



Universiteit
Leiden
The Netherlands

Distribution of Bayes' factor

Amendola, L.; Patel, V.; Sakr, Z.; Sellentin, E.; Wolz, K.

Citation

Amendola, L., Patel, V., Sakr, Z., Sellentin, E., & Wolz, K. (2024). Distribution of Bayes' factor. *Physical Review D - Particles, Fields, Gravitation And Cosmology*, 110(12).
doi:10.1103/PhysRevD.110.123522

Version: Publisher's Version

License: [Leiden University Non-exclusive license](#)

Downloaded from: <https://hdl.handle.net/1887/4198608>

Note: To cite this publication please use the final published version (if applicable).

Distribution of Bayes' factor

Luca Amendola[✉],* Vrund Patel[✉],* and Ziad Sakr[✉]*

Institut für Theoretische Physik, Universität Heidelberg, Philosophenweg 16, 69120 Heidelberg, Germany

Elena Sellentin[✉]

*Mathematical Institute, Leiden University,
Gorlaeus Building, Einsteinweg 55, NL-2333 CA Leiden, The Netherlands
and Leiden Observatory, Leiden University,
Gorlaeus Building, Einsteinweg 55, NL-2333 CA Leiden, The Netherlands*

Kevin Wolz[✉]

*Department of Physics, University of Oxford,
Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, United Kingdom*



(Received 14 August 2024; accepted 18 November 2024; published 16 December 2024)

The ratio of Bayesian evidences is a popular tool in cosmology to compare different models. There are however several issues with this method: Bayes' factor depends on the prior even in the limit of noninformative priors, and the Jeffreys scale, used to assess the test, is arbitrary. Moreover, the standard use of Bayes' factor is often criticized for being unable to reject models. In this paper, we address these shortcoming by promoting evidence ratios to frequentist statistics and deriving their sampling distributions. By comparing the evidence ratios to their sampling distributions, poor fitting models can now be rejected. Our method additionally does not depend on the prior in the limit of very weak priors, thereby safeguarding the experimenter against premature rejection of a theory with a uninformative prior, and replaces the arbitrary Jeffreys scale by probability thresholds for rejection. We provide analytical solutions for some simplified cases (Gaussian data, linear parameters, and nested models), and we apply the method to cosmological supernovae Ia data. We dub our method the FB method, for frequentist-Bayesian.

DOI: [10.1103/PhysRevD.110.123522](https://doi.org/10.1103/PhysRevD.110.123522)

I. INTRODUCTION

Modern cosmology is entering an era where large amounts of data will be forthcoming from cosmic microwave background, large scale structure, distance indicators, and gravitational waves. This requires a thorough understanding of the employed statistical methodology in order to make robust inferences from the data. Often investigations are performed in terms of a Bayesian study where the analysis of competing models is typically performed in three steps: first, the confidence region of a model is produced; second, two or more competing models are compared via Bayes' factors; and third, the quality of fit of the models is assessed via a simple χ^2 test [1–15]. These three tools are based on heterogeneous assumptions: for instance, the χ^2 quality of fit is a frequentist test that does not take into account priors, and the parameter confidence region cannot tell, *per se*, if an alternative point hypothesis outside this region is excluded, until we perform a specific

model comparison. Moreover, concerning Bayes' factor test limitations, its interpretation relies on arbitrary criteria to map its value onto the commonly used the Jeffreys [16,17], or Raftery-Kass scale [18]. Another important issue is that the result depends on the considered priors, even in the weak limit or in the case where an informative flat prior is used, since the evidence is averaged over the prior volume which can impact differently each model. Finally, the Bayesian comparison cannot reject models: it can only determine which model is relatively better, but not whether either model is a good fit to the data on its own.

Several other issues have been discussed in the literature. For instance, Ref. [19] showed that computing Bayesian evidence using cosmological data could yield, wrongly, a statistically insignificant preference for a model over a physically poorly motivated alternative model. Reference [20] showed that when comparing nested models, the Jeffreys scale can fail to penalize extra degrees of freedom when mock data are generated using the simpler model. Reference [21] noted that the performance of Bayes' factor test is strongly affected by

*These authors contributed equally to this work.

the signal to noise ratio in the data along with the non comprehensiveness of the models under consideration.

There have been some attempts, still within the Bayesian framework, to remedy for some of these limitations, in particular for the models comparison or rejection issue, by proposing alternatives such as the Akaike information criterion (AIC) test, the Bayesian information criterion (BIC) test or the deviance information criterion (DIC). In that regard, [22] performed a comparison of the above tests with Bayes' factor using cosmological data and found different levels of agreement between the outcome for certain models and probes. References [23] or [24] for a review, proposed the Bayesian model averaging method where each individual posterior, generated under the assumption of a particular model, is weighted by the model likelihood and then combined with the other posteriors to give a model-independent posterior. This method is especially useful in case the data are not strong enough to distinguish decisively between the models. Reference [25] introduced a statistical measure of the effective model complexity that assesses how many effective parameters a set of data can support and demonstrated that it can be used as a useful complement to Bayes' factor in model selection questions. Reference [26] develops a full set of estimators of internal and mutual agreement and disagreement between datasets whose strengths complement each other. This allows to take into account the effect of prior information and the statistical significance of discrepancies and unknown systematics. Reference [27] studied the sampling properties of Bayesian evidences as a frequentist statistic when compressing the data from its original size down to the number of free parameters. Finally, [28] proposed a method in which the likelihood distribution of the difference in χ^2 between a smoothed function and that of the best-fit of the model being tested is calculated. They then use the values that enclose 95% and 99% of the likelihood volume as criterion to test whether data support a given model, independent of how well other models perform, though they do not assess the effect of the prior on the final decision.

In this work we propose instead the use of Bayes' factor as a statistic to perform a frequentist p -value test. Moreover, we implement a frequentist version of Bayesian model comparison by using the odds ratio, i.e., the ratio of probabilities, between two competing models. These two complementary tools, p -value and frequentist model comparison, have interesting properties: they are entirely based on the same statistic; they generally depend on the prior but not in the limit of weak priors; they can assess both the quality of the fit and the relative probability. Assuming two competing models A and B , the final result can give one of the following answers: the data are compatible with A but not with B ; they are compatible with B but not with A ; they are compatible with both; they are incompatible with both. Finally, regardless of the above, the model comparison test can tell us

whether A is favored over B or viceversa. We illustrate the whole procedure with two complementary types of plots that provide a full analysis of the problem: one comparing the Bayes' factor distributions with the real data Bayes' factor, and another one comparing the model comparison odds ratio to the real data odds ratio.

Perhaps the first suggestion to perform a frequentist test with Bayes' factor is in Ref. [29] (see review in [30]). A study of the frequentist distribution of Internal Robustness ([31]), which is formed with ratio of evidences, was carried out in [32]. In Ref. [33], a Bayesian evidence was also considered as a statistic in the form of a p -value, to assess whether observed exceptional gravitational-wave events are consistent with extremal simulated events. Reference [34] studied the impact on model comparison if evidences are to be interpreted based on the ensemble of possible observations, and derived the mean and variance of Bayes's factor. Several further works studied various aspects of this concept, see for instance [35,36].

In this paper, the approach is quite close to Ref. [37]. They generated a distribution of evidences from datasets assuming a model and showed that the model is ruled out if its evidence lies outside that distribution. Here, in addition, we will assess the impact of the prior on the final decision and put forward a frequentist analog of the Bayesian model comparison. We also provide the analytical Bayes' factor distribution for correlated Gaussian data with linear nested models (thereby extending [27] who studied model rejection under data compression). In this paper, we apply our method, to be denoted as FB method, to a simple polynomial fit and to supernovae Ia data.

II. BAYES' FACTOR FOR GAUSSIAN DATA AND LINEAR MODELS

Bayes' factor is used in general to compare two competing models, where the fiducial model is compared to a new one by computing the ratio b defined as

$$b = \frac{E(D|M_A)}{E(D|M_B)}, \quad (1)$$

where E is the evidence

$$E(D|M) = \int d\theta \mathcal{L}(D|\theta, M) \mathcal{P}(\theta, M), \quad (2)$$

with $\mathcal{L}(D|M_{A,B})$ the probability of the data D given model $M_{A,B}$ with parameters θ_α and a prior $\mathcal{P}$. Model M_A is favored over model B if $b > 1$, and M_B is favored otherwise.

In this and the next section we will find the analytical distribution of Bayes' factor in a particular, but important, case, namely Gaussian data with linear nested models. In this section we start with reviewing the formalism and the general problem (see [38] for a review).

Linear models can be described by a sum $f_i = \theta_\alpha g_{\alpha i}$ where θ is the parameter vector and $g_{\alpha i}$ denotes the linear mapping between the α -th model parameter and the i th datapoint. For example, the parameters may act as the linear coefficients of a polynomial in a deterministic variable x , $f_i = \theta_0 + \theta_1 x_i + \theta_2 x_i^2 + \dots$. We use Greek subscripts referring to parameters, lowercase Latin subscripts referring to data, and capital Latin subscripts to refer to models. Moreover, we adopt the convention that repeated indexes are summed over. The likelihood for N data points d_i is

$$\mathcal{L} = (2\pi)^{-N/2} |\Sigma|^{-1/2} \exp -\frac{1}{2} (d_i - f_i) \Sigma_{ij}^{-1} (d_j - f_j) \quad (3)$$

where Σ is the data covariance matrix. We also assume a Gaussian prior with mean vector $\tilde{\theta}$

$$\mathcal{P} = (2\pi)^{-N_M/2} |P|^{1/2} \exp -\frac{1}{2} (\theta_\alpha - \tilde{\theta}_\alpha) P_{\alpha\beta} (\theta_\beta - \tilde{\theta}_\beta) \quad (4)$$

where N_M is the number of parameters for model M . The posterior is therefore $L\mathcal{P}/E$, where E is the evidence. In the limit of weak prior, $P_{\alpha\beta} \rightarrow 0$. We will often adopt this limit to obtain simpler expressions. Frequently used symbols are listed in Table I.

The best fit (maximum likelihood) parameters are found by imposing $\partial\mathcal{L}/\partial\theta_\alpha = 0$, i.e. solving the equation

$$g_{\alpha i} \Sigma_{ij}^{-1} (d_i - \hat{\theta}_\beta g_{\beta j}) = 0 \quad (5)$$

from which $\hat{\theta}_\beta = D_\alpha L_{\alpha\beta}^{-1}$ where $L_{\alpha\beta} = g_{\alpha i} \Sigma_{ij}^{-1} g_{\beta j}$ is the Fisher matrix, and $D_\alpha = d_i \Sigma_{ij}^{-1} g_{\alpha j}$. In the limit of weak prior, $\hat{\theta}$ coincides with the maximum of the posterior, and therefore gives the Bayesian best fit solution. For general Gaussian priors, the posterior maximum is

TABLE I. Frequently used symbols.

Symbol	Meaning
Σ_{ij}	Data covariance matrix
d_i	Data points
$g_{\alpha i}$	Fit functions matrix (design matrix)
$P_{\alpha\beta}$	Prior precision matrix (inverse of the parameter covariance matrix)
$L_{\alpha\beta}$	Likelihood parameter Fisher matrix
$F_{\alpha\beta} = L_{\alpha\beta} + P_{\alpha\beta}$	Posterior parameter Fisher matrix
N	Number of data points
N_M	Number of parameters for model M
$\hat{\theta}_\alpha$	Vector of best-fit parameters (dependent on the data)
$\tilde{\theta}_\alpha$	Vector of prior means
$\hat{\chi}^2$	Minimum χ^2 (dependent on the data)
$\beta = 2 \log b$	Log of ratio of evidences b (Bayes' ratio)

$$\theta^* = (\hat{\theta}L + \tilde{\theta}P)^{-1}. \quad (6)$$

For linear models and Gaussian data, the maximum likelihood parameters are Gaussian random variables and the Fisher matrix represents their exact inverse covariance. In Sec. IV we briefly discuss the Fisher approximation for nonlinear models.

We define the matrix G as $g_{\alpha i} L_{\alpha\beta}^{-1} g_{\beta j} \equiv G_{ij}$. In index-free notation, $\hat{\theta}^T = D^T L^{-1}$, $D^T = d^T \Sigma^{-1} g^T$, $G = g^T L^{-1} g$ and $L = g \Sigma^{-1} g^T$. We will use both indexes and index-free notation, whichever is clearer in each case. It follows

$$f = \hat{\theta}^T g = D^T L^{-1} g = d^T \Sigma^{-1} g^T L^{-1} g = d^T \Sigma^{-1} G. \quad (7)$$

We see that $C \equiv \Sigma^{-1} G$ is the matrix that projects the information from d to the theory best-fit $g^T \hat{\theta}$. (Note that $C^2 = C$, so it is indeed a projector; it is also a symmetric matrix). Note also that

$$\begin{aligned} \text{Tr} C &= \text{Tr}(\Sigma^{-1} G) = \text{Tr}(\Sigma^{-1} g^T L^{-1} g) \\ &= \text{Tr}(g \Sigma^{-1} g^T L^{-1}) = \text{Tr}(L L^{-1}) = N_M \end{aligned} \quad (8)$$

where N_M is the number of parameters of model M . Then on the best fit one has

$$\hat{\chi}^2 = (d^T - \hat{\theta}^T g) \Sigma^{-1} (d - g^T \hat{\theta}) \quad (9)$$

$$= d^T \Sigma^{-1} (I - C) d. \quad (10)$$

We now rewrite the likelihood as a function of the parameters, by writing $f_i = \hat{\theta}_\alpha g_{\alpha i} + (\theta_\alpha - \hat{\theta}_\alpha) g_{\alpha i}$ and using the fact that

$$\begin{aligned} \frac{1}{2} (d_i - f_i) \Sigma_{ij}^{-1} (d_j - f_j) &= \frac{1}{2} (d_i - \hat{\theta}_\alpha g_{\alpha i}) \Sigma_{ij}^{-1} (d_j - \hat{\theta}_\beta g_{\beta j}) \\ &\quad - (\theta_\alpha - \hat{\theta}_\alpha) g_{\alpha i} \Sigma_{ij}^{-1} (d_j - \hat{\theta}_\beta g_{\beta j}) \\ &\quad + \frac{1}{2} (\theta_\alpha - \hat{\theta}_\alpha) g_{\alpha i} \Sigma_{ij}^{-1} g_{\beta j} (\theta_\beta - \hat{\theta}_\beta) \\ &= \frac{1}{2} \hat{\chi}^2 + \frac{1}{2} (\theta_\alpha - \hat{\theta}_\alpha) L_{\alpha\beta} (\theta_\beta - \hat{\theta}_\beta) \end{aligned} \quad (11)$$

since $(\theta_\alpha - \hat{\theta}_\alpha) g_{\alpha i} \Sigma_{ij}^{-1} (d_j - \hat{\theta}_\beta g_{\beta j}) = 0$ on account of Eq. (5). Therefore

$$\mathcal{L} = (2\pi)^{-N/2} |\Sigma|^{-1/2} e^{-\frac{1}{2} \hat{\chi}^2} \exp -\frac{1}{2} (\theta_\alpha - \hat{\theta}_\alpha) L_{\alpha\beta} (\theta_\beta - \hat{\theta}_\beta). \quad (12)$$

Using this expression, one finds that the evidence for Gaussian data, linear model and Gaussian priors is

$$E = \int \mathcal{L} \mathcal{P} d^N \theta = (2\pi)^{-N/2} |\Sigma|^{-1/2} e^{-\frac{1}{2} \hat{\chi}^2} \frac{|P|^{1/2}}{|F|^{1/2}} \exp \left[-\frac{1}{2} (\hat{\theta} - \tilde{\theta})^T L F^{-1} P (\hat{\theta} - \tilde{\theta}) \right]. \quad (13)$$

In the linear fit case, the data (our random variables) are only in $\hat{\theta}$ and in $\hat{\chi}^2$; all the other quantities are not random variables.

Bayes' ratio for models A (with N_A free parameters) and B (with $N_B > N_A$ free parameters) defined as the ratio of evidences, becomes then (see [39,40] for similar calculation of the Bayes factor with the Fisher matrix approximation)

$$b \equiv \frac{E_A}{E_B} = \frac{|\Sigma|^{-1/2} e^{-\frac{1}{2} \hat{\chi}_A^2} \frac{|P_A|^{1/2}}{|F_A|^{1/2}} \exp \left[-\frac{1}{2} (\hat{\theta}_A - \tilde{\theta}_A)^T L_A F_A^{-1} P_A (\hat{\theta}_A - \tilde{\theta}_A) \right]}{|\Sigma|^{-1/2} e^{-\frac{1}{2} \hat{\chi}_B^2} \frac{|P_B|^{1/2}}{|F_B|^{1/2}} \exp \left[-\frac{1}{2} (\hat{\theta}_B - \tilde{\theta}_B)^T L_B F_B^{-1} P_B (\hat{\theta}_B - \tilde{\theta}_B) \right]} \quad (14)$$

$$= \exp \left[-\frac{1}{2} (\hat{\chi}_A^2 - \hat{\chi}_B^2) - \frac{1}{2} (\chi_{P,A}^2 - \chi_{P,B}^2) \right] \frac{|P_A F_B|^{1/2}}{|P_B F_A|^{1/2}} \quad (15)$$

where

$$\chi_{P,M}^2 = (\hat{\theta}_M - \tilde{\theta}_M)^T L_M F_M^{-1} P_M (\hat{\theta}_M - \tilde{\theta}_M). \quad (16)$$

We have then

$$\beta \equiv 2 \ln b = \chi_B^2 - \chi_A^2 + \chi_{P,B}^2 - \chi_{P,A}^2 + Y_B - Y_A \quad (17)$$

where

$$Y_M = \ln \frac{|F_M|}{|P_M|} \quad (18)$$

with $M = A, B$. In Eq. (17), the first pair of terms gives the distance between the likelihood maxima, the second pair measures how close are the best-fits to the prior mean, and the third pair measures how new data improve upon the prior. In the standard Bayesian approach, model A is preferred if β is positive and large, that is, when it is a good fit (small χ_A^2), which does not depart too much from the prior (small $\chi_{P,A}^2$), and when the gain from the prior to the final uncertainty is minimal (small Y_A). From Eq. (17) we already intuitively see that a “robust” rejection of one model (i.e., a rejection that holds up under repeated draws of the data) depends on the relative magnitude of the χ_M^2 terms with respect to the Y_M terms. If $Y_M \gg \chi_M^2$, then β will hardly scatter under repetitions of the experiment. However, Y_M sets the experiments measurement precision with respect to the experiments prior range. We therefore see that an experiment that fails to strongly update the prior indeed also produces a β that strongly scatters, and hence warrants a comparison to the sampling distribution of β for reliable model selection.

We removed the hat from χ^2 since from now on we interpret β as a frequentist statistic, i.e. as a random variable which can assume any value, while $\hat{\chi}^2$ was the specific value obtained with the real dataset.

For nested models, there will be a number of common parameters, and some additional parameters for model B (we call them “common parameters” and “extra parameters,” respectively, in the following). We assume that the prior means for the common parameters are the same for both models, and zero for the extra ones (this is not restrictive: one can always redefine the extra parameters so that when they vanish model B reduces to model A). The elements g_{ai} will then be the same for the common parameters, and the prior mean vector $\tilde{\theta}$ for model B will be the same as for model A , plus as many zeros as the extra parameters.

In the limit of weak prior, since

$$L_A F_A^{-1} P_A = L_A (L_A + P_A)^{-1} P_A \quad (19)$$

$$= (P_A^{-1} + L_A^{-1})^{-1} \approx P_A \quad (20)$$

vanishes, also $\chi_{P,M}^2$ vanishes, and we have

$$\beta = \chi_B^2 - \chi_A^2 + Y_B - Y_A. \quad (21)$$

In this limit, the statistic β has the same distribution as the log of the likelihood ratio (up to the constant $Y_B - Y_A$).

If, in the same weak prior limit, for simplicity one assumes that both F and P are diagonal matrices with entries σ^{-2} and σ_P^{-2} , respectively, for each parameter, and $\sigma_P = \alpha \sigma$, then one has $|F_M| = \sigma^{-2N}$ and $|P| = \sigma_P^{-2N}$, and therefore

$$\beta = \chi_B^2 - \chi_A^2 + 2r \log(\alpha) \quad (22)$$

where $r = N_B - N_A$, which always favors A for $\alpha \rightarrow \infty$. This argument can be directly extended to any likelihood, not just Gaussian. In fact, if the prior is very broad with respect to the likelihood, then its value can be taken approximately constant over a large volume V_M in parameter space, so that $\mathcal{P} \approx 1/V_M$. Then one has

$$b \approx \frac{\int \mathcal{L} \mathcal{P} d^{N_A} \theta_A}{\int \mathcal{L} \mathcal{P} d^{N_B} \theta_B} \rightarrow \frac{V_B \int \mathcal{L} d^{N_A} \theta_A}{V_A \int \mathcal{L} d^{N_B} \theta_B}. \quad (23)$$

If models are nested, then $V_B = V_A V_e$, where V_e is the volume of the extra parameters. Then again one sees that b increases arbitrarily for large prior volume V_e . The unappealing consequences of this behavior in a simple but extreme example are illustrated in the Appendix.

Note that the artificial suppression of the new model arises because the prior was subjectively chosen as one that is effectively a flat prior with wide boundaries. If the prior is derived by sound physical or mathematical principles (see e.g. [41] who computed an uninformative prior that does not disturb the low signal-to-noise measurements of elusive neutrino physics), then of course the prior is part of the model and no issue arises.

III. ANALYTICAL DISTRIBUTION OF BAYES' RATIO

Our next goal now is to answer this question: given that the data are Gaussian random variables distributed around a given model, what is the distribution of β ? We always assume the models to be nested, i.e. model B has a subset of parameters that coincide with those of model A, and with the same prior, plus a number of extra parameters with their corresponding prior. In other words, model A is a constrained version of model B. Moreover, without lack of generality, model B reduces to model A if the extra parameters are put to zero. This applies, for instance, to Λ CDM and w CDM, which is the real data example we investigate later on, provided one takes as extra parameter $w' = w - 1$.

It is important to realize that the random variables (the data) enter quadratically both in $\chi^2_{A/B}$ and in $\chi^2_{P,B/A}$. A positive-defined quadratic form of Gaussian variables $x_i M_{ij} x_j$, as β in Eq. (21), is distributed as a χ^2 variable. There are however three types of χ^2 variables: standard, noncentral, and generalized. If the Gaussian variables x have zero mean, the χ^2 distribution is the standard one, defined uniquely by the degrees of freedom (d.o.f.) ν . If the Gaussian variables have mean μ_i such that

$$\lambda = \mu_i M_{ij} \mu_j \quad (24)$$

then they are noncentral χ^2 variables, defined by ν and λ . Finally, a linear combination of noncentral χ^2 variables is distributed as a generalized χ^2 variable. This distribution is defined by the ν , λ of each χ^2 variables, and by their coefficients of the linear combination. More details on these distributions will be presented below.

As we will see, the standard case applies to nested models with weak priors when assuming A to be true and the prior mean to coincide with the best fit values. The noncentral case applies to nested models with weak priors

when the prior mean does not coincide with the best fit values, or with any prior but the two models differ by just one parameter. The generalized case, finally, applies in all other situations. We need also consider variables that are shifted and scaled versions of χ^2 variables, i.e. $z \equiv \alpha \chi^2 + a$, and their distributions are trivially obtained by shifting and scaling the argument of the original χ^2 distributions.

We need then to show that the combination $(\chi_B^2 + \chi_{P,B}^2) - (\chi_A^2 + \chi_{P,A}^2)$, which are the terms that depend on data, is a quadratic form and find out its distribution. For each term we have (we temporarily suppress here the subscript A or B)

$$\chi^2 + \chi_P^2 = \Xi + \Pi \quad (25)$$

where

$$\Xi \equiv (d_m - u_m) T_{mk} (d_k - u_k) \quad (26)$$

$$\Pi \equiv -\tilde{\theta}_\gamma (W_{m\gamma} T_{mk} W_{k\alpha} - L_{\gamma\beta} F_{\beta\delta}^{-1} P_{\delta\alpha}) \tilde{\theta}_\alpha \quad (27)$$

$$u_k = T_{kn}^{-1} \Sigma_{n\ell}^{-1} g_{\alpha\ell} (F^{-1} P)_{\alpha\gamma} \tilde{\theta}_\gamma = W_{k\gamma} \tilde{\theta}_\gamma \quad (28)$$

$$T = \Sigma^{-1} (I - g^T F^{-1} g \Sigma^{-1}). \quad (29)$$

However, one can show that $\Pi = 0$. It turns out that

$$T^{-1} = \Sigma + g^T P^{-1} g \quad (30)$$

and therefore

$$u = g^T \tilde{\theta} \quad (31)$$

i.e., $W = g$. Since Ξ is a quadratic form of Gaussian variables, it is distributed as one of the three types of χ^2 variable.

Let us denote with μ_m the expected value of $d_m - u_m$

$$\mu = \langle d - u \rangle = g^T (\hat{\theta} - \tilde{\theta}) \quad (32)$$

We also have $\langle d \rangle = f = \hat{\theta} g$ and, by definition of covariance matrix,

$$\langle (d - g^T \hat{\theta})(d^T - \hat{\theta}^T g) \rangle = \Sigma \quad (33)$$

so that

$$\begin{aligned} \langle \Xi \rangle &= \langle (d - g^T \hat{\theta} + \mu)^T T ((d - g^T \hat{\theta}) + \mu) \rangle \\ &= \langle (d - g^T \hat{\theta})^T T (d - g^T \hat{\theta}) \rangle + \mu^T T \mu - 2\mu^T T \langle (d - g^T \hat{\theta}) \rangle \\ &= T \Sigma + \mu^T T \mu. \end{aligned} \quad (34)$$

From now on we reinsert the subscripts A or B. Now we can consider the combination of A and B, i.e. $\beta = \Xi_B - \Xi_A + Y_B - Y_A$. We assume that the prior mean

vector $\tilde{\theta}$ is the same for the common parameters, and zero for the additional B parameters. This is where the condition of nestedness enters and simplifies considerably the problem. Then clearly $u_A = u_B = u$. We have then

$$(d_m - u_m)T_{B,mk}(d_k - u_k) - (d_m - u_m)T_{A,mk}(d_k - u_k) = (d_m - u_m)T_{mk}(d_k - u_k) \quad (35)$$

where

$$T = T_B - T_A = \Sigma^{-1}(g^T F^{-1} g|_A - g^T F^{-1} g|_B) \Sigma^{-1}. \quad (36)$$

In order to define the mean

$$\mu = \langle d - u \rangle = \langle d \rangle - u \quad (37)$$

we need to specify whether we assume A or B to be true. Then we have

$$\mu_M = g_M^T(\hat{\theta}_M - \tilde{\theta}_M). \quad (38)$$

In the following, μ_M will take different values depending on whether one assumes $M = A$ or $M = B$ to be true. In the general case we have the following mean and variance:

$$\begin{aligned} \langle \beta \rangle &= \langle \Xi_B \rangle - \langle \Xi_A \rangle + Y_B - Y_A \\ &= \text{Tr}(T\Sigma) + \mu_M^T T \mu_M + Y_B - Y_A \end{aligned} \quad (39)$$

$$V[\beta] = \sigma_\beta^2 = 2\text{Tr}(T\Sigma)^2 + 4\mu_M^T T \Sigma T \mu_M. \quad (40)$$

Let us first consider the weak prior limit. In this case, $F \approx L$, and $T_M \Sigma \approx \Sigma^{-1}(I - g_M^T L^{-1} g_M \Sigma^{-1}) \Sigma = I - \Sigma^{-1} G_M = I - C_M$.¹ Since C_M is a projector, its eigenvalues are 1 or 0. Because of Eq. (8), we see that C_M has N_M unit eigenvalues and $N - N_M$ zero eigenvalues. Then the matrix

$$-T\Sigma = -I + C_B + I - C_A = C_B - C_A \quad (41)$$

has $r = N_B - N_A$ unit eigenvalues and the rest are zero. We conclude that $-\beta + (Y_B - Y_A)$ is a noncentral χ^2 variable with $\nu = r$ degrees of freedom and noncentral parameter

$$\lambda_M = -\mu_M^T T \mu_M. \quad (42)$$

That is, β is distributed as

$$\begin{aligned} P(x; a, \nu, \lambda) &= \frac{1}{2} e^{-(a-x+\lambda)/2} \left(\frac{a-x}{\lambda} \right)^{\nu/4-1/2} \\ &\times I_{\nu/2-1}(\sqrt{\lambda(a-x)}) \end{aligned} \quad (43)$$

¹One can also simplify the expression for the variance by noting that $(T_M \Sigma)^2 = T_M \Sigma$ and therefore $T_M \Sigma T_M = T_M \Sigma T_M \Sigma \Sigma^{-1} = (T_M \Sigma)^2 \Sigma^{-1} = T_M \Sigma \Sigma^{-1} = T_M$.

with $a = Y_B - Y_A$, and where I is a modified Bessel function of the first kind. In the special case in which the elements of the prior mean vector $\tilde{\theta}$ coincide with model A's best fit parameters $\hat{\theta}_A$ for the common parameters, and vanishes for the extra B parameters, then $\mu_A = 0$, so that $\lambda_A = 0$, and β is a standard χ^2 variable up to shift and inversion, so its distribution is

$$P(x; a, \nu) = \frac{(a-x)^{\nu/2-1} e^{-(a-x)/2}}{2^{\nu/2} \Gamma(\nu/2)}. \quad (44)$$

Finally, if the prior is not negligible, then the eigenvalues of $-T\Sigma$ are no longer 1 or 0's, and β can be written as a sum of noncentral χ^2 variables. This sum follows a generalized χ^2 distribution, as we discuss next.

We need to find the distribution of a general quadratic form

$$Q = (d_m - u_m)T_{mk}(d_k - u_k) = X^T T X \quad (45)$$

This can be obtained as follows (see [42,43]). The mean of X is $\mu = \langle X \rangle = g^T(\hat{\theta} - \tilde{\theta}) = f - u$. The covariance matrix of X is

$$\Sigma = \langle (d - u - \mu)(d - u - \mu)^T \rangle \quad (46)$$

$$= \langle (d - f)(d - f)^T \rangle \quad (47)$$

Let now C be a matrix such that (Cholesky decomposition)

$$CC^T = \Sigma \quad (48)$$

and P an orthonormal matrix (i.e., $PP^T = I$) that diagonalizes CTC^T , i.e.

$$PCTC^T P^T = D = \text{diag}(\alpha_1, \alpha_2, \dots, \alpha_p) \quad (49)$$

where the eigenvalues of CTC^T α_i are nonzero for $i = 1 \dots r$ and zero for $i = r, \dots, p$. The number of nonzero eigenvalues should be $r = N_B - N_A$. Let also define the two vectors $Y = P(C^T)^{-1}X$ and $\xi = P(C^T)^{-1}\mu$. Then one can see that Y is a vector of Gaussian variables with mean vector ξ and covariance matrix $\Sigma_Y = I$. Therefore Q can be written as a linear combination of noncentral, uncorrelated χ^2 variables,

$$Q = X^T T X = \sum_i^r \alpha_i Y_i^2 = \sum_i^r \alpha_i \chi_1^2(\lambda_i) \quad (50)$$

where $\chi_1^2(\lambda)$ are noncentral χ^2 variables with one degrees of freedom and noncentral parameter

$$\lambda_i = \xi_i^2. \quad (51)$$

Then, in general, Q is distributed as a generalized χ^2 variable. The characteristic function (CF) for the generalized χ^2 variable can be obtained analytically [42]

$$\phi(t) = \frac{\exp\left\{it \sum_i^r \frac{\lambda_i \alpha_i}{1-2i\alpha_i t}\right\}}{\prod_i^r (1-2i\alpha_i t)^{\frac{1}{2}}}. \quad (52)$$

However, one cannot derive analytically neither the probability density function (PDF) nor the cumulative distribution function. One can numerically determine the PDF performing an inverse Fourier transform or utilizing existing codes (e.g. [43]).

We have shown earlier that in the limit of weak prior, the eigenvalues are indeed unity. In this case, Q is a noncentral variable with $\lambda = \xi^T \xi$ and r d.o.f.s. Finally, if the eigenvalues are unity and $\lambda = 0$, Q is a standard χ^2 variable with r d.o.f.s. This shows how the generalized χ^2 distribution reduces to one of the simpler forms.

In the case $r = 1$, Q is a single noncentral variable with one d.o.f. and

$$\lambda = \xi^2. \quad (53)$$

In this case, the distribution of $\beta = \alpha\chi_1^2(\lambda) + a$ is obtained from the noncentral distribution $P(\chi_1^2(\lambda))$ as

$$P_\beta(\beta) = P(\chi_1^2(\lambda)) \left| \frac{d\chi_1^2(\lambda)}{d\beta} \right| = \frac{1}{|\alpha|} P\left(\frac{\beta - a}{\alpha}\right). \quad (54)$$

To summarize, if one assumes a Gaussian prior and Gaussian data with mean model linear in the parameters, then the logarithm of the Bayes factor is a random variable distributed like one of three kinds of χ^2 variables (central, noncentral, or generalized χ^2), depending on the mean and width of the prior relative to the likelihood.²

The Bayes' distributions can now be employed to build our FB method, i.e. to test the hypothesis that, given that model M is true, the data are consistent with M rather than with an alternative M' . Since a p -value test can only reject hypotheses, and not assess whether they are true, we need to make two tests when comparing A versus B . In the first, we assume A to be true, so that the null hypothesis is H_0 : (A is true), and we either reject or nonreject A ³; in the second test, we assume B to be true, so H_0 : (B is true), and again we can reject or nonreject B . Together, the two tests

²Clearly, the logarithm of the evidence is also a χ^2 variable, now with $N - N_M$ d.o.f.s. One could also envisage a frequentist test based on the evidence of a single model, without comparing it to an alternative. In this paper, however, we explore only the properties of Bayes' ratio.

³In this paper, we will adopt the terminology commonly used by in cosmological statistical data analysis for frequentist hypothesis testing, replacing the terms Type I/II error with the expressions like "rejecting" or "not rejecting the null hypothesis" H_0 .

can produce four distinct outcomes: reject both models, reject neither of them, reject only A or only B . One can consider both one-tailed tests, as commonly done for likelihood ratio tests, or two-tailed tests, as is typically the case for χ^2 tests. In this second case, the point of view is that too good fits are rejected because incompatible with the assumed data uncertainties. Here, however, we only discuss one-tailed tests for simplicity. The entire procedure will be illustrated below in Secs. VI and VII.

IV. FISHER APPROXIMATION

In many realistic applications, parameters do not enter linearly in the likelihood. The fitting function f can however be linearized near the best fit as

$$f(x, \theta) \approx f(x)|_{\theta=\hat{\theta}} + \sum_\alpha \frac{\partial f}{\partial \theta_\alpha} \Big|_{\theta_\alpha=\hat{\theta}_\alpha} (\theta_\alpha - \hat{\theta}_\alpha). \quad (55)$$

The Fisher matrix for Gaussian data with parameters only in the mean vector f_i is then

$$F_{\alpha\beta} = \left(\frac{\partial f_i}{\partial \theta_\alpha} \Sigma_{ij}^{-1} \frac{\partial f_j}{\partial \theta_\beta} \right)_{\theta_\alpha=\hat{\theta}_\alpha}. \quad (56)$$

Then we see that we can identify

$$g_{ai} = \frac{\partial f(x_i)}{\partial \theta_\alpha} \Big|_{\theta_\alpha=\hat{\theta}_\alpha}. \quad (57)$$

We can then use a Fisher matrix approximation to obtain the matrix L

$$L_{\alpha\beta} = g_{ai} \Sigma_{ij}^{-1} g_{bj} \quad (58)$$

and the problem is reduced to the one of the previous sections. Whether this approximation is acceptable, however, should be analyzed on a case-by-case basis. We apply this Fisher approximation in Sec. VII.

V. A FREQUENTIST MODEL-COMPARISON TEST

The use of the β distribution to see whether models A or B are compatible with the data is an absolute test of compatibility of data with models (goodness of fit). But we can also directly compare A with B and see which one is to be considered a better fit to the data (model comparison). The result of this model comparison will be a statement about which model better explains the data, even if both are, on absolute terms, unacceptable fits. These two tools, together, form the content of our FB method.

In Bayesian context, $b = e^{\beta/2}$ is not a random variable. Its value expresses the odds ratio of A versus B , i.e. the ratio of the probabilities that, given the observed data D , the model A or the model B is to be preferred,

$$b(D) = e^{\beta/2} = \frac{P(A)}{P(B)} \quad (59)$$

The value of β is then compared to an empirical scale (the Jeffreys scale or one of its variants) to convert the number into an expression of ordinary language. The conditioning over D will be implicit in the following.

Alternatively, we can also convert b into the probability of A given that the possible models are only A or B as follows (the subscript b here stands for Bayesian)

$$\begin{aligned} R_b(A) &\equiv P(A|A \text{ or } B) = \frac{P(A \text{ and } (A \text{ or } B))}{P(A \text{ or } B)} \\ &= \frac{P(A)}{P(A) + P(B)} = \frac{b}{1+b} = (1 + e^{-\beta/2})^{-1} \end{aligned} \quad (60)$$

In this way, Bayes' ratio can be interpreted as the probability of "event" A within a set of possible disjoint events composed only by A and B . This probability over the space of events is well normalized: in fact, one also has

$$R_b(B) \equiv P(B|A \text{ or } B) = \frac{P(B)}{P(A) + P(B)} = 1 - R_b(A) \quad (61)$$

Following this point of view, the Jeffreys scale can be replaced by ordinary probability thresholds. For instance we can accept model A versus B at 68% CL if

$$R_b(A) = \frac{b}{1+b} \geq 0.68 \quad (62)$$

or $b \geq 2.125$ ("barely worth mentioning" in Jeffreys' language), at 95% CL if $b \geq 19$ ("strong") and at 99.7% if $b \geq 332$ ("decisive"). Naturally, if there are other models to consider, and not just A and B (as is almost always the case), then this scale is hardly meaningful because it depends on the arbitrary number of models we decide to include. The main problem, however, as we already emphasized, is that R_b depends on the prior also in the limit of weak priors

$$\begin{aligned} R_b &= (1 + e^{-\beta/2})^{-1} = (1 + e^{-(\chi_B^2 - \chi_A^2 + Y_B - Y_A)/2})^{-1} \\ &= \left(1 + \sqrt{\frac{|F_A||P_B|}{|P_A||F_B|}} e^{-\frac{1}{2}(\chi_B^2 - \chi_A^2)} \right)^{-1} \end{aligned} \quad (63)$$

For instance, let us assume that model B is identical to model A except for an additional parameter that has an uncertainty σ_B and a prior with a very large uncertainty $\sigma_{P,B} \rightarrow \infty$. Then, if both models have a diagonal Fisher matrix, we can write $F_B = F_A \sigma_B^{-2}$ and $P_B = P_A \sigma_{P,B}^{-2}$, and one gets

$$R_b(A) \approx 1 - \left(\frac{\sigma_B}{\sigma_{P,B}} \right)^2 e^{-\frac{1}{2}(\chi_B^2 - \chi_A^2)} \rightarrow 1 \quad (64)$$

This shows that model A is always favored in this limit, regardless of its quality. That is, the addition of a parameter with a very weak prior, maybe because B is a new model with little experimental history, is discarded *a priori*. This carries the risk of artificially enhancing trust on established models when confronted with unbeaten paths (see the Appendix for an example).

In the frequentist interpretation, the Bayesian odds ratio should be replaced by the ratio of the distributions of β when assuming either A or B . We can then form the frequentist (subscript f) analog of R_b as

$$R_f(A) = \frac{P_A(\alpha_A \beta + a)}{P_A(\alpha_A \beta + a) + P_B(\alpha_B \beta + a)} \quad (65)$$

and of course its complement, $R_f(B) = 1 - R_f(A)$. We denoted with $P(\alpha_M \beta + a)$ a distribution (χ^2 or otherwise) for the variable $\chi^2 = \alpha_M \beta + a$ with $M = A, B$. Since in the weak prior limit $\alpha_M \beta + a = \chi_A^2 - \chi_B^2$, we see that R_f , contrary to R_b , becomes independent of the prior. A value of $R_f(A)$ larger than 1/2 means that A is favored with respect to B ; viceversa, $R_f(A) < 1/2$ means that B is favored. However, here we encounter the same caveat as in the Bayesian case, namely, that R_f can be directly compared to a Gaussian scale only if the models A and B are exhaustive of the model landscape.

VI. APPLICATIONS: POLYNOMIAL FITS

Now that we know the distribution of the statistic β , we can form an analytical p -value test. Let us assume as null hypothesis H_0 : model A is better than B in explaining the data, and as alternative hypothesis H_1 : model B is better than A . We can then find a specific value of β from a given dataset and compare it to the β distribution.

As a first illustration of the method, we define two models: model A (a straight line $f = a + bx$) and model B (a quadratic function $f = a + bx + cx^2$, with the same a, b as for model A). We take as prior for a, b, c a Gaussian with mean $a = 1, b = 1$ and two values of c , and diagonal prior covariance with a very broad variance, $\sigma_P = 10$ for each parameter (to be compared to uncertainties between 0.1 and 1 for the three parameters). Then we generate many sets of Gaussian data points distributed around model A with a given data covariance matrix Σ (we put for simplicity $\Sigma = I$, but we tested also correlated cases). Every time a dataset is generated, the best fit parameters $\hat{a}, \hat{b}, \hat{c}$ are obtained by fitting first with a straight line and then with a quadratic function, and the resulting β is calculated. This gives the distribution of β assuming A is correct. Then we repeat generating many datasets around model B (with the same prior and data covariance matrix), and obtain the

distribution of β when B is assumed correct. The parameter c is a measure of how different the two models are, and therefore of how much the two distributions overlap. Putting as means $a = b = 1$, we find the distributions as in Figs. 1 (with $c = 1$) and 2 (with $c = 0.6$).

We now assume that the data give a particular value of β , say $\beta = \beta_i$. As illustrated in the figures, there are four possibilities: A is rejected (β_1), B is rejected (β_3), both models are rejected (β_2), neither model is rejected (β_0 in Fig. 2). In this last case, Fig. 2, if the data give $\beta = 0$, i.e. equal evidences for the two models, the Bayesian test

would be inconclusive, while here we see that A is rejected but B is not. In terms of p -values, we obtain $p \sim 10^{-7}$ for β_1 assuming A is true; we obtain $p \sim 10^{-6}$ for β_3 assuming B is true; and we obtain $p = 0.00063(0.0017)$ for β_2 assuming A (respectively B). Notice that the p -values are obtained integrating the distribution in the interval $(-\infty, \beta_i)$ when assuming A, and in the interval (β_i, ∞) when assuming B.

In Figs. 3 and 4 we illustrate the model comparison test for the same settings as the previous figures. Clearly, β_1 favors B, β_3 favors A, while β_0, β_2 are noncommittal. If we




FIG. 1. Distribution of β assuming model A, a straight line with two parameters (right) or model B, a quadratic polynomial with three parameters (left), with $c = 1$. The prior has been taken to be very broad. The histograms have been obtained by a numerical simulation. Here and in the subsequent similar plots, the red and green curves are the χ^2 analytical results. If an experiment finds a value β_1 , A is rejected; if it finds β_2 , both A and B are rejected; if it finds β_3 , B is rejected.




FIG. 2. Same as the previous plot, but for two models that are closer ($c = 0.6$). If an experiment finds a value β as marked by the dashed vertical line, neither model can be rejected to a high significance.




FIG. 3. For the same settings as in Fig. 1, here we plot R_f (theoretical distribution: red line; numerical distribution: green line) and R_b (blue line) versus β .




FIG. 4. For the same settings as in Fig. 2, here R_f (theoretical distribution: red line; numerical distribution: green line) and R_b (blue line) are shown versus β .

had obtained the value $\beta = 0$ for the case of Fig. 3, then the Bayesian model comparison would have been inconclusive, while in our FB test it would have been quite in favor of B because essentially incompatible with the data, as can be seen in Fig. 2.

Plots (1–4) represent therefore a full analysis of the models. The p -value test gives the absolute quality of the fits, while the probability ratio R_f expresses the relative advantage of a model over another one, regardless of whether they are a good fit or not.

Further, in Fig. 5 we show how the noncentral χ^2 distribution reproduces the numerical distribution in the same case as Fig. 2 but now with a tight prior ($\sigma_P = 1$). Since the prior is strong and centered on model A parameters, the range of β that favors model B is now much reduced with respect to Fig. 2. Finally, in Fig. 6, we modified model B to a cubic function ($f = a + bx + cx^2 + dx^3$), and the corresponding numerical distribution was generated as described earlier. In this case, however, the analytical distribution takes the shape of a generalized χ^2 .

VII. APPLICATIONS: SUPERNOVAE

In this section, we apply the FB approach to real supernovae data. We have used supernovae data from the Union2.0 catalog [44], which consists of a total of 557 supernovae observations. The models we are comparing are Λ CDM and w CDM.

The absolute magnitude of a supernova can be related to the luminosity distance using the equation

$$\mu = 5 \log_{10}(d_L) + 25 = m - M \quad (66)$$

Here, apparent magnitudes (M) and absolute magnitudes (m) are the observables for a supernova, and the luminosity




FIG. 5. Same as Fig. 2, but for a strong prior $\sigma_P = 1$.




FIG. 6. Distribution of β with modified model B, a cubic function, with $c = 0.04$ and $d = 0.06$. The prior used here is $\sigma_p = 0.5$. In this case, the distribution is a generalized χ^2 .

distance d_L is measured in Mpc. The luminosity distance d_L can be given as (we assume flat space and adopt units such that $c = 1$)

$$d_L(z) = \frac{(1+z)}{H_0} \int_0^z \frac{dz'}{E(z')}, \quad (67)$$

where $E^2(z) = \Omega_m(1+z)^3 + \Omega_\Lambda$ for Λ CDM and $E^2(z) = \Omega_m(1+z)^3 + \Omega_{de}(1+z)^{3(1+w)}$ for w CDM and where H_0 is the Hubble constant, w the dark energy equation of state, and $\Omega_m, \Omega_\Lambda, \Omega_{de}$ the present matter, cosmological constant, and dark energy density parameters, respectively.

Now, we can formulate the likelihood Eq. (3) for our analysis as

$$\mathcal{L} = M \exp \left[-\frac{1}{2} (m_{\text{data}} - 5 \log_{10} d_L^* + \alpha)_i \Sigma_{ij}^{-1} (m_{\text{data}} - 5 \log_{10} d_L^* + \alpha)_j \right] \quad (68)$$

where $d_L^* = H_0 d_L$, the indexes i, j run over the data points, and $\alpha = 5 \log_{10} H_0 + 25 - M$. We can marginalize over α and the new likelihood can be written as

$$\mathcal{L} = M^* \exp \left[-\frac{1}{2} (m_{\text{data}} - 5 \log_{10} d_L^*)_i K_{ij}^{-1} (m_{\text{data}} - 5 \log_{10} d_L^*)_j \right] \quad (69)$$

where

$$K_{ij}^{-1} \equiv \Sigma_{ij}^{-1} - \frac{u_p \Sigma_{pi}^{-1} u_n \Sigma_{nj}^{-1}}{S_0} \quad (70)$$

with

$$u_i = \{1, 1, 1, \dots\} \quad (71)$$

$$S_0 = u_i \Sigma_{ij}^{-1} u_j = \sum_{ij} \Sigma_{ij}^{-1} \quad (72)$$

Note that the matrix K^{-1} is singular, so one cannot invert it to obtain K . This causes no problem, however, since we only need K^{-1} . Bayes' ratio was computed for the data using this likelihood and two different priors: a strong prior case, in which the variances were set at $\sigma_{\Omega_m}^2 = 0.005$ and $\sigma_w^2 = 0.08$, and a weak prior case, with $\sigma_{\Omega_m}^2 = 0.5$ and $\sigma_w^2 = 8$. We obtained a value of $\beta_{\text{data}} = 0.79$ for strong priors and $\beta_{\text{data}} = 4.30$ for weak priors (see Figs. 7 and 8). In standard Bayesian analysis, a β_{data} value of 0.79 for strong priors would be considered “barely worth mentioning” evidence in favor of Λ CDM, while a value of 4.30 for weak priors would be deemed “substantial” evidence. Our method aligns with the standard test in the case of strong priors. However, we contend that in the case of weak priors, both models should be regarded on a similar footing, contrary to what the standard test suggests.

For the β distribution, first, a model was supposed to be true; let us say Λ CDM. Subsequently, a synthetic dataset was created by selecting random μ values from a multivariate Gaussian distribution assuming Λ CDM. The mean of this distribution was determined using Eq. (67), and the covariance matrix was taken from the Union2.0 catalogue. From Eq. (69), a likelihood and, consequently, β are computed for this simulated dataset. This procedure was repeated 100000 times, resulting in a distribution of β values denoted as $P(\beta)$. A similar procedure is then replicated, assuming the w CDM model to be true.




FIG. 7. Distribution of β for supernovae Ia. For this simulation, a strong prior configuration was employed with $\sigma_{\Omega_m}^2 = 0.005$ and $\sigma_w^2 = 0.08$.




FIG. 8. Similar to Fig. 7, but for weak priors: $\sigma_{\Omega_m}^2 = 0.5$ and $\sigma_w^2 = 8$.

Our analytical approach works for linear models, but the parameters enter the models Λ CDM and w CDM in a nonlinear way. Linearization of our models can be done by approximating them by first-order Taylor expansion at the parameter best-fit values Eq. (55),

$$d_L^*(z; \theta_i) \approx d_L^*(z)|_{\theta=\hat{\theta}} + \sum_i \left. \frac{\partial d_L^*}{\partial \theta_i} \right|_{\theta_i=\hat{\theta}_i} (\theta_i - \hat{\theta}_i) \quad (73)$$

where θ_i are (Ω_m) and (Ω_m, w) for Λ CDM and w CDM, respectively. This is a good approximation: the relative error between $d_{L,\text{exact}}^*$ [Eq. (67)] and $d_{L,\text{approx}}^*$ [Eq. (73)] for both Λ CDM and w CDM models is less 1% for $z \leq 2.5$.

In this scenario, where we have two linear models with the number of parameters $N_{\Lambda\text{CDM}} = 1$ and $N_{w\text{CDM}} = 2$, the analytical approximation of the β distribution can be represented by a noncentral χ^2 distribution with $\nu = N_{w\text{CDM}} - N_{\Lambda\text{CDM}} = 1$ degrees of freedom and noncentrality parameter λ . This analytical distribution is shown in Figs. 7 and 8.

The best-fit value of the additional parameter in the w CDM model ($\hat{w} = -1.15$) is very close to the value for the Λ CDM model ($w = -1$), making both distributions very similar to each other. This close alignment between the distributions makes it challenging to conclusively comment on the accuracy of either cosmological model at present. However, we anticipate that future supernovae observations, with reduced observational uncertainties, will enable a more distinct separation between two distributions.




FIG. 9. Similar to Fig. 7, but now with the supernovae covariance reduced by a factor of 20 (i.e., $C_{ij}/20$). The two distributions are now quite different, and it would be possible to clearly differentiate the two competing models.




FIG. 10. Dependence of β on different priors. The Gaussian variance of priors for each β is taken as $\sigma_{\Omega_m}^2 = n \times 0.005$ and $\sigma_w^2 = n \times 0.08$, where n = prior breadth.

As a purely illustrative example, if the SN magnitude uncertainties are decreased by a factor of 20, the new $\beta_{\text{data}} = -2.98$ would imply that $w\text{CDM}$ is more likely to be the true model of the Universe than ΛCDM (see Fig. 9). Finally, Fig. 10 highlights how the value of β depends on the prior width. Relying on the Jeffreys scale in such cases might result in inaccurate predictions for the correct cosmological model.

VIII. CONCLUSIONS

The decisive role of statistics in testing physical models cannot be overestimated, especially in view of the very

abundant data that are expected to be generated by current and future experiments. Having a number of alternative and complementary tools at our disposal to compare theoretical models is therefore a pressing issue. In this paper we argued that some shortcomings of standard Bayesian analysis can be overcome in the framework of a mixed frequentist/Bayesian approach, denoted as FB method, in which Bayes' ratio is employed as a frequentist statistic (see also [37] for a similar point of view). In particular, we have emphasized the dependence of Bayes' ratio on the prior, even in the limit of weak priors, and have shown that this issue is solved in our mixed approach.

The frequentist distribution of Bayes' factor can be employed both as a quality-of-fit test and as model comparison, thereby producing an exhaustive overview of competing models. We also found the analytical expression of such distribution in the case of correlated Gaussian data with linear nested models and arbitrary Gaussian priors, and shown that it provides a good approximation also in a realistic setting of interest to cosmology. In such case, the FB method can be carried out without the need of data simulations.

In the case of non-Gaussian posteriors, the analytical results presented in this work are generally not applicable. Recent advances in simulation-based inference and emulators may, however, warrant an extension of the present method beyond the Gaussian approximation while still benefiting from its insight and fast applicability. Such an extension could involve building an emulator for the non-Gaussian correction to the FB distribution that learns from a (potentially small) number of simulations for which the Bayes factor is known exactly (e.g., computed using nested samplers such as PolyChord [45] or Nautilus [46], or simulation-based methods, such as the one proposed in [47].)

ACKNOWLEDGMENTS

L. A. thanks Juan Garcia-Bellido, Valerio Marra, Savvas Nesseris, and Björn M. Schäfer for useful discussions on these topics. Z. S. and L. A. acknowledge support from DFG project 456622116. L. A. acknowledges support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2181/1–390900948 (the Heidelberg STRUCTURES Excellence Cluster). K. W. acknowledges support from the Science and Technology Facilities Council (STFC) under Grant No. ST/X006344/1. More detailed calculation of several expressions can be found in the earlier versions of the arXiv submission.

APPENDIX: A SIMPLE EXAMPLE

Figure 11 shows an extremely simple example of how misleading weak priors can be. Here, model A is a




FIG. 11. This plot appear to show that model B is a better model than A in explaining the data. However, if the extra B parameter has an extremely weak prior, Bayes' factor b would strongly favor model A.

horizontal straight line $y = a$ with just a as free parameter, while model B is a general straight line $y = a + bx$ with two free parameters. The data points have been created by sampling a Gaussian distribution around the line $y = 1 + 0.6x$, with uncertainty $\sqrt{1/2}$. The blue dashed line is the best fit for model A, while the continuous orange line is the best fit for B. The Gaussian prior standard deviation for the common parameter a is 10, while for the extra parameter b is huge, 10^6 . Clearly, model B appears strongly favored by this dataset. However, the very weak prior shifts β to roughly 9.8, which corresponds to a Bayes' ratio ≈ 130 , well into the “decisive” range in favor of model A of the Jeffreys scale. So a pure Bayesian analysis of this simple case would produce the puzzling result that model A is to be much preferred over B. Bayes' factor becomes smaller than unity (i.e., model B is favored) only if the prior uncertainty on b is smaller than 7000.

-
- [1] Bikash R. Dinda, Haveesh Singirikonda, and Subhabrata Majumdar, Constraints on cosmic curvature from cosmic chronometer and quasar observations, [arXiv:2303.15401](#).
 - [2] Alexander Reeves, Laura Herold, Sunny Vagnozzi, Blake D. Sherwin, and Elisa G. M. Ferreira, Restoring cosmological concordance with early dark energy and massive neutrinos?, *Mon. Not. R. Astron. Soc.* **520**, 3688 (2023).
 - [3] David Alonso, Giulio Fabbian, Kate Storey-Fisher, Anna-Christina Eilers, Carlos García-García, David W. Hogg, and Hans-Walter Rix, Constraining cosmology with the Gaia-unWISE quasar catalog and CMB lensing: Structure growth, *J. Cosmol. Astropart. Phys.* **11** (2023) 043.
 - [4] Evan McDonough, J. Colin Hill, Mikhail M. Ivanov, Adrien La Posta, and Michael W. Toomey, Observational constraints on early dark energy, *Int. J. Mod. Phys. D* **33**, 2430003 (2024).
 - [5] Raphaël Kou, Calum Murray, and James G. Bartlett, Constraining $f(R)$ gravity with cross-correlation of galaxies and cosmic microwave background lensing, *Astron. Astrophys.* **686**, A193 (2024).

- [6] Jing Yang, Xin-Yan Fan, Chao-Jun Feng, and Xiang-Hua Zhai, Latest data constraint of some parameterized dark energy models, *Chin. Phys. Lett.* **40**, 019801 (2023).
- [7] Emil Brinch Holm, Laura Herold, Théo Simon, Elisa G. M. Ferreira, Steen Hannestad, Vivian Poulin, and Thomas Tram, Bayesian and frequentist investigation of prior effects in EFT of LSS analyses of full-shape BOSS and eBOSS data, *Phys. Rev. D* **108**, 123514 (2023).
- [8] Ziad Sakr, Ana Carvalho, Antonio Da Silva, Juan Garcia Bellido, Jose P. Mimoso, David Camarena, Savvas Nesseris, Carlos J. A. P. Martins, Nelson J. Nunes, and Domenico Sapone, Constraining LLTB models with galaxy cluster counts from next generation surveys, *Astron. Astrophys.* **683**, A230 (2024).
- [9] Eoin Ó. Colgáin, Saeed Pourojaghi, M. M. Sheikh-Jabbari, and Darragh Sherwin, MCMC marginalisation bias and Λ CDM tensions, [arXiv:2307.16349](https://arxiv.org/abs/2307.16349).
- [10] Ziad Sakr, Stephane Ilic, and Alain Blanchard, Cluster counts—III. Λ CDM extensions and the cluster tension, *Astron. Astrophys.* **666**, A34 (2022).
- [11] P. Serra, A. Heavens, and A. Melchiorri, Bayesian evidence for a cosmological constant using new high-redshift supernova data, *Mon. Not. R. Astron. Soc.* **379**, 169 (2007).
- [12] Aleksandra Kurek and Marek Szydlowski, The Λ CDM model in the lead—A Bayesian cosmological model comparison, *Astrophys. J.* **675**, 1 (2008).
- [13] Moncy V. John and J. V. Narlikar, Comparison of cosmological models using Bayesian theory, *Phys. Rev. D* **65**, 043506 (2002).
- [14] Daniel L. Shafer, Robust model comparison disfavors power law cosmology, *Phys. Rev. D* **91**, 103516 (2015).
- [15] Alan Heavens, Yabebal Fantaye, Elena Sellentin, Hans Eggers, Zafirah Hosenie, Steve Kroon, and Arrykrishna Mootoovaloo, No evidence for extensions to the standard cosmological model, *Phys. Rev. Lett.* **119**, 101301 (2017).
- [16] H. Jeffreys, *Theory of Probability*, 3rd ed., Oxford Classics Series (reprinted 1998) (Oxford University Press, Oxford, UK, 1961).
- [17] Christian P. Robert, Nicolas Chopin, and Judith Rousseau, Harold Jeffreys's theory of probability revisited, [arXiv:0804.3173](https://arxiv.org/abs/0804.3173).
- [18] Robert E. Kass and Adrian E. Raftery, Bayes factors, *J. Am. Stat. Assoc.* **90**, 773 (1995).
- [19] G. Efstathiou, Limitations of Bayesian evidence applied to cosmology, *Mon. Not. R. Astron. Soc.* **388**, 1314 (2008).
- [20] Savvas Nesseris and Juan Garcia-Bellido, Is the Jeffreys' scale a reliable tool for Bayesian model comparison in cosmology?, *J. Cosmol. Astropart. Phys.* **08** (2013) 036.
- [21] C. R. Jenkins and J. A. Peacock, The power of Bayesian evidence in astronomy, *Mon. Not. R. Astron. Soc.* **413**, 2895 (2011).
- [22] Mehdi Rezaei and Mohammad Malekjani, Comparison between different methods of model selection in cosmology, *Eur. Phys. J. Plus* **136**, 219 (2021).
- [23] Jennifer A. Hoeting, David Madigan, Adrian E. Raftery, and Chris T. Volinsky, Bayesian model averaging: A tutorial (with comments by M. Clyde, David Draper and E. I. George, and a rejoinder by the authors), *Stat. Sci.* **14**, 382 (1999).
- [24] David Parkinson and Andrew R. Liddle, Bayesian model averaging in astrophysics: A review, *Stat. Anal. Data Mining* **6**, 3 (2013).
- [25] Martin Kunz, Roberto Trotta, and David Parkinson, Measuring the effective complexity of cosmological models, *Phys. Rev. D* **74**, 023503 (2006).
- [26] Marco Raveri and Wayne Hu, Concordance and discordance in cosmology, *Phys. Rev. D* **99**, 043506 (2019).
- [27] Alan F. Heavens, Arrykrishna Mootoovaloo, Roberto Trotta, and Elena Sellentin, Extreme data compression for Bayesian model comparison, *J. Cosmol. Astropart. Phys.* **11** (2023) 048.
- [28] Hanwool Koo, Ryan E. Keeley, Arman Shafieloo, and Benjamin L'Huillier, Bayesian vs frequentist: Comparing Bayesian model selection with a frequentist approach using the iterative smoothing method, *J. Cosmol. Astropart. Phys.* **03** (2022) 047.
- [29] I. J. Good, Saddle-point methods for the multinomial distribution, *Ann. Math. Stat.* **28**, 861 (1957).
- [30] I. J. Good, The Bayes/non-Bayes compromise: A brief review, *J. Am. Stat. Assoc.* **87**, 597 (1992).
- [31] Luca Amendola, Valerio Marra, and Miguel Quartin, Internal robustness: Systematic search for systematic bias in SN Ia data, *Mon. Not. R. Astron. Soc.* **430**, 1867 (2013).
- [32] Kevin Wolz, Bayesian and frequentist methods to investigate tensions between data sets, M.Sc. thesis, University of Heidelberg, Heidelberg, Germany, 2018.
- [33] Lachlan Passenger, Eric Thrane, Paul D. Lasky, Ethan Payne, Simon Stevenson, and Ben Farr, Are all models wrong? Falsifying binary formation models in gravitational-wave astronomy, *Mon. Not. R. Astron. Soc.* **535**, 2837 (2024).
- [34] B. Joachimi, F. Köhlinger, W. Handley, and P. Lemos, When tension is just a fluctuation: How noisy data affect model comparison, *Astron. Astrophys.* **647**, L5 (2021).
- [35] T. Castro and M. Quartin, First measurement of σ_8 using supernova magnitudes only, *Mon. Not. R. Astron. Soc.* **443**, L6 (2014).
- [36] Andrew Fowlie, Neyman-Pearson lemma for Bayes factors, *Commun. Phys.* **52**, 5379 (2023).
- [37] Ryan E. Keeley and Arman Shafieloo, On the distribution of Bayesian evidence, *Mon. Not. R. Astron. Soc.* **515**, 293 (2022).
- [38] Roberto Trotta, Bayes in the sky: Bayesian inference and model selection in cosmology, *Contemp. Phys.* **49**, 71 (2008).
- [39] George Lazarides, Roberto Ruiz de Austri, and Roberto Trotta, Constraints on a mixed inflaton and curvaton scenario for the generation of the curvature perturbation, *Phys. Rev. D* **70**, 123527 (2004).
- [40] Alan F. Heavens, T. D. Kitching, and L. Verde, On model selection forecasting, dark energy and modified gravity, *Mon. Not. R. Astron. Soc.* **380**, 1029 (2007).
- [41] Alan F. Heavens and Elena Sellentin, Objective Bayesian analysis of neutrino masses and hierarchy, *J. Cosmol. Astropart. Phys.* **04** (2018) 047.
- [42] J. P. Imhof, Computing the distribution of quadratic forms in normal variables, *Biometrika* **48**, 419 (1961).
- [43] Pierre Duchesne and Pierre Lafaye De Micheaux, Computing the distribution of quadratic forms: Further comparisons between the Liu-Tang-Zhang approximation and exact methods, *Computational Statistics and Data Analysis* **54**, 858 (2010).

-
- [44] Amanullah *et al.*, Spectra and Hubble Space Telescope light curves of six type Ia supernovae at $0.511 < z < 1.12$ and the union2 compilation, *Astrophys. J.* **716**, 712 (2010).
- [45] W.J. Handley, M.P. Hobson, and A.N. Lasenby, PolyChord: Next-generation nested sampling, *Mon. Not. R. Astron. Soc.* **453**, 4384 (2015).
- [46] Johannes U. Lange, Nautilus: Boosting Bayesian importance nested sampling with deep learning, *Mon. Not. R. Astron. Soc.* **525**, 3181 (2023).
- [47] A. Spurio Mancini, M.M. Docherty, M.A. Price, and J.D. McEwen, Bayesian model comparison for simulation-based inference, *RAS Tech. Instrum.* **2**, 710 (2023).