An analysis of Bayesian estimates for missing higher orders in perturbative calculations

Claude Duhr,^a Alexander Huss,^a Aleksas Mazeliauskas^a and Robert Szafron^b

^a*Theoretical Physics Department, CERN,
CH-1211 Geneva 23, Switzerland*

^b*Department of Physics, Brookhaven National Laboratory,
Upton, N.Y., 11973, U.S.A*

E-mail: claude.duhr@cern.ch, alexander.huss@cern.ch,
aleksas.mazeliauskas@cern.ch, rszafron@bnl.gov

ABSTRACT: With current high precision collider data, the reliable estimation of theoretical uncertainties due to missing higher orders (MHOs) in perturbation theory has become a pressing issue for collider phenomenology. Traditionally, the size of the MHOs is estimated through scale variation, a simple but ad hoc method without probabilistic interpretation. Bayesian approaches provide a compelling alternative to estimate the size of the MHOs, but it is not clear how to interpret the perturbative scales, like the factorisation and renormalisation scales, in a Bayesian framework. Recently, it was proposed that the scales can be incorporated as hidden parameters into a Bayesian model. In this paper, we thoroughly scrutinise Bayesian approaches to MHO estimation and systematically study the performance of different models on an extensive set of high-order calculations. We extend the framework in two significant ways. First, we define a new model that allows for asymmetric probability distributions. Second, we introduce a prescription to incorporate information on perturbative scales without interpreting them as hidden model parameters. We clarify how the two scale prescriptions bias the result towards specific scale choice, and we discuss and compare different Bayesian MHO estimates among themselves and to the traditional scale variation approach. Finally, we provide a practical prescription of how existing perturbative results at the standard scale variation points can be converted to 68%/95% credibility intervals in the Bayesian approach using the new public code `MiH0`.

KEYWORDS: QCD Phenomenology, NLO Computations

ARXIV EPRINT: [2106.04585](https://arxiv.org/abs/2106.04585)

Contents

1	Introduction	1
2	Bayesian methods for observables in quantum field theory	3
2.1	Incorporating scale information into Bayesian inference	3
2.1.1	Prescription 1: scale-marginalisation	4
2.1.2	Prescription 2: scale-averaging	6
2.1.3	Extension to multiple perturbative scales	7
2.2	Connection to prescriptions for choosing the perturbative scales	7
2.3	Connection to MHO estimates from scale variation	9
3	Modelling probability distributions at fixed scales	13
3.1	The mathematical setup	13
3.2	The geometric model	14
3.3	An asymmetric geometric model: the <i>abc</i> -model	16
3.3.1	Sensitivity to the choice of priors	17
4	Quantities without explicit scale dependence	20
4.1	The cusp anomalous dimension	21
4.2	On-shell and $\overline{\text{MS}}$ -scheme quark masses	25
4.3	Electron anomalous magnetic moment and positronium hyperfine-splitting	27
5	Hadronic observables with explicit scale dependence	30
5.1	Higgs production in gluon fusion	31
5.1.1	Sensitivity to the scale interval and scale prescription	33
5.1.2	Accuracy of Gauss-Legendre quadrature rule	34
5.2	Vector-boson fusion Higgs and di-Higgs production	35
5.3	Drell-Yan processes	37
5.4	Deep inelastic scattering	39
5.5	Di-photon production	40
5.6	Gauge boson production in association with a jet	41
5.7	Higgs boson production in association with a jet	42
5.8	Inclusive jet production	43
6	Discussion	45
A	The position of the peak of the distribution	48
A.1	The position of the peak for the scale-marginalisation prescription	49
A.2	The position of the peak for the scale-averaging prescription	49
B	Analytic solution of the <i>abc</i>-model	50
C	Supplementary plots for sections 3 and 4	51

1 Introduction

Perturbative calculations are currently the only way to obtain Standard Model (SM) predictions for collider experiments. The backbone of almost all calculations for hadron collider observables is the QCD factorisation theorem, which relates an observable Σ in proton-proton collisions (e.g., a total cross-section or differential distribution) to the corresponding partonic quantity,

$$\Sigma = \sum_{i,j} \int_0^1 dx_1 dx_2 f_i(x_1, \mu_F) f_j(x_2, \mu_F) \hat{\Sigma}_{ij}(x_1, x_2, \mu_F) + \mathcal{O}\left(\frac{\Lambda_{\text{QCD}}}{Q}\right), \quad (1.1)$$

where $f_i(x, \mu_F)$ denotes the parton distribution function (PDF) for parton i and $\hat{\Sigma}_{ij}$ denotes the observable computed on a collision of the partons i and j . Here μ_F is the factorisation scale, Q denotes a characteristic hard scale of the process, and $\Lambda_{\text{QCD}} \lesssim 1 \text{ GeV}$ is the scale at which QCD becomes strongly coupled.

The partonic coefficients $\hat{\Sigma}_{ij}$ are computable in perturbation theory,

$$\hat{\Sigma}_{ij}(x_1, x_2, \mu_F) \simeq \sum_{k=0}^{\infty} \alpha_s(\mu_R)^{k_0+k} \hat{\Sigma}_{ij}^{(k)}(x_1, x_2, \mu_F, \mu_R), \quad (1.2)$$

where $k_0 \geq 0$ is a non-negative integer, $\alpha_s(\mu_R)$ denotes the renormalised strong coupling constant evaluated at the renormalisation scale μ_R , and $\hat{\Sigma}_{ij}^{(k)}$ is the (next-to-) k -leading order (N^kLO) contribution to the partonic coefficient.¹ One can combine the QCD factorisation theorem in eq. (1.1) with the perturbative expansion of the partonic coefficients in eq. (1.2) to obtain the N^nLO approximation for the observable Σ within (fixed-order) perturbation theory

$$\Sigma \simeq \Sigma_n(\mu_F, \mu_R) := \sum_{k=0}^n \Sigma^{(k)}(\mu_F, \mu_R). \quad (1.3)$$

Computations are typically performed in the $\overline{\text{MS}}$ renormalisation scheme. Although Σ is formally independent of the choice of the factorisation and renormalisation scales, the truncation of the perturbative series introduces a residual scale dependence into Σ_n .

Since a perturbative result is always an approximation of the exact value Σ , it is important to have robust methods to assess the quality of the approximation and to quantify the uncertainty attached to it. Three of the main sources of uncertainty are coming from the values for the strong coupling constant and other input parameters (we collectively refer to them as the parametric uncertainty), the PDFs² and the truncation of the perturbative series. The parametric and PDFs uncertainties are related to their extraction from

¹The ‘ $\simeq$ ’-sign in eq. (1.2) indicates that the perturbative expansion on the right-hand side is in general an asymptotic series with zero radius of convergence. For most phenomenological predictions only very few terms in the perturbative expansion are known (typically not more than the first 3 or 4 terms), and one does not have to worry about the non-convergent nature of the series.

²Note that the PDF uncertainty can itself be seen as a parametric uncertainty, just like the coupling constants or the masses. Indeed, there is no strong conceptual difference between the PDF and, e.g., α_s uncertainties, and often they are treated on the same footing. For our discussion in later sections, we find it convenient to think of these two sources separately. This does, however, not impact any of our results, as our primary focus are MHO uncertainties.

experimental data and the fitting methodology used. Consequently, the PDF and parametric uncertainties are not directly related to the observable Σ . They are computed using observable-independent recipes, cf., e.g., refs. [1–3] (though care is needed as there may be correlations between Σ and the data used to extract the PDFs and α_s). The uncertainty resulting from the truncation of the perturbative series quantifies the effect of the missing higher orders (MHOs) in perturbation theory. Its estimation is one of the main topics of this paper. Under the assumption that the PDF, parametric, and MHO uncertainties are uncorrelated (see [4–7] for discussion of correlation between PDF and MHO uncertainties), they can be added in quadrature to obtain the uncertainty on the N^n LO prediction.³

Currently, there is no reliable way to estimate the size of the MHOs (other than computing the next order in perturbation theory). Since the effect of varying the scales at N^n LO is an N^{n+1} LO effect, one conventionally estimates the size of the MHO terms by studying the variation of $\Sigma_n(\mu_F, \mu_R)$ as the factorisation and renormalisation scales are varied in a certain range around a central scale $\mu_{F/R,\text{cent}}$ corresponding to a characteristic hard scale of the process. A popular choice is the so-called *7-point variation*, where the MHO uncertainty at N^n LO is estimated by the interval

$$\left[\min_{(k_F, k_R) \in S_7} \Sigma_n(k_F \mu_{F,\text{cent}}, k_R \mu_{R,\text{cent}}), \max_{(k_F, k_R) \in S_7} \Sigma_n(k_F \mu_{F,\text{cent}}, k_R \mu_{R,\text{cent}}) \right]. \quad (1.4)$$

(k_F, k_R) take values in the following set:

$$S_7 = \{(1/2, 1/2), (1/2, 1), (1, 1/2), (1, 1), (1, 2), (2, 1), (2, 2)\}. \quad (1.5)$$

This differs from the *9-point variation* where $k_F, k_R \in \{1/2, 1, 2\}$ by the exclusion of large ratios of the scales, i.e., $k_F/k_R = 4$ or $1/4$. We stress that the procedure of estimating the MHO uncertainty from scale variation is completely ad hoc. There is no reason other than experience from NLO and NNLO calculations why the obtained uncertainty interval at N^n LO should capture the actual size of the N^{n+1} LO corrections $\Sigma^{(n+1)}$ or the true value of Σ . Indeed, the higher-order terms probed by varying the scales are universal and predictable from renormalisation group methods. They do not take into account the size of the genuine higher-order corrections, which are process-dependent and independent of the perturbative scales. With advances in precision calculations over the last 5 years, NNLO results are becoming routinely available for many LHC processes, and even N^3 LO corrections to colourless final states have been recently computed. It is, therefore, time to critically assess the MHO uncertainties and to look for alternative prescriptions. A number of attempts have been made in the past to estimate or predict the truncation errors and MHOs, see, e.g., refs. [8–29]. Various prescriptions for reducing or eliminating scale ambiguity have also been discussed broadly in the literature, e.g., refs. [30–47].

Here we focus on a promising, widely applicable approach to estimate the MHO contributions using Bayesian inference on the known fixed-order computations. This idea was

³For a given observable there may be other uncertainties one needs to take into account, e.g., the choice of the renormalisation schemes used for the input parameters, the truncation of the heavy mass expansion etc. How to estimate these uncertainties and combine them is often a matter of debate for specific processes and observables. We do not discuss these issues in this paper as we intend to keep the discussion general.

pioneered by Cacciari and Houdeau (CH) in ref. [48], and developed further in refs. [49–51]. Bayesian inference is a powerful method to construct probability distributions in which Bayes’ theorem is used to update the probability as new information becomes available iteratively. The procedure can be summarised as follows: one starts from some model and the prior distribution of the model parameters. These inputs are purely subjective. The next step is a construction of a posterior distribution for model parameters based on available data. This is the inference step, where one updates prior belief as new data become available. Finally, one constructs the posterior distribution for the next data point by marginalising over the posterior distribution for the model parameters. The main idea of the CH-method is to use Bayesian inference on the known perturbative orders to model the size of the MHOs. The original CH-model can be applied only after one specifies the choices for the perturbative scales. In this paper, we follow up on the recent work in ref. [51], where the CH-approach was extended by the postulate that unphysical scales should be treated on par with other hidden model parameters. According to the lore of Bayesian inference, one expects that the posterior distribution for the perturbative scales peaks at the ‘optimal value’ for the scale choice, and the dependence on the scales is eliminated by marginalisation. However, this is only one example of possible prescriptions, as there is no reason, physical or mathematical, to interpret the perturbative scales as hidden model parameters. One of the main objectives of this paper is to thoroughly scrutinise the consequences of this postulate and to explore alternative approaches.

This paper is organised as follows: in section 2 we discuss two prescriptions of how to incorporate scale dependence into a model for Bayesian inference of fixed-order calculations, and we investigate implicit biases inherent to them. In section 3 we recap the mathematical formalism behind Bayesian inference at fixed scale and introduce a model to account for asymmetric probability distributions. In section 4 we study extensively the model performance for representative perturbative calculations without scale dependence or at fixed scale choice. Finally, in section 5 we study the performance of the two prescriptions to incorporate scale information for an extensive set of collider processes known at NNLO and/or N³LO. We conclude with a discussion in section 6.

2 Bayesian methods for observables in quantum field theory

2.1 Incorporating scale information into Bayesian inference

In this section we present a way to estimate the MHO uncertainty for an observable Σ via Bayesian methods. To be concrete, let us assume that we have computed the first $n + 1$ orders $\Sigma^{(k)}$, $0 \leq k \leq n$, in the perturbative expansion of Σ . Our goal is to construct a probability distribution $\mathcal{P}(\Sigma|\mathbf{\Sigma}_n)$ for Σ given the sequence of perturbative coefficients $\mathbf{\Sigma}_n := (\Sigma^{(0)}, \Sigma^{(1)}, \dots, \Sigma^{(n)})$. This distribution can be used to construct credibility intervals (CIs) for Σ , which may serve as the uncertainty estimate. In particular, we define the $x\%$ CI as $\text{CI}_x = [\Sigma_x^{\text{low}}, \Sigma_x^{\text{up}}]$ where

$$\int_{-\infty}^{\Sigma_x^{\text{low}}} d\Sigma \mathcal{P}(\Sigma|\mathbf{\Sigma}_n) = \int_{\Sigma_x^{\text{up}}}^{\infty} d\Sigma \mathcal{P}(\Sigma|\mathbf{\Sigma}_n) = \frac{1 - \frac{x}{100}}{2}. \quad (2.1)$$

We will mainly consider $x\% \approx 68\%$ and $x\% \approx 95\%$ CIs.

The original approach of ref. [48], known as the *Cacciari-Houdeau* (CH) method, allows one to construct a probability $\mathcal{P}(\Sigma|\Sigma_n)$ in the case where the Σ_n are *numbers*. In perturbation theory, however, the perturbative coefficients Σ_n are *functions* of the perturbative scales, which are arbitrary quantities introduced by perturbation theory without any physical interpretation. Clearly, $\mathcal{P}(\Sigma|\Sigma_n)$ should be independent of the scale choice, but the value of the NⁿLO prediction may depend strongly on this choice, and there is no reason to prefer one scale over another (as long as the scale does not differ too much from the hard scale of the process).

Reference [51] introduced a possible way of incorporating perturbative scales into a Bayesian approach to estimating MHOs (see also ref. [50] for an alternative approach). The approach of ref. [51]⁴ to include the information on the perturbative scales into a Bayesian model is not the only one. Broadly speaking, we identify two different prescriptions of how one can incorporate information on perturbative scales into the construction of a probability distribution:

Prescription 1: Scale-marginalisation (sm). Despite the perturbative scales not being physical, one assigns a Bayesian degree of belief to them. In this way the scales become (hidden) parameters of a Bayesian model. It is then possible to construct $\mathcal{P}(\Sigma|\Sigma_n)$ through Bayesian inference and marginalisation over the perturbative scales. This is the approach advocated in ref. [51].

Prescription 2: Scale-averaging (sa). Since the perturbative scales are devoid of any physical meaning, they cannot be treated as hidden parameters of a Bayesian model, i.e., they cannot be treated as random variables with a statistical interpretation. Instead, the result of a perturbative calculation through NⁿLO should be interpreted as a family of independent predictions for the true value Σ parametrised by the scales. The probability distribution $\mathcal{P}(\Sigma|\Sigma_n)$ is then defined by adding coherently the probabilities computed for each family member.

As we will show later, the two prescriptions of obtaining scale-independent probabilities differ not only in their philosophy, but in practice they can produce different results due to implicit biases towards some scales over others. In the next section we present the mathematical framework to construct the probability distribution $\mathcal{P}(\Sigma|\Sigma_n)$. In order to keep the notations concise, we focus first on the simplified situation where the perturbative coefficients Σ_n depend on a single perturbative scale μ (which we may think of as the renormalisation scale for now).

2.1.1 Prescription 1: scale-marginalisation

We start by reviewing the *scale-marginalisation* (*sm*) prescription of ref. [51]. Since μ is a parameter of the model, the probability distribution $\mathcal{P}(\Sigma|\Sigma_n)$ is obtained by marginalising over μ ;

$$\mathcal{P}(\Sigma|\Sigma_n) = \int d\mu \mathcal{P}(\Sigma, \mu|\Sigma_n) = \int d\mu \mathcal{P}(\Sigma|\Sigma_n, \mu) \mathcal{P}(\mu|\Sigma_n). \quad (2.2)$$

⁴Reference [51] excludes the factorisation scale from this discussion. We will come back to this point in section 2.1.3.

Let us interpret the probabilities in the integrand. For the first factor, we apply Bayes' formula to write

$$\mathcal{P}(\Sigma|\Sigma_n, \mu) = \frac{\mathcal{P}(\Sigma, \Sigma_n|\mu)}{\mathcal{P}(\Sigma_n|\mu)} = \frac{P(\Sigma, \Sigma_n(\mu))}{P(\Sigma_n(\mu))} = P(\Sigma|\Sigma_n(\mu)). \quad (2.3)$$

The interpretation of this equation is as follows: since μ is a parameter of the model, whenever we evaluate a conditional probability *for a fixed value of μ* , we can replace the perturbative coefficients Σ_n (which are functions) by their values $\Sigma_n(\mu)$ for fixed μ (which are numbers). In the following we always indicate by $\mathcal{P}$ a probability distribution which takes as arguments the perturbative coefficients Σ_n (which are functions), while P denotes a probability distribution on perturbative coefficients computed at a fixed scale μ . For now we assume that P is given. In section 3 we discuss a framework how a distribution P can be constructed.

The second factor in the integrand of eq. (2.2) is the posterior distribution for μ , and we have

$$\mathcal{P}(\mu|\Sigma_n) = \frac{\mathcal{P}(\mu, \Sigma_n)}{\mathcal{P}(\Sigma_n)} = \frac{\mathcal{P}(\Sigma_n|\mu) P_0(\mu)}{\int d\mu' \mathcal{P}(\Sigma_n|\mu') P_0(\mu')} = \frac{P(\Sigma_n(\mu)) P_0(\mu)}{\int d\mu' P(\Sigma_n(\mu')) P_0(\mu')}, \quad (2.4)$$

where $P_0(\mu)$ is an (arbitrary) prior distribution on the perturbative scale μ . In ref. [51] it was argued that a physically well-motivated choice is a flat distribution in the logarithm of the scale,

$$P_0(\mu) := \frac{1}{2\mu \log F} \Theta \left(\log F - \left| \log \frac{\mu}{\mu_0} \right| \right), \quad (2.5)$$

where $F > 1$ is a constant and Θ denotes the Heaviside step function;

$$\Theta(x) = \begin{cases} 1, & \text{if } x > 0, \\ 0, & \text{otherwise.} \end{cases} \quad (2.6)$$

The central scale μ_0 is in principle arbitrary, but assumed to be of the order of the hard scale of the process. Putting everything together, we see that $\mathcal{P}(\Sigma|\Sigma_n)$ can be expressed in terms of the probability distribution P and the prior P_0 on the scale:

$$\mathcal{P}(\Sigma|\Sigma_n) \equiv \mathcal{P}_{\text{sm}}(\Sigma|\Sigma_n) := \frac{\int d\mu P(\Sigma|\Sigma_n(\mu)) P(\Sigma_n(\mu)) P_0(\mu)}{\int d\mu' P(\Sigma_n(\mu')) P_0(\mu')}. \quad (2.7)$$

Equation (2.7) allows one to relate $\mathcal{P}_{\text{sm}}(\Sigma|\Sigma_n)$ to two other probability distributions (in addition to the prior on μ). The first one is the probability distribution $P(\Sigma|\Sigma_n(\mu))$ for the (true) value of the observable Σ given the $n+1$ first perturbative orders calculated at the scale μ . The second one is the probability distribution $P(\Sigma_n(\mu))$, which is related through Bayes' theorem to $P(\Sigma^{(n)}(\mu)|\Sigma_{n-1}(\mu))$. The latter computes the probability distribution for the NⁿLO corrections $\Sigma^{(n)}(\mu)$ evaluated at the scale μ , given the n first perturbative orders evaluated at that scale. The two distributions $P(\Sigma|\Sigma_n(\mu))$ and $P(\Sigma^{(n+1)}(\mu)|\Sigma_n(\mu))$ are the answers to two different mathematical questions, and there is a priori no reason why they should be related. However, if perturbation theory is applicable, then the true value of Σ should be well approximated by $\Sigma_{n+1}(\mu)$ (at least for reasonable choices of

μ). We therefore expect that also the distribution $P(\Sigma|\Sigma_n(\mu))$ is well approximated by $P(\Sigma^{(n+1)}(\mu)|\Sigma_n(\mu))$. Substituting $\Sigma^{(n+1)}(\mu) \approx \Sigma - \Sigma_n(\mu)$ in the latter, we expect that, if perturbation theory is applicable, then

$$P(\Sigma|\Sigma_n(\mu)) \approx P(\Sigma - \Sigma_n(\mu)|\Sigma_n(\mu)). \quad (2.8)$$

For a more detailed derivation of this approximation in a Bayesian framework we refer to ref. [51]. Using this approximation, we can cast eq. (2.7) into the form:

$$\mathcal{P}_{\text{sm}}(\Sigma|\Sigma_n) \approx \frac{\int d\mu P(\Sigma - \Sigma_n(\mu)|\Sigma_n(\mu)) P(\Sigma_n(\mu)) P_0(\mu)}{\int d\mu' P(\Sigma_n(\mu')) P_0(\mu')}. \quad (2.9)$$

We stress that, unlike indicated by the notation in eq. (2.8), the probability distributions in the left and right-hand sides of eqs. (2.7) and (2.8) are actually *distinct*: $P(\Sigma|\Sigma_n(\mu))$ is the conditional distribution for the true value Σ of the observable, while $P(\Sigma - \Sigma_n(\mu)|\Sigma_n(\mu))$ is the distribution for the value of the N^{n+1} LO correction $\Sigma^{(n+1)}(\mu) \approx \Sigma - \Sigma_n(\mu)$ computed at the scale μ . We will always assume that available N^n LO results sufficiently well approximate Σ and that we are in the perturbative regime where eq. (2.8) holds, so we do not distinguish the notation for the two distributions.

2.1.2 Prescription 2: scale-averaging

In this section we present an alternative way to include scale information into the model which does not require the perturbative scale μ to be promoted to a model parameter with a statistical or physical meaning. We call this method the *scale-averaging* (*sa*) prescription. The starting point is to interpret μ as a parameter defining a family of perturbative predictions which approximate Σ at N^n LO in perturbation theory. For each member of this family, i.e., for each fixed value μ of the scale, we can compute the probability $P(\Sigma|\Sigma_n(\mu))$ (as before we assume again that P is given). We then combine this family of predictions into a single distribution $\mathcal{P}(\Sigma|\Sigma_n)$ as a weighted average of probabilities

$$\mathcal{P}(\Sigma|\Sigma_n) \equiv \mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n) := \int d\mu w(\mu) P(\Sigma|\Sigma_n(\mu)). \quad (2.10)$$

The form of the weighting function $w(\mu)$ is arbitrary, but it must be chosen such that $\mathcal{P}(\Sigma|\Sigma_n)$ satisfies all the axioms of a normalised probability distribution. In particular, $w(\mu)$ must itself be normalised:

$$\int d\mu w(\mu) = 1. \quad (2.11)$$

Since one expects that the perturbative scale μ should be chosen close to the characteristic hard scale of the process in order to ensure good perturbative convergence, we postulate a window function around a fixed scale $\mu_0 \sim Q$:

$$w(\mu) := \frac{1}{2\mu \log F} \Theta \left(\log F - \left| \log \frac{\mu}{\mu_0} \right| \right), \quad (2.12)$$

where $F > 1$ is a constant. We note that our choice for $w(\mu)$ is identical to the prior on the scale in the scale marginalisation in eq. (2.5). The interpretation, however, is different:

in the previous section the scale μ is a model parameter with a statistical interpretation that leads to eq. (2.7). In eq. (2.10), the scale has no statistical interpretation, and is simply a parameter defining a family of distributions, which are then combined into a single distribution with weighting function $w(\mu)$ according to eq. (2.10). With the weight given in eq. (2.12), this corresponds to a uniform average in the logarithm of the scale over the interval $\mu_0/F < \mu < F\mu_0$. Using the approximation in eq. (2.8) we obtain the final expression

$$\mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n) \approx \int d\mu w(\mu) P(\Sigma - \Sigma_n(\mu)|\Sigma_n(\mu)). \quad (2.13)$$

2.1.3 Extension to multiple perturbative scales

In the case where the perturbative coefficients Σ_n are independent of the unphysical perturbative scale μ , i.e., $\Sigma_n(\mu) = \Sigma_n(\mu_0)$ for some μ_0 , the distributions $\mathcal{P}$ reduce to P :

$$\mathcal{P}(\Sigma|\Sigma_n) = P(\Sigma|\Sigma_n(\mu_0)), \quad \text{if } \Sigma_n(\mu) = \Sigma_n(\mu_0). \quad (2.14)$$

To extend the discussion to more scales μ_i , $1 \leq i \leq s$ we introduce a multivariate prior $P_0(\mu_1, \dots, \mu_s)$ or weighting function $w(\mu_1, \dots, \mu_s)$, and we integrate over each scale μ_i separately.

We need to comment on hadronic observables, which always depend on the factorisation scale μ_F . In ref. [51] the factorisation scale was explicitly excluded from the marginalisation procedure described in section 2.1.1. Indeed, the factorisation scale also appears in the PDFs. It is not clear what is the correlation and how the PDF uncertainty impacts the probability distribution $\mathcal{P}(\Sigma|\Sigma_n)$, which is supposed to represent the probability distribution for the true value of the observable and should supersede and incorporate all other uncertainties. Here we aim for something much simpler. We want to use the distribution $\mathcal{P}(\Sigma|\Sigma_n)$ to construct credibility intervals that serve as estimators for the MHO uncertainty. In this sense, the role of $\mathcal{P}(\Sigma|\Sigma_n)$ is similar to scale variation as a tool to estimate the size of the MHOs. Like in the case of scale variation, we will assume that the MHO uncertainty is independent of the PDF uncertainty, and the two can be added in quadrature.

2.2 Connection to prescriptions for choosing the perturbative scales

In the previous section we have presented two methods to construct a scale-independent probability distribution $\mathcal{P}(\Sigma|\Sigma_n)$. The two methods differ not only by a different philosophy of how the unphysical perturbative scales are included, but can also lead to different results for the inferred probability distributions. Here we show that under rather general conditions the position of the peak of $\mathcal{P}(\Sigma|\Sigma_n)$ as a function of Σ will be biased towards two different scales μ_{FAC} and μ_{PMS} :

- The point of *fastest apparent convergence* (FAC) [52, 53] of an observable computed at NⁿLO is the scale μ_{FAC} such that

$$\Sigma_n(\mu_{\text{FAC}}) = \Sigma_{n-1}(\mu_{\text{FAC}}), \quad (2.15)$$

where we used the notation $\Sigma_n(\mu) := \Sigma_n(\mu, \mu)$. Note that the value of μ_{FAC} depends on the order at which it is computed, and there may be more than one solution to eq. (2.15) at a given order in perturbation theory.

- The *principle of minimal sensitivity* (PMS) [53] aims at selecting a scale μ_{PMS} at $N^n\text{LO}$ that maximises the stability of the observable under changes of the scale

$$\frac{\partial}{\partial \mu} \Sigma_n(\mu) \Big|_{\mu=\mu_{\text{PMS}}} = 0. \quad (2.16)$$

Just like the FAC point, the PMS point depends on the order on which it is computed and may not be unique at a given order in perturbation theory.

Consider now a broad class of Bayesian inference models satisfying the following conditions:

1. We work with the prior in eq. (2.5) or the weighting function in eq. (2.12).
2. $P(\Sigma^{(n+1)}(\mu) | \Sigma_n(\mu))$ is *symmetric*, i.e., it is insensitive to the signs of the correction:

$$P(\Sigma^{(n+1)}(\mu) | \Sigma_n(\mu)) = P(|\Sigma^{(n+1)}(\mu)| | |\Sigma^{(0)}(\mu)|, \dots, |\Sigma^{(n)}(\mu)|). \quad (2.17)$$

3. The distributions $P(\Sigma^{(n+1)}(\mu) | \Sigma_n(\mu))$ have a single peak in $\Sigma^{(n+1)}(\mu)$ for each fixed value of μ , and the peak is reached for some $\mu \in [\mu_0/F, \mu_0 F]$.
4. The peak is more and more pronounced as n increases, i.e., as more information on the progression of the perturbative series becomes known.

Note that these assumption imply that the peak of $P(\Sigma^{(n+1)}(\mu) | \Sigma_n(\mu))$ is located at $\Sigma^{(n+1)}(\mu) = 0$, for all μ . The aforementioned assumptions are reasonable, and they are satisfied for some of the models considered in ref. [51], in particular by the geometric model and the scale-variation model of ref. [51] (see also section 3.2). In appendix A we show that any model satisfying the assumptions above will also satisfy the following property for the peak of the distribution $\mathcal{P}(\Sigma | \Sigma_n)$:

- If there is a single $\mu_{\text{FAC}} \in [\mu_0/F, \mu_0 F]$, then $\mathcal{P}_{\text{sm}}(\Sigma | \Sigma_n)$ peaks at $\Sigma = \Sigma_n(\mu_{\text{FAC}})$, i.e., at the $N^n\text{LO}$ prediction for the observable evaluated at the point of fastest apparent convergence.
- If there is a single $\mu_{\text{PMS}} \in [\mu_0/F, \mu_0 F]$, then $\mathcal{P}_{\text{sa}}(\Sigma | \Sigma_n)$ peaks at $\Sigma = \Sigma_n(\mu_{\text{PMS}})$, i.e., at the $N^n\text{LO}$ prediction for the observable evaluated at the point of minimal sensitivity.

As a consequence, the sensitivity of the probability distributions and CIs on the choice of the scale variation interval (given by F) depends on the presence (or absence) of FAC and PMS points in the interval for the sm- and sa-prescriptions respectively. We will illustrate this on concrete examples in section 5.

Let us conclude this section with an important comment about the sm-prescription. In the sm-prescription, the perturbative scales are treated as model parameters with a

statistical interpretation. Bayesian inference will then select the ‘best’ value for the perturbative scales according to the model, see also the discussion in ref. [51]. At high enough perturbative orders, the dependence on the priors is expected to be mild, so that the conclusions derived from the inference are robust and (to a certain extent) prior-independent. The results of this section indicate that such a Bayesian interpretation of the perturbative scales should be taken with a grain of salt: for symmetric models with the prior in eq. (2.5) (which applies in particular to some of the models of ref. [51]), as soon as the range $[\mu_0/F, \mu_0 F]$ contains an FAC point, the scale selected by the inference is biased: it will *automatically* select the value of the observable at the FAC point as the most likely value. In other words, such a model will *always* select the FAC point as its ‘best’ estimate for the scales, independently of the actual progression of the perturbative series. Note that the previous discussion does not apply to the sa-prescription, because the perturbative scales are not assigned a probabilistic interpretation (in particular, the weight $w(\mu)$ in eq. (2.12) is not a probability distribution). In fact we will see examples in section 5 where the CIs for the sa-prescription increase monotonically with F .

2.3 Connection to MHO estimates from scale variation

One may wonder about the practicality of the prescriptions in sections 2.1.1 and 2.1.2, because evaluating the integrals of the scale could be rather slow, as it may require to evaluate the perturbative approximation of Σ for many values of the scales μ_i . In many examples of higher-order computations for the LHC, the dependence on the factorisation and renormalisation scales is rather smooth. We can then approximate the integrals over scale in eqs. (2.9) and (2.13) by a low-order quadrature rule.

We focus for now on observables depending on a single scale μ , and we comment on multiple scales later. It is particularly convenient to choose the quadrature points to coincide with the ones used in the conventional scale variation, namely $\mu/\mu_0 \in \{1/2, 1, 2\}$.⁵ For the three-point Gauss-Legendre quadrature rule this can be achieved by using the value of $F = 2\sqrt{5/3} \approx 2.45$ for $w(\mu)$ in eq. (2.12). Then the integral in eq. (2.13) can be approximated as

$$\int d\mu w(\mu) f(\mu) \approx w_{-1} f(\mu_0/2) + w_0 f(\mu_0) + w_{+1} f(2\mu_0), \quad (2.18)$$

where the weights w_i for the Gauss-Legendre quadrature are

$$w_0 = \frac{8}{18} \quad \text{and} \quad w_{\pm 1} = \frac{5}{18}. \quad (2.19)$$

This quadrature achieves high accuracy for smooth functions in the interval $-\log F < \log \mu/\mu_0 < \log F$ (it is in fact exact if $f(\mu)$ is polynomial in $\log \mu$ up to degree 5; note that hadronic observables are never of this form, because some logarithms have been resummed

⁵We note that the functional dependence on μ_R is determined by the running of the coupling constant and the 9-point μ_F, μ_R grid can also be reconstructed from the 7-points given in eq. (1.5) in some cases. For a dynamical scale choice, this reconstruction is typically only feasible for distributions that allow for the determination of the numerical value of the scale that contributes to a specific bin within the uncertainties of the bin size.

Figure 1. Left: the integration regions for different quadrature rules using scale variation points. Right: sketch of probability distributions at scale variation points when the width is smaller than scale variation interval (shaded) and when it is larger (dashed lines).

to all orders into the PDFs and $\alpha_s(\mu)$). Similar formulas can of course be derived for quadrature rules of higher order and multi-dimensional μ integrals. The advantage of the three-point rule in eq. (2.18) is that it allows one to easily convert the available data from the scale-variation approach into an approximation for the Bayesian probability distribution, which is scale-marginalised or scale-averaged over the interval $\mu_0/F < \mu < \mu_0 F$ with $F \approx 2.45$. This approximation is valid as long as the three-point quadrature rule gives a good approximation of the integral. In applications, the dependence on the scale is usually smooth, so that low-order quadrature rules are expected to give reasonable approximations. We will illustrate this on examples in section 5. One could also restrict the integration region to the interval $[\mu_0/2, 2\mu_0]$, i.e. $F = 2$, by using a quadrature rule where the nodes coincide with the end-points of the integration, e.g., the Gauss-Lobatto quadrature rule (accurate at cubic order). We compare the $\log \mu/\mu_0$ ranges of different quadrature rules using the same scale variation points in the left panel of figure 1.

The quadrature rule in eq. (2.18) allows us to obtain an approximate bound on the probability that the true value lies in the range obtained from conventional scale variation. We will first discuss the case of an observable depending on a single perturbative scale μ , and generalise the result later. Define the lower and upper values of the scale variation around some central scale μ_0 as

$$\begin{aligned}\Sigma_{\min} &:= \Sigma_n(k_{\min}\mu_0) = \min_{k \in \{1/2, 1, 2\}} \Sigma_n(k\mu_0), \\ \Sigma_{\max} &:= \Sigma_n(k_{\max}\mu_0) = \max_{k \in \{1/2, 1, 2\}} \Sigma_n(k\mu_0).\end{aligned}\tag{2.20}$$

Consider the sa-prescription and a model such that

1. the model is symmetric,
2. $P(\Sigma^{(n+1)}(\mu)|\Sigma(\mu))$ peaks only at $\Sigma^{(n+1)}(\mu) = 0$, for all values of μ ,
3. the three-point quadrature in eq. (2.18) gives a good approximation of the integral.

If these assumptions are satisfied, then we can show that the probability that $\Sigma \in [\Sigma_{\min}, \Sigma_{\max}]$ is bounded from above:

$$\mathcal{P}_{\text{sa}}(\Sigma \in [\Sigma_{\min}, \Sigma_{\max}] | \Sigma_n) \lesssim 1 - \frac{1}{2}(w_{\min} + w_{\max}), \quad (2.21)$$

where $w_{\min}$ is the quadrature weight associated to the scale $\mu = k_{\min}\mu_0$ through eq. (2.18), and similarly for $w_{\max}$. The $\lesssim$ sign reminds us that the bound is only as accurate as the assumptions above, in particular the validity of eq. (2.18). To understand the origin of this approximate bound, we first note that

$$\mathcal{P}_{\text{sa}}(\Sigma \in [\Sigma_{\min}, \Sigma_{\max}] | \Sigma_n) = \Phi(\Sigma_{\max}) - \Phi(\Sigma_{\min}), \quad (2.22)$$

where $\Phi(\Sigma)$ is the cumulative distribution $\Phi(\Sigma) := \int_{-\infty}^{\Sigma} d\Sigma' \mathcal{P}_{\text{sa}}(\Sigma' | \Sigma_n)$. Using eqs. (2.13) and (2.18), we express $\Phi(\Sigma)$ as a sum of cumulative distributions $\varphi(\Sigma, \mu) := \int_{-\infty}^{\Sigma} d\Sigma' P(\Sigma' - \Sigma_n(\mu) | \Sigma_n(\mu))$ for a fixed choice of the scale:

$$\begin{aligned} \Phi(\Sigma_{\min}) &\approx w_{-1} \varphi(\Sigma_{\min}, \mu_0/2) + w_0 \varphi(\Sigma_{\min}, \mu_0) + w_{+1} \varphi(\Sigma_{\min}, 2\mu_0) \\ &\gtrsim w_{\min} \varphi(\Sigma_{\min}, k_{\min}\mu_0), \end{aligned} \quad (2.23)$$

where the last step follows from the fact that all three terms are positive. An illustration of the three functions $P(\Sigma - \Sigma_n(k\mu_0) | \Sigma_n(k\mu_0))$, $k \in \{1/2, 1, 2\}$, is shown in figure 1 for two scenarios: when the distributions are narrow compared to the scale variation interval $[\Sigma_{\min}, \Sigma_{\max}]$ (shaded areas) and when the distributions are wider than $[\Sigma_{\min}, \Sigma_{\max}]$ (dashed lines). Note that we do not know which value of k corresponds to which distribution in figure 1. However, from the assumptions above it follows that only one of them peaks at $\Sigma_{\min}$, namely the one corresponding to the value of $k = k_{\min}$. Since the distribution is by assumption symmetric around $\Sigma = \Sigma_{\min}$, we must have $\varphi(\Sigma_{\min}, k_{\min}\mu_0) = \frac{1}{2}$. Hence, we obtain

$$\Phi(\Sigma_{\min}) \gtrsim \frac{1}{2} w_{\min}. \quad (2.24)$$

Using a similar argument, we can conclude

$$\Phi(\Sigma_{\max}) \lesssim 1 - \frac{1}{2} w_{\max}. \quad (2.25)$$

Equation (2.21) then follows upon inserting the bounds for $\Phi(\Sigma_{\min})$ and $\Phi(\Sigma_{\max})$ into eq. (2.22).

We can ask when the approximate bound in eq. (2.21) is saturated. It is easy to see that the bound for the cumulative distribution in eq. (2.24) is saturated if the three distributions $P(\Sigma - \Sigma_n(k\mu_0) | \Sigma_n(k\mu_0))$, $k \in \{1/2, 1, 2\}$ are narrow compared to their separation, because in that scenario we have $\Phi(\Sigma_{\min}) \approx \frac{1}{2} w_{\min}$. (see figure 1). A similar argument holds for the bound in eq. (2.25), $\Phi(\Sigma_{\max}) \approx 1 - \frac{1}{2} w_{\max}$. Hence, we conclude that the approximate bounds on the probability for $[\Sigma_{\min}, \Sigma_{\max}]$ are saturated for observables and models for which the distributions $P(\Sigma - \Sigma_n(k\mu_0) | \Sigma_n(k\mu_0))$, $k \in \{1/2, 1, 2\}$ are well separated.

We can repeat exactly the same argument for observables depending on multiple scales $(\mu_1, \dots, \mu_s)$ by using the same quadrature rule s times for scale-averaging for each scale μ_i

Quadrature rule	β_1	β_2
Gauss-Legendre	$13/18 \approx 72\%$	$\approx 92\%$
Gauss-Lobatto	$5/6 \approx 83\%$	$\approx 97\%$
Equidistant	$2/3 \approx 66\%$	$\approx 89\%$

Table 1. Approximate bounds on the probability for interval $[\Sigma_{\min}, \Sigma_{\max}]$.

in eq. (2.23). Each cumulative distribution $\varphi(\Sigma, \mu_1, \dots, \mu_s)$ in the sum is multiplied by a product of s quadrature weights. It is then straightforward to show that the bound becomes

$$\mathcal{P}_{\text{sa}}(\Sigma \in [\Sigma_{\min}, \Sigma_{\max}] | \Sigma_n) \lesssim \beta_s, \text{ with } \beta_s := 1 - \frac{1}{2^s} (w_{\min} + w_{\max})^s. \quad (2.26)$$

Let us discuss some implications of the approximate bound in eq. (2.26). Under the assumptions spelled out above, eq. (2.26) establishes a connection between the interval obtained by varying the scale up and down by a factor of 2 and the Bayesian probability distribution constructed from the sa-prescription. The value of the approximate bound depends on the quadrature weights used to approximate the integral. For the three quadrature rules illustrated in figure 1, the values of the bound are shown in table 1 for $s = 1, 2$. Note that different quadratures correspond to different intervals over which the scale-averaging is done. In addition, the accuracy of the approximation in eq. (2.18) also depends on the order of the quadrature rule. Therefore the approximate bounds in table 1 depend on how well the assumptions stated at the beginning of this section are satisfied. In particular, it *does not* allow one to attach a definite probability that the true value lies in the range given by the scale variation. However, we can show that the $(100\beta_s)\%$ credibility interval $\text{CI}_{100\beta_s} = [\Sigma_{100\beta_s}^{\text{low}}, \Sigma_{100\beta_s}^{\text{up}}]$ will contain $[\Sigma_{\min}, \Sigma_{\max}]$ and the two intervals (approximately) agree whenever the approximate bound is saturated. Indeed, from the definition of the credibility interval in eq. (2.1) it follows that

$$\Phi(\Sigma_{100\beta_s}^{\text{low}}) = \frac{1}{2}(1 - \beta_s) \text{ and } \Phi(\Sigma_{100\beta_s}^{\text{up}}) = \frac{1}{2}(1 + \beta_s). \quad (2.27)$$

Equations (2.24) and (2.25) and their generalisations for $s > 1$ then imply:

$$\Phi(\Sigma_{100\beta_s}^{\text{low}}) \lesssim \Phi(\Sigma_{\min}) \leq \Phi(\Sigma_{\max}) \lesssim \Phi(\Sigma_{100\beta_s}^{\text{up}}), \quad (2.28)$$

and so the monotonicity of the cumulative distribution implies

$$\Sigma_{100\beta_s}^{\text{low}} \lesssim \Sigma_{\min} \leq \Sigma_{\max} \lesssim \Sigma_{100\beta_s}^{\text{up}}, \quad (2.29)$$

and the claim follows.

The previous discussion has an important consequence: assume that we are in a situation with $s = 2$ scales. Then we conclude from table 1 that the 9-point scale variation interval satisfies

$$[\Sigma_{\min}, \Sigma_{\max}] \subset \text{CI}_{100\beta_2} \approx \text{CI}_{89-97}. \quad (2.30)$$

Clearly, the 7-point interval is also contained in the same interval. In the case where the approximate bound is saturated, we expect that for two scales (μ_F, μ_R) the intervals

obtained from scale variation are comparable to $\approx 95\%$ credibility intervals. In practise, we often find that observables are much less sensitive to one of the scales, namely, μ_F , and we are effectively dealing with $s = 1$ case, where $\beta_1 \approx 66\text{--}83\%$. We will illustrate in section 5 that the conventional scale variation intervals are often similar in size to CI_{68} using Bayesian methods. The arguments presented in this section show that this is not a coincidence. Nevertheless we reiterate that the derivation of these bounds relies on specific assumptions. In particular, the sm-prescription and models with asymmetric probability distributions do not need comply with these requirements and the bounds can be broken. However, as we will see in the examples in section 5, in many cases the size of the CIs is similar for symmetric and asymmetric models.

3 Modelling probability distributions at fixed scales

In the previous section we have discussed two ways to define a probability distribution $\mathcal{P}(\Sigma|\Sigma_n)$ in the context of perturbation theory. In both approaches the computation of the distribution $\mathcal{P}(\Sigma|\Sigma_n)$ is reduced to an integral over the perturbative scale μ . The integrand involves, on the one hand, the prior $P_0(\mu)$ in eq. (2.5) or the weighting function $w(\mu)$ in eq. (2.12), and, on the other hand, the probability $P(\Sigma|\Sigma_n(\mu))$ evaluated at a fixed value μ of the perturbative scale. The latter is related through eq. (2.8) to the probability distribution $P(\Sigma^{(n+1)}(\mu)|\Sigma_n(\mu))$ for the size of the next perturbative order at the scale μ , given the $n + 1$ first orders computed at the same scale. Various models for $P(\Sigma^{(n+1)}(\mu)|\Sigma_n(\mu))$ have been presented in the literature, in particular the CH-method of ref. [48] and its generalisations of ref. [51]. All these examples define symmetric distributions in the sense of eq. (2.17). In section 2.2 we have argued that symmetric distributions always lead to very special distributions $\mathcal{P}(\Sigma|\Sigma_n)$ after integration over μ , with peaks that are correlated to the FAC and PMS points. This may introduce unwanted bias into the model.

The goal of this section is to introduce a model that is not symmetric, which is a necessary condition for distributions not necessarily peaked at the FAC or PMS points. We start by presenting a framework to define different models for $P(\Sigma^{(n+1)}(\mu)|\Sigma_n(\mu))$. This framework contains, in particular, the CH-method of ref. [48] and its generalisations of ref. [51] as special cases. We then use this framework to define a new model that takes into account the signs of the corrections and leads to asymmetric distributions.

3.1 The mathematical setup

Consider the set consisting of elements $(\delta, \mathbf{p})$, where $\mathbf{p} \in \mathbb{R}^p$ is a vector and $\delta = (\delta_0, \delta_1, \dots)$ is an infinite sequence of real numbers representing the infinite series $\sum_{k=0}^{\infty} \delta_k$. We prefer to think of the sequence δ rather than the infinite series, because the order of the δ_k 's matters.

We make the following assumptions:

1. The parameters $\mathbf{p}$ are hidden, and we do not have access to them. Instead, they are distributed according to a known prior distribution $P_0(\mathbf{p})$.
2. The sequence δ and the hidden parameters $\mathbf{p}$ are not independent. For a given value of $\mathbf{p}$, we know the probability that the first $n + 1$ elements of δ take a certain value,

i.e., for each $n \geq 0$, if $\boldsymbol{\delta}_n := (\delta_0, \delta_1, \dots, \delta_n)$, we know the conditional probability $P^{(n)}(\boldsymbol{\delta}_n|\mathbf{p})$.

3. The conditional probabilities $P^{(n)}(\boldsymbol{\delta}_n|\mathbf{p})$ for different values of n are related by marginalisation:

$$P^{(n-1)}(\boldsymbol{\delta}_{n-1}|\mathbf{p}) = \int d\delta_n P^{(n)}(\boldsymbol{\delta}_n|\mathbf{p}), \text{ for all } n > 0. \quad (3.1)$$

We refer to a choice of $(P_0, P^{(0)}, P^{(1)}, P^{(2)}, \dots)$ as a *model* for sequences $(\boldsymbol{\delta}, \mathbf{p})$. We always assume that all probabilities are normalised, and we will often keep the superscript on $P^{(n)}$ implicit and use the simplified notation $P(\boldsymbol{\delta}_n|\mathbf{p}) := P^{(n)}(\boldsymbol{\delta}_n|\mathbf{p})$.

Assume that we are given the first $n+1$ terms $\boldsymbol{\delta}_n$ of $\boldsymbol{\delta}$. The main idea is that, within a given model, we can use Bayesian inference to obtain a conditional probability distribution $P(\delta_{n+1}|\boldsymbol{\delta}_n)$ for the next term δ_{n+1} in the sequence $\boldsymbol{\delta}$, given the first $n+1$ terms $\boldsymbol{\delta}_n$. The probability distribution $P(\delta_{n+1}|\boldsymbol{\delta}_n)$ is given by Bayes' formula:

$$P(\delta_{n+1}|\boldsymbol{\delta}_n) = \frac{P(\boldsymbol{\delta}_{n+1})}{P(\boldsymbol{\delta}_n)}. \quad (3.2)$$

The probability distributions for $P(\boldsymbol{\delta}_n)$ and $P(\boldsymbol{\delta}_{n+1})$ are themselves obtained from Bayes' formula:

$$P(\boldsymbol{\delta}_n) = \int d^p \mathbf{p} P(\boldsymbol{\delta}_n|\mathbf{p}) P_0(\mathbf{p}). \quad (3.3)$$

We refer to a model as *symmetric* if it satisfies (cf. eq. (2.17))

$$P(\delta_0, \dots, \delta_n) = P(|\delta_0|, \dots, |\delta_n|). \quad (3.4)$$

A symmetric model always predicts a vanishing median and expectation value for δ_{n+1} given $\boldsymbol{\delta}_n$,

$$\langle \delta_{n+1} \rangle_{\boldsymbol{\delta}_n} = \delta_{n+1}^{\text{median}} = 0, \quad \text{if the model is symmetric,} \quad (3.5)$$

and all CIs computed from a symmetric model are symmetric around $\delta_{n+1} = 0$.

3.2 The geometric model

The CH model assumes a universal bound c for all coefficients $\delta_k(\mu) = \Sigma^{(k)}(\mu)/\alpha_s^k(\mu)$ for a fixed value of μ . In ref. [51] a variation of the CH model, called the *geometric model*, was introduced to take into account the fact that the expansion parameter does not need to be equal to α_s .⁶ The geometric model considers a sequence of coefficients $\delta_k(\mu) = \Sigma^{(k)}(\mu)/\Sigma^{(0)}(\mu)$ (so that $\delta_0(\mu) = 1$), whose magnitude is bounded from above by a geometric series. Consequently, the model has two hidden parameters $\mathbf{p} = (a, c)$, and the coefficients of the sequence $\boldsymbol{\delta}(\mu)$ must satisfy the constraint (for simplicity we suppress the dependence on the scale):

$$\frac{|\delta_k|}{a^k} \leq c, \quad \text{for all } k \geq 0. \quad (3.6)$$

⁶In ref. [50] another modification of the CH model was introduced to take this effect into account. We do not explore the method of ref. [50] in this paper.

The constraints for $k = 0, 1$ imply $c > 1$ and $a > 0$. Since we do not make any other assumption about the coefficients, we choose a flat distribution that takes into account eq. (3.6) [51]:

$$\begin{aligned} P_{\text{geom}}(\boldsymbol{\delta}_n|a, c) &= \prod_{k=0}^m \frac{1}{2a^k c} \Theta\left(c - \frac{|\delta_k|}{a^k}\right) \\ &= \frac{1}{(2c)^{n+1} a^{n(n+1)/2}} \Theta\left(c - \max\left(\frac{|\delta_0|}{a^0}, \dots, \frac{|\delta_n|}{a^n}\right)\right). \end{aligned} \quad (3.7)$$

The model is clearly symmetric. For the prior on the hidden parameters, a and c are assumed independent, $P_0(a, c) = P_0(a)P_0(c)$. The prior on c is given by

$$P_0(c) = \frac{\varepsilon}{c^{1+\varepsilon}} \Theta(c - 1). \quad (3.8)$$

Here $\varepsilon > 0$ is a constant introduced to obtain a normalisable prior. The prior on a is [51]:

$$P_0(a) = (1 + \omega) (1 - a)^\omega \Theta(a) \Theta(1 - a), \quad (3.9)$$

where $\omega > 0$ suppresses the prior at $a = 1$. Note that ε and ω are constants, and they are not considered hidden parameters of the model. For comparisons with the asymmetric model defined in the next section, we will use $\varepsilon = 0.1$ and $\omega = 1$. The form of eq. (3.7) and the priors are simple enough that the integral in eq. (3.3) can be performed in closed form, see ref. [51].

Let us conclude this section with some comments. First, we note that the geometric model contains the original CH model as a special case, where $a = \alpha_s(\mu)$, i.e., the prior for a is

$$P_{0,\text{CH}}(a) = \delta(a - \alpha_s(\mu)), \quad (3.10)$$

so that

$$P_{\text{CH}}(\boldsymbol{\delta}_n|c) = \int da P_{\text{geom}}(\boldsymbol{\delta}_n|a, c) P_{0,\text{CH}}(a) = P_{\text{geom}}(\boldsymbol{\delta}_n|\alpha_s(\mu), c). \quad (3.11)$$

Second, so far we have chosen the coefficients $\delta_k(\mu)$ to be real numbers. We can assume that this sequence was obtained by evaluating some sequence of functions $\Sigma^{(k)}$ at some fixed value μ , i.e., $\delta_k(\mu) = \Sigma^{(k)}(\mu)/\Sigma^{(0)}(\mu)$. Reference [51] proposes a variant where the probability distribution takes input from other values of the functions $\Sigma^{(k)}$. The *scale variation model* of ref. [51] considers sequences $\boldsymbol{\delta}(\mu)$ whose coefficients satisfy a bound of the form:

$$\frac{|\delta_k(\mu)|}{r_k(\mu)} \leq c, \quad \text{for all } k \geq 0, \quad (3.12)$$

where $r_k(\mu)$ encodes information on the derivative of $\Sigma_k(\mu)$ with respect to μ . The probability distribution for *fixed* μ is related to P_{CH} via the relation (with α_s replaced by unity):

$$P_{\text{sc.-var.}}(\boldsymbol{\delta}_n(\mu)|c) = P_{\text{CH}}\left(\frac{\delta_0(\mu)}{r_0(\mu)}, \dots, \frac{\delta_n(\mu)}{r_n(\mu)} \middle| c\right) \prod_{k=0}^n \frac{1}{r_k(\mu)}. \quad (3.13)$$

In this case the prior for c is chosen in ref. [51] to be a Poisson distribution.

3.3 An asymmetric geometric model: the *abc*-model

The symmetry in the geometric model can be traced back to the fact that the constraint in eq. (3.6) defining the model only depends on the magnitude of the coefficients, and it is insensitive to their signs. Alternative models that lead to asymmetric distributions and include information about coefficient's signs are discussed in appendix B of ref. [51] (see in particular B.2 and B.5). Here, we consider a new model where the upper and lower bounds are different in magnitude. More precisely, we have three hidden parameters $\mathbf{p} = (a, b, c)$, and we consider a sequence of coefficients $\delta_k = \Sigma^{(k)}/\Sigma^{(0)}$ such that (again, we suppress the dependence on the scale μ)

$$b - c \leq \frac{\delta_k}{a^k} \leq b + c, \quad \text{for all } k \geq 0. \quad (3.14)$$

We refer to this model as the *asymmetric geometric model*, or the *abc-model*. The conditional probability distribution takes a form very similar to the one in the geometric model in eq. (3.7):

$$\begin{aligned} P_{abc}(\boldsymbol{\delta}_n | a, b, c) &= \prod_{k=0}^n \frac{1}{2|a|^k c} \Theta\left(c - \left|\frac{\delta_k}{a^k} - b\right|\right) \\ &= \frac{1}{(2c)^{n+1} |a|^{n(n+1)/2}} \Theta\left(c - \max\left(\left|\frac{\delta_0}{a^0} - b\right|, \dots, \left|\frac{\delta_n}{a^n} - b\right|\right)\right). \end{aligned} \quad (3.15)$$

Like in the geometric model, our choice of priors assumes that a is independent of the parameters b and c . However, since b and c jointly define the upper and lower bounds, we do not see any reason to consider them independent. Instead, we use a joint probability for (b, c) :

$$P_0(a, b, c) = P_0(a) P_0(b, c), \quad (3.16)$$

where the factors in the right-hand side are given by:

$$\begin{aligned} P_0(a) &= \frac{1}{2} (1 + \omega) (1 - |a|)^\omega \Theta(1 - |a|), \\ P_0(b, c) &= \frac{\varepsilon \eta^\varepsilon}{2\xi c^{2+\varepsilon}} \Theta(c - \eta) \Theta(\xi c - |b|). \end{aligned} \quad (3.17)$$

Unlike in the geometric model, we do not constrain a to be positive and only require that $|a| < 1$. The prior $P_0(b, c)$ is characterised by three constants ε, η and ξ . As in eq. (3.8), $\varepsilon > 0$ is used to make the prior normalisable. The second parameter $\eta > 0$ determines the minimal difference between the upper and lower bounds of the series, see eq. (3.14). Finally, $\xi > 0$ restricts the allowed range of b , with $\xi = 0$ forcing b to vanish and we recover a symmetric geometric model. We note that the constraints in eq. (3.14) for $k = 0$ imply $c > |b - 1|$. Combining this with the prior in eq. (3.17), we obtain $c > \max\{\eta, \frac{1}{1+\xi}\}$. In figure 2 we show the ranges for b and c allowed by the prior in eq. (3.17) and $c > |b - 1|$ for the two cases $\xi > 1$ and $\xi < 1$. We can clearly see from figure 2 that if $\eta < \frac{1}{1+\xi}$, then η no longer affects the model inference. Similarly if $\xi > \frac{\eta+1}{\eta} > 1$, i.e., $\eta > \frac{1}{\xi-1}$, then the inference becomes independent of ξ . Every subsequent term δ_k in the series adds the constraint $c > \left|\frac{\delta_k}{a^k} - b\right|$, whose position depends on the values of δ_k and a .

Figure 2. The shaded area indicates the ranges for b and c allowed by the prior $P_0(b, c)$, in eq. (3.17) and the inference from the universal input $\delta_0 = 1$ in eq. (3.15) for $\xi > 1$ (left) and $\xi < 1$ (right). The position of $c = \eta > 0$ line is arbitrary.

We conclude this section by stressing that the quantities ω , ξ , ϵ and η are not hidden parameters of the model. Rather, they are (positive) constants defining a family of abc -models. In the next section we will study the model sensitivity to the choice of these constants. However, one has to keep in mind that this only explores the robustness for the specific model defined by eqs. (3.14), (3.15) and (3.17). There is an infinite-dimensional functional space of possible models and priors that one might consider. We believe that the abc -model offers a good compromise between the flexibility of the model (it gives asymmetric probability distribution and can be equally well applied to monotonic and simple alternating series) and the transparency and simplicity of the model.

3.3.1 Sensitivity to the choice of priors

Let us illustrate the abc -model and its sensitivity to the choice of priors on the example of the geometric series $\delta_k = x^k$, for $|x| < 1$. Then there exists $a = x$ for which the region of values of (b, c) excluded by eq. (3.15) collapses to $c > |1 - b|$. We see from figure 2 that $\eta > 0$ prevents c from reaching zero. Setting $\eta \rightarrow 0$ and $\xi \rightarrow \infty$ turns off any constraints by the prior $P_0(b, c)$, and the abc -model can describe such a geometric series arbitrary well with $a \rightarrow x, b \rightarrow 1, c \rightarrow 0$ in eq. (3.15).

We compute the conditional probability for the estimated next term in the series $\delta_{n+1}^{\text{est}}$:

$$P_{abc}(\delta_{n+1}^{\text{est}} | \delta_n) = \frac{P_{abc}(\delta_{n+1})}{P_{abc}(\delta_n)}, \quad (3.18)$$

where

$$\begin{aligned} P_{abc}(\delta_n) &= \int da db dc P_{abc}(\delta_n | a, b, c) P_0(a) P_0(b, c) \\ &= \int da \frac{P_0(a)}{|a|^{n(n+1)/2}} \int db \frac{\epsilon \eta^\epsilon}{2^{n+2} \xi (n+2+\epsilon)} \max \left(\eta, \max_{0 \leq k \leq n} \left| \frac{\delta_k}{a^k} - b \right|, \frac{|b|}{\xi} \right)^{-(n+2+\epsilon)} \end{aligned} \quad (3.19)$$

For fixed a , the integral over b can be done analytically and expressed in terms of maximum and minimum functions of $\frac{\delta_k}{a^k}$, η and ξ . In appendix B we show how to evaluate the

Figure 3. Top left panel: the probability distribution from the abc -model for S_{n+1}^{est} for different values of n for the geometric series with $x = 0.7$. Top right panel: the same as the left panel, but we show the probability for the scaled deviation from the known correction $(S_{n+1}^{\text{est}} - S_{n+1})/|S_{n+1} - S_n|$. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of S_{n+1}^{est} , computed from the abc (blue) and geometric (red) models using information on the previous orders. The exact values of S_n are shown as black circles. Bottom right panel: the same as the left panel, but the exact S_n value is subtracted from S_{n+1}^{est} and the difference is normalised by $|S_{n+1} - S_n|$.

distribution $P_{abc}(\delta_n)$ analytically in terms of Gauss' hypergeometric function, allowing for a fast and efficient numerical implementation of the model.

In the top left panel of figure 3 we show the probability distributions for the partial sums $S_n := \sum_{k=0}^n \delta_k = \sum_{k=0}^n x^k$ for $x = 0.7$ and $n \leq 7$ using the abc -model with parameter values $(\epsilon, \omega, \xi, \eta) = (0.1, 1, 2, 0.1)$ (see the discussion below for the choice of these values). For $n = 0$, the probability distribution is symmetric and centred around $S_0 = 1$. For $n > 0$ the distributions are clearly not symmetric and become more and more peaked as n increases. In the top right panel we show the probability distributions for the scaled deviation from the known correction $(S_{n+1}^{\text{est}} - S_{n+1})/|S_{n+1} - S_n| = (\delta_{n+1}^{\text{est}} - \delta_{n+1})/|\delta_{n+1}|$. This allows us to compare different orders without the suppression of the expansion parameter. In this plot the S_{n+1} value corresponds to zero on the x -axis, while S_n corresponds to ± 1 (depending on the sign of δ_{n+1}). Again, for $n = 0$ the distribution is centred around the initial value S_0 , but for each subsequent order, the distribution shifts towards the true value. We note that the shape of the distribution does not change significantly beyond $n = 3$, so that the narrowness of the distributions in the left panel is solely due to the higher

power of the expansion parameter. In the simple case of a geometric series, the scaled distributions are controlled by the prior $P_0(b, c)$. For $\eta = 0.1$ and $\xi = 2$, the minimum of c is at $(b, c) = (\frac{\xi}{1+\xi}, \frac{1}{1+\xi}) \approx (0.6, 0.3)$ (see figure 2), which gives a good approximation for the centre position and the (half) width of distributions seen in figure 3. As explained earlier, increasing ξ and reducing η can narrow down the posterior distribution around the exact geometric series value x^{n+1} .

In the bottom panels of figure 3 we show the 68% and 95% CIs for the *abc* (blue) and geometric models (red) (we use the same values $(\omega, \epsilon) = (1, 0.1)$ for both models). On the left we show the CIs for S_{n+1}^{est} , while on the right we again scale to the known size of the next term in the series. We see that for $n > 0$ the 68% CIs computed with the *abc*-model encompass the next value of S_n . This is not the case for the geometric model, even though the 68% CIs are larger in magnitude. This is due to the fact that the geometric model is symmetric around the last known order, thereby not accommodating the fact that all the coefficients δ_k are positive. In contrast, the CIs of the *abc*-model are systematically shifted towards the exact values for $n > 3$, and the 95% CIs even exclude the previous value S_n . However, one must remember that these features of the *abc*-model may not always be desired in practical applications. and can be controlled by the priors (see below).

For completeness, in appendix C we have also included a similar plot showing the alternating series with $x = -0.7$. Here we just note that *abc*-model correctly predicts the negative sign of the expansion parameter and remains consistently shifted towards the next value in the geometric series δ_{n+1} . In all other respects the distributions and the CIs are identical to the ones shown in figure 3.

In figure 4 we study the dependence of $P_{abc}(\delta_{n+1}|\delta_n)$ on the choice of the values for $(\epsilon, \omega, \eta, \xi)$ that define the priors on (a, b, c) in eq. (3.17). We vary each parameter independently, with the other held fixed at their default values $(\epsilon, \omega, \xi, \eta) = (0.1, 1, 2, 0.1)$. We see that, as expected, the dependence on ϵ , ω , and η reduces as n increases. The ξ -dependence, however, does not disappear completely with increasing n . This is due to the fact that ξ controls the amount of asymmetry of $P_{abc}(\delta_{n+1}|\delta_n)$. For small values of ξ , the distribution approaches a symmetric distribution around the last known value S_n , while for large ξ it becomes much more peaked around the next value S_{n+1} . Empirically, we find that our choice $\xi = 2$ allows distributions to be shifted towards the next term, but keeping the width of the distribution of the same size as the shift. Next we look at ω , which controls the suppression of the prior for a around $|a| = 1$ ($\omega = 0$ corresponds to a flat prior). Because here we consider a rather large value of the expansion parameter $x = 0.7$, increasing ω slightly reduces the expectation value of a and shifts the distribution to smaller values of S_{n+1}^{est} . Our default choice of $\omega = 1$ biases the hidden model parameter inference towards small $|a|$ values. In the bottom left panel of figure 4 we see that larger values of η shift the posterior distribution towards the previous order. η plays a similar role to ξ in controlling the allowed minimum value of c . We use a small value of $\eta = 0.1$, which is in fact superseded by the constraint from ξ (see figure 2). Finally, in contrast to η , for larger values of ϵ the model becomes slightly more peaked and shifted towards the next term in the series. We choose a small value $\epsilon = 0.1$ to maintain fairly flat distribution in the prior of $\log c$.

Figure 4. The dependence of the probability distribution $1/S_0 P_{abc}((S_{n+1}^{\text{est}} - S_n)/S_0 | \delta_n)$ for geometric series for $n = 1$ (dashed) and $n = 3$ (solid) on the parameters $(\epsilon, \omega, \eta, \xi)$ for 4 selected values, with the others held fixed at $(\epsilon, \omega, \eta, \xi) = (0.1, 1, 0.1, 2)$.

4 Quantities without explicit scale dependence

The goal of this section and the next is to illustrate the concepts and models from the previous sections on various examples of observables in QFT. We start by investigating several quantities whose perturbative expansion does not have an explicit dependence on the perturbative scales, or for which scale dependence is not relevant. This allows us to illustrate the application of the *abc*-model to genuine perturbative expansions from QFT. The sequence δ_n is the sequence of the $n + 1$ first perturbative coefficients, normalised such that $\delta_0 = 1$

$$\delta_k := \Sigma^{(k)} / \Sigma_0. \quad (4.1)$$

The rescaling introduces a Jacobian into eq. (2.8), which now takes the form

$$P(\Sigma | \Sigma_n) \approx \frac{1}{\Sigma_0} P_{abc} \left(\frac{\Sigma}{\Sigma_0} - \sum_{k=0}^n \delta_k \middle| \delta_n \right). \quad (4.2)$$

We then use this probability distribution to compute CIs for Σ within the *abc*-model using perturbative input through $N^n\text{LO}$. These intervals estimate the size of the missing $N^{n+1}\text{LO}$ terms, and so they serve as measures of the MHO uncertainty at $N^n\text{LO}$. To assess the validity of this procedure, we show in each case how the size of the CIs computed at

different perturbative orders compares to the actual size of the next order whenever it is available. By abuse of notation, we will use both $P_{abc}(\Sigma_{n+1}^{\text{est}}|\delta_n)$ and $P_{abc}(\delta_{n+1}^{\text{est}}|\delta_n)$, where the relation to eq. (4.2) has to be inferred from the argument.

Before discussing explicit examples of perturbative expansions in QFT, let us make some general comments. First, we stress that the approximation in eq. (4.2) is only valid if we are in the perturbative regime. It is in particular not valid once we enter the regime where the perturbative series starts to diverge. In that situation the right-hand side of eq. (4.2) may still represent a valid probability distribution for the next perturbative coefficient, but this distribution is no longer a good approximation of the left-hand side, which is the probability distribution for the true value of Σ . We note that for most applications only very few perturbative orders are available, and this will not be a problem. Second, it is possible to be in a situation where the model happens to describe the input data extremely well. The *abc*-model strongly favours a perfect geometric series. If the available perturbative data closely follow a geometric progression up to a certain order by some numerical coincidence, then the resulting distributions can be strongly peaked. In such a scenario, the 68% or 95% CIs may even exclude the highest available order Σ_n . This is a *design feature* of the *abc*-model that can be controlled by the priors (see section 3.3).

We stress that regardless of how well the truncated perturbative series is described by a chosen model, there is no guarantee that the actual next order in perturbative QFT or the measured experimental value of Σ will fall into a specific choice of CI. Therefore, the choices of the model, the priors and the CI used to estimate the MHO uncertainty of an $N^n\text{LO}$ computation should be carefully scrutinised for each new observable separately.

4.1 The cusp anomalous dimension

Our first example is the light-like cusp anomalous dimension. It is known through four loops in QCD and $\mathcal{N} = 4$ Super Yang-Mills (SYM) [54–73]. In the planar limit of the $\mathcal{N} = 4$ SYM theory, it is even known to all orders in perturbation theory from integrability [74]. The cusp anomalous dimension is therefore an ideal candidate to study the applicability of Bayesian techniques to derive MHO uncertainties for perturbative field theory quantities.

Planar $\mathcal{N} = 4$ SYM. The cusp anomalous dimension in planar $\mathcal{N} = 4$ SYM admits a perturbative expansion:

$$\gamma_K(\lambda) = \sum_{k=1}^{\infty} \gamma_K^{(k)} \lambda^k. \quad (4.3)$$

Here the expansion parameter is the 't Hooft coupling $\lambda = \frac{g_s^2 N_c}{8\pi^2}$, where g_s is the Yang-Mills coupling and N_c is the number of fundamental $\text{SU}(N_c)$ colours. Note that $\mathcal{N} = 4$ SYM is a conformal field theory (even away from the planar limit), and so the coupling is not UV renormalised and does not depend on any renormalisation scale. The first few expansion coefficients are shown in table 2. It was observed that the series in eq. (4.3) is convergent if $|\lambda| \lesssim \frac{1}{8}$.

k	1	2	3	4	5	6	7
$\gamma_K^{(k)}$	4	-6.57	23.81	-117.17	662.62	-4028.78	25618.7
$\gamma_{\text{QCD}}^{(k)}$	1	1.73	2.80	2.89	—	—	—

Table 2. The first few terms in the perturbative series of the cusp anomalous dimension for planar $\mathcal{N} = 4$ SYM, eq. (4.3), and QCD, eq. (4.5).

Figure 5. Top left panel: the probability distribution from the *abc*-model for the cusp anomalous dimension $(\gamma_K)_{n+1}^{\text{est}}$ in planar $\mathcal{N} = 4$ SYM for $\lambda = 0.06$ and for different values of n . Top right panel: the same distributions normalised to the exact N^{n+1} LO correction. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of $(\gamma_K)_{n+1}^{\text{est}}$, computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact value of $(\gamma_K)_n$ is shown as black circles. Bottom right panel: CIs scaled to the exact N^{n+1} LO correction.

In figures 5 and 6 we show the probability distributions obtained from the *abc*-model for the partial sums

$$(\gamma_K)_n := \sum_{k=1}^{n+1} \gamma_K^{(k)} \lambda^k, \quad (4.4)$$

with $\lambda = 0.06$ and $\lambda = 0.155$. The series δ_n is normalised according to eq. (4.1).

For $\lambda = 0.06 < \frac{1}{8}$, the perturbative series (which is alternating) is expected to converge. In figure 5 we clearly see a convergent pattern in the left top and bottom panels. The distributions quickly become very peaked around the limiting value with very small 68% and 95% CIs. In the right panels, we show the distributions and CIs normalised to the

Figure 6. The probability distributions and CIs for the cusp anomalous dimension $(\gamma_K)_{n+1}^{\text{est}}$ in planar $\mathcal{N} = 4$ SYM for $\lambda = 0.155$ and for different values of n . See the caption of figure 6 for details.

exact $N^{n+1}\text{LO}$ correction. This effectively cancels the suppression due to increasing power of the expansion parameter λ . We see that after such a rescaling, the distributions do not change significantly for $n > 2$ (up to a trivial reflection around the vertical axis, because of the alternating nature of the expansion). We note that in contrast to the case of an exact geometric series (see figure 3), here the posterior distribution for $\delta_{n+1}^{\text{est}}$ is centred around the half of the exact value $|\delta_{n+1}|$. As the series progresses, the scaled distribution reaches a peak for $n = 5$. After that it becomes steeper and broader. In the bottom right panel of figure 5, we see that the 95% CIs shrink more than the 68% CIs. For comparison we also show the CIs for the geometric model, which are significantly wider, but are no better in capturing the next order than the narrower CIs from the *abc*-model. In both cases the 95% CIs appear to give a reasonable estimate for the size of δ_{n+1} , even though the series is clearly not a perfect geometric series.

For $\lambda = 0.155$ the series is asymptotic and starts to diverge after the seventh perturbative order, i.e. $n > 6$. This behaviour is reflected in the probability distributions in the left panel of figure 6: they become peaked around the limiting value up to the inclusion of the sixth order (the sixth and seventh order corrections are numerically almost identical). After that a plateau develops which becomes broader as we increase the perturbative order. Consequently, the width of the CIs increases steadily after the seventh order (bottom left panel). However, the rescaled distributions (right panels) look essentially identical to the

convergent $\lambda = 0.06$ case. This is not surprising, because the two series differ only by the value of the expansion parameter, which is scaled away in the right panels. Therefore, like before the 95% CIs give a good estimate of the size of the next order corrections, even in the case of an asymptotic series. We note, however, that for $n > 6$ we enter the regime where the next order is not expected to give a good approximation of the true value of γ_K . Hence, eq. (4.2) is not valid, and so the credibility intervals are not expected to give a good estimate of the uncertainty attached to the perturbative result.

The cusp anomalous dimension in QCD. In QCD, the cusp anomalous dimension is known up to four loops [54–58, 72]. It depends on the representation of the Wilson loop and the matter content of the theory. Here we focus on the fundamental representation, and we work with 5 massless flavours. In this case we have

$$\Gamma_{\text{cusp}}^{\text{QCD}}(\mu_R) = C_F \sum_{k=1}^{\infty} \left(\frac{\alpha_s(\mu_R)}{\pi} \right)^k \gamma_{\text{QCD}}^{(k)}, \quad (4.5)$$

with $C_F = 4/3$ and the numerical values of coefficients $\gamma_{\text{QCD}}^{(k)}$ are summarised in table 2. We note that the scale dependence enters solely through the coupling constant α_s . Without any context, the scale is completely arbitrary and there exists no physically motivated choice for the central scale without specifying the physical process.

In figure 7 we show (with $\alpha_s(\mu_R = 5 \text{ GeV}) \approx 0.213$) the probability distributions and CIs for the partial sums $(\Gamma_{\text{cusp}}^{\text{QCD}})_n = C_F \sum_{k=1}^{n+1} \left(\frac{\alpha_s}{\pi} \right)^k \gamma_{\text{QCD}}^{(k)}$ obtained from the *abc*-model. In the top left panel we see that the series is rapidly convergent, and the probability distributions become progressively more peaked. From the rescaled distribution in the top right panel we see that the probability distribution for $n = 2$ peaks at the exact value of the correction. We do not show scaled results for $n = 3$, because the exact five-loop result is not known yet. In the bottom left panel we show the 68% and 95% CIs for the *abc*-model and the geometric model. We see that, unlike in the case of planar $\mathcal{N} = 4$ SYM, the 68% CIs always include the next correction, thus giving a reliable estimate of the MHOs (at least for the first few orders considered here). We emphasise that this is particularly non-trivial at four loop order, where Casimir scaling is violated for the first time, and a new quartic Casimir enters. Estimates of the size of the four-loop corrections based on Padé approximants from the first three orders were within 5% of the four-loop contributions that admit Casimir scaling for $n_f = 5$, but were 57% off from the complete result [75]. The 68% CI obtained from the first three loop orders in the *abc*-model correctly captures the size of the four-loop result, including the contribution from the quartic Casimir.

From the bottom right plot we see that the median of the posterior distribution for the *abc* model (marked by a plus) is closer to the exact value than for the geometric model, although the *abc*-model appears to expect larger higher-order correction, and the distributions are skewed towards larger corrections. For $n = 4$ the CIs are not visible in the bottom left panel due to the suppression of the expansion parameter. Using the *abc*-model we estimate that CIs for MHO corrections given the 4-loop order result for $\mu_R = 5 \text{ GeV}$,

Figure 7. Top left panel: the probability distribution from the *abc*-model for the QCD cusp anomalous dimension $(\Gamma_{\text{cusp}}^{\text{QCD}})_{n+1}^{\text{est}}$ evaluated with $\alpha_s(\mu_R = 5 \text{ GeV}) \approx 0.213$ and for different values of n . Top right panel: the same distributions normalized to the exact $N^{n+1}\text{LO}$ correction. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of $(\Gamma_{\text{cusp}}^{\text{QCD}})_{n+1}^{\text{est}}$, computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact values of $(\Gamma_{\text{cusp}}^{\text{QCD}})_n$ are shown as black circles. Bottom right panel: CIs scaled to the exact $N^{n+1}\text{LO}$ correction.

i.e., $\Sigma = \Gamma_{\text{cusp}}^{\text{QCD}} - (\Gamma_{\text{cusp}}^{\text{QCD}})_4$, is

$$\text{CI}_{68} = C_F \left(\frac{\alpha_s(\mu_R)}{\pi} \right)^5 [2.1, 9.5], \quad \text{CI}_{95} = C_F \left(\frac{\alpha_s(\mu_R)}{\pi} \right)^5 [-0.38, 21], \quad (4.6)$$

where we scaled out $C_F \left(\frac{\alpha_s(\mu_R)}{\pi} \right)^5$ for visibility. We see that the *abc*-model at 68% credibility level would expect a positive higher-order correction. In contrast, the geometric model estimates a symmetric interval $\text{CI}_{68} = [-5.4, 5.4]$ with the same normalisation.

4.2 On-shell and $\overline{\text{MS}}$ -scheme quark masses

Our next example is the relation between the heavy quark mass m in the on-shell scheme and the mass $\overline{m}$ in the $\overline{\text{MS}}$ -scheme. The bare mass m_0 is related to the on-shell and $\overline{\text{MS}}$ quark masses via the renormalisation factors Z_m

$$m_0 = Z_m^{\overline{\text{MS}}}(\mu_R) \overline{m}(\mu_R) = Z_m^{\text{OS}} m. \quad (4.7)$$

Figure 8. Top left panel: the probability distribution from the *abc*-model for the on-shell top quark mass $(m_t)_{n+1}^{\text{est}}$ evaluated at $\mu_R = \bar{m}_t$ and for different values of n . Top right panel: the same distributions normalised to the exact $N^{n+1}\text{LO}$ correction. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of $(m_t)_{n+1}^{\text{est}}$, computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact values of $(m_t)_n$ are shown as black circles. Bottom right panel: CIs scaled to the exact $N^{n+1}\text{LO}$ correction.

For heavy quarks it is possible to perturbatively compute these renormalisation factors. They can be used to define the scheme-conversion factor between the two schemes,

$$z_m(\mu_R) = \frac{Z_m^{\text{OS}}}{Z_m^{\overline{\text{MS}}}(\mu_R)} = \frac{\bar{m}(\mu_R)}{m} = \sum_{k=0}^{\infty} \left(\frac{\alpha_s(\mu_R)}{\pi} \right)^k z_m^{(k)}(\mu_R). \quad (4.8)$$

For heavy quarks $z_m(\mu_R)$ is currently known up to four loops [76–78]. We can express the on-shell mass as a series calculated in perturbative QFT

$$m = z_m^{-1}(\mu_R) \bar{m}(\mu_R) = \sum_{k=0}^{\infty} m^{(k)}. \quad (4.9)$$

In practice, one does not consider arbitrary μ_R but rather work with a self-consistent definition of the $\bar{m}$ mass $\bar{m}(\bar{m})$. Thus for the quark masses we consider the scale as fixed number rather than as a free parameter.

In figure 8 we show the probability distributions and CIs of the on-shell top quark mass $(m_t)_{n+1}^{\text{est}} = \sum_{k=0}^{n+1} m^{(k)}$ given the first $n+1$ perturbative orders of the scheme-conversion factor $z_{m_t}^k(\mu_R)$ evaluated at $\mu_R = \bar{m}_t$. In the top left panel we see that the distributions

become more and more peaked as perturbative information becomes available. On the top right we see that the *abc*-model correctly estimates the positive sign of the corrections, but the estimated corrections are in general smaller than the exact value. In the bottom left panel we see that the CIs shrink rapidly as n increases. In the bottom right we see that, unlike in the case of the QCD anomalous dimension, the 68% CIs of both the *abc* and geometric models do not include the next perturbative order, although the 95% CIs are sufficiently wide to include it. We note that the *abc*-model distributions are asymmetric, with a tail towards positive perturbative corrections showing that, although the *abc*-model does not peak at the exact value of $N^n\text{LO}$ contribution, it (correctly) anticipates that larger positive corrections are more likely. This is a result of the fact that the perturbative coefficients $z_m^{(k)}$ are all negative (at least through four loops). Using the *abc*-model we estimate the CIs for the MHOs given the 4-loop result, i.e., $\Sigma = m_t - (m_t)_4$, to be (in GeV):

$$\text{CI}_{68} = [0.008, 0.046], \quad \text{CI}_{95} = [-0.027, 0.112]. \quad (4.10)$$

It is interesting to compare the size of the CIs obtained from the *abc*-model to the range five-loop estimates (in GeV) obtained in ref. [79]:

$$\begin{aligned} (m_t)_0 \delta_5^{\text{ECH}} &= 0.073, & (m_t)_0 \delta_5^{\text{ECH direct}} &= 0.083, & (m_t)_0 \delta_5^{\text{FL}} &= 0.126, \\ (m_t)_0 \delta_5^{\text{FL}, M \rightarrow \bar{m}} &= 0.086, & (m_t)_0 \delta_5^{r-n} &= 0.112. \end{aligned} \quad (4.11)$$

We see that the 68% CI from the *abc*-model are lower than the estimates of ref. [79], but the 95% CI includes all values of ref. [79], except one.

For completeness, in appendix C, we include plots for the bottom and charm quarks. The scaled probability distributions and CIs look similar to the ones in the right panels of figure 8. We note, however, that the MHO estimate using Bayesian inference is less reliable for lighter quarks, as in this case corrections become large and the perturbative expansion itself fails. This can be seen especially in the case of the charm quark mass.

In figure 9 we show the dependence of the probability distributions for the $(m_t)_{n+1}^{\text{est}}$ for $n = 1, 3$ from the *abc*-model on the parameters $(\eta, \xi, \epsilon, \omega)$. As already discussed in section 3.3, we see that the dependence on (η, ϵ, ω) reduces as more perturbative information becomes available, while ξ controls the amount of asymmetry of the distribution. However, after the fourth perturbative order, the distributions become very narrow and peaked, and the shape is largely independent of the values of $(\eta, \xi, \epsilon, \omega)$ if $\xi \geq 2$.

4.3 Electron anomalous magnetic moment and positronium hyperfine-splitting

In this section we illustrate how Bayesian techniques work in the context of two classical results in QED which are known to relatively high order in the coupling, namely the electron anomalous magnetic moment and the positronium hyperfine splitting. QED corrections are typically believed to be rapidly convergent due to the smallness of the QED coupling $\alpha = 1/137.035\,998\,995(85)$ [80]. Further, these results are evaluated in the on-shell renormalisation scheme, thus they do not depend on any unphysical scale.

Figure 9. The sensitivity of the probability distribution $1/(m_t)_0 P_{abc}(((m_t)_{n+1}^{\text{est}} - (m_t)_n)/(m_t)_0 | \delta_n)$ for on-shell top quark mass for $n = 1$ (dashed) and $n = 3$ (solid) on the parameters $(\epsilon, \omega, \eta, \xi)$ for 4 selected values with the others held fixed at $(\epsilon, \omega, \eta, \xi) = (0.1, 1, 0.1, 2)$.

k	1	2	3	4	5	
$C_e^{(2k)}$	0.5	-0.328	1.181	-1.912	6.675(192)	
ij	00	10	21	20	32	31
$C_{\text{HFS}}^{(ij)}$	0.583	-1.24	-2.06	-3.88	-8.64	14.97

Table 3. Known coefficients of mass-independent corrections to electron anomalous magnetic moment and to positronium hyperfine splitting.

Let us start by discussing the electron anomalous magnetic moment. The QED corrections take the form

$$\frac{g-2}{2} \equiv a_e = \sum_{k=1}^{\infty} \left(\frac{\alpha}{\pi} \right)^k C_e^{(2k)}. \quad (4.12)$$

Here we focus only on the mass-independent part of the corrections, which are known up to 5 loops [81–90]. The numerical values for the coefficients $C_e^{(2k)}$ are collected in table 3.

Figure 10 illustrates the *abc*-model estimates for the probability distributions and the CIs for the partial sums $(a_e)_{n+1}^{\text{est}} = \sum_{k=1}^{n+1} \left(\frac{\alpha}{\pi} \right)^k C_e^{(2k)}$. Due to the very small values of the expansion parameter, the absolute distributions and CIs are hardly visible for $n > 1$. The distribution for $n = 1$ is rather broad, resulting in CIs that overestimate the actual size of the higher-order corrections. At $n = 2$, the 68% CI captures the size of the next correction, while for $n = 3$ only the 95% CI includes the next order.

Figure 10. Top left panel: the probability distribution from the *abc*-model for $(a_e)_{n+1}^{\text{est}}$ and for different values of n . Top right panel: the same distributions normalised to the exact $N^{n+1}\text{LO}$ correction. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of $(a_e)_{n+1}^{\text{est}}$, computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact values of $(a_e)_n$ are shown as black circles. Bottom right panel: CIs scaled to the exact $N^{n+1}\text{LO}$ correction.

Our second QED example is the positronium hyperfine-splitting. This is a non-perturbative quantity and admits a double expansion in α and $\ln \alpha$:

$$\Delta E = m_e \alpha^4 \sum_{i,j=0}^{\infty} C_{\text{HFS}}^{(ij)} \left(\frac{\alpha}{\pi} \right)^i \ln^j \alpha. \quad (4.13)$$

Only terms with $j < i$ are non-zero with the exception of the coefficient C_{00} . This introduces a natural ordering among the coefficients, which are ordered first by i then by j . For the known terms this ordering is also equivalent to ordering the contributions by their size, since $\ln \alpha \sim -5$. The currently known values [91–95] are shown in table 3.

The presence of the logarithmic corrections in the series clearly violates the geometric progression of the series assumed in both the geometric and *abc*-models. Our Bayesian inference for such series is therefore less reliable as an estimate of the MHOs in comparison to a purely perturbative quantity such as a_e . This can be also seen in figure 11: at the lowest two perturbative orders the CIs overestimate the size of the perturbative corrections. As the logarithmically enhanced corrections appear, both geometric and *abc* models fail to provide reliable estimates of the next-order correction. This is expected, since the models

Figure 11. The probability distribution and the CIs for the hyperfine-splitting ΔE in units of the electron mass m_e . See the caption of figure 10.

consider only a single expansion parameter a and expect the next-order correction to be suppressed by an additional power of a .

In principle, we can consider even more complicated QED bound-state observables, for example, the bound electron g -factor [96–99]. In this case, one has a triple expansion in powers of $(Z\alpha)$, α and $\ln(Z\alpha)$, where Z is atomic charge. We expect that such series require more attuned models that take into account non-perturbative terms and double or triple expansions. We conclude this section by observing that QED corrections require dedicated studies within Bayesian framework that we leave for future works.

5 Hadronic observables with explicit scale dependence

In this section we discuss the application of Bayesian techniques to estimate MHOs for a selection of hadronic observables known to NNLO, or even N³LO, in the strong coupling constant. In all cases, we compute scale-independent probability distributions and CIs using the scale-marginalisation (sm) and scale-averaging (sa) prescriptions introduced in sections 2.1.1 and 2.1.2. Throughout this section we work with the geometric and *abc*-models of section 3. We start by discussing inclusive cross-sections and discuss differential distributions towards the end of this section. In contrast to the examples studied in ref. [51], we take into account the variation of the factorisation scale.

Figure 12. Top left panel: the probability distributions from the *abc*-model for the inclusive gluon-fusion Higgs production cross-section $(\sigma_{ggH})_{n+1}^{\text{est}}$ evaluated at $\mu_R = \mu_F = m_H/2$ and for different values of n . Top right panel: the same distributions normalised to the exact $N^{n+1}\text{LO}$ correction. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of $(\sigma_{ggH})_{n+1}^{\text{est}}$, computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact values of $(\sigma_{ggH})_n$ are shown as black circles. Bottom right panel: CIs scaled to the exact $N^{n+1}\text{LO}$ correction.

5.1 Higgs production in gluon fusion

We start by discussing the inclusive Higgs production cross-section in gluon fusion:

$$(\sigma_{ggH})_n(\mu_F, \mu_R) = \sum_{k=0}^n \sigma_{ggH}^{(k)}(\mu_F, \mu_R). \quad (5.1)$$

The inclusive cross-section is known through NLO in QCD including all finite top-mass effects [100–109], and through $N^3\text{LO}$ in an effective theory expansion where the top quark is infinitely heavy [110–115]. All results in this section were obtained with *iHixs 2* [116], neglecting finite top-mass effects, with the `PDF4LHC15_nnlo_100` PDF set and the on-shell top quark mass $m_t = 172.5 \text{ GeV}$. We choose as value for the central scale $\mu_0 = m_H/2 = 62.5 \text{ GeV}$ (both for scale-variation results and in the priors in eqs. (2.5) and (2.12)).

First, in figure 12 we show the probability distributions and CIs for $(\sigma_{ggH})_n^{\text{est}}$ for fixed $\mu_R = \mu_F = \mu_0$. We see that for $n = 1, 2$, the CIs for both the *abc* and geometric models include the next order at 68% credibility level. We note, however, that for $\mu_0 = m_H/2$ the $N^3\text{LO}$ corrections are particularly small, and the CIs of the *abc*-model overshoot the

Figure 13. The inclusive Higgs production cross-section as a function of the renormalisation scale (with $\mu_F = m_H/2$ held fixed) at different orders $n = 0, 1, 2, 3$. The errorbars and erroboxes correspond to the NNLO 95% and 68% CIs for the geometric model (light green) and the *abc*-model (light blue), given the $n < 3$ results.

computed value. We will see below that for other scale choices, the CIs of the *abc*-model are centred around the N³LO values.

Next, in figure 13 we plot the renormalisation scale dependence of $(\sigma_{ggH})_n$ at different orders (for this plot only we keep $\mu_F = \mu_0$ fixed). We see that in the conventional scale variation window, $F = 2$, the cross-section depends quite strongly on the scale, except for the highest order. The N³LO corrections vanish for $\mu_R \approx 0.8\mu_0$, and the lines representing the NNLO and N³LO results as a function of the scale cross, i.e., we reach a FAC point. The FAC point at NNLO is at a smaller value, $\mu_R \approx 0.3\mu_0$. Also, the crossing points visible in figure 13 approximately coincide with the local maxima for $n = 2, 3$, i.e., the FAC and PMS points are surprisingly close to each other.

In figure 13 we further show the CIs in the geometric and *abc*-models with LO, NLO and NNLO cross-sections (at fixed scale) used as an input. We observe that the CIs for the geometric model are centred around the $n = 2$ line due to the symmetry of the geometric model, but the *abc*-model assigns a higher probability to a positive correction. Indeed, for $\mu_R > 2\mu_0$ the CIs are almost centred around the N³LO line. For $\mu_R < 0.8\mu_0$, the N³LO corrections become negative, while the *abc*-model still anticipates a positive correction. We note that the CIs for the geometric model shrink dramatically as one approaches the NNLO FAC point at $\mu_R = 0.3\mu_0$, i.e., large corrections are deemed unlikely by the geometric model close to that scale. Once the NNLO corrections become negative, the *abc*-model also becomes centred around the NNLO result, i.e., the *abc*-model no longer expects a positive correction. The CIs in this region do not shrink as dramatically as for the geometric model, although generally the *abc*-model leads to narrower CIs than the geometric model.

Figure 14. Left: the dependence of the CIs on the size F of the scale interval ($\mu_0/F \leq \mu \leq F\mu_0$) for gluon-fusion Higgs cross-section for the geometric and *abc*-models using the sm-prescription. Right: the same plot using the sa-prescription.

5.1.1 Sensitivity to the scale interval and scale prescription

We begin by analysing the dependence of our sm- and sa-prescriptions on the range F that is chosen in eqs. (2.9) and (2.13), i.e., $\mu_0/F \leq \mu \leq F\mu_0$.

Figure 14 shows the CIs for the values $F = 2, 4, 10$ using the geometric and *abc*-models. The left and right panels summarise the results from the sm- and sa-prescriptions respectively. For $n = 2$ the sm-prescription exhibits a somewhat smaller dependence on the choice of F , while the CIs for the sa-prescription grow with F . This can be understood from the discussion in section 2.2: for symmetric models like the geometric model, the sm-prescription tries to adapt to the point where the higher-order corrections are minimised, i.e., the FAC point. Once this point is covered by the range in the marginalisation, a further increase does not have a substantial impact on the uncertainties. From figure 13 we see that this is precisely what happens for the geometric model for $n = 2$, when F is increased from 2 to 4, and even for $F = 10$ the CIs change only very little. The discussion of section 2.2 does not apply to the *abc*-model, which is asymmetric. Recall from figure 13 that the *abc*-model does not have such dramatic reduction in the size of the CIs when reaching an FAC point. Therefore it is also less biased to the inclusion of FAC point into the integration range for μ in the sm-prescription. For $n = 3$ the second FAC point at very small scales becomes relevant only for the largest F values.

The sa-prescription, on the other hand, is biased to the regions where the scale dependence is flat, i.e., it is biased towards the PMS point. However, for this particular process we observe the peculiar situation that FAC and PMS points are very close to each other (at least for the orders we considered). Due to the coherent addition of probabilities, for the sa-prescription increasing the range of the μ -integration inevitably will also increase the CIs.

Let us conclude this discussion with an important point. From a Bayesian perspective, one expects that the predictions of the model become independent of the priors once enough data have become available. In particular, in the sm-prescription, the perturbative scales are treated as model parameters, and so one expects that at high enough orders the model

Figure 15. Left: the coloured lines show the probability distributions for the *abc*-model at different orders using the sa-prescription and the 3-point Gauss-Legendre quadrature on $\mu_F/\mu_0, \mu_R/\mu_0 \in \{1/2, 1, 2\}$. The black dashed line indicates the “exact” numerical integration in the interval with $F = 2.45$. Right: the CIs for the geometric and *abc* models at different orders using the sm- and sa-prescriptions using the Gauss-Legendre quadrature. The result of the “exact” numerical integral are overlaid using black dashed lines. For comparison we also include the 9-point and 7-point variation intervals.

predictions should only mildly depend on the prior in eq. (2.5), and the choice of F , in agreement with our findings in figure 14. However, it would be premature to conclude that the probability distributions are independent of the prior and F : we have shown in section 2.2 that for the geometric model in the sm-prescription, the preferred scale is the FAC point. The fact that the CIs are insensitive to the choice of F is likely to be related to the fact that the distributions will be highly peaked at the values of the cross-section at the FAC point, which is not necessarily related to the prior independence. Note that the argument of section 2.2 does not apply to the *abc*-model in the sm-prescription (because the *abc*-model is not symmetric).

5.1.2 Accuracy of Gauss-Legendre quadrature rule

As already discussed in section 2.3, performing the numerical integral over the scale can become computationally expensive. Quadrature rules, on the other hand, allow one to approximate the full integral using only a small number of input points. In particular, a set of points for a 9-point scale variation may be readily turned into an integral over (μ_R, μ_F) .

The left panel of figure 15 illustrates the accuracy of this approximation for the probability distributions in the *abc*-model using the sa-prescription. The distributions for different values of n (in colour) are obtained using the 3-point Gauss-Legendre quadrature for both μ_R and μ_F . The results obtained by a finely-spaced equidistant numerical integration over the same scale interval with $F \approx 2.45$ is overlaid on top (black dashed lines). We see that, except for some features around the peak, the quadrature rules reproduce the “exact” numerical integration very well. In particular, this approximation has very little effect on the CIs, which are shown in the right panel of figure 15, where we show the 68% and 95% CIs for the geometric and *abc* models using both the sm- and sa-prescriptions. The dashed black lines indicate the results of the finely-spaced numerical integration. We note that

after scale integration the probability distributions are no longer symmetric, even if the distributions at fixed scale are. Overall, we observe that the Gauss-Legendre quadrature rule is able to reproduce both the 68% and 95% CIs from the “exact” numerical integration to a very good degree. Therefore, in the subsequent sections we will only show the results obtained using the Gauss-Legendre quadrature.

For Higgs production in gluon fusion both the sm- and sa-prescriptions produce very similar CIs shown in the right panel of figure 15. The CIs are also very similar in size for both the geometric and the *abc*-models. The results for the latter are, however, systematically shifted upwards due to the sequence of positive corrections at lower orders. For comparison we also display the scale variation intervals using 9- and 7-point rules, which have no probabilistic interpretation. We observe that the scale variation intervals are slightly smaller, but rather close, to the 68% CIs obtained from the Bayesian methods. This illustrates the arguments of section 2.3 in a practical example. In particular, thanks to the slow μ_F variation, the scale variation intervals are close to saturating the $s = 1$ bound in eq. (2.29).

Finally, we conclude with the summary of 68% and 95% CIs and scale-variation intervals for the highest known order $n = 3$ of $(\sigma_{ggH})_n$:

model	prescription	CI ₆₈ (pb)	CI ₉₅ (pb)	7-point (pb)	9-point (pb)
<i>abc</i>	sa	[47.8, 51.2]	[44.4, 56.5]	[46.3, 48.3]	[46.0, 48.3]
<i>abc</i>	sm	[48.3, 51.1]	[46.1, 55.8]		
geo	sa	[46.0, 49.2]	[41.6, 53.0]		
geo	sm	[47.0, 49.2]	[42.7, 53.0]		

We note that as the N³LO cross-section reaches a local maximum near the central scale μ_0 , traditional scale variation produces a one-sided interval. The geometric model has approximately symmetric intervals around the central value, while the *abc*-model anticipates positive MHO terms. Although all 68% CIs and the scale-variation intervals are numerically close in size, we stress that they have completely different statistical meanings. Scale-variation intervals have no probabilistic interpretation and cannot be associated with a particular level of credibility. In contrast, the Bayesian CIs express a mathematically rigorous degree of belief for the MHOs from the posterior distribution for a given model and scale prescription.

5.2 Vector-boson fusion Higgs and di-Higgs production

Our next two examples are Higgs production in vector-boson fusion (VBF), and a similar process for double Higgs production. The inclusive VBF cross-section is known up to N³LO. We denote VBF cross-sections for single Higgs [117–120] and double-Higgs production [120–122] by $(0 \leq n \leq 3)$

$$\begin{aligned}
 (\sigma_{\text{VBF-H}})_n(\mu_F, \mu_R) &= \sum_{k=0}^n \sigma_{\text{VBF-H}}^{(k)}(\mu_F, \mu_R), \\
 (\sigma_{\text{VBF-HH}})_n(\mu_F, \mu_R) &= \sum_{k=0}^n \sigma_{\text{VBF-HH}}^{(k)}(\mu_F, \mu_R).
 \end{aligned}
 \tag{5.2}$$

Figure 16. The 68% and 95% CIs for the VBF cross-sections for Higgs and di-Higgs production for the geometric and *abc*-models using the sa- and sm-prescriptions. The scale variation intervals using 7 and 9 points are shown for comparison.

As before the centre-of-mass energy is $\sqrt{s} = 13$ TeV. The central scale is given by the vector boson momentum [123] and we take into account the dependence on both factorisation and renormalisation scales. Computations were performed with the `proVBFH` code [124].

In the left panel of figure 16 we display the CIs for different models and prescriptions for single Higgs VBF production. For $n < 2$ the Bayesian approach gives a larger uncertainty (68% CIs) than the traditional scale variation. Because the NLO correction is negative, the *abc*-model anticipates an alternating series, and consequently the CIs for $n = 1$ for the *abc*-model are positively shifted compared to the NLO result. However, the NNLO corrections are again negative, and for $n = 2$ all studied models and prescriptions give very similar 68% CIs, although the *abc*-model has much larger 95% CIs than the geometric model. For $n = 3$ the 68% CIs shrink even further and become somewhat smaller than the scale variation intervals. For the single Higgs VBF cross-section $(\sigma_{\text{VBF-H}})_n$ at $n = 3$ these CIs are:

model	prescription	CI ₆₈ (fb)	CI ₉₅ (fb)	7 point (fb)	9 point (fb)
<i>abc</i>	sa	[3.9306, 3.9357]	[3.9287, 3.9478]	[3.9304, 3.9367]	[3.9304, 3.9367]
<i>abc</i>	sm	[3.9304, 3.9337]	[3.9290, 3.9430]		
geo	sa	[3.9305, 3.9343]	[3.9287, 3.9385]		
geo	sm	[3.9304, 3.9324]	[3.9293, 3.9355]		

We note that the sm-prescription gives much smaller CIs than the sa-prescription. In fact, the 95% CIs of the scale-marginalised geometric model is smaller and does not contain the scale-variation interval, demonstrating that the bounds discussed in section 2.3 do not necessarily apply to the sm-prescription. In contrast the 95% CIs for the geometric model in sa-prescription contain the scale variation intervals, as expected.

In the right panel of figure 16 we display the CIs for different models and prescriptions for di-Higgs VBF production. We observe very good convergence of the cross-section, and correspondingly the CIs from Bayesian inference shrink rapidly. We observe that the sa-prescription gives larger CIs than the sm-prescription, which is due to the presence of an

FAC point in the scale integration interval. We also see a significant difference between the 7- and 9-point scale variation intervals. The 9-point interval for $n = 2$ is larger than the 68% CI for the geometric model with the sa-prescription, but still within the 95% CI. This is a sign that both renormalisation and factorisation scale dependencies are comparable and the one dimensional ($s = 1$) bound estimated in section 2.3 is not applicable in this case. For $n = 2$ and $n = 3$ the 95% CI for scale-marginalised geometric model is even smaller than 9-point scale variation interval, while sa-prescription 95% still encompasses it, as expected. The CIs for di-Higgs VBF cross-section $(\sigma_{\text{VBF-HH}})_n$ at $n = 3$ are:

model	prescription	CI ₆₈ (fb)	CI ₉₅ (fb)	7 point (fb)	9 point (fb)
<i>abc</i>	sa	[1.7237, 1.7248]	[1.7229, 1.7280]	[1.7232, 1.7247]	[1.7232, 1.7258]
<i>abc</i>	sm	[1.7241, 1.7248]	[1.7238, 1.7252]		
geo	sa	[1.7237, 1.7247]	[1.7229, 1.7263]		
geo	sm	[1.7239, 1.7247]	[1.7232, 1.7249]		

We see that sa-prescription gives generally more conservative estimates for the MHOs than either the sm-prescription or scale variation.

5.3 Drell-Yan processes

Drell-Yan-type processes play a crucial role in hadron collider phenomenology. The inclusive cross-section for off-shell photon production with virtuality Q^2 , $pp \rightarrow \gamma^{(*)}$, is known though N³LO [125–127]:

$$(\sigma_{\text{DY-NC}})_n(Q^2, \mu_F, \mu_R) = \sum_{k=0}^n \sigma_{\text{DY-NC}}^{(k)}(Q^2, \mu_F, \mu_R). \quad (5.3)$$

N³LO corrections are also known for the charged-current Drell-Yan process $pp \rightarrow W^\pm \rightarrow \ell^\pm \nu_\ell$ [126, 128–133]. We focus here on the lepton charge asymmetry defined as

$$A_W(Q^2) = \frac{\frac{d\sigma_{\text{DY-W}^+}}{dQ^2} - \frac{d\sigma_{\text{DY-W}^-}}{dQ^2}}{\frac{d\sigma_{\text{DY-W}^+}}{dQ^2} + \frac{d\sigma_{\text{DY-W}^-}}{dQ^2}}, \quad (5.4)$$

and the n -th order approximation is

$$(A_W)_n(Q^2, \mu_F, \mu_R) = \sum_{k=0}^n A_W^{(k)}(Q^2, \mu_F, \mu_R). \quad (5.5)$$

For both processes we only consider here $Q = m_W$ with the central scale set to $\mu_0 = Q = m_W$.

In the left panel of figure 17 we show results for the neutral current Drell-Yan cross-section. We see that for $n = 2, 3$ the sa-prescription produces 68% CIs which are similar in size to the regular scale variation intervals. The CIs from the sm-prescription for $n = 2$ are much smaller, again due to the presence of an FAC point in the integration region for μ . The N³LO corrections are known to be sizeable, and they are not covered by the

Figure 17. The 68% and 95% CIs for the neutral-current Drell-Yan cross-section and charged-current lepton-charge asymmetry for the geometric and *abc*-models using the sa- and sm-prescriptions. The scale variation intervals using 7 and 9 points are shown for comparison.

conventional 7-point scale variation at $n = 2$ [127]. The CIs for the neutral-current Drell Yan cross-section $(\sigma_{\text{DY-NC}})_n$ at $n = 3$ are:

model	prescription	CI ₆₈ (nb)	CI ₉₅ (nb)	7 point(nb)	9 point (nb)
<i>abc</i>	sa	[45.6, 46.6]	[44.8, 49.0]	[45.6, 46.4]	[45.5, 46.4]
<i>abc</i>	sm	[45.9, 46.5]	[45.1, 48.3]		
geo	sa	[45.5, 46.4]	[44.6, 47.2]		
geo	sm	[45.8, 46.3]	[45.0, 46.9]		

We observe that the 68% CIs are similar in size among themselves, and to the scale-variation intervals, but the CIs from the *abc*-model are slightly shifted upwards in the anticipation of a positive MHO correction.

In the right panel of figure 17 we show results for the lepton charge asymmetry for $\mu_0 = Q = m_W$. The perturbative expansion for $A_W(m_W^2)$ is quickly convergent with only a mild scale dependence, because some corrections cancel in the ratio. The perturbative coefficients feature a monotonic increase with the perturbative order, and the *abc*-model correctly anticipates positive contributions from MHOs. The CIs from the *abc*-model are slightly smaller than for the geometric model. We do not observe significant differences between the sm- and sa-prescriptions, except for $n = 3$, where scale-marginalisation gives more aggressive CIs. We note that the traditional 7-point scale variation intervals for $n = 0, 1$ fail to include the next correction, but for $n = 2$ they are similar to the 68% CIs obtained from Bayesian inference. The results for the CIs for $A_W(m_W^2)$ at N³LO are:

model	prescription	CI ₆₈	CI ₉₅	7 point	9 point
<i>abc</i>	sa	[0.1489, 0.1494]	[0.1486, 0.1497]	[0.1488, 0.1493]	[0.1485, 0.1494]
<i>abc</i>	sm	[0.1491, 0.1494]	[0.1488, 0.1497]		
geo	sa	[0.1487, 0.1493]	[0.1484, 0.1495]		
geo	sm	[0.1490, 0.1494]	[0.1485, 0.1495]

Figure 18. Deep inelastic scattering distribution $d\sigma/dQ^2$ up to $N^3\text{LO}$ with $\mu_0 = Q$ as the central scale choice. Left: absolute distributions and 9-point scale variation envelopes normalised to the NLO prediction at μ_0 . Right: 68% CIs for the *abc* and geometric models using the scale-marginalisation (top) and scale-averaging (bottom) prescriptions (also normalised to NLO).

We observe that the sm-prescription produces smaller CIs than the sa-prescription, indicating the presence of FAC point, which also pulls the CIs upwards with respect to the $N^3\text{LO}$ result. Scale-averaged CIs are more conservative and more centred around the $N^3\text{LO}$ results for both the geometric and *abc*-models.

5.4 Deep inelastic scattering

The Bayesian procedure of estimating MHOs can be readily applied to differential spectra by considering individual points/bins of the associated observable.⁷ In the following, we will investigate how the geometric and *abc*-models perform for various differential predictions. Results will be further contrasted between the sm- and sa-prescriptions, and compared to the 9-point scale variation intervals. Like for inclusive observables, we will perform scale integrals using the Gauss-Legendre method from section 2.3. This setup allows us to accommodate already existing results based on 7- or 9-point scale variation, without the need for a full re-computation of these predictions. In this section we begin our discussion of differential spectra with the DIS process. We will present results for several benchmark processes at the LHC in the subsequent sections.

Predictions up to $N^3\text{LO}$ for the DIS process were obtained from the calculation in ref. [134], where we consider the distribution $d\sigma/dQ^2$ with respect to the virtuality Q^2 of the intermediate photon. The results are presented in figure 18. The left panels show the

⁷Note that here no correlations are considered, i.e., cross-section entries in each bin of the distribution are assumed to be independent. The incorporation of correlations is left for future studies.

absolute distributions (top) and the standard scale variation envelopes normalised by the NLO prediction evaluated at the central scale (bottom). The four right panels contrast the MHO estimates (68% CIs) obtained using the *abc* and geometric models with the *sa*- (top two panels) and *sm*-prescriptions (bottom two panels). All CIs are normalised to the NLO prediction evaluated at the central scale.⁸ Bands of different colours correspond to estimates at different perturbative orders: grey (LO), blue (NLO), red (NNLO) and dark grey (N³LO). Overall, we can observe that all models and prescriptions perform well in capturing the progression of the perturbative series with either overlapping or touching uncertainty bands. At NLO the Bayesian approaches give rise to 68% CIs which are slightly larger than the 9-point scale variation interval, but are comparable in size at NNLO and N³LO. In the case of the *sm*-prescription shown in the top two panels on the right in figure 18, the CI bands narrow dramatically at NLO for $Q^2 \sim 400$ GeV, as well as at NNLO and N³LO at the lower end of the spectrum. This can be traced back to the presence of a FAC points that are strongly emphasised by the *sm*-prescription. In the case of the geometric model, we see that both scale prescriptions move the median of the NLO probability distribution towards the direction of the higher-order corrections, which slightly improves the overlap between the CIs at NNLO and N³LO. For the *abc*-model, on the other hand, the central value at NLO moves in the opposite direction, with strongly asymmetric uncertainty bands that extend further towards the LO prediction. This is because after the negative NLO correction, the *abc*-model anticipates a positive NNLO correction, as would be the case for an alternating series. However, for this process the NNLO correction is again negative. The geometric model, which ignores the signs of the corrections, appears to perform better. We note that at higher orders the *abc*-model captures the perturbative progression well with only slightly larger CIs than the geometric model.

5.5 Di-photon production

We now consider predictions for di-photon production at the LHC at $\sqrt{s} = 8$ TeV using the results of ref. [135]. The production of a pair of photons, $pp \rightarrow \gamma\gamma$, provides a clean and well-measured final state that is an important background in Higgs measurements and can be used in the search for New Physics resonances. One striking feature of this process is the apparent slow convergence of the perturbative QCD series, as partially induced by the staggered transverse-momentum cuts on the photons. As a consequence, lower-order predictions are found to be inadequate to describe the measurement, and only starting from NNLO a sensible comparison to the data can be performed.

Results for the di-photon invariant-mass distribution $d\sigma/dm_{\gamma\gamma}$ are shown in figure 19 following the same pattern as in figure 18. We note that the sizeable higher-order corrections put the respective predictions far outside from the traditional scale variation intervals of the previous order. Moreover, the scale variation envelope at NNLO is found to be of similar size, or even larger, than at NLO, with no convincing sign of convergence. The Bayesian models, on the other hand, produce 68% CIs that are generally larger, and the CIs do not increase when going to higher orders. With the large positive corrections pro-

⁸For clarity, we do not show 68% CIs at LO, nor any of the 95% CIs, which can be very large.

Figure 19. Di-photon production up to NNLO with $\mu_0 = m_{\gamma\gamma}$ as the central scale choice. Left: absolute $d\sigma/dm_{\gamma\gamma}$ distributions and 9-point scale variation envelopes normalised to the NLO prediction at μ_0 . Right: 68% CIs for the *abc* and geometric models using the scale-marginalisation (top) and scale-averaging (bottom) prescriptions (also normalised to NLO).

hibiting any FAC points, the two scale prescriptions produce CIs that are almost identical. While the CI bands between NLO and NNLO touch in the case of the geometric model, for the *abc*-model the NNLO CIs are outside of the NLO band for $m_{\gamma\gamma} \gtrsim 200$ GeV. This is likely due to large positive corrections encountered at each perturbative order. The *abc*-model anticipates this pattern to continue at higher orders and shifts the probability distribution towards larger cross-section values. Finally, in view of the noticeable differences for the size and position of the bands at the highest order (NNLO) using different approaches, we can conclude that N³LO corrections will be highly relevant for this process to obtain more robust MHO estimates, which the traditional scale variation is likely substantially underestimating.

5.6 Gauge boson production in association with a jet

The production of electroweak gauge bosons that recoil against hard QCD emissions is among the most important standard candle processes at the LHC with a wide range of applications that span from detector calibration, precision QCD studies, to New Physics searches. This process is now known up to NNLO for all gauge boson types. Here we focus on W^+ +jet production using the predictions from ref. [136].

Figure 20 presents the results for the inclusive transverse momentum distribution $d\sigma/dp_{T,W}$ of the W boson following the same layout as in figure 18. We can observe a good perturbative behaviour for this process with NNLO corrections that are well captured by the traditional scale variation at one order lower. The Bayesian approaches give rise to 68% CIs that are wider at NLO, while the size of the bands at NNLO are compara-

Figure 20. W^+ +jet production up to NNLO with $\mu_0 = E_{T,W}$ as the central scale choice. Left: absolute $d\sigma/dp_{T,W}$ distributions and 9-point scale variation envelopes normalised to the NLO prediction at μ_0 . Right: 68% CIs for the *abc* and geometric models using the scale-marginalisation (top) and scale-averaging (bottom) prescriptions (also normalised to NLO).

ble between all approaches. The sm- and sa-prescriptions give rise to almost identical CIs, similar to the observations made for the di-gamma process. In the case of the *abc*-model, we can appreciate the flexibility of the asymmetric probability distribution. The positive offset at NLO centres the CIs at this order around the NNLO band. Finally, we note that once NNLO corrections are included, all the different prescriptions of estimating MHOs considered here give almost identical bands indicating that the uncertainties assigned at this order are very robust.

5.7 Higgs boson production in association with a jet

The study of the Higgs boson and its properties is among the highest priorities of the LHC programme and the transverse momentum distribution, in particular, plays an important role in this endeavour not only for boosted analyses but also to resolve the details of the production dynamics for this process. We focus here on Higgs boson production in association with a jet at $\sqrt{s} = 13$ TeV in the four-lepton decay channel, $H \rightarrow 4\ell$, and use the predictions for the “ATLAS II” setup computed in ref. [137].

The results up to NNLO for the transverse momentum distribution $d\sigma/dp_{T,H}$ of the Higgs boson are shown in figure 21. We observe large positive corrections at each perturbative order with the exception of the lowest p_T bins. In contrast to the other differential distributions considered so far, we notice that the Bayesian models are not giving rise to substantially larger uncertainty bands (68% CI) at NLO compared to the traditional scale variation envelope, but they are similar in size. Going to NNLO, the uncertainty estimates are found to be comparable between the different approaches, with the sm-prescription giv-

Figure 21. Higgs+jet production up to NNLO with $\mu_0 = \frac{1}{2}E_{T,H}$ as the central scale choice. Left: absolute $d\sigma/dp_{T,H}$ distributions and 9-point scale variation envelopes normalised to the NLO prediction at μ_0 . Right: 68% CIs for the *abc* and geometric models using the scale-marginalisation (top) and scale-averaging (bottom) prescriptions (also normalised to NLO).

ing somewhat narrower bands. This compatibility between the different approaches again supports the robustness of the uncertainties assigned to these predictions. With the monotonic pattern of positive corrections at each order, the *abc*-model is again able to adapt to this feature in the intermediate-high p_T range, capturing to a large extent the NNLO corrections in the NLO shift. This is more pronounced in the sm-prescription suggesting that the central scale choice of $\mu_0 = \frac{1}{2}E_{T,H}$ is probing regimes in the scale variation that are close a FAC point.

5.8 Inclusive jet production

Jet production is the most prominent process at hadron colliders and an essential tool for the study of QCD. The inclusive transverse momentum distribution $d\sigma/dp_{T,j}$ is particularly interesting in the context of theory uncertainties, due to its sensitivity on the specific choice of the central scale that further persists at NNLO due to instabilities in the sub-leading jet distribution that impacts the inclusive spectrum [138]. In order to gain insights into this sensitivity and how the different prescriptions for assigning MHO uncertainties compare in these scenarios, we consider two different central scale choices: the jet transverse momentum $p_{T,j}$ and the scalar sum of the transverse momenta of the partons $\hat{H}_T = \sum_{i \in \text{partons}} p_{T,i}$, which are shown in figures 22 and 23, respectively.

We begin our discussion with the scale choice $\mu_0 = p_{T,j}$, which makes the NLO corrections artificially small. From the results presented in figure 22, we see that in the low- p_T region all approaches give rise to sizeable NNLO CIs. This is most clearly visible for the *abc*-model, where no reduction in the size of the CI band is observed as we move from

Figure 22. Inclusive jet production up to NNLO with $\mu_0 = p_{T,j}$ as the central scale choice. Left: absolute $d\sigma/dp_{T,j}$ distributions and 9-point scale variation envelopes normalised to the NLO prediction at μ_0 . Right: 68% CIs for the *abc* and geometric models using the scale-marginalisation (top) and scale-averaging (bottom) prescriptions (also normalised to NLO).

Figure 23. Inclusive jet production up to NNLO with $\mu_0 = \hat{H}_T$ as the central scale choice. Left: absolute $d\sigma/dp_{T,j}$ distributions and 9-point scale variation envelopes normalised to the NLO prediction at μ_0 . Right: 68% CIs for the *abc* and geometric models using the scale-marginalisation (top) and scale-averaging (bottom) prescriptions (also normalised to NLO).

NLO to NNLO. Neither of the bands overlap, because the *abc*-model expects a positive NNLO correction, which is not true for this central scale choice. The geometric model, instead, is able to assign more conservative MHO uncertainties at NLO that allow for some overlap between the last two orders. Overall, we can conclude that the Bayesian models considered here are not able to “un-do” an inappropriate scale choice through either the sm- or sa-prescriptions.

More specifically, no prescription based on a simple re-scaling from a central scale choice will be able to adapt between different *dynamical* scales that encapsulate a different kinematic dependence. This becomes most apparent in observables that are sensitive to infrared physics, because it is the cancellation between real and virtual corrections that strongly impacts such a sensitivity. For two dynamical scale choices that have a different dependence on the kinematics, the difference in the scales assigned to the real and virtual corrections can vary substantially. Such a feature can never be captured by a global scaling, as done in the sa- and sm-prescriptions considered here, that varies the scales for the real and virtual contributions in the same manner.

Using $\hat{H}_T$ as the central scale choice as advocated in ref. [138], we observe in figure 23 a stable progression of the series throughout with the next order, typically lying within the uncertainty estimate of the previous order. The 68% CIs are slightly larger for the Bayesian approaches at NLO, while at NNLO all uncertainty estimates are mutually compatible. This further supports this scale as an appropriate choice given its robustness with respect to the different prescriptions in assigning theory uncertainties. Finally, it is interesting to note that the *abc*-model is able to adapt to the sizeable positive corrections in the tail of the distribution, thus giving a noticeable positive shift at NLO that reduces the overall impact of NNLO corrections.

6 Discussion

The estimation of MHOs is a vital aspect of collider phenomenology. Traditional approaches are based on scale variation, an ad hoc prescription to obtain uncertainty intervals for MHOs without a clear probabilistic underpinning. In this paper, we have analysed Bayesian approaches to estimate the MHO uncertainties from the progression of the perturbative expansion of a physical observable. Building on the recent work in ref. [51], we have focused in particular on the question of how the dependence on the unphysical perturbative scales μ_F and μ_R can be incorporated into Bayesian inference. Our goal was to scrutinise various Bayesian models and to assess their common features by studying their performance on a wide and representative selection of hadron collider observables known at high orders in perturbation theory, namely NNLO and/or N³LO. In figure 24 we show a summary plot of 68% and 95% CIs for MHOs in our Bayesian approach and the traditional scale variation intervals for selected collider processes known up to N³LO. In the following we summarise our main findings and conclusions, pointing out directions for future research.

Biases related to the choice of the scale prescription. An important question is how the dependence on perturbative scales can (or should) be included into Bayesian MHO

Figure 24. Comparison of MHO estimates at N³LO between different models and scale prescriptions for the gluon-fusion Higgs cross-section, single and di-Higgs production in vector boson fusion, neutral current Drell-Yan production and the lepton charge asymmetry. See the corresponding sections for each process for further details.

estimation. In ref. [51] it was proposed to treat the perturbative scales on par with other hidden parameters of a Bayesian model, which we refer to as the scale-marginalisation (sm) prescription. In section 2 we have introduced an alternative prescription, called scale-averaging (sa), which does not require the perturbative scales to be assigned a statistical interpretation. We do not see any compelling physics argument to prefer one prescription over the other.

The two prescriptions do not just differ in their philosophy of how the perturbative scales are interpreted, but we have demonstrated that they may lead to substantially different estimates for MHOs. In particular, in section 2.2 we have shown that (under some additional reasonable assumptions for Bayesian models) each of the two prescriptions has intrinsic biases towards specific scales: the sm-prescription tends to favour FAC points as the ‘best’ value of the perturbative scale, while the sa-prescription is biased towards the values of the observable at the PMS points. These biases are clearly visible in the examples in section 5 (even for models for which the assumptions entering the proofs are not satisfied). Care is thus needed when applying Bayesian techniques to incorporate scale information: the choice of the prescription may bias the predictions in favour of a specific scale choice. It would be interesting to see if one can construct other Bayesian models that incorporate scale information in an unbiased way.

The choice of the Bayesian model and the priors. In order to estimate MHOs in a Bayesian approach, one has to specify a model for the progression of the perturbative series. The geometric model studied in ref. [51] neglected the sign of perturbative corrections, which enforces the posterior distributions for MHOs (at fixed scale) to be symmetric around

zero. To allow for asymmetric MHO probability distributions we introduced the *abc*-model in section 3 (see appendix B of ref. [51] for alternatives). In section 4 we benchmarked the *abc*-model on a number of perturbative QFT quantities known to high orders, for which scale dependence is not relevant. Our results show that the specific choice of the model and its priors can strongly influence the size and position of MHO CIs, especially when the number of input data is small.

From the physics perspective, there is no preference between Bayesian models. While from an aesthetic point of view, one may prefer simple models, a judicious but ad hoc choice has to be made. It might be therefore desirable to develop new models in the future suited for specific classes of perturbative series, e.g., see our QED examples.

The choice of the central scale and the range of scales considered. The prior on the perturbative scale in eq. (2.5) for the sm-prescription (see ref. [51]) (or the weighting function in eq. (2.12) for the sa-prescription) depends on a choice of the central scale μ_0 and the size F of the interval for the scale integration. Our analysis in section 5 shows that the CIs have a substantial dependence on the choice of μ_0 and F . For example, once the range $\mu_F, \mu_R \in [\mu_0/F, \mu_0 F]$ is wide enough to include an FAC point, the CIs in the sm-prescription may become relatively F -independent, but mostly because the probability distributions become strongly peaked and biased towards to the FAC point. The sa-prescription, instead, will typically increase with F , thereby rendering the probability distributions F -dependent.

The study of the inclusive jet cross section at NNLO in section 5.8 shows that the convergence of the CIs in the Bayesian approach (at least for the prescriptions and models studied here) is not any less sensitive to the central scale choice as the traditional scale-variation intervals. It would be interesting to see if there are other prescriptions or choices for the prior $P_0(\mu)$ and the weighting function $w(\mu)$ that lead to probability distributions and CIs that are truly independent of the choice of the central scale μ_0 and interval size F . The main challenge in this regard arises from the limitations of a naive re-scaling of the central scale by constant factors that, in general, will not be able to capture the difference in the kinematic dependence among possible choices for a dynamical scale.

Relation to MHO estimates using scale-variation. Let us point out a very practical connection between the existing theoretical results with the 9-point scale variation and scale prescriptions in the Bayesian approach. The scale dependence is typically mild enough so that a low-order Gauss-Legendre quadrature rule can be used to approximate the scale integrals in the sm- and sa-prescriptions. For the specific choice of integration interval $F \approx 2.45$, the traditional scale variation points are placed at the quadrature nodes. Therefore the Bayesian CIs can be efficiently computed from already existing predictions.

Furthermore, for the observables studied in section 5, we have demonstrated that the traditional 7- or 9-point scale variation intervals are often similar in size to the Bayesian 68% and 95% CIs. Similar behaviour has been observed before [48, 50, 51]. This is not a coincidence, and in section 2.3 we have shown that (under certain assumptions) the 95% CIs in the sa-prescription must include the scale-variation intervals. We stress that scale-variation intervals themselves are devoid of any statistical interpretation.

Finally, we observe that the 68% CIs at different perturbative orders in the Bayesian approach share some qualitative features known from scale-variation. For example, neither the scale variation envelopes nor the 68% CIs overlap for Drell-Yan production at N³LO and NNLO or di-photon production at NNLO and NLO (however some models lead to barely overlapping 68% CIs, e.g., see figure 19). For processes where scale-variation intervals indicate good perturbative convergence, Bayesian models often also show a good, or even better, nesting of 68% CIs at different orders. The similarities of the two approaches are not coincidental, as both prescriptions share assumptions on the convergence of the perturbative series, namely that we are at low enough order so the perturbative series exhibits a convergent behaviour. Despite these similarities, scale-variation intervals and Bayesian CIs are very different objects from a statistical point of view. In particular, within the Bayesian approach, one can choose in a quantitative manner how aggressive or conservative estimates of MHO should be for a given process and for a given choice of model.

Conclusions. Based on our comprehensive analysis, we conclude that Bayesian approaches provide solid alternatives to the traditional scale variation as a means to estimate MHOs, but there are several caveats that one needs to keep in mind (and that possibly need to be improved in the future, e.g., by developing novel Bayesian models). The choice of the Bayesian setup, in particular the scale prescription and the priors, may have a substantial impact on the inference. In a certain sense, the Bayesian approach is not any less subjective or ad hoc than customary scale variation. Therefore care is needed not to introduce any biases by the model choice.

A significant advantage of Bayesian techniques over scale variation is that the CIs computed from Bayesian approach have a clear probabilistic interpretation (at least within the context of a given model). This opens the way to apply statistical techniques, e.g., to systematically combine MHO uncertainties with uncertainties from different sources or to include correlations between observables or bins for differential distributions. It would also be interesting to explore how far one can treat the choice of the model itself in a Bayesian way, in order to minimise bias from the model choice, e.g., using ideas from *Bayesian Model Averaging (BMA)* (see, e.g., ref. [139], and references therein). We leave these studies for future work.

The source code to the computer program `MiHO` that was used to obtain the results shown in the paper is made publicly available at <https://github.com/aykhuss/miho>.

Acknowledgments

RS is supported by the United States Department of Energy under Grant Contract DE-SC0012704.

A The position of the peak of the distribution

In this appendix we present a (heuristic) proof of the position of the peak of the distribution $\mathcal{P}(\Sigma|\Sigma_n)$. We work under the assumptions discussed in section 2.2, which we recall here

for convenience:

1. The model (P, P_0) is symmetric.
2. We work with the prior in eq. (2.5) or the weighting function in eq. (2.12).
3. The distributions $P(\Sigma^{(n+1)}(\mu)|\Sigma_n(\mu))$ have a single peak in $\Sigma^{(n+1)}(\mu)$ for each fixed value of μ , and the peak is reached for some $\mu \in [\mu_0/F, \mu_0 F]$.
4. The peak is more and more pronounced as m increases, i.e., as more information on the progression of the series becomes known.
5. There is a single FAC or PMS point in the interval $[\mu_0/F, \mu_0 F]$.

A.1 The position of the peak for the scale-marginalisation prescription

Using the prior on the scale in eq. (2.5) and the approximation in eq. (2.8), eq. (2.9) takes the form:

$$\mathcal{P}_{\text{sm}}(\Sigma|\Sigma_n) = \frac{\int_{L_-}^{L_+} dL P(\Sigma - \Sigma_n(e^L)|\Sigma_n(e^L)) P(\Sigma_n(e^L))}{\int_{L_-}^{L_+} dL P(\Sigma_n(e^L))}. \quad (\text{A.1})$$

In the previous equation we have changed variables from μ to $L = \log \mu$, and we defined $L_{\pm} = \log(\mu_0 F^{\pm 1})$. In order to determine the position of the peak of $\mathcal{P}_{\text{sm}}(\Sigma|\Sigma_n)$ as a function of Σ , we analyse where the bulk of the integral in the numerator receives its contributions. Bayes' formula implies:

$$P(\Sigma_n(\mu)) = P(\Sigma_0(\mu)) \prod_{k=1}^n P(\Sigma^{(k)}(\mu)|\Sigma_{k-1}(\mu)). \quad (\text{A.2})$$

Since $P(\Sigma^{(k)}(\mu)|\Sigma_{k-1}(\mu))$ is assumed to be a peaked function, and the peak becomes more pronounced as k increases, the product in the right-hand side in eq. (A.2) is dominated by the peak from $k = n$. Since we assume P to be symmetric, the position of the peak is the point where $\Sigma^{(n)}(\mu) = 0$, and by assumption there is a $\mu_{\text{FAC}} \in [\mu_0/F, \mu_0 F]$ such that $\Sigma^{(n)}(\mu_{\text{FAC}}) = 0$, i.e., μ_{FAC} is a FAC point at NⁿLO. Hence, $P(\Sigma_n(\mu))$ is peaked at the FAC point $\mu = \mu_{\text{FAC}}$, and we can approximate the integral by

$$\mathcal{P}_{\text{sm}}(\Sigma|\Sigma_n) \approx P(\Sigma - \Sigma_n(\mu_{\text{FAC}})|\Sigma_n(\mu_{\text{FAC}})) = P(\Sigma - \Sigma_{n-1}(\mu_{\text{FAC}})|\Sigma_n(\mu_{\text{FAC}})). \quad (\text{A.3})$$

Since P is symmetric, $P(\Sigma - \Sigma_{n-1}(\mu_{\text{FAC}})|\Sigma_n(\mu_{\text{FAC}}))$ reaches its peak when $\Sigma - \Sigma_{n-1}(\mu_{\text{FAC}}) = 0$, i.e., the position of the peak is located at the FAC point at NⁿLO:

$$\Sigma_n^{\text{peak}} \approx \Sigma_{n-1}(\mu_{\text{FAC}}) = \Sigma_n(\mu_{\text{FAC}}). \quad (\text{A.4})$$

A.2 The position of the peak for the scale-averaging prescription

We use the same notations and conventions as in the previous section. Using the weighting function in eq. (2.12) and the approximation in eq. (2.8), eq. (2.10) takes the form:

$$\mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n) = \frac{1}{2 \log F} \int_{L_-}^{L_+} dL P(\Sigma - \Sigma_n(e^L)|\Sigma_n(e^L)). \quad (\text{A.5})$$

We proceed in a similar way as in the previous section, and we want to understand where the bulk of the integral comes from. The probability in the integrand is peaked by assumption. To understand the impact on the position of the peak of $\mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n)$, let us in first approximation write the probability distribution in the integrand as a δ -function with support at $\Sigma - \Sigma_n(e^L) = 0$:

$$\mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n) \approx \frac{1}{2 \log F} \int_{L_-}^{L_+} dL \delta(\Sigma - \Sigma_n(e^L)). \quad (\text{A.6})$$

By assumption, the δ -function has support in the range $[L_-, L_+]$, and so there is $L_\Sigma = \log \mu_\Sigma \in [L_-, L_+]$ such that $\Sigma - \Sigma_n(\mu_\Sigma) = 0$ (note that μ_Σ depends on Σ !). This gives:

$$\mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n) \approx \frac{1}{\Sigma'_n(\mu_\Sigma)}. \quad (\text{A.7})$$

We see from this argument that $\mathcal{P}_{\text{sa}}(\Sigma|\Sigma_n)$ reaches its maximum when $\Sigma'_n(\mu_\Sigma) \approx 0$, and so $\Sigma_n^{\text{peak}} = \Sigma_n(\mu_\Sigma) = \Sigma_n(\mu_{\text{PMS}})$ is at the PMS point at $N^n \text{LO}$.

B Analytic solution of the *abc*-model

In this appendix we show the analytical solution of the *abc* model for $\xi = 1$. The general case $\xi \neq 1$ can be solved using similar techniques. Recall that for $\xi = 1$, after integration over c we have

$$P_{abc}(\delta_n) = \int \frac{da db P_0(a)}{|a|^{n(n+1)/2}} \frac{\epsilon \eta^\epsilon}{2^{n+2}(n+2+\epsilon)} \max\left(\eta, \max_{0 \leq k \leq n} \left| \frac{\delta_k}{a^k} - b \right|, |b|\right)^{-(n+2+\epsilon)}. \quad (\text{B.1})$$

We perform the integral over b first. We split factors $\frac{\delta_i}{a^i}$ into two sets according to

$$\Delta^{(+)}(a) = \max\left(\left\{ \frac{\delta_i}{a^i} \mid \frac{\delta_i}{a^i} \geq 0 \right\} \cup \{0\}\right) \geq 0, \quad (\text{B.2})$$

$$\Delta^{(-)}(a) = \min\left(\left\{ \frac{\delta_i}{a^i} \mid \frac{\delta_i}{a^i} \leq 0 \right\} \cup \{0\}\right) \leq 0, \quad (\text{B.3})$$

and we define $b_V \equiv \frac{\Delta^{(+)} + \Delta^{(-)}}{2}$. We distinguish two cases:

(I) $\eta \leq |b_V|$: The integration region splits into two regions

$$\begin{aligned} \int_{-\infty}^{+\infty} db \max\{\dots\}^{-(n+2+\epsilon)} &= \int_{-\infty}^{b_V} db \left[\Delta^{(+)} - b \right]^{-(n+2+\epsilon)} \\ &\quad + \int_{b_V}^{\infty} db \left[b - \Delta^{(-)} \right]^{-(n+2+\epsilon)} \\ &= \frac{2^{n+2+\epsilon}}{n+1+\epsilon} \left(\Delta^{(+)} - \Delta^{(-)} \right)^{-(n+1+\epsilon)}. \end{aligned} \quad (\text{B.4})$$

(II) $\eta > |b_V|$: In this case, we have to consider three regions

$$\begin{aligned}
 \int_{-\infty}^{+\infty} db \max\{\dots\}^{-(n+2+\epsilon)} &= \int_{-\infty}^{-\eta+\Delta^{(+)}} db \left[\Delta^{(+)} - b\right]^{-(n+2+\epsilon)} \\
 &\quad + \int_{-\eta+\Delta^{(+)}}^{\eta+\Delta^{(-)}} db \eta \\
 &\quad + \int_{\eta+\Delta^{(-)}}^{\infty} db \left[b - \Delta^{(-)}\right]^{-(n+2+\epsilon)} \\
 &= \eta(\Delta^{(-)} - \Delta^{(+)} + 2\eta) + \frac{2\eta^{-\epsilon-n-1}}{n+\epsilon+1}. \quad (\text{B.5})
 \end{aligned}$$

Finally, we obtain the probability distribution

$$\begin{aligned}
 P_{abc}(\delta_n) &= \int_{-1}^1 da \frac{(1-|a|)^\omega}{|a|^{n(n+1)/2}} \frac{(1+\omega)\epsilon\eta^\epsilon}{2^{n+3}(n+2+\epsilon)} \\
 &\quad \times \begin{cases} \frac{2^{n+2+\epsilon}}{n+1+\epsilon} \left(\Delta^{(+)} - \Delta^{(-)}\right)^{-(n+1+\epsilon)}, & \text{for } \eta \leq |b_V|, \\ \eta(\Delta^{(-)} - \Delta^{(+)} + 2\eta) + \frac{2\eta^{-\epsilon-n-1}}{n+\epsilon+1}, & \text{for } \eta > |b_V|. \end{cases} \quad (\text{B.6})
 \end{aligned}$$

The result for $\xi \neq 1$ takes a similar functional form, but one has to consider additional cases, as shown in figure 2. In practice, we perform the integral over a numerical in our implementation of the *abc* model in the code `MiHO`. Nonetheless, it is instructive to discuss the analytical solution for a specific choice of ω . We note that the $\Delta^{(\pm)}$ depend on a and are piecewise power functions of a ; $\Delta^{(\pm)} \sim a^{-i}$ for some power i . As such, the a integration domain is divided into disjoint intervals $0 < \dots < a_k < a_{k+1} < \dots < 1$ on which the maximum function can be evaluated (see ref. [51] for how to find the a_k) and the most complicated integral we encounter has the form

$$\int_{a_k}^{a_{k+1}} da (1-|a|)^\omega a^l \left(\frac{\alpha}{a^i} + \frac{\beta}{a^j}\right)^\nu, \quad (\text{B.7})$$

with $\alpha, \beta > 0$, $l \in \mathbb{Z}$, $i, j \in \mathbb{N}$. Restricting ourselves to integer ω , the most general integral that we need is of the form

$$I_{pq}(\mu, \nu; a_k, a_{k+1}, \alpha, \beta) \equiv \int_{a_k}^{a_{k+1}} da a^\mu \left(\frac{\alpha}{a^p} + \frac{\beta}{a^q}\right)^\nu \quad (\text{B.8})$$

$$= \frac{\alpha^\nu}{p-q} \left(\frac{\alpha}{\beta}\right)^{\frac{\mu-\nu p+1}{p-q}} F\left(\frac{\mu-\nu p+1}{p-q} - 1, \nu; \frac{\beta}{\alpha} a_k^{p-q}, \frac{\beta}{\alpha} a_{k+1}^{p-q}\right), \quad (\text{B.9})$$

with F related to Gaussian hypergeometric function

$$F(\rho, \sigma; A, B) \equiv \frac{B^{1+\rho}}{1+\rho} {}_2F_1\left(2+\rho+\sigma, 1+\rho; 2+\rho; \frac{B}{1+B}\right) - (B \leftrightarrow A). \quad (\text{B.10})$$

C Supplementary plots for sections 3 and 4

For completeness we include the plots for probability distributions and CIs for the alternating geometric series, figure 25. Compare to figure 3 and see section 3.3.1 for a detailed discussion.

Figure 25. Top left panel: the probability distribution from the *abc* model for S_{n+1}^{est} for different values of n for the geometric series with $x = -0.7$. Top right panel: the same as the left panel, but we show the probability for the scaled deviation from the known correction $(S_{n+1}^{\text{est}} - S_{n+1})/|S_{n+1} - S_n|$. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of S_{n+1}^{est} , computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact values of S_n are shown as black circles. Bottom right panel: the same as the left panel, but the exact S_n value is subtracted from S_{n+1}^{est} and the difference is normalised by $|S_{n+1} - S_n|$.

We also document the plots for probability distributions and CIs for the on-shell bottom and charm quark masses in figures 26 and 27. Compare to figure 8 and see section 4.2 for a detailed discussion.

Figure 26. Top left panel: the probability distributions from *abc*-model for the on-shell bottom quark mass $(m_b)_{n+1}^{\text{est}}$ evaluated at $\mu_R = m_b$ and for different values of n . Top right panel: the same distributions normalised to the exact $N^{n+1}\text{LO}$ correction. Bottom left panel: the median (plus), 68% CI (errorbox) and 95% CI (errorbar) for the posterior of $(m_b)_{n+1}^{\text{est}}$, computed from the *abc* (blue) and geometric (red) models using information on the previous orders. The exact values of $(m_b)_n$ are shown as black circles. Bottom right panel: CIs scaled to the exact $N^{n+1}\text{LO}$ correction.

Figure 27. The probability distributions and CIs for charm quark mass. See caption of figure 26.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, *J. Phys. G* **43** (2016) 023001 [[arXiv:1510.03865](https://arxiv.org/abs/1510.03865)] [[INSPIRE](#)].
- [2] S. Alekhin et al., *The PDF4LHC Working Group Interim Report*, [arXiv:1101.0536](https://arxiv.org/abs/1101.0536) [[INSPIRE](#)].
- [3] M. Botje et al., *The PDF4LHC Working Group Interim Recommendations*, [arXiv:1101.0538](https://arxiv.org/abs/1101.0538) [[INSPIRE](#)].
- [4] L.A. Harland-Lang and R.S. Thorne, *On the Consistent Use of Scale Variations in PDF Fits and Predictions*, *Eur. Phys. J. C* **79** (2019) 225 [[arXiv:1811.08434](https://arxiv.org/abs/1811.08434)] [[INSPIRE](#)].
- [5] NNPDF collaboration, *Nuclear Uncertainties in the Determination of Proton PDFs*, *Eur. Phys. J. C* **79** (2019) 282 [[arXiv:1812.09074](https://arxiv.org/abs/1812.09074)] [[INSPIRE](#)].
- [6] NNPDF collaboration, *Parton Distributions with Theory Uncertainties: General Formalism and First Phenomenological Studies*, *Eur. Phys. J. C* **79** (2019) 931 [[arXiv:1906.10698](https://arxiv.org/abs/1906.10698)] [[INSPIRE](#)].
- [7] R.D. Ball and R.L. Pearson, *Correlation of Theoretical Uncertainties in PDF Fits and Theoretical Uncertainties in Predictions*, [arXiv:2105.05114](https://arxiv.org/abs/2105.05114) [[INSPIRE](#)].
- [8] J.R. Ellis, E. Gardi, M. Karliner and M.A. Samuel, *Pade approximants, Borel transforms and renormalons: The Bjorken sum rule as a case study*, *Phys. Lett. B* **366** (1996) 268 [[hep-ph/9509312](https://arxiv.org/abs/hep-ph/9509312)] [[INSPIRE](#)].
- [9] E. Gardi, *Why Pade approximants reduce the renormalization scale dependence in QFT?*, *Phys. Rev. D* **56** (1997) 68 [[hep-ph/9611453](https://arxiv.org/abs/hep-ph/9611453)] [[INSPIRE](#)].
- [10] J.R. Ellis, M. Karliner and M.A. Samuel, *A prediction for the four loop β -function in QCD*, *Phys. Lett. B* **400** (1997) 176 [[hep-ph/9612202](https://arxiv.org/abs/hep-ph/9612202)] [[INSPIRE](#)].
- [11] E.J. Weniger, *Performance of superconvergent perturbation theory*, *Phys. Rev. A* **56** (1997) 5165 [[INSPIRE](#)].
- [12] J.R. Ellis, I. Jack, D.R.T. Jones, M. Karliner and M.A. Samuel, *Asymptotic Pade approximant predictions: Up to five loops in QCD and SQCD*, *Phys. Rev. D* **57** (1998) 2665 [[hep-ph/9710302](https://arxiv.org/abs/hep-ph/9710302)] [[INSPIRE](#)].
- [13] M. Karliner, *Precise estimates of high orders in QCD*, *Acta Phys. Polon. B* **29** (1998) 1505 [[hep-ph/9804381](https://arxiv.org/abs/hep-ph/9804381)] [[INSPIRE](#)].
- [14] I. Caprini and J. Fischer, *Accelerated convergence of perturbative QCD by optimal conformal mapping of the Borel plane*, *Phys. Rev. D* **60** (1999) 054014 [[hep-ph/9811367](https://arxiv.org/abs/hep-ph/9811367)] [[INSPIRE](#)].
- [15] G. Cvetič and R. Kogerler, *Towards a physical expansion in perturbative gauge theories by using improved Baker-Gammel approximants*, *Nucl. Phys. B* **522** (1998) 396 [[hep-ph/9802248](https://arxiv.org/abs/hep-ph/9802248)] [[INSPIRE](#)].
- [16] V. Elias, F.A. Chishtie and T.G. Steele, *Pade improvement of hadronic Higgs decays*, *J. Phys. G* **26** (2000) 1239 [[hep-ph/0004140](https://arxiv.org/abs/hep-ph/0004140)] [[INSPIRE](#)].

- [17] I. Caprini and J. Fischer, *Convergence of the expansion of the Laplace-Borel integral in perturbative QCD improved by conformal mapping*, *Phys. Rev. D* **62** (2000) 054007 [[hep-ph/0002016](#)] [[INSPIRE](#)].
- [18] U.D. Jentschura, E.J. Weniger and G. Soff, *On the asymptotic improvement of resummations and perturbative predictions in quantum field theory*, *J. Phys. G* **26** (2000) 1545 [[hep-ph/0005198](#)] [[INSPIRE](#)].
- [19] M. Beneke and M. Jamin, α_s and the τ hadronic width: fixed-order, contour-improved and higher-order perturbation theory, *JHEP* **09** (2008) 044 [[arXiv:0806.3156](#)] [[INSPIRE](#)].
- [20] A. David and G. Passarino, *How well can we guess theoretical uncertainties?*, *Phys. Lett. B* **726** (2013) 266 [[arXiv:1307.1843](#)] [[INSPIRE](#)].
- [21] R.J. Furnstahl, N. Klco, D.R. Phillips and S. Wesolowski, *Quantifying truncation errors in effective field theory*, *Phys. Rev. C* **92** (2015) 024005 [[arXiv:1506.01343](#)] [[INSPIRE](#)].
- [22] D. Boito, M. Jamin and R. Miravitllas, *Scheme Variations of the QCD Coupling and Hadronic τ Decays*, *Phys. Rev. Lett.* **117** (2016) 152001 [[arXiv:1606.06175](#)] [[INSPIRE](#)].
- [23] M. Jamin and R. Miravitllas, *Scalar correlator, Higgs decay into quarks, and scheme variations of the QCD coupling*, *JHEP* **10** (2016) 059 [[arXiv:1606.06166](#)] [[INSPIRE](#)].
- [24] J.A. Melendez, S. Wesolowski and R.J. Furnstahl, *Bayesian truncation errors in chiral effective field theory: nucleon-nucleon observables*, *Phys. Rev. C* **96** (2017) 024003 [[arXiv:1704.03308](#)] [[INSPIRE](#)].
- [25] I. Caprini, *Higher-order perturbative coefficients in QCD from series acceleration by conformal mappings*, *Phys. Rev. D* **100** (2019) 056019 [[arXiv:1908.06632](#)] [[INSPIRE](#)].
- [26] J.A. Melendez, R.J. Furnstahl, D.R. Phillips, M.T. Pratola and S. Wesolowski, *Quantifying Correlated Truncation Errors in Effective Field Theory*, *Phys. Rev. C* **100** (2019) 044001 [[arXiv:1904.10581](#)] [[INSPIRE](#)].
- [27] O. Costin and G.V. Dunne, *Physical Resurgent Extrapolation*, *Phys. Lett. B* **808** (2020) 135627 [[arXiv:2003.07451](#)] [[INSPIRE](#)].
- [28] D. Boito and F. Oliani, *Renormalons in integrated spectral function moments and α_s extractions*, *Phys. Rev. D* **101** (2020) 074003 [[arXiv:2002.12419](#)] [[INSPIRE](#)].
- [29] C. Drischler, R.J. Furnstahl, J.A. Melendez and D.R. Phillips, *How Well Do We Know the Neutron-Matter Equation of State at the Densities Inside Neutron Stars? A Bayesian Approach with Correlated Uncertainties*, *Phys. Rev. Lett.* **125** (2020) 202702 [[arXiv:2004.07232](#)] [[INSPIRE](#)].
- [30] H.D. Politzer, *Stevenson's Optimized Perturbation Theory Applied to Factorization and Mass Scheme Dependence*, *Nucl. Phys. B* **194** (1982) 493 [[INSPIRE](#)].
- [31] M.R. Pennington, *An Island of Ambiguity in Perturbative QCD: Has Stevenson Found the Treasure?*, *Phys. Rev. D* **26** (1982) 2048 [[INSPIRE](#)].
- [32] P. Stevenson, *Sense and Nonsense in the Renormalization Scheme Dependence Problem*, *Nucl. Phys. B* **203** (1982) 472 [[INSPIRE](#)].
- [33] G. Grunberg, *Renormalization Scheme Independent QCD and QED: The Method of Effective Charges*, *Phys. Rev. D* **29** (1984) 2315 [[INSPIRE](#)].
- [34] S.J. Brodsky, G.P. Lepage and P.B. Mackenzie, *On the Elimination of Scale Ambiguities in Perturbative Quantum Chromodynamics*, *Phys. Rev. D* **28** (1983) 228 [[INSPIRE](#)].

- [35] G. Cvetič, *Renormalization scale invariant continuation of truncated QCD (QED) series: An analysis beyond large Beta0 approximation*, *Nucl. Phys. B* **517** (1998) 506 [[hep-ph/9711406](#)] [[INSPIRE](#)].
- [36] C.J. Maxwell and A. Mirjalili, *Complete renormalization group improvement: Avoiding factorization and renormalization scale dependence in QCD predictions*, *Nucl. Phys. B* **577** (2000) 209 [[hep-ph/0002204](#)] [[INSPIRE](#)].
- [37] S.J. Brodsky and L. Di Giustino, *Setting the Renormalization Scale in QCD: The Principle of Maximum Conformality*, *Phys. Rev. D* **86** (2012) 085026 [[arXiv:1107.0338](#)] [[INSPIRE](#)].
- [38] A.L. Kataev and S.V. Mikhailov, *Generalization of the Brodsky-Lepage-Mackenzie optimization within the β -expansion and the principle of maximal conformality*, *Phys. Rev. D* **91** (2015) 014007 [[arXiv:1408.0122](#)] [[INSPIRE](#)].
- [39] A.L. Kataev and S.V. Mikhailov, *The β -expansion formalism in perturbative QCD and its extension*, *JHEP* **11** (2016) 079 [[arXiv:1607.08698](#)] [[INSPIRE](#)].
- [40] M. Czakon, D. Heymes and A. Mitov, *Dynamical scales for multi-TeV top-pair production at the LHC*, *JHEP* **04** (2017) 071 [[arXiv:1606.03350](#)] [[INSPIRE](#)].
- [41] A.V. Garkusha, A.L. Kataev and V.S. Molokoedov, *Renormalization scheme and gauge (in)dependence of the generalized Crewther relation: what are the real grounds of the β -factorization property?*, *JHEP* **02** (2018) 161 [[arXiv:1801.06231](#)] [[INSPIRE](#)].
- [42] X.-G. Wu, J.-M. Shen, B.-L. Du, X.-D. Huang, S.-Q. Wang and S.J. Brodsky, *The QCD renormalization group equation and the elimination of fixed-order scheme-and-scale ambiguities using the principle of maximum conformality*, *Prog. Part. Nucl. Phys.* **108** (2019) 103706 [[arXiv:1903.12177](#)] [[INSPIRE](#)].
- [43] D. Boito, P. Masjuan and F. Olani, *Higher-order QCD corrections to hadronic τ decays from Padé approximants*, *JHEP* **08** (2018) 075 [[arXiv:1807.01567](#)] [[INSPIRE](#)].
- [44] S.J. Brodsky, *The QCD coupling $\alpha_s(Q^2)$ at all momentum scales and the elimination of renormalization scale uncertainties*, [[INSPIRE](#)].
- [45] H.A. Chawdhry and A. Mitov, *Ambiguities of the principle of maximum conformality procedure for hadron collider processes*, *Phys. Rev. D* **100** (2019) 074013 [[arXiv:1907.06610](#)] [[INSPIRE](#)].
- [46] F.A. Chishtie, *Achieving effective renormalization scale and scheme independence via the Principle of Observable Effective Matching (POEM)*, [arXiv:2005.11783](#) [[INSPIRE](#)].
- [47] L. Di Giustino, S.J. Brodsky, S.-Q. Wang and X.-G. Wu, *Infinite-order scale-setting using the principle of maximum conformality: A remarkably efficient method for eliminating renormalization scale ambiguities for perturbative QCD*, *Phys. Rev. D* **102** (2020) 014015 [[arXiv:2002.01789](#)] [[INSPIRE](#)].
- [48] M. Cacciari and N. Houdeau, *Meaningful characterisation of perturbative theoretical uncertainties*, *JHEP* **09** (2011) 039 [[arXiv:1105.5152](#)] [[INSPIRE](#)].
- [49] S. Forte, A. Isgrò and G. Vita, *Do we need N^3 LO Parton Distributions?*, *Phys. Lett. B* **731** (2014) 136 [[arXiv:1312.6688](#)] [[INSPIRE](#)].
- [50] E. Bagnaschi, M. Cacciari, A. Guffanti and L. Jenniches, *An extensive survey of the estimation of uncertainties from missing higher orders in perturbative calculations*, *JHEP* **02** (2015) 133 [[arXiv:1409.5036](#)] [[INSPIRE](#)].

- [51] M. Bonvini, *Probabilistic definition of the perturbative theoretical uncertainty from missing higher orders*, *Eur. Phys. J. C* **80** (2020) 989 [[arXiv:2006.16293](#)] [[INSPIRE](#)].
- [52] G. Grunberg, *Renormalization Group Improved Perturbative QCD*, *Phys. Lett. B* **95** (1980) 70 [*Erratum ibid.* **110** (1982) 501] [[INSPIRE](#)].
- [53] P.M. Stevenson, *Optimized Perturbation Theory*, *Phys. Rev. D* **23** (1981) 2916 [[INSPIRE](#)].
- [54] R.A. Brandt, F. Neri and M.-a. Sato, *Renormalization of Loop Functions for All Loops*, *Phys. Rev. D* **24** (1981) 879 [[INSPIRE](#)].
- [55] G.P. Korchemsky and A.V. Radyushkin, *Renormalization of the Wilson Loops Beyond the Leading Order*, *Nucl. Phys. B* **283** (1987) 342 [[INSPIRE](#)].
- [56] S. Moch, J.A.M. Vermaseren and A. Vogt, *The three loop splitting functions in QCD: The nonsinglet case*, *Nucl. Phys. B* **688** (2004) 101 [[hep-ph/0403192](#)] [[INSPIRE](#)].
- [57] A. Grozin, J.M. Henn, G.P. Korchemsky and P. Marquard, *Three Loop Cusp Anomalous Dimension in QCD*, *Phys. Rev. Lett.* **114** (2015) 062006 [[arXiv:1409.0023](#)] [[INSPIRE](#)].
- [58] A. Grozin, J.M. Henn, G.P. Korchemsky and P. Marquard, *The three-loop cusp anomalous dimension in QCD and its supersymmetric extensions*, *JHEP* **01** (2016) 140 [[arXiv:1510.07803](#)] [[INSPIRE](#)].
- [59] J.M. Henn, A.V. Smirnov, V.A. Smirnov and M. Steinhauser, *A planar four-loop form factor and cusp anomalous dimension in QCD*, *JHEP* **05** (2016) 066 [[arXiv:1604.03126](#)] [[INSPIRE](#)].
- [60] A. von Manteuffel and R.M. Schabinger, *Quark and gluon form factors to four-loop order in QCD: the N_f^3 contributions*, *Phys. Rev. D* **95** (2017) 034030 [[arXiv:1611.00795](#)] [[INSPIRE](#)].
- [61] J. Davies, A. Vogt, B. Ruijl, T. Ueda and J.A.M. Vermaseren, *Large- N_f contributions to the four-loop splitting functions in QCD*, *Nucl. Phys. B* **915** (2017) 335 [[arXiv:1610.07477](#)] [[INSPIRE](#)].
- [62] J. Henn, A.V. Smirnov, V.A. Smirnov, M. Steinhauser and R.N. Lee, *Four-loop photon quark form factor and cusp anomalous dimension in the large- N_c limit of QCD*, *JHEP* **03** (2017) 139 [[arXiv:1612.04389](#)] [[INSPIRE](#)].
- [63] R.H. Boels, T. Huber and G. Yang, *Four-Loop Nonplanar Cusp Anomalous Dimension in $N = 4$ Supersymmetric Yang-Mills Theory*, *Phys. Rev. Lett.* **119** (2017) 201601 [[arXiv:1705.03444](#)] [[INSPIRE](#)].
- [64] R.H. Boels, T. Huber and G. Yang, *The Sudakov form factor at four loops in maximal super Yang-Mills theory*, *JHEP* **01** (2018) 153 [[arXiv:1711.08449](#)] [[INSPIRE](#)].
- [65] S. Moch, B. Ruijl, T. Ueda, J.A.M. Vermaseren and A. Vogt, *Four-Loop Non-Singlet Splitting Functions in the Planar Limit and Beyond*, *JHEP* **10** (2017) 041 [[arXiv:1707.08315](#)] [[INSPIRE](#)].
- [66] S. Moch, B. Ruijl, T. Ueda, J.A.M. Vermaseren and A. Vogt, *On quartic colour factors in splitting functions and the gluon cusp anomalous dimension*, *Phys. Lett. B* **782** (2018) 627 [[arXiv:1805.09638](#)] [[INSPIRE](#)].
- [67] A. Grozin, *Four-loop cusp anomalous dimension in QED*, *JHEP* **06** (2018) 073 [*Addendum ibid.* **01** (2019) 134] [[arXiv:1805.05050](#)] [[INSPIRE](#)].
- [68] R.N. Lee, A.V. Smirnov, V.A. Smirnov and M. Steinhauser, *Four-loop quark form factor with quartic fundamental colour factor*, *JHEP* **02** (2019) 172 [[arXiv:1901.02898](#)] [[INSPIRE](#)].

- [69] J.M. Henn, T. Peraro, M. Stahlhofen and P. Wasser, *Matter dependence of the four-loop cusp anomalous dimension*, *Phys. Rev. Lett.* **122** (2019) 201602 [[arXiv:1901.03693](#)] [[INSPIRE](#)].
- [70] R. Brüser, A. Grozin, J.M. Henn and M. Stahlhofen, *Matter dependence of the four-loop QCD cusp anomalous dimension: from small angles to all angles*, *JHEP* **05** (2019) 186 [[arXiv:1902.05076](#)] [[INSPIRE](#)].
- [71] A. von Manteuffel and R.M. Schabinger, *Quark and gluon form factors in four loop QCD: The N_f^2 and $N_{q\gamma}N_f$ contributions*, *Phys. Rev. D* **99** (2019) 094014 [[arXiv:1902.08208](#)] [[INSPIRE](#)].
- [72] J.M. Henn, G.P. Korchemsky and B. Mistlberger, *The full four-loop cusp anomalous dimension in $\mathcal{N} = 4$ super Yang-Mills and QCD*, *JHEP* **04** (2020) 018 [[arXiv:1911.10174](#)] [[INSPIRE](#)].
- [73] A. von Manteuffel, E. Panzer and R.M. Schabinger, *Cusp and collinear anomalous dimensions in four-loop QCD from form factors*, *Phys. Rev. Lett.* **124** (2020) 162001 [[arXiv:2002.04617](#)] [[INSPIRE](#)].
- [74] N. Beisert, B. Eden and M. Staudacher, *Transcendentality and Crossing*, *J. Stat. Mech.* **0701** (2007) P01021 [[hep-th/0610251](#)] [[INSPIRE](#)].
- [75] S. Moch, J.A.M. Vermaseren and A. Vogt, *Higher-order corrections in threshold resummation*, *Nucl. Phys. B* **726** (2005) 317 [[hep-ph/0506288](#)] [[INSPIRE](#)].
- [76] K. Melnikov and T.v. Ritbergen, *The three loop relation between the $\overline{MS}$ -bar and the pole quark masses*, *Phys. Lett. B* **482** (2000) 99 [[hep-ph/9912391](#)] [[INSPIRE](#)].
- [77] P. Marquard, A.V. Smirnov, V.A. Smirnov and M. Steinhauser, *Quark Mass Relations to Four-Loop Order in Perturbative QCD*, *Phys. Rev. Lett.* **114** (2015) 142002 [[arXiv:1502.01030](#)] [[INSPIRE](#)].
- [78] P. Marquard, A.V. Smirnov, V.A. Smirnov, M. Steinhauser and D. Wellmann, *$\overline{MS}$ -on-shell quark mass relation up to four loops in QCD and a general $SU(N)$ gauge group*, *Phys. Rev. D* **94** (2016) 074025 [[arXiv:1606.06754](#)] [[INSPIRE](#)].
- [79] A.L. Kataev and V.S. Molokoedov, *Multiloop contributions to the on-shell $\overline{MS}$ heavy quark mass relation in QCD and the asymptotic structure of the corresponding series: the updated consideration*, *Eur. Phys. J. C* **80** (2020) 1160 [[arXiv:1807.05406](#)] [[INSPIRE](#)].
- [80] P. Mohr, D.B. Newell, B.N. Taylor and E. Tiesinga, *Data and analysis for the CODATA 2017 special fundamental constants adjustment*, *Metrologia* **55** (2018) 125.
- [81] J.S. Schwinger, *On quantum electrodynamics and the magnetic moment of the electron*, *Phys. Rev.* **73** (1948) 416 [[INSPIRE](#)].
- [82] A. Petermann, *Fourth order magnetic moment of the electron*, *Helv. Phys. Acta* **30** (1957) 407 [[INSPIRE](#)].
- [83] C.M. Sommerfield, *Magnetic Dipole Moment of the Electron*, *Phys. Rev.* **107** (1957) 328 [[INSPIRE](#)].
- [84] S. Laporta and E. Remiddi, *The analytical value of the electron ($g-2$) at order α^3 in QED*, *Phys. Lett. B* **379** (1996) 283 [[hep-ph/9602417](#)] [[INSPIRE](#)].
- [85] T. Aoyama, M. Hayakawa, T. Kinoshita and M. Nio, *Tenth-Order QED Contribution to the Electron $g-2$ and an Improved Value of the Fine Structure Constant*, *Phys. Rev. Lett.* **109** (2012) 111807 [[arXiv:1205.5368](#)] [[INSPIRE](#)].

- [86] T. Aoyama, M. Hayakawa, T. Kinoshita and M. Nio, *Tenth-Order Electron Anomalous Magnetic Moment — Contribution of Diagrams without Closed Lepton Loops*, *Phys. Rev. D* **91** (2015) 033006 [Erratum *ibid.* **96** (2017) 019901] [[arXiv:1412.8284](#)] [[INSPIRE](#)].
- [87] S. Laporta, *High-precision calculation of the 4-loop contribution to the electron $g-2$ in QED*, *Phys. Lett. B* **772** (2017) 232 [[arXiv:1704.06996](#)] [[INSPIRE](#)].
- [88] T. Aoyama, T. Kinoshita and M. Nio, *Revised and Improved Value of the QED Tenth-Order Electron Anomalous Magnetic Moment*, *Phys. Rev. D* **97** (2018) 036001 [[arXiv:1712.06060](#)] [[INSPIRE](#)].
- [89] S. Volkov, *Numerical calculation of high-order QED contributions to the electron anomalous magnetic moment*, *Phys. Rev. D* **98** (2018) 076018 [[arXiv:1807.05281](#)] [[INSPIRE](#)].
- [90] T. Aoyama, T. Kinoshita and M. Nio, *Theory of the Anomalous Magnetic Moment of the Electron*, *Atoms* **7** (2019) 28 [[INSPIRE](#)].
- [91] R. Karplus and A. Klein, *Electrodynamics displacement of atomic energy levels. 3. The hyperfine structure of positronium*, *Phys. Rev.* **87** (1952) 848 [[INSPIRE](#)].
- [92] G.P. Lepage, *Analytic Bound State Solutions in a Relativistic Two-Body Formalism with Applications in Muonium and Positronium*, *Phys. Rev. A* **16** (1977) 863 [[INSPIRE](#)].
- [93] S. Karshenboim, *New logarithmic contributions in muonium and positronium*, *J. Exp. Theor. Phys.* **76** (1993) 541.
- [94] A. Czarnecki, K. Melnikov and A. Yelkhovsky, *Positronium hyperfine splitting: Analytical value at $O(m\alpha^6)$* , *Phys. Rev. Lett.* **82** (1999) 311 [[hep-ph/9809341](#)] [[INSPIRE](#)].
- [95] B.A. Kniehl and A.A. Penin, *Order $\alpha^7 \ln(1/\alpha)$ contribution to positronium hyperfine splitting*, *Phys. Rev. Lett.* **85** (2000) 5094 [[hep-ph/0010159](#)] [[INSPIRE](#)].
- [96] K. Pachucki, A. Czarnecki, U.D. Jentschura and V.A. Yerokhin, *Complete two-loop correction to the bound-electron g factor*, *Phys. Rev. A* **72** (2005) 022108 [[physics/0506227](#)] [[INSPIRE](#)].
- [97] K. Pachucki and M. Puchalski, *One-loop binding corrections to the electron g factor*, *Phys. Rev. A* **96** (2017) 032503 [[arXiv:1707.08518](#)] [[INSPIRE](#)].
- [98] A. Czarnecki, M. Dowling, J. Piclum and R. Szafron, *Two-loop binding corrections to the electron gyromagnetic factor*, *Phys. Rev. Lett.* **120** (2018) 043203 [[arXiv:1711.00190](#)] [[INSPIRE](#)].
- [99] A. Czarnecki, J. Piclum and R. Szafron, *Logarithmically enhanced Euler-Heisenberg Lagrangian contribution to the electron gyromagnetic factor*, *Phys. Rev. A* **102** (2020) 050801 [[arXiv:2008.07550](#)] [[INSPIRE](#)].
- [100] H.M. Georgi, S.L. Glashow, M.E. Machacek and D.V. Nanopoulos, *Higgs Bosons from Two Gluon Annihilation in Proton Proton Collisions*, *Phys. Rev. Lett.* **40** (1978) 692 [[INSPIRE](#)].
- [101] S. Dawson, *Radiative corrections to Higgs boson production*, *Nucl. Phys. B* **359** (1991) 283 [[INSPIRE](#)].
- [102] A. Djouadi, M. Spira and P.M. Zerwas, *Production of Higgs bosons in proton colliders: QCD corrections*, *Phys. Lett. B* **264** (1991) 440 [[INSPIRE](#)].
- [103] D. Graudenz, M. Spira and P.M. Zerwas, *QCD corrections to Higgs boson production at proton proton colliders*, *Phys. Rev. Lett.* **70** (1993) 1372 [[INSPIRE](#)].

- [104] M. Spira, A. Djouadi, D. Graudenz and P.M. Zerwas, *Higgs boson production at the LHC*, *Nucl. Phys. B* **453** (1995) 17 [[hep-ph/9504378](#)] [[INSPIRE](#)].
- [105] R. Harlander and P. Kant, *Higgs production and decay: Analytic results at next-to-leading order QCD*, *JHEP* **12** (2005) 015 [[hep-ph/0509189](#)] [[INSPIRE](#)].
- [106] U. Aglietti, R. Bonciani, G. Degrossi and A. Vicini, *Analytic Results for Virtual QCD Corrections to Higgs Production and Decay*, *JHEP* **01** (2007) 021 [[hep-ph/0611266](#)] [[INSPIRE](#)].
- [107] U. Aglietti, R. Bonciani, G. Degrossi and A. Vicini, *Two loop light fermion contribution to Higgs production and decays*, *Phys. Lett. B* **595** (2004) 432 [[hep-ph/0404071](#)] [[INSPIRE](#)].
- [108] C. Anastasiou, S. Beerli, S. Bucherer, A. Daleo and Z. Kunszt, *Two-loop amplitudes and master integrals for the production of a Higgs boson via a massive quark and a scalar-quark loop*, *JHEP* **01** (2007) 082 [[hep-ph/0611236](#)] [[INSPIRE](#)].
- [109] C. Anastasiou, S. Bucherer and Z. Kunszt, *HPro: A NLO Monte-Carlo for Higgs production via gluon fusion with finite heavy quark masses*, *JHEP* **10** (2009) 068 [[arXiv:0907.2362](#)] [[INSPIRE](#)].
- [110] R.V. Harlander and W.B. Kilgore, *Next-to-next-to-leading order Higgs production at hadron colliders*, *Phys. Rev. Lett.* **88** (2002) 201801 [[hep-ph/0201206](#)] [[INSPIRE](#)].
- [111] C. Anastasiou and K. Melnikov, *Higgs boson production at hadron colliders in NNLO QCD*, *Nucl. Phys. B* **646** (2002) 220 [[hep-ph/0207004](#)] [[INSPIRE](#)].
- [112] V. Ravindran, J. Smith and W.L. van Neerven, *NNLO corrections to the total cross-section for Higgs boson production in hadron hadron collisions*, *Nucl. Phys. B* **665** (2003) 325 [[hep-ph/0302135](#)] [[INSPIRE](#)].
- [113] C. Anastasiou, C. Duhr, F. Dulat, F. Herzog and B. Mistlberger, *Higgs Boson Gluon-Fusion Production in QCD at Three Loops*, *Phys. Rev. Lett.* **114** (2015) 212001 [[arXiv:1503.06056](#)] [[INSPIRE](#)].
- [114] C. Anastasiou et al., *High precision determination of the gluon fusion Higgs boson cross-section at the LHC*, *JHEP* **05** (2016) 058 [[arXiv:1602.00695](#)] [[INSPIRE](#)].
- [115] B. Mistlberger, *Higgs boson production at hadron colliders at N³LO in QCD*, *JHEP* **05** (2018) 028 [[arXiv:1802.00833](#)] [[INSPIRE](#)].
- [116] F. Dulat, A. Lazopoulos and B. Mistlberger, *iHixs 2 — Inclusive Higgs cross sections*, *Comput. Phys. Commun.* **233** (2018) 243 [[arXiv:1802.00827](#)] [[INSPIRE](#)].
- [117] P. Bolzoni, F. Maltoni, S.-O. Moch and M. Zaro, *Higgs production via vector-boson fusion at NNLO in QCD*, *Phys. Rev. Lett.* **105** (2010) 011801 [[arXiv:1003.4451](#)] [[INSPIRE](#)].
- [118] P. Bolzoni, F. Maltoni, S.-O. Moch and M. Zaro, *Vector boson fusion at NNLO in QCD: SM Higgs and beyond*, *Phys. Rev. D* **85** (2012) 035002 [[arXiv:1109.3717](#)] [[INSPIRE](#)].
- [119] F.A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs Production at Three Loops in QCD*, *Phys. Rev. Lett.* **117** (2016) 072001 [[arXiv:1606.00840](#)] [[INSPIRE](#)].
- [120] F.A. Dreyer, A. Karlberg and L. Tancredi, *On the impact of non-factorisable corrections in VBF single and double Higgs production*, *JHEP* **10** (2020) 131 [[arXiv:2005.11334](#)] [[INSPIRE](#)].
- [121] L.-S. Ling, R.-Y. Zhang, W.-G. Ma, L. Guo, W.-H. Li and X.-Z. Li, *NNLO QCD corrections to Higgs pair production via vector boson fusion at hadron colliders*, *Phys. Rev. D* **89** (2014) 073001 [[arXiv:1401.7754](#)] [[INSPIRE](#)].

- [122] F.A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs Pair Production at N^3LO* , *Phys. Rev. D* **98** (2018) 114016 [[arXiv:1811.07906](#)] [[INSPIRE](#)].
- [123] T. Han, G. Valencia and S. Willenbrock, *Structure function approach to vector boson scattering in pp collisions*, *Phys. Rev. Lett.* **69** (1992) 3274 [[hep-ph/9206246](#)] [[INSPIRE](#)].
- [124] M. Cacciari, F. Dreyer, A. Karlberg, G. Salam and G. Zanderighi, *proVBFH*, <https://probfh.hepforge.org/>.
- [125] G. Altarelli, R.K. Ellis and G. Martinelli, *Leptoproduction and Drell-Yan Processes Beyond the Leading Approximation in Chromodynamics*, *Nucl. Phys. B* **143** (1978) 521.
- [126] R. Hamberg, W.L. van Neerven and T. Matsuura, *A complete calculation of the order α_s^2 correction to the Drell-Yan K factor*, *Nucl. Phys. B* **359** (1991) 343 [Erratum *ibid.* **644** (2002) 403] [[INSPIRE](#)].
- [127] C. Duhr, F. Dulat and B. Mistlberger, *Drell-Yan Cross Section to Third Order in the Strong Coupling Constant*, *Phys. Rev. Lett.* **125** (2020) 172001 [[arXiv:2001.07717](#)] [[INSPIRE](#)].
- [128] G. Altarelli, R.K. Ellis and G. Martinelli, *Large Perturbative Corrections to the Drell-Yan Process in QCD*, *Nucl. Phys. B* **157** (1979) 461 [[INSPIRE](#)].
- [129] P. Aurenche and J. Lindfors, *QCD Corrections to Direct Lepton Production in Hadronic Collisions*, *Nucl. Phys. B* **185** (1981) 274 [[INSPIRE](#)].
- [130] J. Kubar-Andre and F.E. Paige, *Gluon Corrections to the Drell-Yan Model*, *Phys. Rev. D* **19** (1979) 221 [[INSPIRE](#)].
- [131] C. Anastasiou, L.J. Dixon, K. Melnikov and F. Petriello, *High precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at NNLO*, *Phys. Rev. D* **69** (2004) 094008 [[hep-ph/0312266](#)] [[INSPIRE](#)].
- [132] K. Melnikov and F. Petriello, *The W boson production cross section at the LHC through $O(\alpha_s^2)$* , *Phys. Rev. Lett.* **96** (2006) 231803 [[hep-ph/0603182](#)] [[INSPIRE](#)].
- [133] C. Duhr, F. Dulat and B. Mistlberger, *Charged current Drell-Yan production at N^3LO* , *JHEP* **11** (2020) 143 [[arXiv:2007.13313](#)] [[INSPIRE](#)].
- [134] J. Currie, T. Gehrmann, E.W.N. Glover, A. Huss, J. Niehues and A. Vogt, *N^3LO corrections to jet production in deep inelastic scattering using the Projection-to-Born method*, *JHEP* **05** (2018) 209 [[arXiv:1803.09973](#)] [[INSPIRE](#)].
- [135] T. Gehrmann, N. Glover, A. Huss and J. Whitehead, *Scale and isolation sensitivity of diphoton distributions at the LHC*, *JHEP* **01** (2021) 108 [[arXiv:2009.11310](#)] [[INSPIRE](#)].
- [136] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and D.M. Walker, *Next-to-Next-to-Leading-Order QCD Corrections to the Transverse Momentum Distribution of Weak Gauge Bosons*, *Phys. Rev. Lett.* **120** (2018) 122001 [[arXiv:1712.07543](#)] [[INSPIRE](#)].
- [137] X. Chen, T. Gehrmann, E.W.N. Glover and A. Huss, *Fiducial cross sections for the four-lepton decay mode in Higgs-plus-jet production up to NNLO QCD*, *JHEP* **07** (2019) 052 [[arXiv:1905.13738](#)] [[INSPIRE](#)].
- [138] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and J. Pires, *Infrared sensitivity of single jet inclusive production at hadron colliders*, *JHEP* **10** (2018) 155 [[arXiv:1807.03692](#)] [[INSPIRE](#)].
- [139] J.A. Hoeting, D. Madigan, A.E. Raftery and C.T. Volinsky, *Bayesian Model Averaging: A Tutorial*, *Statist. Sci.* **14** (1999) 382.