



HAL
open science

Unexpected Topology of the Temperature Fluctuations in the Cosmic Microwave Background

Pratyush Pranav, Robert J. Adler, Thomas Buchert, Herbert Edelsbrunner,
Bernard J.T. Jones, Armin Schwartzman, Hubert Wagner, Rien van de
Weygaert

► **To cite this version:**

Pratyush Pranav, Robert J. Adler, Thomas Buchert, Herbert Edelsbrunner, Bernard J.T. Jones, et al.. Unexpected Topology of the Temperature Fluctuations in the Cosmic Microwave Background. *Astronomy & Astrophysics - A&A*, 2019, 627, pp.A163. 10.1051/0004-6361/201834916 . hal-01975230

HAL Id: hal-01975230

<https://hal.science/hal-01975230v1>

Submitted on 21 Aug 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Unexpected topology of the temperature fluctuations in the cosmic microwave background

Pratyush Pranav^{1,2}, Robert J. Adler², Thomas Buchert¹, Herbert Edelsbrunner³, Bernard J. T. Jones⁵,
Armin Schwartzman⁴, Hubert Wagner³, and Rien van de Weygaert⁵

¹ Univ. Lyon, ENS de Lyon, Univ. Lyon 1, CNRS, Centre de Recherche Astrophysique de Lyon UMR5574, 69007 Lyon, France
e-mail: pratyuze@gmail.com

² Technion – Israel Institute of Technology, 32000 Haifa, Israel

³ IST Austria (Institute of Science and Technology Austria), 3400 Klosterneuburg, Austria

⁴ Division of Biostatistics, University of California, San Diego, CA, USA

⁵ Kapteyn Astronomical Institute, Landleven 12, 9747 AG Groningen, The Netherlands

Received 18 December 2018 / Accepted 27 April 2019

ABSTRACT

We study the topology generated by the temperature fluctuations of the cosmic microwave background (CMB) radiation, as quantified by the number of components and holes, formally given by the Betti numbers, in the growing excursion sets. We compare CMB maps observed by the *Planck* satellite with a thousand simulated maps generated according to the Λ CDM paradigm with Gaussian distributed fluctuations. The comparison is multi-scale, being performed on a sequence of degraded maps with mean pixel separation ranging from 0.05 to 7.33° . The survey of the CMB over \mathbb{S}^2 is incomplete due to obfuscation effects by bright point sources and other extended foreground objects like our own galaxy. To deal with such situations, where analysis in the presence of “masks” is of importance, we introduce the concept of relative homology. The parametric χ^2 -test shows differences between observations and simulations, yielding p -values at percent to less than permil levels roughly between 2 and 7° , with the difference in the number of components and holes peaking at more than 3σ sporadically at these scales. The highest observed deviation between the observations and simulations for b_0 and b_1 is approximately between 3σ and 4σ at scales of 3 – 7° . There are reports of mildly unusual behaviour of the Euler characteristic at 3.66° in the literature, computed from independent measurements of the CMB temperature fluctuations by *Planck*'s predecessor, the *Wilkinson Microwave Anisotropy Probe* (WMAP) satellite. The mildly anomalous behaviour of the Euler characteristic is phenomenologically related to the strongly anomalous behaviour of components and holes, or the zeroth and first Betti numbers, respectively. Further, since these topological descriptors show consistent anomalous behaviour over independent measurements of *Planck* and WMAP, instrumental and systematic errors may be an unlikely source. These are also the scales at which the observed maps exhibit low variance compared to the simulations, and approximately the range of scales at which the power spectrum exhibits a dip with respect to the theoretical model. Non-parametric tests show even stronger differences at almost all scales. Crucially, Gaussian simulations based on power-spectrum matching the characteristics of the observed dipped power spectrum are not able to resolve the anomaly. Understanding the origin of the anomalies in the CMB, whether cosmological in nature or arising due to late-time effects, is an extremely challenging task. Regardless, beyond the trivial possibility that this may still be a manifestation of an extreme Gaussian case, these observations, along with the super-horizon scales involved, may motivate the study of primordial non-Gaussianity. Alternative scenarios worth exploring may be models with non-trivial topology, including topological defect models.

Key words. cosmic background radiation – early Universe – methods: statistical – methods: numerical

1. Introduction

The Λ cold dark matter (or Λ CDM) standard paradigm of cosmology postulates that the Universe consists primarily of cold non-relativistic dark matter, which reveals its presence only through gravitational interactions, and the Universe is currently driven by dark energy, causing accelerated volume expansion in this model. The cosmic microwave background (CMB) radiation, which originates at the epoch of recombination, is the most important observational probe into the validity of the standard paradigm today (Jones 2017). It is the earliest visible light and offers a glimpse into the processes during the nascent stage of the Universe. Fluctuations about the mean in the temperature field of the CMB correspond to the fluctuations in the distribution of matter in the early Universe. Understanding the CMB is therefore crucial to understanding the primordial Universe.

The Λ CDM paradigm together with the inflationary theories in their simplest forms, predict the primordial perturbations to be realizations of a homogeneous and isotropic Gaussian random field (Guth & Pi 1982). This hypothesis is supported experimentally by CMB observations (Smoot et al. 1992; Bennett et al. 2003; Spergel et al. 2007; Komatsu et al. 2011; Planck Collaboration XIII 2016) and theoretically by the central limit theorem. While it has largely been agreed upon that the CMB exhibits characteristics of a homogeneous and isotropic Gaussian field, there are lingering doubts. The pioneering works of Eriksen et al. (2004a) and Park (2004) challenge the assumption of homogeneity, and the alignment of low multipoles (Copi et al. 2015) challenges the assumption of isotropy. Other noted anomalies include the vanishing correlation function at large scales, and the unusually low variance at approximately 3° ; see Schwarz et al. (2016) for a review and possible

interpretations. [Planck Collaboration XXIII \(2014\)](#) independently confirms these anomalies.

The primordial non-Gaussianity remains a topic of ongoing debate. Deviations from Gaussianity, if found, will point to new physics driving the Universe in its nascent stages. The consensus is in relative favour of the absence of non-Gaussianity ([Komatsu et al. 2011](#); [Planck Collaboration XXIII 2014](#); [Planck Collaboration XXIV 2014](#); [Matsubara 2010](#); [Bartolo et al. 2010](#)); see also [Buchert et al. \(2017\)](#) for a review and a model-independent route of analysis. Despite mildly unusual behaviour of the Euler characteristic, pointed out in [Eriksen et al. \(2004b\)](#) and [Park \(2004\)](#), the methods employed until today have not provided compelling evidence of non-Gaussianity in the CMB. In contrast, the topological methods of this paper find the observed CMB maps ([Planck Collaboration IX 2016](#)) to be significantly different from the Full Focal Plane 8 (FFP8) simulations ([Planck Collaboration XXIII 2014](#); [Planck Collaboration XII 2016](#)) that assume the initial perturbations to be Gaussian.

Topology is the branch of mathematics concerned with properties of shapes and spaces preserved under continuous deformations, such as stretching and bending, but not tearing and gluing. It is related to but different from geometry, which measures size and shape. Both geometry and topology have been used in the past to study the structure of the CMB radiation and other cosmic fields. Historically, the predominant tools in this endeavour were the Minkowski functionals, which for a 2-manifold embedded in the three-dimensional space are related to the enclosed volume, the area, the total mean curvature, and the total Gaussian curvature. By the Gauss–Bonnet Theorem, for 2-manifolds, the latter is 2π times the Euler characteristic ([Euler 1758](#)), thus providing a bridge between geometry and topology. Early topological studies of the cosmic mass distribution were based on the Euler characteristic of the iso-density surfaces, which generically are 2-manifolds ([Doroshkevich 1970](#); [Bardeen et al. 1986](#); [Gott et al. 1986](#); [Park et al. 2013](#)). The full set of Minkowski functionals was later introduced to cosmology in [Mecke et al. \(1994\)](#), [Schmalzing & Buchert \(1997\)](#), [Sahni et al. \(1998\)](#), [Schmalzing & Gorski \(1998\)](#). For Gaussian, and Gaussian-related random fields, the expected values of the Minkowski functionals of excursion sets have known analytic expressions ([Adler 1981](#); [Adler & Taylor 2010](#)), which is one of the main reasons they have played a key role in the study of real valued fields arising in cosmology and other disciplines. Analyses based on Minkowski functionals have been used for predicting and quantifying the presence of non-Gaussianity in the CMB maps obtained with the *Planck* satellite ([Ducout et al. 2013](#); [Buchert et al. 2017](#)).

While the Minkowski functionals have been instructive, the topological information contained in them is limited and convolved with geometric information. Moreover, they are not equipped to address the hierarchical aspects of the matter distribution directly, although partial Minkowski functionals ([Schmalzing et al. 1999](#)) may be useful in certain settings. We therefore analyse CMB fluctuations in terms of the purely topological concepts of homology ([Munkres 1984](#)), as quantified by Betti numbers ([Betti 1871](#)) and persistence ([Edelsbrunner et al. 2002](#); [Edelsbrunner & Harer 2010](#)). Following the Euler–Poincaré formula, the Euler characteristic is the alternating sum of the Betti numbers, implying that the latter provide a finer description of topology ([Munkres 1984](#)). A broad exposition of these concepts in a cosmological setting is given in [Pranav et al. \(2017\)](#), [Pranav \(2015\)](#), [van de Weygaert et al. \(2011\)](#); also see [Park et al. \(2013\)](#), [Sousbie \(2011\)](#), [Shivashankar et al. \(2016\)](#), [Adler et al. \(2017\)](#), [Makarenko et al. \(2018\)](#), [Cole & Shiu \(2018\)](#) for some applica-

tions. Related but slightly different methodologies used for the analysis of cosmological datasets, emanating from concepts in Morse theory, maybe found in [Colombi et al. \(2000\)](#), [Novikov et al. \(2006\)](#), [Sousbie et al. \(2008\)](#).

Our main result is an anomaly of the observed CMB radiation when compared with simulations based on Gaussian prescriptions. The χ^2 -test yields a significant difference between the number of components and holes in the observed sky compared to the simulations, with p -values at percent to less than permil levels at scales of roughly 2 – 7° . The differences peak sporadically at more than 3σ at these scales. Non-parametric tests reveal an even more significant difference between the observation and the simulations at almost all scales. The χ^2 -test shows the anomaly at roughly the same scales at which the power spectrum exhibits a dip. [Eriksen et al. \(2004b\)](#) reports a mildly unusual Euler characteristic at approximately 3° in the earlier measurements of the CMB radiation by the *Wilkinson Microwave Anisotropy Probe* (WMAP) satellite, which is related to the anomalous behaviour of components and holes. The noted anomaly motivates a closer look at the standard paradigm. Possible scenarios include but are not limited to primordial non-Gaussianity, topological defect models, and models with non-trivial topology ([Bouchet et al. 2002](#); [Aurich & Steiner 2001](#); [Aurich et al. 2007](#); [Bernui et al. 2018](#)).

The workflow in this paper is as follows. Topological descriptors are computed from cosmology data, and statistical tests based on these descriptors are used to compare the observations with simulations. Section 2 gives a summary of the topological concepts. Section 3 describes the data, the computational pipeline, and a brief account of the statistical tests employed. Section 4 presents the main results of the paper, followed by a summary and conclusions in Sect. 5.

2. Topological background

Since the CMB radiation is observed as a scalar field on the two-dimensional sphere, the topological concepts needed in this paper are elementary, namely the components and the holes of subsets of this sphere. To count them in the presence of regions with unreliable data, we compute the ranks of the homology groups relative to the mask that covers these regions.

2.1. Excursion sets and absolute homology

Writing \mathbb{S}^2 for the two-dimensional sphere and $f: \mathbb{S}^2 \rightarrow \mathbb{R}$ for the temperature field of the CMB, the excursion set at a temperature ν is the subset of the sphere in which the temperature is ν or larger: $\mathbb{E}(\nu) = \{x \in \mathbb{S}^2 \mid f(x) \geq \nu\}$. It is a closed set, and we write $\beta_0(\nu)$ for its number of components. A hole is a component of the complement, $\mathbb{S}^2 \setminus \mathbb{E}(\nu)$. Assuming there is at least one hole, we write $\beta_1(\nu) + 1$ for the number of holes, and we set $\beta_2(\nu) = 0$, because $\mathbb{E}(\nu)$ does not cover the entire sphere. On the other hand, if there is no hole, we set $\beta_1(\nu) = 0$ and $\beta_2(\nu) = 1$; see the left panel of Fig. 1 for an illustration. These definitions are motivated by the more general theory ([Munkres 1984](#)) in which the p -th Betti number is the rank of the p -th homology group: $\beta_p = \text{rank } H_p$ for $p = 0, 1, 2$. These are the basic objects of homology. The Euler characteristic of the excursion set is the alternating sum of the Betti numbers: $\text{EC}(\nu) = \beta_0(\nu) - \beta_1(\nu) + \beta_2(\nu)$.

The Euler characteristic has a long history in the CMB literature, largely due to the fact that a simple analytic formula for its expected value is known when the CMB is modelled as




Fig. 1. *Left:* blue excursion set on the sphere consisting of an upper-left component with a hole, an upper-right component, and a lower component. Its Betti numbers are $\beta_0 = 3, \beta_1 = 1, \beta_2 = 0$, and its Euler characteristic is $EC = 3 - 1 + 0 = 2$. *Middle:* pink mask in which the data are not reliable. The mask covers part of the upper-left component and hole; its hole is fully contained in the upper-right component, and it overlaps the lower component in two disconnected pieces. *Right:* visualization of the relative homology groups obtained by shrinking the mask to a point and pulling the excursion set with it. We have $b_0 = 0$ because all three components connect to the shrunken mask, $b_1 = 2$ because the loop in the upper-left component is preserved and a new loop in the lower component is formed, and $b_2 = 1$ because the upper-right component takes on the shape of a sphere. The (relative) Euler characteristic is therefore $EC_{rel} = 0 - 2 + 1 = -1$.

a Gaussian random field (Adler 1981; Adler & Taylor 2010). While such formulas are not known for the Betti numbers, the information they carry is richer. Pranav et al. (2019) present a numerical study of the Betti numbers of Gaussian random fields, and compare them to the Euler characteristic and Minkowski functionals, and find that Betti numbers present a more detailed account of the topological properties of the field compared to the Euler characteristic (also see Park et al. 2013). In general, near the mean level of ν , one expects the components and holes of the excursion set to be of similar size and number. Accordingly, one expects $\beta_0(\nu)$ and $\beta_1(\nu)$ to be of similar magnitude, combining to give an Euler characteristic close to zero. Such an Euler characteristic tells us nothing about the individual Betti numbers beyond the fact that they are similar.

2.2. Masks and relative homology

We define the mask to be the region in which the data is not reliable, and denote it by $\mathbb{M} \subseteq \mathbb{S}^2$. In our application, the mask includes a belt around the equator corresponding to the thickened disc of the Milky Way, along with other galactic and extragalactic bright foreground objects that interfere with the observation of the CMB radiation. In an effort to exclude the mask from our computations, we consider the reduced excursion set: $\mathbb{E}(\nu) \setminus \mathbb{M}$. Treating \mathbb{M} as a closed set, this difference is not necessarily closed. An appropriate topological measure is the relative homology of a pair of closed spaces, (E, M) , with the second being contained in the first. In our setting, the pair is $E = \mathbb{E}(\nu)$ and $M = \mathbb{M} \cap \mathbb{E}(\nu)$. Just as in the absolute case, we get relative homology groups in dimensions 0, 1, and 2, and we use their ranks for quantification. It is tempting to refer to these ranks as relative Betti numbers, but this is not the traditional terminology, and we simply write $b_p = \text{rank } H_p(E, M)$ for $p = 0, 1, 2$. If $M = \emptyset$, then $b_p = \beta_p$, for all three choices of p , but if the mask overlaps with the excursion set, then there are differences. We explain some of these differences with reference to Fig. 1: If \mathbb{M} overlaps with a component of E , this component is no longer counted because every vertex in it bounds a path connecting it to the mask. If \mathbb{M} overlaps with the component in two disconnected pieces, we count a new loop, namely the path connecting these two pieces. If \mathbb{M} covers part of a hole, this hole is still counted




Fig. 2. Small section of the sphere of directions, with the temperature field visualized by the green landscape that complements the blue mask drawn at lower altitude. We see one closed loop surrounding a relative depression of the temperature field, and two open loops connecting points in the mask along locally highest paths. The visualization is based on the observed CMB maps cleaned using the NILC technique, and smoothed at 4° .

because the part of its boundary curve outside the mask is open, with endpoints in \mathbb{M} . If a hole of \mathbb{M} is contained in the excursion set, we get a surface without a boundary. The relation between absolute and relative homology is compactly expressed by the exact sequence of the pair $M \subseteq E$ (Munkres 1984):

$$0 \rightarrow H_2(M) \rightarrow H_2(E) \rightarrow H_2(E, M) \rightarrow H_1(M) \rightarrow H_1(E) \rightarrow H_1(E, M) \rightarrow H_0(M) \rightarrow H_0(E) \rightarrow H_0(E, M) \rightarrow 0. \quad (1)$$

Briefly, this means that we can assign non-negative integers to the arrows, meaning that the rank of each group is the sum of integers assigned to its incoming and outgoing arrows. For example, in Fig. 1, we have $0 \rightarrow 0 \rightarrow 0 \rightarrow 1 \rightarrow 1 \rightarrow 1 \rightarrow 2 \rightarrow 4 \rightarrow 3 \rightarrow 0 \rightarrow 0$, and it is easy to find the assignment of integers that satisfies the stated property.

2.3. Variationally maximal loops

When we count $\beta_1 + 1$ holes in absolute homology, we in fact count β_1 loops needed to separate them. In relative homology, the connection is not as intuitive because we also have open loops, whose endpoints lie in the mask; see Fig. 2. Generally, there are uncountably many ways to draw a loop, and in homology they are all considered equivalent. The set of equivalent loops is called a homology class, and any one of the loops in the class is a representative. These classes are the elements of the one-dimensional homology group, which is a vector space. The rank of this group counts the classes that are needed to span the vector space.

For visualization, it is desirable to have a unique representative for each class. Similar to the intuitive notion of the rim of a crater, we choose this representative to be as high as possible, alternating between peaks and saddles of f which it connects via ridges within the reduced excursion set. We refer to this loop as the variationally maximal representative of its class; see Fig. 2 for an example. While constructing variational maxima for smooth scalar fields may be problematic, the persistence




Fig. 3. Facets of the rhombic dodecahedron which serve as patches in the HealPix representation of the sphere. In sequence, we show the 12 patches decomposed into 1, 4, 16, and 64 pixels. The final representation is obtained by central projection of the pixel centres and a distortion yielding an approximately equal-area decomposition of the sphere (not shown).




Fig. 4. UT78 mask released by the *Planck* team. It is a conservative mask, that masks the known point sources and other bright foreground objects, in addition to the galactic disc.

algorithm applied to a piecewise linear scalar field produces them as a byproduct of reducing the boundary matrix; see Sect. 3.2.

3. Data and methods

3.1. Data

Decades after the accidental discovery of the CMB, its first space-based observational probe was carried out by the CMB Explorer (COBE) satellite (Smoot et al. 1992), establishing that the CMB is a perfect black-body radiation. Later, the WMAP was launched to study the temperature fluctuations in greater detail (Spergel et al. 2007). Most recently, the high-precision *Planck* mission was launched, measuring temperature fluctuations to an accuracy of 10^{-5} degrees (Planck Collaboration I 2014), and at a resolution of five arc-minutes, giving the most detailed and precise measurement of CMB temperature fluctuations currently available. We use the *Planck* maps for our analyses (Planck Collaboration IX 2016; Planck Collaboration XII 2016).

Format. The CMB sky maps are presented in the HealPix format (Górski et al. 2005), which is an equal-area pixelisation of the sphere, which we denote as \mathbb{S}^2 ; see Fig. 3. Using the faces of the rhombic dodecahedron, we start by decomposing the sphere into twelve patches. Fine resolution is achieved by dividing these patches into N^2 equal area pixels each, meaning that the total number of pixels at this resolution is $12 \times N^2$. At maximum resolution $N = 2048$, the maps have about 50 million pixels.

Observed sky. The *Planck* satellite observes the sky at seven different frequency bands, leading to component-separated maps using four different techniques: Commander-Ruler (C-R), NILC, SEVEM, and SMICA; cf. Planck Collaboration IX (2016). These are the publicly available maps from *Planck* data release 2 (PR2-2015)¹. We use these component-separated CMB maps throughout. These maps are contaminated by noise from various sources, including inherent detector noise, and efforts by the *Planck* team to denoise the data have not been completely successful. Consequently, our analysis is performed on the maps produced by combining the CMB and noise maps for each realisation of the simulation. This is a fairly simple task, given that the map-making exercise is linear in nature:

$$f_{\text{final}} = f_{\text{CMB}} + f_{\text{noise}}. \quad (2)$$

Simulations. In addition to the observed data, the *Planck* team released a set of Full Focal Plane 8 (or FFP8) simulations (Planck Collaboration XII 2016) of both the CMB and noise. We use 1000 NILC simulations for our computational experiments. These simulations assume that the CMB is a Gaussian random field, consistent with the null hypothesis of Gaussianity for the CMB, which is what we wish to check, and we use them to estimate the error-bars for testing the significance of differences between observed and simulated maps. Important to note is that these simulations include the effects of realistic foreground models for gravitational lensing, Rayleigh scattering, and more (Planck Collaboration XII 2016, Sect. 3.3.1).

Degradation. In order to perform a scale-dependent analysis of the CMB maps, we degrade them to resolutions between $N = 1024$ and 8, dividing N by two from one level to the next. The process of degradation amounts to decomposing them into spherical harmonics on the full sky at the input resolution. The spherical harmonics coefficients a_{lm} are then convolved to the new resolution using the formula (Planck Collaboration XXIII 2014):

$$a_{lm}^{\text{out}} = \frac{b_l^{\text{out}} p_l^{\text{out}}}{b_l^{\text{in}} p_l^{\text{in}}} a_{lm}^{\text{in}}, \quad (3)$$

where, b_l is the beam transfer function, p_l is the pixel window function, and in and out denote the input and the output functions at the different resolutions, respectively. These are then synthesised into maps at the output resolution directly.

Masks. The observation of the CMB by the *Planck* satellite is incomplete in some regions of the sky, typically as a result of interference from bright foreground objects such as our own galactic disc and bright point sources. In these regions, the CMB sky map is reconstructed as a constrained Gaussian field. In order to avoid these areas in the analysis, we use the most conservative UT78 mask released by the *Planck* team; see Fig. 4. This mask is a combination of all foreground objects with the least sky coverage and therefore leads to a conservative analysis. The mask is a binary map, where reliable pixels of the CMB map are marked by the value 1, and the unreliable parts by 0.

¹ Available at <https://pla.esac.esa.int/>




Fig. 5. Degraded masks before binarization. For high-enough resolutions, the masks have a similar appearance to the original one, but are distinguishable when zooming into the image.

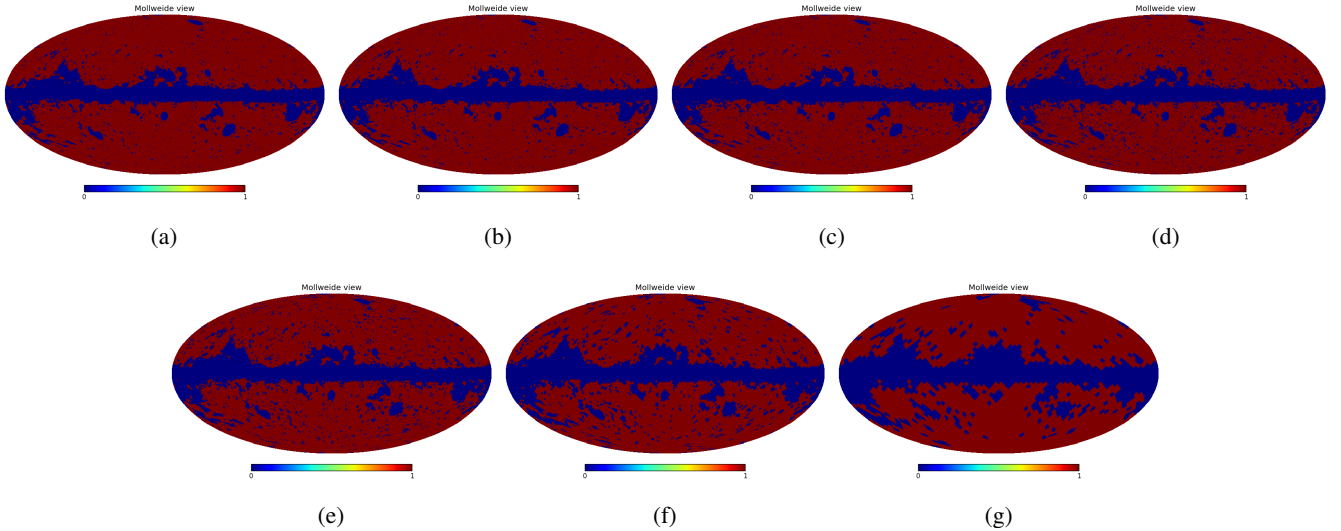


Fig. 6. Degraded masks after binarization, thresholded at 0.9.

Table 1. Percentage of sky area covered by the unmasked regions for the various degraded resolutions between $N = 1024$ and 16, for mask binarization threshold 0.9.

Resolution	% Unmasked
1024	77.19
512	76.52
256	75.50
128	73.37
64	72.41
32	69.39
16	66.24

For the scale-dependent analysis, we also degrade the masks, so that the map and the mask have the same resolution. Degrading the original binary UT78 mask converts it into a non-binary mask in a thickened zone at the boundary separating the reliable part of the mask from the non-reliable part. Figure 5 presents

these yet-to-be-binarized masks. To re-convert them to binary masks, we set a range of binarization thresholds for our experiments: 0.7, 0.8, 0.9, 0.95. Pixels with values above or equal to the binarization threshold are marked as 1, and the rest as 0. Figure 6 presents the binarized maps at various degraded resolutions, for binarization threshold 0.9. Table 1 presents the percentage of sky that is useable for analysis after masking at various resolutions for this threshold. The percentage of usable area drops with decreasing resolution, with only 66% for $N = 16$. Figure 7 presents a visualization of the degraded and masked maps for all the resolutions analysed in this paper in the Mollweide projection view.

3.2. Computational pipeline

The computational pipeline is tailored specifically to the *Planck* data. The preprocessing step involves converting the CMB maps given in absolute units to a dimensionless unit, corrected for mean and scaled by the standard deviation (computed using the non-masked pixels only). We use the HealPix package for the preprocessing step. The output of this operation is the




Fig. 7. Visualization of the masked maps at various degraded resolutions.

normalized temperature values on $12N^2$ pixels, along with their coordinates on the sphere, which is the input to subsequent steps, which we discuss in five sections: (i) triangulating the surface of the sphere with the pixel centres as vertices, (ii) sorting the vertices, edges, and triangles to form an upper-star filtration, (iii) computing the persistence in terms of a reduced boundary matrix, (iv) computing the ranks of the relative homology groups b_p , $p = 0, 1, 2$, and (v) computing the variationally optimal loops from the reduced matrix. The software is written in C++ and designed to handle very large data sets².

Triangulation. The HealPix format stores the data in twelve square arrays of N^2 pixels each, with $N = 2048$ at the finest resolution. The centres of these pixels are points on the faces of a rhombic dodecahedron. With central projection, these points are mapped to the 2-sphere and distorted to achieve an approximately equal-area decomposition. Taking the convex hull of these points in \mathbb{R}^3 , we get a convex polytope whose boundary is homeomorphic to the sphere. Most of the faces will be triangles, and the occasional faces with $k \geq 4$ sides can be subdivided into $k - 2$ triangles to obtain a triangulation of the sphere. This triangulation is the input to all the downstream computations; consisting of $V = 12N^2$ vertices, $3V - 6$ edges, and $2V - 4$ triangles, it represents the temperature field, $f: \mathbb{S}^2 \rightarrow \mathbb{R}$, by storing the temperature value at every vertex. We implicitly assume a piecewise linear interpolation along the edges and the triangles. Figure 8 illustrates such a triangulation, using colours to visualize the temperature field. We use the CGAL library (The CGAL Project 2018) to implement the triangulation.

Upper-star filtration. Given a triangulation K of \mathbb{S}^2 , let $K(\nu) \subseteq K$ contain all simplices (vertices, edges, and triangles) whose temperature values are ν or larger. We use $K(\nu)$ as a proxy for $\mathbb{E}(\nu)$, the corresponding excursion set. Indeed, because of the linear interpolation, there is a deformation retraction from $\mathbb{E}(\nu)$ to $K(\nu)$ (Edelsbrunner & Harer 2010), which implies that the two have corresponding components and holes. To process the sequence of excursion sets, it makes sense to sort the vertices of K in the order of decreasing temperature value. More precisely, we order the simplices of K such that σ precedes τ if (i) $f(\sigma) > f(\tau)$ or (ii) $f(\sigma) = f(\tau)$ and $\dim \sigma < \dim \tau$, in which

$f(\sigma)$ is the minimum temperature value of the one, two, or three vertices of σ . The remaining ties are broken arbitrarily. Assuming any two vertices have different temperature values, then the edges and triangles that immediately follow a vertex are exactly the ones in the upper star of that vertex. We therefore refer to any ordering that satisfies (i) and (ii) as an upper-star filter of K and f . The corresponding upper-star filtration consists of all prefixes of the filter, each representing an excursion set. This filtration is instrumental in computing the persistence of components and holes.

Computing persistence. Given an upper-star filter of the piecewise linear temperature field, there is optimised software available to compute its persistence (Bauer et al. 2014). We base our persistence computation on an adaptation of the software. This software is a sophisticated implementation of the basic algorithm, which we now describe and modify to obtain the variationally optimal loops. We write $\sigma_1, \sigma_2, \dots, \sigma_n$ for the simplices in the triangulation of the sphere, sorted into an upper-star filter. Let $\partial[1 \dots n, 1 \dots n]$ be the corresponding ordered boundary matrix, with $\partial[i, j] = 1$, if σ_i is a face of σ_j and $\dim \sigma_i = \dim \sigma_j - 1$, and $\partial[i, j] = 0$, otherwise. This matrix is sparse and stored as such. The standard persistence algorithm reduces the matrix from left to right. To reduce column j , we subtract columns to the left of j with the goal to move the lowest 1 in column j higher or eliminate it altogether. We use modulo 2 arithmetic, and therefore subtracting is the same as adding: $1 - 1 = 1 + 1 = 0$. We refer to column j as reduced if it is zero or its lowest 1 has only zeros in the same row to its left. We modify the standard algorithm by continuing the reduction even if the lowest 1 can no longer be changed, referring to the final result as fully reduced. To be unambiguous, we explain this algorithm in pseudo-code, where we write $\text{pivot}(j)$ for the row index of the lowest 1 in column j .

```

for  $j = 1$  to  $n$  do
  while  $\exists k < j$  with  $\partial[\text{pivot}(k), j] = 1$  do
    add column  $k$  to column  $j$ 
  endwhile
endfor.

```

The running time of this algorithm is cubic in the number of simplices in the worst case, but the available optimised software is typically much faster.

² All codes, analysed data, and results available from the corresponding author on request.




Fig. 8. Visualization of the temperature field for the NILC observed maps at $N = 16$. Also visible is the corresponding triangulation, for which the pixel centres of the maps serve as the vertices. The temperature values are stored in the vertices of this triangulation.

Ranks of relative homology groups. For computing the ranks of the homology groups relative to the mask, we set the vertices belonging to the mask at $+\infty$, and consider the complex \mathbb{M} induced by the union of these vertices. This mask is closed by definition. We then compute the filtration and persistence diagram corresponding to absolute homology of $\mathbb{E} \cup \mathbb{M}$. Writing $Dgm_p(\mathbb{E} \cup \mathbb{M})$ for the p -dimensional persistence diagram, and recalling that each diagram consists of intervals with real birth and death values, $b > d$, we obtain the ranks of homology groups relative to the mask:

$$\begin{aligned} b_0 &= \#\{[b, d) \in Dgm_0(\mathbb{E} \cup \mathbb{M}) \mid +\infty > b \geq \nu > d\}; \\ b_1 &= \#\{[b, d) \in Dgm_0(\mathbb{E} \cup \mathbb{M}) \mid +\infty = b > d \geq \nu\} \\ &\quad + \#\{[b, d) \in Dgm_1(\mathbb{E} \cup \mathbb{M}) \mid +\infty > b \geq \nu > d\}; \\ b_2 &= \#\{[b, d) \in Dgm_1(\mathbb{E} \cup \mathbb{M}) \mid +\infty = b > d \geq \nu\} \\ &\quad + \#\{[b, d) \in Dgm_2(\mathbb{E} \cup \mathbb{M}) \mid +\infty > b \geq \nu > d\}. \end{aligned} \quad (4)$$

For computing absolute homology, we set the mask pixels at $-\infty$ and consider the union of such vertices as the mask, which is open by definition.

3.3. Statistical tests

The data consist of topological summaries (b_0, b_1, EC_{rel}) obtained from 1000 simulations, as well as of the observed CMB field, processed according to the NILC scheme. The goal is to estimate the probability that the physical model that produced the simulations produces quantities consistent with those from the observed CMB field. Let $\mathbf{x}_i \in \mathbb{R}^m$, $i = 1, \dots, n$, be a sample of i.i.d. m -dimensional vectors drawn from a distribution F . Let $\mathbf{y} \in \mathbb{R}^m$ be another sample point, assumed to be drawn from a distribution G . We wish to test the (null) hypothesis that $F = G$, and give the test results in terms of p -values, which compute the probability that \mathbf{y} is “consistent” with this hypothesis. We consider two methods of testing for statistical consistency. The first is a parametric test based on the Mahalanobis distance

(Mahalanobis 1936), also known as the χ^2 -test. The second is a non-parametric test based on the Tukey depth (Tukey 1975). The χ^2 -test is more standard but has the disadvantage of assuming that the compared quantities follow a Gaussian distribution, while the Tukey depth works without any assumption on the distribution.

Mahalanobis distance or χ^2 -test. Let

$$\bar{\mathbf{x}} = \sum_{i=1}^n \mathbf{x}_i / n$$

and

$$\mathbf{S} = \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T / (n - 1),$$

the sample mean and covariance matrix of the sample $\mathbf{x}_1, \dots, \mathbf{x}_n$, respectively. The squared Mahalanobis distance of \mathbf{y} to $\bar{\mathbf{x}}$ is then

$$d_{\text{Mahal}}^2(\mathbf{y}) = (\mathbf{y} - \bar{\mathbf{x}})^T \mathbf{S}^{-1} (\mathbf{y} - \bar{\mathbf{x}}). \quad (5)$$

If F is assumed to be Gaussian and n is large, then under the hypothesis that $G = F$ the squared Mahalanobis distance (5) is approximately distributed as a χ^2 distribution with m° of freedom. Thus the corresponding p -value is

$$p_{\text{Mahal}}(\mathbf{y}) = P[\chi_m^2 > d_{\text{Mahal}}^2(\mathbf{y})]. \quad (6)$$

Tukey depth. As shown in the data analysis below, the distribution F does not always conform to elliptical contours and therefore may not be Gaussian. In such a setting, p -values computed using the Mahalanobis distance may not be reliable.

The Tukey half-space depth provides a general metric for identifying outliers in a flexible manner and in a non-parametric setting. Take \mathbf{x}_i , $i = 1, \dots, n$ and \mathbf{y} as above, making no assumptions on the structure of F and G , and let \mathbf{z} be any point in \mathbb{R}^m . Then the half-space depth $d_{\text{dep}}(\mathbf{z}; \mathbf{x}_1, \dots, \mathbf{x}_n)$ of \mathbf{z} within the sample of the \mathbf{x}_i is the smallest fraction of the n points $\mathbf{x}_1, \dots, \mathbf{x}_n$ to either side of any hyperplane passing through \mathbf{z} . By definition, the half-space depth is a number between 0 and 0.5. Points that have the same depth constitute a non-parametric estimate of the isolevel contour of the distribution F .

To evaluate a p -value for \mathbf{y} , we first compute $d_j = d_{\text{dep}}(\mathbf{x}_j; \mathbf{x}_1, \dots, \mathbf{x}_n)$ for every point \mathbf{x}_j , $j = 1, \dots, n$, yielding an empirical distribution of depth. The p -value is then computed as the proportion of points whose depth is lower than that of \mathbf{y} :

$$p_{\text{dep}}(\mathbf{y}) = \#\{j \mid d_j > d_{\text{dep}}(\mathbf{y})\} / n. \quad (7)$$

We note that by construction the depth p -value increases in units of $1/n$. For computing half-space depths below, we use the R package `depth`.

4. Results

We use the *Planck* maps for our analyses, which measure fluctuations about the mean in the CMB temperature to an accuracy of 10^{-5} K (Planck Collaboration I 2014). Our primary resources for the comparison between the observations and the simulations are the component-separated observed maps obtained using NILC, C-R, SEVEM, and SMICA techniques, as well as 1000 FFP8 simulations obtained using the NILC technique (Planck Collaboration XII 2016). The simulations are based on the Λ CDM paradigm and assume that the temperature fluctuations have a Gaussian distribution. We perform our analyses for a range of scales between 0.05 and 7.33° , which correspond to resolutions between $N = 1024$ and $N = 8$ in the HealPix format (Górski et al. 2005). Further degradation of the maps destabilises the statistics due to the low number of data points in these cases. We do this for a range of mask binarization thresholds: 0.7,

0.8, 0.9, and 0.95; see Sect. 3 for the details of degradation and masking.

In addition, we also compare the observed maps cleaned using the NILC, C-R, SEVEM, and SMICA techniques with 100 simulations each based on the SEVEM and SMICA techniques. The graphs and the p -value tables for these two cases are presented in the appendix. The motivation for this comparison based on a smaller number of simulations is primarily to ascertain if the trends observed are generally consistent irrespective of the cleaning methods. We confirm that this is indeed the case.

We present our analyses in terms of the ranks of relative homology groups, b_p for $0 \leq p \leq 1$. The relative components and loops are quantified by the relative component function, $b_0 : \mathbb{R} \rightarrow \mathbb{R}$, and the relative loop function, $b_1 : \mathbb{R} \rightarrow \mathbb{R}$. We present the graphs of b_0 , b_1 , and of the (relative) Euler characteristic, EC_{rel} , followed by statistical tests that estimate the significance of results. If $f(x) : \mathbb{S}^2 \rightarrow \mathbb{R}$ is the absolute temperature at a location x , and f_0 the mean temperature of the distribution, the dimensionless temperature is given by: $\nu(x) = (f(x) - f_0)/\sigma(f)$, where $\sigma(f)$ is the standard deviation computed from the non-masked pixels. We then obtain the ranks of relative homology groups as functions of the normalized temperature.

4.1. Ranks of relative homology groups

To carry out omnibus tests, we choose 13 a priori levels, ℓ_{-6}, \dots, ℓ_6 , where $\ell_k = k/2$, meaning that the normalized temperature thresholds run from -3 to $+3$ in steps of 0.5 , and consider collections of random variables $b_0(\ell_k)$, $b_1(\ell_k)$, and $EC_{\text{rel}}(\ell_k)$, for $-6 \leq k \leq 6$.

The top two rows of Figs. 9–11 present the curves of b_0 , b_1 , and EC_{rel} , respectively, for resolutions between $N = 1024$ and $N = 8$, for mask threshold 0.9 . The graphs present the average curve (black) from 1000 NILC simulations, with error-bands drawn up to 3σ . The individual curves from simulations are drawn as dotted black lines, a few of which escape the 3σ band. Also plotted are curves from NILC, C-R, SEVEM, and SMICA observed maps. The bottom two rows present the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds. $b_0(\nu)$, $b_1(\nu)$ and $EC_{\text{rel}}(\nu)$ show a difference from simulations peaking near 2σ for some temperature levels for all resolutions. Additionally, b_0 and b_1 show differences peaking between 3 and 4σ sporadically between $N = 32$ and $N = 8$. Noteworthy is the 4.5σ deviation of b_1 between the observations and simulations at $N = 8$ at the normalized temperature threshold $\nu = -3$ in Fig. 10. For the same case, the numbers based on the SEVEM and the SMICA simulations are approximately 5.5σ . However, the low temperature and resolution at which this deviation occurs entail a small number of topological objects on which the statistics are based. As a result, these numbers should perhaps be regarded with a degree of scepticism.

At $N = 1024$, the observed maps based on SEVEM and SMICA simulations deviate very significantly from the NILC simulations map in the range 4 – 6σ at the mean temperature threshold $\nu = 0$. This may perhaps be attributed to the differences in the cleaning pipelines. However, the fact that for lower resolutions the curves are broadly consistent with each other points to the robustness of the underlying data measurement, as well as the mutual consistency of the cleaning methods. A similarly consistent trend is observed in the graphs based on SEVEM and SMICA simulations in the appendix.

Similar graphs based on a comparison between the four observation maps and 100 of each of the SEVEM and SMICA simulations are presented in Figs. B.1–B.6. Figures B.1–B.3 present the graphs based on the SEVEM simulations, while Figs. B.4–B.6 present the graphs for the SMICA simulations. It is evident that the comparisons based on the NILC (in the main paper), the SEVEM, and the SMICA simulations (in the appendix) show consistent trends.

4.2. Experimental evidence of Euler characteristic suppression

As noted above, the Euler-Poincaré formula states that the Euler characteristic is the alternating sum of the Betti numbers. As a consequence, the signals in Euler characteristic are suppressed by design, due to the cancellation of the constituent Betti numbers. Our experiments provide evidence for such suppressions of the topological signals emanating from the Euler characteristic. As an example, consider the quantities at the degraded resolution $N = 16$, and temperature threshold value $\nu = 0.5$. At this resolution and threshold, there is a significant difference in b_0 between observations and simulations at 3.7σ (Fig. 9), but the corresponding value for the Euler characteristic is 2.4σ (Fig. 11). This is because of the cancellation effects between b_0 and b_1 in determining the Euler characteristic. More instances of such cancellation effects can be seen in the graphs, particularly at $N = 1024$, where even though the graphs for b_0 and b_1 from SEVEM and SMICA observed maps deviate by 4 – 6σ from the NILC simulations at $\nu = 0$, the graph of EC_{rel} shows no significant deviation at this threshold.

4.3. Statistical significance of the results

We consider the two methods detailed in Sect. 3.3, and present p -values of the observed maps for both. We consider the variables $b_0(\ell_{k=0,\dots,6})$, $b_1(\ell_{k=-6,\dots,0})$, and $EC_{\text{rel}}(\ell_{k=-6,\dots,6})$ for estimating the statistical significance of the results. The choice of regions is determined by the fact that $b_0(\nu)$ tends to be small, and carries little information for $\nu < 0$, $b_1(\nu)$ tends to be small for $\nu > 0$, and the Euler characteristic is informative over the full range of levels. We perform summary and specific tests for mask binarization threshold values corresponding to 0.7 , 0.8 , 0.9 , and 0.95 . Appendix A presents the table of p -values and the principal component graphs based on 1000 NILC simulations.

4.3.1. Global or summary tests

As a global test for the evidence of a non-random discrepancy in any of the degraded resolutions analysed, we take the full set of normalised differences – for the eight degraded resolutions and the relevant thresholds – as a single vector for each of the three topological quantities separately. Thus, in terms of the notations in Sect. 3.3, for each of b_0 and b_1 , seven thresholds result in $m = 56$ (8 degraded resolutions and 7 temperature thresholds for each resolution); and for the EC, 13 thresholds result in $m = 104$ (8 degraded resolutions and 13 temperature thresholds for each resolution). The last entry for each mask threshold in Table A.1 presents the summary χ^2 and depth p -values. Overall, there is strong indication that the observations differ from the simulations. While at any given degraded resolution the p -value to test for differences is not always small, the fact that they all point in a consistent direction is captured by the summary statistics, which show a very high level of statistical significance.




Fig. 9. b_0 graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curve obtained using NILC, C-R, SEVEM, and SMICA methods, and the expected (black) curve computed from 1000 NILC simulations, along with bands drawn up to 3σ . Also plotted underneath are the curves from individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds. Maximum noted deviation is at $N = 16$ at 3.7σ . The threshold along the horizontal axis runs from positive to negative, in view of the fact that we analyse superlevel sets of the normalized temperature field.

4.3.2. Tests for specific degraded resolutions

This is followed by specific tests for each resolution. The rest of the entries in Table A.1 present the Mahalanobis and the depth p -values for each resolution, for different mask thresholds. The Mahalanobis distances are particularly small for b_1 at $N = 16$, and very significant at $N = 8$, across all binarization thresholds. Although they are not stable across binarization thresholds, b_0 and EC_{rel} also show significance. The depth p -values are very significant for b_1 for $N = 16$ and $N = 8$, while b_0 shows high significance at $N = 32$ consistently across the range of binarization thresholds. The depth p -values also show high significance at higher resolutions, more often for b_1 than for b_0 , but not at all for EC_{rel} , presumably because of cancellation effects. When considering Tukey depth, b_1 shows significance more often than b_0 and EC_{rel} , and is an order of magnitude more significant compared to b_0 and EC_{rel} when considering the Mahalanobis distance. These

trends are broadly consistent irrespective of the cleaning method. Tables of p -values based on 100 SEVEM and SMICA simulations are presented in Tables B.1 and B.2. The Mahalanobis p -values are consistent with those obtained with the NILC simulations. For the Tukey depth test, 100 simulations are inadequate to resolve the p -values in most cases.

Regardless of the choice of test or the cleaning pipeline, it is evident that the model and observations disagree significantly at least in the number of loops on a range of scales between approximately 1 and 7° . For the Mahalanobis values, the general trend of significance increases up to $N = 8$, providing additional evidence that the deviations are not purely due to chance. For both tests, the p -values for the summary tests tend to be more significant than for individual resolutions. Another general trend is that the non-parametric test shows the difference between the observed and the simulated maps to be starker than the parametric test. To help interpret this difference, Figs. A.1–A.4 present




Fig. 10. b_1 graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curve obtained using NILC, C-R, SEVEM, and SMICA methods, and the expected (black) curve computed from 1000 NILC simulations, along with bands drawn up to 3σ . Also plotted underneath are the 1000 curves from individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds. Maximum noted deviation is at $N = 8$ at 4.5σ . It is difficult to judge the validity of this number, as the low temperature threshold ($v = -3$) entails a low total number of objects on which the statistics are based. The next peak in the curve is located at a moderate threshold ($v = -0.5$), and indicates a deviation at 2.9σ . The threshold along the horizontal axis runs from positive to negative, in view of the fact that we analyse superlevel sets of the normalized temperature field.

plots that visualise to what extent the assumption of a Gaussian distribution for the compared quantities is justified; see also Table C.1 for a comparison with p -values for the absolute homology. The results indicate similar trends to the relative homology case.

4.4. Principal-component graphs

Figure A.1 presents the projection onto the first two principal components for the summary tests, which include results from all resolutions. Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (top) and purple (bottom). Observed CMB points are in red. Examining the diagrams corresponding to the Mahalanobis distance, the hypothesis that the distribution conforms to elliptical contours is questionable.

Figures A.2–A.4 present the projection onto the first two principal components for b_0 , b_1 , and EC_{rel} , respectively, for specific resolutions. Also drawn are the Mahalanobis (top two

rows) and Tukey depth (bottom two rows) contours. In general, the symmetric Mahalanobis contours do not always fit the data. However, as the resolution decreases, the Mahalanobis contours, which are Gaussian in nature, seem to fit the data well, and may be a reasonable approximation after all. Such graphs based on SEVEM and SMICA simulations are presented in Figs. B.7–B.10. It is evident that 100 simulations may not be enough to reliably resolve the p -values.

5. Summary and conclusions

We provide evidence for the deviation of the observed *Planck* CMB maps from the Gaussian predictions of the standard Λ CDM model. Specifically, we find an over-abundance of loops in the observed maps, deviating from the simulations at per cent to less than per mil levels. This is in terms of p -values computed using χ^2 statistics, between the resolutions $N = 32$ and $N = 8$. The difference in the number of components and loops peaks sporadically at more than 3σ from the predictions between

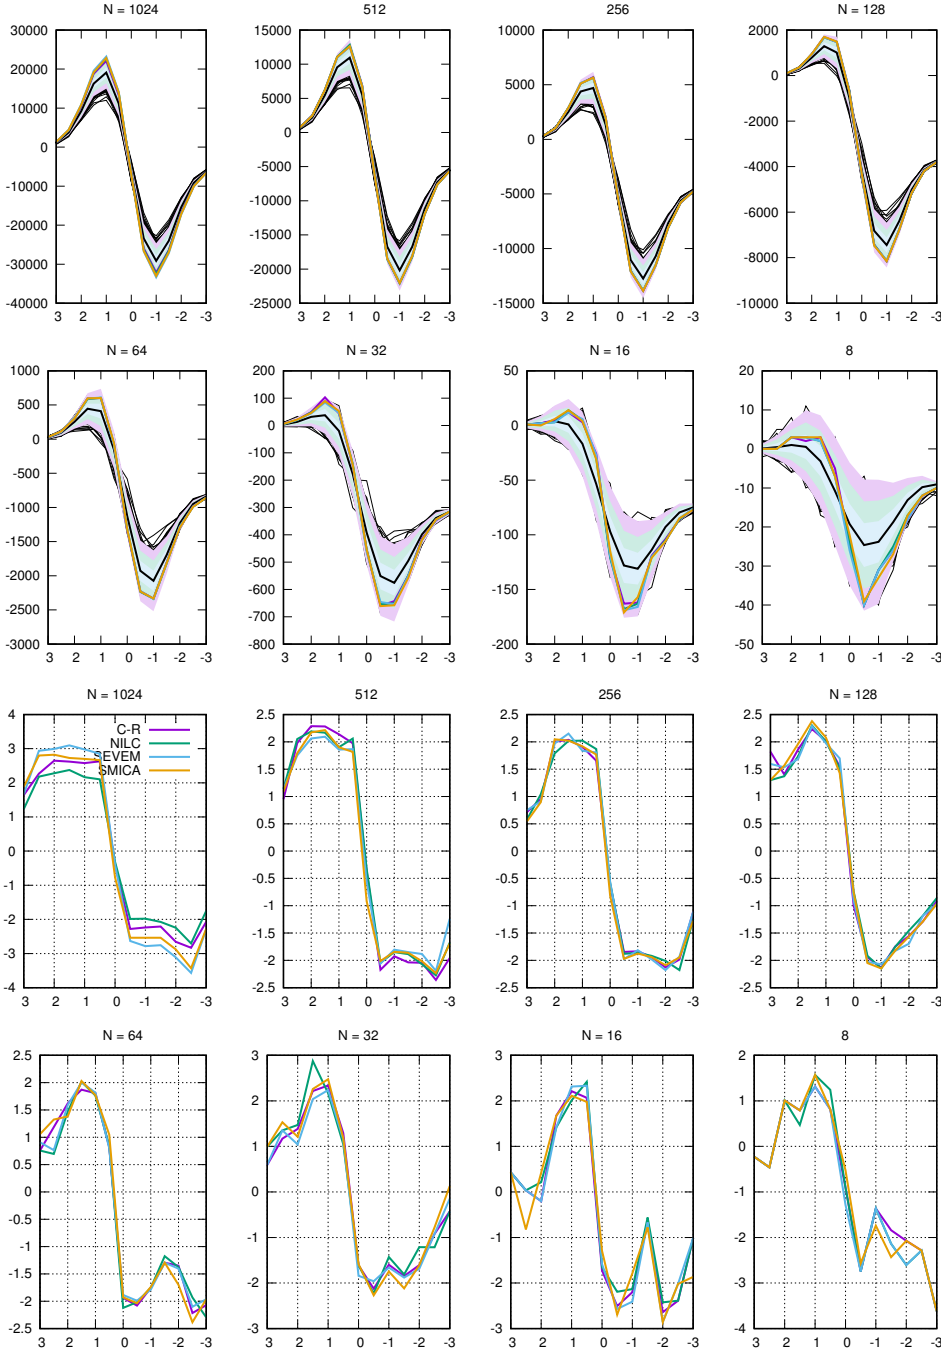


Fig. 11. Euler characteristic graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curve obtained using NILC, C-R, SEVEM, and SMICA methods, and the expected (black) curve computed from 1000 NILC simulations, along with bands drawn up to 3σ . Also plotted underneath are the 1000 curves from individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds. Maximum noted deviation is at $N = 32$ at 2.9σ . The threshold along the horizontal axis runs from positive to negative, in view of the fact that we analyse superlevel sets of the normalized temperature field.

$N = 32$ and $N = 8$. Results based on smoothed maps corroborate with those based on degraded maps in terms of approximate scales at which the anomaly is observed. We also compute the absolute homology for the dataset, and confirm that the results are consistent with those from relative homology. External evidence that these deviations are not a result of overanalysing the data comes from the fact that the variance of the observed CMB is anomalous with respect to the standard model at $N = 16$ (Planck Collaboration XXIII 2014), and the computed power spectrum exhibits a dip roughly at this range of scales. In addition, there are reports of a mildly significant Euler characteristic at 3.66° ($N = 16$) (Eriksen et al. 2004b), computed from independent measurements of the CMB by Planck’s predecessor – Wilkinson Microwave Anisotropy Probe (WMAP) satellite. This can be explained by the significantly high number of

loops and components, together with cancellation effects that the Euler characteristic suffers from. Similar observations by independent satellites suggests that it is unlikely that the source of the anomaly has its origin in instrumental noise or systematic effects. Moreover the medium super-horizon scales at which we observe it, could possibly point to a cosmological origin. The non-parametric Tukey depth test shows the observations to be different from the simulations at almost all resolutions. Regardless of the preferred test, the topological structure of the CMB appears to deviate from the simulations, at least on some scale. This trend is robust and consistent irrespective of the choice of the cleaning method, thus ruling out the possibility that the deviations we observe are merely an artifact of the cleaning method.

We can rule out this anomaly being the effect of the cold spot in the CMB sky, or any previously detected directional




Fig. 12. Visualization of the loops for the largest excursion set, which consists of the entire sphere minus the mask. To improve the visualization, the temperature field has been smoothed by a small amount, and we do not draw very short loops. *From left to right:* the sphere from the top, the bottom, the left, and the right views.

anomalies. Our statistics are based on a large number of loops surrounding the low-density regions, to which the loop generated by the cold spot may contribute at most only a few, and often only one. Moreover, to support this claim, we visually confirm that these loops are scattered all over the sky (see Fig. 12). We also test and confirm that simulations that are based on Gaussian prescriptions and match the characteristics of the observed “dipped” power spectrum cannot resolve this anomaly. Additionally, we present topological methods that are suitable in the presence of obfuscating masks. As such, the results presented in this paper are robust despite lacking full sky coverage, and are model-independent.

In conclusion, we reiterate that we present clear evidence of departure of the observed CMB maps with respect to the simu-

lations based on the Λ CDM paradigm, but make no attempt to address the issue of the physical mechanism behind this phenomenon; a question we leave to the wider cosmological community. Nevertheless, our analysis demonstrates the existence of unexpected topology in the CMB. Possible, but non-exhaustive scenarios worth exploring may be primordial non-Gaussianity, as well as models with non-trivial topology including topological defect models.

Acknowledgements. PP is grateful to Julian Borill from the *Planck* consortium for providing the data, and for the illuminating discussions and inputs. PP also thanks Hans Kristen Eriksen, Anne Ducout, and Francois R. Bouchet for significantly helpful discussions at various stages. The authors collectively thank the anonymous referee for the invaluable comments and suggestions that

have added significant value to the contents of the manuscript. PP and RA acknowledge the support of ERC advanced grant Understanding Random Systems through Algebraic Topology (URSAT) (no: 320422, PI: RA). This work is also part of a project that has received funding for PP and TB from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement ERC advanced grant 740021—Advances in Research on THEories of the dark UniverSe (ARTHUS), PI: TB). HE and HW acknowledge the support by the Office of Naval Research, through grant N62909-18-1-2038, and by the DFG Collaborative Research Center TRR 109, “Discretization in Geometry and Dynamics”, through grant I02979-N35 of the Austrian Science Fund (FWF). PP acknowledges the support and use of resources at the NERSC computing center.

References

- Adler, R. J. 1981, in *The Geometry of Random Fields* (Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM)), Classics Appl. Math.
- Adler, R. J., & Taylor, J. E. 2010, in *Random Fields and Geometry* (Springer), Springer Monographs Math.
- Adler, R. J., Agami, S., & Pranav, P. 2017, *Proc. Natl. Acad. Sci.*, **114**, 11878
- Aurich, R., & Steiner, F. 2001, *MNRAS*, **323**, 1016
- Aurich, R., Lustig, S., Steiner, F., & Then, H. 2007, *Class. Quant. Grav.*, **24**, 1879
- Bardeen, J. M., Bond, J. R., Kaiser, N., & Szalay, A. S. 1986, *ApJ*, **304**, 15
- Bartolo, N., Matarrese, S., & Riotto, A. 2010, *Adv. Astron.*, **2010**, 157079
- Bauer, U., Kerber, M., Reininghaus, J., & Wagner, H. 2014, in *Mathematical Software – ICMS 2014* (Berlin Heidelberg: Springer), 137
- Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, *ApJS*, **148**, 1
- Bernui, A., Novaes, C. P., Pereira, T. S., & Starkman, G. D. 2018, ArXiv e-prints [arXiv:1809.05924]
- Betti, E. 1871, *Ann. Mat. Pura Appl.*, **2**, 140
- Bouchet, F. R., Peter, P., Riazuelo, A., & Sakellariadou, M. 2002, *Phys. Rev. D*, **65**, 021301
- Buchert, T., France, M. J., & Steiner, F. 2017, *Class. Quant. Grav.*, **34**, 094002
- Cole, A., & Shiu, G. 2018, *JCAP*, **3**, 025
- Colombi, S., Pogosyan, D., & Souradeep, T. 2000, *Phys. Rev. Lett.*, **85**, 5515
- Copi, C., Huterer, D., Schwarz, D., & Starkman, G. 2015, *MNRAS*, **449**, 3458
- Doroshkevich, A. G. 1970, *Astrophysics*, **6**, 320
- Ducout, A., Bouchet, F. R., Colombi, S., Pogosyan, D., & Prunet, S. 2013, *MNRAS*, **429**, 2104
- Edelsbrunner, H., & Harer, J. 2010, *Computational Topology: An Introduction, Applied Mathematics* (American Mathematical Society)
- Edelsbrunner, H., Letscher, J., & Zomorodian, A. 2002, *Discrete Comput. Geom.*, **28**, 511
- Eriksen, H. K., Hansen, F. K., Banday, A. J., Górski, K. M., & Lilje, P. B. 2004a, *ApJ*, **605**, 14
- Eriksen, H. K., Novikov, D. I., Lilje, P. B., Banday, A. J., & Górski, K. M. 2004b, *ApJ*, **612**, 64
- Euler, L. 1758, *Novi Commentarii academiae scientiarum Petropolitanae*, **4**, 140
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Gott, III., J. R., Dickinson, M., & Melott, A. L. 1986, *ApJ*, **306**, 341
- Guth, A. H., & Pi, S.-Y. 1982, *Phys. Rev. Lett.*, **49**, 1110
- Jones, B. 2017, *Precision Cosmology: The First Half Million Years* (Cambridge University Press)
- Komatsu, E., Smith, K. M., Dunkley, J., et al. 2011, *ApJS*, **192**, 18
- Mahalanobis, P. C. 1936, *Proc. Natl. Inst. Sci.*, **2**, 49
- Makarenko, I., Shukurov, A., Henderson, R., et al. 2018, *MNRAS*, **475**, 1843
- Matsubara, T. 2010, *Phys. Rev. D*, **81**, 083505
- Mecke, K. R., Buchert, T., & Wagner, H. 1994, *A&A*, **288**, 697
- Munkres, J. 1984, *Elements of Algebraic Topology, Advanced Book Classics* (Perseus Books)
- Novikov, D., Colombi, S., & Doré, O. 2006, *MNRAS*, **366**, 1201
- Park, C., Pranav, P., Chingangbam, P., et al. 2013, *J. Korean Astron. Soc.*, **46**, 125
- Park, C.-G. 2004, *MNRAS*, **349**, 313
- Planck Collaboration I. 2014, *A&A*, **571**, A1
- Planck Collaboration XXIII. 2014, *A&A*, **571**, A23
- Planck Collaboration XXIV. 2014, *A&A*, **571**, A24
- Planck Collaboration IX. 2016, *A&A*, **594**, A9
- Planck Collaboration XII. 2016, *A&A*, **594**, A12
- Planck Collaboration XIII. 2016, *A&A*, **594**, A13
- Pranav, P. 2015, *Persistent Holes in the Universe: A Hierarchical Topology of the Cosmic Mass Distribution* (University of Groningen)
- Pranav, P., Edelsbrunner, H., van de Weygaert, R., et al. 2017, *MNRAS*, **465**, 4281
- Pranav, P., van de Weygaert, R., Vegter, G., et al. 2019, *MNRAS*, **485**, 4167
- Sahni, V., Sathyaprakash, B., & Shandarin, S. 1998, *ApJ*, **507**, L109
- Schmalzing, J., & Buchert, T. 1997, *ApJ*, **482**, L1
- Schmalzing, J., & Gorski, K. M. 1998, *MNRAS*, **297**, 355
- Schmalzing, J., Buchert, T., Melott, A. L., et al. 1999, *ApJ*, **526**, 568
- Schwarz, D. J., Copi, C. J., Huterer, D., & Starkman, G. D. 2016, *Class. Quant. Grav.*, **33**, 184001
- Shivashankar, N., Pranav, P., Natarajan, V., et al. 2016, *IEEE Trans. Vis. Comput. Graph.*, **22**, 1745
- Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJ*, **396**, L1
- Sousbie, T. 2011, *MNRAS*, **414**, 350
- Sousbie, T., Pichon, C., Courtois, H., Colombi, S., & Novikov, D. 2008, *ApJ*, **672**, L1
- Spergel, D. N., Bean, R., Doré, O., et al. 2007, *ApJS*, **170**, 377
- The CGAL Project 2018, *CGAL User and Reference Manual*, 4.11.1 edn. (CGAL Editorial Board)
- Tukey, J. W. 1975, *Proc. 1974 Int. Congr. Math.*, **2**, 523
- van de Weygaert, R., Vegter, G., Edelsbrunner, H., et al. 2011, *Trans. Comput. Sci.*, **14**, 60

Table A.1. continued.

Relative homology								Relative homology							
Resolution	Method	Mahalanobis			Tukey Depth			Resolution	Method	Mahalanobis			Tukey Depth		
		b_0	b_1	EC_{rel}	b_0	b_1	EC_{rel}			b_0	b_1	EC_{rel}	b_0	b_1	EC_{rel}
Threshold = 0.90								Threshold = 0.95							
128	NILC	0.259	0.525	0.486	<0.001	0.160	<0.001	128	NILC	0.309	0.524	0.554	<0.001	0.165	0.446
	C-R	0.355	0.605	0.614	0.180	0.238	<0.001		C-R	0.313	0.601	0.612	0.169	0.165	0.530
	SEVEM	0.275	0.681	0.683	<0.001	0.160	0.344		SEVEM	0.381	0.624	0.795	0.169	0.237	0.674
	SMICA	0.250	0.617	0.528	<0.001	0.160	<0.001		SMICA	0.308	0.597	0.563	<0.001	0.237	0.446
64	NILC	0.611	0.248	0.444	0.547	<0.001	0.428	64	NILC	0.420	0.157	0.430	0.276	<0.001	<0.001
	C-R	0.571	0.369	0.651	0.426	<0.001	0.428		C-R	0.491	0.125	0.440	0.319	<0.001	<0.001
	SEVEM	0.597	0.336	0.630	0.559	<0.001	0.663		SEVEM	0.466	0.082	0.292	0.276	<0.001	<0.001
	SMICA	0.430	0.348	0.458	0.397	<0.001	0.307		SMICA	0.438	0.111	0.259	0.276	<0.001	<0.001
32	NILC	0.053	0.305	0.300	<0.001	0.243	0.340	32	NILC	0.076	0.246	0.174	<0.001	<0.001	<0.001
	C-R	0.248	0.633	0.754	0.190	0.441	0.700		C-R	0.375	0.513	0.767	0.356	0.342	0.715
	SEVEM	0.175	0.638	0.715	<0.001	0.370	0.585		SEVEM	0.217	0.493	0.634	<0.001	0.166	0.522
	SMICA	0.141	0.389	0.514	<0.001	0.333	0.531		SMICA	0.107	0.367	0.431	<0.001	<0.001	<0.001
16	NILC	0.009	0.054	0.081	<0.001	<0.001	<0.001	16	NILC	0.436	0.016	0.027	0.383	<0.001	<0.001
	C-R	0.026	0.038	0.087	<0.001	<0.001	<0.001		C-R	0.358	0.007	0.022	0.339	<0.001	<0.001
	SEVEM	0.024	0.025	0.055	<0.001	<0.001	<0.001		SEVEM	0.582	0.020	0.021	0.562	<0.001	<0.001
	SMICA	0.013	0.014	0.081	<0.001	<0.001	<0.001		SMICA	0.353	0.012	0.062	0.339	<0.001	<0.001
8	NILC	0.408	<0.001	0.007	0.401	<0.001	<0.001	8	NILC	0.871	0.047	0.188	0.846	<0.001	0.517
	C-R	0.610	<0.001	0.023	0.568	<0.001	0.330		C-R	0.773	0.144	0.384	0.525	0.160	0.611
	SEVEM	0.502	<0.001	0.012	0.528	<0.001	0.237		SEVEM	0.735	0.104	0.265	0.451	<0.001	0.564
	SMICA	0.414	<0.001	0.014	0.455	<0.001	<0.001		SMICA	0.887	0.081	0.181	0.816	0.195	0.466
Summary	NILC	0.010	<0.001	0.001	<0.001	0.002	<0.001	Summary	NILC	0.213	0.003	0.010	<0.001	0.009	0.001
	C-R	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001		C-R	<0.001	<0.001	<0.001	<0.001	<0.001	0.001
	SEVEM	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001		SEVEM	<0.001	<0.001	<0.001	<0.001	<0.001	0.001
	SMICA	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001		SEVEM	<0.001	<0.001	<0.001	<0.001	<0.001	0.001

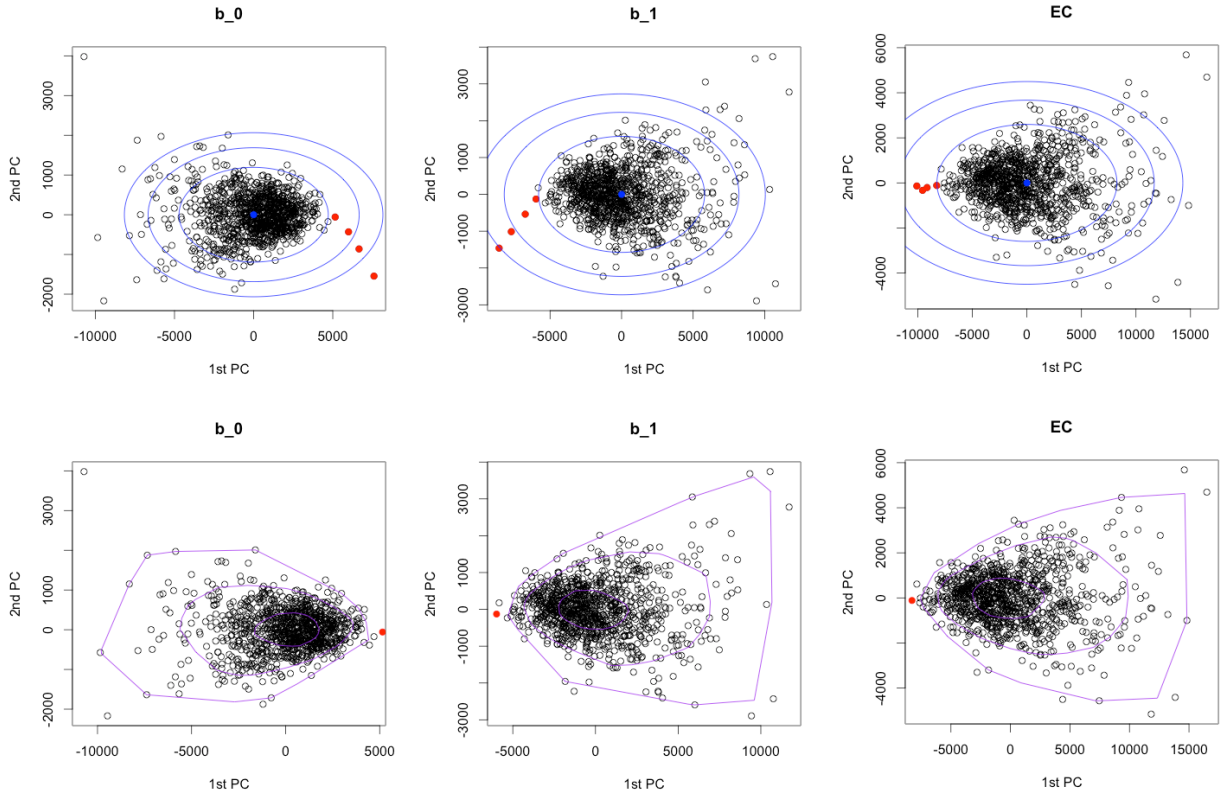


Fig. A.1. Summary test. Projection onto first two principal components. Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top*) and purple (*bottom*). Observed CMB points are in red (filled circles). Points from simulations are denoted by black empty circles.




Fig. A.2. Projection onto first two principal components for specific resolution tests for b_0 . Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top two rows*) and purple (*bottom two rows*). Observed CMB points are in red (filled circles). Points from simulations are denoted by black empty circles.

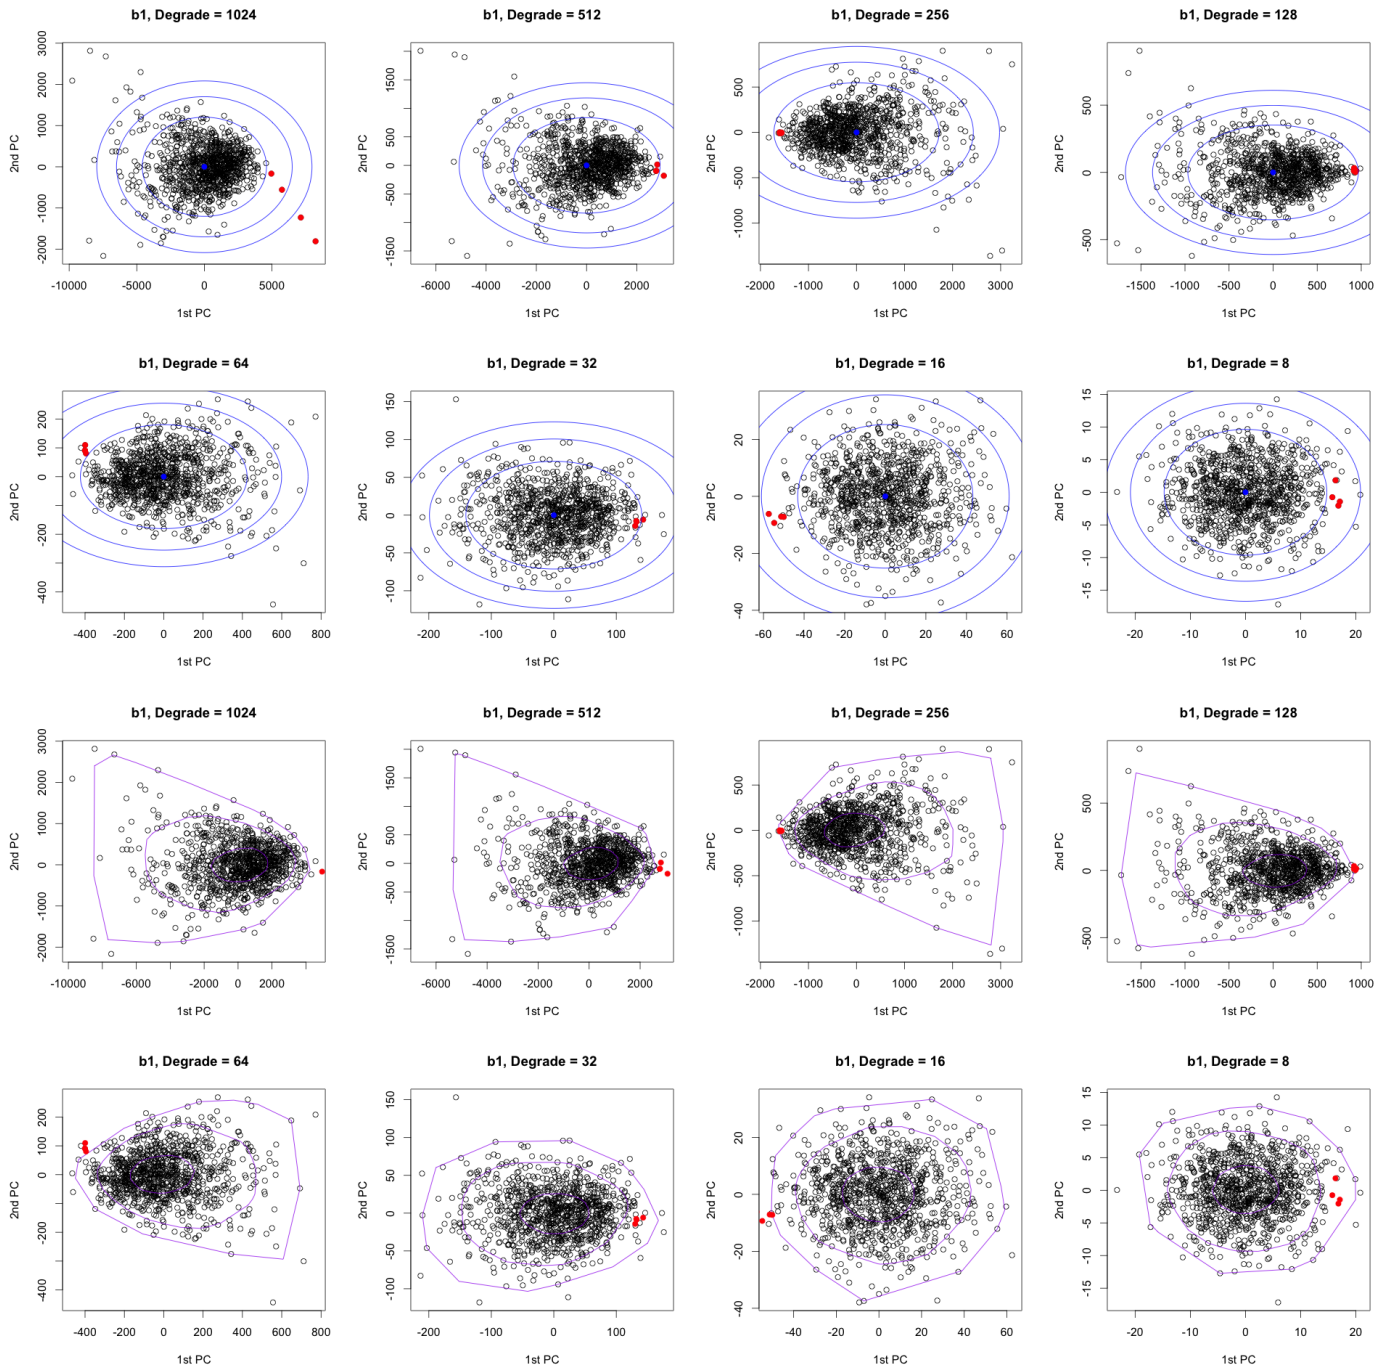


Fig. A.3. Projection onto first two principal components for specific resolution tests for b_1 . Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top two rows*) and purple (*bottom two rows*). Observed CMB points are in red (filled circles). Points from simulations are denoted by black empty circles.

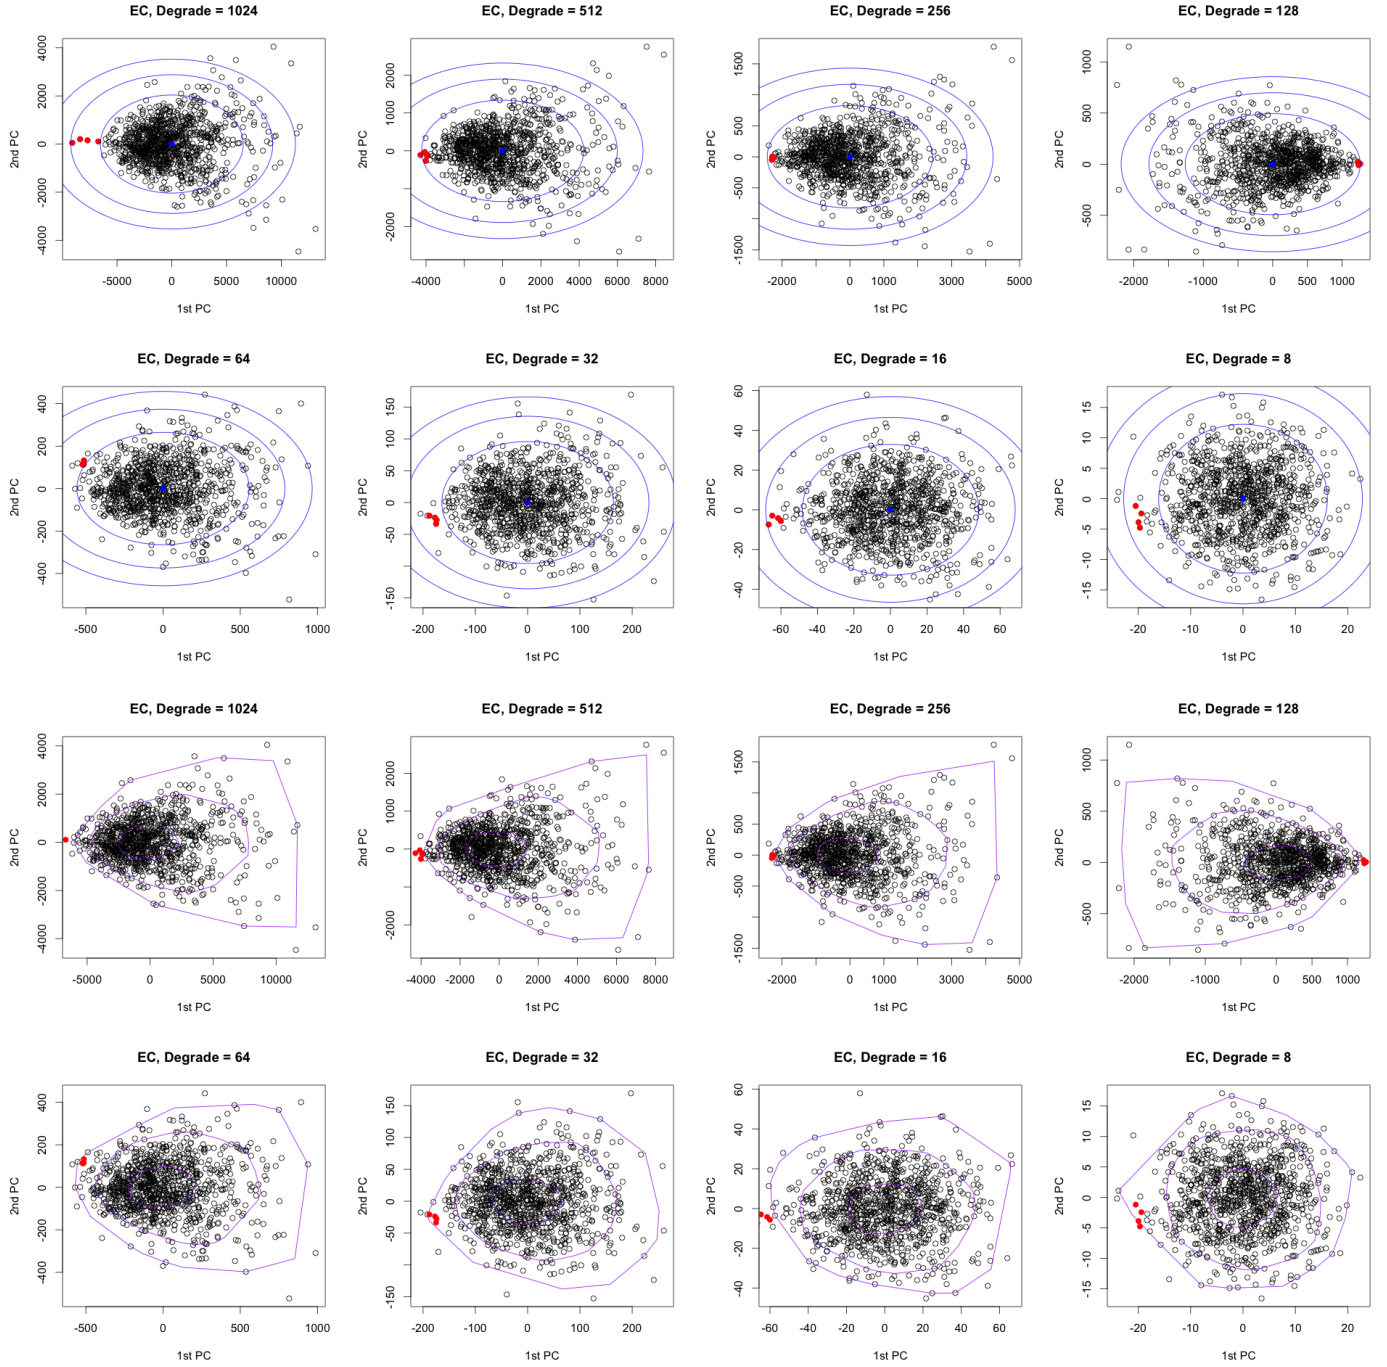


Fig. A.4. Projection onto first two principal components for specific resolution tests for EC_{rel} . Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top two rows*) and purple (*bottom two rows*). Observed CMB points are in red (filled circles). Points from simulations are denoted by black empty circles.

Appendix B: Comparison with SEVEM and SMICA simulations

This appendix presents the results from the comparison of the observed maps with a hundred simulations from each of SEVEM

and SMICA methods. In summary, the results are consistent with those based on the NILC simulations presented in the main paper and Appendix A.




Fig. B.1. b_0 graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curves from all the simulation methods, and the expected (black) curve computed from 100 SEVEM simulations, along with bands drawn up to 3σ . Also plotted underneath are 100 curves from the individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds.




Fig. B.2. b_1 graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curves from all the simulation methods, and the expected (black) curve computed from 100 SEVEM simulations, along with bands drawn up to 3σ . Also plotted underneath are 100 curves from the individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds.




Fig. B.3. EC_{rel} graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curves from all the simulation methods, and the expected (black) curve computed from 100 SEVEM simulations, along with bands drawn up to 3σ . Also plotted underneath are 100 curves from the individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds.

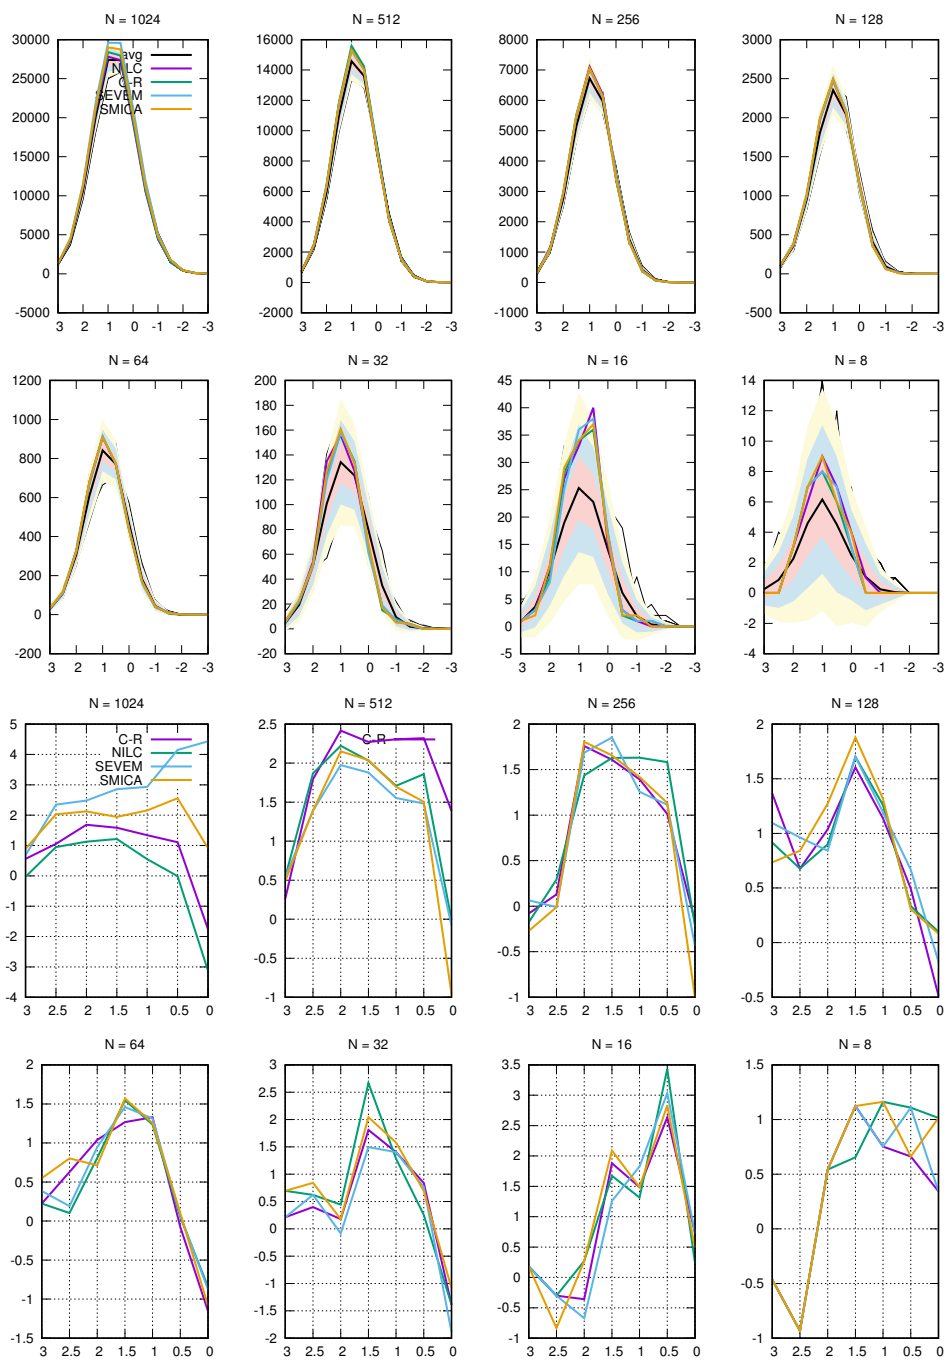


Fig. B.4. b_0 graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curves from all the simulation methods, and the expected (black) curve computed from 100 SMICA simulations, along with bands drawn up to 3σ . Also plotted underneath are 100 curves from the individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds.

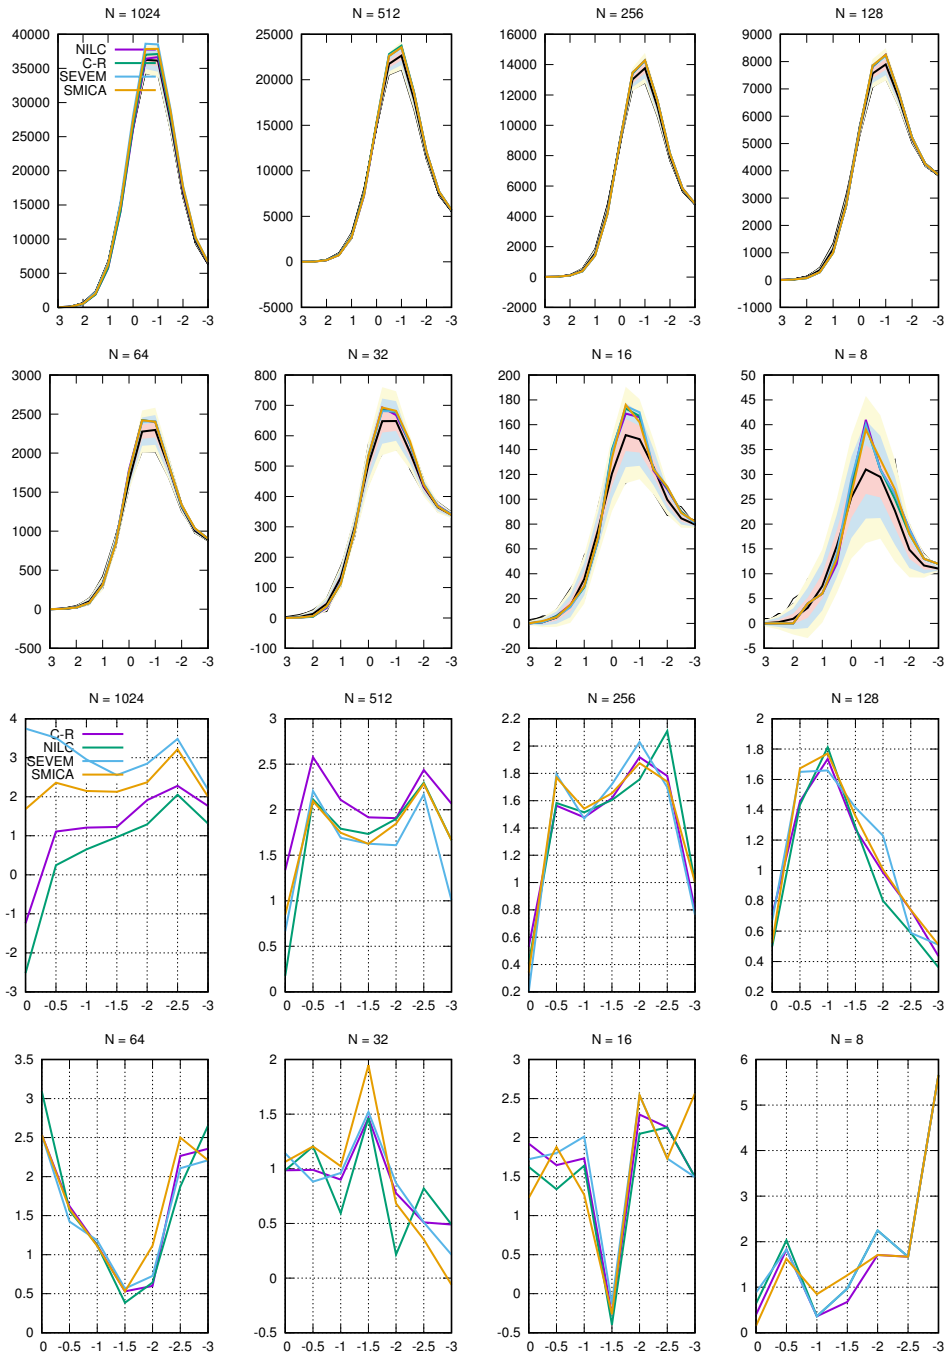


Fig. B.5. b_1 graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curves from all the simulation methods, and the expected (black) curve computed from 100 SMICA simulations, along with bands drawn up to 3σ . Also plotted underneath are 100 curves from the individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds.




Fig. B.6. EC_{rel} graphs for resolutions between $N = 1024$ and $N = 8$. *Top two rows:* observed curves from all the simulation methods, and the expected (black) curve computed from 100 SMICA simulations, along with bands drawn up to 3σ . Also plotted underneath are 100 curves from the individual simulations. *Bottom two rows:* curve presenting the difference between the observations and simulations in terms of the number of standard deviations for the various temperature thresholds.

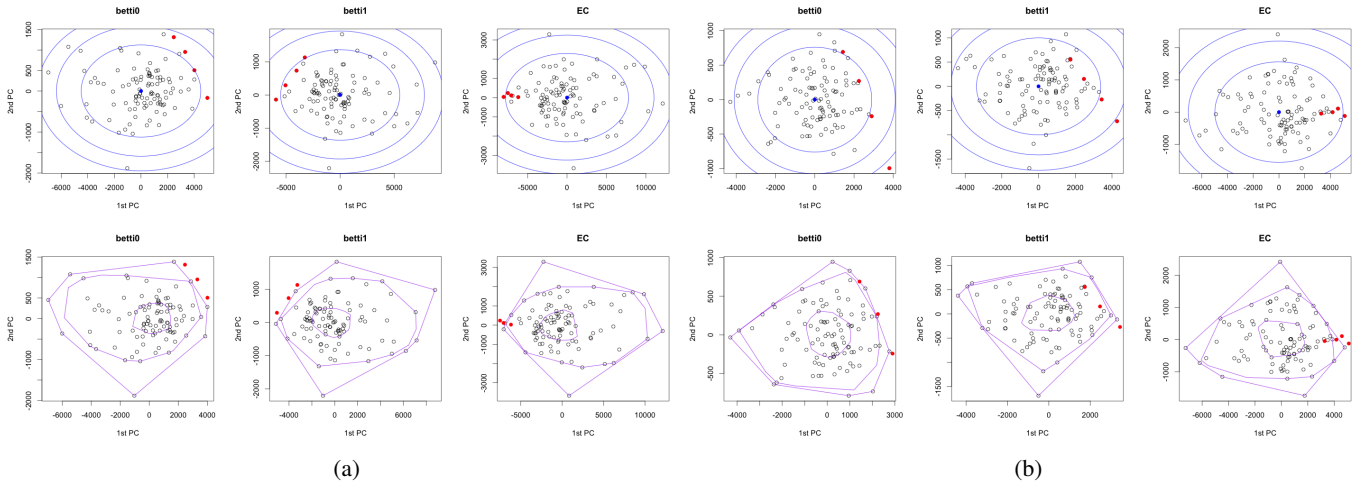


Fig. B.7. Summary test computed from SEVEM (*panel a*) and SMICA (*panel b*) simulations. Projection onto first two principal components. Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top*) and purple (*bottom*). Observed CMB points are in red (filled circles). Points from simulations are denoted by black empty circles.




Fig. B.8. Projection onto first two principal components for specific resolution tests for b_0 for SEVEM and SMICA simulations. Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top two rows*) and purple (*bottom two rows*). Observed CMB points are in red (filled circles). Points from simulations are denoted by black empty circles. *Panels a* and *c*: Mahalanobis and Tukey depth graphs for the SEVEM simulations. *Panels b* and *d*: Mahalanobis and Tukey depth graphs for the SMICA simulations, respectively.




Fig. B.9. Projection onto first two principal components for specific resolution tests for b_1 for SEVEM and SMICA simulations. Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top two rows*) and purple (*bottom two rows*). Observed CMB points are in red. *Panels a* and *c*: Mahalanobis and Tukey depth graphs for the SEVEM simulations, respectively. *Panels b* and *d*: Mahalanobis and Tukey depth graphs for the SMICA simulations, respectively.




Fig. B.10. Projection onto first two principal components for specific resolution tests for EC_{tel} for SEVEM and SMICA simulations. Mahalanobis and depth contours corresponding to p -values of 0.1, 0.01, and 0.001 are shown in blue (*top two rows*) and purple (*bottom two rows*). Observed CMB points are in red. *Panels a* and *c*: Mahalanobis and Tukey depth graphs for the SEVEM simulations, respectively. *Panels b* and *d*: Mahalanobis and Tukey depth graphs for the SMICA simulations, respectively.

Appendix C: Table of p -values for absolute homology

This appendix presents the table of p -values computed from the absolute homology of the data sets.

Table C.1. Two-tailed p -values for absolute homology obtained from parametric (Mahalanobis distance) and non-parametric (Tukey depth) tests, for four mask binarization thresholds, obtained using 1000 NILC simulations.

Absolute homology						
Resolution	Mahalanobis			Tukey Depth		
	β_0	β_1	EC	β_0	β_1	EC
Threshold = 0.70						
1024	0.416	0.318	0.629	<0.001	<0.001	0.501
512	0.492	0.371	0.573	0.241	0.133	0.437
256	0.535	0.308	0.488	0.131	<0.001	<0.001
128	0.490	0.316	0.542	0.163	<0.001	0.445
64	0.413	0.184	0.396	0.180	<0.001	<0.001
32	0.012	0.024	0.035	<0.001	<0.001	<0.001
16	0.220	0.017	0.135	0.164	<0.001	<0.001
8	<0.001	0.584	0.004	<0.001	0.647	<0.001
Summary	0.064	0.276	0.040	0.152	0.502	<0.001
Threshold = 0.80						
1024	0.416	0.318	0.629	<0.001	<0.001	0.501
512	0.492	0.371	0.573	0.241	0.133	0.437
256	0.535	0.308	0.488	0.131	<0.001	<0.001
128	0.490	0.316	0.542	0.163	<0.001	0.445
64	0.413	0.184	0.396	0.180	<0.001	<0.001
32	0.012	0.024	0.035	<0.001	<0.001	<0.001
16	0.220	0.017	0.135	0.164	<0.001	<0.001
8	0.001	0.584	0.004	<0.001	0.647	<0.001
Summary	<0.001	0.204	0.001	<0.001	0.010	0.001
Threshold = 0.90						
1024	0.281	0.339	0.642	<0.001	<0.001	0.301
512	0.523	0.389	0.697	0.135	<0.001	0.520
256	0.402	0.329	0.422	<0.001	<0.001	0.307
128	0.542	0.216	0.342	0.153	<0.001	0.298
64	0.276	0.228	0.421	<0.001	<0.001	0.448
32	0.045	0.172	0.231	<0.001	0.175	<0.001
16	0.058	0.019	0.037	<0.001	<0.001	<0.001
8	0.223	0.475	0.071	0.199	0.510	<0.001
Summary	0.009	0.386	0.024	0.371	0.631	0.206
Threshold = 0.95						
1024	0.281	0.339	0.642	<0.001	<0.001	0.301
512	0.523	0.407	0.698	0.250	<0.001	0.512
256	0.380	0.339	0.302	0.132	<0.001	<0.001
128	0.456	0.226	0.248	<0.001	<0.001	<0.001
64	0.491	0.449	0.794	0.235	0.346	0.431
32	0.261	0.262	0.553	<0.001	0.308	<0.001
16	0.137	0.066	0.050	0.155	<0.001	<0.001
8	0.155	0.024	0.008	<0.001	<0.001	<0.001
Summary	0.049	0.033	0.036	<0.001	<0.001	<0.001

Notes. The last entry for each threshold is the p -value for the summary statistic computed across all resolutions. Marked in boldface are p -values of 0.05 or smaller.