



The parameter-level performance of covariance matrix conditioning in cosmic microwave background data analyses

L. Balkenhol ★ and C. L. Reichardt 

School of Physics, University of Melbourne, Parkville, VIC 3010, Australia

Accepted 2022 February 24. Received 2022 February 20; in original form 2021 October 12

ABSTRACT

Empirical estimates of the band power covariance matrix are commonly used in cosmic microwave background (CMB) power spectrum analyses. While this approach easily captures correlations in the data, noise in the resulting covariance estimate can systematically bias the parameter fitting. Conditioning the estimated covariance matrix, by applying prior information on the shape of the eigenvectors, can reduce these biases and ensure the recovery of robust parameter constraints. In this work, we use simulations to benchmark the performance of four different conditioning schemes, motivated by contemporary CMB analyses. The simulated surveys measure the TT , TE , and EE power spectra over the angular multipole range $300 \leq \ell \leq 3500$ in $\Delta\ell = 50$ wide bins, for temperature map-noise levels of 10, 6.4, and $2 \mu\text{K arcmin}$. We divide the survey data into $N_{\text{real}} = 30, 50, \text{ or } 100$ uniform subsets. We show the results of different conditioning schemes on the errors in the covariance estimate, and how these uncertainties on the covariance matrix propagate to the best-fitting parameters and parameter uncertainties. The most significant effect we find is an additional scatter in the best-fitting point, beyond what is expected from the data likelihood. For a minimal conditioning strategy, $N_{\text{real}} = 30$, and a temperature map-noise level of $10 \mu\text{K arcmin}$, we find the uncertainty on the recovered best-fitting parameter to be $\times 1.3$ larger than the apparent posterior width from the likelihood ($\times 1.2$ larger than the uncertainty when the true covariance is used). Stronger priors on the covariance matrix reduce the misestimation of parameter uncertainties to < 1 per cent. As expected, empirical estimates perform better with higher N_{real} , ameliorating the adverse effects on parameter constraints.

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

1 INTRODUCTION

Current cosmological surveys provide new insight on inflation, dark matter, dark energy, and neutrino properties (Dodelson et al. 2013; Abazajian et al. 2015; Aiola et al. 2020; Planck Collaboration I 2020a; Dutcher et al. 2021). Cosmic microwave background (CMB) and large-scale structure (LSS) observations complement one another by providing snapshots of the universe at different times. The interplay of these two probes allows us to constrain structure formation and modified gravity models (Omori et al. 2019).

A crucial ingredient in many data analyses is the covariance matrix, which encodes the uncertainties on the measured data points. An accurate covariance matrix is essential for producing reliable model constraints; errors in the covariance matrix directly propagate into skewed cosmological parameters (Hartlap, Simon & Schneider 2007; Hamimeche & Lewis 2009; Dodelson & Schneider 2013; Taylor, Joachimi & Kitching 2013; Friedrich et al. 2016).

Since analytic calculations of the covariance matrix may struggle to capture the full complexity of the instrument, estimators are frequently used on simulated or real data to generate a robust estimate of the covariance matrix. While one can use simulations to produce a covariance estimate to any desired precision, this requires

a good understanding of the noise properties of the experiment, and, perhaps more importantly, can be very computationally expensive. Dodelson & Schneider (2013) note that modern surveys underestimate their parameter uncertainties by $\sim 5\text{--}10$ per cent, due to only running $\mathcal{O}(100)$ simulations (e.g. Louis et al. 2017b; Henning et al. 2018; Planck Collaboration V 2020b). Shrinkage estimators (Hamimeche & Lewis 2009; Hall & Taylor 2019), the CARPOOL algorithm (Chartier et al. 2021), data compression via MOPED (Heavens et al. 2017), emulators (Schneider et al. 2008; Morrison & Schneider 2013), and large-scale mode-resampling methods (Schneider et al. 2011) can ameliorate the increasing computational cost of simulations by reducing the number of realizations required to achieve adequate precision.

Alternatively, the measured data itself can be used to estimate the covariance matrix (e.g. Lueker et al. 2010; Das et al. 2014; Crites et al. 2015; Henning et al. 2018; Dutcher et al. 2021). This approach does not make any assumptions about the noise profile of the experiment and can avoid or reduce the required number of computationally expensive simulations. Empirical estimators require that the data is divided into uncorrelated, uniform subsets, i.e. independent realizations. The specific details of how the survey was conducted generally impose an upper limit on how many such subsets can be created, which can limit the accuracy of the covariance matrix estimate.

* E-mail: lbalkenhol@student.unimelb.edu.au

Conditioning can be applied to more accurately estimate the covariance matrix obtained from any of the above methods (e.g. Lueker et al. 2010; Das et al. 2014; Polarbear Collaboration 2014; BICEP2/Keck Collaboration 2015; Crites et al. 2015; Louis et al. 2017b; Henning et al. 2018; Choi et al. 2020; Dutcher et al. 2021). Conditioning schemes introduce prior knowledge of the correlation structure of the covariance matrix or of the values of specific entries, in order to reduce the overall uncertainty on the covariance estimate. The details of the conditioning approach vary drastically between experiments, ranging from no conditioning to strong priors. The approach chosen for a given analysis is often not justified in detail, and the effects on the final parameter constraints are rarely discussed.

This work fills this gap. Focussing on CMB data analysis, we investigate the impact of conditioning at the covariance- and the parameter level. This is done using a suite of simulations, mimicking access to $N_{\text{real}} = 30, 50, \text{ or } 100$ data realizations and temperature map-noise levels of $\sigma_{\text{map}} = 10, 6.4, \text{ and } 2 \mu\text{K arcmin}$. We assume the TT , TE , and EE power spectra are measured over the angular multipole range $300 \leq \ell \leq 3500$ in bins with $\Delta\ell = 50$. We report results for four different conditioning strategies, paying special attention to signs of bias and uncertainty misestimation in the parameter constraints. We also look briefly at how these results shift when changing the combination of temperature and polarization spectra analysed, and the width of band power bins.

This work is structured as follows. In Section 2, we review the necessary background information on the covariance matrix, how to estimate it empirically, and define the four conditioning schemes we benchmark later on. In Section 3, we describe the simulation framework used. We analyse our results in Section 4, first in covariance space, then at the parameter level, before concluding in Section 5.

2 BAND POWER COVARIANCE

2.1 Empirical estimation

The band power covariance matrix, \mathcal{C} , describes the uncertainty on band power measurements and the correlations between them. It is defined as

$$C_{bb'} = \langle (C_b - \bar{C}_b) (C_{b'} - \bar{C}_{b'}) \rangle, \quad (1)$$

where C_b is the binned power spectrum and \bar{C}_b the average thereof (e.g. Hivon et al. 2002; Tristram et al. 2005). Tristram et al. (2005) show that under certain assumptions we can model the covariance matrix as

$$C_{bb'}^{AB,CD} \simeq \frac{1}{\nu_b} (C_b^{AC} C_b^{BD} + C_b^{AD} C_b^{BC}) \delta_{bb'}, \quad (2)$$

where $\{A, B, C, D\} \in \{T, E, B\}$ are any temperature or polarization field of the CMB, and ν_b is the effective number of modes in a given band power bin.

However, modelling the experimental noise sufficiently accurately to use these theoretical prescriptions can be difficult. Equation (2) also assumes the mode-coupling matrix can be treated as diagonal, which is rarely satisfied given filtering choices and the requirement to mask bright radio galaxies and the Milky Way. The mode-coupling reduces the accuracy of straightforward analytic methods to $\mathcal{O}(10 \text{ per cent})$ (Hamimeche & Lewis 2009).

Alternatively, the covariance matrix can be estimated empirically. This approach relies on forming a uniform set of N_{real} realizations of the experimental data, by dividing it into uniform subsets. These realizations form the basis of estimating the noise and signal–noise

interaction term of the covariance matrix, which we refer to together as the noise-variance contribution, $C_{bb'}^N$. Adding the signal term, i.e. the sample variance, $C_{bb'}^S$, to the noise variance gives the full band power covariance matrix: $C_{bb'} = C_{bb'}^S + C_{bb'}^N$.

By calculating the $N_{\text{real}}(N_{\text{real}} - 1)/2$ unique cross spectra of the data realizations, $\hat{C}_\ell^{\alpha\beta}$, $\alpha \neq \beta$, where α and β index the realization, we can use the observed scatter of the distribution of band powers to estimate part of the covariance matrix. As shown by Lueker et al. (2010), the correct noise-variance contribution to the covariance is provided by the estimator

$$\hat{C}_{bb'}^N = \frac{2f(N_{\text{real}})}{N_{\text{real}}^4} \sum_{\lambda} \sum_{\alpha \neq \lambda} \left[\Delta C_b^{\lambda\alpha} \Delta C_{b'}^{\lambda\alpha} + 2 \left(\sum_{\beta \neq \lambda, \alpha} \Delta C_b^{\lambda\alpha} \Delta C_{b'}^{\lambda\beta} \right) \right], \quad (3)$$

where ΔC_b are the mean-subtracted binned power spectra and $f(N_{\text{real}})$ is a correction for the finite number of realizations that approaches unity for large N_{real} . In this limit, $\langle \hat{C}_{bb'}^N \rangle = C_{bb'}^N$. This estimator is easily extended for the cross-frequency case by the using the desired frequencies for any of the two spectra in each product, but not modifying the indices of the data splits, so that the first term in the sum still represents the sample variance of the cross-spectra and the second term encapsulates the variance of cross-spectra with one map in common (cf. A18 in Lueker et al. 2010).

The sample-variance contribution can be obtained from the signal-only simulations typically used for the transfer-function calculation; since these simulations are uncorrelated, the variance of band powers provides an unbiased estimate of $C_{bb'}^S$. Therefore, by adding the noise-variance estimator to the variance of the signal-only simulation spectra, the band power covariance matrix can be estimated.

Empirical estimators provide a noisy estimate of the band power covariance matrix, drawn from a Wishart distribution. The number of realizations N_{real} sets the precision. Lueker et al. (2010) show that each element has a statistical variance of

$$\langle (C_{bb'} - \langle C_{bb'} \rangle)^2 \rangle = \frac{C_{bb'}^2 + C_{bb} C_{b'b'}}{N_{\text{real}}}. \quad (4)$$

Errors in the covariance matrix estimate propagate through to cosmological parameters, widening the posterior distributions and biasing the best-fitting point (Hartlap et al. 2007; Hamimeche & Lewis 2009; Dodelson & Schneider 2013; Taylor et al. 2013). Increasing N_{real} reduces the errors, and thus the magnitude of these effects, although the observation strategy often limits the possible number of data splits for real data. For simulations, the computing cost limits the number of realizations available.

2.2 Conditioning

Conditioning the covariance matrix estimate can reduce the errors for a fixed N_{real} . This step takes place before parameter estimation and uses prior knowledge to constrain noisy elements of the covariance matrix. The details of conditioning schemes vary drastically between analyses. For example, Polarbear Collaboration (2014), Das et al. (2014), and Louis et al. (2017b) drop the off-diagonal elements of all blocks, i.e. ignore mode-coupling. Choi et al. (2020) allow for mode-coupling to the nearest angular multipole bin, but explicitly zero the effects of filtering and masking further than that. Similarly, The BICEP/Keck Collaboration (2018) keep correlations in the first off-diagonal bins, but further reduce the noise of the covariance estimate by explicitly ignoring contributions with an expectation value of

zero. Crites et al. (2015) and Henning et al. (2018) both average band-diagonal elements in the autocorrelation matrices to address masking effects. A direct manipulation of on-diagonal elements is performed by Dutcher et al. (2021) and Reichardt et al. (2021) (as well as preceding analyses targeting secondary anisotropies).

To navigate this broad landscape, we define four different, increasingly strict conditioning strategies, which we will benchmark in this work. These act on the complete covariance matrix, i.e. the sum of the noise-variance and sample-variance estimators.

(i) *No conditioning*: The final covariance is given by the noise-variance plus sample-variance estimators.

(ii) *Diagonal conditioning*: We calculate the diagonal elements of all covariance blocks as before, but set all off-diagonal elements to zero.

(iii) *Moderate conditioning*: Diagonal elements are calculated as before. Off-diagonal elements of the auto-covariance blocks are band-averaged in the correlation matrix for bins within $\Delta\ell \leq 100$ of the diagonal, but set to zero for $\Delta\ell > 100$. All off-diagonal elements of off-diagonal covariance blocks are set to zero.

(iv) *Deep conditioning*: The conditioning approach detailed in Dutcher et al. (2021), which we briefly summarize here. An estimate of the effective number of modes, or any prefactor in equation (2), is extracted by comparing the variance of signal-only simulation spectra to the expected signal-variance based on the input power spectrum to the simulation. This estimate can then be used together with the empirical estimator in equation (3) to obtain an estimate of the noise spectrum. Using the number of modes, the noise spectrum, and the input signal to the simulations, we can replace the diagonal elements of the noisy covariance estimate with the expectation values of equation (2). Off-diagonal elements are replaced by performing second-order polynomial fits to band-diagonal elements in the correlation matrices.

3 SIMULATION FRAMEWORK

We run simulations to gauge to what extent the conditioning schemes outlined in Section 2.2 protect against the adverse effects of a noisy covariance matrix estimate. To reduce the computational cost, we mimic the observation of a modest $\sim 290 \text{ deg}^2$ field, comprised of 512×512 square pixels of 2 arcmin width, using the flat-sky approximation. The underlying CMB signal is based on a power spectrum calculated by CAMB (Lewis, Challinor & Lasenby 2000),¹ with cosmological parameters set to the latest *Planck* results (Planck Collaboration VI 2020d). We add white and $1/f$ noise at the map level, which follows the power spectrum

$$N_\ell^X = N_{\text{white}}^X \left[1 + \left(\frac{\ell}{\ell_{\text{knee}}} \right)^{\alpha^X} \right], \quad (5)$$

where $X \in \{T, E\}$, $\alpha^T = -3.5$, $\ell_{\text{knee}}^T = 1000$, $\alpha^E = -1.4$, $\ell_{\text{knee}}^E = 700$, and N_{white}^X is the white noise level, which is in turn set by the map-noise level, σ_{map} . We assume that the white noise power in polarization is twice that of temperature. This noise curve is chosen to approximate the anticipated noise characteristics of the Simons Observatory (SO) large aperture telescope (cf. equation 1 in Simons Observatory Collaboration 2019). We simulate having access to N_{real} maps of the same CMB signal with independent noise realizations added to each one. We proceed to measure the TT , TE , and EE cross-spectra between all maps across the angular multipole range $300 \leq$

$\ell < 3500$, binned into bins of width $\Delta\ell = 50$. The power spectrum measurements or band powers for each simulation are estimated by averaging across all the cross-spectra. The simulated power spectrum measurements are complemented by $N_{\text{sim}} = 300$ signal-only simulations, which mimic the typical approach to calculating the transfer function.

We run nine sets of simulations, for $N_{\text{real}} = 30, 50, \text{ or } 100$, and $\sigma_{\text{map}} = 10, 6.4, \text{ or } 2 \mu\text{K arcmin}$. The number of realizations is chosen to span the range seen in contemporary analyses, without inflating computation cost (e.g. Louis et al. 2017a; Henning et al. 2018a; Choi et al. 2020; Planck Collaboration V 2020b; Dutcher et al. 2021). The map-noise levels are quoted for the co-added N_{real} maps and are chosen to match the SO baseline configuration, the SO goal, and the projected depth of the complete SPT-3G survey (Bender et al. 2018; Simons Observatory Collaboration 2019). We run 1000, 500, 100 simulations for $N_{\text{real}} = 30, 50, 100$, respectively.

For each simulation, we estimate the band power covariance matrix as detailed in Section 2.1 and condition copies of it as described in Section 2.2, i.e. each simulation is associated with four differently conditioned versions of the same covariance matrix estimate. Furthermore, we calculate the average of all unconditioned covariance matrices for a given N_{real} and σ_{map} combination, which we refer to as the fiducial covariance matrix. We also calculate a theoretical covariance void of mode-coupling using the prescriptions in equation (2).

We note that the sky area chosen here is much less than the coverage of current and next-generation ground-based CMB surveys. However, the sky fraction only serves to scale the overall amplitude of the covariance matrix. Ultimately, we are interested in a comparative study of conditioning schemes, not their absolute performance. Therefore, the shape of the covariance matrix and the noise on the covariance estimate, as set by σ_{map} and N_{real} , respectively, are of greater interest here.

4 RESULTS

4.1 Comparison in covariance space

First, we look at how well the true covariance matrix can be estimated, for a fixed number of realizations, under each conditioning scheme. We assess this by looking at the eigenvalues of differently conditioned covariance matrices for the $N_{\text{real}} = 30$ and $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations in Fig. 1.

We see that the largest ~ 70 eigenvalues are recovered well by all conditioning schemes, though the unconditioned matrices show the largest fluctuations around the theoretical expectation values. We consider fluctuations > 20 per cent away from the fiducial values as failures to retrieve the correct eigenvalues. For the unconditioned matrices, 61 per cent of eigenvalues are within this threshold. The performance of the other conditioning schemes is comparable to one another for the largest ~ 70 eigenvalues, though for smaller eigenvalues they begin to deviate; 14 per cent of the eigenvalues of the diagonal and moderately conditioned covariance deviate by more than 20 per cent from the fiducial case. Deeply conditioned matrices recover all eigenvalues, with no statistically significant deviation from the theoretical covariance.

The estimator given in equation (3) does not guarantee that the resulting covariance matrix is positive definite. Particularly noisy estimates, i.e. using few data splits, are prone to negative eigenvalues, which must be addressed before the data can be analysed further. Different conditioning strategies ameliorate this issue with varying degrees of success, though deep conditioning

¹<https://camb.info/>




Figure 1. Comparison of the eigenvalues of covariance matrices for different conditioning strategies: unconditioned (black), diagonal (green), moderate (blue), and deep (red). We plot the ratio of the average of the absolute eigenvalues across all $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations in increasing order, to the eigenvalues of the theoretical covariance matrix. The shaded regions correspond to the standard deviation around the mean. The eigenvalues of unconditioned and mildly conditioned covariances recover the largest third and half of eigenvalues well, respectively – beyond that the eigenvalues drop off. The dotted grey lines indicate our threshold for acceptable deviations from the fiducial case, ≤ 20 per cent. For the unconditioned covariances 39 per cent of values lie outside this range, whereas 14 per cent of the eigenvalues of diagonal and moderately conditioned matrices show large deviations from the expectation. Deep conditioning recovers the entire range of eigenvalues.

yields positive-definite covariance matrices in all cases tested. As the breakdown in positive definiteness is related to the magnitude of the noise on the covariance estimate, increasing the number of realizations generally improves the situation (cf. equation 4). For example, for the $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations, 24 per cent, 19 per cent, and 12 per cent of the eigenvalues of unconditioned matrices are negative for $N_{\text{real}} = 30$, 50, and 100, respectively. Diagonal and moderate conditioning perform better: at $N_{\text{real}} = 30$ only 0.4 per cent and 0.3 per cent of eigenvalues are negative, which reduces to 0.03 per cent and 0.02 per cent for $N_{\text{real}} = 50$, respectively. For $N_{\text{real}} = 100$, diagonal and moderate conditioning match the performance of deep conditioning and we register no negative eigenvalues.

This work varies the number of realizations used in the empirical covariance estimate for the noise variance, but uses a fixed higher number of realizations to estimate the sample variance. As a result, the number of negative eigenvalues decreases in these tests as the angular multipole range dominated by sample variance increases (i.e. as the map-noise level decreases). The sample-variance contribution to the covariance is always comparatively well constrained with 300 realizations. For example, the average fraction of negative eigenvalues for unconditioned covariance matrices reduces from 24 per cent for $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$, to 20 per cent and 14 per cent for $\sigma_{\text{map}} = 6.4 \mu\text{K arcmin}$ and $\sigma_{\text{map}} = 2 \mu\text{K arcmin}$, respectively, for the $N_{\text{real}} = 30$ simulations.

One can also improve the performance of empirical covariance estimators by reducing the number of entries in the covariance. In this context, the number can be reduced by using a coarser angular multipole binning for the power spectrum measurement. For the $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations, the temperature and E-mode power spectrum is signal dominated for $\ell < 3300$ and $\ell < 1750$, respectively. Doubling the bin size to $\Delta\ell = 100$ decreases

the fraction of negative eigenvalues in the unconditioned covariances from 24 per cent to 17 per cent. Similarly, for diagonal and moderate conditioning the fraction of negative eigenvalues is reduced to 0.02 per cent and 0.01 per cent, respectively. The performance of deep conditioning does not change. While such rebinning can help reduce the noise of the covariance matrix, a coarser binning of the power spectrum can result in wider parameter constraints.

Given the large fraction of negative eigenvalues of the unconditioned matrices, we exclude this scheme from further analysis.

4.2 Impact on cosmological parameters

While errors on the covariance are easy to measure, the real statistics of interest are the effects on the derived parameter constraints during model fitting. Our analysis focuses on the following key questions: is the parameter uncertainty over- or underestimated? Is the best-fitting point biased? Are the derived constraints suboptimal, i.e. do the results improve appreciably if the covariance estimate is improved? Before pursuing these questions and reporting our results, we briefly describe our analysis methodology below.

We benchmark the effects on parameter fitting under the assumption of the lambda cold dark matter (ΛCDM) model, using model power spectra calculated by CAMB (Lewis, Challinor & Lasenby 2000). The six model parameters are: the amplitude of primordial density perturbations, A_s , the tilt of the power spectrum, n_s , defined at a pivot scale of 0.05 Mpc^{-1} ; the density of cold dark matter $\Omega_c h^2$; the baryon density $\Omega_b h^2$; the optical depth to reionization, τ ; and θ_{MC} , the approximation to ratio of the sound horizon to the angular diameter distance used by CAMB.

To reduce the computational cost of fitting cosmological parameters on $\mathcal{O}(5000)$ simulations, we primarily focus on a single parameter: θ_{MC} , the approximation to the ratio of the sound horizon to

the angular diameter distance. We fix other cosmological parameters to the best-fitting value of the signal-only spectrum obtained from combining all transfer-function simulations across the nine sets of simulations. We explore the range $1.046912 \geq \theta_{\text{MC}} \geq 1.034112$ around this fiducial θ_{MC} value, although in practice we filter out any likelihoods that peak close to the edge of the range – outside of the $\sim 7.5\sigma$ region of the least constraining simulation.

We investigate the statistics in the full six-dimensional parameter space for a subset of 100 of the $N_{\text{real}} = 30$ and $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations. We use Cobaya (Torrado & Lewis 2021) to find the minimum of the likelihood for these realizations. Since the simulated surveys lack access to information from large angular scales, we fix the optical depth to reionization τ to the input value. Using the sampler of Lewis & Bridle (2002), we run a full Markov Chain Monte Carlo for a single realization using the fiducial band power covariance to obtain a fiducial parameter covariance matrix.²

We calculate the joint likelihood of the *TT*, *TE*, and *EE* band powers for each simulation using a Gaussian approximation:

$$\log \mathcal{L} = -0.5 \Delta C_b C_{bb'}^{-1} \Delta C_{b'}, \quad (6)$$

where ΔC_b is the difference of the simulation and model band powers and $C_{bb'}$ is the covariance matrix. This procedure is carried out for each realization using the diagonally, moderately, and deeply conditioned covariance matrices, as well as the fiducial covariance.

We set any negative eigenvalues of the covariance matrices to a large, positive number, six orders of magnitude larger than maximum eigenvalue, in order to guarantee positive definiteness for parameter estimation. We analyse the likelihood curves of all simulations, identifying the peak of the likelihood as the best-fitting point and obtaining the uncertainty of the parameter constraint from the curvature of the likelihood around this point. We filter the results for outliers as follows. First, we reject a realization if the best-fitting value of any conditioned matrix peaks towards the edge of the parameter range as mentioned before. Furthermore, we perform sigma-clipping (5σ , five iterations), based on the best-fitting values and parameter uncertainty. For all simulations and analysis, cases >96 per cent of realizations survive these cuts and we therefore do not expect this filtering to bias our results significantly.

We compare parameter uncertainties first and investigate the bias afterwards. We highlight key results in detail below, and point the reader to Tables A1 and A2 in the Appendix for the full results.

4.2.1 Uncertainty

We inspect the uncertainty on parameter constraints for different conditioning schemes for the $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations and compare them to the fiducial covariance. Diagonal covariance matrices underestimate the parameter uncertainty by $5.8_{-1.9}^{+2.1}$ per cent. This underestimate is in-part due to zeroing the off-diagonal entries in the covariance matrix (which are non-zero due to mode-coupling), but also due to the excessive noise on the main diagonals of the covariance. The latter point is confirmed by the performance of the moderately conditioned covariance matrices, which can recover the off-diagonal entries but have the same noisy

²We assume that this covariance matrix provides an accurate description of the parameter error for all other simulations. This is strictly only true in a Gaussian parameter space. However, we are analysing fairly constraining CMB data in a restricted ΛCDM model space (five free parameters). Moreover, the mean size of the deviations observed in Section 4.2.2 is small ($< 0.5\sigma$). We therefore judge the effects of non-Gaussianity to be minimal.

estimate of the diagonals. The moderately conditioned case underestimates the uncertainty by $3.5_{-2.3}^{+2.7}$ per cent. The deep conditioning scheme, which also reduces the uncertainty on the diagonal entries, recovers the fiducial error on θ_{MC} to within 0.3 ± 1.5 per cent. These differences between the conditioning schemes are illustrated in the left-hand panel of Fig. 2. The uncertainty on θ_{MC} is artificially changed by the uncertainty in the covariance matrix for the diagonal and moderate conditioning cases, while the likelihood for deep conditioning closely traces the fiducial likelihood curve.

Increasing the number of realizations used in the covariance estimation will reduce the uncertainties (for a fixed number of band-powers), and decrease their impact on the parameter uncertainties. The right-hand panel of Fig. 2 shows the trend with increasing N_{real} for $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations. As expected, the uncertainties on θ_{MC} converge to the fiducial uncertainty for all conditioning schemes as the number of realizations, N_{real} , increases. At $N_{\text{real}} = 100$, moderate conditioning recovers the fiducial uncertainty well, although diagonal conditioning still underestimates the uncertainty by $2.4_{-1.2}^{+1.5}$ per cent due to zeroing the off-diagonal entries.

By considering the number of noise-variance dominated bins estimated, we can understand how these results change for: different map-noise levels; different band power bin widths; and subsets of the *TT*, *TE*, and *EE* power spectra. Lowering the map-noise level increases the angular multipole range over which the well-measured sample-variance part of the covariance dominates. This boosts the performance of the conditioning schemes; for example, for $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 2 \mu\text{K arcmin}$ simulations, the uncertainty underestimation of moderate conditioning is reduced from $3.5_{-2.3}^{+2.7}$ per cent to $0.5_{-1.2}^{+1.4}$ per cent. Similarly, when doubling the bin size from $\Delta\ell = 50$ to $\Delta\ell = 100$, moderate conditioning recovers the correct parameter uncertainty to $0.1_{-1.3}^{+1.2}$ per cent (at $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$). Since the temperature power spectrum is sample-variance dominated over the largest angular multipole range, we observe more robust constraints from *TT*-only data: moderate conditioning yields the correct parameter uncertainty at $1.3_{-2.5}^{+2.7}$ per cent. The improvement in each case is owed to the reduction in the number of covariance elements estimated from the limited data realizations.

4.2.2 Bias

We next turn to evaluating the level of bias in the recovered parameter constraints due to the noisy covariance estimate. We define the bias $\Delta\theta$ as the median difference of the best-fitting values of a given conditioning scheme compared to the best-fitting value of the fiducial covariance matrix. We also measure the additional scatter (or noise) introduced in the best-fitting values by the noisy covariance. We quote both quantities relative to the expected uncertainty, by dividing the measured bias or scatter by the expected uncertainty from a Fisher forecast, $\sigma_{\mathcal{F}}$.³

We use the median and σ_b , defined as half the central 68 per cent interval, to describe the two key quantities of interest, the bias and the extra noise to the best-fitting point, respectively. The bias is a global shift to the parameter constraint and depends on the specific point in parameter space being explored. In theory, this can be corrected for using a set of simulations. The extra noise on the other hand can be understood as a random shift to the best-fitting value, which is unique for each individual realization. The additional scatter is not captured by the data likelihood and independent of the change to

³We use the Fisher forecast to avoid introducing additional noise from the individual error bar estimates, although we find similar results in both cases.




Figure 2. *Left-hand panel:* We show stacked likelihood curves for $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations for diagonal (green), moderate (blue), and deep conditioning (red) and the fiducial covariance (black) in the top panel. The shaded region around the lines corresponds to the standard deviation across simulations. The curvature of the stacked likelihoods from diagonal, moderate, and deep conditioning deviate by 13 per cent, 8 per cent, and 1 per cent from the fiducial covariance, respectively, which translates into per cent-level differences in the uncertainty on θ_{MC} . The bottom panel shows the ratio of the different conditioning schemes to the stacked likelihood of the fiducial covariance. We mask the region $|\Delta\theta| \leq 0.0003$ in dark grey, because of the numerical artefact in dividing by zero. *Right-hand panel:* We plot the median of the ratio of θ_{MC} uncertainty inferred from the data likelihood, $\sigma_{\mathcal{L}}$, from diagonal (green), moderate (blue), and deep conditioning (red), to the fiducial covariance matrix, σ_0 , across all N_{real} , $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations. The wide and thin markers indicate the 68 per cent and 95 per cent interval, respectively. While shallow conditioning schemes tend to underestimate the parameter uncertainty, increasing the number of independent realizations N_{real} ameliorates this issue.

the curvature of the likelihood explored in Section 4.2.1. This effect needs to be accounted for explicitly through a modification of the likelihood or similar procedure.

Scaled bias distributions are shown in Fig. 3 for the $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations. For $N_{\text{real}} = 30$, diagonal and moderate conditioning lead to a sizeable scatter of parameter constraints, with $\sigma_b/\sigma_{\mathcal{F}} = 0.44$ and 0.47 , respectively. There is little difference in the best-fitting values obtained from these two schemes, which indicates that this extra noise is driven by the uncertainty on the main diagonal elements of the covariance matrix estimates. With medians of 0.16 and 0.17 for diagonal and moderate conditioning, respectively, we detect a bias at $\sim 10\sigma$. However, this bias is drowned out by the random scatter to the best-fitting point. Deep conditioning does not admit a sizeable bias or significant extra noise. The median of the distribution is consistent with zero and $\sigma_b/\sigma_{\mathcal{F}} = 0.11$.

We do not expect to reach zero scatter to the best-fitting point, due to residual noise in the fiducial covariance. We quantify this lower limit by running an additional set of simulations with the same setting and benchmarking the performance of the associated new fiducial covariance using the original simulations. We find that the half width of the 68 per cent interval of the scaled bias distribution spans 0.07 . This is over half of the level of scatter observed for deep conditioning quoted above, indicating that the random shifts admitted by the conditioning scheme are in practice appreciably lower and typically below the $0.1\sigma_{\mathcal{F}}$ level.

Fig. 3 also allows us to observe the evolution of the distributions with N_{real} . As expected, the distributions become tighter as more realizations become available, i.e. the 68 per cent and 95 per cent intervals shrink, and the median values move closer towards zero,

as the covariance estimate is improved. Practically, this means that the scatter to the parameter constraints, σ_b , shrinks and the uncertainty approaches the fiducial value. We focus on the evolution of the extra noise with N_{real} in Fig. 4, which shows this trend more clearly. However, even at $N_{\text{real}} = 100$, diagonal and moderate conditioning both still lead to non-negligible scatter of the best-fitting parameters, with $\sigma_b/\sigma_{\mathcal{F}} = 0.19$ and 0.21 , respectively. We detect no significant change in the performance of deep conditioning with N_{real} .

The extra noise dominates over the change to the curvature of the likelihood seen in Section 4.2.1 for most simulations tested. We combine the two effects by adding the standard deviation of the bias distributions to the parameter uncertainties measured in Section 4.2.1 in quadrature. Generally, the total effect is an increase in the parameter uncertainty compared to the fiducial case. For the $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations, we find the uncertainty on the recovered best-fitting parameter to be a factor of 1.3 larger than the apparent posterior width from the likelihood for diagonal and moderate conditioning. This corresponds to an increase by a factor of 1.2 and 1.3 compared to when the true covariance is used, respectively. Diagonal conditioning leads to a total uncertainty closer to unity, because the scatter of the best-fitting point is offset by the larger uncertainty underestimation from the curvature of the likelihood. This balance changes with increasing N_{real} , as the extra noise shrinks but the likelihood curve obtained from diagonal conditioning remains narrower than the fiducial case, leading to a total underestimation of the parameter uncertainty by $2.4^{+1.5}_{-1.2}$ per cent. For $N_{\text{real}} \geq 100$ the latter is no longer an issue for moderate conditioning, and as the scatter of the best-fitting point is




Figure 3. Noise in the covariance matrix leads to additional scatter in the parameter constraints. We show the probability density function (PDF) of the ratio of the parameter bias to the fiducial uncertainty, $\Delta\theta/\sigma_{\mathcal{F}}$, for $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations for diagonal (green), moderate (blue), and deep conditioning (red). With increasing vertical offset we show the $N_{\text{real}} = 30, 50,$ and 100 simulations and indicate the 68 per cent and 95 per cent interval in faint colour under the curves. Diagonal and moderate conditioning lead to a substantial scatter in the recovered parameters, as shown by the broad 68 per cent and 95 per cent intervals. This extra noise is not accounted for in the data likelihood and independent of the change to the curvature of the likelihood. Increasing N_{real} reduces uncertainty on the covariance estimate and leads to narrower distributions, i.e. the scatter to parameter constraints is decreased. Deep conditioning performs almost as well as the fiducial covariance matrix, even for $N_{\text{real}} = 30$. Moreover, with increasing N_{real} the distributions approach a Gaussian, with KS-test p -values in the 95 per cent interval for all conditioning schemes given $N_{\text{real}} = 100$. The PDFs shown here are estimated by computing a histogram of the scaled bias and smoothing it using a Gaussian kernel.

reduced the total uncertainty converges to the fiducial value. Deep conditioning recovers the fiducial uncertainty to less than 2 per cent in all cases tested.

We expect that the effect of noise in the covariance on the parameter constraints will scale as the ratio of the number of bins estimated to the number of realizations (Dodelson & Schneider 2013). For a fixed number of realizations, reducing the number of bins estimated will also reduce the additional uncertainty incurred due to the covariance estimate for all conditioning schemes. We have demonstrated this on all three axes: changing the map-noise level (varying the number of band powers dominated by noise variance); changing the number of band power bins; and reducing the number of spectra from $TT, TE,$ and EE to TT -only. For instance, for $N_{\text{real}} = 30$ and moderate conditioning, lowering the noise from $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ to $\sigma_{\text{map}} = 2 \mu\text{K arcmin}$ reduces $\sigma_b/\sigma_{\mathcal{F}}$ from 0.47 to 0.17. Similarly, halving the number of bins through a straightforward rebinning from $\Delta\ell = 50$ to $\Delta\ell = 100$ decreases the random scatter to 0.21 for moderate conditioning (at $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$). Lastly, analysing only the temperature power spectrum reduces $\sigma_b/\sigma_{\mathcal{F}}$ to 0.16, since the TT -band powers are sample-variance dominated over the largest angular multipole range.

While the tests so far have focused on single-parameter fits, we are generally interested in the size of the additional scatter due to noise in the covariance matrix in multiple-parameter fits, for instance to the full ΛCDM model. We use the best-fitting points of a subset of 100 $N_{\text{real}} = 30, \sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations and look at the additional scatter σ_b . We compare the size of the additional scatter to the uncertainty based on the fiducial parameter covariance matrix described in Section 4.2, σ_0 . We find that all cosmological parameters show a similar level of additional scatter in the marginalized 1D

posteriors. Taking the geometric mean across the five dimensions, the extra scatter with diagonal, moderate, and deep conditioning is $\sigma_b/\sigma_0 = 0.44, 0.50$ and 0.11 , respectively, statistically consistent with the results of 0.44, 0.47, and 0.11 in the single-parameter fits. Note, however, that this scatter applies to each individual axis in the parameter space, so if one looks at the distribution of the distance between the true and recovered best-fitting across N dimensions, the width expands by up to a factor of \sqrt{N} . Hence large effects can present themselves when analysing multidimensional contours, and when the number of parameters being fit is large (Dodelson & Schneider 2013; Taylor et al. 2013). We conclude that the estimates of the additional scatter due to noise in the covariance estimate in Section 4.2 and Fig. 4 are applicable to the marginalized 1D posterior constraints of each parameter in the full ΛCDM model.

5 CONCLUSION

In this work, we assess how effective conditioning noisy estimates of the covariance matrix is at producing robust parameter constraints. We benchmark the performance of four conditioning schemes imposing increasingly strong priors on the estimated covariance matrix: no conditioning, diagonal conditioning, moderate conditioning, and deep conditioning. We simulate $TT + TE + EE$ power spectrum analyses spanning an angular multipole range of $300 \leq \ell < 3500$ with band powers of width $\Delta\ell = 50$, mimicking access to $N_{\text{real}} = 30, 50,$ or 100 data realizations and map-noise levels of $\sigma_{\text{map}} = 10, 6.4,$ and $2 \mu\text{K arcmin}$. We find that, for all cases tested, conditioning the covariance matrix is necessary to obtain robust parameter constraints.

When examining the resulting parameter constraints, the largest effect of noise in the covariance estimate is a significant scatter in




Figure 4. Errors on the covariance matrix introduce additional (and unaccounted for) scatter in the recovered parameters. We show here the ratio of the additional scatter to the fiducial parameter uncertainty for simulations with $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$, parametrized by σ_b , the half width of the 68 per cent interval of $\theta/\sigma_{\mathcal{F}}$. The wide and thin markers indicate the $1-2\sigma$ range based on the number of samples for diagonal (green), moderate (blue), and deep conditioning (red), respectively. The dotted lines show an inverse power-law fit with slope N_{real}^{-1} and the semitransparent region corresponds to the 1σ region around it. Due to the small dataset, these fits only serve to guide the eye and are not statistically meaningful. The additional scatter to parameter constraints is reduced as the covariance estimate improves with higher N_{real} . We do not expect to be able to reach zero extra noise due to residual noise in the fiducial covariance matrix.

the best-fitting points. Crucially, this additional uncertainty is not reflected in the reported parameter posteriors, from e.g. an MCMC analysis of the data likelihood, and must be explicitly accounted for by some other means. Not accounting for this, additional scatter can significantly underestimate the uncertainty on cosmological parameters. For $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations, we find that the parameter uncertainties are a factor of 1.3 larger than the apparent posterior width from the likelihood for diagonal and moderate conditioning. Importantly, the distribution of this additional scatter is significantly non-Gaussian. The stricter priors of deep conditioning eliminate the increase in uncertainty. Increasing the number of data splits used in the covariance estimate also improves the situation: with 100 data realizations, we report no increase in the parameter uncertainty for diagonal and moderate conditioning. In many cases, we expect it will be necessary to add an explicit correction for this additional scatter in order to accurately capture the parameter uncertainty. While the optimal approach may vary between analyses, we alert the reader to Sellentin & Heavens (2016) and Percival et al. (2021), who offer solutions in the Bayesian and Frequentist inference frameworks.

We also report two other less significant phenomena: a distortion of the curvature of the likelihood and a bias of parameter constraints. For substantial noise in the covariance estimate ($N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K arcmin}$), diagonal and moderate conditioning schemes lead to an artificially tight likelihood; this translates into an underestimation of the parameter uncertainty by a few per cent. More data realizations

ameliorates this underestimate. However, for diagonal conditioning, which does not capture the full correlation structure of the data, the likelihood does not approach the fiducial curve and parameter errors remain modestly too small by $2.4^{+1.5}_{-1.2}$ per cent at $N_{\text{real}} = 100$. While we detect a small bias to parameter constraints induced by the residual noise in the covariance matrix, it is smaller than the additional scatter in all cases. Both the bias and uncertainty underestimation approach zero as N_{real} increases and the covariance estimate improves; the additional scatter is more significant than both effects at all times.

At a fixed number of data realizations, we expect the covariance matrix to become better constrained as the number of band powers decreases as there are fewer matrix elements (with the same fractional error per element). In the tests in this work, where the sample variance contributions were estimated from a much larger number of simulations, this means the effects of uncertainty in the covariance matrix should fall at lower map-noise levels as well as when looking at a brightest selection of TT , TE , and EE spectra, or widening the band power angular multipole bins. We confirm that this intuition is correct and find that these changes indeed lead to a reduction in the additional scatter, misestimation of the likelihood curvature, and bias.

We are interested in expanding upon this work, by exploring parameter constraints for a wider and finer grid of analysis parameters, including N_{real} , σ_{map} , the angular multipole range, the band power bin width, the number of transfer-function simulations, inclusion of the lensing power spectrum, and increasing the sky area observed. This extended set of simulations would allow us to provide scaling relations for the parameter uncertainty misestimation and bias for different conditioning schemes as these analysis parameters are varied. This would ultimately allow us to produce a practical guide for readers to determine what amount of conditioning is necessary for a given analysis. Additionally, we would like to explore further how parameter constraints behave in a multi-dimensional model space. Using interpolation approaches to obtain model CMB power spectra (e.g. Fendt & Wandelt 2007; Spurio Mancini et al. 2021) may make full explorations of the parameter posteriors computationally tractable. Current-generation ground-based CMB experiments have the potential of making exciting discoveries of physics beyond the standard model. Understanding and managing the sources of even small changes to parameter uncertainties and biases is crucial to producing reliable results.

ACKNOWLEDGEMENTS

The authors thank Piaera Lauritz for helpful advice on the nomenclature of conditioning schemes. We acknowledge support from the University of Melbourne and an Australian Research Council Future Fellowship (FT150100074). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under contract no. DE-AC02-05CH11231. The data analysis presented uses the scientific PYTHON stack (Jones et al. 2001; Hunter 2007; van der Walt, Colbert & Varoquaux 2011).

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author.

REFERENCES

- Abazajian K. N. et al., 2015, *Astropart. Phys.*, 63, 66
 Aiola S. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 047

- Bender A. N. et al., 2018, in Zmuidzinas J., Gao J.-R., eds, Proc. SPIE Conf. Ser. Vol. 10708, Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX. SPIE, Bellingham, p. 1070803
- BICEP2/Keck Collaboration, 2015, *Phys. Rev. Lett.*, 114, 101301
- Chartier N., Wandelt B., Akrami Y., Villaescusa-Navarro F., 2021, *MNRAS*, 503, 1897
- Choi S. K. et al., 2020, *J. Cosmol. Astropart. Phys.*, 2020, 045
- Crites A. T. et al., 2015, *ApJ*, 805, 36
- Das S. et al., 2014, *J. Cosmol. Astropart. Phys.*, 2014, 014
- Dodelson S., Schneider M. D., 2013, *Phys. Rev. D*, 88, 063537
- Dodelson S. et al., 2013, preprint ([arXiv:1309.5386](https://arxiv.org/abs/1309.5386))
- Dutcher D. et al., 2021, *Phys. Rev. D*, 104, 022003
- Fendt W. A., Wandelt B. D., 2007, *ApJ*, 654, 2
- Friedrich O., Seitz S., Eifler T. F., Gruen D., 2016, *MNRAS*, 456, 2662
- Hall A., Taylor A., 2019, *MNRAS*, 483, 189
- Hamimeche S., Lewis A., 2009, *Phys. Rev. D*, 79, 083012
- Hartlap J., Simon P., Schneider P., 2007, *A&A*, 464, 399
- Heavens A. F., Sellentin E., de Mijolla D., Vianello A., 2017, *MNRAS*, 472, 4244
- Henning J. W. et al., 2018a, *ApJ*, 852, 97
- Hivon E., Górski K. M., Netterfield C. B., Crill B. P., Prunet S., Hansen F., 2002, *ApJ*, 567, 2
- Hunter J. D., 2007, *Comput. Sci. Eng.*, 9, 90
- Jones E., Oliphant T., Peterson P., 2001, SciPy: Open source scientific tools for Python. Available at: <http://www.scipy.org/>
- Lewis A., Bridle S., 2002, *Phys. Rev. D*, 66, 103511
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, 538, 473
- Louis T. et al., 2017a, *J. Cosmol. Astropart. Phys.*, 6, 031
- Louis T. et al., 2017b, *J. Cosmol. Astropart. Phys.*, 2017, 031
- Lueker M. et al., 2010a, *ApJ*, 719, 1045
- Morrison C. B., Schneider M. D., 2013, *J. Cosmol. Astropart. Phys.*, 2013, 009
- Omori Y. et al., 2019, *Phys. Rev. D*, 100, 043517
- Percival W. J., Friedrich O., Sellentin E., Heavens A., 2021, *MNRAS*, 510, 3207
- Planck Collaboration I, 2020a, *A&A*, 641, A1
- Planck Collaboration V, 2020b, *A&A*, 641, A5
- Planck Collaboration VI, 2020d, *A&A*, 641, A6
- Polarbear Collaboration, 2014, *ApJ*, 794, 171
- Reichardt C. L. et al., 2021, *ApJ*, 908, 199
- Schneider M. D., Knox L., Habib S., Heitmann K., Higdon D., Nakhleh C., 2008, *Phys. Rev. D*, 78, 063529
- Schneider M. D., Cole S., Frenk C. S., Szapudi I., 2011, *ApJ*, 737, 11
- Sellentin E., Heavens A. F., 2016, *MNRAS*, 456, L132
- Simons Observatory Collaboration, 2019, *J. Cosmol. Astropart. Phys.*, 2019, 056
- Spurio Mancini A., Piras D., Alsing J., Joachimi B., Hobson M. P., 2021, *MNRAS*, 511, 1771
- Taylor A., Joachimi B., Kitching T., 2013, *MNRAS*, 432, 1928
- The BICEP/Keck Collaboration, 2018, preprint ([arXiv:1807.02199](https://arxiv.org/abs/1807.02199))
- Torrado J., Lewis A., 2021, *J. Cosmol. Astropart. Phys.*, 2021, 057
- Tristram M., Macías-Pérez J. F., Renault C., Santos D., 2005, *MNRAS*, 358, 833
- van der Walt S., Colbert S., Varoquaux G., 2011, *Comput. Sci. Eng.*, 13, 22

APPENDIX A

We present the uncertainty underestimation and shift to the best-fitting point from joint $TT + TE + EE$ constraints in Table A1. The results cover diagonal, moderate, and deep conditioning for $N_{\text{real}} = 30, 50, 100$, and $\sigma_{\text{map}} = (10, 6.4, 2) \mu\text{K arcmin}$. In Table A2, we show the same statistics for all combinations of the TT, TE , and EE spectra of the $N_{\text{real}} = 30, \sigma_{\text{map}} = 10 \mu\text{K arcmin}$ simulations.

Table A1. Median and 68 per cent interval of the uncertainty underestimation in per cent compared to the performance of the fiducial matrix and shift to the best-fitting point of θ_{MC} constraints scaled by the predicted uncertainty from a Fisher forecast for simulations with $N_{\text{real}} = 30, 50, 100$ and $\sigma_{\text{map}} = (10, 6.4, 2) \mu\text{K arcmin}$ for diagonal, moderate, and deep covariance conditioning.

Noise – level (μ – arcmin)	30			N_{real} 50			100		
	Diagonal	Moderate	Deep	Diagonal	Moderate	Deep	Diagonal	Moderate	Deep
Uncertainty underestimation: $1 - \sigma/\sigma_0$ (per cent)									
10	$5.8^{+2.1}_{-1.9}$	$3.5^{+2.7}_{-2.3}$	$0.3^{+1.5}_{-1.5}$	$4.4^{+1.8}_{-1.6}$	$2.1^{+2.1}_{-1.9}$	$0.2^{+1.6}_{-1.6}$	$2.4^{+1.5}_{-1.2}$	$0.0^{+1.5}_{-1.4}$	$-0.7^{+1.2}_{-1.4}$
6	$4.5^{+1.8}_{-1.6}$	$2.1^{+2.1}_{-1.9}$	$0.0^{+1.3}_{-1.3}$	$3.5^{+1.4}_{-1.2}$	$0.9^{+1.6}_{-1.4}$	$-0.2^{+1.3}_{-1.3}$	$3.4^{+1.0}_{-1.3}$	$0.6^{+1.5}_{-1.3}$	$-0.2^{+1.5}_{-0.9}$
2	$3.3^{+1.1}_{-1.0}$	$0.5^{+1.4}_{-1.2}$	$-0.3^{+1.1}_{-1.1}$	$3.3^{+0.9}_{-1.0}$	$0.5^{+1.1}_{-1.1}$	$-0.1^{+1.1}_{-1.1}$	$2.8^{+0.9}_{-1.1}$	$0.1^{+1.2}_{-1.3}$	$-0.4^{+1.0}_{-1.2}$
Best-fitting shift: $\Delta\theta/\sigma_{\mathcal{F}}$									
10	$0.16^{+0.45}_{-0.44}$	$0.17^{+0.46}_{-0.47}$	$0.00^{+0.10}_{-0.11}$	$0.17^{+0.25}_{-0.26}$	$0.16^{+0.25}_{-0.27}$	$0.03^{+0.11}_{-0.10}$	$0.13^{+0.22}_{-0.17}$	$0.13^{+0.24}_{-0.17}$	$0.00^{+0.14}_{-0.13}$
6	$0.14^{+0.31}_{-0.31}$	$0.13^{+0.31}_{-0.35}$	$0.02^{+0.11}_{-0.11}$	$0.12^{+0.24}_{-0.21}$	$0.11^{+0.24}_{-0.20}$	$0.01^{+0.11}_{-0.12}$	$0.07^{+0.15}_{-0.15}$	$0.07^{+0.16}_{-0.15}$	$-0.01^{+0.12}_{-0.13}$
2	$0.09^{+0.16}_{-0.18}$	$0.08^{+0.16}_{-0.18}$	$0.02^{+0.13}_{-0.12}$	$0.08^{+0.14}_{-0.12}$	$0.08^{+0.14}_{-0.13}$	$0.02^{+0.12}_{-0.12}$	$0.11^{+0.14}_{-0.14}$	$0.11^{+0.15}_{-0.15}$	$0.04^{+0.17}_{-0.15}$

Table A2. Median and 68 per cent interval of the uncertainty underestimation in per cent compared to the fiducial matrix and shift to the best-fitting point of θ_{MC} constraints scaled by the predicted uncertainty from a Fisher forecast for the $N_{\text{real}} = 30$, $\sigma_{\text{map}} = 10 \mu\text{K}$ arcmin simulations for diagonal, moderate, and deep covariance conditioning for any combination of the TT , TE , and EE power spectra.

Spectra used in likelihood	Uncertainty underestimation: $1 - \sigma/\sigma_0$ (per cent)			Best-fitting shift: $\Delta\theta/\sigma_{\mathcal{F}}$		
	Diagonal	Moderate	Deep	Diagonal	Moderate	Deep
TT, EE, TE	$5.8^{+2.1}_{-1.9}$	$3.5^{+2.7}_{-2.3}$	$0.3^{+1.5}_{-1.5}$	$0.16^{+0.45}_{-0.44}$	$0.17^{+0.46}_{-0.47}$	$0.00^{+0.10}_{-0.11}$
TT	$4.8^{+2.2}_{-2.1}$	$1.3^{+2.7}_{-2.5}$	$0.6^{+2.0}_{-2.1}$	$0.07^{+0.15}_{-0.15}$	$0.07^{+0.15}_{-0.16}$	$0.00^{+0.07}_{-0.07}$
TE	$2.8^{+1.9}_{-2.0}$	$1.4^{+2.4}_{-2.5}$	$0.4^{+2.1}_{-2.0}$	$0.02^{+0.15}_{-0.15}$	$0.01^{+0.15}_{-0.15}$	$0.00^{+0.06}_{-0.07}$
EE	$3.9^{+3.3}_{-3.1}$	$3.2^{+4.6}_{-4.3}$	$0.4^{+2.4}_{-2.5}$	$0.10^{+0.27}_{-0.30}$	$0.10^{+0.30}_{-0.34}$	$0.00^{+0.07}_{-0.08}$
TT, TE	$4.7^{+1.9}_{-1.8}$	$2.1^{+2.1}_{-2.1}$	$0.4^{+1.8}_{-1.6}$	$0.10^{+0.23}_{-0.23}$	$0.09^{+0.23}_{-0.23}$	$0.00^{+0.09}_{-0.08}$
TT, EE	$5.4^{+2.4}_{-2.3}$	$3.1^{+3.0}_{-2.7}$	$0.5^{+1.7}_{-1.8}$	$0.10^{+0.31}_{-0.41}$	$0.13^{+0.36}_{-0.43}$	$0.00^{+0.08}_{-0.09}$
EE, TE	$4.1^{+2.4}_{-2.2}$	$2.9^{+3.1}_{-2.8}$	$0.3^{+1.8}_{-1.8}$	$0.09^{+0.31}_{-0.32}$	$0.09^{+0.32}_{-0.35}$	$0.02^{+0.10}_{-0.09}$

This paper has been typeset from a \LaTeX file prepared by the author.