

A unified perspective on modified Poisson likelihoods for limited Monte Carlo data

T. Glüsenkamp

*Erlangen Centre for Astroparticle Physics (ECAP),
Erlangen, Germany*

E-mail: thorsten.gluesenkamp@fau.de

ABSTRACT: Counting experiments often rely on Monte Carlo simulations for predictions of Poisson expectations. The accompanying uncertainty from the finite Monte Carlo sample size can be incorporated into parameter estimation by modifying the Poisson likelihood. We first review previous frequentist methods of this type by Barlow et al, Bohm et al, and Chirkin, as well as recently proposed probabilistic methods by the author and Argüelles et al. We show that all these approaches can be understood in a unified way: they all approximate the underlying probability distribution of the sum of weights in a given bin, the compound Poisson distribution (CPD). The Probabilistic methods marginalize the Poisson mean with a distribution that approximates the CPD, while the frequentist counterparts optimize the same integrand treating the mean as a nuisance parameter. With this viewpoint we can motivate three new probabilistic likelihoods based on generalized gamma-Poisson mixture distributions which we derive in analytic form. Afterwards, we test old and new formulas in different parameter estimation settings consisting of a “background” and “signal” dataset. The probabilistic counterpart to the Ansatz by Barlow et al. outperforms all other existing approaches in various scenarios. We further find a surprising outcome: usage of the exact CPD is actually bad for parameter estimation. A continuous approximation performs much better and in principle allows to perform bias-free inference at any level of simulated livetime if the first two moments of the CPD of each dataset are known exactly. Finally, we also discuss the situation where new Monte Carlo simulation is produced for a given parameter choice which leads to fluctuations in the computed likelihood value. Two of the new formulas allow to include this Poisson uncertainty directly into the likelihood which substantially decreases these fluctuations.

KEYWORDS: Analysis and statistical methods; Data analysis; Data Processing

ARXIV ePRINT: [1902.08831](https://arxiv.org/abs/1902.08831)

Contents

1	Introduction	1
2	Previous likelihood approaches for limited Monte Carlo	3
2.1	Statistics of weighted Monte Carlo	3
2.2	Frequentist approaches	6
2.2.1	Barlow/Beeston (1993)	6
2.2.2	Chirkin (2013)	7
2.2.3	Bohm/Zech (2012)	8
2.3	Probabilistic approaches	8
2.3.1	Glüsenkamp (2018)	8
2.3.2	Argüelles et al. (2019)	9
2.4	Summary	10
3	Exact CPD for equal weights	11
4	Further generalizations	12
4.1	Generalization 1 — discrete marginalization of k_{mc} with a Poisson-gamma mixture	12
4.2	Generalization (2) — modeling the CPD for each dataset	14
4.3	Generalization (3) — marginalization of the unknown Poisson mean	15
4.4	Overview	16
5	Incorporating extra prior information	16
5.1	Reducing bias on μ_n	17
5.2	Reducing bias for $\text{var}[W]$	17
5.3	Handling empty bins	17
5.3.1	Strategy 1	18
5.3.2	Strategy 2	18
6	Comparisons using Toy MC	18
6.1	Situation 1: likelihood ratio bias and coverage	19
6.1.1	Equal weight test	19
6.1.2	Background dataset with limited statistics	20
6.1.3	Other scenarios	22
6.2	A new MC set every minimization step	22
7	Conclusion	23
A	Derivation generalization (1)	26
A.1	Equal weights	26
A.2	General weights	28

B Derivation generalization (3)	30
B.1 Single source dataset	30
C Overview of some generic formulas	31
C.1 Poisson-gamma mixture form	31
C.1.1 Simple generalization	31
C.1.2 Extended generalization	31
C.2 Negative binomial form	32
C.2.1 Simple generalization	32
C.2.2 Extended generalization	33
D Supplementary figures	34

1 Introduction

Monte Carlo (MC) simulations are used throughout science for statistical inference. In high-energy physics, modern experiments simulate physical processes with Monte Carlo generators to approximate intractable functions that map certain physical inputs x to observable outputs y . In experiments related to astro-particle physics, for example, cosmic-ray simulations with MC-event generators like *CORSIKA* [1] generate MC events following a certain cosmic-ray primary flux model with a specific composition and energy spectrum. This flux model defines the generation function $f_{\text{gen}}(x; \theta_0)$ which depends on spectral and composition parameters θ_0 and is used to draw MC events from. The input x here could for example be the true primary energy E_{prim} of a given cosmic-ray nucleus. After a full simulation involving physical interactions in the atmosphere and the experimental detector response, the end products of the shower development are usually recorded in bins of an observable quantity by counting. In the high-energy neutrino detector IceCube [2], such an observable quantity is for example the total charge deposited in all Photo-multipliers from cosmic-ray induced muons [3] or neutrinos [4]. Initially, all recorded MC events have equal weights, which is illustrated in figure 1 (a) in the central panel. In order to save computing power, a MC simulation is often re-weighted with a function $f_{x;\text{new}}(\theta)$, which might yield individual event weights different from unity (figure 1 b) given by $w_i = \frac{f_{\text{new}}(x; \theta)}{f_{\text{gen}}(x; \theta_0)}$. One can then bin the MC samples of the observable quantity and obtain an approximation for the expectation value in each bin i as $\lambda_i = \sum_j w_j(\theta)$ (1 (a,b) in the lower panel). The MC events therefore serve as a mapping from parameters θ to expectation values λ_i , which can be used to write down a Poisson likelihood function for i.i.d. (independent and identically distributed) data as

$$L(\theta) = \prod_{\text{bins } i} p(k; \theta) = \prod_{\text{bins } i} \frac{e^{-\lambda_i} \cdot \lambda_i^k}{k!} \text{ with } \lambda_i = \sum_{j=1}^N w_j(\theta). \quad (1.1)$$

As can be seen in figure 1 (a,b) lower panel, a suitable change in θ results in a better match to the data which can be achieved via maximum-likelihood optimization [5] of L (eq. (1.1)) over θ .

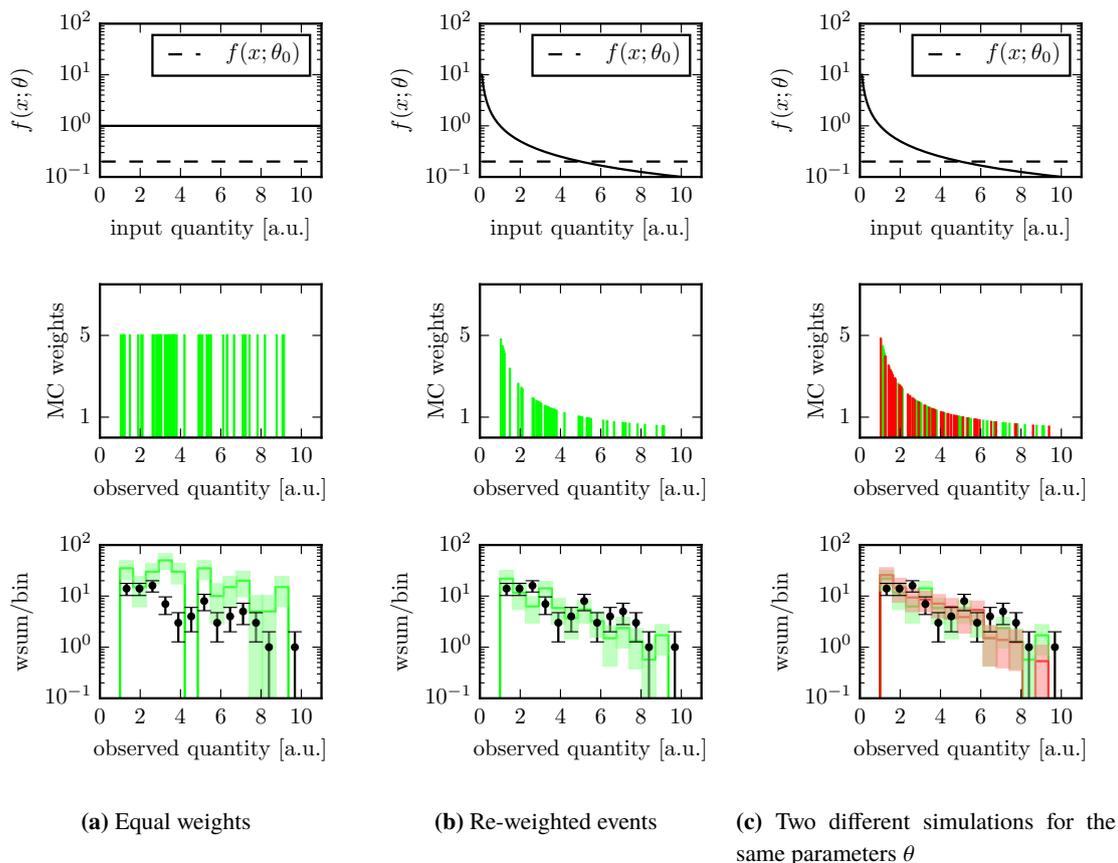


Figure 1. Illustration of the Monte Carlo process for weighted simulation: generation function and re-weighting function (upper panel), resulting sampled and weighted MC events in the unobservable “true” dimension (central panel) and the resulting sum of MC events and some data events (black) in bins of an observable (lower panel). The uncertainty of the MC expectation is indicated by $\sum_i w_i^2$.

It can also happen that re-weighting is not possible, and one has to simulate a whole new set of simulations for a given set of parameters θ . This happens for example in the IceCube experiment for a specific algorithm to reconstruct neutrino directions and energies [6] where the MC events now represent simulated Cerenkov photons emitted by a process parametrized by θ . In addition to the extra computational cost in such a scheme, each MC sample is different from a previous one, even for similar generation parameters θ (see figure 1 c). This generates extra complications for minimization routines due to fluctuations in the resulting likelihood function.

The issues we discuss in this paper arise because the Monte Carlo event count is finite which induces a statistical uncertainty in the evaluation of the related likelihood functions which depend on the MC events and their distribution in the observable bins (see eq. (1.1)). We focus on Poisson-type likelihoods of the form in (eq. (1.1)). The investigation is a continuation of a previous paper [7] which also discussed the more involved multinomial likelihood for unbinned approximations which we omit here.

In the first section, we review previous frequentist and probabilistic approaches for this problem from the last 25 years and indicate how they are related. In the second section, we introduce new

generalized Poisson-gamma mixture distributions that model the underlying statistical distributions more manifestly or take care of extra uncertainty so far neglected. In the third section, we discuss techniques to incorporate extra prior information into the likelihood and a scheme to incorporate empty bins. In the final section, we compare different methods in various benchmark scenarios. All results of this paper are motivated by the use case in high-energy physics, but they are generally applicable in situations where MC simulations are used to calculate Poisson likelihoods.

2 Previous likelihood approaches for limited Monte Carlo

In the following we will revisit the common approaches in the literature that deal with limited Monte Carlo statistics. We will change the notation of some of them from their original publications in order to get a unified understanding.

2.1 Statistics of weighted Monte Carlo

In order to better understand the approaches in the literature, let us first take a look at the actual random variable describing weighted events in a bin i . For a given set of generation parameters θ_0 and re-weighting parameters θ , the weight of a Monte Carlo event ending up in bin i follows a continuous and in general intractable distribution with random variable $W_i(\theta_0, \theta)$, where θ_0 includes a parameter N_{tot} that specifies the total number of events simulated.¹ For a given N_{tot} , there is a resulting Poisson distributed number of MC events N_i for bin i with unknown underlying mean $\mu_i = N_{\text{tot}} \cdot p_i$. The probability p_i is some unknown probability for an event to land in the specific bin i . As long as the number of bins is sufficiently large, say $n_{\text{bins}} > 10$, and p_i sufficiently small, this follows from the usual Poisson limit of the binomial distribution. The actual observed weights are samples from the PDF corresponding to W_i . While the underlying PDF of W_i is unknown, one could imagine repeating the MC experiment with the same parameters θ_0 and θ a large number of times in order to visualize the PDF corresponding to W_i . Figure 2 shows such a PDF of W_i for two different bins. The PDF is shown with and without detector-related smearing in the observable, which could for example be a reconstructed energy. This illustrates that the intractable weight distribution not only depends on the parameters θ and the corresponding bin, but also on detector effects.

From eq. (1.1) we see that we are really interested in a random variable describing the sum of weights Z_i with $Z_i = \sum_{j=1}^{N_i} W_{i,j}$. From the preceding discussion we also see that N_i is not fixed but itself Poisson distributed with unknown mean μ_i . The resulting random variable Z_i is called a compound Poisson distribution (CPD). To our knowledge this explicit statistical description of the problem was first pointed out in [8]. The problematic issue is that not only the PDF of W_i is unknown (see figure 2), but also the value of the underlying mean μ_i . To get a better understanding of the CPD, figure 3 shows two example CPDs and how they are formed from the underlying weight distribution. All approaches described later in this section can be interpreted as different ways to approximate the CPD given the observed MC samples.

¹ N_{tot} could also be itself Poisson distributed, which would result in every N_i be Poisson distributed even without the limit of the binomial distribution.

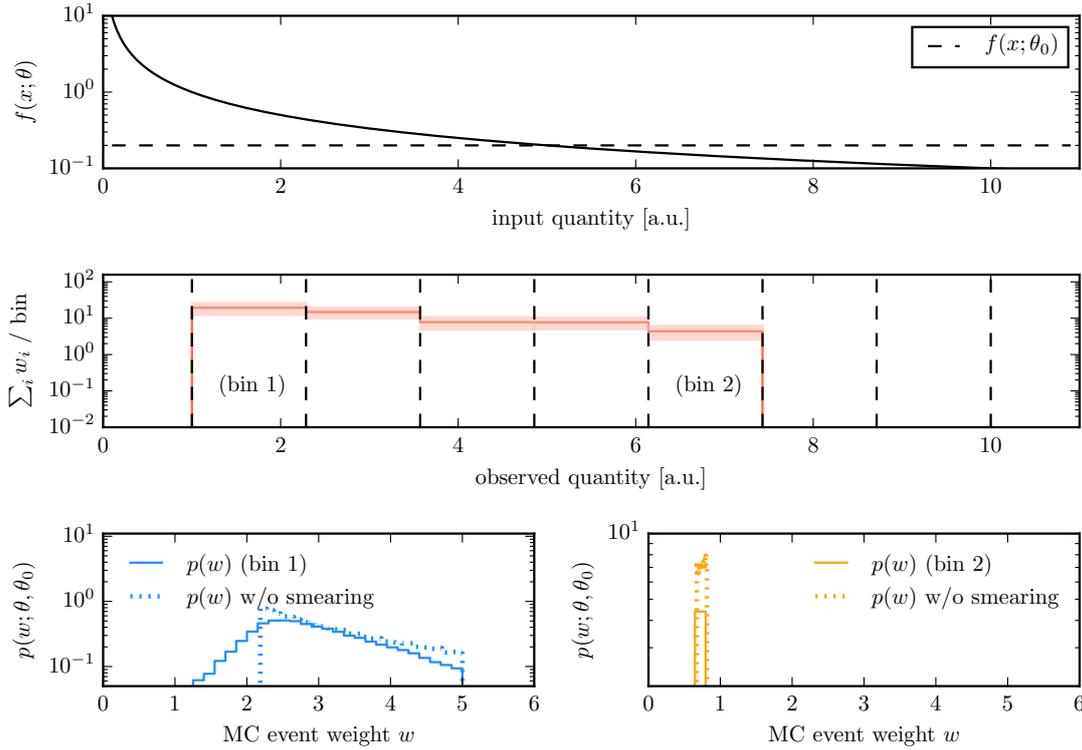


Figure 2. Illustration of the intractable true weight distribution $p(w; \theta, \theta_0)$ of the weight random variable W for given parameters θ_0 and θ in two example bins. The PDF is generated from many MC realizations, and shown with and without reconstruction uncertainty in the observable. The central row shows one particular realization of the sum of weights.

The mean and variance of a compounding Poisson random variable Z is given by

$$E[Z] = \mu_Z = E[N] \cdot E[W] = \mu_N \cdot E[W] \quad (2.1)$$

$$\text{Var}[Z] = E[N] \cdot \text{Var}[W] + (E[W])^2 \cdot \text{Var}[N] \quad (2.2)$$

$$\stackrel{\text{Var}[N]=E[N]}{=} \mu_N \cdot (\text{Var}[W] + (E[W])^2) \quad (2.3)$$

where $Z = \sum_j W_j$, W_j are independent and identical realizations of a continuous random variable W and N is a Poisson random variable defined on \mathbb{N}_0^+ . Relationships (2.1) and (2.3) follow from the law of total expectation and total variance [9], and in the last step we can exploit that mean and variance for a Poisson distribution are the same μ_N . We can estimate the mean and variance using the sample mean $\widehat{\mu}_Z$ and sample variance $\widehat{\text{Var}}_Z$ as

$$\widehat{\mu}_Z = \widehat{\mu}_N \cdot \widehat{\mu}_W = \widehat{\mu}_N \cdot \left(\frac{1}{N} \sum_j w_j \right) = \sum_j w_j \quad (2.4)$$

$$\widehat{\text{Var}}_Z = \widehat{\mu}_N \cdot \widehat{\text{Var}}_W + (\widehat{\mu}_W)^2 \cdot \widehat{\text{Var}}_N \quad (2.5)$$

$$\stackrel{\widehat{\text{Var}}_N = \widehat{\mu}_N}{=} \widehat{\mu}_N \cdot \left(\frac{1}{N} \sum_j (w_j - \widehat{\mu}_W)^2 + \widehat{\mu}_W^2 \right) \quad (2.6)$$

$$= \widehat{\mu}_N \cdot \left(\frac{1}{N} \sum_j w_j^2 \right) = \sum_j w_j^2 \quad (2.7)$$

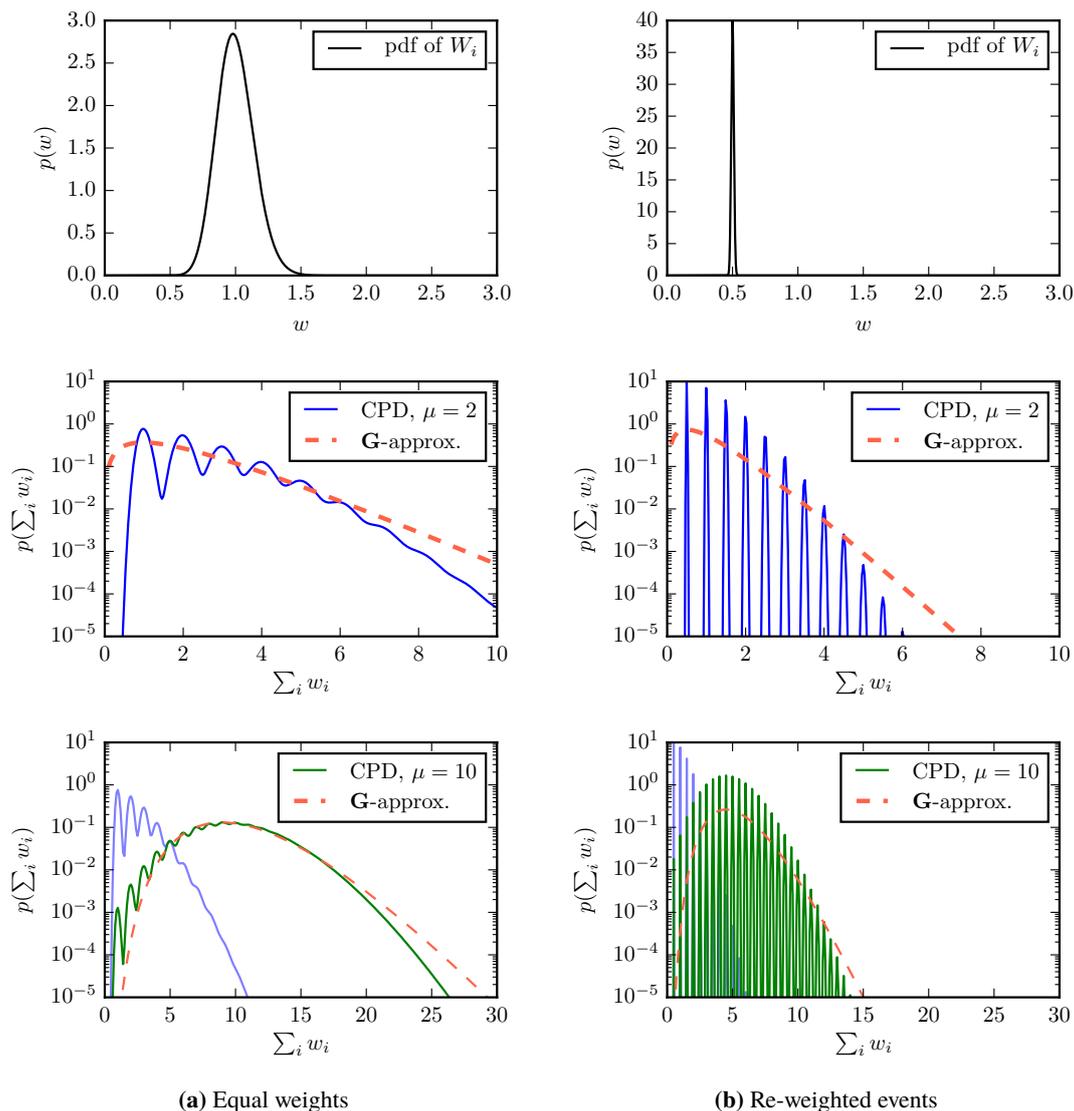


Figure 3. Illustration of the compound Poisson distribution for a weight distribution with a larger spread (a) and smaller spread modeling nearly equal weights (b). The upper row shows the weight distribution $p(w)$, the central row the CPD $p(\sum_i w_i)$ for a mean of $\mu_N = 2$, and the lower row for a larger mean of $\mu_N = 10$. Also illustrated are approximations of the CPD distributions with gamma distributions (red dashed) that encode mean and variance of the CPD as described in eq. (2.1) and eq. (2.3).

where we just replace every term in eqs. (2.1) and (2.3) by the respective sample estimate, using w_j to denote the observed weights. We used the biased sample variance of weights, but could have used the unbiased variance estimator instead. The resulting estimates are incidentally also the sample mean and sample variance of the sum of N weighted Poisson random variables with $\lambda = 1$ each, and they are widely used for plotting purposes to indicate the uncertainty of weighted MC histograms. The mean $\hat{\mu}_N$ is unknown and by approximating it with N there is always some sample uncertainty neglected. In the following we will sometimes use N to denote number of Monte Carlo events or sometimes explicitly call it k_{mc} .

2.2 Frequentist approaches

The approaches in this section are called “frequentist” because they can be thought of as “profile likelihoods” [10] with additional nuisance parameters.

2.2.1 Barlow/Beeston (1993)

The systematic treatment of limited MC statistics for Poisson likelihood applications goes back at least 25 years to Barlow and Beeston [11]. The authors augment the Poisson likelihood with additional Poisson factors for MC datasets, effectively treating MC as additional data. Instead of a per-bin likelihood as described in eq. (1.1), the modified likelihood per bin looks like

$$L_{\text{bin}} = \frac{e^{-\sum_j p_j \widehat{w}_j \lambda_j} \cdot (\sum_j p_j \widehat{w}_j \lambda_j)^k}{k!} \cdot \prod_j^{N_{\text{src}}} \frac{e^{-\lambda_j} \cdot (\lambda_j)^{k_{\text{mc},j}}}{k_{\text{mc},j}!} \quad (2.8)$$

Each factor j enumerates N_{src} individual MC sources,² for example a signal and a background dataset. The weights of Monte Carlo events are averaged per dataset as \widehat{w}_j . The parameters λ_j are nuisance parameters and not of physical interest, and they should be optimized to maximize the values of the individual Poisson factors. The p_j are global all-bin strength factors and can also be optimized as parameters, or fixed if relative strengths are known — in fact they could be incorporated into the weights directly. They are shared for all bins, and an overall scaling by $p_j \rightarrow c \cdot p_j$ of source j would just increase the weights of source j in all bins by the scaling factor c and make p_j redundant if the weights are trusted in an absolute manner. This is actually the case in many modern applications. The nuisance optimization itself effectively handles some of the uncertainty from the finite MC event count. The authors [11] further show how the j parameters λ_j can be reduced to a single one per bin by taking the logarithm and finding the extremal values of $\ln(L_{\text{bin}})$ in λ_j , which results in coupled equations for the λ_j where $j - 1$ of them can be eliminated. The resulting important relations eq. (25) and eq. (26) from [11] look like

$$\frac{k}{(1-t)} = \sum_j \lambda_j \cdot p_j \widehat{w}_j = \sum_j \frac{k_{\text{mc},j} \cdot p_j \widehat{w}_j}{1 + p_j \widehat{w}_j \cdot t} \quad (2.9)$$

in our notation. Taking the logarithm of eq. (2.8), we can write

$$\ln(L_{\text{bin}}) = - \sum_j p_j \widehat{w}_j \lambda_j + k \cdot \ln \left(\sum_j p_j \widehat{w}_j \lambda_j \right) - \ln(k!) + \sum_j \left[-\lambda_j + k_{\text{mc},j} \cdot \ln(\lambda_j) - \ln(k_{\text{mc},j}!) \right] \quad (2.10)$$

$$= - \sum_j \frac{(1 + p_j \widehat{w}_j) \cdot k_{\text{mc},j}}{1 + p_j \widehat{w}_j \cdot t} + k \cdot \ln \left(\frac{k}{1-t} \right) + \sum_j k_{\text{mc},j} \cdot \ln \left(\frac{k_{\text{mc},j}}{1 + p_j \widehat{w}_j \cdot t} \right) - \sum_j \ln(k_{\text{mc},j}!) - \ln(k!) \quad (2.11)$$

$$= -k - k_{\text{mc,tot}} + k \cdot \ln \left(\frac{k}{1-t} \right) + \sum_j k_{\text{mc},j} \cdot \ln \left(\frac{k_{\text{mc},j}}{1 + p_j \widehat{w}_j \cdot t} \right) - \sum_j \ln(k_{\text{mc},j}!) - \ln(k!) \quad (2.12)$$

²We use slightly different notation from the original publication [11]. We write $k_{\text{mc},j}$ instead of a_{ji} , λ_j instead of A_{ji} , and directly use weights as described in section 6 in [11].

where we use eq. (2.9) and insert $\lambda_j = \frac{k_{mc,j}}{1+p_j\widehat{w}_j \cdot t}$ to reduce the j parameters λ_j to a single parameter t except for the first logarithm where we replace $p_j\widehat{w}_j \cdot \lambda_j$ with $\frac{k}{1-t}$. We further simplified eq. (2.11) to eq. (2.12) using the reformulation

$$\sum_j \frac{(1+p_j\widehat{w}_j) \cdot k_{mc,j}}{1+p_j\widehat{w}_j \cdot t} = \sum_j \frac{(1+p_j\widehat{w}_j \cdot t - p_j\widehat{w}_j \cdot t + p_j\widehat{w}_j) \cdot k_{mc,j}}{1+p_j\widehat{w}_j \cdot t} \quad (2.13)$$

$$= k_{mc,tot} + \sum_j \frac{p_j\widehat{w}_j \cdot (1-t) \cdot k_{mc,j}}{1+p_j\widehat{w}_j \cdot t} = k_{mc,tot} + k \quad (2.14)$$

again making use of eq. (2.9). Assuming the p_j are fixed, for example when they can be absorbed into the average weights \widehat{w}_j , we only have to find the value of t that maximizes expression eq. (2.12) for every bin. We can form the derivative of eq. (2.12) with respect to t and obtain

$$\frac{d}{dt} \ln(L_{bin}) = 0 = \frac{k}{(1-t)} - \sum_j \frac{k_{mc,j} \cdot p_j\widehat{w}_j}{1+p_j\widehat{w}_j \cdot t} \quad (2.15)$$

which is just reproducing eq. (2.9). When there is only one source dataset this equation can be solved for t exactly and one obtains

$$t = \frac{k_{mc,1} \cdot p_1\widehat{w}_1 - k}{p_1\widehat{w}_1 \cdot (k + k_{mc,1})} \quad (2.16)$$

which using eq. (2.9) can be used to substitute λ_1 in eq. (2.8) which yields the solution

$$L_{bin,eq.} = \frac{e^{-(k+k_{mc,1})}}{k!k_{mc,1}!} \cdot \left(\frac{k + k_{mc,1}}{1 + 1/(p_1\widehat{w}_1)} \right)^k \cdot \left(\frac{k + k_{mc,1}}{1 + p_1\widehat{w}_1} \right)^{k_{mc,1}} \quad (2.17)$$

with no free parameters. This is also the exact solution when all weights are equal and the average weight \widehat{w}_1 is actually not averaged. In practice, one would insert t into the log-likelihood expression eq. (2.12), but it is instructive to compare eq. (2.17) with other methods later. For more than one dataset we have to find the optimal value of t numerically. The authors also discuss the case when no MC events are present in a bin. In that case they argue that t should only depend on the largest p_j , which then can be used again to calculate the λ_j . While the method by Barlow/Beeston can not be directly interpreted as an approximation of the overall CPD, it can be interpreted as approximating individual CPD's for each dataset, whose random variables are then summed. This will be exploited later in section 4 for a probabilistic generalization.

2.2.2 Chirkin (2013)

In 2013 Chirkin [12] implicitly derived a certain generalization of eq. (2.8). The author starts with a multinomial likelihood and derives a formula that looks very similar to eq. (2.9) in the setting of what he calls “weighted simulation without model errors”. In fact, the resulting equation is the same if we absorb the relative fractions p_j into the weights and if we then treat each MC event independently as coming from its own “source”, i.e. $k_{mc,j} = 1$. From eq. (2.9) we then obtain

$$\frac{k}{(1-t)} = \sum_{i=1}^{k_{mc,tot}} \frac{w_i}{1+w_i \cdot t} \quad (2.18)$$

which is similar to the equation obtained in section 5 in [12].³ The resulting corresponding likelihood formulation equivalent to eq. (2.8) is

$$L_{\text{bin}} = \frac{e^{-\sum_i w_i \lambda_i} \cdot (\sum_i w_i \lambda_i)^k}{k!} \cdot \prod_{i=0}^{k_{\text{mc,tot}}} \frac{e^{-\lambda_i} \cdot (\lambda_i)^1}{1!} \quad (2.19)$$

which now has a product over all individual MC events instead of whole datasets with averaged weights. The method by Chirkin can therefore be interpreted as a certain generalization of the Barlow/Beeston method. It omits weight averaging and gives better results for bias reduction in likelihood scans for a single weighted dataset as demonstrated in [7], and sometimes also for multiple datasets (see later comparison). If all weights are equal the two methods are similar and their likelihoods are given by eq. (2.17) up to constant factors. The method can also be extended to include systematic uncertainties in a log-normal term that is dubbed “model error” or handle bins with no MC using a constant “noise” term. These terms have to be added by hand and are somewhat arbitrary.

2.2.3 Bohm/Zech (2012)

Bohm and Zech discuss another possibility to set up a pseudo-likelihood via

$$L_{\text{bin}} = \frac{e^{-\lambda} \cdot \lambda^k}{k!} \cdot \frac{e^{-\lambda \beta} \cdot (\lambda \cdot \beta)^\alpha}{\Gamma(\alpha + 1)} \quad (2.20)$$

where λ is a nuisance parameter, $\alpha = \frac{(\sum_i w_i)^2}{\sum_i w_i^2}$ and $\beta = \frac{\sum_i w_i}{\sum_i w_i^2}$. A single Poisson factor is used to encode the mean and variance of the sum of weighted MC in terms of effective counts to approximate the CPD. Because of the gamma function instead of the factorial this is not a real Poisson distribution and therefore not a real likelihood, so in general it can not be used for likelihood applications. The authors also only discuss it in the context of goodness of fit tests where they argue it is sensible. For equal weights, however, the gamma function turns into a factorial and the expression is similar to the previous two approaches and eq. (2.20) turns into eq. (2.17). It will still be interesting to compare the general pseudo-likelihood (eq. (2.20)) to its probabilistic counterpart discussed later.

2.3 Probabilistic approaches

In contrast to the previous frequentist approaches this section deals with solutions that integrate nuisance parameters instead of optimizing them.

2.3.1 Glüsenskamp (2018)

In a previous paper [7] we discussed a generalization of the Poisson likelihood of the form

$$L_{\text{bin}}(\theta) = \frac{e^{-\sum_i w_i} \cdot (\sum_i w_i)^k}{k!} = \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot \delta(\lambda - \sum_i w_i) d\lambda \quad (2.21)$$

$$= \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot [\delta(\lambda - w_1) * \dots * \delta(\lambda - w_N)](\lambda) d\lambda \quad (2.22)$$

$$\rightarrow L_{\text{bin}}(\theta) = \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot [\mathbf{G}(\lambda_1; 1 + \alpha/N, 1/w_1) * \dots * \mathbf{G}(\lambda_N; 1 + \alpha/N, 1/w_N)](\lambda) d\lambda \quad (2.23)$$

³Using the notation of [12] we have $d_k = k$, $t_{ki} = 1/w_i$, $t_d = 1$, $\xi_k = t$, $s_{ki} = 1$.

which can be thought of as directly expanding the delta factors into gamma distributions. Each gamma distribution can be derived via Bayesian inference and approximates the underlying weight distribution W_i of the CPD. This means a convolution of these gamma factors is some approximation of the CPD (see figure 4), the difference being a fixed number of terms and different individual PDFs. There is some prior freedom in this approach, which we parametrized by a parameter α which we here call α' . The parameter is shared among individual gamma distributions such that for equal weights it reduces to a single gamma factor with shape parameter $N + \alpha'$. In [7] we observed optimal performance in terms of likelihood ratio bias with respect to infinite statistics for $\alpha' = 0$ for a single weighted dataset. We will later see that this is not true in general. The setting $\alpha = 0$ yields the expected first and second of the CPD as $\mu_{G_1 * \dots * G_N} = \sum_i w_i$ and $\text{var}_{G_1 * \dots * G_N} = \sum_i w_i^2$. It also showed practically identical likelihood ratio behavior as the method by Chirkin [12] for a single source dataset [7]. This is not surprising since a comparison of eq. (2.23) with eq. (2.19) shows that it can be interpreted as the probabilistic counterpart to the approach by Chirkin [12] which enforces the sum constraint via optimization instead of probabilistic convolution (see figure 4). The solution of eq. (2.23) was derived in [7] as an iterative sum via

$$L_{\text{bin}}(\theta) = D_k \cdot \prod_{i=1}^N \left(\frac{1}{1 + w_i} \right)^{1 + \alpha/N} \quad (2.24)$$

where

$$D_k = \frac{1}{k} \sum_{j=1}^k \left[\left(\sum_{i=1}^N (1 + \alpha/N) \cdot \left(\frac{1}{1 + 1/w_i} \right)^j \right) D_{k-j} \right] \quad (2.25)$$

with $D_0 = 1$. The formula allows for an efficient computation.⁴ For equal weights, the result can be written as a simple Poisson gamma mixture as

$$\begin{aligned} L_{\text{bin,eq.}} &= \mathbb{E} \left[\frac{e^{-\lambda} \lambda^k}{k!} \right]_{G(\lambda; N + \alpha, 1/w)} \\ &= \frac{(1/w)^{N + \alpha} \cdot \Gamma(k + N + \alpha)}{\Gamma(N + \alpha) \cdot k! \cdot (1 + 1/w)^{k + N + \alpha}} \end{aligned} \quad (2.26)$$

, where again we use $N = k_{\text{mc}}$. Comparing eq. (2.26) with eq. (2.17) we see that for equal weights the approach has similar dependence on the weights w as the frequentist approaches using the unique prior $\alpha = 0$. Likelihood scans with equal weights therefore give numerically identical confidence intervals as the frequentist methods.

2.3.2 Argüelles et al. (2019)

In Argüelles et al. [13] the authors describe an Ansatz to encode the mean and variance of the CPD in a single gamma distribution which effectively approximates the CPD. It looks like

$$L_{\text{bin}}(\theta) = \int \frac{e^{-\lambda} \cdot \lambda^k}{k!} \cdot \frac{e^{-\lambda \cdot \beta} \cdot (\lambda \cdot \beta)^\alpha}{\Gamma(\alpha + 1)} d\lambda \quad (2.27)$$

⁴See also appendix C for the general form in a slightly different nomenclature.

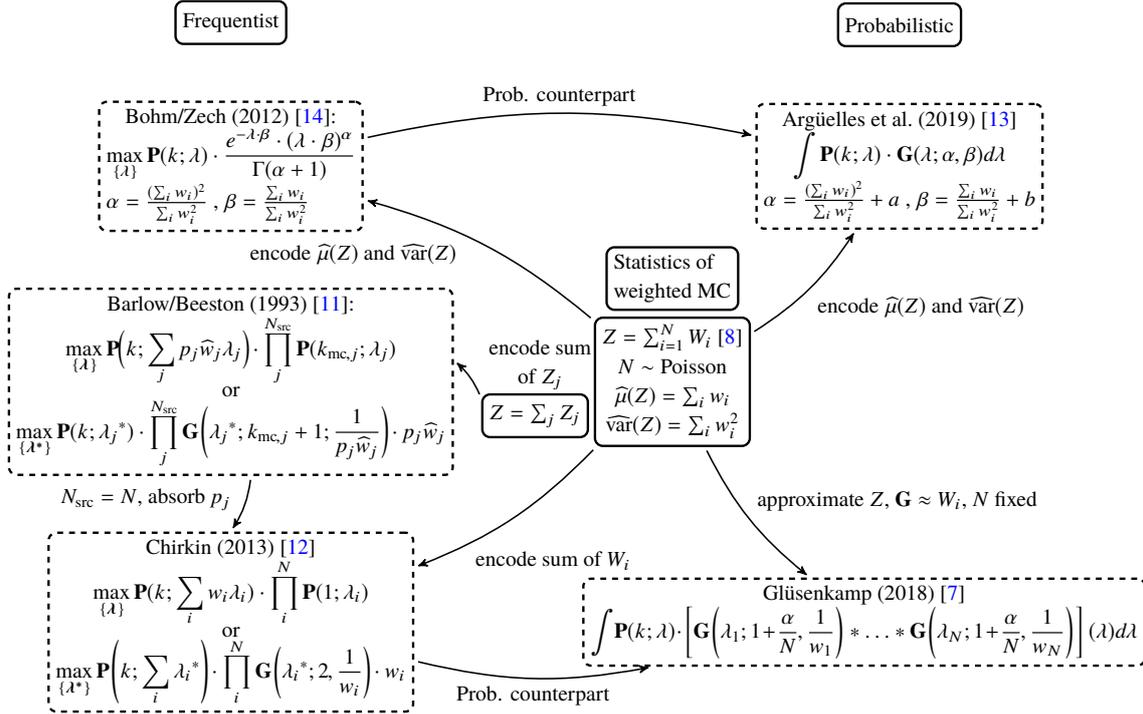


Figure 4. Relationships between existing frequentist and probabilistic approaches in the literature that extend the standard Poisson likelihood for a given bin. N denotes the total number of MC events, N_{src} the total number of source datasets, \mathbf{P} the Poisson distribution, \mathbf{G} the gamma distribution.

where they add some additional effective parameters a and b to expand its possibilities. It can be thought of as a direct probabilistic counterpart of the pseudo-likelihood (eq. (2.20)) described by Bohm and Zech [14] (see figure 4). In contrast to the pseudo likelihood, this Ansatz involves a marginalization with a proper PDF and therefore works for arbitrary weights. The authors argue that the choice of $a = 1$, $b = 0$ works best in their test in terms of proper coverage, in particular better than the choice $a = 0$, $b = 0$ which would just encode the sample mean and sample variance of the CPD in the gamma distribution. We will see later that this is not the case in general, in particular that it depends on the number of datasets combined and the MC statistics per dataset. In section 4.2 we describe an alternative parametrization that directly parametrizes the unknown Poisson mean and unknown weight moments which we later use to incorporate extra Prior information (section 5).

2.4 Summary

The relationship between the different approaches is highlighted in figure 4. It rests on the interpretation that all approaches fundamentally approximate the CPD, an interpretation that has previously only been explicitly made for the formulas put forward by Bohm/Zech [14] and Argüelles et al. [13] and vaguely for our previous construction in [7]. For the frequentist approaches by Barlow/Beeston [11] and Chirkin [12], the prevailing interpretation has been that the Monte Carlo data is simply treated as additional data, as indicated in figure 4 by the upper formulas, respectively. However, by a simple change of variables from λ_i to λ_i^* , one can see that one obtains a formula that almost resembles the probabilistic counterpart, in the case of Chirkin it resembles the probabilistic

formula we proposed earlier in [7]. In this interpretation, the frequentist approaches therefore also approximate the CPD via optimization in λ_i , which now resemble nuisance parameters where each parameter has its own gamma Prior. It rests on the property of the Poisson distribution that it can always be interpreted as a corresponding gamma distribution, whose shape parameter is shifted by one. It turns out that in order for an optimization to give similar results as the corresponding marginalization, the mode of the gamma distribution for optimization has to match the mean during marginalization. This is exactly fulfilled if the shape α is larger by one on the frequentist side, since the mode of the gamma distribution is $\frac{\alpha-1}{\beta}$, while the mean is $\frac{\alpha}{\beta}$. The only difference remaining is the extra scaling factor w_i , that is being pulled out of the gamma factor. This, however, would exactly be the inverse of the Jacobian factor from the variable transform $\lambda_i^* = w_i \cdot \lambda_i$ and cancel, if one actually performed an integration.

3 Exact CPD for equal weights

What if we knew the PDF of the exact CPD, $p_{\text{CPD}}(\lambda)$, and integrate over it

$$L_{\text{bin,exact}} = \int \frac{e^{-\lambda} \cdot \lambda^k}{k!} \cdot p_{\text{CPD}}(\lambda) d\lambda \quad (3.1)$$

instead of using gamma factors in some form? In the special case of equal weights we can actually do this. In [8] it was argued that we can approximate the CPD with a scaled Poisson distribution, but for equal weights the CPD is a scaled Poisson distribution. This results in $p(w) = \delta(\lambda - w)$ and $p_{\text{fix}}(\sum_i w_i) = p(w_1) * \dots * p(w_N) = \delta(\lambda - k_{\text{mc}} \cdot w)$, which yields

$$p_{\text{CPD}}(\lambda) = \sum_{k_{\text{mc}}=0}^{\infty} \mathbf{P}(k; \lambda) \cdot p_{\text{fix}}(\lambda) = \sum_{k_{\text{mc}}=0}^{\infty} \frac{e^{-\mu} \cdot \mu^{k_{\text{mc}}}}{k_{\text{mc}}!} \cdot \delta(\lambda - k_{\text{mc}} \cdot w) \quad (3.2)$$

which is a superposition of the Poisson PDF to see k_{mc} weighted events with the probability distribution p_{fix} to see the weight sum of a given k_{mc} weighted events. So instead of approximating the CPD with gamma factors (see figure 3) we can write down the exact CPD and integrate over it. The exact CPD for equal weights looks roughly equivalent to the example in figure 3 (b) where the weight distribution is very narrow. The final likelihood then can be written as

$$L_{\text{bin,exact,equal}} = \int \frac{e^{-\lambda} \cdot \lambda^k}{k!} \cdot \sum_{k_{\text{mc}}=0}^{\infty} \frac{e^{-\mu} \cdot \mu^{k_{\text{mc}}}}{k_{\text{mc}}!} \cdot \delta(\lambda - k_{\text{mc}} \cdot w) d\lambda \quad (3.3)$$

$$= \sum_{k_{\text{mc}}=0}^{\infty} \frac{e^{-k_{\text{mc}}w} \cdot (k_{\text{mc}}w)^k}{k!} \cdot \frac{e^{-\mu} \cdot \mu^{k_{\text{mc}}}}{k_{\text{mc}}!} \quad (3.4)$$

which can be solved numerically. The parameter μ is an unknown mean, but can be determined by the average of many Monte Carlo runs. The result is a likelihood that does not depend on individual MC realizations anymore and involves the true PDF of the sum of weights without any approximations. It will serve as a crosscheck in the next section.

4 Further generalizations

We will now discuss some further constructions which generalize the previously discussed probabilistic approaches. An overview is shown in figure 5, which we will discuss in detail in the following.

4.1 Generalization 1 — discrete marginalization of k_{mc} with a Poisson-gamma mixture

A first possible generalization starts with the interpretation of eq. (2.23), in which each gamma factor encodes a potential weight distribution. The number of weight distributions is fixed to the number of weighted MC events N . However, we can imagine to marginalize over the shape parameter of each gamma distribution with another discrete distribution, which models the Poisson uncertainty part of the CPD. If all weights are the same, this amounts to integrating out N . For practical reasons, we marginalize with a Poisson-gamma mixture distribution (PG) instead of a Poisson distribution. We can write the result as

$$L_{\text{Gen,(1)}}(\theta) = \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot \left[\left(\sum_{i=0}^\infty G_1(\lambda_1; i + \alpha/N, 1/w_1) \cdot \text{PG}_1(i) \right) * \dots \right. \\ \left. * \left(\sum_{j=0}^\infty G_N(\lambda_N; j + \alpha/N, 1/w_N) \cdot \text{PG}_N(j) \right) \right] (\lambda) d\lambda \quad (4.1)$$

$$= \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot [\text{GPG}_1(\lambda_1) * \dots * \text{GPG}_N(\lambda_N)] (\lambda) d\lambda \quad (4.2)$$

where $\text{PG}(i) = \text{PG}(i; 1 + \beta/N; 1/\nu)$, i.e. the shape parameter of the Poisson-gamma mixture is $1 + \beta/N$ and the rate parameter is $\frac{1}{\nu}$. We also define the mixture of the gamma distribution with the Poisson-gamma distribution as $\text{GPG} = \sum_{i=0}^\infty G_N(\lambda_N; i + \alpha/N, 1/w_N) \cdot \text{PG}_N(i)$. Here we have one mixture for each MC event. If two MC events have the same weight $w_1 = w_2 = w$, the convolution of two GPG distributions behaves similar as the convolution of two gamma distribution with the same rate parameter. We set $\beta = 0$ and $\nu = 1$ for simplicity and can write

$$\sum_{i=0}^\infty G_1(\lambda_1; i + \alpha/N, 1/w) \cdot \text{PG}_1(i) * \sum_{j=0}^\infty G_2(\lambda_2; j + \alpha/N, 1/w) \cdot \text{PG}_2(j) \quad (4.3)$$

$$= \sum_{i=0}^\infty \sum_{j=0}^\infty G(\lambda_1; i + j + 2 \cdot \alpha/N, 1/w) \cdot \text{PG}_1(i) \cdot \text{PG}_2(j) \quad (4.4)$$

$$= \sum_{k=0}^\infty \sum_{j=0}^k G(\lambda_N; k + 2 \cdot \alpha/N, 1/w) \cdot \text{PG}_1(k - j) \cdot \text{PG}_2(j) \quad (4.5)$$

$$= \sum_{k=0}^\infty G(\lambda_N; k + 2 \cdot \alpha/N, 1/w) \cdot \text{PG}(k; 2; 1/\nu) \quad (4.6)$$

When all weights are equal, we can therefore combine all GPG-distributions in eq. (4.2), and obtain

$$\int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot \sum_{i=0}^\infty G(\lambda; i + \alpha, 1/w) \cdot \text{PG}(i; N; 1/\nu) d\lambda \quad (4.7)$$

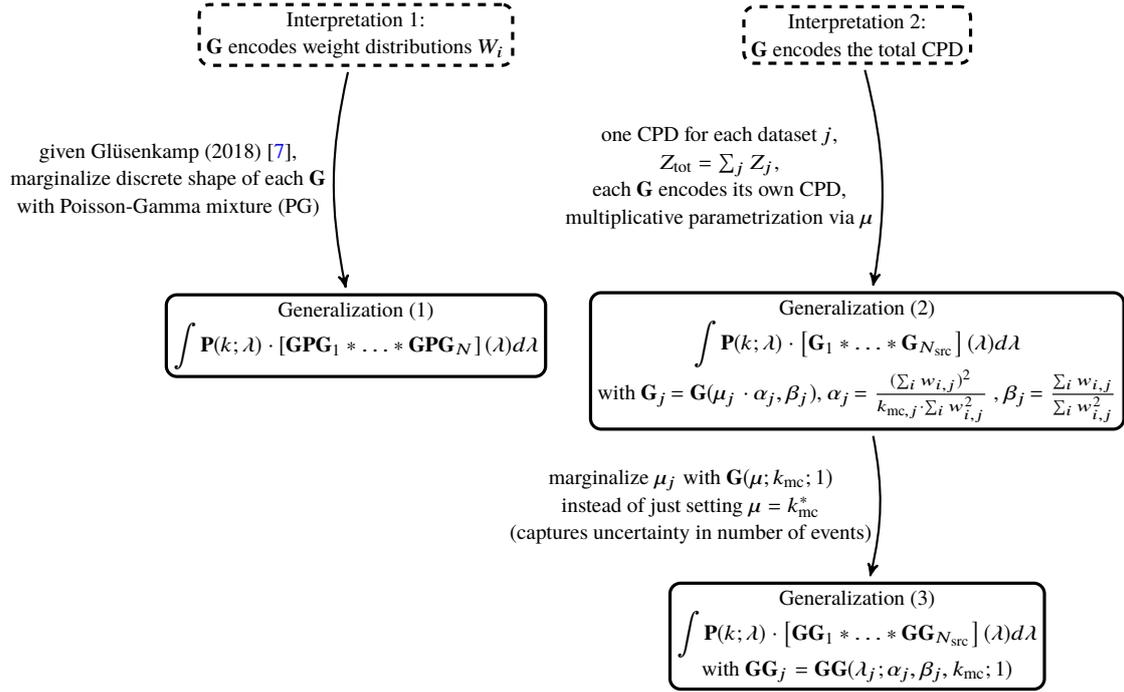


Figure 5. Further generalizations based on two different fundamental interpretations. N denotes the total number of MC events, N_{src} the total number of source datasets, \mathbf{P} the Poisson distribution, \mathbf{G} the gamma distribution, \mathbf{GPG} a gamma-Poisson-gamma mixture, and \mathbf{GG} a gamma-gamma mixture distribution. The mixtures integrate out shape parameters. Z denotes the random variable for the sum of weights, either for all datasets (Z_{tot}), or for individual individual datasets (Z_j).

which represents a direct generalization of the formula involving the exact CPD (eq. (3.3)) for equal weights discussed in section 3. Instead of the delta factor we have a gamma factor, and instead of the Poisson distribution we have a Poisson-gamma mixture distribution. The expression eq. (4.2) therefore involves a certain generalization of a general compound distribution, where each of the continuous distributions W_i is different.

The solution for M different GPG distributions with different weights is derived in appendix A and yields

$$L_{\text{Gen,(1)}}(\theta) \underset{\substack{\alpha=0 \\ \beta=0 \\ \nu_j=1}}{=} \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot \left[\left(\sum_{i=0}^{\infty} \mathbf{G}_1(\lambda_1; i, 1/w_1) \cdot \mathbf{PG}_1(i; 1, 1) \right) * \dots \right. \\ \left. * \left(\sum_{j=0}^{\infty} \mathbf{G}_N(\lambda_N; j, 1/w_N) \cdot \mathbf{PG}_N(j; 1, 1) \right) \right] (\lambda) d\lambda \quad (4.8)$$

$$= \left(\prod_{j=1}^M \cdot \left(\frac{1 + w_j}{1 + 2w_j} \right)^{\langle k_{\text{mc}} \rangle_j} \right) \cdot \Delta_k \quad (4.9)$$

with

$$\Delta_k = \frac{1}{k} \sum_{i=1}^k \left[\left(\sum_{j=1}^M \langle k_{mc,j} \rangle \cdot \left(\frac{2}{2 + 1/w_j} \right)^i - \langle k_{mc,j} \rangle \cdot \left(\frac{1}{1 + 1/w_j} \right)^i \right) \Delta_{k-i} \right] \text{ and } \Delta_0 = 1 \quad (4.10)$$

for the case $\alpha = \beta = 0$ and $\nu = 1$ which we use for all later comparisons. The general solution with arbitrary parameters is given in appendix C. Computationally, the calculation is only a little slower than eq. (2.24) and still scales as $O(N \cdot k^2)$. We tried the same calculation with a Poisson distribution as mixture instead of a Poisson-gamma distribution, but the result is computationally not as efficient. The Poisson-gamma distribution as the Posterior-Predictive distribution of observed counts is also well motivated, because it includes some uncertainty of the unknown underlying mean.

4.2 Generalization (2) — modeling the CPD for each dataset

In the second generalization we model the CPD of each dataset independently, and afterwards join all CPDs via convolution, instead of approximating all datasets with a combined CPD (see figure 5). Before we do so, we first introduce a slightly different parametrization of the CPD compared to the one used by [13]. Encoding mean and variance as defined in eq. (2.1) and eq. (2.3) to be mean and variance of a general gamma distribution with parameters α and β we obtain

$$\alpha = \frac{(\mu_Z)^2}{\text{var}_Z} = \frac{(\mu_N \cdot \text{E}[W])^2}{\mu_N \cdot (\text{Var}[W] + (\text{E}[W])^2)} \stackrel{\text{sample moments}}{=} \frac{\left(\frac{\mu_N}{k_{mc}} \sum_j w_j \right)^2}{\mu_N \cdot \left(\frac{1}{k_{mc}} \sum_j w_j^2 \right)} = \frac{\mu_N \cdot (\sum_j w_j)^2}{k_{mc} \cdot \sum_j w_j^2} \quad (4.11)$$

$$\beta = \frac{\mu_N \cdot \text{E}[W]}{\mu_N \cdot (\text{Var}[W] + (\text{E}[W])^2)} \stackrel{\text{sample moments}}{=} \frac{\sum_j w_j}{\sum_j w_j^2} \quad (4.12)$$

where we can observe that the unknown mean μ_N cancels out in β but remains in α . Choosing $\mu_N = k_{mc} + \frac{k_{mc} \cdot \sum_j w_j^2}{(\sum_j w_j)^2}$ results in $\alpha = \frac{(\sum_j w_j)^2}{\sum_j w_j^2} + 1$ which is the choice that was made in [13]. Instead of an additive nuisance parameter a one can therefore choose a multiplicative parameter that is naturally given by μ_N , which compared to the other choice has an intuitive interpretation as the unknown mean of the Poisson part of the CPD. We also do not have to use the sample based mean and variance for the weight distribution W , but can use other ways to encode the moments. For example a wider variance might be reasonable to model very few events (see section 5) or we can use asymptotic Monte Carlo simulations to get the exact moments for crosschecks (see section 6). The likelihood function with several such CPDs joined by convolution looks like

$$L_{\text{Gen},2} = \int \mathbf{P}(k; \lambda) \cdot [\mathbf{G}(\lambda_1; \alpha_1, \beta_1) * \dots * \mathbf{G}(\lambda_{N_{\text{src}}}; \alpha_{N_{\text{src}}}, \beta_{N_{\text{src}}})] (\lambda) d\lambda \quad (4.13)$$

where each gamma factor has its parameters α and β encoded as described above. Compared to the single CPD Ansatz by Argüelles et al. [13], the advantage is that large differences in statistics between datasets can be explicitly modeled in the respective $\mu_{N,j}$. It also has an interpretation as a direct probabilistic counterpart of the Barlow/Beeston [11] frequentist Ansatz. In contrast to Barlow/Beeston, however, each CPD can take into account the full weight distribution, instead of

averaging the weights. The analytic solution is known from [7] and given in appendix C. It scales computationally as $O(N_{\text{src}} \cdot k^2 + N)$. We can also introduce an effective version

$$L_{\text{Gen},2,\text{eff}} = \int \mathbf{P}(k; \lambda) \mathbf{G}(\lambda; \alpha, \beta) d\lambda \quad (4.14)$$

which is just a single CPD for all weights, similar to Argüelles et al. [13], but with the different parametrization of α and β as introduced above. This allows to model per-dataset mean adjustments or weight distribution approximations as discussed later in section 5.

4.3 Generalization (3) — marginalization of the unknown Poisson mean

Starting with the second generalization, we can go one step further and marginalize the unknown mean $\mu_{N,j}$ of the Poisson part of each CPD j with another gamma distribution. The idea is to capture some of the uncertainty from ignorance about the unknown mean of the number of MC events per bin. We obtain

$$L_{\text{Gen},(3)} = \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot \left[\left(\int G(\lambda_1; \mu_1 \cdot \alpha_j^*, \beta_j) \cdot G\left(\mu_1; \frac{k_{\text{mc},1}}{s}, 1/s\right) d\mu_1 \right) * \dots \right. \\ \left. * \left(\sum_{j=0}^\infty G(\lambda_1; \mu_j \cdot \alpha_j^*, \beta_j) \cdot G\left(\mu_j; \frac{k_{\text{mc},N_{\text{src}}}}{s}, 1/s\right) d\mu_{N_{\text{src}}} \right) \right] (\lambda) d\lambda \quad (4.15)$$

$$= \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot [\text{GG}_1 * \dots * \text{GG}_{N_{\text{src}}}] (\lambda) d\lambda \quad (4.16)$$

$$= \text{PGG}_1 * \dots * \text{PGG}_{N_{\text{src}}} \quad (4.17)$$

because we can write the mixture of a Poisson distribution with a convolution of gamma distributions as a discrete convolution of Poisson-gamma mixtures (see appendix C). The scaling parameter s has a natural choice of $s = 1$. Choosing $s = 2$ results in a gamma distribution with twice the variance but the same mean, and can be used to mimic the variance of a Poisson-gamma mixture distribution. The term $\alpha_j^* = \frac{\mathbb{E}[W_j]^2}{\text{Var}[W_j]}$ is again decoupled from the mean $\mu_{N,j}$, similar to the previous generalization. To simplify terminology we use GG to denote a gamma distribution whose shape is marginalized by another gamma distribution, and PGG to denote a Poisson-gamma mixture distribution whose shape is marginalized with another gamma distribution. Several ways to calculate the solution exist, but we found the convolutional form (eq. (4.17)) to be the most efficient way. Its computational complexity approximately scales as $O(N_{\text{src}} \cdot k^2 + N)$. The calculation of an individual PGG distribution is shown in appendix B and yields

$$\text{PGG}_j(k) = \left(\frac{1}{1+\beta_j} \right)^k \cdot (1/s)^{k_{\text{mc},j}/s} \sum_{n=0}^k \frac{1}{k!} \cdot \alpha_j^{*n} \begin{bmatrix} k \\ n \end{bmatrix} \frac{\Gamma(k_{\text{mc},j}/s + n)}{\Gamma(k_{\text{mc},j}/s) (\gamma - \alpha_j^* \cdot \ln(\frac{\beta_j}{1+\beta_j}))^{k_{\text{mc},j}/s+n}} \quad (4.18)$$

with $\alpha_j^* = \frac{1}{k_{\text{mc},j}} \frac{(\sum_j w_j)^2}{\sum_j w_j^2}$ and $\beta_j = \frac{\sum_j w_j}{\sum_j w_j^2}$ where we inserted sample mean and sample variance of the weight distributions. If not otherwise stated, we use $s = 1$. The unsigned Stirling numbers of the first kind can be precomputed once via the recursion relation [15]

$$\begin{bmatrix} k \\ n \end{bmatrix} = (k-1) \begin{bmatrix} k-1 \\ n \end{bmatrix} + \begin{bmatrix} k-1 \\ n-1 \end{bmatrix} \quad (4.19)$$

and then only require a table lookup.

Table 1. Summary table of different likelihood approaches that modify the Poisson likelihood. The number of Monte Carlo events N , the number of datasets N_{src} and the number of observed data k are to be evaluated per bin. The frequentist approaches (Barlow/Beeston and Chirkin) typically require multiple iterations (“niter”) of a numerical solver when there are more than one source dataset (Barlow/Beeston) or events with different weights (Chirkin).

Name	different weights	extra Poisson uncertainty	Per-dataset mean adjustment	$N \rightarrow \infty$ $\Rightarrow L \rightarrow L_{\text{Poisson}}$ (“probabilistic”)	Computational complexity
Bohm/Zech [14]	✗	✗	✗	✗	$O(N)$
Chirkin [12]	✓	✗	✓(not manifest)	✗	$O(N \cdot \text{niter})$ ($N > 1$)
Barlow/Beeston [11]	✗	✗	✓(not manifest)	✗	$O(N + N_{\text{src}} \cdot \text{niter})$ ($N_{\text{src}} > 1$)
Argüelles et al. [13]	✓	✗	✗	✓	$O(N)$
Glüsenkamp [7]	✓	✗	✓(not manifest)	✓	$O(N \cdot k^2)$
$L_{\text{Gen,(1)}}$	✓	✓	✓(not manifest)	✓	$O(N \cdot k^2)$
$L_{\text{Gen,(2)}}$	✓	✗	✓	✓	$O(N_{\text{src}} \cdot k^2 + N)$
$L_{\text{Gen,(2),eff}}$	✓	✗	✓(not manifest)	✓	$O(N)$
$L_{\text{Gen,(3)}}$	✓	✓	✓	✓	$O(N_{\text{src}} \cdot k^2 + N)$

4.4 Overview

An overview of the properties of all approaches is given in table 1. The first column with weight uncertainty shows that most approaches approximate the underlying weight distribution to some extent. The exceptions are the approaches by Barlow/Beeston and by Bohm/Zech which approximate it essentially with a delta peak by averaging the weights. The second column shows which approaches include extra uncertainty for the unknown mean of the Poisson part of the CPD. Generalizations (1) and (3) are the only ones which do this, either by a discrete marginalization of the observed MC counts or by marginalizing the unknown mean with another probability distribution. The third column discusses the ability to adjust the expected mean for different datasets independently. The frequentist approaches by Chirkin and by Barlow/Beeston in particular do this to some extent, as the profile likelihood fit implicitly varies parameters that can be interpreted as means of different Poisson components. However, in the second interpretation of these approaches (see figure 4), this is not the case anymore. The fourth column indicates that only the probabilistic approaches converge to the absolute value of the Poisson likelihood as more Monte Carlo simulation is used. The frequentist approaches only capture the shape, but their absolute values are meaningless. The last column shows approximate computational complexities. While some approaches scale similarly, for example generalization (2) and (3), one has to remember that individually they require different mathematical operations, so in practice they might differ by a factor of a few.

5 Incorporating extra prior information

While the old and new formulas capture more uncertainty than the standard Poisson likelihood, the approximations are often biased, in particular if only a few MC events are present. This can be seen directly in section 4.2: the mean μ_N is approximated by the observed MC counts and the weight

distribution mean μ_W and variance $\text{var}[W]$ are approximated by the sample mean and variance of observed weights. If only a few events are present in a given bin, these estimates can be biased. For μ_N this happens when the Monte Carlo livetime is so low that several bins are empty, which indicates that the true mean μ_N is smaller than unity per bin. For μ_W and $\text{var}[W]$, this happens when the weight distribution is rather wide and only few events are present. Here we discuss a few approaches that can remedy these biases to some extent.

5.1 Reducing bias on μ_n

If the Monte Carlo livetime is very low it can happen that some bins are empty. Using the observed number of Monte Carlo events in the remaining bins that do see events can lead to an over estimation of the underlying mean. One strategy to circumvent this is to calculate the average number of Monte Carlo events over all bins that are eventually filled by the simulation. If this average is smaller than 1,⁵ we subtract the difference from the sample mean of observed events in every bin. Once the average number of Monte Carlo events per bin is above 1, no correction is performed, so it only affects very low statistics. Importantly, it has to be done for every Monte Carlo dataset separately. If one has two datasets with very different statistics, and calculates this adjustment based on the combined average number of events per bin, it does not work. This is the reason generalizations (2) and (3) make this per-bin mean adjustment natural, because every dataset j has its own mean $\mu_{N,j}$ that gets adjusted. However, it is in principle also possible to do this for most of the other approaches on a per event basis. This is what we call “not manifest” in table 1.

5.2 Reducing bias for $\text{var}[W]$

For very few events the sample variance might give a biased estimate of the true variance of the weight distribution. In particular, for a single event, the sample variance is zero, which is always an underestimation if the true weight distribution has a non-zero width. This happens whenever we have non-equal weights. Simply using the unbiased sample variance often does not help. One possibility to counteract this behavior is to assume that each weight represents a gamma distribution, and the total weight PDF is sum of individual gamma PDFs. The variance of a single gamma distribution with rate parameter $1/w$ is w^2 , so the variance of a superposition of multiple gamma distributions is $\text{var}[W] = \frac{1}{N} \sum_i w_i^2$. We will see in section 6 how this compares to the standard biased sample variance.

5.3 Handling empty bins

Empty bins must be handled with care. In particular, it is important to differentiate between datasets. For one dataset, a bin might be empty, while for another dataset it might be filled with many Monte Carlo events. If nothing is done about that, the resulting parameter estimation, which is likely connected only to one of the datasets, will be biased. In [11], the strategy advocated was to add the dataset with the largest potential contribution based on all datasets that have no Monte Carlo events in the bin to add to the bin. However, this strategy adds discrete jumps to the likelihood, because for different parameters a different dataset might have higher contributions. Another strategy is discussed in [12], which is just to add a constant “noise” contribution to all bins that are empty.

⁵This is an arbitrary choice, but 1 seems sensible and gives reasonable results.

In this strategy, however, no differentiation is made between datasets, which is bad for the reason stated above. The weight of the pseudo event is chosen to be the largest weight from all other bins, similar to the suggestion in [11], but we do it for all datasets without any differentiation. The corresponding CPD approximations of these single “pseudo” Monte Carlo events again typically suffer the problem of an over estimation of the mean, which is discussed in the section 5.1, and are almost by definition bias for the weight distribution. Therefore we typically combine the strategies for empty bins with the mean adjustment in section 5.1 and increased weight variance 5.2. Here we discuss two different strategies to incorporate empty bins. Both strategies basically fill up bins with “pseudo counts” for each dataset, which avoids the problem of jumps in the likelihood. The difference between the strategies involves the type of bins which are filled up.

5.3.1 Strategy 1

In the first strategy we only fill up bins that have at least one Monte Carlo event from any dataset. Once a bin is identified, we add a single “pseudo” Monte Carlo event for every other dataset that has no events in it.

5.3.2 Strategy 2

In the second strategy we fill up all bins that can potentially see Monte Carlo events, even those that see no events from any dataset yet. This is a more aggressive strategy, and leads to potentially larger bias for very low counts. The main aim of this strategy is to reduce likelihood fluctuations in the case when a new Monte Carlo simulation is run every parameter query (section 6.2). These fluctuations tend to be larger when totally empty bins are neglected, which happens by definition in strategy 1.

6 Comparisons using Toy MC

To compare the various approaches, we will use a toy Monte Carlo setting of a falling background energy spectrum with a peaked signal spectrum on top. An artificial detector is simulated with energy smearing and energy-dependent detection efficiency and events are observed in some observable space. We might be interested in determining parameters related to the “signal” distribution like normalization and position, or parameters related to the “background” distribution, like normalization or spectral index. Figure 6 shows likelihood scans for the different standard approaches and the respective Monte Carlo and data distribution in observable space for the best-fit value of the Poisson likelihood. Assuming asymptotic behavior of the log-likelihood ratio $\lambda = -2 \cdot \Delta\text{LLH}(\hat{p}, p_0) \sim \chi^2_\nu$ [16], one can determine the significance of exclusion of a given parameter value p_0 , given \hat{p} the best fit value. In the following tests, instead of using the log-likelihood ratio in the usual way, we show the bias of the log-likelihood ratio λ (in σ equivalents) for a specific data realization compared to infinite Monte Carlo statistics for a given true parameter p_0 . We also show the expected coverage vs. the actual coverage by evaluating 500 Toy experiments for a given simulated livetime.

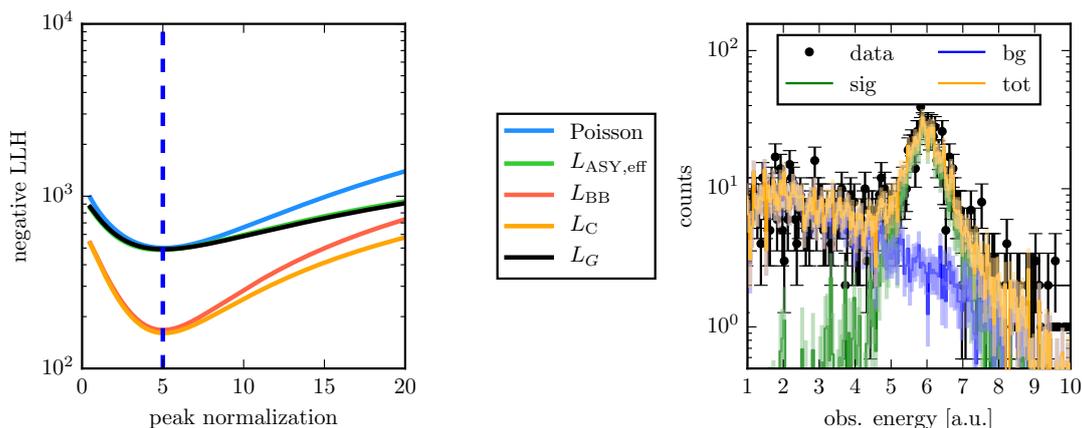


Figure 6. Example likelihood scans for various likelihood approaches from the literature in peak normalization (left). The non-standard likelihoods from other publications include $L_{ASY,eff}$ [13], L_{BB} [11], L_C [12] and L_G [7]. The true parameter is indicated as the dashed vertical line. One data realization and Monte Carlo for the minimum of the Poisson scan is shown on the right.

6.1 Situation 1: likelihood ratio bias and coverage

6.1.1 Equal weight test

The simplest test one can perform is a fit involving only the background dataset. We use a simple weighting function that just scales normalization (see figure 1) and all weights for the background dataset are the same. In this particular case of equal weights, we can calculate the CPD exactly (see section 3). We perform a large number of Monte Carlo simulations at each statistical level to obtain a good estimator of the true mean μ_N , while $\mu_W = w$ and $\text{var}[W] = 0$ are known since we deal with equal weights. We compare this “exact” CPD likelihood with the new generalizations (1), (2), and (3), and a particular form of generalization (2) where we replace the unknown mean μ_N not by the observed number of MC events k_{mc} , but also by the true mean (“exact $E[N]$ ”) obtain from a large number of MC realizations. This can be understood as a continuous approximation of the exact CPD. The comparison of these different likelihoods in terms of λ -bias and coverage is shown in figure 7. The first surprising observation is that the exact CPD ($L_{CPD,exact}$) has a large bias at low average MC events per bin, although one would expect that the exact CPD works at any level of Monte Carlo statistics. This probably has to do with the multi-modal structure of the exact CPD. The continuous approximation ($L_{Gen,2,Z1,exactE[N]}$) on the other hand is basically unbiased everywhere. We find that generalization (1), which is similar in nature to the exact CPD, performs not so good, which is related to the fact that the exact CPD itself performs badly, and we only plot it for the simple case without any modifications. Generalization (2) and generalization (3), which are again continuous approximations, perform better and are almost equivalent in this scenario. We can also observe how the mean adjustment is important to reduce bias, while the strategy for empty bins or using a more generous estimate for the weight variances do not change things too much in this particular scenario. Comparing the approach using the adjusted mean ($L_{Gen,2,M}$) with the exact formula ($L_{Gen,2,Z1,exactE[N]}$), we see that in the intermediate statistics regime there still is some bias

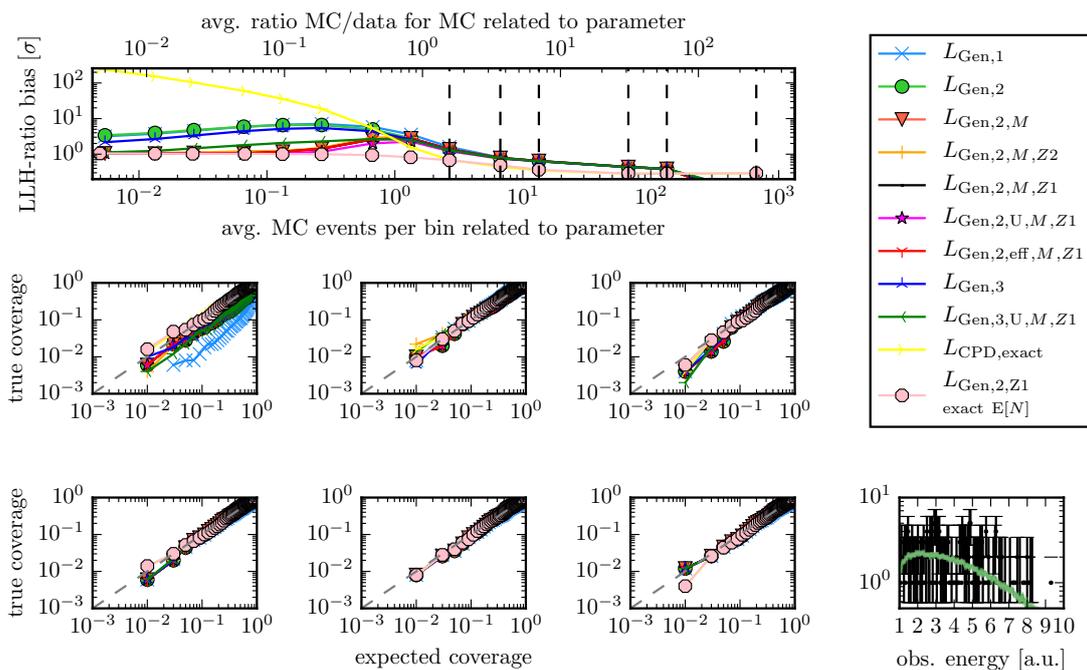


Figure 7. Bias of the log-likelihood ratio λ and coverage in dependence of the background dataset statistics. The signal is not present. The parameter of interest is the normalization of the background dataset. Z1 denotes strategy 1 for “zero MC”, Z2 strategy 2 for “zero MC”, U denotes more “unbiased” weight variance estimation and M denotes shifted mean based on average MC events in all bins. The coverage is shown at 6 statistic levels (in order) which are indicated as dashed vertical lines in the upper plot. In the lower right plot a particular data realization and a MC realization at the highest statistic level is shown for the observable space.

in the likelihood ratio that is not removed by the mean adjustment strategy applied here. One could think about addressing this problem by some modification of the strategy, but we do not further test it here. Figure 8 shows the best performing cases of generalization (2) and generalization(3) in comparison to older likelihood approaches. In general, the new approaches show lower bias and better coverage over a wide range of MC statistics.

6.1.2 Background dataset with limited statistics

In this test setting we fix the background dataset at limited statistics, while the signal statistics are increased step-by-step. We fit the normalization of the signal peak. Such a situation with limited background statistics can happen if hard cuts are applied in the data selection process and computational resources are not sufficient to still have enough simulated background events surviving at final level. We compare various combinations of treating empty bins with incorporating prior information. Again we first compare various strategies with the new formulas as shown in figure 9. We can obtain the exact moments from high statistics simulations and use these moments at the respective level, which results in a unbiased likelihood ratio at all statistic levels for crosschecks ($(L_{\text{Gen},2,Z1,\text{exactmoments}})$). If we use the exact Poisson mean, but use sample estimates for the weight moments ($L_{\text{Gen},2,Z1,\text{exactE}[N]}$), a bias becomes visible. We can compare this result to the likelihood

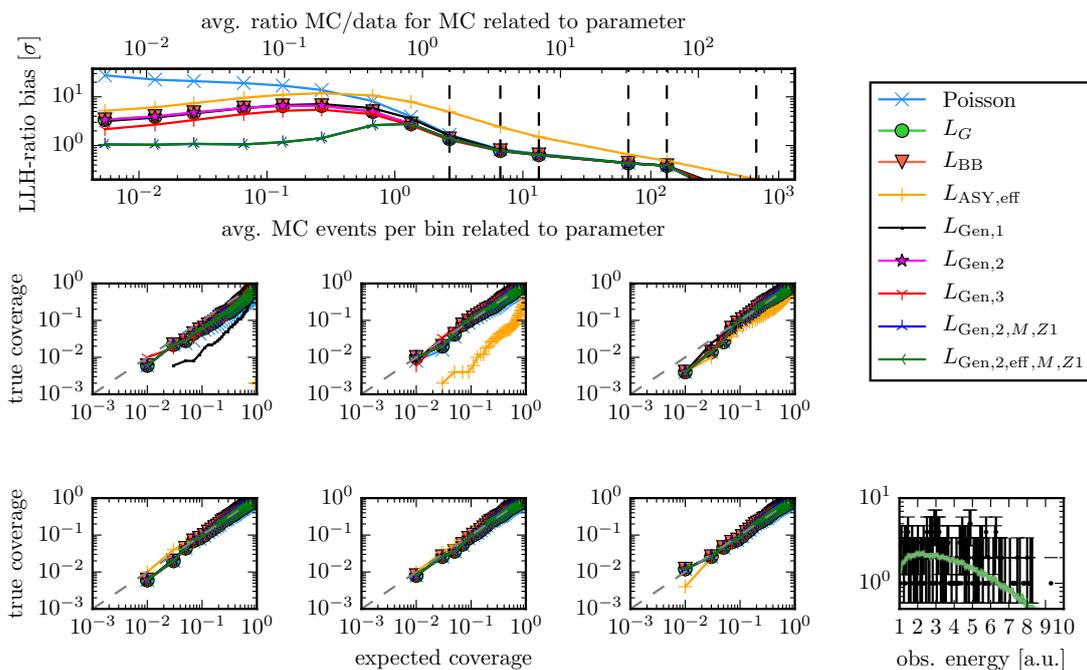


Figure 8. Bias of the log-likelihood ratio λ and coverage in dependence of the background dataset statistics. The signal is not present. The parameter of interest is the normalization of the background dataset. Z1 denotes strategy 1 for “zero MC”, Z2 strategy 2 for “zero MC”, U denotes more “unbiased” weight variance estimation and M denotes shifted mean based on average MC events in all bins. The coverage is shown at 6 statistic levels (in order) which are indicated as dashed vertical lines in the upper plot. In the lower right plot a particular data realization and a MC realization at the highest statistic level is shown for the observable space. The non-standard likelihoods from other publications include $L_{ASU,eff}$ [13], L_{BB} [11], L_C [12] and L_G [7].

of generalization (2) with mean adjustment, which shows the mean adjustment is nearly as good as using the exact Poisson mean. We can also see that filling up the empty bins with pseudo counts is necessary here. All likelihood approaches that do not fill up empty bins (no “ZB”) end up heavily biased even with infinite signal Monte Carlo, only because the background Monte Carlo stays limited. Additionally, usage of a more flexible weight variances estimator (“U”) can be helpful. For generalization (3), though, having a more unbiased estimator (“U”) seems to introduce a slight bias at high statistics, which we have not fully understood yet. Potentially the total variance of the estimators is too large here (compare also figure 12 in appendix D without “U”). Lastly, the effective version of generalization (2) does not seem to perform well in this scenario for very low statistics, but for higher statistics it approaches the standard formula of generalization (2). In figure 12, which we show in appendix D, we compare again generalization (2) against older approaches from the literature and also show the maximum likelihood estimator (MLE) bias and variance. We see again that all approaches from the literature fail when empty bins are neglected. Interestingly, the standard Poisson likelihood performs quite good in a low-statistic certain region, which has to do with very low MLE bias compared to all modifications.

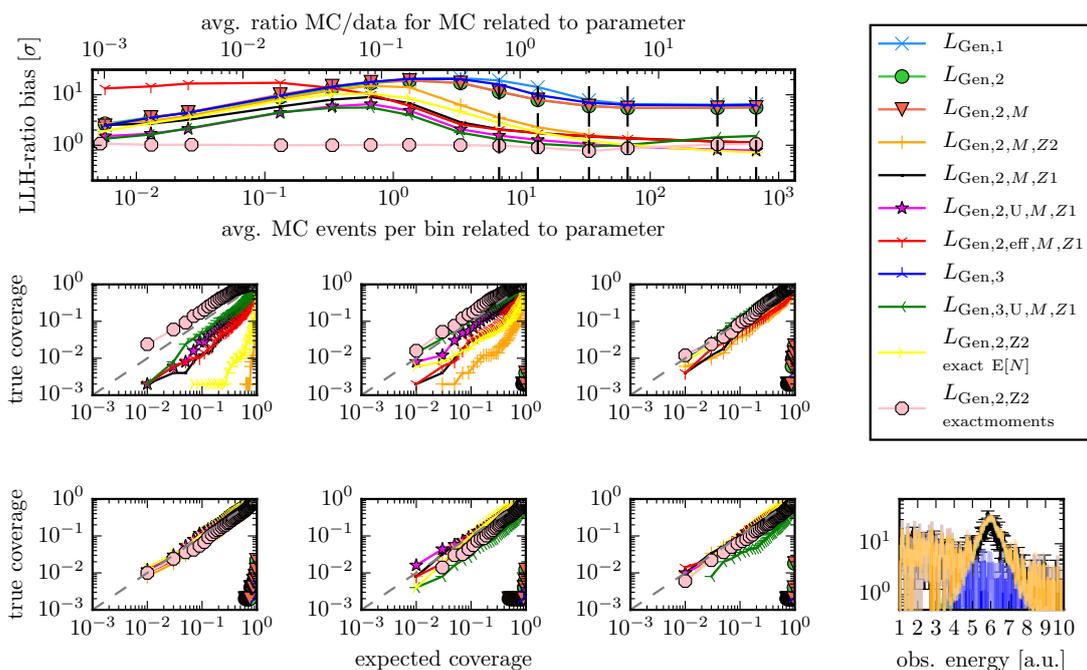


Figure 9. Bias of the log-likelihood ratio λ and coverage in dependence of the signal dataset statistics. The background statistics stays the same. The parameter of interest is the normalization of the peak in the signal dataset. Z1 denotes strategy 1 for “zero MC”, Z2 strategy 2 for “zero MC”, U denotes more “unbiased” weight variance estimation and M denotes shifted mean based on average MC events in all bins. The coverage is shown at 6 statistic levels (in order) which are indicated as dashed vertical lines in the upper plot. In the lower right plot a particular data realization and a MC realization at the highest statistic level is shown for the observable space.

6.1.3 Other scenarios

We also tested several other scenarios, whose results we show in appendix D. One scenario deals with more statistics and a smaller relative peak where we fit the peak normalization. The results are shown in figure 13. This situation is more akin to the scenario that was studied by the authors in [13] (“ $L_{ASY,eff}$ ”). We see that our likelihoods with mean adjustment perform better, in particular the effective version of generalization (2) (“ $L_{Gen2,eff}$ ”), which shows that a different parametrization is really crucial. In yet another scenario we increase the livetime of background and signal simultaneously and fit the spectral index (see figure 14). Again, the results are similar.

6.2 A new MC set every minimization step

In certain applications we might simulate a new set of MC events for every successive parameter call. This is a usual technique in IceCube to determine the ice properties or to perform expensive particle reconstructions [6]. Even when the parameters are the same, a new set of MC events will have different weights and a different simulated count per bin (see section 1 figure 1). In this situation, the uncertainty is directly visible in fluctuations of the corresponding likelihood function. In figure 10 we compare various versions of generalization (1) and generalization (3), which directly model this uncertainty within the PDF, for a likelihood scan of the normalization of the background dataset.

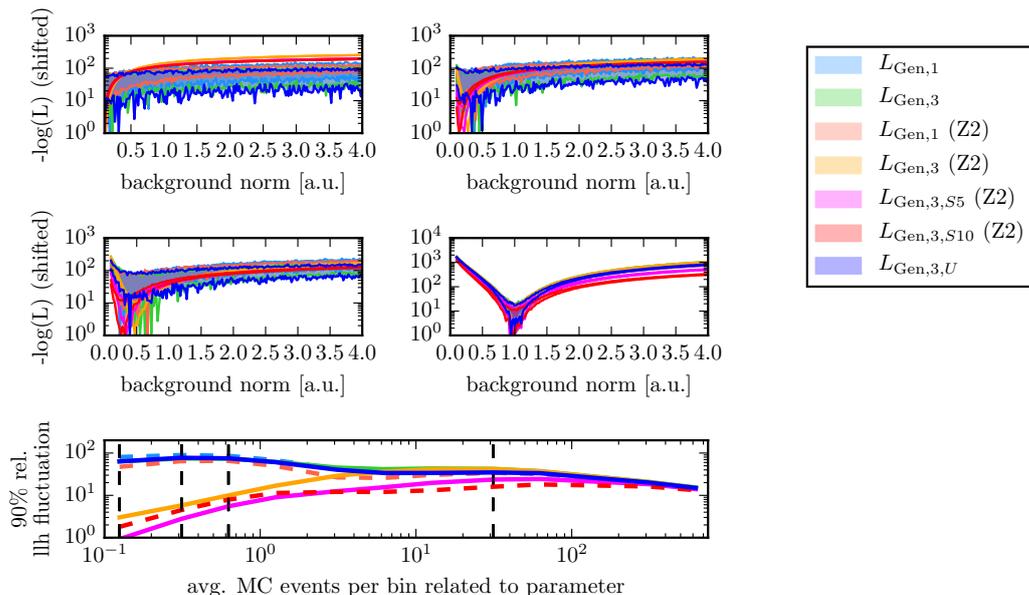


Figure 10. Stochastic scans of the likelihood functions at four different statistic levels (top). Shown is the central 90% interval of 50 realizations, shifted in absolute value for visualization so one can directly compare their shapes. In the bottom part we show the average of such likelihood fluctuations in dependence of the simulated Monte Carlo statistics. Vertical lines indicate which statistical levels are shown in the four plots above. *Z2* denotes strategy 2 for “zero MC”, *U* denotes more “unbiased” weight variance. Different parameters values of *S* correspond to slightly different behavior of generalization (3).

In this scenario all weights are the same. We test different parameter values of the scaling parameter *S* from generalization (3) (see section 4.3). A larger scaling parameter leads to a larger uncertainty being considered in the likelihood function, and correspondingly the fluctuations go down. However, we can observe that the likelihood bias goes up. A useful strategy might therefore be to start with a large parameter value of *S* and decrease it during the optimization process. Similarly, we could use generalization (1) and modify the parameter ν (see section 4.1) and introduce another scaling factor for the respective shape parameters of the Poisson-gamma mixtures. This should have a similar effective behavior. Additionally, we see that filling up the empty bins (*Z2*) is a second ingredient which is important. In figure 11 we compare the best performing generalization (3) with previous approaches. Here we observe that it is also possible to lower relative fluctuations by removing constant parts from the likelihood function which lowers the absolute values, and thereby also the relative fluctuations. The approaches described in [6, 7, 11] yield exactly similar fluctuations in this case. However, they can not compete with a direct uncertainty modeling.

7 Conclusion

The limited statistics of Monte Carlo simulation leads to extra uncertainty that has to be taken into account for correct parameter estimation. In this paper, we discussed several solutions from the literature that modify the Poisson likelihood to tackle this problem. We showed that all the approaches of this type can be interpreted as approximating the underlying probability distribution

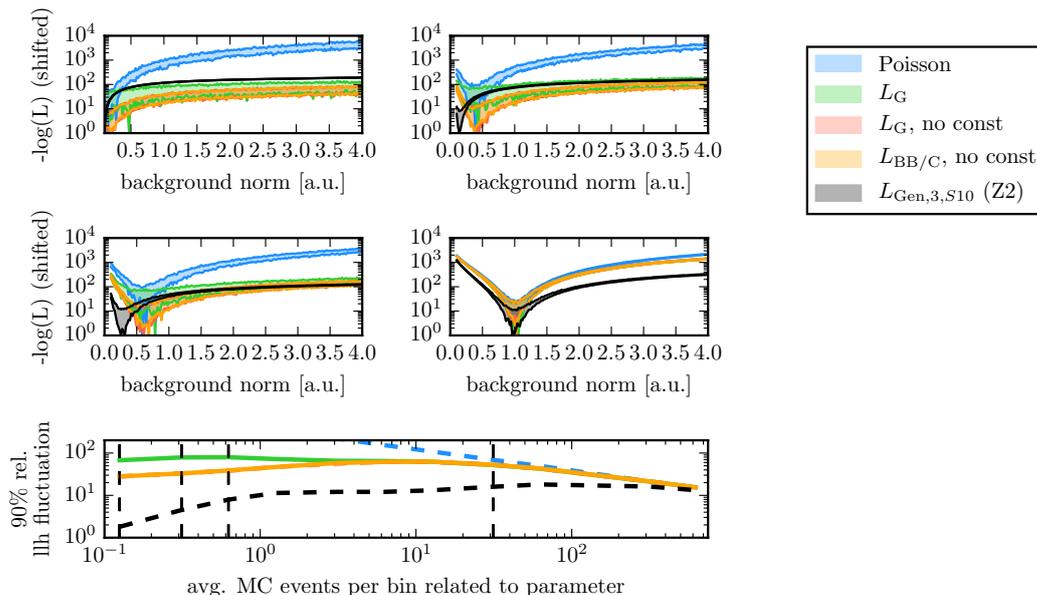


Figure 11. Stochastic scans of the likelihood functions at four different statistic levels (top). Shown is the central 90% interval of 50 realizations, shifted in absolute value for visualization so one can directly compare their shapes. In the bottom part we show the average of such likelihood fluctuations in dependence of the simulated Monte Carlo statistics. Vertical lines indicate which statistical levels are shown in the four plots above. Z2 denotes strategy 2 for “zero MC”. Different parameters of s relate to parameterizations of generalization (3). The older formulas are the approaches from [11, 12] (“ $L_{BB/C}$ ”), which are similar for equal weights, and from [7] (“ L_G ”).

of the sum of weights (the CPD) and then use this approximation to solve a marginalization integral over the Poisson mean. In probabilistic approaches, the integral is calculated exactly. In frequentist approaches, the integrand is optimized over the mean, not marginalized. This is an alternative viewpoint to the prevailing interpretation of frequentist approaches ([11, 12]), which is that these simply treat the Monte Carlo data similar to real data in the likelihood function.

In the new perspective, the central question is really how to approximate the mean μ_N of the Poisson part of the CPD, as well as the mean μ_W and variance $\text{var}[W]$ of the continuous “weight part” of the CPD. We introduced three new generalizations of existing probabilistic approaches that are well-motivated in this context. In generalization (1), we marginalize each gamma distribution from the probabilistic Ansatz in [7] with another Poisson-gamma mixture distribution. For equal weights, this distribution involves a marginalization of a Poisson-gamma mixture compound distribution, which can be viewed as a generalization of the exact CPD with more variance. In generalization (2), we model the CPD of each dataset independently with a gamma distribution, which are then convolved to form the total CPD. It can be viewed as a generalization of the approach by [13] in combination with the convolutional formula introduced in [7]. It is also the direct probabilistic counterpart to the approach by Barlow et al. [11]. Each gamma distribution is parametrized by the relevant CPD moments μ_N , μ_W , $\text{var}[W]$ of a given dataset, which are usually taken to be the sample moments, but they don’t have to be. We also discuss an effective version of generalization (2) which encodes a single CPD for all datasets, similar to [13], but with a different parametrization that allows

to incorporate extra Prior information in a more natural way. In generalization (3), we marginalize the unknown mean μ_N of each CPD from generalization (2) with yet another gamma distribution to incorporate extra sampling uncertainty directly into the PDF. All of the new probability distributions are analytically solvable.

In the last part of the paper we test these new formulas in a Toy MC parameter estimation setting with two datasets, a background dataset and signal dataset. We find that generalization (1) performs quite poor in these tests in terms of coverage and likelihood ratio bias. We also surprisingly find that the exact CPD performs worse than a continuous approximation with a gamma distribution, especially at very low statistics. We suspect that this has to do with the multi-modal form of the exact CPD. This explains why generalization (1) performs so poor for correct coverage, because it is just a slightly more general version of the exact CPD in the equal weights setting and has similar multimodal structure. Generalizations (2) and (3), which are based on continuous approximations of the overall CPD, perform better in all coverage tests.

If the moments are known exactly, a continuous approximation of the CPD as it is done in generalization (2) is in a certain sense optimal: it never undercovers at *any* level of Monte Carlo livetime. Achieving good coverage therefore reduces to finding good approximations to the moments μ_N , μ_W and $\text{var}[W]$ — the better these approximations, the better the coverage properties. Here, the usual sample moments $\mu_N = k_{\text{mc}}$, $\mu_W = \frac{1}{k_{\text{mc}}} \sum_i w_i$, $\text{var}[W] = \frac{1}{k_{\text{mc}}} \sum_i w_i^2$ are often biased which leads to undercoverage.

We can remedy this problem to some extent by incorporating extra Prior information about the Monte Carlo into these moments. One of these is a “mean adjustment” based on the average Monte Carlo count in all bins. The mean, and thereby also the mean adjustment, can be interpreted as a hyperparameter similar to the parameter α in the effective likelihood discussed in [13], and we show how they are precisely related. In contrast to the choice of an effective likelihood with $\alpha = 1$ discussed in [13], the mean adjustment is conceptually easier to interpret and gives better results in all test scenarios.

We also discuss a certain strategy how to treat bins with zero Monte Carlo events for a given dataset. In contrast to previous suggestions from the literature (see [11]) how to treat such scenarios, we find that a good choice is to fill up bins with pseudo counts such that every dataset has at least a single event in every bin, indifferent to the relative weights between datasets. This has the advantage that no jumps in the likelihood can appear.

In a final test we studied the scenario of new simulation Monte Carlo simulation every parameter step, as it happens for example in the application of direct Cerenkov photon simulation in the Icecube observatory [6]. Here, we observed that generalization (1) and generalization (3), which take into account the Poisson uncertainty from repeated sampling steps, greatly reduce fluctuations in successive likelihood evaluations compared to all other approaches. Given that generalization (3) has better properties in terms of log-likelihood ratio bias, it seems to be better suited than generalization (1) in this use case.

Our overall recommendation for parameter estimation situations can be summarized as follows:

1. Normal use case of weighted simulation: use generalization (2) with strategy 1 for empty bins and mean adjustment. If the computational cost is too high for the application, use the effective version of generalization (2).

2. Re-simulation every parameter query: use generalization (3) with strategy 2 for empty bins and a scaling parameter $s > 1$ to minimize fluctuations. Potentially decrease the scaling parameter during the minimization process to unity to decrease bias.

Besides the tests with exact moments using perfect knowledge from “infinite Monte Carlo”, it was never possible to reduce the bias of the log-likelihood ratio completely at all levels of simulated livetime. Blindly trusting a formula can therefore lead to trust in a systematically wrong result if log-likelihood ratio values are directly translated into significances. The way out are ensemble-based test statistic constructions. These might again be problematic because Monte Carlo is expensive to produce. If ensemble-based constructions are possible, however, we recommend to always use the new formulas because they capture the underlying statistics of the CPD. In general, it is always a good idea to calculate the log-likelihood ratio with multiple likelihood functions, at least also the standard Poisson likelihood, and conclude from the differences between the different approaches if enough Monte Carlo statistics have been produced. Probabilistic approaches allow to additionally compare absolute likelihood values. If those absolute values disagree substantially, for example between the Poisson likelihood and the modified likelihood in use, it is a strong hint that one must be careful on the interpretation of the result. Implementations of the new formulas can be found on http://www.github.com/thoglu/mc_uncertainty.

Acknowledgments

We would like to thank Carlos Argüelles, Marty Cohen, Austin Schneider and Tianlu Yuan for stimulating discussions.

A Derivation generalization (1)

This section computes the generalization described in section 4.1.

A.1 Equal weights

We start with equal MC weights for simulated events in a given bin. In this case, the generalized per-bin likelihood function is given by suitable marginalization as [7]

$$\begin{aligned}
 L_{\text{bin,eq.}} &= \mathbb{E} \left[\frac{e^{-\lambda} \lambda^k}{k!} \right]_{G(\lambda; k_{\text{mc}} + \alpha, 1/w)} \\
 &= \frac{\Gamma(k + k_{\text{mc}} + \alpha)}{\Gamma(k_{\text{mc}} + \alpha) \cdot k!} \cdot \frac{(1/w)^{k_{\text{mc}} + \alpha}}{(1 + 1/w)^{k + k_{\text{mc}} + \alpha}} \tag{A.1} \\
 &= \text{PG}(k; k_{\text{mc}} + \alpha, 1/w) \tag{A.2}
 \end{aligned}$$

where the result is a standard gamma-Poisson mixture $\text{PG}(k; k_{\text{mc}} + \alpha, 1/w)$ distribution. For more information on the parameter α , which can be seen as a hyperparameter parametrizing a gamma prior, in particular its special value $\alpha = 0$, see [7]. Let us now extend this construction by further

marginalizing out the discrete counts k_{mc} via

$$\begin{aligned} L_{\text{bin,(1),eq.}} &= \text{E} \left[L_{\text{bin,eq.}} \right]_{\text{PG}(k_{\text{mc}}; \langle k_{\text{mc}} \rangle + \beta, 1/\nu)} \\ &= \sum_{k_{\text{mc}}=0}^{\infty} \text{PG}(k; k_{\text{mc}} + \alpha, 1/w) \cdot \text{PG}(k_{\text{mc}}; \langle k_{\text{mc}} \rangle + \beta, 1/\nu) \end{aligned} \quad (\text{A.3})$$

$$\equiv \text{PGPG}(k; \alpha; 1/w; \langle k_{\text{mc}} \rangle + \beta, 1/\nu) \quad (\text{A.4})$$

and obtain a new likelihood function $L_{\text{P,finite,twice,eq.}}$ that has been marginalized twice, whose probability distribution we abbreviate as PGPG, because we marginalize a Poisson-gamma mixture with another Poisson-gamma mixture. This second marginalization is discrete, and happens with another gamma-Poisson mixture $\text{PG}(k_{\text{mc}}; \langle k_{\text{mc}} \rangle + \beta, 1/\nu)$ which is similar in structure to the Poisson mixture in eq. (A.1), but with different parameters.⁶ We will later set the ‘‘count weight’’ ν to $\nu = 1$, but for now keep it in this general form. A weight of unity makes sense since we are now interested in marginalizing pure counts, which are unweighted by construction. Another nuisance parameter β is introduced to handle Prior freedom, similar to the introduction of α in [7]. In eq. (A.2) the variable k_{mc} denotes the observed number of MC events, while in the marginalized expression (eq. (A.3)) the variable that is integrated out is k_{mc} and the observed MC counts as $\langle k_{\text{mc}} \rangle$. Equation (A.3) can be solved using Egorychev rules ([17], pg. 19) via

$$L_{\text{bin,(1),eq.}} = \sum_{k_{\text{mc}}=0}^{\infty} \text{PG}(k; k_{\text{mc}} + \alpha, 1/w) \cdot \text{PG}(k_{\text{mc}}; \langle k_{\text{mc}} \rangle + \beta, 1/\nu) \quad (\text{A.5})$$

$$= \sum_{k_{\text{mc}}=0}^{\infty} \frac{\Gamma(k + k_{\text{mc}} + \alpha)}{\Gamma(k_{\text{mc}} + \alpha) \cdot k!} \cdot \frac{(1/w)^{k_{\text{mc}} + \alpha}}{(1 + 1/w)^{k + k_{\text{mc}} + \alpha}} \frac{\Gamma(k_{\text{mc}} + \langle k_{\text{mc}} \rangle + \beta)}{\Gamma(\langle k_{\text{mc}} \rangle + \beta) \cdot k_{\text{mc}}!} \cdot \frac{(1/\nu)^{\langle k_{\text{mc}} \rangle + \beta}}{(1 + 1/\nu)^{k_{\text{mc}} + \langle k_{\text{mc}} \rangle + \beta}} \quad (\text{A.6})$$

$$= \underbrace{\left(\frac{1}{1 + 1/w} \right)^k \left(\frac{1}{1 + w} \right)^\alpha \cdot \left(\frac{1}{1 + \nu} \right)^{\langle k_{\text{mc}} \rangle + \beta}}_{\equiv K} \quad (\text{A.7})$$

$$\cdot \sum_{k_{\text{mc}}=0}^{\infty} \underbrace{\left(\frac{\nu}{(1 + \nu)(1 + w)} \right)^{k_{\text{mc}}}}_{\equiv C} \frac{1}{2\pi i} \oint_{|z_1|=\epsilon} \frac{(1 + z_1)^{k + k_{\text{mc}} + \alpha - 1}}{z_1^{k+1}} dz_1 \frac{1}{2\pi i} \oint_{|z_2|=\epsilon} \frac{(1 + z_2)^{k_{\text{mc}} + \langle k_{\text{mc}} \rangle + \beta - 1}}{z_2^{k_{\text{mc}} + 1}} dz_2$$

$$= K \cdot \frac{1}{2\pi i} \oint_{|z_1|=\epsilon} \frac{(1 + z_1)^{k + \alpha - 1}}{z_1^{k+1}} \cdot \sum_{k_{\text{mc}}=0}^{\infty} C^{k_{\text{mc}}} \cdot \frac{1}{2\pi i} \oint_{|z_2|=\epsilon} (1 + z_2)^{\langle k_{\text{mc}} \rangle + \beta - 1} \cdot \left[(1 + z_1)(1 + z_2) \right]^{k_{\text{mc}}} \cdot z_2^{-k_{\text{mc}} - 1} dz_2 dz_1 \quad (\text{A.8})$$

$$= K \cdot \frac{1}{2\pi i} \oint_{|z_1|=\epsilon} \frac{(1 + z_1)^{k + \alpha - 1}}{z_1^{k+1}} \cdot \underbrace{\left(\frac{1}{1 - C(1 + z_1)} \right)^{\langle k_{\text{mc}} \rangle + \beta}}_{\equiv g(z_1)} dz_1 = K \cdot \frac{1}{2\pi i} \oint_{|z_1|=\epsilon} \frac{f(z_1)}{z_1^{k+1}} \cdot g(z_1) dz_1 \quad (\text{A.9})$$

⁶We could have used a standard Poisson distribution for marginalization instead of a Poisson-gamma mixture. However, it turns out that such an Ansatz gives a more complicated result for general weights which are discussed in the next section. In particular, essential singularities appear which make the end result harder to compute.

Using the Taylor expansion of $f(z_1)$ and $g(z_1)$ around $z_1 = 0$, we can combinatorially combine the relevant terms to contribute to the residue of the contour integral at $z_1 = 0$. The result is then readily written down as

$$L_{\text{bin,(1),eq.}} = K \cdot \sum_{n=0}^k \frac{1}{n!} f^{(n)}(0) \cdot \frac{1}{(k-n)!} g^{(k-n)}(0) \quad (\text{A.10})$$

$$= K \cdot \left(\frac{1}{1-C} \right)^{\langle k_{\text{mc}} \rangle + \beta} \cdot \left(\frac{C}{1-C} \right)^k \sum_{n=0}^k \frac{\Gamma(k+\alpha)}{n! \Gamma(k-n+\alpha)} \frac{\Gamma(\langle k_{\text{mc}} \rangle + \beta + k - n)}{(k-n)! \Gamma(\langle k_{\text{mc}} \rangle + \beta)} \cdot \left(\frac{1-C}{C} \right)^n \quad (\text{A.11})$$

where the constants K and C are defined in eq. (A.7) as $K = \left(\frac{1}{1+1/w} \right)^k \left(\frac{1}{1+w} \right)^\alpha \cdot \left(\frac{1}{1+v} \right)^{\langle k_{\text{mc}} \rangle + \beta}$ and $C = \frac{v}{(1+v)(1+w)}$.

A.2 General weights

For general weights we start with eq. (4.2) and rewrite it as

$$L_{\text{bin,(1)}} = \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot [\text{GPG}_1(\lambda_1) * \dots * \text{GPG}_N(\lambda_N)](\lambda) d\lambda \quad (\text{A.12})$$

$$= \sum_{k_{\text{mc},1}=0}^\infty \dots \sum_{k_{\text{mc},M}=0}^\infty \text{PG}_1 \cdot \text{PG}_M \cdot \int_0^\infty \frac{e^{-\lambda} \lambda^k}{k!} \cdot [\text{G}_1(\lambda_1) * \dots * \text{G}_N(\lambda_N)](\lambda) d\lambda \quad (\text{A.13})$$

$$= \sum_{k_{\text{mc},1}=0}^\infty \dots \sum_{k_{\text{mc},M}=0}^\infty \text{PG}_1 \cdot \text{PG}_M \cdot \sum_{\substack{k_1+\dots+k_M=k \\ k_j \geq 0}} \prod_{j=1}^M P_{\text{PG}}(k_j; k_{\text{mc},j} + \alpha_j^*, 1/w_j) \quad (\text{A.14})$$

$$= \sum_{\substack{k_1+\dots+k_M=k \\ k_j \geq 0}} \prod_{j=1}^M \sum_{k_{\text{mc},j}=0}^\infty P_{\text{PG}}(k_j; k_{\text{mc},j} + \alpha_j^*, 1/w_j) \cdot P_{\text{PG},i}(k_{\text{mc},j}; \langle k_{\text{mc}} \rangle_j + \beta_j^*, 1/v_j) \quad (\text{A.15})$$

where we first pull out all the infinite sums, then rewrite the integral over the convolution of gamma factors with a Poisson factor in combinatorial form (see appendix C) and then again combine each summation with the respective terms. The result is a combinatorial sum over PGPG distributions, each of which we calculated the result in the previous section. We can continue and replace every PGPG distribution in eq. (A.15) by its respective contour integral (eq. (A.9)), which results in

$$L_{\text{bin,(1)}} = \sum_{\substack{k_1+\dots+k_M=k \\ k_j \geq 0}} \prod_{j=1}^M \sum_{k_{\text{mc},j}=0}^\infty P_{\text{PG}}(k_j; k_{\text{mc},j} + \alpha_j^*, 1/w_j) \cdot P_{\text{PG},i}(k_{\text{mc},j}; \langle k_{\text{mc}} \rangle_j + \beta_j^*, 1/v_j) \quad (\text{A.16})$$

$$= \sum_{\substack{k_1+\dots+k_M=k \\ k_j \geq 0}} \prod_{j=1}^M K_j \cdot \frac{1}{2\pi i} \oint_{|z_j|=\epsilon} \frac{(1+z_j)^{k_j+\alpha_j^*-1}}{z_j^{k_j+1}} \cdot \left(\frac{1}{1-C_j(1+z_j)} \right)^{\langle k_{\text{mc}} \rangle_j + \beta_j^*} dz_j \quad (\text{A.17})$$

where we use $\alpha_j^* = \frac{\langle k_{\text{mc},j} \rangle \cdot \alpha}{N}$ and $\beta_j^* = \frac{\langle k_{\text{mc},j} \rangle \cdot \beta}{N}$ to share the nuisance parameter among groups of events with different weights and the definition of K_j and C_j is the same as K and C for given a weight group in the previous section. Again we use $\langle k_{\text{mc}} \rangle$ to denote observed number of MC events,

and k_{mc} as a variable to sum over. We can now solve the constrained sum over k_j using Egorychev's rules for constrained summation indices [17]. The constraint here is $\sum k_j = k$, which is encoded with complex weight variables τ of equal magnitude. Equation (A.17) with these weight variables τ and an extra outer contour integral then can be re-written as

$$L_{\text{bin},(1)} = \oint_{|\tau|=\epsilon} \frac{1}{\tau^{k+1}} \sum_{k_1=0 \dots k_M=0}^{\infty \dots \infty} \prod_{j=1}^M K_j \cdot \frac{1}{2\pi i} \oint_{|z_j|=\epsilon} \frac{(1+z_j \cdot \tau)^{k_j+\alpha_j^*-1}}{z_j^{k_j+1}} \cdot \left(\frac{1}{1-C_j(1+z_j \cdot \tau)} \right)^{\langle k_{mc} \rangle_j + \beta_j^*} dz_j \quad (\text{A.18})$$

$$= \oint_{|\tau|=\epsilon} \frac{1}{\tau^{k+1}} \sum_{k_1=0 \dots k_M=0}^{\infty \dots \infty} \prod_{j=1}^M K'_j \cdot \left(\frac{1}{1+1/w_j} \right)^{k_j} \cdot \frac{1}{2\pi i} \oint_{|z_j|=\epsilon} \underbrace{\left(\frac{1}{1-C_j(1+z_j \cdot \tau)} \right)^{\langle k_{mc} \rangle_j + \beta_j^*}}_{\varphi} \cdot \underbrace{(1+z_j \cdot \tau)^{\alpha_j^*-1} \cdot (1+z_j \cdot \tau)^{k_j} \cdot z_j^{-k_j-1}}_f dz_j d\tau \quad (\text{A.19})$$

$$= \oint_{|\tau|=\epsilon} \frac{1}{\tau^{k+1}} \prod_{j=1}^M K'_j \cdot \frac{(1-E_j \cdot \tau)^{\langle k_{mc} \rangle_j + \beta_j^* - \alpha_j^*}}{(1-C_j-E_j \cdot \tau)^{\langle k_{mc} \rangle_j + \beta_j^*}} d\tau \quad (\text{A.20})$$

$$= \left(\prod_{j=1}^M K'_j \cdot (-E_j)^{-\alpha_j^*} \right) \oint_{|\tau|=\epsilon} \frac{1}{\tau^{k+1}} \prod_{j=1}^M \frac{(\tau-1/E_j)^{\langle k_{mc} \rangle_j + \beta_j^* - \alpha_j^*}}{(\tau-(1-C_j)/E_j)^{\langle k_{mc} \rangle_j + \beta_j^*}} d\tau \quad (\text{A.21})$$

$$= \left(\prod_{j=1}^M K'_j \cdot \left(\frac{1}{1-C_j} \right)^{\langle k_{mc} \rangle_j + \beta_j} \right) \cdot \Delta_k \quad (\text{A.22})$$

$$\stackrel{\substack{\alpha=0 \\ \beta=0 \\ \nu_j=1}}{=} \left(\prod_{j=1}^M \left(\frac{1+w_j}{1+2w_j} \right)^{\langle k_{mc} \rangle_j} \right) \cdot \Delta_k \quad (\text{A.23})$$

using $K_j \equiv K'_j \cdot E_j^{k_j}$, $E_j = \frac{1}{1+1/w_j}$, $C_j = \frac{\nu_j}{(1+\nu_j)(1+w_j)}$ and $K'_j = \left(\frac{1}{1+w_j} \right)^{\alpha^*} \cdot \left(\frac{1}{1+\nu_j} \right)^{\langle k_{mc} \rangle_j + \beta_j^*}$ from the definitions in the previous section. The terms we call f and φ are highlighted to indicate the structure used for the main theorem described in [17] (pg. 19, 1.4.2). The term Δ_k represents an efficient solution [18] of the contour integral in eq. (A.21), and is written as

$$\Delta_k = \frac{1}{k} \sum_{i=1}^k \left[\left(\sum_{j=1}^M (\langle k_{mc,j} \rangle + \beta_j^*) \cdot \left(\frac{E_j}{1-C_j} \right)^i - (\langle k_{mc,j} \rangle + \beta_j^* - \alpha_j^*) \cdot E_j^i \right) \Delta_{k-i} \right] \quad (\text{A.24})$$

$$= \frac{1}{k} \sum_{i=1}^k \left[\left(\sum_{j=1}^M (\langle k_{mc,j} \rangle + \beta_j^*) \cdot \left(\frac{w_j(1+\nu_j)}{1+w_j(1+\nu_j)} \right)^i - (\langle k_{mc,j} \rangle + \beta_j^* - \alpha_j^*) \cdot \left(\frac{1}{1+1/w_j} \right)^i \right) \Delta_{k-i} \right] \quad (\text{A.25})$$

$$\stackrel{\substack{\alpha=0 \\ \beta=0 \\ \nu_j=1}}{=} \frac{1}{k} \sum_{i=1}^k \left[\left(\sum_{j=1}^M \langle k_{mc,j} \rangle \cdot \left(\frac{2}{2+1/w_j} \right)^i - \langle k_{mc,j} \rangle \cdot \left(\frac{1}{1+1/w_j} \right)^i \right) \Delta_{k-i} \right] \quad (\text{A.26})$$

with $\Delta_0 = 1$. One can also write the sum going over all individual events N instead over M groups of events that share a similar weight, which results in $k_{mc,j} = 1$ for each summand given by the last equal sign. The nuisance parameters $\alpha_j^* = \frac{\langle k_{mc,j} \rangle \cdot \alpha}{N}$ and $\beta_j^* = \frac{\langle k_{mc,j} \rangle \cdot \beta}{N}$ are typically set to zero since $\alpha = \beta = 0$ is a natural prior choice (see also [7]). The parameter v_j is further set to unity, because it came about as the scaling parameter of the gamma distribution in the Poisson gamma mixture describing pure counts. The general case of convolutions of arbitrarily mixed gamma distributions and gamma-Poisson-gamma mixtures is described in appendix C.

B Derivation generalization (3)

B.1 Single source dataset

For a single source dataset we have a single gamma factor that approximates the CPD similar to the description in section 2.3.2. The corresponding integral looks like

$$\text{PGG}(k) = \int_0^\infty \int_0^\infty \frac{e^{-k} \cdot \lambda^k}{k!} G(\lambda; \mu \cdot Q, \beta) \cdot G(\mu; k_{mc}, \gamma) d\lambda d\mu = \int_0^\infty \text{PG}(k; \mu \cdot Q, \beta) \cdot G(\mu; k_{mc}, \gamma) d\mu \quad (\text{B.1})$$

$$= \int_0^\infty \frac{\Gamma(\mu \cdot Q + k)}{\Gamma(\mu \cdot Q)k!} \cdot \underbrace{\left(\frac{\beta}{1+\beta}\right)^{\mu \cdot Q}}_{\equiv A} \cdot \underbrace{\left(\frac{1}{1+\beta}\right)^k}_{\equiv B} \cdot G(\mu; k_{mc}, \gamma) d\mu \quad (\text{B.2})$$

where the first gamma factor in eq. (B.1) represents the CPD approximation of the source dataset with $Q = \frac{1}{k_{mc}} \frac{(\sum_j w_j)^2}{\sum_j w_j^2}$, $\beta = \frac{\sum_j w_j}{\sum_j w_j^2}$, and the second gamma factor marginalizes over the unknown rate where we leave a scaling parameter γ for generality. We will later set $\gamma = 1$. Additionally we define $A = \frac{\beta}{1+\beta}$ and $B = \frac{1}{1+\beta}$. We can also use the fact that $\frac{\Gamma(x+k)}{\Gamma(x)} = \sum_{n=0}^k \binom{k}{n} x^n$ [19] where $\binom{k}{n}$ are the unsigned Stirling numbers of the first kind and reformulate eq. (B.2) as

$$\text{PGG}(k) = \int_0^\infty \frac{\Gamma(\mu \cdot Q + k)}{\Gamma(\mu \cdot Q)k!} \cdot \left(\frac{\beta}{1+\beta}\right)^{\mu \cdot Q} \cdot \left(\frac{1}{1+\beta}\right)^k \cdot G(\mu; k_{mc}, \gamma) d\mu \quad (\text{B.3})$$

$$= \int_0^\infty \frac{1}{k!} \sum_{n=0}^k \binom{k}{n} (\mu Q)^n \cdot A^{\mu \cdot Q} \cdot B^k \cdot \frac{e^{-\gamma \cdot \mu} \cdot \mu^{k_{mc}-1}}{\Gamma(k_{mc})} \cdot \gamma^{k_{mc}} d\mu \quad (\text{B.4})$$

$$= B^k \cdot \gamma^{k_{mc}} \sum_{n=0}^k \frac{1}{k!} \cdot Q^n \binom{k}{n} \frac{\Gamma(k_{mc} + n)}{\Gamma(k_{mc})(\gamma - Q \ln(A))^{k_{mc}+n}} \times \int_0^\infty \frac{e^{-\mu(\gamma - Q \ln(A))} \cdot \mu^{k_{mc}+n-1}}{\Gamma(k_{mc} + n)} (\gamma - Q \ln(A))^{k_{mc}+n} d\mu \quad (\text{B.5})$$

$$= B^k \cdot \gamma^{k_{mc}} \sum_{n=0}^k \frac{1}{k!} \cdot Q^n \binom{k}{n} \frac{\Gamma(k_{mc} + n)}{\Gamma(k_{mc})(\gamma - Q \ln(A))^{k_{mc}+n}} \quad (\text{B.6})$$

C Overview of some generic formulas

The solutions of generalization (1) 4.1 can be written in various forms which we summarize in the following. We abbreviate the Poisson distribution as P, the gamma distribution as G, the gamma-Poisson mixture as GP, and the Negative-Binomial distribution as NB. When we marginalize the shape parameter of a gamma distribution with a GP distribution, we abbreviate that as GPG, and a marginalization of a GP distribution shape parameter with another GP distribution is abbreviated as GPGP.

C.1 Poisson-gamma mixture form

C.1.1 Simple generalization

The PMF of the generalized Poisson-gamma mixture has already derived in [7]. In [20] it was shown that the convolution of Poisson-gamma mixtures is actually describing the same distribution, and so we have

$$P_{\text{gen}}(k; \boldsymbol{\alpha}, \boldsymbol{\beta}) = [\text{PG}(k_1; \alpha_1, \beta_1) * \dots * \text{PG}(k_N; \alpha_N, \beta_N)](k) \quad (\text{C.1})$$

$$= \int_0^{\infty} \text{P}(k; \lambda) \cdot [\text{G}(\lambda_1; \alpha_1, \beta_1) * \dots * \text{G}(\lambda_N; \alpha_N, \beta_N)](\lambda) d\lambda \quad (\text{C.2})$$

$$= \sum_{\sum_i k_i = k, k_i \geq 0} \prod_i \frac{\Gamma(k_i + \alpha_i)}{k_i! \cdot \Gamma(\alpha_i)} \cdot \beta_i^{\alpha_i} \cdot \left(\frac{1}{1 + \beta_i} \right)^{k_i + \alpha_i} \quad (\text{C.3})$$

$$= \sum_{\sum_i k_i = k, k_i \geq 0} \prod_i \text{PG}(k_i; \alpha_i, \beta_i) \quad (\text{C.4})$$

$$= D_k(\boldsymbol{\alpha}, \boldsymbol{\beta}) \cdot \prod_i \left(\frac{\beta_i}{1 + \beta_i} \right)^{\alpha_i} \quad (\text{C.5})$$

with iterative definition

$$D_k(\boldsymbol{\alpha}, \boldsymbol{\beta}) = \frac{1}{k} \sum_{j=1}^k \left[\left(\sum_{i=1}^N \alpha_i \cdot \frac{1}{1 + \beta_i} \right)^j D_{k-j} \right] \quad \text{and } D_0 = 1. \quad (\text{C.6})$$

The combinatorial form (eq. (C.3)) has been derived in [7] and in [20] in the different context of a convolution of Poisson-gamma mixtures. However, the calculation based on the iterative sum (eq. (2.24)) is much more efficient. The convolution of Poisson-gamma mixtures indicated in eq. (C.1) is an important step in the efficient calculation derived in appendix B. This distribution is very rich in structure, and its relationship to the Carlson-R function [21] and Lauricella function F_D [22] have been explored in [7]. The iterative formula has substantial computational advantages in various applications of these special functions.

C.1.2 Extended generalization

In appendix A we derived a generalization of the extended Poisson gamma mixture by taking a further expectation value (eq. (A.16)) with respect another Poisson-gamma mixture. Here, we

proceed more generally and only take the expectation value of M Poisson-gamma terms while N standard terms remain. Starting with the combinatorial expression similar to eq. (A.16) we can write

$$P_{\text{gen}}(k; \alpha, \beta, \gamma, \delta, \varepsilon) = \sum_{\substack{\sum_i k_i + \sum_j k_j = k \\ k_i \geq 0, k_j \geq 0}} \prod_{i=1}^N \text{PG}(k_i; \alpha_i, \beta_i) \prod_{j=1}^M \sum_{t=0}^{\infty} \text{PG}(k_j; t, \gamma_j) \text{PG}(t; \delta_j, \varepsilon_j) \quad (\text{C.7})$$

$$= \sum_{\substack{\sum_i k_i + \sum_j k_j = k \\ k_i \geq 0, k_j \geq 0}} \prod_{i=1}^N \text{PG}(k_i; \alpha_i, \beta_i) \prod_{j=1}^M \text{PGPG}(k_j; \gamma_j, \delta_j, \varepsilon_j) \quad (\text{C.8})$$

$$= \int_0^{\infty} \text{P}(k; \lambda) \cdot \left[\text{G}(\lambda_1; \alpha_1, \beta_1) * \dots * \text{G}(\lambda_N; \alpha_N, \beta_N) \right. \\ \left. * \text{GPG}(\lambda_{1^*}; \gamma_1, \delta_1, \varepsilon_1) * \dots * \text{GPG}(\lambda_{M^*}; \gamma_M, \delta_M, \varepsilon_M) \right] (\lambda) d\lambda \quad (\text{C.9})$$

$$= \left[\text{PG}(k_1; \alpha_1, \beta_1) * \dots * \text{PG}(k_N; \alpha_N, \beta_N) \right. \\ \left. * \text{PGPG}(k_{1^*}; \gamma_1, \delta_1, \varepsilon_1) * \dots * \text{PGPG}(k_{M^*}; \gamma_M, \delta_M, \varepsilon_M) \right] (k) \quad (\text{C.10})$$

$$= \Delta_k(\alpha, \beta, \gamma, \delta, \varepsilon) \cdot \prod_{i=1}^N \left(\frac{\beta_i}{1 + \beta_i} \right)^{\alpha_i} \prod_{j=1}^M \left(\frac{\varepsilon_j \cdot (\gamma_j + 1)}{\varepsilon_j \cdot (\gamma_j + 1) + 1} \right)^{\delta_j} \quad (\text{C.11})$$

$$\Delta_k(\alpha, \beta, \gamma, \delta, \varepsilon) = \frac{1}{k} \sum_{j=1}^k \left[\left(\sum_{i=1}^N \alpha_i \cdot \left(\frac{1}{1 + \beta_i} \right)^j + \sum_{h=1}^M \delta_h \cdot \left[\left(\frac{1 + \varepsilon_h}{1 + \varepsilon_h \cdot (1 + \gamma_h)} \right)^j - \left(\frac{1}{1 + \gamma_h} \right)^j \right] \right) \Delta_{k-j} \right] \\ \text{and } \Delta_0 = 1 \quad (\text{C.12})$$

where we have introduced additional M -dimensional vectors γ , δ , ε that parametrize the additional M marginalized factors resulting in the PGPG or GPG distributions, depending how we write it. If we put $M = N$, $\delta_i = \langle k_{\text{mc},i} \rangle$, $\gamma_i = \beta_i = 1/w_i$ and $\varepsilon_i = 1$ we obtain eq. (A.23) as a special case.

C.2 Negative binomial form

Poisson-gamma mixtures are often re-parametrized as negative binomial distributions, which can be obtained by setting $\beta = \frac{1-p}{p}$ and using $\alpha = r$. For completion we give the simple and extended negative-binomial generalizations below.

C.2.1 Simple generalization

$$\text{NB}_{\text{simple}}(k; \mathbf{r}, \mathbf{p}) = \int_0^{\infty} \text{P}(k; \lambda) \cdot \left[\text{G}\left(\lambda_1; r_1, \frac{1-p_1}{p_1}\right) * \dots * \text{G}\left(\lambda_N; r_N, \frac{1-p_N}{p_N}\right) \right] (\lambda) d\lambda \quad (\text{C.13})$$

$$= D_k(\mathbf{r}, \mathbf{p}) \cdot \prod_i (1-p_i)^{r_i} \quad (\text{C.14})$$

with iterative definition

$$D_k(\mathbf{r}, \mathbf{p}) = \frac{1}{k} \sum_{j=1}^k \left[\left(\sum_{i=1}^N r_i \cdot p_i^j \right) D_{k-j} \right] \quad \text{and } D_0 = 1. \quad (\text{C.15})$$

If all r_i except one are zero, which effectively means there is only one pair of r_i and p_i , we obtain the standard negative binomial distribution as a special case.

C.2.2 Extended generalization

In order to obtain an extended negative binomial distribution that is sensible, we take eq. (C.7) and set $\beta_i = \frac{1-p_i}{p_i}$, $\gamma_i = \frac{1-p_i^*}{p_i^*}$, $\alpha_i = r_i$, $\delta_i = r_i^*$ and $\varepsilon_i = 1$. The result gives more variance to the gamma factors with p_i^* and r_i^* which are marginalized with an extra Poisson-gamma factor. Using unity for ε_i is the natural choice that is also used for the application in eq. (A.23). We obtain

$$\begin{aligned} \text{NB}_{\text{ext}}(k; \mathbf{r}, \mathbf{p}, \mathbf{r}^*, \mathbf{p}^*) &= \int_0^\infty \text{P}(k; \lambda) \cdot \left[\text{G}\left(\lambda_1; r_1, \frac{1-p_1}{p_1}\right) * \dots * \text{G}\left(\lambda_N; r_N, \frac{1-p_N}{p_N}\right) \right. \\ &\quad \left. * \text{GPG}\left(\lambda_1^*; \frac{1-p_1^*}{p_1^*}, r_1^*, 1\right) * \dots * \text{GPG}\left(\lambda_M^*; \frac{1-p_M^*}{p_M^*}, r_M^*, 1\right) \right] (\lambda) d\lambda \end{aligned} \quad (\text{C.16})$$

$$= \Delta_k(\mathbf{r}, \mathbf{p}, \mathbf{r}^*, \mathbf{p}^*) \cdot \prod_{i=1}^N (1-p_i)^{r_i} \prod_{j=1}^M \left(\frac{1}{1+p_j^*} \right)^{r_j^*} \quad (\text{C.17})$$

with iterative definition

$$\Delta_k(\mathbf{r}, \mathbf{p}, \mathbf{r}^*, \mathbf{p}^*) = \frac{1}{k} \sum_{j=1}^k \left[\left(\sum_{i=1}^N r_i \cdot p_i^j + \sum_{h=1}^M r_h^* \cdot \left[\left(\frac{p_h^*}{1+p_h^*} \right)^j - (p_h^*)^j \right] \right) \Delta_{k-j} \right] \quad \text{and } \Delta_0 = 1. \quad (\text{C.18})$$

D Supplementary figures

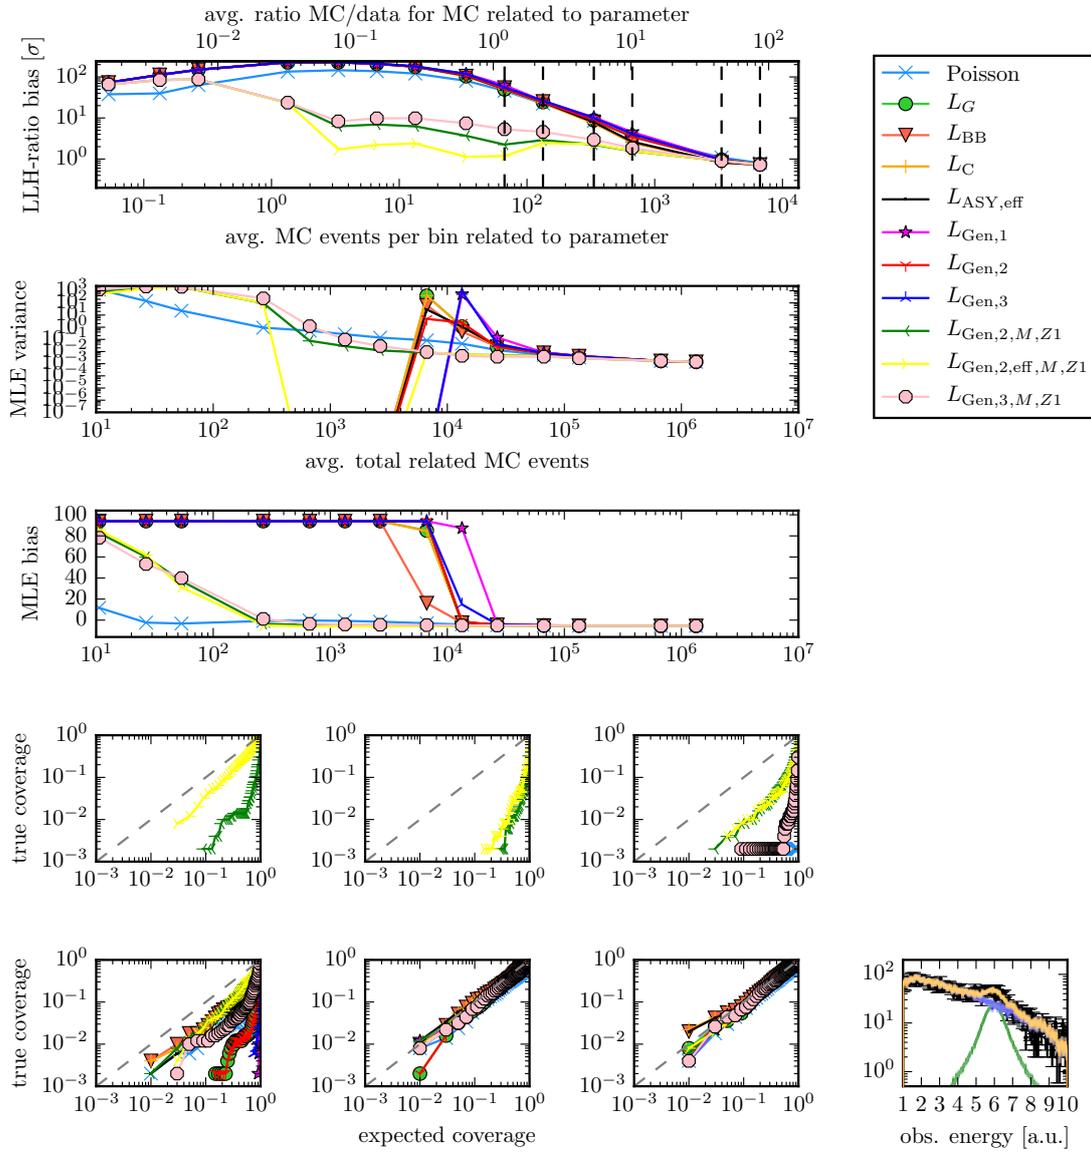


Figure 12. Bias of the log-likelihood ratio λ and coverage in dependence of the signal dataset statistics. The background statistics stays the same. The parameter of interest is the normalization of the peak in the signal dataset. Z1 denotes strategy 1 for “zero MC”, Z2 strategy 2 for “zero MC”, U denotes more “unbiased” weight variance estimation and M denotes shifted mean based on average MC events in all bins. The coverage is shown at 6 statistic levels (in order) which are indicated as dashed vertical lines in the upper plot. In the lower right plot a particular data realization and a MC realization at the highest statistic level is shown for the observable space. The non-standard likelihoods from other publications include $L_{ASY,eff}$ [13], L_{BB} [11], L_C [12] and L_G [7].

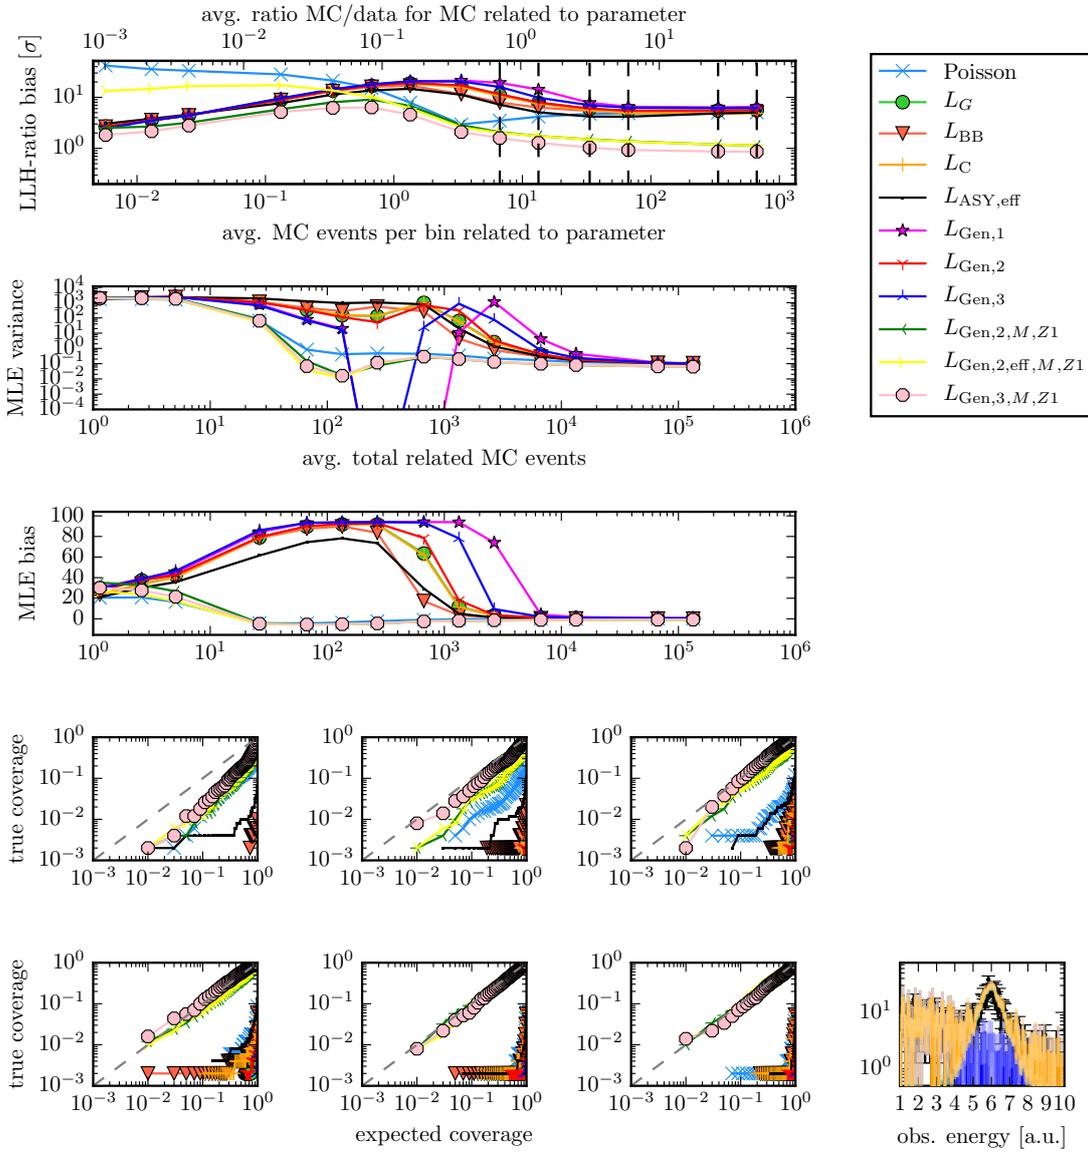


Figure 13. Bias of the log-likelihood ratio λ and coverage in dependence of the signal dataset statistics. Signal and background statistics increase together. The parameter of interest is the normalization of the peak in the signal dataset. $Z1$ denotes strategy 1 for “zero MC”, $Z2$ strategy 2 for “zero MC”, U denotes more “unbiased” weight variance estimation and M denotes shifted mean based on average MC events in all bins. The coverage is shown at 6 statistic levels (in order) which are indicated as dashed vertical lines in the upper plot. In the lower right plot a particular data realization and a MC realization at the highest statistic level is shown for the observable space. The non-standard likelihoods from other publications include $L_{ASY,eff}$ [13], L_{BB} [11], L_C [12] and L_G [7].

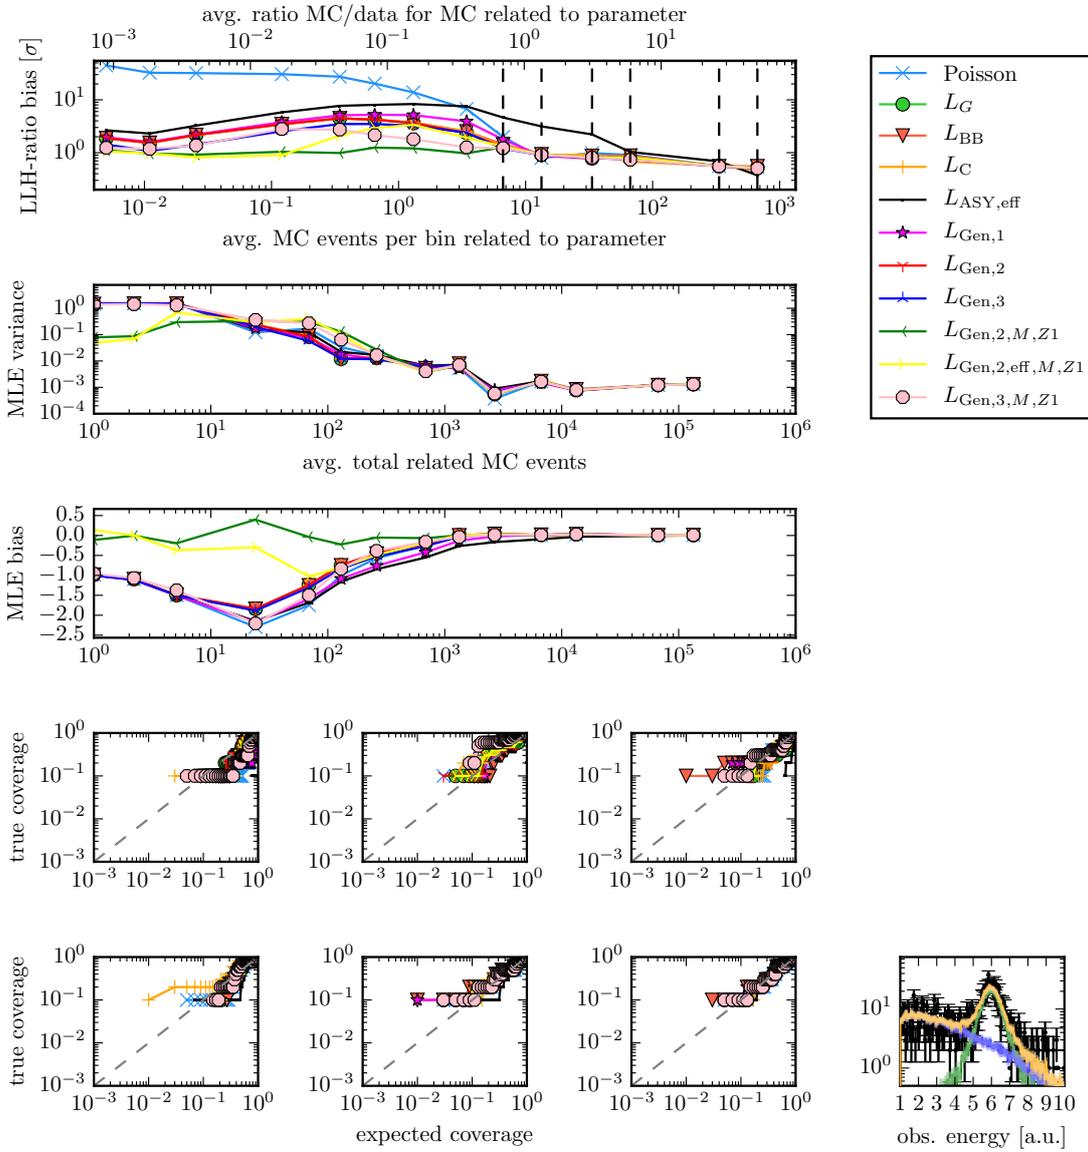


Figure 14. Bias of the log-likelihood ratio λ and coverage in dependence of the background dataset statistics. The signal and background statistics increase together. The parameter of interest is the spectral index of the background dataset. Z1 denotes strategy 1 for “zero MC”, Z2 strategy 2 for “zero MC”, U denotes more “unbiased” weight variance estimation and M denotes shifted mean based on average MC events in all bins. The coverage is shown at 6 statistic levels (in order) which are indicated as dashed vertical lines in the upper plot. In the lower right plot a particular data realization and a MC realization at the highest statistic level is shown for the observable space. The non-standard likelihoods from other publications include $L_{ASY,eff}$ [13], L_{BB} [11], L_C [12] and L_G [7].

References

- [1] D. Heck, J. Knapp, J. N. Capdevielle, G. Schatz and T. Thouw, *CORSIKA: A Monte Carlo code to simulate extensive air showers*, FZKA-6019 (1998)
- [2] F. Halzen and S.R. Klein, *IceCube: An Instrument for Neutrino Astronomy*, *Rev. Sci. Instrum.* **81** (2010) 081101 [[arXiv:1007.1247](#)].
- [3] ICECUBE collaboration, *Characterization of the Atmospheric Muon Flux in IceCube*, *Astropart. Phys.* **78** (2016) 1 [[arXiv:1506.07981](#)].
- [4] ICECUBE collaboration, *Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data*, *Astrophys. J.* **833** (2016) 3 [[arXiv:1607.08006](#)].
- [5] R.A. Fisher, *Statistical methods, experimental design, and scientific inference*, Oxford University Press, Oxford U.K. (1990).
- [6] D. Chirkin, *Event reconstruction in IceCube based on direct event re-simulation*, in *33rd International Cosmic Ray Conference (ICRC2013)*, Rio de Janeiro Brazil (2013), Number 0581.
- [7] T. Glüsenskamp, *Probabilistic treatment of the uncertainty from the finite size of weighted Monte Carlo data*, *Eur. Phys. J. Plus* **133** (2018) 218 [[arXiv:1712.01293](#)].
- [8] G. Bohm and G. Zech, *Statistics of weighted Poisson events and its applications*, *Nucl. Instrum. Meth. A* **748** (2014) 1 [[arXiv:1309.1287](#)].
- [9] N.A. Weiss, *A course in probability*, Addison-Wesley, New York U.S.A. (2006).
- [10] S.A. Murphy and A.W. Van der Vaart, *On profile likelihood*, *J. Am. Stat. Assoc.* **95** (2000) 449
- [11] R.J. Barlow and C. Beeston, *Fitting using finite Monte Carlo samples*, *Comput. Phys. Commun.* **77** (1993) 219.
- [12] D. Chirkin, *Likelihood description for comparing data to simulation of limited statistics*, in *Proceedings, 33rd International Cosmic Ray Conference (ICRC2013)*, Rio de Janeiro Brazil (2013), Number 0582.
- [13] C.A. Argüelles, A. Schneider and T. Yuan, *A binned likelihood for stochastic models*, *JHEP* **06** (2019) 030 [[arXiv:1901.04645](#)].
- [14] G. Bohm and G. Zech, *Comparison of experimental data to Monte Carlo simulation-parameter estimation and goodness-of-fit testing with weighted events*, *Nucl. Instrum. Meth. A* **691** (2012) 171.
- [15] M. Abramowitz and I. Stegun, *Handbook of mathematical functions*, Dover Publications Inc., New York U.S.A. (1970).
- [16] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *Phys. Rev. D* **98** (2018) 030001.
- [17] G.P. Egorychev, *Translations of Mathematical Monographs. Vol. 59: Integral representation and the computation of combinatorial sums*, AMS Press, New York U.S.A. (1984).
- [18] Y. Ma, J. Yu and Y. Wang, *Efficient recursive methods for partial fraction expansion of general rational functions*, *J. Appl. Math.* **2014** (2014) 1.
- [19] R.L. Graham, D.E. Knuth, O. Patashnik and S. Liu, *Concrete mathematics: a foundation for computer science*, *Comput. Phys.* **3** (1989) 106.
- [20] P. Vellaisamy and N.S. Upadhye, *Compound negative binomial approximations for sums of random variables*, *Probab. Math. Statist.* **29** (2009) 205.
- [21] B.C. Carlson, *Lauricella's hypergeometric function F_D* , *J. Math. Anal. Appl.* **7** (1963) 452.
- [22] G. Lauricella, *Sulle funzioni ipergeometriche a piu variabili*, *Rend. Circ. Mat. Palermo* **7** (1893) 111.