EVOLUTION OF THE COSMIC MICROWAVE BACKGROUND POWER SPECTRUM ACROSS *WILKINSON MICROWAVE ANISOTROPY PROBE* DATA RELEASES: A NONPARAMETRIC ANALYSIS

AMIR AGHAMOUSA¹, MIHIR ARJUNWADKAR^{1,2}, AND TARUN SOURADEEP³

¹ Centre for Modeling and Simulation, University of Pune, Pune 411 007, India; amir@cms.unipune.ac.in, mihir@ncra.tifr.res.in

² National Centre for Radio Astrophysics, University of Pune Campus, Pune 411 007, India

³ Inter-University Centre for Astronomy and Astrophysics, University of Pune Campus, Pune 411 007, India; tarun@iucaa.ernet.in

Received 2011 July 20; accepted 2011 October 24; published 2012 January 6

ABSTRACT

Using a nonparametric function estimation methodology, we present a comparative analysis of the Wilkinson Microwave Anisotropy Probe (WMAP) 1-, 3-, 5-, and 7-year data releases for the cosmic microwave background (CMB) angular power spectrum with respect to the following key questions. (1) How well is the power spectrum determined by the data alone? (2) How well is the ACDM model supported by a model-independent, data-driven analysis? (3) What are the realistic uncertainties on peak/dip locations and heights? Our results show that the height of the power spectrum is well determined by data alone for multipole l approximately less than 546 (1-year), 667 (3-year), 804 (5-year), and 842 (7-year data). We show that parametric fits based on the ACDM model are remarkably close to our nonparametric fits in *l*-regions where data are sufficiently precise. In contrast, the power spectrum for an HACDM model is progressively pushed away from our nonparametric fit as data quality improves with successive data realizations, suggesting incompatibility of this particular cosmological model with respect to the WMAP data sets. We present uncertainties on peak/dip locations and heights at the 95% (2σ) level of confidence and show how these uncertainties translate into hyperbolic "bands" on the acoustic scale (l_A) and peak shift (ϕ_m) parameters. Based on the confidence set for the 7-year data, we argue that the low-l upturn in the CMB power spectrum cannot be ruled out at any confidence level in excess of about 10% ($\approx 0.12\sigma$). Additional outcomes of this work are a numerical formulation for minimization of a noise-weighted risk function subject to monotonicity constraints, a prescription for obtaining nonparametric fits that are closer to cosmological expectations on smoothness, and a method for sampling cosmologically meaningful power spectrum variations from the confidence set of a nonparametric fit.

Key words: cosmic background radiation – cosmological parameters – methods: data analysis – methods: statistical

Online-only material: color figures

1. INTRODUCTION

The angular power spectrum of cosmic microwave background (CMB) temperature fluctuations is a measurable physical quantity that depends sensitively on the physics of the early universe. In particular, the shape of the angular power spectrum and the locations and heights of its peaks relate directly to parameters in the underlying cosmological models. As such, it has been used extensively as an acid test of the relative merit of competing cosmological models and as a rich source of information about cosmological parameters themselves.

Traditionally, and almost exclusively, cosmologists have resorted to model-based parametric statistical methods for estimating the CMB angular power spectrum from data. Parametric regression methods require the functional form of the unknown regression function f to be pre-specified.

The adjustable parameters in f, finite in number that is independent of the data size, are usually estimated by maximizing an appropriate likelihood function or a posterior distribution. In the cosmological context, it is conventional to employ parametric models that attempt to capture the essential physics of the problem via the pre-specified functional form, and any pre-existing knowledge about parameters is incorporated in the estimation process via appropriate prior distributions.

Nonparametric function estimation methods, on the other hand, assume no specific functional form for f, except for mild regularity conditions such as smoothness assumptions, membership to a function space, etc. In this approach, the number of parameters that define the unknown regression function f is either infinite or grows proportionally with the data size, and the estimate \widehat{f}_N of f is obtained by balancing bias and variance of f_N via optimal smoothing. Nonparametric methods are therefore model-independent and are based on fewer and less restrictive assumptions about f. This, in turn, implies that any inferences about f made from nonparametric regression analyses tend to be more data-driven as opposed to being primarily model-driven. In other words, to a greater extent nonparametric function estimation methods tend to infer what is rather than what should be. As such, nonparametric regression methods can be meaningfully employed as sanity-enforcing mechanisms on parametric analyses, thereby making the conclusions drawn more conservative. For example, a feature seen in a parametric fit that survives in a nonparametric analysis is likely to be a real and robust feature of the data itself, and not merely an artifact that is seen because a parametric model expects it to be there.

Alternative methodologies, such as the nonparametric methodology (Genovese et al. 2004; Bryan et al. 2007) used in this work, may allow posing inferential questions that are difficult to address using conventional methods. For example, this particular nonparametric methodology allows validating modelbased, parametric fits against the confidence set for the nonparametric fit to the same data, possibly to rule them out as candidates for the true but unknown regression function. This methodology also has the formal advantage of being able to provide realistic uncertainties on *any* number of features of the angular power spectrum that are simultaneously valid at the same level of confidence. Such desirable features are arguably lacking in most mainstream methodologies, including Bayesian ones, that are commonly used in cosmology. (An incisive, insightful, and discerning discussion about the relative merits of this methodology over statistical methods conventionally employed in cosmology can be found in the two references cited above and is best read in the original.)

The four CMB angular power spectrum data sets (Hinshaw et al. 2003b, 2007; Nolta et al. 2009; Larson et al. 2011) released by the Wilkinson Microwave Anisotropy Probe (WMAP; Bennett et al. 2003) mission, representing cumulative observations at the end of 1, 3, 5, and 7 years of operation, present a unique opportunity for statistical analysis. For example, till date these are claimed to represent the most precise and extensive fullsky CMB measurements ever made (only to be superseded by the *Planck* mission (Tauber et al. 2010)). The four data sets represent the evolution of data over a period of about a decade, thereby making it possible to assess progressive and possibly systematic resolution of features of the spectrum. From a statistical perspective, each of these moderately large data sets (minimum of 899 data points for the WMAP 1-year release) is not only heteroskedastic, but also has substantial correlations that arise in typical data pipelines (Tegmark 1997; Hinshaw et al. 2003a, 2009; Jarosik et al. 2007; Jarosik et al. 2011).

In this paper, we present a comparative nonparametric analysis of the WMAP 1-, 3-, 5-, and 7-year data sets for the angular power spectrum. Specifically, for each data realization, we address the following three key questions. (1) How well is the angular power spectrum determined by the data alone? (2) How well is the ACDM model supported by a model-independent, nonparametric, data-driven analysis? (3) What are the realistic uncertainties on peak/dip locations and heights? Our analysis is based on a nonparametric function estimation methodology (Genovese et al. 2004; Bryan et al. 2007), which is discussed in Section 2 together with our extensions; i.e., a numerical formulation for minimization of a inverse-noise-weighted risk function subject to monotonicity constraints, a prescription for obtaining nonparametric fits that are closer to cosmological expectations on smoothness, and a method for sampling cosmologically meaningful power spectrum variations from the confidence set of a nonparametric fit. Results are presented in Section 3, and conclusion in Section 4.

2. METHODOLOGY

The nonparametric function estimation methodology (Genovese et al. 2004; Bryan et al. 2007) used in this work is an extension of the REACT methodology (Beran 2000a, 2000b), which is in turn founded on rigorous formal results in Beran & Dümbgen (1998) and Beran (1996).

Two early papers (The Pittsburgh Institute for Computational Astrostatistics 2003; Miller et al. 2002) used this methodology to analyze, under the assumption of homoskedasticity, a pre-WMAP data set that combined BOOMERanG, MAXIMA, and DASI data sets. A generalization of this formalism for dealing with heteroskedasticity via an inverse-noise-weighted loss function was developed in Genovese et al. (2004). More recently, using the WMAP 1-year data, Bryan et al. (2007) illustrated how confidence intervals on cosmological parameters and boundaries in the cosmological parameter space can be inferred from the confidence set for a nonparametric power spectrum fit.

In this section, we first present an operational outline of this methodology (Sections 2.1 and 2.2). This outline is entirely

based on Genovese et al. (2004) and Bryan et al. (2007) and is included here for completeness. A pedagogic treatment of the central ideas and a simpler variant of the problem can be found in Wasserman (2006). Specific citations to other sources are provided wherever appropriate.

Our own numerical formulation of the monotone risk minimization problem, where the risk function is derived from an inverse-noise-weighted loss function, is presented in Section 2.3. In Section 2.4, we describe a systematic way of obtaining a monotone fit with smoothness that meets cosmological expectations. In Section 2.5, we describe our method for probing the confidence set; this is the basis for the results presented in Sections 3.1 and 3.3.

2.1. The Nonparametric Fit

We are given CMB angular power spectrum data of the form

$$Y_l = C_l + \epsilon_l \tag{1}$$

consisting of *N* data points observed over multipole index range $l_{\min} \leq l \leq l_{\max}$. Here, C_l stands for the value of the true but unknown power spectrum at multipole index *l*. The noise variables ϵ_l are assumed to have a mean-0 normal distribution with *known* covariance matrix Σ . In practice, any reasonable estimate/approximation $\hat{\Sigma}$ of this covariance matrix, such as an inverse Fisher matrix for a pilot parametric fit, can be used in place of Σ .

This nonparametric regression method is based on expanding the unknown regression function f, assumed to belong to an appropriate L_2 function space, in a complete orthonormal basis $\{\phi_i(x)\}$, as

$$f(x) = \sum_{j=0}^{\infty} \beta_j \phi_j(x)$$

A basis that has proven useful in the CMB angular power spectrum context is the cosine basis defined over $0 \le x \le 1$:

$$\phi_j(x) = \begin{cases} 1 & (j=0) \\ \sqrt{2}\cos(j\pi x) & (j=1,2,\ldots) \end{cases}$$
(2)

Assuming that *f* is sufficiently smooth, we take

$$f(x) \approx \sum_{j=0}^{N-1} \beta_j \phi_j(x),$$

and estimate it as

$$\widehat{f}_N(x) = \sum_{j=0}^{N-1} \widehat{\beta}_j \phi_j(x).$$
(3)

While the method is asymptotically (i.e., as the data size $N \rightarrow \infty$) basis-independent, choice of the basis may matter in any finite-*N* application; see Beran (2000a, 2000b) for a detailed discussion. This basis satisfies a discrete orthogonality property when the data Y_i are available over an equispaced grid $\{x_i = (2i + 1)/2N, 0 \le i \le N - 1\}$ consisting of zeros of $\phi_N(x)$. In the CMB context, any contiguous range of *N* integer multipole indices $l_{\min} \le l \le l_{\max}$ can be formally mapped onto this equispaced grid, hence we will not make any categorical distinction between data index *i* and the corresponding multipole index *l*. The true angular power spectrum $C_l \equiv f(x_l)$ is estimated as $\widehat{C}_l \equiv \widehat{f}_N(x_l)$ via coefficient estimates $\widehat{\beta}_j$, which are estimated as $\widehat{\beta}_i = \lambda_i Z_i,$

where

$$Z_{j} = \frac{1}{N} \sum_{i=0}^{N-1} Y_{i} \phi_{j}(x_{i}) = \frac{(U^{T}Y)_{j}}{\sqrt{N}}.$$
 (5)

(4)

Here, U is the orthonormal matrix with elements U_{ij} = $\phi_j(x_i)/\sqrt{N}$ and $Y \equiv (Y_0, \ldots, Y_{N-1})^T$.

The task of obtaining coefficient estimates $\hat{\beta} \equiv$ $(\widehat{\beta}_0, \ldots, \widehat{\beta}_{N-1})^T$, and thereby the fit $\widehat{f}_N(x_i)$, is now relegated to determining the shrinkage parameters λ_i . Assuming smoothness for f (which implies a rapid decay of the true coefficients β_i with j), the shrinkage parameters λ_i are constrained to be monotonically decreasing

$$1 \ge \lambda_0 \ge \lambda_1 \ge \ldots \ge \lambda_{N-1} \ge 0$$
 (Monotone shrinkage). (6)

A special, discrete subset of the monotone shrinkage defined above is the nested subset selection (NSS) shrinkage, defined as

$$\lambda_j = \begin{cases} 1 & \text{for } 0 \leqslant j < J \\ 0 & \text{for } J \leqslant j < N \end{cases}$$
(NSS shrinkage). (7)

Either shrinkage type results in selective damping of highfrequency harmonics in the data Y_i , which results in smoothing of the fit $\widehat{f_N}$. A useful interpretation of shrinkage parameter λ_i is that it represents the effective degree of freedom for the jth coefficient estimate $\hat{\beta}_i$. One can thus define the effective degrees of freedom (EDoF) for the entire fit f_N as

$$\text{EDoF}(\lambda) = \sum_{j=0}^{N-1} \lambda_j.$$
(8)

This definition follows from the fact that for a linear smoother, EDoF of a fit is formally defined as $tr(\mathcal{H})$, where \mathcal{H} is the matrix that connects the fitted values \widehat{Y} to the data Y as $\widehat{Y} = \mathcal{H}Y$. For the present nonparametric regression method, $\mathcal{H} = UDU^T$, where U is the orthonormal basis matrix (Equation (5)), and $D \equiv$ diag $(\lambda_0, \ldots, \lambda_{N-1})$. This implies tr $(\mathcal{H}) = tr(D) = \sum_{i=0}^{N-1} \lambda_i$.

In the present formalism, the discrepancy between the (unknown) regression function f and its estimator \widehat{f}_N is measured by the inverse-noise-weighted squared loss function $L(\lambda)$, defined as

$$L(\lambda) = \int_0^1 \left(\frac{f(x) - \hat{f}_N(x)}{\sigma(x)}\right)^2 dx.$$

Here, $\sigma^2(x)$ is the (known) variance of the data Y at x, which accounts for the heteroskedasticity of the data Y. The loss L is considered a function of the vector of shrinkage parameters $\lambda \equiv (\lambda_0, \dots, \lambda_{N-1})^T$ that determine the regression estimator \widehat{f}_N via Equation (4). Risk $R(\lambda)$, which is the expected value of $L(\lambda)$, can be written as a sum of two non-negative terms; namely,

$$R(\lambda) = \int_0^1 \left(\frac{f(x) - \mathbb{E}(\widehat{f}_N(x))}{\sigma(x)}\right)^2 dx + \int_0^1 \mathbb{E}\left[\left(\frac{\widehat{f}_N(x) - \mathbb{E}(\widehat{f}_N(x))}{\sigma(x)}\right)^2\right] dx.$$

These two terms represent, respectively, the integrated squared bias and the integrated variance of $f_N(x)$, both weighted by $1/\sigma^2(x)$. Optimal smoothing is achieved, in principle, by minimizing $R(\lambda)$ with respect to λ . Generally speaking, too little smoothing leads to a fit \widehat{f}_N with low bias and high variance, and too much smoothing yields a fit with high bias and low variance. Minimal risk or optimal smoothing therefore can be thought of as a balance between the bias of \hat{f}_N and its variance.

The risk function $R(\lambda)$, unfortunately, depends on the unknown regression function f, and therefore needs to be estimated. A particular estimator of this risk, which is of the SURE (Stein's unbiased risk estimator; Stein 1981) kind, takes the following form:

$$\widehat{R}(\lambda) = Z^T \overline{D} W \overline{D} Z + \operatorname{tr}(D W D B) - \operatorname{tr}(\overline{D} W \overline{D} B), \quad (9)$$

where $Z \equiv (Z_0, \dots, Z_{N-1})^T$, $D \equiv \text{diag}(\lambda_0, \dots, \lambda_{N-1})$, $\overline{D} =$ $I_N - D$, $B = U^T \Sigma U / N$ is the covariance of Z, and I_N is the $N \times N$ identity matrix. The positive (semi)definite weight matrix W is defined as

$$W_{jk} = \sum_{l} \Delta_{jkl} w_l, \tag{10}$$

where w_l is the *l*th coefficient in the expansion $(1/\sigma^2(x)) \approx$ $\sum_{i=0}^{N-1} w_j \phi_j(x)$ and, for the cosine basis (Equation (2)),

$$\begin{split} \Delta_{jkl} &= \int_0^1 \phi_j(x) \phi_k(x) \phi_l(x) \, dx \\ &= \begin{cases} 1, & \text{if } \#\{j, k, l=0\} = 3, \\ 0, & \text{if } \#\{j, k, l=0\} = 2, \\ \delta_{jk} \delta_{0l} + \delta_{jl} \delta_{0k} + \delta_{kl} \delta_{0j}, & \text{if } \#\{j, k, l=0\} = 1, \\ \frac{1}{\sqrt{2}} (\delta_{l,j+k} + \delta_{l,|j-k|}) & \text{if } \#\{j, k, l=0\} = 0. \end{cases} \end{split}$$

We denote the optimal shrinkage obtained by minimizing risk $R(\lambda)$ by λ . The best NSS shrinkage $\lambda_{\rm NSS}$ is obtained simply by evaluating risk $\widehat{R}(\lambda)$ for each of the (N + 1) NSS shrinkage vectors and choosing the one with the least risk. Monotone shrinkage usually results in a lower risk than the NSS shrinkage because of the greater freedom available in the monotone set of shrinkages. We will discuss risk minimization subject to monotonicity constraints (Equation (6)) in Section 2.3.

Figure 1 illustrates the contrasting behavior of the nonparametric risk (red curve) and the WMAP 1-year likelihood function (green curve; Verde et al. 2003) for the WMAP 1-year data (Hinshaw et al. 2003b), as a function of the EDoF of all NSS fits. This figure is motivated by the fact that cosmologists, by and large, are better-acquainted with parametric likelihood-based methods. Each integer value on the horizontal axis represents one NSS fit, from the zero function at EDoF = 0 to the fit that simply interpolates through the data (EDoF = N). Optimal smoothing for NSS shrinkage occurs at EDoF = 12 where the nonparametric risk function attains its minimum over the NSS set of fits. The likelihood function, on the other hand, keeps on improving with the EDoF indefinitely.

2.2. Confidence Set Around the Fit

Conventional regression methods provide a confidence band around the fit that quantifies the uncertainty in the fit. In contrast, this nonparametric methodology quantifies the uncertainty surrounding the nonparametric fit in the form of an elegant construct, namely, a $(1 - \alpha)$ confidence set at a pre-specified confidence level $0 \leq (1 - \alpha) \leq 1$. The $(1 - \alpha)$ confidence set for the



Figure 1. Nonparametric risk (red curve) and $-2 \log(\text{likelihood})$ (blue curve) as functions of the EDoF (Equation (8)) for the NSS fits to the *WMAP* 1-year data. This illustrates the contrasting behavior of the two quantities. Optimal smoothing occurs at EDoF = 12 where the nonparametric risk attains its minimum over the NSS set of fits. The likelihood function, on the other hand, keeps on improving with the EDoF indefinitely. log(likelihood) values were computed using the *WMAP* 1-year likelihood code (Verde et al. 2003). The blue (left) and red (right) vertical scales on the plot are associated with the nonparametric risk and the $-2 \log(\text{likelihood})$, respectively.

coefficient vector β is defined as

$$\mathcal{D}_{N,\alpha} = \left\{ \beta : (\beta - \widehat{\beta})^T W(\beta - \widehat{\beta}) \leqslant r_{\alpha}^2 \right\},\tag{11}$$

which is centered at the vector of estimated coefficients $\hat{\beta}$, and the *confidence radius* r_{α} is given by

$$r_{\alpha}^{2} = \frac{\widehat{\tau} z_{\alpha}}{\sqrt{N}} + \widehat{R}(\widehat{\lambda}).$$
(12)

Here, z_{α} is the upper α quantile of the standard normal distribution, and

$$\widehat{\tau}^2/N = 2\mathrm{tr}(ABAB) + Z^T Q Z - \mathrm{tr}(QB), \qquad (13)$$

with Q = 4(ABA + WDBDW - 2ABDW) and A = DW + WD - W. In practice, the risk estimator (Equation (9)) and/or $\hat{\tau}^2$ (Equation (13)) may turn out to be negative for particular data/covariance matrix realizations. In such cases, the squared confidence radius (Equation (12)) may be negative (or may not be 0 for $\alpha = 0$). For minimization purposes, the risk estimator (Equation (9)) is adequate and appropriate (Beran & Dümbgen 1998). For confidence radius purposes, we suggest the following modifications to avoid the negativity problem:

$$\widehat{R}_{+} = Z^{T} \overline{D} W \overline{D} Z + \max\{0, \operatorname{tr}(DWDB) - \operatorname{tr}(\overline{D} W \overline{D} B)\}$$

$$\widehat{\tau}_{+}^{2} / N = 2\operatorname{tr}(ABAB) + \max\{0, Z^{T} Q Z - \operatorname{tr}(QB)\}$$

$$r_{\alpha+}^{2} = \max\left\{0, \frac{\widehat{\tau}_{+} z_{\alpha}}{\sqrt{N}} + \widehat{R}_{+}\right\}.$$
(14)

At worst, this adjustment will make the confidence radius bigger, resulting in, e.g., wider confidence intervals, but more conservative inferences. Similar modifications have been suggested in Beran & Dümbgen (1998), Beran (2000a), and Genovese et al. (2004).

The corresponding confidence set on the true regression function f is given by

$$\mathcal{B}_{N,\alpha} = \left\{ f(x) = \sum_{j=0}^{N-1} \beta_j \phi_j(x) : \beta \in \mathcal{D}_{N,\alpha} \right\}.$$
 (15)

The quadratic form of the inverse noise-weighted loss function and the fact that the weight matrix W is positive (semi)definite implies that both confidence sets $\mathcal{D}_{N,\alpha}$ and $\mathcal{B}_{N,\alpha}$ are ellipsoidal in shape. For any functional T of the spectrum f, such as location or height of a peak or a dip, the $(1 - \alpha)$ confidence interval is defined as

$$\mathcal{I}_{N,\alpha} = \left(\min_{f \in \mathcal{B}_{N,\alpha}} T(f), \max_{f \in \mathcal{B}_{N,\alpha}} T(f)\right).$$
(16)

Moreover, prior information that the true regression function f belongs to a subset $\mathcal{P}_{N,\alpha}$ of the confidence set $\mathcal{D}_{N,\alpha}$ (e.g., f has k peaks over the range of x-values represented in the data) can be incorporated in the analysis by replacing $\mathcal{D}_{N,\alpha}$ with $\mathcal{P}_{N,\alpha} \cap \mathcal{D}_{N,\alpha}$. This methodology further provides the formal assurance that, asymptotically,

- 1. $\mathcal{B}_{N,\alpha}(\mathcal{D}_{N,\alpha})$ will contain the true spectrum f (true coefficient vector β) with probability $\geq (1 \alpha)$, and
- 2. confidence intervals $\mathcal{I}_{N,\alpha}$ on *any number of* functionals T(f), computed from the confidence set $\mathcal{B}_{N,\alpha}$, will be *simultaneously valid* at the same confidence level (1α) , and that these will trap their corresponding true but unknown values with probability $\ge (1 \alpha)$.

2.3. Risk Minimization Subject to Monotonicity Constraints

In this section, we show how the risk function $\widehat{R}(\lambda)$ (Equation (9)) can be minimized subject to the monotonicity constraints $1 \ge \lambda_0 \ge \lambda_1 \ge \ldots \ge \lambda_{N-1} \ge 0$. The risk function corresponding to the unweighted loss function ($W = I_N$) has a simple weighted-sum-of-squares form, and can be minimized exactly and efficiently using the pooled adjacent violators (PAV) algorithm (Robertson et al. 1988). While the risk function corresponding to the inverse-noise-weighted loss function ($W \ne I_N$) is still quadratic in λ , it can no longer be expressed as a weighted sum-of-squares and the PAV algorithm cannot be used to minimize it.

It can be shown that, disregarding terms that do not depend on λ , the risk function $\widehat{R}(\lambda)$ (Equation (9)) can be written as

$$\widehat{R}(\lambda) = \frac{1}{2}\lambda^T H\lambda - \lambda^T h,$$

where $H_{jk} = 2z_j z_k W_{jk}$, $h = (H - V)(1, 1, ..., 1)^T$, and $V_{jk} = 2W_{jk}B_{kj}$. *H* and *V* are both manifestly symmetric. Positive (semi)definiteness of *W* implies that *H* is a positive (semi)definite matrix, implying that $\hat{R}(\lambda)$ is a convex function. The system of linear inequality constraints $1 \ge \lambda_0 \ge \lambda_1 \ge$ $\dots \ge \lambda_{N-1} \ge 0$ implies that the constrained region (Figure 2) has a convex trianguloidal shape determined by flat surfaces. The original risk minimization problem can therefore be formulated as the following equivalent convex quadratic minimization problem:

Minimize
$$\widehat{R}(\lambda) = \frac{1}{2}\lambda^T H\lambda - \lambda^T h$$

subject to $C\lambda \leq (0, 0, \dots, 0)^T$ (17)
and $0 \leq \lambda_i \leq 1$ for all i ,

where C is the $(N - 1) \times N$ matrix

$$C = \begin{bmatrix} -1 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & -1 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & -1 & 1 & 0 \\ 0 & 0 & 0 & \dots & 0 & -1 & 1 \end{bmatrix}.$$



Figure 2. Trianguloid-shaped admissible regions, marked by red lines, for the monotonicity constraint $1 \ge \lambda_0 \ge \lambda_1 \ge ... \ge \lambda_{N-1} \ge 0$ for N = 2 (left) and 3 (right). The (N + 1) vertices of the trianguloid correspond to the (N + 1) NSS fits, with the origin corresponding to the zero function, and the vertex (1, 1, ..., 1) corresponding to the function that exactly interpolates through the data. Surfaces with constant value *p* of the EDoF (Equation (8)) are hyperplanes of the form $\sum_{i=0}^{N-1} \lambda_i = p$. (A color version of this figure is available in the online journal.)

An additional linear inequality or equality constraint of the form

$$\sum_{i=0}^{N-1} \lambda_i \leqslant q \quad \text{or} \quad \sum_{i=0}^{N-1} \lambda_i = q, \tag{18}$$

where *q* constrains the EDoF of the fit $(0 < q \leq N)$, can easily be accommodated in this formulation. This reformulation of the monotone risk minimization problem makes it possible to use standard minimization methods (Pshenichny & Danilin 1978; Powell 1985; Goldfarb & Idnani 1983) and computational tools (Schittkowski 2007; Vanderbei 1999) for convex quadratic minimization problems subject to linear constraints and simple bounds.

2.4. Obtaining a Smoother Monotone Fit That is Closer to Cosmological Expectations

To motivate the discussion in this section, consider the fullfreedom monotone fit to the *WMAP* data sets obtained as above (green curves in Figures 3 and 4). By *full-freedom monotone fit*, we mean the fit that minimizes risk over the entire monotoneadmissible region (Figure 2) without any restriction on the EDoF of the fit. This fit turns out to be quite wiggly especially at the high-*l* end because of the high noise variance here. Such wiggliness implies the presence of high-frequency components in the fit, which in turn implies a large number of EDoF in the fit.

Without cosmological pre-conditioning (i.e., from a completely agnostic and data-driven perspective) and when viewed in relation to the data, it is clear that this fit is not at all unreasonable, given the high noise levels at the high-*l* end. However, all cosmological models suggest far smoother shapes for the angular power spectrum. In the context of the present methodology, one candidate for a smoother fit is the NSS fit (i.e., one that has the minimal risk over the set of (N + 1) NSS fits). Indeed, this possibility has been exploited, e.g., in Genovese et al. (2004). However, the NSS fit (see the blue curve in Figures 3 and 4) may also turn out to be somewhat unsatisfactory with respect to cosmological expectations (and, some times, also with respect to trends reflected in the full-freedom monotone fit). This is primarily because of the limited freedom available in the NSS class. Note again that the NSS fit is not entirely unreasonable from an agnostic viewpoint.

The monotone set, on the other hand, offers the possibility of harnessing local minima in the risk function that are constrained to lie in appropriate "smoother" subsets of the full monotone set. This may be achieved in two distinct ways that may be combined for greater effect.

- 1. By imposing one of the additional constraints (Equation (18)) on the EDoF of the fit. Examples of such restricted-freedom monotone fits are the red curves in Figures 3 and 4.
- 2. By truncating the expansion (Equation (3)) to p number of coefficients (p < N) and then performing monotone risk minimization over this subset of the full monotone set.

In practice, such a smoother restricted-freedom monotone fit can be obtained by gradually reducing the value of q(Equation (18)) starting from the EDoF of the NSS fit until all low-amplitude, high-frequency wiggles in the fit disappear. Generally, the resulting fit has a lower risk than the NSS fit with EDoF = q and is manifestly consistent with trends captured by the full-freedom monotone fit. We find it useful to present (or consider) all three fits (NSS, full-freedom monotone, and restricted-freedom smoother monotone) together; this helps build a realistic picture about estimated trends in the data and thereby about the shape of the underlying true spectrum. Like the NSS fit, the smoother restricted-freedom monotone fit will generally be more biased than the full-freedom monotone fit. This greater bias, however, is partially compensated for by a larger risk which results in a larger confidence radius value (Equation (12)), a larger confidence set, and therefore more conservative inferences.

2.5. Probing the Confidence Set for Uncertainties on Features of the Fit

In this paper, we need to probe the confidence set for a fit for two purposes: (1) for validating cosmological models against a nonparametric fit (see Section 3.2) and (2) for finding the uncertainties on specific features of the fit such as peak heights and locations, and below we describe our method to scan and sample the confidence set for the latter. Our particular method for probing the confidence set for determining uncertainties on features of the fit is based on the following observations.



Figure 3. Nonparametric fits for *WMAP* 1- (top) and 3-year (bottom) data sets. *x*- and *y*-ranges are identical across plots. Green: full-freedom monotone fit (EDoF \approx 80.2, 76.5, respectively); blue: NSS fit (EDoF = 12, 10, respectively); red: restricted-freedom monotone fit (EDoF \approx 9.4, 9.5, respectively); solid gray: best ACDM-based parametric fit; dashed gray: power spectrum for an HACDM model. See Table 1 for details of model-based power spectra. (A color version of this figure is available in the online journal.)

- 1. The confidence set $\mathcal{D}_{N,\alpha}$ on the vector of coefficients $(\beta_0, \ldots, \beta_{N-1})$, by construction, is centered at the vector of estimates $(\hat{\beta}_0, \ldots, \hat{\beta}_{N-1})$.
- 2. The confidence interval defined in Equation (16) requires locating extreme variations in any functional T of the power spectrum f; e.g., location or height of a peak. The largest possible variations in T will be located as far away from the center of the confidence set as possible, i.e., on its surface.
- 3. Cosmologically meaningful and sufficiently smooth variations in the fitted spectrum \hat{C}_l are most likely to be located in the projection of the full confidence set onto the lowest *M* dimensions.

We therefore generate a uniform sample from the projection of the full confidence set surface onto the lowest *d* dimensions, where $2 \le d \le M$, with $M \le 23$. For convenience, we use the smoother restricted-freedom monotone fit for this purpose, with the justification that the confidence set corresponding to the full-freedom monotone fit happens to be nested inside that for this fit, for all four data realizations. The Cholesky factorization $W = u^T u$ is used to transform the original confidence ellipsoid $\mathcal{D}_{N,\alpha}$ (whose principal axes may not be aligned with coordinate directions in the β -space) into a sphere of the form $\{\psi : \|\psi - \widehat{\psi}\|^2 \le r_{\alpha}^2\}$, where $\psi = u\beta$ and $\widehat{\psi} = u\widehat{\beta}$. The surface of this ψ -sphere can be efficiently sampled with uniform density using a standard algorithm (Section 3.4.1.E.6 on p.130 of Knuth 1981), and then transformed back into β -space, preserving uniformity of density because of the linearity of the transformation. From a sufficiently large sample of such variations of the power spectrum, we further selected



Figure 4. Nonparametric fits for *WMAP* 5- (top) and 7-year (bottom) data sets. *x*- and *y*-ranges are identical across plots. Green: full-freedom monotone fit (EDoF \approx 60.4, 102.9, respectively); blue: NSS fit (EDoF = 13, 20, respectively); red: restricted-freedom monotone fit (EDoF \approx 14.4, 14.1, respectively); solid gray: best ACDM-based parametric fit; dashed gray: power spectrum for an HACDM model. See Table 1 for details of model-based power spectra. (A color version of this figure is available in the online journal.)

those functions for which successive peaks and dips are separated by at least 50 multipole moments *l*. This cosmologically motivated selection criterion ensures that (1) the sampled functions are sufficiently smooth, and (2) high-frequency wiggles are not counted as peaks or dips when estimating uncertainties on locations and heights of peaks and dips (see Section 3.3). Based on cosmological considerations, we restricted the search to functions with three peaks (*WMAP* 1-, 3-, and 5-year data) or four peaks (7-year data). The set of functions thus sampled is used to estimate uncertainties on specific features of the fit. As an aside, we note that the confidence set construct and the formal guarantees related to confidence intervals (Equation (14)) do not necessarily imply a uniform density over the confidence set; uniform sampling is used here as a convenient computational device for scanning the confidence set surface in an unbiased fashion.

3. RESULTS AND DISCUSSION

The four *WMAP* angular power spectrum data sets used in this work, ACDM parametric fits for the CMB angular power spectrum, and likelihood codes that produce their respective Fisher (inverse covariance) matrices are obtained from the *WMAP* data archive http://lambda.gsfc.nasa.gov/. For all four data realizations, it turned out that (1) the weight matrix *W* (Equation (10)) is numerically positive definite, and (2) the confidence set for the full-freedom monotone fit is completely nested inside that for the smoother restricted-freedom monotone



Figure 5. Results of a probe of the confidence sets for the *WMAP* 1- (black), 3- (blue), 5- (red), and 7-year (green) nonparametric restricted-freedom monotone fits (Figures 3 and 4) to determine how well the angular power spectrum is determined by the data alone. The quantity plotted for each data realization is the total vertical variation at each *l* within the respective 95% (2σ) confidence set, divided by the absolute value of the fit. This quantity is an approximate measure of how well the angular power spectrum is determined by the data: values $\ll 1$ indicate that the fit is tightly determined by the data, whereas values $\gtrsim 1$ indicate that the data contain little or no information about the height of the angular power spectrum for that *l*. Disregarding the low-*l* region, the color-coded vertical lines indicate the approximate *l*-value at which each curve rises above 1.

fit. It is worth noting that our nonparametric fits and confidence sets are not too sensitive to the details of the covariance matrix (this was also pointed out in Bryan et al. 2007). Most computations were done using the R statistical computing environment (R Development Core Team 2010). We used the QL codes (Schittkowski 2007) for the monotone risk minimization problem (Equation (17)). Our nonparametric fits, obtained using the method outlined in Section 2.3, are shown in Figures 3 and 4 for the *WMAP* 1-, 3-, 5-, and 7-year data sets.

3.1. How Well is the Power Spectrum Determined by Data Alone?

Considering that the noise in all four data sets is highly heteroskedastic and noise levels are especially high for large l, it would be useful to make an assessment of how such noise in the data affects local uncertainties in the fitted spectrum and to quantify how well the angular power spectrum value at each l is determined by the data.

To this end, we compute, for each data set, the approximate 95% confidence interval on each \widehat{C}_l using 5000 function variations from the confidence set as outlined in Section 2.5. The length of this vertical confidence interval at given l, divided by the absolute value of the fit, $|\widehat{C}_l|$, provides an approximate indication of how well each \widehat{C}_l is determined via the following interpretation: a value $\ll 1$ indicates that the fit is well determined by the data, whereas values $\gtrsim 1$ imply that the data contain little or no information about the height of the power spectrum. This approach, which is inspired by the boxcar probe approach of Genovese et al. (2004), has the practical advantage of not having adjustable parameters (i.e., the boxcar width) in the procedure.

In Figure 5, we plot this height, scaled by the value of the fit, as a function of the multipole index l, for all four data realizations. We see that the range of l-values over which the fit is well determined expands consistently between 1-, 3-, and 5-year data realizations, from $l \approx 546$ (1-year), to $l \approx 667$ (3-year), to $l \approx 804$ (5-year). On the other hand, while the l-range of the data expanded substantially between *WMAP* 5 and 7, the information contained in the data does not appear to have grown proportionately beyond $l \approx 842$ for the 7-year data.

3.2. How Well is the Λ CDM Model Supported by a Model-independent, Nonparametric, Data-driven Analysis?

In each of the four plots in Figures 3 and 4, we have also included parametric fits based on the Λ CDM (Hinshaw et al. 2003b, 2007; Nolta et al. 2009; Larson et al. 2011) and H Λ CDM (Primack et al. 1995; Primack & Gross 2001) models (see the figure caption for details). The parametric Λ CDM-based fits turn out to be quite close to the respective nonparametric fits wherever the data are precise. This is remarkable considering that our nonparametric method does not rely on any cosmologically motivated prior information whatsoever. Moreover, the parametric Λ CDM-based fit appears to get closer to the respective nonparametric fit across the four *WMAP* data releases. The closeness of a parametric fit ($\hat{C}_{l_{min}}, \ldots, \hat{C}_{l_{max}}$) to the corresponding nonparametric fit ($\hat{C}_{l_{min}}, \ldots, \hat{C}_{l_{max}}$) can be measured through the distance

$$d(C, \widehat{C}) \approx \sqrt{\frac{1}{N} \sum_{l=l_{\min}}^{l_{\max}} \left(\frac{C_l - \widehat{C}_l}{\sigma_l}\right)^2}.$$

THE ASTROPHYSICAL JOURNAL, 745:114 (12pp), 2012 February 1

 Table 1

 Distances of Two Model-based Power Spectra from Our Nonparametric Fits

Data	ЛСDM	НЛCDM
1-year	0.1540	0.1493
$r_{\alpha} = 0.3679$	16.70%	15.87%
3-year	0.1462	0.2057
$r_{\alpha} = 0.3653$	14.52%	29.03%
5-year	0.1419	0.3552
$r_{\alpha} = 0.3563$	19.66%	94.82%
7-year	0.1238	0.3550
$r_{\alpha} = 0.3551$	9.08%	94.98%

Notes. These are the distances of two model-based (ACDM and HACDM) power spectra from our nonparametric (fullfreedom monotone) fits. The HACDM model (Primack et al. 1995; Primack & Gross 2001) considered here for illustrative purposes is defined by a small neutrino fraction $(\Omega_{\nu}h^2 = 0.00275)$ with corresponding adjustment to the dark energy content ($\Omega_{\Lambda} = 0.729756$), and the rest of the parameters (including zero curvature) being identical to that of the best ACDM model (Larson et al. 2011) for the 7-year data (power spectrum generated using the CAMB software; Lewis et al. 2000). ACDM-based fits used are the best parametric fits for the corresponding data realization (Hinshaw et al. 2003b, 2007; Nolta et al. 2009; Larson et al. 2011). r_{α} is the confidence radius at $\alpha = 0.05$ (i.e., 95% confidence level $\equiv 2\sigma$). Percentages reported are the confidence levels corresponding to these distances; these can be interpreted as asymptotic probabilities with which the corresponding parametric fit is ruled out as a candidate for the true but unknown spectrum. Note the dramatic progression, as the data become precise, of (1) how this HACDM model is pushed to the boundary of the confidence set, and (2) how the ACDM model gets closer to the nonparametric fit.

Using Equation (12), this distance can be further expressed as the smallest confidence level α beyond which the parametric fit is rejected as a candidate for the true spectrum.

Table 1 lists distances of parametric fits based on the two cosmological models (ACDM and HACDM) from the nonparametric full-freedom monotone fit for the corresponding data realization (see the caption for specific details and description). The progression of distance values between a parametric ACDM fit and the corresponding nonparametric fit clearly shows that the two are getting closer as the data become precise. In contrast to this, the angular power spectrum generated by the particular HACDM model considered, which is almost as strong a contender for the true power spectrum as the ACDM fit with respect to the 1-year confidence set, is progressively pushed away to the boundary of the 95% confidence set for the 7-year confidence set. Visually, this trend can be understood on the basis of the differences between the HACDM power spectrum and the nonparametric fit (e.g., differences in the heights of the first peak; see Figures 3 and 4) that result in pushing this particular HACDM model out of the confidence set. Given the formal guarantees of this methodology (Section 2.2), the WMAP 7-year data thus rules out, at $\approx 95\%$ confidence level, the particular HACDM model considered.

3.3. Uncertainties on Locations and Heights of Peaks and Dips

We now consider the problem of determining uncertainties on the locations and heights of peaks and dips in the nonparametrically fitted spectrum. The motivation for this exercise comes from the fact that peak locations and heights contain valuable information about cosmological models and parameters (Doran & Lilley 2002; Durrera et al. 2003).

Following the prescription outlined in Section 2.5, we sampled a set of 5000 function variations from the confidence set for each data realization. Peak and dip locations and heights were recorded for each peak and dip over this set of functions. This results in an empirical scatterplot that is indicative of the joint distribution of location and height for each peak or dip, under the assumption of uniform surface density on the confidence set $\mathcal{D}_{N,\alpha}$.

Figure 6 shows the results of probing the 95% confidence sets for uncertainties on peaks and dips, as outlined above, with 5000 acceptable function variations for each data realization. The box around a peak or a dip represents the largest horizontal and vertical variations in the scatter. In accordance with the confidence interval defined in Equation (16), these form the 95% confidence intervals on the location and height of a peak/ dip. Table 2 lists these confidence intervals together with 95% confidence intervals on peak height ratios.

The following features of these results are worth pointing out. As is well known, the first peak was very clearly resolved in the 1-year data itself. Our results are manifestly consistent with this observation, in the sense that its box does not overlap with any other box. However, the uncertainty on the first peak does not shrink appreciably across the four data realizations. Further, our results clearly indicate that the second peak is resolved cleanly only in the 5-year data, whereas the third and fourth peaks are not resolved completely even in the 7-year data.

3.4. Uncertainties on the Acoustic Scale (l_A) and Peak Shift (ϕ_m) Parameters

Consider the following relationship (Hu et al. 2001; Doran & Lilley 2002) between the location l_m of the *m*th peak, the acoustic scale l_A , and the shift parameter ϕ_m :

$$l_m = l_A(m - \phi_m). \tag{19}$$

Substituting the end-points of the 95% confidence interval for the *m*th peak location, this relationship results in a hyperbolic band of allowed values in the $l_A - \phi_m$ plane. Such bands, derived from 95% confidence intervals on the first three peaks (Table 2), are shown in Figure 7. Additional information from other sources is required to constrain these bands to physically meaningful regions in the $l_A - \phi_m$ plane. For example, if we assume $l_A = 300$ (Page et al. 2003) then, based on the 7-year data, the 95% confidence intervals for ϕ_m will be ϕ_1 : (0.1600, 0.3767), ϕ_2 : (0.0367, 0.3600), ϕ_3 : (-0.2167, 0.7300). Conversely, additional constraints on ϕ_m could be used to generate a confidence interval on l_A . From a model-independent point of view, we note that the (l_A, ϕ_m) bands for different peaks *m* appear to overlap around $\phi_m \approx 0$ and $200 \lesssim l_A \lesssim 400$. We interpret this as a nonparametric realization of the nearly harmonic structure of peaks in the CMB power spectrum.

3.5. The Low-l Upturn from a Nonparametric Viewpoint

Another interesting feature in Figure 6 is the tiny but clearly observable scatter for the very first dip at the low-l end. This scatter corresponds to extreme power spectrum variations that reside on the surface of the 95% confidence set and have an upturn at low l values. In the Λ CDM cosmology, such upturn at the low-l end is primarily the result of the integrated



Figure 6. Nonparametric uncertainties on peak and dip locations and heights for the *WMAP* 1- (top left), 3- (top right), 5- (bottom left), and 7-year (bottom right) data sets. The nonparametric fit displayed for reference is the restricted-freedom monotone fit in Figures 3 and 4. The number of acceptable function variations sampled from the confidence set for each data realization is 5000.

	55% Confidence met vals on bevera i caules of the Angula 1 ower Spectrum							
Data	Peak Location	Peak Height	Dip Location	Dip Height	Peak Location Ratio	Peak Height Ratio		
1-year	l_1 : (186, 252)	h_1 : (4968, 6133)	$l_{1+\frac{1}{2}}$: (378, 507)	$h_{1+\frac{1}{2}}$: (1152, 2029)				
	l_2 : (440, 680)	h_2 : (1879, 3933)	$l_{2+\frac{1}{2}}^{2}$: (494, 772)	$h_{2+\frac{1}{2}}^{2}$: (-2929, 2868)	l_2/l_1 : (1.953, 3.217)	h_2/h_1 : (0.329, 0.724)		
	l_3 : (559, 897)	h_3 : (41, 10733)	$l_{3+\frac{1}{2}}^{2}$: (639, 900)	$h_{3+\frac{1}{2}}^{2}$: (-11376, 9962)	$l_3/l_1:(2.495, 4.385)$	h_3/h_1 : (0.0075, 1.906)		
3-year	l_1 : (184, 254)	h_1 : (5130, 6199)	$l_{1+\frac{1}{2}}$: (384, 456)	$h_{1+\frac{1}{2}}$: (1402, 1917)				
	l_2 : (479, 606)	h_2 : (2172, 2997)	$l_{2+\frac{1}{2}}^{2}$: (593, 856)	$h_{2+\frac{1}{2}}^{2}$: (80, 2145)	l_2/l_1 : (2.028, 3.000)	h_2/h_1 : (0.380, 0.555)		
	l_3 : (646, 977)	h_3 : (1506, 6045)	$l_{3+\frac{1}{2}}^{2}$: (729, 1000)	$h_{3+\frac{1}{2}}^{2}$: (-4501, 5113)	$l_3/l_1:(2.761,4.801)$	h_3/h_1 : (0.264, 1.082)		
5-year	l_1 : (187, 251)	h_1 : (5249, 6301)	$l_{1+\frac{1}{2}}$: (387, 445)	$h_{1+\frac{1}{2}}$: (1489, 1934)				
	l_2 : (485, 596)	h_2 : (2306, 2955)	$l_{2+\frac{1}{2}}^{2}$: (608, 793)	$h_{2+\frac{1}{2}}^{2}$: (971, 2095)	l_2/l_1 : (2.040, 2.963)	h_2/h_1 : (0.381, 0.525)		
	l_3 : (663, 982)	h_3 : (1863, 4635)	$l_{3+\frac{1}{2}}^{2}$: (731, 1000)	$h_{3+\frac{1}{2}}^{2}$: (-2978, 3157)	$l_3/l_1:(2.883,4.672)$	h_3/h_1 : (0.321, 0.833)		
7-year	l_1 : (187, 252)	h_1 : (5177, 6377)	$l_{1+\frac{1}{2}}$: (390, 442)	$h_{1+\frac{1}{2}}$: (1512, 1931)				
	l_2 : (492, 589)	h_2 : (2328, 3015)	$l_{2+\frac{1}{2}}^{2}$: (623, 803)	$h_{2+\frac{1}{2}}^{2}$: (1074, 2063)	l_2/l_1 : (2.060, 2.887)	h_2/h_1 : (0.386, 0.534)		
	l_3 : (681, 965)	h_3 : (1871, 4119)	$l_{3+\frac{1}{2}}^{2}$: (733, 1104)	$h_{3+\frac{1}{2}}^{2}$: (-3111, 2709)	l_3/l_1 : (3.000, 4.553)	h_3/h_1 : (0.323, 0.732)		
	$l_4:(815,1193)$	$h_4:(-1391,7726)$	$l_{4+\frac{1}{2}}^{2}$: (951, 1200)	$h_{4+\frac{1}{2}}^{2}$: (-11102, 6364)	l_4/l_1 : (3.606, 6.047)	$h_4/h_1: (-0.244, 1.340)$		

 Table 2

 95% Confidence Intervals on Several Features of the Angular Power Spectrum

Notes. Negative values for heights and height ratios are a reflection of the high noise at the high-*l* end. Here, l_m (h_m) stands for the location (height) of the *m*th peak, and $l_{m+\frac{1}{2}}$ ($h_{m+\frac{1}{2}}$) denotes the location (depth) of the *m*th dip.

Sachs–Wolfe (ISW) effect and is seen in all parametric ACDM fits in Figures 3 and 4. It would therefore be interesting to see what could be said about the low-*l* upturn (and thereby about the ISW effect) based on the nonparametric confidence set.

Note that our nonparametric fits, which are at the center of their respective confidence sets, do not show a low-l upturn. However, the 7-year parametric Λ CDM fit, e.g., does show a clear upturn at the low-l end. This parametric fit is at a



Figure 7. Confidence "bands" for the acoustic scale l_A and the shift ϕ_m for the *m*th peak, as derived from the 95% confidence intervals on the first three peak locations (Table 2) and Equation (19). Blue: ϕ_1 , red: ϕ_2 , and green: ϕ_3 . Top left: *WMAP* 1-year, top right: 3-year, bottom left: 5-year, and bottom right: 7-year data sets. Note that these (l_A, ϕ_m) bands for different peaks *m* appear to overlap around $\phi_m \approx 0$ and $200 \leq l_A \leq 400$: We interpret this as a nonparametric realization of the nearly harmonic structure of peaks in the CMB power spectrum.

distance corresponding to a confidence level of about 10% ($\approx 0.12\sigma$; see Table 1) from our 7-year nonparametric full-freedom monotone fit. This means that the confidence set for the 7-year nonparametric fit contains spectra with a low-l upturn at most as far away as the 7-year parametric fit. We therefore conclude conservatively that the low-l upturn as a feature of the CMB angular power spectrum cannot be ruled out at any confidence level in excess of about 10% Actually, there are indications in our results (not shown) that such upturned variations of the spectrum may be much closer to the center of the confidence set for the 7-year full-freedom monotone fit; this needs further investigation.

4. CONCLUSION

In this paper, we have presented a comparative nonparametric analysis of the *WMAP* 1-, 3-, 5-, and 7-year data releases for the CMB angular power spectrum, using a nonparametric function estimation methodology (Genovese et al. 2004; Bryan et al. 2007). In the context of this methodology, we have also presented our own numerical formulation for minimization of the inverse-noise-weighted risk function subject to monotonicity constraints, and a prescription for obtaining monotone nonparametric fits that are closer to cosmological expectations on smoothness. For all data realizations, we have presented results pertaining to the following questions. (1) How well is the angular power spectrum determined by the data alone? (2) How well is the ACDM model supported by a model-independent, nonparametric, data-driven analysis? (3) What are the realistic uncertainties on peak/dip locations and heights?

The motivation for the analysis presented here was to explore what could be inferred about the CMB angular power spectrum in a model-independent, data-driven manner. On the other hand, the basic physics of the CMB is quite well established. It would therefore be useful to connect a nonparametric/modelindependent analysis such as ours with the known physics of the CMB angular power spectrum. This is reserved for the future.

To conclude, we have demonstrated in this paper the threefold utility of the nonparametric methodology used here for cosmological function estimation problems: as a method with sound formal guarantees, as a sanity-enforcing mechanism on parametric model-based analyses, and as a method that allows interesting inferential questions to be addressed and answered in a data-driven manner.

M.A. is deeply indebted to Christopher R. Genovese and Larry Wasserman for many enlightening discussions covering all of statistics. Our R codes for computing the nonparametric fit are based on original codes by Christopher R. Genovese. T.S. acknowledges support from the DST Swarnajayanti Fellowship. Insightful questions and suggestions from an anonymous referee helped improve the overall presentation aspects of the paper and, specifically, the method used in Section 3.1.

REFERENCES

- Bennett, C. L., Bay, M., Halpern, M., et al. 2003, ApJ, 583, 1
- Beran, R. 1996, Ann. Inst. Statist. Math., 48, 1
- Beran, R. 2000a, J. Amer. Statist. Assoc., 95, 155
- Beran, R. 2000b, in Asymptotics in Statistics and Probability: Papers in Honor of George Gregory Roussas, ed. M. L. Puri (Leiden: VSP International Science Publishers), 1
- Beran, R., & Dümbgen, L. 1998, Ann. Statist., 26, 1826
- Bryan, B., Schneider, J., Miller, C. J., et al. 2007, ApJ, 665, 25
- Doran, M., & Lilley, M. 2002, MNRAS, 330, 965
- Durrera, R., Novosyadlyj, B., & Apunevych, S. 2003, ApJ, 583, 33
- Genovese, C. R., Miller, C. J., Nichol, R. C., Arjunwadkar, M., & Wasserman, L. 2004, Statist. Sci., 19, 308
- Goldfarb, D., & Idnani, A. 1983, Math. Program., 27, 1
- Hinshaw, G., Barnes, C., Bennett, C. L., et al. 2003a, ApJS, 148, 63
- Hinshaw, G., Nolta, M. R., Bennett, C. L., et al. 2007, ApJS, 170, 288
- Hinshaw, G., Spergel, D. N., Verde, L., et al. 2003b, ApJS, 148, 135
- Hinshaw, G., Weiland, J. L., Hill, R. S., et al. 2009, ApJS, 180, 225
- Hu, W., Fukugita, M., Zaldarriaga, M., & Tegmark, M. 2001, ApJ, 549, 669
- Jarosik, N., Barnes, C., Greason, M. R., et al. 2007, ApJS, 170, 263
- Jarosik, N., Bennett, C. L., Dunkley, J., et al. 2011, ApJS, 192, 14
- Knuth, D. E. 1981, The Art of Computer Programming, Vol. 2, Seminumerical Algorithms (2nd ed.; Reading, MA: Addison-Wesley)
- Larson, D., Dunkley, J., Hinshaw, G., et al. 2011, ApJS, 192, 16
- Lewis, A., Challinor, A., & Lasenby, A. 2000, ApJ, 538, 473
- Miller, C. J., Nichol, R. C., Genovese, C. R., & Wasserman, L. 2002, ApJ, 565, L67

Nolta, M. R., Dunkley, J., Hill, R. S., et al. 2009, ApJS, 180, 296

- Page, L., Nolta, M. R., Barnes, C., et al. 2003, ApJS, 148, 233
- Powell, M. J. D. 1985, in Mathematical Programming Studies, Vol. 25, Mathematical Programming Essays in Honor of George B. Dantzig, Part II, ed. R. W. Cottle et al. (Berlin: Springer), 46
- Primack, J. R., & Gross, M. A. K. 2001, in Current Aspects of Neutrino Physics, ed. D. O. Caldwell (Physics and Astronomy Online Library; Berlin: Springer), 287
- Primack, J. R., Holtzman, J., Klypin, A., & Caldwell, D. O. 1995, Phys. Rev. Lett., 74, 2160
- Pshenichny, B. N., & Danilin, Y. M. 1978, Numerical Methods in Extremal Problems (Moscow: Mir)
- R Development Core Team. 2010, R: A Language and Environment for Statistical Computing (Vienna: R Foundation for Statistical Computing)
- Robertson, T., Wright, F. T., & Dykstra, R. L. 1988, Order Restricted Statistical Inference (New York: Wiley)
- Schittkowski, K. 2007, QL: A Fortran Code for Convex Quadratic Programming - User's Guide, Tech. rep., University of Bayreuth, Department of Computer Science
- Stein, C. M. 1981, Ann. Stat., 9, 1135
- Tauber, J. A., Mandolesi, N., Puget, J.-L., et al. 2010, A&A, 520, A1
- Tegmark, M. 1997, Phys. Rev. D, 55, 5895
- The Pittsburgh Institute for Computational Astrostatistics 2003, in Statistical Challenges in Astronomy, ed. E. D. Feigelson & G. J. Babu (New York: Springer), 221
- Vanderbei, R. J. 1999, Optim. Methods Softw., 11, 451
- Verde, L., Peiris, H. V., Spergel, D. N., et al. 2003, ApJS, 148, 195
- Wasserman, L. 2006, All of Nonparametric Statistics (New York: Springer)