

Reference priors for high energy physics

Luc Demortier

Laboratory of Experimental High Energy Physics, The Rockefeller University, New York, New York 10065, USA

Supriya Jain*

Homer L. Dodge Department of Physics and Astronomy, University of Oklahoma, Norman, Oklahoma 73019, USA

Harrison B. Prosper

Department of Physics, Florida State University, Tallahassee, Florida 32306, USA

(Received 4 February 2010; published 3 August 2010)

Bayesian inferences in high energy physics often use uniform prior distributions for parameters about which little or no information is available before data are collected. The resulting posterior distributions are therefore sensitive to the choice of parametrization for the problem and may even be improper if this choice is not carefully considered. Here we describe an extensively tested methodology, known as reference analysis, which allows one to construct parametrization-invariant priors that embody the notion of minimal informativeness in a mathematically well-defined sense. We apply this methodology to general cross section measurements and show that it yields sensible results. A recent measurement of the single-top quark cross section illustrates the relevant techniques in a realistic situation.

DOI: [10.1103/PhysRevD.82.034002](https://doi.org/10.1103/PhysRevD.82.034002)

PACS numbers: 06.20.Dk, 14.65.Ha

I. INTRODUCTION

The Bayesian approach [1] to inference plays an increasingly important role in particle physics research. This is due, in part, to a better understanding of Bayesian reasoning within the field and the concomitant abating of the frequentist/Bayesian debate. Moreover, the small but growing number of successful applications provide concrete examples of how the Bayesian approach fares in practice.

In spite of these successes the specification of priors in a principled way remains a conceptual and practical hurdle. In the so-called subjective Bayesian approach [2], one is invited to elicit the prior based on one's actual beliefs about the unknown parameters in the problem. If one has well-understood information, based, for example, on subsidiary measurements or simulation studies, one can encode this partial information in an evidence-based prior [3]. Such priors generally occasion little or no controversy. On the other hand, if one knows little about a given parameter, or if one prefers to act as if one knows little, then it is far from clear how one ought to encode this minimal information in a prior probability.

Since there is, in fact, no unique way to model prior ignorance, a viewpoint has evolved in which this lack of knowledge is represented by one's willingness to *adopt* a standard prior for certain parameters [4], just as one has adopted a standard for quantities such as length and weight. In this spirit, our field adopted as a convention a uniform (flat) prior for unknown cross sections and other

parameters (see, for example, Ref. [5]), mainly because this prescription is simple to implement and seems to embody Laplace's principle of insufficient reason. Unfortunately, uniform priors are both conceptually and practically flawed. The conceptual difficulty is with their justification: lack of knowledge about a parameter σ implies lack of knowledge about any one-to-one transform σ' of σ , and yet a prior distribution that is uniform in σ will not be so in σ' if the transform is nonlinear. The practical problem is that careless use of uniform priors can lead to improper posteriors, that is, posteriors whose integrals are infinite and which can therefore not be used to assign meaningful probabilities to subsets of parameter space. An example of this pathology is found in a common method for reporting the exclusion of a new physics signal, where one estimates an upper limit from a posterior distribution for the signal's production cross section. When constructed from a Poisson probability mass function for the observations, a flat prior for the signal cross section, and a truncated Gaussian prior for the signal acceptance, this posterior is actually improper. However, for small acceptance uncertainties the divergence of the upper limit is often concealed by the inevitable truncation of numerical computations [6].

The specification of priors that encode minimal information is of such importance in practice that a large body of literature exists describing attempts to construct priors that yield results with provably useful characteristics. These priors are typically arrived at using formal rules. In this paper, therefore, we refer to them as *formal priors* [7] to distinguish them from evidence-based priors. Many such formal rules exist [4]. In this paper we study, and then recommend, a rule which is arguably the most successful:

*Present address: State University of New York, Buffalo, NY 14260, USA.

that developed by Bernardo [8] and Berger and Bernardo [9–11]. Formal priors constructed according to the Bernardo-Berger rule are called *reference priors*, a somewhat unfortunate name given that the term reference prior is sometimes used as a synonym for what we have called a formal prior.

Reference priors have been shown to yield results with several desirable properties, all of which should appeal to particle physicists. Therefore, in principle such priors could be a foundation for Bayesian inference in particle physics research. However, reference priors and the associated methods collectively referred to as reference analysis [12,13] have yet to enter the field in a significant way. The purpose of this paper is to initiate this process by applying the Bernardo-Berger method to a familiar but important class of problems, namely, that of calculating posterior densities for signal cross sections.

In the next section we describe the general goals of reference prior construction and show how these are implemented via the concept of missing information. For simplicity we limit that discussion to one-parameter problems. Section III then considers the treatment of nuisance parameters about which prior information is available. Examples of such parameters include detector calibration constants, background contaminations, geometrical acceptances, and integrated luminosities. We describe two methods for handling these parameters, depending on the type of information that is available about them. These methods are then applied to counting experiments with uncertain background contamination and effective luminosity. In the simplest cases we have obtained analytical expressions for the marginal posterior for the quantity of interest. For the general case we have developed a numerical algorithm. Some appealing properties of these posteriors are examined in Sec. IV. In Sec. V the reference prior methodology is applied to a recent measurement of the production cross section for single-top quarks at the Tevatron. Final comments are presented in Sec. VI.

II. REFERENCE PRIORS

In 1979, Bernardo [8] introduced a formal rule for constructing what he called a reference prior. The goal was to construct a prior which, in a sense to be made precise, contained as little information as possible relative to the statistical model under consideration. By statistical model he meant a representation of the entire experimental design, including the probability distribution of the data, the sampling space, and the stopping rule. Hence, by construction reference priors depend on all these aspects of a statistical model, and so will inferences derived from data with the help of a reference prior. This may seem to violate the so-called *likelihood principle* [14], according to which all the information about unknown model parameters obtainable from an experiment is contained in the likelihood function, i.e. the probability distribution of the

data, evaluated at the observations and viewed as a function of the parameters. While this is formally true, it should be kept in mind that the likelihood principle applies after data have been observed, whereas reference priors are constructed at the experimental design stage. Their purpose is to approximate a consensus of opinions that is suitable for scientific communication. This is generally unproblematic in large-sample situations, where posterior inferences are dominated by the likelihood function. In small sample cases, however, results obtained with reference priors should be considered preliminary, and a careful study should be conducted of the degree to which inferences about the physics model underlying the observations can be trusted. This can be achieved by examining the sensitivity of the results to changes in the prior and subsequently assessing the need for additional observations.

Reference priors have several desirable properties, including

- (1) *generality*: a well-defined algorithm exists to create a reference prior for almost any type of estimation problem, and the resulting posterior is proper;
- (2) *invariance*: given a one-to-one map from a parameter θ to a parameter ϕ , applying the reference prior construction separately to θ and ϕ yields posteriors that are related by the correct transformation law, $\pi(\phi | x) = \pi(\theta | x)|\partial\theta/\partial\phi|$;
- (3) *sampling consistency*: the posterior densities from an ensemble of experiments tend to cluster around the true values of the parameters; and
- (4) *coherence*: inferences derived from reference priors avoid marginalization paradoxes.

Marginalization paradoxes [15] arise in multiparameter problems when a posterior density can be calculated in different ways that ought to give the same answer but do not (see Fig. 1).

This incoherence does not happen with subjective or evidence-based priors because these priors are always proper. With formal priors, however, it can only be avoided by allowing the joint prior for all the parameters in a given statistical model to depend on the quantity of interest. This is in fact what the reference prior construction does. For a simple illustration, consider n measurements x_i from the normal model with unknown mean μ and standard deviation σ . The likelihood function is

$$p(\vec{x} | \mu, \sigma) = \prod_{i=1}^n \frac{e^{-(1/2)((x_i - \mu)/\sigma)^2}}{\sqrt{2\pi}\sigma} = \frac{e^{-((n-1)/2)(s/\sigma)^2 - (n/2)((\bar{x} - \mu)/\sigma)^2}}{(\sqrt{2\pi}\sigma)^n}, \quad (1)$$

where $\bar{x} = \sum_i^n x_i/n$ and $s^2 = \sum_i^n (x_i - \bar{x})^2/(n-1)$. When μ is the quantity of interest, the reference prior derived from this likelihood is $1/\sigma$. Restricting the remaining calculations to the case $n=2$ for convenience, the joint

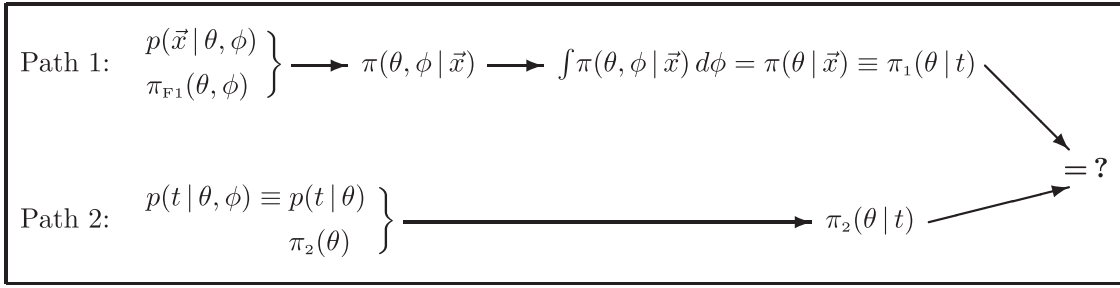


FIG. 1. Let \vec{x} be a data set modeled by the probability density $p(\vec{x} | \theta, \phi)$, where θ and ϕ are unknown parameters, and consider the following two paths to a posterior density for θ . In path 1, we use a formal prior $\pi_{F1}(\theta, \phi)$ to construct the joint posterior for θ and ϕ and then integrate out ϕ . Suppose that the result of this operation only depends on the data \vec{x} through the statistic t ; this gives us $\pi_1(\theta | t)$. For path 2, assume further that the sampling distribution of t only depends on θ . We can then directly construct a posterior for θ , say, $\pi_2(\theta | t)$. A marginalization paradox occurs if $\pi_1(\theta | t) \neq \pi_2(\theta | t)$ regardless of the choice of prior $\pi_2(\theta)$ in path 2.

reference posterior for μ and σ is then

$$\pi(\mu, \sigma | \vec{x}) = \frac{\sqrt{2}s}{\pi\sigma^3} e^{-(1/2)(s/\sigma)^2 - ((\mu - \bar{x})/\sigma)^2}. \quad (2)$$

Integrating out σ yields the marginal μ posterior, which is a Cauchy distribution with location parameter \bar{x} and scale parameter $s/\sqrt{2}$. Suppose, however, that our interest lies in the standardized mean $\theta = \mu/\sigma$. In a nonreference approach one would perform the transformation $(\mu, \sigma) \rightarrow (\theta, \sigma)$ in Eq. (2) and integrate out σ in order to obtain the marginal θ posterior. The latter only depends on the data through the statistic $t \equiv \sqrt{2}\bar{x}/s$:

$$\pi(\theta | \vec{x}) = \frac{e^{-(\theta^2/(1+t^2))}}{\sqrt{\pi(1+t^2)}} \left[1 + \operatorname{erf}\left(\frac{t\theta}{\sqrt{1+t^2}}\right) \right] = p(\theta | t), \quad (3)$$

where erf is the error function. Furthermore, the sampling distribution of t turns out to depend on θ only and is a noncentral Student's t distribution for one degree of freedom and with noncentrality parameter θ :

$$p(t | \theta) = \frac{e^{-\theta^2}}{\pi(1+t^2)} + \frac{\theta t e^{-(\theta^2/(1+t^2))}}{\sqrt{\pi}(1+t^2)^{3/2}} \left[1 + \operatorname{erf}\left(\frac{\theta t}{\sqrt{1+t^2}}\right) \right]. \quad (4)$$

It is clear that there exists no prior (no function of θ only) that, multiplied by the likelihood (4), leads to the posterior (3). Hence the marginalization paradox: someone who is only given the data value of the statistic t will be able to make inferences about the parameter θ , but these inferences are guaranteed to disagree with those previously made by the Bayesian who had access to the full data set. Resolution of this paradox hinges on the realization that lack of information about μ is not the same as lack of information about θ . Therefore, the choice of which quantity is of interest must be done *before* calculating the prior. Since reference priors are derived from the likelihood function, the latter must first be expressed in terms of the relevant parameters, θ and σ :

$$p(\vec{x} | \theta, \sigma) = \frac{e^{-((n-1)/2)(s/\sigma)^2 - (n/2)((\bar{x}/\sigma) - \theta)^2}}{(\sqrt{2\pi}\sigma)^n}. \quad (5)$$

Applying the reference algorithm to this likelihood while treating θ as the quantity of interest yields the prior $\pi(\theta, \sigma) = 1/(\sigma\sqrt{1 + \theta^2/2})$, which is very different from the prior $1/\sigma$ obtained by treating μ as the quantity of interest. The resulting reference posterior suffers no marginalization problems. Further details about this example can be found in Refs. [8,12].

Our discussion of marginalization also helps to clarify the behavior of reference posteriors under transformations in the multiparameter setting. Reference posteriors are invariant under one-to-one transformations of the parameter of interest but not under transformations that redefine the parameter of interest by mixing in one or more nuisance parameters. However, redefining the nuisance parameters is permitted. Suppose, for example, that φ is the parameter of interest and ν the nuisance parameter(s), and consider an invertible transformation of the form $(\varphi, \nu) \rightarrow (\varphi, \lambda)$, where λ is a function of both φ and ν . Then the reference posterior for φ is unchanged by the transformation.

Reference priors on unbounded parameter spaces are usually improper, which invalidates the application of Bayes' theorem. To circumvent this problem one introduces a nested sequence of compact subsets $\Theta_1 \subset \Theta_2 \subset \dots$ of the parameter space Θ , such that $\Theta_\ell \rightarrow \Theta$ as $\ell \rightarrow \infty$. Given an improper prior $\pi(\theta)$, its restriction to Θ_ℓ will be proper, so that Bayes' theorem can be applied to construct the corresponding restricted posterior $\pi_\ell(\theta | x)$. The unrestricted posterior for the entire parameter space is then defined by the limit of the $\pi_\ell(\theta | x)$ as $\ell \rightarrow \infty$. The practical justification for this procedure is that one often knows the shape, but not the size, of the physical region of parameter space where the prior has nonzero weight. As this size is typically very large, the limiting posterior can be viewed as an approximation to the posterior on the physical region.

Interestingly, the limiting posterior can also be obtained by direct, formal application of Bayes' theorem to the improper prior $\pi(\theta)$, provided the marginal distribution of the data,

$$m(x) \equiv \int p(x | \theta) \pi(\theta) d\theta, \quad (6)$$

is finite. It can then be shown that the restricted posteriors $\pi_\ell(\theta | x)$ converge *logarithmically* to their limit $\pi(\theta | x)$:

$$\lim_{\ell \rightarrow \infty} D[\pi(\theta | x), \pi_\ell(\theta | x)] = 0, \quad (7)$$

where

$$D[p(\theta), q(\theta)] \equiv \int q(\theta) \log \frac{q(\theta)}{p(\theta)} d\theta \quad (8)$$

is the Kullback-Leibler divergence between $p(\theta)$ and $q(\theta)$. This divergence is a parametrization-independent, non-negative measure of the separation between two densities; it is zero if and only if the densities are identical. Unfortunately, pointwise logarithmic convergence is not enough to avoid inferential inconsistency in some special cases [16], so that a stronger form of convergence is needed, *expected* logarithmic convergence:

$$\lim_{\ell \rightarrow \infty} \mathbb{E}_m \{ D[\pi(\theta | x), \pi_\ell(\theta | x)] \} = 0, \quad (9)$$

where the expectation is taken with respect to the marginal density $m(x)$.

The above discussion motivates the following terminology [16]. Given a statistical model on a parameter space Θ , a *standard prior* is a strictly positive and continuous function on Θ that yields a proper posterior. A *permissible prior* is a standard prior for which the posterior is the expected logarithmic limit of a sequence of posteriors defined by restriction to compact sets.

A. The concept of missing information

Reference priors make use of the notion of expected intrinsic information. For one observation from a model $p(x | \theta)$, the expected intrinsic information about the value of θ when the prior is $\pi(\theta)$ is given by the functional

$$I\{\pi\} = \mathbb{E}_m \{ D[\pi(\theta), \pi(\theta | x)] \}. \quad (10)$$

The more informative the observation, the greater the expected separation between the posterior and the prior. The larger this separation, the greater the expected intrinsic information $I\{\pi\}$. Thus, $I\{\pi\}$ measures the amount of information about the value of θ that might be expected from one observation when the prior is $\pi(\theta)$.

Suppose next that we make k independent observations $x_{(k)} = \{x_1, x_2, \dots, x_k\}$ from the model $p(x | \theta)$. The definition of expected intrinsic information can be generalized to include all k observations:

$$I_k\{\pi\} = \mathbb{E}_m \{ D[\pi(\theta), \pi(\theta | x_{(k)})] \}, \quad (11)$$

where the expectation is a k -dimensional integral over $x_{(k)}$ weighted by

$$m(x_{(k)}) \equiv \int p(x_{(k)} | \theta) \pi(\theta) d\theta = \int \left[\prod_{i=1}^k p(x_i | \theta) \right] \pi(\theta) d\theta. \quad (12)$$

As the sample size k grows larger, one expects the amount of information about θ to increase, and in the limit $k \rightarrow \infty$, the true value of θ would become exactly known. In this sense, the limit $I_\infty\{\pi\} \equiv \lim_{k \rightarrow \infty} I_k\{\pi\}$ represents the *missing* information about θ when $\pi(\theta)$ is the prior. This concept of missing information is central to the construction of reference priors.

B. Reference priors for one-parameter models

The goal of reference analysis is to construct a prior that maximizes the missing information. This maximization cannot be done directly, however, because I_∞ typically diverges. To avoid this problem, one first constructs the prior $\pi_k(\theta)$ that maximizes I_k and then takes the limit of $\pi_k(\theta)$ as $k \rightarrow \infty$. Additional care is required when the parameter space Θ is unbounded, since in that case the prior that maximizes I_k is often improper, and I_k is undefined for improper priors. The solution is to define reference priors via their restrictions on arbitrary compact subsets Θ_ℓ of Θ . Thus one is led to the formal definition of a reference prior for θ as any permissible prior $\pi_R(\theta)$ that satisfies the so-called maximizing missing information (MMI) property, namely, that

$$\lim_{k \rightarrow \infty} [I_k\{\pi_{R,\ell}\} - I_k\{\pi_\ell\}] \geq 0 \quad (13)$$

for any compact set Θ_ℓ and candidate prior $\pi(\theta)$, where $\pi_{R,\ell}$ and π_ℓ are the renormalized restrictions of π_R and π to Θ_ℓ . A candidate prior is a standard prior that incorporates any prior knowledge about θ .

A key result is the following constructive definition of the reference prior $\pi_R(\theta)$ [16]:

$$\pi_R(\theta) = \lim_{k \rightarrow \infty} \frac{\pi_k(\theta)}{\pi_k(\theta_0)},$$

$$\text{with } \pi_k(\theta) = \exp \left\{ \int p(x_{(k)} | \theta) \right. \\ \left. \times \ln \left[\frac{p(x_{(k)} | \theta) h(\theta)}{\int p(x_{(k)} | \theta) h(\theta) d\theta} \right] dx_{(k)} \right\}, \quad (14)$$

where θ_0 is an arbitrary fixed point in Θ , $h(\theta)$ is any continuous, strictly positive function, such as $h(\theta) = 1$, and $p(x_{(k)} | \theta) = \prod_{i=1}^k p(x_i | \theta)$ is the probability model for a sample of k independent observations. We emphasize that this constructive definition only guarantees that the MMI property (13) is satisfied. The permissibility part of the reference prior definition must be separately verified. However, the proponents of reference priors view the MMI

property as considerably more important than permissibility [16] and also believe that it would be highly unusual for a prior satisfying the MMI property to fail permissibility [11] (counterexamples are known, but they are rather exotic).

A further useful result that we shall exploit is that, when certain regularity conditions are met—essentially those that guarantee asymptotic normality of the posterior—the reference prior for models with one continuous parameter reduces to the well-known Jeffreys prior [4],

$$\pi_R(\theta) = \sqrt{\mathbb{E}\left[-\frac{d^2}{d\theta^2} \ln p(x|\theta)\right]}, \quad (15)$$

where the expectation is taken with respect to the sampling model $p(x|\theta)$. In general, the analytical derivation of reference priors can be extremely challenging. However, Eq. (14) is amenable to numerical integration [16].

We emphasize that the definition and results described in this section apply only to the case where the model of interest depends on a single parameter. A generalization to the multiparameter case has been formulated and shown to have the properties listed at the beginning of Sec. II [12]. We shall not describe it here, however, except for when evidence-based priors are specified for the additional parameters. This is a very common situation in high energy physics and will be discussed next.

III. NUISANCE PARAMETERS

The reference prior algorithm described in Sec. II B pertains to models containing no nuisance parameters. In practice, however, every nontrivial problem must contend with such parameters and the reference prior algorithm must be generalized accordingly. In this paper we restrict our attention to nuisance parameters for which partial information is available, which is often the case in practice.

Depending on the type of partial information that is available, there are two plausible ways one might choose to incorporate nuisance parameters ϕ into the calculation of the reference priors for a parameter of interest θ [17]:

Method 1: Assume that we are given a marginal prior $\pi(\phi)$ for the nuisance parameters; compute the conditional reference prior $\pi_R(\theta|\phi)$ for the interest parameter given a fixed value of ϕ ; the full prior is then $\pi(\theta, \phi) = \pi_R(\theta|\phi)\pi(\phi)$.

Method 2: Assume that we are given a conditional prior $\pi(\phi|\theta)$ for the nuisance parameter given the interest parameter; marginalize the probability model $p(x|\theta, \phi)$ with respect to ϕ in order to obtain $p(x|\theta) = \int p(x|\theta, \phi)\pi(\phi|\theta)d\phi$, and compute the reference prior $\pi_R(\theta)$ for the marginalized model; the full prior is then $\pi(\theta, \phi) = \pi(\phi|\theta)\pi_R(\theta)$.

In many high energy physics measurements there are often sound reasons for assuming that the nuisance parameter is independent of the parameter of interest. Information about

a detector energy scale, for example, is typically determined separately from the measurement of interest, say, of a particle mass, and is therefore considered to be independent *a priori* from one's information about the particle's mass. When an experimenter is willing to make this assumption, he or she can declare that $\pi(\phi|\theta) = \pi(\phi)$ and use method 2. When this assumption does not seem fully justified, and it is too difficult to elicit the θ dependence of $\pi(\phi|\theta)$, then it will seem preferable to use method 1, which only requires knowledge of the marginal prior $\pi(\phi)$. When one is unsure of which method to use, one should use both and treat the results as part of a test of robustness. An important practical advantage of method 1 is that the conditional reference prior is computed once and for all, for a given model, and can be used with any evidence-based prior for the nuisance parameters. In contrast, for method 2 the reference prior must be computed anew every time the priors for the nuisance parameters change. On the other hand, since method 2 reduces the problem to one involving a single parameter, the reference prior algorithm reduces to Jeffreys' rule (15), which is typically easier to implement.

In the next section we introduce the basic model studied in this paper and follow with the application of methods 1 and 2 to that model.

A. The single-count model

A very common model for high energy physics measurements is the following. A number of events N is observed by some apparatus, and it is assumed that N is Poisson distributed with mean count $\epsilon\sigma + \mu$, where σ is the rate of a physics signal process, typically the cross section, which we detect with an effective integrated luminosity ϵ —that is, the integrated luminosity scaled by the signal efficiency—and μ is a background contamination. Thus, σ is the parameter of interest, whereas ϵ and μ are nuisance parameters for which we usually have partial information. For physical reasons none of these three parameters can be negative. We write the likelihood for this model as

$$p(n|\sigma, \epsilon, \mu) = \frac{(\epsilon\sigma + \mu)^n}{n!} e^{-\epsilon\sigma - \mu}, \quad (16)$$

$$\text{with } 0 \leq \sigma < \infty \quad \text{and} \quad 0 < \epsilon, \mu < \infty.$$

Information about ϵ and μ usually comes from a variety of sources, such as auxiliary measurements, Monte Carlo simulations, theoretical calculations, and evidence-based beliefs (for example, some sources of background contributing to μ may be deemed small enough to ignore, and some physics effects on ϵ , such as gluon radiation, may be believed to be well enough reproduced by the simulation to be reliable “within a factor of 2”). It is therefore natural to represent that information by an evidence-based prior. Here we will assume that ϵ and μ are independent of σ and that their prior factorizes as a product of two gamma

densities:

$$\pi(\epsilon, \mu | \sigma) = \pi(\epsilon, \mu) = \frac{a(a\epsilon)^{x-1/2} e^{-a\epsilon}}{\Gamma(x + \frac{1}{2})} \frac{b(b\mu)^{y-1/2} e^{-b\mu}}{\Gamma(y + \frac{1}{2})}, \quad (17)$$

where a , b , x , and y are known constants, related to the means $\bar{\epsilon}$ and $\bar{\mu}$ and coefficients of variation $\delta\epsilon$ and $\delta\mu$ by

$$\begin{aligned} \bar{\epsilon} &= \frac{x + \frac{1}{2}}{a}, & \delta\epsilon &= \frac{1}{\sqrt{x + \frac{1}{2}}}, \\ \bar{\mu} &= \frac{y + \frac{1}{2}}{b}, & \delta\mu &= \frac{1}{\sqrt{y + \frac{1}{2}}}. \end{aligned} \quad (18)$$

The built-in assumption that ϵ and μ are uncorrelated is clearly an approximation, since they share a dependence on the integrated luminosity, which is itself uncertain.

There are two ways of interpreting this prior. The first one is appropriate when information about ϵ and μ comes from one or more nonexperimental sources, such as Monte Carlo studies and theoretical calculations, and takes the form of a central value plus an uncertainty. Since the ϵ and μ components of the prior are each modeled by a two-parameter density, one can fix the shape of this density in each case by matching its mean with the central value of the corresponding measurement and its standard deviation with the uncertainty. It will then be necessary to check the robustness of the final analysis results to reasonable changes in this procedure. For example, one may want to replace the gamma distribution by a log-normal or truncated Gaussian one and the mean by the mode or median.

The second interpretation of prior (17) follows from the analysis of two independent, auxiliary Poisson measurements, in which the observed number of events is x for the effective luminosity and y for the background. The expected numbers of events in these auxiliary measurements are $a\epsilon$ and $b\mu$, respectively. For a Poisson likelihood with mean $a\epsilon$ the reference prior coincides with Jeffreys' prior and is proportional to $1/\sqrt{\epsilon}$. Given a measurement x , the posterior will then be a gamma distribution with shape parameter $x + 1/2$ and scale parameter $1/a$. A similar result holds for the background measurement. In this manner the prior (17) is obtained as a joint reference posterior from two auxiliary measurements.

The problem we are interested in is finding a prior for σ , about which either little is known or one wishes to act as if this is so.

1. Application of method 1 to the single-count model

This section serves two purposes: to illustrate the analytical algorithm for computing reference priors and to apply method 1 to model (16).

In method 1 [17], we find first the conditional reference prior $\pi_R(\sigma | \epsilon, \mu)$ and then multiply by the evidence-based

prior $\pi(\epsilon, \mu)$ to construct the full prior $\pi(\sigma, \epsilon, \mu)$. As will be illustrated in Sec. IV, the single-count model is regular enough to warrant using Jeffreys' rule in the first step of the calculation of $\pi_R(\sigma | \epsilon, \mu)$. We therefore apply Eq. (15) to the σ dependence of the likelihood (16), while holding ϵ and μ constant; this yields

$$\begin{aligned} \pi_J(\sigma | \epsilon, \mu) &\propto \sqrt{\mathbb{E}\left[-\frac{\partial^2}{\partial \sigma^2} \ln p(n | \sigma, \epsilon, \mu)\right]} \\ &\propto \frac{\epsilon}{\sqrt{\epsilon\sigma + \mu}}. \end{aligned} \quad (19)$$

This prior is clearly improper with respect to σ and is therefore only defined up to a proportionality constant. However, this constant could very well depend on ϵ and μ , since we kept these parameters fixed in the calculation. It is important to obtain this dependence correctly, as examples have shown that otherwise inconsistent Bayes estimators may result. Reference [17] proposes a compact subset normalization procedure. One starts by choosing a nested sequence $\Theta_1 \subset \Theta_2 \subset \dots$ of compact subsets of the parameter space $\Theta = \{(\sigma, \epsilon, \mu)\}$, such that $\cup_\ell \Theta_\ell = \Theta$ and the integral $K_\ell(\epsilon, \mu)$ of $\pi_J(\sigma | \epsilon, \mu)$ over $\Omega_\ell \equiv \{\sigma: (\sigma, \epsilon, \mu) \in \Theta_\ell\}$ is finite. The conditional reference prior for σ on Ω_ℓ is then

$$\pi_{R,\ell}(\sigma | \epsilon, \mu) = \frac{\pi_J(\sigma | \epsilon, \mu)}{K_\ell(\epsilon, \mu)}. \quad (20)$$

To obtain the conditional reference prior on the whole parameter space, one chooses a fixed point $(\sigma_0, \epsilon_0, \mu_0)$ within that space and takes the limit of the ratio

$$\pi_R(\sigma | \epsilon, \mu) \propto \lim_{\ell \rightarrow \infty} \frac{\pi_{R,\ell}(\sigma | \epsilon, \mu)}{\pi_{R,\ell}(\sigma_0 | \epsilon_0, \mu_0)}. \quad (21)$$

By taking the limit in this ratio form, one avoids problems arising from $K_\ell(\epsilon, \mu)$ becoming infinite as $\ell \rightarrow \infty$.

The theory of reference priors currently does not provide guidelines for choosing the compact sets Θ_ℓ , other than to require that the resulting posterior be proper. In most cases this choice makes no difference and one is free to base the choice of compact sets on considerations of simplicity and convenience. However, we have found that some care is required with the single-count model. Indeed, suppose we make the plausible choice

$$\Theta_\ell = \{(\sigma, \epsilon, \mu): \sigma \in [0, u_\ell], \epsilon \in [0, v_\ell], \mu \in [0, w_\ell]\}, \quad (22)$$

where $\{u_\ell\}$, $\{v_\ell\}$, and $\{w_\ell\}$ are increasing sequences of positive constants. If we use these sets in applying Eqs. (20) and (21) to the prior (19), we obtain

$$\pi_R(\sigma | \epsilon, \mu) \propto \sqrt{\frac{\epsilon}{\epsilon\sigma + \mu}}. \quad (23)$$

Although this prior is still improper with respect to σ , its

dependence on ϵ is different from that of the conditional Jeffreys' prior, Eq. (19). This demonstrates the potential importance of the compact subset normalization. The prior in Eq. (23) has a serious problem, however. Suppose that the ϵ marginal of our evidence-based prior for ϵ and μ is $\exp(-\epsilon)/\sqrt{\pi\epsilon}$. It is then easy to verify that the resulting posterior is improper, since its ϵ marginal has the non-integrable form $\exp(-\epsilon)/\epsilon$. The cause of this problem is the choice of compact sets (22).

Fortunately it is not difficult to find a sequence of compact sets that will provide a proper posterior. Indeed, the σ dependence of the prior (19) suggests that the compact sets should be based on the parametrization $(\epsilon\sigma, \epsilon, \mu)$ rather than (σ, ϵ, μ) [18]. We therefore set

$$\Theta_\ell = \{(\sigma, \epsilon, \mu): \sigma \in [0, u_\ell/\epsilon], \epsilon \in [1/v_\ell, v_\ell], \mu \in [0, w_\ell]\}, \quad (24)$$

where u_ℓ, v_ℓ , and w_ℓ are as before. Again using Eqs. (19)–(21), we now find

$$\pi_{R1}(\sigma | \epsilon, \mu) \propto \frac{\epsilon}{\sqrt{\epsilon\sigma + \mu}}, \quad (25)$$

which is identical to Jeffreys' prior for this problem and yields well-behaved posteriors. For future use, the subscript $R1$ on the left-hand side indicates that this reference prior was obtained with method 1.

We now have all the ingredients needed to calculate the marginal reference posterior $\pi_{R1}(\sigma | n)$ for the cross section σ : the likelihood (16), the marginal nuisance prior (17), and the conditional reference prior (25). For calculating posterior summaries in terms of intervals and upper limits it is convenient to express the result as a tail probability:

$$\begin{aligned} & \int_\sigma^\infty \pi_{R1}(\tau | n) d\tau \\ &= \int_{\sigma/(a+\sigma)}^1 \frac{u^{n+y}(1-u)^{x-1/2}}{B(n+y+1, x+\frac{1}{2})} \\ & \quad \times \frac{B_{[b/(b+1)]}[1+(1-1/u)\sigma/a](y+\frac{1}{2}, n+\frac{1}{2})}{B_{b/(b+1)}(y+\frac{1}{2}, n+\frac{1}{2})} du, \end{aligned} \quad (26)$$

where

$$B_z(u, v) \equiv \int_0^z t^{u-1}(1-t)^{v-1} dt \quad (27)$$

is the incomplete beta function, and $B(u, v) \equiv B_1(u, v) = \Gamma(u)\Gamma(v)/\Gamma(u+v)$.

2. Application of method 2 to the single-count model

In contrast with method 1, method 2 requires from the start that we specify the evidence-based prior for the effective integrated luminosity ϵ and the background contamination μ . Furthermore, this specification must be done

conditionally on the signal rate σ . As mentioned earlier, we will use expression (17) for this prior.

The next step in the application of method 2 is to marginalize the probability model (16) with respect to ϵ and μ :

$$\begin{aligned} p(n | \sigma) &= \iint p(n | \sigma, \epsilon, \mu) \pi(\epsilon, \mu | \sigma) d\epsilon d\mu \\ &= \iint \frac{(\epsilon\sigma + \mu)^n}{n!} e^{-\epsilon\sigma - \mu} \frac{a(a\epsilon)^{x-1/2}}{\Gamma(x+\frac{1}{2})} \\ & \quad \times e^{-a\epsilon} \frac{b(b\mu)^{y-1/2}}{\Gamma(y+\frac{1}{2})} e^{-b\mu} d\epsilon d\mu \\ &= \left[\frac{a}{a+\sigma} \right]^{x+1/2} \left[\frac{b}{b+1} \right]^{y+1/2} S_n^0(\sigma), \end{aligned} \quad (28)$$

where

$$\begin{aligned} S_n^m(\sigma) &\equiv \sum_{k=0}^n k^m \binom{k+x-\frac{1}{2}}{k} \binom{n-k+y-\frac{1}{2}}{n-k} \\ & \quad \times \left[\frac{1}{b+1} \right]^{n-k} \left[\frac{\sigma}{a+\sigma} \right]^k, \end{aligned} \quad (29)$$

and the binomial coefficients are expressed in terms of gamma functions to accommodate noninteger values of their arguments. Finally, the reference prior algorithm must be applied to the marginalized model $p(n | \sigma)$. As in the case of method 1, the conditions for applying Jeffreys' rule are satisfied here; we therefore obtain

$$\pi_{R2}(\sigma) \propto \sqrt{\sum_{n=0}^{\infty} \frac{[(x+\frac{1}{2})S_n^0(\sigma) - \frac{a}{\sigma}S_n^1(\sigma)]^2}{(a+\sigma)^{x+5/2}S_n^0(\sigma)}}. \quad (30)$$

We will use the notation $\pi_{R2}(\sigma)$ to refer to the marginal reference prior for σ obtained with method 2. Note that the compact subset argument invoked in the construction of the method 1 reference prior is not needed here because all the parameters other than σ have already been eliminated by marginalization.

For method 2 the marginal reference posterior for σ is proportional to the product of the marginal data probability distribution (28) and the marginal reference prior (30):

$$\pi_{R2}(\sigma | n) \propto p(n | \sigma) \pi_{R2}(\sigma). \quad (31)$$

The normalization of $\pi_{R2}(\sigma | n)$ must be obtained numerically.

B. The multiple-count model

An important generalization of the single-count model is obtained by considering M replications of the latter; the likelihood is

$$p(\vec{n} | \sigma, \vec{\epsilon}, \vec{\mu}) = \prod_{i=1}^M \frac{(\epsilon_i\sigma + \mu_i)^{n_i}}{n_i!} e^{-\epsilon_i\sigma - \mu_i}. \quad (32)$$

To obtain the method 1 reference prior for this model, we

first calculate Jeffreys' prior for σ , while keeping $\vec{\epsilon}$ and $\vec{\mu}$ fixed:

$$\pi_J(\sigma | \vec{\epsilon}, \vec{\mu}) \propto \sqrt{\sum_{i=1}^M \frac{\epsilon_i^2}{\epsilon_i \sigma + \mu_i}}. \quad (33)$$

This prior is improper, requiring us to apply the compact subset normalization described in Sec. III A 1. Using a straightforward generalization of the nested compact sets of Eq. (24), we find that the correct reference prior is identical to Jeffreys' prior.

In order to apply method 2, we need to specify a proper conditional prior for the μ_i and ϵ_i given σ . Neglecting correlations, we set

$$\pi(\vec{\epsilon}, \vec{\mu} | \sigma) = \prod_{i=1}^M \frac{a_i (a_i \epsilon_i)^{x_i - 1/2} e^{-a_i \epsilon_i}}{\Gamma(x_i + 1/2)} \frac{b_i (b_i \mu_i)^{y_i - 1/2} e^{-b_i \mu_i}}{\Gamma(y_i + 1/2)}. \quad (34)$$

The marginalized data probability distribution $p(\vec{n} | \sigma)$ is then a product of expressions of the form (28), one for each count i .

Here we no longer attempt to obtain analytical expressions for the method 1 and 2 reference posteriors. Instead, we use the numerical algorithms described below.

C. Numerical algorithms

In this section we describe numerical algorithms that can be used to compute method 1 or 2 reference posteriors for the single- and multiple-count Poisson likelihoods discussed in the previous sections.

For method 1 the algorithm starts by generating $(\sigma, \vec{\epsilon}, \vec{\mu})$ triplets from the ‘‘flat-prior posterior,’’ i.e. the posterior obtained by setting $\pi(\sigma | \vec{\epsilon}, \vec{\mu}) = 1$ (line 3 in the pseudo-code below); the correct reference prior $\pi(\sigma | \vec{\epsilon}, \vec{\mu})$ is then computed at lines 4–7 and is used at line 9 to weight the generated σ values so as to produce the reference posterior:

- 1 Set \vec{n}_o to the array of observed event numbers.
- 2 For $i = 1, \dots, I$:
- 3 Generate $(\sigma_i, \vec{\epsilon}_i, \vec{\mu}_i) \sim p(\vec{n}_o | \sigma, \vec{\epsilon}, \vec{\mu}) \pi(\vec{\epsilon}, \vec{\mu})$.
- 4 For $j = 1, \dots, J$:
- 5 Generate $\vec{n}_j \sim p(\vec{n} | \sigma_i, \vec{\epsilon}_i, \vec{\mu}_i)$.
- 6 Calculate $d^2[-\ln p(\vec{n}_j | \sigma_i, \vec{\epsilon}_i, \vec{\mu}_i)]/d\sigma_i^2$
by numerical differentiation.
- 7 Average the J values of $d^2[-\ln p(\vec{n} | \sigma_i, \vec{\epsilon}_i, \vec{\mu}_i)]/d\sigma_i^2$ obtained at line 6, and take the square root. This yields a numerical approximation to the conditional Jeffreys' prior $\pi_J(\sigma_i | \vec{\epsilon}_i, \vec{\mu}_i)$.
- 8 Histogram the σ_i values generated at line 3, weighting them by $\pi_J(\sigma_i | \vec{\epsilon}_i, \vec{\mu}_i)/p(\vec{n}_o | \sigma_i, \vec{\epsilon}_i, \vec{\mu}_i)$. This yields $\pi_{R1}(\sigma)$, the σ -marginal prior.
- 9 Histogram the σ_i values generated at line 3, weighting them by $\pi_J(\sigma_i | \vec{\epsilon}_i, \vec{\mu}_i)$. This yields $\pi_{R1}(\sigma | \vec{n}_o)$, the σ -marginal posterior.

Although not required for the calculation of the reference posterior, an approximation to the reference prior is provided at line 8. By construction this approximation is only reliable for σ values in the bulk of the flat-prior posterior. The generation step at line 3 is done via a Markov chain Monte Carlo procedure [19]. The particular choice of sampling distribution for the generated $(\sigma, \vec{\epsilon}, \vec{\mu})$ triplets is motivated by the desire to obtain weights with reasonably small variance at steps 8 and 9. However, the flat-prior posterior $p(\vec{n}_o | \sigma, \vec{\epsilon}, \vec{\mu}) \pi(\vec{\epsilon}, \vec{\mu})$ is not always proper with respect to $(\sigma, \vec{\epsilon}, \vec{\mu})$. When $M = 1$ (single-count model), it is improper if $x \leq 1/2$. Propriety can then be restored by multiplying the flat-prior posterior by ϵ and correspondingly adjusting the weights at steps 8 and 9. Another feature of the above algorithm is that it does not implement the compact subset normalization. In the cases that we examined, this procedure made no difference, but this may not be true for more general problems than those our code seeks to solve. Unfortunately the current lack of guidelines in the choice of compact sets limits our ability to address this issue in the code.

The algorithm for method 2 has a simpler structure, since all it does is apply Jeffreys' rule to a marginalized likelihood $p(\vec{n}_o | \sigma)$ provided by the user. The calculation does not require random sampling of the parameters and is done at fixed σ values. For a given σ , the reference prior $\pi_{R2}(\sigma)$ is obtained by Monte Carlo averaging, over an ensemble of vectors \vec{n} generated from $p(\vec{n} | \sigma)$, of an accurate numerical approximation of the second derivative of the negative log-likelihood [20]. As already pointed out, method 2 does not require a compact subset normalization procedure. The reference posterior is thus proportional to the product of $p(\vec{n}_o | \sigma)$ and $\pi_{R2}(\sigma)$, and the normalization with respect to σ must be determined numerically.

IV. VALIDATION STUDIES

We have performed a number of studies to validate inferences from the single-count model, using both the numerical algorithms described in Sec. III C and analytical expressions we obtained for the marginal method 1 and 2 posteriors for σ . To recapitulate, we have two reference priors for this model:

$$\pi_{R1}(\sigma, \epsilon, \mu) = \pi_{R1}(\sigma | \epsilon, \mu) \pi(\epsilon, \mu), \quad (35)$$

$$\pi_{R2}(\sigma, \epsilon, \mu) = \pi_{R2}(\sigma) \pi(\epsilon, \mu | \sigma), \quad (36)$$

and we have assumed that $\pi(\epsilon, \mu | \sigma) = \pi(\epsilon, \mu)$ at Eq. (17). As explained in Sec. III, this extra assumption affects only the definition of π_{R2} , which therefore incorporates more information than π_{R1} . In the present section we study and compare the properties of these two reference priors. To begin, we show some example prior and posterior σ marginals in Fig. 2.




FIG. 2. Left: Marginal method 1 and 2 priors, normalized to 1 at $\sigma = 1$. Right: Marginal method 1 and 2 posteriors for 0, 1, and 4 observed events, together with the posteriors obtained from a flat prior. The ϵ and μ priors have a mean of 1 and a 20% coefficient of variation [corresponding to $x = y = 24.5$ and $a = b = 25$ in Eq. (18)]. Here and in subsequent plots, the units of σ are arbitrary but consistent with those of ϵ ; e.g., if the latter is expressed in pb^{-1} , then σ is given in pb so that, like μ , the product of ϵ and σ is dimensionless.

As expected, posteriors corresponding to a small observed number of events favor small cross sections, and posteriors derived from flat priors put less weight on small cross sections than reference posteriors.

Our derivations of the two reference prior methods made use of Jeffreys' rule (15). As pointed out in Sec. II B, this approach assumes that some regularity conditions are satisfied, such that the resulting posterior is asymptotically normal. We now wish to verify this assumption with a graphical example. If one adopts the objective Bayesian view that the parameters σ , ϵ , and μ have true values, then the asymptotic limit can be defined as the result of a large number N_R of replications of the measurement, in the limit where that number goes to infinity. For the case where each measurement replication i consists of a number of events n_i drawn from a probability mass function $f(n | \sigma, \epsilon, \mu)$, the reference posterior has the form:

$$\begin{aligned} \pi_R(\sigma, \epsilon, \mu | n_1, n_2, \dots, n_{N_R}) \\ \propto \pi_R(\sigma, \epsilon, \mu) \prod_{i=1}^{N_R} f(n_i | \sigma, \epsilon, \mu), \end{aligned} \quad (37)$$

and the reference prior $\pi_R(\sigma, \epsilon, \mu)$ is calculated from the combined likelihood for the N_R measurements; it can also be calculated from a single one of these likelihood functions, since it follows from their constructive definition (14) that reference priors are independent of sample size. For method 1 the prior is given by the product of Eqs. (17) and (25) and the likelihood component $f(n | \sigma, \epsilon, \mu)$ by Eq. (16). Replicating the measurement N_R times is then equivalent to making a single measurement with a Poisson

likelihood whose mean is N_R times the original mean and whose observation is the sum of the N_R original observations n_i . This property of Poisson measurements simplifies the calculations considerably. For method 2 the prior is given by Eq. (30) and the likelihood by Eq. (28). In this case no simplification obtains when considering multiple replications, and numerical calculations must use explicitly the full product of likelihood functions.

Figure 3 illustrates the calculations with the help of so-called Q - Q plots, where recentered quantiles from the reference posterior for σ are plotted against standard normal quantiles. The posterior quantiles Q_π are recentered according to $\tilde{Q}_\pi = (Q_\pi - \langle \sigma \rangle) / \Delta \sigma$, where the posterior mean $\langle \sigma \rangle$ and standard deviation $\Delta \sigma$ are numerically estimated. For method 1 we set the true values of σ , ϵ , and μ to 1. We then randomly generate a sequence of 100 independent measurements from the probability mass function (16) and use the subsequences with $N_R = 1, 10$, and 100 to produce the curves in the left panel. For method 2 we set the true value of σ to 1 and give the priors for ϵ and μ each a mean of 1 and a coefficient of variation of 20%. Measurements are then generated from the probability mass function (28) in order to compute the curves in the right panel. Both panels clearly show that the respective reference posteriors approach a Gaussian shape as the number of measurement replications increases.

Given the almost negligible difference between method 1 and 2 posteriors exhibited in Fig. 2 and the fact that our analytical results for method 1 are computationally more tractable than those for method 2, our considerations in the remainder of this section will focus exclusively on method 1.

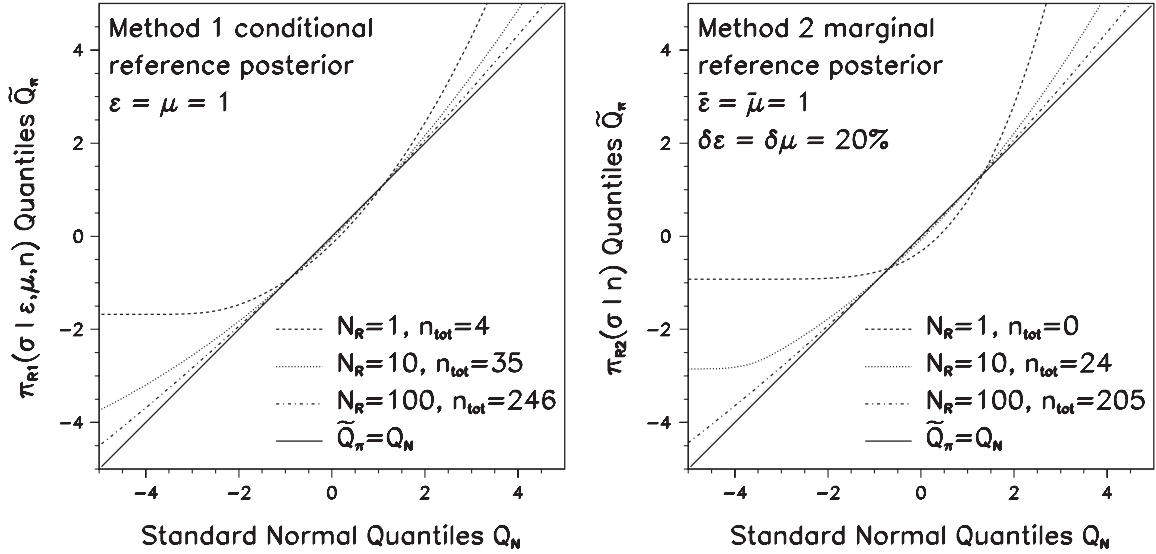


FIG. 3. Q - Q plots of the method 1 conditional reference posterior (left) and the method 2 marginal reference posterior (right) versus the standard normal distribution. The reference posterior quantiles have been recentered (see text). N_R is the number of measurement replications and n_{tot} is the observed number of events summed over all replications.

Among the reference prior properties listed in Sec. II, the ones of generality, invariance, and coherence are true by construction. The property of sampling consistency needs more elaboration, however, since Bayesian inferences do not generally coincide with exact frequentist ones, and a proper evaluation requires first of all the specification of an ensemble of experiments. A well-known property of Bayesian posterior intervals constructed from a proper prior is that their coverage is exact when averaged over the prior [21]. This is an immediate consequence of the law of total probability. Indeed, given a parameter θ , with proper prior $\pi(\theta)$, and a measurement X , the prior-averaged frequentist coverage of a $1 - \alpha$ Bayesian credibility interval $R(X)$ can be written as

$$\begin{aligned} \mathbb{E}_\pi[\mathbb{P}(\theta \in R(X) | \theta)] &= \mathbb{P}(\theta \in R(X)) \\ &= \mathbb{E}_m[\mathbb{P}(\theta \in R(X) | X)] \\ &= \mathbb{E}_m(1 - \alpha) = 1 - \alpha, \end{aligned} \quad (38)$$

where the first expectation is over the prior $\pi(\theta)$ and the second one over the marginal sampling distribution $m(x) = \int p(x | \theta)\pi(\theta)d\theta$. When $\pi(\theta)$ is a reference prior, and especially when it is improper, there is no natural metric over which the coverage can be averaged. The only sensible approach in that case is to study the coverage pointwise, i.e. as a function of the true value of θ . Since the single- and multiple-count models discussed in this paper combine an improper prior for the parameter of interest σ with proper priors for the nuisance parameters $\tilde{\epsilon}$ and $\tilde{\mu}$, we will study interval coverage for a fixed value of σ , but averaged over $\pi(\tilde{\epsilon}, \tilde{\mu})$. Our interest is in how this coverage evolves toward the asymptotic limit. As before, we take this limit in the sense of an ever-increasing number N_R of

experiment replications. For a given value of N_R , the posterior is formed as in Eq. (37) and its coverage is computed.

Figure 4 shows the coverage of 95% credibility upper limits and 68% credibility central intervals as a function of N_R . As the latter increases, the coverage converges to the credibility, confirming the sampling consistency of the method.

Finally, we examine the behavior of reference posterior upper limits on σ as a function of the expected background $\tilde{\mu}$ [defined in Eq. (18)] when the observed number of events n is small. For comparison, when $n = 0$ and there are no uncertainties on signal efficiency and background, frequentist upper limits decrease linearly with background. From a Bayesian point of view this result is surprising. Indeed, when zero events are observed the likelihood function factorizes exactly into background and signal components, indicating that the experiment *actually* performed can be analyzed as the combination of two independent experiments, one to measure background and the other signal. If, in addition, signal and background are *a priori* independent, then posterior inferences about signal will be independent of background. In particular, upper limits on σ will be constant as a function of $\tilde{\mu}$, not linearly decreasing. The reference priors entangle signal and background, however, so that upper limits will not be exactly constant. The $n = 0$ case is illustrated in Fig. 5 for two values of the relative uncertainties on background and signal efficiency.

For $n > 0$ the likelihood function still factorizes approximately since $(\mu + \epsilon\sigma)^n \approx \mu^n(1 + n\epsilon\sigma/\mu) \approx \mu^n$ for $\mu \gg \epsilon\sigma, n$. Thus upper limits will flatten out at large $\tilde{\mu}$, as seen in Fig. 5. A comparison of the left and right panels in that figure also shows that upper limits increase

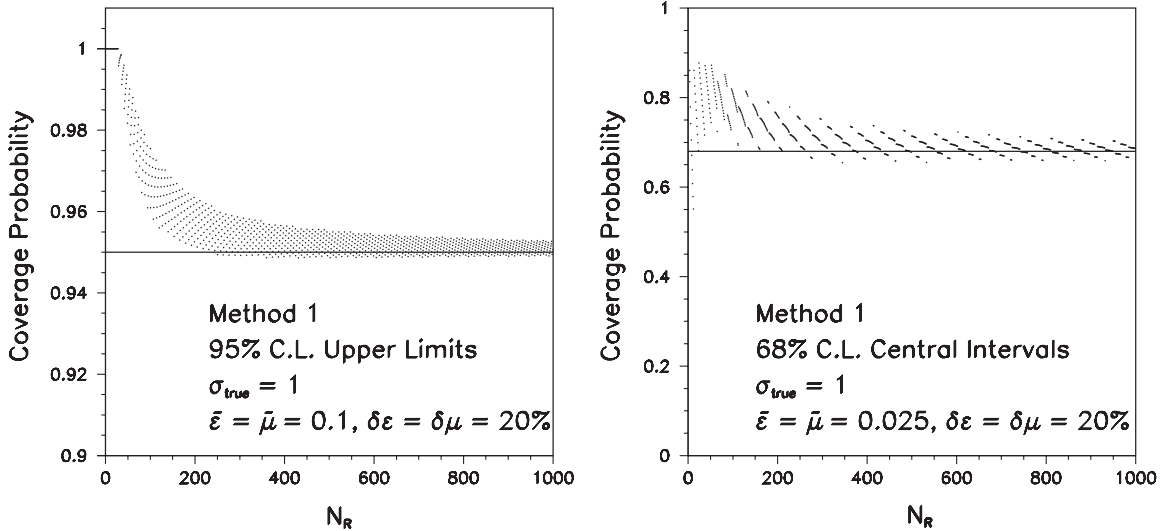


FIG. 4. Coverage probability of method 1 posterior credibility upper limits (left) and central intervals (right), as a function of the number of experiment replications N_R . The solid lines indicate the nominal credibility.

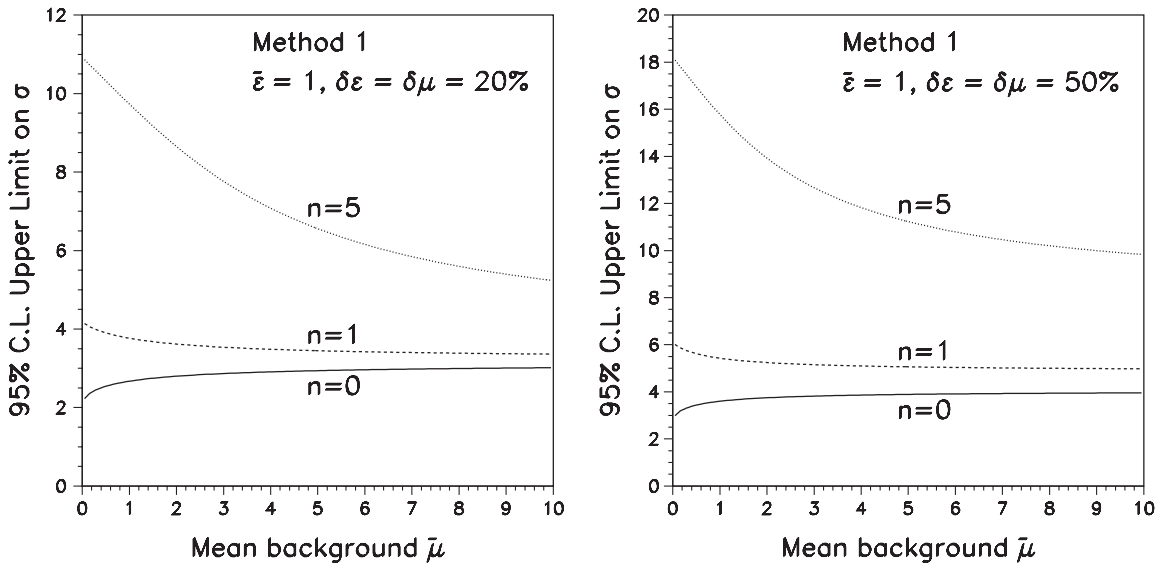


FIG. 5. Variation of the method 1 reference posterior upper limit with mean background for several values of the observed number of events n . The relative uncertainty on the background and on the effective luminosity is 20% for the left plot and 50% for the right one.

with the uncertainty on background and signal efficiency, as expected.

V. MEASUREMENT OF THE SINGLE-TOP QUARK CROSS SECTION

In this section, we demonstrate the computational feasibility of the methods described above by applying them to the recent measurement of the single-top cross section by the D0 and CDF Collaborations [22,23]. Both Collaborations use the same form of likelihood function—a product of Poisson distributions over multiple bins of a multivariate discriminant, the same form of evidence-based priors, namely, truncated Gaussians, and flat priors

for the cross section [24]. As a realistic example, we construct the reference prior for the cross section using one of the data channels considered by D0.

D0 partitioned their data into 24 channels, defined by lepton flavor (electron or muon), jet multiplicity (two, three, or four), number of b -tagged jets (one or two), and two data collection periods. The discriminant distribution is shown in Fig. 3 of Ref. [22]. Here we consider the electron, two-jet, single-tag channel from one of the data-taking periods. The discriminant distribution contains about 500 counts spread over 50 bins, with a maximum bin count of about 40.

We model information about the effective integrated luminosity ϵ and the background μ for each bin with the




FIG. 6 (color online). (a) Prior densities computed using method 1 (histogram) and method 2 (dashed curve), and (b) corresponding posterior densities, for one of the channels in the D0 single-top measurement and using 2.3 fb^{-1} of data. For comparison we show the posterior density computed using a flat prior (dotted curve).

help of the gamma priors of Eq. (17). These evidence-based priors describe the uncertainty due to the finite statistics of the Monte Carlo simulations. We do not include systematic uncertainties in this example. Figure 6(a) shows a comparison of the reference prior for the cross section using methods 1 (the histogram) and 2 (the dashed curve). The jaggedness of the method 1 prior reflects the fluctuations due to the Markov chain Monte Carlo [19] sampling of the parameters. The increased jaggedness at large σ is due to the fact that the numerical algorithm samples from the flat-prior posterior, whose density rapidly decreases in this region. It is noteworthy that the priors computed using the two methods are very similar for this particular example. This is also reflected in the similarity of the posterior densities, shown in Fig. 6(b). In principle, fluctuations in the calculated posterior can be made arbitrarily small by increasing the size of the Monte Carlo sample. For reference, Fig. 6(b) also shows the posterior density using a flat prior for the cross section. An obvious conclusion can be drawn: when the data set is large, here of order 500 events, the precise form of the reference prior is not important. However, for small data sets—which is typical of searches for new or rare phenomena—one should expect the form of the prior to matter. It is then important to use a prior with provably useful properties, such as the ones enumerated in Sec. II.

VI. DISCUSSION

Our main purpose in this paper was to propose a set of formal priors with properties that make them attractive for use in the analysis of high energy physics data. Aside from the theoretical properties of invariance, coherence and sampling consistency, the reference prior method has three important practical advantages: (i) priors can be defined for

almost any problem, regardless of the complexity of the likelihood function and the number of nuisance and interest parameters, (ii) in contrast with flat priors, reference priors have so far always yielded proper posteriors, and (iii) reference priors are computationally tractable, as shown by the single-top example.

Here we have limited our numerical investigations to the class of likelihood functions that are derived from Poisson probability mass functions. For this class the method 1 reference prior agrees with Jeffreys' rule. For other classes the compact subset normalization argument may introduce a difference. A possible generalization of our treatment is to unbinned likelihoods. Since our method 1 and 2 numerical algorithms make no assumptions about the likelihood function, they can be generalized to the unbinned case. However, the method 1 algorithm does not implement the compact subset normalization and is therefore only applicable to cases where this procedure makes no difference. Method 2 requires no compact subset normalization but makes an extra assumption about the conditional nuisance prior.

For problems that involve a single continuous parameter or that can be reduced to this case by a method-2-type integration, Ref. [12] proposes a numerical algorithm that is based directly on Eq. (14) and is therefore very general. However, we found that this algorithm presents some difficulties for the complicated likelihood functions used in high energy physics. One difficulty is the round-off error in the product of large numbers of probability densities. Another difficulty is the assessment of the convergence of the integrals in the formula and of the convergence of the finite-sample priors to the reference prior.

Another possible generalization is to problems with more than one parameter of interest, as, for example, in

the measurement of the individual single-top production cross sections in the s and t channels. For this situation the reference prior algorithm requires one to sort the parameters of interest by order of importance [12], and the results may depend on this ordering. A possible interpretation of this dependence is that it is a measure of the robustness of the result to the choice of prior. This is an area that requires more study.

We have developed a general-use software package that implements the methods described in this paper and have released it to the Physics Statistics Code Repository (phystat.org).

Finally, we note that the main ideas underlying the construction of reference priors, namely, generality, reparametrization invariance, coherence, and sampling consistency,

have motivated the development of methods for summarizing reference posteriors via point estimates, intervals, and hypothesis tests. This subfield of objective Bayesianism is known as reference analysis [12,13].

ACKNOWLEDGMENTS

We thank the D0 Collaboration for its support of this work and for granting us permission to use a channel of its single-top data. We also thank the members of the CMS Statistics Committee for many useful discussions. This work was supported in part by the U.S. Department of Energy under Grants No. DE-FG02-91ER40651, No. DE-FG02-04ER41305, and No. DE-FG02-95ER40896.

-
- [1] C.P. Robert, *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation* (Springer, New York, 2007), 2nd ed.; E.T. Jaynes, *Probability Theory: The Logic of Science*, edited by G.L. Bretthorst (Cambridge University Press, Cambridge, England, 2003); A. O'Hagan, *Kendall's Advanced Theory of Statistics*, Bayesian Inference Vol. 2B (Edward Arnold, London, 1994); H. Jeffreys, *Theory of Probability* (Oxford University Press, Oxford, 1961), 3rd ed.
- [2] See, for example, O'Hagan (Ref. [1]).
- [3] For a discussion of different types of informative priors, see, e.g., D.R. Cox, in *Proceedings of the PHYSTAT LHC Workshop on Statistical Issues for LHC Physics*, edited by H.B. Prosper, L. Lyons, and A. de Roeck (CERN, Geneva, 2008), pp. 3–7 (CERN Yellow Report No. CERN-2008-001, 2008).
- [4] R.E. Kass and L. Wasserman, *J. Am. Stat. Assoc.* **91**, 1343 (1996).
- [5] I. Bertram, G. Landsberg, J. Linnemann, R. Partridge, M. Paterno, and H.B. Prosper, Fermilab Report No. FERMILAB-TM-2104, 2000.
- [6] L. Demortier, in *Proceedings of the Conference on Advanced Statistical Techniques in Particle Physics*, edited by M.R. Whalley and L. Lyons (University of Durham, Durham, 2002), pp. 145–151 (Durham Report No. IPPP/02/39, 2002).
- [7] There is no agreement on what to call such priors. In the literature one finds many names: objective, nonsubjective, noninformative, minimally informative, default, neutral, conventional, reference. However, we prefer the term formal prior because it reflects the manner in which they are constructed.
- [8] J.M. Bernardo, *J. R. Stat. Soc. Ser. B* **41**, 113 (1979).
- [9] J.O. Berger and J.M. Bernardo, *J. Am. Stat. Assoc.* **84**, 200 (1989).
- [10] J.O. Berger and J.M. Bernardo, *Biometrika* **79**, 25 (1992).
- [11] J.O. Berger and J.M. Bernardo, in *Bayesian Statistics* 4, edited by J.M. Bernardo, J.O. Berger, A.P. Dawid, and A.F.M. Smith (Oxford University Press, Oxford, 1992), pp. 35–60, <http://www.uv.es/~bernardo/1992Valencia4Ref.pdf>.
- [12] J.M. Bernardo, in *Handbook of Statistics 25*, edited by D.K. Dey and C.R. Rao (Elsevier, Amsterdam, 2005), pp. 17–90, <http://www.uv.es/~bernardo/RefAna.pdf>.
- [13] L. Demortier, in *Statistical Problems in Particle Physics, Astrophysics, and Cosmology: Proceedings of PHYSTAT05*, edited by L. Lyons and M.K. Ünel (Imperial College Press, London, 2006), pp. 11–14.
- [14] A. Birnbaum, *J. Am. Stat. Assoc.* **57**, 269 (1962); J.O. Berger and R.L. Wolpert, *The Likelihood Principle*, Lecture Notes—Monograph Series Vol. 6 (Institute of Mathematical Statistics, Hayward, California, 1988), 2nd ed.
- [15] A.P. Dawid, M. Stone, and J.V. Zidek, *J. R. Stat. Soc. Ser. B* **35**, 189 (1973).
- [16] J.O. Berger, J.M. Bernardo, and D. Sun, *Ann. Stat.* **37**, 905 (2009), <http://www.uv.es/~bernardo/2009Annals.pdf>.
- [17] D. Sun and J.O. Berger, *Biometrika* **85**, 55 (1998), <http://www.stat.duke.edu/~berger/papers/sun.html>.
- [18] J.O. Berger and D. Sun (private communication).
- [19] A. Caldwell, D. Kollar, and K. Kröninger, *Comput. Phys. Commun.* **180**, 2197 (2009).
- [20] In our calculations we have used the approximation $d^2f/d\sigma^2 \approx [4D(f, h)/3 - D(f, 2h)/12]/h^2 + \mathcal{O}(h^4)$, where $D(f, h) = f(\sigma + h) - 2f(\sigma) + f(\sigma - h)$ and h is the step size.
- [21] J.W. Pratt, *J. R. Stat. Soc. Ser. B* **27**, 169 (1965).
- [22] V.M. Abazov *et al.* (D0 Collaboration), *Phys. Rev. Lett.* **103**, 092001 (2009).
- [23] T. Aaltonen *et al.* (CDF Collaboration), *Phys. Rev. Lett.* **103**, 092002 (2009).
- [24] The improper posterior pathology mentioned in the introduction disappears when data are collected in more than one channel or bin. The CDF and D0 single-top analyses are therefore unaffected by this problem.