

Reweighting QCD matrix-element and parton-shower calculations

Enrico Bothmann^{1,a}, Marek Schönherr^{2,b}, Steffen Schumann^{1,c}

¹ II. Physikalisches Institut, Universität Göttingen, 37077 Göttingen, Germany

² Physik-Institut, Universität Zürich, 8057 Zürich, Switzerland

Received: 6 July 2016 / Accepted: 10 October 2016 / Published online: 28 October 2016

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract We present the implementation and validation of the techniques used to efficiently evaluate parametric and perturbative theoretical uncertainties in matrix-element plus parton-shower simulations within the SHERPA event-generator framework. By tracing the full α_s and PDF dependences, including the parton-shower component, as well as the fixed-order scale uncertainties, we compute variational event weights on-the-fly, thereby greatly reducing the computational costs to obtain theoretical-uncertainty estimates.

Contents

1	Introduction	1
2	Reweighting fixed-order calculations	2
2.1	The leading-order case	2
2.2	The next-to-leading-order case	3
2.3	Validation	4
3	Reweighting parton-shower calculations	6
3.1	Parton-shower dependence structure	7
3.2	Reweighting trial emissions	7
3.3	Next-to-leading-order matching	8
3.4	Validation	9
	The final-state only case: Thrust in $e^+e^- \rightarrow q\bar{q}$ events	9
	The initial-state dominated case: p_\perp^W in $pp \rightarrow W[e\nu]$ events	9
4	Reweighting multijet-merged calculations	11
4.1	Preliminaries	13
4.2	The leading-order case	13
4.3	The next-to-leading-order case	15
4.4	Validation	16
5	Consistent variations	17
6	Conclusions	18
	Appendix A: CPU time measurements	20

Appendix B: Configuring and accessing event-weight variations	21
Scale variations	22
PDF and α_s variations	22
Configuring how variations are calculated	22
References	23

1 Introduction

The first operational run of the LHC collider during the years 2009–2013 was a tremendous success, clearly culminating in the announcement of the discovery of a Higgs-boson candidate by the ATLAS and CMS collaborations in July 2012 [1,2]. Through a large number of experimental analyses, focusing on a variety of final states and observables, the LHC experiments (re)established and underpinned to an unprecedented level of accuracy the validity of the Standard Model of particle physics (SM) [3].

When comparing theoretical predictions with actual collider data, Monte-Carlo event generators prove to be an indispensable tool. In particular parton-shower Monte-Carlo programmes like HERWIG [4,5], PYTHIA [6] and SHERPA [7,8] provide simulations at the level of exclusive particle-level final states [9]. The cornerstones of these generators are their implementations of QCD parton-shower algorithms and their modelling of the non-perturbative parton-to-hadron fragmentation process. With the advent of sophisticated techniques to combine parton-shower simulations with exact higher-order QCD calculations at leading [10,11], next-to-leading [12,13] and even next-to-next-to-leading order [14–17], Monte-Carlo simulations have developed into high-precision tools, encapsulating the best of our current knowledge of perturbative QCD.

With these simulations being widely used for making SM predictions, e.g. of the background expectation in searches for New Physics or the detailed properties of Higgs-boson production final states, a comprehensive and efficient eval-

^a e-mail: enrico.bothmann@phys.uni-goettingen.de

^b e-mail: marek.schoenherr@physik.uzh.ch

^c e-mail: steffen.schumann@phys.uni-goettingen.de

uation of associated theoretical uncertainties is of utmost importance. A comprehensive list of sources for generator uncertainties has been quoted in [18]. Following the categories identified there, when focusing on systematics related to the perturbative phases of event evolution, the following uncertainties might be distinguished:

- *Parametric* uncertainties reflecting the dependence of the prediction on input parameters such as couplings, particle masses or the parton-density functions (PDFs).
- *Perturbative* uncertainties originating from the fact that perturbation theory is used in making predictions, to fixed-order in the matrix elements and resummed to all-orders with a certain logarithmic accuracy in the showers, thereby, however, neglecting higher-order contributions. Similarly, the use of the large- N_c approximation in the showers belongs in this category.
- *Algorithmic* uncertainties corresponding to the actual choices made in the implementation of the shower algorithm, i.e. for the evolution variable, the inclusion of non-singular terms in the splitting functions, or the employed matching/merging prescription. Per construction, for sensible choices, these systematics also correspond to higher-order perturbative corrections, but might be addressed separately.

In addition to the listed categories, generically non-perturbative effects such as hadronisation or the underlying event are described through phenomenological models that feature various generator-specific choices and parameters, typically subject to tuning against experimental data, see for instance [19,20].

This publication focuses on the efficient evaluation of parametric and (some) perturbative uncertainties in matrix-element plus parton-shower simulations within the SHERPA event-generator framework. We present a comprehensive approach to fully trace the α_s and PDF dependences in the matrix-element and parton-shower components of particle-level SHERPA simulations in leading- [21] and next-to-leading [22] order merged calculations based on the SHERPA dipole-shower implementation [23]. Furthermore, we provide the means to quickly evaluate the renormalisation- and factorisation-scale dependence of the fixed-order matrix-element contributions. Our approach is based on event-wise reweighting and allows us to provide with a single generator run a set of variational event weights corresponding to the predefined parameter and scale variations, that would otherwise have to be determined through dedicated re-evaluations. The alternative event weights can either be accessed through the output of a HEPMC event record [24], or directly passed via the internal interface of SHERPA to the RIVET analysis framework [25].

The systematics of leading-order parton-shower simulations with HERWIG 7 have recently been discussed in [18], a

corresponding reweighting procedure has been presented in [26]. A similar reweighting implementation for the PYTHIA 8 parton shower has also appeared recently [27]. A discussion of uncertainty estimates for the VINCIA shower model can be found in [28–30]. A comprehensive comparison of various generators is presented in [31]. The impact of PDFs in parton-shower simulations has been discussed in [32,33].

Our paper is organised as follows. In Sect. 2 we review the dependence structure of leading-order (LO) and next-to-leading-order QCD calculations on α_s , the PDFs and the renormalisation and factorisation scales, and introduce the reweighting approach. In Sect. 3 we extend this to parton-shower simulations and in particular the algorithm employed in the SHERPA framework. In Sect. 4 we present the generalisation of the reweighting approach to multijet-merged calculations, based on leading and next-to-leading-order matrix elements matched to the parton shower. Our conclusions are summarised in Sect. 6. In Appendix A we present CPU time measurements that assess the reduction in computational time when the reweighting is used. The technical details on enabling and accessing the variations considered in SHERPA runs are listed in Appendix B.

Note, while fixed-order reweighting is already available with SHERPA-2.2, the general reweighting implementation described here, including parton showers and multijet merging, will be part of the next release, i.e. SHERPA-2.3.

2 Reweighting fixed-order calculations

In order to re-evaluate a QCD cross-section calculation for a new choice of input parameters, i.e. α_s , PDFs or renormalisation and factorisation scales, it is necessary to understand and trace-out its respective dependences. This is a rather easy task at leading-order (LO) but is already more involved when considering next-to-leading order (NLO) calculations in a given subtraction scheme. However, these decompositions have been presented for Catani–Seymour dipole subtraction and the FKS subtraction formalism in [34,35].

In this section, we briefly review the dependence structure and discuss the corresponding reweighting equations for LO and Catani–Seymour subtracted NLO calculations within the SHERPA framework. With this paragraph we also introduce the notation used in the later sections, which explore the reweighting of more intricate QCD calculations, involving QCD parton showers and merging different final-state multiplicity processes.

2.1 The leading-order case

A LO parton-level calculation of some observable or measurement function of the final-state momenta O is based on Born matrix elements B of $\mathcal{O}(\alpha_s^n)$. It exhibits explicit depen-

dences on the PDFs $f = f_a(x, \mu_F^2)$, the running strong coupling $\alpha_s = \alpha_s(\mu_R^2)$, the renormalisation scale μ_R and the factorisation scale μ_F :

$$\begin{aligned} \langle O \rangle^{\text{LO}} &= \int d\Phi_B \mathbf{B}(\Phi_B) O(\Phi_B) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \sum_{i=1}^N \mathbf{B}(\Phi_{B,i}) O(\Phi_{B,i}) \end{aligned} \tag{1}$$

with $N_{\text{trial}} = \sum_{i=1}^N n_{\text{trial},i}$, n_{trial} denoting the number of attempts to generate an accepted event configuration, and

$$\begin{aligned} \mathbf{B}(\Phi_B) &\equiv \mathbf{B}(\Phi_B; \alpha_s, f; \mu_R, \mu_F) \\ &= \alpha_s^n(\mu_R^2) f_a(x_a, \mu_F^2) f_b(x_b, \mu_F^2) \mathbf{B}'(\Phi_B). \end{aligned} \tag{2}$$

Therein, the \mathbf{B} contains all couplings, symmetry and flux factors, and PDFs, whereas \mathbf{B}' has the PDFs, here for assumed two incoming parton flavours a and b , and the strong coupling stripped off. Note that we have suppressed the event index i here. It is understood that \mathbf{B} depends on the event kinematics and that μ_R and μ_F can be chosen dynamically, i.e. in a momentum (and flavour) dependent way. Changing the input parameters $\mu_R \rightarrow \tilde{\mu}_R$, $\mu_F \rightarrow \tilde{\mu}_F$, and the input functions $f \rightarrow \tilde{f}$, $\alpha_s \rightarrow \tilde{\alpha}_s$ results in

$$\begin{aligned} \mathbf{B}(\Phi_B; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_R, \tilde{\mu}_F) \\ = \tilde{\alpha}_s^n(\tilde{\mu}_R^2) \tilde{f}_a(x_a, \tilde{\mu}_F^2) \tilde{f}_b(x_b, \tilde{\mu}_F^2) \mathbf{B}'(\Phi_B). \end{aligned} \tag{3}$$

From Eq. (3) we conclude that for PDF reweighting it is necessary to know the $x_{a,b}$ values of the event.

For an unweighted event generation, the event weights are uniform initially, i.e. $\mathbf{B}(\Phi_B; \alpha_s, f; \mu_R, \mu_F) = w_{\text{norm}}$, Eq. (1) thus simplifies to

$$\langle O \rangle^{\text{LO}} = \lim_{N \rightarrow \infty} \frac{w_{\text{norm}}}{N_{\text{trial}}} \sum_{i=1}^N O(\Phi_{B,i}). \tag{4}$$

Scale and parameter variations then work the very same way as for weighted events. Applying Eq. (3) then, however, leads to a broader weight distribution and Eq. (1) has to be used again. Partially unweighted events can be treated on the same footing. These conclusions hold irrespective of the type of event generation whenever (partially) unweighted event generation is possible, i.e. when the weight distribution is bounded from above and below. We therefore will not comment further on it.

2.2 The next-to-leading-order case

A full NLO parton-level calculation including real-emission and one-loop corrections of $\mathcal{O}(\alpha_s^{n+1})$ based in Catani–Seymour dipole subtraction [36,37] has the following struc-

ture

$$\begin{aligned} \langle O \rangle^{\text{NLO}} &= \int d\Phi_B \left[\mathbf{B}(\Phi_B) + \text{VI}(\Phi_B) \right. \\ &\quad + \left. \int dx'_{a/b} \text{KP}(\Phi_B, x'_{a/b}) \right] O(\Phi_B) \\ &\quad + \int d\Phi_R \left[\mathbf{R}(\Phi_R) O(\Phi_R) \right. \\ &\quad \left. - \sum_j \text{D}_{S,j}(\Phi_{B,j} \cdot \Phi_1^j) O(\Phi_{B,j}) \right] \\ &= \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \left\{ \sum_{i=1}^{N_B} \left[\mathbf{B}(\Phi_{B,i}) \right. \right. \\ &\quad + \left. \left. \text{VI}(\Phi_{B,i}) + \text{KP}(\Phi_{B,i}, x'_{a/b}) \right] O(\Phi_{B,i}) \right. \\ &\quad + \sum_{i=1}^{N_R} \left[\mathbf{R}(\Phi_{R,i}) O(\Phi_{R,i}) \right. \\ &\quad \left. \left. - \sum_j \text{D}_{S,j}(\Phi_{B,j,i} \cdot \Phi_{1,i}^j) O(\Phi_{B,j,i}) \right] \right\}, \end{aligned} \tag{5}$$

where the new parts have the following dependences

$$\begin{aligned} \text{VI}(\Phi_B) &\equiv \text{VI}(\Phi_B; \alpha_s, f; \mu_R, \mu_F), \\ \text{KP}(\Phi_B, x'_{a/b}) &\equiv \text{KP}(\Phi_B, x'_{a/b}; \alpha_s, f; \mu_R, \mu_F), \\ \mathbf{R}(\Phi_R) &\equiv \mathbf{R}(\Phi_R; \alpha_s, f; \mu_R, \mu_F), \\ \text{D}_{S,j}(\Phi_{B,j} \cdot \Phi_1^j) &\equiv \text{D}_{S,j}(\Phi_{B,j} \cdot \Phi_1^j; \alpha_s, f; \mu_{R,j}, \mu_{F,j}). \end{aligned} \tag{6}$$

Therein, VI combines the renormalised one-loop matrix element with the I-operator of the Catani–Seymour subtraction scheme. This operator gives the flavour-diagonal endpoint contribution of the integrated subtraction terms. VI is thus separately infrared finite and exhibits a common transformation behaviour. Thus, for $\alpha \rightarrow \tilde{\alpha}$, $f \rightarrow \tilde{f}$, $\mu_R \rightarrow \tilde{\mu}_R$ and $\mu_F \rightarrow \tilde{\mu}_F$

$$\begin{aligned} \text{VI}(\Phi_B; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_R, \tilde{\mu}_F) \\ = \tilde{\alpha}_s^{n+1}(\tilde{\mu}_R^2) \tilde{f}_a(x_a, \tilde{\mu}_F^2) \tilde{f}_b(x_b, \tilde{\mu}_F^2) \left[\text{VI}'(\Phi_B) \right. \\ \left. + c_R'^{(0)} l_R + \frac{1}{2} c_R'^{(1)} l_R^2 \right], \end{aligned} \tag{7}$$

with α_s - and PDF-independent coefficients $c_R'^{(i)}$ and $l_R = \log(\tilde{\mu}_R^2/\mu_R^2)$. Again, VI' is stripped of all coupling and PDF factors.

The KP -terms are defined as the remainders of the integrated dipole subtraction terms, containing all flavour chang-

ing and $x'_{a/b}$ -dependent pieces, combined with the collinear counterterms. Here, $x'_{a/b}$ are the ratios of the partonic momentum fractions in the respective dipole before and after radiation. Again, this combination is separately infrared finite and transforms as one unit. When evaluated for the modified set of input parameters, they read

$$\begin{aligned}
 & \text{KP}(\Phi_B, x'_{a/b}; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_R, \tilde{\mu}_F) \\
 &= \tilde{\alpha}_s^{n+1} (\tilde{\mu}_R^2) \tilde{f}_a(x_a, \tilde{\mu}_F^2) \\
 & \quad \times \tilde{f}_b(x_b, \tilde{\mu}_F^2) \text{KP}'(\Phi_B, x'_{a/b}; \tilde{f}; \tilde{\mu}_F) \\
 &= \tilde{\alpha}_s^{n+1} (\tilde{\mu}_R^2) \left[\left(\tilde{f}_a^q c'_{F,a}{}^{(0)} + \tilde{f}_a^q(x'_a) c'_{F,a}{}^{(1)} + \tilde{f}_a^g c'_{F,a}{}^{(2)} \right. \right. \\
 & \quad \left. \left. + \tilde{f}_a^g(x'_a) c'_{F,a}{}^{(3)} \right) \tilde{f}_b(x_b, \tilde{\mu}_F^2) \right. \\
 & \quad \left. + \tilde{f}_a(x_a, \tilde{\mu}_F^2) \left(\tilde{f}_b^q c'_{F,b}{}^{(0)} + \tilde{f}_b^q(x'_b) c'_{F,b}{}^{(1)} + \tilde{f}_b^g c'_{F,b}{}^{(2)} \right. \right. \\
 & \quad \left. \left. + \tilde{f}_b^g(x'_b) c'_{F,b}{}^{(3)} \right) \right] \tag{8}
 \end{aligned}$$

with the coefficients $c'_{F,a/b}{}^{(i)} = \tilde{c}_{F,a/b}{}^{(i)} + \bar{c}_{F,a/b}{}^{(i)} l_F$ for $i \in \{0, \dots, 3\}$, $l_F = \log(\tilde{\mu}_F^2/\mu_F^2)$, and

$$\begin{aligned}
 \tilde{f}_q^q &= \tilde{f}_q(x_q, \tilde{\mu}_F^2), & \tilde{f}_g^q &= \sum_q \tilde{f}_q(x_g, \tilde{\mu}_F^2), \\
 \tilde{f}_q^q(x'_q) &= x'_q \tilde{f}_q\left(\frac{x_q}{x'_q}, \tilde{\mu}_F^2\right), & \tilde{f}_g^q(x'_g) &= x'_g \sum_q \tilde{f}_q\left(\frac{x_q}{x'_g}, \tilde{\mu}_F^2\right), \\
 \tilde{f}_q^g &= \tilde{f}_g(x_q, \tilde{\mu}_F^2), & \tilde{f}_g^g &= \tilde{f}_g(x_g, \tilde{\mu}_F^2), \\
 \tilde{f}_q^g(x'_q) &= x'_q \tilde{f}_g\left(\frac{x_q}{x'_q}, \tilde{\mu}_F^2\right), & \tilde{f}_g^g(x'_g) &= x'_g \tilde{f}_g\left(\frac{x_g}{x'_g}, \tilde{\mu}_F^2\right),
 \end{aligned}$$

for $a, b = \{q, g\}$, respectively. Thereby, the sum over q includes all light-quark flavours, corresponding to all potential quarks emitting a gluon. We note that in order to obtain the reweighted expressions for the VI and KP contributions the additional book-keeping of the $c_R{}^{(i)}$, $\tilde{c}_{F,a/b}{}^{(i)}$ and $\bar{c}_{F,a/b}{}^{(i)}$ (altogether 18)¹ coefficients is required [34]. Due to its composite structure, the KP-terms do not possess a coupling- and PDF-stripped version KP'. Nonetheless, we formally introduce a still PDF-dependent version KP' in Eq. (8) for reference in later sections.

The remaining pieces of Eq. (5) are the Born matrix element B, the real emission contribution R and the differential dipole subtraction terms $D_{S,j}$. The latter defines an underlying Born configuration $\Phi_{B,j}$ through its dipole-dependent phase-space map, employing the phase-space factorisation $\Phi_R = \Phi_{B,j} \cdot \Phi_1^j$. While the transformation of B under

¹ The two parameters $c_R{}^{(i)}$ correspond to the single and double pole coefficients of the loop matrix element while the remaining sixteen coefficients are comprised of eight pairs of coefficients, $\tilde{c}_{F,a/b}{}^{(i)}$ and $\bar{c}_{F,a/b}{}^{(i)}$, corresponding to the μ_F -dependent and -independent parts for all four flavour structures of each beam, respectively.

the exchange of input parameters was detailed in Eq. (3), the transformation of R and the $D_{S,j}$ contributions works identically, merely having to adjust the power of the strong-coupling factor.

2.3 Validation

The reweighting approach outlined above has been implemented in the SHERPA framework for the two matrix-element generators AMEGIC [38] and COMIX [39,40] in conjunction with the corresponding Catani–Seymour dipole-subtraction implementation [41]. The required decomposition of virtual amplitudes is generic and can be used for matrix elements from BLACKHAT [34,42], OPENLOOPS [43], GoSAM [44], NJET [45], the internal library of simple 2 → 2 processes, or, via the BLHA interface [46].

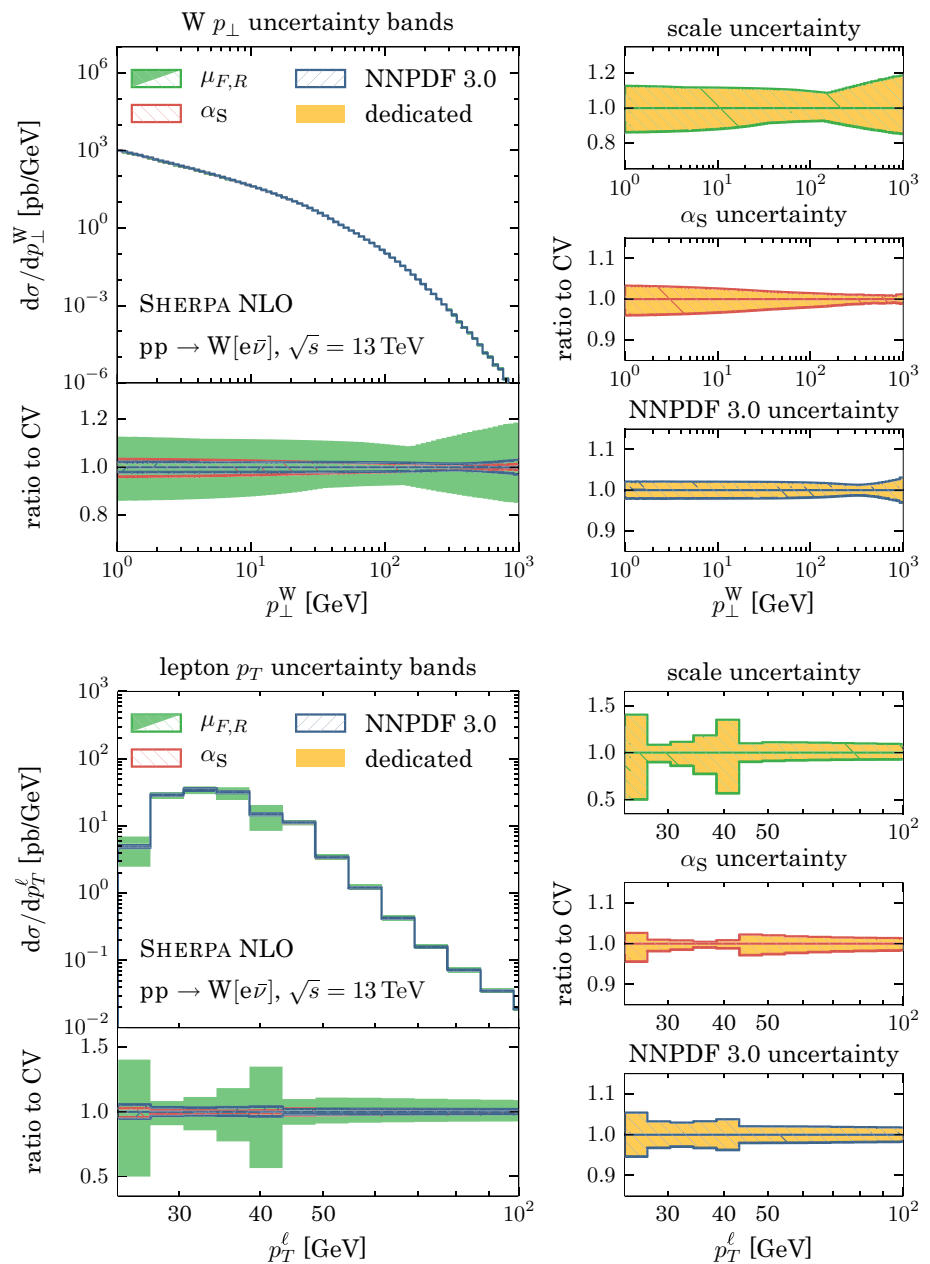
Here we shall present the validation of the reweighting approach in particular of NLO QCD event samples. For that purpose we consider W-boson zproduction in 13 TeV proton-proton collisions at NLO QCD, and focus on the transverse-momentum distribution for the Wand the lepton it decays to. In Fig. 1, the scale, α_s and PDF uncertainty bands for the $W p_\perp$ and the lepton p_\perp distributions are presented. All three bands have been produced for both observables using the internal reweighting of SHERPA from a single event generation run using $\mu_F = \mu_R = H_T^l$ with

$$H_T^l \equiv m_\perp^{\text{ev}} + \sum_j p_\perp^j, \tag{9}$$

a scale choice that has been motivated in [47]. For the PDFs the NNPDF 3.0 NLO set [48] has been used with $\alpha_s(m_Z^2) = 0.118$. The running of $\alpha_s(\mu_R^2)$ is calculated within SHERPA using its renormalisation group equation at NLO with parton thresholds as given by the PDF.

The treatment of partonic thresholds deserves a short discussion. While any flavour thresholds in the running of α_s do not present any challenges to the reweighting algorithm as $\alpha_s(\mu^2) > 0$ for all $\mu^2 > 0$ and any loop order, this is different for the PDFs, where crossing a parton threshold results in a vanishing PDF for that flavour. Hence, the cross section component of the given partonic channel may be zero if no other non-zero contribution exists. Such an event will be discarded and, thus, cannot be reweighted. If now the respective parton threshold of the target PDF is smaller than the target factorisation scale while the one of the nominal PDF is larger than the nominal factorisation scale we are in a region of phase space where the reweighting must fail to reproduce a dedicated calculation. This could be remedied by storing events as well which vanish solely due to crossing PDF thresholds. However, as only observables sensitive to on-threshold production of light quarks (typically bottom quarks) are suscep-

Fig. 1 The gauge-boson and lepton transverse momenta in off-shell Wproduction at the LHC with independent variations of $\mu_{F,R}$ (green), α_s (red) and the PDF (blue). In the right-hand panels, the individual uncertainty bands, calculated via an on-the-fly reweighting, are compared to uncertainty bands from dedicated calculations (yellow). They are found to be equal



tible to these effects, they are of little relevance to the vast majority of LHC observables.²

For the *scale uncertainty band* we employ a 7-point scale variation for μ_R and μ_F : Both scales are varied independently by factors of $\frac{1}{2}$ and 2, omitting the variations with ratios of 4 between the two scales. The uncertainty is then taken as the envelope of all variations. The *α_s uncertainty band* is generated by varying the numerical value of the start-

ing point of the running coupling, $\alpha_s(m_Z^2)$, to the following five values: 0.115, 0.117, 0.118, 0.119 and 0.121. Note that this variation of α_s should also enter the PDF fit, and hence the PDFs are varied consistently. This is expected to extenuate the effect of the α_s variation in most cases, as the PDF of the varied α_s is still fitted to describe the same data as the PDF of the nominal α_s . This consistent α_s +PDF variation is also part of the PDF4LHC recommendations for LHC Run II [49]. The envelope of these α_s +PDF variations is taken as the respective uncertainty. The pure *PDF uncertainty estimate* is generated using the average and the standard deviation over the 100 PDF replicas provided by the NNPDF3.0 set (at a fixed value of $\alpha_s = 0.118$). This corresponds to the

² A typical example for the threshold problem in the reweighting would be the very low- p_{\perp} part of b -jet spectrum in Wb production in a calculation with five massless flavours. Any strong dependence on the bottom quark PDF threshold, however, also indicates the invalidity of a calculation with five massless flavour for this observable.

Table 1 Variations, which are used for studies in this publication, with two variants depending on the PDF choice. Note that each $\alpha_s(m_Z^2)$ value is used with its associated PDF set variant in the context of hadronic collisions

	Nominal	Variations	Error band
PDF sets	CT14 NNPDF3.0	56 Hessian error sets with a 90% CL 100 statistical replicas with a 68% CL	Hessian Statistical
$\alpha_s(m_Z^2)$ value	0.118	0.115, 0.117, 0.119, 0.121	Envelope
μ_R/μ_F factors	(1, 1)	$(\frac{1}{2}, \frac{1}{2}), (1, \frac{1}{2}), (\frac{1}{2}, 1), (2, 1), (1, 2), (2, 2)$	Envelope

68% confidence level. This set-up is repeated for later reference in Table 1, along with a CT14 PDF variant, which is used in later studies.

Comparing the uncertainties for the $W p_{\perp}$, we observe that the scale uncertainties are the largest, with relative deviations of $\mathcal{O}(10\%)$. The relative deviations related to the PDF and the strong coupling do not exceed $\sim 3\%$. The scale uncertainty exhibits a minimum for $100 \text{ TeV} < p_{\perp}^W < 200 \text{ TeV}$. The reason is that the variations of μ_F alone cross the central value prediction in this range, such that only the μ_R variation contributes to the overall scale uncertainty here.

Note that $p_{\perp}^W = 0$ at $\mathcal{O}(\alpha_s^0)$, and therefore only real-emission events contribute to the distribution. Hence, the observable is only described to leading-order. We introduce it here as a reference for our later validations including the parton-shower, which use this observable. For the current validation, we complement the discussion of the W transverse momentum with the one of the lepton it decays to, as the region below $m_W/2$ is already filled at $\mathcal{O}(\alpha_s^0)$, and therefore we have in part a true next-to-leading description for this observable. In fact, the scale uncertainties are much larger in that region, especially towards the $m_W/2$ threshold, and at the lepton p_{\perp} cut at 25 TeV. This gives a more realistic picture of the perturbative uncertainties than in the leading-order region above the threshold.

The small panels on the right of Fig. 1 compare the uncertainty bands calculated using the reweighting approach to uncertainty bands where dedicated calculations have been done for each variation. We observe that all bands overlap perfectly for both observables. This is because the reweighting as presented above is exact and for all runs the same phase-space points could be used: The reweighted and the dedicated predictions for each variation are therefore equal, and so are the uncertainty bands.³

3 Reweighting parton-shower calculations

If parton-showering is added to a LO calculation, the value of the observable is not evaluated at Φ_B any longer, but at

³ The reweighted and the dedicated calculations are implemented independently, such that their predictions can vary within the numerical uncertainties of the calculation. However, these lie several orders of magnitude below the physical uncertainties considered here.

$\text{PS}(\Phi_B)$, which denotes the phase-space point after showering. Applying this modification to Eq. (1) yields

$$\begin{aligned} \langle O \rangle^{\text{LOPs}} &= \int d\Phi_B \text{B}(\Phi_B) \text{PS}(O, \Phi_B) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \sum_{i=1}^N \text{B}(\Phi_{B,i}) \text{PS}(O, \Phi_{B,i}). \end{aligned} \quad (10)$$

Therefore the reweighting for B does not need to be altered, but the parton-shower emissions depend on the PDF, the strong coupling, their respective scale prefactors k_{α_s} and k_f (detailed below) and the starting scale μ_Q , i.e.

$$\begin{aligned} \text{PS}(O, \Phi_B) &\equiv \text{PS}(O, \Phi_B; \mu_Q^2) \\ &\equiv \text{PS}(O, \Phi_B; k_{\alpha_s}, k_f; \alpha_s, f; \mu_Q^2). \end{aligned} \quad (11)$$

In order to reweight the parton-shower emissions, we first need to identify its exact dependence structure. Schematically, it acts on the phase-space element in the following way

$$\begin{aligned} \text{PS}(O, \Phi_n; t') &= \Delta_n(t_{\text{IR}}, t') O(\Phi_n) \\ &+ \int_{t_{\text{IR}}}^{t'} d\Phi_1 \text{K}_n(\Phi_1) \Delta_n(t, t') \text{PS}(O, \Phi_{n+1}; t), \end{aligned} \quad (12)$$

where the Sudakov form factor of the n -parton state, Δ_n , and its splitting kernel K_n have been introduced. While the first term describes the no-emission probability between the starting scale t' and the infrared cut-off t_{IR} and therefore does not change the phase-space element, the second term describes the emission of a parton at scale t in the configuration $d\Phi_1 = dt dz d\phi J(t, z)$ (the integration boundaries are to be understood in this decomposition), leading to a configuration $d\Phi_{n+1} = d\Phi_n \cdot d\Phi_1$. The Jacobian J is not relevant to the discussion here and is subsequently absorbed in the splitting kernel K_n . As the emissions are ordered in t , the Sudakov form factor in the second term ensures that the current emission is the hardest after starting the evolution at t' . Additional emissions may occur at smaller t and are not resolved at this stage – they are described by the parton shower acting on the newly produced state Φ_{n+1} with the new starting scale t . In Eq. (12) the dependences on α_s , the PDFs, and their respective scale prefactors k_{α_s} and k_f have been omitted for brevity. They directly carry over to the splitting kernel and the Sudakov form factor, according to

$$\Delta_n(t_2, t_1; k_{\alpha_s}, k_f; \alpha_s, f) = \exp\left(-\int_{t_2}^{t_1} d\Phi_1 K_n(\Phi_1; k_{\alpha_s}, k_f; \alpha_s, f)\right). \tag{13}$$

When considering parton-shower emissions off NLO QCD matrix elements special emphasis has to be given to the first emission as described in Sect. 3.3 below.

3.1 Parton-shower dependence structure

The default parton shower of SHERPA, dubbed CSSHOWER [23], is based on Catani–Seymour dipole factorisation [36, 37]. Each branching of an emitter parton into two daughters is witnessed by a spectator parton, which takes the recoil, and ensures that on-shell states are transferred into on-shell states and energy–momentum conservation is respected simultaneously. The emitter and spectator partons reside either in the initial-state (I) or final-state (F), such that four dipole types need to be distinguished: II, IF, FI and FF. In this notation, the first letter refers to emitter, and the second to the spectator parton. The no-branching probabilities are given by the four corresponding Sudakov form factors

$$\Delta_n(t_2, t_1; k_{\alpha_s}, k_f; \alpha_s, f) = \prod_{\text{type} \in \{\text{FF}, \text{FI}, \text{IF}, \text{II}\}} \Delta_n^{\text{type}}(t_2, t_1; k_{\alpha_s}, k_f; \alpha_s, f). \tag{14}$$

They share the common form

$$\Delta_n^{\text{type}}(t_2, t_1; k_{\alpha_s}, k_f; \alpha_s, f) = \exp\left(-\sum_{ij} \sum_k \int_{t_2}^{t_1} dt \int_{z^-}^{z^+} dz \alpha_s(k_{\alpha_s} t) K'_{ij,k}(t, z) \frac{f_{c'}(\frac{\eta_c}{x}, k_f t)}{f_c(\eta_c, k_f t)}\right), \tag{15}$$

wherein the kinematics of the splitting are given by the default choice for $t = Q^2 y z(1 - z)$ in the massless case while the $K'_{ij,k}(t, z)$ denote the coupling and PDF stripped splitting kernels incorporating the remaining pieces of the $K_{ij,k}$ and the Jacobian J of the phase-space parametrisation. The precise definitions of the variables for each dipole type are given in Table 2. It directly follows that for FF-type dipole splittings the ratio of PDFs is simply unity. Eq. (15) further details the dependence on the α_s and PDF scale factors k_{α_s} and k_f . These multiplicative factors as well as their variations are assumed to be of order one, such that they do not induce spurious large logarithms. The generalisation to the massive case is straightforward and only involves generalised definitions of t, x, y and z , cf. [23].

Table 2 Definition of the evolution and splitting variables for each dipole type. The fifth column lists the splitting process as seen from the Born process, c and c' refer to the flavour of the initial state before and after the splitting process, respectively. The variables $y_{ij,k}, \tilde{z}_i, x_{ij,a}, x_{jk,a}, x_{j,ab}, u_j$ and v_j are defined in [23,36,37]

Type	z	y	x	$(ij, k) \rightarrow (i, j, k)$	c, c'
FF	\tilde{z}_i	$y_{ij,k}$	1	$(ij, k) \rightarrow (i, j, k)$	a, a
FI	\tilde{z}_i	$\frac{1 - x_{ij,a}}{x_{ij,a}}$	$x_{ij,a}$	$(ij, a) \rightarrow (i, j, a)$	a, a
IF	$x_{jk,a}$	$\frac{u_j}{x_{jk,a}}$	$x_{jk,a}$	$(aj, k) \rightarrow (a, j, k)$	aj, a
II	$x_{j,ab}$	$\frac{\tilde{v}_j}{x_{j,ab}}$	$x_{j,ab}$	$(aj, b) \rightarrow (a, j, b)$	aj, a

3.2 Reweighting trial emissions

To numerically integrate Sudakov form factors typically the Sudakov Veto Algorithm is used [50–55]. Therein the integrands K found in the Sudakov form factors are replaced with integrable overestimates \hat{K} . This is balanced by only accepting a proposed emission with probability $P_{\text{acc}} = K/\hat{K}$. A multiplicative factor in K is therefore equivalent to a multiplicative factor in P_{acc} [52]. This observation is for example used to apply matrix-element corrections [54], where the splitting kernels are replaced with a real-emission-like kernel R/B . This is done a-posteriori, i.e. the event weight is multiplied by $(R/B)/K$, the emission itself is unchanged. The same method is also used in the Vincia parton shower to calculate uncertainty variations for different scales, finite terms of the antenna functions, ordering parameters and sub-leading colour corrections [28]. Here we employ this technique to account for variations of the strong-coupling parameter and the PDFs in the shower evolution of LO and NLO QCD matrix elements.

As has been laid out in the previous section, the emission kernels K depend linearly on α_s and on a ratio of parton densities $f_{c'}(\eta_c/x, k_f t)/f_c(\eta_c, k_f t)$. A change of PDFs $f \rightarrow \tilde{f}$, the strong coupling $\alpha_s \rightarrow \tilde{\alpha}_s$ and the scale prefactors entering both, i.e. $k_{\alpha_s} \rightarrow \tilde{k}_{\alpha_s}$ and $k_f \rightarrow \tilde{k}_f$, is equivalent to modifying the emission probability accordingly:⁴

$$P_{\text{acc}} \rightarrow q_{\text{acc}} P_{\text{acc}}, \quad q_{\text{acc}} \equiv \frac{\tilde{\alpha}_s(\tilde{k}_{\alpha_s} t)}{\alpha_s(k_{\alpha_s} t)} \frac{\tilde{f}_{c'}(\frac{\eta_c}{x}, \tilde{k}_f t)}{f_{c'}(\frac{\eta_c}{x}, k_f t)} \frac{f_c(\eta_c, k_f t)}{\tilde{f}_c(\eta_c, \tilde{k}_f t)}, \tag{16}$$

where the scale dependence and the definition of η_c and x can be read off the Sudakov form factors given in Eq. (15) and

⁴ Although the emission scales can not be reweighted themselves using the presented method, the input scales of the strong coupling and the PDFs can be changed, as indicated in the text. We focus on constant prefactors here, but the functional form can also be changed, although the overall functional form of $k_{\alpha_s} t$ should be restricted to the CMW-like rescaling [56].

Table 2. In case of FF dipoles Eq. (16) simplifies significantly as the ratios of PDF factors reduces to unity. It further follows, that the event weight for each accepted emission needs to be multiplied by the corresponding factor q_{acc} in order to incorporate the new choice of α_s , PDFs and the scales they are evaluated at. Accordingly, the probability to reject an emission is changed to

$$P_{\text{rej}} = 1 - P_{\text{acc}} \rightarrow 1 - q_{\text{acc}} P_{\text{acc}}$$

$$= \left[1 + (1 - q_{\text{acc}}) \frac{P_{\text{acc}}}{1 - P_{\text{acc}}} \right] P_{\text{rej}} \equiv q_{\text{rej}} P_{\text{rej}}. \quad (17)$$

Consequently, for each rejected emission the event weight receives a corrective weight of q_{rej} . Proofs that this treatment indeed results in the correct Sudakov form factors can be found in [26,27,52].

3.3 Next-to-leading-order matching

To match NLO QCD parton-level calculations with subsequent parton-shower evolution SHERPA employs a variant of the original MC@NLO algorithm presented in [12], referred to as S- MC@NLO [54]. Schematically, such a S- MC@NLO calculation has the following structure:

$$\langle O \rangle^{\text{NLOPS}} = \int d\Phi_B \left[B(\Phi_B) + \text{VI}(\Phi_B) \right. \\ \left. + \int dx'_{a/b} \text{KP}(\Phi_B, x'_{a/b}) \right. \\ \left. + \sum_j \int d\Phi_1^j (D_{A,j} - D_{S,j}) (\Phi_B \cdot \Phi_1^j) \right] \\ \times \text{PS}_{\text{NLOPS}}(O, \Phi_B) \\ + \int d\Phi_R \left[R(\Phi_R) - \sum_j D_{A,j}(\Phi_{B,j} \cdot \Phi_1^j) \right] \\ \times \text{PS}(O, \Phi_R) \\ = \int d\Phi_B \bar{B}(\Phi_B) \text{PS}_{\text{NLOPS}}(O, \Phi_B) \\ + \int d\Phi_R H_A(\Phi_R) \text{PS}(O, \Phi_R). \quad (18)$$

Here the real-emission contribution R of the NLO calculation has effectively been split into an infrared-singular (soft) and an infrared-regular (hard) part, the resummation kernel D_A and the finite hard remainder H_A , respectively, such that $R = D_A + H_A$ [54,57]. The \bar{B} -function has the following explicit parameter dependences

$$\bar{B}(\Phi_B) \equiv \bar{B}(\Phi_B; \alpha_s, f; \mu_R, \mu_F) \\ = B(\Phi_B; \alpha_s, f; \mu_R, \mu_F) + \text{VI}(\Phi_B; \alpha_s, f; \mu_R, \mu_F) \\ + \int dx'_{a/b} \text{KP}(\Phi_B, x'_{a/b}; \alpha_s, f; \mu_R, \mu_F)$$

$$+ \sum_j \int d\Phi_1^j (D_{A,j} - D_{S,j}) \\ \times (\Phi_B \cdot \Phi_1^j; \alpha_s, f; \mu_R, \mu_F). \quad (19)$$

From the perspective of parameter reweighting, the resummation kernel D_A behaves the same way as the subtraction term D_S . In fact, in our reweighting implementation the $(D_A - D_S)$ contribution is treated as a single term, as indicated. It is only to note that their PDFs are evaluated at the partonic momentum fraction $x_{a/b,j}$ and external flavours a_j and b_j of their $\Phi_B \cdot \Phi_1^j$ phase-space configuration rather than those of Φ_B . The other parts of the \bar{B} -function can then be reweighted as described in Sect. 2.2, leading to

$$\bar{B}(\Phi_B; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_R, \tilde{\mu}_F) \\ = \tilde{\alpha}_s^n(\tilde{\mu}_R^2) \tilde{f}_a(x_a, \tilde{\mu}_R^2) \tilde{f}_b(x_b, \tilde{\mu}_R^2) \\ \times \left[B'(\Phi_B) + \tilde{\alpha}_s(\tilde{\mu}_R^2) \left(\text{VI}'(\Phi_B) + c_R^{(0)} l_R + \frac{1}{2} c_R^{(1)} l_R^2 \right) \right. \\ \left. + \tilde{\alpha}_s(\tilde{\mu}_R^2) \int dx'_{a/b} \text{KP}'_j(\Phi_j, x'_{a/b}; \tilde{f}; \tilde{\mu}_{F,\text{core}}) \right] \\ + \sum_j \int d\Phi_1^j \tilde{f}_{a_j}(x_{a,j}, \tilde{\mu}_R^2) \tilde{f}_{b_j}(x_{b,j}, \tilde{\mu}_R^2) \tilde{\alpha}_s^{n+1}(\tilde{\mu}_R^2) \\ \times [D'_{A,j} - D'_{S,j}] (\Phi_B \cdot \Phi_1^j). \quad (20)$$

The H_A -function then transforms as

$$H_A(\Phi_R; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_R, \tilde{\mu}_F) \\ = R(\Phi_R; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_R, \tilde{\mu}_F) \\ - \sum_j D_{A,j}(\Phi_{B,j} \cdot \Phi_{1,j}^j; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_{R,j}, \tilde{\mu}_{F,j}) \\ = \tilde{\alpha}_s^{n+1}(\tilde{\mu}_R^2) \tilde{f}_a(x_a, \tilde{\mu}_R^2) \tilde{f}_b(x_b, \tilde{\mu}_R^2) R'(\Phi_R) \\ - \sum_j \tilde{\alpha}_s^{n+1}(\tilde{\mu}_{R,j}^2) \tilde{f}_a(x_a, \tilde{\mu}_{F,j}^2) \tilde{f}_b(x_b, \tilde{\mu}_{F,j}^2) \\ \times D'_{A,j}(\Phi_{B,j} \cdot \Phi_{1,j}^j), \quad (21)$$

wherein each subtraction term $D_{A,j}$ has its own scales $\mu_{R,j}$, $\mu_{F,j}$ defined on its underlying Born configuration $\Phi_{B,j}$. Writing Eq. (18) as a Monte-Carlo sum over events with B-like and R-like structure, which are conventionally called S and H events in MC@NLO calculations, and with $N = N_S + N_H$, we obtain

$$\langle O \rangle^{\text{NLOPS}} = \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \left\{ \sum_{i=1}^{N_S} \bar{B}(\Phi_{B,i}) \text{PS}_{\text{NLOPS}}(O, \Phi_{B,i}) \right. \\ \left. + \sum_{i=1}^{N_H} H_A(\Phi_{R,i}) \text{PS}(O, \Phi_{R,i}) \right\}. \quad (22)$$

Thus, under $\mu_R \rightarrow \tilde{\mu}_R, \mu_F \rightarrow \tilde{\mu}_F, \alpha_s \rightarrow \tilde{\alpha}_s$ and $f \rightarrow \tilde{f}$ both \bar{B} of the S-events and H_A of the H-events transform as composite objects in terms of their constituents, as defined above. This leaves the S- MC@NLO parton shower, PS_{NLOPS} , defined through

$$\begin{aligned}
 PS_{NLOPS}(O, \Phi_B) &\equiv PS_{NLOPS}(O, \Phi_B; k_{\alpha_s}, k_f; \alpha_s, f; \mu_Q^2) \\
 &= \bar{\Delta}_n(t_{IR}, t') O(\Phi_n) \\
 &\quad + \int_{t_{IR}}^{t'} d\Phi_1 \frac{D_A(\Phi_B \cdot \Phi_1)}{B(\Phi_B)} \\
 &\quad \times \bar{\Delta}_n(t, t') PS(O, \Phi_{n+1}; t). \tag{23}
 \end{aligned}$$

It differs from the usual PS of Eq. (12) with respect to the splitting kernel for the first emission and the associated definition of the Sudakov form factor, cf. [54,58]. However, for the purpose of reweighting, all trial emissions can be treated in the same way as in the standard parton shower as the parameter dependences are identical.

3.4 Validation

To validate the reweighting of scale and parameter dependences in CSSHOWER and S- MC@NLO calculations within the SHERPA framework we perform closure tests between reweighting results and dedicated simulations.

Our implementation allows to constrain the maximum number of reweighted shower emissions per event. For a pure leading-order parton-shower run or H-like events in S- MC@NLO calculations this amounts to setting $n_{PS} \in \{0, 1, 2, \dots, \infty\}$. When considering S- MC@NLO simulations in addition the parameter $n_{NLOPS} \in \{0, 1\}$ can be used to disable the reweighting of the $\mathcal{O}(\alpha_s)$ emission for S-events.

Of course the reweighting result will only coincide with a dedicated calculation if all emissions are reweighted, i.e. $n_{NLOPS} = 1$ and $n_{PS} = \infty$. However, by subsequently enabling the reweighting of more and more emissions the relevance of their dependences for the determination of the full uncertainty can be studied. A finite value of n_{PS} can also be useful in production, if the effect of reweighting higher-order emissions becomes negligible. The reduced amount of reweighting per event then allows for a faster event generation. An additional benefit would be that rare high-multiplicity shower histories do not spoil the statistical convergence of the reweighted result, even if their exact kinematics might be irrelevant for the studied observable.

The final-state only case: Thrust in $e^+ e^- \rightarrow q\bar{q}$ events

To validate LOPS reweighting, we consider two observables, which are complementary in their sensitivity to parton-shower emissions. At first, we consider the event-shape variable thrust T [59] in hadronic events in e^+e^- -collisions at

$\sqrt{s} = 91.2$ TeV. In this case QCD emissions are restricted to the final-state. Accordingly, there appear no PDF factors in the shower reweighting, cf. Eq. (16), and thus no factorisation scale dependence. Moreover, as we consider the leading-order matrix element for $e^+e^- \rightarrow q\bar{q}$ only, the renormalisation scale is also absent in the hard-process component. Therefore, we can concentrate on the pure α_s uncertainty in the parton shower here. Leaving the perturbative order of its running invariant it is defined by its value at the input scale m_Z .

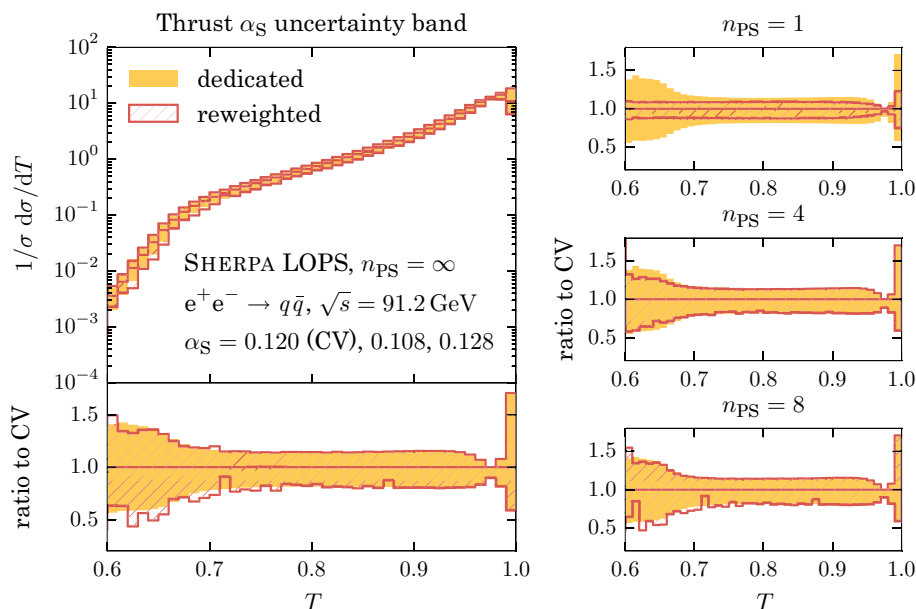
In Fig. 2, we compare α_s uncertainty bands generated by reweighting the nominal prediction with the one generated by dedicated predictions for each variation. As in Fig. 1, the uncertainty band is defined as the envelope over the distributions with different $\alpha_s(m_Z^2)$ input values. The nominal value is taken as $\alpha_s(m_Z^2) = 0.120$, and its up/down variations are 0.128 and 0.108, respectively. Reweighting bands are presented for $n_{PS} = 1, 4, 8, \infty$. The $n_{PS} = 1$ band underestimates the uncertainty, especially for $T \leq 2/3$, where multiple hard emissions are required, and for $T \approx 1$, the region sensitive to multiple soft emissions. For $n_{PS} = 4$, the uncertainty is underestimated only for bins with $T \leq 2/3$, and less so than for the $n_{PS} = 1$ case. The difference between the two choices of $n_{PS} = 8$ and ∞ is merely statistical and both reproduce the dedicated result very accurately.

However, for low values of T , the statistical fluctuations of the reweighting results with higher n_{PS} grow larger, corresponding to a widening of the distribution of reweighting factors. Low values of the thrust observable correspond to the emission of several hard partons, which is less probable in the parton-shower approximation, and more appropriately modelled in multijet-merged calculations, cf. Sect. 4. In this phase-space region it is difficult for the reweighting to compensate the multitude of accepted *soft* emissions off these hard legs, that turn unstable for $P_{acc} \rightarrow 1$, with rejected ones, cf. Eq. (17). This issue can be addressed by introducing a prefactor for the over-estimator function \hat{K} in the reweighting runs, to ensure that P_{acc} does not approach 1, cf. [26,27,52]. This renders the Sudakov Veto Algorithm somewhat less efficient, but is shown to reduce statistical fluctuations in the reweighting.

The initial-state dominated case: p_{\perp}^W in $pp \rightarrow W[e\nu]$ events

The second observable considered to validate our CSSHOWER and S- MC@NLO reweighting implementation is the W-boson transverse-momentum distribution p_{\perp}^W in 13 TeV proton-proton collisions, that has already been used in Sect. 2.3 in the NLO case. The definitions for constructing the uncertainty bands used there are kept the same, and are stated in Table 1. We now use the CT14nlo PDF set, which uses a

Fig. 2 Uncertainty band for the Thrust event shape in dijet production in e^-e^+ annihilation for a variation of α_s . The *left panel* shows the nominal distribution and the ratio to the central value. The uncertainty band calculated using reweighting (including all emissions, i.e. $n_{\text{PS}} = \infty$) is compared to the one obtained from dedicated calculations. The comparison is repeated in the *three panels on the right* for different maximum number of reweighted emission n_{PS}



Hessian error representation at a 90% confidence level [60].⁵ Therefore the PDF error band will be larger than before, as it now corresponds to nearly two standard deviations instead of only one.

Considering a hadronic environment, initial-state emissions are present, which means that our reweighting factors now include PDF double ratios. In Fig. 3, we compare LOPS uncertainty bands for scale, α_s and PDF variations, including comparisons between reweighted and dedicated predictions for a varying maximum number of reweighted shower emissions n_{PS} . Before discussing the bands, we observe that the tail of the p_{\perp}^W spectrum is not populated, in particular in comparison to Fig. 1. This is expected, as the LO configuration is restricted to $p_{\perp}^W = 0$, such that all other bins are filled through recoils against parton-shower emissions only. However, the phase space of parton-shower emissions is restricted to the soft region, and therefore the W boson can not build up a large recoil.

We now turn to the scale uncertainty band – which is entirely due to factorisation scale variations, because the LO matrix element is independent of α_s , and therefore the band underestimates the perturbative uncertainty. We also observe that the band is nearly flat. As we vary only the scales of the matrix-element calculation, the constant spread corresponds to the factorisation-scale uncertainty of the Born configuration at $p_{\perp}^W = 0$, merely propagating to higher p_{\perp}^W bins through the parton shower, which is unaware of the scale

variations. In the matrix-element reweighting, we can guarantee the same phase-space points as in the dedicated run, such that we see perfect agreement between dedicated and reweighted predictions. We therefore omit comparisons for different n_{PS} for the scale-uncertainty band.

Looking at the α_s uncertainty band, we can see that the envelope constricts at the position of the peak of the distributions. This reflects that the variation of α_s shifts the position of the peak, such that variations that are below the nominal distribution on the left side of the peak, are exceeding the nominal distribution on the right side, and vice versa. Comparing the reweighted prediction to the dedicated one, we find a flat band for $n_{\text{PS}} = 0$, corresponding to restricting the reweighting to the fixed-order matrix element. As the LO calculation is independent of α_s , this only reflects the change of the PDFs, which are fitted to $\alpha_s(m_Z)$. The reproduction of the shape of the α_s uncertainty improves a lot when reweighting up to one emission ($n_{\text{PS}} = 1$), and slightly more when adding another emission on top ($n_{\text{PS}} = 2$).

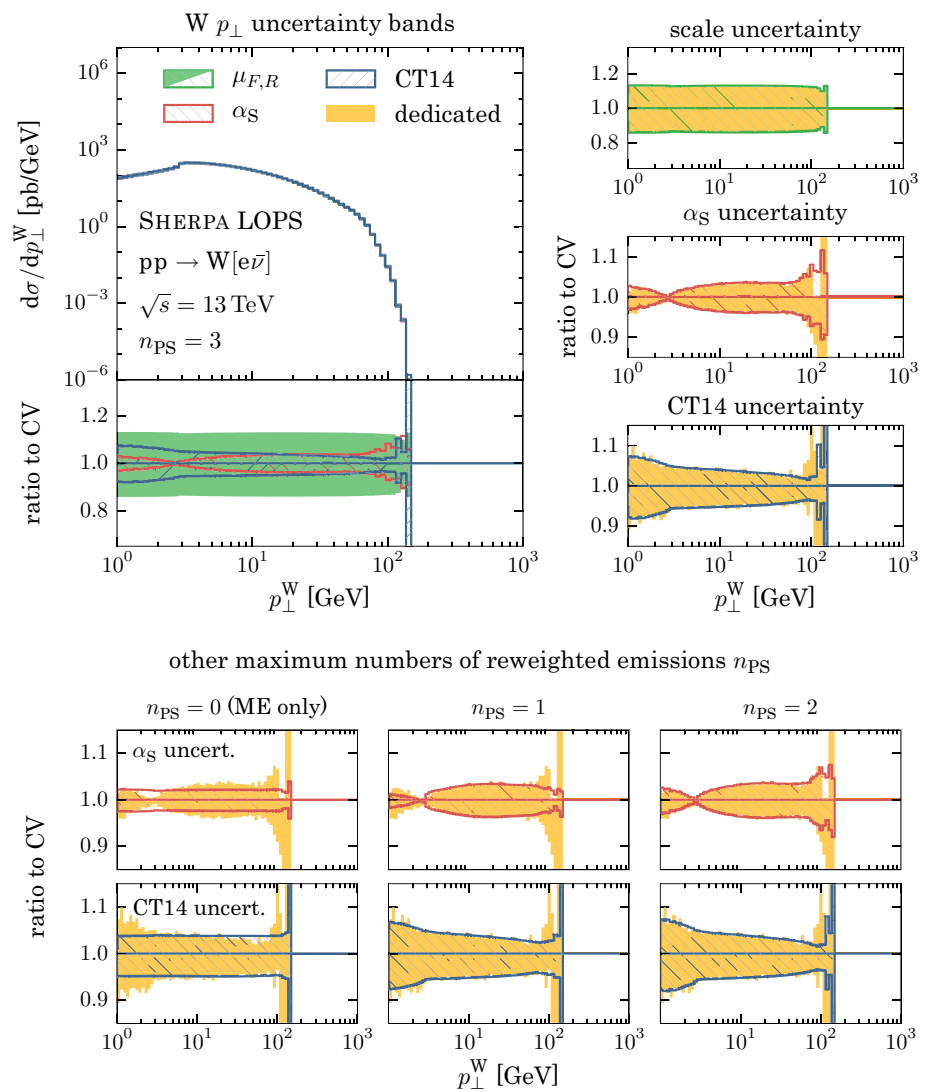
For the PDF uncertainty, we see that the reweighting with $n_{\text{PS}} = 0$ underestimates it by at least 1–5% for small transverse momenta, and overestimate it around the W mass. As for the α_s uncertainty, this improves for $n_{\text{PS}} = 1, 2$.

The last depicted step, i.e. $n_{\text{PS}} = 3$, on the other hand, does not contribute further to the reproduction of the α_s and PDF uncertainties. No significant differences with respect to the $n_{\text{PS}} = 2$ case is observed. It can be concluded that it is sufficient to include up to two emissions to reproduce the uncertainty bands for this observable.

This is to be expected, as the gauge boson recoils against the shower emissions and is therefore mostly affected by the few hardest branchings. These mainly originate from the incoming hard virtual partons, so the generally softer

⁵ The reason for switching from NNPDF to CT14 PDFs is the strict positivity of the latter. The CSSHOWER rejects emissions when negative PDF values are involved, a behaviour which spoils the reweighting in regions where the original and the target PDF do not have the same sign. The deviations seem to be small in practical applications, but here we chose to establish closure in a clean context first.

Fig. 3 The same as in Fig. 1, but for LO + parton-shower (PS) generation. The uncertainty bands are calculated by reweighting the ME and up to n_{PS} shower emissions. In the upper four plots, $n_{PS} = 3$. In the lower plots, n_{PS} is varied for comparison. The scale uncertainties do not change with n_{PS} and are therefore not repeated



final-state emissions barely contribute. Although we do not reproduce this here, we confirmed this by entirely disabling final-state emissions, which showed no effect on the results.

In Fig. 4 we present the validation of NLOPS predictions for the W-boson transverse-momentum distribution. Overall, we get a similar picture as in the LOPS case. The main differences are the increased high- p_{\perp}^W reach and the significantly smaller scale uncertainties in the low p_{\perp}^W range, a consequence of including the complete set of $\mathcal{O}(\alpha_s)$ corrections to the production process. For large p_{\perp}^W , the uncertainty increases again, because we fall back to a LO description again: *Only* the $2 \rightarrow 3$ matrix element still contributes.

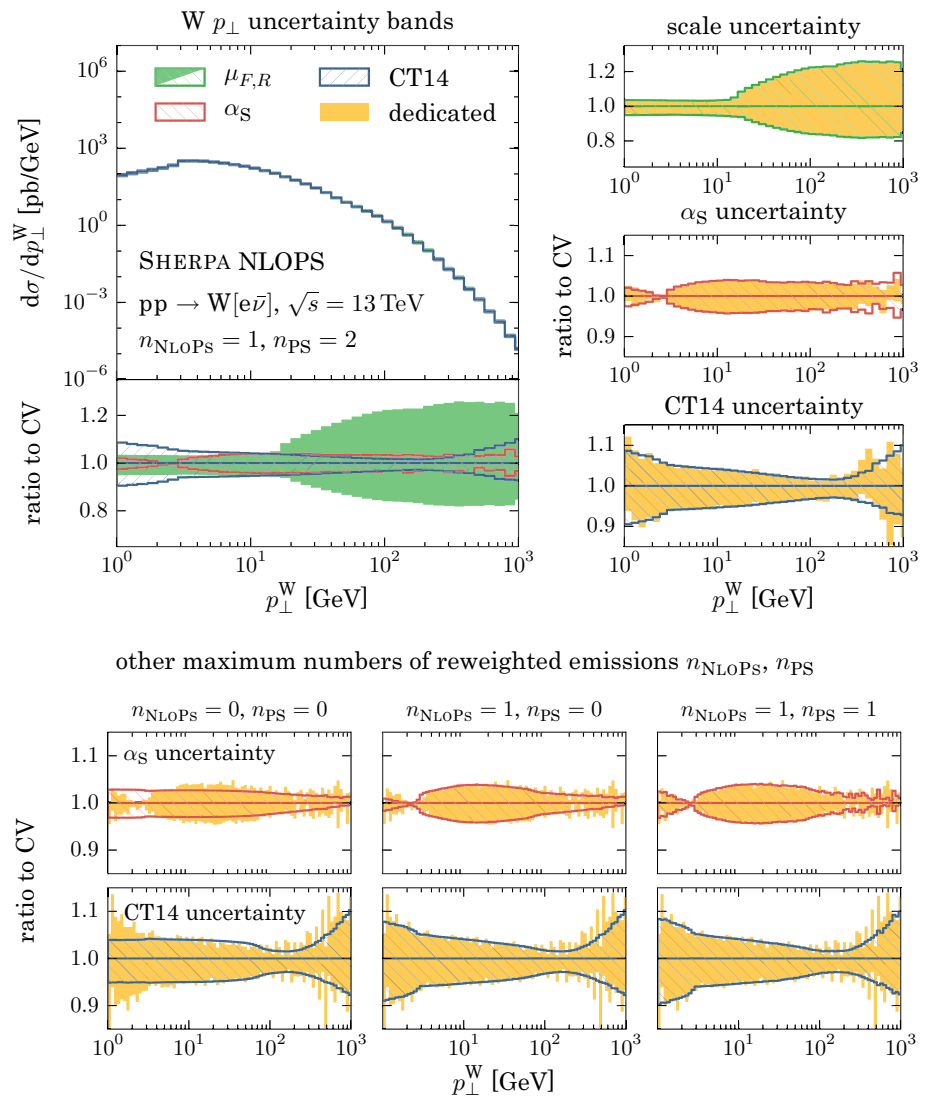
To assess the quality of the reweighting, we consider again different settings for the parameters n_{NLOPS} and n_{PS} . Assuming $n_{NLOPS} = n_{PS} = 0$, only the scale variations of the hard process are considered and the parton-shower contribution to the $\mathcal{O}(\alpha_s)$ correction is *not* reweighted. Furthermore, we present results for $n_{NLOPS} = 1$ and $n_{PS} = 0, 1, 2$. With these settings, the $\mathcal{O}(\alpha_s)$ corrections get properly reweighted,

but the number of subsequent shower emissions off the S- and H-like events treated correctly is varied. We observe a saturation for reproducing the dedicated calculations at $n_{NLOPS} + n_{PS} \geq 2$, with no further improvement when n_{PS} is increased from 1 to 2. This confirms the findings made when considering the LOPS setup in Fig. 3: the gauge-boson transverse-momentum distribution is dominated by the few hardest emissions.

4 Reweighting multijet-merged calculations

In this section we address the reweighting of multijet-merged event generation runs. These approaches allow to combine LO or NLO QCD matrix elements of different multiplicity dressed with parton showers into inclusive samples. Accordingly, the production of jets associating a given core process can be modelled through exact matrix elements rather than relying on the logarithmic approximation of the par-

Fig. 4 The same as in Fig. 3, but for NLO + parton-shower (PS) generation. The uncertainty bands are calculated by reweighting the ME and a maximum number of emissions from the MC@NLO (n_{NLOPS}) and the ordinary PS (n_{PS}). n_{NLOPS} is constrained to 0 or 1, as the MC@NLO prescription only affects the first emission. In the upper four plots we consider $n_{\text{NLOPS}} = 1$ and $n_{\text{PS}} = 2$, thus up to three emissions get reweighted. In the lower plots, we consider variations of n_{NLOPS} and n_{PS}



ton shower only. In particular, when considering hard jet kinematics or angular correlations such techniques prove to be indispensable to properly describe experimental observations, see for instance [61–64].

To first approximation the reweighting as described in the previous sections can be used without change, only that the perturbative order p is no longer a constant across the sample, but varies for each event, corresponding to the considered matrix-element parton multiplicity. However, there are also new algorithm-specific intricacies which complicate the dependence on the input parameters and need to be dealt with to allow for a consistent reweighting. The LO and NLO merging techniques employed within the SHERPA framework are presented in [21,65,66], respectively. They rely on the reconstruction of parton-shower histories for multiparton amplitudes that set the parton-shower initial conditions for their subsequent evolution. This is achieved by running a *backward-clustering* algorithm that identifies a corresponding core process and calculates hard-parton splitting

scales that serve as predetermined shower branchings. In the SHERPA approach the actual parton shower then starts off the reconstructed core process and implements the predetermined hard splittings based on a *truncated shower*. Furthermore, it is the purpose of the truncated-shower evolution to implement possible Sudakov vetoes for shower emissions above the phase-space separation or merging scale Q_{cut} . It should be emphasised here, parton-shower reweighting is vital when using modified input parameters in order to cancel the Q_{cut} dependence to the accuracy of the parton shower. In case only the hard-process matrix element parameters get reweighted, the dependence on Q_{cut} is cancelled to leading-logarithmic accuracy only, however, residual subleading contributions from the running coupling or the PDF evolution remain [65].

For the reweighting of the truncated-shower Sudakov veto probability the methods described in Sect. 3 can be applied. In what follows we will detail the specifics of the reweighting procedure for LO and NLO multijet-merging runs with




Fig. 5 Possible parton-shower histories of a $qg \rightarrow Zgqq'\bar{q}'$ matrix element allowing only QCD splittings (*left*) and also including electroweak splittings (*right*)

SHERPA supplemented by an extensive validation of the implementation.

4.1 Preliminaries

Common to the LO and NLO merging techniques used in SHERPA, cf. [21, 22, 65–67], is the separation of the emission phase space into a soft and a hard region, defined through a suitable m -parton measure Q_m and a separation criterion Q_{cut} . For each parton configuration Φ_m with $Q_m > Q_{\text{cut}}$ a shower history that represents the event as a core process with subsequent $1 \rightarrow 2$ shower splittings is probabilistically build through backward clustering. The resulting sequence of cluster steps is characterised by tuples $\{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}$, recording the (possibly changing) initial-state flavours and momentum fractions as well as the evolution variable of each splitting. We allow for both QCD and EW splitting functions [21, 68] to identify such splitting processes and veto recombinations that would lead to the reduction of configurations which are not present in the matrix elements.⁶ Figure 5 details possible cluster histories for a given $pp \rightarrow Z + 4\text{jets}$ configuration, depending on its kinematics, allowing for QCD splittings only (*left*) or both QCD and EW splittings (*right*).

The sequence $\{t_i\}$ of reconstructed branching scales then may be either ordered or unordered, with an ordered history satisfying $t_j < t_{j-1} < \dots < t_1 < t_0 = \mu_{F,\text{core}}^2$. The recombination probabilities in each clustering step are determined by the forward-splitting probabilities and are therefore dependent on the parton shower and its parameters and choices. This is reflected, step-by-step, in the addition of one factor of α_s (when appropriate) at the reconstructed splitting scale, a ratio of PDFs at the reconstructed initial flavours and their momentum fractions, and a Sudakov form factor describing the evolution of each step.

In the SHERPA implementations the α_s and PDF factors are added explicitly onto the respective matrix elements and can therefore be reweighted directly. The Sudakov form factor, on the other hand, is implemented through a vetoed truncated parton shower [21, 22]. The truncated shower itself,

⁶ An example here is the interpretation of a $e^+e^- \rightarrow g\bar{d}\bar{d}$ configuration. Its matrix element does not contain terms/diagrams that allow the quark-antiquark pair to be clustered.

accounting for the possibility of soft parton-shower emissions between subsequent reconstructed hard emissions, i.e. with $t_m < t < t_{m-1}$ but $Q < Q_{\text{cut}}$, can be reweighted with the methods described in Sect. 3. If, however, an emission with $Q > Q_{\text{cut}}$ occurs the event is vetoed. Practically, this is accounted for through increasing n_{trial} of the next accepted event by n_{trial} of the vetoed event. Thus, n_{trial} becomes dependent on the parton-shower parameters.

A special remark concerning unordered histories, i.e. histories whose sequence of $\{t_i\}$ has at least one pair $t_k \geq t_{k-1}$, is in order. Such histories can be encountered in various configurations, e.g. when the last clustering step produces a splitting scale larger than the nominal starting scale of the core process⁷ or the flavour structure only allows further clusterings at scales t_{k-1} lower than the last identified one t_k .⁸ As such configurations cannot be generated by a strictly ordered parton shower, for each unordered step neither the accompanying PDF ratio nor Sudakov factor is therefore present in the calculation. More than one unordering in a cluster history of a given event is possible and in fact likely at high multiplicities. PDF ratios and Sudakov factors then of course only occur in the ordered subhistories in between the unorderings. For the sake of clarity and brevity we will omit this case from the discussion of the following subsections. Its implications to the algorithm, and therefore to the reweighting, are straightforward. If ordered histories are enforced, core configurations beyond the standard $2 \rightarrow 2$ processes occur. Independent of the presence and number of unorderings, the renormalisation and factorisation scales μ_R and μ_F are always set in the way, as will be detailed below.

4.2 The leading-order case

We start the discussion with the simplest case, where all matrix elements used in the merging are given at leading

⁷ An example here is the interpretation of a $gq \rightarrow Zq$ configuration. In regions of large transverse momenta of the final state parton its identified branching scale t_1 is larger than the starting scale t_0 of the core process $q\bar{q} \rightarrow Z$, usually defined as the Z virtuality.

⁸ An example here is the interpretation of a $e^+e^- \rightarrow d\bar{d}u\bar{u}$ configuration. In a first step there are only two choices to cluster, resulting in an identified branching scale t_2 . There now is a finite region in phase space where the gluon can (only) be clustered with scale $t_1 < t_2$.

order. A LO multijet-merged calculation, with Born matrix elements at $\mathcal{O}(\alpha_s^{n+j})$, containing j additional partons relative to the core process, has the following structure

$$\begin{aligned} \langle O \rangle^{\text{MEPS@LO}} &= \sum_{j=0}^{j_{\max}} \int d\Phi_j B_j^{\text{merge}}(\Phi_j) \Theta(Q_j - Q_{\text{cut}}) \\ &\quad \times PStv(O, \Phi_j) \\ &= \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \sum_{i=1}^N \sum_{j=0}^{j_{\max}} B_j^{\text{merge}}(\Phi_{j,i}) \\ &\quad \times \Theta(Q_j - Q_{\text{cut}}) \text{PS}^{\text{vt}}(O, \Phi_{j,i}). \end{aligned} \tag{24}$$

Note that Φ_j here denotes the entire final-state phase space of the process, including all particles of the core process. As before, Q_j is a suitable infrared-safe distance measure of Φ_j . The Θ -function thus realises a minimum separation of Q_{cut} and acts as an infrared regulator. PS^{vt} is the vetoed truncated parton shower derived from Eq. (12). As the limit in the second line is well defined, it can be transposed with the summation over parton multiplicities. As the ingredient leading-order matrix elements need to incorporate the soft-collinear resummation properties of the parton shower, they have the following parameter dependences:

$$B_j^{\text{merge}}(\Phi_j) \equiv B_j^{\text{merge}}(\Phi_j; \alpha_s, f; \mu_{R,\text{core}}, \mu_{F,\text{core}}, k_{\alpha_s}, k_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}). \tag{25}$$

The cluster steps $\{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}$ denote the identified cluster history of the configuration Φ_j , as discussed above. Therein, the a_i, b_i are the possibly changing initial-state flavours, the $x_{a,i}, x_{b,i}$ their momentum fractions, and the t_i are the reconstructed values of the parton-shower evolution variable at each splitting. Together with the α_s and PDF scale prefactors k_{α_s}, k_f of the parton shower, the cluster steps relate B_j^{merge} to the scale and PDF stripped Born matrix element B'_j encountered in Sect. 2.1,

$$\begin{aligned} B_j^{\text{merge}}(\Phi_j; \alpha_s, f; \mu_{R,\text{core}}, \mu_{F,\text{core}}, k_{\alpha_s}, k_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) \\ = \prod_{i=1}^j \frac{f_{a_i}(x_{a,i}, k_f t_i)}{f_{a_{i-1}}(x_{a,i-1}, k_f t_i)} f_{a_0}(x_{a,0}, \mu_{F,\text{core}}^2) \\ \times \prod_{i=1}^j \frac{f_{b_i}(x_{b,i}, k_f t_i)}{f_{b_{i-1}}(x_{b,i-1}, k_f t_i)} f_{b_0}(x_{b,0}, \mu_{F,\text{core}}^2) \\ \times \alpha_s^{n+j} (\mu_R^2) B'_j(\Phi_j). \end{aligned} \tag{26}$$

In this notation, the core scale is $t_0 = \mu_{F,\text{core}}^2$, it is therefore not multiplied by the prefactors of the parton shower. The partonic momentum fractions of the core process are $x_{a,0}, x_{b,0}$.

The scales of each single α_s within the cluster history vary, but an effective global renormalisation scale can be defined through

$$\alpha_s^{n+j} (\mu_R^2) = \alpha_s^{n+e} (\mu_{R,\text{core}}^2) \prod_{i=1}^j \alpha_s^{1-\epsilon_i} (k_{\alpha_s} t_i), \tag{27}$$

where $\epsilon_i = 0$ if the identified splitting process at branching i is of QCD-type, and 1 otherwise, $e = \sum_{i=1}^j \epsilon_i$. To consistently vary the μ_R scale, we consider variations of the splitting scales t_i and the core scale $\mu_{R,\text{core}}$ on the right-hand side by a common factor, solving for the prefactor of the effective μ_R to be used in the matrix-element calculation. Thus, while up to NLO accuracy μ_R is varied by the same common factor, the full solution of this procedure results in slightly larger variations of the effective renormalisation scale.

Apart from the Sudakov form factors the soft-collinear structure of the B_j^{merge} is now identical to the emission of j partons off a B_0 configuration with the parton shower described in Sect. 3. In case of final-state splittings the ratio of parton distribution functions is simply unity as neither the partonic $x_{a/b,i}$ and $x_{a/b,i-1}$ nor the initial-state flavours a_i, b_i and a_{i-1}, b_{i-1} differ. In principle, with every ratio of PDFs there is also a ratio of flux factors. However, all such factors cancel except for the outermost ones, corresponding to Φ_j and, hence, are regarded as part of B'_j .

Changing the scales $\mu_{R,\text{core}} \rightarrow \tilde{\mu}_{R,\text{core}}, \mu_{F,\text{core}} \rightarrow \tilde{\mu}_{F,\text{core}}, k_{\alpha_s} \rightarrow \tilde{k}_{\alpha_s}, k_f \rightarrow \tilde{k}_f$ as well as $\alpha_s \rightarrow \tilde{\alpha}_s$ and $f \rightarrow \tilde{f}$ results in

$$\begin{aligned} B_j^{\text{merge}}(\Phi_j; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_{R,\text{core}}, \tilde{\mu}_{F,\text{core}}, \tilde{k}_{\alpha_s}, \tilde{k}_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) \\ = \prod_{i=1}^j \frac{\tilde{f}_{a_i}(x_{a,i}, \tilde{k}_f t_i)}{\tilde{f}_{a_{i-1}}(x_{a,i-1}, \tilde{k}_f t_i)} \tilde{f}_{a_0}(x_{a,0}, \tilde{\mu}_{F,\text{core}}^2) \\ \times \prod_{i=1}^j \frac{\tilde{f}_{b_i}(x_{b,i}, \tilde{k}_f t_i)}{\tilde{f}_{b_{i-1}}(x_{b,i-1}, \tilde{k}_f t_i)} \tilde{f}_{b_0}(x_{b,0}, \tilde{\mu}_{F,\text{core}}^2) \\ \times \tilde{\alpha}_s^{n+j} (\tilde{\mu}_R^2) B'_j(\Phi_j). \end{aligned} \tag{28}$$

The scale $\tilde{\mu}_R^2$ is now calculated from Eq. (27) using $\tilde{\mu}_{R,\text{core}}^2$ and \tilde{k}_{α_s} as input.

Equation (28) describes what happens to the matrix-element part of a multijet-merged calculation. This leaves the vetoed truncated shower PS^{vt} . While the truncated and standard shower part is described in Sect. 3, the vetoed shower leads to vetoed events. As vetoed events correspond to events whose weights have been set to zero, their description is equivalent to increasing the number of trials, n_{trial} , by one. Thus, when varying the parameters of the parton shower, also the probabilities of vetoing events are changed. Consequently, n_{trial} acquires a dependence on the parameters of the

variation. Thus, now explicitly stating the dependence on the shower’s starting scale μ_Q^2 and the merging scale Q_{cut} ,

$$\begin{aligned}
 n_{\text{trial}}(\mu_Q^2; Q_{\text{cut}}^2; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) & \\
 \equiv n_{\text{trial}}(\mu_Q^2; Q_{\text{cut}}^2; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}; \alpha_s, f; k_{\alpha_s}, k_f) & \\
 = \frac{1 - \Delta_{n+j}(\mu_{\text{IR}}, \mu_Q^2; Q_{\text{cut}}^2; \tilde{\alpha}_s, \tilde{f}; \tilde{k}_{\alpha_s}, \tilde{k}_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\})}{1 - \Delta_{n+j}(\mu_{\text{IR}}, \mu_Q^2; Q_{\text{cut}}^2; \alpha_s, f; k_{\alpha_s}, k_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\})}, & \quad (29)
 \end{aligned}$$

where

$$\begin{aligned}
 \Delta_{n+j}(\mu_{\text{IR}}, \mu_Q^2; Q_{\text{cut}}^2; \alpha_s, f; k_{\alpha_s}, k_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) & \\
 = \prod_{i=1}^j \exp \left\{ - \int_{t_i}^{t_{i-1}} d\Phi_1 \mathbf{K}_{n+i}(\Phi_1; k_{\alpha_s}, k_f; \alpha_s, f) \right. & \\
 \left. \times \Theta(Q_{n+i} > Q_{\text{cut}}) \right\}. & \quad (30)
 \end{aligned}$$

Thus, n_{trial} corresponds to the survival probability between the unfolding of preexisting splittings when evolving from the n -parton core configuration to the $(n + j)$ -parton configuration. Hence, when changing the parameters of the simulation the truncated showers probability to emit a parton with $Q > Q_{\text{cut}}$ must be re-evaluated following the substitutions $k_{\alpha_s} \rightarrow \tilde{k}_{\alpha_s}, k_f \rightarrow \tilde{k}_f, \alpha_s \rightarrow \tilde{\alpha}_s$ and $f \rightarrow \tilde{f}$ using the methods of Sect. 3. Note, all emissions produced by the truncated shower *prior* to the one that triggers the veto need to be reweighted as they impact the initial conditions for that emission.

4.3 The next-to-leading-order case

The merging of multijet matrix elements at next-to-leading order accuracy proceeds schematically similar as in the leading order case. The input quantity is now the NLOPS matched $(n + j)$ -parton configuration, thus

$$\begin{aligned}
 \langle O \rangle^{\text{MEPS@NLO}} &= \sum_{j=0}^{j_{\text{max}}} \left[\int d\Phi_j \bar{\mathbf{B}}_j^{\text{merge}}(\Phi_j) \Theta(Q_j - Q_{\text{cut}}) \right. \\
 &\quad \times \text{PS}_{\text{NLOPS}}^{\text{V}}(O, \Phi_j) \\
 &\quad \left. + \int d\Phi_{j+1} \mathbf{H}_{A,j}^{\text{merge}}(\Phi_{j+1}, Q_{\text{cut}}) \text{PS}^{\text{vt}}(O, \Phi_{j+1}) \right] \\
 &= \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \left\{ \sum_{i=1}^{N_S} \sum_{j=0}^{j_{\text{max}}} \bar{\mathbf{B}}_j^{\text{merge}}(\Phi_{j,i}) \Theta(Q_j - Q_{\text{cut}}) \right. \\
 &\quad \times \text{PS}_{\text{NLOPS}}^{\text{V}}(O, \Phi_{j,i}) \\
 &\quad \left. + \sum_{i=1}^{N_H} \sum_{j=0}^{j_{\text{max}}} \mathbf{H}_{A,j}^{\text{merge}}(\Phi_{j+1,i}, Q_{\text{cut}}) \text{PS}^{\text{vt}}(O, \Phi_{j+1,i}) \right\}, & \quad (31)
 \end{aligned}$$

keeping the notation of Eqs. (18) and (24). For the S-events the same Θ -function of Eq. (24) is used as an infrared regulator and the S- MC@NLO parton shower of Eq. (23) is replaced by its vetoed version, it only matches the softest emission in t and, thus, does not generate truncated emissions. These are added dressing it with additional emissions through the standard parton shower. As all ingredients of $\bar{\mathbf{B}}_j^{\text{merge}}$ are evaluated at the same phase-space point Φ_j they share a common cluster history $\{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}$. Hence, again suppressing any further Q_{cut} -dependence which is not varied,

$$\bar{\mathbf{B}}_j^{\text{merge}}(\Phi_j) \equiv \bar{\mathbf{B}}_j^{\text{merge}}(\Phi_j; \alpha_s, f; \mu_{R,\text{core}}, \mu_{F,\text{core}}, k_{\alpha_s}, k_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) \quad (32)$$

transforms under the replacements $k_{\alpha_s} \rightarrow \tilde{k}_{\alpha_s}, k_f \rightarrow \tilde{k}_f, \alpha_s \rightarrow \tilde{\alpha}_s$ and $f \rightarrow \tilde{f}$ in the following way:

$$\begin{aligned}
 \bar{\mathbf{B}}_j^{\text{merge}}(\Phi_j; \tilde{\alpha}_s, \tilde{f}; \tilde{\mu}_{R,\text{core}}, \tilde{\mu}_{F,\text{core}}, \tilde{k}_{\alpha_s}, \tilde{k}_f; \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) & \\
 = \prod_{i=1}^j \frac{\tilde{f}_{a_i}(x_{a,i}, \tilde{k}_f t_i)}{\tilde{f}_{a_{i-1}}(x_{a,i-1}, \tilde{k}_f t_i)} \tilde{f}_{a_0}(x_{a,0}, \tilde{\mu}_{F,\text{core}}^2) & \\
 \times \prod_{i=1}^j \frac{\tilde{f}_{b_i}(x_{b,i}, \tilde{k}_f t_i)}{\tilde{f}_{b_{i-1}}(x_{b,i-1}, \tilde{k}_f t_i)} \tilde{f}_{b_0}(x_{b,0}, \tilde{\mu}_{F,\text{core}}^2) & \\
 \times \tilde{\alpha}_s^{n+j}(\tilde{\mu}_R^2) \left[\mathbf{B}'_j(\Phi_j) + \tilde{\alpha}_s(\tilde{\mu}_R^2) \left(\text{VI}'_j(\Phi_j) + c'_{R,j(0)} l_R \right. \right. & \\
 \left. \left. + \frac{1}{2} c'_{R,j(1)} l_R^2 \right) \right. & \\
 \left. + \tilde{\alpha}_s(\tilde{\mu}_R^2) \int dx'_{a/b} \text{KP}'_j(\Phi_j, x'_{a/b}; \tilde{f}; \tilde{\mu}_{F,\text{core}}) \right. & \\
 \left. + \tilde{\alpha}_s(\tilde{\mu}_R^2) \sum_k \int d\Phi_1^k (D'_{A,k} - D'_{S,k}) (\Phi_j \cdot \Phi_1^k) \right] & \\
 - \sum_{i=1}^j \frac{\tilde{\alpha}_s(\tilde{\mu}_R^2)}{2\pi} \log \frac{t_{i-1}}{t_i} & \\
 \times \left(\sum_{c=q,g} \int \frac{dx'_{a,i}}{x'_{a,i}} P_{ac}(x'_{a,i}) \tilde{f}_c\left(\frac{x_{a,i}}{x'_{a,i}}, \tilde{k}_f t_i\right) \right. & \\
 \left. + \sum_{d=q,g} \int \frac{dx'_{b,i}}{x'_{b,i}} P_{bd}(x'_{b,i}) \tilde{f}_d\left(\frac{x_{b,i}}{x'_{b,i}}, \tilde{k}_f t_i\right) \right) & \\
 \times \tilde{\alpha}_s^{n+j}(\tilde{\mu}_R^2) \mathbf{B}'_j(\Phi_j). & \quad (33)
 \end{aligned}$$

In addition to the transformation properties of the $\bar{\mathbf{B}}$ -function of Eq. (19), supplemented with the PDF ratios already encountered in the leading-order case, additional terms appear. They subtract the $\mathcal{O}(\alpha_s)$ expansion of these ratios, in order to retain the NLO accuracy of the merged calculation. Again, please note $t_0 = \mu_{F,\text{core}}^2$.

The same does not hold, however, for the H-events. Their constituent real-emission matrix elements are defined on Φ_{j+1} , while each subtraction term $D_{A,k}$ has its own projection on a phase-space point Φ_j^k . Thus,

$$\begin{aligned}
 & H_{A,j}^{\text{merge}}(\Phi_{j+1}, Q_{\text{cut}}) \\
 & \equiv R_j^{\text{merge}}(\Phi_{j+1}; \alpha_s, f; \mu_{R,\text{core}}, \mu_{F,\text{core}}, k_{\alpha_s}, k_f; \\
 & \quad \{a_i, b_i, x_{a,i}, x_{b,i}, t_i\}) \Theta(Q_j - Q_{\text{cut}}) \\
 & \quad - \sum_k D_{A,k,j}^{\text{merge}}(\Phi_j^k \cdot \Phi_1^k; \alpha_s, f; \mu_{R,\text{core},k}, \mu_{F,\text{core},k}, \\
 & \quad k_{\alpha_s,k}, k_{f,k}; \{a_{i,k}, b_{i,k}, x_{a,i,k}, x_{b,i,k}, t_{i,k}\}) \\
 & \quad \times \Theta(Q_j^k - Q_{\text{cut}}), \tag{34}
 \end{aligned}$$

wherein both R_j^{merge} and the $D_{A,k,j}^{\text{merge}}$ separately transform as the leading-order counterpart B_j^{merge} . While the measure Q on Φ_{j+1} of R_j^{merge} is defined to act on the underlying Φ_j after the first cluster step where the real emission configuration has been reduced to a Born configuration, it is defined directly on each Φ_j^k in each $D_{A,k}$. Infrared safety is guaranteed through the infrared safety of their phase-space maps, the clustering algorithm and the measure Q .

Finally, we consider merging additional multiplicities up to j_{max} described through leading-order matrix elements on top of a next-to-leading order merged calculation with up to $j_{\text{max}}^{\text{NLO}}$ jets, where $j_{\text{max}} > j_{\text{max}}^{\text{NLO}}$. The method of choice was outlined in [22, 65, 66, 69, 70] and is historically referred to as MENLOPS. Its methodology is defined as

$$\begin{aligned}
 & \langle O \rangle^{\text{MEPS@NLO+MENLOPS}} \\
 & = \sum_{j=0}^{j_{\text{max}}^{\text{NLO}}} \left[\int d\Phi_j \bar{B}_j^{\text{merge}}(\Phi_j) \text{PS}_{\text{NLOPS}}^{\text{v}}(O, \Phi_j) \right. \\
 & \quad \left. + \int d\Phi_{j+1} H_{A,j}^{\text{merge}}(\Phi_{j+1}) \text{PS}^{\text{vt}}(O, \Phi_{j+1}) \right] \\
 & \quad + \sum_{j=j_{\text{max}}^{\text{NLO}}+1}^{j_{\text{max}}} \int d\Phi_j k_{j_{\text{max}}^{\text{NLO}}+1}(\Phi_j) B_j^{\text{merge}}(\Phi_j) \\
 & \quad \times \text{PS}^{\text{vt}}(O, \Phi_j) \\
 & = \lim_{N \rightarrow \infty} \frac{1}{N_{\text{trial}}} \left\{ \sum_{i=1}^{N_S} \sum_{j=0}^{j_{\text{max}}^{\text{NLO}}} \bar{B}_j^{\text{merge}}(\Phi_{j,i}) \text{PS}_{\text{NLOPS}}^{\text{v}}(O, \Phi_{j,i}) \right. \\
 & \quad + \sum_{i=1}^{N_H} \sum_{j=0}^{j_{\text{max}}^{\text{NLO}}} H_{A,j}^{\text{merge}}(\Phi_{j+1,i}) \text{PS}^{\text{vt}}(O, \Phi_{j+1,i}) \\
 & \quad \left. + \sum_{i=1}^{N_{\text{LO}}} \sum_{j=j_{\text{max}}^{\text{NLO}}+1}^{j_{\text{max}}} k_{j_{\text{max}}^{\text{NLO}}+1}(\Phi_j) B_j^{\text{merge}}(\Phi_{j,i}) \right\}
 \end{aligned}$$

$$\times \text{PS}^{\text{vt}}(O, \Phi_{j,i}) \Bigg\}, \tag{35}$$

with $N = N_S + N_H + N_{\text{LO}}$. A differential K -factor $k_{j_{\text{max}}^{\text{NLO}}}$ is applied to the higher-multiplicity leading-order matrix elements in order to facilitate a smooth transition across Q_{cut} . It has the form

$$\begin{aligned}
 k_m(\Phi_{m+1}) & = \frac{\bar{B}_m(\Phi_m)}{B_m(\Phi_m)} \left(1 - \frac{H_{A,m}(\Phi_{m+1})}{R_m(\Phi_{m+1})} \right) \\
 & \quad + \frac{H_{A,m}(\Phi_{m+1})}{R_m(\Phi_{m+1})}, \tag{36}
 \end{aligned}$$

and therefore moulds the $B_{j_{\text{max}}^{\text{NLO}}+1}$ into the same form as the $H_{A,j_{\text{max}}^{\text{NLO}}}$ it is replacing. The projection $\Phi_{j_{\text{max}}^{\text{NLO}}+1}(\Phi_j)$ for $j > j_{\text{max}}^{\text{NLO}} + 1$ is defined through its cluster history, as is $\Phi_m(\Phi_{m+1})$ inside k_m itself. When now changing the parameters of the calculation, $\alpha_s \rightarrow \tilde{\alpha}_s, f \rightarrow \tilde{f}, \mu_R \rightarrow \tilde{\mu}_R$ and $\mu_F \rightarrow \tilde{\mu}_F, k_{j_{\text{max}}^{\text{NLO}}}$ transforms as a composite object in terms of its constituents, cf. Sects. 2 and 3. The scales are set directly by the B_j^{merge} process. Of course, in the interest of decreased computational costs one may decide to choose $k_m \equiv 1$ throughout at the cost of larger merging systematics. Similarly, if an electroweak cluster history leads to a changed signature in $\Phi_{m+1}, k_m \equiv 1$ is chosen.

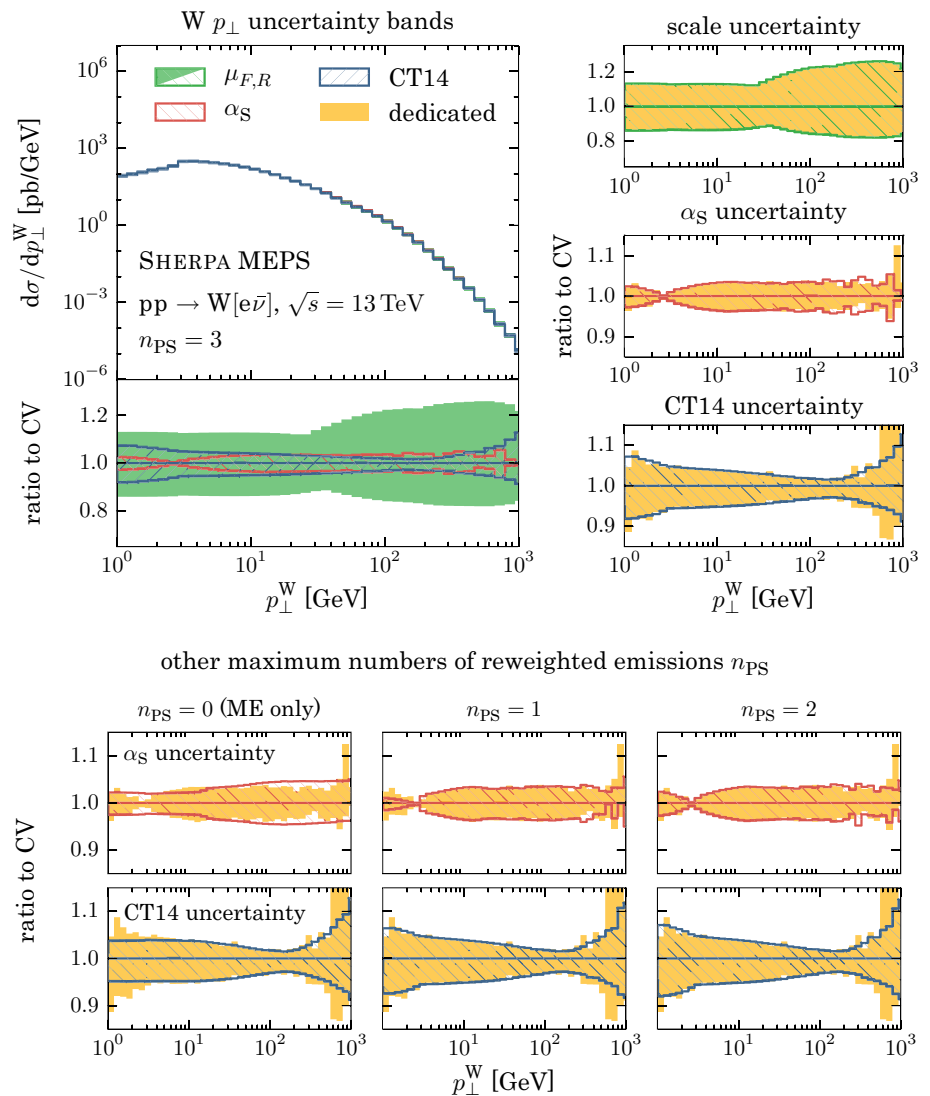
4.4 Validation

The reweighting for multijet-merged calculations as discussed in the previous sections has been implemented within SHERPA with the CSSHOWER for leading-order matrix elements (MEPS@LO), next-to-leading order matrix elements (MEPS@NLO) and next-to-leading-order matrix elements with additional leading-order ones on top (MENLOPS). For the validation, we again perform closure tests between reweighted and dedicated predictions for the transverse momentum of the Wboson in Figs. 6, 7 and 8. As before, the uncertainty bands are the ones defined in Table 1, with the CT14 PDF set. For all merged calculations, we employ a merging cut of $Q_{\text{cut}} = 20$ TeV.

For the MEPS@LO validation in Fig. 6, we combine LO matrix elements for 0-, 1- and 2-jet multiplicities, obtained from COMIX [39]. We can observe that we populate a much larger phase space than for a mere LOPS calculation in terms of p_{\perp}^{W} . Below the merging cut (i.e. $p_{\perp}^{\text{W}} \lesssim 20$ TeV), the scale uncertainty band is equal to the one of the LOPS calculation. For higher p_{\perp}^{W} , the scale uncertainty increases corresponding to the larger uncertainty of the higher-multiplicity matrix elements, that contribute renormalisation scale uncertainties.

In Fig. 7, we consider the MENLOPS case. We combine an NLO matrix element for the 0-jet multiplicities with LO matrix elements for the 1- and 2-jet multiplicities. The scale

Fig. 6 The same as in Figs. 3 and 4, but for a multijet-merged generation with LO matrix elements for 0-, 1- and 2-jet multiplicities. The uncertainty bands are calculated by reweighting the ME and a maximum number of emissions n_{PS} of PS emissions. In the upper four plots, $n_{PS} = 3$, thus up to three emissions are reweighted. In the lower plots, n_{PS} is varied for comparison. Again, we find a saturation when reproducing dedicated calculations for $n_{PS} \geq 2$, with no further improvement when n_{PS} is increased from 2 to 3



uncertainty for low p_{\perp}^W values now features the reduced scale uncertainty, that we already have seen in the NLOPS validation.

The same is true in the MEPS@NLO case depicted in Fig. 8. A direct comparison of the scale uncertainties to the MENLOPS case is not straightforward though, as we combine NLO matrix elements for the 0- and the 1-jet multiplicity, where the virtual amplitudes are obtained from BLACKHAT [42]. Hence, the 2-jet multiplicity is described at leading order through the 1-jet H-events. As such, the set-up is not a simple upgrade from our MENLOPS calculation.

In all multijet-merging validations, we find a similar behaviour with respect to the imprint of including emissions in the reweighting. For $n_{NLOPS} + n_{PS} = 2$, the dedicated calculations are well reproduced, and no further improvement is found for $n_{NLOPS} + n_{PS} = 3$. It is noteworthy, that for the MENLOPS case we find a worse reproduction for $n_{NLOPS} = 1$ and $n_{PS} = 0$ compared to the NLOPS and the MEPS@NLO cases. This originates in the fact that in the latter two cases,


we enable the reweighting of emissions off S-events at all involved multiplicities, whereas in the MENLOPS case only the first of the three multiplicities is affected, because the other two are at LO and therefore do not have S-events. Thus, the overall importance of the S emission reweighting gets restricted to the region below Q_{cut} of the 1-jet configuration in the MENLOPS case.

5 Consistent variations

In general, the renormalisation and factorisation scales, α_s and the PDFs should be varied consistently throughout any of the presented calculations. While at fixed order the situation is clear, the matched and merged approaches allow for some degree of freedom regarding partial variations while still retaining their respective accuracies.

In the simplest case, LOPS, μ_R and μ_F of the short distance cross section and the parton shower may be var-

Fig. 7 The same as in Figs. 3, 4 and 6, but for a multijet-merged generation with one NLO matrix element for the 0-jet multiplicity, and LO matrix elements for the 1- and 2-jet multiplicities. The uncertainty bands are calculated by reweighting the ME and a maximum number of emissions from the MC@NLO (n_{NLOPS}) and the ordinary PS (n_{PS}). In the upper four plots, $n_{\text{NLOPS}} = 1$ and $n_{\text{PS}} = 2$, thus up to three emissions are reweighted. In the lower plots, both n are varied for comparison. Again, we find a saturation when reproducing dedicated calculations for $n_{\text{NLOPS}} + n_{\text{PS}} \geq 2$, with no further improvement when n_{PS} is increased from 1 to 2



ied independently as these variations can be expressed as higher-order terms in a perturbative expansion in the coupling parameter α_s . This is not the case for α_s itself and the PDFs as they are fixed through measured input values and parametrisations. Changes in these input values cannot be expressed as simple higher-order terms. Thus they need to be chosen consistently throughout.

Similarly, in NLOPS calculations, the renormalisation and factorisation scales may be varied in the matrix element (\bar{B} and H_A) or the parton shower (PS_{NLOPS} and PS) separately without losing neither the fixed-order nor the resummation accuracy. As the pseudo-subtraction through the D_A in any case employs different scales in PS_{NLOPS} and the \bar{B} and H_A functions, it always leaves remainders of $\mathcal{O}(\alpha_s^2)$. Hence, further scale variations in either one, the short-distance cross sections or the PS_{NLOPS} , do not worsen the nominal accuracy of the method. Retaining the logarithmic accuracy of the parton shower on the other hand requires identical renormalisation and factorisation scales throughout all resummation-

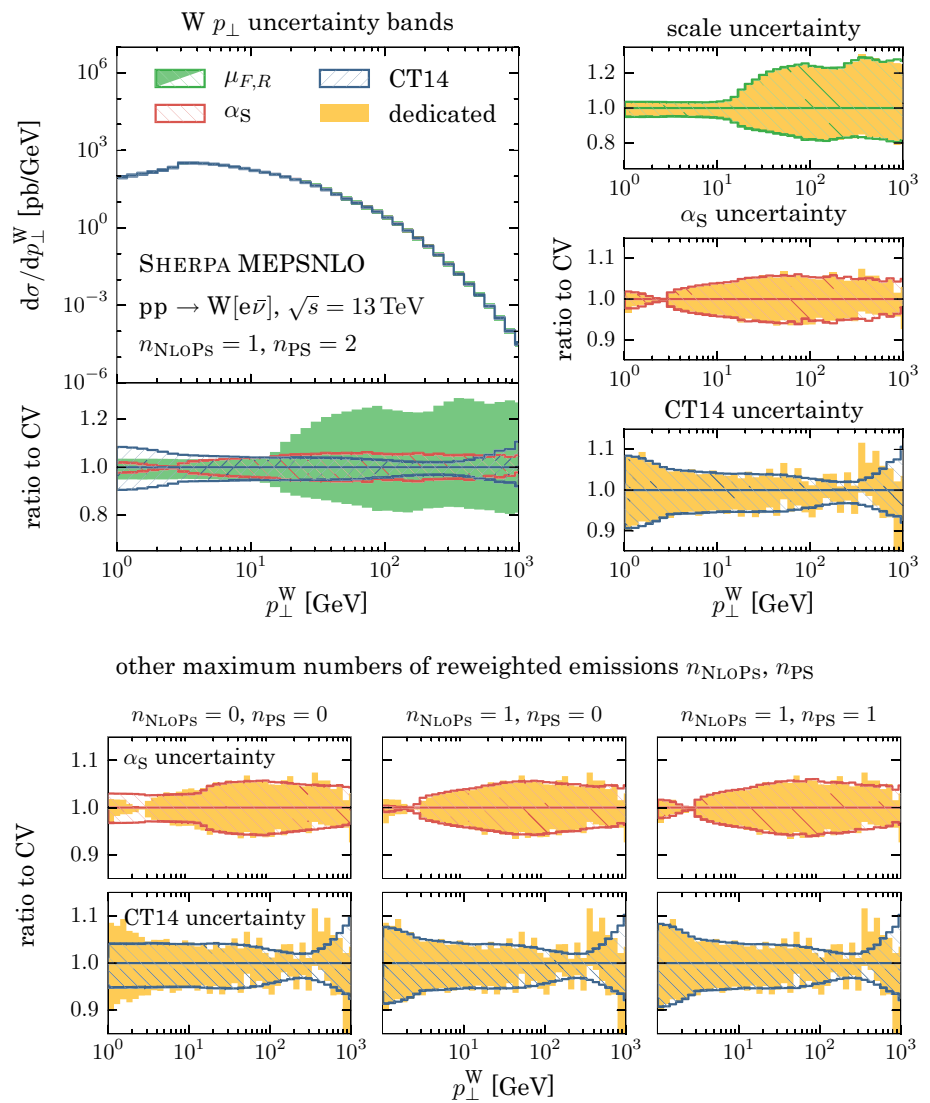
relevant components, i.e. PS_{NLOPS} and PS. Again, variations in α_s or the PDFs need to be consistent throughout the calculation.

The multijet-merged calculations impose further constraints since they treat multijet matrix elements and parton-shower emissions on the same footing. The notation of the scales already reflects this for μ_R and μ_F . In their definitions only the core scales remain as free parameters and may be varied independently. Again, the α_s and PDF parametrisations need to be the same throughout.

6 Conclusions

In this publication we have presented the implementation and validation of reweighting techniques allowing for the fast and efficient evaluation of perturbative systematic uncertainties in the SHERPA event-generator framework. We have lifted the available techniques for the determination of PDF,

Fig. 8 The same as in Figs. 3, 4, 6 and 7, but for a multijet merged generation with NLO matrix elements for the 0- and 1-jet multiplicities. The uncertainty bands are calculated by reweighting the ME and a maximum number of emissions from the MC@NLO (n_{NLOPS}) and the ordinary PS (n_{PS}). In the upper four plots, $n_{\text{NLOPS}} = 1$ and $= n_{\text{PS}} = 2$, thus up to three emissions are reweighted. In the lower plots, both n are varied for comparison. Again, we find a saturation when reproducing dedicated calculations for $n_{\text{NLOPS}} + n_{\text{PS}} \geq 2$, with no further improvement when n_{PS} is increased from 1 to 2



α_s and scale uncertainties in leading- and next-to-leading order QCD calculations to include the respective variations in parton-shower simulations. In turn we provide the means to perform consistent uncertainty evaluations for multijet-merged simulations based on leading- or next-to-leading-order accurate matrix elements of varying multiplicity matched with parton showers. The foundation for our reweighting method is the knowledge of the very dependence structure of the perturbative calculations on the parameters to be varied. For the fixed-order components this amounts to the corresponding decomposition of the Catani–Seymour dipole subtraction terms. This needed to be supplemented by the reweighting of the parametric dependences of the parton shower treated through the Sudakov Veto Algorithm.

With our extensive validation we have been able to prove on the one-hand-side the correctness of the implementation and have, furthermore, been able to illustrate the importance of parton-shower reweighting for reliable uncertainty

estimates. With comparably little additional computational costs this allows for the on-the-fly determination of PDF, α_s and scale uncertainties based on one single generator run, that, otherwise, would require explicit re-computations. The overall reduction in CPU time is by a factor of about 3–20, depending on the event-generation mode used, see Appendix A. The variational event weights provided are easily accessible through the HEPMC event record and are furthermore consistently handed over to the RIVET analysis software by the corresponding SHERPA interface.

The methods presented in this publication are ideally suited for event-wise uncertainty estimates and can readily be used in arbitrary theoretical and experimental analyses. An extension to next-to-next-to-leading-order QCD calculations possibly dressed with parton showers, as presented in [16, 17, 71], is straightforward and planned for the near future.

The decomposition of the fixed-order part of QCD calculations employed here is also a necessary ingredient to produce cross-section grids as provided by the APPLGRID [72] and FASTNLO [73,74] tools. These store the perturbative coefficients for a certain observable calculation discretised in Q^2 and x . Using interpolation methods, this allows for the a posteriori inclusion of PDFs, α_s and variations of the renormalisation and factorisation scales. In turn, such techniques are well suited for (combined) fits of PDFs and α_s that require a multitude of re-computations of the theoretical predictions. Over the last years, tools have been developed that automate the projection of arbitrary next-to-leading-order QCD calculations onto such grids, namely the AMCFast [75] and the MCGRID [76–78] packages. The first one produces APPLGRIDS with MADGRAPH5_AMC@NLO [79], the latter APPLGRIDS or FASTNLO grids from SHERPA events projected on the observables through RIVET. The APFELGRID tool [80] provides an improved convolution method for use with APPLGRID files that furthermore speeds-up the re-evaluations.

However, none of these approaches includes generic parton-shower effects, i.e. the parametric dependence of the shower component on the PDFs and α_s is ignored. With the methods presented in this publication we are confident that we can surmount this limitation and in the future provide interpolation grids that properly reflect shower-resummation effects and allow for the inclusion of the affected phase-space regions in PDF determinations.

Acknowledgements We acknowledge financial support from BMBF under Contract 05H15MGCAA, the Swiss National Foundation (SNF) under Contract PP00P2128552 and from the EU MCnetITN research network funded under Framework Programme 7 Contract PITNGA 2012315877.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. Funded by SCOAP³.

Appendix A: CPU time measurements

The benefit of reweighted calculations is given by the saving of CPU time. In order to evaluate the gain, we shall compare the event generation time of reweighted calculations with the sum of generation times for all corresponding dedicated computations. Here we consider both parton-level calculations, as well as runs including multiple interactions and hadronisation, the typical default in physics analyses applications. For the latter it can be expected that the gain in CPU time by using the reweighting approach is most considerable, as

the CPU intense non-perturbative event generation phases do not need to be re-evaluated. In what follows we compare actual event-generation times, neglecting the set-up times of the individual runs.⁹

In Fig. 9, we consider event generations using LOPS, NLOPS, MEPS@LO and MEPS@NLO calculations for $pp \rightarrow W[e^- \bar{\nu}]$ at 13 TeV. The ratio of CPU time between the reweighting and the dedicated generations is shown for different maximum numbers of reweighted shower emissions $n_{PS} + n_{NLOPS}$. Whether non-perturbative effects are included or not, the time needed for the reweighting calculation is below 10% of the time needed for dedicated calculations if only the matrix element is reweighted ($n_{PS} = n_{NLOPS} = 0$). The ratio then increases for larger numbers of reweighted emissions, as their reweighting needs additional time, asymptotically approaching the value when all parton-shower emissions are reweighted. For parton-level-only calculations, this ratio is around 0.35 for LOPS events, and around 0.3 for NLOPS events. This reduction can be explained due to relatively smaller computational cost of the parton shower as a whole when the rest of the calculation is more complex. Also note that n_{PS} for LOPS is only equivalent to $n_{PS} + n_{NLOPS}$ for \mathcal{S} events. \mathcal{H} events do not feature the S-MC@NLO-emission, and hence for them n_{NLOPS} does not contribute to their reweighting.

For the same reason, when non-perturbative effects are included, that ratio improves to about 0.1: The parton shower (and its reweighting) component plays a relatively smaller rôle in terms of CPU cycles, when multiple interactions and hadronisation are enabled.

If on top of the non-perturbative effects the events are also unweighted, the ratio does not change in the LOPS case, but in the NLOPS case (by about 20%). A reason might be, that only for NLOPS a sizeable number of events gets rejected. For these, the jet evolution and non-perturbative phases are not performed at all, whereas the matrix-element calculation (and its reweighting) is always done, for accepted and rejected events alike. The same is true for the S-MC@NLO emission from \mathcal{S} events. As a consequence, the *relative* cost of the reweighting grows slightly. A future improvement of the implementation would postpone these futile reweightings in unweighted calculations to a time point after the possible rejection. This of course requires that the dependence of the rejection probability is negligible. For the observables studied so far this was found to be true, at least to $\mathcal{O}(10^{-4})$.

⁹ If NLO matrix elements at higher multiplicities are needed for an event generation, the time needed for the integrator optimisation and the process selection weight optimisation can be quite substantial, e.g. a couple of days. In the case of unweighted event generation, this even has to be re-done for every single parameter variation, as the channel weights are used for the unweighting. When reweighting is used, this is not necessary and so even more CPU time is saved.




Fig. 9 The ratio of CPU time needed for a reweighting event generation t_{rew} over the time needed for generating predictions for all variations with dedicated runs t_{ded} . The reweighting includes up to $n_{PS} + n_{NLOPS}$ parton-shower emissions. The sampled variations are listed in Table 1. Parton-level-only results are compared to results for cal-

culations including multiple interactions and hadronisation effects (“+ non-perturbative”), and to calculations where in addition to adding non-perturbative effects the events have also been unweighted (“+ unweighting”). The ratios for reweighting all emissions are indicated with a *horizontal line*

Note that the effective gains will be lower than the results presented in this section, when we take into account the reduced statistical accuracy which comes with the parton-shower reweighting. This requires more events to be generated in a reweighting calculation to reach the same statistical accuracy as in a dedicated calculation.

Appendix B: Configuring and accessing event-weight variations

SHERPA provides a list of pre-calculated alternative event weights, which are automatically output to the HEPMC event record [24] or directly to an interfaced RIVET analysis [25]. For versions of SHERPA later than v.2.2.0, the variations to calculate can be specified with the following line in the (run) section of the SHERPA run card:

```
VARIATIONS muR2 fac1 , muF2 fac1 , PDF1
muR2 fac2 , muF2 fac2 , PDF2 . . . ;
```

Each variation is characterised by up to three arguments

- muR2 fac a prefactor multiplying the nominal (squared) renormalisation scale
- muF2 fac a prefactor multiplying the nominal (squared) factorisation scale
- PDF a parton density and its accompanying α_s parametrisation.

This syntax works for all employed scale setters of SHERPA and both SHERPA’s internal PDFs and PDFs interfaced through LHAPDF5/6 [81,82]. If trailing arguments are omitted from a variation, their default values are used, which is 1.0 for scale factors and the PDF set used by SHERPA for the nominal calculation.

In HEPMC event records (v. 2.06 or later), the alternate weights can be accessed as named weights within the `HepMC::WeightContainer` of each event. The keys are given in one of the following formats:

```
MUR<muR2 fac>_MUF<muF2 fac>_PDF<ID>
MUR<muR2 fac>_MUF<muF2 fac>_PDF<ID>
_PSMUR<muR2 fac>_PSMUF<muF2 fac>
```

The parts in angle brackets are replaced with the respective scale factors and LHAPDF IDs. The second form is used, if a factor is applied to the renormalisation/factorisation scale of parton-shower splittings. This includes splittings within cluster histories determined by the multijet merging procedure, as discussed in Sect. 4. If the scale reweighting with parton-shower splittings has been enabled (we discuss below how to do so), the scale factors for MUR, MUF and PSMUR, PSMUF are always equal, respectively, in the current implementation.

If the internal RIVET interface of SHERPA is used to analyse events during the generation, one histogram file per variation is written to disk, along with the nominal one. The file names follow a pattern resembling the HEPMC weight-container keys as specified above.

Scale variations

The scale argument can also be specified by enclosing it in square brackets: `[mu2fac]`. This syntactic sugar implies both the given factor, its inverse and the default value. For example, `1.0, [4.0]` is equivalent to `1.0, 4.0 1.0, 0.25 1.0, 1.0` and therefore triggers up and down variations of the factorisation scale, along with the central value. If both scale factors are enclosed in brackets, they are expanded individually, keeping the other at its default value of 1.0: Hence, `[4.0], [4.0]` is equivalent to the 5-point scale variation `4.0, 1.0 0.25, 1.0 1.0, 4.0 1.0, 0.25 1.0, 1.0`. To include simultaneous variations in the same direction, both factors can be surrounded by a single pair of brackets. Thus, `[4.0, 4.0]` is equivalent to the 7-point scale variation `4.0, 1.0 0.25, 1.0 1.0, 4.0 1.0, 0.25 4.0, 4.0 0.25, 0.25 1.0, 1.0`.

PDF and α_s variations

PDF and α_s variations both work by specifying a PDF set through the PDF argument of a variation. This is because SHERPA per default uses the value for $\alpha_s(m_Z^2)$ given by the PDF set in use. Therefore an α_s variation can be achieved by using PDF fits for different values of $\alpha_s(m_Z^2)$.

To specify a specific member of a PDF set, its number is given as an additional argument separated by a slash. Thus, `1.0, 1.0, CT14nlo/38` asks for the 38th member of the CT14nlo PDF set, without modifying the renormalisation

and factorisation scales. If the slash and the number are not given, the central PDF member is used, i.e. `CT14nlo` is equivalent to `CT14nlo/0`.

SHERPA can also be asked to do variations for all members of a PDF set by enclosing it in square brackets. Hence, `1.0, 1.0, [CT14nlo]` is equivalent to

```
1.0, 1.0, CT14nlo/0 1.0, 1.0, CT14nlo/1 ...
1.0, 1.0, CT14nlo/56
```

This [PDF]-notation only works with PDFs interfaced through LHAPDF6 [82]. It can be combined with scale factors that are enclosed in square brackets. Again, the expansions are done individually, keeping other arguments at their default values. This means that for example `1.0, [4.0], [CT14nlo]` is equivalent to `1.0, [4.0] 1.0, 1.0, [CT14nlo]`. Hence, a 7-point scale variation and a full CT14nlo PDF variation can be requested by

```
VARIATIONS [4.0, 4.0], [CT14nlo];
```

Configuring how variations are calculated

The following options always affect all variations that are specified by arguments to the VARIATIONS keyword.

REWEIGHT_SPLITTING_ALPHAS_SCALES (default: 0)

If this is set to 1, the renormalisation scale factor is applied to the α_s argument of individual splittings, instead of applying it only to the overall renormalisation scale, see Sect. 4.2. This means that parton-shower splittings are only included in the rescaling, if this option is enabled.

In the notation of Sects. 3 and 4, this sets $\tilde{k}_{\alpha_s} = \tilde{\mu}_R/\mu_R$.

REWEIGHT_SPLITTING_PDF_SCALES (default: 0)

If this is set to 1, the factorisation scale factor is also applied to PDF scale arguments within shower splittings (and intermediate cluster history PDF ratios), and not only to the core-process PDFs. In the notation of Sects. 3 and 4, this sets $\tilde{k}_f = \tilde{\mu}_F/\mu_F$.

REWEIGHT_MAXEM (default: -1) This option specifies the number of ordinary parton-shower emissions included in the reweighting per event. If this is set to 0, no emission is reweighted. The default value -1 means that all emissions should be reweighted.

REWEIGHT_MCATNLO_EM (default: 1) If this is set to 0, the single parton-shower emission within the MC@NLO contribution is not reweighted.

VARIATIONS_INCLUDE_CV (default: 1) If this is set to 0, the behaviour of the square bracket syntax is changed, such that the central-value variation is not included when expanding a parameter in square brackets. It is recommended not to disable it, such that one can do a closure test between the dedicated calculation and the reweighting. However, in CPU intensive applications,

this setting can be used to omit this one obsolete variation while still making use of the convenient square-bracket syntax.

References

- ATLAS Collaboration, G. Aad et al., Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett. B* **716**, 1–29 (2012). [arXiv:1207.7214](#)
- CMS Collaboration, S. Chatrchyan et al., Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett. B* **716**, 30–61 (2012). [arXiv:1207.7235](#)
- I. Schörner-Sadenius (ed.), *The Large Hadron Collider* (Springer, Berlin, 2015). doi:10.1007/978-3-319-15001-7
- J. Bellm et al., Herwig 7.0/Herwig++ 3.0 release note. *Eur. Phys. J. C* **76**, 196 (2016). [arXiv:1512.01178](#)
- M. Bähr et al., Herwig++ physics and manual. *Eur. Phys. J. C* **58**, 639–707 (2008). [arXiv:0803.0883](#)
- T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten et al., An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **191**, 159–177 (2015). [arXiv:1410.3012](#)
- T. Gleisberg, S. Höche, F. Krauss, A. Schälicke, S. Schumann, J.-C. Winter, SHERPA 1. Alpha: a proof of concept version. *JHEP* **02**, 056 (2004). [arXiv:hep-ph/0311263](#)
- T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert et al., Event generation with SHERPA 1.1. *JHEP* **02**, 007 (2009). [arXiv:0811.4622](#)
- A. Buckley et al., General-purpose event generators for LHC physics. *Phys. Rep.* **504**, 145–233 (2011). [arXiv:1101.2599](#)
- S. Catani, F. Krauss, R. Kuhn, B.R. Webber, QCD matrix elements + parton showers. *JHEP* **11**, 063 (2001). [arXiv:hep-ph/0109231](#)
- F. Krauss, Matrix elements and parton showers in hadronic interactions. *JHEP* **08**, 015 (2002). [arXiv:hep-ph/0205283](#)
- S. Frixione, B.R. Webber, Matching NLO QCD computations and parton shower simulations. *JHEP* **06**, 029 (2002). [arXiv:hep-ph/0204244](#)
- P. Nason, A New method for combining NLO QCD with shower Monte Carlo algorithms. *JHEP* **11**, 040 (2004). [arXiv:hep-ph/0409146](#)
- K. Hamilton, P. Nason, E. Re, G. Zanderighi, NNLOPS simulation of Higgs boson production. *JHEP* **10**, 222 (2013). [arXiv:1309.0017](#)
- A. Karlberg, E. Re, G. Zanderighi, NNLOPS accurate Drell-Yan production. *JHEP* **09**, 134 (2014). [arXiv:1407.2940](#)
- S. Höche, Y. Li, S. Prestel, Drell-Yan lepton pair production at NNLO QCD with parton showers. *Phys. Rev. D* **91**, 074015 (2015). [arXiv:1405.3607](#)
- S. Höche, Y. Li, S. Prestel, Higgs-boson production through gluon fusion at NNLO QCD with parton showers. *Phys. Rev. D* **90**, 054011 (2014). [arXiv:1407.3773](#)
- J. Bellm, G. Nail, S. Plätzer, P. Schichtel, A. Siodmok, Parton shower uncertainties with Herwig 7: benchmarks at leading order (2016). [arXiv:1605.01338](#)
- A. Buckley, H. Hoeth, H. Lacker, H. Schulz, J.E. von Seggern, Systematic event generator tuning for the LHC. *Eur. Phys. J. C* **65**, 331–357 (2010). [arXiv:0907.2973](#)
- P.Z. Skands, Tuning Monte Carlo generators: the Perugia tunes. *Phys. Rev. D* **82**, 074018 (2010). [arXiv:1005.3457](#)
- S. Höche, F. Krauss, S. Schumann, F. Siegert, QCD matrix elements and truncated showers. *JHEP* **05**, 053 (2009). [arXiv:0903.1219](#)
- S. Höche, F. Krauss, M. Schönherr, F. Siegert, NLO matrix elements and truncated showers. *JHEP* **08**, 123 (2011). [arXiv:1009.1127](#)
- S. Schumann, F. Krauss, A parton shower algorithm based on Catani–Seymour dipole factorisation. *JHEP* **03**, 038 (2008). [arXiv:0709.1027](#)
- M. Dobbs, J.B. Hansen, The HepMC C++ Monte Carlo event record for high energy physics. *Comput. Phys. Commun.* **134**, 41–46 (2001)
- A. Buckley, J. Butterworth, L. Lönnblad, D. Grellscheid, H. Hoeth, J. Monk et al., Rivet user manual. *Comput. Phys. Commun.* **184**, 2803–2819 (2013). [arXiv:1003.0694](#)
- J. Bellm, S. Plätzer, P. Richardson, A. Siodmok, S. Webster, Reweighting parton showers. *Phys. Rev. D* **94**, 034028 (2016). [arXiv:1605.08256](#)
- S. Mrenna, P. Skands, Automated parton-shower variations in Pythia 8. *Phys. Rev. D* **94**, 074005 (2016). [arXiv:1605.08352](#)
- W.T. Giele, D.A. Kosower, P.Z. Skands, Higher-order corrections to timelike jets. *Phys. Rev. D* **84**, 054003 (2011). [arXiv:1102.2126](#)
- A.J. Larkoski, J.J. Lopez-Villarejo, P. Skands, Helicity-dependent showers and matching with VINCIA. *Phys. Rev. D* **87**, 054033 (2013). [arXiv:1301.0933](#)
- N. Fischer, S. Prestel, M. Ritzmann, P. Skands, Vincia for Hadron colliders. [arXiv:1605.06142](#)
- S. Badger et al., Les Houches 2015: physics at TeV colliders standard model working group report (2016) [arXiv:1605.04692](#)
- S. Gieseke, Uncertainties of Sudakov form-factors. *JHEP* **01**, 058 (2005). [arXiv:hep-ph/0412342](#)
- A. Buckley, Sensitivities to PDFs in parton shower MC generator reweighting and tuning (2016). [arXiv:1601.04229](#)
- Z. Bern, L.J. Dixon, F. Febres Cordero, S. Höche, H. Ita, D.A. Kosower et al., Ntuples for NLO events at hadron colliders. *Comput. Phys. Commun.* **185**, 1443–1460 (2014). [arXiv:1310.7439](#)
- R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, R. Pittau, P. Torrielli, Four-lepton production at hadron colliders: aMC@NLO predictions with theoretical uncertainties. *JHEP* **02**, 099 (2012). [arXiv:1110.4738](#)
- S. Catani, M.H. Seymour, A general algorithm for calculating jet cross-sections in NLO QCD. *Nucl. Phys. B* **485**, 291–419 (1997). [arXiv:hep-ph/9605323](#)
- S. Catani, S. Dittmaier, M.H. Seymour, Z. Trocsanyi, The dipole formalism for next-to-leading order QCD calculations with massive partons. *Nucl. Phys. B* **627**, 189–265 (2002). [arXiv:hep-ph/0201036](#)
- F. Krauss, R. Kuhn, G. Soff, AMEGIC++ 1.0: a matrix element generator in C++. *JHEP* **02**, 044 (2002). [arXiv:hep-ph/0109036](#)
- T. Gleisberg, S. Höche, Comix, a new matrix element generator. *JHEP* **12**, 039 (2008). [arXiv:0808.3674](#)
- S. Höche, S. Kuttimalai, S. Schumann, F. Siegert, Beyond standard model calculations with Sherpa. *Eur. Phys. J. C* **75**, 135 (2015). [arXiv:1412.6478](#)
- T. Gleisberg, F. Krauss, Automating dipole subtraction for QCD NLO calculations. *Eur. Phys. J. C* **53**, 501–523 (2008). [arXiv:0709.2881](#)
- C.F. Berger, Z. Bern, L.J. Dixon, F. Febres Cordero, D. Forde, H. Ita et al., An automated implementation of on-shell methods for one-loop amplitudes. *Phys. Rev. D* **78**, 036003 (2008). [arXiv:0803.4180](#)
- F. Cascioli, P. Maierhöfer, S. Pozzorini, Scattering amplitudes with open loops. *Phys. Rev. Lett.* **108**, 111601 (2012). [arXiv:1111.5206](#)
- G. Cullen et al., GoSam-2.0: a tool for automated one-loop calculations within the standard model and beyond. *Eur. Phys. J. C* **74**, 3001 (2014). [arXiv:1404.7096](#)
- S. Badger, B. Biedermann, P. Uwer, V. Yundin, Numerical evaluation of virtual corrections to multi-jet production in massless QCD. *Comput. Phys. Commun.* **184**, 1981–1998 (2013). [arXiv:1209.0100](#)
- T. Binoth et al., A proposal for a standard interface between Monte Carlo tools and one-loop programs. *Comput. Phys. Commun.* **181**, 1612–1622 (2010). [arXiv:1001.1307](#)

47. C.F. Berger, Z. Bern, L.J. Dixon, F. Febres Cordero, D. Forde, T. Gleisberg et al., Precise predictions for $W + 4$ jet production at the large hadron collider. *Phys. Rev. Lett.* **106**, 092001 (2011). [arXiv:1009.2338](#)
48. NNPDF Collaboration, R.D. Ball et al., Parton distributions for the LHC Run II. *JHEP* **04**, 040 (2015). [arXiv:1410.8849](#)
49. J. Butterworth et al., PDF4LHC recommendations for LHC Run II. *J. Phys. G* **43**, 023001 (2016). [arXiv:1510.03865](#)
50. M.H. Seymour, Matrix element corrections to parton shower algorithms. *Comput. Phys. Commun.* **90**, 95–101 (1995). [arXiv:hep-ph/9410414](#)
51. T. Sjöstrand, S. Mrenna, P.Z. Skands, PYTHIA 6.4 physics and manual. *JHEP* **05**, 026 (2006). [arXiv:hep-ph/0603175](#)
52. S. Höche, S. Schumann, F. Siegert, Hard photon production and matrix-element parton-shower merging. *Phys. Rev. D* **81**, 034026 (2010). [arXiv:0912.3501](#)
53. S. Plätzer, M. Sjöstrand, The Sudakov Veto algorithm reloaded. *Eur. Phys. J. Plus* **127**, 26 (2012). [arXiv:1108.6180](#)
54. S. Höche, F. Krauss, M. Schönherr, F. Siegert, A critical appraisal of NLO + PS matching methods. *JHEP* **09**, 049 (2012). [arXiv:1111.1220](#)
55. L. Lönnblad, Fooling around with the Sudakov Veto algorithm. *Eur. Phys. J. C* **73**, 2350 (2013). [arXiv:1211.7204](#)
56. S. Catani, B.R. Webber, G. Marchesini, QCD coherent branching and semiinclusive processes at large x . *Nucl. Phys. B* **349**, 635–654 (1991)
57. S. Alioli, P. Nason, C. Oleari, E. Re, NLO vector-boson production matched with shower in POWHEG. *JHEP* **07**, 060 (2008). [arXiv:0805.4802](#)
58. S. Höche, M. Schönherr, Uncertainties in next-to-leading order plus parton shower matched simulations of inclusive jet and dijet production. *Phys. Rev. D* **86**, 094042 (2012). [arXiv:1208.2815](#)
59. E. Farhi, A QCD test for jets. *Phys. Rev. Lett.* **39**, 1587–1588 (1977)
60. 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](#)
61. ATLAS Collaboration, G. Aad et al., Measurement of the production cross section of jets in association with a Z boson in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *JHEP* **07**, 032 (2013). [arXiv:1304.7098](#)
62. CMS Collaboration, S. Chatrchyan et al., Measurement of differential cross sections for the production of a pair of isolated photons in pp collisions at $\sqrt{s} = 7$ TeV. *Eur. Phys. J. C* **74**, 3129 (2014). [arXiv:1405.7225](#)
63. CMS Collaboration, V. Khachatryan et al., Differential cross section measurements for the production of a W boson in association with jets in proton-proton collisions at $\sqrt{s} = 7$ TeV. *Phys. Lett. B* **741**, 12–37 (2015). [arXiv:1406.7533](#)
64. ATLAS Collaboration, M. Aaboud et al., Measurement of the $W^\pm Z$ boson pair-production cross section in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B* **759**, 601–621 (2016). [arXiv:1606.04017](#)
65. S. Höche, F. Krauss, M. Schönherr, F. Siegert, QCD matrix elements + parton showers: the NLO case. *JHEP* **04**, 027 (2013). [arXiv:1207.5030](#)
66. T. Gehrmann, S. Höche, F. Krauss, M. Schönherr, F. Siegert, NLO QCD matrix elements + parton showers in $e^+e^- \rightarrow$ hadrons. *JHEP* **01**, 144 (2013). [arXiv:1207.5031](#)
67. A. Schälicke, F. Krauss, Implementing the ME + PS merging algorithm. *JHEP* **07**, 018 (2005). [arXiv:hep-ph/0503281](#)
68. F. Krauss, P. Petrov, M. Schönherr, M. Spannowsky, Measuring collinear W emissions inside jets. *Phys. Rev. D* **89**, 114006 (2014). [arXiv:1403.4788](#)
69. K. Hamilton, P. Nason, Improving NLO-parton shower matched simulations with higher order matrix elements. *JHEP* **06**, 039 (2010). [arXiv:1004.1764](#)
70. S. Höche, F. Krauss, S. Pozzorini, M. Schönherr, J.M. Thompson, K.C. Zapp, Triple vector boson production through Higgs–Strahlung with NLO multijet merging. *Phys. Rev. D* **89**, 093015 (2014). [arXiv:1403.7516](#)
71. S. Dawson, P. Jaiswal, Y. Li, H. Ramani, M. Zeng, Resummation of jet veto logarithms at N^3LL_a + NNLO for W^+W^- production at the LHC. *Phys. Rev. D* (2016) (submitted). [arXiv:1606.01034](#)
72. T. Carli, D. Clements, A. Cooper-Sarkar, C. Gwenlan, G.P. Salam, F. Siegert et al., A posteriori inclusion of parton density functions in NLO QCD final-state calculations at hadron colliders: The APPLGRID Project. *Eur. Phys. J. C* **66**, 503–524 (2010). [arXiv:0911.2985](#)
73. T. Kluge, K. Rabbertz, M. Wobisch, FastNLO: fast pQCD calculations for PDF fits, in *Deep Inelastic Scattering. Proceedings, 14th International Workshop, DIS 2006, Tsukuba, Japan, April 20–24, 2006* (2006), pp. 483–486. [arXiv:hep-ph/0609285](#)
74. fastNLO Collaboration, D. Britzger, K. Rabbertz, F. Stober, M. Wobisch, New features in version 2 of the fastNLO project, in *Proceedings, 20th International Workshop on Deep-Inelastic Scattering and Related Subjects (DIS 2012)* (2012), pp. 217–221. [arXiv:1208.3641](#)
75. V. Bertone, R. Frederix, S. Frixione, J. Rojo, M. Sutton, aMCfast: automation of fast NLO computations for PDF fits. *JHEP* **08**, 166 (2014). [arXiv:1406.7693](#)
76. L. Del Debbio, N.P. Hartland, S. Schumann, MCgrid: projecting cross section calculations on grids. *Comput. Phys. Commun.* **185**, 2115–2126 (2014). [arXiv:1312.4460](#)
77. E. Bothmann, N. Hartland, S. Schumann, Introducing MCgrid 2.0: projecting cross section calculations on grids. *Comput. Phys. Commun.* **196**, 617–618 (2015)
78. E. Bothmann, M. Schönherr, S. Schumann, Fast evaluation of theoretical uncertainties with Sherpa and MCgrid. *PoS DIS2015*, 136 (2015). [arXiv:1507.03908](#)
79. J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP* **07**, 079 (2014). [arXiv:1405.0301](#)
80. V. Bertone, S. Carrazza, N.P. Hartland, APFELgrid: a high performance tool for parton density determinations. *Comput. Phys. Commun.* (2016) (submitted). [arXiv:1605.02070](#)
81. M. R. Whalley, D. Bourilkov, R.C. Group, *The Les Houches Accord PDFs (LHAPDF) and LHAGLUE, in HERA and the LHC: A Workshop on the Implications of HERA for LHC Physics. Proceedings, Part B, 2005*. [arXiv:hep-ph/0508110](#)
82. A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht et al., LHAPDF6: parton density access in the LHC precision era. *Eur. Phys. J. C* **75**, 132 (2015). [arXiv:1412.7420](#)