

Bayesian methods for fitting Baryon Acoustic Oscillations in the Lyman- α forest

Andrei Cuceu^a Andreu Font-Ribera^a Benjamin Joachimi^a

^aDepartment of Physics and Astronomy, University College London, London, UK

E-mail: andrei.cuceu.14@ucl.ac.uk, a.font@ucl.ac.uk, b.joachimi@ucl.ac.uk

Abstract. We study and compare fitting methods for the Lyman- α ($\text{Ly}\alpha$) forest 3D correlation function. We use the nested sampler `PolyChord` and the community code `picca` to perform a Bayesian analysis which we compare with previous frequentist analyses. By studying synthetic correlation functions, we find that the frequentist profile likelihood produces results in good agreement with a full Bayesian analysis. On the other hand, Maximum Likelihood Estimation with the Gaussian approximation for the uncertainties is inadequate for current data sets. We compute for the first time the full posterior distribution from the $\text{Ly}\alpha$ forest correlation functions measured by the extended Baryon Oscillation Spectroscopic Survey (eBOSS). We highlight the benefits of sampling the full posterior distribution by expanding the baseline analysis to better understand the contamination by Damped $\text{Ly}\alpha$ systems (DLAs). We make our improvements and results publicly available as part of the `picca` package.

Contents

1	Introduction	1
2	Bayesian vs. frequentist methods	2
2.1	The best fit model	3
2.2	Quantifying uncertainties	3
2.3	Nuisance parameters	5
2.4	Toy Example	5
3	Testing BAO measurements on mock data sets	6
3.1	Synthetic Correlation Functions	7
3.2	Fitting Methods	7
3.3	Results	8
4	The Full Posterior of eBOSS DR14	10
4.1	The Baseline Analysis	10
4.2	A simple extension	10
5	Conclusions	12
A	Computation of CIs	17
B	Comparison with Chan et al. 2018	18

1 Introduction

Over the last few decades, cosmology has entered a data driven era, with large surveys providing rich data sets. These data sets encode a vast amount of information, which most of the time is non-trivial to extract. Performing such large surveys is very expensive and time consuming, which puts even more emphasis on efficient and accurate extraction of meaningful information. Over the last two decades, the Λ Cold Dark Matter (Λ CDM) model has become widely accepted as the standard cosmological model, however, it has 6 free parameters with extensions adding even more. On top of this, most analyses need extra nuisance parameters to create good models of their data, leading to very high-dimensional parameter spaces (e.g. [1, 2]). It has become essential to have reliable analysis tools, and this has led to an increased focus on statistical methods and interpretation.

The efficient and reliable extraction of Baryon Acoustic Oscillation (BAO) information from large scale structure (LSS) data has been a very active topic of research over the last 15 years. Since the first detections of BAO, using the distribution of galaxies [3, 4], there has been much attention given to optimising statistical methods used on data of discrete tracers, e.g. [5–8]. However, the newer method of measuring BAO using the Lyman- α ($\text{Ly}\alpha$) forest has received comparatively less attention.

The first detection of the BAO scale in the $\text{Ly}\alpha$ forest auto-correlation ($\text{Ly}\alpha \times \text{Ly}\alpha$) function was done by the Baryon Oscillation Spectroscopic Survey (BOSS) using Data Release 9 (DR9) of the Sloan Digital Sky Survey (SDSS) [9–11]. It was also detected in the cross-correlation of $\text{Ly}\alpha$ absorbers with quasar positions ($\text{Ly}\alpha \times \text{QSO}$) using BOSS DR11 [12]. A

physical model for the contaminations was first used by [13, 14], however it has a large number of parameters that model the contamination by Damped Ly α systems (DLAs) and different metal absorption lines. Recent analyses have introduced yet more effects [15, 16], and these have led to a large parameter space and potentially complex behaviour. However, as the main aim has always been the measurement of the BAO peak position, these astrophysical parameters have received little attention.

The BAO scale can be measured using just two parameters: α_{\parallel} and α_{\perp} . These measure the size of the BAO scale relative to a fiducial cosmology along and across the line of sight, respectively. Because of the focus on measuring BAO, the high-dimensional parameter spaces have so far been investigated only using a frequentist methodology. In particular, the profile likelihood (e.g. [7, 17, 18]) has been used to extract the relevant information. This method approximates the probability at some value of α_{\parallel} and α_{\perp} by the maximum likelihood over the nuisance parameters at that point. In this work we use a Bayesian approach to fit the Ly α forest 3D correlation function and we provide a tool for studying this large parameter space and for extracting all the relevant information. While Bayesian methods have been used in Ly α forest analyses before (e.g. [19–22]), they were never used in BAO analyses of the Ly α correlation function.

The principal difference between the Bayesian [23, 24] and frequentist [25] methodologies is their interpretation of the concept of probability. In the Bayesian framework, probability is a degree of belief in an event, while in the frequentist framework, the probability of an event is the limit of its relative frequency in many trials. In the limit of infinite data, the two approaches produce the same results. Our focus, however, is not on philosophical interpretations, but on the practical consequences of the two frameworks when working with real data. One of the reasons for the widespread use of Bayesian methods is the availability of samplers such as Monte Carlo Markov Chain (MCMC) [26, 27] which facilitate the efficient exploration of the complex high-dimensional posterior distributions that often appear in cosmology. It is the efficiency and accuracy of such tools that we want to compare with equivalent frequentist approaches within the context of fitting the Ly α forest correlation function.

The purpose of this work is to investigate the methods used so far and compare them with a Bayesian framework. In particular, we use for the first time a sampler to obtain the full posterior distribution of all parameters. We begin in Section 2, where we discuss the frequentist methods used so far, and compare them with the Bayesian methodology. In Section 3, we use synthetic correlation functions to showcase the similarities and differences of the two approaches when fitting the BAO parameters. Finally, in Section 4, we use a Bayesian framework to analyse the latest extended BOSS (eBOSS) DR14 correlation functions [15, 16], and highlight some potential uses and advantages of the full posterior distribution.

2 Bayesian vs. frequentist methods

We first discuss some of the theoretical differences between the Bayesian and frequentist approaches to statistical data analysis. We include this discussion for completeness and also because, as we will see, the two frameworks answer fundamentally different questions when it comes to the quantification of uncertainties. As such, we believe that these theoretical considerations are important to the interpretation of our results. We also discuss some practical differences and their implications for cosmology and we conclude with a simple toy example.

2.1 The best fit model

The difference in the interpretation of probability between the Bayesian and frequentist approaches leads to a difference in their principal object of study. The frequentist approach treats parameters as fixed quantities and the data as the random variable. The object of study is the probability of obtaining the data, D , given a model, \mathcal{M} , and some parameters, $\vec{\theta} = (\theta_1, \dots, \theta_n)$; this is also known as the likelihood: $P(D|\vec{\theta}, \mathcal{M})$. For normally distributed data the likelihood takes the form:

$$P(D|\vec{\theta}, \mathcal{M}) = \frac{\exp\left[-\frac{1}{2}(D - \mathcal{M}(\vec{\theta}))^T \Sigma^{-1}(D - \mathcal{M}(\vec{\theta}))\right]}{\sqrt{(2\pi)^n |\Sigma|}}, \quad (2.1)$$

where Σ is the covariance matrix of the data. We will also refer to the logarithm of the likelihood which we denote by $\mathcal{L} \equiv \log P(D|\vec{\theta}, \mathcal{M})$.

The object of interest for a Bayesian is the posterior distribution of the parameters $\vec{\theta}$, given the data, and a model: $P(\vec{\theta}|D, \mathcal{M})$. This fully encapsulates our knowledge of the probability of possible values of the parameters of interest, by treating these parameters as random variables. The posterior distribution can be computed through Bayes' Theorem [23, 24]:

$$P(\vec{\theta}|D, \mathcal{M}) = \frac{P(D|\vec{\theta}, \mathcal{M})P(\vec{\theta}|\mathcal{M})}{P(D|\mathcal{M})}, \quad (2.2)$$

where $P(\vec{\theta}|\mathcal{M})$ is the prior probability, and $P(D|\mathcal{M})$ is a constant (for a model \mathcal{M}) known as the Bayesian evidence. The evidence is the normalization of the posterior and requires an n-dimensional integral to be computed. It can be used to perform Bayesian model selection (see e.g. [28–30] for applications in cosmology), however, when the only goal is inference this quantity is not necessary.

If we work with wide flat priors, we can deduce from Equations 2.1 and 2.2 that the two frameworks will produce the same best fit $\vec{\theta}_{best}$, given by the frequentist maximum likelihood $P_{max}(D|\vec{\theta}, \mathcal{M})$, and by the Bayesian maximum posterior probability $P_{max}(\vec{\theta}|D, \mathcal{M})$.

2.2 Quantifying uncertainties

The two approaches diverge again when it comes to finding the uncertainty on $\vec{\theta}_{best}$. In this case it is not only a difference in methodology, but also a fundamental difference in the object of interest. Frequentists quantify uncertainty through confidence intervals¹ (CI), which are defined by the proportion (frequency) of intervals, measured from the ensemble of possible data sets, that contain the true values of the parameters ($\vec{\theta}_{true}$). On the other hand, Bayesian uncertainty is quantified through credible regions (CR), defined as the smallest region of the posterior that encompasses a certain probability (most often 68% and 95% CRs are quoted).

The two questions asked by Bayesians and frequentists are very different. A Bayesian CR is telling us that, given our data and prior, we are e.g. 95% confident that the true values $\vec{\theta}_{true}$ are within that region. Meanwhile, the frequentist CI is telling us that if we repeat our experiment many times, the confidence intervals we obtain will contain the true values $\vec{\theta}_{true}$ in e.g. 95% of the cases. Note that the frequentist CI that we obtain from our data does not state anything about the probability that it contains $\vec{\theta}_{true}$ (this is a common misconception). In fact, there are extreme cases in the literature where a frequentist CI has 0% probability

¹For n-dimensional distributions these are referred to as confidence regions, however we chose to only talk about intervals to clearly distinguish them from Bayesian credible regions

of containing the truth² [31, 32]. This behaviour is due to the fact that in the frequentist methodology, one never conditions a result on the actually observed data D , but instead on its distribution of possible realisations, which we do not always fully understand given limited data. This issue is fundamental in cosmology, because we only have one Universe to observe. Bayesian CRs can also cause problems because of prior choices. For example, flat priors are only flat for the specific parametrisation they are defined on.

We now turn to the computation of these uncertainty intervals and regions. In a frequentist framework, we usually start by computing the best fit parameters $\vec{\theta}_{best}$ by maximising the likelihood. This is called Maximum Likelihood Estimation (MLE). After that, we can compute a covariance matrix for the parameters by taking the second derivative of the likelihood in parameter space around this peak. However, this covariance matrix is only accurate in general if the errors on the data are normally distributed and the model is linear in all parameters (which would correspond to a Gaussian posterior in a Bayesian framework). If this is not the case, it can still be applied around the peak, but a better approach is to compute the likelihood on a grid in parameter space (we will call this a scan).

The confidence intervals can be computed using regions of equal likelihood around $\vec{\theta}_{best}$. To find these regions we need $\Delta\mathcal{L}_p \equiv \mathcal{L}_p - \mathcal{L}_{min}$ values such that the region defined by \mathcal{L}_p corresponds to a certain CI of probability p . In the Gaussian case these values can be computed analytically. In the general case, a large number of Monte Carlo simulations of the data are needed to compute $\Delta\mathcal{L}_p$ values. This procedure is outlined in Appendix A. Using these $\Delta\mathcal{L}_p$ values, frequentists can draw constant likelihood contours using a scan of the parameters and obtain the correct confidence intervals.

We must stress that from a Bayesian perspective the best fit and the credible regions are just special values computed from the posterior. It is this full posterior distribution that is the real object of interest because it contains all the information about the probable values of the parameters. The computation of the posterior distribution can however be a very demanding task. When dealing with low-dimensional spaces, Bayesians can compute the posterior on a grid, similar to the frequentist method. The scan can be used to compute the Bayesian credible regions by finding the smallest region of the scan that contains a certain probability p . This is usually done by ordering the scan in decreasing order of probability and computing their running sum until the result is a fraction p of the total probability of the grid:

$$\sum_{j=1}^M P(\vec{\theta}_j|D, \mathcal{M}) = p \sum_{j=1}^N P(\vec{\theta}_j|D, \mathcal{M}), \quad (2.3)$$

where N is the total number of points on the grid, and the M points obtained through this method cover the CR of probability p . Note that the integral of the posterior over a region is normally required to do this, however, this can be approximated by a sum if the grid is equally spaced because the probability density at each point is proportional to the probability mass for that region. As we will show in the next Section, these two interpretations of the same grid results often produce identical results. However, the scan quickly becomes infeasible with increasing number of parameters, and Bayesians move on to using more effective methods, such as MCMC.

Bayesians can deal with high-dimensional spaces by efficient and accurate sampling, using tools such as MCMC [26, 27] (see e.g. [33, 34] for detailed introductions). Furthermore,

²We must stress however, that even in such extreme cases the frequentist CIs are not wrong. They are just answering a different question.

the efficient computation of the Bayesian evidence has also become possible with the introduction of Nested Sampling [35]. The underlying principle of such methods is the creation of samples from continuous random variables with probability density proportional to a known function, in our case the unnormalized posterior distribution. Once a sufficient number of samples have been generated, they can be used to compute summary statistics. For example, computing the credible regions amounts to finding the smallest region containing a certain fraction of samples. This fraction corresponds to the probability that, given the data and the prior, the true values of the parameters are within that region.

2.3 Nuisance parameters

The handling of nuisance parameters is also an intense topic of debate. In a Bayesian framework, the answer is marginalization. If the parameter vector contains two sets: interesting parameters $\vec{\theta}_i$ and nuisance parameters $\vec{\theta}_n$, then the posterior distribution of the interesting parameters is given by:

$$P(\vec{\theta}_i|D, \mathcal{M}) = \int_{\vec{\theta}_n} P(\vec{\theta}_i, \vec{\theta}_n|D, \mathcal{M}) d\vec{\theta}_n. \quad (2.4)$$

We partition the full posterior by integrating over the nuisance parameter space. This ensures that all of the probability mass contained in the nuisance parameters is accounted for. On the other hand, nuisance parameters are a major problem of non-Bayesian statistical theories [36]. There is no consensus frequentist way of addressing this problem (see [37] for a discussion of some of the methods).

One of the most common frequentist methods is the profile likelihood (e.g. [7, 17, 18]), which involves computing the likelihood of $\vec{\theta}_i$ on a grid, while conditioning on special values of the nuisance parameters, e.g. the best fit values $\vec{\theta}_{n,best}$. In practice, this is done by maximising the likelihood over all nuisance parameters at every point on the grid of interesting parameters:

$$P(D|\vec{\theta}_i, \mathcal{M}) \propto \max_{\vec{\theta}_n} P(D|\vec{\theta}_i, \vec{\theta}_n, \mathcal{M}). \quad (2.5)$$

For a Gaussian distribution this is equivalent to marginalization because low dimensional cuts from a high dimensional Gaussian are also Gaussian and their volume scales exactly the same as their peak. However, this method can go wrong when a lower dimensional cut changes shape depending on the position of the cut (we show an example of this in Section 2.4), because it effectively ignores any possibility that the nuisance parameters can have other values.

The profile likelihood is also useful for computing confidence levels when there are many interesting parameters. The scanning method presented in Section 2.2 becomes infeasible if we have too many parameters, and as such, the recursive application of the profile likelihood can be a useful approximation. We can scan a small subset of the parameters (usually one or two at a time) and treat the other parameters as nuisance by applying the profile likelihood at every point on the grid. If we apply this method recursively, we can compute an approximation of the full CIs. The profile likelihood is the method used so far in analyses of the Ly α forest correlation function by BOSS and eBOSS [10, 12–16, 38].

2.4 Toy Example

We illustrate the profile likelihood using a toy example in Figure 1. We chose a bivariate distribution $P(\alpha, \beta)$ used in [39], which, depending on the parameter of interest, shows where




Figure 1: Toy model of a bivariate distribution $P(\alpha, \beta)$ used to illustrate the profile likelihood, a common frequentist approximation for dealing with high-dimensional problems. A scan over α conditional on the best fit value of β at each point gives identical results to the marginal distribution (right panel). On the other hand, an equivalent scan over β gives the wrong result (top panel), because the conditional distribution $P(\alpha|\beta)$ changes shape depending on the value of β .

the profile likelihood works perfectly as well as where it fails. If the interesting parameter is $\theta_i = \alpha$, a scan over α conditional on the best fit value of β at each point gives a result identical to the marginal distribution (bottom right panel of Figure 1). On the other hand, if $\theta_i = \beta$, the equivalent procedure gives the wrong result because on the right hand side of the β grid, the conditional distribution $P(\alpha|\beta)$ is multimodal. This means the shape of $P(\alpha|\beta)$ changes depending on the value of β , resulting in a failure of the profile likelihood approximation.

This behaviour is not due to the difference in interpretation between a Bayesian and a frequentist. In fact, for this bivariate case, a simultaneous scan of both parameters would produce frequentist results identical to the Bayesian ones, but this is infeasible in high dimensional spaces. If such pathological cases are correctly identified, frequentists can rely on reparameterizations to Gaussianize the problem. However, such cases could appear over arbitrarily many dimensions and be easily missed because usually we can only investigate 1 and 2 dimensional projections.

3 Testing BAO measurements on mock data sets

We begin our investigation by applying the different methods introduced above to the problem of efficient and accurate extraction of BAO information from the Ly α forest correlation

function. To this end, we use a baseline model of the correlation function to create synthetic data sets. The baseline model is based on the Ly α correlation functions measured by the extended Baryon Oscillation Spectroscopic Survey (eBOSS) using SDSS DR14 data [15, 16], and the publicly available modelling package `picca`³. Their main analysis follows the frequentist methods introduced in Section 2. We perform a Bayesian analysis of the synthetic data sets and compare the different methodologies, with a focus on accurate BAO measurement.

3.1 Synthetic Correlation Functions

In order to compare different fitting methods, we produce 100 Monte Carlo simulations of the Ly α forest flux 3D correlation function. This allows us to investigate the differences over the entire population of possible correlation functions given a data set such as SDSS DR14. For simplicity, in this section we focus only on the Ly α ×Ly α auto-correlation function using Ly α absorbers only in the Ly α region, and leave the analysis using the full DR14 data for the next section. We use the measured Ly α correlation function from SDSS DR14 (shown in Figure 8 of [15]). We fit this using the full model from [15], including metal contaminations.

The mock data sets are drawn randomly from a multivariate normal distribution with mean $\xi(\vec{\theta}_{best})$ and covariance C , where $\xi(\vec{\theta})$ is the best fit model of the correlation function measured in the DR14 analysis [15] and C is the covariance matrix of ξ . A new simulated correlation function is then given by:

$$\tilde{\xi} = \xi(\vec{\theta}_{best}) + A\vec{y}, \quad (3.1)$$

where the matrix A comes from the Cholesky decomposition $C = AA^T$, and \vec{y} is a vector of N (the size of the C) independent standard normal variates.

The model used to fit the Ly α ×Ly α correlation function in the DR14 analyses has 11 free parameters, of which nine are considered nuisance parameters and two ($\alpha_{||}$, α_{\perp}) are the parameters of interest [15]. We follow this distinction and leave the analysis and description of the nuisance parameters for the next section. We fit our mocks using this model and the four methods introduced above. We use large flat priors for most parameters, but we follow [15] and set tight Gaussian priors for two of the nuisance parameters that are less constrained by the data (β_{HCD} and b_{CIV}). The choice of uninformative priors means that the shape of the posterior is given only by the likelihood. This means that any differences between Bayesian and frequentist results will be produced by the differences discussed in Section 2, and will not be influenced by prior choices.

3.2 Fitting Methods

Following the discussion in Section 2, we choose four fitting methods to compare over the population of synthetic data sets:

1. **Frequentist MLE:** The likelihood is maximised over all parameters and the uncertainties are given by the covariance around $\vec{\theta}_{best}$, using the Gaussian $\Delta\mathcal{L}_p$ values.
2. **Frequentist scan:** A scan over $\alpha_{||}$ and α_{\perp} using the profile likelihood, with the uncertainties given by confidence intervals which are set using MC simulations.
3. **Bayesian scan:** A scan over $\alpha_{||}$ and α_{\perp} using the profile likelihood, with the uncertainties given by credible regions, computed using Equation 2.3.

³Available at <https://github.com/igmhub/picca>

4. **Bayesian sampler:** The full posterior distribution is computed, and the uncertainties on $(\alpha_{\parallel}, \alpha_{\perp})$ are given by credible regions after marginalization over the nuisance parameters.

$\text{Ly}\alpha$ BAO analyses have so far been frequentist, and used the first two methods. The Bayesian interpretation of the scan was also used whenever scan results were combined with other cosmological probes (e.g. [40]) as part of popular packages such as `CosmoMC` or `MontePython` [41–43]. However, the two interpretations of the scan were never tested together. This method is also a good middle ground between the frequentist scan and the Bayesian sampler, because the contrast with the first is only in the interpretation of uncertainty, while a comparison with the second allows us to directly test the profile likelihood.

The scan requires maximising the likelihood over all nuisance parameters at each point on the grid. As this operation can be performed independently for each point, we implemented a parallel version of the scan code in `picca` to speed up our analysis. The frequentist analysis also requires a large number of MC simulations in order to compute the $\Delta\mathcal{L}_p^2$ as described in Section 2. We also implemented a parallel version of this step which is now available in `picca`.

In order to sample the full posterior distribution, we implemented an interface to the popular Nested Sampler `PolyChord`⁴ [44, 45] in `picca`. Nested samplers were designed to efficiently compute the Bayesian evidence [35], but they also provide accurate sampling of the posterior. In particular, nested samplers are very good at dealing with multimodal and highly degenerate posteriors.

3.3 Results

Noise on the correlation function can conspire to improve or smear the BAO peak. This means that a population of simulated correlation functions produced from the same covariance gives rise to a range of possible BAO constraints, from very tight, closely Gaussian fits, to very degenerate, multimodal ones where the peak is barely detected. In Figure 2, we showcase two examples of results from this population. The one on the left is one of the best BAO constraints in the population, and is very close to Gaussian. On the other hand, the one on the right gives one of the worst constraints and is very non-Gaussian. We note that in the case of DR14-like errors on the correlation function, roughly 70% of constraints (computed using visual inspection of the mock fits) are close to Gaussian, similar to the one on the left. On the other hand, 4% are catastrophic failures, where the constraints are multimodal and strongly non-Gaussian in both parameters (BAO not detected).

Results of the four methods described above are presented in Figure 2. The last three methods agree very well with each other even in the non-Gaussian case. This agreement is remarkable, because as discussed in Section 2, they answer fundamentally different questions. The frequentist scan uses confidence intervals to quantify uncertainty, while the two Bayesian methods use credible regions. The agreement between the frequentist scan and the Bayesian scan shows us that BAO results are robust to different interpretations of uncertainty. On the other hand, the agreement between the Bayesian scan and the Bayesian sampler shows us that the profile likelihood is a good approximation for fitting the BAO peak from the $\text{Ly}\alpha$ forest correlation function.

The Frequentist MLE, which assumes Gaussianity, does not fare as well. It produces results that are close to the ones obtained using the other three methods in the Gaussian mock,

⁴Available at <https://github.com/PolyChord/PolyChordLite>

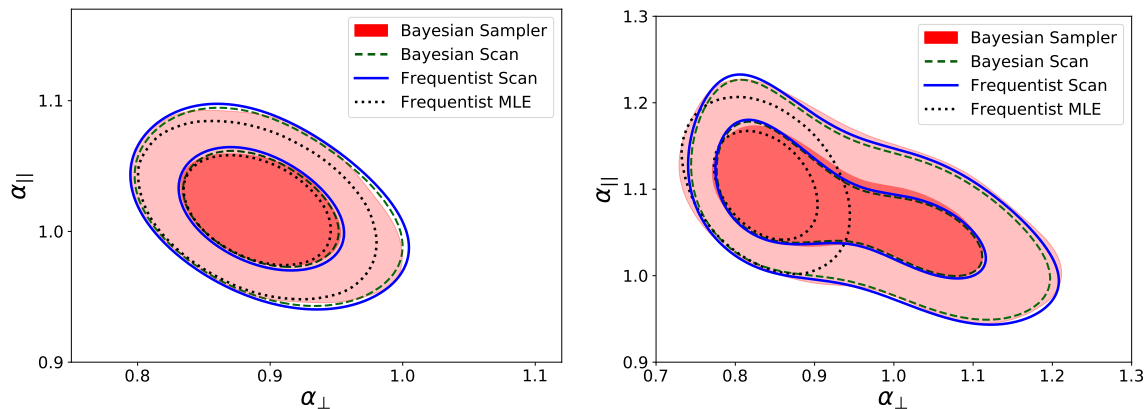


Figure 2: Comparison of BAO parameter constraints on mock correlation functions using the frequentist MLE and frequentist scan confidence intervals, and the Bayesian scan and Bayesian sampler credible regions. We showcase two mocks, one where the constraints are very close to Gaussian (left), and one where the constraints are strongly non-Gaussian (right). Note the different scaling. The last three methods produce remarkably similar results considering they use different methods and very different quantification of uncertainties. On the other hand, the frequentist MLE can fail to properly capture the uncertainty in the results, especially in non-Gaussian cases.

although slightly smaller. On the other hand, it completely fails to capture the uncertainty in the results for the non-Gaussian mock. As these non-Gaussian posteriors make up roughly 30% of the population of possible results, this approximation is inadequate for dealing with current data sets.

The comparison of specific posterior results such as in Figure 2 is Bayesian in nature. However, we can also compare them using a frequentist approach by computing the fraction of the interval population that contains the truth $\vec{\theta}_{true}$ within the 68% and 95% regions. We find that the frequentist scan, the Bayesian scan and the Bayesian sampler are again in remarkable agreement, with roughly 55 – 60 mocks containing the truth within their 68% regions, and roughly 92 of them containing the truth within their 95% regions. On the other hand, the MLE fails this test as well, with only 44 of the 68% intervals, and 74 of the 95% intervals containing the truth. These numbers are affected by sample noise because we only have 100 simulations, and as such are very rough. In particular, if the profile likelihood works perfectly, by construction the fractions for the frequentist scan will tend to 68% and 95% for a large number of samples.

These results do, of course, rely on the large uniform priors assumed for most parameters. Non-uniform or overly restrictive priors would have an impact on our conclusions as some of the differences could come from the choice of priors. In a study similar to our work, [7] arrive at a different conclusion when tight priors are used on the BAO parameters. We discuss the similarities and differences between our study and [7] in Appendix B.

Another relevant question that distinguishes these methods is that of the computational cost. MLE is by far the fastest method, with computation times of order $10^{-1} - 10^0$ CPU hours. However, as we just showed, MLE alone is inadequate for current Ly α forest BAO analyses. For the scan we find computation times of order $10^2 - 10^4$ CPU hours using typical grid sizes of 30×30 up to 50×50 . Furthermore, the frequentist interpretation of the scan

requires a large set of MC simulations to compute the $\Delta\mathcal{L}_p$ necessary for setting the right CIs. In the best case these require an extra 10^3 CPU hours for 10000 MC mocks. `PolyChord` with a typical setup ($n\text{live} = 25 \times$ number of parameters, $\text{numrepeats} = 3 \times$ number of parameters) performs over a timescale similar to the scan ($10^2 - 10^4$ CPU hours), and not only computes the full posterior distribution, but also its integral (the Bayesian evidence).

4 The Full Posterior of eBOSS DR14

We now turn our attention to the full posterior distribution. As we have shown above, when the only interest is measuring the BAO peak position, both the scan and the sampler produce very similar results. However, the sampler also computes accurate distributions for all the other parameters. This wealth of information is generally ignored. Some of these are astrophysical parameters, such as the bias of high column density (HCD) absorbers (b_{HCD}). These parameters tend to be very correlated with each other and have proved to be very sensitive to modelling choices, which makes their measurement less robust compared to that of the BAO peak.

The computation of the full posterior distribution allows us to access this previously ignored information. It allows us to study the complex high-dimensional distribution of these parameters, and to better understand how our modelling choices affect their measurement. Furthermore, our use of the Nested Sampler `PolyChord` means that we can deal with very degenerate and even multimodal distributions. This sampler also computes the Bayesian evidence, thus providing an accurate tool for model comparison even for strongly non-Gaussian posteriors. As such, it allows us to test possible extensions to the baseline model that may have previously appeared daunting due to their complex interaction with other parameters.

4.1 The Baseline Analysis

We start by using `PolyChord` to analyse all Ly α forest eBOSS DR14 correlation functions. In contrast to the last section, we include the correlation of Ly α absorbers in the Ly α region with Ly α absorbers in the Ly β region, $Ly\alpha(Ly\alpha) \times Ly\alpha(Ly\beta)$, and the cross-correlation with quasars Ly $\alpha \times$ QSO. We use the models presented in [15] and [16] to model the auto-correlation and the cross-correlation with quasars respectively. Our analysis follows the same steps, but we do not include relativistic effects, and we use only one parameter to model the bias of HCDs (instead of three). As in the previous section, we use broad flat priors on all parameters (except β_{HCD} and the bias of foreground CIV absorption (b_{CIV}) which have Gaussian priors). The model parameters used and their priors are presented in Table 1.

In Figure 3 we show the projected posterior distributions of α_{\parallel} and α_{\perp} versus the other parameters. This shows the robustness of BAO measurements as none of the parameters are correlated with $(\alpha_{\parallel}, \alpha_{\perp})$. The only exception is the bias of the *SiII*(1260) absorbers ($b_{SiII(1260)}$) which has a small correlation with α_{\parallel} . This is due to the line causing an increased correlation along the line of sight at a separation of $r_{\parallel} \approx 105h^{-1}\text{Mpc}$, which is very close to the BAO peak. However, this metal contamination is barely detected (at $\sim 2\sigma$).

4.2 A simple extension

We now use `PolyChord` to illustrate the advantages of sampling the full posterior. We do this by sampling L_{HCD} as an extension to the baseline analysis. This is a parameter used to

Parameter	Description	Prior
$\alpha_{\parallel}, \alpha_{\perp}$	BAO peak position	$\Pi[0.1, 2]$
$b_{\eta Ly\alpha}$	Ly α velocity bias	$\Pi[-0.5, 0]$
$\beta_{Ly\alpha}$	Ly α RSD parameter	$\Pi[0.1, 5.0]$
β_{QSO}	Bias of HCDs	$\Pi[-0.2, 0]$
$\Delta r_{\parallel}[h^{-1}Mpc]$	Shift due to QSO redshift errors	$\Pi[-10, 10]$
$\sigma_v[h^{-1}Mpc]$	Smoothing parameter for QSO non-linear velocities and redshift precision	$\Pi[2, 15]$
ξ_0^{TP}	Amplitude parameter of quasar radiation	$\Pi[0, 2]$
b_{HCD}	Bias of HCDs	$\Pi[-0.2, 0]$
β_{HCD}	RSD parameter of HCDs	$\mathcal{N}(0.5, 0.2^2)$
$b_{\eta CIV(eff)}$	Velocity bias of metal absorber	$\mathcal{N}(-0.005, 0.0026^2)$
$b_{\eta SiII(1190)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$
$b_{\eta SiII(1193)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$
$b_{\eta SiIII(1207)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$
$b_{\eta SiII(1260)}$	Velocity bias of metal absorber	$\Pi[-0.2, 0]$

Table 1: Sampled parameters and their prior limits. We use flat priors $\Pi[a, b]$ for most parameters, with limits a and b chosen such that the prior is uninformative. Following [15, 16], we use Gaussian priors $\mathcal{N}(\mu, \sigma^2)$ with mean μ and standard deviation σ for β_{HCD} and $b_{\eta CIV(eff)}$.

describe the contamination by HCDs and it corresponds to a typical length scale. HCDs with a length scale above $\sim 14h^{-1}Mpc$ are efficiently identified and masked before computing the correlation function. This means unidentified systems are expected to have a typical length scale below this value, but above the bin width of the correlation function ($4h^{-1}Mpc$). This parameter has so far been fixed to a value of $L_{HCD} = 10h^{-1}Mpc$ following the study by [46]. Different values for this parameter were tested by [15, 16] to confirm there are no biases when measuring the BAO position, however, completely freeing this parameter proved challenging for the minimiser to deal with. We set a flat prior given by: $\Pi[2, 30]h^{-1}Mpc$.

As the full projected posterior is too large to show, we chose a subset of parameters and plot their posteriors in Figure 4. We show α_{\parallel} and α_{\perp} to study the impact of L_{HCD} on measurements of the BAO peak position. From the other parameters we chose those whose posterior is affected by L_{HCD} . We also plot the posterior using the basic model where L_{HCD} is fixed for comparison. Plots of the full posterior are available as part of a Jupyter Notebook at <https://github.com/andreiceuceu/eBOSS-Lya-Posteriors>.

The first two columns in Figure 4 show that even though the posterior is significantly affected by the choice of L_{HCD} , the BAO peak position is very robust to it. In particular, the $\alpha_{\parallel} - \alpha_{\perp}$ posterior of the extended model is in excellent agreement with the results from the baseline analysis. L_{HCD} is constrained within the expected region: $L_{HCD} = 6.6 \pm 1.2, {}^{+2.5}_{-2.3} h^{-1}Mpc$ (68%, 95%). As can be seen in the bottom row, the other parameters shown in Figure 4 are correlated with L_{HCD} , and as such their posterior is significantly affected. The constraints on the biases and RSD parameters of the Ly α forest and HCDs are all wider




Figure 3: Projected posterior distributions of the BAO parameters (α_{\parallel} , α_{\perp}), versus the other model parameters, using the full eBOSS DR14 Ly α forest data. There are no major correlations between $\alpha_{\parallel}/\alpha_{\perp}$ and any of the other parameters, which shows the robustness of BAO parameters to different modelling choices.

than in the baseline analysis, showing the fragility of these parameters to modelling choices. Interestingly, the bias of the SiIII(1207) absorption is also very correlated with L_{HCD} .

Comparisons such as the one presented in Figure 4 are beneficial because they visually show us how model parameters behave, and how our modelling choices influence our results.

5 Conclusions

Analyses of the Ly α forest 3D correlation function have so far focused on measuring the position of the BAO peak at the expense of analyses of other parameters and their interaction. A frequentist methodology has been used so far, despite the BAO results being subsequently used as part of Bayesian analyses when combined with other probes. In this work, we performed for the first time a Bayesian analysis of the Ly α forest correlation function, and we computed the full posterior using the Nested Sampler `PolyChord`.

We started by discussing the different approaches to fitting the correlation function in Section 2, and in particular we focused on the methods used and the difference in the quantification of uncertainty. Frequentists use Confidence Intervals computed from the population of possible data sets, while Bayesians use credible regions computed from the posterior distribution of the parameters given the data and the prior. Furthermore, when dealing with high-dimensional model spaces, Bayesians have access to tools such as MCMC. On the frequentist side, approximations such as the profile likelihood are used. A scan is performed over a few parameters (usually anything above 2 or 3 is computationally infeasible), and the

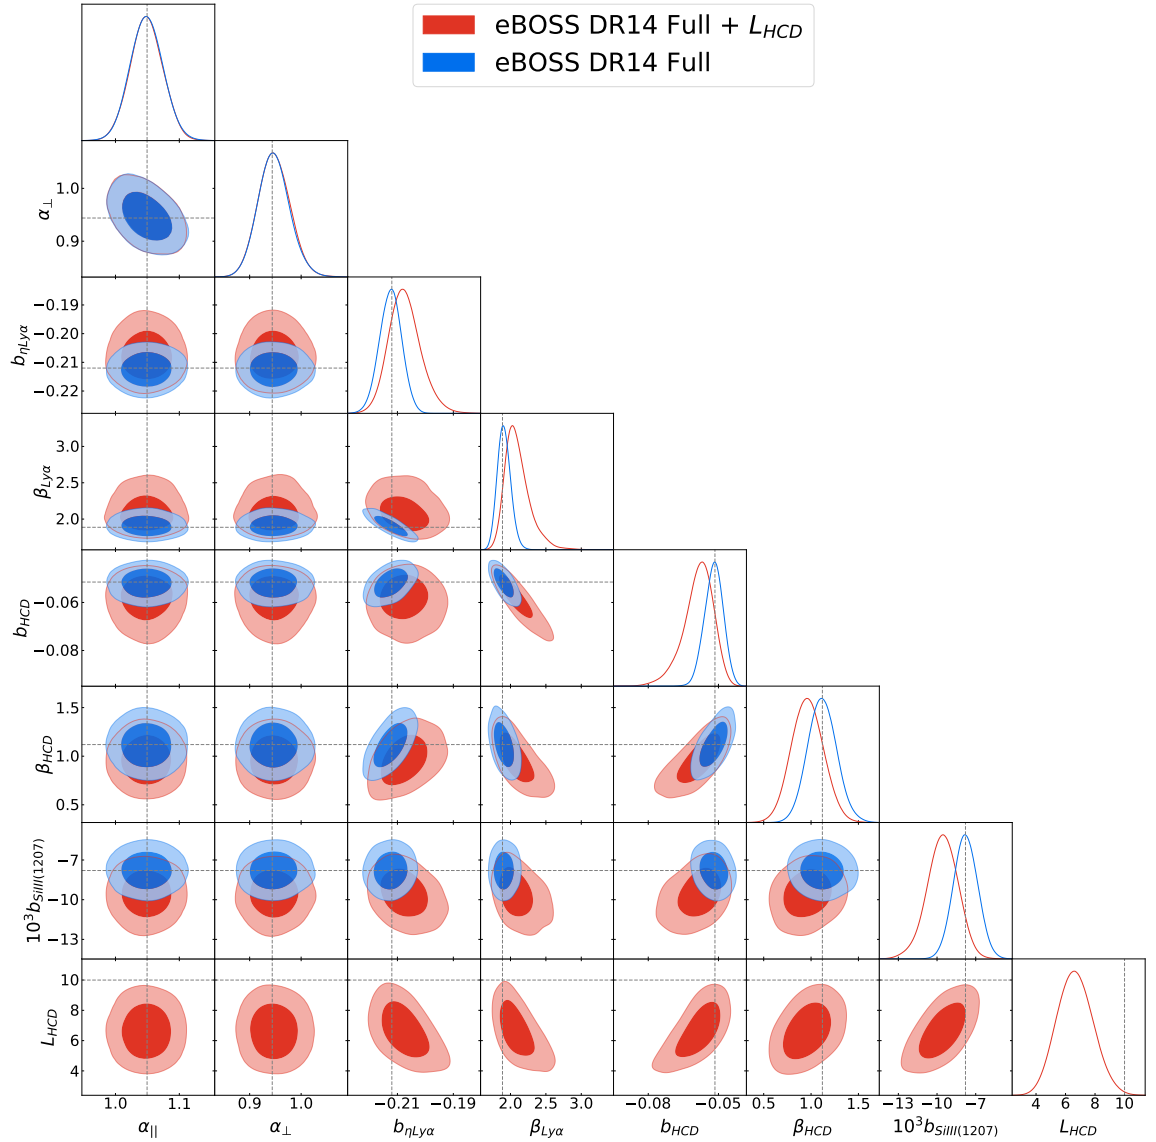


Figure 4: Triangle plot showing projected posteriors of DR14 Ly α forest results using the baseline analysis versus an extension where the typical scale of high column density absorbers, L_{HCD} , is sampled. The best fit results for the baseline analysis are given by the dashed lines. The first two columns show that the measurement of the BAO scale is very robust to this change. However, the other parameters plotted are correlated with L_{HCD} , and as such, their posteriors are significantly affected by this parameter.

likelihood is maximised over other parameters at each point. This is generally a good approximation, however, it can fail in some cases, such as in the toy example presented in Section 2.

We compared the different methodologies on synthetic correlation functions of eBOSS DR14 in Section 3. We used both a Bayesian and a frequentist interpretation of scan results. This allows us to test potential differences from the quantification of uncertainty by comparing the two. We can also test the profile likelihood approximation by comparing the Bayesian scan with Bayesian sampler results. We showed that the three methods agree remarkably well on both a mock with a tight BAO constraint, and one with a large non-Gaussian one. We also plotted results using the frequentist Maximum Likelihood Estimation, where the uncertainties are set using the second derivative around the peak likelihood. This method fails to capture the uncertainty in non-Gaussian cases. Considering these make up roughly $\sim 30\%$ of the population, the MLE is inadequate for current data sets.

By computing the Bayesian interpretation of the scan and showing it agrees very well both with the frequentist scan and the Bayesian sampler, we have confirmed that scan results can safely be combined with other probes as part of Bayesian packages commonly used in cosmology.

In Section 4 we turned our attention to the full posterior distribution of eBOSS DR14 Ly α forest correlation functions. We showed that the BAO peak position parameters, α_{\parallel} and α_{\perp} , do not have any strong correlation with any of the other parameters, and as such, they are very robust to modelling choices. We extended the baseline analysis by sampling the typical length scale of HCD systems, which was previously fixed to $L_{HCD} = 10h^{-1}\text{Mpc}$. We compared the projected posteriors of the two models, and showed that it has no impact on the measured BAO position. We measured $L_{HCD} = 6.6 \pm 1.2 h^{-1}\text{Mpc}$ (68%), and find that this parameter is correlated with many other astrophysical parameters and as such has a significant impact on their posteriors. Plots of the full posterior are available at <https://github.com/andreicuceu/eBOSS-Lya-Posteriors>.

Comparisons such as the one presented in Section 4.2 are now easy and fast to perform. Furthermore, this benefit does not come at a major computational cost. As discussed in Section 3, PolyChord requires similar computational time compared to a 2D scan using the profile likelihood. Beyond constraining the BAO peak parameters, PolyChord also computes the full posterior distribution of all parameters and the Bayesian evidence.

Our improvements are freely available as part of the community code `picca`. We hope that these tools will be used to improve future Ly α forest analyses by using the full posterior distribution to study complex parameter distributions and inform modelling choices.

Acknowledgments

We thank Lorne Whiteway, Pablo Lemos, Anže Slosar, Ashley Ross and Kwan Chuen Chan for their comments and advice on the draft. This work was partially enabled by funding from the UCL Cosmoparticle Initiative. AC was supported by a Science and Technology Facilities Council (STFC) studentship. AFR acknowledges support by an STFC Ernest Rutherford Fellowship, grant reference ST/N003853/1 and further support by STFC Consolidated Grant number ST/R000476/1.

References

- [1] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi et al., *Planck 2018 results. VI. Cosmological parameters, arXiv e-prints* (2018) arXiv:1807.06209 [[1807.06209](#)].
- [2] T. M. C. Abbott, F. B. Abdalla, A. Alarcon, J. Aleksić, S. Allam, S. Allen et al., *Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing, Phys. Rev. D* **98** (2018) 043526 [[1708.01530](#)].
- [3] D. J. Eisenstein, I. Zehavi, D. W. Hogg, R. Scoccimarro, M. R. Blanton, R. C. Nichol et al., *Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies, ApJ* **633** (2005) 560 [[astro-ph/0501171](#)].
- [4] S. Cole, W. J. Percival, J. A. Peacock, P. Norberg, C. M. Baugh, C. S. Frenk et al., *The 2dF Galaxy Redshift Survey: power-spectrum analysis of the final data set and cosmological implications, Mon. Not. Roy. Astron. Soc.* **362** (2005) 505 [[astro-ph/0501174](#)].
- [5] M. Vargas Magaña, S. Ho, X. Xu, A. G. Sánchez, R. O’Connell, D. J. Eisenstein et al., *SDSS-III Baryon Oscillation Spectroscopic Survey: Analysis of Potential Systematics in Fitting of Baryon Acoustic Feature, arXiv e-prints* (2013) arXiv:1312.4996 [[1312.4996](#)].
- [6] L. Anderson, É. Aubourg, S. Bailey, F. Beutler, V. Bhardwaj, M. Blanton et al., *The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples, Mon. Not. Roy. Astron. Soc.* **441** (2014) 24 [[1312.4877](#)].
- [7] K. C. Chan, M. Crocce, A. J. Ross, S. Avila, J. Elvin-Poole, M. Manera et al., *BAO from angular clustering: optimization and mitigation of theoretical systematics, Mon. Not. Roy. Astron. Soc.* **480** (2018) 3031 [[1801.04390](#)].
- [8] S. R. Hinton, C. Howlett and T. M. Davis, *Barry and the BAO Model Comparison, Mon. Not. Roy. Astron. Soc.* (2020) [[1912.01175](#)].
- [9] N. G. Busca, T. Delubac, J. Rich, S. Bailey, A. Font-Ribera, D. Kirkby et al., *Baryon acoustic oscillations in the Ly α forest of BOSS quasars, Astron. Astrophys.* **552** (2013) A96 [[1211.2616](#)].
- [10] A. Slosar, V. Iršič, D. Kirkby, S. Bailey, N. G. Busca, T. Delubac et al., *Measurement of baryon acoustic oscillations in the Lyman- α forest fluctuations in BOSS data release 9, JCAP* **2013** (2013) 026 [[1301.3459](#)].
- [11] D. Kirkby, D. Margala, A. Slosar, S. Bailey, N. G. Busca, T. Delubac et al., *Fitting methods for baryon acoustic oscillations in the Lyman- α forest fluctuations in BOSS data release 9, JCAP* **2013** (2013) 024 [[1301.3456](#)].
- [12] A. Font-Ribera, D. Kirkby, N. Busca, J. Miralda-Escudé, N. P. Ross, A. Slosar et al., *Quasar-Lyman α forest cross-correlation from BOSS DR11: Baryon Acoustic Oscillations, Journal of Cosmology and Astro-Particle Physics* **2014** (2014) 027 [[1311.1767](#)].
- [13] J. E. Bautista, N. G. Busca, J. Guy, J. Rich, M. Blomqvist, H. du Mas des Bourboux et al., *Measurement of baryon acoustic oscillation correlations at $z = 2.3$ with SDSS DR12 Ly α -Forests, Astron. Astrophys.* **603** (2017) A12 [[1702.00176](#)].
- [14] H. du Mas des Bourboux, J.-M. Le Goff, M. Blomqvist, N. G. Busca, J. Guy, J. Rich et al., *Baryon acoustic oscillations from the complete SDSS-III Ly α -quasar cross-correlation function at $z = 2.4$, Astron. Astrophys.* **608** (2017) A130 [[1708.02225](#)].
- [15] V. de Sainte Agathe, C. Balland, H. du Mas des Bourboux, N. G. Busca, M. Blomqvist, J. Guy et al., *Baryon acoustic oscillations at $z = 2.34$ from the correlations of Ly α absorption in eBOSS DR14, arXiv e-prints* (2019) arXiv:1904.03400 [[1904.03400](#)].

- [16] M. Blomqvist, H. du Mas des Bourboux, N. G. Busca, V. de Sainte Agathe, J. Rich, C. Balland et al., *Baryon acoustic oscillations from the cross-correlation of Ly α absorption and quasars in eBOSS DR14*, *arXiv e-prints* (2019) arXiv:1904.03430 [[1904.03430](#)].
- [17] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont et al., *Planck intermediate results. XVI. Profile likelihoods for cosmological parameters*, *Astron. Astrophys.* **566** (2014) A54 [[1311.1657](#)].
- [18] S. Algeri, J. Aalbers, K. Dundas Morà and J. Conrad, *Searching for new physics with profile likelihoods: Wilks and beyond*, *arXiv e-prints* (2019) arXiv:1911.10237 [[1911.10237](#)].
- [19] C. Pichon, J. L. Vergely, E. Rollinde, S. Colombi and P. Petitjean, *Inversion of the Lyman $\hat{\text{I}}\hat{\text{s}}$ forest: three-dimensional investigation of the intergalactic medium*, *Monthly Notices of the Royal Astronomical Society* **326** (2001) 597 [<https://academic.oup.com/mnras/article-pdf/326/2/597/3249991/326-2-597.pdf>].
- [20] F.-S. Kitaura, S. Gallerani and A. Ferrara, *Multiscale inference of matter fields and baryon acoustic oscillations from the Ly α forest*, *Mon. Not. Roy. Astron. Soc.* **420** (2012) 61 [[1011.6233](#)].
- [21] B. Horowitz, K.-G. Lee, M. White, A. Krolewski and M. Ata, *TARDIS. I. A Constrained Reconstruction Approach to Modeling the $z \sim 2.5$ Cosmic Web Probed by Ly α Forest Tomography*, *ApJ* **887** (2019) 61 [[1903.09049](#)].
- [22] N. Porqueres, J. Jasche, G. Lavaux and T. Enßlin, *Inferring high-redshift large-scale structure dynamics from the Lyman- α forest*, *Astron. Astrophys.* **630** (2019) A151 [[1907.02973](#)].
- [23] T. Bayes, *Lii. an essay towards solving a problem in the doctrine of chances. by the late rev. mr. bayes, frs communicated by mr. price, in a letter to john canton, amfr s*, *Philosophical transactions of the Royal Society of London* (1763) 370.
- [24] P. S. Laplace, *Théorie analytique des probabilités*. Courcier, 1820.
- [25] J. Neyman, *Outline of a theory of statistical estimation based on the classical theory of probability*, *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences* **236** (1937) 333.
- [26] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller, *Equation of State Calculations by Fast Computing Machines*, *The Journal of Chemical Physics* **21** (1953) 1087.
- [27] W. K. Hastings, *Monte Carlo sampling methods using Markov chains and their applications*, *Biometrika* **57** (1970) 97 [<https://academic.oup.com/biomet/article-pdf/57/1/97/23940249/57-1-97.pdf>].
- [28] A. R. Liddle, P. Mukherjee and D. Parkinson, *Cosmological model selection*, *arXiv e-prints* (2006) astro [[astro-ph/0608184](#)].
- [29] R. Trotta, *Applications of Bayesian model selection to cosmological parameters*, *Mon. Not. Roy. Astron. Soc.* **378** (2007) 72 [[astro-ph/0504022](#)].
- [30] R. Trotta, *Bayes in the sky: Bayesian inference and model selection in cosmology*, *Contemporary Physics* **49** (2008) 71 [<https://doi.org/10.1080/00107510802066753>].
- [31] E. T. Jaynes and O. Kempthorne, *Confidence intervals vs Bayesian intervals*. Springer, 1976.
- [32] B. L. Welch, *Note on discriminant functions*, *Biometrika* **31** (1939) 218.
- [33] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007.
- [34] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari and D. B. Rubin, *Bayesian data analysis*. Chapman and Hall/CRC, 2013.

- [35] J. Skilling, *Nested Sampling*, vol. 735 of *American Institute of Physics Conference Series*. Nov, 2004, [10.1063/1.1835238](#).
- [36] J. Ghosh, *Statistical information and likelihood: a collection of critical essays by Dr. D. Basu*. Springer, 1988.
- [37] D. Basu, *On the elimination of nuisance parameters*, *Journal of the American Statistical Association* **72** (1977) 355.
- [38] T. Delubac, J. E. Bautista, N. G. Busca, J. Rich, D. Kirkby, S. Bailey et al., *Baryon acoustic oscillations in the Ly α forest of BOSS DR11 quasars*, *Astron. Astrophys.* **574** (2015) A59 [[1404.1801](#)].
- [39] Z. Wang, M. Broccardo and J. Song, *Hamiltonian Monte Carlo Methods for Subset Simulation in Reliability Analysis*, *arXiv e-prints* (2017) arXiv:1706.01435 [[1706.01435](#)].
- [40] A. Cuceu, J. Farr, P. Lemos and A. Font-Ribera, *Baryon Acoustic Oscillations and the Hubble constant: past, present and future*, *JCAP* **2019** (2019) 044 [[1906.11628](#)].
- [41] A. Lewis and S. Bridle, *Cosmological parameters from CMB and other data: A Monte Carlo approach*, *Phys. Rev. D* **66** (2002) 103511 [[astro-ph/0205436](#)].
- [42] B. Audren, J. Lesgourgues, K. Benabed and S. Prunet, *Conservative Constraints on Early Cosmology: an illustration of the Monte Python cosmological parameter inference code*, *JCAP* **1302** (2013) 001 [[1210.7183](#)].
- [43] T. Brinckmann and J. Lesgourgues, *MontePython 3: boosted MCMC sampler and other features*, *arXiv e-prints* (2018) arXiv:1804.07261 [[1804.07261](#)].
- [44] W. J. Handley, M. P. Hobson and A. N. Lasenby, *polychord: nested sampling for cosmology.*, *Mon. Not. Roy. Astron. Soc.* **450** (2015) L61 [[1502.01856](#)].
- [45] W. J. Handley, M. P. Hobson and A. N. Lasenby, *POLYCHORD: next-generation nested sampling*, *Mon. Not. Roy. Astron. Soc.* **453** (2015) 4384 [[1506.00171](#)].
- [46] K. K. Rogers, S. Bird, H. V. Peiris, A. Pontzen, A. Font-Ribera and B. Leistedt, *Correlations in the three-dimensional Lyman-alpha forest contaminated by high column density absorbers*, *Mon. Not. Roy. Astron. Soc.* **476** (2018) 3716 [[1711.06275](#)].

A Computation of CIs

If the model is not a linear function of the parameters, then the $\Delta\mathcal{L}_p$ values, necessary for setting confidence intervals from grid results, usually cannot be computed analytically. In general these values must be computed from the population of the possible realisations of the data. However, this population is inaccessible if we can only perform the experiment once, and we must rely on simulated populations of data sets instead of the real one [33]. In cosmology this could be achieved by running the full analysis many times on simulations, however as these are extremely expensive computationally, we usually rely on Monte Carlo (MC) simulations instead. The idea behind MC is to use the measured correlation functions and their covariance matrices together with a model of the data to produce random realisations of our data. See Chapter 15.6 of [33] for a detailed discussion of this algorithm. Note that the χ^2 is generally used instead of the log-likelihood, and $\Delta\chi^2$ values are computed, however, we use the log-likelihood to maintain consistency and because the two are proportional if the likelihood is Gaussian: $\chi^2 \propto -2 \log \mathcal{L}$.

We want to produce samples from the distribution of possible data sets given the true model and the true parameter values: $P(D|\vec{\theta}_{true}, \mathcal{M}_{true})$. However, finding this distribution is

the goal of our experiment, and as such it is inaccessible at this point. Therefore, we start by assuming that the shape of $P(D|\vec{\theta}_{true}, \mathcal{M}_{true})$ is similar to the shape of $P(D|\vec{\theta}_0, \mathcal{M}_0)$, where we are using our particular model \mathcal{M}_0 , and the best fit parameters given that model and our data: $\vec{\theta}_0$. This means that we can use $\vec{\theta}_0$ as a surrogate for $\vec{\theta}_{true}$ and create a population of synthetic data sets from a multivariate normal with mean $\mathcal{M}_0(\vec{\theta}_0)$ and covariance matrix given by the measurement errors of the data (again assuming our data is normally distributed).

For each mock data set created using this method we find the best fit parameter values and call them $\vec{\theta}_\alpha$. If we compute the quantity $\Delta\mathcal{L} \equiv \mathcal{L}(\vec{\theta}_0) - \mathcal{L}(\vec{\theta}_\alpha)$ for each mock data set, we obtain a chi-square distribution with n degrees of freedom, where n is the number of parameters. By analysing a large population of mocks and computing this quantity, the $\Delta\mathcal{L}_p$ is easily obtained by looking at the $\Delta\mathcal{L}$ value that contains the fraction of mocks corresponding to p for that chi-square distribution.

If we are only interested in a subset of ν parameters, then the quantity of interest is $\Delta\mathcal{L}_\nu \equiv \mathcal{L}_\nu - \mathcal{L}(\vec{\theta}_\alpha)$, where \mathcal{L}_ν is found by fixing the parameters of interest to their best fit values from $\vec{\theta}_0$ and maximising over the other parameters for each mock. $\Delta\mathcal{L}_\nu$ is a chi-square distribution with ν degrees of freedom, and the same procedure as above is used to find $\Delta\mathcal{L}_p$.

B Comparison with Chan et al. 2018

The Dark Energy Survey (DES) collaboration performed a study similar to the one presented in this work, testing different fitting methods with the aim of measuring BAO using angular correlation functions in tomographic bins [7]. They also compared the performance using an MLE approach, the profile likelihood and an MCMC. When they computed the mean and uncertainty from the profile likelihood they used the mean and variance across the scan weighted by the likelihood value at each point. This is in contrast with our method of finding constant χ^2 surfaces, and as such it may lead to different results.

The main difference is that they worked with isotropic BAO which means they only need to measure one parameter, α . Furthermore, they also pruned their population of synthetic data sets to those where the 68% α constraints are entirely contained within the prior range [0.8, 1.2] (where $\alpha = 1$ is the true value). The result is that the remaining mocks are the ones where the BAO peak has a $\Delta\chi^2 = 1$ region within their prior range on α , which ensures the posterior distribution of α can be approximated as Gaussian. They found that all three methods perform very well, and after investigating the population statistics they conclude that the MLE is the best tool in these conditions.

In the particular case of isotropic BAO we find similar results. As the model has less freedom, most constraints are Gaussian and all three methods work very well. However, they found that the MCMC has a larger bias in estimating α compared to the MLE (although both are very small). We believe that this is because they used the median to quantify the MCMC result, while we are using the maximum posterior point. The best fit point given by an MCMC (assuming flat priors) and MLE should be the same if both are run appropriately, and as such there should not be any difference in the bias of the peak for a unimodal posterior.

In Appendix A, [7] worked with a larger range ([0.6, 1.4]) and showed that in this case the MCMC is the best approach, with both the MLE and profile likelihood showing small biases. This is in line with our results, considering that in this case there are some non-Gaussian results that are better fit using an MCMC approach. However, we also found the profile likelihood still works very well in these cases. This difference in results could be either

caused by the difference in quantifying uncertainty as discussed above, or by a failure of the profile likelihood approximation when dealing with the nuisance parameters.