Baryons do trace dark matter 380,000 years after the big bang: Search for compensated isocurvature perturbations with WMAP 9-year data

Daniel Grin,[1,*] Duncan Hanson,[2,3] Gilbert P. Holder,[2] Olivier Doré,[3,4] and Marc Kamionkowski[5]

[1]*School of Natural Sciences, Institute for Advanced Study, Princeton, New Jersey 08540, USA*
[2]*Department of Physics, McGill University, Montreal, Quebec H3A 2T8, Canada*
[3]*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109, USA*
[4]*California Institute of Technology, Pasadena, California 91125, USA*
[5]*Department of Physics and Astronomy, Johns Hopkins University, Baltimore, Maryland 21218, USA*

Primordial isocurvature fluctuations between photons and either neutrinos or nonrelativistic species such as baryons or dark matter are known to be subdominant to adiabatic fluctuations. Perturbations in the relative densities of baryons and dark matter (known as compensated isocurvature perturbations or CIPs), however, are surprisingly poorly constrained. CIPs leave no imprint in the cosmic microwave background (CMB) on observable scales, at least at linear order in their amplitude and zeroth order in the amplitude of adiabatic perturbations. It is thus not yet empirically known if baryons trace dark matter at the surface of last scattering. If CIPs exist, they would spatially modulate the Silk damping scale and acoustic horizon, causing distinct fluctuations in the CMB temperature/polarization power spectra across the sky: this effect is first order in both the CIP and adiabatic mode amplitudes. Here, temperature data from the Wilkinson Microwave Anisotropy Probe (WMAP) are used to conduct the first CMB-based observational search for CIPs, using off-diagonal correlations and the CMB trispectrum. Reconstruction noise from weak lensing and point sources is shown to be negligible for this data set. No evidence for CIPs is observed, and a 95% confidence upper limit of 1.1×10^{-2} is imposed to the amplitude of a scale-invariant CIP power spectrum. This limit agrees with CIP sensitivity forecasts for WMAP and is competitive with smaller-scale constraints from measurements of the baryon fraction in galaxy clusters. It is shown that the root-mean-squared CIP amplitude on 5–100° scales is smaller than ~ 0.07–0.17 (depending on the scale) at the 95% confidence level. Temperature data from the *Planck* satellite will provide an even more sensitive probe for the existence of CIPs, as will the upcoming ACTPol and SPTPol experiments on smaller angular scales.

I. INTRODUCTION

Measurements of primordial density perturbations are consistent with adiabatic initial conditions, for which the ratios of neutrino, photon, baryon, and dark-matter number densities are initially spatially constant. The simplest inflationary models predict adiabatic fluctuations [1–6]. Isocurvature perturbations, on the other hand, are fluctuations in the relative number densities of different species. They are produced in topological-defect models for structure formation [7], multifield inflationary models and curvaton models [8–12], in which two different fields drive inflation and generate curvature perturbations, and in simple inflationary models if the dark matter is composed of axions [8,13,14].

CMB temperature anisotropies limit the contribution of baryon isocurvature perturbations (fluctuations in the baryon-to-photon ratio) [15,16] and CDM isocurvature perturbations (fluctuations in the dark-matter-to-photon ratio) [14,17–19] to the total perturbation amplitude [20–32]. The recent Planck CMB results limit the CDM

isocurvature fraction to be $\lesssim 0.039$ of the total perturbation amplitude.

It is therefore surprising that perturbations in the baryon density can be almost arbitrarily large, as long they are compensated by dark-matter perturbations such that the total nonrelativistic matter density remains unchanged [33,34]. These compensated isocurvature perturbations (CIPs) obey

$$\rho_c \delta_c^{\rm CI} + \rho_b \delta_b^{\rm CI} = 0, \qquad \delta_\gamma^{\rm CI} = 0, \tag{1}$$

where δ_c, δ_b, and δ_γ are fractional energy density perturbations in the cold dark matter, baryons, and photons, respectively, while ρ_c and ρ_b are the homogeneous dark matter and baryon densities. CIPs induce no curvature perturbation at early times, and they therefore leave the photon density, and thus large-angle CMB fluctuations, homogeneous at linear order.

Curvaton models for inflation may generate CIPs [35–38], with amplitudes approaching the regime detectable by the proposed EPIC mission [33], and other inflationary models [39] could generate even larger CIP amplitudes. Recent theoretical ideas [40–45] connecting the baryon asymmetry and dark-matter density could also have implications for

*dgrin@ias.edu

CIPs. In any case, our principal motivation in studying CIPs is curiosity: can we determine empirically, rather than simply assume, that the primordial baryon fraction is homogeneous and traces the dark matter?

CIPs induce baryon motion through baryon-pressure gradients, but these motions occur only at the baryon sound speed. The resulting anisotropies would be imprinted on the baryonic sound horizon, at $l \sim 10^6$ [33,37,46]. Existing measurements at low redshift constrain the CIP perturbation amplitude to be $\lesssim 10\%$ [33,34], while more sensitive proposed measurements of 21-cm absorption during the cosmic dark ages are a ways off in the future [33,46–48].

More recently, it has been shown that CIPs would modulate the CMB anisotropies produced by adiabatic perturbations, both by inducing anisotropies in the optical depth to reionization [34], and more dramatically, by changing the Silk damping length of the CMB in regions of sky containing a CIP [49,50]. This modulation would induce a specific pattern of higher order correlations in the temperature and polarization anisotropies, analogous to those induced by variations of other cosmological parameters [51] and by weak gravitational lensing [52]. This signature can be exploited to construct estimators for the CIP perturbation [50]. CIPs with amplitudes comparable to those generated in curvaton models for the primordial density fluctuation [35–38], lower-energy inflationary models [39], and perhaps other scenarios, are within the range of detectability for a cosmic-variance-limited CMB polarization experiment.

Here, we use WMAP 9-year temperature maps to search for CIPs. Our CIP estimator is based on the full non-Gaussian trispectrum of the observed CMB multipole moments, and can be used to perform either a model-independent reconstruction of the CIP power spectrum $C_L^{\Delta\Delta}$, or to measure the amplitude of a scale-invariant spectrum of CIPs.

We impose a 95% confidence upper limit of 1.1×10^{-2} to the amplitude of a scale-invariant spectrum of CIPs, as well as model-independent constraints of ~ 0.07–0.17 to the root-mean-squared (RMS) amplitude of the CIP power spectrum at angular scales in the range $1 \leq L < 20$, where L is the multipole index of the CIP. We show that secondary contractions of the trispectrum contribute negligibly to the estimator, at least for WMAP experimental parameters. We show that known sources of non-Gaussianity, such as gravitational lensing and unresolved point sources, do not provide a significant bias for our estimates of the CIP power spectrum. The same methodology which we have used here could be applied to the *Planck* data and has the potential to significantly improve on the constraints above.

We begin in Sec. II with a derivation of the temperature anisotropies induced by CIPs, which lead to a coupling of CIP modes of multipole index L with primordial CMB fluctuation modes of index l, and off-diagonal correlations between different multipole moments, absent under the

null hypothesis. In Sec. III, we use these correlations to obtain an optimal CIP estimator, based on that of Ref. [50], but generalized to be run on a partial-sky map with realistic noise properties. We account for bias and estimator normalization using analytic estimates and Monte Carlo simulations. In Sec. IV we present our results (limits to the amplitude of a scale-invariant spectrum of CIPs and to the model-independent CIP amplitude on different scales) and compare with the forecasted WMAP sensitivity, and we conclude in Sec. V. Throughout this work we use a fixed, flat ΛCDM cosmology consistent with the WMAP-9 [53] power spectrum, given by $\Omega_b = 0.045$, $\Omega_c = 0.222$, $h = 0.733$, adiabatic spectral index $n_s = 0.963$, reionization optical depth $\tau = 0.088$, and adiabatic scalar power spectrum normalization $A_s = 2.4 \times 10^{-9}$.

II. CMB TEMPERATURE CORRELATIONS IN THE PRESENCE OF A CIP

In Refs. [49,50], it was shown that a CIP would induce off-diagonal correlations between CMB anisotropy multipole moments. We rederive these results using a clearer method here, applying the line-of-sight (LOS) formalism of Ref. [54]. This approach is more readily generalized to CIPs with wavelength smaller than the thickness of the surface of last scattering and is useful in computing the CIP bispectrum. Before launching into the formalism, we review the physical origin of CIP-induced CMB correlations.

As noted in Refs. [33,37,46], CIPs have no initial potential perturbations (like other isocurvature modes), but they also have no initial radiation pressure gradients. In linear theory, flows then begin at the baryon sound speed, but are only effective in transferring fluctuations to the photons on very small scales ($l \sim 10^6$). Even when baryons do begin to evacuate initial density fluctuations (which would yield a net potential perturbation as the CIP evolved), these flows will be diffusion damped, slowing down the growth of these small-scale potential perturbations. This intuition is confirmed by running the Boltzmann code CAMB [55] with a CIP initial condition, and noting that the induced CMB temperature anisotropy is negligible.

The fluctuating baryon fraction in the presence of a CIP would lead to an inhomogeneous redshift of reionization (when the first sources turn on), leading to a fluctuating optical depth and a distinct patchy reionization signal in the CMB [34]. More dramatically, however, if CIPs are in fact primordial, they will lead to a spatial modulation of coefficients in the early-time (tight-coupling era) equations of motion for fluid perturbations, altering observed CMB anisotropies in a detectable way [49,50].

A. Physical origin of effect

Consider a compensated isocurvature perturbation (CIP) along the line of sight $\Delta(\hat{n})$, which gives local changes in the baryon and CDM energy densities given by

$$\Omega_b \rightarrow \Omega_b[1 + \Delta(\hat{n})], \qquad \Omega_c \rightarrow \Omega_c - \Omega_b\Delta(\hat{n}). \quad (2)$$

A positive (negative) value of Δ corresponds to a more (less) baryon-loaded plasma, which decreases (increases) the sound speed and thus decreases (increases) the physical acoustic horizon. The multipole index l_s of the first CMB acoustic peak thus increases (decreases) as Δ increases (decreases), as shown in Fig. 1 (from Ref. [50]), generated using expressions in Ref. [56]. The other acoustic peak locations behave similarly.

CMB temperature anisotropies are suppressed on angular scales $l > l_d \sim 1000$ due to diffusion damping. Using the expressions in Ref. [57] and the CAMB [55] code, we evaluate $l_d(\Delta)$ and show the results in the top right panel of

Fig. 1. We see that, as photons diffuse over smaller distances, as a result of higher local baryon density in the presence of a CIP with positive Δ, the transition to exponential damping of CMB anisotropies occurs at higher l.

In the bottom panel of Fig. 1 (from Ref. [50]), we show the visibility functions $g(z) = e^{-\tau}d\tau/dz$ for three different values of Δ; τ is the optical depth due to Thomson scattering and z is the redshift. The peak of the visibility function is the redshift $z_{\rm SLS}$, at which most CMB photons last scatter. In the presence of a positive (negative) Δ CIP, decoupling occurs later (earlier) due to higher (lower) baryon density.

The effects described above all result in modifications to the CMB power spectrum. One can therefore imagine
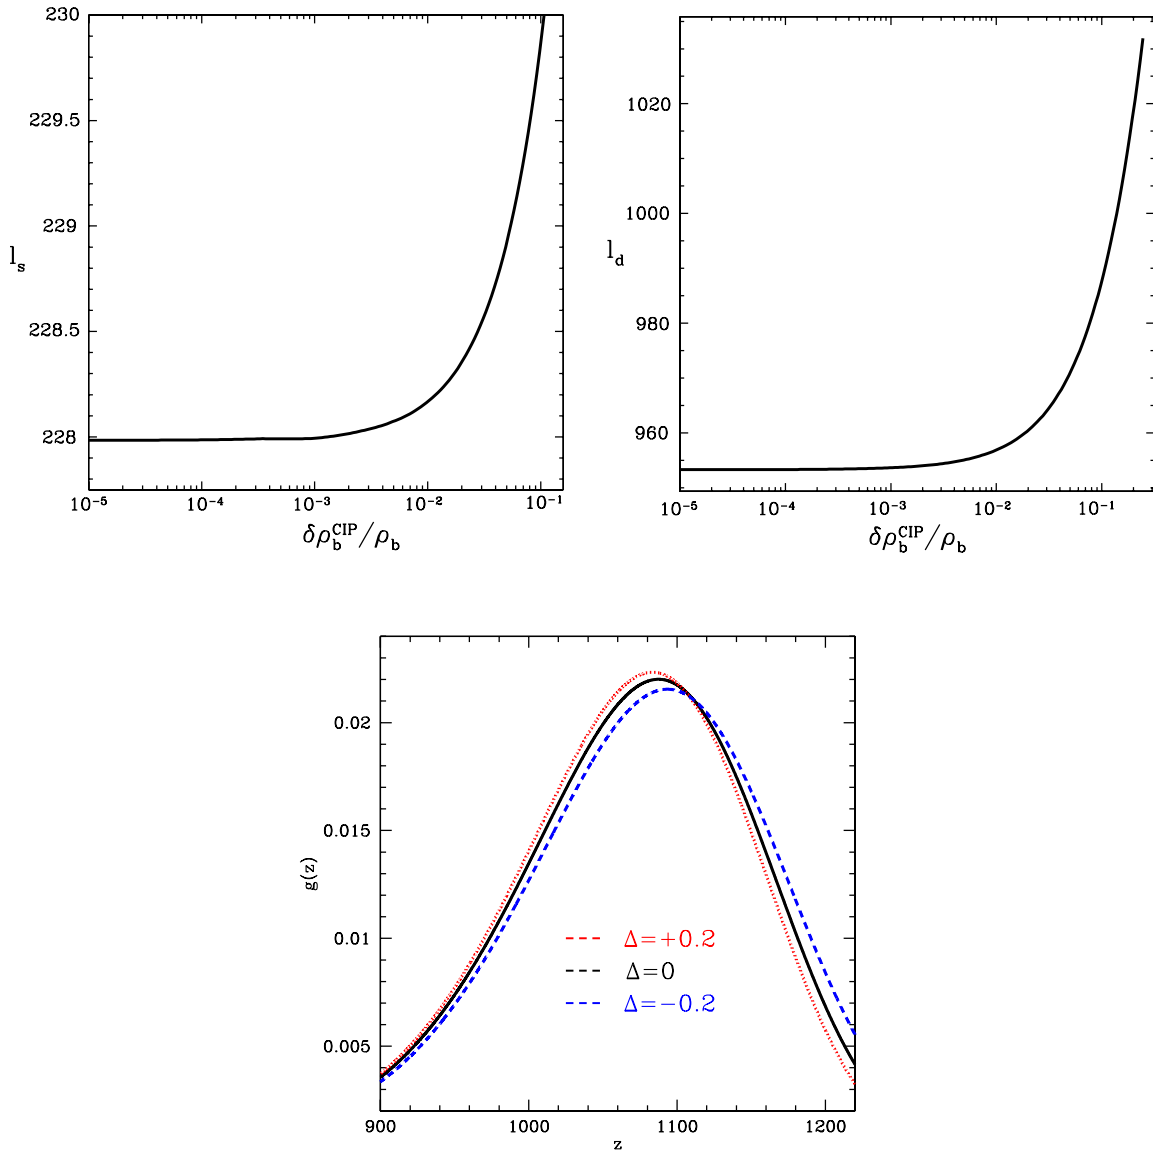

FIG. 1 (color online).   Physical and Thomson scattering visibility function $g(z)$ in the presence of a global CIP Δ. Top left panel shows angular sound horizon l_s as a function of a spatially uniform CIP Δ. Top right panel shows diffusion damping scale l_d as a function of Δ. Bottom panel shows $g(z)$ evaluated for three different values of Δ. Later we will use these physical effects to probe the CMB for spatially varying $\Delta(\hat{n})$.

constructing an estimator for Δ by forming localized power spectrum estimates and fitting for Δ relative to a fiducial model based on the full-sky power spectrum. We will apply this heuristic principle more precisely to build estimators for the CIP realization Δ using off-diagonal CMB correlations in Sec. III, but first we must quantify the change induced by a CIP Δ in the presence of some underlying adiabatic perturbation.

B. Standard line-of-sight solution for CMB temperature anisotropy

The calculation of CMB anisotropies is greatly simplified using the line-of-sight (LOS) approach, introduced in Ref. [54]. At spatial location $\vec{y}$, conformal time η, and for photon momentum direction vector $\hat{p}$ (where $|\hat{p}| = 1$), the photon temperature perturbation $T(\vec{y}, \hat{p}, \eta)$ obeys the equation (derived from the Boltzmann equation)

$$\dot{T}(\vec{y}, \hat{p}, \eta) + \hat{p} \cdot \nabla T(\vec{y}, \hat{p}, \eta) = \mathcal{D}(\hat{p}, \eta)[\vec{u}(\vec{y}, \eta)], \quad (3)$$

where $\vec{u}(\vec{y}, \eta)$ is a vector whose entries are the fluid density/velocity perturbations (as well as higher-order moments of the distribution function, for neutrinos) and metric fluctuations characterizing the system, and $\mathcal{D}$ is a linear differential operator which maps $\vec{u}(\vec{y}, \eta)$ to a source term for temperature perturbations. Defining the LHS of this equation as the operator $\mathcal{B}$ (for Boltzmann), and taking a Fourier transform, we obtain

$$\mathcal{B}_{\vec{k}}[T_{\vec{k}}(\hat{p}, \eta)] = \mathcal{D}_{\vec{k}}(\hat{p}, \eta)[\vec{u}_{\vec{k}}(\eta)], \quad (4)$$

where $\mathcal{D}_{\vec{k}}$ is a matrix operator in Fourier space (as opposed to a differential operator). For the usual adiabatic mode, $\vec{u}_{\vec{k}}(\eta) = \vec{f}_{\vec{k}}(\eta)\Phi_{\vec{k}}$, where $\vec{f}_{\vec{k}}(\eta)$ is a time evolution operator mapping the initial potential perturbation $\Phi_{\vec{k}}$ to the solution for the fluid variables at subsequent times.

The components of the matrix operator $\mathcal{D}_{\vec{k}}$ (which maps the fluid and metric variables to the observed temperature perturbation) and vector-valued function $\vec{f}_{\vec{k}}$ are laid out in detail in Refs. [54,58] and others, so we use the operator notation to keep things simple and general. Equation (4) may be formally integrated to obtain

$$T_{\vec{k}}(\hat{p}, \eta_0) = \int_0^{\eta_0} d\eta\, e^{ik\mu(\eta_0 - \eta)} \tilde{S}[\hat{p}, \vec{k}, \eta] \Phi_{\vec{k}}, \quad (5)$$

where the source function is $\tilde{S}[\hat{p}, \vec{k}, \eta] = \mathcal{D}_{\vec{k}}(\hat{p}, \eta)\vec{f}_{\vec{k}}(\eta)$, η_0 is the conformal time today and $\mu = \hat{p} \cdot \vec{k}/|\vec{k}|$. It turns out that the source function depends on $\vec{k}$ and $\hat{p}$ only through μ (as a polynomial in μ) and $k = |\vec{k}|$, and so Eq. (5) may be integrated by parts to obtain

$$T_{\vec{k}}(\hat{p}, \eta_0) = \int_0^{\eta_0} d\eta\, e^{ik\mu(\eta_0 - \eta)} S[k, \eta] \Phi_{\vec{k}}, \quad (6)$$

in terms of a different source function $S[k, \eta]$, whose terms are specified in Refs. [54,56]. Going back to real space,

following suit with an inverse spherical harmonic transform to derive multipole coefficients, and expanding the exponential using a Fourier-Bessel series, we obtain

$$T_{lm} = \frac{4\pi i^l}{(2\pi)^3} \int d^3k\, T_l(k) \Phi_{\vec{k}} Y_{lm}^*(\hat{k}), \quad (7)$$

$$T_l(k) \equiv \int_0^{\eta_0} S[k, \eta] j_l[k(\eta_0 - \eta)], \quad (8)$$

where $j_l(x)$ is a spherical Bessel function of index l.

The familiar CMB angular power spectrum may then be obtained:

$$\langle T_{lm} T_{l'm'}^* \rangle = \delta_{ll'} \delta_{mm'} C_l, \quad (9)$$

$$C_l = \frac{2}{\pi} \int k^2 dk\, T_l^2(k) P_\Phi(k), \quad (10)$$

where the three-dimensional potential fluctuation power spectrum $P_\Phi(k)$ is defined by

$$\langle \Phi_{\vec{k}} \Phi_{\vec{k}'}^* \rangle = (2\pi)^3 \delta^3(\vec{k} - \vec{k}') P_\Phi(k). \quad (11)$$

We now generalize the LOS solution to compute the off-diagonal temperature correlations induced by a CIP.

C. Line-of-sight solution in the presence of a CIP

The operators $\mathcal{D}_{\vec{k}}(\eta)$ and $\vec{f}_{\vec{k}}(\eta)$ depend on the cosmological parameters and thus also on the amplitude $\Delta(\hat{n})$ of a CIP. In the presence of a CIP, the real-space evolution equation will read

$$\mathcal{B}[\delta T(\vec{y}, \hat{p}, \eta)] = \frac{d\mathcal{D}(\hat{p}, \eta)[\vec{u}(\vec{y}, \eta)]}{d\Delta} \Delta(\vec{y}), \quad (12)$$

where $\delta T(\vec{y}, \hat{p}, \eta)$ is the modulation of the usual temperature anisotropy by the CIP.

As discussed in Sec. I, on the scales of interest, CIPs do not drive matter flows at linear order in perturbation theory, and are thus frozen in time at that level of approximation. In principle, CIP amplitudes still have a dependence on the conformal time along a photon trajectory (as a proxy for LOS distance), if the mode wavelength is shorter than the integration interval. For modes of large angular scale $l \ll l_{silk}$, however, this radial dependence may be neglected, and $\Delta(\vec{y}) = \Delta(\hat{y})$. Fourier transforming Eq. (12), performing a spherical harmonic expansion of $\Delta(\hat{y}) = \sum_{LM} \Delta_{LM} Y_{LM}(\hat{n})$, we see that the evolution equation for $\delta T_{\vec{k}}(\hat{p}, \eta)$ is of the same form as Eq. (4), but with a source term that is a linear superposition of source terms like those in Eq. (4). Since this is a linear system, solutions may be superimposed. Fourier-Bessel expanding and performing an inverse spherical harmonic transform of the solution, we obtain the perturbation to the LOS solution induced by a CIP,

$$\delta T_{lm} = \sum_{LMl_1m_1} \frac{\Delta_{LM}4\pi i^{l_1}\xi^{LM}_{lml_1m_1}K^L_{ll_1}}{(2\pi)^3}$$

$$\times \int d^3q Y^*_{l_1m_1}(\hat{q})\Phi_{\vec{q}}\frac{dT_{l_2}(q)}{d\Delta}, \qquad (13)$$

$$\xi^{LM}_{lml_1m_1} \equiv (K^L_{ll_1})^{-1}\int d\hat{n}Y^*_{lm}(\hat{n})Y_{LM}(\hat{n})Y_{l_1m_1}(\hat{n})$$

$$= (-1)^m\sqrt{\frac{(2L+1)(2l+1)(2l_1+1)}{4\pi}}$$

$$\times \begin{pmatrix} l & L & l_1 \\ -m & M & m_1 \end{pmatrix},$$

$$K^L_{ll_1} \equiv \begin{pmatrix} l & L & l_1 \\ 0 & 0 & 0 \end{pmatrix}, \qquad (14)$$

expressed in terms of the familiar Wigner-3J symbols [59].

D. CIP statistics

As discussed in the previous section, a compensated isocurvature perturbation induces a small fluctuation in the CMB temperature which is proportional to the primordial potential $\Phi_{\vec{k}}$. If we consider a fixed realization of isocurvature perturbations Δ_{LM}, the effect of CIPs is to introduce "statistical anisotropy" into the CMB, which manifests itself through off-diagonal elements in the covariance matrix of the CMB fluctuations, whose elements are then given by

$$\langle T^*_{l'm'}T_{lm}\rangle = \delta_{ll'}\delta_{mm'}C^{TT}_l + \sum_{LM}\begin{pmatrix} l & l' & L \\ m & m' & M \end{pmatrix}W^\Delta_{ll'L}\Delta_{LM}. \qquad (15)$$

Here C^{TT}_l is the usual CIP-free temperature power spectrum. The quantity

$$W^\Delta_{ll'L} \equiv \sqrt{\frac{(2l+1)(2l'+1)(2L+1)}{4\pi}}\times(C^{T,dT}_{l'}+C^{T,dT}_l)$$

$$\times \begin{pmatrix} l & l' & L \\ 0 & 0 & 0 \end{pmatrix} \qquad (16)$$

is a weight function associated with the CIPs and

$$C^{T,dT}_l \equiv \frac{2}{\pi}\int k^2dkP_\Phi(k)T_l(k)\frac{dT_l(k)}{d\Delta}. \qquad (17)$$

The derivative power spectra of Eq. (17) are evaluated using the CAMB code, using numerical methods described in Ref. [55], including a spatial modulation in the optical depth to reionization, τ, as well as the much larger effect from physics near recombination. Analogously to the violations of primordial isotropy considered in Refs. [60,61], CIPs introduce a violation of statistical isotropy in the CMB. As a result, it is no surprise that the off-diagonal correlations described by Eqs. (15) and (16) are similar

to those obtained due a breaking of primordial isotropy using the Bipolar Spherical Harmonic (BiPoSH) formalism [60,61].

We see above that a fixed CIP realization breaks the diagonality of the covariance matrix of the T_{lm} in a very specific way, yielding a unique statistical signature which can be used to reconstruct the Δ_{LM} realization. A similar derivation can reproduce the off-diagonal polarization correlations obtained in Ref. [50].

If the Universe as a whole has no preferred orientation, then the CIP perturbations Δ_{LM} are themselves random variables, with some statistically isotropic distribution. If we assume that Δ_{LM} are independent of the primordial fluctuations and Gaussian, then they are completely characterized by their power spectrum $C^{\Delta\Delta}_L$. In this more realistic picture, the first distinctive statistical signature of CIPs appears in the CMB "trispectrum," or connected four-point function. The connected part of the four-point function is zero for purely Gaussian fluctuations, but in the presence of CIPs it becomes non-zero. Following Ref. [62], the connected four-point function must take the form

$$\langle T_{l_1m_1}T_{l_2m_2}T_{l_3m_3}T_{l_4m_4}\rangle_C = \sum_{LM}(-1)^M T^{l_1l_2}_{l_3l_4}(L)\mathcal{G}^{l_1m_1l_2m_2L}_{l_3m_3l_4m_4M}, \qquad (18)$$

where $T^{l_1l_2}_{l_3l_4}(L)$ is known as the trispectrum, and we have used the Glebsch-Morgan coefficient

$$\mathcal{G}^{l_1m_1l_2m_2L}_{l_3m_3l_4m_4M} = \begin{pmatrix} l_1 & l_2 & L \\ m_1 & m_2 & -M \end{pmatrix}\begin{pmatrix} l_3 & l_4 & L \\ m_3 & m_4 & M \end{pmatrix}. \qquad (19)$$

Symmetry of the four multipoles further requires that the trispectrum may be encoded as

$$T^{l_1l_2}_{l_3l_4}(L) = P^{l_1l_2}_{l_3l_4}(L) + (2L+1)$$

$$\times \sum_{L'}\left[(-1)^{l_2+l_3}\begin{Bmatrix} l_1 & l_2 & L \\ l_4 & l_3 & L' \end{Bmatrix}P^{l_1l_3}_{l_2l_4}(L')\right.$$

$$\left.+(-1)^{L+L'}\begin{Bmatrix} l_1 & l_2 & L \\ l_4 & l_3 & L' \end{Bmatrix}P^{l_1l_4}_{l_3l_2}(L')\right]. \qquad (20)$$

The first term $P^{l_1l_2}_{l_3l_4}(L)$ is referred to as the primary contraction of the trispectrum, while the final two terms are known as secondary contractions. The primary contraction introduced by CIPs is given simply by

$$P^{l_1l_2}_{l_3l_4}(L) = C^{\Delta\Delta}_L W^\Delta_{l_1l_2L}W^\Delta_{l_3l_4L}. \qquad (21)$$

Observationally, we will see that the picture of statistical anisotropy introduced initially above is still a useful *ansatz*; as the Universe shows us only a single realization of both $\Phi_{\vec{k}}$ and Δ_{LM}, a fixed Δ_{LM} model is experimentally indistinguishable from one in which the Δ_{LM} are Gaussian random variables [63].

III. CIP ESTIMATOR

We can construct an estimator for the CIP power spectrum utilizing the connected trispectrum which it induces. If we are interested in a single mode of this power spectrum $C_L^{\Delta\Delta}$, then following Ref. [64], for example, an Edgeworth expansion of the CMB likelihood leads to the following optimal (minimum-variance) trispectrum estimator

$$
\begin{aligned}
\hat{C}_L^{\Delta\Delta} = \frac{\mathcal{N}_L}{2L+1} & \sum_M \sum_{l_i m_i} (-1)^M \mathcal{G}_{l_3 m_3 l_4 m_4 M}^{l_1 m_1 l_2 m_2 L} W_{l_1 l_2 L}^{\Delta} \\
& \times W_{l_3 l_4 L}^{\Delta} (\bar{T}_{l_1 m_1} \bar{T}_{l_2 m_2} \bar{T}_{l_3 m_3} \bar{T}_{l_4 m_4} \\
& - [2\bar{T}_{l_1 m_1} \bar{T}_{l_2 m_2} - \bar{C}_{l_1 m_1, l_2 m_2}^{TT}] \bar{C}_{l_3 m_3, l_4 m_4}^{TT} \\
& - [2\bar{T}_{l_1 m_1} \bar{T}_{l_3 m_3} - \bar{C}_{l_1 m_1, l_3 m_3}^{TT}] \bar{C}_{l_2 m_2, l_4 m_4}^{TT} \\
& - [2\bar{T}_{l_1 m_1} \bar{T}_{l_4 m_4} - \bar{C}_{l_1 m_1, l_4 m_4}^{TT}] \bar{C}_{l_2 m_2, l_3 m_3}^{TT}),
\end{aligned} \tag{22}
$$

where $\mathcal{N}_L$ is a normalization, $\bar{T}_{lm}$ are a set of so called "inverse-variance filtered" multipoles determined from the data map (we will discuss how these are obtained in more detail shortly in Sec. III A), and $\bar{C}_{lm,l'm'}$ is their covariance matrix.

For any modern CMB experiment, which measures millions of modes, the covariance matrices $\bar{C}_{lm,l'm'}$ have trillions of elements and are impossible to work with directly. It is, however, computationally tractable to evaluate the expression above (which distills these covariance matrices down to a single number) using Monte Carlo simulations. We will rewrite Eq. (22) in a form that makes the details of this evaluation clearer, and also makes connection to the discussion of statistical anisotropy in the previous section.

We begin by introducing the "quadratic estimator" $\bar{\Delta}_{LM}$, which is a function of two inverse-variance filtered temperature multipoles as

$$
\bar{\Delta}_{LM}[\bar{T}_{lm}^{(1)}, \bar{T}_{lm}^{(2)}] = \sum_{lm,l'm'} \begin{pmatrix} l & l' & L \\ m & m' & M \end{pmatrix} W_{ll'L}^{\Delta} \bar{T}_{lm}^{(1)} \bar{T}_{l'm'}^{(2)}. \tag{23}
$$

If we were considering the CIP realization to be fixed, this quadratic estimator is precisely the quantity which is required for optimal estimation of Δ_{LM}, following the formalism of quadratic maximum likelihood (QML) estimators [49,50,65–72]. This estimator is related to the actual CIP multipole moment Δ_{LM} by an overall normalization. We will often find it useful to work with such unnormalized quantities, which we will denote with overbars. The reason is that for these maximum-likelihood estimators, the normalization is formed from the inverse of the estimator Fisher matrix. This quantity is also the variance of the estimator, and so these un-normalized estimates are effectively inverse-noise weighted. For this reason, we have denoted the $\bar{\Delta}_{LM}$ with an overbar, in analogy to the inverse variance filtered temperature multipoles. This makes the un-normalized estimators useful for estimating

other parameters, such as the overall amplitude of a scale-invariant spectrum of CIPs, as we shall see below. We note that Δ_{LM} may be evaluated rapidly (with computational cost $\mathcal{O}(l_{\max}^3)$ using fast spherical harmonic transforms) in position space as the product of two filtered maps,

$$
\bar{\Delta}_{LM}[\bar{T}_{lm}^{(1)}, \bar{T}_{lm}^{(2)}] = \int d\hat{n} Y_{LM}^*(\hat{n}) \bar{T}^{(1)}(\hat{n}) S^{(2)}(\hat{n}) + [(1) \leftrightarrow (2)], \tag{24}
$$

where the filtered maps themselves are given by

$$
\bar{T}^{(a)}(\hat{n}) = \sum_{lm} Y_{lm}(\hat{n}) \bar{T}_{lm}^{(a)}, \tag{25}
$$

$$
S^{(a)}(\hat{n}) = \sum_{lm} Y_{lm}(\hat{n}) C_l^{T,dT} \bar{T}_{lm}^{(a)}, \tag{26}
$$

where $a = 1$ or $a = 2$ as appropriate.

Using this notation, we rewrite the $C_L^{\Delta\Delta}$ estimator as $\hat{C}_L^{\Delta\Delta} = \bar{C}_L^{\Delta\Delta} \mathcal{N}_L$, where

$$
\bar{C}_L^{\Delta\Delta} = C_L^{\bar{\Delta}\bar{\Delta}} - D_L^{\bar{\Delta}\bar{\Delta}}. \tag{27}
$$

The naive *un-normalized* power spectrum estimate $C_L^{\bar{\Delta}\bar{\Delta}}$ is given by

$$
\begin{aligned}
C_L^{\bar{\Delta}\bar{\Delta}} = \sum_{l_i m_i} \sum_M \frac{1}{(2L+1)} \langle (\bar{\Delta}_{LM}[\bar{T}_{lm}, \bar{T}_{lm}] \\
- \bar{\Delta}_{LM}[\bar{T}_{lm}^{(g)}, \bar{T}_{lm}^{(g)}])^* (\bar{\Delta}_{LM}[\bar{T}_{lm}, \bar{T}_{lm}] \\
- \bar{\Delta}_{LM}[\bar{T}_{lm}^{(f)}, \bar{T}_{lm}^{(f)}]) \rangle_{g,f},
\end{aligned} \tag{28}
$$

and the "disconnected noise bias" estimate is given by

$$
\begin{aligned}
D_L^{\bar{\Delta}\bar{\Delta}} = \sum_{l_i m_i} \sum_M \frac{1}{(2L+1)} \langle -4\bar{\Delta}_{LM}[\bar{T}_{lm}, \bar{T}_{lm}^{(g)}] \bar{\Delta}_{LM}[\bar{T}_{lm}, \bar{T}_{lm}^{(g)}] \\
+ 2\bar{\Delta}_{LM}[\bar{T}_{lm}^{(g)}, \bar{T}_{lm}^{(f)}] \bar{\Delta}_{LM}[\bar{T}_{lm}^{(g)}, \bar{T}_{lm}^{(f)}] \rangle_{g,f}.
\end{aligned} \tag{29}
$$

For both spectra, the ensemble average is taken over two sets of statistically independent Monte Carlo simulations of $\bar{T}$, labeled g and f. This expression could equivalently be written with just a single set of simulations, however evaluating it as is done here reduces the susceptibility to numerical noise in the evaluation procedure.

The disconnected noise bias arises because even a purely Gaussian (CIP and lensing-free) CMB has a nonzero four-point function (but a vanishing trispectrum). This four-point function has some nonzero overlap with the first term on the right-hand side of Eq. (27), which is equivalent to a more naive CIP power spectrum estimator, $\tilde{C}_L^{\Delta\Delta} = [\mathcal{N}_L/(2L+1)] \sum_M \bar{\Delta}_{LM}^* \bar{\Delta}_{LM}$. If this estimator were used, a spurious CIP signal would arise, even from a CIP-free map. We avoid such contamination through proper use of the Edgeworth expansion of the CMB likelihood function, which yields the noise-bias-subtracted estimators given by Eq. (22) and equivalently Eq. (27).

Rather than a single mode $C_L^{\Delta\Delta}$, we may be interested in the amplitude A of a fiducial power spectrum $C_L^{\Delta\Delta} = AC_L^{\Delta\Delta,(\mathrm{fid})}$, for which the corresponding optimal estimator is

$$\hat{A} = \mathcal{N} \sum_{L=L_{\min}}^{L_{\max}} (2L+1)\bar{C}_L^{\Delta\Delta} C_L^{\Delta\Delta,(\mathrm{fid})}, \qquad (30)$$

where $\mathcal{N}$ is an overall normalization, not to be confused with $\mathcal{N}_L$. Note the use of the un-normalized $\bar{C}_L^{\Delta\Delta}$ here rather than $\hat{C}_L^{\Delta\Delta}$.

A. Filtering

In the case of full-sky coverage with homogeneous noise levels, the inverse-variance filter is given simply by

$$\bar{T}_{lm} = \frac{F_l}{B_l} \sum_{p=0}^{n_{\mathrm{pix}}} \frac{4\pi}{n_{\mathrm{pix}}} Y_{lm}(\hat{n}_p) T_p^{\mathrm{obs}}, \qquad (31)$$

where B_l is the beam- and pixel-transfer function, T_p^{obs} is the observed map indexed by pixel p, and the filter function F_l is given by

$$F_l = \frac{1}{C_l^{\mathrm{TT}} + C_l^{\mathrm{TT,noise}}}. \qquad (32)$$

For the more realistic case of a beam-convolved sky map with inhomogeneous noise, the construction of $\bar{T}_{lm}$ is more involved. To obtain $\bar{T}_{lm}$ from a set of WMAP sky maps, we use an inverse-variance filter which properly accounts for sky-cuts and the inhomogeneity of the map and its noise levels by solving the equation

$$\bar{T}_{lm} = (C_l^{\mathrm{TT}})^{-1} \sum_{l'm'} \sum_{p,\nu} \mathcal{C}_{lm,l'm'}^{-1} \mathcal{Y}_{l'm'}^{*p,\nu} N_{p,\nu}^{-1} T_{p,\nu}^{\mathrm{obs}}, \qquad (33)$$

where the matrix $\mathcal{C}_{lm,l'm'}$ is given by

$$\mathcal{C}_{lm,l'm'} \equiv \sum_{p,\nu} [(C_l^{\mathrm{TT}})^{-1} \delta_{ll'} \delta_{mm'} + \mathcal{Y}_{lm}^{*p\nu} N_{p,\nu}^{-1} \mathcal{Y}_{l'm'}^{p,\nu}]. \qquad (34)$$

Here the p denotes a map pixel, ν denotes a particular channel map (usually a given frequency band), $T_{p,\nu}^{\mathrm{obs}}$ is the observed (beam convolved) sky map in pixel p at frequency ν, the pointing matrix $\mathcal{Y}_{lm}^{p,\nu} \equiv B_l^{\nu} Y_{lm}(\hat{n}_p)$ gives the value of the spherical harmonic at the center of pixel p, convolved with the appropriate beam + pixel transfer function B_l. N is the noise covariance matrix. We use a diagonal noise covariance in pixel space with

$$N_{p,\nu}^{-1} = \frac{N_{\mathrm{hits},p,\nu}}{\sigma_\nu^2} M_p, \qquad (35)$$

where σ_ν is a map-dependent noise level and $N_{\mathrm{hits},p}$ is the number of observations of pixel p for map ν. M_p is a map which is zero for masked pixels and unity elsewhere. Effectively, it sets the noise level to infinity for masked pixels, ensuring that they are ignored in the rest of the

analysis. We have simulated the noise as white and uncorrelated between pixels. There are noise correlations in the WMAP maps [73] on large angular scales ($l < 48$), which can enhance the noise power spectrum. The CMB power spectrum is more than 2 orders of magnitude higher than the noise power spectrum even where the noise power spectrum is enhanced by correlations, and so we neglect them in our simulations.

We evaluate the matrix inverse of Eq. (34) using conjugate descent with the fast multigrid preconditioner of Ref. [74]. Generally, we will find it useful to work with individual WMAP frequency maps (using only a single entry for ν in the equations above), though for our final results we will combine all of the useable bands.

For analytical purposes, it is useful to have a diagonal approximation to the full-blown inverse variance filter. For this we use the F_l function at the top of this section, estimating the average noise power spectrum of the channel-combined, beam-deconvolved map as

$$C_l^{\mathrm{TT,noise}} \approx \frac{4\pi}{n_{\mathrm{pix}} f_{\mathrm{sky}}} \left[\sum_\nu (B_l^\nu)^2 \frac{1}{\sum_{p=0}^{n_{\mathrm{pix}}} N_{p,\nu}} \right]^{-1}, \qquad (36)$$

where $f_{\mathrm{sky}} = \sum_{p=0}^{n_{\mathrm{pix}}} M_p / n_{\mathrm{pix}}$ is the unmasked sky fraction.

B. Normalization and bias

We now analytically derive an approximate normalization $\mathcal{N}$ of our estimator. We will eventually correct for deviations from this normalization using Monte Carlo simulations. The analytical treatment, however, is also useful as a tool to explicitly compute the estimator bias induced by other physical sources of non-Gaussianity at the trispectrum level.

Consider full-sky coverage, with homogeneous noise, in which case the filtered CMB covariance matrix $\bar{C}_{lm,l'm'}$ is diagonal. If we ensemble average over CMB and CIP realizations in Eq. (22), we find that (by construction) $\hat{C}_L^{\Delta\Delta}$ is directly proportional to the connected four-point function of T_{lm},

$$\langle \hat{C}_L^{\Delta\Delta} \rangle_{\mathrm{CIP,CMB}} = \frac{\mathcal{N}_L}{2L+1} \sum_M \sum_{l_i m_i} (-1)^M \mathcal{G}_{l_3 m_3 l_4 m_4 M}^{l_1 m_1 l_2 m_2 L}$$
$$\times F_{l_1} F_{l_2} F_{l_3} F_{l_4} W_{l_1 l_2 L}^{\Delta} W_{l_3 l_4 L}^{\Delta}$$
$$\times \langle \bar{T}_{l_1 m_1} \bar{T}_{l_2 m_2} \bar{T}_{l_3 m_3} \bar{T}_{l_4 m_4} \rangle_C. \qquad (37)$$

Here F_l are the filter functions of the previous section.

For generality, so that we can estimate possible biases to our estimator from known non-Gaussian sources such as unresolved point sources and gravitational lensing by large-scale structure, let us consider the response of this estimator to a trispectrum with the primary form

$$^x P_{l_3 l_4}^{l_1 l_2}(L) = C_L^{xx} W_{l_1 l_2 L}^x W_{l_3 l_4 L}^x. \qquad (38)$$

This covers both the CIP trispectrum of Eq. (21), as well as the trispectra due to gravitational lensing of the CMB by

large-scale structure, and unresolved point sources which pollute the map. The trispectrum for CMB lensing is given by [62]

$$^\phi P^{l_1 l_2}_{l_3 l_4}(L) = C^{\phi\phi}_L W^\phi_{l_1 l_2 L} W^\phi_{l_3 l_4 L}, \tag{39}$$

where

$$W^\phi_{l_1 l_2 L} = \frac{C_{l_2}}{2} \sqrt{\frac{(2l_1+1)(2l_2+1)(2L+1)}{4\pi}} \begin{pmatrix} l_1 & l_2 & L \\ 0 & 0 & 0 \end{pmatrix}$$
$$\times [L(L+1) + l_2(l_2+1) - l_1(l_1+1)] + \{l_1 \leftrightarrow l_2\}. \tag{40}$$

The trispectrum associated with point-source shot noise is given by [75]

$$^{S^4} P^{l_1 l_2}_{l_3 l_4}(L) = \frac{1}{3} \langle S^4 \rangle W^{S^2}_{l_1 l_2 L} W^{S^2}_{l_3 l_4 L}, \tag{41}$$

where $\langle S^4 \rangle$ is the kurtosis of the point sources and the weight function is given by

$$W^{S^2}_{l_1 l_2 L} = \begin{pmatrix} l_1 & l_2 & L \\ 0 & 0 & 0 \end{pmatrix} \sqrt{\frac{(2l_1+1)(2l_2+1)(2L+1)}{4\pi}}. \tag{42}$$

We discuss effective amplitudes for S^4 using the WMAP point source masks in Appendix A.

Both point sources and gravitational lensing represent potential sources of bias in the reconstruction of $C^{\bar\Delta \bar\Delta}_L$ and the estimator $\hat{A}$. It is important to either verify that these sources of bias are negligible, or to construct appropriately debiased estimators (as done, for example, for estimators of patchy reionization optical depth τ, in Ref. [76], or for CMB lensing estimators in [77]).

We propagate the trispectrum described by the primary contraction through to Eq. (37), obtaining the ensemble-averaged contribution of the physical effect x (which can denote CIPs, weak gravitational lensing, or point sources) to the CIP trispectrum estimator $\hat{A}$,

$$\langle \bar{C}^{\Delta\Delta}_L \rangle |_x = \mathcal{P}^{\Delta x}_L + \mathcal{S}^{\Delta x}_L. \tag{43}$$

Here $\mathcal{P}^{\Delta x}_L$ and $\mathcal{S}^{\Delta x}_L$ capture the contributions from the primary and secondary contractions of any trispectrum, respectively.

The primary term is given by $\mathcal{P}^{\Delta\Delta}_L = C^{xx}_L \mathcal{R}^{\Delta x}_L$, where the response function $\mathcal{R}^{\Delta x}_L$ is given by

$$\mathcal{R}^{\Delta x}_L = \left[\frac{1}{2L+1} \sum_{ll'} W^\Delta_{ll'L} W^x_{ll'L} F_l F_{l'} \right]^2. \tag{44}$$

The secondary term $\mathcal{S}$ is more complicated, involving Wigner-6j symbols which are numerically intensive to calculate. We estimate these contributions using flat-sky expressions to evaluate the secondary contractions, given by [78]

$$\mathcal{S}^{\Delta x}_L = \int \frac{d^2 l_1}{(2\pi)^2} \int \frac{d^2 l_2}{(2\pi)^2} F_{|\mathbf{l}_1|} F_{|\mathbf{l}_2|} W^\Delta(\mathbf{l}_1, \mathbf{l}_2) W^\Delta(\mathbf{l}'_1, \mathbf{l}'_2)$$
$$\times \{ C_{|\mathbf{l}_1 - \mathbf{l}'_1|} W^x(-\mathbf{l}_1, \mathbf{l}'_1) W^x(-\mathbf{l}_2, \mathbf{l}'_2)$$
$$+ C_{|\mathbf{l}_1 - \mathbf{l}'_2|} W^x(-\mathbf{l}_1, \mathbf{l}'_2) W^x(-\mathbf{l}_2, \mathbf{l}'_1) \}, \tag{45}$$

where $\mathbf{l}_1$ and $\mathbf{l}_2$ are Fourier space angular multipole vectors, $\mathbf{l}_1 + \mathbf{l}_2 = \mathbf{L}$ and the flat-sky weight function is given by

$$W^\Delta(\mathbf{l}_1, \mathbf{l}_2) = C^{\mathrm{T,dT}}_{|\mathbf{l}_1|} + C^{\mathrm{T,dT}}_{|\mathbf{l}_2|}. \tag{46}$$

The flat-sky weight functions for CMB lensing and point-source shot noise are [75,79]

$$W^\phi(\mathbf{l}_1, \mathbf{l}_2) = C^{\mathrm{TT}}_{l_1}[(\mathbf{l}_1 + \mathbf{l}_2) \cdot \mathbf{l}_1] + (\mathbf{l}_1 \leftrightarrow \mathbf{l}_2),$$
$$W^{S^2}(\mathbf{l}_1, \mathbf{l}_2) = 1. \tag{47}$$

A special case of the bias calculation is for $x = \Delta$, which yields the response of $\bar{C}^{\Delta\Delta}_L$ to CIP fluctuations themselves and therefore the normalization of the estimator. We will see that for CIPs, the primary term is dominant. We can therefore use Eq. (44) as an approximate, analytical normalization for $\hat{C}^{\Delta\Delta}_L$, with

$$\mathcal{N}^{\,\mathrm{approx}}_L = [f_{\mathrm{sky}} \mathcal{R}^{\Delta\Delta}_L]^{-1}. \tag{48}$$

The corresponding approximate normalization for the fiducial power spectrum amplitude estimator $\hat{A}$ is given by

$$\mathcal{N}^{\,\mathrm{approx}} = \left[f_{\mathrm{sky}} \sum_{L=L_{\min}}^{L_{\max}} (2L+1)(C^{\Delta\Delta,(\mathrm{fid})}_L)^2 \mathcal{R}^{\Delta\Delta}_L \right]^{-1}. \tag{49}$$

We will ultimately correct this normalization for our $\hat{A}$ estimates using Monte Carlo simulations, which implicitly include the contribution from the secondary contractions above, as well as cut-sky effects beyond simple f_{sky} scaling. For a scale-invariant power spectrum, we find that the approximation of Eq. (49) is accurate to better than 20%.

In Fig. 2 we show estimates of the contribution of primary and secondary contractions of CIPs for a scale-invariant power spectrum, as well as weak lensing of the CMB and point sources. We see that the primary and secondary contractions of the lensing trispectrum yield negligible contributions to the CIP estimator, compared with the estimator noise power spectrum.

The point-source (discussed at length in Appendix A) shot-noise trispectrum also contributes negligibly to the CIP estimator. We also see that secondary contractions of the CIP trispectrum are negligible compared with primary contractions. As the magnitude of these additional biases is orders of magnitude below the maximum allowed signal and the primary estimator noise level, we do not worry about inaccuracies in the flat-sky approach, although they could be quantified using simulations (as done for lensing in Refs. [80,81], where these inaccuracies are shown to be
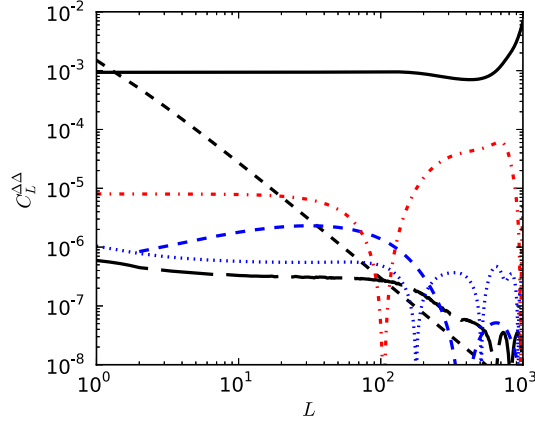
FIG. 2 (color online). Contributions to our trispectrum estimator of the normalized CIP power spectrum $C_L^{\Delta\Delta}$ as a function of scale, for WMAP V-band noise and beam. The estimator noise power spectrum $\mathcal{N}_L^{\mathrm{approx}} D_L^{\Delta\Delta}$ is shown as a black solid curve. For comparison, we plot a scale-invariant power spectrum (black short-dashed) $C_L^{\Delta\Delta} = 0.003 \times [L(L+1)]^{-1}$ with an amplitude comparable to our one standard-deviation constraint. We also calculate the contribution from the secondary contractions for this CIP power spectrum ($\mathcal{N}_L^{\mathrm{approx}} \mathcal{S}_L^{\Delta\Delta}$, black long-dashed). The expected contribution from the primary and secondary contractions of the CMB lensing trispectrum are shown as blue (gray) short-dashed and blue (gray) dotted curves, respectively. In addition, we plot a pessimistic (larger than the maximum in Appendix A) estimate of the V-band contamination from unresolved point sources, with $S^4 = 5 \times 10^{-7}$ μK^4, shown as a red (grey) dot-dashed curve. The primary and secondary contributions from the shot-noise trispectrum have the same shape, and so we only plot their sum. An analogous analysis in the Q and W bands yields the same conclusion, and so we omit those curves here.

negligible at levels of precision comparable to those of the Planck satellite mission).

IV. WMAP ANALYSIS

A. Data and simulations

We analyze the "foreground-reduced" CMB temperature anisotropy maps, with associated beam and noise characterization, from the final, 9-year WMAP data release [82]. We use the Q-, V-, and W-band maps, with nominal central frequencies of 40, 60, and 90 GHz, respectively. To remove known Galactic and bright point source contamination, we use the KQ85 temperature analysis mask produced by the WMAP team.

To compute the mean-field and disconnected noise bias of our estimator [Eqs. (29) and (28)], we use CMB + noise simulations of the data to compute the Monte Carlo averages (300 realizations per WMAP band). We simulate a CMB temperature realization T_{lm}^{sim} with power spectrum C_l^{TT} given by a flat ΛCDM cosmology consistent with the WMAP power spectrum [53], with parameters given in the Introduction, computed using the CAMB [55] code. We apply a beam and pixel transfer function to these temperature multipoles, and then project them onto an $N_{\mathrm{side}} = 512$

HEALPix map with a harmonic transform. To each pixel in this map, we add a Gaussian noise contribution with variance given by Eq. (35). These simple simulations are then filtered in the same way as the real data [that is, using Eq. (31)] to produce the simulated quantities $\bar{T}^{(g)}$ and $\bar{T}^{(f)}$ of Eqs. (29) and (28).

To test the normalization of our CIP estimator, we also require non-Gaussian simulations with $C_L^{\Delta\Delta} \neq 0$. We form these non-Gaussian CMB simulations with power spectrum C_l^{TT} and trispectrum given by Eq. (21) using three pieces $T_{lm} = A_{lm} + \beta_{lm} + C_{lm}$. A_{lm} is drawn from a Gaussian distribution with power spectrum C_l^{AA} (the choice of C_l^{AA} is somewhat arbitrary for our purposes, we will use $C_l^{\mathrm{AA}} = C_l^{\mathrm{TT}}/2$). Then we add the term β_{lm} which ensures that our simulations have the desired trispectrum, with

$$\beta_{lm} = \int d\hat{\mathbf{n}} Y_{lm}^*(\hat{\mathbf{n}})\left[\sum_{LM} Y_{LM}(\hat{\mathbf{n}})\Delta_{LM}\right]$$
$$\times \left[\sum_{l'm'} Y_{l'm'}(\hat{\mathbf{n}}) A_{l'm'} \frac{C_{l'}^{\mathrm{T,dT}}}{C_{l'}^{\mathrm{AA}}}\right]. \quad (50)$$

Here Δ_{LM} are Gaussians drawn for a fiducial power spectrum $C_L^{\Delta\Delta,\mathrm{(fid)}}$. We add a final Gaussian term C_{lm} which sets the overall power spectrum of our non-Gaussian simulations to be C_l^{TT}. The C_{lm} are drawn from the power spectrum $C_l^{CC} = C_l^{\mathrm{TT}} - C_l^{\mathrm{AA}} - C_l^{\beta\beta}$, with

$$C_l^{\beta\beta} = \sum_{ll'} \frac{(2l+1)(2l'+1)}{4\pi}\begin{pmatrix} l & l' & L \\ 0 & 0 & 0 \end{pmatrix}^2 \times \frac{C_L^{\Delta\Delta}(C_{l'}^{\mathrm{T,dT}})^2}{C_{l'}^{\mathrm{AA}}}. \quad (51)$$

B. Results

In Fig. 3 we plot estimated CIP power spectra $\hat{C}_L^{\Delta\Delta}$ from separate analyses of Q-, V-, and W-band maps, as well as an inverse-variance weighted combination of the three maps (denoted QVW). We use the approximate analytical normalization of Eq. (48). This neglects the secondary CIP trispectrum contractions, which we see in Fig. 2 are negligible, as well as cut-sky effects beyond simple f_{sky} scaling. For the question we are asking in this figure ("Is the measured CIP power nonzero with statistical significance?"), the accuracy of the normalization is not important, as corrections to Eq. (48) show up in both the signal and noise, and thus leave the significance of any signal or limit unaffected. We see no evidence for anomalous $C_L^{\Delta\Delta}$ power in the individual Q-, V-, or W-band maps. We may now obtain limits on specific models for the CIP power spectrum, as well as model-independent constraints to the CIP power spectrum at different multipole numbers L.

1. Limits to a scale-invariant power spectrum of CIPs

We now place limits on the amplitude of a scale-invariant power spectrum of CIP fluctuations, which in
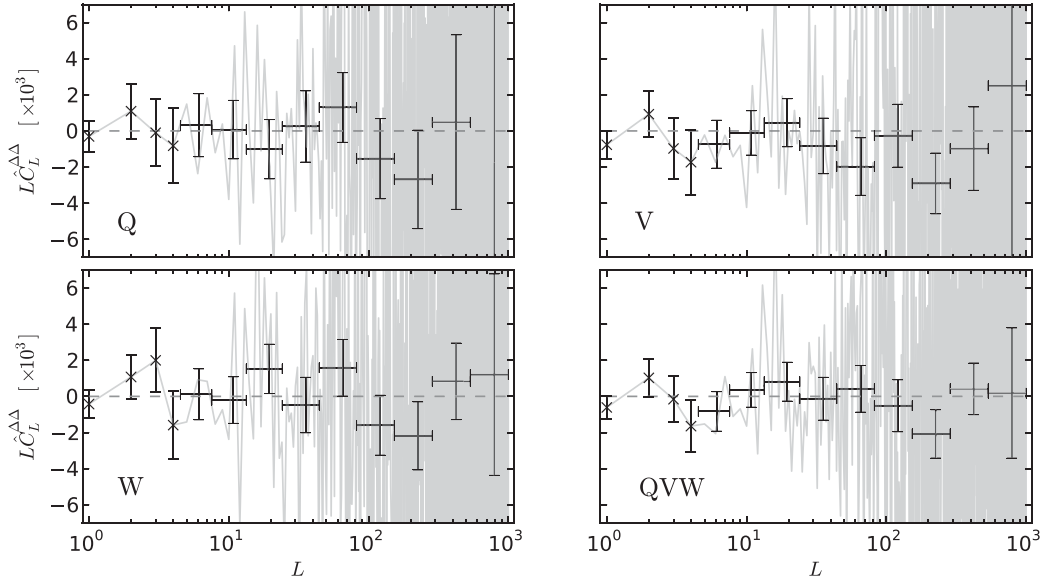
FIG. 3. Estimates of the CIP power spectrum $C_L^{\Delta\Delta}$ using the WMAP Q-, V-, and W-band data. The QVW points give the result when all three bands are combined with an inverse-variance filter. Gray lines give the measurements for individual multipoles.

multipole space projected onto a two-dimensional sky is [50]

$$C_L^{\Delta\Delta} = C_L^{\Delta\Delta,(\text{fid})} \simeq \frac{A}{L(L+1)}, \qquad (52)$$

for some amplitude A. For $A = 1$, corresponding realizations of Δ would have fluctuations with amplitude of approximately unity. For $L \gtrsim 870$, the scaling in Eq. (52) actually turns over to L^{-3} as the wavelength of the CIP falls below the thickness of the last-scattering surface. These modes have very low signal-to-noise ratio using present data and thus contribute negligibly to our estimator of A. It is thus adequate to use Eq. (52). The results of these fits are shown in Fig. 4, and are also consistent with zero. Note that a value of $A < 0$ is nonphysical, but allowed by our measurement as we have not imposed a prior on the positivity of A.

For our measurement, we use Eq. (30), $L_{\min} = 2$, and $L_{\max} = 1000$. We use the normalization $\mathcal{N} = \mathcal{N}^{\text{approx}}$, where $\mathcal{N}^{\text{approx}}$ is given in Eq. (49). Although this normalization is only approximate, it is adequate to answer the question of whether or not the measured value of A deviates from zero with statistical significance. Later, when limits on A are set, we account for corrections to this normalization.

Our constraint on $\hat{A}$ comes mainly from very low-L modes of $\bar{C}_L^{\Delta\Delta}$. The noise variance of $\hat{C}_L^{\Delta\Delta}$ at $L < 100$ goes as $1/(2L+1)$, which falls off more slowly than the scale-invariant spectrum, and so most of our sensitivity to A comes from a small number of low-L modes. Our estimator probability density function (PDF) is thus slightly non-Gaussian (as can be seen from the histograms of Fig. 4), for the same reason that the distribution of low-L power-spectrum estimates is non-Gaussian. This

also occurs for trispectrum-based estimators of the amplitude of local-type non-Gaussianity $|f_{\text{NL}}|$ [83].

We find best-fit values for Q-, V-, and W-bands separately, as well as for the QVW combination, of $\hat{A}_Q \simeq 1.4 \times 10^{-3}$, $\hat{A}_V \simeq -8.4 \times 10^{-4}$, $\hat{A}_W \simeq 2.7 \times 10^{-3}$, $\hat{A}_{QVW} \simeq 3.9 \times 10^{-4}$. In Fig. 4, these values are compared with the distribution of $\hat{A}$ values for simulations in which the real value $A = 0$. We see that our best-fit values are consistent with the null hypothesis at 95% confidence.

To establish upper limits on A, we need the function $P(A|\hat{A})$, the probability distribution of A values, given the estimated value $\hat{A}$. Assuming flat priors for the "data" (estimator value $\hat{A}$) and theoretical parameter (A), we have $P(A|\hat{A}) \propto P(\hat{A}|A)$. From our Monte Carlo simulation of the $A = 0.01$ case, we know that the width of the non-Gaussian function $P(\hat{A}|A)$ depends noticeably on the real value of A, and so it is important to properly determine $P(\hat{A}|A)$.

To this end, we conduct a suite of Monte Carlo simulations, estimating $\hat{A}_{QVW}$ for multiple realizations (300 per band and per real A value), at the series of true parameter values $A = \{0.000, 0.002, 0.005, 0.010, 0.020\}$. The results for $P(\hat{A}|A)$ are shown in Fig. 5. To compute an upper limit to A, we use a model for $P(A|\hat{A})$ [and thus $P(\hat{A}|A)$].

The model for $P(\hat{A}|A)$ is as follows: A χ^2 distribution $P_\chi(\mathcal{E}_L; k_L, s_L)$ is fit to the simulated random variable $\mathcal{E}_L \equiv \mathcal{N}(2L+1)\bar{C}_L^{\Delta\Delta}C_L^{\Delta\Delta,\text{fid}}$ for $2 \leq L \leq 20$. At each L and for any given A value, this distribution has two free parameters, the number of degrees of freedom k_L, and a scale parameter s_L such that $\langle \mathcal{E}_L \rangle = k_L s_L$. The estimator used is given by

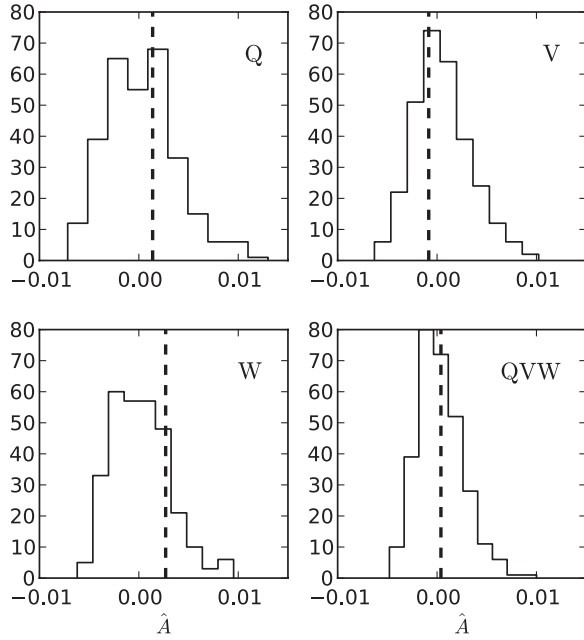$$\hat{A} = \sum_{L=L_{\min}}^{L=L_{\max}} \mathcal{E}_L. \qquad (53)$$

FIG. 4. Fits for the amplitude of a scale invariant power spectrum [Eq. (52)], using the WMAP Q-, V-, and W-band data. Dashed vertical lines give the fits themselves, while the solid histograms give the distribution of $\hat{A}$ measured on simulations with $\hat{A} = 0$.

FIG. 5. Arbitrarily normalized probability distribution for $\hat{A}_{\mathrm{QVW}}$ in the presence of a scale-invariant spectrum of CIPs $C_L^{\Delta\Delta} = A/[L(L+1)]$. The estimated value is determined from an appropriately inverse-variance weighted sum of the three maps, simulated with an ensemble of Monte Carlo simulations. Dashed vertical line shows input value of A. Dashed curves show $P^{\alpha}(\hat{A}|A)$, while solid curves show $P^{\beta}(\hat{A}|A)$. Both models are described in Sec. IV B.

A model PDF $P^{\alpha}(\hat{A}|A)$ is thus obtained by convolution over $P_{\chi}(\mathcal{E}_L; k_L, s_L)$. This model is shown as the dashed curve in Fig. 5 and is clearly a reasonable fit to the data.

Alternatively, we may just fit $P_{\chi}(\mathcal{E}_L; k_L, s_L)$ for the $A = 0$ case and then adjust the scale parameter $s_L \rightarrow s_L \times \{(1 + A/[L(L+1)(2L+1)]\}$. An analogous convolution then yields the semianalytic function $P^{\beta}(\hat{A}|A)$, which is comparably accurate to $P^{\alpha}(\hat{A}|A)$ in fitting the estimator PDF from the simulation, and readily interpolated to obtain quantities of interest.

Our final upper limit is computed by starting with $P^{\beta}(\hat{A}|A)$, obtaining $P(A|\hat{A})$ using Bayes's theorem with a flat prior on A in the simulated domain, computing the cumulative probability function (CDF) associated with $P(A|\hat{A})$, and interpolating this function, incrementing A upwards to define a 95% confidence interval. Like the measurement of $\hat{A}_{\mathrm{QVW}}$, these simulations use $\mathcal{N} = \mathcal{N}^{\mathrm{approx}}$, and thus calibrate the relationship between A and $\hat{A}_{\mathrm{QVW}}$, accounting for errors induced in the approximation $\mathcal{N} = \mathcal{N}^{\mathrm{APPROX}}$. The resulting upper limits on A are thus correctly normalized.

Given the observed multiband value $\hat{A}_{\mathrm{QVW}} = 3.9 \times 10^{-4}$, we compute $P(3.9 \times 10^{-4}|A)$ [and thus $P(A|3.9 \times 10^{-4})$] and find that $A \leq 1.1 \times 10^{-2}$ at 95% confidence. Measurements of the baryon fraction in galaxy clusters [34] impose the limit $A \leq 5.4 \times 10^{-3}$ to the amplitude of a scale-invariant spectrum of CIPs [49,50]. The larger value quoted in Ref. [49] is obtained using a different definition
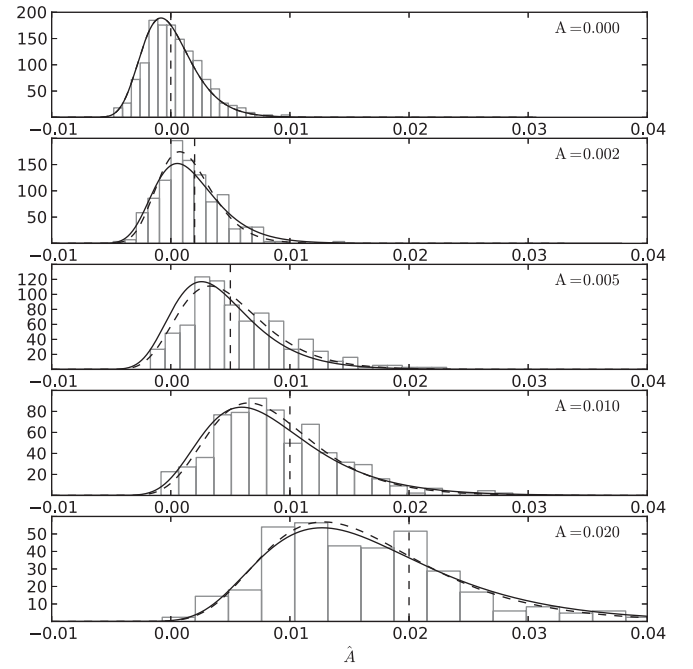
of A, which we have appropriately rescaled here. We may also restate our result in terms of Δ_{cl}, the rms fluctuation in the baryon-dark matter density ratio on galaxy cluster scales, using our definition of A and Eq. (1) of Ref. [49] (with the associated wave-number range $k_{\mathrm{min}} \lesssim k \lesssim k_{\mathrm{max}}$) to obtain

$$\Delta_{\mathrm{cl}}^2 \simeq \frac{A \ln(1000)}{2\pi}. \tag{54}$$

We obtain $\Delta_{\mathrm{cl}} \lesssim 0.11$, to be compared with the result directly obtained from the cluster baryon fraction, $\Delta_{\mathrm{cl}} \lesssim 0.077$.

As predicted in the forecasts of Ref. [50], the CIP sensitivity of WMAP is comparable to measurements of the baryon fraction in galaxy clusters. The physics of this CMB probe of CIPs is, however, completely different than that used in the cluster probe, offering an important and truly primordial constraint on the amplitude of CIPs.

2. Constraints to model-independent CIP amplitude at different angular scales

Aside from placing limits to a scale-invariant CIP power spectrum, we may also place limits on the CIP amplitude at different angular multipole numbers L, without reference to a fiducial model.

The PDF for $P^\beta(\hat{A}|A)$ used to obtain the above upper limits is derived by fitting a χ^2 distribution to the PDF for each $\hat{C}_L^{\Delta\Delta}$, assuming no covariance between $\hat{C}_L^{\Delta\Delta}$ for different L, convolving over different multipole moments, and applying Bayes's theorem. This *ansatz* fits the Monte Carlo histograms well, and so we may use the same simulation to impose limits on $C_L^{\Delta\Delta}$ independently for each scale L, without making any assumptions about the functional form of the CIP power spectrum.

We thus model the histogram of $\hat{C}_L^{\Delta\Delta}$ values using the function $P_\chi(\mathcal{E}_L; k_L, s_L)$, where $s_L = s_L^{A=0} \times \{[1 + C_L^{\Delta\Delta}/(2L + 1)]\}$. Here $s_L^{A=0}$ is the best-fit value of the χ^2 scale-parameter s_L for the null hypothesis simulation. For all multipoles in the range, $1 \leq L \leq 20$, this PDF fits the Monte Carlo results. Applying Bayes's theorem on an L-by-L basis with the resulting PDF, we obtain $P(C_L^{\Delta\Delta}|\hat{C}_L^{\Delta\Delta})$. Using the estimated values $\hat{C}_L^{\Delta\Delta}$ from the minimum-variance-weighted QVW combination, we obtain 95% confidence upper limits on $C_L^{\Delta\Delta}$ in the L-range of interest. The results are shown in Table I and Fig. 6, where the corresponding angular scale and rms CIP amplitude upper limit is shown. We thus see that the relative mass fractions of baryons and dark matter can vary by no more than 7%–17% on angular scales of 5–100° at the
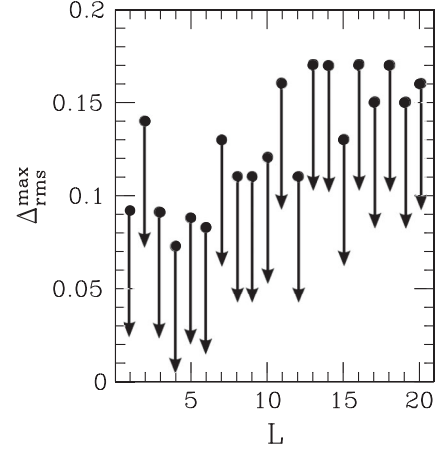

FIG. 6. Upper limits (95% confidence level) to the RMS CIP amplitude as a function of multipole number L, as recovered from an inverse-variance weighted sum of the three WMAP 9-year CMB maps, and the Monte Carlo simulations described in Sec. IV B 1.

SLS. This limit is forced upon us by CMB data alone, with no reliance on galactic abundance measurements, galaxy physics, or knowledge about the nature of cosmic reionization.

V. CONCLUSIONS

Compensated isocurvature perturbations (CIPs) break the usual assumption that the primordial cosmic baryon fraction is spatially homogeneous. Here, building on past work searching for CIPs using variations in the cluster baryon fraction, we use WMAP 9-year maps of the CMB temperature fluctuations to search for CIPs. We apply a trispectrum-based estimator to show that the amplitude of a scale-invariant spectrum of CIPs must obey the constraint $A \lesssim 1.1 \times 10^{-2}$, combining all three CMB-dominated WMAP bands. Independent of the functional form of the CIP power spectrum, the rms CIP amplitude must be no greater than $\sim 10\%$ on scales of 5–100°.

As we saw, the estimator for $C_L^{\Delta\Delta}$ is an appropriately filtered four-point function of the data, probing correlations that are absent under the null hypothesis. Gaussian CIPs uncorrelated with the underlying adiabatic fluctuation can then be thought of as inducing a non-Gaussian signal in the CMB, with vanishing three-point function and nonvanishing four-point function.

It may be that CIPs are not an independent Gaussian random field (as assumed here), but rather, as in some curvaton models [33], correlated with the usual adiabatic fluctuations. In this case, schematically, the induced temperature fluctuation is $\delta T \sim \Delta\Phi$, as before, but now, since $\Delta = A\Phi$ for some correlation coefficient, $\delta T \sim \Phi^2$. The resulting CMB three-point function is then $\langle TTT \rangle \propto \langle \delta TTT \rangle \propto \langle \Phi^4 \rangle$. Curvaton-inspired CIPs will thus cause an additional non-Gaussian signature, a three-point correlation or bispectrum.

TABLE I. Upper limits (95% confidence level) $C_L^{\Delta\Delta,\mathrm{max}}$ to the angular power spectrum of CIPs as a function of multipole number L, as recovered from an inverse-variance weighted sum of the three WMAP 9-year CMB maps, and the Monte Carlo simulations described in Sec. IV B 1. Also shown is the corresponding angular scale in degrees of $\theta^\circ \simeq 100^\circ/L$ and the 95% confidence upper limit to the rms CIP amplitude on that angular scale, $\Delta_{\mathrm{rms}}^{\mathrm{max}} \simeq \sqrt{L(L+1)C_L^{\Delta\Delta,\mathrm{max}}/(2\pi)}$.

L	$C_L^{\Delta\Delta,\mathrm{max}}$	θ (in °)	$\Delta_{\mathrm{rms}}^{\mathrm{max}}$
1	2.7×10^{-2}	100	9.2×10^{-2}
2	2.1×10^{-2}	50	1.4×10^{-1}
3	4.3×10^{-3}	33	9.1×10^{-2}
4	1.7×10^{-3}	25	7.3×10^{-2}
5	1.6×10^{-3}	20	8.8×10^{-2}
6	1.0×10^{-3}	17	8.3×10^{-2}
7	2.0×10^{-3}	14	1.3×10^{-1}
8	1.0×10^{-3}	13	1.1×10^{-1}
9	9.0×10^{-4}	11	1.1×10^{-1}
10	7.8×10^{-4}	10	1.2×10^{-1}
11	1.2×10^{-3}	9.1	1.6×10^{-1}
12	4.6×10^{-4}	8.3	1.1×10^{-1}
13	9.7×10^{-4}	7.7	1.7×10^{-1}
14	7.8×10^{-4}	7.1	1.7×10^{-1}
15	4.5×10^{-4}	6.7	1.3×10^{-1}
16	6.9×10^{-4}	6.3	1.7×10^{-1}
17	4.9×10^{-4}	5.9	1.5×10^{-1}
18	5.3×10^{-4}	5.6	1.7×10^{-1}
19	3.9×10^{-4}	5.3	1.5×10^{-1}
20	4.0×10^{-4}	5.0	1.6×10^{-1}

We note that the CIP estimator described here could also be applied to the recently released nominal mission data of the *Planck* satellite, for which we forecast a sensitivity improvement factor of ∼3 in Δ_{cl} or A [50]. Future ground-based CMB polarization data (such as ACTPol [84] and SPTpol [85]) also have the potential to significantly improve the sensitivity of the CIP probe described here.

It is claimed in Ref. [33] that CIPs would lead to negligible changes in the total matter power spectrum, and would thus not be detectable using galaxy surveys. It stands to reason, however, that star formation efficiency, cooling, and other processes are strongly dependent on the baryon density, and that the statistics of actual galaxies would be drastically altered in the presence of CIPs. Using semianalytic models of galaxy formation, this could be tested. Large fluctuations in the Jeans/cooling masses in the presence of a CIP could also change the dynamics of reionization, leading to an additional probe. We will explore these possibilities in future work.

ACKNOWLEDGMENTS

APPENDIX: POINT SOURCES

Point sources that are bright enough to be detected are removed using the WMAP analysis mask used in this paper. There is, however, a population of residual sources which are too faint to be directly detected individually in the maps, but which contribute non-negligibly to the CMB temperature power spectra, as well as higher-order statistics like the trispectrum which we probe here. When modeling these sources, it is useful to consider an effective limiting flux threshold S_c, above which all sources are assumed to be detected, and below which none are. At WMAP frequencies and sensitivity, the dominant residual contribution is from a small number of radio sources directly below the detection threshold. The statistics of this population are dominated by the one-point (or "shot-noise") component. If the differential number density of sources with flux S is denoted dN/dS, the unresolved shot noise contribution to the reduced trispectrum is given by [86–88]

$$\langle S^4 \rangle \equiv g^4(x) \int_0^{S_c} dS S^4 \frac{dN}{dS}, \qquad (A1)$$

where $x = h\nu/(k_B T) = \nu/56.84$ GHz for experimental channel frequency ν and $g(x)$ is a conversion factor from Jansky flux units to thermodynamic μK, given by

$$g(x) \equiv 2\frac{(hc)^2}{(k_B T)^3}\left[\frac{\sinh^2(x/2)}{x^4}\right] \simeq \frac{(e^x-1)^2}{x^4 e^x}\frac{\mu\mathrm{K}}{99.27\ \mathrm{Jy\ sr}^{-1}}. \qquad (A2)$$

The physical constants here are the Planck constant h, speed of light c, Boltzmann constant k_B, and CMB mean temperature T.

We use two models for dN/dS, to be sure that our conclusions about the contribution of shot noise to the trispectrum are robust. The first ("model 1" hereafter) is given in Ref. [89], and is parametrized by $dN/dS = N_0/S^\beta$, where $\beta = 2.15$, and $N_0 = 12$ Jy$^{1.15}$ sr^{-1}. The second ("model 2" hereafter) is obtained from tables of source counts in Ref. [90]. We fit a power law of the form used above to these tables, and obtain $\beta = 2.74$ and $N_0 = 44.2$ for WMAP Q-band source counts. For the V-band source counts, we obtain $\beta = 2.59$ and $N_0 = 21.8$. W-band source counts are not given in these tables, but we use the Q- and V-band fits (cases Wα and Wβ) as a model for W-band source counts. Flux cuts and conversion functions $g(x)$ appropriate for the W band, however, are used for the W-band trispectrum estimates.

We now use Eq. (A1) to obtain estimates for the unresolved point-source trispectra. We use a flux cut of $S_c \simeq 1.00$ Jy, as it was found in Ref. [91] that this value reproduces observed WMAP point-source bispectra. For the Q-band data with model 1, we obtain $\langle S^4 \rangle = 7.03 \times 10^{-7}\ \mu$K^4. For the Q-band data with model 2, we obtain $\langle S^4 \rangle = 3.29 \times 10^{-6}\ \mu$K^4. For V-band data with model 1, we obtain $\langle S^4 \rangle = 3.60 \times 10^{-8}\ \mu$K^4. For V-band data with model 2, we obtain $\langle S^4 \rangle = 7.71 \times 10^{-8}\ \mu$K^4. For W-band data with model 1, we obtain $\langle S^4 \rangle = 1.89 \times 10^{-9}\ \mu$K^4. For model 2, case W$\alpha$, we obtain $\langle S^4 \rangle = 8.78 \times 10^{-9}\mu$K^4. For model 2, case W$\beta$, we obtain $\langle S^4 \rangle = 4.05 \times 10^{-9}\mu$K^4.

We plot a pessimistic estimate of the point-source contamination in Fig. 2 for the WMAP V-band. As we know from the preceding discussion, the W-band contribution is negligible in comparison. The Q band has the most worrisome amplitude of possible point source bias; however, it is still more than an order of magnitude below the scale-invariant power spectrum shape at the low ($L < 10$) multipoles and amplitudes ($\hat{A} < 5 \times 10^{-3}$), which we probe most sensitively. We do not see evidence for any departure from zero in our model-independent $C_L^{\Delta\Delta}$ measurement, which would be characteristic of point-source contamination.

[1] A. H. Guth and S. Pi, Phys. Rev. Lett. **49**, 1110 (1982).
[2] A. D. Linde, Phys. Lett. **116B**, 335 (1982).
[3] J. M. Bardeen, P. J. Steinhardt, and M. S. Turner, Phys. Rev. D **28**, 679 (1983).
[4] S. Hawking, Phys. Lett. **115B**, 295 (1982).
[5] V. F. Mukhanov and G. Chibisov, Sov. Phys. JETP **56**, 258 (1982).
[6] A. A. Starobinsky, Phys. Lett. **117B**, 175 (1982).
[7] R. H. Brandenberger, Int. J. Mod. Phys. A **09**, 2117 (1994).
[8] A. D. Linde, JETP Lett. **40**, 1333 (1984).
[9] A. D. Linde and V. F. Mukhanov, Phys. Rev. D **56**, R535 (1997).
[10] V. F. Mukhanov, H. Feldman, and R. H. Brandenberger, Phys. Rep. **215**, 203 (1992).
[11] D. Langlois and A. Riazuelo, Phys. Rev. D **62**, 043504 (2000).
[12] D. Langlois, Phys. Rev. D **59**, 123512 (1999).
[13] M. Axenides, R. H. Brandenberger, and M. S. Turner, Phys. Lett. **126B**, 178 (1983).
[14] D. Seckel and M. S. Turner, Phys. Rev. D **32**, 3178 (1985).
[15] P. J. E. Peebles, Astrophys. J. **510**, 523 (1999).
[16] P. J. E. Peebles, Astrophys. J. **510**, 531 (1999).
[17] S. D. Burns, arXiv:astro-ph/9711303.
[18] W. Hu, Phys. Rev. D **59**, 021301(R) (1998).
[19] W. Hu, E. F. Bunn, and N. Sugiyama, Astrophys. J. **447**, L59 (1995).
[20] K. Enqvist, H. Kurki-Suonio, and J. Valiviita, Phys. Rev. D **62**, 103003 (2000).
[21] K. Enqvist and H. Kurki-Suonio, Phys. Rev. D **61**, 043002 (2000).
[22] D. Larson *et al.*, Astrophys. J. Suppl. Ser. **192**, 16 (2011).
[23] E. Komatsu *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **192**, 18 (2011).
[24] P. Ade *et al.* (Planck Collaboration), arXiv:1303.5082.
[25] J. Valiviita and T. Giannantonio, Phys. Rev. D **80**, 123516 (2009).
[26] C. Zunckel, P. Okouma, S. Muya Kasanda, K. Moodley, and B. A. Bassett, Phys. Lett. B **696**, 433 (2011).
[27] M. Bucher, J. Dunkley, P. Ferreira, K. Moodley, and C. Skordis, Phys. Rev. Lett. **93**, 081301 (2004).
[28] M. Kawasaki and T. Sekiguchi, Prog. Theor. Phys. **120**, 995 (2008).
[29] M. Beltrán, J. García-Bellido, J. Lesgourgues, and M. Viel, Phys. Rev. D **72**, 103515 (2005).
[30] U. Seljak, A. Slosar, and P. McDonald, J. Cosmol. Astropart. Phys. 10 (2006) 014.
[31] M. Beltran, J. Garcia-Bellido, J. Lesgourgues, and A. Riazuelo, Phys. Rev. D **70**, 103530 (2004).
[32] T. Takahashi, M. Yamaguchi, and S. Yokoyama, Phys. Rev. D **80**, 063524 (2009).
[33] C. Gordon and J. R. Pritchard, Phys. Rev. D **80**, 063535 (2009).
[34] G. P. Holder, K. M. Nollett, and A. van Engelen, Astrophys. J. **716**, 907 (2010).
[35] D. H. Lyth, C. Ungarelli, and D. Wands, Phys. Rev. D **67**, 023503 (2003).
[36] S. Gupta, K. A. Malik, and D. Wands, Phys. Rev. D **69**, 063513 (2004).
[37] C. Gordon and A. Lewis, Phys. Rev. D **67**, 123513 (2003).
[38] K. Enqvist, S. Nurmi, G. Rigopoulos, O. Taanila, and T. Takahashi, J. Cosmol. Astropart. Phys. 11 (2009) 003.

[39] D. Spolyar, arXiv:1111.3629.
[40] D. E. Kaplan, M. A. Luty, and K. M. Zurek, Phys. Rev. D **79**, 115016 (2009).
[41] M. R. Buckley and L. Randall, J. High Energy Phys. 09 (2011) 009.
[42] R. Allahverdi, B. Dutta, and K. Sinha, Phys. Rev. D **83**, 083502 (2011).
[43] P.-H. Gu, M. Lindner, U. Sarkar, and X. Zhang, Phys. Rev. D **83**, 055008 (2011).
[44] J. J. Heckman and S.-J. Rey, J. High Energy Phys. 06 (2011) 120.
[45] J. McDonald, Phys. Rev. D **83**, 083509 (2011).
[46] A. Lewis and A. Challinor, Phys. Rev. D **76**, 083005 (2007).
[47] R. Barkana and A. Loeb, Mon. Not. R. Astron. Soc. **363**, L36 (2005).
[48] M. Kawasaki, T. Sekiguchi, and T. Takahashi, J. Cosmol. Astropart. Phys. 10 (2011) 028.
[49] D. Grin, O. Doré, and M. Kamionkowski, Phys. Rev. Lett. **107**, 261301 (2011).
[50] D. Grin, O. Doré, and M. Kamionkowski, Phys. Rev. D **84**, 123003 (2011).
[51] K. Sigurdson, A. Kurylov, and M. Kamionkowski, Phys. Rev. D **68**, 103509 (2003).
[52] A. Lewis and A. Challinor, Phys. Rep. **429**, 1 (2006).
[53] G. Hinshaw *et al.* (WMAP Collaboration), arXiv:1212.5226.
[54] U. Seljak and M. Zaldarriaga, Astrophys. J. **469**, 437 (1996).
[55] A. Lewis and A. Challinor, Phys. Rev. D **66**, 023531 (2002).
[56] S. Dodelson, *Modern Cosmology* (Academic Press, New York, 2003).
[57] M. Zaldarriaga and D. D. Harari, Phys. Rev. D **52**, 3276 (1995).
[58] C.-P. Ma and E. Bertschinger, Astrophys. J. **455**, 7 (1995).
[59] D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonsky, *Quantum Theory of Angular Momentum: Irreducible Tensors, Spherical Harmonics, Vector Coupling Coefficients, 3nj Symbols.* (World Scientific, Singapore, 1988).
[60] A. Hajian and T. Souradeep, Astrophys. J. Lett. **597**, L5 (2003).
[61] A. Hajian and T. Souradeep, arXiv:astro-ph/0501001.
[62] W. Hu, Phys. Rev. D **64**, 083005 (2001).
[63] P. G. Ferreira and J. Magueijo, Phys. Rev. D **56**, 4578 (1997).
[64] D. Regan, E. Shellard, and J. Fergusson, Phys. Rev. D **82**, 023520 (2010).
[65] C. M. Hirata and U. Seljak, Phys. Rev. D **67**, 043001 (2003).
[66] M. Kamionkowski, Phys. Rev. Lett. **102**, 111302 (2009).
[67] A. R. Pullen and M. Kamionkowski, Phys. Rev. D **76**, 103529 (2007).
[68] C. Dvorkin, W. Hu, and K. M. Smith, Phys. Rev. D **79**, 107302 (2009).
[69] C. Dvorkin and K. M. Smith, Phys. Rev. D **79**, 043003 (2009).
[70] V. Gluscevic, M. Kamionkowski, and A. Cooray, Phys. Rev. D **80**, 023510 (2009).
[71] R. R. Caldwell, V. Gluscevic, and M. Kamionkowski, Phys. Rev. D **84**, 043504 (2011).

[72] D. Hanson and A. Lewis, Phys. Rev. D **80**, 063004 (2009).

[73] N. Jarosik *et al.*, Astrophys. J. Suppl. Ser. **192**, 14 (2011).

[74] K. M. Smith, O. Zahn, and O. Dore, Phys. Rev. D **76**, 043510 (2007).

[75] S. Osborne *et al.* (to be published).

[76] M. Su *et al.*, arXiv:1106.4313.

[77] T. Namikawa, D. Hanson, and R. Takahashi, arXiv:1209.0091.

[78] M. H. Kesden, A. Cooray, and M. Kamionkowski, Phys. Rev. D **67**, 123507 (2003).

[79] W. Hu and T. Okamoto, Astrophys. J. **574**, 566 (2002).

[80] D. Hanson, A. Challinor, G. Efstathiou, and P. Bielewicz, Phys. Rev. D **83**, 043005 (2011).

[81] M. M. Schmittfull, A. Challinor, D. Hanson, and A. Lewis, Phys. Rev. D **88**, 063012 (2013).

[82] C. Bennett *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **208**, 20 (2013).

[83] T. L. Smith and M. Kamionkowski, Phys. Rev. D **86**, 063009 (2012).

[84] M. D. Niemack *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **7741**, 77411S (2010).

[85] J. J. McMahon *et al.*, AIP Conf. Proc. **1185**, 511 (2009).

[86] F. Argueso, J. Gonzalez-Nuevo, and L. Toffolatti, Astrophys. J. **598**, 86 (2003).

[87] E. Komatsu and D. N. Spergel, Phys. Rev. D **63**, 063002 (2001).

[88] P. Serra, A. Cooray, A. Amblard, L. Pagano, and A. Melchiorri, Phys. Rev. D **78**, 043004 (2008).

[89] K. M. Smith *et al.*, AIP Conf. Proc. **1141**, 121 (2009).

[90] G. Zotti, M. Massardi, M. Negrello, and J. Wall, Astron. Astrophys. Rev. **18**, 1 (2010).

[91] E. Komatsu *et al.* (WMAP Collaboration), Astrophys. J. Suppl. Ser. **148**, 119 (2003).

[92] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, Astrophys. J. **622**, 759 (2005).