

Properties and use of CMB power spectrum likelihoods

Samira Hamimeche* and Antony Lewis†

Institute of Astronomy and Kavli Institute for Cosmology, Madingley Road, Cambridge, CB3 0HA, United Kingdom

(Received 6 February 2009; published 21 April 2009)

Fast robust methods for calculating likelihoods from cosmic microwave background observations on small scales generally rely on approximations based on a set of power spectrum estimators and their covariances. We investigate the optimality of these approximations, how accurate the covariance needs to be, and how to estimate the covariance from simulations. For a simple case with azimuthal symmetry we compare optimality of hybrid pseudo- C_l cosmic microwave background power spectrum estimators with the exact result, indicating that the loss of information is not negligible, but neither is it enough to have a large effect on standard parameter constraints. We then discuss the number of samples required to estimate the covariance from simulations, with and without a good analytic approximation, and assess the use of shrinkage estimators. Finally we discuss how to combine an approximate high- l likelihood with a more exact low- l harmonic-space likelihood as a practical method for accurate likelihood calculation on all scales.

DOI: 10.1103/PhysRevD.79.083012

PACS numbers: 98.70.Vc

I. INTRODUCTION

The cosmic microwave background (CMB) appears to be isotropic and Gaussian to a good approximation and hence allows robust statistical constraints to be placed on a variety of cosmological parameters. However, high-resolution observations such as those from the Planck satellite produce sky maps with many millions of pixels, for which performing an exact likelihood analysis becomes numerically very difficult. Most data analyses therefore use fast robust approximations based on power spectrum estimators [1,2]. While these are expected to be suboptimal at some level, they should be unbiased on average. The main advantage of a fast method is that multiple simulations can be performed to assess in detail the propagation of errors from numerous instrumental processes into the final results, and hence in practice they may also be significantly more reliable than an in-principle more optimal method.

In a previous paper we investigated in detail various approximations for calculating cosmological parameter likelihoods from high-resolution CMB power spectrum estimators and showed that indeed good approximations can be found that produce unbiased results [3]. These approximations effectively transform a set of power spectrum estimators so the likelihood of a theoretical power spectrum can be written in a Gaussian form, and then evaluate this Gaussian function using an estimate of the power spectrum estimator covariance. In simple cases this covariance can be calculated to reasonable accuracy using analytic approximations, though ideally it should be assessed by performing large numbers of full data-analysis simulations, fully accounting for noise, sky cuts, noise correlations, beam effect, map-making errors, etc.

Additional complications may additionally be accounted for by modifying the theory power spectrum, for example, due to beam or foreground uncertainty modes.

In this paper we aim to assess how much information is being lost by using a fast pseudo- C_l method compared to a more optimal method (e.g. maximum likelihood or Gibbs sampling [4]), and whether the increase in error bars from using a suboptimal method has a significant effect on cosmological parameters. We then assess how to estimate the covariance from simulations and how many simulations are required under various assumptions. We show how information from approximate models can be combined with simulations by using shrinkage estimators or by fitting model parameters. Finally we suggest how an approximate power spectrum likelihood at high l can be combined with a more exact low- l likelihood for evaluating the total likelihood, consistently accounting for the contribution of high- l power to the low- l mode variance.

We focus on nearly full-sky observations, for example, from the WMAP or Planck satellites, where only a small fraction ($\sim 15\%$) of the sky is cut out due to foreground or point-source contamination. For simplicity we discuss mainly the CMB temperature, which is what drives the parameter constraints, though in the appendix we give a general harmonic method for calculating the low- l likelihood in harmonic space. Since the paper is only likely to be of interest to experts in the field, we refer to the extensive literature for introductory material.

II. HOW OPTIMAL ARE PSEUDO- C_l LIKELIHOODS?

We focus on hybrid pseudo- C_l power spectrum estimators, constructed by combining sets of pseudo- C_l estimators from maps with different weightings. For details of how the pseudo- C_l estimators are constructed and how the

*samira@ast.cam.ac.uk

†URL: <http://cosmologist.info>

covariance can be estimated, see Refs. [3,5–13]. When the sky is noise dominated, the pseudo- C_l estimator with inverse-noise weighting is optimal; when the noise is negligible a uniform weighting is optimal. Combining different pseudo- C_l estimators with different weightings gives hybrid estimators that interpolate smoothly between the regimes, giving significantly smaller error bars at all l than a single weight function could [9].

Over the years various code comparisons and tests have indicated that pseudo- C_l estimator are in fact rather good at high l , in that their variance is typically within $\mathcal{O}(10\%)$ of that expected for a maximum-likelihood estimator (e.g. Refs. [7,12]). However, the degree of suboptimality will depend on the noise and sky cut under consideration. Because of its scanning strategy, the Planck satellite will have highly anisotropic noise, and we would like to assess whether this is likely to be a problem for pseudo- C_l methods. Here we perform a comparison in the case of a strongly anisotropic noise, but with azimuthal symmetry so that the optimal Fisher errors can be calculated exactly for comparison. Though this situation is clearly unrealistic, it should give a good idea of the amount of suboptimality that can be expected in practice. Our main concern is the anisotropy of the noise, so we shall consider the full sky; the main advantage of doing this is that the covariance of the pseudo- C_l estimators can then be calculated accurately analytically, so that our comparison is assessing the information loss due to suboptimal handling of noise anisotropy, rather than depending on errors in the covariance calculation. We discuss the effects of covariance errors in a later section. We focus on the high- l temperature power spectrum since this dominates the parameter constraints from the Planck satellite; for discussion of how to improve polarization pseudo- C_l estimators see Refs. [13,14].

A. Comparing the Fisher matrix and the pseudo- C_l 's variance in the full sky

The Cramér-Rao inequality states that the inverse of the Fisher information is a lower bound on the variance of any unbiased estimator, so we use the Fisher matrix to quantify the errors on the power spectrum that one could hope to get using an optimal method. For a Gaussian likelihood function \mathcal{L} , the Fisher matrix F , for the power spectrum C_l at some fiducial model is given by (see e.g. Ref. [15])

$$F_{ll'} = -\left\langle \left(\frac{\partial^2}{\partial C_l \partial C_{l'}} \right) \ln \mathcal{L} \right\rangle = \frac{1}{2} \text{Tr} \left[\mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial C_l} \mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial C_{l'}} \right], \quad (1)$$

where $\mathbf{C} = \mathbf{S} + \mathbf{N}$, with \mathbf{S} and \mathbf{N} being the signal and noise covariances, respectively. For a statistically isotropic signal on the full sky \mathbf{S} is diagonal in harmonic space, and for pixel-uncorrelated noise \mathbf{N} can be calculated easily in terms of the pixel noise variance. We make the assumption

of azimuthal symmetry so that $N_{l'm'lm} = \delta_{mm'} N_{l'mlm}$, which makes the Fisher matrix numerically tractable by putting it into block diagonal form (with blocks for each m). The Fisher matrix then reads:

$$F_{ll'} = \frac{1}{2} \sum_m ([\mathbf{C}^{-1}]_{ll'}^{(m)})^2, \quad (2)$$

which is simply half the square of the inverse of the covariance matrix (signal plus noise) summed over m .

We compare this with the errors expected using pseudo- C_l power spectrum estimators. For weight functions labeled by i and j there are a set of pseudo- C_l estimators \hat{C}_l^{ij} and a covariance

$$M_{ll'}^{ijkl} \equiv \langle \Delta \hat{C}_l^{ij} \Delta \hat{C}_{l'}^{kl} \rangle. \quad (3)$$

Explicit expressions for the covariances are given in e.g. Ref. [3]. The hybrid pseudo- C_l estimators of Refs. [3,9] are a combination of pseudo- C_l 's with different combinations of weights. In the low noise regime, the best weight function, $\omega(\Omega)$, is close to uniform in order to minimize the cosmic variance. However, in the high noise regime, the weight function is proportional to the inverse noise to reduce noise itself [9]. Therefore, a natural choice would be to choose a combination of results from uniform and inverse-noise weight functions. Here we consider the simplest case of combining the two estimators calculated separately from the two weighted maps. The full hybrid covariance matrix then reads:

$$M_{ll'}^h \equiv \langle \Delta \hat{C}_l^h \Delta \hat{C}_{l'}^h \rangle = \left(\sum_{AB} H_{ll'}^{AB} \right)^{-1}, \quad (4)$$

where $H_{ll'}^{AB}$ is the inverse of the covariance matrix $\langle \Delta \hat{C}_l^A \Delta \hat{C}_{l'}^B \rangle$, with A and B denoting the different weight functions used for the \hat{C}_l contributing to the hybrid estimator.

We can now proceed to compare the errors on individual C_l and their correlations. However, for parameter estimation we are not just interested in individual l : the C_l are smooth functions of l , and the effects of parameters enter in a smooth way over a range of l . We are therefore more interested in the constraint on the amplitude over some smoothing scale, as this will more closely determine how well we can constrain different parameters. The constraint over a range of scales depends on the correlations of different individual l , so considering a range of l also has the advantage of checking any effects due to correlations.

We therefore also evaluate the error of an amplitude parameter, A , over some range in Δl , so that the power spectrum over the range Δl is given by AC_l^n , and elsewhere by C_l^n , for some fiducial model C_l^n . The Fisher matrix for amplitude A over Δl is then

$$F_{AA} = \frac{1}{2} \sum_m \sum_{l=|m|}^{l_{\max}} \sum_{l'=|m|}^{l_{\max}} [C^{-1}]_{ll'}^2 C_l^{\text{in}} C_{l'}^{\text{in}} \begin{cases} 1 & \text{if } l \text{ and } l' \text{ in } (\Delta l) \\ 0 & \text{otherwise} \end{cases}. \quad (5)$$

This can be compared to using the pseudo- C_l estimators in a Gaussian likelihood approximation:

$$-2 \ln \mathcal{L}_f(C_l) = \sum_{l'} (\hat{C}_l - C_l) [\mathbf{M}_f^{-1}]_{ll'} (\hat{C}_{l'} - C_{l'}), \quad (6)$$

where $[\mathbf{M}_f]$ is the covariance for a fiducial model. Differentiating twice with respect to A gives

$$-\frac{\partial^2}{\partial A^2} (\ln \mathcal{L}_f) = \sum_{l'} C_l^{\text{in}} [\mathbf{M}_f^{-1}]_{ll'} C_{l'}^{\text{in}}, \quad (7)$$

where the sum is over the l in the bin being considered. Note that we are considering the case of generating estimators at each l and then using the covariance to estimate the power over a range of l . This is not the same as making a binned C_l estimator (depending on some window function range of l), which may well have different properties. Full-sky observations such as WMAP and Planck are normally analyzed into individual C_l estimators, the case we consider, though there may be advantages to also considering binned estimators.

To make an error comparison, we use an azimuthally averaged (in ecliptic coordinates) version of the noise

expected for the Planck satellite [16] and consider an isotropic 7arcmin-fwhm Gaussian beam.

We compare the inverse of the hybrid variance (the diagonal elements of Eq. (4)) to the Fisher errors (the diagonals of Eq. (2)) in Fig. 1. The figure also shows the uniform and inverse-weighted estimators separately. Although the variance of the uniform noise weighting is almost identical to the fisher error at low l , it deviates considerably when noise dominates. On the other hand, the variance of the inverse-variance noise weighting case remains poor even at high l where noise is starting to dominate over the signal. The hybrid estimators significantly improve the errors, giving results within a few percent of the ideal Fisher errors at all l .

Figure 2 compares the results for the amplitude over a range $\Delta l = 2, 10, 50$. In this case a larger range of l gives relatively worse hybrid errors compared to the Fisher errors, though the hybrid errors are still within 10% at all l . One interpretation of the effect of bin size might be that the optimal result is combining correlation information between different scales significantly more efficiently than the hybrid estimator. This is perhaps not surprising since the Fisher result “knows” about correlations be-

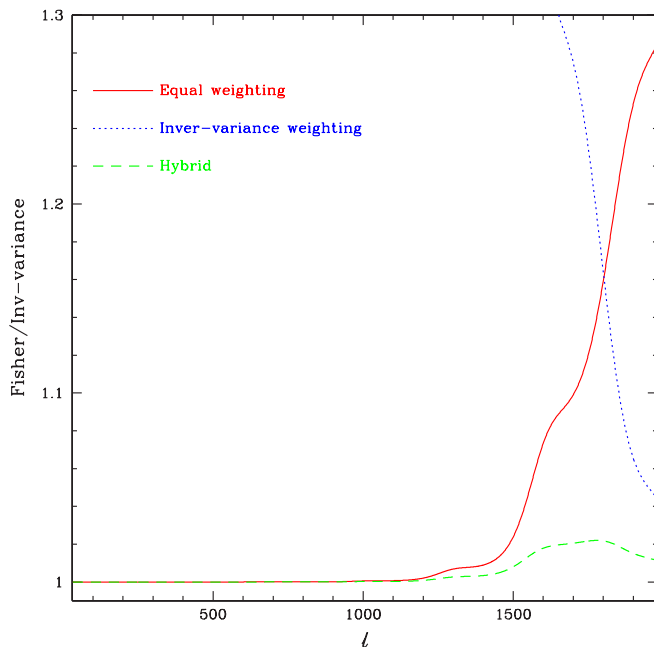


FIG. 1 (color online). A comparison between the Fisher errors and the inverse of the hybrid estimator variance for the different weighting combinations considered. The red (solid), blue (dotted), and green (dashed) curves represent the ratios of Fisher errors to the inverse variance of uniform noise weighting, inverse-variance noise weighting, and hybrid, respectively.




FIG. 2 (color online). A comparison of the errors in the power over a range $\Delta l = 2$ (solid line), $\Delta l = 10$ (dashed line), and $\Delta l = 50$ (dot-dashed line), for bands starting at l_{\min} . The red, blue, and green curves represent the ratios of Fisher errors to the inverse variance of uniform noise weighting, inverse-variance noise weighting, and hybrid, respectively.

tween all the l and m modes individually, whereas the hybrid approximation has compressed the correlation information into a covariance matrix accounting only for l correlations. The hybrid estimator could be further improved by including more pseudo- C_l estimators in the mix if greater optimality is desired, for example, the uniform-inverse weight estimator; however, as we see below the level of optimality found here is likely to be sufficient in most cases.

B. Effect on parameter estimation

Since the hybrid pseudo- C_l may be suboptimal at up to the 10% level, it is also useful to directly check the effect on parameter estimation. We use full-sky Planck-like simulations with azimuthal symmetry as previously described. We consider a standard 6-parameter Λ CDM model (with running of the spectral index, but the optical depth fixed since we exclude $l \leq 30$ where a more optimal likelihood analysis is possible). For simplicity and for a quick check, we consider the fiducial Gaussian likelihood approximation since it depends on a precomputed covariance matrix.

We calculate the likelihood with the Fisher covariance (Eq. (2)) and with the hybrid covariance (Eq. (4)) using the COSMOMC parameter estimation code to sample from the posterior parameter distribution [17]. We use a fiducial analytic model for the \hat{C}_l , corresponding to averaging the log-likelihood over realizations.

Figure 3 shows the results, which are very similar. This is not a surprise since Planck can measure many acoustic peaks in the CMB power spectrum, so most parameters are well constrained from the cosmic-variance limited region where the hybrid estimator is close to optimal. The effect could be somewhat larger in extended parameter spaces where models differ significantly only over the region where the hybrid estimator is significantly suboptimal.

III. HOW TO CALCULATE THE COVARIANCE?

In the simple cases we have considered above the estimator covariance can be calculated accurately analytically. In more realistic cases this is unlikely to be the case: even a simple sky cut renders the commonly used covariance approximates accurate at only the $\mathcal{O}(10\%)$ level, depend-

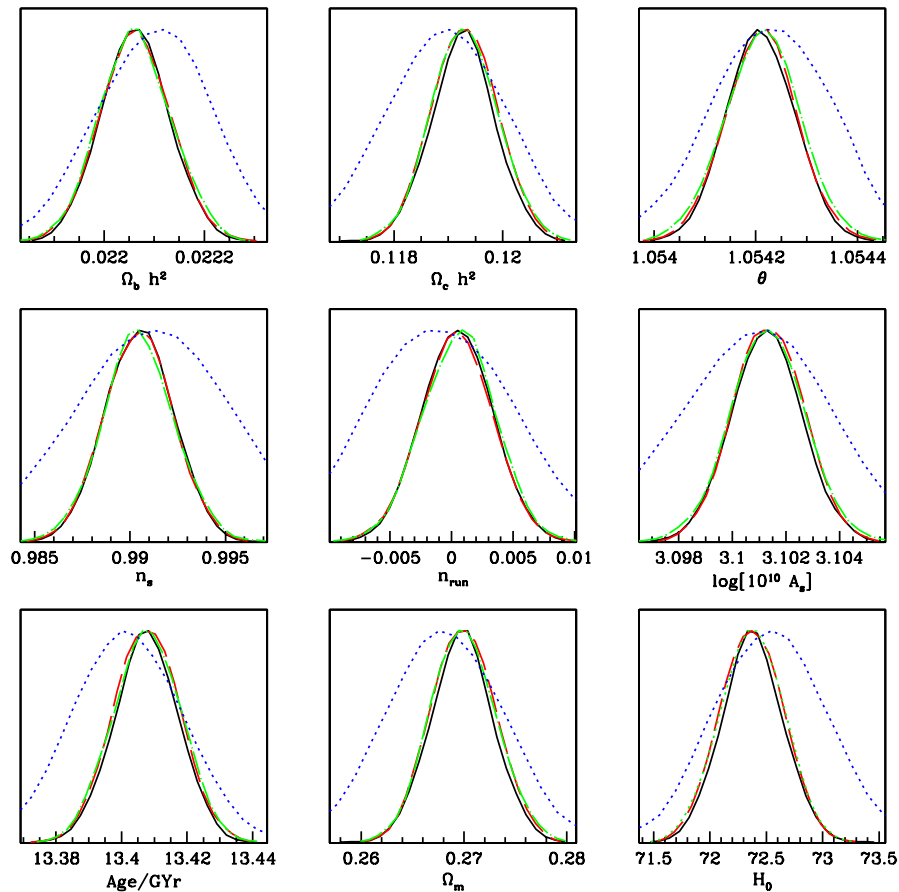


FIG. 3 (color online). Parameter constraints from a single idealized Planck-like simulation with anisotropic noise. The one-dimensional marginalized posteriors are from the fiducial Gaussian approximation with the covariance being from the Fisher matrix or the pseudo- C_l estimator covariance. The black (solid) line uses the Fisher covariance, blue (dotted) line is the inverse-noise weighted estimator, green (dot-dashed) line is with uniform weight, and red (dashed) line uses the hybrid estimator. The results are very consistent as constraints are driven by scales where the uniform weighting is nearly optimal.

ing on apodization (see e.g. [18]). Ideally we would like to be able to evaluate the covariance (in some fiducial model) directly from simulations of the full time stream, including all relevant instrumental effects. This is the approach adopted by the MASTER/xFASTER pipelines used by many CMB observations with data over only a small part of the sky [7,19,20]; the advantage of a Monte Carlo approach is that numerous complicated effects can be included straightforwardly, where an analytic result would be intractable. It should, however, be remembered that if the likelihood approximation is derived to be good for near-isotropic-Gaussian fields, including systematics in the covariance may invalidate the likelihood approximation. For example beam uncertainty modes effectively scale the theory power spectrum coherently over a large range of scales, so beam uncertainties may be better accounted for by including them as extra parameters affecting the theory C_l in the parameter estimation code, rather than as part of the likelihood function for beam-uncertain estimators [18].

The maximum-likelihood estimator of the covariance of a set of zero-mean samples of a Gaussian vector \mathbf{n} is

$$\hat{M} = \frac{1}{n} \sum_i \mathbf{n}_i \mathbf{n}_i^T, \quad (8)$$

where the sum is over the n independent samples. If our only knowledge about the covariance comes from simulations, we should account for the sampling uncertainty in the true covariance by marginalizing over the probability distribution of the true covariance given the Monte Carlo estimate; see Appendix A where we give mathematical details in an idealized case.

If there are a large number of samples, the uncertainty in the true covariance may be negligible, so the Monte Carlo covariance can be used directly, but the question is: how many samples do you need? There is a basic lower limit of $n \geq p$, where p is the dimensionality of the covariance; this is required to have p independent modes sampled, and hence for the covariance estimate to be invertible. Already in the case of the power spectrum estimators this is a nontrivial requirement: for WMAP temperature $p \sim 10^3$, but due to the very high numerical cost they were only able to do a dozen or so full time-stream simulations. For the distribution of the true covariance given the estimator to be normalizable we actually need $n > 2p$, and for fractional accuracy better than α on the error bars, we require $n \geq 2p/\alpha$. Indeed to have the numerical value of the χ^2 accurate, one needs $n \gg p^2$, which is getting to be very time consuming for unbinned C_l estimators, even for simplified map-level simulations. However, in practice it is not usually required to have the χ^2 accurate, as long as the variation under changes in parameters is accurate, and hence the correct parameter constraints are obtained. For further discussion see Appendix A and Refs. [21,22].

In a realistic situation we may have some idea of the covariance from approximate analytical results, but would also like to calibrate it from simulations to account for complications that are not included in the analytical model. An estimator can use both the approximation and the simulations. Perhaps the simplest case to consider is where there is a prior estimate that one considers to be about as accurate as that from q simulations. Then adding the information from n actual simulations, the best estimate of the true covariance will be the weighted sum of the two covariance estimates (see Appendix A). This is a simple example of a shrinkage estimator: it “shrinks” noisy simulated estimators towards some prior target, giving a new estimator that should be better than both individually. In the limit of many simulations the new estimator is simulation dominated; with few simulations it is prior dominated.

A. Shrinkage estimators

Shrinkage estimators were originally introduced for the situation in which there are many fewer samples than numbers of dimensions, so that the maximum-likelihood covariance estimate is not even invertible. This is a common situation in many problems, and shrinkage provides a method to regularize the estimate in a well-defined way so that it is invertible. Here we would ideally also like to have very few simulations, but we also require good accuracy of the answer and may prefer to generate more samples than have an inaccurate answer. In this section we investigate whether shrinkage estimators are useful for CMB likelihoods.

A shrinkage estimator \mathbf{s}^* for a quantity \mathbf{s} is constructed from a linear combination of some “target” \mathbf{t} and an estimator $\hat{\mathbf{s}}$ calculated from samples [23–25]:

$$\mathbf{s}^* = \lambda \mathbf{t} + (\mathbf{1} - \lambda) \hat{\mathbf{s}}, \quad (9)$$

with λ , the shrinkage intensity, being in the range 0 to 1. The target can either be a fixed prior, or it could be some approximate estimate from the data or some combination. If $\lambda = 0$ then the shrinkage estimate equals the unrestricted estimate, $\hat{\mathbf{s}}$. If, however, $\lambda = 1$ then the target estimate, \mathbf{t} , is recovered. Therefore, the main advantage of this weighted combination is to form a regularized estimator that outperforms both estimators individually. The shrinkage intensity is chosen in a way that would optimize the estimator, for example, by minimizing the mean-squared error:

$$\text{MSE}(\lambda) = \left\langle \sum_l (s_l^* - s_l)^2 \right\rangle, \quad (10)$$

with the angle brackets being the expectation value. Ledoit and Wolf [23] derived an analytic solution for λ to minimize this error:

$$\lambda^* = \frac{\sum_l [\text{var}(\hat{s}_l) - \text{cov}(t_l, \hat{s}_l) - \text{bias}(\hat{s}_l)(t_l - \hat{s}_l)]}{\sum_l (t_l - \hat{s}_l)^2}. \quad (11)$$

For a practical use of Eq. (11), Ref. [24] suggests replacing the covariance, variance, and bias with their unbiased samples estimates ($\hat{\text{cov}}$, $\hat{\text{var}}$, $\hat{\text{bias}}$). If \hat{s} is an unbiased estimator, then the bias term may be omitted. The new expression of the shrinkage intensity then reads:

$$\hat{\lambda}^* = \frac{\sum_l [\hat{\text{var}}(\hat{s}_l) - \hat{\text{cov}}(t_l, \hat{s}_l)]}{\sum_l (t_l - \hat{s}_l)^2}. \quad (12)$$

In practice we use $\min(1, \hat{\lambda}^*)$.

The shrinkage method can be applied directly to estimate the covariance matrix when \mathbf{s} is taken to be a vector of the distinct elements of the covariance matrix. If we have n measurements of a data vector \mathbf{x} of size p , and we assume that $x_l^{(k)}$ is the k th measurement of the l th element of \mathbf{x} , then the empirical mean of the variable x_l is

$$\bar{x}_l = \frac{1}{n} \sum_{k=1}^n x_l^{(k)}. \quad (13)$$

The unbiased empirical covariance, $\hat{\mathbf{S}}$, is given by

$$\hat{S}_{ll'} = \hat{\text{cov}}(x_l, x_{l'}) = \frac{n}{n-1} \bar{W}_{ll'}, \quad (14)$$

where we have set

$$\bar{W}_{ll'} = \frac{1}{n} \sum_{k=1}^n W_{ll'}^{(k)}, \quad (15)$$

$$W_{ll'}^{(k)} = (x_l^{(k)} - \bar{x}_l)(x_{l'}^{(k)} - \bar{x}_{l'}). \quad (16)$$

Similarly, the required variance can be estimated as

$$\hat{\text{var}}(S_{ll'}) = \frac{n}{(n-1)^3} \sum_{k=1}^n (W_{ll'}^{(k)} - \bar{W}_{ll'})^2. \quad (17)$$

If the target is fixed (does not depend on the data) then the covariance term vanishes; if not, it should be included to account for the correlation. We estimate the shrinkage intensity by replacing $\hat{\mathbf{s}}$ and \mathbf{t} in Eq. (11) with $\hat{\mathbf{S}}$ and \mathbf{T} , where \mathbf{T} is some arbitrary target matrix. Therefore, the optimized shrinkage covariance now takes the form

$$\mathbf{M} = \hat{\lambda}^* \mathbf{T} + (1 - \hat{\lambda}^*) \hat{\mathbf{S}}. \quad (18)$$

The shrinkage estimator is designed to use information from the sample estimator roughly in proportion to how relatively accurate it is. If we have a target that is say 10% accurate, $\hat{\mathbf{S}}$ needs to have $\leq 10\%$ accuracy to significantly improve on the target. For high precision results, the estimator needs to be accurate, so the shrinkage estimator

requires the same order of number of samples as the direct accurate estimator. It simply has the advantage of being slightly superior and not breaking down rapidly as the number of samples is decreased.

Note that it is not at all clear in the current context that minimizing the Frobenius norm of the error (Eq. (10)) is the best thing to do. Indeed if the diagonal and nondiagonal elements of the covariance have very different distributions, it is not a good idea to use a shrinkage estimator derived by treating them on equal terms. One can, however, equally well apply a shrinkage estimator to subsets of parameters separately. Since the off-diagonal terms are generally small, the number of simulations required to estimate them accurately is very large: for a given number of samples the shrinkage intensity λ^* will be significantly larger for the off-diagonal than the diagonal components. Also note that the shrinkage intensity λ^* is not independent of the scaling of the elements (e.g. using C_l or $l(l+1)C_l$ gives different results). We therefore apply the estimator to the covariance matrix after taking out the approximate $C_l C_{l'}$ scaling of the elements (e.g. normalizing so that the target has unit diagonal).

The effect of the shrinkage estimator on estimated likelihoods is illustrated in Fig. 4, where separate shrinkage estimates are used for the diagonal and off-diagonal parts




FIG. 4 (color online). Example log likelihood for the power spectrum amplitude over $100 < l \leq 700$ in a typical realization when using the exact covariance (thick solid line), a prior guess diagonal target that is off by a factor of 1.3 from the true covariance (dot-dashed line), the covariance estimated from 2000 samples (dashed line), and the shrinkage estimator that combines the estimated and target covariances to get closer to the correct result (thin solid line). The diagonal and off-diagonal parts of the matrix are shrunk separately, giving a shrinkage estimator close to the sample estimate on the diagonal and close to the target on the off diagonal. We approximate the likelihood as Gaussian in the C_l .

of the covariance matrix. Here toy samples were generated from the Gaussian model with a known true covariance, and the shrinkage target is diagonal and 30% wrong. The shrinkage estimator is significantly closer to the true result than using either the direct sample estimate or the diagonal target, as expected. Using a purely diagonal covariance approximation with the variances estimated from the samples would give a similar result. As more samples are used the shrinkage estimator gradually includes more information about the off-diagonal correlations from the simulations.

The division of the covariance into components with different properties can be further generalized. Cut-sky effects tend mostly to correlate nearby l , but the matrix is smooth in l , so, for example, $M_{l(l+1)}$ is similar to $M_{(l+1)(l+2)}$ (see Fig. 5). This suggests slicing the matrix into diagonal strips: M_{ll} (for all l), $M_{l(l+1)}$, $M_{l(l+2)}$, etc., where each strip is given its own shrinkage intensity. In practice $M_{l(l+n)}$ is very small for $n \gg 1$ in most cases (the matrix is nearly band diagonal), so that the combined shrinkage estimator effectively sets $n \gg 1$ elements to the target. Using a




FIG. 5 (color online). Diagonal strips of the maximum-likelihood correlation matrix estimated from 23 000 simulations of the WMAP 5-year temperature maps [18], $(l, l + 1)$ [thin solid line], $(l, l + 2)$ [thick solid line], $(l, l + 3)$ [thin dashed line], and $(l, l + 4)$ [thick dashed line]; the elements become very small when further away from the diagonal. The top figure is the raw estimate, the bottom figure shows the result smoothed with a Gaussian kernel of width $\sigma_l = 10$. The $(l, l + 1)$ and $(l, l + 3)$ elements are much smaller than $(l, l + 2)$, $(l, l + 4)$ due to the near north-south symmetry of the foreground mask in galactic coordinates. The thin black solid lines in the bottom half show an analytic approximation for the corresponding parts of the correlation matrix [18].

shrinkage estimator does, however, allow the possibility of the estimator including strongly off-diagonal correlations if they appear to be needed by the samples. If the target is nearly zero for $n \gg 1$, then all the $n \gg 1$ elements could be lumped together with a single shrinkage intensity, reducing noise in the intensity due to the small size of the strips in the corners of the matrix. The strips with $n \sim 1$ will gradually change from the target to the sample estimator as the number of samples is increased.

Figure 5 compares the covariance estimated from 23 000 realistic simulations of the WMAP 5-year temperature maps [18] with an analytic approximation for off-diagonal parts of the covariance. The analytic result agrees rather well, indicating that in this case an extremely large number of simulations would be required to improve the off-diagonal result significantly from simulations. On the other hand if there was a disagreement between the simulations and analytic result, a shrinkage estimator would account for this and start to correct the result. On the much larger diagonal analytic approximations are only accurate to $\mathcal{O}(5\%)$ (see Fig. 7), so a relatively modest number of simulations can improve on the analytic result.

Figure 6 shows how the shrinkage intensity varies with the number of simulations, both for the case of a significantly wrong target, and the example of WMAP where we have an analytic approximation accurate at the 5% level. The better the target is the more simulations are required to improve on it from the simulations: with a significantly wrong target the simulation estimator rapidly dominates the shrinkage estimator ($\hat{\lambda}^* \rightarrow 0$). With a good analytic result the diagonal estimate is gradually improved over a few thousand simulations, but far more simulations would be needed to correct small inaccuracies in the off-diagonal terms.

A disadvantage of the shrinkage approach is that it does not specify how parameters should be scaled or divided into subsets with different shrinkage intensities. For example, it may also be beneficial to divide in ranges of l if the accuracy of the target varies as a function of l . This is the case for many power spectrum estimators since the covariance can be calculated accurately in the noise-dominated regime, with the main inaccuracies coming from the sample variance over the acoustic peaks. The extent to which covariance errors affect parameters also varies as a function of l : errors in the noise-dominated regime have little effect because there is no signal there.

B. Covariance models

The general shrinkage estimator uses a target covariance and combines it directly with data from simulations, allowing general deviations from the target to be determined with enough simulations. However, in many instances we actually have rather strong priors about the form of the covariance, even without an accurate prior; for example, in the case of the CMB the covariance is expected to be




FIG. 6 (color online). The shrinkage intensity $\hat{\lambda}^*$ against the number of WMAP 5-year simulation samples used, using the analytic approximation as the target (left) and using a target that is significantly wrong (right; the analytic approximation multiplied by 1.3). The shrinkage intensity is calculated separately for each strip of the normalized matrix, the plots show the diagonal of the covariance (solid line), first off-diagonal strip (short-dashed line), and the second strip (long-dashed line).




FIG. 7 (color online). The fractional difference of various approximations to the smoothed diagonal of the WMAP5 C_l covariance estimated from 23 000 samples from Ref. [18]. The thick solid line shows an analytic approximation, which is accurate at the 5% level. Thin lines show the result from fitting a cubic spline to the C_l variance using 100 (dot-dashed line) and 1000 (solid line) simulation samples. Spline nodes are separated by $\Delta_l = 50$.

strongly diagonally dominated and also smooth in l unless sharp changes are introduced by the choice of C_l -estimator¹ (e.g. see Fig. 5). If we are confident that this structure is correct, we should be able to use it to improve our covariance estimator.

The approach adopted by the WMAP team is to use an analytic model of the covariance and then calibrate it from simulations by fitting parameters to a model of the covariance diagonal [2]. If the model has the right shape this can reduce to fitting a very small number of parameters, which can be done accurately from a relatively small number of simulations, making this a very good approach. On the other hand if the actual behavior is more complicated than the assumed model, or the off-diagonal correlations cannot be calculated reliably analytically, this could lead to misestimation of the covariance.

For analysis methods that do not introduce features in l , variations in the C_l covariance diagonal are smooth on a scale determined by the scale of the acoustic peaks. Rescaling the diagonal of the covariance using a smooth model and taking the correlation matrix from the analytic result therefore gives close to the right answer; e.g. using the estimator

¹Note the WMAP team’s analysis does use a sharp cut in l to switch between pseudo- C_l weighting schemes. However, each estimator is separately smooth.

$$M_{ll'}^{\text{model}} = p_l T_{ll'} p_{l'} \quad (19)$$

(no sum), where \mathbf{T} is the analytic model and p_l are fit to the simulation data using cubic splines. Figure 7 shows that fitting only 100 simulation samples using cubic splines for the diagonal can be as accurate as the analytic approximation. With a thousand or more simulations the covariance diagonal can be determined to the subpercent level. If required additional parameters could be introduced to model variations in the off-diagonal components. If the theoretical model dependence of the covariance is not well captured by the approximation of Ref. [3], fits could also be made as a function of simulation parameters and then interpolated as required.

An alternative to fitting a model for the covariance when no good model is available would be using a smoothed version of the estimated covariance, for example, by applying a Gaussian smoothing kernel to each diagonal strip of the matrix. If the true covariance is indeed smooth, this will significantly reduce the sampling noise and hence improve the covariance estimate. For example, applying a smoothing kernel of width $\sigma_l = 5$ to the diagonal strips produces results for the likelihood in Fig. 4 that are similar to the shrinkage estimator, depending on the realization. However, it does not guarantee that the smoothed matrix is invertible.

IV. COMBINING LOW AND HIGH l

We have discussed how the likelihood can easily be calculated at high l without losing very much information by compressing the data into a set of pseudo- C_l estimators. However, at low l one should be able to perform a more optimal analysis, which is desirable since the exact shape of the likelihood function depends on the particular realization of low l modes on the sky. The WMAP data analysis uses a pixel-based likelihood at low l [26], essentially assuming Gaussianity of a set of low-resolution pixel values and using the pixel covariance to calculate the likelihood. An alternative would be to work directly in harmonic space, avoiding pixelization issues, as discussed in Appendix B. However, one calculates the low l likelihood, the question remains of how to combine this with a high- l likelihood, given it will not be possible to do an exact l separation on the cut sky.

Given a set of observed modes on the sky $\tilde{\mathbf{X}}$, a low l likelihood operates on some subset of the modes given by $\mathbf{X}_s = \tilde{\mathbf{M}} \tilde{\mathbf{X}}$, where $\tilde{\mathbf{M}}$ is rectangular and \mathbf{X}_s is sufficiently small that an exact likelihood calculation is numerically feasible. In general \mathbf{X}_s will contain the low l modes we are aiming to analyze exactly ($l \leq l_{\text{exact}}$), along with some leakage from higher $l > l_{\text{exact}}$ due to the sky cut. It follows that in the fiducial model the remaining information is contained in the independent high- l -dominated modes

$$\begin{aligned} \tilde{\mathbf{X}}_{>} &\equiv (\mathbf{I} - \langle \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger \rangle \tilde{\mathbf{M}}^\dagger [\tilde{\mathbf{M}} \langle \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger \rangle \tilde{\mathbf{M}}^\dagger]^{-1} \tilde{\mathbf{M}}) \tilde{\mathbf{X}} \\ &= \tilde{\mathbf{X}} - \langle \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger \rangle \tilde{\mathbf{M}}^\dagger [\langle \mathbf{X}_s \mathbf{X}_s^\dagger \rangle]^{-1} \mathbf{X}_s \\ &\approx \tilde{\mathbf{X}} - \langle \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger \rangle \tilde{\mathbf{M}}^\dagger \mathbf{X}_s, \end{aligned} \quad (20)$$

where the third line follows for the temperature in the true model if the modes have been decorrelated and normalized to be white. For temperature and polarization the same construction can be used on the large vector \mathbf{X}_{TEB} in Eq. (20), though there are simplifying advantages to doing temperature and polarization separately. In a different model these modes should also be approximately independent as long as the fiducial model is reasonably accurate.

From the high- l modes $\tilde{\mathbf{X}}_{>}$, one can construct pseudo- C_l (or more optimal) estimators as normal, however the coupling matrix and covariance are more complicated because




FIG. 8 (color online). Input temperature map with kp2 mask [35] (top), low- l modes $\langle \tilde{\mathbf{X}} \tilde{\mathbf{X}}^\dagger \rangle \tilde{\mathbf{M}}^\dagger \mathbf{X}_s$ with $l_{\text{exact}} = 20$, $l_{\text{low}} = 80$ (middle, see Appendix B), and the low- l modes' contribution from $l > l_{\text{exact}}$ (bottom). Note how the supported modes are going smoothly to zero at the cut edges, and the mixing from $l > l_{\text{exact}}$ is mostly near the edges of the cut. Color scales are not the same.




FIG. 9 (color online). Decomposition of the same map into low- l modes (top), independent high- l modes (middle), and the contribution from $l > l_{\text{exact}}$ to the low- l map. Here $l_{\text{exact}} = 30$, $l_{\text{low}} = 120$. Note a slightly larger mask would remove much of the high- l contribution.

the weight function is no longer local in pixel space. If the leakage from low to high l is small, these should have roughly the same properties as the normal cut-sky estimators, and hence C_l -estimator likelihood approximations can be used. If the leakage is small we could also just ignore the correlation, which is probably fine as long as nothing interesting is happening over the affected l range.

However, as shown in Figs. 8 and 9, the mixing of modes to higher l dominates around the cut edges, where higher- l modes are needed so that the projected map goes smoothly to zero on the cut. This suggests that by slightly enlarging or apodizing the mask used for the $l \geq l_{\text{exact}}$ pseudo- C_l analysis, most of the modes already included in the low- l likelihood would be removed. This enlarged cut would only be needed for some number of $l \geq l_{\text{exact}}$ where the leakage is significant; at very high l the leakage is negligible so the original cut can be used. This is a quick and simple solution to separating the scales to a reasonable




FIG. 10 (color online). Pseudo- C_l temperature power spectra of a simulated map projected into low- l modes ($l_{\text{exact}} = 30$, $l_{\text{low}} = 120$, $\epsilon = 0.001$, see the appendix). The dashed line shows the result using the “kp2” mask (85% of the sky) used to construct the modes. The thin red line shows the estimated C_l s using the enlarged “QK75” temperature mask [36] that includes only 72% of the sky. The leakage to $l > l_{\text{exact}}$ is relatively small but is further decreased by using the larger mask. The thick line shows the power spectrum from the cut sky without projecting out the low- l modes.

approximation: use an accurate likelihood at low l (including some modes from higher l), use pseudo- C_l estimators with an enlarged mask over intermediate scales where the leakage is important, then at high l use pseudo- C_l estimators with the original mask.

Figure 10 shows that the leakage between scales is small but can further be reduced by using a larger mask over a range $\Delta l \sim 10$ above l_{exact} . Using an enlarged mask is suboptimal because some modes in the enlarged cut are being lost: the high- l modes included in the low- l likelihood are largely aligned with the boundary, so by cutting all modes in the boundary region transverse modes are lost that could be included. However, a small suboptimality is not likely to matter if there is little interesting information in the range $l \geq l_{\text{exact}}$. The separation method has the advantage of keeping the low and high- l analyses straightforward while including almost all of the valuable information at low l . The harmonic method described in the appendix is also free from pixelization errors and allows a high-resolution mask to be used: there is no additional loss of information from having to use an enlarged mask and big pixels, as is often the case when using smoothed pixel-based likelihoods. Simply ignoring the correlation between low and high- l would be a good approximation in most models, but is formally incorrect due to double-counting of information in the overlap region.

A recent paper [27] has claimed that for accurate likelihood calculations on our observed CMB sky, a nearly

exact likelihood is required up to $l \sim 100$. Although this is somewhat surprising given that pseudo- C_l methods are unbiased and not far from optimal, it would still be tractable by direct as well as Gibbs-sampling approaches.

V. CONCLUSIONS

We have seen that approximate fast power spectrum likelihoods work well on small scales even with strongly anisotropic noise. The error bars are slightly increased by using a suboptimal method, but since the effect on parameters is small, this is likely to be a price well worth paying for robustness.

The covariance that is required to calculate the high- l likelihood can be calculated from simulations, but in general a very large number of simulations are required to calculate the likelihood accurately. Using approximate analytical results and expected smoothness properties, a more accurate covariance can be calculated with far fewer samples by calibration from simulations. There is a trade-off between being able to see any strange behavior in the simulations and using prior expectations to reduce the number of simulations required. Shrinkage estimators provide one way of combining analytic and simulation results so that if the analytic result is significantly wrong the estimator becomes dominated by the simulation result. Shrinkage estimators, however, require far more samples than a model-fitting approach, which can obtain very accurate results from less than a thousand simulations.

With an accurate covariance and set of power spectrum estimators, the high- l likelihood can be calculated quickly and easily. We showed how this could be combined with a more optimal low- l likelihood function, compensating for the small leakage between high and low l while including almost all of the useful information from low l . We conclude that a combined likelihood function should be a good option for analyzing foreground-cleaned Planck data, though further work is required to model the foreground uncertainties accurately. Fast and robust methods based on power spectra at high l and harmonic near-exact likelihoods at low l are likely to be a very good alternative to globally more accurate (but numerically more difficult) Gibbs-sampling methods, as concluded by other authors [9,11,28]. In the appendix we described a practical harmonic method for calculating the low- l likelihood without introducing pixelization issues, which can be used easily once the noise covariance is calculated in harmonic space. Allowing for marginalization over foreground templates, or similar method, can often take place in any basis, so generalizing the method suggested here to more realistic cases could be straightforward (see e.g. Ref. [29]).

ACKNOWLEDGMENTS

S.H. gratefully acknowledges the support of the Algerian Ministry of Higher Education and Scientific

Research (MESRS) and the British Federation for Women Graduates (BFFWG). A.L. acknowledges support from PPARC/STFC and thanks Anthony Challinor, Steven Gratton, and George Efstathiou for discussion. Some of the results in this paper have been derived using the HEALPix [37] package. We acknowledge the use of the Legacy Archive for Microwave Background Data Analysis (LAMBDA).² Support for LAMBDA is provided by the NASA Office of Space Science.

APPENDIX A: COVARIANCE PRIORS

The maximum-likelihood estimator of the covariance of a set of zero-mean samples of a Gaussian vector \mathbf{n} is

$$\hat{\mathbf{C}} = \frac{1}{n} \sum_i \mathbf{n}_i \mathbf{n}_i^T. \quad (\text{A1})$$

A p -dimensional symmetric positive definite matrix has the Wishart distribution (see e.g. Ref. [21]) $\mathbf{S} \sim W_p(n, \mathbf{C})$ if

$$P(\mathbf{S}) = \frac{|\mathbf{S}|^{(n-p-1)/2}}{2^{(1/2)np} \Gamma_p(n/2) |\mathbf{C}|^{n/2}} e^{-(1/2) \text{Tr}(\mathbf{S}\mathbf{C}^{-1})}, \quad (\text{A2})$$

so the covariance estimator is distributed as $n\hat{\mathbf{C}} \sim W_p(n, \mathbf{C})$. Given the estimator, the true covariance with a flat prior has a distribution $\mathbf{C} \sim IW_p(n, n\hat{\mathbf{C}})$ given by

$$P(\mathbf{C}) = \frac{|\mathbf{S}|^{(n-p-1)/2}}{2^{(1/2)(n-p-1)p} \Gamma_p(\frac{1}{2}(n-p-1)) |\mathbf{C}|^{n/2}} e^{-(1/2) \text{Tr}(\mathbf{S}\mathbf{C}^{-1})}, \quad (\text{A3})$$

where $n > 2p$.

Assume we wish to calculate the likelihood of \mathbf{d} , where \mathbf{d} has a p -dimensional Normal distribution $\mathbf{d} \sim N_p(0, \mathbf{C})$. Using the estimator of the covariance we want to marginalize out the uncertainty in the true covariance to give

$$\begin{aligned} L(\mathbf{d}|\hat{\mathbf{C}}) &= \int d\mathbf{C} \frac{e^{-(1/2)\mathbf{d}^T\mathbf{C}^{-1}\mathbf{d}}}{(2\pi)^{p/2} |\mathbf{C}|^{1/2}} \\ &\times \frac{|n\hat{\mathbf{C}}|^{(n-p-1)/2}}{2^{(1/2)(n-p-1)p} \Gamma_p(\frac{1}{2}(n-p-1)) |\mathbf{C}|^{n/2}} \\ &\times e^{-(n/2) \text{Tr}(\hat{\mathbf{C}}\mathbf{C}^{-1})}. \end{aligned} \quad (\text{A4})$$

The integral is just another inverse-Wishart distribution and can be done, so that

$$-2 \log L = (n-p) \log |n\hat{\mathbf{C}} + \mathbf{d}\mathbf{d}^T| + \dots \quad (\text{A5})$$

$$= (n-p) \log \left(1 + \frac{1}{n} \mathbf{d}^T \hat{\mathbf{C}}^{-1} \mathbf{d} \right) + \dots \quad (\text{A6})$$

²<http://lambda.gsfc.nasa.gov/>

$$\approx \frac{n-p}{n} \left[\hat{\chi}^2 - \frac{1}{n} \frac{1}{2} (\hat{\chi}^2)^2 + \dots \right] + \dots \quad (\text{A7})$$

(ignoring \mathbf{d} -independent terms) and assuming $\hat{\mathbf{C}}^{-1}$ is invertible ($n \geq p$). Here $\hat{\chi}^2$ ($\sim \mathcal{O}(p)$ for typical data) is the naive chi-squared we would calculate taking the covariance estimated from simulations to be the true covariance. We are most interested in how the likelihood varies with parameters, so taking the derivative with respect to a parameter θ gives

$$\frac{-2\partial \log L}{\partial \theta} \approx \frac{n-p}{n+\hat{\chi}^2} \frac{\partial \hat{\chi}^2}{\partial \theta}. \quad (\text{A8})$$

Thus using the likelihood $e^{-\hat{\chi}^2/2}$ gives best-fit values that agree with the marginalized result, but we need $n \gg p$ samples for the error bars to be similar. If this is not the case then using $\hat{\chi}^2$ will underestimate the error bars by a fraction $\mathcal{O}([n-p]/[n+p])$; for a fractional accuracy on the error bars better than α , we need $n \gg 2p/\alpha$. For the numerical value of the log likelihood to be close we would also need $n \gg p^2$.

One approach to correct for the underestimated error bars would be to multiply the log likelihood by an estimate of the underestimation factor. This should at least ensure the error bars are consistent, though they would then be larger than could be obtained by using more simulations to estimate the covariance.

If instead of taking a flat prior for \mathbf{C} we assume an inverse-Wishart prior $IW_p(n_r, \mathbf{C}_r)$, the result above is the same with

$$\hat{\mathbf{C}} \rightarrow \frac{n\hat{\mathbf{C}} + n_r\mathbf{C}_r}{n + n_r}, \quad (\text{A9})$$

and $n \rightarrow n + n_r$, which is essentially a shrinkage estimator of the covariance. A hierarchical Bayesian model might take a flat or log prior on n_r , or often it is fixed at its minimum value for normalizability.

APPENDIX B: CUT-SKY HARMONIC LOW- l LIKELIHOOD

Here we present a harmonic near-exact low- l likelihood method that has no artefacts from using a low-resolution pixelization and does not rely on a map-smoothing procedure that potentially mixes information inside and outside the foreground mask. The method must work with strongly anisotropic cusped noise as expected with Planck—but this is no problem in harmonic space. The main disadvantage is that it may be slower than pixel-based codes. For this discussion we neglect all nonideal complications (foregrounds, etc.), and assume that if the noise is correlated, the noise covariance can be calculated in harmonic space. The method can be applied to harmonic coefficients computed from high-resolution pixelized maps or directly to the output of a harmonic map-making algorithm (e.g.

Ref. [30]), thereby avoiding pixelization issues at all stages.

Cut-sky pseudo harmonics $\tilde{\mathbf{X}}$ are given in terms of the underlying full-sky harmonics \mathbf{X} and a noise vector $\tilde{\mathbf{n}}$ by

$$\tilde{\mathbf{X}} = \mathbf{W}^\infty \mathbf{X} + \tilde{\mathbf{n}}, \quad (\text{B1})$$

where \mathbf{W}^∞ is the coupling matrix, which in general is an $n \times \infty$ matrix where we consider $\tilde{\mathbf{X}}$ only up to the first n modes (ordered in l up to l_{low}). Elements of the coupling matrix can be evaluated easily numerically up to moderate l in terms of the harmonic coefficients of the window (foreground and point-source mask) W_{lm} using

$$W_{(l_1 m_1)(l_2 m_2)} = (-1)^{m_1} \sum_l W_{lm} \sqrt{\frac{(2l_1+1)(2l_2+1)(2l+1)}{4\pi}} \times \begin{pmatrix} l & l_1 & l_2 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} l & l_1 & l_2 \\ m & -m_1 & m_2 \end{pmatrix}, \quad (\text{B2})$$

where $m = m_1 - m_2$.

We wish to model the covariance using only modes up to $l = l_{\text{low}}$ and aim to compute the exact likelihood for $l \leq l_{\text{exact}}$ (where $l_{\text{exact}} \leq l_{\text{low}}$). Since \mathbf{W}^∞ is an $n \times \infty$ matrix, the covariance of $\tilde{\mathbf{X}}$ depends on all l , so we need to project out dependence on $l > l_{\text{low}}$. We define \mathbf{W} to be the $n \times n$ Hermitian matrix that we can easily compute, so that $\mathbf{W}^\infty = (\mathbf{W}, \mathbf{Y})$ where \mathbf{Y} is $n \times \infty$. For window functions that are 0 or 1 everywhere \mathbf{W}^∞ is a diagonal projection matrix in pixel space, and in harmonic space completeness of the spherical harmonics then implies $\mathbf{W}^\infty \mathbf{W}^{\infty \dagger} = \mathbf{W}$ (hence $\mathbf{Y} \mathbf{Y}^\dagger = \mathbf{W} - \mathbf{W}^2$). Eigenmodes \mathbf{e} of \mathbf{W} , with $\mathbf{W} \mathbf{e} = \lambda \mathbf{e}$, then have eigenvalues $0 \leq \lambda \leq 1$. Modes with $\lambda \sim 0$ have no signal variance and hence can be deleted without loss of information. The remaining modes that we want have $\mathbf{e}^\dagger \mathbf{Y} \sim 0$, so that dependence on high l is removed. Since

$$|\mathbf{e}^\dagger \mathbf{Y}|^2 = \mathbf{e}^\dagger (\mathbf{W} - \mathbf{W}^2) \mathbf{e} = \lambda(1 - \lambda), \quad (\text{B3})$$

modes with $\lambda \sim 1$ will have the required properties: we just want the well-supported modes (cf. Ref. [31]). Using sufficiently large $l_{\text{low}} > l_{\text{exact}}$ allows us to construct well-supported modes that contain almost all of the information at $l \leq l_{\text{exact}}$ (but only incomplete information at $l_{\text{exact}} < l \leq l_{\text{low}}$). As $l_{\text{low}} \rightarrow \infty$ we have $\mathbf{W}^2 \rightarrow \mathbf{W}$, and hence $\lambda \rightarrow \{0, 1\}$, so for large enough l_{low} the supported modes are guaranteed to contain all of the information. Whether or not it works in practice depends on how high l_{low} needs to be to obtain reliable nearly optimal results at $l \leq l_{\text{exact}}$.

We therefore define the cut-sky supported modes

$$\mathbf{X}_c = \hat{\mathbf{D}}^{-1/2} \hat{\mathbf{U}}^\dagger \tilde{\mathbf{X}} = \hat{\mathbf{D}}^{1/2} \hat{\mathbf{U}}^\dagger \mathbf{X} + \mathbf{n}_c, \quad (\text{B4})$$

where $\mathbf{W} = \mathbf{U} \mathbf{D} \mathbf{U}^\dagger$, \mathbf{U} is orthogonal, and rows of the diagonal matrix \mathbf{D} corresponding to not well-supported modes ($D_{ii} < 1 - \epsilon_1$) are deleted to form $\hat{\mathbf{D}}$ and $\hat{\mathbf{U}}$; ϵ_1 is a free parameter that determines the tolerance for unmod-

eled mixing of modes with $l > l_{\text{low}}$ and \mathbf{n}_c is noise. Note that the uncorrelated noise covariance can be calculated including noise power from all l , so ϵ_1 is only determining the signal leakage. For isotropic white noise $\langle \mathbf{X}_c \mathbf{X}_c^\dagger \rangle_N = \sigma^2 \mathbf{I}$, and for low noise levels it may be a good idea to add a small amount of fake isotropic noise σ_ϵ^2 so that at $l \geq l_{\text{low}}$ the noise is not negligible and theory leakage becomes small compared to the noise on these scales. Also note that we only need to compute the eigenvectors corresponding to small ϵ_1 . In practice ϵ_1 need not be terribly small since for l_{low} significantly bigger than l_{exact} the power from $l \geq l_{\text{exact}}$ is suppressed when we focus on modes with interesting signal from $l \leq l_{\text{exact}}$.

Neglecting the small coupling from $l > l_{\text{low}}$, the full covariance of the well-supported modes can be written (and then Cholesky decomposed) as

$$\langle \mathbf{X}_c \mathbf{X}_c^\dagger \rangle = \mathbf{S}_+ + \mathbf{N}_s + \mathbf{S}_{01} + \mathbf{N}_c + \sigma_\epsilon^2 \mathbf{I} = \mathbf{L} \mathbf{L}^\dagger. \quad (\text{B5})$$

Assuming pixel noise is uncorrelated, for the temperature the noise variance is given by $N_{lm, l'm'} = \sum_s \Omega_s^2 w(s)^2 \sigma^2(s) Y_{lm}(s) Y_{l'm'}(s)^*$ and $\mathbf{N}_c = \hat{\mathbf{D}}^{-1/2} \hat{\mathbf{U}}^\dagger \mathbf{N} \hat{\mathbf{U}} \hat{\mathbf{D}}^{-1/2}$ (the generalization to correlated noise is trivial if \mathbf{N} can be calculated). We split \mathbf{S} into parts, the bit we want from $2 \leq l \leq l_{\text{exact}}$ (\mathbf{S}_+), contributions from other $l_{\text{exact}} < l \leq l_{\text{low}}$ (\mathbf{N}_s) that can be thought of as part of the noise on the low l modes, and \mathbf{S}_{01} which is the contribution from any X monopole and dipole (set to be a large number so that the next step projects it out). Then

$$\mathbf{X}_L = \mathbf{L}^{-1} \mathbf{X}_c \quad (\text{B6})$$

are uncorrelated with unit variance (in the assumed model), and signal covariance from the target l range is

$$\langle \mathbf{X}_L \mathbf{X}_L^\dagger \rangle_s = \mathbf{L}^{-1} \mathbf{S}_+ (\mathbf{L}^\dagger)^{-1} = \mathbf{U}_s \mathbf{D}_s \mathbf{U}_s^\dagger. \quad (\text{B7})$$

Since the \mathbf{X}_L have variance one in total, modes corresponding to small \mathbf{D}_s will be ‘‘noise’’ dominated. We therefore define the signal to signal plus noise eigenmodes by keeping only modes with $[\mathbf{D}_s]_{ii} > \epsilon$ for some small ϵ (for other discussions of similar procedures see Ref. [32] and the WMAP likelihood code³). These new modes contain the interesting signal

$$\mathbf{X}_s = \hat{\mathbf{U}}_s^\dagger \mathbf{X}_L = \hat{\mathbf{U}}_s^\dagger \mathbf{L}^{-1} \mathbf{X}_c. \quad (\text{B8})$$

There are typically $\sim (l_{\text{exact}} + 1)^2 - l_{\text{min}}^2$ of these per field, corresponding to the number of X_{lm} for $l_{\text{min}} \leq l \leq l_{\text{exact}}$. However, a significant fraction of the modes have important power from $l \geq l_{\text{exact}}$ due to the sky cut. See Fig. 11 for an example of how the leakage between scales depends on the parameters.

Define \mathbf{M} , $\tilde{\mathbf{M}}$ as the coupling matrices to the underlying true and pseudoharmonics (at $l \leq l_{\text{low}}$), so that

³http://lambda.gsfc.nasa.gov/product/map/dr2/likelihood_faster_v2p2/wmap_fasttt.ps




FIG. 11 (color online). Pseudo- C_l temperature power spectra of a simulated map (‘‘kp2’’ mask) with the low- l modes ($l_{\text{exact}} = 30$) projected out using a signal/noise cut $\epsilon = 0.5$ (blue, marked points) and $\epsilon = 0.1$ (red, thin line) for $l_{\text{low}} = 60$. The thick line is the power spectrum without projection. This shows the trade-off between including all the low- l power and increasing leakage to higher l . The leakage between scales is smaller with higher l_{low} so that more supported modes containing low- l signal can be extracted.

$$\mathbf{X}_s = \mathbf{M} \mathbf{X} + \mathbf{n}_s \equiv \hat{\mathbf{U}}_s^\dagger \mathbf{L}^{-1} \hat{\mathbf{D}}^{1/2} \hat{\mathbf{U}}^\dagger \mathbf{X} + \mathbf{n}_s, \quad (\text{B9})$$

$$\mathbf{X}_s = \tilde{\mathbf{M}} \tilde{\mathbf{X}} \equiv \hat{\mathbf{U}}_s^\dagger \mathbf{L}^{-1} \hat{\mathbf{D}}^{-1/2} \hat{\mathbf{U}}^\dagger \tilde{\mathbf{X}}. \quad (\text{B10})$$

Note that \mathbf{M} should project out any underlying monopole and dipole. The covariance in the case of the temperature is then given by

$$\langle \mathbf{X}_s \mathbf{X}_s^\dagger \rangle = \mathbf{N}_s + \mathbf{M} \text{diag}(C_l^{TT}) \mathbf{M}^\dagger, \quad (\text{B11})$$

where

$$\mathbf{N}_s = \tilde{\mathbf{M}} \mathbf{N} \tilde{\mathbf{M}}^\dagger. \quad (\text{B12})$$

If desired, and \mathbf{N}_s is sufficiently nonsingular, we can then write $\mathbf{N}_s = \mathbf{L}_s \mathbf{L}_s^\dagger$ and define the modes $\mathbf{L}_s^{-1} \mathbf{X}_s$ which have unit white noise (in general we can diagonalize even if the matrix is nearly singular). Note that although the signal variance is dominated by modes with $2 \leq l \leq l_{\text{exact}}$, \mathbf{M} couples in power up to l_{low} .

For the polarization we have

$$\tilde{\mathbf{E}} = \mathbf{W}_+ \mathbf{E} + i \mathbf{W}_- \mathbf{B}, \quad (\text{B13})$$

$$\tilde{\mathbf{B}} = \mathbf{W}_+ \mathbf{B} - i \mathbf{W}_- \mathbf{E}, \quad (\text{B14})$$

where

$$W_{\pm(l_1 m_1)(l_2 m_2)} = \frac{1}{2}(-1)^{m_1} \sum_l W_{lm} \sqrt{\frac{(2l_1+1)(2l_2+1)(2l+1)}{4\pi}} \left[\begin{pmatrix} l & l_1 & l_2 \\ 0 & 2 & -2 \end{pmatrix} \pm \begin{pmatrix} l & l_1 & l_2 \\ 0 & -2 & 2 \end{pmatrix} \right] \begin{pmatrix} l & l_1 & l_2 \\ m & -m_1 & m_2 \end{pmatrix} \quad (\text{B15})$$

and $m = m_1 - m_2$. Before continuing we change to using real rather than complex X_{lm} modes, as follows.

1. Real harmonics

It is convenient for numerical work to use real harmonics [33]. For the T , E , and B we define the real harmonic coefficients

$$X_{l|m}^R = \sqrt{2} \Re(X_{l|m}), \quad X_{l-|m}^R = \sqrt{2} \Im(X_{l|m}), \quad (\text{B16})$$

$$X_{l0}^R = X_{l0},$$

($|m| > 0$), which in the full-sky case are uncorrelated and have variance C_l^{XX} . The real coupling matrices then relate the pseudo and true harmonics

$$\tilde{\mathbf{T}}^R = \mathbf{W}^R \mathbf{T}^R, \quad (\text{B17})$$

$$\tilde{\mathbf{E}}^R = \mathbf{W}_+^R \mathbf{E}^R + \mathbf{W}_-^R \mathbf{B}^R, \quad (\text{B18})$$

$$\tilde{\mathbf{B}}^R = \mathbf{W}_+^R \mathbf{B}^R - \mathbf{W}_-^R \mathbf{E}^R, \quad (\text{B19})$$

where \mathbf{W}_+^R is symmetric and \mathbf{W}_-^R is antisymmetric. The real matrices are related to the complex ones by

$$\begin{aligned} \mathbf{W}_{l|m|l'|m'}^R &= \Re(\mathbf{W}_{l|m|l'|m'}) + (-1)^{m'} \mathbf{W}_{l|m|l'|-m'}, \\ \mathbf{W}_{l-|m|l'|m'}^R &= \Im(\mathbf{W}_{l|m|l'|m'}) + (-1)^{m'} \mathbf{W}_{l|m|l'|-m'}, \\ \mathbf{W}_{l|m|l'-|m'}^R &= \Im(-\mathbf{W}_{l|m|l'|m'}) + (-1)^{m'} \mathbf{W}_{l|m|l'|-m'}, \\ \mathbf{W}_{l-|m|l'|-|m'}^R &= \Re(\mathbf{W}_{l|m|l'|m'}) - (-1)^{m'} \mathbf{W}_{l|m|l'|-m'}, \\ \mathbf{W}_{l|m|l'0}^R &= \sqrt{2} \Re(\mathbf{W}_{l|m|l'0}), \\ \mathbf{W}_{l-|m|l'0}^R &= \sqrt{2} \Im(\mathbf{W}_{l|m|l'0}), \\ \mathbf{W}_{l0l'|m'}^R &= \sqrt{2} \Re(\mathbf{W}_{l0l'|m'}), \\ \mathbf{W}_{l0l-|m'}^R &= -\sqrt{2} \Im(\mathbf{W}_{l0l'|m'}), \\ \mathbf{W}_{l0l'0}^R &= \Re(\mathbf{W}_{l0l'0}), \end{aligned} \quad (\text{B20})$$

where $|m|, |m'| > 0$ and \mathbf{W} can be replaced by \mathbf{W}_+ or $i\mathbf{W}_-$ to obtain the equivalent results for \mathbf{W}_\pm^R . The noise on the real harmonics (for equal and uncorrelated noise on Q and U) is given by

$$\langle \mathbf{E}^R (\mathbf{E}^R)^T \rangle_N = \langle \mathbf{B}^R (\mathbf{B}^R)^T \rangle_N = \mathbf{W}_+^{RN}, \quad (\text{B21})$$

$$\langle \mathbf{E}^R (\mathbf{B}^R)^T \rangle_N = -\langle \mathbf{B}^R (\mathbf{E}^R)^T \rangle_N = \mathbf{W}_-^{RN}, \quad (\text{B22})$$

where \mathbf{W}_\pm^{RN} is evaluated with window function $w_N(s) = \Omega_s w(s)^2 \sigma_P^2(s)$.

2. Polarization modes

It is convenient to recomplexify the polarization analysis by defining $\mathbf{P} = \mathbf{E}^R + i\mathbf{B}^R$ so that

$$\tilde{\mathbf{P}} = (\mathbf{W}_+^R - i\mathbf{W}_-^R) \mathbf{P} \equiv \mathbf{W}_P \mathbf{P}, \quad (\text{B23})$$

where \mathbf{W}_P is Hermitian. The mode construction therefore goes through exactly as for the temperature, except now all matrices are complex, with $\langle \tilde{\mathbf{P}}^* \tilde{\mathbf{P}}^\dagger \rangle \neq 0$. The noise covariance under stated assumptions is given by

$$\langle \tilde{\mathbf{P}} \tilde{\mathbf{P}}^\dagger \rangle_N = 2\mathbf{W}_P^N, \quad \langle \tilde{\mathbf{P}}^* \tilde{\mathbf{P}}^\dagger \rangle_N = 0. \quad (\text{B24})$$

The real set of modes we end up with is then $\mathbf{X}_{\text{TEB}} \equiv \{\mathbf{T}_s^R, \Re(\mathbf{P}_s), \Im(\mathbf{P}_s)\}$, which includes the E/B mixed modes. Nearly pure E and B would be obtained by keeping only the well-supported modes of \mathbf{W}_+ rather than the well-supported modes of \mathbf{W}_P (see Ref. [31]). To calculate the signal to noise eigenmodes we can, for example, take $C_l^{EE} = C_l^{BB}$, where C_l^{EE} is a high optical depth model, to ensure that no potentially interesting modes are lost.

3. Implementation

There is some freedom in how Eq. (B11) is calculated. N_s can be precomputed assuming we know the noise model and are not fitting it from the data. The coupling matrix \mathbf{M} has size $\mathcal{O}(l_{\text{exact}}^2) \times \mathcal{O}(l_{\text{low}}^2)$. If there is plenty of memory, $\mathcal{O}(l_{\text{low}})$ matrices $\sum_m M_{i(lm)} M_{j(lm)}^*$ can be precomputed for each l , so that calculating the covariance is quick, but taking up $\mathcal{O}(l_{\text{exact}}^4 l_{\text{low}})$ memory. Calculating the likelihood is then dominated by the cost of Cholesky decomposition, $\mathcal{O}(l_{\text{exact}}^6)$, which is quite fast for $l_{\text{exact}} \sim 30$. Alternatively the covariance can be calculated on the fly at a dominating computational cost of $\mathcal{O}(l_{\text{exact}}^4 l_{\text{low}}^2)$ (and only storing \mathbf{M} of size $\mathcal{O}(l_{\text{exact}}^2 l_{\text{low}}^2)$).

Since information at $l > l_{\text{exact}}$ is subdominant, if l_{exact} is chosen so that in all models of interest the C_l are of well-determined shape at $l > l_{\text{exact}}$, it may be possible to precompute the most time-consuming contribution to the covariance from $l_{\text{exact}} < l \leq l_{\text{low}}$ and simply scale it by some weighted average of the spectrum over that l range. Certainly for $l_{\text{exact}} \sim 30$ all the spectra are expected to be very smooth up to $l_{\text{low}} \sim 100$ and this should work well. It could also be fixed at some fiducial model, but there is then a danger of biasing the likelihood from $l \leq l_{\text{exact}}$ by misestimating the contribution to the variance from higher l .

Using $l_{\text{exact}} = 30$, $l_{\text{low}} = 100$, $\epsilon_1 = 0.01$, $\epsilon = 10^{-3}$ seems to work well (for Planck a significantly larger ϵ_1 can be used for the temperature since the noise is very small). For an optimal tensor mode analysis we may want

to push to $l_{\text{exact}} \sim 150$ to get all the B -mode power, which is just about numerically tractable. The low- l harmonic likelihood described here has been implemented in the COSMOMC⁴ package for parameter estimation since the February 2008 version.

It is possible that the low- l likelihood can be approximated very accurately for very fast subsequent evaluation.

⁴<http://cosmologist.info/cosmomc/>

For example, the likelihood approximation of Ref. [3] could be applied to maximum-likelihood power spectrum estimators, or parameters in a likelihood model could be fit to accurately reproduce a full calculation [34]. In this case a near-exact low- l method would be an important step for testing or calibrating the approximation, and any speed hit of a harmonic-space approach would be much less important compared to possible accuracy advantages over a pixel-based method. Another possible fast approximation could come from fitting Gibbs-sampling results [27].

-
- [1] J. R. Bond, A. H. Jaffe, and L. E. Knox, *Astrophys. J.* **533**, 19 (2000).
- [2] L. Verde *et al.*, *Astrophys. J. Suppl. Ser.* **148**, 195 (2003).
- [3] S. Hamimeche and A. Lewis, *Phys. Rev. D* **77**, 103013 (2008).
- [4] B. D. Wandelt, D. L. Larson, and A. Lakshminarayanan, *Phys. Rev. D* **70**, 083511 (2004).
- [5] M. Tegmark, *Phys. Rev. D* **55**, 5895 (1997).
- [6] B. D. Wandelt, E. Hivon, and K. M. Gorski, *Phys. Rev. D* **64**, 083003 (2001).
- [7] E. Hivon *et al.*, *Astrophys. J.* **567**, 2 (2002).
- [8] F. K. Hansen, K. M. Gorski, and E. Hivon, *Mon. Not. R. Astron. Soc.* **336**, 1304 (2002).
- [9] G. Efstathiou, *Mon. Not. R. Astron. Soc.* **349**, 603 (2004).
- [10] M. L. Brown, P. G. Castro, and A. N. Taylor, *Mon. Not. R. Astron. Soc.* **360**, 1262 (2005).
- [11] G. Hinshaw *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **170**, 288 (2007).
- [12] G. Efstathiou, *Mon. Not. R. Astron. Soc.* **370**, 343 (2006).
- [13] K. M. Smith and M. Zaldarriaga, *Phys. Rev. D* **76**, 043001 (2007).
- [14] S. Smith, A. Challinor, and G. Rocha, *Phys. Rev. D* **73**, 023517 (2006).
- [15] S. P. Oh, D. N. Spergel, and G. Hinshaw, *Astrophys. J.* **510**, 551 (1999).
- [16] M. A. J. Ashdown, C. Baccigalupi, A. Balbi, J. G. Bartlett, J. Borrill, C. Cantalupo, G. de Gasperis, K. M. Górski, V. Heikkilä, E. Hivon, *et al.*, *Astron. Astrophys.* **471**, 361 (2007).
- [17] A. Lewis and S. Bridle, *Phys. Rev. D* **66**, 103511 (2002).
- [18] A. Lewis, *Phys. Rev. D* **78**, 023002 (2008).
- [19] C. L. Reichardt *et al.*, *Astrophys. J.* **694**, 1200 (2009).
- [20] W. C. Jones *et al.*, *Astrophys. J.* **647**, 823 (2006).
- [21] A. Gupta and D. Nagar, *Matrix Variate Distributions* (Chapman and Hall, London, 1999), ISBN 1584880465.
- [22] J. Hartlap, P. Simon, and P. Schneider, *Astron. Astrophys.* **464**, 399 (2007).
- [23] O. Ledoit and M. Wolf, *Journal of Empirical Finance* **10**, 603 (2003).
- [24] J. Schaefer and K. Strimmer, *Statistical Applications in Genetics and Molecular Biology* **4**, 32 (2005), <http://www.bepress.com/sagmb/vol4/iss1/art32>.
- [25] A. C. Pope and I. Szapudi, *Mon. Not. R. Astron. Soc.* **389**, 766 (2008).
- [26] J. Dunkley *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **180**, 306 (2009).
- [27] O. Rudjord *et al.*, *Astrophys. J.* **692**, 1669 (2009).
- [28] A. Slosar, U. Seljak, and A. Makarov, *Phys. Rev. D* **69**, 123003 (2004).
- [29] G. Efstathiou, S. Gratton, and F. Paci, arXiv:0902.4803.
- [30] C. Armitage-Caplan and B. D. Wandelt, *Astrophys. J. Suppl. Ser.* **181**, 533 (2009).
- [31] A. Lewis, *Phys. Rev. D* **68**, 083509 (2003).
- [32] J. R. Bond, *Phys. Rev. Lett.* **74**, 4369 (1995).
- [33] D. J. Mortlock, A. D. Challinor, and M. P. Hobson, *Mon. Not. R. Astron. Soc.* **330**, 405 (2002).
- [34] K. Benabed, J. F. Cardoso, S. Prunet, and E. Hivon, arXiv:0901.4537.
- [35] C. L. Bennett *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **148**, 1 (2003).
- [36] B. Gold *et al.* (WMAP Collaboration), *Astrophys. J. Suppl. Ser.* **180**, 265 (2009).
- [37] K. M. Gorski *et al.*, *Astrophys. J.* **622**, 759 (2005).