

Proceeding Paper

Marginal Bayesian Statistics Using Masked Autoregressive Flows and Kernel Density Estimators with Examples in Cosmology [†]

Harry Bevins ^{1,2,*} , Will Handley ^{1,2} , Pablo Lemos ^{3,4}, Peter Sims ⁵, Eloy de Lera Acedo ^{1,2} and Anastasia Fialkov ^{2,6}

¹ Astrophysics Group, Cavendish Laboratory, J. J. Thomson Avenue, Cambridge CB3 0HE, UK

² Kavli Institute for Cosmology, Madingley Road, Cambridge CB3 0HA, UK

³ Department of Physics & Astronomy, University College London, Gower Street, London WC1E 6BT, UK

⁴ Department of Physics and Astronomy, University of Sussex, Pevensey Building, Brighton BN1 9QH, UK

⁵ Department of Physics and McGill Space Institute, McGill University, 3600 University Street, Montreal, QC H3A 2T8, Canada

⁶ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

* Correspondence: htjb2@cam.ac.uk

[†] Presented at the 41st International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Paris, France, 18–22 July 2022.



Citation: Bevins, H.; Handley, W.; Lemos, P.; Sims, P.; de Lera Acedo, E.; Fialkov, A. Marginal Bayesian Statistics Using Masked Autoregressive Flows and Kernel Density Estimators with Examples in Cosmology. *Phys. Sci. Forum* **2022**, *5*, 1. <https://doi.org/10.3390/psf2022005001>

Academic Editors: Frédéric Barbaresco, Ali Mohammad-Djafari, Frank Nielsen and Martino Trassinelli

Published: 27 October 2022

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: Cosmological experiments often employ Bayesian workflows to derive constraints on cosmological and astrophysical parameters from their data. It has been shown that these constraints can be combined across different probes, such as Planck and the Dark Energy Survey, and that this can be a valuable exercise to improve our understanding of the universe and quantify tension between multiple experiments. However, these experiments are typically plagued by differing systematics, instrumental effects, and contaminating signals, which we collectively refer to as ‘nuisance’ components, which have to be modelled alongside target signals of interest. This leads to high dimensional parameter spaces, especially when combining data sets, with $\gtrsim 20$ dimensions of which only ~ 5 correspond to key physical quantities. We present a means by which to combine constraints from different data sets in a computationally efficient manner by generating rapid, reusable, and reliable marginal probability density estimators, giving us access to nuisance-free likelihoods. This is possible through the unique combination of nested sampling, which gives us access to Bayesian evidence, and the marginal Bayesian statistics code MARGARINE. Our method is lossless in the signal parameters, resulting in the same posterior distributions as would be found from a full nested sampling run over all nuisance parameters, and typically quicker than evaluating full likelihoods. We demonstrate our approach by applying it to the combination of posteriors from the Dark Energy Survey and Planck.

Keywords: Bayesian analysis; normalizing flows; Kullback–Leibler; cosmology

1. Introduction

Bayesian inference is a cornerstone of modern cosmology and astrophysics. It is frequently employed to derive parameter constraints on key signal parameters from data sets such as the Dark Energy Survey (DES, [1,2]), Planck [3], REACH [4], and SARAS2 [5], among others. Often, experiments are sensitive to different aspects of the same physics, and by combining constraints across probes we can improve our understanding of the Universe or reveal tensions between different experiments.

However, this can become a computationally expensive task as many experiments feature systematic, instrumental effects, as well as contamination from other physical signals that need to be modelled alongside the signal or parameters of interest. For individual

experiments, this can lead to high dimensional problems with $\gtrsim 20$ parameters of which the majority can be considered ‘nuisance’ parameters. The problem is compounded when combining different data sets with different models for common nuisance components and different systematics or instrumental effects that have to be modelled.

In this work, we demonstrate that we can use density estimators, such as kernel density estimators [6,7] and masked autoregressive flows [8], to rapidly calculate reliable and reusable representations of marginal probability densities and marginal Bayesian summary statistics for key signal or cosmological parameters. This gives us access to the nuisance-free likelihood functions and allows us to combine parameter constraints from different data sets in a computationally efficient manner given marginal posterior samples from the different experiments. We use the publicly available code (<https://github.com/htjb/margarine> (accessed on 26 October 2022)) MARGARINE [9] to generate density estimators.

In Section 2, we mathematically demonstrate that the application of MARGARINE to the problem of combining the marginal posteriors from two data sets is equivalent to running a full nested sampling run including all ‘nuisance’ parameters. We define in this section the nuisance-free likelihood. Section 3 briefly discusses the methodology behind MARGARINE with reference to a previously published work [9]. Finally, we show the results of combining samples from DES and Planck in Section 4 and conclude in Section 5.

2. Theory

2.1. Notation

Given a likelihood $\mathcal{L}(\Theta) \equiv P(D|\Theta, \mathcal{M})$ representing the probability of data D given some model \mathcal{M} with parameters Θ , Bayesian inference proceeds by defining a prior $\pi(\Theta) \equiv P(\Theta|\mathcal{M})$, and then through Bayes theorem computing a posterior distribution $\mathcal{P}(\Theta) \equiv P(\Theta|D, \mathcal{M})$ for the purposes of parameter estimation and an evidence $\mathcal{Z} \equiv P(D|\mathcal{M})$ in order to perform model comparison. In our notation, we suppress model dependence, but where we wish to refer to the likelihoods derived from different datasets, we denote this with a subscript so for example $\mathcal{L}_A(\Theta) \equiv P(D_A|\Theta, \mathcal{M})$, and $\mathcal{L}_B(\Theta) \equiv P(D_B|\Theta, \mathcal{M})$.

In our setting, the parameter vector is split into two sub-vectors $\Theta = (\theta, \alpha)$, where θ are parameters of scientific interest, and α are nuisance parameters, including for the purposes of data analysis. Such situations are common in astrophysics, where for example θ might be parameters governing the Universe’s evolution, whilst α might be associated with instrument calibration and foreground removal [3,4]. The α parameters are generally “marginalised out” and not considered further in final or future analyses.

2.2. Definitions

With this notation established, the Bayes theorem version (including nuisance parameters) takes the form

$$\mathcal{L}(\theta, \alpha) \times \pi(\theta, \alpha) = \mathcal{P}(\theta, \alpha) \times \mathcal{Z}, \quad (1)$$

where we have placed the inputs of inference (likelihood and prior) on the left-hand side, and the outputs (posterior and evidence) on the right. The evidence is as usual equivalent to the (fully-)marginalised likelihood $\mathcal{Z} = \int \mathcal{L}(\theta, \alpha) \pi(\theta, \alpha) d\theta d\alpha$.

We may marginalise any probability distribution so can straightforwardly define the nuisance marginalised posterior and prior by integrating over α

$$\mathcal{P}(\theta) = \int \mathcal{P}(\theta, \alpha) d\alpha, \quad \pi(\theta) = \int \pi(\theta, \alpha) d\alpha. \quad (2)$$

The nuisance marginalised version of Bayes theorem Equation (1) takes the form

$$\mathcal{L}(\theta) \times \pi(\theta) = \mathcal{P}(\theta) \times \mathcal{Z}. \quad (3)$$

Here \mathcal{Z} is the original evidence, whilst $\mathcal{L}(\theta)$ is non-trivially the nuisance-free likelihood

$$\mathcal{L}(\theta) \equiv \frac{\int \mathcal{L}(\theta, \alpha) \pi(\theta, \alpha) d\alpha}{\int \pi(\theta, \alpha) d\alpha} = \frac{\mathcal{P}(\theta) \mathcal{Z}}{\pi(\theta)}, \tag{4}$$

where the above is motivated by marginalising over α of the full Bayes theorem Equation (1) and substituting the definitions in Equations (2) and (4) recovers the marginalised Bayes theorem Equation (3).

The nuisance-free likelihood Equation (4) is straightforward to compute in our framework since we (uniquely) have NS-computed evidence \mathcal{Z} combined with MARGARINE trained distributions $\mathcal{P}(\theta)$ and $\pi(\theta)$ [9] (see Section 3).

We now explain why Equation (4) is a useful definition

Theorem 1. Let $\mathcal{L}_A(\theta, \alpha_A)$ and $\mathcal{L}_B(\theta, \alpha_B)$ be two likelihoods with distinct datasets, each with their own nuisance parameters. The nuisance-free likelihoods $\mathcal{L}_A(\theta)$, $\mathcal{L}_B(\theta)$ form a lossless compression in θ . This means that we can recover the same (marginal) inference in combination that we would have made when performing a combined analysis with all nuisance parameters:

$$\mathcal{L}_A(\theta, \alpha_A) \mathcal{L}_B(\theta, \alpha_B) \pi_{AB}(\theta, \alpha_A, \alpha_B) = \mathcal{P}_{AB}(\theta, \alpha_A, \alpha_B) \mathcal{Z}_{AB}, \tag{5}$$

$$\Rightarrow \mathcal{L}_A(\theta) \mathcal{L}_B(\theta) \pi(\theta) = \mathcal{P}_{AB}(\theta) \mathcal{Z}_{AB}, \tag{6}$$

if their respective priors $\pi_A(\theta, \alpha_A)$ and $\pi_B(\theta, \alpha_B)$ satisfy the marginal consistency relations:

$$\pi(\theta) = \int \pi_A(\theta, \alpha_A) d\alpha_A = \int \pi_B(\theta, \alpha_B) d\alpha_B, \tag{7}$$

$$\int \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_A = \pi_B(\theta, \alpha_B), \quad \int \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_B = \pi_A(\theta, \alpha_A). \tag{8}$$

This process is represented graphically in Figures 1 and 2.

Proof. Integrating the combined Bayes theorem Equation (5) with respect to α_B , applying the definition of the marginal posterior Equation (2) on the right-hand side, and drawing out terms independent of α_B on the left, yields

$$\mathcal{L}_A(\theta, \alpha_A) \int \mathcal{L}_B(\theta, \alpha_B) \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_B = \mathcal{P}_{AB}(\theta, \alpha_A) \mathcal{Z}_{AB}. \tag{9}$$

From the definition of a nuisance-free likelihood Equation (4), and the marginal consistency Equation (8), we can say that the integral on the left-hand side becomes:

$$\begin{aligned} & \int \mathcal{L}_B(\theta, \alpha_B) \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_B \\ &= \int \mathcal{L}_B(\theta, \alpha_A, \alpha_B) \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_B && [\mathcal{L}_B(\theta, \alpha_B) \equiv \mathcal{L}_B(\theta, \alpha_A, \alpha_B) \text{ since } \mathcal{L}_B \text{ indep of } \alpha_A] \\ &= \mathcal{L}_B(\theta, \alpha_A) \int \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_B && [\text{Using Equation (4) for } \mathcal{L}_B] \\ &= \mathcal{L}_B(\theta) \int \pi_{AB}(\theta, \alpha_A, \alpha_B) d\alpha_B && [\mathcal{L}_B(\theta, \alpha_A) \equiv \mathcal{L}_B(\theta) \text{ since } \mathcal{L}_B \text{ indep of } \alpha_A] \\ &= \mathcal{L}_B(\theta) \pi_A(\theta, \alpha_A). && [\text{Using marginal consistency Equation (8)}] \end{aligned} \tag{10}$$

Substituting Equation (10) back into Equation (9) we find

$$\mathcal{L}_A(\theta, \alpha_A) \mathcal{L}_B(\theta) \pi_A(\theta, \alpha_A) = \mathcal{P}_{AB}(\theta, \alpha_A) \mathcal{Z}_{AB}. \tag{11}$$

Proceeding with a similar manipulation to Equation (10), marginalising with respect to α_A , and applying the definition of the nuisance-free likelihood $\mathcal{L}_A(\theta)$ Equation (4) and the marginal prior consistency Equation (7) we recover Equation (6)

$$\mathcal{L}_A(\theta)\mathcal{L}_B(\theta)\pi(\theta) = \mathcal{P}_{AB}(\theta)\mathcal{Z}_{AB}.$$

□

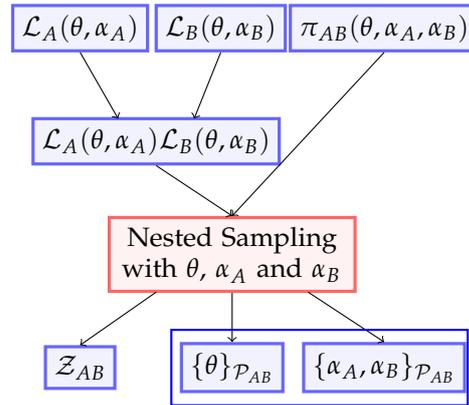


Figure 1. A graphical representation of combining constraints from different data sets via a full nested sampling run over both cosmological and nuisance parameters (Equation (5)).

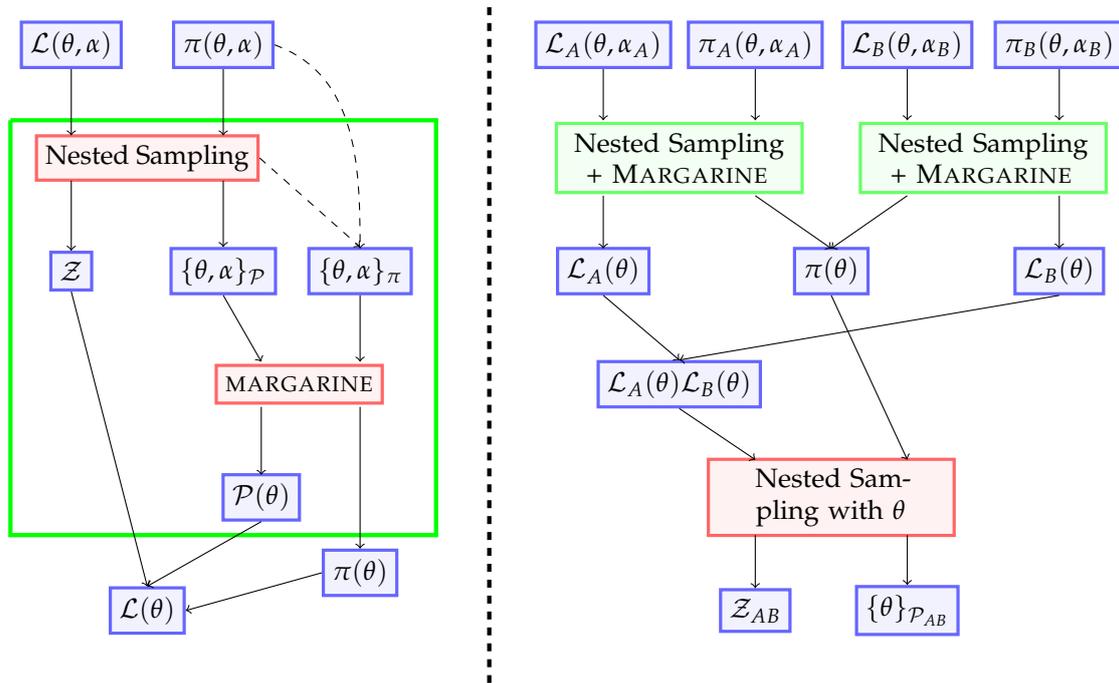


Figure 2. A graphical representation of combining constraints from two data sets via MARGARINE (Equation (6)). Left of the dashed line illustrates the derivation of a nuisance-free likelihood function for one experimental data set.

2.3. Discussion

Equation (5) represents Bayes theorem for the combined likelihood of both datasets $\mathcal{L}_{AB}(\theta, \alpha_A, \alpha_B) = \mathcal{L}_A(\theta, \alpha_A)\mathcal{L}_B(\theta, \alpha_B)$, using the combined prior $\pi_{AB}(\theta, \alpha_A, \alpha_B)$. We assume the combined prior is marginally consistent, Equations (7) and (8), which is reasonable, merely demanding that the priors are identical in the parameter spaces where they overlap. In practice, this would usually be achieved by assuming separability between

signal and nuisance parameter spaces $\pi(\theta, \alpha) = \pi(\theta)\pi(\alpha)$, but Equations (7) and (8) are a slightly less restrictive requirement and therefore more general.

The upshot of this is that if you have performed inference for two datasets separately, such that you are able to compute the nuisance-free likelihoods with MARGARINE, you may discard the nuisance parameters for the next set of analyses when you combine the datasets.

3. Methods

MARGARINE was first introduced in [9] and uses density estimation to approximate probability distributions such as $\mathcal{P}(\theta)$ and $\pi(\theta)$ given sets of representative samples. The code was developed initially to calculate marginal Kullback–Leibler (KL) divergences [10] and Bayesian model dimensionalities (BMD) [11] however as discussed in Section 2 it can be used to calculate the nuisance-free likelihoods. This in turn means that we can use MARGARINE alongside an implementation of the nested sampling algorithm to sample the product $\mathcal{L}_A(\theta)\mathcal{L}_B(\theta)$. In this manner, MARGARINE allows us to combine constraints on common parameters across different data sets.

We refer the reader to [9] for a complete discussion of how MARGARINE works; however, we discuss briefly the density estimation here. MARGARINE uses two different types of density estimators to model posterior and prior samples, namely masked autoregressive flows (MAFs, [8]) and kernel density estimators (KDEs, [6,7]).

MAFs transform a multivariate base distribution, the standard normal, into a target distribution via a series of shifts and scaling, which are estimated by autoregressive neural networks. To improve the performance of the MAF the samples representing the target distribution, in our case $\mathcal{P}(\theta)$ and $\pi(\theta)$, are transformed into a Gaussianized space. We implement the MAFs using TENSORFLOW and the KERAS backend [12].

KDEs use a kernel to approximate the multivariate probability density of a series of samples. In our case, the kernel is Gaussian and the probability density is a sum of Gaussians centred on the sample points with a given bandwidth. Again, we transform the target samples into a Gaussianized parameter space allowing the KDE to better capture the distribution. The KDEs are implemented with SCIPY in MARGARINE [6,7,13].

Since both types of density estimator build approximations to the target distribution using known distributions, the approximate log probabilities of the target distribution can be easily calculated.

The evaluation of normalised log probabilities for the marginal posterior and marginal prior allows us to calculate the nuisance-free likelihoods, as discussed, along with marginal Kullback–Leibler divergences

$$\mathcal{D}(\mathcal{P}||\pi) = \int \mathcal{P}(\theta) \log \frac{\mathcal{P}(\theta)}{\pi(\theta)} d\theta, \quad (12)$$

which quantifies the amount of information gained when moving from the marginal prior to posterior.

4. Cosmological Example

It has previously been demonstrated that MARGARINE is capable of replicating complex probability distributions and approximating marginal Bayesian statistics such as the KL divergence and the BMD [9]. Here, we demonstrate the theory discussed in Section 2 by combining samples from the Dark Energy Survey (DES) Year 1 posterior [1] and Planck posterior [3] using MARGARINE to estimate nuisance-free likelihoods. DES surveys supernovae, galaxies, and large-scale cosmic structures in the universe in an effort to measure dark matter and dark energy densities and model the dark energy equation of state. In contrast, Planck mapped the anisotropies in the cosmic microwave background (CMB) and correspondingly provided constraints on key cosmological parameters.

The constraints from DES and Planck have previously been combined using a full nested sampling run over all parameters including a multitude of ‘nuisance’ parameters in a computationally expensive exercise [14]. This corresponds to the flow chart in Figure 1 and the previous

analysis gives us access to the combined DES and Planck evidence, which is found to have a value of $\log(\mathcal{Z}) = -5965.7 \pm 0.3$. In Figure 3, we show the DES, Planck, and joint posteriors for the six cosmological parameters derived in this work using MARGARINE and the flow chart in Figure 2. The constrained parameters are the baryon and dark matter density parameters, $\Omega_b h^2$ and $\Omega_c h^2$, the angular size of the sound horizon at recombination, θ_{MC} , the CMB optical depth, τ , the amplitude of the power spectrum, A_s , and the corresponding spectral index, n_s . These make up the set $\theta = (\Omega_b h^2, \Omega_c h^2, \theta_{MC}, \tau, A_s, n_s)$. We use the nested sampling algorithm POLYCHORD in our analysis [15,16].

We use a uniform prior that is defined to be three sigma around the Planck posterior mean. This is done to improve the efficiency of our nested sampling run. However, we subsequently have to re-weight the samples and correct the evidence for the difference between the priors used here and in the previous full nested sampling run [14] for comparison. If we define

$$\mathcal{Z}_A = \int \mathcal{L}(\theta) \pi_A(\theta) d\theta, \quad \mathcal{Z}_B = \int \mathcal{L}(\theta) \pi_B(\theta) d\theta, \quad (13)$$

where A is our uniform prior space and B is our target prior space from the previous work [14], then

$$\mathcal{Z}_B = \int \mathcal{L}(\theta) \pi_B(\theta) d\theta = \int \mathcal{L}(\theta) \pi_A(\theta) \frac{\pi_B(\theta)}{\pi_A(\theta)} d\theta = \mathcal{Z}_A \left\langle \frac{\pi_B(\theta)}{\pi_A(\theta)} \right\rangle_{\mathcal{P}_A} \quad (14)$$

giving

$$\mathcal{Z}_B = \mathcal{Z}_A \left\langle \frac{\pi_B(\theta)}{\pi_A(\theta)} \right\rangle_{\mathcal{P}_A}. \quad (15)$$

Following similar arguments, we can transform our posteriors by re-weighting the distributions with the following

$$w_B^{(i)} = w_A^{(i)} \frac{\pi_B(\theta^{(i)})}{\pi_A(\theta^{(i)})}. \quad (16)$$

We see from the figure and corresponding table that with our joint analysis we are able to derive a log-evidence that is approximately consistent with that found in [14] validating the theory discussed and its implementation with MARGARINE. We note that the re-weighting described above relies on the calculations of the two prior log probabilities for which we use MARGARINE, and currently do not have an estimate of the error for. As a result, the error in the combined evidence, \mathcal{Z}_B , is given by the error in \mathcal{Z}_A from the nested samples and is likely underestimated. Using MARGARINE [9] we can also derive the combined KL divergence, also reported in Figure 3, which we find is consistent with the result in the literature of $\mathcal{D} = 6.17 \pm 0.36$ [11]. Similarly, we derive marginal KL divergences for the DES and Planck cosmological parameters using MARGARINE. A full discussion of the implications of combining the two data sets for our understanding of cosmology can be found in the literature (e.g., [11,14]).

By reducing the number of parameters that need to be sampled, we significantly reduce the nested sampling runtime. For POLYCHORD the runtime scales as the cube of the number of dimensions [17]. This can be seen by assessing the time complexity of the algorithm, where $T \propto n_{\text{live}} \times \langle T\{\mathcal{L}(\theta)\} \rangle \times \langle T\{\text{Impl.}\} \rangle \times \mathcal{D}(\mathcal{P}||\pi)$. Here n_{live} scales with the number of dimensions, d , as does the Kullback–Leibler divergence. For POLYCHORD, the specific implementation time complexity factor, $\langle T\{\text{Impl.}\} \rangle$, representing the impact of replacing dead points with higher likelihood live points on the runtime, scales linearly with d . Together this gives $T \propto d^3 \times \langle T\{\mathcal{L}(\theta)\} \rangle$. Therefore, by using nuisance-free likelihoods and sampling over 6 parameters rather than 41 parameters (cosmological plus 20 nuisance parameters for DES and 15 different nuisance parameters for Planck) we reduce the runtime by a factor of $(41/6)^3 \approx 319$ with further improvements in $\langle T\{\mathcal{L}(\theta)\} \rangle$. Using MARGARINE, $\langle T\{\mathcal{L}(\theta)\} \rangle$ is typically reduced since analytic likelihoods are computationally more expensive than emulated likelihoods.

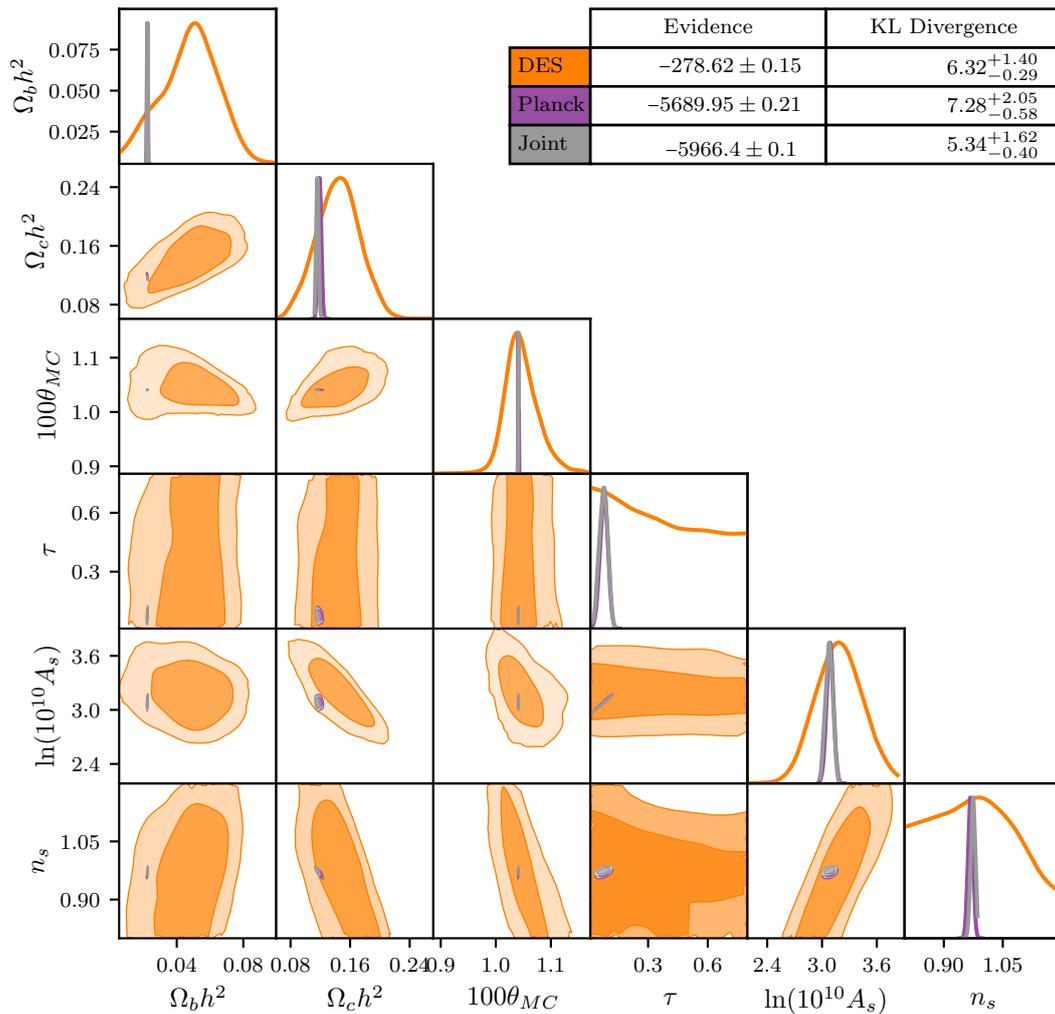


Figure 3. The combined posterior (in grey) found when combining the constraints on the cosmological parameters from DES and Planck using MARGARINE. For DES and Planck, we calculate the marginal KL divergences using MARGARINE, whereas the Bayesian evidence is calculated using ANESTHETIC. The joint evidence and joint KL divergence are calculated with a combination of the two codes and are found to be approximately consistent with those found in the literature [11,14]. Note that the error on the joint evidence is likely underestimated as it relies on evaluations of log probabilities for various distributions, for which MARGARINE does not currently provide errors. The figure was produced with ANESTHETIC [18].

5. Conclusions

In the paper, we demonstrated the consistency between combining constraints from different experiments in a marginal framework using density estimators and the code MARGARINE with a full nested sampling run over all parameters, including those describing ‘nuisance’ components of the data. We have shown this consistency mathematically and with a cosmological example. For the combination of Planck and DES, we found Bayesian evidence and KL divergence that were consistent with previous results [11,14].

The analysis in this paper is only possible because (a) we were able to estimate densities in the (much smaller) cosmological parameter space θ using MARGARINE, and (b) because we have evidence, \mathcal{Z} , from our original nested sampling runs. It is this unique combination that allowed us to compress away or discard nuisance parameters once they have been used. We note also that working in the marginal space results in a compression that is lossless in information on θ as it recovers an identical marginal posterior and total evidence as is found during a full Bayesian inference. Finally, through the nuisance-free likelihood

we can significantly reduce the dimensionality of our problems and since it is faster to emulate a likelihood rather than analytically evaluate, MARGARINE offers a much more computationally efficient path to combined Bayesian analysis.

In principle, our work paves the way for the development of a publicly available library of cosmological density estimators modelled with MARGARINE that can be combined with new data sets using the proposed method in a more efficient manner than currently implemented techniques. However, the work has implications outside of cosmology in any field where multiple experiments probe different aspects of the same physics.

Author Contributions: Conceptualization, W.H. and H.B.; methodology, W.H. and H.B.; formal analysis, H.B.; software, H.B.; writing—original draft preparation, H.B.; writing—review and editing, H.B., W.H., P.L. and P.S.; supervision, W.H., E.d.L.A. and A.F. All authors have read and agreed to the published version of the manuscript.

Funding: H.B. acknowledges the support of the Science and Technology Facilities Council (STFC) through grant number ST/T505997/1 and Fitzwilliam College, Cambridge. W.H. and A.F. were supported by Royal Society University Research Fellowships. PHS acknowledges support from a McGill Space Institute Fellowship and the Canada 150 Research Chairs Program. E.d.L.A. was supported by the STFC through the Ernest Rutherford Fellowship.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The Planck and DES posterior samples used in this paper are available in Zenodo at 10.5281/zenodo.4116393.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. DES Collaboration. Dark Energy Survey year 1 results: Cosmological constraints from galaxy clustering and weak lensing. *Phys. Rev. D* **2018**, *98*, 043526. [[CrossRef](#)]
2. DES Collaboration. Dark Energy Survey year 3 results: Covariance modelling and its impact on parameter estimation and quality of fit. *MNRAS* **2021**, *508*, 3125–3165. [[CrossRef](#)]
3. Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *A&A* **2020**, *641*, A6. [[CrossRef](#)]
4. Anstey, D.; Acedo, E.d.L.; Handley, W. A General Bayesian Framework for Foreground Modelling and Chromaticity Correction for Global 21 cm Experiments. *arXiv* **2020**, arXiv:2010.09644.
5. Bevins, H.T.J.; de Lera Acedo, E.; Fialkov, A.; Handley, W.J.; Singh, S.; Subrahmanyam, R.; Barkana, R. A Comprehensive Bayesian re-analysis of the SARAS2 data from the Epoch of Reionization. *arXiv* **2022**, arXiv:2201.11531.
6. Parzen, E. On Estimation of a Probability Density Function and Mode. *Ann. Math. Stat.* **1962**, *33*, 1065–1076. [[CrossRef](#)]
7. Rosenblatt, M. Remarks on Some Nonparametric Estimates of a Density Function. *Ann. Math. Stat.* **1956**, *27*, 832–837. [[CrossRef](#)]
8. Papamakarios, G.; Pavlakou, T.; Murray, I. Masked Autoregressive Flow for Density Estimation. *arXiv* **2017**, arXiv:1705.07057.
9. Bevins, H.T.J.; Handley, W.J.; Lemos, P.; Sims, P.H.; de Lera Acedo, E.; Fialkov, A.; Alsing, J. Removing the fat from your posterior samples with margarine. *arXiv* **2022**, arXiv:2205.12841.
10. Kullback, S.; Leibler, R.A. On Information and Sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86. [[CrossRef](#)]
11. Handley, W.; Lemos, P. Quantifying dimensionality: Bayesian cosmological model complexities. *Phys. Rev. D* **2019**, *100*, 023512. [[CrossRef](#)]
12. Abadi, M.; Agarwal, A.; Barham, P.; Brevdo, E.; Chen, Z.; Citro, C.; Corrado, G.S.; Davis, A.; Dean, J.; Devin, M.; et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. 2015. Available online: <https://www.tensorflow.org/> (accessed on 26 October 2022).
13. Virtanen, P.; Gommers, R.; Oliphant, T.E.; Haberland, M.; Reddy, T.; Cournapeau, D.; Burovski, E.; Peterson, P.; Weckesser, W.; Bright, J.; et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **2020**, *17*, 261–272. [[CrossRef](#)] [[PubMed](#)]
14. Handley, W.; Lemos, P. Quantifying tensions in cosmological parameters: Interpreting the DES evidence ratio. *Phys. Rev. D* **2019**, *100*, 043504. [[CrossRef](#)]
15. Handley, W.J.; Hobson, M.P.; Lasenby, A.N. PolyChord: Nested sampling for cosmology. *Mon. Not. R. Astron. Soc.* **2015**, *450*, L61–L65. [[CrossRef](#)]
16. Handley, W.J.; Hobson, M.P.; Lasenby, A.N. PolyChord: Next-generation nested sampling. *Mon. Not. R. Astron. Soc.* **2015**, *453*, 4385–4399. [[CrossRef](#)]

17. Petrosyan, A.; Handley, W.J. Supernest: Accelerated nested sampling applied to astrophysics and cosmology. In Proceedings of the International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering, Paris, France, 18–22 July 2022.
18. Handley, W. Anesthetic: Nested sampling visualisation. *J. Open Source Softw.* **2019**, *4*, 1414. [[CrossRef](#)]