

## SAVAGE-DICKEY DENSITY RATIO ESTIMATION WITH NORMALIZING FLOWS FOR BAYESIAN MODEL COMPARISON

KIYAM LIN,<sup>1\*</sup> ALICJA POLANSKA,<sup>1</sup> DAVIDE PIRAS,<sup>2,3</sup> ALESSIO SPURIO MANCINI,<sup>4</sup> AND JASON D. MCEWEN<sup>1</sup>

<sup>1</sup> Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey, RH5 6NT, UK

<sup>2</sup> Centre Universitaire d'Informatique, Université de Genève, 7 route de Drize, 1227 Genève, Switzerland

<sup>3</sup> Département de Physique Théorique, Université de Genève, 24 quai Ernest Ansermet, 1211 Genève 4, Switzerland and

<sup>4</sup> Department of Physics, Royal Holloway, University of London, Egham Hill, Egham, UK

(Dated: 4th June 2025)  
submitted XXX; accepted YYY

### Abstract

A core motivation of science is to evaluate which scientific model best explains observed data. Bayesian model comparison provides a principled statistical approach to comparing scientific models and has found widespread application within cosmology and astrophysics. Calculating the Bayesian evidence is computationally challenging, especially as we continue to explore increasingly more complex models. The Savage-Dickey density ratio (SDDR) provides a method to calculate the Bayes factor (evidence ratio) between two nested models using only posterior samples from the super model. The SDDR requires the calculation of a normalised marginal distribution over the extra parameters of the super model, which has typically been performed using classical density estimators, such as histograms. Classical density estimators, however, can struggle to scale to high-dimensional settings. We introduce a neural SDDR approach using normalizing flows that can scale to settings where the super model contains a large number of extra parameters. We demonstrate the effectiveness of this neural SDDR methodology applied to both toy and realistic cosmological examples. For a field-level inference setting, we show that Bayes factors computed for a Bayesian hierarchical model (BHM) and simulation-based inference (SBI) approach are consistent, providing further validation that SBI extracts as much cosmological information from the field as the BHM approach. The SDDR estimator with normalizing flows is implemented in the open-source `harmonic` Python package. 

### 1. INTRODUCTION

A core aspect of science is to determine which scientific model of reality is the most accurate. Bayesian model comparison provides a principled statistical approach to comparing scientific models that has found widespread application, particularly in astrophysics and cosmology (Trotta 2007). With advances in computational power, it is possible to analyse increasingly complex cosmological and astrophysical models. To describe these more complex and intricate models, additional parameters are often required, including those that are physically interesting as well as those that might be considered nuisance parameters yet must be modelled (e.g. von Wietersheim-Kramsta et al. 2025).

These more complex models that often exist in higher-dimensional parameter spaces, however, can be costly to analyse, with some models being infeasible to explore due to insurmountable computational costs. We can make use of modern machine learning techniques and their underlying technologies to massively accelerate the computational speed of performing Bayesian inference. This can be achieved by combining (i) *emulation of physical models by machine learning techniques*; (ii) *differentiable and probabilistic programming*; (iii) *scalable Markov chain Monte Carlo sampling* (MCMC) that

makes use of gradient information, and (iv) *decoupled and scalable Bayesian model selection*, as advocated in Piras et al. (2024). Gradient-based Markov chain Monte Carlo (MCMC) sampling algorithms that can leverage gradients computed efficiently by automatic differentiation make sampling in higher-dimensional parameter spaces tractable, e.g. the No-U-Turn Sampler (NUTS; Hoffman & Gelman 2014), a highly efficient and adaptive variant of Hamiltonian Monte Carlo (HMC; Duane et al. 1987; Neal 1996).

Bayesian model selection requires the computation of the Bayesian evidence, which is typically performed by nested sampling (Skilling 2006; see Ashton et al. 2022; Buchner 2023 for reviews), where sampling must be performed in a prescribed nested manner for which a variety of algorithms have been developed (Feroz & Hobson 2008; Feroz et al. 2009; Higson et al. 2019; Feroz et al. 2019; Handley et al. 2015a,b; Speagle 2020; Buchner 2021; Cai et al. 2021; McEwen et al. 2023; Lange 2023). Nested sampling is thus not applicable for arbitrary MCMC sampling algorithms. To compute the Bayesian evidence for arbitrary sampling algorithms, such as NUTS, we require methods to compute the Bayesian evidence that are decoupled from the sampling algorithm while being scalable. Methods that require only posterior samples to compute the evidence have become increasingly popular (Dickey 1971; DiCiccio et al.

\* E-mail: kiyam.lin@ucl.ac.uk

1997; Trotta 2007; Heavens et al. 2017; Jia & Seljak 2020; McEwen et al. 2021; Srinivasan et al. 2024; Rinaldi et al. 2024). Amongst these methods, the learned harmonic mean estimator (McEwen et al. 2021; Polanska et al. 2023; Polanska et al. 2024a), available in the `harmonic` software package, has been shown to provide robust and scalable evidence calculations and has been growing in popularity within astrophysics and cosmology (Spurio Mancini et al. 2023; Piras et al. 2024; Polanska et al. 2024b; Spurio Mancini et al. 2024; Stiskalek & Desmond 2024; Du et al. 2025; Carrion et al. 2025; Stiskalek et al. 2025).

Even with such acceleration, the process of performing Bayesian model comparison can still be highly computationally costly as it requires acquiring posterior samples along with their respective likelihood values for every model being evaluated. However, if the models being compared have a relationship where one model is a *nested* model of the larger *super* model, then we can make use of the Savage-Dickey density ratio (SDDR; Dickey 1971; Verdinelli & Wasserman 1995; O’Hagan & Forster 2004) to perform Bayesian model comparison without needing to explicitly acquire samples from the nested model. The term ‘nested’ here refers to a relationship between two models where a super model reduces to a smaller, nested model when some of the super model’s parameters are fixed at particular values, as is often the case in cosmological settings. We refer to the parameters common to both the super and nested model as *common parameters* and the additional parameters of the super model as *extra parameters*. This approximately halves the computational cost of performing Bayesian model comparison as we only need to perform sampling once for the super model. The SDDR has already been applied to cosmology, for instance by Trotta (2007); Marin & Robert (2009); Verde et al. (2013); Leistedt et al. (2014); Salvatelli et al. (2014); Di Valentino et al. (2017, 2020); Kreisch et al. (2020).

Besides requiring nested models, the SDDR also requires there to be sufficient posterior support with high enough sample density around the parameter values where the super model reduces to the nested model to be able to evaluate the marginal posterior accurately. If this requirement is not met, calculating the evidence directly with the learned harmonic mean provides a good alternative.

Traditionally the SDDR has been calculated through the use of classical density estimators, such as histograms or kernel density estimation. For situations where there are only a small number of extra parameters, classical density estimators provide good estimates of the evidence ratio. However, this approach can fail even in settings with just a few extra parameters as, e.g., it becomes computationally difficult to normalise the histogram of the probability distribution of the extra parameters by numerical integration due to the curse of dimensionality. Normalizing flows (Papamakarios et al. 2021), however, are normalized by construction and thus are a natural neural alternative to classical density estimators when considering higher-dimensional problems. Thus, in this work we make use of normalizing flows to estimate the normalized probability distribution of the extra parameters. As such, we are able to perform SDDR calculations with a large number of extra parameters.

The remainder of this article is organised as follows. In Section 2 we provide a brief overview of SDDR and our methodology with normalizing flows. In Section 3 we demonstrate and validate the use of our SDDR methodology in both toy and realistic cosmological examples. We conclude in Section 4.

## 2. SAVAGE-DICKEY DENSITY RATIO WITH NORMALIZING FLOWS

In this section we provide a brief overview of the SDDR and its connection to Bayesian model comparison. We introduce our SDDR methodology with normalizing flows, discuss how uncertainties can be estimated and present an open-source software implementation.

### 2.1. Bayesian model comparison

The Bayesian evidence is a measure of the probability of obtaining the observed data given the model under consideration, defined as

$$z_i = p(\mathbf{d}|M_i) = \int p(\mathbf{d}|\theta, M_i)p(\theta|M_i)d\theta, \quad (1)$$

where  $z_i$  denotes the Bayesian evidence of model  $M_i$  for data,  $\mathbf{d}$ . The evidence ratio, or Bayes factor, between any two models,  $M_1$  and  $M_2$ , can be used to assess which model is favored by the data, expressed as

$$\frac{z_1}{z_2} = \frac{p(\mathbf{d}|M_1)}{p(\mathbf{d}|M_2)}. \quad (2)$$

In the case where we have equal prior probabilities for each model, as is often the case in practice, the Bayes factor reduces to the posterior model odds following Bayes theorem:

$$\frac{z_1}{z_2} = \frac{p(M_1|\mathbf{d})p(M_2)}{p(M_2|\mathbf{d})p(M_1)} = \frac{p(M_1|\mathbf{d})}{p(M_2|\mathbf{d})}. \quad (3)$$

For models with equal prior model probabilities, the Bayes factor can thus be used to inform us of which model is more probable given the observed data.

### 2.2. Savage-Dickey density ratio

The SDDR (Dickey 1971; Verdinelli & Wasserman 1995; O’Hagan & Forster 2004) is a method of calculating the Bayes factor, i.e. ratio of evidences between two models, when one model is a nested model of the other. It is useful when one does not have direct access to the nested model or it is computationally infeasible to explore all possible nested models. The SDDR gives the Bayes factor in terms of the ratio of the normalized marginal posterior to the prior density of the super model in the extra parameter(s). This ratio is evaluated at the point where the super model reduces to the nested model. Model  $M_1$  can be considered a nested model of a larger super model  $M_2$  if:

1. the parameters of  $M_2$  can be expressed as a set of parameters  $\{\boldsymbol{\theta}, \boldsymbol{\eta}\}$ , where  $\boldsymbol{\theta}$  are the common parameters of  $M_1$  and  $M_2$ , whilst  $\boldsymbol{\eta}$  are the extra parameters required to extend  $M_1$  to  $M_2$ ;
2. for some particular value of  $\boldsymbol{\eta} = \boldsymbol{\eta}_1$ , the likelihood of the super model reduces to the likelihood of the nested model;

3. the prior distributions of  $\boldsymbol{\theta}$  are identical for  $M_1$  and  $M_2$ ;
4. the prior distributions of  $\boldsymbol{\eta}$  and  $\boldsymbol{\theta}$  are separable, i.e.  $p(\boldsymbol{\theta}, \boldsymbol{\eta}) = p(\boldsymbol{\theta})p(\boldsymbol{\eta})$ .

The Bayes factor given data and two models  $M_1$  and  $M_2$  is given by

$$\frac{z_1}{z_2} = \frac{p(\mathbf{d}|M_1)}{p(\mathbf{d}|M_2)} = \frac{\int p(\mathbf{d}|\boldsymbol{\theta}', M_1) p(\boldsymbol{\theta}'|M_1) d\boldsymbol{\theta}'}{\int p(\mathbf{d}|\boldsymbol{\theta}'', M_2) p(\boldsymbol{\theta}''|M_2) d\boldsymbol{\theta}'' d\boldsymbol{\eta}''}. \quad (4)$$

As our models are nested, we have  $p(\mathbf{d}|\boldsymbol{\theta}, M_1) = p(\mathbf{d}|\boldsymbol{\theta}, \boldsymbol{\eta}_1, M_2)$ . It is then straightforward to show (Dickey 1971; Verdinelli & Wasserman 1995; O'Hagan & Forster 2004; Trotta 2007; Verde et al. 2013) that the Bayes factor reduces to

$$\frac{z_1}{z_2} = \frac{\int p(\boldsymbol{\theta}', \boldsymbol{\eta}_1|\mathbf{d}, M_2) d\boldsymbol{\theta}'}{p(\boldsymbol{\eta}_1|M_2)}. \quad (5)$$

We can see that the integration of the posterior distribution depicted in the numerator is simply marginalization over the common parameters  $\boldsymbol{\theta}$ , yielding the marginalized posterior in  $\boldsymbol{\eta}_1$ . Hence, the Bayes factor between two nested models is given by

$$\frac{z_1}{z_2} = \frac{p(\boldsymbol{\eta}_1|\mathbf{d}, M_2)}{p(\boldsymbol{\eta}_1|M_2)}. \quad (6)$$

The numerator in Eq. (6) is the marginalised posterior of the extra parameter(s) for the super model, whilst the denominator is the prior of the extra parameter(s) for the super model. A more detailed derivation of the SDDR can be found in Appendix A.

It follows from Eq. (6) that to compute the SDDR one only requires the posterior samples from the super model and the prior distribution of the extra parameter(s). This is a significant advantage as it means that we only need to perform sampling once for the super model. We can then use those samples to calculate the Bayes factor with relation to all possible nested models.

As evaluation of the prior is typically trivial, the challenge lies in the evaluation of the marginal posterior at  $\boldsymbol{\eta}_1$ , as it must be normalized. This requires numerical integration in potentially high-dimensions if there are many extra parameters. In this work, we consider two approaches.

In the first approach we make use of classical density estimators by way of a normalized histogram (as considered previously, e.g. Verde et al. 2013). We also introduce a second approach where we make use of a neural density estimator by way of normalizing flows to learn the marginal posterior distribution. Importantly, it is only with normalizing flows that we are able to calculate the SDDR in high-dimensional marginal posterior spaces.

Care however needs to be taken when using the SDDR as it is only valid when the posterior has sufficient sample density around  $\boldsymbol{\eta}_1$ . If this requirement is not met, calculating the evidence directly with the learned harmonic mean is preferred.

### 2.3. SDDR with classical density estimators

To compute the SDDR with classical density estimators we consider a simple normalized histogram. Following normalization, we evaluate the marginal posterior probability at  $\boldsymbol{\eta}_1$ . Beyond even a few dimensions normalizing a histogram becomes computationally challenging. This is due to the need to numerically integrate over the high-dimensional parameter space alongside the curse of dimensionality. Consequently, this method does not scale well to high-dimensional problems.

### 2.4. SDDR with normalizing flows

Normalizing flows have the advantage of being normalized by construction, and thus are a natural neural density estimator analogue to classical density estimators that scale to higher-dimensional settings. A normalizing flow, in essence, learns the transformation from a known tractable base distribution such as a Gaussian to the target distribution. For a comprehensive review of normalizing flows we refer the reader to Papamakarios et al. (2021).

More specifically, a vector  $\boldsymbol{\theta}$  from an unknown distribution  $p(\boldsymbol{\theta})$  can be expressed through a transformation  $B$  of a latent vector  $\mathbf{u}$  sampled from the base distribution  $q(\mathbf{u})$ :

$$\boldsymbol{\theta} = B(\mathbf{u}), \quad \mathbf{u} \sim q(\mathbf{u}). \quad (7)$$

If  $B$  is invertible and both  $B$  and its inverse  $B^{-1}$  are differentiable, we can calculate the density of the distribution of  $\boldsymbol{\theta}$  through the change of variables formula by

$$p(\boldsymbol{\theta}) = q(\mathbf{u}) |\det J_B(\mathbf{u})|^{-1}, \quad (8)$$

where  $J_B(\mathbf{u})$  is the Jacobian corresponding to  $B$ . As this transformation is composable, we can apply  $B$  multiple times to obtain approximations of more complex target distributions. A normalizing flow thus consists of a series of such transformations.

In this work we make use of rational quadratic spline flows (Durkan et al. 2019). Rational quadratic spline flows split the input domain of  $\mathbf{u}$  into multiple segments and apply monotonic rational quadratic polynomial transformations to each segment. The polynomial segments are restricted to match at bin boundaries such that the transformed distribution is continuous and differentiable.

The rational quadratic splines consist of learned parameters that are optimised during training as well as alternating affine transformation layers to create the normalizing flow. Network training is performed by minimizing the Kullback-Leibler divergence between samples of the marginal posterior and the target distribution parameterised by the normalizing flow. We train the flow to learn the marginal posterior distribution and evaluate the flow at  $\boldsymbol{\eta}_1$ .

### 2.5. Uncertainty estimation

To obtain error estimates for both the classical and neural SDDR methodologies we perform bootstrapping. We randomly subsample the marginalised posterior samples, with replacement, to create several sets of bootstrapped samples. We then make use of each bootstrapped sample set to calculate an SDDR estimate

either with an independently trained flow or a normalized histogram. We take the mean of these estimates as our SDDR estimate and the standard deviation across estimates as the error estimate.

### 2.6. Code implementation

The implementation of the SDDR using normalizing flows can be found within the `harmonic`<sup>1</sup> software package. We have chosen to implement our methods as part of `harmonic` to provide a single code for Bayesian model comparison that is versatile and can handle all situations. The implementation is also consistent with the interface of `harmonic`. `harmonic` is a professional code that is well-tested and documented and can be installed from PyPi. Given posterior samples, computation of the Bayes factor using our implementation of the SDDR using normalizing flows requires very little computational resources to train and run, typically running within a few minutes on a single CPU. In the snippet in listing 1 we demonstrate the use of SDDR within `harmonic`.

```
import harmonic as hm
import numpy as np

# Define the number of marginal dimensions
nparams = 4
# Load the marginalised posterior samples
marginalised_samples = np.load('samples.npy')
# Define the pre-calculated log prior value
log_prior_η_1 = -2
# The point of nesting
η_1 = np.array([-2.0, 0.0, -0.5, 0.5])
# Create the normalizing flow model
model = hm.model.RQSplineModel(nparams, temperature=1.0)
# Create the sddr object
sddr = hm.sddr.sddr(model, marginalised_samples)
# Calculate the Bayes factor
log_bf, std = sddr.log_bayes_factor(log_prior_η_1, η_1)
```

Listing 1: Example code snippet of how to use the SDDR functionality within `harmonic` to calculate the Bayes factor given  $\eta_1$  and its corresponding log prior value as well as the samples of the extra parameters.

## 3. APPLICATIONS

We calculate the SDDR with both classical density estimators (histogram) and neural density estimators (normalizing flows) for a variety of toy and realistic cosmological examples. Furthermore, we validate the Bayes factors computed against alternative approaches such as nested sampling and the learned harmonic mean (which require sampling both the super and nested models), finding excellent agreement. We demonstrate the effectiveness of neural SDDR both in low-dimensional marginal posterior spaces where we have access to the classical density estimator method, but also in higher-dimensional extra parameter spaces where the classical density estimator method breaks down.

### 3.1. Gaussian toy examples

We construct two toy Gaussian likelihood examples with either 1 or 4 extra parameters in the super model compared to the nested model. In all scenarios we calculate the log Bayes factor (evidence ratio) by the SDDR, making use of both the classical histogram method and

**Table 1:** Gaussian toy model with 1 extra parameter.

Method	Log Bayes factor
Nested sampling	0.905 ± 0.013
Learned harmonic mean	0.877 ± 0.002
SDDR (classical)	0.893 ± 0.021
SDDR (flows)	0.890 ± 0.056

the normalizing flow method. We validate our results with a combination of both the learned harmonic mean estimator using `harmonic` with samples obtained with a MCMC sampler via `emcee`<sup>2</sup> (Foreman-Mackey et al. 2013), and nested sampling with the use of `nautilus`<sup>3</sup> (Lange 2023) to have two independent methods of calculating the Bayesian evidence.

We consider the data model  $\mathbf{d} \sim \mathcal{N}(\boldsymbol{\mu}(\boldsymbol{\theta}), \Sigma)$ , where the mean  $\boldsymbol{\mu}$  is dependent on parameters  $\boldsymbol{\theta}$  but the covariance matrix  $\Sigma$  is not. The corresponding Gaussian likelihood is given by

$$\mathcal{L}(\mathbf{d} | \boldsymbol{\theta}) = \frac{\exp\left(-\frac{1}{2}(\mathbf{d} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T \Sigma^{-1} (\mathbf{d} - \boldsymbol{\mu}(\boldsymbol{\theta}))\right)}{(2\pi)^n |\Sigma|^{1/2}}, \quad (9)$$

where  $n$  is the dimension of the data vector.

For the first Gaussian example we consider a data dimension of  $n = 3$  and a super model with parameter dimension  $d = 4$  where

$$\boldsymbol{\mu}(\boldsymbol{\theta}) = \begin{bmatrix} \theta_1 \\ \theta_2 \\ e^{0.5 \times \theta_3} + \theta_4 \end{bmatrix}. \quad (10)$$

In this first Gaussian model considered, the super model (with 4 parameters) reduces to the nested model with 3 parameters when  $\theta_4 = -2.0$ . We initialise a random covariance,  $\Sigma$ , by first creating a diagonal covariance with elements randomly drawn from a 1D Gaussian with mean of 1.0 and standard deviation of 0.1 ensuring all elements are positive. Off diagonal elements in the upper triangle that are only one index away from the diagonal with matrix index  $i$  are then set with  $\Sigma_{i,i+1} = -1^i \times 0.5 \times \sqrt{\Sigma_{i,i} \times \Sigma_{i+1,i+1}}$ . The upper triangle is then symmetrised with the lower triangle making the covariance matrix symmetrical. A mock data vector is generated by evaluating the forward model at  $\boldsymbol{\theta} = [0.0, -0.5, 0.5, -2.0]$ , making the nested model the ground truth. For inference we consider Gaussian priors on all parameters with mean of 0.0 and diagonal covariance of 2.0.

We find that with this simple example both the histogram and normalizing flow methods provide accurate estimates of the Bayes factor that are in close agreement with those computed by nested sampling and the learned harmonic mean, as shown in Table 1. In this case, there is a weak preference for the nested model which is indeed the model used to generate the ground truth.

For the second toy Gaussian example, we consider the same setup as before but we have a super model that has 8 parameters whilst the nested model only has 4 parameters. The likelihood is again of the form shown in Eq. 9, this time with  $n = 4$ , and the super model is

<sup>1</sup> <https://github.com/astro-informatics/harmonic>

<sup>2</sup> <https://github.com/dfm/emcee>

<sup>3</sup> <https://github.com/johannesulf/nautilus>

**Table 2:** Gaussian toy model with 4 extra parameters. Note that there is no entry for the histogram method as already in this simple example it was computationally challenging to normalize the marginal posterior histogram and thus calculate the SDDR.

Method	Log Bayes factor
Nested sampling	$2.227 \pm 0.013$
Learned harmonic mean	$2.248 \pm 0.007$
SDDR (classical)	—
SDDR (flows)	$2.374 \pm 0.088$

defined with

$$\boldsymbol{\mu}(\boldsymbol{\theta}) = \begin{bmatrix} \theta_1 + \operatorname{arcsinh}(\theta_5) \\ \arctan(\theta_2) + \theta_6^2 \\ e^{0.5 \times \theta_2} + \theta_7 \\ e^{0.5 \times \theta_4} + \arctan(\theta_8) \end{bmatrix}, \quad (11)$$

which reduces to the nested model by setting  $\theta_5 = -2.0$ ,  $\theta_6 = 0.2$ ,  $\theta_7 = -0.2$  and  $\theta_8 = 1.5$ . A mock data vector is generated by evaluating the forward model at  $\boldsymbol{\theta} = [0.0, -0.5, 0.5, 1.0, -2, 0.2, -0.2, 1.5]$ , making the nested model the ground truth once again. For inference we again consider Gaussian priors on all parameters with mean of 0.0 and diagonal covariance of 2.0.

Already in this setting we found it computationally challenging to normalize the histogram accurately. As a result we are unable to easily compute the SDDR with the histogram methodology by direct numerical integration. Instead, we rely on the normalizing flow methodology to calculate SDDR and find that it is in good agreement with the Bayes factor calculated from both nested sampling and the learned harmonic mean as shown in Table 2. Fig. 1 depicts the posterior contours with truth markers for both the nested and super model and the Bayes factor indicates that there is a definitive preference for the nested model, which is the model used to generate the ground truth.

Given that the toy Gaussian examples show good agreement between the Bayes factor calculated with the classical and flow SDDR methodologies, nested sampling and the learned harmonic mean, we will now turn our attention to more realistic cosmological examples.

### 3.2. Dark Energy Survey: $\Lambda$ CDM vs. $w$ CDM

The first more realistic cosmological scenario explored is with the Dark Energy Survey (DES; [The Dark Energy Survey Collaboration 2005](#)). We consider simulated DES Year 1 weak lensing and clustering data following the approach described in [Campagne et al. \(2023\)](#).

For this example two models are considered. The super model is the  $w$ CDM cosmological model where dark energy is treated as a fluid with equation of state characterised by the  $w$  parameter. The corresponding nested model is  $\Lambda$ CDM, which corresponds to setting  $w = -1$  in the  $w$ CDM model. In this work we directly make use of the posterior  $w$ CDM samples obtained already by [Polanska et al. \(2024a\)](#) by MCMC sampling. We also compare to the Bayes factors calculated by [Polanska et al. \(2024a\)](#) using the learned harmonic mean and nested sampling.

Table 3 reports the results obtained by our SDDR methodology for the DES cosmology analysis. There is a definitive preference for  $\Lambda$ CDM matching the truth. We

**Table 3:** DES cosmology  $\Lambda$ CDM vs.  $w$ CDM showing correct preference for  $\Lambda$ CDM given mock ground truth.

Method	Log Bayes factor
Nested sampling	$2.23 \pm 0.45$
Learned harmonic mean	$2.15 \pm 0.01$
SDDR (classical)	$2.14 \pm 0.02$
SDDR (flows)	$2.13 \pm 0.02$

**Table 4:** Stage-IV weak lensing  $\Lambda$ CDM vs.  $w_0w_a$ CDM showing correct preference for  $\Lambda$ CDM given mock ground truth.

Method	Log Bayes factor
Nested sampling	$0.78 \pm 0.71$
Learned harmonic mean	$1.53 \pm 0.07$
SDDR (classical)	$1.75 \pm 0.06$
SDDR (flows)	$1.75 \pm 0.06$

find that our methodology is in good agreement with the results obtained in the original analysis. A clear advantage here is that we are able to obtain these evidence ratios without needing to perform additional MCMC sampling of the nested  $\Lambda$ CDM model, approximately halving the computational cost as the SDDR method has negligible computational cost compared to the MCMC sampling of the models.

### 3.3. Stage-IV weak lensing: $\Lambda$ CDM vs. $w_0w_a$ CDM

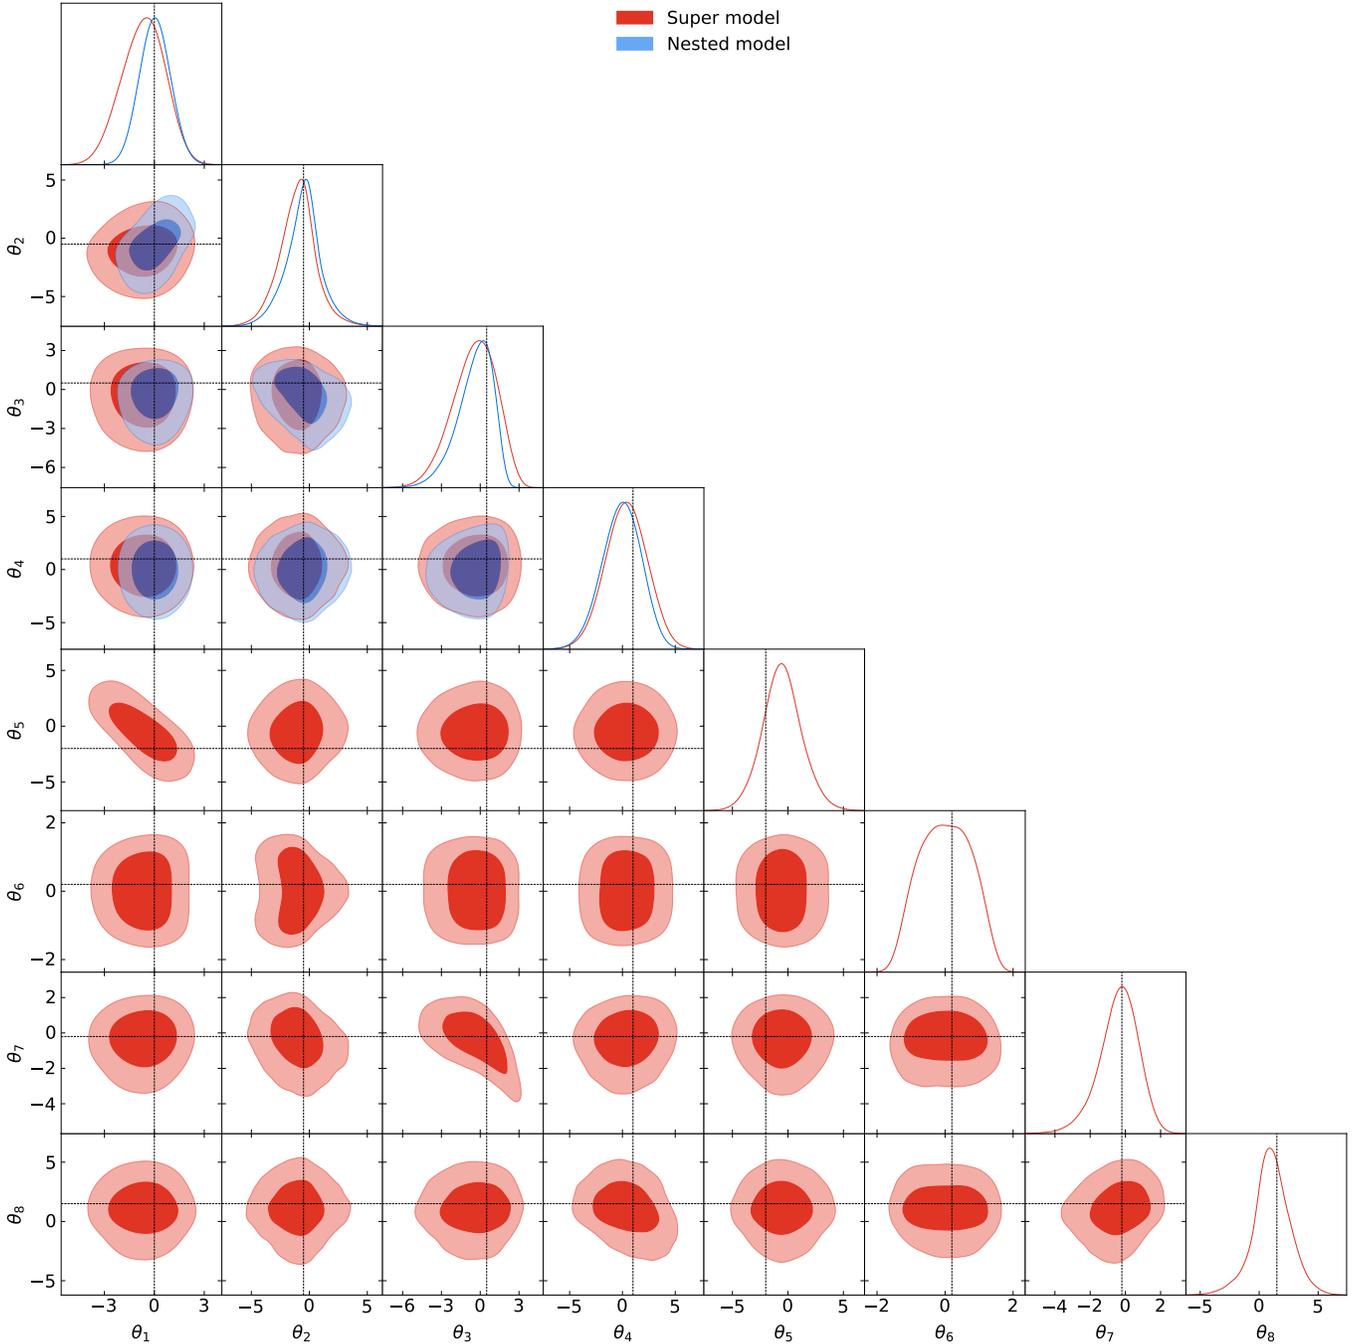
We explore another more realistic cosmological example that was also considered in [Piras et al. \(2024\)](#). The models being compared are those of  $\Lambda$ CDM and  $w_0w_a$ CDM. The super model of  $w_0w_a$ CDM is a phenomenological model where dark energy is not only treated as a fluid, but is also allowed to evolve with time, characterised by the two parameters of  $w_0$  and  $w_a$ . It reduces to  $\Lambda$ CDM when we set  $w_0 = -1$  and  $w_a = 0$ . This analysis followed a next-generation simulated weak lensing setup and featured a super model that had 39 parameters and a corresponding nested model with 37 parameters.

Such higher-dimensional spaces are typically more challenging to explore. If one is only interested in the Bayes factor, however, then the SDDR can alleviate challenges presented from direct evidence calculation as we only need concern ourselves with the extra parameters of the super model,  $\boldsymbol{\eta}$ , and the evaluation of that marginal posterior volume at  $\boldsymbol{\eta}_1$ .

Table 4 shows the results obtained through the classical density (histogram) and neural (normalizing flow) SDDR methodologies. The results we obtain with the SDDR are broadly consistent with the results obtained in [Piras et al. \(2024\)](#), showing correct preference for the  $\Lambda$ CDM model used to generate the ground truth.

### 3.4. Field-level weak lensing inference: $\Lambda$ CDM vs. $w$ CDM

In the field-level setting, where the evidence is typically challenging to compute due to the high-dimensionality of parameters, we can still make use of the SDDR to perform model comparison. The SDDR methodology is only dependent on the marginal posterior samples of parameter  $\boldsymbol{\eta}$  and the ability to evaluate the posterior probability at the point of nesting, i.e. at  $\boldsymbol{\eta} = \boldsymbol{\eta}_1$ , and so even in high-dimensional inference settings can provide



**Figure 1.** Posterior contours of the toy Gaussian model with 4 extra parameters including both the super model (red) and the nested model (blue). The dashed lines depict the ground truth.

a robust and precise estimate of the Bayes factor.

For the field-level setting, we consider the  $w$ CDM analysis presented in [Lanzieri et al. \(2024\)](#); [Zeghal et al. \(2024\)](#). This scenario considers  $w$ CDM with a field-level cosmological analysis constructed using a Bayesian hierarchical model (BHM) with a mock stage-IV weak lensing survey setup. Like the DES Y1 example presented previously, the  $w$ CDM model reduces to the  $\Lambda$ CDM model when setting  $w = -1$ . This is in general a costly setup to run even when harnessing modern hardware and gradient-accelerated methodologies, with [Zeghal et al. \(2024\)](#) writing that the explicit full field inference re-

quired  $\mathcal{O}(10^5)$  to  $\mathcal{O}(10^6)$  forward model evaluations to converge.

Previous works ([Lanzieri et al. 2024](#); [Zeghal et al. 2024](#)) did not explicitly calculate the Bayesian evidence. However, using the SDDR approach we are able to calculate the Bayes factor for the BHM setting between the  $w$ CDM and  $\Lambda$ CDM models for the first time. As we do not have access to the field-level evidence either calculated with the learned harmonic mean or with nested sampling directly for the BHM setting, we validate our results by comparison with a simulation-based inference (SBI) analysis that compressed the field with a convolu-

Method	Log Bayes factor
BHM SDDR (classical)	$0.951 \pm 0.012$
BHM SDDR (flows)	$0.947 \pm 0.011$
SBI learned harmonic mean	$1.093 \pm 0.014$
SBI SDDR (classical)	$0.905 \pm 0.018$
SBI SDDR (flows)	$0.934 \pm 0.020$

**Table 5:** Field-level cosmology example comparing  $w$ CDM vs.  $\Lambda$ CDM.

tional neural network. We leverage recent work presented in [Spurio Mancini et al. \(2024\)](#) which made use of the same simulator setup as in [Lanzieri et al. \(2024\)](#) but performed an SBI analysis based on neural likelihood estimation (in contrast to the neural posterior estimation performed by [Lanzieri et al. \(2024\)](#)). We make use of the SBI chains computed by [Spurio Mancini et al. \(2024\)](#) to calculate the log Bayes factor with the SDDR.

Table 5 shows the results obtained from the SDDR methodology for the field-level BHM and SBI  $w$ CDM analysis with both classical (histogram) and neural (normalizing flows) methods, showing consistent results and correct preference for the  $\Lambda$ CDM model used to generate the ground truth. We compare the Bayes factors computed by the SDDR for the SBI scenario to that computed by the learned harmonic mean estimator, finding close agreement with respect to the interpretation of the Bayes factor on the Jeffreys scale ([Jeffreys 1939](#); [Nesseris & García-Bellido 2013](#)) leading to the same conclusions. Moreover, the Bayes factors computed by the BHM and SBI approaches are also in close agreement, further validating the accuracy of not only the SDDR approach but also the SBI approximate inference framework.

#### 4. CONCLUSIONS

In this work we have developed a neural methodology to calculate the Savage-Dickey density ratio (SDDR) using normalizing flows. We have shown that our SDDR method scales to cases with many extra parameters involving high-dimensional marginal posterior spaces. We subsequently compared it to a classical approach of making use of normalized histograms, which is computationally challenging when there are many extra parameters. In low-dimensional marginal settings, both methods produce accurate and precise Bayes factor estimates. As expected, when extending the analysis to higher-dimensional marginal settings such as the 4D marginal posterior toy Gaussian example presented in [Sec. 3.1](#), the classical method begins to struggle. The normalizing flow method however continues to provide accurate and precise calculations of the Bayes factor. We further validated our normalizing flow methodology by producing consistent results with the classical approach and Bayes factors calculated with both nested sampling and the learned harmonic mean.

For a field-level inference setting, where the number of parameters is very large but the number of extra parameters of the super model is low, we leverage the SDDR to perform Bayesian model comparison for the first time. Furthermore, we show that Bayes factors computed for BHM (Bayesian hierarchical model) and SBI (simulation-based inference) inference approaches are consistent, providing further validation that the SBI approach has extracted as much cosmological information from the field as the BHM approach.

Our neural method also requires very little computational resources to train and run, typically running within a few minutes on a single CPU, given the MCMC chains. Making use of the neural SDDR with normalizing flows can approximately halve the computational cost of performing model comparison when working with nested models since one only needs to obtain posterior samples from the super model. The code implementing the neural SDDR methodology with normalizing flows is included within the latest `harmonic` software package release.

#### ACKNOWLEDGEMENTS

KL is supported by STFC (grant number ST/W001136/1) and the UK Space Agency (grant ST/X00208X/1). AP is supported by the UCL Centre for Doctoral Training in Data Intensive Science (STFC grant number ST/W00674X/1). DP was supported by the SNF Sinergia grant CRSII5-193826 ‘‘AstroSignals: A New Window on the Universe, with the New Generation of Large Radio-Astronomy Facilities’’. JDM is supported in part by STFC (grant number ST/W001136/1) and EPSRC (grant number EP/W007673/1).

#### REFERENCES

- Ashton G., et al., 2022, *Nature Reviews Methods Primers*, **2**, 39
- Buchner J., 2021, *The Journal of Open Source Software*, **6**, 3001
- Buchner J., 2023, *Statistics Surveys*, **17**, 169
- Cai X., McEwen J. D., Pereyra M., 2021, *arXiv e-prints*, p. [arXiv:2106.03646](#)
- Campagne J.-E., et al., 2023, *The Open Journal of Astrophysics*, **6**, 15
- Carrion K., Spurio Mancini A., Piras D., Hidalgo J. C., 2025, *MNRAS*, **539**, 3220
- Di Valentino E., Melchiorri A., Mena O., 2017, *Phys. Rev. D*, **96**, 043503
- Di Valentino E., Melchiorri A., Silk J., 2020, *Nature Astronomy*, **4**, 196
- DiCiccio T. J., Kass R. E., Raftery A., Wasserman L., 1997, *Journal of the American Statistical Association*, **92**, 903
- Dickey J. M., 1971, *The Annals of Mathematical Statistics*, **42**, 204
- Du G.-H., Wu P.-J., Li T.-N., Zhang X., 2025, *European Physical Journal C*, **85**, 392
- Duane S., Kennedy A., Pendleton B. J., Roweth D., 1987, *Physics Letters B*, **195**, 216
- Durkan C., Bekasovs A., Murray I., Papamakarios G., 2019, in *Advances in Neural Information Processing Systems 32 (NeurIPS 2019)*. Neural Information Processing Systems Foundation, Inc, pp 7511–7522, [https://neurips.cc/](#)
- Feroz F., Hobson M. P., 2008, *Monthly Notices of the Royal Astronomical Society*, **384**, 449
- Feroz F., Hobson M. P., Bridges M., 2009, *Monthly Notices of the Royal Astronomical Society*, **398**, 1601
- Feroz F., Hobson M. P., Cameron E., Pettitt A. N., 2019, *The Open Journal of Astrophysics*, **2**
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, **125**, 306
- Handley W. J., Hobson M. P., Lasenby A. N., 2015a, *Monthly Notices of the Royal Astronomical Society: Letters*, **450**, L61
- Handley W. J., Hobson M. P., Lasenby A. N., 2015b, *Monthly Notices of the Royal Astronomical Society*, **453**, 4384
- Heavens A., Fantaye Y., Mootooyaloo A., Eggers H., Hosenie Z., Kroon S., Sellentin E., 2017, *Marginal Likelihoods from Monte Carlo Markov Chains (arXiv:1704.03472)*, [doi:10.48550/arXiv.1704.03472](#)
- Higson E., Handley W., Hobson M., Lasenby A., 2019, *Statistics and Computing*, **29**, 891
- Hoffman M. D., Gelman A., 2014, *J. Mach. Learn. Res.*, **15**, 1593–1623
- Jeffreys H., 1939, *Theory of Probability*

- Jia H., Seljak U., 2020, in Zhang C., Ruiz F., Bui T., Dieng A. B., Liang D., eds, Proceedings of Machine Learning Research Vol. 118, Proceedings of The 2nd Symposium on Advances in Approximate Bayesian Inference. PMLR, pp 1–14, <https://proceedings.mlr.press/v118/jia20a.html>
- Kreisch C. D., Cyr-Racine F.-Y., Doré O., 2020, *Phys. Rev. D*, **101**, 123505
- Lange J. U., 2023, *MNRAS*, **525**, 3181
- Lanzieri D., Zeghal J., Makinen T. L., Boucaud A., Starck J.-L., Lanusse F., 2024, *arXiv e-prints*, p. [arXiv:2407.10877](https://arxiv.org/abs/2407.10877)
- Leistedt B., Peiris H. V., Verde L., 2014, *Phys. Rev. Lett.*, **113**, 041301
- Marin J.-M., Robert C., 2009, *arXiv e-prints*, p. [arXiv:0910.1452](https://arxiv.org/abs/0910.1452)
- McEwen J. D., Wallis C. G. R., Price M. A., Spurio Mancini A., 2021, *arXiv e-prints*, p. [arXiv:2111.12720](https://arxiv.org/abs/2111.12720)
- McEwen J. D., Liaudat T. I., Price M. A., Cai X., Pereyra M., 2023, p. [arXiv:2307.00056](https://arxiv.org/abs/2307.00056) ([arXiv:2307.00056](https://arxiv.org/abs/2307.00056)), [doi:10.48550/arXiv.2307.00056](https://doi.org/10.48550/arXiv.2307.00056)
- Neal R. M., 1996, Bayesian Learning for Neural Networks, Vol. 118 of Lecture Notes in Statistics. Springer-Verlag
- Nesseris S., García-Bellido J., 2013, *J. Cosmology Astropart. Phys.*, **2013**, 036
- O’Hagan A., Forster J., 2004, *Kendall’s advanced theory of statistics. Vol.2B: Bayesian inference, 2nd ed.*, 2B
- Papamakarios G., Nalisnick E., Rezende D. J., Mohamed S., Lakshminarayanan B., 2021, *The Journal of Machine Learning Research*, **22**, 2617
- Piras D., Polanska A., Mancini A. S., Price M. A., McEwen J. D., 2024, *The Open Journal of Astrophysics*, **7**, 73
- Polanska A., Price M. A., Spurio Mancini A., McEwen J. D., 2023, *Physical Sciences Forum*, **9**
- Polanska A., Price M. A., Piras D., Spurio Mancini A., McEwen J. D., 2024a, *arXiv e-prints*, p. [arXiv:2405.05969](https://arxiv.org/abs/2405.05969)
- Polanska A., Wouters T., Pang P. T. H., Wong K. W. K., McEwen J. D., 2024b, in Proceedings of the Machine Learning and Physical Sciences Workshop as part of the 38th International Conference on Neural Information Processing Systems. p. [arXiv:2410.21076](https://arxiv.org/abs/2410.21076) ([arXiv:2410.21076](https://arxiv.org/abs/2410.21076))
- Rinaldi S., Demasi G., Del Pozzo W., Hannuksela O. A., 2024, *arXiv e-prints*, p. [arXiv:2405.07504](https://arxiv.org/abs/2405.07504)
- Salvatelli V., Said N., Bruni M., Melchiorri A., Wands D., 2014, *Phys. Rev. Lett.*, **113**, 181301
- Skilling J., 2006, *Bayesian Analysis*, **1**, 833
- Speagle J. S., 2020, *MNRAS*, **493**, 3132
- Spurio Mancini A., Docherty M. M., Price M. A., McEwen J. D., 2023, *RAS Techniques and Instruments*, **2**, 710
- Spurio Mancini A., Lin K., McEwen J. D., 2024, *arXiv e-prints*, p. [arXiv:2410.10616](https://arxiv.org/abs/2410.10616)
- Srinivasan R., Crisostomi M., Trotta R., Barausse E., Breschi M., 2024, *arXiv e-prints*, p. [arXiv:2404.12294](https://arxiv.org/abs/2404.12294)
- Stiskalek R., Desmond H., 2024, *Research Notes of the American Astronomical Society*, **8**, 281
- Stiskalek R., Desmond H., Devriendt J., Slyz A., Lavaux G., Hudson M. J., Bartlett D. J., Courtois H. M., 2025, *arXiv e-prints*, p. [arXiv:2502.00121](https://arxiv.org/abs/2502.00121)
- The Dark Energy Survey Collaboration 2005, *arXiv e-prints*, pp [astro-ph/0510346](https://arxiv.org/abs/astro-ph/0510346)
- Trotta R., 2007, *Monthly Notices of the Royal Astronomical Society*, **378**, 72
- Verde L., Feeney S. M., Mortlock D. J., Peiris H. V., 2013, *J. Cosmology Astropart. Phys.*, **2013**, 013
- Verdinelli I., Wasserman L., 1995, *Journal of the American Statistical Association*, **90**, 614
- Zeghal J., Lanzieri D., Lanusse F., Boucaud A., Louppe G., Aubourg E., Bayer A. E., The LSST Dark Energy Science Collaboration 2024, *arXiv e-prints*, p. [arXiv:2409.17975](https://arxiv.org/abs/2409.17975)
- von Wietersheim-Kramsta M., Lin K., Tessore N., Joachimi B., Loureiro A., Reischke R., Wright A. H., 2025, *A&A*, **694**, A223

## APPENDIX

## A. SDDR DERIVATION

What follows is a derivation of SDDR based on Verde et al. (2013). Suppose we have two models that are nested, where  $M_1$  is the nested model and  $M_2$  is the super model with  $\theta$  common parameters and  $\eta$  extra parameters that extend the nested model to the super model. The Bayes factor is given by

$$\frac{z_1}{z_2} = \frac{p(\mathbf{d}|M_1)}{p(\mathbf{d}|M_2)} \quad (\text{A1})$$

$$= \frac{\int p(\mathbf{d}|\theta', M_1) p(\theta'|M_1) d\theta'}{\int p(\mathbf{d}|\theta'', \eta'', M_2) p(\theta'', \eta''|M_2) d\theta'' d\eta''} \quad (\text{A2})$$

$$= \frac{\int p(\mathbf{d}|\theta', \eta_1, M_2) p(\theta'|M_1) d\theta'}{\int p(\mathbf{d}|\theta'', \eta'', M_2) p(\theta''|M_1) p(\eta''|M_2) d\theta'' d\eta''}, \quad (\text{A3})$$

where in the second line we have used the fact that as the models are nested, we have  $p(\mathbf{d}|\theta, M_1) = p(\mathbf{d}|\theta, \eta_1, M_2)$ . In the third line we have used the fact that the prior distributions of  $\theta$  are identical for  $M_1$  and  $M_2$  and that the prior distributions of  $\eta$  and  $\theta$  are separable, i.e.  $p(\theta, \eta) = p(\theta)p(\eta)$ .

Multiplying both numerator and denominator by  $p(\eta_1|M_2)$  then gives

$$\frac{z_1}{z_2} = \frac{1}{p(\eta_1|M_2)} \frac{\int p(\mathbf{d}|\theta', \eta_1, M_2) p(\theta'|M_1) p(\eta_1|M_2) d\theta'}{\int p(\mathbf{d}|\theta'', \eta'', M_2) p(\theta''|M_1) p(\eta''|M_2) d\theta'' d\eta''}. \quad (\text{A4})$$

To simplify this expression, we note that the normalized parameter posterior for  $(\theta, \eta)$  under the second model applicable to all values of  $\eta$  is given by

$$p(\theta, \eta|\mathbf{d}, M_2) = \frac{p(\mathbf{d}|\theta, \eta, M_2) p(\theta|M_1) p(\eta|M_2)}{\int p(\mathbf{d}|\theta'', \eta'', M_2) p(\theta''|M_1) p(\eta''|M_2) d\theta'' d\eta''}. \quad (\text{A5})$$

Comparison with Eq. (A4) and setting  $\eta = \eta_1$  then yields

$$\frac{z_1}{z_2} = \frac{\int p(\boldsymbol{\theta}', \boldsymbol{\eta}_1 | \mathbf{d}, M_2) d\boldsymbol{\theta}'}{p(\boldsymbol{\eta}_1 | M_2)}. \quad (\text{A6})$$

We might notice that the integral in the numerator is performed only over a subset of parameters, and so is akin to marginalization, reducing to the marginalized posterior  $p(\boldsymbol{\eta}_1 | \mathbf{d}, M_2)$ . Hence, the evidence ratio or Bayes factor between two nested models is given by

$$\frac{z_1}{z_2} = \frac{p(\boldsymbol{\eta}_1 | \mathbf{d}, M_2)}{p(\boldsymbol{\eta}_1 | M_2)}. \quad (\text{A7})$$