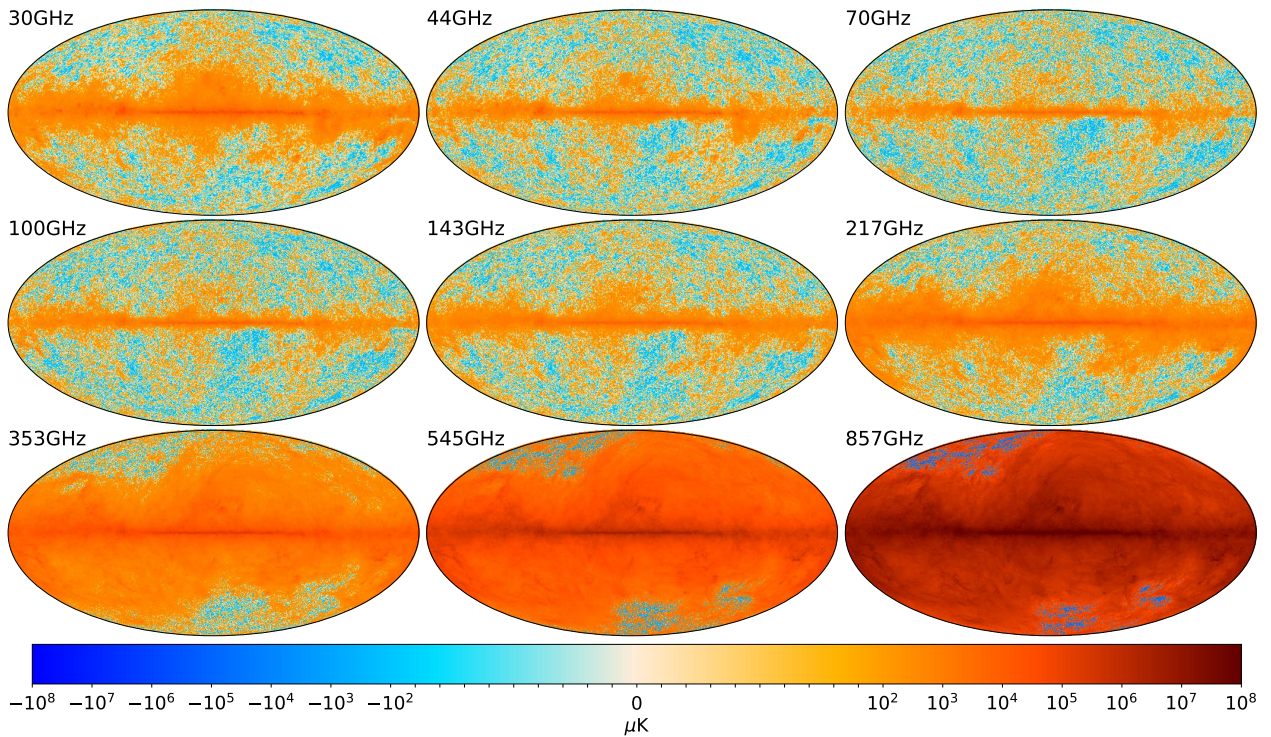


Fig. 15. NPIPE temperature maps, with the *Planck* 2015 dipole removed. The scaling is linear between -100 and 100 μ K.

4.1.5. A/B maps

NPIPE provides only one data split where the systematics between the splits are expected to be uncorrelated (unlike the case for the half-ring splits discussed in Sect. 4.1.2). We have split the horns in the focal plane into two independent sets, A and B, and performed reprocessing independently on each set.

The *Planck* scanning strategy requires at least two polarized horns to solve for a full-sky polarized map. The detector-set splitting was not possible at either 30 or 44 GHz because of the lack of redundant polarized horns. Instead, for these two frequencies, the split is done time-wise: set A comprises operational years 1 and 3; while set B has years 2, 4, and the month of integration time from the final, fifth observing year.




Fig. 16. NPIPE polarization maps. The scaling is linear between -100 and $100\mu\text{K}$. Here $P = (Q^2 + U^2)^{1/2}$.

(A half-mission split is not feasible because the second half mission does not have full sky coverage.) The time-wise split is not expected to have independent systematics, since the two subsets will share initial beam and bandpass mismatch; however, the instrument noise and gain fluctuations will be uncorrelated between the two subsets, making the time-wise split at these

frequencies a reasonable compromise to provide two consistent and maximally disjoint sets of *Planck* data. Full descriptions of the A/B splits are provided in Table 5.

The NPIPE approach of processing maximally-independent subsets of data is orthogonal to the approach adopted by HFI in PR3. There the systematic templates were fitted on the entire




Fig. 17. NPIPE polarization maps smoothed with a Gaussian beam with FWHM 3° . The scaling is linear between -3 and $3 \mu\text{K}$.

frequency data set, and the cleaned data were split only for the purpose of projecting the TOD into subset maps. Such processing introduces correlations between systematic residuals in the subset maps, making it necessary to estimate and correct for

noise bias even in cross-spectra between the subsets. This is part of the reason [Planck Collaboration III \(2020\)](#) emphasizes the importance of using inter-frequency cross-spectra in extracting scientific information from PR3 maps.

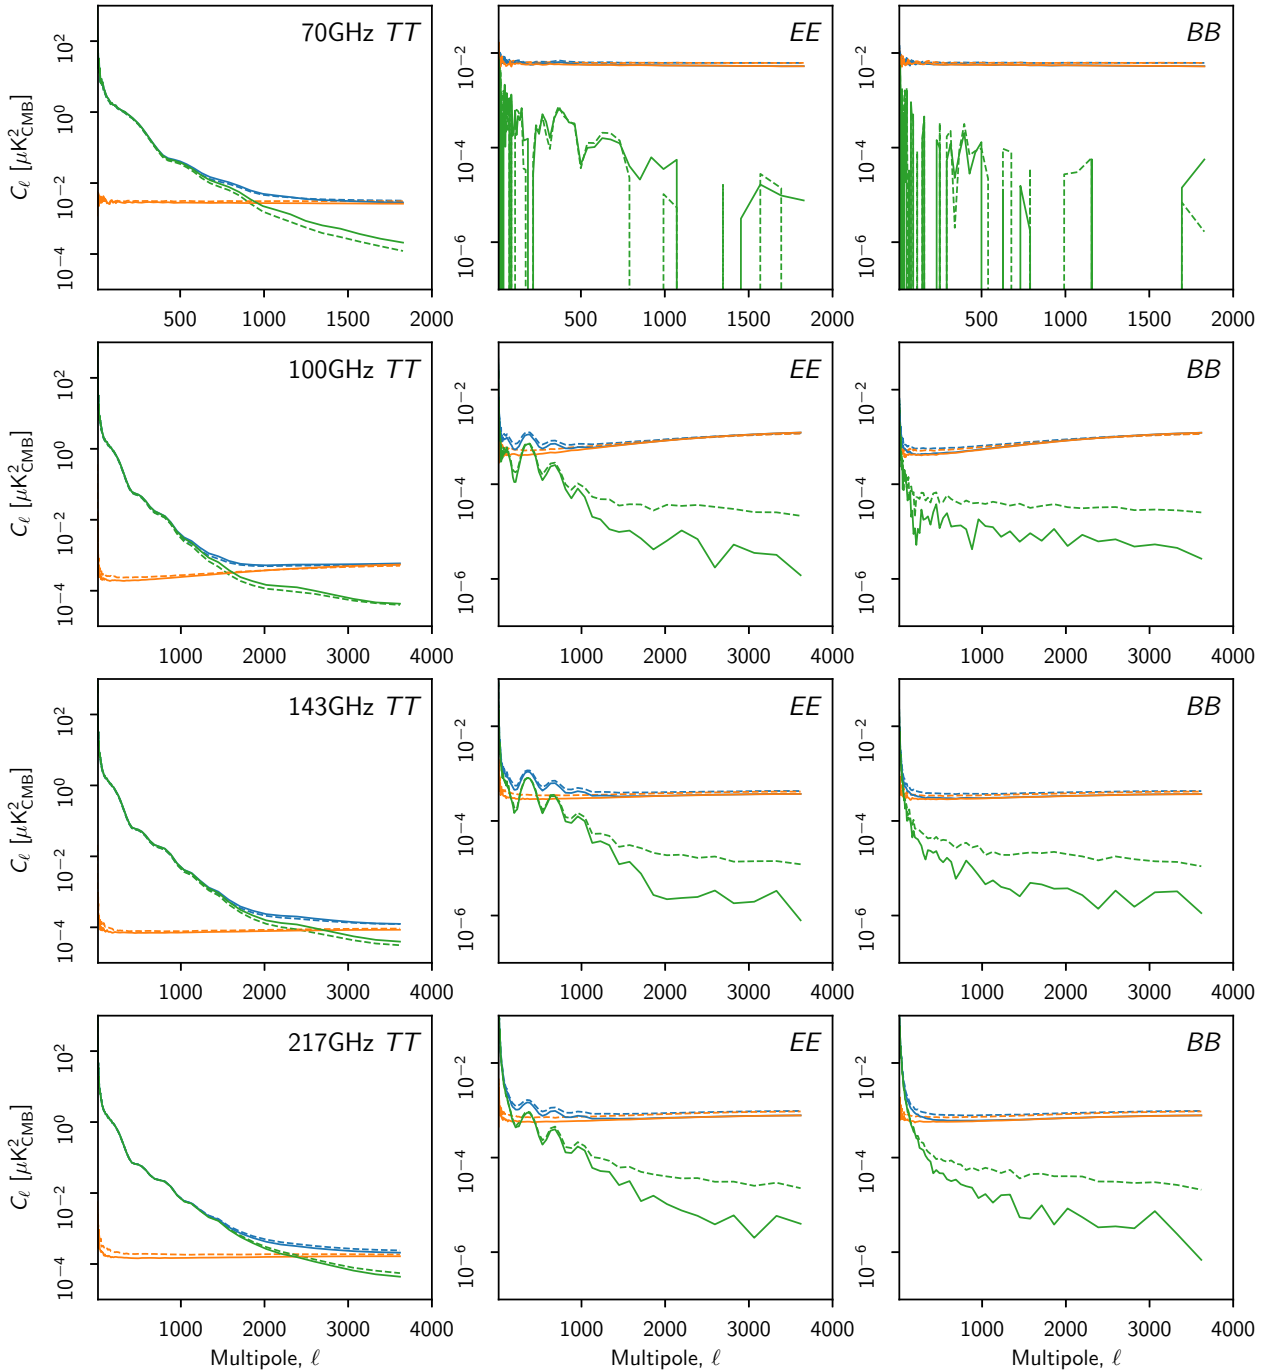


Fig. 18. Half-ring sum (blue), difference (orange), and cross (green) power spectra from NPIPE (solid lines) and PR3 (dashed lines). Improvements in gap-filling and transfer-function deconvolution reduce the small-scale half-ring correlations in NPIPE data from PR3 levels. The power spectra are binned into 100 logarithmically-spaced bins, and corrected for sky fraction. The 70 GHz results are shown as a reference case without half-ring correlations.

4.2. Low-resolution data set

We provide low-resolution versions of the NPIPE maps for maximum-likelihood pixel-domain analysis of the large scales. The maps are downgraded to HEALPix resolution $N_{\text{side}} = 16$ and convolved with a cosine apodizing kernel:

$$b_\ell = \begin{cases} 1, & \ell \leq \ell_1, \\ \frac{1}{2} \left(1 + \cos \left[\pi \frac{(\ell - \ell_1)}{(\ell_2 - \ell_1)} \right] \right), & \ell_1 < \ell \leq \ell_2, \\ 0, & \ell > \ell_2, \end{cases} \quad (15)$$

with the choice of $\ell_1 = 1$, $\ell_2 = 3N_{\text{side}}$. The lower threshold has been decreased from the usual $\ell_1 = N_{\text{side}}$ to reduce ringing in

the 857 GHz maps. In addition to the smoothing kernel, the maps include the standard $N_{\text{side}} = 16$ pixel window function. Our downgrading tool noise-weights each high resolution pixel with the estimated noise level and parallel-transport the polarization parameters between pixel centres. Initial high resolution pixel window function was not deconvolved as it is negligible at the angular scales that are represented in the low resolution maps.

The low-resolution maps are accompanied by pixel-pixel noise-covariance matrices that reflect the baseline uncertainties and the high-frequency instrumental noise approximated as uniform white noise. Since the high-frequency noise in NPIPE is

Table 5. Details of the A/B subset split.

Frequency [GHz]	Set A	Set B
30	Years 1 and 3	Years 2, 4, and start of 5
44	Years 1 and 3	Years 2, 4, and start of 5
70	Horns 18, 20, and 23	Horns 19, 21, and 22
100	Horns 1 and 4	Horns 2 and 3
143	Horns 1, 3, 5, and 7	Horns 2, 4, and 6
217	Horns 1, 3, 5, and 7	Horns 2, 4, 6, and 8
353	Horns 1, 3, 5, and 7	Horns 2, 4, 6, and 8
545	Horn 1	Horns 2 and 4
857	Horns 1 and 3	Horns 2 and 4




Fig. 19. Scaling factors applied to the Madam noise matrices to match the variance in half-ring-difference pixels. We apply different scalings to the temperature and polarization parts of the noise matrices, and scale the detector-set matrices and the full-frequency matrices independently. The 100 GHz channel stands out because of the prominent bump at the high-frequency end of the noise PSD (cf. Fig. 27).

not uniform for any of the *Planck* detectors, the approximation is inexact and, as a result, requires an extra scaling step performed on the diagonal of the matrix to match the half-ring difference maps. The off-diagonal elements correspond to low frequency noise and need not be scaled. We show the scaling factors in Fig. 19. For the CMB channels, these factors are close to unity except at 100 GHz.

4.3. Large-scale polarization transfer function

NPIPE calibration at the CMB frequencies (44–217 GHz) acts like a constrained filter. The polarized sky model that we derive from the extreme frequencies contains, for all practical purposes, only the foregrounds. This leaves NPIPE actively trying to suppress the CMB polarization with the available templates. While obviously undesirable, the effect of this filtering is kept small by the number and structure of the templates. It is also straightforward to measure the amount of suppression using simulations and incorporate the effect into the transfer functions.

Allowing for the non-trivial transfer function in NPIPE calibration is a compromise between measuring very noisy but unbiased large-scale polarization from all large-scale modes, and filtering out the modes that are most compromised by the *Planck* calibration uncertainties left in the data by the *Planck* scan strategy (see Figs. 40 and 41). Based on similarities between LFI

2018 processing and NPIPE it is possible that the LFI 2018 results were also subject to filtering. Unfortunately the accompanying simulations were not sufficient to determine this as only one CMB realization was processed with the calibration pipeline. Algebraically there is no reason to anticipate the HFI 2018 results to contain a similar effect but, again, the 2018 simulations cannot be used to demonstrate this.

The filtering approach is typically adopted by ground and balloon-borne experiments to suppress atmospheric noise in their data. The only requirement for the filtering to be workable is that the resulting ℓ -dependent transfer function be measurable. We show in Fig. 20 that it is well characterized in the case of NPIPE analysis of *Planck* data. Tabulated values of the transfer function can be found in Appendix G.

4.3.1. CMB transfer function

We measure the transfer function from our simulations by comparing the input CMB map to the output foreground-cleaned map. We smooth and downgrade both maps and compare the input (CMB \times CMB) cross-spectrum to the CMB \times cleaned map. If the effect is multiplicative, as is evident from Fig. 22, we expect to find

$$C_{\ell}^{\text{CMB} \times \text{output}} = k_{\ell} C_{\ell}^{\text{CMB} \times \text{CMB}}, \quad (16)$$

where k_{ℓ} is a potentially scale-, frequency-, and mask-dependent suppression factor. The left panel of Fig. 20 shows the k_{ℓ} for each CMB frequency. The right panel of Fig. 20 compares the 100 GHz transfer functions evaluated over sky fractions ranging from 0.1 to 0.9 and shows to what degree the shape of the transfer function depends on the mask. These transfer functions, like the beam window functions, must be squared to obtain the full impact on a power spectrum. We find that the *EE* quadrupole and octupole are significantly suppressed by our calibration, with about 64% of the quadrupole and 73% of the octupole missing in the NPIPE power spectra. Multipoles 4 and 5 are missing about 30% of the power, and $\ell = 6$ is missing about 20%. The simulations do not include enough CMB *B* modes to measure the *B*-mode transfer function and changing the signal content would compromise our interpretation due to the nonlinear calibration process.

The suppression of *EE* power is limited to frequencies that are calibrated with the polarization prior. There is no suppression at 30 GHz or 353 GHz (Fig. 21). Measuring the transfer function using the CMB polarization at these frequencies is much harder because of the foregrounds and noise. This is evidenced by the large upwards fluctuations at 353 GHz, particularly at $\ell = 4$ and 9. These fluctuations are associated with excess power at the level of the CMB *E*-mode fluctuations, rather than suppression of power. Their importance diminishes when these frequency maps are scaled to the CMB frequencies to estimate the foregrounds.

Figure 22 shows individual input-output *EE* multipole pairs for 100 GHz. The data points show a great deal of scatter owing to the limited S/N. Nevertheless, fitting for a linear model between the input and output values robustly identifies the degree to which the input *EE* power is suppressed by the NPIPE calibration. It is also apparent that a simple multiplicative correction will remedy the bias for all realizations.

Frequency and detector-set cross-spectra may be impacted by a different transfer function. The measurement also potentially differs, since we must compare the input spectrum to the foreground-cleaned cross-spectrum:

$$C_{\ell}^{\text{output1} \times \text{output2}} = k_{\ell}^2 C_{\ell}^{\text{CMB} \times \text{CMB}}. \quad (17)$$

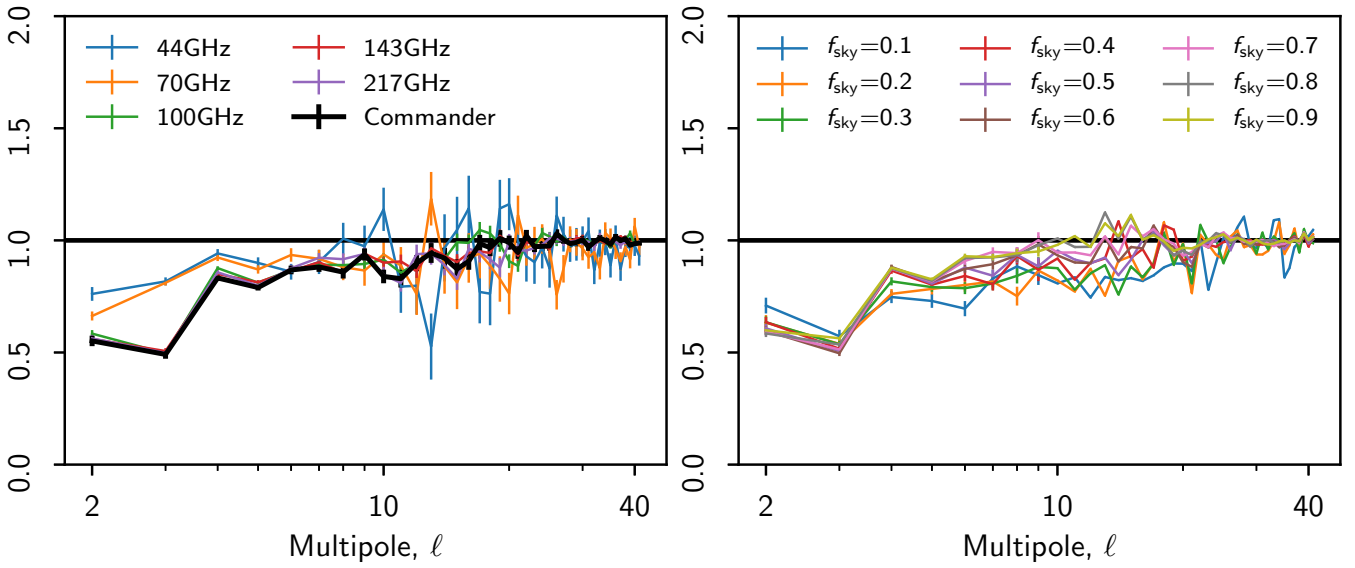


Fig. 20. NPIPE E -mode transfer functions measured by comparing simulated CMB input and foreground-cleaned output maps. *Left:* CMB frequencies and component-separated Commander maps (Sect. 7) over 60% of the sky. The apparent mismatch between the LFI and HFI transfer functions results from the quantity and structure of the template corrections; templates that are specific to HFI, especially the ADC distortion, provide more degrees of freedom to suppress the CMB power. The 44 GHz transfer function is closer to unity because the 30 GHz template shields about 22% of the CMB polarization. The error bars reflect the statistical uncertainty of the measured transfer function, not the total Monte Carlo scatter. Tabulated values of the transfer functions are listed in Table G.1. *Right:* E -mode transfer function for 100 GHz over multiple sky fractions. The error bars at $\ell \geq 10$ were suppressed to show more structure.




Fig. 21. NPIPE E -mode transfer functions, measured by comparing simulated CMB input and foreground-cleaned output maps over 60% of the sky. The 30 and 353 GHz frequencies are not expected to have a measurable transfer function because they are not calibrated with a polarization prior. These two transfer functions are merely of diagnostic value (to demonstrate the absence of signal suppression) and are not applied in any analysis. The error bars reflect the statistical uncertainty of the measured transfer function, not the total Monte Carlo scatter. Tabulated values of the transfer functions are listed in Table G.1.

We find that the use of the simulated cross-spectra to be problematic for two reasons: they are much noisier than the CMB \times output spectra; and they may contain persistent bandpass-mismatch residuals. For these reasons, we instead derive the cross-spectral transfer functions as the geometric mean of the

individual transfer functions. The geometric mean was found to be in agreement with the Eq. (17) and offered superior noise performance. We show the measured E -mode transfer functions for detector-set and frequency cross-spectra in Fig. 23.

We also considered measuring the full, anisotropic transfer function, a complex multiplier applied to every $a_{\ell m}$ mode of the input CMB sky. The results are shown in Fig. 24. We cannot reliably expand the foreground-cleaned CMB map from the simulations without masking out the strongest Galactic residuals. This is problematic for a study of anisotropy, because the mask already includes a preferred orientation. Instead, we write a least-squares minimization problem in terms of a transfer function acting on the input CMB expansion:

$$\chi^2 = \mathbf{r}^T \mathbf{r} \quad \text{with} \quad \mathbf{r} = \mathbf{m} - \sum_{X\ell m} f_{\ell m}^X a_{\ell m}^X Y_{\ell m}^X, \quad X \in [T, E, B] \quad (18)$$

where the residual, \mathbf{r} , is estimated by subtracting a CMB expansion ($a_{\ell m}^X Y_{\ell m}^X$) convolved with a transfer function ($f_{\ell m}^X$). In this shorthand notation each basis function, $Y_{\ell m}^X$, is a full IQU map with either the I or the QU part identically zero. The sum in Eq. (18) is conveniently carried out using the HEALPix `a1m2map` facility. This formulation allows us to use full-sky input expansions and still evaluate the residual over a masked sky.

The anisotropic transfer functions in Fig. 24 show a statistically significant anisotropy, especially at $\ell = 3$. This suggests that fully isotropic methods for power spectrum estimation and transfer function correction are not optimal in analysing the large-scale polarization in the NPIPE maps. Certain modes could be down-weighted to minimize uncertainty. Nevertheless, specializing the methods to utilize this information is complicated and not attempted in this paper.

To gain insight into the magnitude of the effect, we take a simulated CMB sky with a realistic amount of large-scale polarization power, and apply the 143 GHz anisotropic transfer

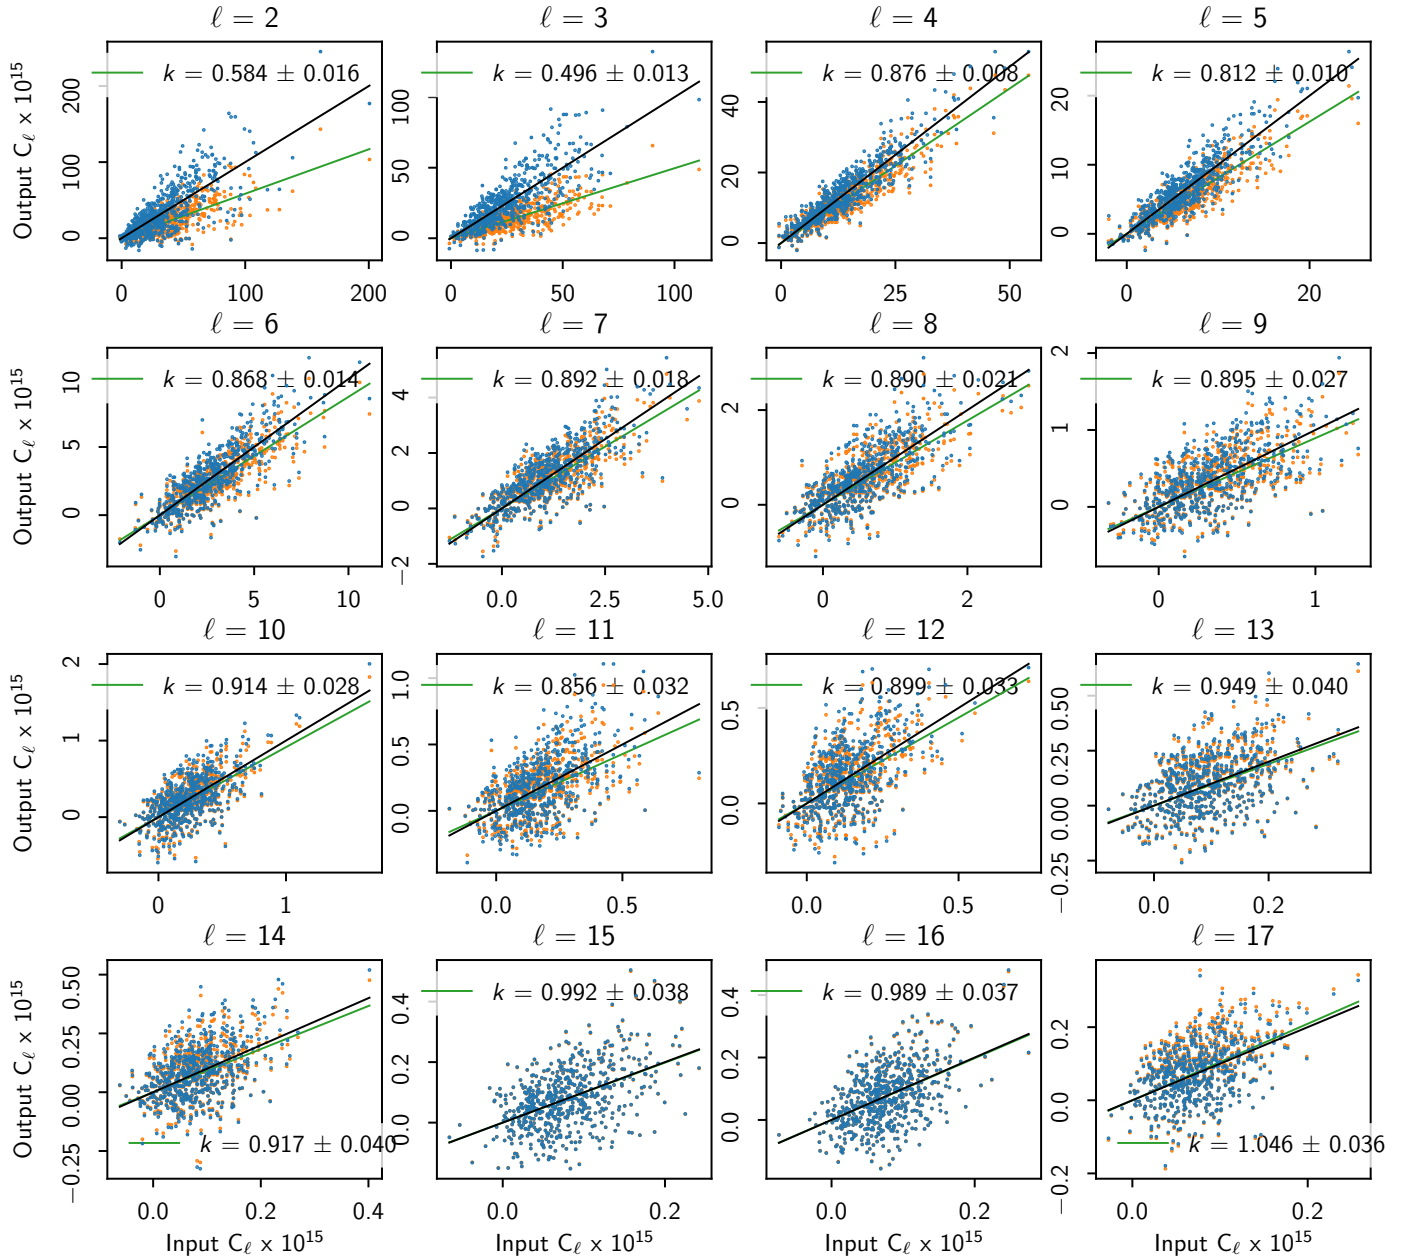


Fig. 22. Measuring the NPIPE large-scale polarization transfer function by comparing the simulated input and output CMB power multipole by multipole at 100 GHz and with 60% sky fraction. The output power spectrum is measured by removing the foreground from the simulated frequency map and then measuring the cross-spectrum with the input CMB map. The orange data points show input versus output before the transfer function correction. Their slope is shown in green. The blue data points show the corrected input versus output. Their (unity) slope is shown in black.

function from Fig. 24 to it. The input, output, and difference maps are shown in Fig. 25.

We provide the measured E -mode transfer functions for all frequency auto-spectra, frequency cross-spectra, and detector-set cross-spectra. The format of the files is exactly the same as the QuickPol beam-window-function files used in Planck Collaboration V (2020). In these files, the transfer function is unity everywhere except for the $\ell < 42$ E modes. These transfer functions were measured with the 60% sky mask. It is impossible to anticipate all sky masks and map combinations that a user may require, so we also provide software that enables measurement of the transfer function for arbitrary sky masks. To ensure fidelity of the science results, any statistic that employs

large-scale ($\ell < 10$) CMB polarization should be corrected for the transfer function or otherwise calibrated with the simulations. This is the baseline for all NPIPE products.

4.3.2. Non-CMB transfer function

There is no algebraic reason to expect that the suppression of large-scale CMB polarization would extend to Galactic or extragalactic foregrounds. Indeed, measurement of these signals is driven by the 30 and 353 GHz channels, neither of which is subject to the polarization prior in the calibration. Furthermore, the presence of polarized foregrounds is accounted for in the calibration process by marginalizing over the polarization templates.




Fig. 23. NPIPE E -mode cross-spectrum transfer functions measured by comparing simulated CMB input and foreground-cleaned output maps over 60% of the sky. *Left:* detector-set cross-spectra. *Right:* frequency cross-spectra.

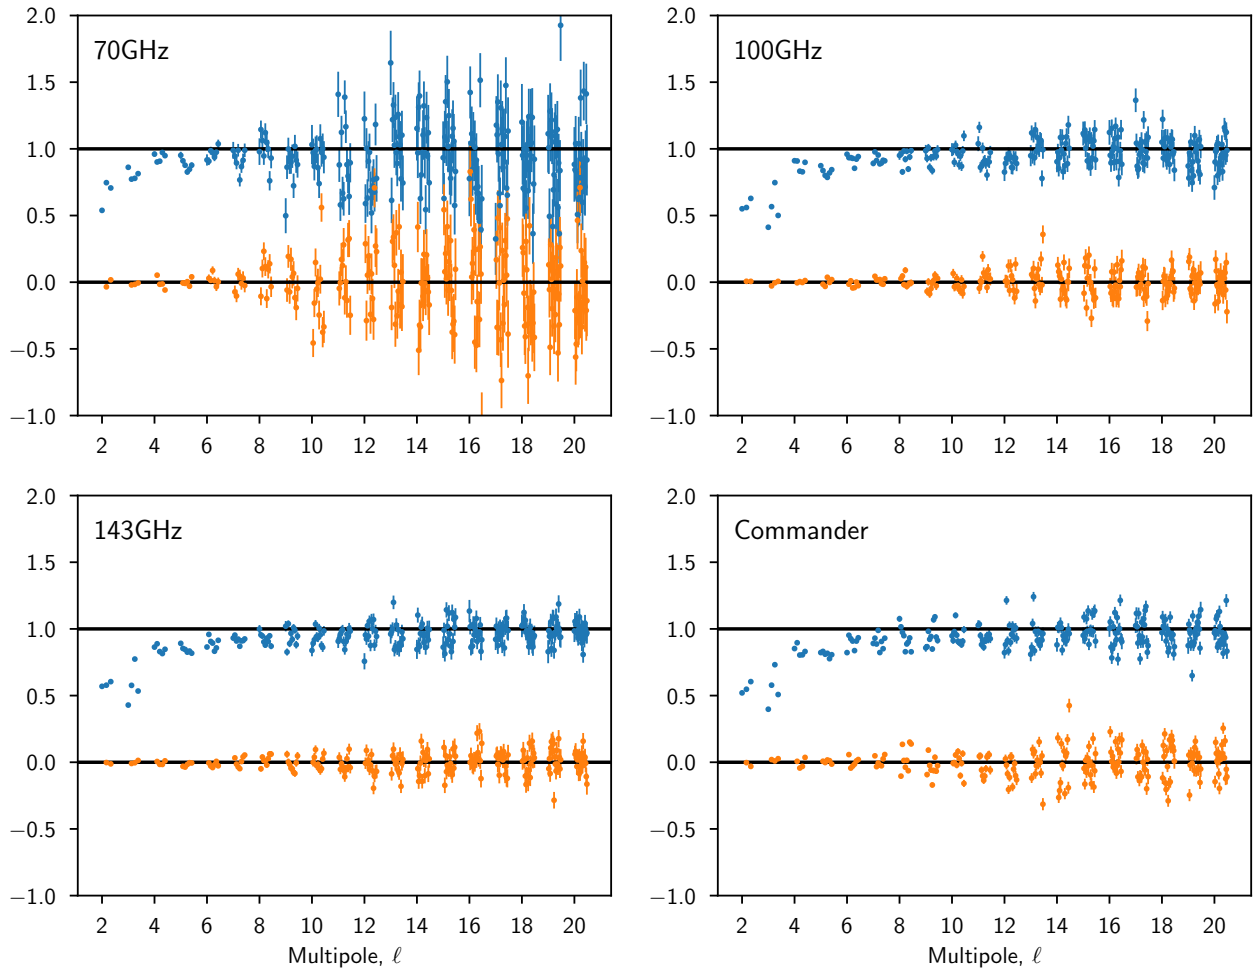


Fig. 24. Full anisotropic transfer functions for 70, 100, and 143 GHz and Commander. These demonstrate that the signal suppression is not completely uniform for the quadrupole and the octupole. The real part of the transfer function is shown in blue. The imaginary part (orange) of the transfer function is consistent with zero. These transfer functions correspond to minimizing Eq. (18) over 80% of the sky. The different m modes are slightly staggered here, with $m = 0 \dots \ell$ plotted left to right, to aid visualization.

Foreground experts may object to the use of foreground templates to model the polarized foregrounds in the calibration process. Such templates are inherently incapable of supporting spatial variation of the spectral index of the foregrounds. When these templates make up the polarization

prior, the calibration process does have the potential to suppress true spatial variation of the spectral index. We are confident that the limitations of our polarization prior have not compromised the foreground polarization for three reasons.




Fig. 25. Impact of the anisotropic NPIPE transfer function (Fig. 24) on a simulated CMB sky smoothed to 3° . The columns (left to right) are the input CMB, the transfer-function-convolved CMB, and the difference.




Fig. 26. Comparing the stable science scan (SCM) and the repointing manoeuvre (HCM) data. We find excellent consistency between the two disjoint data sets. The solid lines are auto spectra computed over 50% of the sky. The dashed lines are corresponding A/B cross-spectra. The spectra are binned into 100 logarithmically-spaced bins. The dashed orange spectrum in the EE panel demonstrates that even the 143 GHz repointing-manoeuvre data alone (which were excluded from all previous releases) are sensitive enough to probe at least the first three acoustic EE spectrum peaks. The agreement in small-scale TT cross-spectra indicates that the HCM pointing reconstruction is accurate enough not to widen the effective beam. The agreement in large-scale EE and BB power suggests that potential thermal effects from the thruster burns do not leak into large-scale polarization. The small scale noise in the HCM data set is higher than in the SCM data set from having only 8% of the integration time.

1. The calibration is performed using a Galactic mask, excluding the more intense Galactic plane and the source of more probable spatial variation of the spectral index.
2. Spatial variation in the dust spectral index has limited support by fitting both 217 and 353 GHz templates.
3. Our simulations that are based on the spatially-varying Commander sky model do not indicate any suppression of the spatially varying spectral indices (see Fig. 74).

4.4. Data taken during repointing manoeuvres

It is not immediately obvious that including data taken during the 4 min repointing manoeuvres is going to improve the resulting frequency maps and hence it is important to check for consistency. The three thruster burns of a given manoeuvre lead to a faint but measurable impact on the HFI thermal baseline, and the attitude reconstruction during the manoeuvre is admittedly less robust due to the dynamic nature of the data. We explored the consistency between “stable science mode” (labelled “SCM” in the *Planck* attitude history files) and the repointing mode (“HCM”) by building 143 GHz full-frequency and detector-set maps exclusively from either SCM or HCM data. We show the

power spectra of such maps in Fig. 26. These power spectra demonstrate excellent consistency between the two data subsets. Accordingly, data taken during pointing manoeuvres are used in all the NPIPE data products. *Only* SCM data were used in previous *Planck* products and results. This additional integration time reduces the small-scale noise uncertainty in NPIPE by approximately 9%.

5. Simulations

NPIPE is released with 600 high fidelity Monte Carlo simulations that include CMB, foreground, noise and systematics. These simulations allow testing and debiasing analysis tools with life-like frequency and detector-set maps with known inputs. For further details of the release, see Appendix A.

Calibrating against the same sky signal we are trying to measure is an inherently nonlinear problem and calibration and other template uncertainties in the NPIPE maps are not negligible. For these reasons we found it necessary to simulate the entire reprocessing part (Sect. 2.4) of the NPIPE processing. Each Monte Carlo iteration begins with a simulated CMB sky, represented as an expansion in spherical harmonics. We convolve

Table 6. Cosmological parameters of the FFP10 simulations.

Parameter	Value
$h^{(*)}$	0.6701904
$\omega_b^{(\dagger)}$	0.02216571
ω_c	0.1202944
τ	0.06018107
A_s	2.119631×10^{-9}
n_s	0.9636852
r	0.01

Notes. $(*)h$ is the dimensionless Hubble parameter: $H_0 = h \times 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$. $(\dagger)\omega \equiv \Omega h^2$.

the CMB with high-order expansions of the final *Planck* scanning beams ($\ell_{\text{max}} = 2048$ for LFI and 4096 for HFI). The convolution is carried out at the sample level (Prezeau & Reinecke 2010), respecting the actual beam orientation as a function of time. Foregrounds are simulated by evaluating the Commander sky model at the target frequencies. Static zodiacal emission is included by adding the same nuisance templates that Commander marginalized over. When components of the model are measured to much higher resolution than the target frequency (e.g., dust at 30 GHz), we smooth the foreground component with an azimuthally-symmetric beam expansion (Hivon et al. 2017) specifically calculated for the NPIPE data set. Bandpass mismatch is included by adding appropriately-weighted maps of foreground frequency derivatives and CO, customizing the foreground for each detector. The foreground map is then sampled into TOD, using the appropriate detector pointing weights, and co-added with the CMB TOD. Finally, we add the expected dipole and far-sidelobe signals. Noise is simulated from the measured noise PSDs as in Planck Collaboration XII (2016), including a correlated noise component in each polarized horn.

The simulated CMB sky includes lensing and frequency-dependent frame-boosting effects, as described in Planck Collaboration XII (2016). The CMB realizations are the FFP10 simulations used in PR3, and the cosmological parameters are listed in Table 6.

The gain fluctuations of several percent seen in LFI detectors are applied to the simulated TOD once all the components are added. Rather than trying to develop a statistical model that captures all the relevant features of the measured fluctuations, we apply the same smoothed version of the measured gain fluctuation to all Monte Carlo realizations.

HFI ADC nonlinearity is included into the simulations using the same tools as for LFI gain fluctuations, i.e., we apply a time-dependent gain fluctuation derived by smoothing the measured apparent gain fluctuations in the flight data. Similarly, we apply the measured transfer-function residuals to the simulated bolometer TOD.

5.1. Instrument noise

The $1/f$ instrument noise is simulated using the inverse fast Fourier transform (FFT) technique. A vector of complex Gaussian random numbers is generated, multiplied by a measured power spectral density (PSD), and then Fourier transformed into the time domain. We treat each pointing period as independent, enforcing no continuity across the pointing period boundaries.

Our noise estimation uses the same technique as described in Planck Collaboration XII (2016), i.e., we subtract an estimate of

the signal by resampling the full-frequency map into the time domain, evaluate the correlation function over unflagged and unmasked samples, and Fourier transform to obtain the PSD. The low-frequency part of the PSD is measured from a down-sampled version of the timestream to lessen the cost of estimating the covariance function. The traditional method of estimating the PSD directly from the TOD by FFT scales as $O(N \log N)$, but is forced to use all of the timestream samples, even the ones that contain a constrained realization of the data (e.g., to fill gaps) or are subject to steep gradients in the sky signal. The covariance function estimation used in NPIPE scales as $O(N^2)$, but allows for omission of the flagged samples and use of only the part of the sky where the signal estimate is robust.

The noise estimates are derived from preprocessed TOD to best match the noise seen during reprocessing. We apply the measured LFI gains during the estimation to remove the several-percent gain fluctuations from the noise estimates.

We have further developed the noise estimation code to use the covariance function method for measuring the correlated noise between detectors. Of particular interest is the correlated noise within each polarized horn (between the two polarization-orthogonal receivers). The correlated noise affects the sum of the two receivers, and thus the estimation of sky temperature; however, it cancels in the difference of the two receivers, leaving the polarization estimate unaffected. The noise estimates in Fig. 27 demonstrate that all polarized horns show significant internal noise correlations at low frequency, but only the HFI horns have significant noise correlations at high frequency. The HFI noise correlations are thought to result from coincident cosmic-ray glitches. The correlated HFI modes also exhibit several 4 K line residuals, which is understandable, given that the changes in the measured line-template amplitudes are highly correlated.

We simulate the correlated modes using the same inverse FFT technique as used for the uncorrelated modes, and then co-add them to the uncorrelated timestreams. Only the intra-horn component is simulated. We ignore the smaller and less relevant correlation between horns.

5.2. What is not included in the simulations

NPIPE simulations make an effort to capture all of the relevant systematics and their coupling in the final maps; however, the prohibitive cost of preprocessing precludes running the preprocessing module on simulated timelines. This leads to some notable omissions:

- ADC nonlinearity residuals beyond the linear gain fluctuation model (the higher-order distortion fluctuations are shown to be small in Fig. 11);
- 1 Hz and 4 K line residuals, except for what is captured by noise estimation;
- cosmic-ray glitch residuals, except for what is captured by noise estimation and modelled as Gaussian noise;
- measuring and updating the detector polarization parameters; and
- errors in the bandpass-mismatch templates, i.e., the sky model and the CO maps that are solved with Commander and fed back in future iterations of NPIPE.

Despite these omissions, we take the general agreement between flight data and simulated residuals in Figs. 28–31 as a validation of our approach.

HFI ADC nonlinearities are simulated in NPIPE using the linear gain model. This notably excludes the higher-order distortion modes discussed in Sect. 2.3.3. These modes are very faint and the fits include a fair amount of statistical uncertainty.




Fig. 27. Averaged noise PSDs for each detector (upper curves) and correlated-noise modes for each polarized horn (lower curves). The total noise power is the sum of the correlated and uncorrelated modes. These noise PSDs are measured from the data by subtracting a signal estimate and then evaluating the sample-sample covariance function. The HFI noise is suppressed near the Nyquist frequency (≈ 90 Hz) by the bolometric transfer function filtering. The PSDs are used for simulating the $1/f$ noise fluctuations, and as inputs to the Madam noise filter for destriping.

Without these higher modes of distortion, the simulated maps are affected only by the statistical uncertainty associated with the distortion correction. The excellent agreement between the large-scale polarization residuals in flight data and simulations suggests that the statistical uncertainties are sufficient to describe the large-scale polarization uncertainties in the NPIPE maps.

The line and glitch residuals would manifest themselves as small-scale excess power in the NPIPE maps. We only model these residuals to the extent that they can be measured in the detector noise PSD. The binned PSD cannot fully capture the Poissonian nature of the glitch residuals, nor the finely localized features associated with the lines. However, since the small-scale power of our simulated maps agrees very well with the real maps, we deduce that the mismatch is negligible.

Updating the instrument model using the single-detector residual maps and polarization estimates has not been formally included in the NPIPE processing plan, but has been done incrementally as post-processing over multiple revisions of the NPIPE maps. Simulating this process was deemed unfeasible given the time constraints. NPIPE simulations are therefore based on the final instrument model, including the adjustments to polarization efficiency and angle determined from flight data processing and shown in Table 4.

NPIPE HFI simulations differ from the FFP10 simulations used in Planck Collaboration III (2020) in some important ways.

FFP10 implemented a light-weight version of the low-level data processing (without glitch or 4 K line removal) allowing for application of the ADCNL directly in the fast, modulated samples. It also allowed simulating the information loss from compression and decompression. A noteworthy complication of this approach was that the instrumental noise also needed to be simulated in the raw data, necessitating an extra noise-alignment step later in the pipeline. Analysis of the FFP10 results showed that the ADCNL was by far the most important systematic included in the simulation. NPIPE uses a more approximate approach, simulating the data directly in their preprocessed form based on noise estimates from the same stage data and adding ADCNL consistent with the linear gain model. Despite their differences, most effects do get included in both pipelines but using very different ways to express them. Our simulations are the first *Planck* simulations to apply the asymmetric scanning beam to each CMB realization separately in the time domain. In FFP10 only one beam-convolved CMB realization was passed through the full processing pipeline.

5.3. NPIPE simulation results

We compare the simulated CMB power to the simulated residuals in Figs. 32 and 33. They demonstrate that we are able to suppress large-scale noise and systematic uncertainties at or below

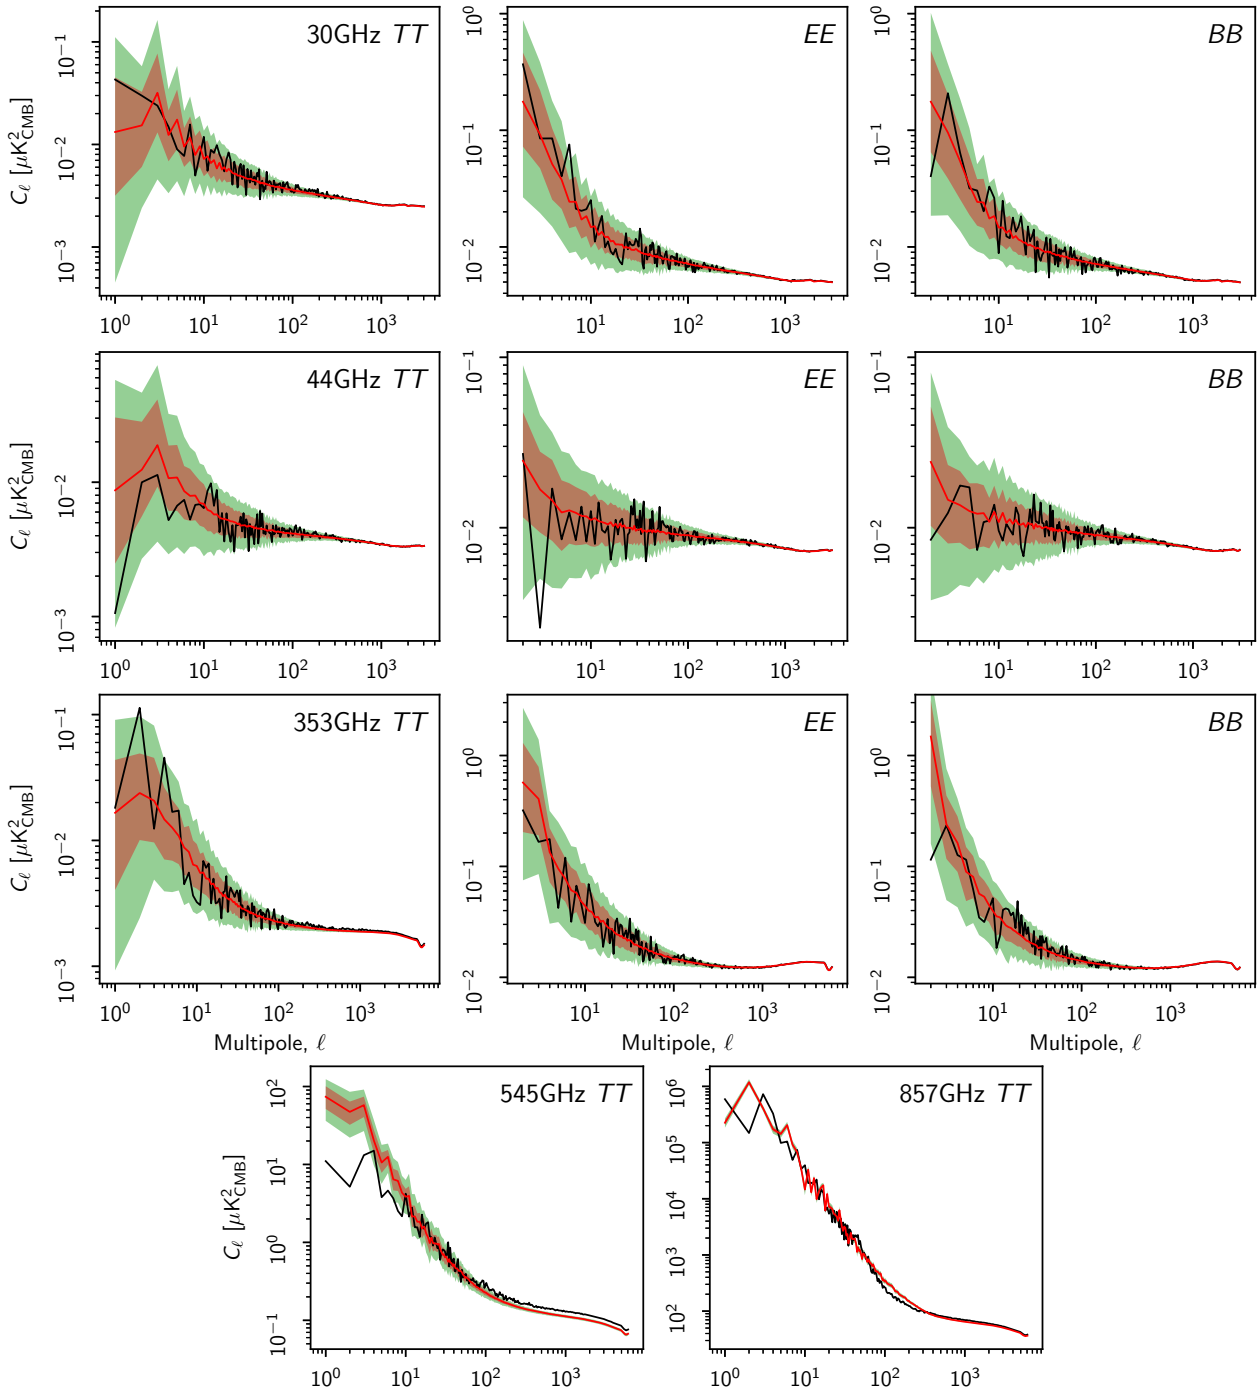


Fig. 28. Simulated A/B difference versus flight data A/B difference at 30, 44, 353, 545, and 857 GHz. The flight data are shown in black, and the median simulated power in each bin is plotted in red. The coloured bands represent the asymmetric 68% and 95% confidence regions. The power spectra are binned into 300 logarithmically spaced bins. These spectra are shown again in Fig. 34, but divided by the median simulated spectrum.

the faint large-scale EE power at the *Planck* CMB frequencies between 70 and 217 GHz.

5.3.1. Null maps and null map power spectra

We can demonstrate that the simulations succeed in capturing the relevant residuals by comparing real and simulated detector-set-difference power spectra in Figs. 28, 29, 34, and 35 (for TE , TB , and EB spectra, which are expected to be null, see Figs. 30 and 31). The detector-set differences exhibit excellent agreement between the flight data and the simulated detector-set

differences in polarization. The temperature difference is not as well matched at the largest scales, owing to three factors. First, at 143 and 217 GHz the solved gain fluctuations seem to converge to a slightly different large-scale temperature sky for the two detector sets. The resulting roughly $1 \mu K_{\text{CMB}}$ mismatch is not fully captured in the simulations, suggesting that it is driven by a systematic that is not simulated. This residual and the associated mismatch are roughly 100 times fainter than the simulated CMB anisotropy at the same angular scales.

Second, at 545 GHz the simulations suggest almost an order of magnitude more large-scale temperature mismatch than is




Fig. 29. Same as Fig. 28, at 70, 100, 143, and 217 GHz. These spectra are shown again in Fig. 35 but divided by the median simulated spectrum.

found in the flight data. This is explained by the crude bandpass-mismatch model (different delta-function bandpasses), coupled with the lack of redundancy at 545 GHz, where the “A” detector set comprises only a single bolometer. Even the amplified detector-set mismatch is 1000 times fainter than the dust in the 545 GHz band.

Third, at 857 GHz the dominant failure mode is the mismatch at intermediate scales between $\ell = 100$ and $\ell = 1000$. This mismatch is not seen in the half-ring differences, so it is not associated with the detector noise, but rather with systematic residuals related to the bandpass mismatch and bolometer transfer function. In the power-spectrum domain, this mismatch is

100–10 000 times fainter than the sky signal over the 50% of sky with the lowest dust emission. The 857 GHz detector set only consist of two SWBs.

5.3.2. Noise alignment

The detector-set difference power spectra in Figs. 34 and 35 testify to a good overall agreement between the simulations and the flight data. Upon closer inspection it is possible to identify a percent level, scale-dependent deficit in simulated noise power at $\ell > 100$ affecting the polarized HFI frequencies (100–353 GHz). The deficit is exclusive to the HFI simulations, suggesting that its

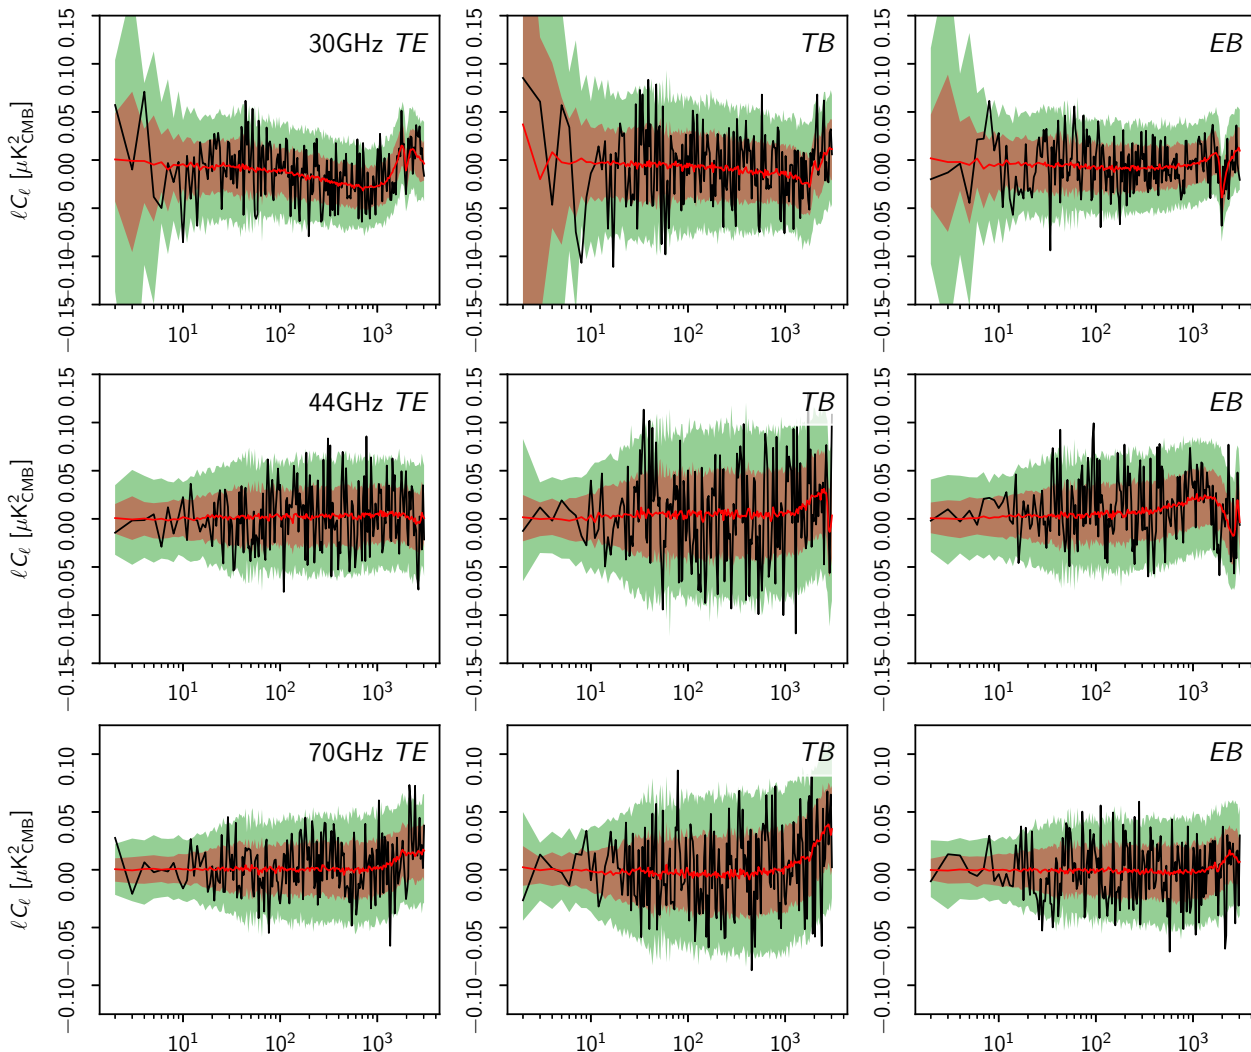


Fig. 30. Simulated A/B difference versus flight data A/B difference for *TE*, *TB*, and *EB* at 30, 44, and 70 GHz. Similar to Fig. 28, except for the ℓ -scaling.

origins lie in HFI specific phenomena, such as Poissonian glitch residuals that are not included in the simulations and could not be reproduced by a Gaussian noise simulation.

Overlooking the noise mismatch implies up to a percent level bias in statistical error measures that are derived from the simulations. For some analyses this degree of uncertainty is acceptable. For estimators that are sensitive to the mismatch in total power between the flight data and simulations, we offer a set of additional, 100–353 GHz, small-scale noise maps that can be added to the simulated NPIPE maps to align the total noise power. We stress that the additional noise maps only adjust the overall noise power, without any effort to match the spatial structure of these missing noise modes. We refer to this as “noise alignment”. We demonstrate the effect of the noise-alignment maps in Fig. 36.

6. NPIPE and the earlier *Planck* releases

NPIPE differs from the earlier *Planck* releases in several ways. We compare NPIPE to both PR2 and PR3, allowing the reader to assess the magnitude of the changes against the 2015/2018 differences, which are themselves substantial. In the 2018 release, the LFI team made a critical change in calibration procedure, taking into account foreground polarization and thus removing

a source of bias that generated spurious polarization. Also in 2018, HFI adopted a new mapmaking and TOD cleaning method called SR011, which included time-domain corrections for band-pass mismatch, far sidelobes, and bolometer transfer-function residuals.

6.1. Calibration

NPIPE photometric calibration is measured by fitting the orbital dipole template (Sect. 2.4.3) to the data. The overall calibration is based solely on matching the far-sidelobe-corrected orbital dipole template against the data, rather than fitting a combination of the orbital dipole and some prior estimate of the total dipole. Our orbital dipole template includes the faint, frequency-dependent, dipole-induced quadrupole.

We follow the LFI approach, correcting the orbital dipole template for subtle suppression and phase shifts induced by the far sidelobes, and extend this treatment to the HFI frequencies. Details of the convolution are presented in Appendix C.

NPIPE processing effectively deconvolves the far sidelobes from the data, which are scaled to have unit response to the main beam rather than the 4π beam. This is reflected in the effective beam window functions, which have unit response to the

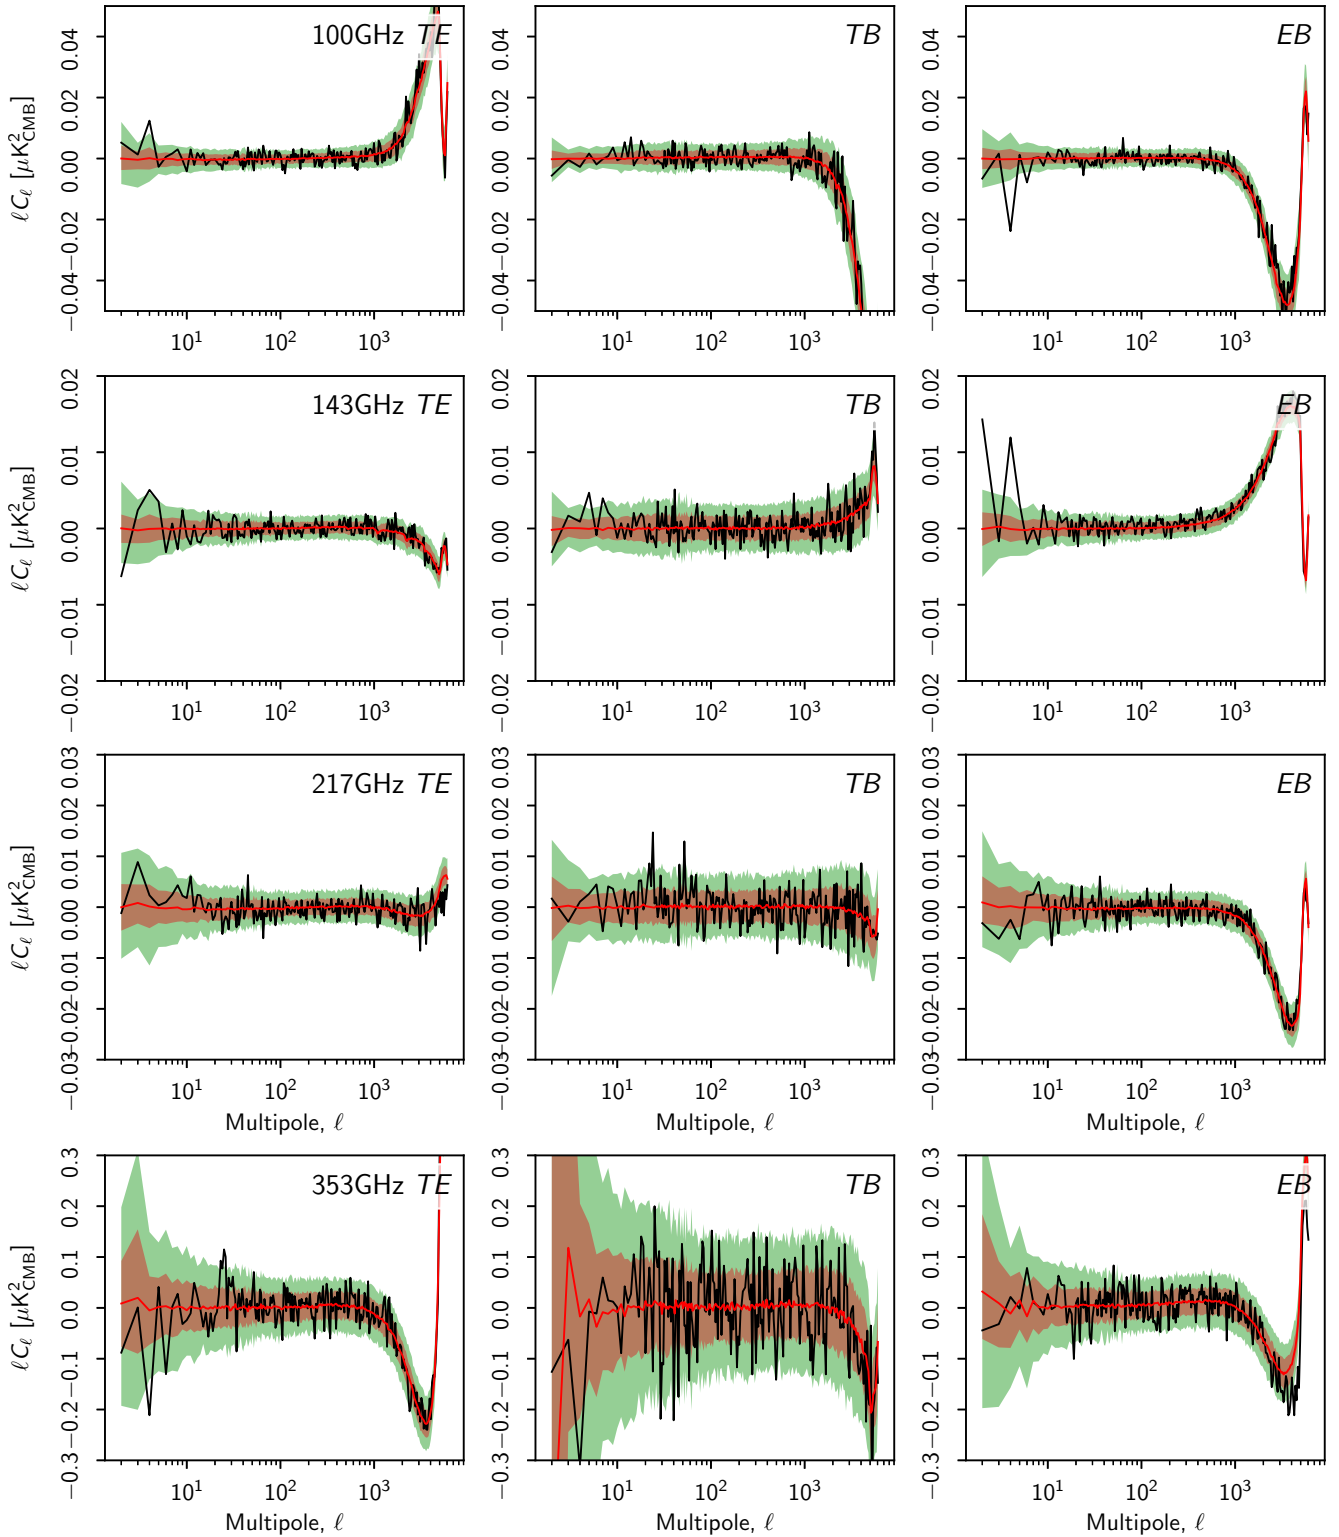


Fig. 31. Same as Fig. 30, for 100, 143, 217, and 353 GHz.

monopole and near unit response to the dipole. The uncorrected main-beam efficiencies are shown in Fig. 37.

We measure the calibration uncertainty and potential bias from template degeneracy, using the full-frequency simulated maps. To measure the calibration of each simulated map, we subtract the input CMB and foreground map, mask out 50% of the sky most affected by foreground residuals, and regress the input

dipole template against the cleaned map. The resulting overall gain distributions are shown and compared to PR3 calibration in Table 7. The simulations indicate that we recover the input dipole at better than $1\mu\text{K}$ (combined statistical and systematic uncertainty) in each of the CMB frequencies. The table also shows that there are small but statistically significant biases in the procedure, which cause the HFI CMB frequencies to be calibrated

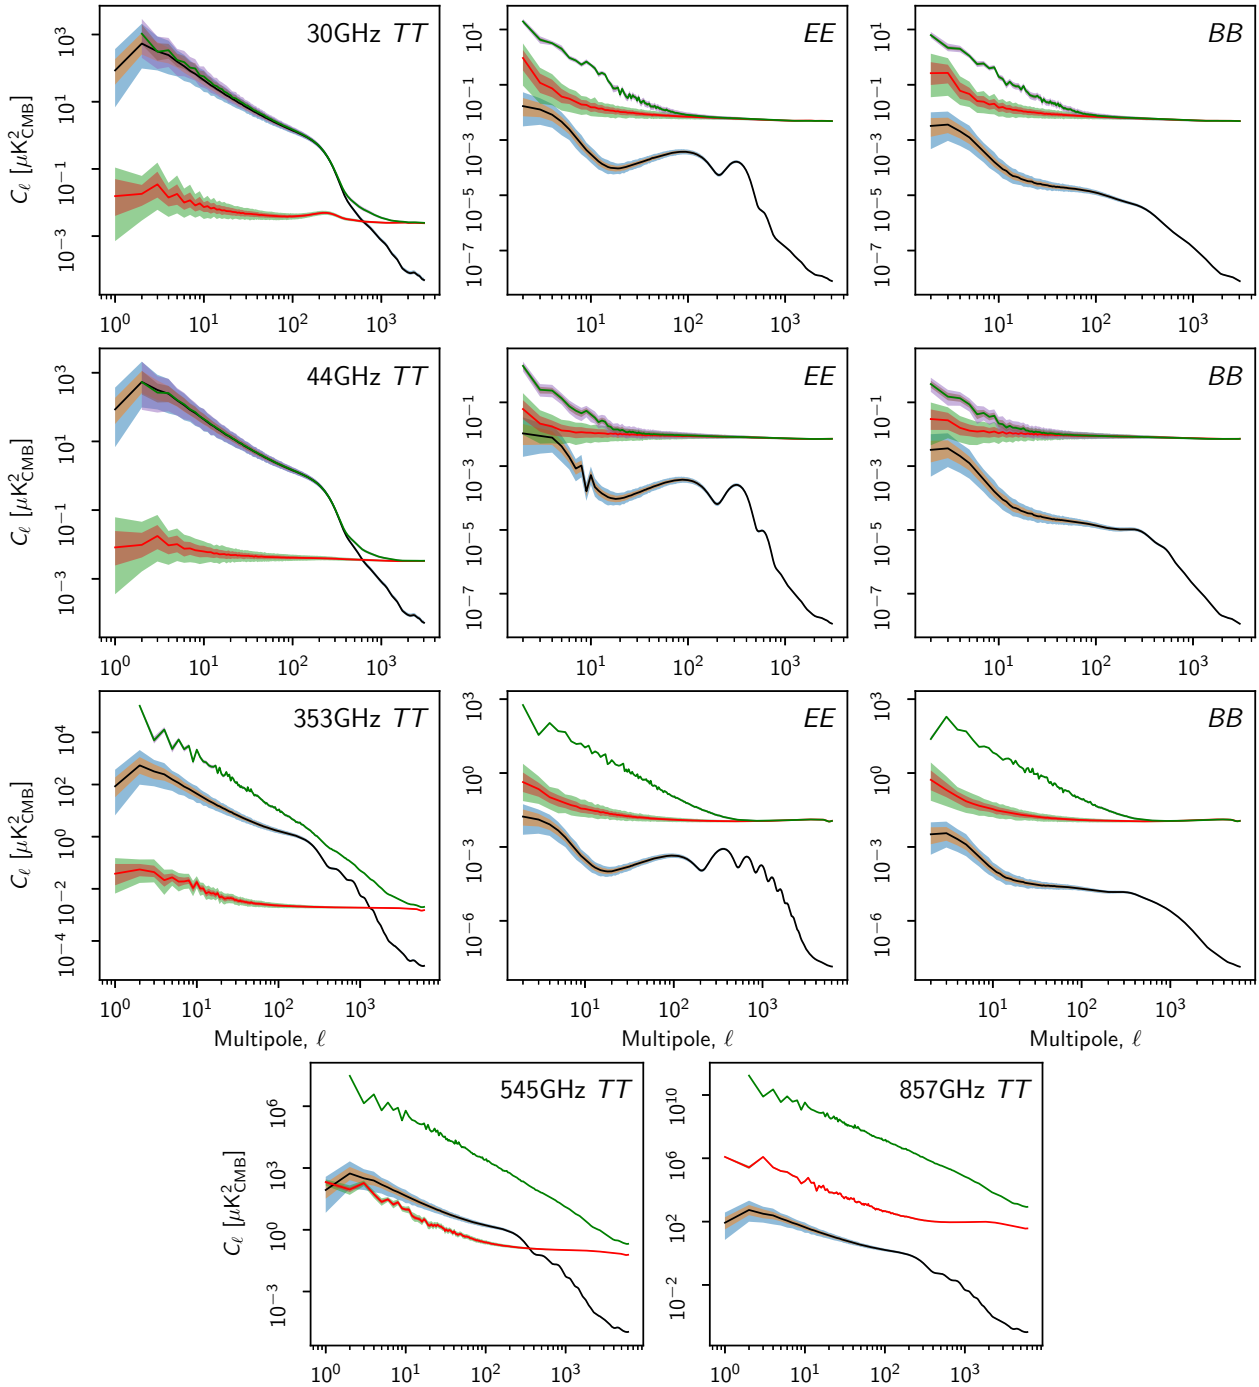


Fig. 32. Simulated CMB, noise, and systematics pseudo-spectra at 30, 44, 353, 545, and 857 GHz. The median CMB power in each bin is plotted in black, and the median noise and systematics power (difference between the input and output maps) in red. The coloured bands represent the asymmetric 68% and 95% confidence regions. The power spectra are binned into 300 logarithmically-spaced bins. The CMB is convolved with the beam and E -mode transfer function (Sect. 4.3). The large-scale CMB B -mode power in these plots follows from mode coupling and is not intrinsic to the simulations.

high by less than 0.03%, amplifying the Solar dipole amplitude by less than $1 \mu\text{K}$.

6.2. Noise and systematics

We measure the level of residual noise and systematics in the *Planck* maps by splitting the available detectors into disjoint sets and performing the same processing on each detector set. The independent processing guarantees that the residuals between the

detector sets are not correlated. Unlike making a split in the time domain (the so-called half-mission split), the detector-set difference does not cancel systematics that vary detector by detector, such as bandpass mismatch, far sidelobes, or the HFI transfer function residuals.

The definition of the detector-set null test has been somewhat ambiguous in past *Planck* publications. In an independent, local processing, the template corrections (calibration, transfer function residuals, etc.) and noise offsets are fitted to each detector set

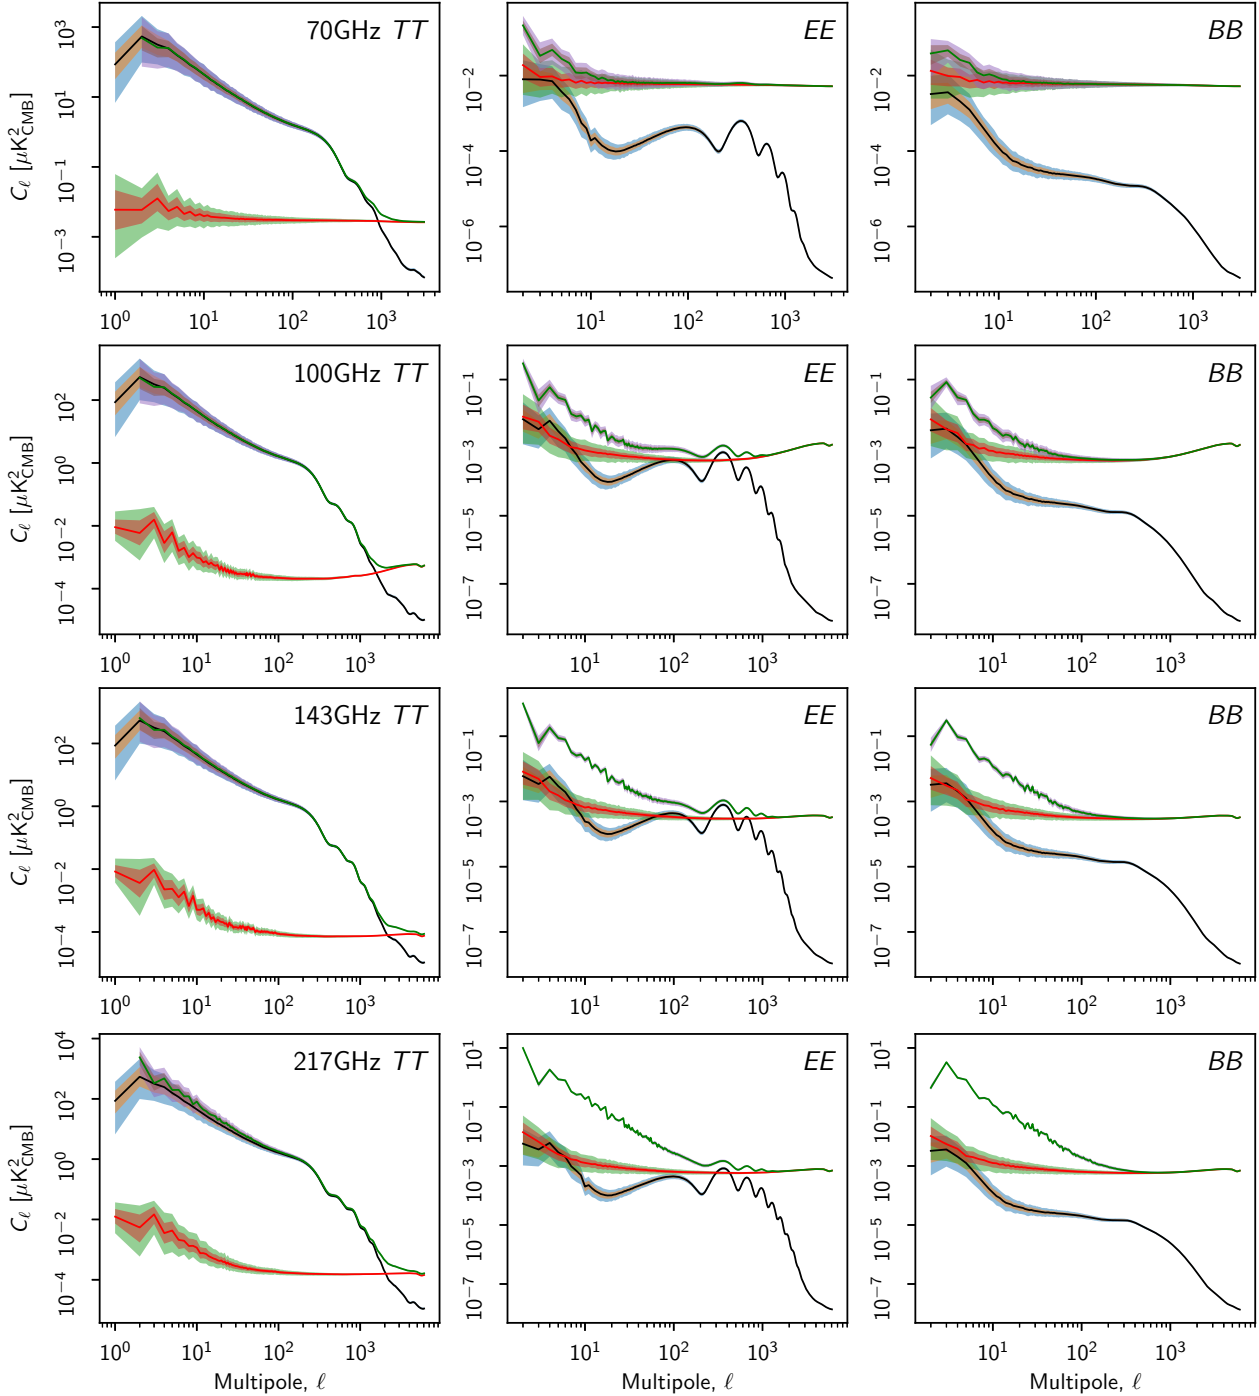


Fig. 33. Same as Fig. 32, for 70, 100, 143 and 217 GHz.

separately. Alternatively, it is possible to fit the templates globally to all of the data and project subsets of the cleaned data into detector-set and half mission maps. The global fit (sometimes referred to as “turbo destripping” because it only requires one run of the destriper) is suitable for measuring the achieved internal consistency in the cleaned TOD. However, residual systematics in turbo-destripped subsets are highly correlated, making it impossible to infer the total level of residuals from their differences. In this paper, all turbo-destripped subset maps are presented with an explicit mention of the global fit being used.

We compare the EE and BB total power in each polarized *Planck* frequency in Fig. 38. In Figs. 39 and 40 we study the

detector-set null maps for the polarized HFI frequencies. These maps and spectra indicate that NPIPE products have less noise than PR3 maps. There are several factors that contribute. The NPIPE maps include data taken during repointing manoeuvres, which account for roughly 8% of the *Planck* mission. LFI noise levels are reduced further by the adaptive sky-load differencing (Sect. 2.3.6). HFI noise levels are improved by the extended flagging kernel during transfer-function deconvolution (Sect. 2.3.9). Large-scale HFI residuals, particularly in polarization, are significantly suppressed by our use of the polarized-foreground prior (Sect. 2.4.13) during calibration, as is demonstrated by the increase in NPIPE large-scale polarization uncertainty when

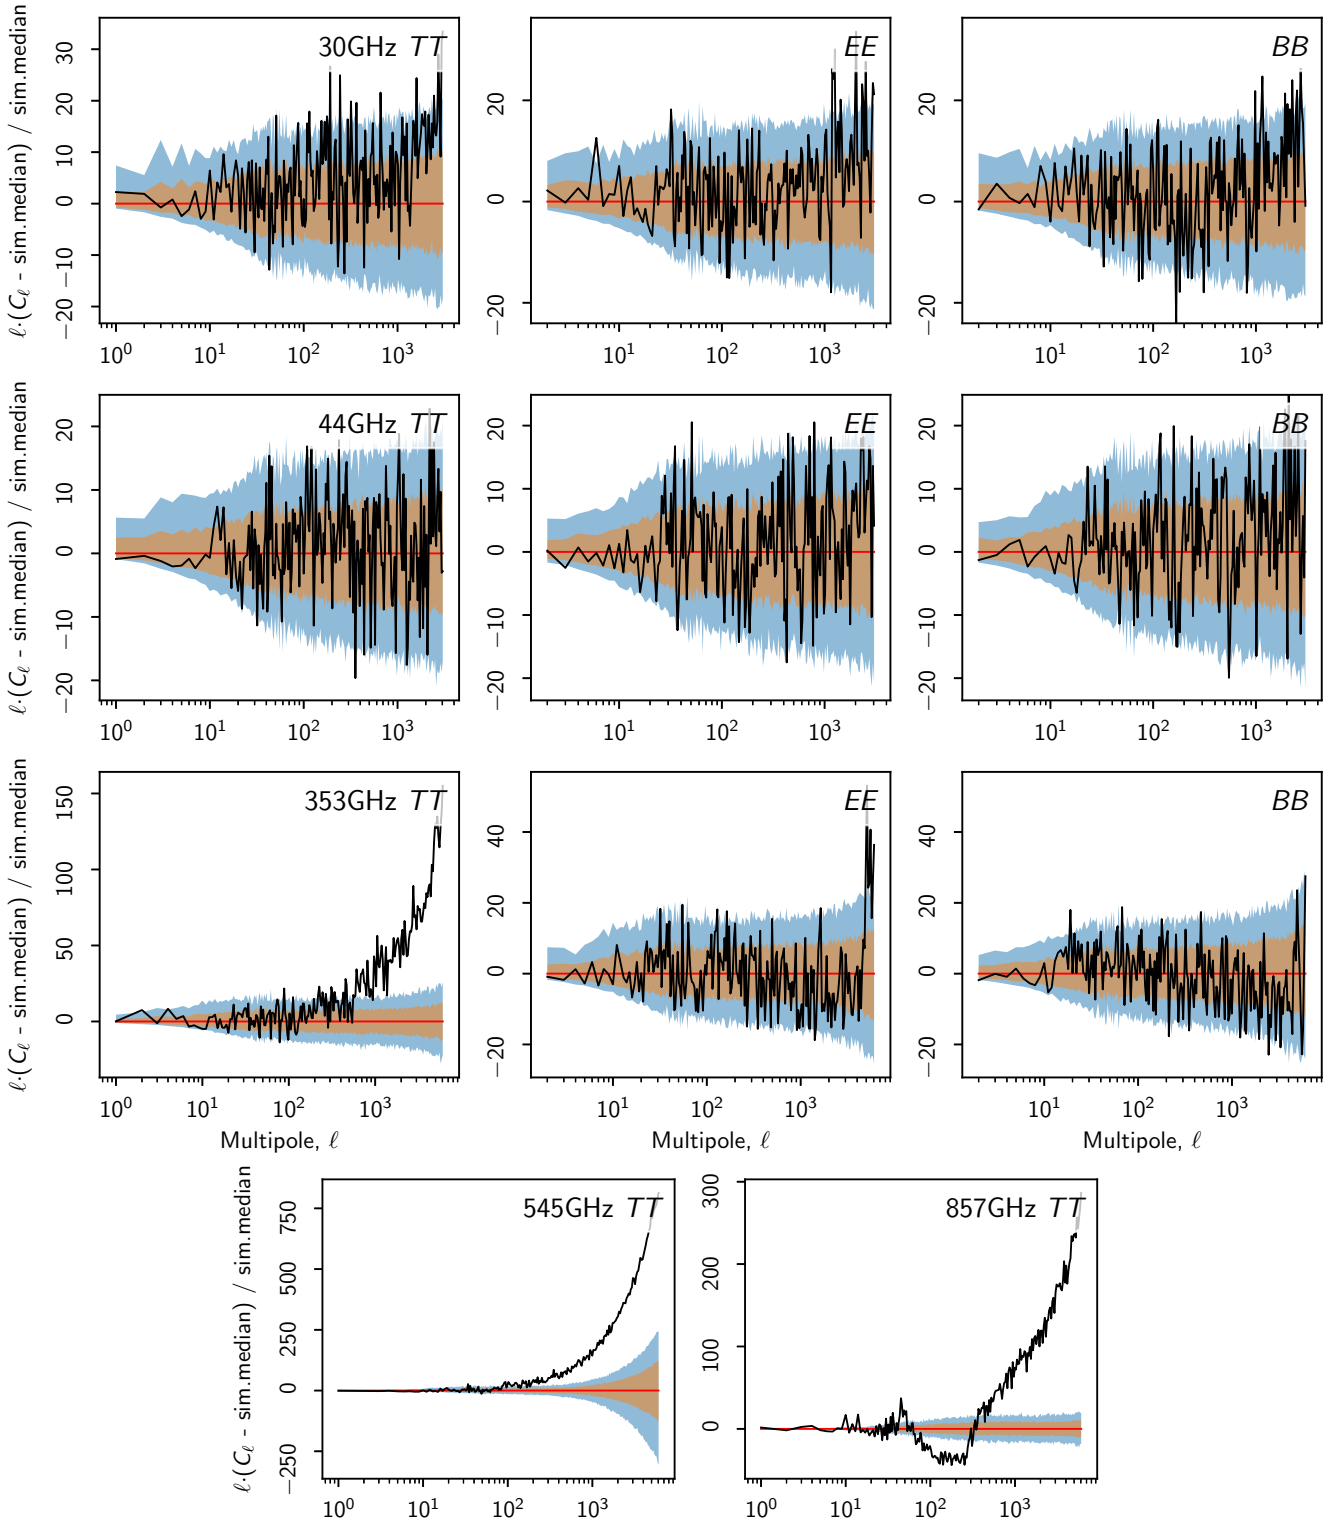


Fig. 34. Simulated A/B difference versus flight data at 30, 44, 353, 545, and 857 GHz. All spectra are differenced and divided by the median simulated spectrum. The flight data are shown in black, and the median simulated power in each bin is plotted in red. The coloured bands represent the asymmetric 68% and 95% confidence regions. The power spectra are binned into 300 logarithmically-spaced bins. The spectra shown here are the same as in Fig. 28. Notice the ℓ -scaling.

the prior is disabled (Fig. 41). These improvements hold even after deconvolution of the E -mode transfer function (Sect. 4.3). Finally, we fit the $1/f$ noise fluctuations with 167 ms baseline

steps, compared to the 0.25 s (30 GHz), 1 s (44 and 70 GHz), and full-pointing-period (35–75 min at HFI frequencies) baseline steps used in PR3.

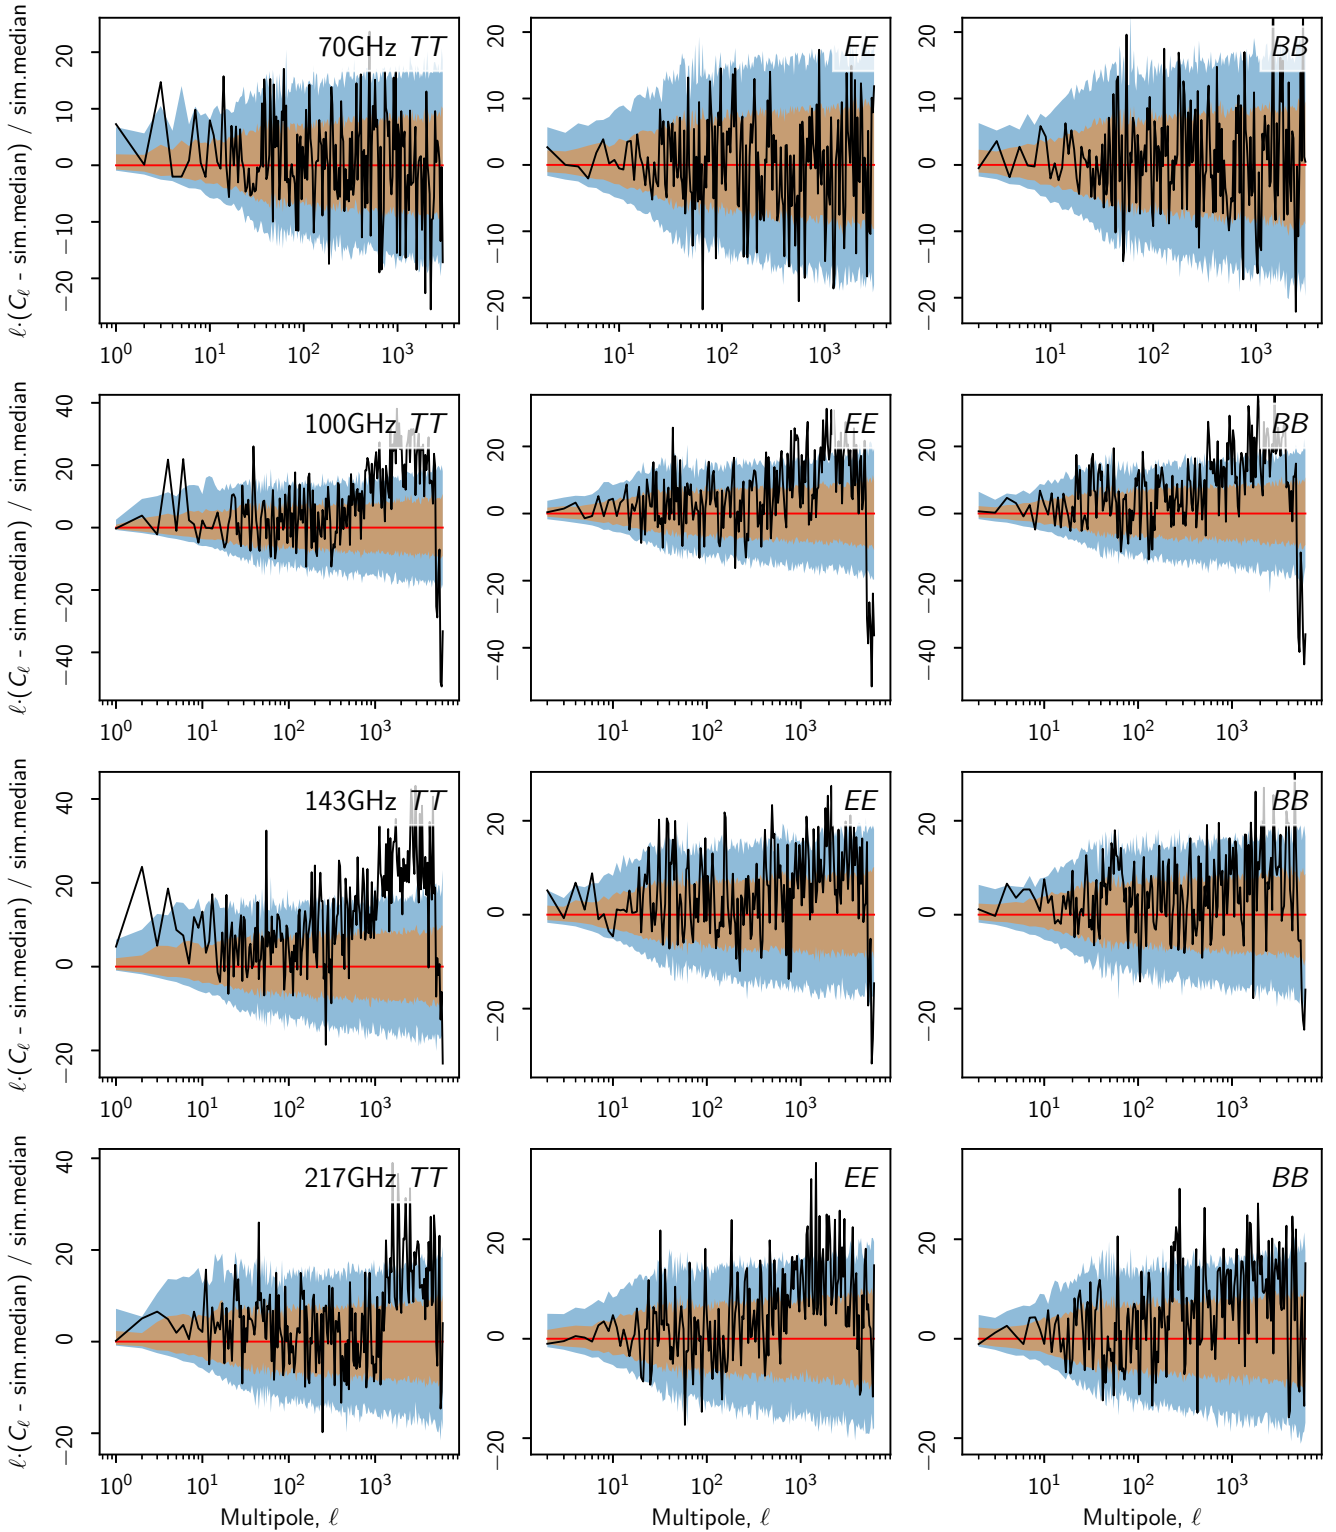


Fig. 35. Same as Fig. 34, but at 70, 100, 143, and 217 GHz. The spectra shown here are the same as in Fig. 29.

6.3. Signal

We present the temperature and polarization difference maps between NPIPE and PR3 in Fig. 42. To make the temperature map differences more informative, we have projected out the relativistic dipole, zodiacal emission, and an overall relative calibration mismatch. The 353 GHz temperature difference has a peculiar ringing pattern, understood to originate from the coarse-grained

transfer function model used in the 2018 processing. For comparison, we show select NPIPE–2015 differences in Fig. 43.

While our work was being prepared for publication, a new processing of *Planck*-HFI polarization data, known as SR0112, was published in Delouis et al. (2019). It features a new model of the ADCNL that does not rely on the linear gain fluctuation approximation. The new ADCNL model is able to capture the bulk of the residual ADCNL with considerable economy, thus

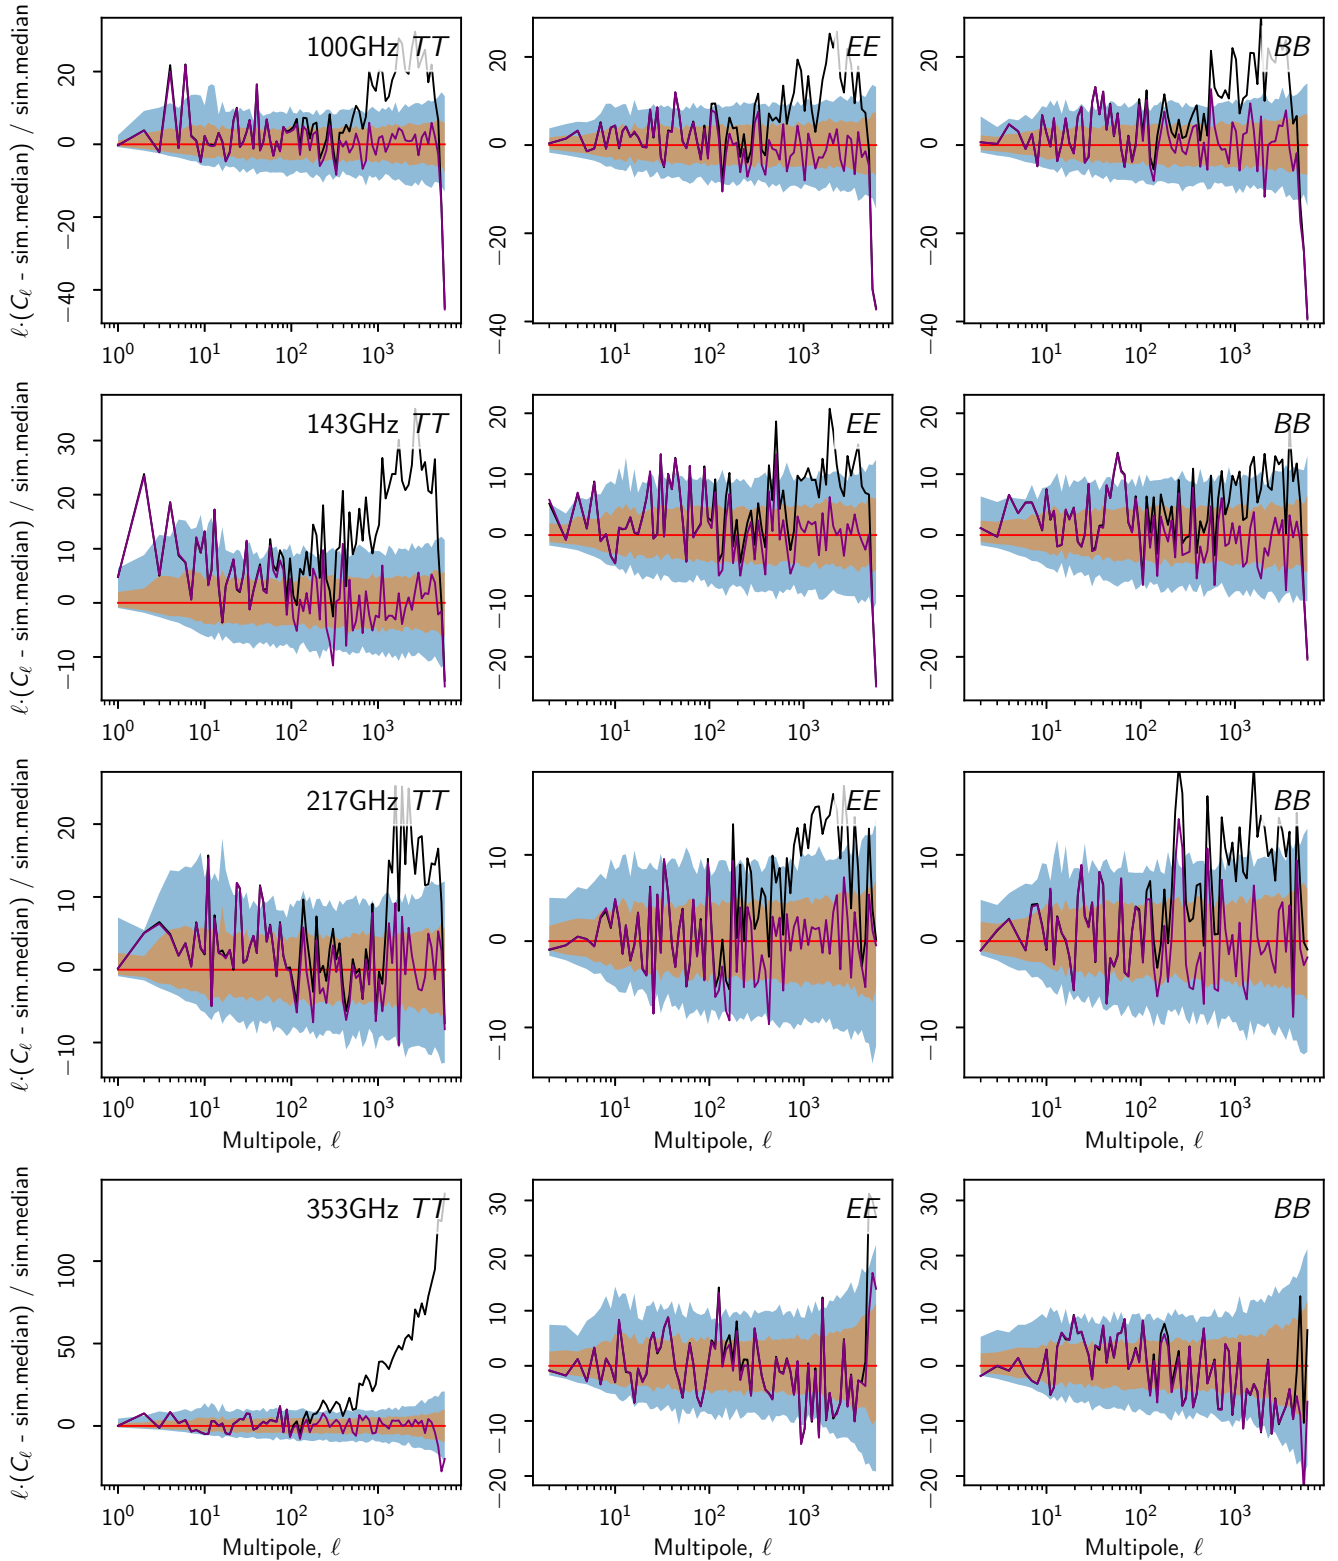


Fig. 36. Simulated A/B difference versus flight data at the HFI frequencies where we identify a percent-level, scale-dependent deficit in simulated noise power. The flight data power spectrum is differenced and divided by the median of the uncorrected simulations (black) and, alternatively, by the simulated maps that include the noise correction (purple). The coloured bands represent the asymmetric 68% and 95% confidence regions. The power spectra are binned into 100 logarithmically spaced bins. Apart from the wider binning, the black curves and the confidence limits are the same as those shown in Figs. 28 and 29.

limiting the degeneracy between the *IQU* map and the ADCNL correction. We show in Fig. 44 the NPIPE–SR0112 difference maps, and observe that for polarization, the NPIPE and SR0112

maps are considerably closer to each other than NPIPE and PR3, except at 143 GHz. But even the NPIPE–PR3 differences are very small. In temperature, maps from 100 to 217 GHz have

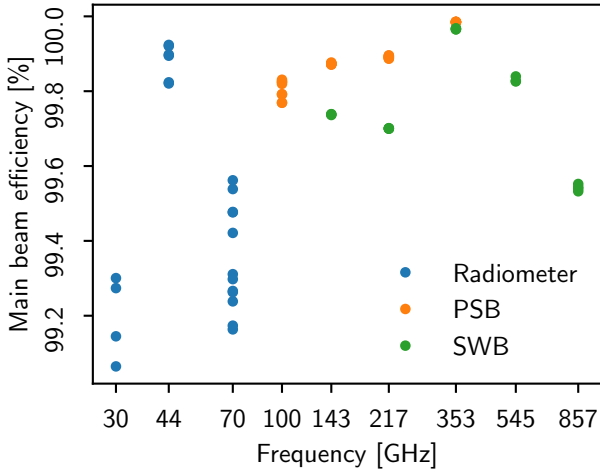


Fig. 37. Uncorrected main-beam efficiencies for all *Planck* detectors. The NPIPE calibration procedure corrects each frequency map to have full main-beam efficiency. The 545 and 857 GHz efficiencies are less certain due to uncertainties in main beam and sidelobe estimates, and are based on fitting the 353 GHz sidelobe templates to the data. The 100–353 GHz sidelobe estimates are based on first-order GRASP simulations, and overestimate the main-beam efficiency by a small but unknown amount.

Table 7. Relative calibration between NPIPE and PR3 maps and the NPIPE gain uncertainty estimated from recovery of the injected dipole in simulated full-frequency maps.

Frequency [GHz]	2018 gain ^(a) [%]	MC gain ^(b) [%]	MC bias ^(c) [%]
30	99.793	99.975 ± 0.045	−0.025 ± 0.005
44	99.763	99.990 ± 0.040	−0.010 ± 0.004
70	99.959	99.997 ± 0.030	−0.003 ± 0.003
100	100.061	100.016 ± 0.023	0.016 ± 0.002
143	100.082	100.016 ± 0.018	0.016 ± 0.002
217	100.005	100.024 ± 0.024	0.024 ± 0.002
353	100.463	100.052 ± 0.053	0.052 ± 0.005
545 ^(d)	115.123	99.586 ± 0.571	0.414 ± 0.057

Notes. ^(a)Relative calibration, NPIPE/2018, measured over 50% of the sky after adding the Solar dipole back into 2018 maps and smoothing the maps with a 1° Gaussian beam. ^(b)Distribution centre and width, measured relative to the input signal. ^(c)This is the average calibration error and the associated uncertainty. A 0.03% bias in calibration corresponds to about $1 \mu\text{K}_{\text{CMB}}$ in Solar dipole amplitude. ^(d)The NPIPE 545 GHz calibration is based on the orbital dipole, and does not suffer from the 10% modelling uncertainty affecting the planet-based calibration in the 2018 release. We also calibrate directly into K_{CMB} , without needing a unit conversion factor with additional uncertainty.

diverged in the Galactic plane, owing to the exclusion of all SWBs from SR0112 processing. The 353 GHz zebra stripe pattern is gone for both NPIPE and SR0112.

6.3.1. Internal consistency

Each *Planck* operational year divides into two sky surveys. The two surveys are complementary in the sense that most sky regions are scanned in approximately opposite directions between the two surveys. This makes the difference between odd and even surveys sensitive to a number of systematic effects, such as pointing error, far sidelobes, and transfer-function residuals. Processing the odd and even surveys into separate maps

would be possible, but not representative of the full-map uncertainty, since the opposite scanning directions are essential in fitting the systematic templates. Instead, we can project the calibrated and template-corrected data into odd and even survey maps, and use their difference to gauge the level of internal consistency achieved. This test reveals if the considered templates have enough freedom to model the systematics that are responsible for the odd–even survey differences.

We compare LFI odd–even survey differences between PR2, PR3, and NPIPE in Fig. 45. This figure shows a clear improvement in internal consistency between the 2015 and 2018 maps, owing to the improvements in calibration. The NPIPE survey differences show further improvement, likely due to the full-frequency calibration scheme that corrects for integrated band-pass mismatches.

We compare the 100–217 GHz survey differences in Fig. 46. These figures contrast the apparent lack of large-scale structure in the NPIPE null map with the visible residuals in the earlier releases. The improvement is due to the polarization prior, short-baseline destriping, and additional bins in the transfer function correction. It is worth pointing out that the NPIPE 143 GHz map has a noticeable residual in the Galactic plane, possibly from the transfer-function correction that down-weights these pixels when fitting the templates. The level of the residual is negligible given the strength of the foreground signal in these pixels. The 217 GHz 2018 map outperforms the 2015 results due to improvements in FSL and zodiacal-light removal.

In Fig. 47 we extend the comparison from 353 to 857 GHz. The 353 GHz results are particularly notable because they show an apparent degradation in the null map between the 2015 and 2018 releases, visible in a pattern referred to as “zebra stripes” in *Planck Collaboration III* (2020). This is a result of insufficient granularity in the measured and applied transfer-function correction. The regression is corrected in NPIPE by increasing the number of spectral bins.

In Fig. 48 we consider polarization difference maps of the form m_{217} to $0.128 m_{353}$ as derived by SR011, SR012, and NPIPE, where the scaling factor corresponds to a modified-blackbody spectral energy distribution (SED) ratio between 217 and 353 GHz evaluated for $\beta_d = 1.6$ and $T_d = 19.6$ K. These maps are thus designed to suppress thermal dust emission without having to resort to component-separation methods, and highlight potential residual systematic effects in the high-frequency channels. Focusing first on high Galactic latitudes, we first note that all three map solutions exhibit notable large-scale fluctuations. Some fluctuations are expected simply from true CMB fluctuations (which are only suppressed by 12.8% in these difference maps), and some fluctuations are expected from random statistical noise. However, some fluctuations are also likely due to residual instrumental systematics. For comparison, we note that a standard Λ CDM CMB sky with $\tau \approx 0.06$ exhibits peak-to-peak fluctuations of about $0.5 \mu\text{K}$, with an overall standard deviation of $0.15 \mu\text{K}$. At low Galactic latitudes, the residuals are dominated by possible spatial variations in the spectral index not captured by the constant scaling factor of 0.128 and temperature-to-polarization leakage. We note that NPIPE exhibits notably lower fluctuations at high latitudes than both SR011 and SR012, and Galactic-plane residuals appear more strongly correlated with the morphology of thermal dust.

In Fig. 49 we show corresponding angular power spectra, evaluated as cross-spectra between detector-split difference maps (for NPIPE) and half-mission split maps (for SR011 and SR012) outside the *Planck* 2018 polarization analysis mask (*Planck Collaboration IV* 2020). Consistent with the above

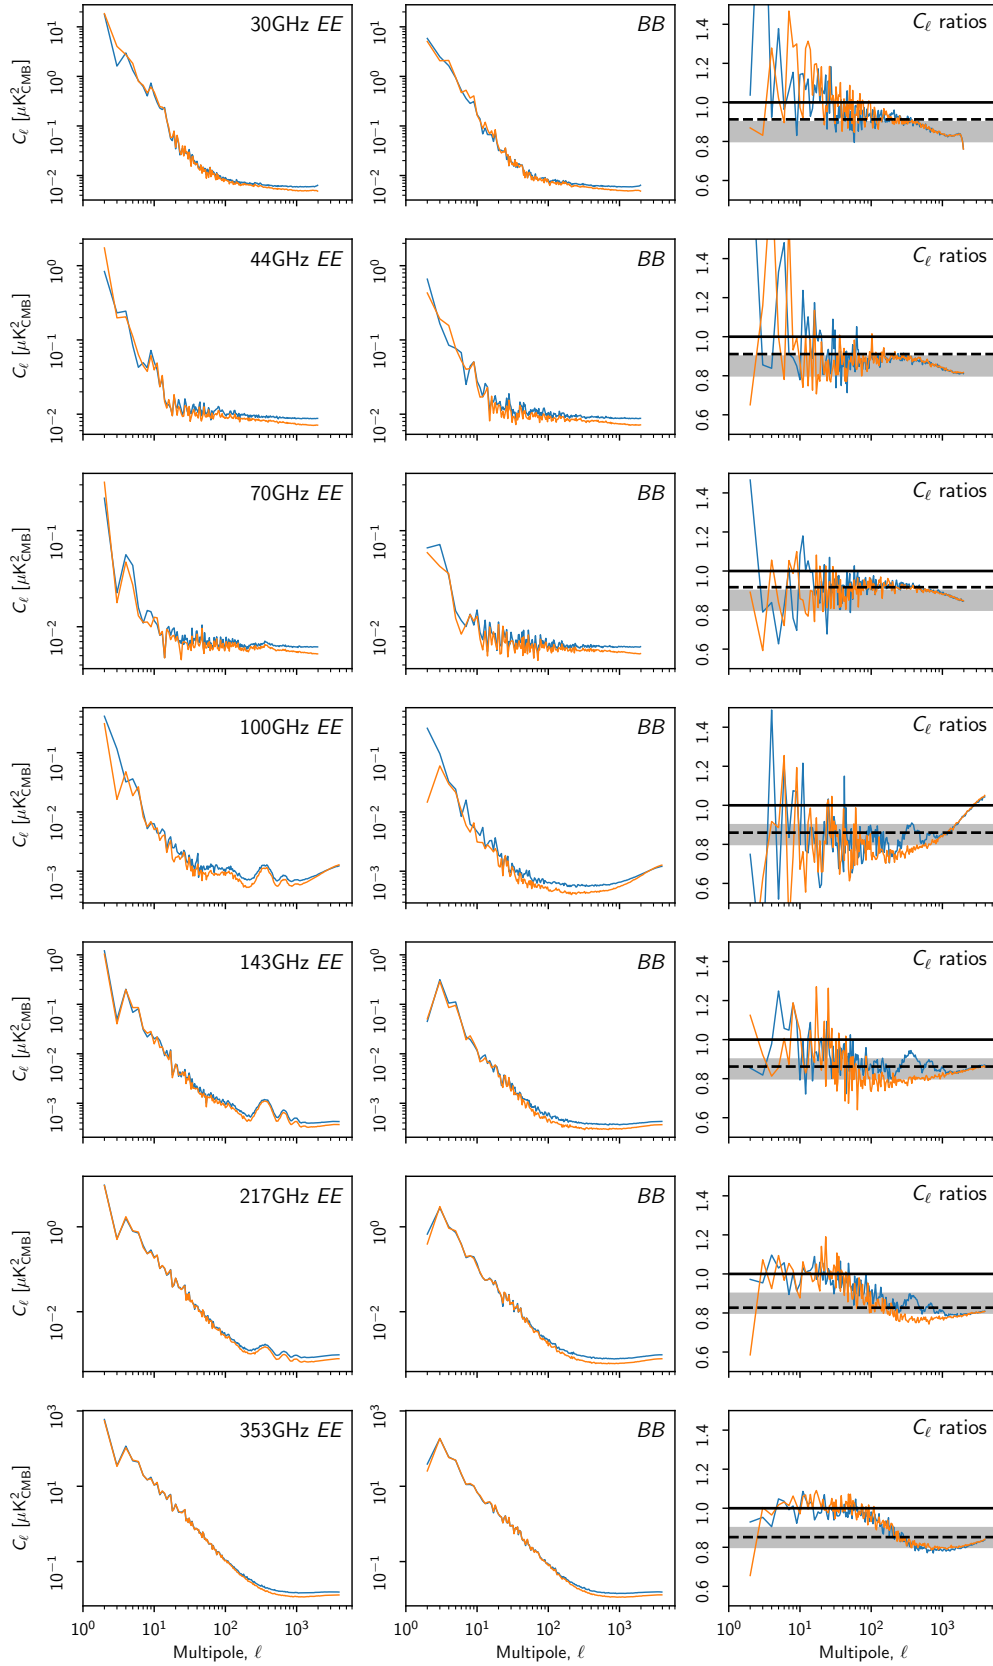


Fig. 38. EE and BB auto-spectra of the polarized frequency maps. The *first two columns* show PR3 (blue) and NPIPE (orange) auto-spectra, while the *third column* shows the ratios (NPIPE/2018) with EE in blue and BB in orange. For noise-dominated angular scales, NPIPE maps have 10–20% lower noise variance, indicated by the grey band in the ratio plot. We show a naive estimate of the ratios, based on the ratio of masked pixel hits in Table 1, as a dashed black line. These spectra are computed over 50.4% of the sky, corrected for the sky fraction and binned into 300 logarithmically-spaced bins.




Fig. 39. Polarization amplitudes of the detector-set difference null maps. The angular power spectra of PR3 and NPIPE maps are shown in Fig. 40. Independent processing of the two detector-sets means that these maps reflect the level of total residuals in the frequency maps.

visual impression, we see that NPIPE exhibits significantly lower EE power at large angular scales, while at small scales all three maps appear broadly consistent.

The observed differences between *Planck* polarization maps illustrated in Figs. 48 and 49 may reflect elements specific to the NPIPE data processing. As discussed in Sect. 4.3.2, polarized Galactic emission is not affected by the same transfer function at low multipoles as the CMB, because polarization templates are part of the data model in Eq. (9). However, this approach introduces an interdependence between the NPIPE frequency maps, which may impact studies of Galactic polarization (see Appendix H). Notably, this is an issue to have in mind when using the NPIPE frequency maps to characterize the frequency correlation of dust polarization (Planck Collaboration XI 2019), an essential question in the search for primordial CMB B -modes.

The bandpass mismatch template used in NPIPE is based on the same sky model as the simulated sky signal, and (as discussed in Sect. 5.2) we do not simulate errors in the template itself. The net effect is that the uncertainty in the polarized emission by foreground Galactic dust is underestimated in the simulations, by the amount that errors in the sky model affect the bandpass mismatch correction. Since the magnitude of this uncertainty is unknown, it may be misleading to rely on

simulations alone to assess uncertainties in polarized Galactic emission on large scales, or to investigate how Galactic polarization decorrelates with frequency. NPIPE is not unique in its approach of not fully sampling the space of template errors. Issues related to bandpass mismatch are a generic feature in *Planck* polarized mapmaking, and caution should be exercised when analysing other *Planck* releases as well.

6.3.2. External consistency

Planck is calibrated without reference to WMAP (Bennett et al. 2013) polarization, although in NPIPE, WMAP temperature data do contribute to the sky model used to derive bandpass-mismatch templates. It is informative to compare the two experiments for agreement in synchrotron polarization. The comparison is particularly interesting due to differences in the two experiments. WMAP's differencing-assembly design allows for gain and bandpass mismatch to be separated into special spurious maps, while *Planck* requires an estimate of the foreground intensity to correct for bandpass mismatch.

We measure the agreement by smoothing the *Planck* and WMAP K -band (23 GHz) maps with a 5° Gaussian beam and then regressing the WMAP K -band and *Planck* 353 GHz maps

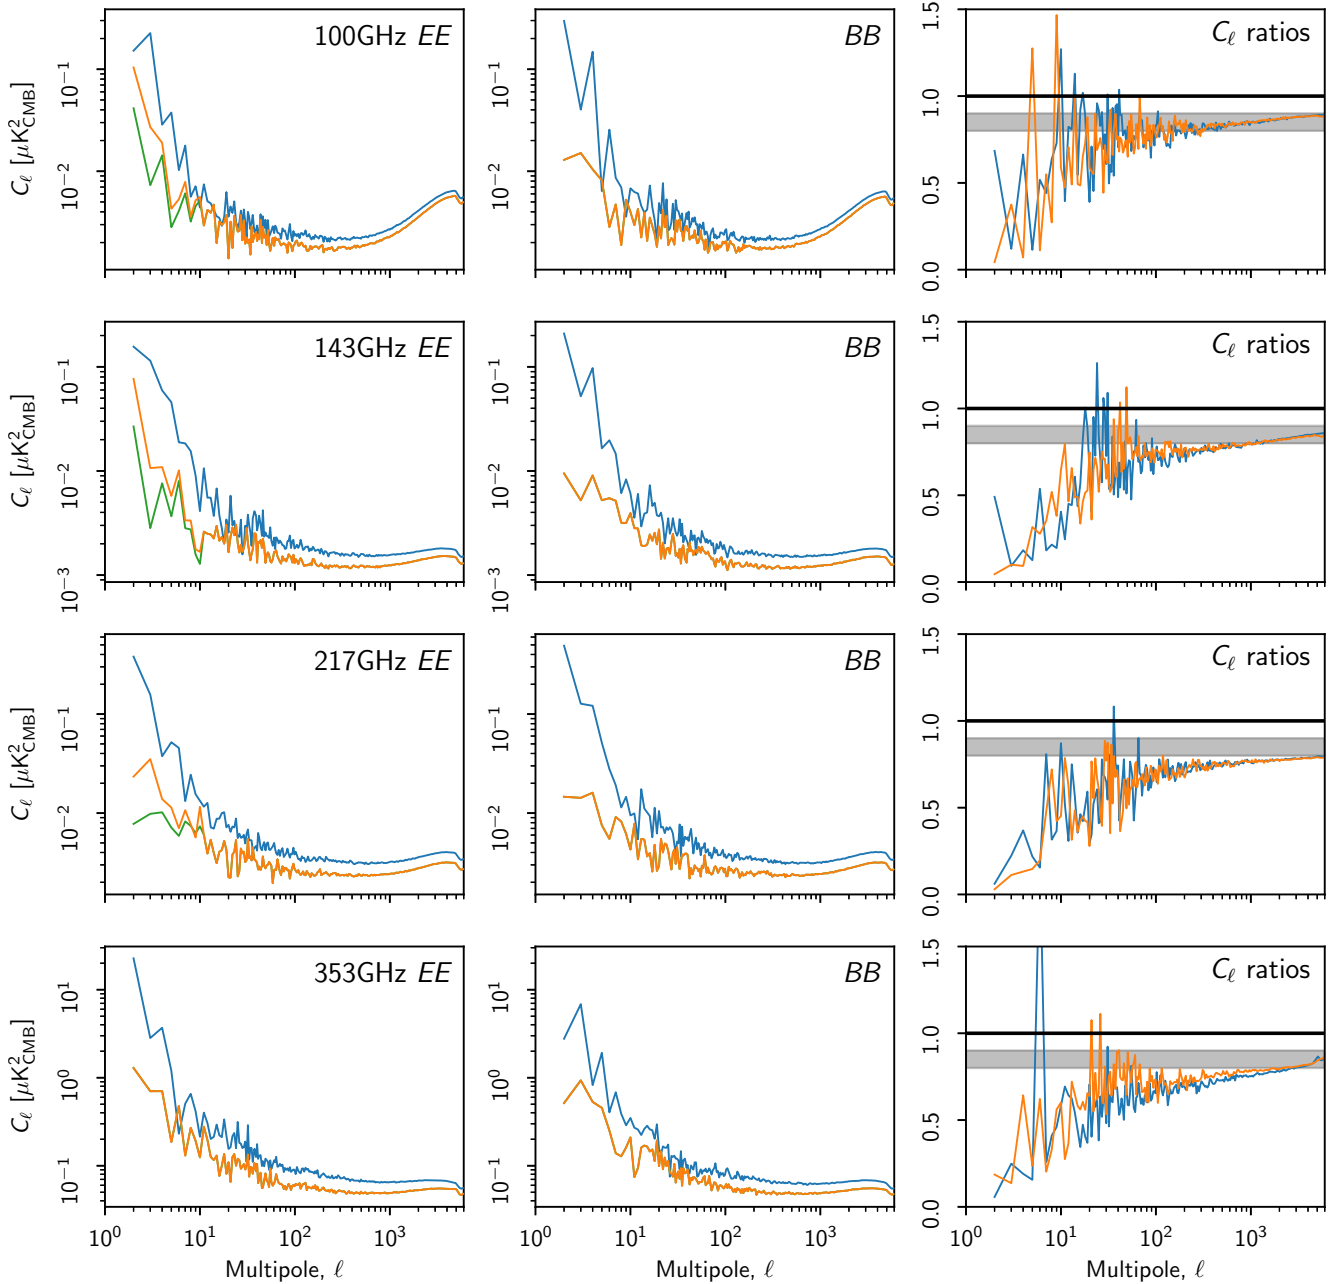


Fig. 40. *EE* and *BB* detector-set difference power spectra. The *first two columns* show PR3 (blue), raw NPIPE (green), and transfer-function-corrected NPIPE (orange) null-map power spectra. Note that PR3 detector sets are *not* the same as were differenced for Fig. 14 in [Planck Collaboration III \(2020\)](#), but rather ones that were destriped independently. The *third column* of panels shows the transfer-function-corrected NPIPE/2018 *EE* and *BB* ratios in blue and orange, respectively. NPIPE has notably less power at all angular scales. The grey band in the third column indicates a 10–20% improvement in power. These spectra are computed over 50.4% of the sky, corrected for the sky fraction and binned into 300 logarithmically-spaced bins. The polarization amplitudes of 2015, 2018, and NPIPE detector-set difference maps are shown in Fig. 39.

from the *Planck*-LFI maps. The residuals for 30 GHz are shown in Fig. 50, and for 44 GHz in Fig. 51. These figures show that improvements in the LFI calibration procedure between the 2015 and 2018 releases significantly improved the agreement between *Planck* and WMAP. They also show that the NPIPE large-scale polarization is more compatible with WMAP than is that of PR3. The fitting was carried out on the 30% of the sky that has the highest polarization amplitude in the smoothed *K*-band map, and we masked out a further 5% of the sky with the highest foreground intensity to reduce bandpass-mismatch-leakage effects, leaving 25.7% of the sky for fitting.

7. Component separation

We now derive new astrophysical component maps from the NPIPE data. For CMB extraction, we employ two different algorithms, *Commander* ([Eriksen et al. 2004, 2008; Seljebotn et al. 2019](#)) and *SEVEM* ([Leach et al. 2008; Fernández-Cobos et al. 2012](#)); for foreground estimation, we only use *Commander*. Both methods have been used extensively in previous *Planck* publications; for full details see [Planck Collaboration IV \(2020\)](#) and references therein. The main motivation for the present analysis is to characterize the internal consistency of the NPIPE

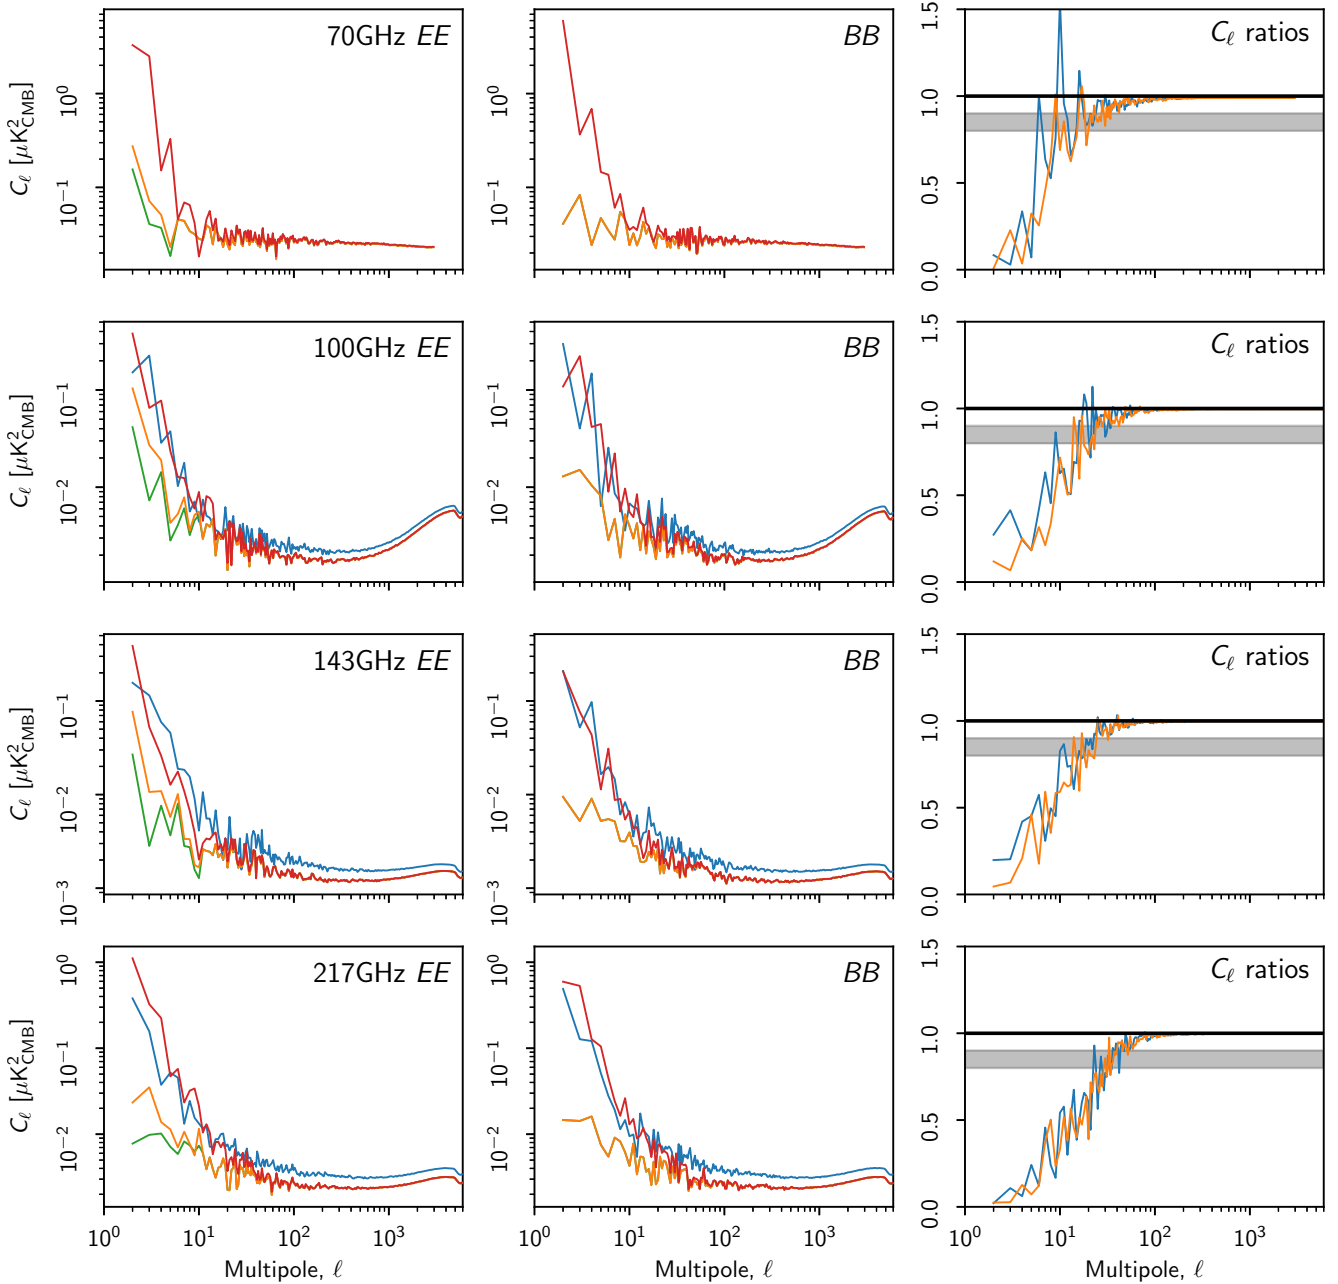


Fig. 41. Effect of the polarization prior on *EE* and *BB* detector-set difference power spectra. For 100–217 GHz, the green, orange, and blue lines on the *EE* and *BB* plots are the same as in Fig. 40 but here we add a red line, showing the power spectra for an alternative version of the NPIPE detector-set maps that are computed *without* the polarization prior. There is no blue line at 70 GHz because there is no comparable detector-set split in PR3, and 353 GHz is not shown because it is always calibrated without the polarization prior.

data themselves, rather than to derive an ultimate NPIPE sky model. In particular, external data sets are not considered in this paper. Instead, combined analyses are deferred to future studies.

7.1. Commander methodology

As described by Eriksen et al. (2008), Commander adopts a parametric approach to component separation. The first step in the process is to specify an explicit parametric model for the data. Since the current analysis considers only *Planck* data, we adopt a model similar to the one employed for the *Planck* 2018 analysis (Planck Collaboration IV 2020). In temperature, this includes

CMB, a power-law low-frequency component, and a single modified-blackbody thermal-dust component. Since NPIPE produces single-detector temperature maps, it can also support individual CO components at all relevant frequencies (i.e., 100, 217, and 353 GHz). This is in contrast to PR3, which includes only integrated full-frequency maps. We therefore model CO emission using three independent components, with single-detector line ratios taking into account the relative bandpass differences between detectors.

One important variation with respect to previous *Planck* releases is the fact that the NPIPE maps retain the CMB Solar dipole. This allows for joint component separation and high-precision relative calibration, eliminates the need to model

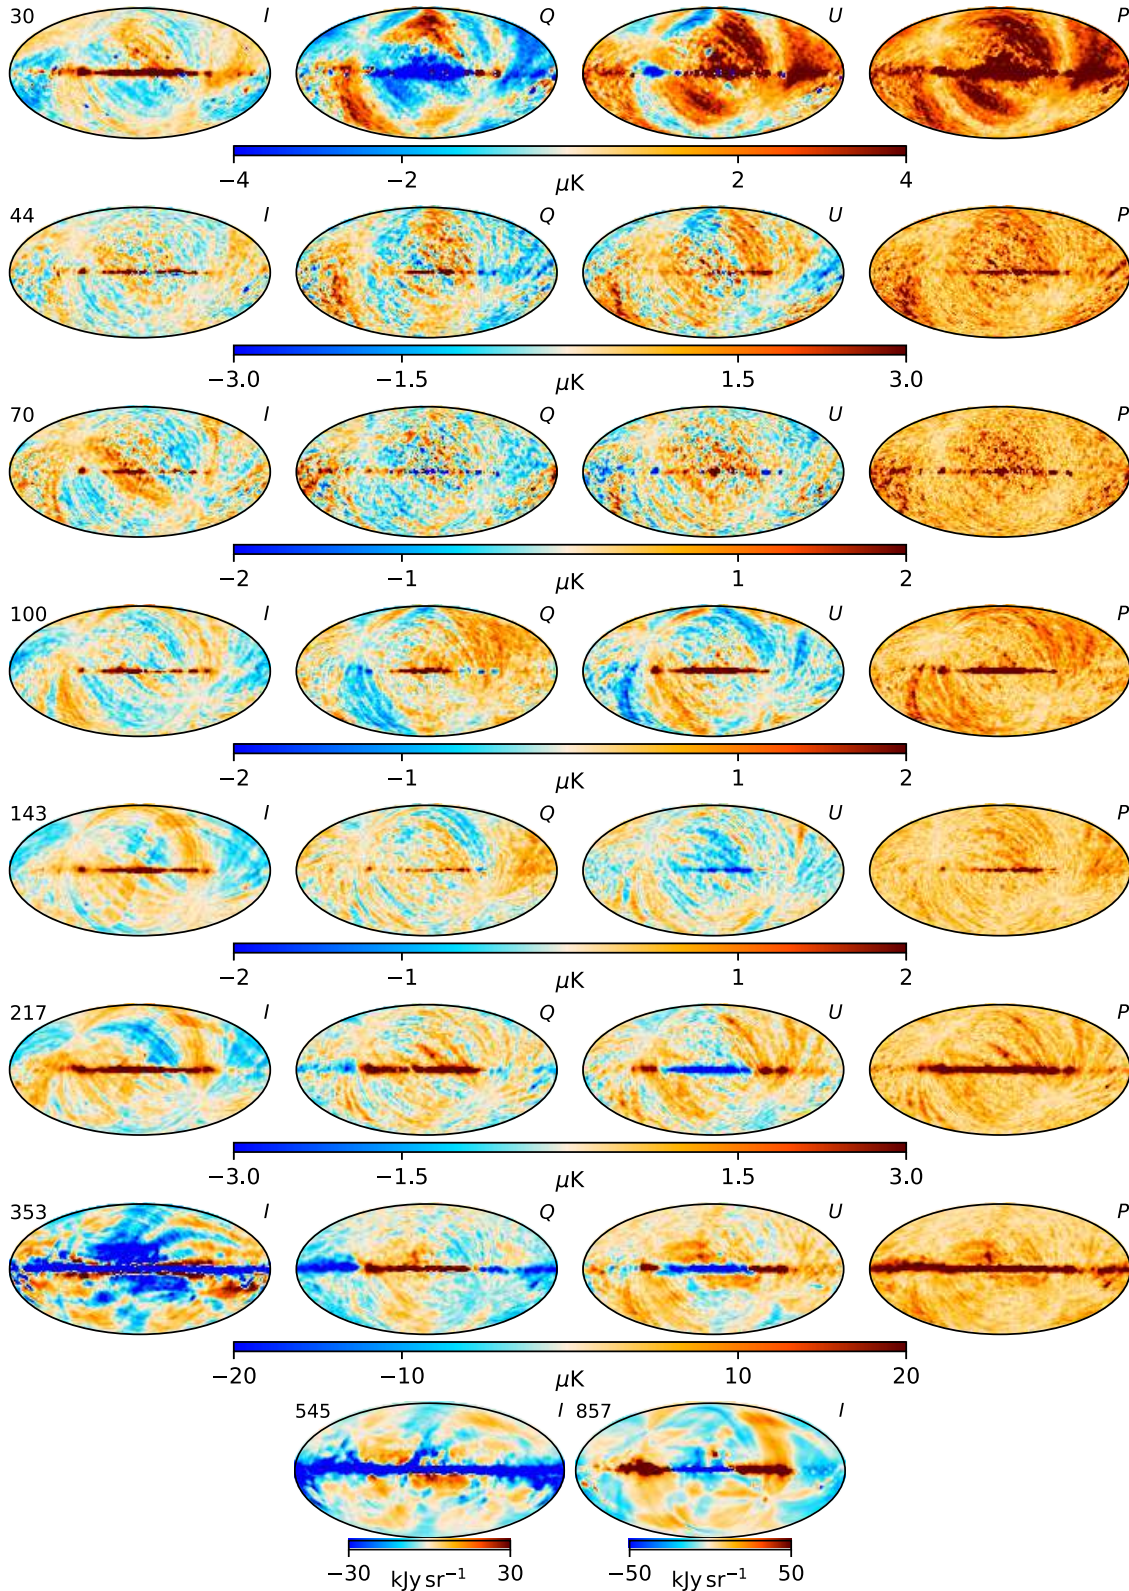


Fig. 42. NPIPE–2018 release difference maps in temperature and polarization. We have projected out the Solar dipole and zodiacal emission templates from the temperature differences, and performed a relative calibration using 50% of the sky to highlight differences beyond these trivial mismatch modes. All maps are smoothed with a 3° Gaussian beam to suppress small-scale noise.

additive dipole components in the frequency maps (Wehus et al. 2014; Planck Collaboration X 2016; Planck Collaboration IV 2020), and removes important large-scale degeneracies from the model. As a consequence, only three global parameters are fitted

per sky map in the following analysis, namely a monopole/zero-level, an absolute calibration, and an overall bandpass correction (Planck Collaboration X 2016), compared to six parameters in previous analyses.




Fig. 43. NPIPE–2015 release difference maps in temperature and polarization to compare to Fig. 42. Note that the Solar dipole model for PR2 did not include the frequency-dependent part of the quadrupole term (Table C.1), so we also omit that correction from the dipole template here. The 353 GHz difference shows that the “zebra” pattern is exclusive to the 2018 temperature map. The polarization differences are indicative of an overall calibration mismatch between polarization-sensitive bolometers, causing substantial temperature-to-polarization leakage from the Solar dipole in the 2015 maps.

The temperature model employed in this paper may be written as a vector in pixel space:

$$\begin{aligned}
 \mathbf{s}_\nu = & u_\nu g_\nu \mathbf{B}_\nu \left[\mathbf{a}_{\text{cmb}} \gamma(\nu) \right. \\
 & + \mathbf{a}_{\text{lf}} \left(\frac{\nu}{\nu_{\text{lf}}} \right)^{\beta_{\text{lf}}(p)} \\
 & + \mathbf{a}_{\text{d}} \left(\frac{\nu}{\nu_{\text{d}}} \right)^{\beta_{\text{d}}(p)+1} \left(\frac{e^{h\nu_{\text{d}}/kT_{\text{d}}(p)} - 1}{e^{h\nu/kT_{\text{d}}(p)} - 1} \right) \\
 & \left. + \mathbf{a}_{\text{co}}^{100} h_\nu^{100} + \mathbf{a}_{\text{co}}^{217} h_\nu^{217} + \mathbf{a}_{\text{co}}^{353} h_\nu^{353} \right] \\
 & + m_\nu, \tag{19}
 \end{aligned}$$

where ν denotes channel and p denotes pixel number. The terms correspond, from top to bottom, to the cosmological CMB signal, a combined power-law low-frequency component with reference frequency ν_{lf} , a single modified-blackbody thermal-dust component with a reference frequency ν_{d} , three CO components ($J=1 \rightarrow 0$, $J=2 \rightarrow 1$, and $J=3 \rightarrow 2$), and, finally, a monopole “offset”. The pre-factors before the astrophysical components are, from right to left: a per-channel beam-convolution operator \mathbf{B}_ν , accounting for the azimuthally symmetric component of the true beam; an overall relative calibration factor g_ν , per channel, with the 143 GHz calibration fixed

to unity; and a unit conversion factor u_ν , scaling from brightness to thermodynamic temperature.

Each component is defined in terms of an amplitude map \mathbf{a}_i at some reference frequency, multiplied by an SED that translates this amplitude map to any other frequency. All SEDs are defined in brightness temperature units. For the CMB component, $\gamma(\nu) = x^2 e^x / (e^x - 1)^2$ is therefore simply the conversion factor from thermodynamic to brightness temperature, where $x = h\nu/k_{\text{B}}T_0$, h is Planck’s constant, k_{B} is Boltzmann’s constant, and T_0 is the mean CMB temperature. For the low-frequency component, the SED is given as a power law with a free spectral index, β_{lf} . For thermal dust emission, the SED has two free parameters, a spectral index β_{d} and a temperature T_{d} . Finally, the multiplicative CO line ratios h_j^i are set to zero for sky maps outside the bandpass range of a given CO component, and set to unity for one arbitrarily-chosen reference sky map inside the bandpass range of the same component (i.e., for the map from one detector at the relevant frequency); all other line ratios for the same component are fitted freely. For instance, h_j^{217} is set to unity for the 217-2 channel, fitted freely for all other 217 GHz channels, and set to zero for all non-217 GHz channels.

For notational simplicity, all unit-conversion and bandpass-integration effects are suppressed in the above model. However, all SEDs are integrated over the full bandpass $\tau(\nu)$ of each detector, using the formulae described in Planck Collaboration V (2014) and Planck Collaboration IX (2014), ensuring that the

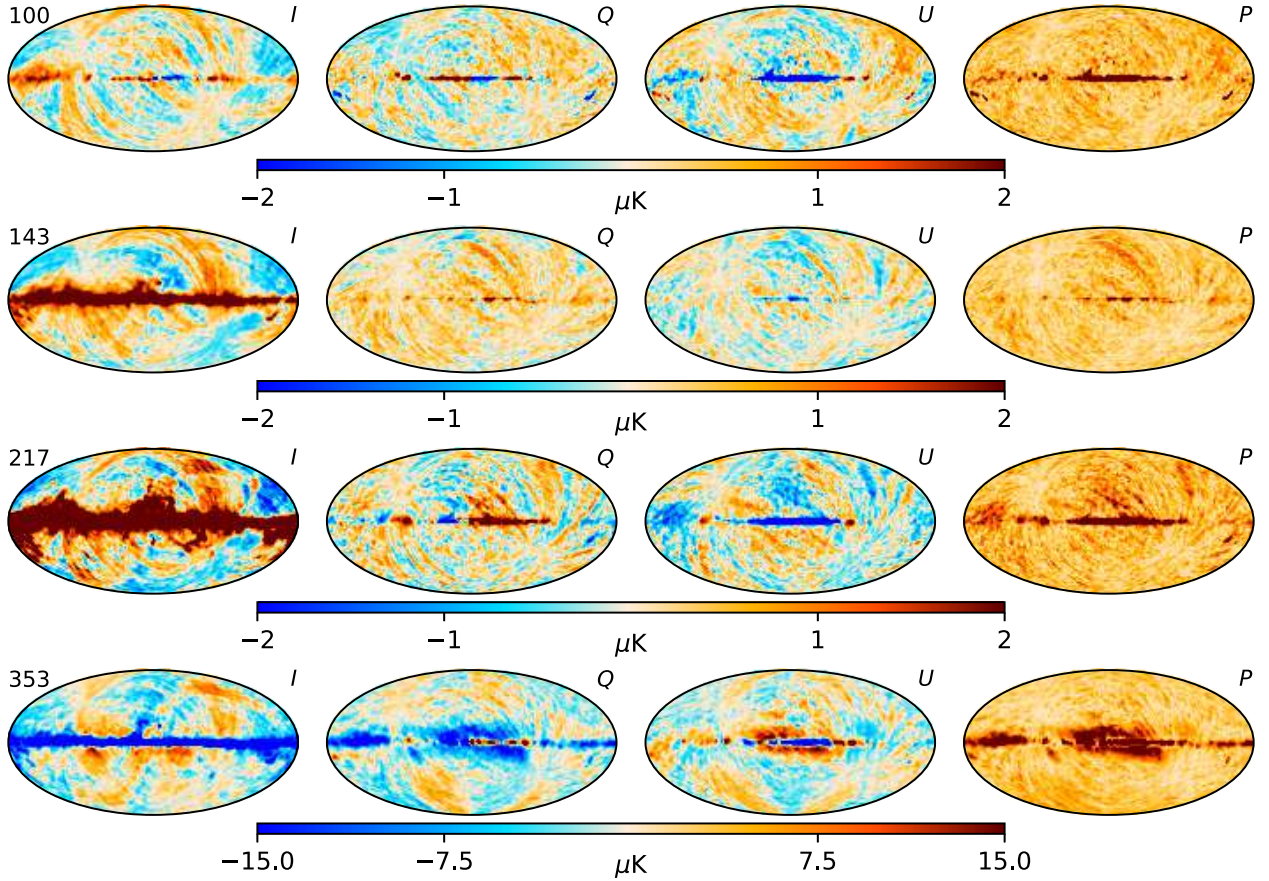


Fig. 44. NPIPE–SRo112 difference maps in temperature and polarization to compare to Fig. 42.

amplitude maps correspond directly to the signal that would be observed at a sharp reference frequency, not as bandpass-integrated quantities. To account (at least partially) for the uncertainties in the bandpasses measured on the ground [Planck Collaboration X \(2016\)](#), we allow for an overall additive shift, Δ_ν , in each bandpass, such that $\tau(\nu) \rightarrow \tau(\nu + \Delta_\nu)$.

For polarization, we adopt a similar, but simpler, model:

$$\begin{aligned}
 s_\nu &= u_\nu g_\nu \mathbf{B}_\nu \left[\mathbf{a}_{\text{cmb}} \gamma(\nu) \right. \\
 &+ \mathbf{a}_s \left(\frac{\nu}{\nu_s} \right)^{\beta_s(p)} \\
 &\left. + \mathbf{a}_d \left(\frac{\nu}{\nu_d} \right)^{\beta_d(p)+1} \left(\frac{e^{h\nu_d/kT_d(p)} - 1}{e^{h\nu/kT_d(p)} - 1} \right) \right]. \quad (20)
 \end{aligned}$$

No CO components or monopoles are fitted for polarization, and the general low-frequency component has been renamed to just “synchrotron”, since neither free-free nor spinning dust emission are expected to be significantly polarized ([Planck Collaboration XXV 2016](#)). Additionally, the thermal dust temperature T_d is fixed to the values derived from the temperature analysis, since *Planck* is not sensitive to polarization above 353 GHz. Note, however, that T_d is spatially varying even though it is fixed throughout the analysis. In contrast, the two spectral indices, β_s and β_d , are fitted directly with polarization data.

All free parameters in these models are fitted jointly using the *Commander* software suite, following the same procedures as detailed in [Planck Collaboration X \(2016\)](#) and [Planck](#)

[Collaboration IV \(2020\)](#). The starting point of these fits is (as usual) Bayes theorem,

$$P(\theta|\mathbf{d}) = \frac{P(\mathbf{d}|\theta)P(\theta)}{P(\mathbf{d})}, \quad (21)$$

where $P(\theta|\mathbf{d})$ is the posterior distribution, $P(\mathbf{d}|\theta) = \mathcal{L}(\theta)$ is the likelihood, $P(\theta)$ is a set of priors, and $P(\mathbf{d})$ is, for our purposes, an irrelevant normalization constant. In most of the following analysis, we report the maximum-posterior solution as our best-fit estimate, and we use corresponding simulations to quantify uncertainties.

We adopt the same set of priors as in [Planck Collaboration IV \(2020\)](#). Most importantly, for SED parameters such as β_i or T_d , we adopt a product of informative Gaussian priors, with central values informed by the high signal-to-noise parts of the data set, and a non-informative Jeffreys prior to account for posterior volume effects due to nonlinear parameterizations. For spatial priors, we adopt angular-power-law-spectrum priors on all foreground amplitudes, but no spatial prior on the CMB amplitudes (for further details, see Appendix A in [Planck Collaboration IV 2020](#)). For global parameters, we impose no active priors on either calibration factors or offsets (g_ν and m_ν , respectively, in Eqs. (19)–(23)), except that we fix one overall calibration factor (at 143 GHz), and one offset per free diffuse component (30 GHz for the low-frequency-component; 100-1a for CO $J = 1 \rightarrow 0$; 143 GHz for CMB; 217-2 for CO $J = 2 \rightarrow 1$; 353-3 for $J = 3 \rightarrow 2$; and 545-2 for thermal dust emission.) The reference calibration factor is fixed to unity, while the reference offsets are determined such that the resulting component maps obtain physically-meaningful zero-levels. Specifically,




Fig. 45. Odd–even survey intensity differences for LFI smoothed to 5° . These maps reflect the internal consistency achieved, not the total residuals, since the calibration errors are correlated between the surveys. To match the 2015 and 2018 processing, the NPIPE 167 ms baseline offsets for this plot are solved using individual survey data.

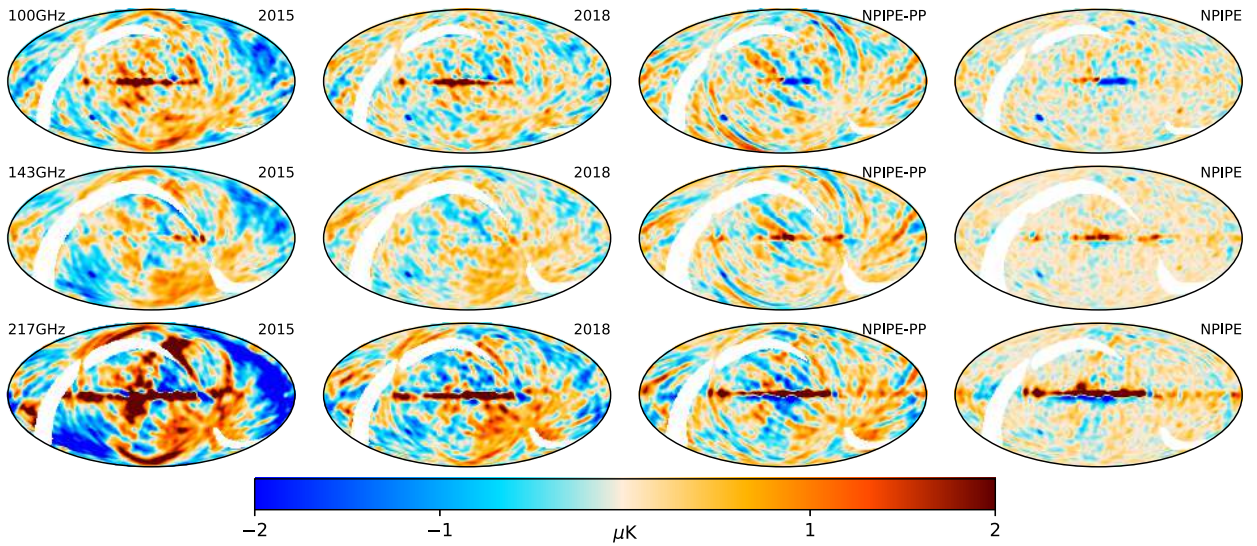


Fig. 46. Odd–even survey intensity differences for 100, 143, and 217 GHz, smoothed to 5° . The 2015 and 2018 maps are the same as in Fig. 12 of *Planck Collaboration III* (2020). The NPIPE-PP column shows the difference obtained if NPIPE is solved only for pointing-period offsets (like PR3), rather than for the 167 ms baseline offsets. The stripes visible in the NPIPE-PP maps are glitch and ADC nonlinearity-correction residuals that are well captured by the short-baseline solution. The comparison is not perfect, because PR2 (2015) baseline offsets were solved using the survey TOD, while the other versions use full-mission baselines. Three variable radio sources can be identified across the frequencies in the NPIPE maps. These maps reflect the internal consistency achieved, not the total residuals, as the calibration errors are correlated between the surveys.

for 100-1a, 143, 217-2, and 353-3, the reference offsets are set such that the CO and CMB maps have vanishing mean offsets at high Galactic latitudes. For 545-2, the offset is set such that a T – T scatter plot between the derived thermal dust map and an HI survey (Lenz et al. 2017) has vanishing intercept. For 30 GHz, it is set such that the derived low-frequency spectral index map does not correlate strongly with the corresponding amplitude map, which is a typical artefact resulting from incorrectly-set offsets (Wehus et al. 2014). The 30 GHz offset determination is the most uncertain among these offsets, since the spectral index information in the map has relatively low S/N.

We approximate the instrumental noise of each detector as Gaussian, and the likelihood therefore is

$$P(\mathbf{d}|\theta) \propto e^{-\frac{1}{2}(\mathbf{d}-s(\theta))^T \mathbf{N}^{-1}(\mathbf{d}-s(\theta))}, \quad (22)$$

where \mathbf{N} is the noise covariance matrix. Since the following analyses are performed at the full angular resolution of the *Planck* instrument, we approximate this matrix by its diagonal, and thereby only take into account the scan-modulated white noise component, not correlated noise features or instrumental systematics. However, since \mathbf{N} affects only the relative weighting between channels, the derived marginal map products remain