

UC Berkeley

UC Berkeley Previously Published Works

Title

Delensing the CMB with the cosmic infrared background: the impact of foregrounds

Permalink

<https://escholarship.org/uc/item/8cg1r5q7>

Journal

Monthly Notices of the Royal Astronomical Society, 514(4)

ISSN

0035-8711

Authors

Baleato Lizancos, Antón
Challinor, Anthony
Sherwin, Blake D
[et al.](#)

Publication Date

2022-07-12

DOI

10.1093/mnras/stac1705

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

Delensing the CMB with the cosmic infrared background: the impact of foregrounds

Antón Baleato Lizancos,^{1,2*} Anthony Challinor,^{1,2,3} Blake D. Sherwin,^{2,3}
Toshiya Namikawa³

¹*Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, UK*

²*Kavli Institute for Cosmology Cambridge, Madingley Road, Cambridge, CB3 0HA, UK*

³*DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge, CB3 0WA, UK*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The most promising avenue for detecting primordial gravitational waves from cosmic inflation is through measurements of degree-scale CMB *B*-mode polarisation. This approach must face the challenge posed by gravitational lensing of the CMB, which obscures the signal of interest. Fortunately, the lensing effects can be partially removed by combining high-resolution *E*-mode measurements with an estimate of the projected matter distribution. For near-future experiments, the best estimate of the latter will arise from co-adding internal reconstructions (derived from the CMB itself) with external tracers such as the cosmic infrared background (CIB). In this work, we characterise how foregrounds impact the delensing procedure when CIB intensity, *I*, is used as the matter tracer. We find that higher-point functions of the CIB and Galactic dust such as $\langle BEI \rangle_c$ and $\langle EIEI \rangle_c$ can, in principle, bias the power spectrum of delensed *B*-modes. To quantify these, we first estimate the dust residuals in currently-available CIB maps and upcoming, foreground-cleaned Simons Observatory CMB data. Then, using non-Gaussian simulations of Galactic dust – extrapolated to the relevant frequencies, assuming the spectral index of polarised dust emission to be fixed at the value determined by Planck – we show that the bias to any primordial signal is small compared to statistical errors for ground-based experiments, but might be significant for space-based experiments probing very large angular scales. However, mitigation techniques based on multi-frequency cleaning appear to be very effective. We also show, by means of an analytic model, that the bias arising from the higher-point functions of the CIB itself ought to be negligible.

Key words: cosmic background radiation – gravitational lensing; weak – polarisation – infrared: diffuse background – dust, extinction

1 INTRODUCTION

The cosmological model best fitting observations features a period of accelerated expansion at very early times known as *cosmic inflation*. In general, the physics of inflation predicts the generation during this period of a stochastic background of gravitational waves (tensor perturbations) and fluctuations in the density (scalar perturbations). Although difficult to detect directly, primordial gravitational waves are expected to have left an imprint on the temperature and polarisation of the cosmic microwave background (CMB; Polnarev 1985).

Scalar perturbations are now known to be responsible for the vast majority of the anisotropy pattern found in the temperature (*T*) and gradient-like *E*-mode component of polarisation. The variance associated with the scalar modes makes it impossible to detect the much smaller contribution from tensor perturbations in temperature or *E*-mode polarisation data. However, it is possible to form a curl-like component (*B*) which, at recombination (and in linear theory), is sourced only by

* E-mail: a.baleatolizancos@berkeley.edu

tensor perturbations (Seljak & Zaldarriaga 1997; Kamionkowski et al. 1997). A detection of a primordial B -mode is therefore widely regarded as a conclusive test of inflation, with the amplitude of the primordial power spectrum being directly related to the energy scale at which inflation occurs. This amplitude is usually parametrised by the tensor-to-scalar ratio, r . Currently, the most stringent experimental upper bounds are those of BICEP2 Collaboration & Keck Array Collaboration (2018), which place $r < 0.06$, or $r < 0.044$ in combination with Planck (Tristram et al. 2020), both with 95% confidence.

Efforts to constrain r below this level face a significant challenge: gravitational lensing of the CMB by the large-scale distribution of matter in the Universe converts part of the primordial E -mode into B -modes (Zaldarriaga & Seljak 1998). The angular power spectrum of the lensing B -modes resembles that of white noise with $\Delta\mathcal{P} = 5 \mu\text{K arcmin}$ on the large angular scales where the primordial signal is expected to be strongest (an effect first detected by SPT Collaboration 2013). Consequently, the variance due to lensing becomes a significant hurdle for detecting small primordial signals. In fact, the lensing noise is at a level where it is already comparable to current experimental noise levels, and thus upcoming surveys, in addition to improving experimental sensitivity, ought to focus on removing the lensing noise.

The removal of lensing B -modes cannot be achieved via multi-frequency cleaning, for lensing is achromatic. However, the lensing B -modes can be estimated by combining observations of E -mode polarisation with estimates of the lensing deflections on the sky in a way that mimics how the lensing operation occurs in nature. This can then be used to remove the lensing effects in a procedure known as *delensing*. This technique has already been demonstrated on real data by several teams (Carron et al. 2017; Manzotti et al. 2017; Planck Collaboration 2018c; Adachi et al. 2019; Han et al. 2020; BICEP/Keck & SPTpol Collaborations 2020).

In order to undo the deflections induced by lensing, an estimate of the projected matter distribution on the sky – which determines the lensing convergence, κ – is required. For the sensitivities and resolution of current and near-future CMB experiments, the best possible estimate is obtained from tracers external to the CMB such as galaxy surveys or the cosmic infrared background (Smith et al. 2012; Sherwin & Schmittfull 2015), as these provide access to small-scale lenses at high redshift that internal reconstruction from the CMB itself (via the quadratic estimators of Okamoto & Hu 2003 or maximum-likelihood approaches in the style of Hirata & Seljak 2003) cannot yet extract with sufficient precision.

In recent years, one tracer has emerged as being particularly useful for delensing – the cosmic infrared background (CIB). The CIB is the integrated emission from dust heated by UV starlight in faraway, star-forming galaxies. It originates from a broad redshift range centred around $z \approx 1-2$, though extending as far back as $z \approx 4$ (Béthermin et al. 2012; Viero et al. 2013a,b). This distribution overlaps extensively with the lenses responsible for CMB lensing; so much so that the CIB and CMB lensing convergence are up to 80% correlated (Song et al. 2003; Holder et al. 2013; Planck Collaboration 2014b). The constituent dust is at a temperature of a few tens of Kelvin so CIB emission peaks in the sub-millimetre part of the electromagnetic spectrum, where it was first detected by Puget et al. (1996), and dominates over other components (at least at high Galactic latitudes) in the high-frequency channels of CMB experiments, with lower frequencies probing higher-redshift sources.

Recent observations at 353, 545 and 857 GHz made by the High-Frequency Instrument (HFI) aboard the *Planck* satellite have led to significant improvements in our understanding of the CIB thanks to the instrument’s high sensitivity and fine angular resolution and large sky coverage. On the one hand, these have allowed for a refinement of models (Planck Collaboration 2011, 2014c; Serra et al. 2014; Mak et al. 2017; Maniyar et al. 2018), and on the other, high-fidelity maps of the CIB anisotropies over large fractions of the sky have been produced (Planck Collaboration 2016e; Lenz et al. 2019). These maps are some of the best tracers of the large-scale structure of the Universe at high redshift that have been produced to date. As such, they hold great promise for upcoming implementations of delensing, particularly in combination with other tracers (Yu et al. 2017). CIB maps from Planck were used by Larsen et al. (2016) in the first demonstration of delensing, and more recently BICEP/Keck & SPTpol Collaborations (2020) used them to show the first improvement in r constraints thanks to delensing. Meanwhile, Manzotti et al. (2017) used *Herschel* CIB data to achieve the first statistically-significant delensing of B -modes.

The aim of delensing is to reduce the variance on estimates of the B -mode power spectrum by subtracting the specific realisation of lensing B -modes present on the sky. The more effective this procedure, the more attention one must pay to potential systematic errors. In spite of this, the impact of foreground emission on delensing has remained largely unexplored, with the exception of Beck et al. (2020)’s recent study of the impact of residual foregrounds on internal delensing of B -modes. In this paper, we focus on the impact of foregrounds when using the CIB to delens B -modes. The principal worry is that, as part of the delensing procedure, residual amounts of dust or CIB emission left over in the CMB maps can couple with residual dust or the CIB in the matter tracer maps. The problem is exacerbated by the fact that both the CIB and Galactic dust emission are the product of radiating dust and therefore display very similar spectral energy distributions, making it difficult to disentangle the two. We show that these couplings are indeed present, even after foreground cleaning, and could be significant on the very largest angular scales. We also discuss mitigation techniques, which fortunately prove very effective.

The purpose of this paper is to gauge how delensing performance is impacted by these considerations. We begin, in section 2, with a brief introduction to B -mode delensing and how foregrounds can impact the procedure. In section 3, the simulations of Galactic dust and CIB used in this work are described (with further details on the Gaussian CIB simulations provided in appendix F), the levels of dust residuals expected from upcoming experiments are estimated, the masking schemes we use are motivated and the overall analysis is laid out in detail. A fast implementation of a curved-sky, linear-order B -mode

template is explained in appendix G. Then, in section 4, we present our results for the magnitude of delensing bias that we expect to arise from higher-point functions of both residual Galactic dust and the CIB itself, with details of the latter provided in appendices A, B, C, D and E. These results are validated in section 5 by comparing to data from Planck and to Gaussian simulations. Finally, we summarise our conclusions in section 6. Throughout this paper, we assume a cosmology with parameters as determined by Planck Collaboration (2018b).

2 MOTIVATION

2.1 B-mode delensing

The lensing-induced B -modes can be approximated, to leading order in the lensing convergence κ and in the flat-sky approximation, by

$$B^{\text{lens}}(\mathbf{l}) = \int \frac{d^2\mathbf{l}'}{(2\pi)^2} W(\mathbf{l}, \mathbf{l}') E(\mathbf{l}') \kappa(\mathbf{l} - \mathbf{l}'), \quad (1)$$

where $E(\mathbf{l})$ is the unlensed E -mode polarisation. The geometric coupling

$$W(\mathbf{l}, \mathbf{l}') = \frac{2\mathbf{l}' \cdot (\mathbf{l} - \mathbf{l}')}{|\mathbf{l} - \mathbf{l}'|^2} \sin 2(\psi_{\mathbf{l}} - \psi_{\mathbf{l}'}), \quad (2)$$

where $\psi_{\mathbf{l}}$ is the angle that \mathbf{l} makes with the x -axis. It was shown by Challinor & Lewis (2005) that the power spectrum of this linear-order approximation is in excellent agreement with the non-perturbative power spectrum at multipoles $l \leq 2000$, which comfortably includes the range we are interested in ($l < 200$) when searching for B -modes from primordial gravitational waves.

Given some tracer I of the lensing convergence, we can construct an estimate of the lensing B -modes as follows:

$$\hat{B}^{\text{lens}}(\mathbf{l}) = \int \frac{d^2\mathbf{l}'}{(2\pi)^2} f(\mathbf{l}, \mathbf{l}') W(\mathbf{l}, \mathbf{l}') E^{\text{obs}}(\mathbf{l}') I(\mathbf{l} - \mathbf{l}'), \quad (3)$$

where E^{obs} is the observed (lensed and noisy but beam-deconvolved) E -mode and $f(\mathbf{l}, \mathbf{l}')$ is a function that can be chosen to minimise the variance of the power spectrum of the delensed B -modes. We shall often refer to equation (3) as the “lensing template”. The delensed B -modes are obtained by subtracting this lensing template from the observed B -modes. Delensing methods structured around the construction and subtraction of a lensing template provide significant advantages over alternative approaches relying on an inverse deflection of the observed fields. This is particularly true for ground-based experiments, for which it is much more cost-effective to carry out the large-angular-scale B -mode science using arrays of small-aperture telescopes (SATs)¹, while the high-resolution polarisation observations required to delens are made by a large-aperture telescope (LAT). The template formalism provides the necessary framework to combine the data from the different telescopes. In the limit of noiseless E -modes and perfect ϕ , such a template built from lensed E -modes (which turns out to be preferable over using delensed E -modes, or forming a non-perturbative template by remapping lensed E -modes) can be used to reduce the residual lensing power down to around 1% of its original amplitude. In practical applications, the method will be effectively optimal even beyond the era of CMB-S4 (Baleato Lizancos et al. 2021).

The matter tracer, I , can be obtained either internally from the observed CMB or externally via observables known to correlate with the distribution of matter in the Universe. Although, ultimately, internal techniques based on quadratic estimators (Okamoto & Hu 2003), maximum-likelihood (Hirata & Seljak 2003) or Bayesian/iterative approaches (Carron & Lewis 2017; Millea et al. 2019) will provide the best means of extracting the information, current and near-future experiments will benefit greatly from combining internal reconstructions with external tracers (Smith et al. 2012; Sherwin & Schmittfull 2015; Yu et al. 2017; Manzotti 2018; BICEP/Keck & SPTpol Collaborations 2020). The reason for this is that large-scale lensing B -modes are produced from lenses over a wide range of multipoles (with 50% of the power coming from lenses at $l > 400$; Smith et al. 2012; Challinor et al. 2018). Internal reconstructions on intermediate and small scales from experiments surveying wide fields at high angular resolution, such as the forthcoming Simons Observatory (SO; The Simons Observatory Collaboration 2019), will be noise-dominated on these intermediate and small scales and so will benefit from combining with external tracers. Looking further ahead, for experiments such as CMB-S4 (Abazajian et al. 2019), the gain in delensing efficiency from adding external tracers is marginal. However, it will be possible to use the CIB in combination with other tracers as a cross-check on internal delensing.

¹ The main challenge for making highly sensitive measurements of CMB polarisation on the largest angular scales ($l < 20$) is control of temperature-to-polarisation leakage of the spatially- and temporally-varying emission from the atmosphere. A leading way to achieve this is by means of devices that modulate the observed polarisation. At the time of writing, it is more cost-effective to increase sensitivity by installing these devices on many SATs rather than on a LAT (Abazajian et al. 2019).

For an external tracer, [Sherwin & Schmittfull \(2015\)](#) showed that the choice of $f(\mathbf{l}, \mathbf{l}')$ in equation (3) that minimises the power spectrum of the delensed B -modes is a product of Wiener filters:

$$f(\mathbf{l}, \mathbf{l}') = \mathcal{W}_l^E \mathcal{W}_{|\mathbf{l}-\mathbf{l}'|}^I, \quad (4)$$

with

$$\mathcal{W}_l^E \equiv \frac{\tilde{C}_l^{EE}}{\tilde{C}_l^{EE} + N_l^{EE}} \quad \text{and} \quad \mathcal{W}_l^I \equiv \frac{C_l^{\kappa I, \text{fid}}}{C_l^{II}}. \quad (5)$$

Here, \tilde{C}_l^{EE} is the lensed E -mode power spectrum, N_l^{EE} is the instrument noise power, $C_l^{\kappa I}$ is the cross-power spectrum between the tracer I and the true CMB lensing convergence κ , and C_l^{II} is the total power spectrum of the tracer including, for example, shot noise. In the remainder of this work, we shall assume that E -mode observations are sample-variance limited up to $l \approx 2000$, so that $\mathcal{W}_{l < 2000}^E = 1$. This is a good approximation for current and upcoming experiments with beams of a few arcmin and polarisation sensitivities better than $10 \mu\text{K arcmin}$.

It is worth noting that delensing performance can be further improved by linearly combining many lensing tracers $\{I_i\}$ in a way that maximises the cross-correlation between the emerging, co-added tracer and the true convergence. [Sherwin & Schmittfull \(2015\)](#) show that the weights that achieve this are the Wiener filter such that

$$\mathcal{W}_l^I I(\mathbf{l}) \rightarrow \sum_{ij} I_i(\mathbf{l}) \left(C_l^{-1}\right)^{ij} C_l^{\kappa I_j}, \quad (6)$$

where $C_l^{\kappa I_i}$ is the cross-power spectrum between the true convergence and the i th tracer, and $\left(C_l^{-1}\right)^{ij}$ is the matrix inverse of the cross-power spectra $C_l^{I_i I_j}$ between the tracers. The correlation coefficient of the linearly-combined tracers with the true convergence can be expressed in terms of the individual correlation coefficients, $\rho_l^{\kappa I_i}$, and the correlation coefficients between tracers, $\rho_l^{I_i I_j}$, as

$$\rho_l^2 = \sum_{ij} \left(\rho_l^{-1}\right)^{ij} \rho_l^{\kappa I_i} \rho_l^{\kappa I_j}. \quad (7)$$

Qualitatively, on a given angular scale, this scheme gives most weight to those tracers that best correlate with the underlying convergence. In practice, this means that internal reconstructions, which accurately reconstruct lensing on the largest angular scales, can be supplemented with external tracers on the small scales where the former are dominated by reconstruction noise. For SO, for example, information from the CIB and photometric galaxy surveys will enable the co-added tracer to maintain a high degree of correlation with the true lensing convergence for multipoles $250 \lesssim l \lesssim 1000$ ([Namikawa et al. 2022](#)), despite internal reconstructions only being signal-dominated for $L \lesssim 250$ ([The Simons Observatory Collaboration 2019](#)). This is of particular importance for delensing, since, as noted above, it is those intermediate and small-scale lenses that are most relevant for this purpose ([Smith et al. 2012](#)). Since the small-scale lenses are located at high redshift (see figure 3 of [Lewis & Challinor 2006](#)), delensing will remain a challenging effort for the foreseeable future.

Though the analysis in this paper will be restricted to the case where delensing is carried out using the CIB as the sole tracer, it will be worth keeping the multi-tracer weights in mind for two reasons. First, it will inform our choice of minimum multipole cutoff, l_{min} , to be applied to the CIB maps – as we will see in later sections, this choice has a significant impact on the expected amplitude of the bias in the delensed spectrum. Second, the magnitude of the bias in an analysis involving multiple tracers, one of them being the CIB, can be recovered from that determined in this work once the weight given to the CIB in the linear combination is known. In this sense, the results in this paper can be regarded as an upper limit on the amplitude of the bias terms.

The power spectrum of residual lensing B -modes after delensing can be expressed to leading order as

$$\tilde{C}_l^{BB, \text{res}} = \int \frac{d^2 \mathbf{l}'}{(2\pi)^2} W^2(\mathbf{l}, \mathbf{l}') C_l^{EE} C_{|\mathbf{l}-\mathbf{l}'|}^{\kappa\kappa} \left[1 - \mathcal{W}_l^E \rho_{|\mathbf{l}-\mathbf{l}'|}^2\right], \quad (8)$$

where ρ_l is the correlation coefficient between the tracer at hand (in our case, the CIB, though possibly a co-added tracer) and the underlying CMB lensing convergence. Much like its lensed counterpart, the delensed spectrum above is almost constant (i.e., like white noise) on large angular scales. Furthermore, this behaviour can be shown to be robust to uncertainties in measurements of the auto- and cross-spectra (with CMB lensing) of the tracers when combined with a precise large-scale internal reconstruction. This means that, in actual analyses, it should be possible to deal with any uncertainties in the residual lensing power by marginalising over the amplitude of a constant spectrum. This will be further discussed in [Namikawa et al. \(2022\)](#). However, the method is rather robust even when the CIB is the only tracer involved, with [BICEP/Keck & SPTpol Collaborations \(2020\)](#) calculating that a misestimation of the cross-spectrum could bias their inference of r by at most 0.2 standard deviations.

2.2 The impact of residual foregrounds on CIB delensing

At high frequencies, Galactic dust emission and CIB radiation both behave as grey-bodies with very similar temperatures and spectral indices. Hence, it is very difficult to separate the two based on multi-frequency information alone. Consequently, CIB maps extracted in this way – such as those obtained by [Planck Collaboration \(2014c\)](#) – suffer considerable contamination from Galactic dust, and vice versa.

Several approaches to separating Galactic dust from the CIB have been explored in the literature. In this paper, we evaluate two such techniques. First, we consider the GNILC CIB maps of [Planck Collaboration \(2016e\)](#) which several groups have recently used in implementations of delensing ([Larsen et al. 2016](#); [Yu et al. 2017](#); [BICEP/Keck & SPTpol Collaborations 2020](#)). This algorithm makes use of differences in the angular power spectrum of the two components – approximately l^{-1} for the CIB ([Planck Collaboration 2011](#)) and $l^{-2.7}$ for the dust ([Planck Collaboration 2014c](#)) – to separate them.

Another method to disentangle CIB and Galactic dust emission is by exploiting the correlation between 21-cm emission from neutral Hydrogen, H I, in the inter-stellar medium (ISM) and Galactic dust ([Planck Collaboration 2011](#)). Recently, [Lenz et al. \(2019\)](#) have harnessed this to clean high-frequency Planck observations using H I data from the HI4PI survey ([HI4PI Collaboration 2016](#)), producing high-quality CIB maps on a wide range of angular scales. As a side note, H I gas in the ISM tends to form filamentary structures aligned with the local Galactic magnetic field ([Clark et al. 2014, 2015](#)), so it correlates with polarised emission from asymmetric dust grains, which also tend to align with their short axes pointing along the local magnetic field. By harnessing this connection, observations of H I have the potential to further our understanding of Galactic foregrounds for CMB polarisation studies.

From the considerations above, it is clear a complete removal of dust from the CIB maps is virtually impossible. Hence, we now seek to identify the ways in which the presence of residual dust could impact the power spectrum of delensed B -modes – the estimator from which we hope to extract information about a primordial component. The most naïve approach would have us model this estimator as the power spectrum of residual lensing B -modes, calculated ignoring any dust or CIB residuals in the B -modes and in the E -modes within the lensing template, plus the power spectra of experimental noise, and residual amounts of dust and CIB B -mode polarisation. That is

$$\begin{aligned} C_l^{BB, \text{del, naive}} &= \tilde{C}_l^{BB, \text{res}} + N_l^{BB} + C_l^{BB, \text{dust, res}} + C_l^{BB, \text{CIB, res}} \\ &= \tilde{C}_l^{BB} - 2g_l \left[\langle \tilde{B} \tilde{E} I^{\text{CIB}} \rangle \right] + h_l \left[\langle \tilde{E} I^{\text{tot}} \tilde{E} I^{\text{tot}} \rangle \right] + N_l^{BB} + C_l^{BB, \text{dust, res}} + C_l^{BB, \text{CIB, res}}, \end{aligned} \quad (9)$$

where I^{tot} includes CIB (I^{CIB} , composed of correlated signal and shot noise), residual Galactic dust (I^{dust}) and experimental noise (I^{noise}); \tilde{E} denotes the lensed E -modes; and \tilde{C}_l^{BB} , N_l^{BB} , $C_l^{BB, \text{dust, res}}$ and $C_l^{BB, \text{CIB, res}}$ are, respectively, the angular power spectra of lensing, instrument noise, residual dust and residual CIB B -modes. For economy of notation, we have denoted the convolution integral of equation (3) as the functional g_l , written the double integral as the functional h_l , and hidden the multipole-dependence of their arguments². We have also assumed that E -mode instrument noise is negligible, but this can easily be accounted for by replacing the lensed E -modes in the h_l term in equation (9) with their noisy counterparts. The inclusion of the dust and noise components along with the CIB in I^{tot} in this term is in order to take into account the decorrelation between CIB and lensing convergence that comes about as a consequence of their presence.

The four-point function in equation (9) can be written in terms of its connected part (denoted $\langle \rangle_c$) and Gaussian contractions as

$$\begin{aligned} \langle \tilde{E} I^{\text{tot}} \tilde{E} I^{\text{tot}} \rangle &= \langle \tilde{E} I^{\text{tot}} \tilde{E} I^{\text{tot}} \rangle_c + \overbrace{\langle \tilde{E} I^{\text{tot}} \tilde{E} I^{\text{tot}} \rangle} \\ &= \langle \tilde{E} I^{\text{CIB}} \tilde{E} I^{\text{CIB}} \rangle_c + \overbrace{\langle \tilde{E} I^{\text{tot}} \tilde{E} I^{\text{tot}} \rangle}, \end{aligned} \quad (10)$$

where we have assumed that all fields are zero-mean. In the Gaussian contractions we have used the fact that the lensed E -modes are uncorrelated (although not independent) of I^{tot} .

In reality, there will also be residual amount of Galactic dust and CIB contained in the E -mode data, as well as residual dust in the CIB map. These will give rise to additional contributions not included in equation (9). These possible additional

² The last two terms in equation (9) are normally calculated to leading order in κ , which amounts to ignoring lensing in \tilde{E} and replacing \tilde{B} by the leading-order expression $\tilde{B} \sim E\phi$. This is actually very accurate (with errors at the 1% level of \tilde{C}_l^{BB} in the case that $I^{\text{tot}} = \phi$, and presumably even better for a partially-correlated tracer). However, the individual ϕ^4 contributions to the three- and four-point terms may be larger (at the 10% level of \tilde{C}_l^{BB}), but they cancel rather precisely ([Baleato Lizancos et al. 2021](#)).

Symbol	Description	Value
T_{dust}	Equilibrium temperature of the thermal dust SED [K]	19.6
β_{dust}	Spectral index of the thermal dust emission	1.53

Table 1. Dust parameters used in this work, obtained from [Planck Collaboration \(2016d\)](#).

terms – which we will be referring to as “biases” – are schematically described by

$$\begin{aligned}
\mathcal{B}_l &\equiv C_l^{BB,\text{del}} - C_l^{BB,\text{del,naive}} \\
&= -2g_l \left[\langle B^{\text{dust}} E^{\text{dust}} I^{\text{dust}} \rangle_c + \langle B^{\text{CIB}} E^{\text{CIB}} I^{\text{CIB}} \rangle_c \right] + h_l \left[\langle E^{\text{dust}} I^{\text{dust}} E^{\text{dust}} I^{\text{dust}} \rangle + \overbrace{E^{\text{dust}} I^{\text{CIB}} E^{\text{dust}} I^{\text{CIB}}} + \overbrace{E^{\text{dust}} I^{\text{noise}} E^{\text{dust}} I^{\text{noise}}} \right] \\
&+ h_l \left[\langle E^{\text{CIB}} I^{\text{CIB}} E^{\text{CIB}} I^{\text{CIB}} \rangle_c + \overbrace{E^{\text{CIB}} I^{\text{tot}} E^{\text{CIB}} I^{\text{tot}}} \right].
\end{aligned} \tag{11}$$

Here, we have ignored terms involving the two-point function $\langle E^{\text{CIB}} I^{\text{CIB}} \rangle$ since this vanishes if we assume, as we shall do, that the polarisation angles of CIB galaxies are uncorrelated (see section 4.2). Of the three terms with the Gaussian contractions shown explicitly, we shall only consider the first (i.e., that with I^{CIB} shown explicitly), since we can ignore instrument noise in the CIB measurement on the scales relevant for delensing and since CIB polarisation residuals are small relative to the polarised dust residuals ([Lagache et al. 2019](#)). This work estimates the remaining terms. Notably, that includes bispectra and trispectra of Galactic dust and CIB given the non-Gaussian statistics of these emission components.

In section 4.1, we use simulations to calculate the bias resulting from $g_l[\langle B^{\text{dust}} E^{\text{dust}} I^{\text{dust}} \rangle_c]$, $h_l[\langle E^{\text{dust}} I^{\text{dust}} E^{\text{dust}} I^{\text{dust}} \rangle]$ and $h_l[\langle E^{\text{dust}} I^{\text{CIB}} E^{\text{dust}} I^{\text{CIB}} \rangle]$. The dust simulations we use are detailed in section 3.1, while those of the CIB are explained in section 3.2. At the time of writing there exist no simulations of the polarised component of the CIB, so an analytic treatment of $\langle B^{\text{CIB}} E^{\text{CIB}} I^{\text{CIB}} \rangle_c$ and $\langle E^{\text{CIB}} I^{\text{CIB}} E^{\text{CIB}} I^{\text{CIB}} \rangle_c$ is required. This calculation is presented in section 4.2.

In addition to introducing a bias, the presence of residual foregrounds will increase the variance of the delensed power spectrum. Part of this additional variance will be associated with tracer fluctuations which can be described using Gaussian statistics; so when actually doing the analysis, this effect could be easily captured using simulations. However, there will be additional contributions to the variance coming from the non-Gaussianity of the foregrounds. When estimating the impact of CIB non-Gaussianities, progress can be made with existing simulations (see, e.g., [BICEP/Keck & SPTpol Collaborations 2020](#)). In the case of polarised Galactic dust, this will also require generating non-Gaussian simulations, or obtaining a faithful template from observational data. Though important, the study of this additional variance is beyond the scope of this work: in what follows, we restrict ourselves to assessing the impact of delensing biases.

3 METHODS

In this section we introduce much of the machinery used in section 4.1 to estimate, based on simulations, the delensing bias arising from residual Galactic dust.

3.1 Simulations of Galactic dust

Given the lack of small-scale polarisation data currently available, we rely on simulations to perform our analysis. In particular, we make use of a single, full-sky, non-Gaussian simulation of the Galactic dust component at 353 GHz produced by [Vansyngel et al. \(2017, hereafter VS\)](#). We focus on this frequency for two reasons: first, it is the highest frequency for which there is polarisation data from Planck, which will enable the validation process described in section 5; and second, that, among the high-frequency Planck channels, the one at 353 GHz is where the CIB is brightest relative to Galactic dust (see, e.g., [Mak et al. 2017](#)), so that one might *a priori* expect CIB maps at this frequency to be particularly useful for delensing.

The phenomenological model underlying this simulation is an extension of that in [Planck Collaboration \(2016d\)](#), which models the local Galactic magnetic field (GMF) by stacking an ordered component together with several layers containing Gaussian random realisations of turbulence drawn from a power-law angular power spectrum. The free parameters of the model are determined by fitting the model spectra to the Planck data. VS then builds on this to allow for a level of TE correlation and E/B power asymmetry that matches the observations of [Planck Collaboration \(2018a\)](#). We emphasise that the total intensity in this simulation is based on dust inferred from Planck observations, and so is representative of (and correlated with) our sky at the map level. In section 5, we will validate this simulation against Planck data, and also compare it to Gaussian simulations.

Assuming that emission from thermal dust is perfectly correlated across microwave frequencies, the dust intensity in the

145 GHz “science channel” (the one where the CMB is brightest relative to the foregrounds) can be obtained by extrapolating the 353 GHz simulated template down as

$$I_{\nu}^{\text{dust}}(145 \text{ GHz}, \hat{n}) = I_{\nu}^{\text{dust}}(353 \text{ GHz}, \hat{n}) \left(\frac{145}{353} \right)^{\beta_{\text{dust}}} \frac{B(145 \text{ GHz}, T_{\text{dust}})}{B(353 \text{ GHz}, T_{\text{dust}})}, \quad (12)$$

and similarly for Q and U . This scaling assumes that the dust spectral energy distribution (SED) is well modelled as a modified blackbody (Thorne et al. 2017) with temperature T_{dust} and spectral index β_{dust} as determined by Planck Collaboration (2018a). The values of T_{dust} and β_{dust} used are detailed in Table 1. The Planck function is $B(\nu, T)$. To convert this specific intensity to the corresponding differential CMB temperature at 145 GHz, we apply a multiplicative unit conversion factor appropriate for infinitesimally-narrow frequency bands; this can be calculated using, for example, equation (8) of Planck Collaboration (2014a).

The extrapolation described above ignores spatial variations in the dust temperature and spectral index that are known to be present in the data (Planck Collaboration 2018a). We deem this to be an acceptable assumption for our purposes, given that our treatment already relies on a number of other approximations and imperfect simulations. However, this issue should be investigated in future work (it is, for instance, an important issue for component separation efforts). In addition to this, we also expect some further decorrelation in polarised dust emission across frequencies for any that passes through more than one dust cloud. The reason, explained by Tassis & Pavlidou (2015), is that, in that direction, the total Q and U Stokes parameters arise as a sum of contributions with different polarisation orientations (to the extent that the GMF direction threading the clouds changes along the line of sight) and whose relative strengths vary with frequency if the temperature of the clouds differs (because then their emission would follow different SEDs). However, Poh & Dodelson (2017) have shown that this effect contributes only a small amount to the current extrapolation error budget for Planck data, and indeed Planck Collaboration (2018a) find no evidence of decorrelation between 150 GHz and 353 GHz at the high Galactic latitudes that we will be restricting our analysis to.

3.2 Simulations of the CIB

The subset of bias terms described in section 2.2 associated with CIB emission can, in principle, be computed using simulations. However, among these terms are $\langle B^{\text{CIB}} E^{\text{CIB}} I^{\text{CIB}} \rangle_c$ and $\langle E^{\text{CIB}} I^{\text{CIB}} E^{\text{CIB}} I^{\text{CIB}} \rangle_c$, and estimating these would require non-Gaussian simulations of the CIB polarisation, which, unfortunately, are not available at the time of writing. We therefore study them analytically in section 4.2. The remaining CIB-related bias terms can be computed using Gaussian simulations. In this section, we explain how those simulations can be generated. For the sake of wider applicability, our method allows for the resulting CIB realisations to be appropriately correlated with an underlying realisation of the lensing convergence. However, we note that having the CIB correlate with κ is not strictly necessary when calculating the subset of bias terms we are after.

Generating correlated Gaussian simulations is straightforward if all the relevant auto- and cross-spectra are known. Suppose κ_{lm} are the spherical harmonic coefficients of a particular realisation of the lensing convergence. We seek to simulate I_{lm} , the coefficients of a simulation of the CIB. (We drop the superscript CIB here to avoid cluttering the notation.) Given that the two are correlated, we can say that, in general,

$$I_{lm} = A_l^{I\kappa} \kappa_{lm} + u_{lm}, \quad (13)$$

where u_{lm} are the harmonic coefficients of the uncorrelated part and $A_l^{I\kappa}$ is a function to be solved for. By requiring the correct auto and cross-spectra, we see that

$$A_l^{I\kappa} = \frac{C_l^{I\kappa}}{C_l^{\kappa\kappa}}, \quad (14)$$

and the power spectrum of the uncorrelated part is

$$C_l^{uu} = C_l^{II} - (A_l^{I\kappa})^2 C_l^{\kappa\kappa}. \quad (15)$$

We calculate $C_l^{\kappa\kappa}$ for the parameters of the best-fit Λ CDM cosmology of Planck Collaboration (2018b). On the other hand, C_l^{II} and $C_l^{I\kappa}$ are obtained following Yu et al. (2017). Briefly, $C_l^{I\kappa}$ is based on fitting the single-SED CIB model of Hall et al. (2010) to the cross-spectrum of the GNILC CIB map (at 353 GHz) and the internal reconstruction of κ from Planck Collaboration (2016b). The same CIB model, plus a shot noise and power-law Galactic dust contribution, is then used to fit the auto-spectrum of the GNILC map. We construct C_l^{II} as the sum of the modelled (clustered) CIB spectrum and shot noise.³

Finally, we can generate I_{lm} as described in equation (13) by drawing the harmonic coefficients u_{lm} from a Gaussian probability distribution with mean zero and variance given by C_l^{uu} , and adding them to a filtered version of the convergence, $A_l^{I\kappa} \kappa_{lm}$.

³ All spectra were kindly provided by Byeonghee Yu.

In appendix F, we generalise this approach to an arbitrary number of tracers and make the implementation publicly available⁴. Along with the weighting scheme of equation (6) – also implemented – this framework can be used to simulate multi-tracer estimates of the lensing potential with which to test delensing procedures.

3.3 Estimating dust residuals in the data

In addition to simulating the raw Galactic dust and CIB emission, we must also determine what fraction of that original radiation is left in the CMB maps after foreground cleaning, as well as how much residual dust is present in the CIB maps used as the matter tracers for delensing. In general, these residuals will be a function of angular scale and of telescope type.

In order to identify what residual dust fraction to expect for the CIB maps, we first investigate the GNILC maps. Their construction from the multi-frequency Planck maps is described in detail in [Planck Collaboration \(2016e\)](#); see also [Remazeilles et al. \(2011b\)](#). For our purposes, the key points are as follows: first, at each frequency, an estimate of the dust component, m_{dust} , is extracted by filtering out the CIB, CMB and instrumental noise components; then, another map, $m_{\text{dust+CIB}}$, is obtained by removing only the CMB and the noise (this is done by withdrawing the prior on the CIB angular power spectrum from the procedure for obtaining m_{dust}). The CIB maps at each frequency are then obtained by subtracting m_{dust} from $m_{\text{dust+CIB}}$. Given that the extent of confusion between dust and noise or CMB emission ought to be relatively small, it follows that any dust that m_{dust} fails to capture will make it into $m_{\text{dust+CIB}}$, and will therefore appear as a dust residual in the CIB map. If we assume that the multipoles of m_{dust} extracted over the region of interest contain a fraction α_l of the true dust multipoles for that region, while $m_{\text{dust+CIB}}$ accurately captures the sum of the dust and CIB, the residual dust in the estimated CIB multipoles will be $1 - \alpha_l$ of the raw dust level. We estimate α_l , and hence the fractional dust residual at a given frequency, multipole and sky region, from the cross-correlation, on the patch of choice, between GNILC maps obtained by applying the GNILC algorithm to simulations containing all sky components (specifically, the FFP8 Planck simulations described in [Planck Collaboration 2016a](#); the results were kindly provided by Mathieu Remazeilles) and the input dust in these simulations. The resulting estimate for the multipole-dependent, fractional residual at 353 GHz is shown in figure 1. We note that the GNILC algorithm is optimised for dust reconstruction, so other methods designed to prioritise the extraction of the CIB might display higher levels of residual dust. We also point out that the GNILC algorithm tends to suppress the true CIB due to residual CIB in m_{dust} , which is subtracted from $m_{\text{dust+CIB}}$. There is evidence of this in the cross-correlation of the GNILC CIB maps with κ from Planck, as discussed in the appendix of [Maniyar et al. \(2019\)](#). However, for the purposes of delensing, provided that correlation coefficients are obtained empirically any suppression of the CIB will not lead to bias but only a reduction in the delensing efficiency.

Another means of mapping dust in our galaxy is via its correlation with the 21 cm emission from the hyperfine transition of neutral hydrogen. This technique was first exploited by [Planck Collaboration \(2011\)](#) to disentangle Galactic dust from extragalactic CIB emission. Subsequently, [Lenz et al. \(2019\)](#) extended the method to larger regions of the sky, and it is those maps that we consider here. In order to estimate the magnitude of dust residuals in them, we focus on the strong correlation observed between higher H I column densities ($N_{\text{H I}}$) – or, equivalently, larger sky areas – and higher dust residuals. This is the reason why angular power spectra of maps obtained using different column density thresholds converge only for the lowest few values of $N_{\text{H I}}$, and grow significantly in amplitude for higher ones (see figure 15 of [Lenz et al. 2019](#)). If we assume that all that additional auto-power is due to dust contamination, we can make a rough estimate of the residual dust power present in any given CIB map by comparing its power spectrum to that of the map with the lowest column density ($N_{\text{H I}} = 1.5 \times 10^{20} \text{ cm}^{-2}$), and assuming that the latter is free of dust. In order to avoid noise biases, auto-spectra are measured by taking the cross-correlation of maps derived from Planck observations with independent instrumental noise (the half-ring maps), masked with the appropriate apodised windows provided. As the maps from [Lenz et al. \(2019\)](#) are provided in intensity units, the resulting absolute residual spectra are converted to temperature units by multiplying by a wide-band unit conversion factor⁵ twice and subsequently compared to the raw power of the VS dust simulation at 353 GHz. The fractional residual shown in figure 1 is the ratio between these two. Away from the small scales where it increases due to the limited resolution of the H I maps, the residual dust power is approximately scale-independent. On the other hand, the raw dust power is very red, hence the shape of the residual fraction curve shown.

We now consider the removal of dust from small-scale polarisation observations of a survey with the characteristics of the Simons Observatory’s (SO) large-aperture telescope survey (LAT; [The Simons Observatory Collaboration 2019](#)) by means of an internal linear combination of channels (ILC) in harmonic space (see, e.g., [Tegmark et al. 2003](#)). Though, a priori, one might think this easier than temperature cleaning, owing to the fact that in polarisation there are only two contaminants to the signal of interest (dust and synchrotron), the situation is complicated by the fact that LAT polarisation noise levels are

⁴ <https://github.com/abaleato/MultitracerSims4Delensing>

⁵ This multiplicative unit correction differs slightly from the one we met in section 3.1 due to the fact that it takes into account the finite width of the Planck bandpasses – see [Planck Collaboration \(2014a\)](#) for an explanation, and [Planck Collaboration \(2018a\)](#) for values appropriate for the 2018 polarisation data we use.




Figure 1. Estimated map-level percentage of residual dust emission after foreground cleaning relative to the unmitigated dust emission in either the 145 GHz channel of the SO LAT (yellow), the 145 GHz channel of the SO SAT (red) or different CIB maps at 353 GHz (blue). These fits to CIB residuals are not to be trusted below $l < 100$; however, this is irrelevant in practical applications, since internal techniques dominate the reconstruction on those scales (see, e.g., Namikawa et al. 2022). For the SO, we plot the *goal* noise levels. The SO LAT residuals are estimated using a global ILC method and thus represent the fraction by which to scale the unmitigated emission anywhere on the sky. On the other hand, the residuals quoted for the CIB maps are appropriate to the specific regions of the sky on which they are estimated.

higher than dust emission for all but the largest angular scales (The Simons Observatory Collaboration 2019). Consequently, in a standard implementation of the ILC method, the most effective way to minimise the variance of the cleaned map while leaving the CMB signal unchanged will be to suppress the noise, with little being done to remove dust. Quantitatively, the fraction, R_l , of the original Galactic dust emission in the 145 GHz channel of SO that remains in the ILC-cleaned map at multipole l is

$$R_l = \sum_{\nu} w_l(\nu) I^{\text{dust}}(\nu) / I^{\text{dust}}(145 \text{ GHz}), \quad (16)$$

where $w_l(\nu)$ are the ILC weights⁶, $I^{\text{dust}}(\nu)$ is the dust SED assumed in the simulations in differential CMB temperature units (such that $\sum_{\nu} w_l(\nu) = 1$) and the sum is over all 13 of the Planck and SO frequency channels expected to be used for component separation. The unit conversions at all these frequencies are calculated under the assumption of infinitely-narrow passbands. Finally, the residual level of dust after ILC-cleaning (with noise levels corresponding to the SO’s *goal*) is shown in figure 1.

The LAT *E*-mode foreground residuals can be further suppressed by “deprojecting” the dust contribution – that is, requiring that the cleaned map have zero response to the SED of the dust component (Remazeilles et al. 2011a). Though in principle one might expect this method completely to neutralise dust residuals, in practice the removal will not be complete due to differences between the dust model used to compute the ILC weights and the true dust SED (arising, for instance, from spatial variation of the dust spectral index). Our calculated dust-deprojected residual, shown in figure 1, illustrates this, as it arises from differences in the dust model used to generate the ILC weights and that which we use to extrapolate the dust template from 353 GHz to other frequencies. This estimate is likely to be rather optimistic given that it appears to be comparable (when measured as a fraction of the unmitigated emission) to that obtained in a parametric foreground cleaning of the SAT – as we shall soon discuss, this is expected to be more extensive since it allows for a degree of variation of foreground parameters across the sky. However, our estimates will suffice given the number of other approximations we are making. As we will see in section 4.1, deprojecting the dust can help mitigate the magnitude of delensing biases. However, doing so comes at the cost of increased variance in the maps, which in turn brings about a degradation of delensing efficiency. In order to estimate the extent of the degradation, we use equation (8) to forecast the residual lensing *B*-mode power after delensing with the GNILC CIB maps and *E*-mode observations after the latter are Wiener-filtered to reflect the polarisation noise level after ILC cleaning. Our forecasts, shown in figure 2, predict a small degradation in delensing efficiency (quantified through the ratio of delensed-to-original lensing *B*-mode power), no greater than one percent, between the dust-deprojected and standard cleaning scenarios. When deciding whether or not to use dust-deprojected maps for delensing, the goal must be to maximise delensing efficiency while keeping the emerging delensing bias to acceptable levels. The newly developed partially-constrained

⁶ Based on the ILC weights provided by the authors of The Simons Observatory Collaboration (2019).




Figure 2. Fraction of lensing B -mode power left over after delensing with the GNILC CIB map and E -mode observations from the SO LAT. In both the baseline and goal scenarios, nulling the dust component in the foreground-cleaned E -mode maps brings about a small degradation in delensing efficiency relative to the standard, minimum-variance solution. This theoretical estimate is obtained from the leading-order approximation of equation (8).

ILC method of [Sultan Abylkairov et al. \(2020\)](#) might be particularly useful in this context, allowing optimisation of the tradeoff between post-cleaning noise (and associated delensing efficiency) and bias on r .

Finally, the cleaning procedure ought to be more extensive (fractionally speaking) for a small-aperture telescope – such as the the SO’s SAT – measuring CMB polarisation on large angular scales. It is on those large scales that polarised foregrounds peak and are brighter than the experimental noise (assuming that systematics are under tight enough control that leakage of atmospheric emission into polarisation is small). Hence, they will be preferentially removed in foreground-cleaning procedures. Furthermore, given the importance of mitigating the foreground variance on large angular scales to facilitate searches for primordial B -modes, map-based component separation techniques have been developed that allow for variation of the foreground modelling parameters across the sky (see, e.g., [Stompör et al. 2016](#) or [Alonso et al. 2017](#) and references therein). This flexibility results in improved cleaning: by simulating the entire foreground-cleaning operation for the SAT, [The Simons Observatory Collaboration \(2019\)](#) forecast a map-level dust residual of around 10 %, shown also in figure 1.

We fit smooth functions to all of these map-level dust residual fractions as a function of multipole and use them to filter the simulated, raw dust maps of section 3.1. The end result is maps of the expected residual dust in each of the E , B and I fields involved in the delensing analysis.

3.4 Masking

When calculating the delensing biases of section 2.2 from simulations, we must adopt masks that mimic the procedures that would be followed and sky areas that would be covered in actual analyses.

In order to reduce the extent of mode coupling induced by masking and facilitate the deconvolution of mode-coupling matrices for power spectrum estimation, all masks are apodised. We obtain Galactic masks apodised with 2° Gaussian kernels from the [Planck Legacy Archive](#). On the other hand, the mask to be used for large-scale B -mode searches was kindly provided to us (in binary form) by the authors of [The Simons Observatory Collaboration \(2019\)](#). Apodisation of this mask is carried out using procedures built into the publicly-available code [NaMaster \(Alonso et al. 2019\)](#), which ensures that regions originally excluded in the binary mask remain so after apodisation. Given that this mask has a relatively small footprint on the sky, we choose a FWHM of 1° for the Gaussian smoothing kernel. The final mask is shown in figure 3a. Henceforth, we refer to it as the “SAT mask”. This will be the mask used to delineate the patch covered by a ground-based, small-aperture telescope going after large-angular-scale B -modes. We shall also consider the case of a space-based experiment targeting B -mode science, in the style of the upcoming LiteBird satellite ([Hazumi et al. 2019](#)). The sky patch covered by such an experiment is approximated by the 60 % sky Galactic mask from Planck shown in figure 3c. We will sometimes refer to this as the “LB patch”.

Though we expect the lensing B -mode template to be quite local – owing to the fact that it is the product of predominantly intermediate and small-scale E -modes and lenses that give rise to lensed B -modes — we saw in section 2.2 that the delensing bias is sensitive to certain bispectra and trispectra. Since these scale as the third and fourth power of pixel values, they are very sensitive to bright structures in the dust residuals lying outside the footprint that we ultimately make B -mode measurements on. This is indeed what we see, already at the level of the template, when bright regions near the Galactic plane are left

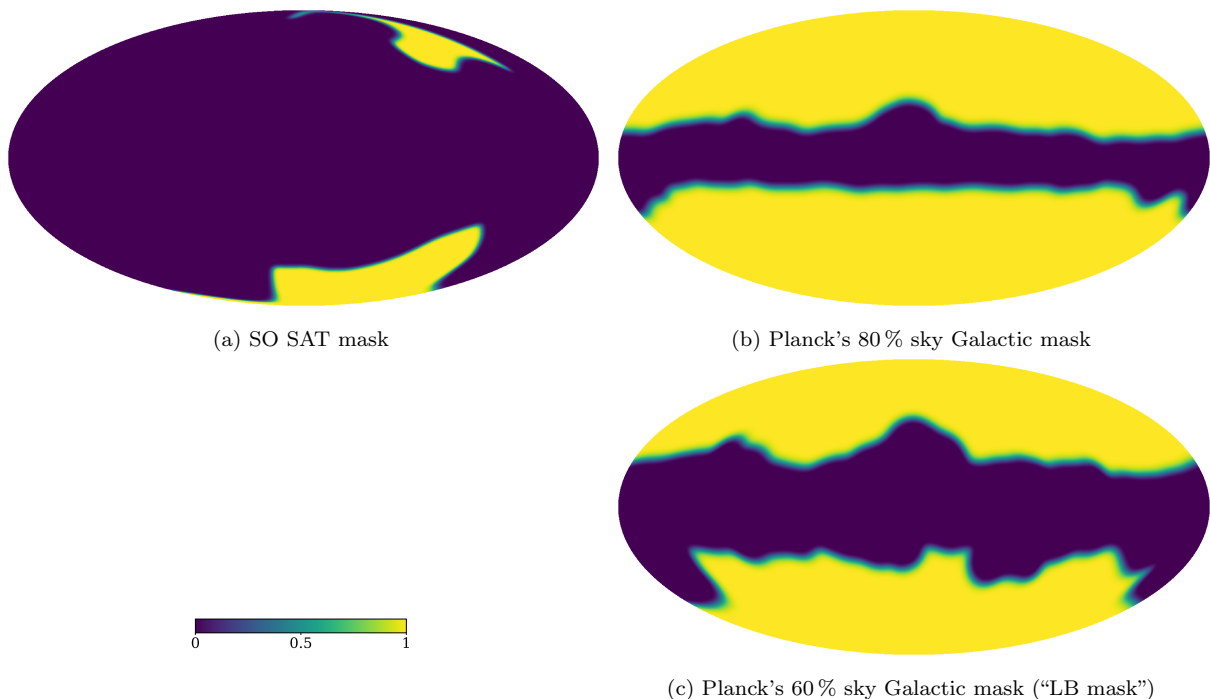


Figure 3. Masks used in our analysis. They have been apodised using Gaussian kernels with a FWHM of 1° in the case of (a) and 2° in the case of (b) and (c).

unmasked in the E and I fields from which these templates are constructed. Figure 4 shows several such templates for which input fields were first masked using Galactic masks of different extents before applying the Wiener filters in equation (5). Bright artifacts outside the masked regions are readily apparent in those involving the largest unmasked fractions. Furthermore, these issues translate to increased power and scatter on the largest scales of the resulting templates, as shown by figure 5.

These phenomena have the potential to hinder efforts to delens by accentuating the biases described in section 2.2 and therefore should be mitigated by masking the Galactic region when constructing templates from wide-area data. We settle on the Planck mask that leaves 80% of the sky available – shown in figure 3b – as the option that minimises the masked fraction (to reduce mode-coupling effects in the construction of the Wiener-filtered fields) while keeping artifacts in the low- l modes of the templates to acceptable levels. This scheme for masking fields going into templates will be implicit in all results quoted in the remainder of this paper. Of course, similar masking arrangements are inevitable in practical applications (at least from the ground) due to the telescopes’ limited sky coverage. Even for space-based experiments, it will be important to mask away the galaxy – and bright sources – before constructing lensing templates.

To check that the power spectrum amplitudes we are seeing in figure 5 are physical in origin, we also show in the figure the results for a statistically-isotropic, Gaussian simulation. In this case, the I and E residual dust fields are simulated on the full sky from power spectra forced to match those of the VS simulation measured over the SO SAT patch. The full-sky Gaussian fields are masked with the 80% sky mask prior to Wiener filtering in the construction of the lensing template. We verify that the amplitude resulting from the Gaussian simulation is consistent with an analytic calculation in the style of the leading-order lensing B -mode power spectrum where C_l^{EE} and $C_l^{\kappa\kappa}$ are replaced with the dust spectra measured from the simulation. The amplitude in this Gaussian limit is significantly lower than that arising from the VS simulations. This is likely due to the power in the non-Gaussian VS simulations being boosted by the $\langle EIEI \rangle_c$ trispectrum of Galactic dust, though, in principle, statistical anisotropy could also be contributing to the difference seen between the VS simulation and the statistically-isotropic Gaussian simulations.

3.5 Summary of simulation-based procedure for estimating delensing biases

We summarise here the procedure used to calculate the delensing biases of section 2.2 from simulations. The steps are as follows:

- (i) Load in full-sky I , Q and U simulated dust maps at 353 GHz. Take them to harmonic space to produce I^{dust} , E^{dust} and B^{dust} .
- (ii) Scale E^{dust} and B^{dust} to 145 GHz as described in section 3.1.
- (iii) Apply the dust residual transfer functions measured in section 3.3 to each of I^{dust} , E^{dust} and B^{dust} .

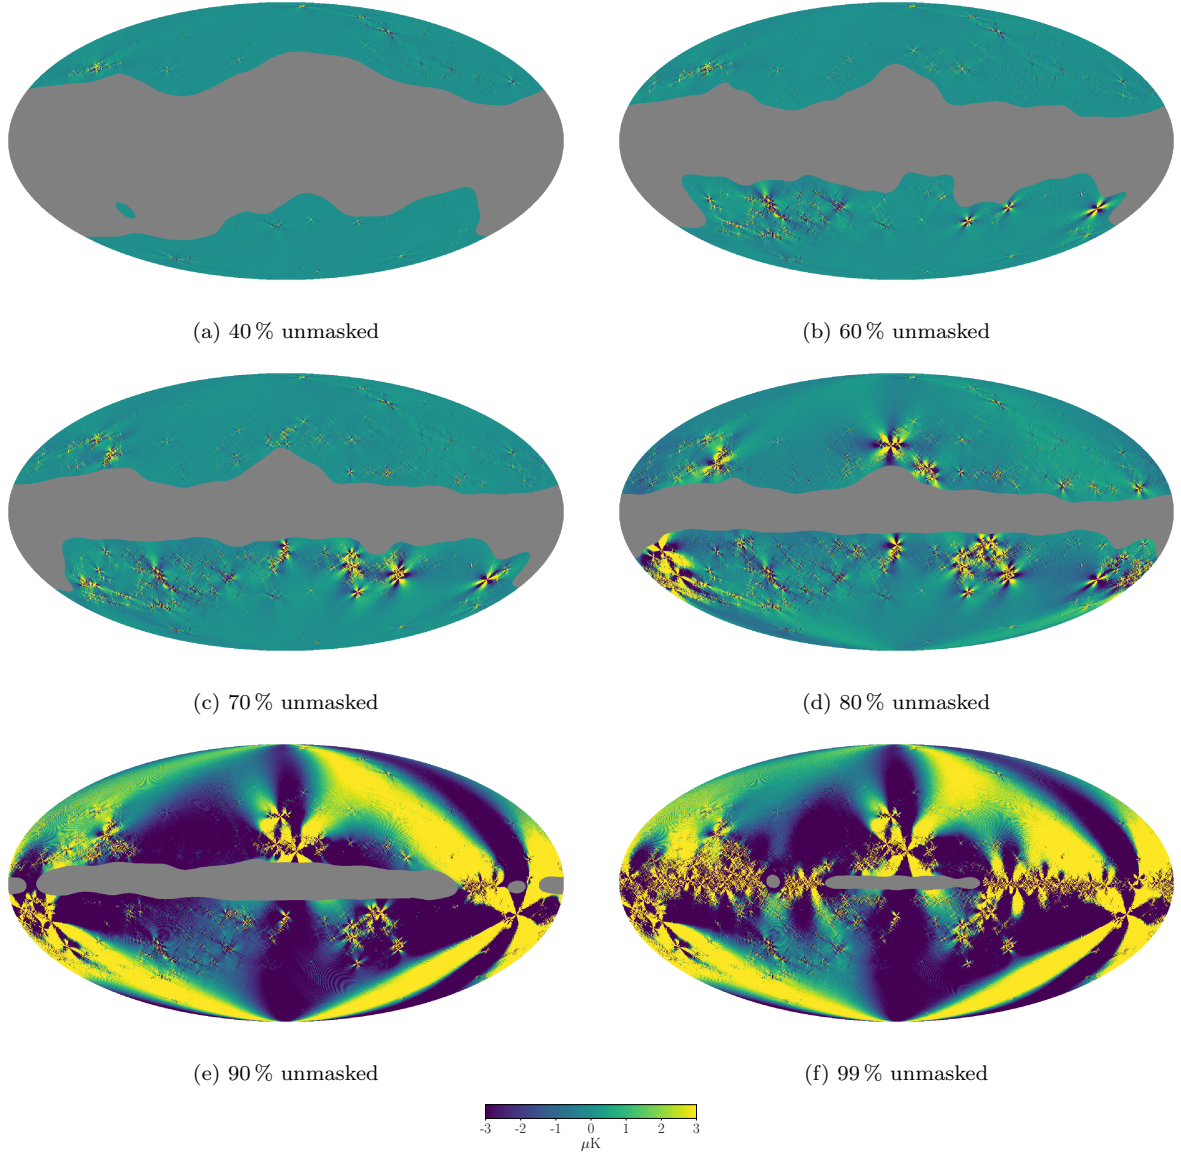


Figure 4. Stokes Q parameter associated with B -mode templates generated from residual dust intensity and E -modes (as determined in section 3.3), for different extents of Galactic masking applied to the input fields prior to Wiener filtering. Here, $X\%$ unmasked is to be interpreted as meaning that $X\%$ of the sky is left unmasked in each of the two fields, E and I , from which the template is built (the masked region is shown in grey). The input dust E -modes are scaled to the residual level expected for standard ILC cleaning of the LAT in the *goal* scenario, while dust total-intensity is scaled to levels appropriate for the residuals in the GNILC CIB maps over the SO SAT patch. Intensity modes below $l_{\min} = 100$ are removed in all cases. The bright, hexadecapolar features are the expected consequence of leaving polarised clumps of Galactic emission unmasked in the inputs.

- (iv) Simulate a CIB map, I^{CIB} , in harmonic space using the method of section 3.2.
- (v) Take I^{dust} , E^{dust} , B^{dust} and I^{CIB} to configuration space.
- (vi) Apply the apodised 80% sky mask of figure 3b to all of the maps.
- (vii) Take maps back to harmonic space.
- (viii) Wiener filter I^{CIB} and I^{dust} according to equation (5), with the total power spectrum in the denominator appropriate to the appropriate residual dust level under study. Mask modes below the l_{\min} cutoff of choice.
- (ix) Using the methods described in appendix G⁷, build template B -modes of the form $\hat{B}(E^{\text{dust}}, I^{\text{dust}})$ and $\hat{B}(E^{\text{dust}}, I^{\text{CIB}})$.
- (x) Convert these templates, as well as B^{dust} from step (vii), to Q and U maps.
- (xi) Apply the apodised SO SAT mask of figure 3a to these Q and U maps.

⁷ These tools are packaged and made publicly-available under https://github.com/abaleato/curved_sky_B_template




Figure 5. Power spectra, measured on the SO SAT patch of figure 3a, of the dust-only B -mode templates shown in figure 4. We emphasise the comparison between the VS dust simulation and a Gaussian simulation (whose amplitudes in intensity and polarisation are made to match the power of the VS simulation measured on the SO SAT patch), after both are masked with the same 80% sky mask of figure 3b prior to Wiener filtering in the construction of the lensing template. We verify that the power in the Gaussian case is consistent with that expected from a leading-order lensing-like calculation based on the E and I dust spectra measured from the masked maps. On the other hand, the VS simulation displays more power, likely due to the relevance of the Galactic dust trispectrum. We also note the approximate convergence of template power between different masking fractions on intermediate and small scales. Power spectra, with the effect of the SO SAT mask deconvolved, are extracted with the code `NaMaster`, using bins with $\Delta l = 90$.

(xii) Use the pseudo- C_l code `NaMaster` (Alonso et al. 2019) to measure the correlations of interest (described in section 2.2) and deconvolve the mode couplings induced by the mask.

(xiii) Fit a power-law to the bandpowers in the range $200 < l < 600$.

Note that, in step (vi), we use a Galactic mask from Planck rather than that from the SO LAT. There are two reasons for this: firstly, this mask is less restrictive and hence leads to less E -to- B leakage in polarisation (Lewis et al. 2002; Bunn et al. 2003); and secondly, due to the local nature of the templates (see figure 4) it should be possible to recover similar results if we were to restrict analysis to the LAT patch. In contrast, in step (xii), the SAT mask is used, which restricts analysis to a much smaller patch than that in step (vi). However, there is no risk of E -to- B leakage here since the polarisation maps do not contain E -modes by construction. For this reason, we do not bother to use the pure- B -mode formalism (Smith 2006) when computing power spectra over the SO SAT patch.

As a means of cross-validating our pipeline, we apply the same filtering and masking scheme to lensed simulations of CMB E -modes and combine these with I^{CIB} to build a lensing template. We verify that, when this template is used to delens B -modes, the theoretically-expected level of power (as determined by equation 8) is obtained.

Except where noted otherwise, we restrict the fields we work with to multipoles $8 < l < 2008$. This is the Planck GNILC bandwidth and, approximately, includes the full range of scales relevant for delensing the large-scale B -mode polarisation.

4 RESULTS

4.1 Bias from residual Galactic dust

With the tools developed in the previous sections, we are now in a position to estimate the delensing bias arising from residual Galactic dust left over in the CMB and CIB maps involved in the delensing analysis. By applying the procedures of section 3.5 to dust simulations from VS, we shall see that there is indeed a bias dominated by the BEI bispectrum of Galactic dust, that it is negative, and that it has a very red power spectrum, peaking on the largest angular scales.

In figure 6, we show this power spectrum bias for several assumptions about the residual dust levels according to the type of CIB maps used, minimum- l -cuts applied to these, sky patch on which the B -mode measurement is carried out and whether or not dust is explicitly deprojected when ILC-cleaning the E -modes. We compare these biases to the power spectrum of residual lensing B -modes after delensing down to 70% of the original lensing power — this is approximately the expected extent of delensing when using the GNILC CIB maps as the only tracer (c.f., figure 2). For reference, we show via shaded regions the Gaussian part of the standard deviation of delensed B -modes when observed on either $\sim 10\%$ or $\sim 60\%$ of the sky, representative of the SO B -mode survey patch or a larger patch observable from space, respectively. We model this delensed B -mode spectrum as comprising 70% of the original lensing power and a white noise component in polarisation

similar to that of the SO SAT’s 145 GHz channel in the *goal* specifications (a beam $\theta_{\text{FWHM}}=17$ arcmin and a sensitivity of $\Delta p = \sqrt{2} \times 2.1 \mu\text{K arcmin}$; [The Simons Observatory Collaboration 2019](#)).

The very ‘red’ spectral tilt displayed by the dust bias means that it will likely be smaller than the statistical uncertainty on the multipole range and sky regions probed by upcoming ground-based telescopes ($l \gtrsim 30$ and $f_{\text{sky}} < 0.1$), as long as modes of the CIB mass-tracer maps are removed for $l \lesssim 100$. From the left panel of figure 6, we see that this is true even in the case where dust is not deprojected as part of the foreground-cleaning procedure. However, one might want to be especially careful and work with dust-deprojected E -modes, as this is shown by the right panel of figure 6 to mitigate the bias very significantly, while only marginally degrading delensing performance (at least for SO, see figure 2).

On the other hand, the bias might pose challenges for space-based experiments (such as LiteBird) for two reasons: first, given their full-sky coverage and immunity to atmospheric noise, satellite missions will be able to probe polarisation down to the lowest multipoles, but it is on those largest scales that the bias is largest; and second, in observing a larger fraction of the sky, they encompass regions where the dust is brighter. These, compounded with the fact that larger sky coverage ($f_{\text{sky}} \approx 0.6$) translates to lower statistical uncertainty, means that this bias could potentially confuse searches for primordial B -modes by space-based missions, unless their foreground cleaning capabilities are improved relative to SO. To illustrate this, we consider a scenario where the foreground cleaning is as for SO, but the analysis is performed on the larger ‘LiteBird’ patch. Figure 6 shows that the resulting bias is well in excess of the statistical errors, particularly when dust is not explicitly deprojected from the E -mode maps. The right panel of figure 6 shows that doing this deprojection is very effective at reducing the amplitude of the bias, but even then it is large enough to significantly compromise constraints on r (as we will quantify shortly). Fortunately, space-based experiments targeting large-scale B -mode science will take advantage of their privileged vantage point above the atmosphere to observe in many more frequency channels than SO – LiteBird, for instance, is expected to have 15 frequency bands ([Hazumi et al. 2019](#)). Hence, foreground cleaning of, at least, the large-scale B -modes should be more extensive than what is assumed in the figure, which has dust residuals appropriate for SO. Since the bias is dominated by the BEI bispectrum of dust, its amplitude is very sensitive to improved cleaning of the CMB polarisation fields.

For a fixed level of foreground residuals, the bias can be mitigated by filtering out the largest scales of the CIB maps, which tend to be relatively more contaminated by dust (this is particularly acute for the GNILC products, as discussed in section 3.3). In all likelihood, upcoming implementations of CIB delensing will be complemented with internal reconstructions, which are able to reconstruct the largest-angular-scale lenses very accurately. Even if the largest scales are discarded altogether, the impact is expected to be small since it is scales smaller than that (chiefly $200 < l < 500$) that are most informative for delensing: for instance, [Sherwin & Schmittfull \(2015\)](#) show that delensing efficiency only worsens by 12% if the CIB fields are high-pass filtered with $l_{\text{min}} = 150$ instead of $l_{\text{min}} = 60$. Consequently, the recommendation we issue here of removing those largest scales from the CIB maps should come at very little cost in terms of delensing efficiency.

In figure 7, we quantify how the power spectrum biases of figure 6 propagate to biases on the inferred value of the tensor-to-scalar ratio, r . We use (e.g., [Roy et al. 2021](#))

$$\Delta \hat{r} = \left(\sum_{l=l_{\text{min}}}^{l_{\text{max}}} \left[C_l^{BB,\text{prim}}(r=1) \right]^2 / \text{Var} \left(C_l^{BB,\text{del}} \right) \right)^{-1} \sum_{l=l_{\text{min}}}^{l_{\text{max}}} C_l^{BB,\text{unmodelled}} C_l^{BB,\text{prim}}(r=1) / \text{Var} \left(C_l^{BB,\text{del}} \right) \quad (17)$$

to translate an improper modelling of the B -mode power spectrum to a shift in r . Here, $C_l^{BB,\text{prim}}(r=1)$ is the angular power spectrum of primordial B -modes with $r=1$, $\text{Var}(C_l^{BB,\text{del}})$ is the variance of the power spectrum of delensed B -modes and $C_l^{BB,\text{unmodelled}}$ is the part of the delensed B -mode spectrum that we have failed to model. We compute $\text{Var}(C_l^{BB,\text{del}})$ assuming the delensed B -modes to be Gaussian, and to feature only experimental noise (at a level appropriate for the SO SAT’s 145 GHz *goal* specifications), primordial signal and 70% of the original lensing power. For constraints on the SO SAT patch, we use $l_{\text{min}} = 20$ and $l_{\text{max}} = 200$, and for the larger ‘LB’ patch we set $l_{\text{min}} = 2$ and $l_{\text{max}} = 200$. We compare the estimated shifts to the standard deviation on inferences of r in the limit $r=0$ expected of experiments covering either 10% or 60% of the sky, with the noise levels of the SO SAT, and in the limit of no foreground BB power and residual lensing power of 70% of the original. Figure 7 shows that the bias discussed here results in inferences of r that are systematically lower than the truth, although for ground-based experiments this is still within the bounds of statistical errors. This is in contrast with the impact of failing to model dust residuals on large angular scales of the B -mode power spectrum, in which case one tends to over-estimate the true value of r . As we show in the figure, failing to model dust B -mode residuals with a level of power as estimated in figure 1 for the SO SAT could result in a bias larger than those considered in this section, but in the opposite direction.

By combining the mitigation techniques described here, we ultimately expect the delensing bias from residual Galactic dust to be contained to levels where it does not compromise constraints on r , for virtually any application of CIB delensing. In particular, deprojection of dust from the E -modes proves to be particularly effective, while incurring only a small degradation in delensing efficiency, as shown in figure 2. However, this work highlights that, if we are to exploit the full power of the CIB for delensing, realistic simulations of Galactic dust, extending to small angular scales, ought to be developed so that they can be used to model and properly account for the expected amplitude and shape of the bias.

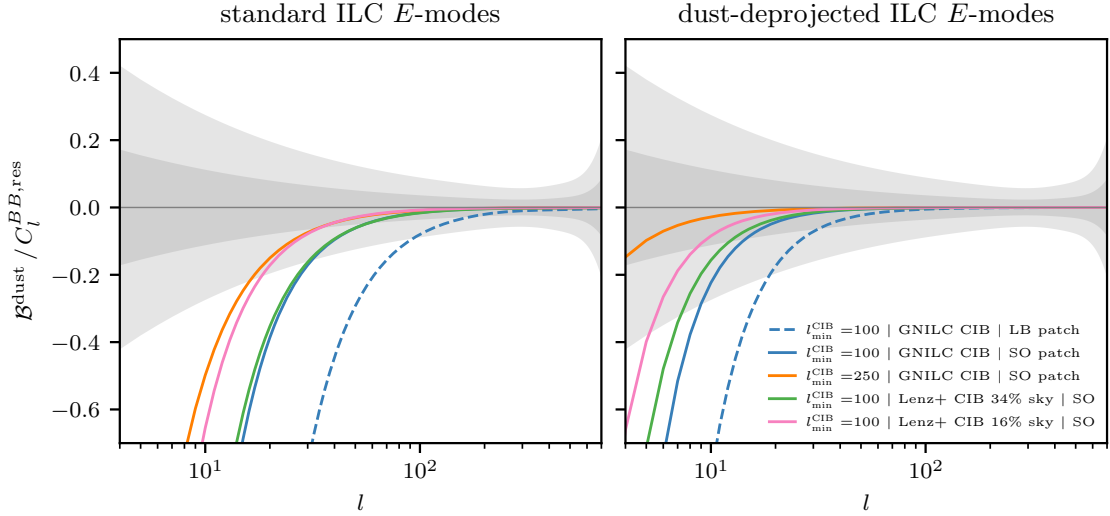


Figure 6. Fractional bias on the angular power spectrum of residual lensing B -modes due to residual Galactic dust, plotted for several scenarios with differing residual dust levels. We model the residual lensing power as $C_l^{BB,res} = 0.7\tilde{C}_l^{BB}$, as expected after delensing with the CIB as the only tracer. *Left panel:* E -mode residual for a standard ILC cleaning of SO's LAT (no dust deprojection and *goal* noise levels assumed). *Right panel:* E -mode residual as for SO's LAT (*goal* noise levels), with deprojection of dust in the ILC cleaning. Foreground cleaning of B -modes is everywhere as for the SO SAT. Dashed lines are for measurements on the large 'LB patch', while solid ones are on the smaller SO SAT patch. The shaded regions show the $\pm 1\sigma$ uncertainty on the power spectrum of delensed B -modes $C_l^{BB,del} = 0.7\tilde{C}_l^{BB} + N_l^{BB}$, where N_l^{BB} is the (*goal*) noise power spectrum of the SO SAT. We calculate this uncertainty by assuming the delensed field is Gaussian, and that the experiment covers either 10% (light grey) or 60% of the sky (dark grey), and uses bins of $\Delta l = 30$. These biases are dominated by the BEI bispectrum of Galactic dust and are negative.

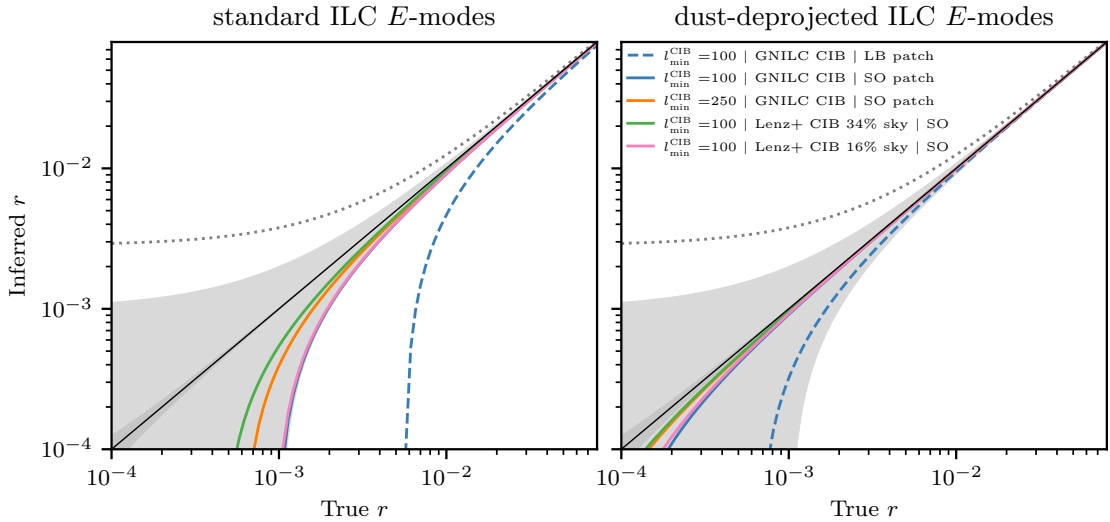


Figure 7. Bias on the inferred tensor-to-scalar ratio due to unmodelled contributions from residual non-Gaussian Galactic dust to the power spectrum of B -modes delensed with the CIB. Dashed lines are for measurements on the 'LB patch' (approximately 60% of the sky), and solid ones are on the smaller SO SAT patch (approximately 10% of the sky). Foreground cleaning of B -modes is everywhere as for the SO SAT. Dust residuals in the E -modes are as for a harmonic ILC cleaning of the SO LAT at *goal* noise levels, either in a standard implementation (left panel) or after deprojection of Galactic dust (right panel). We constrain r using scales $l_{min} = 20$ and $l_{max} = 200$ or $l_{min} = 2$ and $l_{max} = 200$ when working with either the SO SAT patch or the 'LB' patch, respectively. For reference, the shaded regions show the approximate $\pm 1\sigma$ uncertainty for $r = 0$ of experiments covering either 10% (light grey) or 60% (dark grey) of the sky with the noise level of the SO SAT, no foregrounds, delensing down to 70% of the original lensing BB power and when r is constrained over the multipole ranges described above for each of the sky fractions. For comparison also, the dotted lines show the size of the bias on r if residual dust B -modes in the SO SAT maps are as in figure 1 and are not modelled in the BB power spectrum.

4.2 Intrinsic bias from non-Gaussianity of the CIB

Emission from UV-heated dust in star-forming galaxies – the diffuse component of which constitutes the CIB – is linearly polarised. This is thought to arise from the alignment of asymmetric dust grains in the inter-stellar medium of these galaxies with the local magnetic field direction. For observations along a particular line of sight, the stacking of emission from a large number of sources, each with incoherent, randomly-orientated polarisation directions, gives rise to depolarisation of the observed, integrated radiation. This depolarisation is practically complete for low-flux sources, of which there are very many. On the other hand, the number of bright sources is more modest, and their incomplete averaging gives rise to a net polarisation of the observed light. Experimental applications are able to mitigate this effect somewhat by removing the bright sources that they are able to resolve (e.g., [Planck Collaboration 2016c](#)), but the contribution from those that are not resolved ought to be modelled.

We determined in section 2.2 that, if the CIB itself has non-vanishing $\langle BEI \rangle_c$ and $\langle EIEI \rangle_c$ correlations, the power spectrum of delensed B -modes will be biased by the following two terms:

$$\mathcal{B}_l^{BEI} = \frac{4}{(2\pi)^2} \int \frac{d^2 l'}{(2\pi)^2} \frac{l' \cdot (l - l')}{|l - l'|^2} \mathcal{W}_{l'}^E \mathcal{W}_{|l-l'|}^I \sin 2(\psi_{l'} - \psi_l) \langle B(-l) E(l') I(l - l') \rangle'; \quad (18)$$

and

$$\mathcal{B}_l^{EIEI} = -\frac{4}{(2\pi)^2} \int \frac{d^2 l' d^2 l''}{(2\pi)^4} \frac{l' \cdot (l - l')}{|l - l'|^2} \frac{l'' \cdot (l + l'')}{|l + l''|^2} \mathcal{W}_{l'}^E \mathcal{W}_{l''}^E \mathcal{W}_{|l-l'|}^I \mathcal{W}_{|l+l''|}^I \sin 2(\psi_{l'} - \psi_l) \sin 2(\psi_{l''} - \psi_l) \times \langle E(l') I(l - l') E(l'') I(-l - l'') \rangle'_c, \quad (19)$$

where the primes on the expectation values denote that the delta-function has been dropped.

In order to obtain an analytic understanding of these biases and be able to calculate higher-point functions of the CIB, we introduce a minimal model where the source polarisation fraction is constant and the polarisation angles of sources are uncorrelated. A detailed derivation of the BEI bispectrum of the CIB in this limit is provided in appendix A. Following equation (A13), we learn that this sources a bias given in CMB temperature units by

$$\mathcal{B}_l^{BEI} = 2p^2 G_{353 \text{ GHz}} G_{145 \text{ GHz}}^2 \int \frac{d^2 l'}{(2\pi)^2} \frac{l' \cdot (l - l')}{|l - l'|^2} \mathcal{W}_{l'}^E \mathcal{W}_{|l-l'|}^I \sin^2 2(\psi_{l'} - \psi_l) \times \int dz \left(\frac{I^{\text{CIB}}[145(1+z) \text{ GHz}]}{I^{\text{CIB}}[353(1+z) \text{ GHz}]} \right)^2 \left[S_{353 \text{ GHz}}^{(3)}(z) + \frac{H(z)}{cr^2(z)} S_{353 \text{ GHz}}^{(2)}(z) S_{353 \text{ GHz}}^{(1)}(z) P_g(|l - l'|/r(z); z) \right], \quad (20)$$

where $I^{\text{CIB}}(\nu)$ is the CIB SED at frequency ν , G_ν is a conversion factor⁸ to go from specific intensity to differential CMB temperature units at frequency ν , p is the polarisation fraction of CIB galaxies, which we assume to have a constant value of 5%⁹, and $P_g(k; z)$ is the galaxy power spectrum. We have also defined

$$S_\nu^{(n)}(z) \equiv \int_{s_{\min}}^{s_{\max}} dS_\nu S_\nu^n \frac{dN}{dS_\nu dz d\Omega}, \quad (21)$$

that is, an integral over the n th power of source flux density, S_ν , weighted by $dN/dS_\nu dz d\Omega$ – the number of sources per steradian, at a given redshift in a certain flux-density range $[s_{\min}, s_{\max}]$ at frequency ν . These number counts can be obtained by fitting a model to sub-mm data. In this work, we use those of [Béthermin et al. \(2011\)](#).

Since we have assumed that the polarisation angles of different galaxies are uncorrelated, the two polarisation legs of the bispectrum must be sourced by the same galaxy, so there can be no 3-source contribution. Hence, both terms in equation (20) can be interpreted as shot-noise contributions where the same galaxy appears in more than one leg – the first term comes from a single galaxy and the second, from two. If a mechanism exists for aligning galaxy spins (and, hence, polarisations) in the filamentary cosmic web (see, e.g., [Piras et al. 2018](#) and [Codis et al. 2018](#)), then there can, in principle, also be a 3-source contribution dependent on the 3D galaxy bispectrum. However, such a mechanism has been shown by [Feng & Holder \(2020\)](#) to produce contributions to the CIB power spectrum that are several orders of magnitude smaller than the shot noise, so we will be ignoring it henceforth.

In order to evaluate equation (20), we need to compute the power spectrum of the galaxies comprising the CIB emission. Galaxies are known to be biased tracers of the matter distribution ([Kaiser 1984](#)), so additional prescriptions are needed to relate their clustering statistics to those of the matter. This could be done, for instance, by applying the tools of the halo model to the CIB, as was done by [Lacasa et al. \(2014\)](#). However, given the number of approximations we have already made, we pursue a cruder approach and relate the power spectrum of galaxy fluctuations to that of the matter by means of an effective

⁸ For narrow frequency bands, this can be calculated using, for example, equation (8) of [Planck Collaboration \(2014a\)](#).

⁹ This value is more conservative than what has recently been discussed in the literature: [Lagache et al. \(2019\)](#), for example, suggest a smaller $p \approx 1\%$. The rationale behind such low figures is that the complex arrangement of magnetic field lines in dusty, star-forming galaxies is expected to bring about extensive depolarisation when integrated over the whole galaxy.

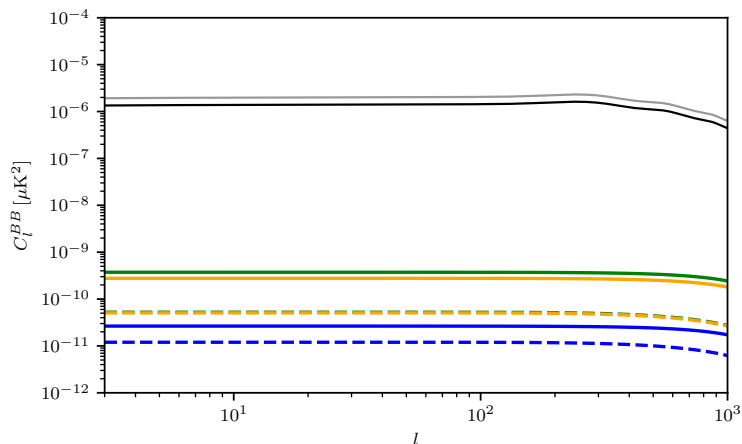


Figure 8. Delensing biases from angular bispectra of the CIB in our minimal model with uncorrelated source polarisation angles. Solid (dashed) lines show the one(two)-source contribution from different point-source masking schemes with $s_{\max} = \{200, 20, 2\}$ mJy (in green, orange and blue, respectively). We evaluate our predictions using $0.2 < z < 6$, $s_{\min} = 10^{-5}$ mJy, a fixed galaxy polarisation fraction of 5%, and the number counts of Béthermin et al. (2011). For comparison, we show also the fiducial power spectrum of lensed B -modes before (grey) and after (black) delensing with the CIB such that the power is reduced by 30%. We note that the two-source contribution is virtually unchanged if the flux cut is lowered from 200 mJy to 20 mJy, so the orange and green dashed lines lie on top of each other.

galaxy bias $b_g(k, z)$, relating the over-densities of galaxy number density and matter as $\delta_g(k, z) = b_g(k, z)\delta_m(k, z)$. We use the value measured by Planck Collaboration (2014c), which assumes that the galaxy bias is scale-independent. If, as evidence suggests, galaxy formation is chiefly a local process, then this ought to be a very good approximation on large scales (Coles 1993), where the two-halo term dominates the galaxy power spectrum. Those are precisely the scales of the power spectrum that we probe in the bispectrum configuration we are dealing with here. To see why, notice that, given \mathcal{W}_l^E decreases quickly below $l > 2000$ due to limited sensitivity, the multipole in the argument of the galaxy power spectrum in equation (20) is at most that large, which in turns forces the wavenumber $k = l/r(z)$ (where $r(z)$ is the comoving distance out to a redshift of a few, i.e., a few 1000 Mpc) to be small enough that we are probing a regime where the linear bias approximation roughly holds.

Finally, we evaluate the delensing bias of equation (20) for different flux cuts typical of current and upcoming experiments, and show the results in figure 8. We see that the one-source contribution dominates over the two-source, but it is everywhere negligible compared to the amplitude of the delensed B -mode spectrum. This is the case even though we have assumed that, beyond masking the brightest sources, no attempt is made to mitigate the CIB contribution to the CMB polarisation maps using, for example, multi-frequency cleaning techniques.

On the other hand, the bias of equation (19), sourced by the $EIEI$ trispectrum of the CIB, can, in principle, receive contributions of one-, two- and three-source type; where in all cases, the two polarisation modes are sourced by the same galaxy. As before, these can be evaluated in the framework of the halo model, this time requiring computation of the 3D halo bispectrum to obtain the three-source term. However, we argue that the trispectrum ought to be subdominant to the bispectrum (which we have shown is negligible) for the level of point-source removal expected of upcoming experiments, since the former scales as the fourth power of the flux density while the latter goes as the third. In appendix C, we motivate this claim by comparing the 1-source contributions to both higher-point functions, showing that, indeed, the relevant trispectrum is smaller than the bispectrum in this even simpler model (see figure C1).

Our conclusion is that biases sourced by higher-point functions of the CIB are negligible for the levels of point-source masking and delensing expected of any foreseen applications of the CIB-based method considered in this paper.

5 VALIDATION

In this section, we aim to validate the results of section 4.1, where we used simulations to quantify the delensing bias associated with residual Galactic dust. For reference, the power spectra of all the fields involved in the delensing procedure are shown in figure 9.

Of course, the results of section 4.1 are only valid in as much as the simulations on which they rely faithfully reproduce the true behaviour of dust on the sky. Given the dearth of small-scale polarisation data at the time of writing, it is not possible to gauge the accuracy of simulations by comparing them to observations. We can, however, test the results in two critical ways: firstly, we can ensure that the delensing bias (which, as we have already established, is dominated by the $\langle BEI \rangle_c$ bispectrum




Figure 9. Comparison of power spectrum amplitudes and biases for a number of fields involved in the delensing procedure. *Red dots*: residual dust intensity after filtering the VS simulation at 353 GHz with the GNILC CIB residual transfer function of figure 1. *Red line*: theoretical CIB intensity at 353 GHz. *Yellow dots*: VS dust E -mode power at 145 GHz scaled to the residual dust amplitude expected for the SO LAT after ILC-cleaning with dust deprojection. *Yellow line*: theoretical, lensed CMB E -mode. *Grey dots*: VS dust B -mode power at 145 GHz, scaled to the residual dust amplitude expected for the SO SAT. *Grey line*: theoretical, lensed CMB B -modes. *Purple dots*: power spectrum of a $\hat{B}(E^{\text{dust}}, I^{\text{dust}})$ template constructed with E^{dust} and I^{dust} with amplitudes as shown in this plot. *Brown dots*: power spectrum of a $\hat{B}(E^{\text{dust}}, I^{\text{CIB}})$ template. *Blue dots*: delensing bias associated with the $\langle B^{\text{dust}} E^{\text{dust}} I^{\text{dust}} \rangle_c$ bispectrum for residual levels as shown in this plot (note that this bias is negative, though here we plot it as a positive quantity). The blue, brown and purple points, along with the terms evaluated in figure 8, comprise all the contributions to equation (11) that we identified as being potentially relevant. Comparison of each of these with the total delensing bias — shown as empty squares where negative, and empty circles where positive — shows that the dust bispectrum term dominates on the scales relevant for primordial B -mode searches. All empirical measurements are carried out on the LB patch, using NaMaster and a binning scheme with $\Delta l = 90$.

of dust) does not appear when our pipeline is applied to Gaussian simulations with the same power spectra; and secondly, we can verify that the biases that ensue after either dust data from Planck or the VS dust simulations are used in the analysis display similar characteristics after both inputs are restricted to the same scales.

5.1 Validation against Gaussian simulations

In order to guarantee that the delensing bias we identified in section 4.1 – associated, recall, with the $\langle BEI \rangle_c$ bispectrum of Galactic dust – is physical in origin and not some artifact of our analysis, we re-run on Gaussian simulations of Galactic dust with the same power spectra as their non-Gaussian counterparts from VS.

We work with full-sky simulations of Galactic dust at 145 GHz. First, we fit smooth curves to the ratios of TT , EE and BB power in these simulations to those measured in the VS simulations on scales $60 < l < 2000$, after the latter have been scaled to the appropriate frequency and residual amplitude, and restricted to the region allowed by the SAT mask of figure 3a. The square root of these smooth transfer functions are then used to filter the Gaussian simulations, obtaining Gaussian maps with the same power spectra as the VS simulations.

Next, we use these re-scaled Gaussian simulations in the analysis steps described in section 3.5. As shown in figure 10, the output is consistent with pure scatter, in contrast with the distinct bias arising from the non-Gaussian simulations.

5.2 Validation against Planck data

Given that the results of section 4.1 are only as accurate as the simulations on which they rely, we set out now to validate those findings against Planck data on the scales where the latter has sufficient signal-to-noise. More specifically, we compare the output of our analysis procedure (detailed in section 3.5) when applied to the VS simulations to the case where the temperature and polarisation maps we use are instead the Planck 2018 full-mission or GNILC products, both at 353 GHz. (The Planck products we use are available on the [Planck Legacy Archive](#).)

The idea is that the dust component of the observations is to be processed in the same ways as the dust in the simulations. For this reason, we apply the same filtering protocol to both, aiming to retrieve the same residual dust fractions. In addition to this, we extrapolate the polarisation observations at 353 GHz to 145 GHz with an SED appropriate for dust, to estimate the dust at this frequency before foreground cleaning. This is done by assuming that the dust behaves as a modified blackbody with parameters given in table 1 and applying the relevant wide-band unit-conversions for the 143 and 353 GHz channels of

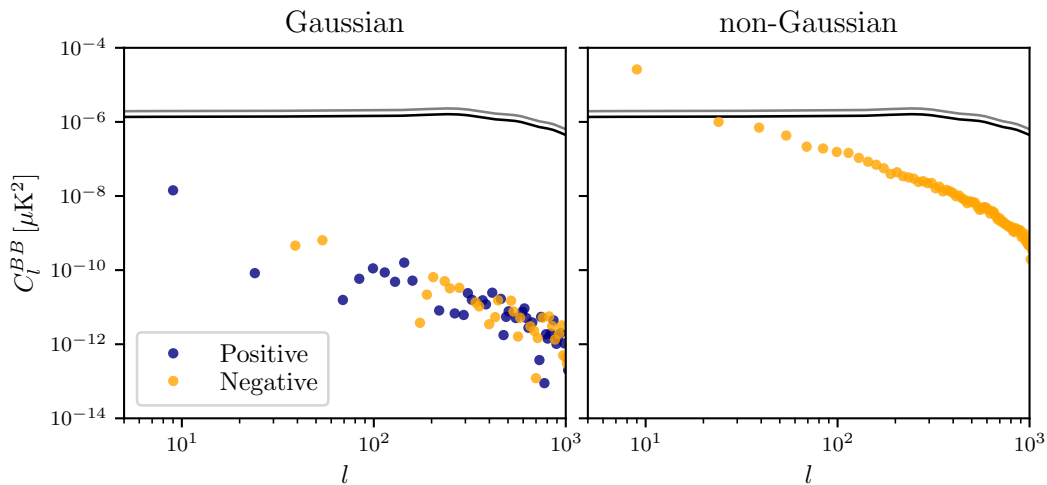


Figure 10. Delensing bias associated with the bispectrum of Galactic dust as calculated from a single realisation of the VS simulation (right panel), compared to that measured from a Gaussian simulation of the dust (left panel) with the same power spectra. The same transfer functions are applied in both cases to induce the scale-dependent dust residuals estimated in section 3.3. Spectra are measured on the “LB patch” shown in figure 3c and binned with $\Delta l = 15$. For comparison, the amplitude of the lensing B -mode spectrum is plotted before (grey) and after (black) delensing using the CIB as the only matter tracer. While the measurement on Gaussian simulations is consistent with scatter around zero, the non-Gaussian case shows a distinct bias-like behaviour.

Planck. Furthermore, specific intensity measurements made by broad-band instruments are necessarily related to a reference frequency and calibration SED profile. In the Planck 353 GHz band, this is the CMB dipole. In order to relate measurements on that band to a thermal dust SED, we apply a multiplicative colour correction as given by Planck Collaboration (2018a). An explanation of the philosophy of such corrections can be found in Planck Collaboration (2014a).

As is the case with all astronomical observations, the Planck data are convolved with the instrument’s beam. In order to ensure that the observed and simulated modes we are comparing have the same normalisation, we filter the VS simulation with approximate beam transfer functions derived by fitting the window function of a symmetric Gaussian beam to the ratios of EE , BB and TT power measured from the Planck data to those measured from the simulations, on scales of $20 < l < 100$ (below the multipoles where experimental noise becomes significant in the Planck polarisation data). This correction enables a like-for-like comparison on scales of approximately $l < 200$, in the case of the GNILC dust data (for which the maps are convolved with a beam with FWHM of 80 arcmin to ensure they have the same resolution everywhere), and $l < 500$ in the case of the raw 353 GHz frequency data (approximately the angular scale where polarisation noise in the 353 GHz channel begins to dominate). The latter is an important range for our purposes, since it encompasses most of the multipole range where Galactic dust dominates over CIB emission in total intensity (see, e.g., Planck Collaboration 2016e). Furthermore, we verify that, in the 353 GHz channel, the effects of beam convolution at $l = 500$ are still relatively limited and signal power is only suppressed by approximately 10%.

When using the raw 353 GHz maps, which include significant amounts of CIB emission, our validation could potentially be confused by its sensitivity to the bispectrum of CIB fields. However, from the calculations in section 4.2, we expect this to be small for the expected extent of point-source removal. In order to ensure this is the case, we apply a mask – also obtained from the Planck Legacy Archive – which removes point sources detected with $S/N > 5$ at 353 GHz. We do, however, see increased scatter in the bispectrum measurements from the CIB (and the CMB and instrumental noise) as discussed below.

The results of the delensing bias from the bispectrum of dust computed from the re-scaled VS simulations or the Planck data are shown in figure 11. The bias that appears when using the raw data as input is approximately two orders of magnitude larger than that produced by the GNILC dust maps, owing to the much restricted multipole range going into the latter.

For a given type of input map, the level of agreement between data and simulations is remarkable. The bias calculation based on the raw Planck frequency maps shows higher scatter than seen in the simulation (particularly on smaller scales), which can be explained as the added variance from the other components present in those maps such as the CIB, the CMB or experimental noise. This explanation is consistent with the fact that such increase in scatter is not seen when the GNILC dust products are used. On that note, it is worth pointing out that the similar amplitude and shape of the biases arising from the raw channel data and the pure-dust simulations is consistent with the claim made in section 4.2 that the BEI bispectrum of the CIB is small once point-sources are masked.

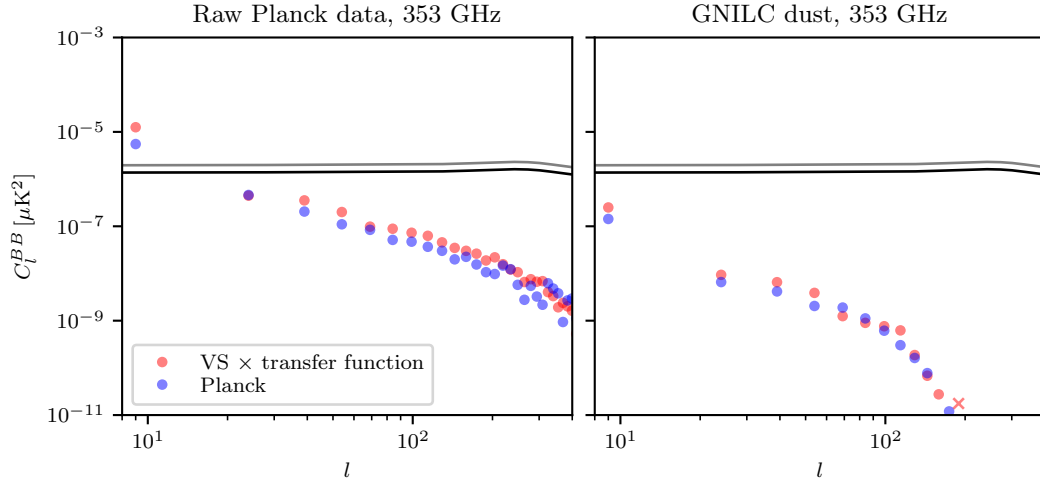


Figure 11. Delensing bias associated with the bispectrum of Galactic dust arising from a restricted range of multipoles of the VS Galactic dust simulation or the Planck data, after both are filtered and scaled to emulate levels of residual dust expected after foreground cleaning. In order to incorporate the effect of beam convolution, the VS simulation is filtered with a smooth transfer function fit to the power spectra of the data. This enables a comparison in the range $l < 200$ when the Planck 353 GHz GNILC dust data is used to model residuals (right panel) and for $l < 500$ when the raw 353 GHz channel data is used for this purpose instead (left panel). In the latter case, we ensure point-sources are masked. For comparison, the lensing B -mode spectrum is plotted before (grey) and after (black) being delensed with the CIB. The amplitude of the dust in each map is scaled down to the level appropriate for the residuals expected in the GNILC estimate of the CIB and the SO LAT and SAT polarisation maps (both after standard ILC-cleaning), as estimated in section 3.3. Furthermore, an $l_{\min} = 100$ cut is applied to the fields. Spectra are measured on the “LB patch” shown in figure 3c and binned with $\Delta l = 15$. In both cases, negative values are plotted as circles and positive ones as crosses.

6 CONCLUSIONS

We have investigated ways in which residual foregrounds can impact the power spectrum of CMB B -modes after they are delensed using the CIB as a tracer of the matter. We identified two sources of bias which could, in principle, affect constraints on the tensor-to-scalar ratio r : residual Galactic dust and residual CIB emission left over in the CMB maps after foreground cleaning.

Given that, at the time of writing, no simulations of polarised CIB emission exist, we calculated the delensing bias arising from higher-point functions of the CIB analytically, in the context of a model where source polarisation angles are uncorrelated. We found that the dominant contribution to this bias ought to be the $\langle BEI \rangle$ bispectrum of the CIB. Our calculation shows that, given the extent of point source masking expected of upcoming experiments, this bias ought to be negligible. As a side note, our model shows that certain bispectra involving CIB polarisation fields can be non-vanishing even if galaxy polarisation angles are uncorrelated, namely those containing two polarisation legs sourced by the same galaxy.

We also presented a first attempt to characterise the delensing bias arising from residual Galactic dust left over in CMB and CIB maps used for delensing. We simulated the delensing procedure as applied to these dust residuals, using a non-Gaussian simulation of Galactic dust from Vansyngel et al. (2017) filtered to model the expected levels of residuals after foreground cleaning appropriate for CMB E - and B -mode polarisation and CIB observations in the coming era of CMB experiments. In particular, we estimated residual dust levels for the Simons Observatory’s (SO) measurements of E -modes with its large-aperture telescope (LAT) and B -modes with its small-aperture telescopes, and for the CIB maps of Planck Collaboration (2016e) and Lenz et al. (2019). We also validated our results directly against the Planck data on large and intermediate scales.

We found that the bias from residual Galactic dust is dominated by the $\langle BEI \rangle_c$ bispectrum of Galactic dust, and that it is negative. It is also very red, peaking on the very largest angular scales. Given the extent of foreground cleaning expected of upcoming ground-based experiments (we focus on the specifications of SO), the bias only becomes relevant for multipoles $l \lesssim 20$, scales likely too large to be probed from the ground. Hence, we find that a failure to model this bias at the level of the power spectrum translates to an inference on r that is systematically lower than the truth, but only by an amount smaller than the 1σ statistical uncertainty. To err on the side of caution, however, we recommend explicitly deprojecting dust from the high-resolution E -mode fields – e.g., with a constrained internal-linear combination approach applied to multi-frequency data Remazeilles et al. (2011a) – used in delensing if the delensed B -modes on scales larger than $l < 30$ are to be considered. Although this might come at the cost of degraded delensing efficiency, we show that for an experiment with the characteristics of the SO LAT this is a small effect (with only a percent-level increase in the residual lensing power).

On the other hand, the bias could be large enough to be a cause of concern for space-based experiments, such as LiteBird, targeting the reionisation peak of the primary CMB in addition to the recombination peak targeted by ground-based experiments. By covering larger sky areas, space-based experiments are prone to including regions where the dust is very bright, causing the bias to grow and become significant on scales that the telescope is sensitive to. Furthermore, their situation above the atmosphere allows them to probe the large-scale signal from reionisation, scales where the bias is largest. In spite of this, we show that a combination of dust deprojection in the foreground cleaning of E -modes and removal of large-scale CIB modes in the construction of the lensing B -mode template is effective in reducing this bias. The latter can always be implemented at little cost in terms of delensing efficiency (at least for $l_{\min}^{\text{CIB}} \lesssim 250$), since internal lensing reconstruction techniques can estimate these large-scale lenses very accurately. Furthermore, the bias is very sensitive to the removal of dust from the CMB polarisation fields involved, so we expect the situation to be less concerning for future satellite missions, for which foreground cleaning will be much improved relative to the levels considered in this paper. Altogether, we are confident that the bias can be kept in check and CIB delensing can remain a valuable tool.

We note also that, in addition to the biases described above, residual foregrounds ought also to impact the variance of the power spectrum of delensed B -modes. However, we defer the study of that problem to future work, given the difficulty of studying it analytically and the lack of realistic simulations to assist a computational treatment.

The results presented in this paper illustrate the community's need for non-Gaussian simulations of Galactic dust extending to small angular scales on which to test the robustness of delensing analyses against the systematic effects identified in this work. Such simulations are also necessary to assess the level of contamination in CMB lensing power spectrum reconstruction from polarisation-based estimators (Challinor et al. 2018; Beck et al. 2020). We emphasise that the approach followed to date of testing against Gaussian simulations (see, e.g., Manzotti et al. 2017 or BICEP/Keck & SPTpol Collaborations 2020) will not suffice, as validation experiments based on those will be insensitive to the non-Gaussian biases discussed here, which can be significant on the largest angular scales.

ACKNOWLEDGEMENTS

The authors would like to thank Alex Van Engelen, Colin Hill, Byeonghee Yu, Will Coulton and Mathieu Remazeilles for useful discussions and their sharing of products that enabled this work. ABL gratefully acknowledges support from an Isaac Newton Studentship at the University of Cambridge and the Science and Technology Facilities Council (STFC). AC acknowledges support from the STFC (grant numbers ST/N000927/1 and ST/S000623/1). TN and BDS acknowledge support from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (Grant agreement No. 851274). BDS further acknowledges support from an STFC Ernest Rutherford Fellowship. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility. It also benefited from use of the publicly-available codes HealPy (Zonca et al. 2019), HealPix (Górski et al. 2005), AstroPy (Astropy Collaboration 2013), CAMB (Lewis et al. 2000) and NaMaster (Alonso et al. 2019). This is not an official SO Collaboration paper.

DATA AVAILABILITY

The data underlying this article will be shared on reasonable request to the corresponding author. Where the data were provided by a third party (as described in the text), the data will be shared subject to permission from the third party.

REFERENCES

- Abazajian K., et al., 2019, arXiv e-prints, p. [arXiv:1907.04473](https://arxiv.org/abs/1907.04473)
 Adachi S., et al., 2019, arXiv e-prints, p. [arXiv:1909.13832](https://arxiv.org/abs/1909.13832)
 Alonso D., Dunkley J., Thorne B., Næss S., 2017, *Phys. Rev. D*, **95**, 043504
 Alonso D., Sanchez J., Slosar A., LSST Dark Energy Science Collaboration 2019, *MNRAS*, **484**, 4127
 Astropy Collaboration 2013, *A&A*, **558**, A33
 BICEP2 Collaboration Keck Array Collaboration 2018, *Phys. Rev. Lett.*, **121**, 221301
 BICEP/Keck SPTpol Collaborations 2020, arXiv e-prints, p. [arXiv:2011.08163](https://arxiv.org/abs/2011.08163)
 Baleato Lizancos A., Challinor A., Carron J., 2021, *Phys. Rev. D*, **103**, 023518
 Beck D., Errard J., Stompor R., 2020, *J. Cosmology Astropart. Phys.*, **2020**, 030
 Béthermin M., Dole H., Lagache G., Le Borgne D., Penin A., 2011, *A&A*, **529**, A4
 Béthermin M., et al., 2012, *A&A*, **542**, A58
 Buchalter A., Kamionkowski M., Jaffe A. H., 2000, *ApJ*, **530**, 36
 Bunn E. F., Zaldarriaga M., Tegmark M., de Oliveira-Costa A., 2003, *Phys. Rev. D*, **67**, 023501
 Carron J., Lewis A., 2017, *Phys. Rev. D*, **96**, 063510
 Carron J., Lewis A., Challinor A., 2017, *Journal of Cosmology and Astro-Particle Physics*, **2017**, 035

- Challinor A., Lewis A., 2005, *Phys. Rev. D*, **71**, 103010
- Challinor A., et al., 2018, *J. Cosmology Astropart. Phys.*, **2018**, 018
- Clark S. E., Peek J. E. G., Putman M. E., 2014, *ApJ*, **789**, 82
- Clark S. E., Hill J. C., Peek J. E. G., Putman M. E., Babler B. L., 2015, *Phys. Rev. Lett.*, **115**, 241302
- Codis S., Jindal A., Chisari N. E., Vibert D., Dubois Y., Pichon C., Devriendt J., 2018, *MNRAS*, **481**, 4753
- Coles P., 1993, *MNRAS*, **262**, 1065
- Feng C., Holder G., 2020, *ApJ*, **897**, 140
- Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, **622**, 759
- HI4PI Collaboration 2016, *A&A*, **594**, A116
- Hall N. R., et al., 2010, *ApJ*, **718**, 632
- Han D., et al., 2020, arXiv e-prints, p. arXiv:2007.14405
- Hazumi M., et al., 2019, *Journal of Low Temperature Physics*, **194**, 443
- Hirata C. M., Seljak U., 2003, *Phys. Rev. D*, **68**, 083002
- Holder G. P., et al., 2013, *ApJ*, **771**, L16
- Kaiser N., 1984, *ApJ*, **284**, L9
- Kamionkowski M., Kosowsky A., Stebbins A., 1997, *Phys. Rev. Lett.*, **78**, 2058
- Lacasa F., Pénin A., Aghanim N., 2014, *MNRAS*, **439**, 123
- Lagache G., Béthermin M., Montier L., Serra P., Tucci M., 2019, arXiv e-prints, p. arXiv:1911.09466
- Larsen P., Challinor A., Sherwin B. D., Mak D., 2016, *Phys. Rev. Lett.*, **117**, 151102
- Lenz D., Doré O., Lagache G., 2019, *ApJ*, **883**, 75
- Lewis A., Challinor A., 2006, *Phys. Rep.*, **429**, 1
- Lewis A., Challinor A., Lasenby A., 2000, *ApJ*, **538**, 473
- Lewis A., Challinor A., Turok N., 2002, *Phys. Rev. D*, **65**, 023505
- Mak D. S. Y., Challinor A., Efstathiou G., Lagache G., 2017, *MNRAS*, **466**, 286
- Maniyar A. S., Béthermin M., Lagache G., 2018, *A&A*, **614**, A39
- Maniyar A., Lagache G., Béthermin M., Ilić S., 2019, *A&A*, **621**, A32
- Manzotti A., 2018, *Phys. Rev. D*, **97**, 043527
- Manzotti A., et al., 2017, *ApJ*, **846**, 45
- Millea M., Anderes E., Wandelt B. D., 2019, *Phys. Rev. D*, **100**, 023509
- Namikawa T., et al., 2022, *Phys. Rev. D*, **105**, 023511
- Okamoto T., Hu W., 2003, *Phys. Rev. D*, **67**, 083002
- Piras D., Joachimi B., Schäfer B. M., Bonamigo M., Hilbert S., van Uitert E., 2018, *MNRAS*, **474**, 1165
- Planck Collaboration 2011, *A&A*, **536**, A18
- Planck Collaboration 2014a, *A&A*, **571**, A9
- Planck Collaboration 2014b, *A&A*, **571**, A18
- Planck Collaboration 2014c, *A&A*, **571**, A30
- Planck Collaboration 2016a, *A&A*, **594**, A12
- Planck Collaboration 2016b, *A&A*, **594**, A15
- Planck Collaboration 2016c, *A&A*, **594**, A26
- Planck Collaboration 2016d, *A&A*, **596**, A105
- Planck Collaboration 2016e, *A&A*, **596**, A109
- Planck Collaboration 2018a, arXiv e-prints, p. arXiv:1801.04945
- Planck Collaboration 2018b, arXiv e-prints, p. arXiv:1807.06205
- Planck Collaboration 2018c, arXiv e-prints, p. arXiv:1807.06210
- Poh J., Dodelson S., 2017, *Phys. Rev. D*, **95**, 103511
- Polnarev A. G., 1985, *Soviet Ast.*, **29**, 607
- Puget J. L., Abergel A., Bernard J. P., Boulanger F., Burton W. B., Desert F. X., Hartmann D., 1996, *A&A*, **308**, L5
- Remazeilles M., Delabrouille J., Cardoso J.-F., 2011a, *MNRAS*, **410**, 2481
- Remazeilles M., Delabrouille J., Cardoso J.-F., 2011b, *MNRAS*, **418**, 467
- Roy A., Kulkarni G., Meerburg P. D., Challinor A., Baccigalupi C., Lapi A., Haehnelt M. G., 2021, *J. Cosmology Astropart. Phys.*, **2021**, 003
- SPT Collaboration 2013, *Phys. Rev. Lett.*, **111**, 141301
- Scherrer R. J., Bertschinger E., 1991, *ApJ*, **381**, 349
- Seljak U., Zaldarriaga M., 1997, *Phys. Rev. Lett.*, **78**, 2054
- Serra P., Lagache G., Doré O., Pullen A., White M., 2014, *A&A*, **570**, A98
- Sherwin B. D., Schmittfull M., 2015, *Phys. Rev. D*, **92**, 043005
- Smith K. M., 2006, *Phys. Rev. D*, **74**, 083002
- Smith K. M., Hanson D., LoVerde M., Hirata C. M., Zahn O., 2012, *J. Cosmology Astropart. Phys.*, **2012**, 014
- Song Y.-S., Cooray A., Knox L., Zaldarriaga M., 2003, *ApJ*, **590**, 664
- Stompor R., Errard J., Poletti D., 2016, *Phys. Rev. D*, **94**, 083526
- Sultan Abylkairov Y., Darwish O., Hill J. C., Sherwin B. D., 2020, arXiv e-prints, p. arXiv:2012.04032
- Tassis K., Pavlidou V., 2015, *MNRAS*, **451**, L90
- Tegmark M., de Oliveira-Costa A., Hamilton A. J., 2003, *Phys. Rev. D*, **68**, 123523
- The Simons Observatory Collaboration 2019, *J. Cosmology Astropart. Phys.*, **2019**, 056
- Thorne B., Dunkley J., Alonso D., Naess S., 2017, *MNRAS*, **469**, 2821
- Tristram M., et al., 2020, arXiv e-prints, p. arXiv:2010.01139
- Vansyngel F., et al., 2017, *A&A*, **603**, A62
- Viero M. P., et al., 2013a, *ApJ*, **772**, 77

Viero M. P., et al., 2013b, *ApJ*, 779, 32

Yu B., Hill J. C., Sherwin B. D., 2017, *Phys. Rev. D*, 96, 123511

Zaldarriaga M., Seljak U., 1998, *Phys. Rev.*, D58, 023003

Zonca A., Singer L., Lenz D., Reinecke M., Rosset C., Hivon E., Gorski K., 2019, *Journal of Open Source Software*, 4, 1298

APPENDIX A: THE BEI BISPECTRUM OF THE CIB

In order to model the *BEI* bispectrum of the CIB, we begin by considering the specific intensity observed at frequency ν , along a line of sight $\hat{\mathbf{n}}$. We shall work in the flat-sky limit, in which case $\hat{\mathbf{n}}$ is close to the pole and the 3D comoving position of a source, \mathbf{x}_{3D} , can be related to its angular position in the plane of the sky, $\boldsymbol{\theta}$, as $\mathbf{x}_{3D} \approx r(z)(\boldsymbol{\theta} + \hat{\mathbf{z}})$, where $r(z)$ is the comoving distance to the redshift where the source is located and $\hat{\mathbf{z}}$ is a unit vector along the z -direction. In this limit, the observed temperature is an integral along the line of sight of j_ν , the emissivity per unit comoving volume and per solid angle :

$$I_\nu(\boldsymbol{\theta}) = \int dz \frac{dr}{dz} a(z) j_{\nu(1+z)}(r(z)(\boldsymbol{\theta} + \hat{\mathbf{z}}), z). \quad (\text{A1})$$

Similarly, the Stokes parameters for linear polarisation can be related to $j_\nu^{(\pm)}$, the polarised emissivity per unit comoving volume along the negative z -direction, via

$$(Q \pm iU)_\nu(\boldsymbol{\theta}) = p \int dz \frac{dr}{dz} a(z) j_{\nu(1+z)}^{(\pm)}(r(z)(\boldsymbol{\theta} + \hat{\mathbf{z}}), z), \quad (\text{A2})$$

where p is the polarisation fraction of the sources, which we will take to be a constant. In the flat-sky limit, the Fourier modes of polarisation can be extracted as

$$\begin{aligned} (Q \pm iU)_\nu(\mathbf{l}) &= \int d^2\boldsymbol{\theta} (Q \pm iU)_\nu(\boldsymbol{\theta}) e^{-i\mathbf{l}\cdot\boldsymbol{\theta}} \\ &= p \int dz \frac{dr}{dz} a(z) \int \frac{d^3\mathbf{k}}{(2\pi)^3} j_{\nu(1+z)}^{(\pm)}(\mathbf{k}, z) \int d^2\boldsymbol{\theta} e^{ir(z)\mathbf{k}\cdot(\boldsymbol{\theta}+\hat{\mathbf{z}})} e^{-i\mathbf{l}\cdot\boldsymbol{\theta}} \\ &= p \int dz \frac{dr}{dz} a(z) \int \frac{d^3\mathbf{k}}{2\pi} j_{\nu(1+z)}^{(\pm)}(\mathbf{k}, z) \delta_D^{(2)}(\mathbf{k}_\perp r(z) - \mathbf{l}) e^{ik_3 r(z)} \\ &= p \int dz \frac{dr}{dz} \frac{a(z)}{r^2(z)} \int \frac{dk_3}{2\pi} e^{ik_3 r(z)} j_{\nu(1+z)}^{(\pm)}(\mathbf{l}/r(z) + k_3 \hat{\mathbf{z}}, z), \end{aligned} \quad (\text{A3})$$

and, similarly, for the intensity:

$$I_\nu(\mathbf{l}) = \int dz \frac{dr}{dz} \frac{a(z)}{r^2(z)} \int \frac{dk_3}{(2\pi)} e^{ik_3 r(z)} j_{\nu(1+z)}(\mathbf{l}/r(z) + k_3 \hat{\mathbf{z}}, z). \quad (\text{A4})$$

Here, \mathbf{k}_\perp is the projection of the wavevector \mathbf{k} perpendicular to the z -direction. We are now in a position to extract *E* and *B* modes. In the flat-sky approximation, these can be constructed as

$$(E \pm iB)(\mathbf{l}) = -(Q \pm iU)(\mathbf{l}) e^{\mp 2i\psi_{\mathbf{l}}}, \quad (\text{A5})$$

where $\psi_{\mathbf{l}}$ is the angle between the Fourier wavevector \mathbf{l} and the x -direction. Substituting in from equations (A3) and (A4), we obtain

$$E_\nu(\mathbf{l}) = -\frac{p}{2} \int dz \frac{dr}{dz} \frac{a(z)}{r^2(z)} \frac{dk_3}{(2\pi)} e^{ik_3 r(z)} \left[j_{\nu(1+z)}^{(+)}(\mathbf{l}/r(z) + k_3 \hat{\mathbf{z}}, z) e^{-2i\psi_{\mathbf{l}}} + j_{\nu(1+z)}^{(-)}(\mathbf{l}/r(z) + k_3 \hat{\mathbf{z}}, z) e^{2i\psi_{\mathbf{l}}} \right] \quad (\text{A6})$$

and

$$B_\nu(\mathbf{l}) = \frac{ip}{2} \int dz \frac{dr}{dz} \frac{a(z)}{r^2(z)} \frac{dk_3}{(2\pi)} e^{ik_3 r(z)} \left[j_{\nu(1+z)}^{(+)}(\mathbf{l}/r(z) + k_3 \hat{\mathbf{z}}, z) e^{-2i\psi_{\mathbf{l}}} - j_{\nu(1+z)}^{(-)}(\mathbf{l}/r(z) + k_3 \hat{\mathbf{z}}, z) e^{2i\psi_{\mathbf{l}}} \right]. \quad (\text{A7})$$

The bispectrum can therefore be constructed as

$$\begin{aligned} \langle B_{\nu_1}(\mathbf{l}_1) E_{\nu_2}(\mathbf{l}_2) I_{\nu_3}(\mathbf{l}_3) \rangle &= \frac{-ip^2}{2(2\pi)^3} \int \prod_{j=1}^3 dz_j dk_{j,3} \frac{dr(z_j)}{dz_j} \frac{a(z_j)}{r^2(z_j)} e^{ik_{j,3} r(z_j)} \left[e^{-2i(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2})} \langle j_{\nu_1(1+z_1)}^{(+)}(\mathbf{l}_1/r(z_1) + k_{1,3} \hat{\mathbf{z}}, z_1) \right. \\ &\quad \times j_{\nu_2(1+z_2)}^{(-)}(\mathbf{l}_2/r(z_2) + k_{2,3} \hat{\mathbf{z}}, z_2) j_{\nu_3(1+z_3)}(\mathbf{l}_3/r(z_3) + k_{3,3} \hat{\mathbf{z}}, z_3) \rangle \\ &\quad - e^{2i(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2})} \langle j_{\nu_1(1+z_1)}^{(-)}(\mathbf{l}_1/r(z_1) + k_{1,3} \hat{\mathbf{z}}, z_1) \\ &\quad \times j_{\nu_2(1+z_2)}^{(+)}(\mathbf{l}_2/r(z_2) + k_{2,3} \hat{\mathbf{z}}, z_2) j_{\nu_3(1+z_3)}(\mathbf{l}_3/r(z_3) + k_{3,3} \hat{\mathbf{z}}, z_3) \rangle \left. \right], \end{aligned} \quad (\text{A8})$$

where we have only retained those terms that survive averaging over independent polarisation directions (see appendix B).

Assuming the 3D bispectrum we have encountered is homogeneous and isotropic, we can write it as

$$\begin{aligned} \langle \prod_{j=1}^3 j_{\nu_j}^{(s_j)}(\mathbf{l}_j/r(z_j) + k_{j,3}\hat{\mathbf{z}}, z_j) \rangle &= (2\pi)^3 \delta_{\mathbf{D}}^{(3)}(\mathbf{l}_1/r(z_1) + \mathbf{l}_2/r(z_2) + \mathbf{l}_3/r(z_3) + (k_{1,3} + k_{2,3} + k_{3,3})\hat{\mathbf{z}}) \\ &\quad \times B_{\nu'_1\nu'_2\nu'_3}^{s_1s_2s_3} \left(|\mathbf{l}_1/r(z_1) + k_{1,3}\hat{\mathbf{z}}|, |\mathbf{l}_2/r(z_2) + k_{2,3}\hat{\mathbf{z}}|, |\mathbf{l}_3/r(z_3) + k_{3,3}\hat{\mathbf{z}}|; z_1, z_2, z_3 \right). \end{aligned} \quad (\text{A9})$$

For the sake of compactness, we have introduced the shorthand s_j to label the (\pm) polarisation and the intensity. Additionally, we have defined $\nu' = \nu(1+z)$. The three-dimensional delta function can be factored as

$$\delta_{\mathbf{D}}^{(3)}(\mathbf{l}_1/r(z_1) + \mathbf{l}_2/r(z_2) + \mathbf{l}_3/r(z_3) + (k_{1,3} + k_{2,3} + k_{3,3})\hat{\mathbf{z}}) = \delta_{\mathbf{D}}^{(2)}(\mathbf{l}_1/r(z_1) + \mathbf{l}_2/r(z_2) + \mathbf{l}_3/r(z_3)) \delta_{\mathbf{D}}(k_{1,3} + k_{2,3} + k_{3,3}). \quad (\text{A10})$$

At this stage, we invoke the Limber approximation (see, e.g., [Buchalter et al. 2000](#) in the context of the projected bispectrum), whereby the radial integrations enforce $k_{j,3} \ll |\mathbf{l}_j|/r(z_j)$ such that the bispectrum is approximately independent of the $k_{j,3}$. Performing the integrals over the $k_{j,3}$ in equation (A8) gives rise to delta functions in r_j , which in turn require that all contributions come from the same redshifts; we find that

$$\begin{aligned} \langle B_{\nu_1}(\mathbf{l}_1)E_{\nu_2}(\mathbf{l}_2)I_{\nu_3}(\mathbf{l}_3) \rangle &= -\frac{p^2}{2} \sin 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2})(2\pi)^2 \delta_{\mathbf{D}}^{(2)}(\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l}_3) \\ &\quad \times \int dr_1 dr_2 dr_3 a(z_1) \frac{a(z_2)}{r_2^2} \frac{a(z_3)}{r_3^2} \delta_{\mathbf{D}}(r_1 - r_2) \delta_{\mathbf{D}}(r_1 - r_3) B_{\nu'_1\nu'_2\nu'_3}^{(+)(-)(0)} \left(\frac{l_1}{r_1}, \frac{l_2}{r_2}, \frac{l_3}{r_3}; z_1, z_2, z_3 \right) \\ &= -\frac{p^2}{2} \sin 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2})(2\pi)^2 \delta_{\mathbf{D}}^{(2)}(\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l}_3) \int dz \frac{dr}{dz} \frac{a^3(z)}{r^4(z)} B_{\nu'_1\nu'_2\nu'_3}^{(+)(-)(0)} \left(\frac{l_1}{r(z)}, \frac{l_2}{r(z)}, \frac{l_3}{r(z)}; z \right), \end{aligned} \quad (\text{A11})$$

where we have used

$$\delta_{\mathbf{D}}^{(2)}(\mathbf{l}_1/r(z) + \mathbf{l}_2/r(z) + \mathbf{l}_3/r(z)) = r^2(z) \delta_{\mathbf{D}}^{(2)}(\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l}_3). \quad (\text{A12})$$

The next step is to calculate the equal-redshift 3D galaxy bispectrum. This is done in appendix B. Using its explicit form from equation (B13) and using $dr/dz = c/H(z)$ to simplify, we obtain

$$\langle B_{\nu}(\mathbf{l}_1)E_{\nu}(\mathbf{l}_2)I_{\nu}(\mathbf{l}_3) \rangle = -\frac{p^2}{2} \sin 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2})(2\pi)^2 \delta_{\mathbf{D}}^{(2)}(\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l}_3) \int dz \left[S_{\nu}^{(3)}(z) + \frac{H(z)}{cr^2(z)} S_{\nu}^{(2)}(z) S_{\nu}^{(1)}(z) P_g \left(\frac{l_3}{r(z)}; z \right) \right], \quad (\text{A13})$$

with $S_{\nu}^{(n)}(z)$ as defined in equation (21). This expression was derived by assuming that the SED of all sources is the same at a given redshift. By introducing in the integral over redshift appropriate ratios of the SED of the galaxies' luminosity at the redshifted frequencies, it can be generalised to a scenario where the different fields are observed at different frequencies.

APPENDIX B: MODELLING THE EQUAL-REDSHIFT 3D BISPECTRUM

We now calculate the equal-redshift 3D bispectrum of polarised emissivities under the assumption that all sources have the same SED. For convenience, we consider the case of all three frequencies being the same, which we can easily generalise to mixed frequencies later on by scaling with appropriate ratios across frequency of the sources' SED.

The emissivity at position \mathbf{x} , redshift z and frequency ν is

$$j_{\nu}(\mathbf{x}, z) = \frac{1}{4\pi} \sum_i L_{\nu,i} \delta_{\mathbf{D}}^{(3)}(\mathbf{x} - \mathbf{y}_i) = \frac{1}{4\pi} \sum_i \int dL_{\nu} L_{\nu} \delta_{\mathbf{D}}(L_{\nu} - L_{\nu,i}) \delta_{\mathbf{D}}^{(3)}(\mathbf{x} - \mathbf{y}_i), \quad (\text{B1})$$

where $L_{\nu,i}$ is the spectral luminosity of the source labeled by the subscript i , and \mathbf{y}_i its position, both at redshift z . Note that the expectation value

$$\langle \delta_{\mathbf{D}}(L_{\nu} - L_{\nu,i}(z)) \delta_{\mathbf{D}}^{(3)}(\mathbf{x} - \mathbf{y}_i) \rangle = n_g(L_{\nu}, z), \quad (\text{B2})$$

where $n_g(L_{\nu}, z) \equiv dN/dVdL_{\nu}$ is the number density of galaxies per comoving volume and specific luminosity range, at redshift z . Similarly, the polarised emissivity is

$$j_{\nu}^{(\pm)}(\mathbf{x}, z) = \frac{1}{4\pi} \sum_i L_{\nu,i} \delta_{\mathbf{D}}^{(3)}(\mathbf{x} - \mathbf{y}_i) e^{\pm 2i\theta_i}, \quad (\text{B3})$$

where θ describes the direction of the linear polarisation in the plane of the sky for emission along the $-z$ direction. As will immediately become apparent, the only possible independent three-point function involving two factors of polarised emissivity

is

$$\begin{aligned}
 (4\pi)^3 \langle j_{\nu}^{(+)}(\mathbf{x}_1, z) j_{\nu}^{(-)}(\mathbf{x}_2, z) j_{\nu}(\mathbf{x}_3, z) \rangle &= \left\langle \sum_{ijk} \int dL'_\nu dL''_\nu dL'''_\nu L'_\nu L''_\nu L'''_\nu \delta_D(L'_\nu - L_{\nu,i}) \delta_D(L''_\nu - L_{\nu,j}) \delta_D(L'''_\nu - L_{\nu,k}) \right. \\
 &\quad \times \delta_D^{(3)}(\mathbf{x}_1 - \mathbf{y}_i) \delta_D^{(3)}(\mathbf{x}_2 - \mathbf{y}_j) \delta_D^{(3)}(\mathbf{x}_3 - \mathbf{y}_k) e^{2i(\theta_i - \theta_j)} \rangle \\
 &= \delta_D^{(3)}(\mathbf{x}_1 - \mathbf{x}_2) \int dL'_\nu dL'''_\nu (L'_\nu)^2 L'''_\nu \\
 &\quad \times \left\langle \sum_i \delta_D^{(3)}(\mathbf{x}_1 - \mathbf{y}_i) \delta_D(L'_\nu - L_{\nu,i}) \sum_k \delta_D^{(3)}(\mathbf{x}_3 - \mathbf{y}_k) \delta_D(L'''_\nu - L_{\nu,k}) \right\rangle. \tag{B4}
 \end{aligned}$$

In going to the last line, we have used the fact that averaging over the polarisation angles θ_i imposes δ_{ij} . This is to be interpreted as saying that, since the polarisation angle is uncorrelated between sources, the only possible contributions come from terms where a polarisation leg can be matched to another polarisation leg from the same source. This logic has consequences for other interesting bispectra such as $\langle BII \rangle$ or $\langle EII \rangle$, which necessarily vanish. Furthermore, we also learn that any bispectrum with three polarisation legs – which we might have naïvely expected to accept one-source contributions – vanish upon averaging as they involve first and third powers of sines and cosines. Of course, this all only holds under the assumption that the polarisation angles of galaxies are uncorrelated.

We can relate the expectation value in equation (B4) to the galaxy density and the two-point correlation function; following Scherrer & Bertschinger (1991), we have

$$\begin{aligned}
 \left\langle \sum_i \delta_D^{(3)}(\mathbf{x}_1 - \mathbf{y}_i) \delta_D(L'_\nu - L_{\nu,i}) \sum_k \delta_D^{(3)}(\mathbf{x}_3 - \mathbf{y}_k) \delta_D(L'''_\nu - L_{\nu,k}) \right\rangle &= n_g(L'_\nu, z) \delta_D^{(3)}(\mathbf{x}_1 - \mathbf{x}_3) \delta_D(L'_\nu - L'''_\nu) \\
 &\quad + n_g(L'_\nu, z) n_g(L'''_\nu, z) [1 + \xi(|\mathbf{x}_1 - \mathbf{x}_3|, L'_\nu, L'''_\nu; z)]. \tag{B5}
 \end{aligned}$$

We shall take the galaxy correlation function, ξ , to be independent of the luminosity of the galaxies involved from here on. Using the relation above in equation (B4), we have that

$$\langle j_{\nu}^{(+)}(\mathbf{x}_1, z) j_{\nu}^{(-)}(\mathbf{x}_2, z) j_{\nu}(\mathbf{x}_3, z) \rangle = (4\pi)^{-3} \delta_D^{(3)}(\mathbf{x}_1 - \mathbf{x}_2) \left(\delta_D^{(3)}(\mathbf{x}_1 - \mathbf{x}_3) L_{\nu}^{(3)}(z) + [1 + \xi(|\mathbf{x}_1 - \mathbf{x}_3|; z)] L_{\nu}^{(2)}(z) L_{\nu}^{(1)}(z) \right), \tag{B6}$$

where we have defined

$$L_{\nu}^{(n)}(z) \equiv \int dL_{\nu} L_{\nu}^n \frac{dN}{dV dL_{\nu}}. \tag{B7}$$

The bispectrum can be easily obtained from the three-point function by expressing $j_{\nu}(\mathbf{x}, z)$ in terms of its Fourier transform. This yields

$$\begin{aligned}
 (4\pi)^3 \langle j_{\nu}^{(+)}(\mathbf{k}_1, z) j_{\nu}^{(-)}(\mathbf{k}_2, z) j_{\nu}(\mathbf{k}_3, z) \rangle &= \int d^3\mathbf{x}_1 d^3\mathbf{x}_2 d^3\mathbf{x}_3 \langle j_{\nu}^{(+)}(\mathbf{x}_1, z) j_{\nu}^{(-)}(\mathbf{x}_2, z) j_{\nu}(\mathbf{x}_3, z) \rangle e^{-i(\mathbf{k}_1 \cdot \mathbf{x}_1 + \mathbf{k}_2 \cdot \mathbf{x}_2 + \mathbf{k}_3 \cdot \mathbf{x}_3)} \\
 &= L_{\nu}^{(3)}(z) (2\pi)^3 \delta_D^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) + L_{\nu}^{(2)}(z) L_{\nu}^{(1)}(z) \left[(2\pi)^6 \delta_D^{(3)}(\mathbf{k}_1 + \mathbf{k}_2) \delta_D^{(3)}(\mathbf{k}_3) \right. \\
 &\quad \left. + \int d^3\mathbf{x}_1 d^3\mathbf{x}_3 \xi(|\mathbf{x}_1 - \mathbf{x}_3|; z) e^{-i(\mathbf{k}_1 + \mathbf{k}_2) \cdot \mathbf{x}_1 - i\mathbf{k}_3 \cdot \mathbf{x}_3} \right]. \tag{B8}
 \end{aligned}$$

The first term in the last line corresponds to the one-source contribution, and the last term to the two-source. The term involving $\delta_D^{(3)}(\mathbf{k}_3)$ arises from a disconnected contraction; we can ignore this term since it will not contribute to the BEI bispectrum for multipoles of l with $l_3 > 0$. Recalling that the galaxy correlation function, $\xi(\mathbf{x})$, is the Fourier transform of the galaxy power spectrum, $P_g(k)$, or more explicitly

$$\xi(|\mathbf{x}_1 - \mathbf{x}_3|; z) = \int \frac{d^3\mathbf{k}}{(2\pi)^3} P_g(k; z) e^{i\mathbf{k} \cdot (\mathbf{x}_1 - \mathbf{x}_3)}, \tag{B9}$$

we can finally write

$$\langle j_{\nu'}^{(+)}(\mathbf{k}_1, z) j_{\nu'}^{(-)}(\mathbf{k}_2, z) j_{\nu'}(\mathbf{k}_3, z) \rangle = (4\pi)^{-3} (2\pi)^3 \delta_D^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \left[L_{\nu'}^{(3)}(z) + L_{\nu'}^{(2)}(z) L_{\nu'}^{(1)}(z) P_g(|\mathbf{k}_3|; z) \right], \tag{B10}$$

where, recall, $\nu' = \nu(1+z)$. The other non-vanishing term, $\langle j_{\nu'}^{(-)}(\mathbf{k}_1, z) j_{\nu'}^{(+)}(\mathbf{k}_2, z) j_{\nu'}(\mathbf{k}_3, z) \rangle$, has the same value.

It will prove convenient to rewrite this expression in terms of observed flux-densities, S_{ν} , rather than intrinsic spectral luminosities, $L_{\nu'}$. These are related via the luminosity distance, d_L , as

$$L_{\nu'} = 4\pi(1+z)^{-1} d_L^2 S_{\nu}. \tag{B11}$$

In a flat cosmology, the luminosity distance is very simply related to the comoving distance as $d_L = (1+z)r$. Hence

$$\begin{aligned} L_{\nu'}^{(n)}(z) &= \int dL_{\nu'} L_{\nu'}^n \frac{dN}{dV dL_{\nu'}} \\ &= r^{-2} \left(\frac{dr}{dz} \right)^{-1} \int dS_{\nu'} \left(4\pi(1+z)r^2 S_{\nu'} \right)^n \frac{dN}{dS_{\nu'} dz d\Omega} \\ &= (4\pi)^n (1+z)^n r^{2n-2} \left(\frac{dr}{dz} \right)^{-1} S_{\nu'}^{(n)}(z), \end{aligned} \quad (\text{B12})$$

with $S_{\nu'}^{(n)}(z)$ as defined in equation (21) – a quantity that can be evaluated from observed number counts.

Finally, we can cast equation (B10) in the following, more convenient, form

$$\langle j_{\nu'}^{(+)}(\mathbf{k}_1, z) j_{\nu'}^{(-)}(\mathbf{k}_2, z) j_{\nu'}(\mathbf{k}_3, z) \rangle = (1+z)^3 r^4(z) (2\pi)^3 \delta_D^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \left(\frac{dr}{dz} \right)^{-1} \left[S_{\nu'}^{(3)}(z) + \frac{1}{r^2(z)} \left(\frac{dr}{dz} \right)^{-1} S_{\nu'}^{(2)}(z) S_{\nu'}^{(1)}(z) P_g(|\mathbf{k}_3|; z) \right]. \quad (\text{B13})$$

APPENDIX C: TOY MODEL FOR CIB HIGHER-POINT FUNCTIONS

In this section, we aim to motivate the claim that, for the extent of point-source removal expected of upcoming CMB experiments, the *EIEI* trispectrum of the CIB ought to result in a smaller bias than the *BEI* bispectrum. We do so by studying the one-source contribution in a minimal model where sources locations on the sky are assumed to be random and uncorrelated. Similarly, we assume that source flux-density, polarisation fraction and polarisation angle are all uncorrelated.

If we consider only one-source contributions where all legs in the correlator come from the same galaxy, the biases of equations (18) and (19) reduce to

$$\mathcal{B}_l^{BEI} = 2p^2 G_{353 \text{ GHz}}^2 G_{145 \text{ GHz}}^2 \int \frac{d^2 l'}{(2\pi)^2} \frac{l' \cdot (l - l')}{|l - l'|^2} \mathcal{W}_{l'}^E \mathcal{W}_{|l-l'|}^I \sin^2 2(\psi_{l'} - \psi_l) \frac{1}{A_{\text{sky}}} \langle \sum_i S_{145 \text{ GHz}, i}^2 S_{353 \text{ GHz}, i} \rangle \quad (\text{C1})$$

and

$$\begin{aligned} \mathcal{B}_{l,1s}^{EIEI} &= -2p^2 G_{353 \text{ GHz}}^2 G_{145 \text{ GHz}}^2 \frac{1}{A_{\text{sky}}} \langle \sum_i S_{145 \text{ GHz}, i}^2 S_{353 \text{ GHz}, i} \rangle \\ &\quad \times \int \frac{d^2 l' d^2 l''}{(2\pi)^4} \frac{l' \cdot (l - l')}{|l - l'|^2} \frac{l'' \cdot (l + l'')}{|l + l''|^2} \mathcal{W}_{l'}^E \mathcal{W}_{l''}^E \mathcal{W}_{|l-l'|}^I \mathcal{W}_{|l+l''|}^I \sin 2(\psi_{l'} - \psi_l) \sin 2(\psi_{l''} - \psi_l) \cos 2(\psi_{l'} - \psi_{l''}), \end{aligned} \quad (\text{C2})$$

where the flux density at 145 GHz, $S_{145 \text{ GHz}}$, is obtained by extrapolating the value at 353 GHz, $S_{353 \text{ GHz}}$, according to the SED of a thermal dust emitter with temperature T_{dust} and spectral index β_{dust} as given in table 1. Details of the derivations of the relevant bispectrum and trispectrum are provided in appendices D and E, respectively.

In practice, it is more convenient to compute the sums over sources as integrals over flux density, weighted by the number of sources per steradian per unit flux density, $dN/dS_{\nu} d\Omega$, in a certain flux-density range $[s_{\nu, \text{min}}, s_{\nu, \text{max}}]$. That is

$$\frac{1}{A_{\text{sky}}} \langle \sum_i S_{\nu, i}^n \rangle \rightarrow \int_{s_{\nu, \text{min}}}^{s_{\nu, \text{max}}} dS_{\nu} S_{\nu}^n \frac{dN}{dS_{\nu} d\Omega}. \quad (\text{C3})$$

The number density $dN/dS_{353 \text{ GHz}} d\Omega$ can be obtained by fitting a model to sub-mm data. Our parametric forms follow those employed by Mak et al. (2017) and are as follows: in the faint end, a double power-law,

$$\frac{dN}{dS_{353 \text{ GHz}} d\Omega} = A \left[\left(\frac{S_{353 \text{ GHz}}}{B} \right)^{n_1} + \left(\frac{S_{353 \text{ GHz}}}{B} \right)^{n_2} \right]^{-1}, \quad (\text{C4})$$

with parameters $A = 2.82 \times 10^8 \text{ Jy}^{-1} \text{ sr}^{-1}$, $B = 0.007 \text{ Jy}$, $n_1 = 6.5$ and $n_2 = 2.42$; and on the bright end, a single power-law,

$$\frac{dN}{dS_{353 \text{ GHz}} d\Omega} = k S_{353 \text{ GHz}}^{-2.5}, \quad (\text{C5})$$

with $k = 17.24 \text{ Jy}^{1.5} \text{ sr}^{-1}$.

The biases from these one-source terms are plotted in figure C1 as a function of the flux cut. Given that \mathcal{B}_{1s}^{EIEI} scales as the fourth power of the flux density, while \mathcal{B}_{1s}^{BEI} goes only as the third power, the former is bigger than the latter if sources brighter than around 1 Jy are left unmasked. However, for the typical flux cut values expected of the next generation of experiments, the bispectrum term dominates over the trispectrum term by several orders of magnitude.

From figure 8, it is clear that, on large scales, the full calculation of \mathcal{B}_{1s}^{BEI} is exceptionally well described by a squeezed




Figure C1. Delensing biases from the BEI bispectrum (solid) and $EIEI$ trispectrum (dashed) of the CIB, including only 1-source contributions, for point-source masking schemes typical of current and upcoming CMB experiments. The squeezed limit of the bispectra are shown as dotted lines. For comparison, we show also the fiducial power spectra of lensed B -modes before (grey) and after delensing with the CIB such that the power is reduced by 30% (black).

limit in which a large-scale B -mode leg couples to two, small-scale T and a E -mode legs. In this limit, equation C1 approaches

$$\mathcal{B}_l^{BEI} \approx -p^2 G_{353\text{ GHz}} G_{145\text{ GHz}}^2 \frac{1}{A_{\text{sky}}} \left\langle \sum_i S_{145\text{ GHz},i}^2 S_{353\text{ GHz},i} \right\rangle \int \frac{d^2 l'}{(2\pi)^2} \mathcal{W}_{l'}^E \mathcal{W}_{l'}^I, \quad (\text{C6})$$

which is independent of l and thus gives rise to the flat, white-noise-like angular power spectrum we see.

APPENDIX D: TOY-MODEL CIB BISPECTRUM: ONE-SOURCE TERM

The one-source contribution to the CIB bispectrum can also be calculated simply by working directly with the flux densities of the sources. Consider a galaxy at angular position θ_i , with flux-density $S_{\nu,i}$ at observed frequency ν . This contributes to the specific intensity of the CIB as

$$I_\nu(\theta) = S_{\nu,i} \delta_D^{(2)}(\theta - \theta_i). \quad (\text{D1})$$

In the Fourier domain, this is

$$I_\nu(l) = S_{\nu,i} e^{-il \cdot \theta_i}. \quad (\text{D2})$$

For N such sources this generalises to

$$I_\nu(l) = \sum_{i=1}^N S_{\nu,i} e^{-il \cdot \theta_i}. \quad (\text{D3})$$

If we allow the sources to be linearly polarised with a polarisation fraction p – which we shall assume is independent of flux and position – the observed polarisation can be described by the Stokes parameters

$$(Q \pm iU)_\nu(l) = \sum_i p S_{\nu,i} e^{\pm 2i\chi_i} e^{-il \cdot \theta_i}, \quad (\text{D4})$$

where χ_i is the angle between the polarisation direction of the source and the x -axis. (We use χ here rather than θ to avoid confusion with the angular position of the source.) The E - and B -modes of the observed polarisation follow from equation (A5):

$$E_\nu(l) = -p \sum_i S_{\nu,i} e^{-il \cdot \theta_i} \cos 2(\psi_l - \chi_i), \quad (\text{D5})$$

$$B_\nu(l) = p \sum_i S_{\nu,i} e^{-il \cdot \theta_i} \sin 2(\psi_l - \chi_i). \quad (\text{D6})$$

The bispectrum at frequency ν can therefore be constructed as

$$\langle B_\nu(l_1) E_\nu(l_2) I_\nu(l_3) \rangle = -p^2 \left\langle \sum_{ijk} S_{\nu,i} S_{\nu,j} S_{\nu,k} e^{-il_1 \cdot \theta_i} e^{-il_2 \cdot \theta_j} e^{-il_3 \cdot \theta_k} \sin 2(\psi_{l_1} - \chi_i) \cos 2(\psi_{l_2} - \chi_j) \right\rangle. \quad (\text{D7})$$

If we assume that the polarisation angles χ_i are independent of the positions and fluxes of sources, averaging over the angles will cause the bispectrum to vanish unless $i = j$. This implies that either $i = j = k$ or $i = j \neq k$. If we ignore source

clustering, so that the positions of the sources are independent of each other, the case $i = j \neq k$ gives vanishing contribution for $\mathbf{l}_3 \neq 0$. Keeping only the one-source term, $i = j = k$, we have

$$\begin{aligned} \langle B_\nu(\mathbf{l}_1)E_\nu(\mathbf{l}_2)I_\nu(\mathbf{l}_3) \rangle &= -p^2 \langle \sum_i S_{\nu,i}^3 e^{-i(\mathbf{l}_1+\mathbf{l}_2+\mathbf{l}_3)\cdot\boldsymbol{\theta}_i} \sin 2(\psi_{\mathbf{l}_1} - \chi_i) \cos 2(\psi_{\mathbf{l}_2} - \chi_i) \rangle \\ &= -\frac{p^2}{2} \sin 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2}) \langle \sum_i S_{\nu,i}^3 e^{-i(\mathbf{l}_1+\mathbf{l}_2+\mathbf{l}_3)\cdot\boldsymbol{\theta}_i} \rangle, \end{aligned} \quad (\text{D8})$$

where we have averaged over the polarisation angles in passing to the second line. Finally, averaging over the source locations across a sky area A_{sky} gives

$$\langle B_\nu(\mathbf{l}_1)E_\nu(\mathbf{l}_2)I_\nu(\mathbf{l}_3) \rangle = -\frac{p^2}{2} \sin 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2}) (2\pi)^2 \delta_{\text{D}}^{(2)}(\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l}_3) \frac{1}{A_{\text{sky}}} \langle \sum_i S_{\nu,i}^3 \rangle. \quad (\text{D9})$$

It can be shown that this spectrum satisfies all the required symmetries. Note also that it agrees with the one-source component of equation (A13) on replacing

$$\frac{1}{A_{\text{sky}}} \langle \sum_i S_{\nu,i}^3 \rangle \rightarrow \int dz \int dS_\nu S_\nu^3 \frac{dN}{dS_\nu dz d\Omega} = \int dz S_\nu^{(3)}(z). \quad (\text{D10})$$

APPENDIX E: TOY-MODEL CIB TRISPECTRUM: ONE-SOURCE TERM

Following the calculation of appendix D, we calculate here the $EIEI$ trispectrum of the CIB. Assuming once again that flux-density, position and polarisation are uncorrelated, the full four-point function is given by

$$\langle E_\nu(\mathbf{l}_1)E_\nu(\mathbf{l}_2)I_\nu(\mathbf{l}_3)I_\nu(\mathbf{l}_4) \rangle = \frac{1}{2} p^2 \cos 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2}) \langle \sum_{ikm} S_{\nu,i}^2 S_{\nu,k} S_{\nu,m} e^{-i(\mathbf{l}_1+\mathbf{l}_2)\cdot\boldsymbol{\theta}_i} e^{-i\mathbf{l}_3\cdot\boldsymbol{\theta}_k} e^{-i\mathbf{l}_4\cdot\boldsymbol{\theta}_m} \rangle. \quad (\text{E1})$$

Once again, we have the restriction that both polarisation legs must come from the same source, though this time there can in principle also be two-source and three-source contributions. The one-source contribution to the *connected* four-point function (trispectrum) is

$$\langle E_\nu(\mathbf{l}_1)E_\nu(\mathbf{l}_2)I_\nu(\mathbf{l}_3)I_\nu(\mathbf{l}_4) \rangle_c = \frac{1}{2} p^2 \cos 2(\psi_{\mathbf{l}_1} - \psi_{\mathbf{l}_2}) (2\pi)^2 \delta_{\text{D}}^{(2)}(\mathbf{l}_1 + \mathbf{l}_2 + \mathbf{l}_3 + \mathbf{l}_4) \frac{1}{A_{\text{sky}}} \langle \sum_i S_{\nu,i}^4 \rangle. \quad (\text{E2})$$

APPENDIX F: SIMULATING GAUSSIAN LSS TRACERS CORRELATED WITH CMB LENSING

Though in this paper we are only concerned with the CIB, we present here a more general framework for generating simulations of large-scale-structure tracers which are appropriately correlated with a given simulation of the CMB lensing potential and amongst themselves. The basic idea — which has been studied extensively in the literature around Cholesky decompositions — is to set up a linear system of equations from which it will be possible to solve for the coefficients involved granted we know the theoretical auto- and cross- spectra of the tracers.

Suppose we are to work with just three tracers: the CIB, LSST galaxies and an internal reconstruction, with spherical harmonic coefficients expressed as $I_{lm}, g_{lm}, \kappa_{lm}^{rec}$, respectively. Given a map of the true convergence with spherical harmonic coefficients κ_{lm} , the different tracers are correlated as

$$\kappa_{lm}^{rec} = \kappa_{lm} + n_{lm}, \quad (\text{F1})$$

$$g_{lm} = A_l^{gK} \kappa_{lm} + u_{lm}, \quad (\text{F2})$$

$$I_{lm} = A_l^{IK} \kappa_{lm} + A_l^{gI} u_{lm} + e_{lm}, \quad (\text{F3})$$

where n_{lm}, u_{lm}, e_{lm} are coefficients of the noise — and as such are presumed to each be uncorrelated with everything else — and κ_{lm} are the coefficients of the true convergence for the particular realisation that we wish to generate correlated tracers for. We can solve for the coefficients

$$A_l^{gK} = \frac{C_l^{gK}}{C_l^{KK}}, \quad (\text{F4})$$

$$A_l^{IK} = \frac{C_l^{IK}}{C_l^{KK}}, \quad (\text{F5})$$

$$A_l^{gI} = \frac{C_l^{gI} - A_l^{gK} A_l^{IK} C_l^{KK}}{C_l^{uu}} \quad (\text{F6})$$

and the noise spectra

$$C_l^{nn} = N_l^{\kappa\kappa}, \quad (\text{F7})$$

$$C_l^{uu} = C_l^{gg} - (A_l^{g\kappa})^2 C_l^{\kappa\kappa}, \quad (\text{F8})$$

$$C_l^{ee} = C_l^{II} - (A_l^{I\kappa})^2 C_l^{\kappa\kappa} - (A_l^{gI})^2 C_l^{uu}, \quad (\text{F9})$$

where $N_l^{\kappa\kappa}$ is the internal reconstruction noise level.

From the above, it is clear that the problem is soluble if we know all the auto- and cross- spectra, in which case all we have to do is draw the coefficients n_{lm}, u_{lm} and e_{lm} from probability distributions with angular power spectra $N_l^{\kappa\kappa}, C_l^{uu}$ and C_l^{ee} .

The solution above can be generalised to any number of tracers as follows. Consider tracer i with auto- spectrum C^{ii} and cross-spectrum C^{ij} with tracer j . Let it be described as a linear combination $\sum_p \sum_{lm} A_l^{ij} a^p(l, m)$ of harmonic coefficients a^p with angular power spectra $C_l^{a^p a^p}$ and scaled with weights given by the square, off-diagonal matrix A_l^{ij} . In the case above, $a^p = \{\kappa_{lm}, u_{lm}, e_{lm}\}$, $C_l^{a^p a^p} = \{C_l^{\kappa\kappa}, C_l^{uu}, C_l^{ee}\}$. The general formula for the weights is

$$A_l^{ij} = \frac{1}{C_l^{a^j a^j}} \left(C_l^{ij} - \sum_{p=0}^{j-1} A_l^{jp} A_l^{ip} C_l^{a^p a^p} \right), \quad (\text{F10})$$

with auxiliary spectra given by:

$$C_l^{a^j a^j} = C_l^{jj} - \sum_{p=0}^{j-1} (A_l^{jp})^2 C_l^{a^p a^p}. \quad (\text{F11})$$

This infrastructure is implemented and made publicly-available on [GitHub](#)¹⁰.

APPENDIX G: CONSTRUCTING A CURVED-SKY B-MODE TEMPLATE

We describe here the construction of a map-level template for lensing B-modes. We work to leading order in lensing and on the curved-sky formalism, a combination that was shown by [Challinor & Lewis \(2005\)](#) to be a very good approximation to the true B-modes on large angular scales. We arrive at a fast position-space implementation which is made publicly-available on [GitHub](#)¹¹. These lensed B-modes can be approximated in harmonic space by¹²

$$\begin{aligned} \tilde{B}_{lm} &= \sum_{(lm)_1} \sum_{(lm)_2} \phi_{(lm)_1} E_{(lm)_2} \frac{1}{2i} \left[{}_2I_{l_1 l_2}^{m m_1 m_2} \quad - {}_2I_{l_1 l_2}^{m m_1 m_2} \right] \\ &= \sum_{(lm)_1} \sum_{(lm)_2} \frac{-i}{2} \frac{1}{2} (-1)^m \phi_{(lm)_1} E_{(lm)_2} \sqrt{(2l+1)(2l_1+1)(2l_2+1)/4\pi} [l_1(l_1+1) + l_2(l_2+1) - l(l+1)] \\ &\quad \times \begin{pmatrix} l_1 & l_2 & l \\ m_1 & m_2 & -m \end{pmatrix} \left[\begin{pmatrix} l_1 & l_2 & l \\ 0 & -2 & 2 \end{pmatrix} - \begin{pmatrix} l_1 & l_2 & l \\ 0 & 2 & -2 \end{pmatrix} \right]. \end{aligned} \quad (\text{G1})$$

Inspired by the implementation of the quadratic estimators for lensing reconstruction in the publicly-available code [Quick-Lens](#)¹³, we write

$$\hat{B}_{lm} = \frac{(-1)^m}{2} \sum_{(lm)_1} \sum_{(lm)_2} \begin{pmatrix} l_1 & l_2 & l \\ m_1 & m_2 & -m \end{pmatrix} W_{l_1 l_2 l} \hat{\phi}_{(lm)_1} \hat{E}_{(lm)_2}, \quad (\text{G2})$$

where the weights

$$W_{l_1 l_2 l} = \frac{-i}{2} \sqrt{(2l+1)(2l_1+1)(2l_2+1)/4\pi} [l_1(l_1+1) + l_2(l_2+1) - l(l+1)] \left[\begin{pmatrix} l_1 & l_2 & l \\ 0 & -2 & 2 \end{pmatrix} - \begin{pmatrix} l_1 & l_2 & l \\ 0 & 2 & -2 \end{pmatrix} \right] \quad (\text{G3})$$

can be cast in separable form as

$$W_{l_1 l_2 l} = \sum_i W_{l_1 l_2 l}^i, \quad (\text{G4})$$

with

$$W_{l_1 l_2 l}^i = \sqrt{(2l+1)(2l_1+1)(2l_2+1)/4\pi} \begin{pmatrix} l_1 & l_2 & l \\ -s_1^i & -s_2^i & s \end{pmatrix} w_{l_1}^i w_{l_2}^i w_l^i. \quad (\text{G5})$$

¹⁰ <https://github.com/abaleato/MultitracerSims4Delensing>

¹¹ https://github.com/abaleato/curved_sky_B_template

¹² In this notation summation is implicit over matching pairs of indices.

¹³ <https://github.com/dhanson/quicklens>, though an amended and extended version can be found in the alternative repository <https://github.com/abaleato/Quicklens-with-fixes>

i	s_1^i	s_2^i	s^i	$w_{l_1}^i$	$w_{l_2}^i$	w_l^i
1	0	2	2	$l_1(l_1 + 1)$	$-1/2$	i
2	0	2	2	$-1/2$	$l_2(l_2 + 1)$	i
3	0	2	2	$1/2$	i	$l(l + 1)$
4	0	-2	-2	$l_1(l_1 + 1)$	$1/2$	i
5	0	-2	-2	$1/2$	$l_2(l_2 + 1)$	i
6	0	-2	-2	$-1/2$	i	$l(l + 1)$

Table G1. Weights for a fast separable implementation of the lensed B -mode template.

The value for the separable weights $w_{l_j}^i$ can be found in Table G1.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.