

CMB spectral distortions revisited: A new take on μ distortions and primordial non-Gaussianities from FIRAS data

F. Bianchini^{1,2,3,*} and G. Fabbian^{4,5,†}

¹*Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, 452 Lomita Mall, Stanford, California 94305, USA*

²*SLAC National Accelerator Laboratory, 2575 Sand Hill Road, Menlo Park, California 94025, USA*

³*Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, California 94305, USA*

⁴*Center for Computational Astrophysics, Flatiron Institute, New York, New York 10010, USA*

⁵*School of Physics and Astronomy, Cardiff University, The Parade, Cardiff CF24 3AA, United Kingdom*



(Received 10 June 2022; accepted 31 August 2022; published 23 September 2022)

Deviations from the blackbody spectral energy distribution of the cosmic microwave background (CMB) are a precise probe of physical processes active both in the early Universe (such as those connected to particle decays and inflation) and at later times (e.g., reionization and astrophysical emissions). Limited progress has been made in the characterization of these spectral distortions after the pioneering measurements of the FIRAS instrument on the COBE satellite in the early 1990s, which mainly targeted the measurement of their average amplitude across the sky. Since at present no follow-up mission is scheduled to update the FIRAS measurement, in this work we reanalyze the FIRAS data and produce a map of μ -type spectral distortion across the sky. We provide an updated constraint on the μ distortion monopole $|\langle\mu\rangle| < 47 \times 10^{-6}$ at 95% confidence level that sharpens the previous FIRAS estimate by a factor of ~ 2 . We also constrain primordial non-Gaussianities of curvature perturbations on scales $10 \lesssim k \lesssim 5 \times 10^4$ through the cross-correlation of μ distortion anisotropies with CMB temperature and, for the first time, the full set of polarization anisotropies from the *Planck* satellite. We obtain upper limits on $f_{\text{NL}} \lesssim 3.6 \times 10^6$ and on its running $n_{\text{NL}} \lesssim 1.4$ that are limited by the FIRAS sensitivity but robust against galactic and extragalactic foreground contaminations. We revisit previous similar analyses based on data of the *Planck* satellite and show that, despite their significantly lower noise, they yield similar or worse results to ours once all the instrumental and astrophysical uncertainties are properly accounted for. Our work is the first to self-consistently analyze data from a spectrometer and demonstrate the power of such instrument to carry out this kind of science case with reduced systematic uncertainties.

DOI: [10.1103/PhysRevD.106.063527](https://doi.org/10.1103/PhysRevD.106.063527)

I. INTRODUCTION

Observations of the cosmic microwave background (CMB) arguably represent the cornerstone of modern cosmology. Over the last three decades, accurate measurements of the CMB temperature and polarization anisotropies have provided us with a snapshot of the Universe at the time of recombination and have yielded stringent constraints on the constituents, dynamics, and geometry of the Universe [1–4]. This picture is complemented by measurements of the intensity spectrum of the CMB which directly probe the thermal history of the Universe, providing access to additional cosmological information not encoded in the spatial anisotropies [5–7]. In particular, departures from a pure blackbody spectral energy distribution, the so-called “spectral distortions” (SD), open a special window on the

physical processes active before recombination as well as at more recent times [e.g., [8]].

SD naturally arise when thermalization is inefficient in keeping matter and radiation in thermodynamical equilibrium. Examples of mechanisms that can drive the photo-baryonic fluid out of equilibrium are dissipation of primordial acoustic waves and energy injections in the form of photons or other electromagnetically interacting particles (see e.g., [9,10] and references therein for a recent review). Based on their spectral shape, we can broadly break down CMB SD into two classes: the μ - and y -type distortions. At redshifts greater than $z_\mu \gtrsim 2 \times 10^6$, the creation and redistribution of photons by Compton scattering, double Compton scattering, and bremsstrahlung restore thermal equilibrium and erase any SD. Between $5 \times 10^4 \lesssim z \lesssim 2 \times 10^6$, when double Compton scattering and bremsstrahlung become inefficient, energy injections result in a Bose-Einstein distribution with a nonzero effective chemical potential μ , giving rise to μ -type

*fbianc@stanford.edu

†gfabbian-visitor@flatironinstitute.org

distortions [e.g., [11,12]]. Around $z_{\mu y} \lesssim 5 \times 10^4$ Compton scattering becomes inefficient too and photons fall out of kinetic equilibrium with electrons, sourcing the y -type distortions. In reality, the transition from the μ - to the y -era is not abrupt and as a result, residual distortions at intermediate times which cannot be fully described by the sum of μ - and y -types (the so-called r -type distortions) form. Their magnitude within Λ CDM is however expected to be smaller than the sensitivity of proposed future missions like PIXIE [e.g., [13]]. The SD sensitivity to different cosmic epochs not only allows us to probe standard and exotic physics, but also the power spectrum of primordial fluctuations over a broad range of scales, $1 \lesssim k \lesssim 10^4 \text{ Mpc}^{-1}$. While y -type distortions are generated both in the early Universe and during the reionization and structure formation epochs [e.g., [14–16]], μ -type distortions are only produced in the prerecombination era, making them a clean and powerful probe of the early Universe physics. For this reason, we shall focus on the latter type of distortions in this paper.

Observational bounds on the monopole component of these distortions have been set by COBE/FIRAS and are at the level of $|\langle \mu \rangle| < 90 \times 10^{-6}$ and $|\langle y \rangle| < 15 \times 10^{-6}$ (95% CL) [6,17]. At lower frequencies than those covered by FIRAS, ARCADE [5] and TRIS [7] have more recently carried out absolute measurements of the CMB spectrum and derived similar constraints. Several follow-up experiments like PIXIE, PRISM, BISOU, COSMO [18–21] have been proposed to improve absolute measurements of the CMB spectrum.

Apart from the Sunyaev-Zel’dovich (SZ) distortion sourced by galaxy clusters on arcminute scales [e.g., [15,22,23]], CMB SD are predicted to be isotropic signals in the simplest cosmological scenarios. There is, however, an intriguing possibility. If the spectrum of primordial perturbations is non-Gaussian, then a spatial modulation of the SD will be induced, leading to a potentially observable anisotropic pattern of the distortions, as first discussed in [24]. In particular, the local-type non-Gaussianities which peak in the so-called squeezed configuration¹ correlate the small-scale primordial power spectrum, traced by the SD, with the long-wavelength modes probed, e.g., by primary CMB, inducing a nonzero cross-correlation between μ and T , E -, and B -modes [e.g., [24–31]]. As such, a measurement of the spatial correlation between the μ -type distortion and primary CMB anisotropies can constrain the amplitude of primordial non-Gaussianity (encoded by the dimensionless parameter f_{NL}) at very small scales with wave numbers of about $k \simeq 10^2\text{--}10^4 \text{ Mpc}^{-1}$.

This information would greatly complement the view on primordial local non-Gaussianities from *Planck*’s measurement of T and E anisotropies bispectrum at much larger

scales ($\sigma(f_{\text{NL}}) \simeq 5$ at $k \lesssim 0.15 \text{ Mpc}^{-1}$ [32]), providing key insights on the scale-dependence of the primordial non-Gaussian signal and hence on the inflationary mechanism.

In this paper, we use archival data from the COBE/FIRAS experiment to reconstruct a sky map of the μ -type distortion fluctuations. We then correlate this map with the primary CMB temperature and, for the first time, the full set of polarization anisotropies observed by the *Planck* satellite. The extracted SD-CMB cross-power spectra are then translated into constraints on the amplitude of primordial non-Gaussianity of the local-type, $f_{\text{NL}}^{\text{loc}}$ (we will drop the superscript hereafter for simplicity), which are the main one.

We point out that reconstructing the μ -type distortion anisotropies does not necessarily require an absolute measurement of the sky. The work of [33] was, to our knowledge, the first one that attempted a reconstruction of the fluctuating part of the μ distortions using a component separation method applied to imaging data; namely, those from the high-frequency instrument on board the *Planck* satellite. However, this approach can be more affected by contaminations from residual primary CMB and other astrophysical foregrounds e.g., [34,35]. In addition, we stress that knowledge of the μ monopole is needed to break the degeneracy between the average level of μ distortions, $\langle \mu \rangle$, and f_{NL} (see [28,33–36] for discussions on measuring μ fluctuations with a relatively calibrated experiment). This possibility is only allowed by instruments such as spectrometers that are sensitive to the absolute sky temperature.

The paper is organized as follows. In Sec. II we briefly introduce the main ingredients of this analysis, the COBE/FIRAS and *Planck* datasets, while the analysis methodology, from the sky modeling to the cosmological inference, is reviewed in Sec. III. Maps of the anisotropic μ -type distortion, its cross-correlation with primary CMB anisotropies, and systematic checks are presented in Sec. IV. We discuss constraints on primordial non-Gaussianities in Sec. V and compare to previous measurements in Sec. VI. Finally, we draw our conclusions in Sec. VII.

II. DATASETS

A. COBE/FIRAS

The Far Infrared Absolute Spectrophotometer (FIRAS) instrument was a cryogenically cooled Martin-Puplett interferometer on board the Cosmic Background Explorer (COBE) satellite. Designed to cover the frequency range between 30 GHz and 2910 GHz in two spectral bands, FIRAS was able to provide accurate measurements of the CMB spectral energy distribution, thermal emission from interstellar dust, and various infrared cooling lines of the interstellar gas [6]. To accomplish this, FIRAS used a differential optical system with two inputs, one collecting radiation from the sky and one from an internal reference calibrator, and two outputs feeding radiation to composite

¹Defined as the bispectrum configuration where two wave numbers are much larger than the third one, $k_1 \ll k_2 \simeq k_3$.

bolometer detectors. The required absolute calibration was achieved through an external blackbody calibrator. The FIRAS horn antenna accepted incoming sky radiation from a 7° circle, resulting in an approximately top-hat beam function.

In this work we use the calibrated FIRAS skymaps of spectra that have had postcalibration offsets removed, a process known as “destriping”. Specifically, we focus on the final delivery of the FIRAS low frequency low spectral resolution destriped sky spectra, which cover the spectral range from 60 up to 630 GHz in frequency bins with a width of about $\Delta\nu \approx 13$ GHz. The twofold reasons for doing so are that the CMB emission becomes almost entirely negligible at higher frequencies while the noise levels significantly increase. The relevant data products are provided as maps in HEALPix pixelization² at a resolution of $N_{\text{side}} = 16$, corresponding to an approximate pixel size of $\delta\theta \sim 3.5^\circ$,³ for each observed frequency. For a detailed discussion on the instrument, the calibration process, and the released data products, we refer the reader to the FIRAS Explanatory Supplement.⁴

B. *Planck* CMB maps

The *Planck* satellite was the third generation of space-based missions, after COBE and WMAP, to study CMB physics and was dedicated to image the temperature and polarization CMB anisotropies at high angular resolution. Launched in 2009, about 20 years after COBE, *Planck* carried out a full-sky survey of the microwave sky in nine bands across the 30 GHz to 857 GHz range, down to $\sim 5'$ angular resolution [37]. The broad spectral coverage of *Planck* enabled an accurate characterization and separation of the diffuse foregrounds. Unlike FIRAS, *Planck* is not sensitive to the absolute brightness temperature of the sky. Its calibration thus assumes the knowledge of the CMB temperature, as measured by FIRAS, and relies on the modulation of the CMB dipole anisotropy induced by the motion of the spacecraft with respect to the CMB reference frame [38,39]. The *Planck* maps from the 2018 release are calibrated with a precision approaching $\sim 10^{-4}$ at frequencies below 500 GHz and more recent analyses extended this calibration procedure at the highest frequency channels [40]. This accurate measurement in multiple frequency channels allowed to separate the CMB from other astrophysical emissions using different approaches.

In this analysis, we use four foreground-cleaned CMB temperature and polarization anisotropies maps publicly released by the *Planck* Collaboration and derived with different component separation algorithm in both real and harmonic domain; Commander, Needlet Internal Linear

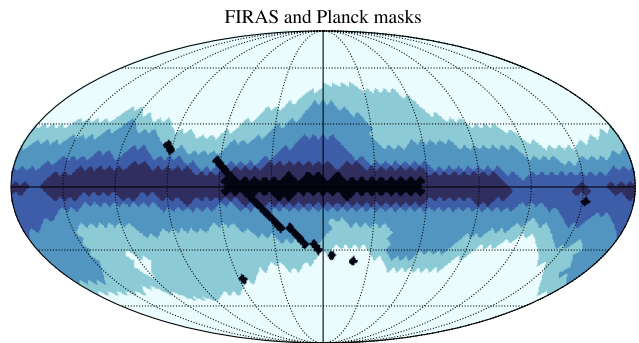


FIG. 1. Sky masks used in this analysis. The black region corresponds to the area masked out by the FIRAS destriper mask while progressively paler colors show the *Planck* masks that retain the 90%, 80%, 60%, and 40% of the sky.

Combination (NILC), Spectral Estimation Via Expectation Maximization (SEVEM), and Spectral Matching Independent Component Analysis (SMICA) [41]. These maps are also provided in the HEALPix format at $N_{\text{side}} = 2048$. We first convolve the *Planck* component-separated maps with the FIRAS scanning beam (which includes the instantaneous optical response and its variation due to the satellite motion as described in Sec. III E) and then downgrade them to a resolution of $N_{\text{side}} = 16$ to produce maps at the native FIRAS resolution.

C. Sky masks

We use different masks to remove pixels close to the Galactic plane where the contamination from foreground emission is high. We adopt the so-called destriper mask of the FIRAS data release (hereafter FDS) that removes the sky pixels not observed by FIRAS as well as those not included in the destriping operation, which are confined to the Galactic center. In addition to this mask, we use the public binary *Planck* Galactic masks that retain 40%, 60%, 80%, and 90% of the sky and that are derived from thresholding the 353 GHz map after CMB subtraction. Hereafter we will refer to these masks, which are shown in Fig. 1 as P40, P60, P80, and P90 respectively and will adopt P60 as our baseline mask.

III. METHODS

In this section we describe the different steps that compose the analysis methodology. We start from the modeling of the FIRAS data cube and foreground emissions; we then discuss the sky-component fitting and power-spectrum estimation, and finally conclude with a description of the cosmological parameter inference framework.

A. Data model

The calibrated FIRAS low frequency sky spectra are a function of frequency ν and position $\hat{\mathbf{n}}$. Effectively, we can

²<http://healpix.sourceforge.net>.

³https://lambda.gsfc.nasa.gov/product/cobe/firas_tpp_all_get.cfm.

⁴https://lambda.gsfc.nasa.gov/product/cobe/firas_exsupv4.cfm.

think of the data as a set of pixelized maps at different frequencies, where the value in each pixel corresponds to the absolute sky emission in MJy/sr. We model the sky emission $I_\nu(\hat{\mathbf{n}})$ as a superposition of several components:

- (1) A blackbody at $T = T_0$ describing the CMB monopole, $B_\nu(T)$;
- (2) A CMB dipole with spectral radiance given by the derivative of the Planck function B_ν with respect to the temperature (assuming it is the result of a Doppler shift of the monopole induced by our peculiar motion);
- (3) Selected foreground emission components, $I_\nu^{\text{FG}}(\hat{\mathbf{n}})$;
- (4) Additional spectral distortions, of μ - or y -type, $I_\nu^{\text{SD}}(\hat{\mathbf{n}})$ (in our notation hereafter SD $\in [\mu, y]$).

Altogether, and similarly to [[17] hereafter F96], the sky model in the direction $\hat{\mathbf{n}}$ reads

$$I_\nu(\hat{\mathbf{n}}) = B_\nu(T_0) + \Delta T(\hat{\mathbf{n}}) \left. \frac{\partial B_\nu}{\partial T} \right|_{T=T_0} + I_\nu^{\text{FG}}(\hat{\mathbf{n}}) + I_\nu^{\text{SD}}(\hat{\mathbf{n}}), \quad (1)$$

where $I_\nu^{\text{FG}}(\hat{\mathbf{n}})$ describes the galactic and extragalactic foreground emissions and $I_\nu^{\text{SD}}(\hat{\mathbf{n}})$ the spectral distortion spectral energy distribution (SED).

In the case of μ -type distortions, we can drop the $I_\nu^{\text{SD}}(\hat{\mathbf{n}})$ term and incorporate it into the generalized Planck's law for blackbody radiation $B_\nu(T, \mu)$

$$B_\nu(T_0, \mu(\hat{\mathbf{n}})) = \frac{2h\nu^3}{c^2} \frac{1}{e^{x+\mu(\hat{\mathbf{n}})} - 1}, \quad (2)$$

with $x = \frac{h\nu}{k_B T_0}$, $\mu(\hat{\mathbf{n}})$ being the chemical potential (at a given sky location), and T_0 the reference CMB monopole temperature. Throughout this work we assume the best fit value of $T_0 = 2.7255$ [42] but we tested that our results are robust with respect to this choice. As a matter of fact, any change in the local temperature of the CMB in a direction $\hat{\mathbf{n}}$ results in a shift in the amplitude of the ΔT parameter. The sum of the first two terms of Eq. (1) represents in fact a Planck blackbody spectrum of temperature $T_0 + \Delta T$. Since we do not make any assumption on its amplitude nor its spatial dependence, the ΔT parameter effectively captures changes in the CMB temperature at linear order and essentially removes residual CMB contaminations in our resulting SD estimate, allowing us to marginalize over this effect in the estimate of μ . We note that including the μ distortion as a nonlinear parameter in the fit, contrary to F96, allows us to minimize degeneracies between CMB and μ (see Fig. 3). In fact, considering μ as a linear deviation around a Planck blackbody spectrum we would have

$$I_\nu^\mu(\hat{\mathbf{n}}) = \mu(\hat{\mathbf{n}}) \frac{-T_0}{x} \frac{\partial B_\nu}{\partial T}. \quad (3)$$

Given that in Eqs. (1) and (3) both the free parameters δT and μ multiply a $\partial B_\nu / \partial T$ term, they become almost degenerate and display a very strong correlation (over 95% according to F96).

A similar approach can be followed to separate y distortion map, where the emission law is

$$I_\nu^{\text{SD},y} = y T_0 \left(x \frac{e^x + 1}{e^x - 1} - 4 \right) \left. \frac{\partial B_\nu}{\partial T} \right|_{T=T_0}. \quad (4)$$

The y -type distortion has a shape that is similar to that of the μ -type, but features a zero-crossing at $\nu \simeq 218$ GHz instead of $\nu \simeq 125$ GHz.

Since the focus of our work is the μ -type distortion and its potentiality to constrain primordial non-Gaussianities, in our baseline analysis we only consider μ as additional spectral distortion components, but in Sec. IV C we explore how including the y -type distortion in the fit affects our results.

B. Foreground modeling

We discuss here the different choices regarding the parametrization of the foreground emission I_ν^{FG} in total intensity.

1. Galactic dust

Dust grains in the interstellar medium absorb UV light from hot stars and reradiate in the submillimeter and infrared bands, dominating the foreground emission at frequencies $\nu \gtrsim 100$ GHz. We model thermal emission from Galactic dust as a modified blackbody described by a dust temperature $T_d = 19.6$ K and a spectral slope of $\beta = 1.6$, which is scaled at a reference frequency of 353 GHz following [43],

$$I_\nu^{\text{dust}}(\hat{\mathbf{n}}) = A_d(\hat{\mathbf{n}}) \left(\frac{\nu}{353} \right)^{\beta_d(\hat{\mathbf{n}})} \frac{B_\nu(T_d)}{B_{353 \text{ GHz}}(T_d)}. \quad (5)$$

We infer the dust brightness in each pixel, $A_d(\hat{\mathbf{n}})$ (expressed in MJy/sr). In the following we will consider two scenarios: one where we keep the dust spectral index fixed to its fiducial value of 1.6, and one where we allow for spatial variations of $\beta_d \equiv \beta_d(\hat{\mathbf{n}})$, i.e., we fit for A_d and β_d separately in each pixel.

2. Galactic synchrotron

Relativistic cosmic-ray electrons accelerated by magnetic fields in our galaxy produce synchrotron radiation. The specific spectral energy distribution depends on the strength of the magnetic fields as well as the energy and flux of the electrons, and typically results in a power law spectrum. Synchrotron radiation represents the dominant foreground contribution for observations of the CMB at low frequencies. We include a synchrotron

component modeled as a power law with $\beta_s = -3.1$ and $\nu_{\text{ref}} = 23$ GHz, i.e.,

$$I_{\nu}^{\text{sync}}(\hat{\mathbf{n}}) = A_s(\hat{\mathbf{n}}) \left(\frac{\nu}{23 \text{ GHz}} \right)^{\beta_s} \quad (6)$$

following [44]. When including synchrotron as a foreground component, we fit for $A_s(\hat{\mathbf{n}})$ in each individual pixel.

3. Free-free

Thermal bremsstrahlung (free-free) emission of free electrons in star-forming regions within our Galaxy is an important foreground to CMB observations at low and intermediate frequencies. To remove this emission, which is mainly concentrated in the Galactic plane, we use the *Planck* Commander free-free template maps for the emission measure $\text{EM}(\hat{\mathbf{n}})$ and electronic temperature $T_e(\hat{\mathbf{n}})$ [44], and we follow the recipe of the same paper to rescale the intensity of the emission across frequencies so that

$$I_{\nu}^{\text{ff}}(\hat{\mathbf{n}}) = 10^6 T_e(\hat{\mathbf{n}}) (1 - e^{-\tau_{\text{ff}}(\hat{\mathbf{n}}, \nu)}), \quad (7)$$

where the free-free optical depth is given by

$$\tau_{\text{ff}}(\hat{\mathbf{n}}, \nu) = 0.05468 T_e(\hat{\mathbf{n}})^{-3/2} \left(\frac{\nu}{1 \text{ GHz}} \right)^{-2} \text{EM}(\hat{\mathbf{n}}) g_{\text{ff}}(\hat{\mathbf{n}}, \nu), \quad (8)$$

and the Gaunt correction factor is [45]

$$g_{\text{ff}}(\hat{\mathbf{n}}, \nu) = 1 + \log \left(1 + e^{4.960 + \frac{\sqrt{3}}{\pi} \log \left[\left(\frac{\nu}{1 \text{ GHz}} \right)^{-1} \left(\frac{T_e(\hat{\mathbf{n}})}{10^4 \text{ K}} \right)^{3/2} \right]} \right). \quad (9)$$

For each sky pixel we then fit an amplitude of the emission template $A_{\text{ff}}(\hat{\mathbf{n}})$.

4. FIRAS Galactic residual

To complement the physically motivated foreground emission models introduced above, we also consider the Galactic residual template from F96. This Galactic spectrum is empirically derived under the assumption that it correlates spatially with the average intensity from the high frequency channels at each pixel.

All foreground components, except for the Galactic thermal dust, are described by one free parameter only, namely the amplitude of the foreground template in each pixel. In our baseline analysis (hereafter referred to as $A_d + \beta_d$ model) we only include emission from Galactic thermal dust where both the amplitude and spectral index are allowed to spatially vary (for a total of four free parameters per pixel) but consider alternative scenarios: Galactic dust with a spectral index fixed to $\beta_d = 1.6$ (A_d model), Galactic dust with $\beta_d = 1.6$ plus synchrotron emission ($A_d + A_s$ model), Galactic dust with $\beta_d = 1.6$ plus free-free emission ($A_d + A_{\text{ff}}$ model), and the FIRAS

model. Finally, we note that in all the scenarios we considered in this work we neglect the molecular lines emission as well as the cosmic infrared background (CIB) emission produced by unresolved dusty star forming galaxies. As far as the CIB is concerned, our baseline model could in principle not only capture the complexity of the dust emission across the sky but also to (at least) partially account for the CIB emission itself. In fact, the CIB itself can be described by a modified blackbody SED with a different T_d compared to the one of the Galactic dust. Given that T_d and β_d are usually highly correlated parameters, our data model can in principle effectively capture the emission of two superposing modified blackbody emission (i.e., Galactic dust and CIB) in a given sky pixel. More complex models should be able to capture this effect more accurately [46]. Moreover, in our analysis we do not include frequencies above 600 GHz where the CIB becomes more important outside of the Galactic plane. For molecular lines, the frequency range we consider covers the CO lines from the ($J = 1-0$) transition at 115.57 GHz to the ($J = 5-4$) transition at 576 GHz as well as the C-I transitions can become important. Since these line emissions are mainly concentrated in the Galactic plane, we do not include them in the analysis. We note that since line emission and β_d are degenerate [see e.g., the discussion in Sec 4.4. of [47]], our fitting could also help partially removing them. Being the signature of μ distortions relatively more important at lower frequencies, the channels potentially contaminated by these line emission have a lower weight in the final component separation fit due to their higher noise, and thus we do not expect them to become important.

We test the sensitivity of our results with respect to the assumed foreground model in Sec. IV C.

C. FIRAS covariance

A key ingredient to carry out cosmological parameter inference from FIRAS data is the sky spectra covariance matrix which describes the correlation structure of observations between different frequencies and sky pixels. As discussed extensively in the FIRAS explanatory supplement [48],⁵ there are six main sources of uncertainties: detector noise (C vectors, using the FIRAS naming), bolometer parameter gain uncertainties (JCJ), emissivity gain uncertainties (PEP), internal calibrator temperature errors (PUP), absolute temperature errors (PTP), and destriper errors β (which include map offsets uncertainties). Note that some of these quantities only vary across pixels, such as β , while others only differ between frequencies (PEP, JCJ, PUP, and PTP).

Denoting an element of the FIRAS data cube $\hat{I}_{\nu}(\hat{\mathbf{n}}) \equiv \hat{I}_{\nu p}$, we assemble the FIRAS covariance matrix as

⁵Available in electronic form at https://lambda.gsfc.nasa.gov/product/cobe/firas_exsupv4.html.




FIG. 2. Noise budget of the FIRAS low frequency low spectral resolution destriped sky spectra. The dotted black line shows the CMB blackbody emission rescaled by a factor 0.001 for visualization purposes. For the JCJ and PEP errors described in Eq. (11) which depend on the sky pixel, we display their average value computed across all the observed pixels.

$$\begin{aligned} C_{\nu p \nu' p'} &= \text{Cov}(\hat{I}_{\nu p}^{\text{FIRAS}}, \hat{I}_{\nu' p'}^{\text{FIRAS}}) \\ &= C^{\nu \nu'} (\delta^{p p'} / N_p + \beta_k^p \beta_{p' k} + 0.04^2) \end{aligned} \quad (10)$$

$$+ S^{p \nu} S^{p' \nu'} (J^\nu J^{\nu'} + G^\nu G^{\nu'} \delta^{\nu \nu'}) \quad (11)$$

$$+ P^\nu P^{\nu'} (U^2 \delta^{p p'} / N_p + T^2). \quad (12)$$

In the equation above, N_p is the map pixel weight, β_k^p is the β matrix described in Sec. 7.2.2 of the FIRAS explanatory supplement and we sum $\beta_k^p \beta_{p' k}$ over all k th orthogonalized stripes used in the data destriping. The $S^{p \nu}$ vectors represent the absolute sky brightness not including the CMB monopole and J^ν is the JCJ gain term described in Sec. 7.3.2 of the FIRAS explanatory supplement while the G^ν terms describe the PEP gain errors. We use the FIRAS data measurement themselves $\hat{I}_{\nu p}$ from which we subtract the CMB monopole adopting a temperature of $T_0 = 2.7255$. We assume $P^\nu = \partial B_\nu(2.728\text{K}) / \partial T$ and U , i.e., the internal calibrator temperature uncertainty, is $180 \mu\text{K}$ as suggested in Sec. 7.4.5 of the FIRAS explanatory supplement. As for the PEP and JCJ terms, we take the PTP terms from the publicly available calibration uncertainty file available on LAMBDA so that $T = 0.002 \text{ K}$. In Fig. 2, we show a breakdown of the different uncertainty terms as function of frequency ν , noting that the main sources of error over the range of frequencies considered here are the absolute temperature errors and detector noise.

D. Map inference

With a data model and an estimate of the measurement uncertainties at hand, we can now reconstruct the spatial

fluctuations of the effective chemical potential $\mu(\hat{\mathbf{n}})$ as well as the other parameters θ that describe the data model.

In the following we will assume that the FIRAS noise covariance is approximately uncorrelated between pixels while still retaining the full frequency-frequency structure. We investigate the validity of this assumption and its impact on the angular power spectrum estimation level in Sec. III E. Assuming pixels are uncorrelated greatly simplifies the analysis since it allows us to perform the parameter inference at the individual pixel level. To this end, we use Bayes' theorem and write the posterior distribution for the model parameters in each pixel as $\propto \ln \mathcal{L}(\hat{I}_\nu | \theta) p(\theta)$, where the likelihood function is given by⁶

$$-2 \ln \mathcal{L}(\hat{I}_\nu | \theta) = \sum_{\nu \nu'} \Delta_\nu^T(\theta) C_{\nu \nu'}^{-1} \Delta_{\nu'}(\theta), \quad (13)$$

with $\Delta_\nu(\theta)$ denoting the residuals between the observed FIRAS spectra and the model spectrum $\Delta_\nu(\theta) = \hat{I}_\nu - I_\nu(\theta)$, and $C_{\nu \nu'}^{-1}$ being the inverse of the covariance matrix defined in Eq. (12) for $p = p'$. We use emcee, an affine-invariant Markov Chain Monte Carlo (MCMC) sampler [49], to explore the posterior distribution and use $N = 100$ walkers for 1000 steps to reach the convergence of the chains (assessed by checking their autocorrelation time). This approach allows us to estimate the full non-Gaussian posterior, the level of degeneracy between different parameters, and to directly marginalize over the impact of foregrounds. In our analysis we do not impose priors on the sampled parameters except for the spectral slope of the Galactic thermal dust, for which we set a uniform prior over $\beta_d \in [0, 3]$. Performing the fit at the individual pixel level also allows to better capture the spatial variations of the foreground emission compared to approaches operating directly on the monopole of single frequency bands, that are often harder to model [50].

In Fig. 3 we show a representative set of credibility contours in the $\{\Delta T, A_d, \beta_d, \mu\}$ parameter space from our baseline data model (Galactic dust with free spectral index) for three different pixels. We can note an anticorrelation between the amplitude and spectral index of the Galactic dust and a positive correlation between μ and the CMB temperature. The extent of the two-dimensional contours reflects the level of noise in a given pixel. The one-dimensional μ posteriors marginalized over the remaining parameters generally follow a Gaussian distribution.

The result of this process is a set of μ values for each pixel in an HEALPix map at $N_{\text{side}} = 16$, along with an estimate of their uncertainty calculated as the standard deviation of the μ posterior.

⁶Omitting the pixel dependence and up to a normalizing constant.




FIG. 3. Constraints on the baseline data model parameters from FIRAS sky spectra. Different colors show results for three different pixels. The darker and lighter shaded regions represent the 1σ and 2σ contours respectively. Typically, our fitting also reduces the correlation between μ and ΔT from the 95% level observed in previous analyses.

E. Power spectrum and uncertainties estimation

After reconstructing the μ distortion anisotropy map from the FIRAS data cube, the next step consists in extracting its angular cross power spectra with maps of primary CMB anisotropies from *Planck*.

A great deal of the information encoded in the SD-CMB cross spectra resides in the lowest multipoles [24,28], therefore it is crucial to build an analysis pipeline that optimally recovers the information on the largest angular scales. For this reason, and given the fact that we work with a dataset at a coarse HEALPix resolution, we use a pixel-based quadratic maximum likelihood (QML) power spectrum estimator $xQML$ [51], an extension of the original method introduced in [52] to cross power spectra. Such methods are computationally more expensive compared to traditional pseudo- C_ℓ methods but we find that they produce measurement of the cross-correlation power spectra with roughly twice smaller error bars for the multipoles most relevant for our analysis. The application of a QML power spectrum estimator requires the knowledge of the covariance of the maps as well as the details of the optical response of the instrument. Below we detail the analysis choices we made.

1. FIRAS beam transfer function and deconvolution

We compute the FIRAS beam response by performing an Hankel transform of the publicly available FIRAS instantaneous beam radial profile. In addition to this smoothing, it is important to take into account the instrument scan motion

during the integration of a FIRAS interferogram. The telescope motion causes the maps to be additionally smoothed in the ecliptic scan direction. We account for this effect by deconvolving an effective transfer function estimated from simulations as described in [53,54]. In this approach, high resolution realizations of CMB maps are first smoothed with the instantaneous FIRAS beam and then smoothed in real space applying a 2.4° boxcar average in the ecliptic direction. The transfer function is then computed by comparing the power spectrum of these maps with the one of the theoretical model used to generate the input CMB maps. We multiply the pixel window function of the HEALPix maps by this transfer function to obtain the total transfer function of the maps that we apply to both the CMB and μ harmonic coefficients when computing auto and cross-power spectra with $xQML$.

2. Covariance matrix of CMB maps

The noise of the *Planck* component-separated CMB maps can be modeled in its full complexity using the publicly available FFP10 simulations which include not only the detector noise in the time-ordered data (TOD), including its correlated $1/f$ part, but also realistic simulations of instrumental effects for all *Planck* frequency channels. These simulated TOD are then processed with the same algorithms as for the flight data, including component separation. We use the data available at NERSC supercomputing center that include 300 realizations of noise and residuals systematic effects for all the component-separated CMB maps and for different data splits (full- and half-mission). After accounting for the *Planck* beam smoothing, we apply the same harmonic domain transfer function that we apply to the data to all 300 realizations for all splits, and downgrade each realization to $N_{\text{side}} = 16$ HEALPix resolution. Given that only 300 noise realizations are available, estimating a full dense pixel-pixel noise covariance matrix for the full-mission dataset is nontrivial. The reason is that the number of matrix elements to be estimated is much larger than 300, making the estimate of the inverse of the covariance matrix poorly conditioned. We therefore consider only a diagonal covariance matrix where the diagonal is given by the variance of the 300 noise realizations for all the T , Q , and U Stokes parameters. While an optimal analysis of the low multipole CMB signal would require an estimate of the off-diagonal terms, for cross-correlations between independent datasets we expect these correlations to be of marginal importance. We also verify that the diagonal term of the QU block of the noise covariance is negligible and we therefore discard it in our analysis. We perform the same operation on the half-mission splits in order to produce a noise model to be employed for null tests analysis and consistency checks. The signal covariance is instead computed on the fly while estimating the power spectrum. To this end, we adopt the fiducial lensed CMB power spectrum of the cosmology from the FFP10

simulations. This assumes a *Planck* 2018 cosmology with a tensor-to-scalar ratio $r = 0$ with the CMB dipole and monopole removed. We stress that an accurate analysis, in particular of CMB polarization, requires the use of the realistic FFP10 simulations since naive estimates of the noise levels relying on the assumption of a Gaussian noise drawn from the noise covariance matrix of the frequency maps largely underestimate the uncertainties on the scales considered in this work.

3. Covariance matrix of μ map

When estimating the μ map, we assume that all pixels are uncorrelated between each other and, under this assumption, for each pixel we obtain an estimate of the error on the inferred μ value in the p th pixel, σ_p^μ , from the standard deviation of the MCMC posterior. However, the covariance definition in Eq. (10) shows a nondiagonal structure in pixel space, thus the estimate of the μ signal in each pixel should show a degree of correlation between pixels similar to the one of the FIRAS data themselves. Therefore, we assumed the μ map is composed by noise only and we adopted two methods to compute its noise covariance matrix:

- (i) We assume the μ map covariance to retain a diagonal structure $\mathbf{C}_{pp'}^\mu = \sigma_p^\mu \delta_{pp'}$. We refer to this approximation as diagonal covariance in the following.
- (ii) We assume the dominant off-diagonal component of the FIRAS pixel-pixel correlation matrix to be inherited by the μ map. Defining the matrix \mathbf{P} from Eq. (11) as

$$P_{pp'} \equiv \delta^{pp'} / N_p + \beta_k^p \beta_{p'k} + 0.04^2, \quad (14)$$

we normalize \mathbf{P} to obtain a correlation matrix $\tilde{\mathbf{P}}$ that we show in Fig. 4 for reference. We then convert $\tilde{\mathbf{P}}$ into a dense pixel-pixel covariance matrix for the μ map $\mathbf{C}_{pp'}^\mu$ using the error estimate from the MCMC as

$$\mathbf{C}_{pp'}^\mu = \frac{P_{pp'}}{\sqrt{\text{diag}(\mathbf{P})_p \text{diag}(\mathbf{P})_{p'}}} \cdot \sqrt{\sigma_p^\mu \sigma_{p'}^\mu}. \quad (15)$$

For each different μ map derived with a different component separation method, we compute the corresponding covariance matrix and use that in the power spectrum estimation step. These two methods describe the noise properties of the maps with different accuracy in different regimes. For FIRAS, no data split that can be used to validate the noise covariance matrix through, e.g., a jackknife approach or a null map is available. As such, we validate the noise model by debiasing the autospectrum of the μ maps and checking its consistency with the null hypothesis. For this purpose, we generate random realizations of the μ map starting from each different covariance model and multiplying the Cholesky decomposition of \mathbf{C}^μ by a stream of Gaussian pseudorandom numbers with zero mean and unit variance. We then compute




FIG. 4. FIRAS pixel-pixel correlation matrix. The pixels are ordered in the ring scheme of HEALPix pixelization, i.e., pixel 0 correspond approximately to the north Galactic pole and the last one to the South Galactic pole. The median correlation coefficient across all measured pixel is about 10% but increase to 20% for pixels far from the Galactic plane.

the noise bias as the mean of the simulations and the error bar as the standard deviation of the same set. The results of this test for different galactic masks are shown in Fig. 5. As we can see, the diagonal μ covariance does not describe the data at large angular scales where we observe a large excess of power that is however greatly reduced if the full covariance model is adopted. The opposite applies when considering angular scales $\ell \gtrsim 10$, where the full covariance model does not deliver an angular power spectrum consistent with zero. In both cases, the χ^2 test yields a low value of the PTEs and a single covariance model is not capable to describe the data on all angular scales. As such, in the following we decide to adopt a hybrid approach for the power spectrum estimation and computed cross spectra using the full covariance matrix for μ for multipoles $\ell < 10$, while the diagonal covariance is adopted for $\ell \geq 10$. The cut at $\ell = 10$ was defined as it is the lowest multipole for which the χ^2 test for consistency with the null hypothesis for $C_\ell^{\mu\mu}$ leads to a PTE higher than 5% for the diagonal covariance model. At large angular scales, the same χ^2 test always leads to a PTE lower than 5%. However, the χ^2 value for the full covariance model is ten times lower than the one obtained for the diagonal-covariance approximation and is driven by the high value of $\ell \leq 4$, pointing to a more accurate description of the data. We will see that this approximation and hybrid approach to power-spectrum estimation of the data is accurate enough to describe the cross-correlation power spectra between μ and the CMB. Finally, we note that an excess of power at large scales might be due to foreground residuals. We test this hypothesis by checking the stability of our estimate as a function of sky fraction used for $C_\ell^{\mu\mu}$ estimation. As we can see in Fig. 5, using a mask covering progressively smaller sky fractions gives to an increased power, which is inconsistent with the

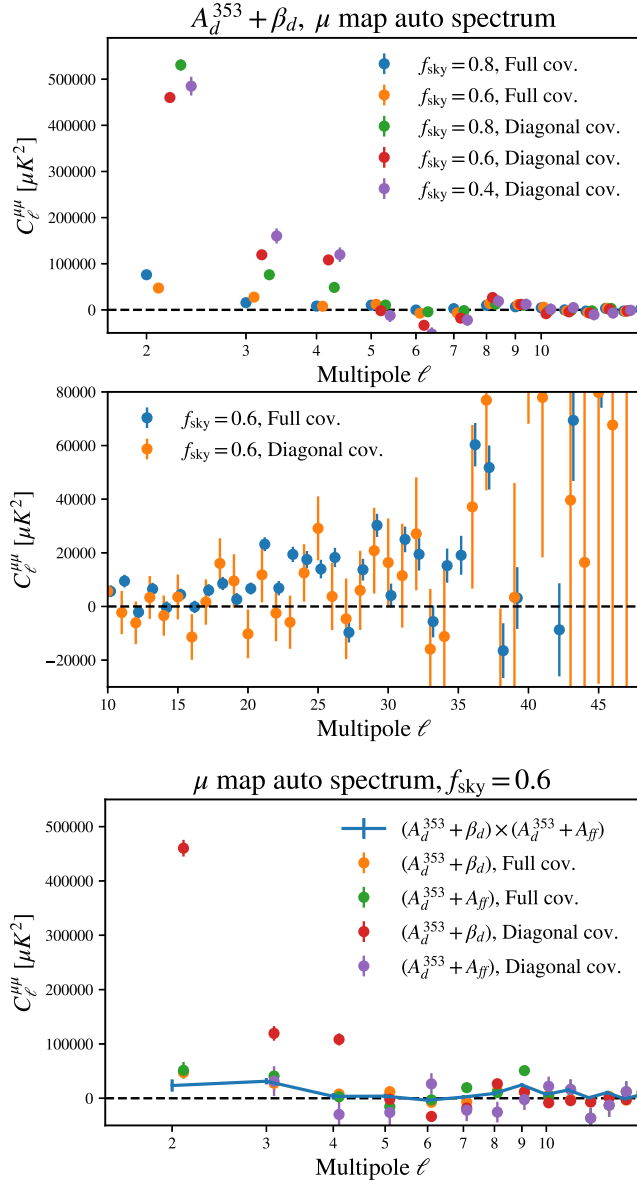


FIG. 5. *Top*: Noise-debiased autospectrum of our μ baseline map on large angular scales for different Galactic masks and different approximation to the pixel-pixel covariance of the μ map. The diagonal covariance does not describe the data accurately enough while the full covariance model provides a more accurate description of the data. *Middle*: Same as above for angular scales $\ell > 10$. In this regime the full covariance model for the μ map is inaccurate and only the diagonal approximation leads to a noise-debiased $C_\ell^{\mu\mu}$ consistent with the null hypothesis. *Bottom*: Same as top panel for μ maps obtained with different component separation methods. Different μ maps have different noise properties and their cross-correlation reduces the residual noise bias in the autospectrum. We show an example of this in solid blue.

expectation of the excess of power being due to foreground contamination as we would expect it to decrease with decreasing sky fraction. As such we conclude that inaccurate noise bias effects dominate the estimate of $C_\ell^{\mu\mu}$.

This conclusion is supported also by the fact that if we take the cross spectra between μ maps obtained with different component separation methods, the noise is only partially correlated between the two maps as the data model that is fitted to the data is different. As such, it is partially free from residual noise bias. In the bottom panel of Fig. 5 we show the bandpowers obtained by cross-correlating the μ maps obtained with the baseline and the $A_d + A_{\text{ff}}$ model adopting our hybrid power spectrum estimation approach, and compare them with the debiased autospectra. As we will discuss further below, A_{ff} does not detect any significant foreground power and as such, the reduced amplitude of the power in the first few multipoles is caused by a lower residual noise bias rather than by an improved foreground subtraction.

F. $C_\ell^{\mu T}, C_\ell^{\mu E}, C_\ell^{\mu B}$ estimation

When correlating the μ map with CMB temperature and polarization maps, we adopt the hybrid approach to evaluate the largest scales outlined in the previous section. In doing so, we also marginalize over spurious monopole and dipole coupling induced by the galactic masking by summing a component proportional to the $\ell = 1$ Legendre polynomial [i.e., $P_1 \propto \cos(\theta_{ij})$, where θ_{ij} is the scalar product between the direction of the i th and j th pixels of the map] and having an amplitude of $1 K^2$. We estimate the error bars on the band powers through Monte Carlo (MC) realizations of the CMB and μ distortion signal. For the CMB signal, we draw random realizations of the T, Q, U Stokes parameters from the theoretical power spectrum of the FFP10 simulations and add the FFP10 official noise realization released by the *Planck* Collaboration. We correlate each of these CMB maps with random realizations of the μ map drawn from either the full or diagonal μ covariance and take the multipoles where each of these approximations to the μ covariance becomes accurate. We use these sets of simulations to verify that the value of the measured $C_\ell^{\mu T}, C_\ell^{\mu E}$, and $C_\ell^{\mu B}$ are consistent with the distribution of the multipoles obtained in the MC simulations.

G. Cosmological inference f_{NL}

The final step of the analysis pipeline is the cosmological inference. Specifically, we are interested in converting the reconstructed SD-CMB cross-power spectra into constraints on the amplitude of the local-type primordial non-Gaussianities, f_{NL} . Assuming that the cross spectra band powers are Gaussianly distributed, we can write

$$\begin{aligned}
 & -2 \ln \mathcal{L}(f_{\text{NL}} | \hat{C}_\ell^{\mu X}) \\
 & = \sum_{\ell \ell'} [\hat{C}_\ell^{\mu X} - f_{\text{NL}} C_\ell^{\mu X}] \Sigma_{\ell \ell'}^{-1} [\hat{C}_{\ell'}^{\mu X} - f_{\text{NL}} C_{\ell'}^{\mu X}], \quad (16)
 \end{aligned}$$

where $\hat{C}_\ell^{\mu X}$ can either be the measured $\mu T, \mu E$, or μB spectra, $C_\ell^{\mu X}$ are the corresponding theoretical templates

calculated for $f_{\text{NL}} = 1$, and $\Sigma_{\ell\ell'}^{-1}$ is the inverse of the band power-band power covariance matrix that we estimated from simulations and includes the effect of correlations between multipoles and correlation between the probes. We verified that the Gaussianity assumption, usually inaccurate for CMB analysis [55], holds for the low, noise-dominated multipoles considered in our work computing the distribution of each of the estimated bandpowers from our simulated dataset and checking their consistency with a Gaussian distribution through a Kolmogorov-Smirnov test. For all the multipoles we indeed obtain p -values above 5%. We account for the finite number of simulations used to estimate the bandpower covariance matrix by debiasing $\Sigma_{\ell\ell'}^{-1}$ following the prescription from [56]. The calculation of the theoretical template spectra is detailed in Sec. V.

From Eq. (16), we can either obtain the maximum-likelihood estimate of f_{NL} and its associated 1σ uncertainty as

$$\hat{f}_{\text{NL}} = \frac{\hat{C}_{\ell}^{\mu X} \Sigma_{\ell\ell'}^{-1} C_{\ell'}^{\mu X}}{C_{\ell}^{\mu X} \Sigma_{\ell\ell'}^{-1} C_{\ell'}^{\mu X}} \quad \sigma(\hat{f}_{\text{NL}}) = \frac{1}{C_{\ell}^{\mu X} \Sigma_{\ell\ell'}^{-1} C_{\ell'}^{\mu X}}, \quad (17)$$

where we recall that the templates are calculated for $f_{\text{NL}} = 1$, or directly sample the posterior assuming an unbound uniform prior on f_{NL} . In this work, we present results based on both approaches. Furthermore, we can jointly analyze the μT and μE power spectra by concatenating both the templates and extracted spectra into two vectors, $\mathbf{C}_{\ell} = [C_{\ell}^{\mu T}(f_{\text{NL}} = 1), C_{\ell}^{\mu E}(f_{\text{NL}} = 1)]^T$ and $\hat{\mathbf{C}}_{\ell} = [\hat{C}_{\ell}^{\mu T}, \hat{C}_{\ell}^{\mu E}]^T$.

IV. MEASUREMENTS

A. Spectral anisotropies, CMB, and foreground maps

We start by showing in Fig. 6 the full-sky maps of our baseline data model parameters inferred from the FIRAS data.

In the top left panel we plot the CMB anisotropy map, ΔT , associated with the component in the data cube that emits as $I_{\nu} \propto \partial B_{\nu} / \partial T$. The map has units of thermodynamic temperature K_{CMB} and is dominated by spatial variations due to the distinct kinematic dipole as expected. The top right panel shows instead the recovered map of the μ -type distortion fluctuations, which is one of the main

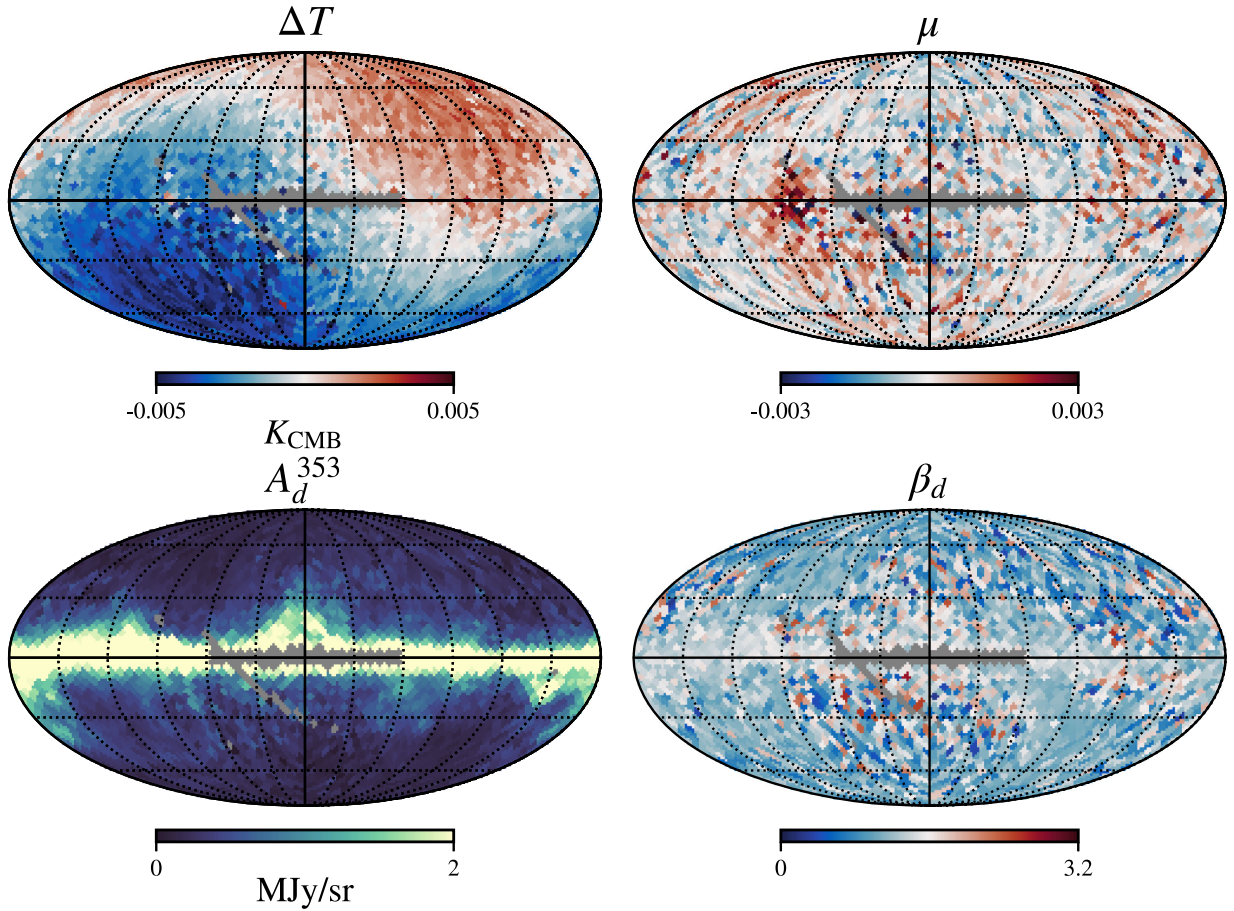


FIG. 6. Maps of the CMB dipole (top left), anisotropic μ -type distortion (top right), amplitude and spectral index (at 353 GHz) of the Galactic dust (bottom left and right respectively) inferred from FIRAS low frequency low spectral resolution destripped sky spectra. Pixels grayed out are removed by the FIRAS destriper mask.




FIG. 7. Scatter plot matrix of the map pixel values after applying the *Planck* $f_{\text{sky}} = 0.6$ Galactic mask. Plots on the diagonal show a Gaussian kernel density estimate of the pixel value distributions in each map. The numbers in the boxes show the Spearman's rank correlation coefficient between the two variables under consideration.

results of this paper. As can be seen, the map exhibits large-scale fluctuations due to the spatially varying noise properties, while the distribution of the pixel values is centered around zero (see the lower right panel of Fig. 7). Finally, the two bottom panels show the maps that describe our baseline foreground model, the amplitude and spectral index of the Galactic thermal dust evaluated at a reference frequency of 353 GHz. The amplitude map A_d , in units of MJy/sr, is recovered at a high S/N and faithfully traces the emission from the Galaxy. The spectral index map is instead noisier, also showing the large-scale noise fluctuations as in the μ map case. The color scale of the β_d map is chosen in a way to show deviations from the reference value of $\beta_d = 1.6$ [57]. We recall that in this work we only use the low frequency data from FIRAS since we are mainly interested in marginalizing over foreground contamination rather than providing their full characterization. For this specific task, the high frequency data would provide additional information. The corresponding uncertainties on the recovered maps are shown in the Appendix.

In Fig. 7, we show a scatter plot of the pixel values for each of the six different map pairs that can be constructed from the parameters in our baseline data model. For this check we apply our nominal fiducial P60 mask before creating the scatter plot. The inferred pixel values are largely uncorrelated between different maps, as quantified by the Spearman's rank-correlation coefficient.

The products of our analysis allow us to provide new estimates of the μ distortion monopole, of the CMB

TABLE I. Statistics of the μ distortion monopole (95% confidence level upper limits, mean and error on the mean) in units of 10^{-6} . We report the values obtained for different choices of Galactic masks.

$[\times 10^{-6}]$	FDS	PL90	PL80	PL60	PL40
$ \langle \mu \rangle < 95\% \text{ C.L.}$	<45	<47	<51	<47	<53
$\langle \mu \rangle$	-16.2	-18.2	-21.7	-15.5	-15.6
$\sigma_{(\mu)}$	14.2	14.3	14.5	15.9	18.6

temperature monopole T_0 , and of the monopole of the dust emission together with its spectral index. We perform an inverse-covariance weighted average of the map pixels and show the results of these estimates in Table I for different Galactic masks. Our improved foreground modeling allows us to use a larger sky fraction compared to the original FIRAS analysis⁷ and sharpen the upper limit on the monopole provided by the original FIRAS analysis $\langle \mu \rangle = (-10 \pm 40) \times 10^{-6}$ [17] by roughly a factor of 2 for our fiducial analysis mask $|\langle \mu \rangle| \lesssim 47 \times 10^{-6}$. Constraints on models predicting energy injections from particle decays or other sources in the μ distortion era should be revised accordingly [e.g., [58,59]]. We obtain stable results using more aggressive Galactic masks. Adopting the same foreground model of the FIRAS analysis and removing a similar sky fraction, we recover an upper limit consistent with their original estimate of $\langle \mu \rangle \lesssim 90 \times 10^{-6}$.

The ΔT template provides information on anisotropic variations of the CMB temperature. The dominating anisotropy is the dipole induced by the motion of the Solar system and the FIRAS satellite with respect to the CMB reference frame. We fit a dipolar emission to the map using the `healpy` dedicated routines on the same fiducial mask of F96 and find an amplitude of the dipole $A_{\text{dipole}} = 3326 \mu\text{K}$ along a direction in Galactic coordinates $(l, b) = (264.13^\circ, 49.21^\circ)$. These values are consistent within $\sim 1.5\sigma$ with the original estimate of the FIRAS team of $(l, b)^{\text{FIRAS}} = (264.14^\circ \pm 0.15^\circ, 48.26^\circ \pm 0.15^\circ)$ and $A_{\text{dipole}}^{\text{FIRAS}} = 3369 \pm 40 \mu\text{K}$ in F96 except for the b coordinate of the dipole. The discrepancy is due to the residual correlation between μ and ΔT parameter. We can restore a complete agreement in the direction of the dipole using more aggressive Galactic masks $(l, b)^{\text{FIRAS}}$ or if we do not fit for the μ distortion on the same mask. In this case we obtain $A_{\text{dipole}} = 3332 \mu\text{K}$, $(l, b) = (264.30^\circ, 48.08^\circ)$. The average of the ΔT template can be added to the pivot value of $T_0 = 2.725$ we used in our data model to estimate the value \hat{T}_0 of T_0 on the sky. By doing so, we obtain $\hat{T}_0 = 2.723$ for our fiducial-component separation method. The value is stable with respect to the mask choice and

⁷97% of the FDS mask compared to the fiducial one of F96 that retained 90% of the sky removing pixels at $|b| < 5^\circ$.

consistent with the ≈ 5 mK temperature reduction compared to the FIRAS original estimates in e.g., F96 $T_0 = 2.728 \pm 0.004$ including statistical and systematics uncertainties as well as the estimates of F96 $T_0 = 2.717 \pm 0.007$ fitting the data to the CMB dipole SED. The shift in our value is expected due to the calibration systematics correction applied to the FIRAS reprocessed data in HEALPIX pixelization.⁸ We also compare the monopole of our dust map with the one obtained by rescaling the Commander dust intensity from the reference 545 GHz to our pivot frequency of 353 GHz using the map of the spectral index and T_d fitted for the low resolution Commander maps and the emission law of Eq. (5). We find a consistent value of 0.29 MJy/sr for both those maps if the *Planck* P90 mask is applied before averaging the pixels. Since our pivot frequency is different than the one chosen by Commander, a direct comparison of the mean spectral index is not straightforward. We obtain a mean dust spectral index across the map $\langle \beta_d \rangle = 1.22 \pm 0.02$. To further validate our component separated products, we compute the cross-correlation coefficient between them and foreground templates released by *Planck*. Specifically, we make use of the dust A_d^{545} and β_d templates derived using the Commander and GNILC algorithms for the 2015 release of *Planck* data. We also correlate the Commander templates for free-free and synchrotron emission against our A_s and A_{ff} maps. For the sake of simplicity, we use a pseudo- C_ℓ power spectrum estimator as implemented in the public code NaMASTER⁹ [60,61] to compute the spectra required to evaluate the cross-correlation coefficient between two X and Y fields, $\rho_\ell = C_\ell^{XY} / \sqrt{C_\ell^{XX} C_\ell^{YY}}$. The results for our baseline $A_d + \beta_d$ foreground model are shown in Fig. 8, where we use a FDS mask and adopt the analytical Gaussian error bars on the cross-correlation coefficient [62]. Our dust amplitude products show a remarkable degree of correlation ($\gtrsim 90\%$) with the *Planck* dust products. The differences observed for β_d are mainly confined to the largest multipoles and are driven by differences in separation in the Galactic plane. The cross-correlation coefficient increases to $\sim 90\%$ at large scales if we exclude the Galactic plane by applying a P90 mask in addition to the FDS mask when computing the cross-correlation coefficient. The CIB does not seem to contaminate our dust templates. The low frequency foregrounds such as free-free and synchrotron are conversely not well constrained by our component separation since the FIRAS data lack a low frequency lever arm to effectively anchor those emissions, with 60 GHz being the lowest available frequency. The cross-correlation coefficient is therefore consistent with the null line for both of these components. This validates our baseline choice to not include those in the reference cleaning method.

⁸https://lambda.gsfc.nasa.gov/product/cobe/firas_tpp_info.html.
⁹<https://github.com/LSSTDESC/NaMaster>.

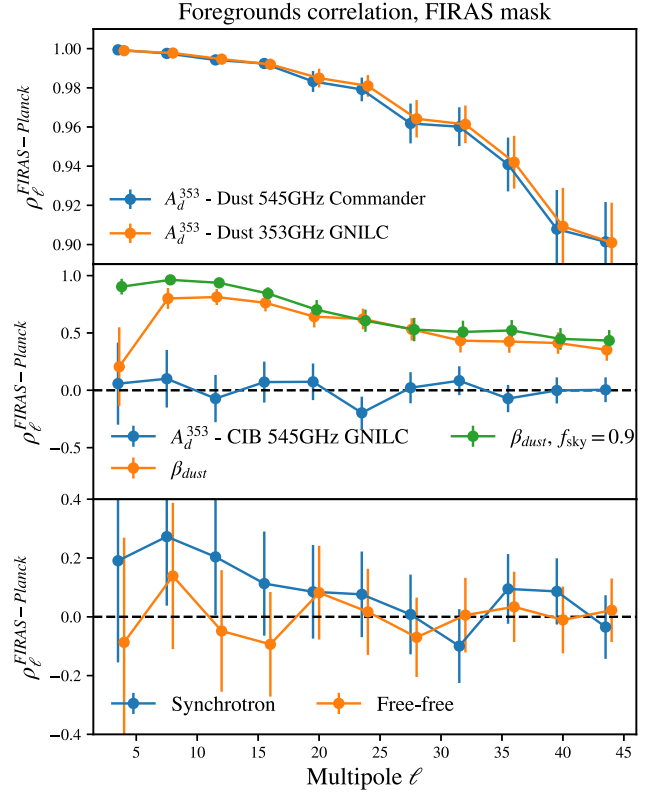


FIG. 8. Cross-correlation coefficient between foreground maps extracted by our baseline component separation analysis and different *Planck* foreground templates. Our foreground maps for the dust intensity and spectral index are highly correlated with the *Planck* data products across all angular scales considered in this work. Discrepancies at the largest scales in β_d are alleviated when excluding the brightest regions of the Galactic plane (removed with a P90 mask). The low frequency foreground components (shown in the bottom panel) are conversely mainly unconstrained and their cross-correlation coefficient is consistent with 0. Points of different templates have been shifted in ℓ for visualization purposes.

B. Power spectra

In Fig. 9 we show the second main result of this analysis, the extracted angular cross-power spectra between the FIRAS μ anisotropy map from FIRAS and the SMICA component-separated CMB temperature and E/B -mode polarization maps from *Planck* (red points). These spectra have been extracted for our baseline foreground model (Galactic thermal dust with free amplitude and spectral index) over about 60% of the sky using a QML approach (see Sec. III E for a detailed discussion of the method). The dark and light blue shaded regions in Fig. 9 show the (14th, 86th) and (5th, 95th) percentiles of the distributions of the SD-CMB cross spectra measured in the FFP10 simulations, respectively. The μT and μE cross-power spectra recovered from FIRAS and *Planck* data are both consistent with zero. Under the null hypothesis that the μ and CMB maps are uncorrelated, we can evaluate

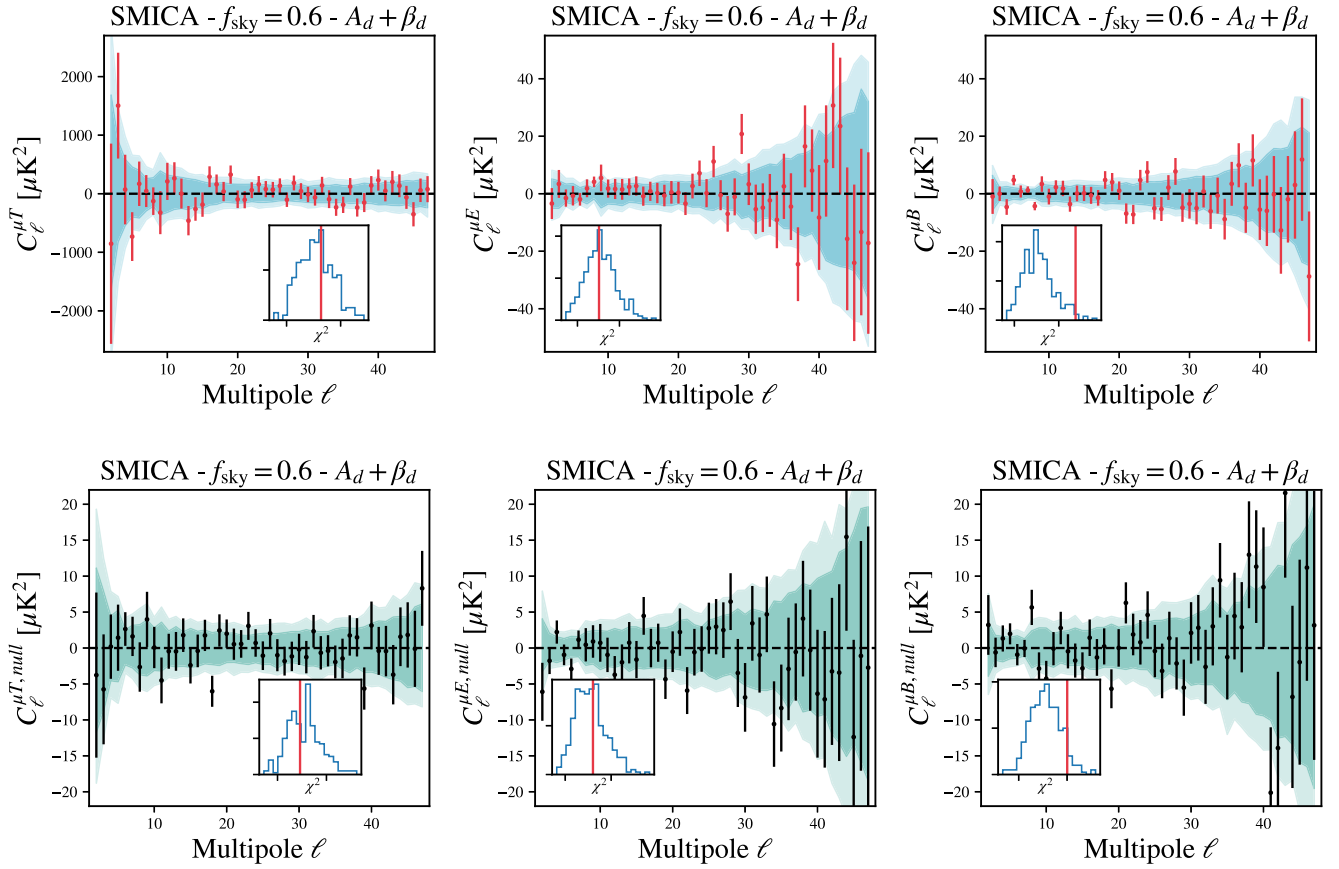


FIG. 9. *Top*: Cross-power spectrum between the anisotropic μ -distortion map from FIRAS and the component-separated CMB temperature (left panel) and E -mode (right panel) maps from *Planck*. This plot shows the baseline results based on the SMICA *Planck* maps, a Galactic mask with $f_{\text{sky}} = 0.6$, and the Galactic dust with free amplitude and spectral index as foreground model. In each panel, the red points indicate the data while the darker and lighter shaded regions represent the 1 and 2σ scatter in simulations that do not include any anisotropic μ . The insets compare instead the χ^2 statistic from simulations (blue histogram) to the value found in data (red vertical line), showing that the extracted spectra are consistent with the null line. *Bottom*: Same as above for the cross-correlation between μ map and the half-mission jackknife map. Data points are shown in black.

$\chi^2_{\text{null}} = \sum_{\ell\ell'} C_{\ell}^{\mu X} \Sigma_{\ell\ell'}^{-1} C_{\ell'}^{\mu X}$ and calculate the probability to exceed (PTE) by counting the number of simulations that have a χ^2_{null} larger than that of the data. We find PTE values of 30% and 67% for μT and μE , respectively: therefore, we cannot rule out the no correlation hypothesis. The distribution of χ^2_{null} for both cases are shown in the inset plots of Fig. 9.

Correlation between μ distortions and B-mode of CMB polarization are potentially a unique probe of tensor non-Gaussianities on scales inaccessible by CMB bispectra. However, any tensor or mixed scalar-tensor bispectra of primordial perturbations that are isotropic are expected to leave either vanishing or strongly suppressed signatures in the observed diagonal $\langle \mu_{\ell_1} B_{\ell_1} \rangle$ correlation that are constrained by $C_{\ell}^{\mu B}$ [29]. Nonetheless we report the amplitude of $C_{\ell}^{\mu B}$ from FIRAS and *Planck* data in Fig. 9. A similar calculation to the one done for $C_{\ell}^{\mu T}$ and $C_{\ell}^{\mu E}$ yields a PTE of about 6%, consistent with a nondetection assumption.

Such null test can be used to constrain parity-violating mechanisms and statistical anisotropies. However, such constraints can be set more naturally in terms of amplitude of the bipolar-spherical harmonics coefficients $C_{\ell_1\ell_2}^{\mu B}$ [63] that capture off-diagonal elements of the μB correlation. A detailed analysis of such correlations together as their implications for non-Gaussian primordial scalar and tensor bispectra is left for future work. In general we note that the use of independent datasets such as FIRAS and *Planck* to constrain μB and μE correlations has the advantage to reduce the impact of large-scale systematics that affect the *Planck* data [40,64].

To further validate the analysis, in the three lower panels in Fig. 9 we show the corresponding null spectra constructed by correlating the FIRAS μ map with the half-mission T , E , B jackknife maps from *Planck*. In this case too, the spectra are statistically consistent with the null line, yielding PTEs of 56%, 42%, and 7% for μT , μE , and μB , respectively. This test also validates the noise model since

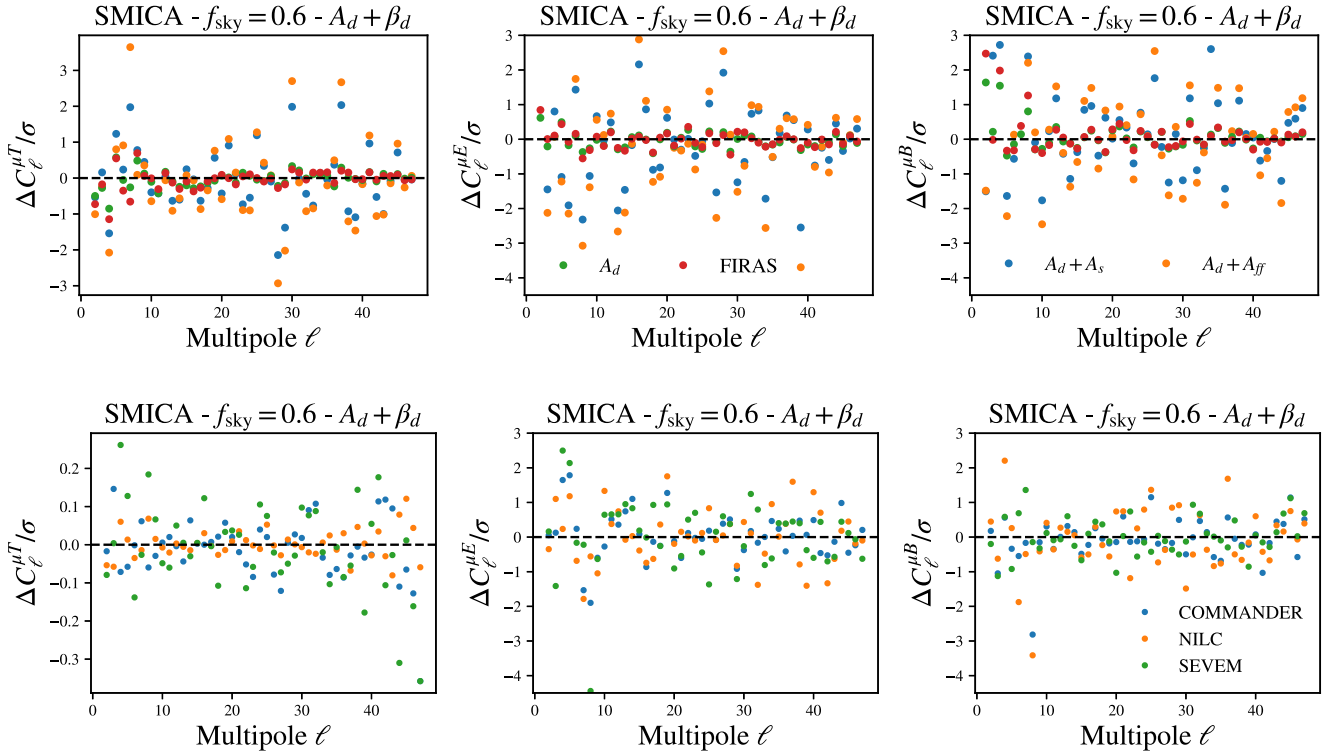


FIG. 10. Top: Changes in the cross-correlation power spectrum between μ distortions and CMB anisotropies of the *Planck* SMICA maps for different component separation methods relative to the statistical uncertainty of our baseline analysis. For the majority of the multipoles the fluctuations are within 1σ . Bottom: Changes in the cross-correlation power spectra used in this work when the reference SMICA maps are swapped with different *Planck* CMB maps obtained adopting different component separation algorithms. Fluctuations in $C_\ell^{\mu T}$ are marginal but some multipoles display fluctuations above 3σ compared to our baseline case for $C_\ell^{\mu E}$, $C_\ell^{\mu B}$. However, these are consistent with noise-induced fluctuations.

the error bars of the null spectra do not contain any CMB cosmic variance.

C. Robustness tests

In this section we perform different systematic tests to validate and assess the robustness of our cross-correlation measurements.

1. Stability against foreground models

We check the stability of our cross-power spectra measurements with respect to different component separation methods both in the CMB and the μ -distortion side. We show the summary of our results in Fig. 10. On the CMB side, we compare the changes in the $C_\ell^{\mu T}$, $C_\ell^{\mu E}$, $C_\ell^{\mu B}$ power spectra obtained using different component separated maps provided by the *Planck* Collaboration (see Sec. II B) relative to the statistical uncertainty determined in our baseline analysis with the SMICA map. Shifts in the $C_\ell^{\mu T}$ values are minor and well below 0.3σ . The cross-power spectra involving polarization show larger variations that remain below 2σ for most of the multipoles. We check that these shifts are expected given the noise level of the *Planck* maps. For this purpose, we generate a set of MC

simulations where, for a specific component separation method, we add a CMB realization (common to all the component separation methods) to each noise realization of the FFP10 suite for that specific method. We then cross-correlate them with a realization of the μ map generated from its pixel-pixel covariance, both using the full and diagonal approximations. This allows us to have a consistent estimate of the power spectrum using the hybrid approach used to analyze the real data. From this set of simulations, we compute the covariance of the differences of the bandpowers between CMB component separation methods that we then use to compute the χ^2 statistics for each of the MC realizations. We then compare the χ^2 value obtained from the shifts in band powers measured on data with the χ^2 distribution obtained from simulations and find that the observed fluctuations between data using different CMB component separation methods are consistent with shifts induced by the noise (PTE > 5%).

We repeat a similar analysis to test whether the SD-CMB spectra obtained using different foreground cleaning assumptions in the construction of the μ map lead to statistically consistent results when correlated with the SMICA map. For this purpose we generate a MC simulation set sharing the same CMB for each foreground

TABLE II. μ distortion monopole upper limits for different choices of foreground separation methods and Galactic masks. We outline in bold values for which the monopole is detected at more than 2σ significance. Since it only happens Galactic masks that include a good fraction of the Galactic plane, this hints for a clear foreground contamination. The last column reports the standard deviation of the values of the upper limits obtained for different masks which we considered as an estimate of systematic uncertainty on the constraint.

$ \langle\mu\rangle < 95\% \text{ C.L. } [\times 10^{-6}]$	PL90	PL80	PL60	PL40	σ_{sys}
$A_d + \beta_d$	<47	<51	<47	<53	2.5
A_d , fixed β_d	<74	<50	<105	<179	48
$A_d + A_s$	< 252	<200	<162	<102	54
$A_d + A_{\text{ff}}$	< 174	<120	<96	<118	47
FIRAS residual	<72	<57	<89	<147	34

cleaning model of Sec. III B and independent realizations of the μ map, since we consider the noise to be mostly uncorrelated between different component separation methods as discussed in Sec. III E 3. We find that for the $A_d + A_{\text{ff}}$ model, which displays deviations up to 3σ for $C_\ell^{\mu E}$ and $C_\ell^{\mu B}$ when compared to our baseline analysis, the PTE for the spectra are consistent at the level of 75% and 73% respectively. This test shows that all the methods clean

TABLE III. 95% upper limits on $|f_{\text{NL}}| \times 10^6$ from the cross-power spectra between anisotropic μ and CMB temperature and E -mode polarization anisotropies. Results are shown for different foreground models and component separation algorithms. The baseline results are highlighted in bold.

	SMICA	COMMANDER	NILC	SEVEM
		$A_d + \beta_d$		
μT	< 5.66	<5.06	<6.53	<5.19
μE	< 5.78	<8.35	<6.95	<7.09
$\mu T + \mu E$	< 3.57	<3.61	<4.18	<3.84
		A_d		
μT	<5.82	<9.58	<9.57	<8.38
μE	<5.37	<6.85	<5.08	<5.50
$\mu T + \mu E$	<3.46	<5.21	<7.15	<5.26
		FIRAS		
μT	<11.49	<10.55	<7.72	<8.27
μE	<7.67	<6.96	<5.63	<5.54
$\mu T + \mu E$	<7.66	<5.45	<5.09	<4.53
		$A_d + A_s$		
μT	<6.56	<6.01	<7.84	<8.61
μE	<11.39	<8.98	<12.57	<9.36
$\mu T + \mu E$	<5.86	<4.17	<5.61	<7.05
		$A_d + A_{\text{ff}}$		
μT	<8.19	<7.36	<9.51	<11.63
μE	<9.88	<10.04	<10.03	<8.71
$\mu T + \mu E$	<6.53	<7.10	<7.77	<7.62

with good consistency the major foreground at this frequency and angular scales, i.e., the Galactic dust, and that residuals of the cleaning do not matter much for the cross-correlation. Different methods however do show differences in the final constraints on f_{NL} and in the overall monopole since the residual noise and signal are different. The baseline model $A_d + \beta_d$ delivers the most stable results against changes in the sky mask or CMB map used in the analysis as shown in Tables II and III, which is why we use this as our baseline analysis.

2. Analysis choices

We check the stability of our analysis choice with respect to several assumptions made in the analysis and in the construction of our reference data model of Eq. (1). In particular we change the value $T = T_0$ where we evaluate the gradient of the blackbody spectrum to the original value used by FIRAS $T_0 = 2.728$; in addition we also try to simultaneously fit for the y and μ distortion. Changing the value of T_0 mainly influences the value of the fitted CMB anisotropy map ΔT but does not lead to any appreciable change in the value of the cross-correlation power spectra nor in the monopole. Averaging the new ΔT maps and summing it to the new value gives consistent results with the estimate we provide in Sec. IV.

Another plausible concern is the fact that we neglect the possible interplay of y and μ distortions as we fit for μ only. A cross-correlation between μ and y distortions could itself be used to constrain non-Gaussianities on different scales and configurations [65]. We do not expect y -type distortions, whose leading term comes from astrophysical objects, to heavily contaminate the measurements at large angular scales in the cross-correlation power spectra. Its anisotropies are in fact mainly sourced by the emission of single massive objects and thus their power spectrum grows at higher multipoles [66]. As a sanity check, we rerun our analysis pipeline when both μ and y distortions are allowed to be present. We find the two parameters to be significantly anticorrelated (by more than 80% on average), suggesting the lack of constraining power of the dataset on both the distortion types simultaneously. We correlate the resulting μ map with the T anisotropy of *Planck*, which is the one potentially contaminated by the tSZ emission, and find no shift in the $C_\ell^{\mu T}$ bandpowers compared to our baseline setup. Finally, we check the stability of our results with respect to the analysis mask. Our monopole constraint are robust to the choice of the mask. We find the $C_\ell^{\mu E}$ and $C_\ell^{\mu B}$ power spectra are still consistent with zero with a reduced statistical uncertainty, however our final cosmological constraint are consistent with the result we present in Sec. V due to a slight shift in the central value that might hint for the presence of mismodeled noise or minor foreground residuals.

3. Foreground deprojection

We test for the presence of residual foreground emission in the μ map by deprojecting external templates of Galactic and extragalactic foregrounds. If we assume that the μ map is contaminated by different foreground emissions described by a set of templates f^i of amplitude α_{fg}^i , we can write

$$\mu = \mu + \sum_i \alpha_{fg}^i f^i. \quad (18)$$

Assuming the templates are uncorrelated, we can estimate the amplitude of the single template as $\alpha_{fg}^i = C_{\ell}^{\mu f^{(i)}} / C_{\ell}^{f^i f^i}$ and derive the foreground-free power spectrum of the μ map from the measured power spectrum $\hat{C}_{\ell}^{\mu\mu}$ as

$$C_{\ell}^{\mu\mu} = \hat{C}_{\ell}^{\mu\mu} \left(1 - \sum_i (\rho_{\ell}^{\mu f^i})^2 \right), \quad (19)$$

where $\rho_{\ell}^{\mu f^{(i)}}$ is the cross-correlation coefficient between the foreground template and the μ map and $\hat{C}_{\ell}^{f^{(i)} f^{(i)}}$ the measured power spectrum of the template. We can extend the formalism to cross spectra with a CMB map X where we assume we have contaminations described by the same template with a different amplitude β_{fg}^i such that

$$X = X + \sum_i \beta_{fg}^i f^i \quad X \in [T, Q, U], \quad (20)$$

so that the foreground deprojected cross-correlation power spectrum reads

$$C_{\ell}^{\mu X} = \hat{C}_{\ell}^{\mu X} - \sum_i \frac{\hat{C}_{\ell}^{\mu f^i} \hat{C}_{\ell}^{X f^i}}{\hat{C}_{\ell}^{f^i f^i}} X \in [T, E, B]. \quad (21)$$

We perform this template deprojection for different types of foreground using *Planck*-based templates that were degraded to the FIRAS resolution and smoothed with the FIRAS effective beam. As for the data pipeline, we compute the cross spectra using the xQML power spectrum estimation pipeline. For the Galactic foregrounds we use the Commander 2015 dust, synchrotron, free-free and CO(2-1) line templates while for the extragalactic emissions we use both Compton- γ maps provided by the *Planck* Collaboration, that are based on different component separation methods, and the GNILC CIB emission templates at 353 GHz. In all cases we do not find significant shifts in the bandpowers, which fluctuate well within the statistical error bars after foreground deprojection.

V. f_{NL} CONSTRAINTS

In this section we translate the extracted SD-CMB cross-power spectra presented in Sec. IV into constraints on the

amplitude of the local-type primordial non-Gaussianities at small scales.

Let us first review the theoretical modeling of the cross-correlation signal between μ fluctuations and primary CMB temperature and polarization anisotropies. To highlight the connection of the SD-CMB cross spectra to primordial non-Gaussianities, we closely follow the earlier works of [24,28,30,31] and define the so-called local-form bispectrum as

$$B(k_1, k_2, k_3) = -\frac{6}{5} f_{\text{NL}}^{\text{loc}} [P_{\zeta}(k_1) P_{\zeta}(k_2) + 2 \text{ perm}], \quad (22)$$

where $P_{\zeta}(k)$ is the primordial spectrum of curvature perturbation, $P_{\zeta}(k) \propto k^{n_s-4}$. This bispectrum peaks in the so-called squeezed triangle configuration ($k_L \equiv k_1 \ll k_2 \approx k_3 \equiv k_s$)¹⁰ and is interesting to study for a twofold reason. First, because measuring a statistically significant deviation of f_{NL} from zero would disfavour single-field inflation models [e.g., [67,68]] and second, because this shape of bispectrum is less prone to contaminations from late-time effects, such as the lensing-integrated Sachs-Wolfe effect bispectrum, which effectively makes it a robust probe of the early Universe [e.g., [69]].

The spherical harmonic coefficients of the primary CMB and μ -type distortion fluctuations are linked to the primordial curvature perturbation $\zeta(\mathbf{k})$ through

$$a_{\ell m}^X = 4\pi i^{\ell} \int \frac{d^3 \mathbf{k}}{(2\pi)^3} T_{\ell}^X(k) Y_{\ell}^{m*}(\hat{\mathbf{k}}) \zeta(\mathbf{k}) \quad (23)$$

$$a_{\ell m}^{\mu} = 4\pi (-i)^{\ell} \int \frac{d^3 \mathbf{k}_1}{(2\pi)^3} \frac{d^3 \mathbf{k}_2}{(2\pi)^3} d^3 \mathbf{k}_3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \times Y_{\ell}^{m*}(\hat{\mathbf{k}}_3) j_{\ell}(k_3 r_{\text{ls}}) f^{\mu}(k_1, k_2, k_3) \zeta(\mathbf{k}_1) \zeta(\mathbf{k}_2), \quad (24)$$

where $X = \{T, E\}$ denotes either temperature or E -mode polarization, T_{ℓ}^X and f^{μ} are the radiation and μ -mode transfer functions respectively, $j_{\ell}(x)$ is the spherical Bessel function, and $r_{\text{ls}} \approx 14$ Gpc is the comoving distance to the last-scattering surface. We calculate the radiation transfer function using CAMB [70] and remind the reader that the μ window function picks up signal from the primordial scalar power spectrum in the $50 \lesssim k \lesssim 12000$ Mpc⁻¹ range [following e.g., [65]].

Then, it is straightforward to calculate the SD-CMB cross-spectrum (in the squeezed limit) as

¹⁰For which we obtain $B(k_1, k_2, k_3) \rightarrow B(k_s, k_s, k_L) \approx -\frac{12}{5} f_{\text{NL}} P_{\zeta}(k_s) P_{\zeta}(k_L)$.

$$\begin{aligned}
 C_{\ell}^{\mu X} &\approx -4\pi \frac{12}{5} \int \frac{k^2 dk}{2\pi^2} T_{\ell}^X(k) j_{\ell'}(kr_{ls}) P_{\zeta}(k) \\
 &\times \int \frac{q_1^2 dq_1}{2\pi^2} f^{\mu}(q_1, q_1, k) P_{\zeta}(q_1). \quad (25)
 \end{aligned}$$

Note that the second integral in the equation above is approximately equivalent to the definition of the monopole of the μ distortion, which we set to its Λ CDM expectation, $\langle \mu \rangle = 2.3 \times 10^{-8}$ [71,72]. From Eq. (25), we can see that the SD-CMB cross-power spectra linearly depend on the product of the f_{NL} parameter and $\langle \mu \rangle$, i.e., $\langle a_{\ell m}^{\mu} a_{\ell m}^X \rangle \propto f_{\text{NL}} \langle \mu \rangle$. Therefore, a larger value of $\langle \mu \rangle$ would translate to a tighter constraint on f_{NL} .

We now have all the tools needed to carry out the cosmological inference. First, we set the amplitude of the primordial scalar perturbation power spectrum to $A_s = 2 \times 10^{-9}$ and its spectral index to $n_s = 0.965$, and calculate the μT and μE template spectra using the theoretical framework outlined above. We then compare the SD-CMB cross spectra measured in our baseline setup over the whole ℓ -range ($2 \leq \ell \leq 47$) to the theoretical curves following the analysis scheme discussed in Sec. III G. The results are presented in Fig. 11.

In the top panel, we show the one-dimensional f_{NL} posteriors obtained by sampling the Gaussian likelihood in Eq. (16) assuming an unbounded uniform prior on f_{NL} . The temperature-only analysis (blue curve) yields a 95% upper limit on the absolute value of $|f_{\text{NL}}|$ of 5.7×10^6 , the E -mode polarization spectra (yellow curve) reveals $|f_{\text{NL}}| < 5.8 \times 10^6$, and their joint analysis suggests $|f_{\text{NL}}| < 3.6 \times 10^6$. The constraints on f_{NL} from SD-CMB cross spectra are robust against changes in the component-separation algorithm used to clean the *Planck* CMB temperature and polarization maps. In Table III we summarize the results based on the alternative cleaning methods which are at the level of few $\times 10^6$, similar to those from SMICA algorithm. If we instead want to be agnostic about the value of the μ monopole, and thus set model-independent constraints, we can directly infer upper limits on the $f_{\text{NL}} \langle \mu \rangle$ product. The corresponding 2σ constraints from μT , μE , and their joint analysis yield $f_{\text{NL}} \langle \mu \rangle < 0.13$, < 0.15 , and < 0.08 respectively. We remind that for high values of f_{NL} , the expressions we used for $C_{\ell}^{\mu T}$, $C_{\ell}^{\mu E}$ become inaccurate as they are based on perturbative expansions in primordial fluctuations. Improving the theoretical modeling of these signals is beyond the scope of this work but we justify their use given the high-noise regime of our measurements and the consistency with 0 of all the measured spectra. To complement the analysis, in the lower panel of Fig. 11 we compare the maximum-likelihood values of $f_{\text{NL}}^{\text{ML}}$ found for the data (vertical lines) to those obtained from the FFP10 simulation cross spectra (histograms). Two things are worth noting. First, the maximum-likelihood values are in agreement with the central values from the Bayesian analysis reported in the top panel.

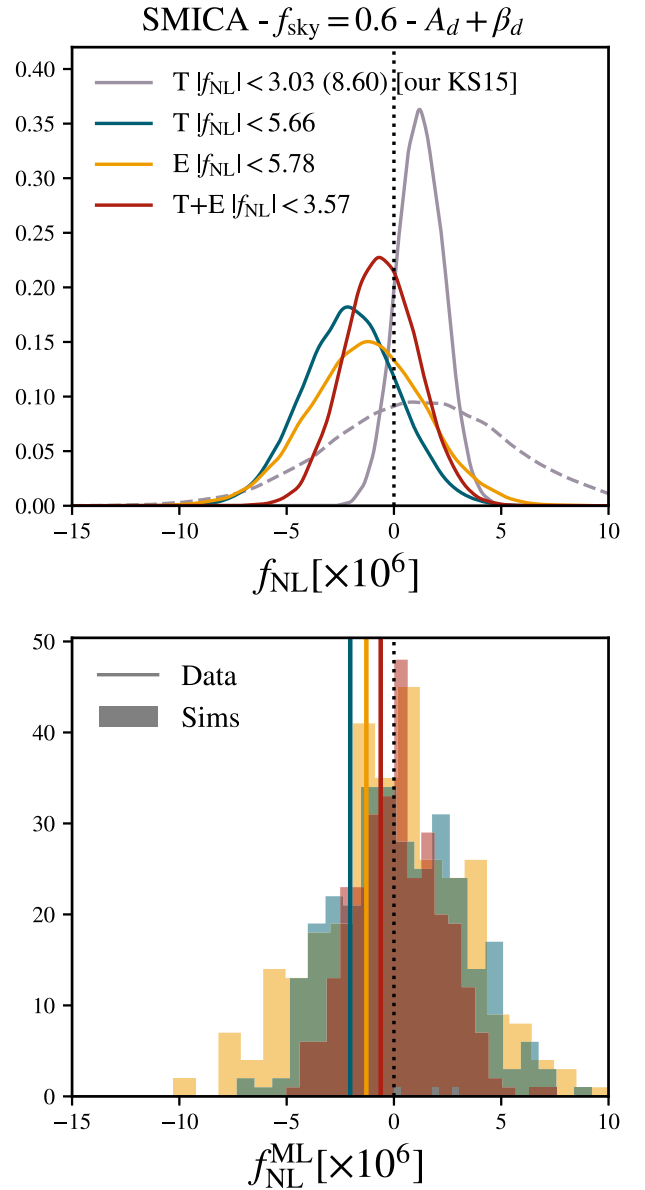


FIG. 11. *Top*: Posteriors on f_{NL} from μT (blue), μE (yellow), and their joint analysis (red). The gray lines show the constraints on f_{NL} from our reprocessing of the [33] data. In particular, the solid one only reflects the effects of the proper band power uncertainties estimation, while the dashed line also accounts for the statistical and systematic uncertainties related to the CIB offset residuals. The numbers in the legend show the corresponding 2σ upper limits on the absolute value of f_{NL} (the number in parenthesis corresponds to the pessimistic constraint from the KS15 reanalysis). *Bottom*: Maximum-likelihood f_{NL} values recovered from the data (vertical lines) and from *Planck* FFP10 simulations (histograms). The color coding is the same as the top panel.

Second, for each of the μT , μE , and $\mu T + \mu E$ cases separately, the best-fit value from the data is consistent with the distribution of f_{NL} from the simulations that do not contain any primordial local-type non-Gaussian signal.

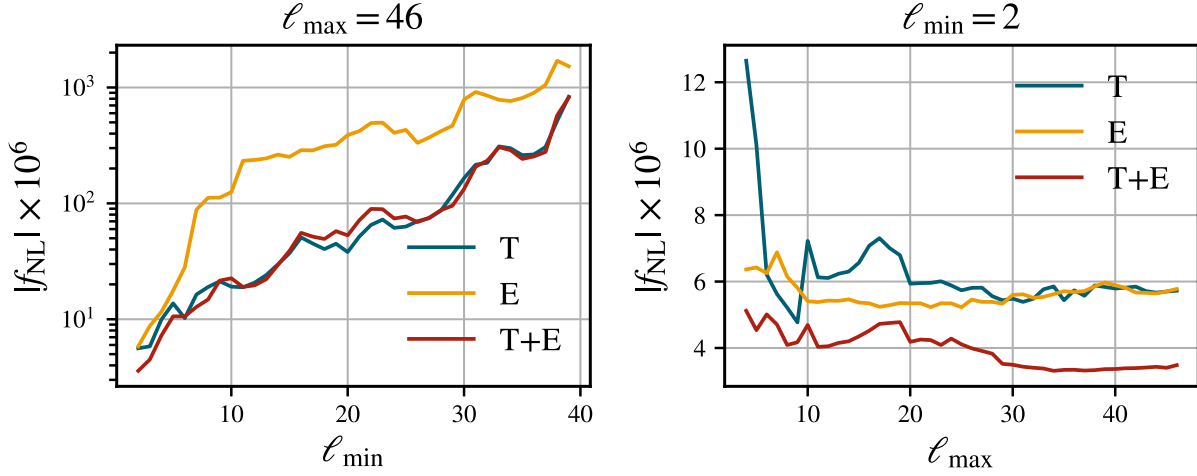


FIG. 12. 2σ upper limits on $|f_{\text{NL}}|$ as function of the minimum multipole ℓ_{min} when fixing $\ell_{\text{max}} = 46$ (left panel) and as function of the maximum multipole when setting $\ell_{\text{min}} = 2$. These curves are reported for our baseline analysis. In each panel, the blue, yellow, and red lines show the constraints from μT , μE , and their joint analysis respectively.

Finally, in Fig. 12 we show the 2σ upper limits on $|f_{\text{NL}}|$ for our baseline case as function of the minimum and maximum multipole included in the likelihood analysis (left and right panel respectively). As we can see, most of the constraining power comes from the first $\ell \lesssim 20$ multipoles, as the upper limit curves flatten out for larger multipoles. For example, reducing the minimum multipole from $\ell_{\text{min}} = 2$ to 10 degrades the 2σ upper limit on $|f_{\text{NL}}|$ by about a factor 6 for both μT and μE , and roughly 3.6 for their combined analysis.

The upper limits on f_{NL} found in this section are significantly weaker than those derived from the analysis of the CMB bispectra measured by *Planck*, $|f_{\text{NL}}| \lesssim 5$ [32], or those obtained from Galaxy surveys using clustering power spectrum or 3D bispectrum, $|f_{\text{NL}}| \lesssim 20$ [e.g., [73–75]]. However, we emphasize that SD-CMB correlations are sensitive to f_{NL} at wave numbers $k_\mu \approx 740 \text{ Mpc}^{-1}$ which are much smaller than those usually probed by CMB temperature and polarization bispectra or large-scale structure. This allows us to place constraints on the scale dependence of non-Gaussianities that can arise in many models of the early Universe such as single-field models having varying sound speed self-interactions, as well as models requiring extra dimensions such as DBI inflation [e.g., [76–80]]. These classes of models could evade the CMB constraints but have potentially observable features in the large-scale structures [25,81]. Considering a simple phenomenological model of the running of primordial non-Gaussianities, $f_{\text{NL}}(k) \simeq f_{\text{NL}}(k_0)(k/k_0)^{n_{\text{NL}}}$ with $k_0 \approx 0.05 \text{ Mpc}^{-1}$ and $f_{\text{NL}}(k_0) \approx 5$, our limit on $|f_{\text{NL}}(k_\mu)| \lesssim 3 \times 10^6$ translates to $n_{\text{NL}} \lesssim 1.4$.

VI. COMPARISON WITH PREVIOUS RESULTS

FIRAS is still a unique dataset to constrain the μ distortion monopole. Being effectively insensitive to the

absolute signal level, modern imagers like *Planck* often rely on the knowledge of the CMB monopole measured by FIRAS and on the annual modulation of the CMB dipole anisotropy (solar dipole) induced by the motion of the spacecraft (orbital dipole), which is known very well, to achieve an accurate intercalibration of the frequency channels [82,83]. At present, the precision of this technique approaches 10^{-4} [40], not far from the precision of the T_0 measurement [42]. Recent works suggested that much better precision needs to be achieved to reach the detection limit of expected cosmological signal of the μ distortion [36]. Nonetheless, some attempts have been made to extract maps of the μ distortion anisotropies from the *Planck* data [[33], hereafter KS15] with dedicated parametric component separation algorithm called LIL [84]. The data of the KS15 analysis are publicly available¹¹ and have been used to constrain $f_{\text{NL}} \lesssim 3.3 \times 10^5$ and $\tau_{\text{NL}} < 2.5 \times 10^{11}$ at 95% confidence level.¹² In order to perform an accurate comparison with our results, we reanalyzed the KS15 data using their fiducial power spectrum estimation pipeline retaining $f_{\text{sky}} = 0.62$, comparable to our baseline setup. We discovered several issues with their original results, that we summarize below, and derive new f_{NL} constraints from the same dataset.

A. Power-spectrum and error-bar estimation

In order to avoid noise correlation leading to a noise bias in the cross-correlation, KS15 created two sets of μ maps from the half-ring splits and used the publicly available

¹¹<https://theory.tifr.res.in/~khatri/muresults/>.

¹²We note that these limits differ from those quoted in the abstract of KS15 because we inserted the exact measured values of the $C_\ell^{\mu T}$ and $C_\ell^{\mu\mu}$ reported in their Eqs. (4.5) and (4.3) in their final formulas of Eqs. (5.32) and (5.33).

TABLE IV. Values of $\mathcal{D}_\ell^{\mu T} = \ell(\ell + 1)C_\ell^{\mu T}/2\pi$ obtained re-analyzing the KS15 μ distortion maps using their same fiducial mask and power spectrum pipeline. The superscript denotes which data split has been used for each μ and CMB field respectively. The error bars are based on the jackknife estimate described in Sec. VI A and exceed by $\sim 30\%$ those of KS15. We report the value obtained by cross-correlating different data splits and different component-separated CMB maps. We use the same units of KS15 to facilitate a direct comparison.

$[\times 10^{-12} K]$	SMICA DR2	SMICA DR3noSZ
μT^{full}	4.1 ± 3.3	4.6 ± 3.5
$\mu T^{(1,2)}$	5.4 ± 3.6	5.9 ± 3.6
$\mu T^{(2,1)}$	2.7 ± 3.5	3.0 ± 3.5
μT^{combined}	4.0 ± 2.6	4.4 ± 2.5

code `polSPICE` [85,86] to compute $C_\ell^{\mu T}$ and $C_\ell^{\mu\mu}$ where the CMB leg is given by the official *Planck* CMB SMICA map of the 2015 release. Given their fiducial analysis mask and maps, we are able to recover their quoted μT cross-correlation amplitude if we correlate the KS15 μ map obtained with the second half-ring split $\mu^{(2)}$ with the *Planck* DR2 SMICA map from the first half-ring split $T^{\text{SMICA,(1)}}$. We report our findings in Table IV. The half-ring maps cross spectra method ensure the absence of any noise bias in the resulting power spectra as the datasets are fully independent. Since all the band powers at scales $\ell > 26$ largely deviate from zero, the authors discarded them and used only the first bin in multipoles which includes angular scales $\ell \in [2, 26]$ to get cosmological constraints. We note that a full analysis of the dataset should use the cross-spectrum computed with the half-ring split 1 for the μ map and the half-ring split 2 for the T map. This is however slightly inconsistent with zero at about 2σ level and would degrade the upper limit of f_{NL} by pushing the measured band power high. The fiducial analysis in KS15 adopted the analytical error bar output of the `polSPICE` code, which assumes uncorrelated Gaussian noise and includes the Gaussian sample variance accounting analytically for the effect of the mask. However, it is unclear whether those approximations are valid for the covariance of component-separated data, in particular for their largest angular scales where the impact of non-Gaussian correlated $1/f$ noise in the data is larger and, in general because the expected distribution of the C_ℓ deviates from the Gaussian approximation for the power spectrum estimation methods they employed [87,88]. We therefore adopt an alternative jackknifing approach to compute the error bars on the measurement. We first identify $N_{\text{JK}} = 192$ patches, corresponding to the pixel area of an HEALPix map of $N_{\text{side}} = 4$ resolution, and then remove one of them before computing the $C_\ell^{\mu T}$ or $C_\ell^{\mu\mu}$ power spectra from the remaining $N_{\text{JK}} - 1$ patches. We repeat the process until every patch has been discarded once from the measurement and then compute the covariance of the measurement from all the measured

C_ℓ -s [89]. We find that the covariance computed for all the combination of the splits exceed significantly the analytical error bars, by as much as 30% for first power spectrum bin in $C_\ell^{\mu T}$ and on average by a factor 50% across the whole range of multipoles considered in KS15. Conversely, the error bars for the $C_\ell^{\mu\mu}$ power spectrum, computed in KS15 as the cross-correlation of the $\mu^{(1)}$ and $\mu^{(2)}$ maps is consistent with the analytical estimate in the first multipole bin but should be inflated by roughly by factor of 5 on average for the smallest angular scales.

B. Foreground contamination

We test for residual foreground contaminations in the maps using the deprojection techniques introduced in Sec. IV C 3 and adopting the same set of Galactic and extragalactic foreground templates. We first compute the cross-correlation coefficient of both the *Planck* SMICA DR2 map used in the KS15 analysis and their μ map with all the templates. We show that the main foreground residuals in the map are represented by Galactic dust, tSZ and, less importantly but non-negligibly, CIB, for which we detect a 60%, 40%, and 30% correlation respectively. The magnitude of the correlation is stable with respect to the choice of different dust (e.g., COMMANDER or GNILC based) or tSZ templates (NILC or MILCA). For the SMICA DR2 map, the correlation coefficient we observe with the same set of templates is well below the 5% on all relevant scales. As such, the deprojection applied to the $C_\ell^{\mu T}$ power spectrum does not bring any statistically significant shift in the first bandpower in particular, and it does not bring the measured power spectra to be consistent with 0. The sharp oscillatory features suggest that the maps are dominated by primary CMB leakage and the tSZ deprojection seems to enhance those features at small angular scale, pointing to a non-negligible residual in the μ map. This is somehow expected since the data model of KS15 does not include y . As such the resulting μ map is $\sim 30\%$ correlated with the y templates in the first bin, while the SMICA map is anticorrelated with y at the 5% level on all angular scales. When correlating the KS15 μ map with the tSZ-deprojected SMICA map of the *Planck* DR3 release, which is designed to be tSZ-free, we find consistent results with the analysis discussed above (see Table IV), supporting the hypothesis of the spectra being dominated by CMB leakage.

C. Map offsets and CIB monopole uncertainties

The sky monopole is not constrained by *Planck* data. The correlated noise component in the raw time-ordered data is in fact modeled by a sequence of baseline offsets of a specific length in time tuned to optimize the noise removal and minimize the computational cost of the map-making step [38,39,90]. Therefore, the monopole cannot be distinguished from a global noise offset and the final map

TABLE V. Estimates of the CIB monopoles based on [92] used to set the zero-level maps of the *Planck* HFI channels and their new data-driven estimates of [54].

ν [GHz]	O19 [MJy/sr]	Planck [MJy/sr]
100	0.007 ± 0.014	0.003 ± 0.003
143	0.010 ± 0.019	0.0079 ± 0.0079
217	0.060 ± 0.023	0.033 ± 0.013
353	0.149 ± 0.017	0.13 ± 0.026

monopoles can assume arbitrary amplitudes. For the *Planck* LFI instrument, these are estimated during the calibration step and removed from the final map, while for the HFI channels used in KS15, these are fixed to reproduce the value of the dust and CIB monopole. While the first set of offset values is estimated from the data through correlation with maps of the HI column density [91], the latter are inserted according to the expected CIB monopole model of [92]. The uncertainties on the offsets depend on the specific astrophysical dataset used to constrain them and can be large [see [54], hereafter O19, and references therein]. The recent analysis of the FIRAS team in O19 provided new estimates of the CIB monopole at the *Planck* HFI frequencies (100, 143, 217, 353, 545, 857 GHz). These are the first data-driven values for frequencies below 200 GHz and are sometimes in tension with the model used by *Planck* for the 217 GHz and 353 GHz channels (we report all relevant values in Table V). Despite the component separation algorithm of KS15 should remove the nominal monopole injected in the maps, we try to quantify how these uncertainties can affect their final result and its uncertainty. For this, we consider a single-pixel toy model where we fit for a μ distortion in a sky model that includes a CMB temperature anisotropy of $200 \mu\text{K}$ and Galactic dust using only the 100 GHz, 143 GHz, 217 GHz, and 353 GHz as in KS15. In the MC simulations, at each frequency we add a random realization of the average *Planck* noise adopting the numbers in Table 12 from [38] before performing the fit. In this case we recover an unbiased estimate of μ . We then proceed to perform a similar test where we add to all simulated frequencies not only the instrumental noise but also a systematic error that represents the uncertainty in the map offset after CIB monopole subtraction. For each frequency, we draw this error from a Gaussian distribution with standard deviation equal to the error of the measurements of O19 and a mean either equal to 0 or equal to the difference between the CIB monopole values adopted by *Planck* and those of O19. These cases cover two different hypotheses; the first accounts only from a data-driven estimate of the statistical uncertainty in the estimate of the CIB offsets residuals, while the second considers also a systematic mean residual in the maps.

In both cases, the overall error on the recovered μ (estimated as the standard deviation of the samples) increases compared to the case with no systematic errors.

When we introduce a nonzero mean residual offset, the recovered distribution of the μ pixels becomes inconsistent with zero at $\sim 1.5\sigma$. The recovered CMB also appear slightly biased at a similar statistical significance. We repeat the analysis by injecting the systematic errors one frequency at a time and find that the major source of the bias and uncertainty are the monopole errors at 143 GHz and 353 GHz. This shows that the limited frequency coverage used in KS15 makes their analysis very sensitive to the CIB monopole uncertainties and their statistical error bars underestimate the overall uncertainty in the measurement. The ratio between the error bars on μ obtained including the CIB monopole uncertainties (that we conservatively take as the 68% upper limit of the recovered μ distribution) and the error bars derived including only the instrumental noise gives us an estimate of the factor that the measured error bars on $C_\ell^{\mu T}$ should be inflated by to account for these effects in the final f_{NL} estimate. This is equal to ~ 3 but can get to ~ 5 for the case where we consider a systematic mean CIB monopole residual in all frequency bands. We also note that the high sensitivity to the CIB monopoles and related potential biases induced by their misestimation appears only when we fit for the μ distortion in the frequency range we are considering. In fact, we do not observe any significant bias in the recovered CMB even when systematic errors are injected if we do not include μ in the data model. Alternative component separation methods working in harmonic domain might be less sensitive to these issues although they might still have an impact through coupling with the analysis mask.

D. Final f_{NL} constraints

With the new band power estimates provided in Sec. VI A we perform a full likelihood analysis of the single power spectrum bin using the pipeline of Sec. V. This includes the full information on the band powers and accounts for the scale dependent shape of $C_\ell^{\mu T}$ and power spectrum binning effects that were neglected in the original KS15 analysis. Including only the statistical uncertainties, we find comparable constraints to what we obtain with FIRAS at the level of $|f_{\text{NL}}| < 3 \times 10^6$ (see the solid gray line in Fig. 11), showing that the original results were underestimated by about an order of magnitude. However, if the uncertainties connected with the removal of the *Planck* HFI maps offsets are accounted for, the constraints significantly degrade and are not competitive with our measurement. Specifically, if we only consider the statistical uncertainty in the estimate of the CIB offset residuals and inflate the error bar on $C_\ell^{\mu T}$ by a factor ~ 3 , the inferred constraint on f_{NL} becomes $|f_{\text{NL}}| < 5.5 \times 10^6$. If we instead also include the systematic mean CIB residual in all frequency bands, and inflate $\Delta C_\ell^{\mu T}$ by about $5\times$, the 2σ upper limit loosens to $|f_{\text{NL}}| < 8.6 \times 10^6$ (dashed gray line in Fig. 11).

VII. CONCLUSIONS

In this paper, we performed the first attempt to self-consistently measure both the monopole and the anisotropic part of the μ -type spectral distortion of the CMB. These provide invaluable information on the thermal history of the Universe, on physics beyond the standard model, and on highly squeezed primordial non-Gaussianities on scales $k \gg 0.05 \text{ Mpc}^{-1}$ that are not constrained by CMB observations both for scalar, tensor, and mixed bispectra of primordial perturbations.

For this purpose, we used the FIRAS spectrometer legacy data to reconstruct an all-sky map of the μ distortion. Being FIRAS sensitive to the absolute sky brightness, we managed to reduce the previous constraint on the monopole of μ distortion by a factor of 2 thanks to a more robust foreground cleaning. We correlated the μ map with CMB temperature and, for the first time, E and B -mode polarization anisotropies measured by the *Planck* satellite on large angular scales $2 \leq \ell \leq 46$. The measured unbinned $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$ cross spectra have been translated to constraints on the amplitude of local-type primordial non-Gaussianity f_{NL} at scales $10^2 \lesssim k \lesssim 10^4 \text{ Mpc}^{-1}$ and to an upper limit on its n_{NL} . Our baseline analysis suggests $|f_{\text{NL}}| < 3.6 \times 10^6$ from the combined analysis of $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$, and $n_{\text{NL}} \lesssim 1.4$ if we assume that $\langle \mu \rangle$ is only sourced by the dissipation of acoustic waves in the primordial plasma. $C_\ell^{\mu B}$, for which we provide the first constraint, was found to be consistent with zero, implying no strong violation of the statistical isotropy of tensor perturbation. We postpone a full analysis and cosmological interpretation of the μB cross-correlation to future work.

We have run an extensive suite of systematic checks to assess the robustness of our results against foreground contaminations and analysis choices. We found that all our results are stable and all the induced shifts are within fluctuations expected from statistical errors. Thanks to their extended spectral coverage and sensitivity to the absolute sky brightness, we show that spectrometers data are very robust against systematic effects. In fact, we reanalyzed previous measurements of $C_\ell^{\mu T}$ based on data from *Planck* and found that they are dominated by effects due to an imperfect component separation due to a reduced frequency coverage and by astrophysical uncertainties. Once all these effects are all taken into account those measurements deliver a constraint comparable to ours or worse. Going forward, spectrometers that will be able to measure the μ distortion monopole accurately are the only way to break the degeneracy with f_{NL} for the interpretation of $C_\ell^{\mu T}$ and $C_\ell^{\mu E}$ in a model independent way. To this end, extending the frequency coverage of spectrometers to $\nu \leq 60$ will be crucial as the signatures of μ distortions are stronger.

We make our maps and results publicly available to the community to explore additional theoretical models.

As this paper was finalized, we became aware of similar work in [93] using the full set of *Planck* channels between 30 GHz and 857 GHz. Their f_{NL} constraint consistently benefits from an improved component separation that reduces foreground contaminations at small scales and allows them to use multipoles $\ell \lesssim 1000$ compared to the previous *Planck* analysis by [33]. In addition to providing the first measurements of $C_\ell^{\mu B}$ and an improved μ distortion monopole constraint, our work is complementary to [93] as it extends the measurements to larger angular scales. We note that the f_{NL} constraints they obtained when considering only angular scales that overlap with this work (which still contribute to $\approx 30\%$ of their sensitivity) are likely optimistic due to a simplistic treatment of the correlation and non-Gaussianity of the noise after component separation, as well as multipole bin-bin correlations in the likelihood. A comparison between the theoretical curves adopted in their analysis and the ones used in this work would also be informative. We expect these effects to push the *Planck* based constraint from large scale closer to ours.

ACKNOWLEDGMENTS

We thank Boris Bolliet, Giovanni Cabass, Tom Crawford, Colin Hill, Reijo Keskitalo, Daan Meerburg, Niels Odegard, Giorgio Orlando, and David Spergel for useful discussions. We thank Niels Odegard and Christopher Anderson for their help with the FIRAS data. We thank Luca Pagano for useful discussion and feedback and his invaluable help with maximum likelihood power spectrum estimation. F.B. would like to thank Simone Aiola and the Center for Computational Astrophysics at Flatiron Institute, where this work was initiated, for hospitality. G.F. acknowledges the support of the European Research Council under the Marie Skłodowska Curie actions through the Individual Global Fellowship No. 892401 PiCOGAMBAS. Part of this research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231. Some of the results in this paper have been derived using the healpy/HEALPix [94,95] and getdist packages [96], and NumPy [97], SciPy [98] and Matplotlib libraries [99].

APPENDIX: UNCERTAINTIES ON THE RECONSTRUCTED MAPS

In Fig. 13 we show the estimated uncertainties on the reconstructed maps in our baseline analysis; ΔT , μ , A_d , and β_d . These uncertainties are calculated as the standard deviation of the posterior, obtained through MCMC sampling, in each pixel.




FIG. 13. Uncertainties on the reconstructed maps calculated as the standard deviation of the posterior in each pixel. The top left plot shows the CMB dipole, the top right refers to the anisotropic μ -type distortion and the bottom ones show the amplitude and spectral index (at 353 GHz) of the Galactic dust. Pixels grayed out are removed by the FIRAS destriper mask.

-
- [1] S. Adachi, M. A. O. Aguilar Faúndez, K. Arnold *et al.*, *Astrophys. J.* **904**, 65 (2020).
- [2] N. Aghanim *et al.*, *Astron. Astrophys.* **641**, A6 (2020).
- [3] S. Aiola *et al.*, *J. Cosmol. Astropart. Phys.* **12** (2020) 047.
- [4] D. Dutcher *et al.*, *Phys. Rev. D* **104**, 022003 (2021).
- [5] D. J. Fixsen, A. Kogut, S. Levin *et al.*, *Astrophys. J.* **734**, 5 (2011).
- [6] J. C. Mather, D. J. Fixsen, and R. A. Shafer, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 2019, Infrared Spaceborne Remote Sensing*, edited by M. S. Scholl (SPIE, San Diego, USA, 1993), pp. 168–179, .
- [7] M. Zannoni, A. Tartari, M. Gervasi, G. Boella, G. Sironi, A. De Lucia, A. Passerini, and F. Cavaliere, *Astrophys. J.* **688**, 12 (2008).
- [8] J. Chluba *et al.*, *Exp. Astron.* **51**, 1515 (2021).
- [9] H. Fu, M. Lucca, S. Galli, E. S. Battistelli, D. C. Hooper, J. Lesgourgues, and N. Schöneberg, *J. Cosmol. Astropart. Phys.* **12** (2021) 050.
- [10] M. Lucca, N. Schöneberg, D. C. Hooper, J. Lesgourgues, and J. Chluba, *J. Cosmol. Astropart. Phys.* **02** (2020) 026.
- [11] C. Burigana, L. Danese, and G. de Zotti, *Astron. Astrophys.* **246**, 49 (1991).
- [12] R. A. Sunyaev and Y. B. Zeldovich, *Astrophys. Space Sci.* **7**, 3 (1970).
- [13] J. Chluba and D. Jeong, *Mon. Not. R. Astron. Soc.* **438**, 2065 (2014).
- [14] W. Hu, D. Scott, and J. Silk, *Astrophys. J. Lett.* **430**, L5 (1994).
- [15] T. Mroczkowski, D. Nagai, K. Basu *et al.*, *Space Sci. Rev.* **215**, 17 (2019).
- [16] A. Refregier, E. Komatsu, D. N. Spergel, and U.-L. Pen, *Phys. Rev. D* **61**, 123001 (2000).
- [17] D. J. Fixsen, E. S. Cheng, J. M. Gales, J. C. Mather, R. A. Shafer, and E. L. Wright, *Astrophys. J.* **473**, 576 (1996).
- [18] P. André *et al.*, *J. Cosmol. Astropart. Phys.* **02** (2014) 006.
- [19] A. Kogut, J. Chluba, D. J. Fixsen, S. Meyer, and D. Spergel, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9904, Space Telescopes and Instrumentation 2016: Optical, Infrared, and Millimeter Wave*, edited by H. A. MacEwen, G. G. Fazio,

- M. Lystrup, N. Batalha, N. Siegler, and E. C. Tong, 99040W (SPIE, Edinburgh, United Kingdom, 2016), .
- [20] B. Maffei *et al.*, in *Proceedings of the 16th Marcel Grossmann Meeting on Recent Developments in Theoretical and Experimental General Relativity, Astrophysics and Relativistic Field Theories* (World Scientific Publishing, 2021), arXiv:2111.00246.
- [21] S. Masi *et al.*, in *Proceedings of the 16th Marcel Grossmann Meeting on Recent Developments in Theoretical and Experimental General Relativity, Astrophysics and Relativistic Field Theories* (World Scientific Publishing, 2021), arXiv:2110.12254.
- [22] J. E. Carlstrom, G. P. Holder, and E. D. Reese, *Annu. Rev. Astron. Astrophys.* **40**, 643 (2002).
- [23] Y. B. Zeldovich and R. A. Sunyaev, *Astrophys. Space Sci.* **4**, 301 (1969).
- [24] E. Pajer and M. Zaldarriaga, *Phys. Rev. Lett.* **109**, 021302 (2012).
- [25] M. Biagetti, H. Perrier, A. Riotto, and V. Desjacques, *Phys. Rev. D* **87**, 063521 (2013).
- [26] G. Cabass, E. Pajer, and D. van der Woude, *J. Cosmol. Astropart. Phys.* **08** (2018) 050.
- [27] R. Emami, E. Dimastrogiovanni, J. Chluba, and M. Kamionkowski, *Phys. Rev. D* **91**, 123531 (2015).
- [28] J. Ganc and E. Komatsu, *Phys. Rev. D* **86**, 023518 (2012).
- [29] G. Orlando, P. D. Meerburg, and S. P. Patil, *J. Cosmol. Astropart. Phys.* **02** (2022) 004.
- [30] A. Ota, *Phys. Rev. D* **94**, 103520 (2016).
- [31] A. Ravenni, M. Liguori, N. Bartolo, and M. Shiraishi, *J. Cosmol. Astropart. Phys.* **09** (2017) 042.
- [32] Y. Akrami *et al.*, *Astron. Astrophys.* **641**, A9 (2020).
- [33] R. Khatri and R. Sunyaev, *J. Cosmol. Astropart. Phys.* **09** (2015) 026.
- [34] M. Remazeilles and J. Chluba, *Mon. Not. R. Astron. Soc.* **478**, 807 (2018).
- [35] M. Remazeilles, A. Ravenni, and J. Chluba, *Mon. Not. R. Astron. Soc.* **512**, 455 (2022).
- [36] S. Mukherjee, J. Silk, and B. D. Wandelt, *Phys. Rev. D* **100**, 103508 (2019).
- [37] N. Aghanim *et al.*, *Astron. Astrophys.* **641**, A1 (2020).
- [38] N. Aghanim *et al.*, *Astron. Astrophys.* **641**, A3 (2020).
- [39] Y. Akrami *et al.*, *Astron. Astrophys.* **641**, A2 (2020).
- [40] Y. Akrami *et al.*, *Astron. Astrophys.* **643**, A42 (2020).
- [41] Y. Akrami *et al.*, *Astron. Astrophys.* **641**, A4 (2020).
- [42] D. J. Fixsen, *Astrophys. J.* **707**, 916 (2009).
- [43] N. Aghanim *et al.*, *Astron. Astrophys.* **596**, A109 (2016).
- [44] R. Adam *et al.*, *Astron. Astrophys.* **594**, A10 (2016).
- [45] B. T. Draine *Annu. Rev. Astron. Astrophys.* **41**, 241 (2003).
- [46] J. Chluba, J. C. Hill, and M. H. Abitbol, *Mon. Not. R. Astron. Soc.* **472**, 1195 (2017).
- [47] B. S. Hensley *et al.*, *Astrophys. J.* **929**, 166 (2022).
- [48] S. Brodd, D. J. Fixsen, K. A. Jensen, J. C. Mather, and R. A. Shafer, *COBE Ref. Pub. No. 97-C* (NASA/GSFC, Greenbelt, MD, 1997).
- [49] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).
- [50] M. H. Abitbol, J. Chluba, J. C. Hill, and B. R. Johnson, *Mon. Not. R. Astron. Soc.* **471**, 1126 (2017).
- [51] S. Vanneste, S. Henrot-Versillé, T. Louis, and M. Tristram, *Phys. Rev. D* **98**, 103526 (2018).
- [52] M. Tegmark, *Phys. Rev. D* **55**, 5895 (1997).
- [53] C. J. Anderson, E. R. Switzer, and P. C. Breysse, *Mon. Not. R. Astron. Soc.* **514**, 1169 (2022).
- [54] N. Odegard, J. L. Weiland, D. J. Fixsen, D. T. Chuss, E. Dwek, A. Kogut, and E. R. Switzer, *Astrophys. J.* **877**, 40 (2019).
- [55] W. J. Percival and M. L. Brown, *Mon. Not. R. Astron. Soc.* **372**, 1104 (2006).
- [56] J. Hartlap, P. Simon, and P. Schneider, *Astron. Astrophys.* **464**, 399 (2007).
- [57] P. Ade, Z. Ahmed, M. Amiri *et al.*, *Phys. Rev. Lett.* **127**, 151301 (2021).
- [58] B. Bolliet, J. Chluba, and R. Battye, *Mon. Not. R. Astron. Soc.* **507**, 3148 (2021).
- [59] V. Poulin, J. Lesgourgues, and P. D. Serpico, *J. Cosmol. Astropart. Phys.* **03** (2017) 043.
- [60] D. Alonso, J. Sanchez, and A. Slosar (LSST Dark Energy Science Collaboration), *Mon. Not. R. Astron. Soc.* **484**, 4127 (2019).
- [61] E. Hivon, K. M. Górski, C. B. Netterfield, B. P. Crill, S. Prunet, and F. Hansen, *Astrophys. J.* **567**, 2 (2002).
- [62] T. Louis, Z. Li, and M. Tristram, *Phys. Rev. D* **100**, 103534 (2019).
- [63] L. G. Book, M. Kamionkowski, and T. Souradeep, *Phys. Rev. D* **85**, 023010 (2012).
- [64] J. M. Delouis, L. Pagano, S. Mottet, J. L. Puget, and L. Vibert, *Astron. Astrophys.* **629**, A38 (2019).
- [65] J. Chluba, E. Dimastrogiovanni, M. A. Amin, and M. Kamionkowski, *Mon. Not. R. Astron. Soc.* **466**, 2390 (2017).
- [66] E. Komatsu and U. Seljak, *Mon. Not. R. Astron. Soc.* **336**, 1256 (2002).
- [67] V. Acquaviva, N. Bartolo, S. Matarrese, and A. Riotto, *Nucl. Phys.* **B667**, 119 (2003).
- [68] J. Maldacena, *J. High Energy Phys.* **05** (2003) 013.
- [69] D. M. Goldberg and D. N. Spergel, *Phys. Rev. D* **59**, 103002 (1999).
- [70] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000).
- [71] G. Cabass, A. Melchiorri, and E. Pajer, *Phys. Rev. D* **93**, 083515 (2016).
- [72] J. Chluba, *Mon. Not. R. Astron. Soc.* **460**, 227 (2016).
- [73] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **106**, 043506 (2022).
- [74] G. D'Amico, M. Lewandowski, L. Senatore, and P. Zhang, arXiv:2201.11518.
- [75] E.-M. Mueller, M. Rezaie, W. J. Percival *et al.*, arXiv:2106.13725.
- [76] C. T. Byrnes, K. Enqvist, S. Nurmi, and T. Takahashi, *J. Cosmol. Astropart. Phys.* **11** (2011) 011.
- [77] C. T. Byrnes, M. Gerstenlauer, S. Nurmi, G. Tasinato, and D. Wands, *J. Cosmol. Astropart. Phys.* **10** (2010) 004.
- [78] X. Chen, *Phys. Rev. D* **72**, 123518 (2005).
- [79] J. Khoury and F. Piazza, *J. Cosmol. Astropart. Phys.* **07** (2009) 026.
- [80] E. Sefusatti, M. Liguori, A. P. Yadav, M. G. Jackson, and E. Pajer, *J. Cosmol. Astropart. Phys.* **12** (2009) 022.
- [81] A. Riotto and M. S. Sloth, *Phys. Rev. D* **83**, 041301 (2011).
- [82] R. Adam *et al.*, *Astron. Astrophys.* **594**, A8 (2016).
- [83] P. A. R. Ade *et al.*, *Astron. Astrophys.* **594**, A6 (2016).

- [84] R. Khatri, *Mon. Not. R. Astron. Soc.* **451**, 3321 (2015).
- [85] A. Challinor and G. Chon, *Mon. Not. R. Astron. Soc.* **360**, 509 (2005).
- [86] G. Chon, A. Challinor, S. Prunet, E. Hivon, and I. Szapudi, *Mon. Not. R. Astron. Soc.* **350**, 914 (2004).
- [87] M. L. Brown, P. G. Castro, and A. N. Taylor, *Mon. Not. R. Astron. Soc.* **360**, 1262 (2005).
- [88] S. Hamimeche and A. Lewis, *Phys. Rev. D* **77**, 103013 (2008).
- [89] R. Makiya, S. Ando, and E. Komatsu, *Mon. Not. R. Astron. Soc.* **480**, 3928 (2018).
- [90] E. Keihänen, R. Keskitalo, H. Kurki-Suonio, T. Poutanen, and A. S. Sirviö, *Astron. Astrophys.* **510**, A57 (2010).
- [91] P. A. R. Ade *et al.*, *Astron. Astrophys.* **571**, A8 (2014).
- [92] M. Béthermin, E. Daddi, G. Magdis *et al.*, *Astrophys. J. Lett.* **757**, L23 (2012).
- [93] A. Rotti, A. Ravenni, and J. Chluba, [arXiv:2205.15971](https://arxiv.org/abs/2205.15971).
- [94] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005).
- [95] A. Zonca, L. Singer, D. Lenz, M. Reinecke, C. Rosset, E. Hivon, and K. Gorski, *J. Open Source Software* **4**, 1298 (2019).
- [96] A. Lewis, [arXiv:1910.13970](https://arxiv.org/abs/1910.13970).
- [97] C. R. Harris *et al.*, *Nature (London)* **585**, 357 (2020).
- [98] P. Virtanen *et al.*, *Nat. Methods* **17**, 261 (2020).
- [99] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).