

Too few spots in the cosmic microwave background

Youness Ayaita, Maik Weber, and Christof Wetterich

Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, D-69120 Heidelberg, Germany

(Received 9 July 2009; published 11 January 2010)

We investigate the abundance of large-scale hot and cold spots in the WMAP-5 temperature maps and find considerable discrepancies compared to Gaussian simulations based on the Λ CDM best-fit model. Too few spots are present in the reliably observed cosmic microwave background (CMB) region, i.e., outside the foreground-contaminated parts excluded by the KQ75 mask. Even simulated maps created from the original WMAP-5 estimated multipoles contain more spots than visible in the measured CMB maps. A strong suppression of the lowest multipoles would lead to better agreement. The lack of spots is reflected in a low mean temperature fluctuation on scales of several degrees (4° – 8°), which is only shared by less than 1% (0.16%–0.62%) of Gaussian Λ CDM simulations. After removing the quadrupole, the probabilities change to 2.5%–8.0%. This shows the importance of the anomalously low quadrupole for the statistical significance of the missing spots. We also analyze a possible violation of Gaussianity or statistical isotropy (spots are distributed differently outside and inside the masked region).

DOI: 10.1103/PhysRevD.81.023507

PACS numbers: 98.80.–k, 98.70.Vc

I. INTRODUCTION

The precise measurement of anisotropies in the cosmic microwave background (CMB) has played a key role in amplifying our knowledge about the structure and evolution of the Universe. The best data available today is provided by the Wilkinson Microwave Anisotropy Probe (WMAP) satellite mission from five years of observation. Its results are powerful enough to put various cosmological models to stringent tests. They helped establishing today's standard model of a spatially flat universe with Gaussian initial perturbations, possibly generated during an early inflationary epoch. According to the standard Λ CDM model, the present Universe is essentially made up from dark energy in the form of a cosmological constant Λ and cold dark matter (CDM). Under the assumptions of Gaussianity and statistical isotropy, all the information about the temperature fluctuations in the CMB are encoded in the angular power spectrum C_l from a harmonic decomposition of the temperature field. A crucial result of the WMAP analysis therefore is an estimate of the multipoles C_l which is in good agreement with the Λ CDM best fit [1] except for the well-known discrepancies of the low multipoles, especially the quadrupole C_2 [2]. Nonetheless, many issues are still under intense discussion. Repeatedly, authors have claimed to detect non-Gaussian signals [3,4] or statistical anisotropy [5–10]. Since the power spectrum is insensitive to these anomalies, it is necessary to perform additional investigations of the temperature sky map. These are done in harmonic, wavelet, and pixel space [11]. Even if Gaussianity holds, it may still give new insights to switch to another representation of the statistical properties of the temperature maps since a phenomenon can be more easily detected in one representation than in another.

The goal of this work is to provide a clear and intuitive analysis in pixel space regarding abundances of large-scale hot and cold spots identified as regions whose mean temperature contrasts exceed some (variable) threshold. We analyze both observed CMB maps and Gaussian simulations based on Λ CDM. The comparison reveals severe deviations. Other authors who worked with statistics of local extrema in the temperature field also observed significant anomalies [12–14].

We start by recalling some basic results that connect pixel-space analyses with the angular power spectrum in Sec. II. A comprehensive description of our method follows in Sec. III including the preparation of adequate Gaussian simulations, the working principle of our spot-searching algorithm, and an error estimation. Our results are presented in Sec. IV. We consider both cut-sky maps (with unreliable pixels excluded by the KQ75 temperature analysis mask) and the Internal Linear Combination (ILC) full-sky map, and quantify deviations from Gaussian simulations. We sum up and conclude in Sec. V.

II. PRELIMINARY CONSIDERATIONS

The most robust comparison between predicted and observed spot abundances of CMB sky maps relies on simulated maps since analytic methods can hardly care for complications due to masking and beam properties. Creating a number of simulated maps and treating them in exactly the same way as the original map therefore is the clearest method. Nonetheless, it is instructive to recall some well-known analytic results that connect the pixel-space analysis to familiar harmonic space.

The spot abundances in a CMB sky map are dictated by the angular correlations of temperature fluctuations. The most popular theories stick to Gaussianity and statistical isotropy. Then, the ensemble average of the angular corre-

lation between two directions (θ, φ) and (θ', φ') only depends on the angle Θ between them. This leads to the definition of the angular correlation function

$$C(\Theta) = \left\langle \frac{\Delta T}{\bar{T}}(\theta, \varphi) \times \frac{\Delta T}{\bar{T}}(\theta', \varphi') \right\rangle. \quad (1)$$

We can switch to harmonic space by decomposing the temperature field into spherical harmonics:

$$\frac{\Delta T}{\bar{T}}(\theta, \varphi) = \sum_{l,m} a_{lm} Y_{lm}(\theta, \varphi), \quad (2)$$

where the crucial assumption of statistical isotropy implies

$$\langle a_{lm} a_{l'm'}^* \rangle = \delta_{ll'} \delta_{mm'} C_l. \quad (3)$$

So, in this case, all the statistical information is in the coefficients C_l , the angular power spectrum. More generally, we may define

$$C_l = \frac{1}{2l+1} \sum_m \langle |a_{lm}|^2 \rangle. \quad (4)$$

When searching for spots of a given size, we will average the temperature fluctuations in regions of that size. These regions are defined by window functions $W(\theta, \varphi)$. The mean temperature contrast in such a region is

$$\Delta T = \int d\Omega \Delta T(\theta, \varphi) W(\theta, \varphi). \quad (5)$$

In our sense, a *spot* is characterized as follows. When a threshold $\Delta \mathfrak{L}$ is fixed, a *hot spot* is found if $\Delta T \geq \Delta \mathfrak{L}$, whereas a *cold spot* is found if $\Delta T \leq -\Delta \mathfrak{L}$. The characteristic scale for ΔT is the mean temperature contrast for these regions $\Delta T_{\text{rms}} = \sqrt{\langle \Delta T^2 \rangle}$. Clearly, if $\Delta \mathfrak{L} \ll \Delta T_{\text{rms}}$, most regions will be spots, if $\Delta \mathfrak{L} \gg \Delta T_{\text{rms}}$, only a few or none.

The transformation to harmonic space can be done by decomposing the window function $W(\theta, \varphi)$ into spherical harmonics with coefficients W_{lm} and defining

$$W_l = \sum_m |W_{lm}|^2. \quad (6)$$

Together with Eqs. (2) and (3), it is straightforward to calculate

$$\Delta T_{\text{rms}}^2 = \sum_l \frac{2l+1}{4\pi} C_l W_l \bar{T}^2. \quad (7)$$

This result shows that the mean temperature fluctuation ΔT_{rms} is given by the multipoles C_l weighted by W_l . The W_l strongly depend on the angular scale of the regions. Their magnitude will suppress large l values corresponding to scales smaller than the window. By virtue of the addition theorem for spherical harmonics, we can write

$$W_l = \int d\Omega \int d\Omega' W(\theta, \varphi) W(\theta', \varphi') P_l(\cos\Theta). \quad (8)$$




FIG. 1 (color online). Coefficients W_l for the top hat circle window function at scales $a = 1^\circ$ (right plot), 6° (left plot). The plots show which multipoles predominantly determine ΔT_{rms} . For smaller angular scale a , higher l values enter the analysis.

This allows us to calculate the W_l for a chosen window. An example is shown in Fig. 1.

In our case, it is adequate to approximate the sphere by the tangent plane at a region, replacing the direction (θ, φ) by points \mathbf{x} on the plane. For our purposes, it is most convenient to work with top hat windows because they have clear boundaries. This is the easiest way to avoid ambiguities arising from overlapping spots. Exemplary choices may be the top hat circle with window function

$$W(\mathbf{x}) = \frac{1}{\pi R^2} \Theta(R - |\mathbf{x}|) \quad (9)$$

or a square with window function

$$W(\mathbf{x}) = \frac{1}{a^2} \Theta(a - x_1) \Theta(a - x_2). \quad (10)$$

Following Durrer ([15], p. 218), we can approximate the W_l by an angular average over the Fourier transform of $W(\mathbf{x})$ which considerably reduces the computational effort:

$$W_l \approx \frac{1}{2\pi} \int_0^{2\pi} d\alpha |\tilde{W}(l)|^2. \quad (11)$$

For the aforementioned window functions, we can use this equation to easily calculate ΔT_{rms} by Eq. (7). The results are plotted for the Λ CDM best-fit power spectrum in Fig. 2. For the sake of comparability, we use the parameter a which equals the square root of the windows' area; in the case of squares, it simply is the side length. We also show the relative deviation due to the different window functions. We conclude that the result is not sensitive to the exact geometry if the covered surface area is the same.

III. METHOD

Our strategy consists of performing an identical analysis of spot abundances both for observational maps and maps generated from simulations of Gaussian fluctuations. For the simulated maps, we use the best-fit Λ CDM model and a




FIG. 2 (color online). Mean temperature fluctuation for various spot sizes and the Λ CDM power spectrum. The plots for circles and squares are visually indistinguishable. The difference between the result for circles and the result for squares is shown in the second figure.

Gaussian fluctuation model based on the C_l quoted by the WMAP collaboration. The comparison with maps from observation tests Gaussianity.

Because of the excellent data products of the WMAP team available at the legacy archive [16] and the comprehensive HEALPIX package [17,18], it is possible to obtain reliable CMB sky maps and to create maps from Gaussian simulations. We summarize the steps in Sec. III A. We developed an algorithm searching for hot and cold spots (in the sense of Sec. II) within these temperature sky maps. Its working principle and properties are presented in Sec. III B. The treatment of statistical errors is described in Sec. III C.

A. Maps and data preparation

Whenever the original signal is to be extracted from CMB data, it is crucial to minimize the influence of foreground contamination. The frequency dependence of the foreground components (e.g., synchrotron emission, free-free emission, and thermal dust) allows us to reduce the contamination with the help of various foreground models [19]. The WMAP team provides foreground-reduced maps for the Q (35–46 GHz), V (53–69 GHz), and W (82–106 GHz) bands. Since the V band has a better signal-to-

noise ratio than the W band and is less foreground contaminated than the Q band [2], it is the natural choice to use the foreground-reduced V map. Further noise minimization by constructing linear combinations of the maps is possible but does not affect our analysis which focuses on large scales. But still, large parts of the temperature map are unreliable and must be excluded from the analysis. We therefore apply the KQ75 mask, cutting out the contaminated galaxy region and point sources [19]. Finally, the residual monopole and dipole are removed with the HEALPIX routine REMOVE_DIPOLE. Figure 3 shows the foreground-reduced V map and the KQ75 mask.

Gaussian simulations based on some input C_l spectrum and a beam window function are achieved with the help of the SYNFAST HEALPIX facility. These input data can be obtained from the legacy archive. The power spectra we used are the Λ CDM best fit and the original WMAP-5 estimate both shown in Fig. 1 of Nolta *et al.* [1]. Subsequently, we will refer to them by “ Λ CDM” and “WMAP-5” power spectrum for short. We take care of treating simulated and original maps as equally as possible. This necessitates the additional simulation of the instruments’ noise, masking, and removal of monopole and dipole. Since the WMAP design minimizes noise correla-




FIG. 3 (color online). The foreground-reduced V map (temperature contrast in mK) and the KQ75 mask cutting out the contaminated galaxy region and point sources.

tion between neighboring pixels in a map [20], it is legitimate to add white noise with the properties described by the WMAP team at the legacy archive.

When studying possible anisotropy of the CMB, we need a full-sky (unmasked) map. Since the foreground contaminations usually force us to mask parts of the sky, it is not a trivial task to reconstruct the full-sky CMB signal. However, the WMAP team tries to tackle this job by combining the measurements of all bands and merge them into a single (ILC) map of the full sky [19]. The applied procedure is independent of foreground models but has the disadvantage of being doubtful on scales below approximately 10° according to the WMAP product description at the legacy archive. But since we are lacking any better alternative, we employ the 5-year WMAP ILC map for full-sky analyses.

B. Spot-searching algorithm

The primary goal of the algorithm is to count hot and cold spots in CMB sky maps on various scales and temperature contrasts. A typical application will be to plot spot abundances against the threshold on the temperature contrast $\Delta\mathfrak{T}$ for a specific angular scale. This application directly imposes several features the algorithm should have:

- (i) It must define *sectors on the sphere of equal surface area* (for some desired scale). Their mean temperature contrasts will decide whether they are counted as spots.
- (ii) The areas must be chosen such that one can *smoothly scan* through the map. Between two distinct areas, there must exist many others allowing for a smooth transition.
- (iii) Double counting of spots has to be excluded. The easiest way to achieve this is working with *top hat windows* which have clear boundaries. Overlapping spots will be counted as a single.
- (iv) For a statistically satisfactory comparison between observed and simulated CMB maps, the algorithm will have to analyze many sky maps. Given the huge amount of data, one has to implement the algorithm carefully in order to make this *numerically tractable*.

The algorithm is designed such that it allows for an approximate pixelization of the sphere into distinct areas of a given scale. Calculating their temperature contrasts determines the mean temperature fluctuation ΔT_{rms} on that scale. By virtue of the ergodic theorem, this is a good estimate for the ensemble average introduced in Sec. II.

1. Working principle

The first task is to define the sectors S of equal surface area on the sphere satisfying the requirements explained above. We choose them to be intersections of latitude and longitude rings. A latitude ring \mathcal{R}^{lat} consists of all points between two latitude angles θ_0 and θ_1 , a longitude ring

\mathcal{R}^{lon} of all points between two longitude angles φ_0 and φ_1 . The rings have two nice properties. First, as needed for spot searching, one can smoothly go from one ring to any other ring by smoothly changing its boundary angles; second, as needed for calculating ΔT_{rms} , it is an easy task to discretize a sphere into distinct rings. Since sectors are intersections $S = \mathcal{R}^{\text{lat}} \cap \mathcal{R}^{\text{lon}}$, they share these properties. We thereby satisfy the requirement of smooth scanning to all directions.

We impose [meeting the requirement (i) above] equal area A for all sectors:

$$A = \int_S d\Omega = \int_{\varphi_0}^{\varphi_1} d\varphi \int_{\theta_0}^{\theta_1} d\theta \sin\theta. \quad (12)$$

Once we have decided to define sectors like this, we still have some freedom to choose the boundaries θ_0 , θ_1 , φ_0 , φ_1 . In order to avoid the influence of small scales, we must reasonably choose the sectors such that they are by no means degenerated. We therefore fix this freedom by adding another constraint. For any sector S , the boundary lines in the north-south direction and the longer boundary in the east-west direction are chosen to be of equal length:

$$(\varphi_1 - \varphi_0) \sin\theta_* = (\theta_1 - \theta_0). \quad (13)$$

On the northern hemisphere $\theta_* = \theta_1$, on the southern hemisphere $\theta_* = \theta_0$. Note that these sectors behave well. In the limiting case near the equator, they correspond to squares in flat space. At the poles ($\theta_0 = 0$ or $\theta_1 = \pi$), they become equilateral triangles.

In practice though, the temperature field is not given as a smooth function of θ and φ . The WMAP temperature sky maps are lists assigning a temperature contrast ΔT_i to each HEALPIX pixel p_i . The mapping $p_i \mapsto (\theta, \varphi)$ is given in the form of a table. But since our approach defines sectors by means of angles, we need the reverse. Given the list $p_i \mapsto (\theta, \varphi)$, finding the appropriate pixel p_i for given angles (θ, φ) corresponds to searching through the list. Whereas searching in an unsorted list is very expensive, an adequate sorting may considerably reduce the effort. The algorithm performs the following steps starting at the north pole $\theta_0 = 0$:

- (1) For given θ_0 and area A , calculate θ_1 and $\Delta\varphi$ by solving Eqs. (12) and (13).
- (2) Collect the pixels $\{p_i\}$ belonging to the latitude ring \mathcal{R}^{lat} between θ_0 and θ_1 . This can be done efficiently if the map was prepared by transforming to sorted latitude angles (HEALPIX RING ordering).
- (3) Using a fast routine, sort the list $\{p_i\}$ with respect to longitude angles. This new sorting allows one to directly identify the pixels out of $\{p_i\}$ belonging to a longitude ring \mathcal{R}^{lon} with boundaries φ_0 and φ_1 —these pixels form the sector $S = \mathcal{R}^{\text{lat}} \cap \mathcal{R}^{\text{lon}}$. Start at $\varphi_0 = 0$ and $\varphi_1 = \Delta\varphi$ and smoothly scan (by increasing φ_0, φ_1 by a small step size h) through all longitude rings. For every sector, calculate the

sector's mean temperature contrast ΔT by averaging over the pixel values ΔT_i and compare it with the threshold $\Delta \mathfrak{T}$. If it exceeds the threshold, count a spot if the sector does not overlap with a previously found spot.

- (4) Choose the next ring by slightly increasing $\theta_0 \mapsto \theta_0 + h$. It is profitable to exploit the fact that the sorting for longitude angles (point 3) need not be repeated completely. The algorithm saves the previous sorting and uses it for a presort such that as much information is transferred as possible.

Having increased the threshold $\Delta \mathfrak{T}$, again searching for spot abundances in a map can be optimized by noticing that spots at a higher threshold cannot be found where there were not spots at a lower threshold. Our algorithm can focus on areas around previously found spots once this becomes advantageous.

If we slightly adapt the algorithm, we can use it to measure ΔT_{rms} . Now, the algorithm jumps between distinct sectors instead of smoothly transforming them. The distinct sectors are visualized in Fig. 4. In every distinct sector, the mean temperature fluctuation is calculated. The squares are averaged to give ΔT_{rms} . Although the shapes of the sectors vary, the results of Sec. II ensure that ΔT_{rms} is only marginally affected.

2. Treatment of masked maps

The sectors defined by our algorithm may include none, some, or many masked pixels. We must define selection rules determining which sectors are to be included in the statistics. We used the following two rules. The most restrictive choice is to only consider sectors with no mask overlap (*strict selection* for short). These sectors will only contain reliable pixels. But since especially on large scales, only a minority of sectors will belong to this group, bad statistics are the price to pay. The alternative choice is to also consider sectors with a slight mask overlap (*tolerant selection*). This is a compromise between good statistics on the one hand and reliable results on the other.




FIG. 4 (color online). Exemplary decomposition of the sky into N_{sec} distinct sectors \mathcal{S}_j for measuring ΔT_{rms} . For searching spots, the algorithm analyzes many more sectors \mathcal{S} (those in between, sharing pixels with the illustrated sectors \mathcal{S}_j). Nonetheless, N_{sec} limits the maximum number of spots since overlapping spots are not multiply counted.

We typically allow for 5% masked area within a sector which guarantees that usually the majority of sectors fall into this group. In any case, we emphasize that masked pixels, even if included in the statistics, are assigned zero temperature fluctuation. This will avoid misinterpreting foregrounds as a CMB signal. Note however, that the pixels of zero temperature fluctuation reduce ΔT_{rms} . For comparisons between observed maps and Gaussian simulations, we employ the tolerant selection for the sake of better statistics; the comparison is still trustworthy.

3. Alternative shapes

The algorithm works with the shapes defined in Sec. III B 1 and illustrated in Fig. 4. But we can easily treat other shapes by embedding them into the previous sectors. This corresponds to a multiplication of the previous window function W_0 with the window function W_1 of the desired shape where W_0 must be large enough to ensure $W_0 \equiv 1$ where W_1 is nonzero. The condition (12) of equal area now concerns the new shape and reads

$$\int_{\varphi_0}^{\varphi_1} d\varphi \int_{\theta_0}^{\theta_1} d\theta \sin\theta W_1(\theta, \varphi) = A. \quad (14)$$

As an example, we compare the standard shape with top hat circles of equal area [cf. Eq. (9)] and plot the result in Fig. 5. For low thresholds, the abundances are systematically higher for the standard window function. This is due to the fact that circles do not exhaust the area without space in between. The effect becomes important where many spots are found and overlapping is frequent but disappears for large thresholds where the results agree.

4. Step size dependence

In the ideal case, the boundary angles of the sectors would vary in a perfectly smooth manner when searching




FIG. 5 (color online). Mean spot abundances in 100 simulated Λ CDM full-sky maps showing the results for different window functions of scale $a = \sqrt{A} = 6^\circ$.

for spots in a map. But numerically, we have to choose a finite step size h (introduced in Sec. III B 1). A good choice of h balances sensitivity and numerical effort. Figure 6 shows detected spot abundances against h in simulated maps. We chose $h = 0.3^\circ$ for which we conclude that our sensitivity to detect spots is satisfactory.

C. Errors and cosmic variance

There are statistical uncertainties simply due to the finite number of Gaussian simulations. Moreover, the CMB signal itself can be regarded as the outcome of a statistical process. It is therefore subject to statistical variation, quantified by the concept of cosmic variance.

Let us assume that N Gaussian maps are analyzed for spots (area and threshold fixed). If $n^{(k)}$ spots are detected in map k , the mean spot abundance is

$$\bar{n} = \sum_{k=1}^N \frac{n^{(k)}}{N}. \quad (15)$$

The statistical uncertainty of the mean value \bar{n} and the statistical deviation of the single values $\bar{n}^{(k)}$ are

$$\sigma_{\bar{n}}^2 = \frac{\sum_{k=1}^N (n^{(k)} - \bar{n})^2}{N(N-1)}, \quad \sigma_{n^{(k)}}^2 = N\sigma_{\bar{n}}^2. \quad (16)$$

The same procedure applies if we measure the mean temperature fluctuations $\Delta T_{\text{rms}}^{(k)}$ in the maps and calculate a mean value $\Delta \bar{T}_{\text{rms}}^{(k)}$.

We now consider cosmic variance. When we observe a spot abundance n , we must expect a certain deviation from the theoretically predicted ensemble average $\langle n \rangle$. The expectation value of this deviation, $\sigma_n^2 = \langle (n - \langle n \rangle)^2 \rangle$, quantifies cosmic variance. For a very large number N of simulated maps, we may replace the ensemble average




FIG. 6 (color online). Mean spot abundances for a fixed threshold ($80 \mu\text{K}$) against a varying step size. 100 masked ΛCDM simulated maps were scanned for $a = 6^\circ$, the error bars quantify the statistical error.

by an averaging over the set of simulations. We then obtain $\sigma_n \approx \sigma_{n^{(k)}}$ with the latter calculated according to Eq. (16). This can be done equally for the mean temperature contrast ΔT_{rms} . Whenever we specify cosmic variance (e.g., in the form of error bars), we estimated it by this method.

IV. RESULTS

The application of the spot-searching algorithm described in Sec. III shows that the standard model ΛCDM together with Gaussianity predicts more large-scale hot and cold spots than are actually present in cut-sky WMAP-5 data (see Sec. IV A). Removing the quadrupole or using the original WMAP-5 C_l (instead of the ΛCDM fit) considerably reduces the discrepancies. While only 0.16%–0.62% of Gaussian ΛCDM simulations fall below the observed mean temperature fluctuations on angular scales of 4° – 8° , this increases to 2.5%–8% if the quadrupole is removed. We also investigate full-sky maps in Sec. IV B and modifications of the first multipoles in Sec. IV C.

A. Cut-sky maps

The spots' size is characterized by their area A . We use the parameter $a = \sqrt{A}$ to specify the angular scale of this size. Since on the one hand, we aim at large scales, and on the other hand, we want reasonable statistics, we are forced to find a compromise. We chose an angular scale $a = 6^\circ$. The spot abundances of the WMAP-5 V map and 500 Gaussian ΛCDM simulations (created as described in Sec. III A) are found for varying threshold $\Delta \mathcal{T}$. The HEALPIX resolution parameter of the maps is 8, corresponding to $N_{\text{pix}} = 12 \times 256^2 = 786,432$ pixels. Statistical uncertainties and cosmic variance are displayed as error bars even though the spot abundances for different thresholds are of course correlated. The results for hot and cold spots are plotted in Fig. 7. The striking feature of the plots is the discrepancy between theory and observation. They only agree in the limit of very small thresholds $\Delta \mathcal{T}$ where it is obvious that almost every area is counted as a spot anyway. The discrepancy is seemingly more drastic for hot spots. In the plot for cold spots, it is seen that there is one considerable cold spot nearly reaching $150 \mu\text{K}$. But even this spot does not surpass the ΛCDM prediction. We note that this spot is localized in the region of the famous Vielva cold spot [21]. It is insightful to look at the spot distributions of single Gaussian simulations in order to get an impression of their typical behavior. Five examples are plotted in Fig. 8.

Because of the strong correlation between the spot abundances n_i for different thresholds $\Delta \mathcal{T}_i$, it is difficult to judge the significance of the discrepancies by eye. A possible quantity that can be used for a comparison of the observed CMB map with Gaussian simulations is obtained by summing up the spot abundances at different thresholds,

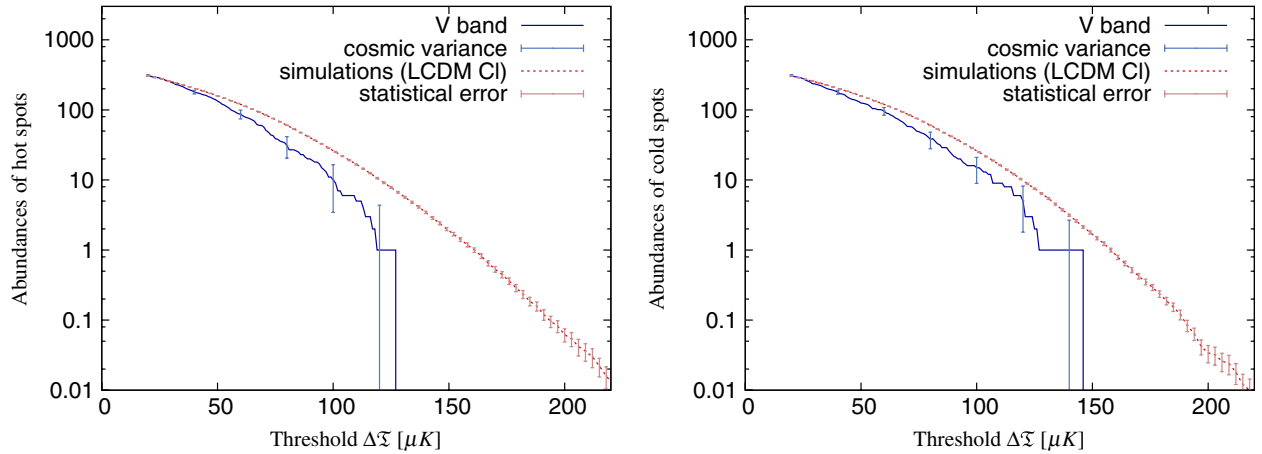


FIG. 7 (color online). Spot abundances in the CMB sky (with cosmic variance) as compared to 500 Λ CDM simulations (with statistical errors) on an angular scale of $a = 6^\circ$. The fractions of Gaussian simulations with smaller values of s [Eq. (17)] are $p_s^{\text{hot}} = 0.2\%$ and $p_s^{\text{cold}} = 1.8\%$.

$$s = \sum_i n_i, \quad (17)$$

where the lowest threshold included is chosen to be the characteristic scale $\Delta\bar{T}_{\text{rms}}$. We denote the fraction of Gaussian simulations k with $s^{(k)}$ smaller than found in the V map by p_s . For the spot abundances shown in Fig. 7, we find $p_s^{\text{hot}} = 0.2\%$ for hot spots and $p_s^{\text{cold}} = 1.8\%$ for cold spots.

The discrepancies are reflected in the mean temperature fluctuation ΔT_{rms} which on large scales is higher in Λ CDM simulations than in the observed CMB sky. We have simulated 5000 Λ CDM maps and compared their mean temperature fluctuations to the value of the V map. We employed the tolerant selection of our algorithm (see Sec. III B 2). For $a = 6^\circ$, we find the mean value $\Delta T_{\text{rms}} = 39.4 \mu\text{K}$ for the V map, as compared to the mean value




FIG. 8 (color online). Spot abundances in five randomly chosen Gaussian simulations based on the Λ CDM best-fit power spectrum and the mean curve from Fig. 7 (hot spots).

$\Delta\bar{T}_{\text{rms}} = 47.9 \pm 0.1 \mu\text{K}$ for Λ CDM, where the error is only statistical while cosmic variance amounts to $4.2 \mu\text{K}$. Only a fraction $p = 0.6\%$ of the simulations had a smaller ΔT_{rms} than the V map. This does not improve at other large angular scales which can be seen in Table I. It is interesting to see how this behavior changes when going to smaller scales. However, the results on smaller scales (approaching 1°) become sensitive to noise and beam properties. Since the WMAP team offers the latter for the differencing assemblies $V1$ and $V2$ [22] instead of the combined V map, it is the easiest to switch to the $V1$ -band map and simulations thereof. The impact on large scales is negligible. Figure 9 shows ΔT_{rms} against the scale a for the $V1$ map and Λ CDM simulations (again with tolerant selection). We see that the deviations decrease when going to smaller scales. This is also suggested by the C_l spectrum which is in good agreement with the Λ CDM fit for large l which dominate on small scales. But still, Monteserin *et al.* [23] find a too low CMB variance which approximately corresponds to ΔT_{rms} on scales even smaller than 1° .

For the results in Fig. 9, we used the highest available HEALPIX resolution 9 corresponding to $N_{\text{pix}} = 12 \times 512^2 = 3,145,728$ pixels in a map. As stated above, the plots are highly influenced by the beam function and noise. The beam function acts as an extra window function which

TABLE I. The fraction p of Gaussian Λ CDM simulations with a ΔT_{rms} smaller than found in the V map on the angular scale a .

Scale a	Fraction p
4°	0.50%
5°	0.62%
6°	0.60%
7°	0.16%
8°	0.36%




FIG. 9 (color online). The mean temperature fluctuation for different angular scales a in 50 Gaussian Λ CDM simulations and the V1 map.

TABLE II. The fraction p of 1000 Gaussian Λ CDM simulations with a ΔT_{rms} smaller than found in the V map on the angular scale a , after removing the quadrupole from the maps.

Scale a	Fraction p
4°	6.5%
5°	8.0%
6°	7.3%
7°	2.5%
8°	6.4%

suppresses the growth of ΔT_{rms} for small scales. The white noise instead leads to a diverging $1/a$ behavior on the smallest scales (with an effective pixel noise amplitude σ_{pix} and the number of pixels $N_a = N_{\text{pix}} \times a^2/4\pi$ within a sector of scale a , the noise contribution will be $\Delta T_{\text{rms}}^{\text{noise}} = \sigma_{\text{pix}}/\sqrt{N_a} \propto 1/a$).

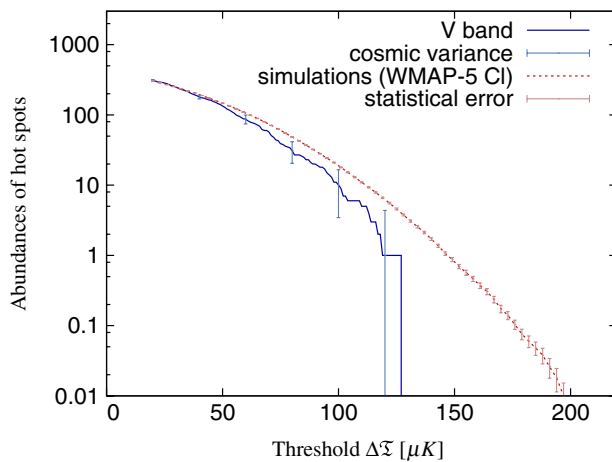


TABLE III. The fraction p of 1000 Gaussian simulations (WMAP-5 C_l) with a ΔT_{rms} smaller than found in the V map on the angular scale a .

Scale a	Fraction p
4°	4.2%
5°	5.4%
6°	5.8%
7°	2.4%
8°	4.1%

On large scales, the first multipoles of the C_l spectrum play an important role (see, e.g., Fig. 1). It is therefore a natural idea to suspect the well-known quadrupole anomaly [2] to be responsible for the observed discrepancies. We check this by repeating the analysis after removing the quadrupole from the Λ CDM simulations as well as the observed CMB map. The results, summarized in Table II, confirm the influence of the quadrupole anomaly. Now, the fractions p of Gaussian Λ CDM simulations reach $p = 7.3\%$ for $a = 6^\circ$. These numbers still do not show good agreement, but they are not statistically significant anymore.

We now investigate whether there are still discrepancies if we compare the observed V map with Gaussian simulations based on the original WMAP-5 C_l spectrum rather than the Λ CDM best fit. This tests whether the observed map is a typical Gaussian realization of the WMAP-5 power spectrum. Figure 10 shows the spot abundances. The effect arising from changing the power spectrum is clearly visible and reduces the discrepancies to some extent. But although closer to the spot abundances in the observed cut-sky CMB map, the numbers of hot and cold spots are still too high. Again, this is reflected in the fact that most simulated maps have a larger ΔT_{rms} than the V map. Even though the values, listed in Table III, are less




FIG. 10 (color online). Spot abundances in the CMB sky (with cosmic variance) as compared to 500 simulations (with statistical errors) based on the WMAP-5 C_l spectrum on an angular scale of $a = 6^\circ$. The fractions of Gaussian simulations with smaller values of s [Eq. (17)] are $p_s^{\text{hot}} = 3.4\%$ and $p_s^{\text{cold}} = 13.2\%$.

drastic, we emphasize that the WMAP-5 estimation of the C_l relies on similar data, i.e., cut-sky CMB maps. If the observed CMB map was a typical Gaussian realization of the extracted C_l spectrum, we would expect agreement.

Bearing in mind, however, that power spectra refer to the full sky whereas we only look at regions outside the mask, an explanation could be that the missing spots were located in the hidden part of the sky. In the next section, we investigate whether the WMAP-5 ILC map indicates this violation of isotropy.

B. ILC full-sky map

The five-year ILC map is the best approximate full-sky CMB map available. We therefore analyze it even though the quality of the reconstruction is not high enough to guarantee robustness of the results (see, also, Sec. III A). We analyze the ILC full-sky map and 100 Gaussian full-sky simulations based on the WMAP-5 power spectrum and separately consider the results in three sky regions. First, we analyze the full sky. Second, we collect the spots

of those regions that have also been studied in the V map, i.e., regions with no or little overlap with the KQ75 mask (tolerant selection). Finally, we consider the remaining spots that consequently lie in sectors completely inside the mask or with considerable mask overlap (rejected by tolerant selection). We loosely refer to the three regions as *full sky*, *outside*, and *inside mask*. The results are plotted in Fig. 11. In the previously analyzed region (outside the mask), we see too few spots, as before. But there are by far too many spots in the complementary region. The variances providing the error bars are, as always, obtained from Eq. (16). Although there are less statistics inside the mask than in the full sky, the error bars in the corresponding figure are smaller. This can be intuitively understood as follows. If, for simplicity, we assumed that the spot abundances outside and inside the mask were statistically independent, the variances σ_{in}^2 , σ_{out}^2 would add to σ_{full}^2 in the full sky, whence $\sigma_{\text{in}} < \sigma_{\text{full}}$. The loss of statistics when counting spots inside the masked region only causes the *relative* fluctuations between two Gaussian simulations to

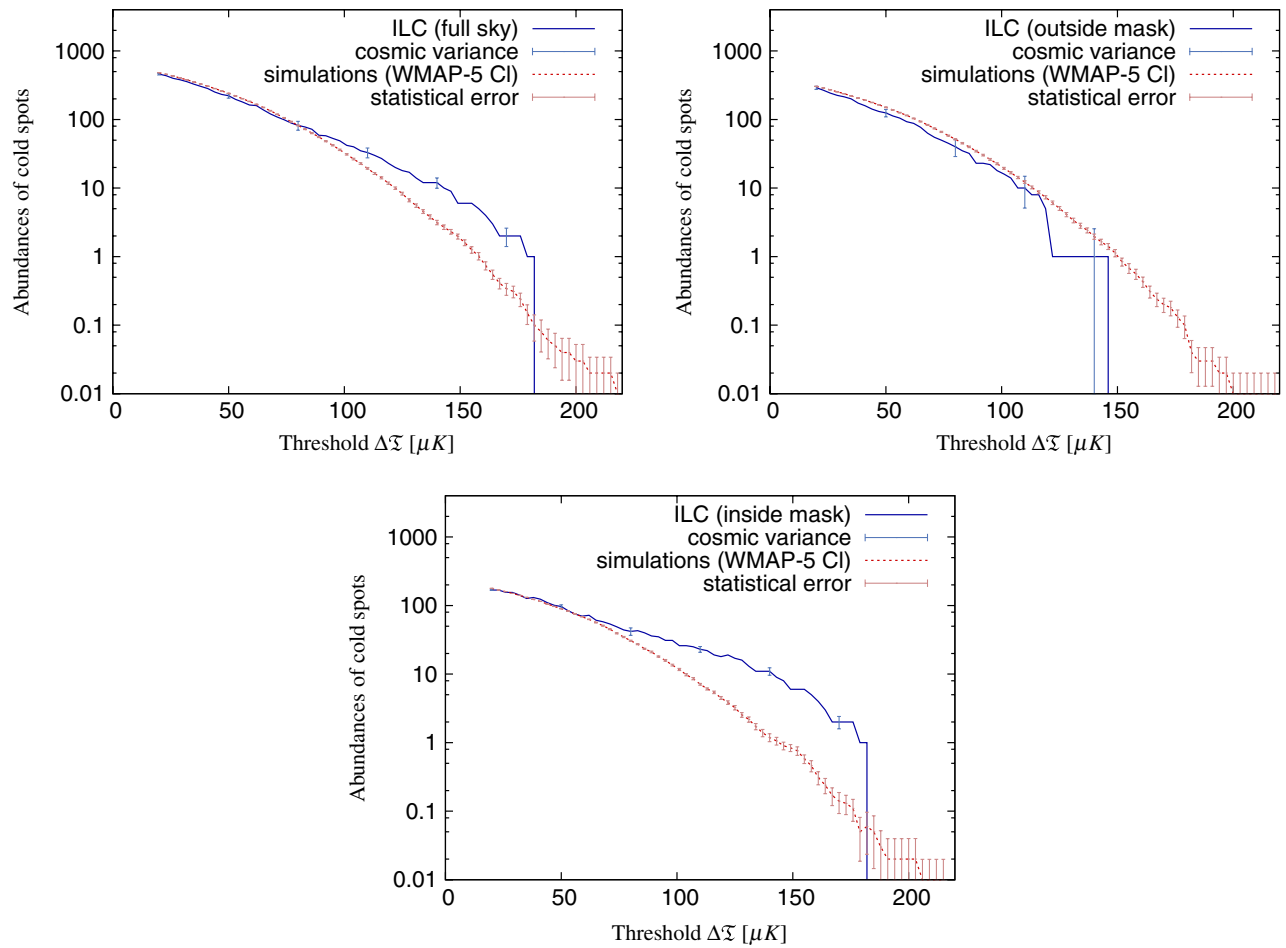


FIG. 11 (color online). Spot abundances in the ILC map (with cosmic variance) compared to simulations (with statistical errors) based on the measured WMAP-5 C_l on an angular scale of $a = 6^\circ$ for three different parts of the sky. The corresponding values of p_s are $p_s^{\text{full}} = 58\%$, $p_s^{\text{out}} = 7\%$, and $p_s^{\text{in}} = 96\%$.

increase. The error bars in the central figure (outside the mask) visually appear larger due to the logarithmic plotting but are in fact smaller than for the full sky.

The values of p_s confirm the uneven distribution of spots in the ILC map. For the full sky, we have $p_s^{\text{full}} = 58\%$ in good agreement with the simulations. Outside the mask, there are too few spots, $p_s^{\text{out}} = 7\%$, whereas inside the mask, we find $p_s^{\text{in}} = 96\%$. The ILC map is clearly anisotropic. Other authors draw the same conclusion [24–26].

Anisotropy of the CMB is a possible explanation of the discrepancies revealed in Sec. IV A and quantified in Table I, and indeed, the ILC map contains this anisotropy. But since there is not enough reliable information about the CMB signal in the galactic plane, we cannot finally judge whether this is the true solution to the problem. We have also studied if, additional to the galactic plane, the orientation of the galactic halo defines a preferred direction. Therefore, we divided the ILC map into two halves, one around the galactic center and one covering the opposite direction. We have seen no signal of anisotropy in this direction.

C. Modified power spectra

We have pointed out that anisotropy is a potential explanation. It is however unsatisfying to assume that so many additional spots lie in the contaminated regions hidden by the KQ75 mask. This would be a surprising coincidence of CMB signals and the orientation of the galactic plane. Alternatively, we may stick to statistical isotropy; then, our results may be due to some non-Gaussian signal.

In this section, we investigate whether our results imply non-Gaussianity or statistical anisotropy by themselves. We do this by analyzing the effect of modifications to the C_l spectrum.

So far, ΔT_{rms} has proved to be a good parameter to quantify the visible effects. We can perform a quick check whether our data supports the hypothesis that ΔT_{rms} is the decisive parameter. Out of the 500 simulations with the WMAP-5 power spectrum used in Sec. IV A, we collect those with a ΔT_{rms} smaller or equal than those found in the V map. Figure 12 shows their spot abundances which agree well with the V map.

If there is a C_l spectrum that produces ΔT_{rms} values similar to the ones found in the V map, our results alone do not imply non-Gaussianity or statistical anisotropy. In order to keep the analysis as generic as possible, we do not use any specific cosmological model but only modify the C_l of the original WMAP-5 spectrum. Figure 9 suggests that only large scales are affected which is why we concentrate on a few low multipoles l . Copi *et al.* [25] found that the correlation function is essentially zero on angular scales above $\approx 60^\circ$. Since this scale is roughly linked to the multipole range $l \leq 3$, our first modification simply consists in setting $C_l \equiv 0$ for $l \leq 3$ (although of course the correlation function does not translate this easily). Another example may be to halve the C_l for $l \leq 5$ (modification II). Figure 13 shows the resulting values of ΔT_{rms} . The plots show the discrepancies between the Λ CDM prediction, the WMAP-5 spectrum, and observation. We also show the results for a combined power spectrum, replacing the first 32 multipoles by the values quoted by WMAP-1 [27]. For this range of multipoles the WMAP analysis changed after the 1-year release, following the suggestion of Efstathiou [28]. The difference between WMAP-5 and WMAP-1 may serve as an illustration that the extraction of reliable C_l for low l is a nontrivial task. Modifications I and II of the power spectrum succeeded in reconciling Gaussian simulations and observed CMB sky. This is confirmed by measuring the spot abundances in simulated maps based on the modified spectra, as

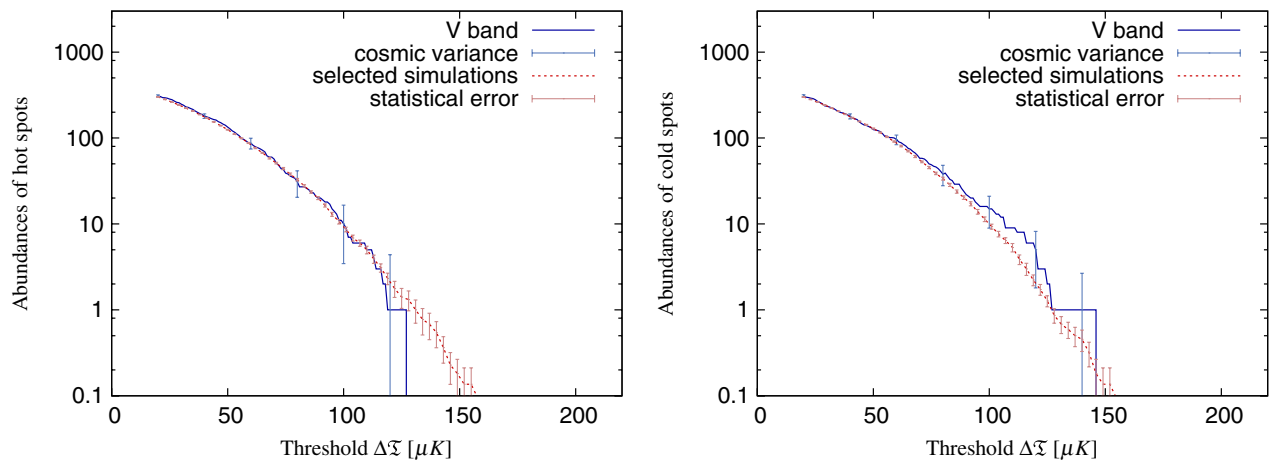


FIG. 12 (color online). Spot abundances of Gaussian simulations k (errors statistical) with $\Delta T_{\text{rms}}^{(k)} \leq \Delta T_{\text{rms}}^{\text{Vmap}}$ in comparison to the V map (with cosmic variance). For these plots, we have $p_s^{\text{hot}} = 68\%$ and $p_s^{\text{cold}} = 86\%$, showing agreement. Since we only consider Gaussian simulations with $\Delta T_{\text{rms}}^{(k)}$ smaller than in the V map, it is no surprise that the p_s values lie above 50%.

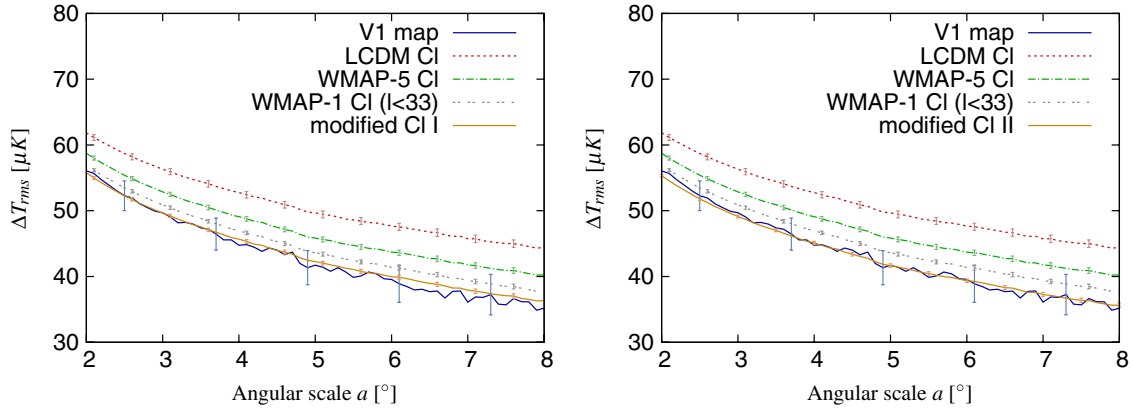


FIG. 13 (color online). The mean temperature fluctuation for large angular scales a . We compare the V1 map with Λ CDM simulations (highest ΔT_{rms}), simulations based on the WMAP-5 C_l (a bit lower), the WMAP-1 C_l for $l < 33$ (still lower), and on two modified spectra. The first modification is created by setting $C_l = 0$ for $l \leq 3$, the second by halving the C_l for $l \leq 5$. The modified spectra agree well with the V1 map.

seen in Fig. 14. We conclude that our results are not incompatible with Gaussianity. However, if we stick to Gaussianity, they favor (although statistically not signifi-

cant, cf. Table III) even lower values of the first multipoles than currently estimated.

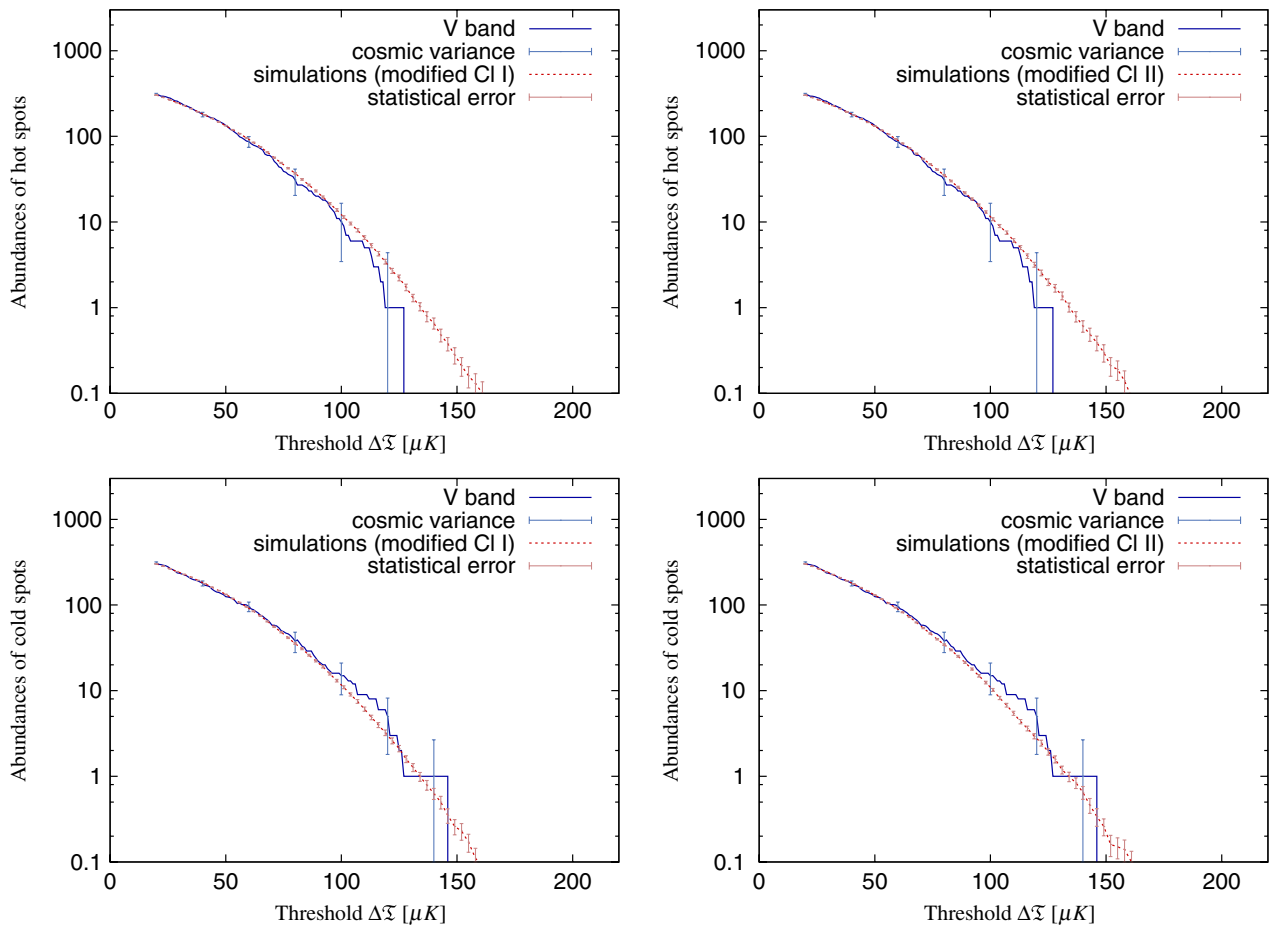


FIG. 14 (color online). Spot abundances in the CMB sky (with cosmic variance) as compared to 100 simulations based on the two modified spectra, respectively, (errors statistical) on an angular scale of $a = 6^\circ$. The first modification yields $p_s^{\text{hot}} = 30\%$ and $p_s^{\text{cold}} = 55\%$, the second modification $p_s^{\text{hot}} = 37\%$ and $p_s^{\text{cold}} = 67\%$.

V. CONCLUSIONS

The study of spot abundances has revealed discrepancies between the cut-sky CMB temperature maps and the standard best-fit Λ CDM model or, but less significant, a Gaussian spectrum for the C_l estimated by WMAP-5. We have shown in Sec. IV C that a good parameter to quantify them is the mean temperature fluctuation ΔT_{rms} which we investigated on large scales. On scales a between 4° and 8° , only 0.16% to 0.62% of Gaussian simulations based on the Λ CDM best-fit power spectrum fall below the ΔT_{rms} value of the observed V map. If this merely was an imprint of the anomalously low quadrupole, we would expect the discrepancies to disappear when removing the quadrupole from the Gaussian simulations and the V map. The difference in fact reduces, the aforementioned fractions change to 2.5% to 8.0%. These numbers are not significant and do not allow for a clear interpretation whether our results go beyond the quadrupole anomaly. Similar fractions are obtained when exchanging the Λ CDM best-fit spectrum by the originally published WMAP-5 C_l , yielding 2.4% to 5.8%. This is difficult to understand if we bear in mind that the C_l themselves are estimated from the cut-sky CMB maps [1].

Non-Gaussianity and also statistical anisotropy are possible explanations. In our case, anisotropy means that many spots have to be hidden behind the masked region. Unfortunately, this hypothesis can hardly be tested as there is currently no method to reliably extract the CMB signal in the highly foreground-contaminated regions. Nonetheless, we have employed the WMAP-5 ILC full-sky CMB map and found evidence for anisotropy in this map. This agrees with results obtained by Hajian [26] and Copi *et al.* [25] who found that most of the power on the largest scales comes from the (masked) galaxy region. Though possible, this unnatural alignment of the CMB signal with the galactic plane would be intriguing and lacks so far any explanation.

Our analysis of Sec. IV C shows that our results for cut-sky maps do not suggest non-Gaussianity or statistical

anisotropy by themselves. They agree well with Gaussian fluctuations if one performs a modification of the lowest multipoles. In doing so, no fine-tuning of the C_l is necessary in order to reconcile the spot abundances from Gaussian simulations and the observed CMB. It is sufficient to lower the first multipoles by a substantial amount. When studying local extrema in the temperature field, Hou *et al.* [12] similarly found discrepancies that disappeared when excluding the first multipoles. We recall, however, that the C_l and the assumption of Gaussianity completely fix the expected spot abundances. If both the extraction of the C_l by WMAP-5 and our analysis of spot abundances are correct, our results may indicate non-Gaussianity or statistical anisotropy.

If the discrepancies are not caused by mere statistical coincidence or unknown secondary effects, we have to leave open the question whether we see the consequence of non-Gaussianity or anisotropy, or whether our results strengthen the evidence for a severe lack of large-scale power. The first case would challenge fundamental assumptions, the second would make it difficult to understand the CMB maps on large scales within standard Λ CDM cosmology. If the discrepancies between the C_l , as determined by WMAP-5, and the spot abundances persist, this can be interpreted as a signal for non-Gaussian fluctuations.

ACKNOWLEDGMENTS

We would like to thank Christian T. Byrnes for useful discussions. We also thank the WMAP team for producing great data products and publishing them on LAMBDA (the Legacy Archive for Microwave Background Data Analysis). Support for LAMBDA is provided by the NASA Office of Space Science. We acknowledge the use of the HEALPIX package [17] that we employed for many tasks, most notably the creation and preparation of Gaussian simulations.

-
- [1] M. R. Nolta *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **180**, 296 (2009).
 - [2] G. Hinshaw *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **170**, 288 (2007).
 - [3] J. D. McEwen, M. P. Hobson, A. N. Lasenby, and D. J. Mortlock, arXiv:0803.2157.
 - [4] A. P. S. Yadav and B. D. Wandelt, *Phys. Rev. Lett.* **100**, 181301 (2008).
 - [5] A. Bernui *et al.*, *Astron. Astrophys.* **454**, 409 (2006).
 - [6] A. de Oliveira-Costa, M. Tegmark, M. Zaldarriaga, and A. Hamilton, *Phys. Rev. D* **69**, 063516 (2004).
 - [7] H. K. Eriksen *et al.*, *Astrophys. J.* **605**, 14 (2004).
 - [8] F. K. Hansen *et al.*, *Astrophys. J.* **704**, 1448 (2009).
 - [9] J. Hoftuft *et al.*, *Astrophys. J.* **699**, 985 (2009).
 - [10] K. Land and J. Magueijo, *Phys. Rev. Lett.* **95**, 071301 (2005).
 - [11] P. Cabella *et al.*, *Phys. Rev. D* **69**, 063007 (2004).
 - [12] Z. Hou, A. J. Banday, and K. M. Gorski, arXiv:0903.4446.
 - [13] D. L. Larson and B. D. Wandelt, *Astrophys. J.* **613**, L85 (2004).
 - [14] D. L. Larson and B. D. Wandelt, arXiv:astro-ph/0505046.
 - [15] R. Durrer, *The Cosmic Microwave Background* (Cambridge University Press, Cambridge, England, 2008).
 - [16] <http://cmbdata.gsfc.nasa.gov>.

- [17] K. M. Gorski *et al.*, *Astrophys. J.* **622**, 759 (2005).
- [18] <http://healpix.jpl.nasa.gov>.
- [19] B. Gold *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **180**, 265 (2009).
- [20] L. Page *et al.*, *Astrophys. J.* **585**, 566 (2003).
- [21] P. Vielva *et al.*, *Astrophys. J.* **609**, 22 (2004).
- [22] R. S. Hill *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **180**, 246 (2009).
- [23] C. Monteserin *et al.*, *Mon. Not. R. Astron. Soc.* **387**, 209 (2008).
- [24] A. Bernui and M. J. Reboucas, *Phys. Rev. D* **79**, 063528 (2009).
- [25] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, *Mon. Not. R. Astron. Soc.* **399**, 295 (2009).
- [26] A. Hajian, arXiv:astro-ph/0702723.
- [27] G. Hinshaw *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **148**, 135 (2003).
- [28] G. Efstathiou, *Mon. Not. R. Astron. Soc.* **348**, 885 (2004).