Simulating and Analyzing Polarization Angle Systematics for CMB-S4 r Measurement

Ben Eichner
Advisor: Dr. Colin Bischoff
University of Cincinnati

April 25, 2025

1 Introduction

The Cosmic Microwave Background (CMB) is residual light left over from the early expansionary era of our universe. Encoded within it is information that reveals the dynamics of the universe's earliest stages. The Lambda-Cold Dark Matter (ΛCDM) model of the universe predicts that the universe underwent a brief period of extremely rapid expansion, also called inflation, shortly after the Big Bang. The ΛCDM model has been generally successful in explaining much of the universe's large-scale behavior. To date, however, there remains no observed evidence of inflation.

Inflation, if it occurred, would have produced two kinds of perturbations in the pre-last scattering universe: Density waves and gravitational waves [8]. The tensor-to-scalar ratio r represents the amplitude of the "primordial" gravitational waves generated by inflation and serves as a key parameter in all models of inflation; the particular model of inflation used has many important consequences for all of astrophysics and particle physics. Certain models are only possible if r is within a particular range of values, so valid models can be constrained by making precise measurements of the value of r. One way of accomplishing this is by making observations of the polarization of CMB light.

The waves produced by inflation would lead to anisotropies in the CMB signal, in particular, density waves would produce a certain pattern of polarized light at last scattering known as "E-modes"; these E-modes are readily visible and constitute the vast majority of the CMB signal across the entire sky. Primordial gravitational waves would produce a different pattern in the CMB signal known as "B-modes". In particular, primordial gravitational waves are the only possible source of B-modes in the CMB, should they exist. An observation of a B-mode signal of nearly any strength that is confirmed to be of CMB origin would therefore be evidence that primordial gravitational waves exist, and by extension, evidence that inflation occurred. The exact strength of the CMB B-mode signal, relative to the E-mode signal, would depend on the amplitude of primordial gravitational waves. So, measurements of the B-mode signal serve as an indirect measurement of r, which itself can be calculated from the power spectrum of the B-mode signal across the CMB. Thus far, observational efforts have placed an upper-bound on r of 0.037 [5] . This has constrained many models of inflation, but many more are valid for smaller r [1]. Lowering this upper-bound further will require even more precise measurements of the polarization of the CMB.

The CMB-S4 project is an ongoing scientific collaboration to conduct precision observations of the CMB using a new generation of small aperture telescopes designed to be ultra-sensitive to polarization. The most pressing challenge faced by these new designs is the potential for systematic errors to affect observational precision. Systematic errors in the polarization detectors, for example, misalignment of these detectors from the expected angle, may lead to biased or imprecise measurements of r, preventing a better upper-bound from being achieved. The current CMB-S4 target for the standard deviation (1-sigma) of the measurement of the upper-bound r value is 5×10^{-4} [1].

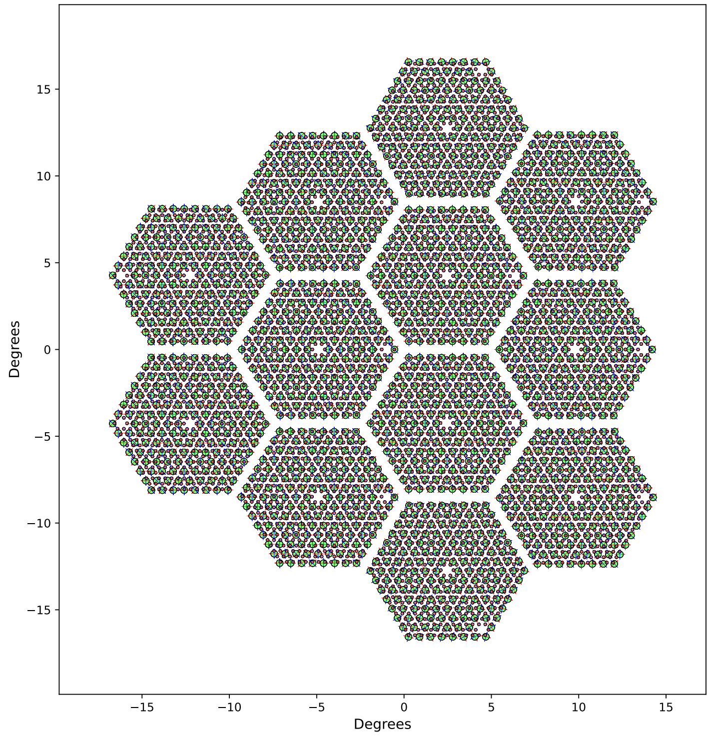

Figure 1: Design of the instrument model used for this simulation. One wafer is simply one of the shown hexagonal formations. Credit: Sara Simon, CMB-S4

The goal of the research described within this paper is to establish the magnitude to which polarization systematics affect the precision of measurements of r taken by a proposed instrument model for CMB-S4, whether they create bias, and how large that bias is. This is achieved by creating a simulation of observations of the CMB affected by polarization systematics inside a small aperture telescope capable of generating realistic maps of the sky as seen by the telescope with randomized systematics. The small aperture instrument being simulated is shown in Figure 1; it has nine optic tubes for observing polarized light at eight frequency bandpasses, each tube has 12 wafers corresponding to it. These simulated maps then have their power spectra calculated and input into a model-search function to calculate the r value predicted by the simulated data. Finally, determination of bias will be made by comparing the r value results for many systematic affected simulations to their non-systematic affected counterparts. Each systematic affected simulation will have a corresponding unaffected simulation that was generated with the same input maps, allowing direct comparison of the shift in r on a simulation-by-simulation basis.

2 Methods And Activities

2.1 Map Generation

The simulations assume noise, dust, and synchrotron foreground emission are present, no frequency bandpass systematics, and no atmospheric interference. The theory power spectrum used to generate the CMB maps includes bandpowers from gravitational lensing foregrounds, but the simulation does not attempt to perform delensing analysis on the data. This will increase the standard deviation of r significantly. This is not considered problematic: Our analysis is focused on the bias on r between systematic and non-systematic affected maps inside of each realization, not the r values themselves. Thus, the shifts of r are the object of interest produced by these simulations, and a delensing analysis will not change the magnitude of those shifts.

Four sky map types are required for each simulation realization to produce a realistic observation map: Foreground maps, CMB maps, noise maps, and inverse variance maps. Additionally, CMB-S4 plans to make observations of the CMB at eight frequency bandpasses identified as ideal for data collection; these bandpasses occur around 30 GHz, 40 GHz, 85 GHz, 95 GHz, 145 GHz, 155 GHz, 220 GHz, and 270 GHz [1]. Each of these bandpasses have different associated noise and beam size parameters, meaning that a new set of sky maps must be generated for each separate bandpass. The resolution of these maps is given in a parameter called NSIDE, which originates from HEALPix, a method for mapping 3D spherical maps to 2D square pixels [6]. The NSIDE used for the simulations was 512, corresponding to approximately 3.1 million pixels per sky map. All maps generated consist of temperature (T), Q-polarization (Q), and U-polarization (U) components. The Q and U polarization components are the only maps of real concern, the T maps are kept but not used after generation as they are not necessary for the analysis.

Foreground maps are generated first using the PySM 3 (Python Sky Model 3) package and stored in a separate array that is accessed by each realization [11]. Every realization will use these same foreground maps. After this is done, the parameters of the systematic are chosen, which will be covered in the next section.

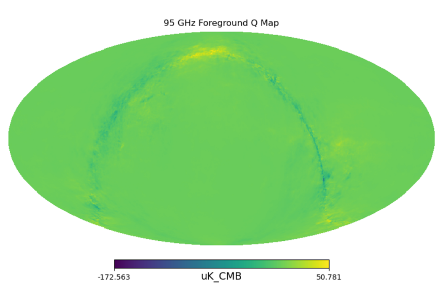

Figure 2: Foreground Map at 95GHz, Q polarization. Note the characteristic galactic band across the sky.

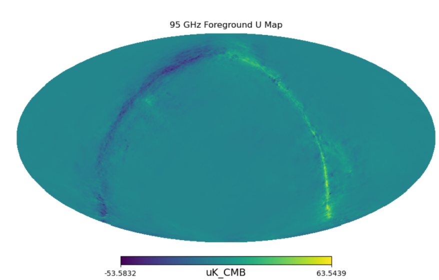

Figure 3: Foreground Map at 95GHz, U polarization. Note the characteristic galactic band across the sky.

Next, a CMB map is generated via the HealPy synfast function using theory spectra from the Planck FFP10 simulation suite [9][10]. These spectra represent the intensity of fluctuation of values of Temperature, Q polarizaton, and U polarization above or below the dominant values of these measurables observed across the CMB. This map is copied eight times, each copy is then modified with the appropriate beam size and noise level of one frequency bandpass, so that eight CMB maps corresponding to simulated observations of the CMB only at each bandpass are generated. These are saved at the end of the generation process.

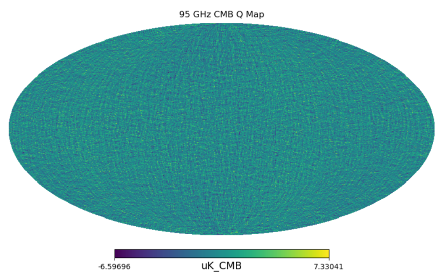

Figure 4: CMB Map at 95 GHz, Q polarization.

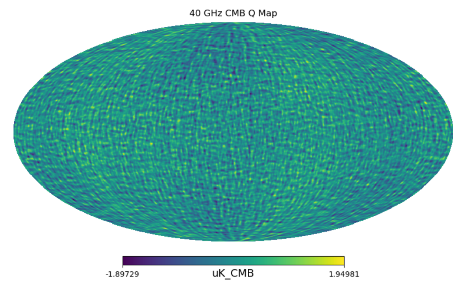

Figure 5: CMB Map at 40 GHz, Q polarization.

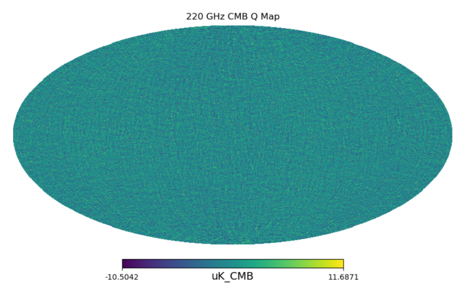

Figure 6: CMB Map at 220 GHz, Q polarization

Sky noise maps are generated in a similar manner using HealPy synfast albeit using a theoretical power spectrum for noise modeling as input:

$$C_\ell = A(1 + (\frac{\ell}{\ell_{knee}})^\alpha) \tag{1}$$

Where A is the noise amplitude, ℓ is the multipole, ℓ_{knee} is the noise knee frequency, and α is the spectral index of the noise. These parameters, with the exception of ℓ, are set prior to the start of the simulation and depend on the frequency bandpass. One initial map is generated from synfast per realization, then copied and adjusted based on bandpass parameters to produce eight bandpass-specific noise maps. At this point, inverse variance maps are introduced and applied to the noise maps, CMB maps, and foreground maps. Like the foregrounds, these are also created and saved before the simulation begins.

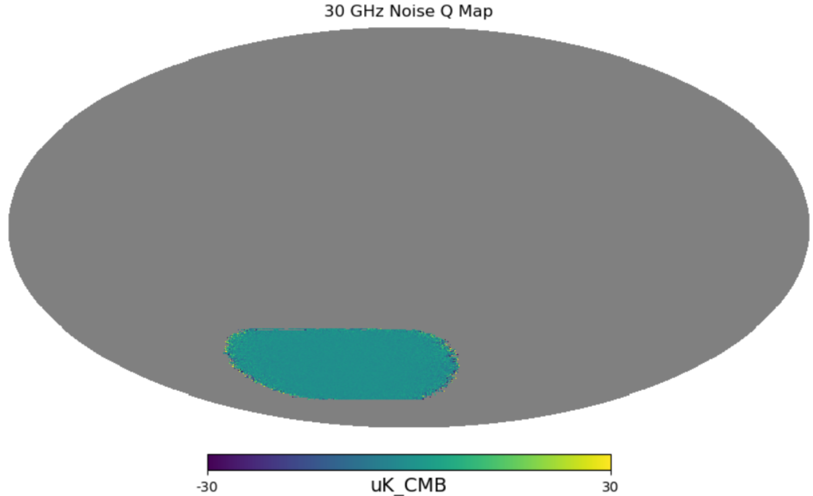

Figure 7: Noise map at 30 GHz, Q polarization.

Inverse variance maps are imported from a separate simulation of the instrument model's observing area of the sky. That simulation emulated a ten-year observation of a particular observing field by each wafer on the instrument model, and from this data generated the inverse variance maps for each wafer. The observing field simulated and used in this simulation corresponds to a previously-identified field ideal for CMB-S4 observation at the South Pole [1]. These maps, when applied to the generated noise, CMB, and foregrounds, weight the generated maps by the inverse of the variance of data collected at that pixel, that is, the instrument noise. Less time is spent by the instrument observing sky near the edge of the field, this leads to high noise, so these pixels are downweighted, while near the center, measurement is more consistent, so these pixels have higher weight. Effectively, inverse variance maps represent the area that each wafer of the instrument model can "see", and to what degree each portion of that area has been observed.

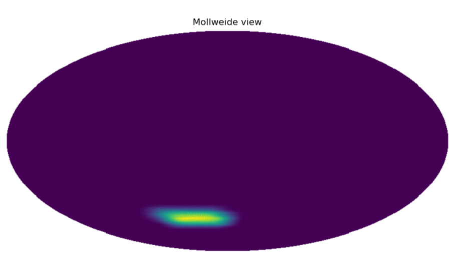

Figure 8: Inverse variance map of 95 GHz wafer near the center of the instrument. Weights are dimensionless, strength of weight indicated by heatmap.

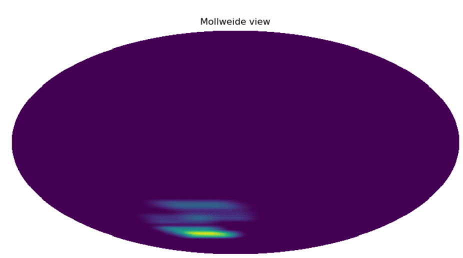

Figure 9: Inverse variance map of 95 GHz wafer near the edge of the instrument. Weights are dimensionless, strength of weight indicated by heatmap.

Once inverse variance maps have been applied to each of the generated map types, noise, CMB, and foreground maps are returned to units of microkelvin CMB by dividing out the inverse variance weights according to the following formula:

$$M^{FB} = \frac{M^{FB}_{weighted}}{\sum_{i=0}^{W} \sigma_i^{-2}} \tag{2}$$

Where M represents the T, Q, or U map of that frequency band, W is the number of wafers in that band, and σ_i^{-2} is the inverse variance. The maps are then combined together per frequency bandpass into a single map representing the simulated sky observation of the instrument in that band. The polarization component maps (Q, U) are saved.

2.2 Systematic Simulation

Polarization systematics occur as a small unaccounted rotation of the instrument polarization detectors off of the expected alignment. This causes the instrument to incorrectly read some amount of incoming Q polarized light as U and vice-versa, which is expected to cause E-mode to B-mode leakage further down the analysis pipeline. The instrument model consists of nine collection tubes with twelve wafers apiece, each tuned to one of the eight frequency bands. These wafers are constructed separately then put together during final assembly, so the simulation treats each wafer as a separate object with a possible randomized misalignment. The possible misalignment of individual cells is considered to be much smaller than the errors that might occur during assembly, so these are neglected. At the start of every simulation, the standard deviation and mean of the random misalignment is specified in degrees. These quantities are used in the numpy.random function to generate a one-dimensional array of random numbers for each frequency bandpass, with a length corresponding to the number of wafers in that bandpass. Each realization of the simulation regenerates this array, ensuring no two realizations have the same exact systematics applied to them. For the simulations analyzed as part of the CMB-S4 project, our choice was to use a mean of 0 degrees and standard deviations of 1 degree and 5 degrees.

The systematic itself is simulated as a randomized rotation of the generated CMB and foreground maps. Noise maps do not have a systematic applied to them as this would not create any meaningful difference in the observed data due to its fully randomized nature. Each Q and U polarization map of the CMB and foreground maps for each frequency bandpass are modified according to the following function:

$$\begin{aligned} Q^{FB}_{syst} &= Q^{FB} cos(2\psi_i) + U^{FB} sin(2\psi_i) \\ U^{FB}_{syst} &= -Q^{FB} sin(2\psi_i) + U^{FB} cos(2\psi_i) \end{aligned}$$

Where ψ_i is the angle of the systematic converted into radians, and i is the index of the entry in that frequency bandpasses's random number array, which corresponds to the wafer that this map will be applied to. This is done separately for each map type, so these maps are summed together afterward with the same noise maps as the non-systematic affected maps to create a new set of simulated sky maps with polarization systematics. Once again, the polarization component maps are saved. This concludes the map generation step of the simulation.

2.3 Power Spectra

Every T, Q, and U observational map of the sky can be decomposed into their respective fields' multipole moments derived from spherical harmonics, which can then be represented by their power spectrum [8]. These power spectra are purely functions of angular multipoles ℓ with the intensity of fluctuations represented by the coefficient C_ℓ, which describes the conjugate square of a particular multipole moment.

$$\delta_{\ell\ell'}\delta_{mm'}C_\ell = \langle a^*_{\ell m} a_{\ell m}\rangle \tag{3}$$

Multipole moments of T, E, and B fields fluctuate according to a Gaussian distribution, but their square does not; this fact will become important during likelihood analysis. Since E-mode and B-mode fields are typically used in CMB analysis, Q and U maps are linearly transformed into E-mode and B-mode power spectra [8]. In this simulation, the Pseudo-C_ℓ framework as implemented by NaMaster is used for bandpower estimation [3].

This calculation is performed using the NaMaster package in Python, with additional help from the prototype S4BB package to streamline the calculation process of power spectra for many realizations [4]. Primordial gravitational waves are predicted to affect B-mode power spectra at the $\ell = 10^2$ level, or around degree-scale on the sky, so the power spectra of the generated sky maps are only calculated for ℓ in the range of 30 to 310 [8]. Noise tends to dominate at greater ℓ, which represent smaller angular scales on the sky. See below.

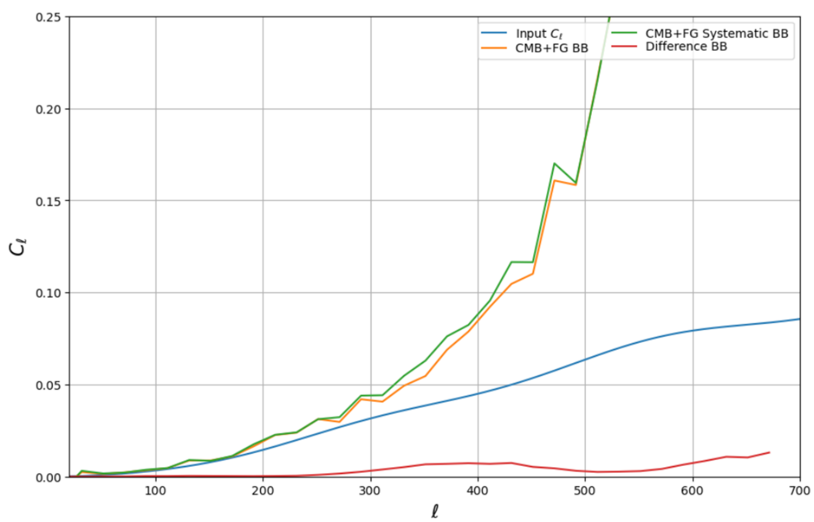

Figure 10: Comparison of B-mode auto-spectra for the 95 GHz frequency band. Plotted are the observation map spectra for non-systematic and systematic simulations, as well as the theory spectrum input and the spectrum of their difference map. Note the exponential increase in bandpowers toward greater ℓ values.

The simulation performs calculations of E-mode and B-mode power spectra on noise, CMB, and final observation maps and stores them as unique objects using the S4BB package framework; these objects contain information about the map type, spectra, and number of realizations stored. Each spectral object contains the bandpower of 36 spectra, each bandpower is grouped into 14 bins comprising a range of 20 ℓ values per realization, representing the auto- and cross-spectra of each frequency band. Noise in one frequency bandpass is generally assumed to be independent of others, so cross-spectra of noise maps are set to 0.

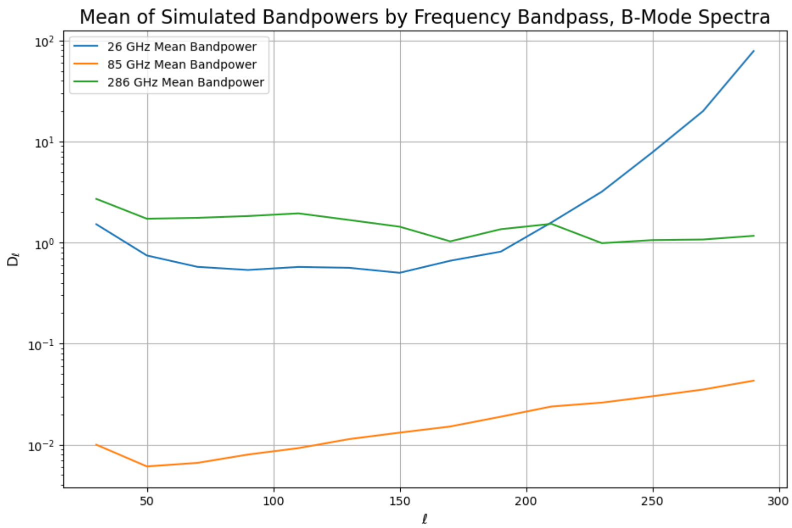

Figure 11: Mean of simulated bandpowers across 45 realizations, no systematics applied. Bandpowers in terms of D_ℓ.

The spectra for the final observation maps represent the data that will be fit by a model to calculate r. Noise and CMB spectra are calculated separately for use in the creation of the bandpower covariance matrix, which is necessary for error estimation and the maximum likelihood search function that will be used to fit models to the data. Noise spectra are also used to calculate B-mode noise bias, which is defined as the mean of the B-mode noise power spectra, per frequency band, across all realizations. The noise bias of one ℓ bin in one frequency band is:

$$N_\ell^{BB} = \frac{1}{S}\sum_{i=0}^{S} N_{\ell i}^{BB} \tag{4}$$

Where S is the number of realizations, and $N_{\ell i}^{BB}$ is the noise bandpower of simulation i in a particular ℓ bin. This bias is used by the likelihood search function to simulate a retrieval of true CMB signal bandpowers from a noise-affected data set such as the simulated final observation maps.

This concludes the calculation of power spectra from generated maps. All spectral objects are saved and used as inputs for likelihood analysis.

2.4 Maximum Likelihood Search and r Values

The calculation of r for each realization is accomplished through maximum likelihood search of the simulated data. This consists of two parts: calculation of a bandpower covariance matrix, and the likelihood analysis itself.

The spectral objects containing noise spectra and CMB spectra are used to calculate the bandpower covariance matrix, whose diagonal elements measure the variance of each ℓ-bin bandpower in auto- and cross-spectra for all frequency bands. While bandpowers are measured from a non-Gaussian distribution, a Gaussian analogy is useful to understand the importance of the matrix. In Gaussian statistics, chi-squared measures the difference between the model and expected data points as compared to the expected variance of the data point:

$$\chi^2 = \sum_{i=0}^{N} \frac{(d_i - m_i)^2}{\sigma_i^2} \tag{5}$$

Since the bandpower covariance matrix provides the expected variance of the data point–in our simulation, this is the bandpower of a particular auto- or cross-spectrum in a certain ℓ-bin–if the covariance matrix is only provided with the bandpowers of one realization, then it is possible for some of the bandpower values of that realization to be very close to the model value, returning a low variance. A particularly low variance will greatly increase the contribution that data point has to the value of chi-squared, possibly well beyond that of every other data point. Since likelihood maximization in Gaussian statistics is effectively the minimization of chi-squared, any maximum likelihood search algorithm will adjust the model parameters to best fit the low-variance data point over all others, since this has the largest impact on chi-squared. This statistical bias can be remedied by providing the matrix with more bandpowers from more realizations, decreasing the chance that an uncharacteristic distribution of data values are taken. This necessitates a high number of realizations, about 50-100, for the simulation to create a reasonably accurate bandpower covariance matrix.

Next, the likelihood analysis takes 14-parameters for two separate models of dust and synchrotron, and CMB emission to calculate expectation values of bandpowers. It also takes the noise bias and the B-mode power spectra of simulated data, which it uses to estimate the true B-mode power spectra of the CMB observed in that realization. The model used for CMB signal is:

$$D_\ell = rD_\ell^{tens} + A_{lens}D_\ell^{lens} \tag{6}$$

Where D_ℓ represents the total bandpower observed, D_ℓ^{tens} is the bandpower contribution from tensor perturbations (gravitational waves), and D_ℓ^{lens} is the bandpower contribution from lensed modes. A_{lens} is a fixed parameter set to a value of 1. The model for dust/synchrotron emission in foregrounds is provided from a 2018 BICEP/Keck collaboration paper[2]. As mentioned previously, bandpowers are drawn from a non-Gaussian distribution. This is because the simulation is measuring bandpowers in the low-ℓ range within a relatively small fraction of the sky; if the observations were full-sky, the distribution would be Gaussian. Thus the maximum likelihood search function utilizes the Hamamiche-Lewis likelihood estimation function which finds $-2log(\mathcal{L})$for non-Gaussian data [7]. Noting that

$$-2log(\mathcal{L}) = log(\frac{1}{\mathcal{L}^2}) \tag{7}$$

The maximum likelihood search seeks to minimize the Hamamiche-Lewis likelihood, effectively maximizing the value of the likelihood squared. If successful, the maximum likelihood function returns the values of the 14 parameters found for the best-fit model per realization, including r. This search is performed twice, once on the non-systematic affected data, and once on the systematic-affected data, realization-by-realization. This allows direct comparison of the shift in r values inside one realization between systematic data and unaltered data. All final parameter values are saved in numerical arrays.

3 Results

Analysis was performed on 45 simulations of a 5 degree systematic and 80 simulations of a 1 degree systematic, each with a mean of 0 degrees. Data on the shift in r values from systematic to non-systematic realizations is shown below.

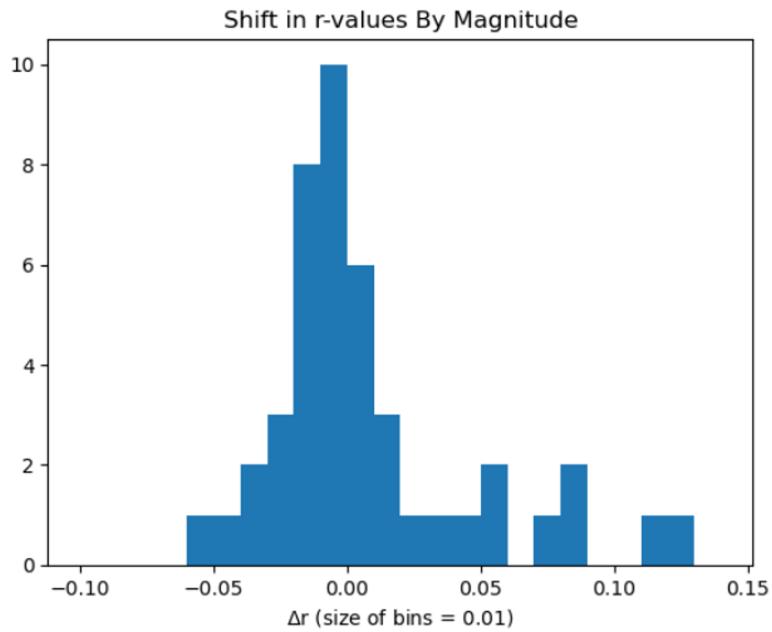

Figure 12: Shift in r by magnitude, histogram of 45 realizations. Size of bins = 0.01.

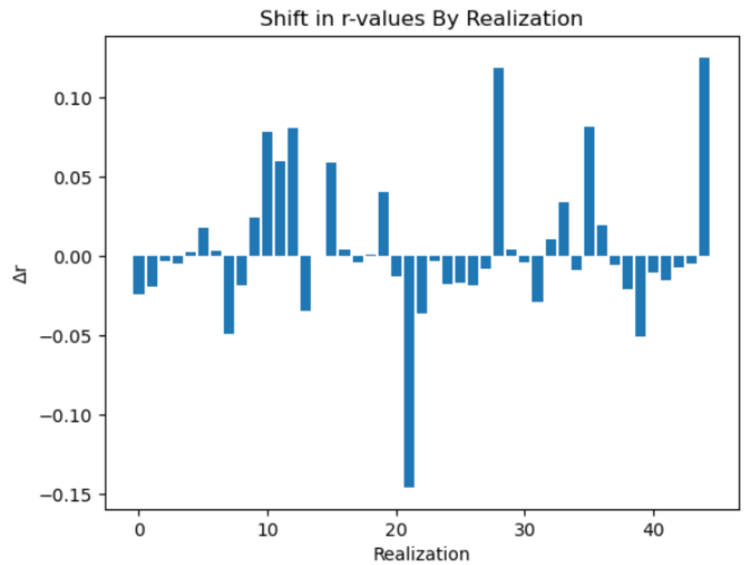

Figure 13: Bar chart of shift in r by realization, 45 realizations.

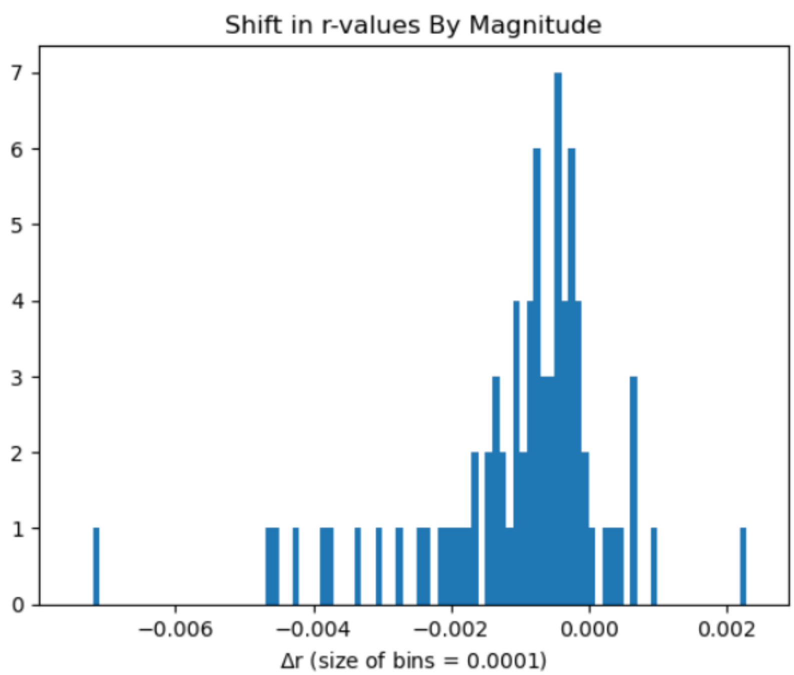

Figure 14: Shift in r by magnitude, histogram of 80 realizations. Size of bins = 0.0001.

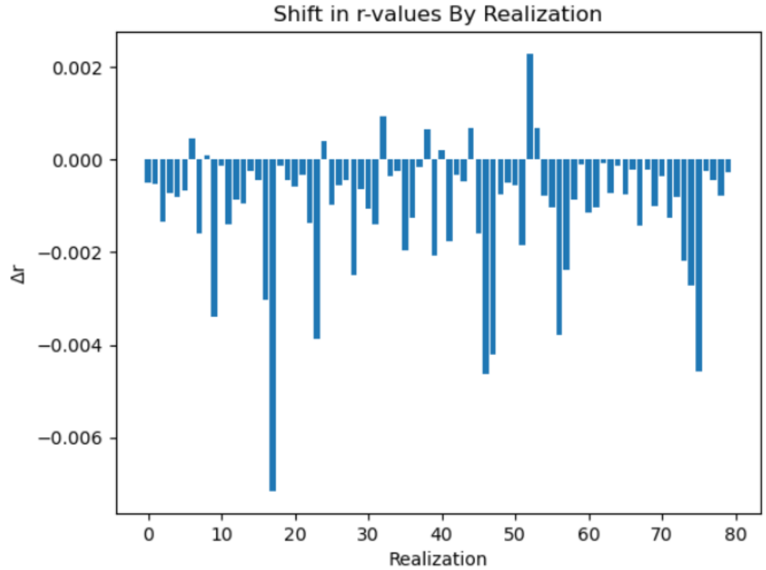

Figure 15: Bar chart of shift in r by realization, 80 realizations.

# of Realizations	Systematic	Mean	$\sigma(r)$
45	OFF	0.0026	0.00368
45	ON; 5 degrees, mean 0 degrees	0.007	0.0451
80	OFF	0.00345	0.00346
80	ON; 1 degrees, mean 0 degrees	0.0024	0.0037

Table 1: Mean and standard deviation of measured r value distributions for 45 and 80 realizations.

# of Realizations	Mean Δr	$\sigma(\Delta r)$
45	0.00436	0.045
80	-0.0011	0.0014

Table 2: Mean and standard deviation of shift in r value for 45 and 80 realizations.

A positive shift value indicates that the systematic best-fit r value is higher than the non-systematic best-fit value for that realization, while a negative value indicates the opposite. In the 45 simulation run, there appears to be no general direction in the shifts; the primary difference between the data sets is the broadening of the standard deviation. Notably, the non-systematic standard deviation of r values is 7.2 times greater than the CMB-S4 science target for $\sigma(r)$, while the systematic standard deviation is almost 90 times. This would seem to indicate that a random set of systematics with standard deviation of five degrees of misalignment significantly impacts the precision of the instrument. However, the 80 simulation run has a clear negative direction in nearly all of its shifts, and the standard deviation of the systematic affected r values is only slightly larger than the non-systematic standard deviation.

Obviously, a smaller systematic will lead to smaller shifts in r, which is reflected in the distribution of shift values seen in the figures, but the standard deviation of the systematic distributions with respect to the non-systematic ones should be similar. It is possible that the 45 simulation run, being slightly below the range of realizations needed to construct a proper bandpower covariance matrix, happened to have a bad distribution of bandpowers which affected the maximum likelihood search function. Why this would cause a significantly broader distribution of r is unknown, and would require more analysis. Of greater interest are the results from the 80 simulation run, as the systematic of 1 degree is currently the expected level of polarization systematics inside this instrument model.

First, given the directionality of the shifts in the 80 simulation run, it is more useful to look at the mean of the shift of r than the difference in standard deviations between distributions, as both are quite similar. The reason for this is simple: If there is an expected or desired standard deviation in the data, say 5×10^{-4}, and in a real observation a value is found that is far outside of the range of values covered by 1-sigma, then that increases the error of the measurement to at least that value. To get precise values, then, shifts due to systematics should be kept at a fraction of the desired error; shifts that greatly exceed the desired error increase error beyond the desired level. Returning to the data, the mean of the shifts is -1.1×10^{-3}, which is 2.2 times the CMB-S4 science target. This is decidedly excessive: According to this data, systematics of even 1 degree standard deviation will prevent the science target from being reached. However, it should be emphasized that the results are very recent and have not undergone a full analysis to confirm that this is a consistent effect due solely to polarization systematics.

The clear negative bias in these results is very interesting, if confirmed. In a real experiment, negative r is impossible and has no physical meaning; recall that r is a measurement of the amplitude of perturbations, which cannot be negative. However, the simulation's models implicitly assume it as a possibility as this simplifies the analysis of bias in the data. If realistic steps were taken to ensure only positive r were generated, the distribution of r seen in Figures 11 and 13 would be decidedly non-normal, encountering a spike towards infinity as r approaches 0. This makes it difficult to calculate a meaningful mean and of course a value for standard deviation. By allowing negative r, our simulation produces a more normal distribution which can be analyzed with these values.

That still begs the question of what a negative bias in the simulated r means. Negative values of r are made possible by the general form of the model of the data assumed by the simulation. This model is used to simulate a realistic calculation of true signal bandpowers from noisy data, such as in the simulated observational maps that combine CMB, foregrounds, and noise. Noting from (2) that we can replace the expression on the right hand side with the values that make up the power spectrum in this model, the noise bandpowers and signal bandpowers:

$$\langle a^*_{\ell m} a_{\ell m} \rangle = \frac{2}{k}(S_\ell + N_\ell)^2 \tag{8}$$

Where k is the number of modes, S_ℓ is the signal bandpower in a ℓ bin, and N_ℓ is the noise bandpower in a ℓ bin; in the simulations this is calculated as the mean of the B-mode noise spectra, called noise bias (see section 2.3). So, then:

$$C_\ell = \frac{2}{k}(S_\ell + N_\ell)^2 \tag{9}$$

And

$$S_\ell = \sqrt{\frac{k}{2}C_\ell} - N_\ell \tag{10}$$

Therefore, if the noise bias exceeds the majority of simulated bandpowers in any given realization, it leads to negative signal values, which themselves could lead to negative values of r, as these are calculated by fitting models to the calculated B-mode signal spectra. In effect, this means that negative r indicate a weaker primordial B-mode signal that is hidden inside the noise of the observation. In a real experiment, these would show up as a measurement of $r = 0.00$. If the systematic simulations have a consistent negative bias, then, that means the polarization systematics decrease the likelihood of a positive detection of primordial B-modes, or similarly, increase the likelihood of a false negative: Detection of no primordial B-modes in a CMB where they do exist.

This is interesting primarily for one reason. As mentioned, the E-mode signal dominates the CMB, and to date measurements have concluded that it must be at least one hundred times stronger than any B-mode signal pattern. Earlier it was hypothesized (see section 2.2) that if systematics are present, this should cause misaligned sensors to read some of the much stronger E-mode signal as B-mode signal; a leakage of E-modes to B-modes. In short, systematic affected simulations should consistently find much stronger B-mode signals, and consequently higher r, than their non-systematic counterparts. But our results seem to indicate the opposite is happening, that polarization systematics are causing the instrument to consistently observe weaker B-mode signals or no B-mode signals.

If these results are confirmed, the leakage of E-modes to B-modes may be far weaker than first suspected, perhaps on the level of the reciprocal possibility of B-modes leaking into E-modes due to the systematics. This leakage was assumed to be negligible compared to its counterpart owing to the tiny magnitude of the B-mode signal by comparison, and the fact that these simulations use CMB theory spectra that assume $r = 0$. If it is approximately equal, this could be due to the presence of lensing modes, though these modes should still be much weaker than the general E-mode signal seen across the CMB.

4 Conclusion

It is possible there is a fault in the simulation itself that leads to this peculiar result. The simulation and its results require more analysis before a definitive answer can be made. In particular, there are multiple refinements and adjustments that could be made to the simulation that could help explain the nature of the negative r bias found in systematic simulations. One is simply running more realizations of the simulation. This helps create a more accurate bandpower covariance matrix that in turn forces the maximum likelihood search function to consider most data points equally when fitting the model. The simulation also may have overestimated noise bias when considering only the mean of the noise spectra bandpowers; it is possible a more refined definition of noise bias may be required. Finally, the addition of a delensing correction, even if not strictly necessary, may be useful in determining if lensed bandpowers were contributing significant amounts of B-modes to leak into E-modes.

Given more time, a conclusive result on the effects of polarization systematics on r could be made using these simulations. The next steps of this analysis require implementing the suggested improvements to the simulation and looking for a change in the results. For the moment, however, our preliminary results suggest that polarization systematics at the level expected by CMB-S4 induce a significant bias in r nearly twice the size of the science target, and that this bias is consistently negative, meaning these systematics increase the likelihood of a false negative. Both of these results would indicate that CMB-S4 needs to revise its tolerance levels for polarization detector systematics to be more stringent.

References

[1] Kevork Abazajian, Graeme Addison, Peter Adshead, Zeeshan Ahmed, Steven W. Allen, and David Alonso, et Al. CMB-S4 Science Case, Reference Design, and Project Plan. *arXiv e-prints*, page arXiv:1907.04473, July 2019.

[2] P.A.R. Ade, Z. Ahmed, R.W. Aikin, K.D. Alexander, D. Barkats, S.J. Benton, C.A. Bischoff, J.J. Bock, R. Bowens-Rubin, J.A. Brevik, I. Buder, E. Bullock, V. Buza, J. Connors, J. Cornelison, B.P. Crill, M. Crumrine, M. Dierickx, L. Duband, C. Dvorkin, J.P. Filippini, S. Fliescher, J. Grayson, G. Hall, M. Halpern, S. Harrison, S.R. Hildebrandt, G.C. Hilton, H. Hui, K.D. Irwin, J. Kang, K.S. Karkare, E. Karpel, J.P. Kaufman, B.G. Keating, S. Kefeli, S.A. Kernasovskiy, J.M. Kovac, C.L. Kuo, N.A. Larsen, K. Lau, E.M. Leitch, M. Lueker, K.G. Megerian, L. Moncelsi, T. Namikawa, C.B. Netterfield, H.T. Nguyen, R. O'Brient, R.W. Ogburn, S. Palladino, C. Pryke, B. Racine, S. Richter, A. Schillaci, R. Schwarz, C.D. Sheehy, A. Soliman, T. St.Germaine, Z.K. Staniszewski, B. Steinbach, R.V. Sudiwala, G.P. Teply, K.L. Thompson, J.E. Tolan, C. Tucker, A.D. Turner, C. Umiltà, A.G. Vieregg, A. Wandui, A.C. Weber, D.V. Wiebe, J. Willmert, C.L. Wong, W.L.K. Wu, H. Yang, K.W. Yoon, and C. Zhang. Constraints on primordial gravitational waves using planck, wmap, and new bicep2/keck array observations through the 2015 season. *Physical Review Letters*, 121(22), November 2018.

[3] David Alonso, Javier Sanchez, and Anže Slosar. A unified pseudo-c_ℓ framework. *Monthly Notices of the Royal Astronomical Society*, 484(3):4127–4151, January 2019.

[4] Colin Bischoff. S4bb package for cmb-s4 spectral analysis. https://github.com/CMB-S4/s4bb.

[5] Paolo Campeti and Eiichiro Komatsu. New constraint on the tensor-to-scalar ratio from the planck and bicep/keck array data using the profile likelihood. *The Astrophysical Journal*, 941(2):110, December 2022.

[6] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759–771, April 2005.

[7] Samira Hamimeche and Antony Lewis. Likelihood analysis of cmb temperature and polarization power spectra. *Physical Review D*, 77(10), May 2008.

[8] Wayne Hu and Martin White. A cmb polarization primer. *New Astronomy*, 2(4):323–344, October 1997.

[9] Planck FFP10 Simulation Suite, https://wiki.cosmos.esa.int/planck-legacy-archive/index.php/Simulation$_data$.

[10] Andrea Zonca, Leo P. Singer, Daniel Lenz, Martin Reinecke, Cyrille Rosset, Eric Hivon, and Krzysztof M. Gorski. healpy: equal area pixelization and spherical harmonics transforms for data on the sphere in python. *Journal of Open Source Software*, 4(35):1298, 2019.

[11] Andrea Zonca, Ben Thorne, Nicoletta Krachmalnicoff, and Julian Borrill. The python sky model 3 software. *Journal of Open Source Software*, 6(67):3783, November 2021.