

Exact inference from noisy simulations in collider physics

Christopher Chang^{1,*} Benjamin Farmer^{2,†,||} Andrew Fowlie^{3,‡} and Anders Kvellestad^{1,§}

¹*Department of Physics, University of Oslo, N-0316 Oslo, Norway*

²*Energy Modelling Group, Gridcog, Victoria 3000, Australia*

³*X-HEP Laboratory, Department of Physics, School of Mathematics and Physics, Xi'an Jiaotong-Liverpool University, Suzhou, 215123, China*



(Received 17 February 2025; accepted 23 March 2026; published 21 April 2026)

We rely on Monte Carlo (MC) simulations to interpret searches for new physics at the Large Hadron Collider (LHC) and elsewhere. These simulations result in noisy and approximate estimators of selection efficiencies and likelihoods. In this context we pioneer an exact-approximate computational method—exact-approximate Markov chain Monte Carlo (MCMC), also known as pseudomarginal MCMC—that returns exact inferences despite noisy simulations. To do so, we introduce an unbiased estimator for a Poisson likelihood. We demonstrate the new estimator and new techniques in examples based on a search for neutralinos and charginos at the LHC using a simplified model. We find attractive performance characteristics—exact inferences are obtained for a similar computational cost to approximate ones from existing methods and inferences are robust with respect to the number of events generated per point. The unbiased estimator uses a Poisson-distributed number of MC events; it is also possible to construct a biased estimator whose bias decays factorially with increasing number of MC events.

DOI: [10.1103/2d39-gsjp](https://doi.org/10.1103/2d39-gsjp)

I. INTRODUCTION

The ATLAS and CMS experiments at the Large Hadron Collider (LHC) are searching for new particles that are predicted by beyond the Standard Model (BSM) theories. The experimental signatures of the new particles typically depend on unknown parameters that control their masses and interactions [1]. There is, however, an underappreciated issue when using data from LHC searches to infer these unknown parameters. The data collected by the LHC and the LHC searches are so complicated that we cannot explicitly compute a likelihood function for use in a statistical analysis [2].¹ Consequently, simplifications and approximations must be made.

Typically, we first make a lossy simplification of the experimental data: we discard the detailed event record and consider only whether an event falls into a judiciously

chosen signal region. The result is thus the number of events observed in the signal region. This is a counting experiment with a Poisson likelihood. However, this is not enough to allow us to compute the likelihood, as we cannot compute directly the expected number of such events from a model. Instead, we compute noisy estimators of the likelihood using Monte Carlo (MC) simulations of the experiment on a computer, e.g., using the Pythia [4], Herwig [5] and Sherpa [6] event-generators, and the Delphes [7] and Geant [8] detector simulators.

There are various statistical methodologies for learning about parameters using experimental data. We focus on Bayesian statistics, in which we learn about parameters by finding posterior distributions for them. The dominant algorithms for sampling from a posterior are variants of Markov chain Monte Carlo (MCMC; see, e.g., Ref. [9]). These algorithms are likelihood based; they require us to be able to compute the likelihood of the data observed by the LHC given a set of model parameters. Here, we pioneer exact-approximate MCMC—an MCMC algorithm that returns exact results from noisy estimates of the likelihood.² This class of algorithms is also known as pseudomarginal MCMC. For details of our numerical implementations in Python, C++ and within the ColliderBit [11] module of the GAMBIT [12] framework, see Appendix C. Our work differs from approximate simulation-based inference methods

* Contact author: c.j.chang@fys.uio.no

† Contact author: benf@gridcog.com

‡ Contact author: andrew.fowlie@xjtlu.edu.cn

§ Contact author: anders.kvellestad@fys.uio.no

|| Formerly Department of Physics, Imperial College London, London SW7 2AZ, United Kingdom.

¹Though see, e.g., the matrix-element method [3].

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

²Though note that other algorithms have exact-approximate variants, e.g., pseudomarginal slice-sampling [10].

such as approximate Bayesian computation (ABC; [13]) in that it returns exact rather than approximate results.

A. Exact-approximate MCMC

Remarkably, in special circumstances approximate estimators of the likelihood can lead to exactly correct results; these algorithms are known as exact-approximate or, in the context of Bayesian computation, pseudomarginal. In this work, we explore exact-approximate MCMC for LHC likelihoods. Exact-approximate MCMC was first applied in Ref. [14] and formalized in Refs. [15,16], where it is also referred to as pseudo-marginal MCMC. Similar ideas, however, were presented earlier in the context of lattice QCD [17].

Exact-approximate MCMC can be seen as MCMC that returns exact inferences despite a noisy estimator of the likelihood, $\hat{L}$. In order for exact-approximate MCMC to work, the estimator must satisfy:

- (1) The estimator is unbiased (at least up to a constant factor) with respect to the process by which it is generated during MCMC,

$$\langle \hat{L} \rangle = CL \quad (1)$$

for constant C . That is, on average, the estimator equals the exact likelihood.

- (2) The estimator is never negative,

$$\hat{L} \geq 0. \quad (2)$$

We later discuss adaptations of MCMC that remove this requirement.

If the estimator satisfies these requirements, the MCMC algorithm converges to the exactly correct stationary distribution. When $\hat{L} = L$, this is standard MCMC. Unfortunately, the usual maximum likelihood estimator (MLE) for the LHC likelihood from MC simulations is biased.³ The cause of this bias is the difference in the expected distribution of events in the experiment versus the MC simulation. Running until a fixed number of events have accumulated results in binomially distributed MC samples, whereas the fixed-time experiment should expect Poisson-distributed samples. This means that MCMC would converge to the wrong stationary distribution. In this work, we develop a new unbiased estimator of the likelihood, and investigate its performance characteristics compared to the traditional estimator.

The fact that we use an estimator in exact-approximate MCMC leads to fluctuations in the likelihood. When an upward fluctuation occurs, it becomes hard to accept subsequent transitions and the chain may become stuck. This leads to a poor MCMC acceptance rate and

efficiency.⁴ Reducing the variance of the estimator, however, could be computationally expensive, and there is thus a trade-off between the variance of the estimator and stickiness in the chain. There are limited theoretical results about the optimal variance or performance of exact-approximate MCMC. In specific cases the optimal variance of $\log \hat{L}$ was found to be around one [19,20] and, as expected, the performance of exact-approximate MCMC converges to that of MCMC as the variance of the estimator goes to zero [21]. Lastly, the optimal scaling of exact-approximate MCMC proposals differs from the well-known MCMC results [22], such that MCMC implementations may need to be retuned for optimal performance.

II. SETTING AND ESTIMATORS

We consider counting experiments where we compare the expected and observed numbers of events in a particular region of phase space, called the signal region; see e.g., Ref. [11] for an introduction. There is a known expected number of background events, b , and an unknown expected number of signal events, s .⁵ The experiment observes o events with likelihood

$$L = \text{Po}(o|\lambda) = \frac{e^{-\lambda} \lambda^o}{o!} \quad (3)$$

where $\lambda = s + b$. The expected number of signal events, s , and thus the likelihood, are intractable functions of a model's parameters, $\mathbf{x}$. That is, we cannot evaluate $s(\mathbf{x})$ directly. We may, however, express it as,

$$s(\mathbf{x}) = n_{\text{LHC}}(\mathbf{x}) \times \epsilon(\mathbf{x}) \quad (4)$$

where n_{LHC} is the number of expected events from our model in the experiment and the selection efficiency, ϵ , is the fraction of them that we expect to fall in the signal region. Fortunately, n_{LHC} is tractable and may be computed through

$$n_{\text{LHC}}(\mathbf{x}) = \mathcal{L} \times \sigma(\mathbf{x}) \quad (5)$$

for an integrated luminosity $\mathcal{L}$ and production cross section σ . The selection efficiency, ϵ , and thus the likelihood, may be inferred through MC. In MC we simulate signal-like events and their behavior in the experimental detector on a computer and count the fraction of events that fall into the signal region. Inferences about the selection efficiency thus usually suffer from MC uncertainty.

There are several existing approaches to handling this MC uncertainty. These existing approaches treat the selection efficiency as a nuisance parameter and construct a joint

³The bias of the MLE for the selection efficiency was explored in Ref. [18].

⁴There are solutions such as clamping randomness [10].

⁵We do not consider uncertainties on the background, b , in this work.

likelihood describing the experiment and the MC simulations. In the classic Barlow-Beeston approach for frequentist inference [23], the selection efficiency is subsequently profiled. This approach is further explored in Refs. [24–26].⁶ In the Bayesian approaches in Refs. [27–32], the selection efficiency is subsequently marginalized over a choice of prior.

Although these existing approaches incorporate MC uncertainty, the resulting inferences are inexact, as they reflect the results of finite-statistics MC simulations. In contrast, in exact-approximate MCMC, an estimator for the likelihood is implicitly averaged in a way that leads to exact inferences if the estimator is unbiased (referred to as pseudomarginalization). In other words, if one uses an unbiased likelihood estimator, the inferences match what one would obtain if one could compute the efficiencies exactly. We now introduce two likelihood estimators: a common biased estimator and our new unbiased estimator.

A. Maximum likelihood estimator (MLE)

Traditionally, every time we wish to estimate a likelihood, we generate a fixed number of MC events from our model, n_{MC} , with specialized software such as Pythia [4]. The computational cost of simulations is substantial and scales as $\mathcal{O}(n_{\text{MC}})$. We count the number of simulated events that fall in the signal region, k . This follows a binomial distribution, $k \sim \text{Binom}(n_{\text{MC}}, \epsilon)$ with probability mass function (pmf)

$$\text{Binom}(k|n_{\text{MC}}, \epsilon) = n_{\text{MC}} C_k \epsilon^k (1 - \epsilon)^{n_{\text{MC}} - k} \quad (6)$$

where we define $n_{\text{MC}} C_k = 0$ for $k > n_{\text{MC}}$ or $k < 0$. From k simulated events and Eq. (6), we construct the maximum likelihood estimator (MLE) for the unknown selection efficiency,

$$\hat{\epsilon} = \frac{k}{n_{\text{MC}}}. \quad (7)$$

We use this in Eq. (3) to construct an MLE estimate of the Poisson likelihood,

$$\hat{L}_{\text{MLE}} = \text{Po}(o|b + \hat{\epsilon} n_{\text{LHC}}) = \frac{e^{-(b + \hat{\epsilon} n_{\text{LHC}})} (b + \hat{\epsilon} n_{\text{LHC}})^o}{o!}. \quad (8)$$

This MLE estimate is noisy and in fact biased, that is,

$$\langle \hat{L}_{\text{MLE}} \rangle \equiv \sum_{k=0}^{n_{\text{MC}}} \hat{L}_{\text{MLE}} \text{Binom}(k|n_{\text{MC}}, \epsilon) \neq L. \quad (9)$$

⁶We plan to explore advantages of our unbiased estimator of the likelihood in optimization and more generally in a frequentist context in future works.

It is, however, always positive and consistent, since $\hat{L}_{\text{MLE}}$ converges in probability to L as $n_{\text{MC}} \rightarrow \infty$.

Typically, the number of events n_{MC} is chosen so that the variance and bias are negligible. For $\epsilon \ll 1$, we may estimate the fractional uncertainty on the estimate of the selection efficiency by

$$\frac{\sigma_{\hat{\epsilon}}}{\hat{\epsilon}} = \frac{\sigma_k}{k} \simeq \frac{1}{\sqrt{k}}, \quad (10)$$

where σ denotes standard deviation. Thus we might assume that simulating events until we collect $k \gtrsim 100$ signal events may be reasonable. The variance on the likelihood, however, depends on the ratio $n_{\text{LHC}}/n_{\text{MC}}$. E.g., consider the simple $o = 0$ case such that

$$\log \hat{L}_{\text{MLE}} = -\frac{n_{\text{LHC}}}{n_{\text{MC}}} k - b \quad (11)$$

with standard deviation,

$$\sigma_{\log \hat{L}_{\text{MLE}}} = \frac{n_{\text{LHC}}}{n_{\text{MC}}} \sigma_k \propto \frac{1}{\sqrt{n_{\text{MC}}}}. \quad (12)$$

Thus while we see the usual statistical scaling $1/\sqrt{n_{\text{MC}}}$, the ratio $n_{\text{LHC}}/n_{\text{MC}}$ plays an important role and we require $n_{\text{MC}} > n_{\text{LHC}}$ for negligible variance.

For LHC searches we are often fitting signals that are just a few events above background, that is, $s \simeq o - b \ll 100$. In this case, generating signallike events until $k \gtrsim 100$ simultaneously implies that $n_{\text{MC}} > n_{\text{LHC}}$.

B. Uniformly minimum variance unbiased estimator (UMVUE)

For use in exact-approximate MCMC, we construct an estimator of the likelihood that is unbiased for any choice of n_{MC} . In order to compute one, we slightly modify traditional MC simulation. For complete mathematical details, see Appendix A. For every likelihood estimate, we first draw the required number of MC events k_{MC} from a Poisson distribution with mean n_{MC} , that is, $k_{\text{MC}} \sim \text{Po}(n_{\text{MC}})$. We then generate our MC events and denote the number that fall into the signal region by k , as before. This allows a simple unbiased estimator of the likelihood. When $b = 0$ [33],

$$\hat{L}_{\text{UMVUE}} = \text{Binom}(o|k, f) = {}^k C_o f^o (1 - f)^{k-o} \quad (13)$$

where $f = n_{\text{LHC}}/n_{\text{MC}}$. That is, the pmf for a binomial $o \sim \text{Binom}(k, f)$. This is not strictly a binomial, however, as the binomial success probability, here f , may be greater than one. In the more realistic $b > 0$ case,

$$\hat{L}_{\text{UMVUE}} = \sum_{i=0}^o \text{Po}(o - i|b) \text{Binom}(i|k, f). \quad (14)$$

```

def likelihood_estimate(n_mc, n_lhc, o, b, model_parameters):
    1. Draw number of MC events from Poisson with mean  $n_{MC}$ 
    k_mc = poisson_rvs(n_mc)
    2. Generate  $k_{MC}$  MC events using e.g. Pythia
    events = generate_events(k_mc, model_parameters)
    3. Count number of MC events that pass selections
    k = len(apply_selections(events))
    4. Compute ratio  $n_{LHC}/n_{MC}$ 
    f = n_lhc / n_mc
    5. Return likelihood estimate using Eq. (14)
    return umvue(k, f, o, b)

```

FIG. 1. Pseudocode for the UMVUE estimator of the likelihood. We write it here as a function of the user-chosen expected number of MC events, n_{mc} ; the expected number of LHC events from Eq. (5), n_{lhc} ; the numbers of observed and background events in the search, o and b , respectively; and any model parameters for the signal model that must be used in MC event generation, $model_parameters$.

When $k \sim \text{Po}(en_{MC})$, these estimators are unbiased, $\langle \hat{L}_{UMVUE} \rangle = L$, as shown in Appendix A by a power-expansion similar to that in Ref. [33]. For clarity, we show pseudocode for producing this estimator in Fig. 1.

When $f = 1$, Eq. (13) reduces to an indicator function that requires the simulated number of events to exactly equal the observed number of events. This resembles ABC with zero tolerance. The fact that Eq. (14) is a binomial probability was remarked upon in Ref. [33]. Significantly, it means that in the limiting case $n_{MC} \rightarrow \infty$, the UMVUE estimator approaches the MLE.

The benefits of using $k_{MC} \sim \text{Po}(n_{MC})$ rather than a fixed number of MC events are twofold. First, as shown in Appendix B, an unbiased estimate does not exist if we simulate a fixed number of MC events. Second, it trivially generalizes to cases where the likelihood is a product of Poisson likelihoods,

$$\langle \prod_i \hat{L}_{UMVUEi} \rangle = \prod_i \langle \hat{L}_{UMVUEi} \rangle = \prod_i L_i, \quad (15)$$

even if the same MC simulations are used to estimate each factor. This trivial generalization happens because the number of simulated events in each search follow independent $\text{Po}(\epsilon_i n_{MC})$ distributions. As discussed in Appendix B, however, using a fixed number of MC events, n_{MC} , it would still be possible to construct an UMVUE of the first n_{MC} terms in an expansion of the Poisson likelihood. The residual bias would decay factorially with increasing n_{MC} .

As shown in Appendix A, the estimators Eqs. (13) and (14) are the uniformly minimum variance unbiased estimators (UMVUE) of the likelihood possible from this procedure. Despite that, they suffer from undesirable properties. E.g., the estimator Eq. (13) can be zero when $k < o$ and even negative when $f > 1$. In these cases the vanishing bias of $\hat{L}_{UMVUE}$ originates from cancellations between overestimates, $\hat{L}_{UMVUE} \gg L$, and underestimates,

$\hat{L}_{UMVUE} \ll L$, of the likelihood. Thus, even though it is the UMVUE, the variance can be enormous. While variance cannot spoil the correctness of exact-approximate algorithms, as discussed in Sec. IA, it could make them inefficient and it cannot be taken for granted that these estimators will work well in practice. Thus, we now explore the performance characteristics of these estimators in numerical experiments.

The existence of these “absurd” UMVUE estimators is discussed in Ref. [34] and attributed to the fact that we are “estimating the probability of an event of interest occurring over a long period of time from observations over a much shorter period.” That is indeed the case here, with the number of events playing the role of time: when $f > 1$, we are estimating the probability of o events in n_{MC} events from observations of only $n_{LHC} < n_{MC}$ events. There are relevant examples in Ref. [35] [Example 3.11] and Ref. [36] [Example 17.26].

III. METHODS

A. Negative estimates of likelihood

When $f > 1$, the UMVUE may be negative. This issue is sometimes known as the “sign problem” [17]. To use this estimator in an exact-approximate MCMC framework, we follow Ref. [37] and take the absolute value of the estimator, $|\hat{L}_{UMVUE}|$, but save the sign,

$$\sigma = \begin{cases} +1 & \hat{L}_{UMVUE} \geq 0 \\ -1 & \hat{L}_{UMVUE} < 0 \end{cases} \quad (16)$$

to the chain. When computing expectations, estimating densities, or computing effective sample size (ESS), we take into account the sign using weighted sums.

Indeed, the corrected ESS for parameter x , denoted $\bar{n}_{\text{eff}}(x)$, depends on the sign through

$$\bar{n}_{\text{eff}}(x) = \left[\frac{\sum_{i=1}^n \sigma_i}{n} \right]^2 n_{\text{eff}}(\sigma x), \quad (17)$$

where n is the number of samples in the chain and $n_{\text{eff}}(\sigma x)$ is the uncorrected ESS for the parameter σx . That is, the ESS for x is the ESS for the weighted quantity σx reduced by a factor accounting for the negative weights. When there are as many negative as positive sign likelihood samples, there are effectively no samples and $\bar{n}_{\text{eff}} = 0$.

B. Performance metrics

In the past even when the maximum likelihood estimator was used in MCMC, it was not knowingly done so in an exact-approximate MCMC framework. Instead, substantial n_{MC} was chosen to make variance and bias negligible and justify an approximation $\hat{L} \approx L$. We know of no general analytic results about the relative performance of these two

estimators in exact-approximate MCMC (though see Ref. [38]). Broadly speaking, we anticipate that the UMVUE may perform better, as we do not need to increase n_{MC} to remove the bias in our estimator; instead, the bias averages out over the whole chain. On other hand, the UMVUE estimator is far noisier, and this degrades the efficiency of exact-approximate MCMC as it results in sticking behavior and increased autocorrelation.

We judge performance by the effective number of posterior samples obtained—effective sample size (ESS)—per MC event simulated,

$$\text{ESS per MC event} = \frac{\text{ESS}}{\text{Total number of MC events}}. \quad (18)$$

We run an affine-invariant MCMC ensemble sampler [39], EMCEE [40], with 10 walkers for at least 100 000 steps. The bulk ESS is computed using ArviZ [41], adjusted via Eq. (17) when necessary.

C. Toy problems

We test the methods on toy problems. We take the luminosity, numbers of observed and background events from the SRWZ_15 signal region of the ATLAS-SUSY-2019-09 search [42]. That is, $o = 5$, $b = 2.8$ and $\mathcal{L} = 139/\text{fb}$ of data. This $\sqrt{s} = 13$ TeV search targeted neutralino-chargino production in final states with three leptons and missing transverse momentum.

We consider three toy models with increasing realism

- (1) First, we construct a simple one-dimensional model in which the selection efficiency, ϵ , is an input parameter. We take a flat prior on ϵ and fix the cross section to $\sigma = 1000$ fb.
- (2) Second, we construct a simplified model based on the TChiWZ topology: mass-degenerate $\chi_1^\pm$ and χ_2 particles are pair produced and decay exclusively through $\chi_1^\pm \rightarrow W\chi_1$ and $\chi_2 \rightarrow Z\chi_1$, respectively.

We fix the production cross section to $\sigma = 1000$ fb and compute a selection efficiency for the SRWZ_15 signal region as a function of (m_1, m_2) through SModelS [43].

There are thus two unknown parameters: $m_1 \equiv m(\chi_1)$ and $m_2 \equiv m(\chi_2) = m(\chi_1^\pm)$. We take flat priors between $0 < m_1 < 300$ GeV and $m_1 + M_Z < m_2 < 300$ GeV.

- (3) Lastly, we use the same simplified model, but compute the production cross section $\sigma(pp \rightarrow \chi_2\chi_1^\pm)$ as a function of (m_1, m_2) using tabulated winolike neutralino-chargino production cross sections [44–46].

The tabulated cross sections were computed at NLO + NLL precision assuming mass-degenerate winolike χ_2 and $\chi_1^\pm$, and a binolike χ_1 and use an envelope of CTEQ6.6 and MSTW2008nlo90cl parton distribution functions [44].

In our toy problems, we never simulate collider events with, e.g., Pythia as this is computationally expensive and unnecessary for our purposes. Instead, we perform white-box experiments in which we internally choose a true selection efficiency ϵ and use it as a parameter in a binomial or Poisson distribution for the MLE and UMVUE estimators, respectively. At every point, we then simulate from the binomial and Poisson distributions directly, rather than perform collider simulations.

IV. RESULTS

In Fig. 2 we show the posterior pdf (upper panel) and likelihood estimates (lower panel) for the unknown selection efficiency, ϵ , for $n_{\text{MC}} = 2n_{\text{LHC}}$. The MLE and UMVUE estimators produce samples at similar rates, around one posterior sample per 10^7 MC events. The MCMC results (stepped histogram) match the expected results (smooth lines). The MLE estimator, however, shows noticeable bias in the posterior. The mean and standard deviation of the likelihood estimates (lower panel) show that the estimators




FIG. 2. The posterior pdf reconstructed by MCMC for the unknown selection efficiency (top) and the mean and standard deviation of the likelihood function (bottom) using MLE and UMVUE Poisson likelihood estimators.




FIG. 3. Effective number of posterior samples for the unknown selection efficiency per MC simulation (top) and bias and variance (bottom) for the MLE and UMVUE Poisson likelihood estimators relative to the exact likelihood as we change n_{MC} .

are similarly noisy; though for the UMVUE, at equilibrium that noise results in the exactly correct distribution.

In Fig. 3 we explore the performance on this problem as we change the number of MC events. First, consider the MCMC efficiency (top panel). For $n_{MC} < n_{LHC}$, the UMVUE estimator suffers from terrible efficiency, due to the possibility of negative likelihood estimates. The MCMC efficiency itself becomes noisy and hard to estimate reliably. Once $n_{MC} \gtrsim n_{LHC}$, however, the noise becomes unimportant and the UMVUE and MLE estimators show remarkably similar MCMC efficiencies that deteriorate as $1/n_{MC}$. Thus, the UMVUE estimator yields unbiased inferences for approximately the same computational cost as biased ones. For the UMVUE, the best choice appears to be about $n_{MC} \approx n_{LHC}$. For the MLE estimate, n_{MC} should be as small as possible while resulting in an acceptable bias.

We show the noise and bias (lower panel) at the best-fit selection efficiency. For the MLE estimator, the noise goes to zero as number of simulations goes to zero, as we always estimate that $k = 0$. Similarly bias reaches a horizontal asymptote: the bias of the $k = 0$ estimate. As the number of simulations increases, the noise and bias decay as $1/\sqrt{n_{MC}}$.




FIG. 4. The posterior pdf for the selection efficiency from the MLE likelihood estimator as we change n_{MC} . We show the exact result (red) for reference.

For the UMVUE estimator, the bias always vanishes. As discussed, the noise starts as enormous when $n_{MC} < n_{LHC}$, and subsequently decays as $1/\sqrt{n_{MC}}$.

Lastly, we consider the impact of the bias on inferences in Fig. 4 by showing the posterior pdf for the selection efficiency from the MLE likelihood estimator as we change n_{MC} . We see that for $n_{MC} < n_{LHC}$, the pdf from the MLE estimator is much broader and flatter than the exact one. When $n_{MC}/n_{LHC} = 0.01$, the pdf is almost flat, resembling the prior for the selection efficiency. Thus, the bias when $n_{MC} < n_{LHC}$ leads to faulty inferences. Indeed, for a posterior that is indiscernible by eye from the true one in Fig. 4, requires $n_{MC}/n_{LHC} \gtrsim 50$. This choice would




FIG. 5. Posterior pdf for the TChiWZ simplified model masses and SRWZ₁₅ signal region of the ATLAS-SUSY-2019-09 search [42] found from MLE and UMVUE Poisson likelihood estimators, and the exact Poisson likelihood.

TABLE I. Mean and MC standard error for the TChiWZ simplified model masses and SRWZ_15 signal region of the ATLAS-SUSY-2019-09 search [42] found from MLE and UMVUE Poisson likelihood estimators, and the exact Poisson likelihood.

Estimator	Posterior mean $\pm$ MCSE	
	m_1 (GeV)	m_2 (GeV)
MLE	76.0 ± 0.1	219.4 ± 0.1
UMVUE	76.4 ± 0.1	218.8 ± 0.1
Exact	76.4 ± 0.1	219.0 ± 0.1

perform a factor of about 25 worse than the UMVUE estimator at $n_{\text{MC}}/n_{\text{LHC}} = 2$.

We now consider the TChiWZ model in which degenerate $\chi_2\chi_1^\pm$ particles are pair produced and decay through $\chi_2 \rightarrow Z\chi_1$ and $\chi_1^\pm \rightarrow W\chi_1$, respectively. There are two unknown parameters, m_1 and m_2 , that determine the selection efficiency. We run 10 walkers for 1 000 000 steps and $n_{\text{MC}} = 2n_{\text{LHC}}$. In each case, we collect an ESS of about 200 000 samples and obtain $\hat{R} = 1.0$ [47]. In Fig. 5 we show 68% and 95% credible regions on the (m_1, m_2) plane, as well as one-dimensional marginal posterior distributions, computed from the exact likelihood and from our two estimators. The MLE estimator and UMVUE estimator result in about 10^{-7} samples per MC event for both the m_1 and m_2 parameters. The UMVUE estimator agrees with the exact result, as expected, though the MLE estimator shows bias. The differences between the MLE and exact result are slight though noticeable, particularly in the 95% contours on the (m_1, m_2) plane. We present estimates of the posterior means and MC standard error (MCSE) in Table I. As expected the UMVUE estimator agrees with




FIG. 6. As in Fig. 5, though with the cross section computed as a function of (m_1, m_2) at NLO + NLL.

TABLE II. As in Table I, though with the cross section computed as a function of (m_1, m_2) at NLO + NLL, and shown as well as the mass parameters.

Estimator	Posterior mean $\pm$ MCSE		
	m_1 (GeV)	m_2 (GeV)	σ (fb)
MLE	74.3 ± 0.1	223.2 ± 0.1	2082 ± 5
UMVUE	75.1 ± 0.1	222.2 ± 0.1	2099 ± 5
Exact	75.2 ± 0.1	222.2 ± 0.1	2103 ± 5

results from the exact likelihood, whereas the MLE shows significant bias. Lastly, we perform Kolmogorov-Smirnov (KS) two-sample tests [48,49] on the m_2 samples thinned to the ESS. The MLE and UMVUE samples are significantly different ($p < 10^{-9}$) whereas the UMVUE and exact likelihood samples have no significant difference ($p \simeq 0.1$).

Finally, we compute the cross section as a function of (m_1, m_2) . As before, we run 10 walkers for 1 000 000 steps and $n_{\text{MC}} = 2n_{\text{LHC}}$, obtaining an ESS of about 200 000 and $\hat{R} = 1.0$ [47]. The MLE estimator and UMVUE estimator result in about 10^{-7} samples per MC event for both the m_1 and m_2 parameters. We show 68% and 95% credible regions on the (m_1, m_2) plane, as well as one-dimensional marginal posterior distributions in Fig. 6. The shapes of the distributions slightly changed, although remain similar to Fig. 5. The MLE estimator shows slightly more bias in the (m_1, m_2) parameters in Fig. 5. The bias can be compared to the MCSE in Table II. Lastly, we again perform a KS test on the m_2 samples. The MLE and UMVUE samples are significantly different ($p < 10^{-9}$) whereas the UMVUE and exact likelihood samples have no significant difference ($p \simeq 0.5$).

V. DISCUSSION AND CONCLUSIONS

The Poisson likelihood in collider searches for new physics can only be estimated through noisy MC simulations. Our inferences are thus at present noisy and approximate, though the noise can be reduced by increasing the number of MC events. With this in mind, we pioneered exact-approximate inference in this setting—exact statistical inference despite an approximate likelihood. The exact-approximate MCMC algorithm (also known as pseudo-marginal MCMC) required us to construct a novel, unbiased estimator for a Poisson likelihood. The estimator is only slightly more complicated than the traditional maximum likelihood one, and we provide public C++, Python and ColliderBit implementations.

We investigated the MCMC efficiency, bias and noise characteristics of the traditional and new estimators on toy problems. The toy problems were based on an ATLAS search for new physics with the TChiWZ topology. We constructed a simplified model with unknown masses of a parent and daughter particle. Lastly, we assumed that these were chargino and neutralino particles and computed realistic cross sections using tabulated NLO + NLL predictions.

We found that, with our new unbiased estimator, exact results could be obtained at a remarkably similar computational cost to biased results. The performance, however, depended strongly on the average number of MC events that were simulated per point. The unbiased estimator suffered from terrible noise and inefficiency when an inadequate number of MC events was chosen. Thus, the most efficient choice of estimator depends on the degree of bias that a user can tolerate. Lastly, the unbiased estimator is undoubtedly a safer choice as it returns exact inferences regardless the number of MC simulations chosen; whereas the MLE estimator returns faulty inferences when the number of MC simulations was inappropriately chosen.

There remains a few open questions. There may be other estimators with advantageous bias-variance characteristics. We anticipate that a degree of bias could be tolerated if the variance was such that it led to improved performance. In particular, it might be possible, though challenging, to construct a non-negative unbiased estimator [50], or there may be ways to reduce the bias in the MLE estimator to a tolerable level [51]. Lastly, if one wished to stick to a fixed number of MC events, it would be possible to construct a biased estimator whose bias decays factorially with an increasing number of MC events. For now, though, the estimator and methods here show acceptable performance and result in exact inferences, and could become standard in Bayesian inference involving collider likelihoods or more generally noisy estimates of Poisson likelihoods.

ACKNOWLEDGMENTS

We would like to thank Andy Buckley for discussions and feedback. A. F. was supported by RDF-22-02-079 and NSFC RFIS-II W2432006. C. C. and A. K. were supported by the Research Council of Norway (RCN) through the FRIPRO Grant No. 323985 PLUMBIN'.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX A: PROOF THAT ESTIMATOR IS UMVUE

We wish to show that our estimator is the uniform minimum variance unbiased estimator (UMVUE). That is, it is expected to equal the true likelihood when $k \sim \text{Po}(\epsilon n_{\text{MC}})$, and that the variance of our estimator is the minimum among estimators with this property for all possible values of the likelihood. The expectation of our estimator may be written,

$$\langle \hat{L}_{\text{UMVUE}} \rangle = \sum_{k=0}^{\infty} \hat{L}_{\text{UMVUE}} \text{Po}(k|\epsilon n_{\text{MC}}), \quad (\text{A1})$$

and the exact likelihood is a Poisson,

$$L = \text{Po}(o|b + \epsilon n_{\text{LHC}}). \quad (\text{A2})$$

First consider the $b = 0$ case; demonstrating that Eq. (13) is the UMVUE is extremely similar to the proof in Ref. [33]. First, we write the Poisson likelihood as a series in s ,

$$\text{Po}(o|s) = \frac{e^{-s} s^o}{o!} = \sum_{i=0}^{\infty} \frac{(-1)^i}{i! o!} s^{o+i}. \quad (\text{A3})$$

We now construct the UMVUE by replacing each power of s by the UMVUE,

$$\text{UMVUE}[\text{Po}(o|s)] = \text{Po}(o|s)|_{s^n \rightarrow \text{UMVUE}[s^n]}. \quad (\text{A4})$$

This is the UMVUE because it is a linear combination of UMVUE estimators [52] [Corollary 4.10]. The UMVUE for the powers of s are [33]

$$\text{UMVUE}[s^n] = \begin{cases} \left(\frac{n_{\text{LHC}}}{n_{\text{MC}}}\right)^n \frac{k!}{(k-n)!} = f^n \frac{k!}{(k-n)!} & k \geq n \\ 0 & k < n \end{cases}. \quad (\text{A5})$$

Performing the substitution in Eq. (A4) yields

$$\text{UMVUE}[\text{Po}(o|s)] = \sum_{i=0}^{k-o} \frac{(-1)^i f^{o+i} k!}{i! o! (k-o-i)!} \quad (\text{A6})$$

$$= {}^k C_{o,f} \sum_{i=0}^{k-o} {}^{k-o} C_i (-1)^i f^i \quad (\text{A7})$$

$$= {}^k C_{o,f} (1-f)^{k-o} = \text{Binom}(o|k, f). \quad (\text{A8})$$

Thus matching Eq. (13), the only difference being that we have f whereas Ref. [33] has $1/n$ and we may have $f > 1$ whereas $1/n \leq 1$.

For the $b > 0$ case, consider the fact that we may write the exact likelihood as

$$L = \text{Po}(o|s+b) = \sum_{i=0}^o \text{Po}(o-i|b) \text{Po}(i|s). \quad (\text{A9})$$

Thus by linearity of the UMVUE we have

$$\begin{aligned} \text{UMVUE}[\text{Po}(o|s+b)] &= \sum_{i=0}^o \text{Po}(o-i|b) \text{UMVUE}[\text{Po}(i|s)] \\ &= \sum_{i=0}^o \text{Po}(o-i|b) \text{Binom}(i|k, f), \end{aligned} \quad (\text{A10})$$

matching Eq. (14). As previously, this is the UMVUE because it is a linear combination of UMVUE estimators.

APPENDIX B: NO UNBIASED ESTIMATOR FOR FIXED NUMBER OF MC EVENTS

Expanding in powers of the unknown efficiency, ϵ , shows that no unbiased estimator can exist if we simulated a fixed number of MC events such that the number of signal events followed a binomial. The Poisson likelihood contains all powers of ϵ through the exponential term $e^{-n_{\text{LHC}}\epsilon}$. The expectation of any estimator (which cannot depend explicitly on ϵ) from a binomial $\text{Binom}(n_{\text{MC}}, \epsilon)$ contains at most $\epsilon^{n_{\text{MC}}}$. That is, we would require an estimator $\hat{L}$ such that

$$\sum_{k=0}^{n_{\text{MC}}} \hat{L} \text{Binom}(k|n_{\text{MC}}, \epsilon) = \text{Po}(o|b + \epsilon n_{\text{LHC}}) \quad (\text{B1})$$

but the right-hand side contains all powers of ϵ , whereas the left-hand side contains at most $\epsilon^{n_{\text{MC}}}$.

We can quantify the bias in an estimator that uses a fixed number of MC events as follows. Suppose we had an estimator that was an unbiased estimator of the Poisson likelihood truncated so that it contained powers no greater than $\epsilon^{n_{\text{MC}}}$. The remaining $\epsilon^{n_{\text{MC}}+1}$ term and higher powers (for which there can be no unbiased estimator) would lead to bias

$$\langle \hat{L} \rangle - \text{Po}(o|b + \epsilon n_{\text{LHC}}) \approx -(\epsilon n_{\text{LHC}})^{n_{\text{MC}}+1} \times \sum_{i=0}^o \frac{b^{o-i} e^{-b}}{i!(n_{\text{MC}} + 1 - i)!(o - i)!}. \quad (\text{B2})$$

Thus, the bias would decay factorially with increasing n_{MC} .

APPENDIX C: PUBLIC COMPUTER CODES

We provide a simple implementation of the UMVUE as a public library at [53].

1. C++ implementation

We provide a simple header-only library `ideal.hpp` providing the function:

```
double umvue_poisson_like(int k, double b, int o, int n_mc, double n_exp)
```

in the namespace `ideal`. This function returns an unbiased estimate of a Poisson likelihood when k events were simulated from an expected n_{mc} simulations, and o events were observed in a sample corresponding to n_{exp} events.

For completeness we provide functionality for drawing the number of MC simulations,

```
int umvue_draw_n_mc(double lambda_mc)
```

and a templated function

```
template<typedef engine_type>
int umvue_draw_n_mc(double lambda_mc, engine_type engine)
```

permitting greater control over random number generation. The inferred `engine_type` should be a random number engine, e.g., an `std::mt19937` from the C++11 standard library header `<random>` [54].

The basic usage would be

```
1 #include <iostream>
2 #include "ideal.hpp"
3
4 int main(int, char**) {
5     std::cout << ideal::umvue_poisson_like(1000, 10, 20, 10000, 100);
6     std::cout << ideal::umvue_draw_n_mc(1000);
7 }
```

See `example.cpp`. This example program can be built by entering `make example`.

2. Python bindings

We supply Python bindings for the C++ functions `umvue_poisson_like` and `umvue_draw_n_mc` using `pybind11` [55]. The `ideal` module can be built and installed by `make`. The equivalent example program would be:

```
1 import ideal
2
3 if __name__ == "__main__":
4     print(ideal.umvue_poisson_like(1000, 10, 20, 10000, 100))
5     print(ideal.umvue_draw_n_mc(1000))
```

3. Implementation in GAMBIT

We use code similar to that in Appendix C 1 as part of the ColliderBit module of GAMBIT and implement an option to enable the UMVUE. The UMVUE may be turned on by setting the estimator in the rules governing the MC generation of collider events in the yaml input file:

```
1 Rules:
2   - capability: RunMC
3     function: operateLHCLoop
4     options:
5       LHC_13TeV: # or LHC_8TeV
6         poisson_estimator: "UMVUE" # by default, "MLE"
7         mean_nEvents: 1000
```

We anticipate that other estimators could be added in the future. The mean number of MC events can be controlled directly by setting `mean_nEvents`, as in the snippet above, or by setting the ratio `mean_relative_nEvents`:

```
1 Rules:
2   - capability: RunMC
3     function: operateLHCLoop
4     options:
5       LHC_13TeV: # or LHC_8TeV
6         poisson_estimator: "UMVUE" # by default, "MLE"
7         mean_relative_nEvents: 2
```

If both `mean_nEvents` and `mean_relative_nEvents` are specified, GAMBIT throws a user input error. The default behavior is to run with `mean_relative_nEvents = 1`. The MLE estimator can also be used with the new `mean_nEvents` and `mean_relative_nEvents` settings, as well as with the existing `min_nEvents` and `max_nEvents` settings.

Our implementation of the `mean_relative_nEvents` option involved two complications. First, we wish to use the same set of MC events to compute selection efficiencies for multiple searches and combine likelihood estimators through Eq. (15). While these different searches would share an identical cross section, they could be associated with different integrated luminosities. We thus set the expected number of MC events through,

$$\begin{aligned} n_{\text{MC}} &= \text{mean_relative_nEvents} \times \max n_{\text{LHC}} \\ &= \text{mean_relative_nEvents} \times \sigma \times \max \mathcal{L}, \quad (\text{C1}) \end{aligned}$$

where we maximize across searches sharing the same MC events.

Using Eq. (C1) to set the expected number of MC events requires computation of the cross section before event generation. Thus, we separate cross section computation and event generation in GAMBIT. Initial cross section estimates are performed by simulating the hard scatter in Pythia without showering or hadronization. Separate events are generated with showering and hadronization during the ColliderBit main event loop for calculating signal estimates in collider analyses. In the near future, we will also support extracting these initial estimates from external cross section calculators.

The UMVUE estimator interacts with other ColliderBit settings as follows. The number of MC events generated by ColliderBit was previously controlled by setting a minimum and a maximum number of MC events through `min_nEvents` and `max_nEvents`, respectively. These settings allow ColliderBit to end event generation when `max_nEvents` events have been generated or when anywhere between `min_nEvents` and `max_nEvents` have been generated and convergence criteria are met. These options are ignored when enabling the UMVUE estimator. By clipping the Poisson distributed number of MC events to lie between `min_nEvents` and `max_nEvents`, they would introduce bias into the estimator. In the event that the UMVUE estimator is used with these settings, ColliderBit will throw an error.

Second, in ColliderBit, the uncertainty in a predicted signal plus background ($s + b$) can be described by a normal or log-normal probability distribution with variance γ^2 for each signal region, set by the combination of experiment-provided background uncertainty and MC signal uncertainty. In ColliderBit, these nuisances (γ) can be either profiled or marginalized. However, since the UMVUE was constructed for Bayesian computation, only marginalization is implemented for the UMVUE estimator. An appropriate error is thrown when attempting to profile these uncertainties and use the UMVUE estimator.

The marginalization of these nuisances in ColliderBit is performed by Monte Carlo integration. The resulting estimate is thus unbiased. As evaluating the integrand is computationally cheap compared to collider simulation, the noise can in any case be made negligible.

[1] S.S. AbdusSalam *et al.*, Simple and statistically sound recommendations for analysing physical theories, *Rep. Prog. Phys.* **85**, 052201 (2022).

[2] J. Brehmer and K. Cranmer, Simulation-based inference methods for particle physics, in *Artificial Intelligence for High Energy Physics* (2020), pp. 579–611, [arXiv:2010.06439](https://arxiv.org/abs/2010.06439).

- [3] I. Volobouev, Matrix element method in HEP: Transfer functions, efficiencies, and likelihood normalization, [arXiv:1101.2259](#).
- [4] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and Peter Z. Skands, An introduction to PYTHIA 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [5] M. Bahr *et al.*, Herwig++ physics and manual, *Eur. Phys. J. C* **58**, 639 (2008).
- [6] Sherpa Collaboration, Event generation with Sherpa 2.2, *SciPost Phys.* **7**, 034 (2019).
- [7] DELPHES 3 Collaboration, DELPHES 3, A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [8] R. Brun, F. Bruyant, F. Carminati, S. Giani, M. Maire, A. McPherson *et al.*, GEANT Detector Description and Simulation Tool, CERN Program Library, Long Wwriteup (1994) [10.17181/CERN.MUHF.DMJ1](#).
- [9] S. Brooks, A. Gelman, G. Jones, and X. Meng, *Handbook of Markov Chain Monte Carlo* (CRC Press, United States, 2011).
- [10] I. Murray and M. M. Graham, Pseudomarginal slice sampling, in *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, edited by A. Gretton and C. C. Robert, Vol. 51 of Proc. Mach. Learn. Res. (Cadiz, Spain, 2016), pp. 911–919, [arXiv:1510.02958](#).
- [11] GAMBIT Collaboration, ColliderBit: A GAMBIT module for the calculation of high-energy collider observables and likelihoods, *Eur. Phys. J. C* **77**, 795 (2017).
- [12] GAMBIT Collaboration, GAMBIT: The global and modular beyond-the-standard-model inference tool, *Eur. Phys. J. C* **77**, 784 (2017).
- [13] S. A. Sisson, Y. Fan, and M. A. Beaumont, Overview of approximate Bayesian computation, in *Handbook of Approximate Bayesian Computation* (Chapman and Hall/CRC, Boca Raton, 2018).
- [14] M. A. Beaumont, Estimation of population growth or decline in genetically monitored populations, *Genetics* **164**, 1139 (2003).
- [15] J. Møller, A. N. Pettitt, R. Reeves, and K. K. Berthelsen, An efficient Markov Chain Monte Carlo method for distributions with intractable normalising constants, *Biometrika* **93**, 451 (2006).
- [16] C. Andrieu and G. O. Roberts, The pseudomarginal approach for efficient Monte Carlo computations, *Ann. Stat.* **37**, 697 (2009).
- [17] L. Lin, K. F. Liu, and J. H. Sloan, A stochastic Monte Carlo algorithm, *Phys. Rev. D* **61**, 074505 (2000).
- [18] H. Dembinski and M. Schmelling, Bias, variance, and confidence intervals for efficiency estimators in particle physics experiments, [arXiv:2110.00294](#).
- [19] A. Doucet, M. Pitt, G. Deligiannidis, and R. Kohn, Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator, *Biometrika* **102**, 295 (2015).
- [20] C. Sherlock, A. H. Thiery, G. O. Roberts, and J. S. Rosenthal, On the efficiency of pseudo-marginal random walk Metropolis algorithms, *Ann. Stat.* **43**, 238 (2015).
- [21] C. Andrieu and M. Vihola, Convergence properties of pseudo-marginal Markov Chain Monte Carlo algorithms, *Ann. Appl. Probab.* **25**, 1030 (2015).
- [22] C. Sherlock, Optimal scaling for the pseudo-marginal random walk metropolis: Insensitivity to the noise generating mechanism, *Methodol. Comput. In Appl. Probab.* **18**, 869 (2015).
- [23] R. J. Barlow and C. Beeston, Fitting using finite Monte Carlo samples, *Comput. Phys. Commun.* **77**, 219 (1993).
- [24] C. A. Argüelles, A. Schneider, and T. Yuan, Tackling limited simulation and small signals, *Proc. Sci. ICRC2019* (2020) 1005 [[arXiv:1907.10636](#)].
- [25] H. P. Dembinski and A. Abdelmotteleb, A new maximum-likelihood method for template fits, *Eur. Phys. J. C* **82**, 1043 (2022).
- [26] C.-A. Alexe, J. Bendavid, L. Bianchini, and D. Bruschini, Under-coverage in high-statistics counting experiments with finite MC samples, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1086**, 171360 (2026).
- [27] D. Chirkin, Likelihood Description for Comparing Data to Simulation of Limited Statistics, in *International Cosmic Ray Conference*, Vol. 33 of International Cosmic Ray Conference (2013), p. 3406, [arXiv:1304.0735](#).
- [28] R. Aggarwal and A. Caldwell, Error bars for distributions of numbers of events, *Eur. Phys. J. Plus* **127**, 24 (2012).
- [29] T. Glüsenskamp, Probabilistic treatment of the uncertainty from the finite size of weighted Monte Carlo data, *Eur. Phys. J. Plus* **133**, 218 (2018).
- [30] T. Glüsenskamp, A unified perspective on modified Poisson likelihoods for limited Monte Carlo data, *J. Instrum.* **15**, P01035 (2020).
- [31] C. A. Argüelles, A. Schneider, and T. Yuan, A binned likelihood for stochastic models, *J. High Energy Phys.* **06** (2019) 030.
- [32] S. Liu and C. McGrew, Binned likelihood including Monte Carlo statistical uncertainty in Bayesian inference, [arXiv:2304.05433](#).
- [33] G. J. Glasser, Minimum variance unbiased estimators for Poisson probabilities, *Technometrics* **4**, 409 (1962).
- [34] E. L. Lehmann, Estimation with inadequate information, *J. Am. Stat. Assoc.* **78**, 624 (1983).
- [35] E. L. Lehmann, *Theory of Point Estimation*, 2nd ed. (Springer, New York, 1998).
- [36] M. Kendall, A. Stuart, and J. Ord, *The Advanced Theory of Statistics*, 4th ed. (MacMillan, London, 1979), Vol. 2.
- [37] A.-M. Lyne, M. Girolami, Y. Atchadé, H. Strathmann, and D. Simpson, On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods, *Stat. Sci.* **30**, 443 (2015).
- [38] C. Andrieu and M. Vihola, Establishing some order amongst exact approximations of MCMCs, *Ann. Appl. Probab.* **26**, 2661 (2016).
- [39] J. Goodman and J. Weare, Ensemble samplers with affine invariance, *Commun. Appl. Math. Comput. Sci.* **5**, 65 (2010).
- [40] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, EMCEE: The MCMC Hammer, *Publ. Astron. Soc. Pac.* **125**, 306 (2013).

- [41] R. Kumar, C. Carroll, A. Hartikainen, and O. Martin, Arviz a unified library for exploratory analysis of Bayesian models in Python, *J. Open Source Software* **4**, 1143 (2019).
- [42] ATLAS Collaboration, Search for chargino–neutralino pair production in final states with three leptons and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector, *Eur. Phys. J. C* **81**, 1118 (2021).
- [43] G. Alguero, J. Heisig, C. K. Khosa, S. Kraml, S. Kulkarni, A. Lessa, H. Reyes-González, W. Waltenberger, and A. Wongel, Constraining new physics with SModelS version 2, *J. High Energy Phys.* **08** (2022) 068.
- [44] LHC SUSY Cross Section Working Group, NLO-NLL winolike chargino-neutralino (N2C1) cross sections, (2017), <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections13TeVn2x1wino>.
- [45] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, Gaugino production in proton-proton collisions at a center-of-mass energy of 8 TeV, *J. High Energy Phys.* **10** (2012) 081.
- [46] B. Fuks, M. Klasen, D. R. Lamprea, and M. Rothering, Precision predictions for electroweak superpartner production at hadron colliders with RESUMMINO, *Eur. Phys. J. C* **73**, 2480 (2013).
- [47] A. Vehtari, A. Gelman, D. Simpson, B. Carpenter, and P.-C. Bürkner, Rank-normalization, folding, and localization: An improved $\hat{r}$ for assessing convergence of MCMC (with discussion), *Bayesian Anal.* **16**, 667 (2021).
- [48] A. N. Kolmogorov, Sulla determinazione empirica di una legge di distribuzione, *Giorn. Dell'inst. Ital. Degli Att.* **4**, 89 (1933).
- [49] N. Smirnov, Table for estimating the goodness of fit of empirical distributions, *Ann. Math. Stat.* **19**, 279 (1948).
- [50] P. E. Jacob and A. H. Thiery, On nonnegative unbiased estimators, *Ann. Stat.* **43**, 769 (2015).
- [51] I. Kosmidis, Bias in parametric estimation: Reduction and useful side-effects, *WIREs Comput. Stat.* **6**, 185 (2014).
- [52] L. Zhenhua, Advanced Statistical Theory I, Lecture Notes (2019), <https://blog.nus.edu.sg/zhenhua/files/2020/08/ST5215-lecture-note.pdf>.
- [53] <https://github.com/xhep-lab/ideal>.
- [54] <https://en.cppreference.com/w/cpp/header/random>.
- [55] W. Jakob, J. Rhineland, D. Moldovan *et al.*, pybind11—Seamless operability between C++11 and Python, (2017), <https://github.com/pybind/pybind11>.