

Precision phenomenology at multi-TeV muon colliders

Stefano Frixione^a, Fabio Maltoni^{b,c,d}, Davide Pagani^d and Marco Zaro^{e,f}

^a*INFN, Sezione di Genova,
Via Dodecaneso 33, Genoa, I-16146, Italy*

^b*Centre for Cosmology, Particle Physics and Phenomenology (CP3),
Université Catholique de Louvain,
Louvain-la-Neuve, B-1348, Belgium*

^c*Dipartimento di Fisica e Astronomia, Università di Bologna,
Via Irnerio 46, Bologna, I-40126, Italy*

^d*INFN, Sezione di Bologna,
Via Irnerio 46, Bologna, I-40126, Italy*

^e*INFN, Sezione di Milano,
Via Celoria 16, Milano, I-20133, Italy*

^f*TIFLab, Università degli Studi di Milano,
Via Celoria 16, Milano, I-20133, Italy*

E-mail: Stefano.Frixione@cern.ch, fabio.maltoni@uclouvain.be,
davide.pagani@bo.infn.it, marco.zaro@unimi.it

ABSTRACT: Future lepton colliders, such as those based on linear e^+e^- or circular $\mu^+\mu^-$ accelerators, are expected to attain centre-of-mass energies in the multi-TeV range. In this regime the impact of QED and of weak radiation, in both the initial and the final state, can become a leading effect. By employing a general framework presented in a companion paper — suitable for any flavour of colliding leptons — we improve next-to-leading order electroweak predictions by including higher-order contributions, which encompass, but are not limited to, vector-boson-fusion processes. We apply this approach to the study of $t\bar{t}$ and W^+W^- production at a muon collider operating at centre-of-mass energies up to 10 TeV. We show that such an approach, where both QED and weak contributions are included at fixed order, in addition to the all-order resummation of initial-state QED effects, can provide predictions for arbitrary observables in all of the phase space which are precise at the percent level.

KEYWORDS: Higher Order Electroweak Calculations, Higher-Order Perturbative Calculations

ARXIV EPRINT: [2506.10733](https://arxiv.org/abs/2506.10733)

Contents

1	Introduction	1
2	Motivation and goals	3
3	VBF-improved NLO electroweak cross sections	12
4	Results	15
4.1	Differential predictions: general considerations	16
4.2	Inclusion of the Sudakov effects	32
4.3	Factorisation scheme dependence	36
5	Summary and conclusions	42

1 Introduction

A new generation of lepton colliders based on novel acceleration technologies — such as plasma wakefield acceleration [1], and muon cooling/rapid acceleration [2] — might be shown to be viable as early as the next decade. If that will be the case, the prospect of colliding leptons in the multi-TeV energy range will become reality before the mid-21st century, providing one with an unprecedented probe that will help deepen our understanding of fundamental interactions, and in particular of the electroweak ones, by opening up a new regime of exploration, which exploits a unique combination of high-precision measurements and of high-energy observables, with a formidable physics potential [2–8].

This exciting possibility has motivated extensive studies of the physics reach of such machines, both for precision Standard Model (SM) measurements and in the searches for resonant physics Beyond the SM (BSM) — see the references in [2] for a recent overview. A key aspect which has attracted significant attention is the achievable precision and the accuracy of SM predictions in this new regime, where approaches employed so far in e^+e^- simulations will most likely become either unpractical or not sufficiently precise and/or accurate. One prominent example of this situation is the proposal to use electroweak (EW) Parton Distribution Functions (PDFs) [9–15] to describe vector boson fusion (VBF) processes, which are expected to play a significant role in high-energy lepton collisions. This approach aims to account for potentially large logarithmic corrections of collinear and soft origin in the initial state, while also simplifying the relevant calculations.¹ However, all techniques of this kind — while *formally* improving certain aspects of fixed-order results, such as the resummation of EW radiation — rely on kinematic approximations rigorously valid only in asymptotic regimes, which are achieved at logarithmic pace. In addition to that, the current state of the art for EW PDFs is based on LL resummation and LO matrix-element computations, and therefore it lacks control of higher orders. At the energy scales currently envisioned for future muon colliders — of the order of 10 TeV — any EW-PDF approach may

¹Similar methodologies have been proposed to incorporate the effects of final-state EW radiation (see e.g. ref. [16]); we shall not be concerned about them in this work.

prove to be neither accurate nor precise enough. Note that the target precision of a multi-TeV muon collider depends strongly on the observable under consideration. For instance, Higgs coupling measurements — relying on vector-boson-fusion production near threshold — are expected to reach per-mille-level precision. Conversely, indirect searches for new physics in the high-energy tails require percent-level precision.² It is therefore desirable to set up a theoretical framework that is systematically improvable and can ultimately achieve the needed level of precision in all of the phase space.

In view of this, in this work and in its companion paper [17] we adopt a different approach. Firstly, we acknowledge the necessity of resumming initial-state collinear QED radiation into suitable PDFs, and build upon recent analytical and numerical advancements that achieve next-to-leading logarithmic (NLL) accuracy for the latter [18–21]. Secondly, rather than attempting to resum the analogous weak-radiation effects, we include them order by order by means of perturbative calculation of fixed-order predictions, at the next-to-leading order (NLO) in the EW coupling. Thirdly, by means of a general (process- and observable-independent) framework we supplement such QED-resummed and NLO EW-accurate predictions with the contributions stemming from a gauge-invariant and finite subset of next-to-next-to-leading order (NNLO) corrections, thus improving the former predictions in a systematic manner. In particular, such a subset leads to the inclusion of the effects due to γ/Z neutral VBF processes that appear at the NNLO at the matrix-element level, which dominate the threshold region.

By analysing two explicit and non-trivial cases — namely $t\bar{t}$ and W^+W^- production at a muon collider with a centre-of-mass energy of several TeVs — we demonstrate that: *i*) percent-level accurate predictions can be obtained across the entire accessible phase space in a single step; *ii*) there is no indication of any perturbative breakdown in fixed-order results that would necessitate the usage of EW PDFs and/or fragmentation functions to recover a well-behaved series;³ *iii*) in the kinematic regions where the approximations underlying the EW PDFs are valid, the cross section remains relatively small.

Furthermore, we point out that the method we propose is not only process- and observable-independent, but also systematically improvable. For instance, charged-current VBF processes can be included currently up to the NLO without introducing any double counting, while Sudakov suppression effects in the high-energy regime can also be incorporated, e.g. by following the approach of ref. [22].

This paper is organised as follows. In section 2 we outline our motivations and set our goals; we also give the most straightforward evidence as to why an EW-PDF approach is unsuitable for high-energy phenomenology. In section 3 we present our approach, and in section 4 we discuss the results obtained with it for $t\bar{t}$ and W^+W^- production at a muon collider, focusing on the scenarios where its centre of mass energy is equal to 3 and 10 TeV. There, we also discuss two important aspects, i.e. the factorisation-scheme dependence and the necessity to include EW Sudakov logarithm resummation effects in some regions of phase space. We draw our conclusions in section 5.

²This corresponds, assuming that new physics effects scale quadratically with energy, to the target sensitivity of up to $O(10^{-5} - 10^{-6})$ for high-precision observables at electroweak-scale lepton colliders.

³Note, however, that while a perturbative breakdown in fixed-order results may occur at multi-TeV energies, it is generally not related to EW PDFs or fragmentation functions. Instead, it arises from large electroweak Sudakov logarithms due to virtual corrections. See the discussion in section 4.2.

This is a phenomenologically-oriented paper. The reader interested in the technical aspects of our approach can find them in ref. [17]. There is a limited amount of overlap between the two papers as far as the introductory material is concerned; this is expressed in a language which reflects the basic nature, technical or otherwise, of the work to which it belongs.

2 Motivation and goals

So far, e^+e^- colliders have been operated at relatively low energies, where most processes are initiated by lepton annihilation, and final states are characterised by low multiplicities. In this regime, theoretical predictions need not account for complex electroweak partonic dynamics, either in the initial or in the final states.⁴ However, proposed future high-energy machines, such as multi-TeV muon colliders, pose new challenges. As the collider centre-of-mass energy \sqrt{s} increases, leptons begin to exhibit partonic substructure at lower momentum fractions, where the contribution of photons is increasingly important and becomes dominant, with a rate of growth only matched, and eventually exceeded, by that of gluons and quarks. Moreover, higher energies allow one to access richer phase spaces, potentially mimicking hadron-like final states with increased multiplicity.

In general, several different mechanisms will contribute to a given final state. In figure 1 we sketch the typical high-energy behaviour of the differential cross section for the production of a final state F (with invariant mass M_F) with a fraction $\sqrt{\tau}$ of the total energy of the collision (where $\tau = M_F^2/s$). Two leading regions, that correspond to local or global maxima of the cross section, can be clearly identified. The threshold region at small values of $\sqrt{\tau}$ is typically dominated by processes that feature collinear emissions of vector bosons from the incoming lepton lines, which bosons then fuse into the final state F — the underlying hard scattering thus being $VV \rightarrow F$, with $V = \gamma, Z, W^\pm$; there, the $V = \gamma$ contributions are increasingly important for decreasing τ values. Conversely, for values of $\sqrt{\tau}$ close to 1 (i.e. at the kinematic limit), the final state carries almost all of the energy of the collider, and therefore it is $\mu^+\mu^-$ annihilation which gives the dominant contribution.

As one moves from the threshold to the kinematic limit, i.e. from a $\gamma\gamma$ -dominated to a $\mu^+\mu^-$ -dominated situation, one enters, traverses, and then leaves, a region (depicted in white in figure 1) where a non-negligible contribution is given by the heavy electroweak vector bosons emitted quasi-collinearly by the incoming beams, which by colliding produce the final state F . As is shown in the figure, this region typically corresponds to a local *minimum* of the cross section; from a diagrammatic viewpoint, the production mechanism can loosely be associated with VBF topologies.

The cross sections for VBF-like configurations are often approximated by using a parton model that treats W and Z bosons as constituents of the lepton, based on the so-called Effective Weak-boson Approximation (EWA) [24–26]. This approach is only valid in the asymptotic limit of infinite energy, where neglected non-VBF contributions (which are *not* logarithmic in

⁴Here we have mostly in mind initial-state and final-state collinear radiation. Soft-radiation effects in the initial state can be very relevant for precise measurements in situations where kinematic configurations are severely constrained, such as for runs at the Z -pole, and have been studied in great detail [23].




Figure 1. Representative shape of the normalised invariant mass distribution in $\sqrt{\tau}$ (where $\tau = M_F^2/s$) of the final state F at a high-energy muon collider (the actual curve is taken from $t\bar{t}$ production at 10 TeV). The threshold region (green) is dominated by underlying vector boson fusion $VV \rightarrow F$ processes (with $V = \gamma, Z, W^\pm$); the $\gamma\gamma$ contribution is dominant, and there are sizeable power-corrections effects when $V = Z$ and $V = W^\pm$. In the $\sqrt{\tau} \simeq 1$ region (red), the dominant process is $\mu^+\mu^- \rightarrow F$ annihilation. The region where the approximations employed in the EW-PDF LO approach are expected to work best (white) is where the cross section is at its minimum.

the mass of the relevant vector boson) and power corrections vanish.⁵ However, for realistic collider energies (up to ten(s) of TeV), these contributions generally remain significant.⁶

Importantly, VBF-like diagrams are just a sub-class of the complete gauge-invariant set of graphs for a given final state. Thus, if one only considers their contributions (as one does in methods that rely on the EWA or the EW PDFs) one is liable to encounter problems since, in order to achieve accurate predictions useful for phenomenology, an approach must be valid across the entire phase space, with smooth transitions between VBF-dominated and VBF-suppressed regions. In any scenario deemed to be realistic in the next century or so, this entails exact, gauge-invariant computations that retain full dependence on the boson masses without any approximation.

More in detail, while simple in essence the use of EW PDFs is affected by intrinsic limitations, which need to be carefully considered when assessing the reliability of such an approach. In particular, the impact of the following characteristics must be determined.

1. Power-suppressed effects: EW PDFs exactly encode only logarithmic mass dependence. Effects which are polynomial in m_V^2/Q^2 (with $V = W, Z$) and/or in its square root

⁵In this paper we focus on initial-state radiation effects. Other analogous effects, such as final-state fragmentation, which are also enhanced at high energy, can be addressed by means of the same techniques.

⁶Some of these, subject to the condition that they be universal, *may* be included in the EWA (but not in the EW PDFs) at the LO. This is expected to improve the accuracy of the approximation, but to degrade its precision owing to the loss of a clearly-defined separation between the classes of contributions which are included in vs excluded from the results. This matter will be further discussed towards the end of this section.

are lost despite the fact that they can be important, especially near threshold (see footnote 6).

2. Theoretical accuracy: EW PDFs are only available at leading logarithmic (LL) accuracy; in keeping with that, the corresponding short-distance coefficients can currently be computed only at the leading order. Higher-order corrections to PDFs stemming from heavy bosons, and the corresponding adjustments at the matrix elements level, are challenging and not well understood.
3. Diagrammatic coverage: only the subset of VBF-like diagrams is associated with the usage of EW PDFs, and the impact of omitted diagrams cannot be quantified in a process- and, crucially, observable-independent way.
4. Phase space coverage: the kinematical approximation on which the EW PDF approach is based is only valid in a very small part of the phase space, i.e. far from both the threshold and the kinematical boundaries; in that region, the rates are small.
5. Simplified kinematics: simulations based on EW PDFs are fully inclusive over weak boson radiation in the initial state, and in boson decay products; it is therefore difficult to faithfully mimic realistic experimental observables.
6. Error estimation: many uncertainties are not parametric, and thus cannot be reliably quantified.
7. Weak resummation effects: EW PDFs differ little from the EWA. Since these differences are typically smaller than the combined intrinsic uncertainties, the possible benefits of these methods are offset by their limitations.

In order to quantify the impact of the issues mentioned above on the overall accuracy of an EW-PDF approach, one would need to make an extensive comparison with an alternative reliable technique. However, since the goal of this paper is to provide the first phenomenological applications, and an assessment on its scope and precision, of the alternative method introduced in a companion paper (ref. [17]), we defer such a comparison to future work. Here, we limit ourselves to giving an indication of the consequences stemming from some of the issues highlighted above. We do so by presenting two simple, yet quite compelling, arguments. Firstly, we show a tree-level comparison between predictions obtained at the matrix-element level and from the EWA, for a selected class of processes leading to $t\bar{t}$ and W^+W^- production. Secondly, by computing parton-parton luminosities, we show that the effects of the EW resummation as is included in the EW PDFs are mild, and in fact much smaller than the size of the intrinsic approximation which underpins the LO EW PDF approach. Taken together, these comparisons support the conclusions that: *a)* predictions based on either the EWA or the EW PDFs may not be accurate enough to be used in phenomenology, and one would need to check their reliability process by process and observable by observable; *b)* the improvement resulting from trading the EWA for the EW PDFs is much smaller than the effects which are neglected by considering only VBF topologies. In that, one finds an additional motivation for the definition and the implementation of an approach alternative to them.

Let us then consider the production⁷ of an (X, \bar{X}) pair with either $(X, \bar{X}) = (t, \bar{t})$ or $(X, \bar{X}) = (W^+, W^-)$, and analyse the differences between differential distributions obtained by means of:

- The $V_1 V_2 \rightarrow X \bar{X}$ matrix elements, where V_i are either W 's or Z 's (kept on their physical mass shell), convoluted with the relevant EWA functions, henceforth referred to simply as “EWA”.
- The corresponding $\mu^+ \mu^- \rightarrow X \bar{X} \mu^+ \mu^-$ matrix elements, henceforth possibly referred to as “ME”.

In figure 2 we present the ratios of the EWA predictions, obtained with three typical choices for the factorisation scale that enters the EWA functions, over those stemming from the $2 \rightarrow 4$ matrix elements, with the c.m. energy set equal to 10 TeV. The results for $t\bar{t}$ and W^+W^- production are on the left and on the right panels, respectively, with the top, middle, and bottom panel relevant to the pair invariant mass $m(X\bar{X})$, the pseudorapidity of the positively-charged heavy particle $\eta(X)$, and its transverse momentum $p_T(X)$. The coloured band on top of each plot is indicative of the differential rate of the matrix-element-based distribution, ranging from red (large cross section) to yellow tones (small cross section). Actual values for the binned cross section (in pb per bin) are written in a few bins. All of the results have been produced with MADGRAPH5_AMC@NLO (MG5_AMC henceforth) [19, 28–30] whose default EWA implementation is that of ref. [31].

More precisely, the $t\bar{t}$ results in the left column of figure 2 have been obtained as follows. The $2 \rightarrow 2$ process which underpins the EWA result is $ZZ \rightarrow t\bar{t}$; for the EWA functions we adopt the parametrisation of section 3.4 of ref. [12], and this in order to be as self-consistent as possible, since later we shall compare luminosities obtained from the EWA with those from the EW PDFs of that same ref. [12]. Likewise, for a comparison which is as close as possible between the $2 \rightarrow 2$ and $2 \rightarrow 4$ results, in the $\mu^+ \mu^- \rightarrow t\bar{t} \mu^+ \mu^-$ matrix elements we exclude the annihilation channels.⁸ In addition to that, in order to remove the contributions of intermediate photons (and thus their interference with the Z bosons, which is not accounted for in the EWA), we veto all QED interactions.⁹ Simple acceptance cuts are imposed (see the labels in figure 2 or eq. (4.4)), which limit the contribution from regions of the phase space close to its borders — regions where the discrepancies between the EWA and ME approaches would otherwise be even more pronounced. In particular, in this setup additional diagrams contributing to the full $2 \rightarrow 4$ process (such as the radiation of a Z/γ^* from one muon line, which then branches into a $t\bar{t}$ pair) are kinematically suppressed, and the dominant contribution stems from an effective $ZZ \rightarrow t\bar{t}$ topology. This is a deliberately academic

⁷For a phenomenological application at hadron colliders, see also ref. [27].

⁸This is a gauge-invariant class. Technically, it can be obtained by requiring the matrix element generator to treat μ^+ and μ^- as if they had different flavours, e.g. by replacing μ^- with e^- both in the initial and final state. In this way, topologies such as $\mu^+ \mu^- \rightarrow t\bar{t}(\gamma/Z \rightarrow) \mu^+ \mu^-$ are safely excluded. We shall use the same method also for W^+W^- production, to be discussed below.

⁹While in general this operation would violate gauge-invariance, as e.g. in the case of $\mu^+ \mu^- \rightarrow W^+W^- \mu^+ \mu^-$, this does not happen in the present case, thanks to the fact that $\mu^+ \mu^- \rightarrow t\bar{t} \mu^+ \mu^-$ diagrams do not feature W bosons. Having said that, we note that at muon colliders with energies up to tens of TeV, the Z -pair induced contribution to $t\bar{t}$ production has a very small rate with respect to the photon-induced channel.

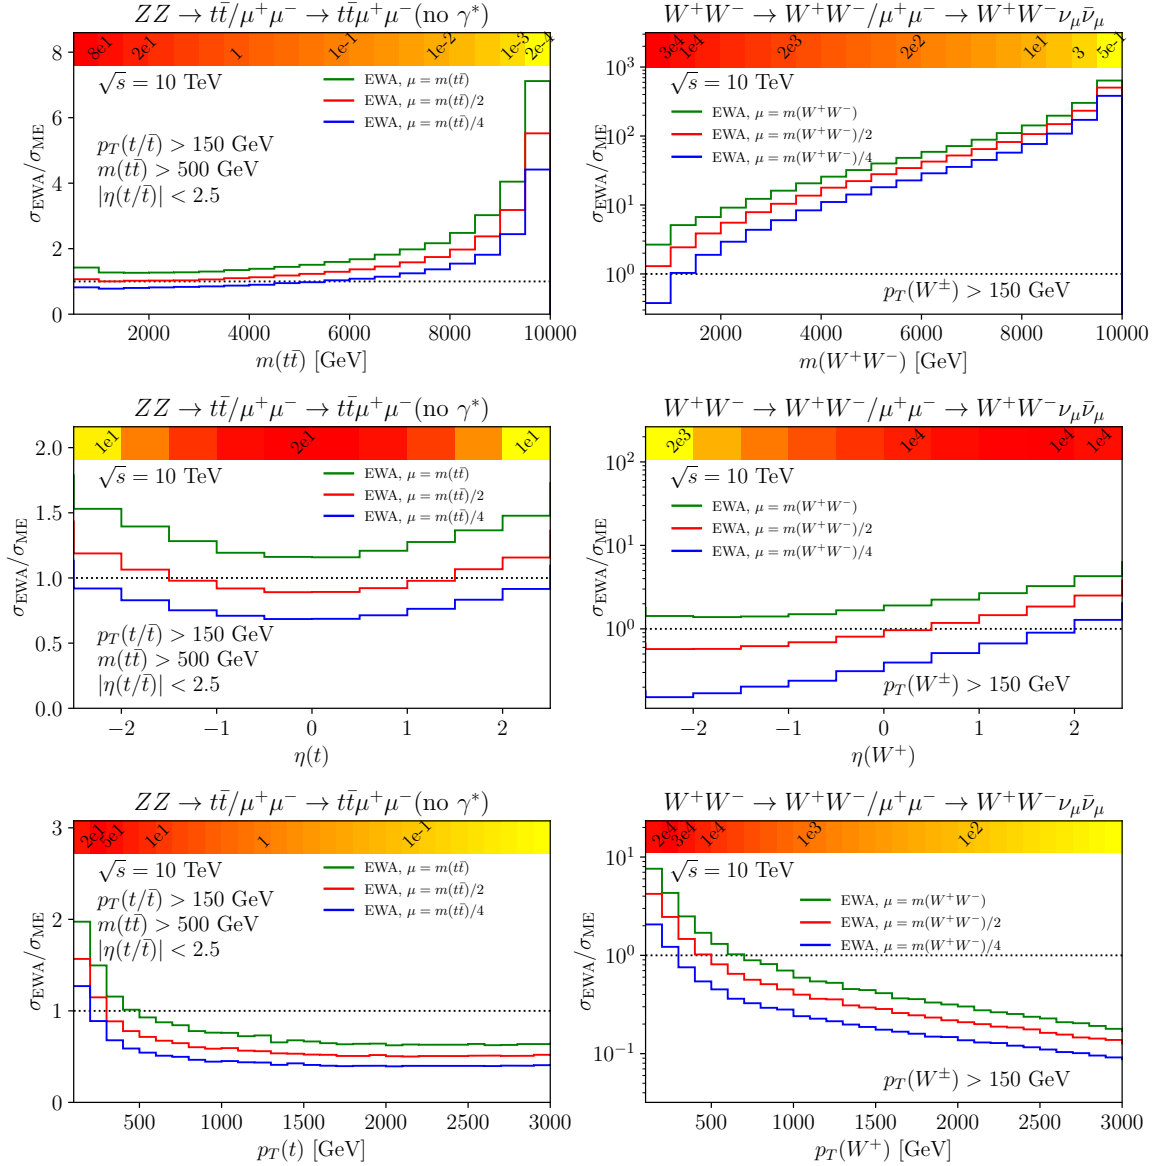


Figure 2. Ratios of the EWA predictions ($2 \rightarrow 2$) over those of the matrix elements ($2 \rightarrow 4$) for $m(t\bar{t})$, $\eta(t)$, and $p_T(t)$ (left panels), and their analogues in W^+W^- production (right panels) — see the text for details.

scenario: both the process definition (diagram selection) and the cuts have been chosen so that differences are minimised; in other words, this is a best-case situation.

One starts by observing that the EWA exhibits a significant scale dependence, which in itself precludes a precise comparison with the full result. This dependence is due to the transverse polarisations of the initial-state Z bosons, since the longitudinal modes do not depend on the scale. Focusing on the top panel on the left column, one sees that for smaller scale choices and at moderate invariant masses (approximately smaller than 6 TeV), the EWA and the ME agree at the 20–30% level. However, the shapes of the pseudorapidity

distributions are considerably different, leading to discrepancies of a similar magnitude. The situation is even worse in the case of $p_T(t)$, where the discrepancies between the EWA and the ME predictions can be as large as 100%, even under the most favourable scale choices.

The bottom line is this: while the EWA and ME descriptions can be made to agree under selective choices of kinematical configurations and scale parameters (and still within significant uncertainties), this is at best true in some one-dimensional projections of the phase space, i.e. for one observable at a time. When one considers the phase space in its true multi-dimensional nature, even this rough agreement breaks down.

We now turn to discussing the case of W^+W^- production, whose results are shown in the right-hand panels of figure 2. In this case, our goal is to compare the EWA and ME descriptions emerging from charged-current (sub)processes. Thus, in the case of the EWA we select W^+W^- fusion, whereas for the full process $\mu^+\mu^- \rightarrow W^+W^-\nu_\mu\bar{\nu}_\mu$ neutral currents contributions involving $\mu^+\mu^-$ or $\nu_\mu\bar{\nu}_\mu$ annihilation and/or creation are excluded (see footnote 8). At variance with the case of $t\bar{t}$ production, we aim to explore a larger fraction of the phase space in order to assess the full spectrum of possibilities. So while we relax the acceptance cuts, these cannot be completely eliminated, in view of the presence of a pole in the $W^+W^- \rightarrow W^+W^-$ scattering amplitude, due to the exchange of a photon in the t channel; this, we avoid by imposing $p_T(W^\pm) > 150$ GeV. Indeed, in general for a given partonic c.m. energy \sqrt{s} , by requiring $p_T(W^\pm) > p_T^{\text{cut}}$, in the limit $p_T^{\text{cut}} \rightarrow 0$ one obtains a cross section proportional to $(\sqrt{s}/p_T^{\text{cut}})^2$. It is important to note that such a pole is actually an artefact of the EWA (and EW PDF) approach, because it is not present in the full $2 \rightarrow 4$ matrix element. There, the W 's stemming from the muon branchings (which enter the sub-diagrams where they play the same role as the initial-state W 's in the $2 \rightarrow 2$ matrix elements employed by the EWA) are space-like, and therefore off-shell, and thus automatically regulate the potential t -channel divergence. In conclusion, if no cut is imposed, the EWA simply leads to a divergent cross section, and therefore cannot be used sensibly. Even by imposing a $p_T(W^\pm)$ cut, as we have done, results are still IR-sensitive, owing to the p_T^{cut} dependence mentioned above. Indeed, the inspection of figure 2 shows extremely large discrepancies between the EWA and the ME approaches in both shapes and rates — the EWA approximation is not under control.

We note that even in the simplified setup we have adopted, where the incoming muons are treated as if they had different flavours (see footnote 8), $\mu^+\mu^- \rightarrow W^+W^-\nu_\mu\bar{\nu}_\mu$ features three main classes of graphs which result in possible kinematic enhancements; these classes can in turn be loosely identified with underlying $2 \rightarrow 2$ reduced processes, namely (see figure 3):

- (a) vector boson scattering $W^+W^- \rightarrow W^+W^-$; this leads to the largest enhancement, owing to the two spin-1 particles exchanged in the t channels of the $2 \rightarrow 4$ diagrams (a.k.a. VBF).
- (b) $fW \rightarrow fW$ scattering, stemming from one (in either the final or initial state) muon branching into a W boson and one branching into one fermion (a.k.a. single- W radiation).
- (c) $f\bar{f} \rightarrow f\bar{f}$, stemming from two (in either the final or initial state) muon branchings into a W boson (a.k.a. double- W radiation).




Figure 3. Representative Feynman diagrams of the three classes of contributions to $\mu^+\mu^- \rightarrow W^+W^-\bar{\nu}_\mu\nu_\mu$ which are kinematically enhanced at high energy, namely: (a) VBF scattering (reduced process $W^+W^- \rightarrow W^+W^-$); (b) single- W radiation, either from initial or final state (reduced process $fW \rightarrow fW$); (c) double- W radiation, either from initial or final state (reduced process $f\bar{f} \rightarrow f\bar{f}$). Here, $f = \mu^\pm, \nu_\mu, \bar{\nu}_\mu$; typically, the impact of these contributions is largest for (a), and smallest for (c).




Figure 4. Comparison of typical parton-parton luminosities computed with the EWA (red curves) and the EW PDFs (black curves).

While the $p_T(W^\pm)$ cut we have imposed is expected to suppress such contributions, some leftover will unavoidably remain, hampering the agreement between the EWA and the full matrix elements results. Again, this is a best-case scenario: indeed, if we considered the physical case where μ^+ and μ^- do have the same flavour, we would get *also* annihilation channels, such as $\mu^+\mu^- \rightarrow W^+W^-(Z \rightarrow)\bar{\nu}_\mu\nu_\mu$. None of these is VBF-like, and therefore the agreement between the EWA and the full matrix elements predictions cannot but degrade further.

In summary, it is clear that in situations such as the one discussed for W^+W^- production (i.e. those which feature other topologies in addition to the VBF one), any approach capable of retaining the full information associated with the exact $2 \rightarrow 4$ matrix element is superior to the EWA based on the $2 \rightarrow 2$ matrix elements; this conclusion extends to the EW-PDF approach, which employs the same matrix-element framework as the EWA.

Conversely, for what concerns cases such as $t\bar{t}$ production, and thus its underlying ZZ -initiated neutral-current contribution, the situation is arguably more favourable, which can

easily be understood as follows. W^+W^- production features particles in the final state (the W 's) which also directly enter the EW-PDF resummation procedure, which is in turn related to the fact that they can be produced (in particular, collinearly) through different parton-parton scatterings as well — see figure 3(c). This implies that a merging of different final-state multiplicities would be mandatory in order for the W 's to be fully accounted for in all of the phase space. On the other hand, the $t\bar{t}$ pair can only be generated in high- Q^2 interactions. One therefore wonders whether by “upgrading” the EWA description of $t\bar{t}$ production by means of an EW-PDF approach, an improvement in the precision of the cross section would follow, whose impact would be larger than the loss of accuracy due to the approximations inherent to the $2 \rightarrow 2$ scattering. In other words we pose the question: is the difference between the EWA and the EW-PDF results larger than that between the EWA and the $2 \rightarrow 4$ matrix elements results? If the answer is negative, the upgrade is phenomenologically irrelevant even in this best-case scenario, and it shows that the resummation included in the EW PDFs is subleading and does not provide one with a tangible improvement.

In order to answer this question (indirectly, given that at present MG5_AMC does not support the publicly-available computer codes for EW PDFs), we estimate the impact of the resummation included by construction in the EW PDFs, but not present in the EWA, on the luminosities relevant to the processes under consideration — this exercise is significant in that the EWA functions can be seen as the initial conditions for EW-PDF evolution, and therefore differ from the latter only owing to higher-order terms. Thus, we compare the luminosities computed by using the EWA functions with those obtained from the EW PDFs. In figure 4 we present the results for two such comparisons, relevant to the transversely-polarised W^+W^- pair (left panel) and ZZ pair (right panel); we remind the reader that, for the self-consistency reason mentioned before, both the EWA functions and the EW PDFs are those of ref. [12]. The plots display the luminosities (again, we have considered a 10 TeV muon collider) as a function of the invariant mass M of a putative final-state system (e.g. the $t\bar{t}$ pair), which are all normalised (i.e. divided by) the luminosity obtained with the EW PDFs evaluated at the scale $\mu = M/2$, taken to be our central choice. The black and red curves give the EW-PDF and EWA results, respectively. The predictions obtained by setting $\mu = M$ and $\mu = M/4$, are shown as well (with the same normalisation as before).

It is manifest that the impact of the resummation of the EW logarithms (the difference between PDFs and EWA curves at the same scale) is minimal compared to the dependence of the luminosities on the factorisation scale, which is consistent with the one shown in figure 2 for the invariant mass distributions. Their hierarchy is reverted only at $m(X) \gtrsim 5$ TeV, where on the other hand weak VBF processes are by far subdominant. Most importantly, in the range considered, the impact of the resummation of the EW logarithms is much smaller than the difference between $2 \rightarrow 4$ matrix elements and the EWA approximation.

Even in the favourable case of $t\bar{t}$ production, the situation of the EW PDFs can be compared to an example relevant to LHC physics, where one would use a pure LL resummed prediction for the Z transverse momentum in comparison to experimental data. Such a comparison would utterly fail, for two reasons: insufficient logarithmic accuracy, and lack of any non-trivial matrix element information; the former prevents the agreement with data in the peak region and below ($p_T(Z) \lesssim 4$ GeV), and the latter in the complementary region (where hard recoils are important). The only positive feature of such a prediction would be

qualitative, in that it would confirm the existence of a small- $p_T(Z)$ peak, and the finiteness of the cross section at $p_T(Z) \rightarrow 0$. The methods needed to improve this prediction and its comparison to data are well known, and all include an increase of the logarithmic accuracy, and a matching with matrix elements.

Exactly as for the LL-accurate $p_T(Z)$ unmatched resummation just mentioned, the current EW PDFs results only retain the LLs, do not include any matching with non-trivial matrix elements, and work for a single observable at a time. Therefore our results, summarised by figure 2, are not surprising:¹⁰ the best one can hope for is a qualitative description. Precisely as for the $p_T(Z)$ example, the situation can be improved by increasing the logarithmic accuracy *and* by matching to matrix elements,¹¹ neither of which options appears to be around the corner.

For what only concerns the EWA, if one stipulates that its functions are not to be used as EW-PDF initial conditions (in which case they would be subject to more stringent constraints), then one has a certain freedom to include in it non-logarithmic terms — in the context of the $p_T(Z)$ example given above, this would amount to comparing the EWA with fixed-order Z -production matrix elements. While it is clear that the inclusion of non-logarithmic terms would improve this comparison, the non-universality of such terms, and/or the presence of analogous terms that do not factorise (and therefore *can not* be employed in the EWA) imply an immediate loss of control over the accuracy of the approximation. This can only be established *a posteriori* through a comparison with the very matrix elements whose computations one seeks to avoid by introducing the EWA in the first place — thus, one ends up with a circular argument.¹²

Given the limitations of the EWA and EW-PDF results outlined above, we are compelled to seek an alternative approach. Let us specify the requirements that such an approach should have.

Firstly, it needs to employ exact matrix elements, incorporating without any approximation mass effects for the weak bosons and other particles (possibly bar the truly light ones, i.e. electrons, muons, and quarks of the first and second families) at the electroweak scale, and thereby including the complete set of contributing topologies and diagrams. Secondly, it must allow, at least in principle, the inclusion of higher-order corrections — specifically, of next-to-leading order (NLO) and beyond. Thirdly, it should feature the resummation of small-lepton-mass logarithms, using available PDFs for leptons, photons, and, optionally, quarks and gluons, up to next-to-leading logarithmic (NLL) accuracy. Finally, it must provide one with the full and smooth coverage of the phase space, ensuring no loss of information across the entire energy range — from the threshold to asymptotic regimes.

¹⁰While figure 2 features the EWA, and not the EW-PDF, results, it is legitimate to use it to draw conclusions on the latter. Specifically, in the context of the present example, while the $p_T(Z)$ resummation is mandatory ($p_T(Z)$ tends to zero), that performed by the EW PDFs is not (it takes a lot of energy for m_Z to be negligible w.r.t. the hard scale of the process).

¹¹We point out that, even if not strictly necessary, higher-logarithmic accuracy would be needed for a better matching with higher-order matrix elements.

¹²A recent implementation of the EWA which includes some subleading universal terms in the vector-boson mass [32] does not appear to improve the quality of the EWA vs matrix-element agreement in a substantial manner, in particular for those processes where additional (non-VBF) topologies are present.




Figure 5. Representative Feynman diagrams for the various classes of contributions to $t\bar{t}$ production. (a) and (b) correspond to the two production channels at LO. Diagram (c) appears as NLO real correction, while diagram (d) is formally an NNLO double-real correction.

With these criteria in mind, we now present a strategy (introduced in ref. [17]) designed to overcome the limitations of the conventional EWA- and EW-PDF-based approaches. This strategy isolates VBF-like configurations that emerge at the NNLO in a gauge-invariant and theoretically controlled manner, framing them as an extension of NLO predictions. Such configurations correspond to double-real diagrams, and are matched to collinear resummation for light particles.

3 VBF-improved NLO electroweak cross sections

In this section we summarise a few key ingredients of the method introduced in ref. [17]; the reader is encouraged to check that paper for more details on the procedure and its numerical validation.

We begin by stating the fact that the automatic computation of the complete NLO EW corrections for an arbitrary process at lepton colliders is available in MG5_AMC [29, 30]. Here, we show how the corresponding predictions can be augmented to consistently account for all NNLO double-real processes that include VBF topologies in an exact manner, while avoiding double counting and resumming collinear lepton-mass logarithms of QED origin. We shall briefly go through these features in the remainder of this section; the reader must bear in mind that the following discussion is kept intuitive intentionally.

In order to be definite, let us consider $t\bar{t}$ production at a muon collider. Lest confusion should arise, we use a different notation to tell the physical muon that constitutes the beam (μ) from the muon parton (μ) which is inside the former. At a mass scale equal to the muon mass, the “probability density of finding a parton muon in the physical muon” is just equal to a Dirac delta, its argument a function of the parton momentum fraction ζ , thus: $\Gamma_{\mu^\pm/\mu^\pm}(\zeta, m_\mu^2) = \delta(1 - \zeta)$. At higher scales Γ_{μ^\pm/μ^\pm} evolves into a different distribution

which has a $\zeta < 1$ tail. In addition to that, a non-zero probability of finding a photon in the physical muon μ , $\Gamma_{\gamma/\mu^\pm}(\zeta)$, exists, which instead peaks at small values of ζ . Initial conditions and explicit expressions for these parton distribution functions are now available at the NLO+NLL accuracy [18–21] for both electrons and muons.

In this framework, at the LO $t\bar{t}$ production proceeds through two partonic channels, $\mu^+\mu^- \rightarrow t\bar{t}$ (annihilation) and $\gamma\gamma \rightarrow t\bar{t}$ (fusion) — see figure 5(a) and (b), respectively. At the NLO in the EW theory, real and virtual diagrams enter. Among the real corrections, a new class of diagrams arises, e.g., $\gamma\mu^- \rightarrow t\bar{t}\mu^-$ with a Z or a γ exchanged in a t -channel, figure 5(c). In the kinematic configurations enhanced by collinear emissions, this contribution can be seen as “half” of a VBF-type contribution underpinned by a $(Z/\gamma^*)(Z/\gamma^*) \rightarrow t\bar{t}$ process, as in figure 5(d). At the NNLO two-loop-virtual, one-loop-real, and double-real diagrams enter; among the latter, one finds the VBF-type contributions mentioned above and depicted in figure 5(d). When squaring the amplitudes, these contributions mix diagrams featuring Z - and γ -exchanges in the t -channels which cannot be disentangled — in other words, Z - and γ -exchange diagrams interfere with each other. However, in the collinear limits, where the outgoing muons are either forward or backward w.r.t. their parent particles, the photon poles can be uniquely identified as coming from two different fermion lines,¹³ and the ensuing collinear divergences can then be reabsorbed into the corresponding muon PDFs. While at the NLO one would need to consider only one such divergence at a time, at the NNLO double divergences appear where there are two incoherent NLO-type singular structures, one for each of the two incoming fermion lines. This implies that the full massless-muon matrix element $\mu^+\mu^- \rightarrow t\bar{t}\mu^+\mu^-$ can be rendered finite by means of a suitable set of insertions of collinear counterterms, and subsequently integrated over the phase space. These insertions are compensated by analogous terms that are included in the PDFs, that thus effectively resum collinear logarithms of QED origin. Conversely, VBF-type contributions mediated by weak charged currents that contribute to $\mu^+\mu^- \rightarrow t\bar{t}\nu_\mu\bar{\nu}_\mu$ do not feature any collinear divergences, and can simply be added (at the squared-amplitude level) to all of the other contributions to the cross section, without needing any counterterms.

With the help of the basic elements listed just above, the essence of our approach is the following: we start with an NLO EW result for the process of interest, and consider the full set of amplitudes that first appear at the NNLO. If these do not feature QED-type initial state divergences, such as those mediated by charged weak currents, one can just square and add them to the NLO result. Otherwise, one includes the complete amplitude squared (thus including Z/γ interference effects) plus a suitable set of collinear counterterms for the photon currents that render the former amplitude squared infrared-finite, while at the same time avoiding double counting with terms of NLO origin. Working out the details of this procedure is not straightforward, and one needs to take care of several technical aspects. As was already mentioned, these are described fully in ref. [17]; here we limit ourselves to presenting the final result, restricting the notation to the case of a muon collider.

Following ref. [17], the differential cross section for producing a given final state in $\mu^+\mu^-$ collisions is written thus:

$$d\sigma(P_1, P_2) = d\zeta_1 d\zeta_2 \sum_{ij} \Gamma_{i/\mu^+}(\zeta_1) \Gamma_{j/\mu^-}(\zeta_1) d\hat{\sigma}_{ij}(\zeta_1 P_1, \zeta_2 P_2), \quad (3.1)$$

¹³This is so owing to the fact that collinear configurations are semi-classical.

where P_i are the momenta of the physical muons (μ^\pm), for the summation indices we use $i = \mu^+, \gamma$ and $j = \mu^-, \gamma$,¹⁴ and the short-distance cross sections can be conventionally written in terms of the electroweak coupling α as follows:

$$d\hat{\sigma}_{ij} = d\hat{\sigma}_{ij}^{[0]} + \frac{\alpha}{2\pi} d\hat{\sigma}_{ij}^{[1]} + \left(\frac{\alpha}{2\pi}\right)^2 d\hat{\sigma}_{ij}^{[2]} + O(\alpha^3), \quad (3.2)$$

where all of the $d\sigma_{ij}^{[k]}$ coefficients are understood to be of the same perturbative order as the Born amplitude squared, i.e. α^b with b a process-dependent integer (e.g., $b = 2$ for $t\bar{t}$ and W^+W^- production).

The keeping of only the two leftmost terms in eq. (3.2) corresponds to an NLO computation, which as we said we perform automatically with MG5_AMC. Conversely, the rightmost term is the NNLO contribution. Its full computation in a process- and observable-independent way within the EW theory is presently unfeasible. However, as was sketched in the discussion above, some of its terms (in particular, those stemming from VBF topologies which, as we shall show, are non-negligible for the phenomenology of high-energy muon collisions) can be singled out and rendered finite [17]. More in detail, a well-defined gauge-invariant¹⁵ subset of the NNLO corrections in the partonic $\mu^+\mu^-$ channel can be identified together with suitable subtractions; this contribution, which essentially consists in the double real-emission from the initial state, we call $d\hat{\Sigma}_{\delta\text{NNLO}_\Gamma}$:

$$\left(\frac{\alpha}{2\pi}\right)^2 d\hat{\sigma}_{\mu^+\mu^-}^{[2]} \ni d\hat{\Sigma}_{\delta\text{NNLO}_\Gamma}(p_1, p_2), \quad (3.3)$$

and we write it thus (see eq. (7.3) of ref. [17])

$$\begin{aligned} d\hat{\Sigma}_{\delta\text{NNLO}_\Gamma}(p_1, p_2) \equiv & \left(\frac{1}{1-y_1}\right)_{\delta_1} \left(\frac{1}{1-y_2}\right)_{\delta_1} \left((1-y_1)(1-y_2)\mathcal{M}_{\mu^+\mu^-}^{(m+2)}\right) d\phi_{m+2}(p_1, p_2) \\ & + \frac{\alpha}{2\pi} \mathcal{Q}_{\gamma\mu}^{(\delta_1)'}(z_2) \left(\frac{1}{1-y_1}\right)_{\delta_1} \left((1-y_1)\mathcal{M}_{\mu^+\gamma}^{(m+1)}\right) d\phi_{m+1}(p_1, z_2 p_2) dz_2 \\ & + \frac{\alpha}{2\pi} \mathcal{Q}_{\gamma\mu}^{(\delta_1)'}(z_1) \left(\frac{1}{1-y_2}\right)_{\delta_1} \left((1-y_2)\mathcal{M}_{\gamma\mu^-}^{(m+1)}\right) d\phi_{m+1}(z_1 p_1, p_2) dz_1 \\ & + \left(\frac{\alpha}{2\pi}\right)^2 \mathcal{Q}_{\gamma\mu}^{(\delta_1)'}(z_1) \mathcal{Q}_{\gamma\mu}^{(\delta_1)'}(z_2) \mathcal{M}_{\gamma\gamma}^{(m)} d\phi_m(z_1 p_1, z_2 p_2) dz_1 dz_2, \end{aligned} \quad (3.4)$$

where p_i are the momenta of the partonic muons.

Equation (3.4) includes all neutral VBF-like contributions, and is structured so that it can be added to an NLO result without double counting. Here m is the number of final state particles of the Born process (e.g. $m = 2$ for $t\bar{t}$ or W^+W^- production). $\mathcal{Q}_{\gamma\mu}^{(\delta_1)'}(z)$ is a function emerging from the phase-space integration of a splitting function, plus a term that removes any double counting (see eq. (5.20) of ref. [17]), while $\mathcal{M}_{ij}^{(m+0,1,2)}$ are the matrix elements for the production of the final-state system of interest, in the collisions of partons

¹⁴This is for simplicity; predictions aiming at reproducing as closely as possible an experimental setup should use all parton types (i.e. all light leptons, all quarks, the photon, and the gluon); the formalism does not require any modifications to account for this extension.

¹⁵Gauge invariance is ensured by not discarding any diagram for the double real-emission process.

i and j with 0, 1, or 2 extra light particles in the final state; by $d\phi_n$ we denote an n -body phase space. Finally, the quantities in the large round brackets are plus distributions (which implement the subtractions stemming from the collinear counterterms), modified in such a way that the subtractions are only performed in a restricted region, defined by the arbitrary parameter δ_1 . We point out that this is the same parameter that enters the $\mathcal{Q}_{\gamma\mu}^{(\delta_1)'}$ functions; by design, eq. (3.4) is independent of this parameter, which constitutes a useful tool to check the correctness of any numerical implementation of this formula. The $\mathcal{Q}_{\gamma\mu}^{(\delta_1)'}$ functions and the subtractions are directly connected with compensating terms included in the PDFs with which the short-distance cross sections are convoluted. This connection, technically called a factorisation scheme, is controlled by certain other functions, which the physical cross section must be independent of — we shall return to this point in section 4.3.

In summary, eq. (3.4) incorporates the $2 \rightarrow 2 + m$ matrix elements where QED initial state collinear divergences have been subtracted, and included in the PDFs with which one convolutes the cross section according to eq. (3.1). Equation (3.4) enables the exact retention of both logarithmic and power-suppressed terms in the Z boson mass at relative $\mathcal{O}(\alpha^2)$, as well as that of all the effects due to Z/γ interference. At the same time, this framework resums all light-fermion mass effects via lepton PDFs, at the logarithmic accuracy at which such PDFs are defined. Overall, eq. (3.4) allows one to significantly improve any NLO EW predictions in those kinematic regions characterised by VBF-like topologies or, more broadly, by partonic interactions underpinned by $\gamma\gamma$, γZ , and ZZ fusion.

We conclude by pointing out that the subtractions that appear in eq. (3.4) only correspond to a subset of the complete NNLO collinear counterterms; therefore, we understand the presence of collinear cutoffs, without which one would still find singularities. This point is discussed at length in ref. [17]: the bottom line is that such cutoffs are harmless physics-wise, and their presence can therefore be ignored in the context of phenomenological simulations.

4 Results

In order to illustrate the practical consequences of the method presented in section 3, we now consider two processes in detail, namely the inclusive production of a $t\bar{t}$ pair, and that of a W^+W^- pair. The parameters employed in our simulations are set as follows:

$$m_t = 173.3 \text{ GeV}, \quad m_W = 80.419 \text{ GeV}, \quad m_Z = 91.188 \text{ GeV}, \quad m_H = 125 \text{ GeV}. \quad (4.1)$$

All decay widths are set equal to zero. We work in the G_μ renormalisation scheme [33], where the EW parameters are determined after one chooses the value of the Fermi constant, which we set thus

$$G_\mu = 1.16639 \cdot 10^{-5} \text{ GeV}^{-2}. \quad (4.2)$$

The muon and photon PDFs inside the physical muon are computed at the NLO+NLL accuracy, following refs. [18, 19, 21, 34], in the so-called Δ factorisation scheme [35]; technically, they are based on EMELA [19, 20]. The factorisation scale is set equal to the invariant mass of the pair:

$$\mu_F = m(X\bar{X}), \quad (4.3)$$

with $(X, \bar{X}) = (t, \bar{t})$ or $(X, \bar{X}) = (W^+, W^-)$. Two different centre-of-mass energies are considered, namely $\sqrt{s} = 3 \text{ TeV}$ and $\sqrt{s} = 10 \text{ TeV}$; in the latter case, in addition to the predictions obtained by integrating over the whole phase space, we also present those that stem from imposing some fiducial cuts, whereby we require:

$$p_T(X), p_T(\bar{X}) > 150 \text{ GeV}, \quad m(X\bar{X}) > 500 \text{ GeV}, \quad |y(X)|, |y(\bar{X})| < 2.5. \quad (4.4)$$

As is discussed at length in ref. [17], some technical cuts are necessary in order to screen uncanceled singularities arising in the δNNLO_Γ corrections; we employ an invariant-mass cut on the outgoing muons:

$$m(\mu^+\mu^-) > 200 \text{ GeV}. \quad (4.5)$$

As is demonstrated in ref. [17], the effect of such a cut (and of a more general class of technical cuts that possibly include (pseudo)rapidity cuts on the outgoing muons) on physical distributions is phenomenologically negligible.

While presenting our results, we focus on different aspects in turn. We begin with a general discussion in section 4.1; process-specific results are then introduced in sections 4.1.1 and 4.1.2 for $t\bar{t}$ and W^+W^- production, respectively.¹⁶ In section 4.2 we show how the predictions for some observables can be further improved by including higher-order, yet sizeable, effects due to EW Sudakov logarithms. Finally, we discuss the impact of the choice of factorisation scheme in section 4.3.

4.1 Differential predictions: general considerations

In this section, we discuss the general features which are common to differential observables in $t\bar{t}$ - and W^+W^- -pair production, while process-specific characteristics are presented in sections 4.1.1 and 4.1.2.

In keeping with eqs. (3.1) and (3.2), and with the notation introduced in ref. [17], we remind the reader that by NLO and NNLO we understand the predictions accurate *up to* those respective orders, i.e. NLO results include both the $\mathcal{O}(\alpha^b)$ and the $\mathcal{O}(\alpha^{b+1})$ contributions, whereas the NNLO results are obtained by adding the $\mathcal{O}(\alpha^{b+2})$ contributions to the NLO predictions. Conversely, we denote by δNLO and δNNLO the sole $\mathcal{O}(\alpha^{b+1})$ and $\mathcal{O}(\alpha^{b+2})$ terms, respectively. Furthermore, the LO, $(\delta)\text{NLO}$, and $(\delta)\text{NNLO}$ contributions which emerge from the procedure of ref. [17] are denoted by LO_Γ , $(\delta)\text{NLO}_\Gamma$, and $(\delta)\text{NNLO}_\Gamma$, respectively.¹⁷

All of the figures of sections 4.1.1 and 4.1.2 have the same layout, thus: there are three panels with the results obtained at 3 TeV (top left), 10 TeV (bottom left), and 10 TeV with the acceptance cuts of eq. (4.4) (bottom right). A fourth panel (top right) reports the labels relevant to the various histograms.

The layouts of the three panels that constitute a figure are also identical to one another, and feature a main frame and three insets. In the main frame, we show the LO, $\text{LO}_{\gamma\gamma}$,

¹⁶Some results at NLO accuracy for these (and other) processes, computed by neglecting any photon-initiated contribution, have been discussed in refs. [36, 37]. In the former, the software WHIZARD [38] has been employed; in the latter, the kinematic region of interest has been restricted to the one dominated by muon-muon annihilation.

¹⁷See section 6 of ref. [17] for precise definitions and further details. Notably, the LO_Γ term coincides with the pure $\gamma\gamma$ -initiated contribution and can therefore also be denoted as $\text{LO}_{\gamma\gamma}$.

NLO, and NNLO_Γ differential distributions, as pink-dotted, purple-dotted, green-dashed and black-solid histograms, respectively. In addition to these, we also show (by using full symbols) the LO-accurate results for processes which lead to the same $X\bar{X}$ final state, but are *not* included in the NNLO_Γ predictions.¹⁸ Such processes are the associated production of the $X\bar{X}$ pair of relevance with either a neutrino pair ($X\bar{X}\nu\bar{\nu}$), or an on-shell Z boson ($X\bar{X}Z$), and are of interest because they account for effects due to EW radiation and to charged-current contributions (i.e., W^+W^- -fusion). For the neutrino-pair associated production, we consider both the case where no cuts are imposed on the neutrinos (cyan circles), and the one where we require the neutrinos to be muonic and that they have an invariant mass larger than 200 GeV (fuchsia diamonds). The Z -associated production results are displayed as yellow stars. The rationale for considering neutrino-pair associated production both without and with invariant mass cuts is because in the latter case the contributions stemming from a resonant Z boson (and which are already included in the $X\bar{X}Z$ results) are effectively discarded, and what is left is expected to be dominated by WW -fusion.

The top inset shows the relative impact of the LO, LO_{γγ}, $X\bar{X}\nu\bar{\nu}$, and $X\bar{X}Z$ contributions, defined as the ratio between these predictions and the NNLO_Γ one — the same plotting patterns as in the main frame are employed here. Similarly, the middle inset presents the relative impact (defined as before) of the δ NLO and δ NNLO_Γ corrections (green-dashed and black-solid histograms, respectively, as is done in the main frame); the latter is multiplied by a factor of 2 in order to enhance its visibility. In addition to those, we also display the δ NNLO_Γ (blue-dotted) and the δ NLO corrections (red-dashed), the latter of which computed by setting the finite part (according to the FKS conventions) of the virtual amplitudes equal to zero (and hence denoted by δ NLO-noV) — the latter two predictions help understand the “origin” of EW corrections, as we shall see in the following. Finally, in the bottom inset, we plot the uncertainty bands obtained by varying the factorisation scale up and down by a factor of two around its central value (see eq. (4.3)) for the LO (dotted pink), NLO (shaded green) and NNLO_Γ (shaded grey) predictions.

4.1.1 $t\bar{t}$ inclusive production

As was anticipated, this section collects our predictions for several differential distributions stemming from the production of a $t\bar{t}$ pair. We start with the invariant mass of the pair (figure 6); we then consider the transverse momentum of the top quark (figure 7) and that of the pair (figure 8), and finally the rapidity of the top quark (figure 9) and that of the pair (figure 10).

The $t\bar{t}$ invariant mass distribution, shown in figure 6, is crucial in order to understand many key features of this production process. In general, the $\gamma\gamma$ -initiated contribution is peaked at small values of this observable, while that emerging from $\mu^+\mu^-$ annihilation peaks at large invariant masses (in fact, at the kinematic limit), a behaviour which is already evident at the LO (the purple-dotted histogram rapidly approaches the pink-dotted one when moving towards the threshold). This is a consequence of the fact that the photon PDF grows with decreasing Bjorken x , while the opposite happens with the muon PDF —

¹⁸We opt for a LO simulation for simplicity, although there is no technical limitation which prevents one from simulating such processes at the NLO with MG5_AMC.

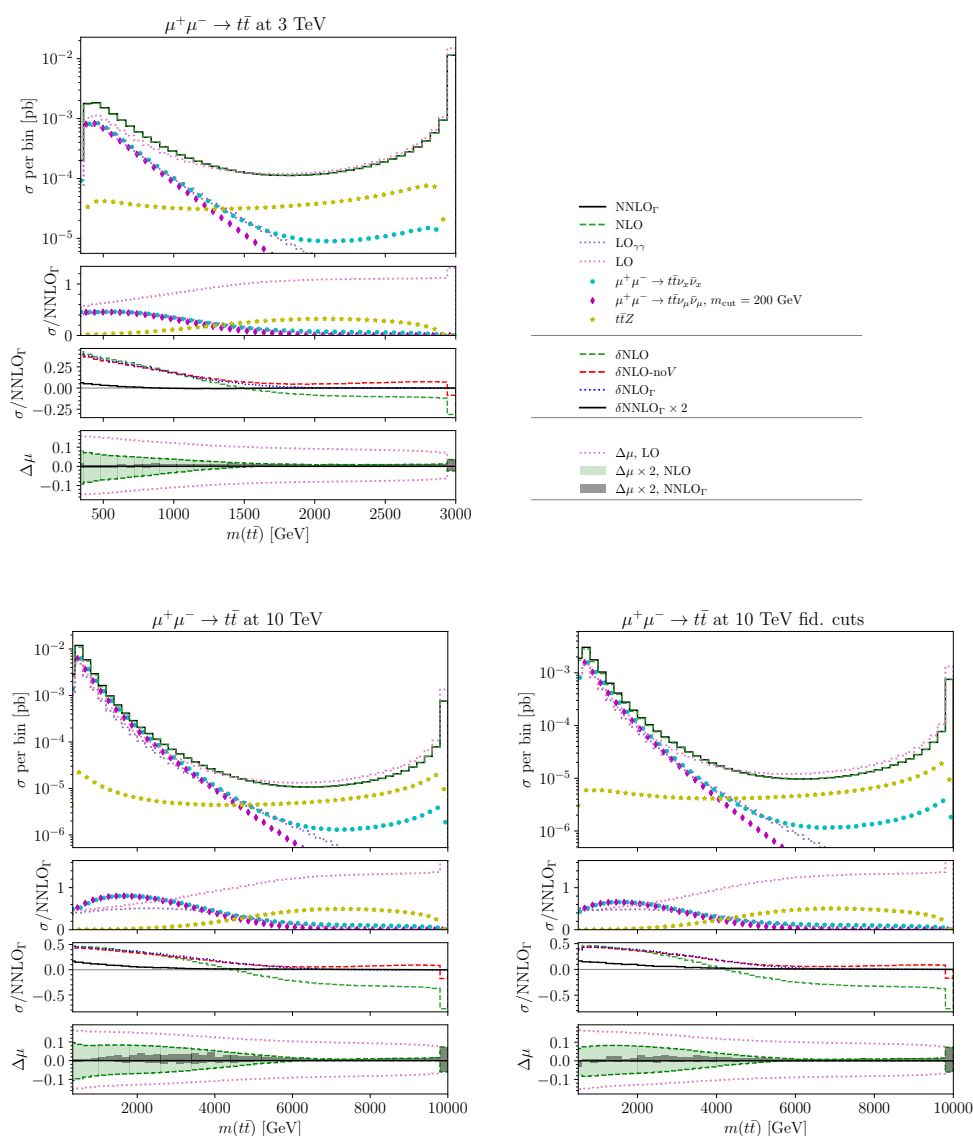


Figure 6. Invariant mass distribution of the top-antitop pair, $m(t\bar{t})$, in $t\bar{t}$ production at a muon collider with $\sqrt{s}=3, 10$ TeV. At 10 TeV, both the cases without and with acceptance cuts defined in eq. (4.4) are displayed. See the text for details.

it grows with the Bjorken x , and has an integrable singularity at $x \rightarrow 1$.¹⁹ Furthermore, the relative impact of the $\gamma\gamma$ -induced contribution grows with the c.m. energy, since smaller values of x become accessible, and at a 10-TeV collider $\text{LO}_{\gamma\gamma}$ essentially coincides with the total LO cross section from the threshold up to $m(t\bar{t}) = 2$ TeV, which in turn is about equal to 50% of the NNLO_Γ result (this is best seen in the top inset, where the agreement

¹⁹Without ISR effects, i.e. by truncating the muon PDF at $\mathcal{O}(\alpha^0)$, the muon PDF is equal to a Dirac delta, $\delta(1-x)$. Therefore, at the LO the entirety of the invariant-mass distribution would be contained in the rightmost bin. Note, however, that higher-order corrections, starting at the NLO, would contribute to the whole invariant-mass range even in the absence of PDF effects, owing to real radiation.

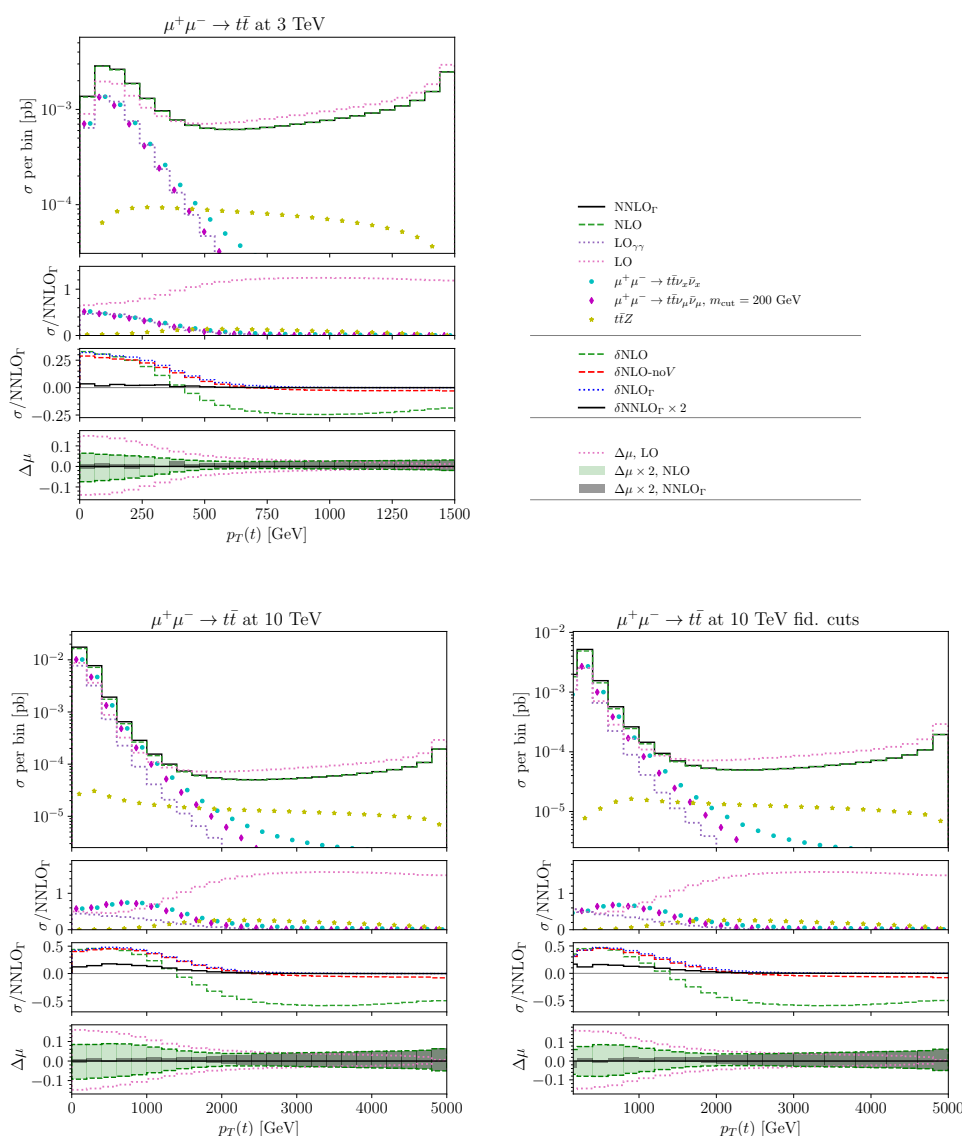


Figure 7. Same as in figure 6, for the top-quark transverse-momentum distribution, $p_T(t)$.

between the pink- and purple-dotted histograms improves with the c.m. energy, with their absolute values equal to about $1/2$). Before discussing the NNLO_Γ prediction in detail, it is appropriate to comment on the features of the associated production modes, $t\bar{t}\nu\bar{\nu}$ and $t\bar{t}Z$. Specifically, we see from the main frame and the top inset that the charged-current contribution is dominant at the threshold, i.e. in the same region as $\text{LO}_{\gamma\gamma}$, and it is of a comparable size (twice as large at its peak) as the latter at 3 TeV (10 TeV). It is worth observing that, even at a c.m. energy of 10 TeV, from $m(tt) = 5$ TeV onwards (i.e. in the region where the ratio of the top mass over the relevant kinematic quantity, the pair invariant mass, is less than 10%) the charged-current channel gives a contribution which is less than 10% of the NNLO_Γ differential cross section. Therefore, in the region where it gives a sizeable

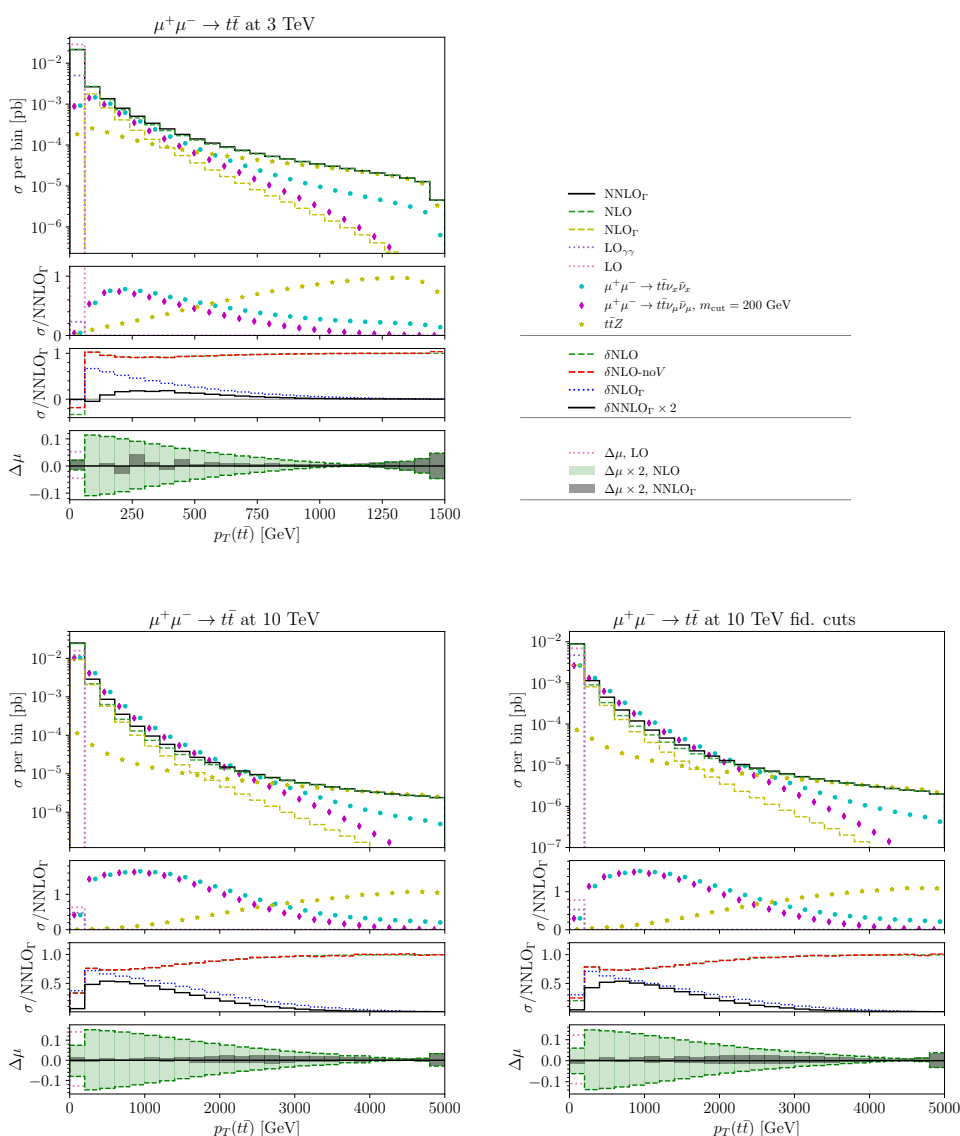


Figure 8. Same as in figure 6, for the top-antitop pair transverse-momentum distribution, $p_T(t\bar{t})$.

contribution (i.e. for $m(t\bar{t}) \lesssim 5$ TeV), Z -mass effects cannot be neglected ($m_Z/m(t\bar{t}) \gtrsim 2\%$). As far as the neutral-current channel is concerned, see footnote 9. Conversely, $t\bar{t}Z$ production induces a flatter spectrum in absolute value, and gives a relative contribution which at its peak is the largest among the associated-production channels. However, while this amounts to about 40% of the NNLO_Γ cross section at $m(t\bar{t}) \simeq 0.7\sqrt{s}$, there one finds oneself at the global minimum of the cross section.

As far as higher-order EW effects are concerned, they are unusually large: δNLO grows in absolute value as one approaches both the small- and large-invariant-mass regions (as is best seen from the green histogram in the middle inset), but this happens owing to two different mechanisms. Namely, at the high-end of the spectrum the growth is driven by the Sudakov negative enhancement of the virtual matrix elements, (this will be discussed in more detail in

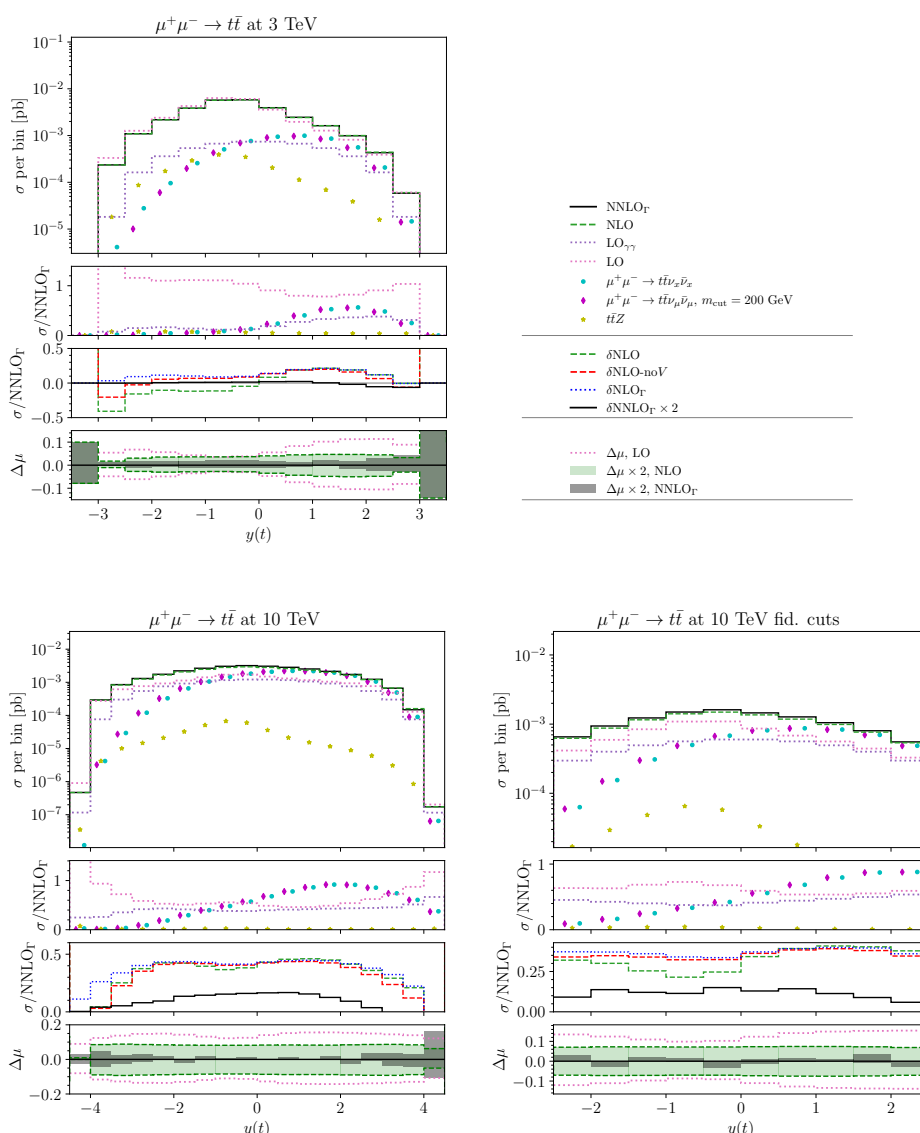


Figure 9. Same as in figure 6, for the top-quark rapidity distribution, $y(t)$.

section 4.2) which is a relative effect ranging from -10% at 3 TeV to -40% at 10 TeV (similar effects have been observed in ref. [37]). Moreover, the rightmost bin of the distribution, owing to its sensitivity to soft emissions and to the steepness of the spectrum, shows even larger effects, almost twice as large as those in the neighbouring bins. Conversely, at small invariant masses, a positive enhancement of the cross section is due to the $\mu^+\gamma$ and $\gamma\mu^-$ partonic channels which open up at the NLO, as real-emission corrections. Indeed, an indirect evidence of this fact is given by the excellent agreement among the δNLO , δNLO_Γ , and $\delta\text{NLO-noV}$ predictions (see the green, blue, and red histograms in the middle inset), which can only happen if one is dominated by γ -initiated real-emission processes. More in detail, since the muon PDF is peaked at $x \simeq 1$, the $\mu\gamma$ channel probes the photon PDF at smaller Bjorken- x

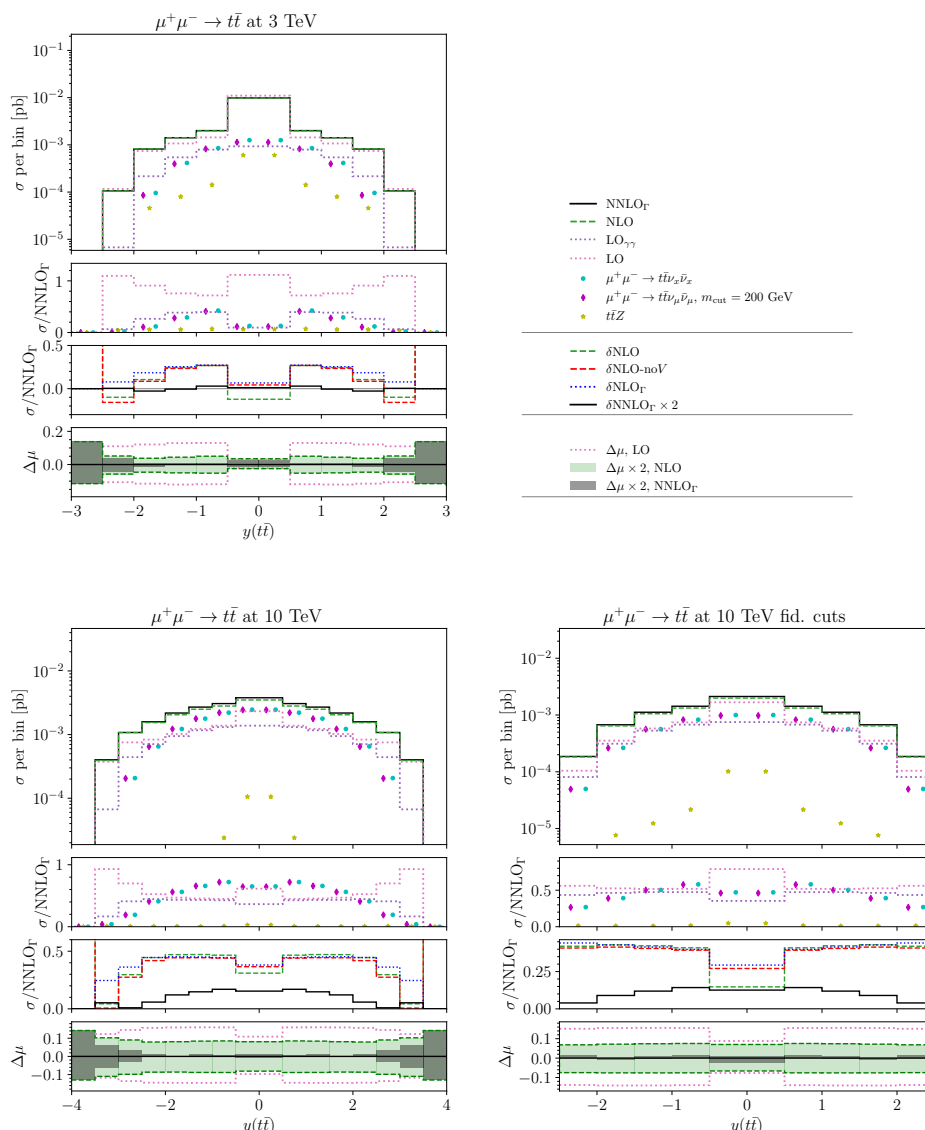


Figure 10. Same as in figure 6, for the top-antitop pair rapidity distribution, $y(t\bar{t})$.

values w.r.t. those relevant to the $\gamma\gamma$ one. Because of this, the NLO correction alone amounts to 40–50% (and grows with the collider energy) of the NNLO_Γ prediction. For what concerns δNNLO_Γ , at small invariant masses it has a similar shape as δNLO (compare the green and black histograms in the middle inset), and is smaller than the latter by roughly a factor of 5, while it vanishes by construction at large invariant masses. At the threshold, it increases the NLO cross section by $\sim 2\%$ and $\sim 10\%$ at 3 TeV and 10 TeV, respectively. These effects are unusually large for (EW) NNLO corrections, and this happens for the same reasons which we have discussed for NLO_Γ : in the case of NNLO_Γ , one predominantly probes both μ^+ and μ^- at large x 's; both of these radiate a photon or Z boson, which eventually fuse and produce the observed $t\bar{t}$ pair. As far as the factorisation-scale dependence of the cross

section is concerned, the bottom inset shows clearly how such a dependence is dramatically reduced when one includes higher-order corrections, from tens of percents at the LO (the band encompassed by the pink-dotted histograms) down to a few percents in NNLO_Γ (the grey band — bear in mind that such a band is multiplied by two in order to improve its visibility). Having said that, we stress that δNNLO_Γ constitutes an improvement of the NLO results only in the small-mass region, which explains why at large values of $m(t\bar{t})$ the NLO and NNLO_Γ bands coincide. We conclude by noting that acceptance cuts have a mild impact on this observable, and on the size of the various contributions.

The top-quark transverse momentum distribution, shown in figure 7, retains most of the features of the pair invariant mass, given that these two observables are closely correlated at the LO. For example, the p_T spectrum has also two peaks at or near the kinematics edges of its range, as was the case for $m(t\bar{t})$. At the LO, small values of p_T correspond to the region where $\text{LO}_{\gamma\gamma}$ is at its largest, while as the p_T grows, one is increasingly dominated by the muon-annihilation contribution — this is best seen from the top inset. For small values of p_T , the size and hierarchy of the various contributions largely follow the pattern observed for $m(t\bar{t})$. Conversely, the large- p_T region, starting from about $p_T = \sqrt{s}/4$, is mostly correlated with the rightmost bin of the $m(t\bar{t})$ distribution, since the large- p_T spectrum is much flatter w.r.t. its large- $m(t\bar{t})$ counterpart. An evidence of this fact can be obtained from the middle inset, where the NLO corrections are seen to reach their asymptotic values, dominated by Sudakov logarithms, already halfway through the kinematics range.²⁰

The transverse momentum of the top-antitop pair, shown in figure 8, is non-trivial starting only at the NLO, since at the LO its value is equal to zero (the top and the antitop are back-to-back in the transverse plane). In view of this, namely because of the outsize importance of real-emission contributions for this observable, in the main frame we also show the NLO_Γ prediction²¹ (yellow-dashed histogram). As an aside, we stress that since larger-than-zero $p_T(t\bar{t})$ values can only be obtained if there are final-state particles which the top pair can recoil against, it is impossible to obtain a non-trivial prediction for this observable in an approach based on EW PDFs (since that would be limited to a LO-accurate kinematic description). The $p_T(t\bar{t})$ spectrum strongly peaks towards small values of this observable, owing to the enhancement associated with the recoil of the $t\bar{t}$ pair against a soft/collinear parton. Indeed, as was discussed above, whenever $p_T(t\bar{t}) > 0$, NLO corrections are entirely due to real-emission processes, which diverge²² for $p_T(t\bar{t}) \rightarrow 0$. As is well known, predictions in this region could be further improved by resumming soft/collinear logarithms, by matching fixed-order results with either a parton shower or an analytical prediction. That said, the soft-photon-emission dominance in the small- $p_T(t\bar{t})$ region is relatively short-lived, as it can be seen by the steep decrease of the NLO_Γ result. Conversely, the relative size of the δNNLO_Γ correction is more significant for this observable w.r.t. to the other cases we have dealt with so far, and it reaches 30% at 10 TeV — this is in keeping with the fact that NNLO contributions

²⁰The large impact of the virtual contribution, which features a LO-like factorisation scale dependence, is such that for large p_T values the LO and NLO scale uncertainty bands are comparable in size.

²¹This is the NLO counterpart of the NNLO term of eq. (3.4), and it is basically the $\mu\gamma$ -initiated NLO cross section of real-emission origin — see ref. [17] for further details.

²²Any IR-sensitive theory such as QED or QCD cannot have infinite-resolution power; the cross section diverges with the logarithm of the bin size at $p_T(t\bar{t}) = 0$, but it is finite for any fixed non-zero bin size.

to all non-zero values for this observable are in fact the second (as opposed to the third) non-trivial ones. At small pair transverse momentum, the charged-current contribution is comparable to (at 3 TeV), or even larger than (by a factor of 2, at 10 TeV) the NNLO_Γ prediction. At the opposite end of the spectrum, up to the kinematic edge $p_T(t\bar{t}) = \sqrt{s}/2$, one finds a remarkably similar NNLO rate as for the $(t\bar{t}Z)$ -induced one. This is of course accidental, and merely tells one that associated-production channels must be accounted for for a sensible comparison with experimental data. Finally, the effective-LO nature of the non-zero $p_T(t\bar{t})$ bins is evident when looking at the factorisation scale dependence (bottom inset), where one sees a very significant reduction when NNLO_Γ corrections are included — clearly, this should and does happen only in regions dominated by initial-state photons, i.e. at small $p_T(t\bar{t})$; elsewhere, NLO and NNLO_Γ predictions behave similarly.

We now turn to longitudinal distributions, starting with the rapidity of the top quark, shown in figure 9. The most prominent feature of this observable is its asymmetry w.r.t. the reflection around zero, $y(t) \rightarrow -y(t)$, which is present in all of the results but LO_{γγ}. This is due to the fact that, in the $\mu^+\mu^-$ -annihilation channels, the interference between vector and axial-vector currents results in an asymmetric distribution, with the top quark emitted preferably in the same direction as the incoming μ^- , i.e. towards $y(t) < 0$ in our simulations; hence, the relative contribution of the $\mu^+\mu^-$ -initiated process grows with $-y(t)$. The asymmetry pattern is non-trivial, since one finds asymmetric histograms also in the top inset (which, we recall, shows contributions relative to the NNLO_Γ prediction) for *both* LO and LO_{γγ}; in the former case, the asymmetry of the numerator is not the same as that in the denominator, while in the latter case only the denominator is asymmetric. At 3 TeV, the LO spectrum is bound in the region $|y(t)| \lesssim 2.85$ by kinematics constraints; these are relaxed at higher orders, although marginally so, since the corresponding rates are very suppressed (by moving from the bin that includes the LO kinematical limit to its nearest neighbour with larger $|y(t)|$ values, one finds a cross section which is about four orders of magnitude smaller). At 10 TeV the LO kinematic limit becomes $|y(t)| \lesssim 4.06$ (still in the absence of acceptance cuts). In addition to this enlarged range, one notes that at 3 TeV the NNLO_Γ prediction is dominated by the $\mu^+\mu^-$ annihilation process (i.e. by large $m(t\bar{t})$ values), while at 10 TeV one finds predominantly $\gamma\gamma$ -initiated contributions (i.e. small $m(t\bar{t})$ values). Another interesting feature to notice, when looking at the NLO results, is that the effect of virtual corrections is largest for $y(t) < 0$, where it suppresses the cross section, a fact which is particularly visible either at 3 TeV or at 10 TeV with acceptance cuts (see the difference between the green- and red-dashed histograms in the middle inset). Again, this is explained by the dominance of the $\mu^+\mu^-$ annihilation channel in this region, which in turn is strictly correlated with large $m(t\bar{t})$ values; since virtual matrix elements are negative there, NLO corrections reduce the asymmetry. At 10 TeV the size of the NLO and the NNLO_Γ corrections is strikingly large (see the middle inset), and this reflects what happens for the top-pair invariant mass at low values of such an observable. Finally, for what concerns the associated-production channels, namely $t\bar{t}\nu\bar{\nu}$ and $t\bar{t}Z$, the key points are as follows. Firstly, they also underpin asymmetric spectra, with the asymmetry of $t\bar{t}Z$ ($t\bar{t}\nu\bar{\nu}$) being in the same (opposite) direction as that of $t\bar{t}$. Secondly, $t\bar{t}Z$ rates are quite suppressed for this distribution, while those for $t\bar{t}\nu\bar{\nu}$ production are chiefly due to the dominant charge-current channel. Lastly, the effect of factorisation-scale variations is quite flat, except at the LO, in which case the corresponding uncertainty is

smaller at $y(t) < 0$, where the competition between the $\gamma\gamma$ and $\mu^+\mu^-$ contributions leads (accidentally) to a compensating effect. NLO corrections reduce the LO uncertainty band by a factor of 2 (3) at 3 (10) TeV w.r.t. its minimum. The inclusion of δNNLO_Γ leads again to a reduction of the band width, by another factor of 2 (4).

The last observable we consider is the rapidity of the top pair, $y(t\bar{t})$, shown in figure 10. At variance with $y(t)$, this observable is symmetric under reflections, but nonetheless displays a rather peculiar shape, in particular at 3 TeV where the muon-annihilation channel has the largest relative impact. Indeed, the distribution shows a marked peak at $y(t\bar{t}) = 0$, which stems from the muon-annihilation channel at the LO, when both the μ^+ and μ^- PDFs are probed at large Bjorken x 's — since both x 's are very close to one, the resulting $y(t\bar{t})$ distribution (which is fairly strictly correlated with the logarithm of the ratio of the two x 's) exhibits a narrow peak, centred at zero. This peak structure is superimposed to the results of other channels, and is thus diluted by them — indeed, NLO and NNLO corrections have much flatter spectra, thanks to the fact that they feature extra final-state particles in addition to the $t\bar{t}$ pair. In particular, the inclusion of NLO corrections reduces the difference between the two central bins and the adjacent ones.²³ The LO prediction for this observable is bound in the range $|y| < -\log(2m_t/\sqrt{s})$, while beyond LO larger absolute values can be accessed, albeit with much suppressed rates, and an increased factorisation-scale dependence (as can be seen in the bottom inset). Similarly to the case of $y(t)$, $y(t\bar{t})$ is dominated by the large invariant-mass muon-induced peak at 3 TeV, and by the small invariant-mass photon-induced contribution at 10 TeV.

4.1.2 W^+W^- inclusive production

We now turn to discussing W^+W^- production. The observables which we consider are the strict analogues of those analysed in section 4.1.1, namely the invariant mass of the W^+W^- pair (figure 11), the transverse momentum of the W^+ (figure 12), the transverse momentum of the pair (figure 13), the rapidity of the W^+ (figure 14), and the rapidity of the pair (figure 15). Having commented at length on the case of $t\bar{t}$ production in section 4.1.1, here we mostly focus on the *differences* that one finds in W^+W^- production w.r.t. the former process. When looking at the distributions and comparing them to the case of $t\bar{t}$ production, it may be helpful to bear in mind some general features of W^+W^- production. Firstly, the relative impact of the photon-initiated contribution is larger than in $t\bar{t}$ production, both because of matrix-element enhancements and because the W boson, being lighter than the top quark, makes it possible to probe smaller Bjorken x 's. Secondly, while in $t\bar{t}$ production the two heavy quarks always belong to the same fermionic line, in turn connected to the other fermionic lines of the processes (the muons') by means of internal spin-1 propagators, this is not the case in W^+W^- production, where the bosons can also be directly radiated by the muon themselves. An immediate consequence of this is that W^+W^- production features a larger number of topologies w.r.t. $t\bar{t}$ production. For example, the muon-induced contribution at the LO for $t\bar{t}$ proceeds only via s -channel diagrams, while for W^+W^- also

²³The contributions that do not feature such an enhancement in the central region have a dip in the ratio over NNLO_Γ plotted in the top and middle insets. Note that such a dip is due to the denominator rather than to the numerator.

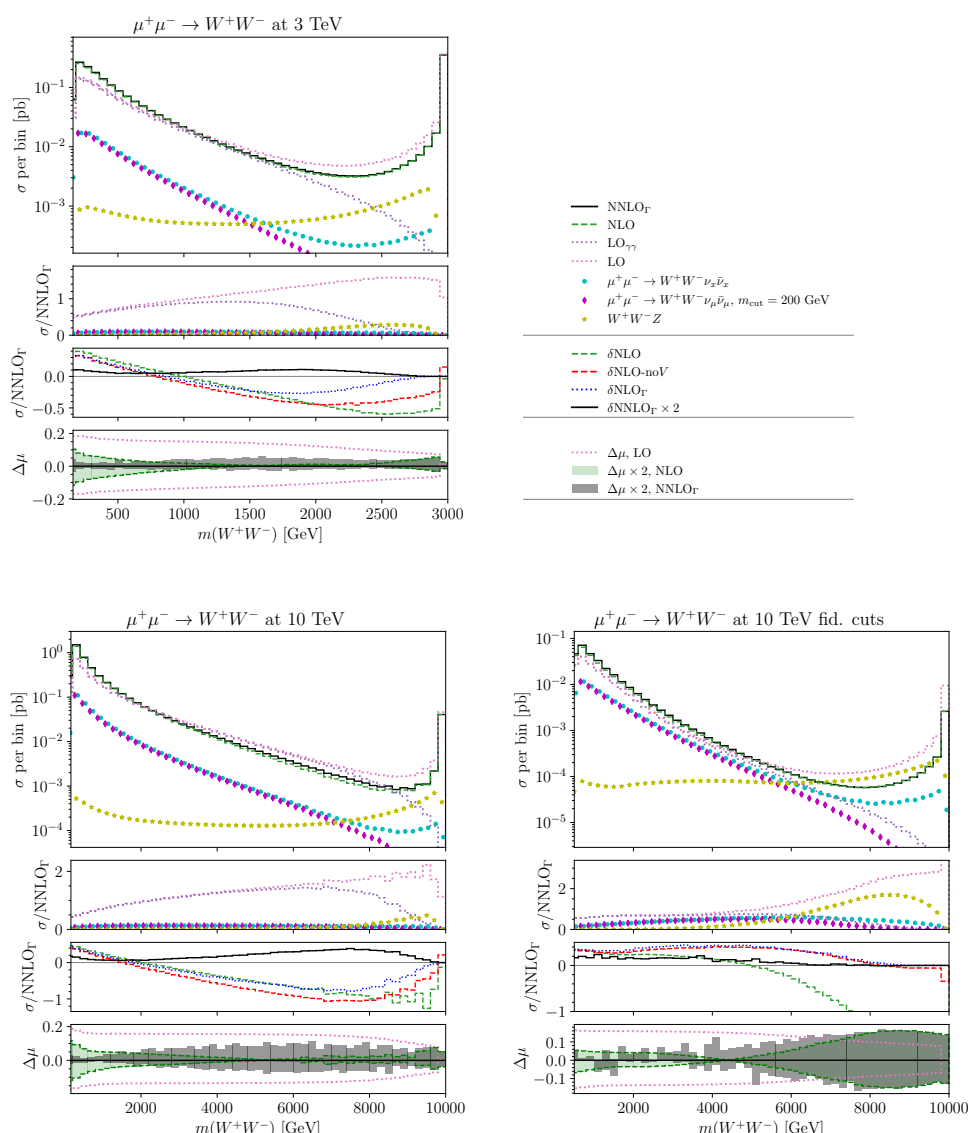


Figure 11. Same as in figure 6, for the W^+W^- pair invariant-mass distribution, $m(W^+W^-)$, in W^+W^- production.

t -channels exist. Thirdly, virtual corrections, in particular when large (w.r.t. the EW scale) energy scales are relevant, become huge in absolute value (almost as large as the LO results) and negative. This implies that, in order to have reliable predictions, effects beyond NLO must be included in the cross section, e.g. via resummation of EW Sudakov logarithms [22]. We shall return to this point in section 4.2.

The W -pair invariant mass, shown in figure 11, is a prototypical observable where the differences w.r.t. its $t\bar{t}$ -analogue, which we have just discussed, are apparent. Indeed, already at 3 TeV, the LO prediction is dominated by the photon-induced contribution up to $m(W^+W^-) \sim 2$ TeV (i.e., γ -initiated channels account for more than 50% of the cross section); at a 10-TeV collider, this dominance extends up to $m(W^+W^-) \sim 8$ TeV. The positive NLO

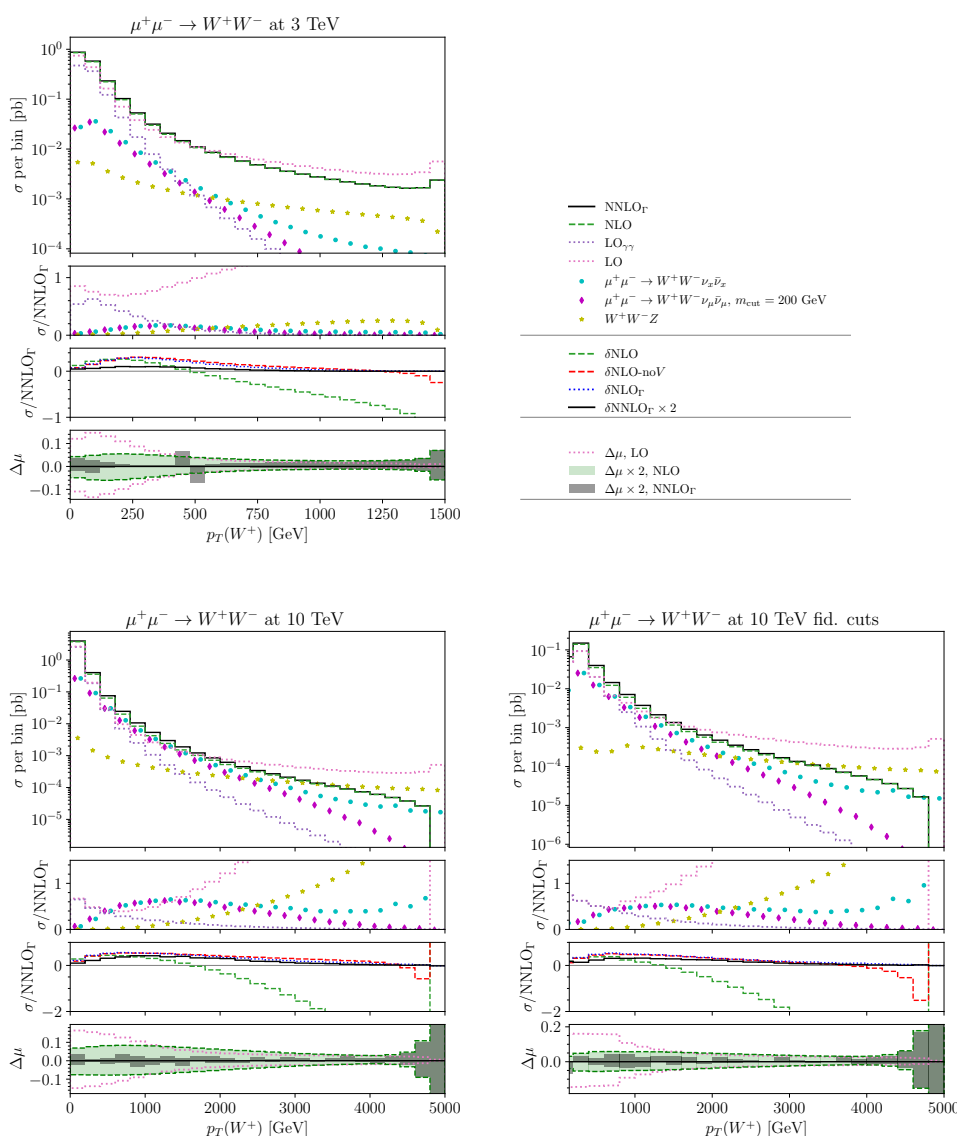


Figure 12. Same as in figure 11, for the W^+ -boson transverse-momentum distribution, $p_T(W^+)$.

corrections which appear at small $m(W^+W^-)$ and are chiefly due to the opening of the $\mu^+\gamma$ and $\gamma\mu^-$ channels, reach 40–50% of NNLO_Γ at the threshold (see the green-dashed histogram in the middle insets). At variance with $t\bar{t}$ production, their contributions become negative and large in absolute value as the invariant mass grows, until one reaches the region dominated by the $\mu^+\mu^-$ -annihilation channel, where instead the Sudakov suppression stemming from the virtual contribution is apparent. From the histograms in the middle insets, one can also see that the contribution of the virtual matrix elements is positive for the $\gamma\gamma$ channel. The impact of the δNNLO_Γ correction is also relatively large: it amounts to 5% of the total NNLO_Γ predictions at 3 TeV, and exceeds 20% at 10 TeV when $m(W^+W^-) \sim 8$ TeV. At both c.m. energies, as is expected, such a correction dies off together with $\text{LO}_{\gamma\gamma}$. Finally, the

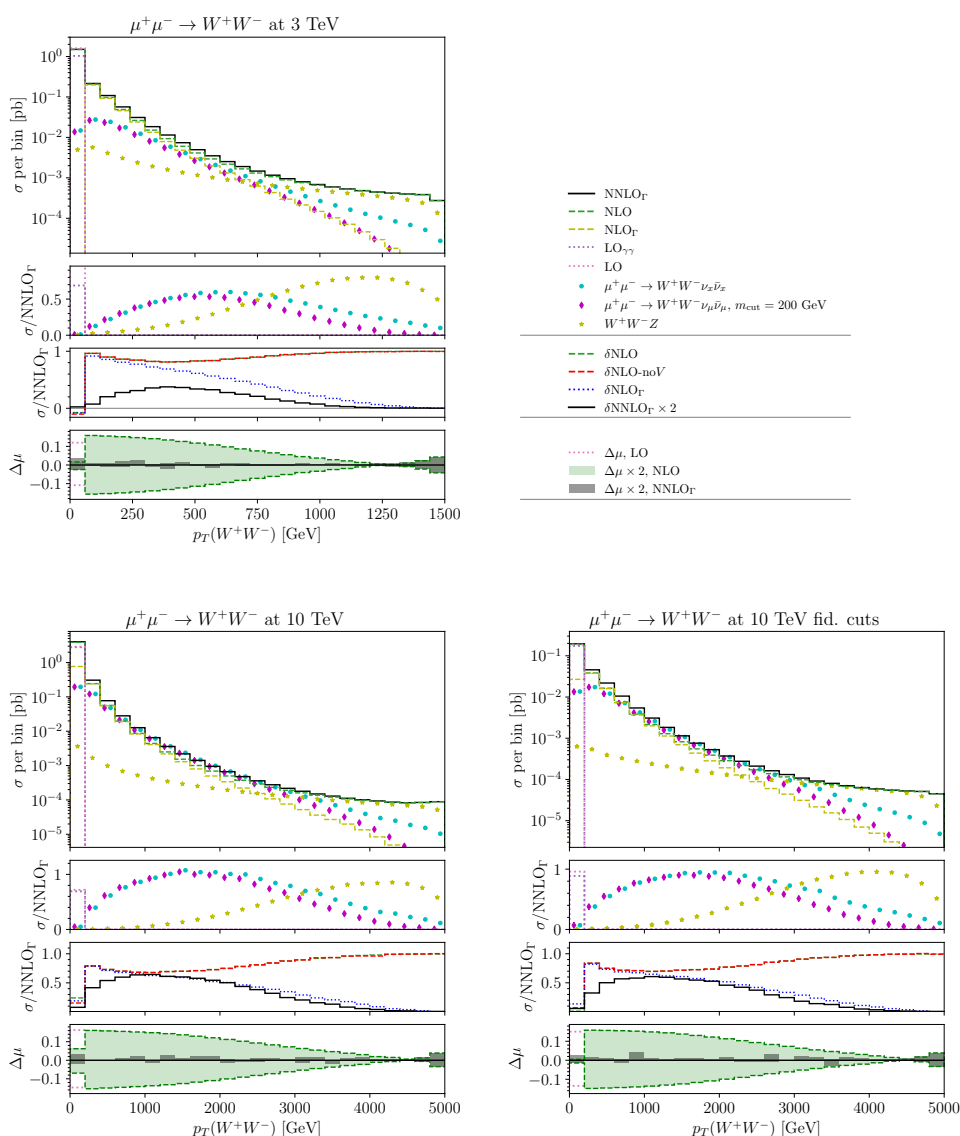


Figure 13. Same as in figure 11, for the W^+W^- pair transverse-momentum distribution, $p_T(W^+W^-)$.

impact of the associated production channels is rather negligible, with the only exception of W^+W^-Z production amounting to some 10 – 20% of NNLO_Γ at large invariant masses.

The presence of acceptance cuts at 10 TeV significantly changes the pattern we have just discussed. The main effect of such cuts is due to the ones in rapidity, as we shall see in detail when considering the W -rapidity distribution. In particular, the rapidity cut suppresses W -boson emissions from the μ^\pm lines in the forward/backward regions, and thus it reduces the rate across the whole $m(W^+W^-)$ range by more than one order of magnitude, with a stronger suppression of the small-mass $\gamma\gamma$ -induced peak than of the large-mass $\mu^+\mu^-$ -induced one. In this case, the dominance of the photon-initiated contribution fades off earlier, at around $m(W^+W^-) = 6$ TeV. The contribution of the virtual matrix elements is negative

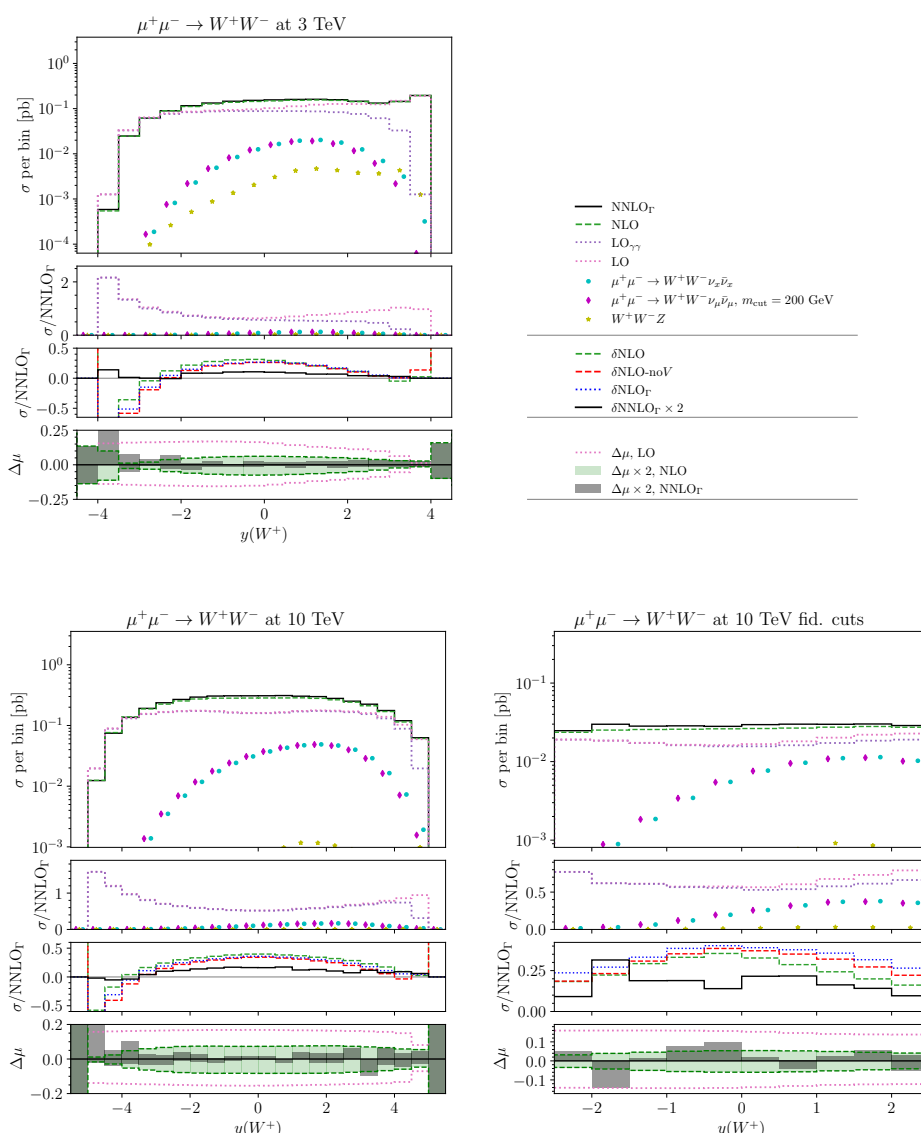


Figure 14. Same as in figure 11, for the W^+ -boson rapidity distribution, $y(W^+)$.

already at 2 TeV, and grows quickly in absolute value with the pair-invariant mass.²⁴ The acceptance cuts have a milder effect on W^+W^-Z associated production; therefore, this channel has a larger relative contribution w.r.t. the case where no cuts are imposed. Last but not least, it is counterintuitive (and at variance with all of the other cases discussed so far) that the NNLO_r factorisation-scale variation bands in the bottom insets are larger than their analogues at the NLO. However, we have explicitly verified that this is due to an accidental reduction (stemming from a cancellation) of the scale dependence at the NLO, in turn due to the large impact of virtual corrections (which have a LO-like scale dependence).

²⁴For better visibility, the range of the middle inset has been limited, and does not show the green histogram in its entirety. A more detailed discussion will follow in section 4.2.

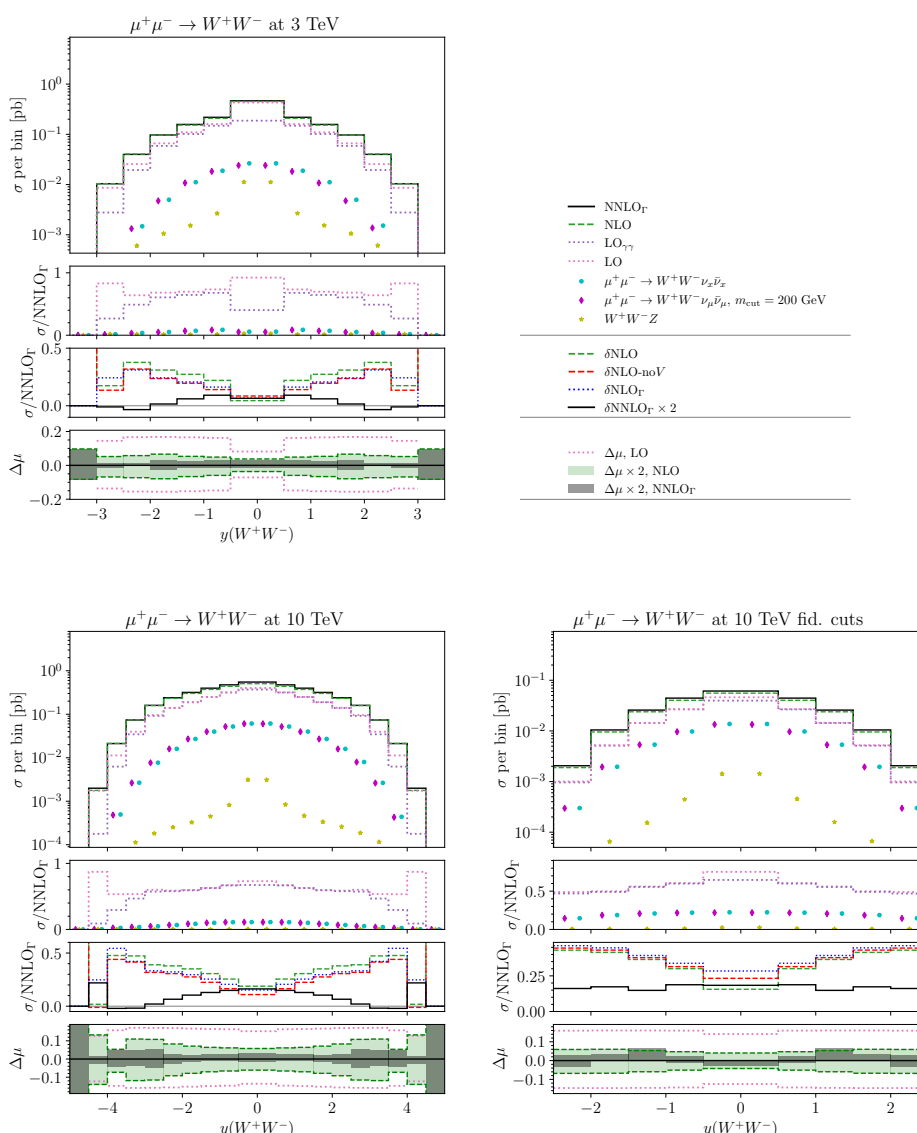


Figure 15. Same as in figure 11, for the W^+W^- pair rapidity distribution, $y(W^+W^-)$.

If the virtual matrix elements are not included in the predictions, the δNNLO_Γ corrections reduce the scale dependence as is expected.

The very large impact of virtual corrections at the NLO is apparent in the $p_T(W^+)$ distribution, shown in figure 12, for large values of the transverse momentum (see the difference between the green and red histograms in the middle panel). This observable displays a much steeper decrease w.r.t. its $t\bar{t}$ -production counterpart which we have discussed earlier, owing to the larger impact of the photon-induced contribution. While at the LO the predictions become less steep once moderate/large values of p_T are attained, the inclusion of NLO corrections changes this pattern, leading to steeper-than-LO results in this kinematic region. It is worth mentioning that NLO corrections grow so large in absolute value that the cross section in the rightmost bin becomes negative. Again, this is a sign of the need to resum the large

effects due to the EW Sudakov logarithms. In the small- p_T range, which is dominated by the photon-induced contribution, higher-order corrections are positive, similarly to what happens with the analogous observable in $t\bar{t}$ production. While at 3 TeV the impact of the associated production processes is small, at 10 TeV this is not the case. In particular, the charged-current channel gives a contribution which is comparable to the complete NNLO $_{\Gamma}$ prediction at around $p_T(W^+) = 1$ TeV; acceptance cuts only slightly reduce its impact. Conversely, W^+W^-Z production contributes mostly towards the large- p_T tail of the spectrum. However, one must take into account that this associated channel is simulated at the LO accuracy only, while NLO corrections are likely to reduce its impact significantly (see e.g. ref. [37]).

Turning to $p_T(W^+W^-)$, shown in figure 13, we remind the reader that, in full analogy with its $t\bar{t}$ counterpart, this observable is trivial in $2 \rightarrow 2$ kinematic configurations, such as at the LO or within an EW-PDF description. Indeed, all processes that do not feature any initial- or final-state radiation contribute to the $p_T(W^+W^-) = 0$ bin. NLO is thus the lowest order which populates the $p_T(W^+W^-) > 0$ region. Similarly to what happens for $p_T(t\bar{t})$, this translates into a very large impact of the δ NNLO $_{\Gamma}$ correction, whose contribution reaches 20% (35%) of the total NNLO $_{\Gamma}$ prediction at 3 TeV (10 TeV) — see the black histogram in the middle inset. The contributions of the associated-production processes follow a similar pattern as in the case of $p_T(W^+)$, with W^+W^-Z becoming comparable to NNLO $_{\Gamma}$ at large values of $p_T(W^+W^-)$ at both 3 TeV and 10 TeV. Conversely, at small and intermediate values of this observable it is the W -annihilation contribution which is comparable to NNLO $_{\Gamma}$. As in the case of $p_T(t\bar{t})$, the accuracy in the neighbourhood of the first bin can be further improved by resumming soft/collinear logarithms. Also, being the first non-trivial order, the NLO prediction has a LO-like (in size) factorisation scale dependence, while the NNLO $_{\Gamma}$ one significantly reduces it, at least for small values of this observable.

As far as rapidity distributions are concerned, we begin by looking at $y(W^+)$, shown in figure 14. Similarly to its $t\bar{t}$ counterpart, this observable is asymmetric, with a much more pronounced asymmetry w.r.t. to the former case. Indeed, because of the presence of the t -channel diagrams that feature $\mu^+\mu^- \rightarrow W^+W^-$ subgraphs, which would diverge at zero angle if the W were massless, the W^{\pm} boson is emitted preferably along the direction of the incoming μ^{\pm} . At a c.m. of 3 TeV, this translates into a strongly-enhanced rate for the W^+ boson at large, positive rapidity, with the distribution featuring a rather curious shape with two maxima. Such a feature is absent at 10 TeV, due to the dominance of the symmetric $\gamma\gamma$ channel. The associated-production processes display an asymmetry with the same sign as that of the NNLO $_{\Gamma}$ cross section. However, their rate is negligible, unless cuts are applied. In that case, the W -fusion channel is roughly 40% of the NNLO $_{\Gamma}$ prediction for $y(W^+) > 0$. The pattern of higher-order corrections is quite non-trivial: towards the negative edge of the LO spectrum (dominated by $\gamma\gamma$ fusion), NLO corrections start with being negative, then become positive around $y(W^+) \simeq -3$ ($y(W^+) \simeq -4$) at 3 (10) TeV. They reach the maximum (about 30% of NNLO $_{\Gamma}$ for either energy) around $y(W^+) = 0$ and start to decrease, approaching zero at the high-end of the distribution. Remarkably, and consistently with what has been observed for the $m(W^+W^-)$ distribution, there is an overall good agreement between NLO $_{\Gamma}$ and NLO without virtual matrix elements. The impact from the virtuals is the largest (and positive) for $y(W^+) < 0$, where the $\gamma\gamma$ initial state is dominant. The δ NNLO $_{\Gamma}$ correction is generally positive, more symmetric, and it grows as one approaches the central region, up

to 10% or 20% respectively at 3 and 10 TeV. At 10 TeV, when acceptance cuts are applied, a large part of the cross section is discarded, and one obtains a very flat distribution. In this case, as was mentioned when discussing $m(W^+W^-)$, cuts have a less severe impact on the $\mu^+\mu^-$ contribution than on the $\gamma\gamma$ one. This enhances the $\mu^+\mu^-$ -initiated fraction of the cross section at positive rapidity, for which NLO corrections receive a negative virtual contribution. Conversely, at negative rapidity, the LO cross section is entirely due to $\gamma\gamma$ fusion. This fact leads to an asymmetric impact of the NLO corrections.

Finally, we look at the rapidity of the W^+W^- pair, shown in figure 15. The kinematics bounds for this observable are $|y(W^+W^-)| < 2.9$ and $|y(W^+W^-)| < 4.1$ at 3 and 10 TeV, respectively. The central peak, due to the $\mu^+\mu^-$ -annihilation channel, is less apparent than for $t\bar{t}$ production, again consistently with the smaller impact of this channel on the cross section. Higher-order corrections are all positive: at the NLO, they grow when moving from the central towards the forward/backward regions, starting from 5%(20%) at 3 TeV (10 TeV), and growing up to 40% (50% at 10 TeV). We note that the δNNLO_Γ corrections display the opposite behaviour: they are very small in the forward/backward region, and grow up to 5% (10%) at $y(W^+W^-) = 0$. The impact of the associated-production processes is similar to the case of $y(W^+)$.

4.1.3 Summary of key phenomenological features

In general, our results for W^+W^- production reinforce the conclusions drawn from the $t\bar{t}$ analysis; at the same time, they also highlight a few differences due to some underlying physics characteristics, as follows.

- For both $t\bar{t}$ and W^+W^- production, the uncertainties associated with missing higher orders, which we have estimated by studying the behaviour of the perturbative series and its scale (and scheme, as we shall see later) dependence, appear to be under control.
- The number of subprocesses and topologies, in particular those which are not VBF-like, that lead to a W^+W^- final state is larger than that relevant to $t\bar{t}$ production. This results in larger corrections across all observables.
- The shape and scale of W^+W^- differential distributions are strongly affected by double-emissions from the muonic lines, which are absent in $t\bar{t}$ production.
- W^+W^- rapidity distributions show a stronger forward shift, consistently with the higher probability of independent boson emissions.

Taken together, the W^+W^- and $t\bar{t}$ production processes provide complementary tests of the robustness of the QED-resummed and fixed-order electroweak framework we have introduced. They both show that our approach is not only consistent with our stated goals, but also that it can reliably account for process-dependent behaviours, and provide a systematic way to evaluate uncertainties, which appear to be under control.

4.2 Inclusion of the Sudakov effects

In inclusive-production processes at multi-TeV muon colliders, the kinematic configuration of high energies *and* large scales is associated with a large fraction of events, at variance

with what happens in hadronic collisions — this behaviour is summarised in the sketch of figure 1, and is underscored by the explicit results of section 4.1. In such a configuration, EW corrections are dominated by large logarithmic terms of the form $\alpha^n \log^k(Q^2/M_W^2)$ with $1 \leq k \leq 2n$; at the NLO, in particular, there are double ($k = 2$) and single ($k = 1$) logarithms. By Q^2 we have denoted any kinematic invariant that can be constructed with a pair of external momenta, and which satisfies the condition $|Q^2| \gg M_W^2$ (in other words, Q^2 is equal to either s , t , or u in a $2 \rightarrow 2$ process, away from the collinear region). These logarithmic terms are commonly referred to as Sudakov logarithms [39], and stem from the calculation of the virtual contribution to EW corrections, where they emerge from mechanisms of both UV origin (associated with the running of the SM parameters) and IR origin (soft and/or collinear). If one would include their real-emission counterparts, due to the weak Heavy-Boson-Radiation (HBR),²⁵ there would be cancellations among these logarithms, but such cancellations would in any case be incomplete, and thus still lead to negative corrections which grow with energy [42–46].²⁶

Therefore, irrespective of whether one includes HBR contributions in inclusive predictions, these electroweak Sudakov logarithms (EWSL henceforth) become sizeable in the multi-TeV regime, and show up in particular in the large-scale regions of kinematic distributions such as invariant masses and transverse momenta, which are central to many new-physics searches. In fact, EWSL can be so large that they can lead to negative (thus unphysical) cross sections at any fixed order in perturbation theory, as is also discussed in ref. [37]. This is precisely the effect observed in the large- $p_T(W^+)$ tail of the 10 TeV result of figure 12. Similar issues, although not as dramatic, are also seen in the $m(W^+W^-)$ distribution, see figure 11.

Our framework, which combines QED-resummed collinear photon radiation with fixed-order NLO EW corrections and VBF-like contributions of NNLO, automatically includes the EWSL that are part of the NLO EW corrections, $\alpha^{b+n} \log^k(Q^2/M_W^2)$ with $n = 1$, and $k = 2$ and $k = 1$; b is the power of α associated with the Born (i.e. $b = 2$ for $t\bar{t}$ and W^+W^- production). In addition to that, these double and single logarithms of relative order α can also be calculated, within the MG5_AMC framework, by means of an automation and extension [49] of the Denner-Pozzorini algorithm [50, 51]. Such an algorithm is based on the factorisation of universal high-energy logarithmic corrections at the amplitude level in a process-independent fashion.

The advantage of having the possibility of computing in a standalone manner the EWSL contribution to NLO EW corrections is that it is not only possible to confirm that the origin of any large effects observed (such as that in figure 12, which we have already mentioned) is indeed due to EWSL, but also to estimate the impact of EWSL from higher orders in a fixed-order expansion. The study of ref. [37] does exactly that, for several processes and

²⁵While this inclusion is mandatory in massless theories such as QCD and QED, lest the cross section diverge, in the weak sector one *can* treat HBR as a separate process, owing to the fact that the heavy bosons are (in principle) taggable. Whether one *should* do it is another matter, which is besides the point here. More detailed discussions can be found in refs. [40, 41] and in ref. [37], for the LHC and the muon-collider cases, respectively.

²⁶The consequences of the incomplete cancellation have been exploited in ref. [47] in order to enhance the sensitivity to new short-distance physics via EW corrections; this is an intriguing strategy which has also been considered in ref. [48].

observables; however, in that paper, the focus was exclusively on the direct $\mu^+\mu^-$ -annihilation process; furthermore, a cut $m(W^+W^-) > 0.8\sqrt{s}$ was systematically imposed. By relaxing both of these conditions we carry out here an analogous study, in order to prove that the negative cross section observed at large $p_T(W^+)$ originates from EWSL. By doing so we also show that, *for this process*, effects due to the EWSL present in the NNLO cross section must be taken into account in order to achieve positive, and phenomenologically sensible, results. Conversely, if one aimed at precise (say, of relative 10^{-2} or smaller) predictions, the resummation of the whole EWSL tower would be necessary.

In order to unambiguously identify the EWSL contributions to the NLO EW results for phenomenological observables, we use the so-called SDK_{weak} scheme, which has been introduced in ref. [49], and is further discussed in detail in ref. [52]. This scheme accounts for the soft and/or collinear logarithms of only weak origin, while UV logarithms of both weak and QED origin are retained. The logic behind it is that, in physical observables, soft and collinear logarithms of QED origin completely cancel when the real radiation of photons is clustered (a.k.a. recombined) with any charged particle (massive or massless, thus including the top quark and the W boson) — this corresponds to defining dressed particles. This is the reason why in this section we also show predictions with photon recombination for charged heavy final states. We consider both the case of $t\bar{t}$ and W^+W^- production at 10 TeV, where the acceptance cuts are imposed, and we address the cases of the $p_T(t)$ and $m(t\bar{t})$ distributions, and of the $p_T(W^+)$ and $m(W^+W^-)$ ones, respectively.²⁷

The effects of the EWSL are illustrated in figure 16, where we also show the potential impact of their resummation. In the main panel, in addition to the results already shown in the bottom-right panels of figures 6, 7, 11, and 12, we present the NLO EW predictions with photon recombination, obtained by using a clustering radius $R = 1$,²⁸ supplemented by the approximate EWSL two-loop results. This is called $\text{NLO-reco}+\delta\text{SDK}_{\text{weak}}^{\text{NNLO}}$ (blue-dashed histograms), with:

$$\delta\text{SDK}_{\text{weak}}^{\text{NNLO}} \equiv \text{LO} \frac{1}{2!} \left(\frac{\delta\text{SDK}_{\text{weak}}}{\text{LO}} \right)^2 = \frac{\delta\text{SDK}_{\text{weak}}^2}{2\text{LO}}. \quad (4.6)$$

In other words, $\delta\text{SDK}_{\text{weak}}^{\text{NNLO}}$ originates from the Taylor expansion of the exponentiation of the corrections induced by EWSL at the NLO, that is:

$$\text{LO} \times \exp \left(\frac{\delta\text{SDK}_{\text{weak}}}{\text{LO}} \right) = \text{LO} + \delta\text{SDK}_{\text{weak}} + \delta\text{SDK}_{\text{weak}}^{\text{NNLO}} + \dots \quad (4.7)$$

In particular, $\delta\text{SDK}_{\text{weak}}$ is the EWSL component of δNLO computed in the SDK_{weak} scheme, and $\delta\text{SDK}_{\text{weak}}^{\text{NNLO}}$ approximates the effects of EWSL at the NNLO. The latter takes into account in an exact manner the dominant EWSL effects at high energy in the NNLO EW corrections, namely those of order $\alpha^{b+2} \log^4(Q^2/M_W^2)$, and estimates by means of a naive exponentiation those of order $\alpha^{b+2} \log^n(Q^2/M_W^2)$, with $n = 3, 2$.²⁹ In the inset of figure 16

²⁷Without any cuts only the p_T distributions can be studied in this context, since for large invariant masses transverse momenta could be small and therefore $|t| \sim M_W^2$, thereby invalidating the Sudakov approximation.

²⁸Any sensible value of R may be employed. Here we choose $R = 1$ in order to clearly isolate the effects of EWSL, since in this way no $\log R$ term affects the cross section.

²⁹This prediction coincides with what is called EXP_{EW} in ref. [37], truncated at relative order α^2 w.r.t. the LO. For a proper treatment of resummation effects, one should consider studies as, e.g., that of ref. [22].

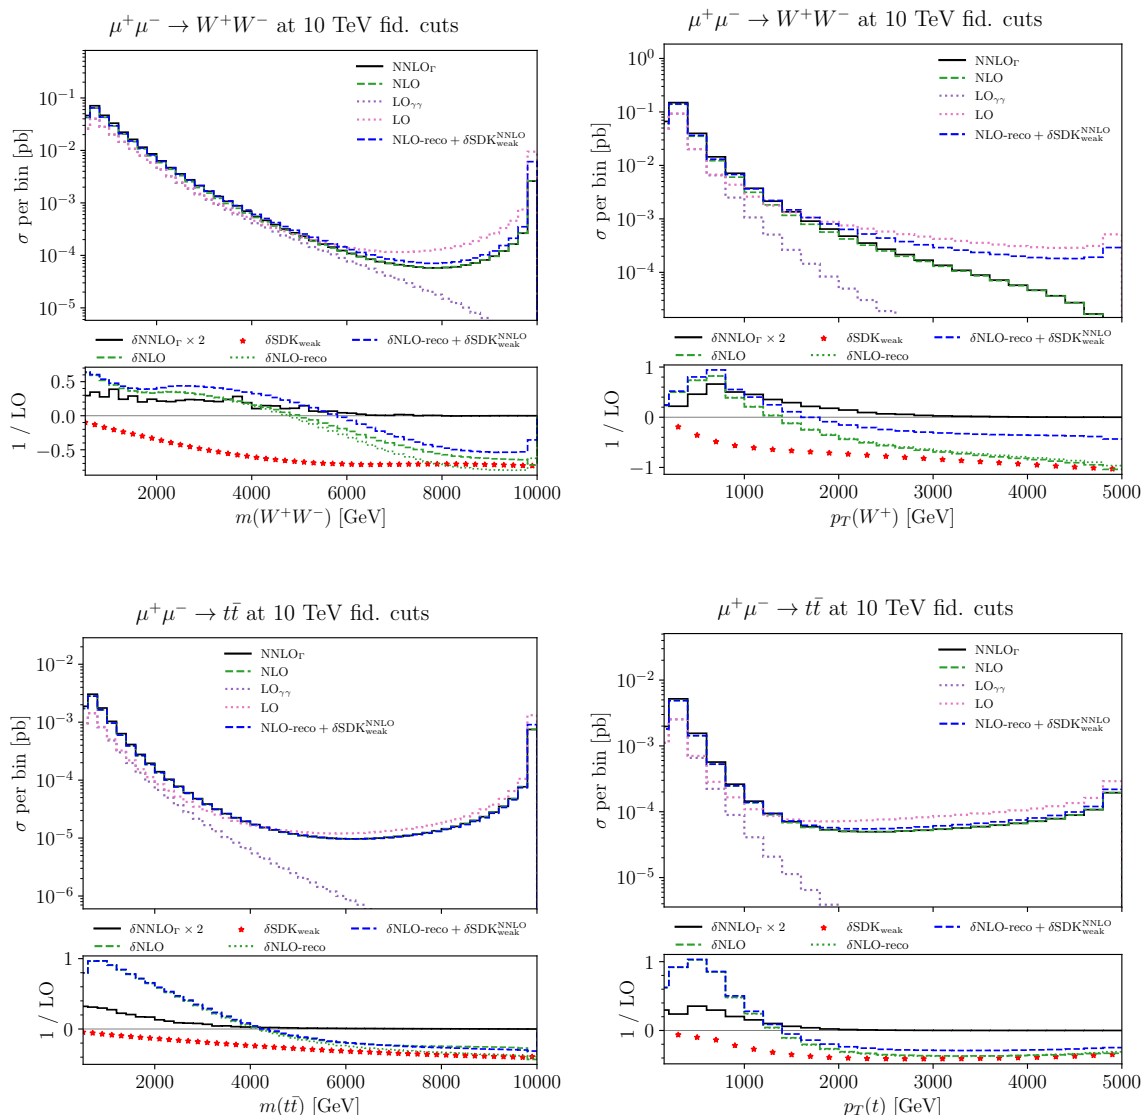


Figure 16. Study of the effects coming from the presence of large Sudakov logarithms in representative distributions for $t\bar{t}$ and W^+W^- final states. See text for details.

some of the same results as in the main panel are shown as relative with respect to the LO.³⁰ In addition to those, we also present the $\delta\text{NLO-reco}$ (i.e. the NLO corrections with photon recombination, green dots) and the $\delta\text{SDK}_{\text{weak}}$ (red stars) predictions.

We start by discussing the case of W^+W^- production, where the EWSL effects are larger, and later briefly move on to the case of $t\bar{t}$ production. For large values of $m(W^+W^-)$, and especially of $p_T(W^+)$, NLO EW corrections are very large in absolute value and negative, as was already noted. Here we also see that, by recombining photons, the NLO EW corrections

³⁰Note that this is different w.r.t. what is done in section 4.1, where the ratios over NNLO_Γ predictions are employed. This is so in view of the fact that potentially large cancellations occurring beyond LO would create artefacts in the ratios, and thus obscure the conclusions here.

very slightly decrease in the tail of the $p_T(W^+)$ distribution. This phase-space region is mostly correlated with the rightmost bin of $m(W^+W^-)$, which also displays the same behaviour. The key message to be taken out of that figure is that the EWSL from δSDK capture a large part of the $\delta\text{NLO-reco}$ result, and that the pathological behaviour observed at fixed order is cured by adding the contribution of $\delta\text{SDK}_{\text{weak}}^{\text{NNLO}}$.

Below $p_T(W^+) \simeq 2$ TeV, we notice how the difference between $\delta\text{NLO-reco}$ and $\delta\text{SDK}_{\text{weak}}$ is not a constant, a fact that signals the presence of other kinematic-dependent and non-negligible effects unrelated to the EWSL. These effects are indeed those already discussed in section 4.1, stemming from the contributions of the $\mu^+\gamma^-$ and $\gamma\mu^-$ -initiated processes, which have a significant impact on NNLO_Γ . The same observation applies to the $m(W^+W^-)$ distribution for $m(W^+W^-) \lesssim 8$ TeV.

In summary, our analysis shows that in W^+W^- production we understand the origin of the pathological behaviours of the $m(W^+W^-)$ and $p_T(W^+)$ distributions at high energy. The implication is that at least the NNLO EWSL contributions (i.e. not only those present in NNLO_Γ), in addition to the full NLO EW corrections, must be taken into account in order to obtain sensible predictions (i.e. positive-defined cross sections) in the presence of corrections as large as -50% w.r.t. the LO. Having said that, further effects of $\mathcal{O}(10\%)$ can still appear due to EWSL of higher orders (e.g., $\alpha^{b+3} \log^6(Q^2/M_W^2)$),³¹ and therefore the resummation of the full EWSL tower would be necessary for precision physics, which in turn would require calculations such as the one of ref. [22]. However, it is important to bear in mind that this kind of conclusions are both observable- and process-dependent. Indeed, by simply looking at the analogous plots for $t\bar{t}$ production (in turn compatible with the results presented in refs. [37, 48]), one ends up with a different overarching message. Namely, that the EWSL NNLO contributions are not necessary for obtaining positive-definite, and generally sensible, predictions. Moreover, they are sufficient for reducing the uncertainty due to higher-order EWSL at the percent level.³² Thus, in $t\bar{t}$ production the resummation of EWSL appears not to be necessary even if one focuses on precision physics at the percent level.

In conclusion, the impact of EWSL must be carefully assessed process-by-process and even observable-by-observable, and it is therefore important to have tools that allow one to do that in a flexible way.

4.3 Factorisation scheme dependence

A potentially important source of theoretical uncertainty in predictions that combine QED resummation with fixed-order electroweak corrections is the choice of the factorisation scheme used to define lepton PDFs. In this section we explicitly show how such a choice (exemplified by the differences between the results obtained with the $\overline{\text{MS}}$ [53] and the Δ [35] schemes) has a huge impact on LO predictions (e.g., inducing $\mathcal{O}(50\%)$ systematics at the threshold

³¹One can easily see that when the impact of EWSL at the NLO is of order -100% , at the NNLO one is liable to see order $(-100\%)^2/2! \simeq +50\%$ effects, which may become of order $(-100\%)^3/3! \simeq -15\%$ at the N^3LO , of order $(-100\%)^2/4! \simeq +2\%$ at the N^4LO , so forth.

³²By employing the same argument as in footnote 31, one can easily see that since EWSL effects at the NLO are at most -40% in this case, those at the NNLO are of order $(-40\%)^2/2! \simeq +8\%$, and therefore important only in order to achieve percent-level accuracy. However, at the N^3LO they are expected to be of order $(-4\%)^3/3! \simeq -1\%$.

of the $m(W^+W^-)$ distribution in 10-TeV muon collisions). Fortunately, we also show how the factorisation-scheme dependence is reduced to the percent level when both NLO EW corrections and the NNLO _{Γ} contribution are taken into account — such a dramatic reduction is particularly effective in the threshold region, and more in general where one is dominated by $\gamma\gamma$ -induced processes. In the following, we briefly discuss some of the technical details underpinning the scheme dependence, and show the improvements that can be obtained by including higher orders.

While the $\overline{\text{MS}}$ scheme is generally assumed to be the default choice in massless theories such as QCD and QED, and indeed it is nowadays universally adopted in hadron-collision simulations, the possibility of calculating analytically lepton PDFs in QED allows one to show that double soft logarithms appear beyond the LO in both the PDF initial conditions and the short-distance cross sections, only to eventually cancel at the level of physical observables. At the NLO, the initial conditions can be written as follows [18]:

$$\Gamma_{\mu/\mu}^{[1]}(z, \mu_0) = \left[\frac{1+z^2}{1-z} \left(\log \frac{\mu_0^2}{m^2} - 2 \log(1-z) - 1 \right) \right]_+ + K_{\mu\mu}(z), \quad (4.8)$$

$$\Gamma_{\gamma/\mu}^{[1]}(z, \mu_0) = \frac{1+(1-z)^2}{z} \left(\log \frac{\mu_0^2}{m^2} - 2 \log z - 1 \right) + K_{\gamma\mu}(z), \quad (4.9)$$

where the functions $K_{a\mu}(z)$ encode all information about the factorisation scheme. Conventionally, in $\overline{\text{MS}}$ one has $K_{\mu\mu}^{\overline{\text{MS}}}(z) = K_{\gamma\mu}^{\overline{\text{MS}}}(z) = 0$, whereby the double logarithmic term in eq. (4.8) is manifest (stemming from the second term within round brackets). The scheme independence of observables becomes intuitively clear when bearing in mind that contributions which feature the K functions must be included (importantly, with a minus sign) in the short-distance cross sections. In the FKS formalism [54–56], this is done in the so-called $(n+1)$ -body degenerate contributions. A residual dependence on the factorisation scheme remains after the convolution with PDFs, but such a dependence is beyond the perturbative order at which the cross section is computed. In this section, we shall present a few explicit examples of factorisation-scheme (in)dependence — the interested reader can find further results in ref. [17].

The Δ scheme [35] has been devised to eliminate the double logarithms from the beginning, by defining:

$$K_{\mu\mu}^{\Delta}(z) = \left[\frac{1+z^2}{1-z} (2 \log(1-z) + 1) \right]_+, \quad (4.10)$$

$$K_{\mu\mu}^{\Delta}(z) = \frac{1+(1-z)^2}{z} (2 \log z + 1), \quad (4.11)$$

so that in the lepton-PDF initial condition the only surviving term is a single soft logarithm, with coefficient proportional to the logarithm of the scale (μ_0) where the initial conditions are imposed, and which thus vanishes with the standard choice $\mu_0 = m$.³³

A very convenient by-product of the usage of the Δ scheme is that observables computed in such a scheme behave much better from the numerical point of view. For example, the muon

³³It has to be stressed that, while the Δ scheme is unambiguously defined in the large- z region, more freedom exists at small z , so that it can effectively be considered as a *class* of different schemes [21].

density in $\overline{\text{MS}}$ can become negative as $z \rightarrow 1$ (this happens for $1-z = \mathcal{O}(10^{-20}) - \mathcal{O}(10^{-30})$)³⁴, while it remains positive in the Δ scheme. This results in significantly faster integration times. To give just one example, the accumulated running time for the integration of W^+W^- production at the NLO with MG5_AMC, at 3 TeV and with a target relative accuracy of 10^{-4} on the total integral is 190 days in $\overline{\text{MS}}$, and 35 days in Δ . Owing to the smaller contribution of the muon-initiated channels, the gain is reduced down to “just” a factor 2 at 10 TeV.

Turning now to the presentation of the results, we comment on our predictions at a c.m. energy of 10 TeV, since there the scheme-dependence effects are larger. Again, we shall consider $t\bar{t}$ and W^+W^- production, in figure 17 and figure 18, respectively, and look at two representative observables, the invariant mass of the pair and the transverse momentum of the positively-charged heavy particle. We shall consider both the case with and that without the acceptance cuts of eq. (4.4). The two figures have the same layout, and show in the top (bottom) row the invariant-mass (transverse momentum) distribution, without (left panels) and with (right panels) cuts. Each plot has a main frame and two insets. The main frame shows absolute rates in the Δ (solid) and $\overline{\text{MS}}$ scheme (dashed), at the LO (blue), NLO (red), and NNLO $_{\Gamma}$ (black) accuracy. The two insets display the same quantity, namely the relative (w.r.t. the Δ -scheme result) difference between predictions in the two schemes, computed at the different accuracies already relevant to the main frame; the bottom inset has a narrower y range w.r.t. to the top one. The colours employed in the insets are the same as those of the corresponding predictions in the main frame, but the patterns may be different in order to improve visibility. In addition to that, in the bottom inset we also show the NLO $_{\Gamma}$ (cyan diamonds) and NLO-noV (yellow stars) relative differences.

We start with $t\bar{t}$ production, shown in figure 17. In the invariant-mass distribution, the scheme dependence reduces dramatically as the accuracy increases. In particular, at the LO, the relative scheme dependence is as large as 50% at small values of $m(t\bar{t})$, while at large values of $m(t\bar{t})$ such a dependence monotonically grows in absolute value up to 10%, only to be essentially equal to zero in the rightmost bin of the distribution. This behaviour, although with asymptotic values slightly smaller than 10%, appears to be a feature common to all considered accuracies. The negligible scheme dependence of the rightmost bin is due to the change of sign of the large- z muon PDF in $\overline{\text{MS}}$, which leads to an accidental cancellation, also driven by the bin size. With the exception of the rightmost bin, the invariant-mass distribution behaves as one would expect: NLO corrections reduce the scheme dependence down to 10% at small values of $m(t\bar{t})$, where the contribution of the virtual matrix elements is small. On the other hand, since the scheme dependence due to the virtual matrix elements at the NLO is LO-like, but (in this particular case) opposite in sign, in the regions where their contribution is sizeable such a reduction does not take place; specifically, this happens at the high-end of the spectrum. This is demonstrated by the results obtained with NLO corrections that do *not* include the virtual matrix elements (see the yellow stars in the bottom inset), which show the expected reduction of scheme dependence. The NLO $_{\Gamma}$ predictions are very close to the NLO ones at small $m(t\bar{t})$, where photon-initiated channels are dominant, and to the LO ones at large $m(t\bar{t})$, since $\delta\text{NLO}_{\Gamma} = 0$ in that region. Finally, the NNLO $_{\Gamma}$ prediction reduces the scheme dependence down to 1% or below at small $m(t\bar{t})$ values, while it behaves similarly to

³⁴See e.g. eq. (5.63) in ref. [34] for the asymptotic behaviour in the $\overline{\text{MS}}$ scheme.

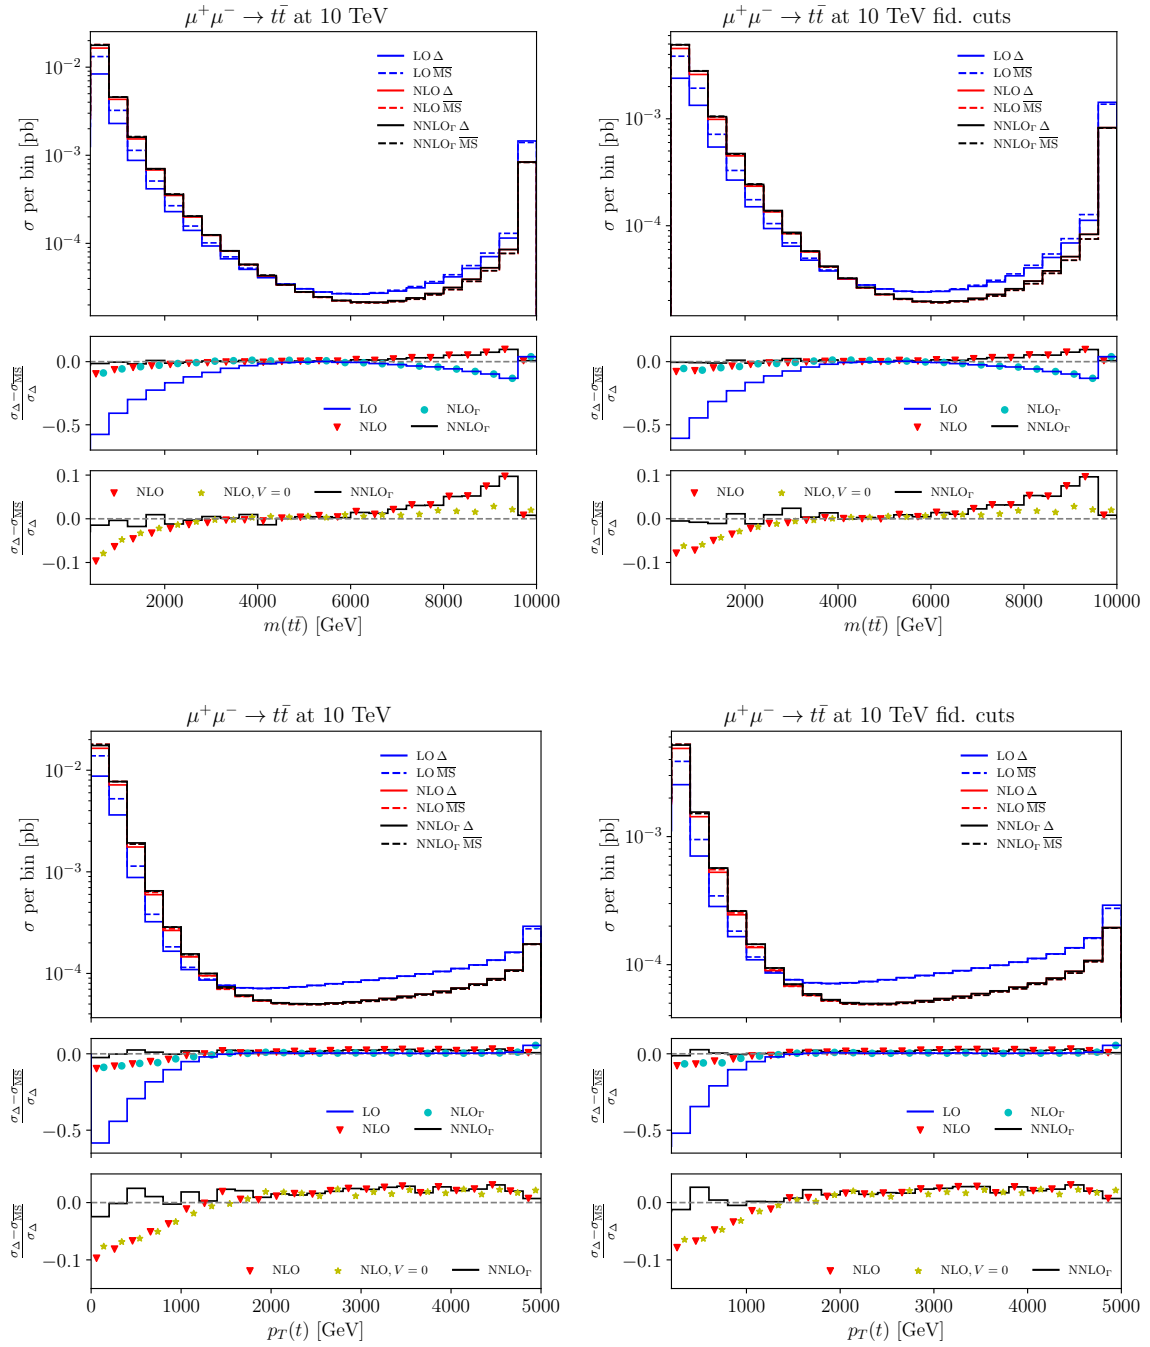


Figure 17. Factorisation-scheme dependence of the invariant-mass and transverse-momentum distributions in $t\bar{t}$ production. See the text for details.

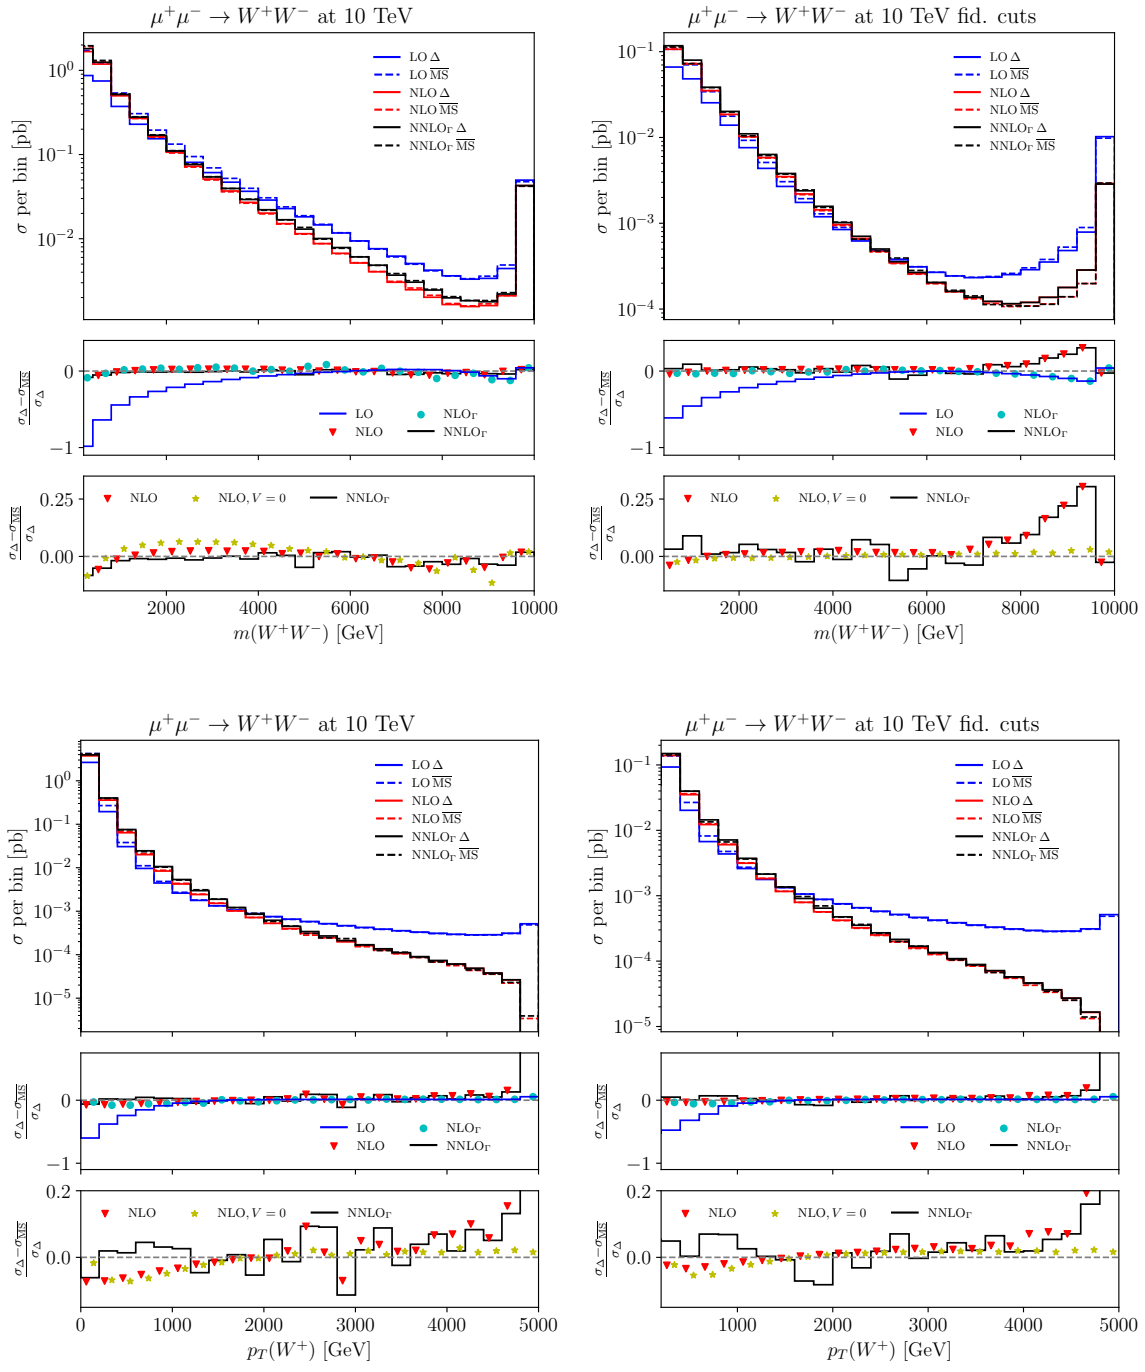


Figure 18. Same as in figure 17 for W^+W^- production.

the NLO one at large invariant masses, where photon-induced contributions are unimportant. We point out that in the latter region, the very large virtual matrix elements are primarily driven by the EWSL, which we have discussed in section 4.2. Thus, they are chiefly of weak origin, whilst the factorisation-scheme dependence pertains only the purely QED sector of EW corrections. This implies that in order to reduce such a dependence at large invariant masses the full NNLO EW results, or at least its mixed weak-QED part, would be necessary.

The $p_T(t)$ distribution, as was already said, correlates strongly with the low- $m(t\bar{t})$ region at small $p_T(t)$, i.e. in the region dominated by $\gamma\gamma$ fusion, and with the rightmost bin of the $m(t\bar{t})$ distribution at moderate/large $p_T(t)$. This is why, while at the low-end of the p_T spectrum one sees very similar features as the corresponding region in the pair invariant mass, at the spectrum high-end all scheme dependences are very small (as for the rightmost bin of $m(t\bar{t})$).

Finally, acceptance cuts have a negligible effect on the scheme-dependence patterns.

Turning to W^+W^- production, shown in figure 18, we observe some striking differences with respect to the $t\bar{t}$ case just considered. These are due to the richer phenomenology of this process (in terms of production modes). The only common feature with $t\bar{t}$ production is the large scheme dependence at the LO for small $m(W^+W^-)$. Regardless of the presence of the acceptance cuts, at small values of the invariant mass all predictions beyond LO display a similar scheme dependence, at the few-percent level. This can be ascribed to different effects, from the presence of sizeable, LO-like contributions in the δNNLO_Γ corrections (e.g. ZZ fusion), to the non-trivial interplay between virtual corrections and e.g. $\mu^+\gamma$ real-emission channels at the NLO. For the case without cuts, this remains true up to the high-end of the spectrum, where forward/backward W emissions play an important role. Acceptance cuts remove these contributions, leading to a behaviour in this region similar to the one observed for $t\bar{t}$, with the large NLO scheme dependence exacerbated by the much larger virtual matrix elements.

As is the case in $t\bar{t}$ production, the p_T distribution display features which correlate quite well with those of $m(W^+W^-)$. It has to be noted that at the upper end of this distribution, the virtual contribution is so large that the cross section becomes negative in the rightmost bin.

In conclusion, from a practical standpoint the Δ scheme appears preferable to the $\overline{\text{MS}}$ one for high-energy lepton colliders. It provides one with a closer correspondence with semi-classical photon emission, and is numerically much more stable w.r.t. the latter. Nonetheless, the residual scheme dependence remains a useful diagnostic tool of the robustness of the calculation, and can be exploited to estimate a part of the theoretical uncertainty associated with the truncation of the perturbative series.

Phenomenologically, our results confirm that while the overall scheme dependence is moderate, it is non-negligible in precision applications, and must be included in the uncertainty budget of future-collider analyses.³⁵ The ability of our framework to consistently account for this dependence, and to switch between schemes transparently, is an important asset in view of its deployment in numerical simulations.

³⁵The reader must pay attention to the fact that EW-PDF-based simulations do not quote such an uncertainty not because it is equal to zero, but because it is ill-defined at the LO accuracy relevant to such simulations unless one follows the prescription of ref. [19]; this is not done currently, and therefore this uncertainty is literally out-of-control there.

5 Summary and conclusions

In this work, we have addressed the challenge of providing accurate and systematically-improvable predictions for inclusive electroweak processes at future high-energy lepton colliders, with a particular focus on multi-TeV muon colliders. We have done this by employing a new approach which we have proposed and implemented in a companion paper [17], that combines QED-resummed collinear radiation with fixed-order EW corrections, thereby overcoming many of the conceptual and practical limitations of the conventional EW-PDF-based approach.

Our strategy consists in improving NLO EW results, obtained in the context of the collinear factorisation theorem and therefore nowadays automated, by adding double-real contributions that embed, but are not limited to, the leading NNLO VBF-like topologies, in a way that preserves gauge invariance, full phase-space coverage, and includes both soft and collinear-enhanced diagrams featuring γ/Z exchange in the t channels, which in the case of the photon are resummed via lepton PDFs up to NLL accuracy. The weak component of these diagrams is treated exactly at fixed order, thus ensuring the inclusion of mass-suppressed and interference effects that are lost in the traditional VBF approximation. Importantly our framework, by avoiding any double counting with NLO EW corrections, opens up a clear path toward future refinements.

We have demonstrated the power and flexibility of this approach by considering two representative processes, namely $t\bar{t}$ and W^+W^- inclusive production at $\sqrt{s} = 3$ and 10 TeV muon colliders. For both processes, we have shown that the improvement we have proposed has a significant quantitative impact across a wide range of kinematic variables, particularly in regions dominated by forward scattering or large invariant masses. In parallel, we have studied the same observables with the EWA, compared the predictions thus obtained with those of the full matrix elements which underpin our procedure, and assessed the differences w.r.t. EW-PDF-based calculations. We have found large discrepancies at the level of both shapes and normalisation, as well as a strong dependence on factorisation scales in the EWA case. This underlines the unreliability of the EW-PDF-based approach at realistic collider energies.

In contrast, our method retains full control over theoretical uncertainties, and enables consistent improvements through the inclusion of both higher-order fixed terms and resummed contributions. In particular, we have shown that, in those phase space regions which receive significant contributions from VBF-like topologies, our predictions exhibit uncertainties akin to those of full NNLO calculations, and have an NLO-type behaviour elsewhere. We have observed that Sudakov logarithmic enhancements can naturally be incorporated within our framework, and combined with the set of NNLO terms which we include, without any ambiguity. Likewise, QCD corrections can be accounted for in a straightforward manner; we point out that this is consistent with the fact that lepton PDFs have a strongly-interacting parton component, i.e. feature quark and gluon densities.

Our results provide robust evidence that the proposed approach is a superior and viable alternative to the current approximations used for high-energy lepton collider physics. By capturing all leading contributions in both the soft and collinear regions, and by maintaining exact mass dependence throughout, our method gives reliable predictions not only for inclusive rates but also for differential distributions relevant to experimental analyses. We point out that while this approach will be superseded by exact EW NNLO results (provided that the

latter will also include the resummation of QED collinear dynamics), these are presently out of reach, and will in any case remain much more challenging computations, to be carried out in a case-by-case manner for the foreseeable future. Conversely, efforts towards automating the method we have proposed are already ongoing (and will be guaranteed to succeed thanks to its process independence), and will be included in future releases of the public event generator MG5_AMC; this will be crucial for easier phenomenological applications.

This work gives a template for a comprehensive, precise, and accurate framework at future lepton colliders. While we have emphasised the case of a multi-TeV muon collider, the method which we have proposed will be equally relevant to the very high-precision targets of a typical sub-TeV e^+e^- machine. In future phenomenological studies, it will be important to apply this methodology to more exclusive observables and to other relevant final states, including those sensitive to new physics.

Acknowledgments

We thank Giovanni Stagnitto for discussions and for having provided us with a still-unreleased NLL- and NLO-accurate version of the muon PDFs of ref. [21]. DP and MZ acknowledge the financial support by the MUR (Italy), with funds of the European Union (NextGenerationEU), through the PRIN2022 grant 2022EZ3S3F; likewise FM, through the PRIN2022 grant 2022RXEZCJ, and by the project “QIHEP–Exploring the foundations of quantum information in particle physics”, which is financed through the PNR and is funded by the European Union — NextGenerationEU, in the context of the extended partnership PE00000023 NQSTI — CUP J33C24001210007. SF thanks the TH division of CERN for the hospitality during the course of this work.

Data Availability Statement. This article has no associated data or the data will not be deposited.

Code Availability Statement. This article has no associated code or the code will not be deposited.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] AWAKE collaboration, *The AWAKE Run 2 Programme and Beyond †*, *Symmetry* **14** (2022) 1680 [[arXiv:2206.06040](https://arxiv.org/abs/2206.06040)] [[INSPIRE](#)].
- [2] INTERNATIONAL MUON COLLIDER collaboration, *The Muon Collider*, [arXiv:2504.21417](https://arxiv.org/abs/2504.21417) [[INSPIRE](#)].
- [3] D. Buttazzo, D. Redigolo, F. Sala and A. Tesi, *Fusing Vectors into Scalars at High Energy Lepton Colliders*, *JHEP* **11** (2018) 144 [[arXiv:1807.04743](https://arxiv.org/abs/1807.04743)] [[INSPIRE](#)].
- [4] A. Costantini et al., *Vector boson fusion at multi-TeV muon colliders*, *JHEP* **09** (2020) 080 [[arXiv:2005.10289](https://arxiv.org/abs/2005.10289)] [[INSPIRE](#)].

- [5] D. Buttazzo, R. Franceschini and A. Wulzer, *Two Paths Towards Precision at a Very High Energy Lepton Collider*, *JHEP* **05** (2021) 219 [[arXiv:2012.11555](#)] [[INSPIRE](#)].
- [6] H. Al Ali et al., *The muon Smasher's guide*, *Rept. Prog. Phys.* **85** (2022) 084201 [[arXiv:2103.14043](#)] [[INSPIRE](#)].
- [7] C. Aime et al., *Muon Collider Physics Summary*, [arXiv:2203.07256](#) [[INSPIRE](#)].
- [8] C. Accettura et al., *Towards a muon collider*, *Eur. Phys. J. C* **83** (2023) 864 [Erratum *ibid.* **84** (2024) 36] [[arXiv:2303.08533](#)] [[INSPIRE](#)].
- [9] J. Chen, T. Han and B. Tweedie, *Electroweak Splitting Functions and High Energy Showering*, *JHEP* **11** (2017) 093 [[arXiv:1611.00788](#)] [[INSPIRE](#)].
- [10] T. Han, Y. Ma and K. Xie, *High energy leptonic collisions and electroweak parton distribution functions*, *Phys. Rev. D* **103** (2021) L031301 [[arXiv:2007.14300](#)] [[INSPIRE](#)].
- [11] T. Han, Y. Ma and K. Xie, *Quark and gluon contents of a lepton at high energies*, *JHEP* **02** (2022) 154 [[arXiv:2103.09844](#)] [[INSPIRE](#)].
- [12] F. Garosi, D. Marzocca and S. Trifinopoulos, *LePDF: Standard Model PDFs for high-energy lepton colliders*, *JHEP* **09** (2023) 107 [[arXiv:2303.16964](#)] [[INSPIRE](#)].
- [13] R. Capdevilla, F. Garosi, D. Marzocca and B. Stechauner, *Testing the neutrino content of the muon at muon colliders*, *JHEP* **04** (2025) 168 [[arXiv:2410.21383](#)] [[INSPIRE](#)].
- [14] D. Marzocca and A. Stanzione, *On the impact of the mixed Z/γ PDF at muon colliders*, *JHEP* **03** (2025) 171 [[arXiv:2408.13191](#)] [[INSPIRE](#)].
- [15] K. Ma and T. Li, *Laser induced Compton scattering to dark photon or axionlike particle*, *Phys. Rev. D* **111** (2025) 055001 [[arXiv:2410.17591](#)] [[INSPIRE](#)].
- [16] C.W. Bauer, D. Provasoli and B.R. Webber, *Standard Model Fragmentation Functions at Very High Energies*, *JHEP* **11** (2018) 030 [[arXiv:1806.10157](#)] [[INSPIRE](#)].
- [17] S. Frixione, F. Maltoni, D. Pagani and M. Zaro, *Double neutral-current corrections to NLO electroweak leptonic cross sections*, [arXiv:2506.10732](#) [[INSPIRE](#)].
- [18] S. Frixione, *Initial conditions for electron and photon structure and fragmentation functions*, *JHEP* **11** (2019) 158 [[arXiv:1909.03886](#)] [[INSPIRE](#)].
- [19] V. Bertone et al., *Improving methods and predictions at high-energy e^+e^- colliders within collinear factorisation*, *JHEP* **10** (2022) 089 [[arXiv:2207.03265](#)] [[INSPIRE](#)].
- [20] S. Frixione and G. Stagnitto, *The muon parton distribution functions*, *JHEP* **12** (2023) 170 [[arXiv:2309.07516](#)] [[INSPIRE](#)].
- [21] M. Bonvini, S. Frixione and G. Stagnitto, *Improved small- x resummation for DGLAP splitting functions: HELL 4.0*, in preparation.
- [22] A. Denner and S. Rode, *Automated resummation of electroweak Sudakov logarithms in diboson production at future colliders*, *Eur. Phys. J. C* **84** (2024) 542 [[arXiv:2402.10503](#)] [[INSPIRE](#)].
- [23] G. Altarelli, T. Sjostrand and F. Zwirner, *Physics at LEP2: Vol. 1*, Geneva, CERN (1996) [[DOI:10.5170/CERN-1996-001-V-1](#)] [[INSPIRE](#)].
- [24] S. Dawson, *The Effective W Approximation*, *Nucl. Phys. B* **249** (1985) 42 [[INSPIRE](#)].
- [25] G.L. Kane, W.W. Repko and W.B. Rolnick, *The Effective W^\pm, Z^0 Approximation for High-Energy Collisions*, *Phys. Lett. B* **148** (1984) 367 [[INSPIRE](#)].

- [26] Z. Kunszt and D.E. Soper, *On the Validity of the Effective W Approximation*, *Nucl. Phys. B* **296** (1988) 253 [INSPIRE].
- [27] S. Dittmaier, P. Maierhöfer, C. Schwan and R. Winterhalder, *Like-sign W -boson scattering at the LHC — approximations and full next-to-leading-order predictions*, *JHEP* **11** (2023) 022 [arXiv:2308.16716] [INSPIRE].
- [28] J. Alwall 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** (2014) 079 [arXiv:1405.0301] [INSPIRE].
- [29] R. Frederix et al., *The automation of next-to-leading order electroweak calculations*, *JHEP* **11** (2018) 085 [Erratum *ibid.* **11** (2021) 085] [arXiv:1804.10017] [INSPIRE].
- [30] S. Frixione, O. Mattelaer, M. Zaro and X. Zhao, *Lepton collisions in MadGraph5_aMC@NLO*, arXiv:2108.10261 [INSPIRE].
- [31] R. Ruiz, A. Costantini, F. Maltoni and O. Mattelaer, *The Effective Vector Boson Approximation in high-energy muon collisions*, *JHEP* **06** (2022) 114 [arXiv:2111.02442] [INSPIRE].
- [32] I. Bigaran and R. Ruiz, *Weak bosons as partons below 10 TeV partonic center-of-momentum*, arXiv:2502.07878 [INSPIRE].
- [33] A. Sirlin, *Radiative Corrections in the $SU(2)_L \times U(1)$ Theory: A Simple Renormalization Framework*, *Phys. Rev. D* **22** (1980) 971 [INSPIRE].
- [34] V. Bertone, M. Cacciari, S. Frixione and G. Stagnitto, *The partonic structure of the electron at the next-to-leading logarithmic accuracy in QED*, *JHEP* **03** (2020) 135 [Erratum *ibid.* **08** (2022) 108] [arXiv:1911.12040] [INSPIRE].
- [35] S. Frixione, *On factorisation schemes for the electron parton distribution functions in QED*, *JHEP* **07** (2021) 180 [Erratum *ibid.* **12** (2012) 196] [arXiv:2105.06688] [INSPIRE].
- [36] P.M. Brecht, W. Kilian, J. Reuter and P. Stienemeier, *NLO electroweak corrections to multi-boson processes at a muon collider*, *JHEP* **12** (2022) 138 [arXiv:2208.09438] [INSPIRE].
- [37] Y. Ma, D. Pagani and M. Zaro, *EW corrections and heavy boson radiation at a high-energy muon collider*, *Phys. Rev. D* **111** (2025) 053002 [arXiv:2409.09129] [INSPIRE].
- [38] W. Kilian, T. Ohl and J. Reuter, *WHIZARD: Simulating Multi-Particle Processes at LHC and ILC*, *Eur. Phys. J. C* **71** (2011) 1742 [arXiv:0708.4233] [INSPIRE].
- [39] V.V. Sudakov, *Vertex Parts at Very High Energies in Quantum Electrodynamics*, *Sov. Phys. JETP* **3,4,5,6** (1956) 65 [INSPIRE].
- [40] S. Frixione et al., *Weak corrections to Higgs hadroproduction in association with a top-quark pair*, *JHEP* **09** (2014) 065 [arXiv:1407.0823] [INSPIRE].
- [41] M. Czakon et al., *Top-pair production at the LHC through NNLO QCD and NLO EW*, *JHEP* **10** (2017) 186 [arXiv:1705.04105] [INSPIRE].
- [42] P. Ciafaloni and D. Comelli, *Sudakov enhancement of electroweak corrections*, *Phys. Lett. B* **446** (1999) 278 [hep-ph/9809321] [INSPIRE].
- [43] P. Ciafaloni and D. Comelli, *Electroweak Sudakov form-factors and nonfactorizable soft QED effects at NLC energies*, *Phys. Lett. B* **476** (2000) 49 [hep-ph/9910278] [INSPIRE].
- [44] M. Ciafaloni, P. Ciafaloni and D. Comelli, *Bloch-Nordsieck violating electroweak corrections to inclusive TeV scale hard processes*, *Phys. Rev. Lett.* **84** (2000) 4810 [hep-ph/0001142] [INSPIRE].

- [45] G. Bell, J.H. Kuhn and J. Rittinger, *Electroweak Sudakov Logarithms and Real Gauge-Boson Radiation in the TeV Region*, *Eur. Phys. J. C* **70** (2010) 659 [[arXiv:1004.4117](#)] [[INSPIRE](#)].
- [46] A. Manohar, B. Shotwell, C. Bauer and S. Turczyk, *Non-cancellation of electroweak logarithms in high-energy scattering*, *Phys. Lett. B* **740** (2015) 179 [[arXiv:1409.1918](#)] [[INSPIRE](#)].
- [47] S. Chen et al., *Learning from radiation at a very high energy lepton collider*, *JHEP* **05** (2022) 180 [[arXiv:2202.10509](#)] [[INSPIRE](#)].
- [48] H. El Faham et al., *Electroweak corrections in the SMEFT: four-fermion operators at high energies*, *JHEP* **06** (2025) 241 [[arXiv:2412.16076](#)] [[INSPIRE](#)].
- [49] D. Pagani and M. Zaro, *One-loop electroweak Sudakov logarithms: a revisit and automation*, *JHEP* **02** (2022) 161 [[arXiv:2110.03714](#)] [[INSPIRE](#)].
- [50] A. Denner and S. Pozzorini, *One loop leading logarithms in electroweak radiative corrections. I. Results*, *Eur. Phys. J. C* **18** (2001) 461 [[hep-ph/0010201](#)] [[INSPIRE](#)].
- [51] A. Denner and S. Pozzorini, *One loop leading logarithms in electroweak radiative corrections. II. Factorization of collinear singularities*, *Eur. Phys. J. C* **21** (2001) 63 [[hep-ph/0104127](#)] [[INSPIRE](#)].
- [52] D. Pagani, T. Vitos and M. Zaro, *Improving NLO QCD event generators with high-energy EW corrections*, *Eur. Phys. J. C* **84** (2024) 514 [[arXiv:2309.00452](#)] [[INSPIRE](#)].
- [53] W.A. Bardeen, A.J. Buras, D.W. Duke and T. Muta, *Deep Inelastic Scattering Beyond the Leading Order in Asymptotically Free Gauge Theories*, *Phys. Rev. D* **18** (1978) 3998 [[INSPIRE](#)].
- [54] S. Frixione, Z. Kunszt and A. Signer, *Three jet cross-sections to next-to-leading order*, *Nucl. Phys. B* **467** (1996) 399 [[hep-ph/9512328](#)] [[INSPIRE](#)].
- [55] S. Frixione, *A general approach to jet cross-sections in QCD*, *Nucl. Phys. B* **507** (1997) 295 [[hep-ph/9706545](#)] [[INSPIRE](#)].
- [56] R. Frederix, S. Frixione, F. Maltoni and T. Stelzer, *Automation of next-to-leading order computations in QCD: The FKS subtraction*, *JHEP* **10** (2009) 003 [[arXiv:0908.4272](#)] [[INSPIRE](#)].