



Available online at www.sciencedirect.com

ScienceDirect

Nuclear Physics B 911 (2016) 623–637

**NUCLEAR
PHYSICS B**

www.elsevier.com/locate/nuclphysb

Taming systematic uncertainties at the LHC with the central limit theorem

Sylvain Fichet

*ICTP South American Institute for Fundamental Research, Instituto de Física Teórica, Sao Paulo State University,
Brazil*

Received 5 June 2016; received in revised form 18 August 2016; accepted 22 August 2016

Available online 27 August 2016

Editor: Tommy Ohlsson

Abstract

We study the simplifications occurring in any likelihood function in the presence of a large number of small systematic uncertainties. We find that the marginalisation of these uncertainties can be done analytically by means of second-order error propagation, error combination, the Lyapunov central limit theorem, and under mild approximations which are typically satisfied for LHC likelihoods. The outcomes of this analysis are *i*) a very light treatment of systematic uncertainties *ii*) a convenient way of reporting the main effects of systematic uncertainties, such as the detector effects occurring in LHC measurements.

© 2016 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>). Funded by SCOAP³.

1. Introduction

The search for physics beyond the Standard Model requires a thorough statistical investigation of the data collected at the LHC. The analyses of LHC samples are often plagued by a number of systematic uncertainties, such as detector resolutions or the imperfect knowledge of physical constants, that should be treated with care in order to obtain reliable conclusions on a putative signal. A consistent implementation of these systematic uncertainties in the likelihood function is done either by summing over all the realisations of the nuisance parameters or by maximising the likelihood with respect to them, respectively within the Bayesian and frequentist frameworks.

E-mail address: sylvain@ift.unesp.br.

<http://dx.doi.org/10.1016/j.nuclphysb.2016.08.029>

0550-3213/© 2016 The Author. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>). Funded by SCOAP³.

While the operations of Bayesian and frequentist marginalisation are conceptually clear, one can distinguish two main issues related to their practical implementation. First, the correct treatment of systematic uncertainties is technically challenging to perform, because it necessitates many multidimensional integrations or maximisations. For example, integrating the $O(4000)$ nuisance parameters of the Higgs likelihood is very difficult, even with the computing power accessible to the LHC Collaborations. Second, transmitting or making public the information of the systematic uncertainties can also be technically challenging. The study we carry out will end up providing new insights into both of these topics.

In the context of the LHC, communicating the full experimental likelihoods via the RooFit/RooStats framework [1,2] has been suggested in [3,4]. The presentation method we will propose is somehow complementary from the proposal of [3,4], in that it is technically straightforward to carry out and leads to a fairly human-readable summary of the systematic uncertainties. Also, the goal of presenting LHC results decoupled from systematic uncertainties has been pursued in [5] in the context of theoretical errors on Higgs cross-sections. While the objectives of [5] are partly similar to the ones of this work, the results obtained are different. In particular, the general marginal likelihood that we will display is derived from first principles, and no discussion about reparametrisation templates is required.

Along this work we are going to adopt the hypothesis of a *large number of small independent sources of uncertainties*. Here by “small” we mean to qualify the *relative magnitude* of the systematic uncertainties. We stress that this is an intrinsic property that does not depend on the magnitude of the statistical uncertainties, *i.e.* on the size of the data sample. The validity of our results will thus not depend on the amount of observed data.

The assumption of small relative magnitude is fairly weak as, to our knowledge, most of LHC systematic uncertainties have a relative magnitude which is much lower than 100%. The assumption of independence follows naturally from the process of describing systematic uncertainty as correctly as possible, as the more one delves into the origin of uncertainty, the more its description becomes a set of elementary sources unrelated to each other. Independence will have the crucial implication that the combination of the elementary uncertainties is mostly described by its *first and second moments*.¹

Our computation consists of two steps of error propagation and error combination, that are laid down in Sec. 2. These steps already lead to a substantially simplified likelihood. In addition, if the relative magnitude of the combined uncertainties is somewhat small with respect to one, we show in Sec. 3 that marginalisation can be done exactly, providing a general, explicit formulation of the marginal likelihood. The cases of signal+background and differential distributions are treated. Finally, a signal strength toy-analysis illustrating the validity of our calculation is displayed in Sec. 4.

2. Taming a large number of small uncertainties

We consider an event-counting likelihood where \hat{n} is the observed event number and n is an expected (*i.e.* theoretical) event number. Assume n depends on parameters of interest θ and on nuisance parameters δ . The likelihood is then defined as

$$L(\theta, \delta) \equiv \Pr(\hat{n} \mid n(\theta, \delta), \theta, \delta). \quad (1)$$

¹ This well-known fact is quantified by the Berry–Esseen theorem [6–8] and is closely related to central limit theorems.

All the variables written in this likelihood $(\theta, \delta, \hat{n}, n)$ should be understood as vectors, whose labels and dimensions will be made explicit below.

Without much loss of generality,² one further assumes that the various measurements of n are statistically independent. In the following we will denote by N the number of independent measurements, *i.e.* $I \in [1 \dots N]$, and by p the number of nuisance parameters. The likelihood of our focus has thus the form

$$L(\theta, \delta) \equiv \prod_{I=1}^N \Pr(\hat{n}_I \mid n_I(\theta, \delta), \theta, \delta). \quad (2)$$

The approach laid down in the present section applies in fact to any likelihood that can be expressed as a function of the expected event numbers, $L[n_I(\theta, \delta)]$.³ However, the subsequent analytical marginalisation presented in Sec. 3 would not hold in general, as the approximate likelihood would not be Poissonian or Gaussian in δ .

2.1. Parametrisation

For a systematic uncertainty spanning a given domain, there exists in principle an infinite number of parametrisations, that are all equivalent under suitable redefinition of the distribution of the uncertainty. Among all possible parametrisations, it is useful to choose one that makes appear the relative magnitude of the uncertainty. We define a standardised representation as follows.

For any quantity A subject to uncertainty, that can take *both signs*, one simply defines

$$A = A_0(1 + \Delta\delta), \quad (3)$$

where the nuisance parameter δ satisfies⁴

$$E[\delta] = 0, \quad V[\delta] = 1, \quad (4)$$

so that Δ corresponds to the relative magnitude of the uncertainty, *i.e.* $V[A] = A_0^2 \Delta^2$. Our working hypothesis of small relative magnitude translates as $\Delta \ll 1$.

We also need to consider quantities that can only be positive, in the first place the expected event number n .⁵ Throughout this paper one defines the standardised form for the error on a positive quantity as

$$n \equiv n_0 e^{\delta\Delta}, \quad (5)$$

² Most of LHC event selections are independent of each other. A notable exception occurs when various overlapping selections of a same dataset are reported, with no information about more elementary, mutually exclusive selections. The correct statistics describing the set of overlapping selections is a multivariate Poisson. This has been described in details in [9] in the context of diboson ATLAS results [10].

³ This includes the case where some systematic uncertainties have a large relative magnitude, and have to be treated exactly instead of using the framework presented in this section. However, the best strategy in that case may be to first implement the small systematic uncertainties, leading to the approximate marginal likelihood of Eq. (29), then marginalise exactly (numerically) over the large ones. Thus, even in that case, it is enough to start our analysis with the form Eq. (2).

⁴ Given a random variable X with density f_X , the expectation and variance operators are defined as $E(X) = \int dx x f_X(x)$, $V(X) = \int dx x^2 f_X(x) - E(X)^2$.

⁵ The linear parametrisation Eq. (4) is not so suitable in that case because it requires to truncate the domain of δ above $-\frac{1}{\Delta}$, implying that the prior itself depends on Δ .

where $E[\delta] = 0$, $V[\delta] = 1$ and n_0 is the nominal value of n in the absence of uncertainty. Below it will become clear that the expansion in Δ has to be done up to second order, so that Eq. (5) can be equivalently taken to be

$$n \equiv n_0 \left(1 + \delta\Delta + \frac{(\delta\Delta)^2}{2} + O(\Delta^3) \right). \tag{6}$$

Compared to the linear form Eq. (4), one can see that the extra quadratic term ensures positivity of n . It also induces a small, positive shift of the mean value of n , as $E(n) = n_0(1 + \Delta^2/2)$ – or similarly $E(n) = n_0 e^{\Delta^2/2}$ without the expansion. The variance is $V(n) = n_0^2(e^{2\Delta^2} - e^{\Delta^2}) = n_0^2\Delta^2(1 + O(\Delta^4))$.

2.2. Error propagation

As a first step, we want to propagate the systematic uncertainties at the level of the event numbers. For an event number n depending on a quantity Q subject to uncertainty, we have

$$n[Q] \equiv n[Q_0(1 + \Delta_Q\delta)]. \tag{7}$$

The propagation amounts to perform a Taylor expansion with respect to Δ_Q . This expansion should be truncated appropriately to retain the leading effects of the systematic uncertainties in the likelihood. For now we take for granted that the expansion should be truncated above *second order*. This order will be justified further below.

In a one-parameter case, second order propagation leads to

$$n \equiv n_0 \exp \left(\frac{n'}{n_0} \Delta_Q \delta + \left(\frac{n''}{n_0} - \frac{n'^2}{n_0^2} \right) \frac{\Delta_Q^2 \delta^2}{2} + O \left(\frac{n^{(3)}}{n_0} \Delta_Q^3 \right) \right). \tag{8}$$

This is most easily obtained by expanding $\log n$. Clearly, the validity of this expansion relies on neglecting higher powers of Δ_Q times the appropriate derivative of $\log n$. As long as n is well-behaved, which should be checked in practice, this expansion is valid for uncertainties that have a small relative magnitude, *i.e.*

$$\Delta_Q \ll 1. \tag{9}$$

For example, for elementary systematic uncertainties that do not exceed $\Delta \sim 10\%$, keeping the expansion up to second degree implies that the neglected higher order terms are $O(0.1\%)$.

In the case were various uncertainties $Q_{1\dots p}$ are propagated into n , it is convenient to use a vector notation for δ and Δ . Assuming p nuisance parameters, one defines

$$\delta = (\delta_1 \dots \delta_p)^t, \tag{10}$$

$$\Delta_1 = \left(\frac{\partial n}{\partial \delta_1} \Delta_{Q,1}, \dots, \frac{\partial n}{\partial \delta_p} \Delta_{Q,p} \right)_{\delta=0}^t n_0^{-1}, \tag{11}$$

$$\Delta_2 = \left(\frac{\partial^2 n}{\partial \delta_i \partial \delta_j} \frac{\Delta_{Q,i} \Delta_{Q,j}}{2n_0} - \frac{\Delta_{1,i} \Delta_{1,j}}{2} \right)_{\delta=0}. \tag{12}$$

The relative uncertainties propagated to n are then written as

$$n \equiv n_0 \exp \left(\Delta_1^t \cdot \delta + \delta^t \cdot \Delta_2 \cdot \delta + O \left(\frac{n^{(3)}}{n_0} \Delta_Q^3 \right) \right). \tag{13}$$

After this step of error propagation, the likelihood takes the form

$$L(\theta, \delta) \equiv \prod_I \Pr(\hat{n}_I \mid n_{0,I} \exp(\Delta_{1,I} \cdot \delta + \delta \cdot \Delta_{2,I} \cdot \delta), \theta, \delta). \tag{14}$$

All the $n_{0,I}$, $\Delta_{1,I}$, $\Delta_{2,I}$ depend in principle on the parameters of interest θ .

Details about the order of truncation. For the sake of determining the truncation order, it is enough to consider a one-parameter case and take limits. Consider a likelihood $\Pr(\hat{n} | n_0 \exp(\Delta_1 \delta + \Delta_2 \delta^2))$ with $\Delta_2 = O(\Delta_1^2)$.

We first study the limit of an infinite amount of data. In that case, the likelihood tends to a Dirac peak,⁶

$$\delta(\hat{n} - n_0 \exp(\Delta_1 \delta + \Delta_2 \delta^2)). \tag{15}$$

If one neglects the Δ_2 term, marginalising this likelihood with a prior $\pi(\delta)$ for the nuisance parameter gives

$$\tilde{L} = \pi(\log(\hat{n}/n_0)/\Delta_1). \tag{16}$$

As δ is $O(1)$ by definition, it follows that $\log(\hat{n}/n_0) = O(\Delta_1)$. Using this fact, one can then verify that $\tilde{L} = \pi(\frac{\log(\hat{n}/n_0)}{\Delta_1} + O(\Delta_1))$ once the Δ_2 term is included. We conclude that the leading effect of the systematic uncertainty comes from Δ_1 , and appears thus at first order in the expansion.

Second we study the case where the amount of data is small enough so that the likelihood itself can be expanded with respect to $n_0 \Delta_1$, $n_0 \Delta_2$. It comes

$$L(\delta) = L(0) + \delta(n_0 \Delta_1 L') + \delta^2 \left(\frac{\Delta_1^2}{2} (n_0^2 L'' + n_0 L') + n_0 \Delta_2 L' \right) + O(n_0^3 \Delta_1^3). \tag{17}$$

One then marginalises this likelihood with respect to the nuisance parameter δ using an arbitrary prior, $\int d\delta L(\delta) \pi(\delta)$. By definition, $E(\delta) = 0$ (see Eq. (4)), so the linear term vanishes. The leading effect of the uncertainties appears thus from the second order term in the expansion. This implies that the expansion has to be done at quadratic order from the beginning, so that the Δ_2 term should not be neglected.

As the truncation has to be done above second order in one of the limiting cases, it is convenient to use this order in all cases to ensure that all leading effects of systematic uncertainties are consistently taken into account.

2.3. Error combination

The previous step of propagation opens up the possibility of combining the nuisance parameters. We define nuisances parameter $\bar{\delta}_I$, associated to every measurement I , so that

$$n_I = \bar{n}_{0,I} \exp(\Delta_I \bar{\delta}_I) \equiv n_{0,I} \exp(\Delta_{1,I} \cdot \delta + \delta^t \cdot \Delta_{2,I} \cdot \delta). \tag{18}$$

These combined nuisance parameters are in general *correlated* with each others, their joint distribution we will denote $\bar{\pi}$. The set of equations (18) is the starting point for the combination

⁶ More precisely, the Dirac limit can be taken when the relative magnitude of the statistical uncertainty – given by $1/\sqrt{n_0}$ in the Poisson case – is small with respect to the inverse Fisher information of all the priors present in the problem.

of uncertainties.⁷ The likelihood expressed with respect to the combined nuisance parameters is written as

$$\bar{L}(\theta, \bar{\delta}) \equiv \prod_I \Pr(\hat{n}_I \mid \bar{n}_{0,I} \exp(\Delta_I \delta_I), \theta, \bar{\delta}). \quad (20)$$

Following our conventions, the combined nuisance parameters have to satisfy $E(\bar{\delta}_I) = \int d\bar{\delta}_I \bar{\delta}_I \bar{\pi}(\bar{\delta}_I) = 0$, $V(\bar{\delta}_I) = \int d\bar{\delta}_I \bar{\delta}_I^2 \bar{\pi}(\bar{\delta}_I) = 1$. The next task is to determine the numbers $\bar{n}_{0,I}$ and Δ_I . This is obtained by taking the expectation and the variance on both sides of Eq. (18).

The central values of the event numbers before and after combination are different because of the nonlinear propagation. It turns out that the diagonal terms of $\Delta_{2,I}$ contribute to the mean value of n_I , so that⁸

$$\bar{n}_{0,I} = n_{0,I} (1 + \text{tr}(\Delta_{2,I})). \quad (21)$$

The relative magnitudes Δ_I are obtained by evaluating the variance on the two sides of Eq. (18). One gets

$$\Delta_I = (\Delta_{1,I}^t \cdot \Delta_{1,I})^{1/2} + O(\Delta^4), \quad (22)$$

where the $O(\Delta^4)$ denotes higher order terms like $\text{tr}(\Delta_{2,I}^t \cdot \Delta_{2,I})$, $(\text{tr} \Delta_{2,I})^2$, $\Delta_{1,I}^t \cdot \Delta_{1,I} \text{tr} \Delta_{2,I}$. One may note the contrast of Eq. (22) with the mean value Eq. (21), where $\Delta_{2,I}$ provides the main correction and cannot be ignored.

The next step is to compute the correlation matrix among the event numbers n_I , n_J induced by the systematic uncertainties. The correlation matrix is found to be⁹

$$\rho_{IJ} = \frac{\Delta_{1,I}^t \cdot \Delta_{1,J}}{\Delta_I \Delta_J} + O(\Delta^2). \quad (23)$$

In the next paragraph it will be made clear that under our working assumptions, this information is enough to describe the entire shape of the combined uncertainties distribution $\bar{\pi}$.

2.4. Shape of the combined prior

Computing the joint distribution $\bar{\pi}$ of the combined uncertainties may seem at first view a very challenging task, as there can be in principle a lot of uncertainty sources. The experimental Higgs likelihood, for example, contains $O(4000)$ nuisance parameters, *i.e.* the vector δ has dimension $O(4000)$. This means that 4000 convolutions would have to be done for each value of the combined nuisance parameters $\bar{\delta}_I$.

⁷ At the level of the likelihood, combination is defined as the variable change

$$\int d\bar{\delta} \bar{L}(\theta, \bar{\delta}) \bar{\pi}(\bar{\delta}) \propto \int d\delta L(\theta, \delta) \prod_{i=1}^p \pi_i(\delta_i), \quad (19)$$

where the π_i are the priors of the elementary nuisance parameters. This is equivalent to Eq. (18).

⁸ This shift can also be observed by evaluating the mean value of the likelihood taken as a function of the $n_{0,I}$. It can be derived explicitly for a Poisson likelihood, and can also be derived for an arbitrary likelihood in the $n_{0,I} \Delta_2 \ll 1$ case, by expanding in $n_{0,I} \Delta_2$ and integrating by part with respect to the $n_{0,I}$.

⁹ In this paper we focus on the case of independent sources of uncertainty. The more general case of correlated nuisance parameters will be treated in [11].

One should however realise that the shape of $\bar{\pi}$ is determined by its central moments of order higher than two, which all depend only on the $\Delta_{1,I}$ at leading order in the Δ -expansion. In fact, at leading order, the Δ_2 term matters only for the mean value and always gives subleading contributions to higher moments. This can be seen by evaluating the central moments of Eq. (18). One can thus safely neglect the $O(\Delta_2)$ term in the combination $\Delta_I \bar{\delta}_I = \Delta_{1,I} \cdot \delta + O(\Delta_2)$, and make the crucial observation that this quantity is as sum of many independent random variables.

Besides, one notices that all the common distributions for nuisance parameters, such as the uniform, normal, log-normal distributions, possess *finite* higher moments. This is enough to invoke the Lyapunov central limit theorem (CLT) [12],¹⁰ which can be stated as follows. If it exists an integer $\kappa > 0$ so that

$$\frac{1}{\left(\sum_{i=1}^n \Delta_{1,I,(i)}^2\right)^{1+\kappa/2}} \sum_{i=1}^n \Delta_{1,I,(i)}^{2+\kappa} E[\delta_{1,I,(i)}^{2+\kappa}] \rightarrow 0 \quad \text{when } n \rightarrow \infty, \tag{24}$$

then the distribution of the combination $\Delta_{1,I} \cdot \delta$ converges in distribution towards a normal law with variance Δ_I^2 . For $\kappa = 1$, for example, the condition involves the third moments of the nuisance parameters. The Lyapunov condition is verified for any kind of prior shape used in LHC analyses, such as normal, log-normal or uniform distributions.¹¹

An estimate of the rate of convergence of the combined prior towards the normal law is given by the Berry–Esseen theorem [6–8]. For the combination of identical nuisance parameters the maximal difference between the combined prior and the Gaussian decreases as $1/\sqrt{n}$. For a combination of arbitrary nuisance parameters, which is our focus, the Berry–Esseen theorem states that the maximal difference between the combined prior and the Gaussian is of order

$$\frac{1}{\left(\sum_{i=1}^n \Delta_{1,I,(i)}^2\right)^{3/2}} \sum_{i=1}^n \Delta_{1,I,(i)}^3 E[\delta_{1,I,(i)}^3]. \tag{25}$$

This can be used in order to get an estimate of the convergence of the combined prior.

The arguments above can be applied separately to every combined nuisance parameter $\bar{\delta}_I$. However, whereas the elementary uncertainties are independent, the various $\bar{\delta}_I$ are correlated between each other – the correlation matrix is given by Eq. (23). The proof that the distribution of the set of $\bar{\delta}_I$ converges towards a multivariate normal is obtained by decomposing ρ_{IJ} as $\rho_{IJ} = A_{IK}^t A_{KJ}$.¹² Provided that the Lyapunov condition is satisfied for every $A_{IK} \delta_K$, one gets by definition a multivariate normal with diagonal correlation matrix. Applying the reverse transformation achieves to prove that the combined prior has asymptotically the form

$$\bar{\pi}(\bar{\delta}) = \frac{1}{(2\pi)^{N/2} |\rho|^{1/2}} \exp\left(-\frac{1}{2} \sum_{IJ=1}^N \bar{\delta}_I (\rho)^{-1}_{IJ} \bar{\delta}_J\right). \tag{26}$$

As a consequence, the combined uncertainty on the expected event numbers n_I asymptotically follows a *multivariate log-normal* distribution (see Eq. (20)).

¹⁰ The Lyapunov CLT does not require identically distributed nuisance parameters, nor identical variances. In a similar fashion, the Lindberg–Feller CLT [8,12,13] applies with a condition weaker than Eq. (24), but maybe less intuitive. The Lindberg condition is implied by the Lyapunov condition.

¹¹ In particular, note that in cases where the distribution of the nuisance parameters is symmetric, Eq. (24) is zero for odd κ and any n , so the Lyapunov condition is automatically satisfied.

¹² This is allowed as ρ_{IJ} is a real symmetric matrix.

Besides these robust arguments, one may also note that many of the elementary systematics uncertainties readily have a Gaussian prior, which further enhances the convergence rate of the combination. This is also true for the log-normal distribution in the limit of relative magnitude, in which case the log-normal is approximately Gaussian. The manifestation of the CLT in the case of theoretical uncertainties for Higgs production and decay rates has been explicitly observed in [14].

2.5. Practical considerations

We claimed above that under the assumption of a large number of small uncertainties, the shape of elementary priors does not matter and the shape of the combined prior is approximately Gaussian. All the information needed to treat the systematic uncertainties is in fact contained in the mean values and the covariance matrix of the combined nuisance parameters, that are obtained through the steps of propagation/combination described above.

Some practical conclusions can already be drawn. It turns out that the approximate treatment proposed above only requires the knowledge of a finite set of numbers:

- The magnitude of the elementary uncertainties, Δ_Q^i , of dimension p .
- The first derivative of the expected event numbers with respect to every nuisance parameters, *i.e.* $\partial n_I / \partial \delta_i$, of dimension $N \times p$.
- The diagonal second derivative of the expected event numbers with respect to every nuisance parameters, *i.e.* $\partial^2 n_I / \partial \delta_i \partial \delta_i$, of dimension $N \times p$.

All the relevant information about systematic uncertainties is thus encoded into $(2N + 1)p$ numbers. The transmission of this information poses no technical challenge. In the context of LHC analyses, it could be an easy and efficient way for the Collaborations of making public the main detector effects.

Besides, as a rule of thumb about the typical number of elementary uncertainties required for the CLT to converge, one can ask for a minimum number of $p = 4\text{--}5$ elementary uncertainties with similar magnitudes and flat priors. In case of Gaussian priors, this constraint does not hold as the combined prior is perfectly Gaussian for any p .

3. Analytic marginalisation for Poisson and Gaussian likelihoods

The previous steps of propagation, combination, and prior simplification can readily be used to reduce the amount of nuisance parameters in any kind of fit. In the case of Higgs experimental uncertainties, the $O(4000)$ -dimensional space of nuisance parameters would be reduced to a $O(100)$ -dimensional space – the amount of statistically independently observed channels. But ultimately, an integration still needs to be carried out over a space of substantially large dimension, for which a Monte Carlo integration is often required.

However, it turns out that one can go further under the extra condition that the *combined* uncertainties are small. Indeed, one can use a Taylor expansion with respect to the magnitude of the combined uncertainties Δ_I up to quadratic order in order to simplify the likelihood. This will render possible a completely analytical marginalisation. In both cases of Poisson and Gaussian statistics, the expansion of the likelihood reads

$$\Pr(\hat{n}|\lambda(1 + \delta\Delta)) = \Pr(\hat{n}|\lambda)e^{(\hat{n}-\lambda)\Delta\delta - \lambda\Delta^2\delta^2/2 + O(\Delta^3)}. \quad (27)$$




Fig. 1. Exact and approximate Poisson likelihood of Eq. (27) (resp. plain and dashed curves) as a function of the nuisance parameter $\Delta\delta$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In practice, the validity range of the approximation depends on the amount of data and on the expected number of events. This is illustrated in Fig. 1 for $\hat{n} = 10, 100$ and for typical values of λ roughly corresponding to 0, 2 and 3 sigma deviations. As a very rough rule of thumb for typical values of \hat{n} and λ , one may keep in mind that the validity is good up to $\Delta \sim 20\%$.¹³

We now plug the approximation Eq. (27) into the general likelihood \bar{L} , given in Eq. (20). The marginalisation of \bar{L} is given by

$$\tilde{L}(\theta) = \int d\bar{\delta} \bar{L}(\theta, \bar{\delta}) \bar{\pi}(\bar{\delta}), \tag{28}$$

where the combined prior $\bar{\pi}$ is given by Eq. (26) and involves the correlation matrix ρ_{IJ} of the $\bar{\delta}$ given in Eq. (23). The approximate marginal likelihood is found to be

$$\tilde{L}(\theta) = L^{\text{stat}}(\theta) L^{\text{sys}}(\theta) = L^{\text{stat}}(\theta) \frac{1}{\sqrt{|\eta\rho + \mathbb{1}|}} \exp\left(\frac{1}{2} \xi \cdot (\eta + \rho^{-1})^{-1} \cdot \xi\right), \tag{29}$$

where “ \cdot ” is matrix multiplication, and one introduced the vector

$$\xi_I = \Delta_I (\hat{n}_I - \bar{n}_{0,I}), \tag{30}$$

the diagonal matrix η with

$$\eta_{II} = (\Delta_I)^2 \bar{n}_{0,I}, \tag{31}$$

and $L^{\text{stat}} = \prod_I \text{Pr}(\hat{n}_I | \bar{n}_{0,I})$.

Both ξ and η depend on the parameters of interest θ via the expected event numbers and the combined uncertainties, *i.e.* one has in general $n_{0,I}(\theta)$, $\Delta_I(\theta)$. The L^{stat} term is almost the likelihood with no nuisance parameters, *i.e.* the piece of likelihood encoding the statistical uncertainty, except that it involves the shifted expected event numbers $\bar{n}_{0,I}$ (see Eq. (21)). In the limit of zero systematic uncertainty, *i.e.* $\Delta \rightarrow 0$, $L^{\text{sys}}(\theta)$ becomes an irrelevant constant so that $\tilde{L}(\theta) \rightarrow L^{\text{stat}}(\theta)$. The L^{sys} encodes most of the effect of the systematic uncertainties. Its effect is to enlarge and shift the preferred regions of the parameters of interest.

¹³ In the presence of large sample $\hat{n} \gg 1$, it is customary to approximate the Poisson likelihood by a Gaussian. It is worth noticing that, while the computation of Eq. (27) is straightforward for the Poisson case, obtaining the same result starting from the Gaussian is a bit more delicate. Depending on how the approximation is done, slightly different expressions can be obtained, that all are close from each other provided that $\hat{n} \gg 1$. The Poisson result Eq. (27) is valid for any \hat{n} , and will be used in the following.

Frequentist marginalisation (*i.e.* profiling), as well as the Bayesian and frequentist bias methods described in [14] can all be treated analytically, substituting Eq. (28) by the appropriate operation. The approach described above and leading to Eq. (29) is general. Nevertheless it is interesting to work out specific cases that are omnipresent in LHC analyses.

3.1. Signal strength fit

The general approach summarised by Eq. (29) can be applied to the very typical case where the expected number of events is split into a signal and background component, $n = s + b$. The signal can be also further parametrised as $s = \mu s_0$, where the parameter of interest μ is a “signal strength modifier” and s_0 is some nominal value for the signal.

In principle both background and signal are plagued by systematic uncertainties, so that one should distinguish the elementary nuisance parameters for signal and background, δ^s, δ^b . After a preliminary step of error propagation, the systematic uncertainty on the expected rates takes the form

$$n_I = s_I \exp(\Delta_{1,I}^s \cdot \delta^s + (\delta^s)^t \cdot \Delta_{2,I}^s \cdot \delta^s) + b_I \exp(\Delta_{1,I}^b \cdot \delta^b + (\delta^b)^t \cdot \Delta_{2,I}^b \cdot \delta^b) \tag{32}$$

In order to obtain the standard form for propagated errors Eqs. (10), (11), (12), one defines the overall vector of elementary uncertainties $\delta = (\delta^s, \delta^b)$, and write $n_I = n_{0,I} \exp(\Delta_{1,I} \cdot \delta + \delta^t \cdot \Delta_{2,I} \cdot \delta)$ where

$$n_{0,I} = s_{0,I} + b_{0,I}, \quad \Delta_{1,I} = \frac{(s_{0,I} \Delta_{1,I}^s, b_{0,I} \Delta_{1,I}^b)}{s_{0,I} + b_{0,I}}, \tag{33}$$

$$(\Delta_{2,I})_{ij} = \frac{1}{s_{0,I} + b_{0,I}} \begin{pmatrix} s_{0,I}(\Delta_{2,I}^s + (\Delta_{1,I}^s)^2/2) & 0 \\ 0 & b_{0,I}(\Delta_{2,I}^b + (\Delta_{1,I}^b)^2/2) \end{pmatrix} - \frac{1}{2}(\Delta_{1,I})_i(\Delta_{1,I})_j. \tag{34}$$

This makes contact with the standard notation of Eq. (13), and the analytic marginal likelihood is readily given by Eq. (29). If the N independent likelihoods correspond to N measurements of a same process, one has $s_{0,I} = s_0, b_{0,I} = b_0$ for every I . This case will be illustrated in Sec. 4.

Here, positive s and b have been assumed. It is also possible to allow s to take negative values if it is dominated by the destructive interference between the SM and BSM matrix elements. In that case a linear modelisation of the error on s is fine, but one should bear in mind that if the support of the δ_I^s is such that $b + s$ can be zero, depending on the prior of δ_I^s the likelihood can blow up above some arbitrary large value of s . This is a general fact that is not specific to the approximations studied in this paper.

3.2. Differential distributions

Another typical measurement at the LHC is the one of a differential distribution. The likelihood with no systematic uncertainties has then the form of Eq. (14), where every measurement corresponds to a different bin I , and \hat{n}_I is the observed event number in the bin I . Denoting by $X \in \mathcal{D}$ the variable along which the events are binned and by \mathcal{D}_I the subdomain of \mathcal{D} defining the bin I , the expected event numbers are given by $n_I = n_{\text{tot}} \int_{\mathcal{D}_I} dx f_X(x)$.

Differential distributions get deformed by detector effects, that typically smear their shape. A general way of modelling the smearing is to write the binning variable as

$$X = X_0(1 + \Delta(X_0)\delta), \tag{35}$$

where $\Delta(X_0)$ is the relative magnitude of the uncertainty at the location X_0 . As a simple example, we assume a model of smearing independent of X – the general case can be treated similarly. The expected number of events in a bin I is given by

$$n_I = n_{\text{tot}} \int_{\mathcal{D}_I} f_X(x(1 + \Delta\delta))dx, \tag{36}$$

where n_{tot} is the expected total number of events. Starting from Eq. (14), one can disentangle the information of shape and total event number,

$$L(\theta, \delta) = L_{\text{tot}}(\theta)L_{\text{shape}}(\theta, \delta), \quad L_{\text{shape}}(\theta, \delta) = \prod_{I=1}^n \left(\frac{n_I(\theta, \delta)}{n_{\text{tot}}(\theta)} \right)^{\hat{n}_I}. \tag{37}$$

Only L_{shape} depends on δ , as this nuisance parameter models a shape deformation. Expanding $\log n$ over Δ at quadratic order gives

$$n_I = n_{I,0} \exp(\Delta_{1,I}\delta + \Delta_{2,I}\delta^2), \tag{38}$$

with

$$n_{I,0} = n_{\text{tot}} \int_{\mathcal{D}_I} f_X(x)dx, \quad \Delta_{1,I} = \Delta \int_{\mathcal{D}_I} x f'_X(x)dx \left(\int_{\mathcal{D}_I} f_X(x)dx \right)^{-1}, \tag{39}$$

$$\Delta_{2,I} = \frac{\Delta^2}{2} \left[\left(\int_{\mathcal{D}_I} x^2 f''_X(x)dx \right) \left(\int_{\mathcal{D}_I} f_X(x)dx \right)^{-1} - \left(\int_{\mathcal{D}_I} x f'_X(x)dx \right)^2 \left(\int_{\mathcal{D}_I} f_X(x)dx \right)^{-2} \right]. \tag{40}$$

As an aside one may notice that when one expands n_I at second order, the quadratic term explicitly shows the effect of smearing. If f is convex (concave) over the bin, then $f''_X > 0$ ($f''_X < 0$), so that the quadratic term fills (depletes) the bin, accordingly to what is expected from a smearing process.

Plugging this expression into the likelihood gives exactly¹⁴

$$L_{\text{shape}} = L_{\text{shape}}^{\text{stat}} e^{\xi\delta - \eta\delta^2/2}, \tag{41}$$

with $\xi = \sum_I \Delta_{1,I}$, $\eta = 2 \sum_I \Delta_{2,I}$. Marginalising with a Gaussian prior for δ gives once again Eq. (29), here in a one-variable version:

¹⁴ Interestingly, no Taylor expansion of the likelihood is needed to get this result.

$$\tilde{L}_{\text{shape}} = L_{\text{shape}}^{\text{stat}} \frac{1}{\sqrt{\eta + 1}} \exp\left(\frac{\xi^2}{2(\eta + 1)}\right). \quad (42)$$

Finally, the unbinned version of the same likelihood is directly obtained by taking the limit of infinitely thin bins. Then \hat{n}_I can be only zero or one, the integrals can be simplified and the observed events end up labelled by their position \hat{X}_I . One gets

$$\Delta_{1,I} = \Delta \hat{X}_I \frac{f'_X(\hat{X}_I)}{f_X(\hat{X}_I)}, \quad \Delta_{2,I} = \frac{\Delta^2}{2} (\hat{X}_I)^2 \left(\frac{f''_X(\hat{X}_I)}{f_X(\hat{X}_I)} - \frac{f_X'^2(\hat{X}_I)}{f_X^2(\hat{X}_I)} \right), \quad (43)$$

from which the marginal likelihood follows. The information that has to be reported to reconstruct this smeared likelihood is:

- The magnitude of the relative uncertainty on the binning variable X ;
- The first and second derivatives of the expected shape f_X .

4. An example of signal strength fit

In order to illustrate our results, we consider a somewhat realistic scenario for the characterisation of a signal. Formally, the scenario considered corresponds to a particular case of the signal strength analysis described in Sec. 3.1. This example will also be used to check the accuracy of the approximate marginal likelihood, \tilde{L} .

To carry out this toy analysis one first has to setup the “observed” data, the expected background and the systematic uncertainties. We assume three independent observation channels $I = (A, B, C)$. An observed event number $\hat{n}_{A,B,C}$ is assumed for each channel. The expected event number is given a signal+background form $n = s + b$, which is taken to be common to all channels, so that $s_I = s$, $b_I = b$, and the nominal value of b is fixed.

We further assume the presence of 3 independent systematic uncertainties labelled a, b, c for the signal and 2 uncertainties labelled d, e for the background. Both signal and background are positive, so that we use the error modelisation of Eq. (5). We are going to consider cases of a flat and Gaussian prior for δ , which imply respectively log-normal and log-flat priors for the s, b components of the expected event number.

Assuming a first step of propagation from more elementary uncertainties, and disregarding the possible Δ_2 terms for simplicity, the leading effect of the systematic uncertainties on the expected event number is characterised by 3×5 numbers $\Delta_{1,I}^{s/b,i}$, given in Table 1. In the notation of Sec. 2.5, one has $N = 3$, $p = 5$. Note that with this starting point, the magnitude of the p elementary uncertainties are already combined with the $N \times p$ derivatives of the expected signal. The uncertainties appear in each channel I as $s = s_0 \exp(\Delta_{1,I}^s \cdot \delta^s)$, $b = b_0 \exp(\Delta_{1,I}^b \cdot \delta^b)$.

Having characterised the main effect of the systematic uncertainties, we can readily use the approximate marginal likelihood \tilde{L} . Moreover, we compute the exact marginal likelihood by numerically integrating over the five nuisance parameters,¹⁵ which provides a way of testing directly the accuracy of \tilde{L} . In practice, plotting the exact marginal likelihood of our example takes about an hour on a laptop with average specifications, while plotting the approximate likelihood is instantaneous.

All the numbers assumed for \hat{n}_I and the $\Delta_{1,I}^{s/b,i}$ are given in Table 1. The observed numbers are chosen so that the statistical uncertainty be of $O(10\%)$. This is for example the case in the

¹⁵ Note that for a flat distribution, $V(\delta) = 1$ implies $\delta \in [-\sqrt{3}, \sqrt{3}]$.

Table 1
Observed data and nuisance parameters in three statistically independent channels A, B, C. The average expected background is taken to be $b_0 = 100$.

Channel	A	B	C
\hat{n}	280	310	320
$\Delta^{s,1}$	10%	−10%	10%
$\Delta^{s,2}$	5%	10%	10%
$\Delta^{s,3}$	0%	−5%	−5%
$\Delta^{b,1}$	−10%	−10%	10%
$\Delta^{b,2}$	5%	10%	−10%




Fig. 2. Likelihood functions drawn from data of Table 1. Plain line: Exact marginal likelihood, evaluated numerically. Dashed line: Approximate marginal likelihood, evaluated analytically. Dotted line: likelihood with no systematic uncertainties, $L^{\text{stat}}(s) \equiv L(s, \delta = 0)$.

global fit of the 8 TeV Higgs signal strengths. The signs of the systematic uncertainties have been chosen so that the likelihood strongly depends on every nuisance parameter. The relative magnitudes of the elementary uncertainties are chosen to be $O(10\%)$.

This illustrative scenario can be used in order to check the accuracy of the two kinds of approximations required to obtain the \tilde{L} likelihood: *a*) the CLT-based Gaussian approximation (see Sec. 2) and *b*) the likelihood expansion (see Sec. 3). The Gaussian and flat priors allow us to disentangle between these two approximations, because for Gaussian priors, approximation *a*) is always satisfied, *i.e.* the CLT is perfectly convergent. Thus for Gaussian priors, the discrepancies between the approximate and local likelihoods come only from approximation *b*). In contrast, in the flat prior case the discrepancies come from both approximations *a*) and *b*).

The exact and approximate marginal likelihoods are shown in Fig. 2. In the case of Gaussian priors for δ (*i.e.* log-normal uncertainties), the two curves agree very well. This shows that approximation *b*) is well under control. In order to test approximation *a*), *i.e.* the CLT convergence, we can now compare with the case of flat priors for δ . It turns out that a mild discrepancy appears in certain regions. This illustrates the degree of convergence of the CLT in a five-parameters case. The two curves agree still fairly well in this flat-prior case, in the sense that the best-fit regions drawn from these likelihoods would be similar. For a larger number of nuisance parameters, this discrepancy is expected to decrease as the CLT convergence should improve.

5. Conclusion

With the goal of simplifying the treatment of systematic uncertainties in typical LHC analyses, we have studied the behaviour of a generic likelihood in the presence of a large number of uncertainties with small relative magnitudes.

Whenever this condition is satisfied, it turns out that well-controlled approximations become available, which provide a way of drastically simplifying the incorporation of systematic uncertainties into the likelihood. Our demonstration is split into steps of error propagation and error combination. In the latter, the Lyapunov central limit theorem applies to the combined uncertainties, thereby approximating their joint distribution as a multivariate normal. This implies that the shape of the priors of the elementary uncertainties is irrelevant – only their magnitudes matter.

Whenever the combined uncertainties are small enough, say $\lesssim 20\%$, the likelihood can be further simplified and the complete marginal likelihood is obtained analytically. This general result is applied to the important cases of signal strength characterisation and differential distribution smearing.

For illustration, we present a toy-analysis of signal strength characterisation including systematic uncertainties on signal and background. The approximate and exact marginal likelihoods are found to be in fairly good agreement in this example, implying that all approximations are well under control.

Beyond the obvious gain of avoiding heavy numerical marginalisation, another practical matter is the communication of systematic uncertainties, for example from an experiment to the public. Our approach implies that all the needed information is encoded into a finite set of numbers, namely the relative magnitude of elementary uncertainties and the derivatives of the expected event numbers. The transmission of this information is straightforward, and gives a fairly human-readable summary of the systematic uncertainties. In principle, this simple method could be used to make public the detector effects that are included in LHC analyses.

The marginal likelihood presented in this paper is purely Bayesian. It is also possible to compute analytically the marginal likelihood in case of a frequentist profiling (described in [App. A](#)), as well as to apply the bias methods formalised in Ref. [14]. Finally, although our study is oriented towards LHC analyses, it could also be readily applied into other experimental contexts.

Acknowledgements

I would like to thank G. Moreau, S. Kraml, E. Ponton, G. von Gersdorff for useful discussions and comments. This work is supported by the São Paulo Research Foundation (FAPESP) under grant 2014/21477-2.

Appendix A. Analytic frequentist marginalisation

In [Sec. 3](#), a frequentist marginalisation of the combined uncertainties can also be done. This is obtained by substituting [Eq. \(28\)](#) with

$$\tilde{L}(\theta)\text{freq} = \max_{\bar{\delta}} \left[\bar{L}(\theta, \bar{\delta}) \bar{\pi}(\bar{\delta}) \right]. \quad (44)$$

The resulting approximate likelihood is

$$\tilde{L}(\theta)\text{freq} = L^{\text{stat}}(\theta) \exp \left(\frac{1}{2} \xi \cdot \left(\eta + \rho^{-1} \right)^{-1} \cdot \xi \right), \quad (45)$$

which resembles very much the Bayesian result up to a $\sqrt{|\eta\rho + 1|}$ factor. Typically, the variation of this factor with respect to the parameters of interest is small compared to the variation of the exponential term. Hence, in practice, the frequentist and Bayesian approximate likelihoods are almost equivalent. The subsequent frequentist and Bayesian best-fit regions obtained from these likelihoods thus differ mostly by the definition of frequentist and Bayesian contours [15].

References

- [1] W. Verkerke, D.P. Kirkby, The RooFit toolkit for data modeling, eConf C0303241 (2003) MOLT007, arXiv:physics/0306116.
- [2] L. Moneta, et al., The RooStats Project, PoS ACAT2010 (2010) 057, arXiv:1009.1003 [physics.data-an].
- [3] S. Kraml, et al., Searches for new physics: Les Houches recommendations for the presentation of LHC results, Eur. Phys. J. C 72 (2012) 1976, arXiv:1203.2489 [hep-ph].
- [4] F. Boudjema, et al., On the presentation of the LHC Higgs results, arXiv:1307.5865 [hep-ph].
- [5] K. Cranmer, S. Kreiss, D. Lopez-Val, T. Plehn, Phys. Rev. D 91 (5) (2015) 054032, arXiv:1401.0080 [hep-ph].
- [6] A.C. Berry, The accuracy of the Gaussian approximation to the sum of independent variates, Trans. Am. Math. Soc. 49 (1) (1941) 122–136.
- [7] C.G. Esseen, On the Liapunoff limit of error in the theory of probability, Ark. Mat. Astron. Fys. A 28 (9) (1942) 1–19.
- [8] W. Feller, An Introduction to Probability Theory and Its Applications, vol. 2, 3rd ed., Wiley, New York, 1971, pp. 257–258.
- [9] S. Fichet, G. von Gersdorff, Effective theory for neutral resonances and a statistical dissection of the ATLAS diboson excess, J. High Energy Phys. 1512 (2015) 089, arXiv:1508.04814 [hep-ph].
- [10] G. Aad, et al., ATLAS Collaboration, Search for high-mass diboson resonances with boson-tagged jets in proton–proton collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector, J. High Energy Phys. 1512 (2015) 055, arXiv:1506.00962 [hep-ex].
- [11] A. Arbey, S. Fichet, F. Mahmoudi, G. Moreau, The correlation matrix of Higgs rates at the LHC, arXiv:1606.00455 [hep-ph].
- [12] P. Billingsley, Probability and Measure, 2nd edition, John Wiley and Sons, Inc., New York, 1995.
- [13] J.W. Lindeberg, Über den zentralen Grenzwertsatz der Wahrscheinlichkeitsrechnung, II, Math. Z. 15 (1) (1922) 211–225.
- [14] S. Fichet, G. Moreau, Anatomy of the Higgs fits: a first guide to statistical treatments of the theoretical uncertainties, arXiv:1509.00472 [hep-ph].
- [15] K.A. Olive, et al., Particle Data Group Collaboration, Review of particle physics, Chin. Phys. C 38 (2014) 090001.