Probabilistic definition of the perturbative theoretical uncertainty from missing higher orders

Marco Bonvini^a

INFN, Sezione di Roma 1, Piazzale Aldo Moro 5, 00185 Rome, Italy

Received: 6 July 2020 / Accepted: 13 October 2020 / Published online: 27 October 2020
 © The Author(s) 2020

Abstract We consider the problem of quantifying the uncertainty on theoretical predictions based on perturbation theory due to missing higher orders. The most widely used approach, scale variation, is largely arbitrary and it has no probabilistic foundation, making it not suitable for robust data analysis. In 2011, Cacciari and Houdeau proposed a model based on a Bayesian approach to provide a probabilistic definition of the theory uncertainty from missing higher orders. In this work, we propose an improved version of the Cacciari–Houdeau model, that overcomes some limitations. In particular, it performs much better in case of perturbative expansions with large high-order contributions (as it often happens in QCD). In addition, we propose an alternative model based on the same idea of scale variation, which overcomes some of the shortcomings of the canonical approach, on top of providing a probabilistically-sound result. Moreover, we address the problem of the dependence of theoretical predictions on unphysical scales (such as the renormalization scale), and propose a solution to obtain a scale-independent result within the probabilistic framework. We validate these methods on expansions with known sums, and apply them to a number of physical observables in particle physics. We also investigate some variations, improvements and combinations of the models. We believe that these methods provide a powerful tool to reliably estimate theory uncertainty from missing higher orders that can be used in any physics analysis. The results of this work are easily accessible through a public code named THunc.

Contents

1 Introduction	2
2 Preliminaries	4
2.1 Basic concepts and assumptions on perturbative expansions	4

2.2 Constructing a probability distribution for a physical observable	6
2.3 The role of unphysical scales	7
2.4 The Cacciari–Houdeau method	9
2.5 A working example: Higgs production at the LHC	10
3 Model-independent features of the inference approach at fixed scale	11
3.1 Basic notations	11
3.2 Definition of the scale-dependence numbers	12
3.3 General features of the models	13
4 Model 1: geometric behaviour model	15
4.1 The hypothesis of the model	15
4.2 The choice of priors	17
4.3 Inference on the unknown higher orders	18
4.4 The posterior of the hidden parameters	19
4.5 Representative results	19
4.6 Dependence on the prior	22
5 Model 2: a new approach using scale variation information	23
5.1 The hypothesis of the model	23
5.2 The choice of prior	24
5.3 Inference on the unknown higher orders	25
5.4 The posterior of the hidden parameter	25
5.5 Representative results	26
5.6 Dependence on the prior	28
6 Dealing with scale dependence	28
6.1 The scale as a model parameter	28
6.2 A pre-example: application to the Cacciari–Houdeau model	31
6.3 Scale-independent geometric behaviour model	31
6.4 Scale-independent scale variation model	33
7 Validation and applications	35
7.1 Convergent series	35
7.2 Divergent series	37
7.3 Anharmonic oscillator in quantum mechanics	40
7.4 Higgs production in the threshold limit	42
7.5 e^+e^- into hadrons at $N^4\text{LO}$	43
7.6 Higgs decay to gg at $N^4\text{LO}$	45

^a e-mail: marco.bonvini@roma1.infn.it (corresponding author)

7.7 A resummed logarithmic-ordered expansion . . .	47
8 Treating correlations between different theory predictions	49
8.1 Correlations for the same observable at different kinematic points	50
8.2 Correlations among observables of different processes	50
9 Conclusions	52
A Algorithms	53
A.1 Computing renormalization scale dependence at fixed order	53
A.2 Computing the coefficients for the geometric behaviour model	55
B Possible ways of improving the models	56
B.1 Scale variation model with violation of the bound	56
B.2 Scale variation model with controlled scale dependence	57
B.3 Mixing models	59
B.4 Custom model	60
B.5 Adding information on the sign of each order	60
References	63

1 Introduction

Computing exactly physical observables in a generic quantum field theory (QFT) is still an open challenge. The closest approximation to such a computation is achieved through numerical simulations of the theory on a discretized space-time (lattice). This approach, however, cannot be used universally. For instance, for some observables (e.g. involving large momenta) the required computing power would be out of reach. A complementary approach, applicable when the coupling of the theory is sufficiently small, is the so-called perturbative approach, or perturbation theory. Namely, physical observables are expressed as a power series in the coupling. Computing the coefficients of this power series can be done systematically, e.g. using Feynman diagram techniques, but the complication of the calculation grows considerably with the perturbative order. Therefore, in practice, only a limited number of coefficients is achievable for a given observable.

This poses an important question: how far is the approximate perturbative result from the unknown exact one? The answer to this question can be cast in an uncertainty, usually called *theory uncertainty from missing higher orders*. Since the exact result is unknown, one can only quantify this uncertainty in terms of a *probability distribution*. How to determine such a distribution is the subject of this paper.

The most standard and widespread approach to estimate this uncertainty is the so-called scale variation method. This approach is based on the observation that physical observables do not depend on unphysical scales (such as the renormalization scale) appearing in QFTs. However, this indepen-

dence is strictly valid only for the exact result. Any approximate result computed in perturbation theory will in fact depend on unphysical scales, the dependence being formally of higher order. The idea is thus to estimate the theory uncertainty by varying the scale, as the result will accordingly change by an amount that is formally of the same order as what this uncertainty wants to quantify. While this idea is certainly valid and powerful, the canonical method used to exploit it has various caveats. Indeed, the canonical recipe consists in varying the unphysical scale μ by a factor of two about a central value μ_0 of choice, and then using the maximal variation of the observable with respect to the value at central scale as a measure of the uncertainty. This is depicted in Fig. 1. In formulae, the canonical way of writing a result based on a perturbative computation of a physical observable Σ is¹

$$\Sigma \stackrel{\text{canonical scale variation}}{\equiv} \Sigma_{\text{pert}}(\mu_0) \pm \max_{\mu_0/2 < \mu < 2\mu_0} \left| \Sigma_{\text{pert}}(\mu) - \Sigma_{\text{pert}}(\mu_0) \right|, \quad (1.1)$$

where $\Sigma_{\text{pert}}(\mu)$ is the scale dependent perturbative result, $\Sigma_{\text{pert}}(\mu_0)$ represents the “central value” of the prediction, and the scale variation is appended as a theory “error”. The caveats of this approach are apparent:

- it is largely arbitrary, in the choice of the central scale μ_0 and of the interval of variation (factor of two);
- in the vicinity of stationary points the uncertainty can become accidentally small;
- the uncertainty has no probabilistic interpretation.

The latter point can be overcome by *assigning* an interpretation (for instance, the “error” could be interpreted as the standard deviation of a gaussian distribution), but any choice would be totally arbitrary. In addition to all this, it is well known that the scale variation uncertainty often underestimates the size of higher order contributions.²

¹ In the literature there exist variants of this formulation. In the most widespread variant the “error” is left asymmetric, as shown in Fig. 1.

² It has to be mentioned that at low orders not all partonic channels contribute to a given QCD process with hadrons in the initial state. Since renormalization scale variation cannot probe different channels, the opening up of new channels certainly contributes to the failure of scale variation in estimating the size of the higher orders. In proton-proton collisions, from next-to-next-to-leading order (NNLO) onwards all partonic channels are active, and indeed for those processes that are known beyond NNLO scale variation typically works better than at previous orders. A remarkable exception is given by the Drell-Yan process, for which the recently computed next-to-next-to-next-to-leading order correction [1, 2] lies outside the NNLO scale variation band in a wide kinematic region.

Fig. 1 Schematic representation of the canonical method to estimate theory uncertainty, namely the scale variation approach. Since the left and right variations of the scale generally lead to different sizes of variation of the perturbative result (shown by the two double-headed arrows), one may either choose to keep the uncertainty asymmetric, or to select the largest and symmetrize it

In 2011, Cacciari and Houdeau [3] proposed a completely different approach to estimate the uncertainty from missing higher orders. In their groundbreaking work they constructed a probabilistic model to define this uncertainty, based on some assumptions on the progression of the perturbative expansion. Roughly speaking, the model assumes that the coefficients of the power expansion in the coupling are bounded by an unknown (hidden) parameter. The knowledge of the first few orders of the perturbative expansion allow to perform (Bayesian) inference on the hidden parameter, whose improved knowledge can in turn be used to make inference on the unknown subsequent perturbative coefficients. Once the likelihood of the perturbative coefficients given the hidden parameter and the prior distribution of the hidden parameter are defined, the model produces probability distributions for the unknown coefficients, and thus for the physical observable itself.

This approach to theory uncertainties is unquestionably superior and more elegant than canonical scale variation, mainly because it is probabilistically founded. However, it also has limitations. For instance, while it performs well for QCD observables at e^+e^- colliders [3], it leads to less reliable results in the case of proton-proton collider observables [4, 5], that are usually affected by larger perturbative corrections. Perhaps for this reason, in conjunction with the simplicity and the deep-rooted attitude of using the scale variation method, the Cacciari–Houdeau (CH) approach is not very popular in the high-energy physics community, with only few applications [5–18], most of which in the context of effective theories. Moreover, the CH approach does not deal with the unphysical scale dependence of the result – the CH prediction is computed for a given choice of the scale, so the final probability distributions is de facto scale dependent.

In this paper, we use a Bayesian approach to build new probabilistic models, similar to the CH model, that overcome the limitations of the previous approach. The inference struc-

Fig. 2 The general structure of the inference in any model considered in this work

ture of any model that we will consider is depicted in full generality in Fig. 2. Starting from the known orders, under some model assumptions one can make inference on the (hidden) parameters characterizing the model, to be used in turn to infer the probability distribution of the unknown higher orders. In formulae, we have, schematically,

$$P(\text{unknown orders}|\text{known orders}) = \int d\text{pars} P(\text{unknown orders}|\text{pars}) P(\text{pars}|\text{known orders}) \quad (1.2)$$

where $P(A|B)$ is the conditional probability distribution of A given B . Eq. (1.2) is given in terms of the posterior distribution of the hidden parameters

$$P(\text{pars}|\text{known orders}) \propto P(\text{known orders}|\text{pars}) P_0(\text{pars}) \quad (1.3)$$

which depends on the prior distribution $P_0(\text{pars})$ of the hidden parameters and on the model assumptions through the likelihood $P(\text{orders}|\text{pars})$, that appears also explicitly in Eq. (1.2). These two model-dependent ingredients are sufficient to let Bayesian inference work, and can be used to eventually construct a probability distribution for the observable, that contains all the information on the uncertainty from missing higher orders.

In this work, we propose two main models: one is an improved version of the CH model that efficiently describes perturbative expansions with large perturbative corrections; the other is a model inspired by the scale variation method, but constructed in such a way to be reliable and probabilistically sound. Both methods outperform the current approaches in terms of reliability. They are also sufficiently general to be used for perturbative expansion that are not necessarily fixed-order expansions in powers of the coupling, but can be for instance resummed expansions or other generalized expansions. Moreover, the first model can be applied also beyond quantum field theory, for instance to quantum mechanics or in general to any expansion that behaves perturbatively. We also explore variants of the methods and combinations of them.

These methods are still applied to the perturbative expansion at a given fixed value of the unphysical scale. An impor-

tant and innovative development proposed in this work is a way to “remove” the scale dependence of the result. This is achieved by dealing with the scale dependence within the probabilistic framework, and leads to a result that is to a large extent scale independent. As a byproduct of this procedure, the central value of the prediction (identified with the mean of the probability distribution for the observable) does not necessarily correspond to the canonical perturbative result at a given “central” scale. Thanks to this feature our method improves not only the reliability of the uncertainty, but also that of the central prediction for the observable.

To facilitate the adoption of the results of this work, a computer code named `THunc` is publicly released. The code is very easy to use: the user provides the perturbative expansion, and the code outputs the probability distribution of the result, together with a number of statistical estimators (mean, mode, median, standard deviation, degree-of-belief intervals). It is also fast and efficient, and flexible as it allows the user to define customized models.

The structure of this paper is the following. In Sect. 2 we give some preliminary information on the perturbative expansions and their scale dependence, we provide some basic concepts on the probabilistic definition of the theory uncertainty as well as a brief recap of the CH method, and we define a working example to be used in the subsequent sections. In Sect. 3 we start defining some notations and describe general features common to all models we will later consider. We then move on presenting our two main models (at fixed scale) in Sects. 4 and 5. We propose a way to construct scale-independent results and uncertainties in Sect. 6. We then validate our methods in Sect. 7, where we also consider some realistic applications. In Sect. 8, we discuss the issue of defining correlations between theory uncertainties. After concluding in Sect. 9, we collect details on numerical implementations in Appendix A and propose a number of variants and possible improvements in Appendix B.

2 Preliminaries

2.1 Basic concepts and assumptions on perturbative expansions

We consider a (renormalizable) quantum field theory (QFT) depending on a single coupling α . We will often refer to practical examples in quantum chromodynamics (QCD), since its coupling is not very small and thus perturbation theory produces somewhat large high-order corrections, which is the case where reliably estimating theory uncertainties is both important and challenging. We focus on a generic physical observable Σ . The theory predicts a unique well-defined value for this observable, that we call Σ_{true} . This is the value we aim to obtain. However, we usually cannot compute it

exactly, and we thus use a perturbative approach to approximate it.

According to the perturbative hypothesis, namely that the dimensionless coupling α of the theory is sufficiently small, it is possible to compute the observable Σ as a power series in the coupling itself. We then write

$$\Sigma_{\text{true}} \simeq \sum_{k=0}^n c_k \alpha^k, \quad (2.1)$$

where c_k are the coefficients of the perturbative expansion, and n is some order at which we stop the expansion. Eq. (2.1) is not an equality. One may be tempted to think that if $n \rightarrow \infty$, then it would become an equality. This limit, however, does not exist, as perturbative expansions are divergent [19, 20]. The divergence of the series is related to the fact that Σ_{true} is a non-analytic function of the coupling α in $\alpha = 0$. One may try to treat the divergent series using e.g. the Borel summation method. For some known (to all orders) divergent contributions, such as those due to renormalons (see e.g. Ref. [21]), one can obtain a finite result through Borel summation.³ However, it is not guaranteed that the Borel-sum of the series captures the full result: there may be intrinsically non-perturbative contributions that cannot be reconstructed from the perturbative expansion.⁴ Moreover, in order to use the Borel summation method, the series should be known to all orders, or at least its asymptotic behaviour should be known.⁵ However, this is usually not the case, so in practice the only information we have about an observable is its (truncated) perturbative expansion, Eq. (2.1).

The fact that the series is divergent and that summation methods cannot be used may suggest that the perturbative result is useless. However, this is in contrast with the well known fact that perturbation theory works, namely it predicts results in decent (or even good) agreement with data, at least when the coupling is sufficiently small (for instance, in QED it works very well). The explanation of this fact relies on the assumption that perturbative series are asymptotic expan-

³ In fact, the renormalon series is not Borel summable, but one can still define a finite sum at the price of introducing an ambiguity proportional to a non-analytic function of the coupling. The ambiguity term is clearly an artefact of the summation method, and it suggests that the information contained in the perturbative expansion is not sufficient to fully reconstruct the true value of the observable.

⁴ It has been pointed out in Ref. [22] (in the context of quantum mechanics) that it is possible to suitably modify the action such that the (Borel-summed) perturbative result gives the whole exact result, without any extra non-perturbative contribution. It is not obvious whether this result can be extended to QFTs.

⁵ If one applies the Borel method to a finite sum, it acts as an identity. However, if the asymptotic behaviour is known, one can mix the Borel method with a conformal mapping, to obtain a “perturbative version” of Borel summation that is effective also on finite sums (see e.g. Refs. [23, 24]).

Fig. 3 The asymptotic expansion of the function Eq. (2.2) truncated at various orders from 0 to 15 (orange dots), normalized to the value of the function itself (blue line). The lower panel shows the absolute differ-

ence between the truncated and exact results. The left plot corresponds to $\alpha = 0.13$, the right plot corresponds to $\alpha = 0.2$

sions of the exact result. To our knowledge, there is no general proof of this statement, but it seems very reasonable and we take it as valid. The asymptotic nature of the perturbative expansion implies that up to some order k_{asympt} adding terms to the expansion improves the accuracy of the prediction, but beyond k_{asympt} the divergent contributions to the series dominate and the sum explodes. A visual example of this fact is shown in Fig. 3, where the following non-analytic function of α is compared with its asymptotic expansion at small α :

$$\frac{1}{\alpha} \exp\left(\frac{1}{\alpha}\right) \Gamma\left(0, \frac{1}{\alpha}\right) \stackrel{\text{asympt}}{=} \sum_{k=0}^{\infty} (-1)^k k! \alpha^k. \quad (2.2)$$

From the figure one sees that for a sufficiently small value of $\alpha = 0.13$ (left plot) the first few (approximately 7) orders give a good approximation of the result, showing an apparently converging behaviour. However, adding extra orders deteriorates the prediction, because the factorial growth of the series wins over the power suppression. The best prediction one can make using the asymptotic expansion is thus obtained truncating it to $k_{\text{asympt}} \sim 7$. This prediction, however, has an irreducible uncertainty due to the truncation itself,⁶ of the size of the last term included in the truncated expansion. The right plot, that is obtained with a larger coupling $\alpha = 0.2$, shows that this irreducible uncertainty grows with the value

of the coupling, and the value of k_{asympt} decreases accordingly.⁷

From these considerations we can conclude that a physical observable can be expressed as the sum

$$\Sigma_{\text{true}} = \sum_{k=0}^{k_{\text{asympt}}} c_k \alpha^k + \Delta_{\text{asympt}} + \Delta_{\text{non-pert}}, \quad (2.3)$$

where the first term is perturbative expansion is truncated at k_{asympt} , which is the best prediction we can make using perturbation theory, while Δ_{asympt} represents the irreducible difference between the truncated expansion and its all-order sum. The $\Delta_{\text{non-pert}}$ term, instead, represents possible intrinsically non-perturbative contributions that cannot be captured by perturbation theory. The general expectation on the size of these contribution is

$$|\Delta_{\text{asympt}}| \sim |\Delta_{\text{non-pert}}| \ll \left| \sum_{k=0}^{k_{\text{asympt}}} c_k \alpha^k \right|, \quad (2.4)$$

at least when the coupling is sufficiently small. This expectation is again not a proof, but a consequence of the goodness of perturbation theory, together with the fact that both Δ terms are known to be exponentially suppressed by $\exp(-a/\alpha)$, $a > 0$, in some established cases (namely, when $\Delta_{\text{non-pert}}$ contains instanton contributions and when Δ_{asympt} is domi-

⁶ Note that this divergent series is in fact Borel summable, and its Borel sum corresponds exactly to the original function Eq. (2.2).

⁷ Indeed, for this specific factorially divergent series, one can show that $k_{\text{asympt}} \sim 1/\alpha$ and the irreducible uncertainty is of order $\exp(-1/\alpha)$ [21].

nated by factorially divergent contributions like renormalons [21]).

Unfortunately, since we do not generally know the asymptotic behaviour of the expansion, we cannot know a priori the value of k_{asympt} , and we thus cannot truncate the expansion at the optimal value. This is not a real issue, as in most cases we know just a rather small number of orders, typically two or three, corresponding to next-to-leading order (NLO) and next-to-next-to-leading order (NNLO) computations. Only in very few cases we know physical quantities at $N^3\text{LO}$ (four terms in the expansion) or beyond. We thus expect (and assume) the number n of known orders to be smaller than k_{asympt} .⁸ Therefore, we can rewrite Eq. (2.3) as

$$\Sigma_{\text{true}} = \sum_{k=0}^n c_k \alpha^k + \Delta_{\text{MHO}}^{(n)} + \Delta_{\text{asympt}} + \Delta_{\text{non-pert}}, \quad (2.5)$$

having defined the contribution from missing higher orders

$$\Delta_{\text{MHO}}^{(n)} = \sum_{k=n+1}^{k_{\text{asympt}}} c_k \alpha^k, \quad (2.6)$$

where n is the highest known perturbative order. Assuming that n is sufficiently smaller than k_{asympt} , using similar considerations to those that led to Eq. (2.4) we can conclude that in most cases the contribution from missing higher orders is larger than the asymptotic and non-perturbative contributions,

$$|\Delta_{\text{MHO}}^{(n)}| \gg |\Delta_{\text{asympt}}| \sim |\Delta_{\text{non-pert}}|. \quad (2.7)$$

This implies that our knowledge of the observable Σ is determined by the perturbative expansion truncated at order n with an uncertainty that is dominated by the missing higher order term $\Delta_{\text{MHO}}^{(n)}$. Quantifying this term, or better determining its probability distribution, is the main task of the rest of this paper.

2.2 Constructing a probability distribution for a physical observable

Defining the theoretical uncertainty from missing higher orders in a probabilistic way may sound impossible or completely arbitrary to many. This would certainly be true in the

context of the so-called frequentist approach to probability, in which the definition of a probability requires the existence of a repeatable event, which is typically the case for an experiment but clearly not the case for a theoretical prediction. However, the frequentist approach to probability is not the only one – actually, the frequentist formulation is mathematically inconsistent (see e.g. [28]), and thus certainly not the best one.

The only mathematically correct formulation of a probability theory is the so-called Bayesian approach, where the probability is defined as the “degree of belief” of an event, which is then intrinsically subjective. Initially, when no information is available, the probability of an event is given by a *prior* distribution, which encodes our subjective and arbitrary prejudices. Acquiring information on the event changes the degree of belief through statistical inference (Bayes theorem). Therefore, any probability will depend on subjective assumptions through the prior distribution, but adding more and more information updates the probability making it less and less dependent on the prior.

In case of repeatable events, one can acquire information on the process by repeating them (and, in the limit of large number of repetitions, one recovers the frequentist result). However, repetition is not the only way of acquiring information, and thus one can use the Bayesian approach also in cases (like ours) when the event is not repeatable. Here “event” means something that can happen in different ways with different likelihoods, which we want to describe through a probability distribution. In our case, the event is “*the observable takes the value Σ* ”, and its probability distribution will be a function of Σ ranging over all possible values. The information on this event that we want to use is the perturbative expansion of the observable. How this will be used in practice depends on the model and will be discussed at length in the rest of this paper. Thanks to this information, we can then use probabilistic inference to improve the knowledge on the observable, namely to update the distribution of Σ .

The goal of this work is thus the construction of a probability distribution of the observable Σ given the perturbative expansion up to order n , namely

$$P(\Sigma|c_0, \dots, c_n, H), \quad (2.8)$$

where we have indicated with H any assumption (hypothesis), including both prior distributions and the model we want to use (we will come back to this point later). $P(A|B)$ indicates the probability distribution of A given the information B . In our case, the information is given by the first $n+1$ coefficients $c_0, \dots, c_n$, and H . This distribution contains all the information we desire about our knowledge of the observable. For instance, we can compute the best estimate of the observable Σ as its expectation value according to such distribution,

⁸ If this is not the case, it would mean that k_{asympt} is a very small number, which in turn implies that the expansion is badly divergent from the very first orders, without showing any apparently converging pattern. This is for instance the case for the perturbative conversion between the pole and the $\overline{\text{MS}}$ charm mass, where k_{asympt} is probably 2, and the number of known orders is 4, see e.g. [25–27]. In this condition, that may occur when the coupling is not small (e.g. in QCD at low energies), it’s clear that perturbation theory is no longer predictive.

$$\langle \Sigma \rangle = \int d\Sigma \Sigma P(\Sigma|c_0, \dots, c_n, H), \quad (2.9)$$

and its uncertainty either as the standard deviation or using degree-of-belief (DoB) intervals. Most importantly, the probability distribution can be used directly in physical analyses, when comparing theory predictions with data.

In the limit of “infinite information”, namely when we know the exact result Σ_{true} , the probability Eq. (2.8) should become

$$P(\Sigma|\Sigma_{\text{true}}) = \delta(\Sigma - \Sigma_{\text{true}}), \quad (2.10)$$

which represents the certainty (not probability) that $\Sigma = \Sigma_{\text{true}}$. In this limit any a priori assumption H does not matter. Eq. (2.10) cannot be seen as the all-order limit of Eq. (2.8), due to the divergent nature of the series and to the non-perturbative contributions discussed in Sect. 2.1. However, it suggests that when adding information (i.e. when increasing the number n of known orders, up to k_{asymp}) the probability distribution should become narrower and more localised. We shall consider this behaviour as a property that a good model for theory uncertainty must satisfy.

Note that knowing the probability distribution for the missing higher order term $\Delta_{\text{MHO}}^{(n)}$ is practically the same as knowing the distribution for Σ . Indeed, in the limit Eq. (2.7) where we neglect the asymptotic and non-perturbative contributions, the distributions for $\Delta_{\text{MHO}}^{(n)}$ and Σ are the same up to a trivial shift given by the perturbative result (Eq. 2.5). In the following, we will always deal directly with the distribution of Σ (Eq. 2.8), but we will compute it by estimating the missing higher orders $\Delta_{\text{MHO}}^{(n)}$.

2.3 The role of unphysical scales

A general feature of renormalizable QFTs is the appearance of an unphysical scale μ as a consequence of the regularization procedure needed to deal with ultraviolet divergences. This is known as the renormalization scale. Physical observables do not depend on it, as this scale is an artefact of the scheme adopted to renormalize the theory. However, in practical computations using perturbation theory, a scale dependence is present in each order, in a way that it is compensated order by order. Any finite-order truncation of the perturbative series will thus have a residual scale dependence, which is formally of higher order. As discussed in the introduction, this observation is at the core of the canonical scale variation method to estimate theory uncertainties.

Because of renormalization scale dependence, perturbative expansions do not uniquely determine a series, but rather a family of series parametrized by the renormalization scale μ . We shall thus rewrite Eq. (2.5) as

$$\Sigma_{\text{true}} = \sum_{k=0}^n c_k(\mu) \alpha^k(\mu) + \Delta_{\text{MHO}}^{(n)}(\mu) + \Delta_{\text{asymp}}(\mu) + \Delta_{\text{non-pert}}(\mu), \quad (2.11)$$

where the left-hand side, the exact result, is scale independent:

$$\frac{d}{d\mu} \Sigma_{\text{true}} = 0. \quad (2.12)$$

Therefore, whenever we want to estimate the value of an observable using perturbation theory, we need to face with the fact that the perturbative result is scale dependent, and also the missing higher orders are.

An immediate consequence of this fact is that also the probability distribution Eq. (2.8) will unavoidably depend on such the choice of scale μ . We can express this by writing the probability Eq. (2.8) as

$$P(\Sigma|c_0(\mu), \dots, c_n(\mu), H), \quad (2.13)$$

where we have emphasised that each coefficient depends on the scale μ ,⁹ or equivalently as

$$P(\Sigma|c_0, \dots, c_n, \mu, H), \quad (2.14)$$

where the coefficients c_k are intended as functions, to be computed at the value μ passed as an extra parameter. This dependence on the scale is clearly undesired, because this probability distribution, which depends on μ , is for the true observable, which is independent of μ . In the limit of infinite knowledge (Eq. 2.10), the distribution should tend to $\delta(\Sigma - \Sigma_{\text{true}})$ irrespectively of the value of μ . This implies that increasing the order, the probability distributions at different values of μ should become more and more similar. While this feature is certainly nice, having an infinite number of different results for the same object is obviously not ideal. Rather, one would like to obtain a probability distribution for Σ that does not depend on the choice of scale. This can be achieved in two ways: either having a criterion for selecting an “optimal value” of the scale, or combining in some way the results at different scales.

The first way is obviously simpler, provided such a criterion exists. In the literature there are various approaches that aim at selecting an optimal scale, e.g. the Brodsky-Lepage-Mackenzie (BLM) method [29–31], the principle of minimal sensitivity (PMS) [32,33], the principle of maximal conformality (PMC) [34–37] and the recent principle of observable effective matching (POEM) [38]. The PMC is probably the most widespread approach. It provides a way to select, order

⁹ Note that there is a μ dependence also in α , so to be rigorous we should also include $\alpha(\mu)$ in the list of known parameters.

by order, an optimal scale that removes non-conformal β -function contributions from the perturbative expansion. This is believed to remove the renormalons from the perturbative expansion, thereby leading to a possibly convergent series (or at least to a less divergent one). For our purposes, this approach could provide a way to select, among the infinitely many probability distributions for Σ , a specific one.¹⁰ Note, however, that the PMC fixes the scale at each known order except the last one, which is free and thus arbitrary. This implies that a residual scale dependence is present also in the PMC approach, even though it is claimed to be much milder than the canonical scale dependence. However, it has been pointed out that a proper study of all the ambiguities in the approach leads to larger uncertainties, comparable to the canonical scale uncertainty [40–43]. We conclude that while the PMC approach is certainly interesting, it cannot provide the full solution to our problem.

The second way to obtain a scale-independent probability distribution for Σ is what we pursue in this work. The treatment of the renormalization scale is addressed within the methodology for computing the probability distribution Eq. (2.8). The actual procedure to combine the probability distributions at different scales, which represents one of the most innovative proposals of this work, will be presented in Sect. 6.

Before moving further, another aspect of scale dependence must be discussed. So far, we have described scale dependence as an *obstacle* to obtain a unique probability distribution for the observable. In fact, scale dependence can be also considered as a *tool*. This relies on the fact, already discussed in the introduction, that the μ -dependence of the finite-order truncation of the perturbative series is of higher order, namely

$$\mu \frac{d}{d\mu} \sum_{k=0}^n c_k(\mu) \alpha^k(\mu) = \mathcal{O}(\alpha^{n+1}) = \mathcal{O}(\Delta_{\text{MHO}}^{(n)}). \quad (2.15)$$

This fact provides additional information on the expansion, which can be very useful as in most cases of interest in particle physics the available information is very limited (typically $n = 2$ or 3). In practice, let us assume for simplicity that the μ dependence at a given order k can be translated in a single number, r_k . It can for instance be the canonical scale uncertainty error, or the slope of the cross section as a function of μ , or something similar (we will provide a precise definition later in Sect. 3.2). We can then generalize the probability distribution as

$$P(\Sigma | c_0(\mu), r_0(\mu), \dots, c_n(\mu), r_n(\mu), H), \quad (2.16)$$

where we have made explicit that also the r_k numbers generally depend on the choice of μ about which the scale dependence is computed. In other words, also in this case we get a family of distributions depending on the value of μ at which the perturbative expansion is computed, however this time we also include in each member of the family some information on the scale dependence. Since these parameters double the previous information¹¹ they are clearly very precious.

Note that the parameters r_k represent the kind of information used in the construction of the canonical scale uncertainty. More precisely, the canonical scale uncertainty is based only on the last one, r_n (assuming r_n is defined according to Eq. (1.1)), and it does not provide a probabilistic interpretation. In Sect. 5 we will instead make use of the scale variation information r_k in a fully fledged probabilistic model, thereby providing a method that is, in some sense, a more reliable and statistically sound version of canonical scale variation.

We finally stress that the renormalization scale is not the only scale appearing in perturbative computations. For instance, in QCD processes involving hadrons in the initial or final states, the factorization of collinear singularities introduces a (perturbative) dependence on another unphysical scale, the so-called factorization scale. Also, in effective field theories widely used in collider phenomenology (e.g. heavy quark effective theory or soft-collinear effective theory), other unphysical scales may appear. The way to deal with these scales strictly depends on the scale itself. For instance, the dependence on the factorization scale cancels between the perturbatively computable coefficients and the non-perturbative parton distribution functions (PDFs). Therefore, if we wish to obtain factorization scale independent probability distribution for a physical observable, we may try to extend the procedure that we propose in Sect. 6 to this scale, with proper caveats due to the fact that PDFs are non-perturbative objects. Instead, if we wish to include our definition of theory uncertainties in the fits used to determine PDFs from data, the situation is completely different, as the PDFs are not physical observables and they are thus scheme and scale dependent. Addressing this issue is beyond the scope of this paper and it is left to future work. Here, we only focus on the renormalization scale dependence, which is universal.

¹⁰ In fact, in the default PMC approach the value of the scale is different for every order, which implies that the PMC perturbative expansion does not correspond to any of the series with a fixed value of μ for all orders. A simplified variant of the PMC approach where the scale is the same at each order has also been proposed [39].

¹¹ According to the definition of the coefficients c_k , the first one, c_0 , is independent of the scale, and therefore r_0 does not provide any information. However, there are cases in which the LO cross section is already scale dependent, for instance because it starts at order α . In these cases also the first coefficient $c_0 = 0$, and the first non-trivial order is c_1 , with r_1 being non-zero.

2.4 The Cacciari–Houdeau method

Before moving to our new proposals, we now present the Cacciari–Houdeau (CH) method for estimating theory uncertainties [3]. Let us forget about the scale dependence (which is not dealt with in the original paper) and consider the perturbative expansion

$$\Sigma_{\text{pert}} = \sum_{k=0}^{\infty} c_k \alpha^k. \quad (2.17)$$

We know that this series is divergent, but for the moment we ignore this fact. The basic assumption made in the CH model is that all the coefficients c_k are bounded in absolute value by a common number $\bar{c}$, namely

$$|c_k| \leq \bar{c} \quad \forall k. \quad (2.18)$$

The coefficient $\bar{c}$ is a parameter of the model, and specifically a hidden parameter, which will disappear (through marginalization) in the final results. Moreover, they assume that all c_k are independent from each other, with the exception for the common bound, which implies that

$$P(c_k, c_j | \bar{c}) = P(c_k | \bar{c}) P(c_j | \bar{c}) \quad \forall k, j, \quad k \neq j. \quad (2.19)$$

The conditional probability $P(c_k | \bar{c})$, that we shall call *the likelihood*, encodes in a probabilistic way the assumption Eq. (2.18). In the CH approach it is given by

$$P(c_k | \bar{c}) = \frac{1}{2\bar{c}} \theta(\bar{c} - |c_k|), \quad (2.20)$$

namely the condition Eq. (2.18) must be strictly satisfied (the probability that the condition is violated is zero), and within the allowed range all values are equally likely (flat distribution). Finally, they provide a prior distribution for the hidden parameter,

$$P_0(\bar{c}) \propto \frac{1}{\bar{c}} \theta(\bar{c}), \quad (2.21)$$

which corresponds to a flat distribution in the logarithm of $\bar{c}$, to encode the idea that the order of magnitude of $\bar{c}$ is a priori unknown. Note that this prior distribution is not normalizable, and thus it requires a regularization procedure to be used.

The ingredients above are sufficient to define the model, and using standard Bayesian inference they allow to compute the sought probability distribution for the observable. In practice, since the starting point is the perturbative expansion Eq. (2.17), to obtain a probability distribution for the full sum it is sufficient to have a probability distribution for the unknown c_k coefficients given the knowledge of the

first $n + 1$ coefficients $c_0, \dots, c_n$. The key ingredient is thus $P(c_k | c_0, \dots, c_n)$, with $k > n$, which can be computed as

$$\begin{aligned} P(c_k | c_0, \dots, c_n) &= \frac{P(c_k, c_0, \dots, c_n)}{P(c_0, \dots, c_n)} \quad (k > n) \\ &= \frac{\int d\bar{c} P(c_k, c_0, \dots, c_n, \bar{c})}{\int d\bar{c} P(c_0, \dots, c_n, \bar{c})} \\ &= \frac{\int d\bar{c} P(c_k, c_0, \dots, c_n | \bar{c}) P_0(\bar{c})}{\int d\bar{c} P(c_0, \dots, c_n | \bar{c}) P_0(\bar{c})} \\ &= \frac{\int d\bar{c} P(c_k | \bar{c}) P(c_0 | \bar{c}) \cdots P(c_n | \bar{c}) P_0(\bar{c})}{\int d\bar{c} P(c_0 | \bar{c}) \cdots P(c_n | \bar{c}) P_0(\bar{c})} \end{aligned} \quad (2.22)$$

where we have used the relation between the joint probability and the conditional probability $P(A, B) = P(A|B)P(B)$ in the first step, introduced the hidden parameter in the second step, used again the definition of the conditional probability in the third step and finally used Eq. (2.19). The last line is written in terms of known functions (the likelihood and the prior), and can thus be easily computed. This result can be easily generalized to the joint probability of more than one unknown coefficient,

$$\begin{aligned} P(c_{k_1}, \dots, c_{k_m} | c_0, \dots, c_n) \\ = \frac{\int d\bar{c} P(c_{k_1} | \bar{c}) \cdots P(c_{k_m} | \bar{c}) P(c_0 | \bar{c}) \cdots P(c_n | \bar{c}) P_0(\bar{c})}{\int d\bar{c} P(c_0 | \bar{c}) \cdots P(c_n | \bar{c}) P_0(\bar{c})}, \end{aligned} \quad (2.23)$$

$k_1, \dots, k_m > n.$

At this point one can also compute, at least formally, the probability distribution for the full sum, which is given by

$$\begin{aligned} P(\Sigma_{\text{pert}} | c_0, \dots, c_n) \\ = \int dc_{n+1} dc_{n+2} \cdots P(c_{n+1}, c_{n+2}, \dots | c_0, \dots, c_n) \\ \times \delta \left(\Sigma_{\text{pert}} - \sum_{k=0}^{\infty} c_k \alpha^k \right). \end{aligned} \quad (2.24)$$

Since this is an infinite-dimensional integration, it is impossible to perform it numerically and too hard to compute it analytically. Therefore, one can approximate the full sum with a truncated sum at some finite order $n + j$ to get

$$\begin{aligned} P(\Sigma_{\text{pert}} | c_0, \dots, c_n) &\simeq \int dc_{n+1} \cdots dc_{n+j} \\ &P(c_{n+1}, \dots, c_{n+j} | c_0, \dots, c_n) \\ &\delta \left(\Sigma_{\text{pert}} - \sum_{k=0}^{n+j} c_k \alpha^k \right), \end{aligned} \quad (2.25)$$

which can be easily handled, at least numerically. The easiest approximation is obtained with $j = 1$, where only the first missing higher order is used to approximate the distribution, and it leads to a simple analytical expression [3].

The CH approach is a breakthrough in the context of estimating the theory uncertainty from missing higher orders, as it provides for the first time a probabilistic way to determine the uncertainty of an observable computed in perturbation theory. Note that this approach only considers the behaviour of the expansion, without using any information from the scale dependence. This is exactly the opposite of the canonical scale variation method, which is based on the scale dependence and does not use any information on the behaviour of the expansion.

Despite the nice properties of the CH approach, there are some caveats that need to be considered. The most obvious one is the assumption Eq. (2.18), that implies that the perturbative expansion is bounded by a convergent (geometric) series,

$$|\Sigma_{\text{pert}}| = \left| \sum_{k=0}^{\infty} c_k \alpha^k \right| \leq \sum_{k=0}^{\infty} |c_k| \alpha^k \leq \sum_{k=0}^{\infty} \bar{c} \alpha^k = \frac{\bar{c}}{1-\alpha}, \quad (2.26)$$

where we have assumed $\alpha < 1$ (which is consistent with the perturbative hypothesis). This is in contrast with the known fact that perturbative expansions are divergent. In a subsequent paper, Ref. [4], the CH approach has been modified to account for the divergence of the series, by modifying the condition Eq. (2.18) into

$$|c_k| \leq \bar{b} k! \quad \forall k, \quad (2.27)$$

with $\bar{b}$ being the new hidden parameter, with the same prior as $\bar{c}$. This condition on the coefficients is much less stringent and compatible with the assumption that the divergence of the series is dominated by a factorial growth such as those due to renormalons. However, with this choice it is no longer possible to use Eq. (2.24) to compute the full sum, as it does not exist. Therefore only the approximation Eq. (2.25) can be considered, and typically with a low value of j otherwise the probability distribution becomes large due to the factorial growth. In Ref. [4] only $j = 1$ is considered.

The second issue is related to the fact that Eq. (2.18) does not account for a possible power growth of the coefficients. In other words, each term of the perturbative expansion is assumed to have a power scaling given just by α . This limitation was stressed already in the first CH paper [3], where they propose to solve it by rescaling α ,

$$\sum_{k=0}^{\infty} c_k \alpha^k = \sum_{k=0}^{\infty} c'_k \left(\frac{\alpha}{\eta} \right)^k \quad (2.28)$$

to obtain new coefficients c'_k that satisfy Eq. (2.18), or Eq. (2.27) in the factorial divergent hypothesis of Ref. [4]. The trouble is how to find such a rescaling factor η . In Ref.

[4], a global survey over a quite large number of observables is proposed to determine an optimal value of η . In this survey they compare the uncertainty computed at the next-to-last known order with the actual (known) next order, to quantify how reliable the uncertainty is for each given value of the rescaling factor. Apart from the details (for which we refer the Reader to Ref. [4]), we stress that this approach assumes that the rescaling factor is the same for all the observables.¹² However, this is hardly the case, as different processes and observables are characterized by different dominant perturbative corrections.¹³ Therefore, it is more appropriate to assume that the rescaling factor η is process and observable dependent. A different way to obtain it has been proposed in Ref. [9], where a fitting procedure is suggested to find an optimal value of η such that the known c_k are all of the same order. This method is observable dependent and uses only the information on the perturbative expansion to obtain the optimal rescaling. However, a fitting procedure to determine the rescaling factor clashes with the probabilistic nature of the rest of the procedure.

Finally, the CH approach or its modified versions do not deal with scale dependence. The CH machinery is applied to the perturbative expansion at a given value of the scale, and if one changes the scale the result changes accordingly. The difference in the final probability distribution at different values of the scales can be sizeable, see e.g. Ref. [5].

2.5 A working example: Higgs production at the LHC

In Sect. 7 we will consider various examples of perturbative expansions, and apply our methods to each of them. Nevertheless, in order to be clearer when discussing our new proposals, we think it is instructive to have a working example to immediately visualize how the various methods work.

The observable we choose is the inclusive cross-section for Higgs production in gluon fusion at the LHC for this purpose. This process has a number of advantages:

- it is known up to N³LO [44–47] (four orders in the perturbative expansion) in the so-called large top mass effective theory (this is a rarity in QCD processes at LHC, most of which are only known to NLO or NNLO);

¹² In Ref. [4] a distinction is made between two classes of observables, namely those involving hadrons in the initial state and those not involving hadrons in the initial state. For each class a value of the rescaling factor is found.

¹³ For instance, the dominant QCD corrections to gluon jets and quark jets are characterized by different color factors, $(C_A \alpha_s)^k$ and $(C_F \alpha_s)^k$ respectively. Also, in differential distributions in certain kinematic conditions there are powers of kinematic logarithms appearing with the powers of the coupling and dominating the perturbative corrections, and depending on the kinematic conditions the logarithm can have different sizes.

Fig. 4 The Higgs production cross section in gluon fusion at LHC $\sqrt{s} = 13$ TeV with $m_H = 125$ GeV for $\mu_F = m_H/2$ as a function of the renormalization scale μ at the four known perturbative orders

- it is characterized by large perturbative corrections;
- canonical scale variation underestimates the impact of (large) higher orders;
- its factorization scale dependence is very mild, so the whole scale dependence is basically fully captured by its renormalization scale dependence;
- because the process starts at $\mathcal{O}(\alpha_s^2)$, the LO is scale dependent;
- it is a real process and not a toy example.

On top of these reasons, the process is interesting also from a phenomenological point of view (see e.g. Ref. [48]), and indeed it was subject of several investigations of theory uncertainties from missing higher orders [4, 5, 8, 46].

Specifically, we consider LHC at $\sqrt{s} = 13$ TeV and set the Higgs mass to $m_H = 125$ GeV. We fix the factorization scale to $\mu_F = m_H/2$ (which is a standard choice [48]), even though changing this value has a negligible effect on the cross section, in particular at high orders. The “raw” result of the computation, which is the cross section as a function of the renormalization scale μ , is plotted in Fig. 4. Note that this process depends on a single hard scale, the Higgs mass m_H . Therefore, it is natural to choose μ of the order of m_H , in order to avoid the presence of large unresummed logarithms of μ/m_H in the perturbative coefficients. Nevertheless, we believe it is instructive to visualize the scale dependence for wide range of scales: the plot covers almost four decades in μ/m_H .

We see that for large values of the scale, where the QCD coupling is smaller, the expansion is characterized by all positive contributions and it progresses very slowly, with large perturbative corrections. Conversely, at small scales where the strong coupling blows up the expansion is highly unstable. In a “central” region, where $\mu \sim m_H$, the expansion behaves in a reasonably perturbative way, even though the perturbative corrections are rather large and it is not at all clear what could possibly be the true cross section. Note also

the presence of a stationary point at NNLO and of two stationary points at N³LO, which could corrupt an estimate of the uncertainty based on canonical scale variation.

We stress that the full plot of Fig. 4 can be constructed from just the sequence of partial sums of the observable at the various orders at a given scale, and the knowledge of the value of the coupling at that scale. In this example, we have at $\mu = m_H$ (and $\mu_F = m_H/2$)¹⁴

$$\alpha_s(m_H) = 0.1126, \Sigma_{\text{pert}}(m_H) = \{13.0, 30.7, 41.8, 46.3\} \text{ pb}, \quad (2.29)$$

where the values in curly brackets correspond to the cross section at LO, NLO, NNLO and N³LO respectively. The way to reconstruct the cross section at any scale from these ingredients is discussed in Sect. A.1.

3 Model-independent features of the inference approach at fixed scale

In this section we start introducing our notation and present some general features of the construction of the models. For the time being we consider only the models at a fixed scale. How to obtain scale-independent probability distributions will be discussed in Sect. 6.

3.1 Basic notations

Let us denote with $\Sigma_n(\mu)$ the partial sum of the perturbative series up to order n depending on the scale μ . If we are considering a standard perturbative expansion in powers of α this is given by

$$\Sigma_n(\mu) = \sum_{k=0}^n c_k(\mu) \alpha^{k+k_0}(\mu), \quad (3.1)$$

where we have explicitly introduced an offset k_0 for observables starting at $\mathcal{O}(\alpha^{k_0})$. This notational change with respect to the previous section allows us to be sure that the first order, $k = 0$, namely the leading order (LO), is non-zero. Note that the information contained in the coefficients c_k is fully contained in the sequence of partial sums

$$\Sigma_0, \Sigma_1, \Sigma_2, \dots \quad (3.2)$$

once the value of α and of k_0 are specified. From now on, we shall consider the partial sums as the basic objects, and forget about the coefficients c_k and even of α . In this way,

¹⁴ These numbers have been computed with the latest version of ggHiggs [5, 49–52]. Note that the N³LO result differs from the one reported in Ref. [5] because in the latest version of the code the N³LO computation through a threshold expansion [46] has been replaced with the complete result [47].

the “perturbative expansion” is more general, as it does no longer necessarily need to be a strict expansion in powers of α . For instance, it can be a logarithmic-ordered expansion of a resummed result, or a (non-)linear transformation of the perturbative expansion. In what follows we simply assume that Σ_n represents the partial sum of a non-specified expansion that behaves perturbatively, defined such that the “LO” Σ_0 is non-zero.

For a number of reasons that will become clear later, it is convenient to introduce a normalized version of the expansion, where the LO is factored out,

$$\Sigma_n(\mu) = \Sigma_0(\mu) \sum_{k=0}^n \delta_k(\mu), \quad (3.3)$$

where we have defined the dimensionless coefficients¹⁵

$$\delta_k(\mu) \equiv \frac{\Sigma_k(\mu) - \Sigma_{k-1}(\mu)}{\Sigma_0(\mu)}. \quad (3.4)$$

According to this definition, $\delta_0 = 1$ always, so we can also write

$$\Sigma_n(\mu) = \Sigma_0(\mu) \left(1 + \sum_{k=1}^n \delta_k(\mu) \right). \quad (3.5)$$

The coefficients δ_k contain the information on the perturbative orders. The fact that $\delta_0 = 1$ sets a common size for all perturbative expansions (useful when defining the model), and also tells us that the LO does not contain any useful information on the behaviour of the expansion and thus on the uncertainty due to missing higher orders (which is obvious, as with just the LO one cannot know how large the perturbative corrections will be). This is true even when the LO is scale dependent, because without another order to compare with one cannot know how to reliably translate the scale dependence into an information on the size of missing higher orders. Only when at least two orders, namely δ_0 and δ_1 , are known one can start making inference. The only role of the LO is to set the dimension and the size of the observable through the prefactor $\Sigma_0(\mu)$.

3.2 Definition of the scale-dependence numbers

Because the expansion is scale dependent, as discussed in Sect. 2.3 one can construct numbers $r_k(\mu)$ that encode this dependence at each order. The importance of these numbers stems from the observation that scale dependence of a

physical observable is formally of higher order, Eq. (2.15), namely¹⁶

$$r_k(\mu) = \mathcal{O}(\alpha^{k+1}). \quad (3.6)$$

The actual construction of these numbers must be such that this equation is satisfied. The simplest measure of the scale dependence is the slope of $\Sigma_k(\mu)$ as a function of μ , or more precisely its derivative with respect to $\log \mu$ (the logarithm is natural because the dependence on the scale is logarithmic). In order to directly compare these numbers with the δ_k coefficients, it is convenient to normalise the derivative to the observable itself to make them dimensionless. The two most natural ways to do this are

$$\frac{1}{\Sigma_0(\mu)} \mu \frac{d}{d\mu} \Sigma_k(\mu) \quad \text{or} \quad \frac{1}{\Sigma_k(\mu)} \mu \frac{d}{d\mu} \Sigma_k(\mu), \quad (3.7)$$

where in the first case we have normalized to the LO Σ_0 , while in the second case we have normalized to the observable at the same order at which the scale dependence is computed. If the observable does not change much at different orders, the two options are equivalent. However, in presence of large perturbative corrections there can be a substantial difference between the two. None of them is better in an absolute sense. However, we argue that the second option has a nicer perturbative behaviour, with higher order r_k being typically smaller than lower order ones, as one would expect from Eq. (3.6).

To prove this, we consider the example of Higgs production introduced in Sect. 2.5, and show in Fig. 5 the (absolute value of) the two options in Eq. (3.7). From the left plot, corresponding to the left option in Eq. (3.7), we see that at NLO (green curve) the slope is always larger than the LO one (black curve), and at the next orders there is a large variability without a precise hierarchy. Conversely, in the right plot, corresponding to the right option in Eq. (3.7), the higher order curves are smaller than the lower order ones over a wide range of scales, which is the expected perturbative behaviour. So the definition of the r_k coefficients that we propose is *morally* given by

$$r_k(\mu) \simeq \left| \frac{1}{\Sigma_k(\mu)} \mu \frac{d}{d\mu} \Sigma_k(\mu) \right| = \left| \mu \frac{d}{d\mu} \log \Sigma_k(\mu) \right|, \quad (3.8)$$

where the absolute value is introduced to keep only the information on the size of the dependence but not the sign. Note that at small scales, where the perturbative expansion is unstable (Fig. 4), the expected hierarchy is violated. This is the case also in proximity of the stationary points. To avoid problems

¹⁵ A possible alternative definition of the dimensionless perturbative coefficients $\delta_k(\mu)$ is obtained dividing by $\Sigma_{k-1}(\mu)$ rather than $\Sigma_0(\mu)$. Given the similarity to the definition used here (Eq. 3.4), the models that we will propose can be straightforwardly modified using this alternative definition. The results will in general differ, the difference being larger if perturbative corrections are large, thus highlighting a source of arbitrariness in the procedure.

¹⁶ We do not include k_0 in the power of α because we assume to define r_k such that they are directly comparable to the normalized δ_k .

Fig. 5 The absolute value of the normalized slopes defined in Eq. (3.7), respectively shown in the left and right plots, for the Higgs production process

arising from stationary points, we propose to define the r_k numbers in a slightly different way. Namely, we replace the derivative with a finite difference, and look for the largest finite difference in a range around the point μ . In formulae, we have¹⁷

$$r_k(\mu) \equiv \frac{1}{|\Sigma_k(\mu)|} \max_{\mu/f \leq v \leq f\mu} \left| \frac{\Sigma_k(v) - \Sigma_k(\mu)}{\log(v/\mu)} \right|, \quad (3.9)$$

that selects the largest finite difference in the range of scales between μ/f and $f\mu$, where $f > 1$ is a factor to be fixed. This definition is more robust than the one based on the derivative.

We show this in Fig. 6, for $f = 2$ (left plot) and $f = 4$ (right plot). Increasing f allows to be less sensitive to stationary points, and indeed in the right plot the expected hierarchy is preserved over a wide range of scales. At low scales, the value of r_k at NNLO and N³LO blows up because the finite difference probes the low-scale region where the perturbative result is unstable (Fig. 4). For larger values of f , the blow-up region obviously moves to larger scales. Therefore, the robustness gained increasing f has to be balanced with the stability loss. We believe that $f = 4$ can be considered as a reasonable compromise.¹⁸

We stress that for an efficient computation of the $r_k(\mu)$ coefficients a fast evaluation of the scale dependence of the observable is needed. Since the renormalization scale dependence is universal and governed by the β -function of the the-

ory, it can be constructed automatically from the knowledge of the observables at the various orders at a single value of the scale. The details are reported in Appendix A.1.

We finally note that if the LO is scale independent, then $r_0 = 0$. Given that we aim at using the r_k numbers to estimate the higher orders, a zero value would be inappropriate. More precisely, $r_0 = 0$ means that the LO cannot be used to probe higher orders through scale variation. We can fix this by assuming that, when the LO is scale independent, we can only say that the NLO corrections may be of the order of the LO itself. Given the definition of the r_k as normalized quantities, this corresponds to assuming $r_0 = \mathcal{O}(1)$. For definiteness, we arbitrarily set $r_0(\mu) = 1/2$ in these cases. We shall see later that this assumption is rather conservative.

3.3 General features of the models

Before discussing the actual models that we propose, we want to give some general features that are common to all of them. Let us recall that the goal of this work is to construct a probability distribution for Σ given the first known $n + 1$ orders $\Sigma_0, \Sigma_1, \dots, \Sigma_n$, or equivalently $\Sigma_0, \delta_1, \dots, \delta_n$ (using our new notation). This was given in Eq. (2.8), or, when working at fixed scale, in Eq. (2.14).

According to our assumption Eq. (2.7) that the uncertainty of a theoretical prediction based on perturbation theory is dominated by the missing higher orders, we shall compute the distribution for Σ through these missing higher orders, ignoring the other contributions from the asymptotic expansion truncation and the non-perturbative part. In other words, we shall approximate the observable Σ as

$$\Sigma \simeq \Sigma_k(\mu), \quad k \leq k_{\text{asympt}}, \quad (3.10)$$

where the best approximation is obtained using $k = k_{\text{asympt}}$, namely including all the missing higher orders up to the point in which the asymptotic expansion starts to grow. Since we

¹⁷ In the numerical implementation, the max over a continuous variable is replaced with a discretised version.

¹⁸ Note that the canonical scale variation method uses the scale dependence over a factor of 2 about the central scale choice. Our definition of r_k with $f = 4$ probes a larger range. However, while in canonical scale variation increasing the factor increases the uncertainty, here since the finite difference is divided by the size of the variation the dependence on the factor is much milder. In a hypothetical linear dependence case, r_k would be independent of f while the canonical scale variation uncertainty would change with $\log f$.

Fig. 6 The coefficients $r_k(\mu)$ (Eq. 3.9), for $f = 2$ (left) and $f = 4$ (right), for the Higgs production process

typically do not know a priori the value of k_{asympt} , the best we can do is to include a few extra orders beyond the known ones, without exaggerating in order not to risk to go beyond k_{asympt} .

The inference on the observable Σ can be obtained in terms of the more fundamental inference of the higher orders from the known ones. This was already done in Sect. 2.4, and we repeat that derivation here in more generality and with our new notation. Assuming we know the coefficients of the expansion up to order n and we approximate the observable with its expansion up to order $n + j$ with $j > 0$, $n + j \leq k_{\text{asympt}}$, the probability distribution is given by (suppressing for ease of notation the implicit dependence on the hypothesis H)

$$\begin{aligned} P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) &= \frac{P(\Sigma, \delta_n, \dots, \delta_1, \Sigma_0|\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)} = \frac{\int d\delta_{n+j} \dots d\delta_{n+1} P(\Sigma, \delta_{n+j}, \dots, \delta_1, \Sigma_0|\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)} \\ &= \frac{\int d\delta_{n+j} \dots d\delta_{n+1} P(\Sigma|\delta_{n+j}, \dots, \delta_1, \Sigma_0, \mu) P(\delta_{n+j}, \dots, \delta_1, \Sigma_0|\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)} \\ &= \int d\delta_{n+j} \dots d\delta_{n+1} P(\Sigma|\delta_{n+j}, \dots, \delta_1, \Sigma_0, \mu) P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu) \\ &\simeq \int d\delta_{n+j} \dots d\delta_{n+1} \delta(\Sigma - \Sigma_{n+j}(\mu)) P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu), \end{aligned} \quad (3.11)$$

where we have used in the last step the $(n + j)$ -th order approximation

$$P(\Sigma|\delta_{n+j}, \dots, \delta_1, \Sigma_0, \mu) \simeq \delta(\Sigma - \Sigma_{n+j}(\mu)), \quad (3.12)$$

which is a direct consequence of Eq. (3.10), and the definition of $\Sigma_n(\mu)$, Eq. (3.5).¹⁹ If Σ is a positive definite observ-

¹⁹ Note that Eq. (3.12) is a rather crude approximation. To be conservative, one could replace the delta function with a smooth distribution, for instance a gaussian. However, in this case one would need some parameters, like the width of the distribution, which would then be arbitrary.

able, like a cross section, one could also impose a positivity constraint through a factor $\theta(\Sigma)$, which however requires the computation of a normalization factor as in general the distribution is no longer normalized. The result Eq. (3.11) is written in terms of $P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu)$, which represents the probability of the higher orders given the known ones, and depends on the model under consideration. Note that we have also included Σ_0 among the known information, even though it is just a prefactor: indeed Σ_0 is required to compute the r_k numbers introduced in Sect. 3.2 that are needed in models that use information on the scale dependence.

The delta function in Eq. (3.11) can be used to perform the integral over one of the higher orders, say δ_{n+j} , which gives

$$\begin{aligned} P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) &\simeq \frac{1}{\Sigma_0(\mu)} \int d\delta_{n+j-1} \dots d\delta_{n+1} \\ &\times P\left(\delta_{n+j} = \frac{\Sigma - \Sigma_{n+j-1}(\mu)}{\Sigma_0(\mu)}, \delta_{n+j-1}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu\right). \end{aligned} \quad (3.13)$$

The other integrations are more complicated, and should be performed numerically (unless the model is particularly sim-

ple). The simplest result is obtained when considering $j = 1$, namely when approximating the observable using just the first unknown order,

$$P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) \stackrel{j=1}{\simeq} \frac{1}{\Sigma_0(\mu)} P\left(\delta_{n+1} = \frac{\Sigma - \Sigma_n(\mu)}{\Sigma_0(\mu)} | \delta_n, \dots, \delta_1, \Sigma_0, \mu\right). \quad (3.14)$$

In practice, since the order $n + j$ that approximates best the observable is not known a priori, a convenient approach to choose properly j is the following. We start considering the simplest approximation, $j = 1$ (Eq. 3.14). Then, we include the next order, namely we use $j = 2$. If the distribution changes visibly, then we further increase j by a unity, and so on until the distribution changes only mildly, up to a tolerance decided by the user.

We now turn our attention to the probability $P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu)$. As we said, this is model dependent. However, we can further write it in terms of more fundamental probabilities, using the relation

$$P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu) = \frac{P(\delta_{n+j}, \dots, \delta_{n+1}, \delta_n, \dots, \delta_1, \Sigma_0|\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)}. \quad (3.15)$$

The numerator and the denominator are the same object, simply with a different number of δ_k terms. In the numerator only some of them are known while the others are unknown, but mathematically this does not make any difference. The joint distribution $P(\delta_m, \dots, \delta_1, \Sigma_0|\mu)$ at fixed scale is the basic object of the model that is needed to make inference on the higher orders. In all the models we will consider, we will assume that there is a number of hidden parameters characterizing the model. Denoting with $\vec{p}$ the vector of such parameters, the joint distribution can be written as

$$P(\delta_m, \dots, \delta_1, \Sigma_0|\mu) = \int d\vec{p} P(\delta_m, \dots, \delta_1, \Sigma_0, \vec{p}|\mu) = \int d\vec{p} P(\delta_m, \dots, \delta_1, \Sigma_0|\vec{p}, \mu) P_0(\vec{p}|\mu), \quad (3.16)$$

where in the second line we have written explicitly the prior of these parameters, as they are hidden and thus can never be known exactly.

To compute the joint distribution, it is sometimes useful to visualize the relation between the different objects using a Bayesian network. The most general network for any model of theory uncertainties is rather simple, and it is depicted in Fig. 7 (left) showing only the first four orders (the generalization to more than four orders is obvious). The various orders depend in general on all previous orders (but not on future ones, otherwise the model cannot be predictive), and

on the hidden parameters $\vec{p}$. Since Σ_0 is the first order and it only sets the size of the observable, it does not depend on anything. In addition to this general structure, we want to consider more explicitly the role of the scale dependence numbers r_k . Since these are functions of the δ_k coefficients and Σ_0 , there is no need to specify them in the network. However, it may be useful to introduce them explicitly in order to better appreciate their role. Therefore, in the same Fig. 7 (right) we also show explicitly a network depending on these r_k . The dashed arrows represent *deterministic links*, namely analytic relations rather than probabilistic ones, and mean that the r_k numbers are computable analytically from the various orders. Note that this network is not as general as the previous one. Indeed now we have made the assumption that each δ_k depend only on the previous δ_{k-1} and r_{k-1} (and on $\vec{p}$), but not on all the previous orders. This simplification is not necessary, but it will be adopted in all our models (except the one of Sect. B.2).

4 Model 1: geometric behaviour model

We now present our first model, that uses only information on the behaviour of the expansion. As such, this model is very general, and not restricted to a QFT application.

4.1 The hypothesis of the model

The first model that we consider is a generalization of the Cacciari–Houdeau model introduced in Sect. 2.4. The main difference is that our model accounts for a possible power growth of the coefficients of the expansion within the probabilistic approach. In the CH model each term of the expansion is bounded by

$$|c_k \alpha^k| \leq \bar{c} \alpha^k \quad \forall k, \quad (4.1)$$

which is Eq. (2.18) in which we have emphasised that the power behaviour of the full $\mathcal{O}(\alpha^k)$ term is entirely described by α^k . As we have discussed in Sect. 2.4, this hypothesis is hardly satisfied, and a variant of the CH method in which the expansion parameter α is rescaled by a factor η is advisable. So far, this has never been done within the context of the probabilistic model.

Now, in our new notation (Eq. 3.5), we have lost information on the coupling α , as the whole information is contained in the δ_k coefficients Eq. (3.4). This was done on purpose and is to be considered as an advantage, as the expansion Eq. (3.5) is more general than a strict expansion in powers of α . If we want to translate the CH condition Eq. (4.1) in the new language, we are forced to introduce back an expansion parameter. Since from our point of view this is a new

Fig. 7 The most general Bayesian network of models of inference of missing higher orders (left), and a more specific one, using explicitly the scale variation numbers r_k , that covers all the cases considered in this work (except the one of Sect. B.2)

parameter (as we have lost information about α), it is natural to consider it as a parameter of the model, rather than an external one. As such, it is not fixed to be α or a fraction of it. The condition that we consider is thus

$$|\delta_k(\mu)| \leq ca^k \quad \forall k < k_{\text{asympt}}, \quad (4.2)$$

where both c and a are positive hidden parameters of the model. This condition implies that the expansion is bounded by a geometric expansion, and we thus call this model a *geometric behaviour model*. We have specified that this bound can only be valid for orders k smaller than k_{asympt} , otherwise it is certainly violated. This is anyway the only region which we are interested in, according to the discussion in Sect. 2.1.

Equation (4.2), though very similar to the CH condition Eq. (4.1), differs from it in a number of very important aspects, that we now list.

- The fact that a is a parameter makes not only the model more general than CH, but it also allows to find, through inference, the most appropriate values (in a probabilistic sense) of the expansion parameter a compatible with the behaviour of the expansion. In other words, a can be interpreted as the rescaled expansion parameter α/η introduced in Sect. 2.4, but with the rescaling factor η being determined through inference from the perturbative expansion itself, as opposed to the approaches of Refs. [4, 9].
- The parameter c is dimensionless, as opposed to $\bar{c}$ which has the dimension of the observable. This is useful as we can legitimately use a universal prior for c without knowing anything about the observable.²⁰
- Since the condition Eq. (4.2) is limited to $k < k_{\text{asympt}}$, the fact that the perturbative expansion is typically factorially divergent does not imply that the geometric bound is unacceptable. Of course one cannot say that the entire series is bounded by a geometric series, but a small portion of it may well be. Therefore, the condition Eq. (4.2)

limited to the first few orders can be considered as perfectly acceptable.²¹

The CH assumption Eq. (2.18) can be recovered from this new approach by fixing $a = \alpha$ (or $a = \alpha/\eta$ in the rescaled variant) and rewriting $c = \bar{c}/\Sigma_0(\mu)$. Note that the hidden parameters c, a depend on the scale μ . This has not been written explicitly, because in a statistical language this information is expressed by saying that c and a are correlated with μ .

In order to construct a probabilistic model to estimate theory uncertainties, we need to translate the condition Eq. (4.2) into a likelihood function. We assume the simple conditional probability

$$P(\delta_k|c, a, \mu) = \frac{1}{2ca^k} \theta\left(ca^k - |\delta_k(\mu)|\right), \quad k < k_{\text{asympt}}, \quad (4.3)$$

which is the straightforward extension of the CH choice, Eq. (2.20). We have considered the idea of allowing violation of the bound, by adding tails to the likelihood. However, the fact of having two hidden parameters already makes the model much more flexible than the original CH model, so adding a violation of the bound would not lead to any substantial improvement to the model stability.

In similarity with the CH model, we assume that all δ_k are independent of each other at fixed c, a and μ , namely

$$P(\delta_k, \delta_j|c, a, \mu) = P(\delta_k|c, a, \mu)P(\delta_j|c, a, \mu) \quad \forall k, j, \quad k \neq j, \quad (4.4)$$

which generalizes to any set of δ_k coefficients. These conditions, together with the prior distributions for the hidden parameters that we are going to discuss, are sufficient to fully define the model. The Bayesian network of this model is a simplified version of the general one introduced in Sect. 3.3, and is depicted in Fig. 8.

²⁰ In the CH approach, the prior for $\bar{c}$ is universal as well. However, in that case, one could have decided to use an observable-dependent prior.

²¹ Whether or not it is also good has to be judged on the basis of the performance of the model. We will discuss this later in Sect. 4.4.

Fig. 8 Bayesian network for the geometric behaviour model. Each δ_k depends only on the hidden parameters (and μ) through the likelihood Eq. (4.3)

4.2 The choice of priors

The two hidden parameters c, a need a prior distribution. We assume that a priori the two variables are uncorrelated

$$P_0(c, a|\mu) = P_0(c|\mu)P_0(a|\mu), \quad (4.5)$$

where we have emphasised that in principle the prior can depend on the scale μ , considered to be externally given for the moment (this will change in Sect. 6).

Before proposing a functional form for the prior, let us comment on the first step of the inference. The information from the LO is encoded in Σ_0 , that appears as a prefactor in our normalized expansion (Eq. 3.5). Since the likelihood Eq. (4.3) looks at the δ_k and not at Σ_0 , the knowledge of the LO does not change our prior: $P(c, a|\Sigma_0, \mu) = P_0(c, a|\mu)$. While this is strictly speaking correct according to our notation, it misses a conceptual point. Indeed, it exists also a δ_k for the LO, which is δ_0 . The fact that $\delta_0 = 1$ makes it a trivial variable, in the sense that it carries no information, which is the reason why it does not appear in our functions. However, from a mathematical viewpoint, it would play a role if we assume that the likelihood Eq. (4.3) is also valid for the LO. Indeed, at order zero, the likelihood becomes

$$P(\delta_0|c, a, \mu) = \frac{1}{2c}\theta(c - |\delta_0|) = \frac{1}{2c}\theta(c - 1). \quad (4.6)$$

This equation implies that the distribution for a is unmodified by the knowledge of the LO, while the distribution for c changes as

$$\begin{aligned} P(c|\delta_0, \mu) &\propto \int da P(\delta_0|c, a, \mu)P_0(c, a|\mu) \\ &\propto \frac{P_0(c|\mu)}{c}\theta(c - 1). \end{aligned} \quad (4.7)$$

This result shows that the requirement that the likelihood Eq. (4.3) applies also at LO implies the constraint $c \geq 1$. Since δ_0 is not explicitly part of our parameters, we will not perform the inference in Eq. (4.7) in our model. In other words, our prior Eq. (4.5) is to be considered as the posterior after the (trivial and universal) knowledge of $\delta_0 = 1$. We will keep track of this result by constructing the prior such that the condition $c \geq 1$ is satisfied.

At this point we are free to choose the functional form of our prior. Note that it is convenient to use simple functional forms, such that analytic computations can be performed. Let us start with the prior for c . Since we do not have any a priori knowledge on the expected size of c , only a monotonic prior is acceptable. We find it reasonable to assume a power law function

$$P_0(c|\mu) = \frac{\epsilon}{c^{1+\epsilon}}\theta(c - 1), \quad \epsilon > 0, \quad (4.8)$$

where we have included the $\theta(c - 1)$ for the reason explained above. Note that we do not include any dependence on the scale, namely for any value of μ we use the same prior distribution. The parameter ϵ is an arbitrary parameter, and can be chosen at will (we will discuss our favourite choices later in Sect. 4.4). We note that for $\epsilon = 1$ we obtain the form $\theta(c - 1)/c^2$, which results from a “pre-prior” (without knowledge of δ_0) proportional to $\theta(c)/c$, as obvious from Eq. (4.7). This is the prior used for $\bar{c}$ in the CH model, Eq. (2.21). Since $\theta(c)/c$ is an improper distribution, a regularization procedure is needed in CH to perform practical computation. Rather, when including the trivial information $\delta_0 = 1$ within the model, the prior is a proper distribution, and the computations are simplified. This is one advantage of using from the start the universal information $\delta_0 = 1$, which is in turn a consequence of using the normalized version of the expansion (Eq. 3.5).

We stress that the choice $\theta(c)/c$ for the “pre-prior” corresponds to a flat (and thus non-informative) distribution for the variable $\log c$, justified in the CH work by the argument that the order of magnitude of the hidden parameter is unknown. The other natural non-informative “pre-prior” is given simply by $\theta(c)$, namely a flat distribution in the hidden parameter c , which is again improper. This choice corresponds to the value $\epsilon = 0$ in Eq. (4.8). With $\epsilon = 0$ also the prior Eq. (4.8) is improper, and indeed in that equation we have assumed ϵ to be strictly greater than zero. In fact, computations can be easily performed also in the $\epsilon \rightarrow 0$ limit with just a little care. We find however that this complication is not necessary: if we wish to mimic the effect of a flat “pre-prior” in c , we can just use a very small positive value for ϵ . This will indeed be our favourite choice, see Sect. 4.4. Variations of this parameter will be explored in Sect. 4.6.

As far as the parameter a is concerned, we have to make an initial choice about the expected behaviour of the expansion. Indeed, the geometric bound is convergent (namely, at finite order, decreasing with the order) only for $a < 1$. In principle, we could allow $a \geq 1$, which would describe a (power) divergent behaviour of the expansion.²² However,

²² The “divergent region” $a \geq 1$ would predict large higher orders, which in turn would enlarge the probability distribution for the observable. The impact of this region can be limited by a strongly suppressed prior for $a \geq 1$.

allowing $a \geq 1$ is in contrast with the asymptotic nature of the expansion that we are assuming, see Sect. 2.1. Therefore, we suggest to limit our interest to the region $a < 1$. We thus propose the functional form

$$P_0(a|\mu) = (1 + \omega)(1 - a)^\omega \theta(a) \theta(1 - a), \quad \omega \geq 0. \quad (4.9)$$

Once again, we assume that this prior is independent of the scale μ . For $\omega = 0$, we obtain a flat distribution in the allowed region $0 \leq a \leq 1$ (in this case, the extreme value $a = 1$ is included), while for $\omega > 0$ we suppress the region $a \rightarrow 1$ to favour small values of a . The actual value of ω that we recommend will be discussed in Sect. 4.4, and its variations will be considered in Sect. 4.6.

4.3 Inference on the unknown higher orders

We have now all the ingredients to perform the inference in this model. The basic probability that we need is the conditional probability of unknown higher orders given the first n known non-trivial orders $\delta_1, \dots, \delta_n$. Note that because of our assumption Eq. (4.2), only a limited number of higher orders, up to order k_{asympt} , can be predicted within this model. Therefore, the most generic probability distribution we need to consider, according to Eq. (3.15), is

$$P(\delta_{n+j}, \dots, \delta_{n+1} | \delta_n, \dots, \delta_1, \mu) = \frac{P(\delta_{n+j}, \dots, \delta_{n+1}, \delta_n, \dots, \delta_1 | \mu)}{P(\delta_n, \dots, \delta_1 | \mu)}, \quad j \geq 1, \quad n + j < k_{\text{asympt}}, \quad (4.10)$$

namely the probability of the unknown orders $n+1, \dots, n+j$ given the known orders $1, \dots, n$. In contrast with Eq. (3.15), we have removed here the explicit dependence on Σ_0 , as it does not appear in the likelihood and thus it does not play any role in the inference procedure.

The numerator and the denominator are the same object, therefore we can focus on the distribution (at fixed scale) for a generic number m of consecutive coefficients, which is given by

$$\begin{aligned} P(\delta_m, \dots, \delta_1 | \mu) &= \int dc \int da P(\delta_m, \dots, \delta_1 | c, a, \mu) P_0(c, a | \mu) \\ &= \int dc \int da P(\delta_m | c, a, \mu) \cdots P(\delta_1 | c, a, \mu) P_0(c, a | \mu) \\ &= \frac{1}{2^m} \int \frac{da}{a^{\frac{m(m+1)}{2}}} P_0(a | \mu) \\ &\quad \times \int \frac{dc}{c^m} P_0(c | \mu) \theta \left(c - \max \left[\frac{|\delta_1(\mu)|}{a}, \dots, \frac{|\delta_m(\mu)|}{a^m} \right] \right) \end{aligned} \quad (4.11)$$

which corresponds to Eq. (3.16) specialized to our case. In the first line we have introduced the hidden parameters, in the second line we have used the independence of the coefficients, Eq. (4.4), and in the third line we have explicitly

written the likelihood Eq. (4.3) and used the prior independence of the hidden parameters Eq. (4.5). The integral in Eq. (4.11) can be computed analytically for our choice of priors Eq. (4.8) and Eq. (4.9). The inner integral is given by

$$\begin{aligned} \int \frac{dc}{c^m} P_0(c | \mu) \theta \left(c - \max \left[\frac{|\delta_1(\mu)|}{a}, \dots, \frac{|\delta_m(\mu)|}{a^m} \right] \right) \\ = \frac{\epsilon}{m + \epsilon} \max \left[1, \frac{|\delta_1(\mu)|}{a}, \dots, \frac{|\delta_m(\mu)|}{a^m} \right]^{-m - \epsilon}. \end{aligned} \quad (4.12)$$

Depending on the value of a , the max function selects a different term with a different a dependence. In order to compute the a integral analytically, it is therefore convenient to partition the integration region $0 \leq a < \infty$ into a finite number of intervals, in each of which the max function returns one of its arguments. Since the arguments of the max function contain powers of a that grow with k , the intervals are ordered with k . More precisely, smaller values of a will select larger powers k , and viceversa. We can thus introduce consecutive decreasing numbers a_k , representing the boundaries of these consecutive intervals, defined such that

$$a_{k+1} < a < a_k \Leftrightarrow \max \left[1, \frac{|\delta_1(\mu)|}{a}, \dots, \frac{|\delta_m(\mu)|}{a^m} \right] = \frac{|\delta_k(\mu)|}{a^k} \quad (4.13)$$

and assuming $a_0 \equiv \infty$ and $a_{m+1} \equiv 0$. An algorithm for extracting the various a_k 's from the knowledge of the δ_k 's is described in Appendix A.2. The a integral is then given by

$$\begin{aligned} P(\delta_m, \dots, \delta_1 | \mu) &= \frac{1}{2^m} \int \frac{da}{a^{\frac{m(m+1)}{2}}} P_0(a | \mu) \frac{\epsilon}{m + \epsilon} \\ &\quad \max \left[1, \frac{|\delta_1(\mu)|}{a}, \dots, \frac{|\delta_m(\mu)|}{a^m} \right]^{-m - \epsilon} \\ &= \frac{1}{2^m} \sum_{k=0}^m \int_{a_{k+1}}^{a_k} \frac{da}{a^{\frac{m(m+1)}{2}}} P_0(a | \mu) \frac{\epsilon}{m + \epsilon} \left(\frac{|\delta_k(\mu)|}{a^k} \right)^{-m - \epsilon} \\ &= \frac{\epsilon(1 + \omega)}{2^m(m + \epsilon)} \sum_{k=0}^m \frac{1}{|\delta_k(\mu)|^{m + \epsilon}} \int_{\min(1, a_{k+1})}^{\min(1, a_k)} da a^{(m + \epsilon)k - \frac{m(m+1)}{2}} (1 - a)^\omega, \end{aligned} \quad (4.14)$$

where in the last line we have used the explicit form of the prior for a (Eq. 4.9), that further restricts the integration region to $a \leq 1$. The general result of this integral can be written in terms of the incomplete Beta function. However, a simpler form is obtained if ω is an integer. Indeed in this case

$$\begin{aligned} P(\delta_m, \dots, \delta_1 | \mu) &= \frac{\epsilon(1 + \omega)}{2^m(m + \epsilon)} \sum_{k=0}^m \frac{1}{|\delta_k(\mu)|^{m + \epsilon}} \\ &\quad \times \sum_{j=0}^{\omega} (-1)^j \binom{\omega}{j} \int_{\min(1, a_{k+1})}^{\min(1, a_k)} da a^{(m + \epsilon)k - \frac{m(m+1)}{2} + j} \\ &= \frac{\epsilon(1 + \omega)}{2^m(m + \epsilon)} \sum_{k=0}^m \frac{1}{|\delta_k(\mu)|^{m + \epsilon}} \sum_{j=0}^{\omega} (-1)^j \binom{\omega}{j} \times \end{aligned}$$

$$\times \begin{cases} \log \frac{\min(1, a_k)}{\min(1, a_{k+1})} & \text{if } (m + \epsilon)k - \frac{m(m+1)}{2} + j + 1 = 0 \\ \frac{\min(1, a_k)^{(m+\epsilon)k - \frac{m(m+1)}{2} + j + 1} - \min(1, a_{k+1})^{(m+\epsilon)k - \frac{m(m+1)}{2} + j + 1}}{(m+\epsilon)k - \frac{m(m+1)}{2} + j + 1} & \text{elsewhere.} \end{cases} \quad (4.15)$$

The advantage of having such a simple analytic form is that the numerical evaluation is very fast. However, nothing prevents one from making more complicated choices for the prior distributions, paying the price that the numerical integration will typically slow down the computation of the distribution.

Equation (4.15) can be used directly in Eq. (4.10) to obtain the probability distribution of the unknown higher orders. Following the derivation of Sect. 3.3 we can then construct the distribution for the observable Σ (Eq. 3.13). A useful property of the result is that the tails of such distributions are dominated by the first missing higher order (this is a consequence of the hierarchy in the arguments of the max function, Eq. (4.13)). Therefore, the asymptotic behaviour of the distribution is given by

$$P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) \sim \frac{1}{|\Sigma - \Sigma_n|^{n+1+\epsilon}}. \quad (4.16)$$

4.4 The posterior of the hidden parameters

Even if the hidden parameters of the model are never part of the final distribution, it is instructive to understand how their distribution changes with the knowledge of the first few orders. The posterior distribution of c and a can be easily computed as

$$\begin{aligned} P(c, a|\delta_n, \dots, \delta_1, \mu) &= \frac{P(\delta_n, \dots, \delta_1, c, a|\mu)}{P(\delta_n, \dots, \delta_1|\mu)} \\ &= \frac{P(\delta_n, \dots, \delta_1|c, a, \mu) P_0(c, a|\mu)}{P(\delta_n, \dots, \delta_1|\mu)} \\ &= \frac{P(\delta_n|c, a, \mu) \cdots P(\delta_1|c, a, \mu) P_0(c, a|\mu)}{P(\delta_n, \dots, \delta_1|\mu)}, \end{aligned} \quad (4.17)$$

where the denominator is given in Eq. (4.15), and also corresponds to the integral over c and a of the numerator. It's clear that even if in our prior c and a were uncorrelated, correlations arise from the model after inference takes place. Using the explicit form of our likelihood (Eq. 4.3), the posterior becomes

$$\begin{aligned} P(c, a|\delta_n, \dots, \delta_1, \mu) &= \frac{P_0(c, a|\mu)}{P(\delta_n, \dots, \delta_1|\mu)} \\ &\times \frac{1}{a^{\frac{n(n+1)}{2}}} \frac{1}{c^n} \theta \left(c - \max \left[\frac{|\delta_1|}{a}, \dots, \frac{|\delta_n|}{a^n} \right] \right) \\ &= \frac{\epsilon(1+\omega)\theta(a)\theta(1-a)}{P(\delta_n, \dots, \delta_1|\mu)} \frac{(1-a)^\omega}{a^{\frac{n(n+1)}{2}}} \\ &\times \frac{1}{c^{n+1+\epsilon}} \theta \left(c - \max \left[1, \frac{|\delta_1|}{a}, \dots, \frac{|\delta_n|}{a^n} \right] \right), \end{aligned} \quad (4.18)$$

where in the second line we have also used the explicit form of our prior (Eqs. 4.8 and 4.9). The theta function cuts out the region of small c and small a , with a boundary given by a sequence of contours identified by $ca^k = |\delta_k|$ in the region $a_{k+1} < a < a_k$ selected by the max function, see Eq. (4.13). On the other hand, the growing negative power of both c and a tend to favour small values of c, a , thus close to this boundary. Since the power of a grows quadratically with n while that of c only linearly, inference tends to favour smaller a at the price of having somewhat larger c . A visual example of this behaviour is given in Fig. 9.

This preference is a nice outcome of the model: inference favours small values of the parameter a that lead to a better behaviour of the expansion, with smaller higher orders. The prediction of the observable will then be more *precise*, namely subject to a smaller uncertainty. Of course one also (and more importantly) wants the prediction to be *accurate*, namely with a reliable uncertainty that does not underestimate the missing higher orders. Judging whether the outcome of the model is reliable is not immediate, and requires explicit examples to verify it. We will come back later to this point in Sects. 4.5 and 7.

Note that all the considerations so far are independent of the prior. The prior has the only role of changing the “starting point” of the inference procedure. With sufficiently many known orders, our choice of the prior will not matter. However, since the number of known orders is typically limited, choosing wisely is important. Since we like the “direction” selected by the inference procedure (it's better to have larger c and smaller a than the opposite) it seems convenient to choose a prior distribution that already favours the same region of parameter space. This is achieved in our Eqs. (4.8) and (4.9) choosing a small value for ϵ and a large value for ω . Note however that a large ω suppresses the region $a \sim 1$, while in the inference there is no such suppression, simply the small a region is enhanced by a negative power of a . Therefore, using a large value of ω may introduce a significant bias. A good compromise, that we advocate as the best choice, is $\omega = 1$. In this way there is a preference for smaller a with only a mild suppression for $a \sim 1$. On the other hand, for ϵ we can choose an arbitrarily small (positive) value, for instance $\epsilon = 0.1$ or $\epsilon = 0.01$. The difference between either choice is relevant only at very low orders: in Fig. 9 only the first plot (the prior) would change visibly. We thus use $\epsilon = 0.1$ in the rest of this work, with the exception of Sect. 4.6 where we will consider variations of the prior parameters.

4.5 Representative results

Before moving further, we now present some representative results of this method. We use our working example of Higgs production to examine the distribution for the cross section. We fix the renormalization scale $\mu = m_H/2$, which is the

Fig. 9 The plots show the probability distribution of the parameters c and a . The first plot (upper left) is the prior, the second (upper right) is the posterior after the knowledge of δ_1 , the third (bottom left) adds the knowledge of δ_2 and the last (bottom right) the one of δ_3 . The observable under consideration is the inclusive Higgs cross section, Sect. 2.5,

for fixed scale $\mu = m_H/2$. Each line represents the contour $ca^k = |\delta_k|$, for increasing values of k from 1 to 3, to clarify the role of the theta function. Note that the c axis is shown in log scale, but the probability distribution is for c and not $\log c$. For this plot, we used $\omega = 1$ and $\epsilon = 0.1$ for the prior parameters

most widely used choice for this process. Of course the result of this model depends on the choice of scale made, and in addition it does not know anything about the scale dependence. So the input for this model are just 4 numbers, the values of the cross section at the chosen scale at LO, NLO, NNLO and N^3 LO. How to deal with scale dependence in this model will be discussed in Sect. 6.

In Fig. 10 we plot the distribution for the observable Σ (the Higgs cross section), Eq. (3.13), using only the first missing higher order (solid lines) or the first two missing higher orders (dashed lines). The four colors correspond to different knowledge: LO (black), NLO (green), NNLO (blue) and N^3 LO (red).

We immediately notice that the solid and dashed curves are basically identical. This implies that, for any given status of knowledge, the uncertainty in this model is dominated by the first missing higher order. This is a consequence of having $a < 1$, that implies that higher and higher orders are smaller and smaller in this model. Moreover, given that the distribution of c and a favours small values of a (Fig. 9), the impact of the next higher orders is significantly smaller than the first missing higher order. The reason why the curves are almost identical also depends on the fact that the tails of these distributions behave as a negative power (Eq. 4.16), and are thus very “long”: therefore, the effect by the next higher orders of “enlarging” the distribution is almost invisible as the

Fig. 10 Probability distributions of the Higgs cross section with different states of knowledge

tails already cover large deviations of the observable. We thus conclude that for this model it is sufficient to consider only the first missing higher order, so we can directly use Eq. (3.14) for all practical applications. This makes the implementation of this model very fast, as no numerical integration is needed.

Let us now comment on the predictions of this model. Of course, every distribution is centered on the cross section at the known order, and it is symmetric, because in our assumption Eq. (4.2) the sign of the missing higher order is treated agnostically (we will consider a possible way of taking the sign into account in Sect. B.5). We consider the four states of knowledge in turn.

- When only the LO is known, the shape of the curves is fully determined by the prior (no inference took place yet), so the black curve is not particularly useful. Indeed, for our choice of prior, the distribution is very broad, so it is basically not predictive. Note also that the distribution is barely normalizable, because it asymptotically behaves as $|\Sigma - \Sigma_0|^{-1-\epsilon}$, Eq. (4.16), with ϵ small. Another consequence of this behaviour is that the variance of this distribution is infinite. Also, the distribution allows with a high probability ($\sim 40\%$) unphysical negative values of the cross section. These can be avoided imposing a positivity constraint on Σ , but it is instructive to see how much information is needed to obtain a distribution sufficiently narrow to have negligible probability of negative cross section.
- The knowledge of the NLO allows to perform the first step of the interference, so the green curve provides the first non-trivial prediction. However, it is still an “immature” prediction, because only one piece of information has been used. Indeed, this distribution is still very broad, with tails asymptotically behaving as $|\Sigma - \Sigma_0|^{-2-\epsilon}$, and a cusp that favours a single point but not really a region. For $0 < \epsilon < 1$, as we advocate, also this distribution has

infinite variance. Also in this case, there is a significant probability ($\sim 4\%$) of negative cross section. Therefore, the NLO alone does not provide sufficient information in this model to make a precise prediction.

- Once the NNLO is known, the situation changes. On top of having tails that decrease more rapidly, the blue curve clearly identifies a more probable region where the distribution has some sort of bump. The prediction is still rather uncertain, but at least the procedure seems to converge. The probability of negative cross section is reduced to less than 0.3% .
- As a confirmation of this, we see that the knowledge of the N^3 LO improves the situation even more. Now the red distribution is well localized, with a clear bump, which is also well compatible with the prediction at the previous order. The uncertainty is clearly reduced, and the probability of negative cross section becomes negligible (less than 0.01%)

We conclude that this method works well, but requires a sufficient number of known orders to be predictive. Two orders (NLO) is the absolute minimum, but three orders (NNLO) are probably necessary to achieve a decent precision. Beyond NNLO (four or more orders) it should work very well. Note also that the distributions do not look like a gaussian, and therefore they *cannot* be approximated with a gaussian distribution in applications.

To further appreciate the results of the approach, we compute from the distributions a number of quantifiers, namely mean (which equals mode and median given the symmetry of the distributions), standard deviation, and degrees of belief (DoB) intervals. We summarize them in Fig. 11. In this summary, we consider both the scale choice used previously, namely $\mu = m_H/2$, and also another one, namely $\mu = m_H$, to emphasise the scale dependence of this approach at this level. A number of comments are in order.

- The uncertainty is clearly reduced visibly and substantially when adding information from perturbative orders. All the uncertainty quantifiers (standard deviation, 68% and 95% DoB intervals) shrink with increasing the knowledge.
- Because of the high tails, the standard deviation is quite large: infinite in the first two cases (knowledge of LO and NLO), and larger than the 68% DoB interval in the other two cases. For the very same reason, the 95% DoB interval is always much larger than the 68% DoB interval.
- In the shrinking of the uncertainty when adding information the results are always well compatible with the previous orders. All 95% DoB intervals are contained in those of the previous orders. The same is true for the standard deviation. The 68% DoB intervals, instead, are contained in the same interval of the previous order only

Fig. 11 Summary of the distributions of Fig. 10 at the various orders given in terms of mean, standard deviation and degree of belief intervals. The canonical scale variation uncertainty bands are also shown.

in some cases, but this is perfectly acceptable, as there is large probability (32%) that the true result is outside that interval.

We conclude once again that the knowledge of the NLO in this method is not sufficient to achieve a decent precision, while when at least the NNLO is known the method works very well. Note that the reliability is manifest, as the uncertainties always cover nicely the next orders. This is achieved at a price, namely having large uncertainties with a poor state of knowledge. But this is perfectly meaningful, because with too few information on the perturbative expansion it is impossible to predict with precision the value of the observable.²³ We note that, despite the large uncertainties at low orders, once the N³LO is known the prediction is very precise, at least for some “good” scale choices (we will come back to this point in Sect. 6).

In the same plot also the conventional scale variation uncertainty is shown for comparison, both its asymmetric version (in black) and its symmetrized version (in grey). We stress once again that these conventional uncertainty bands have no probabilistic interpretation, but they just represent the “error” usually assigned to perturbative results. Since the probability distributions are symmetric, their mean coincides with the “central value” of the conventional scale variation approach. There is little to say about these “error bars”, given

²³ This is the reason for which using additional information on the expansion would be very precious. The information on the scale dependence, that we will exploit in the next section, can also be combined with this model, as we shall consider in Sect. B.3.

The left plot is for the standard scale $\mu = m_H/2$, while the right plot is for the “natural” scale $\mu = m_H$

the absence of a probabilistic meaning. We observe that, accidentally, the width of these “error bars” is similar in size to the 68% DoB intervals, with the exception of that of the LO.

4.6 Dependence on the prior

Before concluding the section we present a study on the impact of the choice of the prior in the final probability distribution for the observable. To do this, we consider variations of the prior parameters (Sect. 4.2) and see how the final distributions change. For definiteness, starting from the default choice of parameters $\epsilon = 0.1$ and $\omega = 1$, we vary one of the two keeping the other at its default value. For the prior for c , we consider $\epsilon = 0.01$ and $\epsilon = 1$, while for a we take $\omega = 0$ and $\omega = 2$. The results of these variations are shown in Fig. 12, both at the level of the distributions for the cross section, and in terms of the statistical estimators (considering in this case only the most accurate results with the knowledge of the N³LO).

At the distribution level, we notice that the shape of the distribution is not particularly affected by the variation of parameters, as they all look qualitatively the same. However, they vary quantitatively, as for instance the “width” and the “height” of each distribution changes with the prior parameters. It has to be noted that, as expected, these differences are more marked at low orders (less information, thus more dependence on the prior), and become less and less relevant increasing the order (more information).

To better appreciate the impact of the different priors in the right plot of Fig. 12 the statistical estimators (standard

Fig. 12 Dependence of the probability distribution of the observable at various orders on the parameters of the prior (left plot), and statistical estimators for the distribution at N³LO for the same choices of parameters

deviation and DoB intervals) are shown for the distribution at N³LO obtained with each choice of the prior parameters. We see that the difference are somewhat marked. In particular, the largest uncertainties are obtained with $\omega = 0$, which is obvious as it allows larger values of the expansion parameter a with larger probability (see discussion in Sect. 4.4). Similarly, larger ϵ also leads to a larger uncertainty. This is due to the fact that, if smaller values of c are favoured, the constraints from the theta function in Eq. (4.18) forces a to be larger (see also Fig. 9). In agreement with these observations, the result with $\omega = 2$ has smaller uncertainty due to the stronger suppression of the $a \rightarrow 1$ region, and also the result with $\epsilon = 0.01$ is more precise, even though here the difference is very mild due to the tiny difference in the parameter value. Note that in all cases the largest difference in the 95% DoB interval, which is dominated by the tails of the distributions that are mostly affected by the choice of prior. The 68% DoB interval, dominated by the peak region, is more stable.

We conclude that there certainly is a non-negligible bias due to the choice of the prior. Nevertheless, we believe that our suggested choice of parameters represents a good compromise between accuracy and precision. One can obtain more conservative results with larger ϵ and/or smaller ω , but if the suggested values already provide accurate results (as we shall verify with other examples in Sect. 7) varying them does not seem convenient. Conversely, one could obtain more precise results with a more aggressive prior for a , namely with larger ω , but this may lead to a loss of accuracy, which is to be avoided.

5 Model 2: a new approach using scale variation information

We now move to our second model, that uses information on the scale dependence of the perturbative expansion. This model is specific for a QFT application.

5.1 The hypothesis of the model

We now consider another model that rather than looking at the behaviour of the perturbative expansion uses only the information coming from the scale variation to infer the size of the missing higher orders, similarly to what is done in the canonical scale variation approach. The foundation of this method relies on the fact already stressed several times that the scale dependence of a perturbative expansion truncated at order n is formally of order $n + 1$, namely of the same order as the missing higher orders, Eq. (2.15). In Sect. 3.2 we have introduced the estimators $r_k(\mu)$, Eq. (3.9) to quantify the scale dependence of the perturbative expansion at order k , which are thus objects of order $k + 1$ (Eq. 3.6). We shall thus expect

$$|\delta_k(\mu)| \sim r_{k-1}(\mu), \quad (5.1)$$

namely, the term of the perturbative expansion at order k should be similar in size to the scale variation estimator of the previous order $k - 1$. From this vague statement we can now propose the hypothesis of our model, which is

$$|\delta_k(\mu)| \leq \lambda r_{k-1}(\mu) \quad k < k_{\text{asympt}}, \quad (5.2)$$

where λ is a hidden parameter of the model, again dependent on (correlated with) μ . This condition does not tell us anything about the behaviour of the expansion, but only relies on the goodness of the scale variation numbers r_k as estimators of the next order, up to a factor λ .

Note that if the scale variation estimator r_k is accidentally small, λ is forced to be large to accommodate the condition Eq. (5.2). This is definitely undesirable, as a large value of λ due to an accident of the scale dependence would enlarge the uncertainty making the model less predictive. This would be the case if r_k was constructed as the derivative of the observable: close to stationary points, the derivative is accidentally small, see Fig. 5 in comparison with Fig. 4. Our definition of r_k , Eq. (3.9), is indeed designed to avoid such problematic behaviours, as one can see from Fig. 6. Since it is not unusual that the leading order is scale independent, one has to redefine r_0 , because in this cases it would be identically zero. As anticipated in Sect. 3.2, in these cases we simply set $r_0 = 1/2$, which represents a rather conservative choice (larger values can also be considered to be even more conservative). Note that in these cases the first non-trivial information comes from the NNLO, so the probability distributions obtained from the knowledge of the NLO only have to be considered biased by the arbitrary choice of r_0 .

It is instructive to understand the connection between our hypothesis Eq. (5.2) and the canonical scale variation approach. The relation between the two is very simple in the case in which the scale dependence is linear in $\log \mu$. In such a case, $r_k(\mu)$ coincides with the derivative of the observable with respect to $\log \mu$. The canonical scale variation Eq. (1.1) in this limit would predict the “error” to be $\log 2$ times the same derivative.²⁴ Therefore, if the canonical scale variation “error” is thought as an absolute limit on the next order, it would coincide with Eq. (5.2) with $\lambda = \log 2$ fixed. This new method can thus be viewed as an improved version of canonical scale variation in which the variation factor is not fixed to be 2, rather it is inferred from the perturbative expansion itself to be such that the uncertainty estimate is reliable.

The condition Eq. (5.2) must be translated into a probability distribution for the coefficients δ_k given λ and r_{k-1} . We assume that the condition is strictly satisfied, and within the allowed range all values are equally likely. This leads to the likelihood

Fig. 13 Bayesian network for the scale variation model

$$P(\delta_k | r_{k-1}, \lambda, \mu) = \frac{1}{2\lambda r_{k-1}(\mu)} \theta(\lambda r_{k-1}(\mu) - |\delta_k(\mu)|), \quad (5.3)$$

$$k > 0, \quad k < k_{\text{asympt}},$$

where we have stressed that this conditional probability makes sense only for $k > 0$, because at LO there is no previous order to be used to compute scale dependence (in other words, r_{-1} does not exist). This is once again in line with the fact that the LO alone does not bring any information on the behaviour of the expansion. Since the likelihood depends on a single hidden parameter, the resulting model is not very “flexible”, and will for example lead to non-smooth distributions for the observable. To add some flexibility, one may allow violations of the bound Eq. (5.2). This possibility will be explored later in Sect. B.1.

The r_{k-1} coefficient is assumed to be a given information in the likelihood Eq. (5.3). In fact, r_{k-1} can be computed from the knowledge of all the previous $\delta_{k-1}(\mu), \dots, \delta_1(\mu)$ and $\Sigma_0(\mu)$. This notation must thus be interpreted as a shorthand notation for the complete expression

$$P(\delta_k | r_{k-1}, \lambda, \mu) \equiv P(\delta_k | \delta_{k-1}, \dots, \delta_1, \Sigma_0, \lambda, \mu). \quad (5.4)$$

Note that, differently from the geometric behaviour model of Sect. 4, here there is an explicit dependence on $\Sigma_0(\mu)$, which is needed to compute the scale variation estimators r_k . Equation (5.4) implies that the δ_k coefficients are not independent (as they were in the geometric behaviour model). The dependence, however, is always only on the previous ones. Thanks to this, the inference procedure is straightforward. The Bayesian network of this model, using explicitly the r_k parameters, is depicted in Fig. 13.

5.2 The choice of prior

This model depends on a single parameter λ . In principle, it can take any positive value. However, Eq. (5.1) suggests that it should be of $\mathcal{O}(1)$. For this reason, we find it convenient to choose its prior distribution such that large values are strongly suppressed. We thus suggest an exponential behaviour

²⁴ This comparison is not exact because in our definition r_k is the slope of the observable normalized to the observable at order k , while the δ_k coefficients are normalized to the LO.

$$P_0(\lambda|\mu) = \frac{1}{\Gamma(1+\gamma)} \lambda^\gamma e^{-\lambda} \theta(\lambda), \quad \gamma \geq 0, \quad (5.5)$$

where γ is a parameter that changes the shape of the distribution in the region of small λ . The mode of the distribution is in

$$\lambda_{\text{mode}} = \gamma. \quad (5.6)$$

Therefore, it makes sense to choose either $\gamma = 1$, so that our expectation $\lambda = \mathcal{O}(1)$ is contained in the prior, or $\gamma = 0$, that represents a more agnostic (less informative) choice. The dependence on γ will be explored in Sects. 5.4 and 5.6. Similarly to the previous model, the prior Eq. (5.5) is chosen to be independent of μ .

5.3 Inference on the unknown higher orders

The probability distribution of the missing higher orders given the known ones, Eq. (3.15), is

$$\begin{aligned} P(\delta_{n+j}, \dots, \delta_{n+1} | \delta_n, \dots, \delta_1, \Sigma_0, \mu) \\ = \frac{P(\delta_{n+j}, \dots, \delta_{n+1}, \delta_n, \dots, \delta_1 | \Sigma_0, \mu)}{P(\delta_n, \dots, \delta_1 | \Sigma_0, \mu)}, \end{aligned} \quad (5.7)$$

$j \geq 1, \quad n+j < k_{\text{asympt}},$

where on the right-hand side we have included $\Sigma_0(\mu)$ as part of the information as it is only needed to compute the scale dependence. Both numerator and denominator can be directly computed from

$$\begin{aligned} P(\delta_m, \dots, \delta_1 | \Sigma_0, \mu) &= \int d\lambda P(\delta_m, \dots, \delta_1, \lambda | \Sigma_0, \mu) \\ &= \int d\lambda P(\delta_m | \delta_{m-1}, \dots, \delta_1, \lambda, \Sigma_0, \mu) P(\delta_{m-1}, \dots, \delta_1, \lambda | \Sigma_0, \mu) \\ &= \int d\lambda P(\delta_m | \delta_{m-1}, \dots, \delta_1, \lambda, \Sigma_0, \mu) \cdots P(\delta_1 | \lambda, \Sigma_0, \mu) \\ &P_0(\lambda | \Sigma_0, \mu) \\ &\equiv \int d\lambda P(\delta_m | r_{m-1}, \lambda, \mu) \cdots P(\delta_1 | r_0, \lambda, \mu) P_0(\lambda | \mu) \\ &= \frac{1}{2^m r_0(\mu) \cdots r_{m-1}(\mu)} \int \frac{d\lambda}{\lambda^m} P_0(\lambda | \mu) \\ &\theta \left(\lambda - \max \left[\frac{|\delta_1(\mu)|}{r_0(\mu)}, \dots, \frac{|\delta_m(\mu)|}{r_{m-1}(\mu)} \right] \right), \end{aligned} \quad (5.8)$$

where in the first step we have introduced the hidden parameter, then in the next two steps we have recursively written the probability of each δ_k in terms of the previous ones, in the fourth step we introduced the shorthand notation Eq. (5.4) and removed Σ_0 from the prior which does not depend on it, and finally we have written explicitly the likelihoods Eq. (5.3). Note that the arguments of the max function depend only on the perturbative coefficients and not on the parameter. Therefore, the max function itself is just a number, so we can conveniently define it

$$\lambda_m(\mu) \equiv \max \left[\frac{|\delta_1(\mu)|}{r_0(\mu)}, \dots, \frac{|\delta_m(\mu)|}{r_{m-1}(\mu)} \right] \quad (5.9)$$

to simplify the expressions. With our choice of prior, Eq. (5.5), it is easy to compute the integral

$$\begin{aligned} P(\delta_m, \dots, \delta_1 | \Sigma_0, \mu) \\ = \frac{1}{2^m r_0(\mu) \cdots r_{m-1}(\mu) \Gamma(1+\gamma)} \int_{\lambda_m(\mu)}^{\infty} d\lambda \lambda^{\gamma-m} e^{-\lambda} \\ = \frac{\Gamma(1+\gamma-m, \lambda_m(\mu))}{2^m r_0(\mu) \cdots r_{m-1}(\mu) \Gamma(1+\gamma)}, \end{aligned} \quad (5.10)$$

where $\Gamma(a, b)$ is the incomplete Gamma function. This result is very simple, thanks to the simplicity of the model (just one parameter) and of the choice of likelihood and prior. The distribution Eq. (5.7) has an equally simple form

$$\begin{aligned} P(\delta_{n+j}, \dots, \delta_{n+1} | \delta_n, \dots, \delta_1, \Sigma_0, \mu) \\ = \frac{\Gamma(1+\gamma-n-j, \lambda_{n+j}(\mu))}{2^j r_n(\mu) \cdots r_{n+j-1}(\mu) \Gamma(1+\gamma-n, \lambda_n(\mu))}. \end{aligned} \quad (5.11)$$

The derivation of the distribution for the observable Σ is then straightforward according to the procedure of Sect. 3.3. There is, however, an important difference with respect to the geometric behaviour model. To make inference on the first missing higher order $n+1$ the knowledge of the scale dependence r_n of the order n is required. This is fine because the order n is known. However, to make inference on the next higher order, it is necessary to have the subsequent coefficient r_{n+1} . This is not known. In principle, this is not a real problem, as r_{n+1} can be computed from δ_{n+1} .²⁵ However, to compute the scale dependence at order m , the m -loop β -function is required. Therefore, the scale dependence can only be computed up to the order at which we know the β -function of the theory. In the case of QCD, the β -function is presently known up to 5 loops [53], which implies that the largest coefficient which can be inferred in this model is δ_7 for observables that are scale independent at LO and δ_6 otherwise (see Sect. A.1 for further detail).

5.4 The posterior of the hidden parameter

We now turn our attention to the posterior distribution of the hidden parameter, that as we have seen in the previous model brings interesting information. The posterior for λ is given by

$$\begin{aligned} P(\lambda | \delta_n, \dots, \delta_1, \Sigma_0, \mu) \\ = \frac{P(\delta_n, \dots, \delta_1, \lambda | \Sigma_0, \mu)}{P(\delta_n, \dots, \delta_1 | \Sigma_0, \mu)} \end{aligned}$$

²⁵ Note that in practice computing r_{n+1} from δ_{n+1} is time consuming because δ_{n+1} is not fixed but rather it is integrated over in the probability of the observable Eq. (3.13).

$$\begin{aligned}
&= \frac{P(\delta_n, \dots, \delta_1 | \lambda) P_0(\lambda)}{P(\delta_n, \dots, \delta_1)} \\
&= \frac{P(\delta_n | r_{n-1}, \lambda, \mu) \cdots P(\delta_1 | r_0, \lambda, \mu) P_0(\lambda | \mu)}{P(\delta_n, \dots, \delta_1)} \\
&= \frac{1}{\Gamma(1 + \gamma - n, \lambda_n(\mu))} \lambda^{\gamma-n} e^{-\lambda} \theta(\lambda - \lambda_n(\mu)), \quad (5.12)
\end{aligned}$$

where in the last line we have already used the specific expressions for our choice of likelihood and prior, with the definition of λ_n Eq. (5.9). This distribution is very simple: it has the same functional form of the prior, but the power of λ is reduced by a unity for each non-trivial known order, and the region of small λ is cut out by the theta function.

In Fig. 14 we show the posterior distribution for our example of Higgs production. In the left plot we use $\gamma = 1$, and in the right plot $\gamma = 0$. We see that the difference is obviously marked for the prior (black dashed curve), but becomes less and less relevant when adding information. Most importantly, the lower limit on λ imposed by the theta function, which plays an important role, is independent of the parameter γ . The figure also shows a striking effect of the inference: the lower limit on λ , imposed by $\theta(\lambda - \lambda_n(\mu))$, is fully determined by the first non-trivial order, δ_1 . Indeed, it's clear that λ_n Eq. (5.9) coincides with $|\delta_1|/r_0 = 2.72$ for all values of $n = 1, 2, 3$ (for $\mu = m_H/2$). Indeed, the next orders give $|\delta_2|/r_1 = 1.53$ and $|\delta_3|/r_2 = 0.69$, both smaller than the first (and decreasing). This means that, in this case, the actual next order is smaller than what estimated by the scale variation parameter r_k times the smallest allowed value of λ . This is the case despite our definition of r_k , which is normalized to Σ_k rather than Σ_0 exactly with the purpose of giving a smaller number. This fact can imply two things:

- either this behaviour is just an accident of the observable under consideration, or
- the model itself is based on a wrong (or at least non-optimal) assumption.

We will see later that this behaviour is shared by other observables, which seems to suggest that it is indeed the model assumption to be problematic. However, the actual pattern of $|\delta_k|/r_{k-1}$ also depends on the scale μ at which they are computed. For instance, for Higgs production, we have verified that for $\mu > 5m_H$ they all become of the same order, in agreement with the model assumption. It is thus difficult to draw a sharp conclusion on the goodness of the model based just on these observations.

It is interesting to note that, since for some scales the $|\delta_k|/r_{k-1}$ values are decreasing, it could seem convenient to modify the assumption by adding a parameter $\tilde{a}$ behaving like a power, namely

$$|\delta_k(\mu)| \leq \lambda \tilde{a}^k r_{k-1}(\mu) \quad k < k_{\text{asympt}}. \quad (5.13)$$

This model would better describe the behaviour of the known orders, but seems unjustified: where should such a power come from? An alternative interpretation could be that the original assumption Eq. (5.2) is meaningful, but it takes a few orders before the bound is homogeneously satisfied, with the first orders being more “unstable”. In this interpretation, it would make sense to allow a violation of the bound. This option will be explored in Sect. B.1.

We stress that using this model as it is can be regarded as a conservative method: indeed the allowed values of λ are larger than actually needed, thus predicting a larger uncertainty. So for the moment we keep using it, but having in mind this conservative interpretation.

5.5 Representative results

We now present some representative results of this method, using Higgs production as an example, as we did in Sect. 4.5. In Fig. 15 we plot the distributions for the observable given by Eq. (3.13), using a single higher order (solid lines) or two higher orders (dashed lines) to approximate the true cross section, with different status of knowledge: LO (black), NLO (green), NNLO (blue), N³LO (red).

The first striking difference with respect to the geometric behaviour model is the fact that the distribution obtained using two unknown higher orders to approximate the cross section is very different from that obtained using only the first unknown higher order. According to the discussion after Eq. (3.13), we should then keep adding unknown orders in our approximation until the distribution stops changing visibly. Unfortunately, this will likely never happen, and the distribution will likely get broader and broader with higher and larger tails. The reason for this comes from the way the model works. When both δ_{n+1} and δ_{n+2} are used in the approximation for Σ , the model needs r_{n+1} to obtain the distribution for δ_{n+2} . But r_{n+1} depends in turn on δ_{n+1} , which is not fixed, but varies among all possible values, in principle from minus infinity to plus infinity. The trouble is that the models relies on the assumption that the scale variation numbers r_k are good estimators of the higher orders, which in turn relies on the fact that they behave in a perturbative way too. For many of the possible values of δ_{n+1} , however, the corresponding r_{n+1} will violate this assumption and will be much larger than expected, thus predicting large values of δ_{n+2} . These large values unavoidably contaminate the distribution for Σ , making it broader with higher tails. This pattern can only get worse using more orders in the approximation of Σ .

One possible solution to this problem would be to impose some constraint also on r_k , requiring that they behave in a perturbative way, and use this constraint in the inference for the unknown δ_k . This approach is certainly interesting and potentially very powerful, but it is much more complicated to implement. We will discuss it in Sect. B.2. For the time being

Fig. 14 The plots show the probability distribution of the parameter λ , for the prior with $\gamma = 1$ (left plot) and $\gamma = 0$ (right plot). The observable under consideration is the inclusive Higgs cross section (Sect. 2.5) for fixed scale $\mu = m_H/2$

Fig. 15 Same as Fig. 10, but for the scale variation model

we stick to another possible solution, namely restricting our attention to the uncertainty coming from the first unknown higher order only. This approach is acceptable, provided the result is interpreted for what it is: not a probability for the true observable Σ , but just the probability for the observable at the next order, Σ_{n+1} .

With this limitation in mind, we now comment the shapes of the distributions in Fig. 15, focussing on those using only the first unknown higher order (solid lines).²⁶ With the exception of the first one (black) based only on the knowledge of the LO, they all feature a plateau, surrounded by exponen-

tially decreasing tails.²⁷ The plateau is a direct consequence of the likelihood being a theta function of the absolute value of the order, and of the model being dependent on a single hidden parameter λ . Indeed the inference on λ sets a lower limit, Fig. 14, which in turn implies that all values of the next order lower (in absolute value) than the lower limit of λ times the previous r_k are equally probable. We also see that the distributions shrink nicely by adding information. Because also the tails die faster and faster increasing the number of known orders, it's clear that these distributions tend to a uniform distribution. Therefore, one can interpret this model as a sort of provider of an “absolute error”, a region where one is almost certain that the next order will lie, but without any clue on where inside that region.

In Fig. 16 we show the results using quantifiers of the distributions, for two values of the scale, $\mu = m_H/2$ (left) and $\mu = m_H$ (right). We see that the standard deviation is very similar to the 68% DoB interval, and the 95% DoB interval is only slightly larger than the 68% DoB interval due to the exponential suppression of the tails, differently from what happened in the geometric behaviour model. Also in this case the bands shrink nicely increasing the order, with a good compatibility with the previous bands (remember however that each band now just represents the uncertainty of the next order, and not of the full result). The only exception is the LO, which has small uncertainties not very compatible with the NLO. Note that in this case the distribution at LO is not fully determined by the prior, but also depends on r_0 , the scale variation number of the LO. Therefore the small

²⁶ The only comment on the dashed curves that we make is that they become asymmetric, and interestingly the asymmetry tends to favour regions in the direction where the actual next orders indeed are. The price to pay is the higher tails. The model of Sect. B.2 will preserve this nice property while removing the issue of the high tails.

²⁷ Because the tails die exponentially, the probability of negative cross section is much lower than in the geometric behaviour model. Only at LO it is significant ($\sim 6\%$), while at NLO it is already below 0.3%, becoming totally negligible ($< 10^{-5}$) at NNLO and beyond.

Fig. 16 Same as Fig. 11, but for the scale variation model

uncertainty may also be due to the fact that r_0 is not a very good representative of the NLO. Anyway, also in this case one must at least know two orders to let the inference work, so the LO uncertainty is not very relevant.

We also compare these results with the canonical scale variation approach. We observe that the “error bars” of the canonical method at NNLO and N^3 LO are close in size (when they are symmetrized) to the standard deviation (or the 68% DoB interval) of our model. This is purely accidental.

5.6 Dependence on the prior

We finally consider variations of the prior parameter γ . On top of our default choice ($\gamma = 1$), we take a smaller value $\gamma = 0$ (more agnostic choice) and a larger value $\gamma = 2$ (assuming a larger prior value for λ of order 2, see discussion in Sect. 5.2). The results of these variations are shown in Fig. 17, for the distributions of the cross section (left plot) and for the statistical estimators at N^3 LO (right plot).

The most striking feature of the first plot is that the width of the plateau is independent of the prior. This is a consequence of the fact that the plateau is determined by the theta function in the likelihood, and is thus a feature that does not depend on the prior. What changes when varying the prior is the functional behaviour of the tails of the distributions, and due to the normalization constraint this also modifies the “height” of the plateau. The difference are however very mild. The only distribution that changes significantly is the one given the knowledge of the LO only, that we recall is fully determined from the prior and it thus does not contain any relevant information.

The stability of the result is also manifest when looking at the summary plot (on the right in the figure). All bands are basically identical. Note that one could obtain a larger difference by changing the functional form of the prior, e.g. replacing the exponential form with a power law. However, even in this case only the tails of the distributions will be affected, thus still leading to small (though perhaps more marked) differences. We conclude that this method, despite its limitations, has the advantage of being essentially independent of the choice of the prior.

6 Dealing with scale dependence

So far we have presented two models that allow to construct a probability distribution for the missing higher orders of a given perturbative expansion *at a fixed scale*. We have already seen in the examples that the same method on the same observable at different scales produces different results. Hopefully with sufficiently many orders the dependence on the scale will be mild, but nevertheless this dependence is undesirable. In this section we propose an approach to produce a prediction for the observable and its uncertainty that is scale independent.

6.1 The scale as a model parameter

The idea to remove the scale dependence from the probability distribution is very simple: promoting the scale μ to be a parameter of the model. The corresponding Bayesian network is depicted in Fig. 18. In this way, the scale dependence is easily removed by simply *marginalizing* over μ :

Fig. 17 Dependence of the probability distribution of the observable at various orders on the parameter of the prior (left plot), and statistical estimators for the distribution at N³LO for the same choices of parameter

$$\begin{aligned}
 P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0) &= \frac{P(\Sigma, \delta_n, \dots, \delta_1, \Sigma_0)}{P(\delta_n, \dots, \delta_1, \Sigma_0)} \\
 &= \frac{\int d\mu P(\Sigma, \delta_n, \dots, \delta_1, \Sigma_0, \mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0)} \\
 &= \frac{\int d\mu P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) P(\delta_n, \dots, \delta_1, \Sigma_0, \mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0)} \\
 &= \int d\mu P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) P(\mu|\delta_n, \dots, \delta_1, \Sigma_0).
 \end{aligned}
 \tag{6.1}$$

The equation above gives the probability of the observable Σ given the knowledge of the first orders up to NⁿLO. The latter are seen not as the values of the perturbative contribution at a given scale, but in a more abstract sense. This probability is expressed as the integral over μ of the same probability but at fixed μ , which was given in Eq. (3.13) and was the object discussed so far in this paper, and the *posterior* distribution for μ given the first $n + 1$ orders, which is in turn given by

$$\begin{aligned}
 P(\mu|\delta_n, \dots, \delta_1, \Sigma_0) &= \frac{P(\delta_n, \dots, \delta_1, \Sigma_0, \mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0)} \\
 &= \frac{P(\delta_n, \dots, \delta_1, \Sigma_0|\mu) P_0(\mu)}{\int d\mu P(\delta_n, \dots, \delta_1, \Sigma_0|\mu) P_0(\mu)},
 \end{aligned}
 \tag{6.2}$$

where $P_0(\mu)$ is the prior distribution for the scale μ . A number of comments are in order.

The first and most important comment is related to the interpretation of this procedure. Indeed, the fact that μ is a parameter of the model implies that it has some sort of “physical” meaning, since it is possible to make inference on it and construct a posterior distribution (Eq. 6.2). That is, inference selects values of the scale that are “better” than oth-

Fig. 18 The most general Bayesian network of models of inference of missing higher orders, including explicitly the scale μ as a parameter over which one can marginalize

ers. This seems to be in direct contradiction with the fact that the renormalization scale is unphysical, and thus any value is equally good. In other words, there is no physical motivation for preferring some values of the scale and disfavoring others.²⁸

This apparent contradiction is solved by noting that the probabilistic inference on μ selects values that are “better” according to the model. Indeed the posterior distribution for μ is proportional to $P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)$, that is the model-dependent probability of the sequence of the first n orders at the scale μ . Inference selects scales for which this probability is higher, namely scales for which the perturbative expansion fits well with the assumptions of the model. Therefore, not

²⁸ This is not entirely true. If the scale μ is chosen to be too different from the hard scale of the process Q , then large logarithms of the ratio μ/Q invalidate the perturbative expansion. This observation however allows only to disfavor some extreme regions of the scale, while leaving a large region of allowed values.

only this procedure is perfectly acceptable, but it also does what we would have desired, namely selecting the scales with the best perturbative expansion (according to the assumptions of the model), within a well defined probabilistic framework.

As a second comment, we observe that since the scale has become a parameter of the model, a prior distribution $P_0(\mu)$ has to be specified. This prior distribution will contain our prejudices on what are the most appropriate scales, and it thus represents in some sense a “residual scale dependence” of the result. However, we will see that if the prior distribution is sufficiently broad and non-informative, the results will depend very little on its precise form and size.

Another comment is related to the fact that the only new ingredient needed to compute the scale-independent distribution is just the prior $P_0(\mu)$. Indeed, the distribution $P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu)$ in Eq. (6.1) was the subject of the previous sections, and the posterior distribution for μ depends on the prior and on $P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)$, which is the basic model-dependent object discussed in Sects. 4 and 5, which is also needed in the computation of $P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu)$, see Sect. 3.3. Therefore, no additional computations are needed to “remove” the scale dependence: scale-independent results can be obtained with just a simple (numerical) integration.

We can use Eq. (3.11) to write explicitly the scale-independent distribution Eq. (6.1) in terms of fundamental objects. We find, approximating the observable with the order $n + j$,

$$\begin{aligned} P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0) &= \int d\mu P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) P(\mu|\delta_n, \dots, \delta_1, \Sigma_0) \\ &\simeq \int d\mu \int d\delta_{n+j} \cdots d\delta_{n+1} \delta(\Sigma - \Sigma_{n+j}(\mu)) \\ &\quad \times \frac{P(\delta_{n+j}, \dots, \delta_{n+1}, \delta_n, \dots, \delta_1, \Sigma_0|\mu) P_0(\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0)}, \quad (6.3) \end{aligned}$$

having used Eq. (3.15) and Eq. (6.2). As we did in Sect. 3.3, we can use the delta function to compute one of the δ_k integrals. Alternatively, we can consider moments of the distribution, and use the delta function to integrate over Σ . We have for the generic p moment

$$\begin{aligned} \langle \Sigma^p \rangle_{\delta_n, \dots, \delta_1, \Sigma_0} &= \int d\Sigma \Sigma^p P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0) \\ &\simeq \int d\mu \int d\delta_{n+j} \cdots d\delta_{n+1} \Sigma_{n+j}^p(\mu) \\ &\quad \times \frac{P(\delta_{n+j}, \dots, \delta_{n+1}, \delta_n, \dots, \delta_1, \Sigma_0|\mu) P_0(\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0)}, \quad (6.4) \end{aligned}$$

from which we can compute for instance the mean of the distribution and its variance. It is interesting to notice that in the cases in which the probability $P(\delta_n, \dots, \delta_1, \Sigma_0|\mu)$ is symmetric for $\delta_k \rightarrow -\delta_k$, as it is the case of the geometric behaviour model, the computation of the mean of the distribution is simplified as the contributions from unknown higher orders integrate to zero. Therefore, in such a symmetric case, the mean reduces to

$$\langle \Sigma \rangle_{\delta_n, \dots, \delta_1, \Sigma_0} = \int d\mu \Sigma_n(\mu) \frac{P(\delta_n, \dots, \delta_1, \Sigma_0|\mu) P_0(\mu)}{P(\delta_n, \dots, \delta_1, \Sigma_0)}, \quad (6.5)$$

which depends on the *known* N^n LO result $\Sigma_n(\mu)$. This is the same result that we would obtain when approximating the observable using only the known orders ($j = 0$). If we were working at fixed scale, this would just be the number that comes out of the computation. But since this number is scale dependent, it generates a distribution, Eq. (6.3) with $j = 0$, whose mean is given by Eq. (6.5). This is a very interesting consequence of the procedure: even a fixed-order computation at order n cannot be regarded as an “exact” prediction at that order, because it depends on the unphysical scale. The distribution Eq. (6.3) with $j = 0$ represents the best we can say about the N^n LO result, ignoring the missing higher orders. Because the N^n LO is a distribution, one can also compute its standard deviation. This standard deviation would represent a true “scale uncertainty” on the *known* finite order, but it would not contain any uncertainty from missing higher orders. This is very different from the usual approach, where the “canonical scale uncertainty” is used as an estimate of the missing higher order uncertainty.

We conclude by discussing the form of the prior distribution for μ . In principle, any value of μ is allowed, and each value is equally valid as they all lead to the same exact result. However, in practice some values have to be avoided. For instance, in theories like QCD at small scales the coupling grows and invalidates the perturbative hypothesis. Therefore, the scale (in QCD) cannot become too low. Similarly, very large scales (in QCD) are not advisable, as they lead to small values of the coupling, slowing down the convergence of the expansion. At the same time, physical processes are characterized by one (or more) physical scale(s), and the various orders δ_k will contain logarithms of the ratio of such scale(s) with the unphysical scale μ . If μ is very different from the physical scale(s) the logarithms became large and invalidate the perturbative expansion. Therefore, the most convenient and meaningful approach is to use the value of the physical scale(s) to determine a “central value” μ_0 of the scale μ , and then vary it in a range that satisfy the previous (or any other) constraints. Considering that the scale dependence is logarithmic, the easiest option is a flat distribution in the logarithm of the scale, namely

$$P_0(\mu) = \frac{1}{2 \log F} \frac{1}{\mu} \theta \left(\log F - \left| \log \frac{\mu}{\mu_0} \right| \right), \quad F > 1, \quad (6.6)$$

where F is factor that sets the size of the interval of allowed scales, assumed to be symmetric about μ_0 (which is not a necessary condition). Alternatively, one could consider a distribution peaked in μ_0 , for instance a gaussian, perhaps with hard limits to avoid entering the low- and large-scale regions. However, in the spirit of letting the model select the scale that leads to the best convergence properties, a flat distribution seems more natural. The factor F should be sufficiently large to contain a region where good convergence is achieved. This may be either selected by eye looking at the behaviour of the expansion as a function of the scale, or with a step-by-step approach where F is enlarged if the posterior for μ turns out to be peaked close or at the boundary of the allowed region.

We stress that, despite the simplicity of this approach to obtain scale-independent results, this method is very innovative and provides for the first time a consistent way to deal with scale dependence which is compatible with physical requirements.

6.2 A pre-example: application to the Cacciari–Houdeau model

Let us now start to investigate the implication of this method to obtain scale-independent results. To begin with, we consider the original Cacciari–Houdeau model, described in Sect. 2.4. We are not really interested in obtaining results within the CH model, since we have proposed a new model that can be seen as an improved version of it. Rather, we want to use it to point out the importance of using the normalized variables δ_k for defining the model.

To see this, let us start by computing the posterior distribution for μ given the knowledge of the LO. In the CH model, indeed, the LO is treated as a source of information. The posterior is given by (using the notation of Sect. 2.4)

$$\begin{aligned} P(\mu|c_0) &= \frac{P(c_0|\mu)P_0(\mu)}{P(c_0)} \\ &= \frac{P_0(\mu)}{P(c_0)} \int d\bar{c} P(c_0|\bar{c}, \mu) P_0(\bar{c}) \\ &\propto P_0(\mu) \int_{|c_0(\mu)|}^{\infty} \frac{d\bar{c}}{\bar{c}^2} = \frac{P_0(\mu)}{|c_0(\mu)|}, \end{aligned} \quad (6.7)$$

where we have used Eqs. (2.20) and (2.21), and assumed that the hidden parameter $\bar{c}$ and the scale μ are uncorrelated a priori. This computation shows that the posterior distribution for μ given the LO is larger for values of μ for which the

LO $c_0(\mu)$ is smaller.²⁹ This is not particularly clever, as very often NLO and higher corrections are positive, so a larger value of the LO should be favoured instead, to reduce the impact of higher orders and lead to a more well-behaved expansion. This is for instance what happens in the Higgs production process that we have used as an example so far.

The underlying reason for the failure of this procedure is that from the LO only it is impossible to decide which scale should be favoured. However, in this model the situation does not improve when adding the NLO. Indeed the posterior becomes

$$P(\mu|c_0, c_1) \propto \frac{P_0(\mu)}{\max(|c_0(\mu)|, |c_1(\mu)|)^2}, \quad (6.8)$$

which is likely still mostly determined by c_0 , as $|c_1|$ will typically be smaller than $|c_0|$ being a perturbative correction, so that it will never win in the max function. Therefore, the problem with using the information from the LO to make inference in the model is not solved by adding orders. In order to avoid this issue one should really treat the LO as just a prefactor, and start the inference procedure from the NLO, as we do in this work. This is another important motivation for using the normalized quantities δ_k in the definition of the model.

6.3 Scale-independent geometric behaviour model

We now move to the geometric behaviour model discussed in Sect. 4. This model is more interesting and we will use the whole machinery introduced in Sect. 6.1 to obtain scale-independent distributions and uncertainties.

To begin with, we want to stress that in this case, making inference on the normalized perturbative expansion, gives the desired posterior distribution for the scale. We start noticing that the knowledge of the LO Σ_0 does not change the distribution of μ ,

$$P(\mu|\Sigma_0) = \frac{P(\Sigma_0|\mu)P_0(\mu)}{P(\Sigma_0)} = P_0(\mu), \quad (6.9)$$

which comes from the fact that Σ_0 does not appear in the likelihood of the model (Eq. 4.3), and therefore $P(\Sigma_0|\mu) = P(\Sigma_0)$. Equivalently, if one considers instead the normalized LO δ_0 for which a likelihood exists, it cannot change the distribution for μ because $\delta_0 = 1$ is scale independent. This is due to the fact, stressed already several times, that the LO does not carry any information on the behaviour of the expansion, not even when scale dependence is taken into account. The first non-trivial information comes from the

²⁹ In this notation we are assuming that if the LO starts at $\mathcal{O}(\alpha_s^{k_0})$ the coefficient c_0 incorporates this factor $\alpha_s^{k_0}$, so that the expansion Eq. (2.17) holds.

Fig. 19 Posterior distribution for the scale μ within the geometric behaviour model with different states of knowledge, for the Higgs production process

NLO coefficient $\delta_1(\mu)$, that changes the distribution for μ . The general computation is somewhat more complicated than in the CH case, so we stick to the simple case in which the prior for a is flat ($\omega = 0$), so that we get

$$P(\mu|\delta_1) \propto P_0(\mu) \left[\frac{1}{1+\epsilon} + \log \frac{1}{|\delta_1(\mu)|} \right]. \quad (6.10)$$

This probability is clearly larger for scales such that $|\delta_1(\mu)|$ is smaller (at this order the dependence is only logarithmic, but at higher orders power behaving contributions appear). This is the expected behaviour, namely inference selects scales for which perturbative corrections are smaller. Note that for such scales the LO $\Sigma_0(\mu)$ may be larger to “compensate”³⁰ for the next orders (this is what happens in the Higgs case).

Understanding the behaviour of the posterior for μ beyond NLO is difficult analytically. Therefore we now switch to numerical results. In Fig. 19 we show the posterior for μ , Eq. (6.2), for our example of Higgs production, for four different states of knowledge: LO (black), NLO (green), NNLO (blue) and N³LO (red). The first curve corresponds to our prior, that has been taken to be flat in $\log \mu$, Eq. (6.6), ranging from $m_H/10$ to $2m_H$. This choice of range is motivated by the fact that scales between $m_H/4$ and $m_H/2$ lead to better convergence properties of the expansion (see e.g. Ref. [5], and also Fig. 4) and are thus used to define the center of the distribution, while the width is set by requiring that the lowest scale ($m_H/10 = 12.5$ GeV) is still characterised by a sufficiently small strong coupling. Note that the whole range of scales allowed by this prior spans a factor of 20, which is way larger than the usual range of scales considered in standard computations (typically, once the hard scale

³⁰ Clearly, it’s the other way round: when the LO is bigger, the higher orders should be consequently smaller.

Fig. 20 Similar to Fig. 10, after removing scale dependence

of the process is identified, the central scale is chosen within a factor two or four). This means that in our computation we consider also scales that are usually considered too far from the hard scale of the process.

We observe that, as commented before through the analytic computation, the knowledge of the NLO leads to a preference towards smaller scales, even though the preference is mild because the dependence on the NLO coefficient δ_1 is logarithmic. The knowledge of the NNLO changes the situation more dramatically, giving a net preference for scales around $\mu \sim 20$ GeV, and suppressing strongly high scales. The peak corresponds to the scale at which the NNLO correction is zero, which incidentally corresponds to the region where the NNLO scale dependence has a plateau, Fig. 4. At N³LO, the distribution is still peaked, but with a smoother bump towards somewhat larger scales, approximately $m_H/3$. The reason is that the N³LO correction is zero at a higher scale, close to $m_H/2$, and so there is a trade off between the preference of the NNLO correction and that of the N³LO correction. The high scale region is further suppressed.

We now move to the distribution for the observable Σ . We use the approximation based on the first unknown higher order, as we have seen that it captures well the full uncertainty. The results are shown in Fig. 20. We immediately observe that these distributions are asymmetric, which is a consequence of the fact that scale dependence is non-trivial. At the first two orders, the distributions are very broad and not predictive. However, we notice that the procedure to remove the scale dependence has the effect of favouring larger values with respect to the fixed-scale results, in particular covering (at NLO) the region favoured by the next orders with a high probability. The distribution at NNLO features a high peak towards large values of the cross section, at $\Sigma \simeq 51$ pb. This peak is a direct consequence of the peak in the posterior of μ , Fig. 19, and indeed it corresponds to the value of the cross section at the plateau of the NNLO, Fig. 4. Note however that the distribution is very asymmetric, and the mean of the

Fig. 21 Left plot: similar Fig. 11, after removing scale dependence. As a reference the conventional scale variation “error” is also shown for $\mu = m_H/2$. Right plot: dependence on the prior for μ of the most precise result (given knowledge of N³LO). The first three results correspond to a larger support $m_H/16 < \mu < 4m_H$, the next three to the

distribution is to the left of the peak, more in line with the next order. Finally, the distribution at N³LO is the narrowest, with a marked peak and a slight asymmetry.

To better understand these results, we show at each order the mean, mode, median, standard deviation and degree of belief intervals in Fig. 21 (left plot). The pattern found is similar to that of the fixed-scale result (Fig. 11). One difference is that the uncertainties here are somewhat larger, especially at low orders. This however implies that the convergence pattern is improved, as for instance this time the 68% DoB band is always contained in that of the previous order (which is mostly due to the fact that the NLO band is increased). The mean and median are rather similar in all cases, while the mode is not always close, due to the asymmetry of the distributions. At N³LO the prediction is rather precise, certainly much more precise than at previous orders. The mean at N³LO is close to the canonical result at $\mu = m_H/2$.

These results clearly depend on the choice of the prior used for the scale μ . Therefore, we now consider variations of the prior, in order to understand to what extent the results depend on it. We consider two types of variations: the support (range of allowed values), and the shape. Previously we used a support $m_H/10 < \mu < 2m_H$, and we now consider a larger support $m_H/16 < \mu < 4m_H$ and a smaller one $m_H/8 < \mu < m_H$. Secondly, on top of the uniform prior used before, we consider a symmetric triangular prior (with the same support), and a gaussian like prior, obtained cutting the tails after two standard deviations, and adjusting the

parameters such that the support is again the same. The results are shown in the right plot of Fig. 21 for the most precise result based on the knowledge of the N³LO. We notice that enlarging the support has the effect of enlarging the uncertainty, and the converse is also true. Also, the flat prior gives obviously the larger uncertainty, while the triangular and gaussian priors give similar results. Overall, we see that the dependence on the prior is rather mild, certainly milder than the dependence on the result on the scale when it is kept fixed. For instance, in this case the mean, mode and median are more or less independent of the prior. We conclude, as anticipated, that the “residual scale dependence” induced by the freedom in choosing the prior for μ is very much reduced with respect to the canonical scale dependence of the result. We can thus consider these results as almost scale independent.

6.4 Scale-independent scale variation model

We now apply the approach to remove scale dependence applied to the scale variation model. Note that it may seem somewhat contradictory to use the information on the scale dependence twice, once for estimating the higher orders and once for removing the scale dependence. However, there is nothing problematic in practice, because as in the previous case the approach to remove the scale dependence just selects, through inference, scales at which the model provides the best performance, independently of how the model itself works.

Fig. 22 Same as Fig. 19, but for the scale variation model

To begin with, we show the posterior for μ in this model. According to the results of Sect. 5.3, from Eq. (6.2) we get

$$P(\mu|\delta_n, \dots, \delta_1, \Sigma_0) \propto \frac{P_0(\mu)}{r_0(\mu) \cdots r_{n-1}(\mu)} \int_{\lambda_n(\mu)}^{\infty} d\lambda \lambda^{\gamma-n} e^{-\lambda}, \quad (6.11)$$

where $\lambda_n(\mu)$ is given in Eq. (5.9). The integral can be computed analytically, but written in this way it's clear that the integral is larger the smaller $\lambda_n(\mu)$. If the numbers r_k were independent of μ , the posterior would thus be larger for scales for which the largest $\delta_k(\mu)/r_{k-1}(\mu)$ in $k \in [1, n]$ is small. Namely, it would select the scale for which the perturbative corrections were small, which is indeed what we expect this approach to do. Of course, the numbers r_k do depend on μ , so this picture is modified by their presence in the prefactor. However, their dependence is mild (see Fig. 6), and typically milder than the dependence of the δ_k 's, so the argument above gives a sufficiently good description of how the marginalization over μ works in this model.

In order to understand exactly what happens, we now move to the numerical analysis. In Fig. 22 we show the posterior of μ for this model, in analogy with the discussion of the previous section. We see that this time, in the range under consideration, both NLO and NNLO tend to favour small scales. At N³LO, too small scales are disfavoured, and the distribution becomes peaked at $\mu \sim m_H/4$. As in the previous case, high scales are strongly disfavoured.

The probability distribution for the cross section is shown in Fig. 23. Again, after removing the scale dependence the distribution becomes asymmetric. Again, since at low orders smaller scales are favoured, larger values of the cross section are more probable, in line with the results of the higher orders. The distribution at N³LO has a sort of plateau, and strongly decreasing tails, giving a well localised result. However, the

Fig. 23 Same as Fig. 20, but for the scale variation model

uncertainty seems rather large, and specifically larger than the result at fixed scale, Fig. 15.

Finally, in Fig. 24 (left), we show the summary of the results using the usual quantifiers: mean, mode, median, standard deviation, degree of belief intervals. With respect to the fixed-scale results, the uncertainties are larger, and the compatibility between different orders consequently improved. The mean of the distribution is rather stable from NLO onwards, suggesting that perhaps these uncertainties are overestimated (as we pointed out also when discussing the model at fixed scale, Sect. 5). At N³LO, because the tails die rapidly, the 95% DoB interval is only slightly larger than the 68% DoB interval, suggesting that also after computing the scale independent result the interpretation of this model as a provider of some sort of “absolute error” is still valid.

In the same figure the right plot shows the dependence of the most precise result on the choice of prior. The variations are the same introduced in Sect. 6.3. We see that for a wider support of the prior (first three results) the uncertainty increases quite significantly, with the exception of the triangular prior (second result). The explanation for this is that when pushing the smallest scale to be $m_H/16$, the calculation of the r_k , that probes scales up to a factor of 4 smaller than the current μ , gets contributions from scales as small as $m_H/64 \sim 2$ GeV, where α_s becomes large. Indeed, from Fig. 6 we see that r_2 and r_3 explode below $\sim m_H/10$, implying that for those scales the uncertainty in this model becomes huge. Therefore, when this region is included in the support for μ , it contaminates the whole distribution with a very broad contribution, thereby leading to the large uncertainties seen in the figure. The triangular distribution is an exception because it goes to zero at the endpoint, thereby suppressing part of the problematic region. Reducing the range or changing the shape, instead, has a much milder, almost negligible effect. We conclude that in this model it is crucial to keep the region of allowed scales within a sensible range, in order to avoid

artificially large contributions. In particular low scales may be problematic because of the large strong coupling.

7 Validation and applications

In this section, we will consider a number of examples of applications of our methods. First, we focus on examples with a known sum, that will serve as a validation of the procedure. Some of them are just mathematical series without a physical meaning, others are physics examples. We then move to some applications where the true sum is not known. We focus on observables that are known to a high order, so that the results are more significant.

7.1 Convergent series

We start by considering the simplest case, namely a convergent series. This example may not be too significant as we know that perturbative expansions are divergent, but it represents a good check of the machinery. We consider a quasi-geometric series, namely

$$\Sigma = \sum_{k=0}^{\infty} \alpha_s^k A^k \cos(Bk) = \frac{1 - A\alpha_s \cos(B)}{1 + A^2\alpha_s^2 - 2A\alpha_s \cos(B)}. \quad (7.1)$$

The presence of the parameter B induces some “oscillations” in the perturbative expansion, and in the limit $B = 0$ (or a multiple of π) we recover a geometric series. This series is convergent for $|A\alpha_s| < 1$, and it is bounded by a geometric

series $\sum_k |A\alpha_s|^k$. In the following, we fix $A = 4$ and $B = 2$, to have sufficiently large perturbative corrections and a sufficiently non-trivial pattern. We take α_s to be $\alpha_s(m_Z) = 0.118$, so that the sum is $\Sigma = 0.740531$. Note that the effective expansion parameter of this series is $A\alpha_s$, but we made a distinction because we want to also induce an artificial scale dependence, to test the scale variation model and the procedure to remove scale dependence. This is obtained by changing the scale of α_s from m_Z to a generic μ , compensating for this by adding the proper scale-dependent contributions to the expansion coefficients, according to the procedure described in Sect. A.1.

Let us start by ignoring the scale dependence, and consider thus the expansion Eq. (7.1) with $\alpha_s = \alpha_s(m_Z)$. In this case, we can only use the geometric behaviour model of Sect. 4. Of course, the series Eq. (7.1) satisfies the assumption of the model Eq. (4.2), and so the model is expected to work well. We assume to know the first five orders, up to N^4 LO. In Fig. 25 (left plot) we show the probability distribution for the next order given the knowledge of the first orders up to N^4 LO. We see that by adding information the distributions become narrower, and they move closer to the exact result, represented by the vertical black thin line. In particular, the last two distributions (given the knowledge of the N^3 LO and N^4 LO) are in excellent agreement with the exact result. In the same figure (right plot) we show the quantifiers (mean, standard deviation, degree of belief intervals) of these distributions (ignoring the one using only the knowledge of the LO, that is purely determined by priors and is thus not significant). We see that all results are well compatible with the exact result, which is always within one standard deviation, and almost always also within the 68% DoB interval. Increasing

Fig. 24 Same as Fig. 21, but for the scale variation model

Fig. 25 Geometric behaviour model applied to the toy convergent series Eq. (7.1). The left plot shows the distributions for the observable, and the right plot shows a summary of the quantifiers of these distributions. The thin solid black lines represent the exact result

ing the knowledge, the uncertainty clearly shrinks, reaching a very high precision at $N^4\text{LO}$ (and also fairly good at $N^3\text{LO}$). These results show that, at least with this somewhat trivial example, the geometric behaviour model works very well, as expected.

Let us now assume that the series Eq. (7.1) is a QCD observable, so that we can consider its scale dependence. We can thus use the scale variation model of Sect. 5. The probability distributions obtained with this model are shown in Fig. 26 (left plot). Here we see something strange and undesired, namely the distributions become broader when adding the knowledge of the NNLO and of the $N^3\text{LO}$. Note that all of them are well compatible with the exact result, so the method is accurate, but the uncertainties are large, even at $N^4\text{LO}$, so the method is not precise. This peculiar pattern can be understood by looking at the scale dependence of the (fake) observable (Fig. 26, right plot). We observe that the convergence pattern of the expansion deteriorates quickly for scales $\mu < Q \equiv m_Z$, and it improves visibly at higher scales. Therefore, for this observable high scales have to be favoured.

Let us take for example $\mu = 4Q$ and look again at the distributions of the scale variation model. In Fig. 27 (left plot) we see that indeed the distributions have a better pattern, and the distribution with the largest amount of information is rather precise (even though not at the level of the geometric behaviour model). However, the distribution at NLO is still narrower than that of the next order. This undesired behaviour is an artefact due to the fact that the LO is scale independent. Indeed, had we used literally the scale independence of the LO, the number r_0 would be zero, forcing λ to be infinite, making the model not predictive. In Sects. 3.2 and 5.1 we have suggested to arbitrarily set $r_0 = 1/2$ when

the LO is scale independent. This value is sufficiently large to give unimportant restrictions to λ , with the drawback that it allows small values of λ leading to an artificially narrow distribution at NLO. This can be appreciated by looking at Fig. 27 (right plot), where we show the posteriors for λ . We see that the knowledge of the NLO gives a very mild restriction on λ , namely $\lambda \geq |\delta_1|/r_0$, which is small because r_0 is large. Conversely, at the next order, when the first real information on the scale dependence (r_1) is used, the smallest allowed value of λ becomes much bigger. This implies that the posterior $P(\lambda|\delta_1, \Sigma_0, \mu)$, used in the construction of the green curve of Fig. 27 (left), is inaccurate and thus leads to an inaccurate prediction. We conclude that, when the LO is scale independent, the knowledge of the NNLO is needed in order to obtain the first meaningful prediction through this model. Note also that the subsequent orders do not push the lowest value of λ forward, similarly to what happens in the Higgs case. This suggest again that this model is not well suited for precise predictions, as the assumption itself is not optimal. For a possibly better model, see Sect. B.1.

We finally use the procedure of Sect. 6 to obtain scale independent results. According to Fig. 26 (right), it is clear that scales larger than Q have to be favoured. We consider a large range, namely $Q < \mu < 50Q$, to be used as the support of our (flat) prior for μ . The distributions obtained after marginalizing over μ are shown in Fig. 28 (left) for both the geometric behaviour model (above) and the scale variation model (below), and the summary of the distributions through the usual quantifiers are also shown (right). In these last plots we also show the conventional scale “error” assuming a central scale $\mu = Q$. For the geometric behaviour model, we see that the convergence pattern is still excellent, and the uncertainty shrinks more rapidly than in the result at

Fig. 26 Left plot: the distributions obtained applying the scale variation model to the toy series Eq. (7.1) assuming $Q = m_Z$ and $\mu = Q$. Right plot: the scale dependence of the observable. The exact result is shown as a thin solid black line

Fig. 27 Left plot: same as Fig. 26 but for $\mu = 4Q$. Right plot: the posterior distribution for λ at $\mu = 4Q$

fixed scale, providing very precise predictions both at N³LO and at N⁴LO, perfectly compatible with the exact result. For the scale variation model the situation is slightly improved with respect to the fixed scale result, but the uncertainties are still rather large, and of course well compatible with the exact result (with the exception of the NLO, for the reasons mentioned above). We conclude that all the methods proposed so far work as expected, and in particular the geometric behaviour model performs extremely well for a series that is compatible with its assumptions.

7.2 Divergent series

We now move to a somewhat opposite case, namely a factorially divergent series. Provided the expansion parameter is sufficiently small, the expansion has the behaviour of an asymptotic series (see Fig. 3). Therefore, it should be a decent

representative of actual perturbative expansions, and thus a good validation test for our models. We consider the expansion

$$\Sigma = \alpha_s \sum_{k=0}^{\infty} \alpha_s^k A^k k! = -\frac{1}{A} \exp\left(-\frac{1}{A\alpha_s}\right) \Gamma\left(0, -\frac{1}{A\alpha_s}\right), \quad (7.2)$$

where the analytic sum (which is the Borel sum, or equivalently the function of which the series is the asymptotic expansion), is strictly speaking valid only for $A < 0$ (or, if A is complex, for values of A that are not real and positive). However, for $A > 0$ the Borel sum can be still computed up to an ambiguous contribution, which is exponentially suppressed in the coupling and we thus ignore it. Note that we have included a factor of α_s in front, so that the expansion

Fig. 28 Results after applying the procedure of Sect. 6 to remove the scale dependence, using a flat prior in the range $Q < \mu < 50Q$. The upper plots show the distributions (left) and the summary (right) for

the geometric behaviour model, while the lower plots are the same for the scale variation model. The thin solid black lines represent the exact result

sion starts at $\mathcal{O}(\alpha_s)$. This choice allows to have a LO that is scale dependent, to avoid the issues discussed in the previous case when using the scale variation model. The series Eq. (7.2) is assumed to be written at $\mu = Q \equiv m_Z$, with $\alpha_s = \alpha_s(m_Z) = 0.118$, and the expansion at generic μ can be obtained with the procedure described in Sect. A.1. In the following, we will consider both cases $A = -1$ and $A = 1$, dubbed respectively as factorially divergent series with alternating signs and with same signs.

Let us start again by ignoring the scale dependence and thus using only the geometric behaviour model. The results (distributions and summary in terms of quantifiers) are shown in Fig. 29. We observe that in the case of alternating signs ($A = -1$) the method works very well, with the distributions shrinking and converging towards the exact result. This is a consequence of the expansion parameter being sufficiently

small to give an approximately convergent pattern at low orders, as we hope is the case for physical observables. The prediction is thus accurate, and increasing the order becomes also very precise. In the case of all same signs ($A = 1$) the situation is different. The distributions shrink and move towards the exact result, without “reaching” it. In this case, it is clear that the shape of the distribution is inadequate, as the exact result always lies in the tail of the distributions. This behaviour is not surprising, as this series violates the assumptions of the model in a “maximal” way (the series can be considered as the bound of a physical expansion).

Let us now include the scale dependence, and consider the scale variation model. In Fig. 30 (left plots) we show the distributions obtained with this model in the two cases, for $\mu = Q \equiv m_Z$. For alternating signs, the accuracy of the model is manifest, with the distributions always covering

Fig. 29 Geometric behaviour model applied to the toy divergent series Eq. (7.2), for $A = -1$ (upper plots) and $A = 1$ (lower plots). The left plot shows the distributions for the observable, and the right plot shows

a summary of the quantifiers of these distributions. The thin solid black lines represent the exact result

the exact result (within the plateau region when present), and getting narrower increasing the order. However, the precision is not at the level of the geometric behaviour model, as usual. When the coefficients of the series have all the same sign, a similar pattern to that of the geometric behaviour model is observed. In the same figure, the right plots show the (artificial) scale dependence of our observable, as we constructed it. Note that the range is different, as in one case (alternating signs) larger scales lead to a better convergence, while in the other case (same signs) lower scales give a better expansion (this peculiar behaviour is just a consequence of the artificial way in which scale dependence is introduced).

According to this observation, we now apply the procedure of Sect. 6 to remove the scale dependence, using flat priors for μ in the range $Q/2 < \mu < 16Q$ for alternating signs and $Q/8 < \mu < 8Q$ for same signs. The results

are shown in Fig. 31. For the alternating signs series (upper plots) removing the scale dependence improves the precision of both the geometric behaviour model and of the scale variation model, while keeping excellent accuracy. This provides an important validation of all our methods. For the same sign series marginalizing over the scale improves a bit the accuracy, especially for the geometric behaviour model, where the exact result becomes compatible at least within 95% DoB also at high orders. While not perfect in this case, our methods are certainly superior to the conventional scale “error”, also shown in the plots.

Note that the methods can improve if the sign pattern is taken into account. The two expansions considered are characterized by either all positive corrections or corrections with alternating signs. If the pattern is recognised within the model, then the prediction of the next orders can be per-

Fig. 30 Left plots: the distributions obtained applying the scale variation model to the toy series Eq. (7.2) assuming $Q = m_Z$ and $\mu = Q$, for alternating signs (above) and same signs (below). Right plots: the

scale dependence of the observable in either case. The exact result is shown as a thin solid black line

formed accordingly. We will investigate this possibility in Sect. B.5.

7.3 Anharmonic oscillator in quantum mechanics

We now consider a physics example, namely a quartic anharmonic oscillator in quantum mechanics. Using the notation of Ref. [54], its hamiltonian is given by

$$H = p^2 + x^2 + \alpha x^4, \quad (7.3)$$

where α is the “coupling”, namely the perturbative expansion parameter. The energy levels of this anharmonic oscillator as a perturbative expansion in α are known to diverge [55]. The ground-state energy is given by

$$E_0(\alpha) = 1 + 2 \sum_{k=1}^{\infty} \frac{A_k}{2^k} \alpha^k, \quad (7.4)$$

where the first 75 A_k coefficients can be found in Ref. [55]. We recall that these coefficients have alternating signs, namely $A_k = (-1)^{k-1} |A_k|$. They grow faster than a power, but slower than a factorial. The exact result for this ground state energy has been computed in Ref. [54] through a Borel-Padé method. We consider the value $\alpha = 0.1$, which is the smallest value reported in that paper, corresponding to the sum $E_0(0.1) = 1.0653$. For this value the series already starts diverging quite soon: the optimal truncation value of the asymptotic expansion is $k_{\text{asympt}} = 6$.

Since this is a quantum mechanical system, there is no unphysical scale dependence, and therefore we can only use the geometric behaviour model of Sect. 4. The distributions for the ground-state energy given different status of knowledge are shown in Fig. 32 (left). We observe that starting from N³LO the distributions are quite narrow, and oscillate around the exact value, because of the alternating signs of the A_k coefficients. In fact, the prediction does not become very accurate and precise, essentially because we are very

Fig. 31 Scale independent results for alternating signs (upper plots) and same signs (lower plots). The thin solid black lines represent the exact result

Fig. 32 Geometric behaviour model applied to the ground-state energy of the quartic anharmonic oscillator in quantum mechanics. The left plot shows the distributions for the observable, and the right plot shows a

summary of the quantifiers of these distributions. The thin solid black lines represent the exact result

close to the asymptotic point $k_{\text{asympt}} = 6$ that sets the limit of validity of the assumptions of the model.

This can be appreciated also by looking at the summary of the distributions through the usual quantifiers in the right plot of the same figure. It seems surprising that the quality of the uncertainty, though perfectly acceptable (the exact result is always within the 95% DoB interval), is not as good as in the similar case of the factorially divergent series discussed in Sect. 7.2. The reason is that the effective coupling, namely the parameter determining the power growth of the expansion, is larger in this case than it was in the factorially divergent case. In fact, we are now in a condition for which the expansion is barely perturbative, and therefore all our assumptions, including the basic ones of Sect. 2.1, are barely satisfied. Nevertheless, with all these limitations in mind, the method works rather well. Note that taking into account the sign pattern (alternating signs) in the model, as described in Sect. B.5, can help improving the accuracy of the results.

7.4 Higgs production in the threshold limit

We conclude the examples used for the validation of the models with a QCD process. Since, to our knowledge, there are no QCD observables in the perturbative regime for which the exact result is known, we use a trick. Namely, we consider a purely all-order resummed result at a finite logarithmic accuracy, and consider it as if it were the full (exact) result. We expand it in powers of α_s to obtain the first few perturbative orders, to which our methods are applied.

This procedure obviously misses contributions that are not logarithmically enhanced, which may constitute an impor-

tant part of the observable. Therefore, it cannot be regarded as a validation of a complete QCD perturbative expansion, but only of a part of it (the logarithmically enhanced contributions) which may in general behave differently from the full result. This limitation can be minimized by considering threshold resummation of an inclusive observable. Indeed, threshold logarithms reproduce a number of structures that are present in the full result (e.g., plus distributions), and indeed they often offer a good approximation to the full result [56]. Additionally, to all orders threshold contributions diverge factorially [57], as the full expansion does. Therefore, a threshold approximation can be regarded as a good representative of a full QCD perturbative expansion.

The process we consider is once again Higgs production in gluon fusion. Its threshold resummation is known at a high logarithmic accuracy, $N^3\text{LL}$ (see e.g. Ref. [58]), and it is well known to provide a good approximation of the full result order by order (see e.g. Ref. [5]). More precisely, we take the so-called ψ -soft₂ resummation [58], with the so-called default choice for the constant terms [5]. Using the same setup of Sect. 2.5, at the scale $\mu = m_H$ (and $\mu_F = m_H/2$) the partial sums of the perturbative expansion of the cross section are given by³¹

$$\Sigma_{\text{pert}}(m_H) = \{13.0, 28.0, 37.8, 42.4\} \text{ pb}. \quad (7.5)$$

These numbers are very similar to the exact ones, Eq. (2.29): they are slightly smaller, but given the large size of the perturbative corrections, this approximation captures well the

³¹ These numbers have been computed using the TROLL code [5, 58].

Fig. 33 Scale dependence of the purely resummed expanded Higgs cross section (settings as in Fig. 4). The thin solid black line represents the “exact” result

behaviour of the expansion. The “exact” result, namely the all-order resummed one, is instead given by

$$\Sigma = 45.0 \text{ pb}. \quad (7.6)$$

Note that this result is in fact scale dependent, as it is not really exact. The scale dependence is induced by terms beyond $N^3\text{LL}$ and beyond the threshold approximation. In practice, this dependence is rather mild, and certainly much milder than the dependence of the fixed-order result. Therefore, to a first approximation, we simply ignore it and use the value of Eq. (7.6), which is computed at $\mu = m_H/2$ where the convergence is faster, as if it were the exact one.

We start by showing the scale dependence of the various orders in this example in Fig. 33. We observe that this plot looks very similar to that of Fig. 4, as expected from the goodness of the threshold approximation. Therefore, the probability distributions that we get from the models are also very similar, and we thus do not report them here. We also observe that the “exact” result is rather close to the plateau of the $N^3\text{LO}$ result.

In Fig. 34 we show the summaries of the uncertainties using the usual quantifiers. We see that all results (with the exception of the scale variation model at LO) are well compatible with the “exact” result, especially after marginalizing over the scale. These plots are very similar to those of Figs. 11, 16, 21 and 24, and thus provide a strong validation of those results. We can also see that the canonical scale variation “error” is acceptable at NNLO and $N^3\text{LO}$, but it underestimates the uncertainty at NLO and strongly at LO, and this even at the “optimal” scale $\mu = m_H/2$ used in the plots: the pattern may become much worse at different (e.g. higher) scales.

7.5 e^+e^- into hadrons at $N^4\text{LO}$

We now consider some applications of our methods to real physical processes of interest for which the exact result is not known. We start from the classical process of $e^+e^- \rightarrow$ hadrons, which is known in QCD up to $N^4\text{LO}$ [59]. We define the observable Σ as the ratio of the function $R(s)$ describing the process³² to its LO $R_0(s)$,

$$\Sigma = \frac{R(s)}{R_0(s)} = 1 + \frac{\alpha_s(s)}{\pi} + 1.40923 \left(\frac{\alpha_s(s)}{\pi} \right)^2 - 12.805 \left(\frac{\alpha_s(s)}{\pi} \right)^3 - 80.434 \left(\frac{\alpha_s(s)}{\pi} \right)^4 + \mathcal{O}(\alpha_s^5) \quad (7.7)$$

where s is the center of mass energy squared of the collision, the renormalization scale is $\mu = \sqrt{s}$, and we have assumed $n_f = 5$ active flavours for the computation of the numerical coefficients of the expansion [59]. We choose the collider energy to be $\sqrt{s} = 10$ GeV, so that the bottom quark is active ($n_f = 5$) and the scale is sufficiently large to be in the perturbative regime, and at the same time sufficiently small to neglect the contribution from a virtual Z boson.

The “raw” result of this process, shown as a function of the renormalization scale, is depicted in Fig. 35 (left). With the exception of the LO, which is scale independent (it corresponds to the lower edge in the plot), we observe that increasing the order the scale dependence flattens out. At $N^4\text{LO}$, the result is very stable upon scale variation. We also observe an apparently convergent pattern at all scales except the smallest ones, where of course the larger value of α_s makes the perturbative expansion badly behaved.

We now apply our probabilistic methods to this expansion. We consider both the geometric behaviour model and the scale variation model, both at fixed scale $\mu = \sqrt{s}$ and after marginalizing over the scale, using a flat prior in the range $\sqrt{s}/3 < \mu < 10\sqrt{s}$. The distributions and their summary in terms of the usual quantifiers are shown in Fig. 36. Let us start from the geometric behaviour model. At fixed scale, the distributions from NNLO onwards are quite narrow and localized, giving rise to precise predictions. However, at least at NNLO, the prediction is not very accurate, as the next orders tend to favour values towards the tail of the NNLO distribution. Indeed, the 68% DoB region is very small, and the next orders are compatible only within the 95% DoB interval. The fact that the precision of the NNLO distribution is not faithful is also seen from the large standard deviation of the distribution, due to the slowly decreasing tails. We observe that after marginalizing over the scale

³² In fact, $R(s)$ is the ratio of the cross section for $e^+e^- \rightarrow$ hadrons to that of $e^+e^- \rightarrow \mu^+\mu^-$, but since we normalize to the LO, the denominator is immaterial.

Fig. 34 The geometric behaviour model (upper plots) and scale variation model (lower plots) at fixed scale $\mu = m_H/2$ (left plots) and after marginalizing over the scale (right plots) for the resummed Higgs pro-

duction cross section at N^3 LL. The thin solid black lines represent the all-order resummed result, used as “exact” result in this example

the pattern improves significantly. The NNLO distribution becomes very asymmetric, covering with higher probability a region of smaller cross section. Indeed, in this case the 68% DoB interval is larger and covers the next orders. We see once again that marginalizing over the scale, on top of removing a bias, leads to more accurate results. A peculiar feature of the scale independent result is the bimodal distribution at N^4 LO. This is a consequence of the scale dependence, shown in Fig. 35 (left), where we see that there are two scales for which the N^4 LO correction vanishes. One value, approximately $\mu \sim 3\sqrt{s}$, leads to a larger cross section which is

also close to the NNLO result, which makes this scale more favourable, thus producing the higher peak of the distribution. The other value, approximately $\mu \sim 0.6\sqrt{s}$, leads to a smaller cross section and produces the secondary peak. This value is less favoured because the N^3 LO is larger here, but at a slightly higher scale all corrections are rather small, and in particular the NNLO correction also vanishes. This implies that the posterior distribution for μ , shown in Fig. 35 (right), is bimodal itself, with the secondary peak in correspondence of the peak of the posterior at NNLO. Note that in the end the fact that the distribution for Σ is bimodal is not a problem,

Fig. 35 Left: scale dependence of the $e^+e^- \rightarrow \text{hadrons}$ process at $\sqrt{s} = 10$ GeV. Right: posterior distribution for the scale for the geometric behaviour model

and indeed the uncertainty shown through the quantifiers is rather small (though slightly asymmetric).

The scale variation model gives, as usual, larger distributions and larger uncertainties, though the tails are exponentially suppressed and therefore the “support” of the distribution is more localized. The accuracy of the results is good, with a mild exception for the NNLO result at fixed scale, that covers the next orders only within the 95% interval. As expected, this is improved after marginalizing over the scale, which in turn also gives more peaked distributions at higher orders. The uncertainty estimate of the scale independent results is certainly more conservative than that obtained with the geometric behaviour model, and consequently very reliable but less precise.

7.6 Higgs decay to gg at $N^4\text{LO}$

Other QCD observables known at $N^4\text{LO}$ are the decay width of the Higgs boson into a $b\bar{b}$ pair [60,61] and into gluons³³ [61]. The former shares various similarities with the $e^+e^- \rightarrow \text{hadrons}$ process discussed in Sect. 7.5, and the application of our methods leads to very similar results. Therefore, we do not consider it and just focus on the decay into gluons. We define the observable Σ as the decay width

$$\Sigma = \Gamma_{H \rightarrow gg}(\mu), \quad (7.8)$$

which, for $\mu = m_H$, leads to the expansion [61]

$$\Sigma = \Gamma_0(m_H) \left[1 + 5.70305\alpha_s(m_H) + 15.5120\alpha_s^2(m_H) + 12.666\alpha_s^3(m_H) + 69.329\alpha_s^4(m_H) + \mathcal{O}(\alpha_s^5) \right], \quad (7.9)$$

³³ In the heavy top effective theory where the interaction vertex between the Higgs and the gluons is pontlike.

where the numerical coefficients have been computed with $n_f = 5$, as appropriate, and the LO at the generic scale μ is given by

$$\Gamma_0(\mu) = \alpha_s^2(\mu) \frac{G_F m_H^3}{36\sqrt{2}\pi^3} = \alpha_s^2(\mu) \times 14.43 \text{ MeV}. \quad (7.10)$$

Differently from the previous example (and from the $H \rightarrow b\bar{b}$ decay) the LO now is scale dependent. When computing this result at another scale, we must then account for the change at LO, using the procedure described in Sect. A.1.

In Fig. 37 (left) we show the “raw” results as a function of μ . We note that the behaviour of this plot is similar to that of Higgs production in gluon fusion, Fig. 4, due to the fact that the two processes are obviously related, and in particular they both start at $\mathcal{O}(\alpha_s^2)$. In this case, one extra order is known, thus providing an interesting and useful validation. Note however that the perturbative correction are somewhat smaller in this case, leading to a better convergence pattern.

We move immediately to the results of our methods, shown in Fig. 38, in the same format as Fig. 36. The scale $\mu = m_H$ used in the plots is a good one, because the perturbative expansion at this scale is very benign. Therefore, the probability distributions at fixed scale are rather good and converge well, with the next order always included within the 68% DoB interval with no exceptions. Moreover, the precision at $N^4\text{LO}$ of both models is comparable.

When we marginalize over the scale, using a flat prior for μ in the range $m_H/8 < \mu < 8m_H$, we obtain slightly larger uncertainties, with improved convergence (in particular the mean of the distributions at lower orders are more in line with the higher order results). We observe that the geometric behaviour model generates again multimodal distributions, this time both at $N^3\text{LO}$ and $N^4\text{LO}$. The explanation is still the same, namely that there is more than one region of scales for which the observable converges well, thus producing multi-

Fig. 36 Fixed scale and scale independent results for the $e^+e^- \rightarrow \text{hadrons}$ process, for both the geometric behaviour model (upper plots) and the scale variation model (lower plots)

Fig. 37 Left: scale dependence of the $H \rightarrow gg$ decay width. Right: posterior distribution for the scale for the geometric behaviour model

modal posterior distributions for the scale, as shown in Fig. 37 (right). The marginalization over the scale produces a prediction that is slightly smaller (though compatible) with the one at $\mu = m_H$, which is an obvious consequence of the fact that for $\mu \sim m_H$ both the $N^3\text{LO}$ and the $N^4\text{LO}$ have a maximum (see Fig. 37, left).

These results allow us to conclude that the estimate of the uncertainties using our methods are robust and reliable. Given the similarity of this process to the case of Higgs production, we can extrapolate that the methods are reliable also in that case.

7.7 A resummed logarithmic-ordered expansion

So far we have applied our methods to fixed-order perturbative expansions. However, we have stressed several times that our methods can work also for more general expansions, provided they behave in a perturbative way. In particular, one can consider the results obtained using all-order resummations, and consider the sequence of increasing logarithmic accuracy, namely LL, NLL, NNLL and so on. Such an expansion is supposed to behave perturbatively, with a structure given by (we consider for simplicity the case of a single-logarithmic enhancement)

$$\Sigma_{\text{pert}} = \sum_{k=0}^{\infty} g_k(\alpha_s L) \alpha_s^k, \quad (7.11)$$

where the g_k coefficients are all-order functions of their argument, L being the resummed logarithm. The resummation assumes $\alpha_s L \sim 1$, so the g_k coefficients are formally coefficients of $\mathcal{O}(1)$, and the explicit power of α_s determines a fully fledged perturbative expansion.

In this section, we consider threshold resummation applied to the Higgs production process, which is known to a high order, $N^3\text{LL}$ [58]. We have already used this result in Sect. 7.4

to construct a fixed-order perturbative expansion by expanding the resummed result. Here, instead, we take the full resummed result at various logarithmic orders, matched (as appropriate for phenomenology) to the corresponding fixed-order result. Namely, the perturbative expansion corresponds to the sequence LO+LL, NLO+NLL, NNLO+NNLL, $N^3\text{LO} + N^3\text{LL}$, whose values at $\mu = m_H$ (and $\mu_F = m_H/2$) are given by

$$\Sigma_{\text{pert}}(m_H) = \{15.9, 38.7, 47.1, 48.5\} \text{ pb} \quad (7.12)$$

(we used the same setting as Sect. 7.4, namely the default of Ref. [5]). Since the coefficients of the perturbative expansion are all-order functions of α_s , their scale dependence cannot be reconstructed using the procedure of Sect. A.1. Therefore, to simplify the treatment of this process, we just consider the geometric behaviour model at fixed scale.³⁴ To investigate the dependence of the results on the scale, we consider also the expansion at another scale, $\mu = m_H/2$, leading to the sequence

$$\Sigma_{\text{pert}}(m_H/2) = \{20.1, 46.2, 50.1, 48.6\} \text{ pb}. \quad (7.13)$$

Note that at this scale the value of the coupling is different, $\alpha_s(m_H/2) = 0.1252$.

We show in Fig. 39 the results of the geometric behaviour model applied at each of the scales considered. Because the resummation speeds up the (apparent) perturbative convergence [5], these results are somewhat more precise (smaller uncertainty) than those coming from the fixed-order expansion. Nonetheless, they remain reliable, as one can appreciate from the fact that the 68% DoB interval always covers the

³⁴ Of course, it is possible to retrieve the correct scale dependence of the resummed result numerically, e.g. from the TROLL code [5,58]. This would however slow down the computation of the uncertainty in THunc.

Fig. 38 Fixed scale and scale independent results for the $H \rightarrow gg$ decay width, for both the geometric behaviour model (upper plots) and the scale variation model (lower plots)

Fig. 39 Probability distributions and their summaries for the Higgs production process at resummed level matched to fixed order, for fixed scale $\mu = m_H$ (upper plots) and $\mu = m_H/2$ (lower plots) using the geometric behaviour model

next order. These results are also fully compatible with those coming from the fixed-order perturbative expansion. A thorough study of the uncertainties of this (or any other) process is beyond the scope of this work and will be done elsewhere.

8 Treating correlations between different theory predictions

The approach and models presented so far can be used to compute the uncertainty on a given perturbative expansion, namely for the prediction of a given observable. In many cases, one needs to consider more than one observable at a time in a physics analysis. In these cases, correlations between the uncertainties of different observables are fundamental.

We stress that when we say different observables we actually also mean the same physical observable at different kinematic points. Consider for example a differential observable, like a transverse momentum distribution. The methods presented so far would work independently for any value of the transverse momentum p_t , as each value of p_t leads to a different perturbative expansion, in principle with very different perturbative behaviours.³⁵ It's clear that in such cases there are strong correlations between different values of p_t , approaching 100% for close values of p_t . But there may be also important correlations for p_t values very far from each other, induced by the knowledge of the full cross section

³⁵ Consider for example the region of small p_t where the presence of logarithms of p_t invalidate the perturbative expansion and requires the all-order resummation of these logarithms.

(integral of the distribution). Similarly, when one consider different distributions for the same physical process, there likely are correlations on the uncertainty from missing higher orders, e.g. due to the fact that all distributions must integrate to the same total cross section.

The situation is somewhat different when considering different processes. Here it is less clear whether the uncertainties should be correlated or not. Note that there is a class of uncertainties, due to the uncertain values of parameters such as couplings or masses, that are clearly (and easily) correlated between different observables of different processes. However, here we are dealing with the uncertainty due to missing higher orders. The presence of correlations in this case implies that the perturbative expansions of the different observables of different processes are somehow related. This is not to be excluded a priori, as indeed perturbative expansions often show some similarities.³⁶ However, quantifying such similarities is hard in general and subject to a large degree of arbitrariness.

Finding a general and satisfactory solution to all these cases is complicated (if not impossible) and it will not be subject of this work. In the following, we limit ourselves to discuss some simple cases and propose some strategies. This section has just to be considered as a starting point for future work in this direction.

8.1 Correlations for the same observable at different kinematic points

Let us start our discussion from the case where there certainly are correlations, namely the case of the same observable at different kinematic points. Let us call by $\vec{v}$ and $\vec{w}$ generic configurations of kinematic variables, and by $\Sigma(\vec{v})$ and $\Sigma(\vec{w})$ the observable that depends on those variables. The information on the correlation between $\Sigma(\vec{v})$ and $\Sigma(\vec{w})$ is contained in the joint distribution

$$P(\Sigma(\vec{v}), \Sigma(\vec{w}) | \delta_n(\vec{v}), \dots, \delta_1(\vec{v}), \Sigma_0(\vec{v}), \delta_n(\vec{w}), \dots, \delta_1(\vec{w}), \Sigma_0(\vec{w})), \quad (8.1)$$

where each term of the perturbative expansion is computer either at $\vec{v}$ or at $\vec{w}$. We recall that, in the limit $\vec{w} \rightarrow \vec{v}$, the distribution must lead to a positive 100% correlation.

Constructing such joint distribution satisfying the aforementioned requirement is far from trivial. The most promising way that we can imagine consists in considering not the values of the observable at given kinematic points, but the observable as a *function* of the kinematics. The probability to be computed is that of the function, and the requirement that the correlation of adjacent points tends to one is automatic

³⁶ This is for instance the case when there is a common dynamical mechanism giving the bulk of the higher order corrections, as it happens in kinematic regions with logarithmic enhancements.

provided obvious conditions of continuity and smoothness are included. Also, the condition that the integral of the distribution is the inclusive cross section is easy to implement. Note that, in general, with this approach the probability distribution of a single point would not necessarily correspond with that obtained from the application of the methods introduced in the previous sections. In particular, one could expect to obtain a smaller uncertainty with this approach, given the additional constraints included.

How to realize in practice this approach is again subject of arbitrary choices. A promising option consists in using the proposal of Ref. [62]. The idea is simple. There are classes of higher order contributions whose functional structure is known to all orders from resummation techniques. These do not capture all the possible missing higher orders, but in some kinematic regions they are quite accurate, and they can be constructed such that in other regions they anyway represent to some accuracy the full result. In this way, the only missing information for constructing the next orders is the value of some numerical coefficients which the functional structure of the resummation depends upon. In other words, rather than using a generic set of functions to model the observable one restricts the attention to a physically motivated subset parametrized by a small number of real parameters. In this way the inference is to be done on the parameters themselves. Moreover, these parameters are typically power series in the coupling, and therefore one can use straight away the probabilistic methods introduced in this work.³⁷ The actual implementation of this approach will be studied elsewhere.

8.2 Correlations among observables of different processes

We now move to discussing the correlations among observables belonging to different physical processes. In this case, it is not clear at all that the uncertainty due to missing higher orders should be correlated: the computations are different, the number of known orders is in general different, the perturbative series describing the various observables are different. Therefore, at first sight, one is tempted to say that the uncertainty from missing higher orders between observables of different process are uncorrelated.

This conclusion is likely correct for very different processes. However, some processes may be characterized by perturbative expansions with common features, for instance when the underlying physical mechanism at the origin of

³⁷ More precisely, since these parameters are not physical observables, they are not scale independent to all orders and therefore the scale variation model of Sect. 5 cannot be used. One can instead use the geometric behaviour model of Sect. 4. The removal of scale dependence of Sect. 6 cannot be applied to the parameters themselves, again because they are not physical and they do legitimately depend on unphysical scales, but it can be applied at the end to the physical observable, which has to be scale independent.

some kind of corrections is the same. This is for instance the case when the bulk of perturbative corrections is given by logarithmically enhanced contributions, which in turn originate from a specific structure of particle emissions. Similarly, processes like Z and W production at proton-proton colliders (or equivalently neutral-current and charged-current structure functions in DIS experiments) are very similar, as they share the same structure of the interaction and thus the form of the perturbative corrections. In these cases it makes sense to assume that correlations between the uncertainties from missing higher orders of the given processes are present.

The tough question is how to quantify such correlations. To keep the discussion very general, we may assume that *some* of the parameters characterizing the probabilistic model are common to two given observables Σ and Σ' of different processes. If this is the case, the joint distribution can be written as

$$\begin{aligned} P(\Sigma, \Sigma' | \delta_n, \dots, \delta_1, \Sigma_0, \delta'_n, \dots, \delta'_1, \Sigma'_0) \\ \propto P(\Sigma, \Sigma', \delta_n, \dots, \delta_1, \Sigma_0, \delta'_n, \dots, \delta'_1, \Sigma'_0) \\ = \int d\vec{p} P(\Sigma, \Sigma', \delta_n, \dots, \delta_1, \Sigma_0, \delta'_n, \dots, \delta'_1, \Sigma'_0, \vec{p}) \\ = \int d\vec{p} P(\Sigma | \delta_n, \dots, \delta_1, \Sigma_0, \vec{p}) P(\Sigma' | \delta'_n, \dots, \delta'_1, \Sigma'_0, \vec{p}) \\ \times P(\delta_n, \dots, \delta_1, \Sigma_0 | \vec{p}) P(\delta'_n, \dots, \delta'_1, \Sigma'_0 | \vec{p}) P_0(\vec{p}) \end{aligned} \quad (8.2)$$

where $\vec{p}$ represents the vector of parameters of the model that are shared between the observables. Of course the model may depend on additional parameters, but these are assumed to be independent for the two observables, and therefore do not enter our decomposition. Equation (8.2) depends on the probability of each observable given its respective perturbative expansion *and* the common parameters $\vec{p}$, and similarly on the probability of the perturbative expansions themselves given $\vec{p}$. The actual form of these probability distributions depend on what is $\vec{p}$, but once the model is specified and the set of common parameters identified, then Eq. (8.2) can be used directly, as it depends on ingredients introduced in the previous sections.

The choice of the common parameters $\vec{p}$ depends on the processes under consideration and it is the trickiest part of the job. For instance, one model-independent parameter that could be used to construct such correlations is the unphysical scale μ , or better its ratio to the hard scale of the process. In principle, the scale dependence of each process is independent, however it is induced in all cases by the dependence of the coupling on the scale. For this reason, one can assume that the variation of the observable induced by a variation of the scale is somehow related, and use it to construct the sought correlation. This reasoning may however lead to some problems, as for instance obtaining such correlations

requires using the same prior $P_0(\mu/Q)$, with Q the hard scale of each process, which makes impossible to adjust this prior to process-dependent conditions.³⁸ In agreement with this conclusion, in recent work [63, 64] on the inclusion of theory uncertainties in PDF fits, the dependence on the renormalization scale is not considered as a source of correlation among different *classes* of processes, while it is considered for similar processes like Z and W production belonging to the same class (Drell–Yan).

An alternative way to construct correlations between different but somehow related processes consists in using again the approach of Ref. [62], as we described in Sect. 8.1. Indeed, some of the parameters used in the construction of the resummed results at the basis of that approach are process dependent, but some others are more universal, and common among different processes. For instance, the cusp anomalous dimension is one of the key ingredients of soft-gluon resummations for any process. Therefore, in Eq. (8.2) one can take $\vec{p}$ to be the model parameters describing the perturbative expansion of the common coefficients (such as the cusp anomalous dimension). The probability distribution would be constructed exactly as one would do for Sect. 8.1. We conclude that the approach of Ref. [62] used in conjunction with our methods is very promising for estimating theory uncertainties from missing higher orders in a way that produces meaningful correlations between different observables, and deserves further studies.

We conclude by mentioning that the situation becomes more complicated for hadron-initiated processes, due to the presence of parton distribution functions (PDFs). Indeed, PDFs induce two additional problems. One is that they depend on a different unphysical scale, the factorization scale μ_F . If the PDFs were known exactly, this would not be too problematic, as one could extend the approach of Sect. 6 to μ_F as well. However, PDFs are non-perturbative objects and cannot be computed from the theory. Therefore, they are determined by a fit to data, using a theoretical computation for the perturbative coefficients describing an observable as input. This means that PDFs indirectly depend on the accuracy of the computation of the perturbative coefficients, which is the second problem. Ideally, one would like to include theory uncertainties with proper correlations in the determination of the PDFs themselves, as studied recently in Refs. [63, 64] (using canonical scale variation). The trouble is that the μ_F dependence is thus entangled with the theory uncertainty, and it is not obvious how to properly deal with it (see e.g. Ref. [65]). This important topic deserves a dedicated study.

³⁸ For instance, if the hard scale Q of one of the (QCD) processes under consideration is small, one does not want to let the prior probe values much smaller than Q , otherwise the observable would get contribution with too large strong coupling.

9 Conclusions

Having a reliable estimate of the uncertainty on theoretical predictions is of fundamental importance for (particle) physics. A particularly important source of uncertainty is the one coming from the finite perturbative accuracy of theoretical computations. Quantifying the uncertainty due to unknown (or missing) higher orders is crucial, especially for those theories like QCD that are characterized by large perturbative corrections. In these cases the missing higher orders may be sizeable, and underestimating them may lead to wrong conclusions when comparing theoretical predictions with precise data.

The standard practice to estimate the size of the missing higher orders is based on the “canonical scale variation” method, where the unphysical scale(s) is varied by a factor of 2 about an arbitrary central scale. Since scale dependence is formally subleading, the effect of the variation is formally higher order. The problem of this approach appears when one tries to promote this information to an uncertainty. Not only the procedure is totally arbitrary and has no probabilistic foundation; it also often underestimates the true uncertainty, and it is thus not reliable.

A groundbreaking approach to define and compute the uncertainty from missing higher orders has been introduced by Cacciari and Houdeau in 2011 [3]. This method uses a Bayesian approach to infer the probability distribution for an observable from its perturbative expansion, under some conditions that define the Cacciari–Houdeau (CH) model. Later, the CH model has been modified to account for the divergence of the perturbative expansion [4]. While this approach is formally superior to canonical scale variation, it also has limitation. Indeed, in some cases the CH model is not very robust and it does not predict reliably the uncertainty from the missing higher orders. This behaviour is due to the assumptions of the model itself, and not to the approach per se.

In this work we built upon the CH model to construct more general, flexible and reliable models, using the same Bayesian approach. Moreover, we have introduced an innovative way to eliminate an ambiguity of theoretical computations in quantum field theory, namely the dependence on unphysical scale(s). We have validated the models using perturbative expansions with a known sum, and tested their quality on some observables that are known to a high perturbative order.

The models proposed in this work are essentially two, with a number of possible variants discussed in the appendices. One model is a generalization of the CH model, and assumes that the various orders k of the perturbative expansion are bounded by a geometric behaviour ca^k , where both c and a are (hidden) model parameters. The main difference with respect to the CH model is that a , that plays the role of the expansion parameter of the geometric bound, in the

CH approach is fixed to be the coupling constant, possibly up to an externally supplied fixed factor. Promoting a to a model parameter allows to select its most appropriate values based only on information from the expansion itself, through a probabilistically valid inference procedure. This improvement alone already makes the model much more reliable and robust. In practical applications, it performs very well, providing precise and accurate uncertainties. These are also pretty stable upon variations of the assumptions, as the bulk of the distributions is rather insensitive to the precise form of the prior, which mostly affects the tails.

The other model uses the information from the scale variation, similarly to the canonical scale variation method, but in a probabilistically well defined way. In particular, it assumes that the next order is bounded by a factor λ times a properly defined scale variation coefficient of the current order. The factor λ is a hidden parameter of the model and it is inferred from the known orders. Oversimplifying, this model can be interpreted as the canonical scale variation method where the size of the variation, rather than being fixed to a factor of 2, is variable and inferred from the expansion itself. We have noticed that this model predicts somewhat larger uncertainties than the geometric behaviour model, thus providing a more conservative estimate of the theory uncertainty. This uncertainty is very stable upon variation of the prior distribution.

There are a number of improvements that can be applied to the models, discussed in the appendices. Two of them are particularly interesting. One suggests to impose an additional condition on the scale dependence of the inferred higher orders: indeed, not only the various perturbative contributions should get smaller and smaller by increasing the order, but also their scale dependence should, as it is always formally of the next order and must thus behave perturbatively. Requesting a reduction of the scale dependence provides strong constraints on the missing higher orders, resulting in smaller uncertainties (i.e., higher precision). Another improvement consists in looking for a sign pattern in the various perturbative orders, to infer the next orders according to the same pattern. Also this procedure allows to reduce the uncertainty, producing more precise results. Moreover, different conditions can be combined together to provide more stringent models that can further improve the precision.

Which model to choose is a duty of the user. A “perfect” or “correct” model does not exist, as the only way to be sure about the size of the missing higher orders is computing them, and any estimate is subject to arbitrary assumptions and biases.³⁹ The Bayesian approach highlights the presence of such subjective assumptions, making them very explicit (the model, the priors). Users, based on their beliefs about the

³⁹ Conversely, “bad” or “wrong” models do exist, and can be discarded based e.g. on an unsuccessful validation.

expansion, can freely choose the model and priors, declaring them. This said, we now give a *personal* preference. From the studies presented in this work, it is clear that the geometric behaviour model provides precise results without sacrificing accuracy. Given the simplicity of the model and the fast implementation provided, we suggest its use as default. Moreover, this is the only model that is applicable beyond a quantum field theory, and it is thus more general and universal. Since the plain scale variation model is simple and fast, for QCD observables we also recommend its application as a “double check”, consistently with its conservative interpretation. Should one be interested in providing more aggressive results, the best performance is obtained by combining all the proposed assumptions (Eqs. 4.2, 5.2, and B.4), and keeping track of the sign pattern as discussed in Appendix B.5. Such an approach makes the most of the known orders, and it is thus particularly useful when only a low number of orders is available. However, we suggest its use only for very specific and dedicated studies and not as a general-purpose approach, also due to its slow numerical implementation. In all cases, variations of the priors should be considered to assess the stability of the results upon changes of the subjective arbitrary input.

Independently of the model used, we have proposed a general way to remove the dependence on the renormalization scale from the probability distribution of the observable (the method can be extended also to other unphysical scales). Indeed, each model uses as basic elements the coefficients of the perturbative expansion, which however depend on the unphysical scale and would thus lead to infinitely many different results depending on the choice of the scale. The solution is actually very simple: promoting the unphysical scale to be a parameter of the model. The inference works exactly as in the general case (Eqs. 1.2 and 1.3), with the difference that among the hidden parameters there is also the scale, with its own prior. In this way, after marginalizing over the scale, the dependence disappears. The original arbitrariness in the choice of scale is almost completely removed – only a mild arbitrariness remains in the choice of the prior. It has to be stressed that the inference on the scale acts differently depending on the model, selecting those values of the scale for which the given model performs better. In this way, the most accurate and precise results are obtained for free, without the need of choosing wisely the scale corresponding to a “good” perturbative expansion. We also stress that with this procedure the “best” prediction for an observable is given by the mean of its probability distribution, in contrast with standard methods, where it corresponds to the value of the perturbative expansion computed at the highest known order at an arbitrary “central” scale. Therefore not only the method provides a reliable uncertainty, but also a better central value for the observable. For scale dependent observables, we always recommend the use of this procedure.

We have finally considered the problem of quantifying the correlations between various theoretical predictions and their uncertainties. We believe that this argument cannot be addressed in full generality, but specific solutions need to be put in place depending on the observables and the processes under considerations. Some proposals have been suggested, but a thorough study is delayed to future work.

The results of this paper can be used to compute the uncertainty on any physical observable through a publicly released code named THunc, available at

<http://www.roma1.infn.it/bonvini/THunc>

The two main models proposed in this work are implemented analytically and lead to a very fast evaluation of the uncertainties. Other models (including those proposed in the appendices or alternatives invented by the user) can be implemented straight away through a custom model feature, at the price of a slower numerical evaluation.

Acknowledgements I am grateful to Enrico Franco, Giulio D’Agostini and Sara Borroni for very useful discussions on probability theory related to this work. I also thank Luca Rottoli and Simone Marzani for useful comments to the manuscript, and Franz Herzog, Juan Rojo and Stefano Forte for discussions. This work was supported by the Marie Skłodowska Curie grant HiPPiE@LHC under the agreement n. 746159.

Data Availability Statement This manuscript has no associated data or the data will not be deposited. [Authors’ comment: The work is of theoretical nature.]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Funded by SCOAP³.

A Algorithms

A.1 Computing renormalization scale dependence at fixed order

In this appendix we show how to reconstruct the renormalization scale dependence of physical observables starting from the perturbative expansion at a given scale μ_0 . We consider the case of QCD for definiteness, but the results can be easily extended to other quantum field theories.

Let us assume we know the expansion of an observable as is given by

$$\Sigma_n(\mu_0) = \sum_{k=0}^n \bar{c}_k \alpha_0^{k+k_0}, \quad (\text{A.1})$$

where $\bar{c}_k \equiv c_k(\mu_0)$ are the (known) coefficients of the expansion computed at the scale $\mu = \mu_0$ up to fixed order $k = n$, and $\alpha_0 = \alpha_s(\mu_0)$ is the value of the strong coupling at this scale, assumed to be known as well. We also included in the expansion an offset k_0 to account for expansions that start at $\mathcal{O}(\alpha_s^{k_0})$. We now wish to rewrite the perturbative expansion at a different scale μ

$$\Sigma_n(\mu) = \sum_{k=0}^n c_k(\mu) \alpha_s^{k+k_0}(\mu), \quad (\text{A.2})$$

which differs from $\Sigma_n(\mu_0)$ (Eq. A.1), by subleading terms, because Σ is scale independent (note that this is true only if Σ is a physical observable).

Because the scale dependence of the coupling is known, one can reconstruct the scale dependence of the expansion coefficients by imposing the scale independence of Σ . Using the definition of the β -function from the Callan-Symanzik equation

$$\mu^2 \frac{d}{d\mu^2} \alpha_s(\mu) = \beta(\alpha_s(\mu)) = -\alpha_s^2(\mu) \left[\beta_0 + \beta_1 \alpha_s(\mu) + \beta_2 \alpha_s^2(\mu) + \beta_3 \alpha_s^3(\mu) + \beta_4 \alpha_s^4(\mu) + \dots \right] \quad (\text{A.3})$$

one can obtain a system of differential equations for the coefficients by imposing

$$0 = \mu^2 \frac{d}{d\mu^2} \Sigma = \sum_{k=0}^n \left[\mu^2 \frac{d}{d\mu^2} c_k(\mu) \alpha_s^{k+k_0}(\mu) + c_k(\mu) (k+k_0) \alpha_s^{k+k_0-1}(\mu) \beta(\alpha_s(\mu)) \right] + \mathcal{O}(\alpha_s^{n+1+k_0}). \quad (\text{A.4})$$

Expanding the term in square brackets in powers of α_s , the system can be solved iteratively order by order, using as initial condition the values $\bar{c}_k$ of the coefficients at μ_0 .

This way to proceed, though perfectly valid, is not very convenient in practice. It is indeed more efficient to consider directly the solution of the Callan-Symanzik equation to express $\alpha_0 \equiv \alpha_s(\mu_0)$ in terms of $\alpha_s(\mu)$, plugging this into Eq. (A.1), and finally expanding in powers of $\alpha_s(\mu)$ and read off the resulting coefficients $c_k(\mu)$ by comparison with Eq. (A.2). The solution of the running coupling equation up to five loops [53] and expanded in powers of α_s as appropriate

$$\begin{aligned} \alpha_0 = \alpha_s(\mu) & \left[1 + \alpha_s(\mu) \beta_0 \ell + \alpha_s^2(\mu) \left(\beta_0^2 \ell^2 + \beta_1 \ell \right) \right. \\ & + \alpha_s^3(\mu) \left(\beta_0^3 \ell^3 + \frac{5}{2} \beta_1 \beta_0 \ell^2 + \beta_2 \ell \right) \\ & + \alpha_s^4(\mu) \left(\beta_0^4 \ell^4 + \frac{13}{3} \beta_0^2 \beta_1 \ell^3 \right. \\ & + \left(3 \beta_0 \beta_2 + \frac{3}{2} \beta_1^2 \right) \ell^2 + \beta_3 \ell \left. \right) \\ & + \alpha_s^5(\mu) \left(\beta_0^5 \ell^5 + \frac{77}{12} \beta_0^3 \beta_1 \ell^4 \right. \\ & + \left(6 \beta_0^2 \beta_2 + \frac{35}{6} \beta_0 \beta_1^2 \right) \ell^3 \\ & + \left. \frac{7}{2} (\beta_0 \beta_3 + \beta_1 \beta_2) \ell^2 + \beta_4 \ell \right] + \mathcal{O}(\alpha_s^6), \end{aligned} \quad (\text{A.5})$$

with

$$\ell = \log \frac{\mu^2}{\mu_0^2}. \quad (\text{A.6})$$

Obtaining the sought result is now a mere exercise of replacement and expansion. For completeness, we report here the results for the first few coefficients $c_k(\mu)$ for three relevant values of k_0 , namely $k_0 = 0, 1, 2$. For $k_0 = 0$ (processes starting at $\mathcal{O}(\alpha_s^0)$ at LO) we have

$$\begin{aligned} c_1(\mu) &= \bar{c}_1 \\ c_2(\mu) &= \bar{c}_2 + \beta_0 \bar{c}_1 \ell \\ c_3(\mu) &= \bar{c}_3 + (\beta_1 \bar{c}_1 + 2 \beta_0 \bar{c}_2) \ell + \beta_0^2 \bar{c}_1 \ell^2 \\ c_4(\mu) &= \bar{c}_4 + (\beta_2 \bar{c}_1 + 2 \beta_1 \bar{c}_2 + 3 \beta_0 \bar{c}_3) \ell \\ &+ \left(\frac{5}{2} \beta_0 \beta_1 \bar{c}_1 + 3 \beta_0^2 \bar{c}_2 \right) \ell^2 + \beta_0^3 \bar{c}_1 \ell^3 \\ c_5(\mu) &= \bar{c}_5 + (\beta_3 \bar{c}_1 + 2 \beta_2 \bar{c}_2 + 3 \beta_1 \bar{c}_3 + 4 \beta_0 \bar{c}_4) \ell \\ &+ \left[\left(\frac{3}{2} \beta_1^2 + 3 \beta_0 \beta_2 \right) \bar{c}_1 + 7 \beta_0 \beta_1 \bar{c}_2 + 6 \beta_0^2 \bar{c}_3 \right] \ell^2 \\ &+ \left(\frac{13}{3} \beta_0^2 \beta_1 \bar{c}_1 + 4 \beta_0^3 \bar{c}_2 \right) \ell^3 + \beta_0^4 \bar{c}_1 \ell^4 \\ c_6(\mu) &= \bar{c}_6 + (\beta_4 \bar{c}_1 + 2 \beta_3 \bar{c}_2 + 3 \beta_2 \bar{c}_3 + 4 \beta_1 \bar{c}_4 + 5 \beta_0 \bar{c}_5) \ell \\ &+ \left[\frac{7}{2} (\beta_1 \beta_2 + \beta_0 \beta_3) \bar{c}_1 + (4 \beta_1^2 + 8 \beta_0 \beta_2) \bar{c}_2 \right. \\ &+ \left. \frac{27}{2} \beta_0 \beta_1 \bar{c}_3 + 10 \beta_0^2 \bar{c}_4 \right] \ell^2 \\ &+ \left[\left(\frac{35}{6} \beta_0 \beta_1^2 + 6 \beta_0^2 \beta_2 \right) \bar{c}_1 \right. \end{aligned}$$

$$\begin{aligned}
& + \frac{47}{3} \beta_0^2 \beta_1 \bar{c}_2 + 10 \beta_0^3 \bar{c}_3 \Big] \ell^3 \\
& + \left(\frac{77}{12} \beta_0^3 \beta_1 \bar{c}_1 + 5 \beta_0^4 \bar{c}_2 \right) \ell^4 + \beta_0^5 \bar{c}_1 \ell^5. \quad (\text{A.7})
\end{aligned}$$

For $k_0 = 1$ (processes starting at $\mathcal{O}(\alpha_s)$ at LO) we have

$$\begin{aligned}
c_1(\mu) &= \bar{c}_1 + \beta_0 \bar{c}_0 \ell \\
c_2(\mu) &= \bar{c}_2 + (\beta_1 \bar{c}_0 + 2 \beta_0 \bar{c}_1) \ell + \beta_0^2 \bar{c}_0 \ell^2 \\
c_3(\mu) &= \bar{c}_3 + (\beta_2 \bar{c}_0 + 2 \beta_1 \bar{c}_1 + 3 \beta_0 \bar{c}_2) \ell \\
& + \left(\frac{5}{2} \beta_0 \beta_1 \bar{c}_0 + 3 \beta_0^2 \bar{c}_1 \right) \ell^2 + \beta_0^3 \bar{c}_0 \ell^3 \\
c_4(\mu) &= \bar{c}_4 + (\beta_3 \bar{c}_0 + 2 \beta_2 \bar{c}_1 + 3 \beta_1 \bar{c}_2 + 4 \beta_0 \bar{c}_3) \ell \\
& + \left[\left(\frac{3}{2} \beta_1^2 + 3 \beta_0 \beta_2 \right) \bar{c}_0 + 7 \beta_0 \beta_1 \bar{c}_1 + 6 \beta_0^2 \bar{c}_2 \right] \ell^2 \\
& + \left(\frac{13}{3} \beta_0^2 \beta_1 \bar{c}_0 + 4 \beta_0^3 \bar{c}_1 \right) \ell^3 + \beta_0^4 \bar{c}_0 \ell^4 \\
c_5(\mu) &= \bar{c}_5 + (\beta_4 \bar{c}_0 + 2 \beta_3 \bar{c}_1 + 3 \beta_2 \bar{c}_2 + 4 \beta_1 \bar{c}_3 + 5 \beta_0 \bar{c}_4) \ell \\
& + \left[\frac{7}{2} (\beta_1 \beta_2 + \beta_0 \beta_3) \bar{c}_0 + (4 \beta_1^2 + 8 \beta_0 \beta_2) \bar{c}_1 \right. \\
& + \left. \frac{27}{2} \beta_0 \beta_1 \bar{c}_2 + 10 \beta_0^2 \bar{c}_3 \right] \ell^2 \\
& + \left[\left(\frac{35}{6} \beta_0 \beta_1^2 + 6 \beta_0^2 \beta_2 \right) \bar{c}_0 \right. \\
& + \left. \frac{47}{3} \beta_0^2 \beta_1 \bar{c}_1 + 10 \beta_0^3 \bar{c}_2 \right] \ell^3 \\
& + \left(\frac{77}{12} \beta_0^3 \beta_1 \bar{c}_0 + 5 \beta_0^4 \bar{c}_1 \right) \ell^4 + \beta_0^5 \bar{c}_0 \ell^5. \quad (\text{A.8})
\end{aligned}$$

For $k_0 = 2$ (processes starting at $\mathcal{O}(\alpha_s^2)$ at LO) we have

$$\begin{aligned}
c_1(\mu) &= \bar{c}_1 + 2 \beta_0 \bar{c}_0 \ell \\
c_2(\mu) &= \bar{c}_2 + (2 \beta_1 \bar{c}_0 + 3 \beta_0 \bar{c}_1) \ell + 3 \beta_0^2 \bar{c}_0 \ell^2 \\
c_3(\mu) &= \bar{c}_3 + (2 \beta_2 \bar{c}_0 + 3 \beta_1 \bar{c}_1 + 4 \beta_0 \bar{c}_2) \ell \\
& + (7 \beta_0 \beta_1 \bar{c}_0 + 6 \beta_0^2 \bar{c}_1) \ell^2 + 4 \beta_0^3 \bar{c}_0 \ell^3 \\
c_4(\mu) &= \bar{c}_4 + (2 \beta_3 \bar{c}_0 + 3 \beta_2 \bar{c}_1 + 4 \beta_1 \bar{c}_2 + 5 \beta_0 \bar{c}_3) \ell \\
& + \left[(4 \beta_1^2 + 8 \beta_0 \beta_2) \bar{c}_0 + \frac{27}{2} \beta_0 \beta_1 \bar{c}_1 + 10 \beta_0^2 \bar{c}_2 \right] \ell^2 \\
& + \left(\frac{47}{3} \beta_0^2 \beta_1 \bar{c}_0 + 10 \beta_0^3 \bar{c}_1 \right) \ell^3 + 5 \beta_0^4 \bar{c}_0 \ell^4 \\
c_5(\mu) &= \bar{c}_5 + (2 \beta_4 \bar{c}_0 + 3 \beta_3 \bar{c}_1 + 4 \beta_2 \bar{c}_2 + 5 \beta_1 \bar{c}_3 + 6 \beta_0 \bar{c}_4) \ell \\
& + [9 (\beta_1 \beta_2 + \beta_0 \beta_3) \bar{c}_0 \\
& + \left(\frac{15}{2} \beta_1^2 + 15 \beta_0 \beta_2 \right) \bar{c}_1 + 22 \beta_0 \beta_1 \bar{c}_2 + 15 \beta_0^2 \bar{c}_3] \ell^2 \\
& + \left[\left(\frac{59}{3} \beta_0 \beta_1^2 + 20 \beta_0^2 \beta_2 \right) \bar{c}_0 + 37 \beta_0^2 \beta_1 \bar{c}_1 + 20 \beta_0^3 \bar{c}_2 \right] \ell^3 \\
& + \left(\frac{57}{2} \beta_0^3 \beta_1 \bar{c}_0 + 15 \beta_0^4 \bar{c}_1 \right) \ell^4 + 6 \beta_0^5 \bar{c}_0 \ell^5. \quad (\text{A.9})
\end{aligned}$$

Note that the LO coefficient remains always the same, $c_0 = \bar{c}_0$, as it is always scale independent. In all cases we have written all the coefficients whose scale dependence can be computed from the known terms of the β -function, namely up to N⁶LO for $k_0 = 0$ and up to N⁵LO for $k_0 > 0$.

A.2 Computing the coefficients for the geometric behaviour model

In this appendix we describe an algorithm to compute the coefficients a_k defined in Eq. (4.13), needed for an efficient implementation of the geometric behaviour model. We report here the definition Eq. (4.13),

$$a_{k+1} < a < a_k \quad \Leftrightarrow \quad \max \left[1, \frac{|\delta_1|}{a}, \dots, \frac{|\delta_m|}{a^m} \right] = \frac{|\delta_k|}{a^k}, \quad (\text{A.10})$$

assuming $a_0 \equiv \infty$ and $a_{m+1} \equiv 0$. This definition means that the range $0 \leq a < \infty$ is partitioned in consecutive intervals and, in each interval, the max function selects one of its arguments. Since the arguments of the max function contain powers of a that grow with k , the a_k coefficients are ordered according to

$$\infty \equiv a_0 > a_1 \geq \dots \geq a_k \geq a_{k+1} \geq \dots \geq a_m \geq a_{m+1} \equiv 0. \quad (\text{A.11})$$

In general, it is possible that two or more consecutive a_k 's are identical. This happens when a given term is never the maximum of the list for any value of a . This behaviour complicates the computation of these coefficients.

To construct an efficient algorithm, let us suppose that this does not happen, namely that all a_k are distinct. If this was the case, then a_k simply represents the value of a for which $|\delta_k|/a^k = |\delta_{k-1}|/a^{k-1}$. Then, they would simply be given by ($\delta_0 = 1$)

$$a_k = \frac{|\delta_k|}{|\delta_{k-1}|} \quad k = 1, \dots, m. \quad (\text{A.12})$$

The first step of the algorithm consists in constructing these “guesses” for the a_k . Then, the ordering condition Eq. (A.11) is checked. If it is satisfied, it means that indeed all a_k are distinct and thus the algorithm can stop and return the list. Instead, if we find that for a given k

$$a_k < a_{k+1}, \quad (\text{A.13})$$

it means that $|\delta_k|/a^k$ can never be the output of the max function. This implies that $|\delta_k|/a^k$ can be removed from the argument of the max function, and the check can be performed

on what remains. In practice, one can perform the comparison directly between the two adjacent terms, $|\delta_{k+1}|/a^{k+1}$ and $|\delta_{k-1}|/a^{k-1}$, which leads to the replacement in the list of “guesses”

$$a_k = a_{k+1} = \sqrt{\frac{|\delta_{k+1}|}{|\delta_{k-1}|}}. \quad (\text{A.14})$$

At this point the list of guesses of a_k is modified, and one can perform again the ordering check. If the ordering Eq. (A.11) is satisfied, then the algorithm stops, otherwise it keeps modifying the list until the ordering is satisfied. Note that, if there is a number of consecutive $|\delta_i|/a^i$, $i = k, \dots, k+j-1$ that cannot be the output the max function, then the corresponding coefficients are given by

$$a_k = a_{k+1} = \dots = a_{k+j} = \left(\frac{|\delta_{k+j}|}{|\delta_{k-1}|} \right)^{\frac{1}{1+j}}. \quad (\text{A.15})$$

Despite the iterative nature of this algorithm, given the small number of known orders in practical perturbative expansion, it performs very well and it provides a fast evaluation of the inference in the geometric behavior model (way faster than performing a numerical integration).

B Possible ways of improving the models

In the main body of the paper we focussed on two main models, with simple assumptions that allowed to compute most of the integrals in the inference procedure analytically. As a results the numerical implementation of these models is very fast.

In this appendix, we collect ideas to go beyond these simple models, to improve their reliability and to provide alternatives that may lead to improved precision. Most of them work at fixed scale, therefore the procedure of Sect. 6 can be applied afterwards without any modification. In the THUNC code, these or other models can be implemented through a “custom model” feature (see Sect. B.4).

B.1 Scale variation model with violation of the bound

We have noticed in Sect. 5.4 that the scale variation model tends to set the lower limit of the parameter λ from the first step of the inference, namely from the knowledge of the first non-trivial scale dependence. We observed that the theta function imposed by the next orders is typically looser than that coming from the first order, thereby possibly overestimating the uncertainty by requiring a too large value of λ .

A possible way to solve this issue is to relax the condition Eq. (5.2) to a less stringent one, namely allowing a violation of that bound. This can be implemented in the model through a modified likelihood (Eq. 5.3). So, rather than having a strict theta function that forbids values of $|\delta_k(\mu)| > \lambda r_{k-1}(\mu)$, we allow them, but with an exponential suppression:

$$P(\delta_k | r_{k-1}, \lambda, \mu) = \frac{\rho_k}{2(1 + \rho_k)\lambda r_{k-1}(\mu)} \exp \left[\rho_k \left(1 - \max \left(1, \frac{|\delta_k(\mu)|}{\lambda r_{k-1}(\mu)} \right) \right) \right]. \quad (\text{B.1})$$

This distribution is flat in the same region $|\delta_k| \leq \lambda r_{k-1}$ that is allowed in the default model, and it decreases exponentially in $|\delta_k|/(\lambda r_{k-1})$ outside this region. The strength of the exponential suppression is governed by the parameter ρ_k . Since this generalization of the model is motivated by the behaviour of the first orders that may not be representative of the expansion, it makes sense to allow more violation at the first orders and make the suppression stronger at higher orders. This is achieved by using a different value of ρ_k for each value of the order k . A reasonable choice that makes the suppression strong enough not to bias too much the method is given by

$$\rho_k = 2k^2. \quad (\text{B.2})$$

The form Eq. (B.1) with ρ_k given in Eq. (B.2) does not allow to perform the inference in a simple analytic way. Therefore, this variant of the model can be implemented numerically through the custom model feature of the THUNC code, see Sect. B.4.

To see how this modification improves the behaviour of the model, we report in Fig. 40 the posterior distribution for λ (left plot) and the distributions for the next order (right plot) for the example process of Higgs production in gluon fusion. These results must be compared with Figs. 14 (left plot) and 15 (solid lines), respectively. We observe that the posterior for λ is dramatically different, now allowing much smaller values of λ . Note that the tail of this posterior at the various orders is the same, up to a factor due to the different normalization, which is a consequence of the structure of the inference described in Sect. 5.4 that does not change at large λ . Similarly, the tails of the distributions for the observable are unchanged, but the central shape is different. In particular the plateau has now smooth corners, and most importantly the distributions at high orders are narrower, leading to a more precise (less conservative) prediction. This modified scale variation method is thus better than the simpler one presented in Sect. 5, but using it requires paying the price of a slower numerical implementation due to the inability of solving the model analytically.

Fig. 40 Posterior distribution for the parameter λ (left) and probability distributions for the observable at the next order (right) for the Higgs production process within the modified scale variation model with violation of the bound

B.2 Scale variation model with controlled scale dependence

We have observed in Sect. 5.5 that making a prediction for the observable using two unknown higher orders within the scale variation model leads to very different results from those obtained using only the first missing higher order. The reason for this comes from the fact that inference on δ_{n+2} requires the knowledge of r_{n+1} , that is computed from the inferred δ_{n+1} . The problem is that there are values of δ_{n+1} allowed by the model condition $|\delta_{n+1}| \leq \lambda r_n$ that generate too large values of r_{n+1} , which in turn allow very large values of δ_{n+2} from the next condition $|\delta_{n+2}| \leq \lambda r_{n+1}$. This behaviour however violates the spirit of the model, that assumes that the various r_k are good estimators of the higher orders, thus implying that they behave perturbatively as well.

The obvious solution to this problem is to require an additional constraint on the values of the r_k coefficients. Physically, this means that not only one requires that the higher order correction behaves perturbatively, but also that its scale dependence is reduced (or at least not increased), as one indeed expects. Therefore, adding this condition to the model is very well justified physically, and would lead to a more stringent condition, thus making such a variant of the scale variation model much more precise than its default version. Additionally, the distribution obtained using two (or more) missing higher orders should be way more stable and similar to that obtained using only the first missing higher order.

To implement it in practice, we shall assume that the r_k satisfy the condition

$$r_k \lesssim r_{k-1}. \quad (\text{B.3})$$

This condition can be made more precise by introducing a new parameter η and writing

$$r_k \leq \eta r_{k-1}. \quad (\text{B.4})$$

Fig. 41 Bayesian network used to implement a constraint on the scale dependence of the missing higher orders

If one does not want to have an extra parameter in the model, it is possible to simply fix η to a given value (for instance 1) by using a prior for it given by a delta function. Incorporating the condition in the model leads to a likelihood that contains the theta function

$$P(\delta_k | \delta_{k-1}, \dots, \delta_1, \Sigma_0, \lambda, \eta, \mu) \propto \theta(\eta r_{k-1}(\mu) - r_k(\mu)). \quad (\text{B.5})$$

Note that in principle this condition could be used alone, namely without also imposing the condition Eq. (5.2), thus defining a new model that should already lead to a decent

constraint on the perturbative expansion. However, a better performance can be obtained by requiring both Eq. (5.2) and Eq. (B.4). In this case the likelihood becomes

$$P(\delta_k | \delta_{k-1}, \dots, \delta_1, \Sigma_0, \lambda, \eta, \mu) \propto \theta(\lambda r_{k-1}(\mu) - |\delta_k(\mu)|) \theta(\eta r_{k-1}(\mu) - r_k(\mu)). \quad (\text{B.6})$$

The trouble of this model is that computing the normalization factor of this likelihood is very complicated because it depends on the scale dependence r_k of the observable. However, this normalization is very important, because its dependence on λ and η determines the structure of the inference.

To circumvent this problem, we use a trick, namely we consider r_k as explicit variables of the model, and we introduce new binary variables I_k defined by

$$I_k = \theta(\eta r_{k-1}(\mu) - r_k(\mu)). \quad (\text{B.7})$$

These new variables are considered as “observables”, namely their values is assumed to be fixed (to 1) by “observation” (in our case, this is an assumption and not an observation). To clarify the structure of these variables and their relations, we show in Fig. 41 the Bayesian network of this model. The joint probability distribution for this model, up to order m , is given by

$$\begin{aligned} P(\delta_m, \dots, \delta_1, \Sigma_0, r_m, \dots, r_0, I_m, \dots, I_1, \lambda, \eta, \mu) \\ = P(I_m | r_m, r_{m-1}, \eta) P(r_m | \delta_m, \dots, \delta_1, \Sigma_0, \mu) \\ P(\delta_m | r_{m-1}, \lambda, \mu) \\ \times \dots \\ \times P(I_1 | r_1, r_0, \eta) P(r_1 | \delta_1, \Sigma_0, \mu) P(\delta_1 | r_0, \lambda, \mu) \\ \times P(r_0 | \Sigma_0, \mu) P_0(\Sigma_0, \lambda, \eta, \mu) \end{aligned} \quad (\text{B.8})$$

where

$$P(I_k | r_k, r_{k-1}, \eta) = \delta_{I_k, \theta(\eta r_{k-1} - r_k)} \quad (\text{B.9})$$

$$P(r_k | \delta_k, \dots, \delta_1, \Sigma_0, \mu) = \delta(r_k - r_k(\mu)) \quad (\text{B.10})$$

$$P(\delta_k | r_{k-1}, \lambda, \mu) = \frac{1}{2\lambda r_{k-1}(\mu)} \theta(\lambda r_{k-1}(\mu) - |\delta_k(\mu)|). \quad (\text{B.11})$$

Equations (B.9) and (B.10) are delta functions because the respective variables I_k and r_k are deterministically defined through Eqs. (B.7) and (3.9) respectively, while Eq. (B.11) is the same likelihood of the default model (Eq. 5.3). From the joint probability we may construct any conditional probability we desire. Let us start from the distribution of the first missing higher order δ_{n+1} given the knowledge of the first orders up to n . If we compute $P(\delta_{n+1} | \delta_n, \dots, \delta_1, \Sigma_0, \mu)$ marginalizing over everything else including the I_k , then we would get the same result of the default model, Sect. 5. Here instead we assume the various I_k to be known and equal

to 1, thus imposing the condition Eq. (B.4). In particular, $I_1 = \dots = I_n = 1$ gives constraints on the hidden parameter η , which in turn puts the desired constraint on δ_{n+1} through the condition $I_{n+1} = 1$. So we compute

$$\begin{aligned} P(\delta_{n+1} | \delta_n, \dots, \delta_1, \Sigma_0, I_{n+1}, \dots, I_1, \mu) \\ \propto P(\delta_{n+1}, \dots, \delta_1, \Sigma_0, I_{n+1}, \dots, I_1, \mu) \end{aligned} \quad (\text{B.12})$$

with

$$\begin{aligned} P(\delta_m, \dots, \delta_1, \Sigma_0, I_m, \dots, I_1, \mu) \\ = \int d\lambda d\eta dr_m \dots dr_0 \\ \times P(\delta_m, \dots, \delta_1, \Sigma_0, r_m, \dots, r_0, I_m, \dots, I_1, \lambda, \eta, \mu) \\ = \int d\lambda d\eta P(I_m | r_m, r_{m-1}, \eta, \mu) \dots \\ \times P(I_1 | r_1, r_0, \eta, \mu) P(\delta_m | r_{m-1}, \lambda, \mu) \dots \\ \times P(\delta_1 | r_0, \lambda, \mu) P_0(\lambda, \eta, \Sigma_0, \mu) \end{aligned} \quad (\text{B.13})$$

and, when setting $I_1 = \dots = I_{n+1} = 1$, it reduces to

$$\begin{aligned} P(\delta_m, \dots, \delta_1, \Sigma_0, I_m = 1, \dots, I_1 = 1, \mu) \\ = \int d\lambda d\eta \theta(\eta r_{m-1}(\mu) - r_m(\mu)) \\ \dots \theta(\eta r_0(\mu) - r_1(\mu)) P(\delta_m | r_{m-1}, \lambda, \mu) \\ \dots P(\delta_1 | r_0, \lambda, \mu) P_0(\lambda, \eta, \Sigma_0, \mu), \end{aligned} \quad (\text{B.14})$$

having used the identity $\delta_{1, \theta(x)} = \theta(x)$. Note that this result could have been obtained in the “standard” approach (without introducing auxiliary variables) if the likelihood Eq. (B.6) depended on λ and η as

$$\begin{aligned} P(\delta_k | \delta_{k-1}, \dots, \delta_1, \Sigma_0, \lambda, \eta, \mu) \\ \propto \frac{1}{\lambda} \theta(\lambda r_{k-1}(\mu) - |\delta_k(\mu)|) \theta(\eta r_{k-1}(\mu) - r_k(\mu)). \end{aligned} \quad (\text{B.15})$$

Note that this cannot be the case, as it’s clear that if λ is large enough the likelihood should become independent of it, because the other constraint is more stringent. This means that the implementation of the constraint on the scale dependence through the procedure proposed here and depicted in the network Fig. 41 is not entirely equivalent to the “standard” approach based on the likelihood Eq. (B.6). Nevertheless, it provides an easy and legitimate way of realizing the same constraint.

In order to implement the model in practice, we need to specify the prior for η . Of course, the smaller η , the better the convergence. We could force η to lie in a finite range, but this would provide an upper bound on the scale variation that seems a bit too restrictive. Therefore, we just use a prior that exponentially suppresses large values of η , namely $P_0(\eta) = \exp(-\eta)$. The results are shown in Fig. 42 for the

Fig. 42 Probability distributions for the observable at the next order for the Higgs production process within the scale variation model with constrained scale dependence

usual ggH process. In comparison with the result of the default scale variation model, Sect. 5.5, the distributions are now narrower, and asymmetric, given that the dependence of r_k on δ_k is not linear. The tails die away faster, and effectively exclude the regions where δ_{n+1} generates a too large scale dependence r_{n+1} . This results in more localized distributions, corresponding to smaller uncertainties. This model thus seems a rather powerful improvement over the default scale variation model, the only caveat being the slow numerical evaluation due to the inability of computing inference analytically.

B.3 Mixing models

So far we have seen two main models, characterized by the ability to compute most integrations analytically and thus leading to a fast numerical implementation, and a number of potentially improved versions of each. We now consider the possibility of combining them.

The simplest but also the least performing option consists in considering the assumptions of two (or more) models as *alternative hypothesis*, each with its own prior, and just “average” the outcome of each model with a weight given by the prior. In this way, each model works independently from the other(s), and the combination is done at the end. This approach is similar to the combination of results from different experiments. However, there is a difference: the experiments perform a measurement, while the output of the two methods is directly the probability distribution for the next orders. So, for the experiments one can apply Bayes theorem using the likelihood of the measurement given the measurement, while in our case there is no such likelihood (it should be the likelihood of the model given the prediction, which is difficult to define in general). Therefore, for each model we just use its prior, so that the result of mixing any number of

models is

$$P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0) = \sum_i P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, H_i) P_0(H_i), \quad (\text{B.16})$$

where H_i represents the hypothesis of each model (now written explicitly in the probability distribution). Obviously, $\sum_i P_0(H_i) = 1$. If one has no preference on a specific model, it is natural to use equal probability for the priors of the models.

This approach, though very simple, is not necessarily the best. Indeed, suppose two models look at different characteristics of the perturbative expansions, and the user trusts both models. In this case, the models are not *alternative*, but *complementary*. In such cases, the user wants both conditions to be satisfied simultaneously. This means, for instance, that if one model excludes (probability zero) or at least makes very unlikely (probability almost zero) a region of the observable Σ , after mixing the models this region should remain excluded. However, in the average process of Eq. (B.16) this does not happen. Indeed, Eq. (B.16) represents the condition in which the user trusts one model *or* the other, while here we want to consider the case in which the user trusts one model *and* the other.

We stress that the two families of models considered so far rely on different assumptions: one looks at the behaviour of the perturbative expansion, and the other looks at the scale dependence.⁴⁰ These assumptions are unrelated, and therefore one is allowed to assume that both are satisfied at the same time. In this case we need to act at the core of the model, and modify the likelihood directly. For definiteness, we consider the mixing of the two basic models of Sects. 4 and 5. Thus, assuming that both conditions Eqs. (4.2) and (5.2) are satisfied, the likelihood becomes

$$P(\delta_k|r_{k-1}, c, a, \lambda, \mu) \propto \theta(ca^k - |\delta_k|) \theta(\lambda r_{k-1} - |\delta_k|). \quad (\text{B.17})$$

The two theta functions can be combined in a single one, from which the normalization is also immediate to compute. The result is

$$P(\delta_k|r_{k-1}, c, a, \lambda, \mu) = \frac{1}{2 \min[ca^k, \lambda r_{k-1}]} \theta\left(\min[ca^k, \lambda r_{k-1}] - |\delta_k|\right). \quad (\text{B.18})$$

The presence of the min function unavoidably entangles the two original approaches, and makes the computation of the

⁴⁰ Note that the assumption considered in Sect. B.2 is in principle independent from the model of Sect. 5, and could therefore be considered as another family.

Fig. 43 Probability distributions for the observable at the next order for the Higgs production process mixing the default geometric behaviour and scale variation models either as alternative hypotheses (dashed lines) or as combined hypotheses (solid lines)

inference difficult from an analytical point of view. Using the custom model feature of the `THunc` code (Sect. B.4), we can easily implement this model numerically.

In order to appreciate the way this mixed model work, we show the distributions obtained with it for the usual example of the ggH process. This is depicted in Fig. 43 (solid lines). For comparison, we also show in the same plot the result obtained by averaging the results according to Eq. (B.16), using equal priors for each model (dashed lines). We see that, with respect to each individual model, these distributions are somewhat more precise: the tails die faster than in the geometric behaviour model, and the central region is more peaked than in the scale variation model. In the dashed curves, corresponding to the combination Eq. (B.16), the tails are higher due to the component coming from the geometric behaviour model. Instead, in the solid curves, corresponding to the combined likelihood Eq. (B.18), the tails are those of the scale variation model (up to a normalization factor), and thus exponentially suppressed. Therefore, as expected, mixing the models assuming both conditions Eqs. (4.2) and (5.2) at the same time gives stronger constraints leading to smaller uncertainties. We stress that, while in the plot of Fig. 43 the two combinations look overall very similar, with different models, e.g. including the constraint on the scale dependence as in Sect. B.2 where the probability in some regions is strictly zero, the two ways of combining the models may lead to results that differ not only quantitatively but also qualitatively.

B.4 Custom model

From the discussion so far it is clear that it is always possible to invent new models to improve some aspects or to adapt them to the behaviour of the specific perturbative expansion under consideration. Since it is not possible to predict and

thus to implement all possible models, we provide in the public code `THunc` a so-called “custom model”. This custom model is defined through the following inputs:

- the number of hidden parameters $\vec{p}$;
- the prior of the hidden parameters $P_0(\vec{p})$;
- the likelihood of the model $P(\delta_k | \delta_{k-1}, \dots, \delta_1, \Sigma_0, \vec{p})$.

These ingredients are sufficient to define the model, and to allow the code to perform the inference and compute the sought distributions.

The likelihood has explicit access to the derived quantities r_k , to simplify an implementation based on these objects (as the scale behaviour model). Note that the likelihood has to be properly normalized, namely

$$\int d\delta_k P(\delta_k | \delta_{k-1}, \dots, \delta_1, \Sigma_0, \vec{p}) = 1. \quad (\text{B.19})$$

If this is not possible, the normalization can be restored at the end by normalizing the final distribution. However, for the inference procedure to work properly, the $\vec{p}$ dependence of the likelihood must be complete. In other words, the likelihood must be correct up to a normalization factor that can be either purely numerical or dependent on just the previous orders.

The inference is made by performing numerical integrations. By default, the parameters $\vec{p}$ are assumed to range between 0 and 1. Therefore, for different ranges, the parameters must be rescaled and the jacobian properly included in the prior. Because of the numerical nature of the integration, using the custom model results in a longer running time of the code. The model variants considered in this appendix have been implemented in the code through this custom model feature.

B.5 Adding information on the sign of each order

So far all the models we have proposed consider only the size of the various orders, but not their sign. Therefore, at fixed scale, higher orders will be predicted with positive and negative sign with equal likelihood.⁴¹ As a consequence, the probability distribution for the observable Σ at fixed scale will be symmetric and centered on the highest known order. This prediction may not be satisfactory in many cases. For instance, if the first known orders are all positive, one is tempted to expect with a greater degree of belief that the next order(s) will also follow the same pattern. Moreover, in asymptotically free theories like QCD, at higher scales the coupling is

⁴¹ There is one exception, namely the model with controlled scale variation of Sect. B.2, where the additional condition on the value of r_k makes the likelihood for δ_k asymmetric.

smaller and the various orders are consequently smaller, and a same-sign pattern is generated even if at smaller scales the various orders have apparently random signs. Keeping track of the sign in these cases where a pattern is present would allow to obtain more realistic and precise predictions.

We propose a way to include information on the sign which is simply based on looking for a pattern. We assume that a given pattern corresponds to an hypothesis, and we make inference on that hypothesis by checking the sign pattern of the known coefficients of the expansion. Specifically, we will only consider three hypotheses:

- H_+ , assuming all positive signs;
- H_- , assuming alternating signs;
- H_0 , assuming a “random sign pattern” (i.e., none of the above).

Remember that using the normalized coefficients δ_k we always have $\delta_0 = 1$, so the hypothesis “all negative signs” is not possible (or better, it would correspond to H_+ as an overall minus sign would be factored out in Σ_0). For the same reason, the hypothesis H_- implies that the sign of δ_k is $(-1)^k$. Of course, there are infinitely many other possible sign patterns, but we believe that only the two considered here, $H_{\pm}$, can be useful, while the hypothesis H_0 can cover efficiently all other cases.⁴² After all, testing a more complicated pattern would require a significant number of known orders, which is never the case.

The assumptions of each hypothesis can be expressed in terms of the likelihood of the sign of the generic order k given one of the hypotheses. These probabilities are given by

$$P(\text{sign } \delta_k = + | H_+) = 1, \quad P(\text{sign } \delta_k = - | H_+) = 0, \quad (\text{B.20})$$

$$P(\text{sign } \delta_k = + | H_-) = \frac{1 + (-1)^k}{2}, \quad P(\text{sign } \delta_k = - | H_-) = \frac{1 - (-1)^k}{2}, \quad (\text{B.21})$$

$$P(\text{sign } \delta_k = + | H_0) = \frac{1}{2}, \quad P(\text{sign } \delta_k = - | H_0) = \frac{1}{2}. \quad (\text{B.22})$$

In the case of H_- , the probability of each sign is either 1 or 0 depending on whether k is even or odd. Note that for H_0 the formula is strictly speaking valid only for $k > 0$, as for $k = 0$ we have $\delta_0 = 1$ and thus the probability of $\delta_0 > 0$ has to be 1 for all hypotheses. At first sight, it may seem that for a generic order each of the three hypotheses are to be considered (in practice, only for k even all three probabilities are non-zero). However, we will always consider a sequence of consecutive orders, so that what we really test is not just the

⁴² There is a simple variation of the aforementioned patterns that could occur, in which the pattern is present but it starts from the δ_1 . This is the case if all δ_k are negative except δ_0 (variant of H_+) or if the signs are alternating as $\text{sign } \delta_k = (-1)^{k+1}$ with the exception of δ_0 (variant of H_-). We have seen an example of each: the toy convergent series at large scale, Sect. 7.1, and the anharmonic oscillator, Sect. 7.3, respectively. We will not cover this case explicitly, but testing these additional hypotheses is straightforward.

sign at a given order, but the sign pattern of the sequence. The sequences generated by H_+ and H_- are clearly incompatible ($+++\dots$ vs $+-+\dots$), so in practice only one of the two has to be considered, and compared with H_0 (which is compatible with any possible sequence). Of course, the sequence may be incompatible with both H_+ and H_- , and in this case only the H_0 hypothesis survives.

Let's construct the inference process in practice step by step. With the only knowledge of the LO, one cannot test any sign hypothesis, as $\delta_0 = 1$ is always positive. The first non-trivial information comes from the NLO. The sign of δ_1 already allows to exclude one hypothesis: if $\delta_1 > 0$, then H_- is excluded, viceversa if $\delta_1 < 0$, then H_+ is excluded.⁴³ We are thus left with two hypotheses, as anticipated.

Let us assume that $\delta_1 > 0$, so that H_- is excluded and only H_+ and H_0 survive (the complementary case with $\delta_1 < 0$ proceeds identically and we will not repeat it). The probability of this sequence is given by

$$P(++ | H_+) = 1, \quad P(++ | H_0) = \frac{1}{2}, \quad (\text{B.23})$$

where we have used a shorthand notation $++$ to indicate the signs of δ_0 and δ_1 , and we have used the fact that $P(\text{sign } \delta_0 = + | H_0) = 1$ and not $1/2$. At this point, if the next order has a minus sign, the probability of that sequence is 0 in the H_+ hypothesis, leaving only H_0 :

$$P(++- | H_+) = 0 \Rightarrow P(H_+ | ++-) = 0 \\ \Rightarrow P(H_0 | ++-) = 1. \quad (\text{B.24})$$

It follows that the only sequence that gives non-trivial results is the one with all plus signs. At this order it gives

$$P(+++ | H_+) = 1, \quad P(+++ | H_0) = \frac{1}{2^2}. \quad (\text{B.25})$$

Now let us assume that we know the first $n + 1$ orders, up to δ_n , and that all signs are positive. We have

$$P(+^{n+1} | H_+) = 1, \quad P(+^{n+1} | H_0) = \frac{1}{2^n}, \quad (\text{B.26})$$

where the meaning of the notation $+^{n+1}$ is obvious. From this, we can infer the probabilities of each hypothesis. Using Bayes theorem, we find

$$P(H_+ | +^{n+1}) = \frac{P(+^{n+1} | H_+) P_0(H_+)}{P(+^{n+1} | H_+) P_0(H_+) + P(+^{n+1} | H_0) P_0(H_0)} \\ = \frac{P_0(H_+)}{P_0(H_+) + \frac{1}{2^n} P_0(H_0)}$$

⁴³ If $\delta_1 = 0$, then one cannot exclude any pattern. In this case, one needs to check the next order, δ_2 , before excluding something. Since it is very unlikely that both δ_1 and δ_2 are zero simultaneously, one can always exclude one of the patterns at this order.

$$\begin{aligned}
&= \frac{1}{1 + \frac{1}{2^n} \frac{P_0(H_0)}{P_0(H_+)}} \\
P(H_0|+^{n+1}) &= 1 - P(H_+|+^{n+1}) \\
&= \frac{1}{1 + 2^n \frac{P_0(H_+)}{P_0(H_0)}}, \tag{B.27}
\end{aligned}$$

where $P_0(H_{0,+})$ are the priors for $H_{0,+}$. The posterior probability of H_+ increases as the number of known orders increases, provided they are all positive, while the one for H_0 dies exponentially. Analogous results hold for the comparison of H_- with H_0 , provided the pattern is replaced with the alternating sign one. Note that this inference procedure is exactly the same that occurs when tossing a coin an testing if it's regular (H_0) or if it has "head" on both faces (H_+).

The result obviously depends on the prior probability of each hypothesis. There is no reason to strongly favour one of the sign hypotheses a priori. If a given sign pattern is present, the H_0 hypothesis becomes very unlikely after a few orders due to the exponential suppression, and if no sign pattern is present, the posterior probability for H_0 is 1 anyway. Therefore, an equal prior for each hypothesis makes perfect sense.

Let us now show how this information on the sign pattern is used in the models for the uncertainty from missing higher orders. First of all, for each hypothesis we need to generate a different variant of the model, accounting for the assumptions on the sign pattern. This is expressed in a different likelihood for each hypothesis. For example, in the geometric behaviour model, the likelihoods are

$$\begin{aligned}
P(\delta_k|c, a, \mu, H_0) &= \frac{1}{2ca^k} \theta(ca^k - |\delta_k|), \\
P(\delta_k|c, a, \mu, H_+) &= \frac{1}{ca^k} \theta(ca^k - \delta_k) \theta(\delta_k), \\
P(\delta_k|c, a, \mu, H_-) &= \frac{1}{ca^k} \theta(ca^k - |\delta_k|) \theta((-1)^k \delta_k). \tag{B.28}
\end{aligned}$$

The likelihood for the H_0 hypothesis corresponds to the original one (Eq. 4.3). Similar changes also hold for the scale variation model. The use of these modified likelihoods have a rather simple effect on the probability for unknown orders given the known ones. We can indeed write

$$\begin{aligned}
&P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_+) \\
&= 2^j \theta(\delta_{n+j}) \dots \theta(\delta_1) P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_0), \\
&P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_-) \\
&= 2^j \theta((-1)^{n+j} \delta_{n+j}) \dots \theta(\delta_2) \theta(-\delta_1) \\
&\quad \times P(\delta_{n+j}, \dots, \delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_0), \tag{B.29}
\end{aligned}$$

which is written in terms of the same distribution in the H_0 hypothesis, which we know already. This result is rather general, and it holds for all models for which the H_0 hypothesis leads to a symmetric likelihood (thus, it would not work for

the model of Sect. B.2). In these cases, the numerical implementation of this procedure is very easy. The last step consists in combining the results from the various hypotheses. For the full observable Σ we simply have

$$\begin{aligned}
&P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu) \\
&= P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_+) P(H_+|\text{sign } \delta_n, \dots, \text{sign } \delta_1) \\
&\quad + P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_-) P(H_-|\text{sign } \delta_n, \dots, \text{sign } \delta_1) \\
&\quad + P(\Sigma|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_0) P(H_0|\text{sign } \delta_n, \dots, \text{sign } \delta_1). \tag{B.30}
\end{aligned}$$

The same structure also holds for the probability of anything given the first $n+1$ orders. Thus, for example, the probability of the first missing higher order, using Eq. (B.29), is given by

$$\begin{aligned}
&P(\delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu) \\
&= P(\delta_{n+1}|\delta_n, \dots, \delta_1, \Sigma_0, \mu, H_0) \\
&\quad \times \begin{cases} \frac{2\theta(\delta_{n+1}) + \frac{1}{2^n} \frac{P_0(H_0)}{P_0(H_+)}}{1 + \frac{1}{2^n} \frac{P_0(H_0)}{P_0(H_+)}} & \text{if all signs are } + \text{'s} \\ \frac{2\theta((-1)^{n+1} \delta_{n+1}) + \frac{1}{2^n} \frac{P_0(H_0)}{P_0(H_-)}}{1 + \frac{1}{2^n} \frac{P_0(H_0)}{P_0(H_-)}} & \text{if signs alternate} \\ 1 & \text{otherwise} \end{cases} \tag{B.31}
\end{aligned}$$

which is given in terms of the H_0 results and a modifying function that depends on the priors of each sign pattern.

To visualize the effect of this procedure, we apply it to the divergent series introduced in Sect. 7.2, one of which has all coefficients with the same sign and the other with alternating signs. Therefore, the sign pattern is present and recognised by the algorithm, thus favouring the hypothesis $H_{\pm}$ over H_0 more and more by increasing the order. The distributions for the observable are shown in Fig. 44, for the geometric behaviour model (similar results hold for the scale variation model). The distributions are clearly asymmetric, due to the fact the hypotheses $H_{\pm}$ predict zero for one of the sign of the next order.⁴⁴ In both cases, the asymmetry favours the half-region where the exact result (thin vertical line) lies, thus showing that accounting for the sign pattern does indeed improve the model prediction. Note that, for a real observable in quantum field theory, after marginalizing over the scale the distributions become smooth (without discontinuities).

⁴⁴ In the case of the hypothesis H_- , using more than one order for approximating the observable would result in a (small) contribution in the region where the prediction based only on the first missing higher order is zero. However, the effect is very mild, especially in comparison with the component from the H_0 hypothesis in the same region.

Fig. 44 Probability distributions in the geometric behaviour model for the observable at the next order for the same sign (left) and alternating sign (right) divergent series discussed in Sect. 7.2, taking into account the sign pattern

References

1. C. Duhr, F. Dulat, B. Mistlberger, *The Drell-Yan cross section to third order in the strong coupling constant*. [arxiv:2001.07717](#)
2. C. Duhr, F. Dulat, B. Mistlberger, *Charged Current Drell-Yan Production at N³LO*. [arxiv:2007.13313](#)
3. M. Cacciari, N. Houdeau, Meaningful characterisation of perturbative theoretical uncertainties. *JHEP* **1109**, 039 (2011). [arxiv:1105.5152](#)
4. E. Bagnaschi, M. Cacciari, A. Guffanti, L. Jenniches, *An extensive survey of the estimation of uncertainties from missing higher orders in perturbative calculations*. [arxiv:1409.5036](#)
5. M. Bonvin, S. Marzani, C. Muselli, L. Rottoli, *On the Higgs cross section at N³LO+N³LL and its uncertainty*. *JHEP* **08**, 105 (2016). [arxiv:1603.08000](#)
6. R.D. Ball, V. Bertone, L. Del Debbio, S. Forte, A. Guffanti et al., Precision NNLO determination of $\alpha_s(M_Z)$ using an unbiased global parton set. *Phys. Lett. B* **707**, 66 (2012). [arxiv:1110.2483](#)
7. S. Gorla, G. Passarino, D. Rosco, The Higgs Boson Lineshape. *Nucl. Phys. B* **864**, 530 (2012). [arxiv:1112.5517](#)
8. A. David, G. Passarino, How well can we guess theoretical uncertainties? *Phys. Lett. B* **726**, 266 (2013). [arxiv:1307.1843](#)
9. S. Forte, A. Isgrò, G. Vita, Do we need N³LO Parton Distributions? *Phys. Lett. B* **731**, 136 (2014). [arxiv:1312.6688](#)
10. S. Dulat, T.-J. Hou, J. Gao, M. Guzzi, J. Huston, P. Nadolsky et al., New parton distribution functions from a global analysis of quantum chromodynamics. *Phys. Rev. D* **93**, 033006 (2016). [arxiv:1506.07443](#)
11. R.J. Furnstahl, N. Klco, D.R. Phillips, S. Wesolowski, Quantifying truncation errors in effective field theory. *Phys. Rev. C* **92**, 024005 (2015). [arxiv:1506.01343](#)
12. E. A. Coello Pérez and T. Papenbrock, *Effective field theory for nuclear vibrations with quantified uncertainties*. *Phys. Rev. C* **92**, 064309 (2015). [arxiv:1510.02401](#)
13. H.W. Griesshammer, J.A. McGovern, D.R. Phillips, Nucleon polarisabilities at and beyond physical pion masses. *Eur. Phys. J. A* **52**, 139 (2016). [arxiv:1511.01952](#)
14. E. A. Coello Pérez, T. Papenbrock, *Effective field theory for vibrations in odd-mass nuclei*. *Phys. Rev. C* **94**, 054316 (2016). [arxiv:1608.02802](#)
15. J.A. Melendez, S. Wesolowski, R.J. Furnstahl, Bayesian truncation errors in chiral effective field theory: nucleon–nucleon observables. *Phys. Rev. C* **96**, 024003 (2017). [arxiv:1704.03308](#)
16. NNPDF collaboration, R. D. Ball, S. Carrazza, L. Del Debbio, S. Forte, Z. Kassabov, J. Rojo et al., *Precision determination of the strong coupling constant within a global PDF analysis*. *Eur. Phys. J. C* **78**, 408 (2018). [arxiv:1802.03398](#)
17. J.A. Melendez, R.J. Furnstahl, D.R. Phillips, M.T. Pratala, S. Wesolowski, Quantifying correlated truncation errors in effective field theory. *Phys. Rev. C* **100**, 044001 (2019). [arxiv:1904.10581](#)
18. A. Ekstrom, *Ab initio models of atomic nuclei: challenges and new ideas*. [arxiv:1912.02227](#)
19. F.J. Dyson, Divergence of perturbation theory in quantum electrodynamics. *Phys. Rev.* **85**, 631 (1952)
20. G. 't Hooft, *Can we make sense out of quantum chromodynamics?* *Subnucl. Ser.* **15**, 943 (1979)
21. M. Beneke, Renormalons. *Phys. Rept.* **317**, 1 (1999). [arxiv:hep-ph/9807443](#)
22. M. Serone, G. Spada, G. Villadoro, The power of perturbation theory. *JHEP* **05**, 056 (2017). [arxiv:1702.04148](#)
23. M. Serone, G. Spada, G. Villadoro, $\lambda\phi^4$ Theory I: the symmetric phase beyond NLO. *JHEP* **08**, 148 (2018). [arxiv:1805.05882](#)
24. I. Caprini, Higher-order perturbative coefficients in QCD from series acceleration by conformal mappings. *Phys. Rev. D* **100**, 056019 (2019). [arxiv:1908.06632](#)
25. P. Marquard, A.V. Smirnov, V.A. Smirnov, M. Steinhauser, Quark mass relations to four-loop order in perturbative QCD. *Phys. Rev. Lett.* **114**, 142002 (2015). [arxiv:1502.01030](#)
26. A. Kataev, V. Molokoedov, On the flavour dependence of the $\mathcal{O}(\alpha_s^4)$ correction to the relation between running and pole heavy quark masses. *Eur. Phys. J. Plus* **131**, 271 (2016). [arxiv:1511.06898](#)
27. P. Marquard, A.V. Smirnov, V.A. Smirnov, M. Steinhauser, D. Wellmann, $\overline{\text{MS}}$ -on-shell quark mass relation up to four loops in QCD and a general $\text{SU}(N)$ gauge group. *Phys. Rev. D* **94**, 074025 (2016). [arxiv:1606.06754](#)
28. G. D'Agostini, *Bayesian reasoning in data analysis: A critical introduction*. (2003)
29. S.J. Brodsky, G. Lepage, P.B. Mackenzie, On the elimination of scale ambiguities in perturbative quantum chromodynamics. *Phys. Rev. D* **28**, 228 (1983)
30. G. Grunberg, A. Kataev, On some possible extensions of the Brodsky–Lepage–MacKenzie approach beyond the next-to-leading order. *Phys. Lett. B* **279**, 352 (1992)
31. S.J. Brodsky, H.J. Lu, Commensurate scale relations in quantum chromodynamics. *Phys. Rev. D* **51**, 3652 (1995). [arxiv:hep-ph/9405218](#)

32. P.M. Stevenson, Optimized perturbation theory. *Phys. Rev. D* **23**, 2916 (1981)
33. P.M. Stevenson, An explicit formula for the renormalization scheme invariants of perturbation theory. *Phys. Rev. D* **33**, 3130 (1986)
34. S.J. Brodsky, L. Di Giustino, Setting the renormalization scale in QCD: the principle of maximum conformality. *Phys. Rev. D* **86**, 085026 (2012). [arxiv:1107.0338](#)
35. S.J. Brodsky, X.-G. Wu, Scale setting using the extended renormalization group and the principle of maximum conformality: the QCD coupling constant at four loops. *Phys. Rev. D* **85**, 034038 (2012). [arxiv:1111.6175](#)
36. S.J. Brodsky, X.-G. Wu, Eliminating the Renormalization Scale Ambiguity for Top-Pair Production Using the Principle of Maximum conformality. *Phys. Rev. Lett.* **109**, 042002 (2012). [arxiv:1203.5312](#)
37. S.J. Brodsky, M. Mojaza, X.-G. Wu, Systematic Scale-Setting to All Orders: The Principle of Maximum Conformality and Commensurate scale relations. *Phys. Rev. D* **89**, 014027 (2014). [arxiv:1304.4631](#)
38. F. A. Chishtie, *Achieving effective renormalization scale and scheme independence via the principle of observable effective matching (POEM)*. [arxiv:2005.11783](#)
39. J.-M. Shen, X.-G. Wu, B.-L. Du, S.J. Brodsky, Novel all-orders single-scale approach to QCD renormalization scale-setting. *Phys. Rev. D* **95**, 094006 (2017). [arxiv:1701.08245](#)
40. A. Kataev, S. Mikhailov, Generalization of the Brodsky–Lepage–Mackenzie optimization within the β -expansion and the principle of maximal conformality. *Phys. Rev. D* **91**, 014007 (2015). [arxiv:1408.0122](#)
41. A. Kataev, S. Mikhailov, The $\{\beta\}$ -expansion formalism in perturbative QCD and its extension. *JHEP* **11**, 079 (2016). [arxiv:1607.08698](#)
42. H.A. Chawdhry, A. Mitov, Ambiguities of the principle of maximum conformality procedure for hadron collider processes. *Phys. Rev. D* **100**, 074013 (2019). [arxiv:1907.06610](#)
43. N. Brambilla, H.S. Chung, J. Komijani, Inclusive decays of η_c and η_b at NNLO with large n_f resummation. *Phys. Rev. D* **98**, 114020 (2018). [arxiv:1810.02586](#)
44. C. Anastasiou, C. Duhr, F. Dulat, F. Herzog, B. Mistlberger, Higgs Boson gluon–fusion production in QCD at three loops. *Phys. Rev. Lett.* **114**, 212001 (2015). [arxiv:1503.06056](#)
45. C. Anzai, A. Hasselhuhn, M. Hörschele, J. Hoff, W. Kilgore, M. Steinhauser, T. Ueda, Exact N^3 LO results for $qq' \rightarrow H + X$. *JHEP* **07**, 140 (2015). [arxiv:1506.02674](#)
46. C. Anastasiou, C. Duhr, F. Dulat, E. Furlan, T. Gehrmann, F. Herzog et al., High precision determination of the gluon fusion Higgs boson cross-section at the LHC. *JHEP* **05**, 058 (2016). [arxiv:1602.00695](#)
47. B. Mistlberger, Higgs Boson production at Hadron Colliders at N^3 LO in QCD. [arxiv:1802.00833](#)
48. LHC Higgs Cross Section Working Group collaboration, D. de Florian et al., *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*. [arxiv:1610.07922](#)
49. R.D. Ball, M. Bonvini, S. Forte, S. Marzani, G. Ridolfi, Higgs production in gluon fusion beyond NNLO. *Nucl. Phys. B* **874**, 746 (2013). [arxiv:1303.3590](#)
50. M. Bonvini, R.D. Ball, S. Forte, S. Marzani, G. Ridolfi, Updated Higgs cross section at approximate N^3 LO. [arxiv:1404.3204](#)
51. M. Bonvini, S. Marzani, *Double resummation for Higgs production*. *Phys. Rev. Lett.* **120**, 202003 (2018). [arxiv:1802.07758](#)
52. M. Bonvini, Small- x phenomenology at the LHC and beyond: HELL 3.0 and the case of the Higgs cross section. *Eur. Phys. J. C* **78**, 834 (2018). [arxiv:1805.08785](#)
53. P.A. Baikov, K.G. Chetyrkin, J.H. Kühn, Five-loop running of the QCD coupling constant. *Phys. Rev. Lett.* **118**, 082002 (2017). [arxiv:1606.08659](#)
54. S. Graffi, V. Grecchi, B. Simon, Borel summability: application to the anharmonic oscillator. *Phys. Lett. B* **32**, 631 (1970)
55. C.M. Bender, T.T. Wu, Anharmonic oscillator. *Phys. Rev.* **184**, 1231 (1969)
56. S. Catani, D. de Florian, M. Grazzini, Higgs production in hadron collisions: soft and virtual QCD corrections at NNLO. *JHEP* **05**, 025 (2001). [arxiv:hep-ph/0102227](#)
57. S. Catani, M.L. Mangano, P. Nason, L. Trentadue, The resummation of soft gluon in hadronic collisions. *Nucl. Phys. B* **478**, 273 (1996). [arxiv:hep-ph/9604351](#)
58. M. Bonvini, S. Marzani, Resummed Higgs cross section at N^3 LL. *JHEP* **1409**, 007 (2014). [arxiv:1405.3654](#)
59. P.A. Baikov, K.G. Chetyrkin, J.H. Kuhn, J. Rittinger, Adler function, sum rules and Crewther relation of order $O(\alpha_s^4)$: the singlet case. *Phys. Lett. B* **714**, 62 (2012). [arxiv:1206.1288](#)
60. P. Baikov, K. Chetyrkin, J.H. Kuhn, Scalar correlator at $O(\alpha_s^{**4})$, Higgs decay into b-quarks and bounds on the light quark masses. *Phys. Rev. Lett.* **96**, 012003 (2006). [arxiv:hep-ph/0511063](#)
61. F. Herzog, B. Ruijl, T. Ueda, J.A.M. Vermaseren, A. Vogt, On Higgs decays to hadrons and the R-ratio at N^4 LO. *JHEP* **08**, 113 (2017). [arxiv:1707.01044](#)
62. F. J. Tackmann, *Beyond scale variations: perturbative theory uncertainties from nuisance parameters*. *DESY-19-021* (2020)
63. NNPDF collaboration, R. Abdul Khalek et al., *A first determination of parton distributions with theoretical uncertainties*. *Eur. Phys. J. C* **79**, 838 (2019). [arxiv:1905.04311](#)
64. NNPDF collaboration, R. Abdul Khalek et al., *Parton distributions with theory uncertainties: general formalism and first phenomenological studies*. *Eur. Phys. J. C* **79**, 931 (2019). [arxiv:1906.10698](#)
65. L.A. Harland-Lang, R.S. Thorne, On the consistent use of scale variations in PDF fits and predictions. *Eur. Phys. J. C* **79**, 225 (2019). [arxiv:1811.08434](#)