



Mitigating Complex Dust Foregrounds in Future Cosmic Microwave Background Polarization Experiments

Brandon S. Hensley and Philip Bull

Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
brandon.s.hensley@jpl.nasa.gov, philbull@gmail.com

Received 2017 September 21; revised 2017 December 5; accepted 2017 December 29; published 2018 January 31

Abstract

Polarized Galactic foregrounds are one of the primary sources of systematic error in measurements of the B-mode polarization of the cosmic microwave background (CMB). Experiments are becoming increasingly sensitive to complexities in the foreground frequency spectra that are not captured by standard parametric models, potentially affecting our ability to efficiently separate out these components. Employing a suite of dust models encompassing a variety of physical effects, we simulate observations of a future seven-band CMB experiment to assess the impact of these complexities on parametric component separation. We identify configurations of frequency bands that minimize the “model errors” caused by fitting simple parametric models to more complex “true” foreground spectra, which bias the inferred CMB signal. We find that: (a) fits employing a simple two-parameter modified blackbody (MBB) dust model tend to produce significant bias in the recovered polarized CMB signal in the presence of physically realistic dust foregrounds; (b) generalized MBB models with three additional parameters reduce this bias in most cases, but non-negligible biases can remain, and can be hard to detect; (c) line-of-sight effects, which give rise to frequency decorrelation, and the presence of iron grains are the most problematic complexities in the dust emission for recovering the true CMB signal. More sophisticated simulations will be needed to demonstrate that future CMB experiments can successfully mitigate these more physically realistic dust foregrounds.

Key words: cosmic background radiation – cosmology: observations – ISM: general – methods: statistical

1. Introduction

B-mode polarization of the cosmic microwave background (CMB) arises from gravitational lensing by large-scale structure and, possibly, from gravitational waves generated during the inflationary epoch. Given the potential for constraining several important aspects of fundamental physics, such as the conditions of the very early universe and the sum of the neutrino masses, the measurement of the B-mode signal is a subject of intense focus for current and planned ground, balloon, and space missions.

In addition to the cosmological signature, Galactic foregrounds such as dust and synchrotron emission are also able to produce B-mode polarization. Disentangling the Galactic and cosmological signatures is already a challenge for current experiments, which have placed upper limits on the tensor-to-scalar ratio r of ~ 0.1 . Achieving constraints on r of order 10^{-3} , as sought by proposed next-generation experiments such as LiteBIRD (Matsumura et al. 2014) and CMB-S4 (Abazajian et al. 2016) will require subtraction of the Galactic signal with unprecedented accuracy.

Component separation in polarization can be significantly more complex than in total intensity. Polarized intensities add vectorially, with directions set by the interstellar magnetic field in the case of dust and synchrotron emission. Changes in the magnetic field direction along the line of sight coupled with spatial variations in the polarization spectra can result in different levels of cancellation/suppression of the polarized signal and a rotation of the polarization angle with frequency (e.g., Tassis & Pavlidou 2015). Thus, the observed polarized emission at one frequency is an imperfect predictor of the polarized emission at a different frequency, an effect termed “frequency decorrelation.” Frequency decorrelation effects,

which arise whenever the foreground spectra have spatial variations, are expected to be present at some level, and some evidence for them has already been noted in the *Planck* data (Planck Collaboration et al. 2017).

Frequency decorrelation poses a serious challenge for component separation methods; template-based methods can no longer rely on being able to factorize the frequency dependence and spatial variation of the foregrounds, while parametric spectral fitting methods require significantly more complex signal models (e.g., Chluba et al. 2017) to account for the extra spectral structure that is induced.

Complexities in the emission physics of dust can also be amplified in polarization relative to total intensity. If the far-infrared (FIR) dust emission arises from two distinct dust components (e.g., silicate and carbonaceous grains) with different polarization properties, the total and polarized dust spectral energy distributions (SEDs) can differ significantly. For example, while both components will contribute to the total intensity signal, one may be significantly less polarized than the other, resulting in far weaker polarized emission. Likewise, magnetic dust grains may contribute negligibly in total intensity at frequencies higher than the microwave, but emit strong, polarized emission at lower frequencies (Draine & Hensley 2013). These scenarios are also challenging for component separation methods, which tend to assume relatively simple spectral models, and extrapolate foreground properties in the (much higher signal-to-noise) total intensity channel into polarization.

The risks of improper dust modeling are well-documented: on intermediate angular scales, the residual foreground emission left after imperfect component separation can easily mimic the cosmological B-mode signal (BICEP2 Collaboration et al. 2014; BICEP2/Keck Collaboration et al. 2015), resulting

in strongly biased cosmological parameter inferences. It is therefore critical to the success of future CMB polarization experiments that they can (a) model and separate Galactic dust emission over a wide range of possible emission physics scenarios, and (b) reliably identify situations in which the modeling and subtraction are inadequate, and hence may bias the recovery of the true CMB signal.

Some recent work has studied the ability of future experiments to reliably recover the polarized CMB signal in the face of the complexities we consider here. Armitage-Caplan et al. (2012) found that neglecting the curvature in the dust SED due to departure from a pure Rayleigh–Jeans spectrum biased the recovered tensor-to-scalar ratio r high by $\sim 1\sigma$. More recently, Remazeilles et al. (2016) evaluated the robustness of parametric component separation to multiple dust components for a suite of proposed CMB satellites. They found that fitting too simple a dust model was sufficient to bias r by more than 5σ while maintaining an acceptable goodness-of-fit criterion, even for the most sensitive experiments. Similarly, Stompopor et al. (2016) found that fitting a single dust component in the presence of multiple dust “layers” biased parametric component separation at the $\Delta r \sim 10^{-3}$ level. Kogut & Fixsen (2016) considered dust SEDs with a distribution of dust temperatures as well as SEDs based on two-level systems, finding biases on r of $\Delta r \sim 3 \times 10^{-3}$. Poh & Dodelson (2017) evaluated the impact of multiple dust components along the line of sight aligned by different magnetic fields, finding that naive extrapolations from 350 GHz to lower frequencies would result in significant bias for $r \lesssim 1.5 \times 10^{-3}$.

In this paper, we consider the ability of future CMB polarization experiments to mitigate dust contamination in a much broader range of physically motivated scenarios. We focus our investigation on parametric component separation methods, which employ physical models of the frequency dependence of each emission component to perform separation pixel-by-pixel in the map domain (e.g., Planck Collaboration et al. 2016a). This type of method produces maps of each foreground component in addition to the CMB, enabling a wide variety of Galactic science. It is also well-suited for cosmological applications at large angular scales, such as measuring the reionization peak, where techniques based on spatial correlations can fail due to lack of modes (Remazeilles et al. 2017). For higher ℓ applications ($\ell \gtrsim 100$), non-parametric techniques will still need to be tested for robustness against the model complexities discussed here.

We first assess which complications to the simplest models lead to the greatest biases in the recovered CMB. Guided by a physical understanding of dust emission, more sophisticated models and techniques can be developed to mitigate these biases. Second, by analyzing a large set of frequency configurations that could be employed in future experiments, we evaluate what frequency coverage is most effective at both mitigating bias and identifying poor model fits via poor goodness-of-fit statistics.

This paper is organized as follows: in Section 2, we motivate and describe the suite of foreground models used in this work; in Section 3, we outline a “single-pixel” (i.e., frequency-space) component separation method; we assess the ability of various mission designs to accurately recover the input CMB as a function of dust model complexity in Sections 4 and 5, and we discuss the implications of this analysis for experiment design and data analysis in Section 6.

2. Foreground Models

In this section, we provide details of all of the foreground component models considered in this paper, with a particular focus on a set of seven dust models that illustrate a range of possible physical effects that would lead to more complex dust spectra than are typically considered in CMB analyses.

In most of this study, we will work in units of CMB brightness temperature. A blackbody of temperature T_{CMB} emits with a specific intensity of $I_{\nu}^{\text{CMB}} = B_{\nu}(T_{\text{CMB}})$, where $B_{\nu}(T)$ is the Planck function and T_{CMB} is taken to be 2.7255 K. A source with specific intensity I_{ν} has a CMB brightness temperature ΔT , satisfying:

$$I_{\nu} \simeq \left(\frac{\partial B_{\nu}}{\partial T} \right)_{T_{\text{CMB}}} \Delta T = \frac{2h\nu^3}{c^2} \frac{xe^x}{(e^x - 1)^2} \frac{\Delta T}{T_{\text{CMB}}}, \quad (1)$$

where $x = h\nu/kT_{\text{CMB}}$ and only the first-order terms are retained. The foreground models presented in the following sections are expressed in terms of specific intensities and converted to CMB temperature units, μK_{CMB} , for the analysis.

It is often convenient, particularly for visualizing the SEDs of the emission components over a broad frequency range, to employ the Rayleigh–Jeans brightness temperature T_{RJ} , which is related to the specific intensity at a given frequency by $T_{\text{RJ}}(\nu) = c^2 I_{\nu} / (2k \nu^2)$.

For polarization, we work with the Stokes parameters I_{ν} , Q_{ν} , and U_{ν} , expressed as specific intensities which can then be converted to the equivalent CMB or Rayleigh–Jeans brightness temperatures. Throughout this work, we neglect any circular polarization (i.e., Stokes $V_{\nu} = 0$). The polarized intensity P_{ν} therefore satisfies $P_{\nu} = |Q_{\nu} + iU_{\nu}|$. The polarization fraction p_{ν} is defined as $p_{\nu} \equiv P_{\nu}/I_{\nu}$.

2.1. Component Amplitudes

In the following sections, we describe models of the frequency dependence of the Galactic foregrounds considered in this work: thermal dust, synchrotron, anomalous microwave emission (AME), and free–free. As we implement a “single-pixel” analysis in which we fit a single representative realization of the CMB and the foregrounds (see Section 3), we must first determine the relative amplitudes of the various components for a typical high-latitude sightline.

To do so, we are guided by the results of the Commander component separation analysis (Planck Collaboration et al. 2016a). In total intensity, the Commander analysis produced full-sky maps at 1° resolution pixellated at a HealPix (Górski et al. 2005) $N_{\text{side}} = 256$. In Figure 1 we plot the cumulative distribution function (CDF) of the best-fit (posterior maximum) amplitudes of thermal dust, synchrotron, free–free, and spinning dust (AME) for all pixels with $|b| > 30^{\circ}$. As the foreground brightness varies strongly with Galactic latitude, the highest values observed in this sky cut give an indication of the brightest foregrounds that need to be mitigated by an experiment wishing to achieve 50% sky coverage.

At 353 GHz, the distribution of posterior maximum dust intensities is $23_{-13}^{+30} \mu\text{K}_{\text{RJ}}$ (68% credible interval). For all of our dust models, we adopt a Stokes I amplitude of $50 \mu\text{K}_{\text{RJ}}$ at 353 GHz, which is at the higher end of this range.

At 30 GHz, the distribution of best-fit synchrotron intensities is $19_{-5}^{+9} \mu\text{K}_{\text{RJ}}$, and we adopt a value of $30 \mu\text{K}_{\text{RJ}}$. Likewise, the AME amplitude distribution at 30 GHz is $19_{-11}^{+21} \mu\text{K}_{\text{RJ}}$ and we adopt a value of $30 \mu\text{K}_{\text{RJ}}$. Free–free emission is not detected

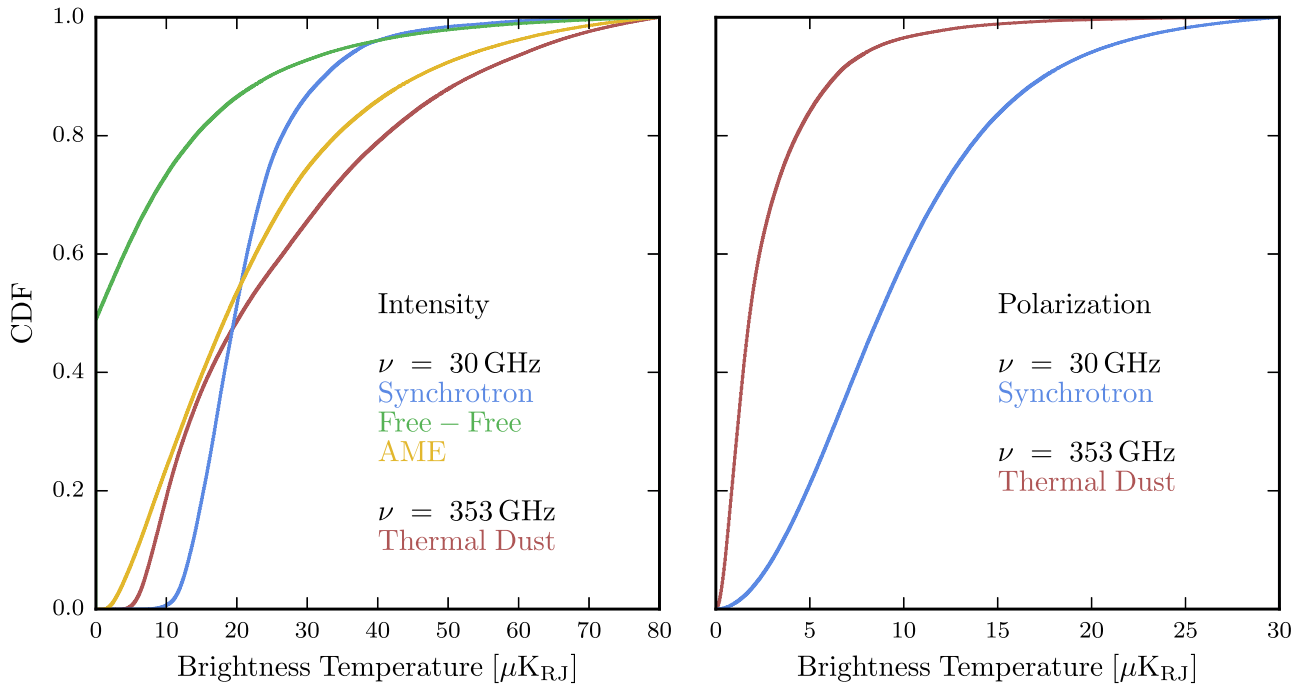


Figure 1. Cumulative distribution function (CDF) of total (left) and polarized (right) intensities for each foreground component in the $|b| > 30^\circ$ sky at $N_{\text{side}} = 256$ as determined by the `Commander` analysis (Planck Collaboration et al. 2016a). We employ these CDFs to select representative amplitudes for each component, which we summarize in Table 1.

Table 1
Assumed Amplitude and Spectral Parameters of the Sky Components in the Simulations, Based on the Distributions from *Planck* Shown in Figure 1

Component	Ref. Freq. [GHz]	I [μK_{RJ}]	Q [μK_{RJ}]	U [μK_{RJ}]	Spectral Parameters
CMB	30	50	0.6	0.6	...
Synchrotron	30	30	10	10	$\beta_s = -1.2$
Free-free	30	30	$\beta_{\text{ff}} = -0.118$
AME	30	30	$\nu_{\text{pk}} = 25$ GHz
Thermal Dust	353	50	3.5	3.5	Various

over a large fraction of the high-latitude sky, as evidenced in Figure 1. A total of 95% of pixels have a 30 GHz free-free amplitude less than $41 \mu\text{K}_{\text{RJ}}$, and so we adopt an amplitude of $30 \mu\text{K}_{\text{RJ}}$. Given the difficulty of separating synchrotron, AME, and free-free with the currently available low-frequency data, we purposefully model these components as equally bright at 30 GHz in total intensity rather than adhere strictly to the `Commander` CDFs.

The `Commander` analysis also extended to polarization, producing full-sky $N_{\text{side}} = 256$ maps of polarized dust and synchrotron emission at $10'$ and $40'$ resolution respectively. We convert the `Commander` Q_ν and U_ν maps to P_ν and plot the resulting histogram over all high-latitude ($|b| > 30^\circ$) pixels in Figure 1. The distributions of dust and synchrotron polarized intensities are $2.7^{+3.1}_{-1.5}$ and $8.9^{+6.8}_{-4.6} \mu\text{K}_{\text{RJ}}$ at 353 GHz and 30 GHz respectively. Assuming a polarization angle of 22.5° for both dust and synchrotron, we adopt $Q_{353 \text{ GHz}} = U_{353 \text{ GHz}} = 3.5 \mu\text{K}_{\text{RJ}}$ and $Q_{30 \text{ GHz}} = U_{30 \text{ GHz}} = 10 \mu\text{K}_{\text{RJ}}$ for the dust and synchrotron, respectively.

The `Commander` analysis assumed that the AME and free-free are unpolarized, following both theoretical and observational arguments which we outline in Sections 2.4 and 2.5. Likewise, we assume these components are unpolarized.

Finally, for the CMB we adopt a temperature of $50 \mu\text{K}_{\text{CMB}}$ and Stokes $Q_\nu = U_\nu = 0.6 \mu\text{K}_{\text{CMB}}$, which are typical values for $\sim 1^\circ$ scales. The adopted amplitudes for all emission components are summarized in Table 1.

While these choices result in a representative, benchmark realization of the foreground components, we note that the ratio of component amplitudes is observed to vary strongly across the sky (Planck Collaboration et al. 2016a). Detailed optimization studies should explore a range of foreground realizations, which is beyond the scope of the present analysis.

2.2. Thermal Dust

In this section, we define a suite of seven dust models that exhibit a range of complex but physically motivated behaviors that could prove challenging for parametric component separation techniques. The dust models, and the parameters that characterize them, are summarized in Table 2, and representative SEDs are plotted in Figure 2.

2.2.1. Generalized Modified Blackbody

Dust grains absorb optical and UV photons and reradiate the absorbed energy in the infrared. The FIR opacity of dust grains

Table 2
Summary of Dust Models

Model	Equations	Components	Constraints
Modified blackbody	(7)–(9)	1	...
Silicate + carbonaceous	(7)–(9)	2	$\arg(A_{d,1}^Q + iA_{d,1}^U) = \arg(A_{d,2}^Q + iA_{d,2}^U)$
MF15	(7)–(9)	2	$\arg(A_{d,1}^Q + iA_{d,1}^U) = \arg(A_{d,2}^Q + iA_{d,2}^U)$
Cloud	(7)–(9)	2	$\beta_{d,1} = \beta_{d,2}$
Silicate + Fe	(7)–(9)	2	$\beta_{d,2} = 0, T_{d,1} = T_{d,2},$ $\arg(A_{d,1}^Q + iA_{d,1}^U) = -\arg(A_{d,2}^Q + iA_{d,2}^U)$
HD	(15)–(17)	2	...
HD with Fe	(18)–(20)	2	...

is often approximated as a power law in frequency,

$$\kappa_\nu^d = \kappa_0 \left(\frac{\nu}{\nu_0} \right)^{\beta_d} \quad (2)$$

where κ_0 is the opacity at reference frequency ν_0 . Assuming this opacity law, the total intensity emitted at frequency ν by a dust grain of temperature T_d is

$$I_\nu^d = A_d^I \left(\frac{\nu}{\nu_0^d} \right)^{\beta_d} B_\nu(T_d), \quad (3)$$

where A_d^I is a dimensionless amplitude parameter and $B_\nu(T)$ is the Planck function. For convenience, we adopt $\nu_0^d = 353$ GHz.

The polarization of the dust emission depends on the angle Ψ between the interstellar magnetic field (along which the grains are aligned) and the line of sight, as well as the material composition and shape of the emitting grains. Variations in the magnetic field direction and grain properties along the line of sight can further influence the polarization signal.

For a single dust grain of temperature T_d and ignoring any frequency dependence of the polarization fraction arising from the frequency-dependence of its dielectric function, we can approximate the polarized intensity as

$$P_\nu^d = A_d^P \left(\frac{\nu}{\nu_0^d} \right)^{\beta_d} B_\nu(T_d), \quad (4)$$

where A_d^P is the specific polarized intensity at a reference frequency ν_0^d . Note that the frequency-independent factor $\sin^2 \Psi$ is subsumed in the amplitude A_d^P .

As the Stokes parameters Q_ν and U_ν are more fundamental observables than P_ν , we construct the amplitudes

$$A_d^Q = A_d^P \cos 2\gamma_d \quad (5)$$

$$A_d^U = A_d^P \sin 2\gamma_d \quad (6)$$

where $0 \leq \gamma_d < \pi$.

Let us now consider emission from an ensemble of grains. If there are N distinct grain types, each with their own β_d , T_d , and polarization angle γ_d , then the total and polarized intensities from the ensemble are given by:

$$I_\nu^d = \sum_{j=1}^N A_{d,j}^I \left(\frac{\nu}{\nu_0^d} \right)^{\beta_{d,j}} B_\nu(T_{d,j}) \quad (7)$$

$$Q_\nu^d = \sum_{j=1}^N A_{d,j}^Q \left(\frac{\nu}{\nu_0^d} \right)^{\beta_{d,j}} B_\nu(T_{d,j}) \quad (8)$$

$$U_\nu^d = \sum_{j=1}^N A_{d,j}^U \left(\frac{\nu}{\nu_0^d} \right)^{\beta_{d,j}} B_\nu(T_{d,j}) \quad (9)$$

$$\gamma_\nu^d = \frac{1}{2} \arg(Q_\nu^d + iU_\nu^d), \quad (10)$$

where \arg denotes the principal value of the argument function with range $[0, 2\pi)$. We denote this model as a “generalized modified blackbody” and employ it as a starting point for a number of simple analytic models which we discuss in the following sections.

2.2.2. Single-component Modified Blackbody

If all interstellar dust has the same β_d (e.g., if all grains are of the same composition) and same temperature T_d (e.g., if all grains are exposed to identical radiation fields), then the dust emission may be modeled as a single-temperature modified blackbody (MBB) described by Equations (7)–(9) with $N = 1$. The polarization angle is frequency independent.

When simulating dust emission with this model, we adopt $\beta_d = 1.6$ and $T_d = 20$ K, consistent with typical values used to fit *Planck* data (Planck Collaboration et al. 2014a, 2016a). To achieve the adopted 353 GHz total and polarized dust intensities, we take $A_d^I = 3.9 \times 10^{-6}$ and $A_d^Q = A_d^U = 2.8 \times 10^{-7}$.

Note that this model is the most commonly used in parametric dust component fits to CMB data, primarily due to its simplicity—there are only two spectral parameters, β_d and T_d , which are the same for the total intensity and polarized emission. We refer to this model as MBB.

2.2.3. Two-component Modified Blackbody

In this model we assume that dust comes in two distinct compositions (e.g., carbonaceous and silicate) with different β_d and T_d . We further assume that these grains are aligned by the same magnetic field and therefore have the same polarization angle. Thus the total and polarized dust emission are given by Equations (7)–(9) with $N = 2$, along with the constraint

$$\arg(A_{d,1}^Q + iA_{d,1}^U) = \arg(A_{d,2}^Q + iA_{d,2}^U), \quad (11)$$

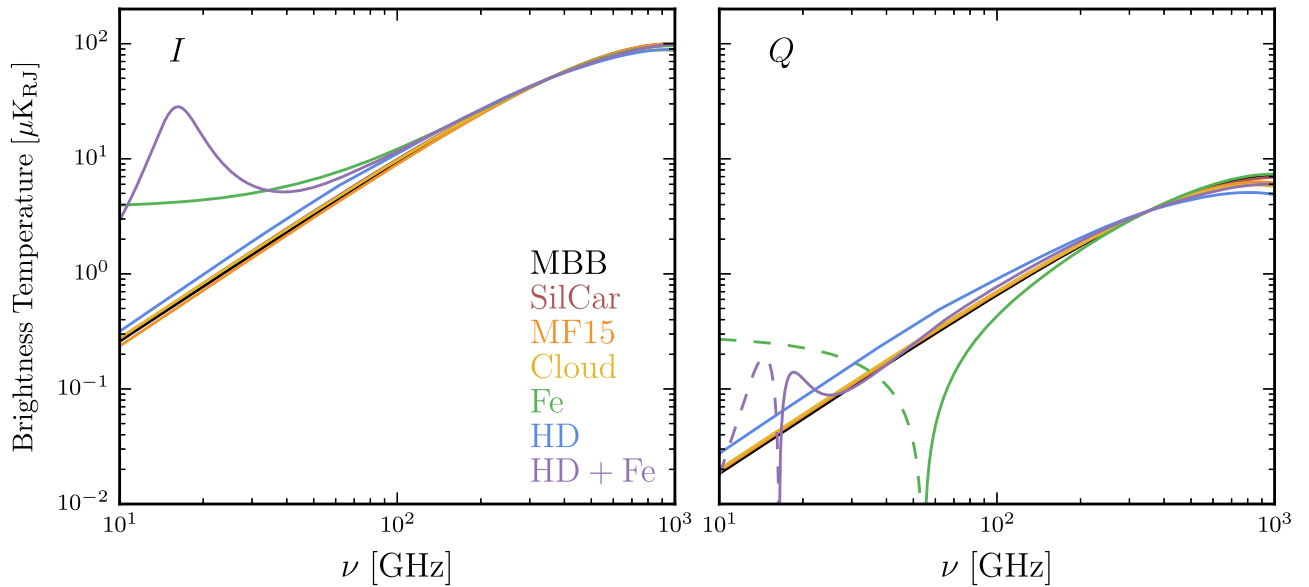


Figure 2. Input SEDs for the various dust models in Stokes I and Q . For most models, Stokes U is nearly identical to Q . The dashed lines indicate where Q is negative.

which enforces consistency between the polarization angles of the two components. The resulting polarization angle is therefore frequency independent.

When performing simulations with this model, we choose two sets of parameters. For the first set, we consider a physically motivated model intended to represent silicate and carbonaceous grains. We choose grain temperatures of $T_{d,1} = 15$ K and $T_{d,2} = 24$ K for silicate and carbonaceous grains, respectively, in agreement with the steady-state temperatures of $0.1 \mu\text{m}$ grains in the model of B. Hensley & B. T. Draine (2018, in preparation). Likewise, we employ spectral indices of $\beta_{d,1} = 1.6$ and $\beta_{d,2} = 1.8$. On the basis of both observational and theoretical results, we assume the carbonaceous grains are unaligned, and so have no net polarization (Chiar et al. 2006; Hoang & Lazarian 2016). Guided by the relative contributions of the silicate and carbonaceous grains to the total infrared emission in the model of B. Hensley & B. T. Draine (2018, in preparation), we adopt $f_I \equiv A_{d,2}^I/A_{d,1}^I = 0.25$, $f_Q \equiv A_{d,2}^Q/A_{d,1}^Q = 0$, and $f_U \equiv A_{d,2}^U/A_{d,1}^U = 0$, $A_{d,1}^I = 3.5 \times 10^{-6}$, and $A_{d,1}^Q = A_{d,1}^U = 3.3 \times 10^{-7}$ to achieve the desired total and polarized intensities. We refer to this model as SilCar.

For the second set, we employ the best-fit parameters of an empirical two-component model developed to fit the observations of Galactic FIR dust emission. Both Finkbeiner et al. (1999), using FIRAS, IRAS, and DIRBE data, and more recently Meisner & Finkbeiner (2015), using *Planck* data, found a statistically significant preference for two-component models over single-component models when fitting thermal dust emission. The best-fit parameters of the Meisner & Finkbeiner (2015) model are $T_{d,1} = 9.75$ K, $T_{d,2} = 15.70$ K, $\beta_{d,1} = 1.63$, $\beta_{d,2} = 2.82$, and $A_{d,1}^I/A_{d,2}^I = 5.35$. To achieve a 353 GHz dust brightness temperature of $50 \mu\text{K}_{\text{RJ}}$, we set $A_{d,1}^I = 9.5 \times 10^{-6}$ and $A_{d,2}^I = 1.8 \times 10^{-6}$. For polarization we adopt $A_{d,1}^Q = A_{d,1}^U = 6.7 \times 10^{-7}$ and $A_{d,2}^Q = A_{d,2}^U = 1.3 \times 10^{-7}$ to yield a total polarized intensity of $5 \mu\text{K}_{\text{RJ}}$ at 353 GHz, and the same ratio between components as in total intensity. We refer to this model as MF15.

Note that we will sometimes use the definitions $\Delta\beta_d \equiv \beta_{d,2} - \beta_{d,1}$ and $f_X \equiv A_{d,2}^X/A_{d,1}^X$ (where $X = I, Q, U$) in our discussion of two-component modified blackbody (2MBB) models.

2.2.4. “Cloud” Model

In the models discussed so far, the polarization angle has been constant with frequency. There are many scenarios in which this will not be the case, such as when both the magnetic field direction and the dust spectrum are varying along the line of sight.

To model this effect, we consider the simple case in which there are two clouds along the line of sight. The dust in each cloud is identical in composition (i.e., has the same β_d), but is heated to different temperatures. Further, the clouds differ in magnetic field direction, yielding different polarization angles. In this case, the total and polarized intensities are given by Equations (7)–(9) with $N=2$ and $\beta_{d,1} = \beta_{d,2}$. Since the dust SEDs of the clouds are different, the polarization angle is frequency dependent.

When simulating data with this model, we choose dust temperatures of $T_{d,1} = 15$ K and $T_{d,2} = 20$ K, well within the observed range of dust temperature variations (e.g., Planck Collaboration et al. 2014a, 2016a; Meisner & Finkbeiner 2015). We assume that the two components are comparably bright in intensity, with $A_{d,1}^I = A_{d,2}^I = 2.4 \times 10^{-6}$. The adopted polarization amplitudes are $A_{d,1}^Q = 1.2 \times 10^{-7}$, $A_{d,1}^U = 2.1 \times 10^{-7}$, $A_{d,2}^Q = 2.4 \times 10^{-7}$, and $A_{d,2}^U = 1.1 \times 10^{-7}$, corresponding to polarization angles of 30° and 12° for components 1 and 2, respectively. The magnetic fields in the clouds are therefore somewhat, but not extremely, misaligned, providing an indication for the typical magnitude of this effect. We refer to this model as Cloud.

2.2.5. Modified Blackbody with Iron Grains (Fe)

Iron is a major constituent of interstellar dust by mass, some of which may be in the form of ferromagnetic nanoparticles (Draine & Hensley 2013). Indeed, embedded metallic iron

nanoparticles have been found in putative interstellar grains collected in the solar system by *Stardust* (Westphal et al. 2014) and *Cassini* (Altobelli et al. 2016). The unique polarization signature of magnetic dipole emission potentially renders these grains an important contaminant for CMB studies—the SED is relatively flat (in CMB temperature units) at low frequencies (see Figure 2). We will consider the case of ferromagnetic inclusions embedded in larger non-magnetic grains.

We can model the composite grains in the context of the generalized modified blackbody model. The non-magnetic matrix of the grain can be described by a simple single-temperature modified blackbody, while the iron inclusions can be modeled as a second modified blackbody of the same temperature but with $\beta = 0$ and polarization angle rotated by 90° with respect to the non-magnetic matrix (Draine & Hensley 2013). Thus, the total and polarized intensities for such composite grains are given by Equations (7)–(9) with $N = 2$, $\beta_{d,2} = 0$, and

$$\arg(A_{d,1}^O + iA_{d,1}^U) = -\arg(A_{d,2}^O + iA_{d,2}^U). \quad (12)$$

We note that when magnetic dipole emission from the iron grains becomes larger than the electric dipole emission from the matrix, the polarization angle flips from $\gamma_{d,1}$ to $\gamma_{d,2} = \gamma_{d,1} + \pi/2$; the two types of emission are polarized perpendicular to one another.

When simulating data with this model, we employ $T_d = 20$ K and $\beta_d = 1.6$, as the presence of iron inclusions does not substantially affect the dust temperature. At the reference frequency of 353 GHz, we assume that the emission from the iron is 5% of that of the silicate matrix in both intensity and polarization, and so $A_{d,2}^I/A_{d,1}^I = 0.05$ and $A_{d,2}^O/A_{d,1}^O = A_{d,2}^U/A_{d,1}^U = -0.05$. This provides a good approximation to the more physical model of B. Hensley & B. T. Draine (2018, in preparation) (see Figure 2). To achieve the adopted 353 GHz total and polarized dust intensities in Table 1, we take $A_{d,1}^I = 3.7 \times 10^{-6}$ and $A_{d,1}^O = A_{d,1}^U = 2.9 \times 10^{-7}$. We refer to this as the Fe model.

2.2.6. A Physical Dust Model (HD and HD+Fe)

The analytic modified blackbody models presented in the previous sections simplify considerably the underlying dust physics. In this section we consider dust models that allow for much greater complexity and which may better reflect the challenges posed by the true dust foreground.

The modified blackbody model of dust emission assumes that the dust opacity κ averaged over all dust sizes and compositions has a frequency dependence that is well-described by a power law of variable index β_d . In a more physical treatment of dust, we compute the dust opacity as a function of grain size for various grain materials based on their complex dielectric function. We designate the opacity of a grain of composition j and radius a at frequency ν as $\kappa_{\nu,j,a}$.

In the modified blackbody model, the dust emission is characterized by a single temperature T_d . In reality, grains of different sizes and compositions will attain different temperatures even when exposed to the same radiation field. Further, very small grains are poorly described by a single temperature. A single optical or UV photon can excite these grains to temperatures exceeding 1000 K, whereas most of the time the grain remains in its very cold ground vibrational state. Hence, it is important to consider the full temperature probability distribution for such grains.

Accounting for these complexities, the total specific intensity from a population of dust grains is

$$I_\nu = \sum_j \int da \frac{dm_j}{da} \int dT \left(\frac{dP}{dT} \right)_{\chi,j,a} \kappa_{\nu,j,a} B_\nu(T), \quad (13)$$

where we have summed over all of the grain compositions j , $(dm_j/da)da$ is the mass in dust of composition j with radius between a and $a+da$, χ is a parameter governing the strength of the radiation field heating the dust, and $(dP/dT)dT$ is the probability of a grain of composition j and radius a in a radiation field χ having temperature between T and $T + dT$. The radiation field is assumed to be a scalar multiple χ of the spectrum of the local interstellar radiation field derived by Mathis et al. (1983), which we set to $\chi = 1$ here.

We employ Equation (13) directly in the context of a physical dust model. By analyzing dust emission and extinction, both total and polarized, from ultraviolet to microwave wavelengths, B. Hensley & B. T. Draine (2018, in preparation) developed a model based on graphitic and silicate grains capable of reproducing the observations. By adopting their opacities, size distributions, and temperature distributions, we can reduce Equation (13) to just a few key parameters:

$$I_\nu = \sum_j A_{d,j}^I I_\nu(j, \chi), \quad (14)$$

where the $A_{d,j}^I$ set the relative contributions of the various grain components as well as the overall amplitude, and the $I_\nu(j, \chi)$ are precomputed quantities based on the physical modeling.

Just as the physical models provide precomputed $I_\nu(\chi)$ based on realistic grain materials, so too can they provide $P_\nu(\chi)$ as long as the grain shapes are specified. In addition to the intrinsic grain properties, the observed Q_ν and U_ν depend on the relative orientation of the interstellar magnetic field and both the line of sight and the reference polarization axes. We subsume these angles into the amplitude parameters A_d^O and A_d^U .

In this work, we explore two such physical dust models. In the first, we assume grains are either carbonaceous or silicate and with no embedded iron inclusions. The total and polarized intensities in this model are

$$I_\nu^{\text{HD}} = A_{d,1}^I I_\nu^{\text{sil}} + A_{d,2}^I I_\nu^{\text{car}} \quad (15)$$

$$Q_\nu^{\text{HD}} = A_{d,1}^O Q_\nu^{\text{sil}} + A_{d,2}^O Q_\nu^{\text{car}} \quad (16)$$

$$U_\nu^{\text{HD}} = A_{d,1}^U U_\nu^{\text{sil}} + A_{d,2}^U U_\nu^{\text{car}}, \quad (17)$$

where the “sil” and “car” superscripts indicate silicate and carbonaceous grains, respectively. The frequency dependence of I_ν and P_ν for each grain type is precomputed in the context of the B. Hensley & B. T. Draine (2018, in preparation) grain model. The relative abundance of silicate and carbonaceous grains is fixed to the default employed by this model and normalized to yield the desired 353 GHz total intensity of $50 \mu\text{K}_{\text{RJ}}$. We note that in this model the carbonaceous grains are unaligned and therefore do not produce polarized emission (i.e., $Q_\nu^{\text{car}} = U_\nu^{\text{car}} = 0$). Therefore, $A_{d,1}^O = A_{d,1}^U$ are adjusted to yield the desired $3.5 \mu\text{K}$ polarized intensity at 353 GHz. We refer to this as the HD model.

Finally, we consider a physical model in which the silicate grains have embedded iron inclusions that constitute 5% of the

grain volume. Denoting this grain type as “sil+Fe,” we have

$$I_{\nu}^{\text{HD+Fe}} = A_{d,1}^I I_{\nu}^{\text{sil+Fe}} + A_{d,2}^I I_{\nu}^{\text{car}} \quad (18)$$

$$Q_{\nu}^{\text{HD+Fe}} = A_{d,1}^Q Q_{\nu}^{\text{sil+Fe}} + A_{d,2}^Q Q_{\nu}^{\text{car}} \quad (19)$$

$$U_{\nu}^{\text{HD+Fe}} = A_{d,1}^U U_{\nu}^{\text{sil+Fe}} + A_{d,2}^U U_{\nu}^{\text{car}}. \quad (20)$$

We note that the signs of $Q_{\nu}^{\text{sil+Fe}}$ and $U_{\nu}^{\text{sil+Fe}}$, and thus the polarization angle, depend on the relative importance of emission from the silicate and iron components. Therefore, this model has a frequency-dependent polarization fraction. We set the parameters of this model in an analogous way to the HD model, assuming the default relative abundances of silicate and carbonaceous grains, assuming that the carbonaceous grains produce unpolarized emission, and normalizing to the desired 50 and $5 \mu\text{K}_{\text{RJ}}$ at 353 GHz in intensity and polarization, respectively. We refer to this model as HD+Fe.

2.3. Synchrotron

As relativistic electrons spiral about the interstellar magnetic field, they emit photons at radio wavelengths known as synchrotron radiation. The spectrum of this radiation depends on the energy spectrum of the electrons. Over the frequency range of interest, the specific intensity is often approximated as a power law:

$$I_{\nu}^s = A_s^I \left(\frac{\nu}{\nu_0^s} \right)^{\beta_s}, \quad (21)$$

where A_s^I is the specific intensity at a reference frequency ν_0^s and β_s is the power-law index.

This parameterization can be extended to polarization by assuming that the polarized intensity spectrum has the same spectral shape as the intensity spectrum:

$$P_{\nu}^s = A_s^P \left(\frac{\nu}{\nu_0^s} \right)^{\beta_s}, \quad (22)$$

where A_s^P is the specific polarized intensity at a reference frequency ν_0^s , taken to be 30 GHz. In this model, we assume the synchrotron polarization angle γ_s is independent of frequency and is equal to the dust polarization angle at 353 GHz.

To produce the adopted total and polarized synchrotron intensities, we adopt $A_s^I = 830 \text{ Jy sr}^{-1}$, $A_s^Q = A_s^U = 280 \text{ Jy sr}^{-1}$, and $\beta_s = -1.2$.

We note that the synchrotron model presented here is a simplification, as it neglects known complexities such as curvature in the synchrotron SED and line-of-sight effects (see Section 2.2.4). We defer investigation of the impact of these complexities to future work, as we focus the present analysis principally on dust emission.

2.4. Anomalous Microwave Emission

Interaction with gas atoms in the interstellar medium causes dust grains to rotate, and grains possessing an electric dipole moment radiate as they rotate (Draine & Lazarian 1998). Ultrasmall grains (radius $a \lesssim 10 \text{ \AA}$) have rotational frequencies of order 30 GHz and thus contribute to Galactic microwave emission. The spinning dust mechanism is believed to be responsible for the AME, a prominent dust-correlated

component of the Galactic microwave emission between $\sim 10\text{--}40$ GHz.

While the AME spectrum has been measured in detail in some Galactic clouds (Planck Collaboration et al. 2014c), the SED of the diffuse Galactic AME appears to peak at lower frequencies and is much more poorly constrained (Miville-Deschênes et al. 2008; Planck Collaboration et al. 2016a). Further, cloud-to-cloud variations suggest large variability in the AME SED, with peak frequencies typically ranging from 20–35 GHz, although as high as $\simeq 50$ GHz (Planck Collaboration et al. 2014c). Theoretical constraints on the spinning dust SED are also weak due to sensitivity of the SED to conditions in the ambient interstellar medium, with different AME SEDs predicted for, e.g., the warm and cold neutral media (Draine & Lazarian 1998; Ali-Haïmoud et al. 2009). The unknown grain size, charge, and dipole moment distributions lead to further uncertainties in the shape of the AME SED (Hensley & Draine 2017).

In light of both the empirical and theoretical uncertainties, we seek a simple parameterization of the AME SED which can capture variability in the amplitude and peak frequency. For the purposes of this work, we adopt the functional form proposed by Draine & Hensley (2012) which includes a parameter governing the peak frequency ν_{pk} of the spectrum:

$$I_{\nu}^{\text{AME}} = A_{\text{AME}}^I \left(\frac{\nu}{\nu_0^{\text{AME}}} \right)^2 \exp[1 - (\nu/\nu_{\text{pk}})^2]. \quad (23)$$

The AME has yet to be detected definitively in polarization, though stringent upper limits ($p \lesssim 1\%$) have been placed in some specific regions (Dickinson et al. 2011; Planck Collaboration et al. 2016c; Génova-Santos et al. 2017) and for the large-scale diffuse emission (Kogut et al. 2007; Planck Collaboration et al. 2015). Theoretical arguments suggest that energy quantization in ultrasmall grains dramatically suppresses the conversion of rotational to vibrational energy, preventing the grains from aligning with the interstellar magnetic field. If this mechanism is acting, the AME will be negligibly polarized ($p \sim 10^{-8}$, Draine & Hensley 2016). We therefore adopt an unpolarized AME component for the purposes of this work on both empirical and theoretical grounds.

To produce the adopted total 30 GHz AME intensity of $30 \mu\text{K}_{\text{RJ}}$, we take $\nu_{\text{pk}} = 25$ GHz and $A_{\text{AME}}^I = 1300 \text{ Jy sr}^{-1}$.

2.5. Free-Free

Free-free emission arises from the acceleration of electrons due to the electric fields of ions. The free-free spectrum is well-known both empirically and theoretically, being well-approximated by a simple power law in the optically thin limit:

$$I_{\nu}^{\text{ff}} = A_{\text{ff}}^I \left(\frac{\nu}{\nu_0^{\text{ff}}} \right)^{-0.12}, \quad (24)$$

where A_{ff}^I is the amplitude at a reference frequency ν_0^{ff} , taken to be 30 GHz.

Free-free emission is inherently unpolarized due to its isotropic nature. However, near the edges of HII regions, Thomson scattering can induce a low level of polarization, estimated to be at the $\sim 10\%$ level when observing in the Galactic plane at high resolution (Keating et al. 1998; Macellari et al. 2011). At the relatively low resolutions analyzed here, the polarization is expected to be much smaller. Indeed,

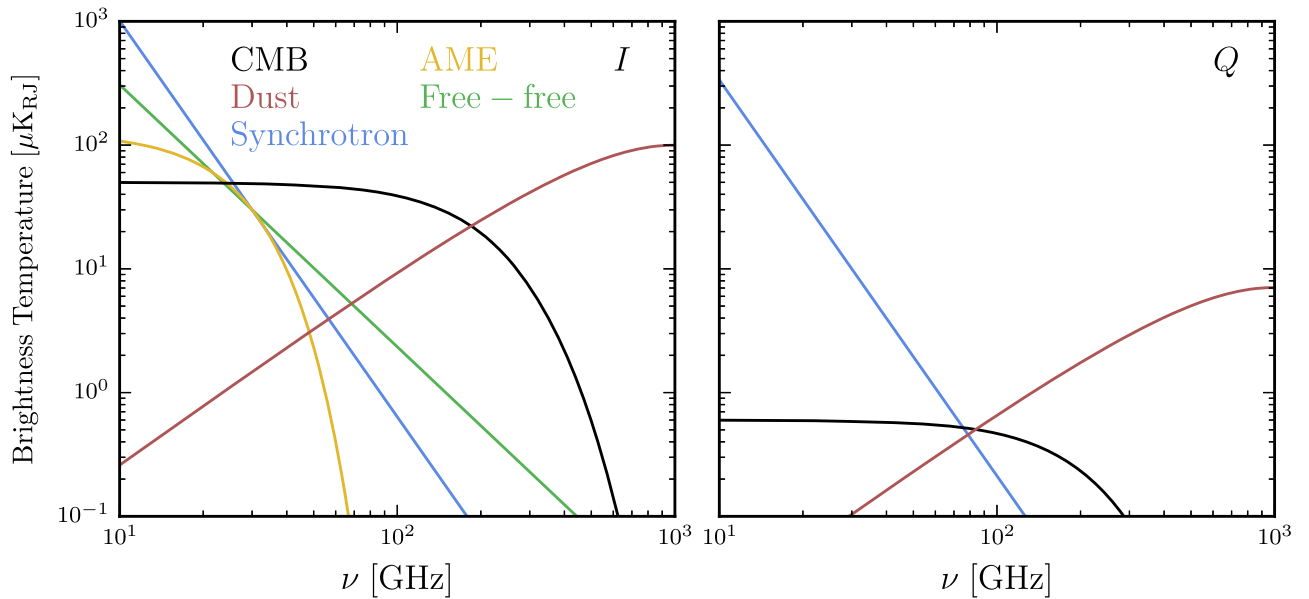


Figure 3. Input SEDs for the CMB, synchrotron, dust (MBB model), AME, and free-free in Stokes I and Q . Stokes U , not shown, is identical to Stokes Q for all components. AME and free-free are assumed to be unpolarized.

empirically, the free-free polarization appears to be less than 1% (Macellari et al. 2011; Planck Collaboration et al. 2016c). Thus, for the purposes of this work, we neglect a potential free-free contribution to the total polarized emission.

To produce the adopted total 30 GHz free-free intensity of $30 \mu\text{K}_{\text{RJ}}$, we take $A_{\text{ff}}^I = 830 \text{ Jy sr}^{-1}$.

2.6. Other Emission Components

The foregoing sections, while extensive, have not been exhaustive of the numerous emission mechanisms in the frequency range of interest. Line emission (notably CO), the cosmic infrared background, and zodiacal light will all contribute to the sky signal observed by a real experiment. As these components are subdominant in intensity and are expected to be negligibly polarized, we neglect them for the purposes of this analysis.

2.7. Summary

We illustrate the SEDs of the emission components employed in this work in Figure 3. In our analysis, all components are kept fixed to these SEDs with the exception of the dust emission. The various dust model SEDs are summarized in Figure 2.

3. Component Separation Simulations

To understand how different types of dust contamination can affect the recovery of CMB temperature and polarization maps, we performed a large number of simulations of single sky pixels across multiple bands. The simulations covered a wide range of assumptions about dust physics and instrumental design, and included multiple (~ 100) noise realizations per dust model and band configuration, so that the statistical properties of the component separation procedure could be studied. In this paper, we consider a single foreground removal algorithm that uses a Bayesian model fitting procedure on multi-band, single-pixel data. Single-pixel foreground fitting is one of the most conservative foreground removal strategies, as

it does not require any assumptions to be made about the spatial distribution of the foreground components. The need to constrain multiple foreground degrees of freedom per pixel is also an important driver of the many-band designs of several future CMB experiments.

In this section, we describe the details of the single-pixel simulations and the Markov chain Monte Carlo (MCMC) model fitting procedure that was subsequently applied to them.

3.1. Single-pixel Simulations

For each dust model, we simulated 100 single-pixel data vectors with different noise realizations over a grid of minimum and maximum frequencies. A seven-band instrumental design was assumed in all cases, with minimum and maximum frequencies of $\nu_{\text{min}} = (20, 30, 40)$ GHz and $\nu_{\text{max}} = (300, 400, 500, 600, 700, 800)$ GHz respectively. The bands were spread across the frequency range with a constant logarithmic interval between them. No bandpass integration was performed, so we have essentially assumed a delta-function bandpass at all frequencies.

The noise at each frequency was obtained by interpolating the CoRE+ Stokes Q noise curve from Remazeilles et al. (2016), shown in Figure 4. We assumed a log-linear extension of the noise curve at frequencies lower than 60 GHz, and a circular, 1° FWHM beam in each frequency channel. This angular resolution should be sufficient for typical B-mode analyses that target the low- ℓ primordial B-mode signal, although higher-resolution experiments have the advantage of being able to estimate the CMB lensing contribution more effectively (Smith et al. 2012). The Q and U polarization channels were assumed to have equal noise levels, with the Stokes I noise rms lower by a factor of $\sqrt{2}$. The noise was assumed Gaussian and uncorrelated between bands. These noise levels are typical of proposed space-based CMB polarization experiments and so provide a benchmark for our analysis. Joint optimization of frequency coverage and the signal to noise in each frequency band is beyond the scope of

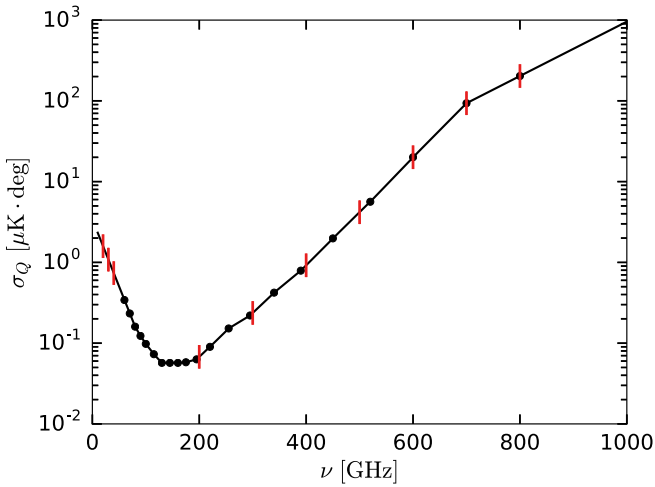


Figure 4. Assumed Stokes Q noise rms as a function of frequency for a future CMB polarization experiment. The U noise rms is identical, and the I noise rms is smaller by a factor of $\sqrt{2}$. We base the noise curve on the CORe+ (Extended) specifications from Remazeilles et al. (2016) (black points), with a log-linear extrapolation to higher and lower frequencies. The minimum/maximum frequencies that were used are marked on the curve as red vertical lines.

this study, but we note that the details of the results presented here would change with the adoption of a different noise curve.

To enable a fair comparison, the model for each component of the sky signal was kept the same between noise realizations and choices of minimum/maximum frequency. The assumptions for each of the non-dust components are summarized in Table 1; I , Q , and U amplitudes were chosen to be representative of the recovered amplitudes from the Planck 2015 Commander–Ruler marginal foreground maps (Planck Collaboration et al. 2016a), smoothed to 1° , and with a sky cut of $|b| \geq 30^\circ$, as discussed in Section 2.1. Spectral parameters were chosen to be broadly consistent with current constraints from the literature. While most of these parameters are expected to vary across the sky (some of them significantly), the chosen values should be representative of a “typical” pixel after excluding the Galactic plane.

Note that we assumed the CMB contribution to be free of lensing contamination, which we would not be able to estimate from a single pixel anyway, and that polarization leakage and other instrumental effects can be ignored.

For the thermal dust component, seven different models were studied, as described in Section 2.2 and summarized in Table 2. The total intensity and polarization amplitudes at 353 GHz were chosen to be the same across all models. The adopted spectral parameters are motivated by the physical effects under study and by models in the literature, as detailed in Section 2.2.

3.2. Single-pixel Foreground Fitting Procedure

With the simulated data in hand, we applied a Bayesian foreground model fitting procedure to each data vector to recover the CMB I , Q , and U amplitudes. This uses an MCMC method to sample from the joint posterior of the amplitude and spectral parameters of all of the component models. The CMB amplitudes can then be obtained by marginalizing over all the other parameters.

Table 3
Prior Ranges of Various Model Parameters

MBB	Min.	Max.
T_d	16	24
β_d	1.4	1.8
2MBB	Min.	Max.
$T_{d,1}, T_{d,2}$	5	30
β_d	1.1	1.8
$\Delta\beta_d$	-1.8	1.8
f_I	0	1
$f_Q = f_U$	-2	2
Other	Min.	Max.
β_s	-1.6	-0.8
$\nu_{\text{peak}}^{\text{AME}}$	15	35
A_X^I	0	∞

Note. Temperatures are in K and frequencies in GHz; all other spectral parameters are dimensionless. Parameters not included in this table were assumed to lie in the range $[-\infty, \infty]$. Note that we set $f_Q = f_U$ to reduce the number of free parameters in the 2MBB model.

The sampling procedure is set up as follows. We first construct a Gaussian likelihood, of the form

$$\log \mathcal{L}(\boldsymbol{\theta}) \sim -\frac{1}{2}(\mathbf{d} - \mathbf{s}(\boldsymbol{\theta}))^T N^{-1}(\mathbf{d} - \mathbf{s}(\boldsymbol{\theta})), \quad (25)$$

where the data vector \mathbf{d} consists of temperature values in each frequency band and polarization. The combined signal vector, \mathbf{s} , is constructed by summing the contributions from each component model in each frequency band and polarization for a given set of parameters, $\boldsymbol{\theta}$. The noise covariance, N , is assumed diagonal, and is identical to the noise covariance that was used when generating the simulations.

We are interested in the effects of assuming the wrong form for the dust model, i.e., what happens when we fit a phenomenological model to one of the more complicated physically motivated dust models described in Section 2.2. For each of the input models described above, we fit two models to the simulated data: the simple MBB model, and the generalized 2MBB model in which both components are assumed to have the same polarization angle. The MBB model has often been used in analyses of CMB data, and only has two free spectral parameters, β_d and T_d . The 2MBB model is expected to better encapsulate complex dust physics, at the cost of introducing several extra parameters. For all other foreground components, we fit the same model that was used to generate the simulated data vector, with all relevant parameters being included in the sampling process.

For each simulated data vector and choice of dust model to be fit, we used the `emcee` affine-invariant ensemble sampler (Foreman-Mackey et al. 2013) to return samples from the joint posterior, using the likelihood from Equation (25). All parameters were taken to have uniform priors over relatively broad ranges, given in Table 3. The sampler was run for 10,000 steps from each of 100 walkers, with a burn-in of 8000 steps being discarded from each. This burn-in period was determined to be sufficient for both the MBB and 2MBB fits, following a series of convergence studies that checked sensitivity of the means and standard deviations of the CMB amplitudes to

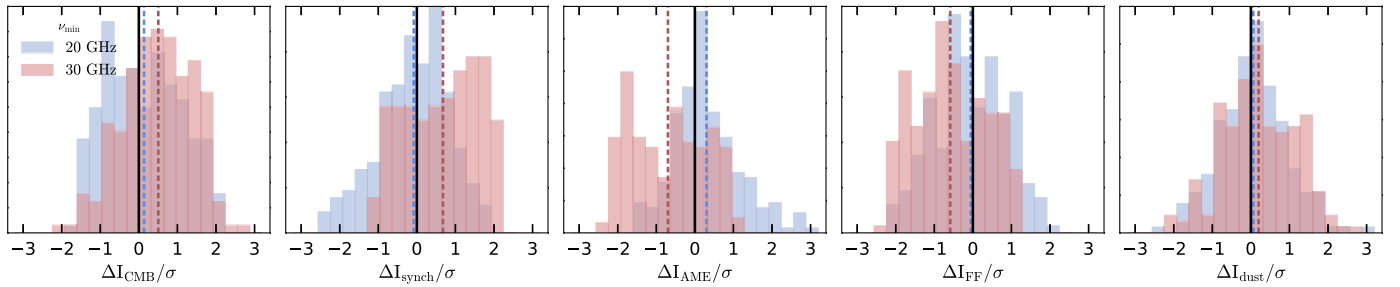


Figure 5. Bias in the maximum likelihood estimates of the Stokes I amplitudes of five components, for a fit to temperature + polarization data containing an MBB dust component, plus CMB, AME, free–free, and synchrotron. Statistics are shown for 200 noise realizations, with $\nu_{\max} = 500$ GHz, and two different values of ν_{\min} . The dashed lines show the mean bias over all 200 realizations. If ν_{\min} is too high, degeneracies between the low-frequency components result in biased amplitude estimates.

burn-in length. Long burn-in periods were needed to allow the walkers time to find the maximum likelihood region on what were sometimes quite complex likelihood surfaces that exhibited multiple degeneracies. Note that, due to the long correlation times in cases where there were strong degeneracies, some of the chains may not have formally converged. Further convergence tests suggest that this is not a serious problem, as only small shifts in the posterior means (a few percent of a standard deviation) were observed when running much longer chains with 100,000 samples per walker. This is subdominant to the error on the posterior mean due to the finite number of noise realizations that were used. Also note that we do not thin the chains, relying instead on combining the samples from the 100 walkers to reduce sample correlation effects.

Each worker was started from an initial position with a small random displacement from the input parameter values used in the simulation. This was done to ensure that the “correct” CMB and foreground parameters would at least be explored in case of multi-modal distributions or parameter degeneracies. The exception was the parameters of the dust model used in the fits, which were set to the same values no matter the input model (but which at least started with the correct amplitude at 353 GHz).

After running and processing the MCMC chains, we calculated a set of summary statistics from the CMB amplitude parameter chains, i.e., the recovered CMB I , Q , and U amplitude pdfs marginalized over all other parameters. This was done for each of the 100 noise realizations for each pair of input and fitting dust models. The distributions of these summary statistics are analyzed in detail in the next two sections. From the MCMC chain for each noise realization, we kept the mean and standard deviation of each parameter, as well as its value at the maximum a posteriori probability (MAP), which is coincident with the point of minimum χ^2 . These were used to calculate the (error-normalized) bias from the true (input) value of each parameter,

$$\frac{\Delta\theta}{\sigma_\theta} \equiv \frac{\bar{\theta} - \theta_{\text{true}}}{\sigma_\theta}, \quad (26)$$

where θ is some parameter, the bar denotes either the mean or MAP estimate from the MCMC chain, and σ_θ is the marginal standard deviation of that parameter estimated from the MCMC chain. In the rest of the paper, we study the distributions of these summary statistics over the 100 noise realizations. Summary statistics that run over noise realizations are denoted

by angle brackets, for example $\langle\Delta\theta\rangle$ denotes the mean of the bias over the 100 noise realizations.

4. Joint Temperature and Polarization Analysis

In this section, we perform an example analysis that uses both temperature (Stokes I) and polarization (Stokes Q and U) data to constrain the CMB and foreground models. Naively, one would expect the inclusion of an I channel to improve the CMB polarization constraints, as it should help to constrain parameters of the foreground models that are common to I and Q/U . As we will see here, however, this is not necessarily the case due to the complexity of the low-frequency temperature foregrounds.

In Figure 5, we show the result of fitting CMB and four foreground components (synchrotron, free–free, AME, and dust) to temperature plus polarization data over seven frequency bands for two different choices of ν_{\min} , up to a ν_{\max} of 500 GHz. In this example, we used a simple MBB model as both the input dust model and the model for the fitting procedure. Since the models to be fitted are the same as the input models, we expect to be able to recover the input model amplitudes and parameters without bias. Figure 5 shows that this is more or less the case for a band configuration with $\nu_{\min} = 20$ GHz, but that there is a significant bias for several components when $\nu_{\min} = 30$ GHz.

This is caused by a degeneracy between several of the low-frequency foreground components, which have similar spectral behaviors around ~ 30 GHz (see Figure 3). Without at least one band below 30 GHz, there is insufficient information to reliably distinguish between synchrotron, free–free, and AME. Incorrect inferences about these components can then lead to them being under- or over-subtracted at higher frequencies, leaving residuals that systematically bias the CMB. In the cases shown in Figure 5, the bias on the CMB Q and U amplitudes is relatively small ($\sim 0.2\sigma$), but this is just the simplest case. For more complex dust models, the bias can be more substantial. Figure 6 compares the bias on the CMB Q amplitude for all seven dust models, for analyses that use temperature and polarization (T+P) information, versus polarization information alone. In four out of the seven cases, the bias is significantly larger for the T + P analysis—and this is for the relatively optimistic scenario where $\nu_{\min} = 20$ GHz. Note that such biases would likely be identifiable, however, as the degeneracies would be apparent in the MCMC chains; the bias shown in Figure 5 is based on a summary statistic, which obscures the existence of degeneracies.

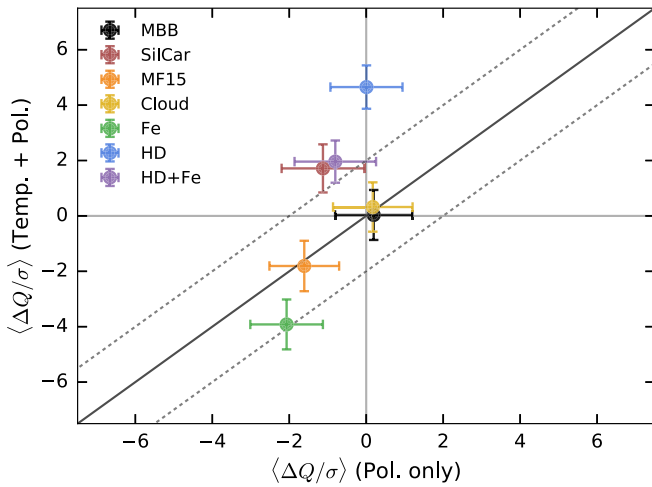


Figure 6. Comparison of the mean bias of the recovered CMB Q amplitude between temperature + polarization and polarization-only analyses. The points shown are for fits with an MBB dust model, and a band configuration with $\nu_{\min}, \nu_{\max} = 20, 500$ GHz. The errorbars show the standard deviation of the bias, while the solid diagonal line marks where the two analyses would have equal bias. Most models show a tendency toward a greater degree of bias when temperature information is included.

We draw two conclusions from this. First, an experiment using joint temperature + polarization foreground fits will likely require frequency coverage below 30 GHz to help break the low-frequency degeneracy. In the absence of low-frequency bands, one can use an amalgamated low-frequency foreground component instead of separate physical synchrotron, free-free, and AME models, as was done in Planck Collaboration et al. (2014b). This is not an unreasonable approach, but runs the risk of introducing subtle model errors, especially if information from the “combined” low-frequency foreground in temperature is used to make inferences about its polarization properties.

Second, if the aim of an experiment is to recover the polarized CMB as accurately as possible, then the value of the additional foreground information provided by the I channel is likely outweighed by the increased complexity of the low-frequency temperature foregrounds. This is especially the case if one considers that we have used quite simplistic models for the synchrotron and AME components here, ignoring potential complications such as spectral curvature, shape of the AME SED, etc.

Additionally, we have assumed the optimistic case that the spectral parameters of the synchrotron and dust components are the same in both temperature and polarization. However, the *Planck* observations indicate that the dust spectral index β_d is systematically different in temperature and polarization (Planck Collaboration et al. 2015). To account for this, one could decouple the spectral parameters used in the I and Q/U channels for these components, but then the I channel provides no extra spectral information about the polarized foregrounds. As such, an analysis that uses only polarization information, without trying to model the temperature components, is likely to be simpler and more robust. This is the strategy that we pursue throughout the rest of the paper.

5. Polarization-only Analysis

In the following sections, we describe the results of fitting the MBB and 2MBB models to the seven dust models

described in Section 2.2, using polarization information alone. In particular, we focus on the bias that is induced in the recovered CMB Q and U amplitudes due to fitting an “incorrect” dust model to the data, and whether the unsuitability of the fitting model can be identified by inspecting the χ^2 goodness-of-fit statistic. While large biases are unwelcome in any situation, it is most problematic if one cannot identify that the results are probably biased and thus take some corrective action (like re-fitting the data with a more complex dust model). As such, models that are strongly biased but still yield low χ^2 values are the most dangerous.

Our main results for the bias and χ^2 are shown in Figures 7 and 8. Numerical values of the bias, χ^2 , and CMB signal-to-noise are given for a few example band configurations in Figure 9.

5.1. Modified Blackbody Models

We begin by fitting the MBB and 2MBB models to simulations that use the simple MBB model described in Section 2.2.2 as the input dust component. As we illustrate in the top panels of Figures 7 and 8, both models recover the input CMB with minimal bias, as expected, with little dependence on the band configuration. The analysis techniques and summary statistics employed therefore appear unbiased and robust. We note however that the extra parameters introduced in the 2MBB model come at the expense of reduced signal to noise on the recovered CMB amplitudes (see Figure 9).

Having performed these consistency tests, we now turn to the two 2MBB models discussed in Section 2.2.3: the model with silicate and carbonaceous grains (SilCar) and the one based on the empirical fits of Meisner & Finkbeiner (2015) (MF15). When fitting either model with the MBB fitting function, the recovered CMB amplitudes are biased by $1-2\sigma$, depending on the band configuration.

In principle, the SilCar model should be perfectly described by the MBB fitting function in polarization only, as the carbonaceous grain component is unpolarized, leaving only a single polarized MBB component associated with the silicate grains. Nevertheless, the fits are biased. This is due to our choice of dust temperature prior. Guided by previous studies fitting the FIR dust emission, we naively selected a uniform prior on the dust temperature of $T_d \in [16, 24]$ K. This is a considerably broader range than was used in the Planck Collaboration et al. (2016a) analysis for example, which assumed a Gaussian prior on T_d with mean 23 K and a standard deviation of 3 K. Our prior is appropriate when fitting the SED in total intensity, but if only the lower-temperature component is polarized, as is the case here (since silicate grains are thought to run cooler than carbonaceous grains; Li & Draine 2001), then it is *not* appropriate for fitting the polarized SED. This is illustrated clearly in Figure 10, in which the maximum a posteriori dust temperatures for each noise realization are shown to cluster at the prior bound.

A similar effect is observed when fitting MF15 with an MBB model. The dust component that dominates at low frequencies has a temperature of 9.75 K, far below the prior bound, and so the same clustering effect occurs at the boundary (see Figure 10). The polarization spectrum is also complicated by the presence of a second component that dominates at high frequencies. The inadequacy of the MBB fit is seen in both the bias in the recovered CMB and, to a lesser degree, in slightly elevated χ^2 values. The bias is reduced when high frequencies, where the

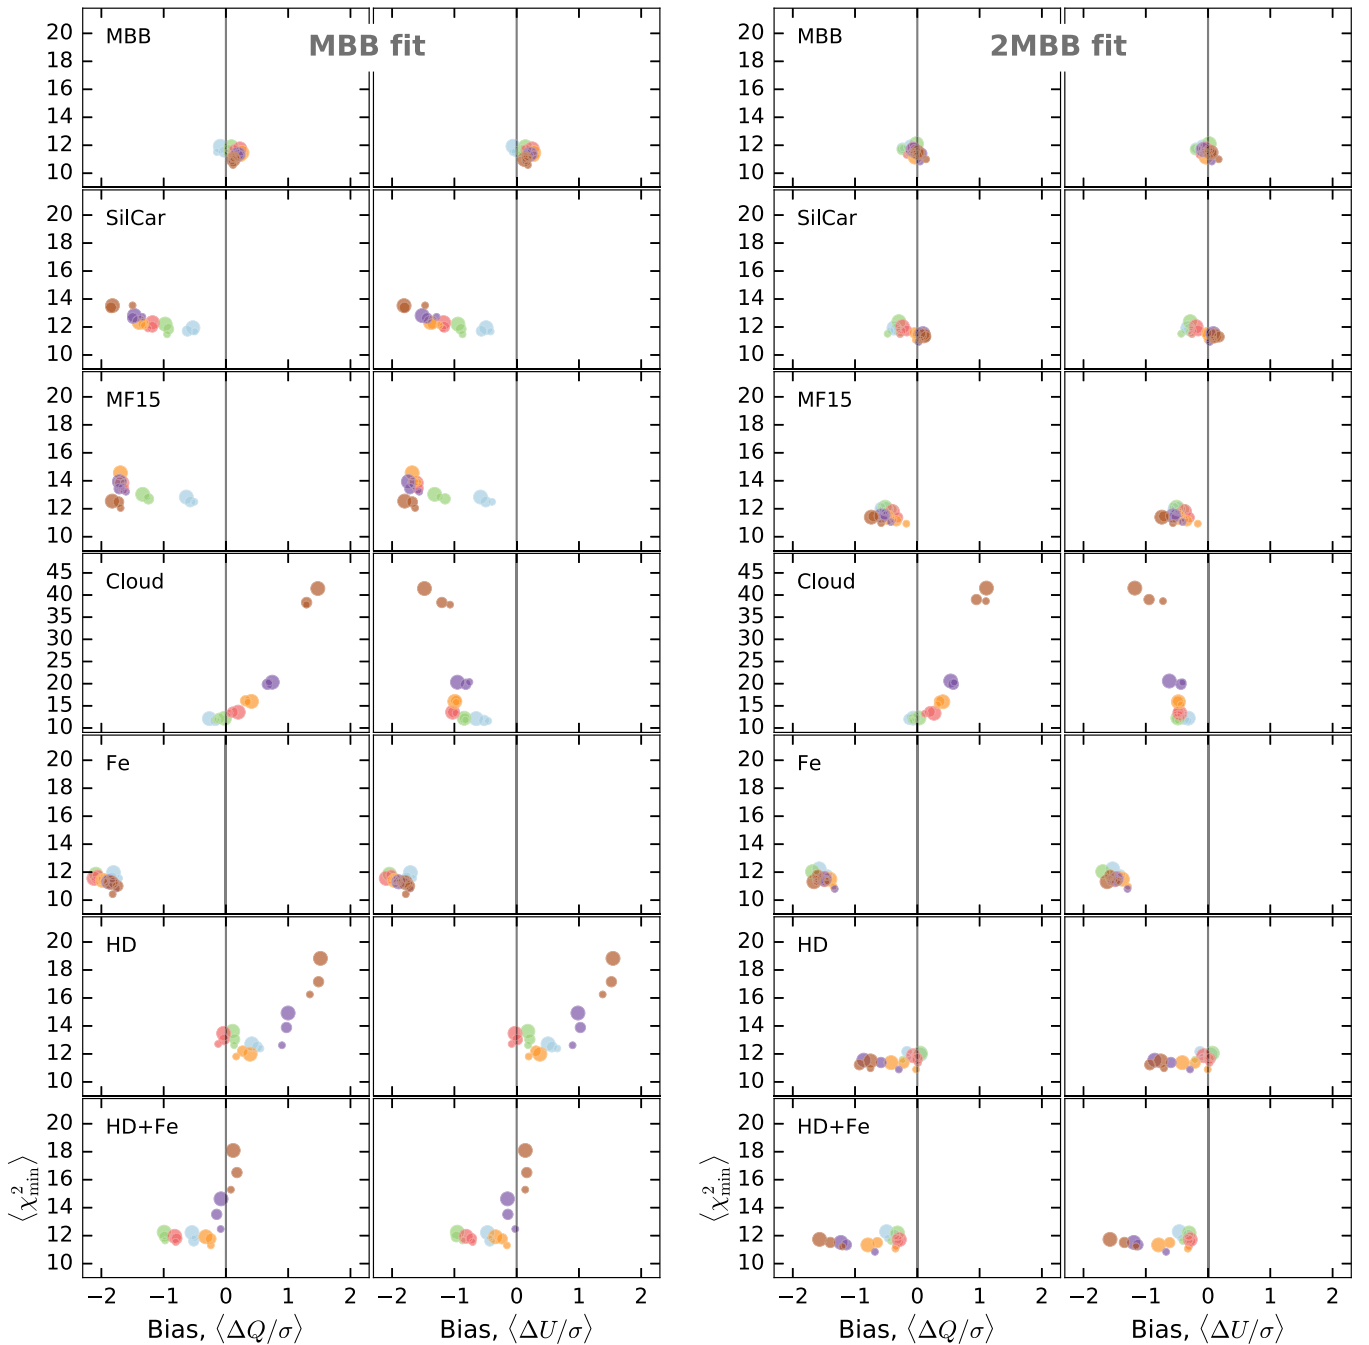


Figure 7. Mean bias of the Stokes Q and U CMB amplitudes (x axis) vs. the median of the minimum χ^2 (y axis) over 100 noise realizations for simulations with several different dust models. The size and color of the points denote different choices of ν_{\min} and ν_{\max} respectively (see the key below), while the left two and right two panels show results for when MBB and 2MBB models are used in the fits respectively. All fits were performed by fitting CMB + synchrotron + dust components to polarization data only. The biases are calculated using the maximum a posteriori probability (minimum χ^2) values for the CMB amplitudes for each noise realization.

complexities induced by the second dust component are most pronounced, are excluded from the analysis.

These problems would be identifiable from inspection of the MCMC chains, and so could be mitigated in a real analysis by expanding the priors. However, this does serve to highlight the risks of using temperature information to infer dust polarization properties. In particular, the dust polarization at CMB frequencies may be dominated by a dust component of significantly lower temperature than that responsible for most of the total intensity at frequencies near the dust peak. Therefore, dust temperature priors based on the total intensity

data may be misleading. Further, at the noise levels expected in future experiments, the effects of the dust temperature on the *shape* of the polarized dust SED (i.e., departures from a pure Rayleigh–Jeans spectrum) are significant enough to impact the recovered CMB even at low frequencies. Joint fits to total and intensity and polarization must be done with care, and models employed in polarization-only fits must be flexible enough to accommodate dust properties significantly different than what are observed in total intensity.

We now turn to fits of these models with a 2MBB model. As illustrated in Figures 7 and 8, the recovered CMB amplitudes

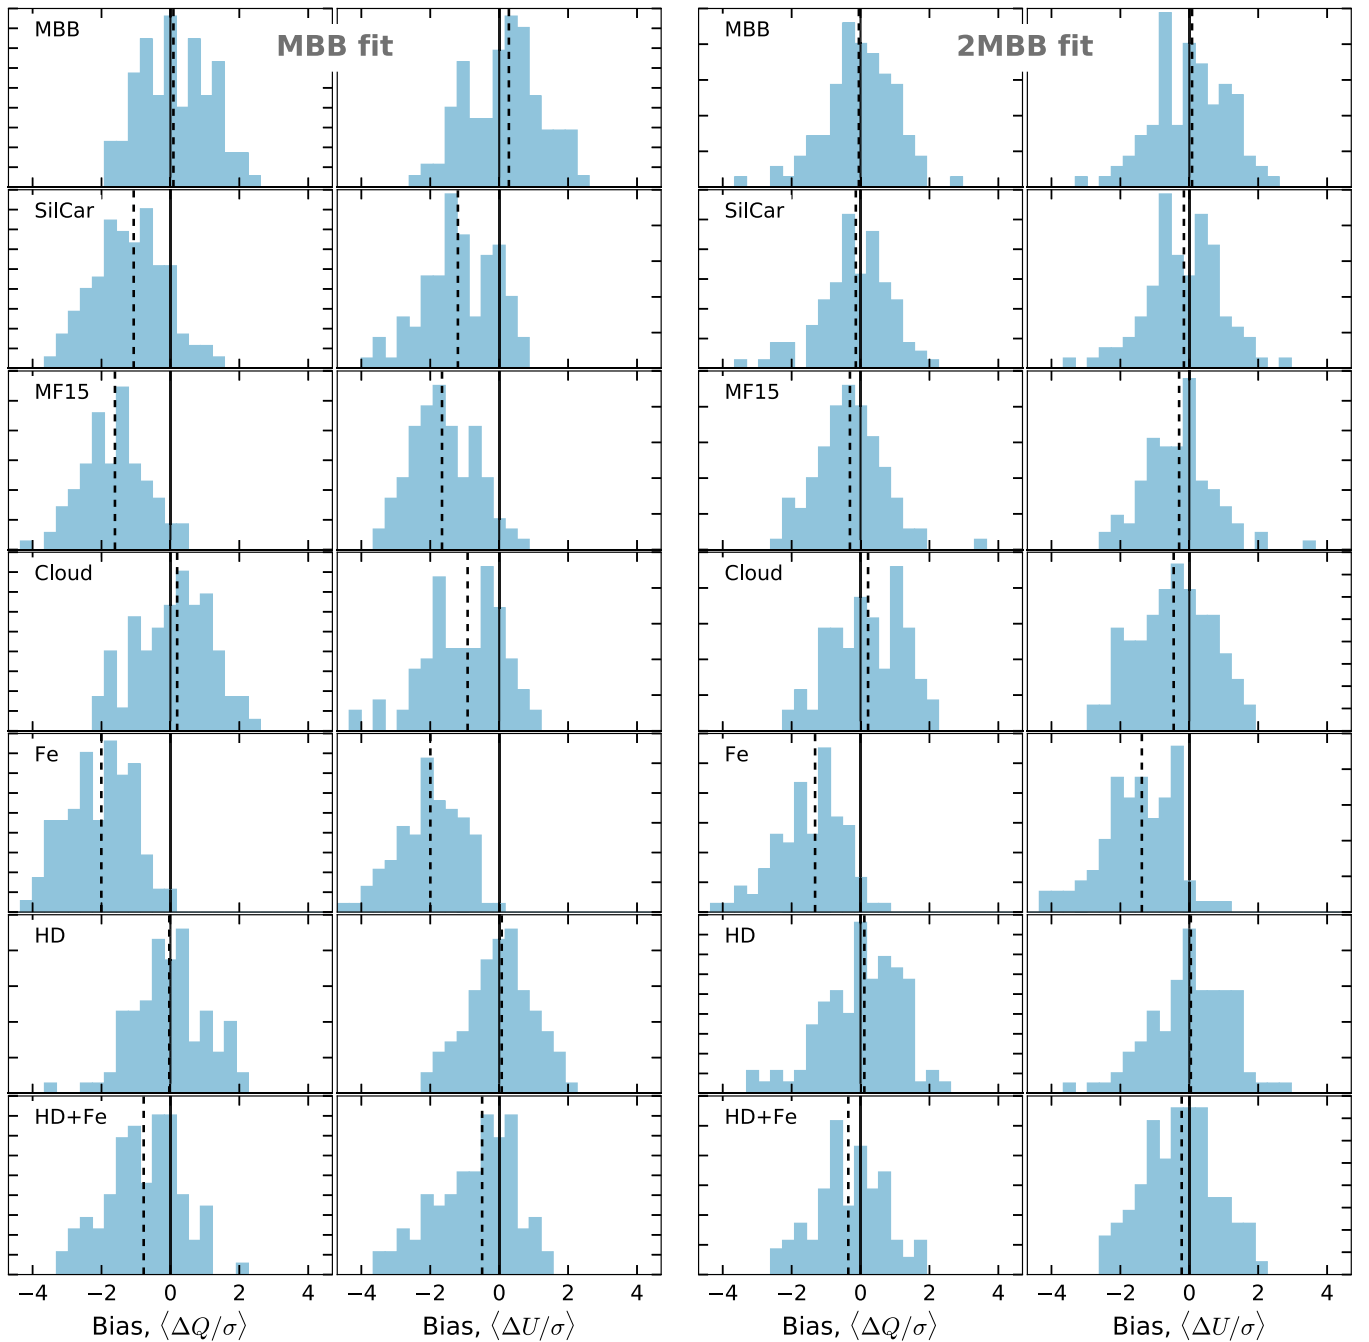


Figure 8. Bias in Stokes Q and U CMB amplitudes over 100 noise realizations, for the same dust models and fitting procedure as in Figure 7. The left two and right two panels again show the results for when MBB and 2MBB models are used to fit the dust component. The dashed vertical lines show the median bias. These results are shown for $\nu_{\min}, \nu_{\max} = 30,500$ GHz.

are much less biased. This is expected, as the SilCar and MF15 models are both fully described by the 2MBB model (which also includes broader priors on its parameters). In the case of MF15, the fits are not completely unbiased. We discuss in more detail the biases that can arise when fitting with the 2MBB model in Section 5.5.

Our conclusion from these results is that future CMB experiments will be sensitive to the additional complexity caused by multiple superposed dust components. Multi-component models of interstellar dust emission are well-motivated by empirical fits to the FIR dust SED (e.g., Finkbeiner et al. 1999;

Meisner & Finkbeiner 2015; Zheng et al. 2017) and by current dust models (e.g., Draine & Li 2007; Siebenmorgen et al. 2014; B. Hensley & B. T. Draine 2018, in preparation; Jones et al. 2017), and so forthcoming experiments should ensure that they can robustly remove multi-component dust contamination, particularly cases in which the components have significantly different properties.

5.2. Models with Line-of-sight Effects

Although the dust models discussed in Section 5.1 each featured multiple dust components, these components were all

MBB fit									
	(20, 300) GHz			(20, 500) GHz			(20, 700) GHz		
	$\langle \Delta Q / \sigma \rangle$	$\langle \chi^2_{\min} \rangle$	$\langle S/N_Q \rangle$	$\langle \Delta Q / \sigma \rangle$	$\langle \chi^2_{\min} \rangle$	$\langle S/N_Q \rangle$	$\langle \Delta Q / \sigma \rangle$	$\langle \chi^2_{\min} \rangle$	$\langle S/N_Q \rangle$
MBB	-0.15	11.5	4.0	+0.10	11.4	5.0	+0.10	10.9	4.7
SilCar	-0.50	11.7	4.2	-1.25	11.9	4.3	-1.34	12.7	3.9
MF15	-0.51	12.5	5.5	-1.64	13.3	6.5	-1.60	13.2	6.4
Cloud	-0.15	11.6	4.1	+0.06	13.4	5.0	+0.69	20.4	5.1
Fe	-1.72	11.6	2.3	-2.10	11.5	3.2	-1.76	10.8	2.9
HD	+0.55	12.4	6.6	-0.12	12.7	7.7	+0.90	12.6	7.8
HD + Fe	-0.50	11.9	4.8	-0.81	11.5	5.0	-0.08	12.5	6.1

2MBB fit									
	$\langle \Delta Q / \sigma \rangle$	$\langle \chi^2_{\min} \rangle$	$\langle S/N_Q \rangle$	$\langle \Delta Q / \sigma \rangle$	$\langle \chi^2_{\min} \rangle$	$\langle S/N_Q \rangle$	$\langle \Delta Q / \sigma \rangle$	$\langle \chi^2_{\min} \rangle$	$\langle S/N_Q \rangle$
	MBB	-0.23	11.8	2.8	-0.17	11.3	2.8	+0.05	10.8
SilCar	-0.43	11.7	2.2	-0.28	11.5	1.6	+0.02	10.9	2.8
MF15	-0.62	12.0	1.6	-0.39	11.2	1.1	-0.43	11.1	2.0
Cloud	-0.10	11.7	2.6	+0.12	13.2	2.4	+0.59	20.3	3.5
Fe	-1.47	11.9	1.6	-1.62	11.3	1.6	-1.33	10.8	1.9
HD	-0.15	11.8	2.1	+0.01	11.4	1.6	-0.30	10.9	2.0
HD + Fe	-0.48	11.8	2.0	-0.34	11.2	1.6	-0.68	10.8	2.0

Figure 9. Summary statistics over 100 noise realizations for 3 selected band configurations, for all 7 input models and both fitting models. The summary statistics are the mean of the error-normalized bias for the CMB Q amplitude; the median of the minimum χ^2 ; and the mean signal-to-noise ratio of the CMB Q amplitude.

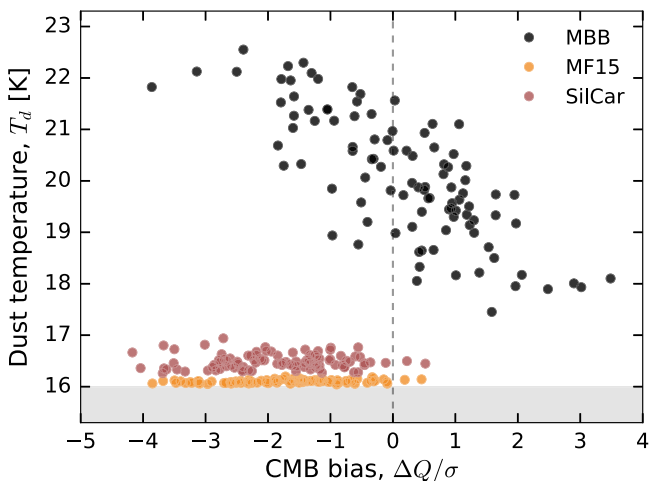


Figure 10. Bias in the maximum a posteriori CMB Q amplitude and dust temperature for MBB fits to polarization-only data, with $\nu_{\min}, \nu_{\max} = 30, 800$ GHz. The gray band shows the region excluded by a prior on T_d . Each point shows the result for one of the 100 noise realizations.

assumed to be aligned by the same magnetic field. In general, the Galactic magnetic field is not uniform along the line of sight, and polarized emission from dust aligned by a field of one direction will add in a vectorial way with emission from dust aligned by a field having a different direction. If the dust properties, such as composition or temperature, are also changing along the line of sight, then the frequency dependence of the resulting polarized emission can be complex and imperfectly correlated between frequencies. As demonstrated by Tassis & Pavlidou (2015) and Poh & Dodelson (2017), failure to account for these line-of-sight effects can lead to biases in CMB fits.

To investigate this in greater detail, we employ the Cloud dust model described in Section 2.2.4, in which the dust emission is assumed to arise from two clouds with different dust temperatures (15 and 20 K) and magnetic field directions. As a result of the components having different polarization angles and SEDs, the ratio of Q_ν to U_ν varies with frequency.

ν_{\min}		ν_{\max}	
● 20 GHz	● 300 GHz	● 600 GHz	
● 30 GHz	● 400 GHz	● 700 GHz	
● 40 GHz	● 500 GHz	● 800 GHz	

The results of fitting the Cloud model with a simple MBB model are presented in Figures 7 and 8. While all frequency configurations are biased to a degree, those with the highest-frequency bands suffer the highest bias. This is because the dust SEDs of the two clouds differ due to their temperature, and temperature effects are most pronounced at high frequencies. Thus, the more the frequency coverage extends to high frequencies, the less the simple model is able to account for the complexity of the emission. However, this is compensated by a dramatically increased χ^2 . Indeed, although models with lower ν_{\max} are less biased, they have more acceptable χ^2 and thus pose a greater potential risk to the analysis.

The effect of the frequency-dependent polarization angle is illustrated in another way in Figure 11, which plots the bias in Q against the bias in U over the 100 noise realizations for the band configuration with $\nu_{\min} = 30$ GHz and $\nu_{\max} = 500$ GHz. When the data are simulated using the MBB or HD models, the biases on the recovered CMB Q and U are roughly equal in each noise realization. For the fits to the Cloud model, however, the Q and U biases are related by a line of different slope and intercept, despite the input CMB signal having $Q = U$ at all frequencies. The Cloud model therefore directly introduces errors on the recovered CMB polarization angle.

The results of fitting the Cloud model with the more complex 2MBB model are also shown in Figures 7 and 8. As we noted in Section 3.2, this model assumes that the two dust components are aligned by the same magnetic field (i.e., $f_Q = f_U$) and thus have the same polarization angle. Due to this assumption, the 2MBB model is unable to account for the frequency-dependent polarization angle that arises in the Cloud model. Allowing $f_Q \neq f_U$ would solve this issue, at the expense of introducing one more free parameter—see Figure 12 for a demonstration. As with the MBB fits, the ($f_Q = f_U$) 2MBB fits recover CMB amplitudes that are biased for all frequency configurations, with those having higher ν_{\max} also having the largest χ^2 . Overall, the bias is less than in the MBB fits due to the greater flexibility of the model.

While dust models with multiple components have been considered in the context of upcoming CMB experiments, the components are almost always assumed to be aligned by the same magnetic field and thus have the same polarization angle (e.g., Armitage-Caplan et al. 2012; Remazeilles et al. 2016; Thorne et al. 2017). However, even for the simple case of a small temperature difference between clouds, we find that line-of-sight effects can induce non-negligible biases in the recovered CMB polarization at the noise levels considered. Additional complications, such as more severely misaligned clouds or different β_d values in each cloud, would exacerbate this effect.

Planck has observed evidence of spatial variation in the spectrum of the polarized dust emission (Planck Collaboration et al. 2017). Depending on the magnitude of the effect, component separation methods relying on persistent spatial correlation of the dust across frequencies may be severely impacted, rendering parametric methods all the more important.

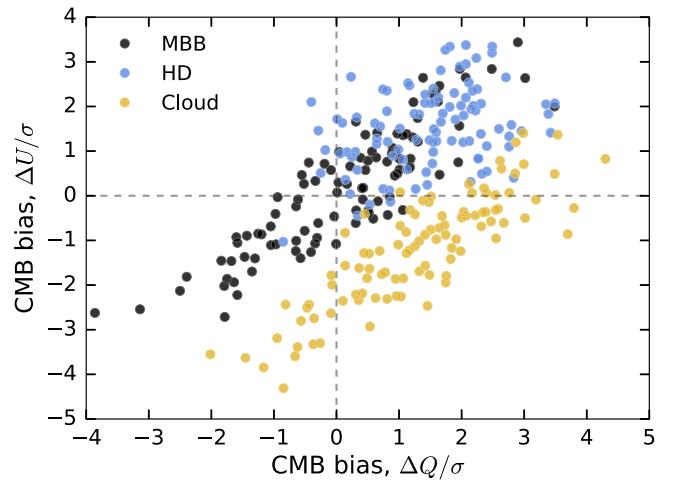


Figure 11. Bias of the CMB Q and U amplitudes for 100 noise realizations with three different input dust models. A simple MBB model was used to fit the data in all three cases. The results shown are for $\nu_{\min}, \nu_{\max} = 30, 800$ GHz, with polarization-only data.

Models such as the Cloud model analyzed here can help test these analysis methods against plausible line-of-sight effects, which we have demonstrated can be at levels of concern for CMB science.

5.3. Models Based on Microscopic Dust Physics

The models discussed so far have been instantiations of the generalized modified blackbody model described in Section 2.2.1. While these models provide a means of investigating a number of physical effects with convenient analytic functions, they likely fail to account in detail for the intrinsic frequency dependence of the dust emission arising from the long-wavelength properties of amorphous materials. Indeed, laboratory studies have demonstrated that interstellar dust analogues can have diverse and complex opacity laws in the FIR that are not well described by simple power laws (Agladze et al. 1996; Demyk et al. 2017a, 2017b). Further, the total microwave emission from dust arises from grains of different sizes, temperatures, and compositions. To evaluate the impact of these complications, we employ the physical dust model of B. Hensley & B. T. Draine (2018, in preparation) (HD), which was described in Section 2.2.6.

We begin by fitting the HD model simulations with the simple MBB model. When high frequencies are included ($\nu_{\max} \geq 700$ GHz), the recovered CMB is biased by more than 1σ , but the χ^2 values are generally elevated, making it easy to identify the poor fit. The remaining configurations have relatively low bias ($<0.5\sigma$) and acceptable χ^2 values. The silicate opacity adopted by B. Hensley & B. T. Draine (2018, in preparation) is not a perfect power law, and thus fitting it as a power law over a wide frequency range induces modeling errors. Additionally, the dust SED is more sensitive to the underlying dust temperature distribution at higher frequencies. Thus, the model failures are more severe for those configurations with coverage extending to high frequencies.

Fits to the HD model with the 2MBB fitting function are presented in the right panels of Figures 7 and 8. In this case, all frequency configurations achieve an acceptable χ^2 goodness of fit, but those with the highest-frequency bands remain significantly biased. Thus, while the extra degrees of freedom

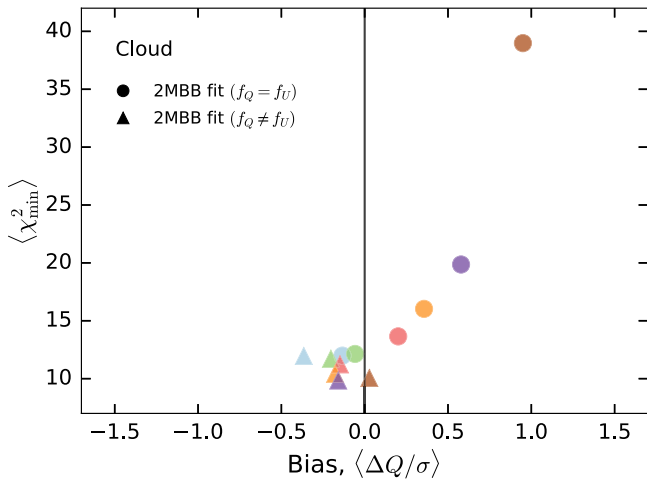


Figure 12. Mean bias in the CMB Q amplitude and median of the minimum χ^2 for 2MBB fits to a Cloud input model, with $f_Q = f_U$ (circles) and $f_Q \neq f_U$ (triangles). The colors denote ν_{\max} , as in Figure 7; only the $\nu_{\min} = 30$ GHz points are shown.

provided by the 2MBB model allow a better fit to be obtained, the inferred parameters are not accurate, and give a misleading picture of the underlying set of components. While having bands at high frequencies is often useful for identifying model errors like this, as seen previously for the Cloud model (Section 5.2), it also requires a reliable parametric model that is valid over a large range of frequencies. The 2MBB model fails in this respect, as it cannot reproduce the non-ideal behavior of the dust opacity in the HD model. Experiments employing high-frequency bands must therefore ensure that their analysis techniques are robust to this effect, as there is a significant risk that it will silently bias the recovered CMB amplitudes.

5.4. Models with an Iron Grain Component

Metallic iron nanoparticles may be a significant component of interstellar dust with strong emissivity in the microwave (Draine & Hensley 2013; B. Hensley & B. T. Draine 2018, in preparation). This component is potentially problematic for CMB experiments due both to its relatively flat spectrum and its polarization signature. In particular, as metallic iron inclusions emit magnetic dipole radiation, their emission is polarized orthogonally to that of the grain in which they are embedded. Therefore, iron inclusions can induce a frequency-dependent change in the dust polarization fraction. We explore their effects in the contexts of two models: one in which the iron component is modeled as a simple graybody (i.e., $\beta_d = 0$; see Section 2.2.5) which we denote “Fe,” and a more physical model which also incorporates magnetic resonance effects (B. Hensley & B. T. Draine 2018, in preparation, see Section 2.2.6), which we denote “HD + Fe.”

In Figure 7, we present the results of fitting the Fe model with both the MBB (left) and 2MBB (right) models. In both cases and for all frequency configurations, the fits are biased by more than 1σ . This is not surprising, as neither model can account for the two orthogonal polarization angles contributing to the total emission. More disconcerting, however, is the overall goodness of the fits in all cases, despite the strong bias. The flat spectrum of the iron grains is readily mimicked by the CMB, leading to substantial bias in the recovered CMB amplitudes.

These effects are illustrated more clearly in Figures 13 and 14, which present the posterior distributions for select model parameters and the best-fit component SEDs, respectively, for an MBB fit to the Fe model with $\nu_{\min} = 30$ GHz and $\nu_{\max} = 500$ GHz. Figure 13 shows that the chains have converged on a seemingly good best-fit model, and that the model parameters are well-determined, with no degeneracies apart from the usual $\beta_d - T_d$ degeneracy. Similarly, Figure 14 shows that the total SED of the best-fit model (blue line) provides an excellent fit to the data, with only a slight deviation from one data point near 80 GHz. Nevertheless, the posterior excludes the true CMB Q amplitude with high significance. This occurs because the decay and sign change of the dust SED induced by the iron component at low frequencies cannot be modeled by the MBB component used in the fits, but is readily compensated by a reduction in CMB amplitude and a shift in the synchrotron spectral index. In fact, with this choice of fitting model, there is no solution that can simultaneously recover the correct CMB amplitudes while also producing a good fit to the data—the assumed MBB dust model is just too inflexible at low frequency. This degeneracy renders iron grains a potentially pernicious foreground for parametric component separation methods, as large biases can be induced without leaving any tell-tale signs, such as a poor goodness of fit.

In the bottom left panel of Figure 7, we present the results of fitting the HD + Fe model with the MBB model. In this case, we find that configurations with $\nu_{\max} \leq 600$ GHz are somewhat biased ($\lesssim 1\sigma$) while having low χ^2 values. In contrast, the two configurations with the highest ν_{\max} have essentially no bias but higher χ^2 . These latter configurations have fewer frequency bands covering the region where the iron emission and the CMB are both significant, and thus the degeneracy between the CMB and the magnetic emission is less severe. However, by having more bands dominated by dust emission that is more complicated than the simple MBB parameterization, the goodness of fit is poorer than for configurations with more bands at lower frequencies. In general, the biases are less than observed in the case of the Fe model as the iron component in this model contributes a smaller fraction of the total polarized signal (see Figure 2).

Finally, in the bottom right panel of Figure 7, we present the results of fitting the HD + Fe model with the 2MBB fitting function. All frequency configurations have a similar and acceptable χ^2 . However, those with high ν_{\max} are significantly biased, whereas the bias is smaller for those experiments concentrated at lower frequencies. The extra degrees of freedom introduced by the 2MBB fit allow much, but not all, of the complexity of the HD + Fe model to be absorbed. Those experiments with high frequencies must correctly model the transition from the dust-polarized emission being dominated by the silicate component to being dominated by the iron component. This cannot be described as the sum of two modified blackbodies having the same polarization angle. In contrast, the low-frequency emission dominated by the iron inclusions can be well-described by a modified blackbody and thus subtracted more effectively. Any residual modeling errors at higher frequencies can be compensated by the second dust component while inducing only minimal bias at frequencies where the CMB dominates.

Iron grains pose a substantial challenge to component separation due to a spectrum that is degenerate with the CMB and their unique polarization signature. Sufficiently

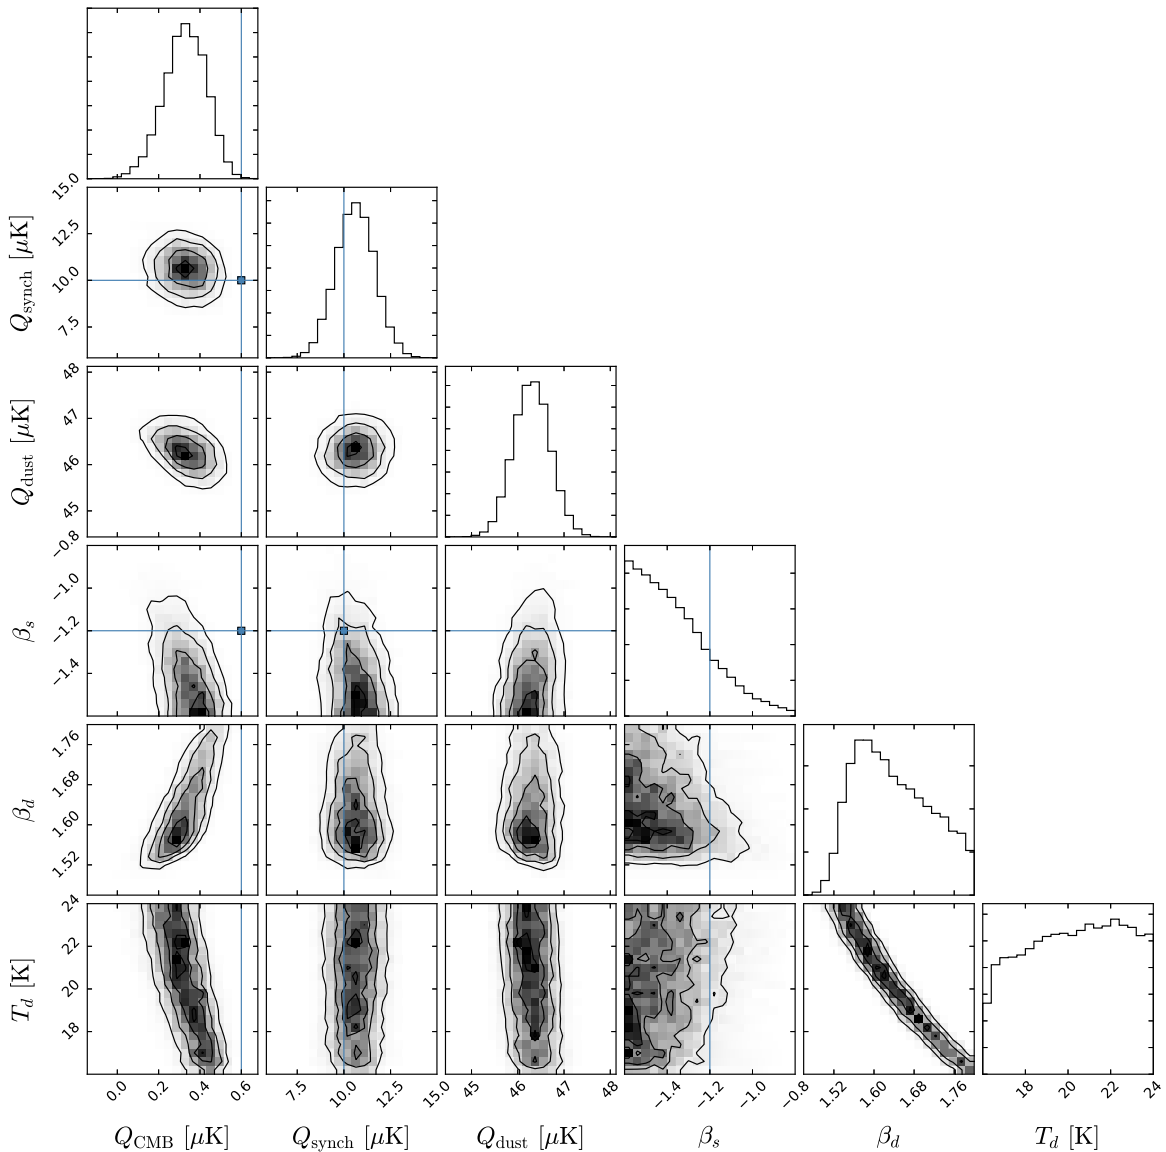


Figure 13. Posterior distributions of selected CMB, synchrotron, and dust parameters for an MBB fit to an Fe input model, for a single noise realization with $\nu_{\min}, \nu_{\max} = 30,500$ GHz. The blue lines show the input values of parameters (where appropriate).

flexible models are required to fit this emission, but the additional complexity must be balanced against an increased ability to fit bad models and still achieve good fits. This problem can be mitigated by increasing the number of frequency bands.

5.5. Causes of Bias in the 2MBB Fits

As discussed above, some of the models return biased results even when the significantly more flexible 2MBB model is used to fit the dust. This is not particularly surprising for the HD and HD + Fe models, which are more complicated than a simple superposition of modified blackbodies (meaning that model errors are expected), while our decision to set $f_Q = f_U$ means that the frequency decorrelation effects in the Cloud model cannot be fully captured. Nevertheless, there are two models that still produce significant biases even though the 2MBB model has the freedom to reproduce them exactly. These are the MF15 and Fe models, which result in $\sim 0.2\text{--}0.8\sigma$ and

$\sim 1.0\text{--}1.5\sigma$ biases respectively, but always with acceptable χ^2 values.

One potential cause of this behavior is the presence of degeneracies between dust model parameters. There is a well-known degeneracy between dust temperature and spectral index even in the MBB case, as illustrated in Figure 13. The more complicated 2MBB parameter space also supports several additional degeneracies, since changes in the parameters of one of the constituent MBB components can often be compensated by changes in the other. As such, entirely different choices of 2MBB parameters can sometimes give almost exactly the same SEDs. If one considers that there is also substantial freedom available in the fits from the CMB and synchrotron components, then it is clear that the parameter space in this fitting problem is likely to be quite complex, with the posteriors possibly exhibiting multiple local maxima. The risk of this happening is even greater for input models which have spectral features similar to the SEDs of other components (e.g., the flat low-frequency part of the Fe SED mimics the CMB). The existence of multiple maxima is problematic; while the correct,

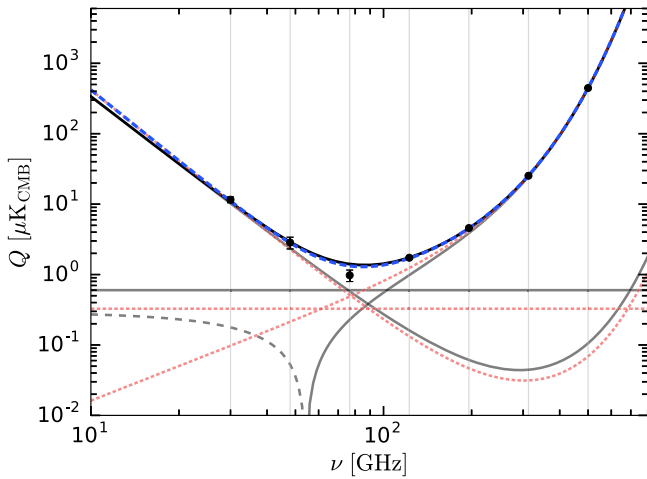


Figure 14. SEDs of the CMB, synchrotron, and Fe dust components, for the input model (gray) and posterior mean values of the fitting parameters shown in Figure 13 (red). The total SEDs are shown in black and blue respectively, while the simulated data are shown as points. The vertical gray lines show the location of each band.

unbiased CMB amplitude may be found in one maximum, other maxima could give a strongly biased result, with no easy way to choose which one is correct.

We do not see obvious multi-modality in the bias histogram of Figure 8, however, and obtain similarly biased results when the chains are started at exactly the correct input parameter values. There is therefore a persistent preference for the incorrect, biased parameter values in some of the 2MBB fits. The recovered best-fit values can change significantly when different prior ranges are chosen for some of the parameters, however. For example, MCMC fits to the MF15 model with the same noise realization but using different prior ranges for β_d resulted in very different best-fit values for the CMB Q amplitude. In our tests, we found biases ranging from 0.3σ for $\beta_d \in [1.4, 1.7]$, to 1.1σ for $\beta_d \in [1.4, 1.8]$, with results in between for different choices of minimum/maximum.

Based on this sensitivity to the priors on some parameters, we suspect that the real issue is that the flat priors we have chosen on the spectral parameters are actually informative. This effect has been pointed out previously in the context of the foreground component separation problem by Eriksen et al. (2008). While flat priors are properly uninformative for linear parameters in the model (e.g., amplitudes or additive constants), they are informative for the nonlinear spectral parameters such as β_d , $\Delta\beta_d$, $T_{d,1}$, and $T_{d,2}$, i.e., some values of these parameters are preferred over others. This can result in biased posterior distributions. When the posterior has multiple maxima, the effect of an informative prior could be enough to single out one of the biased maxima as the global maximum, disfavoring the “true” maximum.

It should be possible to mitigate this effect by using a “Jeffreys prior,” which is proportional to the square root of the determinant of the Fisher matrix. This choice of prior is uninformative even for nonlinear parameters. However, we defer a more detailed study of how to avoid these biases to future work.

5.6. Signal-to-noise Optimization

So far we have focused on how different foreground modeling assumptions and band specifications affect the

likelihood that the recovered CMB polarization will be biased, and whether the bias can be identified. These factors also have an important effect on the sensitivity of the measurements. Since component separation can never be performed perfectly due to the presence of instrumental noise, it will always result in some residual foreground signal being left in the estimated CMB map. This acts as an additional noise term, typically reducing the signal-to-noise ratio compared with an “ideal” scenario with no foregrounds. The component separation procedure can also affect the recovered signal in other ways, such as by over-subtracting the foregrounds (and therefore suppressing the true CMB signal), or by over-fitting, resulting in an artificial reduction in the noise level. As we have seen above, the effectiveness of the component separation procedure can depend strongly on the minimum and maximum frequencies of the bands that are available.

Some example signal-to-noise ratio (S/N) statistics are shown in Figure 9, for three different band specifications and the two different fitting models. There is a clear dependence on the fitting model, with the MBB fits resulting in S/Ns that can be factors of several larger than in the 2MBB fits. This can be understood by counting degrees of freedom—since the 2MBB model has several more parameters than the MBB, most of the information available in the data is being used to constrain those degrees of freedom instead of beating down the noise on the CMB amplitudes. In other words, marginalizing over more nuisance parameters (in this case, parameters of the dust model) reduces the precision with which the CMB signal can be recovered. This is problematic: while simple MBB models are unable to model complex dust scenarios with sufficient accuracy, the more successful 2MBB models are considerably more complex, and result in a significant reduction in sensitivity to the CMB signal. A potential solution is to include several more bands in the instrument design, but we have not considered this possibility here.

Band configurations with higher ν_{\max} have a slight tendency to result in larger S/Ns. This is primarily because the dust signal increases with frequency, and so one can gain a better handle on at least some of the dust parameters by increasing ν_{\max} . This is not always the case, however, as model errors caused by the increasing complexity of the dust spectra with frequency can result in the CMB signal being systematically underestimated, reducing the S/N. A similar effect can also be seen at low frequency for the Fe model, which has a flattened spectrum below ~ 100 GHz that introduces a degeneracy with the primary CMB. Degeneracies can also lead to systematic *over-estimates* of the CMB amplitude, which probably explains why fitting a simple MBB model to the HD model results in a higher S/N than when fitting the MBB model to itself.

As noted previously, increasing ν_{\max} can reduce the risk of bias in some scenarios, while increasing it in others. As such, there is not really an “optimal” band configuration that is robust against any plausible dust complexity. This conclusion carries through when one also considers the S/N—the optimal band configuration depends on what the true dust model is, which we do not know. For a given model the S/N is only mildly sensitive to ν_{\max} , however, suggesting that future experiments should focus on finding configurations that reduce bias first, and then optimize for S/N as a secondary concern.

5.7. Interpretation of Bias Statistics

In order to study a wide range of possible instrumental configurations and dust scenarios (each with a large set of noise realizations), we have restricted ourselves to a single-pixel analysis. The lack of full-sky simulations means that we are unable to translate our findings into implications for measurements of r , the tensor-to-scalar ratio. This is left for future work, which will be able to consider a more restricted set of scenarios informed by the findings in this paper.

Nevertheless, we can make some simplistic statements about the possible effects of the biases that we have identified on a cosmological parameter analysis. Comparing the most idealistic case—MBB fits to a true MBB dust model—with the other scenarios, we see that the expected S/N on the polarized CMB does not change by much except for in the case of the Fe model (see Section 5.6). As such, the basic sensitivity to r is not expected to change by much compared to the simulated MBB component separation analyses that have been performed previously (e.g., Alonso et al. 2017). The problem in this case is the bias, however. For the 2MBB fits, the S/N per pixel is generally at least a factor of 2 worse, which translates to a factor of 4 on the CMB power spectrum.

It is difficult to estimate the effect of the bias without making maps of the residuals, as the contamination of the B-mode signal depends on the pattern of the residuals on the sky, i.e., how much of the residual is in B-modes as opposed to E-modes. Still, a bias of 1σ per pixel means that the power in residual foregrounds is comparable to $\sim 1/(S/N)^2 \approx 1/5^2$ of the power in the CMB polarization on pixel scales. We used pixels of size $\sim 1^\circ$, corresponding to $\ell \sim 200$, where the CMB EE power is $D_{200}^{EE} \approx 1 \mu\text{K}^2$. If one third of the power of the foreground residuals is in B-modes (Planck Collaboration et al. 2016b), and they have an angular power spectrum similar to the dust itself ($D_\ell \sim \ell^{-0.5}$), this would translate to a B-mode contamination of order $\sim 0.1 \mu\text{K}^2$ in angular power around the reionization feature at $\ell \sim 5$. This is about three orders of magnitude larger than the CMB BB power if $r = 10^{-2}$. While we caution against taking this back-of-the-envelope estimate too seriously, typical biases of 1σ per pixel would clearly be a serious matter for forthcoming CMB polarization experiments.

Finally, we note that the precision of our determinations of the mean bias is limited by a couple of effects, of roughly similar magnitude. The first is the limited number of noise realizations that were used to estimate the mean. For a Gaussian random variable, the standard error on the mean scales like σ/\sqrt{N} , where N is the number of realizations. We therefore expect an error on the mean bias of $\sim 0.1\sigma$ for the standard choice of 100 noise realizations. (The number of realizations could be increased, but this would be computationally expensive—the full set of results presented here required several days of wall clock time on 7×64 cores of a shared cluster.) Second, we have used the MAP estimate of the CMB Q and U amplitudes to define the bias, as discussed in Section 3.2. While this seems like the most sensible quantity to use, we could have also chosen other summary statistics, such as the mean. Figure 15 shows how the mean bias changes depending on whether the mean or MAP estimates are used for the bias definition. The difference between the two is around $0.1\text{--}0.5\sigma$.

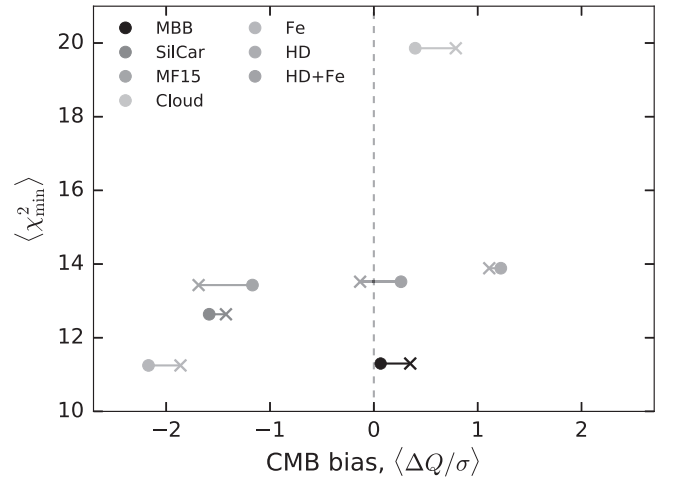


Figure 15. Mean bias in the CMB Q amplitude over 100 noise realizations, shown for two different bias definitions, plotted against the median of the minimum χ^2 . The two bias definitions use the mean (circles) and maximum a posteriori (crosses) values of the marginal posterior distribution for Q from each noise realization. The results are shown for all seven input dust models, assuming polarization-only MBB fits with $\nu_{\min}, \nu_{\max} = 30, 700$ GHz.

6. Discussion and Conclusions

In this paper, we studied the ability of future CMB polarization experiments to remove foreground contamination that includes realistic, physically motivated dust components. Until recently, most studies have assumed that galactic dust emission in the microwave band can be well-described by a simple modified blackbody form with only two spectral parameters. There are many reasons to expect that this may not be the case, however. In Section 2 we enumerated some of the possibilities based on recent *ab initio* attempts to model Galactic dust. These include the following:

1. Separate silicate and carbonaceous dust populations that have significantly different SEDs in temperature and polarization, so that a single MBB model (with spectral parameters that are the same in the Stokes I , Q , and U channels) cannot provide a good description.
2. Models with line-of-sight effects (“cloud”), where the superposition of multiple dust clouds along the line of sight, each with different SEDs and polarization angles, adds complexity to the polarization spectrum, while leaving the total intensity unchanged.
3. The presence of ferromagnetic iron grains, which give rise to magnetic dipole emission that produces flat, CMB-like dust SEDs at low frequencies.
4. More complex “physical” dust models, based on modeling the detailed properties (e.g., size and temperature distributions) and relative abundances of different types of dust grains.

We found that these more realistic dust models can cause significant complications for attempts to remove dust foreground emission and therefore recover the true polarized CMB signal. In particular, biases in the CMB Q and U amplitudes can easily be larger than one standard deviation per 1° pixel when simple MBB models are fitted to polarization-only data. These biases can often be identified by the enlarged residuals (poor χ^2 fits) that they cause, with some important exceptions. For instance, employing MBB

fits, we found that silicate + carbonaceous, MF15 (Meisner & Finkbeiner 2015), and iron-bearing (Fe) dust models can result in $\sim 1\text{--}2\sigma$ biases per pixel while still returning apparently good fits. This is dangerous, even if a good fraction of the residual dust contamination can be removed (e.g., by fitting an angular power spectrum template for the residual during cosmological parameter estimation).

This situation improves somewhat when a more sophisticated two-component MBB model is used. This can model a much wider range of dust physics, including the cloud, iron-bearing, and silicate + carbonaceous scenarios mentioned above, at the cost of introducing three or four additional parameters for a polarization-only analysis, depending on whether the polarization angle is allowed to change with frequency. There are several instances in which the recovered CMB amplitudes remain biased, however. For the Cloud and HD models, the bias can be avoided by restricting the analysis to a relatively low maximum frequency, $\nu_{\text{max}} \lesssim 400$ GHz. This is not a general solution, however: some models, like MF15 and especially Fe, remain significantly biased while still producing a good fit to the data (i.e., low χ^2), regardless of the choice of minimum and maximum frequency. This is a serious cause for concern, as the bias would potentially be difficult to detect. Additionally, the 2MBB parameter space is more complicated than for MBB, and so extra care is needed to handle degeneracies, multi-modality, and biasing due to informative priors.

Contrary to expectations, using temperature data to augment the polarized foreground fits was not helpful even when making the simplifying assumption that the components have the same spectral parameters in both temperature and polarization. We explored this effect in Section 4, where we found that the additional complexity of the low-frequency temperature foregrounds (i.e., including AME and free-free emission) gives rise to degeneracies that have a knock-on effect on the CMB, causing small biases even when the correct (input) dust model is used for the fits. These biases can be mitigated somewhat by including frequency channels below ~ 30 GHz, which help to break degeneracies between the low-frequency temperature components. We note, however, that our assumed AME and synchrotron models are highly idealized, neglecting, for instance, the line-of-sight effects explored in the dust emission which are equally applicable for the low-frequency foregrounds. Multiple sub-30 GHz bands will likely be required to account for the additional spectral degrees of freedom needed to provide a more realistic description of these components.

We did not consider possible spatial variations of foreground spectral parameters in this paper. Instead, we fixed the input parameter values to what should be reasonably typical values at intermediate-to-high latitudes. Future work will consider the effects of spatial variations, in particular frequency decorrelation arising from line-of-sight effects (i.e., due to clouds with a range of different SED parameters being averaged into one beam pointing). Similarly, by performing only a single-pixel analysis, we have excluded the possibility of using morphological information to break degeneracies and identify biases. Dust models that leave significant residuals while still having apparently reasonable goodness-of-fit statistics are certainly dangerous in a single pixel context, but could give rise to residuals with tell-tale morphological signatures that are easier

to identify. A full-sky analysis will be needed to understand whether this is the case or not.



Our conclusions are as follows:

1. A single-temperature modified blackbody model of dust emission, even with spatially varying spectral parameters, does not span the range of possible dust physics to which future CMB experiments will be sensitive. We have provided a library of physically realistic dust models in Section 2.2 that covers a variety of potential complexities. Some subset of these can be used to probe important dust physics effects, such as frequency decorrelation and the presence of ferromagnetic iron grains.
2. High-quality multi-band data are needed to properly constrain the complex foregrounds in Stokes I at low frequencies. Even in the case of idealized low-frequency foregrounds considered in this work, frequency bands below 30 GHz are required to successfully disentangle the contributions of synchrotron, AME, free-free, and the CMB in total intensity. Curvature in the synchrotron SED, variations in the AME SED, and line-of-sight effects are all expected in real data and would exacerbate these degeneracies.
3. Polarization-only analyses using a generalized two-component MBB (2MBB) model are the most robust fitting options that we considered, and should be regarded as a definite improvement to MBB in future analyses. The 2MBB models can still result in model errors and biases, however. Restricting the analysis to relatively low maximum frequencies, $\nu_{\text{max}} \sim 200\text{--}500$ GHz can reduce the risk of bias in many scenarios, as some dust complexities become more severe at high frequencies. On the other hand, we found that high-frequency bands ($\nu \gtrsim 500$ GHz) can be critical for identifying poor model fits in other cases that result in biases on the recovered CMB amplitudes. A detailed study of these trade-offs should be undertaken in the context of specific proposed experiment designs. In particular, future work should examine the information gained by increasing the number of frequency channels.
4. The 2MBB model can be further improved by allowing the two dust components to be polarized with different polarization ($f_Q \neq f_U$ in our notation). Although it introduces an extra parameter into the fit, this more general model is better able to capture the line-of-sight effects that can lead to frequency decorrelation.

The Python code and data files for the summary statistics calculated as part of this analysis are available to download from <http://philbull.com/singlepixel/>.

We are grateful to Geoff Bryden, Curt Cutler, Clive Dickinson, Tim Eifler, Hans Kristian Eriksen, Charles Lawrence, Graca Rocha, and especially Jeff Jewell, for valuable discussions and comments. We thank the anonymous referee for feedback which improved the manuscript. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. P.B.'s research was supported by an appointment to the NASA Postdoctoral Program at the Jet Propulsion Laboratory, California Institute of Technology, administered by Universities Space Research Association under contract with NASA.

ORCID iDs

Brandon S. Hensley  <https://orcid.org/0000-0001-7449-4638>
 Philip Bull  <https://orcid.org/0000-0001-5668-3101>

References

- Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, arXiv:1610.02743
- Agladze, N. I., Sievers, A. J., Jones, S. A., Burlitch, J. M., & Beckwith, S. V. W. 1996, *ApJ*, **462**, 1026
- Ali-Haïmoud, Y., Hirata, C. M., & Dickinson, C. 2009, *MNRAS*, **395**, 1055
- Alonso, D., Dunkley, J., Thorne, B., & Naess, S. 2017, *PhRvD*, **95**, 043504
- Altobelli, N., Postberg, F., Fiege, K., et al. 2016, *Sci*, **352**, 312
- Armitage-Caplan, C., Dunkley, J., Eriksen, H. K., & Dickinson, C. 2012, *MNRAS*, **424**, 1914
- BICEP2 Collaboration, Ade, P. A. R., Aikin, R. W., et al. 2014, *PhRvL*, **112**, 241101
- BICEP2/Keck Collaboration, Planck Collaboration, Ade, P. A. R., et al. 2015, *PhRvL*, **114**, 101301
- Chiar, J. E., Adamson, A. J., Whittet, D. C. B., et al. 2006, *ApJ*, **651**, 268
- Chluba, J., Hill, J. C., & Abitbol, M. H. 2017, *MNRAS*, **472**, 1195
- Demyk, K., Meny, C., Leroux, H., et al. 2017a, *A&A*, **606**, 50
- Demyk, K., Meny, C., Lu, X.-H., et al. 2017b, *A&A*, **600**, A123
- Dickinson, C., Peel, M., & Vidal, M. 2011, *MNRAS*, **418**, L35
- Draine, B. T., & Hensley, B. 2012, *ApJ*, **757**, 103
- Draine, B. T., & Hensley, B. 2013, *ApJ*, **765**, 159
- Draine, B. T., & Hensley, B. S. 2016, *ApJ*, **831**, 59
- Draine, B. T., & Lazarian, A. 1998, *ApJ*, **508**, 157
- Draine, B. T., & Li, A. 2007, *ApJ*, **657**, 810
- Eriksen, H. K., Jewell, J. B., Dickinson, C., et al. 2008, *ApJ*, **676**, 10
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, **524**, 867
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
- Génova-Santos, R., Rubiño-Martín, J. A., Peláez-Santos, A., et al. 2017, *MNRAS*, **464**, 4107
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, **622**, 759
- Hensley, B. S., & Draine, B. T. 2017, *ApJ*, **836**, 179
- Hoang, T., & Lazarian, A. 2016, *ApJ*, **831**, 159
- Jones, A. P., Köhler, M., Ysard, N., Bocchio, M., & Verstraete, L. 2017, *A&A*, **602**, A46
- Keating, B., Timbie, P., Polnarev, A., & Steinberger, J. 1998, *ApJ*, **495**, 580
- Kogut, A., Dunkley, J., Bennett, C. L., et al. 2007, *ApJ*, **665**, 355
- Kogut, A., & Fixsen, D. J. 2016, *ApJ*, **826**, 101
- Li, A., & Draine, B. T. 2001, *ApJ*, **554**, 778
- Macellari, N., Pierpaoli, E., Dickinson, C., & Vaillancourt, J. E. 2011, *MNRAS*, **418**, 888
- Mathis, J. S., Mezger, P. G., & Panagia, N. 1983, *A&A*, **128**, 212
- Matsumura, T., Akiba, Y., Borrill, J., et al. 2014, *JLTP*, **176**, 733
- Meisner, A. M., & Finkbeiner, D. P. 2015, *ApJ*, **798**, 88
- Miville-Deschênes, M.-A., Ysard, N., Lavabre, A., et al. 2008, *A&A*, **490**, 1093
- Planck Collaboration, Abergel, A., Ade, P. A. R., et al. 2014a, *A&A*, **571**, A11
- Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016a, *A&A*, **594**, A10
- Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016b, *A&A*, **586**, A133
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014b, *A&A*, **571**, A12
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014c, *A&A*, **565**, A103
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016c, *A&A*, **594**, A25
- Planck Collaboration, Ade, P. A. R., Alves, M. I. R., et al. 2015, *A&A*, **576**, A107
- Planck Collaboration, Aghanim, N., Ashdown, M., et al. 2017, *A&A*, **599**, A51
- Poh, J., & Dodelson, S. 2017, *PhRvD*, **95**, 103511
- Remazeilles, M., Banday, A. J., Baccigalupi, C., et al. 2017, arXiv:1704.04501
- Remazeilles, M., Dickinson, C., Eriksen, H. K. K., & Wehus, I. K. 2016, *MNRAS*, **458**, 2032
- Siebenmorgen, R., Voshchinnikov, N. V., & Bagnulo, S. 2014, *A&A*, **561**, A82
- Smith, K. M., Hanson, D., LoVerde, M., Hirata, C. M., & Zahn, O. 2012, *JCAP*, **6**, 014
- Stompor, R., Errard, J., & Poletti, D. 2016, *PhRvD*, **94**, 083526
- Tassis, K., & Pavlidou, V. 2015, *MNRAS*, **451**, L90
- Thorne, B., Dunkley, J., Alonso, D., & Naess, S. 2017, *MNRAS*, **469**, 2821
- Westphal, A. J., Stroud, R. M., Bechtel, H. A., et al. 2014, *Sci*, **345**, 786
- Zheng, H., Tegmark, M., Dillon, J. S., et al. 2017, *MNRAS*, **464**, 3486