



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

BAO modulation as a probe of compensated isocurvature perturbations

Chen Heinrich and Marcel Schmittfull

Phys. Rev. D **100**, 063503 — Published 4 September 2019

DOI: [10.1103/PhysRevD.100.063503](https://doi.org/10.1103/PhysRevD.100.063503)

BAO Modulation as a Probe of Compensated Isocurvature Perturbations

Chen Heinrich^{1,*} and Marcel Schmittfull²

¹*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109*

²*Institute for Advanced Study, Princeton, New Jersey 08540*

Compensated isocurvature perturbations (CIPs) are opposite spatial fluctuations in the baryon and dark matter density. They can be generated for example in the curvaton model in the early Universe but are difficult to observe because their gravitational imprint nearly cancels. We therefore propose a new measurement method by searching for a spatial modulation of the baryon acoustic oscillation (BAO) scale that CIPs induce. We find that for a Euclid-like survey the sensitivity is marginally better than the WMAP cosmic microwave background (CMB) constraint, which exploits the CIP-induced modulation of the CMB sound horizon. For a cosmic-variance limited BAO survey using emission-line galaxies up to $z \sim 7$ the sensitivity is between stage 3 and stage 4 CMB experiments. These results include using CIP-galaxy cross-correlations, which improves the sensitivity by a factor of $\sim 2 - 3$ for correlated CIPs. The method could be further improved with an optimal estimator, similarly to the CMB, and could provide a useful cross-check of other CIP probes. Finally, if CIPs exist, they can bias cosmological measurements made assuming no CIPs. In particular, they can act as a super-sample fluctuation of the baryon density and bias measurements of the BAO scale. For modern BAO surveys, the largest 1σ isotropic (2σ radial) CIP fluctuation allowed by Planck's 95% bound could bias BAO measurements of $H(z)$ by 1% (6.6%), partially reducing the tension with the local H_0 measurements from 3.1σ to 2.7σ (0.6σ).

I. INTRODUCTION

Measurements of the cosmic microwave background (CMB) have shown that the Universe started with adiabatic initial conditions [1–3] in which all particle species have the same fractional fluctuations in space. Isocurvature perturbations, on the other hand, are differences between the fractional fluctuations of two species and could arise from additional fields during inflation. Although the total matter-to-radiation isocurvature has been well constrained by *Planck* to be less than a few percents [3], there remains the possibility of an isocurvature mode orthogonal to that of the matter-radiation – the compensated isocurvature perturbations (CIPs) – which are still relatively unconstrained [3–6].

In the CIP mode, the baryon and cold dark matter (CDM) have opposite fluctuations while the radiation has no perturbations [7]. The result is that the total gravitational effects are canceled at linear order, leaving no matter-radiation isocurvature nor gravitational imprints (at least for scales above the baryonic Jeans scale). CIPs could arise from curvaton models in which an additional scalar field, the curvaton, is a spectator during inflation but later decays and seeds most of the curvature or adiabatic perturbations in the Universe [8–11]. Depending on different decay scenarios of the curvaton, CDM or baryons can be produced by, before or after the curvaton decay, giving rise to different amounts of CIPs.

The CIPs in the curvaton model are fully or partially correlated with the adiabatic perturbations depending on how much of the Universe's energy density is dominated by the curvaton at its decay. The largest possible fully

correlated CIPs in the curvaton model would be about $A \approx 16$ times the adiabatic perturbations, corresponding to the curvaton scenario in which baryon number is produced by curvaton decay, and dark matter before the decay (see e.g. [12]). This leaves room for potential detection with a variety of cosmological probes.

A few such probes of CIPs have been proposed so far, e.g. gas fraction of galaxy clusters [13], primordial non-Gaussianities [14] and 21cm measurements [15] (see Ref. [3, 6] for a more comprehensive review). One of the major methods to distinguish curvaton models is to search for the effects of a modulated sound horizon in the CMB. Even though the gravitational effects of the baryon fluctuations are compensated by those of dark matter in the CIP mode, there is still a spatial modulation of the baryon-to-photon ratio, leading to a modulation of the sound speed in the photon-baryon fluid before recombination. Such a modulation can be measured similarly to gravitational lensing of the CMB – via a quadratic estimator or from the smoothing of acoustic peaks in the CMB power spectra.

Applying these methods on WMAP and *Planck* data produced some of the most stringent constraints on CIP fluctuations to date. The quadratic estimator was used on WMAP temperature data, placing an 95% upper limit on the CIP r.m.s. fluctuations at cluster scale (~ 10 Mpc) $\Delta_{\text{cl}}^2 < 0.012$ [4] (or $A^2 < (870)^2$ for correlated CIPs). More recently, the *Planck* CMB power spectra and/or lensing power spectra were also used (e.g. [3, 6, 16]), the latest constraint being from the *Planck* team [3] with $\Delta_{\text{rms}}^2 = 0.0037_{-0.0021}^{+0.0016}$ for CIP r.m.s. on the scale of the last-scattering surface (or $A^2 \approx (580)^2$). The conversion to the correlated CIP amplitude A (defined in Eq. 3) used in this work as well as caveats are described in Appendix B.

We propose in this work a new method of looking for

* chenhe@caltech.edu




FIG. 1. 1D illustration of the effect of a large scale CIP $\Delta(X)$ on the BAO peak in the local correlation function of galaxies $\xi(r)$ in the separate universe limit. The local baryon density in a subvolume is $\Omega_b = \bar{\Omega}_b(1 + \Delta(X))$, and the local CDM density is compensated, $\Omega_c = \bar{\Omega}_c - \bar{\Omega}_b\Delta(X)$, leaving $\Omega_b + \Omega_c$ unchanged. The spatial modulation of the baryon-to-photon ratio changes the plasma sound speed before recombination and leads to a modulation of the BAO scale, observable in the local correlation function of galaxies (insets: dashed red with CIP, solid black without CIP).

CIPs by noting that a modulated sound speed affects not only the acoustic peaks in the CMB, but also the clustering of galaxies: When baryons stop feeling the radiation pressure after recombination, the primordial sound wave stops, leading to an enhanced clustering of baryons (and galaxies) when separated by the distance the sound wave traveled until recombination — the baryon acoustic oscillation (BAO) scale. A modulation of the sound speed therefore leads to a modulation of the BAO scale as a function of 3-dimensional position (see Fig. 1 for a simple illustration). As a result, galaxy surveys can constrain CIPs by searching for such a modulation of the BAO scale. We will outline for the first time a procedure for measuring CIPs from BAO surveys in this way, and forecast the errors on the amplitude of fully correlated CIPs for future surveys.

We show that when the CIP-galaxy cross spectrum is included, a fiducial next-generation galaxy survey yields similar constraints to WMAP [4]. We find that the cosmic-variance-limit (CVL) using $z \leq 7$ emission line galaxies yields a constraint between stage 3 and stage 4 CMB experiments, which is about 3 times worse than the CMB CVL. We also study how the results depend on different properties of the surveys, such as survey volume, cell size, noise level, and inclusion of the cross-correlation with the galaxy density. Similarly to cross-correlating CIPs with CMB anisotropies as adopted in Ref. [12], we find that constraints for large future surveys are improved by a factor of about $\sim 2 - 3$ when cross-correlations are

included.

The BAO method could cross-check CMB measurements, which are contaminated by gravitational lensing signals, and could potentially help to remove biases due to CMB lensing [17, 18]. We also note that since we are only interested in studying the feasibility of the BAO method, we will use a simple and non-optimal estimator for these forecasts that could in principle be improved with an optimal estimator.

We note that the authors of Refs. [19, 20] also proposed to use BAOs to probe CIPs. They measure the luminosity-weighted galaxy density fluctuations and compare them to the unweighted ones, with the expectation that more stars are formed in regions of higher baryon content, indicative of CIPs. An advantage of this method is that it does not require dividing the survey into multiple patches. But two modeling systematics are present in this approach: How galaxy luminosity and star formation depend on the baryon fraction of the halo, and how the baryons in the halos trace the total baryon content of the Universe. Our method is different and does not rely on modeling these nonlinear processes because it only uses the spatial modulation of the BAO scale, which is imprinted in the density fluctuations right after recombination and is transferred to the clustering of galaxies. It would be interesting to combine both methods, especially for future large-volume surveys.

We begin in section II by providing background on CIPs in the curvaton model and describing their effects on BAO measurements. In section III, we describe a reconstruction procedure for measuring CIPs from modulated BAOs. In section IV we use Fisher matrix to forecast the sensitivity of a fiducial BAO survey to correlated CIPs and study the forecast dependencies on different survey parameters. We discuss the results in section V and conclude in section VI.

II. CIPS AND BAO MODULATION

In this section we review what CIPs are and describe how they modulate the BAO scale and the observable galaxy power spectrum.

A. CIPS

The primordial perturbations in the early Universe can be decomposed into two types. The adiabatic fluctuations ζ are fluctuations on constant density slicings, while the isocurvature fluctuations $S_{i\gamma}$ are the differences between the fractional fluctuations of number density n of species i and the photons γ

$$S_{i\gamma} = \frac{\delta n_i}{n_i} - \frac{\delta n_\gamma}{n_\gamma}. \quad (1)$$

Here $i \in \{b, c, \nu\}$ with b for baryons, c for cold dark matter and ν for neutrinos.

The compensated isocurvature mode Δ is a special combination of those perturbations for which the baryon number density fluctuations are exactly compensated by those of CDM

$$S_{b\gamma} = \Delta, \quad S_{c\gamma} = -\frac{\rho_b}{\rho_c} \Delta, \quad S_{\nu\gamma} = 0, \quad (2)$$

where ρ_i is the energy density of species i .

An early-universe model that generates perturbations of this form is the curvaton model. In this model, the curvaton is a spectator scalar field during inflation but later decays into other particles. When the Hubble scale becomes comparable to the mass of the curvaton, the field starts to oscillate and its energy density redshifts like that of matter in the radiation era after inflation. As a result, the curvaton could dominate the energy density of the Universe at its decay and give rise to most of the adiabatic perturbations in the Universe.

Depending on whether it decays into baryon number or CDM, isocurvature perturbations are produced and are correlated with the adiabatic fluctuations. In particular, different amount of CIPs can be produced, so one can use them to distinguish the decay scenarios of the curvaton model (see e.g. [12]).

In the case where the curvaton dominates the Universe's energy density completely at its decay, it would be responsible for generating all the adiabatic perturbations ζ , and any resulting CIPs would be fully correlated with ζ ,

$$\Delta = A\zeta, \quad (3)$$

where A determines the amplitude of the CIP. Our goal is to constrain this CIP amplitude A .

There are two decay scenarios that are particularly interesting because of their relatively large CIP amplitude: $A \approx 3\Omega_c/\Omega_b \approx 16$ if baryon number is produced by the curvaton decay, and CDM before the decay, and $A = -3$ if CDM is produced by curvaton decay, and baryon number before the decay.

B. BAO Modulation from CIPs

The CIP modes described above modulate the baryon-to-photon ratio spatially. As a result, the acoustic waves in the baryon-photon fluid generated by the adiabatic perturbations propagate with spatially varying sound speed before recombination. This modulation of the sound speed leads to a position-dependent BAO scale, which can be used to constrain the CIP amplitude as we are going to show.

How exactly do the CIPs modulate the BAO scale? For CIP modes with wavelength shorter than the sound horizon at recombination, the sound waves speed up and slow down over many wavelengths of the CIP modulation, leading to a negligible net effect at recombination [12]. But for CIP wavelengths longer than the sound horizon, sound waves propagate on a local background that

is slowly modulated. In this limit, the behavior of sub-horizon modes in the patch around a location \mathbf{X} can be approximated as if they were in a separate universe with perturbed cosmological parameters [7]

$$\delta\Omega_b = \bar{\Omega}_b \Delta(\mathbf{X}), \quad \delta\Omega_c = -\bar{\Omega}_b \Delta(\mathbf{X}), \quad (4)$$

where barred quantities denote values in the absence of CIPs. From now on we will use \mathbf{X}, \mathbf{K} to denote the long-wavelength CIP position and wavenumber, and use \mathbf{x}, \mathbf{k} to denote those of the short-wavelength perturbations on the modulated background.

As a consequence, a small, long-wavelength CIP perturbation $\Delta(\mathbf{X}) \ll 1$ changes the local mean baryon density as $\bar{\Omega}_b \rightarrow \bar{\Omega}_b(1 + \Delta(\mathbf{X}))$, and the sound speed of the baryon-photon fluid in the patch around \mathbf{X} changes accordingly as

$$\bar{c}_s \rightarrow c_s(\mathbf{X}) = \sqrt{\frac{1}{3} \left(1 + \frac{3\bar{\rho}_b(1 + \Delta(\mathbf{X}))}{4\bar{\rho}_\gamma} \right)^{-1}} \approx \tilde{\alpha}(\mathbf{X}) \bar{c}_s, \quad (5)$$

where we have Taylor expanded in $\Delta \ll 1$ to obtain

$$\tilde{\alpha}(\mathbf{X}) = 1 + \frac{\Delta(\mathbf{X})}{C}, \quad (6)$$

where

$$C \equiv -\frac{2(1+R)}{R}, \quad R \equiv \frac{3\bar{\rho}_b}{4\bar{\rho}_\gamma}; \quad (7)$$

in the fiducial cosmology described in section IV A, $C \approx -5.23$. The corresponding change in the local BAO scale is

$$\bar{r}_{\text{BAO}} \rightarrow \tilde{\alpha}(\mathbf{X}) \bar{r}_{\text{BAO}}. \quad (8)$$

Note that for the convenience of linearly relating $\tilde{\alpha}$ to Δ , we have defined it differently than the common convention, $\alpha = 1/\tilde{\alpha}$, which would have given $\bar{r}_{\text{BAO}} \rightarrow \bar{r}_{\text{BAO}}/\alpha(\mathbf{X})$ for Eq. 8 and $f_{\text{BAO}}(k/\alpha(\mathbf{X}))$ in Eq. 9 below.

C. Observable Imprints in the Galaxy Power Spectrum

Given the modulation of the BAO scale, the matter power spectrum in the patch around \mathbf{X} is

$$P_m(k, z; \mathbf{X}) = [1 + f_{\text{BAO}}(\tilde{\alpha}(\mathbf{X})k)] P_m^{\text{nw}}(k, z), \quad (9)$$

where we have split the matter power spectrum into a smooth nowiggle part without BAO wiggles, P_m^{nw} , and an oscillatory wiggle part, f_{BAO} , due to baryons [21–23]. Specifically, we use the wiggle power spectrum evaluated with CAMB [24] for the fiducial cosmology, and the corresponding nowiggle power spectrum from the appendix of Ref. [22]. This nowiggle spectrum is constructed to have

the correct low- k and high- k behavior, and the same density normalization σ_8 and velocity dispersion σ_v as the power spectrum with BAO wiggles¹.

In reality, we observe the galaxy rather than the matter power spectrum to measure the BAO scale from galaxy surveys, so additional effects from galaxy bias, redshift space distortion (RSD) and nonlinear evolution need to be taken into account. For extracting the BAO signal in a local patch around \mathbf{X} , we model the galaxy power spectrum as

$$P_{gg}(k, z; \mathbf{X}) = b_g^2(z)R(\mu, z)P_m^{\text{nl}}(k, z, \mathbf{X}), \quad (10)$$

where

$$R(\mu, z) = (1 + \beta(z)\mu^2)^2 \quad (11)$$

is the linear RSD effect with $\mu = \cos(\hat{\mathbf{k}} \cdot \hat{\mathbf{n}})$ being the angle between $\hat{\mathbf{k}}$ and the line of sight direction $\hat{\mathbf{n}}$. We use $\beta(z) = f(z)/b_g(z)$ with $f(z) \approx \Omega_m(z)^{0.55}$ and follow Ref. [25] to model the linear galaxy bias $b_g(z)$ for emission line galaxies as $b_g(z)\tilde{D}(z) = 0.84$, where $\tilde{D}(z=0) \equiv 1$ is the growth normalized at $z=0$.

We model the impact of nonlinear evolution as a Gaussian damping of the oscillations in Fourier space with widths Σ_{\perp} and Σ_{\parallel} respectively for modes perpendicular and along the line of sight [26, 27]. The matter power spectrum including these non-linear effects is therefore

$$P_m^{\text{nl}}(k, z; \mathbf{X}) = [1 + f_{\text{BAO}}(\tilde{\alpha}(\mathbf{X})k)\mathcal{D}(k, \mu, z)]P_m^{\text{nw}}(k, z), \quad (12)$$

where

$$\mathcal{D}(k, \mu, z) = e^{-\frac{1}{2}k^2[\Sigma_{\perp}^2 + \mu^2(\Sigma_{\parallel}^2 - \Sigma_{\perp}^2)]}. \quad (13)$$

Following Ref. [27], we use $\Sigma_{\perp} = \Sigma_0 D(z)$ and $\Sigma_{\parallel} = \Sigma_{\perp}(1 + f(z))$ where $\Sigma_0 = 12.4 \times (\sigma_8/0.9)h^{-1}\text{Mpc} = 11.4h^{-1}\text{Mpc}$, and where $D(z)$ is the growth normalized to $1/(1+z)$ during matter domination. Note that a fraction p_{recon} of the non-linear damping effect can be removed by applying BAO reconstruction which moves observed galaxy positions using the large-scale flow derived from the galaxy density on large scales [28]. To account for this possibility, we let $\Sigma_0 \rightarrow \Sigma_0(1 - p_{\text{recon}})$ where the default $p_{\text{recon}} = 0.5$.

CIP modes Δ therefore affect the galaxy power spectrum P_{gg} by changing $\tilde{\alpha}$, which changes the frequency of oscillations of f_{BAO} ,

$$\frac{\partial P_{gg}}{\partial \tilde{\alpha}} = b^2 R(\mu)\mathcal{D}(k, \mu, z)\frac{\partial f_{\text{BAO}}(k)}{\partial \ln k}P_m^{\text{nw}}(k, z), \quad (14)$$

so that

$$\frac{dP_{gg}}{d\Delta} = \frac{1}{C} \frac{\partial P_{gg}}{\partial \tilde{\alpha}} \quad (15)$$

in linear theory. Note that there are additional signatures of the CIP, corresponding to additional terms in Eq. 15 of the form $\sum_i (\partial P_{gg}/\partial p_i)(\partial p_i/\partial \Delta)$ where p_i 's are other parameters that depend on Δ including $p_i = \Delta$ itself. We choose to ignore those terms as well as derivative terms that arise from varying components in Eq. 12 other than f_{BAO} , as we expect Eq. 15 to be the dominant and cleanest signal for now. For example, the broadband shape of the power spectrum at high k also depends on the baryon fraction (e.g. it is suppressed for high baryon density) and would be modulated by the CIP mode as a function of \mathbf{X} . However, modeling the response of the high- k broadband to CIP modes and accounting for observational systematics that could also modulate the broadband power seems challenging. The difficulty of modeling the broadband can be seen from the fact that current BAO analyses compute the broadband using a polynomial template [29–31]. It could be that the modeling and understanding of systematics will improve in the future, and one could draw additional information from these other terms. But we will choose to ignore them for now, corresponding to a conservative forecast.

D. BAO Modulation from Other Effects

Up to now we have assumed that CIPs are the only effect modulating the BAO scale. In reality, other effects like matter density modes can also modulate the BAO scale. Over- (under-) densities of very long modes $\delta_m^L(\mathbf{K})$ that are longer than the BAO scale, $K \lesssim 1/r_{\text{BAO}}$, correspond to closed (open) separate universes. They contract (dilate) all scales inside that local patch by $\delta_m^L(\mathbf{K})/3$ [32]; due to this dilation the BAO scale measured in these patches is also modulated².

In contrast, displacements from modes smaller than the BAO scale, i.e. modes with $k \gtrsim 1/r_{\text{BAO}}$, perturb galaxies separated by the BAO scale by moving them towards or away from each other. This again stretches or shrinks the BAO scale locally depending on the realization of those modes, but the effect averages out and only smears out the BAO peak in the correlation function. Therefore, the very long modes (longer than BAO scale) can bias the CIP measurement, while the modes shorter than the BAO scale reduce the signal-to-noise ratio.

These effects can be avoided by applying BAO reconstruction [28] to the entire observed volume. The reconstruction removes the BAO modulation by moving galaxies by the Zel'dovich displacement computed from the observed modes that modulate the BAO scale, which removes the bias and smearing of the BAO feature. In

¹ This is achieved by constructing a family of spline curves for the nowiggle spectrum and forming a linear combination such that the integral constraints for σ_8 and σ_v are satisfied [22].

² When measuring the global power spectrum, this dilation together with nonlinear growth leads to a small shift of the inferred BAO scale, because subvolumes with a shrunk BAO scale contribute more to the global power spectrum than subvolumes where the BAO scale is stretched [32].

practice, this process of BAO reconstruction is not perfect, for example due to noise in the estimated Zel'dovich displacement, so that a residual BAO modulation may be left. This can lead to a residual bias on the CIP measurement which would have to be accounted for. We will ignore such a potential bias here and restrict ourselves to computing the statistical power with which the CIP amplitude can be constrained in absence of any bias.

An alternative procedure would be to apply BAO reconstruction locally in each subvolume and account for the dilation and growth from long modes by manually rescaling the density in each subvolume based on the observed long mode on scales larger than the subvolume. We describe this in Appendix A. The two procedures are in principle equivalent in linear theory, but may differ in practice, depending on the reconstruction algorithm used.

An additional modulation of the observed BAO peak in the correlation function is expected from gravitational lensing, which can magnify or shrink the physical BAO scale depending on the structures along the line of sight. This leads to a damping of the BAO feature in the global correlation function at the few percent level at $z = 2$ and less at lower redshift [33]. Compared to the modulation from bulk flows and the induced damping of the global BAO feature, this BAO lensing modulation is therefore subdominant, and we will ignore it for simplicity.³

In sum, the CIPs change the sound horizon as a function of position on the sky and induce a large-scale modulation of the measured BAO scale. The contaminating modulation due to long wavelength modes can be avoided by applying BAO reconstruction, and the modulation due to gravitational lensing is small so that it can be ignored.

III. MEASURING CIPS USING BAO

The modulation of the BAO by CIP modes described above suggests the following procedure for the measurement of CIPs from the modulation of the BAO on the sky.

1. Apply BAO reconstruction in the total survey volume.
2. Divide the total survey volume V_{tot} into N_{cell} cells with the same volume V_{cell} , centered at positions \mathbf{X}_i , where $i = 1, \dots, N_{\text{cell}}$.

³ If the BAO lensing modulation turns out to be sufficiently large to be worrying for CIP constraints, one could delens the BAO modulation, e.g., using external delensing based on observations of low-redshift galaxies, or using internal delensing that removes the lensing-induced non-Gaussianity of the observed BAO modulation.

3. Measure the BAO scale $\hat{\alpha}(\mathbf{X}_i)$ using galaxies in each cell \mathbf{X}_i , where $\hat{\alpha}$ is defined as the measured BAO scale \hat{r}_{BAO} with respect to the unmodulated BAO scale \bar{r}_{BAO} , approximated by that of the fiducial cosmology measured by *Planck*.
4. Estimate the CIP density perturbation $\hat{\Delta}(\mathbf{X}_i)$ from $\hat{\alpha}(\mathbf{X}_i)$.
5. Compute the CIP auto-spectrum $\hat{P}_{\Delta\Delta}(K)$ and its cross-spectrum $\hat{P}_{g\Delta}(K)$ with the long-wavelength galaxy density $\delta_g^L(\mathbf{K})$. Use them to measure the CIP amplitude A in the curvaton model.

In step 1, we apply BAO reconstruction to the entire survey to avoid contamination from BAO scale modulations induced by dilation and bulk flows, as discussed in section IID.

In step 2, the survey volume needs to be divided into sufficiently many N_{cell} cells for the Fourier transform of the later steps to make sense (i.e., to see a modulation). At the same time, we cannot use arbitrarily many cells because each cell must be larger than the BAO scale so that we can detect a BAO signal in each cell. We will explore the effect of cell size in Sec. IV B.

Once the BAO map $\hat{\alpha}(\mathbf{X}_i)$ is obtained in step 3, the CIP estimator is computed in step 4 as

$$\hat{\Delta}(\mathbf{X}_i) = C \left(\hat{\alpha}(\mathbf{X}_i) - 1 \right), \quad (16)$$

where the unmodulated BAO scale \bar{r}_{BAO} can be taken as that measured by the CMB. Taking the power spectrum in step 5 implicitly involves taking the Fourier transform of $\hat{\Delta}(\mathbf{X}_i)$. Since the true CIP map is multiplied by the survey window function in the measurement, the measured $\hat{\Delta}(\mathbf{K})$ is the truth $\Delta(\mathbf{K})$ convolved with the Fourier transform of the window function plus noise.

The CIP measurement noise is directly related to the BAO error per cell. Assuming uncorrelated and homogeneous noise, we can characterize the noise by its power spectrum alone

$$N_{\Delta\Delta}(K) = V_{\text{cell}} \sigma_{\Delta}^2 [W(KR)]^{-2}. \quad (17)$$

The low- K amplitude of the noise is proportional to the cell volume V_{cell} multiplied by $\sigma_{\Delta}^2 = C^2 \sigma_{\hat{\alpha}}^2$, where $\sigma_{\hat{\alpha}}$ is the BAO measurement error in a cell. The product $V_{\text{cell}} \sigma_{\Delta}^2$ is approximately independent of cell size as long as there are enough cells and they are larger than the BAO scale.

The noise increases rapidly at high K because we cannot probe fluctuations smaller than the cell size, so in principle one would low-pass filter the $\hat{\Delta}(\mathbf{X}_i)$ map before applying the Fourier transform in order to avoid aliasing. We model the low-pass filter as a convolution of the true map with a tophat window function with radius $R = (4\pi/3)^{1/3} r_{\text{cell}}$ such that the volume of the sphere is the same as the cell volume, where

$$r_{\text{cell}} = V_{\text{cell}}^{1/3}. \quad (18)$$

This filtering corresponds to multiplying the signal in Fourier space by

$$W(KR) = \frac{3}{(KR)^3} [\sin(KR) - KR \cos(KR)]. \quad (19)$$

We checked that our results do not change significantly if we use a Gaussian convolution instead, where $N_{\Delta\Delta}(K) = V_{\text{cell}} \sigma_{\Delta}^2 e^{(K/(K_{\text{Nyquist}}))^2}$ such that frequencies above the Nyquist frequency $K_{\text{Nyquist}} = \pi/r_{\text{cell}}$ are suppressed.

The procedure outlined above is effectively measuring a position-dependent galaxy correlation-function where the BAO peak shifts depending on the position on the sky. More optimally, a quadratic estimator $\hat{\Delta}^{\text{mv}}(K)$ can be constructed using off-diagonal correlations of the galaxy densities with weights that minimize the estimator variance. By definition the minimum variance estimator would have better noise properties than the one we use in this paper. Computing such an optimal estimator is however beyond the scope of this work. We will leave it for future work and focus on first assessing the method using the simpler estimator described above.

Since we consider curvaton CIPs, which are fully correlated with the adiabatic fluctuations ζ , we can measure cross-correlations of Δ with other probes of ζ in step 5 to improve the sensitivity in the low signal-to-noise regime. Specifically, we consider the cross-correlation of Δ with the galaxy density which traces ζ .

For the CMB, the authors of Ref. [12] found that including cross-correlations with CMB temperature and polarization anisotropies leads to an improvement by a factor of $\sim 2 - 3$ in σ_A for CIP measurements using quadratic estimators of the CMB. We shall see in Sec. IV that a similar improvement applies when we include cross-correlations with galaxies in the BAO modulation method.

IV. FISHER FORECAST

In this section, we describe how we forecast the sensitivity of the BAO method to the amplitude of curvaton CIP modes assuming a fiducial survey (section IV A), and explore how the results depend on various aspects of the data and assumptions (section IV B).

A. Fiducial Survey

We will use as observables the auto- and cross-spectra of the measured CIP mode Δ and the galaxy density on large scales

$$P_{XY}(K, z) = T_X(K, z) T_Y(K, z) P_{\zeta\zeta}(K), \quad (20)$$

where

$$P_{\zeta\zeta}(K) = 2\pi^2 K^{-3} A_s(K/K_0)^{n_s-1}, \quad (21)$$

with $X, Y \in \{g, \Delta\}$ where the CIP transfer function is

$$T_{\Delta}(K) = A, \quad (22)$$

and the galaxy transfer function on scales larger than the BAO scale is

$$T_g^L(K, z) = b_g(z) T_m^{\text{nw}}(K, z). \quad (23)$$

Note that we ignore RSD effects and use $T_m^{\text{nw}} \equiv \sqrt{P_m^{\text{nw}}/P_{\zeta\zeta}}$ from the no-wiggle matter power spectrum because we do not expect BAO wiggles on these scales. Recall that we only need to cross-correlate galaxy density modes on the same scales as the CIPs, and these are larger than the BAO scale by construction (one can only measure the modulation of the BAO scale on scales larger than the BAO scale).

We use the Fisher information matrix F_{ij} to predict the inverse covariance matrix of a set of parameters p_i that determine the reconstructed CIP field and the galaxy number density field. Under the assumption of Gaussian statistics for these fields, it can be approximated as

$$F_{ij} = \int_{K_{\text{min}}}^{K_{\text{max}}} dK \frac{4\pi K^2 V_{\text{survey}}}{(2\pi)^3} \int_{-1}^1 \frac{d\mu}{2} \times \sum_{a,b} \frac{\partial \tilde{P}_a(K, z_c)}{\partial p_i} [\mathcal{C}^{-1}(K, z_c)]_{ab} \frac{\partial \tilde{P}_b(K, z_c)}{\partial p_j}, \quad (24)$$

where we assume a central redshift z_c at which the power spectrum, noise and covariance matrix are calculated (we will suppress the z_c dependence from now on). The covariance matrix is

$$\mathcal{C}_{XX',YY'} = \tilde{P}_{XY} \tilde{P}_{X'Y'} + \tilde{P}_{XY'} \tilde{P}_{X'Y}, \quad (25)$$

where

$$\tilde{P}_a(K) = P_a(K) * W(K) + N_a(K) \quad (26)$$

are the observed power spectra for $a \in \{\Delta\Delta, g\Delta, gg\}$ where $P_a(K)$ are defined in Eq. 20, $N_{\Delta\Delta}$ is defined in Eq. (17), $N_{gg} = 1/\bar{n}$ is the usual Poisson noise, we assume $N_{g\Delta} = 0$ on these large scales, and $W(K)$ is the survey window function in Fourier space. We set $\beta = 0$ in the modeling of P_{gg} for the CIP Fisher matrix, but will include RSD when estimating the error of the BAO scale per cell.

The effect of having a finite-sized survey is that the neighboring K modes in the power spectrum are convolved by a window function with the width roughly equal to $dK \approx K_{\text{min}}$. We therefore approximate the Fisher matrix above as a discrete sum with bin-size $dK = K_{\text{min}}$ so that different K -bins are roughly uncorrelated and the covariance matrix is diagonal in K -space. For the wavenumber range we choose $K_{\text{min}} = 2\pi V_{\text{survey}}^{-1/3}$ which is limited by the survey volume, and $K_{\text{max}} = \pi V_{\text{cell}}^{-1/3}$, the Nyquist frequency set by the sampling frequency due to the finite cell size.

We are interested in the CIP amplitude parameter $p_i = A$ defined in Eq. 3. The information on A comes from the auto-spectrum of the reconstructed CIP, $P_{\Delta\Delta} \propto A^2$, and its cross-spectrum with the galaxy density $P_{g\Delta} \propto A$. Since the other cosmological parameters are well determined, and we do not expect them to be degenerate with A as they do not modulate the BAO scale, we will fix them to the fiducial LCDM values that we assume. The error on A is then

$$\sigma_A = \frac{1}{\sqrt{F_{AA}}}. \quad (27)$$

We compute this error as a function of the fiducial value of A and report the result when $A/\sigma_A = 2$, corresponding to the 2σ detection threshold.

The BAO error $\sigma_{\tilde{\alpha}}$ per cell that goes into the calculation of $N_{\Delta\Delta}$ is obtained with a similar Fisher matrix, but with $a = gg$ and $p_i = \tilde{\alpha}$ (the BAO scale relative to the fiducial value). In addition, the observed galaxy power spectrum includes redshift-space distortion effects as in Eqs. 10 and 11. For the wavenumber range, we use $k_{\min} = 2\pi V_{\text{cell}}^{-1/3}$ and $k_{\max} = 0.6 h \text{Mpc}^{-1}$. To use a diagonal covariance matrix, we again approximate the integral with a discrete sum of width $dk = k_{\min} = 2\pi V_{\text{cell}}^{-1/3}$. Since $\partial \tilde{P}_{gg}(k)/\partial \tilde{\alpha}$ has wiggles with period roughly $1/r_{\text{BAO}}$, it is no longer slowly varying inside each bin for small cell size close to the BAO scale $r_{\text{cell}} \sim r_{\text{BAO}}$. To deal with this, we do not use bin centers for the summand but approximate the observed power spectrum derivative by averaging

$$\frac{\partial P_{gg}(k, \mu)}{\partial \tilde{\alpha}} D(k, \mu) \quad (28)$$

within a k -bin. Note that the power spectrum itself is still slowly varying inside each bin, so it is safe to approximate the covariance matrix with its value at the bin-center. We have checked that using a no-wiggle power spectrum for the covariance computation does not change our results.

For surveys with large redshift range, the sensitivity of the BAO measurement would also vary with redshift, and the approximation of homogeneous CIP reconstruction noise would break down. In addition, the power spectrum and covariances would also evolve with redshift. However, we will take a fiducial survey with a sufficiently small redshift range and ignore this effect in the fiducial forecast. For simplicity we will also assume this when we extrapolate results to larger volumes in Sec. IV B.

Throughout, we use a flat Λ CDM cosmology consistent with *Planck* 2015 results [34], with $\Omega_b h^2 = 0.0223$ and $\Omega_c h^2 = 0.1188$, the adiabatic scalar power spectrum with amplitude $A_s = 2.207 \times 10^{-9}$ at pivot wavenumber $k_0 = 0.05 \text{Mpc}^{-1}$, spectral index $n_s = 0.9645$, Hubble constant $h = 0.6711$, resulting in $\sigma_8 = 0.8249$ assuming one massive neutrino with $m_\nu = 0.06 \text{eV}$.

We run the forecast described above for a fiducial galaxy survey with sky fraction $f_{\text{sky}} = 0.5$, signal-to-noise ratio $\bar{n}P_{0.2} = 5$ where $P_{0.2}$ is evaluated at

$k = 0.2 h \text{Mpc}^{-1}$ and $\mu = 0$, reconstruction fraction $p_{\text{recon}} = 0.5$ and a single redshift bin from $z = 0.75$ to 1.25 with bin center $z_c = 1.0$, corresponding to a survey volume of $V_{\text{survey}} = 27.8 (h^{-1} \text{Gpc})^3$ which we divide into $N_{\text{cell}} = 1000$ cells. The number density \bar{n} of this fiducial survey is similar to what would be expected from next-generation galaxy surveys with flux limit of $1 - 2 \times 10^{-16} \text{erg s}^{-1} \text{cm}^{-2}$ [35], similar to Euclid [36] at $z \sim 1$, but with a smaller redshift range and larger f_{sky} . We find that the BAO method applied to this fiducial survey can detect the CIP amplitude with 2σ if $A = 484$ (see Table I). This is similar to WMAP's sensitivity to CIPs without cross-correlations [4, 12] (see Appendix B for the conversion of WMAP results and the subtleties involved).

TABLE I. $2\sigma_A$ detection threshold for the correlated CIP amplitude A given the auto-spectrum of the CIP modes reconstructed from BAO modulation, the auto-spectrum of the galaxy density on large scales and their cross-spectrum, for the fiducial next-generation survey and a cosmic-variance-limited survey of emission-line galaxies (see section IV A for assumptions). The fiducial survey performs similarly to current WMAP [4] (see Appendix B), whereas the CVL result falls between a stage-3 ($2\sigma_A \sim 60$) and stage-4 ($2\sigma_A \sim 15$) CMB experiment if the CMB lensing bias effect is included.

Survey	$2\sigma_A$
Fiducial	484
CVL ($z = 7$)	30

B. Forecast Dependencies

How can the fiducial BAO survey be improved to increase the sensitivity to CIPs? To examine this we vary various aspects of the data and forecast assumptions: The galaxy number density \bar{n} , the cell volume V_{cell} and the survey volume V_{survey} . Unless specified otherwise, we vary the assumptions one at a time from the fiducial choices described in the last section.

Galaxy number density – We begin by varying the galaxy number density. For a given reconstruction fraction, the CIP detection threshold is close to cosmic variance limited when the galaxy power signal-to-noise ratio exceeds $\bar{n}P \approx$ a few, so there is little gain with improving $\bar{n}P$ in that regime. But for $\bar{n}P$ smaller than a few, improving $\bar{n}P$ makes a bigger difference. This is especially true when cross-correlations are included as improving $\bar{n}P$ decreases the variance of the observed $P_{g\Delta}$ as well.

Cell volume – Next we study the dependence on the volume of the cells into which we divide the survey. We




FIG. 2. 2σ CIP detection threshold in the main panel and BAO S/N per cell $\tilde{\alpha}/\sigma_{\tilde{\alpha}}$ in the inset as a function of cell volume (bottom axis) and the cell side length r_{cell} in units of r_{BAO} (top axis). We vary $K_{\text{max}} = \pi V_{\text{cell}}^{-1/3}$ accordingly and fix all other parameters at their fiducial values (e.g. $V_{\text{survey}} = 27.8 (h^{-1}\text{Gpc})^3$). For these choices, there is a small optimal range of cell volume around side length $r_{\text{cell}} \sim 3 - 4 r_{\text{BAO}}$. Note that the fiducial cell volume $V_{\text{cell, fid}} = 1000$ (blue cross) is already close to optimal. The spikes at 1 and 2 r_{BAO} are not physical but artifacts of an approximation made in the forecast (see text).

vary $K_{\text{max}} = \pi V_{\text{cell}}^{-1/3}$ accordingly but fix all other parameters to their fiducial values. In Fig. 2 we plot the 2σ detection threshold vs cell volume on the bottom axis, and on the top axis, the cell side length r_{cell} in units of r_{BAO} . If the cells are too small, there are not enough galaxy pairs separated by the BAO scale in each cell to make a high S/N measurement of the BAO scale: The inset shows that the BAO S/N per cell $\tilde{\alpha}/\sigma_{\tilde{\alpha}}$ drops below 3 for $r_{\text{cell}} \lesssim 3 r_{\text{BAO}}$. On the other hand, if the cells are too big and approach the survey volume, only a small range of CIP wavenumbers is accessible, and the CIP detectability also degrades. This means that there is an optimal regime in the middle that is best for CIP detectability. For the fiducial choices used, the optimal cell size is around $r_{\text{cell}} \sim 3 - 4 r_{\text{BAO}}$, and our fiducial choice of $V_{\text{cell}} = 0.0278 (h^{-1}\text{Gpc})^3 \approx (3 r_{\text{BAO}})^3$ given $N_{\text{cell}} = 1000$ is nearly optimal. But this would change depending on survey properties. In general, with larger surveys comes more freedom to choose bigger cells instead of smaller but more cells without sacrificing CIP detectability; but for smaller surveys, such flexibility would be limited.

Note that the spikes in Fig. 2 are not physical, but are artifacts of the approximations used for obtaining uncorrelated k -bins of the galaxy power spectrum from which the BAO is measured. These spikes occur for $r_{\text{cell}} = r_{\text{BAO}}$ and $2 r_{\text{BAO}}$, where k -bins are spaced by $2\pi/r_{\text{BAO}}$ or π/r_{BAO} and exactly one or half an oscillation of the signal $\partial P(k)/\partial \tilde{\alpha}$ averages out. We

see this more clearly in the inset where the BAO S/N per cell $\tilde{\alpha}/\sigma_{\tilde{\alpha}}$ drops sharply for the same cell sizes. This is only an artificial consequence of using bins with $dk = k_{\text{min}} = 2\pi(V_{\text{cell}})^{-1/3}$ and the averaged power spectrum derivative for each bin. In reality, the BAO detectability is a smooth function of the cell size as the power spectrum is convolved with the cell window function.

Survey volume – Next we examine the dependence of $2\sigma_A$ on the survey volume. We vary correspondingly $dK = K_{\text{min}} = 2\pi(V_{\text{survey}})^{-1/3}$, but keep the cell volume and K_{max} fixed. The result is shown in Fig. 3. In addition, for the CVL survey, instead of using the fiducial cell volume, we fix it at a slightly more optimal choice of $V_{\text{cell}} = (4 r_{\text{BAO}})^3$. Note also that although the volume is scaled up, all quantities in the Fisher matrix are still being calculated at fixed redshift $z_c = 1$. A few subtleties are involved with this assumption. For a realistic survey with fixed flux limit, one expects the galaxy number density to decrease with redshift, increasing the Poisson noise. On the other hand, the galaxy bias also increases with redshift and enhances the signal. Finally, the percentage of BAO reconstruction achievable grows closer to 100% at higher redshift since structure formation is less nonlinear there.

None of these effects matter for the CVL survey, which has zero Poisson noise and 100% reconstruction by definition. Furthermore, for cross-correlation dominated results the time evolution of the galaxy bias in the signal and the covariance pieces would cancel each other in absence of Poisson noise. Therefore we do not take into account any redshift evolution when we compute the CVL results, which is our main interest here. But we do show the same calculation for the fiducial survey as well, which is not meant to be accurate but only provides a rough comparison with CVL.

In Fig. 3, we show the 2σ detection threshold for the CVL (black dashed) and fiducial (blue solid) experiments using the auto-correlation only (thin lines) and including cross-correlations (thick lines). With the cross-correlation included, the fiducial experiment performs nearly as well as the CVL experiment with the auto-correlation alone. Moreover, we find that most of the improvement from the fiducial to the CVL curve at the same volume comes from improving the BAO reconstruction fraction from 0.5 to 1. So for a survey with fixed volume and $\bar{n}P$, its ability to detect CIPs can still improve with improved BAO reconstruction algorithms (e.g. [37–43]).

While the sensitivity to CIPs keeps improving with larger volume, there are actually no galaxies during the dark ages, so one cannot in reality use all the volume up to the surface of last scattering. For $z_{\text{max}} \sim 7$, corresponding to $V_{\text{survey}} \sim 900 (h^{-1}\text{Gpc})^3$, the BAO CVL with a slightly more optimal cell size $r_{\text{cell}} = 4 r_{\text{BAO}}$ yields $2\sigma_A = 30$. This sensitivity is between stage-3 and stage-4 CMB experiments which would have $2\sigma_A \sim 60$ and




FIG. 3. 2σ CIP detection threshold as a function of survey volume V_{survey} for fixed $V_{\text{cell}} = 0.0278 (h^{-1}\text{Gpc})^3$. CIP auto-correlation only results are shown (thin lines) as well as the total results including the cross-correlation with galaxies (thick lines), for both the CVL (black dashed) and the fiducial survey (blue solid). Note that the CVL survey has a larger cell volume with side $r_{\text{cell}} = 4r_{\text{BAO}}$. For reference, we also show the CMB CVL for $l_{\text{max}} = 2500$ (red dotted) which includes the degradation from lensing bias [18]. Note that the CMB result is for an optimal estimator while the BAO is not. Compared to the CMB, the BAO CVL experiment with $z_{\text{max}} \sim 7$ is about 3 times worse. For reference, the volumes for $z_{\text{max}} = 2, 7, 150$ are $\{0.19, 0.9, 3.0\} \times 10^3 (h^{-1}\text{Gpc})^3$ respectively for $f_{\text{sky}} = 1$, and the right edge is at $z_{\text{max}} = 1100$ in the fiducial cosmology.

$2\sigma_A \sim 15$ respectively assuming they suffer from a similar (if not worse) factor of 1.5 degradation than the CVL experiment due to lensing bias effects. Compared to the CVL CMB experiment, the galaxy BAO CVL at $z \lesssim 7$ is a factor of 2.5 (3.6) worse than the CVL CMB experiment with (without) degradation from lensing bias.

To understand the relative strength of the CMB and BAO methods in their current forms it is useful to make an order of magnitude estimate. As an indication of CIP detectability we consider the S/N of the acoustic scale per patch and the number of patches on the last scattering surface and in the volume up to $z_{\text{max}} = 7$. For the CMB, the fractional uncertainty of the angular size of the sound horizon θ_* is $\Delta\theta_*/\theta_* = 4.4 \times 10^{-4}$ in the *Planck* 2015 measurement, which would be improved by a factor of 7 by a full-sky CVL experiment [44]. Using patches large enough to include the first acoustic peak, say square patches with side length $2.7r_{\text{BAO}}$, the S/N per patch would be ~ 133 using the naive $1/\sqrt{N}$ scaling with $N = 1.4 \times 10^4$. Now a similar number of cubic cells with side length $4r_{\text{BAO}}$ exists in the volume between $z = 0$ and 7. For emission line galaxies, the CVL BAO S/N would be about ≈ 44 per patch, which is a factor of ~ 3 lower than the CMB, so we do expect slightly better performance from the CMB. We keep in mind that this is only an order of magnitude estimate, as the CIP

detectability also depends on other factors such as the optimization of the estimator, cross-correlation coefficient between the spectra, etc.

Despite this expected advantage of the CMB method, the BAO method that we assumed for the CIP forecasts above can be improved in a number of ways. First, similarly to the CMB forecast in Fig. 3, the BAO method could be modified to use an optimal quadratic estimator. This would involve deriving weights that minimize the variance of a quadratic estimator built from the correlation between different Fourier modes of the galaxy density. Another possible improvement unique to the BAO side is to use multiple types of galaxies as biased tracers of the adiabatic perturbations to cancel some sample variance; there are also other BAO probes to take advantage of such as those from Lyman- α measurements. Furthermore, 21cm experiments have the potential to offer competitive BAO measurements in the future, and could extend the range of galaxy surveys to higher redshift [45]. The 21cm signal is there, in principle, all the way until $z \sim 150$, though foregrounds could be more challenging at higher redshifts. The 21cm observations could serve as additional BAO measurements, or as a probe of the adiabatic perturbations. Finally, the forecast could be improved using CIP information from the modulation of the broadband power spectrum.

C. Impact of Including Cross-Correlation

In this section we examine in more detail how the error on the CIP amplitude σ_A benefits from including the cross-correlation between the reconstructed Δ field and the galaxy density, for different values of the CIP amplitude A .

We show in Fig. 4 the results from the CIP auto-spectrum alone, the cross-spectrum alone, and their combination. As was the case with CIP-cross-CMB anisotropies [12], the cross-spectrum with galaxies here entirely dominates the total result at low signal-to-noise ratio A/σ_A . For large signal-to-noise the inclusion of the cross-spectrum helps to eliminate sample variance from the underlying random Gaussian field ζ .

For small CIP amplitude, the constraint from the auto-spectra alone has a subtlety [12]. Since $P_{\Delta\Delta} \propto A^2$, the Fisher error $\sigma_A \propto A^{-1}$ diverges as $A \rightarrow 0$ whereas the Fisher error on A^2 does not, $\sigma_{A^2} = 2|A|\sigma_A$. Of course an upper limit on A^2 would imply an upper limit on A

$$\tilde{\sigma}_A = \sqrt{A^2 + \sigma_{A^2}} - |A|. \quad (29)$$

We show in Fig. 4 (following Ref. [12]) that the two auto-spectrum errors σ_A (orange dashed) and $\tilde{\sigma}_A$ (black dot-dashed) agree well for $A \gtrsim 2\sigma_A$ (unshaded region). Because in this regime we can meaningfully compare the Fisher errors, we choose to report as our main result the 2σ detection threshold.

The behavior of the amplitude error and impact of including the cross-correlation can be understood




FIG. 4. Fisher error σ_A as a function of the true value of the CIP amplitude A using the CIP auto- ($\Delta\Delta$) and cross- ($g\Delta$) spectrum, and their combination for the fiducial survey. Also shown is $\tilde{\sigma}_A$ (Eq. 29), the upper limit on A implied by an upper limit on A^2 . The two auto-spectrum results σ_A and $\tilde{\sigma}_A$ agree well for $A \gtrsim 2\sigma_A$, which defines the regime where the Fisher results can be interpreted meaningfully. The cross-spectrum increasingly dominates the total result at low A/σ_A while also contributing at high signal-to-noise ratio by reducing sample variance.

analytically in the low and high signal-to-noise regime.

The low signal-to-noise regime – We can approximate the auto-spectrum error in the noise-dominated regime $P_{\Delta\Delta}(K) \ll N_{\Delta\Delta}(K)$ as

$$\sigma_A^{-2} \Big|_{\Delta\Delta} = \sum_K \frac{2N_{\text{modes}}}{A^2} \left(\frac{P_{\Delta\Delta}(K)}{N_{\Delta\Delta}(K)} \right)^2 \quad (30)$$

and the cross-spectrum errors as

$$\sigma_A^{-2} \Big|_{g\Delta} = \sum_K \frac{N_{\text{modes}}}{A^2} \frac{P_{\Delta\Delta}(K)}{N_{\Delta\Delta}(K)} \frac{P_{gg}}{P_{gg} + \frac{1}{\bar{n}}} R_{g\Delta}^2, \quad (31)$$

where

$$N_{\text{modes}} = V_{\text{survey}} 4\pi K^2 dK / (2\pi)^3. \quad (32)$$

Let us estimate that the $A = 2\sigma_A$ detection threshold occurs for the auto-spectrum at

$$\frac{P_{\Delta\Delta}(K)}{N_{\Delta\Delta}(K)} \sim \left(\frac{K_{\min}}{K_{\Delta}} \right)^{3/2}, \quad (33)$$

where K_{Δ} be the representative wavenumber when the Fisher sum reaches half its total value. Then, at this level of signal, the constraint from the cross-spectrum alone is better than that from the auto-spectrum alone by a factor of

$$\frac{\sigma_{g\Delta}^{-2}}{\sigma_{\Delta\Delta}^{-2}} \sim \frac{1}{1 + (\bar{n}P)^{-1}} \left(\frac{K_{\Delta}}{K_{\min}} \right)^{3/2} R_{g\Delta}^2 \quad (34)$$

where the signal correlation

$$R_{g\Delta} \equiv P_{g\Delta} / \sqrt{(P_{gg}P_{\Delta\Delta})} = 1 \quad (35)$$

drops out.

Similarly to the CMB, the improvement from the cross-spectrum only depends on a few survey or analysis choices. As can be seen in Eq. 34, the improvement is larger for a survey with smaller shot noise. In addition, a larger cell volume means that a smaller range of CIP wavenumbers is accessible, so K_{Δ} is lower, and the improvement is smaller. For $r_{\text{cell, fid}} \approx 3r_{\text{BAO}}$ used for the fiducial survey in Fig. 3 (we don't expect to use much smaller cell size than this), the improvement is about a factor of 2–3. Finally, the ratio in Eq. 34 scales as $V_{\text{survey}}^{1/2}$ (or $f_{\text{sky}}^{1/2}$ for CMB) – the total number of modes available in the survey. So a survey with a larger volume would see a larger improvement, which is also reflected in Fig. 3.

The high signal-to-noise regime – In the high signal-to-noise regime, we can eliminate sample variance by jointly analyzing the auto- and the cross-spectra, taking their covariance into account. The high level of correlation between the various spectra is what makes the total results of Fig. 4 do much better at high A than would a simple inverse squared sum. Interestingly, this effect seems to be more pronounced for the galaxy probes than for the CMB (see Fig. 6 of Ref. [12]), possibly because in absence of shot noise the galaxy density δ_g is fully correlated with the adiabatic perturbations ζ as well as with the CIP fluctuations $R_{g\Delta} = 1$ (at least on the large scales that we consider), whereas the CMB anisotropies are not fully correlated with ζ nor with Δ because of the different projection effects from a 3-dimensional field to the last-scattering surface. The cross-correlation coefficient between Δ and CMB fields $Y \in \{T, E\}$

$$R_L^{Y\Delta} \equiv C_L^{Y\Delta} / \sqrt{(C_L^{YY}C_L^{\Delta\Delta})} \quad (36)$$

oscillates between -1 and 1 as a function of L (see Fig. 1 of Ref. [12]), and as a result, averages to a factor less than one when averaged over L . However, even in the fully correlated case for the BAO, we still do not get infinite information as the result is still limited by the finite $\bar{n}P$ at high A and by the large CIP measurement noise at low A .

V. DISCUSSION

A. Comments on H_0 tension

If CIPs exist, they can bias our current cosmological measurements that are made assuming no CIPs. In particular, given a BAO survey, CIPs on scales larger than the survey act as a super-sample fluctuation of the baryon density that could lead us to measure a different BAO

scale than expected on average. A natural question to ask is: How much could our current BAO results be affected given the CIPs allowed by current data?

Recently, the *Planck* team [3] used the smoothing of CMB peaks and the lensing potential power spectra to obtain $\Delta_{\text{rms}}^2 = 0.0037_{-0.0021}^{+0.0016}$ for scale-invariant CIPs⁴, ruling out $\Delta_{\text{rms}}^2 \gtrsim 0.0069$ at 95% confidence level. Since CMB measurements are affected by CIP fluctuations on the surface of last scattering (SLS), this measured $\Delta_{\text{rms,CMB}}^2 = \langle |\Delta(\hat{n})|^2 \rangle$ receives contributions from scales above the sound horizon up to the size of the SLS.

Now only CIPs on scales larger than the volume of the galaxy survey would contribute to the supersample effect. Let us assume that the CMB-measured r.m.s. fluctuations are similar to those inside the 3D sphere enclosed by the SLS $\Delta_{\text{rms,CMB}}^2 \approx \Delta_{\text{rms,3D}}^2$. This assumption is reasonable since they share the same range of scales

$$\Delta_{\text{rms,3D}}^2 \propto \int_{k_{\text{min}}}^{k_{\text{max}}} k^2 P_{\Delta\Delta} dk \sim \log\left(\frac{k_{\text{max}}}{k_{\text{min}}}\right), \quad (37)$$

where k_{min} and k_{max} are set by the size of the SLS and the BAO scale respectively, and where in the last step we used a scale-invariant CIP power spectrum $P_{\Delta\Delta}(k) \propto k^{-3}$, which is approximate for the curvaton CIPs. It turns out that for BOSS or any next generation galaxy survey, the survey volume would be roughly at the logarithmic middle between k_{min} and k_{max} , so that about half the contribution to $\Delta_{\text{rms,3D}}^2$ would be from scales smaller than the survey volume, and about half for larger scales.

So taking half of the maximal value allowed by the Planck 95% bound $\Delta_{\text{rms}}^2 = \frac{1}{2}(0.0069)$ for the supersample effect, and assuming a 1σ fluctuation in the field $\Delta(\hat{x}) \approx \Delta_{\text{rms}} = 0.059$, we get a bias on the inferred BAO scale $\tilde{\alpha}$ of $|\Delta/C| \approx 1.1\%$. Given that current experiments such as BOSS [46] have errors at the level of 1%, this corresponds to a bias at the 1σ level. For a 2σ fluctuation $\Delta(\hat{x}) \approx 2\Delta_{\text{rms}}$, we would have a 2.2% bias in α . Suppose now that all the 2σ fluctuations were in the radial direction, then a 2.2% change in $\tilde{\alpha}$ means a 6.6% change in the inferred $H(z)$, because $H(z)$ is measured from the modes parallel to the line of sight (we have $\tilde{\alpha} = (\alpha_{\parallel}\alpha_{\perp}^2)^{-1/3}$ and $\delta H/H = -\delta\alpha_{\parallel}$, which implies $\delta H/H = 3\delta\tilde{\alpha}$ if the angular diameter distance $D_A \propto \alpha_{\perp}$ is kept fixed).

The current inverse distance ladder measurement is $H_0 = (67.9 \pm 0.8) \text{ km s}^{-1} \text{ Mpc}^{-1}$ [47], in tension with the local measurement $H_0 = (73.52 \pm 1.62) \text{ km s}^{-1} \text{ Mpc}^{-1}$ [48] at 3.1σ . A 6.6% change in H at $z = 0$ corresponds to roughly $\Delta H_0 \approx 4.5 \text{ km s}^{-1} \text{ Mpc}^{-1}$. So a 2σ radial fluctuation would reduce the tension to 0.6σ while a 1σ isotropic fluctuation would give 2.7σ with $\Delta H_0 = 0.7 \text{ km s}^{-1} \text{ Mpc}^{-1}$. Although the 2σ radial CIP

maximally allowed by *Planck* at 95% confidence level could reduce the inverse ladder tension significantly, it does not resolve the tension between local and CMB-only H_0 measurements⁵. Finally, we caution the interpretation of CIPs as solving the *Planck* lensing anomaly and possibly (partially) the H_0 tension since there is not yet a known physical mechanism to produce CIPs as high as that allowed by *Planck* ($\Delta_{\text{rms}}^2 = 0.0037$ is about $A^2 \approx (600)^2$). For comparison, the largest possible CIPs in the curvaton model is $A = 16$, giving a fractional bias in the BAO scale at the $\sim 3 \times 10^{-4}$ level for a 1σ fluctuation, which translates to only $\sim 0.1\%$ bias in H_0 in the isotropic case.

B. Measuring δ_m^L using BAO Modulation

The dilation due to long-wavelength matter fluctuations δ_m^L , which we considered as a potential bias of the CIP measurement from BAO, can also be used to measure the long-wavelength dark matter density from the modulation of the BAO scale that it induces — similarly to the CIP estimator from BAO. Since $r_{\text{BAO}} = (1 - \delta_m^L/3)\bar{r}_{\text{BAO}}$, the r.m.s. fluctuation of the BAO scale is $\langle |\delta_m^L|^2 \rangle^{1/2}/3 \simeq 2.5\% D(z)/D(0)$ if the r.m.s. of the long mode is evaluated using the variance at the BAO scale [32].

A BAO modulation at this level is likely too small to be measurable with a survey like BOSS, because the BAO scale in the entire survey is only measured at the 1% level. As a consequence, the super-sample variance of the BAO scale from the long-mode dilation is also relatively small for a BOSS-like survey [50].

However, future surveys that probe larger volumes may be sensitive to a BAO modulation at this level, especially when cross-correlating the measured BAO modulation with the long mode measured directly from the survey galaxy density. This could be useful to estimate the matter power spectrum on very large scales. Such a measurement could in principle be useful to constrain local primordial non-Gaussianity by avoiding potential observational systematics that can affect standard measurements of the large-scale power spectrum, such as star contamination (stars in the Galaxy that are mistakenly identified as extra-galactic galaxies, adding power on large angular scales corresponding to the angular extent of the Galaxy), or catastrophic redshift errors (low-redshift galaxies that are mistakenly attributed to high redshift, adding power on large

⁴ These are not correlated with adiabatic perturbations as in the curvaton case, but their constraint can still be applied to correlated CIPs (see Appendix B).

⁵ A similar reduction of tension would apply for the BAO plus Big Bang nucleosynthesis constraints (e.g. [49]), as long as the light element abundances are measured from a much larger volume than the BAO survey. This is interesting since the BAO + BBN results are sometimes used to check that BAO H_0 measurements made independent of *Planck* are consistent with the CMB H_0 .

angular scales). While potentially reducing the impact of systematics, we expect this method to have larger statistical uncertainty than standard measurements of the large-scale power spectrum because the long mode estimated from BAO modulation is rather noisy.

VI. SUMMARY

In this paper, we explored measuring CIPs from large-scale structure using the spatial modulation of the BAO scale that they induce: As CIPs modulate the baryon-to-photon ratio spatially, the sound speed of acoustic waves before recombination is modulated, leading to a modulation of the BAO scale as a function of position on the sky.

We forecasted the detectability of CIPs assuming a simple but non-optimal estimator: first divide the survey volume into many cells, then measure the local BAO scale deviation in each cell, and make a CIP map assuming a linear relationship. In the forecast, we assumed that the measurement noise is uncorrelated between cells, and that the other effects known to bias the BAO scale can be removed (e.g. dilation effects from long modes δ_m^L), either via BAO reconstruction, or via an alternative removal procedure described in Appendix A.

Using the Fisher matrix formalism, we first compute the BAO error per cell, then the resulting noise of the CIP modes measured from the modulation of the BAO across the cells, and from this the forecasted error of the CIP amplitude A in the curvaton model using the reconstructed CIP and galaxy density auto- and cross-spectra as observables. For a fiducial survey similar to next-generation galaxy surveys like Euclid with $V_{\text{survey}} = 27.8 (h^{-1}\text{Gpc})^3$ obtained from $z = 0.75$ to 1.25 , $f_{\text{sky}} = 0.5$, $\bar{n}P = 5$, and $N_{\text{cell}} = 1000$, we find a 2σ CIP detection threshold of 484, which is similar to WMAP constraints from the auto-correlation alone.

We explored how this result depends on various survey aspects (fixing all other parameters at the fiducial choices except specified otherwise). Varying the size of the cells, we find that there is a sweet spot for the cell size at $r_{\text{cell}} \approx 2 - 3r_{\text{BAO}}$. For smaller cells, the BAO S/N per cell becomes too small, whereas for larger cells one cannot fit enough cells in the survey volume to observe the BAO modulation across cells.

Given the fiducial survey volume, most of the fiducial settings are already close to optimal. When varying the survey volume and using a cosmic variance limited (CVL) survey, we find that it would be competitive with a stage 3 CMB experiment if one used emission line galaxies up to $z_{\text{max}} \sim 7$, giving $2\sigma_A = 30$. These results are compatible with what we would expect from a naive mode counting estimate. If a larger volume is used, say up to $z \sim 150$, then the BAO method in its current form would start to become similar to the CMB CVL. Of course, there are no galaxies up to those high redshifts, so an-

other tracer must be found to fully take advantage of the information in this large volume.

There is hope that 21cm experiments could – in addition to offering BAO measurements at low redshift $z < 7$ – offer a window into the $z > 7$ range where it would be hard to observe galaxies. These 21cm measurements could extend galaxy BAO measurements to higher redshifts, as well as serve as potential tracers of the adiabatic perturbations there, to allow cross-correlations even at high redshifts.

We found that using cross-correlations in the fiducial galaxy survey with $r_{\text{cell}} \approx 3r_{\text{BAO}}$ already improved results by about a factor of 2.1 (from $2\sigma = 1023$ to 484). We expect an improvement by a factor of about 2–3 for large future surveys, given the range of survey volume and cell size we explored, as the improvement decreases with cell size and increases with survey volume.

As in the CMB case, the ultimate limit to the method is the BAO scale. For CIPs much smaller than the sound horizon, the sound waves propagate at a speed that's modulated many times before recombination, so the net effect is too small to be observable. Therefore, the CMB and BAO methods presented here are ultimately limited by the same pre-recombination physics. However, they will see different systematics challenges from post-recombination processes.

On the BAO side, accurate measurements require an accurate BAO reconstruction and knowing the right galaxy bias in order to infer the large-scale δ_m^L from δ_g^L and correct for the effects of δ_m^L . Because the BAO and CMB suffer from different systematics, they can potentially be used as a cross-check for each other and/or be used in a joint analysis. For example, if CIPs are detected in the BAOs but not in the CMB, it could point us to new physics that modulate the BAO scale or to systematics in the BAO survey. On the other hand, if CIPs are detected in the CMB – say from the extra smoothing of peaks in the *Planck* CMB power spectra which is larger than lensing can account for (see the lensing anomaly [3]) – but they are not detected in the BAOs, then we may rule out CIPs as a potential cause of the anomaly and look for other effects.

The presence of CIPs – at least at the level allowed by the *Planck* power spectra measurement – could bias our current measurements of the BAO scale. For BOSS or any next generation survey, a 1σ (2σ) fluctuation can cause a 1.1% (2.2%) bias in α , which could lead to an up to three times larger bias in the $H(z)$ measurement if all fluctuations are radial. This type of supersample fluctuation could partially account for the current tension between the inverse-distance ladder and local measurement of H_0 . We caution that while this is appealing, we do not yet know any physical mechanism that could give rise to the high amount of CIPs corresponding to the maximal value allowed by *Planck* at 95% confidence level.

Finally, we note that unlike the CMB forecast we compare to, the simple method of measuring BAO scale

modulation on the sky is not optimized for measuring CIPs. To do so, a minimum-variance estimator $\hat{\Delta}^{\text{mv}}(\mathbf{K})$ should be built by optimally weighing the correlations between the different wavenumber pairs $k, k' \gg K$ of galaxy densities. This means that there is still room for improvement on the sensitivity reported here that one can explore in future work. Another way to improve upon the current forecast is to use multiple tracers of the adiabatic perturbations to cancel sample variance, for example by using different types of galaxies.

In conclusion, searching for a spatial modulation of the BAO scale represents a new tool to use the BAO feature as a window into primordial physics, specifically to constrain compensated isocurvature perturbations that modulate the sound speed of the baryon-photon fluid. This adds to the previous applications of using the BAO as a standard ruler [51] and as a tool to search for light particles [23, 52]. It also adds to and opens up new possibilities for synergies with other CIP probes such as CMB power spectra and quadratic estimators, luminosity-weighted galaxy correlation functions and 21cm measurements.

ACKNOWLEDGMENTS

We thank Olivier Doré, Marko Simonović and Matias Zaldarriaga for instructive discussions. We are very grateful to Olivier Doré, Daniel Grin, Marko Simonović and Tristan Smith for feedback on the manuscript, as well as Tzu-Ching Chang and Daniel Lenz for other useful feedback. We thank Zvonimir Vlah for providing the nowiggle power spectrum. Part of this work was done at Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2019. All rights reserved. MS gratefully acknowledges support from the Jeff Bezos and Corning Glass Works Fellowships at IAS as well as the National Science Foundation.

Appendix A: Explicit Removal of the Dilation Bias

We have argued in the main text that applying BAO reconstruction globally before dividing the survey into cells takes care of the potential biases of the CIP measurement due to the BAO modulation induced by dilation and bulk flows. In this appendix we present an alternative procedure where we treat separately the dilation effect induced by very long modes, i.e. modes longer than the cell size. This procedure works as follows.

1. Divide the total survey volume V_{tot} into N_{cell} cells with the same volume V_{cell} centered at positions \mathbf{X}_i , where $i = 1 \dots N_{\text{cell}}$.
2. Measure the BAO scale $\hat{\alpha}(\mathbf{X}_i)$ using galaxies in each cell \mathbf{X}_i , where $\hat{\alpha}$ is defined with respect to

fiducial cosmology measured by *Planck*.

3. Perform BAO reconstruction locally in each cell, using only large-scale modes within each cell.
4. Measure the galaxy density $\delta_g^L(\mathbf{K})$ on scales larger than the cell, but smaller than the survey size. Infer the long-wavelength matter density modes using $\hat{\delta}_m^L(\mathbf{K}) \approx \hat{\delta}_g^L(\mathbf{K})/b(z)$ where $b(z)$ is the linear galaxy bias at redshift z , and smoothing on the scale of the cell size is applied.
5. Obtain the CIP estimator using the Fourier transform of the BAO map $\hat{\alpha}(\mathbf{K})$, and remove the dilation bias due to $\delta_m^L(\mathbf{K})$ measured in step 4

$$\hat{\Delta}(\mathbf{K}) = C \left(\hat{\alpha}(\mathbf{K}) - 1 \right) + \frac{C \hat{\delta}_m^L(\mathbf{K})}{3}. \quad (\text{A1})$$

6. Compute the CIP auto-spectrum $\hat{P}_{\Delta\Delta}(K)$ and cross-spectrum $\hat{P}_{\delta\Delta}(K)$ with the long-wavelength galaxy density $\delta_g^L(\mathbf{K})$, and use them to measure the CIP amplitude A in the curvature model.

In step 3, the BAO reconstruction uses only modes inside each cell. This is sufficient to reduce most of the non-linear peak smearing because density modes larger than the BAO scale or cell size move two galaxies separated by r_{BAO} by the same amount, so they do not shift the BAO scale. In fact, reconstructing only modes on scales close to the BAO scale should be sufficient because the non-linear peak smearing, given by [53]

$$\Sigma^2 = \frac{1}{3} \int \frac{dq}{2\pi^2} [1 - j_0(qr_{\text{BAO}}) + 2j_2(qr_{\text{BAO}})] P(q), \quad (\text{A2})$$

where j_n is the n^{th} order spherical Bessel function, is dominated by bulk flows on scales around the BAO scale: for $q \ll 1/r_{\text{BAO}}$ the term in square brackets is suppressed, whereas for $q \gg 1/r_{\text{BAO}}$ the power spectrum term becomes suppressed.

Note that contributions to Σ only impact the uncertainty of the measured BAO scale without leading to a bias. There is an effect, however, that shifts the BAO scale from cell to cell that we propose to mitigate using steps 4 and 5. A long-wavelength mode $\delta_m^L(\mathbf{K})$ varying on scales larger than the cell size (the cells are assumed to be larger than the BAO scale), have effects indistinguishable (at first order in δ_m^L) to a separate universe with local curvature. In this locally curved universe, small scales inside the cell are contracted by a factor of

$$\frac{a_{\text{curved}}(\delta_m^L)}{a_0} = (1 + \delta_m^L)^{-1/3} \approx \left(1 - \frac{\delta_m^L}{3} \right), \quad (\text{A3})$$

shifting the BAO scale in linear theory as [32]

$$r_{\text{BAO}} \rightarrow r_{\text{BAO}} \left(1 - \frac{\delta_m^L}{3} \right). \quad (\text{A4})$$

In Ref. [32], a second effect biasing the global BAO scale measurement was considered: The enhanced growth of local perturbations in a closed universe and the conversion between local and global mean densities resulting in an enhanced contribution from regions of positive δ_m to the global correlation function. We argue, however, that this effect is not relevant for us because we do not combine the measurements in different cells to give one global BAO scale measurement. Furthermore, for δ_m^L on scales smaller than the cell size, we assume that the local BAO reconstruction would revert its effect.

Therefore, in the presence of large-scale CIPs $\Delta(\mathbf{X}_i)$ and matter density fluctuations $\delta_m^L(\mathbf{X}_i)$, the local BAO scale at \mathbf{X}_i is shifted as

$$\begin{aligned} r_{\text{BAO}}(\mathbf{X}_i) &\approx r_{\text{BAO}}^0 \left(1 - \frac{\delta_m^L(\mathbf{X}_i)}{3}\right) \left(1 + \frac{\Delta(\mathbf{X}_i)}{C}\right) \\ &\approx r_{\text{BAO}}^0 \left(1 - \frac{\delta_m^L(\mathbf{X}_i)}{3} + \frac{\Delta(\mathbf{X}_i)}{C}\right) \\ &= r_{\text{BAO}}^0 \left(1 + \frac{\Delta_{\text{wrong}}(\mathbf{X}_i)}{C}\right), \end{aligned} \quad (\text{A5})$$

where r_{BAO}^0 is the unmodulated BAO scale in absence of $\Delta(\mathbf{X}_i)$ and $\delta_m^L(\mathbf{X}_i)$.

If uncorrected, the dilation effect from a positive $\delta_m(\mathbf{X}_i)$ fluctuation would therefore contribute to a negative bias of the CIP measurement

$$\Delta_{\text{wrong}}(\mathbf{X}_i) = \Delta(\mathbf{X}_i) - C\delta_m^L(\mathbf{X}_i)/3, \quad (\text{A6})$$

or correspondingly in Fourier space (when the linear approximation holds)

$$\Delta_{\text{wrong}}(\mathbf{K}) = \Delta(\mathbf{K}) - C\delta_m^L(\mathbf{K})/3. \quad (\text{A7})$$

Note that the above holds for CIP scales larger than the cell size but smaller than the survey size $V_{\text{cell}}^{1/3} \lesssim 2\pi/K \lesssim V_{\text{survey}}^{1/3}$, for we cannot measure Δ nor δ_m^L on scales larger than the survey, and for scales smaller than the cell size, we assume that the BAO reconstruction can mitigate the bias.

In step 4, we measure $\hat{\delta}_m^L(\mathbf{X}_i) = \hat{\delta}_g^L(\mathbf{X}_i)/b(z)$ and use it in step 5 to correct the bias. Since the correction is realization-dependent, there is no cosmic variance that would contribute to a noise bias in the CIP power spectra. However, modeling uncertainties in $\delta_m^L(\mathbf{X}_i)$ and $b(z)$ could lead to a residual bias, but we expect this to be small given that the relevant scales are larger than the BAO scale, and linear theory holds well on these scales.

Finally, one might worry about the dilation bias from $\delta_m^L(\mathbf{X}_i)$ inside the cell when the BAO reconstruction is not perfect. We argue that since the bias would fluctuate many times inside the cell, most of its contributions cancel as the BAO scale measurement is an averaged quantity of the cell (just like the effect of CIP modes with many wavelengths inside the sound horizon).

Appendix B: CIP Conventions and Conversion Rules

There are a few different conventions used in the literature for scale-invariant CIPs. Here we briefly review these conventions and provide approximate rules to convert between these conventions and the correlated CIP amplitude A that we use in this paper.

CMB measurements probe the CIP field on the surface of last scattering

$$\Delta(\hat{n}) = \Delta(\mathbf{x} = \chi_* \hat{n}), \quad (\text{B1})$$

where χ_* is the comoving distance to the SLS. Assuming instantaneous recombination for the field Δ , its angular power spectrum is

$$C_L^{\Delta\Delta} = \int k^2 dk P_{\Delta\Delta}(k) j_L^2(k\chi_*). \quad (\text{B2})$$

For correlated CIPs (e.g. of curvaton origin) we used in this paper

$$P_{\Delta\Delta}^{\text{corr}}(k) = A^2 P_{\zeta\zeta}, \quad (\text{B3})$$

whereas for the scale-invariant CIP power spectrum, a common convention is

$$P_{\Delta\Delta}^{\text{SI}}(k) = A_{\text{CIP}} k^{-3}, \quad (\text{B4})$$

in which case we can approximate

$$C_L^{\Delta\Delta, \text{SI}} \approx \frac{A_{\text{CIP}}}{\pi L(L+1)}. \quad (\text{B5})$$

Note that the correlated CIPs are nearly scale-invariant because of the small tilt $n_s - 1$. We derive approximate conversions rules below by comparing a small range of $C_L^{\Delta\Delta, \text{corr}}$ at low L to $C_L^{\Delta\Delta, \text{SI}}$ such that for that range

$$\frac{A_{\text{CIP}}}{\pi L(L+1)} \approx A^2 \left(\frac{C_L^{\Delta\Delta}}{A^2} \right). \quad (\text{B6})$$

Taking the average over the range $2 \leq L \leq 37$ we get $\langle (C_L^{\Delta\Delta}/A^2) L(L+1) \rangle_{L \in [2, 37]} = 1.462 \times 10^{-8}$ and so

$$A_{\text{CIP}} \approx (4.594 \times 10^{-8}) A^2. \quad (\text{B7})$$

The *Planck* constraints so far have been focused on using the smoothing of the acoustic peaks and the contamination to the lensing potential power spectrum that scale-invariant CIPs would induce. These contributions are expressed in terms of $\Delta_{\text{rms}}(R)$ the root-mean-square CIP amplitude over some length-scale R :

$$\Delta_{\text{rms}}^2(R) = \frac{1}{2\pi^2} \int k^2 dk [3j_1(kR)/(kR)]^2 P_{\Delta\Delta}(k). \quad (\text{B8})$$

For the CMB calculations, the relevant scale is the size of the surface of last scattering R_{CMB}

$$\Delta_{\text{rms}}^2(R_{\text{CMB}}) \equiv \langle \Delta^2(\hat{n}) \rangle = \sum_{L=1}^{100} \frac{2L+1}{4\pi} C_L^{\Delta\Delta}. \quad (\text{B9})$$

Numerically evaluating the sum, we get

$$\Delta_{\text{rms}}^2(R_{\text{CMB}}) \approx 0.2377 A_{\text{CIP}}, \quad (\text{B10})$$

for a scale-invariant power spectrum.

With the caveat that CIP contributions to the CMB and lensing power spectra expressed in terms of Δ_{rms} may differ for scale-invariant and correlated CIPs, we can convert the central value of the latest *Planck* constraint $\Delta_{\text{rms}}^2 = 0.0037$ (or $A_{\text{CIP}} \approx 0.015$) to roughly $A^2 \approx (580)^2$

$$\Delta_{\text{rms}}(R_{\text{CMB}}) \approx (1.045 \times 10^{-4})A. \quad (\text{B11})$$

The WMAP constraints [4] on the other hand used a quadratic CMB temperature estimator to put an upper limit on the scale-invariant CIPs amplitude A_{CIP}/π as well as the individual multipoles of $C_L^{\Delta\Delta}$ without any model assumption. Note that a different usage of A was employed there, $A \equiv A_{\text{CIP}}/\pi$ (as well as in Ref. [54] where $A \equiv A_{\text{CIP}}$.) These usages of A are different from the notation in this paper and Refs. [12, 17, 18] to denote the amplitude of correlated CIPs.

Ref. [4] also stated their constraints in terms of Δ_{cl} , to facilitate comparison with cluster constraints from Ref. [13]. Here $\Delta_{\text{cl}} \equiv \Delta_{\text{rms}}^2(R_{\text{cl}})$, where R_{cl} is the mean separation between galaxy clusters. Taking $k_{\text{min}} \approx (10 \text{ Gpc})^{-1}$ and $R_{\text{cl}} \approx 10 \text{ Mpc}$ in Eq. B8, we have

$$\Delta_{\text{cl}}^2 \approx \frac{A_{\text{CIP}} \ln(1000)}{2\pi^2}. \quad (\text{B12})$$

Using Eq. B7 we get

$$\Delta_{\text{cl}} \approx (1.268 \times 10^{-4})A. \quad (\text{B13})$$

Using these conversion rules, the WMAP 95% upper bound of $A_{\text{CIP}}/\pi < 0.011$ (or $\Delta_{\text{cl}}^2 < 0.012$) corresponds to $A^2 < (870)^2$ for correlated CIPs. Note that this is a 2σ upper bound in A^2 : A 2σ detection in A^2 actually corresponds to a 4σ detection in A , so a 2σ detection in A would occur at a smaller A , probably around $2\sigma_A \sim 500$. This is why we say that the fiducial galaxy survey considered in this work with $2\sigma_A = 484$ is similar in sensitivity to the current WMAP constraint.

-
- [1] P. A. R. Ade *et al.* (Planck), (2015), arXiv:1502.02114.
[2] P. A. R. Ade *et al.* (Planck), (2015), arXiv:1502.01592.
[3] Y. Akrami *et al.* (Planck), (2018), arXiv:1807.06211 [astro-ph.CO].
[4] D. Grin, D. Hanson, G. Holder, O. Doré, and M. Kamionkowski, *Phys.Rev.* **D89**, 023006 (2014), arXiv:1306.4319.
[5] J. B. Muñoz, D. Grin, L. Dai, M. Kamionkowski, and E. D. Kovetz, *Phys. Rev.* **D93**, 043008 (2016), arXiv:1511.04441.
[6] T. L. Smith, J. B. Muoz, R. Smith, K. Yee, and D. Grin, *Phys. Rev.* **D96**, 083508 (2017), arXiv:1704.03461 [astro-ph.CO].
[7] D. Grin, O. Doré, and M. Kamionkowski, *Phys.Rev.* **D84**, 123003 (2011), arXiv:1107.5047.
[8] A. D. Linde, *JETP Lett.* **40**, 1333 (1984).
[9] A. D. Linde and V. F. Mukhanov, *Phys.Rev.* **D56**, 535 (1997), astro-ph/9610219.
[10] D. Langlois and A. Riazuelo, *Phys.Rev.* **D62**, 043504 (2000), astro-ph/9912497.
[11] D. H. Lyth, C. Ungarelli, and D. Wands, *Phys.Rev.* **D67**, 023503 (2003), astro-ph/0208055.
[12] C. He, D. Grin, and W. Hu, *Phys. Rev.* **D92**, 063018 (2015), arXiv:1505.00639.
[13] G. P. Holder, K. M. Nollett, and A. van Engelen, *Astrophys. J.* **716**, 907 (2010), arXiv:0907.3919.
[14] C. Gordon and A. Lewis, *Phys.Rev.* **D67**, 123513 (2003), astro-ph/0212248.
[15] C. Gordon and J. R. Pritchard, *Phys.Rev.* **D80**, 063535 (2009), arXiv:0907.5400.
[16] J. Valiviita, *JCAP* **1704**, 014 (2017), arXiv:1701.07039 [astro-ph.CO].
[17] C. H. Heinrich, D. Grin, and W. Hu, *Phys. Rev.* **D94**, 043534 (2016), arXiv:1605.08439 [astro-ph.CO].
[18] C. Heinrich, *Phys. Rev.* **D97**, 023513 (2018), arXiv:1708.01276 [astro-ph.CO].
[19] M. T. Soumagnac, R. Barkana, C. G. Sabiu, A. Loeb, A. J. Ross, F. B. Abdalla, S. T. Balan, and O. Lahav, *Phys. Rev. Lett.* **116**, 201302 (2016), arXiv:1602.01839 [astro-ph.CO].
[20] M. T. Soumagnac, C. G. Sabiu, R. Barkana, and J. Yoo, (2018), arXiv:1802.10368 [astro-ph.CO].
[21] J. Hamann, S. Hannestad, J. Lesgourgues, C. Rampf, and Y. Y. Y. Wong, *JCAP* **1007**, 022 (2010), arXiv:1003.3999 [astro-ph.CO].
[22] Z. Vlah, U. Seljak, M. Y. Chu, and Y. Feng, *JCAP* **1603**, 057 (2016), arXiv:1509.02120 [astro-ph.CO].
[23] D. Baumann, D. Green, and B. Wallisch, *JCAP* **1808**, 029 (2018), arXiv:1712.08067 [astro-ph.CO].
[24] A. Lewis, A. Challinor, and A. Lasenby, *Astrophys. J.* **538**, 473 (2000), arXiv:astro-ph/9911177 [astro-ph].
[25] N. Mostek, A. L. Coil, M. Cooper, M. Davis, J. A. Newman, and B. J. Weiner, *Astrophys. J.* **767**, 89 (2013), arXiv:1210.6694 [astro-ph.CO].
[26] D. J. Eisenstein, H.-J. Seo, and M. White, *Astrophys. J.* **664**, 660 (2007), arXiv:astro-ph/0604361 [astro-ph].
[27] H.-J. Seo and D. J. Eisenstein, *Astrophys. J.* **665**, 14 (2007), arXiv:astro-ph/0701079 [astro-ph].
[28] D. J. Eisenstein, H.-J. Seo, E. Sirko, and D. N. Spergel, *Astrophys. J.* **664**, 675 (2007), arXiv:astro-ph/0604362 [astro-ph].
[29] N. Padmanabhan, X. Xu, D. J. Eisenstein, R. Scalzo, A. J. Cuesta, K. T. Mehta, and E. Kazin, *Mon. Not. R. Astron. Soc.* **427**, 2132 (2012), arXiv:1202.0090 [astro-ph.CO].
[30] L. Anderson, E. Aubourg, S. Bailey, D. Bizyaev, M. Blanton, A. S. Bolton, J. Brinkmann, J. R. Brownstein, A. Burden, A. J. Cuesta, L. A. N. da Costa, K. S. Dawson, R. de Putter, D. J. Eisenstein, J. E. Gunn, H. Guo, J.-C. Hamilton, P. Harding, S. Ho, K. Honscheid, E. Kazin, D. Kirkby, J.-P. Kneib, A. Labatie, C. Loomis, R. H. Lupton, E. Malanushenko, V. Malanushenko,

- R. Mandelbaum, M. Manera, C. Maraston, C. K. McBride, K. T. Mehta, O. Mena, F. Montesano, D. Muna, R. C. Nichol, S. E. Nuza, M. D. Olmstead, D. Oravetz, N. Padmanabhan, N. Palanque-Delabrouille, K. Pan, J. Parejko, I. Pâris, W. J. Percival, P. Petitjean, F. Prada, B. Reid, N. A. Roe, A. J. Ross, N. P. Ross, L. Samushia, A. G. Sánchez, D. J. Schlegel, D. P. Schneider, C. G. Scóccola, H.-J. Seo, E. S. Sheldon, A. Simmons, R. A. Skibba, M. A. Strauss, M. E. C. Swanson, D. Thomas, J. L. Tinker, R. Tojeiro, M. V. Magaña, L. Verde, C. Wagner, D. A. Wake, B. A. Weaver, D. H. Weinberg, M. White, X. Xu, C. Yèche, I. Zehavi, and G.-B. Zhao, *Mon. Not. R. Astron. Soc.* **427**, 3435 (2012), arXiv:1203.6594 [astro-ph.CO].
- [31] F. Beutler, H.-J. Seo, A. J. Ross, P. McDonald, S. Saito, A. S. Bolton, J. R. Brownstein, C.-H. Chuang, A. J. Cuesta, D. J. Eisenstein, A. Font-Ribera, J. N. Grieb, N. Hand, F.-S. Kitaura, C. Modi, R. C. Nichol, W. J. Percival, F. Prada, S. Rodriguez-Torres, N. A. Roe, N. P. Ross, S. Salazar-Albornoz, A. G. Sánchez, D. P. Schneider, A. Slosar, J. Tinker, R. Tojeiro, M. Vargas-Magaña, and J. A. Vazquez, *Mon. Not. R. Astron. Soc.* **464**, 3409 (2017), arXiv:1607.03149 [astro-ph.CO].
- [32] B. D. Sherwin and M. Zaldarriaga, *Phys. Rev. D* **85**, 103523 (2012), arXiv:1202.3998 [astro-ph.CO].
- [33] A. Vallinotto, S. Dodelson, C. Schimd, and J.-P. Uzan, *Phys. Rev. D* **75**, 103509 (2007), astro-ph/0702606.
- [34] P. A. R. Ade *et al.* (Planck), (2015), arXiv:1502.01589.
- [35] A. Merson, Y. Wang, A. Benson, A. Faisst, D. Masters, A. Kiessling, and J. Rhodes, *Mon. Not. Roy. Astron. Soc.* **474**, 177 (2018), arXiv:1710.00833 [astro-ph.GA].
- [36] R. Laureijs, J. Amiaux, S. Arduini, J. . Auguères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, *et al.*, ArXiv e-prints (2011), arXiv:1110.3193.
- [37] H.-M. Zhu, Y. Yu, U.-L. Pen, X. Chen, and H.-R. Yu, *Phys. Rev. D* **96**, 123502 (2017), arXiv:1611.09638 [astro-ph.CO].
- [38] X. Wang, H.-R. Yu, H.-M. Zhu, Y. Yu, Q. Pan, and U.-L. Pen, *Astrophys. J.* **841**, L29 (2017), arXiv:1703.09742 [astro-ph.CO].
- [39] M. Schmittfull, T. Baldauf, and M. Zaldarriaga, *Phys. Rev. D* **96**, 023505 (2017), arXiv:1704.06634 [astro-ph.CO].
- [40] U. Seljak, G. Aslanyan, Y. Feng, and C. Modi, *Journal of Cosmology and Astro-Particle Physics* **2017**, 009 (2017), arXiv:1706.06645 [astro-ph.CO].
- [41] Y. Shi, M. Cautun, and B. Li, *Phys. Rev. D* **97**, 023505 (2018), arXiv:1709.06350 [astro-ph.CO].
- [42] R. Hada and D. J. Eisenstein, *Mon. Not. R. Astron. Soc.* **478**, 1866 (2018), arXiv:1804.04738 [astro-ph.CO].
- [43] E. Sarpa, C. Schimd, E. Branchini, and S. Matarrese, *Mon. Not. R. Astron. Soc.* **484**, 3818 (2019), arXiv:1809.10738 [astro-ph.CO].
- [44] D. Scott, D. Contreras, A. Narimani, and Y.-Z. Ma, *JCAP* **1606**, 046 (2016), arXiv:1603.03550 [astro-ph.CO].
- [45] R. Ansari *et al.* (Cosmic Visions 21 cm), (2018), arXiv:1810.09572 [astro-ph.CO].
- [46] S. Alam *et al.* (BOSS), *Mon. Not. Roy. Astron. Soc.* **470**, 2617 (2017), arXiv:1607.03155 [astro-ph.CO].
- [47] N. Aghanim *et al.* (Planck), (2018), arXiv:1807.06209 [astro-ph.CO].
- [48] A. G. Riess *et al.*, *Astrophys. J.* **861**, 126 (2018), arXiv:1804.10655 [astro-ph.CO].
- [49] G. E. Addison, D. J. Watts, C. L. Bennett, M. Halpern, G. Hinshaw, and J. L. Weiland, *Astrophys. J.* **853**, 119 (2018), arXiv:1707.06547 [astro-ph.CO].
- [50] Y. Li, M. Schmittfull, and U. Seljak, *Journal of Cosmology and Astro-Particle Physics* **2018**, 022 (2018), arXiv:1711.00018 [astro-ph.CO].
- [51] D. J. Eisenstein, W. Hu, and M. Tegmark, *Astrophys. J.* **504**, L57 (1998), arXiv:astro-ph/9805239 [astro-ph].
- [52] D. Baumann, F. Beutler, R. Flauger, D. Green, M. Vargas-Magaña, A. Slosar, B. Wallisch, and C. Yèche, arXiv e-prints, arXiv:1803.10741 (2018), arXiv:1803.10741 [astro-ph.CO].
- [53] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, *Phys. Rev. D* **92**, 043514 (2015), arXiv:1504.04366 [astro-ph.CO].
- [54] D. Grin, O. Doré, and M. Kamionkowski, *Phys. Rev. Lett.* **107**, 261301 (2011), arXiv:1107.1716.