

OPEN ACCESS

Objective Bayesian analysis of neutrino masses and hierarchy

To cite this article: Alan F. Heavens and Elena Sellentin JCAP04(2018)047

View the [article online](#) for updates and enhancements.

Related content

- [Strong Bayesian evidence for the normal neutrino hierarchy](#)
Fergus Simpson, Raul Jimenez, Carlos Pena-Garay et al.
- [Can we measure the neutrino mass hierarchy in the sky?](#)
Raul Jimenez, Thomas Kitching, Carlos Peña-Garay et al.
- [A tale of two modes: neutrino free-streaming in the early universe](#)
Lachlan Lancaster, Francis-Yan Cyr-Racine, Lloyd Knox et al.

Objective Bayesian analysis of neutrino masses and hierarchy

Alan F. Heavens^a and Elena Sellentin^b

^aImperial Centre for Inference and Cosmology (ICIC), Imperial College, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, U.K.

^bDépartement de Physique Théorique, Université de Genève, 24 Quai Ernest-Ansermet, CH-1211 Genève, Switzerland

E-mail: a.heavens@imperial.ac.uk, elena.sellentin@unige.ch

Received March 4, 2018

Accepted April 5, 2018

Published April 18, 2018

Abstract. Given the precision of current neutrino data, priors still impact noticeably the constraints on neutrino masses and their hierarchy. To avoid our understanding of neutrinos being driven by prior assumptions, we construct a prior that is mathematically minimally informative. Using the constructed uninformative prior, we find that the normal hierarchy is favoured but with inconclusive posterior odds of 5.1:1. Better data is hence needed before the neutrino masses and their hierarchy can be well constrained. We find that the next decade of cosmological data should provide conclusive evidence if the normal hierarchy with negligible minimum mass is correct, and if the uncertainty in the sum of neutrino masses drops below 0.025 eV. On the other hand, if neutrinos obey the inverted hierarchy, achieving strong evidence will be difficult with the same uncertainties. Our uninformative prior was constructed from principles of the Objective Bayesian approach. The prior is called a reference prior and is minimally informative in the specific sense that the information gain after collection of data is maximised. The prior is computed for the combination of neutrino oscillation data and cosmological data and still applies if the data improve.

Keywords: cosmological neutrinos, neutrino masses from cosmology, solar and atmospheric neutrinos

ArXiv ePrint: [1802.09450](https://arxiv.org/abs/1802.09450)



Contents

1	Introduction	1
2	The meaning of priors in Bayesian inferences	2
3	Least-informative priors for a given experiment	3
3.1	Parametrisation of neutrino experiments	5
4	Construction and interpretation of the neutrino mass prior	7
5	Results: neutrino masses and hierarchy	9
5.1	Mass posteriors	9
5.2	Marginal likelihood for the hierarchies	9
6	Future neutrino constraints	11
7	Conclusions	11

1 Introduction

The precise masses m_1, m_2 and m_3 of neutrinos and the ordering of these masses are currently poorly determined features of particle physics. Solar, atmospheric, accelerator and reactor neutrino experiments are able to determine mass-squared differences, and the so-called Mikheyev-Smirnov-Wolfenstein matter effect determines the ordering of two of the mass eigenstates, see e.g. [25]. However, the absolute masses are currently relatively poorly constrained without recourse to cosmological data, giving various 90% credible upper limits for m_1 in the range $0.18 - 0.48$ eV, depending on the datasets used [9, 22]. This leaves open both possibilities of m_3 being larger than m_1 and m_2 (the normal hierarchy, which we abbreviate by ‘NH’), and smaller than m_1 and m_2 (the inverted hierarchy; ‘IH’).

Current particle physics experiments are not able to make strong statements in favour of one ordering or the other [11, 13, 28]. Cosmological observations add important information, since massive neutrinos alter, in a scale-dependent way, the growth of fluctuations. The principal sensitivity arises because the primordial fluctuations are erased on small scales by neutrino free-streaming, largely while the neutrinos are still relativistic. This leads to a suppression of power at all but the lowest wavenumbers ($k < 10^{-2} h\text{Mpc}^{-1}$) in the cosmic structures. In principle, such a suppression could be achieved by other sufficiently abundant relativistic free-streaming particles, but neutrinos are the particles preferred by the data [30]. Since the proportion of the total fluctuation carried by the neutrinos increases as their total mass increases, the observed suppression of power allows us to infer model-dependent bounds on the sum of neutrino masses.

Current data put an upper limit on the sum of the masses that is getting interestingly close to the minimum mass sums that are allowed by neutrino oscillations (0.06 and 0.1 eV for the NH and IH respectively). For example, [27] find an upper limit of 0.12 eV at 95% confidence after having marginalised over other cosmological parameters from BOSS $\text{Ly}\alpha$ clustering and the cosmic microwave background. This result is similar to other studies using different probes and redshift ranges, such as those of [10, 17].

Previous determinations of the relative odds of the normal to the inverted hierarchy usually moderately favour the NH, depending on the dataset used. However, the data are not so constraining, with the effect that the prior chosen can still have a considerable influence on the results, and this is the focus of this paper.

A frequentist view was taken by [8], but most analyses are Bayesian (e.g. [7, 16, 19, 26, 32, 36]), so priors are necessary. In summary, [19] and [16] found an odds ratio in favour of the normal hierarchy of 1.5 with a uniform prior for the lightest neutrino mass. [36] found odds of 1.8 and 3.3 depending on whether conservative or aggressive mass sum limits were chosen. Such limits are again a form of imposing a prior. The shape [36] chosen for their prior was uniform over the sum of neutrino masses. In contrast to these approaches, [32] used the logarithm of the mass as a parameter to specify the prior, and advocated a more complex hierarchical model, with normal priors in the logarithm of the masses, with the mean and variance of the normal being hyper-parameters. As a consequence, [32] found much larger odds favouring the normal hierarchy by 42:1, and [15] similarly found that a logarithmic prior gave strong odds for the NH, and weak preference for the NH with linear priors. An even higher odds ratio was found by [26] where priors were set on the mass matrix itself.

Clearly, the quoted odds ratios still depend on the priors chosen. This indicates that the data are not powerful enough to override the choice of prior. Answering the physical question of which hierarchy is preferred, still depends strongly on the information contributed by the chosen prior. We regard it however as important that physical conclusions depend on prior assumptions as little as possible. In this paper, we shall therefore explicitly construct a prior that *minimizes* the information that the prior contributes to the analysis. We shall see that such an uninformative prior indeed leads to an uninformative odds ratio, meaning better data are needed to infer the neutrino mass hierarchy.

2 The meaning of priors in Bayesian inferences

In Bayesian analysis, there are broadly two schools of thought concerning appropriate priors. In the subjective Bayesian approach, which has been taken to date in Bayesian neutrino mass analysis, one chooses the prior and states it, since any inferences will be dependent on the choice. Without previous data the aim is often to try to choose a ‘non-informative’ prior, such that the posterior is driven by the likelihood and not the prior. One difficulty can be that what constitutes a non-informative prior may not be obvious, and apparently innocuous choices such as uniform priors can be extremely informative in some respects. The alternative school of thought, called Objective Bayesianism, constructs a prior that is, in a mathematically well-defined sense, as uninformative as possible, on information-theoretic grounds. This includes arguments based on the Kullback-Leibler divergence, Haar measures, or the maximum entropy principle [2, 12, 14, 20].

The concept which applies to the situation at hand are so-called reference priors, introduced by [5]. A reference prior ensures that when the probability distribution of model parameters is updated from the prior to the posterior, the information gain is maximised. The prior then has an obvious merit in that it is determined by a mathematical procedure, rather than a subjective view, and it maximises the influence of the data on the inference. One complication is that, in multi-parameter problems, the reference prior can depend on the ordering of parameters [6], so may not yield a unique maximally non-informative prior. However, a very nice feature of the combination of neutrino experiments used in this paper is that the ordering of the parameters is irrelevant, and there is, in this case, a unique reference

prior. This prior is one of the main results of this paper. The approach is outlined in the next section; for a full discussion of reference priors the reader is directed to [3, 6]. The other main results are the ensuing posteriors on the neutrino masses, and the resultant posterior odds of the two possible neutrino hierarchies (where with current data none is strongly preferred).

3 Least-informative priors for a given experiment

To make precise the notion of a ‘least-informative prior’, consider a model M with parameters θ , which can generate data \mathbf{x} . The theory of uninformative priors has progressed in the past thirty years, and the concept which applies to the problem of inferring the neutrino mass hierarchy is the theory of ‘reference priors’. What is nowadays understood as a reference prior is a prior that maximises the expected information gain caused by the arrival of the data \mathbf{x} . In other words, a reference prior maximizes the missing information, in the absence of data. The concept of missing information is thereby central: it defines what is mathematically meant when we request an ‘uninformative’ prior.

For a reference prior, the missing information is assessed as follows. Let there be an experiment, which aims to constrain parameters θ . Being Bayesian, the experiment is only conducted once, and procures a data set $\mathbf{x} \sim p(\mathbf{x}|\theta)$. The likelihood then quantifies how tightly a single outcome of the experiment constrains the parameters θ . In contrast, the missing information before knowing the parameter precisely depends additionally on the prior: if the prior is sharply peaked, it contributes a lot of information. To quantify the missing information, reference prior theory uses k hypothetical repetitions of the experiment, as these would finally overwrite any sensible prior choice. For $k \rightarrow \infty$ repetitions, the parameter is then as precisely constrained as possible by the *data*, conditional on the assumed experimental setup. Note that the experiment does not need to be repeated in reality — the hypothetical repetition is only necessary to quantify the missing information. The reference prior is then the prior which maximizes this missing information.

Denoting the data generated by k repetitions of the experiment as \mathbf{x}^k , the mutual information between a prior $\pi(\theta)$ and posterior $p(\theta|\mathbf{x}^k)$ is [3]

$$I(\theta, \mathbf{x}^k) = \int p(\mathbf{x}^k) \left[\int p(\theta|\mathbf{x}^k) \ln \frac{p(\theta|\mathbf{x}^k)}{\pi(\theta)} d\theta \right] d\mathbf{x}^k, \quad (3.1)$$

where we recognise the inner integral as the Kullback-Leibler divergence between the posterior $p(\theta|\mathbf{x}^k)$ and the prior $\pi(\theta)$. The outer integral runs over realizations of the data, meaning the information here defined is an *expected* information. Maximising the missing information yields the on average least-informative prior $\pi(\theta)$.

The solution for the least-informative prior $\pi(\theta)$ amongst a class of candidate priors given the information measure eq. (3.1), is [3, 6]

$$\pi(\theta) = \lim_{k \rightarrow \infty} \frac{f_k(\theta)}{f_k(\theta_0)}, \quad (3.2)$$

$$f_k(\theta) = \exp \left[\int p(\mathbf{x}^k|\theta) \ln p^*(\theta|\mathbf{x}^k) d\mathbf{x}^k \right]. \quad (3.3)$$

where θ_0 is a point in the parameter space, and $p^*(\theta|\mathbf{x}^k)$ is the posterior with a fiducial choice of a prior $\pi^*(\theta)$, and the limit is loose in that it permits improper priors. The choice of the fiducial point θ_0 and the fiducial prior π^* can facilitate the computation, but (given the

existence of the resulting reference prior and certain regularity conditions) do not influence the outcome. Typically, algorithmic sampling procedures are employed to solve eq. (3.3). See [3] for regularity conditions and proofs.

However, of particular convenience is that for likelihoods which already achieved asymptotic normality (roughly: a precision experiment), the reference priors for single parameters converge to a prior class which can be more easily computed, namely the so-called Jeffreys priors. Certain likelihoods may of course never achieve asymptotic normality. This is for example possible in multi-parameter problems in cases where the data cannot lift parameter degeneracies, see e.g. [29, 31]. The reference prior which is least-informative according to (3.1) does then not coincide with the Jeffreys prior.

The Jeffreys prior in one dimension is defined as

$$\pi_J(\theta) \propto \sqrt{-\mathbb{E} \left[\frac{d^2 \ln p(\mathbf{x}|\theta)}{d\theta^2} \right]}, \quad (3.4)$$

where the expectation operator \mathbb{E} is again the integral over the data, see e.g. [6, 21]. Note that the statistical definition of the Jeffreys prior is more general and also a deeper concept than what is usually called a ‘Jeffreys prior’ in astronomy (where the astronomical ‘Jeffreys prior’ is used to denote the prior $\pi \propto 1/\theta$). The general Jeffreys prior hence uses as information measure the Fisher information

$$F = -\mathbb{E} \left[\frac{d^2 \ln p(\mathbf{x}|\theta)}{d\theta^2} \right], \quad (3.5)$$

rather than the mutual information of eq. (3.1).

For multi-parameter models with a likelihood function $p(\mathbf{x}|\boldsymbol{\theta})$, the reference prior is found sequentially, by first fixing all but one parameter and finding the reference prior $\pi(\theta_1|\boldsymbol{\theta}_{i \neq 1})$ for the remaining parameter θ_1 by treating it as a single parameter problem. Assuming the resulting prior is proper, then θ_1 may be integrated out, and the likelihood is

$$p(\mathbf{x}|\boldsymbol{\theta}_{i \neq 1}) = \int p(\mathbf{x}|\boldsymbol{\theta})\pi(\theta_1)d\theta_1. \quad (3.6)$$

The priors are then obtained sequentially by repeating this procedure, and the final prior is $\pi(\boldsymbol{\theta}) = \pi(\theta_n)\pi(\theta_{n-1}|\theta_n)\dots\pi(\theta_1|\boldsymbol{\theta}_{i \neq 1})$. Note that in some cases the prior obtained this way depends on the ordering of the parameters, and it is then conventional to order by the perceived importance of the parameters. However, if the likelihood is separable into subsets of the data \mathbf{x} , i.e. $\prod_i p(\mathbf{x}_i|\boldsymbol{\theta}_i)$, as in this case, then $\prod_{i \neq 1} p(\mathbf{x}_i|\boldsymbol{\theta}_i)$ may be taken outside the integral in eq. (3.6), and repeated application of this principle shows that the reference prior is the product of the individual 1D reference priors, and is unique.

A reference prior constructed either via (3.3), or if appropriate via (3.4), has multiple attractive features. Firstly, it minimizes on average the information that the prior contributes to the measurement. Secondly, reference priors satisfy multiple sanity criteria. For example, reference priors are preserved when compressing the data onto sufficient summary statistics. Furthermore, reference priors are preserved when conducting independent repetitions of a measurement. This means reference priors do not tip from uninformative to informative with increasing sample size, as many other priors do which are a function of the number of data points [33]. Reference priors are also consistent under invertible reparametrizations and they furthermore hold even if the initially accessible parameter space is retrospectively

constrained to a compact subset [3, 6]. The latter property ensures that reference priors consistently handle unphysical regions.¹

Reference priors and Jeffreys priors are sometimes criticized by pure statisticians for depending on the experimental setup. This dependence ensues since the priors are computed from either infinite repetitions of the data, or from expectation values gained from integrating over the distribution of the data. For a physicist, this dependence on the experimental setup is however natural: different experiments will procure different data sets, and different data sets will overwrite the same prior with different speeds. It is therefore natural that a least-informative prior depends on the measurement design.

In the following sections we construct the reference prior for the investigated set of neutrino experiments, and use it to determine the neutrino mass posteriors, and the odds of the two neutrino hierarchies.

3.1 Parametrisation of neutrino experiments

The mutual information is parameter-independent, which means that the reference prior is also independent of parametrisation. This is convenient, because in order to constrain the neutrino mass hierarchy, it is most natural to parametrise the model directly with the neutrino masses m_1, m_2, m_3 . In contrast, to compute the reference prior, it is simpler to start from the quantities that the experiments are actually sensitive to, namely the mass-squared splittings and the sum of neutrino masses. We therefore order the masses according to $m_H \geq m_M \geq m_L$, where the subscripts H, M, L abbreviate ‘high, medium, low’. The hierarchy is then normal if $m_H^2 - m_M^2 > m_M^2 - m_L^2$, and inverted otherwise.

The quantities that are constrained or measured by experiments are then

$$\begin{aligned}
\phi &\equiv m_M^2 - m_L^2 && \text{Normal} \\
&\equiv m_H^2 - m_M^2 && \text{Inverted} \\
\psi &\equiv m_H^2 - \frac{1}{2}(m_M^2 + m_L^2) && \text{Normal} \\
&\equiv \frac{1}{2}(m_H^2 + m_M^2) - m_L^2 && \text{Inverted} \\
\Sigma &\equiv m_H + m_M + m_L. &&
\end{aligned} \tag{3.7}$$

The parameters ϕ and ψ are the (squared) mass splittings, and the parameter Σ is the sum of neutrino masses. For further discussion, see [26]. These parameters are constrained by solar experiments, a combination of long baseline (LBL) accelerator and reactor experiments, and cosmological neutrino experiments respectively. Using the parametrisation in [26], current constraints are [8, 10, 13, 18, 26, 27]

$$\begin{aligned}
m_M^2 - m_L^2 &= (7.50 \pm 0.18) \times 10^{-5} \text{ eV}^2 \text{ (Normal)} \\
m_H^2 - m_M^2 &= (7.50 \pm 0.18) \times 10^{-5} \text{ eV}^2 \text{ (Inverted)} \\
m_H^2 - (m_M^2 + m_L^2)/2 &= (2.524 \pm 0.04) \times 10^{-3} \text{ eV}^2 \text{ (Normal)} \\
(m_H^2 + m_M^2)/2 - m_L^2 &= (2.514 \pm 0.04) \times 10^{-3} \text{ eV}^2 \text{ (Inverted)} \\
m_L + m_M + m_H &< 0.12 \text{ eV} \quad 95\% \text{ credible region.}
\end{aligned} \tag{3.8}$$

To construct the least informative prior, we first notice that the constraints above are means and their respective standard deviations, which is justified because the posteriors for the

¹This can intuitively be understood by noting that the distribution function in eq. (3.1) will not generate data in unphysical regions, although the actual proof is somewhat more involved [3].

terrestrial experiments are already very close to Gaussian, see also [4, 13]. This means asymptotic posterior normality has already set in. Additionally, ϕ and ψ are well determined, meaning the constraints in eq. (3.8) are insensitive with respect to priors potentially assumed in the original publications (if they were Bayesian analyses). The cosmological probability of the neutrino mass sum is also approximately a Gaussian truncated at zero, but evidently not well measured yet [10]. To compute the reference priors for ϕ, ψ and Σ , we additionally note that the quoted experiments are independent of each other: especially the terrestrial and solar constraints on ϕ and ψ were computed from data which are independent of cosmological data which constrain the sum of neutrino masses. The joint data set of terrestrial, solar and cosmological data therefore factorizes, and hence so does the least informative reference prior. We can therefore compute a total reference prior for the neutrino hierarchy, which is the product of the reference priors for ϕ, ψ, Σ . We will then transform the reference prior for ϕ, ψ, Σ onto the reference prior for m_L, m_M, m_H .

To compute the reference priors for ϕ, ψ, Σ we exploit that asymptotic posterior normality has already set in for these parameters. This means we do not need to solve eq. (3.3) via a sampling algorithm, but can instead take the short-cut and compute the Jeffreys prior $\pi_J(\theta) \propto \sqrt{F(\theta)}$ from eq. (3.4). The Fisher information is given by eq. (3.5). The Gaussian likelihoods are given by

$$L(\hat{\mu}|\theta) = \frac{1}{\sqrt{2\pi\sigma^2(\theta)}} \exp\left(-\frac{1}{2} \frac{[\hat{\mu} - \mu(\theta)]^2}{\sigma^2(\theta)}\right). \quad (3.9)$$

where θ is any of the three parameters ϕ, ψ, Σ and $\hat{\mu}$ is the summary statistic of the data from which the parameters are estimated (here the estimated mean). The Fisher information is then

$$F(\theta) = \frac{[\partial_\theta \mu(\theta)]^2}{\sigma^2(\theta)} + 2 \frac{[\partial_\theta \sigma(\theta)]^2}{\sigma^2(\theta)}. \quad (3.10)$$

How this Fisher information now scales with the parameters determines the shape of the priors. We see that for θ being any of ϕ, ψ, Σ , the likelihoods are already Gaussian, hence the parameters are effectively linear over the domain of relevance and we hence have $\mu(\theta) \propto \theta$ in all three cases. Moreover, we know from the measurement designs that the standard deviations $\sigma(\theta)$ are independent of θ over a wide range. For the cosmological measurements of Σ , this is because the covariance matrices are there taken to be cosmology independent. For the mass splittings ϕ and ψ , the independence of the noise is somewhat harder to see, but arises because of the following reasoning.

For a given neutrino oscillation experiment, the sensitivity S varies periodically with the difference in the squared masses of two neutrino mass eigenstates i, j , according to

$$S \propto \sin^2 \left[\alpha \frac{(m_i^2 - m_j^2)L}{E_\nu} \right] \quad (3.11)$$

where L is the baseline, E_ν the neutrino energy, and α is a constant, containing the speed of light, etc. The sensitivity does however not depend on the absolute masses. We are interested in the prior on ϕ and ψ over the range where the likelihood is high. Ideally L is optimised such that the angle $\alpha\psi L/(4E_\nu) \sim \pi/2$, [35] with the result that the expected error is sensitive to the mass splittings ψ or ϕ only when the shift in mass difference changes the angle by $\mathcal{O}(1)$. This is more than an order of magnitude larger than the width of the likelihood. Therefore,

ϕ and ψ are indeed effectively linear parameters over the width of the likelihood, and the measurement uncertainty is constant over the still admissible range for ϕ and ψ .

We therefore see that given the current measurement precisions, all three parameters ϕ, ψ, Σ can effectively be treated as linear and the noise can be treated as constant. The Fisher information (3.10) is then constant, since the derivative $\partial_\theta \sigma(\theta)$ vanishes, and since the derivative $\partial_\theta \mu(\theta)$ is effectively constant over the domain of high likelihood. Because of the Fisher information being constant, we arrive at uniform Jeffreys priors for all three parameters

$$\pi_J(\phi) = \text{const.}, \quad \pi_J(\psi) = \text{const.}, \quad \pi_J(\Sigma) = \text{const.} \quad (3.12)$$

Let us shortly discuss the approximations in this derivation. These priors were derived using that the parameters are effectively linear, due to the constraining power of the data. In order for this to be not only approximately but rather strictly true, there should be translational invariance of the likelihood also for *less* constraining data. This translational invariance is meant in the sense that the likelihood only shifts as a function of each θ but does not change shape or width. For current data, assuming linear parameters is an excellent approximation, and has been studied especially for the cosmological experiment, where Fisher matrix analyses [1, 34] have shown roughly constant Fisher information as a function of the neutrino masses. For future data, the precision of this approximation will improve further. Note that although cosmological neutrino results are presented truncated at zero mass sum, these should be interpreted as a likelihood multiplied by a hard theoretical prior truncated at zero. The likelihood itself can extend to negative mass sum. For example, an estimator of the mass sum is not required to be positive. We make the assumption that the likelihood itself is translation-invariant, consistent with the Fisher analyses of [1, 34], in which case the Jeffreys prior is constant, regardless of its shape.

In addition to noting that truncating the distribution for the neutrino masses at $\Sigma = 0$ does not affect the least informative reference prior, there are some regions of (ϕ, ψ, Σ) space that are unphysical, since they correspond to negative values of one or more of the masses or mass-squared terms. In conjunction with later imposing an upper limit on the sum of masses, this can again be seen as a restriction of the parameter space. The reference prior handles this, since it can always be retrospectively constrained to a compact subspace. This can be proven, see [3], but in principle arises because independently gained information is additive, and restricting the prior to physically meaningful parameter ranges simply adds physical information to the missing information as assessed by the reference prior.

Having thus determined the reference prior in the (ϕ, ψ, Σ) parametrisation, in the next section we use the invariance of the mutual information to parametrisation (and hence consistency of the reference prior) to find the prior in the (m_L, m_M, m_H) parametrisation. From there, we will compute the posterior mass distributions and the marginal evidence for the two hierarchies.

4 Construction and interpretation of the neutrino mass prior

To compute the posterior odds of the two possible neutrino hierarchies, we take the uniform reference priors in ϕ, ψ, Σ , and transform to m_L, m_M, m_H . This produces the Jacobian

$$\begin{aligned} J(m_L, m_M, m_H) &= \left\| \frac{\partial(\phi, \psi, \Sigma)}{\partial(m_L, m_M, m_H)} \right\| \\ &= 4(m_L m_M + m_L m_H + m_M m_H), \end{aligned} \quad (4.1)$$

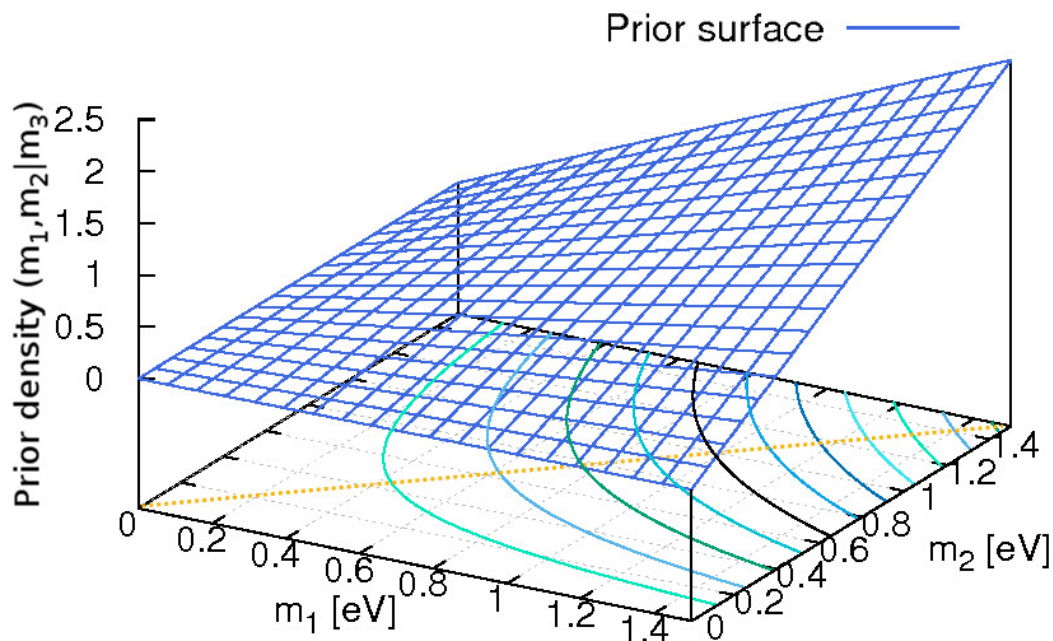


Figure 1. Surface plot of the on average least informative prior density for the neutrino masses (keeping $m_3 = 0.01\text{eV}$ fixed as an example for illustration. This simply determines the intersection with the z -axis; the z -axis can be rescaled for other values of m_3). This prior was constructed from the reference priors for the mass splittings and the sum of neutrino masses, see eq. (4.1). The prior correctly reflects that a priori, a low neutrino mass sum is more difficult to detect (low prior density to the bottom left). It also reflects that keeping one of the three neutrino masses fixed, it is most difficult to detect a neutrino mass splitting if the two other neutrinos have identical masses (minimal prior value along the line $m_1 = m_2$, indicated in yellow dashes). The coloured lines indicate isocontours.

for both hierarchies. J is therefore proportional to the resulting non-informative reference prior in the space of neutrino masses m_i . This prior makes immediate sense. Recall that the prior (4.1) had been constructed from the information which an experiment misses on the parameters.

We now see that the prior is zero, if all neutrino masses are zero. The prior thereby correctly reflects that if the neutrino masses were all zero, then no experiment searching for the masses could ever detect them.

Keeping one of the three masses fixed (m_L say), the prior furthermore takes a minimal value if the two other masses coincide ($m_M = m_H$ in our example). This reflects that an experiment which searches for a mass-splitting between m_M and m_H will then fail, since there is then no mass splitting to be detected.

The prior also increases with increasing masses, which reflects that it is a priori easier to detect the sum of the neutrino masses if it is large. This again makes immediate sense, as highly massive neutrinos will wash out cosmological structures efficiently. The prior hence directly reflects why cosmology is able to yield a sequence of increasingly constraining *upper* bounds. Moreover, the prior increases with the sum of neutrino masses only in a polynomial fashion. This directly implies that it is a sound prior which can be overwritten by the data: for increasingly constraining data, the likelihood will turn towards Gaussian, thereby acquiring exponentially decreasing tails. As an exponential times a polynomial converges, this directly illustrates that the posterior will converge.

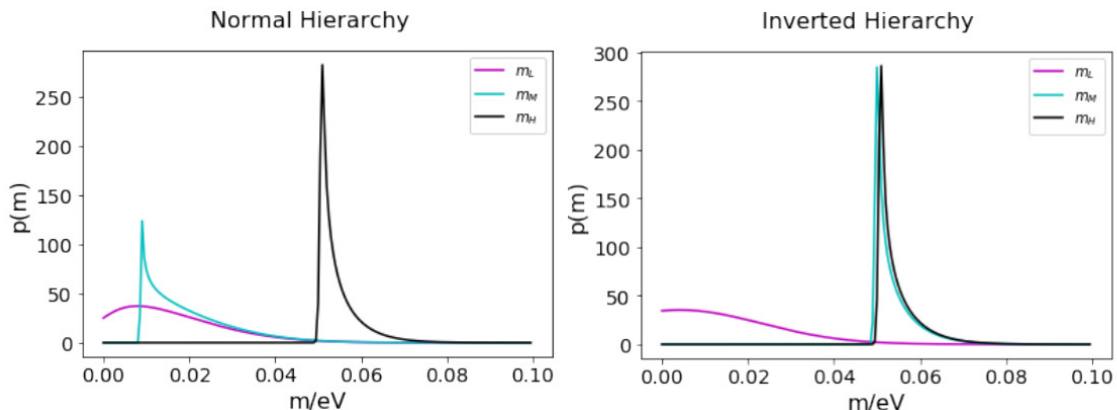


Figure 2. Posteriors for masses for the normal Hierarchy (left), and the Inverted Hierarchy (right).

Comparing the prior (4.1) to the often chosen priors that are flat in mass or proportional to $1/m$, we see a conceptual difference: these priors describe a prior belief that a value for the masses or their logarithm is not preferred. The reference prior eq. (4.1) instead reflects whether the *experimental setup* is designed to be more sensitive with respect to certain combinations of the values for the masses or not. It constructs the prior which is on average most easily overwritten by data. If we would reinterpret the $1/m$ prior as a reference prior, we see that it puts maximal weight on zero masses. This would imply an experiment is a priori most likely to detect the neutrino masses if they were zero (which is of course a contradiction in itself).

5 Results: neutrino masses and hierarchy

5.1 Mass posteriors

Using the constructed least informative prior, eq. (4.1), we show in figure 2 the posteriors for the neutrino masses for the two hierarchies. As expected, for the inverted hierarchy, the two most massive states are tightly constrained by the accelerator/reactor result and the cosmological upper bound. If we wish, we can marginalise over the model (NH or IH), in which case the posteriors become an average of these posteriors, weighted with the Bayes factor. The most massive neutrino state is then tightly constrained, regardless of hierarchy.

5.2 Marginal likelihood for the hierarchies

Having determined the reference prior, constructed to be non-informative for the neutrino masses, we now use it to compute the posterior probability of the normal hierarchy and inverted hierarchy. This model comparison is achieved with the Bayesian evidence, or marginal likelihood. The relative probability of the two models (posterior odds) is

$$\frac{p(\text{NH}|\mathbf{x})}{p(\text{IH}|\mathbf{x})} = \frac{\pi(\text{NH}) \int p(\mathbf{x}|\boldsymbol{\theta}, \text{NH})\pi(\boldsymbol{\theta}|\text{NH}) d\boldsymbol{\theta}}{\pi(\text{IH}) \int p(\mathbf{x}|\boldsymbol{\theta}, \text{IH})\pi(\boldsymbol{\theta}|\text{IH}) d\boldsymbol{\theta}}. \quad (5.1)$$

The first term is the prior model ratio, which we shall take to be unity to express no prior model preference. The second term is the Bayes factor, being the ratio of the evidences, $B = Z_{\text{NH}}/Z_{\text{IH}}$.

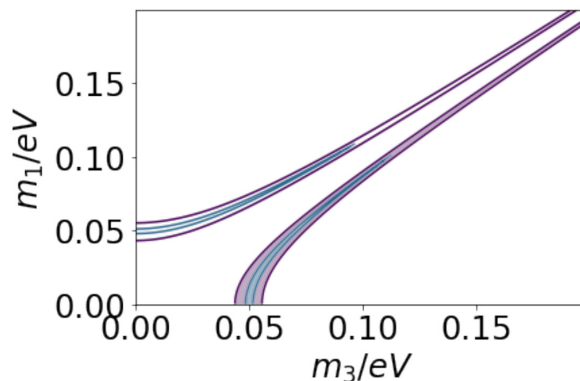


Figure 3. Contours of equal posterior probability density for m_1 and m_3 , marginalised over m_2 , using the reference prior J . These masses have the conventional labelling, with m_1 being the dominant contributor to the electron neutrino. NH is depicted in solid contours at the bottom right (with m_1 identified as m_L) and IH at top left (with open contours and $m_1=m_M$). Contours are set at a very low level to render the posterior visible across the mass range.

The likelihood is given by the product of the three Gaussians for ϕ, ψ and Σ , or their transformed distributions for m_L, m_M, m_H . They are combined with the prior (4.1) such that we can compute the marginal likelihoods for NH and IH respectively:

$$Z_{\text{NH,IH}} = A_{\text{NH,IH}} \int dm_L dm_M dm_H J(m_L, m_M, m_H) \mathcal{N}_\phi(\mu_s, \sigma_s^2) \mathcal{N}_\psi(\mu_a, \sigma_a^2) \mathcal{N}_\Sigma(\mu_\Sigma, \sigma_\Sigma^2), \quad (5.2)$$

with obvious integration limits. $A_{\text{NH,IH}}$ are constants and we insert the appropriate functions of m_i from eq. (3.7). $\mathcal{N}_x(\mu, \sigma^2) = (2\pi\sigma^2)^{-1/2} \exp[-(x-\mu)^2/(2\sigma^2)]$ is the normal distribution, and the appropriate means and variances for the solar (s) and LBL accelerator/reactor (a) experiments are obtained from eq. (3.8). For the cosmological data, we take $\mu_\Sigma = 0$, and $\sigma_\Sigma \simeq 0.06 \text{ eV}$ [27]. The cosmological term is truncated at $\Sigma = 0$. Note that these likelihoods each have a normalisation, which is common for all except the LBL/reactor results, for which the IH peak likelihood is lower by $\Delta\chi^2 = 0.83$ [13], so the NH/IH Bayes factor is increased by $A_{\text{NH}}/A_{\text{IH}} = \exp(0.415) \sim 1.5$.

For illustration, the integrands marginalised over m_2 are shown in figure 3 for the two hierarchies. To compute the Bayes factor, the integrals are taken over the wedge $m_L \leq m_M \leq m_H \leq 0.5 \text{ eV}$. The upper bound is necessary to make the prior proper, but the results are very insensitive to the choice if sufficiently large, because of the requirement from oscillation data that the masses be almost degenerate if large, and the cosmological bound strongly disfavors $\Sigma = 1.5 \text{ eV}$ (i.e. the posterior converges quickly). Furthermore, since the prior volumes are the same for both models, there is essentially no sensitivity of the posterior odds to the chosen prior volume.

When these marginal likelihoods are evaluated, we find that the posterior odds ratio is

$$\frac{p(\text{NH}|\mathbf{x})}{p(\text{IH}|\mathbf{x})} = \frac{Z_{\text{NH}}}{Z_{\text{IH}}} = 5.1 \quad (5.3)$$

in mild favour of the normal hierarchy and with a negligible error from the integration. The result can either be obtained using the J prior and integrating over masses m_i , or by using a constant prior and integrating over (ϕ, ψ, Σ) , paying attention to the physical domain in

the latter case. As expected, the result depends on the cosmological error; for example, increasing the error to $\sigma_\Sigma = 0.07 \text{ eV}$ gives a significantly lower odds ratio of 4.0. Thus we find, in agreement with [19] and [36], but in contrast to [32] and [26], that the reference prior approach still favours the normal hierarchy, but only with low and inconclusive odds.

This makes mathematically robust what could already intuitively be expected: if multiple reanalyses of the same data yield strongly differing posterior odds, and the difference arises from the choice of priors, then it is clear that all additional information must have come from the prior rather than the data. Here, we have shown this by explicitly using a prior that contributes on average minimal prior information to the information inherent in the data, and we indeed find that using such an uninformative prior leads to uninformative posterior odds. This provides a strong case for collecting better data, in order to determine the neutrino hierarchy. Note, that such an improved data set might still reveal that neutrinos obey the inverse hierarchy: all statements in the paper at hand are entirely probabilistic, and the posterior odds present the odds given current data. The odds may well tip once a measurement of the neutrino mass sum has been established.

6 Future neutrino constraints

It is straightforward to forecast the Bayes factor for future cosmological experiments, by reducing the variance on the sum of the masses, and centring the cosmological Gaussian on a fiducial sum of masses. This is shown in figure 4 for the normal hierarchy as a function of the standard deviation of the assumed Gaussian future error. We assume that there is one massless neutrino, so for the normal hierarchy we centre the truncated Gaussian on 0.06 eV , and for the inverted hierarchy, on 0.11 eV . Assuming the normal hierarchy, we see from the figure that the logarithm of the Bayes factor should exceed 5 for $\sigma_\Sigma < 0.025 \text{ eV}$, similar to that found by [19]. Future data could yield uncertainties below 0.02 eV with optimistic Cosmic Microwave Background and galaxy clustering measurements [1], and 3D weak lensing from a Euclid-like survey may give errors $\sigma_\Sigma \simeq 0.03 \text{ eV}$ if the masses are not too low [24], with the advantage of being insensitive to galaxy bias, and there is the prospect of smaller errors in combination with CMB lensing [23]. These results should be achievable by the mid-2020s, on the same timescale as laboratory experiments are expected to yield evidence of the hierarchy (see [28] for a review of these prospects). Improved measurements would be inconclusive if the inverted hierarchy is correct, since the normal hierarchy can increase the minimum mass to find acceptable solutions, and the Bayes factor is always close to unity for very small errors.

7 Conclusions

This paper is an objective Bayesian study analysing the neutrino masses and their hierarchy, using summary information from solar, long baseline accelerator, reactor, and cosmological experiments. Since it was found in former studies that the prior still impacts the physical conclusions strongly, we have determined the on average least informative prior for the described neutrino experiments. This is called a reference prior and for the case at hand coincides with the product of one-dimensional Jeffreys priors. In multi-parameter models, the procedure that leads to the reference prior can depend on the ordering of the parameters [6], but in the case considered here the summary data separate in the ϕ, ψ, Σ parametrisation, and the reference prior is independent of ordering and is unique, given by eq. (4.1). This prior favours larger masses, which is uninformative with respect to the parameters, but reflects

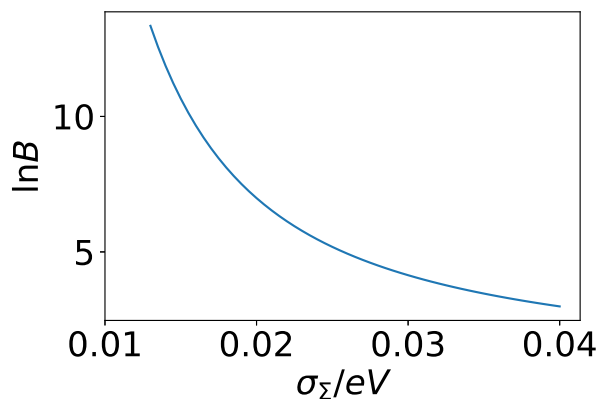


Figure 4. The Bayes factor $B = Z_{\text{NH}}/Z_{\text{IH}}$ for future cosmological experiments where the standard deviation on the sum of neutrino masses reduces from its current value of 0.06eV. One massless neutrino and the normal hierarchy is assumed.

the measurement sensitivity because a given variance in ϕ or ψ corresponds to a smaller uncertainty in mass difference if the masses are larger; see section 4. The uninformative prior is shown in figure 1.

Given this prior, the inverted hierarchy is only mildly disfavoured by current data, despite the minimum total mass being close to the cosmological bound. Using the least informative prior, the posterior odds ratio (assuming equal model priors) is inconclusive, being 5.1:1 in favour of the normal hierarchy.

As has been remarked upon previously, the conclusions on the neutrino hierarchy are prior-dependent. The results using the non-informative reference prior lie between the low odds ratios of [16, 19, 36], based on uniform mass priors, and the conclusive results of [26, 32] which use the logarithm of the masses as parameters for the prior, and a random mass matrix prior respectively. The prospects of future cosmological experiments determining the hierarchy depends on which hierarchy it is. With errors on the sum of masses of the order of 0.025 eV or better, they would find strong evidence ($\ln B > 5$) for the normal hierarchy if it is true and the smallest mass is negligibly small. If, on the other hand, the inverted hierarchy is correct, future improvements in cosmological data will not conclusively determine the hierarchy ($\ln B$ does not vary significantly from zero), since the NH can accommodate the larger mass sum by increasing the smallest mass.

There are a number of refinements that could be made to this first objective Bayesian approach, principally in refining the reference priors to take full account of the small variation in sensitivity of experiments to the mass sum and mass-squared differences. However, over the range where the likelihood is substantial, the approximation is very good for the terrestrial experiments, and reasonable for the cosmological experiment, so it is difficult to see how the results could change substantially. Secondly, from a Bayesian perspective a marginalisation over parameters could be done in the precursor LBL/reactor analysis, rather than a profile likelihood that is used to construct the mass-squared difference posteriors [13]. Finally, any dependence of the solar and accelerator-reactor results on each other should be treated by a joint likelihood, but the dependences are very weak, so again, it seems unlikely that residual correlations could turn the inconclusive results into a strong preference for the normal hierarchy, or favour the inverted hierarchy.

In total, better data are needed before the neutrino masses and their hierarchy can be well determined. The prior here constructed is then still applicable, and stays on average the prior that is most easily overwritten by data.

Acknowledgments

We are grateful to Fergus Simpson, Roberto Trotta, Morgan Wascko, Yoshi Uchida, Raul Jimenez, Licia Verde and Carlos Peña-Garay for helpful discussions and comments on the manuscript.

References

- [1] R. Allison, P. Caucal, E. Calabrese, J. Dunkley and T. Louis, *Towards a cosmological neutrino mass detection*, *Phys. Rev. D* **92** (2015) 123535 [[arXiv:1509.07471](#)] [[INSPIRE](#)].
- [2] J.O. Berger, *Statistical decision theory and Bayesian analysis*, second edition, Springer Verlag, New York U.S.A., (1985).
- [3] J.O. Berger, J.M. Bernardo and D. Sun, *The formal definition of reference priors*, *Ann. Statist.* **37** (2009) 905.
- [4] J. Bergström, M.C. Gonzalez-Garcia, M. Maltoni and T. Schwetz, *Bayesian global analysis of neutrino oscillation data*, *JHEP* **09** (2015) 200 [[arXiv:1507.04366](#)] [[INSPIRE](#)].
- [5] J. Bernardo, *Reference posterior distributions for Bayesian inference (with discussion)*, *J. Roy. Statist. Soc. B* **41** (1979) 113.
- [6] J. Bernardo, *Reference analysis*, *Handbook Statist.* **25** (2005) 17.
- [7] M. Blennow, *On the Bayesian approach to neutrino mass ordering*, *JHEP* **01** (2014) 139 [[arXiv:1311.3183](#)] [[INSPIRE](#)].
- [8] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri and A. Palazzo, *Global constraints on absolute neutrino masses and their ordering*, *Phys. Rev. D* **95** (2017) 096014 [[arXiv:1703.04471](#)] [[INSPIRE](#)].
- [9] L. Chen, Q.-G. Huang and K. Wang, *New cosmological constraints with extended-baryon oscillation spectroscopic survey DR14 quasar sample*, *Eur. Phys. J. C* **77** (2017) 762 [[arXiv:1707.02742](#)] [[INSPIRE](#)].
- [10] A.J. Cuesta, V. Niro and L. Verde, *Neutrino mass limits: robust information from the power spectrum of galaxy surveys*, *Phys. Dark Univ.* **13** (2016) 77 [[arXiv:1511.05983](#)] [[INSPIRE](#)].
- [11] P.F. de Salas, D.V. Forero, C.A. Ternes, M. Tortola and J.W.F. Valle, *Status of neutrino oscillations 2017*, [arXiv:1708.01186](#) [[INSPIRE](#)].
- [12] M.L. Eaton and W.D. Sudderth, *Group invariant inference and right Haar measure*, *J. Statist. Planning Infer.* **103** (2002) 87.
- [13] I. Esteban, M.C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler and T. Schwetz, *Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity*, *JHEP* **01** (2017) 087 [[arXiv:1611.01514](#)] [[INSPIRE](#)].
- [14] D.A.S. Fraser, N. Reid, E. Marras and G.Y. Yi, *Default priors for Bayesian and frequentist inference*, *J. Roy. Statist. Soc. B* **72** (2010) 631.
- [15] S. Gariazzo, M. Archidiacono, P.F. de Salas, O. Mena, C.A. Ternes and M. Tortola, *Neutrino masses and their ordering: global data, priors and models*, *JCAP* **03** (2018) 011 [[arXiv:1801.04946](#)] [[INSPIRE](#)].

- [16] M. Gerbino, M. Lattanzi, O. Mena and K. Freese, *A novel approach to quantifying the sensitivity of current and future cosmological datasets to the neutrino mass ordering through Bayesian hierarchical modeling*, *Phys. Lett. B* **775** (2017) 239 [[arXiv:1611.07847](#)] [[INSPIRE](#)].
- [17] E. Giusarma, M. Gerbino, O. Mena, S. Vagnozzi, S. Ho and K. Freese, *Improvement of cosmological neutrino mass bounds*, *Phys. Rev. D* **94** (2016) 083522 [[arXiv:1605.04320](#)] [[INSPIRE](#)].
- [18] M.C. Gonzalez-Garcia, M. Maltoni and T. Schwetz, *Updated fit to three neutrino mixing: status of leptonic CP-violation*, *JHEP* **11** (2014) 052 [[arXiv:1409.5439](#)] [[INSPIRE](#)].
- [19] S. Hannestad and T. Schwetz, *Cosmology and the neutrino mass ordering*, *JCAP* **11** (2016) 035 [[arXiv:1606.04691](#)] [[INSPIRE](#)].
- [20] D. Heath and W. Sudderth, *On finitely additive priors, coherence, and extended admissibility*, *Ann. Statist.* **6** (1978) 333.
- [21] H. Jeffreys, *Theory of probability*, Oxford University Press, Oxford U.K., (1961).
- [22] KAMLAND-ZEN collaboration, A. Gando et al., *Search for Majorana neutrinos near the inverted mass hierarchy region with KamLAND-Zen*, *Phys. Rev. Lett.* **117** (2016) 082503 [*Addendum ibid.* **117** (2016) 109903] [[arXiv:1605.02889](#)] [[INSPIRE](#)].
- [23] T.D. Kitching, A.F. Heavens and S. Das, *3D weak gravitational lensing of the CMB and galaxies*, *Mon. Not. Roy. Astron. Soc.* **449** (2015) 2205 [[arXiv:1408.7052](#)] [[INSPIRE](#)].
- [24] T.D. Kitching, A.F. Heavens, L. Verde, P. Serra and A. Melchiorri, *Finding evidence for massive neutrinos using 3D weak lensing*, *Phys. Rev. D* **77** (2008) 103008 [[arXiv:0801.4565](#)] [[INSPIRE](#)].
- [25] P. Langacker, S.T. Petcov, G. Steigman and S. Toshev, *On the Mikheev-Smirnov-Wolfenstein (MSW) mechanism of amplification of neutrino oscillations in matter*, *Nucl. Phys. B* **282** (1987) 589 [[INSPIRE](#)].
- [26] A.J. Long, M. Raveri, W. Hu and S. Dodelson, *Neutrino mass priors for cosmology from random matrices*, *Phys. Rev. D* **97** (2018) 043510 [[arXiv:1711.08434](#)] [[INSPIRE](#)].
- [27] N. Palanque-Delabrouille et al., *Neutrino masses and cosmology with Lyman- α forest power spectrum*, *JCAP* **11** (2015) 011 [[arXiv:1506.05976](#)] [[INSPIRE](#)].
- [28] X. Qian and P. Vogel, *Neutrino mass hierarchy*, *Prog. Part. Nucl. Phys.* **83** (2015) 1 [[arXiv:1505.01891](#)] [[INSPIRE](#)].
- [29] E. Sellentin, *A fast, always positive definite and normalizable approximation of non-Gaussian likelihoods*, *Mon. Not. Roy. Astron. Soc.* **453** (2015) 893 [[arXiv:1506.04866](#)] [[INSPIRE](#)].
- [30] E. Sellentin and R. Durrer, *Detecting the cosmological neutrino background in the CMB*, *Phys. Rev. D* **92** (2015) 063012 [[arXiv:1412.6427](#)] [[INSPIRE](#)].
- [31] E. Sellentin, M. Quartin and L. Amendola, *Breaking the spell of Gaussianity: forecasting with higher order Fisher matrices*, *Mon. Not. Roy. Astron. Soc.* **441** (2014) 1831 [[arXiv:1401.6892](#)] [[INSPIRE](#)].
- [32] F. Simpson, R. Jimenez, C. Pena-Garay and L. Verde, *Strong Bayesian evidence for the normal neutrino hierarchy*, *JCAP* **06** (2017) 029 [[arXiv:1703.03425](#)] [[INSPIRE](#)].
- [33] D. Sun and J. Berger, *Objective Bayesian analysis for the multivariate normal model*, in *Proc. Valencia/ISBA 8th World Meeting on Bayesian Statistics*, Benidorm Alicante Spain, (2006).
- [34] Y. Takeuchi and K. Kadota, *Probing neutrinos from Planck and forthcoming galaxy redshift surveys*, *JCAP* **01** (2014) 046 [[arXiv:1310.0037](#)] [[INSPIRE](#)].
- [35] A. Timmons, *The results of MINOS and the future with MINOS+*, *Adv. High Energy Phys.* **2016** (2016) 7064960 [[arXiv:1511.06178](#)] [[INSPIRE](#)].
- [36] S. Vagnozzi et al., *Unveiling ν secrets with cosmological data: neutrino masses and mass hierarchy*, *Phys. Rev. D* **96** (2017) 123503 [[arXiv:1701.08172](#)] [[INSPIRE](#)].