Hubble constant tension between CMB lensing and BAO measurements

W. L. Kimmy Wu,[1] Pavel Motloch,[2] Wayne Hu,[3] and Marco Raveri[4]

[1] *Kavli Institute for Cosmological Physics, University of Chicago, Chicago, Illinois 60637, U.S.A*
[2] *Canadian Institute for Theoretical Astrophysics, University of Toronto, M5S 3H8, ON, Canada*
[3] *Kavli Institute for Cosmological Physics, Department of Astronomy & Astrophysics, Enrico Fermi Institute, University of Chicago, Chicago, Illinois 60637, U.S.A*
[4] *Center for Particle Cosmology, Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA*

We apply a tension metric $Q_{\rm UDM}$, the update difference in mean parameters, to understand the source of the difference in the measured Hubble constant H_0 inferred with cosmic microwave background lensing measurements from the *Planck* satellite ($H_0 = 67.9^{+1.1}_{-1.3}$ km/s/Mpc) and from the South Pole Telescope ($H_0 = 72.0^{+2.1}_{-2.5}$ km/s/Mpc) when both are combined with baryon acoustic oscillation (BAO) measurements with priors on the baryon density (BBN). $Q_{\rm UDM}$ isolates the relevant parameter directions for tension or concordance where the two data sets are both informative, and aids in the identification of subsets of data that source the observed tension. With $Q_{\rm UDM}$, we uncover that the difference in H_0 is driven by the tension between *Planck* lensing and BAO+BBN, at probability-to-exceed of 6.6%. Most of this mild tension comes from the galaxy BAO measurements parallel to the line of sight. The redshift dependence of the parallel BAOs pulls both the matter density Ω_m and H_0 high in ΛCDM, but these parameter anomalies are usually hidden when the BAO measurements are combined with other cosmological data sets with much stronger Ω_m constraints.

I. INTRODUCTION

The standard cosmological model ΛCDM is extremely successful in describing observations over a wide range of scales and redshifts: from the cosmic microwave background (CMB) to the expansion of the universe today. However, increasingly precise measurements of cosmological parameters obtained in the past several years uncovered mild to strong tensions between different data sets. Most notably, *Planck* infers the Hubble constant to be $H_0 = 67.36 \pm 0.54$ km/s/Mpc under ΛCDM [1] while the Cepheid-calibrated Type Ia Supernovae from SH0ES gives $H_0 = 74.03 \pm 1.42$ km/s/Mpc [2] (see [3–7] for other measurements). Such tensions between different data sets could suggest need to extend the ΛCDM model to accommodate the observations or, alternatively, existence of unmodeled systematics in the data sets. It is thus very important that we find independent measurements that can clarify the source(s) of the current H_0 tension.

One such example has been reported recently for the data sets that result from measurements of the weighted gravitational potential integrated along the line of sight (CMB lensing) from the *Planck* satellite and the South Pole Telescope (SPTpol), baryon acoustic oscillation (BAO) parallel and perpendicular to the line of sight in galaxy surveys, and the baryon density inferred from the deuterium abundance (D/H) measurements (denoted by BBN). In [8], the values of H_0 inferred from the BAO+BBN+*Planck* lensing and the BAO+BBN+SPTpol lensing data sets are $67.9^{+1.1}_{-1.3}$ km/s/Mpc and $72.0^{+2.1}_{-2.5}$ km/s/Mpc respectively. This mild discrepancy is intriguing since it is reminiscent of the tension between the SH0ES vs. *Planck* CMB power spectra measurements above.

In addition, the constraints on the lensing amplitude as captured by the parameter combination $\sigma_8\Omega_m^{0.25}$ between the two lensing data sets are completely consistent. This presents somewhat of a puzzle since the mild tension in H_0 appears through adding BAO+BBN to the otherwise consistent lensing data sets. Therefore, we set out to investigate the underlying driver(s) of the differences in the inferred H_0.

To do this, we apply a tension metric developed in [9] to quantify tension between BAO+BBN+*Planck* lensing and BAO+BBN+SPTpol lensing. We find that the difference in H_0 between the two is driven by the different inferences from the shape of the *Planck* lensing spectrum and BAO+BBN. Specifically, it is the line-of-sight BAO measurements that pull H_0 high. This preference is ordinarily hidden when the line-of-sight BAO measurements are combined with other cosmological data sets, since it requires a high matter density in ΛCDM which is strongly ruled out by these other data sets.

This paper is organized as follows: In Sec. II we summarize the tension metric we use, before presenting the data sets used in this work in Sec. III. In Sec. IV, we isolate and quantify the tension between the BAO+BBN+SPTpol lensing and the BAO+BBN+*Planck* lensing data sets and show that it originates from the *Planck* lensing and the parallel BAO measurements. In Sec. V, we show that values of cosmological parameters preferred by the parallel BAO measurements are strongly ruled out by other cosmological data sets. We discuss the results and conclude in Sec. VI.

II. QUANTIFYING TENSIONS

To quantify tensions between uncorrelated data sets, we use the update difference-in-mean (UDM) statistic defined in [9]. This statistic compares the mean parameter values from a data set A alone, θ_α^A, with their "updated" values after adding another data set B to A,

θ_α^{A+B}. The index α here enumerates the individual parameters. Specifically, this statistic computes the square of the difference between the mean parameter values of the two sets

$$\Delta\bar{\theta}_\alpha = \bar{\theta}_\alpha^A - \bar{\theta}_\alpha^{A+B}, \tag{1}$$

in units of its covariance C_Δ,

$$Q_{\rm UDM}^{\rm tot} = \sum_{\alpha,\beta} \Delta\bar{\theta}_\alpha \left(C_\Delta^{-1}\right)_{\alpha\beta} \Delta\bar{\theta}_\beta. \tag{2}$$

C_Δ is inferred from the covariances of A and $A+B$ as

$$C_\Delta = C^A - C^{A+B}. \tag{3}$$

To quantify the tension, we use the fact that if the A and $A+B$ parameter posteriors are Gaussian distributed and drawn from a self-consistent model, $Q_{\rm UDM}^{\rm tot}$ is chi-squared distributed with the number of parameters measured by both A and B as the degrees of freedom.

There are two advantages of using $Q_{\rm UDM}^{\rm tot}$ as opposed to the simpler difference-in-mean statistic. The first is practical. If the parameter posterior for B is highly non-Gaussian due to weak constraints or degeneracies, then the squared difference-in-mean is also far from chi-squared distributed, making its significance difficult to quantify. However, if the parameter posterior for A is more Gaussian, then the posterior of the combined data set $A+B$ is also more Gaussian and $Q_{\rm UDM}^{\rm tot}$ is closer to chi-squared distributed. The second advantage is the ability to pre-identify parameter directions in which the combination of A and B improves the errors over A or B individually and hence can exhibit interesting tension or confirmation bias. This is in contrast to defining the parameter space of investigation by inspecting the parameter differences in mean *a posteriori*, e.g. by picking the most discrepant directions.

An effective method to isolate these directions is to identify the Karhunen-Loève (KL) eigenmodes of the covariance matrices [9]. These eigenmodes ϕ_α^a are the solutions to the generalized eigenvalue problem

$$\sum_\beta C_{\alpha\beta}^A \phi_\beta^a = \lambda^a \sum_\beta C_{\alpha\beta}^{A+B} \phi_\beta^a, \tag{4}$$

and are normalized so that

$$\sum_{\alpha\beta} \phi_\alpha^a C_{\alpha\beta}^{A+B} \phi_\beta^b = \delta^{ab}. \tag{5}$$

The parameters in the KL basis

$$p^a = \sum_\alpha \phi_\alpha^a \theta_\alpha \tag{6}$$

are uncorrelated for both A and $A+B$ with variance λ^a and 1, respectively. From Eqs. (2) and (3),

$$Q_{\rm UDM} = \sum_a \frac{(\Delta p^a)^2}{\lambda^a - 1}. \tag{7}$$

If the sum is over all KL modes then $Q_{\rm UDM} = Q_{\rm UDM}^{\rm tot}$. However, to isolate the directions of interest, we restrict this sum to eigenvalues

$$0.1 < \lambda^a - 1 < 100. \tag{8}$$

Notice that this selection does not involve the actual values of the difference in means, only the expected ability of B to update A and vice versa. For cases where $\lambda_a \approx 1$, data set B does not update the constraints of A appreciably whereas for $\lambda_a \gg 1$, data set A itself becomes irrelevant and cannot update B. In the former case, there are also numerical problems due to the MCMC sampling of the posteriors. This selection also covers cases where there are nuisance parameters that are constrained by only one of A or B.

With only the interesting directions in the parameter space retained, we can then determine the significance of their associated difference in means by noting that $Q_{\rm UDM}$ is chi-squared distributed with the number of remaining parameter directions as the degrees of freedom. From this point forward we refer to $Q_{\rm UDM}$ as defined by Eqs. (7) and (8) as the update difference-in-mean statistic.

We will also be interested in how much information each KL eigenmode contributes to constraining individual cosmological parameters θ_α. Recall that the Fisher information matrix is the inverse of the parameter covariance matrix $F_{\alpha\beta} = (C_{\alpha\beta})^{-1}$ and each diagonal entry corresponds to the inverse variance of the parameter if all other parameters are held fixed. Using Eq. (5), we can express this Fisher information of data set A for the parameter θ_α as

$$F_{\alpha\alpha} = \sum_a F_{\alpha\alpha}^a = \sum_a \phi_\alpha^a \phi_\alpha^a / \lambda^a. \tag{9}$$

The Fisher information of data set $A+B$ is the same expression with $\lambda^a \to 1$, but this will not be needed in our analysis below.

The fractional Fisher information $F_{\alpha\alpha}^a / F_{\alpha\alpha} \in [0,1]$ parameterizes how important KL mode a is in constraining the cosmological parameter θ_α, where low values mean that dropping this mode does not significantly affect its constraints.

When considering correlated data sets, in particular the internal consistency of parallel and perpendicular BAO measurements, we use the generalization of the above discussion as in [10]. Specifically, we duplicate the parameter space of the model, and fit the joint data set with one copy of parameters controlling the theory prediction for the first part of the joint data set (e.g. parallel BAOs) and the other copy controlling the theory prediction for the second part (e.g. perpendicular BAOs). We then assess the confidence intervals of the difference in these two parameter sets by sampling its posterior. Because we fit the joint data set, the correlations are properly accounted for. This technique also has the benefit of applying to non-Gaussian posterior distributions.

III. DATA

The data sets we investigate in Sec. IV include *Planck* 2018 lensing [11], SPTpol lensing [12], BAO from SDSS DR12 BAO consensus sample [DR12, 13], Main Galaxy Sample [MGS, 14], and the 6dF Galaxy Survey [6dF, 15], and baryon density prior motivated by [16] using D/H measurements. In Sec. V we compare BAO constraints with parameter constraints from the *Planck* temperature and polarization power spectra [17] and from the Pantheon supernova sample [18].

The applicability of the $Q_{\rm UDM}$ statistic hinges upon the data sets being uncorrelated. For the data sets we are considering, the BAO measurements are uncorrelated with the lensing measurements. While the *Planck* and the SPTpol measurements have partial sky overlap, the overlap is very small ($\sim 1\%$) and the angular scales of overlap are small as well (with *Planck*'s $L = [8, 400]$ and SPTpol's $L = [100, 2000]$). Therefore, the lensing data sets are nearly uncorrelated. Consequently $Q_{\rm UDM}$ is appropriate for quantifying tension between our data sets. We only require the parameter duplication generalization discussed in the previous section to quantify tension between correlated subsets of a given data set (e.g. BAO).

In all cases, we use CosmoMC [19] to sample the posteriors of these data sets. We impose the following priors for ΛCDM parameters when sampling: uniform priors for the cold dark matter density $\Omega_c h^2 = [0.001, 2.99]$, initial curvature power spectrum amplitude $\ln(10^{10} A_s) = [1.61, 3.91]$, and the effective angular sound horizon scale $\theta_{\rm MC} = [0.5, 10]$. We assume Gaussian priors (mean, σ) for the initial spectrum tilt n_s : $(0.96, 0.02)$ and the baryon density $\Omega_b h^2$: $(0.0222, 0.0005)$ (the latter representing the D/H measurements/BBN data) and we fix the optical depth to recombination τ to 0.055. To draw the contour plots, we use GetDist [20].

IV. SPTPOL AND PLANCK LENSING VS. BAO+BBN

In this section, we start with identifying the key parameter directions that contribute to the apparent tension between SPTpol lensing and *Planck* lensing when both data sets are combined with BAO+BBN using the update difference in mean statistic of Sec. II. Upon finding that the apparent tension is not between SPTpol and *Planck* lensing, but rather between *Planck* lensing and the parallel BAO measurements, we then focus on and quantify tension between those measurements.

A. Parameters

The BAO+BBN+lensing data sets depend on 5 out of the 6 ΛCDM parameters (they are not sensitive to τ). In order to obtain more Gaussian covariance matrices, we perform our investigations in the parameter space that is native to the BAO and CMB lensing measurements. This way, it is also easier to interpret the influence of each. Specifically, we work in the parameter basis

$$\theta = \left[\theta_\perp, \theta_\parallel, \sigma_8 \Omega_m^{1/4}, \Omega_b h^2, n_s\right]. \tag{10}$$

Here

$$\begin{aligned} \theta_\perp &= D_M(z_{\rm BAO}) \frac{r_{\rm fid}}{r_{\rm d}}, \\ \theta_\parallel &= H(z_{\rm BAO}) \frac{r_{\rm d}}{r_{\rm fid}}, \end{aligned} \tag{11}$$

where $D_M(z)$ is the comoving angular diameter distance to redshift z, $H(z)$ is the expansion rate at this redshift, $r_{\rm d}$ is the comoving BAO scale (the sound horizon at the end of the Compton drag epoch), $r_{\rm fid} \equiv 147.78\,{\rm Mpc}$ is the fiducial $r_{\rm d}$, and σ_8 is the root mean square of the linear matter density fluctuations at the $8h^{-1}$ Mpc scale. We choose $z_{\rm BAO} = 0.61$, one of the DR12 points, but other choices of z within the DR12 range do not qualitatively affect our results.

The first three parameters are the most relevant to this work and may be interpreted as a perpendicular BAO, parallel BAO, and CMB lensing amplitude parameter. Given that under ΛCDM $H_0 D_M(z)$ and $H(z)/H_0$ are functions of Ω_m alone, the first two parameters span the same space as do Ω_m and $H_0 r_{\rm d}$. $\Omega_b h^2$ is constrained mainly by the BBN data and n_s is a nuisance parameter that is constrained by the prior given in Sec. III.

The BAO measurements do not depend on the lensing amplitude parameter $\sigma_8 \Omega_m^{1/4}$, while the shape of the lensing power spectrum within ΛCDM does depend on the BAO parameters, in this case mainly supplying extra information on $\theta_\parallel$. The difference between SPTpol and *Planck* lensing can be attributed to the fact that the former mainly constrains the amplitude of the lensing power spectrum whereas the latter constrains both the amplitude and shape [8].

We can see some of these properties in the posterior distributions of these three parameters shown in Fig. 1. In this space, the tension between BAO+BBN+SPTpol lensing and BAO+BBN+*Planck* lensing is confined to the parallel BAO parameter $\theta_\parallel$. From the BAO+BBN result, we can see the shift in this parameter is already present without the SPTpol lensing data, albeit with larger uncertainties. Another notable observation is that the lensing amplitude parameter constrained by SPTpol lensing and *Planck* lensing appears to agree too well with each other. Finally, the perpendicular BAO parameter $\theta_\perp$ is mainly constrained by the BAO data themselves and the addition of either lensing data set does not change its posterior appreciably.

To tie the tension between BAO+BBN+SPTpol lensing and BAO+BBN+*Planck* lensing seen in $\theta_\parallel$ to the tension originally identified in H_0, we show the posteriors of these two parameters in Fig. 2. Note that the distributions and shifts in means follow each other due to the high correlation between the two parameters. We will hereafter use tension in $\theta_\parallel$ as a proxy for this tension

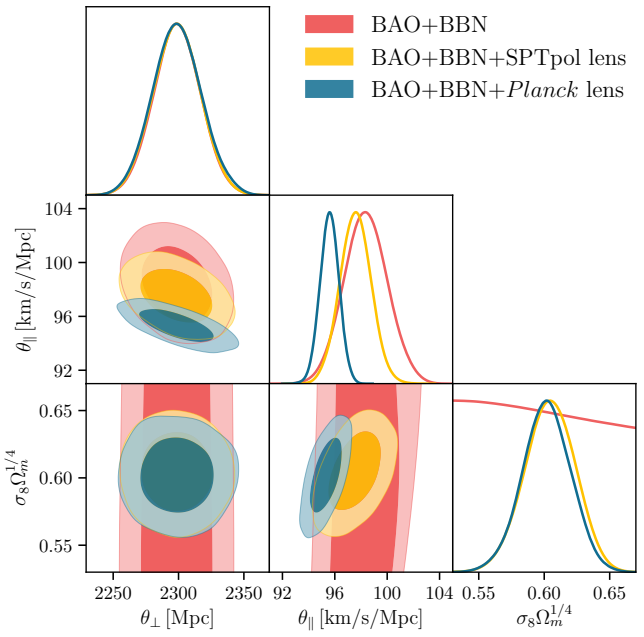

FIG. 1. The posterior distributions of the data sets BAO+BBN, BAO+BBN+SPTpol lensing, BAO+BBN+*Planck* lensing in the parameter basis $[\theta_\perp, \theta_\parallel, \sigma_8\Omega_m^{1/4}]$ which is native to the BAO and CMB lensing measurements. The posteriors for the combined BAO+BBN+*Planck* lensing+SPTpol lensing data set are not shown for clarity as they are qualitatively close to those from BAO+BBN+*Planck* lensing.

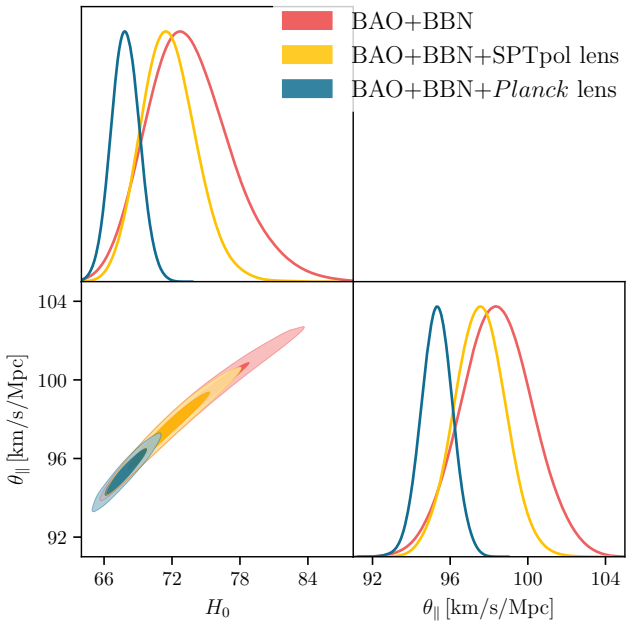

FIG. 2. The posterior distributions of the data sets BAO+BBN, BAO+BBN+SPTpol lensing, BAO+BBN+*Planck* lensing for H_0 and $\theta_\parallel$. The differences observed in H_0 are highly correlated with the differences in $\theta_\parallel$, which suggests a common origin in the parallel BAO measurements.

in H_0. In the following sections, we use the $Q_{\rm UDM}$ analysis to quantify these tensions and further isolate their origin in the various data sets.

B. BAO+BBN+SPTpol lensing vs *Planck* lensing

We first analyze tension between BAO+BBN+SPTpol lensing and its update BAO+BBN+SPTpol lensing+*Planck* lensing through the $Q_{\rm UDM}$ statistic. In this case, two directions satisfy the KL update criteria on eigenvalues (Eq. 8). These two KL eigenmodes are $a = 4, 5$ and, as shown in Fig. 3, they dominate the Fisher information for the lensing amplitude parameter $\sigma_8\Omega_m^{1/4}$ and the parallel BAO parameter $\theta_\parallel$ respectively. We will refer to them below as the amplitude mode and the parallel mode. With these two degrees of freedom, $Q_{\rm UDM} = 4.0$, which corresponds to a probability-to-exceed (PTE) of $\sim 13\%$. We obtain similar values for all DR12 values of $z_{\rm BAO}$.

Taken at face value, the PTE signals that the parameters of BAO+BBN+SPTpol lensing and *Planck* lensing are not particularly in tension. However, as can be seen in Fig. 1 and also [8], the best-fit $\sigma_8\Omega_m^{1/4}$ between SPTpol lensing and *Planck* lensing are almost too consistent with each other despite being nearly uncorrelated in their lensing information. It is therefore interesting to examine the individual contributions to $Q_{\rm UDM}$ of the amplitude and parallel modes. We find that the amplitude mode contributes only 0.11 whereas the parallel mode contributes 3.9 to the total $Q_{\rm UDM}$. Considered separately, these correspond to a PTE of 74% and 4.7% respectively. Because the contribution to $Q_{\rm UDM}$ from the amplitude mode is smaller than expected, the total significance downplays the tension in the parallel mode. This parallel mode reflects the tension originally identified in H_0. Indeed if we computed the update difference in mean for only the marginal H_0 distributions, we would obtain a 6.8% PTE or effectively a $1.8\,\sigma$ tension.

As discussed in the previous section, the BAO+BBN data do not contribute to the $\sigma_8\Omega_m^{1/4}$ constraint and the SPTpol lensing measurements contribute little to the $\theta_\parallel$ constraint. The $Q_{\rm UDM}$ contributions imply that the $\sigma_8\Omega_m^{1/4}$ measurements from SPTpol lensing and *Planck* lensing are slightly too consistent, while the $\theta_\parallel$ parameter from the BAO data set and *Planck* lensing are in mild tension of $\sim 2\,\sigma$. With this information, we now see that the difference between SPTpol lensing and *Planck* lensing when both are combined with BAO+BBN, first identified in H_0, is actually driven by the mild tension in $\theta_\parallel$ between *Planck* lensing and the BAO measurements.

We next focus on quantifying this tension between *Planck* lensing and BAO+BBN.

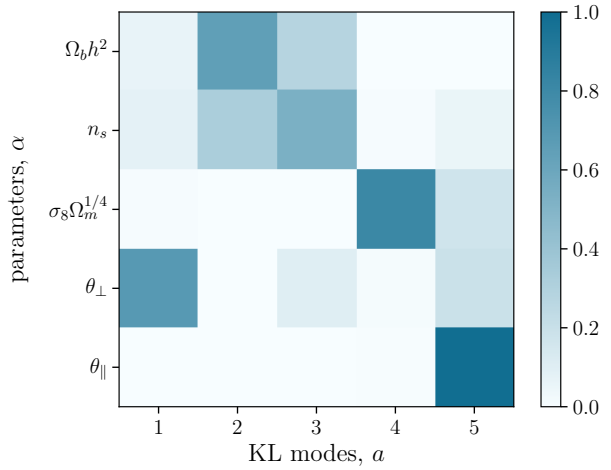

FIG. 3. Fractional Fisher information $F^a_{\alpha\alpha}/F_{\alpha\alpha}$ of the BAO+BBN+SPTpol lensing data set computed using the KL eigenmodes from updating it with *Planck* lensing. The numbers in each row add to one. The KL directions $a = 4$ and $a = 5$ satisfy the KL update criteria (Eq. 8). They contribute the most information to the lensing amplitude parameter $\sigma_8\Omega_m^{1/4}$ and the parallel BAO parameter $\theta_\parallel$ respectively.

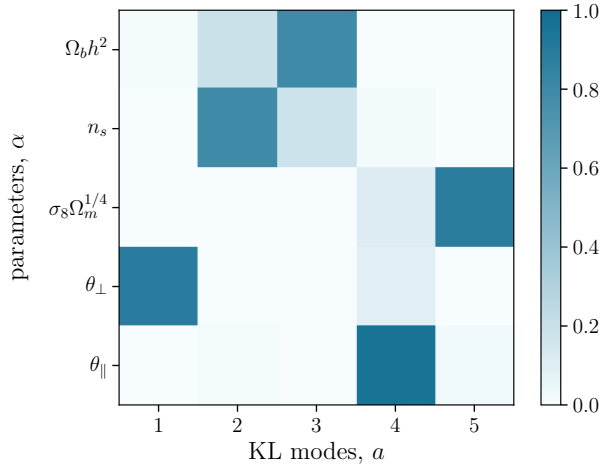

FIG. 4. Fractional Fisher information $F^a_{\alpha\alpha}/F_{\alpha\alpha}$ of the BAO+BBN data set computed using the KL eigenmodes from updating it with *Planck* lensing. Only the KL direction $a = 4$ satisfies the KL update criteria. It contributes the most information to the parallel BAO parameter $\theta_\parallel$.

C. BAO+BBN vs *Planck* lensing

In the previous section, we determined that the parallel BAO parameter $\theta_\parallel$ is the main indicator of tension in the data sets we consider. Given that SPTpol lensing has very little information on this parameter, we now focus on the comparison between BAO+BBN and *Planck* lensing.

In this case, we calculate $Q_{\rm UDM}$ between BAO+BBN and its update BAO+BBN+*Planck* lensing. From the KL decomposition, only one mode satisfies the KL update criteria on eigenvalues. For BAO+BBN, this mode again dominates the information on the parallel BAO parameter $\theta_\parallel$, as shown by the fractional Fisher information of this "parallel" mode $a = 4$ in Fig. 4. In this case $a = 5$ dominates the information in the lensing amplitude parameter but its constraints come almost entirely from *Planck* lensing and so do not satisfy Eq. (8).

With the parallel mode, $Q_{\rm UDM} = 3.37$ for a single degree of freedom and hence the PTE is 6.6%. We obtain similar values for all DR12 values of $z_{\rm BAO}$. If we compute the update difference in means from the marginalized constraints on H_0 alone, we obtain a PTE of 10.5%, underestimating the significance. To confirm that SPTpol lensing does not contribute to $\theta_\parallel$ beyond what *Planck* lensing does in this context, we calculate the $Q_{\rm UDM}$ between BAO+BBN and its update BAO+BBN+*Planck* lensing+SPTpol lensing. The eigenmodes have similar distributions as the BAO+BBN update with *Planck* lensing case and the PTE is also 6.6%, concluding that SPTpol lensing does not affect this result. Again, but now more explicitly, the $Q_{\rm UDM}$ analysis shows that the parallel BAO parameter is in mild tension between *Planck* lensing and BAO+BBN.

For the BAO data we have used in this work, the SDSS DR12 BAO data set have separate parallel BAO and perpendicular BAO measurements. In the next section, we look into the effects on the BAO parameters from the parallel and the perpendicular BAO measurements separately.

D. BAO DR12 parallel vs *Planck* lensing

To identify the origin of tension between BAO+BBN and *Planck* lensing, we examine the posterior constraints on the BAO parameters $\theta_\perp$ and $\theta_\parallel$ from various subsets of the BAO measurements themselves in Fig. 5. Since these individual constraints are themselves too weak to have data-dominated Gaussian posteriors, we do not employ $Q_{\rm UDM}$ here.

First, we see that the BAO constraints mainly come from the DR12 points with little information from other BAO data (MGS, 6dF). Next, we see that parallel and perpendicular DR12 measurements map onto these parameters without significant degeneracy, as expected given the design of the parameters. However, the correspondence between the measurements and parameters is not entirely one-to-one. There is a small amount of constraining power of the perpendicular measurements on the parallel parameter and vice versa due to the combination of the three redshift points in each set. Recall that we take $z_{\rm BAO} = 0.61$ in our fiducial parameter choice, which is the highest of the three DR12 redshifts. We see in Fig. 5 that the 68% CL regions of the two posteriors almost overlap, which suggests that the parallel and perpendicular measurements might be in mild tension. However, the level of tension cannot be inferred

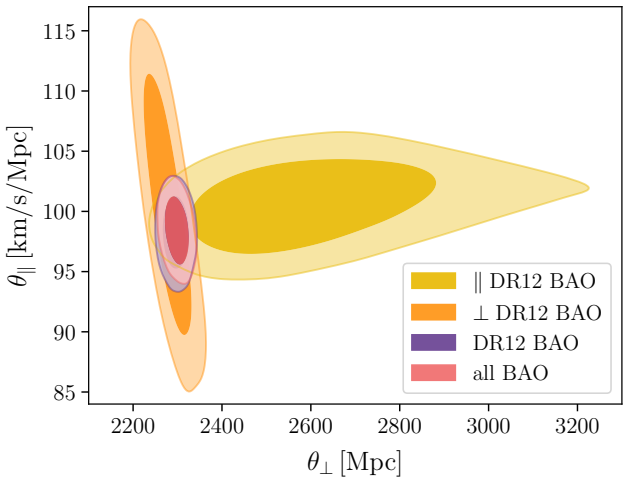

FIG. 5. $\theta_\parallel$ vs $\theta_\perp$ for various combinations of BAO data.

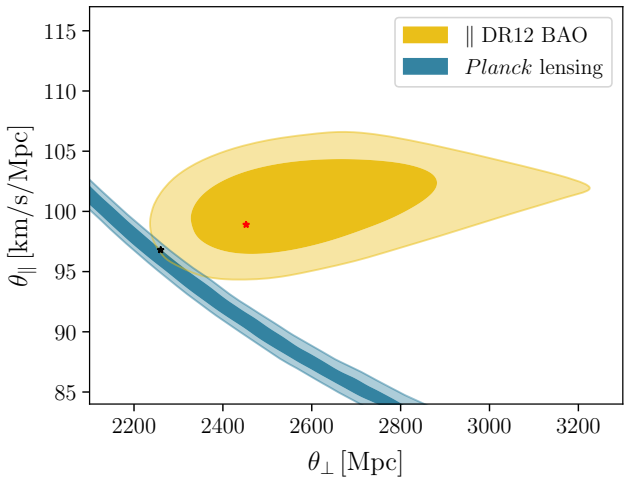

FIG. 6. $\theta_\parallel$ vs $\theta_\perp$ for parallel BAO and *Planck* lensing. The red star is the best-fit model to the parallel BAO measurements, the black star is a point we choose for illustration. It lies on the 95% CL line of the parallel BAO posterior and 68% CL line of the *Planck* lensing posterior. We calculate the $\Delta\chi^2$ between these two models using the parallel BAO likelihood.

from this observation directly because of the correlation between the parallel and perpendicular measurements. To properly account for the correlation, we apply the parameter duplication technique [10] described in Sec. II, and find that the PTE associated with zero parameter difference is 32%, indicating no significant tension.

Comparing the posteriors of the perpendicular and parallel BAO measurements with that of *Planck* lensing in the $\theta_\parallel - \theta_\perp$ plane, we find good overlap between perpendicular BAO and *Planck* lensing but that the 95% CL regions from parallel BAO barely overlaps with the *Planck* lensing posterior, as shown in Fig. 6. However, we caution the reader that these two data sets have non-Gaussian posterior probabilities and intuition of tension based on Gaussian data-dominated posteriors may not apply. Specifically, due to the weak constraining power of the parallel BAO data in the $\theta_\perp$ direction, the shape of the priors (assumed flat in $\theta_{\rm MC}, \Omega_c h^2$) over the constrained range informs the shape of the posterior.

To demonstrate that the prior is informative, we consider a goodness-of-fit statistic for the parallel BAO data. We compare $\chi^2_{\parallel\,\rm BAO}$ between the best-fit parameters to the parallel BAO data (red star in Fig. 6) and a representative test case (black star). The latter sits on the 95% CL line for the parallel BAO data, which for a two dimensional Gaussian likelihood and flat priors in the BAO parameters would correspond to $\Delta\chi^2_{\parallel\,\rm BAO} = 6$. However, actually evaluating the difference in $\chi^2_{\parallel\,\rm BAO}$ between the two cosmological models highlighted in Fig. 6 leads to only $\Delta\chi^2_{\parallel\,\rm BAO} = 3.2$. Because of the increase of the prior volume at high $\theta_\perp$, which we shall see is associated with the large range in Ω_m that it encompasses, the yellow contours in Fig. 6 are shifted to the right compared to the position of the best-fit model.

This shows that, in fact, the two models are not as discrepant as one would normally infer from a 95% CL exclusion in two dimensions. However, this mild discrepancy does account for a large portion of the tension between BAO+BBN and *Planck* lensing in the previous section. This is in part because the *Planck* lensing constraint, and any tension with it, is effectively one-dimensional in the BAO parameters, where a 95% exclusion would correspond to $\Delta\chi^2 = 3.8$.

Correspondingly, we can trace these results back to the fits to the 3 redshift points of the DR12 BAO data themselves. We plot in Fig. 7 the parallel and perpendicular BAO measurements against the parallel BAO best-fit model and the test model, all relative to a reference model chosen as the best-fit to the BAO+BBN+*Planck* lensing data set. On the top panel, we see that the test model does not deviate from the BAO+BBN+*Planck* lensing best-fit very much. The parallel BAO best-fit model does not fit the perpendicular BAO measurements, with χ^2 of 57 for 3 data points. On the lower panel, we show the parallel BAO measurements against the same sets of models. The measurements are very well fit by the parallel BAO best-fit model. The test model reflects the $\Delta\chi^2_{\parallel\,\rm BAO} = 3.2$ shift noted above. Moreover, we can now see that this penalty in the fit comes from its mismatch to the redshift slope of the parallel BAO data points.

Finally, to allow for easier comparison with literature and with other data sets in the next section, in Fig. 8 we re-plot Fig. 6 in the parameter plane of Ω_m and $r_{\rm d}H_0$. As already mentioned, these parameters have a one-to-one mapping with $\theta_\parallel$ and $\theta_\perp$. From the figure, it is clear that the parallel BAOs pull in the direction of very large Ω_m. The best-fit parallel BAO model (red star) has $\Omega_m = 0.64$ and $r_{\rm d}H_0 = 84.0\,$km/s. With the BBN prior, the best-fit also has a high $H_0 = 74\,$km/s/Mpc,

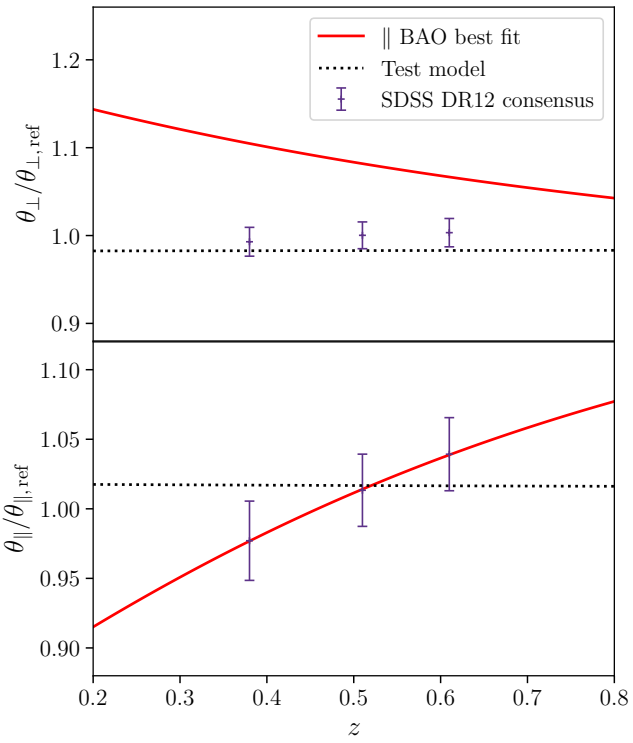

FIG. 7. SDSS DR12 BAO consensus measurements plotted against the $\theta_\perp(z)$ and $\theta_\parallel(z)$ model predictions from the parallel BAO best-fit and the test model, the red and black stars in Fig. 6 respectively. $\theta_\perp$ and $\theta_\parallel$ are normalized by θ_{ref}, the best-fit model to the BAO+BBN+*Planck* lensing data set.

but as we shall show in the next section, there are many other data sets that would exclude the high Ω_m required.

V. COMPARISON WITH OTHER DATA SETS

To put things in context, in Fig. 9 we compare constraints on Ω_m and $r_\mathrm{d}H_0$ from the parallel BAO and *Planck* lensing measurements with those from other cosmological data sets. We plot constraints on Ω_m from the Pantheon supernova sample [18] assuming flat ΛCDM (dashed lines). While these supernova constraints are in good agreement with *Planck* lensing and the combined BAO constraints, they are in mild tension with the parallel BAOs. We then show constraints from the *Planck* primary CMB power spectra measurements [17] (red contours), which are compatible with and even stronger than supernovae.

The best-fit Ω_m value of the parallel BAO+BBN data set (red star) is therefore strongly ruled out by both the supernova sample and the *Planck* primary CMB measurements. This disallowed preference for high Ω_m in the parallel BAO data set is the ultimate origin of the high H_0 preferred by the BAO+BBN+SPTpol lensing data set compared with BAO+BBN+*Planck* lensing.

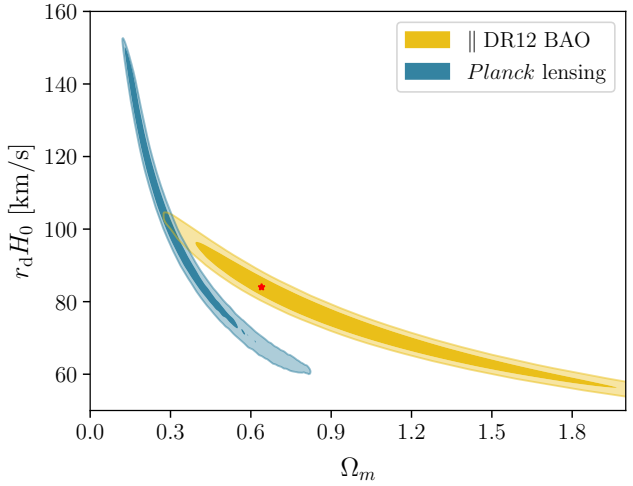

FIG. 8. Constraints as in Fig. 6 but here in terms of Ω_m and $r_\mathrm{d}H_0$. Note the long degeneracy for the parallel BAO data out to high Ω_m and the high value of $\Omega_m = 0.64$ for the parallel BAO best fit (red star).

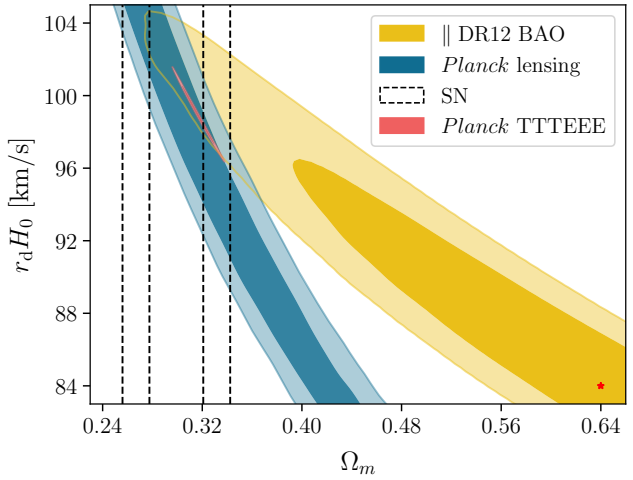

FIG. 9. Constraints on Ω_m and $r_\mathrm{d}H_0$ from parallel BAOs compared with those from *Planck* lensing, the Pantheon supernovae sample, and *Planck* CMB power spectra. Note that the posteriors extend beyond the part of the parameter plane shown here. All of the other data sets strongly disfavor the parallel BAO best fit (red star) due to its high value of Ω_m.

VI. CONCLUSION

In this work, we apply the update difference-in-mean statistic Q_{UDM} to quantify tension between two composite data sets, BAO+BBN+SPTpol lensing and BAO+BBN+*Planck* lensing, and track the origin of the differences in H_0 to the individual data sets that are primarily responsible. We work in a parameter basis that is native to the BAO and the CMB lensing measurements, replacing the cosmological parameters Ω_m, H_0

and A_s with $\theta_\parallel$, $\theta_\perp$, and $\sigma_8\Omega_m^{1/4}$, where the parameter posteriors are nearly Gaussian and their constraints are relatively easy to map back to the measurements. With this setup, we isolate the parameter direction that dominates the tension and isolate its origin in the *Planck* lensing vs. parallel BAO measurements.

We arrive at this conclusion through a process of narrowing down parameter combinations that matter and removing data sets that contribute little to the tension. In calculating the update difference in mean of BAO+BBN+SPTpol lensing updated by *Planck* lensing, the parameter direction that dominates the tension is $\theta_\parallel$ with a PTE of 4.7%. This direction is highly correlated with H_0, which carries a comparable tension. Knowing that SPTpol lensing contributes little to $\theta_\parallel$ constraints, we next check the $Q_{\rm UDM}$ of BAO+BBN updated by *Planck* lensing. This test confirms the tension between BAO+BBN and *Planck* lensing along the $\theta_\parallel$ direction at a PTE of 6.6%.

Both update difference in mean statistics point to the parallel BAO parameter as the source of the tension. We thus divide the BAO measurements into subsets to further our investigation. While the perpendicular BAO measurements are largely compatible with both the parallel BAO measurements and *Planck* lensing, there is tension between the parallel BAO and *Planck* lensing measurements around the 95% CL. This exclusion is exacerbated by our chosen prior, which in particular allows a large range in Ω_m. Independent of this prior, the $\Delta\chi^2$ between the best-fit model to the parallel BAO data and a representative model that is consistent with *Planck* lensing is $\Delta\chi^2 = 3.2$ for effectively 1 degree of freedom. These results indicate that the bulk of the tension between *Planck* lensing and the BAO data set is indeed from the parallel BAO measurements. Finally, we trace the origin of this $\Delta\chi^2$ to a slope in the parallel BAO measurements as a function of redshift, which drives its preference for high Ω_m values. In combination with constraints from BBN, this translates into a preference for high H_0 values.

We note that the Ω_m preferred by the parallel BAO data under ΛCDM is highly excluded by other data sets, including supernova measurements and *Planck* primary CMB measurements. These other measurements tend to have much stronger constraining power on Ω_m than the parallel BAO data. For this reason, the mild tension of the parallel BAO data with the other data sets is hidden when analyzed in combination. While this tension clearly cannot be resolved within ΛCDM, it is useful to bear this in mind when considering alternatives.

Finally, we reiterate the importance of not selecting cosmological parameters *a posteriori* when adjudicating tension between data sets. Had we calculated the difference in mean of BAO+BBN+SPTpol lensing updated with *Planck* lensing on H_0 alone, the PTE would be lower than letting the algorithm reveal that there are two relevant parameter combinations. Conversely, had we calculated the difference in mean of BAO+BBN updated with *Planck* lensing on H_0 alone, the PTE would be higher than letting the algorithm choose the single relevant parameter direction. With parameters selected *a posteriori*, a trials factor is required to accompany the resultant PTE for fair interpretation of the statistic and that selection may still not reflect the true source of tension. Looking forward, as upcoming surveys provide more precise measurements of our universe, it is of utmost importance that we identify the origin and significance of tension accurately to aid the differentiation of the underlying causes of the observed tension—be it unmodeled systematics or new physics.

ACKNOWLEDGMENTS

We thank Georgios Zacharegkas for useful discussions. WLKW is supported in part by the Kavli Institute for Cosmological Physics at the University of Chicago through grant NSF PHY-1125897 and an endowment from the Kavli Foundation and its founder Fred Kavli. WH is supported by by U.S. Dept. of Energy contract DE-FG02-13ER41958 and the Simons Foundation. MR is supported in part by NASA ATP Grant No. NNH17ZDA001N, and by funds provided by the Center for Particle Cosmology. This work was completed in part with resources provided by the University of Chicago Research Computing Center.

[1] N. Aghanim *et al.* (Planck), (2018), arXiv:1807.06209 [astro-ph.CO].
[2] A. G. Riess, S. Casertano, W. Yuan, L. M. Macri, and D. Scolnic, Astrophys. J. **876**, 85 (2019), arXiv:1903.07603 [astro-ph.CO].
[3] W. L. Freedman *et al.*, Astrophys. J. **882**, 34 (2019), arXiv:1907.05922 [astro-ph.CO].
[4] W. L. Freedman, B. F. Madore, T. Hoyt, I. S. Jang, R. Beaton, M. G. Lee, A. Monson, J. Neeley, and J. Rich, Astrophys. J. **891**, 57 (2020), arXiv:2002.01550 [astro-ph.GA].
[5] K. C. Wong *et al.*, (2019), arXiv:1907.04869 [astro-ph.CO].
[6] W. Yuan, A. G. Riess, L. M. Macri, S. Casertano, and D. Scolnic, Astrophys. J. **886**, 61 (2019), arXiv:1908.00993 [astro-ph.GA].
[7] D. Pesce *et al.*, Astrophys. J. **891**, L1 (2020), arXiv:2001.09213 [astro-ph.CO].
[8] F. Bianchini *et al.* (SPT), Astrophys. J. **888**, 119 (2020), arXiv:1910.07157 [astro-ph.CO].
[9] M. Raveri and W. Hu, Phys. Rev. **D99**, 043506 (2019), arXiv:1806.04649 [astro-ph.CO].
[10] M. Raveri, G. Zacharegkas, and W. Hu, (2019), arXiv:1912.04880 [astro-ph.CO].
[11] N. Aghanim *et al.* (Planck), (2018), arXiv:1807.06210 [astro-ph.CO].

[12] W. L. K. Wu *et al.*, Astrophys. J. **884**, 70 (2019), arXiv:1905.05777 [astro-ph.CO].
[13] S. Alam *et al.* (BOSS), Mon. Not. Roy. Astron. Soc. **470**, 2617 (2017), arXiv:1607.03155 [astro-ph.CO].
[14] A. J. Ross, L. Samushia, C. Howlett, W. J. Percival, A. Burden, and M. Manera, Mon. Not. Roy. Astron. Soc. **449**, 835 (2015), arXiv:1409.3242 [astro-ph.CO].
[15] F. Beutler, C. Blake, M. Colless, D. H. Jones, L. Staveley-Smith, L. Campbell, Q. Parker, W. Saunders, and F. Watson, Mon. Not. Roy. Astron. Soc. **416**, 3017 (2011), arXiv:1106.3366 [astro-ph.CO].
[16] R. J. Cooke, M. Pettini, and C. C. Steidel, Astrophys. J. **855**, 102 (2018), arXiv:1710.11129 [astro-ph.CO].
[17] N. Aghanim *et al.* (Planck), (2019), arXiv:1907.12875 [astro-ph.CO].
[18] D. M. Scolnic *et al.*, Astrophys. J. **859**, 101 (2018), arXiv:1710.00845 [astro-ph.CO].
[19] A. Lewis and S. Bridle, Phys. Rev. D **66**, 103511 (2002), arXiv:astro-ph/0205436 [astro-ph].
[20] A. Lewis, (2019), arXiv:1910.13970 [astro-ph.IM].