

CMB polarization at large angular scales: Data analysis of the POLAR experimentChristopher W. O'Dell,^{1,*} Brian G. Keating,² Angelica de Oliveira-Costa,³ Max Tegmark,³ and Peter T. Timbie⁴¹*Department of Astronomy, University of Massachusetts, Amherst, Massachusetts 01003, USA*²*Division of Physics, Math, and Astronomy, California Institute of Technology, Pasadena, California 91125, USA*³*Department of Physics, University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*⁴*Department of Physics, 2531 Sterling Hall, University of Wisconsin–Madison, Madison, Wisconsin 53706, USA*

(Received 20 December 2002; published 29 August 2003)

The coming flood of cosmic microwave background (CMB) polarization experiments, spurred by the recent detection of CMB polarization by the DASI and WMAP instruments, will be confronted by many new analysis tasks specific to polarization. For the analysis of CMB polarization data sets, the devil is truly in the details. With this in mind, we present details of the data analysis for the POLAR experiment, which recently led to the tightest upper limits on the polarization of the cosmic microwave background radiation at large angular scales. We discuss the data selection process, map-making algorithms, offset removal, and likelihood analysis which were used to find upper limits on the polarization. Stated using the modern convention for reporting CMB Stokes parameters, these limits are $5.0 \mu\text{K}$ on both E - and B -type polarization at 95% confidence. Finally, we discuss simulations used to test our analysis techniques and to probe the fundamental limitations of the experiment.

DOI: 10.1103/PhysRevD.68.042002

PACS number(s): 98.70.Vc

I. INTRODUCTION

The detection of polarization in the cosmic microwave background (CMB) has been a long sought goal for cosmology. CMB polarization was recently detected at subdegree angular scales by the DASI instrument, a ground-based interferometer [1], and the temperature-polarization correlation has now been detected at larger angular scales by the Wilkinson Microwave Anisotropy Probe (WMAP) satellite [2]. While information on the CMB polarization at degree and subdegree angular scales further confirms the standard cosmological model, polarization information at larger angular scales has the potential to provide additional information regarding the formation and evolution of the Universe. The process of reionization leaves a characteristic signature on CMB polarization at large angular scales which can be used as a means to determine the epoch of reionization [3–5]. The power spectrum of CMB polarization at angular scales greater than a degree or so is sensitive to inflationary model parameters such as the inflaton potential and the energy scale of inflation, as well as primordial gravitational waves [6–10].

These exciting rewards, taken together with increasingly sensitive receiver technologies, have set the stage for a host of new CMB polarization experiments. These experiments will be faced with new and more challenging analysis tasks than for the simpler case of anisotropy, and it is with that in mind that we now set out to report the details of the data analysis of the POLAR experiment.

POLAR (polarization observations of large angular regions) was designed to measure the polarization of the CMB on large angular scales, in the K_a frequency band from 26–36 GHz. This high electron mobility transistor (HEMT)-based correlation polarimeter operated for a single season in

the spring of 2000 near Madison, Wisconsin. The data from this single season led to simultaneous upper limits on E - and B -type CMB polarization, results that were initially presented in [11] (hereafter K01). The details of the POLAR instrument and its operation were later described in [12] (hereafter K02). In this paper we present the details of the data selection and analysis techniques used to arrive at the results in K01. We also discuss the results of a recent cross-correlation analysis of the POLAR data with anisotropy data from the Cosmic Background Explorer (COBE) Differential Microwave Radiometer (DMR) experiment, described fully in [13].

The rest of this paper is organized as follows. We first briefly discuss the two conventions used to quantify CMB polarization. We next review the basic properties of the POLAR instrument in Sec. III. In Sec. IV, we discuss the data selection procedures that were used to remove large amounts of contaminated data. Section V describes the mapmaking algorithms used to transform the raw data into sky maps of Q and U , and provides a full pipeline simulation in order to test the algorithms. Section VI presents the likelihood analysis used to arrive at the upper limits on CMB polarization, as well as the evaluation of the polarization power spectra and some commentary on the lack of substantial foreground contamination. Finally, in Sec. VII, we discuss the limitations of our experiment that could be improved upon in future projects.

II. Q AND U CONVENTIONS

In recent years there have appeared in the literature two conventions for reporting the Stokes parameters Q and U as applied to the CMB. The first convention takes $Q(\vec{x}) = T_x(\vec{x}) - T_y(\vec{x})$, where T_x and T_y are antenna temperatures measured in orthogonal directions by a single-mode radiometer with unit optical efficiency, and \vec{x} denotes the angle of observation; U is obtained similarly by rotating the radiometer coordinates by 45° . These antenna temperatures are then

*Electronic address: codell@astro.umass.edu

converted to thermodynamic temperatures in the usual way (see, e.g., [14]). This convention was employed by the pioneering experiments of Lubin and Smoot and others [15,16], and was also adopted for POLAR.

The more recent convention (see, e.g., [17,18]) takes $Q = (T_x - T_y)/2$, and also has been adopted by several experiments [1,19–21]. The results herein are presented using the former convention (in thermodynamic units) unless otherwise noted.¹ Specifically, this means that our basic limits on *E*- and *B*-type polarization (presented in K01 and in Sec. VIA of this paper) must be divided by a factor of 2 when directly comparing to experiments or theory employing the latter convention. We originally reported our simultaneous upper limits on *E*- and *B*-type polarization in K01 as 10.0 μK at 95% confidence; stated using the latter convention, these limits become 5.0 μK each.

III. INSTRUMENT

POLAR observed the local zenith from the University of Wisconsin–Madison’s Pine Bluff Observatory in Pine Bluff, Wisconsin (latitude $+43^\circ 01'$, longitude $+89^\circ 45'$) using a simple drift-scan strategy, with a 7° full width at half maximum (FWHM) beam defined by a K_a -band microwave feed horn. A correlation radiometer operated as a polarimeter that was instantaneously sensitive to the *U* Stokes parameter. The full rf band was divided into three subbands, 32–36, 29–32, and 26–29 GHz. Each of the channels was detected by a separate microwave correlator, labeled *J1*, *J2*, and *J3*, respectively. The constant rotation of the instrument about its vertical axis allowed for simultaneous detection of both the *Q* and *U* Stokes parameters.

POLAR employed a phase-sensitive detection technique. The relative phase of the two arms of the correlation radiometer was modulated at approximately 1 kHz. An analog lock-in circuit produced the actual “in-phase” polarization signal. However, the instrument simultaneously locked into the chop frequency shifted by $\pi/2$, yielding pure noise. We obtained one of these “quadrature-phase channels” (QPCs) for each of the three in-phase channels (IPCs). The quadrature-phase channels were a good monitor of the noise in our system, and as such were employed in several places throughout the analysis pipeline. We use the terms *J1i*, *J2i*, and *J3i* to refer to the three in-phase channels, while *J1q*, *J2q*, and *J3q* refer to the corresponding quadrature-phase (noise) channels.

The data were sampled at 20 Hz and continuously recorded to files, each containing precisely 9000 samples (7.5 min) of data; these files are one of the fundamental data units to be discussed throughout this paper. The instrument was calibrated daily with a 3-mil thick dielectric sheet. We did not calibrate the instrument when the weather was poor. The calibration was accurate to about 10%, based upon laboratory measurements of the dielectric sheet properties; see [22] for details. Each *section* of data is defined as the longest

period of observations between calibrations. Approximately 50 such sections were collected throughout the observing season, each some 2–24 h in length.

POLAR’s constant rotation rate allows us to characterize its response to a polarized signal as follows:

$$\mathbf{y}(t) = I_0 + C \cos \omega t + S \sin \omega t + Q \cos 2\omega t + U \sin 2\omega t + n(t), \quad (1)$$

where $\omega = 2\pi f = 0.2055 \text{ rad s}^{-1}$ was the (angular) rotation frequency. The constant offset I_0 is due to coupling of the unpolarized total power signal into the polarization channels via the nonzero cross polarization of the instrument. This offset term was typically 10–100 mK, depending on the channel; during good weather its stability was better than 0.6 mK per hour. *C* and *S* are signals modulated at the rotation frequency (referred to hereafter as 1ϕ signals), and can be caused by various types of ground pickup and other systematic effects. Via Eq. (1), linearly polarized signals will produce a signal in the data stream at twice our rotation frequency, henceforth called the 2ϕ frequency.

IV. OBSERVATIONS AND DATA SELECTION

During the observing season, POLAR observed 24 hours per day over a two-month period, yielding roughly 750 hours of data. However, the observing season contained a large diversity of weather conditions, and this led to a correspondingly large diversity of data quality. Developing robust data selection criteria was one of the most critical tasks in the data analysis pipeline. Because of the data’s diversity, we were not able to arrive at a single selection criterion; rather, we developed a battery of conditions that the data were required to meet before being accepted. Many effects can conspire to contaminate the polarization signals, be they instrumental, atmospheric, or celestial. We wish to flag and remove any data with a noncosmological contribution to our signal that mimics a cosmological signal in a way that we cannot account for and remove.

The rest of this section will describe the various criteria we established in order to robustly separate out contaminated data. There were three basic types of cuts used: instrument based, celestial, and data based. Instrument-based cuts were those in which our systems were not functioning properly or led to unanalyzable data. Celestial cuts were those in which some noncosmological celestial source was in a position to contaminate our data, such as the sun, moon, or galaxy. Data-based cuts were those which used noise properties of the data themselves to assess contamination. Table I lists all of the cut criteria used in the POLAR analysis; the entries in the table will be described in more detail throughout the rest of this section.

We flagged data at the file level; that is, we either kept or cut individual 7.5-min length data files. This proved convenient as well as ensuring good noise stationarity over an individual file, while allowing for a good estimation of the low frequency noise, necessary for our mapmaking algorithms. A related question to that of the cutting time scale is whether or not to cut all channels simultaneously versus

¹We shall switch to the latter convention when discussing the temperature-polarization correlation results.

TABLE I. Amounts of data cut at various steps in the data selection process.

Cut type	Hours surviving ^a	Hours cut ^b	Individual hours cut ^c
Full data set	746.5 (100%)	0 (0%)	0 (0%)
Instrument	629.3 (84.3%)	117.3 (15.7%)	117.3 (15.7%)
Celestial	270.9 (36.3%)	358.4 (48%)	434.3 (58.2%)
Data based	152 (20.4%)	118.9 (15.9%)	478 (64%)
Length based	78.1 (10.5%)	73.8 (9.9%)	N/A

^aThe amount of data left after that cut and all the cuts above it have been applied.

^bAmount of data cut at that stage.

^cAmount of data that would have been cut if that particular criterion were the only one applied.

separately. An important fact of our experiment was that some of the effects contaminating the data were frequency dependent; in particular, channel $J3$ (26–29 GHz) often showed evidence of a spurious signal when the other two channels did not. For this reason, we performed the data-based cuts on each of the channels individually; the instrumental and celestial cuts had no frequency dependence, and hence were applied equally to the three channels.

A. Instrument-based cuts

We performed two instrument-based cuts. The first cut was to remove any files during which the system was not rotating for the entire file, or the rotation was unnaturally slow or jittery. This cut removed 1.8% of the data. The other issue was dew formation on the optics, a common problem in ground-based experiments. When the relative humidity was high for an hour or more, moisture condensed on the vacuum window, leading to a spurious polarization signal. The mechanism was likely that ambient radiation scattered off the layer of water, leading to the spurious signal. Dew contamination necessitated removal of 14% of the data.

B. Celestial cuts

Because the feedhorn accepts some radiation from all angles, a celestial source will contribute a nonzero amount of radiation when it is anywhere above the horizon. We are concerned with celestial objects leaking radiation through a sidelobe in our beam, or the main lobe if the source ever gets close enough to the local zenith. The sun, moon, and planets as well as point sources and diffuse radiation from the galaxy are all possible contaminants.

1. Solar system objects

We used the measured beam profile (see K02) to determine attenuation of the microwave emission from the sun, moon, and planets as a function of angle from the beam (which equals altitude for our zenith-staring experiment). Treating the sun as an 11 000 K blackbody with an angular width of 0.5° FWHM [23], and conservatively assuming 1% polarization, we find that contamination from the sun is less

than $0.4 \mu\text{K}$ for elevations less than 70° . Similarly, the lunar model of [24] shows that the moon can be treated as a blackbody of approximately 240 K in its brightest phase, with an angular diameter of 0.5° and $\leq 1\%$ polarization. Again assuming 1% polarization, lunar emission falls below $0.05 \mu\text{K}$ at zenith angles greater than 40° . Using cut elevations of 70° for the sun and 40° for the moon removed a sizable 42% of the data.

In principle, contamination by the planets was also possible. The largest planetary signal was estimated to be $\sim 0.01 \mu\text{K}$ in total intensity. This was due to Jupiter, which reached a maximum of 65° elevation during our observations. Including the fraction of the source that is polarized will reduce this number even further. Thus in the case of POLAR planetary radiation can be safely neglected.

2. Galactic foreground emission

Finally, there is the question of emission from the galaxy, from either diffuse or point sources. POLAR had a sensitivity to point sources of roughly $2 \mu\text{K}/\text{Jy}$. Using the WOMBAT point-source catalog for 30 GHz [25], we calculate that there are approximate 40 sources that passed through our beam contributing greater than $1 \mu\text{K}$ antenna temperature in intensity. Only three sources contribute greater than $10 \mu\text{K}$. We chose not to perform a point-source-based cut as detecting these sources, while challenging, would be both interesting and relatively easy to identify as galactic versus cosmological in origin.

The observations described were conducted over a wide range of galactic latitudes and therefore there is a potential for diffuse galactic contamination, especially at low latitudes [26,27]. Galactic synchrotron emission can be up to 70% polarized [28]. No maps of polarized synchrotron emission exist at 30 GHz, and extrapolation of measurements at lower frequencies, e.g., [29], is not a reliable probe of synchrotron polarization at 30 GHz due to Faraday rotation. Although the unpolarized intensity is apparently not correlated with the polarized intensity as shown in [13], we attempt to limit our susceptibility to synchrotron emission by using only data corresponding to galactic latitudes $|b| > 25^\circ$.

C. Data-based cuts

Together, the instrument-based and celestial cuts removed approximately 64% of the original 750-hr data set. cursory analysis, however, immediately revealed that spurious signals still remained in this reduced data set at unreasonable levels. This is not surprising given that the weather in Wisconsin is highly variable in the spring, and most of the radiation entering the system was atmospheric emission. Especially important was the presence of clouds, which can mimic a polarized signal to our radiometer [30]. We therefore follow the standard practice of deriving selection statistics based upon the data themselves, to serve as a measure of the quality of observing conditions.

1. The 1ϕ cut

Signals modulated only at the 2ϕ frequency correspond to true polarization signals; thus, a signal that has response at

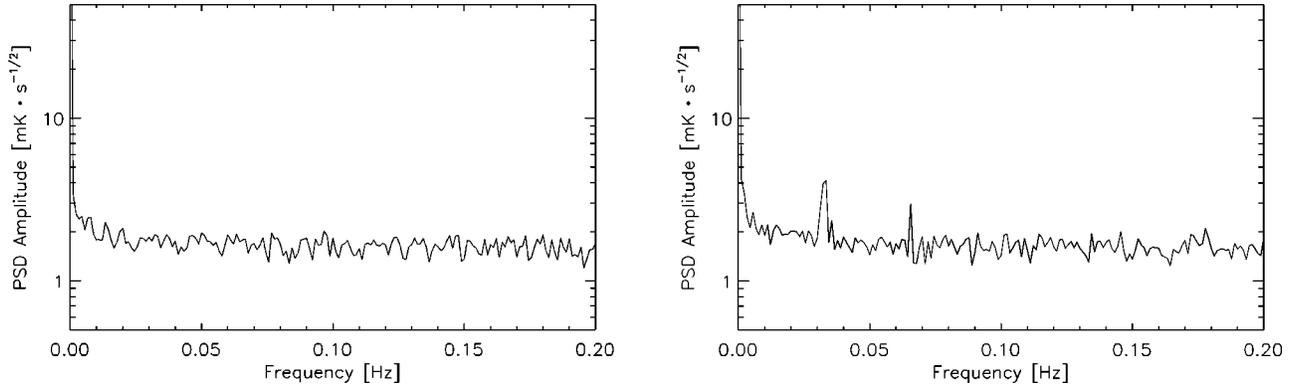


FIG. 1. Low frequency power spectrum of two sections of data for channel $J2i$. (a) shows a power spectrum of the time stream when the weather is good and systematic effects are low. (b) displays a similar period when clouds appeared; there is a dramatic appearance of spikes at both the 1ϕ and 2ϕ frequencies. This motivated our using the strength of the 1ϕ peak as a cut statistic. True celestial polarization signals will appear at the 2ϕ frequency only.

both 1ϕ and 2ϕ cannot correspond to a true celestial signal. This can be seen in the power spectrum of the data (Fig. 1). Panel (a) shows a roughly featureless power spectrum taken during a period of good weather for channel $J2i$. But as the weather worsens, due to clouds and/or increased humidity and water vapor, features at the 1ϕ and 2ϕ frequencies appear [Fig. 1(b)]. This motivates a study of the correlation between 1ϕ and 2ϕ signals in the time stream.

We computed the 1ϕ and 2ϕ values for each data file relative to the noise floor,² denoted as $1\phi_r$ and $2\phi_r$, respectively. Figure 2 displays these data plotted against each other for channel $J2i$, with both instrument-based and celestial cuts applied. The high degree of correlation between the two quantities is striking; this strongly motivated the use of the $1\phi_r$ level as a cut statistic. In order to determine a logical $1\phi_r$ cut level, we compared our data to the behavior of white noise. Monte Carlo simulations showed that this statistic is 1.0 ± 0.27 for white noise, distributed roughly as a Gaussian. We opted to cut whenever $1\phi_r > 2.1$; this is more than 4σ from the mean for white noise. Implementing this requirement cut an additional 15% of the data. However, even after instituting this cut there remained unphysically high values of $2\phi_r$ in the data, motivating additional cuts.

2. The zeta cut

Because the autocorrelation function and power spectrum of any data set form a Fourier transform pair, the information in one is the same as in the other. For instance, a rise in $1/f$ noise leads directly to a higher “floor” in the autocorrelation function. Any signal in the data will lead to a nontrivial autocorrelation function. We used that fact to our advantage and defined the following statistic for each data file:

$$\zeta \equiv \frac{\sum_{\text{lag}=1}^{1000} C(y_{\text{in}})^2}{\sum_{\text{lag}=1}^{1000} C(y_{\text{quad}})^2}, \quad (2)$$

²We found it useful to compare the height of these harmonics relative to the noise floor, to remove the effect of variable radiometer noise.

where y_{in} denotes data from an in-phase channel and y_{quad} denotes data from the corresponding quadrature-phase channel. Simulations showed that a reasonable astrophysical signal (of, say, less than $100 \mu\text{K}$ rms) would have a negligible contribution to ζ for POLAR. Equation (2) has an intuitive explanation; it is roughly the integral of the power spectrum between 0.01 Hz (near our lowest observable frequency) and 10 Hz (our Nyquist frequency), weighted by $1/f$, so low frequency drifts cause ζ to increase rapidly. As such drifts are generally indicative of poor atmospheric conditions, ζ proves to be a sensible cut statistic.

Figure 3 shows a histogram of ζ values throughout the season (with the basic sun, moon, and dew cuts applied). Also shown in the plot is a model of our data stream with no $1/f$ noise or signal of any kind. As can be seen, this distribution is peaked around 1.0, with more than 99.9% of the data lying below 1.2. However, any deviation from white noise will rapidly change the ζ value of the data. We found that requiring $\zeta < 4$ cut about 3% of the data over the other cuts. Varying the actual cut level between $\zeta < 2$ and $\zeta < 5$ had little effect on the final results.

3. Outliers in the time-ordered data

Occasionally, birds, planes, electrical noise, etc., would cause short-lived yet large spikes in the data stream. For each

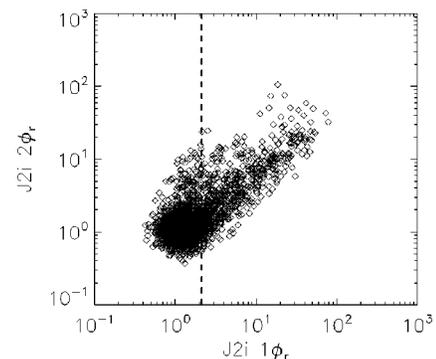


FIG. 2. $2\phi_r$ vs $1\phi_r$ for the data channel $J2i$. Periods contaminated by sun, moon, and dew have been removed. The vertical dashed line shows the $1\phi_r < 2.1$ cut level.

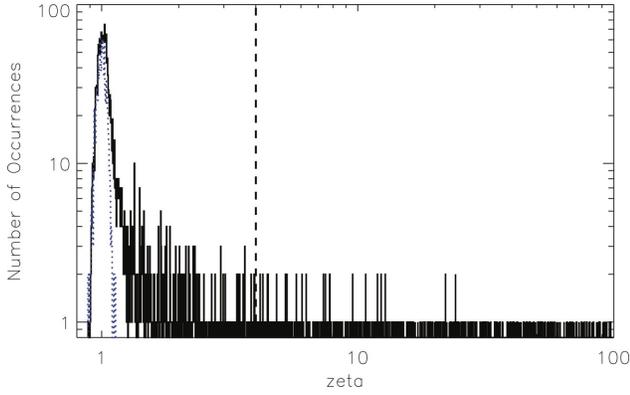


FIG. 3. (Color) Distribution of the ζ variable. The solid (black) curve shows the distribution of ζ for channel $J2$, with the basic sun, moon, and dew cuts applied. The dotted (blue) curve shows the ζ distribution for simulated white noise run through our antialiasing filter. The vertical dashed curve is the cut level of $\zeta=4.0$.

data file and channel, we calculated the mean and standard deviation, and recorded how many standard deviations the first, second, and third strongest outliers were from the mean. We cut any data file whose *second largest outlier* was more than 5σ from the mean of those data. Data files with one strong outlier alone rarely occurred, because our 5 Hz filters in conjunction with our 20 Hz data acquisition frequency guaranteed that even a delta-function signal should have a width of at least a few samples in the data stream; thus the second strongest outlier proved a more robust indicator for such spurious effects. This cut removed $\sim 1.2\%$ of the data in addition to all previous cuts.

D. Duration-based cuts

In order to ensure that our surviving data segments were obtained during long, contiguous periods of good weather, we instituted a battery of three duration-based cuts. First, we required that if a data file were to survive, both its neighbors had to survive as well. Second, we required a minimum of eight consecutive data files (one hour of data) to survive in order to keep any piece of data. Lastly, we required a minimum of three hours total to survive in a *section*³ of data for any data in that section to be retained; in contrast to the cut above, this three hours was not required to be contiguous. This cut was instituted because of the offset removal step of the analysis pipeline, described in Sec. VD. This algorithm takes place on a per section basis and effectively removes all information from very short sections of data; in the end we found it simpler just to remove these short sections of data. Together, the duration-based cuts removed about 10% of the full data set.

V. FROM DATA TO MAPS

We now proceed to construct Q and U maps of the postcut data set, a process that does not remove any cosmological

information [31]. Many authors have written on the map-making problem and all the nuances that can arise during its solution: Wright (1996) [32]; Tegmark (1997), hereafter T97a [31]; Bond, Jaffe, and Knox (1998) [33]; and Stompor *et al.* (2002) [34]. We use the following notational conventions. When discussing a vector, we use a lowercase boldface letter (e.g., \mathbf{y}) to represent it. Similarly, matrices are represented by uppercase boldface letters (e.g., \mathbf{W}).

We employ the technique of *minimum variance map making* throughout this work, defined briefly as follows; the interested reader is referred to T97a for a thorough discussion of this topic. We begin with a set of n observations $\mathbf{y} = y_1, \dots, y_n$ of signal plus noise, such that $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$, where \mathbf{x} is the underlying map, \mathbf{n} represents noise, and \mathbf{A} is the pointing matrix that encodes the mapping from data space to map pixel space (see [32] for examples). Let \mathbf{N} represent the data noise covariance matrix, where $\mathbf{N} \equiv \langle \mathbf{nn}' \rangle$. The best estimate $\tilde{\mathbf{x}}$ of the true map \mathbf{x} is given by⁴

$$\tilde{\mathbf{x}} = [\mathbf{A}'\mathbf{N}^{-1}\mathbf{A}]^{-1}\mathbf{A}'\mathbf{N}^{-1}\mathbf{y}, \quad (3a)$$

$$\Sigma = [\mathbf{A}'\mathbf{N}^{-1}\mathbf{A}]^{-1}, \quad (3b)$$

where Σ is the *map noise covariance matrix* of $\tilde{\mathbf{x}}$. Note that the *total map covariance matrix* \mathbf{C} is the sum of signal plus noise: $\mathbf{C} = \mathbf{S} + \Sigma$, where $\mathbf{S} \equiv \langle \mathbf{xx}' \rangle$ is the theory covariance matrix. The total covariance matrix will be used in Sec. VI. This method is quite general, and can be applied to any type of problem where a linear combination of data is made in order to determine some physical quantity. Whenever the noise of said data is not white, this is the best approach. For POLAR, we exploit this technique no fewer than four different times throughout the analysis pipeline.

A. Constructing the file maps

The data pipeline begins with binning the data into our map coordinates, via Eq. (3). We formed a small *file map* for each 7.5-min (9000-sample) data file, for each channel and in both Q and U . The “map” we calculated for each data file consisted of the Q and U parameters for whatever map pixels were observed during that particular file. Because of our scan strategy, a zenith drift scan at $\delta=43^\circ$, our maps are one dimensional in right ascension (for each of Q and U). We pixelized our maps with 180 pixels of width 2° in right ascension (RA), with the first pixel arbitrarily centered at $RA=0^\circ$. The pointing matrix \mathbf{A} was formed according to Eq. (1); note that this matrix is not sparse and is significantly more complicated in the case of extracting both Q and U rather than the usual intensity T (see [30] for further details).

We performed binning on each data file separately because our noise was stationary only over periods of tens of minutes. In principle, even nonstationary noise can be treated with the standard mapmaking formalism, but the algorithms are unacceptably slow because the covariance matrix \mathbf{N} has no special properties except that it is symmetric, and the

³A section of data is a data segment that was taken without interruption (this was typically about one day of data).

⁴This estimate of $\tilde{\mathbf{x}}$ is typically termed the “COBE-style solution” [32,35].

central rate determining step of the mapmaking algorithm of Eq. (3) is in calculating \mathbf{N}^{-1} . For stationary noise, however, the corresponding noise covariance matrix has the special property that it is both symmetric and *Toeplitz*; that is, each upper left–lower right diagonal is the same. Toeplitz matrices can be inverted in $O(n_t^2)$ time, as opposed to $O(n_t^3)$ for general $n_t \times n_t$ matrices [36]. However, the memory requirements are still $O(n_t^2)$ because the inverse of a symmetric Toeplitz matrix is not in general Toeplitz. For POLAR, $n_t = 9000$, so simply holding the inverse matrix takes ~ 324 Mbytes of memory for double-precision arithmetic. While not prohibitive, this memory requirement motivates further study of approximation techniques.

We use a slight modification of the circulant matrix approximation introduced in [37], which notes that a typical noise covariance matrix \mathbf{N} is nearly *circulant*. In this case, $\mathbf{N} = \mathbf{N}_c + \mathbf{N}_s$, where \mathbf{N}_c is the circulant component of \mathbf{N} and \mathbf{N}_s is the noncirculant component. As \mathbf{N}_c is the major contributor to \mathbf{N} , Ref. [37] lets $\mathbf{N} \rightarrow \mathbf{N}_c$ in Eq. (3a); the equation for the covariance matrix becomes significantly more complicated. This solution is valid in that it still produces an unbiased, nearly optimal approximation of the map, at the expense of slightly increased noise. It has the advantage that circulant matrices can be inverted in $O(n_t \ln n_t)$ time.

However, it is possible to enact the circulant matrix approximation in a more streamlined way. Using the definitions that

$$\mathbf{y}_{\text{pw}} \equiv \mathbf{N}_c^{-1/2} \mathbf{y}, \quad (4a)$$

$$\mathbf{A}_{\text{pw}} \equiv \mathbf{N}_c^{-1/2} \mathbf{A}, \quad (4b)$$

the circulant matrix approximation becomes

$$\tilde{\mathbf{x}} \equiv [\mathbf{A}_{\text{pw}}^t \mathbf{A}_{\text{pw}}]^{-1} \mathbf{A}_{\text{pw}}^t \mathbf{y}_{\text{pw}}, \quad (5a)$$

$$\tilde{\Sigma} \equiv [\mathbf{A}_{\text{pw}}^t \mathbf{A}_{\text{pw}}]^{-1}. \quad (5b)$$

The second equation above is not precisely exact, but for our data it was correct to better than 1%. We state the mapmaking algorithm in this fashion because there is a very fast way to calculate \mathbf{y}_{pw} and \mathbf{A}_{pw} . Because a circulant matrix becomes diagonal in the Fourier domain, it can be shown that

$$\mathbf{b}_{\text{pw}} \equiv \mathbf{N}_c^{-1/2} \mathbf{b} = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{\mathbf{b}\}}{\sqrt{S_N(f)}} \right\}, \quad (6)$$

where \mathbf{b} is any vector of length n_t , S_N is the power spectral density that characterizes the covariance matrix \mathbf{N} , and \mathcal{F} denotes the Fourier transform. This approach is advantageous in that it requires only a simple fit to the noise power spectrum for the time stream of interest; it is unnecessary to explicitly calculate the resulting covariance matrix \mathbf{N} . The resulting vector \mathbf{b}_{pw} is simply the prewhitened version of \mathbf{b} . We fitted our noise to a model containing $1/f$ noise alone:

$$S_N(f) = \sigma^2 \left(1 + \frac{f_{\text{knee}}}{f} \right), \quad (7)$$

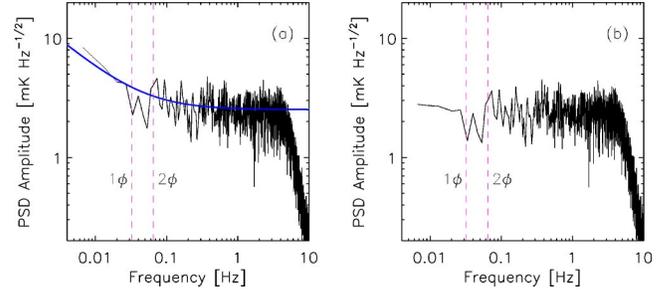


FIG. 4. (Color online) Effect of prewhitening on a sample noise power spectrum for channel *J2i*. (a) Power spectral density (PSD) amplitude for a sample data file, with a knee frequency of ~ 0.07 Hz. The thick (blue) curve is the logarithmic-weighted fit to the PSD. The vertical dashed lines show the 1ϕ and 2ϕ rotational frequencies. (b) Same as (a), but for the \mathbf{y}_{pw} version of the data. The effect was to whiten the power spectrum of the data; it is equivalent to dividing the PSD by the fitted curve.

where f_{knee} is the knee frequency of the $1/f$ noise. It proved unnecessary to include the antialiasing filter in the model, as it has virtually no effect on the power spectrum at the 2ϕ frequency, where the Q and U data live. We used Eq. (6) to then calculate \mathbf{y}_{pw} and \mathbf{A}_{pw} for all data files passing the cuts,⁵ and applied Eq. (5) to generate roughly 1000 1–2 pixel maps in each of Q and U , along with the corresponding covariance matrices (see Fig. 4).

B. The problem of the Q and U offsets

At this point in the analysis, a serious issue revealed itself. The Q and U data were not in general consistent with zero; there were slowly varying offsets in Q and U that varied over the course of the season, and differed between the channels. This fact is illustrated in Fig. 5, which shows the derived time stream of Q and U values for all the surviving data. Notice first that the data lie in “chunks” along the time axis; each such chunk corresponds with a *section* as defined in Sec. III. The distribution of Q 's and U 's for each section is consistent with a Gaussian distribution, but those distributions are in general not centered around zero. Notice also that there is no clear relationship between the Q or U offset among the three channels. This is in contrast to the QPC data, which were offset-free. Thus the offset was not electronic in nature, and also was not consistent with a celestial signal when plotted in sky coordinates.

The underlying cause of the offset was never discovered, but there are several reasonable possibilities. K02 hypothesizes that the quadrupolar shape of the outer groundscreen may be to blame [12]. Reference [30] examines the characteristics of the offsets in great detail and provides several theories for the possible cause. If we had some model of the effects that could make successful predictions about its level, we could use it to safely subtract out the offsets without strongly affecting our signal recovery. However, without

⁵Equation (6) was applied to each row of \mathbf{A} to form the corresponding row of \mathbf{A}_{pw} .

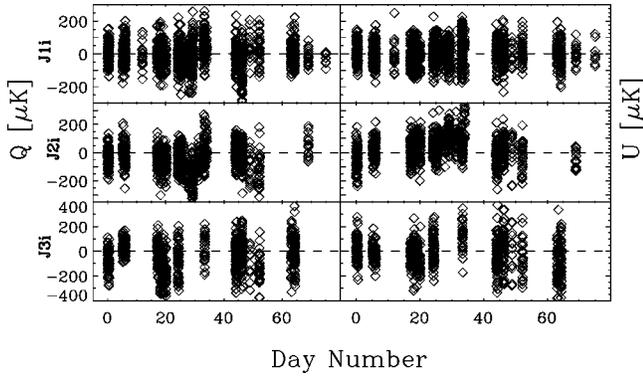


FIG. 5. The mean values of Q and U for the POLAR in-phase data versus time, after the cuts have been applied. Each data point represents one data file. Non-Gaussian behavior is immediately apparent at the 50–100 μK level. This behavior was not seen the quadrature-phase data.

such a model we proceed along a different tack.

C. Submap generation

It is a simple task to combine all the file maps for a single channel and Stokes parameter, within a single section into a *submap* (each corresponding roughly to one night of good data). Given a set of file maps to be combined, we first expand them to the full 180-pixel map by assigning unmeasured pixels infinite noise (and hence zero weight). The correlations between Stokes Q and U were negligible both within a pixel and between pixels, and were therefore neglected. This allowed Q and U to be treated completely independently throughout the analysis. We use the standard map combination prescription to combine the file maps into submaps (see, e.g., de Oliveira-Costa *et al.* [38]). Given a set of m independent maps $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ with corresponding noise covariance matrices $\{\mathbf{N}_1, \dots, \mathbf{N}_m\}$, the best estimates of the underlying map \mathbf{x}_f and its associated covariance matrix \mathbf{N}_f are given by

$$\mathbf{N}_f = \left[\sum_{i=1}^m \mathbf{N}_i^{-1} \right]^{-1}, \quad (8a)$$

$$\mathbf{x}_f = \mathbf{N}_f \left[\sum_{i=1}^m \mathbf{N}_i^{-1} \mathbf{x}_i \right]. \quad (8b)$$

In this way, we combined all file maps into submaps for each section and channel. At this step each submap was still left with a substantial offset in Q and U , as discussed previously.

D. Marginalization and submap combination

We next employed the technique of *marginalization* to each such submap [33], in order to remove sensitivity to the mean of Q and U in each. We performed the algorithm separately for each channel. Note that this technique provides identical results to those obtained using the method of *virtual pixels* for mode removal, described, for instance, in [37]. Detailed recipes for applying the marginalization technique

are described in S01 [34]; we give a cursory review here as to how it was specifically applied for POLAR. Given a submap \mathbf{y} (containing Q or U data) with covariance matrix \mathbf{N} , the covariance matrix is transformed to

$$\mathbf{N}' = \mathbf{N} + \sigma_c^2 \mathbf{Z} \mathbf{Z}', \quad (9)$$

where the matrix \mathbf{Z} contains a normalized column vector corresponding to each mode to be removed, and σ_c is some large number to ensure that the modes get zero weight. For marginalization over the mean we take $\mathbf{Z} = \mathbf{e}_0$, where \mathbf{e}_0 is the column vector of all 1's. Taking the limit as $\sigma_c^2 \rightarrow \infty$, the inverse of \mathbf{N}' still exists and is given by

$$\mathbf{N}'^{-1} = \mathbf{N}^{-1} - (\mathbf{N}^{-1} \mathbf{Z}) [\mathbf{Z}' \mathbf{N}^{-1} \mathbf{Z}]^{-1} (\mathbf{N}^{-1} \mathbf{Z})'. \quad (10)$$

We then applied Eq. (8) to combine all the submaps into a single map of Q and U for each channel, for both the IPC and QPC data. This recipe requires knowing only the inverse covariance matrix corresponding to each submap, given explicitly by Eq. (10). However, we still have to perform the final inversion of

$$\mathbf{\Sigma}^{-1} \equiv \sum_{i=1}^m \mathbf{N}_i'^{-1} \quad (11)$$

to find the final noise covariance matrix $\mathbf{\Sigma}$, and $\mathbf{\Sigma}^{-1}$ is singular due to our marginalization over the means of all the submaps. Following S01, the final covariance matrix is taken to be

$$\mathbf{\Sigma} = (\mathbf{\Sigma}^{-1} + \epsilon \mathbf{Z} \mathbf{Z}')^{-1} - \epsilon^{-1} \mathbf{Z} \mathbf{Z}', \quad (12)$$

where ϵ is any small positive number. In practice, it is best to choose ϵ to be of the same order as the nonzero eigenvalues of $\mathbf{\Sigma}^{-1}$.

The final maps for each of our three frequencies were constructed according to the recipe given above and presented in K01. Qualitatively, there is no visual evidence of a common signal among the three subbands, in either Q or U . The χ^2 values from each map are also not consistent with a statistically significant signal. Let us additionally consider the corresponding maps made from the QPC (pure noise) channels. We do not expect these to contain signals, but they serve as a useful litmus test when viewing the IPC maps: if the IPC maps differ strongly from the QPC maps, that is evidence of either signal or some type of contamination in the IPC maps. However, that is not the case; none of the QPC maps contain strong outliers and all exhibit χ^2 values consistent with no signal.

Combining the channel maps

Based on the fact that there is no apparent correlation between the maps of each channel, we simplify the analysis by combining the three channel maps into one combined map (and covariance matrix) for each of Q and U . Since the CMB signal is independent of frequency (when expressed in thermodynamic units), we are free to do this. We would like to employ the standard map coaddition algorithm of Eq. (8) to perform this task. However, that algorithm assumes that

TABLE II. Interchannel cross-correlation coefficients. The errors are the same within each row, and we assume that the underlying distribution of correlation coefficients is Gaussian.

	$\langle QQ \rangle$	$\langle UU \rangle$	$\langle QU \rangle$	$\langle UQ \rangle$
$\langle J1 J2 \rangle_{\text{IPC}}$	0.144 ± 0.034	0.134	-0.024	0.021
$\langle J1 J2 \rangle_{\text{QPC}}$	0.005 ± 0.034	0.062	-0.011	-0.002
$\langle J1 J3 \rangle_{\text{IPC}}$	0.074 ± 0.041	0.063	-0.069	-0.003
$\langle J1 J3 \rangle_{\text{QPC}}$	0.023 ± 0.041	0.048	-0.009	-0.042
$\langle J2 J3 \rangle_{\text{IPC}}$	0.104 ± 0.041	0.093	-0.024	-0.029
$\langle J2 J3 \rangle_{\text{QPC}}$	0.066 ± 0.041	0.001	-0.054	-0.047

the measurements made of each individual map are *independent*; if there is a systematic effect that introduces correlations between measurements from different channels, then we have less information than we think we do, and this fact must be taken into account in constructing the summed map.

We first evaluated the correlation coefficients among channels $J1i$, $J2i$, and $J3i$ in the time-ordered data. As expected, the correlations in the time stream were on the order of 1%, and were due to the very slight overlap of the three channels in frequency space (see K02 for details). However, it is not so much the time stream correlations that we care about, it is the correlations between Q or U for the channels. For instance, if there were a 10% correlation between $J1i$ Q and $J2i$ U , it could be hidden in the smaller time stream correlations. We must therefore evaluate these correlations directly.

In order to measure these correlations, we used the Q and U data set and determined the Pearson's correlation coefficient in the same way as for the time-ordered data, but because there are so many fewer data, we evaluated only one correlation coefficient for each surviving section. We calculated means and errors of these coefficients by averaging from the distribution of these values, shown in Table II. The numbers in this table are very suggestive. For instance, $\langle QQ \rangle$ for all IPC channels is about the same as $\langle UU \rangle$, suggesting a common source. All QPC correlation coefficients are consistent with zero, as are all correlations of the $\langle QU \rangle$ variety.⁶ As Q and U show no correlation between them (either within a channel or between channels), we can continue treating Q and U as completely independent measurements. This is not too surprising, considering they are essentially the $\sin 2\phi$ and $\cos 2\phi$ projections from each rotation, which are orthogonal functions. However, the correlations between IPC channels (for the same Stokes parameter) are $\sim 10\%$, so we cannot ignore them in constructing a final map.

Let us now combine the maps of Q (or U) from the three nonindependent channels, armed with the knowledge of their mutual correlations. The algorithm of Sec. VC did not deal with adding nonindependent maps together, but it is relatively easy to expand the methods to do so. We treat Q and U separately, as they are completely uncorrelated. Let us consider our situation for Q (U will follow an identical format).

⁶Except perhaps $\langle Q_1 U_3 \rangle$, but because all the other coefficients of this type are consistent with zero, we assume it is an outlier.

We have three channel maps, call them \mathbf{x}_1 , \mathbf{x}_2 , and \mathbf{x}_3 , with their corresponding covariance matrices Σ_1 , Σ_2 , and Σ_3 . Let the correlation coefficient between \mathbf{x}_i and \mathbf{x}_j be ρ_{ij} . We appeal to minimum variance map making to form the best possible map. First, we form the concatenated map and corresponding covariance matrix:

$$\mathbf{x}_{\text{cat}} = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}, \quad (13a)$$

$$\Sigma_{\text{cat}} = \begin{bmatrix} \Sigma_1 & \rho_{12}\Sigma_{12} & \rho_{13}\Sigma_{13} \\ \rho_{12}\Sigma_{12} & \Sigma_2 & \rho_{23}\Sigma_{23} \\ \rho_{13}\Sigma_{13} & \rho_{23}\Sigma_{23} & \Sigma_3 \end{bmatrix}, \quad (13b)$$

where $\Sigma_{ij} \equiv \sqrt{\Sigma_i} \sqrt{\Sigma_j}$. We can take the square roots of the Σ matrices since they are all positive definite, as long as we add a large offset to each matrix (corresponding to the uncertainty in the offset, which is formally infinite).⁷ The final full covariance matrix Σ_f will then also have a large offset, but again this corresponds to our marginalization over the means of all the component submaps and it will not affect our final results.

We now apply the minimum variance mapmaking formalism to the concatenated map [which in this case is our data vector \mathbf{y} from Eq. (3)]. The pointing matrix is given by

$$\mathbf{A}_{\text{cat}} = \begin{bmatrix} \mathbf{I}_n \\ \mathbf{I}_n \\ \mathbf{I}_n \end{bmatrix}, \quad (14)$$

where \mathbf{I}_n is the $n \times n$ identity matrix, and n is the number of pixels in our maps. This matrix points our three individual maps to the same final map. Explicitly, the final joint map and covariance matrix are given by

$$\mathbf{x}_f = \Sigma_f \mathbf{A}_{\text{cat}}^t \Sigma_{\text{cat}}^{-1} \mathbf{x}_{\text{cat}}, \quad (15a)$$

$$\Sigma_f = [\mathbf{A}_{\text{cat}}^t \Sigma_{\text{cat}}^{-1} \mathbf{A}_{\text{cat}}]^{-1}. \quad (15b)$$

Because of the large offset each covariance matrix possesses, the final map \mathbf{x}_f may have some random offset to it, but it is meaningless and can be safely subtracted out. Notice that, in the case of vanishing intermap correlations, we correctly reproduce the algorithm of Eq. (8).

The final joint maps for the IPCs and QPCs are shown in Fig. 6. There is no obvious evidence of an underlying sky signal. In order to determine robustly if they show evidence of a signal, we employ the standard techniques of maximum likelihood analysis and quadratic estimators, described fully in Sec. VI.

⁷For a symmetric, positive-definite, $n \times n$ matrix \mathbf{M} , its square root is given by $\mathbf{P}' \mathbf{D}^{1/2} \mathbf{P}$, where \mathbf{P} is an $n \times n$ matrix such that the i th row of \mathbf{P} contains the i th eigenvector of \mathbf{M} , and $\mathbf{D}^{1/2}$ is a diagonal matrix with the square roots of the eigenvalues of \mathbf{M} along its diagonal. The eigenvectors must be normalized, such that $\mathbf{P}\mathbf{P}' = \mathbf{I}$.

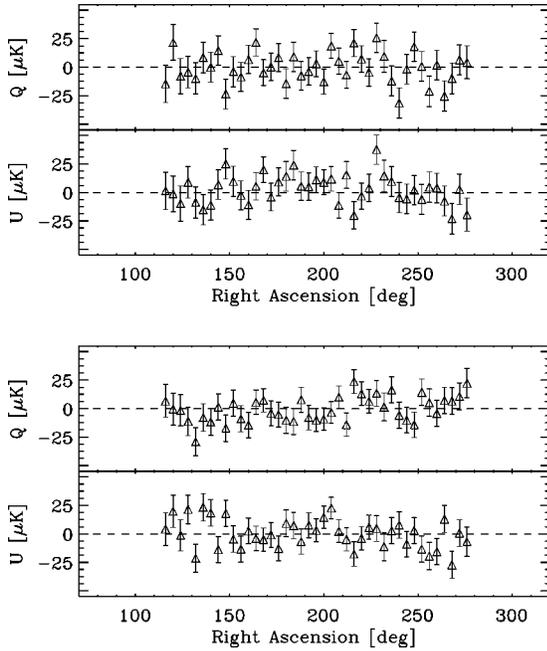


FIG. 6. Final joint-channel sky maps. The joint IPC maps for Q and U are displayed in (a), while (b) shows the QPC maps.

E. Testing the map-making pipeline

We next simulated the data stream in order to test our map reconstruction pipeline. There were three primary steps involved in the simulation process: building the underlying map, letting POLAR “observe” this fake sky and generate data based on these observations, and constructing maps from the resulting data set. We built the underlying sky maps out of simple sine and cosine modes, for both total intensity (I) and Stokes Q and U . We assumed a basic flatband-power model with $\sim 10 \mu\text{K}$ per band, and convolved these signals with our 7° (FWHM) Gaussian beam. This signal level is of course highly unrealistic for true CMB polarization; we used it merely so we could reconstruct the actual map with a good signal-to-noise ratio (instead of merely providing an upper limit).

To shorten the processing time, we made the simulated instrument about a factor of 2 more sensitive than the real POLAR. We included all IPC and QPC channels in the analysis, but no total power channels. We assumed the noise was almost white, with a small amount of $1/f$ noise in each channel in agreement with the true data set. We convolved each data stream with our 5 Hz antialiasing filter to ensure a resulting power spectrum that was similar to the actual data. We also added random offsets in I , Q , and U for each section and channel, of levels consistent with those experienced by POLAR. The results of these simulations showed that our map reconstruction algorithms worked as expected, even when including large random offsets for each section in the Q and U data, and illustrates the amazing utility of the marginalization technique. The Q and U reconstruction of one such “fake sky” is shown in Fig. 7.

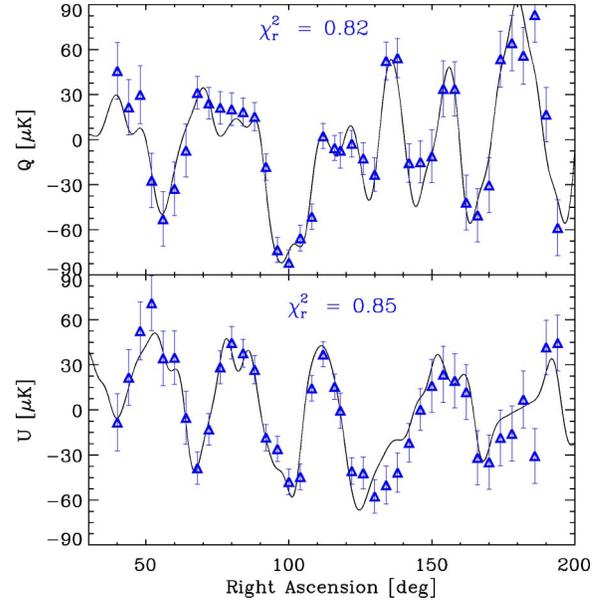


FIG. 7. (Color online) Derived joint-channel maps from simulated data for one example “fake-sky” map. The thin (black) curve is the underlying sky map convolved with a 7° beam. The (blue) data points represent the derived joint-channel map, after 31 h of simulated data in five submaps. The good agreement illustrates the robustness of our mapmaking algorithms and of the marginalization technique used to remove sensitivity to the means of Q and U from each contributing submap.

VI. CONSTRAINING CMB POLARIZATION MODELS

A. Limits on E and B in a flatband-power model

1. Primary likelihood analysis

We now address the question of the level of rms polarization in the joint-channel maps using a standard likelihood analysis. For the beginner, an excellent discussion of how to carry out such an analysis is provided in [39].

As POLAR measured both Stokes Q and U simultaneously, we are able to set limits on both E - and B -type polarization independently. As described previously in K01, we constrain a flat CMB power spectrum characterized by two polarization temperatures T_E and T_B such that $\ell(\ell+1)C_\ell^X/2\pi = T_X^2$ where $X \in \{E, B\}$. We then solve for the likelihood function $\mathcal{L}(\vec{a})$, given by

$$\mathcal{L}(\vec{a}) \propto \frac{e^{-\mathbf{x}^T \mathbf{C}^{-1}(\vec{a}) \mathbf{x} / 2}}{|\mathbf{C}(\vec{a})|^{1/2}}, \quad (16)$$

where \vec{a} is our parametrization vector $\vec{a} \equiv \{T_E, T_B\}$, $\mathbf{x} = \{\mathbf{q}, \mathbf{u}\}$ is the concatenation of the joint channel maps of Q and U given in the previous section, and $\mathbf{C}(\vec{a})$ is the full covariance matrix given by the sum of the data (or noise) and theory covariance matrices:

$$\mathbf{C}(\vec{a}) = \mathbf{S}(\vec{a}) + \mathbf{\Sigma}. \quad (17)$$

The construction of the theory covariance matrix \mathbf{S} is rather complicated [7,40,41], and has been previously presented for

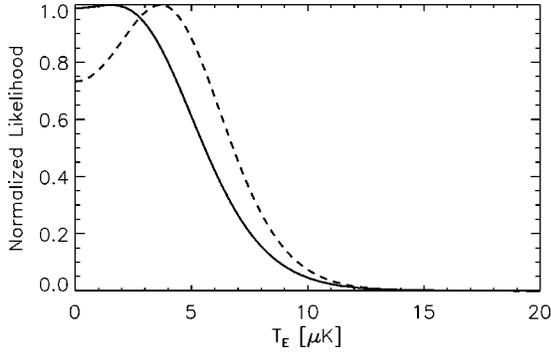


FIG. 8. Normalized likelihood plot of T_E , with the prior constraint that $T_B=0$. The solid line is the result for the in-phase channels, and the dashed line is for the quadrature-phase (null) channels. The resulting upper limit is $T_E < 7.7 \mu\text{K}$ at 95% confidence.

POLAR [11,30]. The data covariance matrix Σ is formed from the covariance matrices for the joint-channel maps of Q and U such that

$$\Sigma = \begin{bmatrix} \Sigma_Q & 0 \\ 0 & \Sigma_U \end{bmatrix}. \quad (18)$$

Note that Σ has a very large offset (and corresponding eigenvalue) due to our original marginalization over the submap means, but its specific magnitude will not affect results of the likelihood analysis.

We calculated the likelihood function for all individual and joint-channel maps, for both the in-phase and quadrature-phase channels. As reported in K01, all cases were consistent with pure upper limits. The combined-channel data yield 95% confidence limits of $10.0 \mu\text{K}$ on both T_E and T_B (or $5.0 \mu\text{K}$ stated using the modern convention for defining Q and U). As the B polarization at large angular scales is assumed to be so much weaker than E polarization (even taking into account lensing), we can set T_B to zero. It yields a 95% confidence limit of $T_E < 7.7 \mu\text{K}$, as shown in Fig. 8.

2. Alternative likelihood analyses

To verify the analysis, we repeated it in two different ways. First, we performed the likelihood analysis *without* the offset removal. The corresponding likelihood analysis yields a spurious detection for the in-phase channels at $(T_E, T_B) = (18, 1) \pm (8, 16) \mu\text{K}$ (95% confidence), while the quadrature channels yield a 95% confidence upper limit of $7 \mu\text{K}$ for both T_E and T_B . We see that the offset removal degraded the upper limit by $\sim 30\%$. This is primarily due to the large width ($\sim 25^\circ$) in right ascension of the central lobe of the theory covariance matrix, as shown in Fig. 9, coupled with our marginalization over the mean of each submap. This is consistent with the fact that each submap averaged roughly 90° of coverage in right ascension. This emphasizes to the experimenter that having long, clean sections of data is key to obtaining the best final noise possible, especially if any kind of mode removal may take place in the analysis.

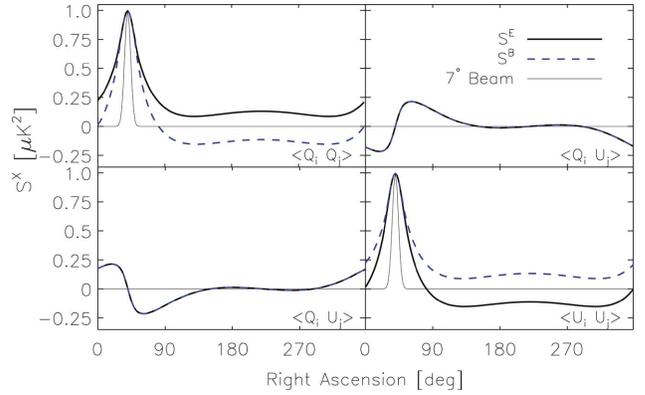


FIG. 9. (Color) Plots of the $\text{RA}=40^\circ$ row of the fundamental signal covariance matrices \mathbf{S}^E and \mathbf{S}^B . The four panels in the plot correspond to the four quadrants of the matrices as labeled. The thick solid line represents \mathbf{S}^E , the thick dashed line represents \mathbf{S}^B , and the thin solid line corresponds to a 7° Gaussian beam for reference.

As a further consistency check, we performed the entire analysis in such a way that the three channels were coadded directly in the time-ordered data. This technique automatically takes into account any correlations that may be present between the channels. One simply coadds the time streams with their inverse noise weightings in order to obtain a time stream with the minimum possible noise. Offset marginalization was still done on the submaps, but since we had already combined all the channels, a single offset was removed for Q and U for each section, rather than one for each channel. When the corresponding likelihood contours were calculated, the upper limits remained, but were degraded to about $12 \mu\text{K}$. This makes sense considering we subtracted only two offsets per section, as compared with six in the primary analysis. Because the offsets were not perfectly correlated among the three channels, there was residual power left over in the maps due to the imperfect coaddition of the channel offsets; it was this phenomenon that led to the slightly degraded upper limits. This result notwithstanding, this analysis is noteworthy in that it shows that our interchannel correlations were not a severe problem.

B. Limits on polarization power spectra

CMB anisotropy and polarization are typically characterized by the temperature fields T , Q , and U . These quantities can be combined to form six measurable power spectra: TT , EE , BB , TE , TB , and EB (see, for instance, [41]). Our analysis of these spectra is presented in [13]; we give a brief summary of it here. Note that, in contrast to the rest of this paper, temperatures in this section are reported using the modern convention of $Q = (T_x - T_y)/2$, as discussed in Sec. II.

As POLAR had very weak sensitivity to temperature anisotropy data, we used the COBE DMR data for the 53 and 90 GHz bands averaged together as discussed in [13]. We analyzed the spectra in five bands of width $\ell=5$, but subsequently averaged them into a single band to increase the

TABLE III. POLAR DMR power spectrum. All values are reported using the convention $Q \equiv (T_x - T_y)/2$, as discussed in Sec. II. Table reproduced from [13].

	$\ell_{\text{eff}} \pm \delta\ell$	$\delta T^2 \pm \sigma$ (μK^2)	δT (μK) ^a
<i>T</i>	15.6 ± 6.6	487.0 ± 270.6	$22.1^{+7.4}_{-5.5}$
<i>E</i>	12.6 ± 4.5	-4.9 ± 16.0	$<3.3(5.2)$
<i>B</i>	12.6 ± 4.5	6.9 ± 16.0	$<4.8(6.3)$
<i>X</i>	14.0 ± 4.8	-18.4 ± 34.3	$<4.0(7.1)$
<i>Y</i>	14.0 ± 4.8	-0.1 ± 34.3	$<5.9(8.3)$
<i>Z</i>	11.4 ± 2.9	-25.0 ± 15.8	$<(2.6)$

^aValues in parentheses are 2σ upper limits.

signal-to-noise ratio. We employed the method of quadratic estimators to evaluate the band powers (see, e.g., [41,42]), using the concordance model ($\Omega_{\text{Lambda}} \sim 0.7, \Omega_{\text{mat}} \sim 0.3, h \sim 0.7$) as input for the theory covariance matrices [43]. All band powers were consistent with upper limits, the results of which are shown in Table III. Note that the 2σ upper limits on *E*- and *B*-type polarization obtained with quadratic estimators are 5.2 and 6.3 μK , respectively, in good agreement with the 2σ upper limits of 5.0 μK obtained from the maximum likelihood analysis.

This technique has the added benefit of explicitly calculating the band-power window functions, while for a maximum likelihood analysis they are less straightforward to evaluate [44]. However, there is a new twist on band powers when it comes to polarization; there is always some amount of leakage into the desired power spectrum from the other five power spectra. In principle, it is possible to choose priors for the quadratic estimators such that 14 of the 15 leakages are zero, but the much discussed *E*-*B* leakage remains [41], although it can be kept to a minimum at the price of a modest increase in error bars. This leakage is a direct result of “ambiguous” modes existing in the scan strategy; the more ambiguous modes, the worse will be the leakage [45]. Sky coverage that is large, two dimensional, and well connected will have few ambiguous modes; but POLAR’s one-dimensional scan strategy leads to virtually *all* of the *E* and *B* modes being ambiguous.

Figure 10 shows the *E* window function for a single ℓ band for $2 < \ell < 8$ for POLAR as an example; the window function was generated using the minimum leakage techniques discussed above. The leakage of *B* into *E* is exactly symmetric for *E* into *B*; thus we show only the *E* window function. It is clear that there is a very large leakage from *B* to *E* and vice versa. Therefore, the upper limits for *E* and *B* cannot be taken separately; we constrain merely the average of the two power spectra.

An early reionization leads to a peak in both the *EE* autocorrelation and *TE* cross-correlation spectra at large angular scales [3,46]; this effect was observed recently in the *TE* power spectrum by WMAP [2]. Unfortunately, our *EE* and *TE* upper limits are still significantly higher than any possible reionization peak. For $\tau=0.4$, a typical concordance model gives $\delta T_{EE} \sim 1 \mu\text{K}$ and $\delta T_{TE} \sim 3 \mu\text{K}$, in comparison

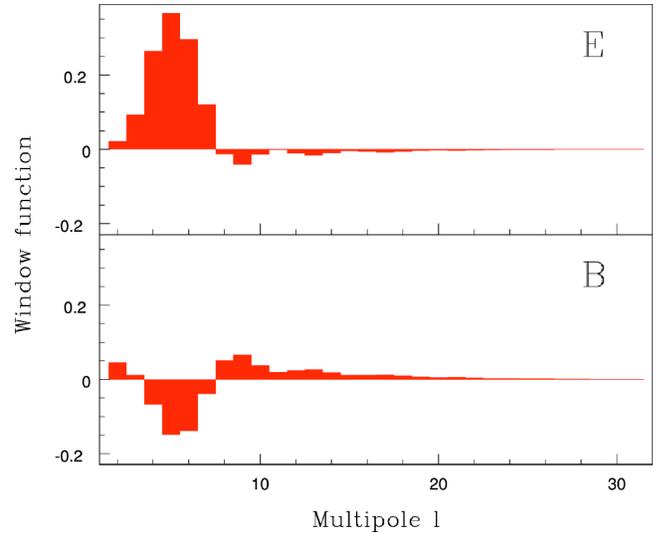


FIG. 10. (Color online) The window function for *E* used to estimate the polarization band power near $\ell=5$, shown as a representative example. The *B* window functions are exactly the same as for *E*, but with *E* and *B* switched. There is significant leakage of *B* power into the *E* estimate because of our one-dimensional scan strategy. In general, a scan done over a two-dimensional region with a large sky coverage will have a much better *E*-*B* separation and narrower window functions. The width of the window function scales with the inverse of the sky patch size in its narrowest dimension.

with our limits⁸ of 5.2 μK on δT_{EE} and 7.1 μK on δT_{TE} . In comparison, WMAP recently reported $\tau=0.17 \pm 0.04$ (68% confidence) [2].

C. Constraining polarized synchrotron emission

Astrophysical synchrotron emission is generally polarized at a level from 10% to 50% [27,29]. Our nondetection of any signal implies that we can place constraints on polarized synchrotron emission at 30 GHz, but this is not as straightforward as it may seem. As discussed in Sec. IV B 2, we cannot simply extrapolate from low frequency polarization measurements such as those of Brouw and Spoelstra [29] because of Faraday depolarization. Extrapolation of the unpolarized Haslam 408 MHz data [47] to our frequencies using a spectral index of -3.1 , assuming a polarization fraction of 10%, and smoothing with a 7° beam yields a rms of about $10 \mu\text{K}$.⁹ Thus, our data favor both a steep spectral index as suggested by recent studies [48] and a polarization fraction of less than 10% at large angular scales. This issue is explored more fully in [13].

VII. DISCUSSION

To summarize our results thus far, using the POLAR data we have placed limits on CMB polarization power spectra in

⁸Stated using the convention $Q=(T_x-T_y)/2$ as discussed in Sec. II.

⁹Stated using the convention $Q=T_x-T_y$ as discussed in Sec. II.

the context of both a flatband-power model and a concordance model, and we have discussed the lack of substantial synchrotron contamination in our scan region. We showed in Sec. VIA that the offset removal technique degraded our limits by only $\sim 30\%$, which is not so bad considering that without it no limits would have been possible.

If we had integrated for a much longer time and/or our detectors had been much more sensitive, how well could we have actually done in determining E and B on large angular scales given our scan strategy? We attempted to answer this question via an additional simulation. We used CMBFAST [49] to generate power spectra for a concordance model universe with no B modes (i.e., only scalar perturbations) and an optical depth to reionization of $\tau=0.2$; ideally, we should recover only E modes from the analysis. We used the HEALPIX package [50] to generate 20 sky realizations of Q and U , and observed each with a nearly perfect POLAR, one that had a sensitivity of $0.01 \mu\text{K}$ for each of $180 2^\circ$ -wide pixels in our $\delta=43^\circ$ RA strip. We then performed our likelihood analysis on these map realizations limited to our scan region, and found that we recovered roughly equal amounts of E and B power. This is consistent with the window functions generated in our quadratic estimator analysis of Sec. VIB. On a positive note, we did recover the correct amount of total E

and B power; that is, the recovered $C_\ell^E + C_\ell^B$ equaled the initial C_ℓ^E placed in the maps.

Ultimately, POLAR was limited by both its instrument sensitivity and atmospheric limitations to integration time. The atmosphere was also most likely responsible for the time-varying polarization offset. Finally, our one-dimensional scan strategy prevented discrimination between E and B modes. Future experiments that effectively deal with these issues may well be able to glean the CMB polarization signals on large angular scales, and hence uncover a wealth of new information with which to increase our understanding of inflation and the details of the big bang.

ACKNOWLEDGMENTS

We are grateful to Josh Gundersen, Lloyd Knox, Ed Wolack, Matias Zaldarriaga, Slade Klawikowski, and Phil Farese for many useful conversations. B.K. and C.O. were supported by NASA. B.K. acknowledges support from the NSF Astronomy and Astrophysics program. POLAR's HEMT amplifiers were graciously provided by John Carlstrom. This work has been supported by NSF grants AST 93-18727, AST 98-02851, and AST 00-71213, and NASA grant NAG5-9194.

-
- [1] J. Kovac, E. M. Leitch, J. E. Carlstrom, C. Pryke, W. L. Holzapfel, and N. W. Halverson, *Nature (London)* **420**, 772 (2002).
 [2] A. Kogut *et al.*, astro-ph/0302213.
 [3] M. Zaldarriaga, *Phys. Rev. D* **55**, 1822 (1997).
 [4] W. Hu, *Astrophys. J.* **529**, 12 (2000).
 [5] M. Kaplinghat *et al.*, *Astrophys. J.* **583**, 24 (2003).
 [6] M. Zaldarriaga and U. Seljak, *Phys. Rev. D* **55**, 1830 (1997).
 [7] M. Kamionkowski, A. Kosowsky, and A. Stebbins, *Phys. Rev. D* **55**, 7368 (1997).
 [8] W. H. Kinney, *Phys. Rev. D* **58**, 123506 (1998).
 [9] L. Knox and Y. Song, *Phys. Rev. Lett.* **89**, 011303 (2002).
 [10] M. Kesden, A. Cooray, and M. Kamionkowski, *Phys. Rev. D* **67**, 123507 (2003).
 [11] B. G. Keating, C. W. O'Dell, A. de Oliveira-Costa, S. Klawikowski, N. Stebor, L. Piccirillo, M. Tegmark, and P. T. Timbie, *Astrophys. J. Lett.* **560**, L1 (2001).
 [12] B. G. Keating, C. W. O'Dell, J. O. Gundersen, L. Piccirillo, N. C. Stebor, and P. T. Timbie, *Astrophys. J., Suppl. Ser.* **144**, 1 (2003).
 [13] A. de Oliveira-Costa, M. Tegmark, C. O'Dell, B. Keating, P. Timbie, G. Efstathiou, and G. Smoot, astro-ph/0212419.
 [14] R. Patridge, *3K: The Cosmic Microwave Background Radiation* (Cambridge University Press, New York, 1995).
 [15] P. M. Lubin and G. F. Smoot, *Astrophys. J.* **245**, 1 (1981).
 [16] G. P. Nanos, *Astrophys. J.* **232**, 341 (1979).
 [17] M. Hedman, Ph.D. dissertation, Princeton University, Princeton, NJ, 2002.
 [18] P. C. Farese *et al.*, in *The Cosmic Microwave Background and its Polarization*, Proceedings of the William I. Fine Theoretical Physics Institute, Minneapolis, 2003.
 [19] M. Hedman, D. Barkats, J. O. Gundersen, S. T. Staggs, and B. Winstein, *Astrophys. J. Lett.* **548**, L111 (2001).
 [20] G. Hinshaw *et al.*, astro-ph/0302222.
 [21] John Cartwright (private communication).
 [22] C. W. O'Dell, D. S. Swetz, and P. T. Timbie, *IEEE Trans. Microwave Theory Tech.* **50**, 2135 (2002).
 [23] H. Zirin, B. M. Baumert, and G. J. Hurford, *Astrophys. J.* **370**, 779 (1991).
 [24] S. J. Keihm, in *Planetary Science Inst. Report* (1983).
 [25] A. Sokasian, E. Gawiser, and G. F. Smoot, *Astrophys. J.* **562**, 88 (2001).
 [26] C. L. Bennett, A. J. Banday, K. M. Gorski, G. Hinshaw, P. Jackson, P. Keegstra, A. Kogut, G. F. Smoot, D. T. Wilkinson, and E. L. Wright, *Astrophys. J. Lett.* **464**, L1 (1996).
 [27] R. D. Davies, R. A. Watson, and C. M. Gutierrez, *Mon. Not. R. Astron. Soc.* **278**, 925 (1996).
 [28] G. B. Rybicki and A. P. Lightman, *Radiative Processes in Astrophysics* (Wiley-Interscience, New York, 1979).
 [29] W. Brouw and T. Spoelstra, *Astron. Astrophys., Suppl. Ser.* **26**, 129 (1976).
 [30] C. W. O'Dell, Ph.D. dissertation, University of Wisconsin-Madison, 2001.
 [31] M. Tegmark, *Astrophys. J. Lett.* **480**, L87 (1997).
 [32] E. L. Wright, astro-ph/9612006.
 [33] J. R. Bond, A. H. Jaffe, and L. Knox, *Phys. Rev. D* **57**, 2117 (1998).
 [34] R. Stompor *et al.*, *Phys. Rev. D* **65**, 022003 (2002).
 [35] D. J. Jansen and S. Gulkis, in *NATO ASIC Proceedings 359: The Infrared and Submillimetre Sky after COBE* (1992), Proceedings of the NATO Advanced Study Institute, Les Houches, France, 1991.
 [36] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed.

- (Johns Hopkins University Press, Baltimore, 1996).
- [37] M. Tegmark, *Phys. Rev. D* **56**, 4514 (1997).
- [38] A. de Oliveira-Costa, M. J. Devlin, T. Herbig, A. D. Miller, C. B. Netterfield, L. A. Page, and M. Tegmark, *Astrophys. J. Lett.* **509**, L77 (1998).
- [39] E. F. Bunn, Ph.D. dissertation, University of California–Berkeley, 1993.
- [40] M. Zaldarriaga, *Astrophys. J.* **503**, 1 (1998).
- [41] M. Tegmark and A. de Oliveira-Costa, *Phys. Rev. D* **64**, 063001 (2001).
- [42] J. R. Bond, A. H. Jaffe, and L. Knox, *Astrophys. J.* **533**, 19 (2000).
- [43] N. A. Bahcall, J. P. Ostriker, S. Perlmutter, and P. J. Steinhardt, *Science* **284**, 1481 (1999).
- [44] L. Knox, *Phys. Rev. D* **60**, 103516 (1999).
- [45] E. F. Bunn, M. Zaldarriaga, M. Tegmark, and A. de Oliveira-Costa, *Phys. Rev. D* **67**, 023501 (2003).
- [46] B. Keating, P. Timbie, A. Polnarev, and J. Steinberger, *Astrophys. J.* **495**, 580 (1998).
- [47] C. Haslam, C. Salter, H. Stoffel, and W. Wilson, NCSA Astronomy Digital Image Library, 1995.
- [48] P. Platania, M. Bensadoun, M. Bersanelli, G. de Amici, A. Kogut, S. Levin, D. Maino, and G. F. Smoot, *Astrophys. J.* **505**, 473 (1998).
- [49] U. Seljak and M. Zaldarriaga, *Astrophys. J.* **469**, 437 (1996).
- [50] K. M. Gorski, B. D. Wandelt, F. K. Hansen, E. Hivon, and A. J. Banday, “The HEALPIX Primer,” astro-ph/9905275, <http://www.eso.org/science/healpix>