

# Cosmic microwave background constraints in light of priors over reionization histories

Marius Millea<sup>1,2</sup> and François Bouchet<sup>2</sup>

<sup>1</sup> Institut Lagrange de Paris (ILP), Sorbonne Universités, 98bis boulevard Arago, 75014 Paris, France  
e-mail: mariusmillea@gmail.com

<sup>2</sup> Institut d'Astrophysique de Paris (IAP), UMR7095, CNRS & Sorbonne Universités, 98bis boulevard Arago, 75014 Paris, France

Received 24 April 2018 / Accepted 24 May 2018

## ABSTRACT

Non-parametric reconstruction or marginalization over the history of reionization using cosmic microwave background data necessarily assumes a prior over possible histories. We show that different but reasonable choices of priors can shift current and future constraints on the reionization optical depth,  $\tau$ , or correlated parameters such as the neutrino mass sum,  $\Sigma m_\nu$ , at the level of 0.3–0.4  $\sigma$ , meaning that this analysis is somewhat prior dependent. We point out some prior-related problems with the commonly used principal component reconstruction, concluding that the significance of some recent hints of early reionization in *Planck* 2015 data has been overestimated. We also present the first non-parametric reconstruction applied to newer *Planck* intermediate (2016) data and find that the hints of early reionization disappear entirely in this more precise dataset. These results limit possible explanations of the EDGES 21cm signal which would have also significantly reionized the universe at  $z > 15$ . Our findings about the dependence on priors motivate the pursuit of improved data or searches for physical reionization models which can reduce the prior volume. The discussion here of priors is of general applicability to other non-parametric reconstructions, for example of the primordial power spectrum, of the recombination history, or of the expansion rate.

**Key words.** cosmic background radiation – dark ages, reionization, first stars – cosmological parameters

## 1. Introduction

The exact details of the epoch of reionization, during which the universe transitioned from a mostly neutral hydrogen gas into the highly ionized state we see today, are still of considerable uncertainty. This transition left several imprints on the cosmic microwave background (CMB) which can be used to constrain the physics of reionization, but also serve as a “nuisance” uncertainty which must be marginalized over when considering constraints on other parameters.

One of the most widely used methods for reconstructing or marginalizing over the reionization history from CMB data has been based on decomposing the history into a set of principal component amplitudes (PCAs), described initially by Hu & Holder (2003; hereafter HH03). This has been applied to a number of datasets, for example the WMAP data (Mortonson & Hu 2008b; Pandolfi et al. 2011, the former hereafter MH08), and most recently the *Planck* 2015 data (Heinrich et al. 2017; Heinrich & Hu 2018; Obied et al. 2018). The latter works argue that the generic reconstruction reveals a  $>2\sigma$  preference for a high-redshift ( $z \gtrsim 15$ ) contribution to the optical depth,  $\tau$ . Other works have used alternate methods, some finding somewhat similar preference for a high-redshift contribution (Hazra & Smoot 2017), while others finding none at all (Villanueva-Domingo et al. 2018).

In this work, we point out two problems with PCA analyses that have to-date been largely or completely overlooked. Correcting these problems will lead to finding reduced evidence of early reionization in the *Planck* 2015 data.

The first problem has been mentioned in Mortonson & Hu (2008b), Heinrich & Hu (2018), and described in some detail in

Lewis et al. (2006), although not explicitly in relation to the PCA model. The issue is that taking a flat prior on amplitudes of the reionization principal components induces a non-flat prior on  $\tau$ . Heinrich & Hu (2018) argue that this effect is unimportant, but do not perform direct or fully conclusive tests. Here we have performed the very direct test of simply calculating the effective  $\tau$  prior induced by the flat mode priors, finding that it is roughly  $\tau \approx 0.20 \pm 0.07$ . Using a maximum entropy procedure, we remove the effects of this prior, finding that it accounts for a significant part of the shift to higher  $\tau$  reported by Heinrich et al. (2017).

Another problem with existing PCA analyses lies in the choice of physicality priors. These priors are necessary to ensure that the reconstructed ionization fraction at any given redshift remains within physical range (i.e. remains positive but less than the maximum corresponding to a fully ionized universe). Due to the nature of the PCA decomposition, the standard physicality priors necessarily allow some unphysical models (this may seem quite counter-intuitive, but we will give a simple geometric picture of why this is indeed the case). However, existing analyses have used a sub-optimal set of physicality priors which allowed more unphysical models than necessary. We will derive the optimal set of priors, and show that switching to these leads to significant changes in constraints. Together with removing the impact of the informative  $\tau$  prior, we find that evidence for early reionization in the *Planck* 2015 data is reduced from  $>2\sigma$  to  $1.4\sigma$ .

Since the PCA results depend significantly on the details of the physicality priors, which will always be necessarily imperfect, we conclude that the PCA procedure is poorly tailored to the problem of generic reionization history reconstruction. We seek a better method, in particular one which does not allow

unphysical models. To this end, we propose what we will call the FlexKnot model. This model interpolates the history using a varying number of knots, with the exact number of knots determined by the data itself via a Bayesian evidence calculation. The model is somewhat similar to the Hazra & Smoot (2017) approach, but improves upon it by not requiring an a posteriori and fixed choice of knot positions, which otherwise creates undesired regions of high prior probability at certain redshifts. The FlexKnot model follows almost exactly the same procedure used previously to generically reconstruct the primordial power spectrum from *Planck* data (Vázquez et al. 2012; Planck Collaboration XX 2016), and is particularly well suited to the problem here as well.

Using these reionization models, we calculate new constraints on the history of reionization coming from *Planck* intermediate data, in particular using the *simlow* likelihood (Planck Collaboration Int. XLVI 2016). We will refer to the combination of *Planck* 2015 data and the *simlow* likelihood as the *Planck* 2016 data. This likelihood has already been used to explore various parametric reionization models (Planck Collaboration Int. XLVII 2016); however, a generic reconstruction fully capable of detecting an early component to reionization (should it be there) has not yet been performed. We report results from this particular dataset here.

The *Planck* 2016 data represents a significant increase in constraining power on reionization as compared to the 2015 data, and we find that the remaining  $1.4\sigma$  hints of an early reionization component are completely erased. Instead, we tightly constrain the  $z > 15$  contribution to the optical depth to be  $< 0.015$  at 95% confidence, and find good agreement with a late and fast transition to a fully ionized universe, consistent with the conclusions of Planck Collaboration Int. XLVII (2016).

Even though the *simlow* data are quite constraining, some dependence on priors remains when performing generic reconstruction. In particular, the difference between a flat prior on  $\tau$  and a flat prior on the knot positions leads to an  $0.4\sigma$  shift in the resulting  $\tau$  constraint. Both of these priors are fairly reasonable, and we consider it difficult to argue very strongly for one or the other of these priors in a completely generic setting. We thus consider this indicative of a level of systematic uncertainty stemming from our imperfect knowledge of the model, which we then propagate into shifts on other parameters which are correlated with  $\tau$ , given both current data and future forecasts.

We specifically focus on the sum of the neutrino masses,  $\Sigma m_\nu$ , which is expected to be measured extremely precisely by next generation CMB experiments (Abazajian et al. 2015). Using Fisher forecasts, we show that a possible  $0.4\sigma$  bias on  $\tau$  from switching priors manifests itself as a  $0.3\sigma$  bias on  $\Sigma m_\nu$  given expect constraints from a Stage 4 CMB experiment (CMB-S4; Abazajian et al. 2016). We comment on the ability of combining with other probes, such as DESI-BAO, or a cosmic-variance-limited CMB large-scale polarization measurement, to reduce this uncertainty.

Given these biases, we conclude that generic reconstructions of the reionization history are likely not enough to achieve the most accurate bounds on  $\Sigma m_\nu$ . This motivates work on obtaining constraints with physical models which have free parameters that can realistically account for our lack of exact knowledge about the physics of reionization. Although not directly used throughout the paper, will also comment on the possibility for other probes of reionization to reduce the dependence on priors, specifically measurements of the patchy kinetic Sunyaev-Zeldovich effect (Gruzinov & Hu 1998; Knox et al. 1998) and

direct probes of the ionization state of the inter-galactic medium (Bouwens et al. 2015).

The paper is outlined as follows. In Sect. 2 we discuss the reionization models we consider, and in Sect. 3 we focus on the implicit  $\tau$  prior induced by these procedures. These sections are a fairly technical description of the methodology used, and those wishing to see the results should skip to Sects. 4–6 where we discuss constraints from *Planck* 2015, intermediate, and future data, respectively.

## 2. Reionization models

### 2.1. TANH model

The reionization history often assumed in CMB analyses because it has a useful parametric form and matches physical expectations somewhat well involves a single smooth step from an almost fully neutral universe<sup>1</sup> to one with hydrogen fully ionized and helium singly ionized. The free electron fraction, in our convention the ratio between the number density of free electrons and hydrogen nuclei,  $x_e \equiv n_e/n_H$ , is taken to be

$$x_e^{\text{TANH}}(z) = \frac{1 + f_{\text{He}}}{2} \left\{ 1 + \tanh \left[ \frac{y(z^*) - y(z)}{\Delta y} \right] \right\}, \quad (1)$$

with  $y(z) = (1+z)^{3/2}$  and  $\Delta y = 3/2(1+z)^{1/2}\Delta z$ , giving a transition centered at redshift  $z^*$  with width  $\Delta z = 0.5$ . Here the factor  $f_{\text{He}}$  is the number density ratio of helium to hydrogen nuclei (we have neglected all other atoms). We will refer to this model as the TANH model.

The contribution to the optical depth between any two redshifts  $z_1$  and  $z_2$  for this (or any) reionization history can be written as

$$\tau(z_1, z_2) = \sigma_T \int_{z_1}^{z_2} dz \frac{n_H(z)(1+z)^2}{H(z)} x_e(z), \quad (2)$$

where  $\sigma_T$  is the Thompson scattering cross-section. This is often used to parametrize Eq. (1) in terms of the total optical depth  $\tau \equiv \tau(0, z_{\text{early}})$  rather than  $z^*$  (where  $z_{\text{early}}$  is some redshift before reionization began, but after recombination ended).

Note that in this convention for  $x_e$ , the maximum ionization fraction can be greater than one, in particular can be as large as  $x_e^{\text{max}} \equiv 1 + f_{\text{He}}$  before the second recombination of helium. We also note that second helium recombination is expected to be a small contribution to  $\tau$ , on the order of  $\approx 0.001$  depending on the exact values of other cosmological parameters. As this is already fairly small, we ignore any model-dependence in helium second reionization and model it as another transition with the same hyperbolic tangent form as in Eq. (1) but with  $z_* = 3.5$  and  $\Delta z = 0.5$ , normalized such that it increases the ionization fraction from  $x_e^{\text{max}}$  to  $1 + 2f_{\text{He}}$ .

### 2.2. PCA model

The TANH model has some parametric freedom, more-so if  $\Delta z$  is also taken as a free parameter. However, we would like a model which can reproduce any arbitrary history, thus allowing us to generically reconstruct  $x_e(z)$  from the data.

<sup>1</sup> There is a small residual ionization level on the order of  $10^{-4}$  remaining after recombination. We ignore this as the data considered here is insensitive to it.

HH03 proposes a model based on decomposing the free electron fraction history into eigenmodes such that

$$x_e^{\text{HH03}}(z) = x_e^{\text{fid}}(z) + \sum_i m_i S_i(z) \quad (3)$$

for some fiducial  $x_e^{\text{fid}}(z)$ , eigenmode amplitudes  $m_i$ , and eigenmode templates  $S_i(z)$ . These templates have support between some  $z_{\text{min}}$  and  $z_{\text{max}}$  and are uniquely defined by three properties:

1. they form a special orthonormal basis<sup>2</sup> for  $x_e(z)$ , implying that  $\int dz S_i(z) S_j(z) = \delta_{ij}$ .
2. they diagonalize the covariance of the amplitude parameters given cosmic-variance-limited large-scale polarization data,  $\langle \Delta m_i \Delta m_j \rangle = \sigma_i^2 \delta_{ij}$ ;
3. they are ordered by increasing  $\sigma_i$ .

The lack of Gunn-Peterson absorption in quasars out to  $z \approx 6$  strongly constrains the universe to be very close to fully reionized by this redshift (Becker et al. 2001; Fan et al. 2002). We implicitly impose these bounds here by setting  $z_{\text{min}} = 6$ . Similarly as in other PCA analyses, we take  $z_{\text{max}} = 30$ .

### 2.3. Geometric view of physicality priors

The HH03 procedure also calls for a set of ‘‘physicality priors’’ on these mode amplitudes. These are necessary because otherwise nothing prevents the mode amplitudes from taking on values such that the bound  $0 < x_e(z) < x_e^{\text{max}}$  fails to hold for some  $z$ , and this would have no physical meaning. Mortonson & Hu (2008a) derive two sets of priors to enforce physicality. The first is an upper and lower bound on each  $m_i$  such that  $m_i^- \leq m_i \leq m_i^+$  with

$$m_i^\pm = \int dz \{ S_i(z) [1 - 2x_e^{\text{fid}}(z)] \pm |S_i(z)| \} / 2. \quad (4)$$

The second is an upper bound on the sum of the squares of the mode amplitudes,

$$\int dz \sum_i m_i^2 < \int dz [x_e^{\text{max}} - x_e^{\text{fid}}(z)]^2. \quad (5)$$

These priors may seem a bit abstruse, but they have a quite simple geometric interpretation which has not been highlighted before. We will give this interpretation here, aided by Fig. 1. This will also aid in spotting some improvements that can be made.

To start, consider the reionization history,  $x_e(z)$ , evaluated at  $N$  redshifts,  $z_i$ , and stacked into a vector  $x_i \equiv x_e(z_i)$ . The physicality region in the  $x_i$  vector space is trivial, each  $x_i$  must individually fall between 0 and  $x_e^{\text{max}}$ . This region is thus an  $N$ -dimensional hypercube with one vertex at the origin, extending into the positive hyper-quadrant, and with edge length  $x_e^{\text{max}}$ .

The decomposition into mode amplitudes is given by applying the transformation  $S$  to the residual between  $x_i$  and the fiducial model,

$$m_i = \sum_j S_{ij} (x_j - x_j^{\text{fid}}), \quad (6)$$

<sup>2</sup> For clarity, we suppress writing the integration limits  $z_{\text{min}}$  and  $z_{\text{max}}$  for the rest of the paper. The exact normalization is purely a definitional choice. Similarly, the distinction of special orthonormal basis, i.e. that  $S$  is purely a rotation, is not usually made, but we choose to do so here for definiteness and for aiding the subsequent geometrical interpretation.

where  $S_{ij} \equiv S_i(z_j)$ , i.e. it is a matrix for which each row is one of the eigenmodes. The physicality region in the  $m_i$  vector space is thus a translated and transformed version of the original hypercube. We note that because  $S$  is special orthonormal and thus can only contain rotations, the region remains a hypercube.

Let us now visualize this transformation in two dimensions, represented in the left panel of Fig. 1. The solid blue square is the physicality region in terms of  $x_i$ , here just  $x_1$  and  $x_2$ . The solid orange square is the same region after transformation by Eq. (6), assuming some arbitrary  $S$  for demonstration purposes, and taking  $x_j^{\text{fid}} = 0.15$  following Mortonson & Hu (2008a) and all other subsequent PCA analyses.

If we were simultaneously fitting  $m_1$  and  $m_2$  in our analysis, the physicality region would be trivial to enforce, and would be given simply by the solid orange square. However, the point of the PCA procedure is that we do not need to simultaneously fit both parameters, as higher order modes like  $m_2$  can be effectively marginalized over by just fixing them, since they have no observable impact and are decorrelated with the lower order modes like  $m_1$ . Usually we fix  $m_2 = 0$ , but any other value should work just as well. In designing the physicality prior for  $m_1$ , we need to ensure that all values that  $m_1$  could take are allowed, not just those allowed when  $m_2 = 0$ , since this was just an artificial method of marginalization<sup>3</sup>. Put another way, one must consider that a model which appears to be unphysical could be brought back into physicality by adjusting the higher order mode amplitudes which were artificially set to zero. This leads to two priors which can be constructed.

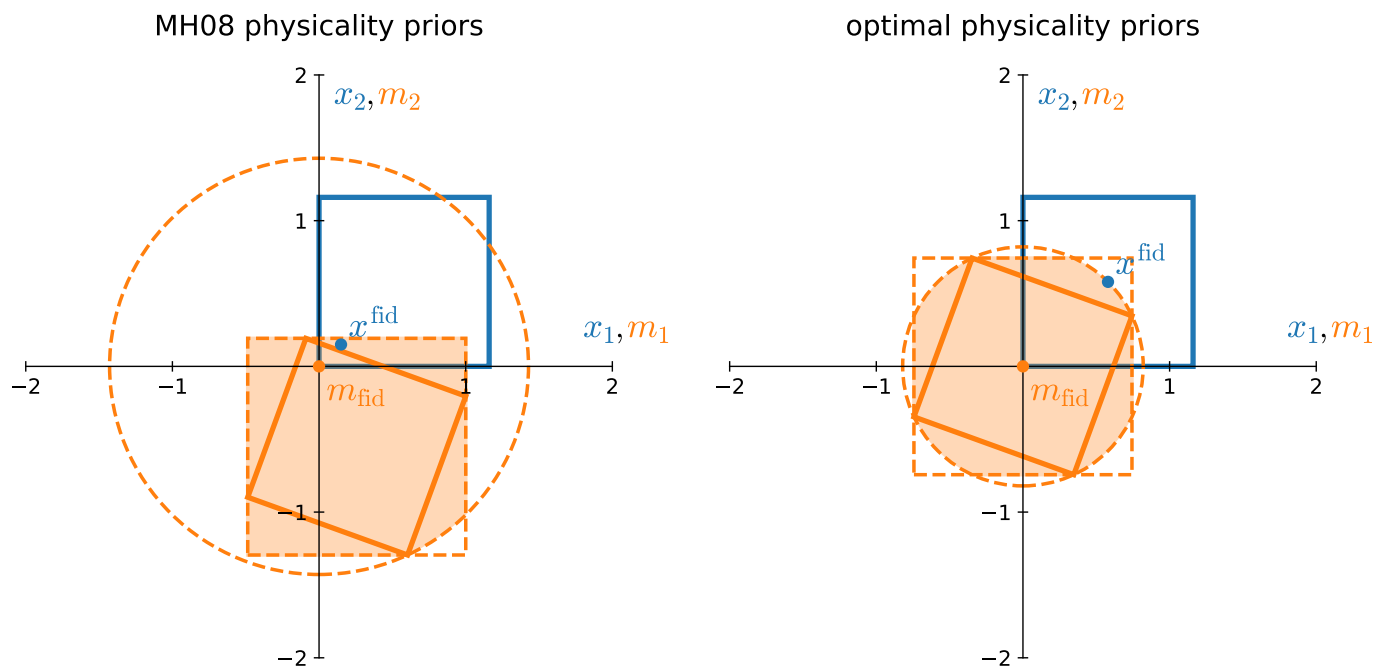
The first prior is simply the bounding box of the solid orange square and is depicted as the dashed orange square. We refer to this as the ‘‘hyperbox’’ prior<sup>4</sup>. This prior just takes a hard bound on each  $m_i$  corresponding to the largest and smallest possible value for that  $m_i$  anywhere within the physicality region. This is indeed exactly what is calculated by Eq. (4). One can see this by noting that the  $x_j$  which maximizes  $m_i$  in Eq. (6) is such that  $x_j = x_e^{\text{max}}$  if  $S_{ij}$  is positive, and zero otherwise, and that Eq. (4) is the transform of this  $x_j$ .

The second prior comes from noting that as long as  $x_e^{\text{fid}} \leq x_e^{\text{max}}/2$ , the top right corner of the physicality region in terms of  $x$  (i.e., the top right corner of the solid blue square in Fig. 1), will always be the furthest point from the origin even after transformation. Its distance from the origin will be unaffected by the rotation  $S$ , only by the translation by  $x_e^{\text{fid}}$ , thus we can exclude points further than  $\|x_e^{\text{max}} - x_e^{\text{fid}}\|$  in the mode parameter space. This leads to the physicality region depicted by the dashed circle, and we will refer to this as the ‘‘hypersphere’’ prior. One can recognize this as the second MH08 prior in Eq. (5) by noting that integral over  $z$  there is analogous to the sum that appears in the vector norm,  $\|x_e^{\text{max}} - x_e^{\text{fid}}\|$ .

The intersection of these two priors is shaded orange, and corresponds to the full MH08 physicality region. Note that although the two priors do remove some of the allowed unphysical regions of parameter space, they can never remove all of them. As we will see in Sect. 5, this unavoidably allowed unphysical region has undesirable consequences.

<sup>3</sup> Pandolfi et al. (2011) use a physicality prior such that the ionization fraction must be physical at all redshifts, which is equivalent to (incorrectly) limiting  $m_1$  assuming  $m_2 = 0$ . This rules out value of  $m_1$  which are not inherently unphysical, and thus potentially biases results.

<sup>4</sup> Note that in dimensions higher than the two depicted in Fig. 1, this region does not necessarily have equal edge length in all directions, and is thus truly a hyperbox as opposed to a hypercube.



**Fig. 1.** Geometric picture of the PCA physicality priors in two-dimensions, demonstrating why they necessarily allow unphysical regions of parameter space. The solid blue square is the physicality region in terms of the ionization fraction,  $x_1$  and  $x_2$ , at individual redshifts. The solid orange square is the physicality region in terms of the mode amplitudes,  $m_1$  and  $m_2$ . The two types of physicality priors corresponding to Eq. (4) and Eq. (5) are shown as the dashed square and dashed circle, respectively. Their intersection, shaded orange, is the full physicality region. The *left panel* takes a fiducial ionization history,  $x^{\text{fid}} = 0.15$ , consistent with [Mortonson & Hu \(2008a\)](#) and subsequent works. The *right panel* shows that the undesired but allowed unphysical parameter space can be reduced (but not fully) with a better choice of fiducial model.

While the amount of unphysical parameter space allowed by the hyperbox prior does not depend on our choice of  $x^{\text{fid}}$ , this is not the case for the hypersphere prior. Indeed, the geometric picture makes it clear that we could reduce the allowed unphysical regions by picking  $x^{\text{fid}} = x^{\text{max}}/2$ , i.e., placing the fiducial model at the center of the original hypercube<sup>5</sup>. This is depicted in the right panel of Fig. 1. This choice is optimal in the sense that no other choice of  $x^{\text{fid}}$  excludes as much of the unphysical region, given the two priors we have constructed. As we will see in Sect. 3, using the optimal physicality region is also useful because the increased symmetries in this case allow an analytic calculation of the induced  $\tau$  prior.

The physicality prior obtained by taking  $x^{\text{fid}}$  to be in the center of the physicality region is objectively better than other choices, and should be used. It does not bias the analysis in any other way, in fact it prevents biases from an over-allowance of unphysical parameters. We will show in Sects. 4 and 5 that the “sub-optimality” of the MH08 prior has a significant impact on constraints from data.

#### 2.4. HS17 model

As discussed in the previous section and elsewhere (e.g. [Mortonson & Hu 2008a](#)), one draw-back of the PCA model is the necessary inclusion of unphysical parameter space in the prior. In some sense this is because the PCA procedure rotates the parameters from a space in which the prior is easy to describe to one in which the likelihood is easy to describe. If we

are in a regime where the prior is completely uninformative, this is a very beneficial rotation; here, however, this is not quite the case.

One solution is thus to not perform the PCA rotation at all, and keep the values of ionization fraction at some given redshifts (the  $x_i$ ) as the parameters. Then we can always fully and sufficiently enforce physicality by requiring that  $0 < x_i < x_e^{\text{max}}$  for all  $z$ . Variations on this avenue have been explored, for example, by [Colombo & Pierpaoli \(2009\)](#), [Lewis et al. \(2006\)](#), and [Hazra & Smoot \(2017\)](#). In this section we consider for definiteness the latter model, which we denote the HS17 model.


The HS17 model takes uniform priors on the values of  $x_e(z)$  at  $z = 6, 7.5, 10$ , and  $20$ , with end-points at  $x_e(5.5) = x_e^{\text{max}}$  and  $x_e(30) = 0$  fixed, and interpolates everywhere in between these knots using piecewise cubic Hermite interpolating polynomials (PCHIP). Although the choice is ad-hoc and made somewhat a posteriori (which makes judging the statistical evidence for this model difficult), it does likely capture much of the observable features of the reionization history. A larger problem, however, is that it creates an undesired prior distribution, shown in the middle panel of Fig. 2. While the prior on  $x_e(z)$  at the  $z$ -values of the knots is uniform, it becomes triangular at the mid-point between knots, leading to the odd patterns seen in this figure. With a simple modification of this model, which we give in the next section, we are able to remove this oddly patterned prior.

#### 2.5. FlexKnots

To remedy the non-idealities of both the PCA and HS17 models discussed thus-far, we develop the following model and analysis procedure, inspired by the primordial power spectrum reconstruction in the *Planck* analysis ([Vázquez et al. 2012](#); [Planck Collaboration XX 2016](#)).

<sup>5</sup> Equivalently, we could modify the hypersphere prior to be centered on the physicality region rather than on the origin, but we choose to discuss this in terms of picking a different  $x^{\text{fid}}$  simply for convenience so that Eq. (5) is unchanged.





**Fig. 2.** One thousand reionization histories sampled from their prior distribution given the different models we consider. In the PCA case, the prior over PCA modes is taken to be uniform within the physicality region; in the two knot cases, the prior is uniform on the position and/or amplitude of each knot (thus none of these priors correspond to a flat  $\tau$  priors). The *left panel* shows the extent to which the PCA procedure allows unphysical models, even when using the optimal physicsity priors. The *middle panel* shows that the HS17 prior creates odd features and tighter peaked priors in the regions between the knot locations (which are indicated with arrows). The *right panel* is the prior for the FlexKnot model and represents our best solution for creating a reasonable and “homogeneous” prior.

The model, which we call the FlexKnot model, has knots which can move left and right in addition to up and down. Our parameters are thus a set of  $x_i$  and  $z_i$ , with uniform priors across the ranges

$$6 < z_i < 30, \quad (7)$$

$$0 < x_i < x_e^{\max}. \quad (8)$$

Additionally, given any set of  $z_i$ , we compute the reionization history by first sorting them before interpolating between the knots (see also Appendix A2 of Handley et al. 2015b). We perform interpolation using the PCHIP scheme, similarly as to what is done in HS17.

Samples from this prior with 5 knots are shown in the right-most panel of Fig. 2. We see that the prior distribution is much more uniform than the HS17 prior, with no clustering around any particular redshift or ionization fraction value. This is the effect of the left/right freedom of knots, coupled with the sorting procedure.

A final question for the FlexKnot model remains, mainly how many knots to allow? We follow Planck Collaboration XX (2016) in marginalizing over the number of knots by computing the Bayesian evidence for each of  $n$  knots,

$$\mathcal{Z}_n = \int d\theta \mathcal{L}_n(\theta | d) \mathcal{P}_n(\theta), \quad (9)$$

where  $\mathcal{L}_n$  and  $\mathcal{P}_n$  are the likelihood and prior given  $n$ , and computing the posterior distribution marginalized over  $n$ ,

$$\mathcal{P}(y) = \sum_{n=1}^N \frac{\pi_n \mathcal{P}_n(y)}{\mathcal{Z}_n}. \quad (10)$$

Here,  $y$  can represent the ionization fraction history,  $\tau$ , or any other derived quantity of interest. The  $\pi_n$  are the prior weights which we give to a model with  $n$  knots; we set this equal to one, i.e. we assume that every number of knots is equally probable. The evidence computation is performed in practice with POLYCHORD (Handley et al. 2015a,b).

Given the necessity of choosing the  $\pi_n$ , one should ask: have we really gained anything by this procedure, or have we just swapped the requirement of making one choice of prior for another? To answer this, note that although some choice still remains, we have reduced the problem of picking a functional prior

over all possible histories to one of simply picking a prior over which integer number of knots to take. This is a massive reduction in how arbitrary a choice we are required to make, and represents the key strength of this procedure. Additionally, as we will show in the next section, the data themselves largely disfavor the need for anything beyond one or two knots anyways, thus we are not particularly sensitive to the  $\pi_n$ , as long as a reasonable choice is made.

### 3. Implicit prior on $\tau$

#### 3.1. Calculating the prior

Although the details of how reionization proceeds are of significant physical interest, it is mostly the total overall integrated optical depth to reionization,  $\tau$ , that plays an important role in terms of CMB constraints. This is because  $\tau$  modulates the amplitude of temperature and polarization power spectra at multipoles greater than a few tens by  $e^{-2\tau}$ , and is thus a source of degeneracy for the scalar amplitude,  $A_s$ , which does the same. The parameter  $A_s$  itself is then degenerate with a number of other physical quantities of interest via various means (Planck Collaboration Int. LI 2017). It is thus important to examine what these different methods of reconstruction have to say in terms of  $\tau$ .

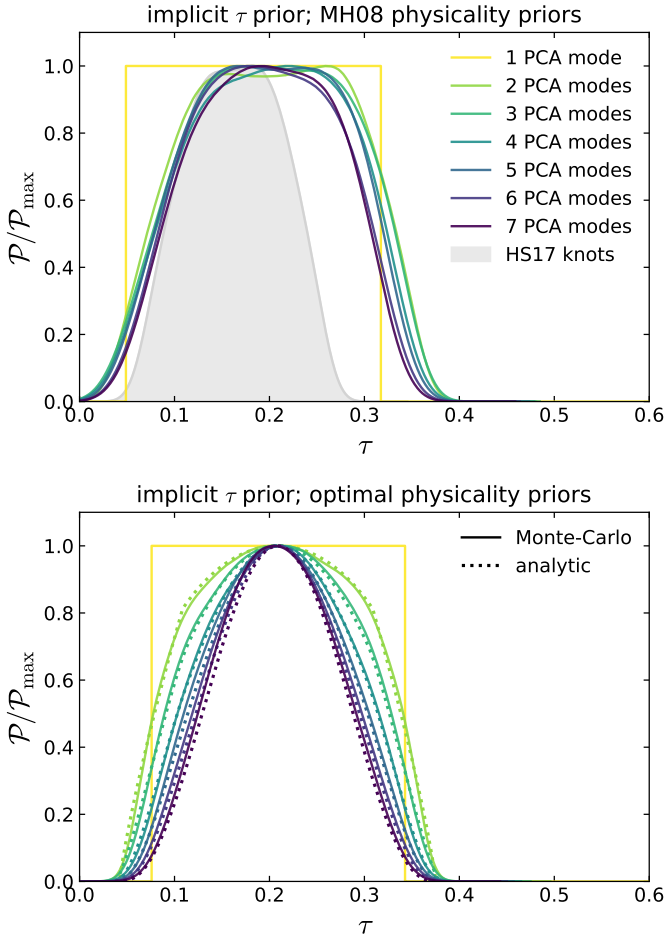
In particular, what prior over  $\tau$  do the different approaches assume? Take, for example, the PCA model. Because the modes contribute linearly to  $x_e(z)$  and because  $\tau$  is in turn a linear functional of  $x_e(z)$ , each mode contributes linearly to the total optical depth,

$$\tau(\{m_i\}) = \tau_{\text{fid}} + \sum_i \left( \frac{d\tau}{dm_i} \right) m_i, \quad (11)$$

where  $\tau_{\text{fid}}$  is the optical depth corresponding to  $x_e^{\text{fid}}$  and

$$\frac{d\tau}{dm_i} = \sigma_T \int dz \frac{n_H(z)(1+z)^2}{H(z)} S_i(z). \quad (12)$$

In all analyses to-date, the prior taken on the  $m_i$  has been uniform inside of the physicality region. We can compute the  $\tau$  prior induced by this choice by propagating the  $m_i$  prior to  $\tau$  via



**Fig. 3.** Solid colored lines: implicit prior on  $\tau$  from the PCA model when taking a flat prior on the mode amplitudes, computed via Monte Carlo. The *top panel* assumes the MH08 physicality region, whereas the *bottom panel* assumes our optimal physicality region. In the case of the *bottom panel*, analytic calculation of the prior is possible via Eq. (15) and shown as dotted lines. The *top panel* also shows the implicit prior for the HS17 model.

Eq. (11). If the  $m_i$  are uncorrelated in the prior (which is a decent approximation if using the MH08 physicality priors where the hypersphere prior plays a sub-dominant role) the induced  $\tau$  prior will be a convolution of a number of “top-hat” functions, which should tend toward some smooth distribution. The exact solution (including both hyperbox and hypersphere) can be obtained numerically by Monte Carlo sampling from the  $m_i$  prior and computing  $\tau$  for each sample. This is shown in the top panel of Fig. 3. Indeed, we find that the prior is centered on  $\tau \approx 0.2$  and has width around  $\sigma(\tau) \approx 0.07$ , tightening and smoothing out as more modes are added. The mean prior value of  $\tau \approx 0.2$  is the optical depth corresponding to an ionization level of  $x_e = x_e^{\max}/2$  throughout the entire  $z = [6, 30]$  reconstruction region, as this is what is given by the mean mode values,  $m_i = (m_i^+ + m_i^-)/2$ , independent of fiducial model (here ignoring the hypersphere prior, which is a good approximation in terms of the mode mean).

It is important to note the existence of this prior when comparing constraints on  $\tau$  from two different models, for example the TANH model (which takes a flat  $\tau$  prior) and the PCA model. In such comparisons, we change not only the model but also implicitly change the prior. Given that the prior we see in Fig. 3 tends toward quite high values of  $\tau$ , it should not be entirely surprising if we find a higher  $\tau$  posterior in the PCA case. In

Sect. 4 we will quantify the impact this has on the Heinrich et al. (2017) analysis.

What about the induced prior in the case of the optimal physicality prior instead of the MH08 one? Here, the extra symmetries of the problem allow us to derive the following useful analytic solution. First, consider the simplified case of computing the  $\tau$  prior induced by only the hypersphere prior. This is given by

$$\mathcal{P}(\tau) = \int_{\mathbb{S}} \prod_{i=1}^N dm_i \delta[\tau - \tau(\{m_i\})], \quad (13)$$

where the region  $\mathbb{S}$  is an  $N$ -dimensional hypersphere with radius  $r \equiv x_e^{\max}/2$  and centered on the origin, and  $\delta$  is the Dirac  $\delta$ -function. In geometric terms, Eq. (13) calculates the volume of the intersection of the hyperplane defined by  $\tau - \tau(\{m_i\}) = 0$  with the hypersphere  $\mathbb{S}$ . This intersection is itself an  $N - 1$  dimensional hypersphere. Its radius,  $\rho$ , is smaller than  $r$  due to the hyperplane having sliced through  $\mathbb{S}$  at a position displaced from its center. The displacement distance is controlled by  $\tau$  and given by  $D = |\tau - \tau_{\text{fid}}|/|\mathbf{g}|$ , where we have defined a vector normal to the hyperplane,  $(\mathbf{g})_i \equiv d\tau/dm_i$ . This allows us to compute the smaller radius,  $\rho^2 = r^2 - D^2$ , and, using the formula for the volume of a hypersphere, we arrive at the final answer,

$$\mathcal{P}(\tau) \propto \frac{\pi^{\frac{N-1}{2}}}{\Gamma(\frac{N}{2} - 1)} \rho^{N-1} \propto \left[ 1 - \frac{(\tau - \tau_{\text{fid}})^2}{r^2 |\mathbf{g}|^2} \right]^{\frac{N-1}{2}}. \quad (14)$$

Note that, as written, neither of these forms are a properly normalized probability distribution (however, this is not a requirement for our purposes, and the proper normalization is straightforward to compute if desired).

We now need to include the fact that the hyperbox prior excludes some of the hypersphere prior, as represented in two dimensions in the right panel of Fig. 1. In general, this breaks the symmetries which made Eq. (14) so simple. We find, however, that because  $d\tau/dm_1$  is much larger than any of the other derivatives, it is predominantly only the exclusion in the  $m_1$  direction which needs to be accounted for. This leaves enough symmetry that the resulting solution is simple enough to be useful.

The above calculation needs to be amended in the following way. Define  $\mathbb{S}'$  to be intersection of the hypersphere with the region defined by  $m_1^- < m_1 < m_1^+$ ; i.e.,  $\mathbb{S}'$  is a hypersphere with two opposite hyper-spherical caps removed. As  $\tau$  is increased or decreased from  $\tau_{\text{fid}}$ , changing where the hyperplane intersects with  $\mathbb{S}'$ , the intersection region initially does not also intersect with the caps, leaving the solution unchanged from before. Once it does intersect with the caps, the intersection region is now no longer an  $N - 1$  dimensional hypersphere, but rather an  $N - 1$  dimensional hypersphere with a single hyper-spherical cap removed. Thus, the induced prior will be instead

$$\mathcal{P}(\tau) \propto \frac{\pi^{\frac{N-1}{2}}}{\Gamma(\frac{N}{2} - 1)} \rho^{N-1} (1 - V), \quad (15)$$

where  $V$  is the volume of the removed hyper-spherical cap relative to the volume of the entire hypersphere. This fraction can be written in terms of the regularized incomplete beta function,  $I_x(a, b)$  (Li 2011), and is given by

$$V = \begin{cases} 0 & h < 0 \\ \frac{1}{2} I_{(2\rho h - h^2)/\rho^2} \left( \frac{N}{2}, \frac{1}{2} \right) & 0 < h < \rho \\ \left( 1 - \frac{1}{2} I_{(2\rho h - h^2)/\rho^2} \left( \frac{N}{2}, \frac{1}{2} \right) \right) & h > \rho, \end{cases} \quad (16)$$

where  $h$  is the height of the cap,  $h = \rho - m^{(+)} \csc \theta + D \cot \theta$ , and  $\theta$  is the angle between the hyperplane and the  $\hat{m}_1$  direction,  $\cos \theta = \mathbf{g}_1 / |\mathbf{g}|$ .

The bottom panel of Fig. 3 shows this analytic result for several values of  $N$ , alongside the Monte Carlo calculation performed similarly as before. The ‘‘kink’’ near  $\tau \sim 0.1$  (visible particularly in the  $N=2$  curve, although present in all of them) corresponds to when the correction,  $V$ , turns on; without this, we would not recover the correct shape in the tails of the distribution. With it, however, we find very good agreement between the Monte Carlo simulations and the analytic result, particularly deep in the low- $\tau$  tail which is most important since this is the region picked out by the data. This analytic result will be useful in the next section where we discuss flattening the prior. Note also that the prior volume for  $\tau$  shrinks noticeably as compared to the MH08 physicality region; we will show in Sect. 4 that this smaller and more optimal prior volume is enough to affect constraints from *Planck* data.

Similarly as to the PCA model, both the HS17 knot model and our FlexKnot model give a non-flat  $\tau$  prior if the prior on the knot amplitudes and/or positions is taken as flat. The prior for the HS17 model is shown in the top panel of Fig. 3. The FlexKnot prior is not shown, but is qualitatively similar to the others.

### 3.2. Flattening the prior

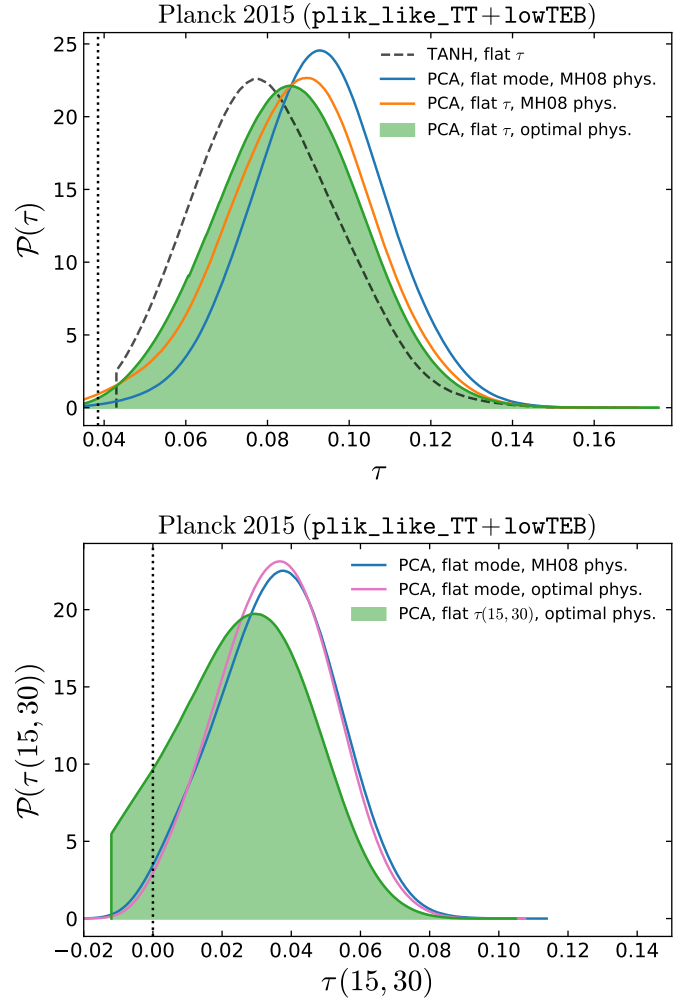
Having ascertained that all of the generic reconstruction methods we have discussed implicitly place a non-flat prior on  $\tau$ , we now discuss how to modify these analyses so that any desired  $\tau$  prior can be used, in particular a uniform one. This will be useful in performing more consistent comparisons between models by using the same prior on  $\tau$  for all models.

Intuitively, if we have an MCMC chain for an analysis that was performed with some particular  $\tau$  prior, we can importance sample the chain to give additional weight to samples with a low prior probability for  $\tau$ , in effect counter-balancing the original prior and forcing it to be flat. More explicitly, we can create a posterior where we have assumed a flat  $\tau$  prior by computing

$$\mathcal{P}^{\text{flat-}\tau}(\theta | d) = \frac{\mathcal{P}^{\text{orig}}(\theta | d)}{\mathcal{P}^{\text{orig}}(\tau(\theta))} \propto \mathcal{L}(d | \theta) \underbrace{\left[ \frac{\mathcal{P}^{\text{orig}}(\theta)}{\mathcal{P}^{\text{orig}}(\tau(\theta))} \right]}_{\mathcal{P}^{\text{flat-}\tau}(\theta)}, \quad (17)$$

where the posterior  $\mathcal{P}^{\text{orig}}(\theta | d)$  given data  $d$  is the original posterior that did not assume a flat  $\tau$  prior,  $\mathcal{P}^{\text{orig}}(\theta)$  and  $\mathcal{P}^{\text{orig}}(\tau(\theta))$  are the original priors in terms of all parameters  $\theta$  and the induced  $\tau$  prior, respectively, and  $\mathcal{L}(d | \theta)$  is the likelihood of the data given parameters. In the case of the PCA model with optimal physicality priors,  $\mathcal{P}^{\text{orig}}(\tau(\theta))$  is conveniently calculable analytically and given in Eq. (15). For other models, it can be obtained via Monte Carlo simulations and a smooth function can be fit (e.g. here we use kernel density estimates with tools from Lewis & Xu 2016).

We should note that this procedure is not unique, and many  $\mathcal{P}^{\text{flat-}\tau}(\theta)$  exist corresponding to a flat prior on  $\tau$ . However, it can be rigorously shown that the choice we have made in Eq. (17) maximizes the possible entropy in  $\mathcal{P}^{\text{flat-}\tau}(\theta)$ , subject to the condition that the resulting  $\tau$  prior is flat (Handley & Millea 2018, we also give a simplified but less general proof of this in Appendix A). This is to say that Eq. (17) is very well motivated because it is the choice that takes the original prior and modifies it in such a way that the induced  $\tau$  prior is flat while introducing the minimal amount of new information. We thus make use of this procedure in the following section.



**Fig. 4.** Constraints on the total optical depth (*top panel*), and on the optical depth between  $z=15$  and  $z=30$  (*bottom panel*), when using 2015 PlanckTT+lowTEB data. Various reionization models and priors are used (labeled in each plot), in addition to marginalization over other  $\Lambda$ CDM parameters. *Top panel*: apparent preference for higher  $\tau$  when using the PCA model as reported in Heinrich et al. (2017) was significantly impacted by the choice of prior. *Bottom panel*: evidence for early reionization was also partly impacted, although some preference remains. We consider the shaded green contour, which corresponds to a flat  $\tau$  prior and optimal physicality region, as the most well-motivated constraint from this dataset.

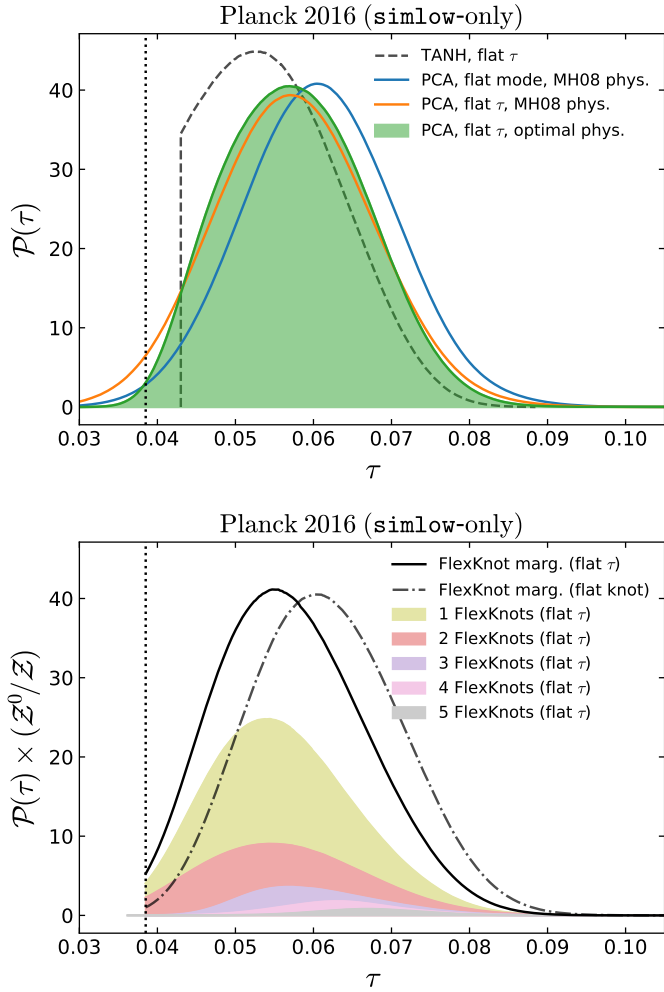
## 4. Hints of early reionization in *Planck* 2015 data?

In this section, we consider the impact of the priors discussed thus far on claims in Heinrich et al. (2017) and Heinrich & Hu (2018) for evidence of early reionization in the *Planck* 2015 data.

Heinrich et al. (2017) show that the value of  $\tau$  inferred from the *Planck* 2015 data is almost  $1\sigma$  higher under the PCA model than with the TANH model. We reproduce<sup>6</sup> this result in the top panel of Fig. 4, where the black and blue lines there correspond to the black and blue lines in Fig. 5 of Heinrich et al. (2017). When we use a flat prior on the modes and the MH08 physicality region as they have, we find very good consistency with their results. We expect that some of the shift to higher  $\tau$  is due to the

<sup>6</sup> One small difference is that Heinrich et al. (2017) use the full TT, TE, and EE data, whereas here we use only the ‘‘robust’’ TT data. We have checked that this makes a very small difference to the conclusions in this section, as expected since the TT data is much more constraining.





**Fig. 5.** Constraints on the total optical depth when using the *Planck* intermediate *simlow* likelihood alone. Various reionization models and priors are used (labeled in each plot), with other  $\Lambda$ CDM parameters fixed to their best-fit values from *Planck*TT data. In the *bottom panel*, the height of each FlexKnot model is proportional to its Bayesian evidence, with an overall scaling so that their sum (the black curve) gives a properly normalized probability distribution. This curve corresponds to marginalization over the number of FlexKnot to use.

high  $\tau$  prior implicit in the flat mode prior; to quantify the exact effect we flatten the  $\tau$  prior using the procedure described in the previous section and show the result in orange. Indeed, the shift to higher  $\tau$  is partly reduced. We also switch to the optimal physicality region, and show the final result in shaded green. This further reduces the inferred value of  $\tau$ , which now agrees much better with the TANH result. We explicitly find

$$\tau = 0.080 \pm 0.017 \quad (\text{TANH, flat } \tau; \text{ P15}) \quad (18)$$

and

$$\tau = 0.085 \pm 0.018 \quad (\text{PCA, flat } \tau, \text{ optimal phys.}; \text{ P15}) \quad (19)$$

which is a shift of only  $0.3\sigma$ . We stress that the difference between this and the nearly  $1\sigma$  shift found by [Heinrich et al. \(2017\)](#) is due purely to a different choice of priors over reionization histories inherent in the PCA procedure.

Although highlighting the shift in  $\tau$ , [Heinrich et al. \(2017\)](#) do not make the claim that one should regard the value of  $\tau$  inferred from the PCA procedure as model independent; the results of

the previous paragraph should make that point very clear. However, a strong claim is made, further argued in [Heinrich & Hu \(2018\)](#) and [Obied et al. \(2018\)](#), that the PCA procedure provides reionization-model-independent proof of early reionization because the value of  $\tau(15, 30)$  is greater than zero at  $2\sigma$ .

The derived parameter  $\tau(15, 30)$  is not fundamentally different from the derived parameter  $\tau$ , and thus its posterior distribution will qualitatively depend on our choice of priors in the same way. In the bottom panel of Fig. 4 we perform a similar set of tests for  $\tau(15, 30)$  as we just have for  $\tau$ . The blue line is the PCA result with flat mode prior and the MH08 physicality region. We find it is higher than zero by  $2.1\sigma$ , again in good agreement with the [Heinrich et al. \(2017\)](#) result. We switch to the optimal physicality region in purple and additionally to a flat prior on  $\tau(15, 30)$  in shaded green, yielding,

$$\tau(15, 30) = 0.027 \pm 0.019 \quad (20)$$

(PCA, flat  $\tau(15, 30)$ , optimal phys.; P15)

This represents a decrease in the evidence for non-zero  $\tau(15, 30)$  from  $2.1\sigma$  to  $1.4\sigma$ . Thus, similarly as before for  $\tau$ , we do not find robustness to choice of priors in this detection of early reionization. Evidently, part of the evidence for early reionization found by [Heinrich et al. \(2017\)](#) is actually due to the choice of prior rather than being driven by the data.

[Heinrich & Hu \(2018\)](#) make the contrary argument, claiming that the PCA evidence for early reionization is robust to the choice of prior. Here, we point out some problems with the arguments therein. First, they show that increasing the  $z_{\text{max}}$  of the reconstruction does not reduce preference for early reionization, arguing that this shows priors are uninformative. In fact, their results actually highlight the opposite. In particular, increasing the  $z_{\text{max}}$  of the reconstruction increases the mean value of the prior on  $\tau(z, z_{\text{max}})$  for all  $z$ . As a result, one would expect the mean value of the posterior of this quantity to increase with higher  $z_{\text{max}}$ , which is exactly what is seen in all cases in Table I of [Heinrich & Hu \(2018\)](#). Furthermore, they perform a likelihood ratio test (which is, by definition, independent of priors), finding that the  $\chi^2$  is decreased by 5–6 with a five-parameter PCA model as compared to the one-parameter TANH model. Without a quantitative argument as to the effective degrees of freedom of the PCA model that are constrained by the data (which is not made), it is impossible to judge the significance of the decrease in  $\chi^2$ . Conversely, we consider the very direct tests performed here of computing and flattening the prior to be more conclusive as to the impact of these priors on *Planck* 2015 data.

Of course, some small hints of non-zero  $\tau(15, 30)$  remain even in the case of the constraints shown in Eq. (20). However, we will show in the next section that these hints are consistent with arising from a noise fluctuation, since they disappear almost entirely with the addition of the newest *Planck* large-scale polarization data.

## 5. Reionization from *Planck* intermediate data

The *Planck* polarization likelihood used in the previous section and by [Heinrich et al. \(2017\)](#) and [Heinrich & Hu \(2018\)](#) is the most recent public *Planck* likelihood available and is based on polarization maps from *Planck*-LFI data ([Planck Collaboration II 2016](#)). Reionization constraints from lower-noise HFI data were presented in an intermediate release ([Planck Collaboration Int. XLVI 2016](#)). Additionally, [Planck Collaboration Int. XLVII \(2016\)](#) also explored various parametric reionization models extending beyond just the TANH



case. In general, good consistency with a near-instantaneous transition was found, and the optical depth constraints tightened and shifted to lower central values of  $\tau \approx 0.055\text{--}0.060$ , with  $\sigma(\tau) \sim 0.01$ . Although no evidence for early reionization was found, no completely generic reconstruction was performed, and the parametric models considered did not fully accommodate an early component should one have been present (Heinrich et al. 2017). We thus give here results from a fully generic reconstruction of the reionization history using the `simlow` likelihood.

So as to most clearly judge the impact of the large-scale polarization data, we compute constraints using only `simlow`. To do so, we fix non-reionization cosmological parameters to their best-fit values from the high- $\ell$  TT data, in particular holding the quantity  $10^9 A_s e^{-2\tau}$  fixed (which better approximates the impact of the TT data as compared to holding just  $A_s$  fixed). The black dashed contour in the top panel of Fig. 5 shows constraints on  $\tau$  from `simlow` assuming the TANH model, as also presented in Planck Collaboration Int. XLVI (2016) and Planck Collaboration Int. XLVII (2016). Note that a significant part of the posterior density is cut off by Gunn-Peterson bound requiring full reionization by  $z=6$ , which for the TANH model translates to  $\tau > 0.0430$ . Due to the finite width of reionization assumed in the TANH case ( $\Delta z = 0.5$ ), this is slightly higher than the absolute minimum possible for instantaneous reionization at  $z = 6$ , which is  $\tau > 0.0385$ . The remaining lines show results using the PCA model and various choices of prior. As before, when flattening the  $\tau$  prior we see a downward shift in  $\tau$ . Switching to the optimal physicality priors now mostly has the effect of removing some weight at  $\tau < 0.0385$ , which before was unphysically allowed by the MH08 region. Switching to the optimal physicality region remedies some of this effect and gives

$$\tau = 0.057 \pm 0.009 \quad (\text{PCA, flat } \tau, \text{ optimal phys.; simlow}) \quad (21)$$

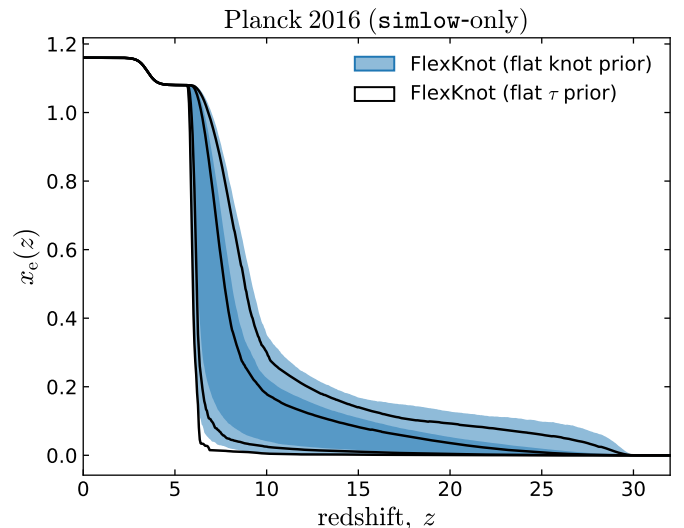
This is in good agreement with the TANH result, as well as those from the parametric models of Planck Collaboration Int. XLVII (2016).

In the bottom panel of Fig. 5, we show constraints from the FlexKnot model. In all cases except for the dot-dashed line, we use a flat prior on  $\tau$ . The posteriors on  $\tau$  with varying numbers of knots are shown as the shaded contours. Each posterior has been normalized to unity then multiplied by  $\mathcal{Z}^0/\mathcal{Z}_i$ , where  $\mathcal{Z}_i$  is the evidence for that number of knots and  $1/\mathcal{Z}^0 = \sum_i 1/\mathcal{Z}_i$ . This means that the height of each curve is proportional to the evidence of that model, and makes it easy to see that the 1-knot model has most evidence, and the evidence decreases for all subsequent models. Summing up these curves produces the posterior on  $\tau$  marginalized over the number of knots, which is shown in black, and is properly normalized to unity by the choice of  $\mathcal{Z}^0$ . In this case we find

$$\tau = 0.058 \pm 0.009 \quad (\text{FlexKnot, flat } \tau; \text{ simlow}). \quad (22)$$

This as well is in very good agreement with the results above, demonstrating the lack of evidence in the `simlow` data for anything other than a single late and near-instantaneous transition to a fully ionized universe.

Even more directly, we show in Fig. 6 the posterior constraints on the history itself using the FlexKnot model, marginalized over the number of knots. Here we see the tight restriction on any amount of non-zero ionization fraction at high redshifts, independent of exactly which  $\tau$  prior we use. Indeed, the posterior on the  $z > 15$  contribution to the optical depth is



**Fig. 6.** Constraints on the ionization fraction as a function of redshift from `simlow`-only data, using either a flat prior on the knot positions or a flat prior on  $\tau$ . The blue and black contours represent the middle 68% and 95% quantile of the posterior at each redshift.

bounded to be

$$\tau(15, 30) < 0.015 \quad (95\%) \quad (23)$$

(FlexKnot, flat  $\tau(15,30)$ ; `simlow`)

Although we consider the FlexKnot result more robust than the PCA one, we also compute the PCA bound on this quantity to stress that the disappearance of the hints of early reionization in the intermediate *Planck* data is more related to the data rather than the model used. Even in the PCA case, we find

$$\tau(15, 30) < 0.028 \quad (95\%) \quad (24)$$

(PCA, flat  $\tau(15,30)$ , optimal phys.; `simlow`)

which is weaker but still largely rules out the central value of even the prior-adjusted results given by *Planck* 2015 data (Eq. (20)).

Of course, the FlexKnot model is not immune to choice of  $\tau$  prior. For comparison, we show as the dot-dashed line in Fig. 5 the FlexKnot result, marginalized over knots, but with a flat prior on the knot positions and amplitudes instead of a flat prior on  $\tau$ . We find

$$\tau = 0.062 \pm 0.009 \quad (\text{FlexKnot, flat knot; simlow}) \quad (25)$$

which is  $0.4\sigma$  higher than the result with a flat  $\tau$  prior (Eq. (22)). We regard this difference as quantifying an unavoidable systematic uncertainty in  $\tau$  due to imperfect knowledge of the model parameter priors which is inherent in fully generic reconstructions of the reionization history. In the next section, we consider the impact of this uncertainty on other cosmological parameters that are correlated with  $\tau$ , both from current and future data.

## 6. Impact on future data

We now turn to forecasting the impact of reionization uncertainty on future data. There are two questions we would like to answer. The first is whether the increased error bars on  $\tau$  when performing a generic reconstruction cause any significant degradation in constraints on other parameters of interest. In the case of Fig. 5

we see around a 25% larger error bar in the FlexKnot case as compared to the TANH case<sup>7</sup>. We will check the impact of this degradation, as well as forecasting whether this can be improved upon with better data. Secondly, we will take the shift in the  $\tau$  posterior between a flat knot prior and a flat  $\tau$  prior and propagate it into shifts on other parameters of interest given future data, to check the extent to which these future constraints are prior-dependent. In our discussion, we focus on the sum of neutrino masses,  $\Sigma m_\nu$ , as it is the cosmological parameter expected to be most impacted by the details of reionization (Allison et al. 2015).

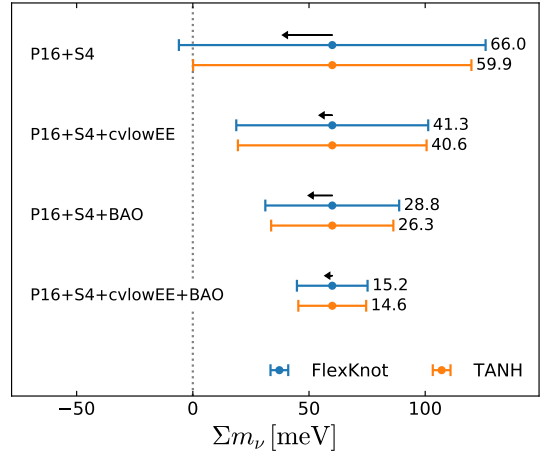
The first dataset we consider is a combination of CMB-S4 with existing *Planck* data (including the intermediate `simLow` likelihood). We perform a Fisher forecast using the standard procedure (Dodelson 2003), using a fiducial *Planck* 2015-like cosmology, with  $\tau = 0.055$  and  $\Sigma m_\nu$  at the minimum value of 60 meV. For CMB-S4, we assume a temperature noise level  $\Delta T = \Delta_P / \sqrt{2} = 1 \mu\text{K-arcmin}$ , a fraction of covered  $f_{\text{sky}} = 0.5$ , and a beam size of  $\theta_{\text{FWHM}} = 3 \text{ arcmin}$ . For the noise levels of the reconstructed lensing deflection power spectrum, we use the noise power spectrum calculated by Pan & Knox (2015)<sup>8</sup> based on an iterative quadratic EB estimator (Okamoto & Hu 2003; Smith et al. 2012). We assume the largest scales measurable by CMB-S4 are  $\ell_{\text{min}} = 50$ . This is combined with the constraints on  $\tau$  from `simLow` discussed in the previous section, which are treated as independent. In combining with the remaining *Planck* 2015 data, to take into account correlations between *Planck* and CMB-S4 due to sky coverage and multipole overlap, we compute a *Planck*-like Fisher forecast rather than use the real *Planck* constraints. The temperature and polarization noise levels assumed for *Planck* are those given by Planck Collaboration (2006), taking an optimal noise weighting of the 100, 143, and 217 GHz, channels, and a sky coverage of  $f_{\text{sky}} = 0.75$ . These reproduce the uncertainties on parameters from the real *Planck* data to within around 15%, which should be good enough for our purposes.

For this first dataset, we forecast a  $1\sigma$  error bar on  $\Sigma m_\nu$  of 59.9 meV assuming the TANH model, or 66.0 meV with the 25% looser  $\tau$  prior coming from the FlexKnot model. This fairly modest degradation is depicted by the error bars labeled P16+S4 in Fig. 7. What about the impact from the  $0.4\sigma$  shift in  $\tau$  depending on the prior assumed? The Fisher methodology allows us to propagate this into a shift in  $\Sigma m_\nu$  via

$$\frac{\Delta m_\nu}{\sigma_{m_\nu}} = \rho \frac{\Delta \tau}{\sigma_\tau}, \quad (26)$$

where  $\rho$  is the correlation coefficient between  $\tau$  and  $\Sigma m_\nu$ . For the P16+S4, we find  $\rho = 0.8$ , thus the shift in  $\Sigma m_\nu$  is  $0.3\sigma$ . This is depicted by the black arrow in Fig. 7. It is arbitrarily chosen to point to lower  $\Sigma m_\nu$ , which is the direction of shift we would expect when going from a flat knot prior to a flat  $\tau$  prior.

The existence of this shift is a main conclusion from this work, and has not been considered before. To the extent that both flat knot and flat  $\tau$  priors are reasonable, this can be considered as an extra source of “model uncertainty” in the future CMB determination of neutrino mass. The magnitude of the effect is not disastrous, but not negligible either. We comment now on a few avenues to reduce its impact.



**Fig. 7.** Fisher forecasted  $1\sigma$  error bars on the sum of neutrino masses,  $\Sigma m_\nu$ , assuming a fiducial value at the minimum allowed sum of 60 meV. Different datasets are labeled on the left, with P16 corresponding to *Planck* 2015 data combined with the intermediate `simLow` likelihood, S4 corresponding to CMB-S4, BAO to DESI-BAO, and `cvLowEE` to a cosmic-variance-limited full-sky EE measurement up to  $\ell = 30$ . Orange error bars use the TANH reionization model, and blue errors bars are slightly degraded due to having marginalized over all possible reionization histories via the FlexKnot model. The black arrows show the expected shift when considering a flat knot prior vs. a flat  $\tau$  prior.

The first comes from adding in other expected future constraints from measurements of the baryon acoustic oscillations (BAO) feature in the galaxy correlation function by DESI (Levi et al. 2013). We use the Fisher matrix for DESI calculated by Pan & Knox (2015), based on the sensitivities presented in Font-Ribera et al. (2014). The addition of this dataset brings us significantly closer to a guaranteed detection of the neutrino mass. Here we find a  $1\sigma$  error bar on  $\Sigma m_\nu$  of 26.3 meV or 28.8 meV, again depending on the reionization model used<sup>9</sup>. Although the error bars shrink, the addition of DESI actually increases the correlation between  $\tau$  and  $\Sigma m_\nu$  slightly to  $\rho = 0.9$ . This arises because DESI is completely insensitive to the  $\tau - A_s$  degeneracy responsible for the correlation, but reduces other physical degeneracies impacting the  $\Sigma m_\nu$  determination, thus leaving the former degeneracy more prominent. In this case, we find a  $0.35\sigma$  shift in  $\Sigma m_\nu$ , depending on the  $\tau$  prior used. While this is a larger relative shift, on an absolute scale it is about half the shift as compared to without DESI.

To reduce  $\rho$  and thus the impact of reionization uncertainty, we need data which directly constrains  $\tau$  and/or  $A_s$ . This could be provided, for example, by more precise large-scale polarization data. There is room for some improvement as the *Planck* data is not yet at the cosmic variance limit. We thus consider the limiting case of a full-sky cosmic-variance limited EE measurement across multipoles 2–30 (which we will label `cvLowEE`). This should be the upper bound of what could be achieved with next-generation satellite missions, e.g. LiteBIRD, PIXIE, CORÉ, or PICO (Matsumura et al. 2016; Kogut et al. 2011; The CORÉ Collaboration 2011). To obtain the most accurate forecasts, we run MCMC chains using the exact full-sky  $C_\ell$  likelihood (e.g. Hamimeche & Lewis 2008). This is more

<sup>7</sup> We compute the TANH error bar here by taking the standard deviation the TANH posterior. Despite that this is clearly non-Gaussian as it is cut off by the Gunn-Peterson prior, it gives us a rough estimate to use in combination with Fisher forecasts.

<sup>8</sup> Provided by the authors upon request.

<sup>9</sup> These predictions are slightly less optimistic than the ones presented in the main body of Pan & Knox (2015), mostly because we use  $\ell_{\text{min}} = 50$  for CMB-S4 as opposed to  $\ell_{\text{min}} = 2$  used there. They are in better agreement with Allison et al. (2015) which take a similar  $\ell_{\text{min}}$ , but still very slightly wider, due mostly to a lower fiducial value of  $\tau$ .

accurate than Fisher forecasts, which do not easily deal with the sharp edges of the physicality priors. Using the TANH model, we forecast an error bar of

$$\sigma_\tau = 0.0020 \quad (\text{TANH, flat } \tau; \text{cvLowEE}) \quad (27)$$

and with the FlexKnot model marginalized over the number of knots, we obtain

$$\sigma_\tau = 0.0024 \quad (\text{FlexKnot, flat } \tau; \text{cvLowEE}). \quad (28)$$

This leads to a 15% degradation of the error bar in  $\Sigma m_\nu$ , somewhat better than the 25% degradation we found previously with `simlow`. Additionally,  $\rho$  is reduced to 0.5, meaning the shift in  $\Sigma m_\nu$  due to the choice of  $\tau$  prior is now only  $0.2\sigma$ . Evidently, these improved large-scale polarization measurements would be quite valuable, not just for reducing the absolute uncertainty on  $\tau$ , but also for reducing our dependence on its prior.

When we combine all datasets discussed thus far, including both `cvLowEE` and DESI-BAO, we find an expected constraint of  $\sim 15$  meV, with the  $\tau$  prior able to shift things only by  $\sim 3$  meV. This would be enough for a  $\sim 4\sigma$  guaranteed detection of non-zero neutrino masses even if the true value is at the minimum.

## 7. Discussion and conclusions

In this work we have examined the impact of priors on generic reconstruction of the reionization history. We have pointed out some problems with PCA procedure, mainly that (1) the priors that are usually taken over the mode amplitudes correspond implicitly to a somewhat informative prior on  $\tau$  and (2) a sub-optimal set of physicality priors has been used to-date. Ultimately, the dependence of the final posterior on the details of these physicality priors argues that the PCA procedure is not well suited to this problem. This is not to say there are issues with PCA analyses in general, simply that in situations where hard priors exist and are informative (such as here), the PCA method is a sub-optimal tool. Indeed, we have shown the significant extent to which some of these priors have impacted the analysis of [Heinrich et al. \(2017\)](#) and [Heinrich & Hu \(2018\)](#), artificially increasing the significance of some hints of a high-redshift reionization component.

Even so, the PCA analysis can still be quite useful in providing a simple and convenient approximate likelihood as described in [Heinrich et al. \(2017\)](#). One can then easily project different physical models onto the principal components to obtain constraints on the physical model parameters, effectively applying the priors of the physical model and alleviating some of the problems that arise when attempting to interpret the PCA results on their own.

We have also described the FlexKnot model, which we argue is better to use in a completely generic reionization reconstruction setting since physicality can be enforced exactly and the question of how many knots to use can be self-consistently dealt with via a Bayesian evidence computation.

Using these models, we have presented the first generic reconstruction results from the *Planck* 2016 intermediate `simLow` likelihood. We find no evidence for early reionization, instead only very tight upper limits on any contribution at  $z \gtrsim 15$ . This is true even when using the as-argued sub-optimal PCA model, thus the qualitative conclusion of preference only for a late and fast reionization is quite robust to modeling choices. These results negate the need for explanations of an early reionization contribution, for example from metal-free stars ([Ahn et al. 2012](#);

[Miranda et al. 2017](#)), and can instead be used to place constraints on these scenarios.

Recently, the EDGES collaboration has detected an absorption feature in the sky-averaged spectrum, presumably due to 21 cm absorption by neutral hydrogen during the epoch of reionization ([Bowman et al. 2018](#)). In the standard picture, the upper and lower frequencies of the feature indicate that photons from early stars were numerous enough to equilibrate the gas temperature and 21cm spin temperature near  $z=20$ , and then raise the spin temperature above the CMB temperature near  $z=15$ . Due to the amplitude and shape of the detected feature, however, some non-standard explanation is required. The results discussed here can limit such possible explanations, as no modifications to the standard picture must be made which would lead to reionizing the universe to a level large enough to violate the bound on the optical depth at  $z > 15$  given in Eq. (23). For example, [Ewall-Wice et al. \(2018\)](#) discuss an explanation of the EDGES signal in which a radio background from accretion onto the black hole remnants of metal-free stars causes the absorption feature to appear deeper against the continuum than in the standard scenario. It is shown that X-rays also produced by the accretion would result in some significant reionization at  $z > 15$ , already in some tension with even the high  $\tau(15, 30)$  obtained by [Heinrich & Hu \(2018\)](#); given the tighter upper bounds we find here, this explanation is now heavy disfavored in its simplest form.

Regardless of the generic reconstruction models we consider, we find there is some sensitivity to whether a flat prior on  $\tau$  is taken, or one which is flat on PCA mode amplitudes or knot positions and/or amplitudes. Other priors for this problem exist as well, for example Jeffreys priors or reference priors ([Jeffreys 1946](#); [Bernardo 1979](#)), and these would be well worth exploring too. Nevertheless, it is unclear that a very strong argument could be made for which particular prior is the correct one. Thus here we instead take the shifts in  $\tau$  we find when switching between the priors we have considered as an unavoidable modeling uncertainty arising from the generic reconstruction.

We showed that this uncertainty can have significant impact on determination of the sum of neutrino masses from future experiments. For example, different priors can cause a shift of  $0.35\sigma$  on  $\Sigma m_\nu$  from future *Planck* + CMB-S4 + DESI-BAO measurements. One avenue for improvement would be to obtain something approaching a cosmic-variance-limited measurement of large-scale CMB polarization. Some other avenues exist as well.

For example, several direct constraints exist on the ionization state of the universe at various redshifts (summarized, e.g. by [Bouwens et al. 2015](#)). Here we have applied only one of them, mainly the stringent bound on the end of reionization from observations of the Gunn-Peterson optical depth in high-redshift quasars. We have done so in a fairly unsophisticated way by simply requiring that  $x_e(6) = x_e^{\text{max}}$ . A more detailed study of the impact that these direct constraints have is warranted.

Another approach is to use physically motivated models of reionization, rather than the generic unphysical reconstruction methods presented here. These could potentially shrink the prior parameter space further to where priors become unimportant. The challenge here is finding models which have parameters that can be marginalized over to accurately represent uncertainty as to the physics of reionization and the types of sources which can contribute ionization radiation. In addition, with physical models, we may better be able to fold in constraints from other sources, such as measurements of the patchy kinetic Sunyaev-Zeldovich effect which can further constrain



reionization (Gruzinov & Hu 1998; Knox et al. 1998; Zahn et al. 2012; Planck Collaboration Int. XLVII 2016; Park et al. 2013). We have not applied these bounds here as it is not straightforward how to do so for the generic reionization histories we consider, but in principle these could reduce our dependence on priors. Additionally, using novel analysis methods (Smith & Ferraro 2017; Ferraro & Smith 2018), experiments like CMB-S4 could potentially measure this effect to high significance. Further down the road, methods described in Liu et al. (2016) or Meyers et al. (2018) could also push beyond the CMB cosmic variance limit we have calculated and remove the uncertainty on other cosmological parameters due to  $\tau$  entirely. Future prospects are thus optimistic, although work remains to be done.

*Acknowledgements.* We thank Zhen Pan for his forecasting code which we adapted to use in this work, Torsten Ensslin, Douglas Scott, and Martin White for their helpful comments on the draft of this paper, and Silvia Galli, William Handley, and Wayne Hu for useful discussions. MM was supported by the Labex ILP (reference ANR-10-LABX-63). This work was aided by the use of several software packages, including CAMB (Lewis et al. 2000), CosmoSlik (Millea 2017), getdist (Lewis & Xu 2016), Julia (Bezanson et al. 2017), and PolyChord (Handley et al. 2015a,b).

## References

- Abazajian, K. N., Arnold, K., Austermann, J., et al. 2015, *Astropart. Phys.*, **63**, 66
- Abazajian, K. N., Adshead, P., Ahmed, Z., et al. 2016, ArXiv e-prints, [arXiv:1610.02743]
- Ahn, K., Iliiev, I. T., Shapiro, P. R., et al. 2012, *ApJ*, **756**, L16
- Allison, R., Caucal, P., Calabrese, E., Dunkley, J., & Louis, T. 2015, *Phys. Rev. D*, **92**, 123535
- Becker, R. H., Fan, X., White, R. L., et al. 2001, *AJ*, **122**, 2850
- Bernardo, J. M. 1979, *J. Royal Stat. Soc. Ser. B (Methodological)*, **41**, 113
- Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. 2017, *SIAM Rev.*, **59**, 65
- Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2015, *ApJ*, **811**, 140
- Bowman, J. D., Rogers, A. E. E., Monsalve, R. A., Mozdzen, T. J., & Mahesh, N. 2018, *Nature*, **555**, 67
- Colombo, L. P. L., & Pierpaoli, E. 2009, *New Ast.*, **14**, 269
- Dodelson, S. 2003, *Modern Cosmology* (Amsterdam: Academic Press)
- Ewall-Wice, A., Chang, T.-C., Lazio, J., et al. 2018, *ApJ*, submitted [arXiv:1803.01815]
- Fan, X., Narayanan, V. K., Strauss, M. A., et al. 2002, *ApJ*, **123**, 1247
- Ferraro, S., & Smith, K. M. 2018, ArXiv e-prints, [arXiv:1803.07036]
- Font-Ribera, A., McDonald, P., Mostek, N., et al. 2014, *JCAP*, **2014**, 023
- Gruzinov, A., & Hu, W. 1998, *ApJ*, **508**, 435
- Hamimeche, S., & Lewis, A. 2008, *Phys. Rev. D*, **77**, 103013
- Handley, W., & Millea, M. 2018, *Bayesian Analysis*, submitted [arXiv:1804.08143]
- Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015a, *MNRAS*, **450**, L61
- Handley, W. J., Hobson, M. P., & Lasenby, A. N. 2015b, *MNRAS*, **453**, 4385
- Hazra, D. K., & Smoot, G. F. 2017, *JCAP*, **11**, 028
- Heinrich, C., & Hu, W. 2018, ArXiv e-prints, [arXiv:1802.00791]
- Heinrich, C. H., Miranda, V., & Hu, W. 2017, *Phys. Rev. D*, **95**, 023513
- Hu, W., & Holder, G. P. 2003, *Phys. Rev. D*, **68**, 023001
- Jeffreys, H. 1946, *Proc. Roy. Soc. London Ser. A, Math. Phys. Sci.*, **186**, 453
- Knox, L., Scoccimarro, R., & Dodelson, S. 1998, *Phys. Rev. Lett.*, **81**, 2004
- Kogut, A., Fixsen, D. J., Chuss, D. T., et al. 2011, *JCAP*, **07**, 025
- Levi, M., Bebek, C., Beers, T., et al. 2013, ArXiv e-prints, [arXiv:1308.0847]
- Lewis, A., & Xu, Y. 2016, *Cmbant/Getdist: 0.2.6*, *Tech. rep.*, Zenodo
- Lewis, A., Challinor, A., & Lasenby, A. 2000, *ApJ*, **538**, 473
- Lewis, A., Weller, J., & Battye, R. 2006, *MNRAS*, **373**, 561
- Li, S. 2011, *Asian J. Math. Stat.*, **4**, 66
- Liu, A., Pritchard, J. R., Allison, R., et al. 2016, *Phys. Rev. D*, **93**, 043013
- Matsumura, T., Akiba, Y., Arnold, K., et al. 2016, *J. Low Temp. Phys.*, **184**, 824
- Meyers, J., Meerburg, P. D., van Engelen, A., & Battaglia, N. 2018, *Phys. Rev. D*, **97**, 103505
- Millea, M. 2017, *Astrophysics Source Code Library* [record ascl:1701.004]
- Miranda, V., Lidz, A., Heinrich, C. H., & Hu, W. 2017, *MNRAS*, **467**, 4050
- Mortonson, M. J., & Hu, W. 2008a, *ApJ*, **672**, 737
- Mortonson, M. J., & Hu, W. 2008b, *ApJ*, **686**, L53
- Obied, G., Dvorkin, C., Heinrich, C., Hu, W., & Miranda, V. 2018, *Phys. Rev. D*, **98**, 043518
- Okamoto, T., & Hu, W. 2003, *Phys. Rev. D*, **67**, 083002
- Pan, Z., & Knox, L. 2015, *MNRAS*, **454**, 3200
- Pandolfi, S., Ferrara, A., Choudhury, T. R., Melchiorri, A., & Mitra, S. 2011, *Phys. Rev. D*, **84**, 123522
- Park, H., Shapiro, P. R., Komatsu, E., et al. 2013, *ApJ*, **769**, 93
- Planck Collaboration 2006, ArXiv e-prints [arXiv:astro-ph/0604069]
- Planck Collaboration II. 2016, *A&A*, **594**, A2
- Planck Collaboration XX. 2016, *A&A*, **594**, A20
- Planck Collaboration Int. XLVI. 2016, *A&A*, **596**, A107
- Planck Collaboration Int. XLVII. 2016, *A&A*, **596**, A108
- Planck Collaboration Int. LI. 2017, *A&A*, **607**, A95
- Smith, K. M., & Ferraro, S. 2017, *Phys. Rev. Lett.*, **119**, 021301
- Smith, K. M., Hanson, D., LoVerde, M., Hirata, C. M., & Zahn, O. 2012, *JCAP*, **2012**, 014
- The CORe Collaboration (Armitage-Caplan, C., et al.) 2011, ArXiv e-prints, [arXiv:1102.2181]
- Vázquez, J. A., Bridges, M., Hobson, M. P., & Lasenby, A. N. 2012, *JCAP*, **06**, 006
- Villanueva-Domingo, P., Gariazzo, S., Gnedin, N. Y., & Mena, O. 2018, *JCAP*, **04**, 024
- Zahn, O., Reichardt, C. L., Shaw, L., et al. 2012, *ApJ*, **756**, 65



## Appendix A: Maximum entropy prior flattening

Here we will show how to construct the maximum entropy prior used in Sect. 3.2. This is the maximally uninformative prior on the input parameters of our model,  $(\theta_1, \theta_2, \dots)$ , for which the induced prior on  $\tau = \tau(\theta_1, \theta_2, \dots)$  is flat. Here, the  $\theta_i$  can represent the PCA model amplitudes, the knot positions and/or amplitudes, or any other parameters.

We first begin with a simplified example: what is the maximum entropy prior on two parameters,  $a$  and  $b$ , each with support on  $0 < a, b < 1$ , for which the prior on the sum,  $a + b$ , is flat between  $0 < a + b < 2$ ? In this analogy,  $a$  and  $b$  are the  $\theta_i$  parameters,  $a + b$  is like  $\tau$ , and the support on  $[0, 1]$  is the physicality region.

The entropy of the probability distribution  $p(a, b)$  is

$$H[p(a, b)] = - \int_0^1 da \int_0^1 db p(a, b) \ln p(a, b). \quad (\text{A.1})$$

We can express the problem of finding the  $p(a, b)$  which maximizes  $H$ , subject to some constraints, as a Lagrange multiplier problem. Our constraint is that the transformed probability distribution for the sum,  $p(a + b)$ , is flat. We will write this constraint in what may seem as an odd form, but which is conducive to solving the Lagrange multiplier problem. We will require that the moments of  $p(a, b)$  with respect to  $a + b$  are those of a flat distribution between 0 and 2. This is to say, that  $\langle a + b \rangle = 1$ ,  $\langle (a + b)^2 \rangle = 8/3$ , etc... By specifying an infinite number of moments, we guarantee our target distribution  $p(a + b)$  is exactly flat. Our constraints are thus that

$$\int_0^1 da \int_0^1 db (a + b)^n p(a, b) = c_n, \quad (\text{A.2})$$

with

$$c_n = \frac{1}{2} \int_0^2 dx x^n \quad (\text{A.3})$$

for all  $n = 0 \dots \infty$ . Note that for  $n = 0$  we simply have that the integral over  $p(a, b)$  is unity, which guarantees that it is a probability distribution. Higher moments fix  $p(a + b)$  to be flat as desired. As a Lagrange multiplier problem, we are maximizing the functional,

$$F[p(a, b)] = \int_0^1 da \int_0^1 db \left\{ -p(a, b) \ln p(a, b) + \sum_{i=0}^{\infty} \lambda_i [(a + b)^i p(a, b) - c_i] \right\}, \quad (\text{A.4})$$

where the  $\lambda_i$  are the Lagrange multipliers. Setting the variation of  $F$  with respect to  $p(a, b)$  to zero gives us that

$$p(a, b) = \exp \left[ -1 + \sum_{i=0}^{\infty} \lambda_i (a + b)^i \right]. \quad (\text{A.5})$$

Substituting this into each of the Eq. (A.2) could then be used to solve for all of the  $\lambda_i$ , giving us the unique maximum entropy solution for  $p(a, b)$ . Rather than performing this process explicitly, we will postulate the answer and show that it is the solution. Consider the probability distribution

$$p(a, b) = 1/q(a + b), \quad (\text{A.6})$$

where  $q(a + b)$  is the transformed probability distribution for  $a + b$  given the initial flat priors on  $a$  and  $b$ ,

$$q(a + b) = \int_0^1 da' \int_0^1 db' \delta((a + b) - (a' + b')). \quad (\text{A.7})$$

It is straightforward to substitute this into Eq. (A.6) and then show that  $p(a + b)$  is indeed flat between 0 and 2. This means it satisfies the infinite number of constraint equations. It remains to check that the variation of  $F$  around this function is zero, which is tantamount to showing that it can be written in the form of Eq. (A.5). Since Eq. (A.5) says that the log of our function is a Maclaurin series in  $a + b$  and since our  $p(a, b)$  is an elementary function of only  $a + b$ , this is also true. Therefore, we have shown Eq. (A.5) is the unique maximum entropy probability distribution of  $a, b$  for which the transformed distribution of  $a + b$  is flat.

The above proof trivially generalizes to  $N$  variables and to any initial support by simply adding in more integrals and changing the limits to something other than 0 and 1. Similarly, one can easily replace  $a + b$  with any desired function for a derived parameter. This proof thus applies to the case of flattening the  $\tau$  prior discussed in Sect. 3.2. For a more rigorous and general proof, see [Handley & Millea \(2018\)](#).