



# Alternative to the application of PDG scale factors

Jens Erler<sup>1,2,a</sup>, Rodolfo Ferro-Hernández<sup>1,b</sup>

<sup>1</sup> Departamento de Física Teórica, Instituto de Física, Universidad Nacional Autónoma de México, 04510 Mexico, Mexico

<sup>2</sup> PRISMA+ Cluster of Excellence and Helmholtz Institute Mainz, Johannes Gutenberg-Universität, 55099 Mainz, Germany

Received: 16 April 2020 / Accepted: 3 June 2020 / Published online: 16 June 2020  
© The Author(s) 2020

**Abstract** The Particle Data Group recommends a set of procedures to be applied when discrepant data are to be combined. We introduce an alternative method based on a more general and solid statistical framework, providing a robust way to include possible unknown systematic effects interfering with experimental measurements or their theoretical interpretation. The limit of large data sets and practical cases of interest are discussed in detail.

## 1 Introduction

In any field of science, it is often the case that a number of data points or data sets need to be combined in order to achieve a greater overall precision. Now, data naturally fluctuate and it is not uncommon that one or several data points may appear discrepant or outlying with respect to the bulk of the data. This is not necessarily a concern, e.g., if the results of the individual measurements or observations are known to be dominated by the statistical uncertainty, or even in the presence of significant systematic effects, as long as their associated uncertainties can be reliably estimated. On the other hand, if the observed discrepancies are suspiciously large or plentiful, one may worry that some unknown systematic effect or unjustified but hidden assumption might have moved the central value of one or more observations. In that latter case, a more conservative handling of the data and its combination would be called for.

Of course, it is impossible to know independently which of the aforementioned situations—larger than expected random fluctuations, unknown systematic effect(s), or both—one is facing, or which of the individual data (sub)sets could be

at fault. As a remedy, the Particle Data Group (PDG) [1] proposed a set of rules according to which the uncertainty of an average is to be enlarged by a scale factor  $S$ , while the central values are to remain unchanged by fiat. Assuming Gaussian errors, in a first step the reduced  $\chi^2$  is computed as twice the log-likelihood of the minimum divided by  $N_{\text{eff}}$ , where  $N_{\text{eff}}$  is the effective number of degrees of freedom given by the number of observations (data points),  $N$ , minus the number of independent fit parameters. Thus, for the most common case of a simple average of one parameter,  $N_{\text{eff}} = N - 1$ :

1. If the reduced  $\chi^2$  is smaller than unity, the results are accepted and there is no scaling of errors.
2. If the reduced  $\chi^2$  is larger than unity, and the experiments are of comparable precision, then all errors are re-scaled by a common factor  $S$ , given by the reduced  $\chi^2$ , i.e.,  $S = \sqrt{\chi^2/N_{\text{eff}}}$ .
3. If some of the individual errors are much smaller than others, then  $S$  is computed from only the most precise experiments. The criterium for these is given with reference to an ad hoc cutoff value.

Given that the rationale for a procedure such as this one, is to err on the conservative side, one immediate objection is that if there is only one data point then no conservative scaling will be applied, even though in this case one is most exposed to a potential problem as there is no control measurement.

Another problem is that the set of individual data points is not well-defined. In principle, one may combine certain data subsets first, such as from different data taking periods or different decay channels obtained by the same experimental apparatus, or combine identical channels obtained by different detectors and average these is a second step. Conversely, one could split up the available results into more but less precise individual entries. While this has no impact on ordinary maximum likelihood analyses, it will generally dilute or enlarge the reduced  $\chi^2$  value on which the  $S$  factors

<sup>1</sup> The PDG collects, evaluates, averages and fits particle physics data world-wide and assesses their implications and interpretations in a large number of dedicated reviews.

<sup>a</sup> e-mail: [erler@fisica.unam.mx](mailto:erler@fisica.unam.mx) (corresponding author)

<sup>b</sup> e-mail: [rferrohernandez00@gmail.com](mailto:rferrohernandez00@gmail.com)

are based upon. In fact, applying PDG scale factors to data points of which some have already undergone the scale factor treatment (typically, by the experimental collaboration) then this kind of iteration *does* generally change the central value of the combination. Also note that the prescription according to which reduced  $\chi^2$  values greater and smaller than unity are being treated differently generates an unnecessary dichotomy.

In this paper we present an alternative which shares some of the features of the PDG recommendation while improving on others. The framework is a hierarchical model within Bayesian parameter inference [2]. The basic idea is that individual data points are not considered independently and identically distributed (*iid*), but rather independently and similarly distributed, in the sense that the parent distributions are permitted to vary to some extent to allow for unknown effects that may or may not be different from one data point (measurement) to another. Thus, we propose a hierarchical model where each measurement is assumed to determine a different parameter, each considered as having arisen as a random draw from a common parent distribution described in turn in terms of *hyper-parameters*.

A similar approach is widely used in the biological sciences when estimating treatment effects by combining several studies performed under similar but not identical conditions [3,4], in what is often referred to as *meta-analysis* [5–7]. In these cases the experimental conditions can vary slightly, so that the individual studies may be affected by different unknown biases.


Several authors within the physics community introduced attempts to incorporate the effects of unknown error sources when combining data. For example, Ref. [8] finds results similar to the ones in our work, but within a frequentist approach. Ref. [9] models the probability of underestimating the experimental error by including a different scale factor for each measurement, which is in turn randomly drawn from a prior distribution. Very recently it was shown [10] that it is even possible to test the shape of the prior distribution, and not just to constrain the values of its parameters. We leave this kind of more complete analysis for the future.

In the next section we summarize the formalism of Bayesian hierarchical modeling using the notation of Ref. [2]. The rest of the paper introduces our approach, illustrated by a number of examples and reference cases.

## 2 Bayesian inference

### 2.1 The non-hierarchical model

Suppose that we want to determine a parameter  $\theta$  from an experimental measurement or observation, and to be specific, that the likelihood for the outcome  $y$  of such an experiment



**Fig. 1** Ordinary averaging. We assume that the  $y_i$  are random outcomes of measurements of the *same* parameter  $\theta$

can be described as a Gaussian with central value  $\theta$  and standard deviation  $\sigma$ ,

$$p(y|\theta, \sigma) = \mathcal{N}(y|\theta, \sigma), \tag{1}$$

where,

$$\mathcal{N}(y|\theta, \sigma) \equiv \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2}(y-\theta)^2}. \tag{2}$$

The posterior distribution for the parameter  $\theta$  can be obtained through Bayes' theorem,

$$p(\theta|y, \sigma) \propto p(y|\theta, \sigma)p(\theta), \tag{3}$$

where  $p(\theta)$  is the *prior probability distribution* of  $\theta$ . It is very convenient to assume  $p(\theta)$  to be a *conjugate prior*, which means that the posterior distribution will fall within the same family of functions as the prior. Thus, in our case we adopt the prior,

$$\theta \sim \mathcal{N}(\tilde{\mu}, \tilde{\tau}), \tag{4}$$

yielding the posterior,

$$p(\theta|y, \sigma, \tilde{\mu}, \tilde{\tau}) = \frac{1}{\sqrt{2\pi}\sigma_{\tilde{\tau}}} e^{-\frac{1}{2\sigma_{\tilde{\tau}}^2}(\theta-\theta_{\tilde{\tau}})^2}, \tag{5}$$

where,

$$\frac{1}{\sigma_{\tilde{\tau}}^2} \equiv \frac{1}{\sigma^2} + \frac{1}{\tilde{\tau}^2}, \tag{6}$$

is the sum of precisions of the prior and the experimental result, while

$$\theta_{\tilde{\tau}} \equiv \left( \frac{1}{\sigma^2} + \frac{1}{\tilde{\tau}^2} \right)^{-1} \left( \frac{y}{\sigma^2} + \frac{\tilde{\mu}}{\tilde{\tau}^2} \right), \tag{7}$$

is the precision averaged central value. Clearly, if the experiment has a small error,  $\sigma \ll \tilde{\tau}$ , it will dominate  $\theta_{\tilde{\tau}}$ . In the limit  $\tilde{\tau} \rightarrow \infty$ , the prior is called *non-informative*.

Now, let us include further such experiments with central values  $y_i$  and total errors  $\sigma_i$ , all measuring the same quantity  $\theta$ , as illustrated in Fig. 1. For simplicity, we assume that the  $\sigma_i$  are mutually uncorrelated. The posterior distribution  $p(\theta|y_i, \sigma_i, \tilde{\mu}, \tilde{\tau})$  is again given by Eq. (5), but now with

$$\frac{1}{\sigma_{\tilde{\tau}}^2} = \sum_{i=1}^N \frac{1}{\sigma_i^2} + \frac{1}{\tilde{\tau}^2}, \tag{8}$$

and

$$\theta_{\tilde{\tau}} = \sigma_{\tilde{\tau}}^2 \left( \sum_{i=1}^N \frac{y_i}{\sigma_i^2} + \frac{\tilde{\mu}}{\tilde{\tau}^2} \right). \tag{9}$$

Obviously, the uncertainty  $\sigma_{\tilde{\tau}}$  in  $\theta$  decreases strictly monotonically with the inclusion of more experiments. Nevertheless, if one or several of the experiments was subject to a number of systematic effects that was neither corrected for, nor accounted for in the individual uncertainties  $\sigma_i$ , then the experiments are (effectively) not measuring the same quantity, and  $\sigma_{\tilde{\tau}}$  would be underestimated. In other words, each experiment can be viewed as measuring different parameters  $\theta_i$ , which are, however, not entirely independent of each other, since after all, the experiments were supposed to constrain the same  $\theta$ . We will now review hierarchical Bayesian modeling, and propose it as a systematic method to interpolate between the extreme and rarely realistic cases of all  $\theta_i$  being either equal or else entirely independent of each other.

### 2.2 The hierarchical model

This is achieved by considering each  $\theta_i$  to be the result of a random draw from a parent distribution,


$$p(\theta_i) = \int p(\theta_i|\mu, \tau)p(\mu, \tau)d\mu d\tau, \tag{10}$$

where  $p(\mu, \tau)$  is the *hyper-prior distribution* for what are now called the hyper-parameters  $\mu$  and  $\tau$ . We sketch this model in Fig. 2. Note that Eq. (10) implies the property of exchangeability between the  $\theta_i$ , i.e. symmetry under  $\theta_i \leftrightarrow \theta_j$ . From Bayes' theorem one has,

$$p(\theta_i, \mu, \tau|y_i, \sigma_i) \propto p(y_i|\theta_i, \sigma_i)p(\theta_i|\mu, \tau)p(\mu, \tau), \tag{11}$$

and explicitly in the Gaussian case,

$$p(\theta_i, \mu, \tau|y_i, \sigma_i) \propto \prod_{i=1}^N \mathcal{N}(y_i|\theta_i, \sigma_i)\mathcal{N}(\theta_i|\mu, \tau)p(\mu, \tau). \tag{12}$$



**Fig. 2** Hierarchical model. Each experimental parameter  $\theta_i$  arises from a random draw from a parent distribution with hyper-parameters  $\mu$  and  $\tau$ , and each experimental central value  $y_i$  is then considered to be the result of a random draw from a Gaussian distribution with central value  $\theta_i$  and error  $\sigma_i$

Marginalizing over  $\theta_i$  one finds the “master” equation,

$$p(\mu, \tau|y_i, \sigma_i) \propto \prod_{i=1}^N \mathcal{N}(\mu|y_i, \sigma_i^2 + \tau^2)p(\mu, \tau). \tag{13}$$

We will use it to compute the posterior distribution of the hyper-parameters, once a hyper-prior is chosen. For example, assuming a flat prior for  $\mu$  and  $\tau$ , we can integrate over  $\mu$  to find,

$$p(\tau|y_i) \propto \left( \sum_{i=1}^N \frac{1}{\sigma_i^2 + \tau^2} \right)^{-\frac{1}{2}} \prod_{i=1}^N \mathcal{N}(\hat{\mu}|y_i, \sigma_i^2 + \tau^2), \tag{14}$$

where,

$$\hat{\mu} = \left( \sum_{i=1}^N \frac{1}{\sigma_i^2 + \tau^2} \right)^{-1} \sum_{i=1}^N \frac{y_i}{\sigma_i^2 + \tau^2}. \tag{15}$$

The parameter  $\tau$  quantifies general differences in the  $\theta_i$ . If  $\tau = 0$ , the experiments measure the same parameter, i.e.,  $\theta_i = \theta_j$ . For  $\tau \rightarrow \infty$ , each one measures a completely independent parameter  $\theta_i$ .

From the master equation one can see that the parameter of interest is  $\mu$ . If  $\tau = 0$  the posterior distribution for  $\mu$  reduces to the ordinary likelihood for parameter estimation given in Eq. (5) with  $\tilde{\tau} \rightarrow \infty$ . The full posterior distribution for  $\mu$  can be obtained integrating Eq. (13) numerically over  $\tau$ . If there are large unknown systematic effects, then the most likely values of  $\tau$  will differ from zero, which leads to the important result of increasing the error in  $\mu$ .

### 2.3 The hyper-prior

We propose a hyper-prior which is  $\mu$ -independent, i.e.,  $p(\mu, \tau) = p(\tau)$ , and that interpolates smoothly between a flat and a sharply peaked  $\tau$  distribution,

$$p(\tau)d\tau^2 \propto \prod_{i=1}^N \left[ \frac{1}{\sigma_i^2 + \tau^2} \right]^{\frac{\alpha}{2N}} d\tau^2. \tag{16}$$

This form will prove to be useful due to the simple interpretation of  $\alpha$  in terms of the number of degrees of freedom, and the possibility to obtain closed analytical formulas for the posterior distribution of  $\mu$ . We remark that in Bayesian methods one needs to specify a prior that cannot be determined from first principles. Here we have chosen a prior with a simple analytical form interpolating between a flat prior and  $\tau = 0$ . Very interestingly, while this prior is only one of many possible choices, it turns out that it coincides with Jeffrey’s prior in a certain limit. We will return to this at the end of Sect. 6.

It is interesting to study the effect of this kind of prior on the tails of the posterior density of  $\mu$ . Integrating Eq. (13) over  $\tau$  produces the posterior density of  $\mu$  given the data,

$$p(\mu|y_i) \propto \int_0^\infty \prod_{i=1}^N (\sigma_i^2 + \tau^2)^{-\frac{1}{2}(1+\frac{\alpha}{N})} e^{-\frac{(\mu-\bar{y}_i)^2}{2(\sigma_i^2+\tau^2)}} d\tau^2. \tag{17}$$

For large  $\mu$ , the exponential suppression factor favors large values of  $\tau$ , so that,

$$p(\mu|y_i) \sim \int_0^\infty \tau^{-(N+\alpha)} e^{-\frac{N\mu^2}{2\tau^2}} d\tau^2, \tag{18}$$


and after a change of variables  $u^2 \equiv \mu^2/\tau^2$ ,

$$p(\mu|y_i) \sim \mu^{-(N+\alpha-2)}. \tag{19}$$

We observe that the usual exponential suppression of  $\mu$  in the tails has turned into a milder power law suppression which increases with the effective number of degrees of freedom, i.e., in our case the number of measurements,  $\nu \equiv N + \alpha - 2$ .

### 3 Experiments with errors of the same size

When all errors are equal,  $\sigma_i = \sigma_j \equiv \sigma$ , we obtain an analytical formula which illustrates how the PDG scale factor re-emerges for large data sets. The master equation reads in



**Fig. 3** Scale factor versus the square root of the reduced  $\chi^2$ . We employed  $\alpha = 0$

this case,

$$p(\mu, \tau|y_i) \propto (\sigma^2 + \tau^2)^{-\frac{\nu+2}{2}} \exp \left[ -\frac{\sum_{i=1}^N (\bar{y}_i - \mu)^2}{2(\sigma^2 + \tau^2)} \right],$$

or simply,

$$p(\mu|y_i) \propto \int_0^\infty (\sigma^2 + \tau^2)^{-\frac{\nu+2}{2}} \exp \left[ -\frac{\sigma^2 \chi^2}{2(\sigma^2 + \tau^2)} \right] d\tau^2, \tag{20}$$

where we defined,

$$\chi^2 \equiv \chi^2(\mu) \equiv \sum_{i=1}^N \frac{(\mu - \bar{y}_i)^2}{\sigma^2}, \tag{21}$$

which is the usual  $\chi^2$  function. Changing variables,

$$u \equiv \frac{\sigma^2 \chi^2(\mu)}{2(\tau^2 + \sigma^2)}, \tag{22}$$


we obtain,

$$p(\mu|y_i) \propto (\chi^2)^{-\frac{\nu}{2}} \int_0^{\chi^2/2} u^{\frac{\nu}{2}-1} e^{-u} du \propto (\chi^2)^{-\frac{\nu}{2}} F^\nu(\chi^2), \tag{23}$$

which is the master formula in this case in terms of the cumulative distribution function  $F$  for a  $\chi^2$  distribution with  $\nu$  degrees of freedom. This equation implies an interesting result. Since  $p(\mu|y_i)$  depends on  $\mu$  only through  $\chi^2(\mu)$ , we have

$$\frac{dp(\mu|y_i)}{d\mu} = \frac{dp(\mu|y_i)}{d\chi^2} \frac{d\chi^2}{d\mu}, \tag{24}$$

so that the mode of the distribution is the same as in the usual case, i.e., at the value of  $\mu$  where  $\chi'(\mu)^2 = 0$ . Thus,



**Fig. 4** Scale factor versus the square root of the reduced  $\chi^2$  for the case  $N = 10$

For  $\sigma_i = \sigma_j$  the posterior distributions of the hierarchical and non-hierarchical models peak at the same location.

From Eq. (23), we can also obtain the scale factor, which we define here as the ratio of the sizes of the 68% highest confidence intervals of the hierarchical and non-hierarchical models. In Figs. 3 and 4, we show the scale factor for several values of  $\alpha$  and  $N$ , from which one can see the similarity to the PDG scale factor for large  $N$ . We now turn to the case of a large number of degrees of freedom and the Gaussian approximation.

### 3.1 Large number of degrees of freedom

We rewrite Eq. (23) by another change of variables,

$$\frac{\chi^2 r}{2} = u, \tag{25}$$


so that

$$p(\mu|y_i) \propto \int_0^1 \exp\left[-\frac{\nu-2}{2} \left(r\chi_{\nu-1}^2 - \ln r\right)\right] dr, \tag{26}$$

where we defined  $\chi_{\nu-1}^2 \equiv \chi^2/(\nu-2)$ . Thus, large values of  $\nu$  suppress the integrand exponentially. Depending on the value  $r_0 = (\chi_{\nu-1}^2)^{-1}$  where  $r\chi_{\nu-1}^2 - \ln r$  has a minimum, we have two cases:

- (1) For  $r_0 > 1$  the minimum falls outside the integration limits, and the integral can be approximated by considering values of  $r$  near 1, which gives

$$p(\mu|y_i) \propto \frac{e^{-\chi^2/2}}{1 - \chi_{\nu-1}^2} \left[1 - e^{-\frac{\nu-2}{2}(1-\chi_{\nu-1}^2)}\right] \sim e^{-\chi^2/2}, \tag{27}$$



**Fig. 5** Comparison of the exact result with the approximate formula for  $\alpha = 0$

We recognize this is the usual likelihood for parameter inference without scaling. Thus,

for  $\sigma_i \approx \sigma_j$ ,  $\nu \rightarrow \infty$  and  $\chi_{\nu-1}^2(\mu_0) < 1$ , the hierarchical model implies no scaling of the errors.

- (2) For  $r_0 < 1$  the minimum resides inside the integration region, and the integral can be approximated by considering values of  $r$  near  $r_0$ . After some algebra,

$$p(\mu|y_i) \propto \left[1 + \frac{2}{\nu-1} \frac{(\mu - \mu_0)^2}{2 \left(\frac{\sigma^2 \chi_{\nu}^2(\mu_0)}{N}\right)}\right]^{-\frac{\nu}{2}}, \tag{28}$$

which is proportional to the Student-t distribution for  $\nu-1$  degrees of freedom, and for very large  $\nu$  it can be further approximated by a Gaussian,

$$p(\mu|y_i) = t_{\nu-1} \left(\mu_0, \frac{\sigma^2 \chi_{\nu}^2}{N}\right) \sim \mathcal{N} \left(\mu_0, \frac{\sigma^2 \chi_{\nu}^2}{N}\right). \tag{29}$$

This yields another important result,

for  $\sigma_i \approx \sigma_j$ ,  $\nu \rightarrow \infty$  and  $\chi_{\nu-1}^2(\mu_0) > 1$ , the hierarchical model implies a re-scaling of the overall error by  $\sigma \rightarrow \sigma \sqrt{\chi_{\nu}^2(\mu_0)}$ .

It is amusing to note that for large  $\nu$  we recovered the PDG scale factor prescription. On the other hand, for low values of  $\nu$  our model implies larger scalings than recommended by the PDG. In the next subsection we approximate the distribution of  $\mu$  as a Gaussian, so as to obtain an analytical formula for the scale factor in terms of  $\nu$  and the value of  $\chi^2$ .

### 3.2 Gaussian approximation

To do so, we expand the logarithm of the posterior distribution  $p = p(\mu|y_i)$  in powers of  $\mu$  around  $\mu_0$ ,

$$\ln p = C + \frac{d \ln p}{d\mu} \Big|_{\mu_0} (\mu - \mu_0) + \frac{d^2 \ln p}{d\mu^2} \Big|_{\mu_0} \frac{(\mu - \mu_0)^2}{2} + \dots$$

The second term on the right hand side is zero because we are expanding around the maximum. The third term can be compared to the corresponding term of the expansion of a Gaussian distribution, which gives

$$\frac{1}{\sigma_{\text{Bayes}}^2} \approx - \frac{d^2 \ln p}{d\mu^2} \Big|_{\mu_0} = - \frac{2N}{\sigma^2} \frac{d \ln p}{d\chi^2} \Big|_{\chi_0^2}. \tag{30}$$

Using Eq. (23) we have,

$$-2 \frac{d \ln p}{d\chi^2} \Big|_{\chi_0^2} = \frac{\nu}{\chi^2} - \frac{(\chi^2/2)^{(\frac{\nu}{2}-1)} e^{-\chi^2/2}}{\gamma(\nu/2, \chi^2/2)}, \tag{31}$$

where  $\gamma$  is the incomplete Gamma function, defined by

$$\gamma(s, x) \equiv \int_0^x t^{s-1} e^{-t} dt. \tag{32}$$

As we mentioned before, the scale factor  $S_{\text{Bayes}}$  is defined as the ratio of the sizes of the 68% highest confidence intervals of the hierarchical and non-hierarchical models. In the Gaussian approximation we find,


$$S_{\text{Bayes}} \approx \sqrt{N} \frac{\sigma_{\text{Bayes}}}{\sigma} \approx \sqrt{\frac{\chi^2}{\nu}} \left[ 1 + \frac{1}{\sum_{k=1}^{\infty} \frac{(\chi^2)^k \nu!!}{(\nu+2k)!!}} \right]^{\frac{1}{2}}, \tag{33}$$

where we have used the power series expansion of the incomplete Gamma function,

$$\gamma(s, x) = x^s \Gamma(s) e^{-x} \sum_{k=0}^{\infty} \frac{x^k}{\Gamma(s+k+1)}. \tag{34}$$

In Fig. 5 we compare the approximate formula with the exact result. As expected, the approximation improves for larger values of  $\nu$ . We are now ready to discuss the general case of unequal errors,  $\sigma_i \neq \sigma_j$ .

Experiments with different errors



**Fig. 6** The blue points with identical errors originate from a Gaussian distribution centered at 10. The last blue point has the same precision as the combination of the previous 10 points, but deviates by about  $5\sigma$ . The red point is the ordinary weighted average after PDG scaling. The black point is obtained using our Bayesian method

### 4 Experiments with unequal precisions

To understand this case, we fix the value of  $\tau$  in Eq. (13). The distribution of  $\mu$  is then Gaussian, with total error,

$$\frac{1}{\sigma_i^2} = \sum_{i=1}^N \frac{1}{\sigma_i^2 + \tau^2}, \tag{35}$$

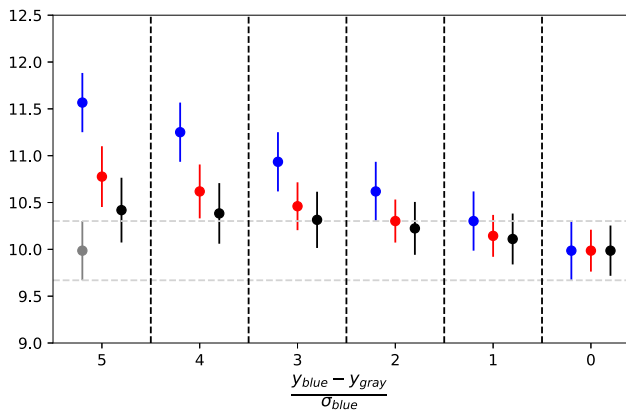
and central value,

$$\mu_0 = \left( \sum_{i=1}^N \frac{1}{\sigma_i^2 + \tau^2} \right)^{-1} \sum_{i=1}^N \frac{y_i}{\sigma_i^2 + \tau^2}. \tag{36}$$

Thus, experiments with smaller errors are more sensitive to  $\tau$  than less precise ones. Suppose that  $M$  of the experiments have an error  $\sigma_M$ , and that  $\sigma_M$  is much smaller than the error  $\sigma$  of the rest of the experiments. Then, for  $\sigma_M \simeq \tau \ll \sigma$  the scaling will mainly affect the experiments with small errors. Since we were unable to find an analytical formula for the peak or mean of  $\tau$ , we proceed with a numerical analysis.

As a first example, we randomly generated eleven fictitious measurement points from a Gaussian with standard deviation  $\sigma = 1$  centered at the value of 10. The last point is from a Gaussian centered at  $10 + 5/\sqrt{10}$  with  $\sigma_M = 1/\sqrt{10}$ , which is chosen so that its precision is the same as the combined precision of the other ten. The results are shown in Fig. 6. The red point denotes the ordinary weighted average with PDG scaling applied, and is pulled away from the horizontal line as a result of the deviating 11th measurement. The black point, on the other hand, is the average obtained as the result of our Bayesian hierarchical model (here we use  $\alpha = 10$  to specify our prior). It is closer to the bulk of





**Fig. 7** The measurement points with small error are shown in blue, the usual averages with the PDG scaling in red, and the hierarchical averages in black. The labels at the horizontal axis show by how many  $\sigma_M$  the blue points deviate from the gray point. The gray band represents the ordinary weighted averages of the bulk of measurements in Fig. 6

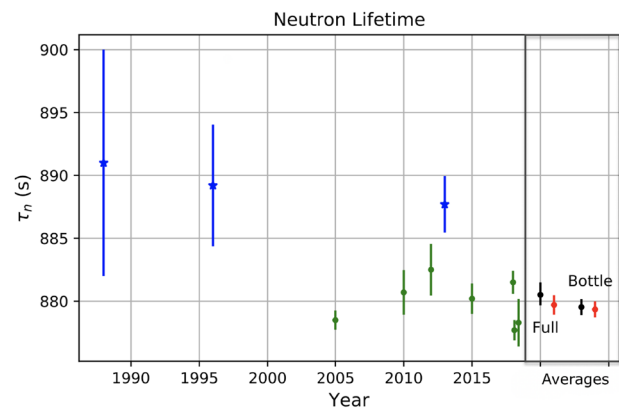
data than to the measurement with the smaller error. This is a reasonable property, since it is less likely that all the measurements in the bulk had a systematic error in the same direction.

In Fig. 7 we show how the two kind of averages change when we move the central value of the 11th measurement (in blue) while leaving the other 10 unchanged. Just for orientation, the gray band represents the ordinary average (non-hierarchical) of the bulk of measurements with the same error. As in Fig. 6, the red points are the usual PDG-scaled averages, while the black points are the hierarchical averages. Clearly, as we approach the bulk the combined error shrinks.

### 5 Neutron lifetime

There is an interesting discrepancy between the two types of experiments measuring the lifetime of the neutron. For a state of the art review of both types and more details, see Ref. [11]. The first type are beam experiments [12–14], which measure the number of protons or electrons from decays of cold neutrons in a beam passing through a magnetic or electric trap. After the beam has passed the trap, some of the neutrons are deposited in a foil at the end of the beam path. The neutron lifetime is proportional to the rate of neutrons deposited and inversely proportional to the rate of decays detected.

The other type of experiment uses bottles [15–21] containing ultra-cold neutrons with a kinetic energy of less than 100 neV. Neutrons with such a low kinetic energy can be confined due to the effective Fermi potential between neutrons and atomic nuclei in many materials. Gravitational forces and magnetic fields can also be used to confine the neutrons within the container. The idea is simply to count the num-



**Fig. 8** Neutron lifetime measurements. The green points are the results of bottle experiments, and the blue ones of beam experiments. The discrepancy can easily be seen. The black point to the left is the Bayesian average of the full data, while the first red point is the usual average with the PDG scaling. Similarly for the right black and red points but restricted to the bottle results. The PDG scaling for beam plus bottle experiments is  $S_{PDG} = 1.96$ , while for bottle only is  $S_{PDG} = 1.56$


ber of surviving neutrons after some time and to deduce the lifetime.

We now apply our method with  $\alpha = 6$  to the results of these experiments which are shown in Fig. 8. PDG  $\chi^2$  scaling ( $S_{PDG} = 1.93$ ), which is shown in red, yields the lifetime  $\tau_n = 879.71 \pm 0.78$  s, while the Bayesian method (black point to the left) gives  $\tau_n^{\text{Bayes}} = 880.51^{+0.98}_{-0.83}$  s. We find that our Bayesian hierarchical method increases the central value when the beam experiments are included. Even when only bottle experiments are considered, our method still gives a slightly larger average value  $\tau_n^{\text{Bayes}} = 879.53^{+0.64}_{-0.63}$  s, than the PDG method  $\tau_n = 879.35 \pm 0.64$  s where  $S_{PDG} = 1.56$ . This is due to the bulk of the bottle experiments that prefer lifetimes longer than 880 s. It is important to recall that the tails of the Bayesian hierarchical model do not fall as fast as a Gaussian, so that there is still a non-negligible probability for  $\tau_n$  to be lower.

### 6 Relations to other models

While this paper was being written, two interesting papers related to our work appeared. The first one [22] discusses the kaon mass in the context of a skeptical combination of experiments, scaling each experimental error independently but correlated. The second one [23] studies the discrepancy that arises when the PDG scaling is applied to sub-sets of experiments and then to the combination of the sets, vs. (for example) applying it to the whole data at the same time. The conclusion is that

the  $\chi^2/\nu$  prescription used to enlarge the standard deviation does not hold sufficiency.



**Fig. 9** Scaling for  $\alpha = 6$

This means that the scaling is not sufficient to properly describe the full probability distribution. Our model would have had the same problem had we used the marginalized (over  $\tau^2$ ) distribution of  $\mu$ . This is because the “correlations” that emerge through  $\tau^2$  would be absent. But it is clear from Eq. (13) that if we use the posterior distribution of  $\mu$  and  $\tau^2$  of a subset of experiments as the prior for the remaining subset, then the updated posterior will be the same as combining the whole data set simultaneously.

Another interesting point made in Ref. [23] is the fact that the PDG scaling treats any value of  $N$  equally, while for fixed  $\chi^2/N$  the  $p$  value decreases with  $N$ . In other words, since the probability distribution of the reduced  $\chi^2$  function peaks around one as the number of degrees of freedom increases, the scaling (given a discrepant value of the reduced  $\chi^2$ ) should be larger when more experiments are included in the average. This is not the case for the PDG description, because the scaling only depends on the reduced  $\chi^2$  value and not on the number of degrees of freedom. Now, it is clear from Fig. 3 that in the Hierarchical Model with  $\alpha$  chosen close to zero this problem would be aggravated, *i.e.*, for any given value of the reduced  $\chi^2$ , there is more scaling for low  $N$ . However, we can use the freedom to choose a value of  $\alpha$  to improve on this issue. First we demand the variance of the  $\tau$  distribution to be finite, which corresponds to  $\alpha > 6$ . In Fig. 9 we show the scaling versus the reduced  $\chi^2$  with  $\alpha = 6 + \epsilon$  (where  $\epsilon$  is an infinitesimal) from which one can see that for large values of the reduced  $\chi^2$  the scaling reduces as  $N$  gets smaller. This is just the desired effect. On the other hand, we still have more scaling for small values of the reduced  $\chi^2$ . This is a natural consequence of the fact that for a low number of experiments  $\tau$  can not be constrained too strongly, which translates into an enlarged error for  $\mu$ .

One can also consider Jeffrey’s prior. *E.g.*, if we specify to the case of uncertainties of equal magnitude,  $\sigma_i = \sigma_j = \sigma$ ,

In the case of a distribution with several parameters (in our case  $\mu$  and  $\tau^2$ ), Jeffrey’s prior is defined as the square root of the determinant of Fisher’s information matrix, which in turn is defined as the average (over  $y_i$ ) of the Hessian of the log-likelihood  $\mathcal{N}(y_i|\mu, \tau^2 + \sigma_i^2)$ .

then Jeffrey’s prior reduces precisely to Eq. (16) with  $\alpha = 3$ . This would lead to a plot very similar to the one shown in Fig. 9.

## 7 Conclusions and outlook

We proposed a Bayesian hierarchical model as a strategy to compute averages of several uncorrelated experimental measurements, specifically with the possibility in mind that unaccounted for systematic effects might be present, leading to underestimates of the quoted uncertainties. We should stress that the point is not that (some part of) the systematic error has been underestimated or assessed too aggressively. If this is suspected then a strategy should be developed to increase the systematic error component(s), which would imply—among other things—that statistics limited measurements would not be questioned. Here, we rather addressed the generic situation in which unknown effects or human errors may be present, and which therefore could affect even ostensibly clean determinations.

We have shown that our methodology resembles the recommendation of the Particle Data Group whenever the number of degrees of freedom (data points) is large. Our approach connects smoothly to cases with fewer degrees of freedom, though. Another important advantage is that it makes the underlying assumptions in the averaging process transparent. *E.g.*, a large value of the parameter  $\alpha$  appearing in our proposed form of the prior, implies a strong believe that the experiments do not have an unknown systematic error, while a small value corresponds to a more agnostic point of view. Our method can be extended to experiments with correlated errors, but we leave this generalization for the future.

Due to the additive form,  $\sigma_i^2 + \tau^2$ , of the denominator in the exponential part of the distribution, our model has the drawback that it tends to penalize experiments with high precision more strongly. This relative issue is already seen in the  $\tau_n$  example, where the most recent beam measurement which has a larger error than most bottle experiments and a higher central value tends to push the combined value up. On the other hand, the natural power suppressed tails of the posterior distribution help to mitigate possible strong shifts in the central value.

We also would like to point out that to apply our method to the PDG, it has to be studied, discussed and compared with other approaches in more detail, to confirm that it can be used within the PDG framework.

In closing, we remark that we also envision an application of this model in the context of new physics searches within the Standard Model Effective Field Theory (SMEFT) framework [24, 25], in which thousands of a priori independent operator (Wilson) coefficients need to be determined. Yet, many of these operators are almost certainly generated at some com-



mon energy scale, and are consequently not entirely independent. Thus, the idea is to assume that (classes of) the Wilson coefficients are random samples generated at a common ultra-violet energy scale, lending itself to a hierarchical approach. This can be particularly useful when estimating the sensitivity of a hypothetical future experiment to physics beyond the Standard Model. This is another direction for future work.

**Acknowledgements** We are happy to thank Glen Cowan and Giulio D'Agostini for discussions and comments and Marumi Kado for pointing us to relevant references. This work was supported by CONACyT (Mexico) project 252167-F, and also the German-Mexican research collaboration grant SP 778/4-1 (DFG) and 278017 (CONACyT).

**Data Availability Statement** This manuscript has no associated data or the data will not be deposited. [Authors' comment: We do not give data since the equations that were used for the analysis are clearly stated in the paper, and the data used to compute the plots is available either in the references or the text itself.]

**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 SCOAP<sup>3</sup>.

## References

1. M. Tanabashi, Particle Data Group, et al., *Phys. Rev. D* **98**, 030001 (2018)
2. A. Gelman et al., *Bayesian Data Analysis* (Chapman and Hall, London, 2013)
3. R.E. Tarone, *Biometrics* **38**, 215 (1982)
4. A.P. Dempster, M.R. Selwyn, B.J. Weeks, *J. Am. Stat. Assoc.* **78**, 221 (1983)
5. L.V. Hedges, I. Olkin, *Statistical Methods for Meta-Analysis* (Academic Press, New York, 1985)
6. T. Friede, C. Röver, S. Wandel, B. Neuenschwander, *Res. Synth. Methods* **8**, 79 (2017)
7. T. Friede, C. Röver, S. Wandel, B. Neuenschwander, *Biom. J.* **59**, 658 (2017)
8. G. Cowan, *Eur. Phys. J. C* **79**, 133 (2019)
9. G. D'Agostini, Sceptical combination of experimental results: general considerations and application to  $\epsilon'/\epsilon$ . [arXiv:hep-ex/9910036](https://arxiv.org/abs/hep-ex/9910036)
10. S. Mukhopadhyay, D. Fletcher, *Sci. Rep.* **8**, 9983 (2018)
11. F.E. Wietfeldt, *Atoms* **70**(4), 6 (2018)
12. P.E. Spivak, *JETP* **67**, 1735 (1988)
13. J. Byrne et al., *Eur. Phys. Lett.* **33**, 187 (1996)
14. A.T. Yue et al., *Phys. Rev. Lett.* **111**, 222501 (2013)
15. A. Pichlmaier et al., *Phys. Lett. B* **693**, 221 (2010)
16. A. Steyerl et al., *Phys. Rev. C* **85**, 065503 (2012)
17. A.P. Serebrov et al., *Phys. Rev. C* **97**, 055503 (2018)
18. A. Serebrov et al., *Phys. Lett. B* **605**, 72 (2005)
19. V.F. Ezhov et al., *JETP Lett.* **107**, 671 (2018)
20. R.W. Pattie Jr. et al., *Science* **360**, 627 (2018)
21. S. Arsumarov et al., *Phys. Lett. B* **745**, 79 (2015)
22. G. D'Agostini, Sceptical combination of experimental results using JAGS/rjags with application to the  $K^\pm$  mass determination. [arXiv:2001.03466](https://arxiv.org/abs/2001.03466)
23. G. D'Agostini, On a curious bias arising when the  $\sqrt{\chi^2/\nu}$  scaling prescription is first applied to a sub-sample of the individual results. [arXiv:2001.07562](https://arxiv.org/abs/2001.07562)
24. W. Buchmüller, D. Wyler, *Nucl. Phys. B* **268**, 621 (1986)
25. B. Grzadkowski, M. Iskrzynski, M. Misiak, J. Rosiek, *JHEP* **1010**, 085 (2010)