



Statistical interpretation of sterile neutrino oscillation searches at reactors

Pilar Coloma^{1,2,a}, Patrick Huber^{3,b}, Thomas Schwetz^{4,c}

¹ Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, Calle Nicolás Cabrera 13-15, 28049 Madrid, Spain

² Instituto de Física Corpuscular UV/CSIC, Parque Científico, Calle Catedrático José Beltrán 2, 46980 Paterna, Spain

³ Physics Department, Center for Neutrino Physics, Virginia Tech, Blacksburg, USA

⁴ Institut für Astroteilchenphysik, Karlsruher Institut für Technologie (KIT), 76021 Karlsruhe, Germany

Received: 31 August 2020 / Accepted: 16 December 2020

© The Author(s) 2020

Abstract A considerable experimental effort is currently under way to test the persistent hints for oscillations due to an eV-scale sterile neutrino in the data of various reactor neutrino experiments. The assessment of the statistical significance of these hints is usually based on Wilks' theorem, whereby the assumption is made that the log-likelihood is χ^2 -distributed. However, it is well known that the preconditions for the validity of Wilks' theorem are not fulfilled for neutrino oscillation experiments. In this work we derive a simple asymptotic form of the actual distribution of the log-likelihood based on reinterpreting the problem as fitting white Gaussian noise. From this formalism we show that, even in the absence of a sterile neutrino, the expectation value for the maximum likelihood estimate of the mixing angle remains non-zero with attendant large values of the log-likelihood. Our analytical results are then confirmed by numerical simulations of a toy reactor experiment. Finally, we apply this framework to the data of the Neutrino-4 experiment and show that the null hypothesis of no-oscillation is rejected at the 2.6σ level, compared to 3.2σ obtained under the assumption that Wilks' theorem applies.

1 Introduction

Despite the strong evidence for the existence of physics beyond the Standard Model (SM), searches for new particles at high-energy colliders have been so far unsuccessful. A possible explanation is that the new physics lies at relatively low scales, and that the dark and the visible sectors communicate through a very weakly-interacting particle. A

particularly appealing scenario in this context is the addition of right-handed neutrinos to the SM, which could also source the SM neutrino masses. While right-handed neutrinos are singlets of the SM gauge group (and therefore *sterile* under the SM interactions), they may lead to observable signatures through their mixing with the SM neutrinos at a variety of experiments, depending on their masses. In particular, and motivated by the LSND anomaly [1], considerable experimental effort was directed towards the search for eV-scale right-handed neutrinos using short-baseline oscillation experiments over the past two decades, see Ref. [2] for a recent review. In this article we focus on reactor neutrino experiments, which have been playing a crucial role in neutrino physics since the discovery of the neutrino by Cowan et al. [3], and have been central to the search for eV-scale sterile neutrinos.

Calculations of the anti-neutrino fluxes emitted from nuclear reactors performed in 2011 [4,5] have led to the so-called “reactor anti-neutrino anomaly” [6], providing a hint for the existence of sterile neutrinos due to an observed mis-match between predicted and observed rates. This hint remains controversial up to today, see Refs. [7,8] for discussions of the latest developments on this issue. Therefore, modern experiments focus on the relative comparison of measured spectra at different baselines [9–16], which is more robust against uncertainties in flux predictions. Recent global fits of reactor data still find indications for sterile neutrino oscillations at the level of $2 - 3\sigma$, even when using only spectral ratios [17–20]. Those indications are based purely on spectral distortions which may feature oscillatory patterns.

Note, however, that the necessary conditions to apply Wilks' theorem [21] are typically not fulfilled for sterile neutrino searches in oscillation experiments, which can lead to wrong results when evaluating significance or confidence lev-

^a e-mail: pilar.coloma@ift.csic.es

^b e-mail: pahuber@vt.edu

^c e-mail: schwetz@kit.edu (corresponding author)

els based on χ^2 values.¹ Therefore, one should wonder if the hints may be coming from a mis-interpretation of the data. This has recently been pointed out in Refs. [27–29], where the authors have shown by way of Monte Carlo simulations that there are indeed corrections and that the statistical significance of the hints is reduced. In this paper we attempt to go one step further and, besides providing analytical arguments that allow to understand the expected distribution for the test statistics, we also study the dependence of the observed corrections on relevant experimental parameters and numerical details of the analysis.

The outline of the paper is as follows. In Sect. 2 we start with some general remarks, introduce a test statistic to evaluate the significance of the presence of sterile neutrino oscillations and give some qualitative arguments why it is likely that experiments find a hint for sterile neutrinos even if there are none. In Sect. 3 we consider an idealized disappearance experiment and derive the expected distribution of the test statistics. We show that the above statement is a consequence of the fact that in a Fourier composition of white noise some frequency will appear with largest amplitude. This will allow us to make predictions for the expected distribution of the best fit points for $\sin^2 2\theta$. In Sect. 4 we perform numerical simulations of toy reactor experiments, and we study in detail the distribution of the test statistic as well as the location of the best-fit points. We also investigate the impact of various parameters, such as restriction to the physical region, the impact of systematics, or alternative χ^2 definitions. In Sect. 6 we consider as a case study the recent results for the Neutrino-4 experiment [16,30], which has reported a $\sim 3\sigma$ hint for sterile neutrino oscillations. We perform a Monte-Carlo study of the Neutrino-4 data and show that the significance for the presence of sterile neutrinos is somewhat lower than expected under the assumption of a χ^2 distribution of the test statistics, and we present the confidence regions obtained by explicit Neyman–Pearson construction based on the Feldman–Cousins prescription [31]. We conclude in Sect. 6.

2 General remarks

Reactor neutrino experiments look for the disappearance of electron anti-neutrinos. In this work we assume that a single sterile neutrino is relevant for the phenomenology and we consider only experimental setups where the baseline is short enough such that oscillations due to the standard three-flavor mass-squared differences can be safely neglected. In this limit, sterile neutrino oscillations are described by an

effective two-flavor survival probability:

$$P^{\text{osc}} = 1 - \frac{1}{2} \sin^2 2\theta \left(1 - \cos \frac{\Delta m^2 L}{2E} \right), \quad (1)$$

where E is the neutrino energy, L is the baseline, θ is an effective neutrino mixing angle, and Δm^2 stands for the mass-squared splitting between the eV-scale mass state and the light SM neutrinos.

In order to analyze the results of a given experiment a least-squares function of binned spectral data is considered, $\chi^2(\sin^2 2\theta, \Delta m^2)$. A common test statistic T for evaluating the hypothesis of the presence of sterile neutrino oscillations is the $\Delta\chi^2$ (or, equivalently, the likelihood ratio) between the best-fit point and the no-oscillation case:

$$\begin{aligned} T &= \chi^2(\text{no osc}) - \chi^2(\text{best fit}) \\ &= \chi^2(0, 0) - \chi^2(\widehat{\sin^2 2\theta}, \widehat{\Delta m^2}), \end{aligned} \quad (2)$$

where $\widehat{\sin^2 2\theta}$ and $\widehat{\Delta m^2}$ indicate the parameter values at the χ^2 minimum. If Wilks' theorem [21] applies, T should be distributed as a χ^2 -distribution for 2 degrees of freedom (DOF), corresponding to the two minimized parameters.

Indeed, for the problem at hand, there are several reasons to suspect that the necessary conditions for Wilks' theorem to apply are not fulfilled: First, there is a physical boundary for the mixing angle, $\sin^2 2\theta \geq 0$. Second, the parameter Δm^2 becomes undefined for $\sin^2 2\theta \rightarrow 0$ and $\sin^2 2\theta$ becomes unphysical for $\Delta m^2 \rightarrow 0$. Third, the cosine dependence on Δm^2 of the oscillation probability in Eq. 1 leads to a strong non-linear behavior. Therefore, significant deviations of the distribution of T from a χ^2 -distribution are expected *a priori*, see also [27,31]. A recent review discussing the applicability of Wilk's theorem can be found in Ref. [32].

As we will show in Sect. 3, for an idealized situation the distribution of \sqrt{T} is the one of the maximum of N standard normal random variables, where N corresponds to an effective number of bins. We will give physical arguments, as to why for this type of experiments a non-vanishing value for $\sin^2 2\theta$ at the best-fit is likely, with a relatively large value of T . In fact, its typical value is set by the size of the relative statistical uncertainty of the sample. In Sect. 4 we will compute the distribution of T for more realistic configurations and will always confirm rather large deviations from a χ^2 -distribution. This suggests that reliable statements about significance and confidence levels require explicit Monte-Carlo simulations, in agreement with previous results [27,28]. We will demonstrate this explicitly using the recent results from Neutrino-4 in Sect. 5.

¹ For similar studies in the context of three-flavor neutrino oscillations see for instance [22–26].

3 Expected distribution of the test statistics

In this section we derive the expected test statistic T for a toy model. After making a series of assumptions that allow us to write the randomly fluctuated events in each bin as a discrete Fourier transform, we will proceed and minimize the χ^2 function analytically. This will provide us with an expression in terms of the Fourier coefficients of the expansion, which we can then substitute into Eq. 2 to get an analytical expression for the expected test statistic.

3.1 Derivation of the test statistic for a toy model

Let us consider a toy model for an oscillation disappearance experiment. We consider N bins in L/E and write the predicted event number in each bin as

$$p_i \approx p_i^0 P_i^{\text{osc}}, \tag{3}$$

where p_i^0 is the predicted number of events in case of no oscillations, P_i^{osc} is given in Eq. 1, and the index i labels the bins in L/E . Let us now adopt the following assumptions:

- (a) We assume that there is *no sterile neutrino* in Nature, i.e., the observed data in bin i is given by the no-oscillation prediction plus a statistical fluctuation with variance σ_i^2 :

$$d_i = p_i^0 + \delta d_i \quad \text{with} \\ \langle \delta d_i \rangle = 0, \quad \langle \delta d_i \delta d_j \rangle = \sigma_i^2 \delta_{ij}. \tag{4}$$

For Poisson statistics we would have $\sigma_i = \sqrt{p_i^0}$. For simplicity we assume that δd_i are Gaussian.

- (b) We assume that *only shape information in L/E* is used, but not on the absolute normalization. This applies to experiments where energy spectra are fitted leaving the overall normalization free, but also to setups or combinations of experiments where relative spectra at different baselines are considered.

The second assumption is implemented in a χ^2 by introducing a free pull parameter ξ , such that

$$\chi^2 = \sum_{i=1}^N \left[\frac{d_i - (1 + \xi)p_i}{\sigma_i} \right]^2, \tag{5}$$

where minimization with respect to ξ is understood. Working to linear order in ξ and $\sin^2 2\theta$, and assuming that terms involving a sum over δd_i or $\cos(\Delta m^2 L/2E_i)$ average to zero, one finds that $\xi \approx (1/2) \sin^2 2\theta$ minimizes the χ^2 . Using this

together with assumption (a) above we find

$$\chi^2 = \sum_{i=1}^N \left[n_i - \frac{p_i^0}{2\sigma_i} \sin^2 2\theta \cos \frac{\Delta m^2 L}{2E_i} \right]^2, \tag{6}$$

where $n_i \equiv \delta d_i/\sigma_i$ are independent standard normal random variables with $\langle n_i \rangle = 0$, $\langle n_i n_j \rangle = \delta_{ij}$, see Eq. 4.

Let us now adopt the additional simplifying assumptions to build a mathematical toy model:

- (c) We assume that the relative statistical uncertainty σ_i/p_i^0 has the same value for each bin and define the new parameter

$$s \equiv \frac{p_i^0}{2\sigma_i} \sin^2 2\theta. \tag{7}$$

Although this is not strictly the case for a reactor experiment, in Sect. 4 we will see that it works relatively well for the experimental setups under consideration in this work. Furthermore we assume that bins have equal width in L/E and define

$$\frac{\Delta m^2}{2} \left(\frac{L}{E} \right)_j = \frac{2\pi}{N} \kappa j \equiv \varphi_{\kappa j}. \tag{8}$$

Hence, j labels bins in L/E while the index κ labels discrete frequencies proportional to Δm^2 . With this idealization, Eq. 6 becomes

$$\chi^2(s, \kappa) = \sum_{i=1}^N [n_i - s \cos \varphi_{\kappa i}]^2. \tag{9}$$

We see that in this limit the sterile neutrino search is equivalent to fitting Gaussian white-noise with a cosine function with the amplitude s and the frequency κ as free parameters. This form suggests to consider the discrete Fourier transform of the N random variables n_i :

$$n_i = \sum_{\kappa=1}^N (a_\kappa \cos \varphi_{\kappa i} + b_\kappa \sin \varphi_{\kappa i}) \tag{10}$$

with $a_\kappa, b_\kappa \in \mathbb{R}$. Focusing on the cosine term, the coefficients a_κ can be computed as

$$a_\kappa = \frac{2}{N} \sum_{i=1}^N n_i \cos \varphi_{\kappa i}. \tag{11}$$

Since n_i are independent standard Gaussian variables, it is clear that a_κ are random Gaussian variables as well, with

$$\langle a_\kappa \rangle = 0, \quad \langle a_\kappa a_\lambda \rangle = \frac{2}{N} \delta_{\kappa\lambda}, \tag{12}$$

where we have assumed that sums over $\cos \varphi_{\kappa i}$ ($\cos^2 \varphi_{\kappa i}$) average to 0 ($N/2$).

Let us now look for the best fit point $(\hat{s}, \hat{\kappa})$. We start by minimizing the χ^2 in Eq. 9 with respect to s , for fixed κ . This gives

$$\hat{s}(\kappa) = \frac{\sum_{i=1}^N n_i \cos \varphi_{\kappa i}}{\sum_{i=1}^N \cos^2 \varphi_{\kappa i}} = a_\kappa, \tag{13}$$

where in the last step we have inserted n_i from Eq. 10, using the fact that all terms average to zero except the one containing $\cos^2 \varphi_{\kappa i}$. This implies that, for fixed κ , $\hat{s}(\kappa)$ follows a Gaussian distribution with its mean and variance as given by Eq. 12. Also, we see that, for fixed κ , the χ^2 is minimized by choosing s as the Fourier coefficient corresponding to the frequency κ . Inserting the Fourier transform from Eq. 10 as well as the solution from Eq. 13 into the χ^2 in Eq. 9, we find

$$\begin{aligned} \chi^2(\hat{s}(\kappa), \kappa) &= \sum_{i=1}^N \left[\sum_{\lambda \neq \kappa} a_\lambda \cos \varphi_{\lambda i} + \sum_{\lambda} b_\lambda \sin \varphi_{\lambda i} \right]^2 \\ &= \frac{N}{2} \sum_{\lambda \neq \kappa} a_\lambda^2 + C. \end{aligned} \tag{14}$$

where in the second step we have expanded the square and used the fact that only terms of the form $\cos^2 \varphi_{\lambda i}$ or $\sin^2 \varphi_{\lambda i}$ survive, while all mixed terms average to zero. Here, the constant C contains the b_λ terms and is independent of κ .

Next, we minimize with respect to κ . Since Eq. 14 is a sum of positive terms, the χ^2 would be minimal for $\kappa = \hat{\kappa}$, such that $a_{\hat{\kappa}}$ is the Fourier coefficient with the largest absolute value. However, considering the definition of s in Eq. 7, we see that the physical requirement $\sin^2 2\theta \geq 0$ implies $s \geq 0$. Therefore, if the minimization is restricted to the physically allowed region we obtain

$$\hat{s} \equiv \hat{s}(\hat{\kappa}) = a_{\hat{\kappa}} = \max \left[0, \max_{\kappa} a_\kappa \right]. \tag{15}$$

For N Gaussian variables with $\langle a_\kappa \rangle = 0$, the probability that all a_κ are negative is $(1/2)^N$. Hence, for sufficiently large N it is very likely to obtain at least one positive a_κ , such that Eq. 15 leads to a positive best-fit point for the parameter s (and therefore for $\sin^2 2\theta$). For simplicity we neglect hereafter the unlikely case that none of the a_κ is positive.

Finally, let us consider the test statistic $T = \chi^2(0, 0) - \chi^2(\hat{s}, \hat{\kappa})$ defined in Eq. 2. Using Eqs. 9 and 10 we find the χ^2 for the SM point as

$$\chi^2(0, 0) = \sum_i (n_i)^2 = \frac{N}{2} \sum_{\lambda} a_\lambda^2 + C, \tag{16}$$

where C is the same constant as in Eq. 14. Evaluating now the minimum of the χ^2 , $\chi^2(\hat{s}, \kappa)$ as given in Eq. 14 at $\kappa = \hat{\kappa}$, we finally obtain

$$T = \frac{N}{2} a_{\hat{\kappa}}^2 = \left[\max_{\kappa} \tilde{a}_\kappa \right]^2, \tag{17}$$

where $\tilde{a}_\kappa \equiv \sqrt{N/2} a_\kappa$ are standard normal random variables, $\langle \tilde{a}_\kappa \rangle = 0$, $\langle \tilde{a}_\kappa \tilde{a}_\lambda \rangle = \delta_{\kappa\lambda}$ (see Eq. 12).

3.2 Discussion

Equations 15 and 17 are the main results of this section. The latter shows that the square-root of the test statistic T has the distribution of the maximum of N standard normal variables. It is proportional to the best fit amplitude \hat{s} , and hence, up to a normalization factor, the best-fit point in Eq. 15 follows the same distribution. Distributions of this type are considered in the field of ‘‘extreme value statistics’’, see e.g., [33,34].

For the case of Gaussian variables of interest here, there exists a limiting distribution for $N \rightarrow \infty$. It is based on the so-called *Gumbel* distribution $e^{-e^{-z}}$. Let $x = \max_i \tilde{a}_i$, where \tilde{a}_i are N standard normal variables. For finite N the cumulative probability distribution (CDF) $F(x)$ can be approximated by [34]:

$$F(x) = \exp \{ - \exp [-A_N(x - B_N)] \}, \tag{18}$$

with

$$A_N = \sqrt{2 \log N}, \quad B_N = A_N - \frac{\log \log N + \log 4\pi}{2A_N}. \tag{19}$$

In Fig. 1 (left) we show 1–CDF for the maximum of N standard normal variables obtained by numerical calculations (solid) compared to the approximate formula in Eq. 18 (long-dashed) for $N = 30$ and 60. We see that they agree reasonably well for $1 - \text{CDF} \gtrsim 0.1$, but start to deviate for smaller values. Indeed, the convergence to the Gumbel distribution goes only as $1/\log N$ [34]. Therefore, since the distribution can be easily calculated numerically, we will for the rest of the paper stick to the numerical method and denote this distribution by ‘‘Max. Gauss’’ in the following.

The important property of this distribution is, that small values of T are rather unlikely. In Fig. 1 we compare 1–CDF as well as the probability density function (PDF) for \sqrt{T} to the one for the square-root of a χ^2 distribution.² Indeed, if Wilks’ theorem was applicable, T should be distributed as χ^2 with 2 DOF. Obviously, the conditions for Wilks’ theorem to hold are badly violated in this case, for the reasons mentioned in Sect. 2. The peak at $\sqrt{T} \sim 2$ and the small probability to

² Note that the Jacobian of the variable transformation has to be taken into account when transforming the PDF for χ^2 into the PDF for $\sqrt{\chi^2}$.




Fig. 1 Expected distribution for the square-root of the test statistic $\sqrt{T} = \max_{\kappa} \tilde{a}_{\kappa}$, where \tilde{a}_{κ} are N standard normal random variables. The left panel shows 1-CDF, while the PDF is shown in the right panel. Dark (light) blue curves correspond to $N = 30$ (60). Solid curves are

obtained by numerical simulations, whereas long-dashed curves correspond to the approximation in Eq. 18. For comparison, the short-dashed gray and black curves show the distributions obtained if T would follow a χ^2 distribution for 2 and 5 DOF, respectively

obtain $\sqrt{T} \lesssim 1$ indicates that, even if there is no sterile neutrino present in Nature, it is very likely to obtain a best-fit point with finite $\sin^2 2\theta$ as well as relatively large value of T . This would lead to claiming a signal at relevant statistical significance, if evaluated with a χ^2 -distribution. The physical reason for this behavior can be understood from Eq. 9: in a white noise spectrum it is very likely to find some frequency with sizable amplitude that is able to fit the data.

The expectation value for a random variable z with the CDF $F(z) = e^{-e^{-z}}$ is given by [34] $\langle z \rangle = \gamma$, where $\gamma = 0.57721 \dots$ is the Euler–Mascheroni constant. From Eqs. 17 and 18 follows then

$$\langle \sqrt{T} \rangle = B_N + \frac{\gamma}{A_N} \approx 2 \dots 2.4, \tag{20}$$

where the numbers hold for $N \approx 30 \dots 60$. These values agree to a good accuracy with the mean values obtained numerically, and depend only weakly (logarithmically) on N . We can use these results to estimate the expectation value for $\sin^2 2\theta$. Let \mathcal{N} be the total number of observed events. According to assumption (c) above we have $p_i^0 \approx \mathcal{N}/N$ and $\sigma_i = \sqrt{p_i^0}$. Then Eq. 7 leads to

$$\sin^2 2\theta = 2\sqrt{\frac{2}{\mathcal{N}}}\sqrt{T} \quad \text{and} \quad \langle \sin^2 2\theta \rangle \approx \frac{6.2}{\sqrt{\mathcal{N}}}, \tag{21}$$

where in the second relation we have used the numerical values from Eq. 20. We see that up to a numerical factor, the expected best fit value for $\sin^2 2\theta$ is set by the relative statistical uncertainty of the event sample. We will find this behavior

in the simulations discussed in the following sections. From Fig. 1 we see that there is a lower bound of $\sqrt{T} \gtrsim 1.5$ at 99% CL for $N = 30 \dots 60$, which translates into a lower bound on $\sin^2 2\theta$ according to Eq. 21.

To conclude this section, we remark that the idealized situation considered here is certainly an over-simplification, and especially assumption (c) will not be satisfied in a realistic oscillation experiment. Nevertheless, these considerations capture the most relevant features and the results obtained here allow an intuitive understanding of the numerical results we are going to present below. In particular, the preference for the presence of sterile neutrino oscillations even in case of no true signal is predicted from those arguments, and allows a qualitative (in some cases even quantitative) understanding of the more realistic simulations discussed in the remainder of this paper.

4 Numerical simulations for a toy experiment

4.1 Description of the simulation

In order to verify the validity of the analytical approach presented in the previous section, we now proceed to perform a numerical simulation for a toy experiment. For this purpose, we choose a reactor disappearance experiment which aims to set a constraint on the $\sin^2 2\theta - \Delta m^2$ parameter space from the observation of $\bar{\nu}_e \rightarrow \bar{\nu}_e$ oscillations. We consider generic shapes of the anti-neutrino flux and inverse beta-decay detection cross section. The distance between the reactor core and

the detector is set to $L = 10$ m. In order to account for the finite size of the reactor core, the probability is averaged over a window $\Delta L = \pm 1$ m:

$$\langle P^{\text{osc}}(\theta, \Delta m^2) \rangle = 1 - \sin^2 2\theta \frac{\int_{L-\Delta L}^{L+\Delta L} dL' \sin^2(\frac{\Delta m^2 L'}{4E})/L'^2}{\int_{L-\Delta L}^{L+\Delta L} dL' \frac{1}{L'^2}} \tag{22}$$

This ensures that fast oscillations are averaged-out at the detector. Unless otherwise stated, the exposure is set such that the total number of events is 1.5×10^4 . A binned χ^2 analysis is performed, using 43 bins in energy of equal size distributed between 2 and 8 MeV, and a Gaussian energy resolution of the form $\sigma(E) = 0.03\sqrt{E/\text{MeV}}$ is applied to the event distributions.

The experimental details outlined above have been chosen to lie in the same ballpark as for some of the running short baseline reactor disappearance experiments [12–16]. However, we have explicitly checked that changing any of these parameters does not qualitatively affect our results. Finally, we have assumed negligible backgrounds in our analysis for simplicity, in order to ease the interpretation of our results. Again in this case, we have checked that the inclusion of a sizable background component does not alter qualitatively our conclusions.

The results presented in this section have been obtained by simulating a large sample of pseudo-experiments, applying random statistical fluctuations to the expected event rates for the reactor experiment setup outlined above. Since here we are mostly interested in evaluating the significance of a potential positive signal, throughout this section we will generate random data under the null-hypothesis, that there is no sterile neutrino in Nature, i.e., for no oscillations. These are generated on a bin-per-bin basis, sampling a normal distribution with its mean set to the expected number of events in a given bin for the SM hypothesis $p_i^0 \equiv p_i(\theta = 0)$, and its width set to the associated statistical uncertainty $\sqrt{p_i^0}$. For the large number of events considered here the Gaussian approximation to the Poisson distribution is well justified. Unless otherwise stated, the number of pseudo-experiments simulated is set to 20,000 for each of the cases studied in this section.

The sample of pseudo-experiments will then be used to determine the distribution of our test statistics T defined in Eq. 2. In order to do so, for each pseudo-experiment a Poisson χ^2 function is built. For a set of parameters $(\theta, \Delta m^2)$, it reads [35]:

$$\chi_{\text{stat,Poisson}}^2(\theta, \Delta m^2) = 2 \sum_i \left[(1 + \xi) p_i(\theta, \Delta m^2) - d_i - d_i \log \frac{(1 + \xi) p_i(\theta, \Delta m^2)}{d_i} \right], \tag{23}$$

where p_i is the expected number of events in the i -th bin for θ and Δm^2 (in the absence of statistical fluctuations), while d_i is the ‘‘observed’’ number of events, i.e., the pseudo-data generated as described above. Here, ξ is a nuisance parameter, introduced in order to account for the systematic uncertainty in the prediction of the expected event rates. Once Eq. 23 has been computed, a pull-term is added and the result is minimized over the nuisance parameter ξ :

$$\chi^2(\theta, \Delta m^2) = \min_{\xi} \left[\chi_{\text{stat,Poisson}}^2(\theta, \Delta m^2) + \left(\frac{\xi - \bar{\xi}}{\sigma_{\text{sys}}} \right)^2 \right] \tag{24}$$

where σ_{sys} stands for the prior uncertainty on the signal normalization. Here, $\bar{\xi}$ is a parameter introduced to account for the fact that the normalization of the signal is typically obtained from previous experimental data, which is also subject to statistical fluctuations. In order to account for the associated uncertainty, for each pseudo-experiment the value of $\bar{\xi}$ is drawn from a normal distribution centered at zero and with a width equal to σ_{sys} , as for instance in Ref. [24]. For each pseudo-data realization we minimize the χ^2 in Eq. 24 with respect to $\sin^2 2\theta$ and Δm^2 and calculate a value for the test statistic $T = \chi^2(0, 0) - \chi_{\text{min}}^2$. From the ensemble of all simulated realizations we obtain then the expected distribution of T under the null-hypothesis of no oscillations.

Equation 23 will be our default χ^2 definition. But we have also studied the case where the Poisson χ^2 in Eq. 23 is replaced by other commonly used χ^2 definitions. The classical definition going back to Pearson [36] is

$$\chi_{\text{stat,Pearson}}^2(\theta, \Delta m^2) = \sum_i \left[\frac{(1 + \xi) p_i(\theta, \Delta m^2) - d_i}{\sqrt{p_i(\theta, \Delta m^2)}} \right]^2. \tag{25}$$

In order to avoid the parameter dependence in the denominator, the variance is often estimated by the data itself. We denote this version in the following as Gauss χ^2 :

$$\chi_{\text{stat,Gauss}}^2(\theta, \Delta m^2) = \sum_i \left[\frac{(1 + \xi) p_i(\theta, \Delta m^2) - d_i}{\sqrt{d_i}} \right]^2. \tag{26}$$

As we will show below, the Gaussian χ^2 can lead to different results for the distribution of T , while the Pearson definition leads to the same result as the Poisson case, Eq. 23, for sufficiently large event numbers per bin (as it is the case for the situations considered here).

For the results presented in the following we assume a single baseline setup. However, the arguments presented

in Sect. 3 apply also to multi-baseline configurations, for instance when ratios of spectra at different baselines are considered, or in case of segmented detectors with additional L information. The derivation in Sect. 3 relies only on general binning in (L/E) including also bins in L . We have verified explicitly that the simulation of a setup combining energy spectra at two different baselines leads to very similar results as the single-baseline configuration. This is also confirmed in Sect. 5, when we consider the Neutrino-4 experiment.

4.2 Results

Figure 2 presents the results of two simulations with different exposures: our default setup, with $\mathcal{N} = 1.5 \times 10^4$ total number of events (dark blue), and the same setup with 100 times more events, $\mathcal{N} = 1.5 \times 10^6$ (light blue). The left panel shows the distribution of the test statistic T . We observe a clear deviation from the χ^2 distribution, and a good agreement with the max. Gauss distribution derived in Sect. 3.

The agreement is excellent for the high statistics case and, in particular, we obtain the best match when the number of bins for the max. Gauss distribution is set at $N = 45$, to be compared with the 43 spectral bins used in the simulation. The reason for this (small) difference is that for a more realistic spectrum, some of the assumptions from Sect. 3 are only approximately fulfilled. In particular, assumption (c) (defined in Sect. 3.1) requires that relative statistical uncertainties are equal in all bins, which is obviously not true for a peaked spectrum as in reactor experiments. Therefore, $N = 45$ should be considered as the *effective* number of random standard normal variables, which leads to the best representation of the T distribution from simulation.

In Table 1 we show, for various values of the test statistic T , the significance which would be obtained by assuming a χ^2 distribution for 2 DOF (as we would expect if Wilks' theorem held) compared to the correct result following from the max. Gauss distribution for $N = 45$. For example, if a value of $T = 11.83$ is observed, we would exclude the SM at 3σ (p -value 0.27%) under the assumption of Wilks' theorem, while the correct significance would be only 2.48σ (p -value 1.3%). As a rule of thumb, we can see from the table that p -values are under-estimated by about a factor 5, and the number of σ gets reduced by roughly 0.5σ (except for low CL, where the difference is close to 1σ).

The right panel in Fig. 2 shows the distribution of the best-fit points obtained in the simulations. Although the pseudo-data has been generated under the no-oscillation hypothesis, we observe a clear preference for a non-vanishing value of $\sin^2 2\theta$. Note that actually *none* of the best fit points is located near the “true value” ($\sin^2 2\theta = 0$). Obviously $\widehat{\sin^2 2\theta}$ and $\widehat{\Delta m^2}$ are biased estimators in this case. Comparing the light and dark blue bands we confirm the scaling of the value of the

mixing angle at the best fit with the relative statistical uncertainty $1/\sqrt{\mathcal{N}}$, as expected from the discussion in Sect. 3. Indeed, in the region $0.5 \text{ eV}^2 \lesssim \Delta m^2 \lesssim 3 \text{ eV}^2$ the mean value of the $\sin^2 2\theta$ best-fit points agrees rather well with the prediction from Eq. 21, as indicated by the vertical dashed lines.³ Interestingly, this region also contains the highest concentration of best-fit points (indicated by the darker shading in each case), while it is more difficult to obtain a result favoring larger/smaller values of the mass splitting. However, outside this region of Δm^2 the best-fit points lie at larger values for the mixing angle. The reason is that for extreme values of Δm^2 the idealizations assumed in Sect. 3.1 do not apply. For example, the feature around $\Delta m^2 \sim 0.4 \text{ eV}^2$ corresponds to $\Delta m^2 L/(2E) \simeq \pi$ at $E \simeq 3 \text{ MeV}$ and, as a result, the first minimum of the survival probability is located at the peak of the event spectrum. This corresponds roughly to the case where half an oscillation period fits into the effective energy range, and therefore corresponds to the minimal frequency which can be sampled by the data. In contrast, for high mass-squared differences the frequency becomes much higher than the bin width can capture and therefore corresponds to over-sampling of the data.⁴ Hence, in both cases we are leaving the domain of the discrete parameterization of Δm^2 in terms of the index $\kappa = 1, \dots, N$ adopted in Sect. 3.1, see Eq. 8, which leads to the observed deviations with respect to the estimate in Eq. 21.

For comparison, we also show in the right panel of Fig. 2 the expected sensitivity at 95% CL under the assumption that Wilks' theorem holds (solid lines). They are obtained by using as “data” the no-oscillation prediction without statistical fluctuations (“Asimov data”) and considering contours of $\Delta\chi^2 = 5.99$. As can be seen from the figure, the best-fit points always lie very close to the expected sensitivity limit in this case and, for a sizeable fraction of the pseudo-experiments simulated, they lie *to the right* of the sensitivity curve, if naively computed assuming a $\Delta\chi^2$ for 2 DOF (as is usually the case in the literature).

Let us now discuss the impact of a systematic uncertainty on the overall normalization of the spectrum for the distribution of the test statistic T . Figure 3 shows the results obtained for different assumed priors on the systematic error for the signal, $\sigma_{\text{sys}} = 20\%$, 1% as well as the no-systematics case. In all cases we assume a total number of events of $\mathcal{N} = 1.5 \times 10^4$. We see from the figure that the distribution for the no-systematic case is somewhat χ^2 -like, with a number of DOF between 1 and 2. Although some deviations from this behaviour are observed (due to the effect of the physical boundary $\sin^2 2\theta \geq 0$ as well as the non-

³ We have verified that the range of values of Δm^2 where this is satisfied scales with the baseline as expected from $\Delta m^2 L = \text{const}$.

⁴ Let us note that the optimal bin width should be determined by the energy resolution of the detector.




Fig. 2 Left panel: Distribution of the test statistics obtained from numerical simulations for the toy reactor experiment described in the text. For comparison, the red-dashed curves shows the max. Gauss distribution for $N = 45$. Right panel: location of the best-fit points in the $\sin^2 2\theta - \Delta m^2$ plane, after minimization over nuisance parameters. In both panels, darker (lighter) blue lines/points correspond to the results obtained for $\mathcal{N} = 1.5 \times 10^4$ (1.5×10^6) events, using a sample of 20,000 pseudo-experiments simulated under the no-oscillation hypothesis. In the right panel, the regions with a higher density of best-fit

points are indicated by the darker shades. The dotted gray lines in the left panel show the χ^2 -distributions as the number of degrees of freedom is increased from 1 (lightest gray line to the left) to 5 (darkest gray line to the right). In the right panel, the vertical lines indicate the predicted value of $\langle \sin^2 2\theta \rangle$ from Eq. 21. The solid curves show the expected sensitivity at 95% CL assuming that Wilks' theorem holds. These results have been obtained using a Poisson χ^2 , with no background, for 10% signal systematics, and restricting $0 < \sin^2 2\theta < 1$ in the fit

Table 1 Comparison of the confidence level (CL), p -value, and corresponding number of standard deviations (σ), for several values of T , obtained for a χ^2 distribution with 2 DOF and for the max. Gauss distribution for $N = 45$

T	CL (%)		p -value (%)		Number of σ	
	$\chi^2(2)$	Max. G.	$\chi^2(2)$	Max. G.	$\chi^2(2)$	Max. G.
4.61	90.00	48.55	10.0	51.4	1.64	0.65
6.18	95.45	74.73	4.55	25.3	2.00	1.14
9.21	99.00	94.72	1.00	5.27	2.58	1.94
9.49	99.13	95.45	0.87	4.55	2.62	2.00
11.83	99.73	98.69	0.27	1.31	3.00	2.48
14.78	99.938	99.73	0.062	0.27	3.42	3.00

linearity of the model), the distribution is clearly different from a max. Gauss. The reason is that in Sect. 3 we assumed that only shape information is used (assumption (b)), whereas in the absence of systematic errors the information on the total event rate is also available. In this case the model of fitting white noise, Eq. 9, does not fully correspond to fitting the disappearance probability in Eq. 1, which can only reduce the event numbers. In contrast, for the case $\sigma_{\text{sys}} = 20\%$ the systematic uncertainty is much larger than the statistical one for the assumed event sample, $\sigma_{\text{sys}} \gg 1/\sqrt{\mathcal{N}}$. This corresponds effectively to a free normalization in the fit and assumption (b) is satisfied. Correspondingly, we observe in Fig. 3 a very good agreement with the max. Gauss distribution for this case. Note that also for the 10% systematic assumed in Fig. 2 we have $\sigma_{\text{sys}} \gg 1/\sqrt{\mathcal{N}}$, such that assump-

tion (b) defined in Sect. 3.1 is fulfilled. For the 1% case we have $\sigma_{\text{sys}} \simeq 1/\sqrt{\mathcal{N}}$, which corresponds to an intermediate situation between fixed and free normalization.

Figure 3 also shows the impact due to the treatment of systematics when simulating the random pseudo-data. Solid curves show the results obtained randomizing the central value of the pull parameter, i.e., for each realization of the pseudo-data we draw $\bar{\xi}$ in Eq. 24 from a Gaussian distribution with width σ_{sys} (as outlined in Sect. 4.1). In contrast, for the dash-dotted curves the central value for the pull parameter is not randomized and kept fixed at $\bar{\xi} = 0$ for all pseudo-data samples. We see that this has a rather large impact on the \sqrt{T} distribution as long as systematic and statistical uncertainties are comparable ($\sigma_{\text{sys}} = 1\%$), whereas for an effectively free normalization ($\sigma_{\text{sys}} = 20\%$) the difference is




Fig. 3 Impact of a systematic uncertainty on the overall normalization. We show the distribution of the square-root of the test statistic, \sqrt{T} , for no systematic uncertainty, 1, and 20% uncertainty. The total number of events is $\mathcal{N} = 1.5 \times 10^4$. For the solid curves we randomize the central values for the systematic for each draw of the pseudo-data, whereas for the dash-dotted curves the central value is kept fixed. The black-dashed curves corresponds to the max. Gauss distribution for $N = 45$. The gray dotted lines show the χ^2 -distributions as the number of degrees of freedom is increased from 1 (lightest gray line to the left) to 5 (darkest gray line to the right)

largely reduced. The reason is that in the latter case the fit can always adjust the normalization within the statistical uncertainty, with negligible impact of the penalty term for the pull parameter.

4.3 Further studies of the properties of the T distribution

In this subsection we investigate in some detail additional properties of the distribution of the test statistic T . We start by discussing the impact of the two χ^2 implementations from Eq. 23 (Poisson χ^2) versus Eq. 26 (Gauss χ^2). Naively one expects that they should give similar results if the number of events per bin is $\gtrsim 10$. The differences on the resulting distributions for the two χ^2 implementations are shown in Fig. 4.

In the left panel of Fig. 4 we adopt different choices for the normalization uncertainty. Interestingly, we find for $\sigma_{\text{sys}} \lesssim 2\%$ notably differences, despite the rather large event number of $\mathcal{N} = 1.5 \times 10^4$. For the 43 bins in our simulation this corresponds to about 350 events per bin on average, where the bin with the smallest number of events has a mean above 20 events. We have checked that for the no-systematic case the differences between Poisson and Gauss disappear only for $\mathcal{N} \gtrsim 10^5$. From the figure we also see that for $\mathcal{N} = 1.5 \times 10^4$ the differences between the Gauss and Poisson implementations disappear for large enough systematic uncertainty, when both cases approach the max. Gauss distribution. The origin of the different behavior is related to the assumption $\sigma_i = \sqrt{d_i}$ in Eq. 26. We have confirmed that when we use instead the Pearson definition, Eq. 25, with

$\sigma_i = \sqrt{p_i^0}$, the Pearson and Poisson χ^2 implementations lead to identical results.

In the right panel of Fig. 4 we study the impact of the physical boundary $\sin^2 2\theta \geq 0$ in the case of $\mathcal{N} = 1.5 \times 10^4$ and no systematic uncertainty on the normalization. Note that, while only $\sin^2 2\theta \geq 0$ makes sense from the mathematical point of view, this parameter controls the amplitude of the oscillation and in principle it is possible to try to fit data without taking this requirement into account. Thus, in the figure we compare three cases: $\sin^2 2\theta \geq 0$ (blue lines), $\sin^2 2\theta \leq 0$ (green lines), and a third case where no restriction is imposed on $\sin^2 2\theta$ (red lines). For the Poisson χ^2 implementation we see that restricting the sign of $\sin^2 2\theta$ has a notable impact on the distribution, but the effect is similar regardless of the sign of $\sin^2 2\theta$. In contrast, for the Gauss χ^2 implementation we observe significant differences between the cases $\sin^2 2\theta \geq 0$ and $\sin^2 2\theta \leq 0$. Again this is a consequence of using $\sqrt{d_i}$ as the statistical uncertainty in the Gauss χ^2 , Eq. 26: as d_i includes statistical fluctuations, $\sqrt{d_i}$ is not symmetric between upward and downward fluctuations, which leads to the asymmetric behavior with respect to the sign of $\sin^2 2\theta$. In contrast, if the theoretical prediction is used as variance as in the Pearson definition, Eq. 25, the χ^2 becomes symmetric between upward and downward fluctuations. We have explicitly checked that in this case the dependence on the sign of $\sin^2 2\theta$ disappears and we recover the result from the Poisson χ^2 . Surprisingly, these second order effects are not negligible even for $\mathcal{N} = 1.5 \times 10^4$ events.

Overall, we observe that the results agree with the Poisson χ^2 if we use the Pearson χ^2 , where the square-root of the prediction to calculate the statistical uncertainty, while sizeable deviations occur for the Gaussian χ^2 , where the square-root of the data is used as statistical uncertainty. Let us remark, however, that as long as the distribution of the test statistic is numerically evaluated by Monte Carlo simulation, of course any reasonable χ^2 definition can be used (including also the Gauss χ^2 as defined in Eq. 26).

To summarize the results found in this section, we find that as long as $\sigma_{\text{sys}} \lesssim 1/\sqrt{\mathcal{N}}$, the distribution of the test statistic T is sensitive to details of the analysis, such as size of systematics, treatment of systematics during randomization, χ^2 variants, physical boundaries. However, once $\sigma_{\text{sys}} \gg 1/\sqrt{\mathcal{N}}$ (i.e., for experiments where only shape information is used) the max. Gauss distribution seems to be a rather robust result.

5 Application: Neutrino-4 as a case study

The Neutrino-4 experiment [16,30] has recently claimed a possible indication of sterile neutrino oscillations with $\Delta m^2 \simeq 7.2 \text{ eV}^2$ and $\sin^2 2\theta \simeq 0.26$. They report a statistical significance using their combined phase 1 and 2 data

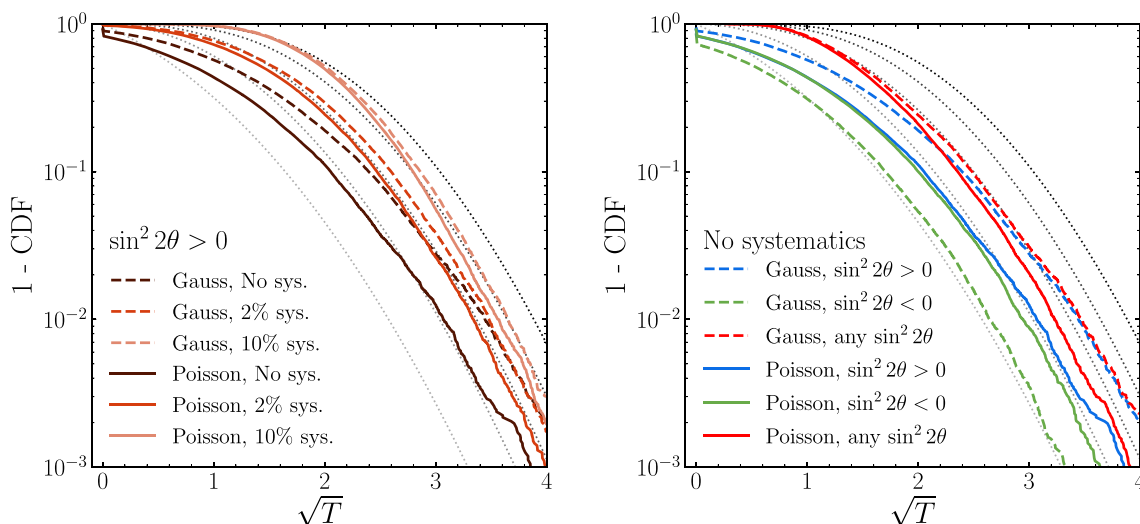


Fig. 4 Impact on the distribution of \sqrt{T} due to the choice of the χ^2 implementation. Solid curves correspond to a Poisson χ^2 , while dashed lines correspond to a Gauss χ^2 . The left panel shows the effect of increasing the systematic uncertainties from no systematics to an overall 10% signal normalization uncertainty. In the right panel we show

the results by imposing different restrictions on the allowed range for $\sin^2 2\theta$, as indicated by the labels. In all cases, a total of 1.5×10^4 events are simulated for each pseudo-experiment. Dotted curves indicate the χ^2 distribution for 1 to 5 DOF from left to right

of 3.2σ [30]. Then an estimate of their systematic uncertainty is quoted, leading to a combined statistical/systematical significance of a positive $\sin^2 2\theta$ of 2.8σ . In this section we use the Neutrino-4 results to illustrate the arguments presented above on a real-life example. We concentrate on the purely statistical aspect and will show that, in light of the discussions in the previous sections, the significance from statistical uncertainties alone is already lower than the quoted 3.2σ .

Neutrino-4 uses a segmented detector, which allows to bin their data in both L and E . The data is binned using 9 bins in energy with width $\Delta E = 0.5$ MeV starting at 2.3 MeV, and 24 bins in baseline with $\Delta L = 0.235$ m starting at 6.25 m, resulting into a total of 216 bins in L/E . The bin width in energy of 0.5 MeV corresponds to the energy resolution of the detector [30]. Eventually, each consecutive group of 8 bins are combined together, leading to $N = 27$ data points. The observed data correspond to the ratio

$$R_i^{\text{obs}} = \frac{d_i}{\frac{1}{N} \sum_{i=1}^N d_i}, \tag{27}$$

where d_i stands for the observed number of events in bin i . Out of these, the first 19 bins are shown for the combined phase 1 and phase 2 data sets in Fig. 47 of [30] (blue points). Following the Neutrino-4 collaboration, we fit these 19 data points with the survival probability for a given L/E bin over the averaged probability:

$$R_i^{\text{pred}} = \frac{1 - \sin^2 2\theta \left\langle \sin^2 \frac{\Delta m^2 L}{4E} \right\rangle_i}{1 - \frac{1}{2} \sin^2 2\theta}. \tag{28}$$

Here $\langle \cdot \rangle_i$ indicates the average over an energy interval ΔE , with the value of L/E set at the bin center of the corresponding (L/E) -bin i . The fit is performed with a simple Gaussian χ^2 definition, using the statistical uncertainties read off from Fig. 47 of [30]. Note that, due to the particular way the fit is performed by Neutrino-4, using the ratios in Eqs. 27, 28 the analysis is only sensitive to spectral distortions in L/E , and therefore assumption (b) from Sect. 3.1 is fulfilled.

With our fit we can reproduce to good accuracy the results from Ref. [30]. Our best-fit point is located at $\Delta m^2 = 8.84$ eV², $\sin^2 2\theta = 0.42$; however, we find a quasi-degenerate local minimum with $\Delta\chi^2 = 5 \times 10^{-3}$ at $\Delta m^2 = 7.28$ eV², $\sin^2 2\theta = 0.34$, close to the best-fit point obtained by Neutrino-4. We explain this slight difference by the fact that the fit reported in Ref. [30] uses more information in L/E than available to us. This additional information seems to somewhat disfavor the local minimum around $\Delta m^2 \simeq 9$ eV² compared to the one at $\simeq 7.25$ eV². Furthermore, we obtain for the χ^2 minimum and the test statistic T , i.e., the $\Delta\chi^2$ between no oscillations and the best-fit point:

$$\chi_{\text{min}}^2 = 16.05, \quad T = 12.94 \quad (\text{our result}), \tag{29}$$

$$\chi_{\text{min}}^2 = 17.11, \quad T = 12.87 \quad (\text{Fig. 47 of Ref. [30]}), \tag{30}$$

showing good agreement, especially for T .

If evaluated under the assumption of Wilks' theorem with a χ^2 distribution with 2 DOF we would get from $T = 12.9$ a p -value of 1.58×10^{-3} , corresponding to 3.16σ . In order to check this reasoning, we have generated a large sample of artificial data sets for Neutrino-4, under the null-hypothesis

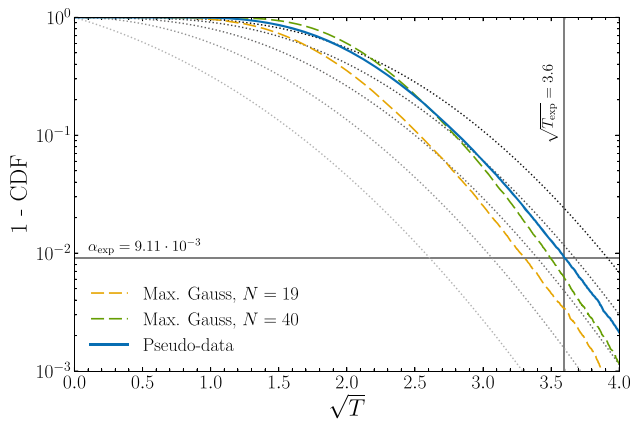


Fig. 5 Distribution of the square-root of the test statistics \sqrt{T} for the Neutrino-4 experiment, obtained from the simulation of 100,000 pseudo-data sets under the assumption of no oscillations. The numerical result (solid blue curve) is compared to the expectation for a χ^2 distribution with 1...5 DOF (dotted curves, from left to right), as well as to the distribution for the max. Gauss distribution with 19 and 40 bins (dashed curves). The vertical line indicates the value of $\sqrt{T} = 3.6$ obtained from the observed Neutrino-4 data, whereas the horizontal line shows the corresponding p -value

of no oscillations, in order to calculate the distribution of T explicitly. The result is shown in Fig. 5, which shows significant deviations from a χ^2 -distribution. In agreement with the discussions above, the distribution of \sqrt{T} is found to be more similar to the max. Gauss distribution. In this case we find that the effective $N = 40$ for the max. Gauss distribution providing the closest fit to the numerical distribution deviates substantially from the actual number of bins in the data, $N = 19$. Based on the numerical T distribution we obtain a p -value of 9.1×10^{-3} (or 2.6σ), indicating that the actual statistical significance is clearly lower. Note that here we study the significance based on statistical uncertainties only. Systematic effects as mentioned in Ref. [30] will reduce the significance even further. See also the discussion in Ref. [29] in this context.

In addition, we have calculated confidence regions in the plane of $\sin^2 2\theta$ and Δm^2 by performing a Feldman–Cousins analysis [31]. For given values of $\sin^2 2\theta$ and Δm^2 we have generated many artificial data sets, assuming that these values are the true values in Nature. This allows us to compute the correct distribution of

$$\Delta\chi^2(\sin^2 2\theta, \Delta m^2) = \chi^2(\sin^2 2\theta, \Delta m^2) - \chi_{\min}^2 \quad (31)$$

for each point in the parameter space. Comparing the value of $\Delta\chi_{\text{exp}}^2$ obtained from the actual experimental data to the numerical distribution for $\Delta\chi^2(\sin^2 2\theta, \Delta m^2)$, we obtain the confidence level (CL) at which a particular point can be rejected. Repeating this procedure for the whole parameter space we obtain confidence regions at a given CL as the set of all points which are accepted at that CL.

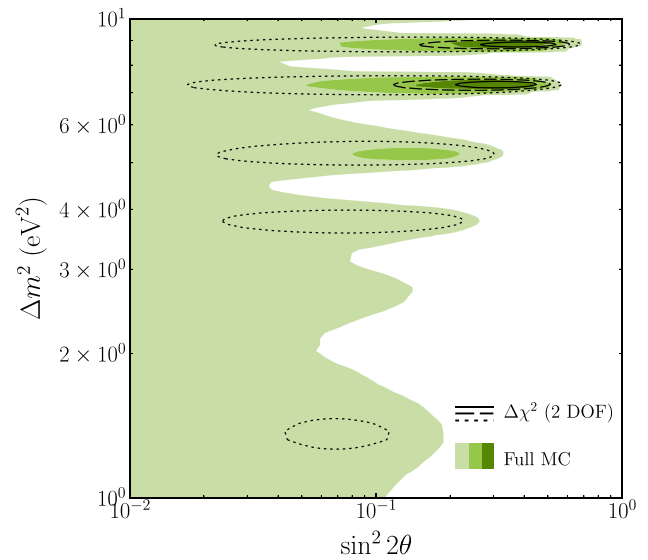


Fig. 6 Confidence regions from our re-analysis of Neutrino-4 data [30] at 68.3% (dark green), 95.45% (medium green), and 99.73% CL (light green). Shaded regions correspond to the confidence regions constructed by Monte-Carlo simulations following the Feldman–Cousins prescription [31], whereas black curves show the corresponding CL contours in $\Delta\chi^2$ assuming it follows a χ^2 -distribution for 2 DOF, that is, $\Delta\chi^2 = 2.3$ (solid), 6.18 (dashed), 11.83 (dotted)

The results of this analysis are shown in Fig. 6 as shaded regions for 68.3% (dark green), 95.45% (medium green), and 99.73% (light green) CL. Our regions are also compared to $\Delta\chi^2$ contours obtained under the assumption that $\Delta\chi^2$ follows a χ^2 distribution with 2 DOF, which would be the case if Wilks’ theorem held. We clearly observe that true confidence regions are substantially larger than the ones based on the χ^2 distribution. In particular, the no-oscillation case is contained in the 3σ contour for the Monte-Carlo calculation, in agreement with the discussion of the test statistic T above.⁵

6 Summary and conclusions

In this paper we have studied the statistical interpretation of sterile neutrino oscillation searches in the disappearance mode, specifically when no information on the absolute normalization of the signal is used. A priori there are several good reasons to expect that Wilks’ theorem does not apply in this case: the presence of a physical boundary, the fact that the parameter space changes dimension if either $\Delta m^2 \rightarrow 0$

⁵ Let us note that our $\Delta\chi^2$ contours are also somewhat larger than the ones shown in Fig. 45 of Ref. [30]. We believe that the reason for this difference is that contours in Ref. [30] are drawn for a χ^2 distribution with 1 DOF, while ours are shown for 2 DOF. We have checked that using the same prescription we can reproduce their regions with good accuracy.

or $\sin^2 2\theta \rightarrow 0$, and the highly non-linear dependence of the number of events on Δm^2 . Not surprisingly, we do indeed find significant deviations. Although in this work we decided to focus on short-baseline reactor experiments as a case study, our results are more general. We find that this situation, under some assumptions, is equivalent to fitting Gaussian white noise with a single frequency of free amplitude. This allows us to express the distribution of the test statistic T to be the maximum of N Gaussian random variables where N is the effective number of bins (“max. Gauss distribution”). Therefore, this class of oscillation searches will always find a best-fit for a non-zero signal even if there is no oscillation in the data, with a non-negligible statistical significance if interpreted as if Wilks’ theorem would apply. In other words, the parameters obtained at the minimum of the χ^2 are biased estimators in this case.

We then perform Monte Carlo simulations of a toy reactor disappearance experiment, to confirm that our analytic understanding carries over to a more realistic setting. The test statistic T we consider is equivalent to the log-likelihood either for a Gaussian likelihood or Poissonian likelihood. We find that, if the systematic uncertainty on the event normalization is comparable to (or smaller than) the statistical uncertainty of the event sample, the distribution function of T is rather sensitive on fine details of the chosen simulation and, in particular, on: whether the central value of the nuisance parameter is randomized or not, and whether a Gaussian or a Poissonian log-likelihood is used (despite the fairly large number of events per bin). Conversely, for experiments relying on shape information only (that is, when no information on the absolute normalization is used) the max. Gauss distribution is a rather robust prediction for the distribution of the test statistic. Although the shape of the max. Gauss does not depend strongly on the value of N (only logarithmically), we have not found a simple way to predict the value that provides the best description of the distribution of the test statistic.

Finally, we apply our understanding to the actual data of the Neutrino-4 experiment. We are able to reproduce the quantitative details of their analysis quite well if we assume that Wilks’ theorem applies. However, in agreement with our arguments presented above we find that the test statistics shows significant deviations from a χ^2 distribution. In particular, we show by explicit Monte Carlo simulation that the significance of the claimed oscillation signal is reduced from 3.2σ ($p = 1.58 \times 10^{-3}$) to 2.6σ ($p = 9.1 \times 10^{-3}$), that is, the probability that this is a mere statistical fluctuation is about 6 times larger than that expected if Wilks’ theorem were to hold. It should be noted that our Neutrino-4 analysis is based on statistical uncertainties only, and that the inclusion of systematic effects may reduce the significance even further.

In summary, our results provide a simple, intuitive understanding on why and how shape-only oscillation searches

are different from the usual case. Applied to Neutrino-4 we find a reduced significance for sterile neutrino oscillation, but not to the extent to completely dismiss this indication as a pure statistical fluctuation. It would be interesting to see how this type of analysis would play out in a global fit of all short-baseline reactor data, but this is beyond the scope of the present work.

Acknowledgements The authors warmly thank Mattias Blennow and Enrique Fernandez-Martinez for useful discussions. PH acknowledges support from the U.S. Department of Energy Office of Science under contract DE-SC0018327. PC acknowledges support from the Grant PROMETEO/2019/083, from the Spanish MICINN through the “Ramón y Cajal” program under Grant RYC2018-024240-I, and from the Spanish Agencia Estatal de Investigación through Grant “IFT Centro de Excelencia Severo Ochoa SEV-2016-0597”. The authors also acknowledge use of the HPC facilities at the IFT (Hydra cluster) and IFIC (SOM cluster). This work was partially supported by the European projects H2020-MSCA-ITN-2015//674896-ELUSIVES and 690575-InvisiblesPlus-H2020-MSCA-RISE-2015.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: Data will be shared by the authors upon request.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. Funded by SCOAP³.

References

1. LSND, A. Aguilar-Arevalo et al., Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Phys. Rev. D* **64** (2001) 112007. <https://doi.org/10.1103/PhysRevD.64.112007>. [arXiv:hep-ex/0104049](https://arxiv.org/abs/hep-ex/0104049)
2. S. Böser, C. Buck, C. Giunti, J. Lesgourgues, L. Ludhova et al., Status of light sterile neutrino searches. *Prog. Part. Nucl. Phys.* **111**, 103736 (2020). <https://doi.org/10.1016/j.pnpnp.2019.103736>. [arXiv:1906.01739](https://arxiv.org/abs/1906.01739)
3. C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, A.D. McGuire, Detection of the free neutrino: a confirmation. *Science* **124**(3212), 103–104 (1956). <https://doi.org/10.1126/science.124.3212.103>
4. T.A. Mueller et al., Improved predictions of reactor antineutrino spectra. *Phys. Rev. C* **83**, 054615 (2011). <https://doi.org/10.1103/PhysRevC.83.054615>. [arXiv:1101.2663](https://arxiv.org/abs/1101.2663)
5. P. Huber, On the determination of anti-neutrino spectra from nuclear reactors. *Phys. Rev. C* **84**, 024617 (2011). <https://doi.org/10.1103/PhysRevC.84.024617>. [arXiv:1106.0687](https://arxiv.org/abs/1106.0687) [Erratum: *Phys. Rev. C* **85**, 029901 (2012). <https://doi.org/10.1103/PhysRevC.85.029901>]

6. G. Mention, M. Fechner, T. Lasserre, T.A. Mueller, D. Lhuillier et al., The reactor antineutrino anomaly. *Phys. Rev. D* **83**, 073006 (2011). <https://doi.org/10.1103/PhysRevD.83.073006>. arXiv:1101.2755
7. J.M. Berryman, P. Huber, Reevaluating reactor antineutrino anomalies with updated flux predictions. *Phys. Rev. D* **101**(1), 015008 (2020). <https://doi.org/10.1103/PhysRevD.101.015008>. arXiv:1909.09267
8. C. Giunti, Y. Li, B. Littlejohn, P. Surukuchi, Diagnosing the reactor antineutrino anomaly with global antineutrino flux data. *Phys. Rev. D* **99**(7), 073005 (2019). <https://doi.org/10.1103/PhysRevD.99.073005>. arXiv:1901.01807
9. Daya Bay, F.P. An et al., Improved search for a light sterile neutrino with the full configuration of the Daya Bay experiment. *Phys. Rev. Lett.* **117**(15), 151802 (2016). <https://doi.org/10.1103/PhysRevLett.117.151802>. arXiv:1607.01174
10. RENO, J. Choi et al., Search for sub-eV sterile neutrino at RENO. *Phys. Rev. Lett.* **125**(19), 191801 (2020)
11. Double Chooz, H. de Kerret et al., First Double Chooz θ_{13} measurement via total neutron capture detection. *Nat. Phys.* **16**(5), 558–564 (2020). <https://doi.org/10.1038/s41567-020-0831-y>. arXiv:1901.09445
12. Y. Ko et al., Sterile neutrino search at the NEOS experiment. *Phys. Rev. Lett.* **118**(12), 121802 (2017). <https://doi.org/10.1103/PhysRevLett.118.121802>. arXiv:1610.05134
13. DANSS, I. Alekseev et al., Search for sterile neutrinos at the DANSS experiment. *Phys. Lett. B* **787**, 56–63 (2018). <https://doi.org/10.1016/j.physletb.2018.10.038>. arXiv:1804.04046
14. PROSPECT, M. Andriamirado et al., Improved short-baseline neutrino oscillation search and energy spectrum measurement with the PROSPECT experiment at HFIR. arXiv:2006.11210
15. STEREO, H. Almazán Molina et al., Improved sterile neutrino constraints from the stereo experiment with 179 days of reactor-on data. *Phys. Rev. D* **102**(5), 052002 (2020)
16. NEUTRINO-4, A. Serebrov et al., First observation of the oscillation effect in the Neutrino-4 experiment on the search for the sterile neutrino. *Pisma Zh. Eksp. Teor. Fiz.* **109**(4), 209–218 (2019). <https://doi.org/10.1134/S0021364019040040>. arXiv:1809.10561
17. M. Dentler, A. Hernandez-Cabezudo, J. Kopp, M. Maltoni, T. Schwetz, Sterile neutrinos or flux uncertainties?—Status of the reactor anti-neutrino anomaly. *JHEP* **11**, 099 (2017). [https://doi.org/10.1007/JHEP11\(2017\)099](https://doi.org/10.1007/JHEP11(2017)099). arXiv:1709.04294
18. M. Dentler, A. Hernandez-Cabezudo, J. Kopp, P.A.N. Machado, M. Maltoni et al., Updated global analysis of neutrino oscillations in the presence of eV-scale sterile neutrinos. *JHEP* **08**, 010 (2018). [https://doi.org/10.1007/JHEP08\(2018\)010](https://doi.org/10.1007/JHEP08(2018)010). arXiv:1803.10661
19. S. Gariazzo, C. Giunti, M. Laveder, Y.F. Li, Model-independent $\bar{\nu}_e$ short-baseline oscillations from reactor spectral ratios. *Phys. Lett. B* **782**, 13–21 (2018)
20. J.M. Berryman, P. Huber, Sterile neutrinos and the global reactor antineutrino dataset. arXiv:2005.01756
21. S. Wilks, The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math. Stat.* **9**(1), 60–62 (1938). <https://doi.org/10.1214/aoms/1177732360>
22. T. Schwetz, What is the probability that theta(13) and CP violation will be discovered in future neutrino oscillation experiments? *Phys. Lett. B* **648**, 54–59 (2007). <https://doi.org/10.1016/j.physletb.2007.02.053>. arXiv:hep-ph/0612223
23. M. Blennow, P. Coloma, P. Huber, T. Schwetz, Quantifying the sensitivity of oscillation experiments to the neutrino mass ordering. *JHEP* **03**, 028 (2014). [https://doi.org/10.1007/JHEP03\(2014\)028](https://doi.org/10.1007/JHEP03(2014)028). arXiv:1311.1822
24. M. Blennow, P. Coloma, E. Fernandez-Martínez, Reassessing the sensitivity to leptonic CP violation. *JHEP* **03**, 005 (2015). [https://doi.org/10.1007/JHEP03\(2015\)005](https://doi.org/10.1007/JHEP03(2015)005). arXiv:1407.3274
25. J. Elevant, T. Schwetz, On the determination of the leptonic CP phase. *JHEP* **09**, 016 (2015). [https://doi.org/10.1007/JHEP09\(2015\)016](https://doi.org/10.1007/JHEP09(2015)016). arXiv:1506.07685
26. I. Esteban, M. Gonzalez-Garcia, M. Maltoni, I. Martínez-Soler, T. Schwetz, Updated fit to three neutrino mixing: exploring the accelerator-reactor complementarity. *JHEP* **01**, 087 (2017). [https://doi.org/10.1007/JHEP01\(2017\)087](https://doi.org/10.1007/JHEP01(2017)087). arXiv:1611.01514
27. M. Agostini, B. Neumair, Statistical methods applied to the search of sterile neutrinos. *Eur. Phys. J. C* **80**(8), 750 (2020)
28. C. Giunti, Statistical significance of reactor antineutrino active-sterile oscillations. *Phys. Rev. D* **101**(9), 095025 (2020). <https://doi.org/10.1103/PhysRevD.101.095025>. arXiv:2004.07577
29. PROSPECT, STEREO, M. Andriamirado et al., Note on arxiv:2005.05301. Preparation of the Neutrino-4 experiment on search for sterile neutrino and the obtained results of measurements. arXiv:2006.13147
30. A. Serebrov et al., Preparation of the Neutrino-4 experiment on search for sterile neutrino and the obtained results of measurements. arXiv:2005.05301 (arXiv version 2)
31. G.J. Feldman, R.D. Cousins, A unified approach to the classical statistical analysis of small signals. *Phys. Rev. D* **57**, 3873–3889 (1998). <https://doi.org/10.1103/PhysRevD.57.3873>. arXiv:physics/9711021
32. S. Algeri, J. Aalbers, K. Dundas Morã, J. Conrad, Searching for new phenomena with profile likelihood ratio tests. *Nat. Rev. Phys.* **2**(5), 245–252 (2020). <https://doi.org/10.1038/s42254-020-0169-5>
33. S. Coles, *An Introduction to Statistical Modeling of Extreme Values* (Springer, Berlin, 2001)
34. D. Pfeifer, Einführung in die Extremwertstatistik. Teubner-Skripten zur mathematischen Stochastik (1989)
35. S. Baker, R.D. Cousins, Clarification of the use of chi square and likelihood functions in fits to histograms. *Nucl. Instrum. Methods* **221**, 437–442 (1984). [https://doi.org/10.1016/0167-5087\(84\)90016-4](https://doi.org/10.1016/0167-5087(84)90016-4)
36. K. Pearson, X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philos. Mag.* V **50**(302), 157–175 (1900). <https://doi.org/10.1080/14786440009463897>