

38. STATISTICS

Revised September 2013 by G. Cowan (RHUL).

This chapter gives an overview of statistical methods used in high-energy physics. In statistics, we are interested in using a given sample of data to make inferences about a probabilistic model, *e.g.*, to assess the model’s validity or to determine the values of its parameters. There are two main approaches to statistical inference, which we may call frequentist and Bayesian.

In frequentist statistics, probability is interpreted as the frequency of the outcome of a repeatable experiment. The most important tools in this framework are parameter estimation, covered in Section 38.2, statistical tests, discussed in Section 38.3, and confidence intervals, which are constructed so as to cover the true value of a parameter with a specified probability, as described in Section 38.4.2. Note that in frequentist statistics one does not define a probability for a hypothesis or for the value of a parameter.

In Bayesian statistics, the interpretation of probability is more general and includes *degree of belief* (called subjective probability). One can then speak of a probability density function (p.d.f.) for a parameter, which expresses one’s state of knowledge about where its true value lies. Bayesian methods provide a natural means to include additional information, which in general may be subjective; in fact they *require* prior probabilities for the hypotheses (or parameters) in question, *i.e.*, the degree of belief about the parameters’ values before carrying out the measurement. Using Bayes’ theorem (Eq. (37.4)), the prior degree of belief is updated by the data from the experiment. Bayesian methods for interval estimation are discussed in Sections 38.4.1 and 38.4.2.4.

For many inference problems, the frequentist and Bayesian approaches give similar numerical values, even though they answer different questions and are based on fundamentally different interpretations of probability. In some important cases, however, the two approaches may yield very different results. For a discussion of Bayesian vs. non-Bayesian methods, see references written by a statistician [1], by a physicist [2], or the more detailed comparison in Ref. 3.

Following common usage in physics, the word “error” is often used in this chapter to mean “uncertainty.” More specifically it can indicate the size of an interval as in “the standard error” or “error propagation,” where the term refers to the standard deviation of an estimator.

38.1. Fundamental concepts

Consider an experiment whose outcome is characterized by one or more data values, which we can write as a vector \mathbf{x} . A *hypothesis* H is a statement about the probability for the data, often written $P(\mathbf{x}|H)$. (We will usually use a capital letter for a probability and lower case for a probability density. Often the term p.d.f. is used loosely to refer to either a probability or a probability density.) This could, for example, define completely the p.d.f. for the data (a simple hypothesis), or it could specify only the functional form of the p.d.f., with the values of one or more parameters not determined (a composite hypothesis).

2 38. Statistics

If the probability $P(\mathbf{x}|H)$ for data \mathbf{x} is regarded as a function of the hypothesis H , then it is called the *likelihood* of H , usually written $L(H)$. Often the hypothesis is characterized by one or more parameters θ , in which case $L(\theta) = P(\mathbf{x}|\theta)$ is called the likelihood function.

In some cases one can obtain at least approximate frequentist results using the likelihood evaluated only with the data obtained. In general, however, the frequentist approach requires a full specification of the probability model $P(\mathbf{x}|H)$ both as a function of the data \mathbf{x} and hypothesis H .

In the Bayesian approach, inference is based on the posterior probability for H given the data \mathbf{x} , which represents one's degree of belief that H is true given the data. This is obtained from Bayes' theorem (37.4), which can be written

$$P(H|\mathbf{x}) = \frac{P(\mathbf{x}|H)\pi(H)}{\int P(\mathbf{x}|H')\pi(H') dH'}. \quad (38.1)$$

Here $P(\mathbf{x}|H)$ is the likelihood for H , which depends only on the data actually obtained. The quantity $\pi(H)$ is the prior probability for H , which represents one's degree of belief for H before carrying out the measurement. The integral in the denominator (or sum, for discrete hypotheses) serves as a normalization factor. If H is characterized by a continuous parameter θ then the posterior probability is a p.d.f. $p(\theta|\mathbf{x})$. Note that the likelihood function itself is not a p.d.f. for θ .

38.2. Parameter estimation

Here we review *point estimation* of parameters, first with an overview of the frequentist approach and its two most important methods, maximum likelihood and least squares, treated in Sections 38.2.2 and 38.2.3. The Bayesian approach is outlined in Sec. 38.2.4.

An *estimator* $\hat{\theta}$ (written with a hat) is a function of the data used to estimate the value of the parameter θ . Sometimes the word 'estimate' is used to denote the value of the estimator when evaluated with given data. There is no fundamental rule dictating how an estimator must be constructed. One tries, therefore, to choose that estimator which has the best properties. The most important of these are (a) *consistency*, (b) *bias*, (c) *efficiency*, and (d) *robustness*.

(a) An estimator is said to be *consistent* if the estimate $\hat{\theta}$ converges to the true value θ as the amount of data increases. This property is so important that it is possessed by all commonly used estimators.

(b) The *bias*, $b = E[\hat{\theta}] - \theta$, is the difference between the expectation value of the estimator and the true value of the parameter. The expectation value is taken over a hypothetical set of similar experiments in which $\hat{\theta}$ is constructed in the same way. When $b = 0$, the estimator is said to be unbiased. The bias depends on the chosen metric, *i.e.*, if $\hat{\theta}$ is an unbiased estimator of θ , then $\hat{\theta}^2$ is not in general an unbiased estimator for θ^2 .

(c) *Efficiency* is the ratio of the minimum possible variance for any estimator of θ to the variance $V[\hat{\theta}]$ of the estimator $\hat{\theta}$. For the case of a single parameter, under rather general

conditions the minimum variance is given by the Rao-Cramér-Fréchet bound,

$$\sigma_{\min}^2 = \left(1 + \frac{\partial b}{\partial \theta}\right)^2 / I(\theta), \quad (38.2)$$

where

$$I(\theta) = E \left[\left(\frac{\partial \ln L}{\partial \theta} \right)^2 \right] = -E \left[\frac{\partial^2 \ln L}{\partial \theta^2} \right] \quad (38.3)$$

is the *Fisher information*, L is the likelihood, and the expectation value in (38.3) is carried out with respect to the data. For the final equality to hold, the range of allowed data values must not depend on θ .

The *mean-squared error*,

$$\text{MSE} = E[(\hat{\theta} - \theta)^2] = V[\hat{\theta}] + b^2, \quad (38.4)$$

is a measure of an estimator's quality which combines bias and variance.

(d) *Robustness* is the property of being insensitive to departures from assumptions in the p.d.f., e.g., owing to uncertainties in the distribution's tails.

It is not in general possible to optimize simultaneously for all the measures of estimator quality described above. For example, there is in general a trade-off between bias and variance. For some common estimators, the properties above are known exactly. More generally, it is possible to evaluate them by Monte Carlo simulation. Note that they will often depend on the unknown θ .

38.2.1. Estimators for mean, variance, and median :

Suppose we have a set of n independent measurements, x_1, \dots, x_n , each assumed to follow a p.d.f. with unknown mean μ and unknown variance σ^2 . The measurements do not necessarily have to follow a Gaussian distribution. Then

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n x_i \quad (38.5)$$

$$\hat{\sigma}^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \hat{\mu})^2 \quad (38.6)$$

are unbiased estimators of μ and σ^2 . The variance of $\hat{\mu}$ is σ^2/n and the variance of $\hat{\sigma}^2$ is

$$V[\hat{\sigma}^2] = \frac{1}{n} \left(m_4 - \frac{n-3}{n-1} \sigma^4 \right), \quad (38.7)$$

where m_4 is the 4th central moment of x (see Eq. (37.8b)). For Gaussian distributed x_i , this becomes $2\sigma^4/(n-1)$ for any $n \geq 2$, and for large n the standard deviation of $\hat{\sigma}$ (the “error of the error”) is $\sigma/\sqrt{2n}$. For any n and Gaussian x_i , $\hat{\mu}$ is an efficient estimator for

4 38. Statistics

μ , and the estimators $\hat{\mu}$ and $\hat{\sigma}^2$ are uncorrelated. Otherwise the arithmetic mean (38.5) is not necessarily the most efficient estimator; this is discussed further in Sec. 8.7 of Ref. 4.

If σ^2 is known, it does not improve the estimate $\hat{\mu}$, as can be seen from Eq. (38.5); however, if μ is known, one can substitute it for $\hat{\mu}$ in Eq. (38.6) and replace $n - 1$ by n to obtain an estimator of σ^2 still with zero bias but smaller variance. If the x_i have different, known variances σ_i^2 , then the weighted average

$$\hat{\mu} = \frac{1}{w} \sum_{i=1}^n w_i x_i, \quad (38.8)$$

where $w_i = 1/\sigma_i^2$ and $w = \sum_i w_i$, is an unbiased estimator for μ with a smaller variance than an unweighted average. The standard deviation of $\hat{\mu}$ is $1/\sqrt{w}$.

As an estimator for the median x_{med} , one can use the value \hat{x}_{med} such that half the x_i are below and half above (the sample median). If the sample median lies between two observed values, it is set by convention halfway between them. If the p.d.f. of x has the form $f(x - \mu)$ and μ is both mean and median, then for large n the variance of the sample median approaches $1/[4nf^2(0)]$, provided $f(0) > 0$. Although estimating the median can often be more difficult computationally than the mean, the resulting estimator is generally more robust, as it is insensitive to the exact shape of the tails of a distribution.

38.2.2. The method of maximum likelihood :

Suppose we have a set of measured quantities \mathbf{x} and the likelihood $L(\boldsymbol{\theta}) = P(\mathbf{x}|\boldsymbol{\theta})$ for a set of parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_N)$. The *maximum likelihood* (ML) estimators for $\boldsymbol{\theta}$ are defined as the values that give the maximum of L . Because of the properties of the logarithm, it is usually easier to work with $\ln L$, and since both are maximized for the same parameter values $\boldsymbol{\theta}$, the ML estimators can be found by solving the *likelihood equations*,

$$\frac{\partial \ln L}{\partial \theta_i} = 0, \quad i = 1, \dots, N. \quad (38.9)$$

Often the solution must be found numerically. Maximum likelihood estimators are important because they are asymptotically unbiased and efficient for large data samples, under quite general conditions, and the method has a wide range of applicability.

In general the likelihood function is obtained from the probability of the data under assumption of the parameters. An important special case is when the data consist of *i.i.d.* (independent and identically distributed) values. Here one has a set of n statistically independent quantities $\mathbf{x} = (x_1, \dots, x_n)$, where each component follows the same p.d.f. $f(x; \boldsymbol{\theta})$. In this case the joint p.d.f. of the data sample factorizes and the likelihood function is

$$L(\boldsymbol{\theta}) = \prod_{i=1}^n f(x_i; \boldsymbol{\theta}). \quad (38.10)$$

In this case the number of events n is regarded as fixed. If however the probability to observe n events itself depends on the parameters $\boldsymbol{\theta}$, then this should be included in

the likelihood. For example, if n follows a Poisson distribution with mean μ and the independent x values all follow $f(x; \boldsymbol{\theta})$, then the likelihood becomes

$$L(\boldsymbol{\theta}) = \frac{\mu^n}{n!} e^{-\mu} \prod_{i=1}^n f(x_i; \boldsymbol{\theta}). \quad (38.11)$$

Equation. (38.11) is often called the *extended likelihood* (see, *e.g.*, Refs. [6–8]). In general μ is a function of $\boldsymbol{\theta}$, and including the probability for n given $\boldsymbol{\theta}$ in the likelihood provides additional information about the parameters and thus leads to a reduction in their statistical uncertainties.

In evaluating the likelihood function, it is important that any normalization factors in the p.d.f. that involve $\boldsymbol{\theta}$ be included. However, we will only be interested in the maximum of L and in ratios of L at different values of the parameters; hence any multiplicative factors that do not involve the parameters that we want to estimate may be dropped, including factors that depend on the data but not on $\boldsymbol{\theta}$.

Under a one-to-one change of parameters from $\boldsymbol{\theta}$ to $\boldsymbol{\eta}$, the ML estimators $\hat{\boldsymbol{\theta}}$ transform to $\boldsymbol{\eta}(\hat{\boldsymbol{\theta}})$. That is, the ML solution is invariant under change of parameter. However, other properties of ML estimators, in particular the bias, are not invariant under change of parameter.

The inverse V^{-1} of the covariance matrix $V_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$ for a set of ML estimators can be estimated by using

$$(\hat{V}^{-1})_{ij} = - \left. \frac{\partial^2 \ln L}{\partial \theta_i \partial \theta_j} \right|_{\hat{\boldsymbol{\theta}}}; \quad (38.12)$$

for finite samples, however, Eq. (38.12) can result in an underestimate of the variances. In the large sample limit (or in a linear model with Gaussian errors), L has a Gaussian form and $\ln L$ is (hyper)parabolic. In this case, it can be seen that a numerically equivalent way of determining s -standard-deviation errors is from the hypersurface defined by the $\boldsymbol{\theta}'$ such that

$$\ln L(\boldsymbol{\theta}') = \ln L_{\max} - s^2/2, \quad (38.13)$$

where $\ln L_{\max}$ is the value of $\ln L$ at the solution point (compare with Eq. (38.68)). The minimum and maximum values of θ_i on the hypersurface then give an approximate s -standard deviation confidence interval for θ_i (see Section 38.4.2.2).

38.2.2.1. ML with binned data:

If the total number of data values x_i , $i = 1, \dots, n_{\text{tot}}$, is small, the unbinned maximum likelihood method, *i.e.*, use of equation (38.10) (or (38.11) for extended ML), is preferred since binning can only result in a loss of information, and hence larger statistical errors for the parameter estimates. If the sample is large, it can be convenient to bin the values in a histogram with N bins, so that one obtains a vector of data $\mathbf{n} = (n_1, \dots, n_N)$ with expectation values $\boldsymbol{\mu} = E[\mathbf{n}]$ and probabilities $f(\mathbf{n}; \boldsymbol{\mu})$. Suppose the mean values $\boldsymbol{\mu}$ can be determined as a function of a set of parameters $\boldsymbol{\theta}$. Then one may maximize the likelihood function based on the contents of the bins.

6 38. Statistics

As mentioned in Sec. 38.2.2, the total number of events $n_{\text{tot}} = \sum_i n_i$ can be regarded as fixed or as a random variable. If it is fixed, the histogram follows a multinomial distribution,

$$f_{\text{M}}(\mathbf{n}; \boldsymbol{\theta}) = \frac{n_{\text{tot}}!}{n_1! \cdots n_N!} p_1^{n_1} \cdots p_N^{n_N}, \quad (38.14)$$

where we assume the probabilities p_i are given functions of the parameters $\boldsymbol{\theta}$. The distribution can be written equivalently in terms of the expected number of events in each bin, $\mu_i = n_{\text{tot}} p_i$. If the n_i are regarded as independent and Poisson distributed, then the data are described by a product of Poisson probabilities,

$$f_{\text{P}}(\mathbf{n}; \boldsymbol{\theta}) = \prod_{i=1}^N \frac{\mu_i^{n_i}}{n_i!} e^{-\mu_i}, \quad (38.15)$$

where the mean values μ_i are given functions of $\boldsymbol{\theta}$. The total number of events n_{tot} thus follows a Poisson distribution with mean $\mu_{\text{tot}} = \sum_i \mu_i$.

When using maximum likelihood with binned data, one can find the ML estimators and at the same time obtain a statistic usable for a test of goodness-of-fit (see Sec. 38.3.2). Maximizing the likelihood $L(\boldsymbol{\theta}) = f_{\text{M/P}}(\mathbf{n}; \boldsymbol{\theta})$ is equivalent to maximizing the likelihood ratio $\lambda(\boldsymbol{\theta}) = f_{\text{M/P}}(\mathbf{n}; \boldsymbol{\theta})/f(\mathbf{n}; \hat{\boldsymbol{\mu}})$, where in the denominator $f(\mathbf{n}; \boldsymbol{\mu})$ is a model with an adjustable parameter for each bin, $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)$, and the corresponding estimators are $\hat{\boldsymbol{\mu}} = (n_1, \dots, n_N)$. Often one minimizes instead the equivalent quantity $-2 \ln \lambda(\boldsymbol{\theta})$. For independent Poisson distributed n_i this is [9]

$$-2 \ln \lambda(\boldsymbol{\theta}) = 2 \sum_{i=1}^N \left[\mu_i(\boldsymbol{\theta}) - n_i + n_i \ln \frac{n_i}{\mu_i(\boldsymbol{\theta})} \right], \quad (38.16)$$

where for bins with $n_i = 0$, the last term in (38.16) is zero. The expression (38.16) without the terms $\mu_i - n_i$ also gives $-2 \ln \lambda(\boldsymbol{\theta})$ for multinomially distributed n_i , *i.e.*, when the total number of entries is regarded as fixed. In the limit of zero bin width, maximizing (38.16) is equivalent to maximizing the unbinned extended likelihood function (38.11) or in the multinomial case without the $\mu_i - n_i$ terms one obtains Eq. (38.10).

A smaller value of $-2 \ln \lambda(\hat{\boldsymbol{\theta}})$ corresponds to better agreement between the data and the hypothesized form of $\boldsymbol{\mu}(\boldsymbol{\theta})$. The value of $-2 \ln \lambda(\hat{\boldsymbol{\theta}})$ can thus be translated into a p -value as a measure of goodness-of-fit, as described in Sec. 38.3.2. Assuming the model is correct, then according to Wilks' theorem, for sufficiently large μ_i and providing certain regularity conditions are met, the minimum of $-2 \ln \lambda$ as defined by Eq. (38.16) follows a χ^2 distribution (see, *e.g.*, Ref. 9). If there are N bins and m fitted parameters, then the number of degrees of freedom for the χ^2 distribution is $N - m$ if the data are treated as Poisson-distributed, and $N - m - 1$ if the n_i are multinomially distributed.

Suppose the n_i are Poisson-distributed and the overall normalization $\mu_{\text{tot}} = \sum_i \mu_i$ is taken as an adjustable parameter, so that $\mu_i = \mu_{\text{tot}} p_i(\boldsymbol{\theta})$, where the probability to be in the i th bin, $p_i(\boldsymbol{\theta})$, does not depend on μ_{tot} . Then by minimizing Eq. (38.16), one obtains that the area under the fitted function is equal to the sum of the histogram contents, *i.e.*, $\sum_i \mu_i = \sum_i n_i$.

38.2.2.2. Frequentist treatment of nuisance parameters:

Suppose we want to determine the values of parameters $\boldsymbol{\theta}$ using a set of measurements \boldsymbol{x} described by a probability model $P_x(\boldsymbol{x}|\boldsymbol{\theta})$. In general the model is not perfect, which is to say it can not provide an accurate description of the data even at the most optimal point of its parameter space. As a result, the estimated parameters can have a systematic bias.

One can in general improve the model by including in it additional parameters. That is, it is extended to $P_x(\boldsymbol{x}|\boldsymbol{\theta}, \boldsymbol{\nu})$, which depends on parameters of interest $\boldsymbol{\theta}$ and *nuisance parameters* $\boldsymbol{\nu}$. The additional parameters are not of intrinsic interest but must be included for the model to be accurate for some point in the enlarged parameter space.

Although including additional parameters may eliminate or at least reduce the effect of systematic uncertainties, their presence will result in increased statistical uncertainties for the parameters of interest. This occurs because the estimators for the nuisance parameters and those of interest will in general be correlated, which results in an enlargement of the contour defined by Eq. (38.13).

To reduce the impact of the nuisance parameters one often tries to constrain their values by means of control or calibration measurements, say, having data \boldsymbol{y} . For example, some components of \boldsymbol{y} could represent estimates of the nuisance parameters, often from separate experiments. Suppose the measurements \boldsymbol{y} are statistically independent from \boldsymbol{x} and are described by a model $P_y(\boldsymbol{y}|\boldsymbol{\nu})$. The joint model for both \boldsymbol{x} and \boldsymbol{y} is in this case therefore the product of the probabilities for \boldsymbol{x} and \boldsymbol{y} , and thus the likelihood function for the full set of parameters is

$$L(\boldsymbol{\theta}, \boldsymbol{\nu}) = P_x(\boldsymbol{x}|\boldsymbol{\theta}, \boldsymbol{\nu})P_y(\boldsymbol{y}|\boldsymbol{\nu}). \quad (38.17)$$

Note that in this case if one wants to simulate the experiment by means of Monte Carlo, both the primary and control measurements, \boldsymbol{x} and \boldsymbol{y} , must be generated for each repetition under assumption of fixed values for the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\nu}$.

Using all of the parameters $(\boldsymbol{\theta}, \boldsymbol{\nu})$ in Eq. (38.13) to find the statistical errors in the parameters of interest $\boldsymbol{\theta}$ is equivalent to using the *profile likelihood*, which depends only on $\boldsymbol{\theta}$. It is defined as

$$L_p(\boldsymbol{\theta}) = L(\boldsymbol{\theta}, \widehat{\boldsymbol{\nu}}(\boldsymbol{\theta})), \quad (38.18)$$

where the double-hat notation indicates the profiled values of the parameters $\boldsymbol{\nu}$, defined as the values that maximize L for the specified $\boldsymbol{\theta}$. The profile likelihood is discussed further in Section 38.3.2.1 in connection with hypothesis tests.

38.2.3. The method of least squares :

The *method of least squares* (LS) coincides with the method of maximum likelihood in the following special case. Consider a set of N independent measurements y_i at known points x_i . The measurement y_i is assumed to be Gaussian distributed with mean $\mu(x_i; \boldsymbol{\theta})$ and known variance σ_i^2 . The goal is to construct estimators for the unknown parameters $\boldsymbol{\theta}$. The likelihood function contains the sum of squares

$$\chi^2(\boldsymbol{\theta}) = -2 \ln L(\boldsymbol{\theta}) + \text{constant} = \sum_{i=1}^N \frac{(y_i - \mu(x_i; \boldsymbol{\theta}))^2}{\sigma_i^2}. \quad (38.19)$$

8 38. Statistics

The parameter values that maximize L are the same as those which minimize χ^2 .

The minimum of Equation (38.19) defines the least-squares estimators $\hat{\boldsymbol{\theta}}$ for the more general case where the y_i are not Gaussian distributed as long as they are independent. If they are not independent but rather have a covariance matrix $V_{ij} = \text{cov}[y_i, y_j]$, then the LS estimators are determined by the minimum of

$$\chi^2(\boldsymbol{\theta}) = (\mathbf{y} - \boldsymbol{\mu}(\boldsymbol{\theta}))^T V^{-1} (\mathbf{y} - \boldsymbol{\mu}(\boldsymbol{\theta})) , \quad (38.20)$$

where $\mathbf{y} = (y_1, \dots, y_N)$ is the (column) vector of measurements, $\boldsymbol{\mu}(\boldsymbol{\theta})$ is the corresponding vector of predicted values, and the superscript T denotes the transpose.

Often one further restricts the problem to the case where $\mu(x_i; \boldsymbol{\theta})$ is a linear function of the parameters, *i.e.*,

$$\mu(x_i; \boldsymbol{\theta}) = \sum_{j=1}^m \theta_j h_j(x_i) . \quad (38.21)$$

Here the $h_j(x)$ are m linearly independent functions, *e.g.*, $1, x, x^2, \dots, x^{m-1}$ or Legendre polynomials. We require $m < N$ and at least m of the x_i must be distinct.

Minimizing χ^2 in this case with m parameters reduces to solving a system of m linear equations. Defining $H_{ij} = h_j(x_i)$ and minimizing χ^2 by setting its derivatives with respect to the θ_i equal to zero gives the LS estimators,

$$\hat{\boldsymbol{\theta}} = (H^T V^{-1} H)^{-1} H^T V^{-1} \mathbf{y} \equiv D \mathbf{y} . \quad (38.22)$$

The covariance matrix for the estimators $U_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$ is given by

$$U = D V D^T = (H^T V^{-1} H)^{-1} , \quad (38.23)$$

or equivalently, its inverse U^{-1} can be found from

$$(U^{-1})_{ij} = \left. \frac{1}{2} \frac{\partial^2 \chi^2}{\partial \theta_i \partial \theta_j} \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} = \sum_{k,l=1}^N h_i(x_k) (V^{-1})_{kl} h_j(x_l) . \quad (38.24)$$

The LS estimators can also be found from the expression

$$\hat{\boldsymbol{\theta}} = U \mathbf{g} , \quad (38.25)$$

where the vector \mathbf{g} is defined by

$$g_i = \sum_{j,k=1}^N y_j h_i(x_k) (V^{-1})_{jk} . \quad (38.26)$$

For the case of uncorrelated y_i , for example, one can use (38.25) with

$$(U^{-1})_{ij} = \sum_{k=1}^N \frac{h_i(x_k)h_j(x_k)}{\sigma_k^2}, \quad (38.27)$$

$$g_i = \sum_{k=1}^N \frac{y_k h_i(x_k)}{\sigma_k^2}. \quad (38.28)$$

Expanding $\chi^2(\boldsymbol{\theta})$ about $\hat{\boldsymbol{\theta}}$, one finds that the contour in parameter space defined by

$$\chi^2(\boldsymbol{\theta}) = \chi^2(\hat{\boldsymbol{\theta}}) + 1 = \chi_{\min}^2 + 1 \quad (38.29)$$

has tangent planes located at approximately plus-or-minus-one standard deviation $\sigma_{\hat{\theta}}$ from the LS estimates $\hat{\boldsymbol{\theta}}$.

In constructing the quantity $\chi^2(\boldsymbol{\theta})$ one requires the variances or, in the case of correlated measurements, the covariance matrix. Often these quantities are not known *a priori* and must be estimated from the data; an important example is where the measured value y_i represents the event count in a histogram bin. If, for example, y_i represents a Poisson variable, for which the variance is equal to the mean, then one can either estimate the variance from the predicted value, $\mu(x_i; \boldsymbol{\theta})$, or from the observed number itself, y_i . In the first option, the variances become functions of the fitted parameters, which may lead to calculational difficulties. The second option can be undefined if y_i is zero, and in both cases for small y_i , the variance will be poorly estimated. In either case, one should constrain the normalization of the fitted curve to the correct value, *i.e.*, one should determine the area under the fitted curve directly from the number of entries in the histogram (see Ref. 8, Section 7.4). As noted in Sec. 38.2.2.1, this issue is avoided when using the method of extended maximum likelihood with binned data by minimizing Eq. (38.16). In that case if the expected number of events μ_{tot} does not depend on the other fitted parameters $\boldsymbol{\theta}$, then its extended ML estimator is equal to the observed total number of events.

As the minimum value of the χ^2 represents the level of agreement between the measurements and the fitted function, it can be used for assessing the goodness-of-fit; this is discussed further in Section 38.3.2.

38.2.4. *The Bayesian approach :*

In the frequentist methods discussed above, probability is associated only with data, not with the value of a parameter. This is no longer the case in Bayesian statistics, however, which we introduce in this section. For general introductions to Bayesian statistics see, *e.g.*, Refs. [22–25].

Suppose the outcome of an experiment is characterized by a vector of data \boldsymbol{x} , whose probability distribution depends on an unknown parameter (or parameters) $\boldsymbol{\theta}$ that we wish to determine. In Bayesian statistics, all knowledge about $\boldsymbol{\theta}$ is summarized by the

posterior p.d.f. $p(\boldsymbol{\theta}|\mathbf{x})$, whose integral over any given region gives the degree of belief for $\boldsymbol{\theta}$ to take on values in that region, given the data \mathbf{x} . It is obtained by using Bayes' theorem,

$$p(\boldsymbol{\theta}|\mathbf{x}) = \frac{P(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\int P(\mathbf{x}|\boldsymbol{\theta}')\pi(\boldsymbol{\theta}') d\boldsymbol{\theta}'}, \quad (38.30)$$

where $P(\mathbf{x}|\boldsymbol{\theta})$ is the likelihood function, *i.e.*, the joint p.d.f. for the data viewed as a function of $\boldsymbol{\theta}$, evaluated with the data actually obtained in the experiment, and $\pi(\boldsymbol{\theta})$ is the prior p.d.f. for $\boldsymbol{\theta}$. Note that the denominator in Eq. (38.30) serves to normalize the posterior p.d.f. to unity.

As it can be difficult to report the full posterior p.d.f. $p(\boldsymbol{\theta}|\mathbf{x})$, one would usually summarize it with statistics such as the mean (or median) values, and covariance matrix. In addition one may construct intervals with a given probability content, as is discussed in Sec. 38.4.1 on Bayesian interval estimation.

38.2.4.1. Priors:

Bayesian statistics supplies no unique rule for determining the prior $\pi(\boldsymbol{\theta})$; this reflects the analyst's subjective degree of belief (or state of knowledge) about $\boldsymbol{\theta}$ before the measurement was carried out. For the result to be of value to the broader community, whose members may not share these beliefs, it is important to carry out a sensitivity analysis, that is, to show how the result changes under a reasonable variation of the prior probabilities.

One might like to construct $\pi(\boldsymbol{\theta})$ to represent complete ignorance about the parameters by setting it equal to a constant. A problem here is that if the prior p.d.f. is flat in $\boldsymbol{\theta}$, then it is not flat for a nonlinear function of $\boldsymbol{\theta}$, and so a different parametrization of the problem would lead in general to a non-equivalent posterior p.d.f.

For the special case of a constant prior, one can see from Bayes' theorem (38.30) that the posterior is proportional to the likelihood, and therefore the mode (peak position) of the posterior is equal to the ML estimator. The posterior mode, however, will change in general upon a transformation of parameter. One may use as the Bayesian estimator a summary statistic other than the mode, such as the median, which is invariant under parameter transformation. But this will not in general coincide with the ML estimator.

The difficult and subjective nature of encoding personal knowledge into priors has led to what is called *objective Bayesian statistics*, where prior probabilities are based not on an actual degree of belief but rather derived from formal rules. These give, for example, priors which are invariant under a transformation of parameters, or ones which result in a maximum gain in information for a given set of measurements. For an extensive review see, *e.g.*, Ref. 26.

Objective priors do not in general reflect degree of belief, but they could in some cases be taken as possible, although perhaps extreme, subjective priors. The posterior probabilities as well therefore do not necessarily reflect a degree of belief. However one may regard investigating a variety of objective priors to be an important part of the sensitivity analysis. Furthermore, use of objective priors with Bayes' theorem can be

viewed as a recipe for producing estimators or intervals which have desirable frequentist properties.

An important procedure for deriving objective priors is due to Jeffreys. According to *Jeffreys' rule* one takes the prior as

$$\pi(\boldsymbol{\theta}) \propto \sqrt{\det(I(\boldsymbol{\theta}))}, \quad (38.31)$$

where

$$I_{ij}(\boldsymbol{\theta}) = -E \left[\frac{\partial^2 \ln P(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_i \partial \theta_j} \right] \quad (38.32)$$

is the *Fisher information matrix*. One can show that the Jeffreys prior leads to inference that is invariant under a transformation of parameters. One should note that the Jeffreys prior depends on the likelihood function, and thus contains information about the measurement model itself, which goes beyond one's degree of belief about the value of a parameter. As examples, the Jeffreys prior for the mean μ of a Gaussian distribution is a constant, and for the mean of a Poisson distribution one finds $\pi(\mu) \propto 1/\sqrt{\mu}$.

Neither the constant nor $1/\sqrt{\mu}$ priors can be normalized to unit area and are therefore said to be *improper*. This can be allowed because the prior always appears multiplied by the likelihood function, and if the likelihood falls to zero sufficiently quickly then one may have a normalizable posterior density.

An important type of objective prior is the reference prior due to Bernardo and Berger [27]. To find the reference prior for a given problem one considers the Kullback-Leibler divergence $D_n[\pi, p]$ of the posterior $p(\boldsymbol{\theta}|\mathbf{x})$ relative to a prior $\pi(\boldsymbol{\theta})$, obtained from a set of i.i.d. data $\mathbf{x} = (x_1, \dots, x_n)$:

$$D_n[\pi, p] = \int p(\boldsymbol{\theta}|\mathbf{x}) \ln \frac{p(\boldsymbol{\theta}|\mathbf{x})}{\pi(\boldsymbol{\theta})} d\boldsymbol{\theta}. \quad (38.33)$$

This is effectively a measure of the gain in information provided by the data. The reference prior is chosen so that the expectation value of this information gain is maximized for the limiting case of $n \rightarrow \infty$, where the expectation is computed with respect to the marginal distribution of the data,

$$p(\mathbf{x}) = \int p(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}) d\boldsymbol{\theta}. \quad (38.34)$$

For a single, continuous parameter the reference prior is usually identical to the Jeffreys prior. In the multiparameter case an iterative algorithm exists, which requires sorting the parameters by order of inferential importance. Often the result does not depend on this order, but when it does, this can be part of a robustness analysis. Further discussion and applications to particle physics problems can be found in Ref. 28.

38.2.4.2. Bayesian treatment of nuisance parameters:

As discussed in Sec. 38.2.2, a model may depend on parameters of interest θ as well as on nuisance parameters ν , which must be included for an accurate description of the data. Knowledge about the values of ν may be supplied by control measurements, theoretical insights, physical constraints, etc. Suppose, for example, one has data \mathbf{y} from a control measurement which is characterized by a probability $P_y(\mathbf{y}|\nu)$. Suppose further that before carrying out the control measurement one's state of knowledge about ν is described by an initial prior $\pi_0(\nu)$, which in practice is often taken to be a constant or in any case very broad. By using Bayes' theorem (38.1) one obtains the updated prior $\pi(\nu)$ (*i.e.*, now $\pi(\nu) = \pi(\nu|\mathbf{y})$), the probability for ν given \mathbf{y}),

$$\pi(\nu|\mathbf{y}) \propto P(\mathbf{y}|\nu)\pi_0(\nu) . \quad (38.35)$$

In the absence of a model for $P(\mathbf{y}|\nu)$ one may make some reasonable but ad hoc choices. For a single nuisance parameter ν , for example, one might characterize the uncertainty in a nuisance parameter ν by a p.d.f. $\pi(\nu)$ centered about its nominal value with a certain standard deviation σ_ν . Often a Gaussian p.d.f. provides a reasonable model for one's degree of belief about a nuisance parameter; in other cases, more complicated shapes may be appropriate. If, for example, the parameter represents a non-negative quantity then a log-normal or gamma p.d.f. can be a more natural choice than a Gaussian truncated at zero. Note also that truncation of the prior of a nuisance parameter ν at zero will in general make $\pi(\nu)$ nonzero at $\nu = 0$, which can lead to an unnormalizable posterior for a parameter of interest that appears multiplied by ν .

The likelihood function, prior, and posterior p.d.f.s then all depend on both θ and ν , and are related by Bayes' theorem, as usual. Note that the likelihood here only refers to the primary measurement \mathbf{x} . Once any control measurements \mathbf{y} are used to find the updated prior $\pi(\nu)$ for the nuisance parameters, this information is fully encapsulated in $\pi(\nu)$ and the control measurements do not appear further.

One can obtain the posterior p.d.f. for θ alone by integrating over the nuisance parameters, *i.e.*,

$$p(\theta|\mathbf{x}) = \int p(\theta, \nu|\mathbf{x}) d\nu . \quad (38.36)$$

Such integrals can often not be carried out in closed form, and if the number of nuisance parameters is large, then they can be difficult to compute with standard Monte Carlo methods. *Markov Chain Monte Carlo* (MCMC) techniques are often used for computing integrals of this type (see Sec. 39.5).

38.2.5. Propagation of errors :

Consider a set of n quantities $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$ and a set of m functions $\boldsymbol{\eta}(\boldsymbol{\theta}) = (\eta_1(\boldsymbol{\theta}), \dots, \eta_m(\boldsymbol{\theta}))$. Suppose we have estimated $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \dots, \hat{\theta}_n)$, using, say, maximum-likelihood or least-squares, and we also know or have estimated the covariance matrix $V_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$. The goal of *error propagation* is to determine the covariance matrix for the functions, $U_{ij} = \text{cov}[\hat{\eta}_i, \hat{\eta}_j]$, where $\hat{\boldsymbol{\eta}} = \boldsymbol{\eta}(\hat{\boldsymbol{\theta}})$. In particular, the diagonal elements $U_{ii} = V[\hat{\eta}_i]$ give the variances. The new covariance matrix can be found by expanding the functions $\boldsymbol{\eta}(\boldsymbol{\theta})$ about the estimates $\hat{\boldsymbol{\theta}}$ to first order in a Taylor series. Using this one finds

$$U_{ij} \approx \sum_{k,l} \left. \frac{\partial \eta_i}{\partial \theta_k} \frac{\partial \eta_j}{\partial \theta_l} \right|_{\hat{\boldsymbol{\theta}}} V_{kl}. \quad (38.37)$$

This can be written in matrix notation as $U \approx AVA^T$ where the matrix of derivatives A is

$$A_{ij} = \left. \frac{\partial \eta_i}{\partial \theta_j} \right|_{\hat{\boldsymbol{\theta}}}, \quad (38.38)$$

and A^T is its transpose. The approximation is exact if $\boldsymbol{\eta}(\boldsymbol{\theta})$ is linear (it holds, for example, in equation (38.23)). If this is not the case, the approximation can break down if, for example, $\boldsymbol{\eta}(\boldsymbol{\theta})$ is significantly nonlinear close to $\hat{\boldsymbol{\theta}}$ in a region of a size comparable to the standard deviations of $\hat{\boldsymbol{\theta}}$.

38.3. Statistical tests

In addition to estimating parameters, one often wants to assess the validity of certain statements concerning the data's underlying distribution. Frequentist *hypothesis tests*, described in Sec. 38.3.1, provide a rule for accepting or rejecting hypotheses depending on the outcome of a measurement. In *significance tests*, covered in Sec. 38.3.2, one gives the probability to obtain a level of incompatibility with a certain hypothesis that is greater than or equal to the level observed with the actual data. In the Bayesian approach, the corresponding procedure is based fundamentally on the posterior probabilities of the competing hypotheses. In Sec. 38.3.3 we describe a related construct called the Bayes factor, which can be used to quantify the degree to which the data prefer one or another hypothesis.

38.3.1. Hypothesis tests :

A frequentist *test* of a hypothesis (often called the null hypothesis, H_0) is a rule that states for which data values \boldsymbol{x} the hypothesis is rejected. A region of \boldsymbol{x} -space called the critical region, w , is specified such that there is no more than a given probability under H_0 , α , called the *size* or *significance level* of the test, to find $\boldsymbol{x} \in w$. If the data are discrete, it may not be possible to find a critical region with exact probability content α , and thus we require $P(\boldsymbol{x} \in w | H_0) \leq \alpha$. If the data are observed in the critical region, H_0 is rejected.

The critical region is not unique. Choosing one should take into account the probabilities for the data predicted by some alternative hypothesis (or set of alternatives)

H_1 . Rejecting H_0 if it is true is called a *type-I error*, and occurs by construction with probability no greater than α . Not rejecting H_0 if an alternative H_1 is true is called a *type-II error*, and for a given test this will have a certain probability $\beta = P(\mathbf{x} \notin w|H_1)$. The quantity $1 - \beta$ is called the *power* of the test of H_0 with respect to the alternative H_1 . A strategy for defining the critical region can therefore be to maximize the power with respect to some alternative (or alternatives) given a fixed size α .

In high-energy physics, the components of \mathbf{x} might represent the measured properties of candidate events, and the critical region is defined by the cuts that one imposes in order to reject background and thus accept events likely to be of a certain desired type. Here H_0 could represent the background hypothesis and the alternative H_1 could represent the sought after signal. In other cases, H_0 could be the hypothesis that an entire event sample consists of background events only, and the alternative H_1 may represent the hypothesis of a mixture of background and signal.

Often rather than using the full set of quantities \mathbf{x} , it is convenient to define a scalar function of \mathbf{x} called a *test statistic*, $t(\mathbf{x})$. The critical region in \mathbf{x} -space is bounded by a surface of constant $t(\mathbf{x})$. Once the function $t(\mathbf{x})$ is fixed, a given hypothesis for the distribution of \mathbf{x} will determine a distribution for t .

To maximize the power of a test of H_0 with respect to the alternative H_1 , the *Neyman–Pearson lemma* states that the critical region w should be chosen such that for all data values \mathbf{x} inside w , the ratio

$$\lambda(\mathbf{x}) = \frac{f(\mathbf{x}|H_1)}{f(\mathbf{x}|H_0)}, \quad (38.39)$$

is greater than a given constant, the value of which is determined by the size of the test α . Here H_0 and H_1 must be simple hypotheses, i.e., they should not contain undetermined parameters.

The lemma is equivalent to the statement that (38.39) represents the optimal test statistic where the critical region is defined by a single cut on λ . This test will lead to the maximum power (*i.e.*, the maximum probability to reject H_0 if H_1 is true) for a given probability α to reject H_0 if H_0 is in fact true. It can be difficult in practice, however, to determine $\lambda(\mathbf{x})$, since this requires knowledge of the joint p.d.f.s $f(\mathbf{x}|H_0)$ and $f(\mathbf{x}|H_1)$.

In the usual case where the likelihood ratio (38.39) cannot be used explicitly, there exist a variety of other multivariate classifiers that effectively separate different types of events. Methods often used in HEP include *neural networks* or *Fisher discriminants* (see Ref. 10). Recently, further classification methods from machine-learning have been applied in HEP analyses; these include *probability density estimation (PDE)* techniques, *kernel-based PDE (KDE or Parzen window)*, *support vector machines*, and *decision trees*. Techniques such as “boosting” and “bagging” can be applied to combine a number of classifiers into a stronger one with greater stability with respect to fluctuations in the training data. Descriptions of these methods can be found in [11–13], and *Proceedings of the PHYSTAT* conference series [14]. Software for HEP includes the *TMVA* [15] and *StatPatternRecognition* [16] packages.

38.3.2. Tests of significance (goodness-of-fit) :

Often one wants to quantify the level of agreement between the data and a hypothesis without explicit reference to alternative hypotheses. This can be done by defining a statistic t , which is a function of the data whose value reflects in some way the level of agreement between the data and the hypothesis. The analyst must decide what values of the statistic correspond to better or worse levels of agreement with the hypothesis in question; the choice will in general depend on the relevant alternative hypotheses.

The hypothesis in question, H_0 , will determine the p.d.f. $f(t|H_0)$ for the statistic. The significance of a discrepancy between the data and what one expects under the assumption of H_0 is quantified by giving the p -value, defined as the probability to find t in the region of equal or lesser compatibility with H_0 than the level of compatibility observed with the actual data. For example, if t is defined such that large values correspond to poor agreement with the hypothesis, then the p -value would be

$$p = \int_{t_{\text{obs}}}^{\infty} f(t|H_0) dt , \quad (38.40)$$

where t_{obs} is the value of the statistic obtained in the actual experiment.

The p -value should not be confused with the size (significance level) of a test, or the confidence level of a confidence interval (Section 38.4), both of which are pre-specified constants. We may formulate a hypothesis test, however, by defining the critical region to correspond to the data outcomes that give the lowest p -values, so that finding $p \leq \alpha$ implies that the data outcome was in the critical region. When constructing a p -value, one generally chooses the region of data space deemed to have lower compatibility with the model being tested as one having higher compatibility with a given alternative, such that the corresponding test will have a high power with respect to this alternative.

The p -value is a function of the data, and is therefore itself a random variable. If the hypothesis used to compute the p -value is true, then for continuous data p will be uniformly distributed between zero and one. Note that the p -value is not the probability for the hypothesis; in frequentist statistics, this is not defined. Rather, the p -value is the probability, under the assumption of a hypothesis H_0 , of obtaining data at least as incompatible with H_0 as the data actually observed.

When searching for a new phenomenon, one tries to reject the hypothesis H_0 that the data are consistent with known (*e.g.*, Standard Model) processes. If the p -value of H_0 is sufficiently low, then one is willing to accept that some alternative hypothesis is true. Often one converts the p -value into an equivalent significance Z , defined so that a Z standard deviation upward fluctuation of a Gaussian random variable would have an upper tail area equal to p , *i.e.*,

$$Z = \Phi^{-1}(1 - p) . \quad (38.41)$$

Here Φ is the cumulative distribution of the Standard Gaussian, and Φ^{-1} is its inverse (quantile) function. Often in HEP the level of significance where an effect is said to qualify as a discovery is $Z = 5$, *i.e.*, a 5σ effect, corresponding to a p -value of 2.87×10^{-7} . One's actual degree of belief that a new process is present, however, will depend in

general on other factors as well, such as the plausibility of the new signal hypothesis and the degree to which it can describe the data, one's confidence in the model that led to the observed p -value, and possible corrections for multiple observations out of which one focuses on the smallest p -value obtained (the "look-elsewhere effect", discussed in Section 38.3.2.2).

38.3.2.1. Treatment of nuisance parameters for frequentist tests:

Suppose one wants to test hypothetical values of parameters θ , but the model also contains nuisance parameters ν . To find a p -value for θ we can construct a test statistic q_θ such that larger values constitute increasing incompatibility between the data and the hypothesis. Then for an observed value of the statistic $q_{\theta,\text{obs}}$, the p -value of θ is

$$p_\theta(\nu) = \int_{q_{\theta,\text{obs}}}^{\infty} f(q_\theta|\theta, \nu) dq_\theta, \quad (38.42)$$

which depends in general on the nuisance parameters ν . In the strict frequentist approach, θ is rejected only if the p -value is less than α for all possible values of the nuisance parameters.

The difficulty described above is effectively solved if we can define the test statistic q_θ in such a way that its distribution $f(q_\theta|\theta)$ is independent of the nuisance parameters. Although exact independence is only found in special cases, it can be achieved approximately by use of the *profile likelihood ratio*. This is given by the profile likelihood from Eq.(38.18) divided by the value of the likelihood at its maximum, *i.e.*, when evaluated with the ML estimators $\hat{\theta}$ and $\hat{\nu}$:

$$\lambda_p(\theta) = \frac{L(\theta, \hat{\nu}(\theta))}{L(\hat{\theta}, \hat{\nu})}. \quad (38.43)$$

Wilks' theorem states that, providing certain general conditions are satisfied, the distribution of $-2 \ln \lambda_p(\theta)$, under assumption of θ , approaches a χ^2 distribution in the limit where the data sample is very large, independent of the values of the nuisance parameters ν . Here the number of degrees of freedom is equal to the number of components of θ . More details on use of the profile likelihood are given in Refs. [36–37] and in contributions to the PHYSTAT conferences [14]; explicit formulae for special cases can be found in Ref. 38. Further discussion on how to incorporate systematic uncertainties into p -values can be found in Ref. 17.

Even with use of the profile likelihood ratio, for a finite data sample the p -value of hypothesized parameters θ will retain in general some dependence on the nuisance parameters ν . Ideally one would find the the maximum of $p_\theta(\nu)$ from Eq. (38.42) explicitly, but that is often impractical. An approximate and computationally feasible technique is to use $p_\theta(\hat{\nu}(\theta))$, where $\hat{\nu}(\theta)$ are the profiled values of the nuisance parameters as defined in Section 38.2.2.2. The resulting p -value is the correct if the true values of the nuisance parameters are equal to the profiled values used; otherwise it could be either too high or too low. This is discussed further in Section 38.4.2 on confidence intervals.

One may also treat model uncertainties in a Bayesian manner but then use the resulting model in a frequentist test. Suppose the uncertainty in a set of nuisance parameters ν is characterized by a Bayesian prior p.d.f. $\pi(\nu)$. This can be used to construct the marginal (also called the prior predictive) model for the data \mathbf{x} and parameters of interest θ ,

$$P_m(\mathbf{x}|\theta) = \int P(\mathbf{x}|\theta, \nu)\pi(\nu) d\nu . \quad (38.44)$$

The marginal model does not represent the probability of data that would be generated if one were really to repeat the experiment, as in that case one would assume that the nuisance parameters do not vary. Rather, the marginal model represents a situation in which every repetition of the experiment is carried out with new values of ν , randomly sampled from $\pi(\nu)$. It is in effect an average of models each with a given ν , where the average is carried out with respect to the prior p.d.f. $\pi(\nu)$.

The marginal model for the data \mathbf{x} can be used to determine the distribution of a test statistic Q , which can be written

$$P_m(Q|\theta) = \int P(Q|\theta, \nu)\pi(\nu) d\nu . \quad (38.45)$$

In a search for a new signal process, the test statistic can be based on the ratio of likelihoods corresponding to the experiments where signal and background events are both present, L_{s+b} , to that of background only, L_b . Often the likelihoods are evaluated with the profiled values of the nuisance parameters, which may give improved performance. It is important to note, however, that it is through use of the marginal model for the distribution of Q that the uncertainties related to the nuisance parameters are incorporated into the result of the test. Different choices for the test statistic itself only result in variations of the power of the test with respect to different alternatives.

38.3.2.2. *The look-elsewhere effect:*

The “look-elsewhere effect” relates to multiple measurements used to test a single hypothesis. The classic example is when one searches in a distribution for a peak whose position is not predicted in advance. Here the no-peak hypothesis is tested using data in a given range of the distribution. In the frequentist approach the correct p -value of the no-peak hypothesis is the probability, assuming background only, to find a signal as significant as the one found or more so anywhere in the search region. This can be substantially higher than the probability to find a peak of equal or greater significance in the particular place where it appeared. There is in general some ambiguity as to what constitutes the relevant search region or even the broader set of relevant measurements. Although the desired p -value is well defined once the search region has been fixed, an exact treatment can require extensive computation.

The “brute-force” solution to this problem by Monte Carlo involves generating data under the background-only hypothesis and for each data set, fitting a peak of unknown position and recording a measure of its significance. To establish a discovery one often requires a p -value less than 2.9×10^{-7} , corresponding to a 5σ or larger effect. Determining

this with Monte Carlo thus requires generating and fitting a very large number of experiments, perhaps several times 10^7 . In contrast, if the position of the peak is fixed, then the fit to the distribution is much easier, and furthermore one can in many cases use formulae valid for sufficiently large samples that bypass completely the need for Monte Carlo (see, e.g., [38]). But this fixed-position or “local” p -value would not be correct in general, as it assumes the position of the peak was known in advance.

A method that allows one to modify the local p -value computed under assumption of a fixed position to obtain an approximation to the correct “global” value using a relatively simple calculation is described in Ref. 18. Suppose a test statistic q_0 , defined so that larger values indicate increasing disagreement with the data, is observed to have a value u . Furthermore suppose the model contains a nuisance parameter θ (such as the peak position) which is only defined under the signal model (there is no peak in the background-only model). An approximation for the global p -value is found to be

$$p_{\text{global}} \approx p_{\text{local}} + \langle N_u \rangle, \quad (38.46)$$

where $\langle N_u \rangle$ is the mean number of “upcrossings” of the the statistic q_0 above the level u in the range of the nuisance parameter considered (e.g., the mass range).

The value of $\langle N_u \rangle$ can be estimated from the number of upcrossings $\langle N_{u_0} \rangle$ above some much lower value, u_0 , by using a relation due to Davis [19],

$$\langle N_u \rangle \approx \langle N_{u_0} \rangle e^{-(u-u_0)/2}. \quad (38.47)$$

By choosing u_0 sufficiently low, the value of $\langle N_u \rangle$ can be estimated by simulating only a very small number of experiments or even from the observed data, rather than the 10^7 needed if one is dealing with a 5σ effect.

38.3.2.3. Goodness-of-fit with the method of Least Squares:

When estimating parameters using the method of least squares, one obtains the minimum value of the quantity χ^2 (38.19). This statistic can be used to test the *goodness-of-fit*, i.e., the test provides a measure of the significance of a discrepancy between the data and the hypothesized functional form used in the fit. It may also happen that no parameters are estimated from the data, but that one simply wants to compare a histogram, e.g., a vector of Poisson distributed numbers $\mathbf{n} = (n_1, \dots, n_N)$, with a hypothesis for their expectation values $\mu_i = E[n_i]$. As the distribution is Poisson with variances $\sigma_i^2 = \mu_i$, the χ^2 (38.19) becomes *Pearson’s χ^2 statistic*,

$$\chi^2 = \sum_{i=1}^N \frac{(n_i - \mu_i)^2}{\mu_i}. \quad (38.48)$$

If the hypothesis $\boldsymbol{\mu} = (\mu_1, \dots, \mu_N)$ is correct, and if the expected values μ_i in (38.48) are sufficiently large (or equivalently, if the measurements n_i can be treated as following a Gaussian distribution), then the χ^2 statistic will follow the χ^2 p.d.f. with the number of

degrees of freedom equal to the number of measurements N minus the number of fitted parameters.

Alternatively, one may fit parameters and evaluate goodness-of-fit by minimizing $-2 \ln \lambda$ from Eq. (38.16). One finds that the distribution of this statistic approaches the asymptotic limit faster than does Pearson's χ^2 , and thus computing the p -value with the χ^2 p.d.f. will in general be better justified (see Ref. 9 and references therein).

Assuming the goodness-of-fit statistic follows a χ^2 p.d.f., the p -value for the hypothesis is then

$$p = \int_{\chi^2}^{\infty} f(z; n_d) dz, \quad (38.49)$$

where $f(z; n_d)$ is the χ^2 p.d.f. and n_d is the appropriate number of degrees of freedom. Values are shown in Fig. 38.1 or obtained from the ROOT function `TMath::Prob`. If the conditions for using the χ^2 p.d.f. do not hold, the statistic can still be defined as before, but its p.d.f. must be determined by other means in order to obtain the p -value, *e.g.*, using a Monte Carlo calculation.

Since the mean of the χ^2 distribution is equal to n_d , one expects in a “reasonable” experiment to obtain $\chi^2 \approx n_d$. Hence the quantity χ^2/n_d is sometimes reported. Since the p.d.f. of χ^2/n_d depends on n_d , however, one must report n_d as well if one wishes to determine the p -value. The p -values obtained for different values of χ^2/n_d are shown in Fig. 38.2.

If one finds a χ^2 value much greater than n_d , and a correspondingly small p -value, one may be tempted to expect a high degree of uncertainty for any fitted parameters. Poor goodness-of-fit, however, does not mean that one will have large statistical errors for parameter estimates. If, for example, the error bars (or covariance matrix) used in constructing the χ^2 are underestimated, then this will lead to underestimated statistical errors for the fitted parameters. The standard deviations of estimators that one finds from, say, Eq. (38.13) reflect how widely the estimates would be distributed if one were to repeat the measurement many times, assuming that the hypothesis and measurement errors used in the χ^2 are also correct. They do not include the systematic error which may result from an incorrect hypothesis or incorrectly estimated measurement errors in the χ^2 .

38.3.3. Bayes factors :

In Bayesian statistics, all of one's knowledge about a model is contained in its posterior probability, which one obtains using Bayes' theorem (38.30). Thus one could reject a hypothesis H if its posterior probability $P(H|\mathbf{x})$ is sufficiently small. The difficulty here is that $P(H|\mathbf{x})$ is proportional to the prior probability $P(H)$, and there will not be a consensus about the prior probabilities for the existence of new phenomena. Nevertheless one can construct a quantity called the Bayes factor (described below), which can be used to quantify the degree to which the data prefer one hypothesis over another, and is independent of their prior probabilities.

Consider two models (hypotheses), H_i and H_j , described by vectors of parameters $\boldsymbol{\theta}_i$ and $\boldsymbol{\theta}_j$, respectively. Some of the components will be common to both models and others




Figure 38.1: One minus the χ^2 cumulative distribution, $1 - F(\chi^2; n)$, for n degrees of freedom. This gives the p -value for the χ^2 goodness-of-fit test as well as one minus the coverage probability for confidence regions (see Sec. 38.4.2.2).




Figure 38.2: The ‘reduced’ χ^2 , equal to χ^2/n , for n degrees of freedom. The curves show as a function of n the χ^2/n that corresponds to a given p -value.

may be distinct. The full prior probability for each model can be written in the form

$$\pi(H_i, \boldsymbol{\theta}_i) = P(H_i)\pi(\boldsymbol{\theta}_i|H_i) . \quad (38.50)$$

Here $P(H_i)$ is the overall prior probability for H_i , and $\pi(\boldsymbol{\theta}_i|H_i)$ is the normalized p.d.f. of its parameters. For each model, the posterior probability is found using Bayes' theorem,

$$P(H_i|\mathbf{x}) = \frac{\int P(\mathbf{x}|\boldsymbol{\theta}_i, H_i)P(H_i)\pi(\boldsymbol{\theta}_i|H_i) d\boldsymbol{\theta}_i}{P(\mathbf{x})} , \quad (38.51)$$

where the integration is carried out over the internal parameters $\boldsymbol{\theta}_i$ of the model. The ratio of posterior probabilities for the models is therefore

$$\frac{P(H_i|\mathbf{x})}{P(H_j|\mathbf{x})} = \frac{\int P(\mathbf{x}|\boldsymbol{\theta}_i, H_i)\pi(\boldsymbol{\theta}_i|H_i) d\boldsymbol{\theta}_i}{\int P(\mathbf{x}|\boldsymbol{\theta}_j, H_j)\pi(\boldsymbol{\theta}_j|H_j) d\boldsymbol{\theta}_j} \frac{P(H_i)}{P(H_j)} . \quad (38.52)$$

The *Bayes factor* is defined as

$$B_{ij} = \frac{\int P(\mathbf{x}|\boldsymbol{\theta}_i, H_i)\pi(\boldsymbol{\theta}_i|H_i) d\boldsymbol{\theta}_i}{\int P(\mathbf{x}|\boldsymbol{\theta}_j, H_j)\pi(\boldsymbol{\theta}_j|H_j) d\boldsymbol{\theta}_j} . \quad (38.53)$$

This gives what the ratio of posterior probabilities for models i and j would be if the overall prior probabilities for the two models were equal. If the models have no nuisance parameters, *i.e.*, no internal parameters described by priors, then the Bayes factor is simply the likelihood ratio. The Bayes factor therefore shows by how much the probability ratio of model i to model j changes in the light of the data, and thus can be viewed as a numerical measure of evidence supplied by the data in favour of one hypothesis over the other.

Although the Bayes factor is by construction independent of the overall prior probabilities $P(H_i)$ and $P(H_j)$, it does require priors for all internal parameters of a model, *i.e.*, one needs the functions $\pi(\boldsymbol{\theta}_i|H_i)$ and $\pi(\boldsymbol{\theta}_j|H_j)$. In a Bayesian analysis where one is only interested in the posterior p.d.f. of a parameter, it may be acceptable to take an unnormalizable function for the prior (an improper prior) as long as the product of likelihood and prior can be normalized. But improper priors are only defined up to an arbitrary multiplicative constant, and so the Bayes factor would depend on this constant. Furthermore, although the range of a constant normalized prior is unimportant for parameter determination (provided it is wider than the likelihood), this is not so for the Bayes factor when such a prior is used for only one of the hypotheses. So to compute a Bayes factor, all internal parameters must be described by normalized priors that represent meaningful probabilities over the entire range where they are defined.

An exception to this rule may be considered when the identical parameter appears in the models for both numerator and denominator of the Bayes factor. In this case one can argue that the arbitrary constants would cancel. One must exercise some caution, however, as parameters with the same name and physical meaning may still play different roles in the two models.

Both integrals in equation (38.53) are of the form

$$m = \int P(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}) d\boldsymbol{\theta} , \quad (38.54)$$

which is the marginal likelihood seen previously in Eq. (38.44) (in some fields this quantity is called the *evidence*). A review of Bayes factors can be found in Ref. 30. Computing marginal likelihoods can be difficult; in many cases it can be done with the nested sampling algorithm [31] as implemented, *e.g.*, in the program `MultiNest` [32].

38.4. Intervals and limits

When the goal of an experiment is to determine a parameter θ , the result is usually expressed by quoting, in addition to the point estimate, some sort of interval which reflects the statistical precision of the measurement. In the simplest case, this can be given by the parameter's estimated value $\hat{\theta}$ plus or minus an estimate of the standard deviation of θ , $\hat{\sigma}_{\hat{\theta}}$. If, however, the p.d.f. of the estimator is not Gaussian or if there are physical boundaries on the possible values of the parameter, then one usually quotes instead an interval according to one of the procedures described below.

In reporting an interval or limit, the experimenter may wish to

- communicate as objectively as possible the result of the experiment;
- provide an interval that is constructed to cover the true value of the parameter with a specified probability;
- provide the information needed by the consumer of the result to draw conclusions about the parameter or to make a particular decision;
- draw conclusions about the parameter that incorporate stated prior beliefs.

With a sufficiently large data sample, the point estimate and standard deviation (or for the multiparameter case, the parameter estimates and covariance matrix) satisfy essentially all of these goals. For finite data samples, no single method for quoting an interval will achieve all of them.

In addition to the goals listed above, the choice of method may be influenced by practical considerations such as ease of producing an interval from the results of several measurements. Of course the experimenter is not restricted to quoting a single interval or limit; one may choose, for example, first to communicate the result with a confidence interval having certain frequentist properties, and then in addition to draw conclusions about a parameter using a judiciously chosen subjective Bayesian prior.

It is recommended, however, that there be a clear separation between these two aspects of reporting a result. In the remainder of this section, we assess the extent to which various types of intervals achieve the goals stated here.

38.4.1. Bayesian intervals :

As described in Sec. 38.2.4, a Bayesian posterior probability may be used to determine regions that will have a given probability of containing the true value of a parameter. In the single parameter case, for example, an interval (called a Bayesian or credible interval) $[\theta_{\text{lo}}, \theta_{\text{up}}]$ can be determined which contains a given fraction $1 - \alpha$ of the posterior probability, *i.e.*,

$$1 - \alpha = \int_{\theta_{\text{lo}}}^{\theta_{\text{up}}} p(\theta|\mathbf{x}) d\theta . \quad (38.55)$$

Sometimes an upper or lower limit is desired, *i.e.*, θ_{lo} or θ_{up} can be set to a physical boundary or to plus or minus infinity. In other cases, one might be interested in the set of θ values for which $p(\theta|\mathbf{x})$ is higher than for any θ not belonging to the set, which may constitute a single interval or a set of disjoint regions; these are called highest posterior density (HPD) intervals. Note that HPD intervals are not invariant under a nonlinear transformation of the parameter.

If a parameter is constrained to be non-negative, then the prior p.d.f. can simply be set to zero for negative values. An important example is the case of a Poisson variable n , which counts signal events with unknown mean s , as well as background with mean b , assumed known. For the signal mean s , one often uses the prior

$$\pi(s) = \begin{cases} 0 & s < 0 \\ 1 & s \geq 0 \end{cases} . \quad (38.56)$$

This prior is regarded as providing an interval whose frequentist properties can be studied, rather than as representing a degree of belief. For example, to obtain an upper limit on s , one may proceed as follows. The likelihood for s is given by the Poisson distribution for n with mean $s + b$,

$$P(n|s) = \frac{(s + b)^n}{n!} e^{-(s+b)} , \quad (38.57)$$

along with the prior (38.56) in (38.30) gives the posterior density for s . An upper limit s_{up} at confidence level (or here, rather, credibility level) $1 - \alpha$ can be obtained by requiring

$$1 - \alpha = \int_{-\infty}^{s_{\text{up}}} p(s|n) ds = \frac{\int_{-\infty}^{s_{\text{up}}} P(n|s) \pi(s) ds}{\int_{-\infty}^{\infty} P(n|s) \pi(s) ds} , \quad (38.58)$$

where the lower limit of integration is effectively zero because of the cut-off in $\pi(s)$. By relating the integrals in Eq. (38.58) to incomplete gamma functions, the solution for the upper limit is found to be

$$s_{\text{up}} = \frac{1}{2} F_{\chi^2}^{-1} [p, 2(n + 1)] - b , \quad (38.59)$$

where $F_{\chi^2}^{-1}$ is the quantile of the χ^2 distribution (inverse of the cumulative distribution). Here the quantity p is

$$p = 1 - \alpha \left(F_{\chi^2} [2b, 2(n + 1)] \right) , \quad (38.60)$$

where F_{χ^2} is the cumulative χ^2 distribution. For both F_{χ^2} and $F_{\chi^2}^{-1}$ above, the argument $2(n+1)$ gives the number of degrees of freedom. For the special case of $b=0$, the limit reduces to

$$s_{\text{up}} = \frac{1}{2} F_{\chi^2}^{-1}(1 - \alpha; 2(n+1)). \quad (38.61)$$

It happens that for the case of $b=0$, the upper limit from Eq. (38.61) coincides numerically with the frequentist upper limit discussed in Section 38.4.2.3. Values for $1 - \alpha = 0.9$ and 0.95 are given by the values μ_{up} in Table 38.3. The frequentist properties of confidence intervals for the Poisson mean found in this way are discussed in Refs. [2] and [21].

As in any Bayesian analysis, it is important to show how the result changes under assumption of different prior probabilities. For example, one could consider the Jeffreys prior as described in Sec. 38.2.4. For this problem one finds the Jeffreys prior $\pi(s) \propto 1/\sqrt{s+b}$ for $s \geq 0$ and zero otherwise. As with the constant prior, one would not regard this as representing one's prior beliefs about s , both because it is improper and also as it depends on b . Rather it is used with Bayes' theorem to produce an interval whose frequentist properties can be studied.

If the model contains nuisance parameters then these are eliminated by marginalizing, as in Eq. (38.36), to obtain the p.d.f. for the parameters of interest. For example, if the parameter b in the Poisson counting problem above were to be characterized by a prior p.d.f. $\pi(b)$, then one would first use Bayes' theorem to find $p(s, b|n)$. This is then marginalized to find $p(s|n) = \int p(s, b|n)\pi(b) db$, from which one may determine an interval for s . One may not be certain whether to extend a model by including more nuisance parameters. In this case, a Bayes factor may be used to determine to what extent the data prefer a model with additional parameters, as described in Section 38.3.3.

38.4.2. Frequentist confidence intervals :

The unqualified phrase "confidence intervals" refers to frequentist intervals obtained with a procedure due to Neyman [29], described below. These are intervals (or in the multiparameter case, regions) constructed so as to include the true value of the parameter with a probability greater than or equal to a specified level, called the *coverage probability*. It is important to note that in the frequentist approach, such coverage is not meaningful for a fixed interval. A confidence interval, however, depends on the data and thus would fluctuate if one were to repeat the experiment many times. The coverage probability refers to the fraction of intervals in such a set that contain the true parameter value. In this section, we discuss several techniques for producing intervals that have, at least approximately, this property.

38.4.2.1. The Neyman construction for confidence intervals:

Consider a p.d.f. $f(x; \theta)$ where x represents the outcome of the experiment and θ is the unknown parameter for which we want to construct a confidence interval. The variable x could (and often does) represent an estimator for θ . Using $f(x; \theta)$, we can find for a pre-specified probability $1 - \alpha$, and for every value of θ , a set of values $x_1(\theta, \alpha)$ and $x_2(\theta, \alpha)$ such that

$$P(x_1 < x < x_2; \theta) = 1 - \alpha = \int_{x_1}^{x_2} f(x; \theta) dx . \quad (38.62)$$

This is illustrated in Fig. 38.3: a horizontal line segment $[x_1(\theta, \alpha), x_2(\theta, \alpha)]$ is drawn for representative values of θ . The union of such intervals for all values of θ , designated in the figure as $D(\alpha)$, is known as the *confidence belt*. Typically the curves $x_1(\theta, \alpha)$ and $x_2(\theta, \alpha)$ are monotonic functions of θ , which we assume for this discussion.




Figure 38.3: Construction of the confidence belt (see text).

Upon performing an experiment to measure x and obtaining a value x_0 , one draws a vertical line through x_0 . The confidence interval for θ is the set of all values of θ for which the corresponding line segment $[x_1(\theta, \alpha), x_2(\theta, \alpha)]$ is intercepted by this vertical line. Such confidence intervals are said to have a *confidence level* (CL) equal to $1 - \alpha$.

Now suppose that the true value of θ is θ_0 , indicated in the figure. We see from the figure that θ_0 lies between $\theta_1(x)$ and $\theta_2(x)$ if and only if x lies between $x_1(\theta_0)$ and $x_2(\theta_0)$. The two events thus have the same probability, and since this is true for any value θ_0 , we can drop the subscript 0 and obtain

$$1 - \alpha = P(x_1(\theta) < x < x_2(\theta)) = P(\theta_2(x) < \theta < \theta_1(x)) . \quad (38.63)$$

In this probability statement, $\theta_1(x)$ and $\theta_2(x)$, *i.e.*, the endpoints of the interval, are the random variables and θ is an unknown constant. If the experiment were to be repeated a large number of times, the interval $[\theta_1, \theta_2]$ would vary, covering the fixed value θ in a fraction $1 - \alpha$ of the experiments.

The condition of coverage in Eq. (38.62) does not determine x_1 and x_2 uniquely, and additional criteria are needed. One possibility is to choose *central intervals* such that the probabilities excluded below x_1 and above x_2 are each $\alpha/2$. In other cases, one may want to report only an upper or lower limit, in which case the probability excluded below x_1 or above x_2 can be set to zero. Another principle based on *likelihood ratio ordering* for determining which values of x should be included in the confidence belt is discussed below.

When the observed random variable x is continuous, the coverage probability obtained with the Neyman construction is $1 - \alpha$, regardless of the true value of the parameter. If x is discrete, however, it is not possible to find segments $[x_1(\theta, \alpha), x_2(\theta, \alpha)]$ that satisfy Eq. (38.62) exactly for all values of θ . By convention, one constructs the confidence belt requiring the probability $P(x_1 < x < x_2)$ to be *greater than or equal to* $1 - \alpha$. This gives confidence intervals that include the true parameter with a probability greater than or equal to $1 - \alpha$.

An equivalent method of constructing confidence intervals is to consider a test (see Sec. 38.3) of the hypothesis that the parameter's true value is θ (assume one constructs a test for all physical values of θ). One then excludes all values of θ where the hypothesis would be rejected in a test of size α or less. The remaining values constitute the confidence interval at confidence level $1 - \alpha$. If the critical region of the test is characterized by having a p -value $p_\theta \leq \alpha$, then the endpoints of the confidence interval are found in practice by solving $p_\theta = \alpha$ for θ .

In this procedure, one is still free to choose the test to be used; this corresponds to the freedom in the Neyman construction as to which values of the data are included in the confidence belt. One possibility is to use a test statistic based on the *likelihood ratio*,

$$\lambda = \frac{f(x; \theta)}{f(x; \hat{\theta})}, \quad (38.64)$$

where $\hat{\theta}$ is the value of the parameter which, out of all allowed values, maximizes $f(x; \theta)$. This results in the intervals described in Ref. 33 by Feldman and Cousins. The same intervals can be obtained from the Neyman construction described above by including in the confidence belt those values of x which give the greatest values of λ .

If the model contains nuisance parameters ν , then these can be incorporated into the test (or the p -values) used to determine the limit by profiling as discussed in Section 38.3.2.1. As mentioned there, the strict frequentist approach is to regard the parameter of interest θ as excluded only if it is rejected for all possible values of ν . The resulting interval for θ will then cover the true value with a probability greater than or equal to the nominal confidence level for all points in ν -space.

If the p -value is based on the profiled values of the nuisance parameters, *i.e.*, with $\nu = \hat{\nu}(\theta)$ used in Eq. (38.42), then the resulting interval for the parameter of interest

will have the correct coverage if the true values of ν are equal to the profiled values. Otherwise the coverage probability may be too high or too low. This procedure has been called *profile construction* in HEP [20](see also [17]) .

38.4.2.2. Gaussian distributed measurements:

An important example of constructing a confidence interval is when the data consists of a single random variable x that follows a Gaussian distribution; this is often the case when x represents an estimator for a parameter and one has a sufficiently large data sample. If there is more than one parameter being estimated, the multivariate Gaussian is used. For the univariate case with known σ , the probability that the measured value x will fall within $\pm\delta$ of the true value μ is

$$1 - \alpha = \frac{1}{\sqrt{2\pi}\sigma} \int_{\mu-\delta}^{\mu+\delta} e^{-(x-\mu)^2/2\sigma^2} dx = \text{erf} \left(\frac{\delta}{\sqrt{2}\sigma} \right) = 2\Phi \left(\frac{\sigma}{\delta} \right) - 1, \quad (38.65)$$

where erf is the Gaussian error function, which is rewritten in the final equality using Φ , the Gaussian cumulative distribution. Fig. 38.4 shows a $\delta = 1.64\sigma$ confidence interval unshaded. The choice $\delta = \sigma$ gives an interval called the *standard error* which has $1 - \alpha = 68.27\%$ if σ is known. Values of α for other frequently used choices of δ are given in Table 38.1.




Figure 38.4: Illustration of a symmetric 90% confidence interval (unshaded) for a measurement of a single quantity with Gaussian errors. Integrated probabilities, defined by $\alpha = 0.1$, are as shown.

We can set a one-sided (upper or lower) limit by excluding above $x + \delta$ (or below $x - \delta$). The values of α for such limits are half the values in Table 38.1.

The relation (38.65) can be re-expressed using the cumulative distribution function for the χ^2 distribution as

$$\alpha = 1 - F(\chi^2; n), \quad (38.66)$$

Table 38.1: Area of the tails α outside $\pm\delta$ from the mean of a Gaussian distribution.

α	δ	α	δ
0.3173	1σ	0.2	1.28σ
4.55×10^{-2}	2σ	0.1	1.64σ
2.7×10^{-3}	3σ	0.05	1.96σ
6.3×10^{-5}	4σ	0.01	2.58σ
5.7×10^{-7}	5σ	0.001	3.29σ
2.0×10^{-9}	6σ	10^{-4}	3.89σ

for $\chi^2 = (\delta/\sigma)^2$ and $n = 1$ degree of freedom. This can be seen as the $n = 1$ curve in Fig. 38.1 or obtained by using the ROOT function `TMath::Prob`.

For multivariate measurements of, say, n parameter estimates $\hat{\boldsymbol{\theta}} = (\hat{\theta}_1, \dots, \hat{\theta}_n)$, one requires the full covariance matrix $V_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]$, which can be estimated as described in Sections 38.2.2 and 38.2.3. Under fairly general conditions with the methods of maximum-likelihood or least-squares in the large sample limit, the estimators will be distributed according to a multivariate Gaussian centered about the true (unknown) values $\boldsymbol{\theta}$, and furthermore, the likelihood function itself takes on a Gaussian shape.

The standard error ellipse for the pair $(\hat{\theta}_i, \hat{\theta}_j)$ is shown in Fig. 38.5, corresponding to a contour $\chi^2 = \chi_{\min}^2 + 1$ or $\ln L = \ln L_{\max} - 1/2$. The ellipse is centered about the estimated values $\hat{\boldsymbol{\theta}}$, and the tangents to the ellipse give the standard deviations of the estimators, σ_i and σ_j . The angle of the major axis of the ellipse is given by

$$\tan 2\phi = \frac{2\rho_{ij}\sigma_i\sigma_j}{\sigma_j^2 - \sigma_i^2}, \quad (38.67)$$

where $\rho_{ij} = \text{cov}[\hat{\theta}_i, \hat{\theta}_j]/\sigma_i\sigma_j$ is the correlation coefficient.

The correlation coefficient can be visualized as the fraction of the distance σ_i from the ellipse's horizontal center-line at which the ellipse becomes tangent to vertical, *i.e.*, at the distance $\rho_{ij}\sigma_i$ below the center-line as shown. As ρ_{ij} goes to +1 or -1, the ellipse thins to a diagonal line.

It could happen that one of the parameters, say, θ_j , is known from previous measurements to a precision much better than σ_j , so that the current measurement contributes almost nothing to the knowledge of θ_j . However, the current measurement of θ_i and its dependence on θ_j may still be important. In this case, instead of quoting both parameter estimates and their correlation, one sometimes reports the value of θ_i , which minimizes χ^2 at a fixed value of θ_j , such as the PDG best value. This θ_i value lies along the dotted line between the points where the ellipse becomes tangent to vertical, and has statistical error σ_{inner} as shown on the figure, where $\sigma_{\text{inner}} = (1 - \rho_{ij}^2)^{1/2}\sigma_i$. Instead of the correlation ρ_{ij} , one reports the dependency $d\hat{\theta}_i/d\theta_j$ which is the slope of the dotted line. This slope is related to the correlation coefficient by $d\hat{\theta}_i/d\theta_j = \rho_{ij} \times \frac{\sigma_i}{\sigma_j}$.




Figure 38.5: Standard error ellipse for the estimators $\hat{\theta}_i$ and $\hat{\theta}_j$. In this case the correlation is negative.

Table 38.2: Values of $\Delta\chi^2$ or $2\Delta\ln L$ corresponding to a coverage probability $1 - \alpha$ in the large data sample limit, for joint estimation of m parameters.

$(1 - \alpha)$ (%)	$m = 1$	$m = 2$	$m = 3$
68.27	1.00	2.30	3.53
90.	2.71	4.61	6.25
95.	3.84	5.99	7.82
95.45	4.00	6.18	8.03
99.	6.63	9.21	11.34
99.73	9.00	11.83	14.16

As in the single-variable case, because of the symmetry of the Gaussian function between $\boldsymbol{\theta}$ and $\hat{\boldsymbol{\theta}}$, one finds that contours of constant $\ln L$ or χ^2 cover the true values with a certain, fixed probability. That is, the confidence region is determined by

$$\ln L(\boldsymbol{\theta}) \geq \ln L_{\max} - \Delta \ln L, \quad (38.68)$$

or where a χ^2 has been defined for use with the method of least-squares,

$$\chi^2(\boldsymbol{\theta}) \leq \chi_{\min}^2 + \Delta\chi^2. \quad (38.69)$$

Values of $\Delta\chi^2$ or $2\Delta\ln L$ are given in Table 38.2 for several values of the coverage probability and number of fitted parameters.

For non-Gaussian data samples, the probability for the regions determined by equations (38.68) or (38.69) to cover the true value of $\boldsymbol{\theta}$ becomes independent of $\boldsymbol{\theta}$ only in the large-sample limit. So for a finite data sample these are not exact confidence regions according to our previous definition. Nevertheless, they can still have a coverage probability only weakly dependent on the true parameter, and approximately as given in Table 38.2. In any case, the coverage probability of the intervals or regions obtained

according to this procedure can in principle be determined as a function of the true parameter(s), for example, using a Monte Carlo calculation.

One of the practical advantages of intervals that can be constructed from the log-likelihood function or χ^2 is that it is relatively simple to produce the interval for the combination of several experiments. If N independent measurements result in log-likelihood functions $\ln L_i(\boldsymbol{\theta})$, then the combined log-likelihood function is simply the sum,

$$\ln L(\boldsymbol{\theta}) = \sum_{i=1}^N \ln L_i(\boldsymbol{\theta}) . \quad (38.70)$$

This can then be used to determine an approximate confidence interval or region with Eq. (38.68), just as with a single experiment.

38.4.2.3. Poisson or binomial data:

Another important class of measurements consists of counting a certain number of events, n . In this section, we will assume these are all events of the desired type, *i.e.*, there is no background. If n represents the number of events produced in a reaction with cross section σ , say, in a fixed integrated luminosity \mathcal{L} , then it follows a Poisson distribution with mean $\mu = \sigma\mathcal{L}$. If, on the other hand, one has selected a larger sample of N events and found n of them to have a particular property, then n follows a binomial distribution where the parameter p gives the probability for the event to possess the property in question. This is appropriate, *e.g.*, for estimates of branching ratios or selection efficiencies based on a given total number of events.

For the case of Poisson distributed n , the upper and lower limits on the mean value μ can be found from the Neyman procedure to be

$$\mu_{\text{lo}} = \frac{1}{2} F_{\chi^2}^{-1}(\alpha_{\text{lo}}; 2n) , \quad (38.71a)$$

$$\mu_{\text{up}} = \frac{1}{2} F_{\chi^2}^{-1}(1 - \alpha_{\text{up}}; 2(n + 1)) , \quad (38.71b)$$

where the upper and lower limits are at confidence levels of $1 - \alpha_{\text{lo}}$ and $1 - \alpha_{\text{up}}$, respectively, and $F_{\chi^2}^{-1}$ is the quantile of the χ^2 distribution (inverse of the cumulative distribution). The quantiles $F_{\chi^2}^{-1}$ can be obtained from standard tables or from the ROOT routine `TMath::ChisquareQuantile`. For central confidence intervals at confidence level $1 - \alpha$, set $\alpha_{\text{lo}} = \alpha_{\text{up}} = \alpha/2$.

It happens that the upper limit from Eq. (38.71b) coincides numerically with the Bayesian upper limit for a Poisson parameter, using a uniform prior p.d.f. for μ . Values for confidence levels of 90% and 95% are shown in Table 38.3. For the case of binomially distributed n successes out of N trials with probability of success p , the upper and lower limits on p are found to be

$$p_{\text{lo}} = \frac{n F_F^{-1}[\alpha_{\text{lo}}; 2n, 2(N - n + 1)]}{N - n + 1 + n F_F^{-1}[\alpha_{\text{lo}}; 2n, 2(N - n + 1)]} , \quad (38.72a)$$

$$p_{\text{up}} = \frac{(n + 1) F_F^{-1}[1 - \alpha_{\text{up}}; 2(n + 1), 2(N - n)]}{(N - n) + (n + 1) F_F^{-1}[1 - \alpha_{\text{up}}; 2(n + 1), 2(N - n)]} . \quad (38.72b)$$

Table 38.3: Lower and upper (one-sided) limits for the mean μ of a Poisson variable given n observed events in the absence of background, for confidence levels of 90% and 95%.

n	$1 - \alpha = 90\%$		$1 - \alpha = 95\%$	
	μ_{lo}	μ_{up}	μ_{lo}	μ_{up}
0	–	2.30	–	3.00
1	0.105	3.89	0.051	4.74
2	0.532	5.32	0.355	6.30
3	1.10	6.68	0.818	7.75
4	1.74	7.99	1.37	9.15
5	2.43	9.27	1.97	10.51
6	3.15	10.53	2.61	11.84
7	3.89	11.77	3.29	13.15
8	4.66	12.99	3.98	14.43
9	5.43	14.21	4.70	15.71
10	6.22	15.41	5.43	16.96

Here F_F^{-1} is the quantile of the F distribution (also called the Fisher–Snedecor distribution; see Ref. 4).

38.4.2.4. Parameter exclusion in cases of low sensitivity:

An important example of a statistical test arises in the search for a new signal process. Suppose the parameter μ is defined such that it is proportional to the signal cross section. A statistical test may be carried out for hypothesized values of μ , which may be done by computing a p -value, p_μ , for all μ . Those values not rejected in a test of size α , *i.e.*, for which one does not find $p_\mu \leq \alpha$, constitute a confidence interval with confidence level $1 - \alpha$.

In general one will find that for some regions in the parameter space of the signal model, the predictions for data are almost indistinguishable from those of the background-only model. This corresponds to the case where μ is very small, as would occur, *e.g.*, if one searches for a new particle with a mass so high that its production rate in a given experiment is negligible. That is, one has essentially no experimental sensitivity to such a model.

One would prefer that if the sensitivity to a model (or a point in a model’s parameter space) is very low, then it should not be excluded. Even if the outcomes predicted with or without signal are identical, however, the probability to reject the signal model will equal α , the type-I error rate. As one often takes α to be 5%, this would mean that in a large number of searches covering a broad range of a signal model’s parameter space,

there would inevitably be excluded regions in which the experimental sensitivity is very small, and thus one may question whether it is justified to regard such parameter values as disfavored.

Exclusion of models to which one has little or no sensitivity occurs, for example, if the data fluctuate very low relative to the expectation of the background-only hypothesis. In this case the resulting upper limit on the predicted rate (cross section) of a signal model may be anomalously low. As a means of controlling this effect one often determines the mean or median limit under assumption of the background-only hypothesis using a simplified Monte Carlo simulation of the experiment. An upper limit found significantly below the background-only expectation may indicate a strong downward fluctuation of the data, or perhaps as well an incorrect estimate of the background rate.

One way to mitigate the problem of excluding models to which one is not sensitive is the CL_s method, where the measure used to test a parameter is increased for decreasing sensitivity [34,35]. The procedure is based on a statistic called CL_s , which is defined as

$$\text{CL}_s = \frac{p_\mu}{1 - p_b}, \quad (38.73)$$

where p_b is the p -value of the background-only hypothesis. In the usual formulation of the method, both p_μ and p_b are defined using a single test statistic, and the definition of CL_s above assumes this statistic is continuous; more details can be found in Refs. [34,35].

A point in a model's parameter space is regarded as excluded if one finds $\text{CL}_s \leq \alpha$. As the denominator in Eq. (38.73) is always less than or equal to unity, the exclusion criterion based on CL_s is more stringent than the usual requirement $p_\mu \leq \alpha$. In this sense the CL_s procedure is conservative, and the coverage probability of the corresponding intervals will exceed the nominal confidence level $1 - \alpha$. If the experimental sensitivity to a given value of μ is very low, then one finds that as p_μ decreases, so does the denominator $1 - p_b$, and thus the condition $\text{CL}_s \leq \alpha$ is effectively prevented from being satisfied. In this way the exclusion of parameters in the case of low sensitivity is suppressed.

The CL_s procedure has the attractive feature that the resulting intervals coincide with those obtained from the Bayesian method in two important cases: the mean value of a Poisson or Gaussian distributed measurement with a constant prior. The CL_s intervals overcover for all values of the parameter μ , however, by an amount that depends on μ .

The problem of excluding parameter values to which one has little sensitivity is particularly acute when one wants to set a one-sided limit, *e.g.*, an upper limit on a cross section. Here one tests a value of a rate parameter μ against the alternative of a lower rate, and therefore the critical region of the test is taken to correspond to data outcomes with a low event yield. If the number of events found in the search region fluctuates low enough, however, it can happen that all physically meaningful signal parameter values, including those to which one has very little sensitivity, are rejected by the test.

Another solution to this problem, therefore, is to replace the one-sided test by one based on the likelihood ratio, where the critical region is not restricted to low rates. This is the approach followed in the Feldman-Cousins procedure described in Section 38.4.2.1. The critical region for the test of a given value of μ contains data values characteristic

of both higher and lower rates. As a result, for a given observed rate one can in general obtain a two-sided interval. If, however, the parameter estimate $\hat{\mu}$ is sufficiently close to the lower limit of zero, then only high values of μ are rejected, and the lower edge of the confidence interval is at zero. Note, however, that the coverage property of $1 - \alpha$ pertains to the entire interval, not to the probability for the upper edge μ_{up} to be greater than the true value μ . For parameter estimates increasingly far away from the boundary, *i.e.*, for increasing signal significance, the point $\mu = 0$ is excluded and the interval has nonzero upper and lower edges.

An additional difficulty arises when a parameter estimate is not significantly far away from the boundary, in which case it is natural to report a one-sided confidence interval (often an upper limit). It is straightforward to force the Neyman prescription to produce only an upper limit by setting $x_2 = \infty$ in Eq. (38.62). Then x_1 is uniquely determined and the upper limit can be obtained. If, however, the data come out such that the parameter estimate is not so close to the boundary, one might wish to report a central confidence interval (*i.e.*, an interval based on a two-sided test with equal upper and lower tail areas). As pointed out by Feldman and Cousins [33], however, if the decision to report an upper limit or two-sided interval is made by looking at the data (“flip-flopping”), then in general there will be parameter values for which the resulting intervals have a coverage probability less than $1 - \alpha$. With the confidence intervals suggested in [33], the prescription determines whether the interval is one- or two-sided in a way which preserves the coverage probability (and are thus said to be *unified*).

The intervals according to this method for the mean of Poisson variable in the absence of background are given in Table 38.4. (Note that α in Ref. 33 is defined following Neyman [29] as the coverage probability; this is opposite the modern convention used here in which the coverage probability is $1 - \alpha$.) The values of $1 - \alpha$ given here refer to the coverage of the true parameter by the whole interval $[\mu_1, \mu_2]$. In Table 38.3 for the one-sided upper limit, however, $1 - \alpha$ refers to the probability to have $\mu_{\text{up}} \geq \mu$ (or $\mu_{\text{lo}} \leq \mu$ for lower limits).

A potential difficulty with unified intervals arises if, for example, one constructs such an interval for a Poisson parameter s of some yet to be discovered signal process with, say, $1 - \alpha = 0.9$. If the true signal parameter is zero, or in any case much less than the expected background, one will usually obtain a one-sided upper limit on s . In a certain fraction of the experiments, however, a two-sided interval for s will result. Since, however, one typically chooses $1 - \alpha$ to be only 0.9 or 0.95 when setting limits, the value $s = 0$ may be found below the lower edge of the interval before the existence of the effect is well established. It must then be communicated carefully that in excluding $s = 0$ at, say, 90% or 95% confidence level from the interval, one is not necessarily claiming to have discovered the effect, for which one would usually require a higher level of significance (*e.g.*, 5σ).

Another possibility is to construct a Bayesian interval as described in Section 38.4.1. The presence of the boundary can be incorporated simply by setting the prior density to zero in the unphysical region. More specifically, the prior may be chosen using formal rules such as the reference prior or Jeffreys prior mentioned in Sec. 38.2.4.

In HEP a widely used prior for the mean μ of a Poisson distributed measurement

Table 38.4: Unified confidence intervals $[\mu_1, \mu_2]$ for a the mean of a Poisson variable given n observed events in the absence of background, for confidence levels of 90% and 95%.

n	$1 - \alpha = 90\%$		$1 - \alpha = 95\%$	
	μ_1	μ_2	μ_1	μ_2
0	0.00	2.44	0.00	3.09
1	0.11	4.36	0.05	5.14
2	0.53	5.91	0.36	6.72
3	1.10	7.42	0.82	8.25
4	1.47	8.60	1.37	9.76
5	1.84	9.99	1.84	11.26
6	2.21	11.47	2.21	12.75
7	3.56	12.53	2.58	13.81
8	3.96	13.99	2.94	15.29
9	4.36	15.30	4.36	16.77
10	5.50	16.50	4.75	17.82

has been uniform for $\mu \geq 0$. This prior does not follow from any fundamental rule nor can it be regarded as reflecting a reasonable degree of belief, since the prior probability for μ to lie between any two finite limits is zero. It is more appropriately regarded as a procedure for obtaining intervals with frequentist properties that can be investigated. The resulting upper limits have a coverage probability that depends on the true value of the Poisson parameter, and is nowhere smaller than the stated probability content. Lower limits and two-sided intervals for the Poisson mean based on flat priors undercover, however, for some values of the parameter, although to an extent that in practical cases may not be too severe [2,21]. Intervals constructed in this way have the advantage of being easy to derive; if several independent measurements are to be combined then one simply multiplies the likelihood functions (cf. Eq. (38.70)).

In any case, it is important to always report sufficient information so that the result can be combined with other measurements. Often this means giving an unbiased estimator and its standard deviation, even if the estimated value is in the unphysical region.

It can also be useful with a frequentist interval to calculate its subjective probability content using the posterior p.d.f. based on one or several reasonable guesses for the prior p.d.f. If it turns out to be significantly less than the stated confidence level, this warns that it would be particularly misleading to draw conclusions about the parameter's value from the interval alone.

References:

1. B. Efron, *Am. Stat.* **40**, 11 (1986).
2. R.D. Cousins, *Am. J. Phys.* **63**, 398 (1995).
3. A. Stuart, J.K. Ord, and S. Arnold, *Kendall's Advanced Theory of Statistics*, Vol. 2A: *Classical Inference and the Linear Model*, 6th ed., Oxford Univ. Press (1999), and earlier editions by Kendall and Stuart. The likelihood-ratio ordering principle is described at the beginning of Ch. 23. Chapter 26 compares different schools of statistical inference.
4. F.E. James, *Statistical Methods in Experimental Physics*, 2nd ed., (World Scientific, Singapore, 2007).
5. H. Cramér, *Mathematical Methods of Statistics*, Princeton Univ. Press, New Jersey (1958).
6. L. Lyons, *Statistics for Nuclear and Particle Physicists*, (Cambridge University Press, New York, 1986).
7. R. Barlow, *Nucl. Instrum. Methods* **A297**, 496 (1990).
8. G. Cowan, *Statistical Data Analysis*, (Oxford University Press, Oxford, 1998).
9. For a review, see S. Baker and R. Cousins, *Nucl. Instrum. Methods* **221**, 437 (1984).
10. For information on neural networks and related topics, see *e.g.*, C.M. Bishop, *Neural Networks for Pattern Recognition*, Clarendon Press, Oxford (1995); C. Peterson and T. Rönvaldsson, An Introduction to Artificial Neural Networks, in *Proceedings of the 1991 CERN School of Computing*, C. Verkerk (ed.), CERN 92-02 (1992).
11. T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning* (2nd edition, Springer, New York, 2009).
12. A. Webb, *Statistical Pattern Recognition*, 2nd ed., (Wiley, New York, 2002).
13. L.I. Kuncheva, *Combining Pattern Classifiers*, (Wiley, New York, 2004).
14. Links to the *Proceedings of the PHYSTAT* conference series (Durham 2002, Stanford 2003, Oxford 2005, and Geneva 2007, 2011) can be found at phystat.org.
15. A. Höcker *et al.*, *TMVA Users Guide*, physics/0703039(2007); software available from tmva.sf.net.
16. I. Narsky, *StatPatternRecognition: A C++ Package for Statistical Analysis of High Energy Physics Data*, physics/0507143(2005); software avail. from sourceforge.net/projects/statpatrec.
17. L. Demortier, *P-Values and Nuisance Parameters, Proceedings of PHYSTAT 2007*, CERN-2008-001, p. 23.
18. E. Gross and O. Vitells, *Eur. Phys. J.* **C70**, 525 (2010); [arXiv:1005.1891](https://arxiv.org/abs/1005.1891).
19. R.B. Davis, *Biometrika* **74**, 33 (1987).
20. K. Cranmer, *Statistical Challenges for Searches for New Physics at the LHC*, in *Proceedings of PHYSTATS 2005*, L. Lyons and m. Karagoz Unel (eds.), Oxford (2005); [arXiv:physics/0511028](https://arxiv.org/abs/physics/0511028).
21. B.P. Roe and M.B. Woodroffe, *Phys. Rev.* **D63**, 13009 (2000).
22. A. O'Hagan and J.J. Forster, *Bayesian Inference*, (2nd edition, volume 2B of *Kendall's Advanced Theory of Statistics*, Arnold, London, 2004).
23. D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial*, (Oxford University Press, 2006).

24. P.C. Gregory, *Bayesian Logical Data Analysis for the Physical Sciences*, (Cambridge University Press, 2005).
25. J.M. Bernardo and A.F.M. Smith, *Bayesian Theory*, (Wiley, 2000).
26. Robert E. Kass and Larry Wasserman, *J. Am. Stat. Assoc.* **91**, 1343 (1996).
27. J.M. Bernardo, *J. R. Statist. Soc.* **B41**, 113 (1979); J.M. Bernardo and J.O. Berger, *J. Am. Stat. Assoc.* **84**, 200 (1989). See also J.M. Bernardo, *Reference Analysis*, in *Handbook of Statistics*, 25 (D.K. Dey and C.R. Rao, eds.), 17-90, Elsevier (2005) and references therein.
28. L. Demortier, S. Jain, and H. Prosper, *Phys. Rev. D* **82**, 034002 (2010); [arXiv:1002.1111](#).
29. J. Neyman, *Phil. Trans. Royal Soc. London, Series A*, **236**, 333 (1937), reprinted in *A Selection of Early Statistical Papers on J. Neyman*, (University of California Press, Berkeley, 1967).
30. R. E. Kass and A. E. Raftery, *J. Am. Stat. Assoc.* **90**, 773 (1995).
31. J. Skilling, *Nested Sampling, AIP Conference Proceedings***735**, 395–405 (2004).
32. F. Feroz, M.P. Hobson, and M. Bridges, *Mon. Not. Roy. Astron. Soc.* **398**, 1601-1614 (2009); [arXiv:0809.3437](#).
33. G.J. Feldman and R.D. Cousins, *Phys. Rev. D***57**, 3873 (1998). This paper does not specify what to do if the ordering principle gives equal rank to some values of x . Eq. 21.6 of Ref. 3 gives the rule: all such points are included in the acceptance region (the domain $D(\alpha)$). Some authors have assumed the contrary, and shown that one can then obtain null intervals.
34. A.L. Read, *Modified frequentist analysis of search results (the CL_s method)*, in F. James, L. Lyons, and Y. Perrin (eds.), *Workshop on Confidence Limits*, CERN Yellow Report 2000-005, available through [cdsweb.cern.ch](#).
35. T. Junk, *Nucl. Instrum. Methods* **A434**, 435 (1999).
36. N. Reid, *Likelihood Inference in the Presence of Nuisance Parameters, Proceedings of PHYSTAT2003*, L. Lyons, R. Mount, and R. Reitmeyer, eds., eConf C030908, Stanford, 2003.
37. W.A. Rolke, A.M. Lopez, and J. Conrad, *Nucl. Instrum. Methods* **A551**, 493 (2005); [physics/0403059](#).
38. G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Eur. Phys. J.* **C71**, 1554 (2011).