

Non-Gaussian estimates of tensions in cosmological parameters

Marco Raveri¹ and Cyrille Doux¹

*Center for Particle Cosmology, Department of Physics and Astronomy,
University of Pennsylvania, Philadelphia, Pennsylvania 19104, USA*



(Received 17 May 2021; accepted 28 June 2021; published 2 August 2021)

We discuss how to efficiently and reliably estimate the level of agreement and disagreement on parameter determinations from different experiments, fully taking into account non-Gaussianities in the parameter posteriors. We develop two families of scalable algorithms that allow us to perform this type of calculations in an increasing number of dimensions and for different levels of tensions. One family of algorithms rely on kernel density estimates of posterior distributions while the other relies on machine learning modeling of the posterior distribution with normalizing flows. We showcase their effectiveness and accuracy with a set of benchmark examples and find both methods agree with each other and the true tension within 0.5σ in difficult cases and generally to 0.2σ or better. This allows us to study the level of internal agreement between different measurements of the clustering of cosmological structures from the Dark Energy Survey and their agreement with measurements of the cosmic microwave background from the Planck satellite.

DOI: [10.1103/PhysRevD.104.043504](https://doi.org/10.1103/PhysRevD.104.043504)

I. INTRODUCTION

Despite the success of the Λ cold dark matter (Λ CDM) model at fitting a vast range of cosmological observations, the emergence of increasingly statistically significant tensions on cosmological parameters has sparked interest in the problem of their precise calculation. In particular, the value of the Hubble constant, H_0 , inferred from observations of the cosmic microwave background [1,2] is lower than that inferred from the distance ladder at low redshifts using type Ia supernovae [3], at 4 to 6σ , which has prompted searches for hints of new physics [4][see][for a recent review]. Of less significance, but still puzzling, is the 1 to 2σ lower amplitude of structure found by weak lensing surveys [5,6] compared to that inferred from the CMB.

Many tension metrics and numerical tools have been developed and characterized in the literature (e.g., [7–26]). Most commonly Gaussian approximations at either the data or parameter posterior level need to be made when computing statistical significance. One open problem is how to precisely compute the statistical significance of a tension in the presence of significant non-Gaussianities in the posterior distributions of cosmological parameters. While the problem is conceptually straightforward, the large number of dimensions usually involved when comparing datasets of interest in cosmology makes its calculation challenging in practice.

In this paper, we first derive and characterize the posterior distribution of *differences* between parameters determined from two, possibly correlated, experiments. We then discuss two methods to compute the statistical

significance of a parameter difference, showcasing their reliability on a set of toy problems and two examples from cosmology. One class of methods relies on kernel density estimates with fixed and variable data-driven smoothing scales. Another method relies on machine learning modeling of the posterior distribution with normalizing flows. Both methods show remarkable accuracy on a set of benchmark examples that cover a wide range of dimensions (from 2 to 30), different known input tension levels (from 1 to 4σ) and different ways in which non-Gaussianities in the parameter posteriors might arise. We comment on the advantages and weaknesses of both strategies and develop efficient algorithms for their computation. Without these algorithms the calculation of statistical significance would be practically impossible in realistic scenarios. All the methods we discuss have to give the same result and are expected to converge for increasing posterior sample size so their spread can be used as a reliable global error estimate. Finally, we implement these non-Gaussian tension estimators in `tensiometer`,¹ that we use to perform all calculations in this paper and that we make publicly available.

As a first worked example we study the mutual consistency of measurements of galaxy lensing and clustering from the Dark Energy Survey year 1 data [DES Y1, [27]], finding them in excellent agreement. We then compute the level of agreement, within Λ CDM, of DES Y1 measurements and cosmic microwave background power spectrum measurements from the Planck satellite [1,2], finding them

¹<https://github.com/mraveri/tensiometer>

in disagreement at about the $p = 0.3\%$ (or 3σ) level. The former example shows the reliability of these estimates in a case with significant correlations and high dimensionality but low statistical significance, while the latter example showcases reliability in the case of uncorrelated measurements but higher statistical significance.

This paper is structured as follows: in Sec. II, we discuss how to obtain the posterior distribution of parameters differences, calculate their statistical significance and characterize its properties; in Sec. III, we show how to efficiently calculate statistical significance with kernel density estimates (KDE)-based methods; in Sec. IV, we discuss how to obtain statistical significance of parameter differences with normalizing flows; in Sec. V, we discuss a set of benchmark examples that are tailored to gauge the accuracy of both calculation methods; in Sec. VI, we apply both the KDE and normalizing flows estimators to data from Planck and DES; we discuss the results and conclude in Sec. VII.

II. PARAMETER DIFFERENCES

In this section we introduce the distribution of parameter differences and comment on its properties.

We start by defining general notations and writing the posterior probability of parameters θ , within model \mathcal{M} given data d as

$$\mathcal{P}(\theta) \equiv P(\theta|d, \mathcal{M}) = \frac{P(\theta|\mathcal{M})P(d|\theta, \mathcal{M})}{P(d|\mathcal{M})}, \quad (1)$$

where we indicate the prior probability as $\Pi(\theta) \equiv P(\theta|\mathcal{M})$, the likelihood as $\mathcal{L}(\theta) \equiv P(d|\theta, \mathcal{M})$ and the evidence as $\mathcal{E} \equiv P(d|\mathcal{M})$. We will drop \mathcal{M} hereafter for conciseness and indicate the support of the prior as V_{Π} . This is defined as the set of parameters where the prior is strictly nonvanishing.

We now consider two different datasets d_1 and d_2 , labeled here and throughout with 1 and 2 subscripts, described by a joint likelihood $\mathcal{L}(\theta) \equiv P(d_1, d_2|\theta)$ with shared parameters θ and prior $\Pi(\theta)$. We denote $\mathcal{L}_1(\theta) \equiv P(d_1|\theta)$ and $\mathcal{P}_1(\theta) \equiv P(\theta|d_1)$ the joint likelihood/posterior that has been marginalized over d_2 (likewise for d_1). The first step to build the parameter difference distribution consists in duplicating all shared parameters, i.e., create two copies of the parameter set, θ_1 and θ_2 , and then choose a prescription for the joint likelihood with duplicated parameters, $\mathcal{L}(\theta_1, \theta_2) \equiv P(d_1, d_2|\theta_1, \theta_2)$. As an example, if the joint likelihood is a multivariate Gaussian distribution over (d_1, d_2) with parameter-independent covariance, a typical choice is to use θ_1 (respectively θ_2) in the computation of the mean of d_1 (d_2). In general, however, there is not necessarily a unique solution and different data splits correspond to different aspects that are tested. We may simply impose the following constraints: (a) $\mathcal{L}(\theta_1 = \theta, \theta_2 = \theta) = \mathcal{L}(\theta)$ such that the likelihood with duplicated parameters coincides with

the joint likelihood on the subset where parameters are equal, and (b) $P(d_1|\theta_1, \theta_2) = P(d_1|\theta_1)$ such that the likelihood for d_1 is independent of θ_2 once marginalized over d_2 (and similarly even switching indices 1 and 2). These conditions ensure that, in the case where the data are conditionally independent given θ , i.e., $P(d_1, d_2|\theta) = P(d_1|\theta)P(d_2|\theta)$, then the only choice is to use the product of likelihoods, such that $\mathcal{L}(\theta_1, \theta_2) = \mathcal{L}_1(\theta_1)\mathcal{L}_2(\theta_2)$. We further assume that the joint prior distribution factorizes, $\Pi(\theta_1, \theta_2) = \Pi(\theta_1)\Pi(\theta_2)$. We can then write the posterior of duplicated parameters, in the general case, as $\mathcal{P}(\theta_1, \theta_2) \equiv P(\theta_1, \theta_2|d_1, d_2) \propto \mathcal{L}(\theta_1, \theta_2)\Pi(\theta_1)\Pi(\theta_2)$.

To obtain the probability density of parameter differences we can change variables and consider $\Delta\theta \equiv \theta_1 - \theta_2$. This gives a joint posterior $\mathcal{P}(\theta_1, \Delta\theta) = \mathcal{P}(\theta_1, \theta_1 - \Delta\theta)$, and we can integrate out the base parameters,

$$\mathcal{P}(\Delta\theta) = \int_{V_{\pi}} \mathcal{P}(\theta, \theta - \Delta\theta) d\theta. \quad (2)$$

We distinguish two cases based on whether the two datasets at hand are correlated or not,

- (i) If the datasets are conditionally independent given parameters θ , the posteriors \mathcal{P}_1 and \mathcal{P}_2 may be sampled independently, and differences of samples $\Delta\theta \equiv \theta_1 - \theta_2$, with $\theta_1 \sim \mathcal{P}_1$ and $\theta_2 \sim \mathcal{P}_2$, are samples of the difference distribution $\mathcal{P}(\Delta\theta)$, which is given by

$$\mathcal{P}(\Delta\theta) = \int_{V_{\pi}} \mathcal{P}_1(\theta)\mathcal{P}_2(\theta - \Delta\theta) d\theta. \quad (3)$$

This type of integral is known in the signal processing literature as the cross-correlation of \mathcal{P}_1 and \mathcal{P}_2 .

- (ii) If the datasets are correlated, we can obtain samples of $\mathcal{P}(\Delta\theta)$ by sampling the joint posterior with duplicated parameters and differentiating sample by sample.

Once the density of parameter differences is obtained, one can quantify the probability that there is in fact a parameter shift, given by

$$\Delta = \int_{\mathcal{P}(\Delta\theta) > \mathcal{P}(\mathbf{0})} \mathcal{P}(\Delta\theta) d\Delta\theta, \quad (4)$$

which is the posterior mass above the isocontour of no shift, $\Delta\theta = \mathbf{0}$. In the following we indicate the parameter set contained by the isocontour of no shift as $\Omega_0 \equiv \{\Delta\theta: \mathcal{P}(\Delta\theta) > \mathcal{P}(\mathbf{0})\}$.

Several basic properties can be easily shown from the definition.

- (i) If the posterior distributions are normalized, then so is the parameter difference distribution.
- (ii) The support of the parameter difference distribution is the support of the convolution of the characteristic

function of V_{Π} with itself. In particular if the support of the prior is an interval $[a, b]$ then the support of the prior for the parameter difference posterior is $[a - b, b - a]$.

- (iii) The distribution of parameter differences is always continuous at the prior boundary, even if the two single posteriors are not. Smoothness or continuity of derivatives is not guaranteed.
- (iv) Δ is symmetric for changes of datasets $1 \leftrightarrow 2$.
- (v) When datasets are uncorrelated and we have n_1 and n_2 samples, we can build up to $n_1 n_2$ samples of the parameter difference distribution. If the samples are weighted the weights of the parameter difference samples are the product of the two weights.
- (vi) If the two posteriors are uncorrelated Gaussian distributions with mean θ_1 and θ_2 and covariance C_1 and C_2 respectively, then the distribution of parameter differences is Gaussian with mean $\theta_1 - \theta_2$ and covariance $C_1 + C_2$. If the two parameters are correlated and C_{12} indicates their covariance, then the covariance of parameter differences is given by $C_1 + C_2 - C_{12} - C_{12}^T$ and the mean is unchanged.

In Fig. 1 we show an example of how the calculation of the parameter difference tension estimator proceeds. We start in the upper panel showing two Gaussian posterior distributions, with different parameter degeneracies and different means. We then obtain and show in the lower panel, the distribution of parameter differences, as the convolution of the two single distributions, as in Eq. (3). Since the two initial distributions are Gaussian the distribution of parameter differences is Gaussian as well. The support of the prior for the two separate posteriors is the $[-1, 1]$ square; hence the support of the prior of parameter differences is the $[-2, 2]$ square. We then find the isocontour that touches the origin. Integrating the posterior inside that isocontour indicates that it contains 0.995 of the posterior mass, which is the estimate of the probability that the parameters are actually different.

In the following sections we comment on other properties of parameter differences that are worth further discussion.

A. Unconstrained parameters

A couple of general results can be shown from the definition of the parameter difference distribution and help its characterization in cases where data do not constrain some of the parameters of the model.

The first consists in showing that parameters that are prior constrained for both datasets can never contribute a shift. If there is no data information on a parameter then the likelihood of both experiments is constant and the parameter shift integral becomes

$$P(\Delta\theta) \propto \int \Pi(\theta)\Pi(\theta - \Delta\theta)d\theta, \quad (5)$$




FIG. 1. A toy example of how the parameter difference tension calculation proceeds, starting from the two single parameter distributions in the upper panel, to the calculation of the distribution of parameter differences and its isocontour of zero parameter shift in light blue in the lower panel. In both panels different lines of the same color show the 68% and 95% C.L. regions of parameter space for different distributions, as shown in legend.

up to an irrelevant normalization constant. In signal processing literature this integral is called the autocorrelation of Π . By Cauchy inequality this is maximum for $\Delta\theta = \mathbf{0}$ so that the posterior can have no mass above it.

The second noticeable property is that parameters constrained by flat priors for one dataset but not the other can never contribute a shift. If the prior is flat across the support and that the second dataset provides no constraint (i.e., has a constant likelihood), we have

$$P(\Delta\theta) \propto \int \Pi(\theta)\mathcal{L}(\theta)\Pi(\theta - \Delta\theta)d\theta, \quad (6)$$

up to an irrelevant normalization constant. Equation (6) is integrating the likelihood over the intersection between the prior volume and the shifted prior volume. A nonvanishing $\Delta\theta$ reduces the integration volume of a positive function and hence reduces the integral. This means that Eq. (6) is maximum at zero shift.

As a consequence of this second property nonshared parameters between two experiments can never contribute

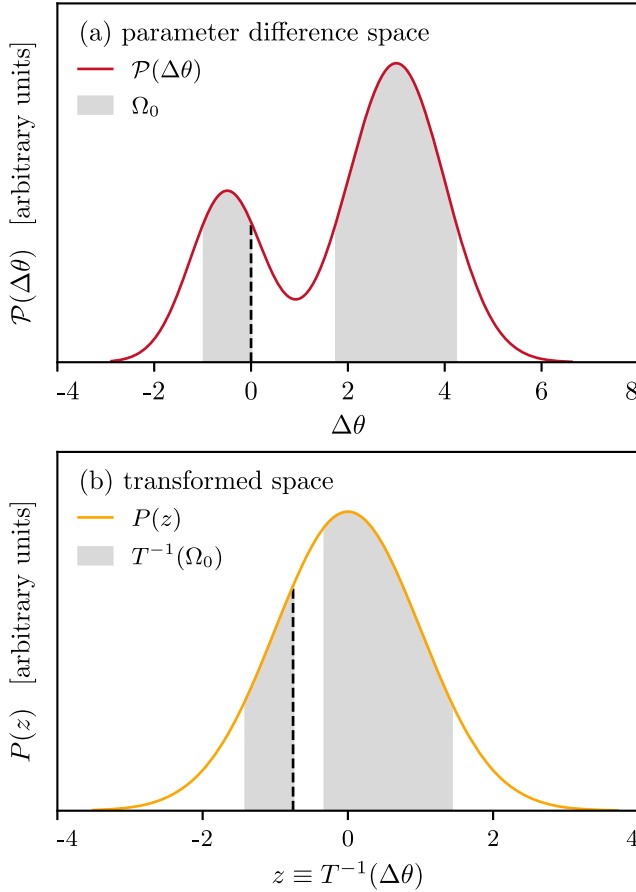


FIG. 2. A toy example of a Gaussianizing transformation showing the mapping of the isocontour at a given point.

any shift and can be premarginalized over before doing any parameter shift calculation.

B. Properties under changes of parameters

Since Eq. (4) is the integral of a probability density it is invariant under reparametrizations. However, when working directly with transformations of parameter differences some care must be exercised.

We consider a change of variables $\mathbf{z} = \mathbf{T}^{-1}(\Delta\theta)$ with an invertible, differentiable transformation \mathbf{T} . The probability density function of \mathbf{z} is given by

$$P(\mathbf{z}) = \mathcal{P}(\Delta\theta) |\det \nabla \mathbf{T}|, \quad (7)$$

computed with the determinant of the Jacobian of \mathbf{T} . The shift probability is then given by

$$\Delta = \int_{\Omega_0} \mathcal{P}(\Delta\theta) d\Delta\theta = \int_{T^{-1}(\Omega_0)} P(\mathbf{z}) d\mathbf{z}, \quad (8)$$

where both the posterior and integration measures change by the same factor that then cancels; but the domain of

integration is modified to the preimage of Ω_0 , which may not necessarily correspond to the zero isocontour of $P(\mathbf{z})$.

In Fig. 2, we consider the one-dimensional case where $\mathcal{P}(\Delta\theta)$ is a mixture of Gaussian distributions with two components. We apply the Gaussianizing transformation given by the successive application of the cumulative distribution function (CDF) of the Gaussian mixture and the inverse CDF of the standard normal distribution. The resulting transformation is highly nonlinear. The top panel shows the posterior $\mathcal{P}(\Delta\theta)$ and the bottom panel shows $P(\mathbf{z})$. The iso contour at zero-shift, Ω_0 , shown in gray in the top panel, has two connected components. Its preimage, shown in gray in the bottom panel, also has two components, and therefore $\mathbf{T}^{-1}(\Omega_0)$ clearly does not correspond to an isocontour of the Gaussian distribution $P(\mathbf{z})$. Note, however, that the posterior mass in the two gray areas is conserved and the same in the two cases.

III. KDE ESTIMATE OF PARAMETER DIFFERENCES

In this section we discuss how to compute the integral in Eq. (4) with kernel density estimates (KDE) of the parameter difference posterior. For an overview of KDE techniques we refer the reader to [28,29].

As we have highlighted in Sec. II when building the parameter difference chain we do not have likelihood values for samples. In this section we comment on how to obtain likelihood values through KDE. In this section only we use the shorthand notation of $P(\mathbf{x}) \equiv \mathcal{P}(\Delta\theta)$.

Given a kernel smoothing function, $K(\mathbf{x})$, that is positive, smooth, spherically symmetric and normalized we consider a bandwidth matrix \mathbf{H} which we assume to be symmetric and positive definite. The bandwidth matrix scales the kernel coordinates $K_{\mathbf{H}}(\mathbf{x}) \equiv |\mathbf{H}|^{1/2} K(\mathbf{H}^{-1/2} \mathbf{x})$ and controls the extent of smoothing in different directions. In practice we always assume that the smoothing kernel $K_{\mathbf{H}}(\mathbf{x})$ is Gaussian.

Given our n weighted samples, \mathbf{X}_i , of P , we can compute its KDE estimate as

$$\hat{P}_{\mathbf{H}}(\mathbf{x}) \equiv \frac{1}{w_{\text{tot}}} \sum_{i=1}^n w_i K_{\mathbf{H}}(\mathbf{x} - \mathbf{X}_i), \quad (9)$$

where w_i denotes the weights of the samples and $w_{\text{tot}} \equiv \sum_{i=1}^n w_i$. In abstract terms our problem consists in estimating the posterior mass inside the isocontour that touches a particular point, \mathbf{x}_0 , whose KDE estimate we indicate with

$$\hat{\Delta} = \int_{\hat{P}(\mathbf{x}) > \hat{P}(\mathbf{x}_0)} \hat{P}(\mathbf{x}) d\mathbf{x}. \quad (10)$$

To the best of our knowledge this problem has not been specifically addressed in KDE literature.

We notice that since the distribution of parameter differences is always continuous at the boundary we do not need to be concerned about correcting for KDE biases at the boundary.

When we consider high dimensional problems it is useful to slightly modify the KDE estimator to have adaptive bandwidth. This leads to the sample point density estimators,

$$\hat{P}_{\text{SP}}(\mathbf{x}) \equiv \frac{1}{w_{\text{tot}}} \sum_{i=1}^n w_i K_{\mathbf{H}_i}(\mathbf{x} - \mathbf{X}_i), \quad (11)$$

which considers a different bandwidth matrix for every sampled point \mathbf{H}_i . As discussed in [30] we can consider different strategies to fix the sampled point bandwidth. In particular we test the case where $\mathbf{H}_i = \delta_{(k)}(\mathbf{x}_i)^2 \mathbf{I}$, where $\delta_{(k)}$ is the distance to the k nearest neighbor of a given point. Note that for sample point density estimators there is no smoothing parameter to optimize. We tested other variable smoothing scale choices, in particular nearest neighbor ellipsoidal smoothing, but find that they never exceed the performances of the simpler nearest neighbor distance smoothing scale that we use.

A. Monte Carlo algorithms

Since we have samples from P we can perform the integration of Eq. (10) as a Monte Carlo (MC) volume integral,

$$\hat{\Delta} = \frac{1}{w_{\text{tot}}} \sum_{i=1}^n w_i S(\hat{P}(\mathbf{x}_i) - \hat{P}(\mathbf{X}_0)), \quad (12)$$

where $S(x)$ is the Heaviside step function that is unity when $x > 0$ and zero otherwise. We note that we have here assumed that the samples from the parameter difference chain are uncorrelated. This is generally easily achieved by suitably thinning the input parameter samples in order to obtain uncorrelated samples. Since we are counting samples above a threshold, the confidence intervals for the result can be estimated from the binomial distribution, through the Clopper–Pearson formula for the sum of weights above threshold and total sum of weights. We notice that the Clopper–Pearson formula applies to non-integer values of trials and successes so we can use it with noninteger weights.

We can compute $\hat{\Delta}$ with a *brute force* algorithm that iterates through the samples, computes KDE probability of a sample, compares its value to the precomputed reference and counts the sample or not. The drawback of this algorithm is that it clearly requires $\mathcal{O}(n^2)$ operations making it unfeasible for high statistically significant tensions that require a large number of samples to be accurately estimated.

For high statistically significant tension, we develop an *ad hoc* algorithm that we call *neighbor elimination*. We notice that Eq. (12) involves a threshold crossing problem, whether $\hat{P}(\mathbf{X}_i) > \hat{P}(\mathbf{x}_0)$ or not, but does not rely on the actual value of $\hat{P}(\mathbf{X}_i)$. Each of the $\hat{P}(\mathbf{X}_i)$ is computed from Eq. (9) and is a sum of positive quantities that decrease with distance from \mathbf{X}_i , since the kernel is positive. Then, after precomputing $\hat{P}(\mathbf{x}_0)$, we start by organizing samples in a k -d tree, which costs $\mathcal{O}(n \log n)$ operations. We then iterate through the samples and query the tree for the nearest neighbor of each point, which globally costs $\mathcal{O}(n \log n)$ operations. We build the sum in Eq. (9) going through the nearest samples, and since the kernel is positive we can discard a point in the tree as soon as it crosses the threshold of $\hat{P}(\mathbf{x}_0)$. For a large tension, where $\hat{P}(\mathbf{x}_0)$ is small and in the tails of the distribution, it only takes a few nearest neighbors to be able to discard a point. If the vast majority of points are discarded after a search of α nearest neighbor then the calculation of Eq. (10) only requires $\mathcal{O}((\alpha + 1)n \log n)$ operations. Notice that in the worst case scenario, where all points are above threshold, this algorithm requires $\mathcal{O}(n^2 \log n)$ operations so it is only suited for the calculation of statistical significance of large tensions.

Problems in one or two dimensions are, in this respect, special, since KDE evaluation can be performed with fast Fourier transforms (FFT) on a discrete grid with n_{grid} points that is then interpolated. This decreases the cost of operations to be globally $\mathcal{O}(n_{\text{grid}} n \log n)$ and makes it clearly faster than *neighbor elimination* below three dimensions. For and above dimension three, we expect the *neighbor elimination* algorithm to outperform FFT methods mostly because of the increasing size of the grids involved in the FFT calculation.

Notice that the *brute force* and the *neighbor elimination* algorithms have to agree exactly on the result while the FFT algorithm can differ with dependence on the size of the grid.

B. KDE bandwidth selection

The last ingredient we need to estimate Eq. (4) with KDE is a selection of the bandwidth matrix \mathbf{H} . We refer the reader to [29] for a comprehensive discussion of multidimensional band selectors. Here we review the main ideas and outline the strategies that we compare in the next sections.

Although we phrase the following in general terms, in practice, we always work with prewhitened samples, scaling all samples as $\tilde{\mathbf{x}} = \Sigma^{-1/2} \mathbf{x}$, where Σ is the sample parameter covariance. Doing this ensures that sample variance is the identity matrix. This cannot generally be done in all applications and we exploit here explicitly the independence of our resulting integral from the choice of coordinates.

The first family of band selectors that we consider aim at minimizing the mean integrated square error (MISE),

$$\text{MISE}(f) \equiv \int \mathbb{E}[(f(\mathbf{x}) - \hat{f}(\mathbf{x}))^2] d\mathbf{x}, \quad (13)$$

where \mathbb{E} denotes expectation over independent and identically distributed samples realizations. Minimizing this quantity is a very common approach as it seeks to minimize differences between the true distribution and its KDE approximation over the whole domain, resulting in a KDE estimate that is overall good for general purposes.

As noted in the two seminal papers [31,32], and can be verified by direct calculation, Eq. (13) can be written exactly in terms of convolutions between the true distribution f and the smoothing kernel $K_{\mathbf{H}}$,

$$\begin{aligned} \text{MISE}(f) = & R(K_{\mathbf{H}} * f) - 2 \int (K_{\mathbf{H}} * f)(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} + R(f) \\ & + n_{\text{eff}}^{-1} |\mathbf{H}|^{-1/2} R(K) - n_{\text{eff}}^{-1} R(K_{\mathbf{H}} * f), \end{aligned} \quad (14)$$

where $n_{\text{eff}} \equiv (\sum_i w_i)^2 / (\sum_i w_i^2)$, $R(f) \equiv \int f^2(\mathbf{x}) d\mathbf{x}$ and we indicate convolutions with $f * g \equiv \int f(\mathbf{x} - \mathbf{y}) g(\mathbf{y}) d\mathbf{y}$.

To make progress we assume that both the kernel and the target distributions are Gaussian, with covariance \mathbf{H} and Σ respectively, that allows to write Eq. (14) as

$$\begin{aligned} \text{MISE}(\mathbf{H}) = & (2\pi)^{-d/2} [2|\mathbf{H} + 2\Sigma|^{-1/2} - 2|\mathbf{H} + \Sigma|^{-1/2} \\ & + |2\Sigma|^{-1/2}] \\ & + n_{\text{eff}}^{-1} (2\pi)^{-d/2} [2|\mathbf{H}|^{-1/2} - |\mathbf{H} + 2\Sigma|^{-1/2}]. \end{aligned} \quad (15)$$

The matrix \mathbf{H}_{MISE} that minimizes Eq. (15) cannot be expressed analytically. Hence our strategy is to minimize Eq. (15) numerically over the space of positive symmetric matrices [33].

On the other hand if we consider the asymptotic expansion of Eq. (14) for large n and small \mathbf{H} , usually referred to as asymptotic MISE (AMISE), we can solve the corresponding minimization problem exactly in the Gaussian case by

$$\mathbf{H}_{\text{AMISE}} = \left(\frac{4}{(d+2)n_{\text{eff}}} \right)^{\frac{2}{d+4}} \Sigma. \quad (16)$$

This is known as the rule of thumb selector that we test and generally use as the starting guess for numerical minimization of MISE.

We test other possible choices for bandwidth selectors, including the maximum bandwidth and cross validation. In practice we find that their performances do not exceed the MISE/AMISE ones at the specific problem we are considering.

IV. NORMALIZING FLOW ESTIMATE OF PARAMETER DIFFERENCES

We now turn to a radically different method to compute the probability of no shift, Δ , based on building an analytic approximation of $\mathcal{P}(\Delta\theta)$. This method relies on normalizing flows (NF), that were introduced as generative models with tractable likelihoods, i.e., analytic approximation to the likelihood enabling fast sampling and inference [34–38]. NF have recently found applications in astrophysics, both for modeling [39] and Bayesian inference tasks [40–42].

A. Approximating difference distributions with normalizing flows

In this approach, one introduces a bijective, differentiable mapping \mathbf{T}_{φ} parametrized by neural networks with weights φ , between D -dimensional input data \mathbf{x} with distribution $P(\mathbf{x})$ and abstract variables $\mathbf{z} = \mathbf{T}_{\varphi}(\mathbf{x})$. The mapping is optimized over parameters φ with (stochastic) gradient descent to maximize the log-probability of \mathbf{z} under a simple distribution, typically the standard D -dimensional normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$, with density φ_D . One finally approximates the data distribution P with $P(\mathbf{x}) \approx q(\mathbf{x}) = \varphi_D(\mathbf{z}) |\det \nabla_{\mathbf{x}} \mathbf{T}_{\varphi}|$, using the determinant of the Jacobian of the transformation, as explained in Sec. II B. The density of the approximated distribution $q(\mathbf{x})$ can be computed analytically using automatic differentiation. One can also easily sample from q , by sampling from the D -dimensional standard Gaussian and transforming samples with the inverse mapping, $\mathbf{x} = \mathbf{T}_{\varphi}^{-1}(\mathbf{z})$.

For the problem at hand in this work, we approximate the parameter difference distribution $\mathcal{P}(\Delta\theta)$ with $q(\Delta\theta) = \varphi_D(\mathbf{z}) |\det \nabla_{\Delta\theta} \mathbf{T}_{\varphi}|$, where $\mathbf{z} = \mathbf{T}_{\varphi}(\Delta\theta)$ and \mathbf{T}_{φ} has been optimized to Gaussianize \mathbf{z} , as in Sec. II B. This is obtained by minimizing the loss function given by $L(\Delta\theta, \varphi) = -\log \varphi_D(\mathbf{T}_{\varphi}(\Delta\theta))$. In order to estimate the shift probability Δ , we sample $\Delta\theta \sim q$, compute the probability density of those samples, $q(\Delta\theta)$, and compare the results to $q(\Delta\theta = \mathbf{0})$, computed the same way. The fraction of samples for which $q(\Delta\theta) > q(\mathbf{0})$ provides a Monte Carlo estimate of the shift probability Δ , with an error that can be made arbitrary low by generating more samples and estimated with the Clopper-Pearson formula. The accuracy of the estimate of Δ is therefore limited by the accuracy of the normalizing flow model, rather than the number of samples drawn during this final step.

This technique transforms the problem of estimating Δ into an optimization problem. This means that it is not guaranteed to converge to the true value and instead, the result will depend on the choice of NF and neural network architecture, as well as drawbacks inherent to neural networks such as initialization, underfitting and overfitting. In particular, if the tension is strong and the no-shift point is far from the bulk of samples, the approximation to $\mathcal{P}(\mathbf{0})$

may be a poor one. Nonetheless, the computation—which involves training a neural network—now benefits from a large initial sample and suffers less than KDE from dimensionality [43]. Moreover, we find it efficient for small to mild tensions, which are interesting to characterize precisely, and thus this method complements the KDE method presented in the previous section.

B. Normalizing flows implementation

In this work, we use masked autoregressive flows [MAF, [34]] to approximate the parameter difference distribution $\mathcal{P}(\Delta\theta)$. MAFs are composed of a series of masked autoencoders for density estimation [MADE, [44]] that each implement autoregressive transformations of the input variables encoded by a single neural network. If the input variable is $\mathbf{x} = (x_1, \dots, x_n)$, the output \mathbf{y} of an autoregressive transformation has components $y_i = \mu(\mathbf{x}_{1:i-1}) + \sigma(\mathbf{x}_{1:i-1})x_i$, where μ and σ are parametrized by unconstrained neural networks that receive masked inputs, $\mathbf{x}_{1:i-1} = (x_1, \dots, x_{i-1}, 0, \dots, 0)$. MAFs allow for estimation of more complex distributions by stacking such transformations and introducing random permutations of the components between them.

In practice, we first compute the mean $\langle\Delta\theta\rangle$ and covariance Σ of the difference samples, $\Delta\theta$, to shift and rescale parameters, $\Delta\theta' = \Sigma^{-1/2}(\Delta\theta - \langle\Delta\theta\rangle)$, thus defining a bijective mapping into a prewhitened (but still non-Gaussian) space. We then apply MAFs to learn the mapping to the fully Gaussianized space and apply chain rules to approximate the difference parameter distribution \mathcal{P} . We find that MAFs are generally sufficiently flexible for typical parameter difference distributions, requiring only few hundreds to few thousands of trainable weights in the examples considered below.

Our code relies on the implementation of MAFs within the TensorFlow-Probability package [TFP, [45]], which benefits from the automatic differentiation of the TensorFlow² framework to compute the Jacobian of the transformation and other derivatives needed for gradient descent. We provide utility functions to quickly instantiate a default (but easily tweakable) architecture of MAFs composed of $2d$ stacked MADEs, each with two dense layers of $2d$ nodes each with sinh activation functions, and with random permutations of variables between MADEs (where d is the dimension of the parameter difference space). We found this architecture to work well for all the examples detailed below. However, the code is not restricted to using MAFs and can accommodate any user-provided normalizing flow implemented as a TFP `Bijector` object. We also found that the uniform Glorot initializer (of the weights and biases of MAFs) provided a good balance between learning speed and stability, although the user can also easily opt for any other initializer. We provide plotting utilities to monitor the

evolution of a number of metrics during training (see Fig. 8), in addition to the standard training and validation losses, to assert the quality of the approximation to the parameter distribution \mathcal{P} . In particular, since NFs learn a mapping to the standard multivariate Gaussian, we compute the squared norms of the transformed samples and verify that their distribution converges to a χ_d^2 distribution, e.g., with a Kolmogorov-Smirnov test.

V. BENCHMARK EXAMPLES

In this section, we perform numerical experiments on benchmark examples designed to test and compare our methods in limit cases.

We first consider, in Sec. VA, examples where the difference parameter distributions are Gaussian in an increasing number of dimensions. In these examples, the input tension is known analytically. We then consider, in Secs. VB–VD, two-dimensional models where the two posteriors, \mathcal{P}_1 and \mathcal{P}_2 , have analytic probability densities that are strongly non-Gaussian. The density of the difference distribution, $\mathcal{P}(\Delta\theta)$, can be computed on a dense grid and is given by the convolution of the posterior densities evaluated over dense grids covering the prior volume V_Π . We then compute the expected probability of no shift by numerical integration of $\mathcal{P}(\Delta\theta)$ over Ω_0 . In all cases, we compare the probability of no shift obtained with the numerical integration to the one obtained with the KDE and NF methods as well as the Gaussian approximation.

For the Gaussian approximation we compute parameter shifts both in standard form and in update form, as in [17]. The standard form of parameter shift, Q_{DM} , requires us to compute

$$Q_{\text{DM}} \equiv (\theta_{p1} - \theta_{p2})(\mathcal{C}_{p1} + \mathcal{C}_{p2})^{-1}(\theta_{p1} - \theta_{p2})^T, \quad (17)$$

where θ_{p1}/θ_{p2} and $\mathcal{C}_{p1}/\mathcal{C}_{p2}$ denote the posterior mean parameters and posterior covariance respectively for the dataset 1/2. Q_{DM} is χ^2 distributed with $\text{rank}(\mathcal{C}_{p1} + \mathcal{C}_{p2})$ degrees of freedom. The update parameter shifts estimator is defined as

$$Q_{\text{UDM}} \equiv (\theta_{p1} - \theta_{p12})(\mathcal{C}_{p1} - \mathcal{C}_{p12})^{-1}(\theta_{p1} - \theta_{p12})^T, \quad (18)$$

where θ_{p12} and \mathcal{C}_{p12} are the joint posterior mean parameters and covariance. Q_{UDM} is χ^2 distributed with $\text{rank}(\mathcal{C}_{p1} - \mathcal{C}_{p12})$ degrees of freedom. We refer the reader to [17] for the details of these estimators and their computation. Two noteworthy properties of these estimators, in this context, are that they have to coincide for Gaussian distributions with uninformative priors and that the Q_{UDM} estimator is symmetric when swapping datasets $1 \leftrightarrow 2$. In the non-Gaussian case though, the update estimator allows us to use the most Gaussian of the two distributions at hand hence providing mitigation of non-Gaussian features.

²<https://www.tensorflow.org/>

For KDE methods we avoid using the FFT algorithm in two dimensions since we want to focus on the properties of the Monte Carlo algorithm. We also use the *neighbor elimination* algorithm in all calculations since we find that it always exceeds the performances of the *brute-force* algorithm. We focus on the dependence of the KDE result on the number of samples that are used in the calculation and different bandwidth choices.

For the NF method, we focus on choices of hyperparameters, the impact of the number of samples used to train the model and model uncertainties. As previously explained, the stochasticity in the initialization of the neural networks and that of the gradient descent during training introduce extra variance in the estimated shift probability which is difficult to evaluate. Therefore, for each example, we perform ten runs and report the uncertainty of the results. For Gaussian examples in dimension d , we use MAFs made of $\min(d, 12)$ stacked MADEs, each with two layers of $2d$ units. We focus on the dependence of the shift probability estimates and their error as a function of the number of samples. For two-dimensional examples, we focus on hyperparameters and compare results for MAFs made of either four or eight stacked MADEs, each with two hidden layers of eight units (for a total of 528/1056 parameters), with fixed training samples with a size of 10^5 . Additionally, the learning rate needs some tuning for these non-Gaussian examples: if it is too large, it may prevent convergence, while if it is too small, the optimizer might fall into a local minimum far from the global one. Therefore, we repeat these tests for learning rates of 10^{-2} , 10^{-3} and 10^{-4} . For all tests, the estimated shift probability is obtained by sampling the approximate posterior q until the confidence interval on the result, estimated from the Clopper-Pearson formula, has a length smaller than 0.05.

In the following, we always report results as effective number of standard deviations. Given an event of probability P , it is given by

$$n_\sigma \equiv \sqrt{2}\text{Erf}^{-1}(P). \quad (19)$$

This corresponds to the number of standard deviations that an event with the same probability would have had were it drawn from a Gaussian distribution. This does not imply Gaussianity of the underlying statistics and should be regarded as a logarithmic scale for probabilities. Note that, when gauging accuracy of estimates, requirements of a precise determination of n_σ is a much stricter requirement than a threshold on absolute accuracy of the P value.

A. High-dimensional Gaussian

We start by considering a Gaussian distribution in an increasing number of dimensions. Since Eq. (2) preserves Gaussianity, if the two original distributions are Gaussian, so is the parameter difference distribution. We then generate a random parameter difference distribution by

generating a random mean vector on the sphere, a random covariance, and rescaling the mean vector length to match the desired benchmark tension.

1. Results with KDE

In Fig. 3 we show the error Δn_σ on the determination of a 1σ tension ($n_\sigma = 1$) using KDE estimates, and we compare the KDE estimate with the MISE smoothing scale to the KDE sample point (SP) estimate. For each example, we draw several batches of increasingly larger samples from the random parameter difference distribution in order to evaluate the bias of the estimator and its variance, as a function of number of samples n_{samples} and dimension d . We find that the MISE estimate converges more rapidly than the SP estimate with the number of samples. As it is known, however, its bias is, in all cases, larger than its variance. However, when increasing dimension at fixed sample size, samples become increasingly sparse and the MISE estimate cannot be computed, while the SP estimator retains workable accuracy in all dimensions, even with a relatively small number of samples.

This trend is confirmed in Fig. 4 where we show examples with higher tensions. Increasingly, samples

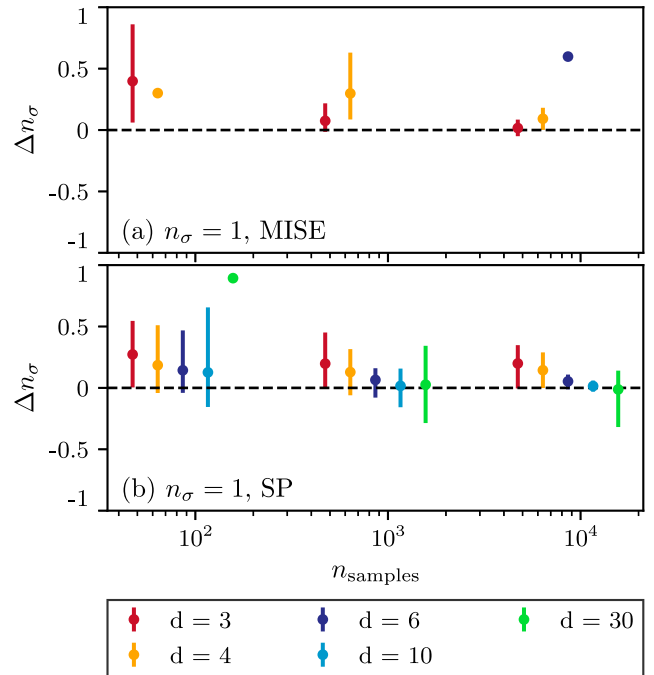


FIG. 3. The mean error Δn_σ on the determination of a 1σ tension ($n_\sigma = 1$) from KDE estimates in increasing number of dimensions as a function of number of samples n_{samples} and dimension d . Different panels show different KDE algorithms while different colors show different number of dimensions, as reported in legends. Error bars are obtained by drawing multiple samples for each case. Note that for high dimensions, samples become sparser and the MISE estimate cannot be computed, which explains missing points.




FIG. 4. The mean error Δn_σ on the determination of 2σ , 3σ and 4σ tensions from the KDE SP estimate in increasing number of dimensions as a function of number of samples n_{samples} and dimension d . Different panels show different input tension levels while different colors show different number of dimensions, as reported in legends. For high tensions, small samples lead to highly biased results that are not visible in the lower plot.

around the zero shift contour become sparse so the MISE estimate would require a very large number of samples to be meaningfully computed. On the other hand the SP estimate shows remarkable performances with error Δn_σ rarely exceeding 0.5σ independently of dimension and sample size. We note that the performances of the SP estimator improve as dimension increases and especially in high dimensions the variance of the estimator over pools of samples is of the same order as its bias.

For large number of samples the variance of the estimator is so small that it cannot be seen in the figure. Some results for low number of samples also show no variance or largely underestimated variance because they are generally borderline feasible, and the calculation fails many times because there is no sample outside the isocontour of no shift.

In all cases we note that the error scaling is worse than what we would expect from binomial trials. The reason why this happens is that the zero shift isocontour which is

the decision boundary in the binomial trial for each sample is inferred from the samples. Sampling uncertainty on its estimate then contribute to the total error budget.

We test several different algorithms and reach to qualitatively similar conclusions. The AMISE bandwidth leads to similar results as the MISE one, while cross-validation band estimators generally lead to a smaller smoothing scale, which then requires more samples to converge. Adaptive KDE with ellipsoidal smoothing scale show performances similar to the SP estimator.

2. Results with normalizing flows

In Fig. 5, we show the bias and variance of the NF estimate over ten runs, for increasing levels of tension and increasing dimensions, as a function of the number of training samples. In order to obtain a fair comparison, we fix the number of epochs during training to 100, each with 100 batches per epoch (possibly using samples several times per epoch). This sets the total number of model

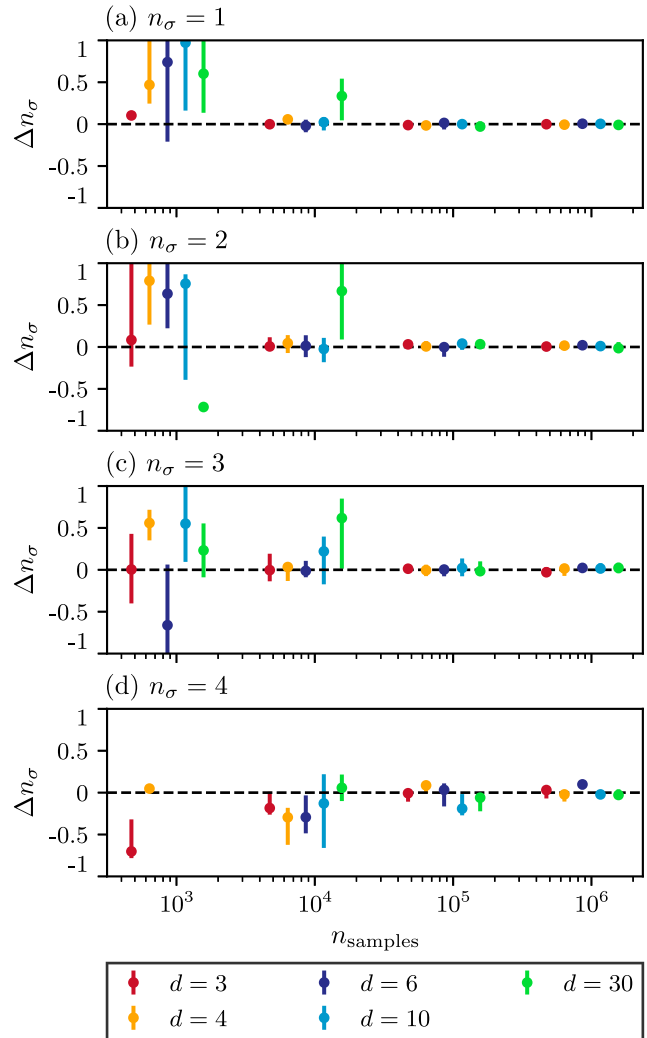


FIG. 5. Similar to Figs. 3 and 4 for the NF estimate.

updates, while the total number of samples impacts the diversity of training batches.

Overall, we observe that having more training samples, even at fixed number of updates and thus computing time, drastically improves the accuracy. Specifically, both the average error and the variance decrease with increasing sample size, down to a variance of about 0.1 on n_σ in our tests. Note that this residual variance includes errors from the NF models as well as the uncertainty from the Monte Carlo estimate (required to be smaller than 0.05). As expected, higher levels of tensions require more training samples, although with a relatively mild dependence on dimensionality, confirming the NF method to be fairly robust to dimensionality. Overall, the NF method, when provided with sufficiently many samples, performs extremely well, with biases below 0.2σ , even for a 4σ tension. However, we also find that, when using only 1000 training samples, MAF models tend to overfit (as can be observed by comparing the training and validation losses), even though the same models, trained on the same number of batches, perform well on larger training samples. In these cases, the method often is unable to produce a shift probability using the Monte-Carlo sampling technique, except in very low dimension. This occurs when all samples drawn from the approximate posterior q fall within the $q(\mathbf{0})$ isohypersurface, which is itself poorly determined due to overfitting. In the next section, we provide some practical tools to monitor training and mitigate for under- and overfitting.

B. Strong non-Gaussianity

We now consider a first example where one of the two distributions has a strong, banana-shaped, non-Gaussian degeneracy. We show in Fig. 6 the two distributions that we use and notice that the strong non-Gaussianity would not be clearly visible in one-dimensional projections. The parameter difference distribution, shown in Fig. 6, inherits the strong non-Gaussianity that causes the input tension not to appear appreciably in one-dimensional marginalized posteriors.

In this example the first distribution is a multivariate Gaussian with zero mean and diagonal covariance equal to $\sigma^2 = 0.01$ while the second distribution is given by $P_2 \sim \mathcal{N}(\sqrt{\theta_1^2 + 20(2\theta_1^2 - \theta_2 - 1/2)^2}, 1/4)$, up to an irrelevant normalization constant. The prior is flat for both parameters in the interval $[-1, 1]$.

The exact input tension for this example is 2.77σ as computed on a very fine grid that makes the error on this determination irrelevant. The Gaussian parameter shift Q_{DM} estimator applied to the two datasets reports a tension of 0.12σ which is largely underestimated. This happens because the long tails of the second distribution move the Gaussian approximation toward the first distribution. From the two single distributions we can compute the joint distribution as the product of the two. We can then use this joint distribution to compute the significance of update

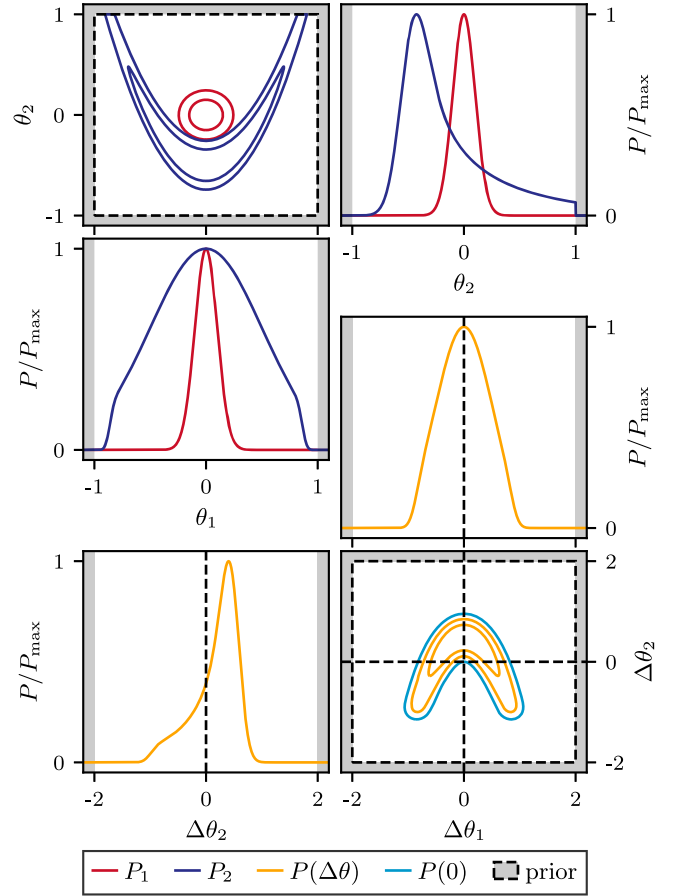


FIG. 6. The relevant posterior distributions of the first benchmark example in Sec. V B. The upper triangle plot shows the posteriors for the two single datasets while the lower triangle shows the posterior of parameter differences. We highlight both the prior boundary and the isocontour that is compatible with zero parameter shift.

parameter shift, Q_{UDM} , applied to the most Gaussian of the two distributions and the joint distribution. This leads to an estimate of 2.81σ tension which is remarkably accurate. The reason for this is that the effect of the non-Gaussian tails is mitigated in the joint distribution.

In the upper panel of Fig. 7 we show the scaling of KDE estimates with number of samples. As we can see the MISE/AMISE estimates increasingly agree as the number of samples increases and show a convergent behavior. While the power law scaling of the MISE estimate follows the binomial one, $\Delta n_\sigma \propto n^{-1/2}$, it does suffer from a factor ten penalty. As we can clearly see in Fig. 7 these two estimators always have a systematic bias that is larger than their variance. Hence the best computation strategy is to use the largest sample size available.

On the other hand the SP estimate is extremely precise in the regime where the number of samples is very limited and samples are sparse. In this regime the accuracy of the SP estimator scales like the binomial but with better overall

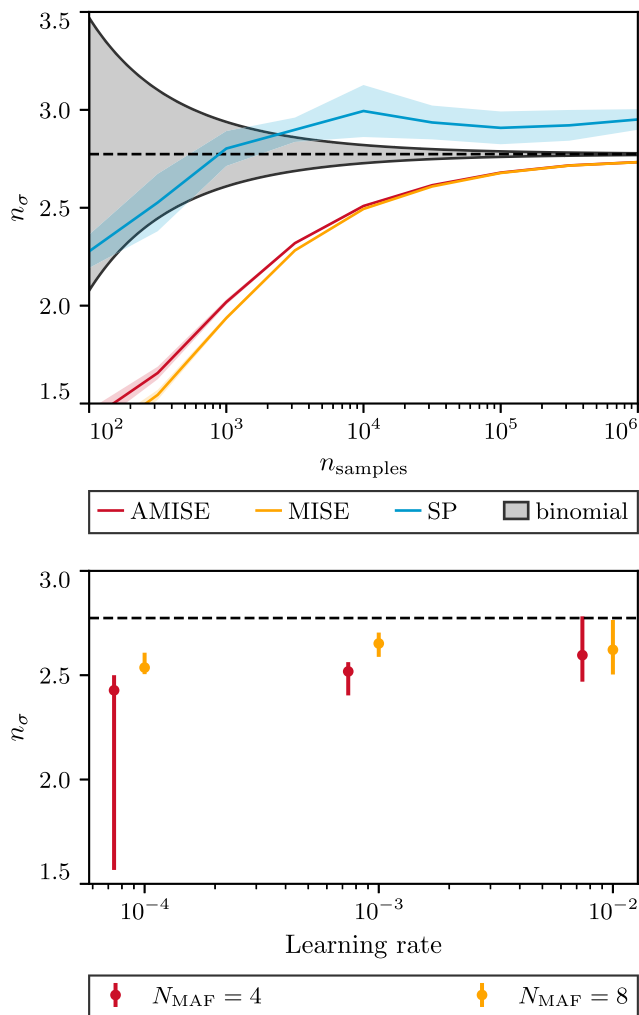


FIG. 7. Determination of the tension in the non-Gaussian example in Sec. VB as a function of number of samples for KDE estimates in the upper panel and as a function of learning rate for NF estimates in the lower panel. The exact result is a 2.77σ tension as shown by the black dashed line. In the upper panel the continuous black lines shows the limiting case of binomial scaling of errors.

performances. Once the sample size exceeds 10^3 and the sparsity of samples decreases, the estimator shows no further improvement. In this regime, we notice that the variance of the SP estimator is of the same order of magnitude as its bias. This suggests that, when a large sample is available, one can increase the precision of the estimator by computing it on smaller batches and averaging. The lack of asymptotic improvement at increasing sample size, for the SP estimator, is due to the absence of a tuning parameter, the smoothing scale, that in the MISE/AMISE case changes as sample size increases.

In the lower panel of Fig. 7 we show the error of NF estimates as a function of learning rate. Note that here, we maintain the value of the learning rate fixed, which determines the sensitivity to local optima, while for realistic

cases, we recommend decreasing it progressively during training (see Sec. VIB). We observe that if the learning rate is set too low, the NF model tends to underfit the banana-shaped distribution, leading to an underestimation of the tension. However, setting the learning rate at 10^{-2} , the maximum value we used in our test, increases the variance in the estimate, as can be seen by the wider boxes in Fig. 7. We also note that in this example, the more flexible models with eight stacked MADEs generally show less variance than models with only four. A possible reason is that MADEs, which implement autoregressive operations, are sensitive to the ordering of variables, which is why we randomly permute them between MADEs. Therefore, since the difference distribution shows strong non-Gaussianity in only one of the two dimensions, models with more MADEs are likely to operate on the most efficient orderings of variables, at the price of some redundancy.

Finally, we illustrate training diagnostics in Fig. 8. Before training, the sample is randomly split into a training sample and a smaller validation sample, which is not used for training, to watch for overfitting. The upper panel shows the evolution, as a function of training epochs, of the training and validation losses (given by the average log-probability of the transformed samples, see Sec. IVA) in blue as well as the evolution of the Kolmogorov-Smirnov statistics, D_n , in yellow. This quantity is testing that the squared norms of samples in the Gaussianized space follows a χ_2^2 distribution. This Gaussianization can be visualized in the middle panel of Fig. 8 where we compare the distribution of the squared norms of samples in the Gaussian approximation space (blue histogram) and in the fully Gaussianized space (yellow histogram) to the χ_2^2 distribution. Finally, we illustrate in the lower panel of Fig. 8 how the model evolves during training and approaches the true distribution (black), by showing the learned distribution at the end of the twenty first epochs (red to yellow).

C. Informative prior

We now consider an example in which the prior is informative. We show in Fig. 9 the two distributions that we use. This example has the highest tension we consider, yet that can hardly be guessed from one-dimensional marginalized posteriors. In this example the first distribution is a multivariate Gaussian with mean $(1/2, -1/2)$ and diagonal covariance equal to $\sigma^2 = 0.01$, while the second distribution is given by $P_2 \sim \mathcal{N}(x - y^3, 1/2)$. The second distribution has a direction which is fully degenerate and constrained by the prior on the unity square. Figure 9 also clearly shows that, while the prior boundary introduces a discontinuity in the two posterior distributions it does not do so for the parameter difference distribution.

The exact input tension for this example is 4.03σ computed as in the previous example. The Gaussian estimator applied to the two datasets reports a tension of

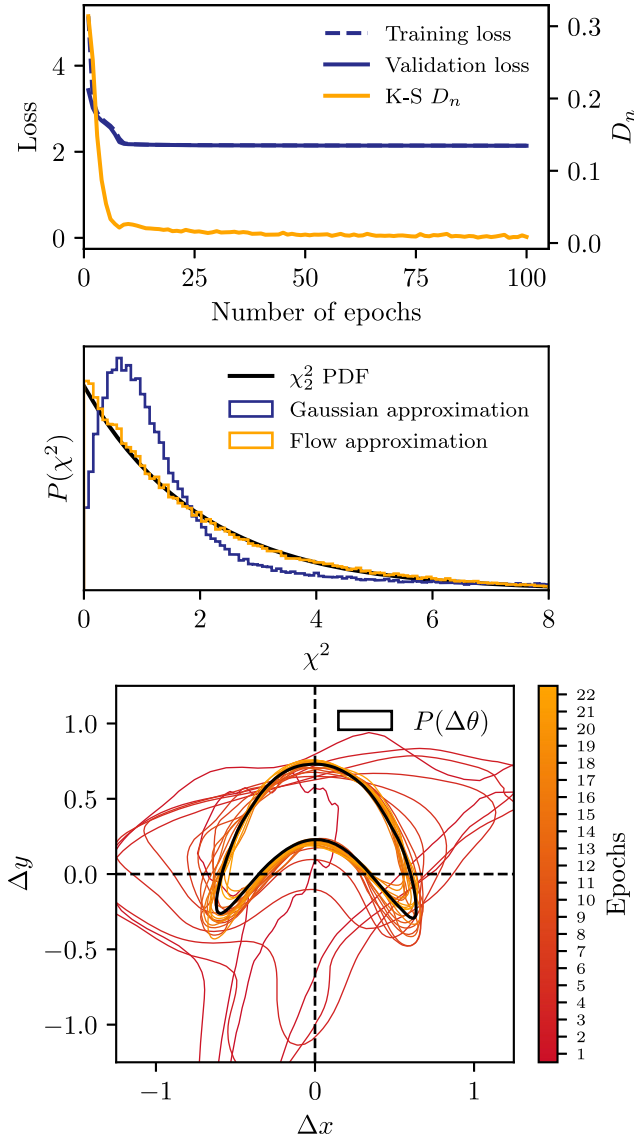


FIG. 8. Illustration of training diagnostics for the NF method. The upper panel shows the evolution of training (solid) and validation (dashed) losses, in blue, as a function of the number of epochs (the two almost completely overlap). At the end of every epoch, we compute the squared norms of the transformed samples and perform a Kolmogorov-Smirnov (K-S) test, for which the null hypothesis is that these norms are χ_d^2 -distributed (here $d = 2$). We show the evolution of the K-S statistic D_n , in yellow, in the upper panel on a different y-axis. In the middle panel, we show the distributions of squared norms of transformed samples, first through the Gaussian approximation (blue) and through the trained MAF (yellow). In practice, we draw these plots at the end of every epoch, allowing the user to monitor convergence. The lower panel illustrates the learning the difference distribution, by showing the 95% contours of the approximate distribution at the end of the first twenty epochs of training (colored contours). The true distribution is shown in black.

3.93σ while the Gaussian update estimator reports 3.89σ . Both Gaussian estimates agree fairly well with the exact result mostly because the tension in the correlated direction

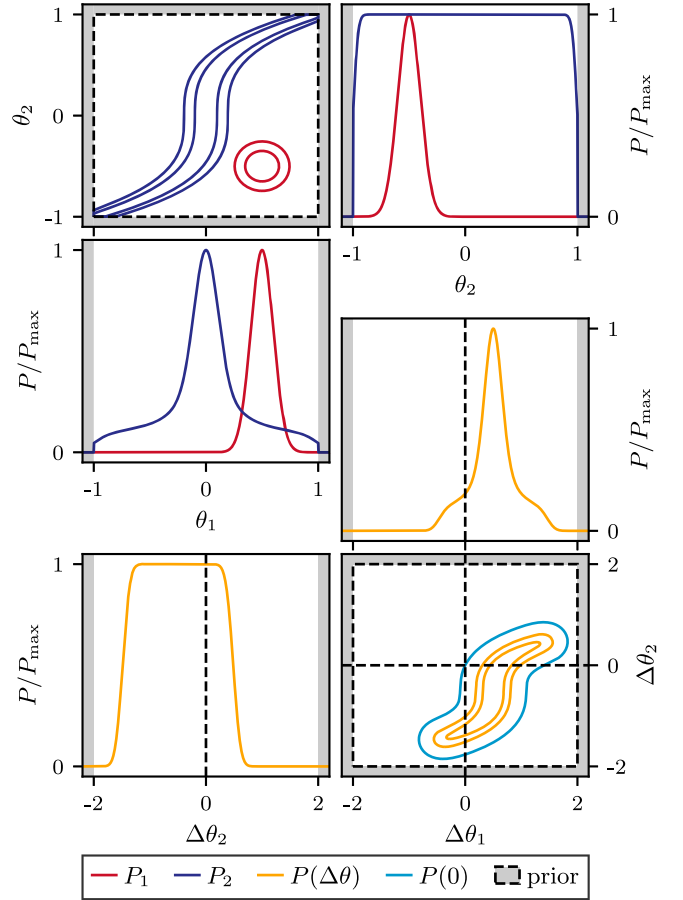


FIG. 9. The relevant posterior distributions of the second benchmark example in Sec. V C. The upper triangle plot shows the posteriors for the two single datasets while the lower triangle shows the posterior of parameter differences. We highlight both the prior boundary and the iso-contour that is compatible with zero parameter shift.

is so high that details of the degeneracy or the prior hardly change the end result.

In Fig. 10 we show the absolute error scaling of KDE estimates with number of samples. As we can see a meaningful determination of the result requires significantly more samples with respect to the previous examples, compatibly with the input tension which is sensibly larger.

The number of samples is so large that the MISE and AMISE estimates entirely agree. In this example the tension is large so the SP estimator is expected to have good performances as can be clearly seen in Fig. 10. In particular samples are so sparse around the zero isocontour, due to the high tension, that the SP estimator has similar performances to the MISE one with an order of magnitude less samples.

NF estimates show good performance, somewhat independent of the learning rate. However, NFs with eight stacked MADEs generally show lower errors, especially at lower learning rate, because of the more complex

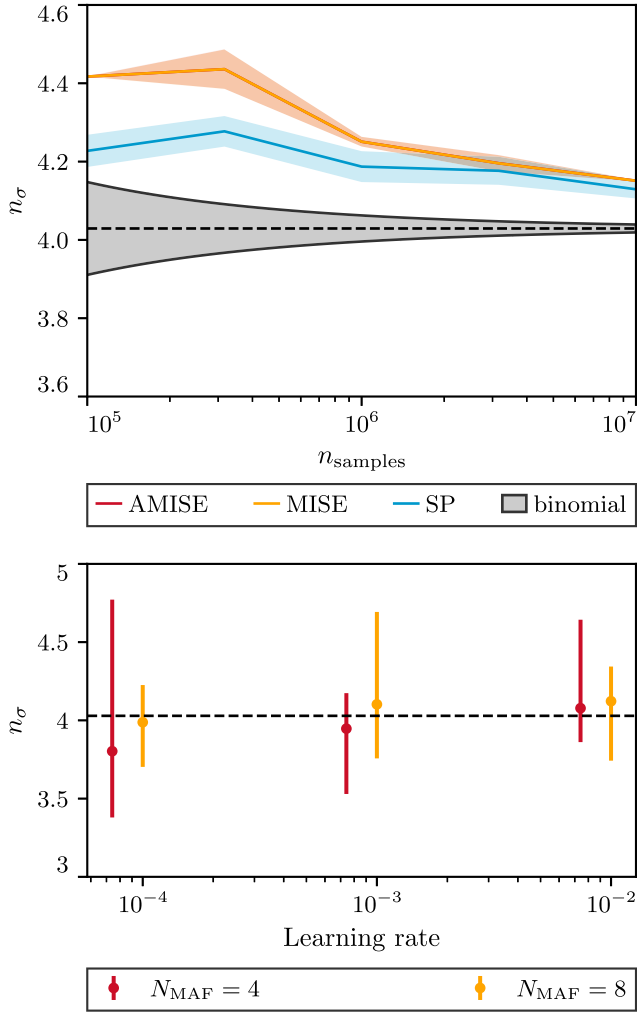


FIG. 10. Determination of the tension in the informative prior example in Sec. V C as a function of number of samples for KDE estimates in the upper panel and as a function of learning rate for NF estimates in the lower panel. The exact result is a 4.03σ tension as shown by the black dashed line. In the upper panel the continuous black lines shows the limiting case of binomial scaling of errors.

morphology of the difference distribution. The best case is found for MAFs made of eight MADEs and the smallest learning rate, in contrast with the previous case, prone to overfitting. In this case, we find errors of order 0.2 with negligible bias (over ten runs). We note that the spread is larger for this example than the previous one. Indeed, we observed more important fluctuations of the predicted isocontour of zero shift over the ten runs, consistent with the fact that training samples are sparser in this case where the true tension is higher.

D. Multimodal posterior

We last consider an example that entails a multimodal posterior. We show in Fig. 11 the two distributions that we use. The first distribution is a multivariate Gaussian with

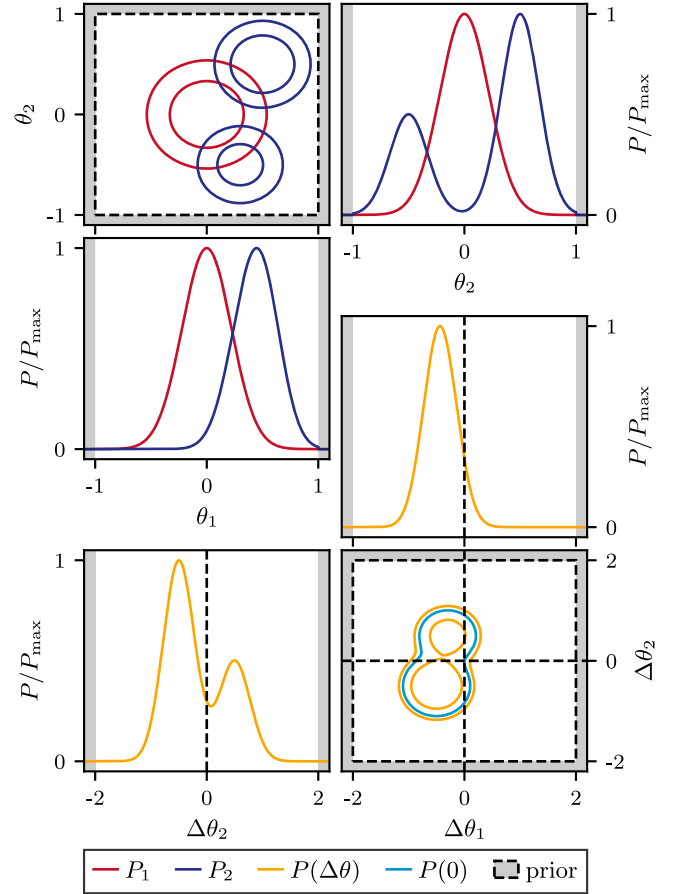


FIG. 11. The relevant posterior distributions of the third benchmark example in Sec. V D. The upper triangle plot shows the posteriors for the two single datasets while the lower triangle shows the posterior of parameter differences. We highlight both the prior boundary and the isocontour that is compatible with zero parameter shift.

zero mean and diagonal covariance equal to $\sigma^2 = 0.05$ while the second distribution is a mixture of two Gaussian distributions, with weights (2, 1), means (0.5, 0.5) and (0.3, -0.5) and diagonal covariances both equal to $\sigma^2 = 0.03$.

The exact input tension for this example is 1.66σ computed as in the previous examples on a grid. The Gaussian estimator applied to the two datasets reports a tension of 1.07σ while the Gaussian update estimator gets closer to the correct result and reports a tension of 1.40σ .

Figure 12 shows the behavior of the KDE estimators that mostly follows previous examples. The large number of samples and the relatively low tension result in fairly poor performances of the SP estimator while the MISE/AMISE estimators show an absolute error decay that is slower than the binomial power law.

Multimodal distributions can be challenging to model with normalizing flows, especially with simple autoregressive operations. However, we observe good performance of




FIG. 12. Determination of the tension in the multimodal example in Sec. VD as a function of number of samples for KDE estimates in the upper panel and as a function of learning rate for NF estimates in the lower panel. The exact result is a 1.66σ tension as shown by the black dashed line. In the upper panel the continuous black lines shows the limiting case of binomial scaling of errors.

our NF method in this example, as shown by the small errors in Fig. 12, albeit with slightly underestimated results. We find that using larger learning rates provides more accurate results, similar to the first example, and models with eight MADEs show lower variance.

VI. APPLICATION TO DES Y1 AND PLANCK

In the previous section we have thoroughly investigated the behavior of the two techniques to compute the significance of parameter differences with benchmark examples where the answer is known.

In this section we apply these estimators to real data where the correct answer is not known, showing a use case with realistic challenges. We consider two datasets: the

clustering and lensing of galaxies from the Dark Energy Survey (DES) year 1 dataset [27]; the Planck 2018 measurements of CMB temperature and polarization at small (Planck 18 TTTEEE) and large angular scales (lowl + lowE) [1,2].

With these two datasets we investigate two cases. In Sec. VIA, we quantify the internal consistency of DES cosmic shear measurements with the DES measurements of the galaxy-galaxy lensing and galaxy clustering two-point functions. This case shows appreciable data correlation. In Sec. VIB, we compute the tension between DES Y1 and Planck 2018, as an example that does not have appreciable correlation between the two datasets.

A. DES Y1 cosmic shear vs galaxy-galaxy lensing and clustering

In this section, we investigate the agreement between cosmic shear measurements from DES Y1 with those of galaxy-galaxy lensing and galaxy clustering.

These datasets are strongly correlated, which means that we must apply the parameter duplication scheme detailed in Sec. II to sample the parameter difference posterior. Here, we duplicate the six cosmological parameters used in the DES Y1 analysis, including the sum of neutrino masses, as well as the two parameters describing intrinsic alignments (amplitude and redshift dependence). We do not duplicate galaxy clustering biases as they only impact the galaxy-galaxy lensing and clustering measurements, and we opt not to duplicate other nuisance parameters that are strongly prior dominated and would considerably increase the computing time needed to sample the parameter difference posterior. To sample the parameter difference distribution, we slightly modify the DES Y1 analysis pipeline within cosmosis [46] to compute the theoretical predictions for the two subsets of the data vector with their respective copies of the parameters, and then combine them in the full likelihood. We sample the posterior with Multinest [47] and show 68% and 95% contours in Fig. 14 in red. We find that the maximum correlation coefficient between all parameters is 58% (see Appendix A of [19]), meaning that it is indeed crucial to account for the correlation between datasets. We note that this analysis in parameter space is complementary to the extensive analysis performed in data space in [25].

We start by computing Gaussian estimators and find that standard parameter difference yields a tension of $n_\sigma = 0.08$ while its update form gives $n_\sigma = 0.5$. In this case we apply the Gaussian estimators to the duplicate parameters to account for the data correlation, as discussed in [19].

In this example the number of samples in the parameter difference chain is limited by the number of samples in the nested sampling chain, and it is sensibly lower than what can be obtained in case of uncorrelated samples. For this reason we vary the sample pool for the KDE estimation from 10^2 samples to all samples in the chain. The number

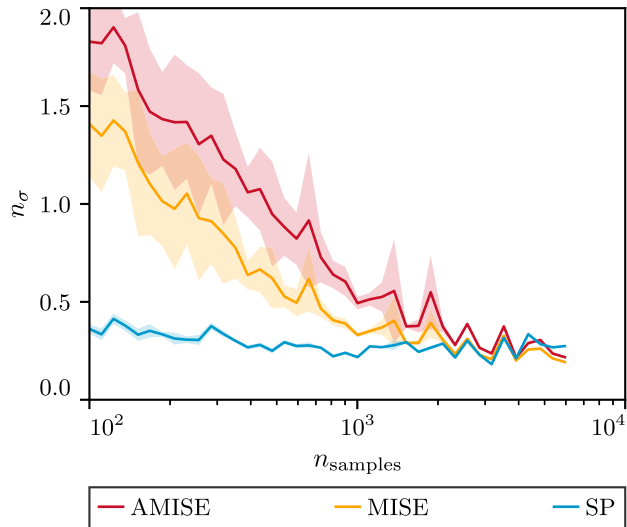


FIG. 13. Internal tension between DES Y1 lensing and galaxy clustering, including their cross-correlation, as in Sec. VI A, as a function of number of samples, according to KDE estimates. Different lines show different algorithms, as shown in legend.

of repetitions of the calculation is then limited by the random nonoverlapping samples that we can draw. As we can see in Fig. 13, similarly to the previous case, for very small sample sizes, all KDE estimates are largely spread and converge toward $n_\sigma = 0.4$ as sample size increases. In this case we notice that the bias of the SP estimator is very low even at small sample size.

For NF, we also employ the default MAF architecture with 16 MADEs, for a total of 11008 trainable weights. With 10 runs, we find $n_\sigma = 0.23 \pm 0.02$, with a small variance that can be attributed to the fact that the difference distribution is very close to Gaussian in this case, as can be seen on Fig. 14. This result is consistent with the results obtained with KDE using the full sample. This test shows good agreement between the subsets of the DES Y1 data used here, consistent with [25]. Using a comparison in data space (asymmetric in the two datasets), they found a calibrated p -value of 0.396 for the “cosmic shear vs galaxy-galaxy lensing and clustering” consistency test, which may be compared to our probability of no shift of $\Delta = 0.29$.

B. DES Y1 3x2pt vs Planck

For this example, we start by sampling the posterior for each dataset, varying the six parameters of the Λ CDM model with the noninformative priors used in the Planck 2018 analysis [2]. We also include all the recommended parameters and priors describing systematic effects in the different datasets we consider. In this case, the sum of neutrino masses is fixed to the minimal value of 0.06 eV (e.g., [48]).

We start by computing the Gaussian tension between the two datasets we consider. As we have seen in the previous

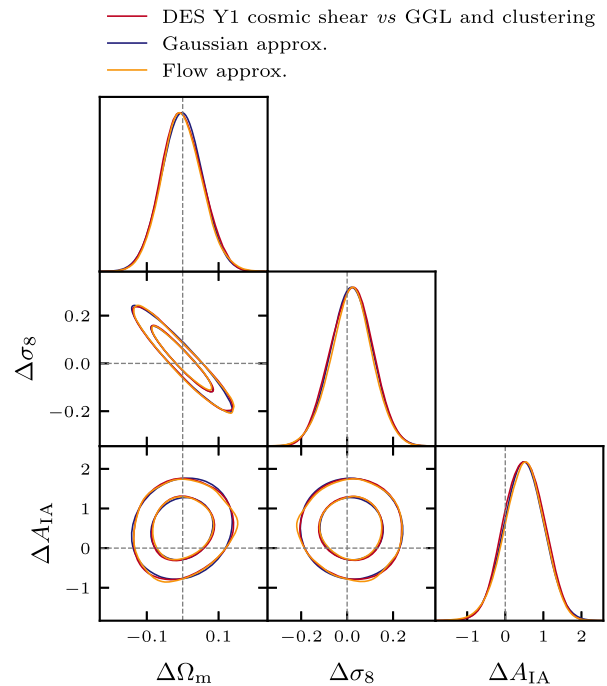


FIG. 14. Marginalized parameter difference distribution between DES Y1 cosmic shear and the combination of galaxy-galaxy lensing (GGL) and clustering, including their cross-correlation. We only show the subspace of the most relevant parameters.

section, the standard Gaussian formula generally misestimates tensions and in this case reports a tension of $n_\sigma = 0.5$. The update Gaussian parameter shift on the other hand, in presence of at least one Gaussian dataset




FIG. 15. Tension between DES and Planck as in Sec. VI B as a function of number of samples, as estimated through the KDE algorithm. Different lines show different algorithms, as shown in legend.

(Planck), is expected to be precise. In this case the update difference in mean (UDM) estimator reports $n_\sigma = 3.0$, which we can use to gauge the order of magnitude of the tension. As such this allows us to orientate algorithmic choices like the number of samples that we would need. The UDM estimator also reports that only two of the six parameters we are considering are significantly contributing to the cosmological parameter constraints and to this tension determination while the other are prior dominated. The data constrained parameters are

combinations of σ_8 and Ω_m that are given by $\sigma_8\Omega_m^{0.7}$ and Ω_m/σ_8 . Even though prior dominated parameter combinations should not contribute to the end result, as we have shown in Sec. II A, we still include them in this calculation.

From the binomial distribution, we would expect to need at least 5×10^3 samples to quantify a 3 sigma tension with an accuracy of $\Delta n_\sigma = 0.1$. We then increase the number of samples from 10^3 to 10^6 and compute the KDE estimator. Figure 15 shows the behavior of the KDE estimator at

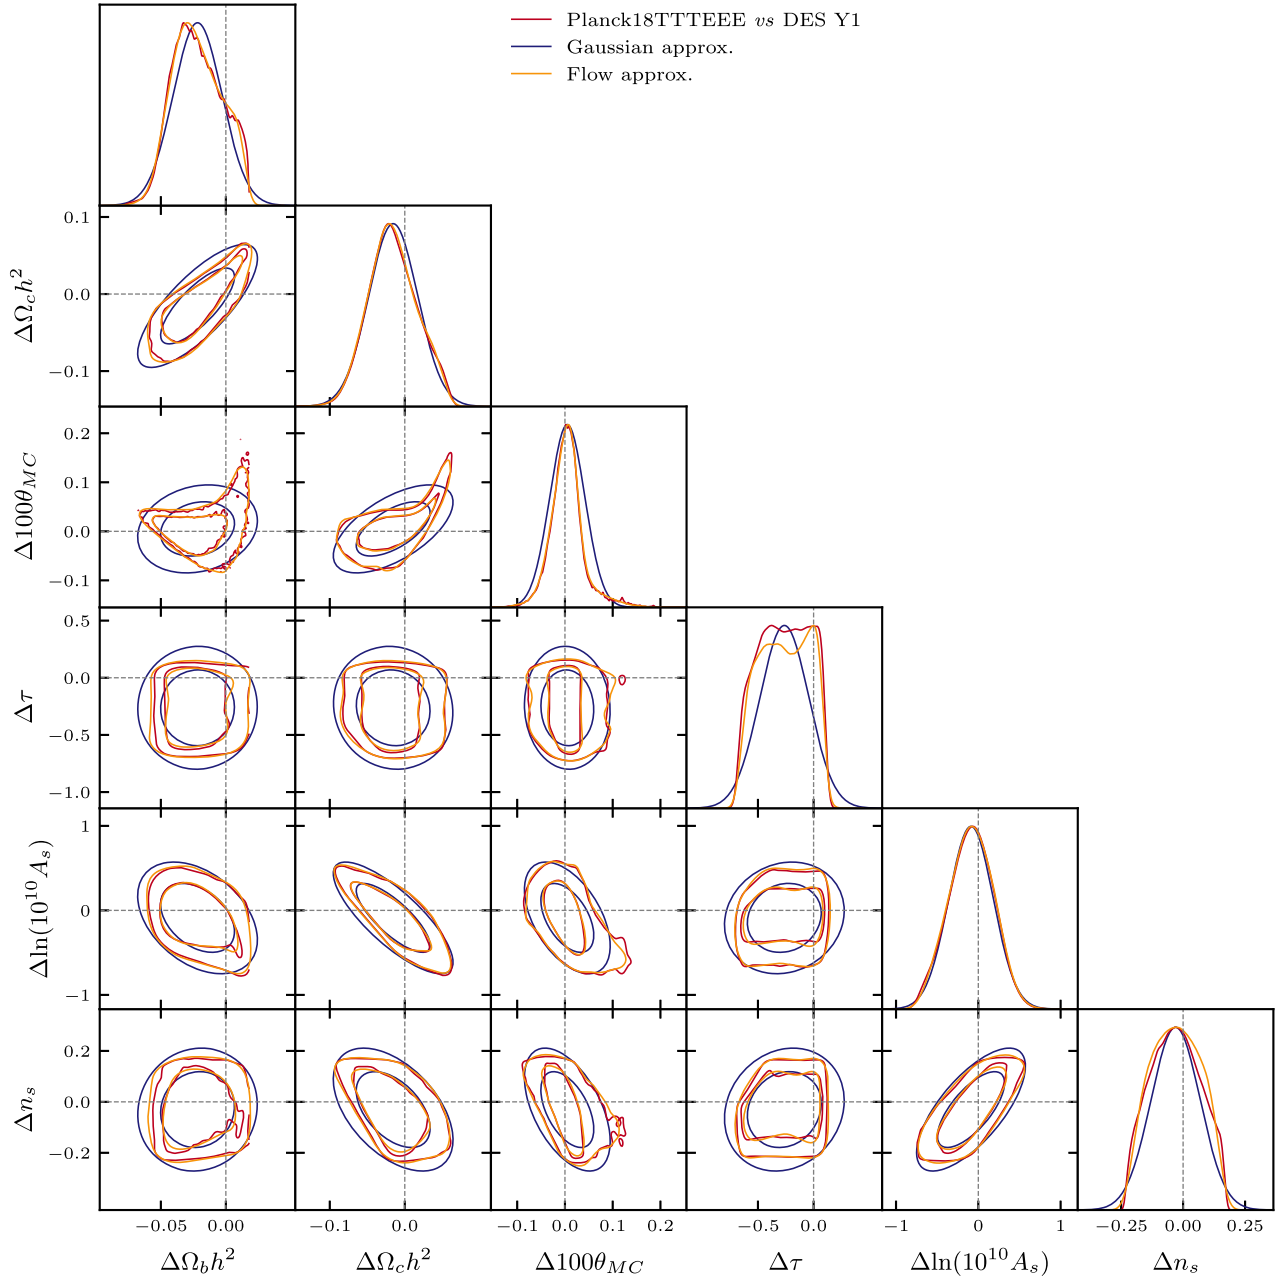


FIG. 16. Parameter difference distribution between DES Y1 and Planck 18 TTTEEE. The true distribution is shown in red, its Gaussian approximation in blue and the normalizing flow approximation in yellow. The NF approximation captures non-Gaussian features of the distribution, including those due to the prior (e.g., on the τ parameter).

increasing number of samples, averaging the result over ten pools of samples randomly drawn. As we can see, at low number of samples, all KDE estimators are largely biased with respect to their end result, and with a significant spread of about 0.5σ . As the number of samples increases the spread between the different estimates reduces and the different estimators converge to about $n_\sigma = 3.1$. This slow convergence is due to the large number of prior constrained directions that are hard for the KDE estimator to resolve accurately. We verified that, using only the parameters that are actually involved in the physical origin of this tension, the result would give a determination of $n_\sigma = 3.1$ with small variation across sample size.

For NF, we employ the default MAF architecture suggested in Sec. IV with 12 MADEs, for a total of 4752 trainable weights. Following guide lines drawn from benchmark examples, we initialize the learning rate at a sufficiently high value of 10^{-2} and decrease it during training using the `ReduceLRonPlateau` callback implemented in `tensorflow.keras.callbacks`. We perform ten runs to assert the level of stochasticity of the estimator and find a average tension of $\langle n_\sigma \rangle = 3.0$ with a standard deviation of 0.12 (with extrema at 2.78 and 3.13). We conclude that the suggested architecture provides sufficiently accurate estimates of tension in this real-data example. Finally, in order to demonstrate the quality of the approximation q to the parameter difference distribution \mathcal{P} , we overlay in Fig. 16 the 68% and 95% contours for the real chain samples (in red) and for a sample of 10^4 points drawn from the approximation q . The agreement between the two sets of contours shows that the NF model can learn non-Gaussian features of the distribution, including those induced by the complex prior (e.g., for the τ parameter).

VII. CONCLUSIONS

As the precision of different cosmological datasets increases, so does the need to reliably check and monitor their consistency. In this paper we have studied a method for doing so that relies on and exploits the full distribution of differences between the posteriors of two experiments. We have characterized the properties of this *parameter difference* estimator showing that it allows to characterize tensions between probes, making no assumption about the underlying distribution of parameters.

We have developed two complementary algorithms for the efficient calculation of the statistical significance of a parameter difference and shown that they generally achieve fully workable performances through a series of toy examples. These show that a tension ranging from 1σ to 4σ in dimension 2 to 30 can be accurately calculated with an error that is below a fraction of a sigma.

The first method we develop relies on kernel density estimates (KDE) of the posterior with fixed and variable smoothing scale. We found that estimates with fixed

smoothing scale work best in low dimensions and show fast convergence. Variable smoothing scale estimators, on the other hand, show slower convergence but largely outperformed fixed smoothing scale estimator in high dimensions with limited number of posterior samples. We have developed an algorithm that solves the problem of the KDE estimator that would naively scale as the number of posterior samples squared, while our *neighbor elimination* algorithm scales as $\mathcal{O}(n \log n)$, outperforming FFT based methods in dimension larger than two.

The second method we develop uses normalizing flows to learn and model the parameter difference posterior distribution from its samples and provides a first use case for machine learning techniques at the problem of estimating tensions between datasets. This technique shows good scaling to large number of samples and high dimension, benefiting from efficient implementations of normalizing flows algorithms in `tensorflow-probability`. In particular, the running time for the normalizing flow estimate is determined by the number of updates during training rather than the number of samples. The more samples are available, the more diverse minibatches are and the faster the model converges.

The two calculation methods we presented have to agree exactly on the statistical significance of a tension. Since they are profoundly different in the way statistical significance is obtained, their spread can be used as a reliable global error estimate, as both results can be obtained in reasonable time on a laptop (from few minutes to few hours). Our benchmark experiments show that the two methods agree with each other and the true tension within 0.5σ in difficult cases, and generally within $0.1-0.2\sigma$ or better, provided enough samples are used.

We have showcased their use with two examples from cosmology. In the first one we have quantified the internal agreement at parameter level of galaxy clustering and lensing from the first year release of DES, finding them in excellent agreement, at the 0.3σ level. We have then quantified the level of agreement between CMB power spectrum measurements from Planck and DES finding them in a significant tension equivalent to 3σ . This is slightly larger than what reported in [26] since we fix neutrino masses to their minimal value, thus increasing the precision of other parameter determinations and hence their tension.

Looking forward, several improvements to the techniques we discuss would be worth pursuing. For KDE methods these would include the development of smoothing scale selectors optimized at the specific isocontour problem at hand rather than general purpose performances across all parameter space. For normalizing flow estimates, an improvement will be to simultaneously train an ensemble of models and average the results, with weights based on the quality of the fit of each model.

ACKNOWLEDGMENTS

We thank Bhuvnesh Jain, Mike Jarvis, Pablo Lemos, Youngsoo Park, Vivian Poulin for helpful comments. M. R. and C.D. are supported in part by NASA ATP Grant No. NNH17ZDA001N, and by funds provided by the Center for Particle Cosmology. C.D. is supported in part

by the US Department of Energy Grant No. de-sc0007901. Computing resources are in part provided by the University of Chicago Research Computing Center through the Kavli Institute for Cosmological Physics at the University of Chicago.

-
- [1] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A6 (2020).
- [2] N. Aghanim *et al.* (Planck Collaboration), *Astron. Astrophys.* **641**, A5 (2020).
- [3] A. G. Riess, S. Casertano, W. Yuan, J. B. Bowers, L. Macri, J. C. Zinn, and D. Scolnic, *Astrophys. J. Lett.* **908**, L6 (2021).
- [4] E. Di Valentino, O. Mena, S. Pan, L. Visinelli, W. Yang, A. Melchiorri, D. F. Mota, A. G. Riess, and J. Silk, *Classical Quantum Gravity* **38**, 153001 (2021).
- [5] T. M. C. Abbott *et al.* (DES Collaboration), [arXiv:2105.13549](https://arxiv.org/abs/2105.13549).
- [6] M. Asgari *et al.* (KiDS Collaboration), *Astron. Astrophys.* **645**, A104 (2021).
- [7] M. P. Hobson, S. L. Bridle, and O. Lahav, *Mon. Not. R. Astron. Soc.* **335**, 377 (2002).
- [8] P. Marshall, N. Rajguru, and A. Slosar, *Phys. Rev. D* **73**, 067302 (2006).
- [9] L. Amendola, V. Marra, and M. Quartin, *Mon. Not. R. Astron. Soc.* **430**, 1867 (2013).
- [10] J. Martin, C. Ringeval, R. Trotta, and V. Vennin, *Phys. Rev. D* **90**, 063501 (2014).
- [11] M. Raveri, *Phys. Rev. D* **93**, 043522 (2016).
- [12] S. Seehars, S. Grandis, A. Amara, and A. Refregier, *Phys. Rev. D* **93**, 103507 (2016).
- [13] S. Joudaki *et al.*, *Mon. Not. R. Astron. Soc.* **465**, 2033 (2017).
- [14] T. Charnock, R. A. Battye, and A. Moss, *Phys. Rev. D* **95**, 123535 (2017).
- [15] W. Lin and M. Ishak, *Phys. Rev. D* **96**, 023532 (2017).
- [16] J. Luis Bernal and J. A. Peacock, *J. Cosmol. Astropart. Phys.* **07** (2018) 002.
- [17] M. Raveri and W. Hu, *Phys. Rev. D* **99**, 043506 (2019).
- [18] S. Adhikari and D. Huterer, *J. Cosmol. Astropart. Phys.* **01** (2019) 036.
- [19] M. Raveri, G. Zacharegkas, and W. Hu, *Phys. Rev. D* **101**, 103527 (2020).
- [20] W. Handley and P. Lemos, *Phys. Rev. D* **100**, 043504 (2019).
- [21] P. Lemos, F. Köhlinger, W. Handley, B. Joachimi, L. Whiteway, and O. Lahav, *Mon. Not. R. Astron. Soc.* **496**, 4647 (2020).
- [22] W. Lin and M. Ishak, *J. Cosmol. Astropart. Phys.* **05** (2021) 009.
- [23] Y. Park and E. Rozo, *Mon. Not. R. Astron. Soc.* **499**, 4638 (2020).
- [24] V. Miranda, P. Rogozenski, and E. Krause, [arXiv:2009.14241](https://arxiv.org/abs/2009.14241).
- [25] C. Doux *et al.* (DES Collaboration), *Mon. Not. R. Astron. Soc.* **503**, 2688 (2021).
- [26] P. Lemos *et al.* (DES Collaboration), [arXiv:2012.09554](https://arxiv.org/abs/2012.09554).
- [27] T. M. C. Abbott *et al.* (DES Collaboration), *Phys. Rev. D* **98**, 043526 (2018).
- [28] M. Wand and M. Jones, *Kernel Smoothing*, Chapman & Hall/CRC Monographs on Statistics & Applied Probability (Taylor & Francis, London, 1994).
- [29] J. Chacón and T. Duong, *Multivariate Kernel Smoothing and Its Applications*, Chapman & Hall/CRC Monographs on Statistics and Applied Probability (CRC Press, Boca Raton, 2018).
- [30] G. R. Terrell and D. W. Scott, *Ann. Statist.* **20**, 1236 (1992).
- [31] M. Rosenblatt, *J. Ann. Math. Stat.* **27**, 832 (1956).
- [32] E. Parzen, *Ann. Math. Stat.* **33**, 1065 (1962).
- [33] L. Bagnato and A. Punzo, *Comput. Stat.* **36**, 1177 (2021).
- [34] G. Papamakarios, T. Pavlakou, and I. Murray, [arXiv:1705.07057](https://arxiv.org/abs/1705.07057).
- [35] D. P. Kingma and P. Dhariwal, [arXiv:1807.03039](https://arxiv.org/abs/1807.03039).
- [36] W. Grathwohl, R. T. Q. Chen, J. Bettencourt, I. Sutskever, and D. Duvenaud, [arXiv:1810.01367](https://arxiv.org/abs/1810.01367).
- [37] I. Kobyzev, S. J. D. Prince, and M. A. Brubaker, *IEEE Trans. Pattern Anal. Mach. Intell.* **1** (2019).
- [38] C. Meng, Y. Song, J. Song, and S. Ermon, [arXiv:2003.01941](https://arxiv.org/abs/2003.01941).
- [39] D. M. Reiman, J. Tamanas, J. X. Prochaska, and D. Ďurovčíková, [arXiv:2006.00615](https://arxiv.org/abs/2006.00615).
- [40] A. Diaz Rivero and C. Dvorkin, *Phys. Rev. D* **102**, 103507 (2020).
- [41] A. Moss, *Mon. Not. R. Astron. Soc.* **496**, 328 (2020).
- [42] N. Jeffrey, J. Alsing, and F. Lanusse, *Mon. Not. R. Astron. Soc.* **501**, 954 (2020).
- [43] Y. Bengio and Y. Lecun, Scaling learning algorithms towards ai, in *Large-Scale Kernel Machines*, edited by L. Bottou, O. Chapelle, D. DeCoste, and J. Weston (MIT Press, Cambridge, MA, 2007).
- [44] M. Germain, K. Gregor, I. Murray, and H. Larochelle, [arXiv:1502.03509](https://arxiv.org/abs/1502.03509).

- [45] J. V. Dillon, I. Langmore, D. Tran, E. Brevdo, S. Vasudevan, D. Moore, B. Patton, A. Alemi, M. Hoffman, and R. A. Saurous, [arXiv:1711.10604](https://arxiv.org/abs/1711.10604).
- [46] J. Zuntz, M. Paterno, E. Jennings, D. Rudd, A. Manzotti, S. Dodelson, S. Bridle, S. Sehrish, and J. Kowalkowski, *Astron. Comput.* **12**, 45 (2015).
- [47] F. Feroz and M. P. Hobson, *Mon. Not. R. Astron. Soc.* **384**, 449 (2008).
- [48] A. J. Long, M. Raveri, W. Hu, and S. Dodelson, *Phys. Rev. D* **97**, 043510 (2018).